

Getting Started with

Dalma

High Performance Computing NYUAD



```
benoit — ssh -Y benoit@dalma.abudhabi.nyu.ec
[login-0-3.local 89: which matlab
/usr/bin/which: no matlab in (/home/benoit/bin:/share/app
mins/tools/scripts:/share/apps/admins/tools/bin:/usr/loc
al/sbin:/usr/sbin:/opt/ibutils/bin:/opt/slurm/default/bi
n:/share/apps/dalma/tools:/home/benoit/bin)
[login-0-3.local 90: module avail matlab
----- /share/apps/NYUAD/modules/ALL -----
matlab/R2015b matlab/R2016a matlab/R2016b matlab/R2017a
[login-0-3.local 91: module load matlab/R2016a
Loading module 'matlab/R2016a'
[login-0-3.local 92: which matlab
/share/apps/NYUAD/matlab/R2016a/bin/matlab
login-0-3.local 93: █
```

HPC systems differ from desktop/laptop computers in how software is installed. In a desktop/laptop you install software in a pre-defined centralized location where all users can access it. For instance the command "yum install python" will install python version 2.7.5 (on CentOS 7) under /usr/bin.

But in a large system hosting a multitude of users it is not practical to have a single version of "python" installed. Some users may need access to the latest version, while others have hard dependencies on an earlier version. And of course, this applies to all applications. So in an HPC system software is installed as independent modules which users must explicitly load into their shell environment.

Here's an example with matlab.

When you login your shell environment is empty of all software modules – **we strongly discourage loading modules automatically in your .bashrc!!**

"module avail matlab" shows you all available versions of matlab on Dalma. "module load matlab/R2016a" loads the R2016a version into your shell environment. So when you invoke "matlab" you actually use the correct version.



Modules: Commands

Software module environment are set on a per shell instance. So if you open a new terminal on a login node the environment you may have set in another terminal will not be propagated.

You can control and monitor the contents of your software modules environment with the following commands:

<code>module list</code>	- Display contents of your software modules environment
<code>module purge</code>	- Reset your software modules environment
<code>module avail</code>	- Display the list of all software modules available
<code>module load XXX</code>	- Load software module XXX into your environment
<code>module unload XXX</code>	- Unload software module XXX from your environment

As a guideline you should:

- Never add "module load" into your .bashrc file
- Use "module purge" before loading a new module environment to run an application
- Use a shell script to run an application and include the necessary modules in the script:

```
#!/bin/bash
module purge
module load matlab
matlab -nodisplay -nodesktop -r "mycode.m"
```

If you use "csh" then you need to initialize the software module environment before using "module" commands:

```
#!/bin/csh
source /usr/share/Modules/init/csh
module purge
module load matlab
matlab -nodisplay -nodesktop -r "mycode.m"
```

Dalma implements a one-of-a-kind software module environment – developed by the HPC team in Abu Dhabi – which is designed to make using software modules simpler and safer.

By simpler we mean that you need not load dependent software modules. For instance if you load netcdf into your environment, you need not load hdf5, nor zlib. All dependencies are automatically handled. So everytime you use netcdf you get the same software modules and get the same results.

Things are also simpler as when compiling applications and libraries you do not need to worry about specifying the directories from which to pick up libraries (`-L`) or include files (`-I`), and you do not need to list all dependent libraries. For example when linking with netcdf you do not need to list the dependent libraries, such as hdf5 and zlib; it's as simple as `"mpicc -o wrf wrf-main.o wrflibs.a -lnetcdf"`!

We also simplify things by making nearly all libraries compiler independent. So math libraries such as "lapack" can be used with both GNU and Intel compilers equally.

Finally, we simplify things further by deploying computationally intensive applications and libraries using an NYU Abu Dhabi designed method which automatically switches to the most efficient implementation based on the processor where you are currently executing. For instance if you load the "fftw2" math library at execution time the module will run the "sse", "avx", or "avx2" compiled version based on the processor you are currently using.



Modules: Dalma Specific

The following chart shows some software modules available on Dalma – the list is way to long to include it here...

