Resource Management

Convener Wes Bethel (ewbethel@lbl.gov)
Editors

Questions

Questions to guide Discussion Rounds and Panel Discussions:

- ❏ What *resources* (compute, data) does your project need/use?
- ❏ How do you go about making use of them?
- ❏ What challenges do you face now/in the future in using them?
- ❏ if you could wave a magic wand and make the world how you want it what would it look like (w.r.t. *resources*).

Notes

- ❏ (Author) Your note for the report

Report

Draft for Workshop Report:

Group 1 discussion

Jerome, Mohammad, Torre, Alexander, Gene, Wouter, Sergei, Paul, Mike, James

- FAIR/ALICE
    - 100k cores for online and a few 100k for offline analysis for ALICE and for FAIR, a center @ 300k cores with CBM/PANDA 60k cores per experiments online. Tracking uses GPU.
    - Plan is to use online exclusively for online during beam but change the scheme and reconfigure to use the resources for analysis off beam time.
- LHCb - stick to CPU and use off run time for simulations.
- STAR - similar usage (or re-usage) of online resources off-run use now.
  STAR use of dedicated offline resources at BNL, Grid resources (OSG or dedicated sites from Europe) and NERSC/Cori HPC resources. All of what we heard about HPC and

that enables it (Containers, etc ...) are in use. Usage follows an event parallelization model (for real data reconstruction) and can be done due to a relative high memory / core. Later infrastructure may not be possible without solving software challenges.

- ATLAS - also HPC use on ACLF. Challenges discussed by Salman were all found and needed to be resolved. Moving data especially a real challenge.
- EIC R&D: Usage of OSG 4k cores maximum (average 2k cores). Generally HTC is more useful for the experiments (EIC).
- JLab - Mid-scale experiments, with data rates of 20-50 MB/sec at most. Almost no data reduction possible. Processed data is the same as the raw data. Every event is equally important. Challenges are cultural one
  - Small datasets have been used, set of users who want to look at the data as in the past (although the volume have changed). Will approach PByte scale of total data volumes.
  - Need rapid turn-around, need local resources + processed in real-time.
  - Analysis does not involve track fitting or similar - linear path, can be resolved in many ways. Due to the small size nature of the experiment, mostly a human resources.
- Resources being used by the JLab experiments 10 years old. Clever ideas of today's may mean that they will be the model in 10 years.
- LHCb - lots of reconstruction and calibration done in real-time. Triggerless readout in 4 years. 5 TB/sec into the software (i.e. reco + calib). This is a real challenge, model thought today are still off comparing to needs. Running reconstructions on those data-rates in real-time (10 ms per events scale). Vectorization of the algorithms to make tracking faster.
- Can any parts of the software that can be pulled out of LHCb re-usable by other experiments?
  - In principle yes …
- LHCb Speed: Simplified geometry found to be enough for some cases for example and speed up tremendously. Rate for the future - comes out as shooting for 30 PBytes / year, a few GB /sec. Reduction without physics.
- ATLAS - heard a lot of HPC ... and we heard that Clouds are not cost effective (different opinion in experimental side as per the cost). On the experimental side, Clouds are extremely effective - if all HPC hours would move to Cloud hours, we would be well placed.
- Innovation - need some area to prototype. At CERN, OpenStack was possible to leverage (at low cost). Having several sites to have OpenStack resources are valuable. Latencies are near 0 versus HPC resources where the learning curve is steep …
- STAR similar experience. NERSC/Cori as HPC worked well in the sense that the same prototyping was possible on a fast turn-around because we used "common" tools (Containers, OSG interface, small change on workflow for event parallelization). Beyond, would indeed be a large human investment.
- Where did the notion of Virtual Data-Centers go? Cloud/Grid - loan of resources is possible. Not really an HPC notion at the moment.

- If I start to need to move PBytes of data however, when is the breaking point at which we change the color of the resources? Issue of moving data on the network came up (but briefly): do we need to? Can we move data out of HPC?
- Data analysis foot-prints may also be the threshold for the choice of platform.
- Coming back on Cloud/Grids - Data ingress and outgress of Clouds (cost of IN/OUT). However, DOE has a deal with Amazon. Future program for Grid projects (OSG) also planning to create super-facilities (HPC, HTC, Cloud altogether behind a HT-Condor like interface). Does this change the game? The general feel is that it does. Whether HT-Condor or some other interface, a unified interface is not only easier but what end-users desire.
    - Using a new interface brings new Jargon and concepts (increase difficulty)
- What is the magick wand?
    - A well identified one is that everything would become a generic interface and accept the same workflow input.
- Is archival important? Not treated in depth - One opinion was that we never know if we are going to look at the data again (and how / does this fold with the idea of workflow preservation?)
- More important storage issues was thought to be caching system (like SSD cache). Create an environment (analysis or otherwise) where data is accessed fast. IO acceleration (on data that is accessed often). Is this being designed or generalized?
- Moving to container seems to be a generic approach and trend.
    - Container for small experiments may also help in making sure that no matter what, the workflow may still work on tomorrow's hardware.
    - Containers allowing to run a test on a laptop (no need to spend workforce for additional support).
- [Additional question] Imagine HPC are the only resources, restate the problems and questions (what are the challenges, ...).
    - Key point - What NERSC presented may be providing interfaces to steer the jobs, do analysis and so on is going in the right direction [that would be a challenge and requirement].
    - The Data challenge is a huge issue however - resolving this issue may be a human scale challenge.

