

LCG Tier2 Setup at LRZ

J. Kennedy, O. Biebel, G. Duckeck, J. Elmsheuser, G. Krobath, B. Ruckert, D. Schaile, and C. Serfon

1. Introduction

To ensure that the computing requirements for the upcoming ATLAS experiment at the LHC are met the ATLAS community has turned to distributed computing and the grid. The ATLAS computing model describes a multi-tier hierarchical structure with different centers supporting different services. The main center at CERN, the T0, is supported by regional T1 centers. The Tier1 centers in turn support a number of associated Tier2 centers and form a so called “Cloud”. The Tier-2’s are responsible for Monte-Carlo data production and providing a user analysis service.

2. A Munich Tier 2

For the Munich Tier-2 a joint effort combining the installations at LRZ and RZG is underway. As a baseline assumption an equal sharing of resources is foreseen between MPI/RZG on one hand, and the LMU/LRZ on the other. Together the Munich Tier-2 should provide the average ATLAS Tier-2 capacity of about 1/3 of the Tier-2 resources for Germany.

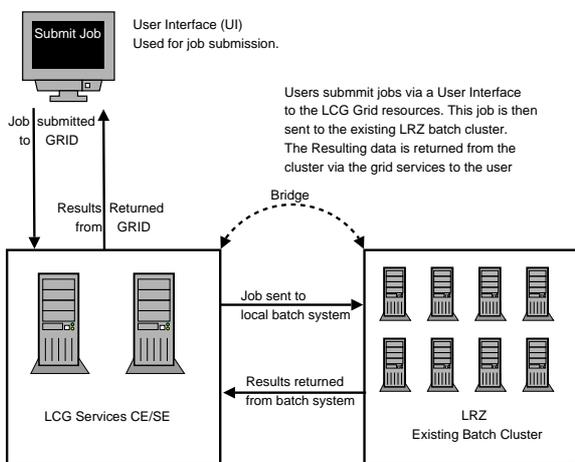


Fig. 1: Schema of Job submission to the LRZ-LCG system

3. gLite and dCache

The gLite grid middleware [1] is used by the LCG to enable a site to contribute both computing and storage resources to the grid.

The model for using the LRZ center as an LCG resource is non standard and requires some modifications of the LCG setup. At LRZ we deploy LCG by bridging into an existing batch cluster see fig 1.

Although bridging into an existing cluster requires this modification of the std LCG setup we benefit since existing resources can be more readily added to the LCG.

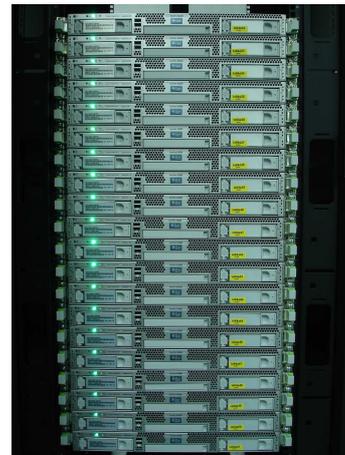


Fig. 2: Worker Nodes within the LRZ cluster

For the management of the disk storage at LRZ we use the dCache system [2]. dCache has been developed in the last years at Desy/Hamburg and FNAL/USA. It is designed to manage large data sets and storage systems, from few TeraBytes to several PetaBytes. It provides a unique namespace for the files and datasets stored and it offers a high flexibility to tune it to the specific requirements of a certain site. An existing dCache setup can be easily extended with new disk-storage, dCache provides access to the data via several interfaces via the local or wide area network and supports standard Grid access protocols (SRM, gridftp).

For our installation at LRZ important features are the division of the storage into several pools for different tasks and user-groups and the option to use the LRZ tape system as a transparent archive for datasets. The detailed setup and configuration is complicated by the LRZ firewall setup, private IP sub-nets for cluster nodes and conflicting requirements for local and long-distance access.

4. Status

The LCG center at LRZ was deployed during the last quarter of 2006 and was certified as a production quality site in early 2007. The site currently provides 148 cpu cores giving a total of 190KSi2K (see fig 2) and 40TB of disk storage.

The ATLAS software has been installed and the site has been integrated into the ATLAS Distributed Data Management System (DDM). Preliminary tests have begun with both MC production/distributed analysis jobs and data storage via DDM.

5. Outlook

The immediate aims are to aid with the production of ATLAS MC data and to store the data thus enabling physics users to perform analysis. In addition we hope to continue to work on improving the integration of LCG into our site to ensure it functions in an optimal manner. The deployment of service and node monitoring software at the LRZ is also a goal for the short term future.

References

- [1] <http://glite.web.cern.ch/glite/>
- [2] <http://www.dcahce.org>