

Computational software available at the CMMCQC IBM HPC cluster:

Gaussian 09

Submitting jobs to the CMMCQC IBM HPC cluster:

Your jobs should be submitted ONLY via the queuing system, as detailed below. No interactive/background jobs are allowed.

1. Connection and file transfer:

- open a Putty or other ssh communication console and connect to 193.231.20.70 (the fix IP address of the CMMCQC cluster), using the *ssh* port 22.
- authorized users should logon with their own usernames and passwords.
- copy the input file (which must have the .gaussian extension) in the folder *Zeus/home/user/demon* or *Zeus/home/user*. In order to do that, you can use any Secure File Transfer Protocol-client software (like WinSCP, use the same setting shown above).

If you are within the *quant2k* network, you can transfer directly the files from and to your *planck* home folder using midnight commander (the folder *NOD-home* on *planck* is a shortcut to Zeus).

2. Launching the calculation

If the user is launching a calculation for the first time, then he/she should use the command *q-g09 start*. In all other instances, the queuing system will periodically check the folder for Gaussian input files and automatically submit the job in a few minutes or place it in the queue if no blade is available. The command *killall sleep* can be used so that the system checks immediately for new input files.

Beginners should check the online Gaussian09 technical support (http://www.gaussian.com/g_tech/1.htm) for help with editing their input file. Go to <http://chem.ubbcluj.ro/pagini/anorganica/isi/cmmcqc%20calculations/water.gaussian> to download an example of an input file, requesting a single point calculations on a water molecule, at the HF/6-31G level. Notice **you don't have to type Link 0 commands**; the current setup will automatically do that, using the maximum amount of disk space, memory and number of cores available.

Running jobs on more than one node:

The line *! num_blades [no. of requested nodes]* should be inserted in the Gaussian input file. However, you should check on http://www.gaussian.com/g_tech/g_ur/m_linda.htm that your job is Linda-parallel, before requesting more than one node.

Verifying running jobs:

The command *q-stat2* will display the current load of the HPC cluster: the number of running jobs (detailed for each user running calculations), the number of pending jobs, the processor usage and the number of free blades.

If you want to check the progress of running calculations:

- type the *bjobs -u username* command, this will display the number of the blade on which your job is running and also the ID of your job
- ssh to the blade (with *ssh compute-0xx*)
- change the folder to the scratch directory (with the *cd /scr/username/jobs/* command)
- use *ls* to display the jobs running, change the directory to the job and you can visualize the output with the command *vi filename.log* or find a specific keyword by using the *grep* command (*grep "word" filename.log*)
- you can go back to Zeus by typing *logout*

Killing your job:

For stopping Gaussian 03 jobs, the *bkill* command should normally work. However, at this time, the *bkill* command does not lead to a clean halting of your Gaussian job, if the job is already running. Therefore, please follow these steps for stopping a job that is already running:

- use the *q-stat* command to find the JOBID (it will be displayed in the first column);
- using the *bjobs JOBID* command (where JOBID is the number obtained in the previous step) or the *bjobs -w JOBID* command find out which nodes are used by your job, and make a note of the first one in the list as this is the 'headnode'
- ssh into the headnode; then, use the *top* command to find the process ID of the Gaussian process (it will be the process using most of the resources on that machine, and the PID will be listed in the first column)
- use this PID to finally kill the job, using the command *kill -9 ID*.

If you use the *bkill* command, please ensure that the killed job was deleted by checking the *scr/username/jobs/* folder on the headnode for the name of your job. If the folder is present you can delete it with the command *rm -rf foldername*.

Retrieving a finished job:

The current settings are that at the end of a calculation (either by normal termination or ending with an error), the files are packed in a .tgz archive *gXXrez.filename.tgz* and placed in the folder in which the input file was copied. After you copy the results files, you are required to remove the archive from Zeus (again by using WinSCP), so harddisk space does not become an issue.

If your calculation ended in error related to memory or disk shortage, please contact the system administrator. If that is not the case, then check out the www.ccl.net for helpful discussions on almost all errors that you can encounter with Gaussian.

Other commands are available from the *Lava user manual*.