



**INTERFACING TO GLITE  
TOOLS AND OTHER  
OPERATIONAL ISSUES**

S. Bagnasco (ALICE/INFN)

bagnasco@to.infn.it

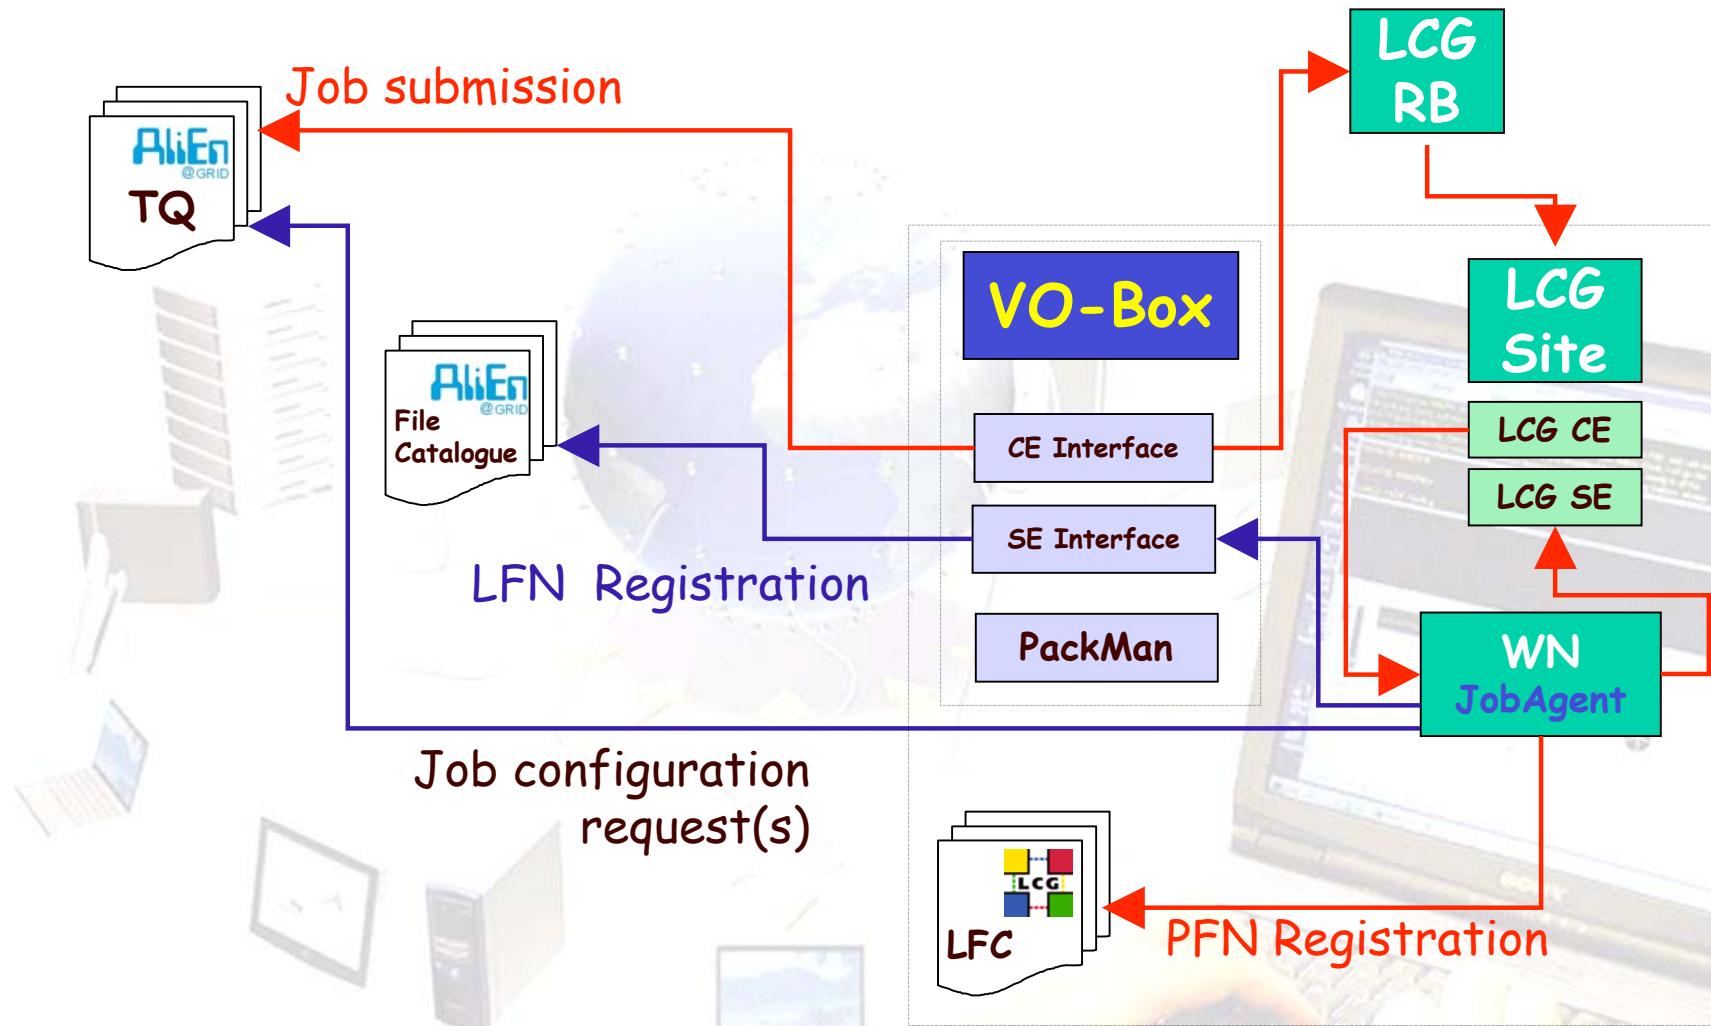
# OUTLINE OF THE TALK

VASTLY INCOMPLETE

- **VO-Box configuration**
  - Where is it kept?
  - What to configure?
- **Migration to gLite tools**
  - VO-Box & WMS tests at CERN
  - Switching configuration
- **Other operational issues**
  - Some new features of v2-12
  - RB Failover mechanism
  - Another go at proxy management
  - Hints for debugging: where are the logs?
- **Questions!**



# THE ALICE VO BOX



# WHAT'S IN A VO BOX

- **Standard LCG User Interface**
  - Either LCG or gLite (but wait for details later)
  - Restricted login access: only one local user, logs in via gsisssh
- **Proxy Renewal mechanism**
  - Vobox-proxyrenewd (more detail later)
- **GRIS** to publish the VO-Box itself to the BDII
- **NFS mount of the experiment's software area**
  - `$VO_ALICE_SW_DIR`
- **ALICE interfaces and services**

