



Contribution ID: 449

Type: poster presentation

Integrating grid and cloud resources at the RAL Tier-1

Today the primary method by which the LHC and other experiments run computing work at WLCG sites is grid job submission. Jobs are submitted to computing element middleware which in turn submits jobs to a batch system managing the local compute resources. With the increasing interest and usage of cloud technology, a new challenge facing sites which support multiple experiments in recent years is a need to provide both traditional grid as well as cloud interfaces, but without partitioning the underlying resources. When the batch system is busy but the cloud is idle, it should be possible for the unused cloud resources to be included in the batch system. Similarly, when the batch system is idle but the cloud is busy, the unused batch resources should be available for users within the cloud. At the RAL Tier-1 a cloud based on OpenNebula has been under development for some time and will be made available to the LHC experiments and others, as well as being used internally by staff for activities such as testing and development. Here we present our experience unifying the cloud with our production HTCondor batch system in a way that avoids static partitioning, ensures that resources are used efficiently and that allocations are respected.

Primary author: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Co-author: COLLIER, Ian Peter (STFC - Rutherford Appleton Lab. (GB))

Presenter: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Track Classification: Track7: Clouds and virtualization