
Lmod Documentation

Release 7.0

Robert McLay

Aug 17, 2018

CONTENTS

1	PURPOSE	1
2	OVERVIEW	3
3	Introduction to Lmod	5
4	Installing Lmod	23
5	Advanced Topics	45
6	Topics yet to be written	103
7	Indices and tables	105

PURPOSE

Lmod is a Lua based module system that easily handles the MODULEPATH Hierarchical problem. Environment Modules provide a convenient way to dynamically change the users' environment through modulefiles. This includes easily adding or removing directories to the PATH environment variable. Modulefiles for Library packages provide environment variables that specify where the library and header files can be found.

OVERVIEW

This guide is written to explain what Environment Modules are and why they are very useful for both users and system administrators. Lmod is an implementation of Environment Modules, much of what is said here is true for any environment modules system but there are many features which are unique to Lmod.

Environment Modules provide a convenient way to dynamically change the users' environment through modulefiles. This includes easily adding or removing directories to the `PATH` environment variable.

A modulefile contains the necessary information to allow a user to run a particular application or provide access to a particular library. All of this can be done dynamically without logging out and back in. Modulefiles for applications modify the user's path to make access easy. Modulefiles for Library packages provide environment variables that specify where the library and header files can be found.

Packages can be loaded and unloaded cleanly through the module system. All the popular shells are supported: `bash`, `ksh`, `csh`, `tcsh`, `zsh`. Also available for `perl` and `python`.

It is also very easy to switch between different versions of a package or remove it.

INTRODUCTION TO LMOD

If you are new to Lmod then please read the User Guide and possibly the Frequently Asked Questions Guide. Users who wish to read about how to have their own personal modules should read the Advanced User Guide.

3.1 User Guide for Lmod

The guide here explains how to use modules. The User's tour of the module command covers the basic uses of modules. The other sections can be read at a later date as issues come up. The Advance user's guide is for users needing to create their own modulefiles.

3.1.1 User's Tour of the Module Command

The module command sets the appropriate environment variable independent of the user's shell. Typically the system will load a default set of modules. A user can list the modules loaded by:

```
$ module list
```

To find out what modules are available to be loaded a user can do:

```
$ module avail
```

To load packages a user simply does:

```
$ module load package1 package2 ...
```

To unload packages a user does:

```
$ module unload package1 package2 ...
```

A user might wish to change from one compiler to another:

```
$ module swap gcc intel
```

The above command is short for:

```
$ module unload gcc  
$ module load intel
```

A user may wish to go back to an initial set of modules:

```
$ module reset
```

This will purge all currently loaded modules then load the list of modules specified by `LMOD_SYSTEM_DEFAULT_MODULES`. There is a related command:

```
$ module restore
```

This command will also purge all currently loaded modules and then load the system default unless the user has a default collection. See [User Collections](#) for more details.

If a module is not available then an error message is produced:

```
$ module load packageXYZ
Warning: Failed to load: packageXYZ
```

It is possible to try to load a module with no error message if it does not exist:

```
$ module try-load packageXYZ
```

Modulefiles can contain help messages. To access a modulefile's help do:

```
$ module help packageName
```

To get a list of all the commands that module knows about do:

```
$ module help
```

The module avail command has search capabilities:

```
$ module avail cc
```

will list for any modulefile where the name contains the string “cc”.

Users may wish to test whether certain modules are already loaded:

```
$ module is-loaded packageName1 packageName2 ...
```

Lmod will return true if all modules are loaded and false if one is not.

Users also may wish to test whether certain modules can be loaded with the current `$MODULEPATH`:

```
$ module is-avail packageName1 packageName2 ...
```

Lmod will return true if all modules are available and false if one can not be loaded.

Modulefiles can have a description section known as “whatis”. It is accessed by:

```
$ module whatis pmetis
pmetis/3.1   : Name: ParMETIS
pmetis/3.1   : Version: 3.1
pmetis/3.1   : Category: library, mathematics
pmetis/3.1   : Description: Parallel graph partitioning..
```

Finally, there is a keyword search tool:

```
$ module keyword word1 word2 ...
```

This will search any help message or whatis description for the word(s) given on the command line.

Another way to search for modules is with the “module spider” command. This command searches the entire list of possible modules. The difference between “module avail” and “module spider” is explained in the “Module Hierarchy” and “Searching for Modules” section.:

```
$ module spider
```

ml: A convenient tool

For those of you who can’t type the *mdoule*, *moduel*, err *module* command correctly, Lmod has a tool for you. With **ml** you won’t have to type the module command again. The two most common commands are *module list* and *module load <something>* and **ml** does both:

```
$ ml
```

means *module list*. And:

```
$ ml foo
```

means *module load foo* while:

```
$ ml -bar
```

means *module unload bar*. It won’t come as a surprise that you can combine them:

```
$ ml foo -bar
```

means *module unload bar; module load foo*. You can do all the module commands:

```
$ ml spider
$ ml avail
$ ml show foo
```

If you ever have to load a module name *spider* you can do:

```
$ ml load spider
```

If you are ever forced to type the **module** command instead of **ml** then that is a bug and should be reported.

SAFETY FEATURES

(1): Users can only have one version active: The One Name Rule

If a user does:

```
$ module avail xyz
----- /opt/apps/modulefiles -----
xyz/8.1   xyz/11.1 (D)   xyz/12.1

$ module load xyz
$ module load xyz/12.0
```

The first load command will load the 11.1 version of xyz. In the second load, the module command knows that the user already has xyz/11.1 loaded so it unloads that and then loads xyz/12.0. This protection is only available with Lmod.

This is known as the *One Name* rule. This feature is core to how Lmod works and there is no way to override this.

(2) : Users can only load one compiler or MPI stack at a time.

Lmod provides an additional level of protection. If each of the compiler modulefiles add a line:

```
family("compiler")
```

Then Lmod will not load another compiler modulefile. Another benefit of the modulefile family directive is that an environment variable “LMOD_FAMILY_COMPILER” is assigned the name (and not the version). This can be useful specifying different options for different compilers. In the High Performance Computing (HPC) world, the message passing interface (MPI) libraries are important. The mpi modulefiles can contain a family(“MPI”) directive which will prevent users from loading more than one MPI implementation at a time. Also the environment variable “LMOD_FAMILY_MPI” is defined to the name of the mpi library.

Module Hierarchy

Libraries built with one compiler need to be linked with applications with the same compiler version. If sites are going to provide libraries, then there will be more than one version of the library, one for each compiler version. Therefore, whether it is the Boost library or an mpi library, there are multiple versions.

There are two main choices for system administrators. For the XYZ library compiled with either the UCC compiler or the GCC compiler, there could be the xyz-ucc modulefile and the xyz-gcc module file. This gets much more complicated when there are multiple versions of the XYZ library and different compilers. How does one label the various versions of the library and the compiler? Even if one makes sense of the version labeling, when a user changes compilers, the user will have to remember to unload the ucc and the xyz-ucc modulefiles when changing to gcc and xyz-gcc. If users have mismatched modules, their programs are going to fail in very mysterious ways.

A much saner strategy is to use a module hierarchy. Each compiler module adds to the MODULEPATH a compiler version modulefile directory. Only modulefiles that exist in that directory are packages that have been built with that compiler. When a user loads a particular compiler, that user only sees modulefile(s) that are valid for that compiler.

Similarly, applications that use libraries depending on MPI implementations must be built with the same compiler - MPI pairing. This leads to modulefile hierarchy. Therefore, as users start with the minimum set of loaded modules, all they will see are compilers, not any of the packages that depend on a compiler. Once they load a compiler they will see the modules that depend on that compiler. After choosing an MPI implementation, the modules that depend on that compiler-MPI pairing will be available. One of the nice features of Lmod is that it handles the hierarchy easily. If a user swaps compilers, then Lmod automatically unloads any modules that depends on the old compiler and reloads those modules that are dependent on the new compiler.

```
$ module list

1) gcc/4.4.5 2) boost/1.45.0

$ module swap gcc ucc

Due to MODULEPATH changes the follow modules have been reloaded: 1) boost
```

If a modulefile is not available with the new compiler, then the module is marked as inactive. Every time MODULEPATH changes, Lmod attempts to reload any inactive modules.

Searching For Modules

When a user enters:

```
$ module avail
```

Lmod reports only the modules that are in the current MODULEPATH. Those are the only modules that the user can load. If there is a modulefile hierarchy, then a package the user wants may be available but not with the current compiler version. Lmod offers a new command:

```
$ module spider
```

which lists all possible modules and not just the modules that can be seen in the current MODULEPATH. This command has three modes. The first mode is:

```
$ module spider

lmod: lmod/lmod
Lmod: An Environment Module System

ucc: ucc/11.1, ucc/12.0, ...
Ucc: the ultimate compiler collection

xyz: xyz/0.19, xyz/0.20, xyz/0.31
xyz: Solves any x or y or z problem.
```

This is a compact listing of all the possible modules on the system. The second mode describes a particular module:

```
$ module spider ucc
-----
ucc:
-----

Description:
Ucc: the ultimate compiler collection

Versions:
ucc/11.1
ucc/12.0
```

The third mode reports on a particular module version and where it can be found:

```
$ module spider parmetis/3.1.1
-----
parmetis: parmetis/3.1.1
-----

Description:
Parallel graph partitioning and fill-reduction matrix ordering routines

This module can be loaded through the following modules:
ucc/12.0, openmpi/1.4.3
ucc/11.1, openmpi/1.4.3
gcc/4.4.5, openmpi/1.4.3

Help:
The parmetis module defines the following environment variables: ...
The module parmetis/3.1.1 has been compiled by three different versions of the ucc_
  ↪ compiler and one MPI implementation.
```

Controlling Modules During Login

Normally when a user logs in, there are a standard set of modules that are automatically loaded. Users can override and add to this standard set in two ways. The first is adding module commands to their personal startup files. The second way is through the “module save” command.

To add module commands to users’ startup scripts requires a few steps. Bash users can put the module commands in either their `~/.profile` file or their `~/.bashrc` file. It is simplest to place the following in their `~/.profile` file:

```
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
fi
```

and place the following in their `~/.bashrc` file:

```
if [ -z "$BASHRC_READ" ]; then
    export BASHRC_READ=1
    # Place any module commands here
    # module load git
fi
```

By wrapping the module command in an if test, the module commands need only be read in once. Any sub-shell will inherit the `PATH` and other environment variables automatically. On login shells the `~/.profile` file is read which, in the above setup, causes the `~/.bashrc` file to be read. On interactive non-login shells, the `~/.bashrc` file is read instead. Obviously, having this setup means that module commands need only be added in one file and not two.

Csh users need only specify the module commands in their `~/.cshrc` file as that file is always sourced:

```
if ( ! $?CSHRC_READ ) then
    setenv CSHRC_READ 1
    # Place any module command here
    # module load git
endif
```

User Collections

User defined initial list of login modules:

Assuming that the system administrators have installed Lmod correctly, there is a second way which is much easier to setup. A user logs in with the standard modules loaded. Then the user modifies the default setup through the standard module commands:

```
$ module unload XYZ
$ module swap gcc ucc
$ module load git
```

Once users have the desired modules load then they issue:

```
$ module save
```

This creates a file called `~/.lmod.d/default` which has the list of desired modules. Once this is set-up a user can issue:

```
$ module restore
```

and only the desired modules will be loaded. If Lmod is setup correctly (see *Providing A Standard Set Of Modules for all Users*) then the default collection will be the user’s initial set of modules.

If a user doesn’t have a default collection, the Lmod purges all currently loaded modules and loads the list of module specified by LMOD_SYSTEM_DEFAULT_MODULES just like the `module reset` command.

Users can have as many collections as they like. They can save to a named collection with:

```
$ module save <collection_name>
```

and restore that named collection with:

```
$ module restore <collection_name>
```

A user can print the contents of a collection with:

```
$ module describe <collection_name>
```

A user can list the collections they have with:

```
$ module savelist
```

Finally a user can disable a collection with:

```
$ module disable <collection_name>
```

If no `collection_name` is given then the default is disabled. Note that the collection is not remove just renamed. If a user disables the `foo` collection, the file `foo` is renamed to `foo~`. To restore the `foo` collection, a user will have to do the following:

```
$ cd ~/.lmod.d; mv foo~ foo
```

User Collections on shared home file systems

If your site has a shared home file system, then things become a little more complicated. A shared home file system means that your site has a single home file system shared between two or more clusters. See *Lmod on Shared Home File Systems* for a system administrators point of view.

If you have a collection on one cluster it needs to be independent of another cluster. Your site should set `$LMOD_SYSTEM_NAME` uniquely for each cluster. Suppose you have cluster A and B. Then `$LMOD_SYSTEM_NAME` will be either A or B. A default collection will be named “default.A” for the A cluster and “default.B” for the B cluster. The names a user sees will have the extension removed. In other words on the A cluster a user would see:

```
$ module savelist

1) default
```

where the default file is named “default.A”.

3.2 An Introduction to Writing Modulefiles

This is a different kind of introduction to Lmod. Here we will remind you what Lmod is doing to change the environment via modulefiles. Then we will start with the four functions that are typically needed for any modulefile.

From there we will talk about intermediate level module functions when things get more complicated. Finally we will discuss the advanced module functions to flexibly control your site via modules. All the Lua module functions available are described at [Lua Modulefile Functions](#). This discussion shows how they can be used.

3.2.1 A Reminder of what Lmod is doing

All Lmod is doing is changing the environment. Suppose you want to use the “ddt” debugger installed on your system which is made available to you via the module. If you try to execute ddt without the module loaded you get:

```
$ ddt
bash: command not found: ddt

$ module load ddt
$ ddt
```

After the ddt module is loaded, executing **ddt** now works. Let’s remind ourselves why this works. If you try checking the environment before loading the ddt modulefile:

```
$ env | grep -i ddt
$ module load ddt
$ env | grep -i ddt

DDTPATH=/opt/apps/ddt/5.0.1/bin
LD_LIBRARY_PATH=/opt/apps/ddt/5.0.1/lib:...
PATH=/opt/apps/ddt/5.0.1/bin:...

$ module unload ddt
$ env | grep -i ddt
$
```

The first time we check the environment we find that there is no **ddt** stored there. But the second time there we see that the **PATH** and **LD_LIBRARY_PATH** have been modified. Note that we have shortened the path-like variables to show the important changes. There are also several environment variables which have been set. After unloading the module all the references for ddt have been removed. We can see what the modulefile looks like by doing:

```
$ module show ddt

help([[
For detailed instructions, go to:
  https://...

]])
whatis("Version: 5.0.1")
whatis("Keywords: System, Utility")
whatis("URL: http://content.allinea.com/downloads/userguide.pdf")
whatis("Description: Parallel, graphical, symbolic debugger")

setenv(      "DDTPATH",      "/opt/apps/ddt/5.0.1/bin")
prepend_path( "PATH",      "/opt/apps/ddt/5.0.1/bin")
prepend_path( "LD_LIBRARY_PATH", "/opt/apps/ddt/5.0.1/lib")
```

Modulefiles state the actions that need to happen when loading. For example the above modulefile uses **setenv** and **prepend_path** to set environment variables and prepend to the **PATH**. If the above modulefile is unloaded then the **setenv** actually unsets the environment variable. The **prepend_path** removes the element from the **PATH** variable. That is unload causes the functions to be reversed.