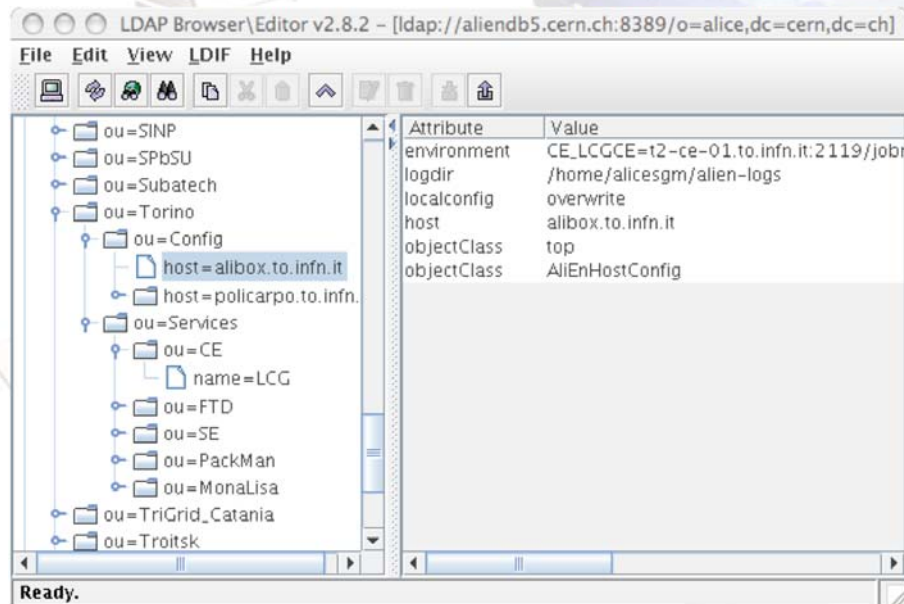
# VO BOX HIGHLIGHTS

- Implement as much as possible "thin" interface services
  - To (stable) LCG standard services
  - Be "good citizens" of the Grid - no backdoors
- Use the VO-Box manager's certificate
  - All jobs in a site share the same LCG user
  - As requested by some sites, an enhancement for security: glexec (to de-anonymize the actual job owner on Worker Nodes)
- Service interfaces on the VO-Box:
  - Job Submission
  - Storage
  - Data transfer
- Non-LCG services:
  - Package Manager
  - Cluster Monitor
  - MonALISA monitoring



# ALIEN CONFIGURATION 1

- LDAP database
  - `ldap://aliendb5.cern.ch:8389/o=alice,dc=cern,dc=ch`
- Local configuration file
  - On the VO-Box
  - `~alicesgm/.alien/alice.conf`



```
LOG_DIR /home/alicesgm/alien-logs
<SE LCG>
host vobox.ct.infn.it
port 8092
mss File
savedir /home/alicesgm/ALIEN_SE
ioDaemons xrootd:port=5443
</SE>
<CE LCG>
Type LCG
HOST vobox.ct.infn.it
TTL 100000
</CE>
CE_LCGCE grid012.ct.infn.it:2119/jobmanager-lcglsf-alice
CE_LCG_GFAL_INFOSYS ldap://egee-bdii.cnaf.infn.it:2170
CE_EDG_WL_UI_CONF /home/alicesgm/.alien/edg_wl_ui.cnaf.conf
CE_SITE_BDII ldap://grid012.ct.infn.it:2135/mds-vo-name=local,o=grid
<PACKMAN LCG>
host vobox.ct.infn.it
port 9991
installDir /opt/exp_soft/alice/packages
</PACKMAN>
```



# ALIEN CONFIGURATION 2

- Priority mode set by admins in central LDAP:
- Normal mode
  - LDAP has priority
  - Local file is ignored
- “Overwrite” mode
  - Local file has priority
- “Add” mode
  - Local file can only define objects not in central LDAP
  - “Blocks” are ignored



# SOME MORE CONFIGURATION

## Environment

### ■ Site-wide

- `$VO_ALICE_SW_DIR/alien/.Environment`
- Seen also from WNs

### ■ On the VO-Box

- `~alicesgm/.alien/Environment`
- Has priority over the site-wide one
- Most important entry: `ALIEN_USER`

### ■ From LDAP (for the services)

- 'environment' attribute in `host=alibox.to.infn.it, ou=Config, ou=Torino, ou=Sites, o=alice, dc=cern, dc=ch`
- Host-wide and service-wide (?)
- Allows LCG-related config items to be kept in LDAP without changing the schema all the time.



# EVEN MORE CONFIGURATION

## Services startup

- `~alicesgm/.alien/etc/aliend/startup.conf`  
`ALIEN_ORGANISATIONS="ALICE"`
- `~alicesgm/.alien/etc/aliend/ALICE/startup.conf`  
`AliEnUser=alicesgm`  
`AliEnServices="Monitor SE CE PackMan MonaLisa"`  
`AliEnCommand="$VO_ALICE_SW_DIR/alice/alien/scripts/lcg/lcgAlien.sh"`

# LCG RELATED CONFIG ITEMS

- **CE\_LCGCE**
  - List of CEs for submission
    - They end up in OR in the JDL, the RB decides
    - This is the only mandatory one
- **CE\_SITE\_BDII**
  - Where to ask for informations about the site
    - Default is using the LCG CE
- **CE\_EDG\_WL\_UI\_CONF**
  - RB config file to use
    - Passed to submit command via --config-vo
    - If not defined, use the default (/opt/edg/etc/alice/edg\_wl\_ui.conf)
    - Could use CE\_SUBMITARG instead



# WMS: GLITE VS. LCG

- **Command names changed**
  - From `edg-job-*` to `glite-job-*`
- All other syntax we use is unchanged
- Everything should be backwards compatible, but...
  - Some `edg-job-*` commands are very slow if used with gLite WMS

▪ ...