The compiler agnostic column shows software modules can be used with GNU (GCC) or Intel compilers.

"SSE", "AVX", and "AVX2" tags show whether a software modules can automatically switch between a variety of processor architectures to maximize performance – eg can the tool benefit from floating-point calculation enhancements.

MPI tools can't be made compiler agnostic as they are tightly integrated with compilers and they include compiler specific Fortran90 components.

NYUAD 3									
NAME	VERSION	Compiler Agnostic	GCC (4.9.3)			Intel (16.1)			
			serial	openmpi 1.10.2	mvapich 2.2b	serial	openmpi 1.10.2	mvapich 2.2b	
admirable	2015.01	serial	serial	openmpi 1.10.2	mvapich 2.2b	serial	openmpi 1.10.2	mvapich 2.2b	
allinea	5.0	SSE							
amber	14	SSE							
atlas	3.10.2	AVX2	AVX2	AVX2	AVX2	AVX2	AVX2	AVX2	AVX2
autoconf	2.69	SSE							
autodock	4.2.6	SSE							
automake	1.15	SSE							
bedtools	2.25.0	SSE							
bias	3.6.0	SSE-AVX-AVX2							
boost	1.57.0			AVX2	AVX2			AVX2	AVX2
bzip2	1.0.6	SSE							
cbias	3.5.0	SSE-AVX-AVX2							
cfitsio	3.380	SSE-AVX2							

Modules: Dalma Specific

Dalma's one-of-a-kind software module environment is also safer to use than normal software module environments because it .

The use of software modules is not without risk. In a classical HPC system if you forget to load a module, change the order in which you load modules, or even merely load extra unrequired modules you may end up combining software components which weren't intended to be used together. Or maybe you will unintentionally pick up a different version of a library which will produce slightly different results. Software cross-contamination is usually the result of careless practices by users.

Dalma's software environment literally eliminates such risks by hardcoding within each application and library their software dependencies and by disabling the use of LD_PRELOAD and LD_LIBRARY_PATH. So users can't override the software modules.

Finally, all software modules installed on Dalma undergo testing prior to be deployed. In all, millions of tests are performed to verify the quality of the software environment at your disposal.

The result is a simple-to-use software module environment which gives you piece-of-mind as you are guaranteed to be able to reproduce computational results at will.



Modules: Dalma Specific

The software module environment on Dalma exists in 2 flavors. The default environment has automatically switches between "sse", "avx", and "avx2" versions based on the processor used, and it prevents cross-contamination between software modules. The extended flavor ("all") disables all safety measures and automatic switching of processor optimization and puts you in total control. This latter mode of use is for experts and is intended for specific non-standard problems.

