

PAPER • OPEN ACCESS

CERN openlab: Engaging industry for innovation in the LHC Run 3-4 R&D programme

To cite this article: M Girone *et al* 2017 *J. Phys.: Conf. Ser.* **898** 072049

View the [article online](#) for updates and enhancements.

Related content

- [News](#)
- [The breaking point of modern processor and platform technology](#)
Sverre Jarpe, Alfio Lazzaro, Julien Leduc et al.
- [Defence Research: Parallel lines to meet](#)
Philip Campbell

CERN openlab: Engaging industry for innovation in the LHC Run 3-4 R&D programme

M Girone¹, A Purcell, A Di Meglio, F Rademakers, K Gunne, M Pachou, S Pavlou
CERN, 1211 Geneva 23, Switzerland.

maria.girone@cern.ch

Abstract. LHC Run3 and Run4 represent an unprecedented challenge for HEP computing in terms of both data volume and complexity. New approaches are needed for how data is collected and filtered, processed, moved, stored and analysed if these challenges are to be met with a realistic budget. To develop innovative techniques we are fostering relationships with industry leaders. CERN openlab is a unique resource for public-private partnership between CERN and leading Information Communication and Technology (ICT) companies. Its mission is to accelerate the development of cutting-edge solutions to be used by the worldwide HEP community. In 2015, CERN openlab started its phase V with a strong focus on tackling the upcoming LHC challenges. Several R&D programs are ongoing in the areas of data acquisition, networks and connectivity, data storage architectures, computing provisioning, computing platforms and code optimisation and data analytics. This paper gives an overview of the various innovative technologies that are currently being explored by CERN openlab V and discusses the long-term strategies that are pursued by the LHC communities with the help of industry in closing the technological gap in processing and storage needs expected in Run3 and Run4.

1. Introduction

CERN openlab is a unique public-private partnership that accelerates the development of cutting-edge 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 complex 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.

The CERN openlab team consists of three complementary groups of people: young engineers hired by CERN and funded by the partners, technical experts from partner companies involved in the projects, and CERN management and technical experts working partly or fully on the joint activities. Each project team is supervised by a project coordinator, who liaises with the collaborating company. This distributed team structure permits close collaboration with computing experts in the LHC experiments, as well as with engineers and scientists from CERN openlab collaborators, who contribute significantly to our activities. Valuable contributions are also made by the students participating in the CERN openlab Summer Student Programme.

¹ Author to whom whom any correspondence should be addressed.



The work described in this paper was carried out in conjunction with CERN openlab's partners (Intel, Oracle, Siemens); contributors (Brocade, Cisco, IDT, Rackspace, Seagate); associates (Comtrade Software, Huawei, Yandex); and research members (EMBL-EBI, GSI, INFN, Innopolis University, Kazan Federal University, King's College London, Newcastle University) [1]. A complete overview of CERN openlab's activities can be found on the CERN openlab website [2].

2. Results

CERN openlab's current three-year phase (2015-2017) is tackling 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 [3]. This section details the main project milestones from 2016.

2.1 Data Acquisition (online)

Existing and emerging large-scale research projects are producing increasingly high amounts of data at ever-faster rates. A prime example of this comes from CERN's LHC, which produces millions of particle collisions every second in each of its detectors, thus generating approximately 1 PB of data per second.

2.1.1 Intel High-Throughput Computing Collaboration. In 2016, the High-Throughput Computing Collaboration (HTCC) continued its studies of three new Intel data-centre technologies: Intel Omni-Path, Intel® Xeon™/FPGA, and Intel® Xeon Phi™.

Intel Omni-Path is used for data-exchange at 100 gigabits per second (Gb/s) between servers for data-acquisition systems. We performed a first large-scale test at the Marconi supercomputer hosted by the Italian inter-university computing consortium 'Cineca'.

