Annual Report
2013

CERN openlab




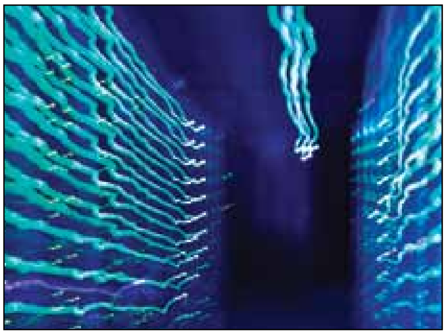

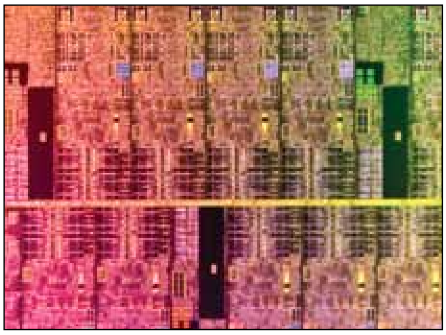

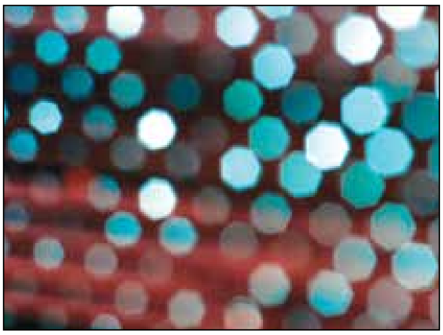
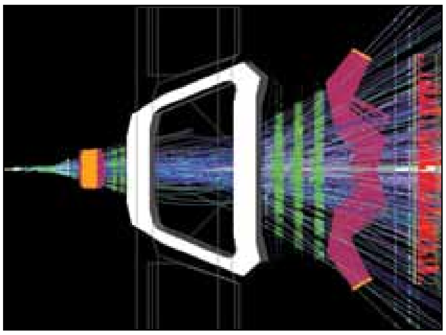

Contents

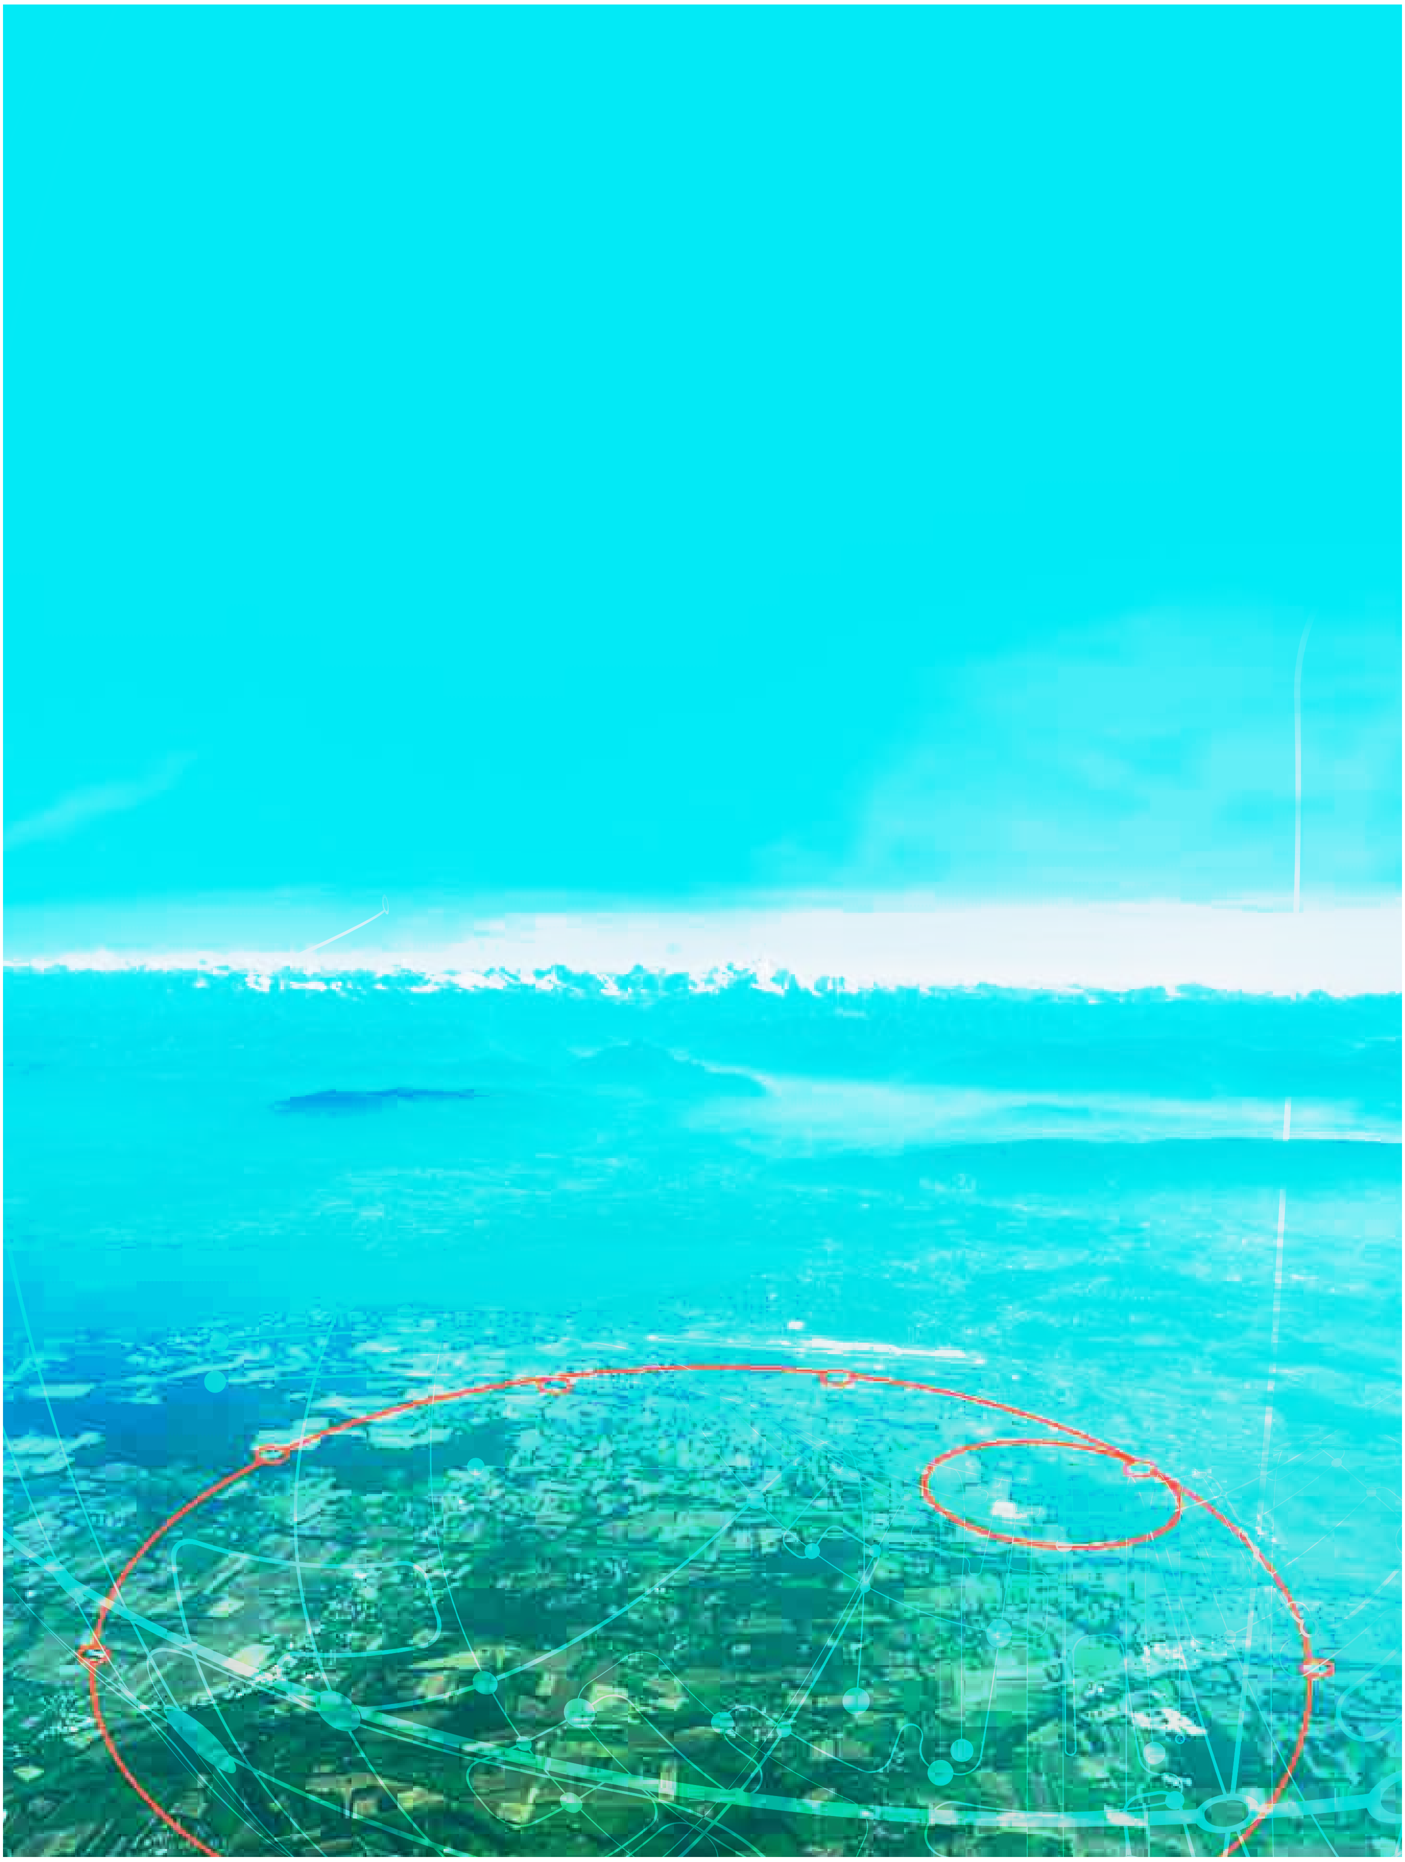

A Word from the DG

CERN openlab entered the second year of its fourth phase in 2013, and continues to go from strength to strength. Founded in 2001 to develop the innovative IT systems needed to cope with the unprecedented computing challenges of the LHC, openlab unites science and industry at the cutting edge of research and innovation.

CERN openlab has generated a long list of technical achievements over the years and played a vital role in the recent discovery of the Higgs boson. This flexible collaboration enables partnerships to flourish when time is right for both CERN and industry. This year, new projects started with Rackspace, which joined CERN openlab as a contributor, and Yandex, which joined as an associate. Meanwhile, the ViSION (Virtual Services In OpenFlow Networks) project with HP came to an end. Huawei, which initially joined CERN openlab in 2012 as a contributor for one year, also became a full partner in the Storage Architecture Competence Centre with a three-year programme of work.

The results presented in this report provide tangible evidence of the virtuous circle linking basic and applied science. The work carried out in 2013 with these newcomers, as well as with partners HP, Intel, Oracle and Siemens, has borne fruits that will be instrumental in ensuring our IT capacity meets the challenges of the LHC new run in 2015. I therefore take this opportunity to thank all our current and past industry members for their support.

Education is an essential part of CERN openlab's DNA and I would like to highlight the fundamental role played by the young researchers hired by CERN and funded by the partner and contributor companies. These young researchers and students demonstrate remarkable curiosity and ingenuity; they truly are the catalyst for many creative ideas within CERN openlab and beyond.

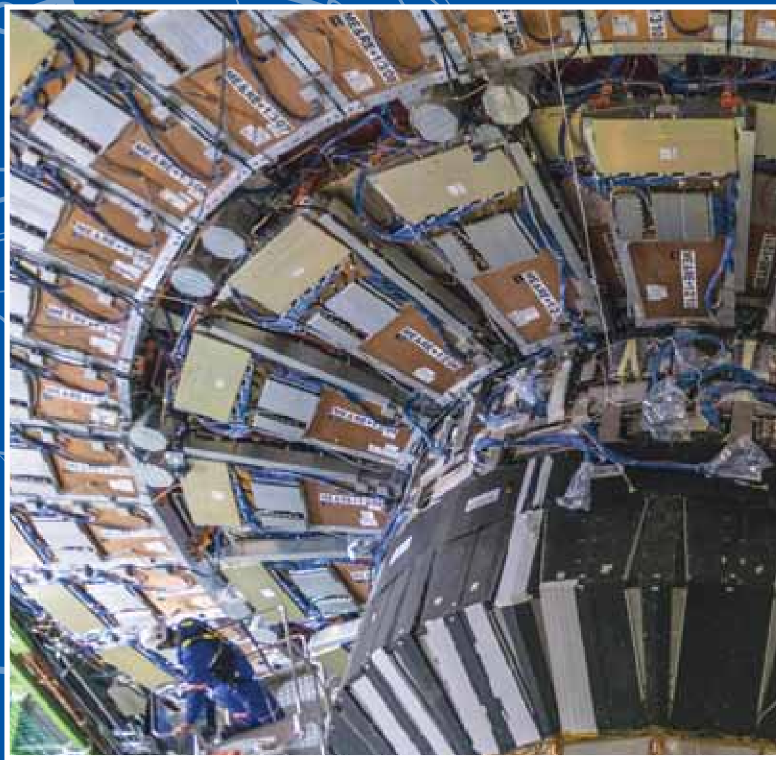

The Context

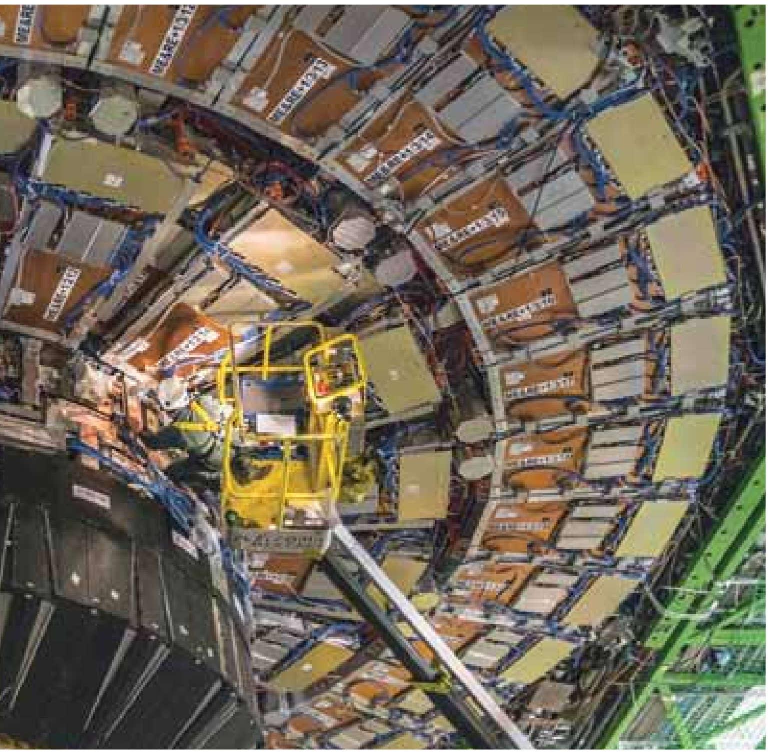

2013 saw a busy programme of maintenance and consolidation across the accelerator chain (on this picture an engineer is working on the CMS detector).

A pivotal year

First three-year LHC running period concludes on a high note.

The Large Hadron Collider (LHC) project started about 30 years ago, with the aim of preparing the next major phase in the on-going 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 buried 50-175 m below ground, it accelerates particles to more than 99.9% the speed of light to energies never reached before in a particle accelerator. Some 9600 superconducting magnets, operating at -271.3°C (just 1.9 degrees above absolute zero), 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-megapixel 3D cameras, record these collisions up to 600 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 LHC's 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. In April 2012, the LHC experiments could start taking data at the new collision energy of 8 TeV (up from 7 TeV in 2011).

By 4 July 2012, the ATLAS and CMS collaborations had crunched enough data to announce that they had both observed a new particle, consistent with the long-sought Higgs boson. The discovery, confirmed this year at the Moriond Conference in March, was made possible by the excellent performance of the LHC, the experiments' computing systems, and the Worldwide LHC Computing Grid (WLCG). The year 2013 ended on a high note with the Nobel Prize in Physics being awarded jointly to François Englert and Peter Higgs "for the theoretical discovery of

Two redundant high-speed links (100 gigabits per second) connect CERN data centre to its extension hosted in Hungary at the Wigner Research Centre for Physics.

a mechanism that contributes to our understanding of the origin of mass of subatomic particles, and which recently was confirmed through the discovery of the predicted fundamental particle, by the ATLAS and CMS experiments at CERN's Large Hadron Collider". Other prizes for the laboratory included the Prince of Asturias award, the UNESCO Niels Bohr gold medal, and the Edinburgh medal awarded to CERN and Peter Higgs.

2013 has indeed been a fruitful year for CERN. The LHC came to the end of its highly successful first three-year run in February, giving way to a busy programme of maintenance and consolidation of CERN's entire accelerator chain that will enable the LHC to restart at higher energies in 2015. CERN's experimental teams have continued to analyse data during this period, and many of them have embarked on upgrades of their own. ALICE took data from lead-proton collisions and continued its steady programme of upgrade and consolidation work. LHCb observed a new matter-antimatter asymmetry in the decays of the B^0_s meson and measured one of the rarest processes in physics. The year was also busy for the rich research programme beyond the LHC. Teams at ISOLDE showed that atomic nuclei can be pear-shaped and measured the ionization potential of the rarest element on Earth; antiproton experiments made world-beating measurements on the antiproton and looked at the effects of gravity on antimatter and the CLOUD collaboration shed new light on the effect of cosmic rays on the formation of clouds in the atmosphere.

Computing is at the heart of all of these discoveries. The new phenomena that scientists are probing are extremely rare, hidden deep in already known physics. The LHC has therefore been designed to produce a very high rate of collisions (600 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, produced over 25 petabytes in 2012, the equivalent of 5.3 million DVD movies, which would take a thousand years to watch. During the last weeks of the run until February 2013, the remarkable figure of 100 petabytes of data stored in the CERN mass-storage systems was surpassed. This data volume is roughly equivalent to 700 years of full HD-quality movies.

To store, share and analyse these data, tens of thousands of computers 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 170 computer centres in 36 countries.

CERN is the central 'hub' of WLCG, the Tier-0 centre, providing approximately 15% of the total computing capacity. 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. In 2013, CERN and the Wigner Research Centre for Physics inaugurated the Hungarian data centre in Budapest, marking the completion of the facility hosting the extension for CERN computing resources. It adds up to 2.5 MW capacity to the 3.5 MW IT load of the Geneva data centre, which has already reached its capacity limit. About 500 servers, 20,000 computing cores, and 5.5 petabytes of storage were already operational at the site. The capacity in Budapest will gradually ramp-up following CERN's needs. The dedicated and redundant 100 gigabits per second circuits connecting the two sites are among the first transnational links at this distance. Operating this capacity remotely from CERN helps build knowledge and creates expertise and solutions with cloud computing to face big data challenges linked to exponential computing needs in all fields of research.

The Tier-0 is connected to about a dozen 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 centres then make the data available to Tier-2 sites (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.

WLCG runs more than 250,000 jobs at any moment, about 1.5 million jobs per day (corresponding to a single computer running for more than 600 years), and moves 20 petabytes of data (33 million files) per month. With data available to a community of 10,000 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 push back technological boundaries. CERN openlab partners contributed in a tangible way to the development of these systems and, as shown in this report, continue to collaborate in various domains on new solutions with success.

The Concept

Second CERN openlab IV annual Board of Sponsors meeting, in the presence of the sponsors and the CERN openlab team members.

Catalysing collaboration

CERN openlab is a framework for multilateral, multi-year projects between CERN and the IT industry.

Within the CERN openlab framework, 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.

Sponsorship can be at the associate, contributor, or partner level. Each type of sponsorship represents a different level of investment. The sponsors engage a combination of cash and in-kind contributions, the cash being used to hire young IT specialists dedicated to the projects. The associate status formalises a one-year collaboration based on independent and autonomous projects that do not require a presence on CERN site. The contributor status is a one-year collaboration based on tactical projects which includes a contribution to hire a young IT specialist supervised by CERN staff to work on the common project. The partners commit to a three-year programme of work and provide three kinds of resources: salaries for young researchers, products and services, and engineering capacity.

The successful CERN openlab concept was formulated in 2001 and stayed basically unchanged throughout the last decade. 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 III (2009-2011) not only capitalised on but also extended the successful work carried out in openlab II with the aim of hosting several major projects with a particular focus on technologies and services relevant to CERN and its partners. This annual report covers the second year of the CERN openlab fourth phase (2012-2014).