# GLITE VO BOX AND WMS TESTS

- SLC4 test required by the GD Group for all 4 LHC Experiments
- Second VO-Box with gLite and SLC4 at CERN
  - [Voalice03.cern.ch](http://Voalice03.cern.ch)
- PPS gLite-WMS, LCG queue with SLC4 WNs behind
  - lxb2006 first with 10 nodes behind
  - ce110 afterwards, with more than 100 WNs, much more stable and reliable

Patricia Mendez Lorenzo, CERN



# GLITE VO BOX AND WMS TESTS

- Everything went well!
  - All available CPUs covered in few minutes by ALICE Jobs which were running for 20 hours
  - Small compatibility issue between SLC3 and SLC4 in AliRoot (already fixed in new AliRoot release)
  - All AliEn commands ran perfectly, in general the test has succeeded

Patricia Mendez Lorenzo, CERN





# SWITCHING TO GLITE COMMANDS

- Easy: change config items for the CE
  - From the edg-job-\* commands to glite-job-\*
- **CE\_SUBMITCMD + CE\_SUMBITARG**
  - Default: edg-job-submit
- **CE\_STATUSCMD + CE\_STATUSARG**
  - Default: edg-job-status
- **CE\_KILLCMD + CE\_KILLARG**
  - Default: edg-job-cancel
- **CE\_MATCHCMD + CE\_MATCHARG**
  - Default: empty (may be edg-job-list-match)

# SWITCHING TO GLITE COMMANDS

OK, it's easy, however...

- We are not yet ready for migration!
  - gLite system still under test at CERN
  - VO-Box (UI) and WMS (RB & CE)
- **DO NOT UPGRADE RBs/CEs!**

# RB FAILOVER CONFIGURATION

- New in v2-12, still to be configured in most sites
- A drop-in replacement for edg-job-submit
  - But not yet for gLite!
  - Slow failure, no memory (will be improved shortly)

- To configure:

CE\_SUBMITCMD

```
'$VO_ALICE_SW_DIR/alien/bin/alien-perl  
$VO_ALICE_SW_DIR/alien/scripts/lcg/failoverSubmit.pl'
```

CE\_SUBMITARG

```
'-rb egee-rb-01.cnaf.infn.it,rb105.cern.ch,grid014.ct.infn.it'
```



# THE INFAMOUS PROXIES 1

**Q:** "We don't understand why it's not working"

**A:** "So it's a proxy problem"

- The **login proxy**, which is used by the manager to login (via `gssh`) to the VO-Box
- The **myProxy** registered on the MyProxy server (`myproxy.cern.ch`)
- The **machine proxy**, which is used by the VO-Box to authenticate to the myproxy server
- The **user proxy**, which is used by AliEn to submit jobs to LCG





# THE INFAMOUS PROXIES 2

1. The user generates a **login proxy** on some LCG UI, which he will use to `gssh` to the VO-Box
2. The user stores a long-lived (e.g. one month) **myProxy** on `myproxy.cern.ch`.
3. From the VO-Box, the user registers the login proxy to the VO-Box Proxy Renewal Service, thus generating, the **user proxy**.
4. The VO-Box **Proxy Renewal Service** keeps the user proxy alive by periodically getting a new 12-hours one from the server. To authenticate to the latter, it uses the root-owned **machine proxy**.
5. The AliEn CE running on the VO-Box uses the user proxy to submit jobs to the LCG RB. Just before submitting a bunch of job agents, the CE will itself renew the user proxy, with a longer duration.





