

Nord III User's Guide



Barcelona Supercomputing Center

Copyright © 2017 BSC-CNS

March 25, 2020

Contents

1	Introduction	2
2	System Overview	2
3	Connecting to Nord III	2
3.1	Password Management	3
3.2	Transferring files	3
3.3	Active Archive Management	5
3.4	Repository management (GIT/SVN)	6
4	File Systems	7
4.1	Root Filesystem	7
4.2	GPFS Filesystem	7
4.3	Local Hard Drive	8
4.4	Quotas	8
5	Running Jobs	8
5.1	LSF Commands	8
5.2	Interactive Sessions	8
5.3	Job directives	9
5.4	MPI particulars	10
5.5	Jobscript examples	11
5.6	Queues	13
6	Software Environment	14
6.1	C Compilers	14
6.2	FORTRAN Compilers	15
6.3	Modules Environment	16
6.4	BSC Commands	16
6.5	TotalView	17
6.6	Tracing jobs with BSC Tools	17
7	Getting help	18
7.1	Frequently Asked Questions (FAQ)	19
8	Appendices	19
8.1	SSH	19
8.2	Transferring files	19
8.3	Using X11	21

1 Introduction

This user's guide for the Nord III cluster is intended to provide the minimum amount of information needed by a new user of this system. As such, it assumes that the user is familiar with many of the standard features of supercomputing as the Unix operating system.

Here you can find most of the information you need to use our computing resources and the technical documentation about the machine. Please read carefully this document and if any doubt arises do not hesitate to contact us (Getting help (chapter 7)).

2 System Overview

Nord III is a supercomputer based on Intel SandyBridge processors, iDataPlex Compute Racks, a Linux Operating System and an Infiniband interconnection.

The current Peak Performance is 251,6 Teraflops. The total number of processors is 12,096 Intel SandyBridge-EP E5-2670 cores at 2.6 GHz (756 compute nodes) with at least 24.2 TB of main memory. See below a summary of the system:

- 9 iDataPlex compute racks. Each one composed of:
 - 84 IBM dx360 M4 compute nodes
 - 4 Mellanox 36-port Managed FDR10 IB Switches
 - 2 BNT RackSwitch G8052F (Management Network)
 - 2 BNT RackSwitch G8052F (GPFS Network)
 - 4 Power Distribution Units
- All IBM dx360 M4 node contain:
 - 2x E5-2670 SandyBridge-EP 2.6GHz cache 20MB 8-core
 - 500GB 7200 rpm SATA II local HDD
 - One of the following RAM configurations:
 - * Default nodes: 32 GB/node
 - * Medium memory nodes: 64 GB/node
 - * High memory nodes: 128 GB/node
- 15 PB of GPFS disk storage
- Interconnection Networks
 - Infiniband Mellanox FDR10: High bandwidth network used by parallel applications communications (MPI)
 - Gigabit Ethernet: 10GbitEthernet network used by the GPFS Filesystem.
- Operating System: Linux - SuSe Distribution 11 SP3

As we said before, there are 3 different types of nodes, each one with a different amount of memory available.

More information about the request of the different node types can be found in the "Job Directives" section.

3 Connecting to Nord III

The first thing you should know is your username and password. Once you have a login and its associated password you can get into the cluster through one of the following login nodes:

- nord1.bsc.es
- nord2.bsc.es
- nord3.bsc.es

You must use Secure Shell (ssh) tools to login into or transfer files into the cluster. We do not accept incoming connections from protocols like telnet, ftp, rlogin, rcp, or rsh commands. Once you have logged into the cluster you cannot make outgoing connections for security reasons.

To get more information about the supported secure shell version and how to get ssh for your system (including windows systems) see the Appendices (chapter 8). Once connected to the machine, you will be presented with a UNIX shell prompt and you will normally be in your home (\$HOME) directory. If you are new to UNIX, you will need to learn the basics before doing anything useful.

3.1 Password Management

In order to change the password, you have to login to a different machine (dt01.bsc.es). This connection must be established from your local machine.

```
% ssh -l username dt01.bsc.es

username@dttransfer1:~> passwd
Changing password for username.
Old Password:
New Password:
Reenter New Password:
Password changed.
```

Mind that that the password change takes about 10 minutes to be effective.

3.2 Transferring files

There are two ways to copy files from/to the Cluster:

- Direct scp or sftp to the login nodes
- Using a Data transfer Machine which shares all the GPFS filesystem for transferring large files

Direct copy to the login nodes.

As said before no connections are allowed from inside the cluster to the outside world, so all scp and sftp commands have to be executed from your local machines and never from the cluster. The usage examples are in the next section.

On a Windows system, most of the secure shell clients come with a tool to make secure copies or secure ftp's. There are several tools that accomplish the requirements, please refer to the Appendices (chapter 8), where you will find the most common ones and examples of use.

