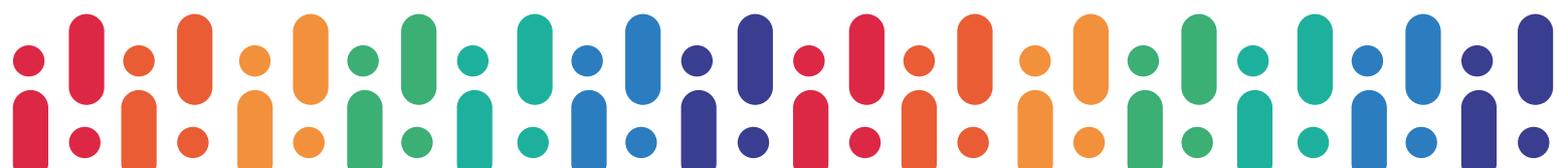




# ANNUAL REPORT

---

2017



<b>INTRO 01</b>	
A word from the DG	4
<b>CONTEXT 02</b>	
Background	6
<b>HIGHLIGHTS 03</b>	
2017 at CERN	10
<b>ABOUT 04</b>	
The concept	12
<b>RESULTS 05</b>	
Extreme Flow Optimizer	20
RapidIO for data acquisition	22
EOS productisation	24
REST services, JavaScript, and JVM performance	26
Oracle WebLogic on Kubernetes	28
OpenStack clouds	30
OpenStack containers	32
Oracle cloud	34
High-throughput computing collaboration	36
Code modernisation: GeantV	38
Industrial control and monitoring	40
Intel big-data analytics	42
Oracle big-data analytics	44
Data popularity and anomaly detection	46
BioDynaMo	48
GeneROOT	50
<b>KNOWLEDGE 06</b>	
Education, training, and outreach	52
<b>FUTURE 07</b>	
Next steps	58

# 01

## INTRO

## A word from the DG

CERN openlab is now embarking upon a new three-year phase. Together with its collaborators from research and industry, CERN openlab is working to accelerate the development of the innovative systems needed to meet the unprecedented ICT challenges posed by CERN's current and future scientific programme.



*Fabiola Gianotti, Director-General of CERN.*

Since its foundation in 2001, CERN openlab has run in successive three-year phases. This model offers the possibility for both CERN and the collaboration members to take stock at regular intervals, and helps to ensure that the most pertinent challenges faced by the worldwide high-energy physics community are being addressed. The fact that several leading ICT companies have signed up as CERN openlab members for multiple successive phases is testament to the value this unique public-private partnership offers.

CERN openlab's fifth phase came to a close at the end of 2017. This three-year period marked a significant time of growth for CERN openlab, with the number of industry collaborators roughly doubling. Also, for the first time, the CERN openlab family included other research institutes. By working together with laboratories and universities facing similar ICT challenges to those encountered by the LHC community, CERN openlab has been able to further increase its impact and ensure maximum relevancy of its work.

During its fifth phase, CERN openlab ran 20 R&D projects, related to ICT topics such as data acquisition, networking, storage, cloud computing, code modernisation, and data analytics. The insights gained through these investigations are already helping to inform the development work of teams across the laboratory as they prepare for the ICT challenges posed by future planned upgrades to both the LHC and the experiments.

Throughout 2017, much of CERN openlab's work was dedicated to preparations for its sixth phase, which runs from the start of 2018 to the end of 2020. A series of workshops and other meetings was held in the first half of the year to discuss the ICT challenges faced by the LHC research community — and other 'big science' projects — over the coming years. The white paper CERN openlab published in September (<http://openlab.cern/whitepaper>) is the culmination of these investigations, and sets out specific challenges that are ripe for tackling through collaborative R&D projects with leading ICT companies.

With the publication of this white paper, CERN openlab has laid the groundwork for another highly successful three-year phase. In the coming years, I expect CERN openlab to capitalise on the potential areas for joint R&D that have been identified and to build on its collaborations with research communities beyond high-energy physics.

# 02

## CONTEXT

## Background

Founded in 1954, the CERN laboratory sits astride the Franco-Swiss border near Geneva. It was one of Europe's first joint ventures and now has 22 member states.

### The laboratory

At CERN, physicists and engineers are probing the fundamental structure of the universe. They use the world's largest and most complex scientific instruments to study the basic constituents of matter — the fundamental particles. The particles are made to collide at close to the speed of light. This process gives the physicists clues about how the particles interact, and provides insights into the fundamental laws of nature.

The instruments used at CERN are purpose-built particle accelerators and detectors. Accelerators boost beams of particles to high energies before the beams are made to collide with each other or with stationary targets. Detectors observe and record the results of these collisions.

CERN is home to the Large Hadron Collider (LHC), the world's largest and most powerful particle accelerator. It consists of a 27-kilometre ring of superconducting magnets, with a number of accelerating structures to boost the energy of the particles along the way.

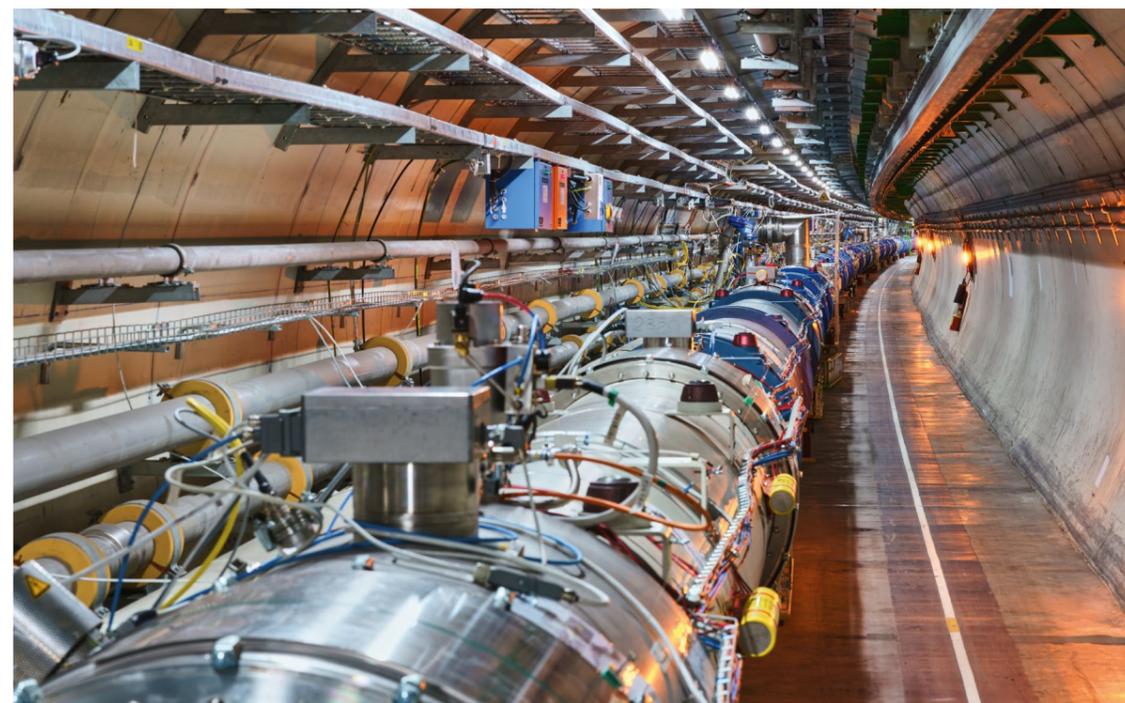
### The LHC

The accelerator complex at CERN is a succession of machines that accelerate particles to increasingly high energy levels. Each machine boosts the energy of a beam

of particles, before injecting the beam into the next machine in the sequence. In the LHC — the last element in this chain — particle beams are accelerated up to the record energy of 6.5 teraelectronvolts (TeV) per beam. Most of the other accelerators in the chain have their own experimental halls where beams are used for experiments at lower energies.

The proton source is a simple bottle of hydrogen gas. An electric field is used to strip hydrogen atoms of their electrons to yield protons. Linac 2, the first accelerator in the chain, accelerates the protons to the energy of 50 MeV. The beam is then injected into the Proton Synchrotron Booster (PSB), which accelerates the protons to 1.4 GeV, followed by the Proton Synchrotron (PS), which pushes the beam to 25 GeV. Protons are then sent to the Super Proton Synchrotron (SPS) where they are accelerated to 450 GeV.

The protons are finally transferred to the two beam pipes of the LHC. Inside the accelerator, two high-energy particle beams travel at close to the speed of light before they are made to collide. The beams travel in opposite directions in separate beam pipes: two tubes kept at an ultra-high vacuum. They are guided around the accelerator ring by a strong magnetic field maintained by superconducting electromagnets.



*The LHC is the world's largest and most powerful particle accelerator.*



Seven experiments on the LHC use detectors to analyse the myriad particles produced by collisions.

The beam in one pipe circulates clockwise while the beam in the other pipe circulates anticlockwise. It takes 4 minutes and 20 seconds to fill each LHC ring, and 20 minutes for the protons to reach their maximum energy of 6.5 TeV. Beams circulate for many hours inside the LHC beam pipes under normal operating conditions. The two beams are brought into collision inside four detectors — ALICE, ATLAS, CMS and LHCb — where the total energy at the collision point is equal to 13 TeV. The particles are so tiny that the task of making them collide is akin to firing two needles 10 kilometres apart with such precision that they meet halfway.

Protons are not the only particles accelerated in the LHC. Lead ions for the LHC start from a source of vaporised lead and enter Linac 3 before being collected and accelerated in the Low Energy Ion Ring (LEIR). They then follow the same route to maximum energy as the protons. Colliding lead particles makes it possible to recreate conditions similar to those just after the big bang, known as 'quark-gluon plasma'.

#### The experiments

Seven experiments on the LHC use detectors to analyse the myriad particles produced by collisions. These experiments are run by collaborations of scientists from institutes all over the world. Each experiment is distinct, and characterised by its detectors.

The biggest of these experiments, ATLAS and CMS, use general-purpose detectors to investigate the largest range of physics possible. Having two independently designed detectors is vital for cross-confirmation of any new discoveries made. ALICE and LHCb have detectors specialised for focusing on specific phenomena. These four detectors sit underground in huge caverns on the LHC ring.

The smallest experiments on the LHC are TOTEM and LHCf, which focus on 'forward particles': protons or heavy ions that brush past each other rather than meeting head on when the beams collide. TOTEM uses detectors positioned on either side of the CMS interaction point, while LHCf is made up of two detectors which sit along the LHC beamline, at 140 metres either side of the ATLAS collision point. MoEDAL uses detectors deployed near LHCb to search for a hypothetical particle called the magnetic monopole.

It is important to note that while the main focus of research at CERN has moved in recent years towards the LHC, experiments at other accelerators and facilities both on-site and off also remain an important part of the laboratory's activities. In fact, CERN has a very diverse research programme, covering a wide range of physics topics, from the Standard Model to supersymmetry, and from dark matter to cosmic rays. Supporting all the experiments hosted at

CERN is a very strong theory programme, which carries out world-leading research in theoretical particle physics.

#### The CERN data centre

The CERN data centre processes hundreds of petabytes of data every year. The data centre in Meyrin, Switzerland, is the heart of CERN's entire scientific, administrative, and computing infrastructure. All services — including email, scientific data management and videoconferencing — use equipment based in the data centre. The centre hosts around 10,000 dual-CPU servers with approximately 300,000 processor cores.

The servers undergo continual maintenance and upgrades to make sure that they will operate in the event of a serious incident. Critical servers are held in their own room, powered and cooled by dedicated equipment. Around 150 petabytes of data are stored on disk at the Meyrin site, with over 200 PB on tape.

#### The Wigner Research Centre for Physics

In June 2013, the Wigner data centre in Budapest was inaugurated. It serves as an extension to the main data centre in Meyrin. Today, the Wigner data centre hosts 3500 dual-CPU servers, with approximately 100,000 processor cores. It also offers around 100 PB of disk space.

The equipment in Wigner is managed and operated from CERN, in the same way as the equipment in the CERN data centre. Only activities requiring physical access to the

equipment are performed by the Wigner data centre staff, such as installation of equipment into racks, repairs to the servers, etc.

The CERN and Wigner centres are connected via three independent and dedicated 100 Gb/s fibre-optic lines, with a bandwidth equivalent to the transfer of eight full DVDs per second. Network latency (the time taken between sending data and receiving on the other end) between the two sites, which are 1800 km apart, is about 25 milliseconds.

#### The Worldwide LHC Computing Grid (WLCG)

Physicists must sift through the 30-50 petabytes (PB) of data produced annually by the LHC experiments to determine if the collisions have thrown up any interesting physics. CERN does not have the computing or financial resources to crunch all of the data produced by the LHC experiments on site, so in 2002 it turned to grid computing to share the burden with computer centres around the world. The Worldwide LHC Computing Grid (WLCG) — a distributed computing infrastructure arranged in tiers — gives a community of thousands of physicists near real-time access to LHC data.

With 170 computing centres in 42 countries, the WLCG is the most sophisticated data-taking and analysis system ever built for science. It runs more than two million jobs per day. The CERN data centre — working in unison with its extension in Budapest — forms 'Tier-0' of the WLCG, the first point of contact between experimental data from the LHC and the grid.



The CERN data centre in Meyrin is the heart of CERN's entire scientific, administrative, and computing infrastructure.

# 03

## HIGHLIGHTS

### 2017 at CERN

In 2017, scientists and engineers at CERN used their knowledge and experience gained to date to push the LHC to its limits. The LHC experiments were thus able to carry out more precise measurements than ever before and push our current theories about the universe to their limits, too.

By the start of 2017, the LHC had already been operating at a collision energy of 13 teraelectronvolts (TeV) for two years. During this time, scientists and engineers at CERN were able to build up an excellent understanding of how the LHC works, thus making it possible to optimise operations even further in 2017.

The LHC's physics season kicked off in May. Throughout the year, physicists at the LHC experiments worked on two different broad areas: improving their knowledge of known phenomena and probing the unknown. The known phenomena constitute the 'Standard Model' of particles and forces, a theory that encompasses all our current knowledge of elementary particles. A particular highlight was the observation at the LHCb experiment of a new particle (containing two 'charm quarks' and one 'up quark') from the baryon family, to which the proton also belongs.

The LHC far exceeded its targets for 2017, producing over 5 million billion collisions. Over the course of the year, the LHC controllers optimised the machines' operating parameters. Using a new system, they were able to notably reduce the size of the particle beams when they meet at the centre of the experiments. The more squeezed the beams, the more collisions occur each time they meet. In 2016, the operators managed to obtain 40 collisions at each bunch crossing, with each bunch containing 100 billion particles. In 2017, up to 60 collisions were produced at each crossing.

This helped lead to new records for data collection: during October alone, the CERN data centre stored over 12 petabytes (PB) of data. By the end of the year, the total amount of data archived on tape stood at 230 PB; this is the equivalent of around 2000 years of HD video.

For an in-depth overview of highlights from the LHC experiments in 2017, please visit the CERN website: [cern.ch/go/XFq6](http://cern.ch/go/XFq6).

Beyond the LHC, several highlights from the year related to neutrino experiments. These included the installation of prototype neutrino detectors for the DUNE experiment, the Baby MIND neutrino detector being shipped to Japan following construction at CERN, and the ICARUS neutrino detector (originally installed at INFN's Gran Sasso National Laboratory in Italy) being sent to Fermilab in the US following refurbishment work at CERN.

Other highlights in 2017 related to research into antimatter: the first antiproton beam was circulated at ELENA (a new decelerator at CERN's antimatter factory), the BASE experiment measured the magnetic moment of the antiproton with greater precision than ever before, and the ALPHA experiment reported the first observation of the hyperfine structure of antihydrogen.

In addition, the CERN-MEDICIS facility began producing radioisotopes for medical research in late 2017. These radioisotopes are destined primarily for hospitals and research centres in Switzerland and across Europe. Great strides have been made recently in the use of radioisotopes for diagnosis and treatment; the facility will enable researchers to devise and test unconventional radioisotopes, with a view to developing new approaches to fight cancer.



