

**Test of the DGAS Accounting System with a unified gatekeeper log at
INFN**

DRAFT

1.	Introduction	2
2.	Goals.....	2
3.	Preliminary Tests.....	3
4.	Further Tests.....	4
5.	Release and Deployment	5
6.	Preliminary results and issues	5
7.	Conclusions	5
8.	Appendix A	7
9.	Appendix B.....	9
10.	Appendix C.....	10

1. Introduction

This document describes the activity carried on in the Italian ROC to test the DGAS Accounting System together with a patch for the LCG Computing Element (patch 898) that implements a unified log for the gatekeeper on the LCG Computing Element.

The InfnGrid middleware release, based on the EGEE gLite middleware release, is currently deployed on the entire Italian production grid (up to now there are about 40 sites).

The InfnGrid release contains two additional components with respect to the EGEE gLite release: the support for DAG jobs and the DGAS accounting system.

DGAS has been used by the Italian ROC to collect accounting information for the grid resources of the Italian production grid for more than one year.

Unfortunately the usage records were incomplete, i.e. they contained all the grid-related information, such as the grid job id, the user's DN and the user's FQAN, only for those jobs submitted through DGAS-aware Resource Brokers, while only a reduced set of information (i.e. those belonging to the local LRMS) could be retrieved for the jobs dispatched by the standard RBs.

Since the number of DGAS-aware Resource Brokers is small (they have been deployed only in Italy) for most jobs the complete DGAS functionalities cannot be used.

Patch 898 (<http://savannah.cern.ch/patch/?898>) to the gatekeeper assures the presence of complete grid information for all jobs in the gatekeeper log files, independently from the Resource Broker they have been submitted through.

This is the reason why the Italian ROC had a strong interest in the patch and therefore we started testing a new DGAS version that supports it.

The combination of the two mentioned elements (patch and new DGAS version) eliminates the described accounting limitation allowing the use of the complete DGAS functionalities for all jobs.

At the moment patch 898 is deployed on all the Italian production Grid resources.

2. Goals

The main goals of this test activity were to verify that:

1. The gatekeeper patch 898 could be applied in a running production farm without the need of putting the sites in downtime status.
2. The new gatekeeper log, containing the grid-related information, was correctly produced and that DGAS sensors were able to parse it and to produce a complete usage record to be inserted in the Home Location Register (HLR).

Moreover, as a second step in testing the DGAS system performances, after the deployment in the production Grid, the LRMS accounting data belonging to a sample of sites were analyzed using a set of independent simple scripts and the results compared with those provided by the site HLRs.

3. Preliminary Tests

In order to verify that the new gatekeeper patch could be “hot-plugged” into a running production farm, a test environment with two LCG CEs has been used.

The first one was located in Torino and used PBS, while the second one, running LSF, was installed in Padova.

The test consisted in the following steps:

- 1) a given number of jobs were submitted to both the CEs, in order to have the maximum number of jobs running and also some queued jobs;
- 2) the new gatekeeper packages were installed and the services reconfigured while the farm was in this state;
- 3) we verified that neither the running jobs nor the queued ones were affected by the operation and all the jobs completed normally. As expected, the unified gatekeeper log was not written for these jobs;
- 4) a new bunch of jobs was submitted to the patched CEs, and we verified that the unified log contained the corresponding entries with the correct values (Grid job ID, submitter credentials, LRMS job id etc..).

As a second step, in order to verify the capability of the DGAS sensors to use the new gatekeeper log, we prepared a small testbed including the two aforementioned CEs and two Resource Brokers; the first one was an Italian DGAS-aware RB which is currently used for site certification purpose, and the second was a standard gLite RB. Bunch of jobs (100; 200; 500; 1000) were submitted to both the PBS and LSF CEs with the patch 898 and the new DGAS sensors, using both the standard and DGAS Resource Brokers.

The result was that all the jobs which terminated normally (few jobs were aborted because of the poor performance of one of the two CEs) were correctly accounted with all the correct information as in both the local LRMS and in the unified gatekeeper logs, regardless of the RB used.

