




WHITE PAPER

FUTURE ICT CHALLENGES IN SCIENTIFIC RESEARCH

September 2017





EXECUTIVE SUMMARY
INTRODUCTION
R&D TOPIC 1

R&D TOPIC 2

R&D TOPIC 3

R&D TOPIC 4

ABOUT CERN OPENLAB
ADDITIONAL INFORMATION
CONTRIBUTORS
CONTACTS

04
06
10

DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES

NETWORKING

- High-bandwidth links from detectors to the data centre
- Automation of network configuration and 'white-boxing'
- IoT for FRU tracking, data centre environmental monitoring...
- Integration of Wi-Fi and 5G: data security and protection

DATA-CENTRE ARCHITECTURES

- Rack disaggregation: rack-scale design
- Hierarchical storage buffers
- Software-defined infrastructure and tool-chain integration

DATA STORAGE

- Investigation of models for expansion of storage-capacity
- 'Cold storage' evolution

DATABASE TECHNOLOGIES

- Data size and rates
- Hardware evolution and consolidation
- Technologies for developer productivity
- Data-stream processing
- Time-series database workloads
- Scale-out databases and cloud resources

CLOUD INFRASTRUCTURES

- Orchestration and automation of compute provisioning
- Scalable clouds and global scientific clouds

20 COMPUTING PERFORMANCE AND SOFTWARE

CODE MODERNISATION

- Storage-layer optimisations for low-latency NVRAM
- Performance-analysis tools for software
- Use of standard library facilities in C++
- Use of alternative concurrency models
- Verifying code and checking quality automatically

HETEROGENEOUS PLATFORMS AND ALTERNATIVE ARCHITECTURES

- Hybrids CPUs
- Optimising code distribution using lightweight containers

DEDICATED HARDWARE AND CO-PROCESSING SYSTEMS

- Optimising code performance using coprocessors and GPUs

26 MACHINE LEARNING AND DATA ANALYTICS

DATA ACQUISITION

- Monitoring of accelerators and detectors
- Monitoring data quality
- Fast inference technology for 'trigger' systems
- Anomaly detection and the search for new physics

DATA PROCESSING

- Simulation
- Jet identification and image-based event identification

BIG DATA

- Data reduction and refresh for analysis
- Optimisation of computing infrastructure

DATA ENGINEERING

- Solutions from industry, challenges and opportunities

34 APPLICATIONS IN OTHER DISCIPLINES

PLATFORMS FOR OPEN COLLABORATION

- A smart data-analysis platform

LIFE SCIENCES AND MEDICAL APPLICATIONS

- Simulating biological systems in the cloud
- Large-scale analysis of genomic data
- Large-scale analysis of healthcare data

ASTROPHYSICS

- Exascale data processing at future astrophysics infrastructures

SMART EVERYTHING

- Environmental monitoring
- Traffic and mobility

40
42
43
44

EXECUTIVE SUMMARY



CERN openlab is a unique public-private partnership. It was established in 2001 to provide a framework through which CERN can collaborate with leading ICT companies to accelerate the development of the cutting-edge ICT solutions needed by the high-energy physics community.

CERN is home to the Large Hadron Collider (LHC), the world's most powerful particle accelerator. The complexity of the scientific instruments at the laboratory throw up extreme ICT challenges, and make it an ideal environment for carrying out joint R&D projects and testing.

Throughout its 16-year history, CERN openlab has worked to develop and test new ICT technologies and techniques that help to make the ground-breaking physics discoveries at CERN possible. CERN openlab runs in three-year phases, with around 20 projects — addressing a wide range of IT topics — being run in its current, fifth phase.

With CERN openlab's sixth three-year phase set to begin at the start of 2018, work has been carried out throughout the first half of 2017 to identify key areas for future collaboration.

A series of workshops and discussions was held to discuss the ICT challenges faced by the LHC research community — and other 'big science' projects — over the coming years. This white paper is the culmination of these investigations, and sets out specific challenges that are ripe for tackling through collaborative R&D projects.

With the LHC and the experiments set to undergo major upgrade work in 2019 and 2020, CERN openlab's sixth phase offers a clear opportunity to develop ICT solutions that will make a tangible difference for researchers when the upgraded LHC and experiments come back online in 2021.

This white paper identifies 16 ICT challenge areas, which have been grouped into four different R&D topics (further details on the next page). Tackling these challenges — through a public-private partnership that brings together leading experts from each of these spheres — has the potential to positively impact on a range of scientific and technological fields, as well as wider society. For each challenge area in this white paper, specific use cases have been identified.



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

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

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

R&D TOPIC 2: COMPUTING PERFORMANCE AND SOFTWARE

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

R&D TOPIC 3: MACHINE LEARNING AND DATA ANALYTICS

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

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

R&D TOPIC 4: APPLICATIONS IN OTHER DISCIPLINES

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

CERN openlab's mission rests on three pillars: technological investigation, education, and dissemination. Collaborating with research communities and laboratories outside the high-energy physics community brings together all these aspects. Challenges related to the life sciences, medicine, astrophysics, and urban/environmental planning are all covered in this section, as well as scientific platforms designed to foster open collaboration.



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

INTRODUCTION



This chapter provides a brief introduction to CERN, the European Organization for Nuclear Research. It gives an overview of the research infrastructures in place at the laboratory, as well highlighting plans for upgrades to these infrastructures. Future upgrades to the Large Hadron Collider (LHC) and the experiments on it will provide enormous scientific value, but will also pose major ICT challenges.

In this white paper, CERN openlab sets out challenges to tackle together through joint R&D projects with our industry collaborators over the coming years. This unique public-private partnership between research and leading ICT companies is ideally placed to tackle these challenges, producing results that are of benefit to all parties, and of interest to the wider scientific community beyond high-energy physics.

ABOUT US

CERN

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

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

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

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

CERN OPENLAB

CERN openlab is a unique public-private partnership that accelerates the development of cutting-edge ICT solutions for the worldwide LHC community and wider scientific research. Through CERN openlab, which was established by the CERN IT department in 2001, the laboratory collaborates with leading ICT companies and other research institutes.

Within this 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 ICT industry partners 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.

Since its foundation in 2001, CERN openlab has run in successive three-year phases. The collaboration's fifth phase (2015-2017) is now coming to a close. This phase tackles ambitious challenges covering the most critical needs of IT infrastructures in domains such as: data acquisition, computing platforms, data storage architectures, compute provisioning and management, networks and communication, and data analytics.

In 2018, CERN openlab begins its sixth three-year phase. This white paper identifies ICT challenges (primarily related to the LHC research programme, but also in other fields) that are ripe for tackling as part of our joint work over this period.

THE ACCELERATOR COMPLEX AND THE DETECTORS

THE LHC

The accelerator complex at CERN is a succession of machines that accelerate particles to increasingly high energy levels. Each machine boosts the energy of a beam of particles, before injecting the beam into the next machine in the sequence. In the Large Hadron Collider (LHC) — the last element in this chain — particle beams are accelerated up to the record energy of 6.5 TeV (tera-electron-Volts) per beam. Most of the other accelerators in the chain have their own experimental halls where beams are used for experiments at lower energies.

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

The protons are finally transferred to the two beam pipes of the LHC. Inside the accelerator, two high-energy particle beams travel at close to the speed of light before they are made to collide.

The beams travel in opposite directions in separate beam pipes — two tubes kept at an ultra-high vacuum. They are guided around the accelerator ring by a strong magnetic field maintained by superconducting electromagnets.

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

Protons are not the only particles accelerated in the LHC. Lead ions for the LHC start from a source of vaporised lead and enter Linac 3 before being collected and accelerated in the Low Energy Ion Ring (LEIR). They then follow the same route to maximum energy as the protons.

THE EXPERIMENTS

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

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

phenomena. These four detectors sit underground in huge caverns on the LHC ring.

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

It is important to note that while the main focus of research at CERN has moved in recent years towards the LHC, experiments at other accelerators and facilities both on-site and off also remain an important part of the laboratory's activities.

COLLISIONS AND 'TRIGGER' SYSTEMS

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

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



COMPUTING AT CERN

THE CERN DATA CENTRE

Even once the data coming from the LHC experiments has been reduced by the trigger systems, the CERN data centre processes hundreds of petabytes of data annually. The CERN data centre in Meyrin is the heart of CERN's entire scientific, administrative, and computing infrastructure. All services, including email, scientific data management and videoconferencing use equipment based in the data centre. The centre hosts around 10,000 dual-CPU servers with approximately 300,000 processor cores.

The servers undergo continual maintenance and upgrades to make sure that they will operate in the event of a serious incident. Critical servers are held in their own room, powered and cooled by dedicated equipment.

Around 150 PB of data are stored on disk at the Meyrin site, with over 200 PB on tape.

THE WIGNER RESEARCH CENTRE FOR PHYSICS

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

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


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

The CERN and Wigner centres are connected via three independent and dedicated 100 Gb/s fibre-optic lines, with a bandwidth equivalent to the transfer of eight full DVDs per second. Network latency (the time taken between sending data and receiving on the other end) between the two sites, which are over a 1000 km apart, is about 25 milliseconds.

THE WORLDWIDE LHC COMPUTING GRID (WLCG)

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

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



THE LHC UPGRADE PROGRAMME

RUN 4

The LHC has been designed to follow a carefully set out programme of upgrades. The LHC typically produces particle collisions for a period of around three years (known as a 'run'), followed by a period of about two years for upgrade and maintenance work (known as a 'long shutdown', which is commonly abbreviated to LS). This is illustrated below.

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

RUN 3

By Run 3, which is planned to start in 2021, the ALICE and LHCb experiments will markedly increase their trigger rates, meaning that much less data will be filtered out. Collecting many more collision events — and therefore much more data — than ever before means that the events will have to be reconstructed (turned from raw sensor data to representations of the particle tracks crossing the detectors) shortly after being collected, rather than first being stored and then processed 'offline'.

