VOs and ARC

Florido Paganelli, Lund University

Jan 2016, ARC Nordforsk Developers Camp

What is a VO?

- Virtual Organization
  - In a broad sense: a group of people working together
  - In practice: Organizations, research groups, experiments
- Examples:
  - Nordugrid
  - ATLAS
  - ALICE
  - CMS

How it is used on the GRID

- **Authorization**: A person belonging to a VO can access several resources.
  - For example in ATLAS a member of the ATLAS VO gets access to ATLAS resources.
- **Information system**:
  - **Resource discovery**: what queues serve ALICE?
  - **Brokering**: which CMS resources are free to be used? How many free job slots allocated for ALICE?
  - **Statistics/job discovery**: How many LHCb jobs are running in cluster X? Which cluster is running ALICE jobs?
- Accounting?

How is this concept implemented?

- A **VOMS server** associates a VO credential (typically in the form of an X.509 attribute certificate) to a member user
- The server can be queried for users belonging to a VO
- The proxy certificate can contain such attribute information (multiple VO attributes)
  - Generation of such credential requires to contact a VOMS server

How does it look like in GLUE1 For the experiments (CREAM)

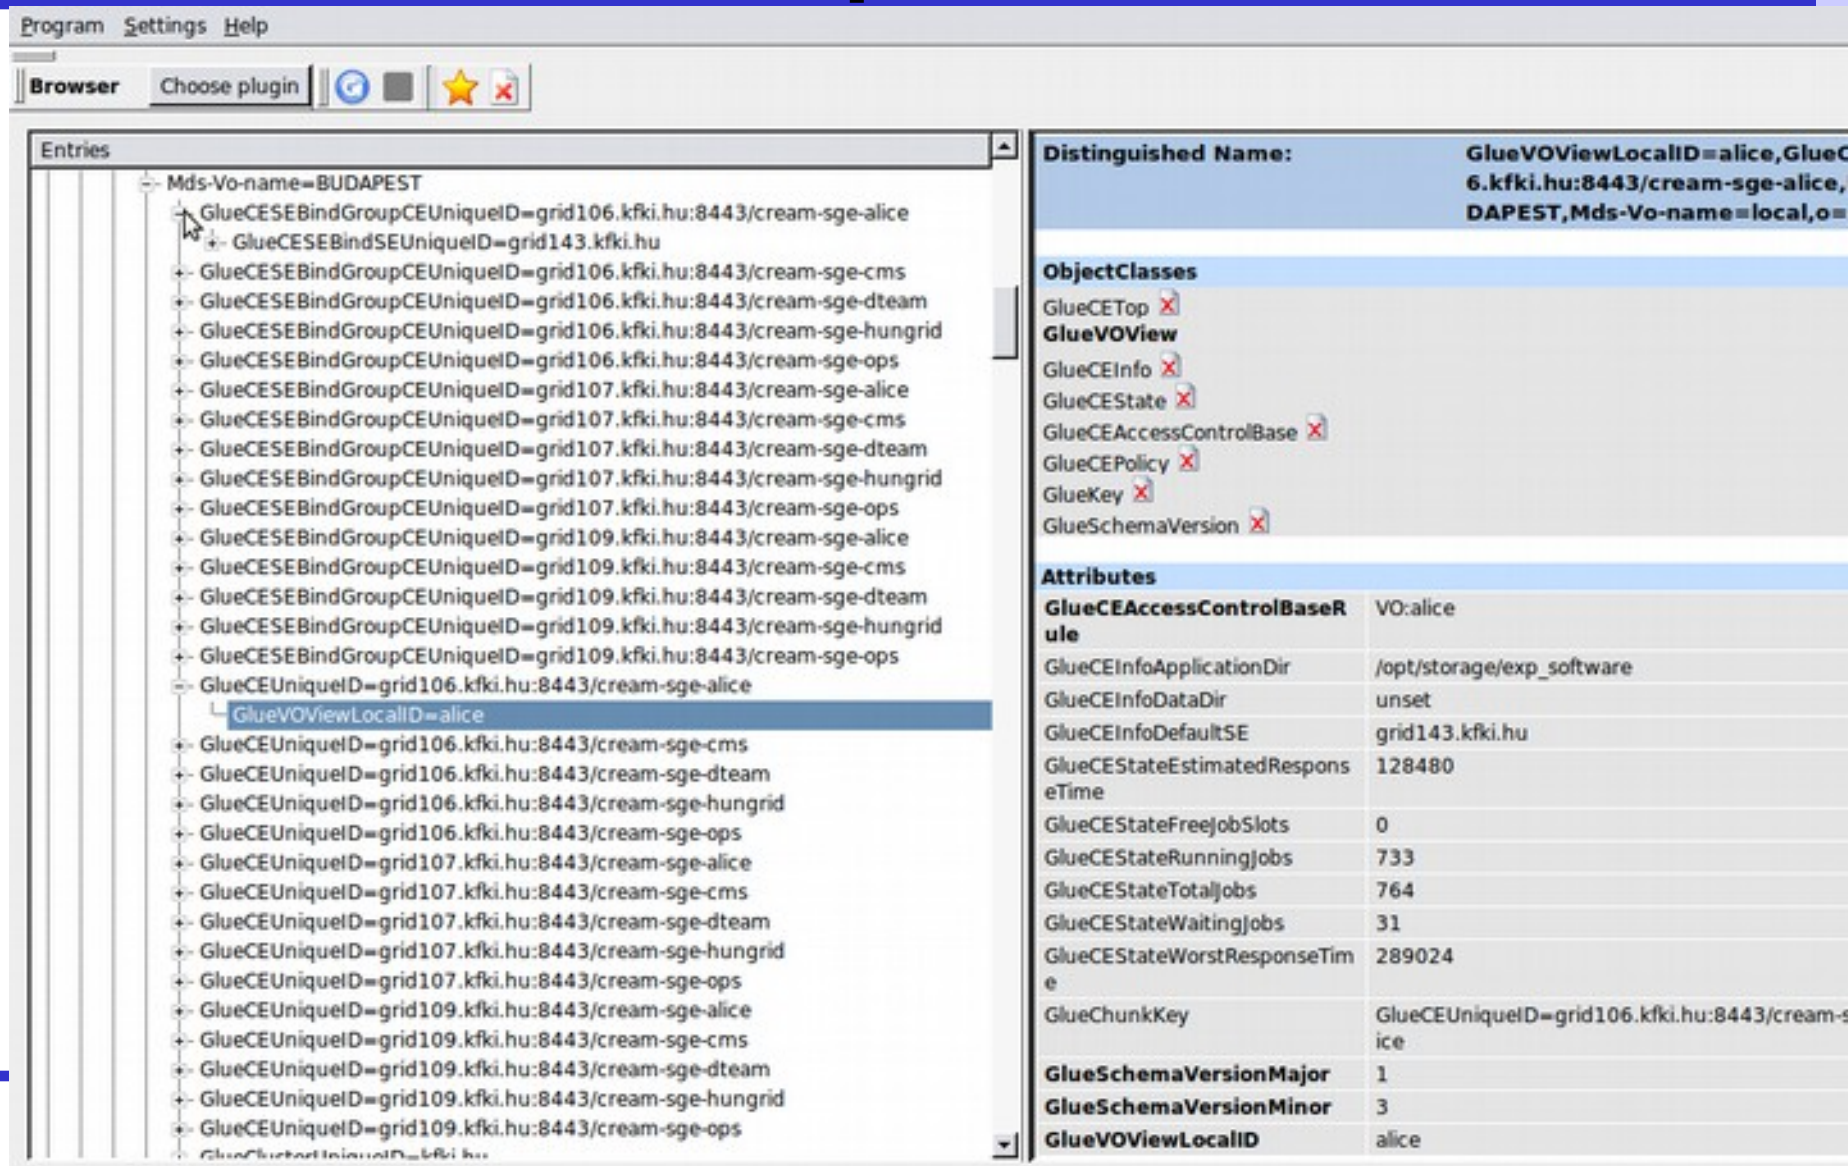

How does it look like in GLUE2 For the experiments (CREAM)

