DISCOVERY MACHINES

CERN operates a unique complex of eight accelerators and one decelerator. The accelerators speed particles up to almost the speed of light before colliding them with each other or launching them at targets. Detectors record what happens during these collisions, and all the data collected is stored and analysed using a worldwide computing grid. Hundreds of physicists, engineers and technicians contribute to the operation and maintenance of these sophisticated machines.

The LHC performed remarkably in 2016, delivering far more collisions than expected. (CERN-GE-1101021-08)

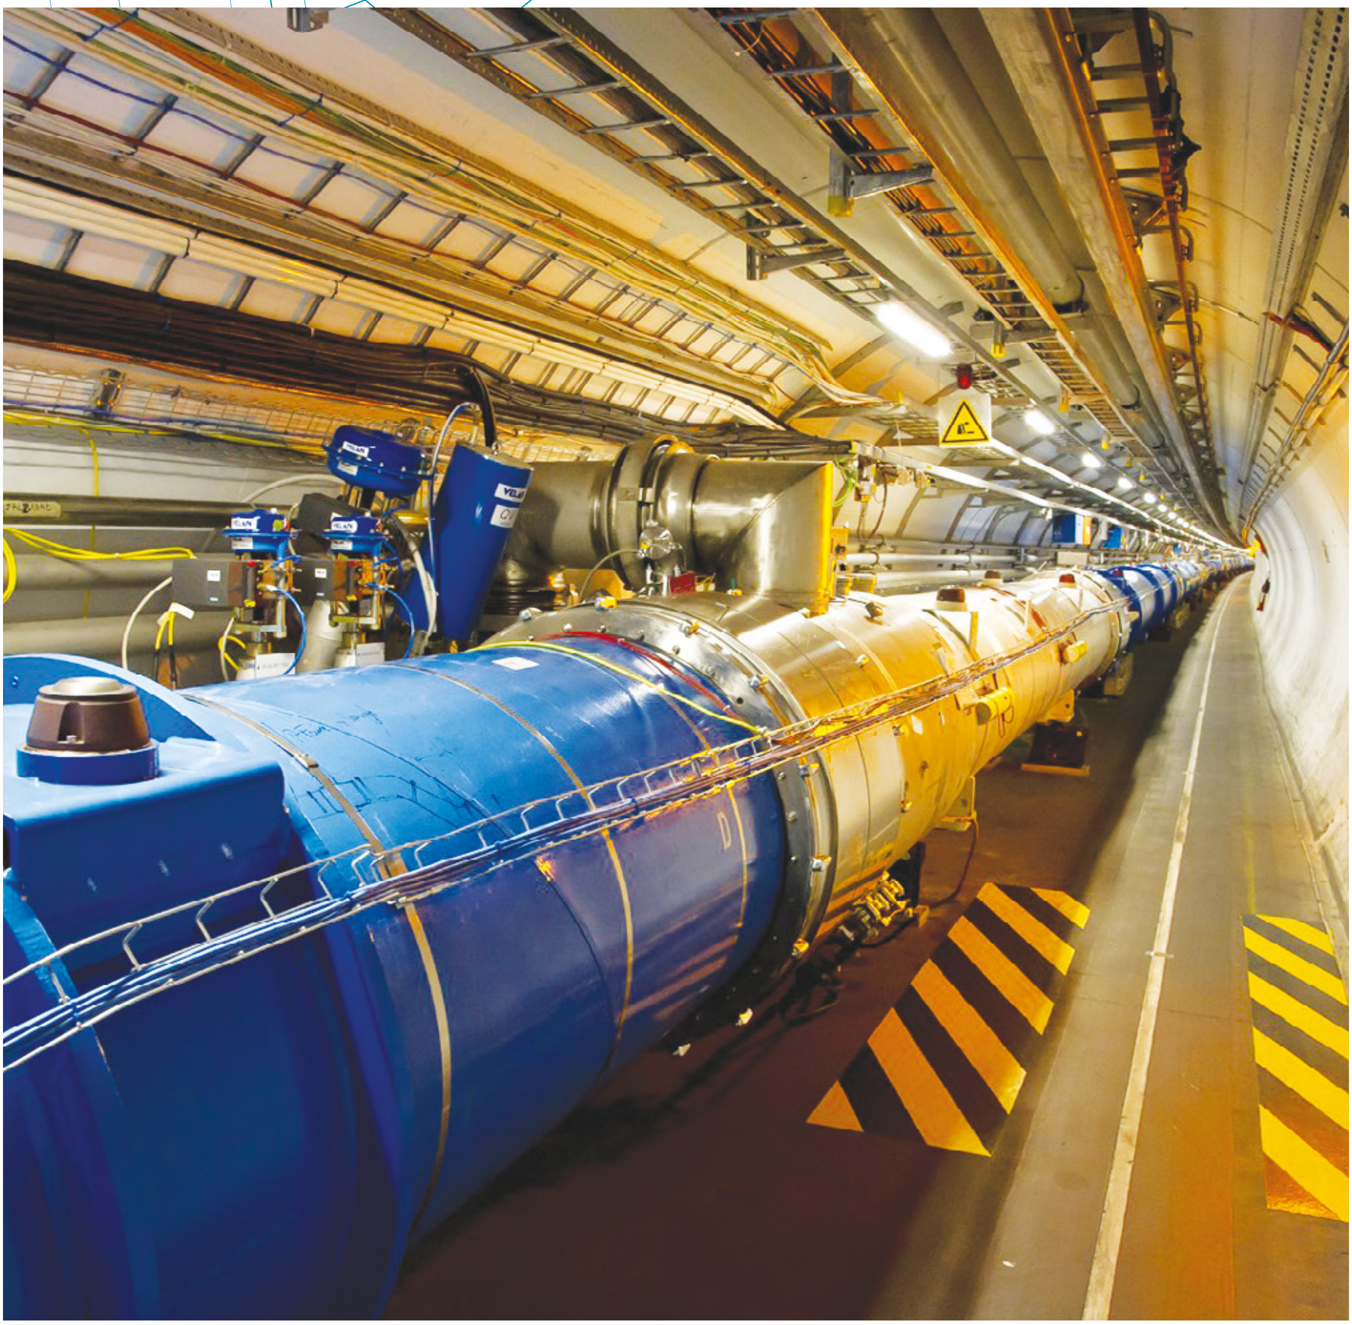

LHC operators at the controls in July 2016.
(CERN-PHOTO-201607-171-3)

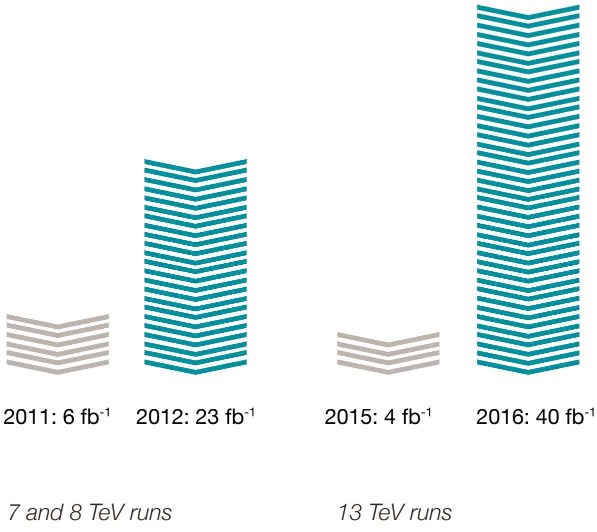

Above, the amount of data delivered by the LHC to the ATLAS and CMS experiments over the course of its proton runs. These quantities are expressed in terms of integrated luminosity, which refers to the number of potential collisions during a given period.

LARGE HADRON COLLIDER IN FULL SWING

2016 was an exceptional year for the Large Hadron Collider (LHC). In its second year of running at a collision energy of 13 teraelectronvolts (TeV), the accelerator's performance exceeded expectations.

The LHC produced more than 6.5 million billion collisions during the proton run from April to the end of October, almost 60% more than originally anticipated. It delivered to the two major experiments, ATLAS and CMS, an integrated luminosity of almost 40 inverse femtobarns, compared with the target of 25. Luminosity, which measures the number of potential collisions per surface unit in a given time period, is a crucial indicator of an accelerator's performance.

The impressive availability of the LHC and its injectors was one of the keys to this success. The LHC was in operation 75% of the time during the physics run (25% for beam commissioning and preparation and 50% providing collisions for the experiments, compared with 33% in 2015). What makes this performance so remarkable is that the LHC is such a complex machine, relying on a chain of four accelerators and thousands of items of equipment.

