

Table of Contents

VOMS FQANs, fair shares gridmapfiles, YAIM.....	1
VOMS FQANs.....	1
How user mapping works.....	2
Format of gridmap files.....	2
Additional comments.....	3
How to configure users.....	4
Questions to clarify.....	4
About fair shares.....	4
Torque configuration.....	4
Publishing it in the infosys.....	5
How this information is generated.....	6
The lcg-info-dynamic-scheduler.....	6
Configuration file.....	7
Divers.....	8
Some info on YAIM.....	8
New structure.....	8

VOMS FQANs, fair shares gridmapfiles, YAIM

VOMS FQANs

- VOMS (Virtual Organisation Membership Service) proxies are straightforward and backward compatible extensions of simple grid-proxies.
- A simple `voms-proxy-init` will result a simple grid-proxy **without** VOMS extension

```
-bash-2.05b$ voms-proxy-init
Enter GRID pass phrase:
Your identity: /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
Creating proxy ..... Done
Your proxy is valid until Thu Jun 14 06:11:54 2007
-bash-2.05b$ voms-proxy-info -all
WARNING: Unable to verify signature! Server certificate possibly not installed.
Error: VOMS extension not found!
subject      : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
issuer       : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
identity     : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
type         : proxy
strength     : 512 bits
path         : /tmp/x509up_u9437
timeleft    : 11:59:56
```

- You have to explicitly ask for VOMS extensions !

```
-bash-2.05b$ voms-proxy-init -voms dteam
Enter GRID pass phrase:
Trying next server for dteam.
Creating temporary proxy ..... Done
Contacting voms.cern.ch:15004 [/DC=ch/DC=cern/OU=computers/CN=voms.cern.ch] "dteam" Done
Creating proxy ..... Done
Your proxy is valid until Thu Jun 14 06:13:52 2007
-bash-2.05b$ voms-proxy-info -all
subject      : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
issuer       : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
identity     : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
type         : proxy
strength     : 512 bits
path         : /tmp/x509up_u9437
timeleft    : 11:59:21
==== VO dteam extension information ====
VO          : dteam
subject      : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
issuer       : /DC=ch/DC=cern/OU=computers/CN=voms.cern.ch
attribute   : /dteam/Role=NULL/Capability=NULL
attribute   : /dteam/ce/Role=NULL/Capability=NULL
attribute   : /dteam/ce/HU/Role=NULL/Capability=NULL
attribute   : /dteam/ce/HU/BUDAPEST/Role=NUL/Capability=NULL
timeleft    : 11:59:21
```

- Extension comes in two form:

- ◆ The old **deprecated** form:

```
/VO=dteam/GROUP=/dteam/ce/HU/BUDAPEST/ROLE=NULL/Capability=NUL
```

- ◆ And the new form:

```
/dteam/ce/HU/BUDAPEST/Role=NULL/Capability=NULL
```

The VO and GROUP attribute merged together the first part of the new group string is always the VO.

- No qualitative difference between attributes, they are handled in the same way. Their usage is matter of convention and not an inner property of the VOMS architecture.
- It is possible to specify a requested attribute on the command line:

```
-bash-2.05b$ voms-proxy-init -voms dteam:/dteam/Role=production
Enter GRID pass phrase:
Trying next server for dteam.
Creating temporary proxy ..... Done
Contacting lxb1928.cern.ch:15002 [/DC=ch/DC=cern/OU=computers/CN=lxb1928.cern.ch] "dteam"
Creating proxy ..... .
Your proxy is valid until Thu Jun 14 06:24:22 2007
-bash-2.05b$ voms-proxy-info -all
subject      : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
issuer       : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
identity     : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
type         : proxy
strength     : 512 bits
path          : /tmp/x509up_u9437
timeleft     : 11:59:36
==== VO dteam extension information ====
VO           : dteam
subject      : /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=gdebrecz/CN=546241/CN=Gergely Debr
issuer       : /DC=ch/DC=cern/OU=computers/CN=lxb1928.cern.ch
attribute    : /dteam/Role=production/Capability=NULL
attribute    : /dteam/Role=NULL/Capability=NULL
attribute    : /dteam/bitface/Role=NULL/Capability=NULL
timeleft     : 11:59:35
```

