



Testbed1 Software

Ron Trompert

ron @ sara.nl

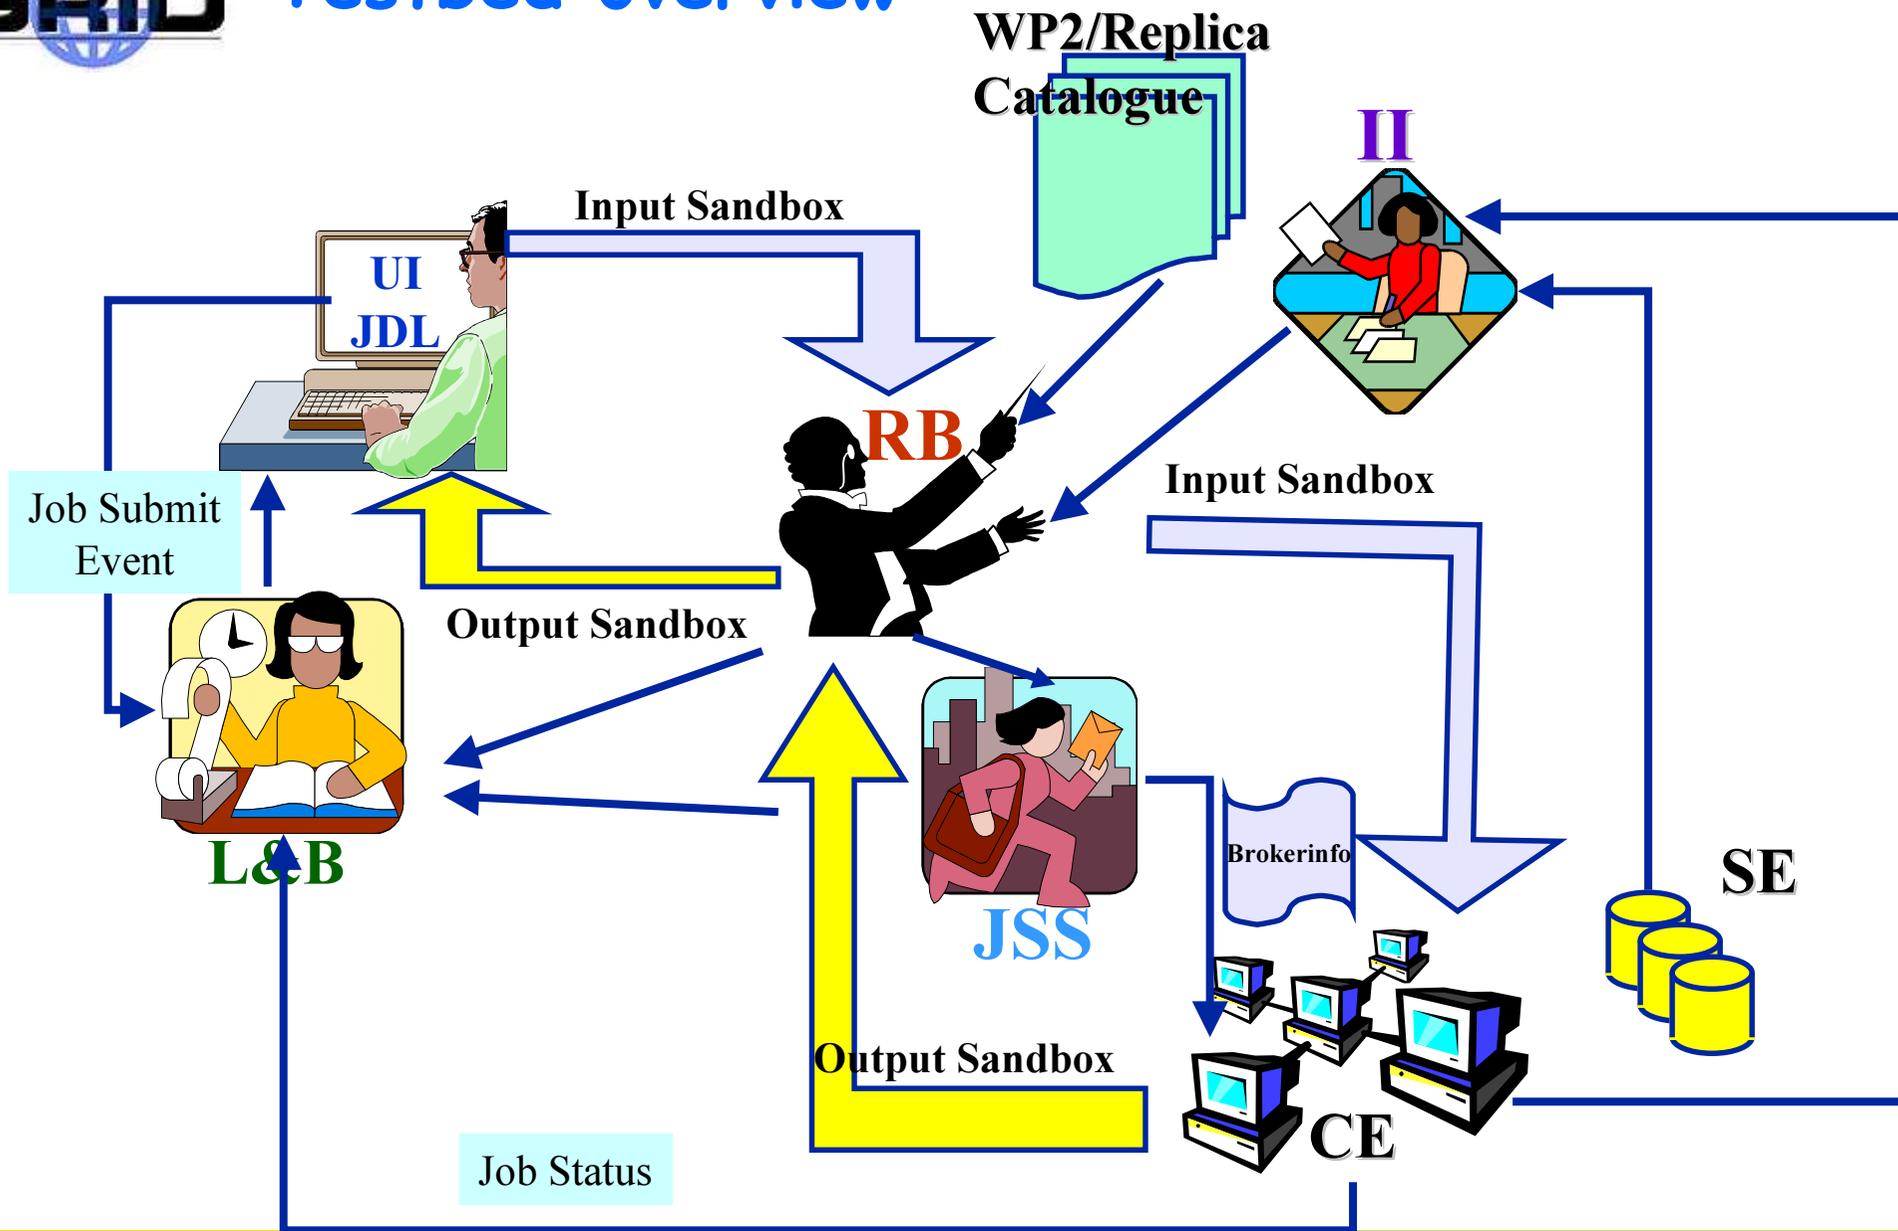


Data GRID Summary

- ◆ Testbed1 software presentation at CERN
 - Testbed overview
 - Job submission
 - User interface
 - Information system
 - WP4 stuff
 - EDG Globus Configuration
- ◆ Globus2.0 installation on a Compaq cluster
- ◆ SARA's contribution to testbed1



