DISCOVERY MACHINES

To study the infinitesimally small, CERN operates a unique complex of machines. Accelerators speed up particles to almost the speed of light, before making them collide. Detectors record what happens during these collisions. All the resulting data 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.

View of the LHC. In 2017, the accelerator collected more data than expected. (CERN-PHOTO-201802-030-6)

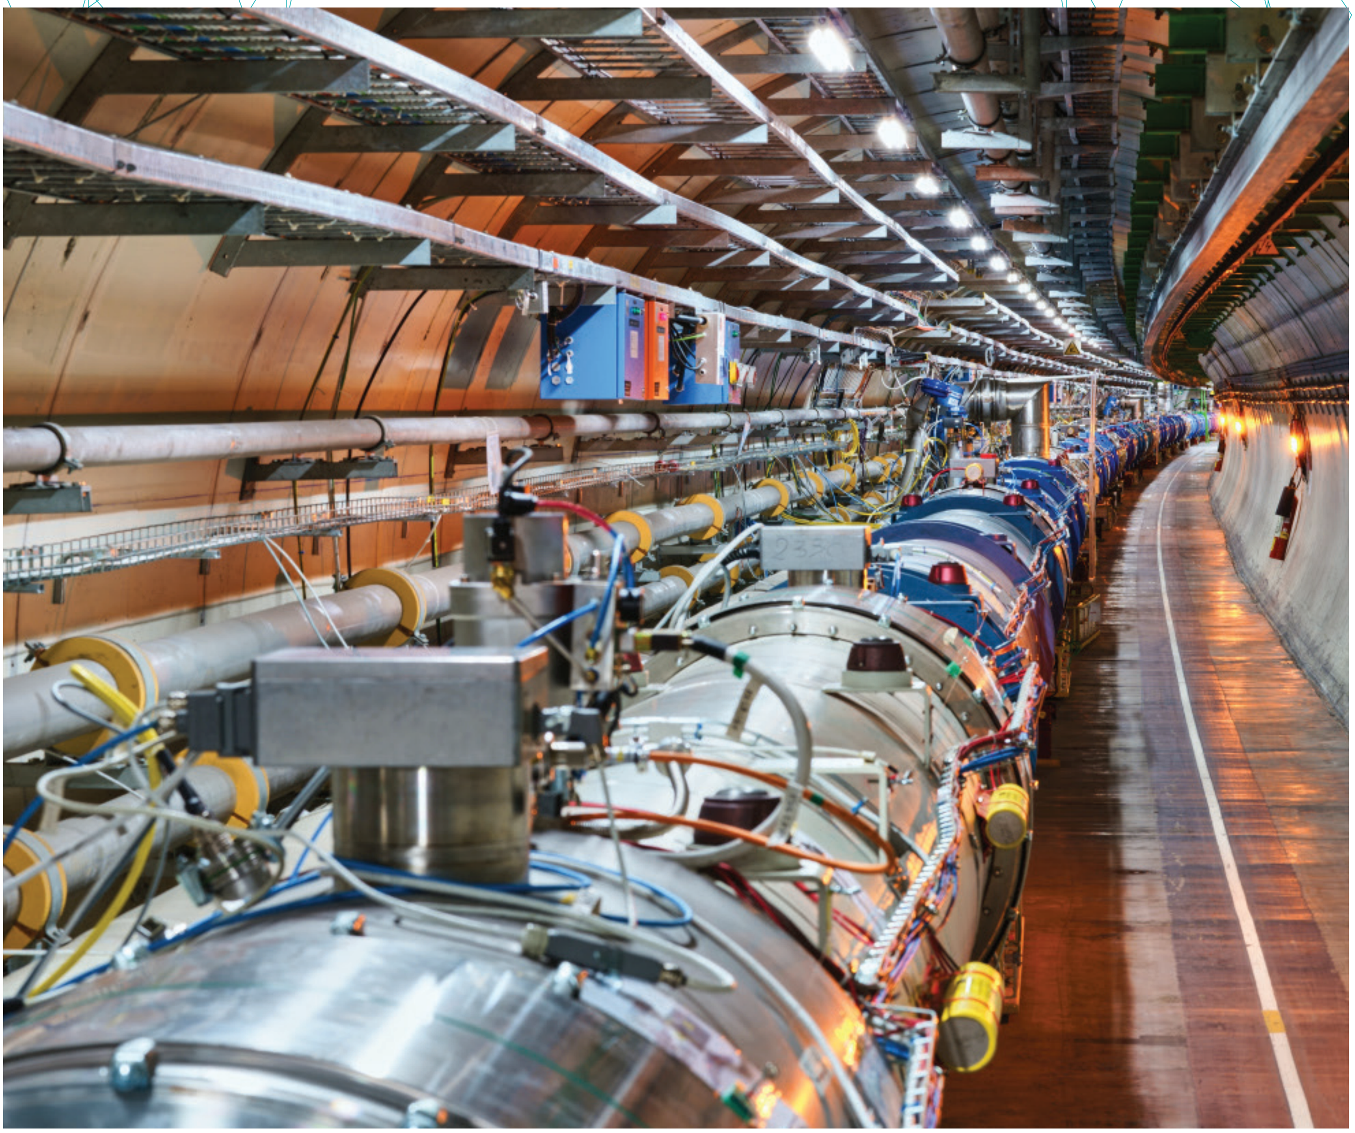

LHC operators in the CERN Control Centre. In 2017, operators tweaked many beam parameters to optimise the LHC's operation. (CERN-PHOTO-201709-214-10)

THE LHC PROVIDED THE EXPERIMENTS WITH 10 MILLION BILLION COLLISIONS, 25% MORE THAN IN 2016.

A FANTASTIC YEAR FOR THE LHC

The Large Hadron Collider (LHC) is the world's most powerful particle accelerator, smashing particles together at unprecedented energies. These collisions enable physicists to study the phenomena that govern particles and forces.

2017 was the LHC's third year at an energy of 13 TeV. The huge ring, measuring 27 kilometres in circumference, produced roughly 10 million billion collisions for the LHC experiments. Two large experiments, ATLAS and CMS, each recorded an integrated luminosity of around 50 inverse femtobarns (fb^{-1}), compared to the 40 expected. Luminosity, which gives the number of potential collisions per surface unit over a given period of time, is the key indicator of an accelerator's performance. The integrated luminosity is measured in inverse femtobarns, 1 fb^{-1} corresponding to around 100 million million (potential) collisions.

This fantastic performance was notably due to the excellent availability of the LHC and its injectors. The LHC was in operation 81% of the time (compared to 75% in 2016), delivering collisions 49% of the time.

To increase luminosity, the operators tweaked various parameters that allow the concentration of the beams to be increased prior to collisions. A new configuration of the accelerator optics, known as Achromatic Telescopic Squeezing (ATS), was used to reduce the size of the proton bunches at the collision points. Instead of using just the quadrupole magnets either side of the experiments to squeeze the bunches, the ATS scheme also makes use of magnets located further away in the machine, transforming seven kilometres of the accelerator into a giant focusing system. ATS was actually developed for the High-Luminosity LHC (see p. 45), but has been successfully tested at the LHC. As a result, each time the proton bunches crossed at the heart of ATLAS and CMS, up to 60 collisions occurred, compared to 40 in 2016.

