LMOD: a hierarchical system for software modules

Ramses van Zon

SciNet HPC Consortium

December 9, 2015

SciNet compute | calcul canada

Content

1. What are software modules and why do we need them?
2. Current setup at SciNet
3. LMOD
4. How to use LMOD setup on the GPC
5. Demonstration

SciNet compute | calcul canada

1.What are software modules and why do we need them?

SciNet compute canada | calcul canada

Why modules?

- In Linux, there are standard paths for libraries and header files.
  /bin, /usr/bin, /usr/include, /usr/local/bin, ...
- Only one version could be installed.
- Different users have different needs. On a shared system, can served everyone only by not installing almost any software in standard spots.
- Instead, they are installed in non-standard directories. Environment variables can be set to point the os and build tools to where stuff is.

Modules...

Modules are a mechanism to set and unset these environment variables and to express the prerequisites and conflicts.

SciNet compute canada | calcul canada

Example

- C compiler is in folder `/strangepath`
- Compiler executables in `/strangepath/bin` must be added to `PATH` variable.
- It is called `scc`, so we'd set the `CC` variable to `scc`
- Applications and libraries produced by this compiler need to load dynamic libraries in `/strangepath/lib` must be added to `LD_LIBRARY_PATH` .
- A short decription would suffice:

```
# I am a pseudo-module called scc
addto  PATH /strangepath/bin
addto  LD_LIBRARY_PATH /strangepath/lib
setvar CC scc
```

This is not the real syntax!

SciNet compute canada | calcul canada

Loading, unloading

```
# I am a pseudo-module called scc
addto  PATH /strangepath/bin
addto  LD_LIBRARY_PATH /strangepath/lib
setvar CC scc
```

SciNet compute | calcul canada

Loading, unloading

```
# I am a pseudo-module called scc
addto  PATH /strangepath/bin
addto  LD_LIBRARY_PATH /strangepath/lib
setvar CC scc
```

- This is a mockup of what a real module file could look like.

SciNet compute | calcul canada

Loading, unloading

```
# I am a pseudo-module called scc
addto  PATH /strangepath/bin
addto  LD_LIBRARY_PATH /strangepath/lib
setvar CC scc
```

- This is a mockup of what a real module file could look like.
- When the module is 'loaded', the file is read line by-line; the specified paths are added to and environment variables are set.

SciNet compute | calcul canada

Loading, unloading

```
# I am a pseudo-module called scc
addto  PATH /strangepath/bin
addto  LD_LIBRARY_PATH /strangepath/lib
setvar CC scc
```

- This is a mockup of what a real module file could look like.
- When the module is 'loaded', the file is read line by-line; the specified paths are added to and environment variables are set.
- When it is 'unloaded', the reverse action is attempted.

SciNet compute | calcul canada

Dependencies and conflicts

- If an application or library compile with this compiler becomes module itself, it requires the scc to be loaded to, e.g.

```
# I am a pseudo-module called scclib/1.0
prereq scc
addto  PATH            /otherpath/bin
addto  CPATH           /otherpath/include
addto  LD_LIBRARY_PATH /otherpath/lib
```

SciNet compute | calcul canada

Dependencies and conflicts

- If an application or library compile with this compiler becomes module itself, it requires the scc to be loaded to, e.g.

```
# I am a pseudo-module called scclib/1.0
prereq scc
addto  PATH            /otherpath/bin
addto  CPATH           /otherpath/include
addto  LD_LIBRARY_PATH /otherpath/lib
```

- Conflicts: Could have two modules that are not allowed to be used at the same time:

```
# I am a pseudo-module called scclib/2.0
prereq scc
conflict scclib/1.0
addto  PATH            /otherpath/2.0/bin
addto  CPATH           /otherpath/2.0/include
addto  LD_LIBRARY_PATH /otherpath/2.0/lib
```

Current setup at SciNet

Current module system at SciNet

- Using "Environment Modules", i.e. `modulecmd`.
- Uses the `module` command.
- Module files written in `tcl`

Module commands

module load NAME	Load module named NAME
module unload NAME	Undo loading of NAME
module purge	Unloading everything
module avail [NAME]	List all modules
module list	List all loaded modules
module help NAME	Show info about module NAME
module find [KEYWORD]	Find module containing KEYWORD
module advice NAME	Undo loading of NAME

Note: the latter two are SciNet extensions.

