CERN openlab

Annual Report
May 2010 – May 2011

Contents

Cover illustration: Fibre channel switches used in Oracle Real Application Clusters (RAC) to support databases for LHC experiments

Third annual report of CERN openlab phase III
Editor: Mélissa Gaillard
Photography: CERN Photo Service – Maximilien Brice (p 3), Luigi Gallerani (p 12, 15, 16, 17, 20, 21, 22, 23), Sylvain Fasel (p 1), Mélissa Gaillard (p 27), Intel (p 22), Claudia Marcelloni (p 6), Andrzej Nowak (p 26), Wolfgang von Rüden (p 30), Siemens (p 8), Society for Science & the Public (p 24)
Fons Rademakers/CERN Photo Club (p29)
Production: André-Pierre Olivier
ISBN 978-92-9083-359-8

CERN openlab: 10 years of partnership

Celebrating its 10th anniversary this year, the CERN openlab is coming of age, and it gives me great pleasure to see this unique organization continue to flourish as its growing band of alumni make their way in the world. In May, I took part in the openlab's 10th anniversary celebrations, held during the Board of Sponsors annual meeting.

That occasion gave me the opportunity to thank the present openlab partners, HP, Intel, Oracle and Siemens for their continuing engagement with CERN: an engagement that has proven itself time and again to be of mutual benefit. It also gave me the opportunity to thank the people: staff of CERN and the openlab partners, CERN Fellows and interns, all of whom have contributed to making the openlab concept such a success.

CERN openlab provides tangible evidence of the virtuous circle linking basic and applied science. Here at CERN, our mission is to push back the frontiers of knowledge. The openlab partners have a very different role. Their job is to provide cutting edge IT products to demanding customers around the world. These may be very different missions, but what makes the partnership work is that CERN is a blue-sky laboratory that uses very real world techniques to unlock the mysteries of nature. And we are perhaps the most demanding customer of them all.

The principle is simple, the requirements of particle physics research push IT to its limits, which leads to new products that in turn feed back into particle physics research: a virtuous circle. Thanks to advances in IT, driven partly by the openlab, the LHC experiments are able to digest and analyse the prodigious quantities of data being delivered by the world's most powerful accelerator. Roughly 25 petabytes of data will be archived and shared every year - the equivalent of 5.3 million DVD movies, which would take a thousand years to watch. Dealing with such huge quantities of data would have been impossible just a few years ago. As a result, the experiments have been turning raw data into published results at a breathtaking pace, thereby adding to the sum of human knowledge. Major breakthroughs cannot be far away, both in terms of blue-sky physics and information technology. I look forward to accompanying the openlab partnership through this exciting phase.

Rolf Heuer
Director-General of CERN

THE CONTEXT

LHC first physics results

The Large Hadron Collider (LHC) project started more than 20 years ago, with the aim of preparing the next major phase in the ongoing quest for a deeper understanding of the fundamental laws of nature. Now that it is up and running, it is the world's most powerful particle accelerator and also the largest and most complex scientific instrument ever built. Located in a 27 km long circular tunnel 100 m underground, it accelerates particles to energies never reached before. Some 9600 superconducting magnets operating at just 1.9 degrees above absolute zero (-271.3˚C), colder than outer space, provide the very strong magnetic fields needed to keep the particles on the right orbit. Two beams of particles travel inside two ultra-high vacuum pipes in opposite directions and are brought to collision in four well-defined points, recreating the conditions that existed a fraction of a second after the Big Bang. Four very large detectors, comparable to huge high-resolution 100 megapixels 3D cameras, record these 'mini Big Bangs' up to 40 million times per second once the machine is running at its full potential.

The LHC saw its first beam in September 2008, but stopped operating for slightly over a year following a severe incident caused by a faulty magnet interconnect. After the repair and the installation of additional protection systems, the accelerator started operation again on 20 November 2009. Milestones were passed extremely quickly, and the first world record beam energy was set on 30 November, promptly followed by many others. On 30 March 2010, beams collided in the LHC with an energy of 7 TeV, marking the start of the LHC research programme. The ALICE, ATLAS, CMS, and LHCb experiments immediately observed and recorded events in their detectors. The accelerator, the experiments' computing systems, and the Worldwide LHC Computing Grid (WLCG) operated perfectly, enabling the physicists to share the results instantly and to publish their first scientific papers promptly.

After a few months of proton-proton collisions, it took only four days for the LHC operations team to switch the LHC from colliding protons to lead ions. After extracting the final proton beam of 2010 on 4 November, commissioning the lead-ion beam was underway by afternoon and first collisions were recorded by the three multi-purpose detectors, ALICE, ATLAS, and CMS, at 00:30 CET on 7 November with stable running conditions marking the start of physics with heavy ions on 8 November.

Lead-ion collisions have opened up an entirely new avenue of exploration for the LHC programme, probing matter as it would have been in the first instants of the Universe's existence. One of the main objectives for lead-ion running was to produce tiny quantities of such matter, which is known as quark-gluon plasma, and to study its evolution into the kind of matter that makes up the universe today. This exploration intends to shed further light on the properties of the strong interaction, which binds the particles called quarks into bigger objects, such as protons and neutrons. Taken together, the results from this first heavy ion run have already ruled out some theories about how the primordial Universe behaved.

Following the winter technical stop, operation of the collider restarted with protons in February 2011. This year saw the fanstastic transition from commissionning to routine operation (accelerator complex, infrastructure, experiments,

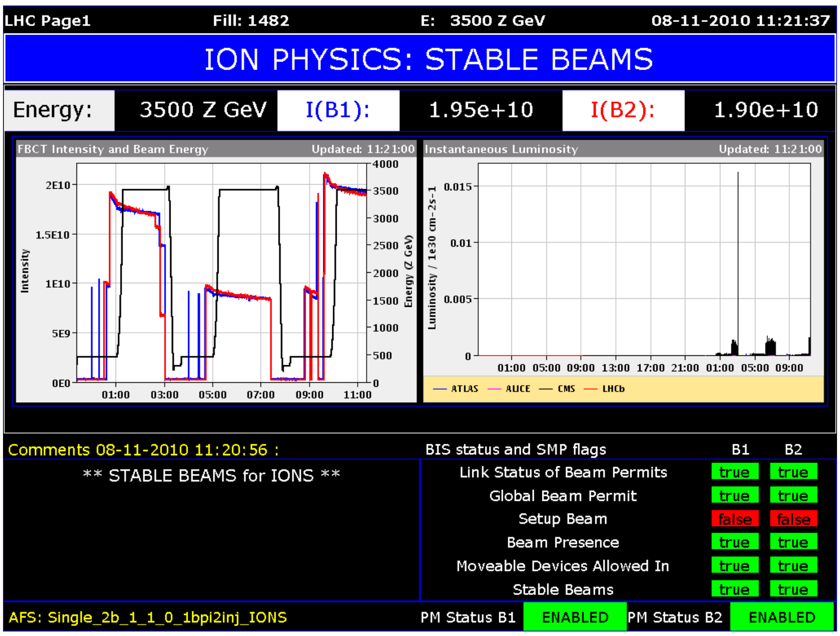

1/ 8 November 2010, CMS control centre screen monitoring the beams for ions.

2/ Display of an event recorded by ATLAS with a Z boson producing two muons.

3/ Event from the first lead ion collisions, reconstructed offline thanks to the WLCG, showing tracks from the Inner Tracking System and the Time Projection Chamber of ALICE.

4/ WLCG activity on 11 November 2011, during the LHC heavy ion run.

2
1 | 3 4

computing and all other services) in a particularly smooth way as the target integrated luminosity for 2011 had already been achieved by mid-year.

A global collaboration powered by the WLCG

The new phenomena that scientists hope to find are extremely rare, hidden deep in already known physics. The LHC has therefore been designed to produce a very high rate of collisions (40 million per second) such that rare events can be found within a reasonable time. The amount of raw data produced per second, once the machine is up to its full potential, is in the order of one million gigabytes. None of today's computing systems are capable of recording such rates. Sophisticated selection systems, called first level triggers, allow most of the events to be rejected after one millionth of a second, followed by a higher level of selection applying more sophisticated criteria. This drops the data rate per experiment to below one gigabyte per second. Even after such a drastic data reduction, the four big experiments, ALICE, ATLAS, CMS and LHCb, will produce over 25 million gigabytes per year, the equivalent of 5.3 million DVD movies, which would take a thousand years to watch.

To store, share and analyse these data, tens of thousands of computers distributed worldwide are being harnessed in a distributed computing network called the Worldwide LHC Computing Grid (WLCG). This ambitious project supports the offline computing needs of the LHC experiments by connecting and combining the IT power of more than 200 computer centres in 36 countries. CERN is the central 'hub' of the WLCG, the Tier-0 centre. This is where a first copy of all data from the LHC experiments is held and where the first reconstruction and data quality checks are performed. This central hub, which provides less than 20% of the total computing capacity, is connected to 11 other major computing centres (Tier-1) using dedicated optical fibre links working at multiples of 10 gigabits per second.

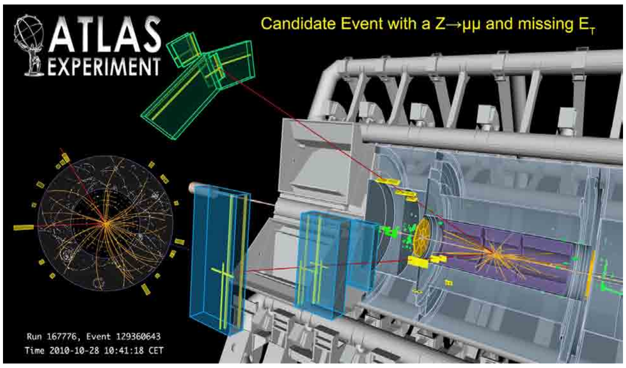

These sites are large computer centres with sufficient storage capacity and round-the-clock support. They provide distribution networks, processing of raw data, data analysis, and storage facilities. The Tier-1 sites then make the data available to Tier-2 centres, each consisting of one or several collaborating computing facilities, which can store sufficient data and provide adequate computing power for specific analysis tasks. The Tier-2 sites, grouped in Federations, cover most of the globe. Individual scientists access these facilities through local (also sometimes referred to as Tier-3) computing resources, which can consist of clusters in a University Department or even individual PCs.

During the periods of data-taking, the WLCG achieved impressive rates of data transfer, reaching peaks of more than 200000 gigabytes per day. With data available to a community of 10000 scientists within hours, the speed of analysis has been dramatic, with results ready for publication within weeks. To rise to such unprecedented computing challenges, new and advanced systems were needed requiring the joint forces of science and industry to expand technological boundaries. CERN openlab partners contributed in a tangible way to their development and, as shown in this report, continue to collaborate on new solutions with success.

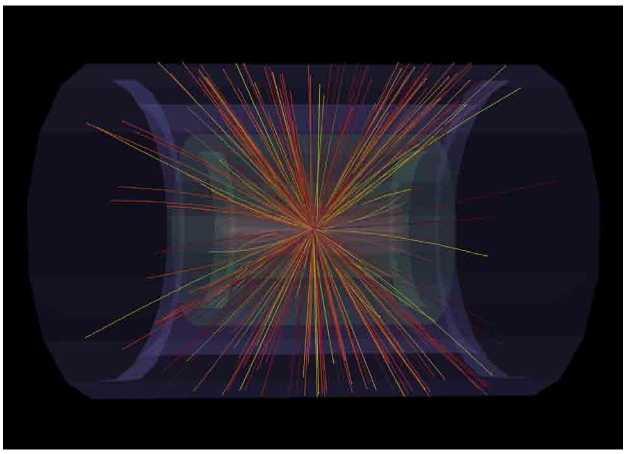

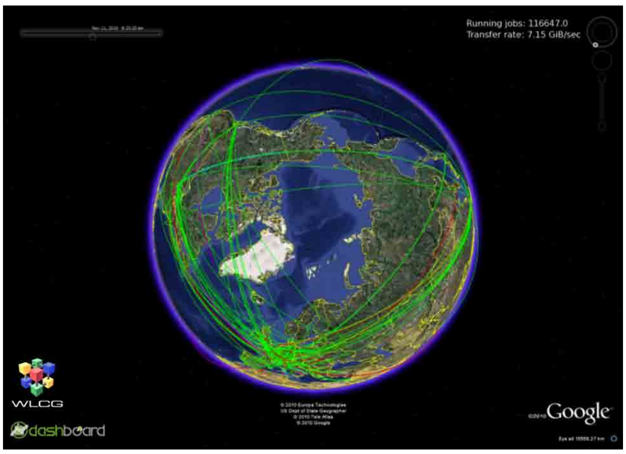

THE CONCEPT

Addressing future challenges

With the LHC reaching its cruising speed in terms of luminosity, reliable computing is becoming even more important. The attention of the scientists is now fully directed towards new discoveries with computing being the key enabler. Stability takes precedence over new features. Nonetheless, there is a constant need to get more out of the installed equipment and to make the right choices for the coming years to cope with the ever-increasing amount of data. Therefore, innovation is a must, while backward compatibility needs to be preserved. CERN openlab allows us to embrace these challenges with confidence, building on our ten-year experience.

The successful CERN openlab concept was formulated in 2001 and stayed basically unchanged throughout the last decade. CERN openlab is a framework for multilateral, multi-year projects between CERN and the IT industry. Within this structure, CERN provides access to its complex IT infrastructure and its engineering experience, in some cases even extended to collaborating institutes worldwide. Testing in CERN's demanding environment provides the partners with valuable feedback on their products while allowing CERN to assess the merits of new technologies in their early stages of development for possible future use. This framework also offers a neutral ground for carrying out advanced R&D with more than one company. The openlab partners commit to a three-year programme of work and provide three kinds of funding: salaries for young researchers, products and services, and engineering capacity.

Wolfgang von Rüden handing over the leadership of CERN openlab to Bob Jones

CERN openlab has been organised into successive three-year phases. In openlab-I (2003–2005), the focus was on the development of an advanced prototype called opencluster. CERN openlab-II (2006–2008) addressed a range of domains from platforms, databases and Grid, security and networking with HP, Intel and Oracle as partners and EDS, an HP company, as a contributor. The combined knowledge and dedication of the engineers from CERN and the companies have produced exceptional results leading to significant innovation in many areas.

CERN openlab is now in the last year of its third phase (2009–2011). It has not only capitalised on but also extended the successful work carried out in openlab-II with a similar aim of hosting a few major projects with a particular focus on technologies and services relevant to CERN and its partners. The technical activities are organised in four Competence Centres (CC): the Automation and Controls CC, the Database CC, the Networking CC, and the Platform CC.

The technical achievements, training sessions and other events of the past year are described later on in this report.