Since the first test phase was successful, we decided to proceed applying the patch to two CEs on the production grid, in order to verify the system on real production farms.

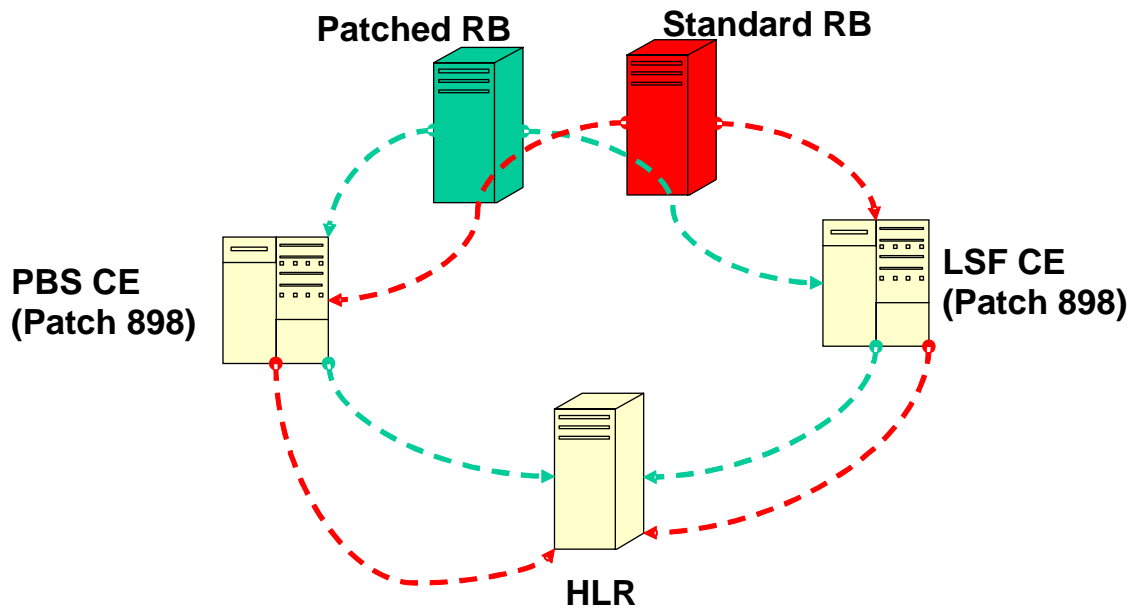


Figure 1: Set up of the small testbed used for the DGAS + patch 898 tests.

4. Further Tests

In order to verify the functionality of the DGAS system + patch 898 on a real production environment, we installed the new packages on two production grid sites and kept them monitored for some time, before releasing the patches to all the sites.

These CEs were located in Torino and Padova; the first one used PBS and the second one LSF as LRMS, and both had a rather standard configuration, with one master LRMS node and one CE as a gatekeeper.

We verified that the DGAS sensors correctly parsed the gatekeeper log to integrate the grid-related job information into its usage records.

Additionally we verified that the so-called out-of-band jobs (jobs that have been submitted with qsub, etc. without passing through the gatekeeper) can be accounted as well, although their usage records contain only local information.

In this case the user VO can still be determined from the pool account or, in the next DGAS version that will be deployed in a few weeks, also by mapping local users/groups to grid VOs.

Since also this test was successful the new packages were considered ready to be deployed on the Italian production Grid.

5. Release and Deployment

In order to deploy the new rpms on the italian production grid in an automatic way the INFN GRID Release Team carried out the following tasks:

- build of new metapackages for some profiles that include dependencies on the packages mentioned in the patch and on the new version of the **DGAS** packages
- build of a new ig-yaim, that includes:
 - new configuration functions: config_dgas_patch, config_dgas_ce and config_dgas_glite_ce (see Appendix A)
 - a new site-info.def file containing new variables to configure various dgas services
 - a new node-info.def file containing a new profile for the HLR (Home Location Register) server, and new functions for other profiles (lcg-CE, glite-CE, UI)

Instructions for the deployment of the new rpms on the production grid have been prepared and provided to the site managers

(<https://grid-it.cnaf.infn.it/checklist/modules/dokuwiki/doku.php?id=rel:20061219> – only the italian version is available at the moment); they are also summarized in Appendix B.

