



Advanced School in High Performance Computing Tools for e-Science
5 March - 16 March 2007

gLite Job Management

Gilberto Diaz
Riccardo Di Meo



Job Submission


Egee submission

Resembles very closely the batch submission in a cluster.

- Can select resources automatically ("Grid like" submission).
- Resources can be selected manually (more "cluster oriented" submission)

Job Submission

After obtaining a digital certificate from a trusted Certification Authority, registering in a VO and obtaining an account on a User Interface, the user is ready to use the Grid.



Job Submission

The user logs in to the UI and creates a proxy certificate to authenticate himself in subsequent secure interactions.


User workstation



remote login



User Interface (UI)





Job Submission

- *edg-job-submit* [-r <res_id>] [-n <user_e-mail address>] [-c <config_file>] [-o <output_file>] **<job.jdl>**

- -r the job is submitted by the RB directly to the computing element identified by <res_id>
- -n an e-mail message containing basic information regarding the job (status and identifier) is sent to the specified <e-mail address> when the job enters one of the following status: READY, RUNNING, DONE or ABORTED



Job Submission

- *edg-job-submit* [-r <res_id>] [-n <user_e-mail address>] [-c <config_file>] [-o <output_file>] **<job.jdl>**
 - -c the configuration file <config_file> is pointed by the UI instead of the standard configuration file
 - -o the generated *dg_jobld* is written in the <output_file> Useful for other commands like *dg-job-status* which display the status information about dg_jobld contained in the <input file>



Job Submission

- *JDL* a file that specifies the resources needed in order to execute the job
- *Job ID* an unique string that identify the job in the grid
- JDLs are **written once** per job type.
- A single JDL specify all jobs that use the same resources.



Job Submission

- *Job IDs are provided by the grid (spec. the RB)*
- A Job ID identify a **single job** and is **the only** reference the user has about it, once submitted.
- The job ID is the only reference the user has for the job.
- Losing the ID is the same as losing the job.
- The job will keep running on the grid, however you will be unable to retrieve data from the output sandbox at the end, or to stop it.



Job Submission

JDL example file:

Executable = "gridTest";

StdInput = "event1.dat";

StdError = "stderr.log";

StdOutput = "stdout.log";

InputSandbox = {"/home/test/gridTest", /tmp/event1.dat"};

OutputSandbox = {"stderr.log", "stdout.log"};

InputData = {"LF:testbed0-00019", "PF:se1.pd.infn.it/data/filex"};

**ReplicaCatalog = "ldap://sunlab2g.cnaf.infn.it:2010/
lc=test, rc=WP2 INFN Test, dc=infn, dc=it";**

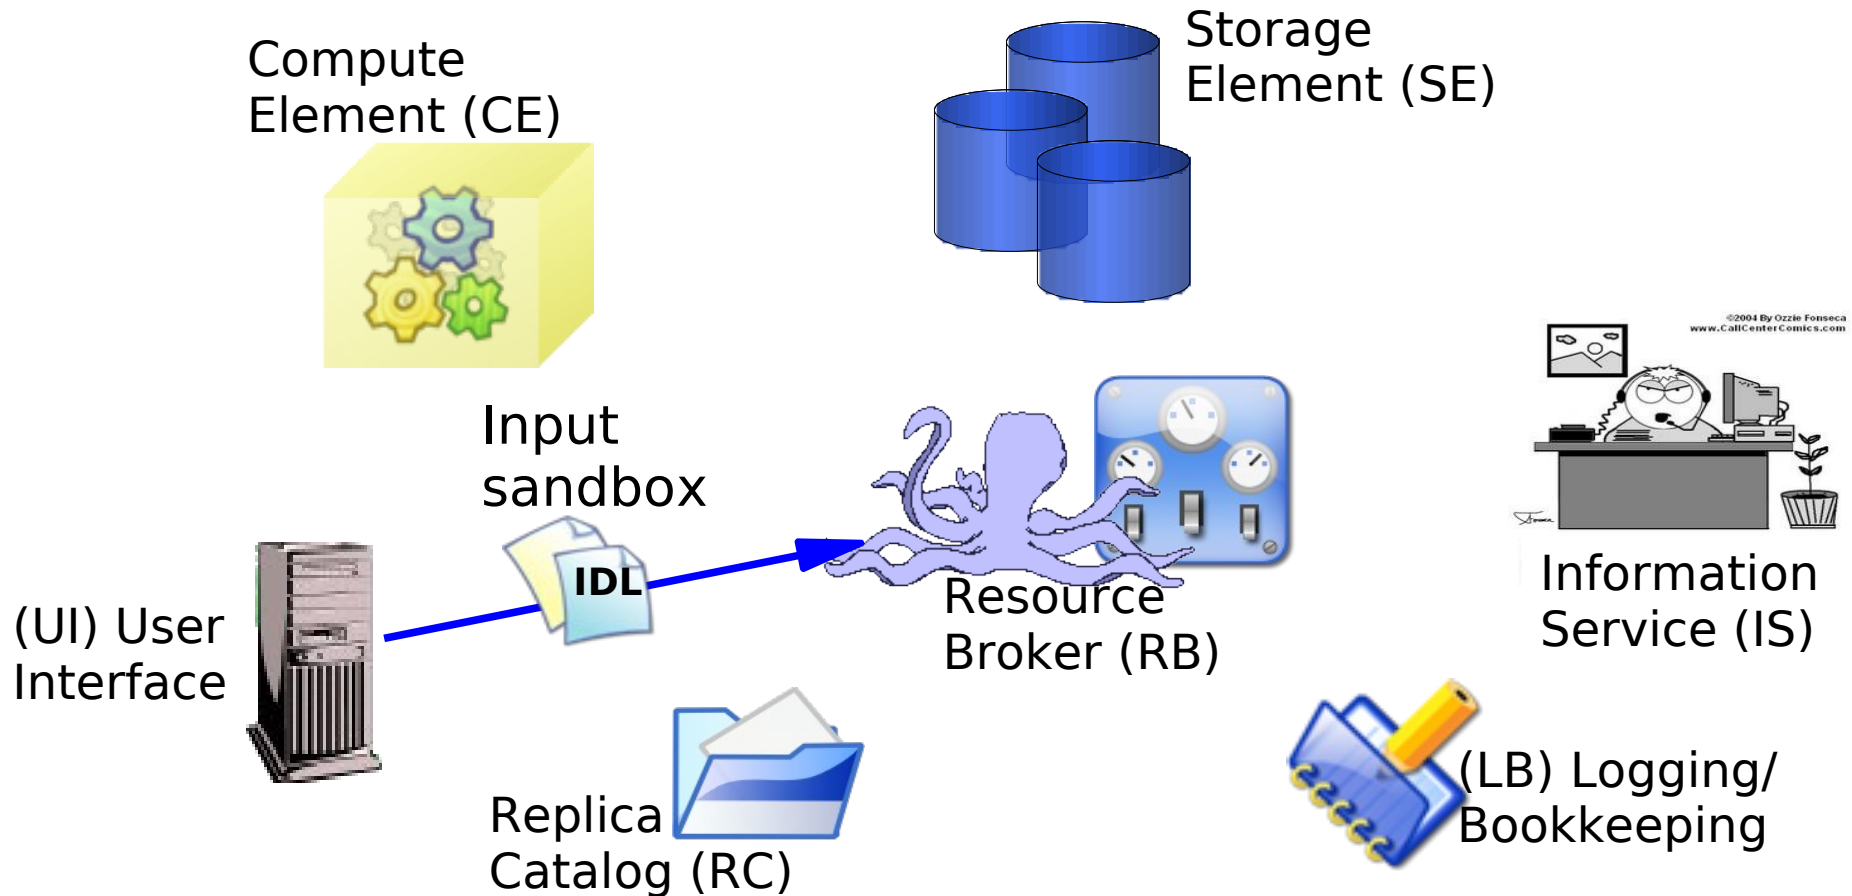
DataAccessProtocol = {"gridftp", "file"};

**Requirements = other.Architecture=="INTEL" && \
other.OpSys=="LINUX" && other.TotalCpus >= 4;**

Rank = other.MaxCpuTime;


