Evolution of Scientific Computing in the next decade: HEP and beyond

Introduction

High Energy Physics (HEP) has demonstrated a unique capability with the global computing infrastructure for LHC, achieving the management of data at the many-hundred-Petabyte scale, and providing access to the entire community in a straightforward manner. This is still a unique facility in science, but aspects of it are more and more needed for other scientific endeavours. Fifteen years ago, when the work started, there were no examples of how to build such a system, and no experience from industry or others. If we were to design the system today, of course, we would benefit from the tools and expertise of the global internet companies. HEP has a challenge for the foreseeable future – which is how to achieve a scale of computing and data management that is orders of magnitude greater than that of today, while maintaining a reasonable cost envelope. HL-LHC is the most immediate such challenge, but we also have other high-rate experiments, and future potential facilities that must be considered. The desire is to build on the existing global structures built for LHC, leverage the experience and capabilities now available in the internet sector, and plan to evolve a HEPwide scientific data and computing environment for the future of our field. Importantly, in addition, we observe similar needs arising in related fields (astronomy, astro-particle) with many of the HEP facilities often directly involved. In planning for the future, we must take into account compatibility and synergy at the facility level. Taking the success of the Worldwide LHC Computing Grid (WLCG) as a starting point it could be envisaged to evolve the infrastructure and tools as a basis for computing for HEP for the coming years, while challenging the concerns of cost (both in terms of equipment and operationally), organization, and community needs. Can we build on what we have today, evolving and providing tools and infrastructure services to be used or adapted to future needs of the HEP community?

S.W.O.T analysis of the current WLCG Organization

Strengths: the WLCG organization has demonstrated and implemented a distributed computing model for the LHC experiments, which has played a crucial role in their scientific mission. The WLCG community is strongly established and a network of trust exists between stakeholders. This allows the implementation of a very lightweight decision making process, based on consensus at various levels of the organization; particularly at the Overview Board, the Management Board and the Grid Deployment board. The infrastructure is also strongly consolidated and builds on common middleware blocks, such as services and clients, and policies in matters of security, identity management, resource sharing, monitoring and accounting.

Weaknesses: the WLCG organization has a strong focus on computing services, infrastructure and policies.

The importance of application software was recognized from the very beginning and an LCG application area was established in the WLCG organization. The Architects Forum holds the responsibility to coordinate the development of common application software and reports to the WLCG Management Board. The activity of the Architects Forum has however decreases in the last years, as the stack consolidated around well established components. Preparing for scientific computing in the 2020s, application software will play a critical role. In fact, innovative solutions will need to be considered and implemented: experts with different skill-sets, such as parallel programming and machine learning, will need to complement the more physics-oriented expertise today available in our community. In addition, improving software performance will be a critical aspect of the way to reduce the cost of computing in the future. The absence of an initiative with a strong focus on software weakens the overall strategy addressing the future needs.

From the point of view of computing services, WLCG has little development effort, as this has been historically driven by middleware initiatives, sponsored by the Funding Agencies and the European Union. WLCG has historically integrated solutions from various sources and built a coherent infrastructure. The absence of internal development resources and a reliance on external providers requires a mechanism to identify the future needs and influence the directions of the various development efforts. Such a mechanism, which should emerge from the strategic discussion in the project is only partially effective, with the consequence that several funded middleware initiatives have focused on marginal aspects of the WLCG needs.

Threats: the funding for scientific computing worldwide is unlikely to increase. Already during LHC Run-2 (2015-2018) the funding agencies of the LHC experiments made clear that funding beyond a constant budget should not be expected during Run-2 and beyond. This constraint is purely economical and has no grounding in terms of computing needs for science. Such a level of flat funding enable the provisioning of enough resources for physics at Run-2 and we expect it to be also adequate for Run-3, but will unlikely be sufficient for HL-LHC unless a major evolution in the application software, the computing services and the infrastructure happens. Without such an evolution and because of the flat funding constraints, WLCG would not be able to cope with the needs of the LHC experiments.

WLCG has so far been the major player among high energy physics and many other sciences in terms of data volume and compute capacity. This allowed it to steer the evolution of the infrastructure and services in the direction of the LHC experiment needs. In the 2020s we expect other high energy physics experiments, such as DUNE, and other sciences, such as astronomy (e.g. the SKA organization) to require a similar level of resources to LHC. In order to maximise the return on investment of the Funding Agencies it would be advantageous to foresee a common infrastructure and set of tools serving the needs of the set of sciences they support. WLCG might therefore lose influence in the process of shaping such infrastructure.

