Review on current workflows and production on HPC centers

Dirk Hufnagel (FNAL)

2019 Joint HSF/OSG/WLCG Workshop

19th March 2019

Why HPC ?

- HL-LHC (Run4) starting in 2026 will be something very different
- Trigger rates increase x10
- Reco times x20 due to pileup
- Event sizes go up x5

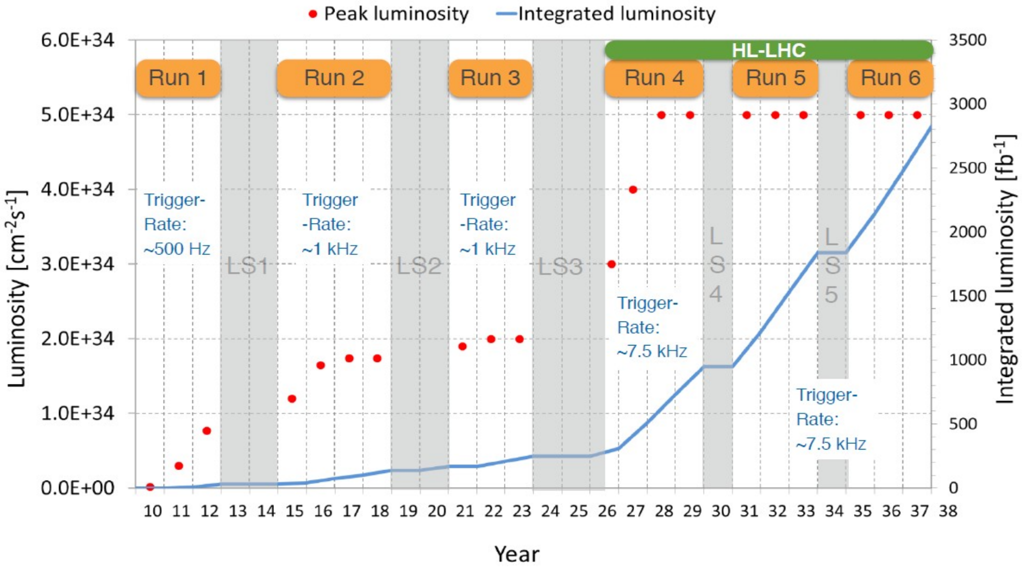

- On the other hand, the push to Exascale leads to rapid advancements in HPC capacity, with planned Exascale HPC deployments in the early 2020s (US, China, Europe).
- Encouragement by funding agencies to use these HPC resources
- HPC provide challenges, but also opportunities

Focus of this talk

- I am concentrating on LHC computing, both ATLAS and CMS.
- Will say some words about Neutrino science (at Fermilab) later. Other sciences also look at HPC, either to complement their own resources and/or because they don't have much owned resources in the first place.
- Can't go into details for every HPC, more overview.
- Mostly ignoring very grid friendly HPC (they do exist) and site extension schemes to local or quasi-local HPC.

CMS/ATLAS Workflows

- Both ATLAS and CMS follow a set of production workflows for both data and MC
  - For MC : Generation / Simulation / Digitization / Reconstruction
  - For Data : Prompt and Re-Reconstruction
- ATLAS and CMS differ somewhat in which workflows they spend most of their cpu budget on now. Reconstruction will just get more and more expensive compared to simulation with increasing pileup, so for HL-LHC both should be dominated by reconstruction (assuming (N...N)LO generator creep can be controlled).
- Other workflows exist (skims...), but this is the bulk of the production activity.
- Different interaction with site if chaining workflow elements within same job

HPC are different

- HPC come in all shapes and sizes
  (“they are like snow flakes, everyone is different”)

- From something very similar to an LHC grid site cluster with some of our middleware missing (best case even a willingness to install some of that middleware) to exotic hardware architecture with alien (to what we are used to from grid computing) software environment and system integration.

- Usually the latter are the ones with the most resources and consequently the most interesting for us to target.

Common solutions

- Common solutions for integration of HPC are difficult, due to their unique nature
- Still have common issues, so common approaches are possible, even if implementation details will differ
- CMS recently compiled a list of requirements for HPC sites, a list of challenges if they are not fulfilled and a list of technical solutions (the challenges are generic, the technical solutions are somewhat CMS spefic)

Tech Doc   Exec Summary / Political Introduction

Integration issues

- My view, which contributed to the document mentioned before, but is not necessarily 100% identical. In the interest of time also not talking about every single integration issue, this is an overview.

