

Computing at the HL-LHC and Beyond

Robert W. Gardner Jr^{a,*}

^a*Enrico Fermi Institute, University of Chicago,
933 East 53rd St., Chicago, Illinois, USA*

E-mail: robert.w.gardner@cern.ch

The High-Luminosity Large Hadron Collider (HL-LHC) is set to introduce unprecedented data volumes and computational demands, necessitating significant enhancements in the current LHC computing infrastructure. We summarize efforts by the experiments to integrate high-performance computing clusters and public cloud resources into their processing frameworks. We also examine the adoption of cloud technologies for implementation of advanced service infrastructure which are finding applications in Tier 2 centers and prototyping of future analysis facilities. We highlight the crucial role of scalable networking capabilities and challenge exercises to prepare for the expected increased data throughput.

12th Edition of the Large Hadron Collider Physics Conference (LHCP2024)

3-7 June 2024

Northeastern University, Boston, USA

*Speaker

1. The LHC Computing Landscape

1.1 Today's WLCG Infrastructure

For nearly two decades, the Worldwide LHC Computing Grid (WLCG) [1] has sustained a robust global infrastructure that, as of 2024, supports access to over a million CPU cores and an exabyte of storage capacity [2] for the four major LHC experiments. The experiments have developed customized job execution systems and policy driven data management services that leverage the shared grid infrastructure while allowing the flexibility to implement unique, collaboration-specific processing pipelines and data management priorities. Tasks from the experiments can utilize up to 170 computing and storage facilities that offer a heterogeneous mix of computational resources, including traditional grid systems, high-performance computing (HPC) clusters, and public cloud platforms. The integrated computing “facility” infrastructure has proven its capability and flexibility for continuous LHC computing operations for years, distinguishing it among scientific computing ecosystems. With the impending era of the High-Luminosity LHC (HL-LHC), there emerges a crucial need to significantly enhance this infrastructure or augment it with alternatives to meet the required resource projections. The integration of diverse facility technologies – from grid systems to HPC centers to cloud computing platforms – highlights the challenges and strategic decisions that must be made by the collaborations in preparation for HL-LHC computing.

1.2 Evolving the Infrastructure for HL-LHC

The concerted efforts of the LHC collaborations [3, 4], alongside the broader High Energy Physics (HEP) community [5], are categorized into several critical areas: software research and development as highlighted at this conference [6], distributed computing R&D aimed at systems capable of managing exabyte-scale data across distributed, heterogeneous resources [7], and computing facility R&D activities focused on creating flexible, reproducible infrastructures that utilize cloud native technologies [8]. These advancements in distributed computing and the development of adaptable facility infrastructures are crucial for introducing new capabilities that will effectively meet HL-LHC requirements for bulk processing and analysis. These initiatives extend beyond merely increasing capacity (such as more jobs, CPU-hours, and storage capacity); they also significantly enhance the adaptability and long term sustainability of both software and infrastructure.

2. Computing Resources for the HL-LHC Scale

2.1 Resource Estimates

The HL-LHC, now scheduled to commence in 2030 [9], is expected to increase the data recording rate for ATLAS by 7–10 times compared to current Run 3 levels, with the number of collisions per bunch crossing potentially rising from 30–60 to 200. This significant increase in data volume and complexity will challenge the existing computational model, necessitating advancements in event generation, fast simulation techniques, lossless data compression, and more efficient data storage strategies with compact data formats. ATLAS has a resource modeling framework to incorporate both conservative and aggressive software R&D scenarios, the conservative assuming potential improvements with current staffing levels, and an aggressive scenario with promising advanced

solutions. These scenarios can be compared to computing and storage resource extrapolations assuming 10% or 20% increases from flat budgets. Key areas for development have been identified, including multi-threading, data format optimizations, and GPU integration for machine learning (ML) training tasks. Similarly, CMS is modernizing its physics software and developing innovative algorithms to leverage machine learning to make optimal use of hardware accelerators, one component of a broad strategy [4] which includes building infrastructure for exabyte scale datasets, leveraging industry advances in data science, and modernizing facility and software infrastructure.

2.2 Technology and Market Risks for On-Premise Facilities

Technology supporting computing in the HL-LHC era will be significantly influenced by the advancements and market dynamics within the semiconductor industry and other IT manufacturing sectors [10]. The concentration of cutting-edge chip production among a few companies poses substantial risks including vulnerabilities to supply chain disruptions and potential stagnation in technological advancements. Projected developments in computing components indicate moderate improvements in CPU performance and power efficiency. Network upgrades, including the transition to 25 Gbps and 100 Gbps interfaces on worker nodes, will entail investments throughout the distributed infrastructure in the lead up to HL-LHC. In storage, the pace of advancements in current disk technologies is slow but recent developments suggest progress [11]. Regardless, the increase in data volumes will necessitate significant expansions in tape storage, prompting investments in additional libraries and drives at CERN and the Tier 1 centers. Overall no insurmountable technological barriers are foreseen, but diligent monitoring of developments in CPU and GPU technologies, energy efficiency, and storage I/O is imperative to remain within existing budget constraints.

