

Common challenges for HPC integration into LHC computing ¹

Motivation

With the detector upgrades for ALICE and LHCb in Run3 (2021-2023) and the HL-LHC accelerator upgrades in Run4 (2026-2028), which will increase the data rates in ATLAS and CMS, the LHC experiments are facing unprecedented computing challenges in the near future. With no changes to the computing models of the experiments the resource gap would be at least a factor of 10 over the gains expected from technology alone and a constant budget: the experiments are working on advances to the software and operating models that will improve the situation, still significant additional processing and storage resources will be needed on the time scale of Run3 and Run4.

High Performance Computing Centers (HPC) are some of the largest processing resources accessible to science applications. They are centers of expertise for computing at large scale with low latency local networking. The efficient usage of HPC facilities may provide a substantial contribution to the success of future LHC data processing by providing much needed computing capacity. R&D investigations are being performed in order to harness the power provided by these facilities and evolve the experiments' computing models to include their usage. HPC facilities represent a unique challenge and opportunity as they are early adopters of technology including heterogenous accelerated computing architectures.

The experiments have compiled HPC-related documents, including the summary of a joint meeting on this subject [1][2][3]. This document intends to extract the commonalities between experiments with the aim of developing a joint roadmap and strategy for enabling the exploitation of HPC resources. To develop common approaches between experiments and HPC sites, a foundation and understanding of the problems is needed. This is built on a summary of technical challenges, described in section 2. They are broken into two main categories: computing resource challenges and software and architecture challenges. Computing resource challenges describe issues related to operations, facility access, provisioning, and monitoring; while software and architecture challenges are related to adapting HEP applications to make effective use of alternative architectures often found on HPC. In order to explore potential solutions a number of pilot demonstrators are proposed in section 3 below.

1. Status

All LHC experiments report using some HPC resources with varying degrees of both success and technical difficulty. Accessing HPC sites with both workflows and data needs a level of customization. Development of applications for HPC centers has been more successful when the site architectures are the most similar to the generic x86 systems used on the WLCG. Some LHC experiments have successfully used KNL processors at CORI at NERSC. There has been

¹ **WLCG-MB-2019-01** October 2019. Author: Maria Girone (maria.girone@cern.ch).

application development on ARM and PowerPC processors. The MIRA facility at Argonne was used by ATLAS for a specific event generation program ported to Power8. Details of the progress and workflows executed can be found in [1].

No LHC experiment has successfully demonstrated efficient use of GPUs at scale on HPC sites. Five of the ten largest HPC centers on the Top500 list derive the bulk of their processing capacity from GPU accelerators.

2. HPC Challenges

There is significant overlap in the challenges facing the LHC experiments as they try to exploit the capabilities of HPC centers. The technical experience is more explicitly described in the experiment specific documents [1],[2],[3].

2.1 Computing Resource Challenges

Data Processing at HPC: HEP workflows deal primarily with data. Simulation workflows produce data, while reconstruction and analysis workflows process input data and generate sometimes large output samples. HPC centers are primarily designed for large scale, tightly coupled calculations where there are thousands of high-performance interprocessor communication channels open, storage is more transient and designed for smaller final results. Cybersecurity is often tighter at HPC sites than at smaller centers, which results in stricter firewall rules and less opportunity for external data access.

Storage at HPC centers is architected to ingest final data products, serve input samples and record checkpoints of large-scale applications, but is rarely designed to handle the level of IO expected of high throughput computing on a large number of processors.

One of the accomplishments of the HEP community is a set of effective tools for large data organization, management and access (DOMA) [4]. If HPC centers are going to evolve to be world class data processing facilities, significant work is needed in organised data movement, real-time data delivery and scalable storage.

Wide area networking: Due to their size and often heightened security, network capacity and functionality need to be properly provisioned. HPC sites are not generally expected to host data samples persistently, so the incoming and outgoing network needs to have sufficient capacity to accept and export multi-petabyte datasets on the timescale of a few days. A network capable of reaching the NREN networks of several hundred gigabit per second would be sufficient depending on the size of the HPC.

Local networking: The LHC experiments are using pilot workflows, where work is requested from a central queue when the resources are available. For this to function most require outgoing connectivity from the worker node, at least sufficient to be able to contact the partner-site batch system. Depending on the pattern used, some services may need to be run at the edge of the HPC site, either directly or containerised (unprivileged). Depending on the technology used and trust relationship with the partner site, the partner site could also operate these components remotely.

