Multi-Grids Interoperation: Job Submission, Auditing and Tracking

Version 0.21: 17/1/06

Editor: Steven Newhouse

Aims

To examine large scale inter-grid interoperation and interoperability – from both a management and a software perspective – with a view to supporting application driven demonstrations at SC06. In addition it is envisaged that it will provide practical experience on emerging GGF standards and identify gaps for future work. This group will focus on job submission, auditing and tracking.

Proposed Activity

From the responses to the initial email there appears to be an existing underlying fabric of GT4 GRAM and WS-GRAM services (distinguishing between the pre and post web service implementations) at many of the sites. Some exploration of the issues of co-existent GT2 GRAM and GT4 GRAM systems will be needed within EGEE and initially within OSG (although they are already looking at migrating within the next 6 months). Our first goal should be to verify that we can interoperate with these systems. These systems may not be the same systems that the large scale grid demonstrations take place on at SC06 but should be representative. For instance, we may manually configure the access control lists (e.g. gridmap files) for these systems while other groups (within this 'project') select a process for doing this in a more scalable manner.

Demonstrating interworking of identical GT4 based infrastructures should not be too difficult! Our ultimate goal in this area should be to show interoperability of a protocol between two (or more) implementations. The first stage of this is at the job description level – through the JSDL specification. There are several systems available in this area:

- GridSAM: Provides a plain web service that accepts compliant JSDL 1.0 which can interface into GRAM, directly to schedulers/execution environments such SGE or Condor.
- NAREGI: The jobs submission system accepts JSDL (with extensions) specified jobs that are submitted into the NAREGI super scheduler for eventual execution.
- UNICORE: Through the Unigrids project UNICORE is looking to interoperate with GT4 execution environments or directly into scheduling systems such as LSF. This will be achieved through a set of 'atomic web services' that will stage files in & out, use JSDL (converted to an Abstract Job Object – AJO) to specify jobs invoked through a BPEL specified workflow environment.
- Globus: The introduction of JSDL into WS-GRAM has been introduced as work item but has not yet been scheduled for a specific release.
- DEISA: Have a production HPC infrastructure across Europe (based around IBM systems) linked through Load-Leveller – jobs submitted at one site may be executed at another. Evaluation and deployment of GT4 to start in May. A web service layer DESHL (DEISA Services for Heterogeneous Layer) will accept JSDL specified jobs for submission into UNICORE middleware layer.
- VPAC have an implementation of a JSDL like system.
- CREAM (planned – date?)

Just converging on a job description language alone will not provide full interoperability, for that we need convergence of the web service interface. While the OGSA-BES specification has not yet entered public comment (planned for GGF 16) there are several implementations available:

- GridSAM: Provides an implementation of the current BES specification (January 2006) in release 1.1.0.
- NAREGI: Has a BES-like interface that submits jobs into their super scheduler.
- University of Virginia: Has a BES implementation using .NET
- UNICORE: An experimental BES interface may be provided within a timescale of a few months.

Of which GridSAM and the University of Virginia implementation have already demonstrated interoperability.

Auditing and job resource tracking can be undertaken through the Usage Record (UR) schema and the Resource Usage Service specification. The UR schema is just (or has completed) a public comment period. The RUS specification has completed its public comment period and is undergoing minor changes in response to these comments. There are implementations of RUS/UR (like) systems from:

- University of Manchester (not supported)
- GridSAM (coming Feb/March 2006)
- NAREGI
- CREAM (planned?)
- Open Science Grid (Internal use?)
- DEISA: Resource usage information exchanged between internal sites using a project specific schema based around CIM & GLUE.
- EGEE using UR within its accounting system.

Milestones

1. Decide on this plan (15/1/06)
2. Build list of target systems for testing with DN's (& CA certificates) for the testers. (30/1/06)
3. Verify low-lying interoperation (16/2/06)
4. Report back and discuss progress at the workshop co-located with GGF16 in Athens.
5. Further planned milestones
    a. JSDL interoperability tests
    b. BES interoperability tests
    c. UR exchanges

Contributors

Steven Newhouse (OMII) s.newhouse@omii.ac.uk
Paolo Malfetti (DEISA) p.malfetti@cineca.it.
Kazushige Saga (NAREGI) saga@grid.nii.ac.jp
Mark Green (OSG) mlgreen@ccr.buffalo.edu
Massimo Sgaravatto (EGEE) massimo.sgaravatto@pd.infn.it
Stuart Martin (TeraGrid) smartin@mcs.anl.gov

Li Zha (CNGrid) char@software.ict.ac.cn
Nate Mueller (Condor) nmueller@cs.wisc.edu
Erwin Laure (EGEE) Erwin.Laure@cern.ch
David Bannon (D.Bannon@vpac.org)

Telecon Notes (Some information has been placed directly in this document)

5/1/06
Present: Steven, Paulo, David, Stuart, Kazushige, Mark

SN provided the background to this activity which emerged out of meeting organised by Charlie Catlett at SC05. The report and slides from the meeting are to be circulated to this list. There was general agreement and support that this was a useful activity in understanding how to manage the challenges to supporting multiple application communities on existing resources.

There was a general discussion soliciting information relating to the currently available middleware solutions based around GGF specifications. Most of the groups on the call would be represented at the GGF 16 workshop related to this activity.

Actions:
All: Review revised document and provide comments/corrections.
All: Identify technical staff at each site that will be installing systems and will be configuring gridmap files.
All: Identify the staff who will need DNs (and CA certificates) to be provided access to the systems.
All: Identify the computing resources that will be running the JSDL accepting services or will be hosting a 'gateway' installation to the compute resource. NB: We are not after capability at this point just enough functionality to run 'HelloWorld'!