Again, teaching played an important role with participation in the CERN School of Computing, the CERN openlab Summer Student Programme and the visit of a group of the ISEF special award winners. During the 2010 Board of Sponsors meeting we initiated the process of preparing openlab-IV (2012-2014), we held a workshop in December, followed by regular bilateral discussions. We will soon be ready to conclude and sign the contracts for the next round, helping to shape the future of computing for CERN and its users.

Looking back over these ten years of CERN openlab, the results achieved, seeing what is on the table for the next round and knowing the quality of both the people and the companies committed to it, I am confident that there are many more years of successful activities ahead.

The time has come for me to say good-bye to CERN. I am indebted to all the industry partners and contributors, and of course to the numerous colleagues at CERN, who, together, made the undertaking a success and a pleasure to lead. I pass the leadership to my successor Bob Jones, along with my best wishes, knowing that the project will be in good hands. I wish CERN openlab and its brilliant team a bright and exciting future.

Wolfgang von Rüden
Head of CERN openlab

The openlab team: a broad partnership

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), DB (Database Services) and DI (Departmental Infrastructure) groups host openlab activities, as does the ICE (Industrial Controls and Electronics) group in the EN (Engineering) department.

A list of the IT and EN department staff most closely involved in the CERN openlab activities is given below. The distributed team structure 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. In addition, significant contributions are made by students participating in the CERN openlab student programme, both directly to openlab activities (6 students during summer 2009) and more widely to WLCG and other Grid and CERN related activities in the IT department.

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
Bill Johnson	HP
Christian Morales	Intel
Juan F. Rada	Oracle
Thomas Hahn	Siemens

CERN openlab Management Unit

François Flückiger	CERN openlab Manager
Mélissa Gaillard	Communications Officer
Kristina Gunne	Administrative Assistant
Sverre Jarp	Chief Technology Officer

CERN openlab Fellows and Staff (sponsor indicated)

Zbigniew Baranowski	Fellow (Oracle)
Marcin Bogusz	Fellow (Siemens)
Carlos Garcia Fernandez	Fellow (Oracle)
Milosz Hulbój	Staff (HP)
Svetozar Kapusta	Fellow (Oracle)
Omer Khalid	Fellow (Siemens)
Vlad Lăpădătescu	Fellow (HP)
Alfio Lazzaro	Fellow (Intel)
Julien Leduc	Fellow (Intel)
Andrzej Nowak	Staff (Intel)
Daniel Filipe Rodrigues	Staff (Siemens)
Yngve Sneen Lindal	Technical student (Intel)
Filippo Tilaro	Fellow (Siemens)

CERN Liaison

Renaud Barillère	Liaison with Siemens
Sverre Jarp	Liaison with Intel
Jean-Michel Jouanigot	Liaison with HP
Tony Cass	Liaison with Oracle

Other IT and EN Departments Staff Contributors to CERN openlab

Sébastien Ceuterickx	CS Group
Brice Copy	ICE Group
Eva Dafonte Pérez	DB Group
Dirk Düllmann	DSS Deputy Group Leader
Philippe Gayet	ICE Group Leader
Manuel Gonzales Berges	ICE Group
Eric Grancher	DB Deputy Group Leader
Jose Carlos Luna Duran	CS Group
Alberto Pace	DSS Group Leader
Anton Topurov	DB Group
Artur Wiecek	DB Group

Industry Partner Liaisons with CERN

Stéphane Laroche	HP
Claudio Bellini	Intel
Herbert Cornelius	Intel
Stephan Gillich	Intel
C. Gregory Doherty	Oracle
Monica Marinucci	Oracle
Ulrich Eger	Siemens
Thomas Hahn	Siemens
Georg Trummer	Siemens

THE RESULTS

The Automation and Controls Competence Centre

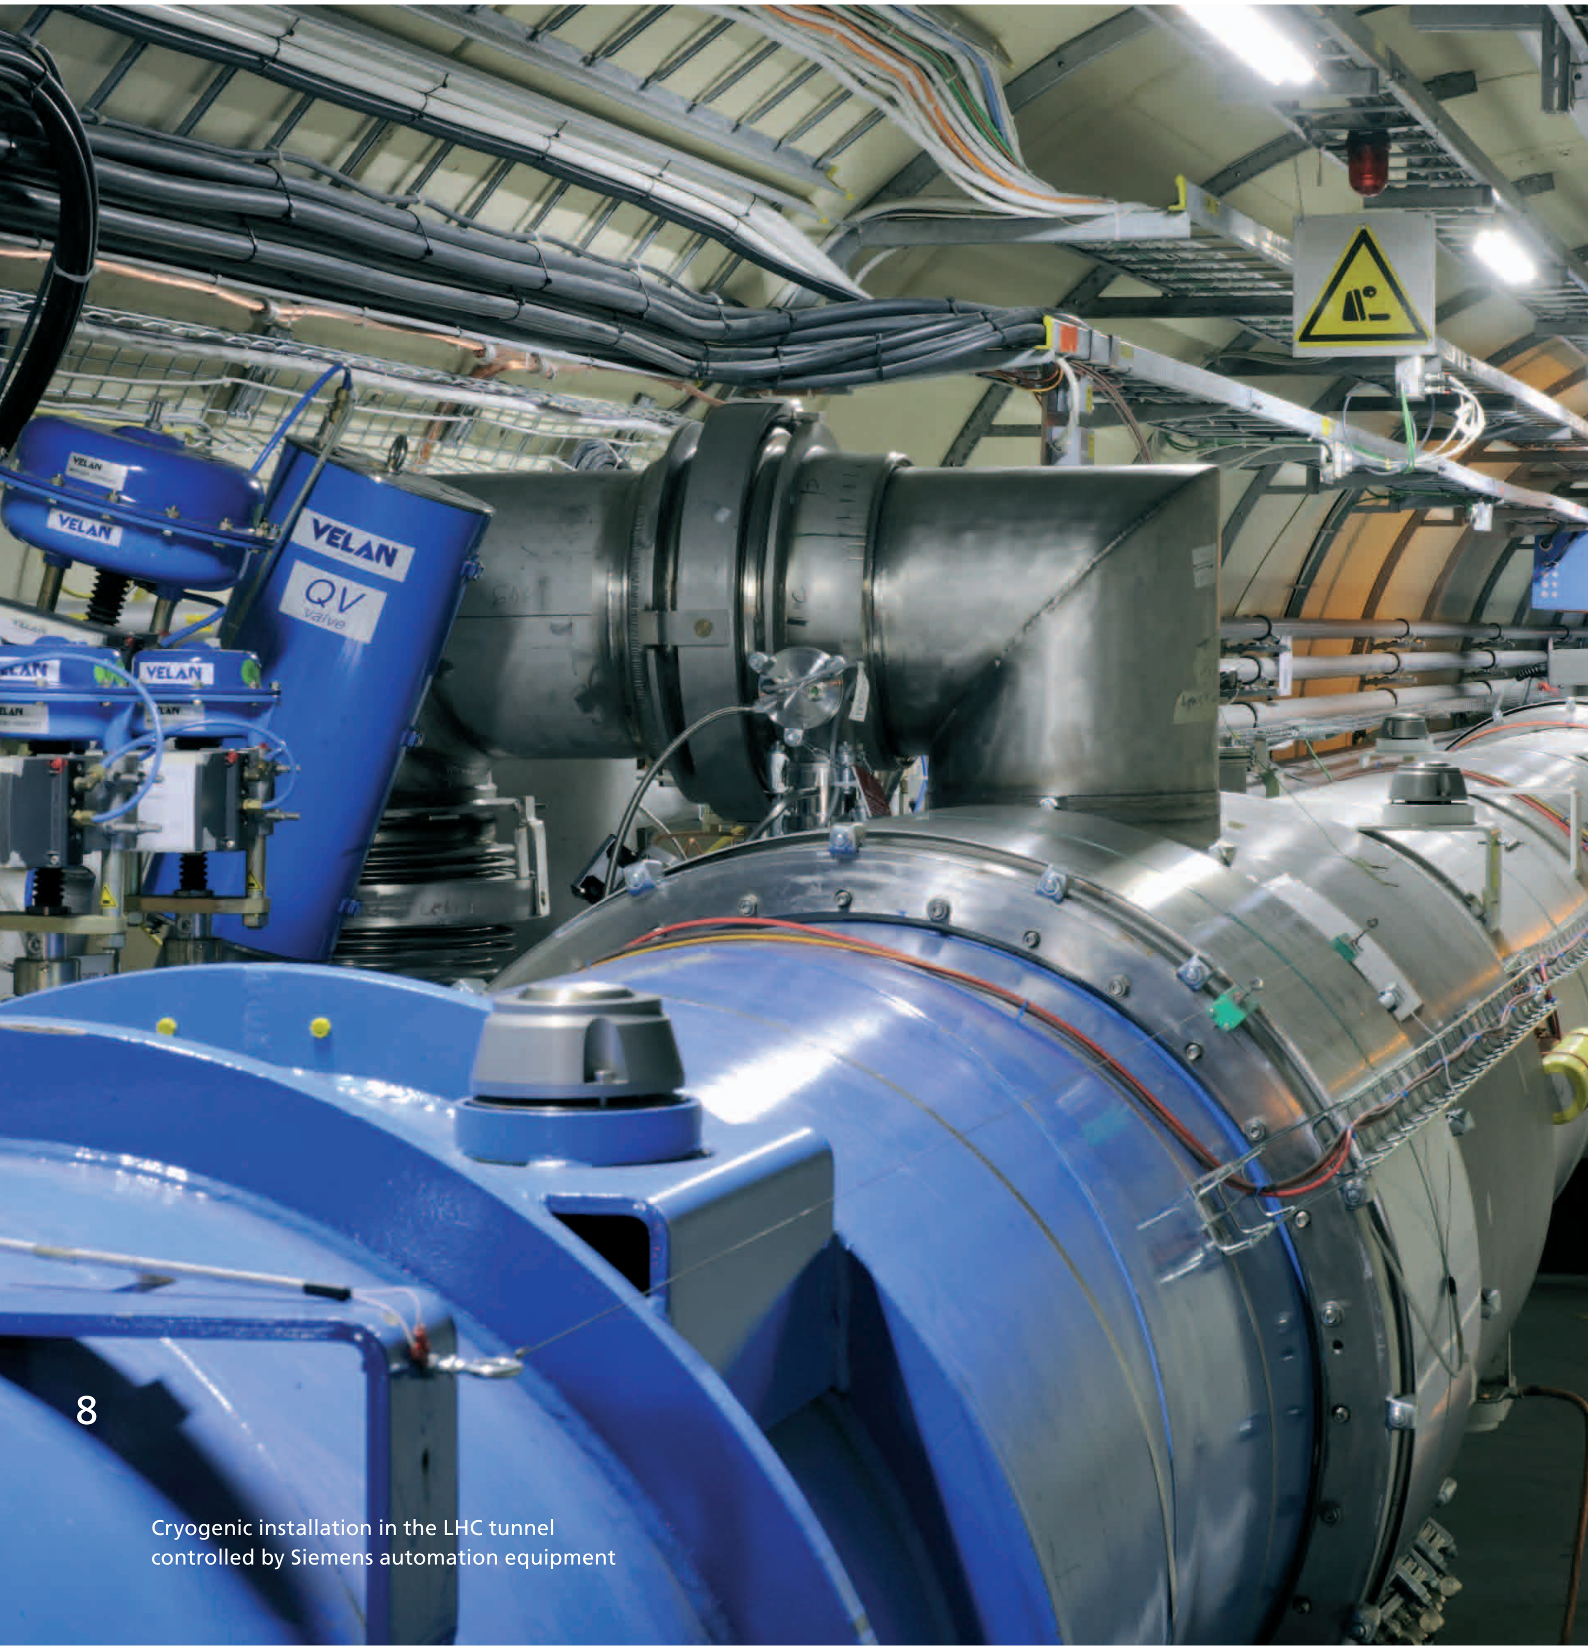

Cryogenic installation in the LHC tunnel
controlled by Siemens automation equipment

The CERN environment is highly reliant on industrial control systems such as control actuators, remote Process Field Bus (PROFIBUS) Input/Output modules, Supervision Control And Data Acquisition (SCADA) systems, and Programmable Logic Controllers (PLCs). Siemens provides a large set of solutions for Automation and Controls and has been collaborating for two years on these topics within the CERN openlab framework. For this second year of joint work, cyber security has been one of the major focuses of the Automation and Controls Competence Centre. A study was primarily carried out at the level of the PLCs, and then extended to the STEP 7 programming environment and the SCADA level. Another key focus was the addition of features to the STEP 7 and PVSS development environments, as well as the test and definition of new facilities (ORACLE Archiver and Web clients) for the PVSS run-time environment.

Industrial Control System security

In the last few years, the growing interconnectivity between Industrial Control Systems (ICS) and enterprise networks has brought numerous benefits, but it also exposed critical installations to new sources of possible cyber threats. Such infrastructures are subject to targeted attacks aimed at disrupting SCADA systems, industrial software, or in the worst cases industrial devices, which are the core entities of every industrial installation. Since ICS have different characteristics and functionalities from traditional IT systems, IT-designed security solutions cannot be directly deployed.

PLC security project

The PLC security project aims at studying and designing new testing and cyber-attack methodologies, in view of supporting Siemens testers in detecting - and subsequently eliminating - potential PLC security issues. Keeping the international standard for security technologies for manufacturing and control systems, ISA-99, as the reference standard, a new customized security model has been defined to analyse and evaluate the robustness and the security level offered by individual industrial devices. For this purpose, a testing framework, TRoIE (Test-bench for Robustness of Industrial Equipment), was designed and implemented to automate the security analysis of IP-based embedded systems. This testing platform has been devised to comply with ISCI CRT (Communication Robustness Tests) specifications. These requirements have been translated into a set of communication protocol tests to certify the device's capability to adequately maintain indispensable process control services while being subjected to abnormal and malformed network traffic. The testing approach makes use of the combination of "protocol fuzzing" techniques relying on "grammar definitions": invalid messages and malicious packet sequences are generated according to a specific set of rules (called grammars) and afterwards injected to the device, with the purpose of determining potential security vulnerabilities related to the implementations of the devices' protocols. Over the last two years of proficient collaboration with Siemens in this area, the Automation and Controls Competence Centre (ACCC) team directly contributed to improving the robustness of Siemens PLC equipment.