Finally, some countries have invested and intend continue to to invest capital in computing facilities that historically were not particularly suited for HEP, such as High Performance

Computers. While in the past HEP was able to downplay the reliance on the use of such resources and obtain processing capacity in conventional High Throughput systems, we expect this to change in the next years. The provisioning of resources at HPC facilities comes with challenges at the level of application software, data access, as well as access policies and resource scheduling.

Opportunities: WLCG has the opportunity to leverage its strengths and play a central role in the evolution of scientific computing, as it is the community with the largest experience in large scale distributed scientific computing. While the infrastructure will inevitably need to diversify to accommodate the needs of other major players, it is expected that the WLCG model is taken as reference for the future. Changing the governance model and splitting the roles of WLCG as a project from the role of WLCG as an organization and infrastructure will allow to open up the latter to more communities in and outside HEP, create consensus among such sciences. It will also allow WLCG to retain the current aspects of its governance which demonstrated so far their effectiveness and those aspects might be cloned by other sciences if they wish.

The scientific computing evolution strategy articulates around three pillars: leverage the existing HEP computing infrastructure and evolve it to serve as common computing system for HEP and sciences beyond HEP; evolve the facilities and services to build a HEP Data Cloud[1]; invest in common software and software techniques, including training, dissemination and recognition. As part of the process in achieving such a strategy we will propose an evolution of the current WLCG governance: factorize the current WLCG organization into a project specific to LHC needs and an collaboration which serves a common infrastructure for HEP and sciences beyond HEP. The next sections will elaborate the details of these ideas

1. General Infrastructure

The fundamental components of a HEP infrastructure are the essential building blocks of the computing, and probably today one of the major successes of WLCG. For LHC we have the computing resources at close to 200 sites using this infrastructure. Baseline grid services are supported by mature monitoring, operational and support processes and teams, including worldwide collaboration on security and incident response. The “grid” that enables the coherent use of those resources must evolve over the coming years, and be capable of supporting continually evolving computing models, and being agile to technology changes.

We have in place global networking infrastructures, not only those provided by the National Research and Education Networks (NRENs) and their coordinating bodies, but HEP-specific structures such as the private LHC Optical Private Network (LHCOPN), and the very successful LHC Open Network Exchange (LHCONE) overlay network, which provide the ease of

---

[1] Here we refer to a “Data Cloud”, also often described as a “Data Lake”. The intention is the same - a distributed data repository serving data to compute resources and clients.

management and connectivity that will be essential for the future. Today this is already used by more than the LHC experiments. While the LHCOPN is a private network, it is nevertheless a good model for specific situations in the future – as will be discussed below. Network resources, thoughts as the main limiting factor at WLCG inception time, emerged as probably the most solid one, both in terms of capacity growth and reliability. They will play a central role in the evolution of the infrastructure.

We also have a global Authentication, Authorisation, Accounting (AAA) service, and associated identity management, trust and policy networks. This is extremely valuable, and very unique. However, it is clear that the X.509 underpinnings are too specific to our infrastructure and are not the best for the future as they diverge from the widely adopted open source trends. Federated identity mechanisms (e.g. eduGAINii) and token-based authentication are being introduced as core components of the AAA infrastructure.

The WLCG infrastructure already integrates heterogeneous resources such as commercial or academic clouds, clouds, HPC facilities and volunteer computing. Such resources will play a more significant role in the future and further harmonization in their adoption is a key element of the strategy. Significant effort has been invested in developing sophisticated data management tools to deal with the processing and distribution of the huge volumes of data. One of the concerns today for LHC is the cost of the computing resources, and in particular the cost of storage which accounts for close to 70% of the overall hardware cost. That cost is in large part driven by the need to distribute data globally, with consequently many copies of the data, and consequent costs both in storage and in operations. For the future, we must consider an alternative model to reduce those costs, such as the HEP DataCloud described below.

The WLCG infrastructure is recognized as being of high value, and we have seen other HEP experiments, such as Belle-II and non LHC CERN recognised experiments, asking to be able to benefit from it already. The needs of such experiments are inline with the priorities and the evolutions strategy of the infrastructure and there is therefore a strong motivation for it to be a shared resource. We should note that this paper is not proposing to use the same resources for all experiments, but rather to try and use the same infrastructure, tools, software, and support as far as possible so that new projects are easier to support on existing facilities. Of course, this helps opportunistic use and sharing, but does not impose it.

