



Contribution ID: 254

Type: oral presentation

Unified Storage Systems for Distributed Tier-2 Centres

Thursday, September 6, 2007 4:30 PM (20 minutes)

The start of data taking this year at the Large Hadron Collider will herald a new era in data volumes and distributed processing in particle physics. Data volumes of 100s of Terabytes will be shipped to Tier-2 centres for analysis by the LHC experiments using the Worldwide LHC Computing Grid (WLCG).

In many countries Tier-2 centres are distributed between a number of institutes, e.g., the geographically spread Tier-2s of GridPP in the UK. This presents a number of challenges for experiments in terms of the use of such centres, as CPU and storage resources may be sub-divided and exposed in smaller units than the experiment would ideally want to work with. In addition, unhelpful mismatches between storage and CPU at the individual centres may be seen, which make efficient exploitation of a Tier-2's resources difficult.

One method of addressing this is to unify the storage across a distributed Tier-2, presenting the centres' aggregated storage as a single system. This greatly simplifies the data management for the VO, which then can access a greater amount of data across the Tier-2. However, such an approach will lead to scenarios where analysis jobs on one site's batch system must access data hosted on another site.

We investigate this situation using the Glasgow and Edinburgh clusters, which are part of the ScotGrid distributed Tier-2. In particular we look at how to mitigate the problems associated with "distant" data access and discuss the security implications of having LAN access protocols traverse the WAN between centres.

Primary authors: Mr ELWELL, Andrew (University of Glasgow); Dr STEWART, Graeme A (University of Glasgow); Dr COWAN, Greig A (University of Edinburgh)

Presenter: Dr COWAN, Greig A (University of Edinburgh)

Session Classification: Grid middleware and tools

Track Classification: Grid middleware and tools