CERN-LHCC-2022-007; LHCC-G-183 25 January 2022

HL-LHC Software and Computing Review Panel Report
2nd Meeting – November 1, 2021

Authors:
Amber Boehnlein (Jefferson Lab) Chair

Catherine Biscarat (LHCC, IN2P3), Andrea Bressan (INFN), David Britton (Glasgow), Rosie Bolton (Cambridge, SKA), Frank Gaede (DESY), Claudio Grandi (INFN), Fabio Hernandez (IN2P3), Thomas Kuhr (LMU) , Gonzalo Merino (CIEMAT), Frank Simon (LHCC, MPP), Gordon Watts (U Washington)

CERN-LHCC-2022-007 / LHCC-G-183
06/03/2022

Introduction

This report summarizes the outcome of a review of ATLAS, CMS and the Common Software Activities, performed as the second stage of a standing review series under the direction of the LHCC. The mandate stems from a request from the Resource Review Board to the CERN Research Director, as one mechanism to validate the HL-LHC Computing Strategies of ATLAS and CMS from a technical perspective. The series of reviews has a progression and is intended to cover the main areas and activities of relevance. The final report from the initial step of this process, a baseline review of the experiments' plans, held May 18-20, 2020, is available at https://cds.cern.ch/record/2725487.

Preparation for this second review began in February 2021 with the delivery of the Review Charge (https://cds.cern.ch/record/2777575). The charge focused on updates from ATLAS and CMS in response to the observations from the previous review, and on a set of Common Application Areas determined in discussions with representatives from the major stakeholders, as follows:

- Data Organization, Management, Access (DOMA)
- Geant4
- ROOT Foundation
- ROOT Analysis
- Data Science Tools for Analysis
- Physics Generators

The input documentation was delivered by October 1, 2021 and the review took place in a hybrid mode from November 1-5, 2021. The format consisted of short presentations followed by an extended discussion period led by questions from the reviewers. The committee was uniformly impressed by the quality of the documentation and the presentations. The effort taken to produce the documentation, which spanned multiple groups, some of which are collaborations or projects and others that are communities-of-practice, was much appreciated.

Key Observations, Comments and Recommendations:

From a technical perspective, all activities and the experiments are showing excellent progress.

We appreciate the agreement on the machine parameters for the start of the HL-LHC, which enables the more consistent projection of resource requirements. Both ATLAS and CMS demonstrated an appropriate understanding of the R&D needed to achieve the performance gains required to deliver the computing within the flat budget projections for HL-LHC. We note that increased investments in R&D activities are essential to contain costs and to sustain operations. The experiments have also made strides in using heterogeneous architectures and in using HPC centers.

Software and Computing activities in the experiments take place in a continuum between daily operations and R&D activities. This underpins the integration of successful R&D into production activities without the need for major shutdowns. However, operational issues must have priority, which can negatively impact R&D timelines, especially with limited effort. This need to support on-going production activities while pursuing aggressive R&D, is further complicated by external dependencies on other projects and activities. We emphasize the importance of an appropriate level of formal project management, including risk characterization and management, to enable the experiments to maintain an appropriate level of control and overview.

The success of the ATLAS-developed data management framework Rucio is impressive and the adoption of Rucio as a common tool by CMS (and others) is very positive. We note that additional contributions from CMS and from experiment-independent effort, would further strengthen the Rucio project.

The HEP Software Foundation discussion forums were particularly useful in fostering positive collaboration around the event generator community and the data science tools for analysis community. Generator projects are making impressive progress despite a mismatch in the underlying incentive structure that we mention later.

Detailed high-fidelity Geant4-based simulations have been crucial for the LHC experiments and will remain a key tool for the success of the HL-LHC physics program. Good progress is being made on fast simulation and the Geant4 collaboration is demonstrating progress on using GPUs for electromagnetic processes. The experiments have strong relationships with Geant4 that should continue. More coordination and collaboration on fast simulation would be beneficial.

ROOT is essential for HL-LHC computing. The new RNTuple data format is a major change in the foundation layer that can significantly reduce the data volume and the data access time. We strongly encourage the ROOT and experiment (ATLAS / CMS) teams to continue prioritizing work on RNTuple rollout and adoption. The developing PyHEP ecosystem demonstrates a wide interest in analysis in the HEP community and underlines the growing importance of the analysis area for HL-LHC. The ROOT and PyHEP approaches have complementary philosophies regarding analysis tools, but the community would most benefit from a better interaction of ROOT and PyHEP, allowing co-existence and the extraction of the best from both worlds. A joint forum with both teams would be highly beneficial. For the HL-LHC, the analysis model will continue to evolve, driven by significantly more data, and the experiments should consider becoming more involved in strategies for end-user analysis including the development of tools.