Issue 1

- Large number of modules: nearly 500 active ones, and about 200 deprecated or experimental ones.

```
$ module avail
------------- /scinet/gpc/Modules6/Modules/versions -------------
3.2.8 3.2.9

--------- /scinet/gpc/Modules6/Modules/3.2.9/modulefiles ---------
dot                    modules                 use.deprecated
module-cvs             null                    use.experimental
module-info            use.defaults/2013-09    use.own

------------ /scinet/gpc/Modules6/Modules/modulefiles ------------
BGW-paratec/1.0.4-2.0.0436          intelmpi/4.1.2.040(default)
CPLEX/12.6.2                        intelmpi/5.0.3.048
EIGENSOFT/6.0.1                     inteltools/2013(default)
GEOS/3.2.0                          inteltools/2015
GEOS/3.5.0                          ipm/0.983-gcc-ompi-mpiio
ImageMagick/6.6.7(default)          ipm/0.983-gcc-ompi-posix
ImageMagick/6.9.1-4                 ipm/0.983-intel-intelmpi-mpi
```

. . .

```
MCR/2012b                                  ipm/0.983-intel-intelmpi-posi
MCR/2014a                                  ipm/0.983-intel-ompi-mpiio
Minimac3/1.0.11                            ipm/0.983-intel-ompi-posix(de
R/2.13.1                                   jasper/1.900.1-intel
R/2.14.1                                   java/6.0
R/2.15.1(default)                          java/7.1
R/3.0.0                                    java/8.0
R/3.0.1                                    java/ibm-java-x86_64-60
R/3.1.1                                    java/ibm-java-x86_64-71
ROOT/5.30.03(default)                      java/oracle-java-x86_64-80_45
ROOT/5.32.00                               libconfuse/2.7(default)
ROOT/5.34.03                               libint/1.1.5
ROOT/6.02.02                               libint/2.0.3(default)
Xlibraries/X11-64                          libxc/1.1.0
abinit/7.10.4                              libxc/2.0.1
abyss/1.3.2                                libxc/2.0.2(default)
adios/131-openmpi-gcc(default)             libxc/2.2.0
allpathslg/47998                           luajit/2.0.3
amber/10.0.30                              luarocks/2.2.0
```

. . .

```
amber/14.0                                  make/3.81
antlr/2.7.7                                 mc/4.7.0(default)
armadillo/3.6.2                             mc/4.8.14
armadillo/3.910.0(default)                  meep/1.1.1-openmpi
arpack-ng/3.1.5                             meep/1.1.1-openmpi-shdf5
autoconf/2.68                               meep/1.1.1-serial
autoconf/2.69                               meep/1.2-intelmpi-shdf5
automake/1.11.2                             meep/1.2.1-intelmpi-shdf5
bbcp/14.07.01.00.0(default)                 mesa/7.11(default)
bedtools/2.21.0                             mesa/CentOS4
binutils/2.25                               mesa/CentOS6.4
blast/2.2.23+                               metis/5.0.2
blast/2.2.30+                               mono/3.12.0
blat/35                                     mothur/1.24(default)
bowtie2/2.2.6                               mothur/1.27
cactus/0.0                                  mpb/1.4.2-openmpi-shdf5(defa
caf/intel/any                               mpb/1.4.2-serial
cairo/1.12.8-intel                          mpe/1.3.0-gcc-ompi
casacore/1.7.0                              mpe/1.3.0-intel-intelmpi
```

. . .

```
casacore/2.0.1                          mpe/1.3.0-intel-ompi(default)
cdo/1.5.1                               mpich1/pgi
cdo/1.5.4(default)                      mysql/5.6.12
cdo/1.6.1-intel                         namd/2.10-ibverbs-smp
cdo/1.6.9-intel                         namd/2.10-ibverbs-smp-CUDA
centos5-compat/lib64                    namd/2.10-multicore-CUDA
centos6-compat/lib                      namd/2.8-ibverbs
cfitsio/3.370                           namd/2.8-ibverbs-smp
cmake/2.8.12.2                          namd/2.9-ibverbs-smp(default)
cmake/2.8.6                             namd/2.9-ibverbs-smp-CUDA
cmake/2.8.8(default)                    namd/2.9-multicore-CUDA
cmake/3.1.0                             nano/2.2.4
cmake/3.4.0                             ncl/6.0.0
coreutils/8.15                          ncl/6.1.0(default)
cp2k/2.4.0                              ncl/6.2.0
cp2k/2.4.0-intel                        ncl/6.3.0
cp2k/Trunk-25Feb2014-intel              ncl/6.3.0-intel
cpmd/3.13.2                             nco/4.0.8-gcc
cuda/3.2                                nco/4.0.8-intel-nocxx
```

