

Basic job scripts for LSF

There are many possible choices for executing a job using LSF. Here is one approach. Do these steps:

Phase 1: Create an executable file

1) Write your program.

2) If you are writing parallel code, set your path to point to your choice of parallel library.

** PVM paths are set by default.

** To use MVAPICH Topspin MPI on Red Diamond, set the paths by adding the following lines to your `.bashrc` file:

```
export PATH=/usr/local/topspin/mpi/mpich/bin:$HOME/bin:$PATH
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/topspin/mpi/mpich/lib64
```

** To use MPICH-P4 MPI on either Red Diamond or Prospero, set the paths by adding the following lines to your `.bashrc` file:

```
export PATH=/opt/mpich/gnu/bin:$HOME/bin:$PATH
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/mpich/gnu/lib
```

3) Compile your program.

You can compile using any available compiler on RedD.

After your paths are set, you can compile using MVAPICH or MPICH-P4 using `mpicc`. Type something like:

```
mpicc -o trapezoid trapezoid.c
```

4) Make sure that the job executable is in your path. If this is not the case, give the full path to the executable.

Phase 2: Submit your program using LSF

1) write a script similar to one of the following files

2) type: `./yourlsfscriptfile` to execute the script file

3) Learn how to use some simple queue management commands:

```
bjobs
bjobs -l                #for a long listing)
bjobs -u all            #to look at the queues
bjobs -u all -q normal  #to look at the default queue
bkill <jobid>          #to kill a specific job
bpeek <jobid>          #to get debug info on a specific job
bhosts                  #to check host status
lsload                  #to check the load on the nodes. This
                        #also shows nodes that are unavailable
```

4) Output will be in the output file that you specify in the script file. If you don't specify an output file then, by default, the output will be sent you in an email. You will receive an email when the job completes, which you can suppress if you want to.

**5) There is a sample execution trace at very end of this file.
Have fun!**

```
#####
# MVAPICH - Topspin Example Script File      #
# compiled with mpicc -o trapezoid trapezoid.c #
#####

#!/bin/sh

#Number of cpus
NCPUS=8

#Output file
OUTFILE=%J.out

#A simple cpi program compiled with topspin

bsub -a mvapich -n $NCPUS -o $OUTFILE mpirun.lsf /home/aapon/lst/trapezoid 0 4

#####
# End of MVAPICH - Topspin Example Script File #
#####

#####
# PVM Example Script File #
#####

#!/bin/sh

#Number of cpus
NCPUS=8
```

```
#User output file
OUTFILE=%J.out
```

```
#The master1 example is included in pvm examples, it will spawn slave1
bsub -n $NCPUS -o $OUTFILE -a pvm mpirun.lsf master1
```

```
#####
# End of PVM Example Script File #
#####
```

```
#####
# MPICH-P4 Example Script File          #
# Use the /opt/mpich/gnu paths          #
#####
#!/bin/sh
```

```
#Number of cpus
NCPUS=16
```

```
#Output file, will get written in CWD, %J will be translated into the
#actual jobid
OUTFILE=%J.out
```

```
export P4_RSHCOMMAND=/usr/bin/ssh
```

```
#Run a sample job from the lsfadmin account, compiled cpi.c
bsub -o $OUTFILE -n $NCPUS -a mpichp4 mpirun.lsf /home/lsfadmin/MPICH_P4/mpichcpi
```

```
#####
# End of MPICH-P4 Example Script File #
#####
```

```
#####
# Example LSF job script for Dr. Spearot #
#####
```

```
#!/bin/sh
#Number of cpus
NCPUS=16
```

```
#Output file
OUTFILE=%J.out
```

```
#Input file
INPUT_FILE=in.shear
```

```
#Run the job
bsub -o $OUTFILE -n $NCPUS -a mpichp4 -i $INPUT_FILE mpirun.lsf $HOME/MPICH_P4/Imp_redd
```

```
#####
# End of Example LSF job script for Dr. Spearot#
#####
```

```
#####
#####
# Sample trace of execution using Topspin MPI          #
#####
#####
```

```
[aapon@reddiamond lsexamples]$ mpicc -o trapezoid trapezoid.c
```

```
[aapon@reddiamond lsexamples]$ ls -l
```

```
total 36
-rw-r--r--  1 aapon  aapon    552 May 19 10:27 lsfscrip
-rw-r--r--  1 aapon  aapon   3972 May 19 10:34 README
-rwxr-xr-x  1 aapon  aapon   9754 May 19 10:36 trapezoid
-rw-r--r--  1 aapon  aapon   4277 May 19 10:23 trapezoid.c
-rw-r--r--  1 aapon  aapon   1121 May 19 10:39 trapezoid.d
-rw-r--r--  1 aapon  aapon   3488 May 19 10:36 trapezoid.o
```

```
[aapon@reddiamond lsexamples]$ ./lsfscrip
Job <1381> is submitted to default queue <normal>.
```

```
[aapon@reddiamond lsexamples]$ bjobs
```

```
JOBID USER  STAT QUEUE   FROM_HOST EXEC_HOST JOB_NAME  SUBMIT_TIME
1381  aapon  RUN  normal  reddiamond 2*compute-0 *ezoid 0 4 May 19 10:39
                2*compute-0-2
                2*compute-0-3
                2*compute-0-8
```

```
[aapon@reddiamond lsexamples]$ bjobs
```

```
No unfinished job found
```

```
[aapon@reddiamond lsexamples]$ ls
```

```
1381.out lsfscrip README trapezoid trapezoid.c trapezoid.d trapezoid.o
```

```
[aapon@reddiamond lsexamples]$ more 1381.out
```

```
Sender: LSF System <lsfadmin@compute-0-0>
Subject: Job 1381: <mpirun.lsf /home/aapon/lsfexamples/trapezoid 0 4> Done
```

```
Job <mpirun.lsf /home/aapon/lsfexamples/trapezoid 0 4> was submitted from host <reddiamond> by user <aapon>.
```

```
Job was executed on host(s) <2*compute-0-0>, in queue <normal>, as user <aapon>.
```

```
<2*compute-0-2>
```

```
<2*compute-0-3>
```

```
<2*compute-0-8>
```

```
</home/aapon> was used as the home directory.
```

```
</home/aapon/lsfexamples> was used as the working directory.
```

```
Started at Fri May 19 10:39:13 2006
```

```
Results reported at Fri May 19 10:39:30 2006
```

```
Your job looked like:
```

LSBATCH: User input
mpirun.lsf /home/aapon/lsfexamples/trapezoid 0 4

Successfully completed.

Resource usage summary:

CPU time : 6.93 sec.
Max Memory : 2 MB
Max Swap : 15 MB

Max Processes : 1
Max Threads : 1

The output (if any) follows:

LSF_PJL_TYPE=mvapich
a=0.000000, b=4.000000, Local number of trapezoids=128
With n = 1024 trapezoids, our estimate
of the integral from 0.000000 to 4.000000 = 21.333326
Job /nfsutil/LSF/6.2/linux2.4-glibc2.2-x86_64/bin/mvapich_wrapper /home/aapon/lsfexample
s/trapezoid 0 4

TID	HOST_NAME	COMMAND_LINE	STATUS	TERMINATION_TIME
00000	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:35
00001	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:35
00002	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:57
00003	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:57
00004	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:51
00005	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:51
00006	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:27
00007	compute-0-	/home/aapon/lsfe	Done	05/19/2006 10:39:27