

# Submitting batch jobs

## Slurm on ecgate

### Solutions to the practicals

Xavi Abellan  
xavier.abellan@ecmwf.int  
User Support Section

# Practical 1: Basic job submission

---

- Practicals must be run on **ecgate**, so make sure you log in there first!

```
$> ssh ecgate
$> cd $SCRATCH
$> tar xvzf ~trx/intro/batch_ecgate_practicals.tar.gz
$> cd batch_ecgate_practicals/basic
```

1. Have a look at the script “**env.sh**”
2. Submit the job and check whether it is running.
  - What QoS is it using? What is the time limit of the job?
3. Where did the output of the job go? Have a look at the output
4. Submit the job again and then once it starts cancel it.
5. Check the output



# Practical 1: Basic job submission

```
$> cat env.sh
#!/bin/bash
# This is an example of a job dumping the SLURM related variables in the environment
[ -z "$SLURM_JOBID" ] && echo "This job should be run in batch!" >&2 && exit 1
echo "Current time is `date`"
echo "This should go to the error file..." >&2
echo "I am job id $SLURM_JOBID"
echo "Going to sleep for a couple of minutes..."
sleep 120
echo "Waking up!"
echo "Current time is `date`"
echo "Bye!"
$> ./env.sh
This job should be run in batch!
$> sbatch env.sh
Submitted batch job 1269989
$> squeue -j 1269989
  JOBID      NAME      USER      QOS      STATE      TIME  TIMELIMIT  NODELIST (REASON)
1269989    env.sh    trx      normal  RUNNING    0:54  1-00:00:00          ecgb04
$> ls
env.sh  slurm-1269989.out  work
```



# Practical 1: Basic job submission

```
$> cat slurm-1269989.out
Current time is Sun Mar  2 15:27:17 GMT 2014
This should go to the error file...
I am job id 1269989
Going to sleep for a couple of minutes...
Waking up!
Current time is Sun Mar  2 15:29:17 GMT 2014
Bye!
$> sbatch env.sh
Submitted batch job 1270043
$> scancel -i 1270043
Cancel job_id=1270043 name=env.sh partition=batch [y/n]? y
$> cat slurm-1270043.out
Current time is Sun Mar  2 15:35:24 GMT 2014
This should go to the error file...
I am job id 1270043
Going to sleep for a couple of minutes...
slurmd[ecgb09]: *** JOB 1270043 CANCELLED AT 2014-03-02T15:35:35 ***
```



# Practical 1: Basic job setup

---

- Can you modify the previous job so it...
  1. ... runs in the express QoS, with a wall clock limit of 5 minutes?
  2. ... uses the subdirectory work/ as the working directory?
  3. ... sends the...
    - a) ... output to the file work/env\_out\_<jobid>.out ?
    - b) ... error to work/env\_out\_<jobid>.err?
  4. ... sends you an email when the job starts?
- Try your job after the modifications and check if they are correct
  - You can do the modifications one by one or all at once...



# Practical 1: Basic job submission

```
#!/bin/bash
# This is an example of a job dumping the SLURM related variables in the environment
#SBATCH --qos=express
#SBATCH --time=5
#Replace trx by your user id
#SBATCH --workdir=/scratch/ectrain/trx/batch_ecgate_practicals/basic/work
#SBATCH --output=env_out_%j.out
#SBATCH --error=env_out_%j.err
#SBATCH --mail-type=BEGIN

[ -z "$SLURM_JOBID" ] && echo "This job should be run in batch!" >&2 && exit 1
echo "Current time is `date`"
echo "This should go to the error file..." >&2
echo "I am job id $SLURM_JOBID"
echo "Going to sleep for a couple of minutes..."
sleep 120
echo "Waking up!"
echo "Current time is `date`"
echo "Bye!"
```



## Practical 2: reviewing past runs

---

- How would you...

- retrieve the list of jobs that you ran today?

```
$> sacct
```

- retrieve the list of all the jobs that were cancelled today by user trx?

```
$> sacct -u trx -s CANCELLED -S 00:00
```

- ask for the submit, start and end times for a job of your choice?

```
$> sacct -j 1234567 -o jobid,submit,start,end
```

- find out the output an error paths for a job of your choice?

```
$> sacct -j 1234567 -o jobid,comment%190
```





# Practical 3: Fixing broken jobs

- What is wrong in job1? Can you fix it?
  - No shebang specified at the beginning of the files
  - Spaces between directive and value
  - Job name must not contain any space

```
#!/bin/bash
#SBATCH --job-name=job_1
#SBATCH --output=job1-%J.out
#SBATCH --error=job1-%J.out
#SBATCH --qos=express
#SBATCH --time=00:05:00

# This is the job
echo "I was broken!"
sleep 30
```





# Practical 3: Fixing broken jobs

---

- What is wrong in job2? Can you fix it?
  - The qos name does not exist
  - The time limit is wrongly set to 10 days instead of 10 hours

```
#!/bin/bash
#SBATCH --job-name=job2
#SBATCH --output=job2-%J.out
#SBATCH --error=job2-%J.out
#SBATCH --qos=normal
#SBATCH --time=10:00

# This is the job
echo "I was broken!"
sleep 30
```



# Practical 3: Fixing broken jobs

- What is wrong in job3? Can you fix it?
  - The output directory does not exist. You must create it before submitting
  - Slurm will not expand shell variables in directives. You must replace `$$SCRATCH` with the right value

```
#!/bin/bash
#SBATCH --job-name=job3
#SBATCH --output=output/job3-%J.out
#SBATCH --error=output/job3-%J.out
#SBATCH --workdir=$$SCRATCH/batch_ecgate_practicals/broken
#SBATCH --qos=normal
#SBATCH --time=00:01

# This is the job
echo "I was broken!"
sleep 30
```



# Bonus: Migrating from LoadLeveler

```
$> ll2slurm -i ll_job.sh -o slurm_job.sh
WARNING: ignoring directive 'environment'. The submitting environment is exported by
default in Slurm.
WARNING: ignoring directive 'job_cpu_limit'. Slurm uses wall clock instead of cpu time.
Please use --time option.
WARNING: No workdir (initialdir) set. Setting it to '/scratch/ectrain/trx'
```

```
#!/bin/ksh
#@ shell = /usr/bin/ksh
#@ job_name = limits
#@ output = $(job_name).$(jobid).out
#@ error = $(job_name).$(jobid).out
#@ environment = COPY_ALL
#@ job_cpu_limit = 00:03:00,00:02:55
#@ wall_clock_limit= 00:05:00,00:04:50
#@ class=normal
#@ queue
```

```
echo "Hello World!"
```

```
#!/bin/ksh
#SBATCH --workdir=/scratch/ectrain/trx
#SBATCH --job-name=limits
#SBATCH --output=limits.%j.out
#SBATCH --error=limits.%j.out
#SBATCH --time=00:05:00
#SBATCH --qos=normal
```

```
echo "Hello World!"
```