The LHC far exceeded its targets for 2017, producing over 5 million billion collisions.

# 04

## ABOUT

### The concept

CERN openlab is a unique public-private partnership that accelerates the development of cutting-edge ICT solutions for the worldwide LHC community and wider scientific research. Through CERN openlab, CERN collaborates with leading ICT companies and research institutes.

Within the CERN openlab framework, CERN provides access to its advanced ICT infrastructure and its engineering experience – in some cases even extended to collaborating institutes worldwide. Testing in CERN’s demanding environment provides the collaborating companies with valuable feedback on their products, while enabling 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.



Members of the CERN openlab management team. From left to right: Fons Rademakers, chief research officer; Alberto Di Meglio, head of CERN openlab; Maria Girone, chief technology officer; Georgios Vasilopoulos, web developer; Kristina Gunne, finance and administration officer; Andrew Purcell, communications officer; Marion Devouassoux, junior administration officer; Orestis Galanis, junior communications officer.

Industry collaboration can be at the associate, contributor, or partner level. Each status represents a different level of investment, with projects lasting typically between one and three years. The collaborating companies engage in a combination of cash and in-kind contributions, with the cash being used to hire young ICT specialists dedicated to the projects. The associate status formalises a collaboration based on independent and autonomous projects that do not require a presence on the CERN site. The contributor status is a collaboration based on tactical projects, which includes a contribution to hire an early-career ICT specialist supervised by CERN staff to work on the common project, in addition to the hardware and software products needed by the projects. The partners commit to a longer-term, strategic programme of work and provide three kinds of resources: funding for early-career researchers, products and services, and engineering capacity. The partners receive the full range of benefits of membership in CERN openlab, including extensive support for communications activities and access to dedicated events.

of an advanced computing-cluster prototype called the ‘opencluster’. The second phase (2006–2008) addressed a wider range of domains. The combined knowledge and dedication of the engineers from CERN and the collaborating companies produced exceptional results, leading to significant innovation in areas such as energy-efficient computing, grid interoperability, and network security. CERN openlab’s third phase (2009–2011) capitalised and extended upon the successful work carried out in the second phase. New projects were added focusing on virtualisation of industrial-control systems and investigation of the then-emerging 64-bit computing architectures. The fourth phase (2012–2014) addressed new topics crucial to the CERN scientific programme, such as cloud computing, business analytics, next-generation hardware, and security for the ever-growing number of networked devices.

CERN openlab was established in 2001, and has been organised into successive three-year phases. In the first phase (2003–2005), the focus was on the development



Attendees at the 2017 CERN openlab Collaboration Board.

This annual report covers the final year of CERN openlab's fifth phase (2015-2017). This phase tackled ambitious challenges covering the most critical needs of ICT infrastructures in domains such as data acquisition, computing platforms, data storage architectures, compute provisioning and management, networks and communication, and data analytics. The work carried out in 2017 built yet further on the progress made in these domains during the first two years of CERN openlab's fifth phase. Work in 2017 also focused on preparations for CERN openlab's sixth three-year phase. This culminated in the production of a white paper identifying future ICT challenges, which was presented at a special 'open day' event in September. More information about this white paper can be found in the final section of this annual report, dedicated to future plans.

Throughout its fifth phase, CERN openlab has grown to include more collaborating companies, thus enabling a wider range of challenges to be addressed. New research institutes have also joined CERN openlab in its endeavour to accelerate the development of cutting-edge ICT solutions for science.

During our fifth-phase, work was carried out on 20 R&D projects. Each project is supervised by a member of CERN staff, who also acts as a liaison between CERN and the

company collaborating on the project. At monthly review meetings, the project coordinators update the CERN openlab management team on the progress of projects, thus ensuring timely follow-up of any issues that may arise.

At CERN openlab's annual technical workshops, representatives of the collaborating companies and research institutes meet with the teams, who provide in-depth updates on technical status. Collaborating companies and research institutes also elect representatives for CERN openlab's annual Collaboration Board meeting, which is an opportunity to discuss the progress made by the project teams and to exchange views on the collaboration's plans.

The CERN openlab team consists of three complementary groups of people: young engineers hired by CERN and funded by our collaborators, technical experts from the companies involved in the projects, and CERN management and technical experts working partly or fully on the joint activities.

The names of the people working on each project can be found in the results section of this report. The members of CERN openlab's management team are shown in the image on the previous page.

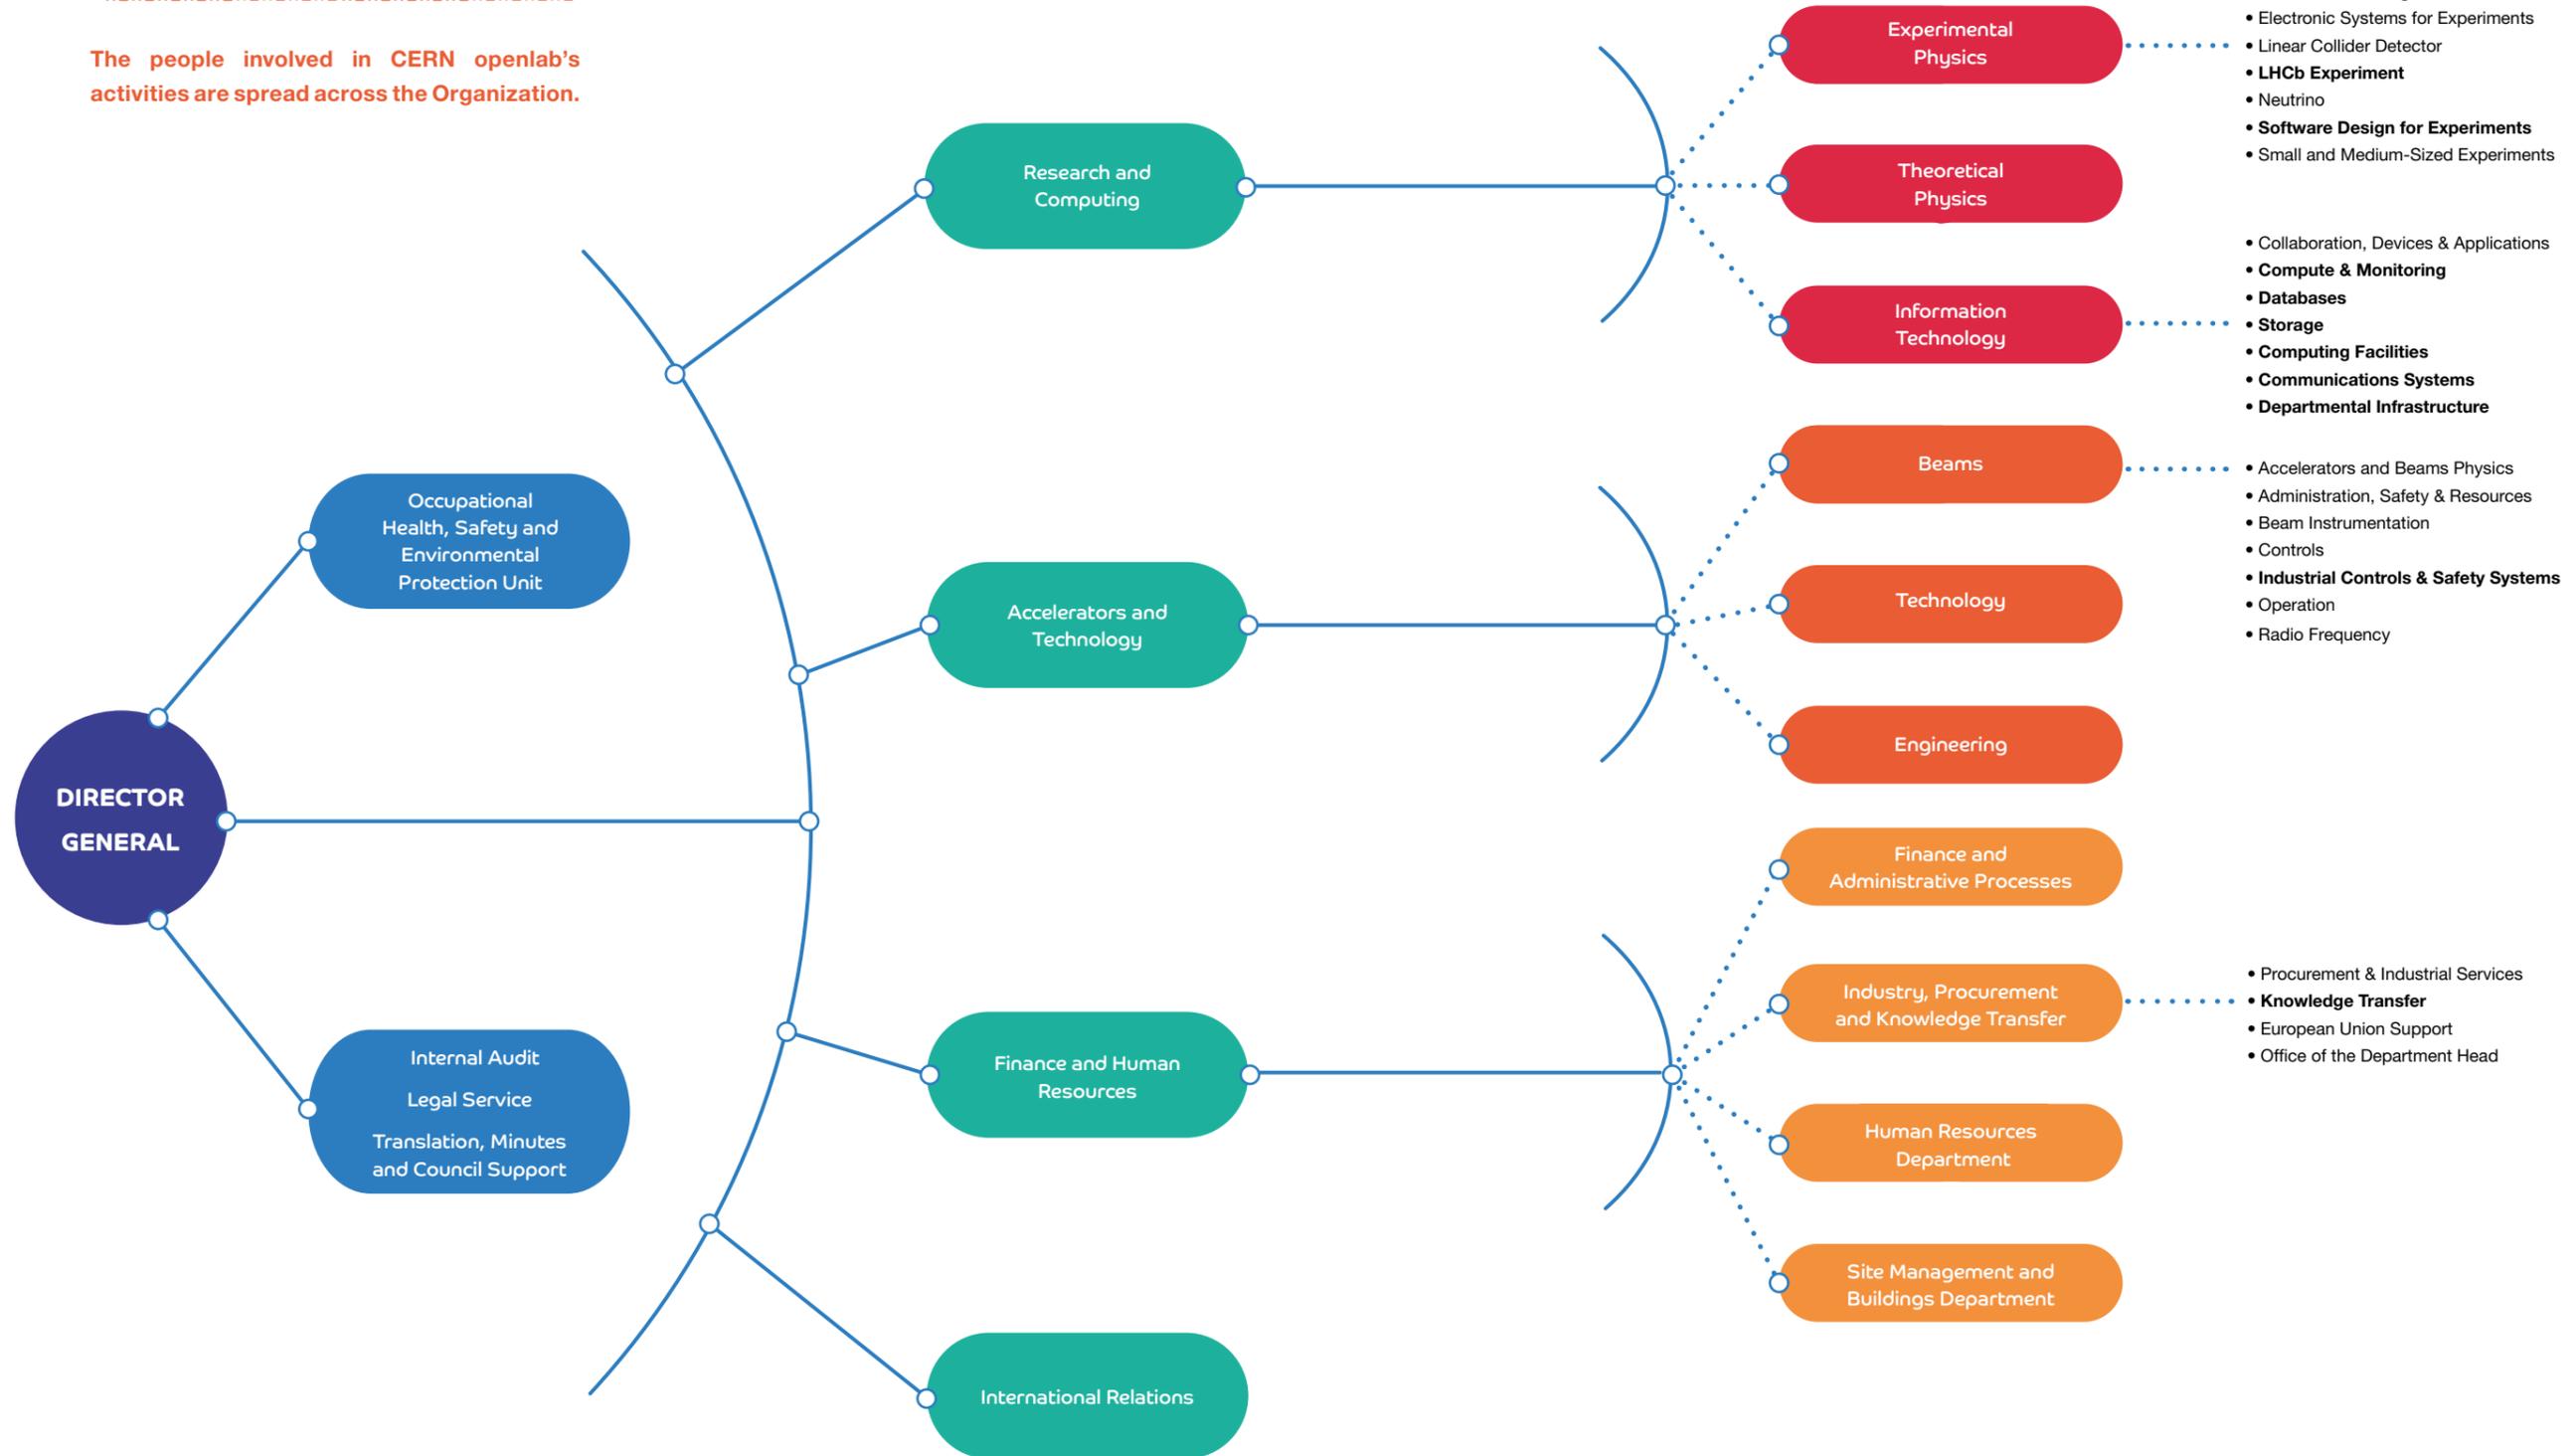


Maria Girone, CERN openlab CTO, presenting at the 2017 CERN openlab Collaboration Board.