2. HEP DataCloud

Because the majority of the cost in WLCG, both in terms of hardware and operational personpower, is on storage and data management, these aspects deserve special attention in when defining a strategy for the future. The currently envisaged model builds on the experience of large commercial cloud providers, as well as the LHC expertise in many-hundred-Petabyte scale data management. The idea is to connect the large HEP data centres between themselves with a dedicated and private multi-Tb/s network. This “virtual data centre” would store all of an experiment’s data, and by policy replicate it between the data centres. This network would probably be managed with Software Defined Networking (SDN) to ensure the data flows and reliability. In this way, we would achieve reliability, and availability. Into this data cloud we would plug compute resources. These resources may be co-located at the data centres, or may be other facilities, such as commercial centres and other large-scale HEP-owned resources. The model also allows for inclusion of commercially procured storage. Policy would prevent reliance on those for non-reproducible data sets, and should be redundant enough that a commercial centre could “unplug” without loss of data. This clearly relies on a very strong collaboration with the networking community, with adequate policies and capabilities to agilely connect to commercial partners.

A key concept in this vision is that data can be processed directly inside the data cloud or externally through a content delivery system, minimizing the possible impacts due to network latency or capacity. In the LHC case all reconstructed data beside the final analysis sets would be kept in the cloud. Having all of the data virtually co-located in this way may open the way to radically new analysis models, and strongly supports today’s models such as analysis trains. It would also permit the increased use of economical high latency media, such as tapes, as an active store for organized analysis, again helping with cost. This type of model also allows cost optimization through the use of hybrid centres: HEP owning compute resources at a level that is guaranteed to be fully used is very cost effective, and supplementing this with elastically provisioned resources as needed, presumably with some form of cheap spot-market pricing. This would allow an agile control of the cost, and can evolve as the commercial markets evolve.

Building and operating a HEP DataCloud may require new funding models and management methods and several aspects will need to be investigated and clarified: the capability to easily procure commercial resources at large enough scale to get economy together with the political implication of purchasing from the largest cloud vendors. At the moment, it is clear that the real cost efficiency and elasticity require the use of spot-market style pricing and what are the implications for the procurement process need to be understood. The data centres we address here are at the scale of a few Mega Watts, much less than large scale commercial centres, but larger than most University solutions. Today we have many Tier-2 centres that also provide significant resources, together with other opportunistic resources. We should distinguish in terms of scale rather than role. Large scale centres could participate as “compute plug ins” in the above model, while others primarily provide simulation resources. It should be remembered that some 50% of LHC compute loads is simulation, and the same will be true for forthcoming large experiments. Depending on the type of resource, some centres may be best suited to specific types of workload (e.g. HPC for event generation). This model will also easily accommodate the anticipated case of a funding agency moving academic computing internally to cloud credits rather than in-house facilities. The data cloud model also may be very interesting for other sciences, e.g. SKA regional centres, as it provides resiliency and long term preservation capabilities. Scaleout is also inherent to the concept, although for practical reasons we might imagine loosely coupled US, European, and Asia-Pacific instances.

3. Software

Software, for the future of HEP, is as important as the infrastructure itself. Software and infrastructure must be considered together as the separation is a source of inefficiency and thus cost. WLCG started a Cost Model working group to understand the interplays between software, workflows and infrastructure and optimize the overall spending, including hardware, staffing and operational cost.

The HSF Community White Paper identified key areas in the software domain which should be first priority in the future strategy of scientific computing with particular impact on WLCG on the timescale of HL-LHC. The software frameworks and algorithms of the WLCG experiments were designed many years ago and today can not leverage efficiently features and architectures of the modern hardware (e.g. vectorization and use of accelerators). Modernising the software in this direction requires skill-sets not broadly available in the HEP community. Building such know-how requires dedicated effort in terms of training. It also needs the right form of recognition in terms of career opportunity of the software developers for which HEP is hardly competitive with industry and therefore has a problem in retaining expertise. It also requires a set of tools and procedures facilitating the process, such as elements of the build systems, tools for documentation and advising on licensing among several others.