Data Transfer Machine

We provide special machines for file transfer (required for large amounts of data). These machines are dedicated to Data Transfer and are accessible through ssh with the same account credentials as the cluster. They are:

- dt01.bsc.es
- dt02.bsc.es

These machines share the GPFS filesystem with all other BSC HPC machines. Besides scp and sftp, they allow some other useful transfer protocols:

- scp

```
localsystem$ scp localfile username@dt01.bsc.es:
username's password:

localsystem$ scp username@dt01.bsc.es:remotefile localdir
username's password:
```

- rsync

```
localsystem$ rsync -avzP localfile_or_localdir username@dt01.bsc.es:
username's password:

localsystem$ rsync -avzP username@dt01.bsc.es:remotefile_or_remotedir localdir
username's password:
```

- sftp

```
localsystem$ sftp username@dt01.bsc.es
username's password:
sftp> get remotefile

localsystem$ sftp username@dt01.bsc.es
username's password:
sftp> put localfile
```

- BSCP

```
bbcp -V -z <USER>@dt01.bsc.es:<FILE> <DEST>
bbcp -V <ORIG> <USER>@dt01.bsc.es:<DEST>
```

- FTPS

```
gftp-text ftps://<USER>@dt01.bsc.es
get <FILE>
put <FILE>
```

- GRIDFTP (only accessible from dt02.bsc.es)
- SSHFTP

```
globus-url-copy -help
globus-url-copy -tcp-bs 16M -bs 16M -v -vb your_file sshftp://your_user@dt01.bsc.es/~ /
```

Data Transfer on the PRACE Network

PRACE users can use the 10Gbps PRACE Network for moving large data among PRACE sites. To get access to this service it's required to contact "support@bsc.es" requesting its use, providing the local IP of the machine from where it will be used.

The selected data transfer tool is Globus/GridFTP¹ which is available on dt02.bsc.es. In order to use it, a PRACE user must get access to dt02.bsc.es:

```
% ssh -l pr1eXXXX dt02.bsc.es
```

Load the PRACE environment with 'module' tool:

```
% module load prace globus
```

Create a proxy certificate using 'grid-proxy-init':

¹<http://www.globus.org/toolkit/docs/latest-stable/gridftp/>

```
% grid-proxy-init
Your identity: /DC=es/DC=irisgrid/O=bsc-cns/CN=john.foo
Enter GRID pass phrase for this identity:
Creating proxy ..... Done
Your proxy is valid until: Wed Aug 7 00:37:26 2013
pr1eXXXX@dttransfer2:~ >
```

The command 'globus-url-copy' is now available for transferring large data.

```
globus-url-copy [-p <parallelism>] [-tcp-bs <size>] <sourceURL> <destURL>
```

Where:

- -p: specify the number of parallel data connections should be used (recommended value: 4)
- -tcp-bs: specify the size (in bytes) of the buffer to be used by the underlying ftp data channels (recommended value: 4MB)
- Common formats for sourceURL and destURL are:
 - file://(on a local machine only) (e.g. file:///home/pr1eXX00/pr1eXXXX/myfile)
 - gsiftp://(e.g. gsiftp://supermuc.lrz.de/home/pr1dXXXX/mydir/)
 - remember that any url specifying a directory must end with /.

All the available PRACE GridFTP endpoints can be retrieved with the 'prace_service' script:

```
% prace_service -i -f bsc
gftp.prace.bsc.es:2811
```

More information is available at the PRACE website²

3.3 Active Archive Management

Active Archive (AA) is a mid-long term storage filesystem that provides 3.7 PB of total space. You can access AA from the Data Transfer Machine (section 3.2) (dt01.bsc.es and dt02.bsc.es) under /gpfs/archive/your_group.

NOTE: There is no backup of this filesystem. The user is responsible for adequately managing the data stored in it.

To move or copy from/to AA you have to use our special commands:

- dtcp, dtmv, dtsync, dttar

These commands submit a job into a special class performing the selected command. Their syntax is the same than the shell command without 'dt' prefix (cp, mv, rsync, tar).

- dtq, dtcancel

dtq shows all the transfer jobs that belong to you. (works like mnq)
dtcancel works like mncancel (see below) for transfer jobs.

- *dttar*: submits a tar command to queues. Example: Taring data from /gpfs/to /gpfs/archive

```
% dttar -cvf /gpfs/archive/usertest/outputs.tar ~/OUTPUTS
```

- *dtcp*: submits a cp command to queues. Remember to delete the data in the source filesystem once copied to AA to avoid duplicated data.

²<http://www.prace-ri.eu/Data-Transfer-with-GridFTP-Details>

```
# Example: Copying data from /gpfs to /gpfs/archive
% dtcp -r ~/OUTPUTS /gpfs/archive/usertest/
```

```
# Example: Copying data from /gpfs/archive to /gpfs
% dtcp -r /gpfs/archive/usertest/OUTPUTS ~/
```

- *dtrsync*: submits a rsync command to queues. Remember to delete the data in the source filesystem once copied to AA to avoid duplicated data.

```
# Example: Copying data from /gpfs to /gpfs/archive
% dtrsync -avP ~/OUTPUTS /gpfs/archive/usertest/
```

```
# Example: Copying data from /gpfs/archive to /gpfs
% dtrsync -avP /gpfs/archive/usertest/OUTPUTS ~/
```

- *dtmv*: submits a mv command to queues.

```
# Example: Moving data from /gpfs to /gpfs/archive
% dtmv ~/OUTPUTS /gpfs/archive/usertest/
```

```
# Example: Moving data from /gpfs/archive to /gpfs
% dtmv /gpfs/archive/usertest/OUTPUTS ~/
```

Additionally, these commands accept the following options:

```
--blocking: Block any process from reading file at final destination until transfer completed.
--time: Set up new maximum transfer time (Default is 18h).
```

It is important to note that these kind of jobs can be submitted from both the ‘login’ nodes (automatic file management within a production job) and ‘dt01.bsc.es’ machine. AA is only mounted in Data Transfer Machine (section 3.2). Therefore if you wish to navigate through AA directory tree you have to login into dt01.bsc.es

3.4 Repository management (GIT/SVN)

There’s no outgoing internet connection from the cluster, which prevents the use of external repositories directly from our machines. To circumvent that, you can use the “sshfs” command in your local machine.