Group 2 discussion

- What resources does your project need/use?
    - Small Experiments:
        - TB/week (though some more)
        - Larger experiments (PB/year)
        - Tape archive in the end

- § Calibration data often stored in other things like MySQL
- § Compute – what is supplied by facility is generally thought it was enough
- § Computing division at the labs – in charge of infrastructure, interact with the experiment groups to get everything setup

- o Ad-hoc notes
  - § Lots of small experiments – one or two months long, but at a national lab, collaborations are not good at figuring out the resources… Users look to staff at lab, but
  - § Small experiments – lab infrastructure, computing center, interactive analysis, temp disk space (few years), compute power (but generally enough is available), lab does work for DAQ etc., collaborate with new detector folks
  - § Labs/Facility staff have learned how to onboard new groups, 1 TB/week
  - § But smaller groups are starting to generate even more data
  - § Large experiments (like in the new halls): they are more careful. 10k core farm, few PB/year.
  - § Tape archive in the end, which is even less accessible
  - § Calibration data is done in MySQL

· How do you go about making use of them?
- o Collaborate with facility staff to get access
- o Persistent storage – the raw data is kept local in the facility for the long term. Some facilities don’t allow data to be accessed externally – need local account
  - § Do all computing locally at the lab on the data
- o All software is built locally, in a facility custom machine (with the facility configuration and software packages)
- o Calibration data (MySQL) can be shipped around
- o Have students and post-docs that come a good time before they show up for the run (with small experiments)

· What challenges do you face now/in the future using them
- o Networking is getting better – could reconstruction be done somewhere else and the raw data shipped around?
- o University collaborators often have clusters and resources; hard to use with the data
- o OSF – would like to be able to run on data
- o Even at JLab, there are 4 different infrastructures so things are quite different as you move around, even if they do similar things.

· If you could wave a wand and make the world how you wanted
- o Enough people to wield the wand.
- o It should be free and infinite bandwidth to move data around.
- o Compute is in pretty good shape
- o Compute memory, however – with KN, etc.
- o Infrastructure Reproduce your analysis
- o Most infrastructure differences are historical – you could do it better.
- o Different flavor of Linux, compilers, etc. Containers/VMs more?

- Interchange your data with another experiment – so your analysis can use another experiment’s data.
- Need people and resources to think about the future – not enough forward looking people.
- Experimenters know how to guess the required resources

Group 3 discussion

Wes, Gagik (presenter), Nathan, Mario, Salman, Laurel, Ethan, Eric, Ying, Varden, James

Resources need/use
Computing
HPC centers
“Farms” and departmental clusters located hither and yon

Data movement
ESnet for WAN movement seems adequate for telescope-based

Data storage
Storage at HPC centers (but comes with constraints)
Grow-your-own solution

How use resources

INCITE allocation for both storage and computing
INCITE storage: get multiple PB of storage, but only for the duration of the project
Buying own storage/drives as a long-term solution due to limits imposed by HPC facilities
$60K/PB, attached to a cannibalized, repurposed cluster
Good examples of collaborations with facilities

Challenges using resources
Time-limited nature of allocations at HPC facilities
NP doesn’t have as good a model for sharing resources as in, say, HEP (large collaborations)
Access to remote resources: each different institution has different access mechanism: one place requires use of VNC while other places use a different means
Variation in environments across different resources
Individuals/groups scavenge for access to resources
Mario mentions an example where an experiment is spec’d out to include adequate computing and storage.

Sysadm always takes more resources than expected to run a cluster: delegated to a grad student (who is not a professional IT person)
Absence of consistent authorization/authentication procedures creates barriers to use of distributed resources
Predicting resource usage: peak usage vs. steady state. Predicting peak loads, and time/duration of peak loads, is difficult.
Difficult for centers to provide resources to meet peak demands
Some communities, like bio, are able to make good use of cloud-based resources

Magic wand
Security but without intrusiveness or impediments
Public access to data (can't do it at JLab,ODU)
I/O bottlenecks on HPC platforms; want for better I/O capacity
Salman's slides: wanting compute platforms that are balanced for data-intensive workloads (existing systems optimized for numerical, compute intensive, batch-mediated workloads)
Making HPC resources an integral part of distributed workflows. E.g., real-time databases used in conjunction with simulations
Programming models/environments that better support changing computer architectures: early multi-threaded code was done by hand, didn't perform very well
Solution to problem where increasing complexity of software environments/tools is an impediment in terms of presenting a high learning curve
Want to make it easier for students/newcomers to quickly get up to speed, to make contributions