Three months after the LHC restart, 2556 proton bunches were circulating in the machine – a record. But this smooth operation came to an abrupt halt in August, due to a vacuum issue. To prevent the protons from encountering any obstacles in their path, a high level of vacuum (10^{-10} millibars) is maintained inside the beam pipes. An accidental ingress of air, which froze and condensed on the vacuum chamber wall, disrupted operation for several weeks. The teams found a solution by changing the beam composition (see p. 23), so from the beginning of September onwards a beam comprising 1920 higher density bunches was used.

This new operating mode kept performance levels high. A new peak luminosity record was even set on 2 November, at $2.05 \times 10^{34} cm^{-2}s^{-1}$, more than twice the nominal value.

For the experiments, however, higher density bunches meant more simultaneous collisions, making analysis more difficult. To limit pile-up and to level the luminosity from the beginning to the end of a run, the operators varied the crossing angle and overlap of the beams. These two changes also resulted in an overall luminosity increase of 8%.

The proton collisions ended on 11 November to allow some time for special runs. The first of these consisted of proton collisions at 5.02 TeV, the same energy as planned

The SPS was fitted with a new beam dump during the technical stop and operated at full throttle in 2017. (CERN-PHOTO-201802-048-8)

for the lead-ion run in 2018. This enabled physicists to collect reference data. The second special run, at very low luminosity, was carried out for the benefit of the TOTEM and ATLAS/ALFA experiments, which study the elastic scattering that occurs when two protons interact without colliding. For these studies, the beams were de-squeezed as much as possible and their energy was restricted to 450 GeV. Other tests were carried out over the course of the year, mainly in preparation for the High-Luminosity LHC.

THE ACCELERATOR COMPLEX IN FULL SWING

CERN operates a complex of eight accelerators and one decelerator, which together supply dozens of experiments (see p. 12). The complex also accelerates particles for the LHC. The protons collided in the LHC are produced in four accelerators in series: Linac2, 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. A total of 1.51 x 10^{20} protons were accelerated in the complex in 2017, which is roughly the number of protons in a grain of sand. The LHC uses less than 0.084% of these protons.