Doing that, you can mount a desired directory from our GPFS filesystem in your local machine. That way, you can operate your GPFS files as if they were stored in your local computer. That includes the use of git, so you can clone, push or pull any desired repositories inside that mount point and the changes will transfer over to GPFS.

Setting up sshfs

- Create a directory inside your local machine that will be used as a mount point.
- Run the following command below, where the local directory is the directory you created earlier. Note that this command mounts your GPFS home directory by default.

```
sshfs -o workaround=rename <yourHPCUser>@dt01.bsc.es: <localDirectory>
```

- From now on, you can access that directory. If you access it, you should see your home directory of the GPFS filesystem. Any modifications that you do inside that directory will be replicated to the GPFS filesystem inside the HPC machines.
- Inside that directory, you can call “git clone”, “git pull” or “git push” as you please.

4 File Systems

IMPORTANT: It is your responsibility as a user of our facilities to backup all your critical data. *We only guarantee a daily backup of user data under /gpfs/home. Any other backup should only be done exceptionally under demand of the interested user.*

Each user has several areas of disk space for storing files. These areas may have size or time limits, please read carefully all this section to know about the policy of usage of each of these filesystems. There are 3 different types of storage available inside a node:

- *Root filesystem:* Is the filesystem where the operating system resides
- *GPFS filesystems:* GPFS is a distributed networked filesystem which can be accessed from all the nodes and Data Transfer Machine (section 3.2)
- *Local hard drive:* Every node has an internal hard drive

4.1 Root Filesystem

The root file system, where the operating system is stored doesn't reside in the node, this is a NFS filesystem mounted from one of the servers.

As this is a remote filesystem only data from the operating system has to reside in this filesystem. It is NOT permitted the use of /tmp for temporary user data. The local hard drive can be used for this purpose as you could read in Local Hard Drive (section 4.3).

4.2 GPFS Filesystem

The IBM General Parallel File System (GPFS) is a high-performance shared-disk file system providing fast, reliable data access from all nodes of the cluster to a global filesystem. GPFS allows parallel applications simultaneous access to a set of files (even a single file) from any node that has the GPFS file system mounted while providing a high level of control over all file system operations. In addition, GPFS can read or write large blocks of data in a single I/O operation, thereby minimizing overhead.

An incremental backup will be performed daily only for /gpfs/home.

These are the GPFS filesystems available in the machine from all nodes:

- */apps:* Over this filesystem will reside the applications and libraries that have already been installed on the machine. Take a look at the directories to know the applications available for general use.
- */gpfs/home:* This filesystem has the home directories of all the users, and when you log in you start in your home directory by default. Every user will have their own home directory to store own developed sources and their personal data. A default quota (section 4.4) will be enforced on all users to limit the amount of data stored there. Also, it is highly discouraged to run jobs from this filesystem. **Please run your jobs on your group's /gpfs/projects or /gpfs/scratch instead.**
- */gpfs/projects:* In addition to the home directory, there is a directory in /gpfs/projects for each group of users. For instance, the group bsc01 will have a /gpfs/projects/bsc01 directory ready to use. This space is intended to store data that needs to be shared between the users of the same group or project. A quota (section 4.4) per group will be enforced depending on the space assigned by Access Committee. It is the project's manager responsibility to determine and coordinate the better use of this space, and how it is distributed or shared between their users.
- */gpfs/scratch:* Each user will have a directory over /gpfs/scratch. Its intended use is to store temporary files of your jobs during their execution. A quota (section 4.4) per group will be enforced depending on the space assigned.

4.3 Local Hard Drive

Every node has a local hard drive that can be used as a local scratch space to store temporary files during executions of one of your jobs. This space is mounted over `/scratch/tmp` directory and pointed out by `$TMPDIR` environment variable. The amount of space within the `/scratch` filesystem is about 500 GB. All data stored in these local hard drives at the compute nodes will not be available from the login nodes. Local hard drive data are not automatically removed, so each job has to remove its data before finishing.

4.4 Quotas

The quotas are the amount of storage available for a user or a groups' users. You can picture it as a small disk readily available to you. A default value is applied to all users and groups and cannot be outgrown.

You can inspect your quota anytime you want using the following command from inside each filesystem:

```
% bsc_quota
```

The command provides a readable output for the quota. Check BSC Commands (section 6.4) for more information.

If you need more disk space in this filesystem or in any other of the GPFS filesystems, the responsible for your project has to make a request for the extra space needed, specifying the requested space and the reasons why it is needed. For more information or requests you can Contact Us (chapter 7).

