

# New strategies of the LHC experiments to meet the computing requirements of the HL-LHC era

---

**Dagmar Adamova\***

*NPI AS CR Prague/Rez*

*E-mail: [adamova@ujf.cas.cz](mailto:adamova@ujf.cas.cz)*

**Maarten Litmaath**

*CERN*

*E-mail: [Maarten.Litmaath@cern.ch](mailto:Maarten.Litmaath@cern.ch)*

The performance of the Large Hadron Collider (LHC) during the ongoing Run 2 is above expectations both concerning the delivered luminosity and the LHC live time. This resulted in a volume of data much larger than originally anticipated. Based on the current data production levels and the structure of the LHC experiment computing models, the estimates of the data production rates and resource needs were re-evaluated for the era leading into the High Luminosity LHC (HL-LHC), the Run 3 and Run 4 phases of LHC operation. It turns out that the raw data volume will grow  $\sim 10$  times by the HL-LHC era and the processing capacity needs will grow more than 60 times. While the growth of storage requirements might in principle be satisfied with a 20 per cent budget increase and technology advancements, there is a gap of a factor 6 to 10 between the needed and available computing resources.

The threat of a lack of computing and storage resources was present already in the beginning of Run 2, but could still be mitigated, e.g., by improvements in the experiment computing models and data processing software or utilization of various types of external computing resources. For the years to come, however, new strategies will be necessary to meet the huge increase in the resource requirements.

In contrast with the early days of the LHC Computing Grid (WLCG), the field of High Energy Physics (HEP) is no longer among the biggest producers of data. Currently the HEP data and processing needs are  $\sim 1$  per cent of the size of the largest industry problems. Also, HEP is no longer the only science with very large computing requirements.

In this contribution, we will present new strategies of the LHC experiments towards the era of the HL-LHC, that aim to bring together the desired requirements of the experiments and the capacities available for delivering physics results.

*55th International Winter Meeting on Nuclear Physics*

*23-27 January, 2017*

*Bormio, Italy*

---

\*Speaker.

## 1. Introduction

The year 2016, the second year of Run 2, was a memorable year for the LHC and CERN [1, 2]. LHC performance was above expectations: the design luminosity was delivered and then exceeded and the stable beams efficiency reached a new record.

The excellent performance of the LHC in 2016 resulted in increased demands for data storage and processing resources. The Worldwide LHC Computing Grid (WLCG [3]) and the computing systems of the LHC experiments were able to accommodate and deliver the requested data thanks to a continuous effort to implement all kinds of optimizations. These changes have zero or minimal impact on the physics output. There was a new record concerning the number of used cores within WLCG [4].

In the coming years, the experimental program of the LHC will focus on two topics: probing the Standard Model with increasing precision and searching for new physics. This program is about identification of rare signals in huge backgrounds (High Luminosity LHC: HL-LHC) and therefore it will require collecting and processing of significantly larger datasets than in Run 1 and Run 2 and it will also be dealing with increased levels of data complexity. This will result in even larger needs of resources and increased demands on the evolution of the whole chain of experiment software tools.