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

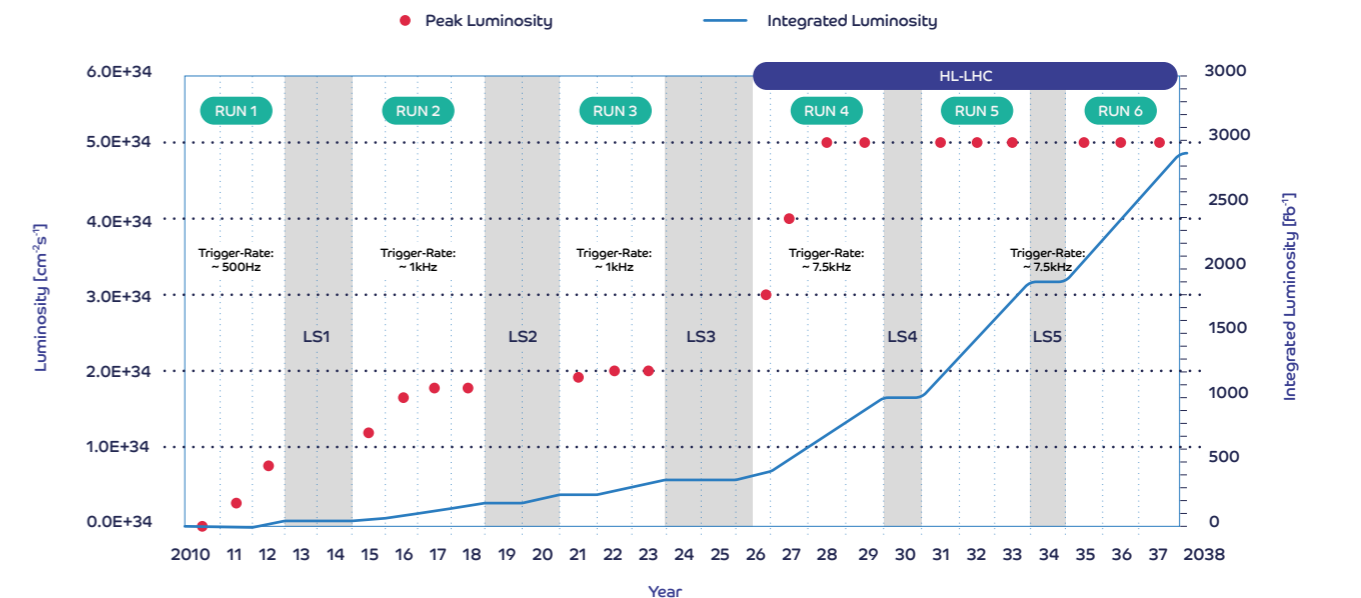
The High-Luminosity Large Hadron Collider (HL-LHC) project aims to crank up the performance of the LHC in order to increase the potential for discoveries starting in Run 4. The objective is to increase integrated luminosity by a factor of ten beyond the LHC's design value.

Luminosity is an important indicator of the performance of an accelerator: it is proportional to the number of collisions that occur in a given amount of time. The higher the luminosity, the more data the experiments can gather. This enables physicists to observe rare processes and study new particles with greater precision.

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

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

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





R&D TOPIC 1

DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES

Designing and operating distributed data infrastructures and computing centres poses many challenges. These challenges are amplified and added to when operating at the extremely large scales required by major scientific endeavours (sometimes even at the global scale).

Overall, the computing needs of the LHC experiments are growing. In addition to increasing the capacity of the systems used for traditional types of data processing and storage, explorations are being carried out into a number of alternative architectures and specialised capabilities. These will add heterogeneity and flexibility to the data centres, and should enable advances in resource optimisation (improved resource provisioning, data processing, accounting, monitoring, etc.).

In order to accommodate the expected needs of the LHC experiments over the next decade — with ever more particle-collision events being reconstructed in quasi-real-time using just-in-time calibration — CERN must evaluate different models for increasing its computing and data-storage capacity. This evaluation needs to cover a range of scenarios, from federating multiple data centres under a common management and access model, to procuring resources from commercial providers, to building new computing centres on CERN's premises. All models present different technological challenges. For example, the construction of a new data centre at CERN would require network connections capable of transporting multiple terabits of data per second, and would involve ambitious goals in terms of energy efficiency.



NETWORKING

Modern 'big science' requires continuing developments to be made in nearly all aspects of networking; these will affect the design and implementation of future computing centres. Wide-area networks are following a road map of increasing performance, with 400-Gb/s connections expected to become available for use in production networks in the next few years. Rapid increases in network speeds will enable data centres to become more interconnected than ever before, thus improving the efficiency of proactive data sharing and resource provisioning.

Developments in local inter-processor communication are also changing how clustered computers can work together, blurring the traditional boundaries between high-throughput computing (HTC) and high-performance computing (HPC). Low-latency interconnects are being introduced to server CPUs, and high-performance low-latency interconnect cards are coming to the market. This expands the possibilities for independent nodes to work on the same operations, and has implications for how scientific software should be written in the future. These technologies also have the potential to expand the capabilities of data-

acquisition farms: increased use of off-the-shelf resources could make them more easily shareable and improve operability.

Devices in data centres and scientific instruments of all sizes — as well as many everyday objects — are getting connected. This enables them to share information about their state, efficiency, energy consumption, and more, thus leading to increased efficiency. This is underpinned by high-speed networks capable of carrying vast amounts of information, with in-built provisioning of data security and protection.

The globalisation of science means that research centres are often becoming community hubs for worldwide collaboration and the exchange of knowledge and expertise. This is driving a range of new needs. Local campus infrastructures must become more reliable and secure for internal and external users reliant on wireless mobile communication devices. Equally, remote collaboration needs to become simpler and more efficient, lowering barriers related to language and increasing participation from many different regions of the world.



USE CASE 1.1.1

HIGH-BANDWIDTH LINKS FROM DETECTORS TO THE DATA CENTRE

CERN is currently investigating the possibility of building a data centre that would consolidate the computing resources required by the online data-acquisition systems of all the major CERN experiments.

Today, the LHC experiments' detectors produce an aggregated data rate of about 1 PB/s before any filtering. A multi-stage sequence of triggers and filters is then applied in dedicated processing centres to reduce the data to a manageable size. The remaining data is finally transferred into the main CERN data centre for further processing.

Consolidating this 'online processing' into a single data centre would therefore pose significant challenges in terms of high-speed connectivity. Among the different possible strategies and technologies applicable to this problem, CERN is interested in investigating DWDM (Dense Wavelength Division Multiplexing) technologies to enable the high-speed transmission (in the order of tens of Tb/s) of data over a single pair of fibres.

USE CASE 1.1.2

AUTOMATION OF NETWORK CONFIGURATION AND 'WHITE-BOXING'

Over the last 20 years, CERN has developed a framework that enables the automatic configuration of its networks. The network topology and its services — as well as the connected devices — are modelled in a structured-query-language (SQL) database.

Based on the information in this database, a suite of tools performs the automatic configuration of network devices (routers, switches, firewalls, etc.). The degree of automation varies, and some of the critical tasks require human supervision/confirmation before configuration changes are deployed in production network equipment.

CERN is investigating alternatives for modernising its network-management platform. The new system should provide additional, specialised capabilities, such as the generation of the desired configuration of a network device, 'zero-touch' provisioning starting from the default factory configuration, alignment of previously configured devices with a new configuration, tracking of the configuration history for all devices, etc.

Ongoing trends in terms of software-defined networking (SDN) are also of interest. CERN sees potential in investigating moves towards further separation of the hardware used for the forwarding plane from the data-plane logic and operating system, as based on the use of 'white-box' devices (i.e. 'bare metal' devices that can run a third-party operating system). If a single operating system could be used on a range of 'commodity' white-box switches, the network-management tools would have a common interface, regardless of the physical hardware deployed in the network.

USE CASE 1.1.3

IoT FOR TRACKING FIELD-REPLACEABLE UNITS (FRUs), MONITORING THE DATA-CENTRE ENVIRONMENT, AND PROVIDING 'SMART CAMPUS' SERVICES

As the use of connected devices and sensors rapidly increases to form what is referred to as 'the internet of things (IoT)', a wide range of potential, specialised uses at CERN is starting to emerge. Information collected from instruments and appliances scattered across technical facilities and office buildings could be used to improve energy efficiency, optimise processes, and reduce costs.

The data generated by a network of thousands (or even millions) of sensors — connected either to distributed or centralised processing units — has a considerably different profile to the data generated by the interaction between humans and machines, or that generated by the LHC experiments. The associated requirements, in terms of real-time processing, storage, security, and networking, might therefore have a non-negligible impact on the CERN data centre and other parts of the laboratory's infrastructure.

Consequently, CERN sees value in investigating various possible network architecture models to support IoT applications in the laboratory. A number of specific uses are now being considered, such as optimising computing provisioning, developing 'smart building' systems to improve environmental efficiency, using on-site road traffic information to plan proactive modifications, and creating other 'intelligent campus' services (e.g. for parking management).

USE CASE 1.1.4

INTEGRATION OF WI-FI AND 5G: DATA SECURITY AND PROTECTION

Today's mobile communication services across the laboratory campus rely on a combination of Wi-Fi and 4G-based telephony services. These two networks are completely separate and do not allow seamless roaming of devices from one to the other without disconnection and a change of network address. Additionally, connections over 4G rely on transmission of voice and data back to home networks, even when roaming across different networks.

In addition to the increased data rates and number of connections it makes possible, 5G offers several interesting new features that CERN would like to explore further. The possibility of connecting simultaneously to different wireless-access technologies would enable seamless roaming across the campus; machine-type communications (MTCs) would natively support campus-wide IoT applications; and the ability to create virtual RAN (Radio Access Networks) would enable traffic to be kept local, thus improving data security and protection.