5 Running Jobs

LSF³ is the utility used at Nord III for batch processing support, so all jobs must be run through it. This document provides information for getting started with job execution at the Cluster.

5.1 LSF Commands

These are the basic commands to submit, control and check your jobs:

```
bsub < job_script
```

submits a “job script” passed through standard input (STDIN) to the queue system. Job directives (section 5.3) explains the available options to write a jobscript

```
bjobs [-w][-X][-l job_id]
```

shows all the submitted jobs.

```
bkill <job_id>
```

remove the job from the queue system, canceling the execution of the processes, if they were still running.

```
bsc_jobs
```

shows all the pending or running jobs from your group. Check BSC Commands (section 6.4) for more information.

5.2 Interactive Sessions

Allocation of an interactive session in the interactive partition has to be done by executing:

```
bsub -q interactive -W 01:00 -n 1 -Is /bin/bash
```

³<http://www.bsc.es/support/LSF/9.1.2>

5.3 Job directives

A job must contain a series of directives to inform the batch system about the characteristics of the job. We encourage that you read the bsub command's manual from any of Nord's terminals:

```
% man bsub
```

Here we provide a short summary of most common directives:

```
#BSUB -J job_name
```

Specify the name (description) of the job.

```
#BSUB -q debug
```

Specify the queue (section 5.6) for the job to be submitted. The *debug* queue is only intended for small tests, so there is a limit of 1 job per user, using up to 64 cpus (4 nodes), and one hour of wall clock limit. The queue might be reassigned by LSF internal policy, as with the *sequential* (section 5.6) queue.

```
#BSUB -W HH:MM
```

Specify how much time the job will be allowed to run. This is a **mandatory** field. NOTE: take into account that you can not specify the amount of seconds in LSF. You must set it to a value greater than the real execution time for your application and smaller than the time limits granted to the user. Notice that your job will be killed after the elapsed period.

```
#BSUB -cwd pathname
```

The working directory of your job (i.e. where the job will run). If not specified, it is the current working directory at the time the job was submitted.

```
#BSUB -e/-eo file
```

The name of the file to collect the stderr output of the job. You can use %J for job_id. -e option will APPEND the file, -eo will REPLACE the file.

```
#BSUB -o/-oo file
```

The name of the file to collect the standard output (stdout) of the job. -o option will APPEND the file, -oo will REPLACE the file.

```
#BSUB -n number
```

The number of tasks for the job. In MPI executions corresponds to the number of MPI processes and for sequential executions the number of cores.

```
#BSUB -x
```

Use the nodes exclusively. This is the default behaviour except for sequential (section 5.6) executions.

```
#BSUB -M number
```

Specify maximum amount of memory necessary for each task, specified in MB. This option is used to better schedule the jobs and select the compute nodes adequate to fulfill the job's needs. By default 1800 MB are reserved per task. Exclusive jobs will still get all available memory.

The default option could choose any kind of node. You can request the minimum amount of necessary memory with "-M" memory per task (in MB) and "-n" number of tasks. If the multiplication of these values is more than 32 GB, then we will get nodes of 64 GB or 128 GB of memory. Therefore, if we set a value higher than 64 GB, we will get nodes of 128 GB of memory.

Here is a list of the common parameters needed to request each type of node:

- 32 GB/node (Default: `#BSUB -M 1800`)
- 64 GB/node (Medium memory: `#BSUB -M 3000`)
- 128 GB/node (High memory: `#BSUB -M 7000`)

You may check this FAQ⁴ for a more in-depth explanation. **Note: For non-LowMem requests you must specify a ptile of 16.**

```
#BSUB -U reservation_ID
```

The reservation where your jobs will be allocated (assuming that your account has access to that reservation). In some occasions, node reservations can be granted for executions where only a set of accounts can run jobs. Useful for courses.

```
#BSUB -R "span[ptile=number]"
```

The number of processes assigned to a node. Note that full nodes will be allocated except for sequential (section 5.6) executions. Examples:

```
# if you want to use 4 processes per node and 4 threads:
```

```
#BSUB -R "span[ptile=4]"
export OMP_NUM_THREADS=4
```

```
# if your program has high memory
# consumption you can reduce the number
# of processes per node
```

```
#BSUB -R "span[ptile=14]"
```

5.4 MPI particulars

There are different MPI implementations available for usage. Some of them present special requirements to be used.

OpenMPI

OpenMPI is an Open Source MPI implementation and is loaded by default on our machines. You can use it like this:

```
% module load openmpi
% mpicc / mpif90 your_files
```

Once compiled, your jobscript must invoke `mpirun`:

```
#!/bin/bash
#BSUB -n 128
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err
#BSUB -J openmpi_example
#BSUB -W 00:05

module load openmpi
mpirun binary.exe
```

⁴<http://www.bsc.es/user-support/faq.php#therearesomespecialrequirementsforjobs.howshouldiputthem>

IntelMPI

IntelMPI is an MPI implementation developed by Intel. You can use it like this:

```
% module load intel impi
% mpicc / mpif90 your_files
```

Once compiled, your jobscript must invoke mpirun:

```
#!/bin/bash
#BSUB -n 128
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err
#BSUB -J impi_example
#BSUB -W 00:05

module load impi
mpirun binary.exe
```

