

Harnessing the Grid

A tutorial for using the CrossGrid/DataGrid testbed
and the local site at the HPCL

George Tsouloupas
HPCL

Contents

- Architecture
- Grid Security/New Accountns
- The Testbed
- Our Site (status)
- Usage (submitting/status)

Architecture

Local Applications

Local Database

Local Computing

Grid

Grid
Application
Layer

Job
Management

Data
Management

Metadata
Management

Object toFile
Mapping

Collective
Services

Information
&
Monitoring

Replica
Manager

Grid
Scheduler

Underlying Grid Services

SQL
Database
Service

Computing
Element
Service

Storage
Element
Service

Replica
Catalog

Authentication
Authorization
Accounting

Service
Index

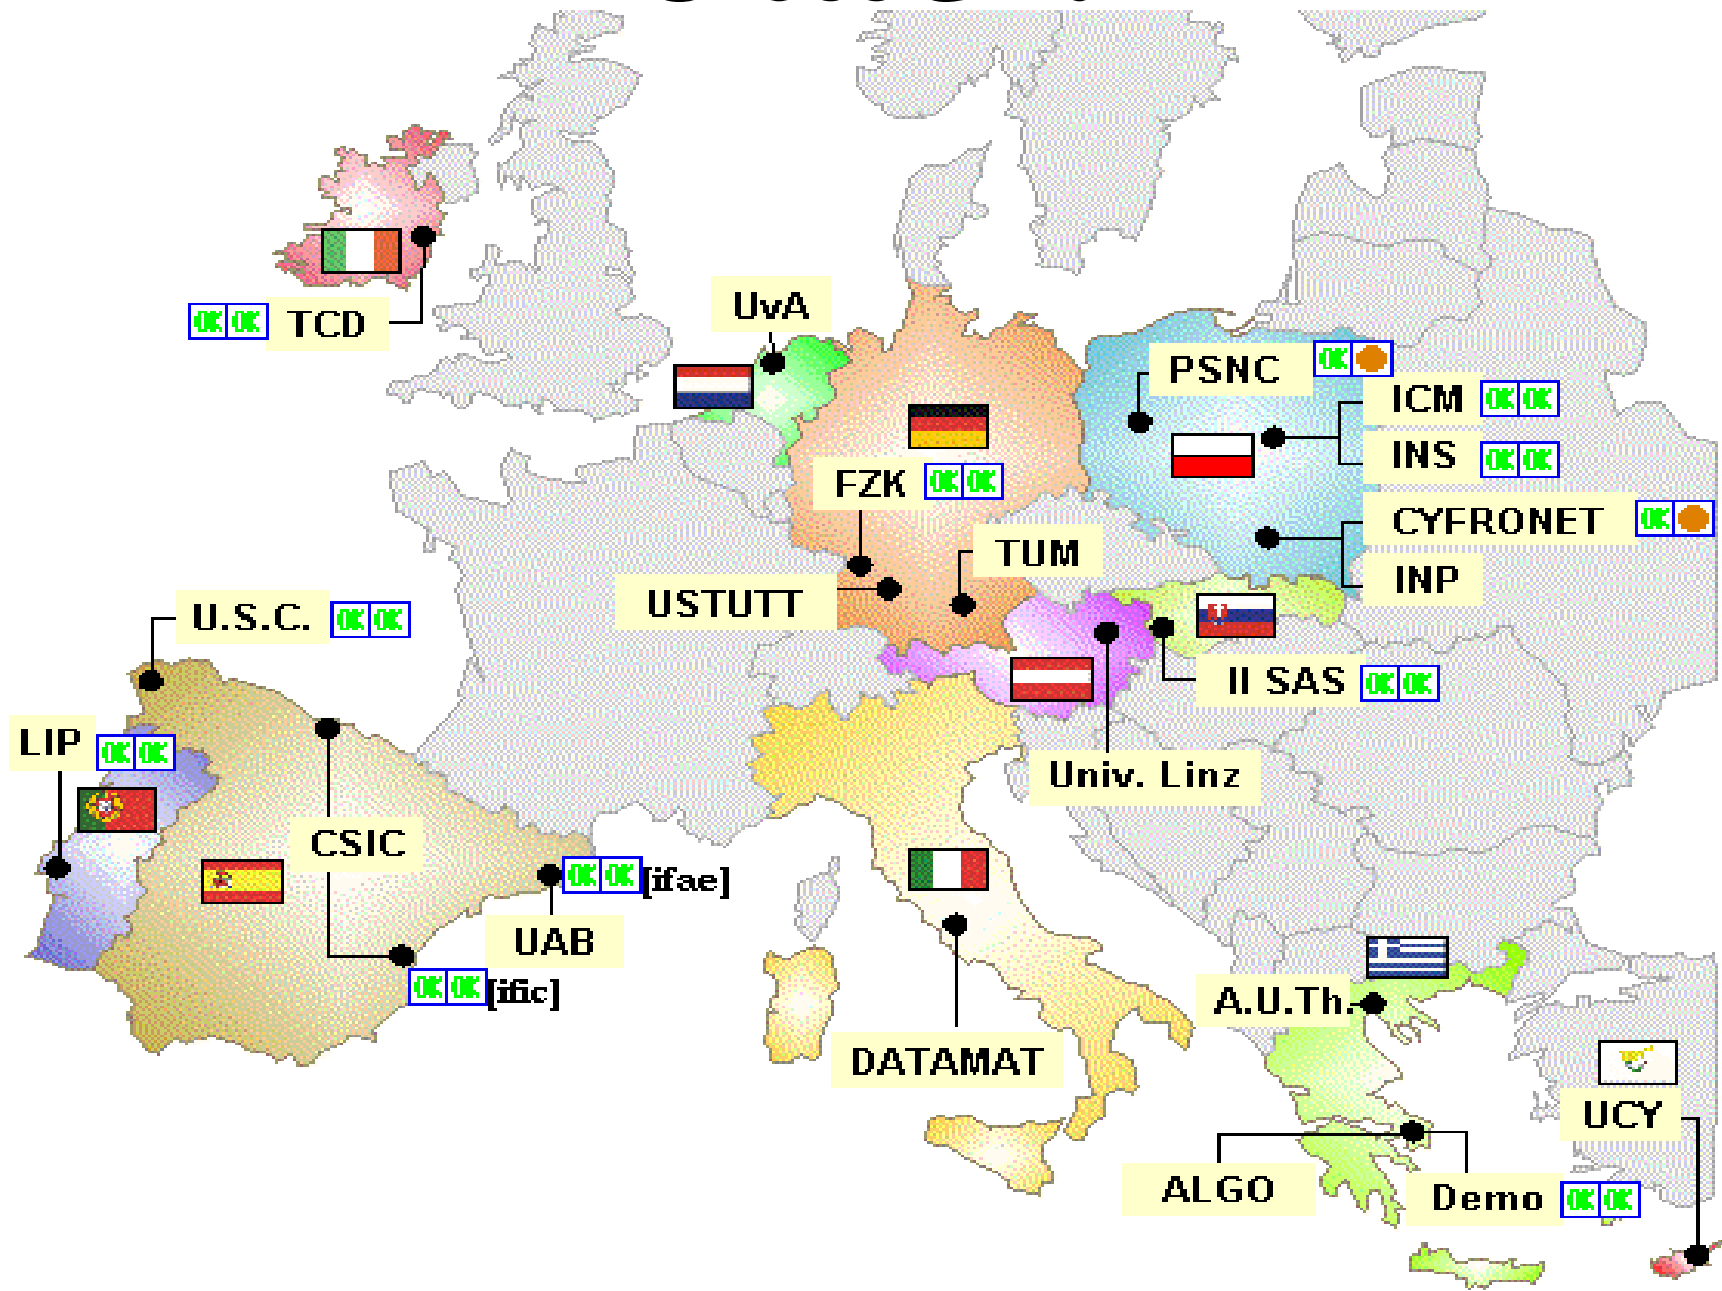
Grid

Fabric

Fabric Services

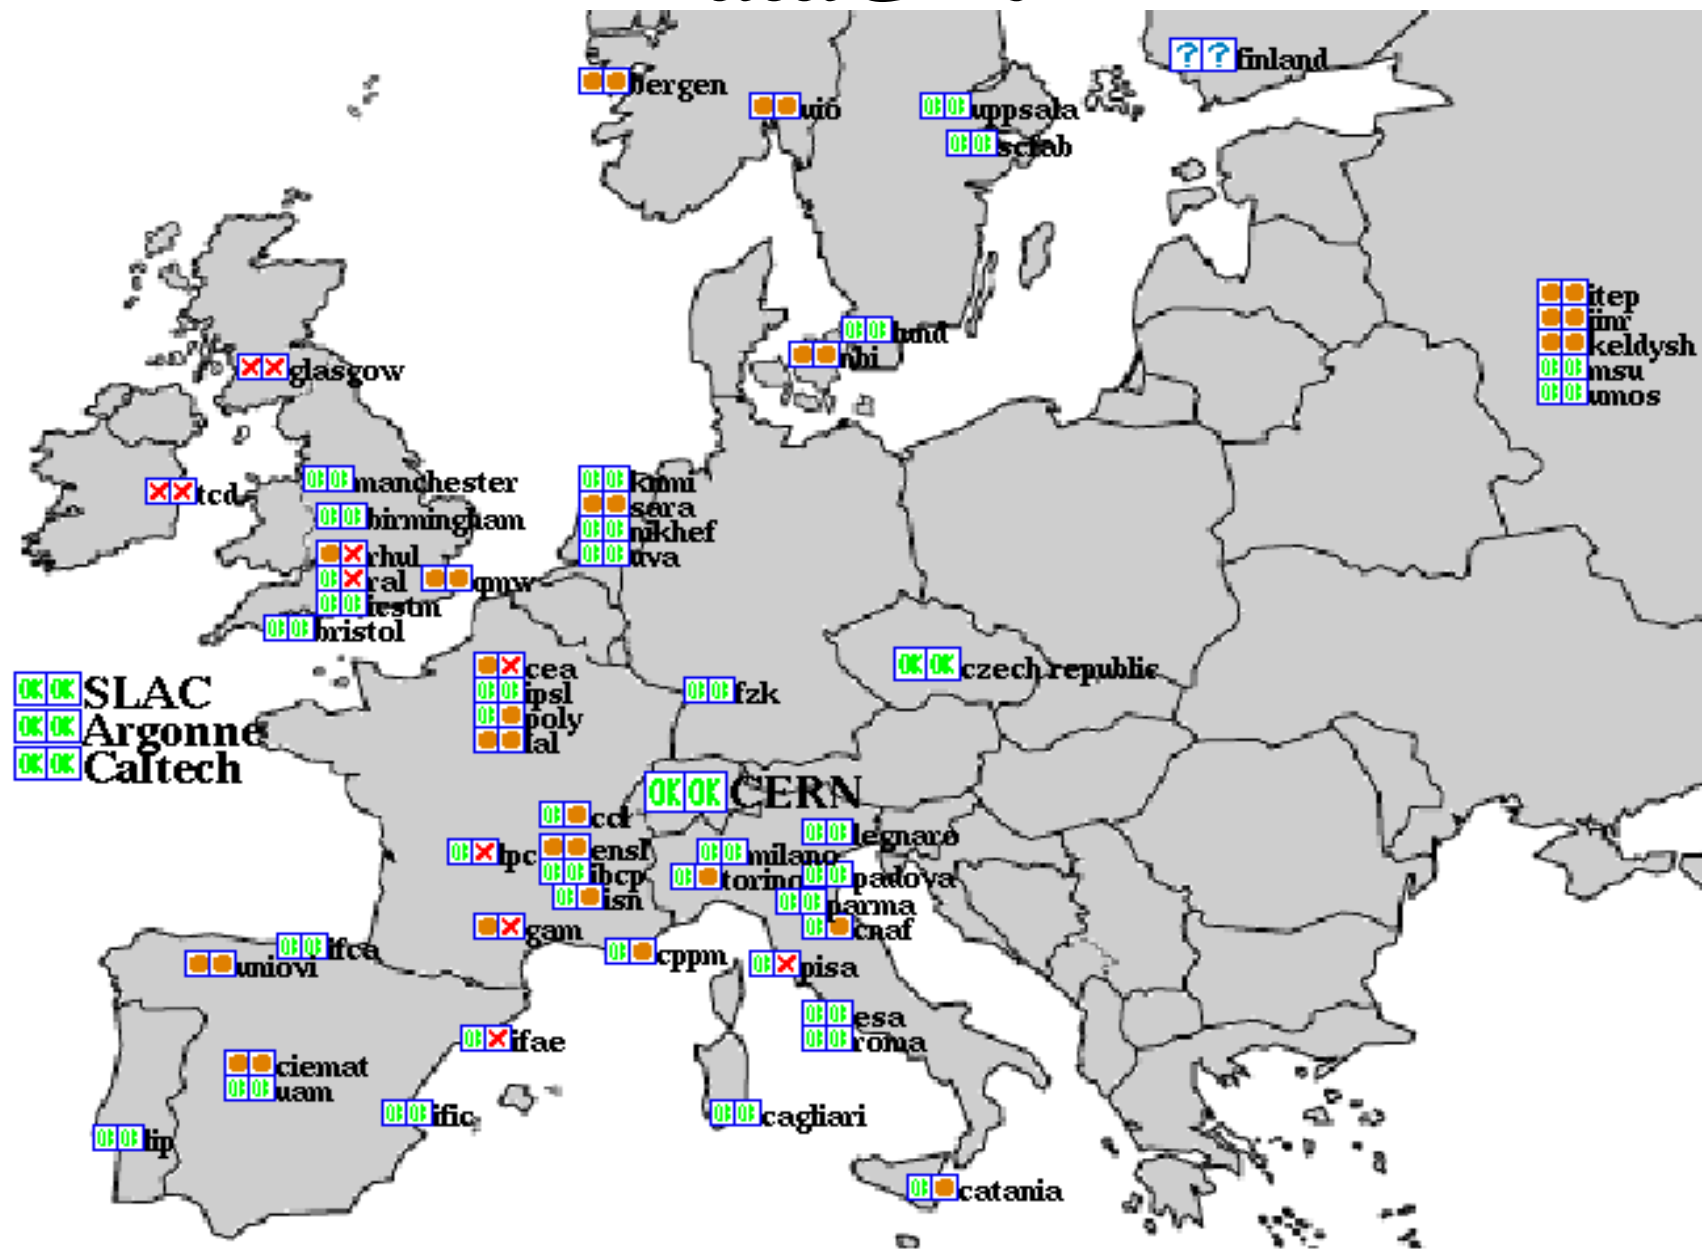
(Installation, Configuration, Fabric Storage management)