In 2017, the accelerator chain achieved an average availability of over 90%, rising to 99% for Linac2 and 97% for the PS Booster. This performance is especially remarkable given that the youngest of these accelerators, the SPS, is over 40 years old, and the oldest, the PS, is approaching the ripe old age of 60! The accelerator fault-tracking system (AFT), installed in the LHC in 2015, is now used across the entire complex, supplying a constant stream of data on the availability of the accelerators, indicating the origin of any faults and helping identify areas for improvement.

The PS Booster supplies particles to the PS and the ISOLDE nuclear physics facility, which turned fifty in 2017. ISOLDE received 9.28 x 10^{19} protons, or just over 61% of the total number injected into the accelerator chain. Moreover, thanks to its new HIE-ISOLDE superconducting accelerator, ISOLDE was able to produce radioactive ion beams at unprecedented energies. Equipped with a third cryomodule, HIE-ISOLDE delivered beams with energies of up to 7 MeV per nucleon to 12 experiments. The fourth and final cryomodule has been assembled and will be installed in 2018.

The next link in the accelerator chain, the PS, redistributes the particle bunches and accelerates them before dispatching them to various facilities. Most of the protons prepared by the PS were sent to the nuclear physics facility n_TOF. The accelerator also feeds the Antiproton Decelerator, which supplied antiprotons to five antimatter experiments over the course of 5500 hours of operation.

The SPS was fitted with a new beam dump at the beginning of the year and was thus able to operate at full throttle in 2017, supplying the LHC, the experiments in the North Area, the HiRadMat test facility and the AWAKE experiment (see p. 49).

Linac3 and LEIR, the two accelerators that prepare heavy ions, demonstrated their flexibility by preparing their first xenon nuclei (see p. 24).

PROTON DISTRIBUTION

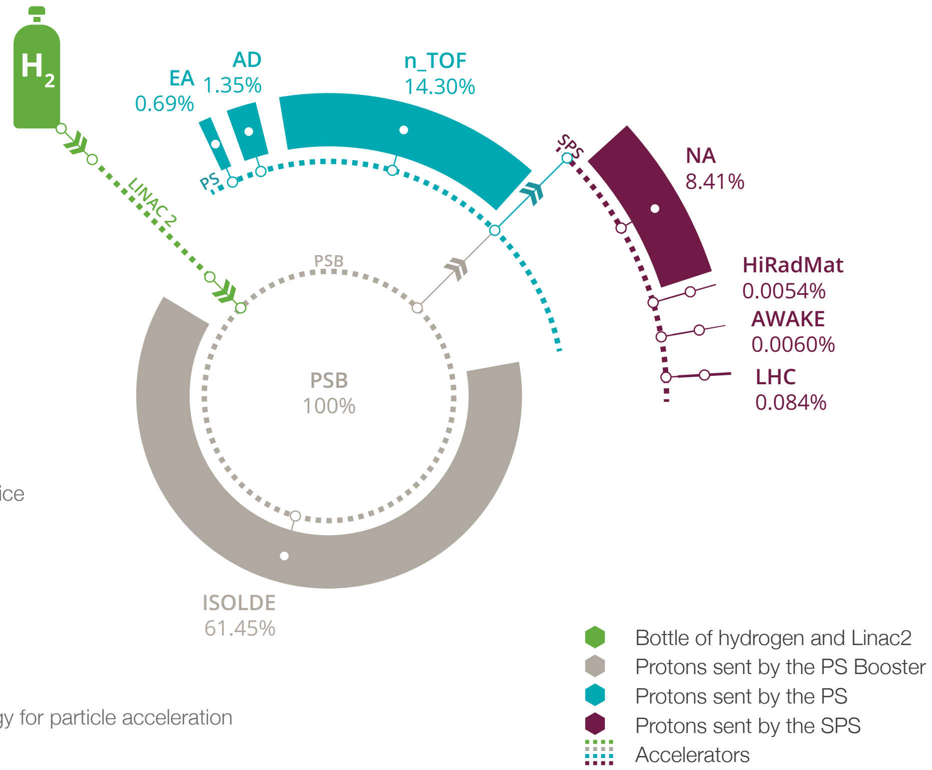

H_2	Hydrogen bottle
Linac2	Linear accelerator
PSB	PS Booster
ISOLDE	Isotope Separator On-Line Device
PS	Proton Synchrotron
EA	East Experimental Area
AD	Antiproton Decelerator
SPS	Super Proton Synchrotron
HiRadMat	Testing facility
n_TOF	Neutron Time-of-Flight facility
AWAKE	Experiment on a new technology for particle acceleration
LHC	Large Hadron Collider
NA	North Experimental Area
...	Other uses, including accelerator studies (machine development)