The start-up of the machine was, moreover, hampered by several incidents. An incident at the end of April involving a transformer stopped the supply of electricity to the accelerator for one week. Then, the supply system of the Proton Synchrotron (PS), the third link in the injector chain, was damaged. Furthermore, a vacuum leak in the beam dump of the Super Proton Synchrotron (SPS), the fourth injector, restricted the quantity of particles that could be injected throughout the run. The number of proton bunches per beam had to be limited to 2220, compared with the 2808 protons per bunch planned. Finally, a fault in an LHC kicker magnet, which injects the particle bunches into the machine, limited the number of particles per bunch.

In spite of these interruptions, the performance of the machines was excellent. Numerous adjustments had been made to all systems over the previous two years to improve their availability. The cryogenics system, for example, which cools the accelerator to - 271°C, achieved 98% availability.

Adjusting the operational parameters of the LHC and its injectors also resulted in an increase in luminosity. The way in which particle bunches are assembled in the injectors was improved (see p. 22) and the beam crossing angle at the point of collision was reduced. These two parameters made it possible to achieve an average of 25 collisions per bunch crossing, compared to 14 in 2015.

The result of these optimisation efforts was a new record, which the operations teams celebrated on 26 June: the LHC surpassed the nominal peak luminosity of $10^{34}cm^{-2}s^{-1}$ defined when the machine was designed 20 years ago. This nominal value was subsequently surpassed repeatedly throughout the year, by as much as 40%.

The LHC thus provided 153 days of physics, several of which were dedicated to running with de-squeezed beams for the forward experiments TOTEM and ATLAS/ALFA.
On 5 December, the last particles circulated in the machine before the technical stop. The first few days of the stop were devoted to magnet training in two of the machine's eight sectors, with a view to reaching a collision energy of 14 TeV. Using the results of this training campaign, the LHC teams will investigate the possibility of increasing the LHC's energy during Run 3, which is scheduled to start in 2021.

