



Job Priorities update

Farida Naz
Andrea Sciabà

Grid Deployment Board
February 6, 2008

Status of testing on PPS

- The JP mechanism has been installed on a CERN PPS CE
 - lxb1937.cern.ch
 - gLite 3.1 LCG CE on SLC4
 - PBS, Torque Server on SLC3
 - three queues with VOViews attached
 - **short** with **6** VOViews
 - **VO:cms**, DENY:/cms/Role=lcgadmin, DENY:/cms/Role=production
 - **VOMS:/cms/Role=lcgadmin**
 - **VOMS:/cms/Role=production**
 - **VO:atlas**, DENY:/atlas/Role=lcgadmin, DENY:/atlas/Role=production
 - **VOMS:/atlas/Role=lcgadmin**
 - **VOMS:/atlas/Role=production**
 - **cms** with **1** VOView
 - **VO:cms**, DENY:/cms/Role=lcgadmin, DENY:/cms/Role=production
 - **atlas** with **1** VOView
 - **VO:atlas**, DENY:/atlas/Role=lcgadmin, DENY:/atlas/Role=production
- Job submission with a CMS certificate using a gLite 3.1 WMS



Tests performed

- Goals

- Verify that the jobs are submitted to the queues where the FQAN in the proxy is authorized
- Verify that the number of waiting and running jobs are correctly reported in the VOView information
- Verify that the WMS takes the information used to calculate the rank from the correct VOView

- Method

- Submit jobs with different VOMS FQANs and check

Results



- All checks *eventually* succeeded
- It took some time to make things work
 - Some misconfigurations
 - The experience will come useful to prepare the documentation

Configuration

- Preliminary documentation in a Wiki page
 - <https://twiki.cern.ch/twiki/bin/view/EGEE/JPTest>
- The CE was configured by an administrator (Farida) with no previous experience with job priorities stuff
 - Important to understand the problems a site administrator could run into
- Notes
 - Need to define `FQANVOVIEWS=yes` in YAIM
 - Need to define for each queue a variable `{queue}_GROUP_ENABLE` listing the VOViews (1-1 correspondence with FQANs for now) whose value is something like

```
{queue}_GROUP_ENABLE="atlas /VO=atlas/GROUP=/atlas/ROLE=production  
/VO=atlas/GROUP=/atlas/ROLE=lcgadmin cms /VO=cms/GROUP=/cms/ROLE=lcgadmin  
/VO=cms/GROUP=/cms/ROLE=production"
```

- Need to install two rpms with respect to the latest gLite release
 - `lcg-info-dynamic-scheduler-pbs-2.0.0-1`
 - `lcg-info-dynamic-scheduler-generic-2.1.0-1`
 - Without doing so, the number of running and waiting jobs published was wrong



Still to do

- Repeat the tests on a LSF CE
- Test with 2 or more VOs at the same time
 - Probably things will work, but...
- Check that the fair share assignment works
 - But this is really a batch system configuration thing
- Prepare the documentation for the release
 - Including clear examples for the batch system configuration



Conclusions

- The Job Priorities mechanism has been successfully tested on the PPS at CERN
 - The certification process can start
- The documentation will be extremely important
 - Unthinkable to spend 2-3 days per site to debug misconfigurations, if the probability of a misconfiguration is more than a few percent!