Testbed overview



Data GRID Job Submission

```
#
# ----- Job Description File -----
#
Executable = "WP1testC";
StdInput = "sim.dat";
StdOutput = "sim.out";
StdError = "sim.err";
InputSandbox = {"home/wp1/HandsOn-0409/WP1testC","home/wp1/HandsOn-0409/file*",
                "/home/wp1/DATA/*"};
OutputSandbox = {"sim.err","test.out","sim.out"};
Rank = other.AverageSI00;
Requirements = (other.OpSys == "Linux RH 6.1" || other.OpSys == "Linux RH 6.2") &&
(other.RunTimeEnvironmnet == "CMS3.2");
InputData = "LF:test10096-0009";
ReplicaCatalog = "ldap://sunlab2g.cnaf.infn.it:2010/rc=WP2 INFN Test Replica Catalog,dc=sunlab2g,
                 dc=cnaf, dc=infn, dc=it";
DataAccessProtocol = "gridftp";
```



Data GRID User interface

- ◆ **dg-job-submit** allows the user to submit a job for the execution on remote resources
 - **-r, -resource res_id** the job is submitted by the Broker to the resource identified by res_id
 - **-i, -input input_file** the user must choose a resource id from a list of resources
 - **-n, -notify e_mail_address** an e-mail message is sent to the specified e-mail address when the job enters in one of the following status: READY, RUNNING, ABORTED or DONE
 - **-o, -output out_file** the generated dg_jobid is written in the file out_file



- ◆ **dg-job-get-output** requests to the Broker the job output files, specified by the OutputSandbox attribute of the job-ad, and stores them on the submitting machine local disk
- ◆ **dg-job-list-match** returns the list of resources which fulfills job requirements
- ◆ **dg-job-cancel** cancels one or more submitted jobs
- ◆ **dg-job-status** displays bookkeeping information about submitted jobs
- ◆ **dg-job-get-logging-info** displays logging information about submitted jobs

Information system

- ◆ There are two to pick from
 - Globus MDS 2.1
 - Integrated in Globus
 - Memory caching, LDAPv3, GSI authentication
 - Multiple VO's on one node
 - GRIS and GIIS use the same slapd and listen to the same port. They are only distinguished through their DN's:
 - GIIS: ... Mds-Vo-Name=**VO_NAME**, o=Grid
 - GRIS: ... Mds-Vo-Name=**local**, o=Grid
 - OpenLDAP Ftree
 - Runs independently of globus
 - Memory caching, LDAPv3, GSI authentication
 - GRIS and GIIS on different ports



