

# Octopus Best Practices for Intel® Cluster Ready



BEST PRACTICES

## 1. Introduction

The following best practices document is provided as courtesy of the HPC Advisory Council.

## 2. Application Description:

Octopus is a scientific program aimed at the ab initio virtual experimentation on a hopefully ever-increasing range of system types. Electrons are described quantum-mechanically within density-functional theory (DFT), in its time-dependent form (TDDFT) when doing simulations in time. Nuclei are described classically as point particles. Electron-nucleus interaction is described within the pseudopotential approximation.

Octopus is one of selected 22 applications for the PRACE application benchmark suite

Octopus is free software, released under the GPL license and it is downloadable from

[http://www.tddft.org/programs/octopus/wiki/index.php/Main\\_Page](http://www.tddft.org/programs/octopus/wiki/index.php/Main_Page)

## 3. Version Information:

Version of this date is used: octopus 3.2.0

## 4. Prerequisites:

The instructions from this best practice have been tested with the following configuration:

### 4.1 Hardware:

- HP ProLiant SL2x170z G6 16-node cluster
- Intel Xeon X5670 CPUs @ 2.93 MHz
- Memory: 24GB per node @ 1333MHz
- Mellanox ConnectX-2 QDR InfiniBand Adapters
- Mellanox QDR InfiniBand Switch

### 4.2 Software:

- Intel® Cluster Ready running CentOS5U5
- Mellanox OFED 1.5.3 InfiniBand Software Stack
- Application: octopus 3.2.0
- Compilers: Intel compiler 11.1.064
- Math library: Intel MKL composerxe-2011.3.174, gsl-1.12
- MPI: Intel MPI 4, Open MPI 1.5.3, Platform MPI 8.0.1, MVAPICH2-1.6rc1

- Benchmark workload: benzene

## 5. Obtaining the source code

Octopus requires gsl (GNU Scientific Library) package on the testing system. It can be downloaded from <http://www.gnu.org/software/gsl/>

Octopus can be found from this URL:

[http://www.tddft.org/programs/octopus/wiki/index.php/Main\\_Page](http://www.tddft.org/programs/octopus/wiki/index.php/Main_Page)

## 6. Building Octopus

Extract gsl package and install it on each system or some shared directory

### Extract octopus-3.2.0.tar.gz

#### % tar -vzxf octopus-3.2.0.tar.gz

Run configure under octopus source directory, gsl, fftw, and lapack need to be specified. MPI support should be enabled by appending "--enable-mpi" during configure.

```
%.configure CC="mpicc" CFLAGS="-O3 -xSSE3 -axSSE4.2" FC="mpif90" FCFLAGS="-O3 -xSSE3 -axSSE4.2" --prefix=/app/octopus/octopus-3.2.0-openmpi-intel --with-gsl-prefix=/app/gsl-1.12-intel --with-fftw=fftw3 --with-fftw-lib=/usr/lib64/libfftw3.a --with-blas=/intel/mkl/lib/em64t/libmkl.so --with-lapack="/intel/mkl/lib/em64t/libmkl.so -lguid -lpthread" --enable-mpi
```

#### %make

At the end of the compilation, the executable named "octopus\_mpi" will be generated under bin directory.

## 7. Running Octopus

To run Octopus benchmark, create a temporary directory and copy benchmark files into it. Octopus will recognize inp as input file.

Create **inp** file as below:

```
CalculationMode = gs
Units = eV_Angstrom
radius = 8
spacing = 0.10
```

```
XYZCoordinates = "benzene.xyz"
EigenSolverInitTolerance = 1.0e-4
EigenSolverFinalTolerance = 1.0e-7
EigenSolverFinalToleranceIteration = 5
```

Then create benzene.xyz with following parameters:

```
12
  Geometry of benzene (in Angstrom)
C 0.000 1.396 0.000
C 1.209 0.698 0.000
C 1.209 -0.698 0.000
C 0.000 -1.396 0.000
C -1.209 -0.698 0.000
C -1.209 0.698 0.000
H 0.000 2.479 0.000
H 2.147 1.240 0.000
H 2.147 -1.240 0.000
H 0.000 -2.479 0.000
H -2.147 -1.240 0.000
H -2.147 1.240 0.000
```

Make sure remove all temporary files from previous run like exec, restart, and static files before a new run will be started.

#### Running with Intel MPI

```
% mpdboot --parallel-startup -r ssh -f < PATH_TO_
HOSTFILE > -n 16
```

```
% mpiexec -perhost 12 -env I_MPI_DEVICE rdssm
-np $i /octopus/octopus-3.2.0-impi-intel/bin/octopus_
mpi
```

#### Tune Intel MPI:

```
% mpiexec -perhost 12 -genv I_MPI_RDMA_
TRANSLATION_CACHE 1 -genv
```

```
I_MPI_ADJUST_ALLTOALL 1 -genv I_MPI_EXTRA_
FILESYSTEM on -genv
```

```
I_MPI_EXTRA_FILESYSTEM_LIST lustre -env I_MPI_
DEVICE rdssm -np 192 /octopus/octopus-3.2.0-impi-
intel/bin/octopus_mpi
```

```
%mpdallexit
```

#### Running with Open MPI

```
% mpirun --mca mpi_paffinity_alone 1 -np 48 -hostfile
~/hosts /octopus/octopus-3.2.0-openmpi-intel/bin/
octopus_mpi
```

Benchmark timing information is contained in output file. SCF time per step is the info should be captured.

