Contribution ID: **5** Type: **oral presentation**

Supporting MPI applications on the EGEE Grid

Friday 11 May 2007 09:40 (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).

This work has been undertaken to improve support for parallel applications using MPI.
Such applications are common in many fields of computational science including earth
sciences, computational chemistry, astrophysics and climate modelling. Existing users
of high-performance clusters are accustomed to MPI support and in many cases it is a
pre-requisite for migrating their applications to the Grid.
Better MPI support would
greatly increase the potential user-base of the EGEE Grid.

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.

It has been technically possible to run MPI jobs on the Grid for some time. However,
support has been lacking in these areas: i) standards for advertising availability of
MPI libraries; ii) recipes for configuring sites for MPI; iii) clear user
instructions for locating MPI sites and submitting jobs; iv) middleware limitations
that assume (and force) inflexible methods for submitting MPI jobs that are not
acceptable to many sites.
Based on discussion within EGEE and with the int.eu.grid (I2G) project, we have
formulated simple solutions that should greatly ease the use of MPI code on the Grid.
The approach is for users to submit their jobs wrapped in a script that performs any
necessary setup (e.g. compilation), using site-defined environment variables and
I2G's MPI-start package to execute their binary using the desired version of MPI.
Script templates will be made available for users to customise.
We have also produced
guidelines for configuring sites to support MPI.

With a forward look to future evolution, discuss the issues you have encountered (or that you expect) in using the EGEE infrastructure. Wherever possi-

ble, point out the experience limitations (both in terms of existing services or missing functionality)

The solutions we propose are workable and can be rapidly implemented by sites. Even without changes to the core middleware, we will be able to provide users with a sensible methodology for submitting MPI jobs and site admins with a practical recipe for configuring their sites. There are issues that will need to be addressed in the future including methods for selecting custom interconnects and compilers and flexible support for parallel jobs in the WMS.

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

The first phase of the Grid has focussed on very loosely-coupled applications such as those common in high-energy physics. However, there is a significant community of computational scientists who need to run more tightly-coupled applications using standards such as MPI. Most of these scientists currently run their jobs directly on clusters. If they were to migrate to the Grid, they would be able to access multiple clusters with a single sign-on. This would be of particular benefit to existing multi-site collaborations. To date, the lack of good MPI support has been one of the factors preventing greater adoption of the Grid by such scientists. Improvements would open up a range of application areas.

Authors: Dr LOOMIS, Charles (LAL); Dr CHILDS, Stephen (Trinity College Dublin)

Co-authors: COSTANTINI, Alessandro (University of Perugia); Dr COGHLAN, Brian (Trinity College Dublin); HEYNMANN, Elisa (UAB); DIJKSTRA, Fokke (University of Groningen); BORGES, Goncalo (LIP); CAMPOS, Isabel (CSIC); STERZEL, Mariusz (Cyfronet); GERVASI, Osvaldo (University of Perugia); KELLER, Rainer (HLRS); STORK, Sven (HLRS)

Presenter: Dr CHILDS, Stephen (Trinity College Dublin)

Session Classification: Workflow

Track Classification: Workflow