3.2.2 Basic Modulefiles

There are two main module functions required, namely **setenv** and **prepend_path**; and two functions to provide documentation **help** and **whatis**. The modulefile for ddt shown above contains all the basics required to create one. Suppose you are writing this module file for ddt version 5.0.1 and you are placing it in the standard location for your site, namely */apps/modulefiles* and this directory is already in **MODULEPATH**. Then in the directory */apps/modulefiles/ddt* you create a file called *5.0.1.lua* which contains the modulefile shown above.

This is the typical way of setting a modulefile up. Namely the package name is the name of the directory, *ddt*, and version name, *5.0.1* is the name of the file with the *.lua* extension added. We add the lua extension to all modulefile written in Lua. All modulefiles without the lua extension are assumed to be written in TCL.

If another version of ddt becomes available, say *5.1.2*, we create another file called *5.1.2.lua* to become the new modulefile for the new version of *ddt*.

When a user does *module help ddt*, the arguments to the **help** function are written out to the user. The **whatis** function provides a way to describe the function of the application or library. This data can be used by search tools such as **module keyword search_words**. Here at TACC we also use that data to provide search capability via the web interface to modules we provide.

3.2.3 Intermediate Level Modulefiles

The four basic functions describe above is all that is necessary for the majority of modulefiles for application and libraries. The intermediate level is designed to describe some situations that come up as you need to provide more than just packages modulefile but need to set up a system.

Meta Modules

Some sites create a single module to load a default set of modules for all users to start from. This is typically called a meta module because it loads other modules. As an example of that, we here at TACC have created the TACC module to provide a default compiler, mpi stack and other modules:

```
help([[
The TACC modulefile defines ...
]])

-- 1 --
if (os.getenv("USER") ~= "root") then
    append_path("PATH", ".")
end

-- 2 --
load("intel", "mvapich2")

-- 3 --
try_load("xalt")

-- 4 --
-- Environment change - assume single threaded.
if (mode() == "load" and os.getenv("OMP_NUM_THREADS") == nil) then
    setenv("OMP_NUM_THREADS", "1")
end
```

This modulefile shows the use of four new functions. The first one is **append_path**. This function is similar to **prepend_path** except that the value is placed at the end of the path-like variable instead of the beginning. We add *.* to our user's path at the end, except for root. This way our new users don't get surprised with some programs in

their current directory that do not run. We used the **os.getenv** function built-in to Lua to get the value of environment variable “USER”.

The second function is **load**, this function loads the modulefiles specified. This function takes one or more names. Here we are specifying a default compiler and mpi stack. The third function is **try_load**, which is similar to **load** except that there is no error reported if the module can’t be found.

The fourth block of code shows how we set **OMP_NUM_THREADS**. We wish to set **OMP_NUM_THREADS** to have a default value of 1, but only if the value hasn’t already been set and we only want to do this when the module is being loaded and not at any other time. So when this module is loaded for the first time **mode()** will return “load” and **OMP_NUM_THREADS** won’t have a value. The **setenv** will set it to 1. If the TACC module is unloaded, the **mode()** will be “unload” so the if test will be false and therefore the **setenv** will not be reversed. If the user changes **OMP_NUM_THREADS** and reloads the TACC modulefile, their value won’t change because **os.getenv(“OMP_NUM_THREADS”)** will return a non-nil value, therefore the **setenv** command won’t run. Now this may not be the best way to handle this. It might be better to set **OMP_NUM_THREADS** in a file that is sourced in */etc/profile.d/* and have all the important properties. Namely that there will be a default value that the user can change. However this example shows how to do something tricky in a modulefile.

Typically meta modules are a single file and not versioned. So the TACC modulefile can be found at */apps/modulefiles/TACC.lua*. There is no requirement that this will be this way but it has worked well in practice.

Modules with dependencies

Suppose that you have a package which needs libraries or an application. For example the octave application needs gnuplot. Let’s assume that you have a separate applications for both. Inside the octave module you can do:

```
prereq("gnuplot")
...
```

So if you execute:

```
$ module unload gnuplot
$ module load octave
$ module load gnuplot octave
$ module unload octave
```

The second module command will fail, but the third one will succeed because we have met the prerequisites. The advantage of using **prereq** is after fourth module command is executed, the gnuplot module will be loaded.

This can be contrasted with including the load of gnuplot in the octave modulefile:

```
load("gnuplot")
...
```

This simplifies the loading of the octave module. The trouble is that when a user does the following:

```
$ module load    gnuplot
$ module load    octave
$ module unload  octave
```

is that after unloading *octave*, the *gnuplot* module is also unloaded. It seems better to either use the **prereq** function shown above or use the **always_load** function in the octave module:

```
always_load("gnuplot")
...
```

Then when a user does:

```
$ module load  gnuplot
$ module load  octave
$ module unload octave
```

The *gnuplot* module will still be loaded after unloading *octave*. This will lead to the least confusion to users.

Fancy dependencies

Sometimes an application may depend on another application but it has to be a certain version or newer. Lmod can support this with the **atleast** modifier to both **load**, **always_load** or **prereq**. For example:

```
-- Use either the always_load or prereq but not both:

prereq(      atleast("gnuplot", "5.0"))
always_load(atleast("gnuplot", "5.0"))
```

The **atleast** modifier to **prereq** or **always_load** will succeed if the version of *gnuplot* is 5.0 or greater.

Assigning Properties

Modules can have properties that will be displayed in a *module list* or *module avail*. Properties can be anything but they must be specified in the *lmodrc.lua* file. You are free to add to the list. For example, to specify a module to be experimental all you need to do is:

```
add_property("state", "experimental")
```

Any properties you set must be defined in the **lmodrc.lua** file. In the source tree the properties are in *init/lmodrc.lua*. A more detailed discussion of the *lmodrc.lua* file can be found at [Module Properties](#)

Pushenv

Lmod allows you to save the state in a stack hidden in the environment. So if you want to set the **CC** environment variable to contain the name of the compiler.:

```
-- gcc --
pushenv("CC", "gcc")

-- mpich --
pushenv("CC", "mpicc")
```

If the user executes the following:

	SETENV	PUSHENV
#		
\$ export CC=cc;	echo \$CC # -> CC=cc	CC=cc
\$ module load gcc;	echo \$CC # -> CC=gcc	CC=gcc
\$ module load mpich;	echo \$CC # -> CC=mpicc	CC=mpicc
\$ module unload mpich;	echo \$CC # -> CC is unset	CC=gcc
\$ module unload gcc;	echo \$CC # -> CC is unset	CC=cc

We see that the value of **CC** is maintained as a stack variable when we use *pushenv* but not when we use *setenv*.

Setting aliases and shell functions

Sometimes you want to set an alias as part of a module. For example the visit program requires the version to be specified when running it. So for version 2.9 of visit, the alias is set:

```
set_alias("visit", "visit -v 2.9")
```

Whether this will expand correctly depends on the shell. While C-shell allows argument expansion in aliases, Bash and Zsh do not. Bash and Zsh use shell functions instead. For example the ml shell function can be set like this:

```
local bashStr = 'eval `$(LMOD_DIR/ml_cmd "$@"`'  
local cshStr = "eval `$(LMOD_DIR/ml_cmd $*`"  
set_shell_function("ml", bashStr, cshStr)
```

Please note that aliases in bash are not expanded for non-interactive shells. This means that it won't work in bash shell scripts. Please change the shell alias to use the `set_shell_function` instead. Shell functions do work in both interactive and non-interactive shells.

3.3 Frequently Asked Questions

How does the module command work?

We know that the child program inherits the parents' environment but not the other way around. So it is very surprising that a command can change the current shell's environment. The trick here is that the module command is a two part process. The module shell function in bash is:

```
$ type module  
module() { eval `$(LMOD_CMD bash "$@")` }
```

Where `$LMOD_CMD` points to your lmod command (say `/apps/lmod/lmod/libexec/lmod`). So if you have a module file (foo/1.0) that contains:

```
setenv("FOO", "BAR")
```

then "`$LMOD_CMD bash load foo/1.0`" produces:

```
export FOO=BAR  
...
```

The eval command read that output from stdout and changes the current shell's environment. Any text written to stderr bypasses the eval and written to the terminal.

What are the environment variables `_ModuleTable001_`, `_ModuleTable002_`, etc doing it in the environment?

The module command remembers its state in the environment through these variables. The way Lmod does it is through a Lua table called `ModuleTable`:

```
ModuleTable = {  
  mT = {  
    git = { ... }  
  }  
}
```

This table contains quotes and commas and must be store in environment. To prevent problems the various shells, the table is encoded into base64 and split into blocks of 256 characters. These variable are decoded at the start of Lmod. You can see what the module table contains with:

```
$ module --mt
```

How does one debug a modulefile?

There are two methods. Method 1: If you are writing a Lua modulefile then you can write messages to stderr with and run the module command normally:

```
local a = "text"
io.stderr:write("Message ",a,"\n")
```

Method 2: Take the output directly from Lmod. You can put print() statements in your modulefile and do:

```
$ $LMOD_CMD bash load *modulefile*
```

Why doesn't `%module avail |& grep ucc` work under tcsh and works under bash?

It is a bug in the way tcsh handles evals. This works:

```
% (module avail) |& grep ucc
```

However, in all shells it is better to use:

```
% module avail ucc
```

instead as this will only output modules that have “ucc” in their name.

Can I disable the pager output?

Yes, you can. Just set the environment variable `LMOD_PAGER` to **none**.

Why are messages printed to standard error and not standard out?

The module command is an alias under tcsh and a shell routine under all other shells. There is an lmod command which writes out commands such as `export FOO="bar and baz"` and messages are written to standard error. The text written to standard out is evaluated so that the text strings make changes to the current environment. See next question for a different way to handle Lmod messages.

Can I force the output of list, avail and spider to go to stdout instead of stderr?

Bash and Zsh user can set the environment variable `LMOD_REDIRECT` to **yes**. Sites can configure Lmod to work this way by default. However, no matter how Lmod is set-up, this will not work with tcsh/csh due to limitations of this shell.

How can I use grep easily with the module command?

If your site doesn't send the output of stdout, you can still use this trick when you need to grep the output of module command. Here are some examples:

```
$ module -t --redirect avail | grep foo
$ module --raw --redirect show foo | grep bar
$ module -t --redirect spider | grep baz
```

Can I ignore the spider cache files when doing module avail?

Yes you can:

```
$ module --ignore_cache avail
```

or you can set:

```
$ export LMOD_IGNORE_CACHE=1
```

to make Lmod ignore caches as long as the variable is set.

I have created a module and “module avail” can’t find it. What do I do?

Assuming that the modulefile is in MODULEPATH then you have an out-of-date cache. Try running:

```
$ module --ignore_cache avail
```

If this does find it then you might have an old personal spider cache. To clear it do:

```
$ rm -rf ~/.lmod.d/.cache
```

If “module avail” doesn’t find it now, then the system spider cache is out-of-date. Please ask your system administrator to update the cache. If you are the system administrator then please read [System Spider Cache](#) and [User Spider Cache](#)

Why doesn’t the module command work in shell scripts?

It will if the following steps are taken. First the script must be a bash script and not a shell script, so start the script with `#!/bin/bash`. The second is that the environment variable `BASH_ENV` must point to a file which defines the module command. The simplest way is having `BASH_ENV` point to `/opt/apps/lmod/lmod/init/bash` or wherever this file is located on your system. This is done by the standard install. Finally Lmod exports the module command for Bash shell users.

How do I use the initializing shell script that comes with this application with Lmod?

The short answer is you don’t. Among the many problems is that there is no way to unload that shell script. If the script is simple you can read it through and create a modulefile. To simplify this task, Lmod provides the `sh_to_modulefile` script to convert shell scripts to modulefiles. See [Converting shell scripts to modulefiles](#) for more details.

Why is the output of `module avail` not filling the width of the terminal?

If the output of `module avail` is 80 characters wide, then Lmod can’t find the width of the terminal and instead uses the default size (80). If you do `module --config`, you’ll see a line:

```
Active lua-term true
```

If it says **false** instead then lua-term is not installed. One way this happens is to build Lmod on one computer system that has a system lua-term installed and the package on another where lua-term isn’t installed on the system.

Why isn’t the module defined when using the **screen** program?

The screen program starts a non-login interactive shell. The Bash shell startup doesn’t start sourcing `/etc/profile` and therefore the `/etc/profile.d/*.sh` scripts for non-login interactive shells. You can patch bash and fix `/etc/bashrc` (see [Issues with Bash](#) for a solution) or you can fix your `~/.bashrc` to source `/etc/profile.d/*.sh`

You may be better off using **tmux** instead. It starts a login shell.

Why does `LD_LIBRARY_PATH` get cleared when using the **screen** program?

The screen program is a *guid* program. That means it runs as the group of the program and not the group associated with the user. For security reason all of these kinds of program clear `LD_LIBRARY_PATH`. This unsetting of `LD_LIBRARY_PATH` is done by the Unix operating system and not Lmod.

You may be better off using **tmux** instead. It is a regular program.

How can you write TCL files that can be safely used with both Lmod and Tmod?

For example the `hide-version` command only works Lmod and could be found in `~/.modulerc`. This could be read by both Tmod and Lmod. You can prevent Tmod from executing Lmod only code in the following way:

```
##Module
global env
if { [info exists env(LMOD_VERSION_MAJOR)] } {
    hide-version CUDA/8.8.8
}
```

Lmod defines the environment variable `LMOD_VERSION_MAJOR` during its execution. This trick can also be used in a TCL modulefile to set the family function:

```
##Module
...
global env
if { [info exists env(LMOD_VERSION_MAJOR)] } {
    family compiler
}
```

How can I get the shell functions created by modules in bash shell scripts such as job submission scripts?

First, please make sure that shell functions and alias works correctly in bash interactive subshells. If they don't then your site is not setup correctly.

Once that works then change the first line of the shell script to be:

```
#!/bin/bash -l
```

Note that is a minus ell not minus one. This will cause the startup scripts to be sourced before the first executable statement in the script.