CrossGrid



Map Link Symbol No Status Normal TCP/UDP/URL failed Ping failed

DataGrid













Wed, 11 Dec 2002 15:06:10 GMT (refresh=10min)

Map Link Symbol ? No Status Normal TCP/UDP/URL failed Ping failed

Grid Security

- Security is a primary concern when it comes to Grids.
- The Grid Security Infrastructure is based on Public Key Infrastructure.
- A user is authenticated based on possession of a *private key* (known **only** to the user).
- Authorization depends on whether the user's certificate is signed by an *accepted* Certification Authority (In our case the CyGrid CA)
- The gridmap-file at the target resource will determine the local user account that will run the job

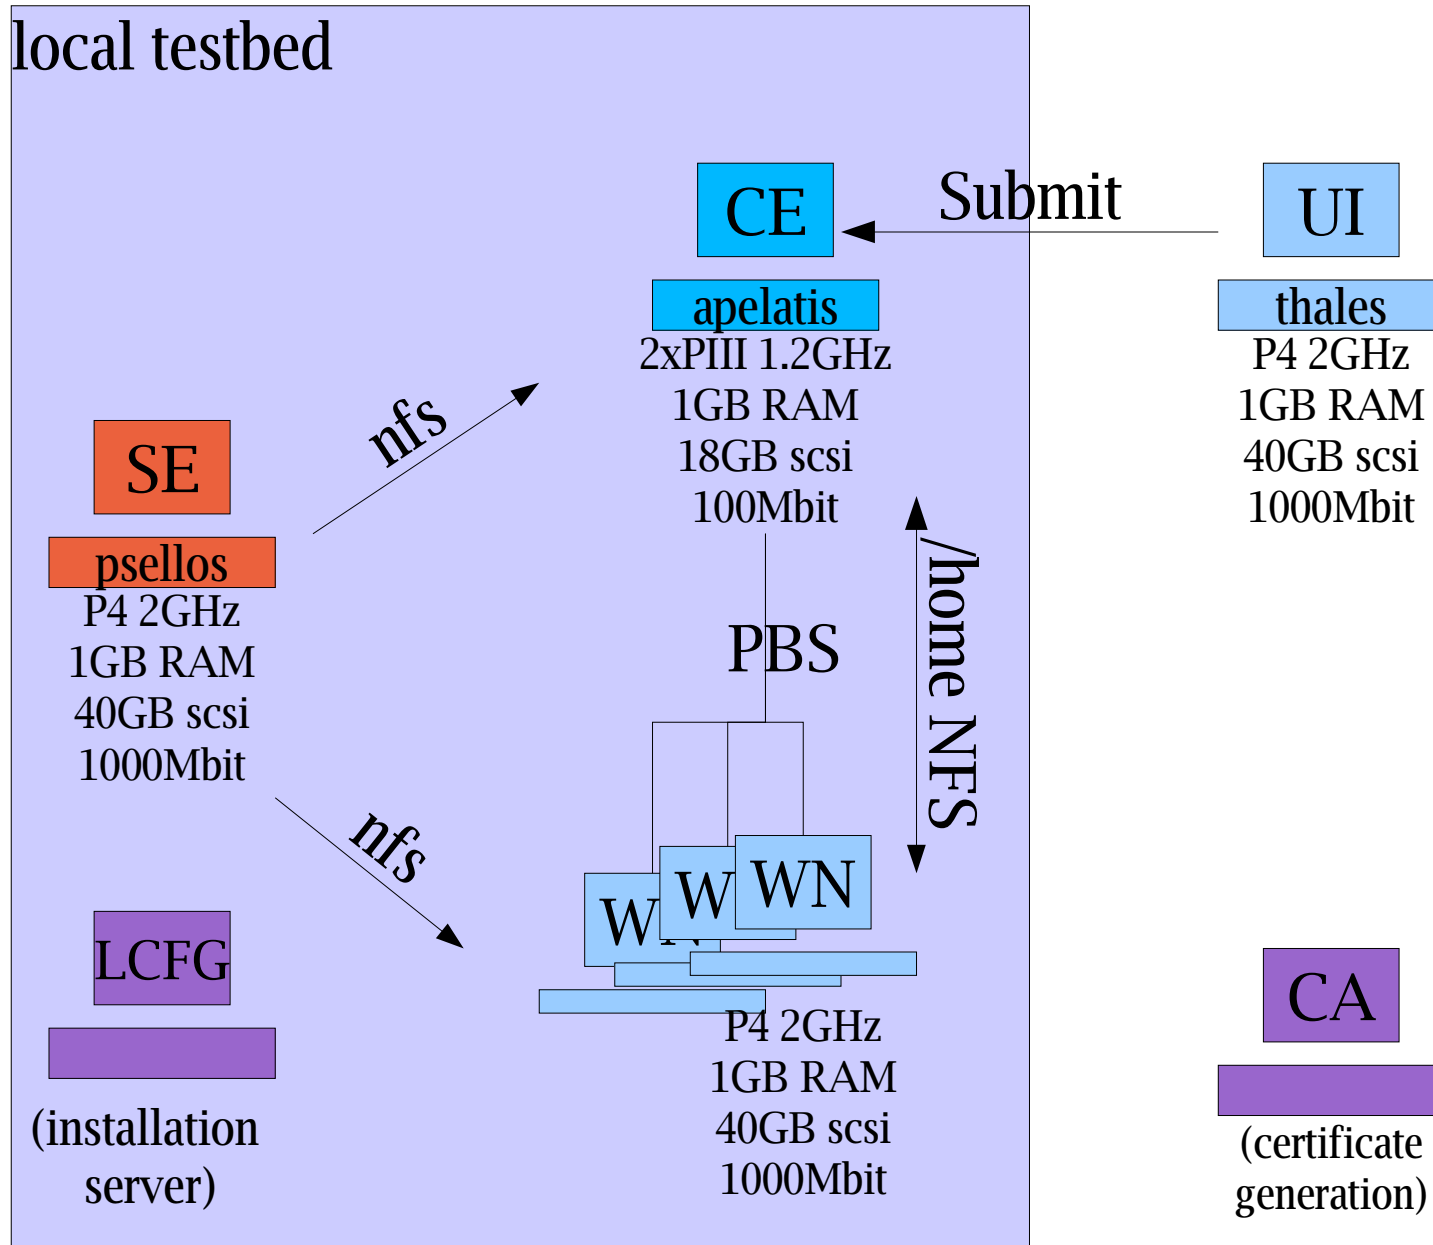
Crossgrid CA's

COUNTRY 	CPS	sLDAP	CA Cert	CA CRL	RPMS	TEST	STATUS	CENTER	Contact Person
 Slovakia	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	II SAS	Jan Astalos
 Cyprus	pdf,ps	sLDAP	CA Cert	CA CRL		CERTICRL	IN PROGRESS	UCY	Wei Xing
 Germany	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	FZK	Ursula Epting
 Greece	CPS	sLDAP	CA Cert	CA CRL	RPMS	CERTICRL	OK	AUTH	Christos Kanellopoulos
 Netherlands	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	NIKHEF(*)	David Groep
 Ireland	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	TCD	Brian Coghlan
 Poland	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	PSNC	Pawel Wolniewicz
 Portugal	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	LIP	Jorge Gomes
 Spain	CPS	sLDAP	CA Cert	CA CRL		CERTICRL	OK	IFCA	Rafael Marco

