EGI SCIENTIFIC USE CASE TEMPLATE:

COMPUTING FOR LHC

Compiled by Jesús Marco, IFCA, CSIC

Collaboration Name: CMS Tier-2 Spain & Ibergrid

Contact details marco@ifca.unican.es

18 Nov 2013

1 Overview

The Large Hadron Collider at CERN is one of the best examples of international research project where e-Infrastructures, and in particular distributed computing infrastructures like EGI, have proved to be critical to its success.

The discovery of the Higgs boson in 2012 by the ATLAS and CMS collaborations at LHC, a major outstanding milestone in science, was possible thanks to the excellent performance and large capabilities of the Worldwide LHC Computing Grid (WLCG) project, relying on EGI infrastructure and services.

By February 2013 LHC completed a very successful data taking period, and entered in what is known as Long Shutdown 1 (LS1) period, that will be devoted to the upgrade of the accelerator to reach a proton-proton collision energy of 13 TeV by 2015.

Since then, the four LHC experiments, ALICE, ATLAS, CMS and LHCb, have relied in the WLCG and in EGI to continue the processing of the data collected, and progress on the physics analysis that are resulting in an impressive collection of new results, as reflected by more than 500 publications in 2013, covering a wide range of topics.

When the upgraded LHC will restart operations in Run 2 by 2015, the computing needs of the collaborations will substantially increase, due to the higher complexity of the collisions with an energy almost doubled, and the collision frequency and pile-up conditions. Processing this large stream of new data will require an upgrade in the capacities of the e-Infrastructure, more than doubling current computing power and data storage resources. The possibility of exploring new physics beyond the Standard Model, like for example complex multi parametric models as those based on supersymmetry, will require even larger resources and services.

There is a clear challenge for the next years in the evolution of the current WLCG related infrastructure, organized in a tiered structure, into a dynamic, flexible and distributed framework, likely based on federated cloud resources, where EGI should play a key role.

2 Scientific Case

The Large Hadron Collider (LHC) [1] is the world's largest and most powerful particle accelerator. It first started up on 10 September 2008, and remains the latest addition to CERN's accelerator complex. The LHC 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. Inside the accelerator, two high-energy particle beams travel at close to the speed of light before they are made to collide. All the controls for the accelerator, its services and technical infrastructure are housed under one roof at the CERN Control Centre. From here, the beams inside the LHC are made to collide at four locations around the accelerator ring, corresponding to the positions of four particle detectors – ATLAS, CMS, ALICE and LHCb.

These experiments use detectors to analyse the myriad of particles produced by collisions in the accelerator and are run by collaborations of scientists from institutes all over the world. Each experiment is distinct, and characterized 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 specialized for focussing on specific phenomena. These four detectors sit underground in huge caverns on the LHC ring.

ALICE (A Large Ion Collider Experiment) is a heavy-ion detector on the Large Hadron Collider (LHC) ring. It is designed to study the physics of strongly interacting matter at extreme energy densities, where a phase of matter called quark-gluon plasma forms. All ordinary matter in today's universe is made up of atoms. Each atom contains a nucleus composed of protons and neutrons (except hydrogen, which has no neutrons), surrounded by a cloud of electrons. Protons and neutrons are in turn made of quarks bound together by other particles called gluons. No quark has ever been observed in isolation: the quarks, as well as the gluons, seem to be bound permanently together and confined inside composite particles, such as protons and neutrons. This is known as confinement. Collisions in the LHC generate temperatures more than 100,000 times hotter than the centre of the Sun. For part of each year the LHC provides collisions between lead ions, recreating in the laboratory conditions similar to those just after the big bang. Under these extreme conditions, protons and neutrons "melt", freeing the quarks from their bonds with the gluons. This is quark-gluon plasma.