Why do modules get sometimes get loaded when I execute `module use <path>`?

A main principal is that when `$MODULEPATH` changes, Lmod checks all the currently loaded modules. If any of thoses modules would not have been chosen then each is swapped for the new choice.

3.4 Advanced User Guide for Personal Modulefiles

This advanced guide is for users wishing to create modulefiles for their own software. The reasons are simple:

1. Install a newer version of open source software than is currently available.
2. Easily change the version of applications or libraries under their own development.
3. Better documentation for what software is available.

One could create new version of some software and place it in your personal `PATH` and forget about it. At least when it is in a module, it will be listed in the loaded modules. It will also appear in the list of available software via `module avail`

3.4.1 User Created Modules

Users can create their own modules. The first step is to add to the module path:

```
$ module use /path/to/personal/modulefiles
```

This will prepend `/path/to/personal/modulefiles` to the `MODULEPATH` environment variable. This means that any modulefiles defined here will be used instead of the system modules.

Suppose that the user creates a directory called `$HOME/modulefiles` and they want a personal copy of the “git” package and they do the usual “tar, configure, make, make install” steps:

```
$ wget https://www.kernel.org/pub/software/scm/git/git-2.6.2.tar.gz
$ tar xf git-2.6.2.tar.gz
$ cd git-2.6.2
$ ./configure --prefix=$HOME/pkg/git/2.6.2
$ make
$ make install
```

This document has assumed that 2.6.2 is the current version of git, it will need to be replaced with the current version. To create a modulefile for git one does:

```
$ cd ~/modulefiles
$ mkdir git
$ cd git
$ cat > 2.6.2.lua
local home      = os.getenv("HOME")
local version   = myModuleVersion()
local pkgName   = myModuleName()
local pkg       = pathJoin(home, "pkg", pkgName, version, "bin")
prepend_path("PATH", pkg)
^D
```

Starting first from the name: `git/2.6.2.lua`, modulefiles with the `.lua` extension are assumed to be written in lua and files without are assumed to be written in TCL. This modulefile for git adds `~/pkg/git/2.6.2/bin` to the user’s path so that the personal version of git can be found. Note that the use of the functions **myModuleName()** and **myModuleVersion()** allows the module to be generic and not hard-wired to a particular module file. We have used the `cat` command to quickly create this lua modulefile. Obviously, this file can easily be created by your favorite editor.

The first line reads the user’s `HOME` directory from the environment. The second line uses the “pathJoin” function provided from Lmod. It joins strings together with the appropriate number of “/”. The last line calls the “prepend_path” function to add the path to git to the user’s path.

Finally the user can do:

```
$ module use $HOME/modulefiles
$ module load git
$ type git
~/pkg/git/2.6.2/bin/git
```

For git to be available on future logins, users need to add the following to their startup scripts or a saved collection:

```
$ module use $HOME/modulefiles
$ module load git
```

The modulefiles can be stored in different directories. There is an environment variable `MODULEPATH` which controls that. Modulefiles that are listed in an earlier directory are found before ones in later directories. This is similar to command searching in the `PATH` variable. There can be several versions of a command. The first one found in the `PATH` is used.

3.4.2 Finding Modules With Same Name

Suppose the user has created a “git” module to provide the latest available. At a later date, the system administrators add a newer version of “git”

```
$ module avail git
----- /home/user/modulefiles -----
git/2.6.2

----- /opt/apps/modulefiles -----
git/1.7.4  git/2.0.1  git/3.5.4 (D)

$ module load git
```

The load command will load `git/3.5.4` because it is the highest version.

If a user wishes to make their own version of git the default module, they will have to mark it as a default. Marking a module as a default is discussed in section [Marking a Version as Default](#).

INSTALLING LMOD

Anyone wishing to install Lmod on a personal computer or for a system should read the Installation Guide as well as the Transitioning to Lmod Guide. The rest of the guides can be read as needed.

4.1 Installing Lua and Lmod

Environment modules simplify customizing the users' shell environment and it can be done dynamically. Users load modules as they see fit. It is completely under their control. Environment Modules or simply modules provide a simple contract or interface between the system administrators and users. System administrators provide modules and users get to choose which to load.

There have been environment module systems for quite a while. See <http://modules.sourceforge.net/> for a TCL based module system and see <http://www.lysator.liu.se/cmod> for another module system. Here we describe Lmod, which is a completely new module system written in Lua. For those who have used modules before, Lmod automatically reads TCL modulefiles. Lmod has some important features over other module system, namely a built-in solution to hierarchical modulefiles and provides additional safety features to users as described in the User Guide.

The hierarchical modulefiles are used to solve the issue of system pre-built libraries. User applications using these libraries must be built with the same compiler as the libraries. If a site provides more than one compiler, then for each compiler version there will be separate versions of the libraries. Lmod provides built-in control making sure that compilers and pre-built libraries stay matched. The rest of the pages here describe how to install Lmod, how to provide the module command to users during the login process and some discussion on how to install optional software and the associated modules.

The goal of installing Lmod is when completed, any user will have the module command defined and a preset list of modules will be loaded. The module command should work without modifying the users startup files (`~/.bashrc`, `~/.profile`, `~/.cshrc`, or `~/.zshenv`). The module command should be available for login shells, interactive shells, and non-interactive shells. The command `ssh YOUR_HOST module list` should work. This will require some understanding of the system startup procedure for various shells which is covered here.

4.1.1 Installing Lua

In this document, it is assumed that all optional software is going to be installed in `/opt/apps`. The installation of Lmod requires installing lua as well. On some system, it is possible to install Lmod directly with your package manager. It is available with recent fedora, debian and ubuntu distributions.

Install lua-X.Y.Z.tar.gz

One choice is to install the lua-X.Y.Z.tar.gz file. This tar ball contains lua and the required libraries. This can be downloaded from <https://sourceforge.net/projects/lmod/files/>:

```
$ wget https://sourceforge.net/projects/lmod/files/luarocks-5.1.4.5.tar.gz
```

The current version is 5.1.4.5 but it may change in the future. This can be installed using the following commands:

```
$ tar xf lua-X.Y.Z.tar.gz
$ cd lua-X.Y.Z
$ ./configure --prefix=/opt/apps/lua/X.Y.Z
$ make; make install
$ cd /opt/apps/lua; ln -s X.Y.Z lua
$ mkdir /usr/local/bin; ln -s /opt/apps/lua/lua/bin/lua /usr/local/bin
```

The last command is optional, but you will have to somehow put the `lua` command in your path. Also obviously, please replace `X.Y.Z` with the actual version (say 5.1.4.5)

Using Your Package Manager

You can use your package manager for your OS to install Lua. You will also need the matching packages: `lua` Filesystem (`lfs`) and `luaposix`. On Ubuntu Linux, the following packages will work:

```
liblua5.1-0
liblua5.1-0-dev
liblua5.1-filesystem-dev
liblua5.1-filesystem0
liblua5.1-posix-dev
liblua5.1-posix0
lua5.1
```

Note; Centos may require looking the EPEL repo. At TACC we install the following rpms:

```
$ rpm -qa | grep lua

lua-posix-5.1.7-1.el6.x86_64
lua-5.1.4-4.1.el6.x86_64
lua-filesystem-1.4.2-1.el6.x86_64
lua-devel-5.1.4-4.1.el6.x86_64
```

You will also need the `libtcl` and `tcl` packages as well.

Using Luarocks

If you have installed `lua` but still need `luafilesystem` and `luaposix`, you can install the `luarocks` program from your package manager or directly from <https://luarocks.org/>. The `luarocks` programs can install many lua packages including the ones required for Lmod.

```
$ luarocks install luaposix; luarocks install luafilesystem
```

Now you have to make the lua packages installed by luarocks to be known by lua. On our Centos system, Lua knows about the following for `*.lua` files:

```
$ lua -e 'print(package.path)'
./?.lua;/usr/share/lua/5.1/?.lua;/usr/share/lua/5.1/?/init.lua;/usr/lib64/lua/5.1/?
lua;/usr/lib64/lua/5.1/?/init.lua;
```

and the following for shared libraries:

```
$ lua -e 'print(package.cpath)'
./?.so;/usr/lib64/lua/5.1/?.so;/usr/lib64/lua/5.1/loadall.so;
```

Assuming that luarocks has installed things in its default location (/usr/local/...) then you'll need to do:

```
LUAROCKS_PREFIX=/usr/local
export LUA_PATH="$LUAROCKS_PREFIX/share/lua/5.1/?.lua;$LUAROCKS_PREFIX/share/lua/5.1/?
↪/init.lua;;"
export LUA_CPATH="$LUAROCKS_PREFIX/lib/lua/5.1/?.so;;"
```

Please change LUAROCKS_PREFIX to match your site. The exporting of LUA_PATH and LUA_CPATH must be done before any module commands. It is very important that the trailing semicolon are there. They are replaced by the built-in system path.

4.1.2 Why does Lmod install differently?

Lmod automatically creates a version directory for itself. So, for example, if the installation prefix is set to /apps, and the current version is X.Y.Z, installation will create /apps/lmod and /apps/lmod/X.Y.Z. This way of configuring is different from most packages. There are two reasons for this:

1. Lmod is designed to have just one version of it running at one time. Users will not be switching version during the course of their interaction in a shell.
2. By making the symbolic link the startup scripts in /etc/profile.d do not have to change. They just refer to /apps/lmod/lmod/... and not /apps/lmod/X.Y.Z/...

4.1.3 Installing Lmod

Lmod has a large number of configuration options. They are discussed in the Configuring Lmod Guide. This section here will describe how to get Lmod installed quickly by using the defaults:

Note: If you have a large number of modulefiles or a slow parallel filesystem please read the Configure Lmod Guide on how to set-up the spider caching system. This will greatly speed up `module avail` and `module spider`

To install Lmod, you'll want to carefully read the following. If you want Lmod version X.Y installed in /opt/apps/lmod/X.Y, just do:

```
$ ./configure --prefix=/opt/apps
$ make install
```

The installation will also create a link to /apps/lmod/lmod. The symbolic link is created to ease upgrades to Lmod itself, as numbered versions can be installed side-by-side, testing can be done on the new version, and when all is ready, only the symbolic link needs changing.

To create such a testing installation, you can use:

```
$ make pre-install
```

which does everything but create the symbolic link.

In the `init` directory of the source code, there are `profile.in` and `cshrc.in` templates. During the installation phase, the path to lua is added and `profile` and `cshrc` are written to the `/apps/lmod/lmod/init`

directory. These files are created assuming that your modulefiles are going to be located in `/apps/modulefiles/$LMOD_sys` and `/apps/modulefiles/Core`, where `$LMOD_sys` is what the command “uname” reports, (e.g., Linux, Darwin). The layout of modulefiles is discussed later.

Note: Obviously you will want to modify the `profile.in` and `cshrc.in` files to suit your system.

Sites can also use `/apps/lmod/lmod/init/.modulespath` or configure Lmod to use `-with-ModulePathInit=...` to point to any file. The format of this file is:

```
# comments are allowed as well as wildcards
/apps/modulefiles/*
/apps/other_modulefiles
```

If this file exists then `MODULEPATH_ROOT` method is not used.

The `profile` file is the Lmod initialization script for the bash and zsh shells, the `cshrc` file is for tcsh and csh shells, and the `profile.fish` file is for the fish shell. Please copy or link the `profile` and `cshrc` files to `/etc/profile.d`, and optionally the fish file to `/etc/fish/conf.d`:

```
$ ln -s /opt/apps/lmod/lmod/init/profile      /etc/profile.d/z00_lmod.sh
$ ln -s /opt/apps/lmod/lmod/init/cshrc       /etc/profile.d/z00_lmod.csh
$ ln -s /opt/apps/lmod/lmod/init/profile.fish /etc/fish/conf.d/z00_lmod.fish
```

To test the setup, you just need to login as a user. The module command should be set and `MODULEPATH` should be defined. Bash or Zsh users should see something like:

```
$ type module
module ()
{
    eval $( $LMOD_CMD bash $* )
}

$ echo $LMOD_CMD
/opt/apps/lmod/lmod/libexec/lmod

$ echo $MODULEPATH
/opt/apps/modulefiles/Linux:/opt/apps/modulefiles/Core
```

Similar for csh users:

```
% which module
module: alias to eval `/opt/apps/lmod/lmod/libexec/lmod tcsh !*`

% echo $MODULEPATH
/opt/apps/modulefiles/Linux:/opt/apps/modulefiles/Core
```

If you do not see the module alias then please read the next section.

4.1.4 Integrating module Into Users’ Shells

Bash:

On login, the bash shell first reads `/etc/profile`, and if `profiles.d` is activated, that in turn should source all the `*.sh` files in `/etc/profile.d` with something like:

```
if [ -d /etc/profile.d ]; then
  for i in /etc/profile.d/*.sh; do
    if [ -r $i ]; then
      . $i
    fi
  done
fi
```

Similarly, the system BASHRC file should source all the *.sh files in /etc/profile.d as well. Here is where things can get complicated. See the next section for details.

Bash Shell Scripts:

Bash shell scripts do not source any system or user files before starting execution. Instead it looks for the environment variable BASH_ENV. It treats the contents as a filename and sources it before starting a bash script.

Bash Script Note:

It is important to remember that all bash scripts should start with:

```
#!/bin/bash
```

Starting with:

```
#!/bin/sh
```

won't define the module command, even if sh is linked to bash. Bash will run those scripts in shell emulation mode and won't source the file that BASH_ENV points to.

Csh:

Csh users have an easier time with the module command setup. The system cshrc file is always sourced on every invocation of the shell. The system cshrc file is typically called: /etc/csh.cshrc. This file should source all the *.csh files in /etc/profile.d:

```
if ( -d /etc/profile.d ) then
  set nonomatch
  foreach i (/etc/profile.d/*.csh)
    source $i
  end
  unset nonomatch
endif
```

Zsh:

Zsh users have an easy time with the module command setup as well. The system zshenv file is sourced on all shell invocations. This system file can be in a number of places but is typically in /etc/zshenv or /etc/zsh/zshenv and should have:

```
if [ -d /etc/profile.d ]; then
  setopt no_nomatch
  for i in /etc/profile.d/*.sh; do
    if [ -r $i ]; then
```

(continues on next page)

(continued from previous page)

```
. $i
fi
setopt nomatch
done
fi
```

Fish:

Fish users have [several standard places](#) searched for scripts. The system location is usually `/etc/fish/conf.d` and the user location is usually “`~/.config/fish/conf.d/`”. Fish users are provided a special profile file, ```init/profile.fish`, that should be linked into one of these places with a suitable name. For example, a local user for fish might want:

```
$ ln -s /opt/apps/lmod/lmod/init/profile.fish ~/.config/fish/conf.d/z00_lmod.fish
```

4.1.5 Issues with Bash

Interactive Non-login shells

The Bash startup procedure for interactive non-login shells is complicated and varies between Operating Systems. In particular, Redhat & Centos distributions of Linux as well as Mac OS X have no system `bashrc` read during startup whereas Debian based distributions do source a system `bashrc`. One easy way to tell how bash is set up is to execute the following:

```
$ strings `type -p bash` | grep bashrc
```

If the entire result of the command is:

```
~/.bashrc
```

then you know that your bash shell doesn't source a system `bashrc` file.