The DOMA activity has played a key role in ensuring the Common Software Activities have made progress towards a common goal by coordinating their activities with the experiments. Detailed requirements have not yet been established for the HL-LHC era, but the activities coordinated by DOMA are working towards meeting the goal of increasing the volume of data stored and transported by an order of magnitude, with respect to Run-3. Given that exact details on how this increase will materialize are yet unknown, regular, realistic, and increasingly ambitious exercises involving the

DOMA activities, the experiments and WLCG sites, are of great importance. These should be pursued to help identify the current limitations of the systems as well as understand how the requirements can be met.

Detailed Observations, Comments and Recommendations:

Common Observations

The experiments and other activities are working to the current flat budgetary guidance for HL-LHC for computing and storage resources at the WLCG sites, and meeting these constraints motivates many of the research and development activities. We are pleased to note a significant reduction in the projected resource requirements has been achieved over the last couple of years, partly due to the targeted investment of effort by ATLAS, CMS and the Common Software Activities. We also note that Software and Computing activities take place on a continuum between daily operations and R&D and that there is genuine tension between maintaining the resources for operations whilst trying to perform essential R&D. Investment in additional effort for R&D by institutes and funding agencies is vital to constrain future hardware requirements. This will lead to significantly reduced costs, better physics outputs, less environmental impact, and increased training in skills that are valued in the wider societal context.

Software sustainability is a common challenge. When core developers on permanent positions retire, there is a high risk of losing critical expertise. We also note that it is more difficult to attract people and get funding for the optimization of existing software stacks than for the development of new projects.

Without additional effort, increasing levels of formal project management will be a burden on those people who are also critical for the project's success. We urge the experiments to consider whether their existing project management processes are sufficient to address the following:

1) The (evolving) degree of certainty that the overall project will deliver the required performance in time for HL-LHC.
2) Recognizing the level of risk to the project from internal and external dependencies, assumptions, and unexpected changes.
3) Establishing the decision point when choices need to be made, priorities established or changed.
4) Person-power requirements and projections.

There is value in appropriate Project Management, both internally to the project, but also as part of the case for engagement from collaborators and funding from external agencies. The excellent progress made in presenting a milestone-based view should be supported by a comprehensive Risk Register. This should be a priority going forward.

The importance of network monitoring should be stressed, both at the site level and centrally, to track network usage and flag any unexpected usage patterns. Although it is non-trivial to develop, comprehensive network monitoring may become a key factor in mitigating risks for HL-LHC data processing.

The success of the ATLAS developed data management framework Rucio is impressive and adoption of Rucio as a common tool by CMS is very positive. This should contribute to improving the sustainability and overall quality of this core framework. However, we are concerned about pressure on the Rucio development team, which appears to be severely understaffed. ATLAS and CMS should develop a common plan (together with other stakeholders) to ensure the core development team is adequately staffed.

Geant4 is a critical part of the software infrastructure for exploiting the HL-LHC data and the experiments and Geant4 have a long record of working together. Despite this, the existing official communication mechanism, with a single liaison person per experiment which seems insufficient for a project that is so complex and dynamic, the experiments and Geant4 have additional collaboration mechanisms. We encourage Geant4 and the experiments to reinforce such collaborative efforts that not only streamline communication but also facilitate sharing effort towards common goals such as, for example, AI-based fast Monte Carlo projects.

ATLAS

ATLAS presented a comprehensive list of milestones that represent a first version of the roadmap for R&D activities towards the HL-LHC, which can be used to develop the planning process. The process of building this first roadmap has been a bottom-up exercise and all the software and computing groups have been engaged. We congratulate everybody within ATLAS who has contributed and encourage ATLAS to implement a process where these milestones are periodically reviewed, and the plan updated. During the review discussions, a few corrections and additions to the figures in the document were requested. We thank ATLAS for generating an updated version of the document including these changes.

The current projected resource estimates for Run4 show that investing in R&D projects can have a big impact. Although the ATLAS conservative estimates are still well above flat-budget projections, the experiment has shown that a more aggressive R&D programme can bring the resource needs back inside the flat-budget envelope. We encourage ATLAS to continue quantifying the impact of the identified R&D projects and the effort needed to carry them out, in order to present a consistent planning framework to funding agencies that shows the importance of funding software and computing R&D activities to tackle the HL-LHC data challenge.