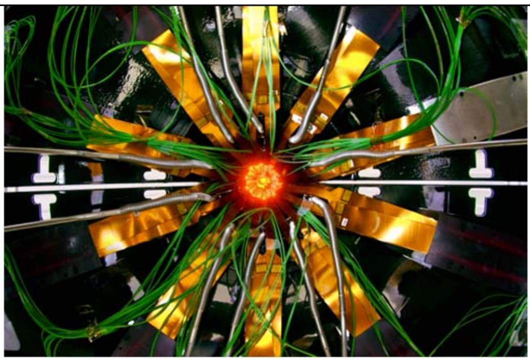

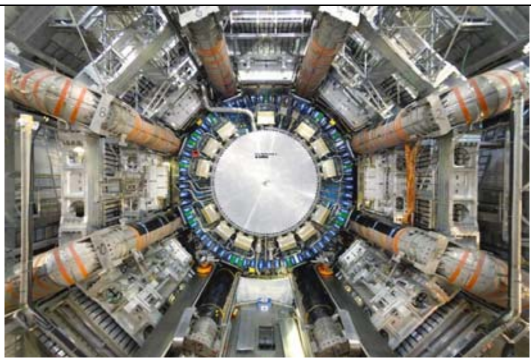

The existence of such a phase and its properties are key issues in the theory of quantum chromodynamics (QCD), for understanding the phenomenon of confinement, and for a physics problem called chiral-symmetry restoration. The ALICE collaboration studies the quark-gluon plasma as it expands and cools, observing how it progressively gives rise to the particles that constitute the matter of our universe today. The ALICE collaboration uses the 10,000-tonne ALICE detector – 26 m long, 16 m high, and 16 m wide – to study quark-gluon plasma. The detector sits in a vast cavern 56 m below ground close to the village of St Genis-Pouilly in France, receiving beams from the LHC. The collaboration counts more than 1000 scientists from over 100 physics institutes in 30 countries.

ATLAS is one of two general-purpose detectors at the Large Hadron Collider (LHC). It investigates a wide range of physics, from the search for the Higgs boson to extra dimensions and particles that could make up dark matter. Beams of particles from the LHC collide at the centre of the ATLAS detector making collision debris in the form of new particles, which fly out from the collision point in all directions. Six different detecting subsystems arranged in layers around the collision point record the paths, momentum, and energy of the particles, allowing them to be individually identified. A huge magnet system bends the paths of charged particles so that their momenta can be measured. The interactions in the ATLAS detectors create an enormous flow of data. To digest the data, ATLAS uses an advanced “trigger” system to tell the detector which events to record and which to ignore. Complex data-acquisition and computing systems are then used to analyse the collision events recorded. At 46 m long, 25 m high and 25 m wide, the 7000-tonne ATLAS detector is the largest volume particle detector ever constructed. It sits in a cavern 100 m below ground near the main CERN site, close to the village of Meyrin in Switzerland. More than 3000 scientists from 174 institutes in 38 countries work on the ATLAS experiment.

The Compact Muon Solenoid (CMS) is a general-purpose detector at the Large Hadron Collider (LHC). It is designed to investigate a wide range of physics, including the search for the Higgs boson, extra dimensions, and particles that could make up dark matter. Although it has the same scientific goals as the ATLAS experiment, it uses different technical solutions and a different magnet-system design. The CMS detector is built around a huge solenoid magnet. This takes the form of a cylindrical coil of superconducting cable that generates a field of 4 Tesla, about 100,000 times the magnetic field of the Earth. The field is confined by a steel “yoke” that forms the bulk of the detector’s 12,500-tonne weight. An unusual feature of the CMS detector is that instead of being built in-situ like the other giant detectors of the LHC experiments, it was constructed in 15 sections at ground level before being lowered into an underground cavern near Cessy in France and reassembled. The complete detector is 21 metres long, 15 metres wide and 15 metres high. The CMS experiment is one of the largest international scientific collaborations in history, involving 4300 particle physicists, engineers, technicians, students and support staff from 179 universities and institutes in 41 countries.

The Large Hadron Collider beauty (LHCb) experiment specializes in investigating the slight differences between matter and antimatter by studying a type of particle called the "beauty quark", or "b quark". Instead of surrounding the entire collision point with an enclosed detector as do ATLAS and CMS, the LHCb experiment uses a series of subdetectors to detect mainly forward particles - those thrown forwards by the collision in one direction. The first subdetector is mounted close to the collision point, with the others following one behind the other over a length of 20 metres. An abundance of different types of quark are created by the LHC before they decay quickly into other forms. To catch the b quarks, LHCb has developed sophisticated movable tracking detectors close to the path of the beams circling in the LHC. The 5600-tonne LHCb detector is made up of a forward spectrometer and planar detectors. It is 21 metres long, 10 metres high and 13 metres wide, and sits 100 metres below ground near village of Ferney-Voltaire, France. About 700 scientists from 66 different institutes and universities make up the LHCb collaboration

2.1 The Scientific Challenge

Since the start of LHC, an impressive progress has been made in the pursuit of its core mission, elucidating the laws of nature at the most fundamental level. A giant leap, the discovery of the Higgs boson, has been accompanied by many experimental results confirming the Standard Model beyond the previously explored energy scales.