IMPORTANT The only differentiation between GROUP and ROLE happens when generating the proxy. *One will be automatically member of all his/her GROUP but =ROLE=s has to be defined explicitely.!!* See the LHCb example, they are using the following FQANs:

```
"/VO=lhcb/GROUP=/lhcb/sgm":::sgm:
"/VO=lhcb/GROUP=/lhcb/lcgprod":::prd:
"/VO=lhcb/GROUP=/lhcb":::
```

So whenever they generate a proxy they will be always Production managers ! ***This is dangerous , not recommended.*** Either

- ◆ they have to use a second certificate for non-production submissions
- ◆ or explicitely exclude group membership upon proxy generation. None of which is convenient. For job submission it's not a problem but for storage it implies risk, since storage (for example DPM > 1.6.4-3) already interprets secondary groups.

How user mapping works

Format of gridmap files

A request can arrive with VOMS proxy or grid-proxy.

- For grid-proxy the /etc/grid-security/grid-mapfile is used,
- For VOMS proxy the /opt/edg/etc/lcmaps/gridmapfile and /opt/edg/etc/lcmaps/groupmapfile is used to do the mapping.

Their format is the same:

- An entry starting with '.' stands for a prefix
- Entries without '.' stand for static accounts.

Extract from

```
"/VO=atlas/GROUP=/atlas/ROLE=lcgadmin/Capability=NULL" .atlassgm
"/VO=atlas/GROUP=/atlas/ROLE=lcgadmin" .atlassgm
"/VO=atlas/GROUP=/atlas/ROLE=production/Capability=NULL" .atlasprd
"/VO=atlas/GROUP=/atlas/ROLE=production" .atlasprd
"/VO=atlas/GROUP=/atlas/Role=NULL/Capability=NULL" .atlas
"/VO=atlas/GROUP=/atlas" .atlas
```

Additional comments

- Due to a bug all entry in lcmaps files is repeated twice, with and without the

Role=NULL/Capability=NULL

tag.

- The Capability tag is ***deprecated***, not used, not interpreted.
- The accounts are not recycled
- If there is no sufficient pool account then it may happen that an ordinary user will be mapped to special user, if their prefix is the same:

```
dteam - for ordinary user
dteamsgm - for lcgadmin user
```

to protect against this:

- ◆ ***one should define sufficient number of pool accounts***, or
- ◆ ***use not overlapping prefixes*** like:

```
dteam - for ordinary dteam user
sgmdteam - for dteam lcgadmin user
```

- Static mapping has precedence over pool users !

- ◆ Old situation: If a DN was found for pool user and production user as well, then it's production mapping got written into the gridmapfile.
- ◆ New situation: Since production users are not static accounts any more, they have no precedence over ordinary pool users, so in the config file special users has to be defined first !

```
# ATLAS
# Map VO members  (sgm)
group vomss://lcg-voms.cern.ch:8443/voms/atlas?/atlas/Role=lcgadmin .atlassgm
group vomss://voms.cern.ch:8443/voms/atlas?/atlas/Role=lcgadmin .atlassgm
group vomss://lxb1928.cern.ch:8443/voms/atlas?/atlas/Role=lcgadmin .atlassgm

# Map VO members  (prd)
group vomss://lcg-voms.cern.ch:8443/voms/atlas?/atlas/Role=production .atlasprd
group vomss://voms.cern.ch:8443/voms/atlas?/atlas/Role=production .atlasprd
group vomss://lxb1928.cern.ch:8443/voms/atlas?/atlas/Role=production .atlasprd

# Map VO members  (root Group)
group vomss://lcg-voms.cern.ch:8443/voms/atlas?/atlas/lcg1 .atlas
group vomss://voms.cern.ch:8443/voms/atlas?/atlas/lcg1 .atlas
group vomss://lxb1928.cern.ch:8443/voms/atlas?/atlas/lcg1 .atlas
```

How to configure users

Advices on user configuration:

<http://glite.web.cern.ch/glite/packages/R3.0/deployment/glite-known-issues.asp>

Site admins should ensure that a sufficient number of software and production manager pool accounts are created on the nodes that need them (this does not apply to VOBOX and SE_castor). Here is how to estimate a lower bound on the number necessary per VO:

```
awk '$NF ~ /(prd|sgm)$/ { print $NF }' /etc/grid-security/grid-mapfile | sort | uniq -c
```