Even if the current conservative extrapolations of future resource needs are still high, the situation has improved with respect to the last review, so we are confident that this trend will continue as R&D projects keep delivering incremental improvements. Given the many uncertainties associated with the resource predictions, which are not only of technological nature but also related to market trends and hardware costs, we consider it important that experiments develop “last resort” emergency strategies to reduce the size of their computing requirements if the available resources are not sufficient by the time the detector starts taking data. We encourage ATLAS to have this discussion internally.

Accounting for the effort dedicated to the R&D computing activities is a complex matter: people with very different skill levels contribute at different rates to the development, which makes a quantitative assessment of contributions difficult. Nevertheless, we still encourage ATLAS to implement mechanisms to gather information on effort investments that, with the appropriate corrections or associated uncertainties, can be used in the planning process.

Analysis Facilities, understood as specialized infrastructures that facilitate interactive analysis on very large datasets, may play an important role in enabling effective analysis in the future. There are interesting initiatives exploring the possibility of exploiting technologies such as Dask or Spark, columnar data formats, etc. ATLAS has so far taken a hands-off approach towards these activities, encouraging free exploration of new ideas. We think this is a right approach in general, however we encourage ATLAS to find a way to track the progress of these activities and incorporate them into their global picture for an end-to-end computing infrastructure.

The migration of the AAI from x509 certificates to tokens is a major project which touches almost every part of the infrastructure. The complete process is expected to take several years, but important changes in the production system are already happening now or will do so in the near future. ATLAS appears to have a good handle on the process and detailed steps have been defined to ensure changes are transparent for users. Given the importance of the change, we encourage ATLAS to keep managing it with short planning cycles and keeping sites engaged at the appropriate level.

We appreciate the efforts from several ATLAS groups towards meeting the HL-LHC challenge and the coordination work from the management team. Several activities are already yielding encouraging results, and we believe that the presented roadmap is an excellent starting point for managing the complex R&D program ahead. We encourage ATLAS to continue on this track and look forward to follow up in future reviews.

CMS

CMS presented a credible and substantial outline of how they plan to address the HL-LHC computing challenges of Run-4 and Run-5 and should be congratulated on their progress since the last review. Building on their existing infrastructure, their emphasis on flexibility and innovation is both commendable and appropriate at this stage of the project. As time progresses, choices will need to be made and effort will need to become focused on the areas identified as most critical. We hope the current optimism is supported fully by the progress that is achieved, but believe that there are still significant challenges ahead.

Although supportive of the flexible and innovative nature of CMS' computing model, we are cautious of overly depending on future network headroom, particularly for the 'Any Data, Anytime, Anywhere' (AAA) approach. This mechanism requires some degree of synchronicity, but it is also recognized that the use of caches is part of the model, which could allow more reliance on asynchronous data movement.

CMS has made excellent progress in constraining storage needs through the size of RAW and nanoAOD formats. We note and encourage the potential to further reduce the former, and further promote the uptake of the latter. We are reassured that CMS has considered "an emergency brake" in the form of Partial Prompt Reconstruction and, of course, share CMS' hope that this will not have to be deployed.

There are good, and in some cases excellent, connections between CMS and the external common projects, which provide tangible benefits to CMS. As HL-LHC gets closer, many aspects of the plan will need to become better defined, such as the choice of generators and the degree to which NNLO needs to be incorporated. This, in turn, relies on progress in the Generator communities, particularly in the suppression of negative weight events. The balance between fast and full simulation, and the use of AI techniques, will need to be established. The choice of analysis tools and possible use of analysis facilities will need to be defined. The timescales and risks associated with all these aspects should be considered in the CMS project management process, enabling an end-to-end scale test of the system well in advance of HL-LHC data.

CMS is congratulated on the recent adoption of Rucio and the currently excellent relationship with the Rucio team is noted. However, we also understand the concern that the core Rucio developers are completely embedded in ATLAS and that, hypothetically, this could lead to tension when setting priorities in the future. This risk would be mitigated by complementing the current Rucio development effort with additional experiment-independent effort and, indeed, by increased CMS contributions to the core Rucio components, leading to a more equitable ownership.

We note the steady progress on speeding up the RECO step and the anticipation that this continues. We caution that the linear progress made to date may be hard to maintain in the longer term; the new HGCAL is noted as a particular challenge. Similarly, the 50% reduction in time targeted in the CMS Geant4 simulation step, coming from a combination of Geant4 and CMS-specific improvements, may be challenging but it is not an unreasonable aspiration at this stage.