CHALLENGE AREA 1.2

DATA-CENTRE ARCHITECTURES

The CERN data centre has undergone many transformations over the years. Today, the main goal is to operate an efficient facility that is able to host a diverse set of equipment. However, the evolving computing and storage requirements of CERN and the LHC experiments will drive many changes in the coming years: for example, it is becoming increasingly important to be able to rapidly adapt the infrastructure to diverse workflow types, intelligently optimise resource use, and strike an appropriate balance between internally and externally hosted resources.

From both a strategic and an operational point of view, CERN has an interest in continuing to operate its own data centre. The increasing requirements in terms of storage capacity and computing resources can be addressed in different, non-exclusive ways: for example, by installing additional capacity at CERN, installing additional capacity off-site (as currently happens with the Wigner facility near Budapest), or by procuring additional capacity from commercial cloud providers (and integrating the resources in the existing processes and tool-chains).

Should CERN decide to build a new data centre, ambitious efficiency goals will have to be set. These include a ‘power usage effectiveness’ (PUE) ratio of less than 1.1, a relatively high power density of 15-20 kW per rack, and a modular design that enables staged power growth over time (from an initial 4 MW to 12-16 MW later). The future growth-rate and replacement-cycle needs would also have to be re-evaluated based on the expected technological evolution. In addition, the trend towards increased use of specialised hardware to run computationally intensive tasks is likely to result in any potential future data centre being required to accommodate a heterogeneous range of equipment. In this case, the expertise and skills to support these diverse resources will also have to be developed.



USE CASE 1.2.1

RACK DISAGGREGATION: RACK-SCALE DESIGN

Utilisation of the data centre changes over time, depending on the specific requirements of different research projects, which often vary significantly. A weakness in the architecture of many of today’s data centres is that they cannot easily accommodate rapid variations of cycles and workloads, or accommodate peak utilisation without having significant periods of underutilisation. These issues are relevant today for the operational models of the LHC experiments.

The possible applications and benefits of rack disaggregation are currently being investigated. This holds promise for rapidly allocating the correct amount of storage and computing resources (as well as possibly memory and bandwidth) to incoming requests. This would help to minimise underutilisation, thus reducing costs and promoting vendor-independence.

USE CASE 1.2.2

HIERARCHICAL STORAGE BUFFERS

The future computing models of the LHC experiments will require much faster processing of data. CERN is therefore interested in the potential of solutions based on rack-scale design for storage buffers capable of handling hundreds of petabytes of data. These would ideally have a hierarchical structure, including volatile/non-volatile memory and high-bandwidth solid-state storage drives (possibly as well as high-capacity solid-state storage drives and/or spinning disks). These would all need to be interconnected with an efficient fabric, such as iSCSI (Internet Small Computer Systems Interface), NVMe-oF™ (NVMe Express™ over Fabrics), or even a proprietary solution that supports industry-standard interfaces.

USE CASE 1.2.3

SOFTWARE-DEFINED INFRASTRUCTURE AND TOOL-CHAIN INTEGRATION

The potential benefits of rack disaggregation can only be fully realised if efficient configuration and orchestration tools are available. As part of the ongoing interest in evaluating disaggregated rack systems, CERN is interested in assessing the functionality and interoperability of software-defined infrastructure (SDI). The goal is to better understand the potential efficiency improvements and cost savings offered by increased automation of data-centre operations. An important aspect to investigate is the potential impact on the existing procedures and tools in place. The definition of a practical roadmap for deployment and integration of new automation tools in the existing tool-chain would also be required.

CHALLENGE AREA 1.3

DATA STORAGE

By 2025 (when the HL-LHC is planned to begin operation), it is expected that the amount of data produced by the experiments will have increased by between one and two orders of magnitude. It will be necessary to significantly increase the storage volume available within existing budget constraints.

The elevated data-processing requirements are expected to generate dramatic increases in the input/output (I/O) needs. Consequently, more performant storage solutions — beyond what is possible today even with state-of-the-art SSD disks or flash memories — will be necessary.

Until now, all the LHC experiments have relied on traditional, hierarchical mass-storage systems. These are primarily based on spinning disks, with tape systems as the final long-term archival layer. In recent years, all experiments have added data-federation capabilities, which enables files to be discovered and streamed to processing resources. This continues, however, to constitute a small fraction of data delivery.

The consolidation happening in the high-end tape market, with fewer manufacturers now playing a role, may also prompt a possible change of strategy. Such a potential move would also be reinforced by the continuing trend observed towards tape-based storage systems being treated primarily as archival resources used for disaster recovery, rather than as active archival systems. The decreasing costs and increasing longevity of solid-state devices are also driving a migration from a hierarchy based on spinning disks and tapes to one based on different media with multiple speeds and quality of service.

USE CASE 1.3.1

INVESTIGATION OF MODELS FOR EXPANSION OF STORAGE-CAPACITY

The CERN IT department is interested in evaluating different strategies for capacity increase — based on a range of scenarios with varying requirements in terms of total storage volume — and assessing the associated TCO (total cost of ownership). The strategies to be evaluated include the following: capacity extension through the use of storage media directly attached to networked servers, the deployment of a distributed file system directly accessible from the computing nodes, and increasing the elasticity of the system through ‘thin provisioning’ (entirely integrated within the existing automated tool-chain).

USE CASE 1.3.2

‘COLD STORAGE’ EVOLUTION

The need to reassess CERN’s current strategy for long-term archival and cold storage (triggered primarily by consolidation in the tape market, as outlined above), is driving interest in investigating alternative solutions. One possibility being investigated is the use of low-cost flash memory — coupled with increased storage bandwidth and connectivity — to reduce the amount of online storage needed for short-term analysis. This would be supported with an agile infrastructure that enables data to be efficiently moved to the most suitable repository at any given time.



DATABASE TECHNOLOGIES

Traditional database applications will continue to play a key role in the HL-LHC era: it will be necessary to store and process accelerator controls logs, as well as information on experiment conditions, detector calibrations, etc. Challenges are expected to be thrown up by the increased use of database applications to index events and the use of analysis applications based on relational databases and NoSQL technologies.

Databases for the LHC, the experiments, and the associated computing workloads require real-time and batch ingestion at high rates of throughput. The main challenges and opportunities can be classified into these main areas: streams processing, consolidation, cloud resource usage, data size and rates, hardware evolution, scale-out databases, data analytics, and machine learning (together, these last two areas form R&D topic 3 of this whitepaper). Solving these challenges would help to boost developer productivity.

USE CASE 1.4.1

DATA SIZE AND RATES

In a broad range of application areas, data is already being collected at an unprecedented scale. The high ingest rate of data is a major challenge: time-series workloads are expected to keep playing an important role, as billions of individual data points are produced. High write- and read-throughputs will be required by the many millions of sensors at CERN. The logging of LHC control information, for instance, is expected to generate up to 1 TB per day by the beginning of Run 3, and is expected to grow further beyond this. The high-frequency queries in the LHC experiments will be also particularly challenging: the file catalogue for the ALICE experiment currently has a read access rate of 11 kHz and is expected to increase by a factor of ten during Run 3.

USE CASE 1.4.2

HARDWARE EVOLUTION AND CONSOLIDATION

The evolution of hardware components — in particular, new types of memory — is expected to have considerable impact on database applications; it presents opportunities to solve a range of challenges currently faced. However, in order to profit fully from these new architectures, adaptations will have to be made.

CERN is interested in exploring the use of high-bandwidth, low-latency non-volatile memory (NVRAM) solutions — like the recently announced 3D XPoint technology — for very large in-memory

databases and data-analytics workloads. In terms of physics research, this includes applications for low-latency event indexing. The challenges to be faced are algorithmic complexity, the need for low-level programming skills, and the deployment and operation of heterogeneous platforms within the existing infrastructure and tool-chains.

Another aspect of the evolution in hardware is that servers are becoming increasingly capable of handling loads for several applications, thus making them well suited for consolidation efforts. This is, of course, beneficial in terms of reducing cost. It does, however, present additional challenges typical of multi-tenant systems, such as those related to resource management, isolation of workloads, and security.

USE CASE 1.4.3

TECHNOLOGIES FOR DEVELOPER PRODUCTIVITY

In addition to the challenges related to infrastructure and hardware, demands from the developer community have increased dramatically. The modern developer expects fast deployments, simple debugging, and support for ‘continuous integration’ development practices.

At CERN there are over 19,000 code repositories used by over 9000 developers. It is therefore difficult to find a platform suitable for running all these applications, and a balance has to be carefully struck between speeding up ‘time to production’ and system reliability. Technologies like Docker and Kubernetes can play an important role.

CERN, like many other organisations, is investigating the use of public clouds. This throws up further challenges for application developers. It will be important to thoroughly test and carefully select the right technologies for building hybrid cloud systems.

USE CASE 1.4.4

DATA-STREAM PROCESSING

As with many other large organisations running scientific or industrial infrastructures, a huge volume of data is created each day from sensors and logging systems at CERN. Data streams play a vital role in the control systems for the LHC, in the monitoring of the computing infrastructure, and in many other areas. As the laboratory adopts the IoT paradigm more closely, this volume of data is increasing rapidly. Each data stream requires sophisticated management infrastructures and needs to be properly processed to gain strategic benefits.

CERN has been successfully managing, processing and analysing its ‘big data’ for many years. Nevertheless, it is not easy to do so ‘online’, with low latency (almost real time), in a scalable way, with fault tolerance, and with the possibility to ingest machine-learning approaches. This represents an important challenge for CERN’s engineering teams across the laboratory.

The CERN Accelerator Logging Service — used to store logging data related to the complete accelerator complex at CERN — is currently being redesigned. The main objective is to seamlessly combine live data streams with persistent data sources. The accelerator control data is generated by approximately 1 million pre-defined signals, coming from a heterogeneous range of sources. This currently produces about 2 TB of data per week, and is increasing rapidly. The engineers in charge of this system need to handle this torrent of information in such a way that makes it possible to access the data in an efficient and transparent manner, as well as unlocking online machine-learning analytics capabilities.