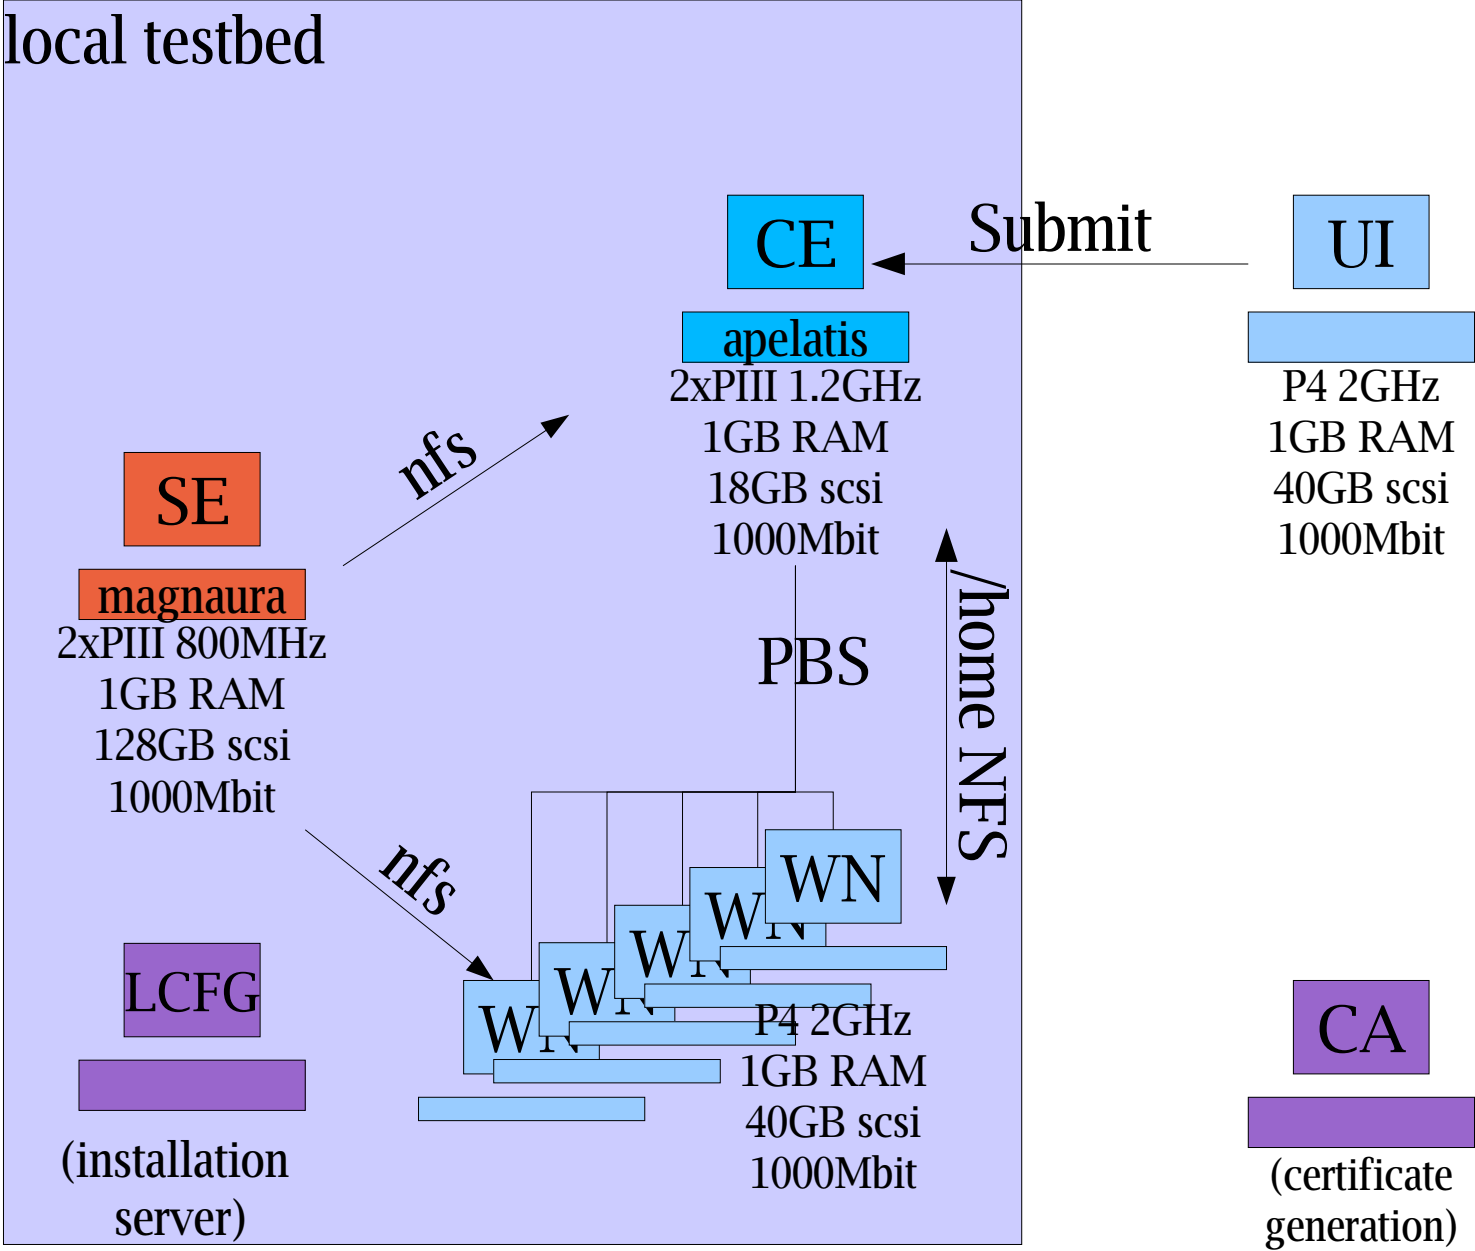
CURRENT local testbed status

- We currently have: (hostnames are on grid.ucy.ac.cy)
 - 1 CE (apelatis) 2xPIII 1.2GHz 1GB
 - 1 SE (psellos) P4 2GHz 1GB 40GB
 - soon will be magnaure 2xPIII 1GB 128GB scsi
 - 1 UI (thales) P4 2GHz 1GB 40GB
 - soon will be another machine (amalthea?)
 - 3 WN (soon 5+) P4 2GHz 1GB 40GB
 - CA
 - LCFG Intallation server
- Pending
 - RB subscription
 - MPICH-G2
 - PBS queues

CURRENT local testbed status



The local testbed status (planned)



New Accounts:

How do I get an “account” for the CrossGrid/Datagrid testbed?

