



Contribution ID: 486

Type: oral presentation

A Data Grid for the Analysis of Data from the Belle Experiment

Thursday 30 September 2004 17:50 (20 minutes)

We have developed and deployed a data grid for the processing of data from the Belle experiment, and for the production of simulated Belle data. The Belle Analysis Data Grid brings together compute and storage resources across five separate partners in Australia, and the Computing Research Centre at the KEK laboratory in Tsukuba, Japan.

The data processing resources are general purpose, shared use, compute clusters at the Universities of Melbourne and Sydney, the Australian Partnership for Advanced Computing (APAC), the Victorian Partnership for Advanced Computing (VPAC) and the Australian Centre for Advanced Computing and Communications (AC3).

This system is in use for the Australian contribution to the production of simulated data for the Belle experiment, and for physics analyses.

The Storage Resource Broker (SRB), from the San Diego Supercomputing Centre, is used to provide a robust underlying data repository. A federation of SRB servers has been established to share and manage Belle data between the KEK laboratory, the mass data store at the Australian National University (ANU) and satellite storage at each of the compute clusters.

The globus toolkit is the underlying technology for the management of the computing resources, and the despatching of jobs. A network aware job scheduler has been developed. The scheduler queries the SRB servers for location of data replicas, and arranges scheduling of processing and production jobs on the compute resources according to a static model of the network connectivity and dynamic assessment of the relative system loads.

Authors: MOLONEY, G R.; SMILLIE, J.; WINTON, L J.; MCMAHON, S.

Presenter: MOLONEY, G R.

Session Classification: Distributed Computing Systems and Experiences

Track Classification: Track 5 - Distributed Computing Systems and Experiences