2.3 High Performance Computing Clusters

HPC clusters have presented significant opportunities for the LHC experiments due to their comparatively vast numbers of cores and accelerators relative to the WLCG grid [12]. Accessing these resources efficiently has posed challenges, although recent advancements have facilitated the integration process through site-level support for container images and software caches, availability of large x86 partitions, and enhanced access to the public network from compute nodes. Despite these improvements, data management—specifically, the transfer of large datasets into local storage and the delivery of outputs back to the experiments’ dedicated storage facilities—continues to be hindered by access policies and the design of these facilities, given their scale and diverse science communities supported. CMS has significantly expanded its use of HPC clusters, with demonstrated scheduling of up to 100K concurrent CPU cores [13]. LHCb primarily dedicates its offline computational resources to Monte Carlo simulations on the WLCG. Recently, however, they have expanded their use to include various HPC clusters in these tasks [14]. To overcome challenges related to software access and network connectivity, a variety of technical solutions have been developed. As a result, LHCb has successfully utilized cores from several HPC clusters in Europe, including Piz Daint, Santos Dumont, and MareNostrum. Similarly, ALICE has successfully integrated Perlmutter into its grid environment [15].

2.4 Burst Capability and Versatility with Cloud Computing Platforms

ATLAS has explored using the Google Cloud Platform either as an independent WLCG site or as an extension of a Tier-2 center [16], given its clear advantages in scale and versatility. All standard ATLAS workflows have been successfully demonstrated in production. The project uses a subscription-based pricing model, allowing for detailed cost breakdowns for various tasks [17]. A total cost of ownership study has shown improved cost-effectiveness, though it remains lower than that of traditional grid sites for most workloads [18]. The on-demand flexibility of cloud platforms is particularly beneficial for handling simulation workloads that require rapid scaling, often referred to as “bursting”. This capability is invaluable for time-sensitive production campaigns. Moreover, the cloud platform’s adaptability is evident in its ease of setting up specialized GPU and ARM queues, a process that can be more challenging for statically procured and configured grid sites.

3. Facility Evolution

3.1 Networks and Data Challenges

The wide area network (WAN) infrastructure, crucial for LHC computing, has consistently proven its reliability as a key component of the computing ecosystem, effectively supporting data flows across globally distributed sites. The LHCONE project [19] provides an effective logical network which optimizes connectivity between LHC computing sites. Many computing sites are now connected to the WAN at 400 Gbps and higher. In anticipation of the HL-LHC’s data demands, ongoing R&D efforts are focused on optimizing data throughput performance and quality of service. The WLCG Data Challenge 2024 (DC24) [20], conducted to test 25% of the expected HL-LHC throughput, aimed to stress-test data transfer tools, optimize network configurations, and validate the scalability of file transfer tools. This challenge successfully achieved rates of up to 2.4 Tbps over several hours. DC24 identified several performance bottlenecks, particularly with token refresh operations and database overloads, which provided significant insights for future improvements. Despite these challenges, token-based transfers were successfully implemented, and new network technologies for load balancing and congestion control were evaluated. The challenge also enhanced monitoring capabilities and identified minor discrepancies between different monitoring systems, laying a solid foundation for future upgrades and optimizations. The next data challenge, scheduled for late 2026 or 2027, will target 50% of the anticipated HL-LHC network traffic.

3.2 Towards More Flexible and Reliable Infrastructure

Cloud-native technologies like Kubernetes are being adopted at CERN and innovative WLCG sites [21] to improve resource management and reduce deployment complexity. Declarative deployment methods build in more reliability and speed up expertise transfer from one facility to the next. Analysis facility and advanced data delivery systems prototyping [22, 23] have demonstrated this approach for both ATLAS and CMS at multiple sites. Analysis challenges [24] serve as a framework to assess technology and scale readiness for HL-LHC. ALICE will evolve its Run 3 O2 system for Run 4 [25].

4. Conclusions

As the LHC transitions into the High-Luminosity era, the evolving demands on its computing infrastructure are substantial yet manageable with strategic planning and continuous R&D. The integration of resources from HPC and cloud platforms during LHC Run 3 is helping inform strategy and decisions for HL-LHC. Upgrades in network capabilities and data throughput exercises are revealing bottlenecks and identifying needed services development. Leveraging cloud technologies and declarative integration tools are improving infrastructure building efficiency and reducing the reliance on specialized expertise. Analysis challenge exercises are proving useful for facility prototyping for the HL-LHC scale.

Acknowledgments

This work was supported in part by the following grants from the National Science Foundation: OAC-2115148, OAC-2029176, OAC-1836650 and PHY-2120747.