Distribution of protons to the various experimental facilities, expressed as a percentage of the total number of protons prepared in the PS Booster

The accelerator complex sends particles (mainly protons) to a dazzling array of experiments. In 2017, 151 billion billion protons (1.51×10^{20}) were prepared, which is actually a very small amount of matter, equivalent to the number of protons in a grain of sand. Most of these particles are used by the ISOLDE and n_TOF facilities. The LHC used just a tiny fraction of these: 0.084%. Around 14% of the particles are used for operating tests (machine development) or are not used at all (beam dumps, losses, etc.)

Clearing the air

A machine as complex as the LHC, with tens of thousands of components working together, is in constant need of adjustments and repairs. Sometimes, more serious issues require the teams to be creative. In August, the operators had to deal with mysterious beam losses at a specific point in the ring.

The cause? The presence of air in the vacuum chamber where the beams circulate. Seven litres of air had seeped in: something and nothing compared to the 110 000-litre volume of the vacuum chambers, but enough to disrupt the flow of the protons. Frozen molecules trapped on the walls of the pipe were stripped off when the beam passed, generating showers of electrons that became greater as the intensity of the beam increased. A working group hurried to the LHC's bedside to find a cure, knowing it was impossible to open the vacuum chamber to extract the gas.

Installing a solenoid magnet to circulate the electrons resulted in a small improvement. In the end, the accelerator specialists solved the problem by changing the composition of the beam in the injectors, namely the PS Booster and the PS, splitting the beams into bunches of varying intensity, density and structure. Thanks to the great flexibility of these two machines, the operators were able to propose several alternative beam configurations. Ultimately, a less intense beam was used, made up of sequences of eight proton bunches followed by four empty slots.

This new configuration limited the heat reaching the vacuum chamber walls and thus the phenomenon of electron clouds. The number of bunches had to be reduced to 1920, but since they were denser, excellent luminosity was maintained in the LHC. During the year-end technical stop, the teams partially warmed up the sector and pumped out 8.4 g of gas.

As part of the consolidation programme, a new accelerating cavity has been produced for the LHC, with performance levels higher than specification. (CERN-PHOTO-201803-084-1)

A MAJOR CONSOLIDATION PROGRAMME

The success of the physics programme relies upon the excellent availability of the accelerators. This is why the technical teams closely monitor every nut and bolt of this huge complex. A major consolidation programme is under way on all the machines: the LHC, the injectors and the associated experimental facilities.

In 2017, maintenance and consolidation work took place during the extended year-end technical stop, which lasted until the end of April, and during a few scheduled stops over the course of the year. The teams carried out routine maintenance activities, fixed faults and upgraded certain systems. For example, an LHC dipole magnet and the SPS beam dump were replaced. Another key job was the replacement of warm magnet power converters, thereby increasing the LHC's resistance to electrical faults, such as network voltage variations.

Consolidation work for all the accelerators focused on the beam instrumentation and on the injection, kicker, radiofrequency and vacuum systems. The technical infrastructures were also a hive of renovation activity, from the electrical distribution system to the lifts, overhead travelling cranes and ventilation system. Programmable

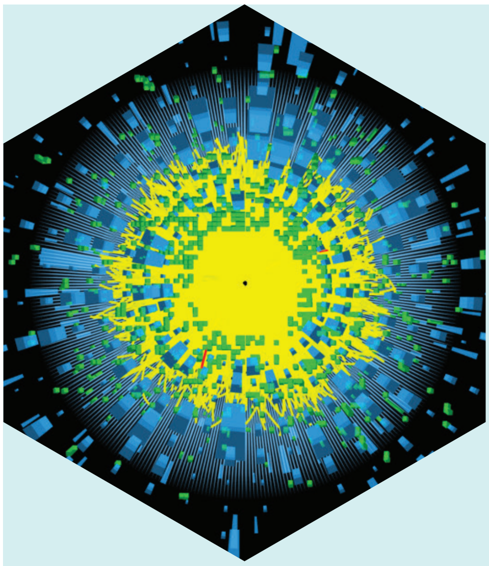

A special run gave the LHC experiments an opportunity to collect data from xenon-ion collisions, like this one, recorded by the CMS detector. (CMS-PHO-EVENTS-2017-007-1)

A spoonful of xenon

On 12 October, the LHC had a taste of something rather different. For eight hours, the accelerator collided xenon nuclei for the first time. The collisions (xenon atoms with 54 protons and 75 neutrons) were similar to the heavy-ion collisions that are regularly performed at the LHC. The LHC experiments were able to collect a brand new type of data.