Finally, we note the difficulties in constraining projections for tape storage into the projected resource bands and the fact that this is largely deterministic from the agreed HL-LHC parameters and the RAW data size. In many countries, it may be possible to trade off CPU and disk storage resources, for increased tape storage, but that national planning for tape facilities can typically be quite long-term due to the discrete nature of tape-robots and the fact that these are sometimes shared with other communities. Therefore, it is important that tape requirements are established and communicated well in advance, particularly if they are higher than expected from flat-cash.

Overall, CMS presented a carefully thought out and credible outline of plans to meet the computing challenges of the HL-LHC. A significant reduction in the projected resource requirements has been achieved over the last couple of years, partly due to the targeted investment of effort by CMS and the Common Software Activities. As noted above, investment in additional effort by institutes and funding agencies is vital to constrain future hardware requirements.

DOMA

The WLCG DOMA project (for Data Organization, Management and Access) covers activities related to data access, content delivery and caching, third party copy and quality of service. It includes Rucio for data cataloging and replication; FTS and GFAL for data transfer; storage technologies (XRootD and XCache, EOS, dCache, Echo, StoRM, CTA); and network technologies. The latter depend on, or complement, common technologies for identification and authentication (IAM); information systems

(CRIC); and network monitoring (perfSONAR). Each DOMA activity is under the responsibility of its own team, often composed of individuals affiliated to several organizations.

DOMA plays a key role ensuring that the ten DOMA Common Software Activities make progress towards a common goal, by taking an agile approach of continuous interactions with the experiments. Although detailed requirements have not been established, the Common Software Activities are working towards meeting the goal for the HL-LHC era of increasing the volume of data stored and transported by an order of magnitude, with respect to Run-3. Given that details on how this increase will materialize are yet unknown, we believe that regular, realistic exercises of increasing ambition and complexity, involving the Common Software Activities, the experiments and WLCG sites are of great importance and should be pursued. Those exercises will help identify the current limitations of the systems as well as establish how the increase can actually be achieved (e.g., number of files, individual file size, data access patterns, etc.).

We perceive the adoption by CMS of Rucio for managing data distribution as an important step forward towards the harmonization of tools and practices among the experiments. Not only is this move likely to contribute to the long-term sustainability of Rucio as a tool, but it may also help WLCG sites supporting the HL-LHC experiments as well as other sciences. Similarly, the initial successful steps of replacing X509 certificates by bearer tokens as a mechanism for authenticating users and services are very encouraging. We understand that work is needed to generalize their usage and that there are still issues to be solved by the storage systems if tokens are to be used for authorization purposes.

We note that the teams responsible for the data management and storage technologies have already identified specific components which may be a limiting factor for coping with the expected data increase (e.g., databases) and are working on understanding the current limitations and addressing the issues. More generally, establishing milestones common to the DOMA Common Software Activities could help ensure that progress is made towards meeting the experiments' expectations. This would build a shared perspective with the WLCG sites that will play a central role operating storage and data transport services used by the experiments.

It is apparent that developers of some Common Software Activities are also involved in operating the services they develop. Whilst this is important for ensuring development teams have a good grasp of the use made of their tools, operational priorities could threaten long-term development work if staffing effort is not sufficient. We note that for several teams the estimated effort available is not sufficient to deliver the additional work required to meet the HL-LHC challenge.

Emerging R&D activities aimed at making the network a managed resource, related both to detailed usage monitoring and dynamic bandwidth provisioning, are important and appropriate to pursue. Given the increasing reliance on the wide area network for delivering data to experiment's applications, detailed network observability is a key factor to understand the behavior of the resource, to identify and mitigate associated risks.

We congratulate each of the DOMA activities on the progress made towards identifying their risks and recommend that risks associated across the Common Software Activities be recorded and regularly assessed in a broader context.

Geant4

The Geant4 toolkit is a framework for the detailed simulation of the passage of particles through matter, developed and maintained by the long-established Geant4 Collaboration. It is widely used in high energy, nuclear and accelerator physics, as well as for studies in medical and space science. The Geant4 collaboration counts ~30 FTE (provided by 130 collaborators) of which 40% are part of the high energy and nuclear physics community.