# THE INFAMOUS PROXIES 3

- **Where are they?**
  - **The login proxy**
    - Something like `/tmp/x509up_p13506.fileh7rqRw.1`
    - Check where `$X509_USER_PROXY` points upon login
  - **The myProxy**
    - Create: `myproxy-init -s myproxy.cern.ch -d -n -t 48 -c 720`
    - Inspect: `myproxy-info -s myproxy.cern.ch`
  - **The machine proxy**
    - `/opt/vobox/alice/renewal-proxy.pem`
  - **The user proxy (there may be more than one...)**
    - Registered in `/opt/vobox/alice/_registerer_proxies.db`
    - Kept in `/opt/vobox/alice/proxy_repository`
    - Create: `vobox-proxy --vo alice [...] register`
    - Inspect: `vobox-proxy --vo alice --dn all query`



# THE INFAMOUS PROXIES 3

- How to startup services picking the right proxy
  - Check the right AliEn user
    - Grep ALIEN\_USER .alien/Environment
  - Check whether there's a corresponding user proxy in the database
    - Myproxy-info --dn <something> query
  - Pick the right proxy file
    - Export X509\_USER\_PROXY=/opt/vobox/alice/proxy\_repository/...
  - Start the services
    - \$VO\_ALICE\_SW\_DIR/alien/bin/alien StartMonitor
  - ...or use lcgAlien.sh, that does all of the above for you:
    - \$VO\_ALICE\_SW\_DIR/alien/scripts/lcg/lcgAlien.sh StartMonitor



# WHERE ARE THE LOGFILES?

- AliEn services logs
  - Under `~alicesgm/alien-logs/`
  - (or wherever the configuration item `LOG_DIR` points to)
  - `CE.log`, `ClusterMonitor.log`, `SE.log`, `PackMan.log`, `MonLisa.log`
- Job logs
  - `~alicesgm/alien-logs/proc/<jobId>`
  - Registered in the AliEn catalogue under `/proc/<user>/<jobId>/job-output`
- JobAgent LCG logs
  - LCG OutputSandbox is retrieved after the job is finished
  - Kept in `/tmp/AliEn/tmp/7200/...` for three days

**New in v2-12!**



# WHAT'S IN THE LOGFILES?

- **ClusterMonitor.log**

```
Jan 25 22:51:32 info Doing 'edg-job-status --noint https://egee-rb-01.cnaf.infn.it:9000/TjPKP9PosDUnJ2ggi-O7Fw'  
Jan 25 22:51:34 info Agent https://egee-rb-01.cnaf.infn.it:9000/9s4anIhsQubbA3XDatFmNA is dead!!  
Jan 25 22:51:34 info Creating a temporary file  
Jan 25 22:51:34 info Setting filename of edg-job-get-output.log  
Jan 25 22:51:34 info Setting ttl of 24 hours  
Jan 25 22:51:34 info Getting status from LCG for https://egee-rb-01.cnaf.infn.it:9000/9s4anIhsQubbA3XDatFmNA  
Jan 25 22:51:34 info Doing 'edg-job-status -noint "https://egee-rb-01.cnaf.infn.it:9000/9s4anIhsQubbA3XDatFmNA" | grep "Curren  
† Status:"'  
Jan 25 22:51:35 info Will retrieve OutputSandbox for Done(Success) job https://egee-rb-01.cnaf.infn.it:9000/9s4anIhsQubbA3XDat  
FmNA  
Jan 25 22:51:35 info Doing 'edg-job-get-output --noint --logfile /tmp/ALICE/tmp/86400/13538/edg-job-get-output.log --dir /tmp/  
ALICE/tmp/86400/13538 https://egee-rb-01.cnaf.infn.it:9000/9s4anIhsQubbA3XDatFmNA'  
Jan 25 22:51:39 info OK, will remove DB entry for https://egee-rb-01.cnaf.infn.it:9000/9s4anIhsQubbA3XDatFmNA  
Jan 25 22:51:45 info Command 1793675 changed to SAVING (and 0)  
Jan 25 22:51:46 info Updating done!!  
Jan 25 22:51:52 info Set Procinfo for 1793723: 04:46:16 17176 79.17 8.1 1606 344772 1194464 4 6 2000.181 26434.39 486075  
Jan 25 22:51:52 info Updating the jobagent  
Jan 25 22:52:04 info Set Procinfo for 1793690: 04:53:12 17592 83.77 9.9 2320 414832 1265236 4 6 2000.175 25750.25 486394  
Jan 25 22:52:04 info Updating the jobagent  
Jan 25 22:52:05 info Asking the Manager how many jobs we can run  
Jan 25 22:52:05 info We can run 50 and 150  
Jan 25 22:52:10 info Set Procinfo for 1794121: 03:51:46 13906 94.63 52.1 13579 535744 1419488 2 15 2395.958 32534.71 496662  
Cannot encode 'LOG_OBJECTS' element as 'hash'. Will be encoded as 'map' instead  
Jan 25 22:52:10 info Updating the jobagent  
Jan 25 22:52:10 info Set Procinfo for 1793675: 04:56:06 17766 89.34 4.5 21 820964 1794348 4 6 2000.204 26738.73 517190 1349086  
Jan 25 22:52:54 info Getting a request to put a file  
Jan 25 22:52:54 info File 1792731.out saved in /home/alicesgm/alien-logs/proc/1792731
```