This innovation came about due to a special run organised for the NA61/SHINE fixed-target experiment, which received xenon ions at six different energies from the Super Proton Synchrotron (SPS), over a period of eight weeks.

This also gave the operators the opportunity to try out some new ideas, injecting partially ionised xenon atoms into the SPS, conserving 15 of their 54 electrons. These beams are very fragile. The teams managed to accelerate a beam, reaching an energy of 81.6 GeV per nucleon. The aim was to test the idea of a high-intensity gamma-ray source with energies of up to 400 MeV, as part of the Physics Beyond Colliders study (see p. 19).

logic controllers are widely used across CERN, and several groups, including those working on cryogenics and ventilation, renovated their systems. The group responsible for accelerator controls carried out maintenance work on its infrastructures.

In offices and workshops, teams manufactured and ordered replacement parts and prepared the upgrades. For example, additional accelerating cavities for the LHC are now being manufactured – work that is crucial to preserving this expertise at CERN. New collimators are also being developed for the High-Luminosity LHC (see p. 45).

Groups from all technical fields also prepared for the major programme of work to be carried out during Long Shutdown 2.

ONE BILLION COLLISIONS PER SECOND WERE PRODUCED AT THE HEART OF THE LARGE LHC DETECTORS.

Open-heart surgery for CMS: the pixel tracker, the heart of the experiment, was replaced with a superior system before the 2017 LHC restart. (CERN-PHOTO-201703-062-43)

A fresh crop of collisions

Every year, the LHC's large detectors achieve extraordinary feats in their quest to collect more and more data. These enormous devices, comprising millions of components, operate at mind-boggling speed to identify the particles that emerge from collisions. In 2017, the accelerator's performance pushed the detectors harder than ever before. The LHC has produced an average of almost 40 collisions with each beam crossing (up to 30 million times per second). By the end of the run, there were even 60 collisions per crossing, two and a half times the value for which ATLAS and CMS were designed. Some one billion collisions per second were produced at the heart of the two large detectors. The two experiments managed to record more than 90% of the data delivered by the LHC, i.e. more than 45 inverse femtobarns each.

At the start of the year, CMS underwent a heart transplant – the detector at its very heart, the pixel tracker, was replaced with a new system comprising an additional pixel layer and 124 million instead of 66 million pixels. The new tracker increases the detection precision, enabling the experiment to cope better with the data pile-up. Thanks to its additional layer, the tracker continued to perform well despite the failure of 5% of the power converters. At the end of the year, all the converters were replaced as a precaution. The CMS detector also benefited from improvements made to the electronics and readout systems of several of its subdetectors.

The ATLAS collaboration also made several improvements before the restart, repairing its calorimeters and muon detection system, replacing the gas distribution system of its transition radiation tracker (TRT) and reconfiguring the front-end electronics of the inner layer (IBL) of its pixel detector. These and other improvements enabled ATLAS to take data with unprecedented quality and efficiency in 2017.

LHCb, having reinforced its trigger and real-time event reconstruction systems, took 1.8 fb^{-1} of data. Its PC farm was also upgraded. As well as the proton–proton and xenon–xenon collisions, LHCb also took data in "fixed-target" mode with neon. It does this by introducing the noble gas into its vacuum chamber, generating collisions between the injected atoms and the proton beam.

ALICE, which specialises in the physics of the quark-gluon plasma obtained in heavy ion collisions, recorded 986 million proton–proton events at an energy of 5.02 TeV, exceeding its goal. This data taking enabled the experiment to collect reference measurements in preparation for the 2018 lead-ion run. Previously, ALICE had collected one billion proton–proton events (with a "minimum bias" trigger), a fraction of them with a reduced magnetic field for specific studies. The experiment also recorded "high-multiplicity" events, data with specific triggers and xenon–xenon events.

The CERN Data Centre houses servers, data-storage systems and networking equipment, not only for the Worldwide LHC Computing Grid, but also for systems critical to the everyday operations of CERN. (CERN-CO-1008289-04)

COMPUTING: DEALING WITH THE DATA DELUGE

Once again this year, the CERN computing infrastructure faced a data deluge from the LHC and dealt with it successfully. Indeed, even though overall the experiments have made tremendous progress on reducing the number of derived data files written to tape, several data-storage records were broken. The CERN Advanced STORage system (CASTOR) reached the milestone of 200 petabytes of data stored on tape by the end of June, and by the end of the year over 230 petabytes had been accumulated. Towards the end of the year, the performance of the LHC exceeded expectations and delivered a higher integrated luminosity to ATLAS and CMS than in any previous year. In October alone, 12.3 petabytes were written to tape, setting a new monthly record.