From left to right, CERN openlab management and CTO office group picture: Ioannis Georgopoulos, Kristina Gunne, Alberto Di Meglio, Mélissa Gaillard, Bob Jones, and Andrzej Nowak. Missing on the picture: Sverre Jarp.

The current phase is addressing new topics crucial to the CERN scientific programme, such as cloud computing, business analytics, the next generation of hardware, and security for the myriads of networks devices. The technical activities are organised in five Competence Centres (CC): the Automation and Controls CC with Siemens as a partner, the Database CC with Oracle as a partner, the Networking CC with HP as a partner, the Platform CC with Intel as a partner, the newly created Storage Architecture CC with Huawei as a partner. Rackspace joined CERN openlab in 2013 as a contributor while Yandex did join as an associate.

Each CERN openlab team is supervised by the CERN staff ensuring the liaison with its sponsor company. At the monthly minor review meetings, the teams are updated about the progress of the on-going projects, which fosters exchanges and ensures timely follow-up. At the bi-annual major review meetings, the sponsors meet with the teams who present their last results, and consider possible synergies. At the occasion of its annual meeting, the board receives information and exchanges views on the progress and medium term plans of CERN openlab.

The CERN 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. A list of the IT, EN and PH departments people most closely involved in the CERN openlab activities is given on page 13, while the positioning of CERN openlab activities within CERN is detailed on pages 14 and 15.

The distributed team structure permits close collaboration with computing experts in the LHC experiments, as well as with engineers and scientists from CERN openlab partners who contribute significant efforts to these activities. Principal liaisons with partners and contributors are listed on page 13. In addition, valuable contributions are made by students participating in the CERN openlab student programme, either directly to openlab activities or more widely to WLCG, and other Grid and CERN related activities in the IT department.

• CERN openlab Management

Rolf Heuer	CERN DG, Chair of CERN openlab Board of Sponsors
Frédéric Hemmer	Head of CERN IT Department
Bob Jones	Head of CERN openlab
Kristina Gunne	Administrative Officer

• CERN openlab CTO Office

Alberto Di Meglio
Sverre Jarp
Andrzej Nowak

• CERN openlab Communication Office

Mélissa Gaillard
Ioannis Georgopoulos

• CERN openlab Fellows and Staff

Dan Savu	Staff (HP)
Stefan Stancu	Fellow (HP)
Grzegorz Jereczek	Fellow (ICE-DIP)
Przemyslaw Karpinski	Fellow (ICE-DIP)
Aram Santogidis	Fellow (ICE-DIP)
Srikanth Sridharan	Fellow (ICE-DIP)
Marcel Zeiler	Fellow (ICE-DIP)
Georgios Bitzes	Technical student (Intel)
Andrzej Nowak	Staff (Intel)
Pawel Szostek	Fellow (Intel)
Liviu Vâlsan	Staff (Intel)
Ignacio Coterillo Coz	Fellow (Oracle)
Andrei Dumitru	Fellow (Oracle)
Maaike Limper	Fellow (Oracle)
Manuel Martin Márquez	Fellow (Oracle)
Lorena Lobato Pardavila	Fellow (Oracle)
Stefano Alberto Russo	Fellow (Oracle)
Marcin Bogusz	Fellow (Siemens)
Pavel Fiala	Fellow (Siemens)
Kacper Szkudlare	Fellow (Siemens)
Filippo Tilaro	Staff (Siemens)
Seppo Heikkila	Fellow (Huawei)
Maitane Zotes Resines	Fellow (Huawei)
Marek Denis	Fellow (Rackspace)

• CERN openlab Liaisons

Jean-Michel Jouanigot/Tony Cass	Liaison with HP
Alberto Pace	Liaison with Huawei
Sverre Jarp	Liaison with Intel
Tony Cass/Eric Grancher	Liaison with Oracle
Tim Bell	Liaison with Rackspace
Manuel González Berges	Liaison with Siemens
Marco Cattaneo	Liaison with Yandex

• Other IT and EN Departments Staff Contributors to CERN openlab

Edoardo Martelli	CS Group
Zbigniew Baranowski	DB Group
Luca Canali	DB Group
Eva Dafonte Pérez	DB Group
Dirk Düllmann	DSS Group
Brice Copy	ICE Group
Philippe Gayet	ICE Group
Piotr Golonka	ICE Group
Axel Voitier	ICE Group
Fernando Varela Rodríguez	ICE Group

• Industry Partner Liaisons with CERN

Rex Pugh	HP
James Prescott Hughes	Huawei
Davis Wu	Huawei
Claudio Bellini	Intel
Herbert Cornelius	Intel
Marie-Christine Sawley	Intel
Stephan Gillich	Intel
C. Gregory Doherty	Oracle
Monica Marinucci	Oracle
Toby Owen	Rackspace
Thomas Hahn	Siemens
Bernard Reichl	Siemens (ETM)
Guenther Zoffmann	Siemens (ETM)
Andrey Ustyuzhanin	Yandex

Positioning CERN openlab activities at CERN

The people involved in the CERN openlab activities are not concentrated in a single group or department at CERN but on the contrary, they span multiple units. In the Information Technology (IT) department, the Computing Facilities (CF), Communication Systems (CS), Database Services (DB), Data and Storage Services (DSS), Departmental Infrastructure (DI), Operating Systems & Infrastructure Services (OIS) groups host openlab activities, as does the Industrial Controls and Electronics (ICE) group in the Engineering (EN) department and the LHCb Computing (LBC) group in the Physics (PH) department.

Organization chart as from 1st January 2013

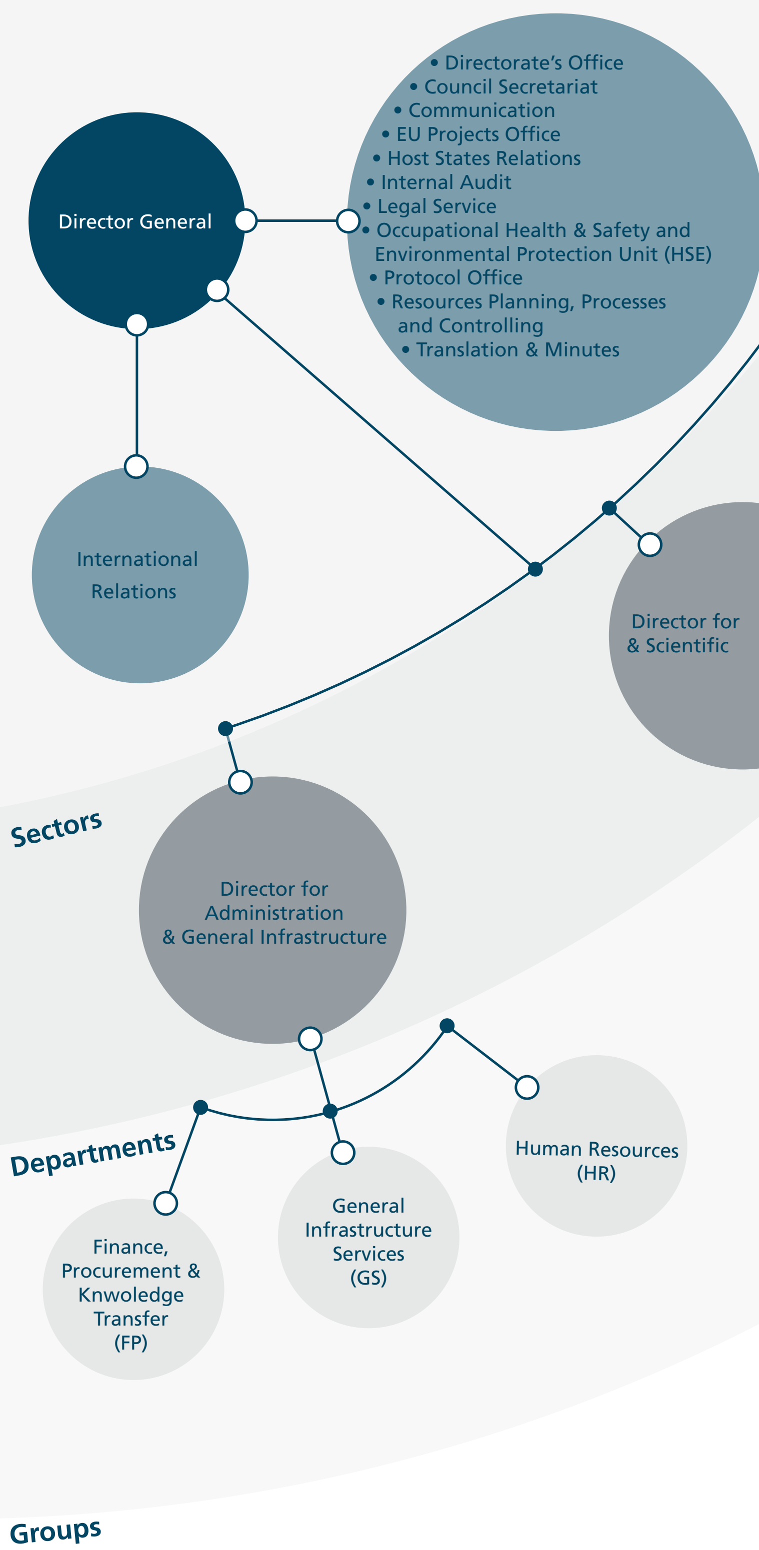

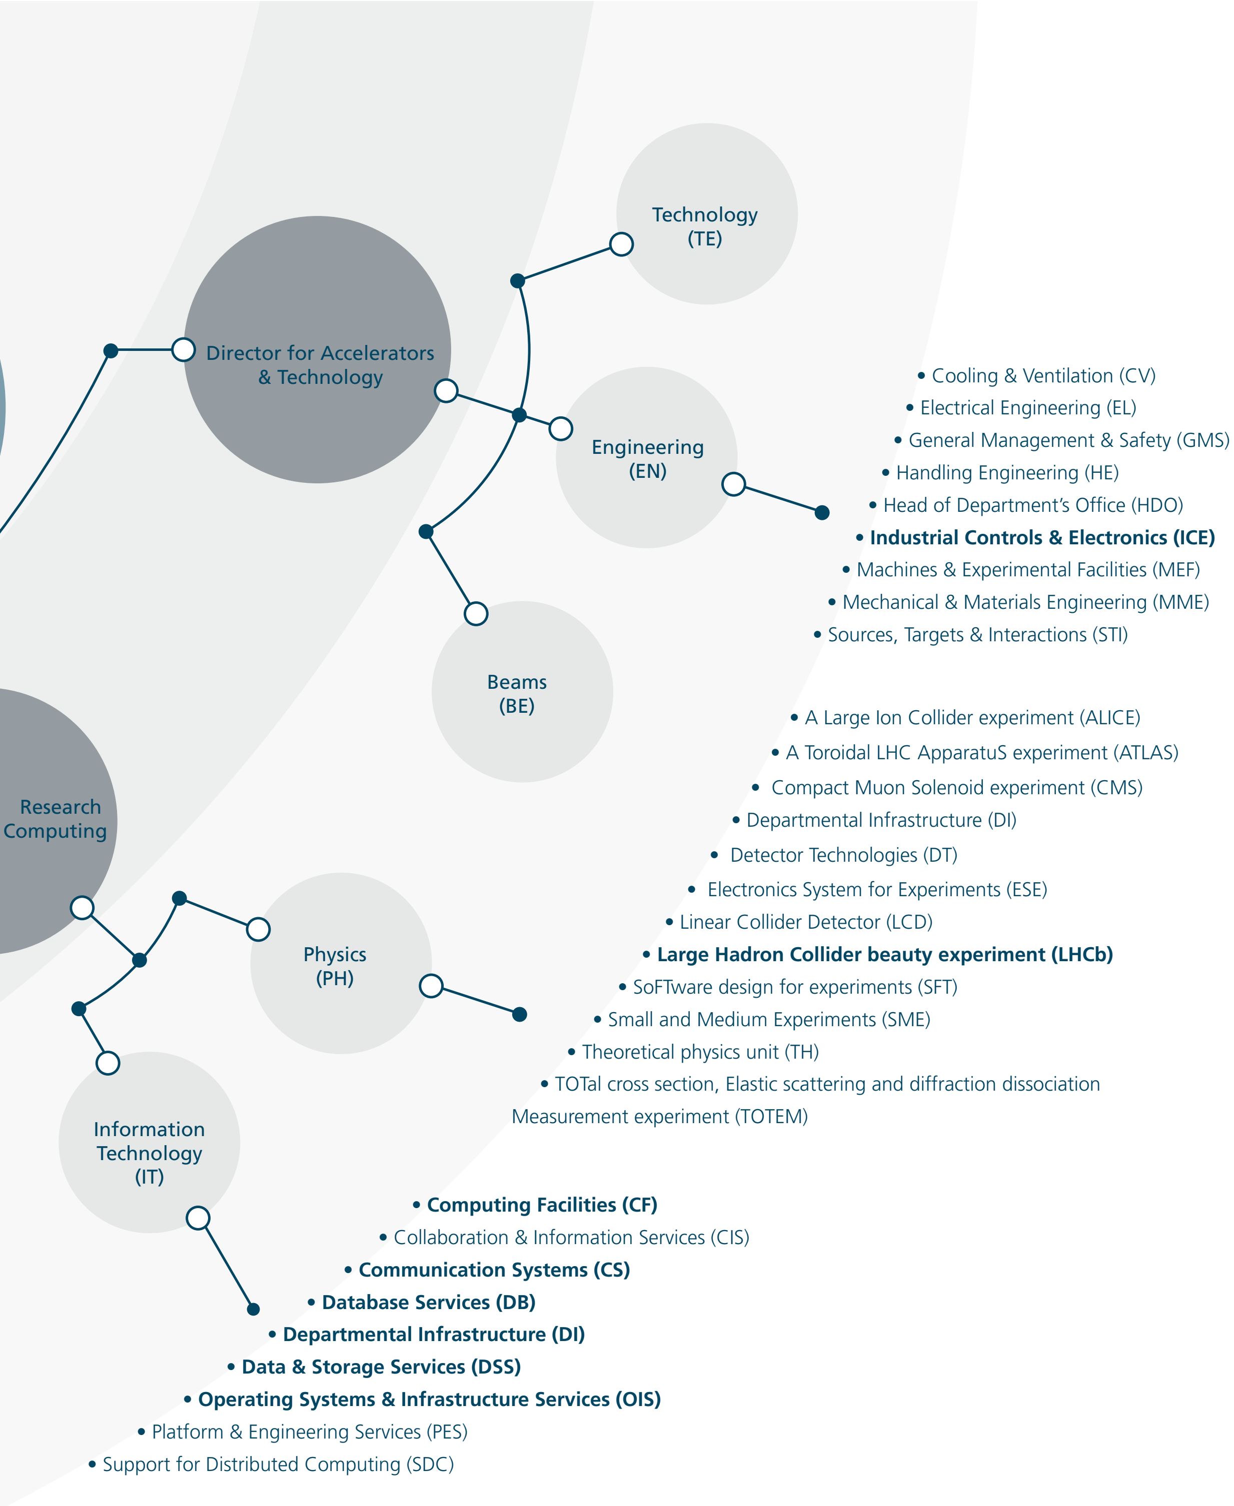

Technology
(TE)
Director for Accelerators
& Technology
Engineering
(EN)
• Cooling & Ventilation (CV)
• Electrical Engineering (EL)
• General Management & Safety (GMS)
• Handling Engineering (HE)
• Head of Department's Office (HDO)
• **Industrial Controls & Electronics (ICE)**
• Machines & Experimental Facilities (MEF)
• Mechanical & Materials Engineering (MME)
• Sources, Targets & Interactions (STI)
Beams
(BE)
Research
Computing
Physics
(PH)
• A Large Ion Collider experiment (ALICE)
• A Toroidal LHC ApparatuS experiment (ATLAS)
• Compact Muon Solenoid experiment (CMS)
• Departmental Infrastructure (DI)
• Detector Technologies (DT)
• Electronics System for Experiments (ESE)
• Linear Collider Detector (LCD)
• **Large Hadron Collider beauty experiment (LHCb)**
• SoFTware design for experiments (SFT)
• Small and Medium Experiments (SME)
• Theoretical physics unit (TH)
• TOTal cross section, Elastic scattering and diffraction dissociation Measurement experiment (TOTEM)
Information
Technology
(IT)
• **Computing Facilities (CF)**
• Collaboration & Information Services (CIS)
• **Communication Systems (CS)**
• **Database Services (DB)**
• **Departmental Infrastructure (DI)**
• **Data & Storage Services (DSS)**
• **Operating Systems & Infrastructure Services (OIS)**
• Platform & Engineering Services (PES)
• Support for Distributed Computing (SDC)

The Results

Cryogenic installation in the LHC tunnel controlled by Siemens automation equipment.

Automation and Controls Competence Centre

Control Systems are at the heart of all CERN facilities, accelerators, experiments and infrastructure.

Control systems are crucial to run CERN facilities. These include accelerators (such as the Large Hadron Collider), experiments (such as ATLAS, CMS, ALICE and LHCb) and technical infrastructure (such as the electrical network and the cooling and ventilation systems). The long-term collaboration between CERN and Siemens has played an important role in the successful running of all these facilities. Within CERN openlab, the Automation and Controls Competence Centre (ACCC) team is looking at the next generation of control systems for large installations. In 2013, the ACCC team focused specially on data analytics, without forgetting other ongoing projects (such as data archiving, large-scale deployment, and industrial security.

Data analytics

The Large Hadron Collider (LHC) has one of the largest and most complex industrial control systems ever built, and the volume of data generated is growing year after year. Analysing this data is vital for better understanding the entire control system, and for developing data-driven strategies that can improve its efficiency, functionality and predictability. In order to achieve this, innovative software solutions need to be developed to overcome the challenges surrounding the capturing, storing and processing of huge amounts of data within a tolerable period of time.

The main goal of the CERN openlab collaboration between Siemens and CERN is to design and implement a software framework, which can be used as a common solution to match the different data analytics requirements of the various control sub-systems (cryogenics, gas, vacuum, machine protection, etc.). CERN groups could potentially use this software framework to perform custom analysis based on the knowledge of the system experts. The research activities include both online analysis (analysis of current values directly retrieved from the control system under analysis) and offline analysis (analysis of the historical data stored into several repositories). An internal software tool provided by Siemens Corporate Technology has been used to perform root-cause analysis. This has enabled the ACCC team to identify the causes of distinct control faults and issues through the analysis of large volumes of data (mainly control alarms and machine logs) that otherwise could not be analysed. An emphasis was put on the gas systems of the LHC experiments to understand the failure modes of each of their components.

In addition, another Siemens data analytics framework has successfully been integrated with the current CERN control systems; this was achieved through OPC-based intercommunication (see OPC Foundation) with the SCADA system WinCC Open Architecture (OA). Finally, several research activities have been started to investigate and develop new data-mining methodologies and predictive models that can be used to extract valuable information from the huge amount of control data stored in the CERN repositories. Despite this project still being in a relatively early stage, we are confident that this fruitful and exciting collaboration between Siemens and CERN will lead towards new data analytics solutions aimed at improving the efficiency and predictability of control systems. Work on data analytics will continue throughout 2014 with the aim to offer a common framework to facilitate the tasks of many CERN engineers.

Relational database archiver

Significant progress was made in 2013 in the activities related to the Oracle database logging system for WinCC OA — both in the current production version (3.11 SP1), as well as for the future version of SCADA developed by Siemens. For the next generation SCADA system, an important milestone was achieved: the new Oracle plugin developed at CERN was integrated into a full vertical-slice test setup and recorded data into a database. This was actually the first time the new version of SCADA was run at CERN. In addition, some initial performance measurements of the logging subsystem (with the Oracle plugin) were completed by looking at the maximum data-write rate and then comparing this to the performance of the 3.11 SP1 version, as shown on the 'Archiving performance comparison' graph. The results were gathered in a report that was sent to Siemens for further analysis.