Job Submission: Data flow

Once the user has open a remote session in the UI, he can submit jobs to the Resource Broker (RB).
 Job Description Language format files




Job Submission: Data flow

An event is loggen in LB and the state of the job is ***SUBMITTED***




Job Submission: Data flow

The RB looks for the best available CE
interrogating to IS. Job state is *WAITING*



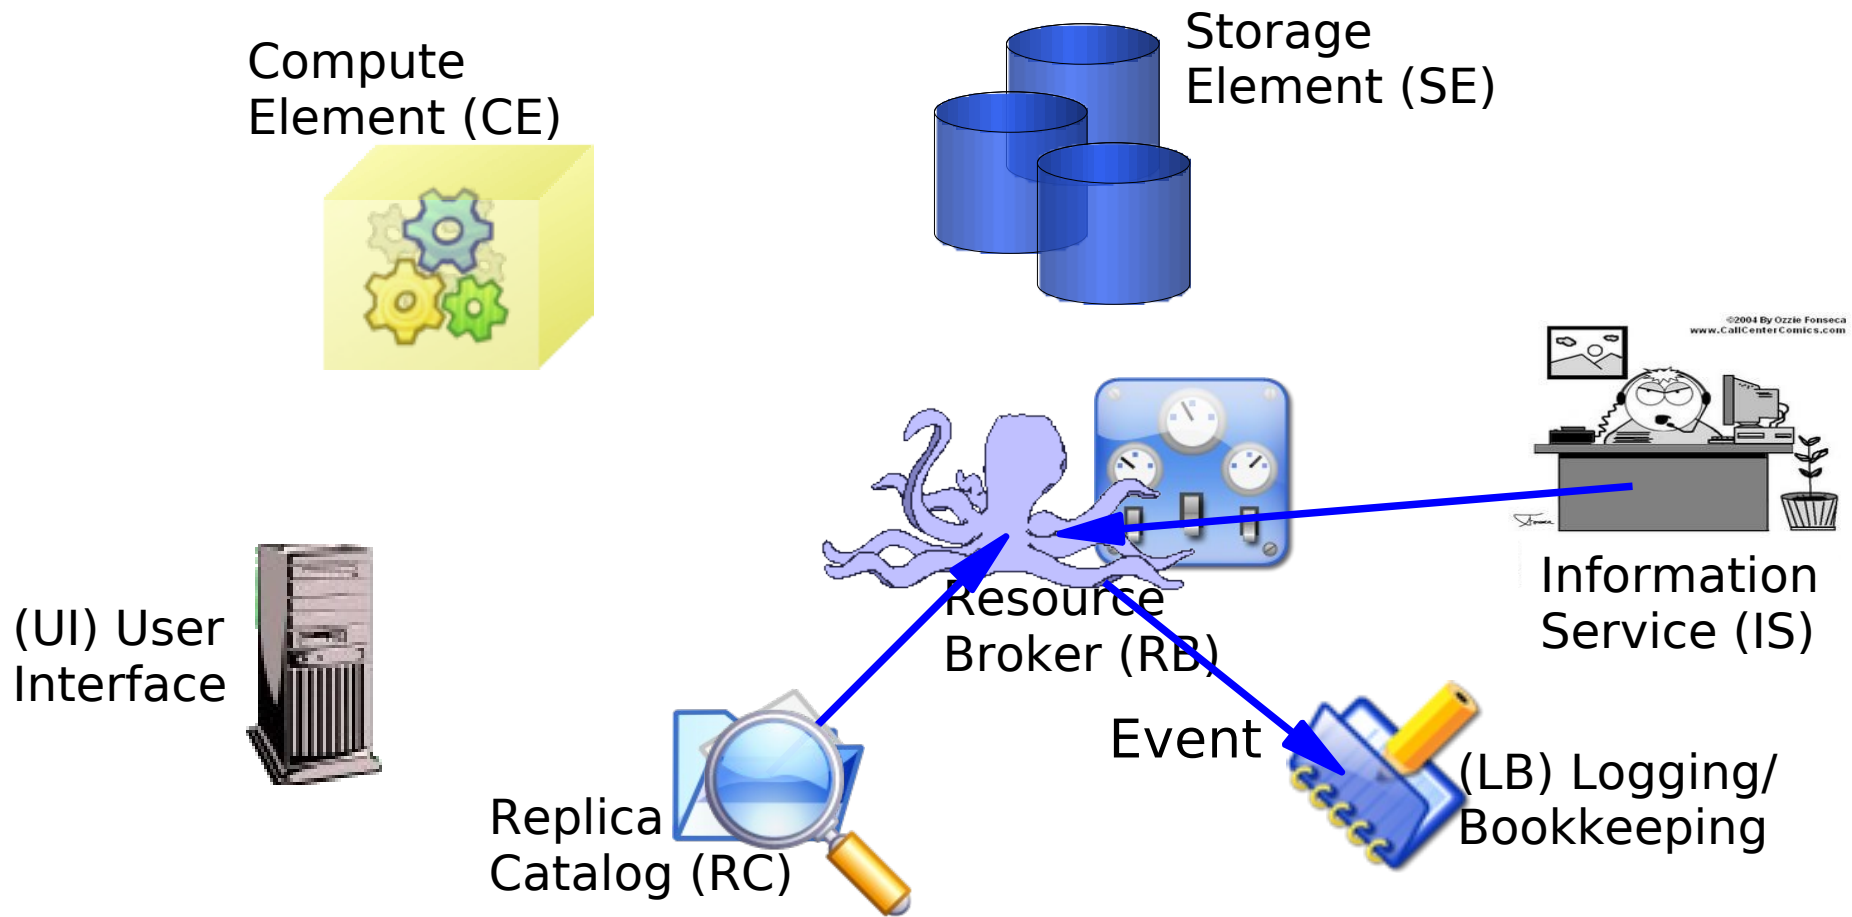
Job Submission: Data flow

IS obtain information about the CE and SE status




Job Submission: Data flow

RC looks for any input file required. IS send the response and another event is logged in LB



Job Submission: Data flow


RB prepares the job for submission. Another event is logged. Job state is *READY*





