

# HP-SEE PARADOX Cluster job management

[www.hp-see.eu](http://www.hp-see.eu)

**Nikola Grkic**  
Institute of Physics Belgrade, Serbia  
[ngrkic@ipb.ac.rs](mailto:ngrkic@ipb.ac.rs)



# HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities



- ❑ Access to login node
- ❑ Preparing job submitting scripts
- ❑ Submitting jobs:
  - ❑ Serial jobs
  - ❑ OpenMP jobs
  - ❑ MPI jobs
  - ❑ Hybrid jobs
  - ❑ Application specific jobs (NAMD)
- ❑ Batch system job control and monitoring

# Access to login node



**HP-SEE**

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ PARADOX login node : **ui.ipb.ac.rs**
- ❑ Access via **ssh** (remote login program)
  - ❑ **\$ ssh username@ui.ipb.ac.rs**  
If you need a graphical environment you have to use the -X option:
    - ❑ **\$ ssh username@ui.ipb.ac.rs -X**

# Access to login node



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

## □ File transfer

To transfer files between **ui.ipb.ac.rs** and your local machine, you can use the **scp** command. Create an archive with the directories you want to copy (it will be faster to transfer):

□ `$ tar -cvzf archivename.tgz directoryname1 directoryname2`

or in the case of a file:

□ `$ tar -cvzf archivename.tgz filename`

Transfer the archive to your home directory at ui.ipb.ac.rs:

□ `$ scp archivename.tgz username@ui.ipb.ac.rs:`

Login on **ui.ipb.ac.rs**:

□ `$ ssh username@ipb.ac.rs`

Uncompress the archive in your target directory:

□ `$ tar -xvzf archivename.tgz destinationdirectory`

# Access to login node



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ There are two file systems available to users at **ui.ipb.ac.rs**:
  - ❑ `/home`
  - ❑ `/nfs`
  
- ❑ File systems have directories like `/home/<USERNAME>` or `/nfs/<USERNAME>`
- ❑ `/nfs` is shared between all nodes on the cluster
- ❑ It is required to put all executables and data used by jobs in this directory.
- ❑ There is another local file system available on each worker node: `/scratch` which is used for temporary storage of running jobs

# Preparing job submitting scripts



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ Portable batch system
  - ❑ Job submissions
  - ❑ Resources allocations
  - ❑ Jobs launching
  - ❑ PBS server **ce64.ipb.ac.rs**
  - ❑ command **qsub** (**ui.ipb.ac.rs**)
  - ❑ Queue **hpsee** is available for user's job submission
  
- ❑ To submit a job, you first have to write a shell script which contains:
  - ❑ A set of directives (beginning with #PBS) which describe needed resources for your job
  - ❑ Lines necessary to execute your code

# Preparing job submitting scripts



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

## □ PBS script

```
#!/bin/bash
#PBS -q hpsee
#PBS -l nodes=1:ppn=1
#PBS -l walltime=00:10:00
#PBS -e ${PBS_JOBID}.err
#PBS -o ${PBS_JOBID}.out
cd $PBS_O_WORKDIR
chmod +x job.sh
./job.sh
```

## □ Job script

```
#!/bin/bash
date
hostname
pwd
sleep 10
```

# Preparing job submitting scripts



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ `#!/bin/bash` - Specifies the shell to be used when executing the command portion of the script.
- ❑ `#PBS -q <queue>` - Directs the job to the specified queue. Queue **hpsee** should be used.
- ❑ `#PBS -o <name>` - Writes standard output to `<name>`. `$PBS_JOBID` is an environment variable created by PBS that contains the PBS job identifier.
- ❑ `#PBS -e <name>` - Writes standard error to `<name>`.



# Preparing job submitting scripts



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ `#PBS -l walltime=<time>` - Maximum wall-clock time. `<time>` is in the format HH:MM:SS.
- ❑ `#PBS -l nodes=1:ppn=1` - Number of nodes to be reserved for exclusive use by the job and number of virtual processors per node (ppn) requested for this job.
- ❑ `cd $PBS_O_WORKDIR` - Change to the initial working directory
- ❑ `chmod +x job.sh` - setting execute permission on job.sh file
- ❑ `./job.sh` - execute job

# Preparing job submitting scripts



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ Job submitting is done by executing `qsub` in `/nfs/username/somefolder`
  - ❑ `$ qsub job.pbs`
- ❑ The `qsub` command will return:
  - ❑ `<JOB_ID>.ce64.ipb.ac.rs`