Data analysis framework

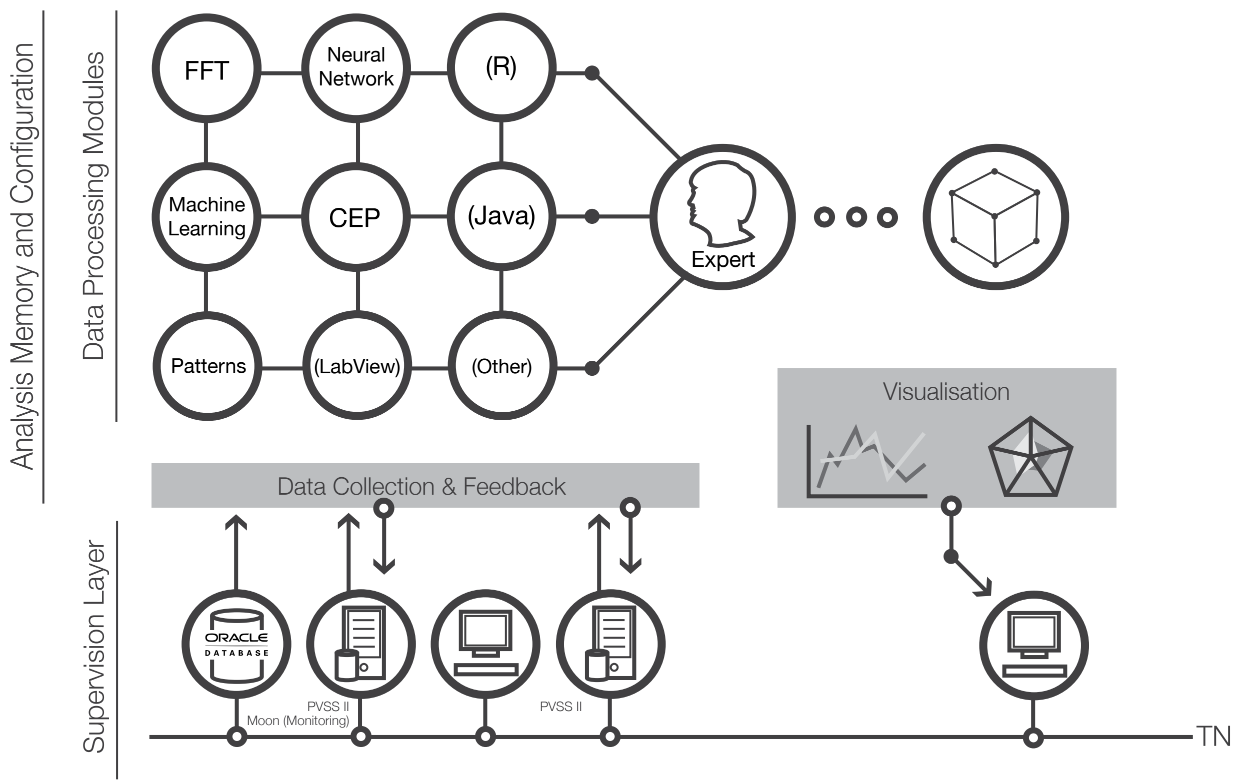

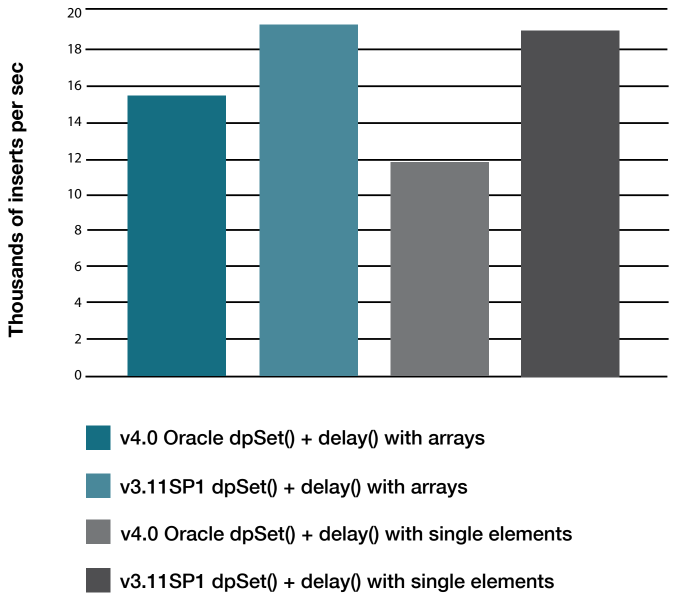

Archiving performance comparison

For the current version, there was a large improvement in performance. This was necessary to sustain the unprecedented data rates of the upgraded LHC magnet protection system (Quench Protection System, QPS). This improvement is crucial, as the upcoming LHC run will double the beam energy compared to the previous run. The data retrieved from the sensors needs to be stored to a database at a constant rate of 150,000 values per second and then queried to perform online analysis. The overall data volume is about 13 billion rows generated daily. Numerous optimisations have enabled reductions in the space required for a single data row and the disk-throughput has been halved. This was combined with careful tuning of the data readout process that enabled the required performance level to be achieved. The modifications prototyped and tested at CERN will now be implemented in the mainstream version of WinCC OA to benefit to all Siemens customers.

ACCC team picture, from left to right: Fernando Varela Rodríguez, Kacper Szkudlarek, Filippo Tilaro, Manuel Gonzalez Berges, Axel Voitier, Brice Copy, and Pavel Fiala. Missing on the picture: Marcin Bogusz and Piotr Golonka.

Large-scale deployment

The Siemens WinCC OA SCADA package is especially well-suited to building very large and highly distributed control systems, like the ones of the LHC experiments and those of many accelerator services. The unprecedented size of these control systems necessitates the use of tools to efficiently manage their evolution over the lifetime of these applications (over 20 years). In the framework of CERN openlab, CERN is collaborating with Siemens ETM on the development of the Central Deployment Tool (CDT) for a new version of the SCADA system. This should ease the initial setup of new controls applications and provide a powerful way to push upgrades to software components onto multiple sets of computers in a centralised fashion. The CDT will make use of the so-called WinCC OA ASCII Manager to import/export the run-time database of a project from/to ASCII files. However, the internal changes to WinCC OA envisaged following the adoption of the next generation of the SCADA platform require the development of a completely new ASCII manager. The new ASCII manager will need to handle the two significantly different data models exposed at different layers of the new SCADA (CHROM and OA) and the support media for the new ASCII Manager will be XML files.

WinCC open architecture

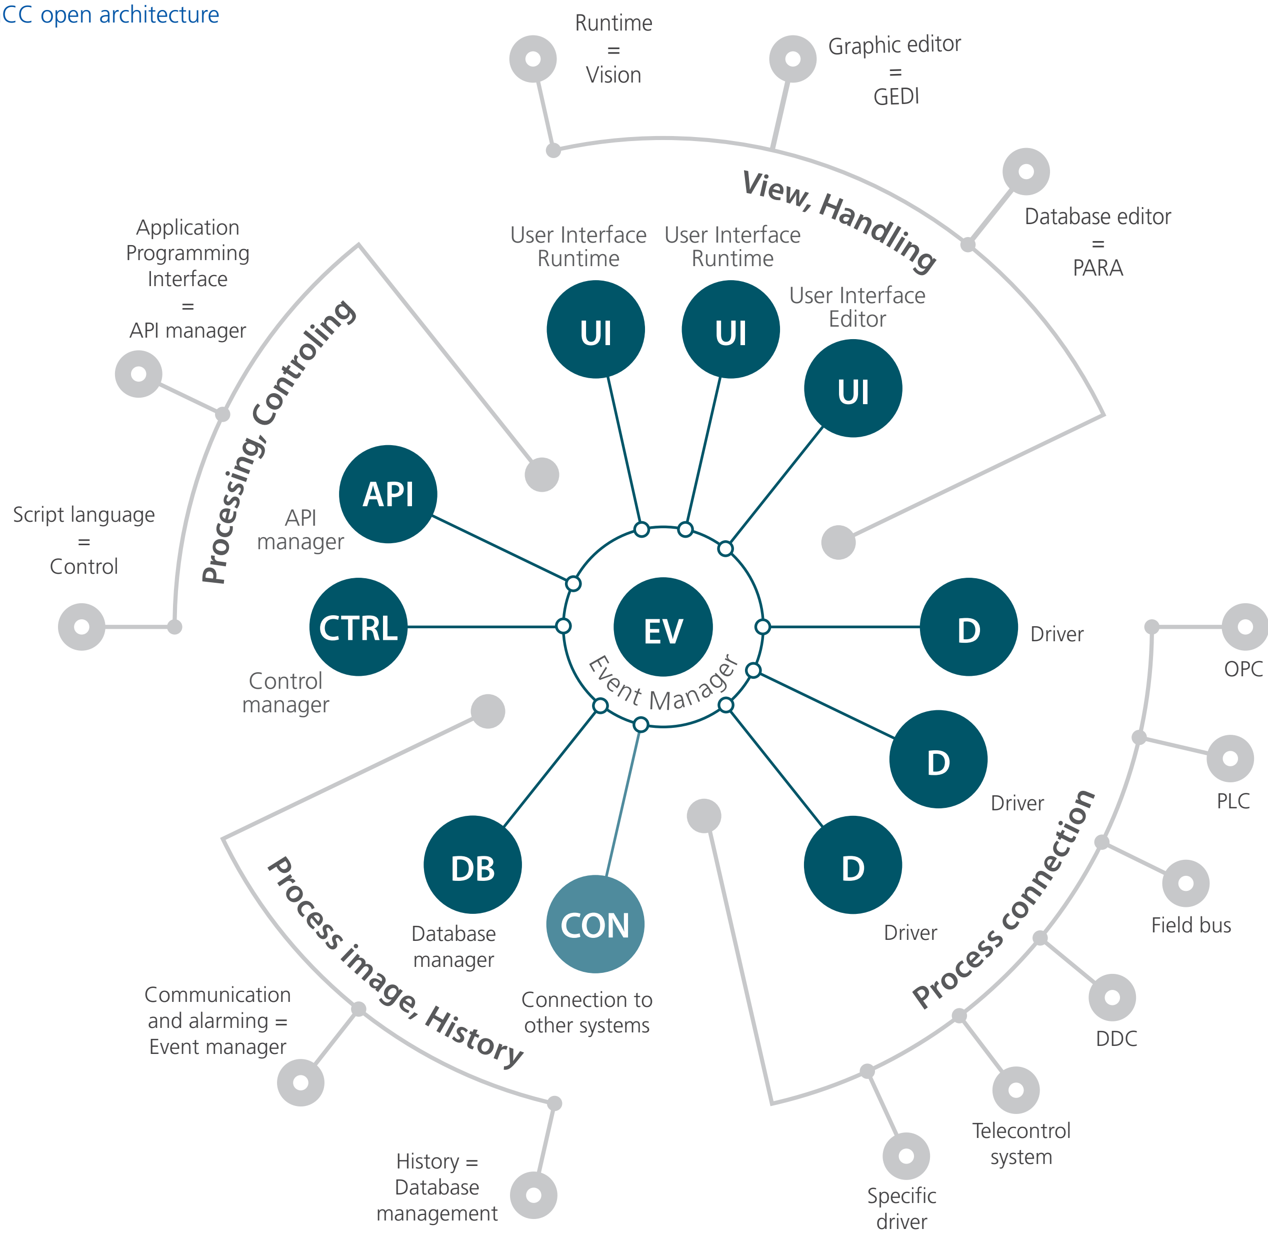

In this initial phase of the project, the work has focused on defining the requirements and functionality of the new ASCII manager. A comparison between the two data models implemented in the next generation SCADA has also been performed. This study has identified significant differences that have a major impact on existing applications and frameworks developed at CERN. Further discussion is therefore needed. The work performed so far has also covered the definition of the XML schema used for import/export, as well as the integration of various XML parsers into the WinCC OA development workspace and their benchmarking. This point is of special relevance as the possible use of SCADA on mobile devices calls for the minimisation of the CPU and memory footprints of the processes. The work is now moving towards an architectural design and the validation of the model proposed through prototypes.

Industrial security

Within the CERN openlab collaboration, CERN and Siemens have combined their efforts to design and implement a cybersecurity model for evaluating the robustness of the protocols implementations defined in the IEC-61850 standards. These protocols are widely deployed in advanced electrical power systems, called 'smart grids', which use digitised information and communication technology to drive industrial process operations on the basis of consumers' needs. The replacement of proprietary industrial control networks with open and standardised TCP/IP-based Ethernet networks enables an easier and more cost-effective integration of all industrial control system levels. However, this also exposes the entire infrastructure to internal and external cyber-attacks, which could destabilise the grid in unpredictable ways. This requires a proper design which takes into account not only the typical functional aspects, but also the cybersecurity requirements. The developed testing methodology, based on fuzzing and grammars techniques, has already proven to be effective at detecting communication weaknesses and at improving the overall robustness of those systems.

CERN industrial control systems architecture

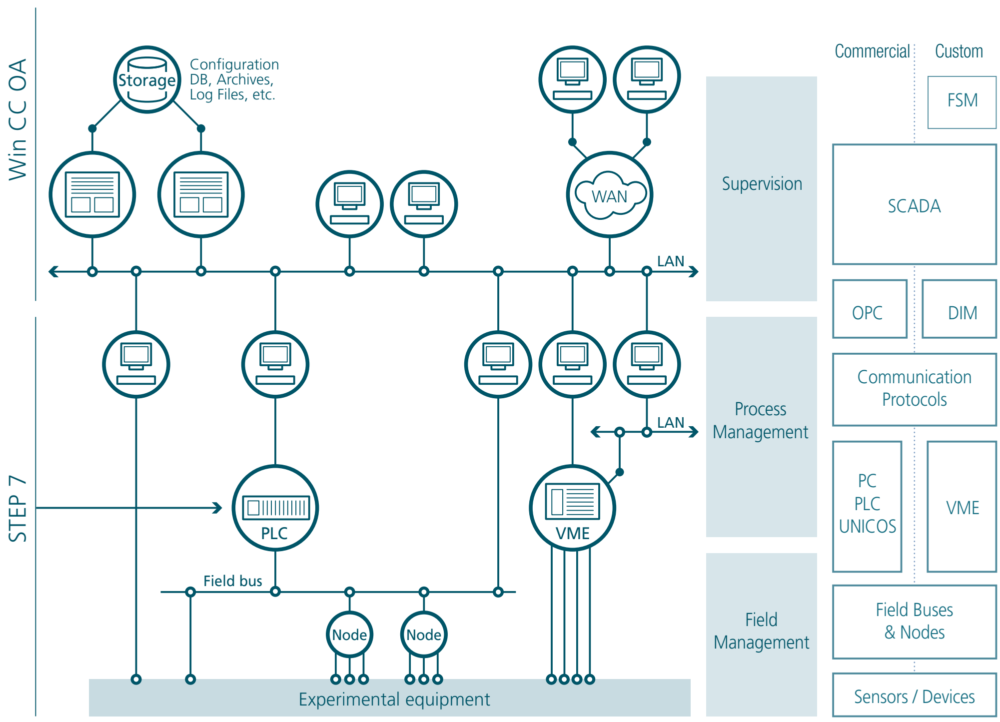

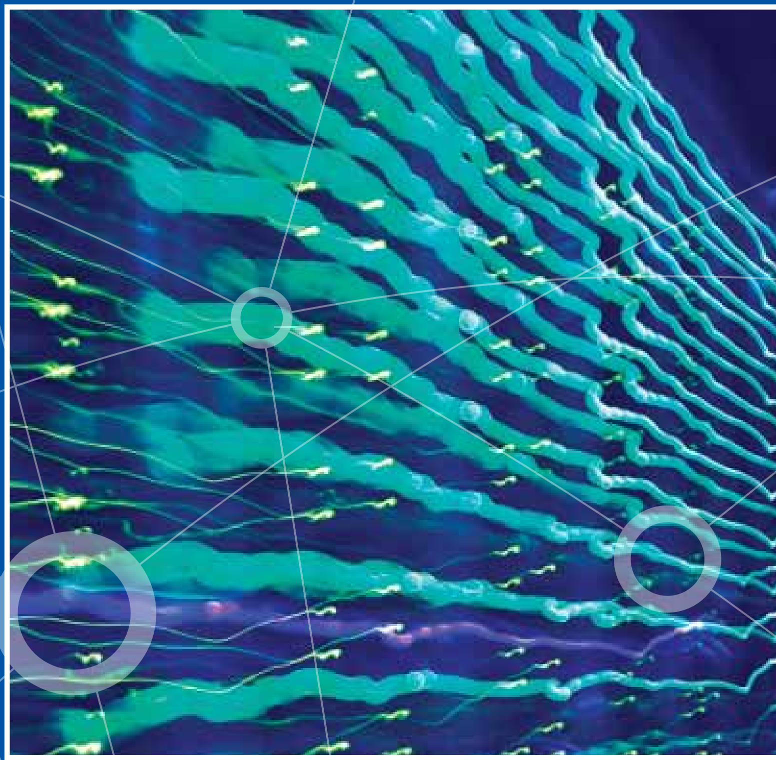

The Results

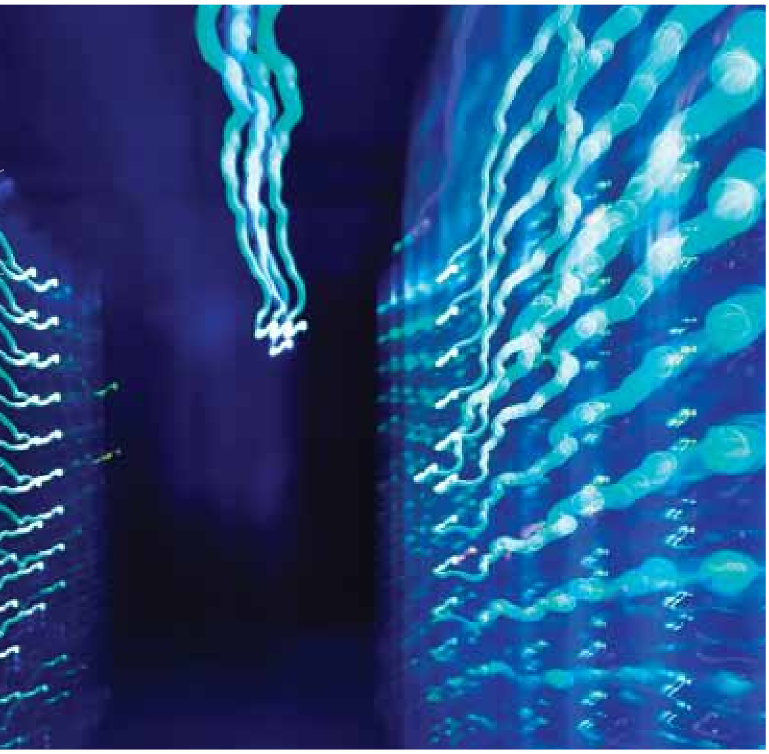

Database Competence Centre

Working with Oracle on key technologies for accelerator operations, physics and administration.

In 2013, the Database Competence Centre (DCC) team continued the work started as part of the CERN openlab phase IV in six different areas: database technology, database monitoring, replication, virtualisation, operational data analytics, and physics data analysis in an Oracle database.

Database technology and monitoring

The DCC team has been very active in the area of database technologies (notably Data Guard, Application Continuity and Multitenant architecture) and has continued to work on the monitoring aspects in 2013. The work related to database technology involved a wide variety of tests performed as part of the 12c beta programme and in relation to the Oracle Multitenant Architecture. Application Continuity has also been a major focus for the team. Oracle Data Guard is used extensively at CERN for replicating data for functionality and disaster recovery. During the 12c beta program, functional and performance tests were performed. The objective of the functional tests was to check that basic operations work as documented and that there are no critical issues in the core functionality of the product. The performance tests looked at redo-apply lag statistics with different redo transport modes using a simulated workload. In the scope of the 12c beta testing, the DCC team validated Data Guard related components, with a special attention given to the new 12c features. In particular, the team thoroughly tested new Active Data Guard functionality including the Far Sync feature, the Fast SYNC redo transport mode, and the Data Guard broker. Thanks to the work done and the collaboration with Oracle, the production version of Oracle database 12c has implemented a number of Data Guard features matching CERN's expectations for long-distance efficient replication.

The Oracle Multitenant architecture enables the use of pluggable databases — one of the key features of the Oracle Database 12c release — which could integrate perfectly in the existing Database on Demand service. The DataBase on Demand service at CERN is a Platform as a Service (PaaS) oriented service providing database instances to users. The pluggable databases tests were concentrated on validating functionality, backup and recovery. The Multitenant Architecture

DCC team picture, from left to right: Manuel Martin Márquez, Andrei Dumitru, Zbigniew Baranowski, Eva Dafonte Perez, Luca Canali, Ignacio Coterillo Coz, Lorena Lobato Pardavila, Eric Grancher, and Antonio Romero Marin. Missing on the picture: Maaike Limper and Stefano Russo.

is envisaged as part of the consolidation work planned at CERN. Work has been done to prepare the integration inside the Database on Demand service.

The new Application Continuity feature in the Oracle Database 12c solves a long-time complex issue with database applications: how to enable applications to recover from issues happening at the transport or database layer. Inversion 12c, Oracle has implemented a way to have the client and server sides keep a copy of the state so that it can be replayed in case it is required. Application Continuity tests were done in collaboration with the Controls group of the Beams department (BE CO) at CERN, in order to take advantage of an existing application — running in production — that could benefit from this new feature. The Java development team from the BE CO group created a replica of the production environment and tested Application Continuity by relocating services while the application was working. This has enabled the DCC to work together with the Oracle development team on real workloads and has proved that the feature can be used with very minimal changes in the application.

In the monitoring area, the DCC team kept working in 2013 on the integration of Enterprise Manager 12c with the existing infrastructure, the automation of tasks and extended the user base. The team upgraded to the latest release and integrated the authentication with the CERN Active Directory. They also took part in the Enterprise Manager Customer Advisory Board at the Oracle HQ in September and in the virtual Customer Advisory Boards that followed in order to share broadly within Oracle the work done at CERN through the openlab framework.