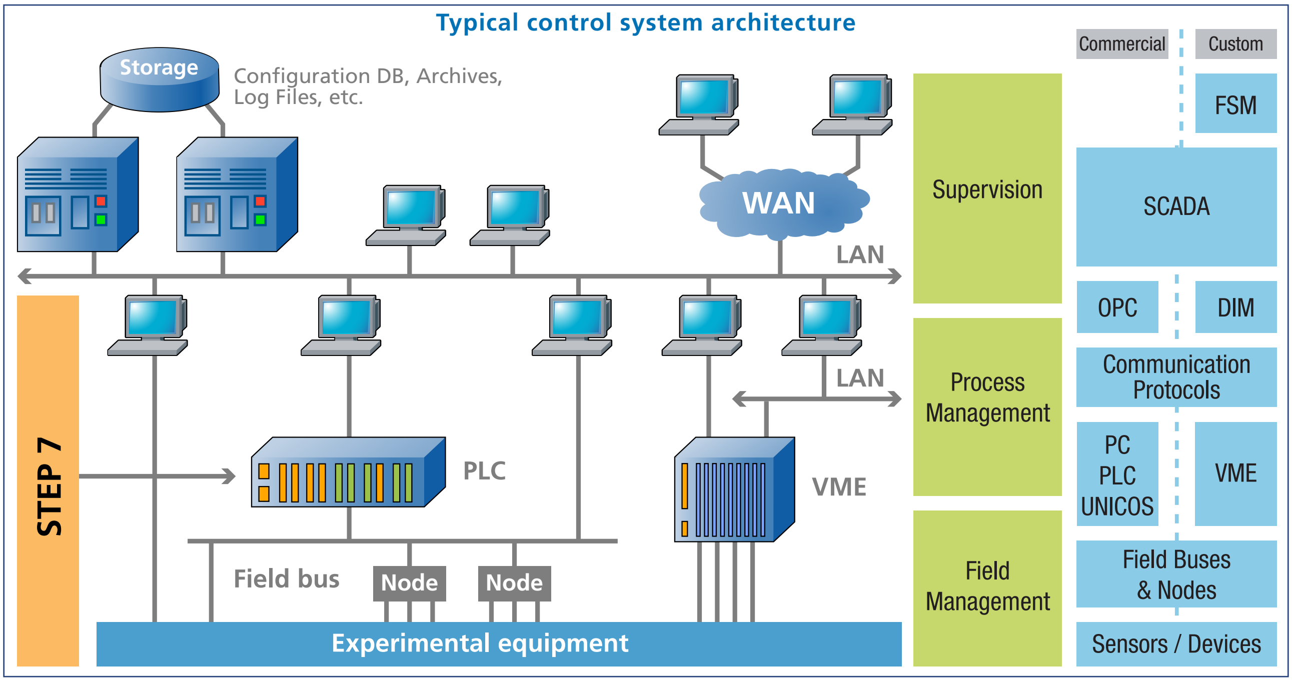

Engineering software security

Historically, ICS has not been the explicit target of malware attacks, but last year the STUXNET malware made the headlines. This case shows the growing importance of security solutions for ICS. Initial investigations were carried out in conjunction with the CERN security team to identify potential vulnerabilities. As a consequence, a security-oriented software development cycle was proposed to the Siemens´ Engineering Software development team. These recommendations have been provided to Siemens so that they could be included in the software development process.

PVSS security

PVSS, the SCADA of reference at CERN, is used to control complex and heterogeneous systems, such as the LHC accelerator and experiments. It is composed of a number of software components distributed over a network that interface with field-level industrial components, such as PLC devices, to gather operational data and enable system operators to monitor and manage the control systems. Hence, it is an important part of a distributed industrial system and it needs to be included in an organization's security system strategy. Based on CERN's extensive experience with PVSS, a best-practices strategy for securing SCADA systems, in particular for PVSS, was devised and documented based on ISA security standards. Now, these recommendations are being extended to a concrete list of potential vulnerabilities to evaluate the security risks and robustness of PVSS software components.

Development environments

STEP 7 automated deployment

Most organizations use some type of software inventory management system to manage software installation, deployment, and updating on the machines in their environment. One of the projects carried out with Siemens consisted of enabling STEP 7 software installations to work with software inventory management systems, such as CMF (Computer Management Framework) at CERN. This required the analysis of the available tools (CFEngine), an automation framework for system administration or IT management, a Microsoft Software Management Server, and to identify how the STEP 7 installation process could be modified to work with these tools. The STEP 7 installation software, called SIA engine, has been adapted after architectural changes, and can now manage automated installations in silent mode for a given set of target machines distributed over the network. The recommended architecture relies on deploying Siemens software on a shared server. Target machines have network access thanks to the light-weight Microsoft Software Installer (MSI) which is setup on the central deployment server. When the first installation is needed, the user sends a request to the deployment server which installs the light-weight MSI which, in turn, installs SIA engine on the target machine. Then SIA engine connects to the shared server which has access to Siemens software, and manages all the pre-requisite, installation and post-installation steps.

In the last year, the STEP 7 deployment project in large-scale environments has been successfully completed, and Siemens have implemented the recommended strategy in the next versions of the STEP 7 installer.

PVSS installation tools

The goal of the Installation Tool work package is to benefit from an integrated tool facilitating the management of large PVSS Distributed Control Systems (DCS). At CERN, the size and complexity of PVSS projects and systems requires such tools to be in place. The experience in tackling this

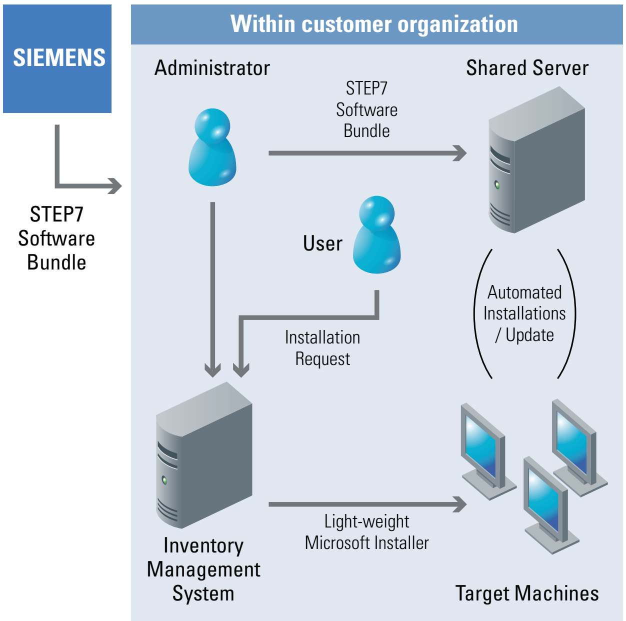

STEP 7 deployment architecture

issue is also of interest to ETM, now incorporated into and strategically managed in Siemens' Automation and Drives – Industrial Automation Solutions division. The topic list includes the packaging of the project components, the remote deployment, and the monitoring of current configurations of PVSS systems. In the context of the openlab collaboration, the CERN solution, named Installation Tool, was studied and analysed, and a proposal of requirements for a similar tool that would be fully integrated in PVSS is under discussion.

Source code management in PVSS

Apache Subversion (SVN) is a very well known software revision control system used to keep track of changes to source code and project files. The PVSS development environment did not feature integrated SVN support, forcing users to control their project revisions using external tools. An SVN Plugin was developed with the aim of providing tools to manage project versions within GEDI (PVSS's Graphical Editor). Its architecture is based on a programming interface written in the native CTRL scripting language, and a set of PVSS panels which can be used by developers to visually control the versioning of files. Several functionalities were developed, namely a project viewer which includes SVN information, a browser of the SVN repository, and a number of standard dialogs to execute SVN operations. A first version was made available and valuable feedback was received from initial users. The SVN Plugin was presented to ETM in the course of a visit and discussions were held on its applicability to future versions of the PVSS system.

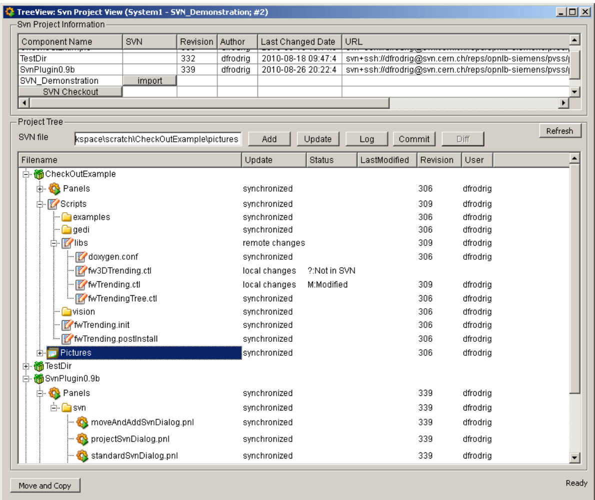

SVN Plugin project view panel

PVSS run-time

RDB archiver

CERN has always been one of the biggest users of PVSS, and, notably, of its key system functionality, the RDB Archiver. Due to the scale and complexity of control systems at CERN, a very high data rate is required to be archived. The work between CERN EN and IT departments and ETM prior to the CERN openlab collaboration had already resulted in a major performance and scalability optimization for the archiving rate. In the tests performed in laboratory conditions, an accumulated stream of over 250 000 Data Point Element (DPE) value changes per second, generated by 150 PVSS systems, was recorded to an Oracle database with Real Application Clusters (RAC).

Currently, ETM and Siemens are working on the major future version of the SIMATIC WinCC Open Architecture. It will feature enhanced archiving services. CERN has become a part of this development effort by starting a new project within the CERN openlab framework. Its aim is to develop the archiving software module for the Oracle RDBMS (Object-Relational DataBase Management System) according to the new architectural guidelines. ETM has developed storage modules for other relational database products, whereas CERN is responsible for the development of the Oracle component to make sure that its architecture meets the Organization's requirements – such as keeping excellent performance and scalability while delivering improved flexibility. An initial version has already been implemented and a special manager bridging the existing PVSS 3.X system and the new logging service is being developed for the initial functionality and performance testing.

Web clients

The Web Access work package included two topics - WebPlugin and Pocket Client. The PVSS Web Plugin is a solution for enabling access to PVSS panels from web browsers. A new version was made available by ETM, removing limitations previously reported by the CERN openlab ACCC team. This new version was re-tested and the following improvements were observed: the installation of the Plugin is now more intuitive, and the performance - measured as the maximum frequency of value changes - is comparable to the PVSS-based RemoteUI. After confirming the improvements on previously analysed metrics, the testing setup was extended to include security and proxy aspects. The preliminary results were presented in the context of the JCOP (Joint Controls Project) Framework, an integrated set of guidelines and software tools which is used by Developers of the Control System to develop their part of the Control System application, and suggestions for further enhancing the testing led to the development of CPU (Central Processing Unit) usage monitoring of PVSS processes. The Pocket Client is another solution for enabling web access, using a different approach based on JavaScript and HTML. Although in a beta stage, it was also analysed as a potential future tool for CERN, and the concepts were also presented in the context of the JCOP Framework.

The Database Competence Centre

The year 2010 was somewhat of a thrilling moment for the Database Competence Centre (DCC): would all of the preparations—especially for the recording of controls configuration and logging data for the accelerator itself and for the recording and export of conditions data for the experiments—be up to the challenges of the first "production" run of the LHC and the data-taking and analysis for the experiments? A year later, the answer is a resounding "Yes". As shown on the figures, the replication rate for conditions data from the ATLAS experiment to the different WLCG Tier-1 sites, and the statistics concerning data volumes logged for the LHC accelerator both reflect the change to production running for the LHC during 2010.

Securing constant access to the physics data is paramount. As a consequence, priority is given to service stability in terms of database requirements. This, together with the developing multi-year schedule for accelerator operations, led to a postponement of the deployment of Oracle Database 11g Release 2 for large scale production. Hence, the numerous new features offered by this release—reported on by the DCC team previously- do not benefit the users as yet. Nevertheless, interesting development and evaluation work has continued, notably in the areas of database replication, virtualisation and monitoring.

Number of Logical Change Records (LCRs) in Millions, per Month, by Tier-1 Site

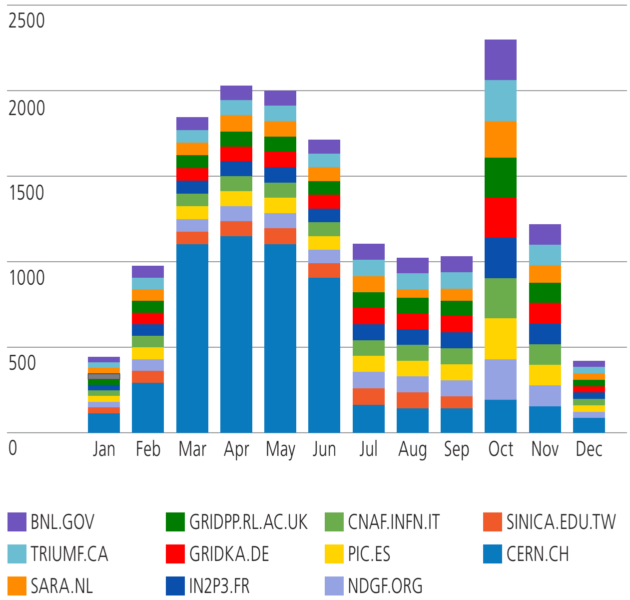

Replication rate for conditions data from the ATLAS experiment to the different WLCG Tier-1 sites in 2010