. . .

```
cuda/4.0                                  nco/4.3.2-intel
cuda/4.1(default)                         nco/4.4.8-intel
cuda/4.2                                  ncview/2.1.1(default)
cuda/5.0                                  ncview/2.1.2
cuda/5.5                                  ndiff/2.0.0
cuda/6.0                                  nedit/5.5
cuda/6.5                                  netcdf/4.1.3_hdf5_intelmpi-i
cxxlibraries/boost/1.47.0-gcc             netcdf/4.1.3_hdf5_openmpi-gc
cxxlibraries/boost/1.47.0-intel           netcdf/4.1.3_hdf5_openmpi-in
cxxlibraries/boost/1.54.0-gcc4.8.1        netcdf/4.1.3_hdf5_serial-gcc
cxxlibraries/boost/1.54.0-intel           netcdf/4.1.3_hdf5_serial-int
cxxlibraries/boost/1.55.0-intel           netcdf/4.2.0_hdf5_openmpi-gc
cxxlibraries/rlog/1.4-gcc                 netcdf/4.2.1.1_intelmpi-inte
dcap/2.47                                 netcdf/4.2.1.1_openmpi-intel
ddd/3.3.12                                netcdf/4.2.1.1_serial-gcc
ddt/4.0                                   netcdf/4.2.1.1_serial-intel
ddt/4.1                                   netcdf/4.3.2_openmpi-intel
ddt/4.2.1                                 netcdf/4.3.3.1_intelmpi-int
ddt/5.0(default)                          netcdf/4.3.3.1_openmpi-inte
```

. . .

```
discovar/47982
emacs/23.3
emacs/24.4
emi2workernode/2.5.1
encfs/1.7.4
erlang/r14b
espresso/4.3.2(default)
espresso/5.0.3
espresso/trunk
extras/64
extras/64_6.3
extras/64_6.4(default)
ferret/6.72
ffmpeg/2.1.1
ffmpeg/2.1.3
fftw/2.1.5-gcc-openmpi
fftw/2.1.5-intel-intelmpi4
fftw/2.1.5-intel-openmpi
fftw/3.3.0-gcc-openmpi
netcdf/4.3.3.1_serial-gcc
netcdf/4.3.3.1_serial-intel
nlopt/2.4.2
nwchem/6.0
nwchem/6.5
octave/3.4.3(default)
octave/3.8.1
octave/4.0.0
octopus/4.1.1
octopus/4.1.2(default)
openbabel/2.3.2
openblas/0.2.13-gcc-openmp
openblas/0.2.13-gcc-serial
openblas/0.2.13-gcc-threaded
openblas/0.2.13-intel-openmp
openblas/0.2.13-intel-serial
openblas/0.2.13-intel-thread
openblas/1.13-multithreaded
openblas/1.13-singlethreaded
```

. . .

```
fftw/3.3.0-intel-impi                          openfoam/2.1.0
fftw/3.3.0-intel-openmpi                       openfoam/2.3.0
fftw/3.3.3-gcc-openmpi                         openmpi/1.4.4-gcc-v4.6.1
fftw/3.3.3-intel-impi                          openmpi/1.4.4-intel-v12.1(de
fftw/3.3.3-intel-openmpi(default)              openmpi/gcc/1.6.4
fftw/3.3.4-gcc-openmpi                         openmpi/gcc/1.8.3
fftw/3.3.4-intel-impi                          openmpi/intel/1.6.4
fftw/3.3.4-intel-openmpi                       openspeedshop/2.1-gcc-openmp
gamess/1May2013                                padb/3.2
gamess/Aug1811R1                               papi/4.1.3
gcc/4.4.6                                      parallel-netcdf/1.2.0_intelm
gcc/4.6.1(default)                             parallel-netcdf/1.2.0_openmp
gcc/4.7.0                                      parallel-netcdf/1.2.0_openmp
gcc/4.7.2                                      parallel-netcdf/1.3.1_intelm
gcc/4.8.1                                      parallel-netcdf/1.3.1_openmp
gcc/4.9.0                                      paraview/3.12(default)
gcc/5.2.0                                      paraview/3.14.1
gcclib/4.4.6                                   paraview/4.1.0
```

