

Joint USATLAS-USCMS HPC/Cloud Workshop Summary

F. Barreiro Megino, L. Bryant, D. Hufnagel, K. Hurtado
WLCG Workshop, 8 Nov 2022, Lancaster University

Based on summary slides prepared by Dirk Hufnagel and Lincoln Bryant

Overview

Why did we do this?

- US LHC Operations Program review (Feb 2022) recommendation
 - Blueprint document about Cloud/HPC usage and evolution in U.S. ATLAS and U.S. CMS
 - Internal document deadline: Dec 2022

Workshop

- Organized workshop as part of the process, held Sept 26-28 (3 half days)
 - <https://indico.cern.ch/event/1183995>
- Hybrid format: meeting room at UChicago, but majority of people connected remote
- Good attendance, about 30-40 participants across ATLAS, CMS, OSG-LHC and WLCG
- Deliberately light on presentations and heavy on discussions
- Few invited contributions: HPC facilities, ESNet, CERN, Vera Rubin

Discussions about HPC

HPC Landscape in US

- User Facilities (NERSC, ACCESS, Frontera)
- Leadership Class Facilities (LCF), more restricted, but also (much) larger

Integration

- Integration frequently depending on HPC architecture/restrictions
- Streaming to worker nodes or intermediate storage within HPC
- Raises network/storage requirements at associated sites
- Large(r) storage allocations at LCF should be explored (incl. full DM integration)
 - Also potentially at User Facilities, but streaming is easy there, so no pressing need

Workflows

- Currently do what is easiest considering storage integration
 - ATLAS: Simulation
 - CMS: almost all MC production through StepChains
- Workflow “specialization” doesn’t help on the pledging discussion (more on this later)

Discussions about HPC

Cost

- No direct cost to experiment for computing resources
 - Indirect effects at the funding agency levels about replacing our own resources with HPC were outside the scope of the blueprint process (but we had a short discussion during the international HPC discussion)
- Cost at personpower level
 - Proposals time and effort consuming
 - Commissioning and support costs (have to provide interface between experiment and HPC)
 - Software porting to new architectures/accelerators outside the scope of the blueprint process
 - In some cases manual operation (e.g. task assignment)

Proposals

- Planning horizon mismatch between allocation process and experiment need
- LCF proposals currently require selling our use cases, not very interested in generic T2 production activities

Discussions about HPC

Edge Services

- Many HPC seem to experiment with platforms for edge services (often Kubernetes based)
 - No overall (visible) coordination in this area
- Can we take advantage? Can there be consolidation to simplify integration?

CVMFS

- Less of a worry these days: *cvmfsexec* works on any modern kernel/OS
- Needs edge service squid (which so far seems to be possible everywhere, even LCF)

General comments about our HPC usage

- We are not treating HPCs for what they were designed (e.g. MPI, GPUs)
- We use HPCs like another T2 and we are an average user for HPCs (vs e.g. LatticeQCD)
- Community needs to make jump to GPU or become legacy computing and accept limitations

Discussions about Cloud

Cloud Landscape

- Considering mostly global, large providers (Google, Amazon, Azure)
- Briefly discussed Lancium
 - Low cost due to use of cheap renewable energy, but also limited features

Integration

- Site extension vs. standalone cloud site
- Storage integration done, but policy to use only for transient replicas (no custodial data)

Workflows

- No technical restrictions. In principle we can run everything
 - However some workflows can be more interesting from cost perspective
- Interactive analysis setups

Elasticity

- Allows evaluation of “exotic” types of hardware (ARM, GPU, FPGA) in a cost effective way
- Allows using (and paying for) capacity as needed
 - “Scheduling for peak”: possibility of high intensity / short runs vs. spreading out use over time on fixed resources

Discussions about Cloud

Cost

- Main discussion topic for cloud
- Cost comparison for CMS Reconstruction Scale Tests ~5 years ago
- ATLAS Cost Simulation
 - Plugs in publicly available numbers for on-demand vs. pre-commitment, spot vs. regular, AWS vs. Google Cloud
 - Allows to estimate costs for various combinations of compute, storage and egress
- ATLAS Google Cloud Subscription
 - Fixed price for certain amount of resources for a certain time period (15 months here)
 - Negotiated by each customer
 - Flexible resource use is still possibly (could use the total compute in subscription in first month in principle)
 - Keeps a running total of “charges” using list prices, but everything is included in subscription (also egress)
- Facilities also looking at this
 - Programs can buy Cloud cycles directly or T1/T2 sites could decide to replace on-prem with Cloud cycles
- Outlook
 - Prices have dropped over the years, should be ready to use them once the cost becomes attractive
 - Cloud providers want to make money, but they are also in competition with each other
 - Cloud providers have economies of scale we cannot achieve on our own

Discussions about Networking

Invited talk from ESnet

● DOE HPC

- ESnet6 built physical network into each DOE national lab
- Site infrastructure (border routers, security stuff, storage) will be a more limiting factor than WAN
 - Wide disparity in HPC support for data-centric workflows

● NSF HPC

- Some very well connected, others not so much, depending on where they are

● Cloud providers

- Multiple 100Gbit interconnects
- Enables capabilities. Network cost still needs to be understood

● LHCONE

- Participation in LHCONE defacto based on IP address, effectively excludes HPC/Cloud traffic
- Possible workaround is to keep traffic partitioned (Global LHC <-> US LHC <-> US HPC/Cloud)
 - Puts more load on US sites and also means less flexibility in using HPC/Cloud

Discussions about Pledging

- **Current HPC/Cloud use is opportunistic / “beyond pledge”**
 - One view point: can just keep doing this and not care about pledging / getting credit
 - Resources can't really be considered for planning purposes
 - Problematic if HPC/Cloud use increases a lot
- **Current MoU policies require pledging a fixed #cores**
 - Incompatible with US HPC resource provisioning policies
 - Just a customer, we get in the queue with all other users and our jobs run ‘eventually’
 - Would need a change to allow time integrated pledge (kWh vs kW)
- **Pledges currently are for universally usable resources**
 - How restrictive can we be with workflow selection for them to be eligible as pledges?
- **How to pledge GPU resources? (not just a HPC/Cloud problem)**
- **Indication to resolve the technical questions (benchmarking, accounting) will open up the discussion on more flexible policies**

Discussions about Benchmarking/Accounting

Benchmarking

- Invited talk from CERN about HEP benchmarking on HPC
 - One conclusion is that increased heterogeneity affects benchmarking
 - Not just one number anymore to describe performance, at minimum CPU vs. GPU
 - Lots of variations in GPU architectures, how to compare them among each other?
 - Even more complex once more use cases like AI/ML get added

Accounting

- HPC/Cloud usage would have to get reported (via APEL)
 - U.S. ATLAS doesn't do this currently
 - U.S. CMS spent some effort last year on this, some HPC still not correct
- Only talking about traditional CPU accounting here
- Have nothing for GPU accounting (not just an HPC/Cloud problem)

Discussions about R&D

AI/ML

- Question whether this is a large enough use-case to matter
 - Hyperparameter optimizations potentially very expensive
 - Invited talk from CERN about training and hypertuning on HPC
 - Continuous training models were mentioned
 - Could matter, depending on where we go with AI/ML

Other R&D topics

- Unique cloud offerings (special accelerators, FPGA etc)
- LCF integration work (to bring operational costs down)
 - This includes work on how to use the edge service platforms
- Full-chain: minimizing egress by running all steps in the cloud
- Function-as-a-service (funcX, parl)
 - Primarily Cloud, but potentially also applicable to HPC
 - Not much usage now, will have to see how this evolves
- External inference servers on GPU (SONIC)