CENTRE DE SERVEIS CIENTÍFICS
I ACADÈMICS DE CATALUNYA

CESCA | 20 ANYS 1991·2011

Computing at CESCA: A Quick Guide

System description

The computing facilities at CESCA are composed of three different systems, but configured as a unique LSF cluster. With this setup, the batch system is in charge of deciding which computing node will execute the jobs, depending on the demanded resources. Therefore, jobs can be submitted from any machine, and will be dispatched at the proper node as soon as resources will be available.

Machine	Description
HP CP4000	132 cores (Opteron 275) 528 GB of main memory (distributed memory) 9.41 TB of disk
Bull NovaScale	224 cores (Xeon E5472) 112 cores (Xeon X5550) 1.6 TB of main memory (distributed memory) 52.6 TB of disk
SGI Altix UV 1000	1344 cores (Xeon X7542) 6.14 TB of main memory (shared memory) 112 TB of disk

Log in

Access to the system can be achieved by using ssh protocol, **through port number 2122**, (i.e.: ssh user@machinename.cesca.cat –p 2122) with the following names or IP addresses:

Machine	Name	IP Address
CP4000	cadi.cesca.cat	84.88.8.30
NovaScale	prades.cesca.cat	84.88.8.66
Altix UV 1000	pirineus.cesca.cat	84.88.8.106

If you have problems to login or do not remember the password, please contact CESCA support staff at *consultes@cesca.cat* or at phone number 932055133.

Transferring files

In order to transfer files to CESCA systems (and not from it, as outgoing ssh connections are not allowed) you can use either scp or sftp command:

scp –P2122 local_file user@hostname.cesca.cat

sftp –oPort=2122 user@hostname.cesca.cat

where hostname can be *cadi*, *prades* or *pirineus*.

Working areas

The working areas available at CESCA facilities are summarized below:

Name	Variable	Availability	Quota	Time limit	Backup
/home/user	$HOME	global	4 GB (*)	unlimited	Yes
/cescascratch/user	$SCRATCH	global	-	30 days	No
/dades/user	$DADES	global	variable	unlimited	No
/tmp/user/jobid.date	$TMPDIR	Local to each node	Job file limit	Duration of the job	No

*Additionally there is a group quota, depending on the number of users. The limits are 4, 8 and 16 GB, for groups up to 5, 10 and 20 users, respectively.

Areas globally available are shared by all the machines and they are offered by a NetApp FAS3170 NFS server.

Submitting jobs

In order to improve scheduling and reduce job's waiting time, a structure of queues has been defined, used in the framework of the LSF batch manager. The following table summarizes all the queues features, with their cpu time, number of cores, memory and disk limits:

Queue	Time (h)	Memory (GB)	Disk (GB)	Processors	Priority
short	0,5	4	40	4	70
medium	120	4	40	1	20
large	360	4	40	1	20
xlarge	2.400	4	40	1	20
memory	360	16	200	1	30
disk	360	8	500	1	30
parallel4	240 x n	4 x n	40 x n	2-4	40
parallel8	240 x n	4 x n	40 x n	8	40

Queue	Time (h)	Memory (GB)	Disk (GB)	Processors	Priority
parallel16	240 x n	4 x n	40 x n	16	60
parallel32	240 x n	4 x n	60 x n	16-32	60
parallel256	240 x n	4 x n	60 x n	32	60
services	∞	288	3.000	224	50

To submit jobs with more than 256 cores or very high resources demanding ones (in terms of memory, disk or time) you should use the *services* queue. This is a special queue, and you need to be authorized in order to use it.

Jobs are submitted to the batch system with the bsub command:

bsub **<** script-file (**the redirection symbol is mandatory**)

There exist some examples of script files for different scientific software at /usr/local/examples on each machine.

Other useful commands and options to monitor jobs are:

bjobs -x --- prints information about sent jobs, where ***x*** can be:

-a	all working jobs
-l	long format
-d	jobs finished in the last hour
-q queue	only selected queue

bhist -x --- historical information about jobs, where ***x*** can be:

-a	all working and finished jobs
-l	long format
-d queue	only finished jobs
-u username	only selected user

bkill jobId --- cancel sent jobs

bstop jobID --- temporarily stop, job is not cancelled

bresume jobID --- resuming stopped jobs

More information about the batch system options can be obtained at *consultes@cesca.cat.*