

Monitoring, accounting and automated decision support for the ALICE experiment based on the MonALISA framework

Friday 11 May 2007 14:40 (20 minutes)

Describe the scientific/technical community and the scientific/technical activity using (planning to use) the EGEE infrastructure. A high-level description is needed (neither a detailed specialist report nor a list of references).

We are developing a general purpose monitoring system for the ALICE experiment, based on the MonALISA framework. MonALISA (Monitoring Agents using a Large Integrated Services Architecture) is a fully distributed system with no single point of failure that is able to collect, store monitoring information and present it as significant perspectives and synthetic views on the status and the trends of the entire system. Furthermore, agents can use it for taking automated operational decisions.

Report on the experience (or the proposed activity). It would be very important to mention key services which are essential for the success of your activity on the EGEE infrastructure.

The system monitors all the components: computer clusters (all major parameters of each computing node), jobs status and consumed resources (CPU, both in time and SpecInt2k units, memory, disk usage), jobs network traffic while reading/writing files with xrootd, services availability with details in case of failures (both AliEn and LCG services, proxies lifetime), storage monitoring with detailed information on number of files, available space, or staging and migrating operations, FTD/FTS transfers. The system is reliable and functional for more than two years, representing the main view towards the ALICE Grid.

With a forward look to future evolution, discuss the issues you have encountered (or that you expect) in using the EGEE infrastructure. Wherever possible, point out the experience limitations (both in terms of existing services or missing functionality)

Our focus is now on using the monitoring information for the development of higher level services that can take more intelligent operational decisions.

Describe the added value of the Grid for the scientific/technical activity you (plan to) do on the Grid. This should include the scale of the activity and of the potential user community and the relevance for other scientific or business applications

Monitoring information is gathered locally from all the components running in each site. The entire flow of information is aggregated on site level by a MonALISA service and then collected and presented in various forms by a central MonALISA Repository. Based on this information, other services take operational decisions such as alerts, triggers, service restarts and automatic production job or transfer submissions.

Authors: Mr CIRSTOIU, Catalin (CERN); Mr GRIGORAS, Costin (CERN); Dr BETEV, Latchezar (CERN)

Co-authors: Mr MURARU, Adrian (CERN); Dr PETERS, Andreas Joachim (CERN); Dr LEGRAND, Iosif (Caltech); Mr SAIZ, Pablo (CERN); Mr VOICU, Ramiro (Caltech)

Presenter: Mr CIRSTOIU, Catalin (CERN)

Session Classification: Grid Monitoring and Accounting

Track Classification: Grid Monitoring and Accounting