

DDM Operations for the ATLAS Experiment at the LHC

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

According to the LHC computing model it is planned to distribute the data and the workload to many regional computing centers in the participating countries and to connect them with high speed networks. A sophisticated Distributed Data Management system (DDM) [1], which is based on Grid software packages, shall automatically manage the data and provide a transparent and unique view of the connected resources.

DDM uses the concept of a dataset, which is essentially a set of logically connected files, as the basic block of data that will be used for transactions. These datasets are located on certain sites and can be copied to others via a subscription. DDM relies on the one hand on a set of central catalogs based at CERN, where informations about datasets are summarised (dataset contents, replicas, subscriptions...) as well as other central services.

On the other hand it is complemented by a catalog (LGC File Catalog or LFC) which is stored on regional computing facilities called Tier-1 (T1). The LFC contains the list of files present on the Storage Element (SE) of this site as well as those present in some smaller computing facilities called Tier2s¹ (T2). The Tier1s also operate additional services required by DDM.

Figure 1 depicts the DDM and its components.

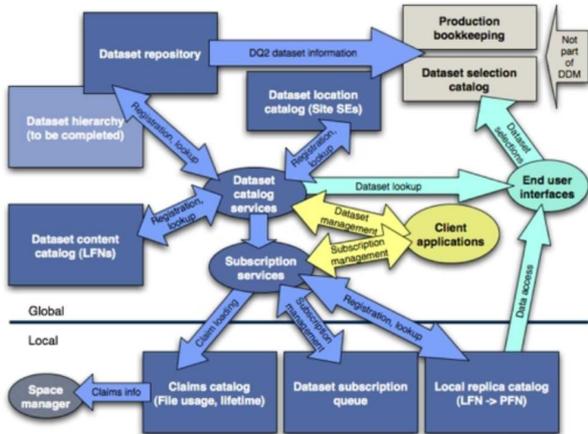


Fig. 1: Schematic view of ATLAS DDM.

During LHC running, a constant data flow will go from the CERN to each Tier1, depending on the Tier1's resources. In case of GridKa, the expected flow is about 88 MB/s. Moreover, transfers will occur between Tier1 and the associated Tier2s (copy of Monte Carlo events from T2 to T1 and copy of physics data from T1 to T2). These transfers involve less data, but are nevertheless crucial in

the computing model developed by ATLAS. To be ready at LHC start-up, several tests have been performed. At first, the basic functionality of the transfers had been checked. Then longer-term performance tests are needed, to ensure that the sustained transfer rates meet the requirements. Figure 2 summarises one of these tests between GridKa and some Tier2s.

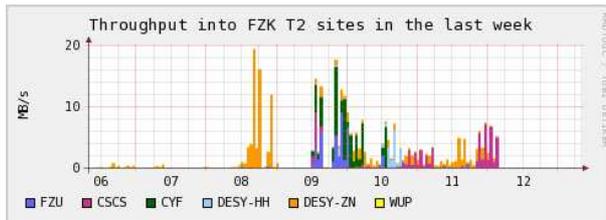


Fig. 2: Monitoring of the transfer of a few GBs of data from GridKa to some of the associated Tier2s.

In addition to these tests, a replication of all csc11 AODs² to GridKa is done in order to allow physicists to perform their analysis on this facility. More than 2 TB of data was accumulated during this replication. During this exercise, various problems have been highlighted and fixed. One of these problems is the consistency between the data stored on a given site and what is stored on the LFC. Transfers between the Tiers rely on various services and sometimes discrepancies occur between the content of the LFC and what is actually stored on disk. A tool to check these problems has been written and a cleaning procedure developed. Table 1 details the list of problems for GridKa and a few of its associated Tier2s. All these problems have been fixed in the meantime.

Site	In LFC not in SE	In SE not in LFC	Different size	0 length in LFC	0 length on SE
GridKa	12 766	31	56	627	18
CYF	1	n.a.	0	0	2
DESY-ZN	599	n.a.	0	0	0
WUP	12	n.a.	0	0	0

Table 1: List of problems seen by the integrity check script on GridKa and some associated Tier2s. In the case of the Tier2s a piece of information is not available.

To ensure a good functionality of DDM, a daily monitoring of all services is crucial and rapid reactions (contact of responsible on the site, or experts of DDM system) are needed to have a system running without too much interruptions. This daily work is performed by a DDM operation team in association with experts in GridKa.

References

[1] ATLAS Distributed Data management Operations, D. Barberis *et al.*, ATL-SOFT-PUB (2006) 006

¹Each Tier1 is associated with several Tier2s in a structure called cloud. The German cloud, with GridKa as Tier1, has 8 Tier2s.
²Analysis Object Data corresponding to the ATLAS csc data challenge.