If you want to have the same behavior between both interactive shells (login or non-login) and your system doesn't source a system `bashrc`, then you have two choices:

1. Patch bash so that it does source a system `bashrc`. See `contrib/bash_patch` for details on how to do that.
2. Expect all of your bash users to have the following in their `~/.bashrc`

```
if [ -f /etc/bashrc ]; then
. /etc/bashrc
fi
```

As a side note, we at TACC patched bash for a different reason which may apply to your site. When an MPI job starts, it logs into each node with an interactive non-login shell. When we had no system `bashrc` file, many of our fortran 90 programs failed because they required `ulimit -s unlimited` which makes the stack size unlimited. *By patching bash, we could guarantee that it was set by the system on each node.* Sites will have to choose which of the two above methods they wish to deal with this deficiency in bash.

You may have to also change the `/etc/bashrc` (or `/etc/bash.bashrc`) file so that it sources `/etc/profile.d/*.sh` for non-login shells.

Bash Shell Scripts

Bash shell scripts, unlike Csh or Zsh scripts, do not source any system or user files. Instead, if the environment variable, `BASH_ENV` is set and points to a file then this file is sourced before the start of bash script. So by default Lmod sets `BASH_ENV` to point to the bash script which defines the module command.

It may seem counter-intuitive but Csh and Zsh users running bash shell scripts will want `BASH_ENV` set so that the module command will work in their bash scripts.

A bash script is one that starts as the very first line:

```
#!/bin/bash
```

A script that has nothing special or starts with:

```
#!/bin/sh
```

is a shell script. And even if `/bin/sh` points to `/bin/bash` bash runs in a compatibility mode and doesn't honor `BASH_ENV`.

To combat this Lmod exports the definition of the module command. This means that even `/bin/sh` scripts will have the module command defined when run by a Bash User. However, a Csh or Zsh user running a bash script will still need to set `BASH_ENV` and run bash scripts. They won't have the module command defined if they run a sh script.

4.2 How to Transition to Lmod (or how to test Lmod without installing it for all)

In the *Installing Lua and Lmod* document, we described how to install Lua and Lmod for all. Sites which are currently running another environment module system will likely wish to test and then transition from their old module system to Lmod. This can be done smoothly by switching over all users on a set date.

It is important to remember the following facts:

- Lmod reads modulefiles written in TCL. There is typically no need to translate modulefiles written in TCL into Lua. Lmod does this for you automatically.
- Some users can run Lmod while others use the old environment module system.
- However, no user can run both at the same time in the same shell.

Obviously, since you are installing Lmod in your own account, this is a good way to test Lmod without committing your site to switch. Part of this document will describe TACC's transition experience.

4.2.1 Steps for Testing Lmod in your account

1. Install Lua
2. Install Lmod in your account
3. Build the list modules required
4. Purge modules using old module command
5. Reload modules using Lmod

Install Lua

The previous document described how to install Lua. If your system doesn't provide a package for Lua, then it is probably easiest to install the lua tarball found at [sourceforge.net](https://sourceforge.net/projects/lmod/files/lua-W.X.Y.Z.tar.gz) using the following command:

```
$ wget https://sourceforge.net/projects/lmod/files/lua-W.X.Y.Z.tar.gz
```

where you replace the *W.X.Y.Z* with the current version (i.e. 5.1.4.8).

Many Linux distributions already have a lua package and it may even be installed automatically. For example, recent Centos and other Redhat based distributions automatically install Lua as part of the rpm tools.

Once you have lua installed and in your path, you'll need the luafilesystem and luaposix libraries to complete the requirements. See the previous document on how to install these libraries via your package manager or luarocks.

Install Lmod

Please follow the previous document on how to install Lmod. Let's assume that you have installed Lmod in your own account like this:

```
$ ./configure --prefix=$HOME/pkg
$ make install
```

This will install Lmod in **`$HOME/pkg/lmod/x.y.z`** and make a symbolic link to **`$HOME/pkg/lmod/lmod`**.

Build the list of modules required

Many sites provide a default set of modules. When testing, you'll want to be able to load those list of modules using Lmod. Using your old module system, login and do:

```
$ module list

Currently Loaded Modules:

1) a1          3) A          5) b2          7) B
2) a2          4) b1          6) b3
```

It turns out that both the latest version of TCL/C modules and the pure TCL script list a module that loads other modules later in the list. In this made up case we unload module B and notice that unloading the B module also unloads modules b3, b2 and b1. Then unloading the A module also unloads modules a2 and a1. In this case then we would set:

```
export LMOD_SYSTEM_DEFAULT_MODULES=A:B
```

Purge modules using old module command

Execute:

```
$ module purge
```

to unload the currently loaded modules **using the old module command**.

Reload modules using Lmod

Once all modules have been purged and the environment variable `LMOD_SYSTEM_DEFAULT_MODULES` has been set, all that is required is to redefine the module command to use Lmod and to restore the default set of modules by:

```
$ export BASH_ENV=$HOME/pkg/lmod/lmod/init/bash
$ source $BASH_ENV
```

This will define the module command. Finally the default set of modules can be loaded:

```
$ module --initial_load restore
```

This command first looks to see if there is a default collection in `~/.lmod.d/default`. If that file isn't found then it uses the value of variable `LMOD_SYSTEM_DEFAULT_MODULES` as a list of module to load.

If you have gotten this far then you have installed Lmod in your account. Congratulations!

Please test your system. Try to load your most complicated modulefiles. See if **module avail**, **module spider** works and so on.

If you have trouble loading certain TCL modulefiles then read the **How Lmod reads TCL modulefiles** to see why you might have problems.

4.2.2 An example of how this can be done in your bash startup scripts

All the comments above can be combined into a complete example:

```
if [ -z "$_INIT_LMOD" ]; then
    export _INIT_LMOD=1          # guard variable is crucial, to avoid breaking_
    ↪existing modules settings
    type module > /dev/null 2>&1
    if [ "$?" -eq 0 ]; then
        module purge >2 /dev/null    # purge old modules using old module command.
        clearMT                    # clear the stored module table (wipe _
    ↪ModuleTable001_ etc.)
    fi

    export MODULEPATH=...          # define MODULEPATH
    export BASH_ENV=$HOME/pkg/lmod/lmod/init/bash # Point to the new definition of Lmod

    source $BASH_ENV              # Redefine the module command to_
    ↪point
                                # to the new Lmod
    export LMOD_SYSTEM_DEFAULT_MODULES=... # Colon separated list of modules
                                # to load at startup
    module --initial_load --no_redirect restore # load either modules listed above_
    ↪or the
                                # user's ~/.lmod.d/default module_
    ↪collection
else
    source $BASH_ENV              # redefine the module command for_
    ↪sub-shell
    module refresh                # reload all modules but only_
    ↪activate the "set_alias"
                                # functions.
fi
```

Obviously, you will have to define **MODULEPATH** and **LMOD_SYSTEM_DEFAULT_MODULES** to match your site setup. The reason for the guard variable **_INIT_LMOD** is that the module command and the initialization of the modules is only done in the initial login shell. On any sub-shells, the module command gets defined (again). Finally the **module refresh** command is called to define any alias or shell functions in any of the currently loaded modules.

4.2.3 How to Transition to Lmod: Staff & Power User Testing

Once you have tested Lmod personally and wish to transition your site to use Lmod, I recommend the following strategy for staff and friendly/power users for testing:

1. Install Lua and Lmod in system locations
2. Install */etc/profile.d/z00_lmod.sh* to redefine the module command
3. Load system default modules (if any) after previous steps
4. Only users who have a file named *~/.lmod* use Lmod
5. At TACC, we did this for 6 months.

Using this strategy, you can have extended testing without exposing Lmod to any user which hasn't opted-in.

How to Deploy Lmod

Once Staff testing is complete and you are ready to deploy Lmod to your users it is quite easy to switch to an opt-out strategy:

1. Change */etc/profile.d/z00_lmod.sh* so that everyone is using Lmod
2. If a user has a *~/.no.lmod* then they continue to use your original module system
3. At TACC we did this for another 6 months
4. We broke Environment Module support with the family directive
5. We now only support Lmod
6. Both transitions generated very few tickets (2+2)

4.3 Lua Modulefile Functions

Lua is an efficient language built on simple syntax. Readers wanting to know more about lua can see <http://www.lua.org/>. This simple description given here should be sufficient to write all but the most complex modulefiles.

It is important to understand that modulefiles are written in the positive. That is, one writes the actions necessary to activate the package. A modulefile contains commands to add to the PATH or set environment variables. When loading a modulefile the commands are followed. When unloading a modulefile the actions are reversed. That is the element that was added to the PATH during loading, is removed during unloading. The environment variables set during loading are unset during unloading.

prepend_path ("PATH";"/path/to/pkg/bin"): prepend to a path-like variable the value.

prepend_path ("PATH";"/path/to/pkg/bin", "sep"): prepend to a path-like variable the value. It is possible to add a third argument to be the separator. By default is ":", the separator can be any single character for example " " or ";;"

append_path ("PATH";"/path/to/pkg/bin"): append to a path-like variable the value.

append_path (“PATH”,*“/path/to/pkg/bin”*, *“sep”*): append to a path-like variable the value. It is possible to add a third argument to be the separator. By default is “:”, the separator can be any single character for example ” ” or “;”

remove_path (“PATH”,*“/path/to/pkg/bin”*): remove value from a path-like variable. This command is a no-op when the mode is unload.

remove_path (“PATH”,*“/path/to/pkg/bin”*, *“sep”*): remove value from a path-like variable. This command is a no-op when the mode is unload. It is possible to add a third argument to be the separator. By default is “:”, the separator can be any single character for example ” ” or “;”

setenv (“NAME”, *“value”*): assigns to the environment variable “NAME” the value.

pushenv (“NAME”, *“value”*): sets NAME to value just like **setenv**. In addition it saves the previous value in a hidden environment variable. This way the previous state can be returned when a module is unloaded. **pushenv** (“FOO”,false) will clear “FOO” and the pop will return the previous value.

unsetenv (“NAME”): unset the value associated with “NAME”. This command is a no-op when the mode is unload.

whatis (“STRING”): The **whatis** command can be called repeatedly with different strings. See the Administrator Guide for more details.

help ([*help string*]): What is printed out when the help command is called. Note that the *help string* can be multi-lined.

pathJoin (“/a”,*“b/c/”*,*“d/”*): builds a path: “/a/b/c/d”. It combines any number of strings with one slash and removes excess slashes. Note that trailing slash is removed. If you need a trailing slash then do **pathJoin**(“/a”,*“b/c”*) .. “/” to get “/a/b/c/”.

depends_on (“pkgA”, “pkgB”, “pkgC”): loads all modules. When unloading only dependent modules are unloaded. See [Dependent Modules](#) for details.

load (“pkgA”, “pkgB”, “pkgC”): load all modules. Report error if unable to load.

try_load (“pkgA”, “pkgB”, “pkgC”): load all modules. No errors reported if unable to load.

always_load (“pkgA”, “pkgB”, “pkgC”): load all modules. However when this command is reversed it does nothing.

set_alias (“name”,*“value”*): define an alias to name with value.

unload (“pkgA”, “pkgB”): In both load and unload mode, the modulefiles are unloaded. It is not an error to unload modules that were not loaded.

family (“name”): A user can only have one family “name” loaded at a time. For example family(“compiler”) would mean that a user could only have one compiler loaded at a time.

prereq (“name1”, “name2”): The current modulefile will only load if **all** the listed modules are already loaded.

prereq_any (“name1”, “name2”): The current modulefile will only load if **any** of the listed modules are already loaded.

conflict (“name1”, “name2”): The current modulefile will only load if all listed modules are NOT loaded.

4.3.1 Extra functions

The entries below describe several useful commands that come with Lmod that can be used in modulefiles.

os.getenv (“NAME”): Ask for environment for the value of “NAME”. Note that if the “NAME” might not be in the environment, then it is probably best to do:

```
local foo=os.getenv("FOO") or ""
```

otherwise `foo` will have the value of `nil`.

capture (“string”): Run the “string” as a command and capture the output. This function uses the value of `LD_PRELOAD` and `LD_LIBRARY_PATH` found when Lmod is configured. Use **subprocess** if you wish to use the current values.

subprocess (“string”) Run the “string” as a command and capture the output.

isFile (“name”): Returns true if “name” is a file.

isDir (“name”): Returns true if “name” is a directory.

splitFileName (“name”): Returns both the directory and the file name. `local d,f=splitFileName("/a/b/c.ext")`. Then `d="/a/b"`, `f="c.ext"`

LmodMessage (“string”,...): Prints a message to the user.

LmodError (“string”,...): Print Error string and exit without loading the modulefile.

mode (): Returns the string “load” when a modulefile is being loaded and “unload” when unloading.

isloaded (“NAME”): Return true when module “NAME” is loaded.

LmodVersion (): The version of lmod.

execute {`cmd=<any command>`},`modeA={“load”}`} Run any command with a certain mode. For example **execute** {`cmd="ulimit -s unlimited"`},`modeA={“load”}`} will run the command **ulimit -s unlimited** as the last thing that the loading the module will do.

4.3.2 Modifier functions to prereq and loads

atleast (“name”,“version”): This modifier function will only succeed if the module is “version” or newer.

between (“name”,“v1”,“v2”): This modifier function will only succeed if the module’s version is equal to or between “v1” and “v2”.

latest (“name”): This modifier function will only succeed if the module has the highest version on the system.

4.3.3 Introspection Functions

The following functions allow for more generic modulefiles by finding the name and version of a modulefile.

myModuleName (): Returns the name of the current modulefile without the version.

myModuleVersion (): Returns the version of the current modulefile.

myModuleFullName (): Returns the name and version of the current modulefile.

myModuleUsrName (): Returns the name the user specified to load a module. So it could be the name or the name and version.

myFileName (): Returns the absolute file name of the current modulefile.

myShellName (): Returns the name of the shell the user specified on the command line.

myShellType (): Returns the shellType based on the name of the shell the user specified on the command line. It returns sh for sh, bash, zsh, csh for csh, tcsh. Otherwise it is the same as **myShellName** ().

hierarchyA (“fullName”, level): Returns the hierarchy of the current module. See the section on Generic Modules for more details.