canada | canada

. . .

```
gcclib/4.6.1                          petsc/3.1_intel_impi
gcclib/4.7.0                          petsc/3.1_intel_impi_cxx
gcclib/4.7.2                          petsc/3.1_intel_openmpi(defau
gcclib/4.8.1                          petsc/3.1_intel_openmpi_cxx
gcclib/4.9.0                          petsc/3.2_intel_impi
gcclib/5.2.0                          petsc/3.2_intel_impi_cxx
gdal/1.9.2                            petsc/3.2_intel_openmpi
gdb/7.3.1                             petsc/3.2_intel_openmpi_cxx
gdb/7.6(default)                      petsc/3.3_intel_openmpi
git/1.7.1                             petsc/3.4.4_intel_openmpi
git/1.7.10(default)                   petsc/3.4.4_intel_openmpi_cx
git/1.9.5                             pfft/1.0.7-alpha
git-annex/4.20130827                  pgi/12.5
git-annex/5.20150219                  pgi/12.6(default)
glib/2.44.1                           pgi/13.2
globus/5.2.5(default)                 pgplot/5.2.2-gcc
gnu-parallel/20130422                 pgplot/5.2.2-intel(default)
gnu-parallel/20140622(default)        plink/1.07
gnu-parallel/20150822                 plumed/2.0.3
```

. . .

```
gnuplot/4.2.6                                  plumed/2.1.3
gnuplot/4.6.1(default)                         py27-h5py/2.5.0-intelmpi-inte
gotoblas/1.13-multithreaded                    python/2.7.2(default)
gotoblas/1.13-singlethreaded(default)          python/2.7.3
grace/5.1.22                                   python/2.7.5
graphics/graphics                              python/2.7.8
gromacs/4.5.5(default)                         python/3.3.4
gromacs/4.5.5-double                           qt/4
gromacs/4.5.7                                  quake/0.3.5
gromacs/4.6.2                                  ray/1.7-large
gromacs/4.6.3                                  ray/1.7-small
gromacs/4.6.7                                  ray/2.1.0-large
gsl/1.13-gcc                                   ray/2.1.0-small(default)
gsl/1.13-intel(default)                        ray/2.1.1-devel-large
gsl/1.15-gcc                                   ray/2.1.1-devel-small
gsl/1.15-intel                                 ray/2.2.0-large
gsl/1.16-gcc                                   ray/2.2.0-small
gsl/1.16-intel                                 ray/2.3.1-large
guile/1.8.8(default)                           ray/2.3.1-small
```

. . .

```
guile/2.0.9                                rsync/3.0.6
harminv/1.3.1                              rsync/3.1.0(default)
haskell/6.12.3                             ruby/1.9.1(default)
haskell/7.4.2(default)                     ruby/1.9.3
hdf4/4.2.6-gcc(default)                    samtools/0.1.19
hdf4/4.2.6-intel                           scalapack/2.0.1-intel-intelmp
hdf5/1811-v18-intelmpi-intel               scalapack/2.0.1-intel-intelmp
hdf5/1811-v18-openmpi-intel                scalapack/2.0.1-openblas-open
hdf5/1811-v18-serial-gcc                   scalapack/2.0.1-openblas-open
hdf5/1811-v18-serial-intel                 scalasca/1.3.3-gcc-openmpi
hdf5/1813-v18-openmpi-intel                scalasca/1.3.3-intel-openmpi
hdf5/1814-v18-intelmpi-intel               scalasca/1.4.3-gcc-openmpi
hdf5/1814-v18-openmpi-gcc                  scalasca/1.4.3-intel-openmpi
hdf5/1814-v18-openmpi-intel                scons/2.0.1
hdf5/1814-v18-serial-gcc                   scotch/5.1.12
hdf5/1814-v18-serial-intel                 siesta/3.1
hdf5/187-v16-intelmpi-intel                siesta/3.2-pl-5
hdf5/187-v16-openmpi-gcc                   silo/4.8-bsd
hdf5/187-v16-openmpi-intel                 spark/1.0.2
```

