Intel MPI Cluster Edition on Graham – A First Look!

Doug Roberts SHARCNET / COMPUTE CANADA

Intel® Parallel Studio XE 2016 Update 4 Cluster Edition for Linux

1. Intel(R) MPI Library 5.1 Update 3 Cluster Ed
2. Intel(R) Trace Analyzer and Collector 9.1 Update 2 Cluster Ed
3. Intel(R) Cluster Checker 3.1 Update 2 Cluster Ed
4. Intel(R) VTune(TM) Amplifier XE 2016 Update 4 Professional Ed
5. Intel(R) Inspector XE 2016 Update 3 Professional Ed
6. Intel(R) Advisor XE 2016 Update 4 Professional Ed
7. Intel(R) C++ Compiler 16.0 Update 4
8. Intel(R) Fortran Compiler 16.0 Update 4
9. Intel(R) Math Kernel Library 11.3 Update 4
10. Intel(R) Integrated Performance Primitives 9.0 Update 4
11. Intel(R) Threading Building Blocks 4.4 Update 6
12. Intel(R) Data Analytics Acceleration Library 2016 Update 4
13. Intel(R) Debugger for Heterogeneous Compute 2016 Update 4
14. GNU GDB 7.8

Intel® Parallel Studio XE 2017 Update 4 Cluster Edition for Linux

1. Intel(R) Trace Analyzer and Collector 2017 Update 3 Cluster Ed
2. Intel(R) Cluster Checker 2017 Update 2 Cluster Ed
3. Intel(R) VTune(TM) Amplifier XE 2017 Update 3 Professional Ed
4. Intel(R) Inspector 2017 Update 3 Professional Ed
5. Intel(R) Advisor 2017 Update 3 Professional Ed
6. Intel(R) C++ Compiler 17.0 Update 4
7. Intel(R) Fortran Compiler 17.0 Update 4
8. Intel(R) Math Kernel Library 2017 Update 3 for C/C++
9. Intel(R) Math Kernel Library 2017 Update 3 for Fortran
10. Intel(R) Integrated Performance Primitives 2017 Update 3
11. Intel(R) Threading Building Blocks 2017 Update 6
12. Intel(R) Data Analytics Acceleration Library 2017 Update 3
13. Intel(R) MPI Library 2017 Update 3 Cluster Ed
14. Intel(R) Debugger for Heterogeneous Compute 2017 Update 4
15. GNU GDB 7.10
16. Intel(R) Debugger for Intel(R) MIC Architecture 2017 Update 4

Initialization of Components -Argument Specification

- mpivars.sh to specify mpi library config:

release	Single threaded optimized library
debug	Single threaded debug library
release_mt	Multi-threaded optimized library (Default)
debug_mt	Multi-threaded debug library

- mklvars.sh to specify mkl library configuration:

lp64	4 bytes integer
ilp64	8 bytes integer

Initialization of Components - The Intel Way

```
[roberpj@gra-login4:/opt/software/intel/17.0.4] find . -name \*vars.sh

./compilers_and_libraries_2017/linux/bin/compilervars.sh
./compilers_and_libraries_2017.4.196/linux/bin/iccvars.sh
./compilers_and_libraries_2017.4.196/linux/bin/ifortvars.sh
./compilers_and_libraries_2017.4.196/linux/daal/bin/daalvars.sh
./compilers_and_libraries_2017.4.196/linux/ipp/bin/ippvars.sh
./compilers_and_libraries_2017.4.196/linux/mkl/bin/mklvars.sh
./compilers_and_libraries_2017.4.196/linux/mpi/intel64/bin/mpivars.sh
./compilers_and_libraries_2017.4.196/linux/mpi/mic/bin/mpivars.sh
./compilers_and_libraries_2017.4.196/linux/tbb/bin/tbbvars.sh
```

