



Contribution ID: 26

Type: Oral

Deployment and Management of Grid Services for a data intensive complex infrastructure

Thursday 15 April 2010 09:20 (20 minutes)

Services needed to access scientific data and computational resources in a distributed computing infrastructure are getting more and more complex. Our Atlas Tier2/Tier3 users, and National e-Science and GRID-CSIC initiatives researchers, have different requirements but all need ubiquitous and efficient access to these resources. We will provide details on how we use WLCG and Glite middleware, Storm, Lustre filesystem and our developments within this software to meet our specific necessities, to allow users direct access to the data, and to reach high availability in the provided services.

Detailed analysis

We are currently providing 590 TB of online data and potential capacity of 560 TB in offline data tapes, available to several communities which users in some cases intersect. For example we have Atlas Tier2 users that are also members of our Tier3, and we need to apply group and user quotas for the different spaces. National Grid Initiative and Grid-CSIC communities have also common members and applications, in some cases parallel computational jobs that also need direct access to data. Authorization for these data is done at the filesystem level with the usage of ACLs, and to be respected by higher levels of middleware we had to develop custom plugins for Storm for example. Lustre is mounted in the WNs to directly access with posix file calls, and local users have direct read-only access from the UI. We also provide X.509 secured web access for commodity. The data access patterns for newer applications outside the HEP scope is being analyzed, as well as optimal parameters like the read ahead for different types of workloads. As we are reaching higher levels of served data, we are scaling the number of Lustre OSTs, and distributing grid middleware services among several machines.

Conclusions and Future Work

The presented approach has proven to work fine and meets the required quality levels for the operation of the different supported virtual organizations, but in order to constantly improve, we are analyzing the usage of Lustre capabilities to access the geographically distributed data stored at the 3 sites of our distributed Tier-2. We are also planning to support 3rd level data tape storage, using the future developments of the Lustre/CEA with HSM when mature. We also intend also to continue supporting our developments and making it available to the community.

Impact

Recent tests as the STEP09 and UAT tests successfully exercised our infrastructure for the Atlas Tier2 in conjunction with other LHC experiments. The usage of Lustre as the backend cluster filesystem has proven to reach excellent performance levels, in addition to the configuration that allow jobs at the WN to directly access data in the servers. These tests also showed some problems mainly related with underestimation of the amount of stored data, which are being easily solved adding more storage servers thanks to the scalability of the services. The inclusion of several newer e-science Vos in the context on NGI and GRID-CSIC projects

have reinforced the need for group and user quotas, and it has been adopted the new pool paradigm of Lustre 1.8 to group several OST, proving to work fine. Development of newer software plugins for Storm to allow local authorization policy based on ACLs for the different spaces, was also adopted in the main source code, making it available for other Grid Sites to use it and watched as a proof of concept. Also Grid services are improving, and storm team is aware of the need to support group and user quotas at the srm level.

Keywords

Grid, Lustre, Data Access, SRM, Middleware development

URL for further information

<http://ific.uv.es/grid/e-science/activeprojects.htm>

Author: FERNANDEZ CASANI, Alvaro (Instituto de Fisica Corpuscular (IFIC) UV-CSIC)

Co-authors: LAMAS DAVIÑA, Alejandro (Instituto de Fisica Corpuscular (IFIC) UV-CSIC); AMOROS VICENTE, Gabriel (Instituto de Fisica Corpuscular (IFIC) UV-CSIC); NADAL DURA, Javier (Instituto de Fisica Corpuscular (IFIC) UV-CSIC); SANCHEZ MARTINEZ, Javier (Instituto de Fisica Corpuscular (IFIC) UV-CSIC); SALT CAIROLS, Jose (Instituto de Fisica Corpuscular (IFIC) UV-CSIC); MUÑOZ MENDEZ, Victor (Instituto de Fisica Corpuscular (IFIC) UV-CSIC)

Presenter: FERNANDEZ CASANI, Alvaro (Instituto de Fisica Corpuscular (IFIC) UV-CSIC)

Session Classification: Data Management

Track Classification: Software services exploiting and/or extending grid middleware (gLite, ARC, UNICORE etc)