



Contribution ID: 131

Type: oral presentation

HEP cloud production using the CloudScheduler/HTCondor Architecture

Tuesday, April 14, 2015 4:30 PM (15 minutes)

The use of distributed IaaS clouds with the CloudScheduler/HTCondor architecture has been in production for HEP and astronomy applications for a number of years. The design has proven to be robust and reliable for batch production using HEP clouds, academic non-HEP (opportunistic) clouds and commercial clouds. Further, the system is seamlessly integrated into the existing WLCG infrastructures for the both ATLAS and BelleII experiments. Peak workloads continue to increase and we have utilized over 3000 cores for HEP applications.

We show that the CloudScheduler/HTCondor architecture has adapted well to the evolving cloud ecosystem. Its design preceded the introduction of OpenStack clouds, however, the integration of these clouds has been straightforward. Key developments over the past year include the use of CVMFS for application software distribution. CVMFS together with the global array of Squid web caches and the use of the Squid-discovery service, Shoal, have significantly simplified the management and distribution of virtual machine images. The recent deployment of micro-CernVM images have reduced the size of the images to a few megabytes and make the images independent of the application software. The introduction of Glint, a distributed VM image management service, has made the distribution of images over multiple OpenStack clouds an automated operation. These new services have greatly simplified management of the application software and operating system, and enable us to respond quickly to urgent security issues such as the ShellShock (bash shell) vulnerability.

HEP experiments are beginning to use multi-core applications and this has resulted in a demand for simultaneously running single-core, multi-core or high-memory jobs. We exploit HTCondor's feature that dynamically partitions the slots within a VM instance. As a result, a new VM instance can run either single or multi-core jobs depending on the demand.

The distributed cloud system has been used primarily for low I/O production jobs, however, the goal is run higher I/O production and analysis application jobs. This will require the use of a data federation and we are integrating the CERN UGR data federator for managing data distributed over multiple storage elements.

The use of distributed clouds with the CloudScheduler/HTCondor architecture continues to improve and expand in use and functionality. We review the overall progress, and highlight the new features and future challenges.

Author: Dr SOBIE, Randy (University of Victoria (CA))

Co-authors: LEAVETT-BROWN, Colin Roy (University of Victoria (CA)); BERGHAUS, Frank Olaf (University of Victoria (CA)); GABLE, Ian (University of Victoria (CA)); PATERSON, Michael (U); Dr DESMARAIS, Ron (University of Victoria); TAYLOR, Ryan (University of Victoria (CA))

Presenter: GABLE, Ian (University of Victoria (CA))

Session Classification: Track 7 Session

Track Classification: Track7: Clouds and virtualization