Intel Xeon/FPGA integration was used to decompress and reformat packed binary data from the detectors and to accelerate key kernels of the event-filtering algorithm base. To demonstrate the utility of FPGA acceleration, a floating-point heavy, complex algorithm was used to perform photon reconstruction in a large detector. To enable better comparison, the algorithm was implemented in both Verilog hardware-description language and in Intel OpenCL for FPGAs. In both cases, very competitive acceleration results (up to a factor of 35) over the reference were achieved. The HTCC was also able to demonstrate that FPGAs are extremely power-efficient compared to other programmable technologies (up to a factor of four). In addition, the tight Xeon integration of Intel's new product effectively removes the 'PCIe-bottle-neck', thus overcoming one of the main challenges in using FPGAs as accelerators: the need to 'feed' them with sufficient amounts of data.

In 2016, Intel released the next generation of its Xeon Phi processors, which has the code name 'Knights Landing'. Some of our most computationally expensive algorithms have been tested on these processors, with significant speedups being observed in certain cases. Work is now being carried out to further optimise our algorithms for running on this platform. Many of the improvements we have obtained during this process are also applicable to running on other Intel Xeon processors — this is particularly true in the case of very-wide vector registers (AVX and AVX2).

2.1.2 IDT RapidIO for Data Acquisition. 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, started in June 2015.

In terms of hardware installations, our machine pool saw extensions and improvements: a 32-port top-of-rack switch was installed in the CERN Data Centre, as well as 16 server nodes equipped with RapidIO-PCIe network interface controllers. In terms of software installations, the IDT team made regular updates to our Linux modules and libraries. Two of these updates were especially relevant: (i) the introduction of reserved memory, which enables applications to allocate any amount of memory for RapidIO operations; (ii) the introduction of the riosocket drivers, which expose a standard TCP/IP

interface running on top of RapidIO, thus enabling us to use any TCP/IP compatible application (such as iperf or Hadoop) without any porting work at all.

The work of the project this year has been focused on the use cases of data analytics and data acquisition. Within the data analytics use case, the porting of ROOT to RapidIO has undergone big changes. We have gone from using a simple transfer protocol at the beginning of the year to a more efficient implementation that makes use of circular buffers and RapidIO doorbells, meaning quick notifications can be sent between nodes. With the introduction of riosockets and the performance improvement they provide, the Hadoop installation was also able to move from the test cluster to the server cluster in the CERN Data Centre. Significant strides were also made for the data acquisition use case. 2016 saw the main work start for the LHCb experiment's benchmark DAQPIPE, which is an interconnect-agnostic application that emulates the behaviour in the LHCb data-acquisition network. We ported DAQPIPE to RapidIO and are conducting benchmark runs to evaluate observations related to the RapidIO technology's core.

2.2 Networks and Connectivity

Today, the ever-increasing external data traffic (more than 100 Gb/s) is putting pressure on the CERN firewalls. CERN and other large organisations require robust, scalable network solutions that provide high bandwidth for data transfers while maintaining appropriate levels of security. Networking also plays a vital role in data acquisition from the experiments, and is paramount to the success of the WLCG.

2.2.1 Brocade Flow Optimizer. The Brocade Flow Optimizer (BFO) project aims to enhance and generalise Brocade's BFO application, and to use it to build an intelligent steering system for network traffic. The system will be capable of supporting use cases that are not easily handled using traditional networking approaches. These use cases include intrusion detection system (IDS) mirroring, a firewall load-balancer, and an advanced policy-based routing engine. BFO is a software-defined networking (SDN) application designed to improve visibility of network traffic and to enable flow-steering by using the OpenFlow protocol to 'program' the fast, specialised application-specific integrated circuits (ASICs), which serve as the hardware-forwarding engines in network devices. BFO can be deployed to control a set of OpenFlow-enabled devices and thus enhance network flexibility and programmability. BFO features a user interface and a REST API, which can be used for consuming its services in an automated manner.

In 2016, significant progress was made on the IDS mirroring use case. The designed system receives traffic intercepted at the CERN firewall system and load-balances it across a pool of IDS servers, each running the open-source Bro Network Security Monitor system. When a security threat is detected, the malicious traffic is temporarily mirrored to a dedicated set of servers for storage and later analysis. The use of BFO is key to achieving the desired intelligence and automation in the system. The prototype setup features a Brocade MLXe router, controlled by a BFO instance that has been extended with a specially developed Bro plugin. This plugin enables the IDS software to call the BFO REST API directly and thus achieve the desired automation. The prototype has now been deployed in the CERN Data Centre for testing.