Data storage is an essential and critical component of CERN's computing infrastructure, which needs to be continuously updated to cope with ever-increasing demand while remaining within the confines of the premises. The Organization currently manages the largest scientific data archive in the field of High-Energy Physics (HEP) and is at the forefront of the field's data-preservation effort; CERN is a founding member of the DPHEP (Data Preservation in High Energy Physics) collaboration. A consequence of the greater data volumes in 2017 was increased demand for data transfer and thus a need for a higher network capacity. Since early February, a third 100-gigabit-per-second fibre-optic circuit has linked the CERN Data Centre with its remote extension hosted 1200 km away at the Wigner Research Centre for Physics (RCP) in Budapest, Hungary. The additional bandwidth and redundancy provided by this third link allowed CERN to benefit reliably from the computing power and storage at the remote extension.

WLCG: TOWARDS THE FUTURE

The mission of the Worldwide LHC Computing Grid (WLCG) is to provide global computing resources for the storage, distribution and analysis of the data generated by the LHC. As in previous years, WLCG performed well in 2017 and adapted to the increasing performance needs of the experiments, supporting the timely delivery of high-quality physics results. The CERN Data Centre continued to carry out essential processing and quality-checking of the data, although, due to the excellent performance of the LHC, the experiments started to see some processing backlogs in October, clearly illustrating the need for additional resources for the coming year when similar performance is expected. The export of data to the WLCG data centres proceeded smoothly, thanks to upgrades of network bandwidths in the preceding years. Finally, the collaboration

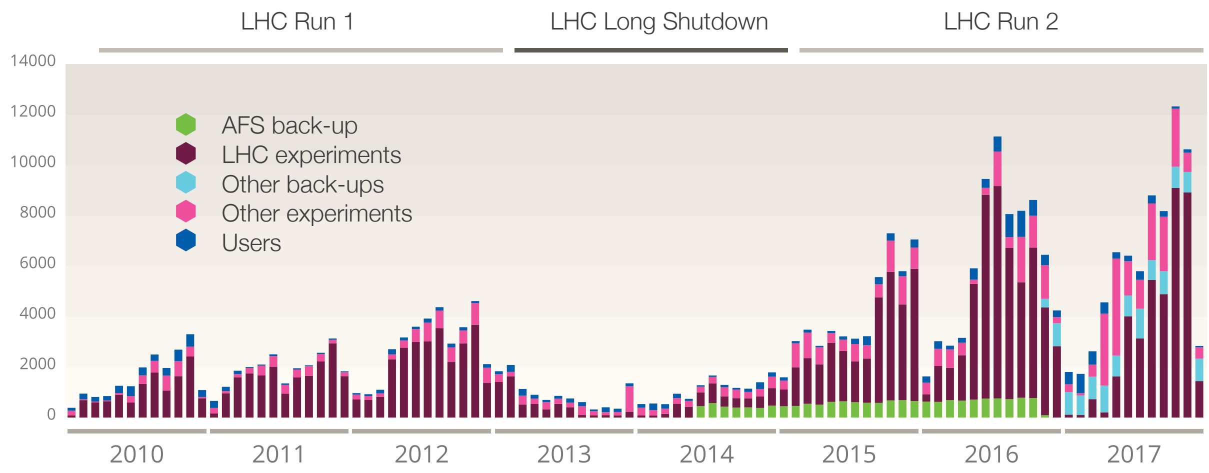

Data (in terabytes) recorded on tape at CERN month-by-month

This plot shows the amount of data recorded on tape generated by the LHC experiments, other experiments, various back-ups and users. In 2017, 72 petabytes of data in total (including 40 petabytes of LHC data) were recorded on tape, with a record peak of 12.3 petabytes in October.

of 170 participating data centres has provided the required combined capacity to process and analyse these data in an effective way and allowed service delivery to be stepped up on multiple occasions. All volunteer computing activities at CERN were consolidated in 2017 under the LHC@home project umbrella and have continued to grow, providing peaks of 400 000 simultaneously running tasks.

Looking to the future, the High-Luminosity LHC (HL-LHC) presents a major and significant increase in computing requirements compared with the current Run 2 and the upcoming Run 3; in fact, the demand exceeds the extrapolations of technology progress within a constant investment budget by several factors, in terms of both storage and computing capacity. In order to address this shortfall, in 2017 the HEP Software Foundation (HSF) produced a roadmap Community White Paper (see p. 54) that explores these software challenges and describes a number of potential mechanisms by which the community may overcome the shortfalls in capacity and efficiency it will face during the next decade.