Here's the list of ever growing software modules supported on Dalma in the "default" flavor.

```

----- /share/apps/NYUAD/modules/SOFTWARE -----
CCfits/2.5                cegma/2.5                gaussian/09.D.01         jemalloc/4.1.0          mvapich2/2.2bmxm        pio/2.2.3                szzip/2.1
Desmond-Maestro/2014.02  cfitsio/3.380            gaussview/5.0.9        julia/0.5.0            namd/2.10               pipelines/3.4.0         tangram/0.3.1
R/3.1.2(default)        charm/6.7.0              gcc/4.9.3              lammps/2014.12         nbo/6                   platanus/1.2.4         tcltk/8.5.17
R/3.4.1                  clapack/3.2.1            geant4/10.3.2          lapack/3.6.0           ncap2/4.4.4            pnetcdf/1.7.0         tedna/1.2.2
admirable/2015.01       cmake/3.4.2              genblasta/1.0.4        libffi/3.2.1           ncl/6.3.0              presto/2017.04        tempo/2017.05
allinea/5.0              comsol/4.4               gencore/1              libgd/2.1.1            nco/4.6.3              psrchive/6.0          thrax/1.2.3
allpaths1g/52488        cp2k/2.6.0               glib/2.44.0            libpng/1.6.24          ncurses/6.0            pyferret/7.0.0        trimomatic/0.32
amber/14                 cpmc/3.17.1              gmp/6.0.0              libtool/2.4.6          ncview/2.1.7           python/2.7.11         udunits/2.2.24
anaconda/2-4.1.1        crystall4/v1.0.1         gnuplot/5.0.0          libunwind/1.2          netcdf/4.4.0           qchem/4.4             valgrind/3.11.0
ansys/v172               cuda/7.0                 gperftools/170515      libxml2/2.9.2          netcdf-fortran/4.4.2   readline/6.3          vapor/2.5.0
apbs/1.5                 cuda/7.5                 grads/2.0.1            m4/1.4.17              ngsTools/1.0.0         root/6.06.02          vcftools/0.1.14
arpack/96                cufflinks/2.2.1          greasy/2.1             maps/4.0                oasis/3.0               ruby/2.3.0            velvet/1.2.10
atlas/3.10.2            curl/7.54.1              gromacs/5.0.4          mathematica/10.3.0     openfoam/2.4.0         scalapack/2.0.2       velvet-k127/1.2.10
autoconf/2.69           discover/52488            gsl/2.1                matlab/R2015b          openfst/1.6.1          scotch/6.0.4          wgrib/1.8
autodock/4.2.6          discovardenovo/52488     hdf5/1.8.16           mbedtls/2.6.0          openmpi/1.10.2        screen/4.3.1          wien2k/13(default)
automake/1.15            expat/2.1.0              help2man/1.47.3        montjoie/2017.04      openssl/1.0.1g         soapdenovo2/r240     wien2k/16
bedtools/2.25.0         ferret/7.0               hmmer/3.1b2            mosesdecoder/3.0      orca/3.0.3             soapdenovo2-prepare/2.0 workbench/1.0
bison/3.0.4             fftw/2.1.5              idl/8.6                mpfr/3.1.5             oxdna/2.2.2            sofia-ml/2015.01     xconv/1.93
blas/3.6.0              fftw3/3.3.4             idv/5.3                mpilacs/3              papi/5.5.1            spades/3.7.1          xios/1.0
blast/2.3.0             fftw3_mpi/3.3.4         intel/16.1             mumps/5.0.1            parallel/20160222     sparsehash/2.0.2     zlib/1.2.8
boost/1.57.0            fregrid/08082013        jags/3.4.0             musket/1.1             paraview/5.3          stata/14              svmtool/1.3.1
bzip2/1.0.6            fsl/5.0.6               jdk/1.8.0_31           mvapich2/2.2b         parmetis/4.0.3        superlu/5.0
castep/16.1            fuse/2.9.3              jellyfish/2.2.5        mvapich2/2.2b         pdb2pqr/2.1.1        perl/5.22.1
cblas/3.6.0             gapcloser/1.12r6        jdk/1.8.0_31           musket/1.1             perl/5.22.1           perl/5.22.1
cdo/1.7.2              gapcloser/1.12r6        jellyfish/2.2.5        mvapich2/2.2b         pgplot/5.2