In another domain, CERN’s computing security team is building a large-scale security operations center. Management and analysis of a variety of data-streams will be used for the detection, containment, and remediation of threats. At CERN, there are close to half a million ‘indicators of compromise’, and there are more than 160,000 more on the WLCG.

The LHCb experiment is also building a data-streaming system to monitor and analyse research-data workflows in real time. This needs to handle all the distributed computing activities (several hundred thousand jobs per week generated by thousands of researchers across the globe), such as data processing (reconstruction, stripping, and streaming), Monte-Carlo simulation, and data replication.

Despite the three examples outlined above coming from different fields, they all face common challenges: (i) the need to define unified, centralised systems for the management and processing of data streams; (ii) finding ways to enable transparent data access and to empower data pipelines; (iii) identifying machine-learning models and establishing new deployment processes that can be applied to data streams.

USE CASE 1.4.5

TIME-SERIES DATABASE WORKLOADS

At CERN, time-series data has traditionally been stored and processed in the same platforms as the rest of the related services- and applications-data (namely, relational databases). The use of specialised databases for time-series data is, however, growing in popularity. These provide easier analysis, improved storage- and ingestion-rates, as well as other benefits. There is a growing interest in such databases from teams across the laboratory.

The systems-monitoring data for the CERN data centre is currently being migrated to a multi-language database environment, whereby diverse data sources (such as logs or systems metrics) are stored, aggregated, and analysed in several specialised sys-

tems. Today, system metrics for approximately 30,000 servers are stored in several time-series databases, totalling 15,000-30,000 writes per second and spread over approximately 5 million series. With the use of retention policies and down-sampling, the whole dataset is currently kept limited to around a few terabytes.

Another example from the IT department is the migration of the monitoring system for the ‘database on demand’ — currently based on a third-party tool and a relational database — to a time-series database. This system supports the collection of hundreds of service metrics from around 500 different database systems, and currently sees around 8000 points written per second on 150,000 series. The challenge for this application would be to support high numbers of series (several million), in order to support performance metrics down to the level of individual tables and indexes.

There is also interest from engineering teams at CERN in using time-series databases to store and analyse digital signals. This would involve being able to sustain frequency insertions as high as 1 MHz, at least during limited-duration peaks. It would also be necessary to have a high degree of time precision (up to nanoseconds), while keeping the system simple to use for both ingestion and consumption of the data.

USE CASE 1.4.6

SCALE-OUT DATABASES AND CLOUD RESOURCES

As the complexity of infrastructures and workloads increases, ensuring scale-out capability is a major challenge. For example, teams at CERN foresee having to run database workloads on hundreds — or maybe even thousands — of nodes. Analytics workloads like Spark, which are tied to databases, might span over even larger number of machines.

Resource management is becoming increasingly complex: it is important to ensure consistent availability and reproducibility of services despite the expected failure of individual resources. A systematic approach to disaster recovery and backup capabilities is also required.

New approaches are already being worked on, such as a ‘platform-as-a-service’ (PaaS) cloud-based architectures for services like transactional databases and scale-out databases. The increasing use of clouds for both databases and data-analytics services must be evaluated in terms of impact on users and service managers (particularly in terms of its effects on data-transfer efficiency, performance, and cost).

CHALLENGE AREA 1.5


CLOUD INFRASTRUCTURES

More and more resources have been moved into CERN's private cloud infrastructure — run using OpenStack — over the past several years (see graph on opposite page). Today, more than 90% of CERN's compute resources are virtualised, with tens of thousands of new cores being added each year.

Future infrastructure usage — as depicted by the models for the HL-LHC provided by the experiments — may see computing capacity increase by a factor of 50 to 100 compared to today, as well as the addition of new types of heterogeneous computing platforms. Achieving this level of scalability requires extensive investigation of both a range of cloud technologies and usage policies.

There are ongoing investigations into simplifying the user experience by providing virtual machines, containers, and bare-metal nodes via the same user interfaces and administrative systems. At the same time, the goal is to enable a wider range of potential services, such as 'file-system-as-a-service', 'secrets-as-a-service', bare-metal management, a single inventory for all computing resources, etc. It is important to simplify the user experience yet further, while providing users with more flexibility in selecting the right combination of hardware resources for their activities (as alluded to in use case 1.2.1).

Total cores in IT OpenStack environment at CERN



USE CASE 1.5.1

ORCHESTRATION AND AUTOMATION OF COMPUTE PROVISIONING

CERN's private cloud currently offers more than 20 different hardware configurations. It is important to offer users a simpler experience for selecting the most appropriate resource configuration for their needs, without the overhead of maintaining a growing number of predefined sets. In this context, there is an interest in investigating how the emerging concepts of rack disaggregation and software-defined infrastructure could be used — behind the existing OpenStack provisioning and user interfaces — to enable this new level of service.

USE CASE 1.5.2

SCALABLE CLOUDS AND GLOBAL SCIENTIFIC CLOUDS

Budgets for large-scale scientific endeavours are under significant pressure in many regions across the globe, thus making it vital to improve efficiency in all areas — including the underpinning ICT systems. The HL-LHC requires the deployment of vastly increased computing capacity over the next decade (as outlined in the introduction of this challenge area). Several procedural and technical improvements are therefore being investigated to expand the CERN's private cloud within the existing constraints. These include OpenStack Cells, software-defined networking (OpenStack Neutron), better capacity planning, utilisation optimisation through more efficient accounting and monitoring, and the application of machine-learning techniques to understand and optimise the infrastructure for the most common usage patterns.

The deployment of global scientific platforms is also a topic of great interest. It has significant potential impact for LHC computing models, as well as for other major scientific communities.

R&D TOPIC 2



COMPUTING PERFORMANCE AND SOFTWARE

Code modernisation is the most promising avenue of investigation for closing the computing 'resource gap' for the HL-LHC era. Software performance can be increased significantly by making better use of the hardware features offered by modern CPUs, GPUs, and coprocessors.

The LHC experiments already have an established record of successfully improving code. As luminosity has increased during the first two runs of the LHC, various code has been optimised to run several times as quickly — without compromising physics performance. There are even examples where code modernisation efforts have led, within a single year, to software running two or three times faster. Additionally, there are many examples of new hardware architectures and coprocessors where large improvements have been observed for properly optimised code.



CODE MODERNISATION


Code used in high-energy physics is typically written by a large number of people — with a variety of skill levels — over long periods of time. The lifetime of the code is usually much longer than the hardware ‘refresh time’, with contributors to the code also changing over time. The physics performance of the code is paramount; optimising the code in terms of speed, memory usage, and I/O is often considered a secondary priority.

The last several years have seen significant changes in the traditional server architecture used in the WLCG. There are many more cores per CPU, and the width of the vector registers is also increasing. The memory per core has remained roughly constant on machines that were bought for the purpose of supporting LHC, but the memory and disk space per core is lower on some of the largest shared resources to which we have access.

Both the growing core count per CPU and the desire to take full advantage of HPC centres is driving a push towards more

parallel code. The parallelised code progresses at event level, with multiple events being processed simultaneously; there is some gain in memory because common read-only elements are only stored once. This has been the easiest form of parallelisation to implement. In contrast, development has taken much longer for the thread-safe reconstruction algorithms: writing modern, highly parallel code requires significant expertise and a high level of software-engineering skills.

The increase in the width of the vector registers — and the ability to execute more instructions in parallel — could potentially lead to significant gains in software performance. However, careful optimisation and tuning of the code to individual hardware platforms is necessary. Analysing and improving the efficiency of the code, as well as ensuring the code makes use of multiple cores simultaneously, would be a large step toward closing the resource gap alluded to in the introduction of this chapter.



USE CASE 2.1.1

STORAGE-LAYER OPTIMISATIONS FOR LOW-LATENCY NVRAM

Software uses caching techniques to reduce differences in latencies between different system components (e.g. disk, RAM, CPU). With the advent of very-low-latency and high-bandwidth NVRAM technology, the caching from disk to RAM becomes superfluous. To make good use of this new NVRAM storage medium, the key storage components in the code need to be redesigned, possibly by removing caching layers and changing data layouts. Early access to special NVRAM support libraries would be very useful.

USE CASE 2.1.2

PERFORMANCE-ANALYSIS TOOLS FOR SOFTWARE

There is clear potential for the code used in high-energy physics to be modified to make good use of recent advances in CPU technologies. However, the sheer size of the code base can often make modernisation difficult. It is therefore important to make full use of available development tools, to ensure code is optimised as efficiently as possible.

USE CASE 2.1.3

USE OF STANDARD LIBRARY FACILITIES IN C++

Where possible, ‘home-grown’ code should be replaced by centrally maintained, highly optimised library code, such as C++ standard library facilities. These are highly optimised and have been very well debugged, and are maintained by the compiler providers, which can lead to more robust scientific software.

USE CASE 2.1.4

USE OF ALTERNATIVE CONCURRENCY MODELS

Classical, task-based concurrency has enabled substantial performance improvements for many applications used in high-energy physics. Nevertheless, it is important to continue to pay attention to new developments in concurrency, which enable parallel semantics to be expressed more efficiently.

New advances in coroutines are of particular interest: these generalise subroutines for non-preemptive multitasking, by enabling multiple entry points for suspending and resuming execution at certain locations.

High Performance ParalleX (HPX) and ReactiveX are also of interest. HPX is a general-purpose C++ runtime system for parallel and distributed applications of any scale, which offers potential benefits in terms of optimisation, portability, and ease of use. ReactiveX, meanwhile, is a library that makes it possible to compose asynchronous and event-based programs using observable sequences.

USE CASE 2.1.5

VERIFYING CODE AND CHECKING QUALITY AUTOMATICALLY

Starting with LHC Run 3, some of the experiments will do most of the processing for particle-collision events ‘online’. This means that it will not be possible to undo any mistakes made. It is therefore vital to have tools and techniques in place to check the quality of the processing. As well as achieving a high-level of performance, it is important that code is comprehensible and can be verified to be functioning in the desired manner.