The large-scale data challenges of the High-Luminosity LHC are being addressed in collaboration with other scientific communities. For example, WLCG and CERN have been working with the Square Kilometre Array (SKA) radio astronomy community to explore some of these common challenges. In June 2017, CERN and the SKA Organisation signed a collaboration agreement for joint work on computing and data management.

Evolution of the global core processor time delivered by the Worldwide LHC Computing Grid (WLCG)

As seen on the graph, the global central processing unit (CPU) time delivered by WLCG (expressed in billions of HS06 hours per month, HS06 being the HEP-wide benchmark for measuring CPU performance) shows a continual increase. In 2017, WLCG combined the computing resources of about 800 000 computer cores.

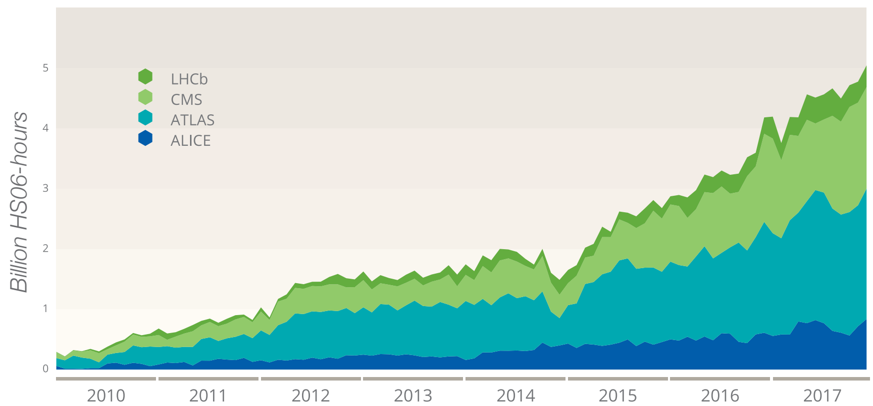

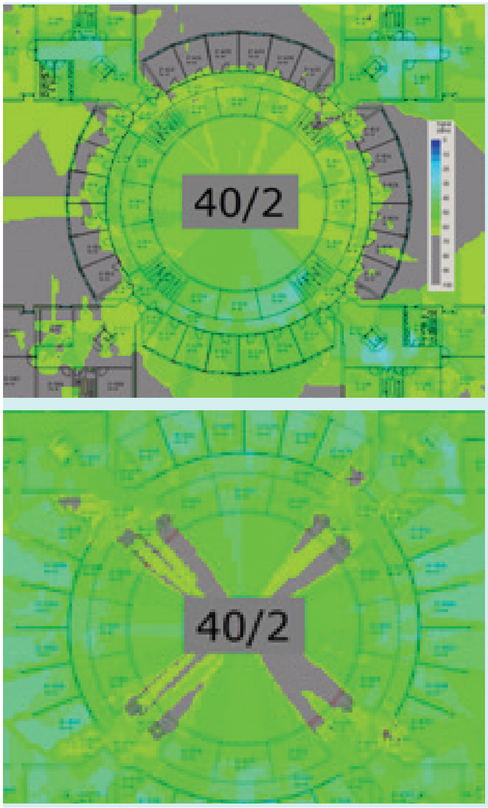

The new campus-wide Wi-Fi service reaches key milestones

The new campus-wide Wi-Fi service project started a few years ago, as part of an overall effort to modernise the Laboratory's communication infrastructure to meet today's demands for mobility, flexibility and accessibility. After a pilot phase in 2016, 2017 saw the start of deployment of the new Wi-Fi infrastructure in some 70 buildings at CERN. The new access points are centrally managed, which allows users to move around without losing their network connection.

Thanks to the new Wi-Fi service, it is now also possible to offer a visitor service similar to the one available in many public spaces, while keeping the visitors' devices isolated from CERN's internal network.

Wi-Fi service coverage before (top) and after (bottom) the deployment of the new infrastructure in Building 40, which hosts a large number of physicists.

SCIENCE IN THE CLOUDS

Over 90% of the resources for computing in the Data Centre are provided through a private cloud based on OpenStack, an open-source project to deliver a massively scalable cloud operating system. With the growth of the computing needs of the CERN experiments and services, this has now reached more than 280 000 computer cores running in the data centres at CERN and Wigner. Migrations of system software this year have simplified the configuration and allowed the capacity to continue to grow. A collaboration project was also started with SKA to enhance OpenStack for scientific use cases and to share experiences. Many of the contributions to OpenStack have been made through CERN openlab in collaboration with industry.

