VOMS deployment in GridPP and NGS

Alessandra Forti[1],

Mike Jones[2], Sergey Dolgobrodov[1]

[1)] School of Physics and Astronomy, [2)] Manchester Computing,
The University of Manchester, Manchester UK

Abstract

We describe our experience in practical deployment of the gLite VOMS (Virtual Organization Management Service). The formation of all sizes of groups with similar research agenda can be viewed as the emergence of Virtual Organisations (VOs). The deployment of centralised VOMS servers at a national level as part of a grid infrastructure facilitates the creation of these groups and their subsequent registration to the existing grid infrastructures. It also helps grid resources to engage user communities by dealing with them in a more scalable way. In the following we will describe the use cases, some of the technical aspects, and the deployment and administration of such a VOMS server. The evaluation of robustness of the VOMS releases 1.4, 1.5 and 3.0 and known problems are also described.

1. Introduction

In a grid environment there is a fundamental difficulty in the granting access rights to resources: users are no longer able to be recognised through their institutional login procedures nor do they belong to well defined local entities which are recognised outside those institutes.

Not long after the emergence of the Grid paradigm came the necessary concept of the Virtual Organisation (VO). VOs were created to solve the problem of identifying groups of abstract entities in online communities. These communities have, in some cases, become identified with groups of user. Access control policies combined with this type of VO have enabled grids to construct middleware to decide who gets access to what on their grid. VOs don't have to have geographical, administrative or even national boundaries. For this reason a strict definition of a VO is difficult to reach consensus upon.. For clarity here we will define a VO as a set of users who have signed an AUP (Acceptable Usage Policy) and that that VO in its own right is able to be given the authority to use a certain percentage of resources under different administrative control.

In the particle physics community for example it has long been understood that each High Energy Physics (HEP) experiment maps neatly onto a VO. Each experiment gets access to its resources and takes its share in the total number of cycles on the HEP's grid when available.

There are many other ways to form VOs besides this all of which would carry huge overheads if all VO members needed to register with all resources available in the emerging grids today.

2. Authorisation on a Grid and the choice VOMS

Authorisation in a grid context is the process which determines whether an entity may gain access to a resource within that grid. Usually by the time a resource is making an authorisation decision for an incoming request the authenticity of that request has been determined. This separation is key to the Public Key Infrastructures used in a number of grids today. It allows a user to identify themselves using a robust and universally acceptable token, and authorisation decisions to be made against that identity. This provides a useful separation aligning an identity (an entity will usually only have one of these) with an assertion from an authority which is well known, at the same time keeping that identity's properties and attributes (e.g. group membership, rolls, capabilities, etc.) separate from this assertion. Allowing a more scalable processes to later assert necessary properties by which authorisation can be determined.

It is for Authorisation that VOMS has been chosen as the provider of authorisation attributes in the GridPP and NGS grids. It provides all the modern security techniques like Single-sign on, delegation, non-repudiation and many more.

Thus, unlike most other security projects, VOMS does not focus on developing password based and local account based security applications or services. Instead, VOMS is meant to act as a server capable of securing all

services in a Grid environment, and exposing their functionality without putting resources and services in jeopardy. This enables the resource providers and grid enabled service provider to share and host their service and resource with full confidence which will help Computational Grid to grow.

3. VOs support in UK

3.1 Examples of local use cases

- Small site group with no resources.
- Small national experiment or research projects like MINOS, CEDAR and RealityGrid with no resources.
- A local group of a bigger VO that owns local resources and do not want to share them but want to access them with the same set of tools. (Typical of this case are data servers for local users.)
- Grid application development test machines.
- Different grids might want to cooperate to support each other resources.
- The local funding situation imposes to share resources with groups not belonging to any VO; however they might be willing to access their portion of resources through the grid.
- Distributed Tier 2 sites[1] might want to unify certain categories of users in one VO.
- A site might want to give temporary access to one group resources to another local group.

3.2 GridPP and NGS

In UK two different national grid organizations GridPP and NGS (National Grid Service) have decided jointly support gLite VOMS servers.

- Common infrastructure to maintain the VOMS servers
- Common VOs support
- Common distribution of information
- Enable each other VOs on each other systems

The pool consists of two front end servers one for NGS and one for GridPP, two backup servers and a test server. The servers are hosted at the in Manchester as part of the Tier2 and NGS infrastructures. They have been running since January 2006 and they now host 9 national VOs and 2 local.

[1] A tier 2 site provides regional Compute Power within the EGEE and GridPP.

3.3 Enabling a VO

A formal request has to be made to the management. The following information about the VO has to be supplied in the request.