# POSITIONING CERN OPENLAB'S ACTIVITIES AT CERN

The people involved in CERN openlab's activities are spread across the Organization.



# 05

## RESULTS

## Our projects

Information about each of the 16 technical projects that ran in 2017 can be found in this section. These projects have been organised into four overarching R&D topics.

### R&D topic 1:

#### Data-centre technologies and infrastructures

Designing and operating distributed data infrastructures and computing centres poses challenges in areas such as networking, architecture, storage, databases, and cloud. These challenges are amplified — and added to — when operating at the extremely large scales required by major scientific endeavours.

CERN is evaluating different models for increasing computing and data-storage capacity, in order to accommodate the growing needs of the LHC experiments over the next decade. All models present different technological challenges. In addition to increasing the capacity of the systems used for traditional types of data processing and storage, explorations are being carried out into a number of alternative architectures and specialised capabilities. These will add heterogeneity and flexibility to the data centres, and should enable advances in resource optimisation.

### R&D topic 2:

#### Computing performance and software

Modernising code plays a vital role in preparing for future upgrades to the LHC and the experiments. It is essential that software performance is continually increased by making use of modern coding techniques and tools, such as software-optimising compilers, etc. It is also important to ensure that software fully exploits the features offered by modern hardware architecture, such as many-core GPU platforms, acceleration coprocessors, and innovative hybrid combinations of CPUs and FPGAs. At the same time, it is of paramount importance that physics performance is not compromised in drives to ensure maximum efficiency.

### R&D topic 3:

#### Machine learning and data analytics

Members of CERN's research community expend significant efforts to understand how they can get the most value out of the data produced by the LHC experiments. They seek to maximise the potential for discovery and employ new techniques to help ensure that nothing is missed. At the same time, it is important to optimise resource usage (tape, disk, and CPU), both in the online and offline environments. Modern machine-learning technologies — in particular, deep-learning solutions applied to raw data — offer a promising research path to achieving these goals.

Deep-learning techniques offer the LHC experiments the potential to improve performance in each of the following areas: particle detection, identification of interesting events, modelling detector response in simulations, monitoring experimental apparatus during data taking, and managing computing resources.

### R&D topic 4:

#### Applications in other disciplines

The fourth R&D topic is different to the others in this report, as it focuses on applications in other disciplines. By working with communities beyond high-energy physics, we are able to ensure maximum relevancy for CERN openlab's work, as well as learning and sharing both tools and best practices across scientific fields. Today, more and more research fields are driven by large quantities of data, and thus experience ICT challenges comparable to those at CERN.

CERN openlab's mission rests on three pillars: technological investigation, education, and dissemination. Collaborating with research communities and laboratories outside the high-energy physics community brings together all these aspects.

For each of these R&D topics, we have also identified a number of key ICT challenge areas, ripe for addressing through future collaborative projects in our sixth three-year phase. More information can be found in our latest white paper: <http://openlab.cern/whitepaper>

# EXTREME FLOW OPTIMIZER

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES



### Project coordinator:

Tony Cass

### Technical team:

Adam Krajewski, Stefan Stancu,  
Edoardo Martelli, David Gutierrez Rueda

### Extreme Networks liaisons:

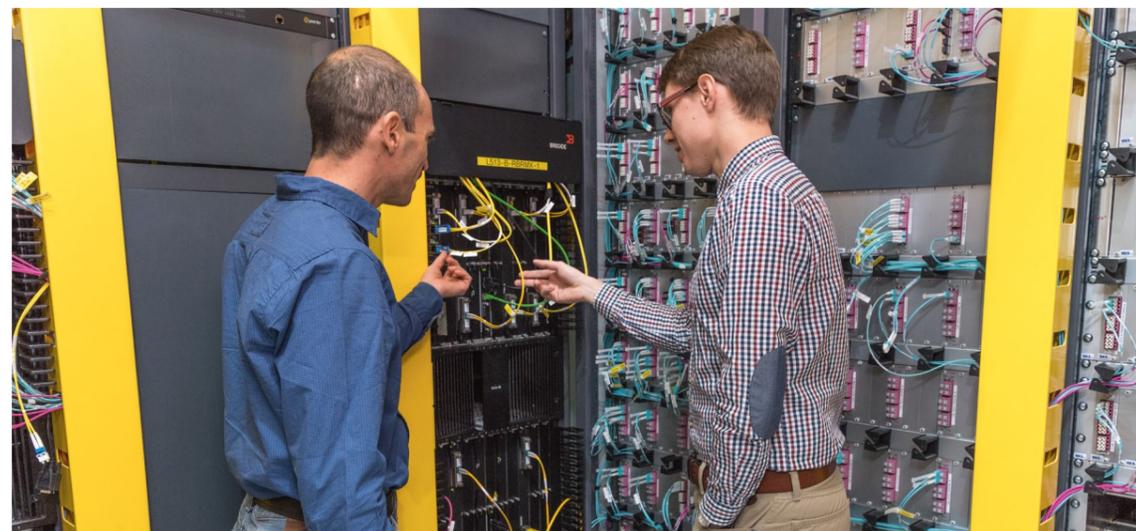
Manjunath Gowda, Murtaza Bawahir,  
Mythil Raman, Yousuf Hasan, Christoph Kaelin,  
Giacomo Bernardi, Salvador Ferrer

### Project goal

The Extreme Flow Optimizer (EFO) project aims to enhance Extreme Networks' EFO application, and to use it to build an automated network-traffic steering system. EFO is a software-defined networking (SDN) application designed to improve visibility of network traffic, manage volumetric DDoS threats, and enable dynamic flow-steering by applying SDN and network-automation principles.

### Background

As technology evolves, traditional static network configurations often prove too rigid. Through this project, we are therefore using EFO's dynamic flow-steering capabilities to provide increased programmability and flexibility in our networks. The project currently focuses on orchestrating traffic for building a scalable intrusion-detection system (IDS), with advanced features enabled by the EFO software (e.g. offloading the inspection of bulk data transfers).



Now that the evolved IDS prototype has been deployed in the CERN data centre, it will be thoroughly evaluated to ensure that all the technical requirements are met for the final production deployment.

### Progress in 2017

During 2017, we made good progress on enhancing the IDS prototype at CERN. The IDS receives a copy of the traffic that crosses CERN's network boundary and load-balances it across a pool of servers, each running the open-source Bro Network Security Monitor system. The setup was enhanced using Extreme Networks' SLX 9540 hardware platform, a high-end data-centre switch with advanced hardware capabilities for traffic orchestration. Furthermore, we used the Extreme Workflow Composer (EWC) software to provide increased automation capabilities through modular and configurable workflows, thus abstracting the configuration of network devices. This upgraded technology stack plays a key role in fulfilling both current and future system requirements.

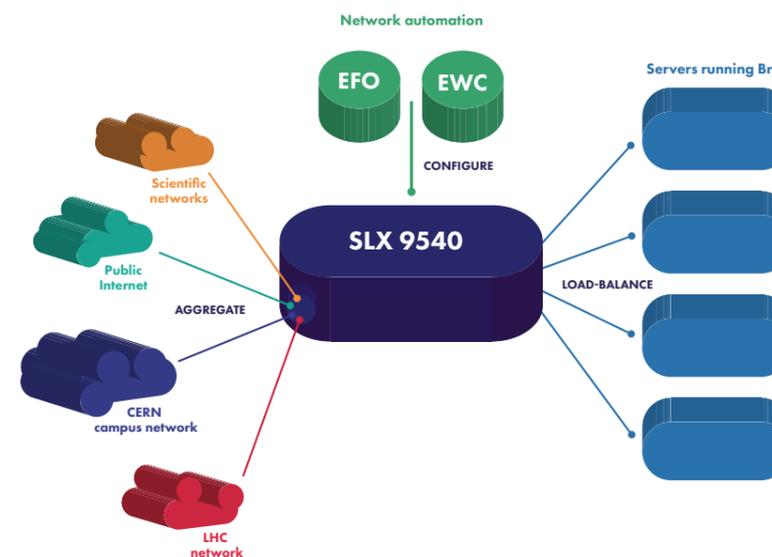
Throughout the year, we also made significant contributions to the general EFO product development; the research fellow funded by Extreme Networks and based at CERN is integrated in the EFO development team. Several features developed at CERN were included in two official EFO releases. In addition, we provided expert consultancy to two Internet service providers (ISPs) in Switzerland that are considering deploying EFO in their production systems.

### Next steps

Now that the evolved IDS prototype has been deployed in the CERN data centre, we will thoroughly evaluate it to ensure that all the technical requirements are met for the final production deployment. Contributions to EFO software development will continue, including the implementation of specific features required for the IDS and other possible CERN use cases.

### Presentations

- M. Abdullah, Network Automation with Brocade Workflow Composer (15 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- A. Krajewski, Lightning talk: Brocade Flow Optimizer (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- A. Krajewski, Network Automation for Intrusion Detection System (18 October), Presented at HEPIX Fall 2017 Workshop, KEK, Tsukuba, Japan, 2017.



High-level architecture of the intrusion-detection system.

# RAPID IO FOR DATA ACQUISITION

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES



### Project coordinators:

Olof Barring

### Technical team:

Sima Baymani, Konstantinos Alexopoulos

### IDT liaisons:

Devashish Paul, Barry Wood, Ian Thomas,  
Mohammad Akhter, Alexandre Bounine, Trevor Hiatt

### Project goal

We are exploring RapidIO interconnect technology and working to determine its suitability for data-analysis and data-acquisition applications.

### Background

RapidIO is an open-standard system-level interconnect, which is today used in all 4G/LTE base stations worldwide. RapidIO is often used in embedded systems that require high reliability, low latency, low energy consumption, and scalability in a heterogeneous environment. The collaboration with IDT, a major provider of RapidIO technology, ran from mid-2015 to mid-2017.

High-bandwidth, low-latency interconnects are currently not used much at CERN. However, developments with the online data-acquisition systems used at the LHC experiments have driven our desire to investigate this technology. It also holds potential for use in the analysis of large sets of unstructured data (text files, logs, monitoring data, etc.) in the CERN data centre.

### Progress in 2017

During 2016, we ported ROOT — a data-processing framework created at CERN and used widely in the high-energy physics community — to RapidIO. We also ported the LHCb experiment's benchmark DAQPIPE (an interconnect-agnostic application that emulates the behaviour in the LHCb data-acquisition network) to RapidIO. In early 2017, we finalised our benchmarking and scalability measurements for each of these areas of work.

### Next steps

Through our investigations, we have found RapidIO to be a promising technology in terms of delivering high-bandwidth, low-latency fabric interconnects. As such, based on the roadmaps currently set out for this technology, it could potentially play a future role in the demanding data-

acquisition infrastructure of the LHC experiments. However, it is important to note that at the time we carried out our investigations, the effective bandwidth provided by available and planned routers/cards was still lower (< 100 Gb/s) than that expected to be required by the LHC experiments in coming years (> 200 Gb/s).

### Publications

- S. Baymani, K. Alexopoulos, S. Valat, Exploring RapidIO Technology Within a DAQ System Event Building Network, *RealTimeConference(RT)*, 2016IEEE-NPSS(2017), Vol. 64, <https://doi.org/10.1109/RTC.2016.7543124>.
- S. Baymani, K. Alexopoulos, S. Valat, RapidIO as a multi-purpose interconnect, *Journal of Physics: Conference Series* (2017), Vol. 898, <https://doi.org/10.1088/1742-6596/898/8/082007>.

### Presentations

- S. E. Fitsimones, A monitoring and management utility for RapidIO clusters (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- O. Barring, IDT RapidIO data acquisition (21 September), Presented at CERN openlab Open Day, Geneva, 2017.



RapidIO also holds potential for use in the analysis of large sets of unstructured data in the CERN data centre.

# EOS PRODUCTISATION

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES



### Project coordinator:

Xavier Espinal Curull

### Technical team:

Elvin Sindrilaru, Luca Mascetti

### Comtrade liaisons:

Goran Garevski, Boris Podboj, Igor Doko,  
Ivan Arizanovic, Branko Blagojevic

### Project goal

The scope of this project — undertaken in collaboration with CERN openlab associate member Comtrade Software — is the evolution of CERN's EOS large-scale storage system. The goal is to simplify the usage, installation, and maintenance of the system, as well as adding support for new platforms.

### Background

EOS is a disk-based, low-latency storage service developed at CERN. With its highly scalable hierarchical namespace, and data access possible via the XROOT protocol (to provide scalable, fault-tolerant access), it was initially designed for the storage of physics data. Access to physics data among the LHC community is typically characterised by many concurrent users, a significant fraction of random data access, and a large file-open rate. Today, however, EOS provides storage for both physics and other use cases.

### Progress in 2017

In 2016, we created a robust installation kit to enable rapid installation of EOS both for evaluation purposes and for fast

deployment in production. In 2017, testing of this beta version of EOS continued at a number of external institutes beyond CERN, such as the European Commission's Joint Research Centre in Ispra, Italy, the University of Vienna, Austria, and the University of Trieste, Italy. At CERN, we also worked on a number of further developments to EOS throughout the year.

### Next steps

We will continue our evolving work on EOS testing, installation, and documentation. We aim to provide users with the full set of EOS functions: sync-and-share (through a tool we call CERNBox), erasure-coding, and the use of geographically distributed multi-site instances. An additional goal for us is to provide a testing framework that will run simultaneously after every release has been built, so as to certify each EOS version as accurately and as quickly as possible.

On top of this, we will also work to integrate new Comtrade Software engineers into the development and operations team at CERN, helping them to gain experience and autonomy in operations, maintenance, and development for EOS, so as to be able to provide direct support.

### Publications

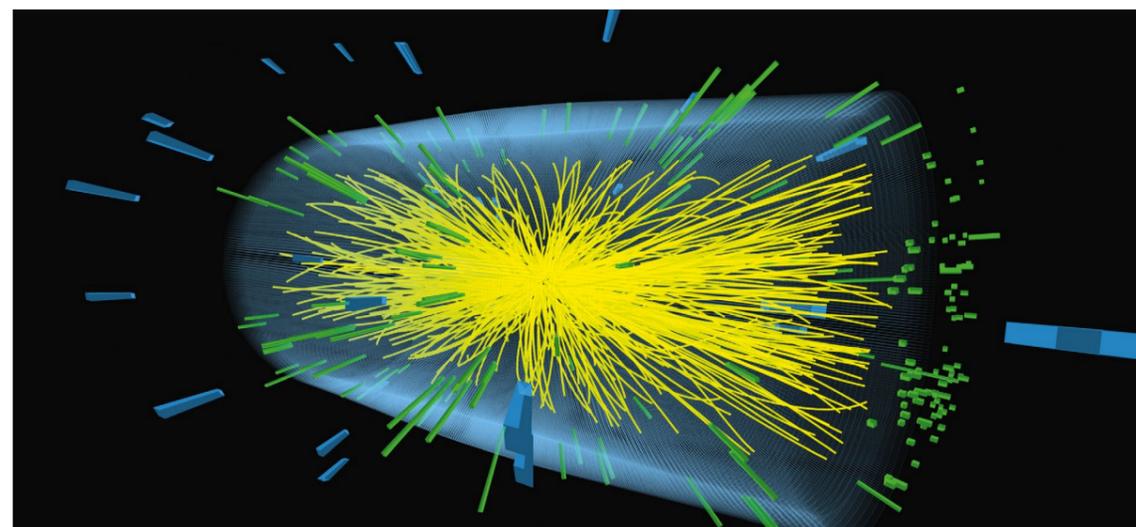
- X. Espinal, M. Lamanna, From Physics to industry: EOS outside HEP, *Journal of Physics: Conference Series* (2017), Vol. 898, <https://doi.org/10.1088/1742-6596/898/5/052023>.

### Presentations

- X. Espinal, Comtrade Software EOS Productisation (21 September), Presented at CERN openlab Open Day, Geneva, 2017.



EOS — the system at the heart of this project — is a disk-based, low-latency storage service developed at CERN.



A particle collision at the CMS experiment.