1) Obtain a certificate from the CyGrid CA

- requests go to `cygrid-ca@ucy.ac.cy`

2) Get an account on the User Interface machine (from where you will submit job), `thales.grid.ucy.ac.cy`

3) To run jobs on:

1) Local site: get an account on the CE

2) Crossgrid: Become part of the crossgrid VO (currently you must be related to crossgrid for this)

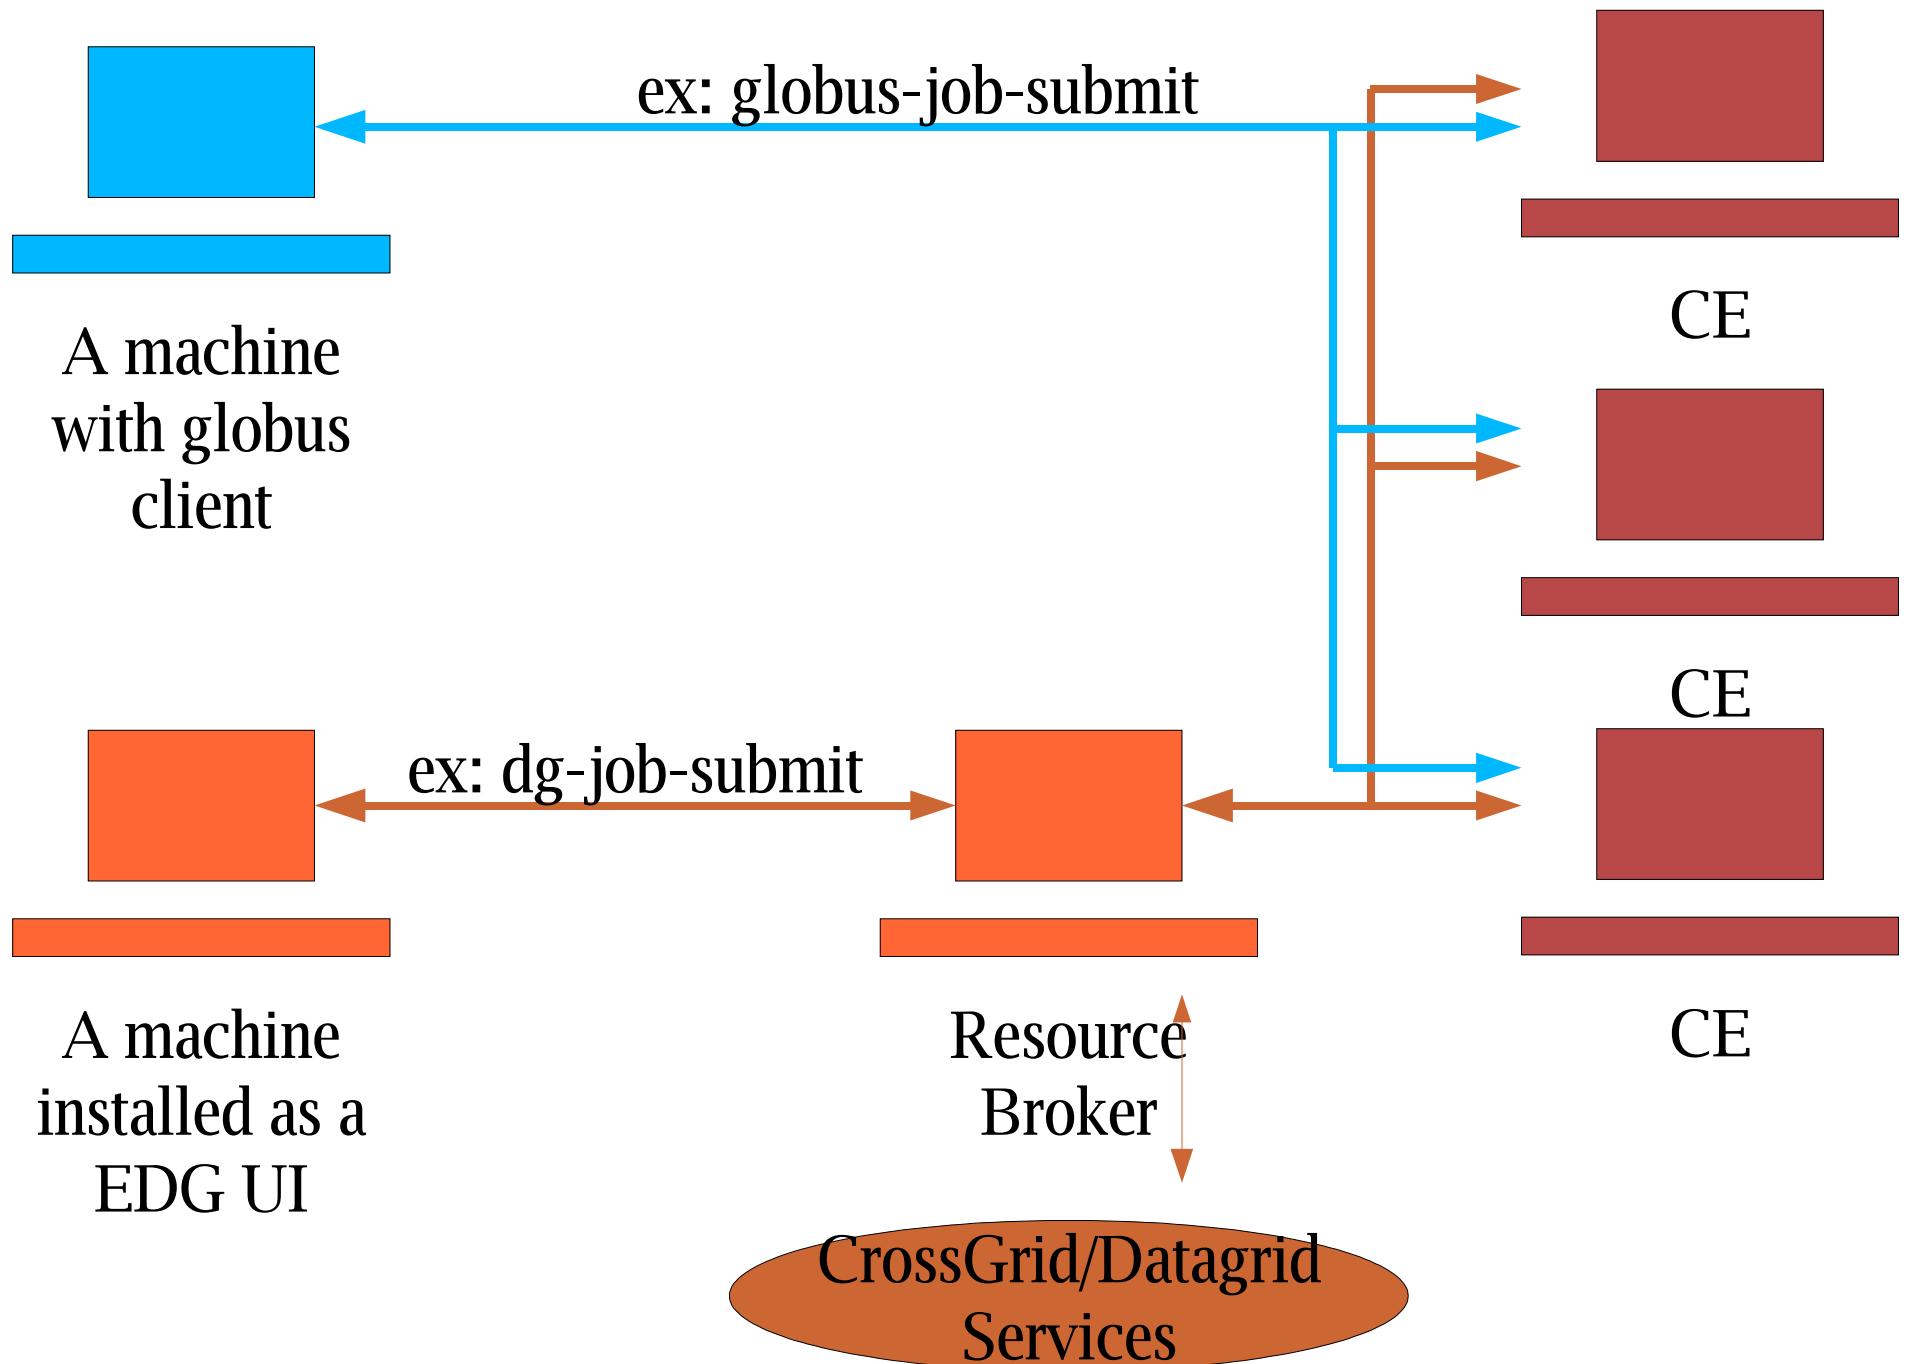
(Requests for both go to `hpcl@ucy.ac.cy`)

