

# Service Incident Report for PIC Computing Service failure after the Scheduled Downtime on July 20

Incident Start: 20th July 2010 18:00 UTC

Incident End: 20th July 2010 around 20:30 UTC

## Description

After the Scheduled Downtime intervention on July 20th, no computing jobs were able to start at Tier-1 for few hours because the user mapping did not work due to a problematic migration of gridmapdir directory on the CEs.

## Impact

Farm was empty of jobs during 2-3 hours after the Scheduled Downtime finished.

## Time line of the incident

- 2010/07/20 17:00 gridmapdir mount point on CEs changes from nfssrv01.pic.es to nfssrv02.pic.es
- 2010/07/20 18:00 Some sgm jobs enter to run in farm (some jobs that were kept in CE).
- 2010/07/20 19:00 First SAM test failed
- 2010/07/20 20:00 Change rolled back: all CEs mount gridmapdir again from old server (nfssrv01.pic.es)

## Analysis

We saw errors in CE's messages like:

[...]

Jul 20 19:02:03 ce05 GRAM gatekeeper[2889]: Got connection 131.169.223.51 at Tue Jul 20 19:02:03 2010

Jul 20 19:02:04 ce05 GRAM gatekeeper[2889]: Authenticated globus user:

/O=GermanGrid/OU=LMU/CN=Johannes Elmsheuser

Jul 20 19:02:04 ce05 GRAM gatekeeper[2889]: globus\_gss\_assist\_gridmap() failed authorization.

globus\_gss\_assist: Error invoking callout globus\_callout\_module: The callout returned an error an unknown error occurred

[...]

Which indicates a problem with gridmapdir. Then, after a quick ls on gridmapdir:

```
#ls -lsrti gridmapdir
```

[...]

```
20048 0 -rw-r--r-- 1 root root 0 Jul 20 10:44 cmprd007
```

```
36722 0 -rw-r--r-- 1 root root 0 Jul 20 10:44 %2fd%3dorg%2fd%3ddoegrids%2fou%3dservices%2fcn%3duscmspilot%2dtest%2fcmssrv38%2efnal%2egov:cmprd
```

```
20044 0 -rw-r--r-- 1 root root 0 Jul 20 10:44 cmprd003
```

```
34363 0 -rw-r--r-- 1 root root 0 Jul 20 10:44 %2fd%3dch%2fd%3dcern%2fou%3dorganic%20units%2fou%3dusers%2fcn%3dasciaba%2fcn%3d430796%2fcn%3dandrea%20sciaba:cmprd
```

[...]

accounts are not correctly set.

Normally, when a request for access is received, the Globus gatekeeper (or another Grid aware server such as sshd-gsi) attempts to find the corresponding local unix username in the grid-mapfile. This file lists pairs of certificate subjects (eg "/O=Grid/O=Globus/OU=hep.man.ac.uk/CN=Andrew McNab") and usernames (eg mcnab.) If no line is found for the current subject, access is denied. For each username in the pool, whether leased or not, an empty file exists in this directory with the same name as the user. When a username is leased, a hard link is made to it, with the same name as the subject. Subject names are stored using the "URL encoding" method - letters and numbers are unchanged, other characters are replaced by %HH where HH is their value in hexadecimal. (This is necessary since subjects include slashes and other unfriendly characters.) To find which username corresponds to which subject and vice versa, the directory can be searched for links pointing to the same inode.

In Gridmapdir structure, each pool account should point to same inode as it's mapped userDN.

We do realize that we did a copy from gridmapdir to gridmapdir without keeping inode relationship.

## Follow up actions

- On future gridmapdir migrations we won't keep old gridmapdir mappings. It's not necessary.
- Review Change Control procedures and workflows to ensure only planned and tested changes are implemented in an SD.