# REST SERVICES, JAVASCRIPT, AND JVM PERFORMANCE

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES

ORACLE

### Project coordinators:

*Eric Grancher, Eva Dafonte Perez*

### Technical team:

*Artur Wiecek, Luis Rodríguez Fernández, Antonio Nappi*

### Oracle liaisons:

*Pauline Mahrer, Simone Indelicato, Cris Pedregal*

### Project goal

The aim of this project is to test several Oracle technologies and evaluate their potential for use in production systems at CERN. The technologies currently being tested within the framework of this project are Oracle REST Data Services (ORDS), Oracle JavaScript Extension Toolkit (JET), and Oracle Management Cloud (OMC). A key part of our investigation into these technologies involves providing Oracle's engineering teams with important feedback that can be used to further improve them.

### Background

ORDS is widely used at CERN; we use it to serve more than 450 Oracle APEX applications. Many developers are very interested in exploring its REST capabilities and the potential for dramatically simplifying the implementation of APIs. Initial evaluation of Oracle JET has also been carried out at CERN; it is a promising candidate for potentially replacing some legacy front-end applications used by the CERN IT department's group for database services. Equally, OMC's adaptability, in terms of use with both on-premises and cloud resources, makes it a promising monitoring technology.

### Progress in 2017

Thanks to the collaboration of several engineering groups at CERN, three production systems now make use of ORDS's REST capabilities. These are CERN's access-management and distribution system, CERN's system for handling digital documents, and the system CERN uses for car sharing. We carried out much work in this area in 2017, particularly in relation to setup automation, configuration, and upgrading the various ORDS installations.

The results of our evaluation of Oracle JET have been summarised in an extended post by one of the CERN openlab summer students on the blog of the CERN IT department's group for database services. Regarding OMC, CERN engineers primarily focused efforts on extending capabilities

related to the monitoring of application performance. Valuable feedback was compiled in a range of reports, with results being discussed with Oracle engineers and product managers at the Oracle Open World 2017 event in October.

### Next steps

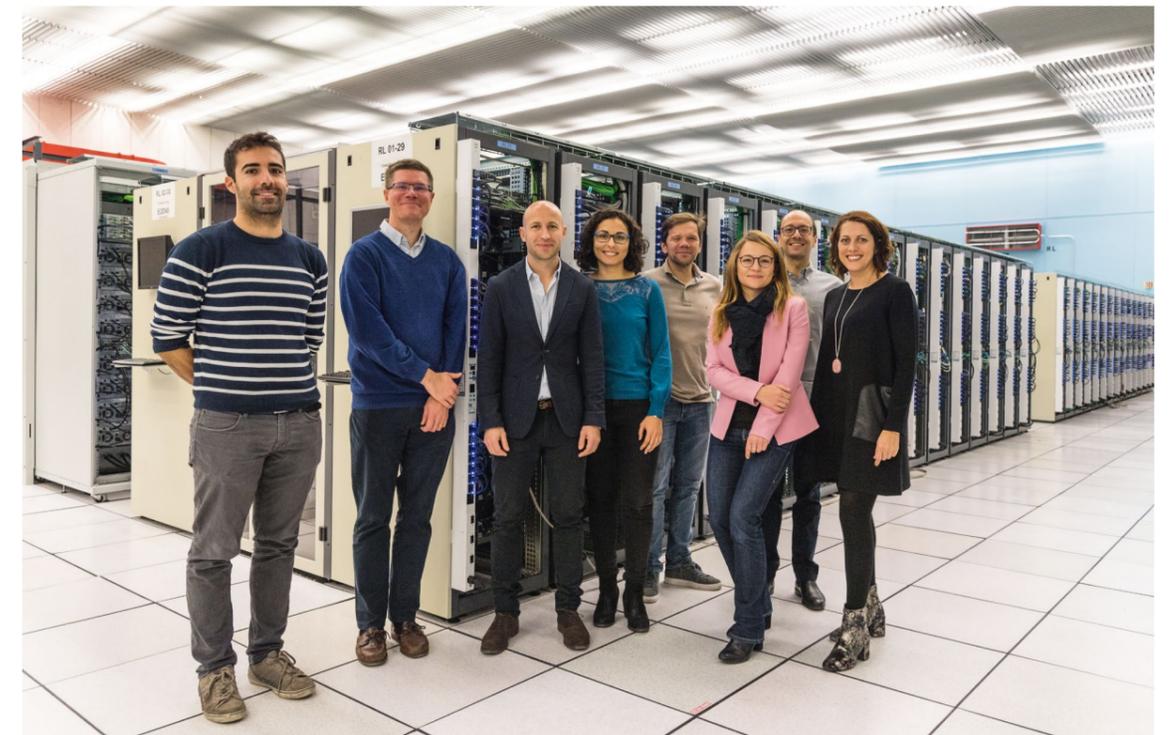
While ORDS is now fully in production, there are still some aspects we plan to investigate further, particularly related to security. We will also evaluate the promising logging, monitoring, and analytics features offered by OMC. Finally, although other strong competitors exist in terms of JavaScript, it would still be valuable to carry out further investigation into Oracle JET.

### Publications

- R. Alsvaag. My experiences from the CERN openlab Summer Student Programme 2017. Databases at CERN blog. 2017.
- L. Rodríguez Fernández. Oracle JET, ORDS & OAUTH2. Databases at CERN blog. 2017.
- L. Rodríguez Fernández, 1000 Things you always want to know about SSO but you never dare to ask, (8 October), Presented at 16th International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS 2017), Barcelona, 2017.

### Presentations

- R. Alsvaag, Refactor of Rota application using modern developer techniques: Oracle Restful Data Services & Oracle Jet (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- L. Rodríguez Fernández, REST services, JavaScript, and JVM performance with Oracle (21 September), Presented at CERN openlab open day, Geneva, 2017.
- L. Rodríguez Fernández, 1000 Things you always want to know about SSO but you never dare to ask (8 October), Presented at 16th International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS 2017), Barcelona, 2017.



*Members of the CERN IT department's database services group are collaborating with Oracle on four separate projects through CERN openlab.*

# ORACLE WEBLOGIC ON KUBERNETES

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES

# ORACLE

### Project coordinator:

Artur Wiecek

### Technical team:

Artur Wiecek, Antonio Nappi, Luis Fernandez Rodriguez, Borja Aparicio Cotarelo, Lukas Gedvilas, Nicolas Bernard Marescaux, Aimilios Tsouvelekakis

### Oracle liaisons:

Paul Jenkins, Martin Hogg, Thom Leggett, Monica Riccelli, Pauline Mahrer, Cris Pedregal

### Project goal

With the number of applications in use at CERN ever increasing, we need to extend our disaster recovery plan – originally conceived for databases – to also include our infrastructure based on Oracle WebLogic. We are therefore working to evaluate technologies like Kubernetes, so as to make it quicker and easier to build our middleware services at CERN. Kubernetes is an open-source system for automating deployment, scaling, and management of containerised applications. By using its functions for federating resources, we can also potentially simplify any migration from private, on-premises cloud resources to commercial, public clouds.

### Background

In a world where users always want to be connected and able to access their resources at any moment, reliability and 'zero downtime' are becoming first-class requirements. Today, thanks to the diverse range of cloud technologies available, achieving this is becoming easier than ever before. Nevertheless, it is important that we fully explore issues related to privacy, security, logging, monitoring, etc. Following an initial phase of thorough investigation and testing with Kubernetes, we believe that Oracle Cloud and Oracle WebLogic are good candidates for fulfilling our needs.

### Progress in 2017

After starting with a Kubernetes cluster prototype, our current testing environment now uses OpenStack Magnum as an on-premises solution for managing CERN's private cloud. In 2017, we built a CERN image of Oracle WebLogic, selected ELK stack (Elastic, Logstash, Kibana) as a solution for log collection and visualisation, and implemented TIG stack (Telegraf, InfluxDB, Grafana) as solution for gathering metrics. We also conducted investigations into HAProxy, an open-source load-balancer; these suggest that it would be an efficient and reliable solution. Security is of vital importance too: all of the external traffic uses transport-layer security, so we are working hard to implement this internally.

In addition, we are collaborating with various engineering teams from Oracle, providing feedback on solutions (like the beta version of Oracle's Kubernetes Engine) that could be used to help us to set up Oracle WebLogic infrastructure and test CERN applications in the Oracle Cloud. Initial investigation into live migration returned positive results: we managed to run and switch an Nginx web server from the CERN private cloud to the Oracle Cloud.

### Next steps

In the first months of 2018, we will deliver on-premises clusters to our community of developers at CERN for testing; with additional testing to be carried out in Oracle Cloud later in the year. These tests will be crucial in informing the development of our disaster recovery plan for Oracle WebLogic at CERN. In parallel, we will begin investigation of Spinnaker; this is an open-source delivery platform that could be used to help us to deploy applications and create Kubernetes Clusters across different cloud platforms.

### Publications

- A. Nappi. HAProxy High Availability Setup. Databases at CERN blog. 2017.
- A. Nappi. HAProxy Canary Deployment. Databases at CERN blog. 2017.

### Presentations

- A. Nappi, L. F. Rodriguez, WebLogic on Kubernetes (17 January), Presented at CERN openlab meeting with Oracle in Geneva, Geneva, 2017.
- A. Nappi, L. F. Rodriguez, WebLogic on Kubernetes (15-17 August), Presented at Oracle Workshop Bristol, Bristol, 2017.
- S. A. Monsalve, Development of WebLogic 12c Management Tools (15 August), Presented at CERN openlab summer students' lightning talks, Geneva.
- A. Nappi, WebLogic on Kubernetes (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- A. Nappi, L. F. Rodriguez, Oracle Weblogic on Containers: Beyond the frontiers of your Data Centre Openday (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- A. Nappi, WebLogic on Kubernetes (20 October), Presented at Official Visit of Oracle PD Monica Ricelli in Geneva to discuss about WebLogic and Kubernetes, Geneva, 2017.



In a world where users always want to be connected and able to access their resources at any moment, reliability and 'zero downtime' are becoming first-class requirements.

# OPENSTACK CLOUDS

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES



**Project coordinator:**  
Tim Bell

**Technical team:**  
Theodoros Tsioutsias, Surya Seetharaman,  
Spyridon Trigazis, Belmiro Moreira

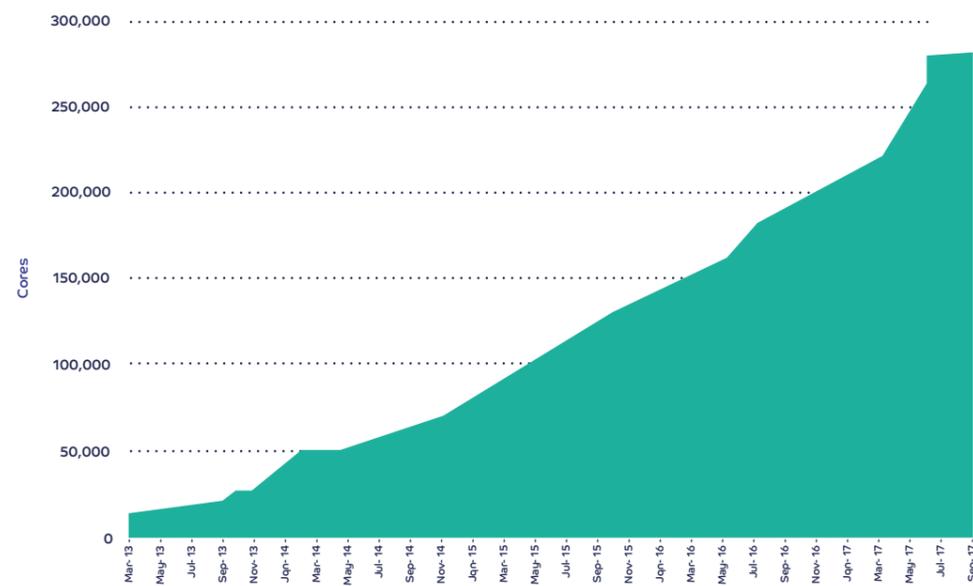
**Huawei liaison:**  
Stephan Keuneke

### Project goal

CERN's cloud infrastructure is among the world's largest private OpenStack installations. CERN and Huawei therefore have a common interest in driving and enhancing the development of OpenStack for large-scale applications. Work on this project is being carried out through the community processes of OpenStack, with the results being made fully open source.

### Background

It is important to ensure that the management system used for allocating CERN's computing resources is as efficient as possible and enables scientists to gain access to resources on a simple, self-service basis. By running common software based on OpenStack — both in on-premises and commercial clouds — we can meet the end users' needs for LHC Run 3 and beyond in an efficient and seamless manner. The enhancements we make will be submitted upstream and included in future releases of OpenStack, meaning that they



Total cores in the OpenStack environment at CERN.

will be available for use both in CERN's private, on-premises cloud and in products from Huawei based on OpenStack.

### Progress in 2017

The project was launched in March at the CeBit conference in Hannover, Germany, with development work then beginning during the third quarter of 2017. The two main focus areas have been OpenStack Cells v2 and pre-emptible instances.

Cells provides a way to build large clouds out of pools of resources; an initial, experimental implementation of this tool has been used at CERN since the in-house cloud went into production in 2013. The new version being developed by the community, Cells v2, will become the default for all OpenStack clouds. It is therefore important to carry out careful design work and rigorous testing, to ensure the code is production ready.

The use of pre-emptible instances in OpenStack makes it possible for users to create virtual machines that can be rapidly stopped if higher priority work arrives, thus enabling CERN's computing resources to be managed more efficiently.

### Next steps

With development activities now underway, the plan is to complete Cells v2 functionality and to deploy it in production at CERN during 2018. For pre-emptible instances, a demonstration is planned for the OpenStack summit in Vancouver, Canada, in May 2018. These enhancements would also be potentially available for inclusion in commercial cloud products, such as Huawei's Fusionsphere.

### Presentations

- T. Bell, OpenStack at scale (20 March), Presented at CeBit 2017, Hannover, 2017.
- V. E. Araujo Soto, Application Catalog for Openstack (15 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- J. van Eldik, Clouds at CERN (5 September), Presented at Huawei Connect, Shanghai, 2017.
- T. Bell, Huawei OpenStack Clouds (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- T. Bell, Clouds at CERN (26 September), Presented at OpenStack days UK, London, 2017.
- T. Bell, Clouds at CERN (26 October), Presented at Huawei Connect Europe, Berlin, 2017.



Enhancements made through this project will be submitted upstream for inclusion in future releases of OpenStack.

# OPENSTACK CONTAINERS

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES



### Project coordinator:

Tim Bell

### Technical team:

Ricardo Rocha, Spyridon Trigazis

### Rackspace liaisons:

Brian Stein, Jason Cannavale

### Project goal

We are working with Rackspace and the open-source community to enhance OpenStack's container support for use in high-energy physics. CERN has been using OpenStack since 2013 to provide virtual-machine resources. We are now adding containers to the standard production service; this enables rapid prototyping of applications and scaling of these services as they grow within the common accounting and resource pledges agreed with the WLCG.

### Background

Container technologies are an area of great interest, both across the ICT industry and within the physics community. Application developers are looking for more flexible ways of delivering components based on micro-services, such as SWAN (a platform created at CERN for performing interactive data analysis in the cloud) and Jupyter notebooks (an open-source web application that enables users to create and share documents containing live code, equations, visualisations, and narrative text).

### Progress in 2017

During 2017, OpenStack Magnum deployments grew significantly: they now represent 14% of all deployments. The CERN collaboration with Rackspace contributed over 60 updates to Magnum, including significant improvements related to cluster monitoring, online upgrades, and heterogeneous clusters of different node sizes.

One area of focus in 2017 was working on scaling large Kubernetes clouds; this was important for investigations into the future design of the ATLAS experiment's upgraded high-level trigger system for Run 4 of the LHC. This work required us to make improvements to both Magnum itself and the underlying orchestration component.

Given these significant contributions, the CERN openlab fellow working on the project has recently been elected to

lead the Magnum project team. This is an international role with responsibility for coordinating development plans, as well as overseeing the next release of Magnum.