New Accounts:

How do I get an “account” for the CrossGrid/Datagrid testbed?

- 4) Setup your account on the UI machine with your cert and key (provided by the CyGrid CA)
 - ~/.globus/usercert.pem (444 permissions)
 - ~/.globus/userkey.pem (400 permissions)
- 5) SSH to `thales.grid.ucy.ac.cy` and start submitting jobs (see section on Job submission)

Job Submission



Basic (Globus) Job Submission

grid-proxy-init

globus-job-run host[/service] <executable [params] >

globus-job-run apelatis /bin/hostname

globus-job-run apelatis/jobmanager-pbs /bin/hostname

grid-proxy-destroy

```
[georget@thales georget]$ globus-job-run apelatis /bin/date
Thu Dec 12 01:36:38 EET 2002
[georget@thales georget]$ globus-job-submit apelatis /bin/date
https://apelatis.grid.ucy.ac.cy:4290/30892/1039649832/
[georget@thales georget]$ globus-job-get-output \
    https://apelatis.grid.ucy.ac.cy:4290/30892/1039649832/
Thu Dec 12 01:37:16 EET 2002
```

(Globus) Submission with RSL

ls.rsl:

```
&(executable = "/bin/ls" )  
  
(arguments= "/flatfiles" "-l")  
  
(stdout= "/home/georget/ls.output")
```

```
globusrun -b -s -r apelatis.grid.ucy.ac.cy -f hello.rsl
```

```
globusrun -status <jobID>
```

Staging the executable

hello.rsl:

```
&(executable=$ (GLOBUSRUN_GASS_URL) #"/home/georget/longjob.sh")
```

More Job Submission

```
[georget@thales georget]$ globusrun -b -s -r apelatis.grid.ucy.ac.cy -f hello.rsl
```

```
https://apelatis.grid.ucy.ac.cy:4499/2755/1039656224/
```

```
[georget@thales georget]$ globus-job-status /
```

```
https://apelatis.grid.ucy.ac.cy:4499/2755/1039656224/
```

```
ACTIVE
```

```
[georget@thales georget]$ globus-job-status /
```

```
https://apelatis.grid.ucy.ac.cy:4499/2755/1039656224/
```

```
ACTIVE
```

```
[georget@thales georget]$ globus-job-status /
```