Virtualisation and replication

Virtualisation and replication technologies are important areas of work in the DCC. In 2013, Oracle VM has been tested for deployment against the CERN IT DB group requirements, while new features of Oracle products relevant to a database-oriented workload were also evaluated. In addition to traditional functionality testing and evaluation, special effort was dedicated to testing integration of the newest releases of both Oracle VM and Oracle Enterprise Monitor 12c. The main findings of

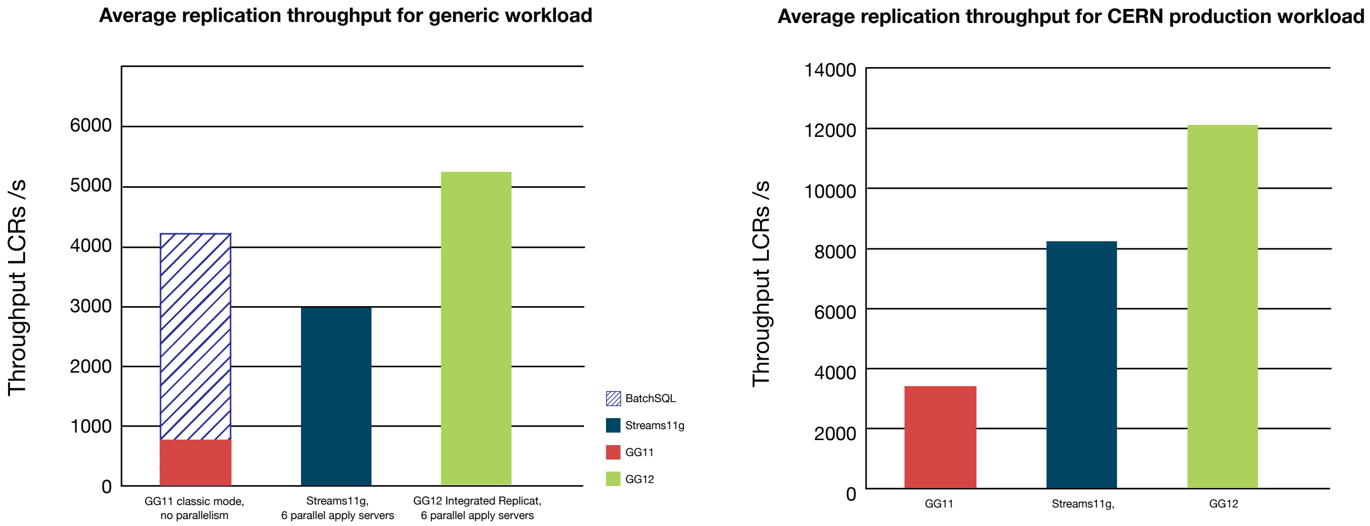

Oracle Streams and Oracle GoldenGate replication throughput comparison using generic and CERN production workload.

this effort were a noticeable improvement in the variety and depth of monitoring metrics for Oracle VM when monitored using the Oracle Enterprise Monitor; and a complete study and analysis of virtual resource provisioning using the OEM 12c controlling interface. The suitability of the Oracle VM hypervisor as a virtualisation resource provider in an OpenStack deployment has also been studied. This is of high interest in the CERN IT environment since CERN now uses OpenStack as a cloud provider for its central computing services. Furthermore, Oracle products on virtual environments running on Oracle VM infrastructure are granted a certification. With this in mind, integration tests have been carried out targeting the current releases of the OpenStack platform (named 'Grizzly' and 'Havana') and the development versions of what will become the next Oracle VM version.

In the database replication area, 100 MB of data are being exchanged daily between CERN and the Worldwide LHC Computing Grid Tier-1 data centres with Oracle Streams. Supervising the correct functionality of this complex replication environment is a difficult task and has been the focus of significant development effort within CERN openlab in the previous years. Following the line of continuous quality improvement, new options for replication have been developed by Oracle, such as Oracle GoldenGate. The DCC team conducted various tests on the new release's features to compare Oracle GoldenGate to Oracle Streams, with notable improvements obtained in monitoring, as well as for the overall throughput for replication. In addition, the team has been working with Oracle on new GoldenGate tools in order to provide dataflow monitoring and to reduce reaction time when facing possible issues. Veridata, GoldenGate Director, GoldenGate Monitor and OEM plugin have joined the replication area to achieve the maximum potential, stability and consistency of databases, which will be reflected in the production environment migration from Streams to GoldenGate programmed in 2014.

Data analytics

In the past years, huge amounts of monitoring and control data have been collected at CERN facilities and much more is expected in the coming years Gathering, storing and visualising this data was a challenge by itself, which has been successfully addressed. The challenge now is to obtain additional value from all this data, developing intelligent, pro-active and predictive monitoring and control systems driven by the analysis. From a theoretical point of view, understanding which models can be applied to extract value from the data for different use cases is crucial since there is no 'magic black box' which can perform this step automatically. As the literature suggests, application of the proper data-analytics approach can result in great improvements. Models for better understanding the performance, spotting the cause of the errors and detecting anomalies are just a few out of the short-term features which can be added on top of the monitoring and control infrastructures. The outcome of these models is the starting point for extremely interesting challenges in the medium term. In fact, by cross checking and correlating the results it becomes possible to find patterns and predict potential issues, thus enabling the development of early warning systems. However, infrastructure is key, as applying these models cannot be done using what was — until recently — the standard approach: a workstation running analysis software (such as in the R project for statistical computing) is today not powerful enough, given the volume of data. Moreover, given the growth of the monitoring and control data foreseen for the next years — in particular for the second run of the LHC — it is necessary to conceive a system capable of scaling out over the long term.

The first investigation stage, carried out in the context of the CERN openlab collaboration, has involved both open-source and Oracle tools to evaluate the methodology, the models and the software solutions for performing advanced data analytics. The DCC studied the current CERN use cases and selected two of them as references: data monitoring of the CERN Advanced STORage manager (CASTOR) and real-time and in-database analytics system for the LHC control systems. For the data monitoring of CASTOR, the field-test permitted a first exploration of the tools and of the methodology, and provided positive results, which have been presented on several occasions at CERN to show the potential of advanced data analytics. The work with CASTOR is now moving to the predictive level, and an early warning system is being studied. Concerning the real-time and in-database analytics system for the LHC control systems, after a review of the requirements and use cases, the team selected and implemented a Proof of Concepts (PoC) based on the automatic detection of failures on the cryogenic valves. They achieved levels of accuracy above 80% and were able to detect overheating magnets by mean of an ensemble of machine learning and statistical models. The next step is to move from the PoC to the production environment and to address new use cases.

Attention has been paid on potential future challenges in terms of data analytics. Consequentially, in the near future other use cases will be embraced, in collaboration with various CERN departments: Engineering, Physics, Beams, Information Technology and potentially General Infrastructure Services. Network monitoring and intelligent workload distribution around the Worldwide LHC Computing Grid sites, IT monitoring, and CERN Accelerator Logging Service are just some of the possible fields of application of advanced data analytics at CERN facilities.

Physics analysis in an Oracle database

This study looks at the possibility of performing analysis of the data collected by the LHC experiments through SQL-queries on data stored in a relational database. LHC analysis data is normally stored in, and analysed from, centrally produced 'ROOT' n-tuple files that are distributed through the Worldwide LHC Computing Grid. In the previous year, it was shown that a simplified physics analysis could be fully performed within an Oracle database, by re-writing the query in the form of SQL and using C++ libraries on the DB-machines linked via PL/SQL to be called within the analysis-query.

In March 2013, Oracle gave access for one week to an Exadata half-rack, which enabled testing of Exadata-specific features with the In-Database Physics Analysis. Hybrid Columnar Compression (HCC) was shown to reduce the overall data-size by a factor of 3.2. Interesting challenges were identified when external functions are required in the selection.

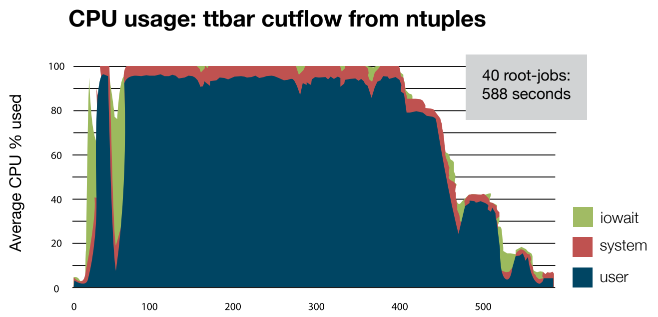

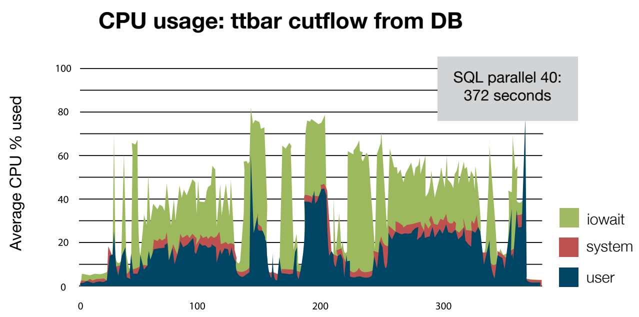

CPU usage comparison for ttbar cutflow within ROOT and an Oracle database.

The benchmark analysis on two terabytes of data took 210 seconds on Exadata without compression and 135 seconds on the same data compressed with HCC.

A five-machine test-cluster was prepared that could run both Oracle and Hadoop. A Hadoop-version of the physics analysis benchmark was made by storing the analysis data in comma-delimited text-files in the Hadoop file storage system and by writing the analysis in the form of Map- and Reduce-classes in Java. CPU-overhead from the Hadoop framework caused relatively poor analysis times in this approach. The setup was also used to participate in the beta-testing of the new Oracle In-Database MapReduce feature. This feature enabled the MapReduce-code prepared for the Hadoop version of the analysis to be used to run on data stored in the Oracle database.

In order to provide a more realistic benchmark, the actual analysis code for performing ttbar cutflow, as written by the ATLAS top physics group, was converted into a SQL-based version. Compared to the simplified benchmark, this analysis used more analysis variables (277 variables compared to 35 variables in the simplified benchmark) and required more calculations through external libraries. On the five-machine test cluster running both ROOT and Oracle, the results depending on some of the parameters, with a degree of parallelism of 40, the analysis took 588 seconds for the ROOT-ntuple analysis and 372 seconds for the Oracle DB-version. The work was presented as a poster at the Computing in High-Energy Physics (CHEP) conference in Amsterdam in October 2013 and won the best-poster vote in its designated track. The conference proceedings that describe this work in detail have been published in a special online volume of the Journal of Physics: Conference Series by IOP Publishing (http://iopscience.iop.org/1742-6596/513/2/022022).

The Results

Networking Competence Centre

Performing research and development with HP on cutting-edge networking technologies.

In 2013, the CERN openlab Networking Competence Centre (NCC) activities focused on Software-Defined Networking (SDN) through the ViSION (Virtual Services In OpenFlow Networks) project, a joint research and development collaboration with HP. This project started in 2012 and leverages HP's SDN platform and expertise for scaling out network resource utilisation by building a flexible traffic orchestrator. The research phase was completed in the project's first year and both parties approved the conceptual design. In 2013, the development phase reached an important milestone through the implementation of the core functionality required for traffic orchestration.

A game-changing technology

SDN is a new technology that promises to change the way networks work. It decouples the forwarding decision logic (the control plane) from the underlying infrastructure performing data transmission (the data plane).

Traditional networking relies on each network device taking its own forwarding decisions. Thus, devices must interact with their peers through complex distributed protocols in order to have enough information about the state of the network. In contrast with the traditional approach, SDN controller exports the decision logic from the networking equipment to an external software controller, which has a global view of the network and can make educated traffic-engineering decisions.

OpenFlow is the protocol that boosted SDN by providing an open standardised API (Application Programming Interface) between the data and control planes. It enables the fast specialised hardware forwarding engines from network devices to be 'programmed' by an external software controller, thus allowing full flexibility for designing and implementing the control plane of the network. SDN started gaining momentum after the first release of the OpenFlow protocol in 2009. The trend has continued and 2013 has been the year in which several SDN development platforms have become available, both from industry (e.g. HP, Juniper) and from the open-source community.

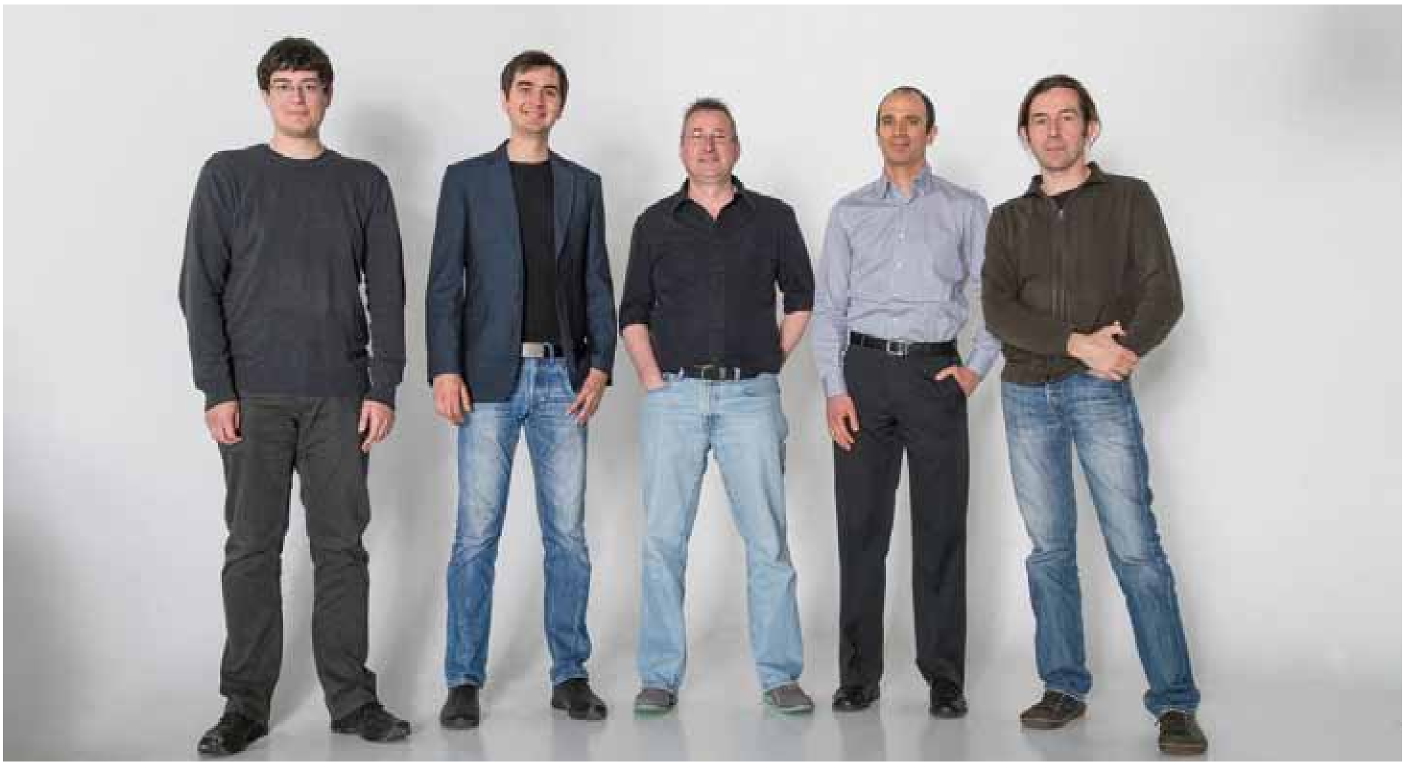

NCC team picture, from left to right: Andrei Radu Patrascoiu, Dan Savu, Tony Cass, Stefan Stancu, and Edoardo Martelli.

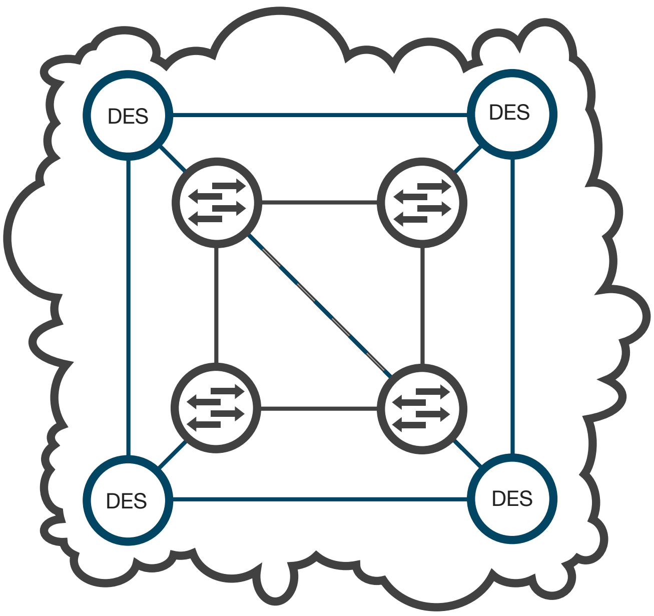

*DES = Device Embedded Software

Traditional networking: the control plane is distributed across devices, which run embedded software.

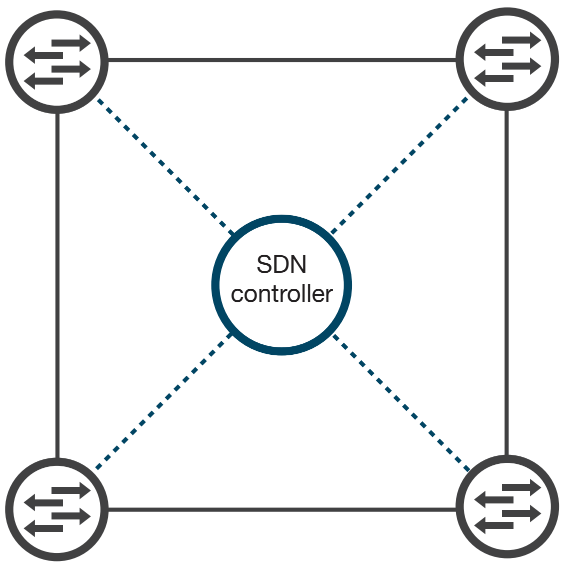

SDN: the control plane is centralised and control decisions are synchronised to devices through a standard API (e.g. OpenFlow).

ViSION researchers, Dan Savu and Stefan Stancu, in the CERN openlab section of the CERN data centre.

ViSION traffic orchestrator

The ViSION traffic orchestrator is an SDN application that aims at enabling resource scale-out by distributing network traffic to multiple resources of the same type. Client networks access sets of similar resources through a set of OpenFlow fabrics, which are 'programmed' by the ViSION controller. The single centralised controller enables correlated traffic orchestration in multiple OpenFlow fabrics, a feature that cannot be achieved with traditional network equipment.

To optimally distribute the load, the ViSION traffic orchestrator needs a feedback loop with information about the load and traffic patterns experienced by each individual resource. A configurable health-monitor module has been developed for this purpose. Using the health-monitor feedback, the controller can optimally distribute the load over the available resources. Furthermore, it can dynamically redistribute the load in case some resources become overloaded or unavailable.

The SDN centralised controller enables synchronised traffic-orchestration decisions to be made in multiple OpenFlow fabrics. This in turn enables scaling out of stateful resources. For example, a scalable firewall system can be built by deploying a properly sized tier of firewalls in-between two OpenFlow fabrics controlled by the same ViSION application. The SDN centralised controller ensures that both directions of each flow are handled by the same firewall, a requirement that cannot be easily fulfilled with traditional load-balancing network appliances.

ViSION application overview

The ViSION traffic orchestrator features a stacked SDN architecture, comprising three layers: infrastructure, control and application. The infrastructure layer consists of a fabric of OpenFlow enabled switches. In the control layer, HP's Virtual Application Networks (VAN) SDN framework is leveraged to accelerate the software development process and to ensure robustness. The ViSION orchestrator is implemented at the application layer, and uses HP's VAN northbound API for interacting with the controller.

One of the challenges of developing SDN applications is the lack of established validation and troubleshooting methodologies. To address this issue, the NCC team developed a prototype traffic injector. It supports the generation of a deterministic configurable sequence of flows, which enables the regressive testing of the ViSION traffic orchestrator.

Status and outlook

2013 saw the completion of the last two phases of the ViSION project: the core framework design phase and its development phase. At the end of each of these phases, the CERN openlab ViSION team visited the HP Networking team in Roseville, California, USA. On this occasion, technical and planning aspects were covered through meetings and brainstorming sessions. The fruitful interactions with HP engineers and technologists helped to better align the project's roadmap with HP's vision.

The ViSION core framework offers a platform for implementing traffic orchestration. The ViSION project has now been completed, enabling CERN and HP to independently assess the appropriateness of applying this technology to fit their own use-cases. From CERN's perspective, possible applications are scaling out its firewall system and data centre flow optimisation, while HP can leverage the solution and know-how to expand their SDN platform.