2.3 Data Storage Architectures

Every year, the four large-scale LHC experiments create tens of petabytes of data, which need to be reliably stored for analysis in the CERN Data Centre and many partner sites in the WLCG [4]. As the user demands are increasing in terms of both data volume and aggregated speed of data access, CERN and its partner institutes are continuously investigating new technological solutions to provide their user communities with more scalable and efficient storage solutions.

2.3.1 Oracle Database Technology and Monitoring. Throughout 2016, we continued to investigate Oracle Database In-Memory, with the key focus placed on performance. We evaluated the query response time under different database configurations for business intelligence applications. In

parallel, we tested the new ‘Active Data Guard’ features for Oracle Database In-Memory. In addition, we upgraded our Oracle Enterprise Manager monitoring systems — both those used for testing purposes and those used in production — to the latest version, 13c. This work included configuration and integration of CERN’s active directory authentication system in the upgraded version of Oracle Enterprise Manager, and required configuration of custom certificates.

2.3.2 Seagate Alternative Storage Architecture: ‘Kinetic’. The collaboration with Seagate made considerable progress in 2016. Building on the software foundation laid in 2015 with the Kinetic I/O module (open-source library available on GitHub), the project was able to show first usage in a prototype service. Using the existing 1-PB Kinetic cluster (with first generation ‘Lombard’ drives), a storage pool was created for a related CERN openlab project with the European Bioinformatics Institute. This project focuses on the application of the ROOT analysis system to genomic data, with around 400 TB of genomic data having been served to a group of researchers throughout 2016.

A web interface was built using AngularJS, enabling efficient configuration and management of the Kinetic cluster and the monitoring of individual drive parameters. A REST API has also been added to all administrator and user commands provided by the EOS management server, so as to enable web front ends. The EOS scheduling software has been adapted to allow multi-path access to the Kinetic cluster via dynamically selected proxy nodes, thus providing high-availability. In addition, the I/O interface of the EOS storage server has been re-factored to better facilitate the requirements of having both locally attached and remote disks. In parallel to the ongoing software developments, work was also carried out to prepare for the installation of a new cluster with second-generation Kinetic drives, which will complement the existing installation. This work also included pre-production testing and firmware optimisation — in particular for non-sequential access patterns, which are typical of some analysis uses cases at CERN.

2.3.3 Comtrade Software EOS Productisation. The scope of this project — undertaken in collaboration with CERN openlab associate member Comtrade Software — is the evolution of the 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. The main target of the project’s initial phase was to provide a robust installation kit to enable rapid installation of EOS both for evaluation purposes and for fast deployment in production. This has now been accomplished. The kit includes the necessary installation instructions and tools for operations (admin guide) and for users (user guide), as well as a first version of the EOS whitepaper.

In addition to its use at CERN, EOS is now in production at both Australia’s Academic and Research Network (AARNet) and the European Commission’s Joint Research Centre in Ispra, Italy. It is also being evaluated for use at several other institutes. The next phase of the project is focused on several items. The first goal is the integration of Comtrade Software engineers into the development and operations team at CERN to gain experience and autonomy in operating, maintaining, and developing EOS. The second goal is to continue the evolution of the installation framework and the documentation, in order to provide all EOS functions at installation time, including ‘Sync&Share’ capabilities, erasure-coding, and support for geographically distributed multi-site instances. An additional milestone for this next phase is to develop a simultaneous testing framework to run after the build of every release. This will be used to certify each EOS version.

2.4 Compute Management and Provisioning

CERN, as infrastructure and service provider for the high-energy physics community, has been very actively involved in grid and cloud computing since the early days. As the use of virtualisation has become an increasingly viable solution for instantiating computing nodes, the concept of ‘the cloud’ or cloud computing has gradually established itself as an efficient and cost-effective solution for scientific computing.