These results raise further questions on the origin of elementary particle masses and on the role of the Higgs boson in the more fundamental theory underlying the Standard Model, which may involve additional particles to be discovered around the TeV scale. The discovery of the Higgs boson is the start of a major programme of work [2] to measure this particle's properties with the highest possible precision for testing the validity of the Standard Model and to search for further new physics at the energy frontier. The LHC is in a unique position to pursue this programme. Europe's top priority should be the exploitation of the full potential of the LHC, including the high-luminosity upgrade of the machine and detectors with a view to collecting ten times more data than in the initial design, by around 2030. This upgrade programme will also provide further exciting opportunities for the study of flavour physics and the quark-gluon plasma.

2.2 State of the art

By 2013 the four LHC experiments are completing the processing of Run 1 data, and the analysis driving to publication of results in topics including Standard Model tests, Top quark physics, B physics, Higgs analysis, search for supersymmetric particles, Exotica, Forward Physics, and Heavy Ion collisions, among others.

The resources available are devoted to the processing of the collected data, from RAW format to ESD/AOD/DST format and skimming to final analysis n-tuples, to simulation of different signal and backgrounds, and to the estimation of limits or discovery significance. These tasks require more than 50 million jobs per month, running on more than 250.000 processors, and a total storage capacity exceeding currently 180 Petabytes on disk and 150 Petabytes on tape.

After completion of the processing and analysis of Run 1 data, the collaborations are starting the detailed simulation for Run 2 that must be ready by 2015, in time for the first analysis with data collected at 13 TeV, to be presented in 2015 summer conferences. Given the potential discovery associated to the large increase in center of mass energy, from 8 to 13 TeV, and the also large increase in backgrounds and signal event complexity, this is a clear challenge for LHC computing.

2.3 Going beyond the state of the art

Two relevant scientific challenges in the area than can be studied in the next years at the LHC are described below [6]:

Properties of the Higgs boson: since the discovery of the 126 GeV Higgs-like particle in Summer 2012, the LHC experiments have focused on the measurements of its production rates and couplings. Both ATLAS and CMS have released results strongly suggesting that the new particle is a Higgs boson and its properties are consistent with the expectations of the SM Higgs boson. After the two-year shutdown, the LHC is scheduled to operate again in 2015 at 13 TeV and it is expected to deliver 300 fb^{-1} to each experiment by 2022. With the planned high luminosity (HL-LHC) upgrade, an integrated luminosity of 3000 fb^{-1} could be foreseen by 2030. The increased luminosity will significantly increase the measurement precision of the Higgs boson properties. LHC at 13 TeV with 300 fb^{-1} of data is essential to firmly establish the major production mechanisms of a Higgs boson and its main bosonic and fermionic decay modes. This will also lead to about 100 MeV precision on the Higgs boson mass and the measurement of the boson spin. HL-LHC also provides unique capabilities to measure rare statistically limited SM decay modes, and increase the precision on the couplings with a high discovery potential for heavy Higgs bosons.

Searches for supersymmetry (SUSY) encompass a wide range of strategies aimed at different particles of the spectrum. The most generic searches assume that supersymmetric partners of the SM particle carry a conserved quantum number, called R-parity. If the lightest supersymmetric particle is neutral, it will typically be weakly interacting and will not be observed in a collider detector. Events are then characterized as containing several hadronic jets, associated with decay to the lightest particle plus missing transverse momentum. No significant excess of such events has yet been observed. The results of the searches are then parametrized by limits on the gluino mass and on a squark mass, assumed common to all squark flavors. Current LHC results exclude events up to gluino masses of 1.0 TeV and, independently, up to squark masses of 1.3 TeV. For the future stages of the LHC, we expect to be able to discover such events up to gluino masses of 1.9 TeV and squark masses of 2.3 TeV with 300 fb^{-1}, and up to 2.7 TeV with 3000 fb^{-1}. It is possible that the first signal of SUSY would not be given by the generic search just described, but would require a more specialized analysis. Special search techniques are needed in models in which mass gaps in the SUSY spectrum are relatively small so that hard jets are not emitted in particle decays, and models in which only the partners of top quarks, or perhaps only color-singlet supersymmetric particles, are produced at accessible energies.

3 E-Infrastructure

3.1 Current e-Infrastructure Activity

The Worldwide LHC Computing Grid [3] is a global collaboration linking grid infrastructures and computer centres worldwide to distribute, store and analyse the data generated by the LHC at CERN.