6. Preliminary results and issues

After the deployment of patch 898 and of the new DGAS release on the italian production grid, further checks have been carried out at some sites (Tier-1 and Tier2s) in order to validate the INFN GRID release update.

Cross-checks have been carried out by comparing accounting data stored into the HLRs and the corresponding data extracted by the LRMS log files by means of appropriate scripts.

Checks have not been completed yet, but the preliminary results are really good.

In particular for all the Tier-2s farms with Torque+Maui and a unique Computing Element no remarkable differences have been found. In some cases there are very small differences (some units over thousands jobs); they are currently being investigated but they probably depend on some particular local configurations.

Checks have not been trivial for sites with LSF since the information stored in the log files, in particular the timestamp fields, may be confusing, in particular for some jobs (e.g. jobs canceled before being executed).

In addition some LSF farms have been configured with multiple CEs and this requires a special DGAS configuration in order to correctly account for local jobs. As for the Torque+Maui sites the preliminary results are good but the cross-check activity has not been completed yet.

7. Conclusions

Although the validation of the sites where patch 898 + DGAS have been deployed has not been completed yet, the preliminary results are very good.
This document will be updated as soon as the activity will be completed.

8. Appendix A

- **config_dgas_patch:**
https://forge.cnaf.infn.it/plugins/scmsvn/viewcvs.php/tags/ig-yaim-3_0_0-31/ig-yaim/func_added/config_dgas_patch?rev=1829&root=igrelease&view=markup
-
- **config_dgas_ce:**
https://forge.cnaf.infn.it/plugins/scmsvn/viewcvs.php/tags/ig-yaim-3_0_0-31/ig-yaim/func_added/config_dgas_ce?rev=1829&root=igrelease&view=markup
-
- **config_dgas_glite_ce:**
https://forge.cnaf.infn.it/plugins/scmsvn/viewcvs.php/tags/ig-yaim-3_0_0-31/ig-yaim/func_added/config_dgas_glite_ce?rev=1829&root=igrelease&view=markup
-
- **config_dgas_hlr:** https://forge.cnaf.infn.it/plugins/scmsvn/viewcvs.php/tags/ig-yaim-3_0_0-31/ig-yaim/func_added/config_dgas_hlr?rev=1829&root=igrelease&view=markup
-
- **site-info.def:**

```
# INFN-GRID: DGAS - HLR resource
#           See at http://grid-
it.cnaf.infn.it/index.php?deployment
#           Only the hostname and the port is needed
#           Ex.: HLR_RESOURCE="prod-hlr-01.pd.infn.it:56568"
HLR_RESOURCE="please put your site RESOURCE HLR"
# INFN-GRID: DGAS - HLR user. Leave default settings
HLR_USER="hlr-user-01.cnaf.infn.it:56568"
# INFN-GRID: write "yes" to send accounting data to the
HLR_USER server,
#           otherwise the data will be sent only to the
HLR_RESOURCE server
#           For now leave default settings
HLR_USER_ENABLED="no"
# INFN-GRID: DGAS - HLR resource server configuration
parameters
#           set them only if you are installing a resource HLR
server
HLR_SQL_USER=root                # default
HLR_SQL_PASSWORD=set_this_to_a_good_password
HLR_SQL_DBNAME=hlr              # default
HLR_TMP_SQL_DBNAME=hlr_tmp      # default

# INFN-GRID: DGAS - urCollector configuration
#           The following option defines whether urCollector
shall consider
#           "grid": ONLY grid jobs;
#           "local": ONLY local jobs;
#           "all": all jobs;
#           (Default is "all")
# ATTENTION: For sites with 1 MASTER node for batch system,
shared by more CEs,
#           you must set this variable to "all" on only ONE CE
(the best one),
#           and to "grid" on all the others
HLR_JOBS_T0_PROCESS="all"
```

```
# INFN-GRID: DGAS - HLR resource
#           Conversion Log Database parameters
CONV_DB_USER=root
CONV_DB_PASSWORD=set_this_to_a_good_password
```