References

- [1] I. Bird, P. Buncic, F. Carminati, M. Cattaneo, P. Clarke, I. Fisk et al., *Update of the Computing Models of the WLCG and the LHC Experiments, CERN-LHCC-2014-014, LCG-TDR-002* (2014) .
- [2] M. Barisits, T. Beermann, F. Berghaus, B. Bockelman, J. Bogado, D. Cameron et al., *Rucio: Scientific data management, Computing and Software for Big Science* **3** (2019) 11.
- [3] ATLAS Collaboration, “ATLAS Software and Computing HL-LHC Roadmap.” <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/UPGRADE/CERN-LHCC-2022-005>, 2022.
- [4] O. Gutsche, T. Bose, M. Votava, D. Mason, A. Melo, M. Liu et al., *The U.S. CMS HL-LHC R&D Strategic Plan, EPJ Web of Conf.* **295** (2024) 04050.
- [5] “The HEP Software Foundation.” <https://hepsoftwarefoundation.org/>.
- [6] D. Shope, “Software Upgrades for HL-LHC (LHC2024).” <https://indico.cern.ch/event/1253590/contributions/5814391/>, 2024.
- [7] D. Cameron, M. Lassnig and D. South, *ATLAS Distributed Computing Evolution: Developments and Demonstrators Towards HL-LHC, EPJ Web of Conf.* **295** (2024) 04034.
- [8] D. Ciangottini, A. Forti, L. Heinrich, N. Skidmore, C. Alpigiani, M. Aly et al., “Analysis Facilities White Paper.” <https://arxiv.org/abs/2404.02100>, 2024.
- [9] <https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm>, 2024.
- [10] B. Panzer-Steindel, “Computing Technology and Market Evolution (WLCG/HSF Workshop).” <https://indico.cern.ch/event/1369601/contributions/5908541/>, 2024.

- [11] W.-H. Hsu and R. Victora, *Heat-assisted magnetic recording — Micromagnetic modeling of recording media and areal density: A review*, *Journal of Magnetism and Magnetic Materials* **563** (2022) 169973.
- [12] F.B. Megino, L. Bryant, D. Hufnagel and K.H. Anampa, “US ATLAS and US CMS HPC and Cloud Blueprint.” <https://arxiv.org/abs/2304.07376>, 2023.
- [13] A. Pérez-Calero Yzquierdo, M. Mascheroni, E. Kizinevic, F.A. Khan, H. Kim, M. Acosta Flechas et al., *HPC resources for CMS offline computing: An integration and scalability challenge for the Submission Infrastructure*, *EPJ Web of Conferences* **295** (2024) 01035.
- [14] A.F. Boyer, F. Stagni, C. Haen, C. Burr, V. Romanovskiy and C. Bozzi, *Integrating LHCb Offline Workflows on Supercomputers State of Practice*, *EPJ Web of Conf.* **295** (2024) 10005.
- [15] S. Weisz, “Integrating the Perlmutter HPC system in the ALICE Grid (CHEP 2024).” <https://indico.cern.ch/event/1338689/contributions/6011007/>, 2024.
- [16] ATLAS collaboration, *Using Kubernetes as an ATLAS computing site*, *EPJ Web Conf.* **245** (2020) 07025.
- [17] F.B. Megino, K. De, J. Elmsheuser, A. Klimentov, M. Lassnig, M. Euell et al., *Operational experience and R&D results using the Google Cloud for High-Energy Physics in the ATLAS experiment*, *Int. J. Mod. Phys.* **39** (2024) .
- [18] ATLAS Collaboration, “Total cost of ownership and evaluation of Google cloud resources for the ATLAS experiment at the LHC.” <https://arxiv.org/abs/2405.13695>, 2024.
- [19] E. Martelli, *Evolving the LHCOPN and LHCONE networks to support HL-LHC computing requirements*, *EPJ Web of Conf.* **295** (2024) 07016.
- [20] M. Lassnig and C. Wissing, “WLCG/DOMA Data Challenge 2024 Final Report.” <https://doi.org/10.5281/zenodo.11444180>, 2024.
- [21] R.P. Taylor, J.R. Albert and F.H. Barreiro Megino, *A grid site reimagined: Building a fully cloud-native ATLAS Tier 2 on Kubernetes*, *EPJ Web Conf.* **295** (2024) 07001.
- [22] M. Adamec, G. Attebury, K. Bloom, B. Bockelman, C. Lundstedt, O. Shadura et al., *Coffea-casa: an analysis facility prototype*, *EPJ Web Conf.* **251** (2021) 02061.
- [23] O. Rind, D. Benjamin, L. Bryant, C. Caramarcu, R. Gardner, F. Golnaraghi et al., *The Creation and Evolution of the US ATLAS Shared Analysis Facilities*, *EPJ Web of Conf.* **295** (2024) 07043.
- [24] B. Bockelman, “IRIS-HEP 200 Gbps Challenge (WLCG/HSF Workshop).” <https://indico.cern.ch/event/1369601/contributions/5924000/>, 2024.
- [25] M. Richter and for the ALICE Collaboration, *A design study for the upgraded ALICE O2 computing facility*, *Journal of Physics: Conference Series* **664** (2015) 082046.