Site Interfaces: HPC centers are in general rather heterogeneous. There are important variations in the methods to access them, the interfaces to the workload management systems, the allocation policies, the availability of services, the access and types of local storage and the network connectivity of the resources to the outside world. The differences can mean that each site needs dedicated experiment effort to interface and adjust to successfully execute its workflows, and for each experiment using the same site this dedicated effort might be needed again. One of the successes of WLCG has been the adoption of a common set of protocols and policies for interfacing to site specific services and resource discovery. Similar tools and interfaces that can scale to the needs of HPC sites are needed. Currently the experiments are using a traditional CE if it's available, but have fallen back to ssh interfaces for some sites [2],[3]. Local "pilot factories" can alleviate some of the authentication load by generating the pilot jobs internally with a service. This model scales better and can give some more transparency to the receiving site.

Equally important to how external services access the site is how internal workflows communicate with the external world. As LHC computing is a globally distributed infrastructure, most experiments request some amount of outgoing network access from the running processes. Pilots call out for workflows and workflows call out for conditions and data.

Runtime Environment and Containers: The LHC experiments have relied on a consistent software and hardware environment from early in Run1. CVMFS [5] is used for all software distribution. Either HPC centers could support CVMFS centrally (and some have) or solutions to support lightweight and dynamically deployable versions of the existing infrastructure are needed. All experiments are investigating methods to completely bundle workflows into containers stored in CVMFS which isolates the running payload from the underlying operating system provided by the site. Common approaches are being coordinated within WLCG - CMS and Atlas now use Singularity containers served from CVMFS in production.

Monitoring services: they help to tie hundreds of independent processing and storage facilities together and allow them to function as a coherent resource for LHC computing. The Site Availability Monitors (SAM) were critical during the commissioning phase of WLCG. Similar services will be needed to monitor the health and availability of the interfaces and services at HPC sites: the infrastructure that accepts processing requests and data transfer infrastructure. It will be also necessary to expose internal monitoring information on the current activity level at each site, the CPU/GPU efficiency, and the IO rates. While most WLCG sites are dedicated facilities and expected to be continuously available, the HPC centers may be opportunistically accessed or used for a defined period based on an allocation. The monitoring services will need to be aware of site resource scheduling.

Accounting/benchmarking services: they are used by all the experiments to verify the accuracy of resource planning, to verify that pledges are met, and to properly account site contributions. HPC sites represent an accounting challenge because the normal metrics of HEPSpec06 hours delivered are less meaningful on heterogeneous architectures, in particular do not benchmark the accelerator and GPU components. Development effort is ongoing [6] and is needed to establish metrics for all compute resources available for LHC: 'standard' multi-core CPUs, GPUs, FPGAs, and any other highly parallel architecture. The metrics under consideration are based on the experiments' workloads that can run on GPUs and/or highly parallel architectures. As we introduce more heterogeneity from the

contributing sites, more work is needed to properly account the different contributions. The experiments' workflow management systems maintain information about events processed, data produced, and job exit codes so detailed throughput of all participating sites can be extracted. Equally, if, as suggested below, batch integration is provided via a partner site, the existing accounting mechanisms of that partner site may be used to report the HPC site's contributions to the WLCG.

Authentication: The LHC collaborations have thousands of active submitting users and even the central operations activities are performed by large teams. A Virtual organization (VO) support infrastructure was a requirement and development priority from the beginning. One of the early security decisions was to adopt the "Trust the VO" model. Sites supporting a VO agree to use an external service to establish membership and authenticate. If there is suspicious activity, the VO is trusted to deal with it centrally and restrict the user. The model of pilot job submission means that the site may only see generic VO credentials, without being able to authenticate or identify the actual user making the submission. The VO is trusted to have sufficient logging to identify the source of any suspicious workflow. HPC sites may have much stricter requirements on authentication. During the demonstration phase, the needs for stronger authentication have been handled by having individuals associated with access credentials and made responsible for all work, but this is difficult to scale. In the case where the "trust the VO" model, together with a partner site, is insufficient, development work would be needed for authentication services that meet the site requirements for security and still allow the experiments to distribute the operations workflows across large operations teams. Historically, users have been authenticated by generated short lived proxy certificates from X.509 end-user certificates issued by IGTF (<https://www.igtf.net>) Certificate Authorities. Work is ongoing to enable OpenId Connect Authentication, with VOs acting as the OpenId Provider (token issuer).

Authorization: The authorization needs of the LHC experiments have traditionally been modest. The initial infrastructure was designed with fine grained authorization infrastructure for roles and groups. However, in reality, the dominant authorization use-case is the access provided by being a collaboration member. There is a selected group of people and services that have authority to create and delete data in centrally controlled areas, which most collaboration members are restricted to reading. Most of the authorization needs have been simple enough to be satisfied by basic unix user and group permissions. Historically, users have been authorised by registering to VOMS (the Virtual Organisation Membership Service), which in turn enables users to add VOMS extensions containing group and role information to their proxy certificates. Work is ongoing to enable OAuth2.0 Authorisation, both as groups and as capabilities (i.e. authorization to perform a specific action).