While part of this work is experiment and community specific, a large part of it can achieved through a common effort, at different level: from common tools and procedures, to sharing methodologies, to actual software libraries shared by multiple experiments. We proposed in the section above the idea of a common infrastructure for scientific computing, intended as common set of tools, services and support for experiments to use, with no imposition on a particular architecture. Here we propose the same model to be applied to software.

It is essential that this is recognised and supported. We have made significant progress in this area by setting up the HEP Software Foundation (HSF). It must be understood that there is no one-size-fits-all, but rather we need a community wide coordination of available tools covering the full stack from general workflow and data management tools to the application level. Common tools and libraries that can be used to build up the needs of an experiment are required. The HSF has made steps in this direction, covering many of the aspects of software for the HEP community, with the aim of collecting a set of tools contributed, developed, maintained and evolved by the community. It is also a mechanism for pursuing common R&D efforts in software, and for coordinating things like technology tracking, and developing software tooling for development and performance analysis. While the HSF is a good framework, this does not remove the need to engage appropriate levels of investment in this area. HEP must recognise that software efficiency and performance will be key to maintaining an affordable infrastructure. We must get ourselves into the position of being able to evolve our codes to make efficient and best use of the evolving computer architectures. This is not a one-off effort but will require sufficient and on-going investment in people and skill development and retention.

Scientific Computing Infrastructure governance: organization and steering

We propose to evolve the existing computing infrastructure for LHC into one for the entire HEP community: the **Scientific Computing Infrastructure** (SCI). For such a process to happen we

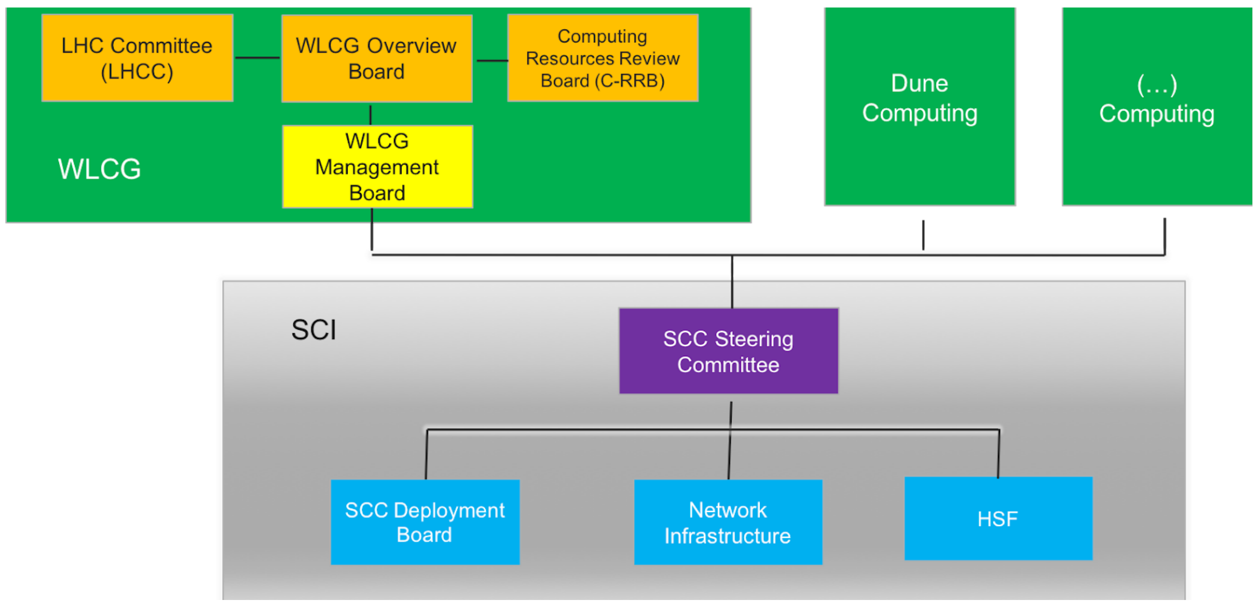

clearly need the buy-in of all of the major stakeholders, and the community itself. What works best is a lightweight steering mechanism rather than strong governance. The experience in setting up the HSF along those lines is very clear as an example of effectiveness. The WLCG collaboration, while structured with a more formal governance and decision making process, agreed in its Memorandum Of Understanding, de facto also reaches decisions by consensus. The proposed organization of the Scientific Computing Infrastructure and its interplay with the major computing projects of which it is part of are visualized in Fig.1.

