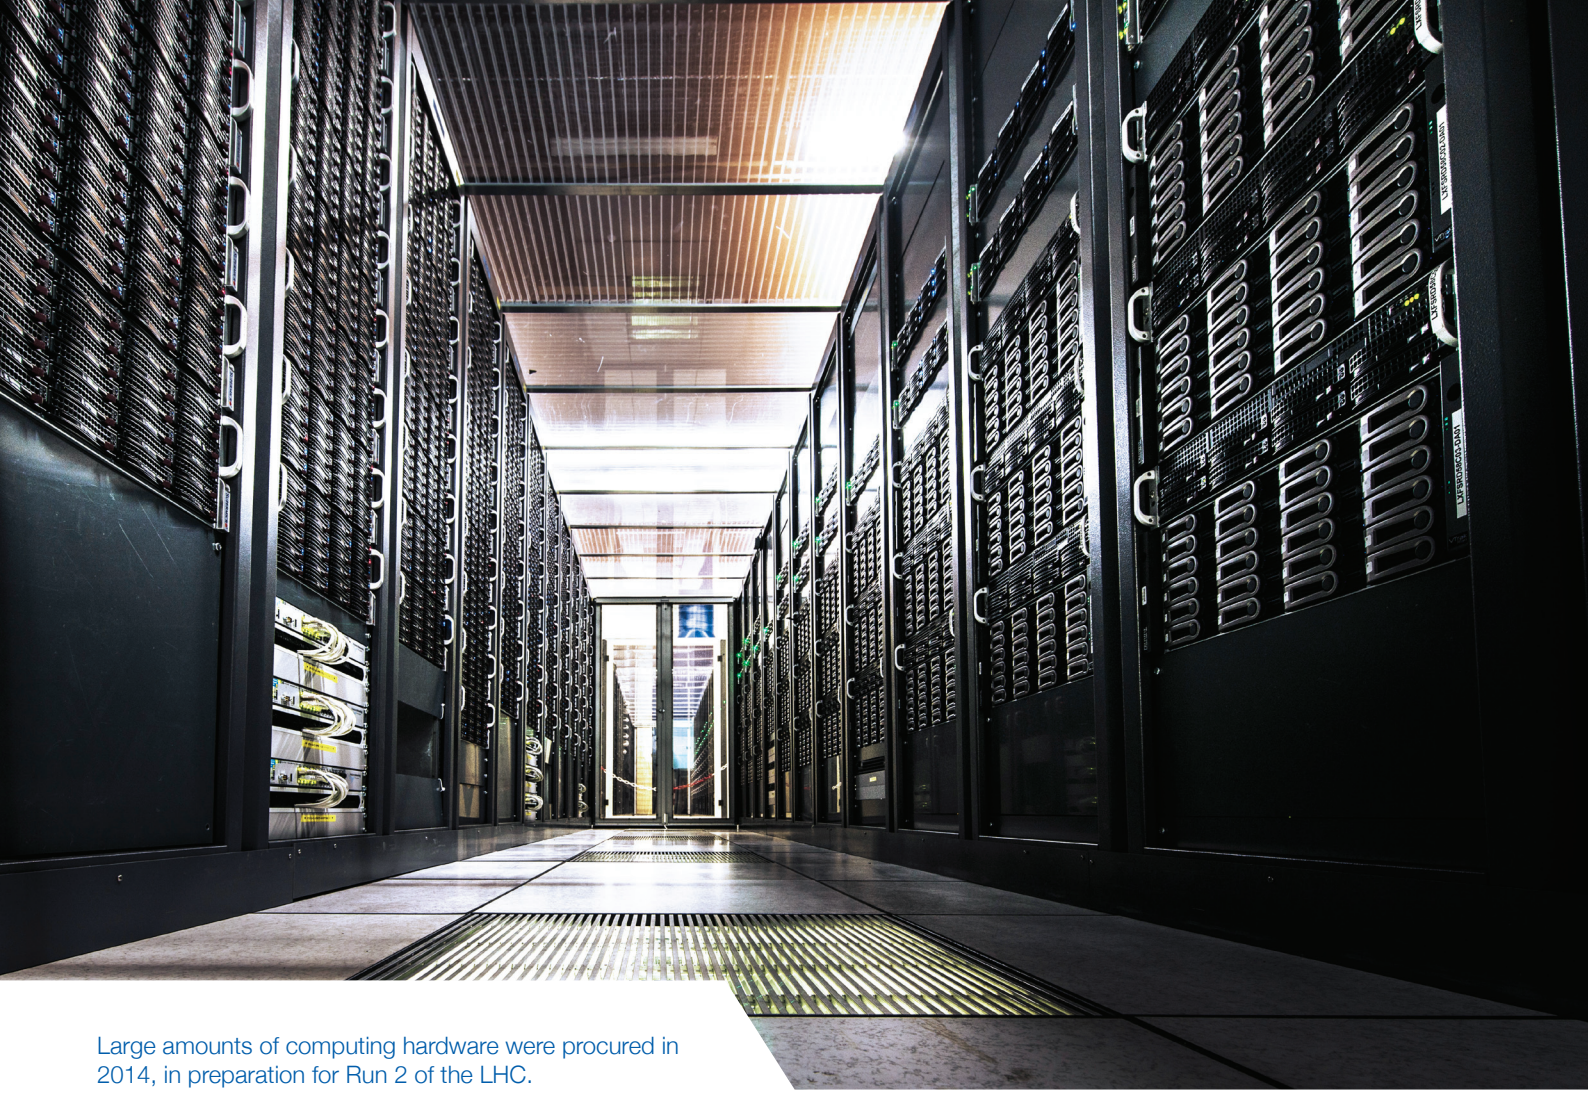

Large amounts of computing hardware were procured in 2014, in preparation for Run 2 of the LHC. (CERN-IT-1209226 – 01)

Computing

New run, new models

To prepare for Run 2, the LHC experiment teams and the Worldwide LHC Computing Grid (WLCG) collaboration have upgraded the computing infrastructure and services during LS1. Large-scale tests in 2014 validated these changes. The experiment collaborations also invested significant effort to improve the performance and efficiency of their core software, with extensive work to validate the new software and frameworks in readiness for the expected increase in data.

To optimize computing and storage resources in Run 2, the experiments have adopted new computing models, in which they move away from the strict hierarchical roles of the Tier-1, -2 and -3 centres described in the original WLCG models, and make more effective use of the capabilities of all sites. This is coupled with significant changes in data-management strategies, away from explicit placement of data-sets globally to a much more dynamic system that replicates data only when necessary. Remote access to data is also now allowed under certain conditions. These 'data federations', which optimize the use of expensive disk space, are possible because of the excellent networking capabilities made available to the WLCG over the past few years.

To improve software performance, a major new long-term activity has been initiated: the HEP Software Foundation. This seeks to address the optimal use of modern CPU architectures and encourage more commonality in key software libraries. The initiative will provide underlying support for the significant reengineering of experiment core software that will be necessary in the coming years.

Increasing computing capacity

The IT Department regularly makes large purchases of computing and storage servers to replace ageing hardware and to increase the capacity available to the LHC experiments. However, 2014