Provisioning: Modern HEP experiments are normally in preparation for at least a decade and the LHC experiments will operate for several decades. The current arrangements between the LHC experiments and many of the supporting computing facilities have been in place for more than fifteen years. The experiments rely on the predictability of resources through a detailed planning and pledging process that extends several years into the future. For HEP collaborations to be able to rely on HPC centres for critical workflows it will be necessary to have an allocation and award process that does not have the uncertainty of annual independent proposals. In the absence of a longer-term arrangement there may be additional interesting physics that could be processed with an independent allocation or lower priority workflows that could be scheduled opportunistically, but workflows critical to the core mission would need to remain on more predictable dedicated resources.

The specific requirements needed (or desired) for the experiment payloads are discussed in detail in [1] [2] [3].

2.2 Software and Architecture Challenges

A number of HPC sites have decided to adopt specialized accelerated HW architectures, which contribute the majority of the TFlop/s performance of those sites. These specialized systems can have excellent performance on applications optimized to use them. HPC sites will often only accept applications that take advantage of such capabilities. Some of the most powerful supercomputers will not be available to HEP applications until the code can be rewritten for common accelerators or methods to interface the existing applications with other architectures are developed.

All HEP experiments are engaged in R&D activities to reengineer their software to use accelerated hardware architectures, in particular GPUs. Solutions are being pursued across the entire processing pipeline to close the resource gap expected at the HL-LHC [7]. Several experiments are assessing their software frameworks to be more modular and support the offloading of suitable parts of their workload to accelerators (GPUs/FPGAs). Full detector simulation is one of the largest consumers of processing resources. The focus has been on parallelization and vectorization of the existing code, and while progress has been made there is much more to do. There is a pilot project on exascale GEANT, which includes GPU optimisation for full detector simulation, but faster simulation (based on combining parameterized approaches with full simulation, which may also include Deep Learning techniques[7]) are being investigated now. Event reconstruction algorithms offloaded to GPUs is an active R&D area for all the experiments. The load of event generation varies by experiment, but some event generation packages are candidates for GPU acceleration. More details on the software R&D efforts can be found in [1],[7].

HPC centers may provide expert support and training opportunities in porting applications and making the most efficient use of the hardware; we hope that we can engage these teams in a mutually beneficial collaboration. High performance GPUs have a sufficient market share of both the HPC community and the community that is developing tools for Big Data and Machine learning. Therefore GPUs will be a large share of the available computing resources for the foreseeable future. Changes in technology require planning and effort, but have been executed many times in HEP, as the landscape of hardware architectures and relevant programming languages have evolved.

All experiments are investigating machine learning applications to augment or replace some traditional workflows. The advantage of exploring ML applications on HPC system is that industry and the open source community have invested heavily in optimizing support libraries for accelerated hardware, in particular GPUs. Any workflow where the bulk of the computing time is in the training and inference steps and can therefore use standard ML frameworks is already well positioned. Machine learning is still in the early stages of development in the experiments and most of the current activities are exploratory and can be accommodated with a few systems. If other sciences are any indication, this activity will increase in HEP and it may be a dominant use case for HPC.

3 Outlook

This document summarizes the common challenges identified to integrating HPC centres into the WLCG computing system.. Technical challenges, limitations, and obstacles are identified. In order to overcome the challenges a series of technical demonstrators will be needed.

Access to small scale testbed environments to port and validate the experiments' software on the architectures deployed by the HPC centers would be very useful. Site interfaces, data access techniques, monitoring and accounting services, authorization and authentication systems will all need to be demonstrated at scale.

References

[1] <https://indico.cern.ch/event/811997/attachments/1862943/3062278/HPCLHCC.docx.pdf> (Cross-experiment HPC workshop)

[2] https://docs.google.com/document/d/1Wu_jvizQASdsfxEi3Twe2n_QQWC38wsZYvDv-sReU14/edit?usp=sharing (HPC-LHCb doc)

[3] (https://cds.cern.ch/record/2707937/files/NOTE2020_003.pdf) (CMS NOTE-2020/003)
https://cds.cern.ch/record/2707936/files/NOTE2020_002.pdf (CMS NOTE-2020/002)

[4] DOMA <https://indico.cern.ch/event/587955/contributions/2936867/>

[5] CVMFS J Blomer et al.; "Delivering LHC Software to HPC Compute Elements with CernVM-FS", In: Kunkel et al. (eds) High Performance Computing. ISC High Performance 2017. Lecture Notes in Computer Science, vol 10524. Springer. https://link.springer.com/chapter/10.1007/978-3-319-67630-2_52

[6] https://indico.cern.ch/event/739878/contributions/3380141/attachments/1840586/3017523/GDB_8_5_19_giordano.pdf

[7] <https://hepsoftwarefoundation.org/cwp-whitepapers.html>