For the LHC VOs the current numbers are the following:

```
4 aliceprd
36 alicesgm
35 atlasprd
33 atlassgm
50 cmsprd
25 cmssgm
18 dteamprd
60 dteamsgm
9 lhcbprd
5 lhcbsgm
```

Questions to clarify

- Is the absence of a tag means it is equal NULL ? Is the following two same *by definition* or *by convention* ?:

```
/VO=atlas/GROUP=/atlas/ROLE=lcgadmin/Capability=NULL
VO=atlas/GROUP=/atlas/ROLE=lcgadmin/
```

- Is the order of tags does matter ? Are they really the same ?:

```
/VO=atlas/GROUP=/atlas/ROLE=lcgadmin/Capability=NULL
/VO=atlas/GROUP=/atlas/Capability=NULL/ROLE=lcgadmin
```

About fair shares

Torque configuration

- Since by now sgm users has different primary groups they have to be explicitly allowed to be able to use a queue.
- YAIM's "queuename_GROUP_ENABLE" variable is for this purpose. Setting for example for the 'atlas' queue:

```
ATLAS_GROUP_ENABLE="atlas /VO=atlas/GROUP=/atlas/ROLE=lcgadmin /VO=atlas/GROUP=/atla
```

will result

```
[root@lxr2018 cert-TB-config]# qmgr
Max open servers: 4
Qmgr: print server
#
# Create queues and set their attributes.
#
```

```

#
# Create and define queue atlas
#
create queue atlas
set queue atlas queue_type = Execution
set queue atlas resources_max.cput = 48:00:00
set queue atlas resources_max.walltime = 72:00:00
set queue atlas acl_group_enable = True
set queue atlas acl_groups = atlas
set queue atlas acl_groups += atlassgm
set queue atlas acl_groups += atlasprd
set queue atlas enabled = True
set queue atlas started = True
#

```

So 3 group of user is allowed to submit to this queue.

Publishing it in the infosys

Different users need different priorities. The CE has to advertize how the priorities are handled, and during the matchmaking the WMS should be aware of this. This is realized with the usage of VoViews. **VoView** information is used by WMS only, the lcg-RB ignores it.

- Each queue is present in the information system as a GluCEUniqueID.

```

# lxb2034.cern.ch:2119/jobmanager-lcgpbs-atlas, local, grid
dn: GlueCEUniqueID=lxb2034.cern.ch:2119/jobmanager-lcgpbs-atlas,mds-vo-name=local,o=grid
objectClass: GlueCETop
.
.
.
GlueCEAccessControlBaseRule: VO:atlas
GlueCEAccessControlBaseRule: VOMS:/VO=atlas/GROUP=/atlas/ROLE=lcgadmin
GlueCEAccessControlBaseRule: VOMS:/VO=atlas/GROUP=/atlas/ROLE=production
.
.
.
```

For each GlueCEAccessControlBaseRule there is a GlueVOView block defined, which contains the information valid for the given group.

```

# /VO=atlas/GROUP=/atlas/ROLE=lcgadmin, lxb2018.cern.ch:2119/jobmanager-lcgpbs-atlas, local, grid
dn: GlueVOViewLocalID=/VO=atlas/GROUP=/atlas/ROLE=lcgadmin,GlueCEUniqueID=lxb2018.cern.ch:2119/jo
objectClass: GlueCETop
objectClass: GlueVOView
objectClass: GlueCEInfo
objectClass: GlueCEState
objectClass: GlueCEAccessControlBase
objectClass: GlueCEPolicy
objectClass: GlueKey
objectClass: GlueSchemaVersion
GlueVOViewLocalID: /VO=atlas/GROUP=/atlas/ROLE=lcgadmin
GlueCEAccessControlBaseRule: VOMS:/VO=atlas/GROUP=/atlas/ROLE=lcgadmin
GlueCEStateRunningJobs: 0
GlueCEStateWaitingJobs: 0
GlueCEStateTotalJobs: 0
GlueCEStateFreeJobSlots: 2
GlueCEStateEstimatedResponseTime: 0
GlueCEStateWorstResponseTime: 0
GlueCEInfoDefaultSE: lxb1921.cern.ch
GlueCEInfoApplicationDir: /opt/exp_soft/atlas