The TOTEM spokesperson in front of one of the experiment's detectors in the LHC tunnel, not far from the CMS detector. A special LHC run was provided for TOTEM and ATLAS/ALFA, with the beams as de-squeezed as possible. TOTEM upgraded its detectors in 2016. (CERN-PHOTO-201609-210-1)

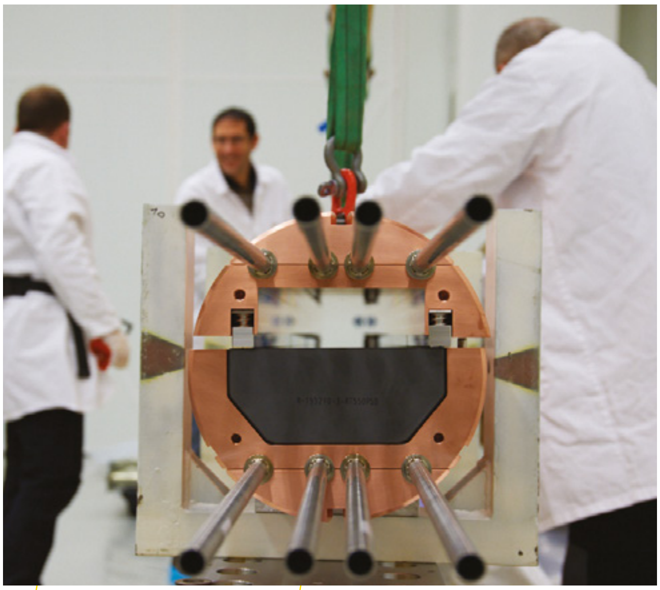

A new SPS beam dump was designed and built in just a few months for installation in the accelerator during the extended year-end technical stop. (CERN-PHO-201704-084-15)

ENCOUNTERS OF THE THIRD KIND

On 10 November, the LHC started colliding protons with the nuclei of lead atoms – the second run of this kind to take place since 2013. This time, the collisions reached the unprecedented energy of 8.16 TeV. The experiments recorded more than 380 billion collisions at this energy, far exceeding expectations.

In order to meet the requirements of the experiments, the LHC and its injectors had to perform some complicated gymnastics. Two energy levels and three machine configurations were employed over four weeks of running. A special run was also conducted for the LHCf experiment. The luminosity reached seven times higher than the nominal value set a few years ago for these runs, and the collision periods were exceptionally long. This performance is particularly noteworthy because colliding protons with lead ions, which have a mass around 206 times greater and a charge 82 times higher than protons, requires numerous painstaking adjustments to the machine. The LHC was able to rely on a chain of highly efficient injectors.

CHAIN ACCELERATION

CERN operates a complex of eight accelerators and one decelerator, which together supply dozens of experiments (see p. 13). These accelerators also serve to propel particles into the LHC. The protons collided in the LHC are first bunched and accelerated by four injectors in series: Linac2,

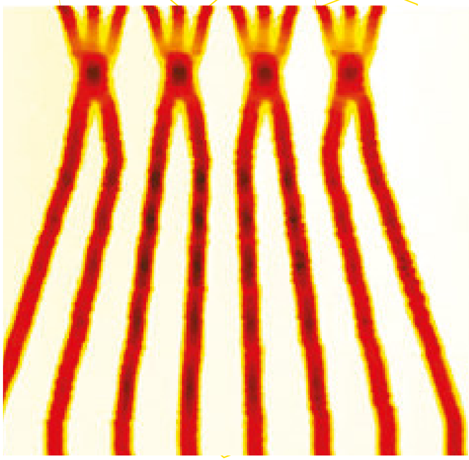

Bright bunches

This image shows how the particle bunches are joined together and pulled apart in the Protron Synchrotron (PS) using the BCMS (batch compression merging and splitting) method. This process, implemented for the first time in 2016, uses the radio-frequency cavities to form denser particle bunches, which increases the probability of collisions occurring in the LHC. These "brighter" bunches, in accelerator jargon, brought about a 20% increase in luminosity in 2016. (OPEN-PHO-ACCEL-2017-008-1)

