

# How to fix the broken job-scheduling system

Niels Bassler <[bassler@phys.au.dk](mailto:bassler@phys.au.dk)> for IFJ-PAN, Krakow; May 2014.

A few things are needed to prepare your virtual machine for parallelization. (Yes, I did not check for persistency of the initial setup after reboot. My bad, sorry. This is a mess.)

- The general recipe for setting up stuff on clusters is shown here:  
<http://willworkforscience.blogspot.dk/2012/01/quick-and-dirty-guide-for-parallelizing.html>
- First, change to the root user, using the “su” command and enter the password.
- Next figure out what your IP number is, by entering  
`ifconfig eth0`  
and look for the number after “inet addr:”. it should be starting with 10.16.XXX.XXX
- Use your favourite editor (e.g. pico) to correct the IP address (10.16.xxx.xxx) in /etc/hosts file.
- Then enter this line:  
`chmod 755 /usr/local/fluka/qfluka.sh`

Now make a little test if there is a job queue (for some reason mine got deleted upon reboot).

- Enter `qstat -q`  
it should respond with a queue. If not (e.g. if there is just a blank answer, there is no queue).

Note, another mistake is that qfluka.sh expects a queue named “long”. You can either edit the /usr/local/fluka/qfluka.sh and change the “long” into “batch”, or create a new queue by executing the next step.

## SETUP A NEW QUEUE

***If there is no queue or you wish to setup the “long” queue, do the following:***

Download the script which will setup two queues “batch” and “long”, and execute it:

```
wget http://neptun.phys.au.dk/~bassler/TUTORIALS/MC/setpbs.sh  
chmod 755 setpbs.sh  
./setpbs.sh
```

If you see something bad or error messages, wave with your hands to catch my attention.

## TEST YOUR SYSTEM

Let's see if all servers are running:

```
ps -e | grep pbs
```

... should say something like:

```
1286 ?          00:00:00 pbs_mom
1293 ?          00:00:00 pbs_server
2174 ?          00:00:00 pbs_sched
```

Anything in the queue?

```
$ qstat
```

```
$
```

Nope, it's empty.

Enter

```
exit
```

to become a normal user again.

Try to submit a simple job

```
echo "sleep 20" | qsub
```

and within the next 20 seconds you can test, if its in the queue:

```
qstat
```

which answers something like:

Job id	Name	User	Time Use	S	Queue
0.kepler	STDIN	ulam	0	R	batch

Now, when this finally works, you are ready to run fluka in parallel.