```

```

GlueCEInfoDataDir: unset
GlueChunkKey: GlueCEUniqueID=lxr2018.cern.ch:2119/jobmanager-lcgpbs-atlas
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

and for example

```

# atlas, lxr2018.cern.ch:2119/jobmanager-lcgpbs-atlas, local, grid
dn: GlueVOViewLocalID=atlas,GlueCEUniqueID=lxr2018.cern.ch:2119/jobmanager-lcg
    pbs-atlas,mds-vo-name=local,o=grid
objectClass: GlueCETop
objectClass: GlueVOView
objectClass: GlueCEInfo
objectClass: GlueCEState
objectClass: GlueCEAccessControlBase
objectClass: GlueCEPolicy
objectClass: GlueKey
objectClass: GlueSchemaVersion
GlueVOViewLocalID: atlas
GlueCEAccessControlBaseRule: VO:atlas
GlueCEAccessControlBaseRule: DENY: VOMS:/VO=atlas/GROUP=/atlas/ROLE=production
GlueCEAccessControlBaseRule: DENY: VOMS:/VO=atlas/GROUP=/atlas/ROLE=lcgadmin
GlueCEStateRunningJobs: 0
GlueCEStateWaitingJobs: 0
GlueCEStateTotalJobs: 0
GlueCEStateFreeJobSlots: 2
GlueCEStateEstimatedResponseTime: 0
GlueCEStateWorstResponseTime: 0
GlueCEInfoDefaultSE: lxr1921.cern.ch
GlueCEInfoApplicationDir: /opt/exp_soft/atlas
GlueCEInfoDataDir: unset
GlueChunkKey: GlueCEUniqueID=lxr2018.cern.ch:2119/jobmanager-lcgpbs-atlas
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

The presence of the DENY tag results that the

```
VOMS:/VO=atlas/GROUP=/atlas/ROLE=production
```

user won't match with the general atlas VOView.

Comments:

- The GlueAccessControlBaseRule variable is inclusive, i.e. for example /atlas will allow also the atlas production manager.
- The name of the VOView (GlueVOViewLocalID) can be freely chosen it should not necessarily match the name of the VO or the VOMS FQAN. The only restriction that it should be unique inside a GlueCEUniqueID block.
- Under SLC4 due to a stricter schema checking the GlueVOViewLocalID cannot contain the '=' sign.

How this information is generated

The lcg-info-dynamic-scheduler

The dynamic scheduler framework provides an interface to the infosystem, plugin modules can be used and adjusted for various batch systems.

The dynamic scheduler itself is a plugin for the GIP, its wrapper script resides in /opt/lcg/var/gip/plugin directory.

It comes via several rpms:

```
lcg-info-dynamic-scheduler-generic-2.1.0-1
lcg-info-dynamic-scheduler-condor-0.2.0-1 -- old one (Laurence), but new version is already avail
lcg-info-dynamic-scheduler-pbs-2.0.0-1
lcg-info-dynamic-scheduler-lsf-1.0.1-1.noarch.rpm
lcg-info-dynamic-lsf-2.0.34-1.noarch.rpm
lcg-info-dynamic-sge -- used in production, will be included into the release.
```

There are two parts to the system:

- One part (lcg-dynamic-info-scheduler program) contains the algorithm that computes the response times. This part doesn't know the details of the underlying batch system, so that the estimated times are as independent as possible from the various LRMS.
- The second part is the LRMS specific part. This gathers information from the LRMS and writes it out in a LRMS-independent format. There are two of these critters; one for the LRMS state, and one for the scheduling policy.
 - ◆ the lrms_backend_cmd provides information about the status of the batch system, while
 - ◆ the vo_max_jobs_cmd's output contains the max job slots defined per groups

The lrms_backend_cmd script output is

```
nactive      342
nfree        0
now          1164968613
schedCycle   26
{'queue': 'qlong', 'state': 'queued', 'qtime': 1164799042.0, 'group': 'atlas', 'user': 'atlas082'}
{'queue': 'qlong', 'state': 'queued', 'qtime': 1164822293.0, 'group': 'biome', 'user': 'biome050'}
```

while the vo_max_jobs_cmd gives:

```
{
'biome': 171,
'pvier': 4,
'users': 50,
'geant': 2,
'ops': 32,
'DEFAULT': 330,
'zeus': 132,
'cms': 10,
'esr': 32
}
```

Using this information the scheduler tries to guess the Estimated Response Time and Free Slots for a VOView.