```

Here's the extended list of modules. To switch from "default" to "extended" use "module purge;module load all" and to switch back to "default" use "module purge".

```
----- /share/apps/NYUAD/modules/ALL -----
CCfits/avx2/2.5
Desmond-Maestro/2014.02
R/avx2/3.1.2
R/avx2/3.4.1
R/sse4.2/3.1.2
R/sse4.2/3.4.1
admirable/avx2/2015.01
allinea/5.0
allpaths1g/52488
amber/openmpi_1.10.2/sse4.2/14
amber/sse4.2/14
anaconda/2-4.1.1
anaconda/3-4.3.1
ansys/gcc_4.9.3/openmpi_1.10.2/v172
arpack/1.5
arpack/96
atlas/avx2/3.10.2
autoconf/2.69
autodock/4.2.6
automake/1.15
bedtools/2.25.0
bison/3.0.4
blas/avx/3.6.0
blas/avx2/3.6.0
blas/sse4.2/3.6.0
blast/2.3.0
boost/gcc_4.9.3/openmpi_1.10.2/avx2/1.57.0
boost/gcc_4.9.3/openmpi_1.10.2/sse4.2/1.57.0
bzip2/1.0.6
castep/openmpi_1.10.2/avx2/16.1
cblas/avx/3.6.0
cblas/avx2/3.6.0
cblas/sse4.2/3.6.0
cdo/1.7.2
cegm/2.5
cfitsio/avx2/3.380
cfitsio/sse4.2/3.380
charm/openmpi_1.10.2/sse4.2/6.7.0
clapack/avx/3.2.1
clapack/avx2/3.2.1
clapack/sse4.2/3.2.1
cmake/3.4.2
cmake/3.9.4
comsol/4.4
cp2k/avx2/2.6.0
cp2k/openmpi_1.10.2/avx2/2.6.0
cpm/2.0
cpm/openmpi_1.10.2/avx2/3.17.1
crystal14/v1.0.1
cuda/gcc_4.9.3/single_node/7.0
cuda/gcc_4.9.3/single_node/7.5
cufflinks/2.2.1
curl/7.54.1
discover/52488
discoverdenovo/52488
eigen/3.2.8
expat/2.1.0
ferret/7.0
fftw2/avx2/2.1.5
fftw2/avx2/2.1.5
fftw2/gcc_4.9.3/mvapich2_2.2b/avx2/2.1.5
fftw2/gcc_4.9.3/mvapich2_2.2b/avx2/2.1.5
fftw2/gcc_4.9.3/mvapich2_2.2b/sse4.2/2.1.5
fftw2/gcc_4.9.3/openmpi_1.10.2/avx2/2.1.5
fftw2/gcc_4.9.3/openmpi_1.10.2/sse4.2/2.1.5
fftw2/intel_16.1/mvapich2_2.2b/avx2/2.1.5
fftw2/intel_16.1/mvapich2_2.2b/sse4.2/2.1.5
fftw2/intel_16.1/mvapich2_2.2b/sse4.2/2.1.5
fftw2/intel_16.1/openmpi_1.10.2/avx2/2.1.5
fftw2/intel_16.1/openmpi_1.10.2/avx2/2.1.5
fftw2/intel_16.1/openmpi_1.10.2/sse4.2/2.1.5
fftw2/intel_16.1/openmpi_1.10.2/sse4.2/2.1.5
fftw3/avx/3.3.4
fftw3/avx2/3.3.4
fftw3/gcc_4.9.3/mvapich2_2.2b/avx/3.3.4
fftw3/gcc_4.9.3/mvapich2_2.2b/avx2/3.3.4
fftw3/gcc_4.9.3/mvapich2_2.2b/sse4.2/3.3.4
fftw3/gcc_4.9.3/openmpi_1.10.2/avx/3.3.4
fftw3/gcc_4.9.3/openmpi_1.10.2/avx2/3.3.4
fftw3/gcc_4.9.3/openmpi_1.10.2/sse4.2/3.3.4
fftw3/intel_16.1/mvapich2_2.2b/avx/3.3.4
fftw3/intel_16.1/mvapich2_2.2b/avx2/3.3.4
fftw3/intel_16.1/mvapich2_2.2b/sse4.2/3.3.4
fftw3/intel_16.1/openmpi_1.10.2/avx2/3.3.4
fftw3/intel_16.1/openmpi_1.10.2/sse4.2/3.3.4
fftw3/sse4.2/3.3.4
fregrid/gcc_4.9.3/openmpi_1.10.2/08082013
fsl/sse4.2/5.0.6
fuse/2.9.3
gapcloser/1.12r6
gaussian/09.D.01
gaussview/5.0.9
gcc/4.9.3
geant4/sse4.2/10.3.2
genblast/1.0.4
gencore/1