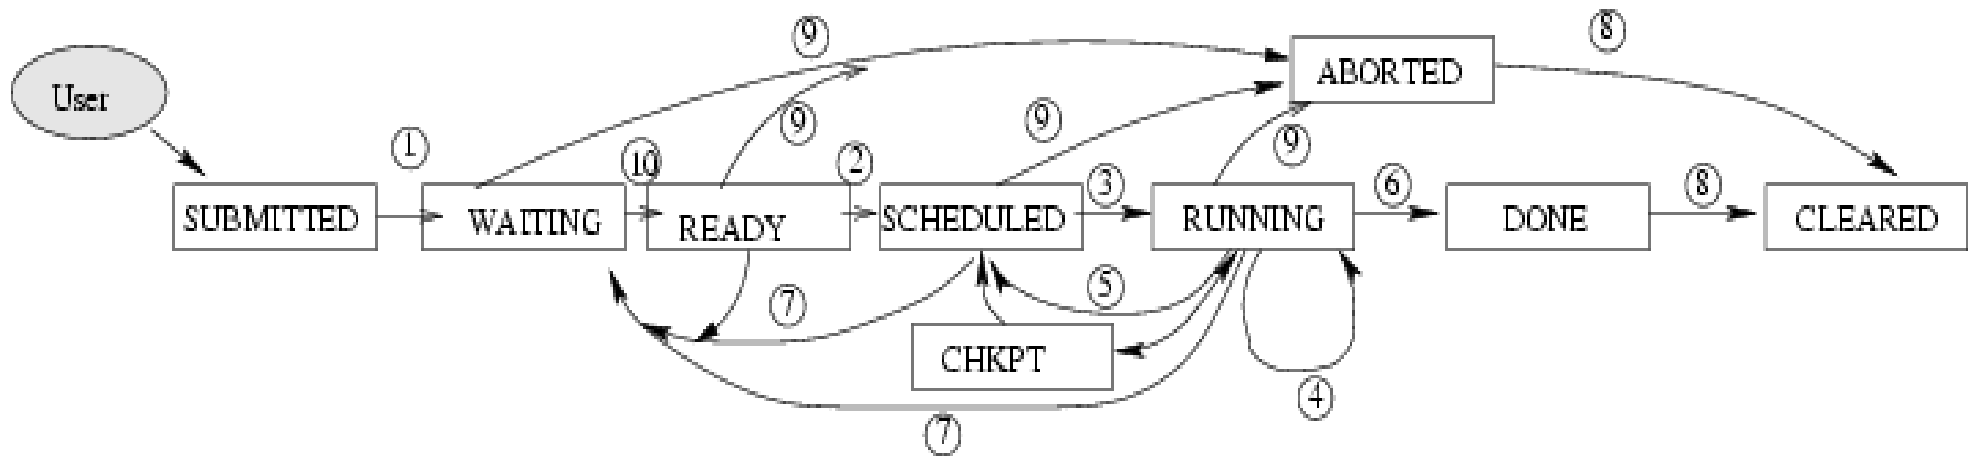
```
https://apelatis.grid.ucy.ac.cy:4499/2755/1039656224/
```

```
DONE
```


Using the Datagrid Tools and services TO BE DONE

Examples up to now made no use of EDG tools or
middleware.

EDG extends the Globus middleware to



READY A Computing Element matching job requirements has been selected.

SCHEDULED The Computing Element has received the job.

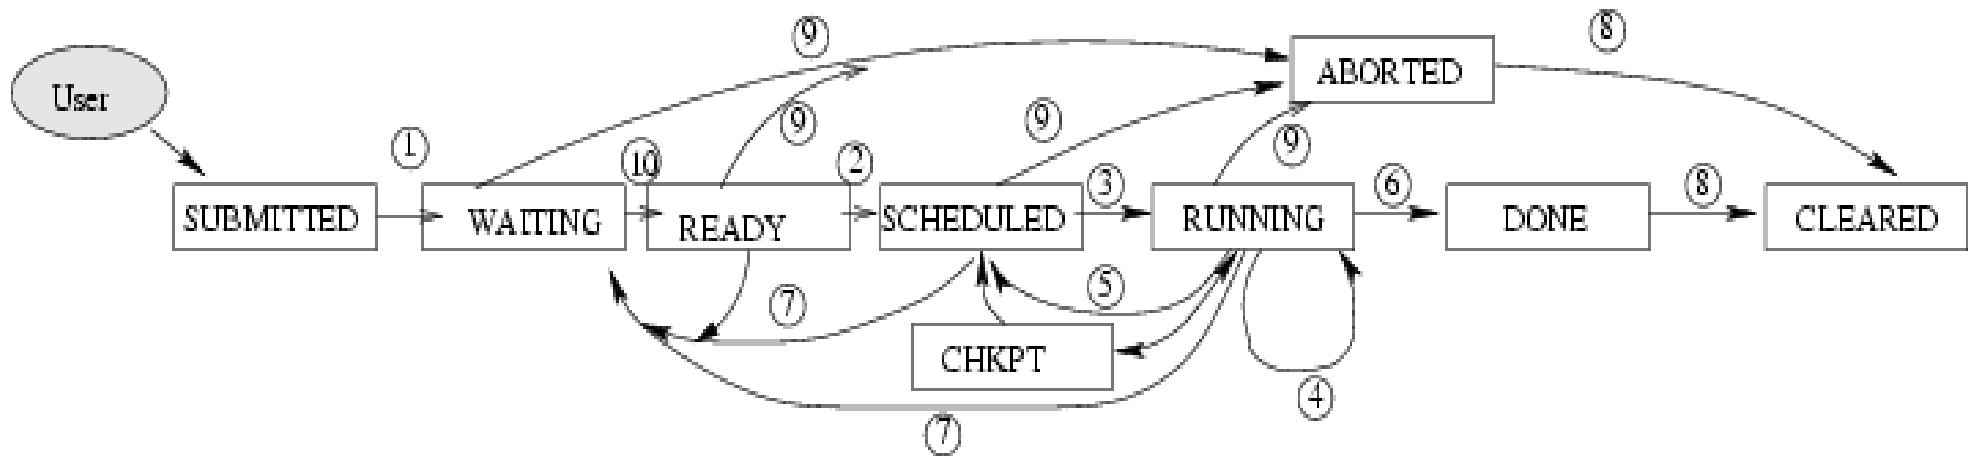
RUNNING The job is running on a Computing Element.

CHKPT The job has been suspended and check-pointed on a Computing Element.

DONE The execution of the job has completed.

ABORTED The job has been terminated.

CLEARED The user has retrieved all output files successfully. Bookkeeping information is purged some time after the job enters this state.



3 . The job is started on a CE.

4 . The job has been affected by an event caused by the CE local resource management system, e.g. it has been suspended or its priority has been lowered.

5 . The job is stopped and queued again in the CE.

6 . The job has completed.

7 . The job is passed again to the RB.

8 . The bookkeeping information about the job is cleared.

9 . The job is terminated.

10 . A suitable CE has been found where the job can run.