Notes from panel discussion:

- Minor course corrections now might result in major improvements in ten years from now
- future of HPC: solving technical bottle necks main remaining talks
- Magic Wand Item by Wes: Sociological barrier for mobilization persons to work towards common data formats and common tools, process has happen from bottom up (and not top down rom institutions and agencies),
    - Wes: Magic Wand Item: Sociological barrier for mobilization persons to work towards common data formats and common tools, process has happen from bottom up (and not from institutions and agencies),
    - Hardware barrier (Gordon): things are highly bound to hardware, hardware differences between research projects important, hardware abstraction layer difficult
    - Time barrier (Gagik): no time to look around what other persons are doing / have done, strongly related to the fact that it takes so long finding that information, persons would do that if that could be done in two hours, complex software hard to understand without good documentation, documentation is typically bad
    - Positive examples mentioned by Kyle: LarSoft, Geant Why have those projects succeeded?
    - reward structure: is finding a tool that solves a problem as rewarded as developing a solution for that problem

- Other Resources:
    - training data for deep / machine learning -> requirement for curated data, community-wide resources
    - quicker to write work-flow tool than re-using an existing one
- Types of data resources: raw data, calibration data, processed data, efficiency data, etc.

Additional Notes by Graham

Notes

- HPC systems are compute intensive but not good sources or sinks of data.
- Gordon Watts - DASPOS.
- Katie - NERSC

Discussion

Group 1

- Most experiments want to reduce data in real time or near time.
- Analysis after reconstruction is modest.
- Large experiments 50 to 100k cores needed to keep up with the data flow.
- Large HPC datacenters.
- Difficulty is moving data to these places since they are not designed with massive in/out data flows
- GPUs have been used successfully.
- People moving to take advantage of vectorization
- Efforts to speed up reconstruction which is where the bulk of the processing of experimental data comes in
- Small experiments have limited expertise with these large resources
- Simulations tend to be done opportunistically, example grid.
- Data resource needs vary widely depending on experiment, from TB to PB x100 per year
- Future generations of HPC may become severe challenges because of a trend to move away from HTC. (What is HTC?)
- Cloud does work but commercial data transfer is an issue that needs work.

Group 3 - Gagik

- Sharing computational resources with universities. Creating a system to make use of them is difficult because of manpower issues 1/2 grad student rather than an IT professional.
- No good model for NP for collaborating universities to contribute computing resources.
- Shifting to programming models etc that better make use of shifting computing architectures. Multi-threaded code is difficult to do efficiently.

Group 2 - Gordon

- Small experiments, disk space etc a TB each week or two, larger ones PB per year. Long term storage, tape etc was not thoroughly discussed.
- Calibration data stored in databases.
- Compute did not come up much, most people have enough compute power at their disposal for current needs.
- National lab computing groups are a valuable resource that is often underused.
- Collaboration between user groups and facility staff is important.
- Accessibility of data from offsite is spotty.
- Training of post-docs etc. People need to show up and get trained, not just dive in and mess with code.
- Disk storage is more expensive than the associated computing.
- OSF is a place where work could be offloaded to but data volumes often restrict that to simulation
- Inconsistencies between infrastructures within facilities.
- It would be great to have enough people to do all that needs to be done.
- Need more inter site bandwidth so that data can be moved.
- memory, models are evolving the wrong way for us, not enough memory per core for current algorithms.
- Automation.
- User - software interface should be more similar between experiments and sites.
- So much “now” being done that the future is being forgotten, always need to look ahead but people don’t have time to make minor course changes.
- Estimations of resource requirements need to be better understood.

- Jerome - How do you see seamless interfaces and existing systems converging?
    - Salman - Two things, HPC environment and the interface to it. This is not a technical challenge it is a policy challenge. People don’t communicate, security protocols differ, computing policies differ. The right decisions need to be made in the right places
- Wes - Think about resources, sw, hw, tools, data etc. If I had a magic wand I’d eliminate the sociological issues, such as “not invented here”.
    - Gordon - How do you make software reusable? Don’t want to force everyone to have identical hw because diversity at some levels is good. The sociological issue is that everyone thinks that they are the only ones who know.
- Gagik - Physicists do things because they have a job to do. A task is the focus not writing nice software that can be shared etc.
- Gordon - Today it is often harder to find and adapt a solution than to write something from scratch.
- Martin - A deterrent is that it is hard to assess what someone else’s software does and if it works well. Lego™ again.

- Gagik - Part of the issue is documentation. People write software until it works then use it but nobody likes to write the documentation.
- Audience - In neutrino community they do have various shared packages.
- Gordon - Once a piece of software reaches the point of common knowledge and use you are inspired to invest the time needed to use it because you have confidence in it’s usability. For small packages that are less well known having them well packaged and of well defined functions that are quickly understandable and testable so that people can become quickly confident and comfortable.
- Wes - Next session is machine learning. There is a lot in there that is germane to this discussion. You need to train AI systems, curated datasets and strong definitions are required to give the AI something to work on.