### Next steps

The current phase of our collaboration with Rackspace comes to a close at the start of 2018. The container service is now in production in over 100 clusters, and the community will continue to develop functionality. Following the success of this project, further areas for collaboration with Rackspace are being explored.

### Presentations

- S. Trigazis, OpenStack @ CERN: Status update (4 February), Presented at CentOS dojo, Brussels, 2017.
- S. Trigazis, OpenStack Magnum Pike and the CERN cloud (28 February), Presented at Containers Meetup, Paris, 2017.
- S. Trigazis, From swarm to swarm-mode in the CERN container service (6 April), Presented at Docker Meetup, Zurich, 2017.
- R. Rocha, CERN Container experience (5 May), Presented at OpenStack Summit, Boston, 2017.
- S. Trigazis, OpenStack/Magnum and the CERN container service (28 June), Presented at German OpenStack Days 2017, Munich, 2017.
- S. Oak, Cloud Log Data Analysis (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- M. Sommer, Stateful services in containers (15 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- C. F. Souza, Kubernetes as a batch scheduler (and more...) (15 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- S. Trigazis, Rackspace Containers at Scale (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- T. Bell, Clouds at CERN (26 September), Presented at OpenStack days UK, London, 2017.



The CERN collaboration with Rackspace contributed over 60 updates to OpenStack Magnum.

# ORACLE CLOUD

## R&D TOPIC 1: DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES

# ORACLE

### Project coordinators:

*Eva Dafonte Perez,  
Katarzyna Maria Dziejniewicz-Wojcik*

### Technical team:

*Ana Lameiro Fernandez, Emil Pilecki,  
Artur Zygodlo, Ben Jones, Marteen Litmaath*

### Oracle liaisons:

*David Ebert, Paul Jenkins, Don Mowbray,  
Pauline Mahrer, Brian Spadolini, Cris Pedregal*

### Project goal

In recent years, Oracle has significantly expanded the range of services and products it offers related to cloud computing. Engineered systems like Oracle Exadata are now offered over the cloud. The aim of this project is to test these systems, including when using high-end CPUs and GPUs, in order to assess the feasibility of integrating them into the workflows of the LHC experiments. At the same time, we are also providing Oracle's engineering teams with feedback that can be used to further improve these products.

### Background

When the HL-LHC becomes operational in the mid-late 2020s, the required computing capacity — using current hardware, software, and analysis techniques — is expected to be 50 to 100 times higher than today. It is therefore important to use the most cost-effective computing solutions, which can often entail using resources that become temporarily available in an opportunistic manner. In terms of databases, cloud technology offers a potential disaster-recovery solution, makes it technically feasible to offload read-only workloads and backups to the cloud, and opens the possibility of replacing some private data-centre capacity if deemed advantageous.

### Progress in 2017

In 2017, we worked closely together with multiple LHC experiments and engineering teams from Oracle. Workloads from the LHC experiments were executed using Oracle Cloud Infrastructure resources. We set up infrastructure that makes it possible to run HTCondor jobs inside Docker containers and to integrate cloud resources into the management of workloads for the LHC experiments.

We also focused on functional tests of cloud-provisioned databases, both single-instance databases and Oracle Real Application Clusters. Additionally, we worked together with Oracle engineering teams to conduct performance

tests comparing engineered systems in the cloud and on-premises databases (including Oracle In-Memory). We gained important insights and shared valuable feedback with the Oracle engineers and product managers related to all areas of work within the project.

### Next steps

We plan to set up a test site for disaster recovery using Oracle Database Cloud services, thus helping us to improve our understanding of the requirements related to this (in terms of personnel, networking, data protection, etc.). We will also evaluate the use of Oracle Exadata to help improve the performance of the most complex workloads. In addition, we intend to repeat the experimental tests carried out this year using the new generation of computing hardware available via the cloud, including making use of HPC resources and machine-learning approaches based on GPUs.

### Presentations

- K. M. Dziejniewicz-Wojcik, Lightning talk: Oracle Cloud (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- K. M. Dziejniewicz-Wojcik, Oracle Cloud (29 November), Presented at Suisse Romande 'DB & Infra' Tech Community, Geneva, 2017.



*Oracle's Sebastian Solbach presents at the CERN openlab Open Day in September.*

# HIGH-THROUGHPUT COMPUTING COLLABORATION

## R&D TOPIC 2: COMPUTING PERFORMANCE AND SOFTWARE



### Project coordinators:

*Olof Barring, Niko Neufeld*

### Technical team:

*Alexandru Amihalachioaei, Luca Atzori, Omar Awile, Daniel Hugo Cámpora Perez, Paolo Durante, Christian Färber, Placido Fernandez, Andrea Luiselli (Intel), Jonathan Machen (Intel), Christina Quast, Rainer Schwemmer, Sébastien Valat, Balázs Voneki*

### Intel liaisons:

*Claudio Bellini, Marie-Christine Sawley*

### Project goal

The high-throughput computing collaboration (HTCC) evaluates upcoming Intel technologies for potential use in the ‘online’ computing infrastructure (the ‘trigger’ and data-acquisition systems) of the LHC experiments. The investigations, which are being carried out with the LHCb experiment, are split into four main areas: (1) assessing future generations of Intel® Xeon Phi processor chips; (2) comparing the merits of different types of FPGA-based accelerators, especially the new Intel Xeon+FPGA prototypes; (3) understanding the potential benefits of Intel® Quick Assist Technology (QAT) as a hardware accelerator for encryption or compression; and (4) testing high-speed network fabrics, such as Intel® Omni-Path.

### Background

As the luminosity of the LHC is ramped up, the rate of particle collisions will significantly increase. This, coupled with upgrades to the LHC experiments, will lead to far higher data rates coming from the particle detectors. At the LHCb experiment, for example, it is expected that the data throughput following upgrades made in 2021 will be 30 times higher than it is today. Thus, to ensure that the necessary ‘trigger’ and data-acquisition systems can be put in place, it is vital to test different types of computing architecture, accelerators, and network fabrics. It is also important to consider how algorithms used in these systems can be adapted to take advantage of modern hardware technologies.

### Progress in 2017

In area 1, we developed a proto-application that implements a new Kalman filter algorithm, which is an important component in LHCb’s data-analysis framework. By fully taking advantage of the vectorisation technologies available and ensuring both optimal data alignment and operation scheduling, we were able to achieve close to a seven-fold speedup on an Intel® Xeon-Phi 7210 (using a production dual-socket Haswell system as a baseline). Based on this

work and the lessons learnt, we created a new Kalman filter module for the LHCb experiment’s current software framework. During 2017, we also explored ways of distribute workloads across the cores of modern NUMA architectures.

In area 2, we ported a key particle-identification algorithm used by LHCb, called ‘RICH’, to a prototype hybrid system based on an Intel® Broadwell CPU and an Intel® Aria10 FPGA. Tests run on this system yielded positive results, both in terms of speed-up and energy efficiency. We also tested the algorithm used to decode raw data from LHCb’s calorimeter on this system, running over 100 times faster than on a single Intel® Xeon thread.

In area 3, we studied the use of Intel Quick Assist Technology for ‘on the fly’ data compression in LHCb’s trigger and data-acquisition system. An aggregate throughput of 120 Gb/s was achieved.

In area 4, we mostly focused on improving our understanding of the cabling and routing considerations related to the use of InfiniBand and Intel® Omni-Path. This work plays a key role in helping us to develop the most appropriate architecture for the LHCb experiment’s updated trigger and data-acquisition systems. We also began work to study aspects related to failure recovery.

### Next steps

Our work on the Kalman filter and RICH proto-applications has shown that the Intel® Knights Landing platform offers a compelling alternative to Intel® Xeon processors. This needs to be further benchmarked for other large cycle consumers within LHCb’s framework. The results on QAT compression are also exciting and warrant further investigation. In addition, we will further explore the use cases for the Intel Xeon+FPGA prototypes in terms of particle identification and raw-data decoding for the detector, as well as further exploring networking technologies.

### Publications

- A. Amihalachioaei, Calorimeter RAW data decoding using Intel Xeon + FPGA computing platform, LHCb Online note, In publication, under review, 2017.
- D. H. Cámpora Pérez, O. Awile and C. Potterat, A high-throughput Kalman filter for modern SIMD architectures, Proc. Euro-Par2017: Parallel Processing Workshops, (Lecture Notes in Computer Science, Springer, 2017), in print.
- D. H. Cámpora Pérez, O. Awile and O. Bouizi, Cross-architecture Kalman filter benchmarks on modern hardware platforms, Proc. ACAT, (Journal of Physics: Conference Series, IOP, 2017), submitted.
- D. H. Cámpora, Pérez and O. Awile, An Efficient Low-Rank Kalman Filter for Modern SIMD Architectures, Concurrency and Computation Practice and Experience, 2017, submitted.
- C. Faerber et al., Particle identification on a FPGA accelerated compute platform for the LHCb Upgrade, (2017), DOI: 10.1109/TNS.2017.2715900, <http://ieeexplore.ieee.org/document/7949150/>.
- C. Faerber, Acceleration of a Particle Identification Algorithm used for the LHCb Upgrade with the new Intel®Xeon®-FPGA, TIPP17, In publication, review finished, 2017.
- P. Fernandez, D. del Rio Astorga, M.F. Dolz, J. Fernandez O. Awile and J.D. Garcia, Parallelizing and optimizing LHCb-Kalman for Intel Xeon Phi KNL processors, Parallel, Distributed, and Network-Based Processing, 2018, submitted.
- M. Manzali et al., “Large-Scale DAQ Tests for the LHCb Upgrade,” in IEEE Transactions on Nuclear Science, vol. 64, no. 6, pp. 1486-1493, June 2017. doi: 10.1109/TNS.2017.2693422.
- C. Quast, A. Pohl, B. Cosenza, B. Juurlink and R. Schwemmer, Accelerating the RICH Particle Detector Algorithm on Intel Xeon Phi, Parallel, Distributed, and Network-Based Processing, 2018, submitted.
- B. Voneki, S. Valat, R. Schwemmer, N. Neufeld, J. Machen and D. H. C. Pérez, “Evaluation of 100 Gb/s LAN networks for the LHCb DAQ upgrade,” 2016 IEEE-NPSS Real Time Conference (RT), Padua, 2016, pp. 1-3. doi: 10.1109/RTC.2016.7543169.
- B. Voneki, S. Valat, R. Schwemmer, N. Neufeld, “RDMA optimizations on top of 100 Gbps Ethernet for the upgraded data acquisition system of LHCb”, Technology and Instrumentation In Particle Physics 2017.

Presentations are listed on the CERN openlab website.

# CODE MODERNISATION: GEANT V

## R&D TOPIC 2: COMPUTING PERFORMANCE AND SOFTWARE



**Project coordinator:**  
*Federico Carminati*

**Technical team:**  
*Andrei Gheata, Andrea Luiselli (Intel),  
Sofia Vallecorsa, Andrea Zanetti (Intel)*

**Intel liaisons:**  
*Claudio Bellini, Laurent Duhem*

### Project goal

Through the GeantV project, we're working to develop the next-generation simulation software used for describing the passage of particles through matter. We aim to recast classical particle-transport simulation in a form that enhances both code- and data-locality and its vectorisation potential, using instruction-level parallelism to improve performance. Another important goal is to integrate generic 'fast simulation' techniques based on machine-learning approaches.

### Background

The main driving force behind this work is the LHC experiments' vital need to increase the throughput of their simulated data samples for the HL-LHC era. Conservative projections suggest that simulation needs are likely to increase by a factor of ten compared to today. We expect that a speed-up of a factor of up to five can be achieved through code modernisation, with the additional speed-up being driven by 'fast simulation' approaches. Our research into new vectorisation techniques and the development of vectorised modules also has benefits in many other areas of computing for high-energy physics beyond simulation.

### Progress in 2017

During 2017, we prototyped a new scheduling approach that splits the stepping procedure for particle tracks into stages, accumulating several particles before actually performing the actions involved at each stage. This approach makes it possible to vectorise additional components of the framework, such as magnetic-field propagation and the physics models. This new version has a much smaller memory footprint, is topology aware, and can be configured to benefit from memory locality. In addition, we carried out work to vectorise the modules for magnetic field propagation, as well as those for geometry navigation.

We also began work to investigate novel machine-learning techniques for a generic approach to 'fast simulation'. We achieved promising preliminary results by implementing an approach based on generative adversarial networks to simulate the calorimeter model proposed for the linear collider.

### Next steps

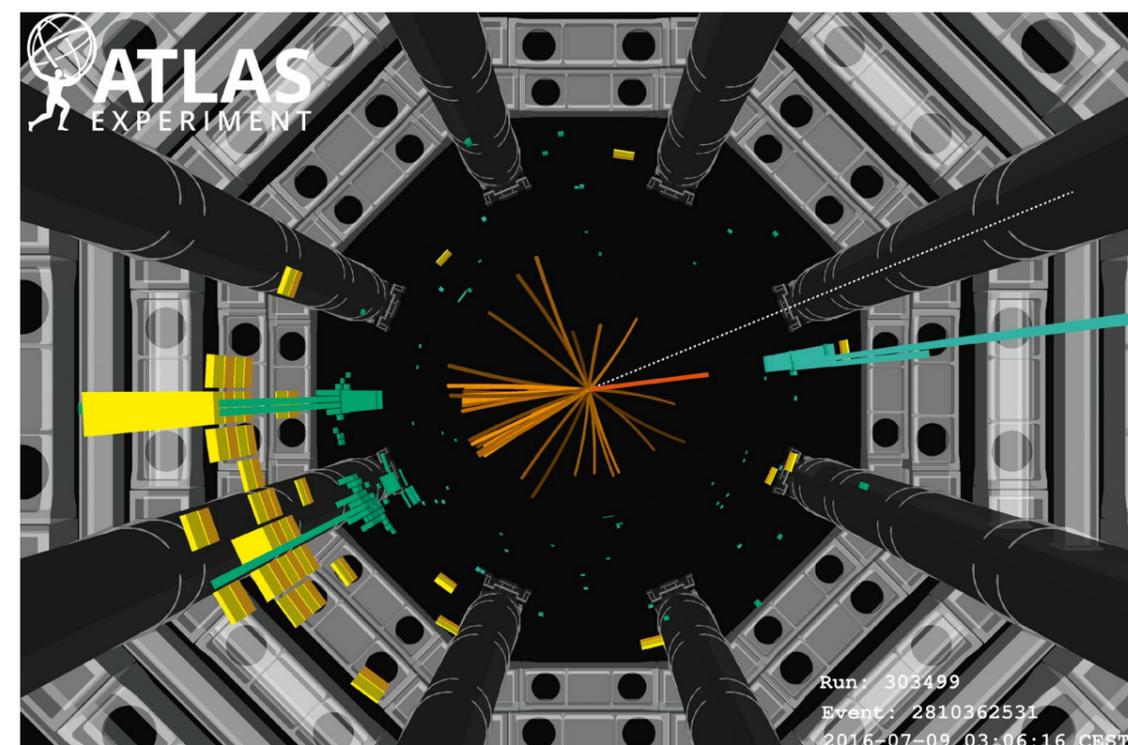
Development work on technologies to accelerate physics simulation workloads will continue; this work is being pursued as part of an international collaboration and is one of the main R&D activities of the 'software development for experiments' group at CERN's Experimental Physics department. The project is now targeting possible applications in the existing LHC experiment frameworks, as well as extension to use cases in other fields of scientific research, such as image analysis or biological simulation.

### Publications

- G. Amadio et al., GeantV alpha release, Proc. Advanced Computing and Analysis Techniques in Physics Research, Seattle, USA, 2017.

### Presentations

- F. Carminati, A. Gheata and S. Vallecorsa, The GeantV prototype on KNL 2017 Intel Xeon Phi User's Group (IXPUG) (11 March), Presented at Annual Spring Conference, Cambridge, 2017.
- S. Vallecorsa, Machine Learning for Fast Simulation 2017 (June 24), Presented at ISC High Performance, Frankfurt, 2017.
- E. Orlova, Deep learning for fast simulation: development for distributed computing systems (15 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- A. Gheata, GeantV (Intel Code Modernisation) (21 September), Presented at CERN openlab Open Day, Geneva, 2017.