- Runtime environment

  Hardware architecture, OS+packages, VO software, conditions access ?

- Workflow Management integration

  How to get production jobs onto resource, how to track them, how to get results back ?

- Data Management integration

  How to get data to jobs, how to get job output back to VO ?

Hardware architecture

- LHC computing is built on x86(-64), (almost) exclusively used for >10years.

- While there are still x86(-64) HPC and will likely remain for the foreseeable future, x86-64 is not (by far) where the biggest additions in terms of compute power are expected.

- There are also other general purpose CPU (Power, ARM), which are ~easy to adjust to. But the biggest additions in compute power are expected in accelerators (for now this mostly means nVidia GPU).

- This is simply a FLOPS/Watt issue. Even if in a few years GPU go out of favor, they'll likely be replaced with some other accelerator flavor or specialized co-processor.

Hardware architecture – How to use GPU ?

- To take full advantage of the push to Exascale, we need to learn how to leverage the compute power of GPU. Otherwise we won't be able to use the full power of the newest HPC.
- There are indications that running CPU-only workflows on such HPC will not be allowed (but unclear how that is policed and what the thresholds are).
- ATLAS : Current activities mostly for ML, using GPU in ML to then extract trained model and speed up production on CPU by use of inference. Unclear whether this is sufficient use of GPU to speed up production. Likely not.
- CMS : Fully integrated into multi-threaded framework, prototype exists. Tracking inside reconstruction being ported to GPU. Question remains if enough work can be offloaded to GPU to make sufficiently efficient use of them.
- Similar to non x86-64, GPU will also cause physics validation challenges.

Runtime Environment – OS+packages and VO software

- Both ATLAS and CMS use containers (shifter and singularity seem to cover most US and European HPC).

- VO software is available via cvmfs (some successes in convincing HPC to provide it, unclear if this will work everywhere, even on a HL-LHC time scale). Large containers (100GB to multiple 100GB) are another option. Copying software to shared FS also possible, but often has performance/stability problems.

- Non-x86-64 architecture can cause problems with VO software availability and that can cause problems with availability of workflows to run. For instance, ATLAS only ran alpgen on Mira, due to the PowerPC architecture. If we want to expand to fully use Power or ARM, we will have to compile and validate our software stacks on these architectures.

Workflow Management Integration

- How do we get our pilots onto HPC nodes and have them acquire work?
- Highly dependent on HPC what options are available
- Grid-friendly HPC with CE (CSCS/Switzerland for instance)
- Worker nodes have access to internet
  - Integration with HTCondor possible via remote ssh through Bosco
- Worker nodes have no internet
  - More difficult, need edge services running on gateway nodes

Workflow Management Integration - ATLAS

- HPC with CE are used like normal grid sites
- Standardizing on Harvester edge services
  - Even for HPC where worker nodes have internet access
- Event service (to improve efficient use of all cores until end of batch jobs)
- Yoda for MPI “jumbo-jobs” (improves efficiency at some HPC)
  - Event-level processing in MPI master-client model
  - Integrates seamlessly with ATLAS event service

Workflow Management Integration - CMS

- HPC with CE are used like normal grid sites

- Going for full integration with HTCondor Global Pool
  - HEPCloud portal at FNAL to access US HPC (see Burt’s talk)
    - Some of the US HPC connected to HEPCloud via OSG supported Hosted-CE
  - HPC without internet access from worker nodes will require edge services at HPC
    - HPC HTCondor pool federated with HEPCloud or HTCondor Global Pool
    - If needed shared FS at HPC can be used for internal HTCondor communication
    - Commissioning work has started (MareNostrum via PIC)
  - CMS only runs normal jobs so far. No event service planned at the moment.