4.4 Module names and module naming conventions

The modulefiles contain commands that change the user's environment to make an application or library available to a user. So loading a modulefile named `gcc/7.1` should hopefully make the programs Gnu Compiler Collection (i.e. `gcc`, `gfortran` and `g++` compilers) available to a user.

So the `gcc/7.1` module is a file named either `7.1` or `7.1.lua` depending on whether the commands are written in TCL or Lua respectively and this file is in the `gcc` directory.

To see a more complete module layout, let's assume your site has has a `MODULEPATH` set to:

```
/apps/modulefiles/Core:/apps/modulefiles/Other
```

And the directory layout from `/apps/modulefiles` is:

```
$ cd /apps/modulefiles; tree -a
```

```
.
├── Core
│   ├── A
│   │   ├── 1.0
│   │   └── 2.0.lua
│   ├── gcc
│   │   ├── 5.4
│   │   └── 7.1.lua
│   └── StdEnv.lua
└── Other
    ├── C
    │   ├── 3.3.lua
    │   └── 3.4.lua
    └── D
        └── 4.0.lua
```

In the above structure there are 8 modulefiles. So the 5 modulefiles under `Core` are: `A/1.0`, `A/2.0`, `gcc/5.4`, `gcc/7.1` and `StdEnv`. Note that Lmod's name of a modulefile doesn't include the `.lua` extension. Also note that both `A` and `gcc` have two versions where as `StdEnv` has a name with no version.

4.4.1 fullName == sn/version

The `fullName` of a modulefile is complete name which is made of two parts: `sn` and `version`. The `sn` is the shortname and represents the minimum name that can be used to specify a module. So for the `gcc/7.1` modulefile. The full-name is `gcc/7.1` and the `sn` is `gcc` and the `version` is `7.1`. This naming convention is known as `NAME/VERSION` and is abbreviated `N/V`. There are two more complicated naming schemes known as `Category/Name/Version` (a.k.a. `C/N/V`) and `Name/Version/Version` (a.k.a. `N/V/V`). In all three naming schemes, a modulefile has a `fullName` which is split into `sn / version`. The split between the `sn` and `version` will be different depending on which naming scheme is used. Sometimes the `version` doesn't exist like for `StdEnv` module. In this case the `fullName` is the same as the `sn`.

In next section *How Lmod Picks which Modulefiles to Load*, we explain how a Lmod takes an `sn` and determines which modulefile to load. But in short, Lmod picks the marked default or if there is no default then Lmod picks the "highest" version.

4.4.2 Special Names for modulefiles

If the `version` has a leading dot then this modulefile is hidden. That is, it can be loaded but it won't be shown by `module avail`. If the `sn` name has a leading dot, such as `.X` then every version of this module will be hidden. So `.X/1.0` is hidden and so is `Y/.2.0`.

Because of the way that Lmod passes information between module command invocations, a modulefile **CANNOT** have two or more leading underscores. So naming a modulefile `_A/1.0` is acceptable but having two underscore such as `__B/1.0` is not!

4.4.3 Module Naming scheme: Category/Name/Version (C/N/V)

Sites may wish to group similar modules into categories. For example, all the biology packages grouped together, like bowtie and tophat packages be named: `bio/bowtie/3.4` or `bio/tophat/2.7`. When the `fullName` is `bio/bowtie/3.4` then the `sn` for is `bio/bowtie` and the version is `3.4`.

Sites should think carefully about choosing to using C/N/V. This can make it easier for users to know which modules provide say physics, chemistry or biology applications but it does lead to great deal more typing of which tab completion provided by the bash or zsh shells can only do so much.

Sites can have meta-module inside a category. A meta-module is a module that has no version. For example, suppose you have the following C/N/V modulefile structure:

```
$ cd /apps/modulefiles/Other; tree -a
.
├── bio
│   ├── bowtie
│   │   └── 3.1
│   ├── genomics.lua
│   └── tophat
│       └── 7.2
```

In this case there are three modules `bio/bowtie/3.1`, `bio/tophat/7.2`. The `bio/genomic` module is a meta module where the `fullName` is the same as the `sn` and it has no version.

The rule that Lmod uses to determine which modules are meta modules and which have a version is the following. If a file is in a directory that has other sub-directories (other than `.` and `..`), the file is a meta-module. So the `genomic.lua` file is part of the `sn` for `bio/genomics` because the `genomics.lua` file is in a sub-directory that has the directories `bowtie` and `tophat`.

Sites can mix N/V and C/N/V layouts, Lmod will be able to decide the `sn` and versions by walking directory tree. In general, the `fullName`, will be divided into directories names to become the `sn` and the version will be the file. So for the `fullName` `bio/tophat/7.2` the directories `bio` and `tophat` become the `sn`, `bio/tophat` and the version is `7.2`.

Lmod supports as many directory levels as site likes. For example, a site could have a modulefile named `A/B/C/D/1.1` where the `sn` name is `A/B/C/D` and the version is `1.1`.

4.4.4 Module Naming scheme: Name/Version/Version (N/V/V)

A site many wish to have directories has part of the version. This could be used to have 32 or 64 bit versions of the same package. So for example a site might wish to have a modules named `acme/32/4.2` and `acme/64/4.2` where `sn` is `acme` and the versions are `32/4.2` and `64/4.2`.

By default, Lmod assumes that the version is just the file, so lmod needs to be told where the `sn` / version split is. This can be done by creating an empty `.version` or `.modulerc` file where a site wants the split to be. For the above example, the following layout make `acme` be the `sn`:

```
$ tree -a
.
├── acme
│   └── .version
```

(continues on next page)

(continued from previous page)

```

├── 32
│   └── 4.2
└── 64
    └── 4.2

```

because there is a `.version` at the same level as the 32 and 64 directories. With the `.version` file, the `fullName` is `acme/64/4.2` and the `sn` is `acme` and the version `64/4.2`. If the `.version` file was removed then the `sn` would be `acme/64` and the version would be `4.2`.

Sites are can name modules with as many directories as they like. For example a site can have a module named:

```
mpi/mpich/64/3.1/048
```

If there is an empty file at `mpi/mpich/.version` then the `sn` would be `mpi/mpich` and the version would be `64/3.1/048`.

4.5 How Lmod Picks which Modulefiles to Load

Lmod uses the directories listed in `MODULEPATH` to find the modulefiles to load. Suppose that you have a single directory `/opt/apps/modulefiles` that has the following files and directories:

```

/opt/apps/modulefiles

StdEnv.lua  ucc/  xyz/

./ucc:
8.1.lua  8.2.lua

./xyz:
10.1.lua

```

Lmod will report the following directory tree like this:

```

----- /opt/apps/modulefiles -----
StdEnv    ucc/8.1    ucc/8.2 (D)    xyz/10.1

```

We note that the `.lua` extension has not been reported above. The `.lua` extension tells Lmod that the contents of the file are written in the Lua language. All other files are assumed to be written in TCL.

Here the name of the file or directory under `/opt/apps/modulefiles` is the name of the module. The normal way to specify a module is to create a directory to be the name of the module and the file(s) under that directory are the version(s). So we have created `ucc` and `xyz` directories to be the names of the module. There are two version files under `ucc` and one version file under `xyz`.

The `StdEnv.lua` file is another way to specify a module. This file is a module with no version associated with it. These are typically used as a meta-module: a module that loads other modules.

4.5.1 N/V: Picking modules when there are multiple directories in MODULEPATH

The follow rules apply when the module layout is either Name/Version (N/V) or Category/Name/Version (C/N/V). The rules are a little different for Name/Version/Version (N/V/V) as described in the section below. When there are multiple directories specified in `MODULEPATH`, the rules get more complicated on what modulefile to load. Lmod uses the following rules to locate a modulefile:

1. It looks for an exact match in all `MODULEPATH` directories. It picks the first match it finds. This is true of hidden modules. Specifying the `fullName` of a module will load it.
2. If the user requested name is a full name and version, and there is no exact match then it stops.
3. If the name doesn't contain a version then Lmod looks for a marked default in the first directory that has one. A marked default which is also hidden will be loaded.
4. Finally it looks for the "Highest" Version in all `MODULEPATH` directories. If there are two or more modulefiles with the "Highest" version then the first one in `MODULEPATH` order will be picked. Note that hidden modulefiles are ignored when choosing the "Highest".
5. As a side note, if there are two version files, one with a `.lua` extension and one without, the lua file will be used over the other one. It will be like the other file is not there.

As an example, suppose you have the following module tree:

```
----- /home/user/modulefiles -----
xyz/11.1

----- /opt/apps/modulefiles -----
StdEnv    ucc/8.1    ucc/8.2    xyz/10.1

----- /opt/apps/mfiles -----
ucc/8.3 (D)    xyz/12.1 (D)
```

If a user does the following command:

```
$ module load ucc/8.2 xyz
```

then `ucc/8.2` will be loaded because the user specified a particular version and `xyz/12.1` will be loaded because it is the highest version across all directories in `MODULEPATH`.

4.5.2 Marking a Version as Default

Suppose you have several versions of the mythical UCC compiler suite:

```
$ module avail ucc
----- /opt/apps/modulefiles/Core -----
ucc/8.1    ucc/9.2    ucc/11.1    ucc/12.2 (D)
```

and you would like to make the 11.1 version the default. Lmod searches three different ways to mark a version as a default in the following order. The first way is to make a symbolic link between a file named "default" and the desired default version.:

```
$ cd /opt/apps/modulefiles/Core/ucc; ln -s 11.1.lua default
```

A second way to mark a default is with a `.modulerc.lua` file in the same directory as the modulefiles.:

```
module_version("ucc/11.1", "default")
```

A third way to mark a default is with a `.modulerc` file in the same directory as the modulefiles.:

```
##Module
module-version ucc/11.1 default
```

There is a fourth method to pick the default module. If you create a `.version` file in the `ucc` directory that contains:

```

#%Module
set    ModulesVersion    "11.1"

```

Please note that a `.modulerc.lua`, `.modulerc` or `.version` file must be in the same directory as the `11.1.lua` file in order for Lmod to read it.

Using any of the above three ways will change the default to version 11.1.

```

$ module avail ucc
----- /opt/apps/modulefiles/Core -----
ucc/8.1   ucc/9.2   ucc/11.1 (D)   ucc/12.2

```

4.5.3 Highest Version

If there is no marked default then Lmod chooses the “Highest” version across all directories:

```

$ module avail ucc

----- /opt/apps/modulefiles/Core -----
ucc/8.1   ucc/9.2   ucc/11.1   ucc/12.2

----- /opt/apps/modulefiles/New -----
ucc/13.2 (D)

```

The “Highest” version is by version number sorting. So Lmod “knows” that the following versions are sorted from lowest to highest:

```

2.4dev1
 2.4a1
2.4beta2
 2.4rc1
  2.4
2.4.0.0
 2.4-1
2.4.0.0.1
 2.4.1

```

4.5.4 N/V/V: Picking modules when there are multiple directories in MODULEPATH

The rules are different when the module layout is Name/Version/Version (N/V/V). The rules for N/V can be called `Find Best` whereas N/V/V is `Find First`. Note that if any one of the directories in `MODULEPATH` are in N/V/V format, the whole tree is searched with N/V/V rules. Below are the rules that Lmod uses to locate a modulefile when in N/V/V mode:

1. It looks for an exact match in all `MODULEPATH` directories. It picks the first match it finds.
2. If there is no exact match then Lmod finds the first match for the names that it has. It matches by directory name. No partial matches are done.
3. In the directory that is found above the first marked default is found
4. If there are no marked defaults, then the “highest” is chosen.
5. The two above rules are followed at each directory level.

For example with the following module tree where `foo` is the name of the module and rest are version information:

```
----- /apps/modulefiles/A -----
foo/2/1  foo/2/4    foo/3/1    foo/3/2 (D)

----- /apps/modulefiles/B -----
foo/3/3    foo/3/4
```

Then the commands `module load foo` and `module load foo/3` would both load `foo/3/2`. The command `module load foo/2` would load `foo/2/4`.

When searching for `foo`, Lmod finds it in the `A` directory. Then seeing a choice between 2 and 3 it picks 3 as it is higher. Then in the `foo/3` directory it chooses 2 as it is higher than 1. To load any other `foo` module, the full name will have to be specified.

4.5.5 Marking a directory as default in an N/V/V layout

There are three ways to mark a directory as a default: Using a default symlink, or the use of either the `.modulerc` or `.version` files. Since it is possible (but not recommended) to have all three possibilities, This is the same technique that was used before to mark a particular version file when in an N/V layout. Lmod chooses the setting of the default directory in the following order:

1. default symlink
2. `.modulerc.lua`
3. `.modulerc`
4. `.version`

Suppose that you have the following architecture split with (32,64,128) bit libraries and you want the 64 directory to be the default. With the following structure:

```
----- /apps/modulefiles/A -----
foo/32/1    foo/64/1    foo/128/1
foo/32/4    foo/64/2 (D)  foo/128/2
```

You can have a symlink for `/apps/modulefiles/A/foo/default` which points to `/apps/modulefiles/A/foo/64`. Or you can have the contents of `/apps/modulefiles/A/foo/.modulerc` contain:

```
##Module
module-version 64 default
```

or you can have the contents of `/apps/modulefiles/A/foo/.modulerc.lua` contain:

```
module_version("64","default")
```

or you can have the contents of `/apps/modulefiles/A/foo/.version` contain:

```
##Module
set ModulesVersion "64"
```

Normally the 128 directory would be chosen as the default directory as 128 is higher than 64 or 32 but any one of these files forces Lmod to pick 64 over the other directories.

4.5.6 Why do N/V/V module layouts use Find First over Find Best?

The main problem here is that of the default directories. There is no sane way to pick. Suppose that you have the following structure:

```

----- /apps/modulefiles/A -----
foo/32/1    foo/64/1    foo/128/1
foo/32/4    foo/64/2 (D)  foo/128/2

----- /apps/modulefiles/B -----
foo/32/5    foo/64/3    foo/128/3
foo/32/6    foo/64/4    foo/128/4

```

And where the default directory in A in 64 and in B it is 32. When trying to load `foo/64` the site has marked 64 the default in A where as it is not in B. Does that mean that `foo/64/2` is “higher” than `foo/64/4` or not. There is no clear reason to pick one over the other so Lmod has chosen Find First for N/V/V module layouts.