CHALLENGE AREA 2.2

HETEROGENEOUS PLATFORMS AND ALTERNATIVE ARCHITECTURES

Up to now, the WLCG has operated with a highly distributed — yet homogeneous — computing environment. There are around 170 computing centres, but the hardware and operating systems are essentially the same.

Work is now being carried out to enable flexible expansion of resources in an opportunistic manner (to meet fluctuating requirements), via dynamically provisioned commercial cloud resources and academic resources intended to be shared with across fields. It is therefore important to expand the range of resources that can be effectively used by the LHC community. Greater adoption of platform-based systems, as well as making more extensive use of virtualisation technologies (e.g. containers), should help to achieve this.

USE CASE 2.2.1
HYBRID CPUs

Although the code used in high-energy physics is typically very large, there are a number of algorithms that are frequently used across different pieces of software (examples include algorithms for generating random numbers, for compression and decompression, and for encryption and decryption). Implementing these algorithms in hardware, directly on the CPU, could lead to significant increases in speed. CERN is interested in collaborating in this area with hardware providers, and working to test the effectiveness of such implementations.

USE CASE 2.2.2
OPTIMISING CODE DISTRIBUTION USING LIGHTWEIGHT CONTAINERS

From virtual machines to Docker and Singularity, container technology is rapidly improving. Companies that widely use container technology to run many thousands of build jobs on the cloud are also developing interesting container-based technologies. Working with the latest container technology will enable CERN to integrate and optimise code distribution and job execution in the cloud.




CHALLENGE AREA 2.3

DEDICATED HARDWARE AND COPROCESSING SYSTEMS

One of the most promising areas of development for helping to close the resource gap is the adoption of dedicated, specialised hardware. High-performance coprocessors, GPUs, FPGAs, and many-core coprocessors all have the potential to dramatically increase performance for specific sections of applications. A systematic assessment of the workflow is needed to identify where each new technology could make the largest impact, how to get the most from each component, and how the pieces could best function together.


USE CASE 2.3.1
OPTIMISING CODE PERFORMANCE USING COPROCESSORS AND GPUS

For a number of years, CERN has been investigating how to make best use of general-purpose GPUs and coprocessor systems. We are now hitting the limitations of these technologies in terms of memory bandwidth and latency, and are interested in exploring developments that would help to overcome these issues.



R&D TOPIC 3

MACHINE LEARNING AND DATA ANALYTICS



Members of the research community are already expending significant efforts to understand how to get the most value out of the data produced by the LHC experiments. The goal is to maximise the potential for discovery and employ new techniques to help ensure that nothing is missed. At the same time, it is important to optimise resource usage (tape, disk, and CPU), both in the online and offline environments. Modern machine-learning technologies — in particular, deep-learning solutions applied to raw data — offer a promising research path to achieving these goals within the time scale required by the plans for the upgrade to the HL-LHC.

The four main LHC experiments are investigating how to benefit from deep-learning techniques. Possible applications include improving precision in particle detection, better identification of interesting events, modelling detector response in simulations, monitoring experimental apparatus during data acquisition, and more efficiently managing computing resources. The growing field of 'big data analysis' has led to the development of tools and techniques for extracting information from datasets in a variety of spheres (both in industry and academia) that now rival high-energy physics in terms of volume.

High-energy physics has worked with automated identification and analysis techniques for more than 20 years, beginning with basic neural networks. Driven by the significant challenges posed by the HL-LHC programme, the proposed joint investigations with industry outlined over the following pages are much more ambitious. They represent opportunities in all segments of the data-collection and -processing chains. Some of the ideas presented would represent a significant departure from the approaches that have previously been employed — in some cases, over decades.



DATA ACQUISITION

Wherever it is possible to build empirical models of systems from large volumes of data, there is a potential role for machine-learning techniques to be employed. Through either supervised or unsupervised training, correlations and patterns in a given dataset can be learnt, thus making it possible to automatically perform specific tasks, such as event classification or clustering. Over recent years, 'deep learning' has played an essential role in speeding up research into artificial intelligence, with the development of new network architectures mimicking vision, memory, etc.

Industry has already demonstrated the utility of machine-learning techniques for monitoring, automation, and anomaly detection in complex systems. For instance, learning algorithms are used in fraud detection and network security. As the home of the world's largest and most powerful particle accelerator, CERN produces extremely large amounts of highly complex data. Machine-learning techniques can be used in industrial monitoring, to optimise resource utilisation and anticipate failures and maintenance needs. Thus, it is possible to maximise 'uptime' and minimise operational costs at large industrial sites.

USE CASE 3.1.1

MONITORING OF ACCELERATORS AND DETECTORS

The accelerators and detectors used in high-energy physics are sensitive, expensive, and highly complex. They include hundreds of millions of data channels, with thousands of monitoring measurements collected continuously. To maximise the scientific value of the LHC, the status of all components needs to be carefully monitored, with interventions scheduled accordingly. Large investments are made in this area, both in terms of human and computing resources. Learning from the experience collected during the first two LHC runs, it would be possible to develop predictive algorithms for hardware failure in various accelerator and detector systems. These algorithms could be based on machine-learning applications such as recursive neural networks, long-short-term memory networks, etc.

As with human-based monitoring, machine-learning algorithms for anomaly detection learn from good data and identify anomalies by difference. Unlike human-based monitoring, however, such algorithms do not need pre-defined lists of failure modes and ad-hoc-defined values to monitor. Accessing the raw data, 'deep neural networks' can perform effective high-level feature engineering, thus enabling them to define what constitutes normal behaviour, and spot anomalous events in real-time.

A research-and-development project is needed to put automatic monitoring software — with an efficient inference phase — into production. It would also be important to put a suitable training infrastructure into place, with close information exchange between the teams responsible for the accelerator chain and the various experiments' detectors.

USE CASE 3.1.2

MONITORING DATA QUALITY

The quality of the collision data coming from the LHC is currently evaluated 24/7 by teams working in shifts; they judge the agreement between the incoming events and a reference sample of good data. Instead of directly accessing raw data, a list of pre-defined quantities are computed in real time and represented through histograms or other graphs. Sometimes, statistical tests are performed to analyse the agreement between a collected histogram and the reference distribution. Often, the judgments made are left to the intuition of the shift-team members.

While this strategy is quite effective on average (>90% of the collected events are eventually classified as good), the strategy has two weak points:

- The level of expertise of the shift-team members is highly variable, creating inhomogeneous quality assessments.
- The monitoring strategy is often tailored to specific failure modes, and thus fails to reveal new problems induced by new experimental environments (e.g., after a change in the LHC's running conditions).

Both problems can be addressed through anomaly-detection techniques. Proof-of-principle tests have already shown the potential of such an approach. Systematic use of machine-learning techniques could make it possible to predict failures by monitoring patterns over time.

Collaboration with industry would enable the transfer of experience and best practices related to data-quality monitoring, thus helping to design a large-scale infrastructure for a global, automatic monitoring application. To achieve this, it will be necessary to carry out an extensive comparison between human-based quality assessment and the outcomes of the automatic applications; this can be done through an intense phase of research and development during the current LHC run.

USE CASE 3.1.3

FAST INFERENCE TECHNOLOGY FOR 'TRIGGER' SYSTEMS

The development of fast inference technology would have a dramatic impact on the operation of the LHC detectors. Evaluating the outcomes of machine-learning techniques for regression or classification within a few milliseconds is possible with currently available parallel computing architectures, such as GPUs and FPGAs. Within this time limit of a few milliseconds, neural networks could be used in the high-level trigger system, so as to improve the accuracy of the filter algorithms (reducing the rate at which false positives are generated), while preserving a high-level of efficiency.

Studies have shown that recursive neural networks are capable of identifying the topology of a given event from the list of input particles. This would enable the samples selected to be cleaned up using inclusive triggers systems.

Being able to train algorithms to isolate potentially interesting events from the background would have a significant impact on offline processing systems, too. More events could be evaluated by the trigger system, with events identified with a high probability of being 'background' (i.e. not of interest) might only be summarised in distributions, rather than fully reconstructed and stored. Thus, the requirements for offline processing and reprocessing, as well as for storage, could be significantly reduced. Such an approach might also make it possible to identify events of potential interest earlier in the data pipeline, thus further reducing the total resources needed for analysis.

Additional benefits could come from bringing deep learning to the level-1 trigger. There, the selection is based on a local, coarse reconstruction of the collision event, typically implemented using FPGAs. Instead of running the coarse reconstruction, one could use machine-learning algorithms to model the relation between the raw input data and the outcome of the detailed offline reconstruction. Such an algorithm could provide better answers on average. It would also reduce the differences between online and offline selection, as well as enhancing the purity of the trigger-selected sample. Such a paradigm shift requires the development of fast inference technology (within a few milliseconds) for the hardware commonly employed in the level-1 trigger, such as FPGAs (or equivalent parallel architectures).

An even more dramatic use case is provided by the LHCb experiment, which will eliminate its level-1 triggers in 2021 and send up to 5 TB of data per second to a commodity processing farm. It is not at all clear that current event-reconstruction techniques (such as tracking based on Kalman filters or reconstruction of Cerenkov detectors) can scale to these data volumes within the foreseen budget. It may therefore be necessary to make use of machine-learning techniques that are more adapted to highly parallel computing architectures, allowing some of the inference cost to be shifted onto the training of the machine-learning algorithm (for example, convolutional or recurrent neural networks deployed on

GPUs or FPGAs). It is important to demonstrate that such algorithms can run fast enough and produce acceptable physics performance.


The companies collaborating in CERN openlab were asked to give their views on whether existing tools could make decisions within the latencies required. The current high-level trigger farms are reportedly slower than those already being used in some industrial applications today (the development of technologies for self-driving cars is pushing advances in this area). Nevertheless, the proposed evolution of the level-1 hardware-trigger represents an interesting technical challenge that the research and industry communities could address together.