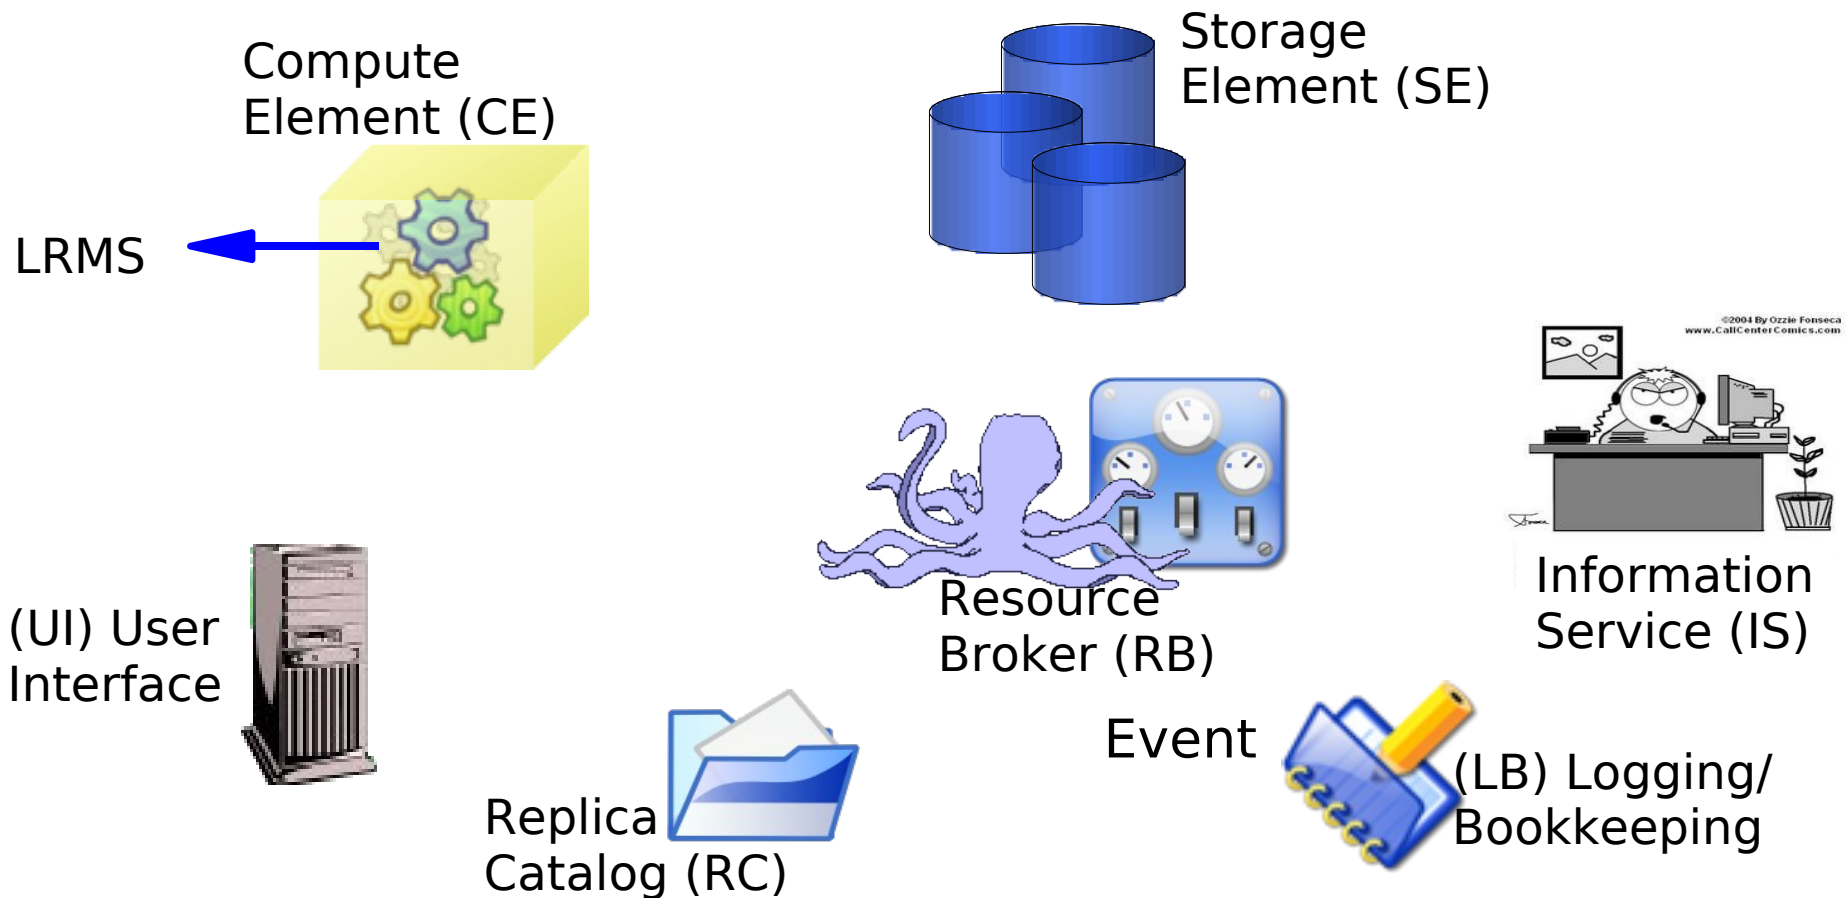
Job Submission: Data flow

CE receives the job



Job Submission: Data flow

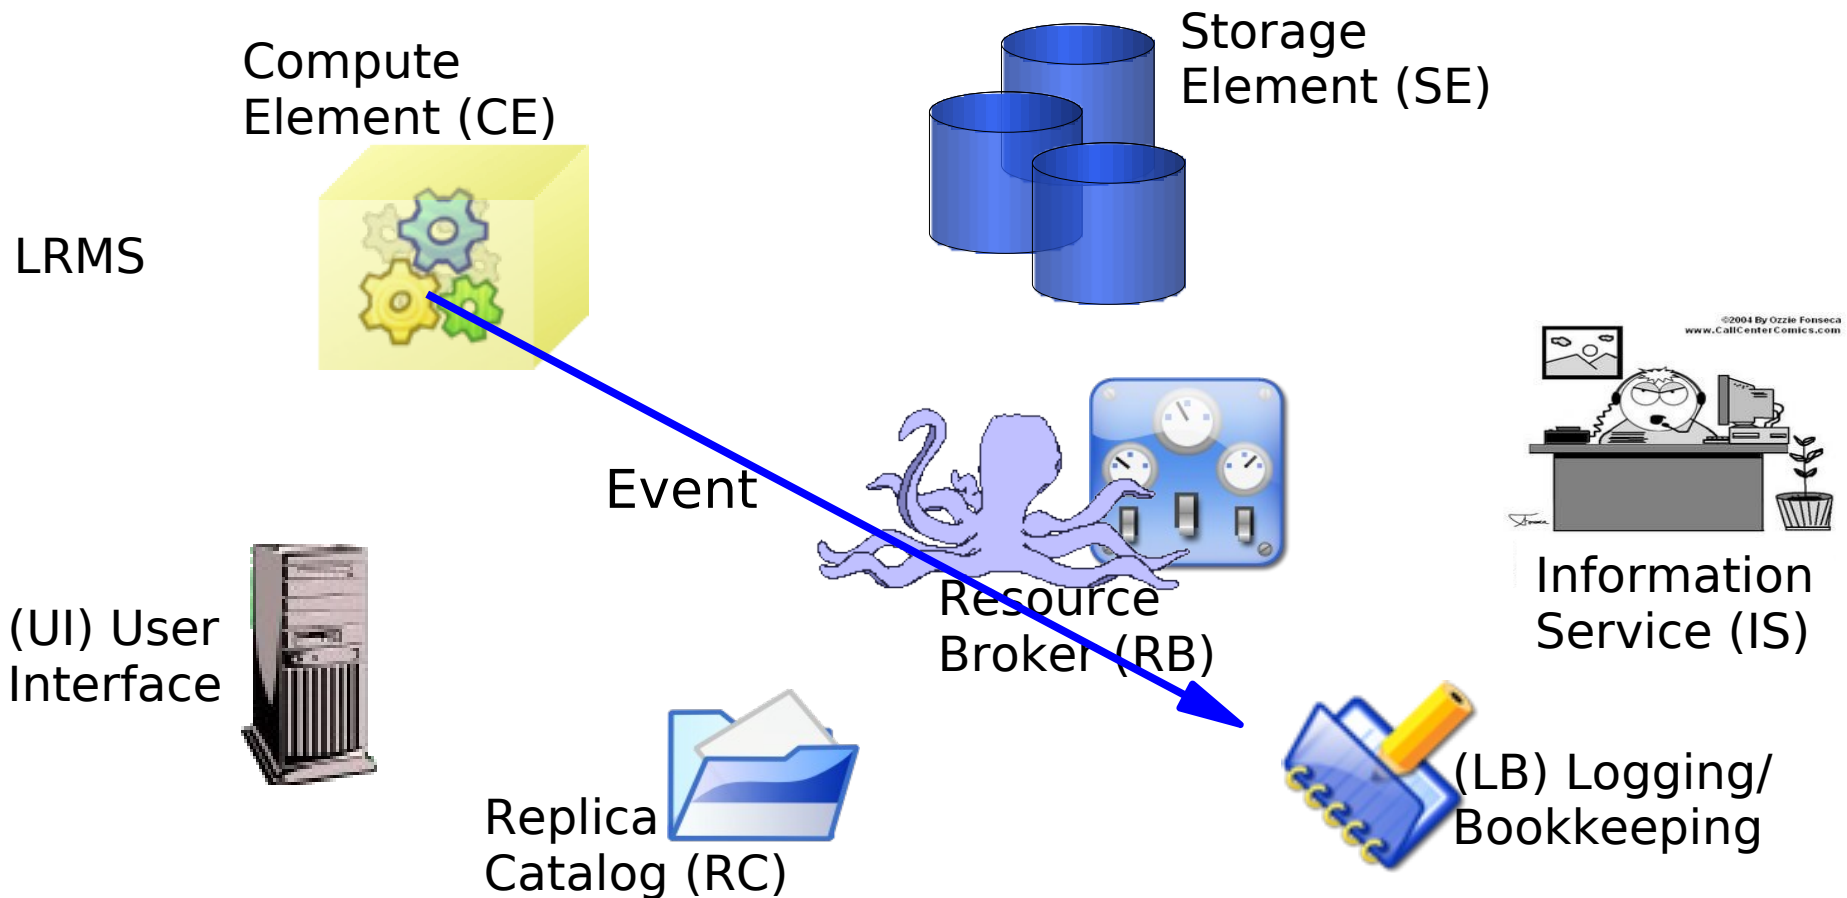
CE send the job to the Local Resource Manager System (PBS,PBSpro,LSF,BQS,Condor,Maui/Torque)