. . .

```
hdf5/187-v16-serial-gcc                 spark/1.5.2
hdf5/187-v16-serial-intel               sparsehash/2.0.2
hdf5/187-v18-intelmpi-intel             splash/2.6.0
hdf5/187-v18-openmpi-gcc                stacks/1.28
hdf5/187-v18-openmpi-intel              stacks/1.29
hdf5/187-v18-serial-gcc(default)        stacks/1.30
hdf5/187-v18-serial-intel               stacks/1.35
hpnssh/5.8p1-hpn13v11                   suitesparse/4.2.1
intel/12.1                              trilinos/11.4.2
intel/12.1.2                            trilinos/11.4.2.petsc
intel/12.1.3(default)                   trilinosml/6.2
intel/12.1.5                            udunits/2.1.11
intel/13.1.1                            udunits/2.1.24
intel/14.0.0                            udunits/2.1.24-intel
intel/14.0.1                            upc/berkeley-2.12.2
intel/14.0.2                            valgrind/3.7.0(default)
intel/15.0.1                            valgrind/3.9.0_intelmpi
intel/15.0.2                            valgrind/3.9.0_openmpi
intellib/12.1                           vim/7.4.5
```

. . .

```
intellib/12.1.2                    visit/2.10.0-bin
intellib/12.1.3(default)           visit/2.10.0-bin_mesa
intellib/12.1.5                    visit/2.6.3
intellib/13.1.1                    visit/2.6.3-parallel
intellib/14.0.0                    vmd/1.8.6
intellib/14.0.1                    vmd/1.9(default)
intellib/14.0.2                    vnc/Xfvb+x11vnc
intellib/15.0.1                    wcslib/5.5
intellib/15.0.2                    xemacs/21.4.22
intelmpi/4.0.2.003                 xml2/2.7.8(default)
intelmpi/4.0.3.008                 yt/2.2
intelmpi/4.1.0.027                 zkcm/0.4.1
intelmpi/4.1.1.036
```

SciNet compute canada | calcul canada

Issue 1 - SciNet patch 1

```
$ module find

Top level modules:

abinit       dcap             harminv      nano       ROOT
abyss        ddd              haskell      ncl        rsync
adios        ddt              hdf4         nco        ruby
allpathslg   discovar         hdf5         ncview     samtools
amber        EIGENSOFT        hpnssh       ndiff      scalapack
antlr        emacs            ImageMagick  nedit      scalasca
armadillo    emi2workernode   intel        netcdf     scons
arpack-ng    encfs            intellib     nlopt      scotch
autoconf     erlang           intelmpi     nwchem     siesta
automake     espresso         inteltools   octave     silo
bbcp         extras           ipm          octopus    spark
bedtools     ferret           jasper       openbabel  sparsehash
BGW-paratec  ffmpeg           java         openblas   splash
binutils     fftw             libconfuse   openfoam   stacks
blast        gamess           libint       openmpi    suitesparse
```

...

blat	gcc	libxc	openspeedshop	trilinos
bowtie2	gcclib	luajit	padb	trilinosml
cactus	gdal	luarocks	papi	udunits
caf	gdb	make	parallel-netcdf	upc
cairo	GEOS	mc	paraview	valgrind
casacore	git	MCR	petsc	vim
cdo	git-annex	meep	pfft	visit
centos5-compat	glib	mesa	pgi	vmd
centos6-compat	globus	metis	pgplot	vnc
cfitsio	gnu-parallel	Minimac3	plink	wcslib
cmake	gnuplot	mono	plumed	xemacs
coreutils	gotoblas	mothur	py27-h5py	Xlibraries
cp2k	grace	mpb	python	xml2
CPLEX	graphics	mpe	qt	yt
cpmd	gromacs	mpich1	quake	zkcm
cuda	gsl	mysql	R	
cxxlibraries	guile	namd	ray	

compute canada | calcul canada

Issue 2

- Lot of versions depend on each other or conflict.
- Switching versions has become hard

