McStas on the ESS cluster

1. Get an account on the cluster - and consult DMSC computing cluster for how to access it!
2. Please **ensure that your simulation runs on your own desktop / laptop first** - and estimate a required runtime to achieve your necessary statistics!
3. When transferring input data to the cluster - please **ensure that all necessary files are present**, i.e. your
    a. instrument file
    b. any extra needed components
    c. any extra c-codes
    d. any extra datafiles
4. Load the **mpi module of choice,** one of the openmpi ones are recommended:

```
[willend@r1n2 ~]$ module load openmpi<tab>

openmpi/3.0_gcc1020 openmpi/3.0_intel17 openmpi/4.0_gcc920
openmpi/3.0_gcc831 openmpi/4.0_gcc1020 openmpi/4.0_intel17
openmpi/3.0_gcc920 openmpi/4.0_gcc831

[willend@r1n2 ~]$ module load openmpi/4.0_gcc831
```

    Recommendation is module load openmpi/4.0_gcc831

5. Load the **mcstas of choice** using e.g. the below command (which will give you the newest installed mcstas)

```
[willend@r1n2 ~]$ module load mcstas/
```

6. **Compile your instrument** at the compile node using a command like

```
[willend@r1n2 ~]$ mcrun -c --mpi=1 -n0 BNL_H8.instr
```

    or

```
willend@r1n2 ~]$ mcrun.pl -c --mpi -n0 BNL_H8.instr
```

    You may experience some error output like below, but this can be safely ignored when compiling, since mpi processes are only allowed to run via the slurm scheduler.

```
Error obtaining unique transport key from ORTE (orte_precondition_transports not present in

the environment).

Local host: r1n2.esss.dk
--------------------------------------------------------------------------
*** The MPI_Init() function was called before MPI_INIT was invoked.
*** This is disallowed by the MPI standard.
--------------------------------------------------------------------------
It looks like MPI_INIT failed for some reason; your parallel process is
likely to abort. There are many reasons that a parallel process can
fail during MPI_INIT; some of which are due to configuration or environment
problems. This failure appears to be an internal failure; here's some
additional information (which may only be relevant to an Open MPI
developer):
```

```
PML add procs failed
 --> Returned "Error" (-1) instead of "Success" (0)
--------------------------------------------------------------------------
*** Your MPI job will now abort.
[r1n2.esss.dk:11016] Abort before MPI_INIT completed successfully; not able to guarantee that all other
processes were killed!
```

7. On the ESS cluster, and optionally also on your own laptop, you can use the **mcsub_slurm** script **to help you** generate a batch file for the cluster. It takes a few input arguments:

   ```
   [willend@r1n2 ~]$ mcsub_slurm --help
   Usage: mcsub_slurm [options] [mcrun params]
    -h --help Show this help
    -rN --runtime=N Specify maximum runtime (hours) [default 1]
    -qQNAME --queue=QNAME Specify wanted SLURM queue [default 'quark']
    -e<mail> --email=<mail> Specify address to notify in reg. sim status [default none]
    --nodes=NUM Specify wanted number of nodes [default 1]
    --name=NAME Specify openPBS job name [default "McSub_<USERNAME>_<TIMESTAMP>"
   ```

   Usage: /mnt/lustre/apps/mcstas/2.4.1/bin/mcsub_slurm [options] [mcrun params]
   -h --help Show this help
   -rN --runtime=N Specify maximum runtime (hours) [default 1]
   -qQNAME --queue=QNAME Specify wanted SLURM queue [default 'express']
   --mpimodule=MODULE Specify wanted MPI module [default 'openmpi/3.0_gcc540']
   -e<mail> --email=<mail> Specify address to notify in reg. sim status [default none]
   --nodes=NUM Specify wanted number of nodes [default 1]
   --name=NAME Specify slurm job name [default "McSub_<USERNAME>_<TIMESTAMP>"]

   ]

   ```
   After running mcsub_slurm NAME.batch is ready for submission using the sbatch command
   ```

   (On McStas installations v. 2.4.1 and newer, the template batchfile writing is also available via File Configuration in mcgui on Linux / Python Preferences on macOS - meta-comma is a shortcut.)

8. To **generate a batch** file repeating a simulation that worked on your Desktop, use the **script "in front" of the mcrun command run at your Desktop**, e.g.:

   ```
   [willend@r1n2 ~]$ mcsub_slurm -qquark mcrun BNL_H8.instr -n1e6 Lambda=2.37
   ```

9. Having **seen the output from 8**, you should be able to request a **production queue and a longer runtime using the -q and -r options** of the script. (Or simply edit the script output to your taste and need!)

10. It is useful to receive output from the cluster jobs on start/termination etc. - please **use the -e option for setting a relevant recipient address**.

In case of issues with any of the above steps - feel free to contact support@esss.dk and peter.willendrup@esss.se