WP4 stuff

◆ Interim Installation System

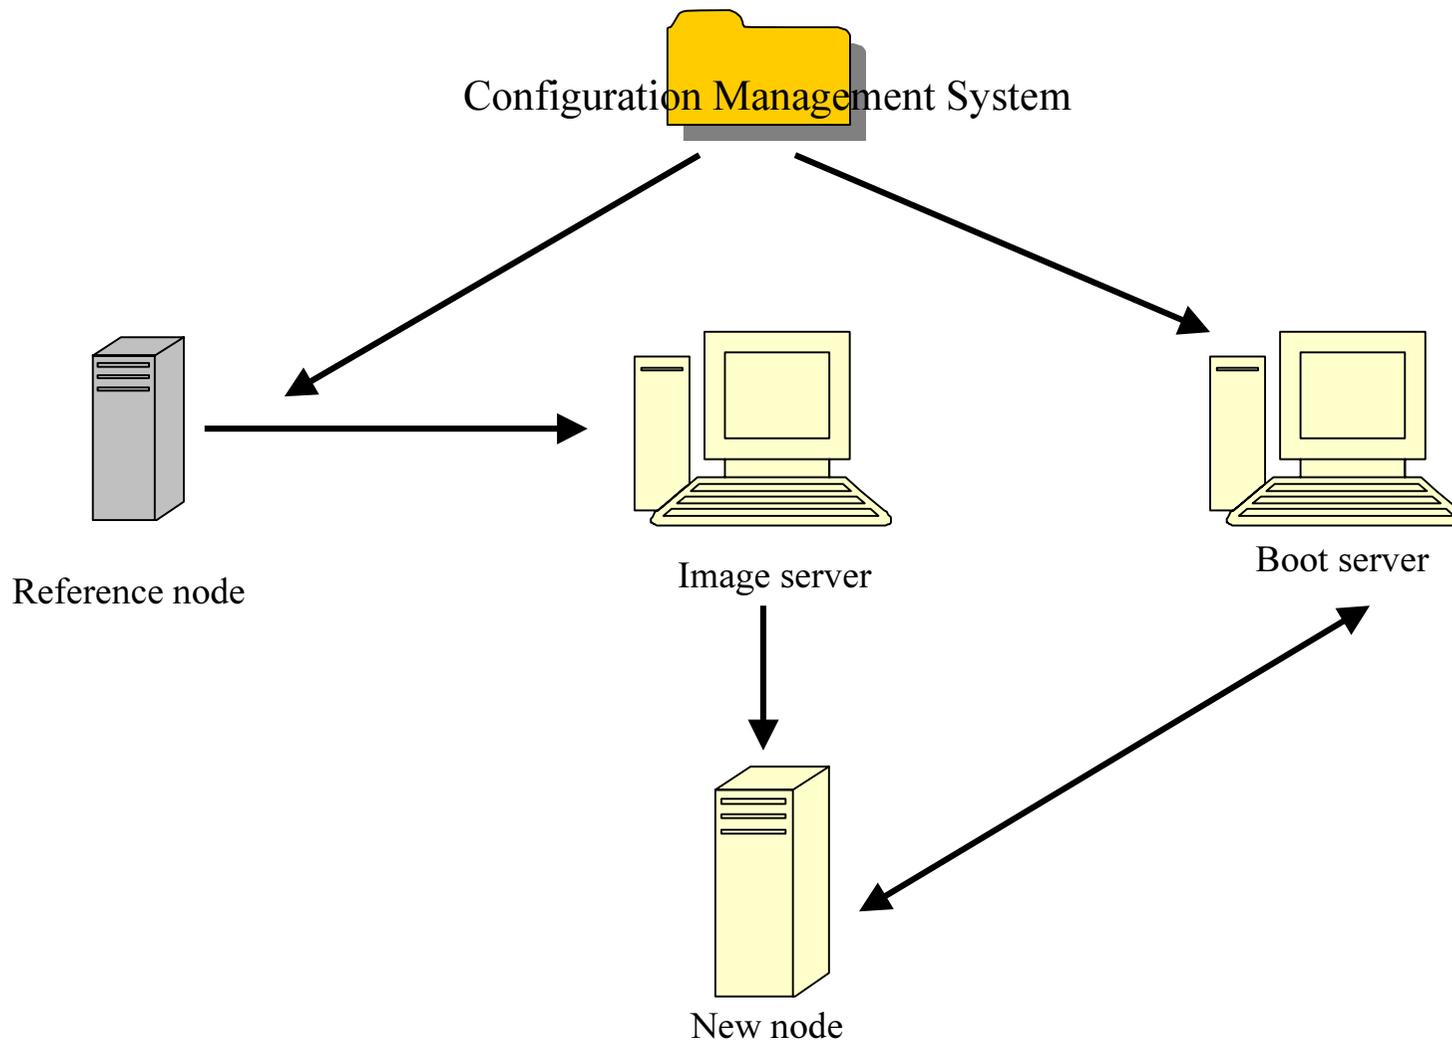
■ Image Cloning

