MAGIC Meeting Minutes
January 4, 2017

Attendees

Ken Bloom	FermiLab
Peter Couvalis	LIGO
Burt Holzman	FermiLab
Shantenu Jha	Rutgers
Dan Katz	U. Ill.
Kate Keahey	ANL
Miron Livny	U. Wisc
Grant Miller	NCO
Claire Mizumoto	UCSD
Mike Nelson	CloudFlare
Valerie Polashar	UCSD
Don Riley	U. Md
Alan Sill	Texas Tech
Derek Simmell	PSC
Mike Sokoloff	UC/McMicken
Jack Wells	Oak Ridge

Action Items

Proceedings
This MAGIC meeting was coordinated by Grant Miller of the NCO. This meeting focused on cloud environments for users using access to commercial providers.

FermiLab HEPCloud: Burt Holzman

There are rapidly changing computer architectures and increasing data volumes in HEP that require effective cross-cutting solutions. Computing needs in HEP will soon be 10-100 x current capacity. Industry resources currently equal or exceed research resources in computation and commercial costs are decreasing significantly. Computation requirements fluctuate significantly over time depending on instrument and modeling schedules.

To address these needs there are resources available in Grids (VOs of users in trust federations), Cloud resources (community clouds in trust federations), and HPC resources (Researchers granted access to HPC installations. In this environment HEPCloud is a portal to the ecosystem of diverse computing resources (commercial and academic). HEPCloud provides complete solutions to users with agreed-upon levels of service. The facility routes to local or remote resources based on workflow requirements, cost, and efficiency of accessing the resources. HEPCloud also manages allocations of users to target compute engines. HEPCloud's goal is to move into production during FY2018. Seed money is provided by industry. HEPCloud collaborators include: worldwide LHC Computing Grid, BNL and ATLAS, ANL, HT Condor, CMS/IF experiments, and CERN.

c/o National Coordination Office for Networking and Information Technology Research and Development
Suite II-405 · 4201 Wilson Boulevard · Arlington, Virginia 22230
Phone: (703) 292-4873 · Fax: (703) 292-9097 · Email: nco@nitrd.gov · Web site: www.nitrd.gov

The HEPCloud architecture includes the submitted workflow, facility interface, Authentication and Authorization, Decision engine, facility pool, Provisioner, resources, and monitoring. Early HEPCloud use cases include NOVA Processing (processing the 2014/2015 dataset and CMS Monte Carlo Simulation of simulated events. The NOVA processing ran on 7300 cores and used general-purpose data handling tools. The Compact Muon Solenoid (CMS) experimental support simulated many billions of events using up to 60,000 cores on AWS. At SC16 they aimed to generate 1 billion events in 48 hours, doubling the size of global CMS computing resources.

Costs of commercial cloud resources include:

- Compute charges over time (per hour)
- Data input/output charges
- Persistent storage for large input data sets
- Ancillary support services
- Per-operation API charges

AWS excess capacity is released to the free (spot) market at a fraction of the on-demand price. Preemptible Google VMs are available at a significantly smaller fixed cost- 1 cent per core hour for a standard candle. Costs for computation were:

- On-premises resources: 0.9 cents per core hour
- Off-premises AWS: 1.4 cents per core hour (ranging up to 3 cents per core hour at smaller scale)
- Premptible Google costs are being estimated but are estimated to be : 1.6 cents per core-hour.

In the exascale era HPC facilities are potential compute resources for HEPCloud and HEPCloud facility services are potential resources for HPC facilities. Early use cases include:

- MicroBooNE production on Cori@NERSC
- Pythia on Mira at ALCF
- CMS production on Edison, Cori@NERSC

There are plansin 2017 to enable HEPCloud at NERSC.

For additional resources on HEPCloud please see: http://hepcloud.fnal.gov
For the full briefing please see the MAGIC WebSite at: https://www.nitrd.gov/nitrdgroups/index.php?title=MAGIC_Meetings_2017#January_2017

HTCondor and (Commercial) Clouds: Miron Livny

HTCondor provides services to users including:

- User submits jobs to the HTCondor SchedD and the SchedD delegates the job for execution on one of many resources including HTCondor StartD, Batch System, Grid Compute Element, or a Cloud.

Commercial clouds are a natural element of a distributed high throughput environment. A cloud is an autonomous computing (processing, storage, networking) resource with an interface that supports remote invocation of jobs and staging input/output data.

Condor manages resources of my cloud, provisions resources from a cloud, or manages my resources in a cloud.

- Jobs run as a VM on a server in the HTCondor pool
- Jobs run as a VM in a cloud
- Jobs run on a VM in an HTCondor pool that was dynamically deployed in a cloud

When user resources include money, HTCondor enables the translation of the money into seamless HTCondor capacity by expending an existing HTCondor pool or flocking with a dynamically deployed Condor pool. For the user, Condor:

- Maximizes capacity of the user resources via a single interface
- Minimizes overhead of accessing remote capacity
- Preserves local computation environments

Future Condor capabilities will include:

- HTCondor Annex: interface with advanced cloud APIs, control and manage spending, provide automated integrity
- - Cloud deployment and migration of HTCondor job managers
- Integration with Workflow Management Systems
- Provide for extreme scaling

For the full briefing, please see:
https://www.nitrd.gov/nitrdgroups/index.php?title=MAGIC_Meetings_2017#January_2017

Meetings of Interest
January 23: Washington DC: Congressional Internet Caucus

Next MAGIC Meeting
February 1, 12:00-2:00 Eastern, NSF, Room TBD