IBM POE

IBM POE is an alternative MPI library binary compatible with IntelMPI. Use POE for executions with a large number of cores. First, in order to compile your MPI code using IBM POE:

```
% module load poe
% mpicc / mpif90 your_files
```

Once compiled, your jobscript must mention it explicitly:

```
#!/bin/bash
#BSUB -n 128
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err
#BSUB -J poe_example
#BSUB -W 00:05
#BSUB -a poe

module load poe
poe binary.exe
```

5.5 Jobscript examples

Purely sequential

```
#!/bin/bash

#BSUB -n 1
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err
#BSUB -J sequential
#BSUB -W 00:05

./serial.exe
```

Sequential with OpenMP threads

```
#!/bin/bash

#BSUB -n 16
#BSUB -R "span[ptile=16]"
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err
#BSUB -J sequential_OpenMP
#BSUB -x
```

```
#BSUB -W 00:05
export OMP_NUM_THREADS=16
./serial.exe
```

Parallel using MPI

```
#!/bin/bash

#BSUB -n 128
#BSUB -o output_%J.out
#BSUB -e output_%J.err
# In order to launch 128 processes with
# 16 processes per node:
#BSUB -R "span[ptile=16]"
#BSUB -J WRF.128-4
#BSUB -W 02:00

# You can choose the parallel environment through modules
module load intel openmpi

mpirun ./wrf.exe
```

Parallel using MPI and OpenMP threads

```
#!/bin/bash

# The total number of processes:
# 128 MPI processes + 2 OpenMP threads/process = 256 cores
#BSUB -n 128 # processes
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err

#####

# This will allocate 8 processes per node so we have #
# 8 cores per node for the threads #

#####

#BSUB -R "span[ptile=8]"

# exclusive mode (enabled by default)
#BSUB -x

#####
# Then (128 MPI tasks) / (8 tasks/node) = 16 nodes #
# reserved #
# 16 nodes * 16 cores/node = 256 cores allocated #
# (Matches the amount of cores asked for above) #
#####

#BSUB -J WRF.128-4
#BSUB -W 02:00

# Clean your environment modules
module purge

# You can choose the parallel environment through
# modules
module load intel openmpi

# 8 MPI processes per node and 16 cpus available
# (2 threads per MPI process):
export OMP_NUM_THREADS=2

mpirun ./wrf.exe
```

Parallel using MPI and extra memory

The following example shows how to ask for 3000 MB/task instead of the default 1800 MB/task. As no '-R' options are included it will still put 16 processes on each node, so the Common (2 GB/core) nodes won't be able to execute this job. This effectively guarantees that either MedMem (4 GB/core) or HighMem (8 GB/core) nodes will be used, or a mixture of both.

```
#!/bin/bash

# Total number of tasks
#BSUB -n 128
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err

# Requesting 3000 MB per task
# As no ptille is specified, only Medium and High memory nodes
# are eligible for this execution (max 128 nodes, 2048 cores)
#BSUB -M 3000

#BSUB -W 01:00

module purge
module load intel openmpi

mpirun ./my_mpi.exe
```

The following example requests the same memory per task as the preceding one but it is set to be executed on any nodes. This is done by reducing the amount of tasks per node so the total memory requested by all the tasks in the same node is below the LowMem memory threshold. This would be necessary to execute jobs that need more nodes or cpus than the MedMem and HighMem nodes can provide.

```
#!/bin/bash

# Total number of tasks
#BSUB -n 128
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err

# Requesting 3000 MB per task
#BSUB -M 3000

# Only 9 task per node.
# All compute nodes are available for execution
#BSUB -R "span[ptile=9]"

#BSUB -W 01:00

module purge
module load intel openmpi

mpirun ./my_mpi.exe
```

5.6 Queues

There are several queues present in the machines and different users may access different queues. All queues have different limits in amount of cores for the jobs and duration. You can check anytime all queues you have access to and their limits using:

```
% bsc_queues
```

Check BSC Commands (section 6.4) for more information.

Sequential executions

For any job that requires a node or less resources, the sequential queue is automatically applied. This queue is the only one that uses the same node for more than one job at a time. It also has the least

priority in the machine and the number of concurrently executed sequential jobs is limited to avoid disturbing large jobs' execution.

If you request 16 processes per node or use -x option, the full node will be assigned to your job. If you have memory problems, mind to specify the exclusive flag because just specifying a ptile may still share the node with other users.

6 Software Environment

All software and numerical libraries available at the cluster can be found at /apps/. If you need something that is not there please contact us to get it installed (see Getting Help (chapter 7)).

6.1 C Compilers

In the cluster you can find these C/C++ compilers :

icc /icpc -> Intel C/C++ Compilers

```
% man icc
% man icpc
```

gcc /g++ -> GNU Compilers for C/C++

```
% man gcc
% man g++
```

All invocations of the C or C++ compilers follow these suffix conventions for input files:

```
.C, .cc, .cpp, or .cxx -> C++ source file.
.c -> C source file
.i -> preprocessed C source file
.so -> shared object file
.o -> object file for ld command
.s -> assembler source file
```

By default, the preprocessor is run on both C and C++ source files.