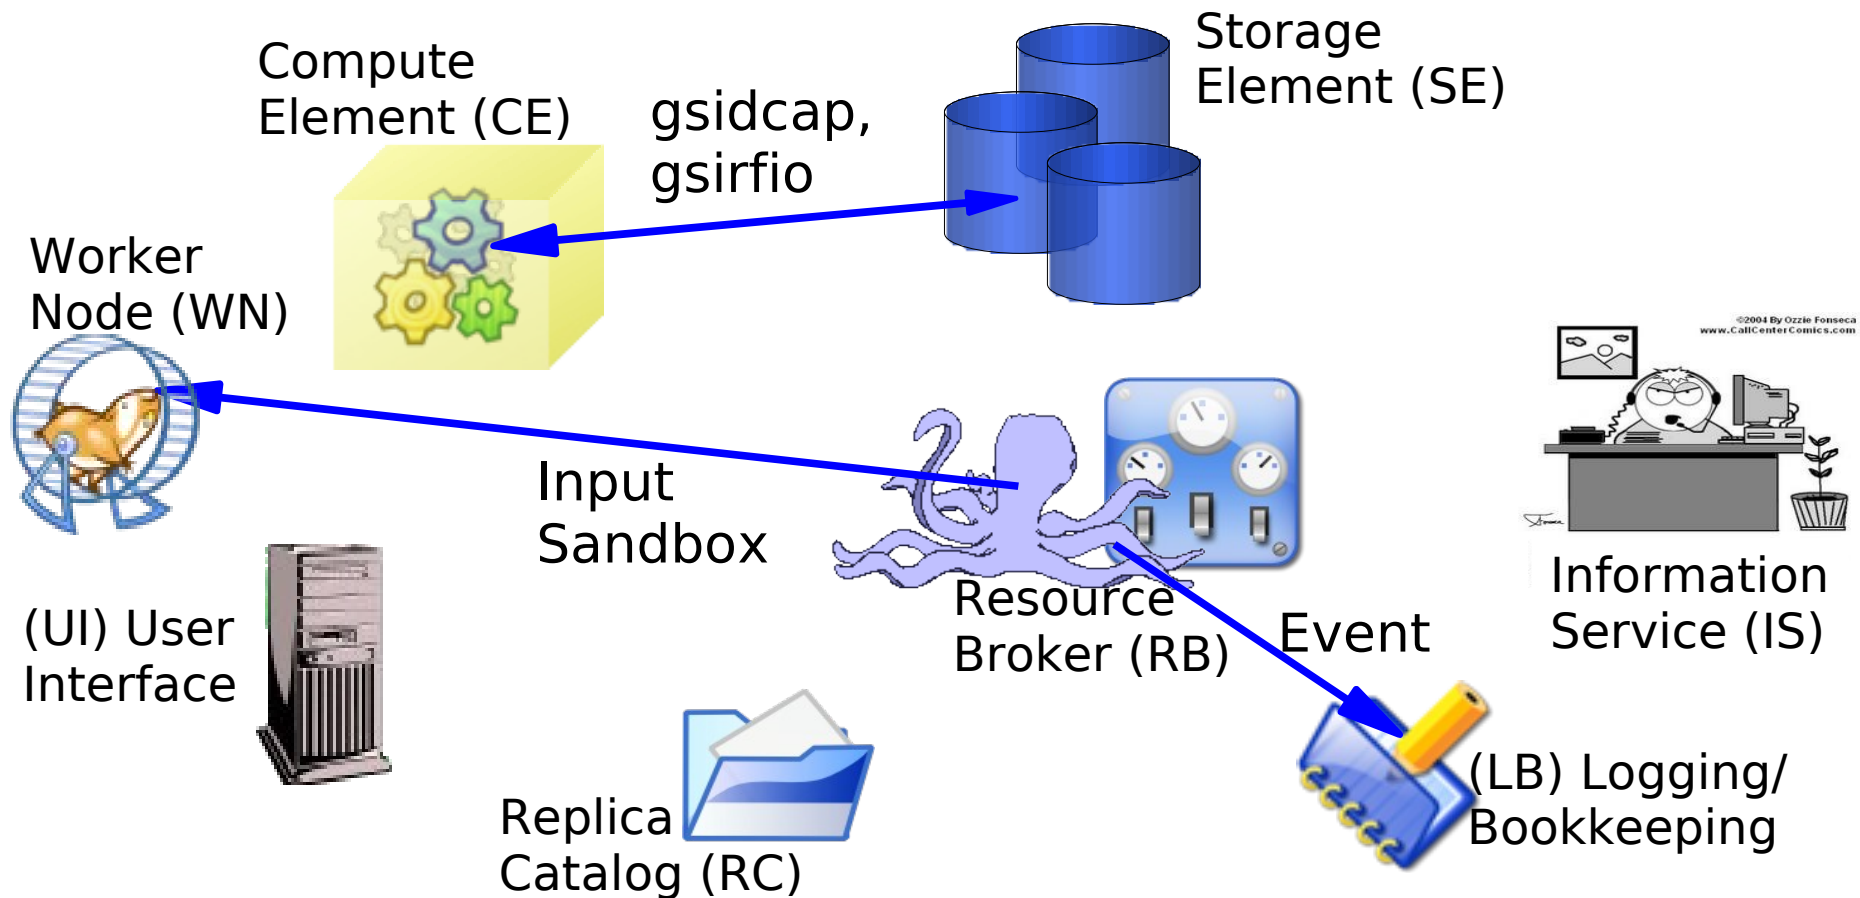
Job Submission: Data flow

An event is sent to LB and the state is *SCHEDULED*



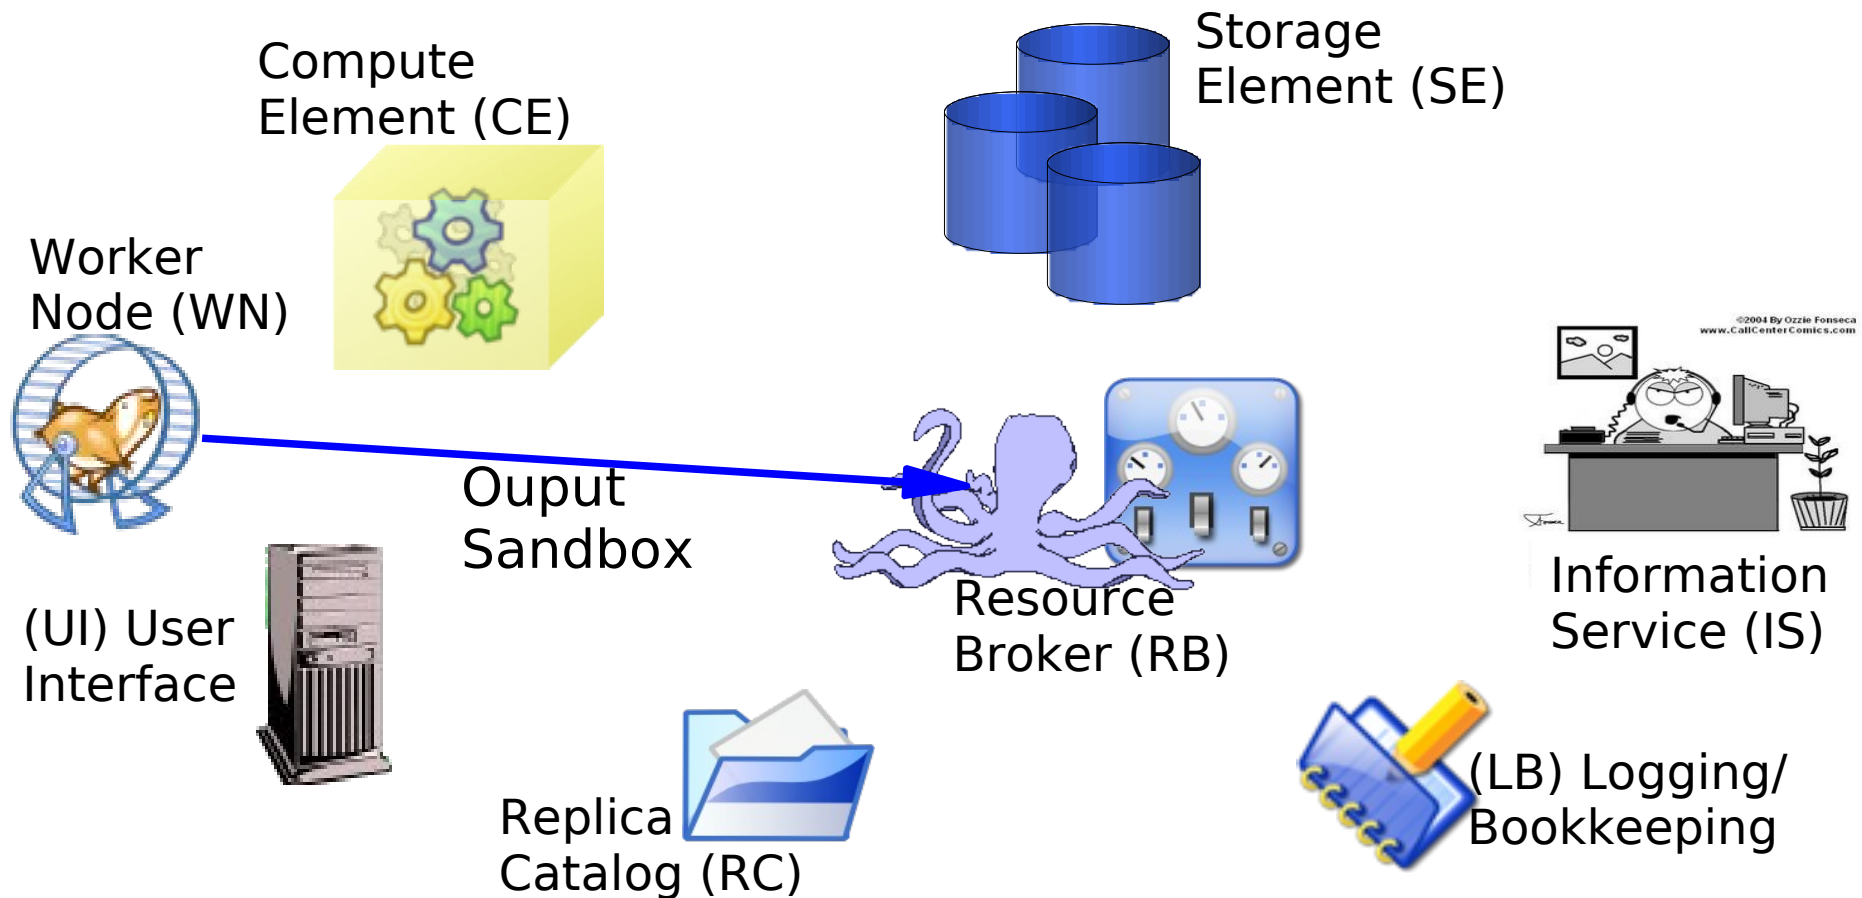
Job Submission: Data flow

When the LRMS send the job to WN, RB send the Input Sandbox to that WN. Now the job is ***RUNNING***



Job Submission: Data flow


If the job finish without errors the output files are sent to RB





Job Submission: Data flow


An event is logged and the state is *DONE*






Job Submission: Data flow

The output sandbox is sent to UI



Job Submission: Data flow

At this point the user can retrieve the results. An event is sent and the state is *CLEARED*





Job Management

Others job operations

- *edg-job-list-match* Lists resources matching a job description ordered by rank. Performs the matchmaking without submitting the job
- *edg-job-cancel* Cancels a specific job
- *edg-job-status* Displays the status of the job plus some additional information