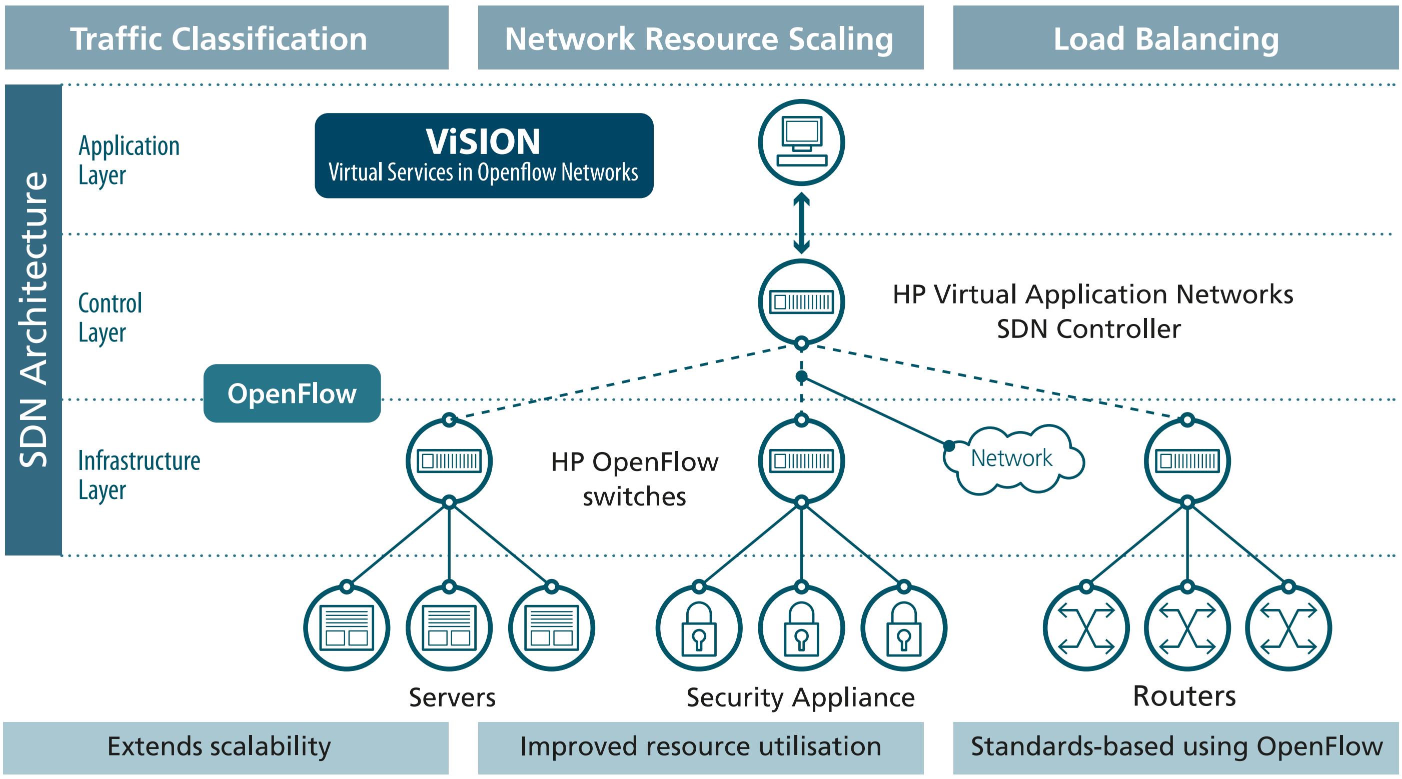

ViSION software architecture

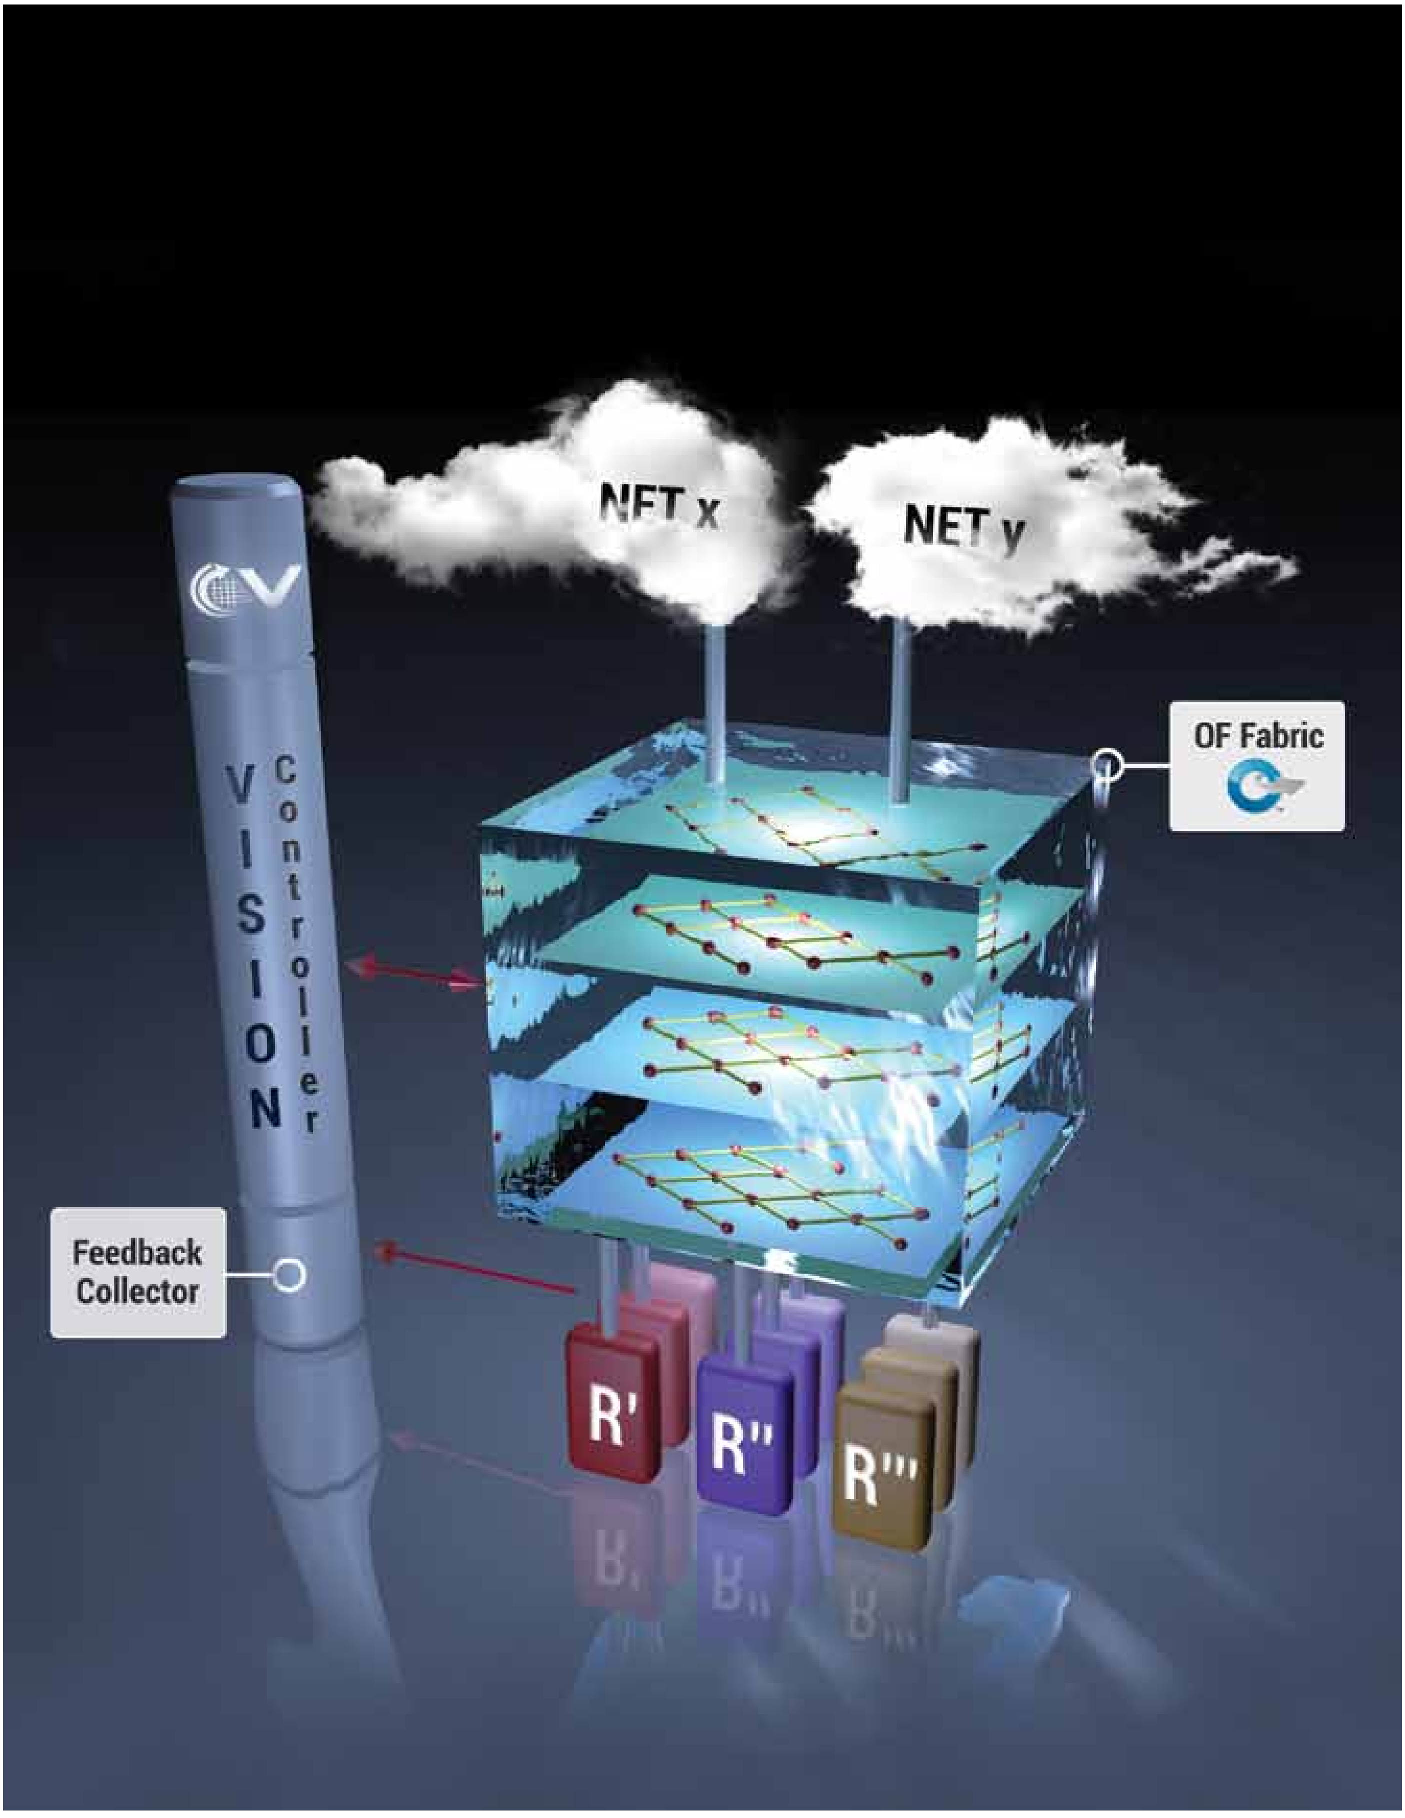
NET x
NET y
OF Fabric
VISION
Controller
Feedback Collector
R'
R''
R'''

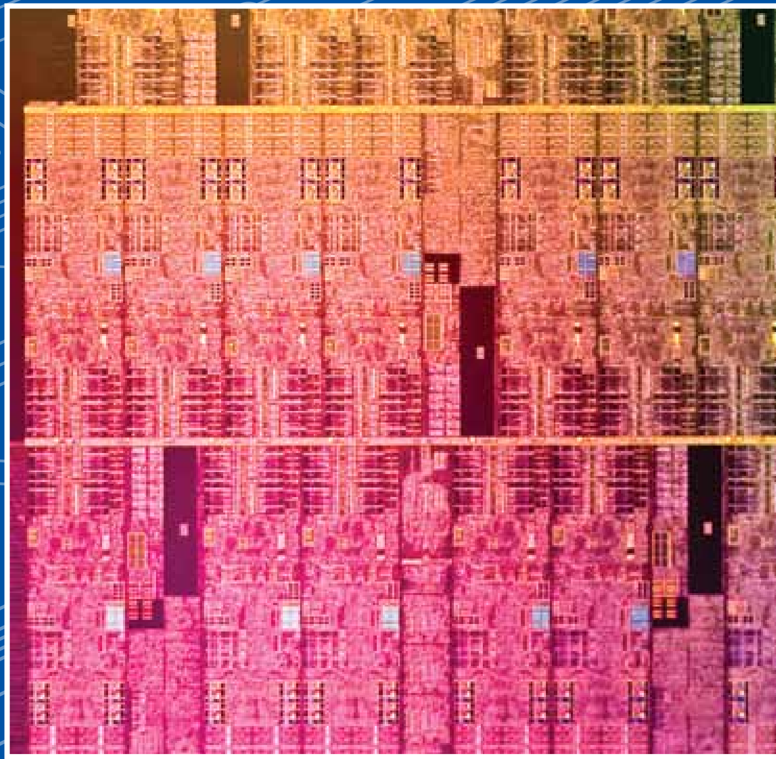

The Results

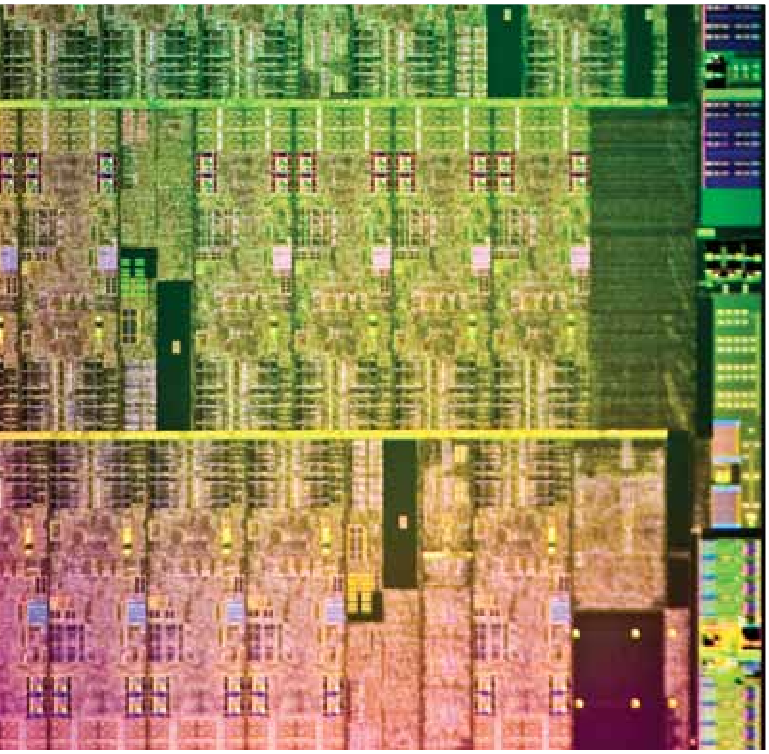

Intel MIC architecture co-processor.

Platform Competence Centre

Towards next-generation computing efficiency, together with Intel.

In collaboration with Intel, the CERN openlab Platform Competence Centre (PCC) continues to focus on hardware efficiency, software optimisations and performance measurements, as well as acceleration. This year again, the strong emphasis on teaching and knowledge dissemination has enabled a broader audience to enjoy the fruits of the PCC's work. 2013 also saw the start of ICE-DIP, the Intel-CERN European Doctorate Industrial Program, which is a European industrial doctorate scheme hosted by CERN and Intel Labs Europe with two Early Stage Researchers (ESRs) joining the PCC. Overall, the last 12 months constituted yet another year of intensive studies and development with tangible effects.

Several changes in the staffing of the PCC team occurred in 2013, which was indeed a dynamic year of growth. In January, Andrzej Nowak took over the PCC leadership responsibilities. At the end of her contract, Mirela-Madalina Botezatu departed for an industrial PhD at ETH Zurich with a major IT company. Liviu Valsan, the PCC hardware manager, joined the CERN IT procurement team, and was replaced by Pawel Szostek, a fellow with strong experience in software and data mining. Out of the five researchers recruited in 2013 for the ICE-DIP project, two are part of the PCC: Aram Santogidis and Przemyslaw Karpinski. Finally, Georgios Bitzes was hired as a Technical Student.

Benchmarking and evaluations

2013 saw the release of the dual socket Intel® Xeon® E5-2600v2 family platform 'Ivy Bridge', as well as the release of the 'Haswell' microarchitecture from Intel. The PCC team manually upgraded around 40 openlab systems to the new Intel Xeon E5-2695v2 processors, benchmarked the dual-socket platforms and disseminated its experience at a major physics conference. The paper was presented at the International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, held in Amsterdam, Netherlands, in October. In the same paper, the team also included experimental results from a desktop system equipped with an Intel® Xeon® E3-1280 v3 CPU, a representative of the new 'Haswell' microarchitecture whose most interesting feature is the addition of the Advanced Vector Extensions 2 (AVX2). Most notably, this new generation of processors is capable of single instruction multiply and adds a highly valued function in many physics calculations. A separate,

From left to right, CERN openlab PCC team members: Aram Santogidis, Georgios Bitzes, Liviu Vâlsan, Paweł Szostek, Andrzej Nowak, and Przemyslaw Karpinski. Missing on the picture: Olof Bärring, Sverre Jarp, Srikanth Sridharan, Marcel Zeiler, and Grzegorz Jereczek.

comprehensive report covering this feature, as well as vectorisation opportunities and other performance benefits brought by AVX2, was published on the CERN openlab website and is listed in the Education section of this annual report.

Low-power computing is still a point of high importance on the agenda, which was a reason for the team to take interest in a new pre-production platform from the Intel® Atom™ family ('Avoton'). The impressive performance-per-Watt results obtained on single-blade servers were seen as an optimistic message from the semiconductor world, and the team is looking forward to reproduce them in larger, full-chassis configurations throughout 2014.

Software

A major part of the PCC efforts is dedicated to software, in particular the evaluation of new software versions or new methodologies. In preparation for new capabilities of platforms, the standard benchmark portfolio has been completely revised and new micro benchmarks have been developed. MLFit, the highly parallel prototype data-analysis application underwent a major review in collaboration with an expert from the CERN Physics department. The new, leaner version sacrifices some generality for much improved scalability and is now used in addition to its predecessor.

The multi-threaded Geant4 prototype, on which the team has collaborated with the authors since 2008, saw the light of day with a major Geant4 release towards the end of 2013. It offers significant memory savings in multi-threaded mode, as well as the possibility of running on the Intel®Xeon Phi™ coprocessor. Several scientific communities are already making use of this new functionality.

The PCC also participated in some of the CERN Physics department's work on Geant V, which is the porting of Geant paradigms and logic to a vector-compatible prototype. With the aim of achieving high throughput even at the lowest levels of the architecture, the project has already demonstrated good speed-ups. This activity was carried out as part of the PCC's close involvement in the Concurrency Forum, which is a virtual organisation consisting of High Energy Physics

Georgios Bitzes and Pawel Szostek performing power-consumption studies in the CERN Data Centre.

stakeholders. Its aim is to ensure that the latest developments in computing architectures are exploited to the maximum extent.

Advanced development and tuning work on large software packages would not be possible without the right tools. In 2013, the pilot compiler project established by the PCC enjoyed another year of growth, as Intel tools made available CERN-wide gained more regular users. Amongst other publicly available software, the PCC also maintains its own builds of the Linux kernel. The package offers specially patched tools for cutting-edge performance tuning, to which the PCC contributes in collaboration with key industry players, as well as sophisticated satellite libraries. It was also reported that these custom kernel builds showed improved double-digit performance on some experimental frameworks. To further support tuning activities, the PCC prepared a report covering the performance impact of various common performance-measurement methods. Along these lines, the PCC also started new work on its own performance analysis packages and disassembly interfaces, where custom measurement tools can now be scripted in Python with ease, and perform better than those commonly available.

In another software-related highlight, the PCC team evaluated Cilk Plus, a promising parallelisation technology. It is currently an extension of the Intel compiler, and is being implemented in the Gnu C Compiler (GCC). Cilk Plus offers simple, yet powerful and much awaited parallelisation aids on two levels: thread (or task) parallelism is implemented using a spawn function and elemental functions, while data parallelism is expressed through familiar Matlab- or Python-like syntax. The results of the evaluation are contained in a report published on the CERN openlab website and listed in the Education section of his annual report. While programmability is apparently very high, some improvements on the performance front can still be made.