A collision event in the ATLAS experiment.

# INDUSTRIAL CONTROL AND MONITORING

## R&D TOPIC 3: MACHINE LEARNING AND DATA ANALYTICS

## SIEMENS

### Project coordinator:

*Fernando Varela*

### Technical team:

*Filippo Tilaro, Jakub Guzik, Rafal Kulaga, Piotr Golonka, Piotr Seweryn, Brad Schofield, Peter Sollander, Fernando Varela, Lauri Sainio, Urshita Puri*

### Siemens/ETM liaisons:

*Thomas Hahn, Alexander Egorov, Mike Roshchin, Elisabeth Bakany, Ewald Sperrer*



*From the CERN Control Centre, experts can monitor and control the particle beams that circulate throughout CERN's accelerator complex.*

### Project goal

Our aim is to make control systems used for the LHC more efficient and 'smarter'. We are working to enhance the functionality of WinCC OA (a SCADA tool used widely at CERN) and to apply data-analytics techniques to the recorded monitoring data, in order to detect anomalies and systematic issues that may impact upon system operation and maintenance.

### Background

The HL-LHC project aims to increase the integrated luminosity — and hence the rate of particle collisions — by a factor of ten beyond the LHC's design value. Monitoring and control systems will therefore become increasingly complex, with unprecedented data throughputs. Consequently, it is vital to further improve the performance of these systems, and to make use of data-analytics algorithms to detect anomalies and to anticipate future behaviour. This project aims to achieve these objectives through three closely related areas of work:

- Developing a modular and future-proof archiving system (NextGen Archiver) that supports different SQL and NOSQL technologies to enable data analytics. It is important that this can be scaled up to meet our requirements beyond 2020.
- Developing a data-analytics platform that combines the benefits of cloud and edge computing.
- Developing a reporting system to feed the results of such data analysis directly into the operators' control consoles.

By applying data-analytics techniques in this manner, our goal is to improve operation and diagnostics, with preventive maintenance leading to greater efficiency and reduced costs.

### Progress in 2017

We developed a first prototype for the front end of the NextGen Archiver, and successfully carried out a series of functional and performance tests. We are also developing and testing back-end modules for Apache Kudu (a column-oriented data storage system) and Oracle databases.

A number of machine-learning algorithms were developed to detect faulty measurements in the cryogenics systems, as well as in other cooling and ventilation systems. We also implemented the control model's estimation for what is known as 'the electron-cloud effect' (a phenomenon that occurs in particle accelerators and reduces the quality of the particle beam) as a job distributed over multiple computing nodes. We used Apache Spark, a cluster-computing framework, to achieve this. We were thus able to reduce the computational time required by a factor of 100.

Two of Siemens's cloud-based solutions for big-data analytics, ELVis and WatchCAT, were integrated to create a platform combining cloud and edge computing (using 'Internet of things' devices). A domain-specific language (DSL) was implemented to develop user scripts that interact with the data.

Finally, a prototype reporting system based on Apache Impala (a massively parallel processing SQL query engine) was developed so that data-analytics results stored in the CERN Hadoop cluster could be injected into WinCC OA applications running on the operators' control consoles.

Even if still in an early stage, the deployment of some of these technologies at CERN has already enhanced control systems and has led to processes being optimised significantly. We have also been able to use these technologies to extend the operational life of some devices.

### Next steps

In the first quarter of 2018, we will carry out a large-scale testing campaign to help optimise the performance of the NextGen Archiver modules. These are due to be officially released toward the end of 2018.

Regarding the data-analytics framework, the tighter integration between ELVis and WatchCAT will have a number of benefits. Through a single user interface, it will be possible to define complex event-processing rules, configure the WatchCAT infrastructure (i.e. push the rules to the analytics processes running on the IoT devices), and to monitor the execution of the analyses.

A new version of the DSL will also be implemented to integrate event and signal processing, as well as to extend its semantics to allow for more powerful and flexible rules. On top of this, we will develop new algorithms to cover other CERN use-cases, thus extending our current portfolio.

### Presentations

- F. Tilaro, M. Gonzalez, B. Bradu, M. Roshchin, An expert knowledge based methodology for online detection of signal oscillations (26 June), Presented at International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA 2017), Annecy, 2017.
- U. Puri, Simplified frontend for data generation and testing purposes in WinCC OA NextGen Archiver project (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- L. M. Sainio, Web reporting framework for control data analysis (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- B. Schofield, F. Varela Rodriguez, F. M. Tilaro, J. Guzik, P. Golonka, P. J. Seweryn, Siemens Industrial Control and Monitoring (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- B. Bradu, E. Blanco, F. Tilaro, R. Marti, Automatic PID Performance Monitoring Applied to LHC Cryogenics (8 October), Presented at International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS 2017), Barcelona, 2017.
- P. Golonka, M. Gonzalez, J. Guzik, R. Kulaga, Future Archiver for CERN SCADA Systems (8 October), Presented at International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS 2017), Barcelona, 2017.
- P. J. Seweryn, M. Gonzalez-Berges, J. B. Schofield, F. M. Tilaro, Data Analytics Reporting Tool for CERN SCADA Systems (8 October), Presented at International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS 2017), Barcelona, 2017.
- F. Tilaro, B. Bradu, F. Varela, M. Roshchin, Model Learning Algorithms for Faulty Sensors Detection in CERN Control Systems (8 October), Presented at International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS 2017), Barcelona, 2017.

# INTEL BIG-DATA ANALYTICS

## R&D TOPIC 3: MACHINE LEARNING AND DATA ANALYTICS



### Project coordinators:

*Maria Girone, Luca Canali*

### Technical team:

*Evangelos Motesnitsalis, Ian Fisk,  
Matteo Cremonesi, Viktor Khristenko, Jim Pivarski,  
Bian Bianny (Intel), Rangarajan Radhika (Intel)*

### Liaisons:

*Claudio Bellini (Intel), Illia Cremer (Intel) Oliver Gutsche  
(Fermilab), Marco Manca (SCImPULSE), Mike Reiss (Intel)*



*Time is critical in analysing the ever-increasing volumes of LHC data.*

### Project goal

We are exploring scalable solutions that can be used to satisfy a variety of different needs related to the analysis of both physics data and data from industrial control systems. Big-data technologies like Spark show great potential in speeding up existing analysis procedures.

Through this project, we are working to optimise the analytics solutions at CERN in the following areas: data integration, data ingestion and transformation, performance, scalability, benchmarking, resource management, data visualisation, and hardware utilisation.

### Background

This project is split into four main areas of work:

- Accelerator controls: CERN runs a large number of industrial control systems based on SCADA tools, PLCs, etc. We are working on a proof-of-concept system to process the controls data using big-data platforms, such as Apache Spark.
- Physics data analysis: The LHC experiments continue to produce valuable physics data, offering numerous possibilities for new discoveries to be made. We are working on benchmarking ROOT, the CERN-created data-processing framework used for LHC physics data.
- Physics data reduction: Physics data reduction plays a vital role in ensuring researchers are able to gain valuable insights from the vast amounts of data produced by the LHC experiments. Our goal is to develop a new system — using industry-standard big-data tools — for filtering many petabytes of heterogeneous collision data to create manageable, but rich, datasets of a few terabytes for analysis.

Personalised medicine, epidemiology and diagnostics: We are also planning to work with the SCImPULSE foundation to explore how big-data-analysis technologies can be applied to medical data collected, thus informing efforts to improve practices related to safety and prevention.

### Progress in 2017

In 2017, we made significant progress in the area of physics data analysis. We first worked to make the format used for files in ROOT accessible via Spark SQL data frames, so as to avoid having to perform format conversions. We then worked to connect the Hadoop-related systems to the existing storage system of CERN, called EOS, thus making it possible to perform physics analysis without having to move large amounts of data into the Hadoop Distributed File System. We also managed to create the first fully functioning analysis examples using Apache Spark; these were tested using 1 TB of open data from the CMS experiment.

Progress was also made in the area of physics data reduction: Fermilab, the USA's premier particle physics and accelerator laboratory, joined CERN openlab in November. Researchers from the laboratory will collaborate with members of the CMS experiment and the CERN IT Department on efforts to improve technologies related to physics data reduction. More information on the scope of this project — as well as plans for proceeding — can be found in a news article published on the CERN openlab website.

Finally, in terms of the accelerator control systems, we began testing Apache Kudu as a candidate solution for data storage in 2017.

### Next steps

Work will be carried out in all four areas next year. In terms of the physics analysis work described above, our next step will be to investigate scaling options for larger inputs (we aim to scale up testing to a sample of 1 PB of open data during 2018). We also plan to investigate the possibility of using Spark over OpenStack, capitalising on the capabilities offered by Intel® CoFluent™ technology for cluster simulation.



### Publications

- O. Gutsche et al., CMS Analysis and Data Reduction with Apache Spark. Proceedings for 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (21 August), Seattle, 2017, (International Workshop on Advanced Computing and Analysis Techniques in Physics Research 'ACAT' 2017).

### Presentations

- O. Gutsche, Status of CMS Big Data Project (April 04), Presented at R&D meeting of CMS Spring Offline and Computing Week 2017, Geneva, 2017.
- O. Gutsche, Data Analytics in Physics Data Reduction (April 27), Presented at CERN openlab workshop on Machine Learning and Data Analytics, Geneva, 2017.
- M. Cremonesi, Infrastructure for Large Scale HEP data analysis (May 11), Presented at DS@HEP 2017 at Fermilab, Illinois, 2017.
- S. Sehrish, A path toward HEP data analysis using high performance computing (May 11), Presented at DS@HEP 2017 at Fermilab, Illinois, 2017.
- O. Gutsche, Status and Plans of the CMS Big Data Project (May 29), Presented at CERN Database Futures Workshop, Geneva, 2017.
- O. Gutsche, CMS Analysis and Data Reduction with Apache Spark (August 22), Presented at 18th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2016), Seattle, 2017.
- E. Motesnitsalis, Intel Big Data Analytics (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- E. Motesnitsalis et al., Physics Data Analytics and Data Reduction with Apache Spark (10 October), Presented at Extremely Large Database Conference 'XLDB' 2017, Clermont-Ferrand, 2017.
- V. Khristenko, HEP Data Processing with Apache Spark (December 6), Presented at CERN Hadoop User Forum, Geneva, 2017.

*Big-data technologies like Spark show great potential in speeding up existing analysis procedures.*

# ORACLE BIG-DATA ANALYTICS

## R&D TOPIC 3: MACHINE LEARNING AND DATA ANALYTICS

ORACLE

### Project coordinators:

*Eric Grancher, Eva Dafonte Perez*

### Technical team:

*Manuel Martin Marquez, Antonio Romero Marin*

### Oracle liaisons:

*Chris Lysnkey, Ryan Stark, Mark Hornick,  
Davide Basilio Bartolini, Cris Pedregal*

### Project goal

CERN's accelerator chain — culminating in the LHC — is extremely complex. We are working with Oracle to assess cutting-edge technologies for big-data analytics. Our goal is to use these technologies to help us gain new insights from the large amounts of monitoring data the accelerator chain's systems generate.

### Background

Both the accelerators at CERN and the experiments' cathedral-sized detectors produce huge amounts of systems-monitoring data. The effective exploitation of this data is vital in helping us to ensure the smooth and efficient running of this highly complex infrastructure. Technologies such as Hadoop, Apache Spark, and Apache Kafka, are already playing important roles. They are helping us to improve analytics performance on bigger datasets, do real-time analytics, do graph-based analysis, carry out 'data discovery' via advanced visualisation techniques, and more. Nevertheless, important challenges remain in terms of both the integration of new technologies and in terms of training users.

### Progress in 2017

During 2017, we contributed to work aimed at assessing the feasibility of the Future Circular Collider, a proposed successor to the LHC. For this, we defined a technology infrastructure using Oracle Big Data Discovery; this enables the analysis of terabytes of technical engineering data produced by approximately 50,000 sensors and monitoring devices in CERN's accelerator complex.

We also worked to develop a real-time analytics framework based on Apache Kafka. We then integrated this with Oracle Stream Explorer as an interface with which to interact with the data streams and process complex events.

Another area of work involved the definition and development of a proof-of-concept analytics system using Oracle Parallel Graph Analytics (PGX). We explored the integration of this with Apache Spark, as well as running it over the cloud, to help us extract important value out of the complex data generated by the accelerator complex at CERN.

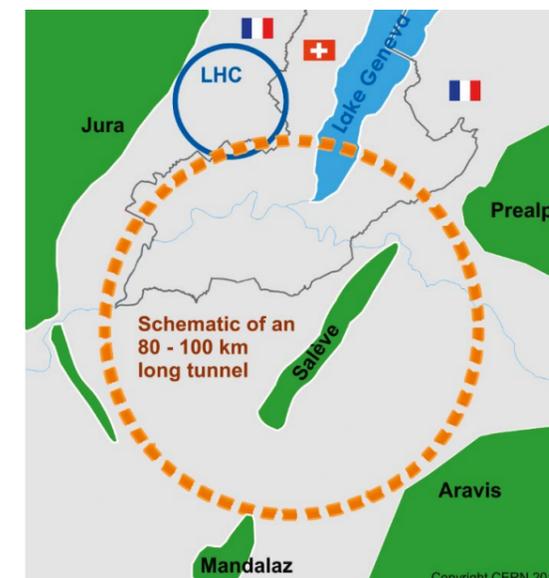
In addition, we carried out significant outreach work in 2017: we shared information about big-data analytics with a range of organisations, participated in a number of important conferences (including the Oracle Global Leaders Summit in Nicosia, Cyprus, and the Kafka Summit, in San Francisco, California, USA), and lectured as part of ESADE Business School's 'Big Data Analytics for Executives' programme in Madrid, Spain. Finally, we also defined new analytics approaches for the scholarly publication system used at CERN in 2017.

### Next steps

Going forwards, we are keen to enhance the analytics services described above by making it possible to run them over the cloud. We will also work to help expand the use of these services across CERN, while supporting the evolution of PGX and its integration with Apache Spark and various cloud technologies. A final area ripe for investigation is the evolution of real-time analytics frameworks coupled with data-source-agnostic interfaces, such as Oracle Big Data SQL.

### Presentations

- M. Martin Marquez, CERN's Journey into Big Data and Analytics: An Opportunity Full of Challenges (1 February), Presented at Oracle Modern Business Experience, London, 2017.
- M. Martin Marquez, Neil Sholay, Michael Connaughton, Big Data and Analytics Analyst Summit (1 February), Presented at Oracle Modern Business Experience, London, 2017.
- M. Martin Marquez, CERN IoT Systems and Predictive Maintenance (23 March), Presented at Oracle Road to Big Data, Madrid, 2017.
- M. Martin Marquez, G. Gabriel, F. Amalfi, How Your University Can Accelerate Enterprise Research at Scale: CERN's Experience with Oracle Big Data Platforms (24 May), webinar, 2017.
- A. Paragas, Zenodo Keyword Auto-Suggest using Parallel Graph Analytics (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- M. Martin Marquez, Oracle Analytics-as-a-Service (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- M. Martin Marquez, Accelerating Particles to Explore the Mysteries of the Universe and How Kafka can Help on that (28 August), Kafka Summit, San Francisco, 2017.
- M. Martin Marquez, Analytics customer reception Q&A panel (1 October), Presented at Oracle Open World, San Francisco, 2017.
- M. Martin Marquez, Franco Amalfi, Enterprise Research at Scale: CERN's Experience With Oracle Big Data Platform (3 October), Presented at Oracle Open World, San Francisco, 2017.



A map showing the location of the proposed Future Circular Collider.

# DATA POPULARITY AND ANOMALY DETECTION

## R&D TOPIC 3: MACHINE LEARNING AND DATA ANALYTICS



### Project coordinators:

Andrey Ustyuzhanin

### Technical team:

Stefan Roiser

### Yandex liaison:

Andrey Ustyuzhanin

### Project goal

The goal of the project is to improve LHCb operations by means of machine learning. There are two particular areas of focus: (1) data certification and (2) data management.