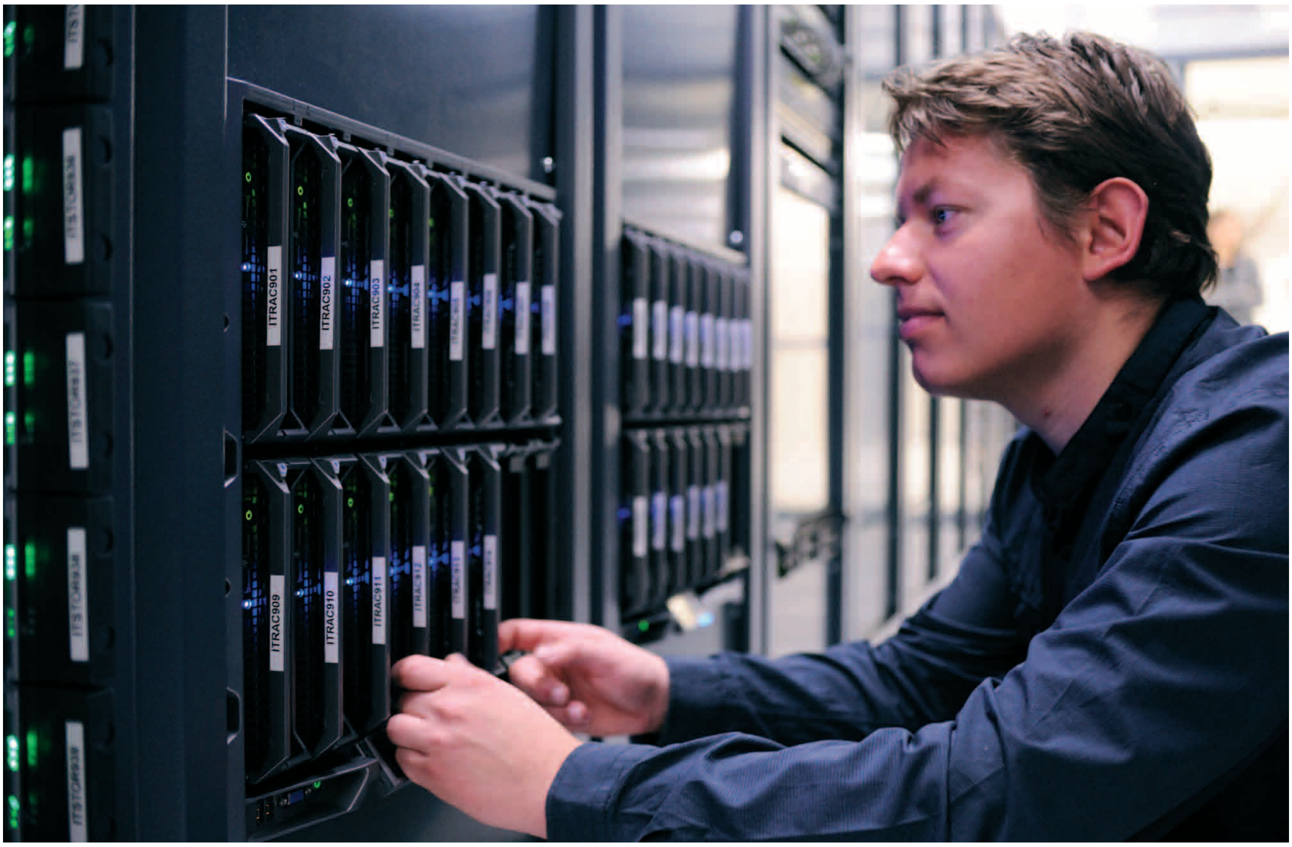

LHC Logging Service - Logged records

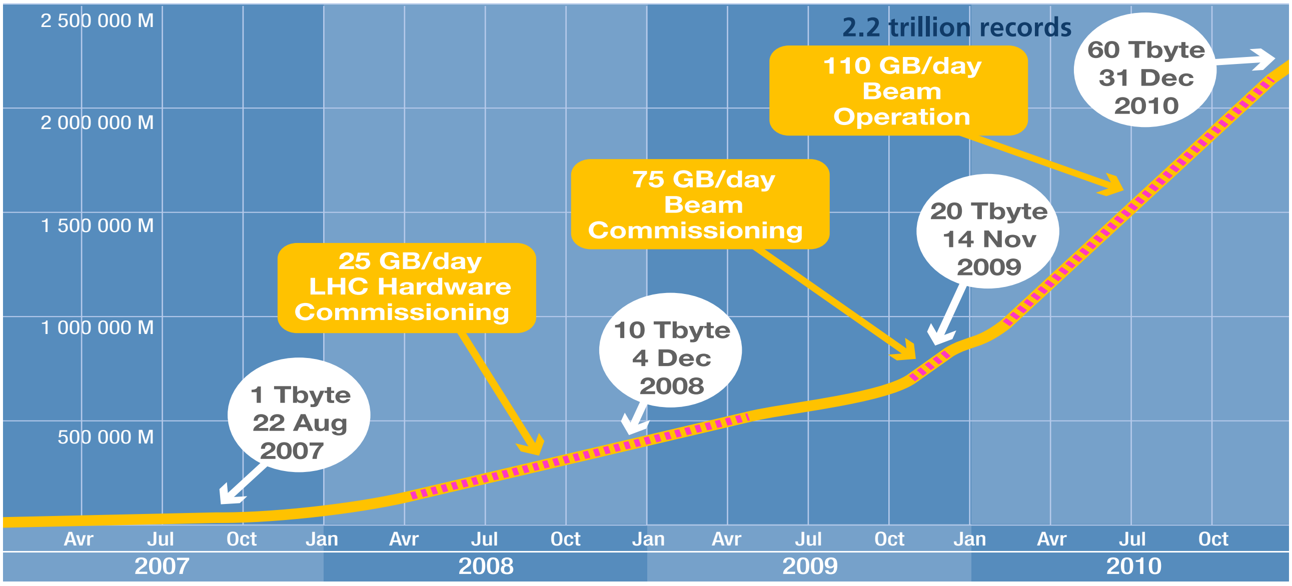

Statistics concerning data volumes logged for the LHC accelerator reflecting the change to production running during 2010

Database replication

Background

Oracle Streams is an essential tool for the distribution of experiment metadata to the Tier-1 sites as well as within CERN and has been the focus of significant development effort within CERN openlab. However, in CERN's loosely controlled environment, user changes to the source schema can easily disrupt replication unless the equivalent changes are applied manually to the target(s). Since two new options for replication, Oracle Active Data Guard and Oracle GoldenGate, are now available and look to be more robust against user errors, it was natural for these to be the subject of the DCC team's study this year.

Oracle Active Data Guard is an extension in Oracle 11g to the previous Data Guard replication tool. Whilst the Data Guard database replicas are inactive secondary copies of a database system, waiting to be called into service in the event of a failure on the primary system, Active Data Guard replicas support read access. Active Data Guard can thus potentially address CERN and the Worldwide LHC Computing Grid's (WLCG) needs for the distribution of metadata as only read access is required at the Tier-1 sites. However, discussions with Oracle replication experts confirmed the need for a tight coupling of database patching and upgrade schedules at the source and target sites. As this is not achievable in our multi-organisation collaboration, Active Data Guard cannot be exploited in this way. However, it remains of definite interest for the replication of data within CERN—for example replication of data from database systems at the experiment sites to systems housed in the computer centre.

Fortunately, the other newly available replication technology, GoldenGate, does not require such tight coupling between the source and target sites—indeed, it is even designed to enable replication between different RDBMS (Relational Database Management System) implementations, for example between Oracle and MySQL—and has thus been investigated extensively over the past year.

Oracle GoldenGate evaluation

A dedicated testbed configuration was established to enable tests of two different GoldenGate versions against two database releases. Initial performance tests with the default software configurations established that GoldenGate performs slightly better than Oracle Database 10g Streams but that it cannot reach the throughput offered by the latest version of Streams provided with the Oracle Database 11g. However, exploiting GoldenGate's "BatchSQL" optimisation mode enabled the team to demonstrate throughput similar to that of Streams 11g for generic data. Tests with the latest GoldenGate version revealed significant improvements in data delivery parallelism, resulting in higher overall throughput, as seen on the related figure.

As well as raw replication performance, stability and reliability aspects are essential. Here, a long term stability test gave remarkable results as the team achieved over four months of continuous data propagation without any negative impact on the source database. Additionally,

Maximum Logical Change Records (LCRs) throughput measured at CERN

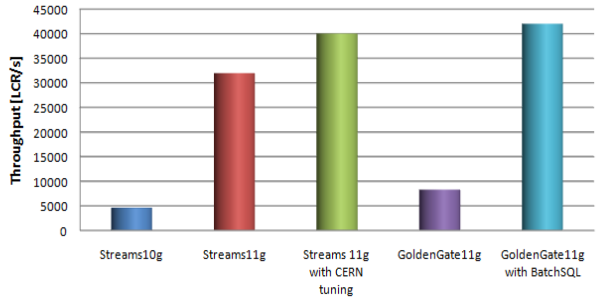

GoldenGate's recording of data changes from master database (in files called trail files) delivers great advantages in terms of data recovery. As trail files are copied to all replicas, these local trail files can be used to re-apply any lost transactions in the event of loss of a slave (replica) database to restore full consistency between master and replica. No additional data retransmission from source database is needed—a very beneficial feature for such a widely distributed environment.

Whilst GoldenGate delivers fast and reliable replication of data in cases where the structure of the database schema is static—representing the majority of industry use cases—replication of schema structure updates is of interest to CERN. Indeed, the LHC experiment supervisory controls and data acquisition systems have data management optimisations which update schema structure in order to improve data collection performance. Unfortunately, the team was unable to achieve the needed performance when simulating a real LHC workload as this use of schema updates prevents the metadata collection applications from leveraging optimisations within the GoldenGate product. Nevertheless, both CERN and Oracle remain hopeful that this may evolve in the future with Jagdev Dhillon, Oracle GoldenGate's Senior Director of Product Development, saying:

«We appreciate CERN's feedback on the GoldenGate product and look forward to continuing our joint efforts to further improve the performance of Oracle GoldenGate with the LHC application workload.»

Based on expertise gained during the development of monitoring tools for Streams, the CERN openlab team also gave some feedback on Director, the GoldenGate monitoring interface, which may help to extend the utility of this effective and intuitive interface when used in a complex distributed environment.

Virtualisation

Previous work in openlab has delivered the seamless integration of Oracle's virtualisation platform, OracleVM, with CERN's Extremely Large Fabric management system, ELFms. In particular, the DCC team demonstrated the uniform treatment of virtual images and physical machines by the provisioning systems at the end of 2009.
Building on this initial work, in 2010 the focus switched to delivering a production-quality environment with, notably, support for Gigabit Ethernet, Ethernet bonding and high availability support through provision of machine pools for different purposes (development, test, production) and exploitation of live migration. Specifically, a command line interface combining XEN and OracleVM commands was developed to simplify complex operations. As a result, it takes just one command to create a virtual machine instance from scratch, and monitoring ongoing instance behaviour has been greatly simplified. Advances such as these have helped to increase the utilisation and reliability of Oracle WebLogic Server on JRockit Virtual Edition environment, now approaching 8 months continuous up-time without any interruption.

The virtualisation team hosted two openlab summer students. The first student explored memory ballooning, a technology that allows running virtual machines with more allocated memory than the real memory of the hosts. These tests demonstrated the possibility to allocate, and stress, up to 30 dual-core virtual machines with 4 GiB of Random Access Memory (RAM) in an eight-core host with 48 GiB of RAM, while no more than 12 hosts could have been allocated without memory ballooning. In cases where the applications running on the virtual machines made little use of the memory allocation, the team was even able to operate with up to 120 virtual machines running on the server.

The task of the second summer student was to integrate Oracle WebLogic Server on JRockit Virtual Edition into the management and distribution system. This work led to a presentation of this technology at Oracle OpenWorld and to a press release. The DCC team also had the pleasure of hosting a visit from the Oracle WebLogic/JRockit team at CERN.

Monitoring

The work in the monitoring area during the year aptly reflected CERN openlab's motto: "You make it, we break it". Practical, careful testing prior to deployment cannot guarantee a zero rate for incidents in the production phase. Some anomalies were reported after the deployment of the 11g release of Oracle Enterprise Manager, but the close collaboration between the Oracle Enterprise Manager Development team and CERN engineers, enabled through CERN openlab contacts, made a swift solution possible. The investigations successfully identified a small memory leak, the effects of which were vastly magnified at CERN—partly due to the deployment scale and partly due to the use of a wide range of Web browsers for connecting to Enterprise Manager.

Once the problem was solved, the DCC team could fully exploit the fact that Enterprise Manager 11g runs on and provides monitoring for WebLogic, Oracle's new application server family. Actually, this feature was a strong motive to deploy the 11g release in the context of the migration from Oracle iAS to WebLogic. Indeed, the team had developed a very effective automatic discovery process for WebLogic domains using the Enterprise Manager 11g command line interface. Furthermore, the Oracle JRockit Mission Control monitoring and profiling capabilities proved to be very useful to CERN application developers, helping them resolve memory leaks on Java Virtual Machine level.

A third CERN openlab summer student joined the DCC team and tested the Corrective Actions feature of Enterprise Manager, which enables the automatic execution of corrective scripts when exceptions are detected. This feature proved to be very useful during the system instability described above: the memory leak invariably led to problems during the night, leading to some delays in the debugging. By exploiting the Corrective Actions feature, the information needed to debug the issue could be automatically made available.

This ability to take corrective action automatically is related to an idea which has often been discussed in the openlab context: the possibility to link Enterprise Manager to external, non-database events. The DCC team made progress on this topic during the year by enabling the import of information related to the LHC beam energy into Enterprise manager. This should facilitate the analysis and correlation of LHC databases performance and workload with the state of the machine and thus enable to act—automatically—in an appropriate manner depending on the machine operation phase. More developments in this area are expected in the coming year and as we look ahead to openlab IV.

From left to right: Anton Topurov, Carlos Garcia Fernandez, Tony Cass, Eva Dafonte Pérez, Eric Grancher, Zbigniew Baranowski

