



Contribution ID: 200

Type: Talk

An implementation of cloud-based grid CE and SE for ATLAS and Belle II

Tuesday 22 October 2024 13:48 (18 minutes)

The economies of scale realised by institutional and commercial cloud providers make such resources increasingly attractive for grid computing. We describe an implementation of this approach which has been deployed for

Australia's ATLAS and Belle II grid sites.

The sites are built entirely with Virtual Machines (VM) orchestrated by an OpenStack [1] instance. The Storage Element (SE) utilises an xrootd-s3 gateway [2][3] with back-end storage provided through an S3-compatible object store from a commercial provider. The provisioning arrangements required the deployment of some site-specific helper modules to ensure all SE interfacing requirements could be met. OpenStack hosts the xrootd redirector and proxy servers in separate VMs.

The Compute Element (CE) comprises virtual machines (VM) within the Openstack instance. Jobs are submitted and managed by HTCondor [4]. A CloudScheduler [5][6] instance is used to coordinate the number of active OpenStack VMs and ensure that VMs run only when there are jobs to run.

Automated configuration of the individual VMs associated with the grid sites is managed using Ansible [7]. This approach was chosen due to its low overheads and the simplicity of deployment.

Performance metrics of the resulting grid sites will be presented to illustrate the viability of this cost-effective approach to resource provisioning for grid computing.

References

- [1] OpenStack: <https://www.openstack.org/>
- [2] Xrootd: <https://xrootd.slac.stanford.edu/>
- [3] Andrew Hanushevsky and Wei Yang: "Xrootd S3 Gateway for WLCG Storage", 26th International Conference on Computing in High Energy & Nuclear Physics (CHEP 2023), <https://doi.org/10.1051/epjconf/202429501057>
- [4] HTCondor: <https://htcondor.org/htcondor/overview/>
- [5] CloudScheduler: <https://github.com/hep-gc/cloudscheduler>
- [6] Randall Sobie, F. Berghaus, K. Casteels, C. Driemel, M. Ebert, F. F. Galindo, C. Leavett-Brown, D. MacDonell, M. Paterson, R. Seuster, S. Tolkamp, J. Weldon: "cloudScheduler a VM provisioning system for a distributed compute cloud", 24th International Conference on Computing in High-Energy and Nuclear Physics (CHEP 2019), <https://doi.org/10.1051/epjconf/202024507031>
- [7] Ansible: <https://www.ansible.com/>

Primary author: Dr WOITHE, Jonathan Mark (University of Adelaide (AU))

Co-authors: DOSSETT, David; Dr EBERT, Marcus (University of Victoria); SEVIOR, Martin (University of Melbourne (AU)); JACKSON, Paul (University of Adelaide)

Presenter: Dr WOITHE, Jonathan Mark (University of Adelaide (AU))

Session Classification: Parallel (Track 7)

Track Classification: Track 7 - Computing Infrastructure