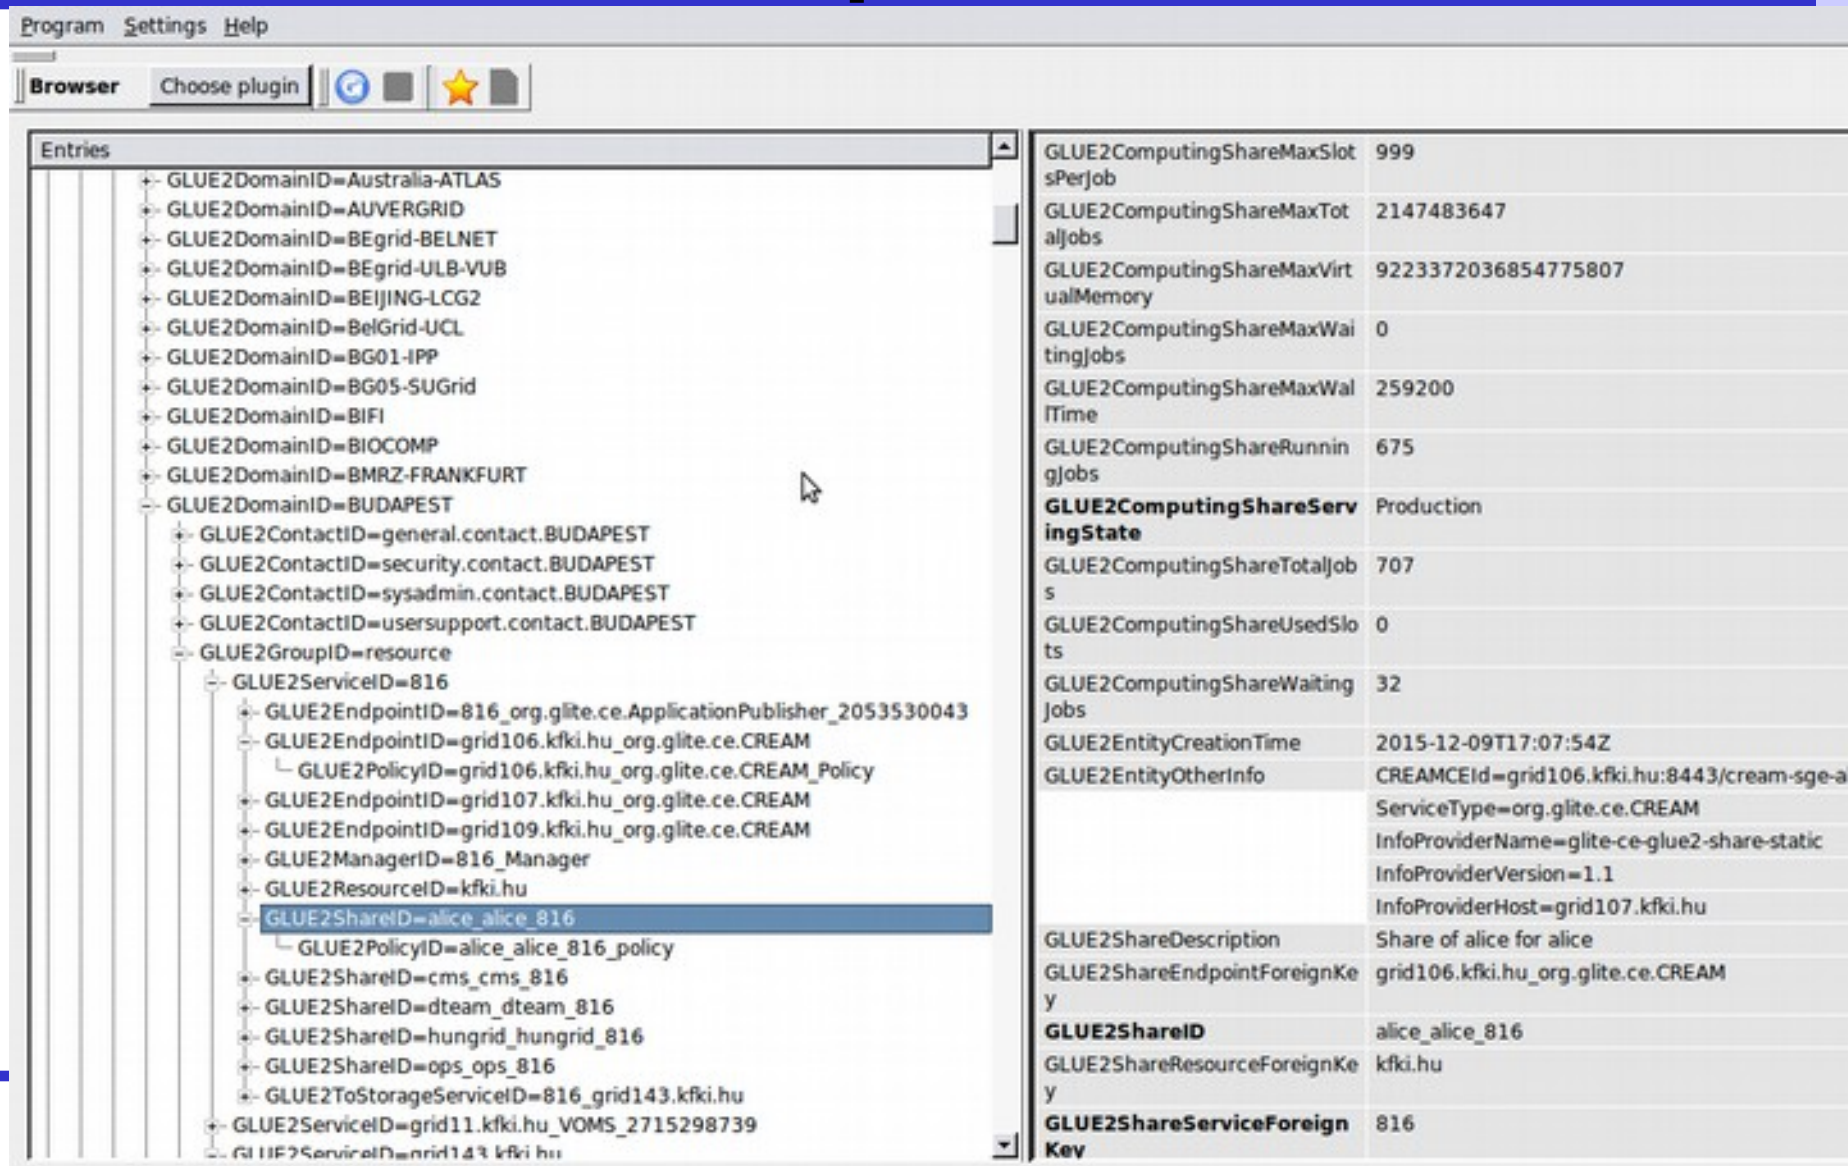

How does it look like in GLUE2 For the experiments (CREAM)