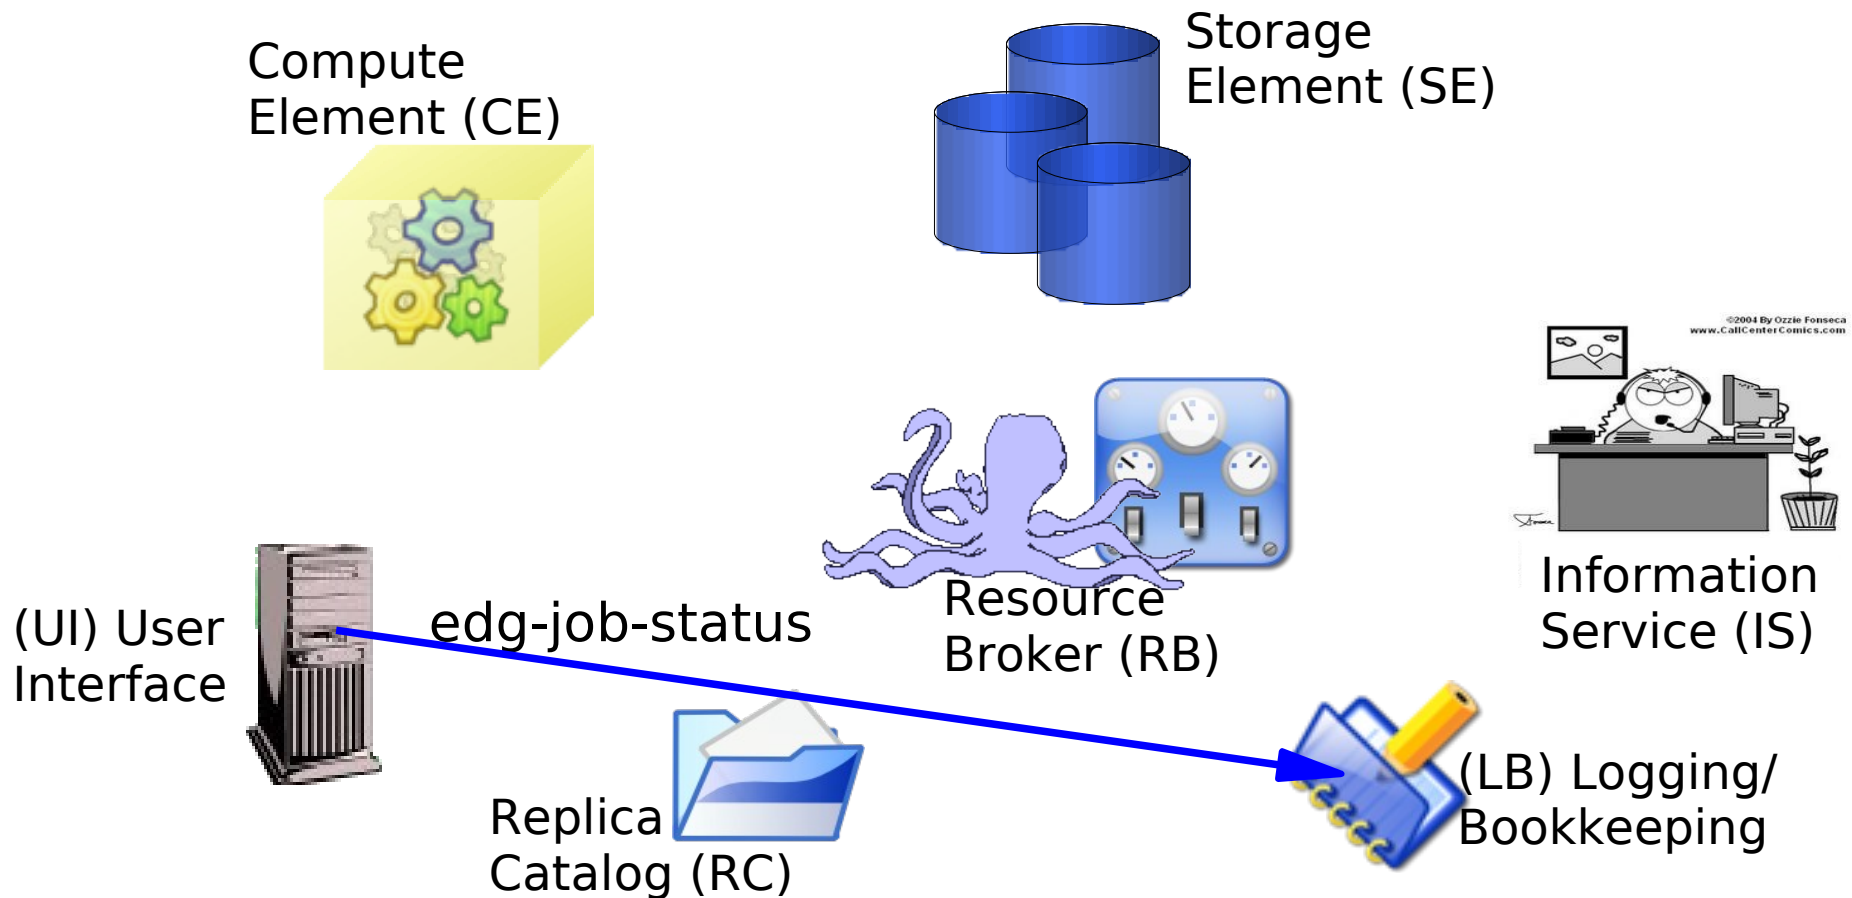
Job Management

Others job operations:

- *edg-job-get-output* Returns the job-output (the OutputSandbox files) to the user
- *edg-job-get-logging-info* Displays logging information about submitted jobs (all the events “pushed” by the various components of the WMS to the LB) Very useful for debug purposes
- *edg-job-id-info* A utility for the user to display job info in a formatted style

Job Management

At any time the user can get information about his/her jobs





```
#####
# Author: Giuseppe La Rocca (giuseppe.larocca@ct.infn.it) #
# ----- #
# Istituto Nazionale di Fisica Nucleare - Sez. di CATANIA #
# Via Santa Sofia,64 - 95100 ITALY #
# ----- #
# +39.095.378.54.46 +39.095.378.54.69 #
#####
```

```
Type = "Job";
JobType = "MPICH";
```

```
Executable = "start_helloworldmpi.sh";
Arguments = "helloworld 2";
NodeNumber = 2;
```

```
StdOutput = "test.out";
StdError = "test.err";
```

```
InputSandbox = {"start_helloworldmpi.sh","helloworld"};
OutputSandbox = {"test.err","test.out","executable.out"};
```

```
Requirements = ((other.GlueCEInfoLRMSType == "PBS") || (other.GlueCEInfoLRMSType == "LSF"));
```

```
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
```

1,1

All



```

| else by yourself.
| In case of any problems, please send an email
| to grid-prod@ct.infn.it
+-----+
|
|           ATTENTION!
|
| Since the 1st of January 2006 GILDA supports
| VOMS authentication only! In order to login on
| the grid, please use the command:
|
| voms-proxy-init --voms gilda
|
| instead of:
|
| grid-proxy-init
+-----+
[glite-tutor] /home/amuoi > ls
glite MessageToSpheres openfoam_examples plot.C trailers
[glite-tutor] /home/amuoi > cd glite
[glite-tutor] /home/amuoi/glite > ls
CODESA-3D edglog.log geant Other Raster vlc
DAG EGEODE mmcs pathsearch scilab
[glite-tutor] /home/amuoi/glite > cd Other/
[glite-tutor] /home/amuoi/glite/Other > voms-proxy-init --voms gilda
Cannot find file or dir: /home/amuoi/.glite/vomses
Your identity: /C=IT/O=GILDA/OU=Personal Certificate/L=INFN/CN=annamaria muoi/E
mail=amuoi@yahoo.it
Enter GRID pass phrase:
Your proxy is valid until Wed Jan 25 23:53:39 2006

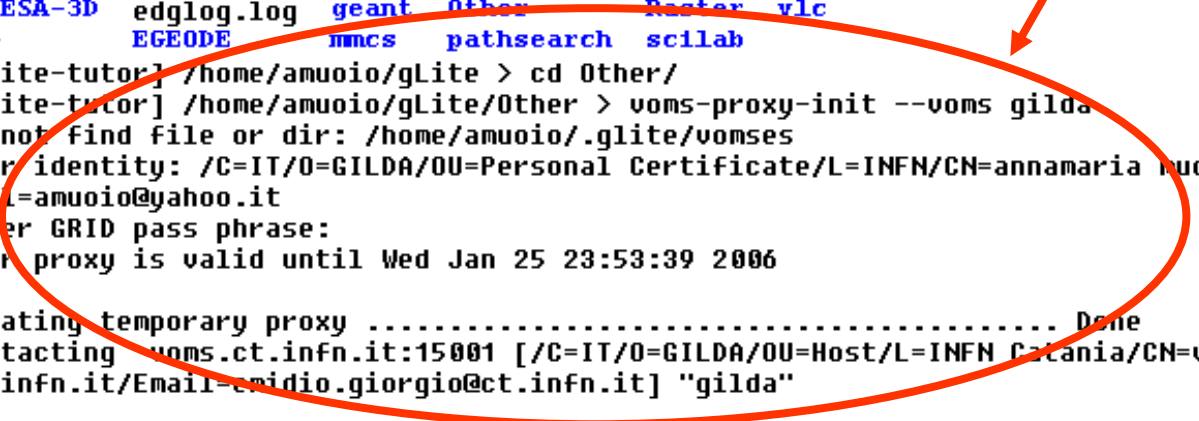
Creating temporary proxy ..... Done
Contacting voms.ct.infn.it:15001 [/C=IT/O=GILDA/OU=Host/L=INFN Catania/CN=voms.
ct.infn.it/Email=midio.giorgio@ct.infn.it] "gilda"

Warning: Cannot find file or dir: /home/amuoi/.glite/vomses Done
Creating proxy .....
.. Done
Your proxy is valid until Wed Jan 25 23:53:40 2006

[glite-tutor] /home/amuoi/glite/Other >

```

Voms-proxy-init





Warning: Cannot find file or dir: /home/amuoiio/gLite/vomses Done
Creating proxy
.. Done
Your proxy is valid until Wed Jan 25 23:53:40 2006

[glite-tutor] /home/amuoiio/gLite/Other > glite-job-list-match helloworld.jdl

Selected Virtual Organisation name (from proxy certificate extension): gilda
Connecting to host glite-rb.ct.infn.it, port 7772

COMPUTING ELEMENT IDs LIST

The following CE(s) matching your job requirements have been found:

- *CEId*
- dgt01.ui.savba.sk:2119/jobmanager-lcgpbs-infinite
- dgt01.ui.savba.sk:2119/jobmanager-lcgpbs-long
- dgt01.ui.savba.sk:2119/jobmanager-lcgpbs-short
- gilda-ce-01.pd.infn.it:2119/jobmanager-lcgpbs-infinite
- gilda-ce-01.pd.infn.it:2119/jobmanager-lcgpbs-long
- gilda-ce-01.pd.infn.it:2119/jobmanager-lcgpbs-short
- grid004.iucc.ac.il:2119/jobmanager-lcgpbs-infinite
- grid004.iucc.ac.il:2119/jobmanager-lcgpbs-long
- grid004.iucc.ac.il:2119/jobmanager-lcgpbs-short
- gildace.oact.inaf.it:2119/jobmanager-lcgpbs-infinite
- gildace.oact.inaf.it:2119/jobmanager-lcgpbs-long
- gildace.oact.inaf.it:2119/jobmanager-lcgpbs-short
- egee008.cnaf.infn.it:2119/blah-pbs-infinite
- egee008.cnaf.infn.it:2119/blah-pbs-long
- egee008.cnaf.infn.it:2119/blah-pbs-short
- glite-ce.ct.infn.it:2119/blah-pbs-infinite
- glite-ce.ct.infn.it:2119/blah-pbs-long
- glite-ce.ct.infn.it:2119/blah-pbs-short
- lxcde01.pd.infn.it:2119/blah-pbs-infinite
- lxcde01.pd.infn.it:2119/blah-pbs-long
- lxcde01.pd.infn.it:2119/blah-pbs-short

