

Opportunistic, Cloud, HPC

There were many talks on this topic, sometimes in parallel in different tracks. Tracks with most talks on this topic were in

Track 3 : Distributed Computing

Track 5 : Infrastructure

but some other tracks also had a few talks related to this topic. IMO the general underlying theme of CHEP 2016 was the upcoming HL-LHC run on the horizon and the computing challenge associated with it. People are crunching the numbers in terms of expected computing needs and compare it to expected funding and expected algorithmic software performance increases and these numbers don't match up. Therefore there is a lot of interest in pursuing non-traditional resources in the hopes that this will help with the problem.

The many talks about Opportunistic/HPC resources can be understood in this context. On this topic there many more ATLAS talks than CMS talks. It also seems there is a lot of internal competition in ATLAS since there were multiple talks on Opportunistic/HPC that overlapped partially, where different groups are kind of but not quite doing the same thing.

On the Cloud side things look a bit different. Everyone is using it in some way and both ATLAS and CMS are looking into it to be ready for when it becomes cheaper to do computing in the cloud and/or to handle load spikes. Facilities also look into expanding their compute resources seamlessly into the cloud. There were also some talks about Meta-Clouds, efforts to provide interoperable clouds for research communities with common API and such.

The site extension schemes, notably HEPCloud, also look at extending into HPC resources, not just clouds.

Talks:

- Exploiting Opportunistic Resources for ATLAS with ARC CE and the Event Service
 - Using ARC CE and an arcControlTower interface to Panda to integrate resource that have no outbound network. Data staging in and out handled by the ARC CE. Events go individually into an Object Store, merging is done on regular grid resources.
 - <https://indico.cern.ch/event/505613/contributions/2230710/>
- ATLAS and LHC computing on CRAY

- From a Swiss ATLAS group, also using ARC CE and arcControlTowers, but on Cray, with some other technical challenges on top of it (no cvmfs, using shifter like CMS is doing at NERSC) . Not using the event service.
 - <https://indico.cern.ch/event/505613/contributions/2227411/>
- Integration of the Chinese HPC Grid in ATLAS Distributed Computing
 - Talk on ATLAS using Chinese HPC sites. These are extremely locked down, was interesting to see the approach to make them usable. Again using ARC CE and arcControlTower with additional tweaks to work around the peculiarities of the environment (software install only on request by sysadmins after vetting and approval for instance).
 - <https://indico.cern.ch/event/505613/contributions/2227415/>
- The OSG Open Facility: An on-ramp for opportunistic scientific computing
 - Bo talking not so much about LHC computing, but about the OSG providing an inter-experiment and even inter-science computing infrastructure.
 - <https://indico.cern.ch/event/505613/contributions/2227426/>
- Opportunistic Computing with Lobster: Lessons Learned From Scaling up to 25k Non-Dedicated Cores
 - Opportunistic resource use, but only within Notre Dame, so while interesting not that useful for others.
 - <https://indico.cern.ch/event/505613/contributions/2230715/>
- The role of dedicated computing centers in the age of cloud computing
 - Brookhaven looking at the resource evolution, both for on-premise resources and external cloud resources and how the balance between them might evolve.
 - <https://indico.cern.ch/event/505613/contributions/2227410/>
- Extending the farm on external sites: the INFN Tier-1 experience
 - These external sites include commercial clouds.
 - <https://indico.cern.ch/event/505613/contributions/2227438/>
- Networking - View from ESnet (Plenary)
 - Many other things in this talk, but one question related to cloud. How will we integrate all the commercial cloud providers worldwide into our networking infrastructure ?
 - <https://indico.cern.ch/event/505613/contributions/2323222/>
- The HEP Cloud Facility: elastic computing for High Energy Physics
 - Fermilab HEPCloud project
 - <https://indico.cern.ch/event/505613/contributions/2230739/>
- Opportunistic resource usage at Argonne Facility for CMS
 - Looking at using HPC resources to do NLO MC generation for CMS. Not at all close to have something usable, really more an exploratory effort for now.
 - <https://indico.cern.ch/event/505613/contributions/2230732/>
- Production Experience with the ATLAS Event Service
 - This looks at the Yoda iteration of the ATLAS event service, the one that uses MPI on NERSC HPC resources.
 - <https://indico.cern.ch/event/505613/contributions/2227438/>

