ANSYS CFX – Application **Execution Guidelines for vSMP Foundation Aggregated Virtual Machine**

Overview

ANSYS CFX supports the use of PVM, HPMPI and MPICH1 on Linux. This document includes execution guidelines only for HPMPI. HPMPI has a built-in mechanism for assigning MPI processes to specific CPUs. Process placement is controlled by environment variables named **MPI_BIND_MAP** and **MPIRUN_OPTIONS**. When these variables are not set, process placement will not be performed, and the operating system will choose where to run the different MPI ranks, per the OS scheduler algorithms.

NOTE: CFX is memory-bandwidth bound; this means that the capability/performance of the memory controller in the system is of higher importance than that of the CPU. On pre-Nehalem platforms, executing CFX processes at a quantity equal to the number of CPU cores on a system or board, may not yield any performance improvement, and potentially degrade performance, as the additional cores simply add to the saturation of the memory controller(s).

Running CFX with HP-MPI

Environment variables – HP-MPI

If you are running CFX with HPMPI, you should set the following environment variables prior to running CFX to yield the optimal performance:

export MPI_BIND_MAP=0,1,2,3,4,5,6,7 (For example)

export MPIRUN_OPTIONS="-cpu_bind=map_cpu,v"

export HPMP_FRAGSIZE=131072

export MPI_SHMEMCNTL=16,24000000,4000000

MPI_BIND_MAP specifies a list of CPUs to which MPI ranks will be bound. You should replace the list above with a list of integers, zero to #cpus-1.

For more information on HP-MPI CPU affinity settings, refer to the HP-MPI user's guide available from "http://docs.hp.com/en/B6060-96022/B6060-96022.pdf".

Process memory limitations

If you see a message similar to:

mid: Cannot create shared memory segment of 225255424 bytes

when you run CFX, then you should execute a command similar to:

sudo sh -c "echo 99999999999 > /proc/sys/kernel/shmmax"

to extend the amount of shared memory allowed per process.

Sample script for HP-MPI

```
---------------------------------------------------------------------------------------------------------------------------------
export PATH=/opt/ansys_inc/v110/CFX/bin:/opt/ansys_inc/v110/CFX/tools/hpmpi-2.2-5/Linux-amd64/bin:$PATH

export HPMP_FRAGSIZE=131072

export MPI_SHMEMCNTL=16,24000000,4000000

export MPI_BIND_MAP=23,22,21,20,19,18,17,16,15,14,13,12,11,10,9,8,7,6,5,4,3,2,1,0

export MPIRUN_OPTIONS="-cpu_bind=MAP_CPU,v"

list=”1 2 4 8 16 ...”

cfx5solve -part 16 -start-method “HP MPI Local Parallel" -def a.def -size-ni 2x -size-nr 2x

---------------------------------------------------------------------------------------------------------------------------------
```

Sample run with GUI

First set some HPMPI environment variables as described above.

```
---------------------------------------------------------------------------------------------------------------------------------
export PATH=/opt/ansys_inc/v110/CFX/bin:/opt/ansys_inc/v110/CFX/tools/hpmpi-2.2-5/Linux-amd64/bin:$PATH

export HPMP_FRAGSIZE=131072

export MPI_SHMEMCNTL=16,24000000,4000000

export MPI_BIND_MAP=23,22,21,20,19,18,17,16,15,14,13,12,11,10,9,8,7,6,5,4,3,2,1,0

export MPIRUN_OPTIONS="-cpu_bind=MAP_CPU,v"

---------------------------------------------------------------------------------------------------------------------------------
```

Next, launch “cfx5solve”, the GUI version of CFX solver.

Finally, click “File” -> “Define Run”, which will pop up a window as shown below.

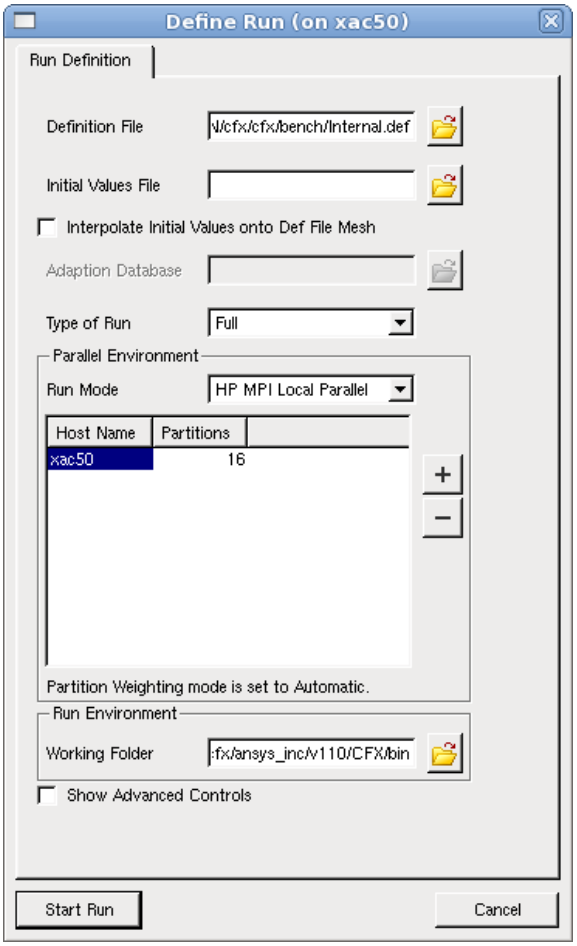

Specify the path to definition file, choose “HP MPI Local Parallel” in Run Mode tab, and then select the total number of processors to use.

References:

[1]. http://www.ansys.com/Products/cfx.asp