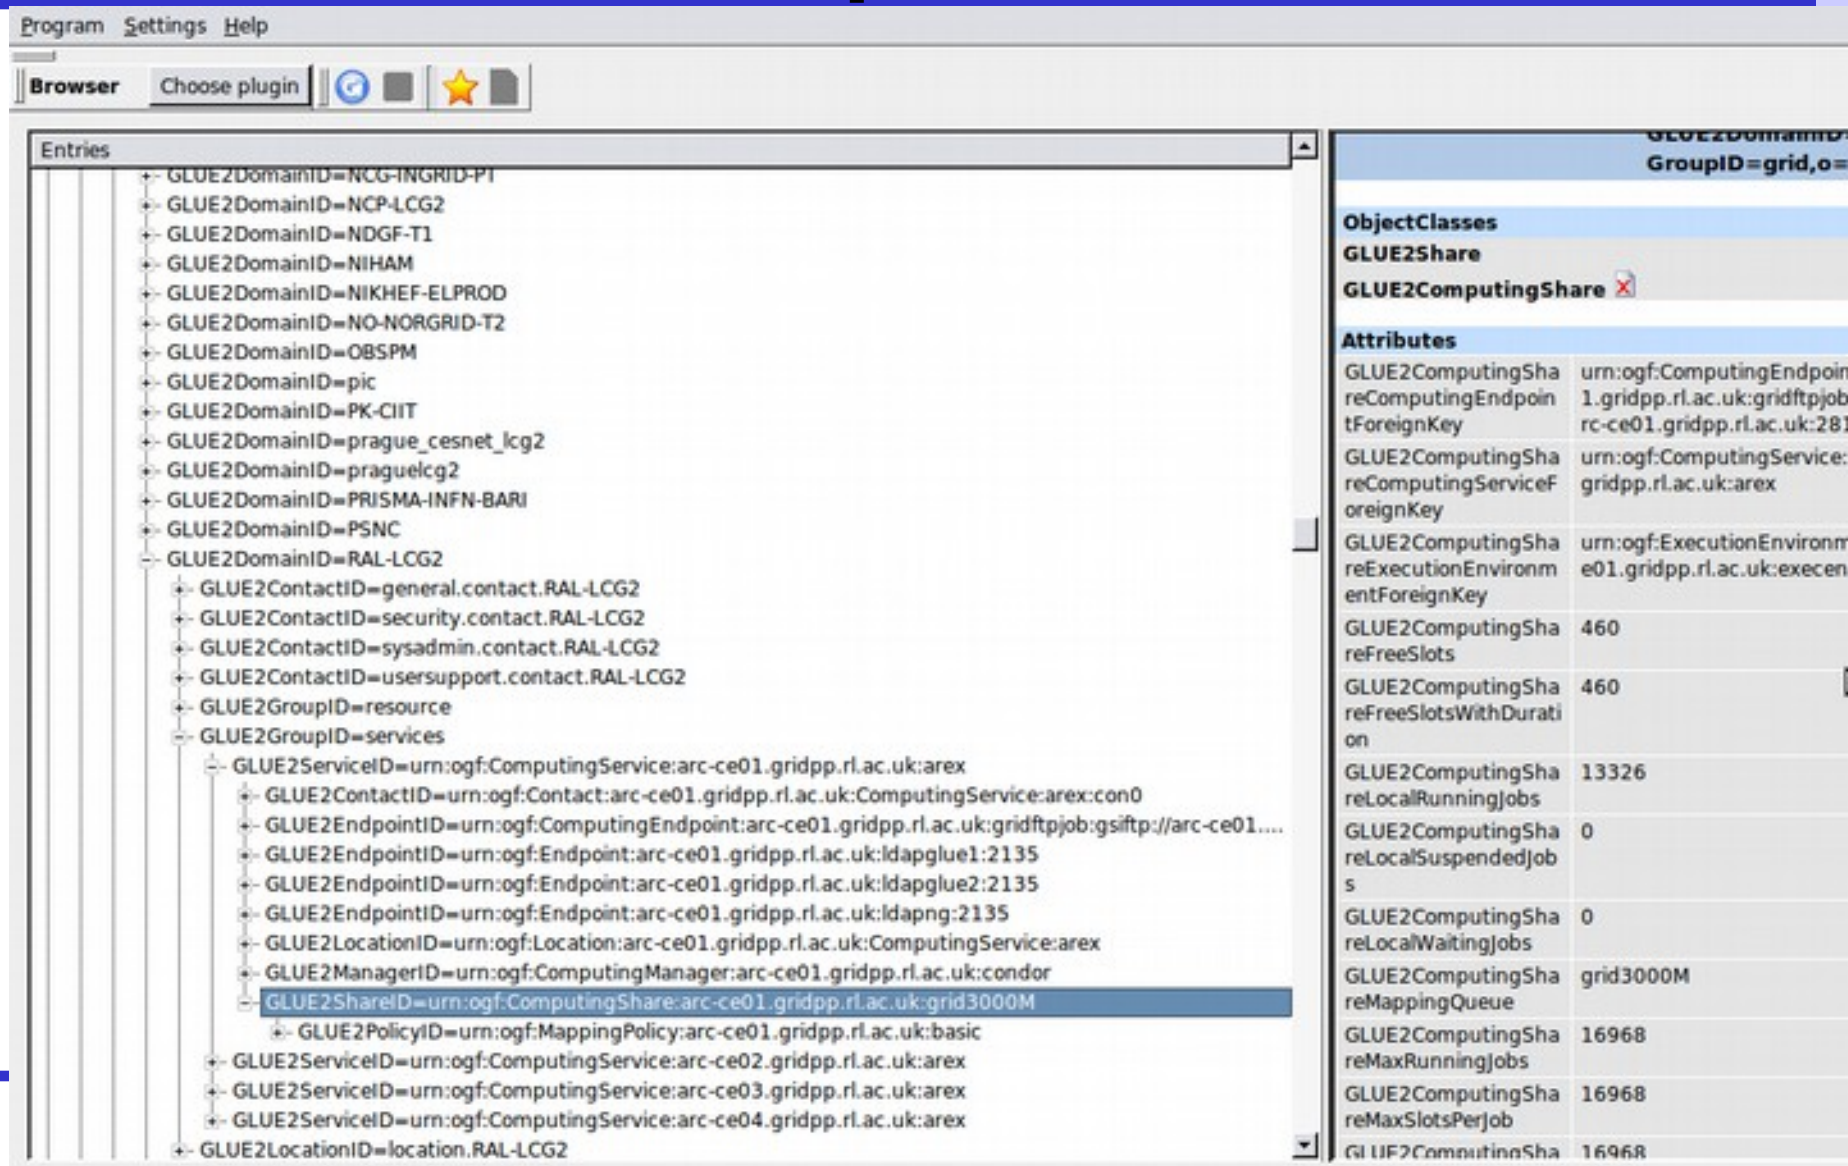

How does it look like in GLUE2 For the experiments (ARC)