# Submitting jobs



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

## Serial jobs

[http://wiki.ipb.ac.rs/index.php/Serial\\_job\\_example](http://wiki.ipb.ac.rs/index.php/Serial_job_example)

```
#!/bin/bash
#PBS -q hpsee
#PBS -l nodes=1:ppn=1
#PBS -l walltime=00:10:00
#PBS -e ${PBS_JOBID}.err
#PBS -o ${PBS_JOBID}.out
cd $PBS_O_WORKDIR
chmod +x job.sh
./job.sh
```

# Submitting jobs



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

## □ OpenMP

□ [http://wiki.ipb.ac.rs/index.php/OpenMP\\_job\\_example](http://wiki.ipb.ac.rs/index.php/OpenMP_job_example)

```
#!/bin/bash
#PBS -q hpsee
#PBS -l nodes=1:ppn=6
#PBS -l walltime=10:00:00
#PBS -e ${PBS_JOBID}.err
#PBS -o ${PBS_JOBID}.out

cd $PBS_O_WORKDIR
chmod +x job
export OMP_NUM_THREADS=6
./job
```

# Submitting jobs



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

## □ MPI jobs

□ [http://wiki.ipb.ac.rs/index.php/MPI\\_job\\_examples](http://wiki.ipb.ac.rs/index.php/MPI_job_examples)

```
#!/bin/bash
#PBS -q hpsee
#PBS -l nodes=3:ppn=2
#PBS -l walltime=10:00:00
#PBS -e ${PBS_JOBID}.err
#PBS -o ${PBS_JOBID}.out

cd $PBS_O_WORKDIR
chmod +x job
cat $PBS_NODEFILE
${MPI_MPICH_MPIEXEC} ./job           # If mpich-1.2.7p1 is used
#${MPI_MPICH2_MPIEXEC} --comm=pmi ./job # If mpich2-1.1.1p1 is used
#${MPI_OPENMPI_MPIEXEC} ./job       # If openmpi-1.2.5 is used
```

# Submitting jobs



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

## □ Hybrid jobs:

- [http://wiki.ipb.ac.rs/index.php/Hybrid\\_job\\_example](http://wiki.ipb.ac.rs/index.php/Hybrid_job_example)

```
#!/bin/bash
#PBS -q hpsee
#PBS -l nodes=4:ppn=8
#PBS -l walltime=10:00:00
#PBS -e ${PBS_JOBID}.err
#PBS -o ${PBS_JOBID}.out
export OMP_NUM_THREADS=8

cd $PBS_O_WORKDIR
chmod +x job

${MPI_OPENMPI_MPIEXEC} -np 4 -npernode 1 ./job
```

# Submitting jobs



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- Application specific jobs (NAMD):
  - [http://wiki.ipb.ac.rs/index.php/Application\\_specific\\_job\\_example\\_\(NAMD\)](http://wiki.ipb.ac.rs/index.php/Application_specific_job_example_(NAMD))

```
#!/bin/bash
#PBS -q hpsee
#PBS -l nodes=4:ppn=8
#PBS -l walltime=10:00:00
#PBS -e ${PBS_JOBID}.err
#PBS -o ${PBS_JOBID}.out

cd $PBS_O_WORKDIR
chmod +x script.sh
./script.sh
```

# Batch system job control and monitoring



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

- ❑ To check the status of job use the following command:
  - ❑ `$ qstat <JOB_ID>`
- ❑ Alternatively check the status of all your jobs using the following syntax of the qstat command:
  - ❑ `$ qstat -u <user_name>`
- ❑ To get detailed information about your job use the following command:
  - ❑ `$ qstat -f <JOB_ID>`
- ❑ If, for some reason, you want to cancel a job following command should be executed:
  - ❑ `$ qdel <JOB_ID>`



# Batch system job control and monitoring



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

qstat	list information about queues and jobs
qstat -q	list all queues on system
qstat -Q	list queue limits for all queues
qstat -a	list all jobs on system
qstat -au <i>userID</i>	list all jobs owned by user <i>userID</i>
qstat -s	list all jobs with status comments
qstat -r	list all running jobs

# Batch system job control and monitoring



HP-SEE

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

qstat- f jobID	list all information known about specified job
qstat -n	in addition to the basic information, nodes allocated to a job are listed
qstat -Qf <queue>	list all information about specified queue
qstat -B	list summary information about the PBS server
qdel jobID	delete the batch job with <i>jobID</i>
qalter	alter a batch job
qsub	submit a job

# Questions



**HP-SEE**

High-Performance Computing Infrastructure  
for South East Europe's Research Communities

