

Job Priorities: status

Simone Campana (CERN)

www.eu-egee.org
www.glite.org

- **Several layers involved**
 - VOMS
 - Batch system configuration
 - Information System
 - Workload Management System
- **User gets a VOMS proxy with some group/role**
 - i.e. /atlas/Role=production
- **The VOMS proxy gets delegated UI->WMS->CE**
- **At the site (CE), LCAS/LCMAPS translates the user FQAN into a local user depending on a configuration file**
 - the first rule in the LCMAPS configuration file that matches the FQAN is used to do the mapping
 - Mapping of VOMS credentials to scheduler credentials
- **The LRMS is configured to grant different priorities/shares to different users/groups**
 - Assignment of scheduler share to these credentials
- **VOViews provide a way for publicizing the relevant characteristics of a CE on a FQAN-based way**
 - # of batch slots, # of waiting jobs, etc... published independently for different FQAN
 - The WMS considers the VOViews matching the VOMS FQAN for the matchmaking

- (1) `dn:GlueCEUniqueID=ce01.ific.uv.es:2119/jobmanager-pbs-short,mds-vo-name=local,o=grid`
`GlueCEAccessControlBaseRule: VO:atlas`
`GlueCEAccessControlBaseRule: VO:cms`
`GlueCEAccessControlBaseRule: VOMS:/cms/Higgs`
`GlueCEAccessControlBaseRule: VOMS:/atlas/Role=production`
`GlueCEStateWaitingJobs: 1251`

- (2) `dn:GlueVOViewLocalID=atlas,GlueCEUniqueID=ce01.ific.uv.es:2119/jobmanager-pbs-short,mds-vo-name=local,o=grid`
`GlueCEAccessControlBaseRule: VO:atlas`
`GlueCEStateWaitingJobs: 0`

- (3) `dn:GlueVOViewLocalID=atlas_role_production,GlueCEUniqueID=ce01.ific.uv.es:2119/jobmanager-pbs-short,mds-vo-name=local,o=grid`
`GlueCEAccessControlBaseRule: VOMS:/atlas/Role=production`
`GlueCEStateWaitingJobs: 1000`

- **Result of the list-match (Rank = -other.GlueCEStateWaitingJobs;):**

	CEId	*Rank
(3)	- ce01.ific.uv.es:2119/jobmanager-pbs-short	-1000
(2)	- ce01.ific.uv.es:2119/jobmanager-pbs-short	0

- ***A VOMS with /atlas/Role=production matches both (2) and (3), LCMAPS will map the FQAN to the ACR VOMS:/atlas/Role=production (3) only***

- **The FQAN VOViews should disappear from the information system.**
 - The VO:atlas view will then show inclusive information for ATLAS jobs submitted with any role
 - Very Short Term means **no more** than 2 weeks
- **Status: DONE but...**
 - The situation prior to FQAN VOViews has been restored but left some dirty
 - The “dirty” is in the configuration of the information provider. Example: `ce02.marie.hellasgrid.gr:2119/jobmanager-pbs-atlas`
 - Publishes 4 running jobs in the inclusive view (ALL VOs)
 - Publishes 1 running job in the VO:atlas view
 - But this CE ONLY accepts the ATLAS VO!! (i.e. ALL Vos = ATLAS)
- **Being followed up by EGEE operations. Some progress but a bit slow (still 60 queues have problems, they were 200)**

- **Make the VOViews mutually exclusive**

- on a CE not more than one VOView can match a FQAN
- Introduce a DENY rule on the Information system:

```
GlueVOViewLocalID: atlas
GlueCEAccessControlBaseRule: DENY:/atlas/Role=production
GlueCEAccessControlBaseRule: VO:atlas
[...]
GlueVOViewLocalID: /atlas/Role=production
GlueCEAccessControlBaseRule: VOMS:/atlas/Role=production
```

- **Required simple changes in the WMS**
- **Required careful publication of the rule on in the GIP**

- **The *FULL* stack has been tested in the SA3 certification testbed.**
 - Log of the test has been distributed to the TCG
- **Results of the test are positive:**
 - The WMS understands the DENY tags.
 - The queue is matched only once
 - The information used for the matchmaking are the correct ones, based on the VOMS FQAN
- **What has not been tested**
 - That priorities at the batch system level are respected
 - What happens with more than 2 FQANs (different roles/groups/VOs).
 - Other batch systems rather than Torque+Maui
- **We decided to consider the test on the certification testbed successful and test the remaining points in the PPS**
 - Larger scale
 - Different Batch system

- **Next actions:**
 - The version of YAIM configuring the JP should remain a branch
 - Not the official version
 - PPS coordinator should deliver as soon as possible 3 PPS sites configured with the JP mechanism, installed using the branched YAIM version
 - Please report any difficulty
 - CERN should be delivered by the end of today (but only with PBS)
 - Remember we need LSF somewhere
 - No timescale for the other two volunteering sites (not identified yet).
- **Tests should be EXTENSIVELY carried on in PPS. In case of success, YAIM branch will be re-merged in the main branch and certified.**
 - And re-deployed in PPS, but this time testing should be straightforward
- **Then we have to deploy. Carefully.**