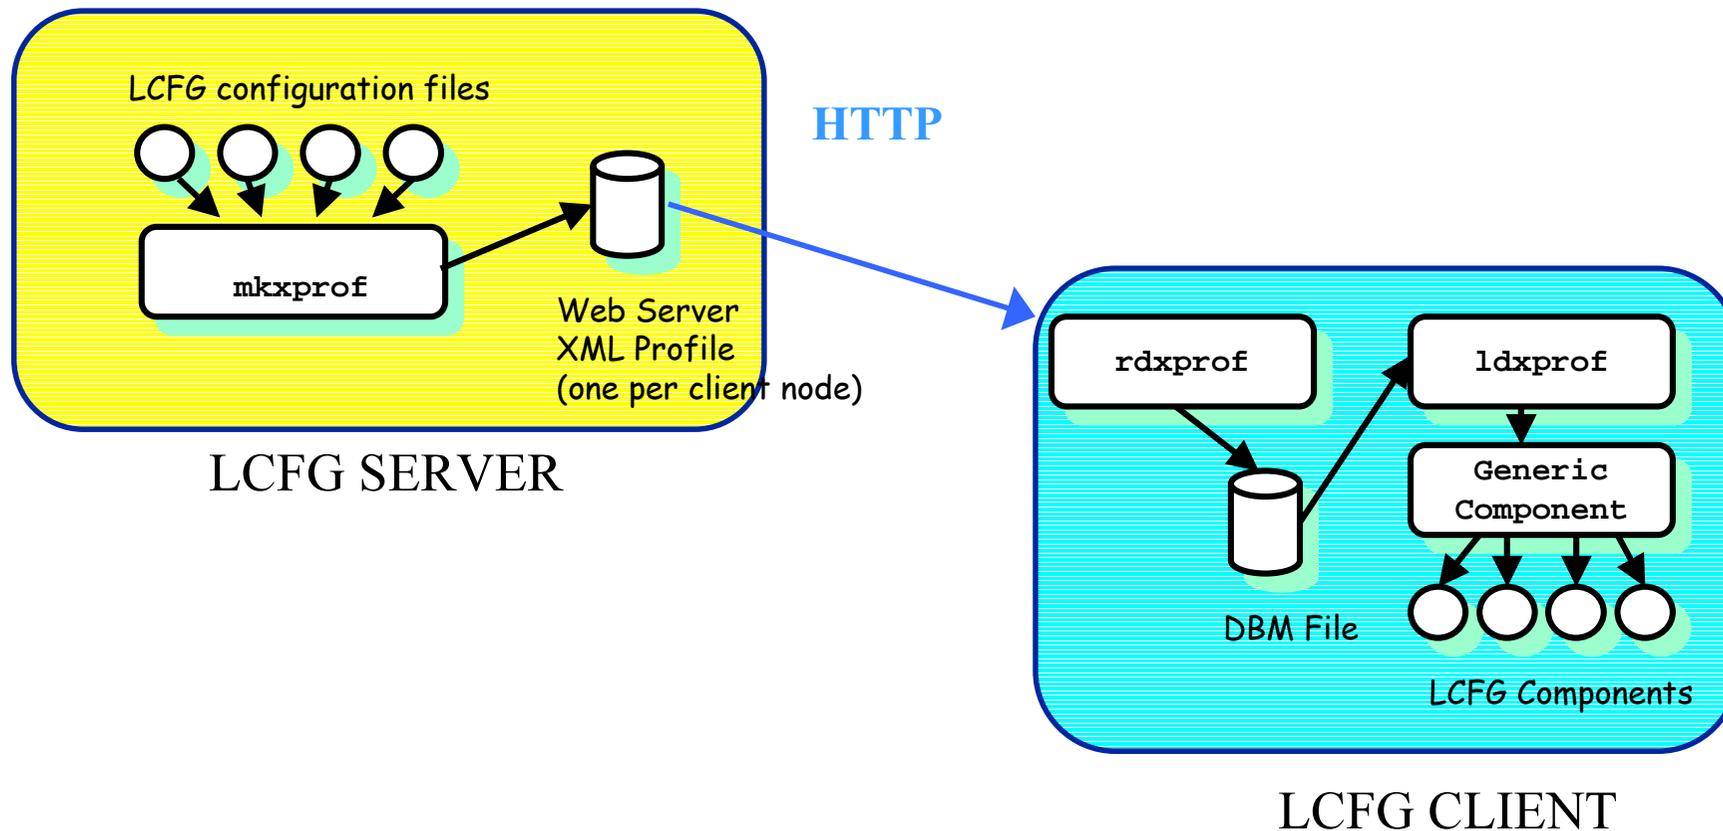
- *Quick and easy installation of identical machines in large clusters*