### Background

Ensuring data quality is essential for the LHCb experiment. Checks are done in several steps, both offline and online. Monitoring is based on continuous comparison of histograms with references, which have to be regularly updated by experts. Over the past two years, we have worked to create a novel, autonomous monitoring service for data collection. The service is capable of identifying deviations from normal operational mode to help personnel responsible for monitoring data quality to find the reason for such deviations.

The data collected by the LHCb experiment is stored across multiple datasets on both disks and tapes in the LHCb data storage grid. The storage systems used vary in terms of cost, energy consumption, and speed of use. We are also continuing our work on creating a 'data popularity estimator service' to analyse the usage history of each dataset, predict future usage patterns, and provide an optimal scheme for data placement and movement.

### Progress in 2017

We investigated how machine-learning approaches can be used to ensure automatic detection of anomalies in the data collected at the LHCb experiment. Our anomaly-detection algorithm was embedded into the LHCb experiment's web-based monitoring system for data quality, called Monet. This helped operators to identify when subsystems of the detector behave abnormally. Using sample data for training, the accuracy of our algorithm was shown to be quite high. However, as the scope of these samples was rather limited, there is still work for us to do to improve the generalisability of our algorithm.

We also carried out work to identify the least popular LHCb datasets in the storage system, so that they can be removed from fast, expensive storage media. In the first half of the year, we published a paper showing that our algorithm offers significant improvement over regular caching techniques. It is now being used in production at LHCb.

### Next steps

This work is close to completion. We are now keen to understand how the algorithms can be generalised and used at other experiments, potentially both at the LHC and beyond.

### Publications

- M. Hushchyn et al., GRID Storage Optimization in Transparent and User-Friendly Way for LHCb Datasets, *Journal of Physics: Conference Series*, Vol. 898 (2017) <https://doi.org/10.1088/1742-6596/898/6/062023>.

### Presentations

- A. Ustyuzhanin, Yandex Data Popularity and Anomaly Detection at LHCb (21 September), Presented at CERN openlab Open Day, Geneva, 2017.



Ensuring data quality is essential for the LHCb experiment.

# BIODYNAMO

## R&D TOPIC 4: APPLICATIONS IN OTHER DISCIPLINES



innopolis  
UNIVERSITY



### Project coordinators:

Fons Rademakers

### Technical team:

Lukas Breitwieser, Ahmad Hesam

### Collaborator liaisons:

Roman Bauer (Newcastle University), Claudio Bellini (Intel),  
Marcus Kaiser (Newcastle University), Manuel Mazzara  
(Innopolis University), Klaus-Dieter Oertel (Intel),  
Max Talanov (Kazan Federal University)

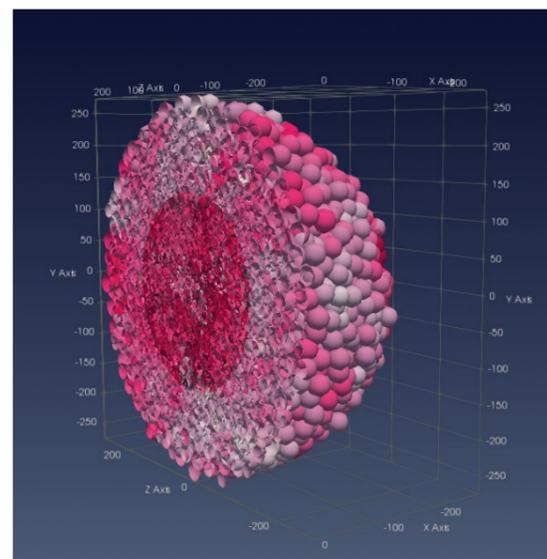
*Tumour concept, simulated in less  
than 10 seconds on a laptop (starting from a single cell).*

### Project goal

We are aiming to create a platform through which life scientists can easily create, run, and visualise three-dimensional biological simulations. Built on top of the latest computing technologies, the BioDynaMo platform will enable users to perform simulations of previously unachievable scale and complexity, making it possible to tackle challenging scientific research questions.

### Background

Within the life-sciences community, computer simulation is being used more and more to model increasingly complex biological systems. Although many specialised software tools exist, establishing a high-performance, general-purpose platform would be a major step forward. CERN is therefore contributing its deep knowledge in large-scale computing to this collaboration with Newcastle University in the UK, Innopolis University in Russia, and Kazan Federal University in Russia, and Intel. Together, we are working to develop a unique platform.



### Progress in 2017

Throughout 2017, we worked on improving the code base to harness fully the capabilities of modern hardware. With its improved modular architecture, BioDynaMo is now ready for use in a range of different biological fields. The integration of ParaView — a popular scientific visualisation engine — into the platform was an important milestone. This enables life scientists to visualise and interact with their simulation, even as it is running.

Backup-and-restore functionality is essential, particularly for simulations that take longer to run. We therefore integrated ROOT — a data-processing framework created at CERN and used widely in the high-energy physics community — into the BioDynaMo platform. At CERN, ROOT is used to store the large amount of data produced by the LHC experiments; it also enables fast post-processing and analysis of data.

It is also very important that the platform is user friendly: we spent significant effort improving the installation and usage experience. It is now possible to set up BioDynaMo and run a simulation in just three simple steps. During 2017, we also laid the foundations for running BioDynaMo in the cloud. Using distributed computing will make it possible to harness many thousands of computer processors to simulate very large biological structures.

### Next steps

Being able to run simulations on multiple compute nodes in parallel is essential for complex modelling. Consequently, our next milestone is to complete the distributed runtime for BioDynaMo. We will also investigate ways to mitigate the I/O bottleneck in low-bandwidth, high-latency environments and to take advantage of heterogeneous computing.

### Publications

- L. Breitwieser, BioDynaMo: A New Platform for Large-Scale Biological Simulation (Master's thesis), Graz University of Technology, Austria, 2016.

### Presentations

- K. Kanellis, Scaling a biological simulation platform to the cloud (15 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- L. Breitwieser, BioDynaMo (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- L. Breitwieser & A. Hesam, BioDynaMo: Biological simulation in the cloud (1 December), Presented at CERN IT technical forum, Geneva, 2017.



*The goal of this project is to  
create a platform through which  
life scientists can easily create,  
run, and visualise three-dimen-  
sional biological simulations.*

# GENEROOT

## R&D TOPIC 4: APPLICATIONS IN OTHER DISCIPLINES



### Project coordinators:

Fons Rademakers

### Technical team:

Taghi Aliyev

### Liaison:

Mario Falchi (King's College London),

Alessia Visconti (King's College London),

Marco Manca (SCimPULSE)

### Project goal

Through this project, we are developing a multi-science data-analysis platform; this is being built up around CERN-developed technologies like Zenodo (a research data repository), REANA (a system for instantiating data-analysis in the cloud), and CVMFS (a scalable, reliable and low-maintenance software distribution service). When finished, our platform will support a complete data-analysis life cycle, from data discovery through to access, processing, and end-user data analysis. The first use case for this platform relates to genomics, with ROOT — a data-processing framework created at CERN and used widely in the high-energy physics community — being used to store and process genomics data sequences.

### Background

In many research communities today, reproducibility, communication, and data pipelines are implemented in suboptimal ways. Through the GeneROOT project, we are therefore working to create a powerful system to capture and facilitate the habits of researchers. Our platform will allow for negotiation and sharing of common 'values' among scientists within a given field and will help us to understand the reasoning behind why certain choices are made. Rather than providing a simple toolkit for researchers, we are creating a rich system through which researchers can challenge the 'value chains' within their own respective fields and potentially enhance their approach to performing research.

### Progress in 2017

The GeneROOT project has now been running since mid-2016. Following the establishment of the project team, work gathered pace in 2017. To date, we have focused on the design and implementation of initial prototypes, mock ups, and the development of a minimal viable platform. Initial use cases for the platform were generated in collaboration with the biomedical community, with significant effort dedicated

to understanding the requirements of the community for such a system.

King's College London in the UK and SIDRA Medicine in Qatar are working with us to carry out initial tests with the platform. We are also exploring opportunities for collaboration with other research institutes, including Maastricht University Hospital and Maastricht University in the Netherlands, the European Bioinformatics Institute (EMBL-EBI) and the University of Cambridge in the UK. Initial prototype systems and minimal viable platforms are set to be shared with the communities for testing and feedback in early 2018.

### Next steps

In the coming year, our primary focus will shift to the implementation and architectural design of the first prototypes. The understanding we have gained from our engagement with the research communities — as well as the technical requirements compiled — will feed into the development of these prototypes. Based on the feedback we receive, we will then work to improve the platform further and to deploy the first product to the wider research community.

### Presentations

- J. J. Gonzalez Ortiz, ROOT based Genomics project (11 August), Presented at CERN openlab summer students' lightning talks, Geneva, 2017.
- T. Aliyev, GeneROOT (21 September), Presented at CERN openlab Open Day, Geneva, 2017.
- T. Aliyev, Big Data in Healthcare (29 September), Presented at Innovation in Healthcare: Futuro Summit, Florence, 2017.
- T. Aliyev, F. Rademakers, A Multi-Science Data Analysis Platform and the GeneROOT Use Case (8 December), Presented at CERN IT Technical Forum, Geneva, 2017.



Taghi Aliyev and his project colleagues are developing a multi-science data-analysis platform based on a range of CERN-developed technologies.

# 06

## KNOWLEDGE

# Education, training, and outreach

CERN openlab is designed to create and share knowledge through a wide range of activities and programmes.

CERN openlab is a knowledge factory. We work to disseminate this knowledge through both outreach activities and educational programmes. As well as promoting our technical work among a variety of stakeholders, we are working to train the next generation of ICT specialists. Thus, CERN openlab provides a means for its collaboration members to share a joint vision of the future of scientific computing. This vision is communicated to a wide audience, including partner clients, policy makers, members of the press, and the general public. Together, we can shape the future of scientific computing for the benefit of both research and wider society.

### Visits and workshops

Top delegations from governments and industry frequently tour CERN: over 130 protocol visits were organised in 2017. The CERN openlab concept and projects are systematically presented to the guests visiting the CERN IT department. CERN openlab partners have the opportunity to organise customer and press visits, too. Visiting groups are briefed about CERN openlab in a dedicated VIP meeting room, known as the CERN openlab 'openspace'. Around 190 press articles were published about our work over the course of the year, with over 300,000 people visiting the CERN openlab website.

Over 50 workshops, lectures, visits and other events related to CERN openlab's work were held throughout the year. Further information about these can be found on our website, including many recordings. On our website, you can also find a complete collection of press coverage, member case studies, press releases, videos, technical documents, presentations, and reports.

### New promotional tools

In 2017, CERN openlab launched its new logo and design package. These were created to modernise the CERN openlab brand (reflecting the wide array of research endeavours we now support), with the colours of the logo being used to help organise the various complimentary strands of CERN openlab's activities. In 2017, this new design package was used to create a number of new promotional materials, including an updated membership brochure, press pack, flyers, posters and much more. All materials can be accessed from the CERN openlab website; print copies can be obtained by contacting the CERN openlab communications office.

### Future ICT challenges

2017 also saw the launch of CERN openlab's new white paper on future ICT challenges in scientific research.



*In 2017, 37 students – representing 22 different nationalities – took part in the CERN openlab summer-student programme.*



*In 2017, the students went on a two-day trip to Zurich, where they visited Google, ETH Zurich, and Open Systems.*

The workshops held in early 2017 to collate input for this paper — as well as the document itself — were promoted across a broad range of channels.

The white paper was officially launched at the 2017 CERN openlab 'open day' in September. In addition to setting out the key challenges we plan to tackle in CERN openlab's upcoming sixth three-year phase, the event provided an excellent opportunity to review the progress made across the around 20 R&D projects carried out in CERN openlab's fifth phase.

#### Summer-student success

During our fifth phase, over 100 students came to CERN as part of the CERN openlab summer-student programme. In 2017, 37 students — representing 22 different nationalities — took part in the programme. They each spent nine weeks at CERN, working on advanced computing projects with applications in high-energy physics and beyond.

As part of the programme, the students attended a series of lectures given by IT experts on advanced CERN-related topics. Topics covered in 2017's programme included machine learning, computing security, grids and clouds, data acquisition, and code modernisation.

#### Educational excursions

Within the CERN openlab summer-student programme, the students also visit CERN facilities and experiments, as well

as other organisations. In 2017, the students went on a two-day trip to Zurich, where they visited Google, ETH Zurich, and Open Systems. As with the previous six years, this trip was organised by the team at Open Systems, who have been awarded the status of 'CERN openlab associate member' in recognition of their fantastic support.

In addition, two of the summer students also took a two-week trip to China, where they continued their project work in Beijing and Shenzhen, collaborating with Tsinghua University. These students, whose projects focused on citizen science and crowd computing, were both working with the SDG (Sustainable Development Goals) Summer School at the University of Geneva.

#### Webfest creativity

Another highlight of the summer was the CERN Summer Student Webfest. The event is a hackathon, through which bright and creative minds meet over a weekend to build exciting science projects using open web technologies. 2017's Webfest, which was supported by CERN openlab for the second year running, featured over 70 participants collaborating on 14 projects over a weekend.

A diverse range of ideas were developed over the weekend, including augmented-reality apps, educational games, expert chat bots, puzzles, and more. Tony Al Najjar, a summer student working with the CMS experiment, was selected as the winner of the competition by a panel of

judges. He created an interactive and educational LED-based game called CERcle, which puts players in control of the LHC. Al Najjar won a trip to London to participate in the Mozilla Festival.

#### Projects presented

The CERN openlab summer students had lots of fun — and learnt much — during the trips, the webfest, and the lecture series. However, the main focus of their time at CERN was undoubtedly their projects. These covered a diverse range of topics, including high-performance computing, big data, visualisation, machine learning, and much more.

The projects enabled the students to gain hands-on experience with some of the latest ICT solutions, working under the supervision of leading experts in the field.

On 11 and 15 August, the students presented their work in two dedicated public 'lightning talk' sessions. In 5-minute presentations, each student explained the technical challenges they have faced and described the results of what they have been working on for the nine weeks they have spent at CERN.

The best presentations from each of the two sessions were selected by a panel of judges. The winners from the first session were as follows:

- 1<sup>st</sup>: **Sharad Agarwal**, IoT Security
- 2<sup>nd</sup>: **Agrima Seth**, Anomaly Detection using Machine Learning for Data Quality Monitoring in the CMS Experiment
- 3<sup>rd</sup>: **Alastair Cuthbert Paragas**, Zenodo Keyword Auto-Suggest using Parallel Graph Analytics

The winners from the second session were as follows:

- 1<sup>st</sup>: **Markus Sommer**, Stateful Services in Containers
- 2<sup>nd</sup>: **Yasmine Nasri**, Building Effective Database Backup and Recovery Monitoring Using Elastic Stack
- 3<sup>rd</sup>: **Clenimar Filemon Souza**, Kubernetes as a batch scheduler

Markus Sommer, a student from Philipps University in Marburg, Germany, was selected as this year's overall winner.

#### Student challenge

Five of this year's summer students were also selected to take part in the Modern Code Developer Challenge with Intel. This competition saw the students' blog about their projects — all of which either make use of Intel projects or are connected to broader collaborative initiatives between Intel and CERN openlab — on a dedicated website: <http://cern.ch/go/xJ9M>. This website also features audio interviews and videos with the students discussing their projects.

At the start of the challenge, the plan was for a panel of judges to select just one of the five students as the winner and to invite this person to present their work at the 2017 International Conference for High Performance Computing, Networking, Storage and Analysis (SC17). However, owing to the high quality of the students' work, the judges decided to select two winners, both of whom received full funding from Intel to travel to the USA and present their work.



*As with the previous six years, the students' trip to Zurich was organised by the team at Open Systems, who have been awarded the status of 'CERN openlab associate member' in recognition of their fantastic support.*

### Cells in the cloud

Konstantinos Kanellis, a final-year undergraduate in the department of electrical and computer engineering at the University of Thessaly, Greece, was selected as one of the two winners of the Modern Code Developer Challenge with Intel. His work is related to the BioDynaMo project.

