

# Munich WLCG Tier2 Status

G. Duckeck, P. Calfayan, J. Elmsheuser, J. de Graat, J. Ebke, C. Galea, J. Kennedy, C. Kummer, C. Mitterer, D. Schaile, C. Serfon, and R. Walker

## 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 regional “Cloud”.

## 2. Munich Tier-2

The Munich Tier-2 center is a federated Tier-2 center, which combines installations 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 operation. The Munich Tier-2 will provide approximately 1/3 of the Tier-2 resources within the German ATLAS community. In this report we focus on the LRZ/LMU side of the Tier-2.

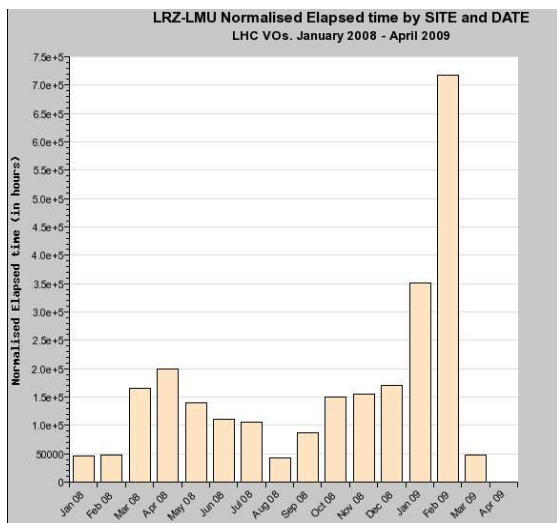


Fig. 1: CPU time per month for ATLAS Grid Jobs at LRZ.

## 3. Status

In order to provide sufficient CPU and storage resources for the start of LHC operation in summer 2008 and the expected data volume the capacity at LRZ/LMU center has been substantially expanded at the beginning of 2008 and a second upgrade took place in December 2008. The compute nodes provide about 1100 CPU cores corresponding to a capacity of about 1800,000 SpecInt2000 (standard CPU benchmark). The disk storage adds up to 500 TByte.

The use of the resources increased in parallel as illustrated in Fig. 1. A large fraction of the CPU resources are used for the central Monte Carlo production in ATLAS, but the resource usage by simulation and analysis jobs of individual users has steadily increased in the last year.

The storage systems are integrated in the ATLAS data-management operations. A continuous flow of datasets pro-

duced in ATLAS simulation jobs worldwide reaches the storage systems. Datasets produced at the LRZ/LMU is also stored and sent to the other ATLAS centers worldwide.

## 4. Analysis Site Tests

A new aspect of the Tier-2 operation are dedicated site tests for user analysis jobs. Until recently the use of the resources was focussed on Monte Carlo production, which is a rather CPU intensive task with relatively little I/O usage. In contrast, analysis jobs process large amounts of data, typically a 10 h job reads several 100 GB of data. Several 100 of such jobs running at the same time produce an extreme load on the storage systems and the networks, which can easily lead to I/O blocking and very inefficient CPU usage. It requires non-trivial tuning and optimizations of access protocols, data distribution, storage configuration, etc., to achieve satisfying performance. Figure 2 shows such a test with 200 jobs starting in parallel. The analysis of this test indicates that the bottleneck is caused by many concurrent requests on the filesystems which reduces the effective I/O rate far below the nominal rate for sequential I/O.

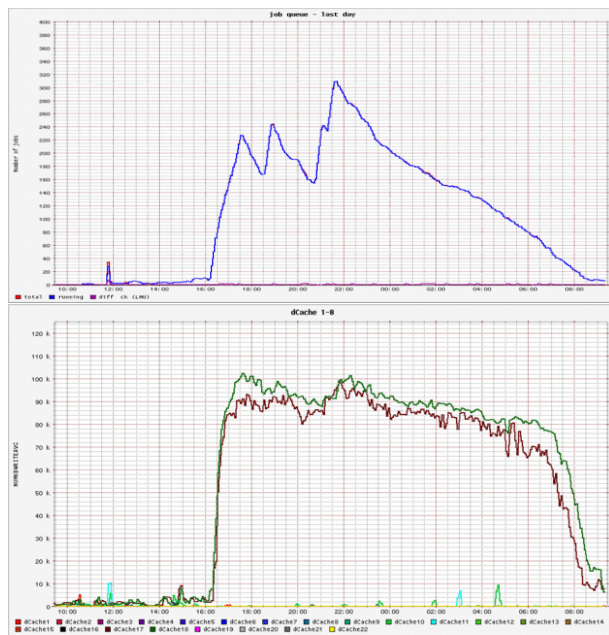


Fig. 2: Analysis test: Number of jobs and I/O rate vs time.

## 5. Outlook

Further optimisations of the storage system setup are planned to improve the performance of analysis jobs in time before the expected continuation of LHC operation in the fall 2009. Further increases in both storage and CPU are foreseen after the LHC turn on.