For sites that are mixing N/V and N/V/V module layouts they may wish to change Lmod to use the find first rule in all cases. See [Configuring Lmod for your site](#) to see how to configure Lmod for find first.

4.5.7 Autoswapping Rules

When Lmod autoswaps hierarchical dependencies, it uses the following rules:

1. If a user loads a default module, then Lmod will reload the default even if the module version has changed.
2. If a user loads a module with the version specified then Lmod will only load the exact same version when swapping dependencies.

For example a user loads the intel and boost library:

```

$ module purge; module load intel boost; module list

Currently Loaded Modules:
1) intel/15.0.2  2) boost/1.57.0

```

Now swapping the Intel compiler suite for the Gnu compiler suite:

```

The following have been reloaded with a version change:
1) boost/1.57.0 => boost/1.56.0

```

Here boost has been reloaded with a different version because the default is different in the gcc hierarchy. However if the user does:

```

$ module purge; module load intel boost/1.57.0; module list

Currently Loaded Modules:
1) intel/15.0.2  2) boost/1.57.0

```

And:

```

$ module swap intel gcc;

Inactive Modules:
1) boost/1.57.0

```

Since the user initially specified loading `boost/1.57.0` then Lmod assumes that the user really wants that version. Because version 1.57.0 of boost isn’t available under the gcc hierarchy, Lmod marks this boost module as inactive. This is true even though version 1.57.0 is the default version of boost under the Intel hierarchy.

4.6 Providing A Standard Set Of Modules for all Users

Users can be provided with an initial set of modulefiles as part of the login procedure. Once a list of modulefiles has been installed, please create a file called `StdEnv.lua` and place it in the `$MODULEPATH` list of directories, typically `/opt/apps/modulefiles/Core/StdEnv.lua`. The name is your choice, the purpose is provide a standard list of modules that get loaded during login. In `StdEnv.lua` is something like:

```
load("name1", "name2", "name3")
```

Using the `/etc/profile.d` directory system described earlier to create a file called `z01_StdEnv.sh`

```
if [ -z "$__Init_Default_Modules" ]; then
    export __Init_Default_Modules=1;

    ## ability to predefine elsewhere the default list
    LMOD_SYSTEM_DEFAULT_MODULES=${LMOD_SYSTEM_DEFAULT_MODULES:-"StdEnv"}
    export LMOD_SYSTEM_DEFAULT_MODULES
    module --initial_load --no_redirect restore
else
    module refresh
fi
```

Similar for `z01_StdEnv.csh`:

```
if ( ! $?__Init_Default_Modules ) then
    setenv __Init_Default_Modules 1
    if ( ! $?LMOD_SYSTEM_DEFAULT_MODULES ) then
        setenv LMOD_SYSTEM_DEFAULT_MODULES "StdEnv"
    endif
    module --initial_load restore
else
    module refresh
endif
```

The `z01_Stdenv.*` names are chosen because the files in `/etc/profile.d` are sourced in alphabetical order. These names guarantee they will run after the module command is defined.

The `z01_Stdenv.sh` now includes `--no_redirect`. This option prevents sites that configure Lmod to write messages to stdout to write them to stderr instead. This is important as any messages written to stdout during shell startup causes scp copies to fail. Csh/tcsh cannot write messages to stdout due to limitations in that shell.

The first time these files are sourced by a shell they will set `LMOD_SYSTEM_DEFAULT_MODULES` to `StdEnv` and then execute `module restore`. Any subshells will instead call `module refresh`. Both of these statements are important to get the correct behavior out of Lmod.

The `module restore` tries to restore the user's default collection. If that doesn't exist, it then uses contents of the variable `LMOD_SYSTEM_DEFAULT_MODULES` to find a colon separated list of Modules to load.

The `module refresh` solves an interesting problem. Sub shells inherit the environment variables of the parent but do not normally inherit the shell aliases and functions. This statement fixes this. Under a "refresh", all the currently loaded modules are reloaded but in a special way. Only the functions which define aliases and shell functions are active, all others functions are ignored.

The above is an example of how a site might provide a default set of modules that a user can override with a default collection. Sites are, of course, free to set up Lmod any way they like. The minimum required setup (for bash with `z01_StdEnv.sh`) would be:

```

if [ -z "$__Init_Default_Modules" ]; then
    export __Init_Default_Modules=1;

    module --initial_load restore
else
    module refresh
fi

```

The module restore command still depends on the environment variable LMOD_SYSTEM_DEFAULT_MODULES but that can be set somewhere else.

4.6.1 Lmod, LD_LIBRARY_PATH and screen

In general, it is probably better to NOT use `screen` and use `tmux` instead. The problem with `screen` is that it is a *set-group-id* (SGID) program (`tmux` is not). That means it uses the group associated with the executable and not the user's group. The main consequence of this is that the operating system removes `LD_LIBRARY_PATH` from the environment. This is a security feature built into the operating system.

A site could change `z01_StdEnv.sh` to have:

```

if [ -z "$__Init_Default_Modules" -o -z "$LD_LIBRARY_PATH" ]; then
    export __Init_Default_Modules=1;

    module --initial_load restore
else
    module refresh
fi

```

to help with the situation. This will force Lmod restore the initial set of modules (or the user's default collection). This works fine as long as the initial set of modules actually sets `LD_LIBRARY_PATH`. If it doesn't every interactive sub-shell will do a module restore, which is probably not what you want. For example, if you see the following then you probably want to remove the test for an empty `LD_LIBRARY_PATH`:

```

$ module list
Currently Loaded Modules:
  1) git/2.7      2) StdEnv

$ module load bowtie
$ bash
$ module list
Currently Loaded Modules:
  1) git/2.7      2) StdEnv

```

Running the bash shell caused the module restore to run which unloaded all modules and restored the modules back to the initial set.

4.7 Converting from TCL/C Environment Modules to Lmod

Sites converting from the TCL/C based Environment Modules (a.k.a Tmod) should be aware of some of the differences between Tmod and Lmod. Lmod is a complete re-implementation of the environment module concept with no code reuse from the original program. There are some differences between that can affect how easy your transition to Lmod.

One major difference between the two tools is that Lmod is written in Lua and not TCL. Lmod has to translate TCL into Lua. This means that pure TCL statements are evaluated by a program called `tcl2lua.tcl`. This program outputs lua

statements. This can lead to difference between Tmod and Lmod because different order that statements are evaluated. The details are discussed at [How does Lmod convert TCL modulefile into Lua](#).

Another important difference between the tools is that Lmod has the *One Name* rule. That is you can only have one “name” loaded at a time. Suppose a user tries to load gcc/5.4 and gcc/7.1 at the same time:

```
$ module load gcc/5.4
$ module load gcc/7.1
```

The following have been reloaded with a version change:

```
1) gcc/5.4 => gcc/7.1
```

Lmod is telling you that the gcc/5.4 has been replaced by gcc/7.1. This happens automatically without anything special in the modulefiles.

4.7.1 Module Naming rules

The fullname of a module is split into name/version. Normally the version is just the string after the last slash. So gcc/5.4 has a name of gcc and a version of 5.4. A module named compiler/gcc/5.4 would have a name of compiler/gcc and a version of 5.4.

With Lmod 7+, sites can change the name - version split if they like. For example a site might name their modules bowtie/64/3.4 and bowtie/32/3.4 where the name of the module would be bowtie and the version would be either 64/3.4 or 32/3.4. Please see [N/V/V: Picking modules when there are multiple directories in MODULEPATH](#) for details.

ADVANCED TOPICS

5.1 How to report a bug in Lmod

Lmod has some built-in tools to make debugging possible on your site. The first feature of Lmod is the configuration report:

```
$ module --config
```

This reports how Lmod has been configured at build time as well as any `LMOD_*` environment variables set. The second tool is the debug output also built-in to Lmod:

```
$ module -D load foo 2> load.log
```

The `-D` option turns on the debug printing and will report all the steps that Lmod took to load a module called `foo`. Note that the configuration report is at the top of every debug output.

5.1.1 Steps to report a bug

1. Test your bug against the latest release from github. Please pull the HEAD branch.
2. Try to reduce the problem to the fewest number of modules. Shoot for 1 or 2 modulefiles if you can.
3. Run the command that fails. i.e. `module -D cmd module ... 2> lmod.log`
4. Combine the `lmod.log` file, the modulefiles from step 2, and possibly the spider cache file into a tar file.
5. Send the tar file to mclay@tacc.utexas.edu

5.2 Rules for PATH-like variables

Lmod provides great flexibility in handling prepending and appending to path-like variables. This is especially true when there are duplicate entries. A modulefile can modify a variable like `PATH` using `append_path()` or `prepend_path()` or their TCL equivalents. For example, suppose that `PATH=/usr/bin:/usr/local/bin` then:

```
prepend_path("PATH", "/bin")
```

would change `PATH` to `/bin:/usr/bin:/usr/local/bin`. The interesting question is what happens when the following is executed:

```
prepend_path("PATH", "/usr/bin")
```

That is, when `/usr/bin` is already in `$PATH` or any other duplicate entry.

5.2.1 LMOD_DUPLICATE_PATHS == yes

Lmod supports two main styles for dealing with duplicates. If `$LMOD_DUPLICATE_PATHS` is yes (or Lmod is configured that way). Then duplicates entries are allowed (assume `PATH` is empty):

```
prepend_path("PATH", "/A")  --> PATH = /A
prepend_path("PATH", "/B")  --> PATH = /B:/A
prepend_path("PATH", "/A")  --> PATH = /A:/B:/A
```

When unloading a modulefile with `prepend_path()`, Lmod removes the first matching entry it finds. Reversing an `append_path()`, Lmod removes the last matching entry.

5.2.2 LMOD_DUPLICATE_PATHS == no

The default setting of Lmod is that duplicates are not allowed. When prepending, Lmod pushes the directory to be first even if it is a duplicate (assume `PATH` is empty):

```
append_path("PATH", "/A")   --> PATH = /A
prepend_path("PATH", "/B")   --> PATH = /B:/A
prepend_path("PATH", "/A")   --> PATH = /A:/B
```

When duplicates are not allowed, Lmod maintains a reference count on each entry. That is, Lmod knows that “/A” has appended/prepended twice and “/B” once. This means that two `prepend_path(“PATH”, “/A”)` will be required to completely remove “/A” from `$PATH`.

5.2.3 LMOD_TMOD_PATH_RULE == yes

If this env. var is set (or configured), then Lmod does not change the order of entries but it does increase the reference count (assume `$PATH` is empty):

```
append_path("PATH", "/A")   --> PATH = /A
prepend_path("PATH", "/B")   --> PATH = /B:/A
prepend_path("PATH", "/A")   --> PATH = /B:/A
```

Here we see that prepending “/A” does not change the order of directories in `$PATH`. Obviously if `LMOD_TMOD_PATH_RULE` is yes then duplicates are not allowed.

5.2.4 Special treatment for \$MODULEPATH

The `MODULEPATH` environment variable is treated special. No duplicates entries are ever allowed even if `LMOD_DUPLICATE_PATHS == yes`. It always uses reference counting for `PATH` entries. In order to not confuse users. The command:

```
$ module unuse /foo/bar
```

will always remove the path entry, even if the reference count is greater than 1. Also a user can always clear `$MODULEPATH` with:

```
$ module unuse $MODULEPATH
```

5.2.5 Specifying Priorities for PATH entries

There are rare occasions where a site might want a directory to at the beginning of the PATH. This can be done by giving a priority to a path:

```
prepend_path("PATH", "/foo/bar", priority=100)
```

Note the use of curly braces instead of parentheses and setting priority to a number. Lmod groups the entries of the same priority together. This means that `/foo/bar` will likely be at the beginning of `$PATH` as long as no other entry has a higher priority.

Assuming that PATH is initially empty, here is an example:

```
prepend_path("PATH", "/foo", priority=100)  --> PATH = /foo
prepend_path("PATH", "/A")                  --> PATH = /foo:/A
prepend_path("PATH", "/B")                  --> PATH = /foo:/B/A
```

Lmod remembers the priority between invocations, meaning that you'll get the same results even if the following where in three separate modulefiles.

5.2.6 An Example of Loading and Unloading a Module

Above we showed that there are three modes for path like variables:

1. LMOD_DUPLICATE_PATH=no
2. LMOD_DUPLICATE_PATH=no, LMOD_TMOD_PATH_RULE=yes
3. LMOD_DUPLICATE_PATH=yes

Let's assume that `$PATH = /A:/B:/C` and the module FOO is:

```
prepend_path("PATH", "/C")
```

then the following table shows what happens for each of the three modes when loading and unloading FOO. Note that `/A (2)` is the path entry `/A` a reference count of 2:

Action	1	2	3
original PATH	/A(1):/B(1):/C(1)	/A(1):/B(1):/C(1)	/A:/B:/C
module load FOO	/C(2):/A(1):/B(1)	/A(1):/B(1):/C(2)	/C:/A:/B:/C
module unload FOO	/C(1):/A(1):/B(1)	/A(1):/B(1):/C(1)	/A:/B:/C

For mode (1) where no duplicates are allowed, upon loading FOO path `/C` is moved to the beginning and stays there when unloading FOO. For mode (2), If a directory is already in `$PATH`, it is not moved, only the ref count is increased on load and decreased upon unload. Finally in mode (3), loading causes `/C` to be placed at the beginning and unloading removes it from the beginning.

When duplicates are allowed and unloading a module, Lmod does not remember which module inserted which directory where, it just removes the first or last entry depending on whether it was a `prepend_path()` or `append_path()` respectively. Also there is no reference counting when duplicates are allowed. It is not necessary and doesn't make sense.

5.3 How to use a Software Module hierarchy

Libraries built with one compiler need to be linked with applications built with the same compiler version. For High Performance Computing there are libraries called Message Passing Interface (MPI) that allow for efficient communication between tasks on a distributed memory computer with many processors. Parallel libraries and applications must be built with a matching MPI library and compiler. To make this discussion clearer, suppose we have the intel compiler version 15.0.1 and the gnu compiler collection version 4.9.2. Also we have two MPI libraries: mpich version 3.1.2 and openmpi version 1.8.2. Finally we have a parallel library HDF5 version 1.8.13 (phdf5).