These are the default sizes of the standard C/C++ datatypes on the machine

Table 1: Default datatype sizes on the machine

Type	Length (bytes)
bool (c++ only)	1
char	1
wchar_t	4
short	2
int	4
long	8
float	4
double	8
long double	16

Distributed Memory Parallelism

To compile MPI programs it is recommended to use the following handy wrappers: mpicc, mpicxx for C and C++ source code. You need to choose the Parallel environment first: module load openmpi /module load impi /module load poe. These wrappers will include all the necessary libraries to build MPI applications without having to specify all the details by hand.

```
% mpicc a.c -o a.exe
% mpicxx a.C -o a.exe
```

Shared Memory Parallelism

OpenMP directives are fully supported by the Intel C and C++ compilers. To use it, the flag `-qopenmp` must be added to the compile line.

```
% icc -qopenmp -o exename filename.c  
% icpc -qopenmp -o exename filename.C
```

You can also mix MPI + OPENMP code using `-openmp` with the mpi wrappers mentioned above.

Automatic Parallelization

The Intel C and C++ compilers are able to automatically parallelize simple loop constructs, using the option “`-parallel`” :

```
% icc -parallel a.c
```

6.2 FORTRAN Compilers

In the cluster you can find these compilers :

ifort -> Intel Fortran Compilers

```
% man ifort
```

gfortran -> GNU Compilers for FORTRAN

```
% man gfortran
```

By default, the compilers expect all FORTRAN source files to have the extension “.f”, and all FORTRAN source files that require preprocessing to have the extension “.F”. The same applies to FORTRAN 90 source files with extensions “.f90” and “.F90”.

Distributed Memory Parallelism

In order to use MPI, again you can use the wrappers `mpif77` or `mpif90` depending on the source code type. You can always `man mpif77` to see a detailed list of options to configure the wrappers, ie: change the default compiler.

```
% mpif77 a.f -o a.exe
```

Shared Memory Parallelism

OpenMP directives are fully supported by the Intel Fortran compiler when the option “`-qopenmp`” is set:

```
% ifort -qopenmp
```

Automatic Parallelization

The Intel Fortran compiler will attempt to automatically parallelize simple loop constructs using the option “`-parallel`”:

```
% ifort -parallel
```

6.3 Modules Environment

The Environment Modules package (<http://modules.sourceforge.net/>) provides a dynamic modification of a user's environment via modulefiles. Each modulefile contains the information needed to configure the shell for an application or a compilation. Modules can be loaded and unloaded dynamically, in a clean fashion. All popular shells are supported, including bash, ksh, zsh, sh, csh, tcsh, as well as some scripting languages such as perl.

Installed software packages are divided into five categories:

- Environment: modulefiles dedicated to prepare the environment, for example, get all necessary variables to use openmpi to compile or run programs
- Tools: useful tools which can be used at any time (php, perl, ...)
- Applications: High Performance Computers programs (GROMACS, ...)
- Libraries: Those are typically loaded at a compilation time, they load into the environment the correct compiler and linker flags (FFTW, LAPACK, ...)
- Compilers: Compiler suites available for the system (intel, gcc, ...)

Modules tool usage

Modules can be invoked in two ways: by name alone or by name and version. Invoking them by name implies loading the default module version. This is usually the most recent version that has been tested to be stable (recommended) or the only version available.

```
% module load intel
```

Invoking by version loads the version specified of the application. As of this writing, the previous command and the following one load the same module.

```
% module load intel/2017.1
```

The most important commands for modules are these:

- *module list* shows all the loaded modules
- *module avail* shows all the modules the user is able to load
- *module purge* removes all the loaded modules
- *module load <modulename>* loads the necessary environment variables for the selected modulefile (PATH, MANPATH, LD_LIBRARY_PATH...)
- *module unload <modulename>* removes all environment changes made by module load command
- *module switch <oldmodule> <newmodule>* unloads the first module (oldmodule) and loads the second module (newmodule)

You can run “module help” any time to check the command's usage and options or check the module(1) manpage for further information.

6.4 BSC Commands

The Support team at BSC has provided some commands useful for user's awareness and ease of use in our HPC machines. These commands are available through a special module (*bsc/current*) loaded at the beginning of any session. A short summary of these commands follows:

- *bsc_acct*: Displays accounting information on the project's allocation usage.
- *bsc_jobcheck*: Returns a comprehensive description of the resources requested by a jobscript.
- *bsc_jobs*: Show a list of your submitted jobs and those from other members of your group.

- `bsc_load`: Show the performance and resources usage of all the nodes of a specific running job.
- `bsc_queues`: Show the queues the user has access to and their time/resources limits.
- `bsc_quota`: Show a comprehensible quota usage summary for all accessible filesystems.
- `bsc_micwrapper`: Automatized script to execute MPI jobs using the Intel Xeon Phi Accelerators.

You can check more information about these commands through any of the following manpages:

```
% man bsc_commands
% man bsc_jobs
% man bsc
```

6.5 TotalView

TotalView is a graphical portable powerful debugger from Rogue Wave Software designed for HPC environments. It also includes MemoryScape and ReverseEngine. It can debug one or many processes and/or threads. It is compatible with MPI, OpenMP, Intel Xeon Phi and CUDA.

Users can access to the latest version of TotalView 8.13 installed in:

```
/apps/TOTALVIEW/totalview
```

Important: Remember to access with `ssh -X` to the cluster and submit the jobs to `x11` queue since TotalView uses a single window control.