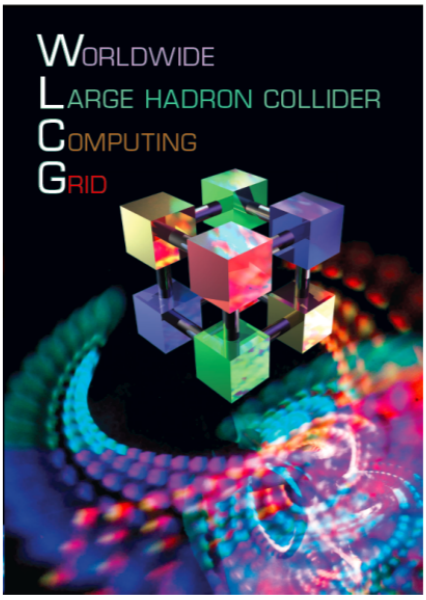

The WLCG is organized in three layers, or "Tiers", which are made up of computer centres which contribute to different aspects of the WLCG.

Tier-0 is one site: the CERN Computing Centre. All data from the LHC passes through this central hub but it provides less than 20% of the total compute capacity. It is connected to other major tiers and Grid services using dedicated 10 gigabits per second optical wide area links.

Tier-1 consists of eleven sites. These are located in Canada, France, Germany, Italy, the Netherlands, the Nordic countries, Spain, Taipei, the UK, and two sites in the USA. These sites provide distribution networks, processing of raw data, data analysis, and storage facilities.

Tier-2: there are around 140 sites covering most of the globe. Together, these sites provide approximately half of the capacity needed to process the LHC data.

Massive, multi-petabyte, storage systems and computing clusters with thousands of nodes connected by high speed networks are the building blocks of the WLCG centres. The

WLCG centres use specialised tools to manage the immense disk and magnetic tape mass storage systems needed for LHC data, and allow applications to access the data for simulation and analysis, independent of the storage medium (tape or disk) that the data resides on. WLCG also uses dedicated software components, called middleware, to link up the varying hardware resources across the Grid in a compatible way. This allows scientists and other users to access these resources in a uniform and secure way from anywhere in the world, turning the diverse and locally managed computing centres into a single massive virtual resource.

• European resources are integrated through EGI infrastructure, and use the Unified Middleware Distribution (UMD).

• The US contribution to WLCG is provided through the Open Science Grid (OSG).

Exchanging data between WLCG centres is managed by the Grid File Transfer Service that was developed by the EGEE project. This is the proven method for securely and reliably transferring large volumes of data across distributed computing grids. It has been tailored to support the special needs of grid computing, including authentication and confidentiality features, reliability and fault tolerance, and third party and partial file transfer.

3.2 Future e-Infrastructure Challenges

After the success of Run 1, ending in February 2013, the LHC is in its first long shutdown, LS1, extending through 2013 and 2014. The accelerator is scheduled to restart in April 2015 for Run 2, at a centre of mass energy of 13 TeV, where cross sections for interesting physics processes are factors of two or more above those at 8 TeV reached in the last part of Run 1. The corresponding requests of computing resources for 2014 and a discussion of requests for 2015 and beyond were recently reported to the LHC RRB[3], and are summarized below.

The anticipated stable LHC beam times for Run 2 will vary from 3 to 7 times 10^6 seconds per year, from 2015 to 2017. In 2015 the machine must be commissioned, while in 2017 it is assumed that LHC will be available for longer. The expected pileup (number of collisions in each beam-crossing) will increase from 25 (2015) to 40 (2017) for ATLAS and CMS for pp running. This is a crucial parameter since event sizes and reconstruction times increase as pileup increases, with concomitant increases in simulation times and sizes. The resources requests assume 25 ns running for Run 2 that reduces pileup by allowing fewer particles in each bunch for the same luminosity. The LHC luminosity is expected to be 10^{34} cm^{-2} s^{-1} in 2015 and will increase by 50% for 2016 and 2017. The increase in centre of mass energy for Run 2 leads to greater track multiplicity in each interaction, increasing event sizes and subsequent processing times.

A recent update of the computing models [4] has reviewed the expected evolution of resource requirements over the next years of LHC Run 2. The figures below show the evolution expected from 2014 to 2017. Both CPU power and Disk storage capacity are expected to double by 2017.