The Networking Competence Centre

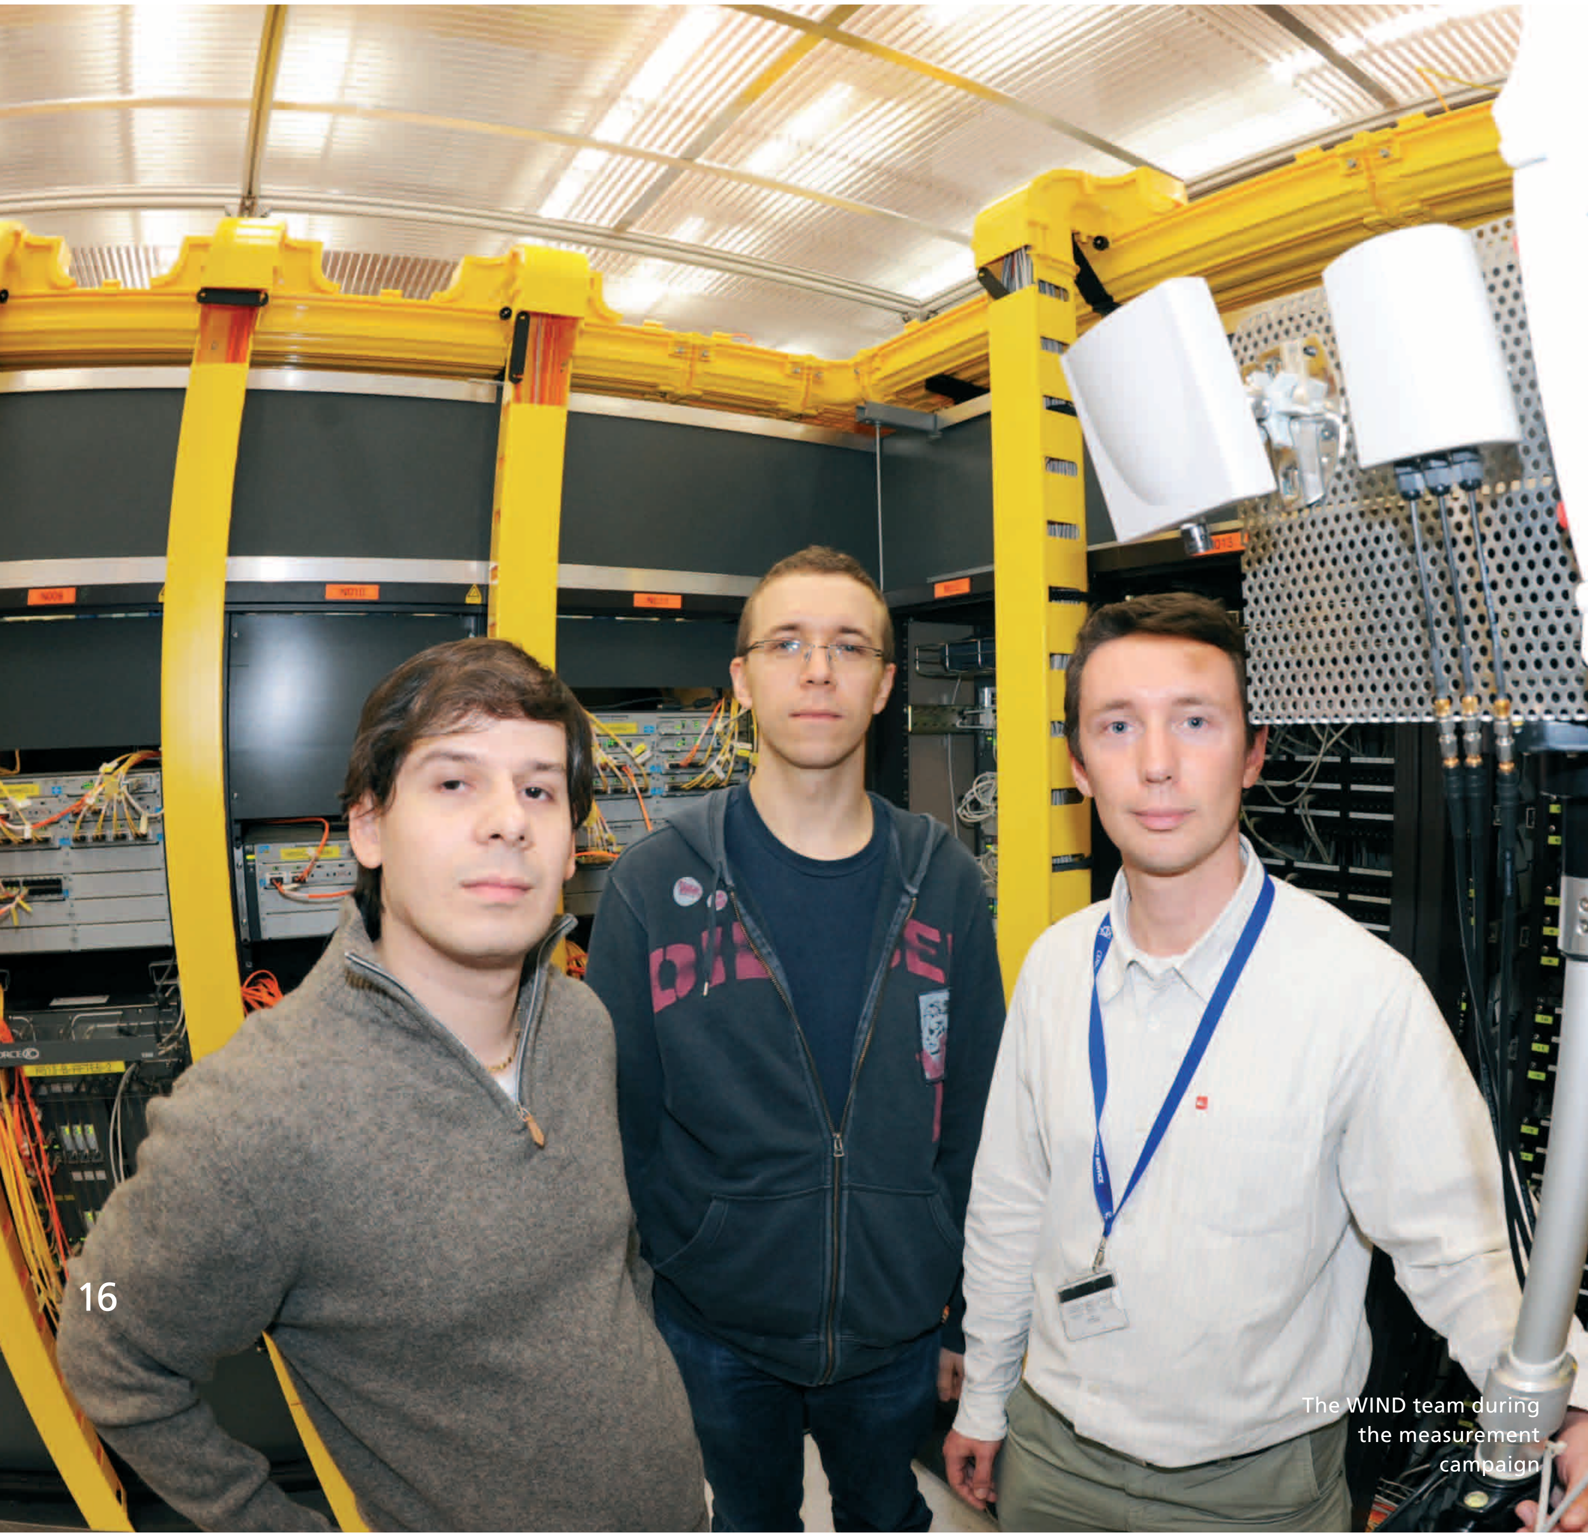

The WIND team during the measurement campaign

Following the success of CINBAD (CERN Investigation of Network Behaviour and Anomaly Detection), a project launched in 2007 in collaboration with HP which ended in 2010, HP and CERN started on a new joint research project in 2010, under the codename WIND (Wireless Infrastructure Network Deployment). WIND focuses on Wi-Fi networks in large campus environments.

Wi-Fi challenges

In a perfect world, the Wi-Fi networks would be as efficient as the wired ones. One could deploy such networks without arduous planning. The need for laborious site surveys would be significantly reduced. Once deployed, the lightweight control logic would adapt to the environment's changes in real-time to squeeze the best possible performance out of the given setup. The network would react to unpredictable congregations of users, radio interference, and other events. Furthermore, the network itself would be able to suggest deployment changes that could improve the performance, reliability, and overall user experience.

Reality is regrettably much harsher and Wi-Fi networks will probably never be as efficient as the wired ones. Indeed, painstaking planning and simulations precede the typical installation phase. Site surveys consume time and money and do not result in a perfect network as they provide only a static snapshot of the environment.Reinforced fire doors, microwave ovens, Bluetooth users to name but a few are very likely to cause interferences and, nevertheless, are generally not taken into account in site surveys.

As many network operators neglect the task of monitoring and tuning network parameters to adapt to new requirements and changes, a large number of deployments suffer from performance problems. Network operators cannot easily tell about the current health of the wireless network from the end-user point of view. No clear procedures help pinpointing potential problems or even find the root cause of existing ones. In many cases, the diagnosis and remedies are part of a random trial-and-error process. Last but not least, the end users do not understand and do not want to know about the intricacies of the radio wave propagation, wireless deployment, and operation. For all these reasons, the objective of providing wire-like quality proves to be extremely challenging.

CERN network

Today, CERN provides Wi-Fi service to more than 2 600 concurrent wireless users. This service is available in various types of environments, from big auditoria, to offices, warehouses, crowded control centres, and underground caverns. The installation currently consists of over more than 700 dual channel access points.

Challenges are foreseen as the current installation is expected to grow in the near future: in terms of coverage as the goal is to have most of the CERN campus covered, but more importantly also in terms of complexity. For instance, over the last few years, the number of wireless devices literally exploded. It is no longer uncommon for a single user to have several Wi-Fi clients such as a laptop and a smartphone, for instance. The number of devices alone is not a serious challenge but the way the devices are being used is definitely a critical one. Indeed, users expect Wi-Fi networks to be natural extensions of wired networks and thus demand similar performance (e.g. sufficient to stream HD video) and reliability (e.g. no disconnections, no roaming problems).

Reliability is taken for granted for wired networks and when performance is at risk, simple solutions such as adding more cables can solve the issue. Solving issues with wireless networks is far more complex. Indeed, the number of available wireless channels is limited, multiple interferers are present, and the radio propagation parameters are by no means constant. These are just a few examples of the limitations Wi-Fi network users are facing. More information on this topic is available on the CERN openlab website, in the presentations sub section, and in the Networking Competence Centre section.

From left to right: Milosz Hulboj, Sébastien Ceuterickx and Vlad Lăpădătescu working in one of the networking labs

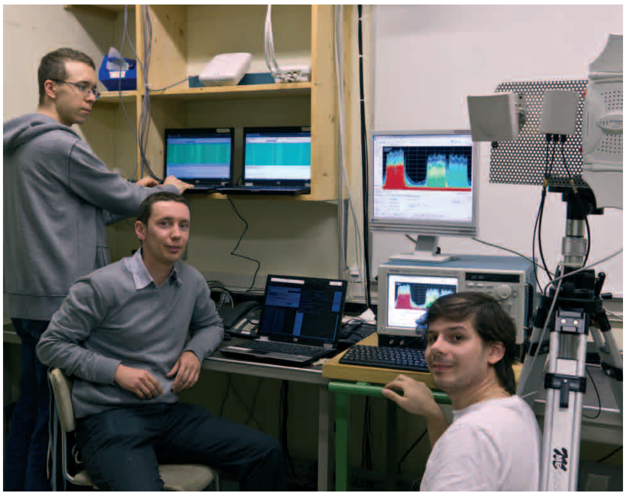

WIND project and first results

Given the variety of its premises and the number of Wi-Fi users on the large CERN campus, the Organization is keen on collaborating in a research activity focusing on wireless networks and provides the perfect environment to carry it out. The WIND project attempts to offer new algorithms, guidelines, and solutions for large-scale installations. In particular, it aims at monitoring, optimizing, and troubleshooting the existing installations as well as providing feedback for Wi-Fi design. Results of the joint research could possibly be incorporated into HP hardware and software to provide even more robust and efficient networking solutions.

In the first phase of the project and its first year, the WIND team investigated the state-of-the-art in the area of wireless networking. The results of this study demonstrated that the research on the operation and optimization of large-scale wireless networks is still in its infancy and that current standards have so far largely overlooked the issues related to monitoring and client control. The report called 'Wireless Control and Optimisation' was published at the end of 2010 and is available in the technical documents sub section of the CERN openlab website.

In parallel, the WIND project team worked on identifying the data sources that could provide information about the current state of the network. These data sources proved to be rather scarce, the only data sources available at the moment being the access points since the current standard does not allow for querying the clients for monitoring purposes. However, HP access points do an adequate job in providing detailed information about the associations, connected clients, and radio visibility by means of the SNMP (Simple Network Management Protocol).

Following this research, a dedicated system for collecting and storing the Wi-Fi related information was created. This scalable system combines the information from SNMP tables (e.g. client and visibility statistics) with the association events signalled by SNMP traps. Stored data is not only used for testing new algorithms but also for helping CERN's network engineers in their daily operations. New features can easily be integrated to the collector. It is envisioned to add an sFlow input module as soon as the access points start providing the additional information.

Several tools using the collected data have already been created. One of them attempts to detect the clients that are most likely to cause network slowdowns by correlating the transmission time of each client with the client's average transfer rate. Another tool tackles a different but equally important aspect of wireless networks by providing a visual representation of roaming events between different