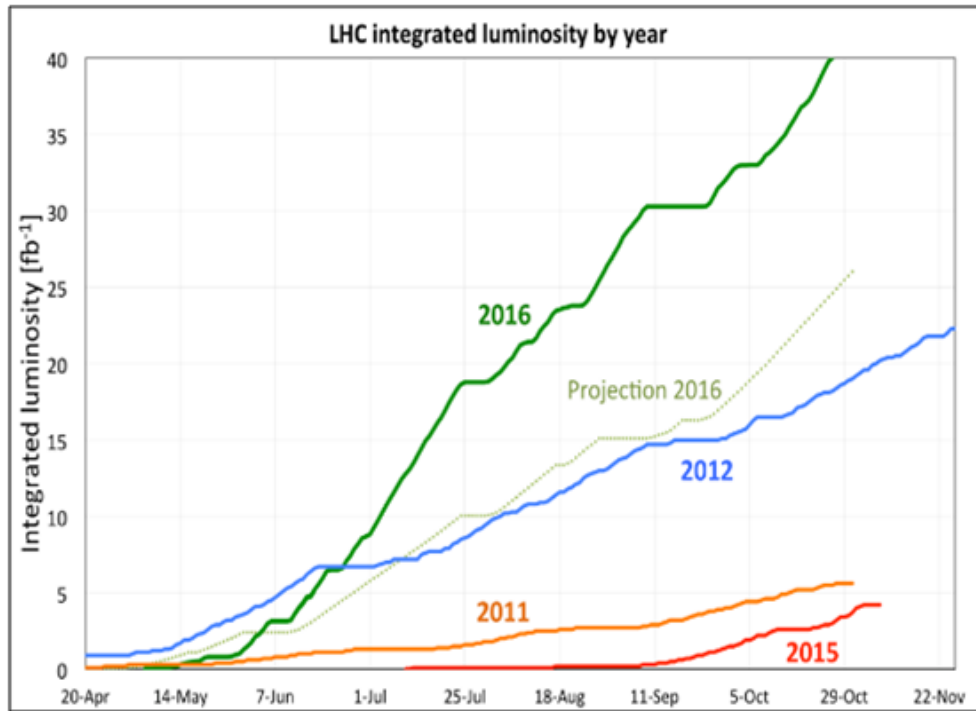
In this article, we will present the performance of WLCG and computing systems of the LHC experiments in 2016 and new strategies and mitigation measures being developed and implemented to allow the current WLCG to evolve toward dealing with future requirements on data collection and processing for the desired physics output.

## 2. Performance of LHC and LHC Computing grid (WLCG) in 2016


The second year of the Run 2 has set new records both in the performance of the LHC and in the capabilities of the WLCG to process and manage the unexpectedly high production of data. The LHC design luminosity was exceeded by about 40 per cent and reached  $L_{\text{peak}} = 1.4 \cdot 10^{34} \text{cm}^{-2}\text{s}^{-2}$ . The stable beams efficiency scored almost 60 per cent. Also very impressive was the level of integrated luminosity delivered by the LHC to the experiments [2]:  $\sim 40 \text{fb}^{-1}$  for ATLAS and CMS (see Fig. 1; 2016 target was  $25 \text{fb}^{-1}$ ).

The volume of raw data from the LHC experiments recorded in 2016 was  $\sim 50$  PB and of derived data (1 copy)  $\sim 80$  PB [5]. The graph on Fig. 2 shows progression in storing the LHC data at the CERN tape facility CASTOR [6]. It became the biggest physics data repository worldwide, with a total archive of about 180 PB and 500 million files. It is anticipated that approximately 100 PB of raw data will be stored for the remainder of Run 2, in 2017 and 2018.

The better than expected performance of the LHC in 2016 made the experiments look for ways to mitigate the increased demand on computing and storage resources. All experiments were able to find additional opportunistic CPU resources. For example, in 2016 the ATLAS experiment used 50 per cent more CPU resources than pledged, also thanks to access to High Performance Computing (HPC) centers and commercial clouds [5], see also Fig. 3. This required investments in software and computing tools to integrate non-Grid resources.




**Figure 1:** The graph shows the integrated luminosity delivered by the LHC to the ATLAS and CMS experiments in 2011, 2012, 2015 and 2016. There were more collisions in 2016 than in all previous years together.



**Figure 2:** Data stored in CASTOR since the beginning of the LHC data taking, in TB/month. The peak in 2016 stands out. The shares of the experiments in 2016 are approximately: ALICE 7.6 PB, ATLAS 17.4 PB, CMS 16.0 PB and LHCb 8.5 PB.

Other activities focused on improving the performance of the computing models and include, e.g., overlay of minimum bias events instead of simulated background, technology improvement like GPUs or KNL microarchitecture, optimization of algorithms and processing steps, use of fast simulation, reduction of data format sizes, parking of data for later processing [4]. Also, a new high-precision calibration schema was developed, dynamic disk cache management was used aggressively, data management and data access layer were optimized, the LHCb experiment has im-



**Figure 3:** Usage of external CPU resources by the ATLAS experiment in 2016.


plemented a first version of a multi-threaded framework and all the experiments actively used their HLT farms for offline processing.