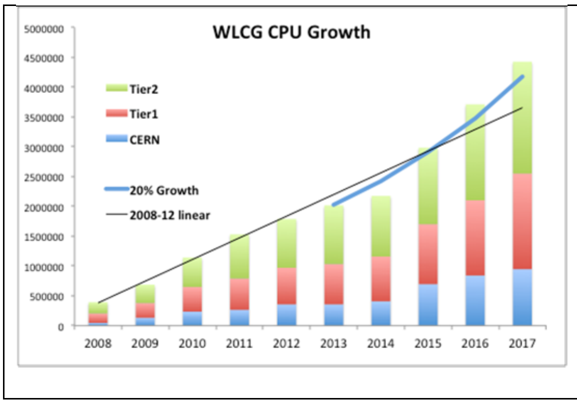

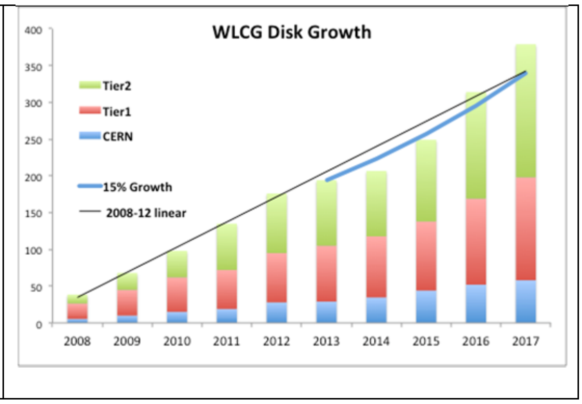

4 Future Plans

Along this last months a major review of the experiment and WLCG computing models has been undertaken, analysing the period of LHC Run 2, extending to 2017. The key points were summarised in a recent report to RRB[4] and are reproduced in what follows:

Regarding the Computing Model, all four experiments will take a more pragmatic view of which tasks should be run at the different Tiers compared to the original model, making use of the capabilities of the sites (capacities, connectivity, etc.) rather than a strict allocation of functions…for all experiments, there is a general assumption now that data can be accessed between any peers, rather in the original hierarchical model. This evolution had already begun in Run 1. For data management, all four experiments now use so-called data federations based on xrootd, which facilitates an optimised access of data from jobs, even if the data is initially remote from the site at which the job runs. Significant work has been invested in more intelligent data placement and caching, to optimise the number of files that need to be pre-placed or dynamically cached when (or just before) they are required. Data popularity services are introduced to better determine which data sets should be widely available and which data can be cleaned from disk caches...

Significant efforts are also being invested by the experiments and the WLCG applications area in order to improve the overall efficiency of use of modern multi-core CPU, and in future to be ready to make better use of parallelism and vector abilities of newer processors…

Regarding the Distributed Computing environment, there are several factors that are driving the simplification of the distributed computing environment. These are: the need to minimise the operational (staff) cost of running the grid sites; the need to be able to simply make use of opportunistic resources (clusters, clouds, HPC resources) with as little as possible set up and configuration; the need to reduce as far as possible the cost of maintaining complex grid middleware.

These factors are all leading towards a simplification of the grid middleware layer, which is realistic as the computing systems have evolved and complexity has moved towards the application layer. The ubiquitous use of pilot job frameworks have also helped simplify the system. At the same time new technologies such as open source cloud management software provides a natural way of implementing a simpler job management layer. An evolution towards a more cloud-like model for job management over the period of Run 2 is anticipated…

5 References & Bibliography

[1] All information on LHC and experiments compiled on November 2013 from http://home.web.cern.ch/about/accelerators/large-hadron-collider and from subsections in http://home.web.cern.ch/about/experiments . All images and text under CERN copyright.

[2] Updated European Strategy for Particle Physics (2013): http://council.web.cern.ch/council/en/EuropeanStrategy/ESBrochure-Strategy_Report2013.pdf

[3] CRSG Computing Resources Scrutiny Group, CERN–RRB–2013–094

[4] I Bird, P Buncic, F Carminati, M Cattaneo, P Clarke, I Fisk, J Harvey, B Kersevan, P Mato, R Mount and B Panzer-Steindel (Eds), Update of the Computing Models of the WLCG and the LHC Experiments, to be released in December 2013.

[5] I.Bird, S.Foffano, CERN-RRB-2013-089

[6] Extracted from Snowmass 2013 report: http://www-public.slac.stanford.edu/snowmass2013/SnowmassWorkingGroupReports.html