There is a Quick View of TotalView⁵ available for new users. Further documentation and tutorials can be found on their website⁶ or in the cluster at:

```
/apps/TOTALVIEW/totalview/doc/pdf
```

6.6 Tracing jobs with BSC Tools

In this section you will find an introductory guide to get execution traces in Nord. The tracing tool `Extræ` supports many different tracing mechanisms, programming models and configurations. For detailed explanations and advanced options, please check the complete `Extræ User Guide`⁷

The most recent stable version of `Extræ` is always located at:

```
/apps/CEPBATTOOLS/extrae/latest/default/64
```

This package is compatible with the default MPI runtime in Nord (`OpenMPI`). Packages corresponding to older versions and enabling compatibility with other MPI runtimes (`IntelMPI`, `MVAPICH`) can be respectively found under this directory structure:

```
/apps/CEPBATTOOLS/extrae/<choose-version>/<choose-runtime>/64
```

In order to trace an execution, you have to load the **module `extrae`** and write a script that sets the variables to configure the tracing tool. Let's call this script `trace.sh`. It must be executable (`chmod +x ./trace.sh`). Then your job needs to run this script before executing the application.

Example for MPI jobs:

```
#!/bin/bash
#BSUB -n 128
#BSUB -o output_%J.out
#BSUB -e output_%J.err
#BSUB -R "span[ptile=16]"
#BSUB -J job_name
#BSUB -W 00:10

module load extrae

mpirun ./trace.sh ./app.exe
```

⁵<https://www.bsc.es/support/TotalView-QuickView.pdf>

⁶<http://www.roguewave.com/products/totalview.aspx>

⁷https://tools.bsc.es/tools_manuals

Example for threaded (OpenMP or pthreads) jobs:

```
#!/bin/bash
#BSUB -n 1
#BSUB -oo output_%J.out
#BSUB -eo output_%J.err
#BSUB -J job_name
#BSUB -W 00:10

module load extrae

./trace.sh ./app.exe
```

Example of trace.sh script:

```
#!/bin/bash

export EXTRAE_CONFIG_FILE=./extrae.xml
export LD_PRELOAD=${EXTRAE_HOME}/lib/<tracing-library>
$*
```

Where:

- `EXTRAE_CONFIG_FILE` points to the Extrae configuration file. Editing this file you can control the type of information that is recorded during the execution and where the resulting trace file is written, among other parameters. By default, the resulting trace file will be written into the current working directory. Configuration examples can be found at: `${EXTRAE_HOME}/share/examples`
- `<tracing-library>` depends on the programming model the application uses:

Job Type	Tracing library	An example to get started
MPI	libmpitrace.so (C codes) libmpitracef.so (Fortran Codes)	MPI/ld-preload/ job.lsf
OpenMP	libomptrace.so	OMP/ run_ldpreload.sh
Pthreads	libpttrace.so	PTHREAD/ README
OmpSs	-	OMPSS/job.lsf
Sequential job (manual instrumentation)	libseqtrace.so	SEQ/ run_instrumented.sh
*Automatic instrumentation of user functions and parallel runtime calls	-	SEQ/ run_dyninst.sh

* Jobs that make explicit calls to the Extrae API do **not** load the tracing library via `LD_PRELOAD`, but link with the libraries instead.

** Jobs using automatic instrumentation via Dyninst **neither** load the tracing library via `LD_PRELOAD` **nor** link with it.

For other programming models and their combinations, check the full list of available tracing libraries at section 1.2.2 of the Extrae User Guide.

7 Getting help

BSC provides users with excellent consulting assistance. User support consultants are available during normal business hours, Monday to Friday, 09 a.m. to 18 p.m. (CEST time).

User questions and support are handled at: `support@bsc.es`

If you need assistance, please supply us with the nature of the problem, the date and time that the problem occurred, and the location of any other relevant information, such as output files. Please contact BSC if you have any questions or comments regarding policies or procedures.

Our address is:

7.1 Frequently Asked Questions (FAQ)

You can check the answers to most common questions at BSC's Support Knowledge Center⁸. There you will find online and updated versions of our documentation, including this guide, and a listing with deeper answers to the most common questions we receive as well as advanced specific questions unfit for a general-purpose user guide.

8 Appendices

8.1 SSH

SSH is a program that enables secure logins over an insecure network. It encrypts all the data passing both ways, so that if it is intercepted it cannot be read. It also replaces the old an insecure tools like telnet, rlogin, rcp, ftp, etc. SSH is a client-server software. Both machines must have ssh installed for it to work.

We have already installed a ssh server in our machines. You must have installed an ssh client in your local machine. SSH is available without charge for almost all versions of UNIX (including Linux and MacOS X). For UNIX and derivatives, we recommend using the OpenSSH client, downloadable from <http://www.openssh.org>, and for Windows users we recommend using Putty, a free SSH client that can be downloaded from <http://www.putty.org>. Otherwise, any client compatible with SSH version 2 can be used.