High-level overview of the collection system

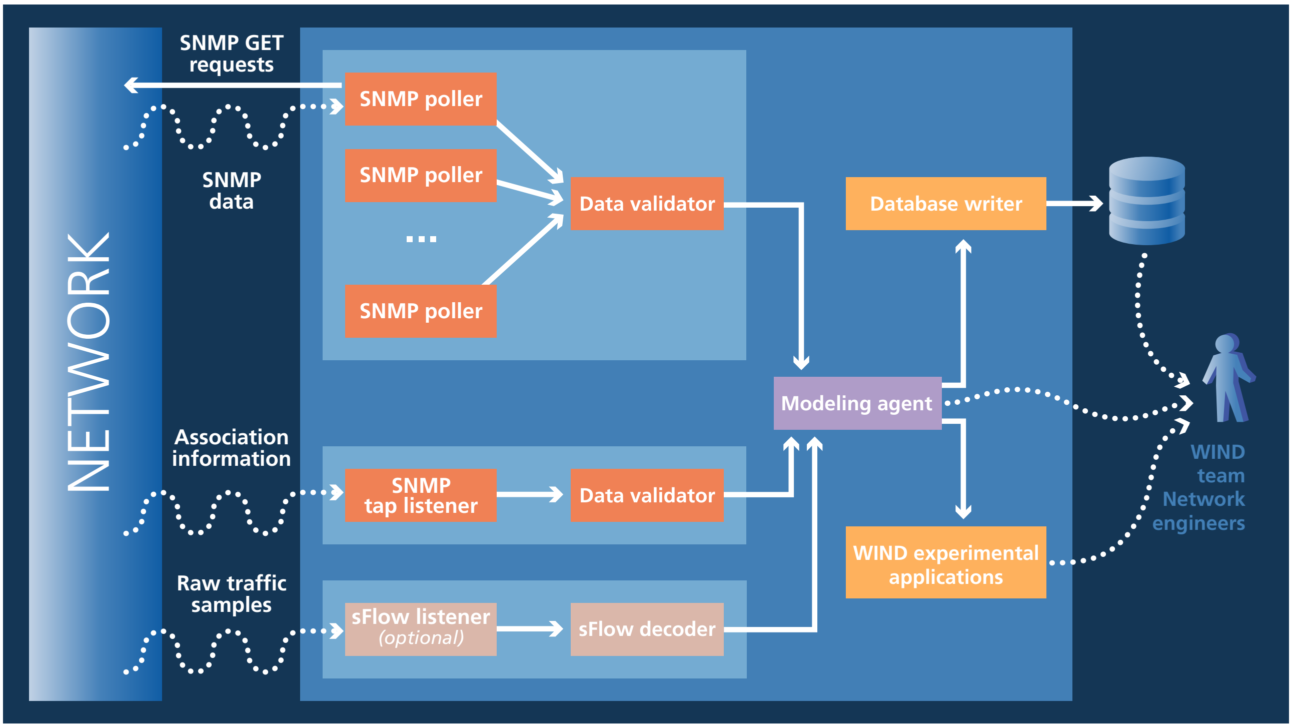

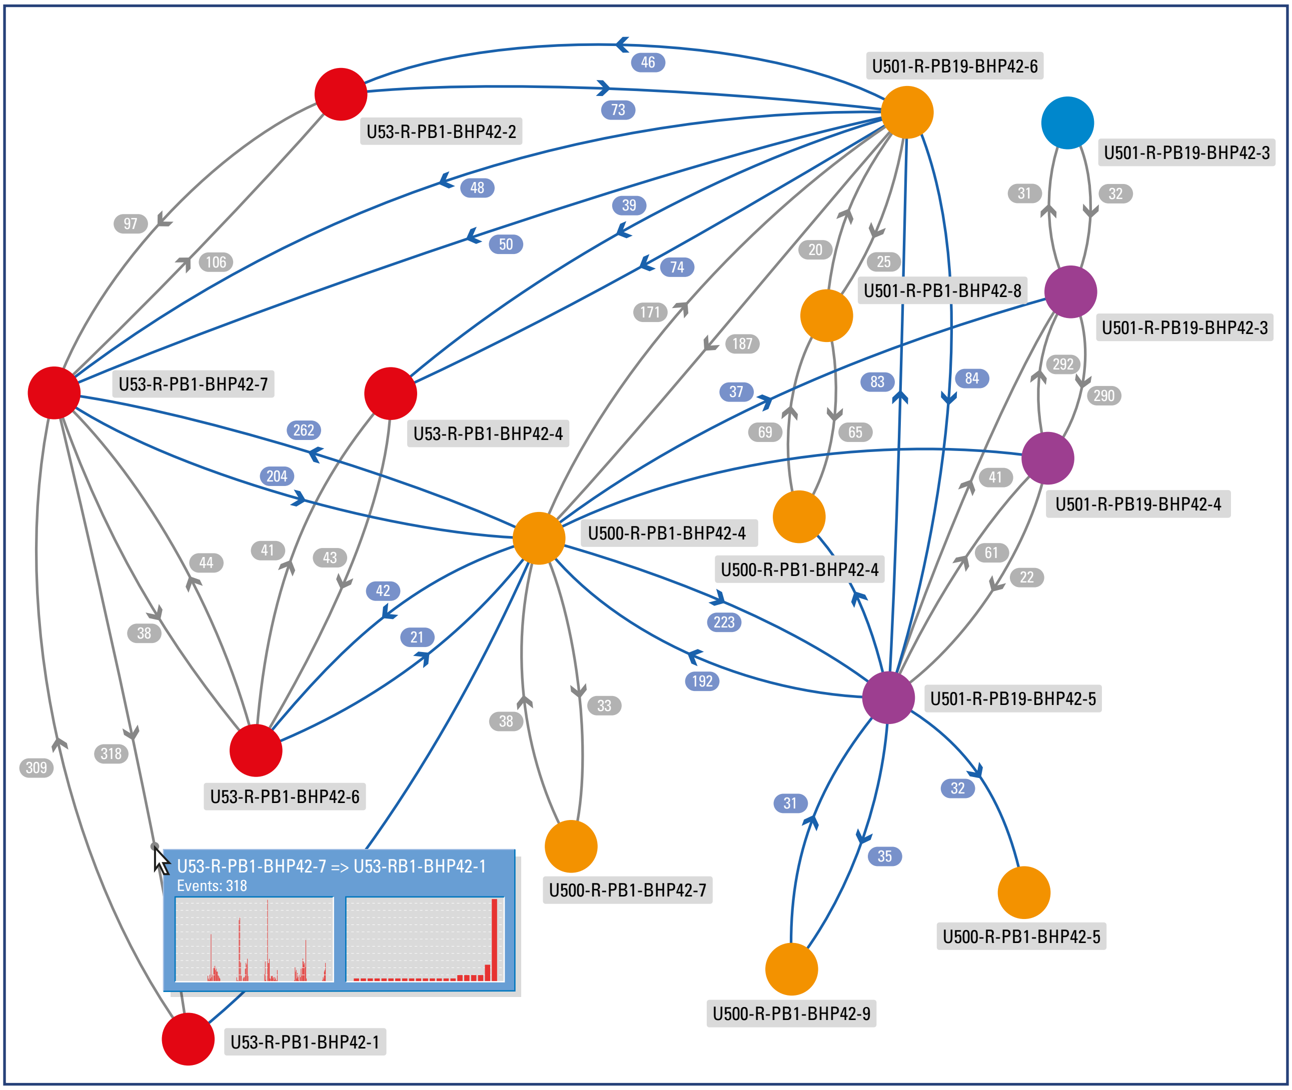

One of the tools developed to visualize roaming behaviour at CERN

access points installed on the campus. This helps to quickly identify roaming «hotspots» or troublesome devices.

During the second year of the project, the WIND team will continue working on algorithms for analysing the state of Wi-Fi networks and for detecting anomalies such as slow devices and deployment problems. Another aim of the WIND project will be to identify the key performance indicators that can reliably characterize the state of the network from the end-users' point of view. Finally, in the coming year, the WIND team will carry out some experiments on selected fragments of live networks to test chosen algorithms for optimizing performance (e.g. by forcing clients to associate with a given access point or on the 5 GHz band). The project promises interesting findings while the tools created up till now are already providing valuable information to CERN's network engineers.

The Platform Competence Centre

Front row, left to right: Alfio Lazzaro, Julien Leduc - Second row, left to right: Yngve Sneen Lindal, Andrzej Nowak, Sverre Jarp

In collaboration with Intel, the CERN openlab Platform Competence Centre (PCC) continues to address crucial fields such as thermal optimisation, multi-core scalability, application tuning and benchmarking. The strong emphasis on teaching and knowledge dissemination allows a broader audience to enjoy the fruit of the PCC's work. The last twelve months constituted yet another year of such intensive studies and development – with tangible effects.

Optimisation and benchmarks

A solid frame of reference is indispensable when evaluating computing systems in terms of their performance, thermal and energy consumption characteristics. Such studies allow for significant optimizations and are indispensable when making choices in the computing domain at CERN. Thus, the PCC maintains and continuously enhances its own set of benchmarks, sourced from various domains of High Energy Physics (HEP), as well as from industry standard suites.
HEPSPEC06, a subset that represents the C++ benchmarks from the SPEC2006 suite, still serves as a fundamental measurement tool. With its resource consumption pattern representative of typical High Energy Physics jobs, it yields comparable baseline numbers for whole families of systems evaluated by the PCC. Other sub-programs, such as the multi-threaded Geant 4 particle simulation prototype with the "FullCMS" detector geometry, the parallel maximum likelihood fitter based on the ROOT analysis package or the ALICE/CBM particle detector trackfitter and trackfinder, are extracted at a significant effort from large real-life software frameworks with millions of lines of code. These handpicked portions are customized and developed by the PCC team in order to allow for seamless throughput, scalability and performance studies of hardware. Extracting a benchmark can go as far as cutting out just a few lines or a few loops with surgical precision. Such "snippets" can later be easily recompiled, distributed and used to demonstrate, for instance, a specific point regarding compiler performance or correctness.

Myriads of systems

All of PCC's benchmarks are run on a wide variety of systems, which are often still in the alpha or beta stage when assessed. All this hardware needs to be deeply understood along with its wide spectrum of configuration options. This process is made easier by collaborating Intel engineers, such as Mark Myers, a server platform architect who visited CERN openlab in 2010. The feedback and numerous PCC-produced reports cover a vast number of indicators, ranging from performance and thermal properties to specific comments about the microarchitecture and its relation to the software that runs on it. Tested systems are split into several distinct families. The core of these evaluation activities is in the dual

As the LHC is entering a land of potential discoveries, a correct interpretation of the data collected by the experiments is paramount. Since the degree of belief used to claim a discovery has its foundation in the probability and statistical concepts deployed in the data analysis, there is a common effort in the HEP community to develop software tools for data analysis. CERN openlab team members (Yngve Sneen Lindal and Alfio Lazzaro in the picture) collaborate in this activity, in particular for developing efficient code for new computing platforms.

x86 processor segment, represented by the Intel Xeon® "EP" family. These have long been the main workhorse at the CERN Computing Centre, and are often chosen for their performance to cost ratios and predictable performance.

The four-socket segment, with Intel Xeon "EX" processors, is also of importance. These systems allow primarily for advanced scalability studies, with up to 80 hardware threads per tested system, as in the case of the Intel Xeon E7-4870 "Westmere-EX" based machines. The behaviour and characteristics of these many-core systems are closely monitored by the PCC and the Database Competence Centre (DCC), as they are good candidates for database deployment and also indicative of possible future developments in the dual-socket space.

A third group of systems includes additional and non-routine evaluations. One example is the desktop segment, which often makes a specific microarchitecture available before the dual-socket parts are launched. In this context, a "Sandy Bridge" microarchitecture-based Intel Core™ i7-2600k processor was evaluated, and was found to have an average 10% clock for clock improvement with respect to its "Westmere"-based predecessors, and a 9-19% performance per Watt improvement. Some existing "Montecito" Itanium servers are still in use by the PCC Compiler Project, which has been active since 2003, evaluating compiler performance on both the Intel64 and the IA64 architecture. This year, the AMS (Alpha Magnetic Spectrometer) experiment searching for antimatter and dark matter on the International Space Station, requested - and was granted - access to these Itanium servers for some of their software.

The PCC receives early versions of several of Intel's products. Installation, configuration and testing of the new hardware represent a significant part of the work of Julien Leduc (left). As an early adopter, the PCC team proactively sends feedback about the hardware: processors, firmware, platform architectures, power consumption, and peripherals such as SSDs. These evaluations aim at improving the efficiency of CERN's Tier-0 and the WLCG.

Since the CERN Computer centre facilities are severely limited both in terms of electrical input and cooling output, the need for intelligent power optimisation is paramount. Hence, the continued quest for better performance to power ratios prompts the investigation of some non-standard avenues. For example, the idea of power efficient microservers, notably supported years ago by the PCC, is being tried in an ongoing pilot project with Intel hardware. The PCC's interest in the low power Intel Atom architecture has been revived by new parts built in smaller processes. The PCC team is eagerly awaiting new offerings from Intel that will show further advantages with respect to their predecessors.

Finally, the PCC's keen interest in future technologies and the development of the x86 architecture has been fuelled even further by the emergence of the Intel® Many Integrated Core (MIC) architecture and development system. This 32-core part, with wide 512-bit vectors and 4-wide hardware threading, is not just an enticing development vehicle, but also an interesting indication of the direction in which the x86 microarchitecture is headed. The PCC maintains a high level of involvement in this project, being one of the first Intel collaborative teams worldwide to work with the part, and providing an on-stage testimonial during the official launch of the MIC technology. This activity, as well as all of the above, helps directly to prepare CERN software for future hardware, and give software developers at CERN a privileged peek into the future of computing.

Support software

Advanced performance and scalability studies would not be possible without the proper support software. The Intel Software tools have for a long time been amongst the market leaders in their segments, regularly bringing additional functionality, productivity and performance. The PCC's long involvement with these tools has led to their general availability at CERN and to an increased, common interest in their usage. The state of the art XE 2011 suite is composed, amongst others, of the Intel 12.0 Compiler, a next-generation performance tuning tool called VTune Amplifier and a next-generation correctness and memory checking tool called Inspector. All of the above have been exhaustively evaluated by the PCC, to the point where a senior CERN developer at CERN stated about VTune Amplifier that "this is the first tool that actually works out of the box with our software". This was good feedback to

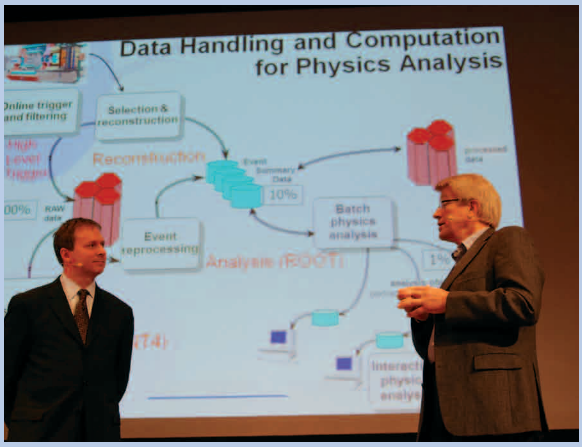

The PCC, given its broad experience and close relationship to Intel, was amongst the first external teams world-wide to receive a "Knights Ferry" (Many Integrated Core) accelerator for testing and evaluation. Consequently, Sverre Jarp (right) was able to comment on early experience when the chip was announced by Kirk Skaugen, Vice President and General Manager Intel Data Center Group (left), at the International Supercomputing Conference in Hamburg, in May 2010.

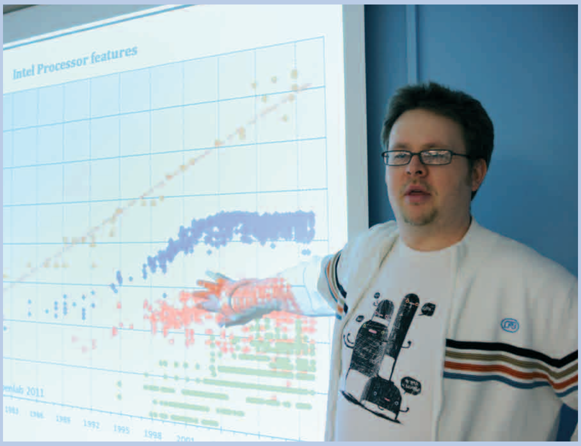

Performance monitoring is one of the essential areas of the PCC and one of the key working domains for Andrzej Nowak (in his teaching activities in the picture). CERN openlab co-developed and improved pfmon, a performance monitoring tool for Linux. These efforts evolved into a closer relationship with the developers of Intel's performance tuning tools such as PTU and VTune Amplifier XE –including extensive pre-release evaluations. Instant access to Intel experts allows openlab to efficiently use and teach these new technologies as soon as they are available.

the PCC team who had been actively participating in the testing of the product.

2010 was a very satisfying year for Intel Compiler studies. Over 50 compiler incidents were filed; many of them concerning the new 12.0 compiler which was closely examined before release. The new release was made available centrally and is now being picked up by developers across CERN, wishing to use it with their software for optimized performance.

Workshops, teaching and dissemination

The combined activities of the PCC regularly produce a large amount of expertise which is broadly shared with the scientific community. In this spirit, the PCC continues to offer training sessions both for intermediate and advanced programmers. Regular courses cover parallelism (twice a year) and computer performance and architecture (also twice a year). As the slide sets evolve with the technology that they describe, some attendees even decide to revisit the training. A participant satisfaction ratio of over 90% allows a steady stream of subscribers to be maintained. Since their inception, over 200 students have attended these workshops. In addition, visits of Intel engineers working on the tools mentioned in the previous section, as well as on other technologies, allowed the PCC to organise special seminars and "expert to expert" training sessions, attended by CERN's most senior programmers. Amongst such Intel visitors were David Levinthal –a world renowned x86 performance expert teaching low-level performance optimization-, and Levent Akyil –an experienced engineer driving the software tools effort demonstrating the new products. Jeff Arnold, a Senior Software Engineer in the Intel compiler team, continues teaching at regular workshops and gave two IT seminars on floating point and on the Intel C++ compiler.