Intel Xeon Phi coprocessor

The PCC activities focusing on the Intel Xeon Phi started in 2008 and have borne plenty of fruit in 2013. The hardware laboratory now boasts 16 Xeon Phi systems, made available to active collaborators from CERN and its experiments. For example, the PCC was consulted by the scientists of the Alpha Magnetic Spectrometer (now launched into space) to double the performance of some test benchmarks. The multi-threaded Geant4 prototype was benchmarked in collaboration with the Geant4

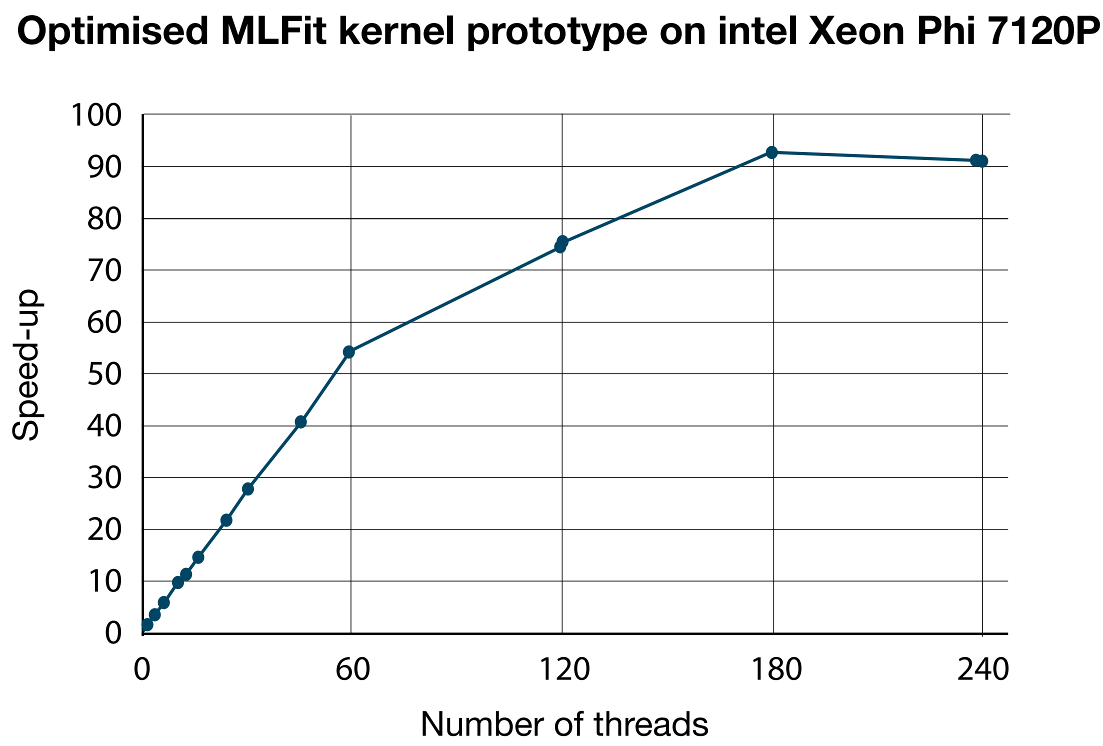

Results obtained on the Intel® Xeon Phi™ demonstrate good scalability for optimised benchmarks (as shown at CHEP 2013)

team, delivering very good scalability. The outcome of these five years of the PCC's experience with the Intel Xeon Phi were summarised at the Annual Meeting of the Concurrency Forum at Fermilab, Chicago, USA, in February, and later in a comprehensive overview paper published at CHEP 2013 (also listed in the in Education section of this annual report and available on the CERN openlab website).

Education

The PCC remains very active on the educational front: in 2013 alone, nine events were organised. The traditional workshops covering 'Parallelism, Compilers and Performance' as well as 'Numerical Computing' attracted well over 100 attendees and multiple expert speakers. In addition, two special workshops took place: one focused on the 'Intel Xeon Phi Platform', and another on 'Advanced Performance Tuning'. The latter event lasted three days and welcomed world-class speakers from major industry players, as well as CERN experts (full list available in the Education section of this annual report).

As in previous years, the CERN School of Computing (CSC) invited the PCC to deliver a teaching programme covering Computer Architecture and Performance Tuning. Given the popularity of the topic and the expanding knowledge of the CSC audience, the PCC co-organised the first thematic CSC as a related (but new) activity. In this week-long school, which took place in Split, Croatia, in June, 18 highly qualified participants were trained on advanced technologies and practices for efficient computing – including acceleration and co-processing. At the International Symposium on Computer Architecture, a major international computer science conference in Tel Aviv, Israel, the PCC collaborated with Intel to share its experience concerning the performance tuning of large workloads. The PCC expertise is well recognised by the IT sector, but also by a larger spectrum of industries: for instance, a large organisation from the finance industry expressed strong interest in the expertise developed by the team, which resulted in the organisation of a special training session.Overall, about two dozen guests from various institutions visited the PCC over the course of the year, offering public talks and expertise, while taking back home observations and feedback.

The ICE-DIP project organisers and supervisors at the kick-off meeting in Dublin, Ireland.

ICE-DIP

In February, the ICE-DIP project kicked off. The EU FP7-sponsored grant was co-authored by the PCC and the LHCb experiment in 2011 and is a part of the Marie Sklodowska Curie actions. From 2018, the LHC and its experiments will undergo a second round of upgrades, which could increase raw data rates by as much as a factor of 100. The objective is, therefore, to develop technical and scientific capabilities enabling efficient data taking under such conditions. The five young fellows, two of whom are associated with the PCC, have been recruited and enrolled in PhD programs in participating Irish universities: National University of Ireland, Maynooth and Dublin City University. After spending some initial time at CERN and at their universities in 2013, all students are going on extended 18-month secondments to Intel starting in 2014.

The Results

Storage Architecture Competence Centre

Evaluating and improving cloud storage architectures with Huawei.

Huawei joined CERN openlab as a contributor in 2012 and as a full partner in 2013 to cover the rapidly expanding area of cloud storage. The storage and management of LHC data is one of the most crucial and demanding activities in the LHC computing infrastructure at CERN and also at the many collaborating sites within the Worldwide LHC Computing Grid (WLCG). The four large-scale LHC experiments create tens of petabytes (PB) of data every year. This needs to be reliably stored for analysis in the CERN data centre and in many partner sites in WLCG. Today, most physics data is still stored with custom storage solutions, which have been developed for this purpose within the High-Energy Physics (HEP) community. As user demands increase (in terms of both data volume and aggregated speed of data access), CERN and its partner institutes are continuously investigating new technological solutions to provide more scalable and performant storage solutions to the user community. At the same time, CERN closely follows the larger market trends on the commercial side in order to continuously evaluate new solutions and to be ready for their adoption as soon as they have matured sufficiently for large-scale deployment.

The recently emerged cloud storage architecture and its implementations may provide scalable and potentially more cost-effective alternatives. Native cloud storage systems, such as the Amazon Simple Storage Service (S3), are typically based on a distributed key-value store, and divide the storage namespace up into independent units called buckets. This partitioning increases scalability by insuring that access to one area (bucket) is unaffected by the activity in other parts of the distributed storage system. In addition, the internal replication and distribution of data replicas over different storage components provides intrinsic fault-tolerance and additional read performance: multiple data copies are available to correct storage media failures and to serve multiple concurrent clients. On the larger scale (across site boundaries), the HTTP-based S3 protocol has become a de-facto standard among many commercial and open-source storage products. It may, therefore, become an important integration technology for consistent data access and exchange between science applications and a larger group of sites. One of the advantages of S3 is that the decision between operating a private storage cloud or using commercial cloud services is still left to the site, based on its size and local cost evaluation.

SACC team picture, from left to right: Dirk Düllmann, Seppo Heikkila and Alberto Pace. Missing on the picture: Maitane Zotes Resines.

During recent years, Huawei has invested significantly in complementing their traditional storage offering with a cloud storage implementation. Their Universal Distributed Storage (UDS) system is based on a key-value store which clusters large numbers of inexpensive CPU-disk pairs to form a scalable, redundant storage fabric. Recent versions of the product provide several redundancy mechanisms ranging from basic replication to erasure encoding. This functionality is aimed at enabling storage providers to choose among different levels of data integrity, access speed and media overhead to match the requirements of different applications.

Evaluation

The CERN openlab Storage Architecture Competence Centre (SACC) team is evaluating a Huawei cloud storage prototype that is deployed in the CERN data centre with a storage capacity of 768 terabytes (TB). The system has two main functional components: front-end nodes and storage nodes. The front-end nodes are user-facing nodes which implement the S3 access protocol and delegate the storage functions to storage nodes. The storage nodes manage data and metadata on local hard disks. The setup at CERN consists of seven front-end nodes and 384 storage nodes, each storage node being a disk-processor pair comprised of a 2 TB disk coupled to a dedicated ARM processor and memory.

The Huawei cloud storage system is designed for handling large amounts of data. All stored objects are divided into one megabyte (MB) chunks, and spread and stored on the storage nodes. The evaluated setup uses three replicas to ensure the data availability and reliability. Data replicas are distributed to different storage nodes such that a loss of one complete chassis (i.e. 16 storage nodes) will not have impact on data availability. In case of a disk failure, an automated self-healing mechanism ensures that the data on the faulty disk is handled by the other storage nodes. Corrupted or unavailable data is replaced using the remaining replicas.

The SACC team evaluated the Huawei cloud storage performance with a S3 benchmark that was developed for this purpose. The benchmark master process deploys and monitors many parallel client processes via dedicated client machines. A large number of benchmarks were executed in

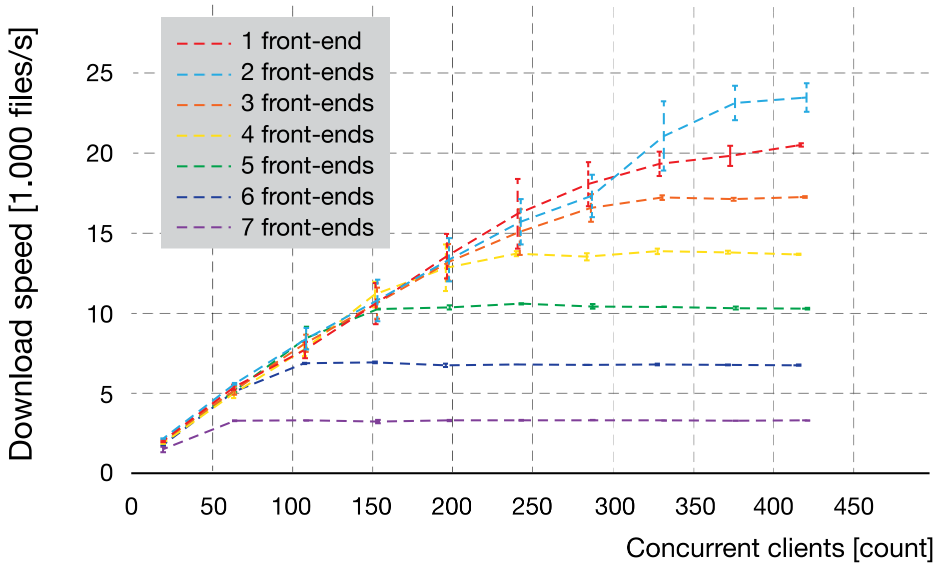

Huawei cloud storage metadata download performance.

different configurations to measure the aggregated throughput and rate of metadata operations. In order to study the metadata performance very small 4 kilobytes (KB) files were used. The small amount of payload data meant that any throughput-related constraints and stress were avoided, particularly with regards to metadata handling. In addition, the number of front-end nodes used was varied. The aim of this was to verify the capacity of each additional front-end node to add a similar amount of processing capability to the storage system.

Performance

The metadata download performance results are shown in the figure above. Each front-end node adds linearly around 3500 files per second to the total download rate. The achieved maximum 4 KB download performance could likely be further increased by simply adding more front-end nodes. The front-end nodes have not, however, been a limiting factor with 4 KB uploads. The maximum metadata upload performance was 2500 files per second when using around 1200 concurrent clients. **The cloud storage throughput performance was measured using 100 MB files.** With this payload size the upload performance is no longer strongly affected by the number of concurrently used buckets because the rate of metadata operations is several orders of magnitude lower. The upload and download throughput results show that the available network bandwidth of 20 gigabits could be filled easily. Each additional front-end node was able to download 1100 MB per second or upload 550 MB per second — until the network bandwidth limit was reached.

Failure recovery

The capability of the cloud storage system to recover transparently from sudden storage failures is one of the most desirable features for CERN. Indeed, if the system is able to recover automatically — including from major failures — the number of maintenance personnel readily available could be reduced. In order to study this recovery capability, the system was kept uploading and downloading actively while a power failure was simulated on one chassis by unplugging its power connection. The recovery test was executed for 500 seconds until the power was disconnected from one chassis containing 16 disks. The power was reconnected 300 seconds later.

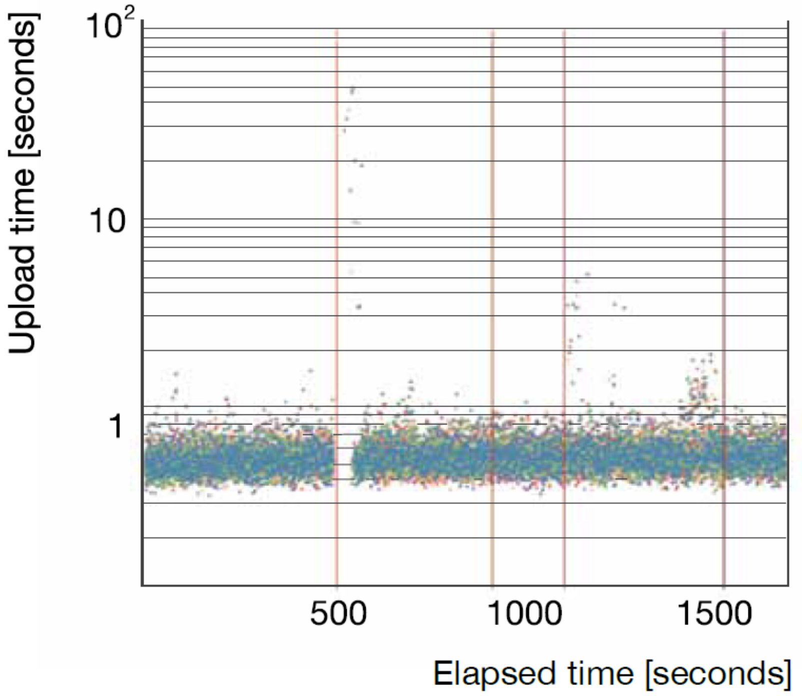

Upload times during Huawei cloud storage failure recovery

The cloud storage failure recovery results are shown in the diagram. The affected clients, doing uploads and downloads on this chassis, experienced delays up to about 60 seconds, but no errors could be detected as all the read and write operations completed successfully. Further in the automated recovery process, a third vertical red line shows the point at which the first nodes became available again after the chassis was rebooted. Again, only delays and no errors were observed.

Another important aspect for CERN is the flexibility of the cloud storage system for serving the existing CERN services. It was important to understand if the required performance level could be achieved and if all the necessary storage system features were present. The first evaluation focused on CernVM File System (CVMFS), which is a file system that, among other applications, is widely used to distribute HEP software. The CVMFS system — with its cloud-storage back-end — was tested by simulating a publishing step of a software release consisting of 30,000 small files. The results show that the Huawei cloud storage back-end was able to publish around 1200 new files per second. This means that a new software release of this size could be achieved in just a few tens of seconds. The cloud storage has been behaving as expected and no major problems were found that would prevent the use of this cloud storage system for HEP data storage.

A second Huawei storage cloud was also recently installed in the CERN data centre. This new and more compact cloud storage system has 1.2 PB of storage space, consisting of 300 4 TB disks. The cloud storage's S3 interface had been updated to support the latest S3 features, such as multi-part uploads. Next year, CERN openlab will focus on testing the performance of this new cloud storage system together with the older Huawei cloud storage installed in the data centre.

Second Huawei cloud storage setup in the CERN data centre.

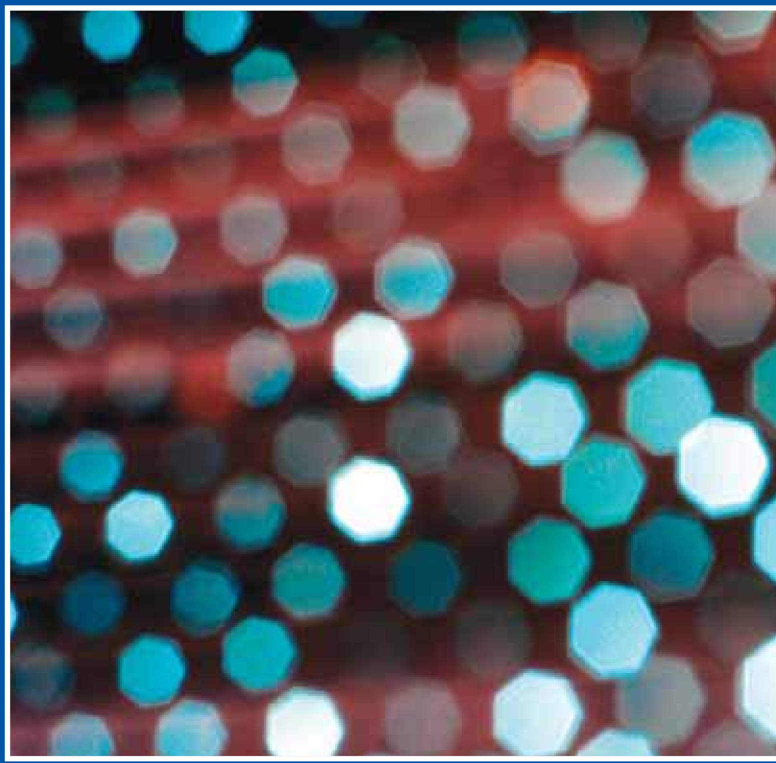

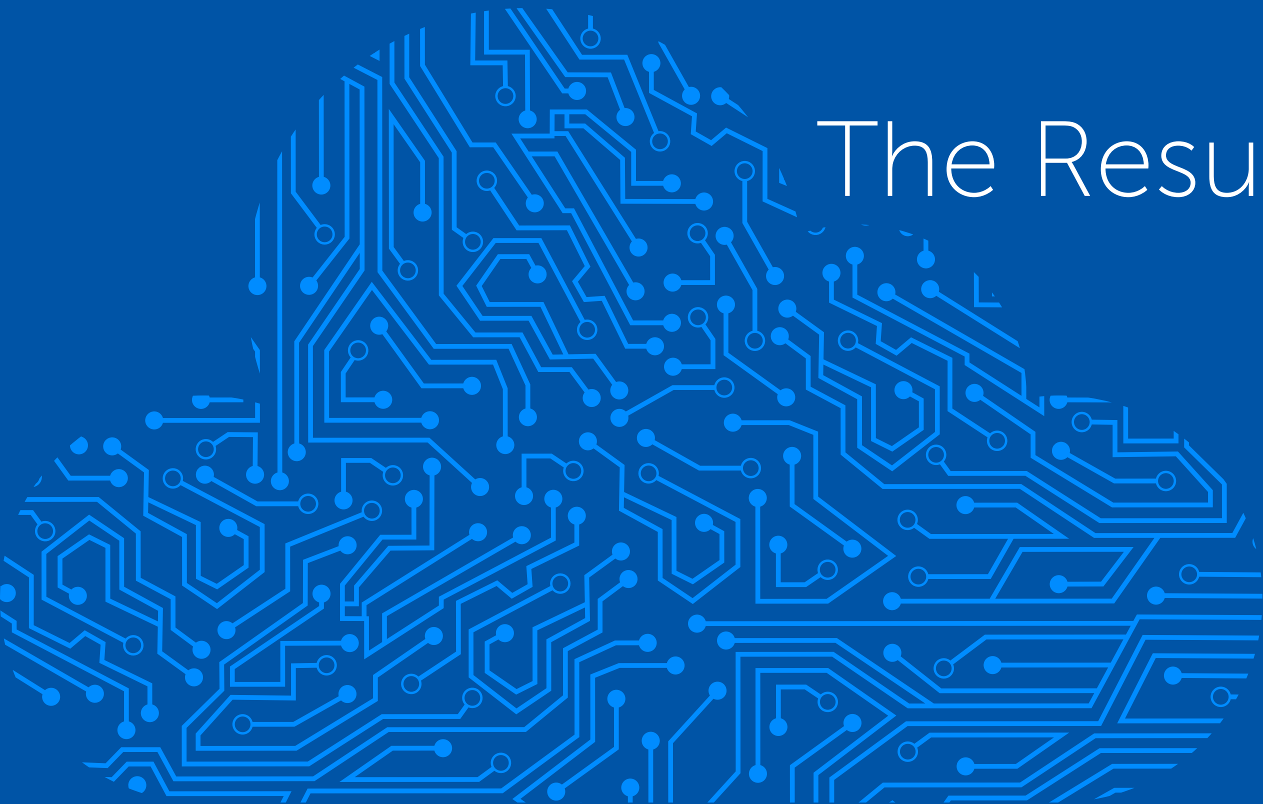

The Results

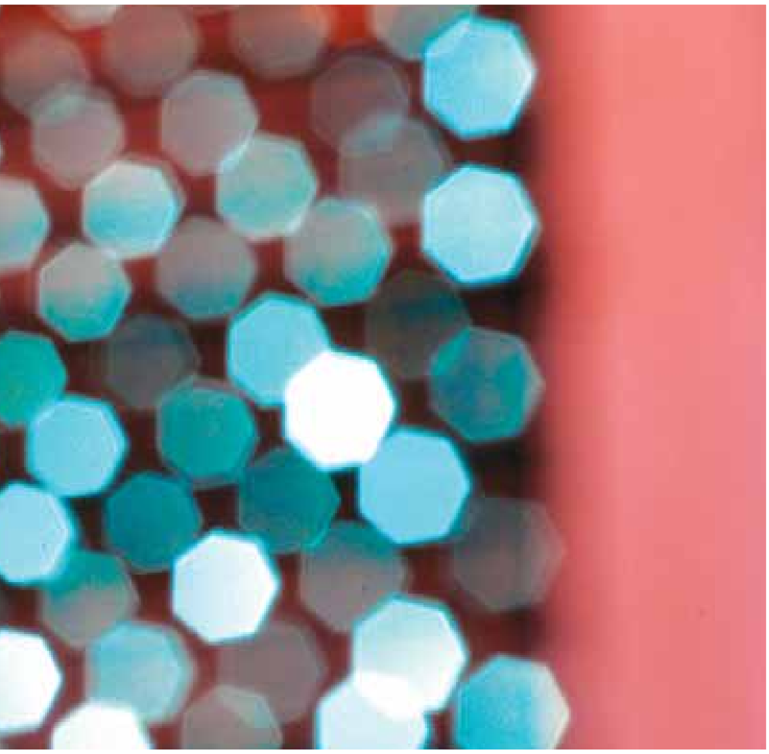

Project with Contributor

Collaborating with Rackspace for hybrid cloud-powered research.