# WHAT'S IN THE LOGFILES?

- **CE.log**

```
Jan 25 10:49:37 info Setting filename of dg-submit.29050.sh
Jan 25 10:49:37 info Creating a temporary file
Jan 25 10:49:37 info Setting filename of dg-submit.29050.jdl
Jan 25 10:49:37 info Renewing proxy for 172800 seconds
Jan 25 10:49:38 info Proxy timeleft is 172546 (162800)
Jan 25 10:49:38 info Submitting to LCG with ''.
Jan 25 10:49:38 info Creating a temporary file
Jan 25 10:49:38 info Setting filename of job-submit.1169718578.log
Jan 25 10:49:38 info Doing 'edg-job-submit --noint --nomsg --logfile /tmp/ALICE/tmp/7200/162460/job-submit.1169718578.log
/tmp/ALICE/tmp/7200/162460/dg-submit.29050.jdl'
Jan 25 10:49:47 info LCG JobID is https://rb115.cern.ch:9000/6woPIPjN-1lptE3wVAQFWw
Jan 25 10:49:47 info Submission took 10 sec.
Jan 25 10:49:47 info Inserting https://rb115.cern.ch:9000/6woPIPjN-1lptE3wVAQFWw in the list of agents
Jan 25 10:49:47 info Creating a temporary file
Jan 25 10:49:47 info Setting filename of dg-submit.29050.sh
Jan 25 10:49:47 info Creating a temporary file
Jan 25 10:49:47 info Setting filename of dg-submit.29050.jdl
Jan 25 10:49:47 info Renewing proxy for 172800 seconds
Jan 25 10:49:48 info Proxy timeleft is 43192 (162800)
A credential has been received for user /DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=pmendez/CN=477458/CN=Patricia Mendez Lorenzo
in /tmp/tmpfile.29050.
Jan 25 10:49:49 info Submitting to LCG with ''.
Jan 25 10:49:49 info Creating a temporary file
Jan 25 10:49:49 info Setting filename of job-submit.1169718589.log
Jan 25 10:49:57 info LCG JobID is https://rb115.cern.ch:9000/40wmw9CvKQWsRzOiYVWew
Jan 25 10:49:57 info Submission took 10 sec.
Jan 25 10:49:57 info Inserting https://rb115.cern.ch:9000/40wmw9CvKQWsRzOiYVWew in the list of agents
```



# OTHER USEFUL FILES

- **Local Databases:**

- **~alicesgm/alien-logs/CE.db/JOBAGENT**

- Queued and active jobs
- Holds the correspondence between jobId, agentId and lcgId

- **~alicesgm/alien-logs/CE.db/TOCLEANUP**

- lcgIds of finished jobs whose OutputSandbox was not yet downloaded

- **~alicesgm/alien-logs/CE.db/JOBID**

- All lcgIds of jobs submitted by this VO-Box



# DOCUMENTATION

- **Wiki pages** (work always in progress...)
  - <http://alien.cern.ch/twiki/bin/view/AliEn/HowTo...>
  - **HowToInstallLcgVoBox**
    - LCG VO-Box tests
    - How to prepare for VO-Box admin role
    - How to install and configure AliEn on the VO-Box
  - **HowToDebugLcgVoBox**
    - Some tips for debugging
    - Collection of common problems, with solutions
    - Even more "in progress" than the rest
  - **HowToManageVoBoxProxies**
    - All you wanted to know about the proxies used in the VO-Box, and possibly more.

**Please send feedback!**

