

# Report from US HPC Integration Meeting

Ian Fisk  
September 5, 2024

# Meeting

The US meeting was held January 22nd and 23rd, 2024, in New York City

Representatives from

- US computing projects of ATLAS and CMS
- WLCG
- CERN IT
- Argonne National Lab (Aurora)
- Berkeley Lab (NERSC)
- SDSC

Was intended as mostly discussion, no slides were allowed, though we did have a google doc for notes and to guide the discussion

13 people came

It took a long time, but there is a draft report

- People are still commenting and responding, so it's a draft

[https://docs.google.com/document/d/1V-vy5y9JUswCtiq3f6q3DmOQlzTJj6eKlwprzPW7\\_7w/edit#heading=h.g65tt9rswf8c](https://docs.google.com/document/d/1V-vy5y9JUswCtiq3f6q3DmOQlzTJj6eKlwprzPW7_7w/edit#heading=h.g65tt9rswf8c)

## Strategy For HPC Integration (US) Whitepaper

Authors: Debbie Bard, Doug Benjamin, Paolo Calafiura, Simone Campana, Taylor Childers, Ian Fisk, Maria Girone, Oliver Gutsche, Dirk Huffnagel, Alexei Klimentov, Eric Lançon, Verena Martinez Outschoorn, Frank Wuerthwein  
Indico: <https://indico.cern.ch/event/1356688/>  
Passcode: HPCS12024

### Introduction

High-Performance Computing (HPC) sites represent the most substantial computational resources available to the scientific community. Since the inauguration of the Large Hadron Collider (LHC), these HPC facilities have experienced an astonishing 1000-fold increase in computing power, now reaching the exascale level. As upcoming High Energy Physics (HEP) experiments and other data-intensive scientific endeavors are projected to generate exabytes of data annually, there is an urgent need to efficiently collect, distribute, process, and analyze this vast amount of information.

HEP, with its extensive datasets and intricate connections to various scientific fields, has the potential to stand at the forefront of this integration effort. To fully harness the physics potential of this data and enable speculative exploration, it is imperative to explore and adopt new processing techniques, technologies, and resources. The convergence of Artificial Intelligence (AI) with HPC further amplifies the potential for groundbreaking discoveries and drives innovation across these domains.

# Topics

Provisioning

Interfaces

Software

Data

Workflows

# Provisioning

## Topics:

- **Adaptive Allocation Models**
- **Proposal-Based Access**
- **Resource Scaling and Elasticity**

**Open Issues:** Lack of long-term provisioned allocations on HPC.

**Status:** Tolerable, with the current approach of using proposals and exploring dynamic and elastic scheduling.

**Impact:** The absence of predictable allocations may render HPC facilities unreliable for the most critical workflows.

**Next Steps:** Engage with funding agencies, WLCG, HPC organizations, and HPC resource providers to explore allocation models. Assess the feasibility of longer-term commitments and the steps required to treat HPC allocations as WLCG pledges.

# Interfaces

## Topics:

- **Standardized Interface Protocols**
- **Interoperability and Compatibility**
- **Automated Integration Tools**
- **Cyber Security**

**Open Issues:** Lack of a common set of interfaces and protocols deployed at HPC sites.

**Status:** Needs to be resolved.

**Impact:** Without a common foundation, the ability to integrate additional HPC sites into a distributed computing environment will be severely limited.

**Next Steps:** Collaborate with funding agencies, existing programs like IRI, WLCG, experiments, HPC organizations, and resource providers to document the performance, security, and technical requirements for interfaces. Form a community to develop a minimum set of common interfaces for HPC integration, focusing on standardization and reuse where possible.

# Software Strategy

## Topics:

- **Software Design Enhancements.**
- **Portability Libraries**
- **AI/ML Workflow Adoption**

**Open Issues:** Most HEP software has been developed for the x86 platform and does not benefit from accelerated linear algebra calculations available with GPUs.

**Status:** Needs to be resolved, but can be mitigated through portability libraries and reliance on GPU-supported applications like AI/ML algorithms.

**Impact:** HEP applications are currently limited to x86 CPU architectures, which are not advancing as rapidly as other hardware platforms.

**Next Steps:** Conduct a survey among experiments to assess:

- Readiness to use accelerated architectures like GPUs.
- Needs and prospects for adopting portability libraries and unified programming models.
- Expected evolution of software applications and the number of supported platforms by the time of HL-LHC.

# Data and Storage

## Topics:

- **Scalable Data Ingestion and Export**
- **Data Caching and Staging**
- **Data Management Protocols**

**Open Issues:** HPC sites lack persistent on-site data storage, necessitating dynamic staging and export of all input and produced data. The data management systems must scale with the available HPC resources.

**Status:** Tolerable, given that HEP data management systems are advanced. With the development of common interfaces for data transfer and access at HPC sites, the necessary scale should be achievable.

**Impact:** Processing, simulating, and analyzing data require access to vast amounts of data. For HPC sites to contribute, data will need to be moved efficiently.

**Next Steps:** Survey experiments on workflow requirements across three categories of access patterns. Focus on developing data challenges involving HPC sites to demonstrate the feasibility of completing workflows at the scale expected for HL-LHC.

# Workflows

## Topics:

- **Prioritizing AI/ML and Simulation Workflows**
- **Code Reengineering for Reconstruction Workflows**
- **Task Splitting for Partial Reconstructions**
- **Workflow Orchestration**

**Open Issues:** HEP workflows have limited modularity for operational efficiency, and most are still modeled around an event loop. This makes it challenging to enable parts of event reconstruction on an HPC farm. In the near term, HPC facilities are unlikely to efficiently handle all steps of the processing chain.

**Status:** Should be resolved, but progress can be made in the interim.

**Impact:** HPC facilities will be most beneficial to HEP when they can handle a large fraction of workflows. While some workflows can currently run on HPC, further development is needed to expand this list.

**Next Steps:** Work with experiments to develop a consensus on the number of workflows that could be run on HPC resources and assess the potential impact under the current provisioning, interfaces, software, and data/storage conditions. Identify the most limiting factors and study how resolving them would enhance the impact of HPC sites.



# Architectural Issues

**An element that was discussed at the meeting was the thought that HEP computing was being left behind as so much investment was being made in GPUs and accelerators we only use under controlled conditions**

- In theory, all the systems designed for AI are reasonable linear algebra accelerators and useful for our calculations**
- In practice the biggest improvements in performance are in lower precision calculations**

**Our workflows that can be rethought as AI/ML will benefit, but classical calculations won't**

- CPUs improvements are slowing and GPU performance is growing in ways that are not helpful to classical calculations**

# Takeaways

**The software and the application development will be a critical issue**

**There are not show stopping technical or political issues preventing more active participation by HPC centers**

**There are many common issues**

**A common document and strategy would be beneficial**