USE CASE 3.1.4

ANOMALY DETECTION AND THE SEARCH FOR NEW PHYSICS

Most of the particle collisions collected by the experiments' detectors are well-understood events, originating from known physics processes. When it comes to searches for new phenomena, these events form a large amount of background information, which has to be dealt with appropriately. The signal could be a very tiny fraction of the selected events or, even worse, might end up in the fraction of events that are rejected by the trigger system.

Given the current state of understanding and expectations regarding new physics, it is not inconceivable that rare events might produce signatures in the detectors that are different from those commonly accepted. Thus, in this context, one could regard such events as anomalies. Semi-supervised and supervised algorithms — running as part of the high-level trigger application — could search for such anomalies and store them in a separate stream of data, for offline visual inspection. The starting point for developing such an application is similar to that of the data-quality-monitoring use case highlighted in section 3.1.2.



DATA PROCESSING

Once data is acquired from the particle detectors, it is reconstructed to determine the physics processes occurring in the collisions. In parallel, samples of equivalent simulated events of known physics processes are produced, thus making it possible to compare and evaluate the presence of new signals. Machine-learning techniques are potentially useful in both reconstruction and simulation, which represent the majority of the computing activities in high-energy physics.

For both real and simulated data, object identification is currently performed via resource-heavy reconstruction of the position and energy information, so as to calculate the energy flow through the detector. Object identification using computer-vision techniques — trained through human supervision — is one of the most common applications of machine learning in industry. The goal is to sort large quantities of data and identify the contents based on images. In high-energy physics, the readout of the detector is similar to a 3D image composed of position and energy information. Using 3D vision techniques and training — similar to those used in 2D image recognition — may make it possible to use machine-learning tools to distinguish physics objects based on topology, without time-consuming reconstructions and other calculations.

USE CASE 3.2.1

SIMULATION

Simulation plays a key role in all high-energy physics experiments: from the development of the detector, to the design of strategies for data analytics, and the comparison of results to theoretical models. With the future HL-LHC runs, a major increase in luminosity and dataset size is expected, thus offering the possibility to probe rarer phenomena more accurately. The need for simulated data will increase enormously: the expectation of the LHC experimental collaborations is that simulation throughput will need to increase by a factor of around 100 compared to today. However, the resources available for this are not expected to increase.

It is therefore essential that the performance of the necessary applications increases dramatically. As well as employing new methods and algorithms, it is vital to optimise existing code — so as to take advantage of the increased parallelism of the new computing architectures. Throughout its fifth phase, CERN openlab has supported a number of collaborative projects to modernise scientific software, including for the GeantV simulation toolkit, which is used to simulate how particles interact with matter.

Recent developments in machine-learning techniques could even offer the possibility of replacing complex algorithms with machine-learning approaches (for example, a deep-learning neural network) that produce the results at a much higher speed. Traditional Monte Carlo simulation could be treated as a ‘black box’ that can be replaced by a deep-learning algorithms trained on a range of information related to particle type, momentum, position, etc.

Initial interest is centred on the testing of two different approaches: multi-objective regression and state-of-the-art computer-vision techniques, such as generative adversarial networks. Inputs would be provided by the elementary particles produced by collisions, or coming directly from particle beams. More complex input types could also be considered. On the output side, there is also a large range of possibilities to consider: the most basic is the energy deposition in various parts of the detector. More complex outputs, such as response signals from the detectors’ data-acquisition systems can also be studied, as well as the final reconstructed collision events.

In the multi-objective regression approach, each particle is characterised by type, energy, and direction, as well as the position of impact within the detector’s layers. This typically leads to around eight or nine parameters being recorded for each particle. Depending on the choices made in terms of input and output, the number of parameters may be very large; techniques such as feature reduction would then need to be applied to reduce dimensionality (and reduce learning time) while preserving correlations.

By contrast, in the approach based on computer-vision techniques, detector response is considered as a whole: it is treated directly as a 3D image that deep networks are trained to reproduce as accurately as possible.

USE CASE 3.2.2

JET IDENTIFICATION AND IMAGE-BASED EVENT IDENTIFICATION

Identifying physics objects in sections of the detector would be useful. However, an alternative approach based on looking at complete collisions and performing high-speed event classification based on topology might enable the future low-latency requirements of the trigger systems to be met. Investigations are planned into classification based on image recognition without pre-computed reconstructed physics objects. The techniques and hardware used in image recognition are rapidly improving and evolving, so applications that are not feasible today may become possible in time for start of the HL-LHC.

Accurately identifying jets — narrow cones of hadrons and other particles produced under certain conditions — is complex. It is important to investigate the amount of information that would need to be provided to a machine-learning network set up to carry out this identification. Such an investigation would enable potential performance to be evaluated and compared with traditional techniques.

Investigating the feasibility of classifying complete events in real-time — based on detector ‘hit’ information (using image-recognition techniques) — is also recommended.

BIG DATA

By 2025 (when the HL-LHC is planned to begin operation), it is expected that the amount of data produced by the experiments will have increased by between one and two orders of magnitude. There will be more collisions (about three times higher than today), and these will become markedly more complex. Extracting signals from an enormous set of background information continues to be the single largest computing challenge in physics analysis.

The huge growth in datasets from social media, e-commerce, and other web applications has driven the development of new tools to access, process, select, and reduce datasets. In collaboration with our industrial members, we would like to explore how such tools could be applied to the processing and analysis of scientific data.

USE CASE 3.3.1

DATA REDUCTION AND REFRESH FOR ANALYSIS

One of the most resource-heavy aspects of data analysis is reducing the bulk of a large sample to produce a concentrated dataset for analysis. In the past, this reduction was often performed through a series of cuts applied to a large sample of events, with users and groups looping over the input samples individually. Agreeing on selection algorithms to make the cuts — and indexing the results — offers the possibility of significantly increasing the efficiency of the analysis. Tools like Spark and Hadoop, as well as techniques like MapReduce, have been successfully deployed in industrial and scientific applications to select objects based on defined and reusable selection criteria.

In addition to reducing data to form analysis datasets, data regularly needs to be recalculated under new conditions and using new software. Normally, data-reprocessing campaigns are largely centralised endeavours that involve recalculating the entire input sample. The large amount of computing resources required for such a reprocessing ‘pass’ means they cannot be performed often. At the same time, it is important that the output data meets the needs of as many users as possible without customisation.

By combining the concepts of data reduction with reprocessing, it may be possible to refresh subsets of the data only. This would enable more targeted reprocessing passes that make more efficient use of resources. And it would make it possible to focus on the needs of individual physicists, or small groups of users.

The goal is to reduce an input dataset by a factor of 1000 in a few hours: a 1-PB input sample would be reduced to a 1-TB relocatable output sample, with reusable cuts and indexed events. Once this has proven successful for one user, the next step is to demonstrate that hundreds of users can scan through data using an infrastructure based on big-data tools to produce analysis datasets.

It is important that after selecting a subset of events, the reprocessing of physics objects under new conditions and using new software can be performed to make a consistent and reproducible dataset. The intermediate results and provenance information should be stored, so that future requests for a data refresh can be fulfilled without expending the processing resources a second time.


USE CASE 3.3.2

OPTIMISATION OF COMPUTING INFRASTRUCTURE

Many of today’s big science endeavours are underpinned by distributed computing infrastructures (the WLCG in the case of the LHC research community). Such infrastructures are often highly heterogeneous, with large numbers of distributed computing and storage components connected via complex network topologies. The optimisation of the scientific output from these infrastructure can be very complex, but can potentially lead to increased efficiency in terms of resource use.

CERN has a dedicated IT Analytics Working Group, which has started to apply statistical and machine-learning methods to large sets of metrics collected from main system components: CPU and batch, disk and archive storage, network topology and flows, and application throughput. A first study into the clustering of batch jobs by their I/O patterns already shows promising potential for optimisation by matching the different job I/O profiles to the different hardware capabilities offered by the existing system.

More generally, the metrics collected over several years at the CERN data centre could be used to construct infrastructure models that — after confirmation using real data — enable quantitative reasoning regarding options for optimisation and future investment. Such studies could be of significant value for many stakeholders, both at CERN and beyond. The set of component metrics currently available already makes it possible to carry out quantitative analysis of the time-dependent reliability of different hardware components in the CERN data centre (e.g. magnetic disks, solid-state drives, RAM modules). It also enables optimisation of operational procedures (e.g., proactive hardware replacement) and purchasing decisions.



DATA ENGINEERING

Data engineering is the backbone of many machine-learning systems. The majority of efforts expended in machine-learning projects often relate to data-intensive tasks, as demonstrated by D. Sculley and colleagues from Google, one of the pioneering companies in this area. In a paper entitled ‘Hidden Technical Debt in Machine Learning Systems’, presented at the 2015 Neural Information Processing Systems Conference, Sculley and his colleagues argue that it is a mistake to think of machine-learning systems as providing easy ‘quick wins’. They make the case that it is common to incur massive ongoing maintenance costs in real-world machine-learning systems.

USE CASE 3.4.1

SOLUTIONS FROM INDUSTRY: CHALLENGES AND OPPORTUNITIES

Many big-data solutions from industry and open-source communities enable users to employ machine-learning techniques. Several of these solutions have seen a quick rise in adoption over the past few years; many now have large user communities that are growing rapidly. Based on the CERN IT department’s experience delivering big-data services, several key components for these solutions have been identified. These include the following:

- Apache Hadoop clusters with YARN (‘Yet Another Resource Negotiator’) cluster manager and the Hadoop distributed file system.
- Apache Spark, a general purpose engine for analytics, machine learning, and stream processing.
- Apache Kafka, a messaging system, which provides the foundations for many streaming solutions.
- Apache HBase and Apache Kudu are also components of interest.

Another key technology available to the CERN community is hosted Jupyter notebooks, notably with the CERN-developed platform SWAN (‘Service for Web based Analysis’). Following adoption patterns in industry, notebooks have proven of significant use to the CERN community, providing an easy way for users to access analytics platforms and software.