Data Management Integration

- How can we get data to jobs, how can we get job output back?
- CMS uses remote streaming for input and remote stageout
  - Also do static placement of certain samples (pileup)
  - Adopting Rucio, so hope to benefit from common ATLAS/HPC/Rucio integration
- ATLAS ARC Data Access : aCT and ARC CE cache data in HPC local storage
- ATLAS Harvester : does transfer itself (in different ways) or uses Rucio

Scale of operations CMS

- Only 2 US HPC in regular production use so far (NERSC and PSC)
- In 2018 used ~42M hours for CMS at NERSC and ~5M at PSC

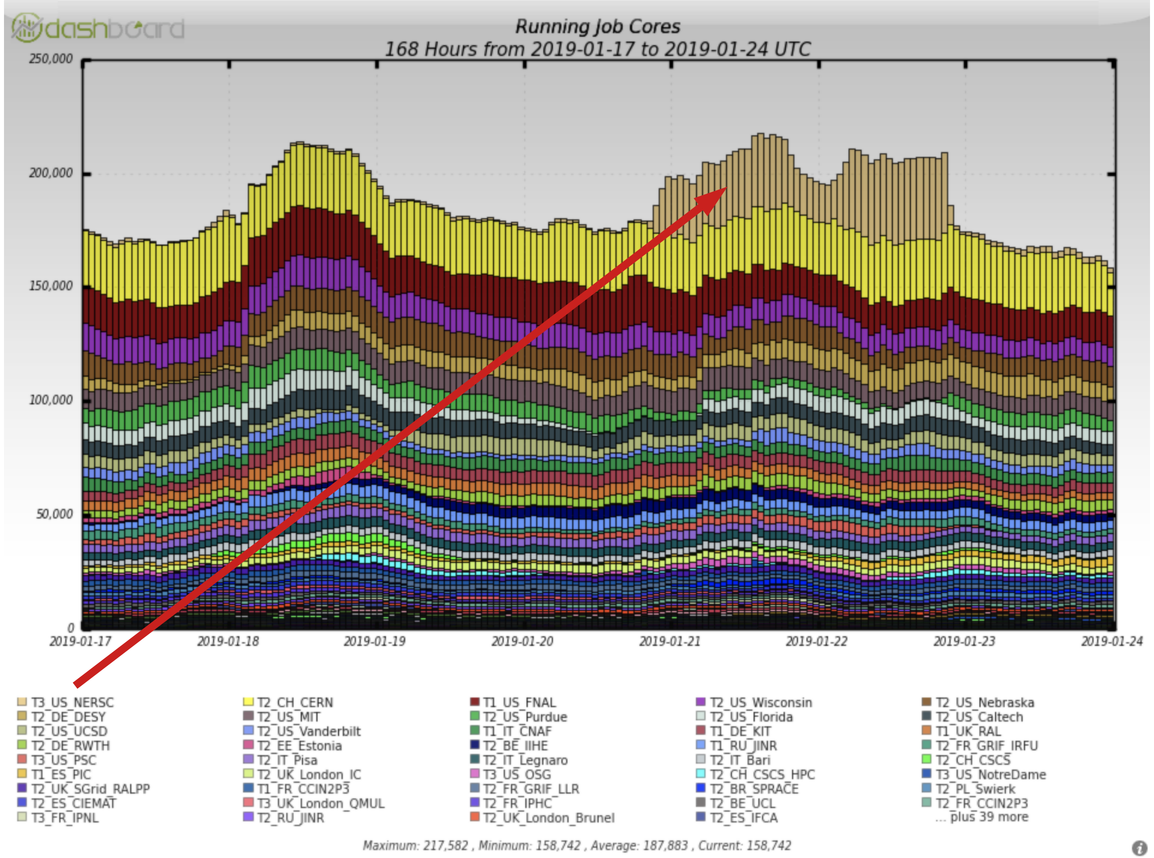

Fermilab

Scale of operations ATLAS

- In Europe using a variety of HPC, altogether 10k to few 10k cores
- In US using NERSC, ALCF, OLCF and others, total 2018 use multiple 100M hours