BioDynaMo is one of CERN openlab's knowledge-sharing projects (another part of CERN openlab's collaboration with Intel on code modernisation). The project's goal is to develop methods for ensuring that scientific software makes full use of the computing potential offered by today's cutting-edge hardware technologies. This is a joint effort by CERN, Newcastle University, Innopolis University, and Kazan Federal University to design and build a scalable and flexible platform for rapid simulation of biological tissue development.

The project focuses initially on the area of brain tissue simulation, drawing inspiration from existing, but low-performance, software frameworks. By using the code to simulate development of both normal and diseased brains, neuroscientists hope to learn more about the causes of — and identify potential treatments for — disorders such as epilepsy and schizophrenia.

Late 2015 and early 2016 saw algorithms already written in Java code ported to C++. Once porting was completed, work began to optimise the code for modern computer processors and co-processors. However, to address ambitious research questions, more computational power was needed. Ongoing work is therefore focused on adapting the code for high-performance computing resources over the cloud.

Kanellis's work focused on adding network support for the single-node simulator and prototyping the computation management across many nodes. More information about this project can be found on pages 48 and 49 of this report.

### Smash-simulation software

Elena Orlova, a third-year student in applied mathematics from the Higher School of Economics in Moscow, Russia, was selected as the other winner of the Modern Code Developer Challenge with Intel. Her work focused on teaching algorithms to be faster at simulating particle-collision events.

Physicists widely use a software toolkit called GEANT4 to simulate what will happen when a particular kind of particle hits a particular kind of material in a particle detector. This toolkit is so popular that researchers use it in other fields to predict how particles will interact with other matter, such as in assessing radiation hazards in space, in commercial air travel, in medical imaging, and in optimising scanning systems for cargo security.

An international team, led by researchers at CERN, is developing updates to this toolkit; their goal is to improve physics accuracy and boost performance on modern computing architectures. This work is supported by a CERN openlab project with Intel on code modernisation, called Geant V.

The project team is implementing a deep-learning tool intended to make simulations faster. Orlova worked to write a flexible mini-application to help train the deep neural network on distributed computing systems. More information about this project can be found on pages 38 and 39 of this report.



The CERN Summer Student Webfest was supported by CERN openlab for the second year running. It featured over 70 participants collaborating on 14 projects over a weekend.

### ICT stars of the future

The other three projects featured in the challenge run with Intel focused on using machine-learning techniques to better identify the particles produced by collision events in the LHC, integrating IoT devices into the control systems for the LHC, and helping computers get better at recognising objects in satellite maps created by UNOSAT, a UN agency hosted at CERN.

Alberto Di Meglio, head of CERN openlab, and Michelle Chuaprasert, director of the developer relations division at Intel, presented the challenge at the SC17 conference. In total, the social media campaign around this project reached over 1.7 million people across Facebook and Twitter.

### Alumni interactions

As well as working to attract and train new talent through its various programmes for students, graduates, and others, CERN is also keen to maintain a strong connection with those who have left the Organization. In 2017, CERN launched a new alumni network. It targets those who have

left CERN, but also welcomes employed and associated members of personnel. It has been designed to make the global CERN community more inclusive, and has been set up to provide those who have left with a means of keeping in touch with CERN and with each other. It also fosters ambassadorship for the mission and values of CERN, and will help colleagues, in particular younger ones, with their future career development.

At the heart of this network is an online alumni platform designed for interaction: <http://alumni.cern/>. Much effort has gone into providing substance in the form of career opportunities, learning opportunities, and other benefits for our members.

We have created a CERN openlab group on the alumni platform and strongly encourage those of you reading this report who have a connection with us to join the discussion there.



During our fifth phase, over 100 students came to CERN as part of the CERN openlab summer-student programme.

# 07 FUTURE

## Next steps

We are working with leaders from both industry and a range of scientific fields to identify — and tackle — the ICT challenges of the future.

CERN openlab's fifth three-year phase came to a close at the end of 2017. For most of this time, the LHC ran like a Swiss clock, delivering large amounts of collision data to the experiments, beyond even researchers' best expectations. With higher luminosity, leading to more collisions and more data, physicists are now able to explore the most fundamental interactions between particles with unprecedented sensitivity and precision.

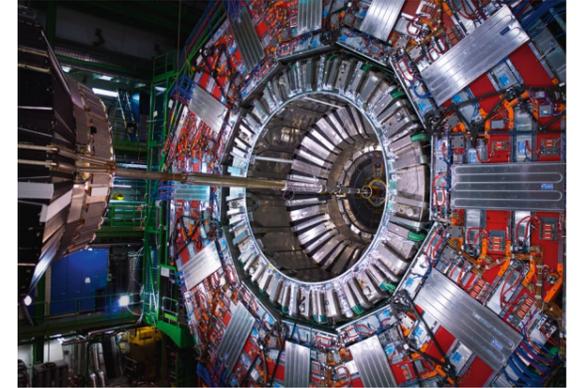
Five years after the announcement of the discovery of the Higgs boson, physicists are now able to look at this very special particle with a magnifying glass, gaining deeper insight into the way it interacts with other particles.

The Higgs boson was the last missing particle in the puzzle that forms the 'Standard Model' of physics. Developed in the early 1970s, the Standard Model has successfully explained almost all experimental results and precisely predicted a wide variety of phenomena. Over time and through many experiments, it has become established as a well-tested physics theory.

### The LHC upgrade programme

However, the Standard Model only explains about 5% of the universe. Planned upgrades to the LHC will help physicists to explore phenomena beyond this, such as dark matter and dark energy.

The LHC has been designed to follow a carefully set out programme of upgrades. The LHC typically produces particle collisions for a period of around three years (known as a



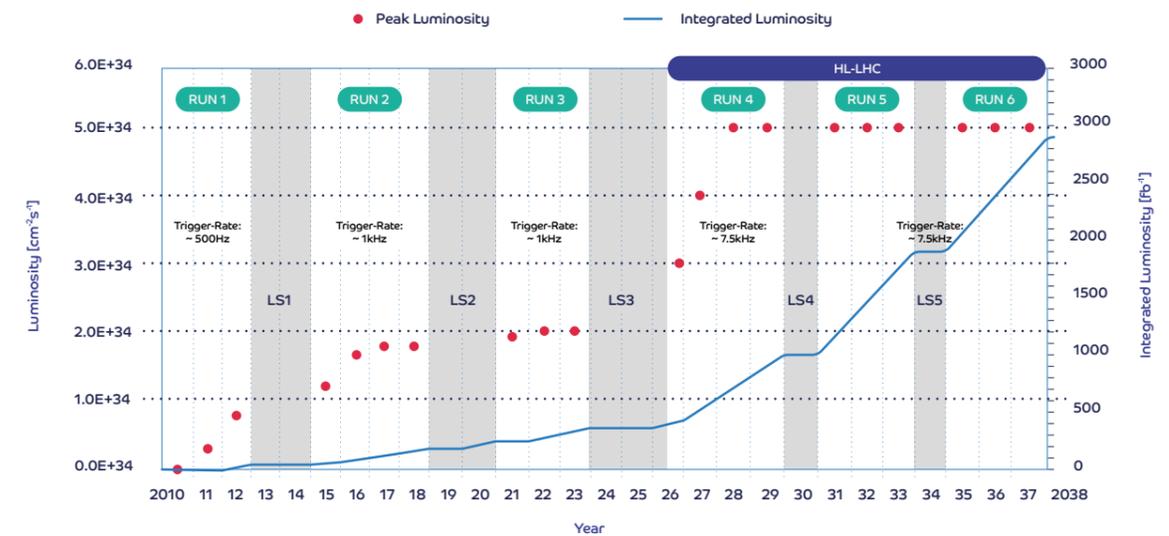
The two large multi-purpose detectors, ATLAS and CMS, will be significantly upgraded for the HL-LHC.

'run'), followed by a period of about two years for upgrade and maintenance work (known as a 'long shutdown', which is commonly abbreviated to LS).

Over the next decade, two significant upgrade periods are foreseen for the LHC, known as LS2 and LS3. Improvements to both the accelerator complex and the detectors during each of these periods will result in a dramatic increase in the computing needs of the laboratory, so as to ensure that the full physics potential is exploited.

### Run 3

By Run 3, which is planned to start in 2021, the ALICE and LHCb experiments will markedly increase their 'trigger' rates, meaning that much less data will be filtered out. Collecting many more collision events — and therefore much more data — than ever before means that the events will



The planned upgrade programme for the LHC.

## Collisions and 'trigger' systems

Collisions in the LHC generate particles that often decay in complex ways into even more particles. Electronic circuits record the passage of each particle through a detector as a series of electronic signals, and send the data to the CERN data centre for digital reconstruction. The digitised summary is recorded as a 'collision event'.

Up to about 1 billion particle collisions can take place every second inside the LHC experiment's detectors. It is not possible to read out all of these events. A 'trigger' system is therefore used to filter the data and select those events that are potentially interesting for further analysis. The LHC experiments' trigger systems usually consist of multiple levels, and are based on both hardware and software. Future plans include modifying trigger systems to make them more flexible, which may pose significant technical challenges.

have to be reconstructed (turned from raw sensor data to representations of the particle tracks crossing the detectors) shortly after being collected, rather than first being stored and then processed 'offline'. Within the budgets and technologies available today, no other complete reprocessing will be possible.

Faster trigger rates at the experiments mean it will be necessary to increase the size of the server farms used for so-called 'online processing', in order to perform real-time calibration and near-real-time reconstruction of collision events. It is also important to investigate specialised hardware and networking technologies that enable compute-intensive tasks to be carried out faster and with higher energy efficiency.

Specific projects have been carried out during CERN openlab's fifth phase to help the ALICE and LHCb experiments address the ICT challenges posed by the upgrade work they will carry out.

### Run 4

The High-Luminosity Large Hadron Collider (HL-LHC) project aims to crank up the performance of the LHC in order to increase the potential for discoveries starting in run 4. The objective is to increase integrated luminosity by a factor of ten beyond the LHC's design value.

Luminosity is an important indicator of the performance of an accelerator: it is proportional to the number of collisions that occur in a given amount of time. The higher the luminosity,

the more data the experiments can gather. This enables physicists to observe rare processes and study new particles with greater precision.

The two large multi-purpose detectors, ATLAS and CMS, will be significantly upgraded for the HL-LHC. The time required for the reconstruction of collision events will increase exponentially as the number of collisions increases; both experiments expect to collect and reconstruct five to ten times more events than during the current run (run 2).

The increased complexity of the collision events — collected at a higher trigger rate — will drive a huge increase in computing needs for run 4. Using current software, hardware, and analysis techniques, it is estimated that the computing capacity required would be around 50-100 times higher than today. The data storage needs are expected to be in the order of exabytes by this time.

It is reasonable to expect that technology improvements over the next seven to ten years will yield an improvement of approximately a factor ten in both the amount of processing and storage available at the same cost. This may address the issue of the need for storage capacity, but leaves a very large gap in terms of processing. There is no expectation that budgets will increase in the manner required to close this 'resource gap' by simply increasing the total ICT resources available. For this reason, and to ensure maximum efficiency of resource use, it is vital to explore new technologies and methodologies together with industry leaders. We need to learn to do more, with greater accuracy, and in ever-smarter

ways. In particular, as the need for yet further precision leads to increased computing capacity requirements, we need to find ways of rapidly scaling ICT resources — both up and down — to match.

### White paper on future ICT challenges

With CERN openlab's sixth three-year phase starting in 2018, we have been busy working to identify the key ICT challenges that will face the LHC research community over the coming years. Throughout the first half of 2017, we carried out an in-depth consultation process involving workshops and discussions with representatives of the LHC experiments, the CERN IT department, our collaborators from industry, and other 'big science' projects. Based on our findings, we published 'the CERN openlab white paper on future ICT challenges in scientific research' in September.

The white paper identifies 16 ICT 'challenge areas' that are ripe for tackling together with industry. These have been grouped into overarching 'R&D topics' (also used to organise the projects in the results section of this report).

The first of these focuses on data-centre technologies. It addresses issues such as ensuring that data-centre architectures are flexible and cost effective; that cloud computing resources can be used in a scalable, hybrid manner; that new technologies for solving storage capacity issues are thoroughly investigated; and that long-term data-storage systems are reliable and economically viable.

The second major R&D topic highlighted in the white paper relates to the modernisation of code, so as to ensure that the maximum performance can be achieved on the new hardware platforms available.

The third R&D topic focuses on machine learning and the importance of successfully translating its huge potential into concrete solutions that will play a role in monitoring the accelerator chain, optimising the use of ICT resources, and even hunting for new physics.

In the white paper, we also identified ICT challenges that are common across research disciplines. With ever more



We published 'the CERN openlab white paper on future ICT challenges in scientific research' in September.

research fields adopting methodologies driven by 'big data', it is vital that we collaborate with research communities, such as astrophysics, biomedicine, and Earth sciences. As well as sharing tools and learning from one another's experience, working together to address common ICT challenges can increase our ability to ensure that leading ICT companies are producing solutions that meet our common needs.

Our white paper identifies ICT challenges that must be tackled over the coming years in order to ensure that physicists across the globe can exploit CERN's world-leading experimental infrastructure to its maximum potential. However, the white paper also serves to demonstrate the emergence of a whole new set of technology paradigms, from pervasive ultra-fast networks of smart sensors in the 'Internet of things', to machine learning and 'the optimisation of everything'. These technologies have the power to revolutionise the way big science is done, particularly in terms of data analysis and the control of complex systems. They have enormous potential, not just for research, but also for wider society. The knowledge and expertise at CERN can play a key role in ensuring this potential is realised, with these technologies being put to use for the benefit of both science and society. CERN openlab — with its unique collaboration with several of the world's leading IT companies — is ideally positioned to help make this a reality.

The white paper sets the stage for our joint R&D activities — together with our collaboration members from both industry and research — over the coming years. It is available on the CERN openlab website: <http://openlab.cern/whitepaper>.

**ALBERTO DI MEGLIO**  
**Head of CERN openlab**  
[Alberto.Di.Meglio@cern.ch](mailto:Alberto.Di.Meglio@cern.ch)

**FONS RADEMAKERS**  
**Chief research officer**  
[Fons.Rademakers@cern.ch](mailto:Fons.Rademakers@cern.ch)

**ANDREW PURCELL**  
**Communications officer**  
[Andrew.Purcell@cern.ch](mailto:Andrew.Purcell@cern.ch)

**MARIA GIRONE**  
**Chief technology officer**  
[Maria.Girone@cern.ch](mailto:Maria.Girone@cern.ch)

**FEDERICO CARMINATI**  
**Chief innovation officer**  
[Federico.Carminati@cern.ch](mailto:Federico.Carminati@cern.ch)

**KRISTINA GUNNE**  
**Administrative / finance officer**  
[Kristina.Ulrika.Gunne@cern.ch](mailto:Kristina.Ulrika.Gunne@cern.ch)



## CERN openlab 2017 Annual Report

**Editor:** Andrew Purcell

**Design and production:** CERN Design and Visual Identity Service

**All images provided by the CERN Audio-Visual Production Service, except:**

CERN Design and Visual Identity Service and CERN Audio-Visual Production Service (page 11);  
CERN Design and Visual Identity Service (page 21, 30, 59); Thomas Mc Cauley (page 24);  
Veronika McQuade (page 29); ATLAS Collaboration (page 39); CERN openlab communications office  
(page 14, 35, 42, 53, 54, 55, 57, 61); Mona Schweizer (page 8); Panagiotis Charitos (page 45);  
BioDynaMo project (page 48); CERN Audio-Visual Production Service, Michael Hoch, and Joseph Gobin (page 59)

**An electronic version of this report is available on the CERN openlab website:**

[www.cern.ch/openlab/resources/annual-reports](http://www.cern.ch/openlab/resources/annual-reports)

ISBN 978-92-9083-481-6

Published by CERN

© CERN 2018