- **node-info.def:**

```
CE_FUNCTIONS="${BASE1_FUNCTIONS} ${BASE2_FUNCTIONS}
.....
config_dgas_patch
config_dgas_ce"
```

```
gliteCE_FUNCTIONS="
.....
config_dgas_glite_ce"
```

```
# INFN-GRID: HLR Server (Home Location Register, DGAS
accounting server)
HLR_FUNCTIONS="
config_ntp
config_java
config_rgma_client
config_glite_env
config_ldconf
config_dgas_hlr
config_apel_rgma
config_fmon_client"
```


9. Appendix B

Instructions provided to site managers:

- update local install-server repository (yam server, if exists)
- upgrade ig-yaim (ig-yaim-3.0.0-31.noarch.rpm)
- **update lcg-CE (LSF, torque) profile:**
 - back-up configuration files (dgas_gianduia.conf, dgas_atmClient.conf, dgas_pushd.conf)
 - remove old **DGAS** packages: edg-wl-dgas-hlr-ATMClient_slc_3_0_3-2.4.1-0, edg-wl-dgas-hlr-pingClient_slc_3_0_3-2.4.1-0
 - update:

```
# apt-get update
# apt-get -f install
# apt-get --assume-yes dist-upgrade
# apt-get install edg_gatekeeper-gcc32dbg_pgm
```

(the last step is needed because of a packaging problem for edg_gatekeeper gcc32dbg_pgm)

- configure:
 - modify site-info.def:

```
# INFN-GRID - urcollector configuration
# The following option defines whether urcollector shall
consider
# "grid": ONLY grid jobs;
# "local": ONLY local jobs;
# "all": all jobs;
# Default is "all"
# ATTENTION: For sites with 1 MASTER node (batch system),
shared by more CEs,
# you must set this variable to "all" on only ONE CE (the
best one), and to "grid" on all the others
HLR_JOBS_TO_PROCESS="all"
```

- run the following configuration functions:

```
# <path>/ig_run_function <your_UPDATED_site_info>
config_dgas_patch
# <path>/ig_run_function <your_UPDATED_site_info>
config_globus
# <path>/ig_run_function <your_UPDATED_site_info>
config_dgas_ce
```

- **update the WNs:**

```
# apt-get update
# apt-get --assume-yes dist-upgrade
# rpm -e edg-wl-dgas-wn-ceServiceClient_slc_3_0_3
```