The CERN team receiving the Superuser award at the OpenStack summit in Paris. (OpenStack)

was exceptional as the experiment collaborations had requested a doubling of capacity before the start of Run 2. This meant the procurement of some 100 petabytes (PB) of disk storage and almost 60 000 new cores.

This was a major challenge for the procurement team and for operations teams both at CERN and at the Wigner Data Centre in Hungary, where about half of the computing capacity and two thirds of the storage capacity has been installed. Four large tenders were initiated resulting in eight contracts: three for computing and three for disk-storage servers, one for intelligent power-distribution units and one for a blanket contract to allow for the purchase of more specialized configurations. Additionally, the Communication Systems Group organized a tender for the purchase of the switches necessary to connect all of this equipment to the CERN network.

Virtual machines and storage

OpenStack is an open-source software project that provides on-demand cloud computing. At CERN, the IT OpenStack service enables users to request virtual machines and storage in a few minutes, either through a web portal or from applications. It also provides the computing infrastructure powering most of the grid services and enables efficient management of the increased computing capacity installed for Run 2.

With around 10 000 virtual machines, more than 1000 people use the OpenStack cloud for many different purposes, including production IT services, compute and analysis applications, and personal test and development servers. In addition, the LHC experiments have deployed OpenStack on their high-level trigger farms for use when the accelerator is not running.

At the OpenStack community summit in Paris in November, the OpenStack Foundation presented CERN with the first Superuser award in recognition of the team's contribution to the OpenStack community.