E.g

```
$ module load gsl
gsl/1.13-intel(11):ERROR:151: Module 'gsl/1.13-intel' depends on
one of the module(s) 'intel/15.0.2 intel/15.0.1 intel/14.0.2
intel/14.0.1 intel/14.0.0 intel/13.1.1 intel/12.1.5 intel/12.1.3 in
gsl/1.13-intel(11):ERROR:102: Tcl command execution failed:
prereq intel
```

SciNet compute | calcul canada

Issue 2 - SciNet patches

```
$ module find gsl/

 Looking for 'gsl'... 6 hits

+  gsl/1.15-gcc   adds the gsl 1.15 library, compiled with the GNU c

+  gsl/1.13-gcc   adds the gsl 1.13 library, compiled with the GNU c
?  gsl/1.16-intel adds the gsl 1.15 library, compiled with the intel
?  gsl/1.16-gcc   adds the gsl 1.16 library, compiled with the GNU c
?  gsl/1.15-intel adds the gsl 1.15 library, compiled with the intel
?  gsl/1.13-intel adds the gsl 1.13 library, compiled with the Intel

(L = loaded,  + = loadable,  - = not loadable, ? = unresolved depenc

$ module advise gsl/1.16

The following set of commands would load the requested module 'gsl/1
module load intel/15.0.2
module load gsl/1.16-intel
```

LMOD

LMOD

- LMOD is another implementation of the module system
- It uses `lua` as the language for modules, but understands a good bit of `tcl`.
- LMOD intends to solve the "module hierarchy problem".
  (The existence of this problem is debatable)
  It essentially does so by hiding modules that you cannot load.
- `lmod` purports to be a drop-in replacement for `modulecmd`
  (Warning: It isn't; and definitely not when using a hierarchy.)

SciNet compute | calcul canada

LMOD improvements and changes

- Initially, `module avail` only shows loadable modules
- When you load, say, a compiler module, it increases that list of modules by adding all modules that depend on it.
- When you then load, say, an mpi modules, it increases it further with modules that use that compiler and that mpi library
- This is what's meant by "module hierarchy".
  It's implementated by having a module directory for each combination of compiler and mpi-library.
- This leads to saner module names. No more

```
$ module load intel/13.1.1 openmpi/intel/1.6.4 petsc/3.4.4_intel_openmpi
```

just do:

```
$ module load intel/13.1.1 openmpi/1.6.4 petsc/3.4.4
```

SciNet compute canada | calcul canada

LMOD improvements and changes

- Want to find a particular modules that is not listed but may exist: use the new commmand `module spider`.
- `spider` should be faster than the old `avail`, `find` or `advice` because it uses a cache.
- Loading another version of petsc? Old one get unload automatically.
- Loading another compiler? Other modules are replaced by their equivalents.
- Lua modules have a few nice features over tcl modules, such as allowing for a range of versions.
- However, lua modules do not support 'or', i.e., you cannot have a modules that will work with either gcc/4.6.1 or gcc/4.9.0. The only solution is to create separate module trees for gcc/4.6.1 and 4.9.0.

SciNet compute | calcul canada

LMOD setup on the GPC

SciNet compute | calcul canada

General setup

You can use lmod on the GPC now!

- Some modules have been renamed.
- Most have gotten a place in the module hierarchy
- Only some modules have been incorporated yet in the hierarchy; let us know which ones are urgent for you.

SciNet compute canada | calcul canada

LMOD on GPC development nodes

Activating lmod

- Create a file .lmod in your home directory on scinet:

```
$ touch $HOME/.lmod
```

- Takes affect upon login, so

```
$ ssh gpc
```

- To undo, remove .lmod

```
$ rm $HOME/.lmod
```

- Does not affect queued jobs!
- Interactive jobs? Next slide!

Using LMOD in queued jobs

Using LMOD in already queued but not yet running jobs

Place the `JOBID` of a non-running job to use LMOD in the hidden file `.jobids` in your home directory. This takes effect upon job start.

```
$ echo JOBID >> $HOME/.jobids
```

Using lmod at submission time

Replace `qsub` with `lmodqsub`

```
$ lmodqsub submissionscript.pbs
```

Could still change your mind my editing `$HOME/.jobids`.

Using lmod for an interactive debugjob

Use `lmoddebugjob` instead of `debugjob`

```
$ lmoddebugjob
```

canada | canada

Demonstration

Demonstration

- Let's login and see.

SciNet compute canada | calcul canada