DISTRIBUTION OF PROTONS DELIVERED BY THE ACCELERATOR CHAIN TO THE DIFFERENT INSTALLATIONS

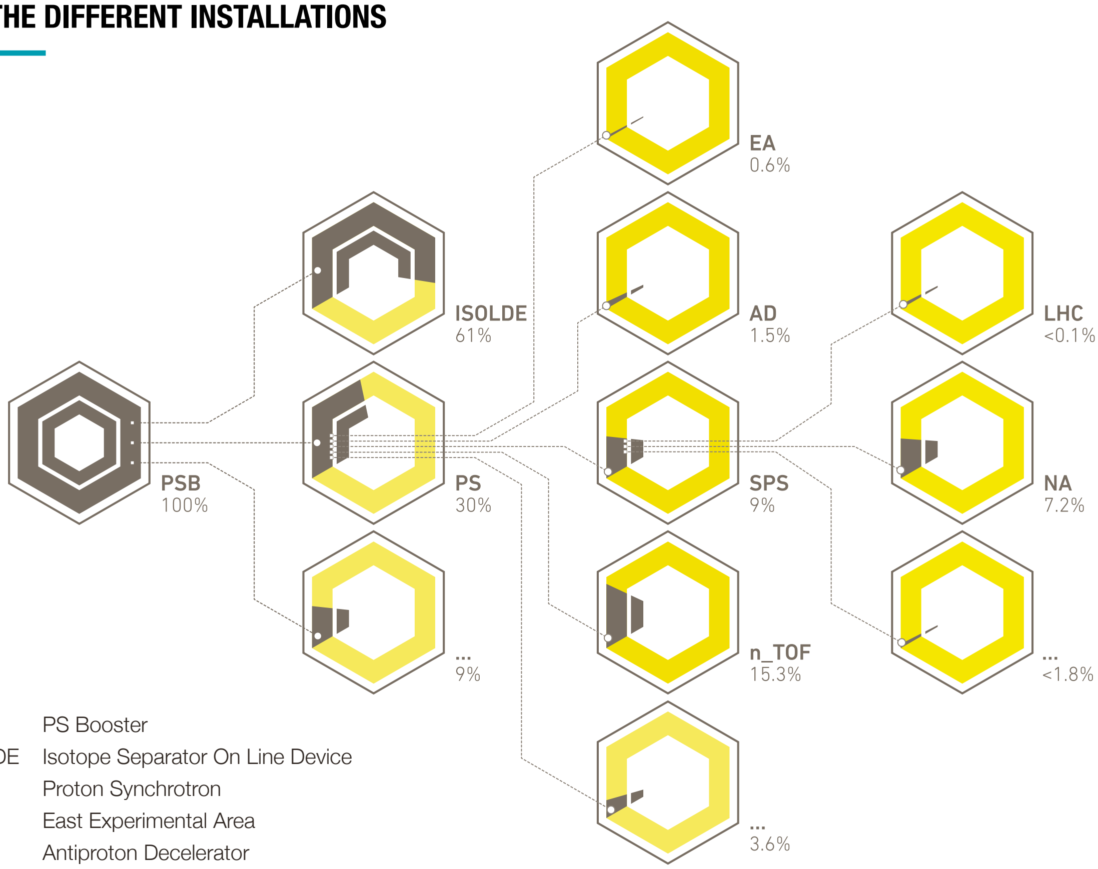

PSB	PS Booster
ISOLDE	Isotope Separator On Line Device
PS	Proton Synchrotron
EA	East Experimental Area
AD	Antiproton Decelerator
SPS	Super Proton Synchrotron
n_TOF	Neutron Time-of-Flight facility
LHC	Large Hadron Collider
NA	North Experimental Area
...	Other uses, including accelerator studies (machine development)

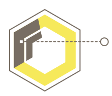

Quantity of protons used in 2016 by each accelerator and experimental facility, shown as a percentage of the number of protons sent by the PS Booster

***1.34×10^{20} protons** were accelerated in the accelerator complex in 2016. This might sound like a huge number, but in reality it corresponds to a minuscule quantity of matter, roughly equivalent to the number of protons in a grain of sand. In fact, protons are so small that this amount is enough to supply all the experiments. The LHC uses only a tiny portion of these protons, less than 0.1%, as shown in the diagram.*

the PS Booster, the Proton Synchrotron (PS) and finally the Super Proton Synchrotron (SPS). Heavy ions are prepared in Linac3 and the Low-Energy Ion Ring (LEIR) before being sent to the PS and the SPS.

The injector chain performed impressively in 2016, achieving availability of over 90% for all the accelerators. For example, the PS Booster, which groups the particles into bunches, achieved 96% availability, supplying particles to the PS and the nuclear physics facility ISOLDE.

The next link in the chain, the PS, redistributes the particle bunches and accelerates them before sending them to various facilities. Half of the protons prepared by the PS go to CERN's other nuclear physics facility, n_TOF. A spare cavity was commissioned in the PS, improving the transmission of particles to the SPS.

The SPS not only supplied the LHC but also provided 80% of the protons required by the experiments in the North Experimental Area. This performance is especially impressive considering that the number of particles in the SPS was limited by a vacuum leak detected in the beam dump. An urgent campaign to produce a new beam dump in just a few months was launched in response to this issue. At the end of the year, the accelerator demonstrated its flexibility by providing lead nuclei at three different energies to nine experiments in the North Experimental Area as well as providing the LHC with protons and lead nuclei.

Linac3 and LEIR, the two accelerators that prepare heavy ions upstream, recorded excellent performances, bearing witness to the success of the upgrade programme that began in 2015. They delivered bunches at very high intensities, even higher than those required by the LHC

A spy in the big tunnel

TIM, the Train Inspector Monorail, is an autonomous mini-vehicle used for real-time monitoring and inspections of the 27-kilometre tunnel of the Large Hadron Collider (LHC). It checks the tunnel structure, the oxygen percentage, the communication bandwidth and the temperature. Suspended from the tunnel's ceiling, it can move at up to 6 km/h.
In 2016, TIM performed several autonomous missions to carry out inspections and radiation-level measurements. The robot has already caught the eye of industry, in particular for autonomous monitoring of utility infrastructures such as underground water pipelines. (OPEN-PHO-TECH-2017-004-1)