Furthermore, the PCC participates actively in external conferences and symposiums. PCC representative were present and both ACAT 2010 and CHEP 2010, with four papers submitted and accepted for the latter alone. The Intel Developer Forum 2010 in San Francisco in September was a very effective mutual learning arena. Teaching sessions were held once more at international computing schools, such as the CERN School of Computing, which has been organised regularly for over 40 years.

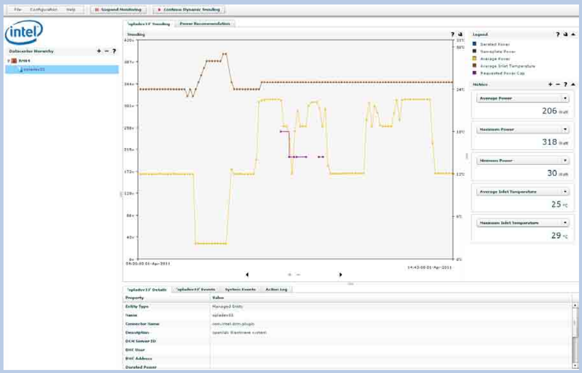

The team just started to evaluate Intel® Data Center Manager (DCM) on its compatible servers in the CERN Computer Centre. Intel DCM is a web service aimed at easing and optimising data centre management by monitoring two crucial metrics: power consumption and inlet temperature. It aggregates those metrics and takes action on the nodes to limit their individual power consumption, which allows maximising the data centre node density while preserving nominal conditions for all the servers.

Events and Outreach

As well as the excellent technical results that CERN openlab provides, the collaboration 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.

Top delegations from governments and industry regularly tour CERN. In addition, customer and press visits are organised by openlab partners. These groups are briefed about CERN openlab in a dedicated VIP meeting room known as the CERN openlab openspace. When visiting CERN openlab to discuss common projects with the team, some of our guests also give computing seminars in the IT auditorium.

Oracle Visit to CERN, Dr. C. Gregory Doherty, Vice President, Collaboration Suite, Server Technology Division, Oracle, 8-10 June 2010

Siemens Visit to CERN, Dr. Hermann Requardt, member of the Managing Board of Siemens AG and CEO of the Healthcare Sector, and his Management Team, 3 August 2010

Intel Visit to CERN, Dr. Richard Dracott, General Manager of High Performance Computing in the Server Platforms Group, Dr. Rajeeb Hazra, Director of Supercomputing Architecture and Planning, 13 October 2010

Intel Executive Summit, **Open Data Center Alliance (ODCA) Launch,** CERN's Globe for Science and Innovation, 26 October 2010

Intel Press Event and Visit, Open Data Center Alliance (ODCA) Launch, CERN's Globe for Science and Innovation, 27 October 2010

SITRA (Finnish Innovation Fund) Visit, CERN, 28 October 2010

Oracle-CERN Get-Together, CERN, 23 November 2010

CERN openlab IV Preparatory Workshop (followed by Meetings at CERN), CERN Representatives and all CERN openlab Partners, Lavey-les-Bains - CH 5-7 December 2010

Nature and NPR Visits and Interviews, CERN, 8-9 December 2010

Helsinki Metropolia University of Applied Sciences Visit, CERN, 15 December 2010

City of Meyrin Visit, CERN, 20 January 2011

Siemens Visit to CERN, Dr. Albert Wick, Vice President of Advanced Technology and Standards , and his Management Team, 21 March 2011

Intel Labs Europe Directors Visit, CERN, 16 May 2011

Intel Exascale Research Summit, CERN's Globe for Science and Innovation, 18 May 2011

Seminars by CERN openlab Guests

Concurrent Collections (CnC): A new approach to parallel programming, CERN, 7 May 2010 Given by Dr. Kathleen Knobe, Scalable Tools Engineer, Intel Corporation, as part of her visit.

Software Aspects of IEEE Floating-Point Computations for Numerical Applications in High Energy Physics, CERN, 11 May 2010 Given by Dr. Jeff Arnold, Senior Software Engineer, Intel Corporation, as part of his visit.

Oracle JRockit: Advanced JVM technologies, CERN, 14 June 2010 Given by Tuva Palm, Java Runtime Platform Group, Oracle, Erik Bergenholtz and Mark Prichard, Java Platform Group, Oracle, as part of their visit.

Siemens: Innovation Concepts in Healthcare, CERN, 3 August 2010 Given by Prof. Dr. Hermann Requardt, member of the Managing Board of Siemens AG and CEO of the Healthcare Sector, Siemens AG, as part of his visit.

Design Patterns: Establishing a Discipline of Parallel Software Engineering', CERN, 27 September 2010 Given by Dr. Tim Mattson, Director of Supercomputing Architecture and Planning, Intel, as part of his visit.

Advanced Features of Intel C++ Composer XE for Linux, CERN, 18 February 2011 Given by Dr. Jeff Arnold, Senior Software Engineer, Intel, as part of his visit.

Intel ISEF Students Visit, CERN, 12 June to 18 July 2010

In 2010, for the second year, CERN hosted the visit of Intel International Science and Engineering Fair (ISEF) award winners. The 12 pre-college students had won the "CERN Special Award" at the Intel ISEF in the spring 2010 in San Jose, USA. Intel ISEF is the world's largest international pre-college science competition and annually provides a forum for more than 1,600 high school students from over 60 countries to showcase their independent research. The students spent five days at CERN, visited the LHC facility and enjoyed presentations from various prominent scientists.

CERN Special Award winners and Alberto Pace (CERN) at the Intel ISEF in San Jose, USA

Publications and Presentations

CERN openlab results have been disseminated in a wide range of international conferences, listed below. These publications and presentations can be consulted on the openlab website, as well as a large number of press articles coming out in the general press, IT-specific press and on the Web.

Presentations:

- **L. Canali/CERN**, Overview of the CERN DB Services for Physics, Orcan Swedish Oracle Users Group Conference, Stockholm, Sweden, 18-19 May 2010
- **L. Canali/CERN**, Data Life Cycle Management at CERN with Oracle, Stockholm, Sweden, 18-19 May 2010
- **T. Cass/CERN**, Data and Database Challenges at CERN, Oracle HPC Consortium, Hamburg, Germany, 29 May 2010
- **S. Jarp/CERN**, **K. B. Skaugen/Intel**, Keynote: HPC Technology – Scale-up & Scale-out, International Super Computing Conference (ISC'10), Hamburg, Germany, 30 May 2010
- **R. Dracott/Intel**, The Path to Exascale – Extending Intel's HPC Commitment, High Performance Computing Symposium (HPCS), Toronto, Canada, 6 June 2010
- **T. Tauber/CERN,** openlab Summer Student, Virtualization: Beyond the Physical Machine Limits, CERN Summer Student lecture, 18 August 2010
- **A. Nowak/CERN**, Protons, Punched Cards and Clouds – Computing at CERN and the Intel Angle, Intel Developper Forum (IDF), Intel Campus, Santa Clara, USA, 16 September 2010
- **C. Garcia Fernandez/CERN**, Application Aware Virtualization, Oracle OpenWorld, San Francisco, USA, 21 September 2010
- **E. Grancher/CERN**, CERN Achieves Database Scalability and Performance with Oracle and NetApp, Oracle OpenWorld, San-Francisco, USA, 23 September 2010
- **S. Jarp/CERN,** How to Harness the Performance Potential of Current Multicore CPUs and GPUs, Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, 18 October 2010
- **CERN openlab Platform Competence Centre team/CERN,** presented by A. Lazzaro/CERN, Evaluating the Scalability of High Energy Physics Software and Multicore Hardware, Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, 18 October 2010
- **CERN openlab Platform Competence Centre team/CERN**, presented by A. Lazzaro/CERN, Manycores Accelerators Evaluation at CERN openlab, Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, 21 October 2010
- **CERN openlab Platform Competence Centre team/CERN,** presented by A. Lazzaro/CERN, Maximum LikelihoodFits on GPUs, Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, 21 October 2010
- **L. Canali/CERN, D. Wojcik/CERN,** ACFS Under Scrutiny, UK Oracle User group (UKOUG) Conference 2010, Birmingham, U K, 30 November 2010
- **L. Canali/CERN, J. Wojcieszuk/CERN,** Data Lifecycle management Challenges and Techniques, a User's Experience, UK Oracle User Group (UKOUG) Conference 2010, Birmingham, UK, 1 December 2010
- **A. Lazzaro/CERN**, High Performance Computing on Commodity PCs, University of Milan, Department of Physics, Milan, Italy, 14 January 2011
- **A. Nowak/CERN,** Evaluating Program Correctness and Performance with new Software Tools from Intel, IT Technical Forum, CERN, 18 March 2011
- **A. Lazzaro/CERN,** CERN openlab: Optimization, Parallelization, Evaluation, Ireland's High Performance Computing Centre, Dublin, Ireland, 5 April 2011
- **CERN openlab Platform Competence Centre team/CERN,** presented by S. Jarp/CERN, KNF experience in CERN openlab, Intel EMEA High Performance Computing Round-table, Paris, France, 7 April 2011
- **CERN openlab Platform Competence Centre team/CERN,** presented by A. Lazzaro/CERN, Evaluation of Likelihood Functions for Data Analysis on Graphics Processing Units, The 12th International Workshop on Parallel and Distributed Scientific and Engineeing Computing (IPDPS11) , Anchorage, Alaska, USA, 20 May 2011
- **W. von Rüden/CERN,** CERN and the LHC project: a short introduction, Intel Labs Europe Directors visit, CERN, 16 May 2011
- **S. Jarp/CERN,** Today's Worldwide Computing Grid for the Large Hadron Collider (WLCG): A Petascale Facility - Moving to Exascale?, Intel Exascale Leadership Conference, CERN Globe for Science and Innovation, 18 May 2011

Publications:

- **S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN, F. Pantaleo/CERN openlab Summer Student**, Parallelization of Maximum Likelihood Fits with OpenMP and CUDA, Conference on Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, October 2010
- **S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN,** The Breaking Point of Modern Processor and Platform Technology, Conference on Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, October 2010
- **S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN,** How to Harness the Performance Potential of Current Multicore Processors, Conference on Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, October 2010
- **S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN,** Evaluating the Scalability of High Energy Physics Software and Multicore Hardware, Conference on Computing in High Energy and Nuclear Physics (CHEP), Taipei, Taiwan, October 2010
- **S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN, F. Pantaleo/CERN openlab Summer Student,** Evaluation of Likelihood Functions for Data Analysis on Graphics Processing Units, The 12th International Workshop on Parallel and Distributed Scientific and Engineering Computing (IPDPS11), Anchorage, Alaska, USA, May 2011

Posters:

- **A. Lazzaro/CERN,** Evaluation of Parallel Applications used in High Energy Physics, Intel European Research and Innovation Conference (ERIC), Braunschweig, Germany, 14 September 2010
- **A. Nowak/CERN,** The Breaking Point of Modern Processor and Platform Technology, CHEP, Taipei, Taiwan, 20 October 2010

CERN openlab Reports:

- **E. Grancher/CERN, S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN**, Evaluation of the Intel nehalem-EX Server Processor, May 2010
- **T. Tauber/Summer Student**, Automation and Optimisation in IT-DB OracleVM Virtualisation Systems, August 2010
- **C. N. Micu/Summer Student**, Shared library implementation for the BDII Information System, August 2010
- **D. A. Popescu/Summer Student**, Performance Monitoring of the Software Frameworks for the LHC Experiments, August 2010
- **M. Sandu Popa/Summer Student**, Automating Management for Weblogic Server on JRockit VE, August 2010
- **G. Bajrami/Summer Student**, TRoIE (Test-bench for Robustness of Industrial Equipments) Reporting System, August 2010
- **J. Wu/Summer Student**, BOSS and LHC Computing using CernVM and BOINC, August 2010
- **A. Beche/Summer Student**, Implementation of Corrective Actions for Policy Violations in Oracle Enterprise Manager, August 2010
- **R. Veznaver/Summer Student**, Design and implementation of Probes for the WLCG SAM Framework, September 2010
- **M. Hulboj/CERN, V. Lăpădătescu/CERN,** Wireless Control and Optimization, October 2010

Education

CERN openlab is a structure designed to create knowledge. This is done through the evaluation of solutions as well as genuine research and development of IT technologies. This knowledge is then disseminated through multiple channels. One of the major ones is the publication of reports and articles, as described in the preceding section, but this may be viewed as passive dissemination. The openlab education programme, which provides active dissemination, is currently implemented through several lines of actions.

Workshops or seminars are regularly organised at CERN on advanced topics directly connected to the openlab projects. This year more than 180 participants attended the tried-and-tested two-day workshop formula on "Multithreading and Parallelism" and "Computer Architecture and Performance Tuning". These workshops have a special feature: they involve a mix of lecturers from both industry and CERN, thus exemplifying the openlab principle of two-way knowledge transfer through active collaboration. Several of the workshops combine hands-off theory with hands-on practice. In addition to the regular quarterly courses, special courses have also been organised this year for advanced CERN users. These classes touched on numerous topics, including future tools and optimisations and were taught by Intel specialists. However, openlab experts also contribute to off-site education activities such as the CERN School of Computing, where eight hours of lectures and exercises were delivered in 2010. These lectures and workshops are listed below.

- **Multithreading and Parallelism Workshop**, 4-5 May 2010, CERN, J. Arnold/Intel, S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN
- **3rd CERN openlab/Intel Special Workshop**, 6-7 May 2010, CERN, J. Arnold/Intel, K. Knobe/Intel, D. Levinthal/Intel
- **International Tracking Workshop,** 9-11 June 2010, GSI, Darmstadt, Germany. One lecture: Harnessing Future Hardware, S. Jarp/CERN,
- **2nd Workshop on Adapting Applications and Computing Services to Multicore and Virtualization,** 21-22 June, CERN. One lecture: Performance Monitoring Tools – Scaling to Manycore, A. Nowak/CERN
- **2nd Workshop on Adapting Applications and Computing Services to Multicore and Virtualization,** 21-22 June, CERN. One lecture: Adapting High Energy Physics Applications to run on GPU, A. Lazzaro/CERN
- **Advanced Session on Performance Monitoring**, 21-22 July 2010, CERN, D. Levinthal/Intel
- **CERN School of Computing**, August 2010, Uxbridge, United Kingdom. Three lectures: Computer Architecture & Performance Tuning, S. Jarp/CERN, A. Nowak/CERN.
- **CERN School of Computing**, August 2010, Uxbridge, United Kingdom. Two lectures: Networking QoS and Performance, F. Flückiger/CERN.
- **Computer Architecture and Performance Tuning Workshop,** 22-23 September 2010, CERN, J. Arnold/Intel, S. Jarp/CERN, J. Leduc/CERN, A. Nowak/CERN
- **Multithreading and Parallelism Workshop**, 10-11 November 2010, CERN, J. Arnold/Intel, S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN
- **Distributed Database Operations Workshop**, 16 November 2010, CERN. One lecture: Experiments Databases, L. Canali/CERN
- **Intel S/W products Workshop**, 29-30 November 2010, CERN, L. Akyil/Intel, H. Pabst/Intel
- **Computer Architecture and Performance Tuning Workshop**, 15-16 February 2011, CERN, J. Arnold/Intel, S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN
- **ESA-GAIA DB Workshop**, 15 March 2011, ISDC, Versoix, Switwerland. One lecture: CERN IT-DB Deployment, L. Canali/CERN
- **Multithreading and Parallelism Workshop**, 9-10 May 2011, CERN, J. Arnold/Intel, S. Jarp/CERN, A. Lazzaro/CERN, J. Leduc/CERN, A. Nowak/CERN

Intel and CERN openlab have organised many workshops over the years. Jeff Arnold (in his teaching activities in the picture), Senior Software Engineer at Intel and Compiler Expert, participated already in ten of them since the beginning of the third phase of CERN openlab. Our warm thanks go to Jeff for his outstanding contribution.