2.4.1 Oracle Java EE. 2015 saw significant work carried out to consolidate our Java EE platform-as-a-service, known as ‘the Middleware on Demand’ (MWOD). All sites were migrated from our previous platform, ‘the J2EE Public Service’. During 2016, we focused on developing new functions for the MWOD, such as support for WebSockets and fine-grained authorisation for application paths.

Given that the developer community is demanding ever more agile and dynamic environments for developing and testing their applications, we also started to assess technologies such as Docker, HAProxy, and Kubernetes in 2016. We created a proof-of-concept platform based on these technologies for provisioning Apache Tomcat application servers.

We ran our applications on Oracle WebLogic with more than 250 clusters. We studied a number of container technologies in 2016, as part of our drive to continuously improve our platform. We created custom WebLogic-Docker images and integrated our CERN management and monitoring tools.

2.4.2 Oracle Database Cloud. In order to find the best ways of overcoming the computing challenges the high-energy physics community is set to face in the coming years, many projects and experiments at CERN are looking at various cloud solutions from a range of vendors. Oracle’s cloud portfolio has grown dramatically in recent years; the company now offers a wide range of products in this area. In late 2016, the Database Services Group in the CERN IT Department started some preliminary tests with Oracle Cloud, focusing primarily on aspects related to security (single sign-on integration) and backup management. In December, we started an evaluation of Oracle Cloud as a disaster-recovery solution, focusing on the feasibility of its use in the CERN environment, as well as the effort needed in case of future migration.

2.4.3. Rackspace Containers at Scale. Since its first releases, OpenStack has supported virtual machines to dynamically provision self-service resources for end users. New application frameworks such as Jupyter and Kubernetes rely on containers, a lightweight abstraction with lower overhead than virtual machines. OpenStack’s support for containers is an area of major interest for industry and research.

We have been working with Rackspace and the open-source communities to enhance OpenStack’s container support for use in high-energy physics. 72 patches have been submitted, supporting new storage technologies, documentation, local container repositories, and the usability of the command-line tools. These patches have been made available to the open-source community. The service is now in production for the physicists and engineers at CERN. There are over 40 CERN projects underway, supporting new approaches to analysis and application architectures. An example of such a project is Jupyter SWAN, which is a web application that enables users to create and share documents containing live code, equations, visualisations, and explanatory text. With the dynamic provisioning of containers from OpenStack (and integration of ROOT), these notebooks can make it possible for researchers to interactively collaborate, publish, and preserve their results. With container usage growing rapidly, it is important to be sure that the service can scale to meet the needs of the experiments at CERN. A 1000-node scale test was able to handle over 7 million requests/second.

2.5 Computing Platforms (Offline)

The success of existing and future scientific and experimental programmes depends — among other factors — on efficient exploitation of the recent and future advances in computing technology. Existing software needs to be revised, optimised, or completely redesigned to fully exploit the performance gains provided by newer multi-core platforms, fast co-processors, and graphical processors.

As well as the projects described below, work with Cisco to enhance the performance of distributed applications through the elimination of kernel processing in the data path came to a close in mid-2016.

2.5.1 Intel Code Modernisation. Across research fields, code optimisation is of paramount importance in ensuring that available hardware is used as efficiently as possible. The increased computing requirements of the LHC Run 2 mean it is more important than ever to optimise high-energy physics codes for new computing architectures. As part of this project, Intel experts have delivered a number of

workshops at CERN addressing the latest Intel software tools and providing training on code-vectorisation technologies. In addition to the work on the Geant software simulation toolkit described below, this project involves work to optimise code to be used at the GSI Helmholtzzentrum für Schwerionenforschung, as well as code used in other research domains, through the BioDynaMo and GeneROOT projects [2].