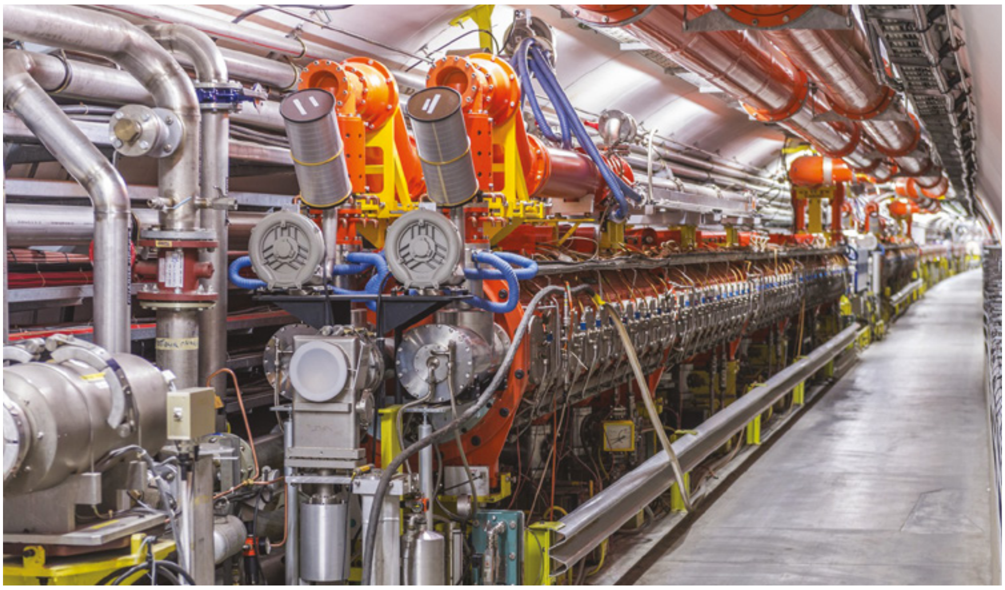

The Super Proton Synchrotron (SPS), CERN's second-largest accelerator, celebrated its 40th birthday in 2016.
An essential link in CERN's accelerator chain, the SPS supplies different kinds of particle to myriad experiments. It accelerates protons and lead ions for the LHC, while also supplying experiments in the North Experimental Area. (CERN-GE-1311288-04)

Injector Upgrade programme being implemented to prepare the injectors for the High-Luminosity LHC (see p. 45).
The Antiproton Decelerator (AD), which sends antiprotons to five experiments, provided 5400 hours of physics in 2016, setting a new record.

While the machines were in operation, the teams were at work behind the scenes to prepare for maintenance and upgrades. For example, several electrical sub-stations, in particular those at the SPS, are being progressively renovated. Studies were also launched with a view to modernising the East Experimental Area during Long Shutdown 2, due to start at the end of 2018.

Once the machines had been stopped in early December, the maintenance and upgrade work began. The 2016-17 year-end technical stop was extended by seven weeks to allow more major work than usual to be carried out.

HIGHER-ENERGY NUCLEAR PHYSICS

The ISOLDE nuclear physics facility received an energy boost. Since September 2016, the new superconducting accelerator HIE-ISOLDE (High-Intensity and Energy ISOLDE) has been increasing the energy of its radioactive ion beams.

Equipped with two cryomodules each containing five superconducting cavities, this machine accelerated six different varieties of radioactive ion to energies ranging from 4.3 to 6.8 MeV per nucleon (compared to a maximum of 3 MeV per nucleon previously). HIE-ISOLDE provided 837 hours of beam time, supplying two experimental areas.

This upgrade makes ISOLDE the only facility in the world capable of studying medium to heavy ions in this energy range. The second phase of the project will involve the installation of two further cryomodules and a third experimental area so that the facility will be capable of reaching up to 10 MeV per nucleon for medium and heavy ions in 2018.

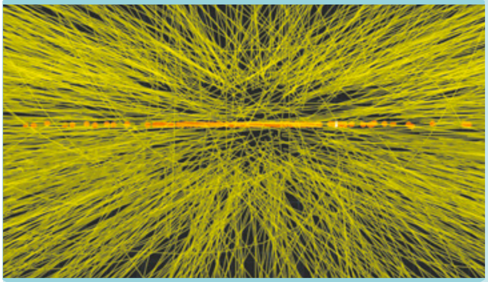

A pile-up of several dozen collisions recorded in one bunch crossing by the CMS experiment in 2016. Each orange dot represents one collision. (CMS-PHO-EVENTS-2016-008-5)

Experiments bombarded with data

The large LHC detectors are extremely complex machines, each made up of millions of parts that have to work together in harmony in order to identify the particles created by the collisions occurring in the accelerator. The unprecedented luminosity of the LHC posed significant challenges for the detectors' triggers, which select the collisions to record. These systems "decide" whether or not to keep the data for each individual collision, on the basis of information transmitted by dedicated sub-systems.

In 2016, the LHC generated an average of 25 collisions simultaneously, 25 million times per second – almost twice as many as in 2015. The experiments adapted their triggers to cope with this pile-up of events. ATLAS optimised its trigger's algorithms and adjusted the way in which the data is handled by the computing grid. The CMS experiment recorded data using a completely upgraded trigger system. These improvements allowed ATLAS and CMS to include more than 90% of the data delivered in their analyses.

The LHCb experiment also employed an improved trigger, as well as a real-time event reconstruction system. This experiment recorded 1.7 inverse femtobarns of proton-proton collision data, five times more than in 2015, and took heavy-ion data for the first time.

The ALICE experiment, which specialises in heavy-ion physics, recorded almost ten times more proton-lead events at an energy of 5.02 TeV than during the previous such campaign in 2013. ALICE used two selection modes: a "minimum bias" trigger that records all types of event without distinction, and a second trigger that selects rare events for specific studies. ALICE also accumulated proton data in line with the goals set.

The third cryomodule was assembled in 2016 and will be installed in 2017.

ISOLDE offers its users a wide range of beams. The facility can produce up to 1000 different isotopes of 75 chemical elements. In 2016, ISOLDE supplied 46 experiments, ranging from those studying the properties of atomic nuclei to biomedical research projects and astrophysics experiments (see p. 18).