Rackspace joined CERN openlab as a contributor in July 2013 to tackle a challenge crucial for both teams: how to get compute clouds cooperating and sharing resources. Together with NASA, Rackspace was one of the companies who founded OpenStack in 2010, an open source project to deliver a massively scalable cloud operating system. OpenStack has now grown to have more than 1300 developers from 200 companies contributing millions of lines of code. This code is licensed under Apache open-source conditions, enabling flexible use by companies and researchers. OpenStack powers Rackspace's public and private cloud solutions to deliver flexible and dynamic compute resources on demand.

OpenStack is also used extensively at CERN. The CERN IT department runs one OpenStack cloud while two others are hosted at the ATLAS and CMS experiments' pits; these make use of the compute resources of the high-level trigger farms when the accelerator is not running. Other high-energy physics sites, such as IN2P3 in France and Brookhaven National Laboratory in the USA, also deploy OpenStack.

CERN and Rackspace had the objective of sharing resources between these CERN clouds, the partner labs' ones and the public cloud resources at Rackspace's UK-based data centre. In the same way as Twitter or Facebook accounts can be used to access other web sites, physicists need to be able to log in to clouds using their identity from their home institution. Marek Denis, a CERN fellow, started to work on the project in October 2013 to develop the required enhancements to OpenStack to meet these needs.

The OpenStack development process is tightly related to its biannual summits, where thousands of developers, solution providers and users gather to share experiences and to design the next release. At the Hong Kong summit in November 2013, Toby Owen from Rackspace and Tim Bell from CERN presented the CERN openlab hybrid cloud project. Its main objectives are to provide a reference architecture for the federation of OpenStack clouds, blueprints and code to the open-sourcecommunities, as well as documentation (presentations and white papers) to enable other members of the OpenStack community to build on the CERN openlab findings.

Tim Bell and Marek Denis.

The design work started immediately with over 50 developers spending several hours debating in the federated identity design room at the summit. It was decided that initial efforts should focus on enabling the users of one cloud to be trusted by another cloud through identity federation. Following the OpenStack development process, a set of blueprints for how the design could be implemented was developed and assigned to Marek for implementation.

Meanwhile, the Rackspace technical support team had started to install the latest release of Rackspace's private cloud in the CERN data centre. The CERN engineering team was, therefore, able to investigate new ways of deploying applications in a cloud environment thanks for instance to the software-defined networking feature of the latest version of OpenStack. In the future, this cloud will be used to test the interconnection between CERN's private cloud and Rackspace's public cloud. The federation code was included in the OpenStack 'Icehouse' open source release in April 2014. This will be deployed at CERN during October 2014 and demonstrated in hybrid clouds between CERN and Rackspace's data centres at the OpenStack summit in Paris during November 2014.

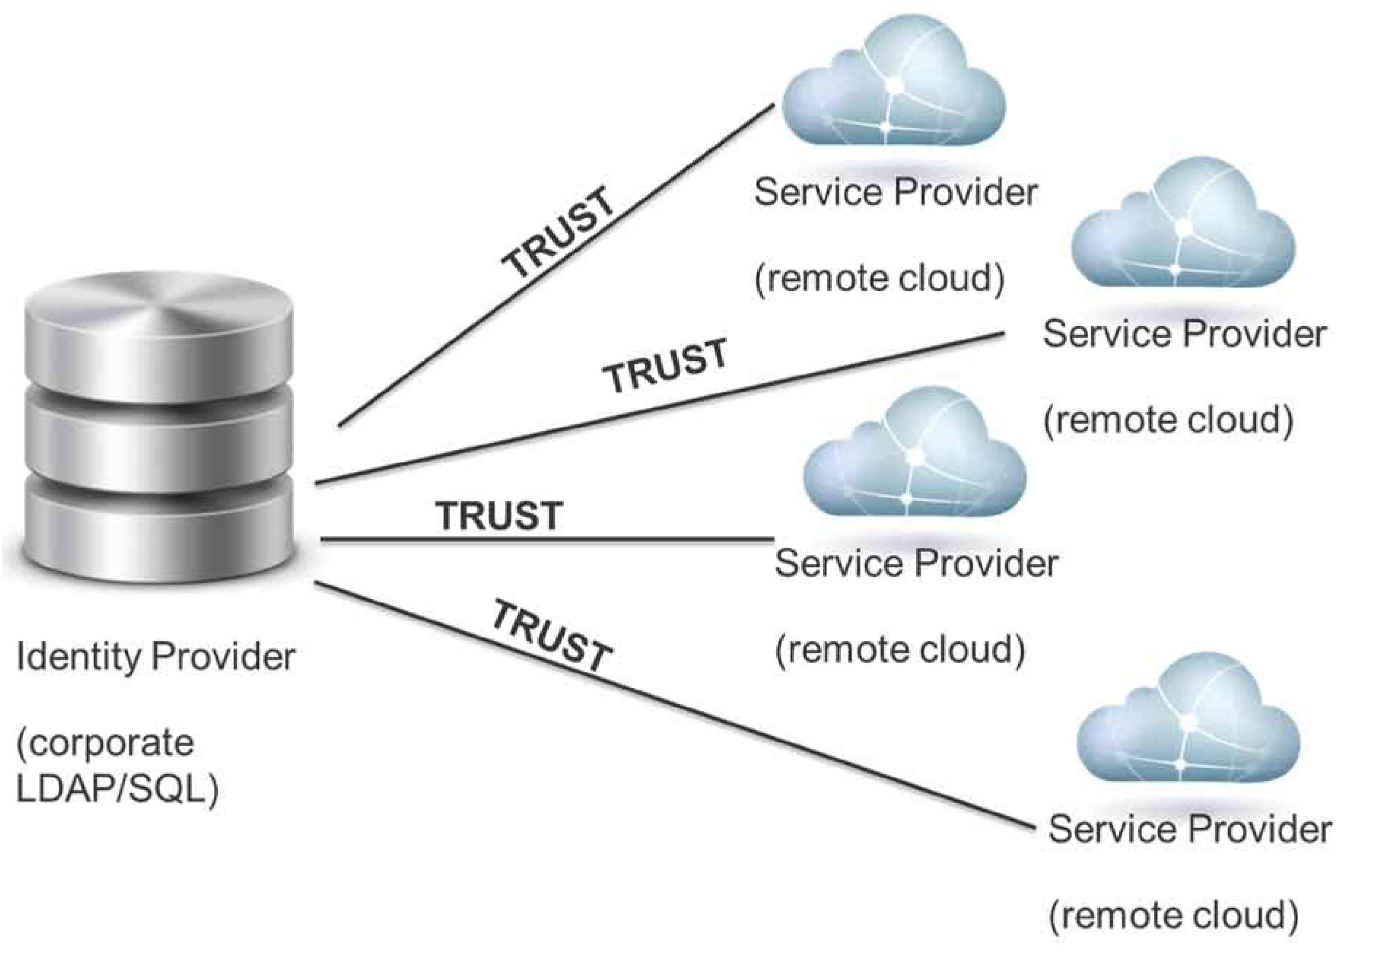

The project aims to make it possible for a single local account to access multiple cloud services.

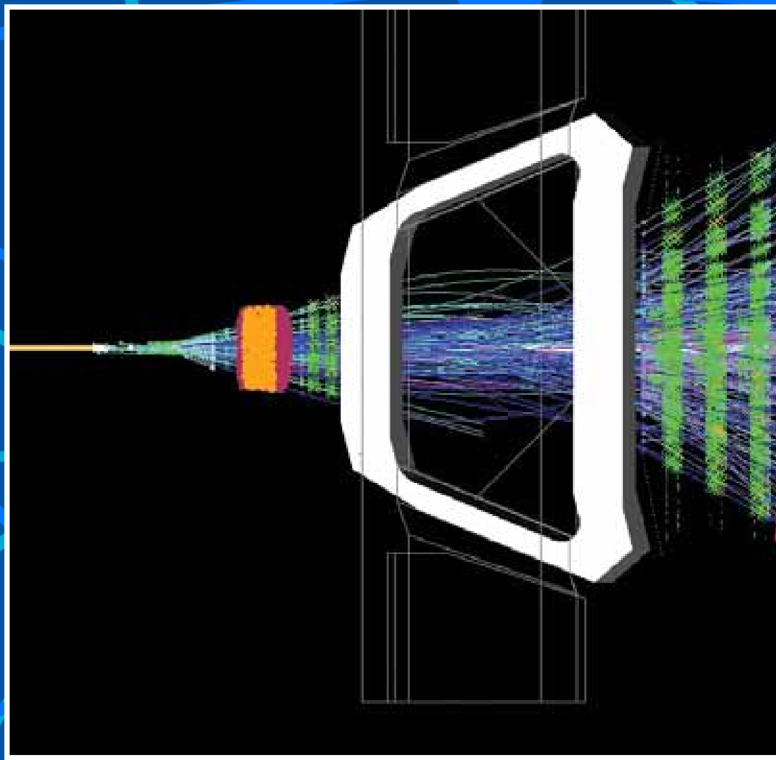

The Results

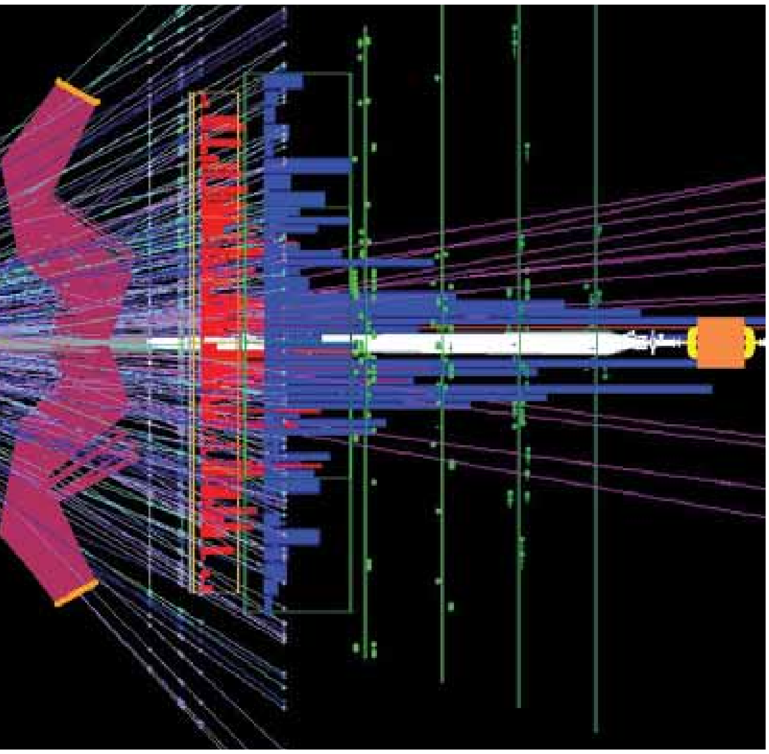

A proton-lead ion collision, as observed by the LHCb detector during the 2013 data-taking period.

Project with Associate

Working with Yandex to provide extra resources for the Grid model.

Yandex, Russia's leading search engine, joined CERN openlab as an Associate in early 2013 to provide computing resources and proprietary data-processing technologies — including the machine-learning technology MatrixNet — to CERN researchers. MatrixNet was already used for analysing B-meson decay data at the LHCb experiment, one of the four big LHC detectors. While this was just the first case of the technology's application at CERN, it encouraged Yandex to provide MatrixNet for analysis of a variety of other decays and further complex physical processes. The technology enables physicists to perform filtering of huge datasets and find extremely rare events.

During 2013, Yandex provided 5.2% of the total worldwide computing resources available to the LHCb experiment, making it the seventh largest provider of such resources, after the LHCb High-Level Trigger farm; CERN; and the Italian, French, German and British national computer centres for High Energy Physics.

Yandex used its search engine technology to provide a prototype index of LHCb events. The large LHCb datasets are typically analysed sequentially to select events satisfying specific criteria. Selecting particular types of events using traditional methods can take several hours on large computer farms. The index developed by Yandex enables LHCb physicists to find within seconds the rare events of interest for their analyses. Geared specifically to LHCb needs, this search system saves both time and resources and can greatly speed up the analysis.

In a second project, Yandex applied its MatrixNet technology as an alternative method of selecting interesting events from the LHCb datasets. By comparing the Yandex software with the algorithms that were used in the past, LHCb physicists managed to significantly improve their science output. The improved understanding of the difference between these old and new methods helped the physicists refine both techniques.


With great power comes great responsibility
—Uncle Ben
Yes, and also some really big heat sinks.
—John G

Education

CERN openlab is designed to create and disseminate knowledge.

Building human capital

Knowledge is created through the evaluation of solutions as well as genuine research and development of IT technologies. It is disseminated through multiple channels.

The CERN 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. In 2013, 10 topical workshops promoting the work of the competence centres were organised. Most of these workshops have a special feature: they involve a mix of lecturers from both industry and CERN, thus exemplifying the CERN openlab principle of two-way knowledge transfer through active collaboration. Several of them combine hands-off theory with hands-on practice. Special courses have also been organised for advanced CERN users, in addition to the regular quarterly courses. These classes touched on numerous topics, including data analytics, future tools and optimisations and were taught by Intel, Oracle and Siemens specialists. CERN openlab experts also contribute to off-site education activities such as the CERN School of Computing. These lectures and workshops are listed on page 54.

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 184 young computer scientists have participated in this programme. In 2013, the programme accepted 22 computer science and physics students of 13 nationalities for two months, during the period June to September.

The students worked on cutting-edge computing technologies supervised by CERN openlab staff, other groups in the IT, EN and PH departments as well as staff from WLCG. Visits were organised to the CERN facilities and to various experiments on site. In addition, the students toured Zurich in Switzerland: they were given talks at EPFZ, Google, and Open Systems. They also visited ILL in Grenoble, France and EPFL, in Lausanne, Switzerland. A dedicated lecture series for the students was given on site by CERN experts. For the first year, they were invited at the end of the summer to give 'Lightning Talks'. In six minutes and four slides maximum, they introduced the audience to their project, explained the technical challenges they faced, and analysed the results of their work. Several of this year's students were co-funded by CERN and HP, Huawei, Intel, Oracle, Siemens, or Yandex. Full details are given on page 55 and on the CERN openlab website.

This year, ICE-DIP, the Intel–CERN European Doctorate Industrial Programme, was officially launched in March. The five newly recruited early-stage researchers (ESRs) met at CERN for the first time in December. This Marie Curie Actions project within the European Union's 7th Framework Programme builds on CERN's long-standing relationship with Intel in the CERN openlab project. It brings together CERN and industrial partners, Intel and Xena Networks, to train the five ESRs. The researchers are funded by the European Commission and are granted CERN fellowships while enrolled in doctoral programmes at partner universities, Dublin City University and the National University of Ireland Maynooth. During their three-year training programme, the researchers will go on extended secondments to Intel Labs at locations across Europe. The project aims to develop unparalleled capabilities in the domain of high throughput, low latency and online data-acquisition. These topics form a solid foundation for future cutting-edge data-processing technologies of high interest to a broad range of industries. The research themes addressed by ICE-DIP are critical to the sustained forward momentum of European high-energy physics and numerous other disciplines of science and technology.

CERN openlab results have been disseminated in a wide range of international conferences. These publications, presentations and reports 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. The full list of presentations, publications, posters, and reports for 2013 is available on pages 56 and 57 of this annual report.

CERN openlab Topical Workshops

IT Requirements for the Next Generation Research Infrastructures Workshop, 1 February 2013, CERN, S. Bertolucci/CERN, C. Clark/HP, H. Cornelius/Intel, L. Field/CERN, K. Glinos/European Commission, J. Hughes/Huawei, B. Jones/CERN, M. Krisch/ESRF, V. Lindenstruth/Frankfurt University, J. de la Mar/T-Systems, M. Marinucci/Oracle, H. Mouren/Teratec, J. Perrin/ILL

Parallelism, Compilers and Performance Workshop, 25-27 March 2013, CERN, J. Arnold/Intel, S. Jarp/CERN, A. Nowak/CERN, L. Valsan/CERN

CERN openlab-Intel MIC / Xeon Phi Training, 11-12 April 2013, CERN, A. Nowak/CERN, K. Oertel/Intel, H. Pabst/Intel

CERN/Intel Workshop on Numerical Computing, 27-28 May 2013, CERN, J. Arnold/Intel, M. Corden/Intel, F. Dinechin/ENS Lyon, V. Innocente/CERN, L. Moneta/CERN, D. Piparo/CERN

Thematic CERN School of Computing (tCSC), 3-7 June 2013, Split, Croatia. S. Jarp/CERN, A. Nowak/CERN, H. Pabst/Intel

CERN/Intel Big Data Workshop, 9-10 July 2013, CERN, F. Carminati/CERN, P. Mato/CERN, N. Neufeld/CERN, A. Nowak/CERN, C. Schwick/CERN

CERN School of Computing (CSC), 19-30 August 2013, Nicosia, Cyprus, S. Jarp/CERN, A. Nowak/CERN

CERN openlab/Intel Workshop on Parallelism, Compilers and Performance, 6-8 November 2013, CERN, J. Arnold/Intel, V. Innocente/CERN, A. Nowak/CERN, G. Zitzlsberger/Intel

CERN openlab Workshop on Data Analytics, 20 November 2013, CERN, G. Anders/CERN, P. Andrade/CERN, L. Burdzanowski/CERN, S. Campana/CERN, M. Coelho dos Santos/CERN, A. Di Meglio/CERN, B. Doering/Intel, V. Garonne/CERN, D. Giordano/CERN, M. Gonzalez Berges/CERN, M. Herberts/Cityzen Data, M. Lamanna/CERN, L. Magnoni/CERN, M. M. Marquez/CERN, S. Ponce/CERN, C. Roderick/CERN, S. A. Russo/CERN, M. Salusti/SAS, F. M. Tilaro/CERN, A. Voitier/CERN

2nd CERN Advanced Performance Tuning Workshop, 21-22 November 2013, CERN, S. Bratanov/Intel, A. S. Charif-Rubial/Intel Versailles Exascale Computing Lab, M. Chynoweth/Intel, M. Dimakopoulou/Google, S. Eranian/Google, V. Innocente/CERN, D. Levinthal/Google, A. Nowak/CERN, R. Richter/Calxeda, M. Williams/ARM, A. Yasin/Intel

The CERN openlab summer students 2013 together with technical students from openlab and the CERN IT department

CERN openlab Summer Student Programme

CERN openlab Summer Student Programme Teaching Series, July-August 2013

General Security, S. Lüders/CERN
Software Security, S. Lopienski/CERN
Web Application Security, S. Lopienski/CERN
Cloud Computing in Large Computer Centers, B. Moreira/CERN
Physics Computing, H. Meinhard/CERN
Machine Learning for LHCb with Yandex MatrixNet, A. Ustyuzhanin/Yandex
Worldwide LHC Computing Grid Overview, F. Furano/CERN
Data Reliability at CERN and Ideas on How to Improve it, A. Pace/CERN
Oracle Databases at CERN, E. Grancher/CERN
Big data for Biomedical Sciences (The European Bioinformatics Institute), J. Dana/EMBL-EBI, T. Hancocks/EMBL-EBI
Systemtap: Patching the Linux Kernel on the Fly, T. Oulevey/CERN
Invenio Technology - Selected Practical Software Development Lessons from a Large Digital Library System, T. Simko/CERN
Control System Cyber Security, S. Lüders/CERN