glib/2.44.0
gmp/6.0.0
gnuplot/5.0.0
gperftools/170515
grads/2.0.1
grads/2.1.0
greasy/openmpi_1.10.2/2.1
gromacs/openmpi_1.10.2/avx2/5.0.4
gromacs/openmpi_1.10.2/sse4.2/5.0.4
gs/avx/2.1
gs/avx2/2.1
gs/sse4.2/2.1
gvim/gcc_4.9.3/single_node/7.4
hd5f/gcc_4.9.3/1.8.16
hd5f/gcc_4.9.3/mvapich2_2.2b/1.8.16
hd5f/gcc_4.9.3/openmpi_1.10.2/1.8.16
hd5f/intel_16.1/1.8.16
help2man/1.47.3
hmm/3.1b2
idl/8.6
idv/5.3
intel/16.1
jags/3.4.0
jdk/1.8.0_31
jellyfish/2.2.5
jemalloc/4.1.0
julia/0.5.0
lammps/omp/avx2/2014.12
lammps/openmpi_1.10.2/omp/avx2/2014.12
lammps/openmpi_1.10.2/omp/avx2/2015.12
lapack/avx/3.6.0
lapack/avx2/3.6.0
lapack/sse4.2/3.6.0
libffi/3.2.1
libgd/2.1.1
libpng/1.6.24
libtool/2.4.6
libunwind/1.2
libxml2/2.9.2
m4/1.4.17
maps/4.0
mathematica/10.3.0
matlab/R2015b
matlab/R2016a
matlab/R2016b
matlab/R2017a
mbedtls/2.6.0
mct/2.9.0
mct/gcc_4.9.3/mvapich2_2.2b/2.9.0
mct/gcc_4.9.3/openmpi_1.10.2/2.9.0
mct/intel_16.1/mvapich2_2.2b/2.9.0
mct/intel_16.1/openmpi_1.10.2/2.9.0
metis/avx2/5.1.0
metis/sse4.2/5.1.0
mgjiza/2015.01
montjoie/avx2/2017.04
mosesdecoder/3.0
mpfr/avx2/3.1.5
mpfr/sse4.2/3.1.5
mpiblacs/gcc_4.9.3/openmpi_1.10.2/avx2/3
mpiblacs/gcc_4.9.3/openmpi_1.10.2/sse4.2/3
mumps/avx2/5.0.1
mumps/openmpi_1.10.2/avx2/5.0.1
mumps/sse4.2/5.0.1
mumps/openmpi_1.10.2/sse4.2/5.0.1
mumps/sse4.2/5.0.1
musket/1.1
mvapich2/gcc_4.9.3/2.2b
mvapich2/gcc_4.9.3/2.2bmxm
mvapich2/intel_16.1/2.2b
namd/openmpi_1.10.2/avx2/2.10
nbo/6
ncap2/4.4.4
ncl/6.3.0
nco/4.5.2
nco/4.6.3
ncurses/6.0
ncview/2.1.7
netcdf/4.4.0
netcdf/4.4.0test
netcdf/gcc_4.9.3/mvapich2_2.2b/4.4.0
netcdf/gcc_4.9.3/openmpi_1.10.2/4.4.0
netcdf/intel_16.1/mvapich2_2.2b/4.4.0
netcdf/intel_16.1/openmpi_1.10.2/4.4.0
netcdf-fortran/gcc_4.9.3/4.4.2
netcdf-fortran/gcc_4.9.3/mvapich2_2.2b/4.4.2
netcdf-fortran/gcc_4.9.3/openmpi_1.10.2/4.4.2
netcdf-fortran/intel_16.1/4.4.2
netcdf-fortran/intel_16.1/mvapich2_2.2b/4.4.2
netcdf-fortran/intel_16.1/openmpi_1.10.2/4.4.2
ngsTools/1.0.0
oasis/3.0
openfoam/gcc_4.9.3/openmpi_1.10.2/2.4.0
openfst/1.6.1
openmpi/gcc_4.9.3/1.10.2
openmpi/gcc_4.9.3/1.10.2mxm
openmpi/gcc_4.9.3/1.6.5
openmpi/gcc_4.9.3/2.1.1mxm
openmpi/intel_16.1/1.10.2
openssl/1.0.1q
orca/3.0.3
oxdna/2.2.2
papi/5.5.1
parallel/20160222
paraview/5.3
parmetis/openmpi_1.10.2/avx2/4.0.3
parmetis/openmpi_1.10.2/sse4.2/4.0.3
pdb2pqr/2.1.1
perl/5.22.1
pgplot/avx2/5.2
pio/gcc_4.9.3/openmpi_1.10.2/2.2.3
pipelines/3.4.0
platanus/1.2.4
pnetcdf/gcc_4.9.3/mvapich2_2.2b/1.0.2
pnetcdf/gcc_4.9.3/mvapich2_2.2b/1.7.0
pnetcdf/gcc_4.9.3/openmpi_1.10.2/1.0.2