Detailed high-fidelity Geant4-based simulations have been crucial for the LHC experiments and will remain a key tool for the success of the HL-LHC physics program. However, the Geant4 toolkit must also evolve. The increase in the number of events that will be accumulated by the experiments and the necessity to keep systematic uncertainties at the current level, calls for the availability of even more accurate simulations for precision measurements, while decreasing the average time per event. For channels not limited by systematic uncertainties, the use of fast simulations, at the expense of reduced accuracy, for the production of part of the required samples will be needed to constrain the overall consumption of computing resources. The benefits of fully exploiting accelerators, increasingly available in the computing facilities used by the LHC experiments, must be explored. The strategy of the Geant4 collaboration takes into consideration these factors.

By the start of the HL-LHC, the Geant4 collaboration estimates at most a factor of two speed-up in the detailed “standard” CPU-based simulation can be achieved by the incremental refinement of algorithms and their implementation. However, new challenges also need to be addressed, such as the efficient simulation of the new CMS HGCAL. We note that improvements in the speed of Geant4 do not automatically translate into equal gains within the experiments’ workflows and we encourage Geant4 experts to reinforce their involvement in the tuning of Geant4 in the experiments’ software stacks to maximize the benefits.

Geant4 has significant experience with fast simulation via the BaBar GFlash electromagnetic shower parametrization and such techniques potentially offer a large gain. Different approaches, based on new machine learning techniques, are also extremely promising. We encourage Geant4 and the experiments to increase coordination on these developments. It is noted that the accuracy of fast techniques still relies on the high-fidelity Geant4 simulations that are used to tune them, and are not a full replacement.

Another key area of R&D is the porting of detailed simulation to GPUs. Whilst the hadronic transport part is extremely challenging, the G4HepEm library release has already demonstrated promising results for the electromagnetic component. Together with the VecGeom development, the very active AdePT and Celeritas projects working towards a detailed GPU-based simulation, are noted. Both projects are currently focusing on the electromagnetic physics and the Celeritas project also plans to integrate the hadronic physics into experimental workflows through Geant4 by the start of the HL-LHC.

The Geant4 Collaboration has responded to the new challenges by adding an R&D task force to coordinate and evaluate activities on fast simulations and accelerators, complementing the existing working groups. In addition, a new role of Geant Contributor has been created to facilitate contributions from people who are not full members of the collaboration. We are keen to see if this new initiative adds flexibility and enables new engagement. It is also crucial that the physics description of the detailed simulation continues to improve, and care should be taken in maintaining a sufficient core developer team with appropriate expertise.

The communication channels between Geant4 and the experiments go beyond the Geant4 Technical Forum with close, but informal, links between them. The HSF simulation working group plays a complementary role, bringing together experts from simulation and other areas, in an agile and informal way. The collaborative spirit between the experiments and Geant4 gives us confidence that milestones in the design of the computing models for the HL-LHC will be met on time and the mix of fast and full simulation will enable the bulk production of simulated samples ahead of the start of the HL-LHC.

The inclusion of FLUKA models in Geant4 may improve the overall precision but is a time-consuming task. We suggest that the experiments are involved in the prioritization of this activity in the overall project schedule.

Physics Generators

Event Generators are absolutely critical to the success of the HL-LHC and we congratulate all involved with the preparation and writing of the material for the review. There appears to be good communication between the experiments, the HSF, and the event generator and MC tool authors. Plans for both physics and performance improvements were presented. We note that generator-authors' career goals are mostly focused on developing new physics calculations, but not on the optimization of the new or existing code needed for LHC experiment production. New physics calculations and optimized code are both required for a successful physics program at the HL-LHC.

The precision required of the event generators to complete the HL-LHC's physics program is daunting. The event generator authors have a long list of physics improvements that could be made, whilst the experiments have a prioritized list of what they need. However, the latter is only one input into the authors' overall prioritization of new feature work, along with funding and publication pressures. Steps to address this by the community would help make sure the event generators were ready for the HL-LHC, especially as MCNet funding is coming to a close.

The steps taken to automate the inclusion of new features and bug-fix releases into the experiments' frameworks are important: being able to quickly test fixes and new features and provide feedback to the authors is crucial in improving both the physics and the performance of the event generators. We encourage the experiments, authors, and the HSF to work together both to make this process smoother and to include as many of the generators and tools as possible.

Negative weight events can have a dramatic effect on the computing resources required as negative weight events cancel out positive weight events. This requires more events to be generated, increasing the running time of the generators and the downstream experiment simulation and reconstruction code. Techniques have been developed to mitigate this issue at NLO but the issue is less understood at NNLO where the number of negative weight events will be intrinsically higher, and will require theoretical advances. Both CMS and ATLAS are depending on the negative weight issue being well under control by the time HL-LHC NNLO production starts.