In addition, all experiments performed a number of file deletion campaigns over their distributed storage and reduced the number of file replicas. Further reduction of the resource usage in Run 2 would pose risks for the physics output. But with all the mitigation measures implemented and considered by the experiments the data taken in Run 2 should be safely processed and there will be further development in resource saving arrangements in the future.

### 3. Worldwide LHC Computing Grid: Current status

The distribution and processing of the data produced by the LHC is managed by a distributed computing infrastructure, the Worldwide LHC Computing Grid (WLCG) [3]. Ever since its start-up in 2002, WLCG has been under steady development, extension and stress-testing and since the first LHC collisions it provided reliable processing, storage and access to the LHC data. It has a hierarchical tier-like structure with a central, largest Tier-0 center at CERN plus an extension at the Wigner Research Center for Physics, Budapest.

In the first decade of this century WLCG was the largest computing infrastructure worldwide, which changed after the massive rise of commercial computational clouds [7]. In 2017, WLCG involves 167 computing centers in 42 countries worldwide (see Fig. 4). It can provide up to 700 thousand computer cores and 985 PB of storage: 395 PB of disks and 590 PB of tapes. The network connectivity between CERN and the Tier-1 sites (currently 14) is managed within the LHC Optical Private Network project (LHCOPN [9]), where some of the links have a capacity of 100 Gb/s, while most have 10 - 40 Gb/s. The interconnectivity between Tier-1 and Tier-2 sites dedicated almost exclusively to the traffic of the LHC data is provided by the LHC Open Network Environment project (LHCONE [10]).




**Figure 4:** Map of the Worldwide LHC Computing Grid, showing individual computing sites, see [8].

WLCG has been a well-working and understood system providing all the required services for the LHC community. With the prospect of HL-LHC ahead a further evolution of the system and computing models is necessary which will be embedded in the new computing Technical Design Report (TDR) expected to appear in 2020.

#### 4. Challenges over the next years

The LHC delivering increased data volumes and levels of complexity in 2016 sharply highlighted the disparity between the anticipated requirements for storage, computing and network resources and what the LHC community can afford. The computing models of the experiments which were developed more than 10 years ago, before the start-up of the LHC, were based on extensive descriptions of types of available computing facilities, services which should be provided, how the software tools and computing systems should be built and used etc. In our post-Run-1 era any system which does not allow for changes and evolution in accordance with technology developments would not be usable for the HL-LHC computing.

During Run 1 and Run 2 the original computing models gradually evolved and incorporated changes under pressure from resource exhaustion, technical developments changing the market, as well as new physics needs. In preparation for HL-LHC (see also Fig. 5), a number of additional challenges and questions will arise. For example, what is the scalability of the current system? Is it possible to scale it up 10 times or more? Is the LHC community able to gain from new technologies like GPUs, KNL, AVX? Is it possible to optimize the models to decrease the overall cost of the resources? Should all resources be considered dynamic? Can the data processing, data and storage management be done on the level of events rather than jobs and files (which would enable more dynamic use of resources)? And also, should the possibility be considered that the



**Figure 5:** Very rough estimate of RAW data per year of running using a simple extrapolation of current data volumes scaled by the output rates (in PB, follow-up of 2015). To be added: derived data (ESD, AOD), simulation, user data ...

computing becomes commercially provisioned, would the funding agencies be willing to pay for the use of resources and services instead of putting them into local computing sites?

Other aspects to be clarified include: understanding of technology development, when advances over 2 - 3 years are difficult to foretell and 10-year predictions are impossible; evolution in networking and data streaming; future role of WLCG computing facilities.

There are many more challenges to be addressed and the high energy physics (HEP) community will provide a Community White Paper, a document setting a roadmap to a new Computing TDR for the HL-LHC era [11].