[glite-tutor] /home/amuoiio/gLite/Other > █

glite-job-list-match helloworld.jdl
It gets a list of available CE



```

dgt01.ui.savba.sk:2119/jobmanager-lcgpbs-long
dgt01.ui.savba.sk:2119/jobmanager-lcgpbs-short
gilda-ce-01.pd.infn.it:2119/jobmanager-lcgpbs-infinite
gilda-ce-01.pd.infn.it:2119/jobmanager-lcgpbs-long
gilda-ce-01.pd.infn.it:2119/jobmanager-lcgpbs-short
grid004.iucc.ac.il:2119/jobmanager-lcgpbs-infinite
grid004.iucc.ac.il:2119/jobmanager-lcgpbs-long
grid004.iucc.ac.il:2119/jobmanager-lcgpbs-short
gildace.oact.inaf.it:2119/jobmanager-lcgpbs-infinite
gildace.oact.inaf.it:2119/jobmanager-lcgpbs-long
gildace.oact.inaf.it:2119/jobmanager-lcgpbs-short
egee008.cnaf.infn.it:2119/blah-pbs-infinite
egee008.cnaf.infn.it:2119/blah-pbs-long
egee008.cnaf.infn.it:2119/blah-pbs-short
glite-ce.ct.infn.it:2119/blah-pbs-infinite
glite-ce.ct.infn.it:2119/blah-pbs-long
glite-ce.ct.infn.it:2119/blah-pbs-short
lxcde01.pd.infn.it:2119/blah-pbs-infinite
lxcde01.pd.infn.it:2119/blah-pbs-long
lxcde01.pd.infn.it:2119/blah-pbs-short

```

glite-job-submit helloworld.jdl
 It assigns one id to job

```

[glite-tutor] /home/amuoiio/gLite/Other > glite-job-submit helloworld.jdl
Selected Virtual Organisation name (from proxy certificate extension): gilda
Connecting to host glite-rb.ct.infn.it, port 7772
Logging to host glite-rb.ct.infn.it, port 9002

```

```

*****
JOB SUBMIT OUTCOME
The job has been successfully submitted to the Network Server.
Use glite-job-status command to check job current status. Your job identifier is:
- https://glite-rb.ct.infn.it:9000/oq13raDQn79U2U4IwUFI-g
*****

```

```

[glite-tutor] /home/amuoiio/gLite/Other > ~

```



Copyright (C) 2003 by DATAMAT SpA

Usage: glite-job-status [options] <job Id(s)>

Options:

--help
--version

--all
--input, -i <input file>
--verbosity [0|1|2|3]
--from [MM:DD:]hh:mm[:[CC]YY]
--to [MM:DD:]hh:mm[:[CC]YY]
--config, -c <config file>
--user-tag <tag name>=<tag value>
--status, -s <status value>
--exclude, -e <exclude value>
--config-vo <config-vo file>
--vo <vo value>
--output, -o <output file>
--noint
--debug
--logfile <logfile file>