A third cryomodule was assembled in 2016 in preparation for the second phase of the HIE-ISOLDE project. HIE-ISOLDE is a superconducting linear accelerator that increases the energy of radioactive ions before they are sent to ISOLDE. (CERN-PHOTO-201603-057-20)

CERN's Data Centre houses servers and data-storage systems not only for the Worldwide LHC Computing Grid, but also for systems critical to the daily functioning of the Laboratory. (OPEN-PHO-CCC-2017-001-1)

COMPUTING: PUSHING THE LIMITS

2016 saw unprecedented volumes of data acquired by the four big LHC experiments, due in large part to the outstanding performance and availability of the LHC itself. Expectations were initially for around 5 million seconds of stable beams, while the final total was around 7.5 million seconds, a very significant and welcome 50% increase. At a higher energy, the collisions themselves are more complex, and at a higher intensity, many collisions overlap, requiring increasingly sophisticated reconstruction and analysis, which has a strong impact on computing requirements. Consequently, 2016 saw records broken in many aspects of data acquisition, data rates and data volumes, and exceptional levels of use of computing and storage resources.

MULTIPLE RECORDS BROKEN

The Worldwide LHC Computing Grid (WLCG) project is a global collaboration of more than 170 computing centres in 42 countries, linking up national and international grid infrastructures. Its mission is to provide global computing resources to store, distribute and analyse the data generated by the LHC. Overall, the performance of the WLCG infrastructure responded very well to the increased needs and levels of use, and enabled, as in previous years, the production of high-quality physics results in a very short time. For example, analyses presented at the major ICHEP physics conference in early August included data acquired only two weeks earlier. In 2016, more than 49 petabytes of LHC data were recorded at the CERN Data Centre, with a striking 11 petabytes in the month of July alone, both being exceptional new records. The large computing capacity of the 170 WLCG sites has been used by the experiments very effectively and efficiently, and they have often made significant use of additional opportunistic computing resources, for example in the form of additional cycles on supercomputer facilities, and from volunteer computing access thanks to LHC@home.

THE NETWORK CHALLENGE

Perhaps one of the most impressive components of the WLCG grid is its networking and connectivity. It can initiate the distribution of data to the hundreds of collaborating institutes worldwide thanks to the excellent connectivity and dedicated networking infrastructure set up at CERN and subsequently worldwide. In 2016, the data transfer rates around the globe also reached new peak rates – between 30 and 40 gigabytes per second continuous rates, around

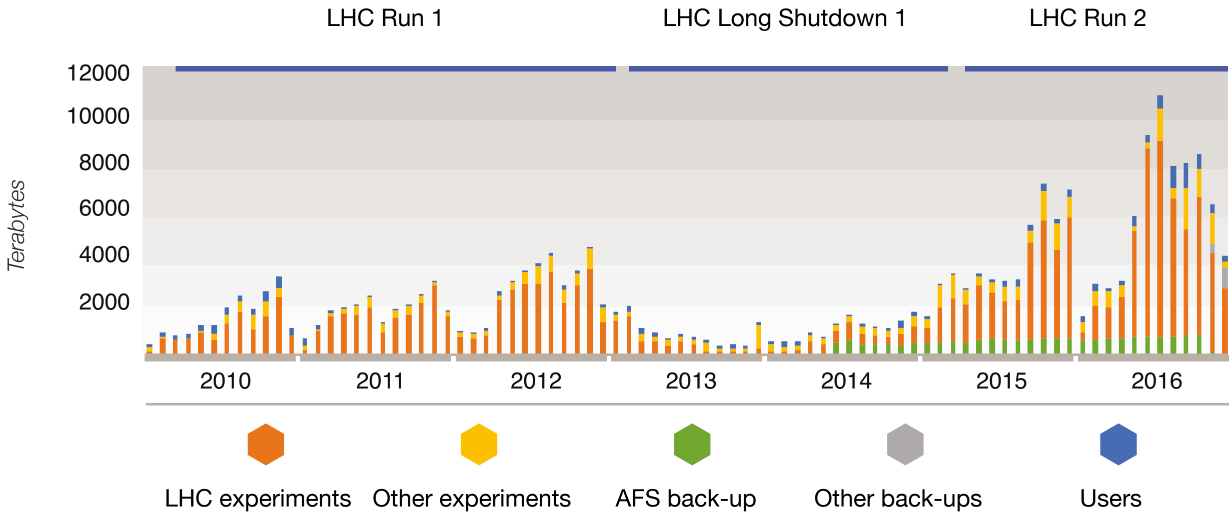

Data recorded on tapes at CERN on a monthly basis

This plot shows, on a monthly basis, the amount of data recorded on tape generated by the LHC experiments, the other experiments, various back-ups, and the users. 2016 was a record year with over 49 petabytes (49 467 terabytes exactly) of LHC data, and a peak of 11 petabytes in July.

a factor of two higher than had been typical during Run 1. The full WLCG collaboration of data centres managed these new rates seamlessly, in some cases thanks to specific adjustments. The increased data rates led several sites to increase the bandwidth of their connection to CERN in order to be able to manage the higher rates. In particular, the increased transatlantic bandwidth put in place for Run 2 was essential.

In addition, significantly increased traffic was observed on the network that allows the National Research and Education Networks (NRENs) to manage LHC data traffic. Consequently, significant additional capacity was deployed into the core network to deal with this increased network use.

