Report from US HPC Integration Meeting

Paolo Calafiura

(see Ian Fisk's talk Sep 5, 2024)

Meeting

BERKELEY LAB

The US meeting was held January 22/23, 2024,
in New York City @ the Flatiron Institute

13 people came, with 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.

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.

[Draft of the Meeting report]

Topics

Provisioning

Interfaces

Software

Data

Workflows

Provisioning

Topics:

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

Open Issues:

1. Lack of long-term provisioned allocations on HPC.
2. Frequency of scheduled and unscheduled downtimes

Status: Tolerable

1. submitting proposals yearly
2. exploring dynamic and elastic scheduling.

Impact: The absence of multi-year allocations and WLCG-level availability may render HPC facilities “unpledgeable”, at list for certain time-critical workflows.

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

Interfaces

BERKELEY LAB

Topics:

- **Standardized Interface Protocols**
- **Interoperability and Compatibility**
- **Automated Integration Tools**
- **Cybersecurity**

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

Status: Needs to be resolved. Encouraging signs on the federated identity and resource brokering fronts.

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 DOE HEP-CCE and Integrated Research Infrastructure, 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

BERKELEY LAB

Topics:

- **Software Parallelization**
- **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 through expensive re-engineering. 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 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.
  → HL-LHC TDRs, WLCG TCB, HSF CAF, HEP-CCE

Data and Storage

Topics:

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

Open Issues:

1. Current HPC sites offer limited on-site data storage, necessitating dynamic stage-in and –out. Data management systems must scale with the available HPC resources.
2. HPC WAN, edge nodes and file systems may not scale to HL-LHC requirements

Status: Tolerable, given that HEP data management systems are advanced. With the development of common interfaces for data transfer and access to future “data-intensive” 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**
    - **IaaS, FaaS**
  - **Workflow Orchestration**
    - **Distributed grid, cloud, HPC**

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, making progress with "as a Service" paradigm.

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.

Computer Architecture 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

BERKELEY LAB

The software and the application development will be a critical issue

There are no show stopping technical or political issues preventing more active participation by HPC centers

There are many common issues

A HEP-wide common document and strategy would be beneficial