- VO name that conforms to the LCG/EGEE guidelines. The latest format was DNS format to avoid names for different VOs clashing. For example minos.gridpp.ac.uk is an acceptable name. However short UNIX like user names are still used for practical reasons.
- VO support contacts. Specific people and mailing lists.
- Security contacts - two people who can respond quickly in the event of a security incident relating to a member of the VO, or to the VO as a whole.
- VO services end points like file catalogs.
- Hardware requirements - memory size, disk space etc.
- Software requirements - any software beyond the basic Linux tools/libraries, including things which are part of standard distributions as they may not be installed by default.
- Typical usage pattern - expected job frequency and variation over time, job length, data read and written per job etc.
- Glue schema fields used - this would give an idea of what is really used in the information system and needs to be ensured to be properly set and maintained.
- General procedures – like VO software installation
- Roughly the number of users expected to use the resources, to give a guide to how many pool accounts to create.

The request has to be approved by the GridPP/NGS management. After approval the VO gets created on the VOMS server and the VO manager is enabled to add users. Sometimes the VO is too small and the VO administration is done by the VOMS manager under request. The information to enable the VO at sites is then downloadable from the GridPP/NGS WEB sites. VOs are considered responsible for the maintenance of the information in their own interest.

4. Technical aspects of VOMS

4.1 gridmap-files to VOMS awareness; from individual authorization to virtual organisation

Production grid services providing simple access to data and compute resources have, to date, dealt with authorization on an individual by individual base. There has been a trend to supply some level of delegation to this process.

VOMS mechanisms have the potential to provide a level of delegation such that authorisation decisions may be taken based on membership to a virtual organization, removing the requirements of pre-registration of individuals. This section describes the evolution of these mechanisms.

4.2 gridmap-files

The purpose of a gridmap-file is to translate an incoming request whose originator's identity is known as a string representation of an X509 namespace (their distinguished name, DN) to a local system identity: the username of a local account. It is a flat file, each line containing a DN in double quotes followed by a comma separated list of local user identities. The service gateway having obtained the user's DN through the context of the SSL negotiation searches the gridmap file line by line until it finds a match. Both the account and the mapping must exist before the user is authorised to use the service.

4.3 Pool Accounts and LDAP directories

The prerequisite of a system account forces each prospective user to register with each system. In a grid comprising of many resources and many users this registration process and account generation is not scalable. The Pool account mechanism goes some way to providing a mechanism with which to address this. Pool accounts system simply leases system accounts to end grid users for a specified duration.

Having only partially addressed the registration issues (accounts are available but authorization to use them has still to be granted) a system of populating the gridmap file automatically is used. This system creates mappings to sets of (pools of) accounts. To date this has been achieved by regularly retrieving users' DNs which have been published by various organisations in secure LDAP servers. With all of the above in place access to resources can be granted to predefined communities; membership of these communities can be maintained by the communities themselves. At the time of writing, this is where authorization policies are realised in production grids like the UK National Grid Services. The authorization mechanisms as described above have a number of drawbacks:

.

On line dependence for maintenance of the gridmapfile,

- Denial of service attacks changes the behaviour of the communities membership,
- No ability to deal with users who are members of multiple communities (Group membership),
- No ability to select differing levels of access within a recognised communities (roles and capabilities),
- Data Protection.

Out of the LCG and its sister project gLite has emerged more elegant authorization mechanisms. While the concept has been around for a while stable implementations are still relatively new.

4.4 lcas and lcmaps

LCAS (Local Centre Authorization Service) and LCMAPS (Local Credential Mapping Service) replace the gridmap file system with something a little more sophisticated. As the names suggest they split the process into separate authorization and the mapping mechanisms. LCAS is a plug-in[2] which is called from within a modified service, (currently there exists modified versions of the Globus gatekeeper and an earlier version of the GridFTP server). LCAS makes Authorization decisions based upon three distinct sets of policy data: users/groups allowed, users/groups disallowed[3], and service availability times. LCMAPS (Local Credential Mapping Service) handles the assignment of local accounts and credentials. LCMAPS is the enforcement point of the authorization decision made by the LCAS.

4.5 VOMS Awareness

The VOMS provides trustable assertions relating group memberships, roles and capabilities to the owner of X509 certificates. It maintains a database of these relationships and a means to administer them. It also provides a web service through which these assertions can be obtained. These assertions can be obtained by an authorised resource or by the individual to whom the assertions belong.

In the First case, where the resource obtains these assertions and applies them, provides little more than today's pool account/LDAP system. The full advantages of the VOMS appears when

[2] A service implementation of LCAS is also being developed.

[3] Rather like the hosts.allow and hosts.deny file used by tcp wrappers.

the user is able to download assertions and hand them over to the requested service, thus providing a mechanism by which the user can assert their choice of group membership, role and capability from those granted to them. It is this ability that fully addresses the five drawbacks listed above.