SCIENCE IN THE CLOUD

Over 90% of the compute resources in the CERN Data Centre are provided through a private cloud based on OpenStack, an open-source software project that provides on-demand cloud computing. In 2016, CERN was a major contributor to the OpenStack container service development in collaboration with Rackspace and the Indigo DataCloud project, with 40 bug fixes and enhancements in the latest release, which has been recognised by the community with one team member now having achieved core reviewer status, reflecting the quality of the submissions.

With the growth of the computing needs of the CERN experiments and services, this CERN private cloud has now reached over 190 000 compute cores running across the two CERN Data Centres in Meyrin and Budapest. With some of the hardware being retired, over 5000 virtual machines were migrated to new hardware during the year.

For the past three years, investigations have been made to see if physics applications can be run on public cloud resources. With the growth of cloud computing, this approach may be interesting in the future for short-term increases or as a more cost-effective approach for providing compute resources. In addition, this year, we have seen significant and large-scale tests of the use of commercial cloud resources in conjunction with the grid resources. These resources have been provided either in the form of research projects by some of the very large cloud vendors, or through real procurement exercises on a smaller scale.

At the start of the year 2016, CERN also successfully concluded its coordination of the PICSE project addressing Procurement Innovation for Cloud Services in Europe. The Helix Nebula Science Cloud (HNSciCloud) Pre-Commercial Procurement (PCP) project was also kicked off by CERN early on in 2016. HNSciCloud is driven by 10 leading research organisations and brings Europe's technical development, policy and procurement activities together to remove fragmentation and maximise exploitation.

Four consortia engaging a total of 16 companies and organisations were awarded contracts during a ceremony hosted by CNRS in Lyon on 2 November 2016. The Helix Nebula hybrid cloud architecture and procurement model represent significant advances for the sustainability of e-infrastructures in Europe and the planned European Open Science Cloud.

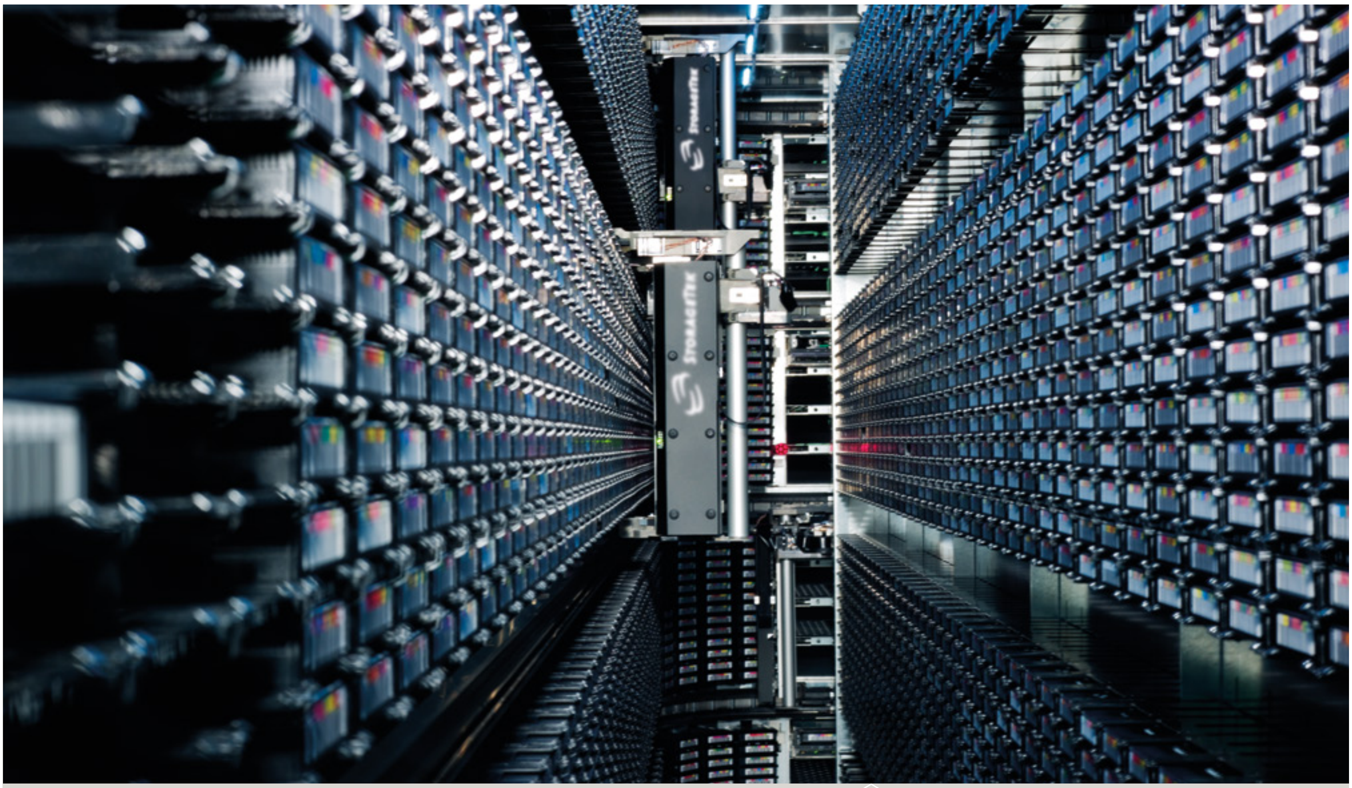

Magnetic tapes, retrieved by robotic arms, are used for long-term storage. (CERN-GE-0809016-01)

Data preservation for future generations

As an organisation with more than 60 years of history, CERN has created large volumes of data of many different types. This involves not only scientific data – to date more than 185 petabytes of data from past and present high-energy physics experiments – but also many other types, including photographs, videos, minutes, memoranda, web pages, etc. CERN hence faces the challenge of preserving its digital memory. Data formats and the tools to access them change constantly, and constant effort is required to tackle the issue, but interestingly, many of the tools that are relevant for preserving data from the LHC and other experiments are also suitable for other types of data. CERN is at the forefront of this effort and participates in the DPHEP (Data Preservation in High Energy Physics) collaboration as a founder member. The status report of the DPHEP collaboration detailing progress during the period 2013-2015 inclusive was published this year in February and is publicly available.