Configuration file

```
static_ldif_file: /opt/lcg/etc/static-file-CE.ldif
vomap :
    alicevo:alice
    aliceprd:/VO=alice/GROUP=/alice/ROLE=production
    astr:astrop
    atlas:atlas      -- OPTIONAL
    lhcbsgm:/VO=lhcb/GROUP=/lhcb/ROLE=lcgadmin
    atlassgm:/VO=atlas/GROUP=/atlas/ROLE=lcgadmin
    biome:biomed
    cms:cms          -- OPTIONAL
lrms_backend_cmd : cat /opt/lcg/libexec/lrm-backend-output.example
vo_max_jobs_cmd : cat /opt/lcg/libexec/vo-max-slot.example
cycle_time : 0
```

The lcg-info-dynamic-scheduler

Divers

- VOView is called VOView but is a generic concept, groups and individuals can be reported using VOView blocks.
- With the increasing number of groups the information system will grow and the present architecture will face serious problems.
- The concept of queues should not be visible in the information system. A queue should be inner to the site, especially because there are batch systems having no concept of queue.

Some info on YAIM

New structure

Presently there exist three branches of YAIM:

- YAIM 3.0.1 – for SLC3 services in production
- YAIM 3.1 – for SLC4 UI and WN in (pre)production
- YAIM 3.1.1 – for SLC4 and SLC3 services – in preparation (to be released in ~3 weeks) and will obsolete 3.0.1 and 3.1.

New structure: YAIM will come via several rpm:

- glite-yaim-core – to provide the framewrok, common functions and utilities.
- glite-yaim-clients – to configure the UI, WN and VOBOX
- glite-yaim-lfc – to configure the LFC
- glite-yaim-fts – to configure the FTS
- glite-yaim-dpm – to configure the DPM
- glite-yaim-myproxy – to configure the MyProxy
- glite-yaim-ce – to configure the lcg-CE and glite-CE
- glite-yaim-wms – to configure the WMS
- glite-yaim-lb – to configure the LB
- glite-yaim-dcache – to configure dCache

New coordination:

- YAIM coordination is now leaded by Maria Allandes Pradillo
- The aim is to achive a more distributed and opened development process.
- Faster reaction time is enabled by the splitting of the modules, no patches will hold up each other, YAIM modules can be released independently.
- Support for other batchsystems is foreseen
- YAIM moves away from delivering configuration files. Only examples will be provided and the correct settings concerning VOs, users, VOMS servers will be availabe in YAIM format through the CIC portal.

Contact:

- ROC managers and experiment people: Use the yaim-contact@cernSPAMNO.ch mailing list to send your request or to ask about the status of some request. This is not a support list, just a shortcut to help the communication with the YAIM people.
- Everyday user is encouraged to use GGUS and/or Savannah

Links:

- YAIM main page: <https://twiki.cern.ch/twiki/bin/view/EGEE/EGEEConfiguration>
- YAIM guide 3.1: <https://twiki.cern.ch/twiki/bin/view/EGEE/LCG.YaimGuide310>
- YAIM guide 3.0.1: <https://twiki.cern.ch/twiki/bin/view/LCG/YaimGuide301>
- Generic Install Guide 3.1: <https://twiki.cern.ch/twiki/bin/view/LCG/GenericInstallGuide310>
- Generic Install Guide 3.0.1: <https://twiki.cern.ch/twiki/bin/view/LCG/GenericInstallGuide301>
- YAIM developer's Guide: <https://twiki.cern.ch/twiki/bin/view/EGEE/YaimDevelopersGuide>
- YAIM planning: <https://twiki.cern.ch/twiki/bin/view/LCG/YaimPlanning>