Job has terminated its execution



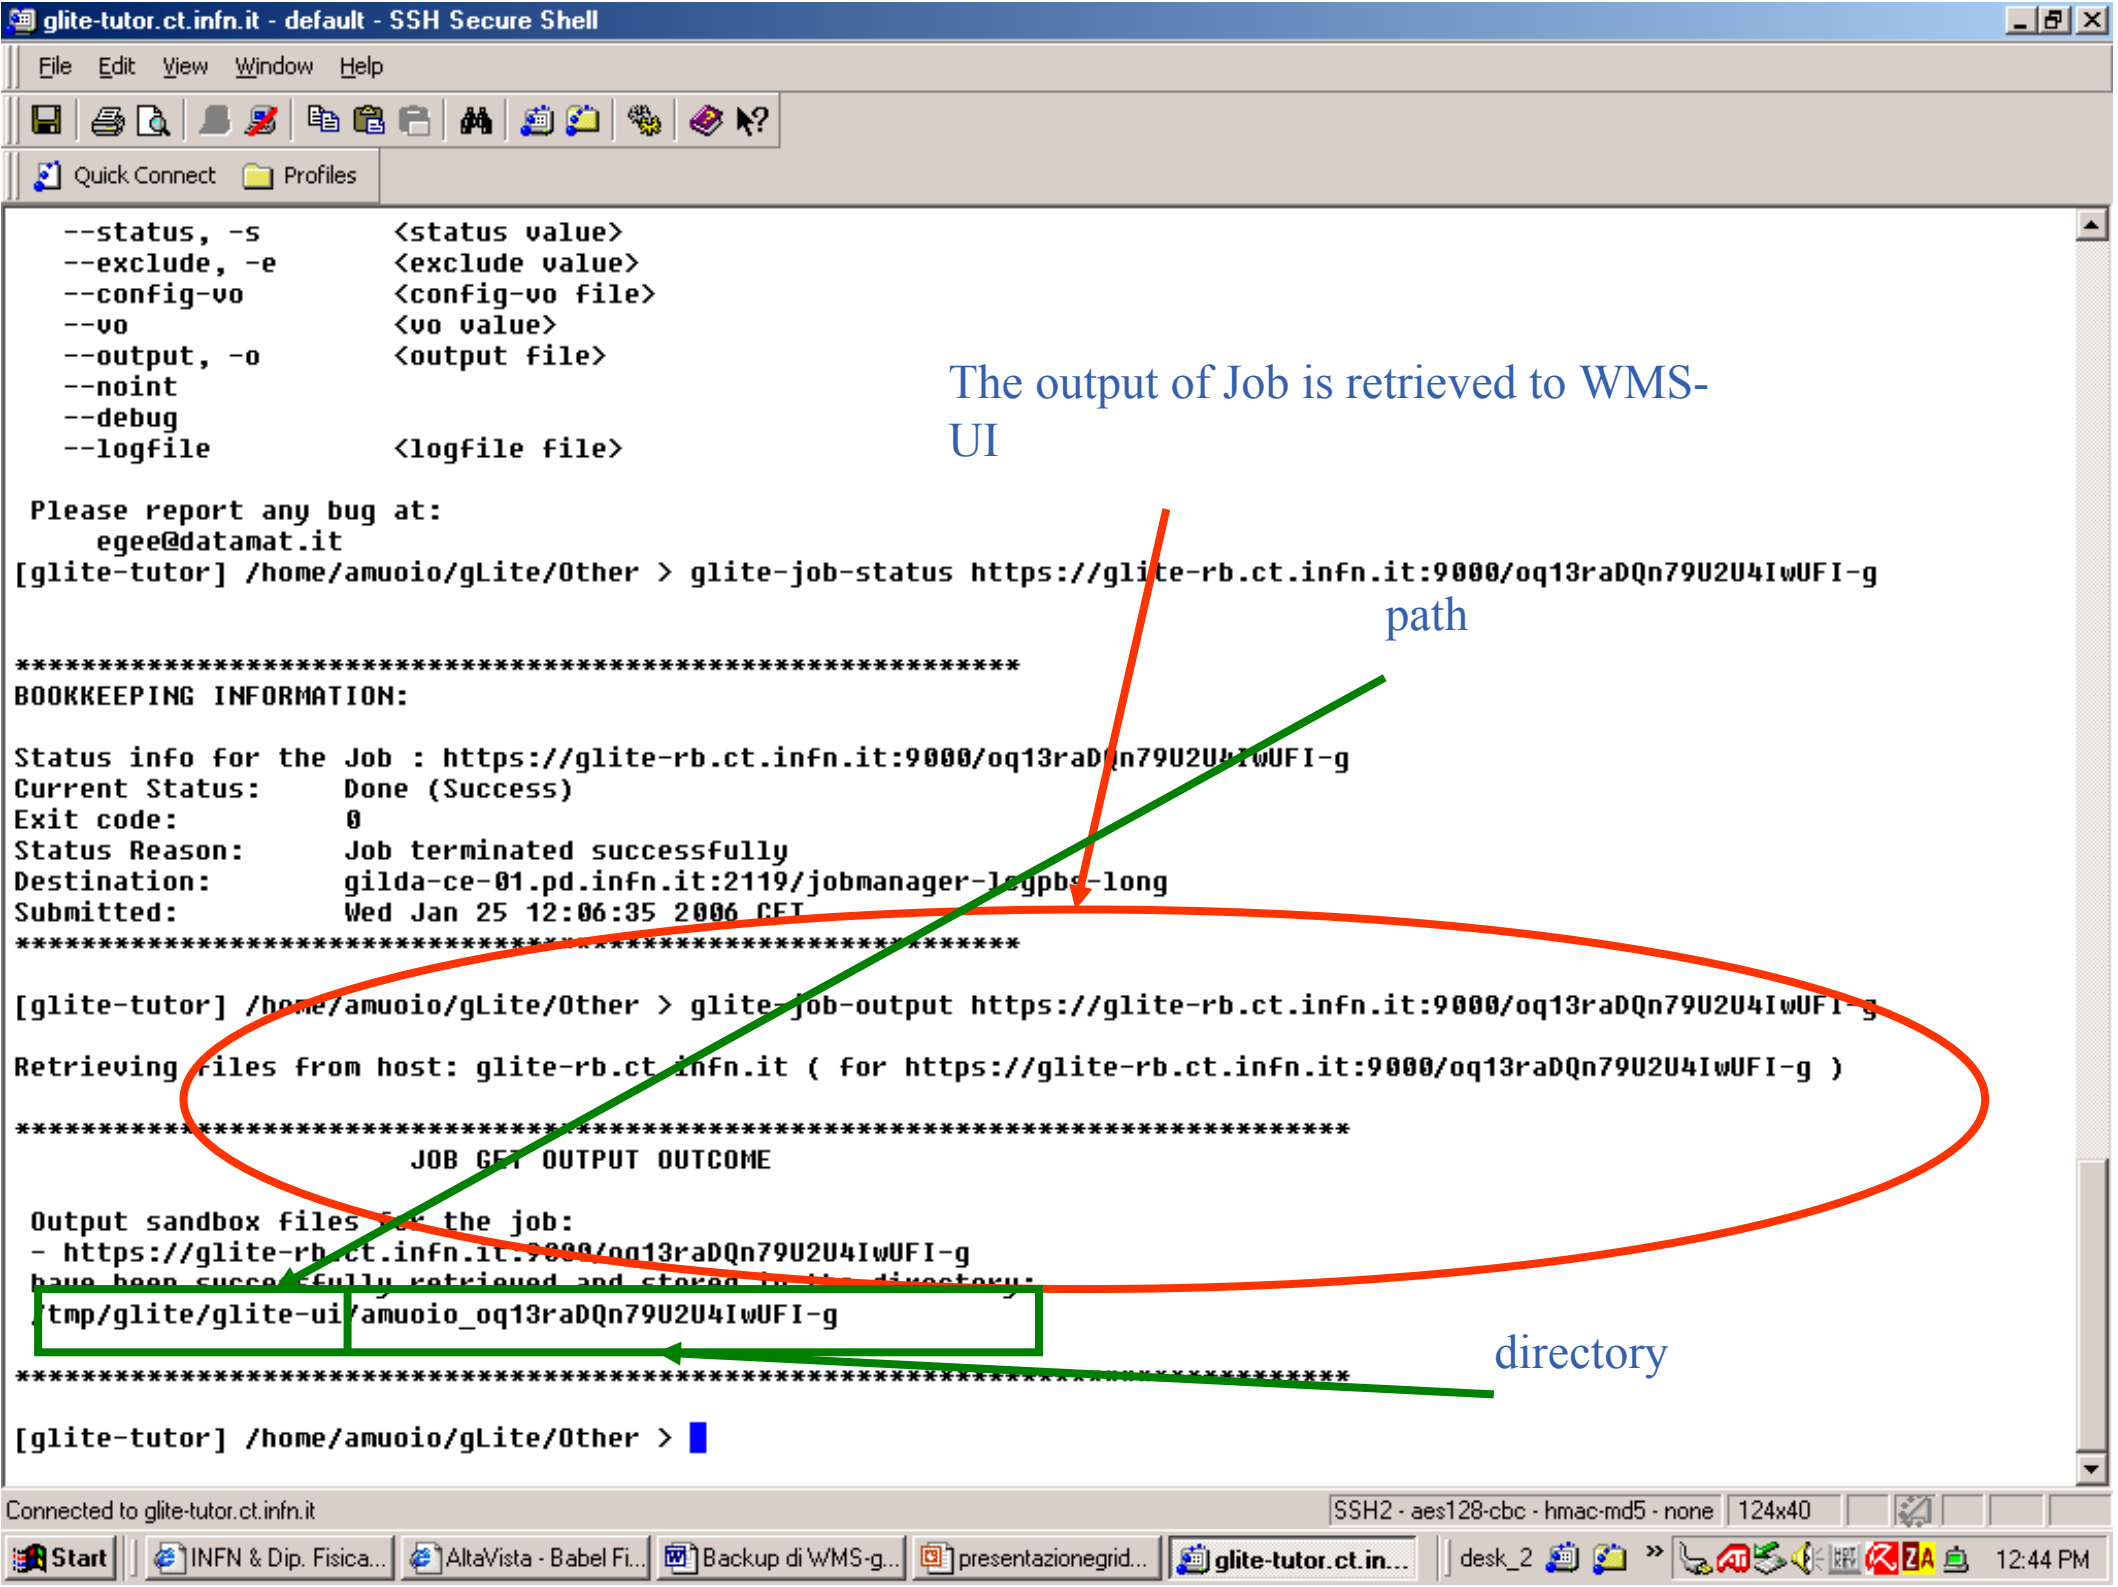
Please report any bug at:
ege@datamat.it

[glite-tutor] /home/amuoiio/gLite/Other > glite-job-status https://glite-rb.ct.infn.it:9000/oq13raDQn79U2U4IwUFI-g

BOOKKEEPING INFORMATION:

Status info for the Job : https://glite-rb.ct.infn.it:9000/oq13raDQn79U2U4IwUFI-g
Current Status: Done (Success)
Exit code: 0
Status Reason: Job terminated successfully
Destination: gilda-ce-01.nd.infn.it:2119/jobmanager-lcgpbs-long
Submitted: Wed Jan 25 12:06:35 2006 CET

[glite-tutor] /home/amuoiio/gLite/Other > █



The output of Job is retrieved to WMS-UI

path

directory


```
WMS User Interface version 1.0.9
Copyright (C) 2003 by DATAMAT SpA
```

```
Usage: glite-job-submit [options] <jdl file>
```

```
Options:
```

```
--help
--version

--input, -i      <input file>
--resource, -r   <resource value>
--chkpt          <chkpt file>
--nolisten
--nogui
--nomsg
--lrms           <lrms value>
--to             [MM:DD:]hh:mm[:[CC]YY]
--valid, -v     <hours>:<minutes>
--config, -c    <config file>
--config-vo     <config-vo file>
--vo            <vo value>
--output, -o    <output file>
--noint
--debug
--logfile       <logfile file>
```

If there is a problem it's possible to cancel it



```
Please report any bug at:
  egee@datamat.it
```

```
[glite-tutor] /home/amuoiio/gLite/Other > glite-job-cancel https://glite-rb.ct.infn.it:9000/TB_atiEDYsnCt2Ab_mP3yg
```

```
Are you sure you want to remove specified job(s)? [y/n]n :y
```

```
----- glite-job-cancel Success -----
```

```
The cancellation request has been successfully submitted for the following job(s):
```

```
- https://glite-rb.ct.infn.it:9000/TB_atiEDYsnCt2Ab_mP3yg
```

```
[glite-tutor] /home/amuoiio/gLite/Other > █
```



Advanced School in High Performance Computing Tools for e-Science
5 March - 16 March 2007

Thank you!