

WLCG Tier-2 Deployment in Munich

J. Kennedy, O. Biebel, C. Cummer, G. Duckeck, J. Elmsheuser, G. Kroboth, B. Ruckert, T. Mahmoud, D. Schaile, M. Schott, and C. Serfon

1. Introduction

The ATLAS computing model describes a multi-tier hierarchical structure with different centers supporting different services. The main center at CERN, the Tier-0, is supported by regional Tier-1 centers. The Tier-1 centers in turn support a number of associated Tier-2 centers and form a so called “Cloud”.

2. A Munich Tier-2

A so called distributed Tier-2 center is deployed in Munich by combining installation at the LRZ and RZG computing centers. A close collaboration between these two centers and the physics groups at the LMU/MPI provides a good basis for a solid and flexible Tier-2 deployment. The Munich Tier-2 will provide approximately 1/3 of the Tier-2 resources within the German ATLAS community. Here we report on the development of the LRZ/LMU side of the Tier-2.

3. Site Monitoring

During 2007 one of the main targets was the deployment of a monitoring framework to aid with debugging and help provide an overview of the usage at the Tier-2 center. A project was undertaken to deploy the Lemon [1] monitoring tool at the LRZ/LMU center. Lemon is a powerful and flexible monitoring system developed within the EGEE/CERN community and is in popular usage across the grid. Information from several sources is collected via a central server and is made available via a web page interface. At LRZ/LMU the standard Lemon monitoring was deployed and a site specific development was undertaken to provide dedicated additional information.

Figure 1 shows a plot of the number of active and queued grid jobs at the LRZ/LMU center.




Fig. 1: Monitoring of jobs within the LRZ/LMU Tier-2 center

An informative and flexible monitoring tool is essential for a site to allow monitoring of usage and the identification of problems.

4. Data Storage and dCache

The dCache system at LRZ consists of a powerful administration server and a set of storage pool nodes. The admin node hosts the namespace databases and services for various access protocols. The pool nodes are powerful storage servers consisting of Raid-6 systems with typically 16 hard-

disk drives. In early 2007 ten such storage pool nodes were in operation, each with a net capacity of 4 TB and a 1 Gb-connection to the LAN. For our applications in ATLAS the important feature of the dCache system is not only good overall performance but in particular a good scaling behaviour when many clients access the storage system simultaneously. Figure 2 shows the result of a dedicated throughput test when up to 40 clients on up to 10 remote worker nodes randomly accessed files on dCache. The overall performance was very satisfactory with more than 500 MB/s read rate. End of 2007 the dCache system had been further upgraded with additional 28 storage pool nodes of 9.6 TB each with 10 GB LAN connection.




Fig. 2: dCache throughput tests at LRZ

5. Status

During 2007 the LRZ/LMU center has undergone an expansion in both CPU and storage and has also been developed to provide a reliable level of service.

From the view of ATLAS computing the center has been fully certified.

- We ensure that we provide the latest atlas software installations
- Several tests of the site within the ATLAS Distributed Data Management System (DDM) have been successfully undertaken
- MC data generation is ongoing with increasing rates and efficiencies
- Data replication for user analysis is underway
- User analysis is ongoing with increasing frequency

6. Outlook

2008 is a mixed year for EGEE sites, as the LHC turns on and data is generated the sites need to ensure a solid and reliable operations level service, however future development and middleware updates need to occur. Major goals for this year involve an upgrade of our grid middleware, increases in both cpu and storage and resulting scaling tests and the deployment of services within a virtual environment for both central services and testing purposes.

References

- [1] <http://lemon.web.cern.ch/lemon>