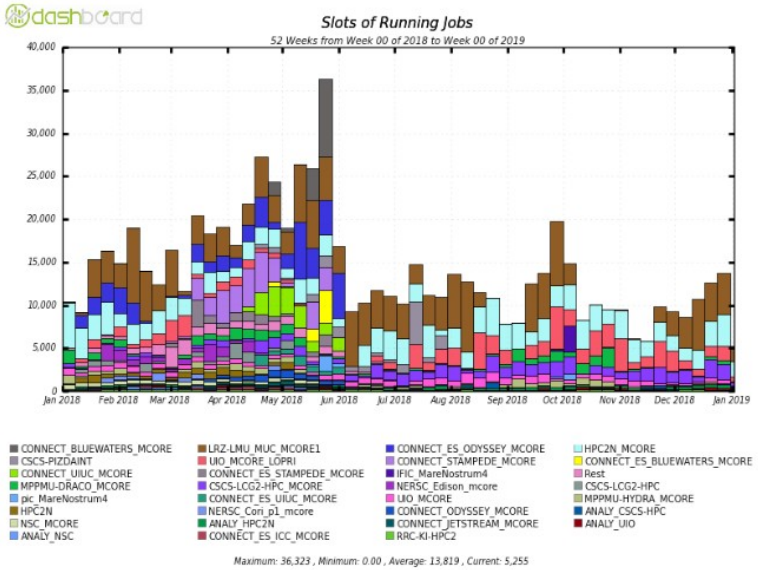

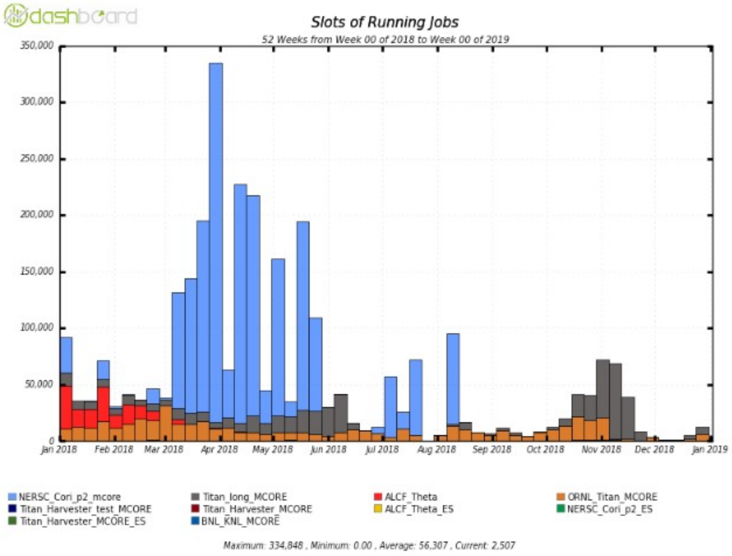

Workflows being run

- ATLAS is (so far) almost exclusively running Generation and Simulation
  - This is where ATLAS spends a large fraction of their cpu budget
  - These workflows are well suited to run on HPC resources

- CMS always tries to commissioning resources to be able to run all it’s workflows (with some exceptions like classical pileup digitization, pileup library generation, skimming etc)
  - Ran Generation, Simulation, Digitization and Reconstruction (NERSC/PSC)
  - NERSC(2018): ~1200 unique workflows, ~1800 unique workflow steps

Future efforts

- ATLAS is trying to allow a larger number of workflows to use the event service
- ATLAS is working on ML on HPC and other GPU friendly workflows
- CMS is trying to commission the ability to integrate HPC with private network nodes into our Global Pool. Currently the largest roadblock to integrate the DOE LCF’s and also some large European HPC.
- CMS is also working on using GPU at HPC
  - NERSC NESAP grant to port algorithms to (Perlmutter) GPU

Other sciences

- The Neutrino experiments at Fermilab will be able to use the HEPCloud portal to access the same HPC that CMS has access to. Nova already used 30M hours last year (reservation for 2M hardware threads).

- For smaller experiments HPC can have a huge impact, allowing new science use cases that weren't possible before. Unfortunately, smaller experiments often don't have much offline&computing manpower/expertise and rely on others to provide tools and/or resource access methods.

- Dune's current plans for HPC/GPU are mostly for ML. Want to use ML on HPC/GPU, extract the models and use inference in regular processing. Builds on work done in previous neutrino experiments
  https://science.energy.gov/hep/highlights/2019/hep-2019-01-c/

Conclusion

- Lots of activities, many different approaches (sometimes all used in parallel)
  - Makes it hard to summarize this topic
  - Somewhat unavoidable given the “snow flake like” nature of HPC

- Given the unique nature of HPC, we don’t know if the integration methods we develop today will be fully applicable to tomorrow’s HPC. Hardware architecture can also change (see for example the Aurora21 press release). But what we learn now is still useful, it helps build relationship with the HPC facilities and can also give us arguments to try to influence technical choices for future HPC.

  https://www.alcf.anl.gov/articles/us-department-energy-and-intel-deliver-first-exascale-supercomputer