During 2017, four projects co-funded by the European Commission (EC) in which CERN was involved were successfully completed: ICE-DIP (Intel CERN Industrial Doctorate Programme), EGI-Engage (Engaging the EGI Community towards an Open Science Commons), INDIGO DataCloud (INtegrating Distributed data Infrastructures for Global ExplOitation of PaaS/SaaS level services), and AARC (Authentication and Authorisation for Research and Collaboration). The positive results of these projects have led to CERN's engagement in new projects that will contribute to the European Open Science Cloud (EOSC).

The Helix Nebula Science Cloud (HNSciCloud) Pre-Commercial Procurement (PCP) project led by CERN is addressing the cloud services needs of a group of 10 research organisations (procurers) that serve data-intensive research communities. In 2017, HNSciCloud successfully completed design and prototype phases and launched the pilot phase of the pre-commercial procurement process. HNSciCloud has been instrumental in promoting and demonstrating the advantages of procuring commercial cloud services in the public research sector. The impact of the project on European-level policy is apparent in the development of the EOSC, where commercial cloud services are now recognised as playing an important role.

EDUCATION AND SHARING

Since the early seventies, the CERN School of Computing (CSC) has been promoting learning and knowledge exchange in scientific computing among young scientists and engineers involved in particle physics or other sciences. It is now made up of three separate schools and this year celebrated the 40th edition of its main school, which took place in Madrid, Spain. Since its inception, the CSC has been attended by 2600 students from 80 countries on five continents.

The CERN and Wigner data centres together host around 15 000 servers, which are replaced every four to five years as they become inefficient for the purposes of CERN's research. However, they remain suitable for less demanding applications. About 1990 boxes of new servers and data storage equipment were deployed in 2017, while 1000 boxes of old equipment were retired. Nine pallets of IT equipment were donated by CERN in 2017: to CERIST in Algeria (in February), to Sofia University in Bulgaria (in August), and to SESAME in Jordan (in September, see p. 10).

CERN inaugurates its second network hub

The day-to-day operation of CERN is heavily reliant on IT services, and preventing any potential long-term interruption is therefore crucial. To this end, a project to create a second network hub was approved in 2014 and, three years later, the new network hub started operations. With fibre connectivity to the outside world, the CERN Control Centre, the CERN Data Centre and its extension at the Wigner Research Centre for Physics in Hungary, as well as to all the major network technical rooms on the CERN campus, this second network hub provides the necessary redundancy for the CERN data network.

The second CERN network hub, located near the CERN Control Centre in Prévessin, was inaugurated on 19 July.

The CERN Open Data portal is a testimony to CERN's policy of Open Access and Open Data. The portal allows the LHC experiments to share their data with a double focus: for the scientific community, including researchers outside the CERN experimental teams, as well as citizen scientists, and for the purposes of training and education through specially curated resources. This year, a major new release was made containing over one petabyte of CMS data, corresponding to around half of the data collected by the detector in 2012. In addition to the datasets themselves, the CMS Data Preservation and Open Data team have also assembled a comprehensive collection of supplementary material, including sample code for performing relatively simple analyses, as well as metadata such as information on how data were selected and on the LHC's running conditions at the time of data collection. The first papers based on data from the CERN Open Data portal have been published.

In January 2017, a consortium of European companies, research laboratories, universities and education networks came together to form Up to University (Up2U), a project co-funded by the European Commission (EC). CERN is playing an active part in this three-year project, which aims to create a bridge between high schools and higher education. The objective is to provide tools and technology commonly used in academia and big science to secondary-school students in order to prepare them for their future careers.

FOSTERING COLLABORATION

CERN, with co-funding from the EC, has long invested in Zenodo, a free repository for storing data, software and other research artefacts. It is intended for use beyond the HEP community and taps into CERN's long-standing tradition and know-how in sharing and preserving scientific knowledge for the benefit of all. Zenodo is hosted at CERN and provides the wider scientific community with the option of storing its data in a non-commercial environment and making it freely available to society at large. In 2017, Zenodo incorporated Digital Object Identifier (DOI) versioning. This new and highly requested feature allows users to update a record's files after it has been made public and allows researchers to reference one or all the versions of a record.

CERN openlab's fifth three-year phase came to a close at the end of 2017. Through this unique public-private partnership, CERN has collaborated with leading ICT companies and other research organisations to accelerate the development of cutting-edge ICT solutions for the research community. Through no fewer than 20 R&D projects, CERN openlab tackled ambitious challenges covering the most critical needs of ICT infrastructures. Throughout 2017, CERN openlab's work also focused on the preparations for its sixth phase, beginning in 2018. The white paper on IT Challenges in Scientific Research (see p. 54) published in September 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.