CERN openlab Summer Students 2013, with Nation State, Home Institute and Project Topic

U. Argawal, India, Vellore Institute of Technology (VIT), India, "Linux Firewall State Synchronization"
Y. Almalioglu, Turkey, Bogazici University, Turkey, "Performance Improvements for the ATLAS Detector Simulation Framework"
A. Azzarà, Italy, Scuola Superiore Sant'Anna, Italy, "CERN Storage Systems for Large-scale Wireless Sensor Networks"
V. Doneva, former Yugoslav Republic of Macedonia, University Ss. Cyril & Methodius, former Yugoslav Republic of Macedonia, "MatrixNet: Using a new Multivariate Technique in High Energy Physics"
T. Dul, Poland, Uniwersytet Wroclawski, Poland, "Oracle AutoTask Enhancement for Multi-database Environment"
J. J. F. Alfonso, Spain, Universidad de La Laguna, Spain, "Vectorization with Haswell and Cilk+"
B. Kolobara, Croatia, University of Zagreb, Croatia, "Electronic Ticket and Check-in System for Indico Conferences"
A. Krajewski, Poland, Warsaw University of Technology, Poland, "Regressive Network System Testing"
Z. Kraljevic, Croatia, University of Zagreb, Croatia, "Integrating Cloud Services into the Invenio Digital Library"
C. Lindqvist, Finland, University of Helsinki, Finland, "Improved Metrics Collection and Visualization for the CERN Cloud Storage Test Framework"
G. McGilvary, United Kingdom, University of Edinburg, United Kingdom, "The Implementation of OpenStack Cinder and Integration with NetApp and Ceph"
D. Michelino, Italy, Università degli Studi di Napoli Federico II, Italy, "Implementation and Testing of OpenStack Heat"
A. M. Lopez, Spain, University of Vigo, Spain, "Improvement of the IT-PES-PS Section Services Statistics Page"
K. Panagidi, Greece, National and Kapodistrian University of Athens, Greece, "DBaaS with Enterprise Manager 12c and Oracle VM"
A. Radu Patrascoiu, Romania, Polytechnic University of Bucharest, Romania, "The ViSION Health Monitor"
R. Rama-Ballesteros, Spain, University of Granada, Spain, "Integration of Network Performance Monitoring Data in FTS3"
A. R. Peón, Spain, Universidad de Oviedo, Spain, "Advanced

Visualization Tools for CERN Institutional Data"
A. Sharma, India, Amity University, India, "Development of Personal Collections in CERN Document Server (CDS)"
E. A. Simon, Hungary, Budapest, University of Technology and Economics, Hungary, "Evaluation of In-memory Database TimesTen"
A. Tsikiridis, Greece, Athens University of Economics and Business, Greece, "Integration of LHC Experiments Resource and Tools with Helix Nebula"
V. Vintila, Romania, Spiru Haret University, Romania, "Evaluation of Standard Monitoring Tools (including log analysis) for Control Systems at CERN"
J. Zhang, China, University of Minnesota - Twin Cities, USA, "CERN Webfest and Citizen Science Challenges"

CERN openlab Presentations and Publications

Presentations

M. Limper/CERN, How to Discover the Higgs Boson in an Oracle Database, IT Technical Forum, CERN, Switzerland, 25 January 2013
B. Jones/CERN, CERN openlab: A Successful Public-Private Partnership, IT Requirements for the Next Generation Research Infrastructures Workshop, CERN, 1 February 2013
A. Nowak/CERN, Vectors: the Trends in Hardware, 2nd Annual Concurrency Forum Meeting, Chicago, USA, 5 February 2013
A. Nowak/CERN, Practical Results of the Intel MIC / Xeon Phi Project at CERN openlab, 2nd Annual Concurrency Forum Meeting, Chicago, USA, 5 February 2013
A. Nowak/CERN, A Brief Correlation Study of x86 Compiler Flags and Performance Events, 2nd Annual Concurrency Forum Meeting, Chicago, USA, 5 February 2013
A. Nowak/CERN, The Search for the Higgs Boson at CERN - Is There Such a Thing as Too Much Compute Power?, Intel Corporation, New Hampshire, USA, 7 February 2013
A. Nowak/CERN, Practical Code Optimization, SFI Poland, Krakow, Poland, 16 March 2013
A. Nowak/CERN, Software Optimization in the Many-core Era – the CERN Case, SFI Poland, Krakow, Poland, 16 March 2013
A. Nowak/CERN, The Search for the Higgs Boson at CERN – A Story of One Particle and 250,000 Cores, Nicholas Copernicus University, Toru, Poland, 18 March 2013
S. Jarp/CERN, What might be Good Software Designs for the Complexity of Current CPUs, Accelerators and GPUs, Workshop on GPUs in High Energy Physics, DESY, Hamburg, Germany, 16 April 2013
S. Jarp/CERN, What might be Good Software Designs for the Complexity of Current CPUs, Accelerators and GPUs, LHCb Many-core Working Group, CERN, 19 April 2013
S. Jarp/CERN, Solving the Mysteries of The Universe With Big Data, Big Data Innovation Summit, London, UK, 30 April 2013
S. Jarp/CERN, The Struggle to Design Software that fills the Performance Dimensions of Modern CPUs, Parallel 2013, Karlsruhe, Germany, 16 May 2013
A. Nowak/CERN, Opportunities and Choice in a new Vector Era, 15th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2013), Beijing, China, 18 May 2013
A. Nowak/CERN, Big Data at the Service of Big Science at CERN, BDigital Global Congress 2013, Barcelona, Spain, 13 June 2013
A. Nowak/CERN, Big Science and Bigger Data - the Growth of Computing at CERN, International Supercomputing Conference (ISC'13), Leipzig, Germany, 18 June 2013
A. Nowak/CERN, Profiling of Large Scientific Applications, International Symposium on Computer Architecture 2013, Tel Aviv, Israel, 23 June 2013
B. Jones/CERN, Big Science Meets Big Data, Forum TERATEC 2013, Bruyères-le-Châtel, France, 25 June 2013
Y. Almalioglu/CERN, A. Azzarà/CERN, V. Doneva/CERN, T. Dul/CERN, J. J. Fumero Alfonso/CERN, B. Kolobara/CERN, A. Krajewski/CERN, Z. Kraljevic/CERN, C. Lindqvist/CERN, G. McGilvary/CERN, D. Michelino/CERN, A. Montes Lopez/CERN, A. R. Pătrăşcoiu/CERN, R. Rama-Ballesteros/CERN, A. Rodriguez Peón/CERN, A. Sharma/CERN, V. Vintila/CERN, J. Zhang/CERN, CERN openlab Summer Student Lightning Talks, CERN, 27 August 2013
S. Jarp/CERN, Solving the Mysteries of the Universe with Big Data, Big Data Innovation Summit, Boston, USA, 12 September 2013
Z. Baranowski/CERN, Hadoop Map Reduce vs Oracle Parallel Query for Sequential Data Access, IT Technical Forum, CERN, 13 September 2013
Alberto Di Meglio/CERN, The Challenges of Scientific Computing, Huawei IT Leaders Forum – Amsterdam, The Netherlands, 18 September 2013
T. Cass/CERN, Capture, Organize, and Analyze Big Data for Research, Oracle OpenWorld 2013, San-Francisco, USA, 25September 2013
E. Grancher/CERN, Consolidation without Tears! Oracle 12c and DB Replay, Oracle OpenWorld 2013, San-Francisco, USA, 24 September 2013
M. Limper/CERN, An SQL-based Approach to Physics Analysis, International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, Amsterdam, The Netherlands, 14 October 2013
S. Heikkila/CERN, Evaluation of the Huawei UDS Cloud Storage System for CERN Specific Data, International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, Amsterdam, The Netherlands, 14 October 2013
Z. Baranowski /CERN, Sequential Data Access with Oracle and Hadoop: a Performance Comparison, International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, Amsterdam, The Netherlands, 15 October 2013
P. Szostek/CERN, Beyond Core Count: a Look at new Mainstream Computing Platforms for HEP Workloads, International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, Amsterdam, The Netherlands, 15 October 2013

S. A. Russo/CERN, Running a Typical ROOT HEP Analysis on Hadoop/MapReduce, International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, Amsterdam, The Netherlands, 17 October 2013

A. Nowak/CERN, Does the Intel Xeon Phi Processor fit HEP Workloads?, International Conference on Computing in High Energy and Nuclear Physics (CHEP) 2013, Amsterdam, The Netherlands, 17 October 2013

A. Nowak/CERN, An Update on Software for Parallelism and Heterogeneity, ATLAS Software Week, CERN, 23 October 2013

S. Jarp/CERN, Exascale in 2020 - Which Components might be Relevant for the CERN Experiments and the LHC Computing Grid?, Intel Exascale Meeting 2013, Nice, France, 23 October 2013

T. Bell/CERN, T. Owen/Rackspace, Towards Hybrid OpenStack Cclouds in the Real World, Fall OpenStack Summit, Hong Kong, 5 November 2013

F. Flückiger/CERN, Free software at CERN: where are we, where are we going?, CERN, 15 November 2013

Z. Baranowski /CERN, Next Generation GoldenGate vs. Streams for Physics Data, UK Oracle User Group Technology Conference 2013, Manchester, England, 2 December 2013

M. Blaszczyk/CERN, L. Canali/CERN, Storage Latency for Oracle DBAs, UK Oracle User Group Technology Conference 2013, Manchester, England, 2 December 2013

M. Blaszczyk/CERN, L. Canali/CERN, Lost Writes, a DBA's Nightmare?, UK Oracle User Group Technology Conference 2013, Manchester, England, 2 December 2013

L. Rodriguez Fernandez/CERN, WebLogic as a Service Provider for CERN Web Applications: APEX & Java EE, UK Oracle User Group Technology Conference 2013, Manchester, England, 4 December 2013

A. Dumitru/CERN, Oracle Database 12c Application Continuity, Swiss Oracle User Group, Lausanne, Switzerland, 5 December 2013

Publications

M. Zotes Resines/CERN, S. S. Heikkila/CERN, D. Duellmann/ CERN, G. Adde/CERN, R. Toebbicke/CERN, J. Hughes/CERN, L. Wang/CERN, Evaluation of the Huawei UDS Cloud Storage System for CERN Specific Data, 20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013), Amsterdam, Netherlands, October 2013

Posters

M. Gaillard/CERN, I. Georgopoulos/CERN, B. Jones/CERN, S. Medykowski/CERN, A. Nowak/CERN, The ICE-DIP Project, EU Commissioner for Education, Culture, Multilingualism and Youth Visit to CERN, CERN, 15 April 2013

M. Limper/CERN, An SQL-based Approach to Physics Analysis, International Conference on Computing in High Energy and Nuclear Physics (CHEP), Amsterdam, The Netherlands, 18 October 2013

CERN openlab Reports

M. Hulboj/CERN, V. Lapadatescu/CERN, Wireless Control and Optimisation, February 2013

U. Argawal/Summer Student, Linux Firewall State Synchronization, September 2013

Y. Almalioglu/Summer Student, Performance Improvements for the ATLAS Detector Simulation Framework, September 2013

A. Azzarà/Summer Student, CERN Storage Systems for Large-scale Wireless Sensor Networks, September 2013

V. Doneva/Summer Student, MatrixNet: Using a new Multivariate Technique in High Energy Physics, September 2013

T. Dul/Summer Student, Oracle AutoTask Enhancement for Multi-database Environment, September 2013

J. J. F. Alfonso/Summer Student, Vectorization with Haswell and Cilk+, September 2013

B. Kolobara/Summer Student, Electronic Ticket and Check-in System for Indico Conferences, September 2013

A. Krajewski/Summer Student, Regressive Network System Testing, September 2013

Z. Kraljevic/Summer Student, Integration of Cloud Services with Invenio Digital Library, September 2013

C. Lindqvist/Summer Student, Improved Metrics Collection and Correlation for the CERN Cloud Storage Test Framework, September 2013

G. McGilvary/Summer Student, The Implementation of OpenStack Cinder and Integration with NetApp and Ceph, September 2013

D. Michelino/Summer Student, Implementation and Test of OpenStack Heat, September 2013

A. M. Lopez/Summer Student, Improvement of the IT-PES-PS Section Services Statistics Page, September 2013

K. Panagidi/Summer Student, DBaaS with Enterprise Manager 12c and Oracle VM, September 2013

A. Radu Patrascoiu,/Summer Student, The ViSION Health Monitor, September 2013

R. Rama/Summer Student, Integration of Network Performance Monitoring Data in FTS3, September 2013

A. R. Peón/Summer Student, Advanced Visualization Tools for CERN Institutional Data, September 2013

A. Sharma/Summer Student, Development of Personal Collections in CERN Document Server (CDS), September 2013

E. A. Simon/Summer Student, Evaluation of In-memory Database Times Ten, September 2013

A. Tsikiridis/Summer Student, Integration of LHC Experiments Resource and Tools with Helix Nebula, September 2013

V. Vintila/Summer Student, Evaluation of Standard IT Tools for Industrial Controls at CERN, September 2013

J. Zhang/Summer Student, CERN Webfest and Citizen Science Challenges, September 2013

Events and Outreach

The 12 pre-college students who won the CERN Special Award at the Intel International Science and Engineering Fair (ISEF) 2013, in Phoenix, Arizona, USA, accompanied by CERN representatives, Jan Iven and Markus Schulz, who interviewed and selected them.

Creating and disseminating knowledge

The CERN openlab framework gives CERN and its sponsors a means to share a vision of the future of scientific computing.

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 2013, 181 protocol visits were organised at CERN. The CERN openlab concept and the CERN openlab sponsors' projects are systematically presented to the guests visiting the CERN IT department.

CERN openlab partners regularly organise customer and press visits. These groups are briefed about CERN openlab in a dedicated VIP meeting room known as the CERN openlab openspace. In addition, the general public is also widely introduced to the CERN openlab activities. For example, 6000 of the 70,000 visitors who came to the laboratory over the week-end of the CERN Open Days on 28 and 29 September 2013 visited the CERN data centre and were introduced to computing at CERN and to CERN openlab (via guided tours, a video http://cds.cern.ch/record/1604210 and posters).

Our sponsors' top management come to visit CERN openlab to discuss common projects with the team. As part of their visit, our guests also give computing seminars in the IT auditorium, which are sometimes webcast (www.cern.ch/webcast) and made accessible later via the CERN Document Server (CDS) and the CERN openlab website. Six of these talks were given in 2013 on topics as diverse as reproducibility of research, machine -learning, next- generation hyper-scale software and hardware systems for big-data analytics, high-performance embedded computing on the MPPA single-chip many-core processor, exascale challenges and general-purpose computing, and performance productivity challenges and research on the future of computing.

Andrzej Nowak presents the ICE-DIP poster to Androulla Vassiliou, the European Commissioner for Education, Culture, Multilingualism and Youth, during her visit to CERN on 15 April 2013.

HP visit to CERN, January 2013
Charles Clark, Director of Research HP Networking.

European Commissioner visit to CERN, 15 April 2013
Androulla Vassiliou, European Commissioner for Education, Culture, Multilingualism and Youth. Introduction to the ICE-DIP project through a poster session.

Visit of the Milanese winners of the Junior Achievement Competition (sponsored by Intel), 09 April 2013
Award winners: Valentina Boldo, Lorenzo Bonni, Eric Di Francia, Silvia Mariotti, Vittoria Pecoriello, Valeria Serenthà.

Huawei visit to CERN, 18 April 2013
Leon He, President Huawei Western Europe Enterprise Business Unit, Davis Wu, Peter Balsiger, Sales Manager Storage & Data Centre of Huawei.

Intel ISEF students visit to CERN, 22 June to 28 June 2013
CERN Award Winners 2013: Colin Aitken, Elisabeth Ashmore, Ionut Budisteanu, Vincent Cao, Jennifer Csele, Zeyu Liu, Dhaivat Pandya, Ema Parker, Ashwin Ramachandran, Valerie Sarge, Yousuf Soliman, Fabian Tschopp.

Oracle staff visit to CERN (20 people), 26 June 2013
Introduction to CERN and CERN openlab, visit of the ATLAS visitor centre and the underground facilities.

Google students visit to CERN, 31 July 2013
Introduction to CERN openlab and visit of the CERN facilities.

HP visit to CERN, September 2013
Rex Pugh, R&D Manager at HP Networking.

CERN Open Days public visits to CERN, 28 to 29 September 2013
Out of the 70,000 visitors, 6,000 people came to visit the CERN data centre over the week-end and were introduced to computing at CERN and CERN openlab.

Siemens ETM visit to CERN, 29 October 2013
Introduction to CERN openlab and visit of the CERN facilities.

Out of the 70,000 people who visited the laboratory during the CERN Open Days, 6000 came to visit the CERN data centre over the weekend and were introduced to computing at CERN and to CERN openlab.

Huawei visit to CERN, 9 December 2013
Peter Balsiger, Sales Manager Storage & Data Centre, WK Leung, Chief Technology Officer Solution - Sales Department, Jorge Rodrigues, Account Manager.

Seminars by CERN openlab guests

Towards Reproducibility of Research by reuse of IT Best Practices, CERN, 7 August 2013. Given by Andrey Ustyuzhanin, Yandex Liaison Person with CERN openlab, as part of his visit.

Performance Productivity Challenges and Researches for the Future of Computing, CERN, 30 August 2013. Given by Victor Lee, Research Scientist at Intel's Parallel Computing Lab, as part of his visit.

Increasing Predictability of Machine-learning Research, CERN, 16 September 2013. Given by Andrey Ustyuzhanin and Artem Vorozhtsov, Yandex, as part of their visit.

Next Generation Hyper-scale Software and Hardware Systems for Big Data Analytics, CERN, 19 September 2013. Given by Rich Friedrich, Director of Strategic Innovation and Research Services (SIRS) at HP Labs, as part of his visit.

High Performance Embedded Computing on the MPPA Single Chip Manycore, CERN, 1 October 2013. Given by Benoit Dupont de Dinechin, Chief Technology Officer at KALRAY, as part of his visit.

Exascale Challenge and General Purpose Computing, CERN, 24 October 2013. Given by Avinash Sodani, Intel Chief Architect of the future Xeon-Phi processor, as part of his visit.

The Future

CLEX (the CLIC Experimental area).

Towards the next chapters

High-Luminosity LHC project, post-LHC projects and next CERN openlab phase get under way.

The LHC's performance has been fantastic and since the first physics run in 2009 it has provided physicists with a huge quantity of data to analyse. After a short run in 2013, the LHC was shut down for two years to prepare the machine for operation at higher energy. Despite the complexity of the work on the 27-km-long LHC, there has been excellent progress in 2013: CERN's accelerator complex, the LHC experiments, and the injectors are on course to be ready to start the next LHC run in 2015 at 13 TeV.

With the LHC restart just a year away, the future of the accelerator is as hot a topic as ever. In terms of CERN's short-term future, the special Council session for the approval of the European Strategy for Particle Physics, held in Brussels in May 2013, gave the go-ahead for a major construction project, the High-Luminosity LHC (HL-LHC), to begin in the next decade. This ambitious project is designed to increase the LHC's integrated luminosity by a factor of ten.

In addition to endorsing the HL-LHC construction programme, the European Strategy session also gave the green light for feasibility studies for two projects that will be fundamental to the future: the Compact Linear Collider (CLIC) and the Future Circular Collider (FCC), a circular machine with a circumference of around 100 kilometres that would be capable of reaching an energy of 100 TeV.

In this exciting context, a workshop on future IT challenges in scientific research was organised by CERN openlab in December 2013 to pave the way for its next three-year phase. For the first time in its history, the preparation of the next CERN openlab phase was extended beyond the physics community. The workshop brought together more than 40 participants from CERN; the LHC experiments; as well as a number of European laboratories, such as EMBL-EBI, ESA, ESRF, and ILL; researchers from the Human Brain Project; and leading IT companies. The participants shared their views and input to identify use cases to consider during CERN openlab phase V, due to start in 2015.

Twelfth Annual Report of CERN openlab - Second Annual Report of phase IV
Design and Production: CERN Graphic Design Service
Editor: Mélissa Gaillard
Photo CERN Photo Service, except: Sylvain Fasel (p 29), Ioannis Georgopoulos (p 55), Huawei (p 45), HP (p 33), Intel (p 35), Vlad Lapadatescu (p 31), LHCb (p 51), Veronika McQuade (p 23), David Merle (p 41), Andrzej Nowak (p 37, p 53), Fons Rademakers (p 11), Shutterstock (p 1, p 2, p 4, p 6, p 10, p 16, p 22, p 28, p 34, p 40, p 46, p 50, p 52, p 58, p 62, p 64, p 67, p 68), Siemens (p 17), Society for Science & the Public (p 59)
For an electronic version see www.cern.ch/openlab/resources/annual-reports

ISBN 978-92-9083-409-0

CERN openlab

www.cern.ch/openlab