These direct training activities are complemented by the CERN openlab Student Programme, which itself is a genuine educational undertaking. 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. A total of 132 students have participated so far. In 2010, the programme accepted 15 computer science and physics students of 12 nationalities 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 WLCG.

Visits were organised to the CERN Computer Centre, the CERN Control Centre, the ATLAS experiment, the LHC magnets test hall, the anti-matter factory and the LINAC-LEIR accelerators. In addition, the students toured the Geneva Observatory in Versoix. A dedicated lecture series for the students was given by CERN experts. Several of this year's students were co-funded by CERN and Intel, Oracle or Siemens.

CERN openlab Summer Student Programme Teaching Series, July-August 2010:

- Control System Cyber-Security, S. Lüders/ CERN
- Software Security, S. Lopienski/ CERN
- Web Application Security, S. Lopienski/ CERN
- Physics Computing at CERN, H. Meinhard/ CERN
- Size and complexity of the CERN network, R. Jurga/ CERN
- Worldwide LHC Computing Grid overview, M. Schulz/ CERN
- Overview of the gLite middleware, M. Schulz/ CERN
- Data Reliability at CERN and ideas on how to improve it, A. Pace/ CERN
- Invenio Technology, T. Simko/ CERN

The CERN openlab summer students 2010 are listed below, with home institute and project topic:

- Gazmend BAJRAMI, Albania, Gjovik University College, Norway, "The TRoIE (Test-bench for Robustness of Industrial Equipments) Reporting System"
- Imon BANERJEE, India, National Institute of Technology, Durgapur, India, "Performance Benchmarking"
- Alexandre BECHE, France, SUPINFO, Grenoble, France, "Implementation of corrective actions for policies violations in Oracle enterprise manager"
- Cornel MICU, Austria, Salzburg University of Applied Science, Austria, "A shared library for information system access"
- Nikolche MIHAJLOVSKI, Macedonia, University of Skopje, Macedonia, "New template engine for Indico"
- Felice PANTALEO, Italy, Instituto Nazionale si Fisica Bari, Italy, "Graphics Processing Units as accelerators for data analysis applications"
- Diana-Andreea POPESCU, Romania, Politechnica University of Bucharest, Romania, "The extraction and performance monitoring of the latest benchmark representatives of the four main LHC experiments' software"
- Cristian REGEP, Romania, University of Edinburgh, UK, "CERN School of Computing (CSC) website management and technical assistance"
- Marius George SANDU POPA, Romania, Politechnica University of Bucharest, Romania, "Automating management for web logic server on JRockit-VE"
- Fabio SOUTO MOURE, Spain, University of Vigo, Spain, "CERN Document Server (CDS) development"
- Tomas TAUBER, Czech Republic, University of Edinburgh, UK, "Automation and optimisation in IT-DB OracleVM virtualisation systems"
- Robert VEZNAVER, Croatia, University of Zagreb, Croatia, "Grid Technologies, Tools for Operations and Monitoring"
- Jie WU, China, IHEP, Beijing, China, "CERN virtual machine technology for clusters and distributed computing"
- Wenjing WU, China, IHEP, Beijing, China, "CERN virtual machine technology for clusters and distributed computing"
- Boris YORDANEV ROBEV, Bulgaria, University of Sofia, Bulgaria, "Automation of CERN accounting data collection for WLCG"

CERN openlab students at CERN, August 2010

Positioning CERN openlab activities at CERN

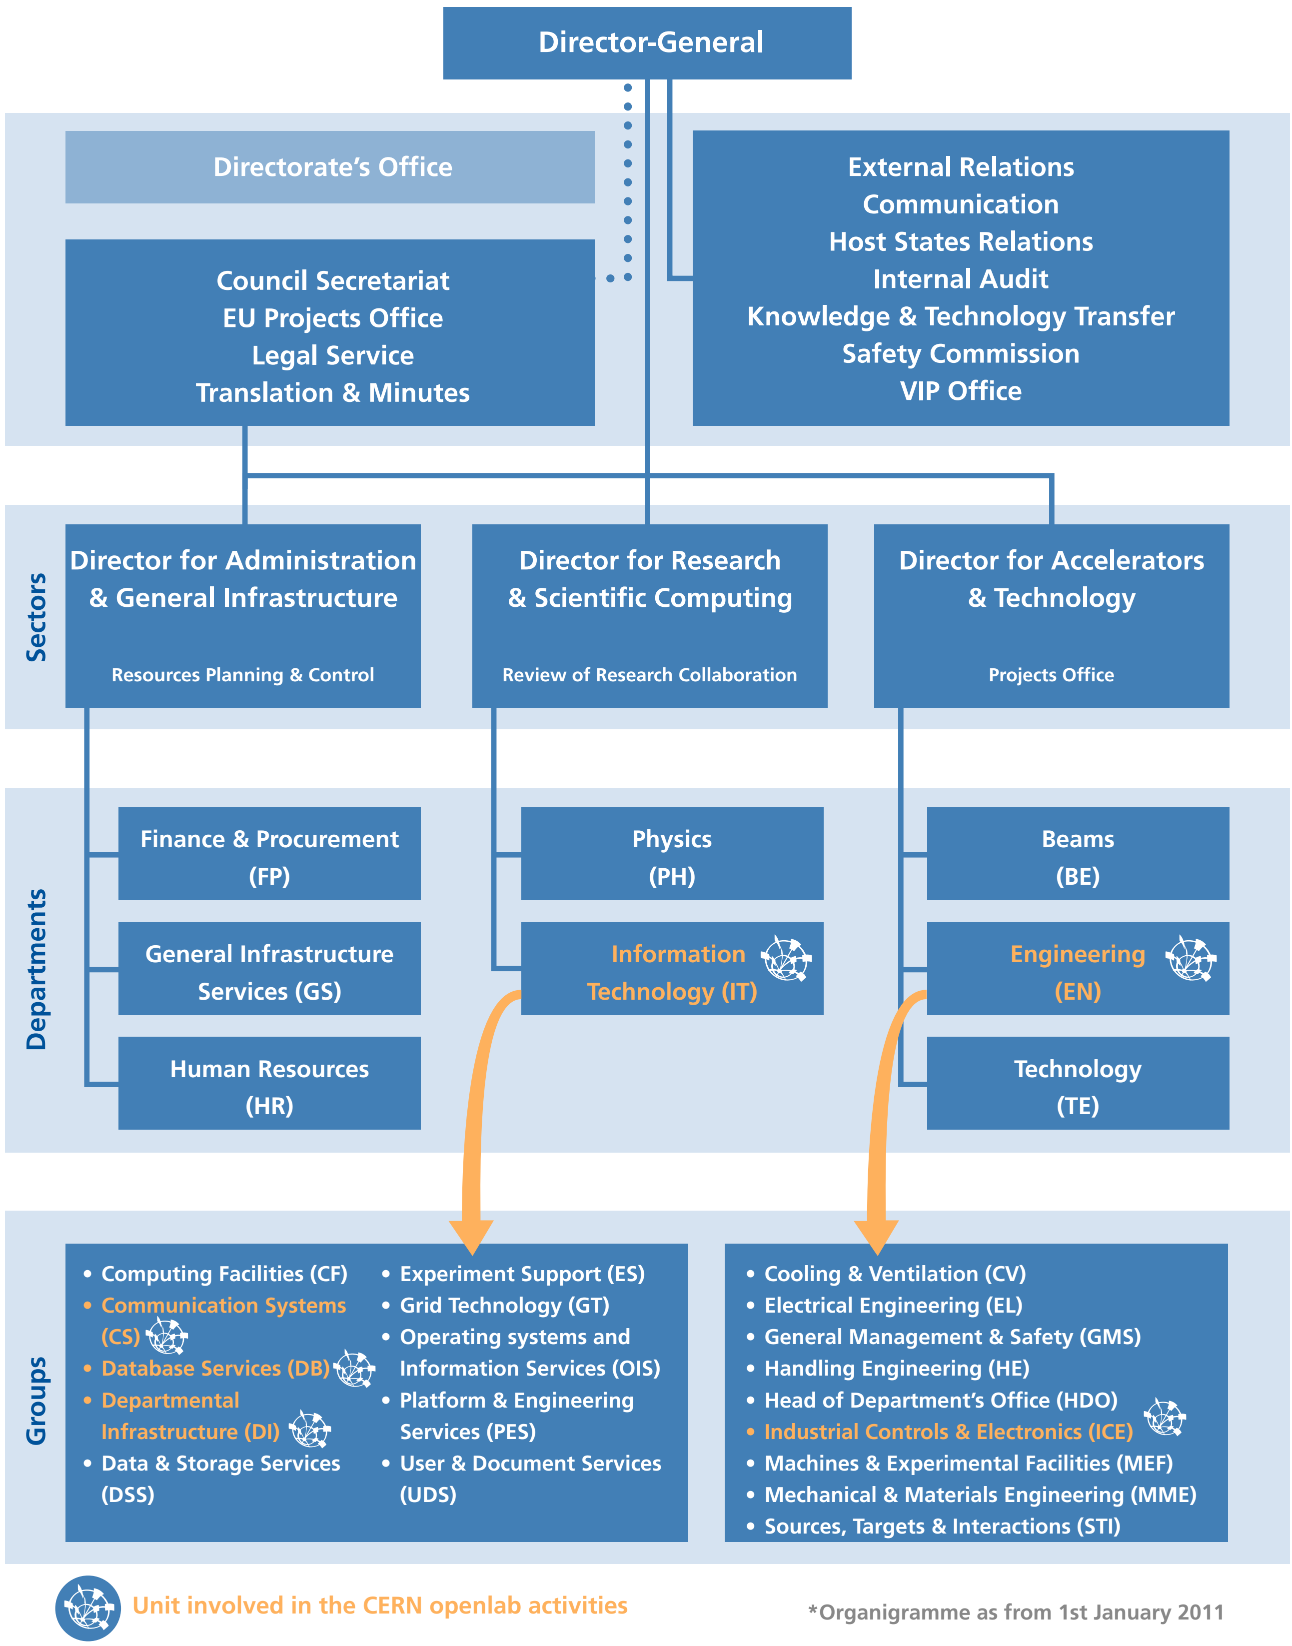

*Organigramme as from 1st January 2011

The Board of Sponsors

CERN openlab celebrated its 10th birthday this year, marking the occasion during its annual Board of Sponsors meeting, in the presence of the CERN Director-General, the partners, and past and present team members.

THE FUTURE

The LHC's performance this year has been fantastic, achieving the target luminosity for the year by June. Thanks to this stellar performance, the present run has been extended, and will continue throughout 2011 and 2012.

Making discoveries in particle physics is often a long and painstaking process, requiring large quantities of data to be carefully sifted for rare processes. In June, the amount of data accumulated by LHC experiments ATLAS and CMS clicked over from 0.999 to 1 inverse femtobarn, signaling another important milestone in the experiments' quest for new physics. The number represents a quantity physicists call integrated luminosity, which is a measure of the total number of collisions produced. One inverse femtobarn equates to around 70 million million collisions, and in 2010 it was the target set for the 2011 run. That it has been achieved just three months after the first beams of 2011 is testimony to how well the LHC is running.

On 21 April, the LHC set a new world record for beam intensity at a hadron collider when its beams collided with a peak luminosity of 4.67 × 1032 m-2s-1. This exceeds the previous world record of 4.024 × 1032 cm-2s-1, which was set by Fermilab's Tevatron collider in 2010, and marks an important milestone in LHC commissioning.

In the meantime, 54 physics papers have been published based on 2010 data, and more than 1140 conference presentations were made by the LHC experiments.

The speed with which the experiments have been able to analyse the data is unprecedented. The Worldwide LHC Computing Grid, which links up computer centers around the world, has proved itself well up to the task, routinely processing up to 200,000 physics analysis jobs concurrently.

The data rates achieved exceeded those planned by a factor 4, requiring extreme agility of the computing services. The LHC is on track to at least double the amount of data delivered to the experiments so far by the end of the year. Many of these achievements are thanks to continuing collaborations between CERN and its partners, including the openlab project. Preparations for the next phase of development have already begun, with existing partners exhibiting enthusiasm, and a number of potential new contributors who could accompany openlab into new domains of research and development. CERN openlab IV will officially begin in January 2012, with more intense cross-partner activities and new projects which will build on the latest technological developments.

Looking ahead, the redesign of the LHC high-current interconnects is now complete, and the machine will be prepared for running at 7 TeV per beam during the shutdown in 2013-14. The interaction of openlab's different communities will help CERN and the LHC community prepare for the increases in data rates forecast from both accelerator and experiments over the next three years.

CERN openlab IV preparatory workshop participants in Lavey Les Bains (December 2010)

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 is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. Additional information about Intel is available at www.intel.com.

Oracle provides the world's most complete, open, and integrated business software and hardware systems, with more than 370,000 customers—including 100 of the Fortune 100—representing a variety of sizes and industries in more than 145 countries around the globe. Now, with Sun server, storage, operating-system, and virtualization technology, Oracle is the only vendor able to offer a complete technology stack in which every layer is integrated to work together as a single system. 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.

For more information

Executive Contact:
Bob Jones, Head of CERN openlab,
Bob.jones@cern.ch

Technical Contact:
Sverre Jarp, openlab Chief Technology Officer,
Sverre.Jarp@cern.ch

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.

www.cern.ch ▪ www.youtube.com/cern ▪ www.twitter.com/cern

CERN openlab

www.cern.ch/openlab

PARTNERS

hp invent

intel

ORACLE

SIEMENS

CERN openlab

Annual Report
May 2010 – May 2011