From the current WLCG Organization, the **WLCG project** would factor out the aspects specific to the LHC experiments. It would also manage the aspects of the infrastructure specific to LHC such as the Optical Private Network. The WLCG experiments would continue to negotiate with their funding agencies for both pledged resources and access to opportunistic ones. The WLCG project would conserve its Memorandum Of Understanding (possibly amended with the modifications proposed here) and continue with the current resource management process: the **WLCG Overview Board** would oversee the functioning of the WLCG project and its role in the SCI; the project would report and respond to the **Computing Resources Review Board** about legal and resource matters, and to the **LHC Committee** concerning the technical matters and scientific aspects. The **WLCG Project Leader,** appointed by the CERN Director General in consultation with the Overview Board, would continue representing the project in front of the above mentioned committees, the SCI and the outside. The **WLCG Management Board,** chaired by the WLCG project leader would retain its role of managing the day-by-day aspects of the project and represent the WLCG interests to the SCI.

HEP experiments or projects others than WLCG would have their own Computing Management and Organization (see the example of DUNE in Fig. 1) which may or may not resemble that of WLCG. They would have also full autonomy on resource negotiations, usage policies and decisions about which services and tools to use or not to use. Clearly, they would need to fund the computing resources they will need. We do not expect nor propose to use the same resources for all experiments, but rather to try and use the same infrastructure, tools, software, and support as far as possible so that new projects are easier to support on existing facilities. Of course, this helps opportunistic use and sharing, but does not impose it. Experiments joining the SCI would likely produce their own MoU and will need to identify their own reporting lines to the SCI bodies.

The **Scientific Computing Infrastructure** would be driven by the major HEP sciences with a stake on the common infrastructure. Such sciences are represented by the WLCG project and the computing projects of other major HEP experiments, as just described above. Major HEP experiments should be considered the ones negotiating, with their own funding agencies, computing needs of the same order of the LHC ones. A **Scientific Computing Collaboration Steering Group** should be set up. The role would be to ensure that the physical and software infrastructures evolve in the direction that is suitable for the community and its projects. It should also be a mechanism to obtain or encourage funding and contributions of effort, through direct feedback to the Funding Agencies and laboratories. Finally, such a steering group would be an ideal forum within which to broker community-wide needs, such as licensing, joint procurements, agreements, and policies. For example, HEP-wide agreements with cloud vendors to get scale economies. It should also address political concerns, for example how to evolve the funding models.

The composition of the group should be discussed more widely, but most likely would include the heads of Information and Technology of the major High Energy and Nuclear Physics laboratories worldwide, the computing project leaders from the major projects, facilities and experiments. This group could receive a mandate from a body such as the International Committee for Future Accelerators (ICFA), which is key to obtaining recognition in some countries. The governance should be very lightweight and decisions should be taken through consensus. The steering group would report on its activities to the involved HEP sciences, through their representatives. The day by day activities would be organized leveraging existing mechanisms and initiatives, such as the HSF, HEPIX, existing working groups and task forces.

The **HSF** would be the vehicle by which the foreseen improvements in the area of software are addressed, as explained in the sections above. It would replace the role of the Architects Forum in WLCG and broaden its scope to provide a common set of tools, libraries and techniques for the different projects in the SCI. It would respond to the SCI Steering Group and inherently to the SCI projects.

The **SCI Deployment Board** would cover, in the new organization, the functions of the Grid Deployment Board in WLCG and therefore replace it. The representation would be broadened to experiments and computing centers of the full SCI, to discuss, and prepare the decisions and plan the deployment and operations of the SCI services.

While we explicitly discuss this for HEP here, there is potential interest from other related scientific collaborations, for example astronomy/astro-particle and 3rd generation gravitational wave community. Negotiating such a broadened scope would also be welcome, and help address concerns of funding agencies and large scientific data centres of being able to more uniformly support a set of sciences with significant requirements.

Conclusions

The WLCG organization has successfully provided a global computing service for the LHC experiments for more than a decade. In this paper we have considered the future HEP ecosystem and the likely evolution of HEP computing needs in the coming years. We propose a model where the current WLCG infrastructure evolves into a Scientific Computing Infrastructure, covering the needs of major HEP experiments. Such an infrastructure would be driven by the major HEP and related sciences, reducing cost by leveraging economies of scale, common tools, services and operations. Each science would retain a high degree of autonomy in negotiating resources, setting policies and adopting tools and services.