Initialization of Components – The Module Way

```
[roberpj@gra-login2:~] module avail intelcluster

---------------------------------------- /opt/software/modules ----------------------------------
   intelcluster/2016.4                    intelcluster/mpi-debug/2016.4
   intelcluster/2017.4                    intelcluster/mpi-debug/2017.4      (D)
   intelcluster/daal/2016.4               intelcluster/mpi-debug_mt/2016.4
   intelcluster/daal/2017.4      (D)      intelcluster/mpi-debug_mt/2017.4   (D)
   intelcluster/ipp/2016.4                intelcluster/mpi-release/2016.4
   intelcluster/ipp/2017.4       (D)      intelcluster/mpi-release/2017.4    (D)
   intelcluster/mkl-ilp64/2016.4          intelcluster/mpi-release_mt/2016.4
   intelcluster/mkl-ilp64/2017.4 (D)      intelcluster/mpi-release_mt/2017.4 (D)
   intelcluster/mkl-lp64/2016.4           intelcluster/tbb/2016.4
   intelcluster/mkl-lp64/2017.4  (D)      intelcluster/tbb/2017.4            (D)
```

```
[roberpj@gra-login3:~] module avail intel

---------------------------------------- Core Modules ----------------------------------------
   intel/2016.4 (L,t,D:16:2016)    intel/2017.1 (t,17:2017)
```

* Where ILP64 means integer, long, and pointer data entities all occupy 8 bytes. Differs from the conventional LP64 model where only long and pointer data entities occupy 8 bytes while integer entities occupy 4 bytes.

Loading the Intel® Parallel Studio XE 2016/2017 Cluster Edition

export MODULEPATH=/opt/software/**module**s:$MODULEPATH; module unload intel imkl;

[roberpj@gra-login2:~] module load intelcluster/2017

Looking into the PATH environment variable whats new ?

Compilers:
[**roberpj**@gra-login2:/opt/software/intel/17.0.4/compilers_and_libraries_2017.4.196/linux/bin/intel64] ls i*

icc icc.cfg icc_libFNP.so icpc icpc.cfg icpc_libFNP.so ifort ifort.cfg ifort_libFNP.so

MPI Wrappers:
[roberpj@gra-login2:/opt/software/intel/17.0.4/compilers_and_libraries_2017.4.196/linux/mpi/intel64/bin] ls mpi*

mpicc mpicxx mpiexec.hydra mpif90 mpigcc mpiicc mpiifort
mpitune mpivars.sh mpicleanup mpiexec mpif77 mpifc mpigxx
mpiicpc mpirun mpivars.csh

[roberpj@orc-login1:/opt/sharcnet/intel/17.0.4/compilers_and_libraries_2017.4.196/linux/mpi/intel64/bin**]** ls mpi*

mpicleanup mpiexec mpiexec.hydra mpirun mpitune mpivars.csh mpivars.sh

The Two Sets of Intel MPI Wrappers

mpiifort mpiicc mpiicpc	Intel ifort Intel icc Intel icpc
mpif90 mpicc mpicxx	GNU gfortran GNU gcc GNU g++

To check: mpicc -show

Caveat! OpenMPI mpicc uses icc ...

Starting MPI Jobs

THE SRUN COMMAND (SLURM, RECOMMENDED)

This advanced method is supported by the Intel Library 4.0 Update 3. This method is the best integrated with Slurm and supports process tracking, accounting, task affinity, suspend/resume and other features.

Use the following commands to allocate a Slurm session and start an MPI job in it, or to start an MPI job within a Slurm session already created using the sbatch or salloc commands:

Set the the Slurm Process Management Interface (PMI) library:

```
export  I_MPI_PMI_LIBRARY=/path/to/slurm/pmi/library/libpmi.so
```

Launch the MPI job: `srun -n <num_procs> a.out`

https://slurm.schedmd.com/mpi_guide.html#intel_mpi

'Some' Intel References

Intel® MPI Library for Linux OS - Users Guide

https://software.intel.com/sites/default/files/User_Guide_0.pdf

Intel® MPI Library for Linux* OS - Reference Manual

https://software.intel.com/sites/default/files/Reference_Manual_1.pdf

Intel® MPI Library Developer Guide for Linux* OS

https://software.intel.com/en-us/mpi-developer-guide-linux-pdf

Intel® MPI Library Release Notes

https://software.intel.com/en-us/articles/intel-mpi-library-release-notes