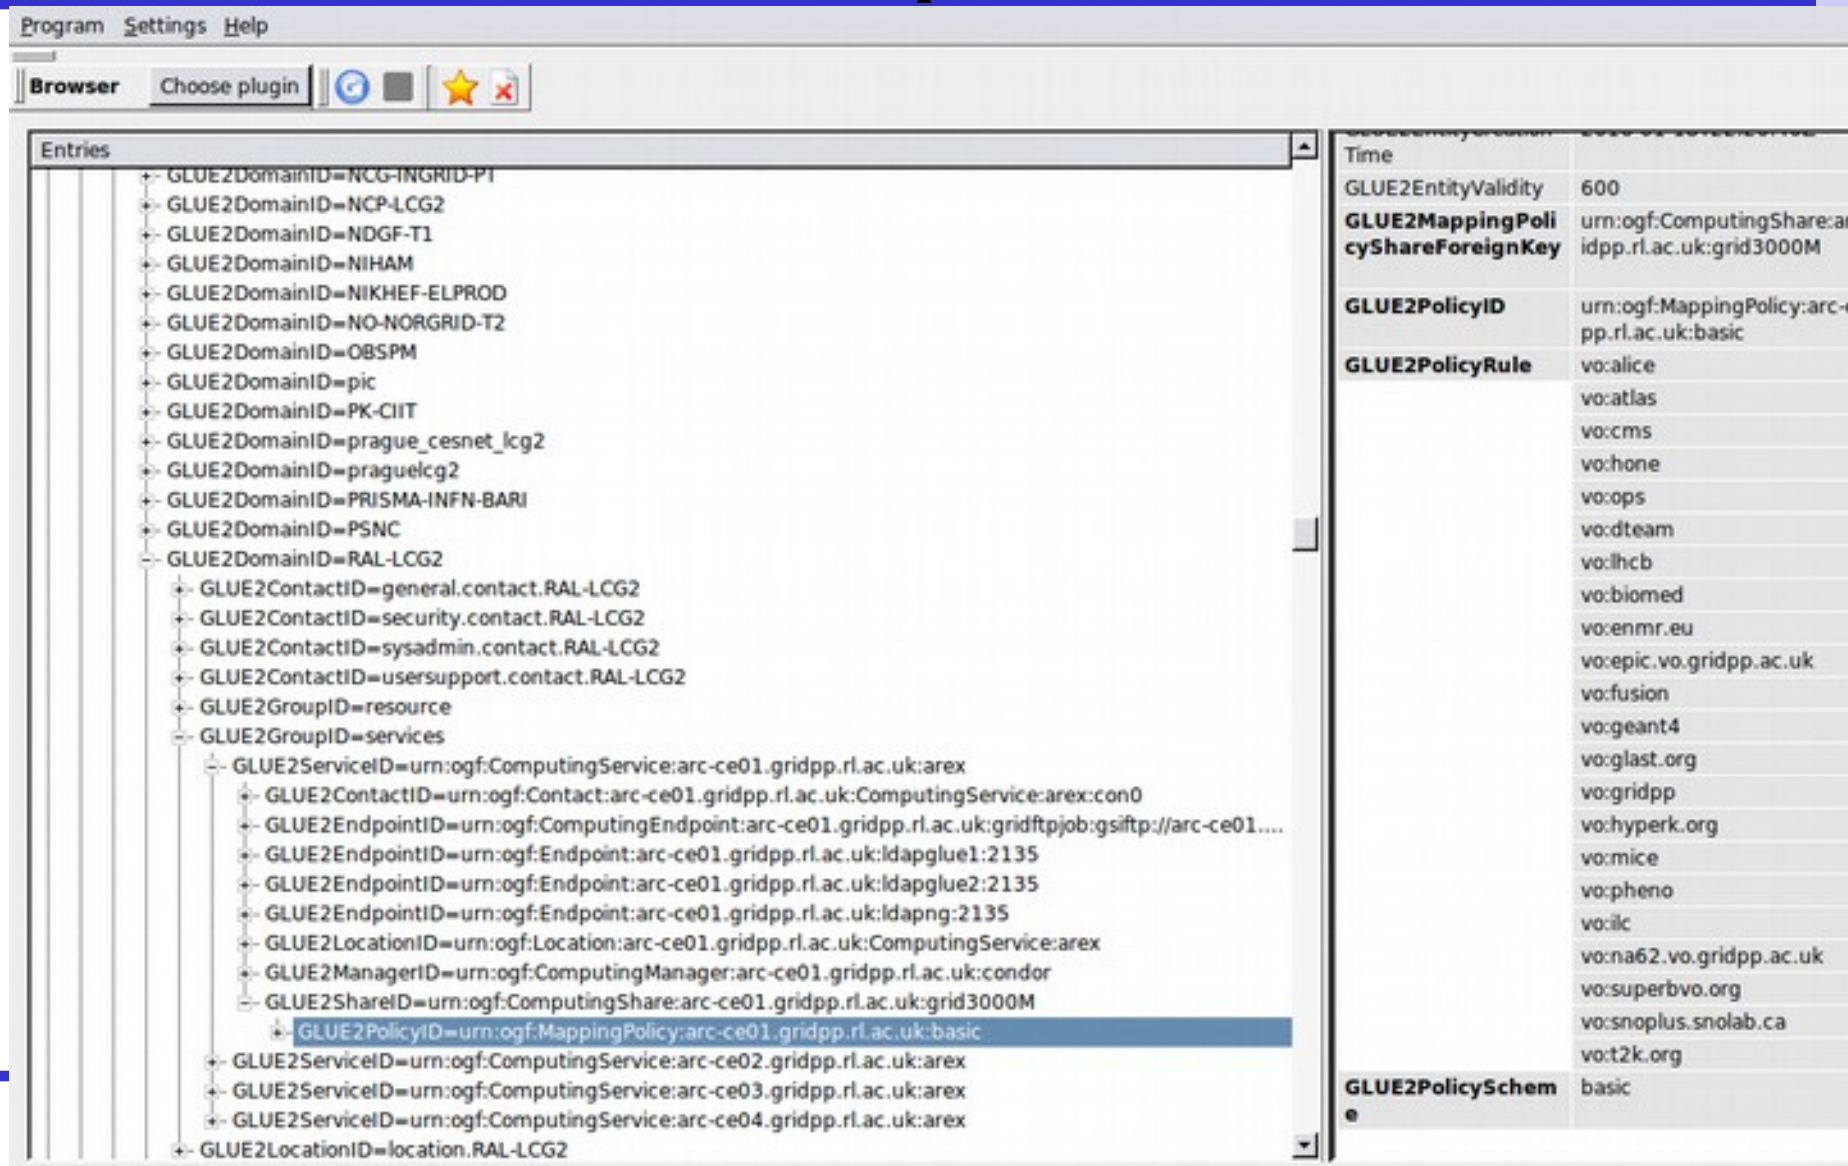

How does it look like in GLUE2 For the experiments (ARC)