10. Appendix C

Submission of jobs, verification of creation of gatekeeper log file, exemple of content of log file:

```
[caifti@prod-ui-01 caifti]$ ./storm_dgas 10 simple_test.jdl
test_cert12_patch_10_2_griditcertrb
* submitting 10 simple_test.jdl in sequences of 5
[caifti@prod-ui-01 caifti]$ cat test_cert12_patch_10_2_griditcertrb
###Submitted Job Ids###
https://gridit-cert-rb.cnaf.infn.it:9000/ubbwnGsFH5wKOzOMVvZxCA
.....
https://gridit-cert-rb.cnaf.infn.it:9000/0UloHRP29R-I0xA4aHuj7g
.....
```

```
[root@prod-ce-01 root]# bjobs -uall |grep cert-12
890747 dteam00 RUN cert cert-12 prod-wn-014 *ort.NS975
Dec 1 17:54
.....
890759 dteam00 PEND cert cert-12 *ort.F1131
Dec 1 17:54
```

```
[root@cert-12 root]# less /opt/edg/var/gatekeeper/grid-
jobmap_20061201
"localUser=18118" "userDN=/C=IT/O=INFN/OU=Personal
Certificate/L=Padova/CN=Cristina Aiftimiei"
"userFQAN=/dteam/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/Role=NULL
/Capability=NULL"
"userFQAN=/dteam/see/RO/NIHAM/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/INFN-PADOVA/Role=NULL/Capability=NULL"
"jobID=https://gridit-cert-rb.cnaf
.infn.it:9000/tAPHYkVqdsMBfxlR057bXQ" "ceID=cert-
12.pd.infn.it:2119/jobmanager-lcglsf-cert" "lrmsID=890747"
"timestamp=2006-12-01 16:54:23"
"localUser=18118" "userDN=/C=IT/O=INFN/OU=Personal
Certificate/L=Padova/CN=Cristina Aiftimiei"
"userFQAN=/dteam/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/Role=NULL
/Capability=NULL"
"userFQAN=/dteam/see/RO/NIHAM/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/INFN-PADOVA/Role=NULL/Capability=NULL"
"jobID=https://gridit-cert-rb.cnaf
.infn.it:9000/ubbwnGsFH5wKOzOMVvZxCA" "ceID=cert-
12.pd.infn.it:2119/jobmanager-lcglsf-cert" "lrmsID=890748"
"timestamp=2006-12-01 16:54:23"
"localUser=18118" "userDN=/C=IT/O=INFN/OU=Personal
Certificate/L=Padova/CN=Cristina Aiftimiei"
"userFQAN=/dteam/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/Role=NULL
/Capability=NULL"
"userFQAN=/dteam/see/RO/NIHAM/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/INFN-PADOVA/Role=NULL/Capability=NULL"
"jobID=https://gridit-cert-rb.cnaf
.infn.it:9000/--Vh72LoPenyn2AjpFKzyw" "ceID=cert-
12.pd.infn.it:2119/jobmanager-lcglsf-cert" "lrmsID=890751"
"timestamp=2006-12-01 16:54:29"
.....
"localUser=18118" "userDN=/C=IT/O=INFN/OU=Personal
Certificate/L=Padova/CN=Cristina Aiftimiei"
"userFQAN=/dteam/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/Role=NULL
/Capability=NULL"
"userFQAN=/dteam/see/RO/NIHAM/Role=NULL/Capability=NULL"
"userFQAN=/dteam/italy/INFN-PADOVA/Role=NULL/Capability=NULL"
"jobID=https://gridit-cert-rb.cnaf
```

```
.inf.n.it:9000/NBeAQmxx2M5Qxz1fjZADqA" "ceID=cert-  
12.pd.inf.n.it:2119/jobmanager-lcglsf-cert" "lrmsID=890759"  
"timestamp=2006-12-01 16:54:39"
```

Submission of 10 jobs, verification of creation of gatekeeper log file:

```
[caifti@prod-ui-01 caifti]$ for i in `seq 1 10`; do subm -o  
test_cert12_patch_10_5_griditcertrb simple_test.jdl; done  
[caifti@prod-ui-01 caifti]$ cat test_cert12_patch_10_5_griditcertrb  
###Submitted Job Ids###  
https://gridit-cert-rb.cnaf.infn.it:9000/RP5GgO4b0EYDwB3vjMHfrA  
.....  
https://gridit-cert-rb.cnaf.infn.it:9000/\_e0dQMFRZ1VAPJwqfsonmA
```

```
[root@prod-ce-01]# bjobs -uall |grep cert-12  
896672 dteam00 RUN cert cert-12 prod-wn-035 *ort.e1399  
Dec 5 14:53  
.....  
896681 dteam00 PEND cert cert-12 *ort.q1697  
Dec 5 14:54
```

```
[root@cert-12 dgas_patch]# ll /opt/edg/var/gatekeeper/grid-  
jobmap_20061205  
-rw-r--r-- 1 edguser dteam 5810 Dec 5 14:54  
/opt/edg/var/gatekeeper/grid-jobmap_20061205
```