Of the many possible ways of specifying a module layout, this flat layout of modules is a reasonable way to do this:

```
$ module avail

----- /opt/apps/modulefiles -----
gcc/4.9.2                phdf5/gcc-4.9-mpich-3.1-1.8.13
intel/15.0.2             phdf5/gcc-4.9-openmpi-15.0-1.8.13
mpich/gcc-4.9-3.1.2      phdf5/intel-15.0-mpich-3.1-1.8.13
mpich/intel-15.0-3.1.2   phdf5/intel-15.0-openmpi-15.0-1.8.13
openmpi/gcc-4.9-1.8.2
openmpi/intel-15.0-1.8.2
```

In order for users to load the matching set of compilers and MPI libraries, they will have to load the matching set of modules. For example this would be correct:

```
$ module load gcc/4.9.2 openmpi/gcc-4.9-1.8.2 phdf5/gcc-4.9-openmpi-15.0-1.8.13
```

But it is quite easy to load an incompatible set of modules. Now, it is possible that the system administrators at your site might have setup `conflict s` to avoid loading mismatched modules. However, using conflicts can be fragile. What happens if a site adds a new compiler such as clang or pgi or a new mpi stack? All those module file conflict statements will have to be updated.

A different strategy is to use a software hierarchy. In this approach a user loads a compiler which extends the **MODULEPATH** to make available the modules that are built with the currently loaded compiler (similarly for the mpi stack).

Our modulefile hierarchy is stored under `/opt/apps/modulefiles/{Core,Compiler,MPI}`. The Core directory is for modules that are not dependent on Compiler or MPI implementations. The Compiler directory is for packages which are only Compiler dependent. Lastly, the MPI directory is packages which dependent on MPI-Compiler pairing. The modulefiles for the compilers are placed in the Core directory. For example the gcc version 4.9.2 file is in `Core/gcc/4.9.2.lua` and contains:

```
-- Setup Modulepath for packages built by this compiler
local mroot = os.getenv("MODULEPATH_ROOT")
local mdir = pathJoin(mroot,"Compiler/gcc", "4.9")
prepend_path("MODULEPATH", mdir)
```

This code asks the environment for **MODULEPATH_ROOT** which is `/opt/apps/modulefiles` and the last two lines prepend `/opt/apps/modulefiles/Compiler/gcc/4.9` to the **MODULEPATH**.

The modulefiles for the MPI implementations are placed under the Compiler directory because they only depend on a compiler. The openmpi module file for the gcc-4.9.2 compiler is then stored at `/opt/apps/modulefiles/Compilers/gcc/4.9/openmpi/1.8.2.lua` and it contains:

```
-- Setup Modulepath for packages built by this MPI stack
local mroot = os.getenv("MODULEPATH_ROOT")
local mdir = pathJoin(mroot,"MPI/gcc", "4.9", "openmpi", "1.8")
prepend_path("MODULEPATH", mdir)
```

The above code will prepend `/opt/apps/modulefiles/MPI/gcc/4.9/openmpi/1.8` to the **MODULEPATH**.

The above description is a suggested way to handle the modulefile software hierarchy. The software packages themselves can be stored in many ways. For software packages, but not the modulefiles, we store them in another software hierarchy as follows:

1. Core packages: **`/opt/apps/pkgName/version`**
2. Compiler dependent packages: **`/opt/apps/compilerName-version/pkgName/version`**
3. MPI-Compiler dependent packages: **`/opt/apps/compilerName-version/mpiName-version/pkgName/version`**

When **MODULEPATH** changes, Lmod unloads any modules which are not currently in the **MODULEPATH** and then tries to reload all the previously loaded modules. Any modules which are not available are marked as inactive. Those inactive modules become active if found with new **MODULEPATH** changes.

Note: In all of the examples above, we used only the first two version numbers. In other words, we used 4.9 instead of 4.9.2 and similarly 1.8 instead of 1.8.2. It is our view that for at least compilers and MPI stacks, the third digit is typically a bug fix and doesn't require rebuilding all the dependent packages. Y.M.M.V.

5.4 Configuring Lmod for your site

Sites can control the behavior of Lmod at configuration time. After Lmod is configured and installed, user can also modify the behavior through environment variables. To see all the configuration options you can execute:

```
$ ./configure --help
```

in the Lmod source directory. There are a few behavior options that do not have a configuration option.

There two kinds of variables: (1) An explicit values; (2) a yes/no variable. An example of first kind is `LMOD_SITE_NAME`. This variable controls the site name (e.g. TACC). This value of variable is used directly. There is no change of case or any other changes in that string.

The second kind of variable is an yes/no variable. One example of this is `LMOD_IGNORE_CACHE`. When this variable is "yes", Lmod ignores any cache files and walks **MODULEPATH** instead.

The following settings are considered "no". Note that the string value is lowercased first, so NO, No, and nO are the same as no. ALL OTHER VALUES are treated as "yes".

1. `export LMOD_IGNORE_CACHE=""`
2. `export LMOD_IGNORE_CACHE=0`
3. `export LMOD_IGNORE_CACHE=no`
4. `export LMOD_IGNORE_CACHE=off`

5.4.1 Environment variables only

The following variables set actions that can only be controlled by environment variables. The actions can not be controlled through the configuration step.

LMOD_ADMIN_FILE: [path] If set this will be a file to specify the nag message. If this variable has no value, then Lmod looks for `<install_dir>/../etc/admin.list`

LMOD_IGNORE_CACHE: [yes/no] If set to yes then Lmod will bypass all cachefile and walk the directories in MODULEPATH instead.

LMOD_PAGER: [string] Lmod uses a pager when not using redirect. It defaults to less. Site/Users can turn off the pager if it is set to “None”.

LMOD_RTM_TESTING: [any value] If this variable has any value it means that Lmod does nothing. This is useful when testing a personal copy of Lmod and your site has the SHELL_STARTUP_DEBUG package installed so that the invoking of the module command in the system startup will a no-op.

LMOD_SYSTEM_NAME: [string] This variable is used to where a site is using shared home files systems. See [Lmod on Shared Home File Systems](#) for more details.

LMOD_TRACING: [yes/no] If set to yes then Lmod will trace the loads/unloads while the module command is running.

5.4.2 Configuration time settings that can be overridden by env. vars.

The following settings are defined by configure but can be overridden by environment variables. The brackets show the following values [kind, default: value, configuration option] where kind is either yes/no, string, path, etc, value is what the default will be. Finally the configuration option which will set the action.

LMOD_ALLOW_TCL_MFILES: [yes/no, default: yes, `–with-tcl`]. Allow tcl modulefiles. Note that .version and .modulerc files still use the tcl interpreter. So setting this to no means that your site will have to use either the “default” symlink or “.modulerc.lua” to specify defaults.

LMOD_ANCIENT: [number, default:86400, `–with-ancient`]. The number of seconds that the user’s personal cache is considered valid.

LMOD_AUTO_SWAP: [yes/no, default: yes, `–with-autoSwap`] Allows Lmod to swap any modules that use the family function such as compilers and mpi stacks.

LMOD_CACHED_LOADS: [yes/no, default:no, `–with-cachedLoads`] If “yes” then Lmod will use the spider cache to load modulefiles and produce a terse avail instead of walking all the directories in MODULEPATH as long as LMOD_IGNORE_CACHE is not set.

LMOD_CASE_INDEPENDENT_SORTING: [yes/no, default: no, `–with-caseIndependentSorting`] Make avail and spider use case independent sorting.

LMOD_COLORIZE: [yes/no, default: yes, `–with-colorize`] Let lmod write colorize message to the terminal.

LMOD_DISABLE_NAME_AUTOSWAP: [yes/no, default: no, `–with-disableNameAutoSwap`] Setting this to “yes” disables the one name rule autoswapping. In other words, “module load gcc/4.7 gcc/5.2 will fail when this is set.

LMOD_DUPLICATE_PATHS: [yes/no, default: no, `–with-duplicatePaths`] Allow duplicates directories in path-like variables, PATH, LD_LIBRARY_PATH, ... Note that if LMOD_TMOD_PATH_RULE is “yes” then LMOD_DUPLICATE_PATH is set to “no”.

LMOD_EXACT_MATCH: [yes/no, default: no, `–with-exactMatch`] Requires Lmod to use fullNames for modules. This disables defaults.

LMOD_HIDDEN_ITALIC: [yes/no, default: no, `–with-hiddenItalic`] Use italics for hidden modules instead of faint.

LMOD_MPATH_AVAIL: [yes/no, default: no, `–with-mpathSearch`] If this is set then module avail <string> will search modulepath names.

LMOD_OVERRIDE_LANG: [string, default: en, `–with-lang`] Override \$LANG for Lmod error/messages/warnings.

LMOD_PIN_VERSIONS: [yes/no, default: no, `–with-pinVersions`] If yes then when restoring load the same version that was chosen with the save, instead of the current default version.

LMOD_PREPEND_BLOCK: [normal/reverse, default: normal, `--with-prependBlock`] Treat multiple directories passed to prepend in normal order and not reversed.

LMOD_REDIRECT: [yes/no, default: no, `--with-redirect`]. Normal messages generated by “module avail”, “module list”, etc write the output to stderr. Turning redirect to “yes” will cause these messages to be written to stdout. Note this only works for bash and zsh. This will not work with csh or tcsh as there is a problem with these shells and not Lmod.

LMOD_SHORT_TIME: [number, default: 2, `--with-shortTime`]. If the time to build the spider cache takes longer than this number then write the spider cache out into the user’s account. If you want to prevent the spider cache file being written to the user’s account then set this number to be large, like 86400.

LMOD_SITE_MSG_FILE: [full path, default: <nil> `--with-siteMsgFile`] The Site message file. This overrides the messageDir/en.lua file so that sites can replace some or all Lmod messages.

LMOD_SITE_NAME: [string, default: <nil>, `--with-siteName`]. This is the site name, for example TACC, and not the name of the cluster. This is used with the family function.

LMOD_SYSHOST: [string, default: <nil>, `--with-syshost`]. This variable can be used to help with module tracking. See [Tracking Module Usage](#) for details.

LMOD_TMOD_FIND_FIRST: [yes/no, default: no, `--with-tmodFindFirst`]. Normally Lmod uses the FIND BEST rule to search for defaults when searching C/N/V or N/V module layouts. A site can force FIND_FIRST for C/N/V or N/V module layouts to match the FIND_FIRST rule for N/V/V module layout. See [N/V/V: Picking modules when there are multiple directories in MODULEPATH](#) for more details.

LMOD_TMOD_PATH_RULE: [yes/no, default: no, `--with-tmodPathRule`]. Normally Lmod prepend/appends a directory in the beginning/end of the path like variable. If this is true then if path entry is already there then do not prepend/append. Note that if LMOD_TMOD_PATH_RULE is “yes” then LMOD_DUPLICATE_PATH is set to “no”.

LMOD_USE_DOT_FILES: [yes/no, default: yes, `--with-useDotFiles`] If yes then use `~/.lmod.d/.cache`, if no then use `~/.lmod.d/__cache__`

5.4.3 Configuration only settings

`--with-silentShellDebug`: [yes/no, default: no] If yes then the module command will silence its output under shell debug.

5.5 How does Lmod convert TCL modulefile into Lua

Lmod uses a TCL program called **tcl2lua.tcl** to read TCL modulefiles and convert them to lua. The whole TCL modulefile is run through. However instead of executing the “module functions” they are converted to Lua. For example, suppose you have the following simple TCL modulefile for git:

```
#%Module
set appDir          $env(APP_DIR)
set version         2.0.3

prepend-path        PATH "$appDir/git/$version/bin"
```

Assuming that the environment variable APP_DIR is */apps* then the output of the **tcl2lua.tcl** program would be:

```
prepend_path("PATH", "/apps/git/2.0.3/bin")
```

Note that all the normal TCL code has been evaluated and the TCL **prepend-path** command has been converted to a lua **prepend_path** function call.

Normally this works fine. However, because Lmod does evaluate the actions of a TCL module file as a two-step process, it can cause problems. In particular, suppose you have two TCL modulefiles:

Centos:

```
#!/Module
setenv SYSTEM_NAME Centos
```

And B:

```
#!/Module
module load Centos

if { $env(SYSTEM_NAME) == "Centos" } {
    # do something
}
```

When Lmod tries to translate the B modulefile into lua it fails:

```
load("Centos")
LmodError([[/opt/mfiles/B: (???): can't read "env(SYSTEM_NAME)": no such variable]])
```

This is because unlike the TCL/C Module system, the **module load Centos** command is converted to a function call, but it won't load the module in time for the test to be evaluated properly.

The only solution is convert the B modulefile into a Lua modulefile (B.lua):

```
load("Centos")
if (os.getenv("SYSTEM_NAME") == "Centos") then
    -- Do something
end
```

The Centos modulefile does not have to be translated in order for this to work, just the B modulefile.

As a side note, you are free to put Lua modules in the same tree that the TCL/C Module system uses. It will ignore files that the first line is not **#!/Module** and Lmod will pick B.lua over B.

5.6 Dependent Modules

Let's assume that module "X" depends on module "A". There are several ways to handle module dependency. Inside the "X" modulefile you could have one of the following choices:

1. Use `prereq("A")`
2. Use `load("A")`
3. Use `always_load("A")`
4. Use `RPATH` to make "X" know where the libraries in "A" can be found.
5. Use `if (not isloaded("A")) then load("A") end`
6. Use `if (not isloaded("A")) then always_load("A") end`
7. Use `depends_on("A")`

Let's examine these choices in order. The main issue for each of these choices is what happens when module "X" is unloaded.

5.6.1 `prereq("A")`

This choice is the one you give for sophisticated users. If a user tried to load module "X" without previously loading "A" then the user will get a message telling the user that they must load "A" before loading "X". This way the dependency is explicitly handled by the user. When the user unloads "X", module "A" will remain loaded.

5.6.2 `load("A")`

This choice will always load module "A" on the users behalf. This is true even if "A" is already loaded. When module "X" is unloaded, module "A" will be unloaded as well. This may surprise some users who might want to continue using the "A" package. At least with `prereq()`, your users won't be surprised by this. Another way to handle this is the next choice.

5.6.3 `always_load("A")`

A site can chose to use `always_load()` instead. This command is a shorthand for:

```
if (module() == "load") then
    load("A")
end
```

The TCL equivalent is:

```
if { [ module-info mode load ] } {
    module load A
}
```

This choice will always load module "A" on the users behalf. This is true even if "A" is already loaded. When module "X" is unloaded, module "A" will remain loaded.

5.6.4 `if (not isloaded("A")) then load("A") end`

This choice will load module "A" on the users behalf if it not already loaded. When module "X" is unloaded, module "A" will be unloaded as well.

5.6.5 `if (not isloaded("A")) then always_load("A") end`

This choice will load module "A" on the users behalf if it not already loaded. When module "X" is unloaded, module "A" will remain loaded.

Use RPATH

We have switched to using RPATH for library dependencies at TACC. That is when we build package X, we use the RPATH linking option to link libraries in package A as part of the X rpm. This has the disadvantage that if the A package is removed then the X package is broken. This has happened to us occasionally. In general, however, we have found that this has worked well for us.

