



Contribution ID: 47

Type: **not specified**

DDM Site Services: A Solution for Global Replication of HEP Data

Tuesday, 27 March 2012 11:00 (20 minutes)

Description of the Work

ATLAS Distributed Data Management is the system that manages the experiment's detector, simulated and user data while enforcing the policies defined in the Computing Model. It provides functionality for data placement, deletion, bookkeeping and access on a hierarchic grid model composed of around 100 sites with heterogeneous storage technologies.

The DDM Site Services are the agents responsible for the data discovery and replication and are optimized for the usage of common SRM, FTS and LFC middleware. They are being used to achieve aggregated throughput rates far beyond the initial requirement of 2GB/s, having reached throughput peaks of over 10GB/s. To ensure further scalability, the core of the DDM Site Services has been designed as a set of independent agents, which work around an internal, independent database to store their state. In this way indefinite instances of DDM Site Services can work in parallel as long as the central bookkeeping system is able to sustain the load.

DDM Site Services is in production since the beginning of the LHC data taking period and, based on the production experience, have been adapted to the evolution of the network infrastructure, new source selection policies, changes in the needs of particular storage elements or to include the requirements of emerging "off-grid" Tier3 sites.

In the topic of monitoring we have worked in improving the service monitoring by sending the health reports through the Messaging System for Grids (MSG) and publishing them from a server to CERN IT's Service Level Status (SLS) infrastructure. We will also present the second version of the DDM Dashboard, which provides a very powerful and highly customizable user interface for the data transfer activity monitoring.

This presentation will give a detailed insight into the architecture of the above-mentioned systems and show how they have adapted to the needs of the data taking era.

Conclusions

DDM Site Services is the service that has been in charge of the ATLAS data discovery and replication. This critical service has been adapted and improved following the evolution and needs of the different grid elements and is successfully coping with the needs of the LHC data-taking era. An insight into the system, its deployment model and the available monitoring can help other VOs to get new ideas and compare it to how they are using the grid.

Impact

DDM Site Services can be considered one of the heaviest clients that throttle data transfers on the grid. It is used to export beam data from CERN immediately to avoid its loss and to transfer real and simulated data across the grid to facilitate its analysis and reprocessing. At the large scales of HEP data processing, it is important to provide a reliable system that is resilient to failures in the distributed environment so that the

operations workload is minimized; the goal of the system is a transfer error rate of less than one fault per million transfers. The implementation is based on common middleware like SRM, FTS and LFC and therefore can be interesting for any grid enabled community.

Overview (For the conference guide)

The ATLAS Distributed Data Management (DDM) is the project built on top of middleware currently maintained by EMI and is responsible for the organization of the multi-Petabyte ATLAS data across more than 100 distributed grid sites. One particular component of the system - the DDM Site Services - is the set of agents responsible for the discovery and placement of ATLAS data between sites. DDM Site Services manage aggregated throughputs of over 6GB/s or one million file transfers a day and have to work with extremely high reliability and availability. This contribution will build upon the production experience acquired during the last 2 years of LHC data taking and show the changes, adaptations and improvements that we implemented on the system to guarantee a flawless service. In a second part we will give an update on the service and activity monitoring frameworks that publish the information needed by shifters and experts.

Primary authors: Dr DI GIROLAMO, Alessandro (CERN); Mr TUCKETT, David (CERN); Mr HARALD BARREIRO MEGINO, Fernando (CERN)

Presenter: Mr HARALD BARREIRO MEGINO, Fernando (CERN)

Session Classification: Services for Data Management and Messaging

Track Classification: Software services for users and communities