HEP Computing Model Evolution/Revolution

By Oliver Gutsche

"A white paper collecting facts/ideas/concepts about HEP Computing Models and their possible evolution or revolution from CHEP 2016 mixed with a healthy dose of own opinion."

The traditional HEP computing model in the LHC era consists of the three main forms of computing resources CPU, disk and tape deployed in various distributed sites, interconnected through strong networks and used through GRID infrastructure. Workflows are orchestrated to use all the distributed CPU resources, data is either automatically or manually distributed across the sites or streamed to the jobs. Newer resource forms like commercial clouds or HPC centers are integrated as of lately. Centrally written software is used to simulate, digitize and reconstruct events. Analysis code is written by physicists and run on the events, either centralized in trains or individually using the GRID. All is based on the deployment of x86 redhat-cloned linux.

What we currently see and will see in the near and medium term future is the dissolving of the strict tiers of resources of the LHC thanks to the increased reliability of resources on lower tiers. We will also see the availability of elastic CPU resources with the possibility to burst to high scales quickly for a limited amount of time. We will though see the end of the simple life of one architecture and will have to start supporting multiple platforms (like ARM, PowerPC, etc.). Virtual machines are a way of easily supporting multiple architectures. They are also used to encapsulate and transport the environment in which scientific software is running. They are widely used on commercial clouds. The trend though goes to containers, which are smaller and easier to deploy. In a containerized world, networking becomes an issue, which will be solved through SDNs merging LAN and WAN.

The next stage of optimization will come (again) through the network. Storage and compute will be more and more loosely coupled, caching will become important as the trend goes to fewer but larger storage facilities to lower the maintenance and operations cost. The network will need to become faster still, but we need a full optimization of WAN and LAN infrastructure and all used I/O layers in between. Multi-level caching should be built in the networking infrastructure rather than bolted on top of it. We will also need to learn how to share the networking bandwidth with other disciplines whose data needs grow significantly. Also here SDNs are expected to be the solution, although one can be sceptical because the SDN revolution proclaimed for many years hasn't arrived yet.

In the long term, we cannot rely anymore on technology to solve our computing challenges. The device markets (smartphones, tablets, PCs, notebooks, servers, HPC) are saturated and show negative growth. New developments don't reach return rates as they used to, so industry development is slowing down. As a consequence, market dominance of few companies is increasing and competition diminishing.

The community is very good at data management and distribution, and this will be the key to tackle the future challenges. We will see the switch to the utility provider model for CPU, where our strength takes care to store data and provide access to it. We will concentrate storage in a few large facilities and use the network to orchestrate the data movement in between them. The CPU from our utility providers accesses the data through streaming and caching. The utility model will provide elastic access to CPU resources. This model will not be transferred to our data. First of all no external provider should be willing to accept custodial responsibility for our data because it is very valuable, also once we couple the storage and processing we lock into providers and loose leverage.

We also have to look closer at our analysis workflows. The estimated data volumes of the DUNE and HL-LHC era will like make our current analysis model obsolete. New concepts favor to use the few data storage facilities as data reduction centers where large amounts of data are quickly reduced to small analysis datasets that are made accessible through the caching and streaming infrastructure. We will see the emergence of new analysis technologies based on industry standards, giving HEP access to the large community of data science emerging in industry and academia alike.

We might also look into replacing our traditional event loop based processing model. Transaction models that can schedule (manually or automatically) the transformation of parts or all of the data at our storage facilities promise to increase the efficiency and lower the operational cost of our workflows.

References:

- Brian Bockelman's WLCG talk:
    - https://indico.cern.ch/event/555063/contributions/2319943/attachments/1350810/2039114/OSG-WLCG-Workshop-Data-2016.pdf
- Shawn McKee's WLCG talk:
    - https://indico.cern.ch/event/555063/contributions/2315805/attachments/1350764/2039165/LHC_Network_Evolution_WLCG_workshop.pdf
- Helge Meinhard's and Bernd Panzer's WLCG talk:
    - https://indico.cern.ch/event/555063/contributions/2285842/attachments/1350817/2039143/2016-10-08-WLCGWorkshop-TechnologyEvolution.pdf
- Ian Fisks CHEP plenary talk
    - https://indico.cern.ch/event/505613/contributions/2323233/attachments/1350914/2039442/Outline_IMF_CHEP16_V1.pdf
- Track 3 summary talk:
    - https://indico.cern.ch/event/505613/contributions/2331046/attachments/1355033/2047461/summary_track3_chep_2016_lb.pdf
- Atlas Hl-LHC computing model parallel talk:
    - https://indico.cern.ch/event/505613/contributions/2241721/attachments/1344208/2032301/Oral-90.pdf