

Orion - Software - Fluent

Fluent on Orion Cluster

There are three main components to run a Fluent job on the cluster:

- Journal File (generated automatically in Fluent or can be made manually)
- Case and Data Files (Made in Fluent contains mesh, geometry, parameters, etc..)
- Bash File (File used to submit the case to the job scheduler on Orion)

Journal File

The journal file contains the commands which control the behavior of the program. An example is a simple read and run sequence:

```
/file/read-case-data "CASE_FILE"  
/solve/iterate 10000  
/file/write-case-data CASE_FILE  
exit
```

The first line reads the case and data. Followed by the solve, which tells Fluent to run for 10000 iterations. If your case was already setup in Fluent, this is the extent of the commands needed to run the case. The sequence ends with a write and exit command. There are two ways to generate this file. The first is to start the journalfile in Fluent

File->Write->Start Journal

This command can be implemented in the Console or by using the menu items. When you start the journaling process, it will record all the commands entered during the session, including the GUI commands. Once the session is over or you want to stop recording your commands you do the following.

File->Write->Stop Journal

If you know the commands you wish to run, writing your own journal file is also an option. The generated file can be used to run a session automatically:

File->Read->Journal

Or in this case, this file will be included in the bash script that is submitted to the cluster. More information on the journal file can be found here: https://www.sharcnet.ca/Software/Ansys/17.0/en-us/help/ai_sinfo/ai_sinfo.html

Case and Data File

I will assume prior knowledge of Fluent. These files hold all the important information for the simulation.

Bash File

The bash file is submitted the to the job scheduler, it provides the number of processors to use, the time required for the simulation, and can output a log file and an error file.

Here is an example of a bash file:

```

# Use the Bash Shell
#PBS -S /bin/bash
# Put a unique job name here
# so that you can monitor it in the queue
#PBS -N name_of_job
# Define the queue you're submitting to
#PBS -q production.q
# Notifying user by email at the beginning,end,abort,suspensions of the job run
#PBS -M youremail@clarkson.edu
#PBS -m eab
#PBS -V
# Uncomment to pass a single environment variable
# #PBS -v VAR_TO_PASS
# Redirecting standard output / error to files
# named "output" and "errors"
#PBS -o output
#PBS -e errors
# The max walltime for this job is 1 hour 31 minutes
2
# Set this longer than required for run!!!
# Set number of nodes/ppn to number of cores required: 8..1 is 8 nodes,
# 1 core per node
#PBS -l walltime=999:00:00,nodes=2:ppn=12
# Keep this.
cd $PBS_O_WORKDIR
# PUT SOME BASIC INFO INTO OUTPUT FILE
local_dir=$(pwd | cut -c 23-)
echo "This job was started in ${local_dir}"
echo "The job id is $JOB_ID"
echo "The start time was $(date)"
NSLOTS=`cat $PBS_NODEFILE | wc -l`
PE='mpich'

/share/apps/ansys_inc/v171/fluent/bin/fluent -r17.1.0 3ddp -pe $NSLOTS -pinfiniband -cnf=$PBS_NODEFILE -ssh -g -
i YOURJOURNAL.jou -mpi=pcmpi
# if a fluent job aborts, then you should use the kill script to delete the leftover # chmod +x <kill-fluent-
file>
# ./<kill-fluent-file>

```

Much of this doesn't need to be change. You will need to provide a name to the job, your email to be notified when the job starts and stops.

The important sections are the walltime and number of processors you want to run

```
#PBS -l walltime=999:00:00,nodes=2:ppn=12
```

The walltime is displayed in hrs:mins:secs The processes are split into nodes (there are 11 on the Orion cluster) and cores (ppn- there are 12 cores on each node).

Fluent is executed with the following command:

```
/share/apps/ansys_inc/v171/fluent/bin/fluent -r17.1.0 3ddp -pe $NSLOTS -pinfiniband -cnf=$PBS_NODEFILE -ssh -g -
i YOURJOURNAL.jou -mpi=pcmpi
```

This is a continuous line. The `r` is the version of Fluent and `3ddp` lets Fluent know that the job is three dimensional with double precision. The Journal file that is shown in this line has the extension 'jou'. This file will be in the same directory as the bash file (as well as the case and data files).

Access the cluster from the P&W lab

To run Fluent on the cluster from the P&W lab, you will need to install Putty, which is a ssh client for Windows machines <http://www.putty.org> (This might be already installed on the computers)

Assuming you have an account for the Orion cluster (if not you can email the Clarkson Help Desk (helpdesk@clarkson.edu) and ask for an account. You may need to let them know what Clarkson faculty you are working with.)

To access the cluster with Putty use the following command:

```
ssh username@orion.clarkson.edu
```

To transfer files to the cluster you can use the rsync command:

```
rsync -avzhe ssh file_to_upload username@orion.clarkson.edu:/Some_Directory
```

To download files from the cluster you can use the rsync command:

```
rsync -avzhe ssh username@orion.clarkson.edu:/Some_Directory/file_to_dow
```

More info about how to use rsync can be found here <https://www.tecmint.com/rsync-local-remote-file-synchronization-commands/>

You can also use Filezilla <https://filezilla-project.org> which is a ftp solution with a GUI.

Run the Job

To submit the job to the scheduler, provided you have the journal, case, data files.

```
qsub Name_of_Batch_file.sh
```

To check the status of the job:

```
qstat  
or  
qstat -a 'username'
```

(Displays the job info: name, ID, duration, etc...)

To look at the residuals:

```
tail -f output
```

This will allowing you to continuously follow the output. To stop the job manually:

```
qdel 'jodid'
```

Additional information about the cluster and the how to run Fluent jobs can be found here: <https://confluence.clarkson.edu/display/OITKB/Orion+Cluster>