pnetcdf/gcc_4.9.3/openmpi_1.10.2/1.7.0
pnetcdf/intel_16.1/mvapich2_2.2b/1.0.2
pnetcdf/intel_16.1/mvapich2_2.2b/1.7.0
pnetcdf/intel_16.1/openmpi_1.10.2/1.0.2
pnetcdf/intel_16.1/openmpi_1.10.2/1.7.0
presto/openmpi_1.10.2/avx2/2015.02
presto/openmpi_1.10.2/avx2/2017.04
parchive/6.0
pyferret/7.0.0
python/2.7.11
python/3.5.1
qchem/gcc_4.9.3/openmpi_1.10.2/4.4
readline/6.3
root/sse4.2/5.34.36
root/sse4.2/6.06.02
ruby/2.3.0
scalapack/gcc_4.9.3/openmpi_1.10.2/avx2/2.0.2
scalapack/gcc_4.9.3/openmpi_1.10.2/avx2/2.0.2
scalapack/gcc_4.9.3/openmpi_1.10.2/sse4.2/2.0.2
scotch/openmpi_1.10.2/avx2/6.0.4
scotch/openmpi_1.10.2/sse4.2/6.0.4
screen/4.3.1
soapdenovo2/r240
soapdenovo2-prepare/2.0
sofia-ml/2015.01
spades/3.7.1
sparsehash/2.0.2
sqlite/3.11.0
srilm/1.6.0
stata/14
superlu/avx2/5.0
superlu/openmpi_1.10.2/avx2/dist_4.1
superlu/openmpi_1.10.2/sse4.2/dist_4.1
superlu/sse4.2/5.0
svmtool/1.3.1
szip/2.1
tangram/0.3.1
tcltk/8.5.17
tedna/1.2.2
tempo/avx2/2015.02
tempo/avx2/2017.05
thrax/1.2.3
trimomatic/0.32
udunits/2.2.24
valgrind/3.11.0
vapor/2.5.0
vftools/0.1.14
velvet/1.2.10
velvet-kl27/1.2.10
wgrib/1.8
wien2k/openmpi_1.10.2/avx2/13
wien2k/openmpi_1.10.2/avx2/16
wien2k/openmpi_1.10.2/sse4.2/13
wien2k/openmpi_1.10.2/sse4.2/16
workbench/1.0
xcom/1.93
xios/gcc_4.9.3/openmpi_1.10.2/1.0
zlib/1.2.8
```




You can find documentation, tutorials, presentations, and examples on how to use Dalma on the login nodes, and on Dalma's support web site.

Our support is divided into 2 categories:

1. User Support : for your everyday questions about how to do things on Dalma. You can reach us through raising a ticket nyuad.it.help@nyu.edu or walking in to A2-132B.
2. Application Support : application porting, tuning, development, and research collaboration. We work on a project basis, you can reach us in A2-131.

```
DALMA()                               Dalma Manual                               DALMA()

SYNOPSIS
  Here you will find a list of all the available topics for help related
  to Dalma and how to call for help

DESCRIPTION
  Users can request help on any specific topic using the following com-
  mand:

  man dalma-<topic>

MAIN TOPICS
  where 'topic' can be one of:
  o storage : Display information about the different filesystems and
             how to use them
  o queues : Display the different login information and when to use
             them
  o modules : NYUME login description. How to access to the available
             ...
--More--
```

- Login Nodes: "man dalma"
- Web Site: <https://wikis.nyu.edu/display/ADRC/HPC>