2.5.1.1 Geant software simulation toolkit. The GeantV work aims to develop the next-generation simulation software used for describing the passage of particles through matter. It started as a research-and-development effort aiming to develop an alternative path to detector simulation software, using a multi-event approach and multi-particle ‘basket’ parallel processing to achieve higher efficiency with simulation-specific computations on modern hardware. This was achieved by expressing algorithms in a type-agnostic manner using either scalar or vector interfaces. In turn, this was enabled by the implementation of a new library called ‘VecCore’, using dedicated backends to support different SIMD architectures. It demonstrated excellent results on both Xeon and Xeon Phi (Knights Corner). A new AVX-512-aware backend called ‘UME::SIMD’ enabled GeantV to be deployed ‘out of the box’ and to demonstrate SIMD gains on the Intel Knights Landing architecture.

One of the key focus areas of the GeantV development in 2016 was the adaptation of the core architecture to offer improved scalability on many-core platforms. We introduced and tested a multi-propagator model where the work of transporting particle baskets in the detector is split among several managers, each taking charge of a limited number of threads and having weak communication with one another for workload balancing. This approach has shown good scalability on the Intel Knights Landing architecture, while preserving vector gains. It will be the basis of the third version of the GeantV core architecture, which is scheduled for an alpha release around the fourth quarter of 2017.

A number of high-performance components developed for GeantV have already been made available to the HEP community. The aim is to integrate GeantV with experiment frameworks by the end of 2018.

2.5.2 ARM Porting and Benchmark Studies. The LHC experiments and the CERN computing and data infrastructure make use of a large number of software frameworks for simulation and reconstruction of collision events. It is important to continuously monitor advances and trends in technology and evaluate software on different computing platforms as they evolve. A new project was therefore launched midway through 2016 to port several widely used codes for running on ARMv8-A 64-bit architecture. As part of the project, studies are being carried out to test and measure performance, energy consumption, and other operational aspects, so as to understand the strengths and weaknesses of the architecture.

A cluster comprised of ARMv8-A 64-bit evaluation prototype servers was delivered in June. Each server contains 32 Cortex-A57 cores with 128 GB RAM connected in four fast memory channels. The project started with the successful porting of two key software packages used at CERN, Geant4 and ROOT. These programs are each of the order of around three million lines of C++ code and are both fundamental to the research carried out at CERN. Following this porting work, the project’s efforts focused on use cases for the two large multi-purpose experiments on the LHC: ATLAS and CMS. Members of the ATLAS collaboration first worked to port their software framework and validate the output from the ARMv8-A 64-bit servers against current technology used. Once the output was found to be in agreement, the team then worked to benchmark the ARMv8-A 64-bit servers. In terms of energy efficiency, the cluster was found to deliver a maximum of approximately 1150 event simulations per kilowatt hour. Members of the CMS collaboration carried out similar benchmarking work using tools pertinent to their experiment. They found that the new ARMv8-A 64-bit cluster demonstrates up to roughly five times increased performance over comparable clusters consisting of previous generations of ARM platforms, which they have been experimenting with since 2013.

2.6 Data Analytics

During the past decades, CERN and other international research laboratories have been gathering not only enormous amounts of scientific data, but also very large quantities of systems-monitoring data from their instruments. Curating, enriching, and managing this data enables its exploitation. The main challenges in data analytics for scientific and engineering applications involve technology, integration, and education.

2.6.1 Oracle Analytics-as-a-Service. During the first half of 2016, the team carried out a detailed evaluation of Oracle Big Data Discovery. This involved analysis of terabytes of technical engineering data produced by approximately 50,000 sensors and other monitoring devices in CERN's accelerator complex. In the second half of the year, we focused our efforts in two main areas: (i) the integration of Oracle Big Data Discovery into our production environment; (ii) establishing an architecture for reliability and availability analysis of the systems within the proposed Future Circular Collider (FCC). This infrastructure has enabled the collaboration working on the FCC studies to perform a variety of analyses for accelerator conditions and modes, such as cooling down, warming up, and injecting beams of energy in a scalable manner.

We also implemented a 'proof of concept' using Oracle R Advanced Analytics for Hadoop (ORAAH) for the reliability and availability studies carried out for the FCC. In addition, the team involved in the specification, deployment, and validation phases for a scalable streaming analytics platform based on Apache Kafka.