One of the challenges of long-term digital preservation with tape libraries is contamination by environmental hazards, such as dust or any particle that can interfere with the read or write process on tapes. CERN prototyped and built custom environmental sensors for the Data Centre, based on a Raspberry Pi board and an Arduino processor. The sensors behave comparably to proprietary systems in terms of precision and reaction time, but at a small fraction of the cost (about 50 times less than those currently available on the market with similar specifications) and with no maintenance required. In 2016, features were added to the sensor, which can now distinguish between small and large dust particles and can detect brief particle emission in high airflows. The updated sensor effectively prevented a major contamination in the CERN Data Centre in 2016. The use of the sensor is now also being evaluated at CERN for other uses. It is freely available under the CERN Open Hardware License.

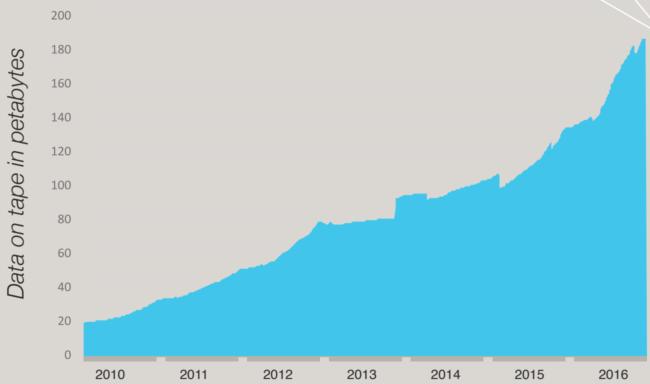

Evolution of the total amount of data stored on tape at CERN
As shown on the graph, the amount of data recorded on tape at CERN is steadily increasing over time, with this trend accelerating in 2016 (+40 % data stored in 2016 compared to what had been accumulated by the end of 2015).

The CERN openlab open day in June marked 15 years of collaboration with industry in support of the LHC research community. (OPEN-PHO-TECH-2016-002-4)

CERN openlab celebrates 15 years of collaboration

For CERN openlab, 2016 marked 15 years of its unique public-private partnership, through which CERN collaborates with leading information and communication technology (ICT) companies and research institutes. Throughout the year, work was carried out to tackle ambitious challenges including activities in domains such as data acquisition, computing platforms, data-storage architectures, compute provisioning and management, networks and communication and data analytics. Work has now begun to identify the ICT challenges that will be tackled in CERN openlab's sixth phase, which will run from 2018 to 2020.

OPEN SOURCE FOR OPEN SCIENCE

The cornerstone of the open source philosophy is that the recipients of technology should have access to all its building blocks, such as software code, schematics for electronics and mechanical designs, in order to study it, modify it and redistribute it to others. Ever since releasing the World Wide Web software under an open-source model in 1994, CERN has continuously been a pioneer in this field, supporting open-source hardware (with the CERN Open Hardware Licence), open access (with the Sponsoring Consortium for Open Access Publishing in Particle Physics - SCOAP3, see p. 35) and open data (with the Open Data Portal for the LHC experiments).

Several CERN technologies are being developed with open access in mind. Invenio is an open-source library management package, now benefiting from international contributions from collaborating institutes, typically used for digital libraries. Invenio 3 was launched in 2016, featuring a whole new concept and full rewrite of the software. CERN, with co-funding from the European Commission, has also long invested in a free open data repository, for use beyond the high-energy physics community: Zenodo.

Zenodo taps into CERN's long-standing tradition and know-how in sharing and preserving scientific knowledge for the benefit of all, giving the scientific community the choice to store its data in a non-commercial environment to be freely available for society at large.

In September 2016, Zenodo was improved with a new release based on Invenio 3: searches became ten times faster, uploads up to 100 GB were also fast, and twice as many visitors and three times more records were handled. The CERN Open Data Portal, which provides seamless access to experimental data, is also built on Invenio. In collaboration with CERN's Scientific Information Service, 300 TB of CMS 2011 data were released. The release received press and media attention and over 210 000 unique site visits.

PLATFORM TECHNOLOGIES FOR OPEN COLLABORATION

The CERN storage system, EOS, was created for the extreme LHC computing requirements. In 2016, EOS instances at CERN approached one billion files, matching the exceptional performances of the LHC machine and experiments. EOS (via the CERNBox project) already fully supports disconnected operations as well as file access and sharing via browsers. It is hence now expanding for other data storage needs across CERN, with about 7000 individual users, and beyond high-energy physics, with AARNET, the Australian Academic and Research Network, and the EU Joint Research Centre for Digital Earth and Reference Data adopting it for their big-data systems.

The Indico conferencing package is another open-source tool developed at CERN and is used by more than 200 sites worldwide, including the United Nations. In 2016, three new versions of Indico were released, improving timetable and category management and the abstract review process.