



Contribution ID: 153

Type: Poster

Configuration and optimization for a gLite grid computing farm.

Monday, April 12, 2010 6:48 PM (3 minutes)

We will illustrate the optimization work performed on a mid-sized gLite grid-computing farm, testing various solutions, to improve the services offered to both grid and local users. Particular attention was devoted to the storage infrastructure and its configuration trying to improve the overall farm performance, reliability and to minimize the manpower required for its management. Improvements were also pursued by optimizing the configuration of the batch system. Experiences from both LHC users and non LHC ones will be reported.

Detailed analysis

Details on the tests performed and the results achieved while attempting to improve the performance of a typical tier2 farm while serving both LHC and non LHC VO's. Important improvements on the storage performance have been achieved through a patient optimization of several layers starting from the hardware to the storage manager software. In particular the performance and the functionalities obtained using two different storage manager solutions, dCache and StoRM/Lustre, will be presented.

Requirements coming from the different users groups' force also an optimization of batch system, Torque/MAUI, in order to allow the final user to exploit the farm resources with grid, local, interactive and MPI jobs. New users also require a clear framework in order to simplify the porting of new applications to grid. Examples of applications ported (coming from different community: theoretical physics, astro-physics, bioinformatics, etc) and the performance achieved will be provided.

We will report also about the performance and the scaling of MPICH based applications, and on experience about using Worker Nodes of the batch farm in order to support interactive sessions of the end users.

Conclusions and Future Work

Through this infrastructure the users could start achieving know how on how to submit their application to a batch farm. This is often the first step needed to help users to migrate their application to a distributed grid infrastructure such as EGI production grid, as soon as the workload for a given application increases and could not be executed within a single batch farm.

Impact

We will show how this configuration could improve the overall usage of the available CPUs in a typical batch farm.

In this work we will give also some short information about the scientific results obtained.

The work spent into optimizing the farm configuration and performance will improve the user satisfaction and increase the number of researchers that could solve their problems by means of using such computational infrastructure and will allow them to solve problem that could not be solved with typical desktop computers.

Keywords

gLite, batch farm, LHC, MPI, storage, Lustre, dCache

URL for further information

<http://webcms.ba.infn.it/cms-software/index.html/index.php/Main/GridBariFarm>

Primary authors: DONVITO, Giacinto (INFN-Bari); Prof. MAGGI, Giorgio Pietro (INFN and Politecnico di Bari); SPINOSO, Vincenzo (INFN-Bari)

Presenter: DONVITO, Giacinto (INFN-Bari)

Session Classification: Poster session

Track Classification: Experiences from application porting and deployment