■ Local ConFiGuration system

- *Handle automated installation and configuration in a very diverse and evolving environment*
- *Cfengine-like*

Data GRID Image cloning







LCFG: RPM configuration

- ◆ The list of RPMs to be installed on a node is defined in text files stored in the LCFG server
- ◆ These RPM lists can be shared by several nodes; the LCFG configuration associates one list to each node
- ◆ These files must be accessible to the clients using a remote file system (NFS)
- ◆ The RPMs themselves must be also accessible via NFS
- ◆ Any RPM installed manually in the client without being included in the files will be **AUTOMATICALLY REMOVED**



EDG Globus configuration

- ◆ Everything is installed in `$GLOBUS_LOCATION` i.e. `/opt/globus`
- ◆ Security files reside in `/etc/grid-security`
 - CA certificates, CRLs and signing policy files in `/etc/grid-security/certificates`
 - Grid-mapfile, host key and host certificate in `/etc/grid-security`
- ◆ Globus setup packages are not used. They are replaced by the EDG config packages `globus_<package>-edgconfig` where `<package>` is the name of the globus setup package it replaces



- ◆ A single configuration file `/etc/globus.conf`
 - Local (internal) configuration files are (re)created by each restart of service with parameters from `/etc/globus.conf`
 - Simple format:
 - `#` at start of line means a comment
 - `MACRO=VALUE`
 - Example:
 - `GLOBUS_LOCATION=/opt/globus`



Globus2.0 installation on a Compaq cluster

◆ Build_from_source_tarballs

- Perl script to build the distribution
- Site specific information should be set:
 - `my $SRC_TAR_LOC = " ... ";`
 - `my $SRC_LOC = " ...";`
 - `my $GLOBUS_LOC = " ... ";`
 - `my $FLAVOR = "gcc32dbg";`
 - Although the Compaq is a 64-bit machine 32-bit was specified here because when you use 64, the scripts think that you are on a SGI machine and invoke the MIPSpro `--enable-64bit` compiler option. 32 does nothing.
 - Haven't tried vendorcc (ccc) yet
 - `my $LOG_LOC = " ... ";`
 - `my $TMP_DIR = " ... ";`



◆ Compilation

■ Globus-ssl-utils-2.1

- Error message: **undefined reference to '_OtsRemainder64Unsigned'**
 - Linking with **-lots** by editing the configure script
- Warning: **gnu_regex.c:cast from pointer to integer of different size**
 - On the Compaq this did not seem to cause any problems. The machine only has 512Mb of memory and tests showed that the high 32 bits of 64-bit pointers did not change.

■ Globus_gatekeeper.c & globus_gram_job_manager.c

- Error: **redefinition of sys_errlist[]**
 - For some reason "TARGET_ARCH_LINUX" was not **#define**-ed. A dirty hack solved the problem.

◆ Globus-script-pbs-submit

- **#PBS -l ncpus=\$grami_count** should be **#PBS -l nodes=\$grami_count**
- Set **is_cluster=true**



- ◆ Tell Globus where MPI is in `globus-job-manager-tools.sh` and `globus-sh-tools.sh`
- ◆ `LD_LIBRARY_PATH` containing `$GLOBUS_LOCATION/lib` for `globus-gatekeeper` and `in.ftpd`
 - Solved by wrapper



SARA's contribution to testbed1

◆ Resources

■ Compute Elements

- ◆ Compact
 - . 16 nodes 600MHz alpha
 - . Myrinet 2000 interconnect
 - . Fast Ethernet interconnect
 - . OpenPBS
 - . GridFTP
 - . Jobmanagers fork and pbs
 - . MPI over myrinet and fast Ethernet. MPI over myrinet is the default, MPI over Fast Ethernet only through RSL.

- ◆ Itanium cluster?





SARA's contribution to the testbed1

- Storage Element
 - Teras with StorageTek NearLine storage
 - . 10TB disk
 - . 120 TB tape
 - . GridFTP
 - . No jobmanagers