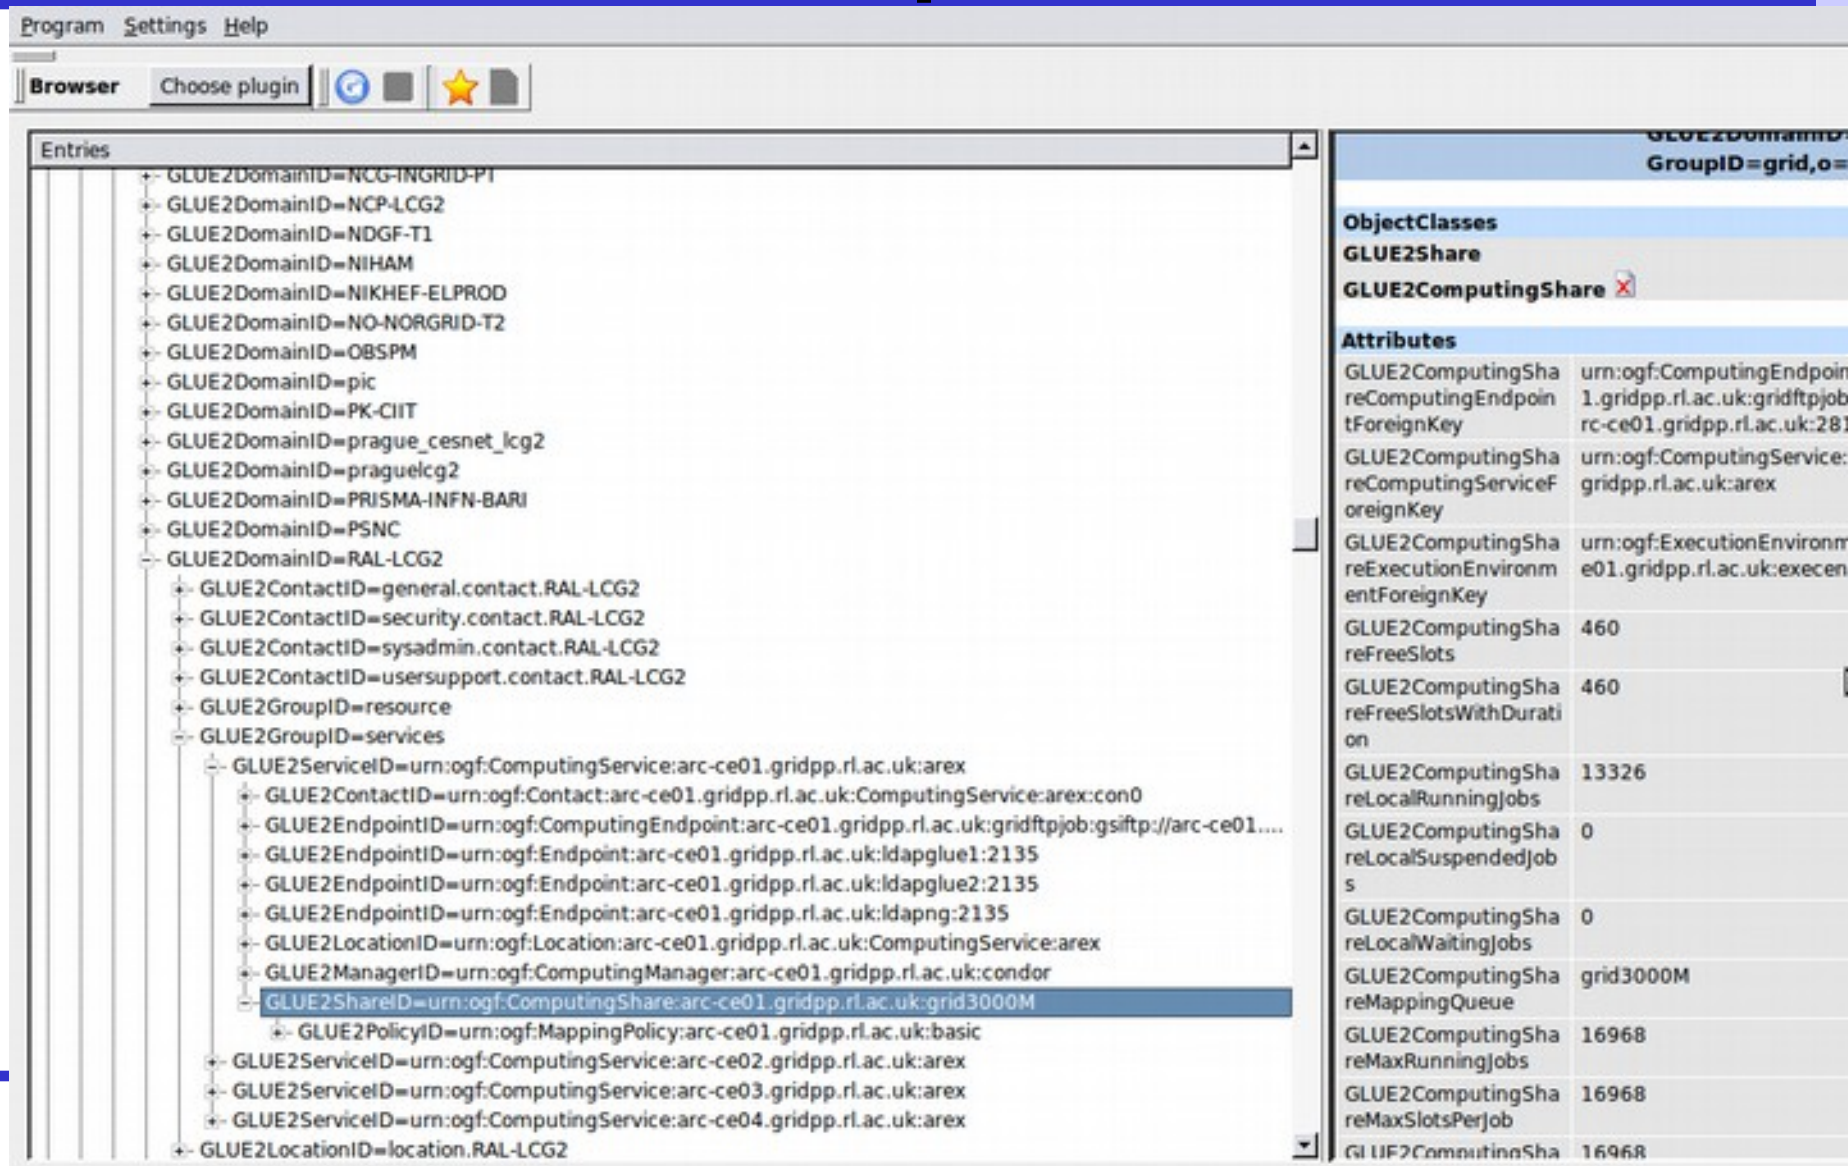

How does ARC use it

- ARC can be configured to get a list of users depending on their VO using the [group] [vo] blocks in `arc.conf`
- The list is retrieved from VOMS servers defined in the above blocks
- The list entries are added to grid-mapfile by the **nordugridmap** utility (runs independently from A-REX)
- A-REX just reads the grid-mapfile. It has no idea of VOs. Only users

How does ARC use it

- The only source of this information is either the [group] [vo] blocks or the proxies
- The above stuff is generally free-form strings defined on a best effort community based practice.
- i.e. ATLAS or atlas or AtLaS – in most cases are the same. Not to mention ATLAS_VO or the like.
- Moreover, a sysadmin is free to put in the [vo] blocks the names he/she wants. (VOATLAS, atlasvirtualorg...)

How does ARC use it

- Before ARC 5, it was not possible to assign different VOs to a queue (maybe it was in ARC 0.8). All queues had the same Vos.
- Now it is possible, though there is no authorization enforcement on such queues, i.e. the users can submit even if they don't belong (A-REX does not consider certificate attributes per queue afaik)
- The number of jobs is calculated per queue, not per VO.
  - Got a GGUS ticket about it!
    https://ggus.eu/index.php?mode=ticket_info&ticket_id=112543

Why ARC uses it like that

- The VO concept was never well defined
- VO Strings may vary from one VOMS server to another
- A user can belong to multiple VOs
- The batch system has no idea of VOs. The jobs themselves do not have knowledge of what VO they belong.
  Only the user belongs to some, but A-REX does not check for VO itself.

Main VO issues

- Experiments want job statistics per VO.
- A-REX used to discard such information.
- BUG 3495 has a patch for A-REX, contains a discussion with a plan how to fix it.
  - There is a patch by Aleksandr that saves the most important VO name in the proxy certificate into .local file for a job, but it doesn't work for me!