Software-defined storage

Having significant data-storage requirements, CERN has long been a leader in the field of software-defined storage — a technology that enables the creation of open-source alternatives to traditional high-performance network storage appliances. In 2014, the IT Department's Data and Storage Services Group initiated a close collaboration with Inktank, Inc. (later acquired by Red Hat, Inc.) to evaluate their storage solution, Ceph. The initial objective of this investigation was to build a block storage-service for the CERN OpenStack cloud, but it has since expanded to include research and development of Ceph-based solutions to solve future LHC data-storage challenges.

The IT Department deployed and now operates a 3 PB Ceph cluster, one of the largest in the world. CERN's developers have contributed significant new features to this open-source project, including erasure-coding libraries for efficient use of space and an object-striping library for high-performance data-analysis applications.

Data preservation

Given the long lifetime of the LHC experiments — measured in decades — and the significant volumes of data involved, special attention has to be paid to cost-effective data storage and 'bit preservation'. Here, CERN's unique knowledge and experience in large-scale bit preservation, as well as other areas, allows the Organization to make valuable contributions to the global effort to improve long-term preservation of data.

At CERN and the WLCG Tier-1 sites, data continued to be proactively migrated to new generations of storage media on a regular basis, with 'data scrubbing' occurring in parallel. In January, CERN presented a data-curation cost model based on known industry trends and a simplified estimation of LHC data growth. This has attracted a great deal of attention for non-high-energy-physics data-preservation initiatives and is available for

download as an Excel spreadsheet, allowing others to modify the basic parameters and assumptions to suit their needs.

Blue skies ahead for cloud computing

In May 2014, 'Helix Nebula — the Science Cloud Initiative' launched its pioneering cloud-computing marketplace, which offers a first-of-its-kind production service to meet the needs of researchers and facilitate innovation in science. The marketplace delivers easy and large-scale access to a range of cloud-computing services using innovative broker technology. The European Commission's Framework Progamme 7 that supported Helix Nebula finished at the end of 2013. By then, the project had found that the process to procure commercial cloud services is quite different from the existing IT-service procurement models used by many public research organizations. A follow-up Horizon 2020 project was therefore launched in October 2014 with the aim of creating a procurement network of public research organizations. Named PICSE for Procurement Innovation for Cloud Services in Europe, it will investigate the feasibility of joint cross-border pre-commercial procurement and public procurement of innovative services across public organizations.

New times for CERN openlab

CERN openlab is the public-private partnership between CERN and leading IT companies that accelerates the development of cutting-edge solutions for the worldwide LHC community and collaborating research institutes. It completed its fourth three-year phase at the end of 2014. During openlab IV, the CERN openlab partners had addressed topics crucial to CERN's scientific programme, such as cloud computing and storage, data analytics, the next generation of computer processors, and controls for complex engineering systems.

In May, together with a number of European laboratories and leading IT companies, CERN openlab published a whitepaper on future IT challenges in scientific research. This gives a detailed overview of the potential future needs of IT infrastructures supporting a wide range of scientific research fields, and serves as the basis for the fifth phase, openlab V.

As part of its extended support for more scientific research areas, CERN openlab hosted a major workshop on IT in healthcare. In addition, the summer-student programme continued to go from strength to strength, with 23 students of 17 different nationalities coming to CERN.

Getting connected

To provide reliable network coverage for teams working during LS1, almost 300 Wi-Fi base stations were installed in the LHC tunnel early in 2013. As LS1 began to draw to a close nearly 2 years later, the time came for these base stations to be removed — work that had to be carefully choreographed with the cooling of the magnets.

The installation and removal of the Wi-Fi base stations bookended a complete renovation of the networking infrastructure for the LHC and the experiments. Ageing switches were replaced, with nearly 7500 interconnections remade in more than 300 network starpoints. This upgrade and renewal, together with the installation of 28 new fibre-optic trunks and new uninterruptible power supplies, should improve the reliability and redundancy of the network for Run 2 of the LHC. In addition, the relocation of equipment during LS1 led to the installation of more than 340 km of network cables and the complete remodelling of network installations at five of the eight LHC points.

Other work in 2014 involved a rigorous check of 'the leaky feeder' — a special cable that ensures the distribution of mobile telephony signals in the LHC tunnel. These signals include those supporting the TETRA digital radios for the Fire and Rescue Service. LS1 saw the installation of more than 1300 beacons that work with the TETRA system to locate people in the underground areas. This innovative system led to the award to CERN of two prizes by the international TETRA industry forum.