This section describes installing, configuring and using the client on Windows machines. No matter your client, you will need to specify the following information:

- Select SSH as default protocol
- Select port 22
- Specify the remote machine and username

For example with putty client:

This is the first window that you will see at putty startup. Once finished, press the **Open** button. If it is your first connection to the machine, you will get a *Warning* telling you that the host key from the server is unknown, and will ask you if you are agree to cache the new host key, press Yes.

IMPORTANT: If you see this warning another time and you haven't modified or reinstalled the ssh client, please do *not* log in, and contact us as soon as possible (see Getting Help (chapter 7)).

Finally, a new window will appear asking for your login and password:

8.2 Transferring files

To transfer files to or from the cluster you need a secure ftp (sftp) o secure copy (scp) client. There are several different clients, but as previously mentioned, we recommend using of Putty clients for transferring files: **psftp** and **pscp**. You can find it at the same web page as Putty (<http://www.putty.org>⁹).

Some other possible tools for users requiring graphical file transfers could be:

- WinSCP: Freeware Sftp and Scp client for Windows (<http://www.winscp.net>)
- SSH: Not free. (<http://www.ssh.org>)

⁸<http://www.bsc.es/user-support/>

⁹<http://www.putty.org/>




Figure 1: Putty client




Figure 2: Putty certificate security alert

```

mn3.bsc.es - PuTTY
login as:
Using keyboard-interactive authentication.
Password:
Last login: Fri May 9 13:19:35 2014 from

-----
                welcome to MareNostrum III

                [ BSC ]

- All home directories are in GPFS and quotas are enabled
- Applications are located at /apps
- To change password, please login from your local machine
to:
  dt01.bsc.es
- Active Archive and transfer management machine:
  dt01.bsc.es
- For further information read MareNostrum III user Guide:
  http://www.bsc.es/support/MareNostrum3-ug.pdf
- BSC SUPPORT COMMANDS:
  See 'man bsc' for more information
- [NEW] LSF extended Documentation:
  http://www.bsc.es/support/LSF/9.1.2
  Please contact support@bsc.es for questions
-----

load openmpi/1.5.4 (PATH, MANPATH, LD_LIBRARY_PATH)
login3:~>

```

Figure 3: Cluster login

Using PSFTP

You will need a command window to execute psftp (press start button, click run and type cmd). The program first asks for the machine name (mn1.bsc.es), and then for the username and password. Once you are connected, it's like a Unix command line.

With command **help** you will obtain a list of all possible commands. But the most useful are:

- get file_name : To transfer from the cluster to your local machine.
- put file_name : To transfer a file from your local machine to the cluster.
- cd directory : To change remote working directory.
- dir : To list contents of a remote directory.
- lcd directory : To change local working directory.
- !dir : To list contents of a local directory.

You will be able to copy files from your local machine to the cluster, and from the cluster to your local machine. The syntax is the same that cp command except that for remote files you need to specify the remote machine:

```

Copy a file from the cluster:
> pscp.exe username@mn1.bsc.es:remote_file local_file
Copy a file to the cluster:
> pscp.exe local_file username@mn1.bsc.es:remote_file

```

8.3 Using X11

In order to start remote X applications you need and X-Server running in your local machine. Here is a list of most common X-servers for windows:

- Cygwin/X: <http://x.cygwin.com>
- X-Win32 : <http://www.starnet.com>
- WinaXe : <http://labf.com>

- XconnectPro : <http://www.labtam-inc.com>
- Exceed : <http://www.hummingbird.com>

The only Open Source X-server listed here is Cygwin/X, you need to pay for the others. Once the X-Server is running run putty with X11 forwarding enabled:

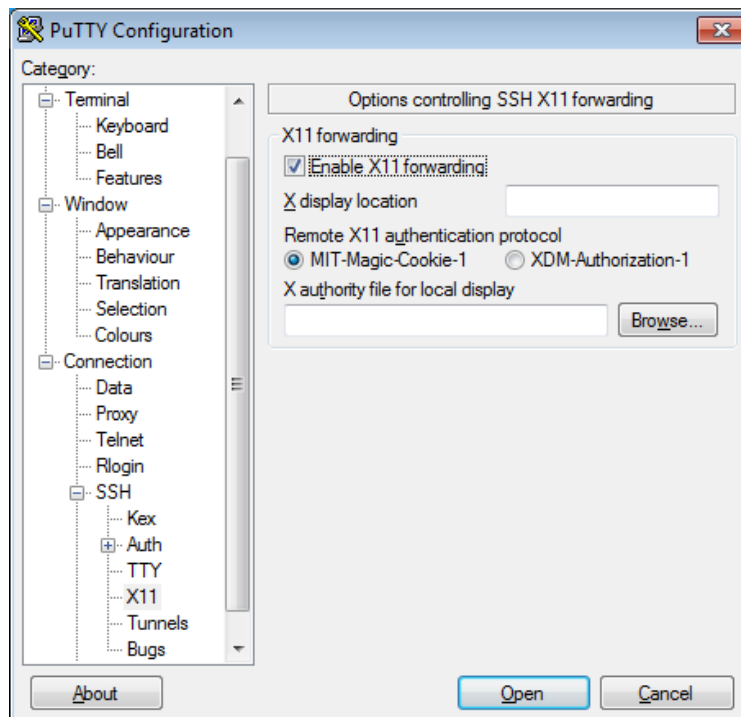


Figure 4: Putty X11 configuration