ISTITUTO NAZIONALE DI FISICA NUCLEARE

CNAF

INFN/TC-02/022
9 Settembre 2002

Authorization with multiple Virtual Organizations

Vincenzo Ciaschini[1], Fabio Spataro[1]
[1)] *INFN, CNAF*

Abstract

In the grid world, users expect to be able to log into a grid-enabled machine without the need of having a personal account on the machine itself. To this end, users have organized themselves into Virtual Organizations (VOs) and access to the machines is (excepting differences in local policy, like for example the banning of a specific user) generally granted on the bases of VO membership. Up to now, users who are members of more than a Virtual Organization (VO) (or different subgroups of a single VO) have had no way to specify this while logging into the grid. This paper describes a possible solution to this problem.

Published by **SIS-Pubblicazioni**
Laboratori Nazionali di Frascati

1 The current situation

VO, or Virtual Organizations, are a way to bind together a large and constantly changing number of users working on a single project and sharing a set of resources.

When a user logs into a CE, he passes the CE the proxy certificate that he has generated with `grid-proxy-init`, and then the CE checks if the user is registered within a VO.

Right now, the VO is in fact a LDAP server which contains the subjects of the personal certificates of all its members, and so checking for membership in a VO means checking if the subject of the proxy certificates is one of those known to a VO. This check is done by looking for the subject of the proxy into the `grid-mapfile`, that is a file generated daily by the `mkgridmap` program and containing the subjects of all the users of all the VOs that the CE is configured to accept, in the order specified into its configuration file, `mkgridmap.conf`.

2 The problem

Since a user can belong to an arbitrary number of VOs, the approach detailed presents a problem: even if a user is a member of more than a single VO, for the CE he will always be identified as a member of only one VO, the one which appears first in `mkgridmap.conf`. Furthermore, there is no way for a user to override this default, to explicitly declare himself a member of a specific VO. While this issue may seem unimportant, it is actually fundamental for issues like accounting and billing, e.g. 'Who should I bill for the use of my farm?'.

3 A solution

In accordance with the Authorization WG, the authors have devised and implemented the following solution.

We have decided to modify `grid-proxy-init`, `mkgridmap` and `grid-mapfile` to allow the user to specify a VO during the creation of his proxy cartificate. The modification of `grid-proxy-init` required in turn some modifications to the globus library `libglobus_ssl_utils`.

However, while doing it we strived to mantain the highest possible compatibility with the already installed software, keeping in mind that this is a temporary fix to tide us over until the new authorization methods (non `grid-mapfile`-based) will be available, and so it wasn't worth to actually overhaul Globus to mantain perfect compatibility.

3.1 grid-proxy-init

3.1.1 Modifications

When the user executes grid-proxy-init to create a new proxy certificate, the subject of the new certificate is the same subject of the original one, with an extra CN=proxy (or CN=limited proxy[1]) added to the end. When this certificate reaches the gatekeeper, it is recognized as a proxy solely on the base of this relation between subjects, and the the subject of the proxy, stripped of the CN=... field at the end is matched against the grid-mapfile to decide if the user is authorized to log into the farm.

To mantain compatibility with the installed software base, we decided to continue to use this subject–grid-mapfile system to authorize users, and this meant that VO membership had to be coded in some way into the subject of the proxy certificate.

Our choice was to add a new field, D=<voname> right before CN=proxy during proxy creation, and consequently to change grid-mapfile in such a way that the subjects contain this field at the end. The fields are added by mkgrimap during the creation of the file.

3.1.2 Usage

The grid-proxy-init command now has two new options: '-vo <voname>' lets a user specify the VO he intends to use, and '-novo' and in this case it explicitly refuses to specify a VO and creates an old-style certificate. In case these options are both specified or are specified multiple times, only the last of them is considered. Furthermore, since many applications call grid-proxy-init internally, there is a new environment variable, VO, that, if defined, contains the default VO that grid-proxy-init uses if not overrided by different options. It should be noted that no difference is made between an undefined variable and an empty variable.

The table (1) exemplifies the relationships between VO, -vo and -novo.

3.1.3 Compatibility

It is immediately evident that a proxy of this new type is not compatible with versions of Globus that predates this change. However, since it is always possible to generate an old-style certificate, this hasn't been considered a problem.

[1] For here on, we will always write only CN=proxy; however, everything is still valid if it is substituted with CN=limited proxy.

Table 1: Relationships betwen options and variable in `grid-proxy-init`

VO=atlas	-vo cms	-novo	resulting vo
Yes	Yes	No	cms
No	Yes	No	cms
Yes	No	No	atlas
No	No	No	none
Yes	No	Yes	none
No	No	Yes	none
No	No	No	none

3.1.4 Changes to libglobus_ssl_utils

To implement this changes, It has been necessary to modify libglobus_ssl_utils, because this library contains the code that both creates the subject of a proxy certificate and recognizes a certificate as such. Without modifications, our new certificates won't be recognized as valid proxies, and in fact couldn't even be created. This means that a new version of this library must be installed in all systems who have Globus installed. However, since this can be done in a fully automated way using LCFG, the datagrid farm installation and maintenance tool, it hasn't been considered a big deal.

3.2 `mkgridmap`

3.2.1 Modifications

Starting from `edg-mkgridmap` version 1.0.7 we have chosen to associate the members of a VO to a pool of accounts characterized by a username created with a fixed prefix (the VO name), and a progressive number. Supposing that INFN's VO is named `infngrid`, this result can be obtained by writing in `mkgridmap.conf` lines like:

```
group ldap://grid-vo.cnaf.infn.it/ou=testbed1,o=infn,c=it .infngrid
```

that will result in grid-mapfile lines like:

```
"/C=IT/O=INFN/OU=Personal Certificate/CN=..." .infngrid
```

This means that all the members of INFN's VO will be associated to the account pool that has `infngrid` as username prefix. The effective association happens during job submission using McNab's patch.

Recently we have introduced the support for multiple VOs. If in the configuration file we have lines like:

```
group ldap://grid-vo.cnaf.infn.it/ou=testbed1,o=infn,c=it .infngrid
group ldap://grid-vo.nikhef.nl/ou=testbed1,o=alice,dc=eu-datagrid,dc=org .alice
group ldap://grid-vo.nikhef.nl/ou=testbed1,o=atlas,dc=eu-datagrid,dc=org .atlas
```

We obtain grid-mapfile lines like:

```
"/C=IT/O=INFN/OU=Personal Certificate/CN=..." .infngrid
"/C=IT/O=INFN/OU=Personal Certificate/CN=.../D=alice" .alice
"/C=IT/O=INFN/OU=Personal Certificate/CN=.../D=atlas" .atlas
"/C=IT/O=INFN/OU=Personal Certificate/CN=.../D=infngrid" .infngrid
```

The subject of the certificate is written once as it would have been written by version 1.0.7 of `mkgridmap`; the following ones are specific for the different VOs to whom the user belongs. For every VO the prefix is duplicated in both the `D=` extension and in the account pool name.

3.2.2 Compatibility

There should be no compatibility problem with previous releases.

4 Conclusions

The system described in these notes has been already developed, and is now undergoing an extensive internal test.

It seems to be an effective way to solve the problems detailed in this document, and we hope to be able distribute it soon. It should however be noted that the widespread implementation of this mechanism would require that this change become a part of the whole authorization mechanism, since inserting the VO membership information is only a part of the whole process, the other part being the gatekeeper and broker, who should be able not only to accept it, but also to interpret it and to act accordingly.