  - Andrii thinks that limiting to the first VO is not sufficient
  - The experiments claim that in the end only one VO per user is actually used
- Our GLUE1 implementation is not flexible wrt to this

How was it solved for CREAM?

- In the CREAM world, sites used to write their own infocollector to target a specific VO
- My attempts to suggest ways to configure ARC to get the same behavior were rejected (see GGUS ticket)
- Implementation for GLUE1 was native, infocollector architecture more flexible than ours (but they only work with BDII/LDAP)

How we plan to solve it

- Make A-REX aware of what VO does a job (or rather, its user) belong to (partially done, need to interact with Aleksandr to understand why it doesn't work for me)
- Make arc.conf configurable wrt which VO is served by each queue (done!)
- Create a PERL module that can work both for GLUE1 and GLUE2 to calculate jobs statistics
- Generate GLUE1/2 objects in a totally different way
  - ! Might create a new bottleneck: additional two GLUE1 and GLUE2 objects for each VO, each object needs the logic to calculate jobs per VO

How to assign a VO to a job?

- Choosen:
  - Read the user proxy
  - Configure authorization in arc.conf using [group]
  - Have an arc.conf configuration where VOMS credentials are checked and enforced first
  - The VO information with which the user was authorized is written in the job.id.local

Rejected solutions

- Add an option to xrsl to specify the VO
- Add an option to xrsl to specify a VOMS server
- Add an external service that tells what VO the job belongs to
- Use the infosys rendering to assign VO to a job, i.e. the selected queue/share determines the VO, A-REX must authorize only if the user is allowed to run on that VO

Reference Bugs and GGUS tickets

- https://ggus.eu/index.php?mode=ticket_info&ticket_id=112543
- https://bugzilla.nordugrid.org/show_bug.cgi?id=3495
- https://bugzilla.nordugrid.org/show_bug.cgi?id=3433
- https://bugzilla.nordugrid.org/show_bug.cgi?id=3464
-