The performance improvements that have been made are also impressive: whilst many participants called these gains “low-hanging fruit”, we note that some of them required optimization of the generator algorithms - not just changing the way a loop was run or using intrinsic transcendental functions. A number of the improvements were achieved through collaboration with experts in code optimization and hot-spot detection. We believe this collaborative model should be encouraged, and that future progress will require contributions from both generator authors and code optimization experts. In many cases the exact way a generator is configured (the process, tune, etc.) can have a large impact on the event generator’s performance. The experiments are encouraged to continue to collaborate and share explicit run-time configurations, so that authors and optimization experts can be sure they are solving the experiment’s problems in the experiment’s software configuration environment. This will be especially important as GPU and SIMD vectorization improvements are rolled into production.

The HSF Generator working group has functioned as a convenient meeting place for generator authors, software and optimization experts, and technical experts from the experiments. We are very appreciative that authors were willing to share technical aspects of improvements - hopefully improving all generators and tools - and encourage this collaboration to continue. We also encourage the experiments to continue to understand the significant differences reported in CPU needs for generation of events. The committee urges tracking the negative weights issue at higher orders carefully.

ROOT Foundation

The ROOT team presented a detailed plan of current and future developments. Their priority is the development of a new data format, called RNTuple, which should reduce the size of the stored data and read access time compared to TTree. ATLAS and CMS plan to use RNTuple to store derived data, CMS even for their raw data. The choice is well motivated by the expected resource savings, and we strongly support the prioritization of the RNTuple development. This includes the development and support of the RNTupleLite library as well as detailed documentation for its use to facilitate the access to RNTuple data in other tools or libraries.

Persistence is a hard problem and long-term support for the data format storage is essential. The ROOT project has an excellent track record and has shown that it can keep a sustainable level of development and support for decades.

We believe the current level of human resources is appropriate for the proposed plan. Maintaining a deep understanding of the system and the required high level of technical skills in an evolving team, including retirements, is identified as a major risk. The ROOT team has an incredible level of technical expertise that not only underpins the ROOT project but also benefits, and is relied upon by, the wider community. This should be preserved.

In conclusion, we believe that ROOT is vital for the HL-LHC and well on track to deliver what the experiments need. Efforts to improve the interoperability with other libraries or tools are encouraged.

Analysis

For many years the analysis of HEP data was done largely with ROOT, but in recent years the use of common data science tools in a Python ecosystem has increased significantly. The PyHEP working group of the HSF supports the exploitation of such data science tools in a Python based HEP data analysis environment.

ROOT and PyHEP have different philosophies: the ROOT team aims for an integrated framework addressing HEP needs. This has the advantage that optimizations can be implemented for HEP use-cases across all software layers. It may also reduce the fragmentation of analysis activities, which then helps with the reuse of analysis code and with analysis preservation. On the other hand, PyHEP supports a toolbox of data science tools, developed partly outside the community, that can be used for HEP analyses. This has the advantage that one can exploit developments by others and there is more flexibility and agility in the choice and use of tools. This can foster innovations and promote careers of people that leave HEP. Both approaches have obvious benefits.

The ROOT project has an established management structure and a detailed and credible work plan. The ROOT team does advanced resource planning and anticipates the needs of the community; it has been reasonably good at this in the past. PyHEP is a collection of smaller projects with less structured organization and decentralized planning. The smaller project size can make it easier to attract new contributors. The interaction of the experiments with ROOT and PyHEP is via personal contacts and official channels

We are convinced that the HL-LHC program and the HEP community in general, would benefit from a closer interaction between ROOT and PyHEP projects. Both share common challenges, such as the handling of systematic uncertainties; analysis preservation; and the discovery of tools. Dissemination and training are important, particularly for analysis tools, and synergies may be exploited if ROOT, PyHEP, and the experiments join forces in this area.

The resources needed for analyses are expected to increase considerably for the HL-LHC. Leaving the responsibility for providing hardware and software for analyses to individuals or institutes is unlikely to be optimal and we recommend that the experiments consider a more active role in solving the HL-LHC data analysis challenge. To assess the advantages and disadvantages of the available analysis tools for different use-cases, we suggest the development of benchmarks or performance indicators such as analysis turn-around time, resource requirements, ease-of-use, reusability, etc.