Several challenges — to be tackled by the teams that provide the CERN Hadoop and Spark services — have already emerged from the projects in this area:

- Integration with existing libraries and data-processing formats of interest to CERN communities is important. Libraries for the integration of Spark with the ROOT data-processing framework and EOS (a disk-based low-latency storage service used at CERN) are currently being developed.
- Machine-learning and deep-learning solutions need to be tested at scale with Spark. Efforts in this area have already started at CERN with the development of a library integrating Keras (a high-level API for neural networks) with Spark. In the future, further development of this work is foreseen. This will be carried out with the interested user communities at CERN, and will follow industry trends (advancements in this area are being made rapidly with several libraries for integrating Spark and deep-learning frameworks currently available).
- Additional investigations need to be carried out to improve understanding of performance and scaling. CERN currently has experience with a Hadoop cluster of 10-50 nodes and just a few petabytes of data. To more meaningfully represent physics use cases, it is important to perform testing at scales of tens of petabytes and hundreds of nodes.


The evolution of the platform architecture also poses important challenges in this area. Many platforms provided by the CERN IT department are based on hosted Hadoop clusters. This is for historical reasons, with many benefits provided by Hadoop (notably, the ability to take advantage of data locality). However, the bulk of physics data in high-energy physics is stored on dedicated



storage systems (notably EOS). This makes it possible to deploy Spark clusters, accessing data via the network. This is similar to what several cloud providers in industry offer today (with the notable difference that the storage system would be EOS rather than Amazon S3, for example).

Analytics workloads also benefit greatly from CPU resources and large memory caches (both dynamic and non-volatile random-access memory). Investigations in this area are an important priority.

Streaming solutions at CERN provide another possible area for innovation. They are currently used in many data pipelines at the laboratory. Such solutions are also popular in industry, often successfully replacing batch ingestion and traditional ETL (extract, transform, load) processing. Combining streaming with machine learning has great potential (see, for example, 3.1.3). There is rapid innovation in the solutions provided by both industry and open-source communities, with significant advances being made in terms of simplifying APIs, reducing latency, and improving robustness. Of particular interest is the work carried out by the Apache Kafka community, as well recent developments for Apache Spark.



R&D TOPIC 4



APPLICATIONS IN OTHER DISCIPLINES

Until recently, high-energy physics was the only research discipline handling extremely large quantities of data. However, life sciences, medical research, Earth sciences, astrophysics, and environmental sciences have all started embracing the paradigm of large-scale, data-driven, distributed research. These communities will soon generate comparable — or even greater — amounts of data.

Oncology, genomics, and other ‘omics’ fields already deal with ‘big data’. They each have data that fulfils the criteria laid out in the prominent ‘Five Vs’ model (expanded from the ‘Three Vs’ model popularised by Gartner): volume, variety, velocity, variability, and veracity. The total volume of cancer patient data in the world today is estimated to be anywhere between 14 and 1400 PB. Storage requirements for global human genome-sequence data are expected to reach as much as 40 exabytes (EB) by 2025. Nevertheless, most of the data in these fields is still largely inaccessible to broader communities of researchers.

On top of this, vast amounts of data are being produced by growing efforts to simulate the development and behaviour of increasingly complex biological systems. The Human Brain Project (HBP) is a prime example of this, but there are also many other international collaborations in this area.

The Square Kilometre Array (SKA) is another major international project, this time in the field of astronomy. It is defining ambitious goals to collect and process — in real time — unprecedented quantities of data at very high rates. The data will also be distributed across international collaborative infrastructures, as happens today in high-energy physics.

In addition to such large-scale research initiatives, there are a number of emerging new technologies that have the potential to transform the way we interact with our environment and to optimise the way services are delivered. Examples include the ‘Internet of things’ (IoT), decentralised protocols (block chains), high-performance mobile data networks (5G), and autonomous systems (including self-driving cars). These technologies could deeply affect the way large research institutions like CERN manage their infrastructure and users, as well as how they impact on the environment and the local communities around them. However, they do also pose significant technical challenges (particularly related to networks, sensors, data quality, etc.), as well as raising ethical considerations (particularly around data privacy and protection).

Early collaborations formed during the construction of the WLCG opened the way to partnerships between different communities, with researchers working together across communities to adapt tools to meet specific needs. New common investigations on the effective use of emerging general-purpose technologies, such as machine learning and new accelerated computing platforms, would be of great benefit across research domains. Today, new projects and initiatives are being defined to develop distributed-computing and data platforms that take into account the specific needs of research communities such as biomedicine and astrophysics, while also incorporating some of the data-sharing and open-collaboration principles from the high-energy-physics community.



CHALLENGE AREA 4.1

PLATFORMS FOR OPEN COLLABORATION

New data/computing platforms common across research communities could provide many advantages:

- A general entry point for access to data-analysis and machine-learning tools.
- Libraries of community-moderated software.
- Integrated access to open data and publication repositories (like the Zenodo system already in use at CERN).
- Channels for sharing information and best practices.

Such platforms have the potential to increase the reproducibility of experiments by linking publications, software, data, and workflows with validated identifiers from reputable, community-curated repositories. However, the path towards efficient implementation of such platforms is not without challenges. Much effort is required to improve semantic definitions, to integrate and validate the quality of heterogeneous data, and to address the need for confidentiality of the data. At the same time, it is vital that the validity of the models generated is not compromised.

The development of user-friendly interfaces to easily describe workflows, submit/retrieve data, and share information, is of paramount importance. The training required to use these platforms should be kept to a minimum.

USE CASE 4.1.1

A SMART DATA-ANALYSIS PLATFORM

Over recent years, engineers and developers at CERN — and in the wider community around high-energy physics — have developed a set of tools to help manage virtualised infrastructures. These include tools to manage aspects such as deploying virtual machines and containers, analysing and moving data, storing diverse digital assets of many different formats.

Generally, these tools have been designed to ease the tasks of sharing data and creating large-scale collaborations. Examples include ROOT, SWAN, Zenodo, and CERNVM — all mentioned previously in this white paper. These tools — combined with other mainstream open-source tools for data analytics and machine learning — could provide the building blocks for the development of open, accessible, smart data-analysis platforms for scientific research. Such a platform could be designed for deployment on international cloud platforms for science.

CERN is interested in collaborating on the integration of different tools ('blocks') in this manner. Work should be carried out to benchmark and test the relevant tools, to automate software installation and deployment over distributed infrastructures, to assess experimental reproducibility, and to define appropriate scientific use cases.



CHALLENGE AREA 4.2

LIFE SCIENCES AND MEDICAL APPLICATIONS

New large-scale collaborative projects in the life sciences are generating increasing amounts of data, as is the extension of electronic health records to more and more people across the globe. The collaboration between the high-energy physics community and the medical community on developing common platforms is therefore becoming increasingly important.

Some of the tools and methods developed in the context of physics experiments to select, store, and analyse very large amounts of complex data can be — and in some case already have been — adapted or generalised for use in the life sciences and for medical applications. Examples of such collaboration over the past 15 years include projects like Health-e-Child, MammoGrid, and WeNMR. These all explored the possible medical applications of grid computing, as developed for the LHC research community. Tools used in high-energy physics for simulation, data analytics, and data-quality assurance could potentially become effective instruments for harnessing the wealth of data being produced in medicine and other life sciences today.

USE CASE 4.2.1

SIMULATING BIOLOGICAL SYSTEMS IN THE CLOUD

In order to understand what happens when particles collide in the detectors of the LHC experiments, very precise simulations of the underlying physics process are required. Today, this is done with a sophisticated tool called Geant4. A new version of this tool is now being developed; this will exploit modern computing platforms and machine-learning techniques to accelerate simulation.

Similar techniques can be applied to the simulation of biological tissue, including how interaction rules between cells (mechanical, electrical, and chemical) affect development. In collaboration with other research institutions, CERN is investigating how to efficiently deploy simulations of tightly coupled, dynamic systems on commodity cloud platforms. For this, it is important to take into account the necessary requirements regarding connectivity across nodes, performance, I/O, user interactions and more.

USE CASE 4.2.2

LARGE-SCALE ANALYSIS OF GENOMIC DATA

Advances in genomic sequencing techniques — and an associated rapid decrease in the cost of carrying out sequencing — has led to the production of increasingly large amounts of data.

However, the sensitive nature of genomic data means that sharing and establishing cross-community collaborations can be difficult.

In addition, many of the tools traditionally used for genomic data analysis were not originally designed to cope with large-scale, collaborative analysis. It would, therefore, be valuable to investigate whether techniques and tools used in high-energy physics (such as the ROOT data-processing framework) can be adapted to handle genomic data and to assess whether this could lead to improvements over the performance of existing tools.

Furthermore, systems like the smart data-analysis platform outlined in 4.1.1 could enable scientists to deploy software in easy, reproducible ways, as well as share data pipelines and best practices. Such platforms could even be used to share data and results under appropriate conditions of privacy and confidentiality.

USE CASE 4.2.3

LARGE-SCALE ANALYSIS OF HEALTHCARE DATA

The use of personal healthcare and fitness devices is becoming more widespread. General interest in information related to personal well-being and medicine is also growing, with much information available freely online. While such information can be very useful in terms of making people more aware of risks and benefits related to their behaviour, there is also a risk of misinterpretation (especially where information is handled in the absence of trained professionals).

The challenge of collecting, analysing, and interpreting health information from different sources (medical records, personal devices, nutritional information, activity patterns, etc.) under the appropriate confidentiality conditions requires collaborative multi-disciplinary initiatives, with aspects related to computing and data-analysis playing a fundamental role.

Machine-learning and data-quality monitoring systems developed by the LHC research community could potentially be exploited to play a role in this. It would be interesting to explore how diverse data ecosystems (with varying levels of data quality and provenance information) could be used to inform new experimental hypotheses, help trained professionals make informed diagnoses, and drive the evolution of safety and prevention practices — including those applicable to the CERN community.