2.6.2 Siemens Industrial Control and Monitoring. Most of the CERN installations in the experiments, accelerators, and other technical infrastructures rely on a multitude of heterogeneous industrial control systems for proper functioning. These control systems produce enormous amounts of data related to both the systems they control and their own internal state. Together with our partners Siemens, we are seeking to apply big-data analytics techniques to this data, in order to improve the operational behaviour of the entire system. Our goal is to increase efficiency and develop new control models that improve reliability. In particular, specific algorithms have been designed and implemented for use-cases related to each of the following areas:

- Detection of faulty sensor measurements for cryogenics systems
- Performance measurement of process control systems using proportional–integral–derivative (PID) controllers
- Detection of alarm flooding and identification of responsible control devices
- Development of a recommendation system for users of WinCC OA (a SCADA tool widely used at CERN).

2.6.3 Yandex Data Popularity at LHCb. 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. The goal of this project is to design, develop, and deploy a 'data popularity estimator service' that would analyse the usage history of each dataset, predict future usage patterns, and provide an optimal scheme for data placement and movement. During 2015, the project team developed an algorithm that is able to reduce the total amount of disk space needed by 40%. In 2016, work focused on determining the optimal number of replicas for each file. This investigation was divided into two main parts.

Firstly, a classifier based on decision trees was used to predict the likelihood of datasets being used in the coming six months. For this purpose, the access histories of the datasets — and their metadata for the last two and a half years — were fed into the system. Each dataset is described in terms of the following features: how recently it was last accessed, time elapsed between consecutive accesses, time of first access, creation time, access frequency, type, and size. Comparison with the 'Least Recently Used' (LRU) algorithm shows that our approach significantly decreases the number of files that are wrongly removed. Secondly, Brown's model for exponential smoothing was used to predict the number of times each dataset would be accessed over the next month. This was then used to prioritise the datasets

in terms of the number of replicas they each need. Thus, the service is able to make appropriate decisions about where replica datasets can and cannot be removed in order to maximise available disk space — while keeping the risk of data loss to a minimum.

2.6.4 Yandex Anomaly Detection in LHCb Online Data Processing. 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. In 2016, work continued 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. It will, therefore, increase the efficiency of the system by reducing the amount of spoiled data that is erroneously stored for further analysis.

An automated data-certification assistant, called ‘Roboshifter’ was developed in 2016. A data pipeline was set up for feeding LHCb run history into the Roboshifter machine-learning algorithm, which predicts whether given data is good or bad. Roboshifter was also integrated into the web service for data-quality monitoring that is used by personnel to help them find the detector components responsible for any problems with the data.

3. Conclusion

CERN openlab is continuing its work to support CERN’s research community, with a particular focus on the upgrades to the LHC and the detectors that will be carried out during LS2 and LS3. With the data rates from the experiments set to increase significantly, efforts have been focused on supporting the work to overhaul and modernise their data-acquisition systems, while also ensuring that the maximum benefits are gained from the hardware available to CERN’s teams by making sure the software running on it has been fully optimised. Efforts have now begun to identify the ICT challenges that will be tackled in CERN openlab’s sixth phase (2018 to 2020), through which we seek to align our work with the needs of the WLCG [4] and HSF [5].

4. Acknowledgements

The CERN openlab management team would like to acknowledge the work of the project teams who are responsible for the results detailed in this paper. We also extend our thanks to each of the companies and research institutes collaborating in CERN openlab for their continued commitment to this unique public-private partnership.

5. References

- [1] CERN openlab. *Collaboration Members*. Available from: http://openlab.cern/about/collaboration_members [Accessed 20 February 2017].
- [2] CERN openlab. *Technical Areas*. Available from: <http://openlab.cern/technical-areas> [Accessed 20 February 2017].
- [3] CERN openlab. *Annual Report 2015*. Geneva: CERN; 2016. Available from: http://openlab.cern/resources/annual-reports_areas [Accessed 20 February 2017].
- [4] WLCG. *About*. Available from <http://wlcg-public.web.cern.ch/about> [Accessed 20 February 2017].
- [5] HSF. *Goals of the HEP Software Foundation*. Available from: <http://hepsoftwarefoundation.org/organization/goals.html> [Accessed 20 February 2017].