## 5. Technical Design Report for the HL-LHC computing

The final version of the Technical Design Report for the HL-LHC computing is planned to appear in 2020. As the original TDR was issued in 2005, the changes in the technologies, the computing and the physics environment since that time are enormous. The new TDR should define strategies for evolution toward the computing models of the future, plans for further development of the existing WLCG infrastructure and, generally, a blueprint to try and guarantee the necessary means for management and processing of the LHC data and delivery of physics results.

In the meantime, the work towards the HL-LHC TDR is ongoing in many areas and is building on the Community White Paper of the HEP Software Foundation (HSF) [12]. The HSF started 3 years ago to coordinate common efforts of HEP software and computing groups worldwide.



According to the many topics to be scrutinized, a number of working groups (WG) have been set up and as a first step they formulated the challenges and goals to be reached. Currently there are the following WGs: Computing models; Physics generators; Math libraries; Detector simulation; Software trigger and event reconstruction; Data access and management; Event processing frameworks; Conditions database; Security and access control; Machine learning; Data analysis; Workflow and resource management; Data and software preservation; Careers, staffing, and training. The programs of all WGs can be accessed from [11].

## 6. Summary and Outlook

The LHC physics discoveries and measurements are only possible through careful processing of massive amounts of data. Therefore already in 2001 the concept of a distributed computing infrastructure for management and processing of the LHC data was first designed. In the following years the LHC community were pioneers in building such a kind of data processing system.

The continuous breathtaking development of computing and information technologies towards the end of the first decade of this century resulted in the establishment and fast evolution of commercial computing facilities of a distributed character, the commercial clouds. The amount of data processed in such clouds quickly exceeded what is processed within WLCG and also the technology leadership moved into that new area.

In the WLCG TDR from 2005, the estimate of data taken in one year of LHC operations was  $\sim 20$  PB, which was exceeded already at the end of Run 1 and more than doubled in 2016. This disparity between the real situation and what was anticipated in that TDR forced WLCG and the LHC experiments to adapt on the way and mitigate the lack of affordable resources. The LHC community demonstrated great skills and flexibility and despite all the constraints the experiments have been able to deliver timely physics output so far.

At present, it is impossible to make predictions concerning the development of technologies for more than 2 - 3 years ahead. Still, the community needs a research and development project to lay out a concept for the HL-LHC computing. There is ongoing work in many areas, coordinated through the HEP Software Foundation, where the community tries to define the challenges and find solutions. Due to the great uncertainty concerning future technology developments, there will be thrills not only in the search for new physics, but also in the future of LHC computing!

**Acknowledgement:** This work has been supported by the grant LG 15052 of the Ministry of Education of the Czech Republic.

## References

- [1] The Large Hadron Collider at CERN; <http://lh.web.cern.ch/lhc/>
- [2] G. Iadarola : Accelerator performance and plans; CERN Council Open Session, CERN, 16.12.2016, Worldwide LHC Computing Grid: <http://wlcg.web.cern.ch/>
- [3] Worldwide LHC Computing Grid: <http://wlcg.web.cern.ch/>
- [4] Executive summary of the 129th meeting of the LHC Committee, February 2017, <https://cds.cern.ch/record/2253239/files/>
- [5] I.Bird: Status of the WLCG project, including Financial Status; LHC RRB Meeting April 2017, CERN.
- [6] CASTOR, <http://castor.web.cern.ch/>
- [7] The NIST Definition of Cloud Computing (Technical report). National Institute of Standards and Technology: U.S. Department of Commerce. doi:10.6028/NIST.SP.800-145. Special publication 800-145.
- [8] WLCG Public Page, <http://wlcg-public.web.cern.ch/>
- [9] Large Hadron Collider Optical Private Network , <https://twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome>
- [10] Large Hadron Collider Open Network Environment , <https://twiki.cern.ch/twiki/bin/view/LHCONE/WebHome>
- [11] Community White Paper (CWP), <http://hepsoftwarefoundation.org/cwp/cwp-working-groups.html>
- [12] The HEP Software Foundation (HSF) , <http://hepsoftwarefoundation.org/index.html>