CERN has over 3000 members of personnel working across its two main sites (in Meyrin, Switzerland, and in Prévessin, France), as well as many more visiting scientists, engineers, technicians, students, and trainees. The laboratory has a dedicated medical service, which carries out examinations, runs prevention campaigns, provides information regarding risks at work, and generally assists those at the laboratory in achieving and maintaining a good health.

Any investigation in this area must consider the need for appropriate and flexible pseudo-anonymisation tools, so as to adequately protect sensitive data. At the same time, such tools must preserve enough information in the anonymised data to enable useful scientific insights to be gained.

ASTROPHYSICS

New large-scale projects in astrophysics are facing computing challenges at the 'exascale' level, leading to increased interest from other scientific communities. Researchers from astrophysics projects now in the design/construction phase (such as the Square Kilometre Array and the Large Synoptic Survey Telescope) are working together with experts from the high-energy physics community to understand and define strategies for coping with the huge volumes of data these infrastructures will produce. It is important to collaborate closely with industry to explore the limits of technologies and methods for acquiring, handling, and analysing this data. CERN openlab can play a leading role in this.

USE CASE 4.3.1


EXASCALE DATA PROCESSING AT FUTURE ASTROPHYSICS INFRASTRUCTURES

Future astrophysics infrastructures are set to produce never-before-seen data rates. A prime example of this is the Square Kilometre Array (SKA) project, which is an international effort to build the world's largest radio telescope; it will eventually have thousands of dishes and up to a million antennas, covering over a square kilometre of collecting area. The SKA will enable astro-

nomers to monitor the sky in unprecedented detail and survey the entire sky much faster than any system currently in existence. Phase 1 of the project — representing only approximately 10% of the whole SKA — will generate around 300 PB of data per year. CERN and the SKA recently signed a cooperation agreement that establishes a framework for collaborative projects; these will address joint challenges related to data storage and the approaching exascale computing era.

The SKA faces computing challenges similar to those in high-energy physics, particularly in terms of data selection and processing. The data rates from the SKA's dishes and antennas will need to be intelligently reduced from petabits to terabits in real time.

Analysing the acquired data will require (on average) roughly 25 petaFLOPS of processing power to be delivered. At peak times, this may be as high as 250 petaFLOPS. Also, due to the size of the dataset, memory bandwidth of 200 PB/s will be needed. The high data rates to the processing elements — and the large processing requirements — make this a good potential match for hardware accelerators like FPGAs. The ongoing investigations within CERN openlab for exploiting new hardware architectures for the HL-LHC have the potential to offer benefits for the SKA.



SMART EVERYTHING

In the past few years, we have witnessed the emergence of a completely new set of technology paradigms, from pervasive ultra-fast networks of smart sensors in the 'Internet of things', to machine learning and artificial intelligence applied to the 'optimisation of everything'. These technologies have the power to revolutionise the way big science is done (particularly in terms of data collection and sharing), and could have major implications for the operation of the engineering process supporting such endeavours.

These technologies also have enormous potential in terms of wider societal applications. However, they also raise many questions in terms of data privacy and protection, as well as having implications for human interactions — both with one another and with machines.

Understanding these wider applications — and their implications — is an important aspect of the ICT research activities carried out at international laboratories like CERN. As part of this work, it is important to consider the impact of such applications on service provision, work efficiency, and personal well-being, thus helping to define the limits that should be imposed on such applications.

USE CASE 4.4.1

ENVIRONMENTAL MONITORING

The widespread availability of sophisticated sensors enables information to be collected from a large variety of technical equipment. Such information can be used to help understand and

optimise usage patterns for equipment, as well as improving energy consumption. This is relevant for various aspects of CERN's infrastructure, including data-centre operations (as described in R&D topic 1), the management of lights, heating, and ventilation systems in the buildings, water consumption, and waste disposal. CERN is interested in investigating state-of-the-art technologies in this area.

USE CASE 4.4.2

TRAFFIC AND MOBILITY

CERN's Meyrin site straddles the Franco-Swiss border, with its Prévessin site located just three kilometres away in France. The sites form part of a complex local traffic system, which includes elements such as people, cars, bicycles, trams, buses, roads, and parking spaces. The optimisation of transport through smart technologies would bring benefits to CERN, as well as to the local population. Shared commuting policies, park-and-ride systems, and communal electrical vehicles could all potentially play a role.

There is growing interest in collaborating on the technical aspects of setting up new and smarter transports solutions, including understanding and optimising traffic patterns, as well as designing sustainable mobility strategies for the future. The design, deployment, and operation of the networks of sensors required to realise this vision — and the related data-analysis activities — are areas where CERN could contribute significant expertise.



ABOUT CERN OPENLAB



CERN openlab is a unique public-private partnership that works to accelerate 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 this 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 ICT industry collaborators 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 research-and-development activities with more than one company.

Since its foundation in 2001, CERN openlab has run in successive three-year phases. In 2018, CERN openlab begins its sixth three-year phase.

Visit the CERN openlab website to find out more about our mission, history, and opportunities for collaboration.



www.cern.ch/openlab

OUR COLLABORATORS



ORACLE

SIEMENS



BROCADE



Yandex



INNOVATION
UNIVERSITY



EMBL-EBI



ADDITIONAL INFORMATION



RELATED EVENTS

CERN OPENLAB TECHNICAL WORKSHOP

8-9 December 2016

<http://cern.ch/go/Bqk8>

CERN OPENLAB WORKSHOP ON DATA-CENTRE TECHNOLOGIES AND INFRASTRUCTURES

1 March 2017

<http://cern.ch/go/66ks>

CERN OPENLAB WORKSHOP ON COMPUTE PLATFORMS AND SOFTWARE

23 March 2017

<http://cern.ch/go/Nc9m>

CERN OPENLAB WORKSHOP ON MACHINE LEARNING AND DATA ANALYTICS

27 April 2017

<http://cern.ch/go/7ShK>

CERN OPENLAB 'OPEN DAY'

21 September 2017

<http://cern.ch/go/T8ff>

RELATED DOCUMENTS

CERN OPENLAB: A PUBLIC-PRIVATE PARTNERSHIP TO DRIVE ICT INNOVATION IN SCIENCE

<http://cern.ch/go/Z6Pl>

CERN OPENLAB ANNUAL REPORT 2016

<http://cern.ch/go/7rQ7>

CERN ANNUAL REPORT 2016

<http://cern.ch/go/6LJS>

CERN OPENLAB WHITE PAPER ON FUTURE IT CHALLENGES IN SCIENTIFIC RESEARCH 2014

<http://cern.ch/go/JR8g>

CONTRIBUTORS



MAIN AUTHORS

Alberto Di Meglio
Maria Girone
Andrew Purcell
Fons Rademakers

ADDITIONAL CONTRIBUTORS

David Abdurachmanov
Eric Acquaronne
Piero Altoè
John Apostolakis
Luca Atzori
Omar Awile
Virginia Azzolini
Olof Barring
Tim Bell
Claudio Bellini
Latchezar Betev
Ian Bird
Jean-Phillipe Breyse
Luca Canali
Federico Carminanti
Tony Cass
Marco Clemencic
Ignacio Coterillo Coz
Xavier Espinal Curull
Dirk Duellmann
Laurent Duhem
Katarzyna Maria Dziejziniewicz-Wojcik
Luis Rodriguez Fernandez
Ian Fisk
Michele Floris
David Francis
Orestis Galanis
Cosimo Damiano Gianfreda
Tom Gibbs
Sergei Gleyzer
Vladimir Gligorov
Eric Grancher
Oliver Gutsche
John Harvey
Frédéric Hemmer
Vincenzo Innocente
Sverre Jarp
Bob Jones
Michael Kagan

Sami Kama
Walter Lampl
Alison Lowndes
Andrea Luiselli
Jon Machen
Marco Manca
Antonio Romero Marin
Manuel Martin Marquez
Edoardo Martelli
Frans Meijers
Emilio Meschi
Lorenzo Moneta
Antonio Nappi
Axel Naumann
Niko Neufeld
Felice Pantaleo
Bernd Panzer-Steindel
Gerardo Pelosi
Eva Dafonte Perez
Andreas-Joachim Peters
Jean-Laurent Philippe
Maurizio Pierini
Fernando Varela Rodriguez
David Rohr
David Rousseau
Marie-Christine Sawley
Thomas Schoenemeyer
Luigi Scibile
Liz Sexton-Kennedy
Stefan Stancu
Graeme Stewart
David Strom
Filippo Maria Tilaro
Sofia Vallecorsa
Pere Mato Vila
Jean-Roch Vlimant
Barthelemy Von Haller
Artur Wiecek

CONTACTS



ALBERTO DI MEGLIO

Head of CERN openlab
alberto.di.meglio@cern.ch

MARIA GIRONE

Chief technology officer
maria.girone@cern.ch

FONS RADEMAKERS

Chief research officer
fons.rademakers@cern.ch

ANDREW PURCELL

Communications officer
andrew.purcell@cern.ch

KRISTINA GUNNE

Administrative/finance officer
kristina.gunne@cern.ch

CERN openlab white paper on future IT challenges in scientific research (September 2017)

Editor: Andrew Purcell

Design and production: CERN Design and Visual Identity Service

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

ATLAS Collaboration (page 2); Sylvain Fasel (page 12); Federico Ronchetti (page 20);
Veronika McQuade (pages 22, 36); Diego Giol (page 24); CMS Collaboration (page 26);
LHCb Collaboration (page 29); Mark Ellisman, National Center for Microscopy and Imaging
Research (NCMIR), CC BY-NC 2.0 (page 34); © SKA Organisation/Swinburne Astronomy
Productions, CC BY 3.0 (page 38); © ESA/NASA (page 39).

An electronic version of this white paper is available on the CERN openlab website:

www.cern.ch/openlab/publications/technical-documents

ISBN: 978-92-9083-465-6

Published by CERN

© CERN 2017