- Extreme I/O on HPC for HEP using the Burst Buffer at NERSC
 - Describes a new fast flash based IO layer at NERSC to support very IO intensive jobs (ATLAS makes good use of it).
 - <https://indico.cern.ch/event/505613/contributions/2227423/>
- Clouds and Competitors (Plenary)
 - General overview over clouds, also talks from Google and Microsoft, followed by a plenary discussion.
 - <https://indico.cern.ch/event/505613/contributions/2323243/>
- The Cloud Area Padovana: from pilot to production
 - Describes a common OpenStack cloud infrastructure between two Italian sites.
 - <https://indico.cern.ch/event/505613/contributions/2230717/>
- On-demand provisioning of HEP compute resources on cloud sites and shared HPC centers
 - Another talk about site extension to external resources, both clouds and HPC. This time from KIT in Germany.
 - <https://indico.cern.ch/event/505613/contributions/2230729/>
- Context-aware distributed cloud computing using CloudScheduler
 - A cloud scheduler that can send jobs to on-premise and off-premise (multiple cloud) resources which has been in production for years already.
 - <https://indico.cern.ch/event/505613/contributions/2230405/>
- Using Shifter to Bring Containerized CVMFS to HPC
 - Describes the NERSC shifter system, a software solution that supports deploying custom software stacks via (docker) containers at NERSC.
 - <https://indico.cern.ch/event/505613/contributions/2227429/>
- Improved Cloud resource allocation: how INDIGO-Datacloud is overcoming the current limitations in Cloud schedulers
 - Trying to do smarter scheduling on top of cloud resources
 - <https://indico.cern.ch/event/505613/contributions/2227934/>
- Design and Implementation of Elastic Computing Resource Management System with HTCondor on OpenStack
 - Using HTCondor to schedule OpenStack resources
 - <https://indico.cern.ch/event/505613/contributions/2227931/>
- Interfacing HTCondor-CE with OpenStack
 - Using HTCondor to submit work directly into OpenStack resources
 - <https://indico.cern.ch/event/505613/contributions/2227921/>
- Contributing opportunistic resources to the grid with HTCondor-CE-Bosco
 - Integrating BOSCO style submission directly into the HTCondor-CE
 - <https://indico.cern.ch/event/505613/contributions/2227941/>
- SCEAPI: A Unified Restful Web APIs for High-Performance Computing
 - API to interact with resources and manage software installation and other things. Used for the ATLAS Chinese HPC tests mentioned in an earlier talk.
 - <https://indico.cern.ch/event/505613/contributions/2227922/>
- Experience in using commercial clouds in CMS

- FNAL AWS CMS tests using HEPCloud
 - <https://indico.cern.ch/event/505613/contributions/2230724/>
- The HNSciCloud project
 - Helix Nebula Science Cloud : European/CERN project to provide multi-provider and multi-science cloud
 - <https://indico.cern.ch/event/505613/contributions/2230727/>
- Technical challenges of HNSciCloud
 - More on the Helix Nebula Science Cloud
 - <https://indico.cern.ch/event/505613/contributions/2230728/>
- Volunteer Computing Experience with ATLAS@Home
 - ATLAS using volunteer computing for MC simulation using BOINC
 - <https://indico.cern.ch/event/505613/contributions/2230707/>
- CERN Computing in Commercial Clouds
 - CERN is evaluating commercial cloud providers in Europe. Resource mix for Run3 could be Meyrin resources plus a variety of cloud resources all over Europe all provided under a common CERN site umbrella transparent to the user. These efforts have to be seen in the wider context of the before mentioned Helix Nebula Science Cloud as well.
 - <https://indico.cern.ch/event/505613/contributions/2227325/>
- INDIGO-Datacloud: a Cloud-based Platform as a Service oriented to scientific computing for the exploitation of heterogeneous resources
 - European project to provide a computing platform for scientific communities. Designed to be deployed on public or private Clouds, integrated with existing resources or e-infrastructures.
 - <https://indico.cern.ch/event/505613/contributions/2227435/>
- TOSCA-based orchestration of complex clusters at the IaaS level
 - Technical provisioning layer on top of different cloud providers (OpenStack, OpenNebula etc) to make them accessible via the same mechanism.
 - <https://indico.cern.ch/event/505613/contributions/2227439/>
- Stealth Cloud: How not to waste CPU during grid to cloud transitions
 - How to seamlessly transition grid resources into the cloud but making the grid jobs run on cloud without them even knowing it. Using ARC CE, glideInWMS and OpenStack.
 - <https://indico.cern.ch/event/505613/contributions/2227436/>
- Mixing HTC and HPC Workloads With HTCondor and Slurm
 - Combining traditional HTC resources and RHIC/ATLAS with HPC resources and use HTCondor and SLURM to facilitate opportunistic use both ways
 - <https://indico.cern.ch/event/505613/contributions/2227416/>
- CMS use of allocation based HPC resources
 - CMS strategy and progress with allocation based HPC resources.
 - <https://indico.cern.ch/event/505613/contributions/2227945/>
- Integration of the Titan supercomputer at OLCF with the ATLAS Production System

- ATLAS is running on the Titan HPC facility at OLCF in “backfill” mode. They do this via query the system for idle resources, ie. resources that are in drain mode to schedule a large multi-node job. They then construct an ATLAS job that fits exactly into the hole, submit it and more often than not are picked to run at the next scheduler cycle (even without any allocation). Talk shows the progress over the last year. They are up to about 15M produced event per month !
- <https://indico.cern.ch/event/505613/contributions/2227926/>