5. VOMS Deployment

5.1 The Virtual Organization Membership Service. Overview

The Virtual Organization Membership Service (VOMS) has been developed in the framework of EDG and DataTAG collaborations to solve the problems of granting users authorisation to access the resources at VO level, providing support for group membership, roles and capabilities. The VOMS system consists of the following parts

- User server: receives requests form client and returns information about the user.
- User client: contacts the server presenting a user's certificate and obtain a list of groups, roles and capabilities[4] of the user.
- Administration client: used by VO administrator to add user, create new groups, change roles.
- Administration server: accept the request from the admin client and updates the database.

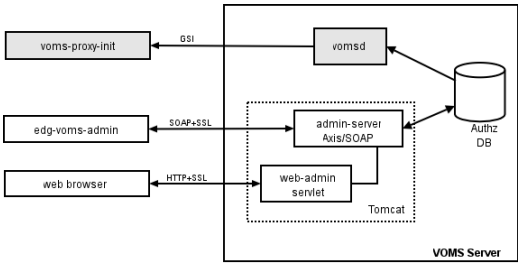

GridPP and NGS Grid environments in Manchester University

5.2 gLite VOMS (MySQL)

The recent release version 3.0 based on MySQL database is rolled out with the EGEE gLite 3.0 Middleware in Manchester HEP group to support several local VOs. The URI for the GridPP VOMS service is https://voms.gridpp.ac.uk . This contains an alias name for two Intel(R) Pentium(R) 4 CPU 3.20GHz with 1GB RAM and 160GB HDD. The one is front-end voms01.gridpp.ac.uk and the second one is backup voms02.gridpp.ac.uk. This name scheme allow to keep service running by moving main alias name to the backup server in case of any serious problems with the main one. The hosts are running under Linux SL3, kernel 2.4.21. Both hosts are now in the "production stage".

Currently there are 10 VOs defined on the GridPP server:

- gridpp – for the GRIDPP project, higher energy physics community
- ltwo – Teaching purposes within the London Tier 2
- t2k – Next Generation Long Baseline Neutrino Oscillation Experiment
- minos – Main Injector Neutrino Oscillation Search
- cedar – Combined e-Science Data Analysis Resource for high-energy physics
- gridcc – for the GRID project on remote control and monitoring of instrumentation such as distributed
- manmace – for the Manchester MACE engineers to run on the resources maintained at the Manchester EGEE Tier2 centre
- babar – for running babar experiment simulation and analysis on the EGEE/LCG grid
- pheno – dedicated to developing the phenomenological tools necessary to interpret the events produced by the LHC
- ralpp – VO for local tests in the RAL Particle Physics department

Another VO for Mathematicians in ScotGrid Tier2 is under process of approval.

NGS has chosen to support a global NGS VO for the moment.

The data reside on the same host in MySQL data base (version 4.11 server and client). There over 25 entries across the VOs in this data base now.

The machine has been very stable; it presents today an uptime over 120 days with one occasional reboot due to exceeded the maximum of process threads, although this may be explained by a modest number of entries and queries. The average number of connections for each VO was 1800 per day.

A similar set up hosts the NGS VOMS server. In this case there is currently one group, that describing the current list of NGS members.

Apart of exploiting version 3.0 we did run the 1.4 and 1.5 releases in the past and found the new version much more stable, many problems have been fixed.

[4]The Capability attribute although still valid in the VOMS attribute certificates is deprecated and no longer able to be produced through the VOMS interfaces

5.3 Known problems

There was a difficult situation with VOMS of versions 1.4/1.5 development and support. Bug samples:

1. VOMS -admin (on MySQL) doesn't list users with more than one Role (https://savannah.cern.ch/bugs/index.php?func=detailitem&item_id=14398);
2. Move the data base from older version of the VOMS based on MySQL to a newer one (e.g. from 1.3 to 1.4 or from 1.4 to 1.5), leads to hangs the VOMS and VOMS-admin tool. Sometimes it happens because of different database frame for different versions of VOMS;
3. Bad tomcat performance affecting VOMS access and gridmap file generation (https://savannah.cern.ch/bugs/?func=detailitem&item_id=14057).
4. VOMS-admin hangs due to tomcat "OutOfMemoryError" (https://savannah.cern.ch/bugs/?func=detailitem&item_id=16250).

The last problem with the tomcat is still presented in version 3.0 and has to be addressed in a new release.

6. Discussion

Our experience in deployment of gLite VOMS service shows that the current state of the middleware is acceptable to serve our needs. An increasing number of VO creation requests are being submitted to satisfy the most diverse necessities of local and regional groups.

7. Acknowledgemnts

This work was funded by the Particle Physics and Astronomy Research Council through their GridPP end e-Science programmes and by NGS. We would also like to thank other members of the EGEE security working group for providing much of the wider environment into which this work fits.

References

1. gLite VOMS Core User and Reference Guide: *https://edms.cern.ch/file/571991/1/voms-guide.pdf*
2. gLite VOMS Admin Tools User and Reference Guide: *https://edms.cern.ch/file/572406/1/user-guide.pdf*