5.6.6 depends_on ("A")

This choice loads module “A” on the users behalf if it not already loaded. When module “X” is unloaded, module “A” will be unloaded if it is a dependent load. Imagine the following scenario with `depends_on ("A")`:

```
$ module purge; module load X; module unload X           => unload A
$ module purge; module load A; module load X; module unload X => keep A
```

Lmod implements reference counting for modules loaded via `depends_on()` and only `depends_on()`. So if “X” and “Y” depend on “A” then:

```
$ module purge; module load X Y; module unload X   => keep A
$ module purge; module load X Y; module unload X Y => unload A
```

5.6.7 Complex uses of depends_on ()

Sites can have complex dependencies that they might wish to express using `depends_on()`. Let’s assume that module X depends on the openblas package but only when using gcc. A site might try to do the following in the X modulefile:

```
-- DO NOT DO THIS!!
if (isloaded("gcc")) then
    depends_on("openblas")
end
-- DO NOT DO THIS!!
```

Let’s also assume that your site is using the hierarchy where X is a compiler dependent module and you wish to use the same modulefile for X for both the gcc and intel compiler modules. The above code in the X modulefile works correctly when loading but will fail when unloading in this common scenario: if a user tries to swap gcc for intel then the openblas module will likely be left loaded or inactive.

To simplify the discussion, let’s have the user start with the following modules loaded:

```
$ module list

1) gcc/7.1  2) X/1.0  3) openblas/0.2.20
```

And lets assume that openblas is a Core module and X is a compiler dependent module. Then executing:

```
$ module swap gcc intel
```

causes gcc to be unloaded. When gcc is unloaded it removes the path in `MODULEPATH` to the gcc dependent modules which means that X will be unloaded and marked as inactive. The way that a module is unloaded is that the contents of the module is evaluated and most action requested are reversed. So load statements cause a module to unload and a `depends_on()` function is told to `forgo()` the modules. The `isloaded()` is not reversed. But as you can see since the gcc modulefile is not loaded the if statement then clause is not evaluated. This means that openblas will still be loaded.

In the case where openblas is a compiler-dependent module then it will be unloaded and marked as inactive. Either way this probably not what the site wants to happen. The trouble here is that environment that happens on load is not the case on unload.

There is another way to determine which compiler and/or mpi stack a module is in and that is its filename. This assumes that you have a rational naming convention for module locations. Using a similar technique to the one describe in *Introspection*. We can determine which compiler is in use. So if the module file is located in `/apps/mfiles/Compiler/<compiler>/<compiler-version>/<app-name>/<app-version>` then we can do the following and use the `hierarchyA()` function in the X modulefile:

```

local hierA = hierarchyA(myModuleFullName(), 1)
if (hierA[1]:find("^gcc/")) then
    depends_on("openblas")
end

```

This will work correctly for both loading and unloading. This, of course, assumes that the location of the X modulefile is something like:

```
/apps/mfiles/Compiler/gcc/7.1/X/1.0.lua
```

5.7 Modulefile Examples from simple to complex

Most of the time a modulefile is just a collection of setting environment variables and prepending to PATH or other path like variables. However, modulefiles are actually programs so you can do a great deal if necessary.

Here we show some of the techniques that site's or user's might use in their modulefiles. In addition, there are many examples in the Lmod source tree. The `rt` directory contains the regression testing suite. Each subdirectory in `rt` is a separate test and below that there will be many modulefiles. For example see all the modulefile associated with the load test can be found in `rt/load/mf`.

5.7.1 Some simple modulefiles

An application modulefile might add to \$PATH, set a few other environment variables and provide help message as well as `whatis()` strings. For example the valgrind memory usage tester might look like:

```

help([[
To use the valgrind utility on an executable called a.out:

valgrind ./a.out
]])

local version = "3.7.0"
local base    = pathJoin("/apps/valgrind", version)
prepend_path("PATH", pathJoin(base, "bin")) -- /app/valgrind/3.7.0/bin
setenv("SITE_VALGRIND_DIR", base)
setenv("SITE_VALGRIND_INC", pathJoin(base, "include"))
setenv("SITE_VALGRIND_LIB", pathJoin(base, "lib"))

whatis("Name: .. pkgName")
whatis("Version: " .. fullVersion)
whatis("Category: tools")
whatis("URL: http://www.valgrind.org")
whatis("Description: memory usage tester")

```

A library module might look like:

```

help([[
...
]])
whatis("Name: boost")
whatis("Version: 1.64")
whatis("Category: Lmod/Modulefiles")
whatis("Keywords: System, Library, C++")

```

(continues on next page)

(continued from previous page)

```

whatis("URL: http://www.boost.org")
whatis("Description: Boost provides free peer-reviewed portable C++ source libraries.
↪")

setenv("SITE_BOOST_DIR", "/apps/intel17/boost/1.64")
setenv("SITE_BOOST_LIB", "/apps/intel17/boost/1.64/lib")
setenv("SITE_BOOST_INC", "/apps/intel17/boost/1.64/include")
setenv("SITE_BOOST_BIN", "/apps/intel17/boost/1.64/bin")
setenv("BOOST_ROOT", "/apps/intel17/boost/1.64")
prepend_path("LD_LIBRARY_PATH", "/apps/intel17/boost/1.64/lib")
prepend_path("PATH", "/apps/intel17/boost/1.64/bin")

```

5.7.2 Introspection

Lmod provides inspection functions that describe the name and version of a modulefile as well as the path to the modulefile. These functions provide a way to write “generic” modulefiles, i.e. modulefiles that can fill in its values based on the location of the file itself.

These ideas work best in the software hierarchy style of modulefiles. For example: suppose the following is a modulefile for Git. Its modulefile is located in the “/apps/mfiles/Core/git” directory and software is installed in “/apps/git/<version>”. The following modulefile would work for every version of git:

```

local pkg = pathJoin("/apps", myModuleName(), myModuleVersion())
local bin = pathJoin(pkg, "bin")
prepend_path("PATH", bin)

whatis("Name:      ", myModuleName())
whatis("Version:   ", myModuleVersion())
whatis("Description: ", "Git is a fast distributive version control system")

```

The contents of this modulefile can be used for multiple versions of the git software, because the local variable `bin` changes the location of the bin directory to match the version of the used as the name of the file. So if the module file is in `/apps/mfiles/Core/git/2.3.4.lua` then the local variable `bin` will be `/apps/git/2.3.4`.

5.7.3 Relative Paths

Suppose you are interested in modules where the module and application location are relative. Suppose that you have an \$APPS directory, and below that you have modulefiles and packages, and you would like the modulefiles to find the absolute path of the package location. This can be done with the `myFileName()` function and some lua code:

```

local fn      = myFileName()           -- 1
local full    = myModuleFullName()     -- 2
local loc     = fn:find(full, 1, true) - 2 -- 3
local mdir    = fn:sub(1, loc)          -- 4
local appsDir = mdir:gsub("(.)", "%1") -- 5
local pkg     = pathJoin(appsDir, full) -- 6

```

To make this example concrete, let’s assume that applications are in `/home/user/apps` and the modulefiles are in `/home/user/apps/mfiles`. So if the modulefile is located at `/home/user/apps/mfiles/git/1.2.lua`, then that is the value of `fn` at line 1. The `full` variable at line 2 will have `git/1.2`. What we want is to remove the name of the modulefile and find its parent directory. So we use Lua string member function on `fn` to find where `full` starts. In most cases `fn:find(full)` would work to find where the “git” starts in `fn`. The trouble is

that the Lua find function is expecting a regular expression and in particular `.` and `-` are regular expression characters. So here we are using `fn:find(full, 1, true)` to tell Lua to treat each character as is with no special meaning.

Line 3 also subtracts 2. The find command reports the location of the start of the string where the “g” in “git” is, We want the value of `mdir` to be `/home/user/apps/mfiles` so we need to subtract 2. This makes `mdir` have the right value. One note is that Lua is a one based language, so locations in strings start at one.

It was important for the value of `mdir` to remove the trailing `/` so that line 5 will do its magic. We want the parent directory of `mdir`, so the regular expressions says greedily grab every character until the trailing `/` and the `%1` says to capture the string found in and use that to set `appsdire` to `/home/user/apps`. Finally we wish to set `pkg` to the location of the actual application so we combine the value of `appsdire` and `full` to set `pkg` to `/home/user/apps/git/1.2`.

The nice thing about this Lua code is that it figures out the location of the package no matter where it is, as long as the relation between apps directories and modulefiles is consistent.

Creating modules like this can be complicated. See [Debugging Modulefiles](#) for helpful tips.

5.7.4 Generic Modules with the Hierarchy

This works great for Core modules. It is a little more complicated for Compiler or MPI/Compiler dependent modules but quite useful. For a concrete example, lets cover how to handle the boost C++ library. This is obviously a compiler dependent module. Suppose you have the gnu compiler collection (gcc) and the intel compiler collection (intel), which means that you’ll have a gcc version and an intel version for each version of booth.

In order to have generic modules for compiler dependent modules, there must be some conventions to make this work. A suggested way to do this is the following:

1. Core modules are placed in `/apps/mfiles/Core`. These are the compilers, programs like git and so on.
2. Core software goes in `/apps/<app-name>/<app-version>`. So git version 2.3.4 goes in `/apps/git/2.3.4`
3. Compiler-dependent modulefiles go in `/apps/mfiles/Compiler/<compiler>/<compiler-version>/<app-name>/<app-version>` using the **two-digit** rule (discussed below). So the Boost 1.55.0 modulefile built with gcc/4.8.3 would be found in `/apps/mfiles/Compiler/gcc/4.8/boost/1.55.0.lua`
4. Compiler-dependent packages go in `/apps/<compiler-version>/<app-name>/<app-version>`. So the same Boost 1.55.0 package built with gcc 4.8.3 would be placed in `/apps/gcc-4_8/boost/1.55.0`

The above convention depends on the **two-digit** rule. For compilers and mpi stack, we are making the assumption that compiler dependent libraries built with gcc 4.8.1 can be used with gcc 4.8.3. This is not always safe but it works well enough in practice. The above convention also assumes that the boost 1.55.0 package will be placed in `/apps/gcc-4_8/boost/1.55.0`. It couldn’t go in `/apps/gcc/4.8/...` because that is where the gcc 4.8 package would be placed and it is not a good idea to co-mingle two different packages in the same tree. Another possible choice would be `/apps/gcc-4.8/boost/1.55.0`. It is my view that it looks too much like the gcc version 4.8 package location where as `gcc-4_8` doesn’t.

With all of the above assumptions, we can now create a generic module file for compiler dependent modules such as Boost. In order to make this work, we will need to use the `hierarchyA` function. This function parses the path of the modulefile to return the pieces we need to create a generic boost modulefile:

```
hierA = hierarchyA(myModuleFullName(), 1)
```

The `myModuleFullName()` function returns the full name of the module. So if the module is named **boost/1.55.0**, then that is what it will return. If your site uses module names like `lib/boost/1.55.0` then it will return that correctly as well. The `1` tells Lmod to return just one component from the path. So if the modulefile is located at `/apps/mfiles/Compiler/gcc/4.8/boost/1.55.0.lua`, then `myModuleFullName()` returns **boost/1.55.0** and the `hierarchyA` function returns an array with 1 entry. In this case it returns:

```
{ "gcc/4.8" }
```

The rest of the module file then can make use to this result to form the paths:

```
local pkgName      = myModuleName()
local fullVersion  = myModuleVersion()
local hierA        = hierarchyA(myModuleFullName(), 1)
local compilerD    = hierA[1]:gsub("/", "-"):gsub("%.", "_")
local base         = pathJoin("/apps", compilerD, pkgName, fullVersion)

whatis("Name:  ..pkgName)
whatis("Version ..fullVersion)
whatis("Category: library")
whatis("Description: Boost provides free peer-reviewed "..
      " portable C++ source libraries.")
whatis("URL: http://www.boost.org")
whatis("Keyword: library, c++")

setenv("TACC_BOOST_LIB", pathJoin(base, "lib"))
setenv("TACC_BOOST_INC", pathJoin(base, "include"))
```

The important trick is the building of the *compilerD* variable. It converts the *gcc/4.8* into *gcc-4_8*. This makes the *base* variable be: */apps/gcc-4_8/boost/1.55.0*.

Creating modules like this can be complicated. See *Debugging Modulefiles* for helpful tips.

A proposed directory structure of */apps/mfiles/Compiler* would be:

```
.base/      gcc/  intel/

.base/
boost/generic.lua

gcc/4.8/boost/

1.55.0.lua -> ../../../../base/boost/generic.lua

intel/15.0.2/boost/

1.55.0.lua -> ../../../../base/boost/generic.lua
```

In this way the *.base/boost/generic.lua* file will be the source file for all the boost version build with gcc and intel compilers.

The same technique can be applied for modulefiles for Compiler/MPI dependent packages. In this case, we will create the *phdf5* modulefile. This is a parallel I/O package that allows for Hierarchical output. The modulefile is:

```
local pkgName      = myModuleName()
local pkgVersion    = myModuleVersion()
local pkgNameVer    = myModuleFullName()

local hierA        = hierarchyA(pkgNameVer, 2)
local mpiD         = hierA[1]:gsub("/", "-"):gsub("%.", "_")
local compilerD    = hierA[2]:gsub("/", "-"):gsub("%.", "_")
local base         = pathJoin("/apps", compilerD, mpiD, pkgNameVer)

setenv("TACC_HDF5_DIR", base)
setenv("TACC_HDF5_DOC", pathJoin(base, "doc"))
```

(continues on next page)

(continued from previous page)

```

setenv(      "TACC_HDF5_INC",    pathJoin(base, "include"))
setenv(      "TACC_HDF5_LIB",    pathJoin(base, "lib"))
setenv(      "TACC_HDF5_BIN",    pathJoin(base, "bin"))
prepend_path("PATH",            pathJoin(base, "bin"))
prepend_path("LD_LIBRARY_PATH", pathJoin(base, "lib"))

whatis("Name: Parallel HDF5")
whatis("Version: " .. pkgVersion)
whatis("Category: library, mathematics")
whatis("URL: http://www.hdfgroup.org/HDF5")
whatis("Description: General purpose library and file format for storing scientific_
↳data (parallel I/O version)")

```

We use the same tricks as before. It is just that since the module for phdf5 built by gcc/4.8.3 and mpich/3.1.2 will be found at `/apps/mfiles/MPI/gcc/4.8/mpich/3.1/phdf5/1.8.14.lua`. The results of `hierarchyA(pkgNameVer,2)` would be:

```
{ "mpich/3.1", "gcc/4.8" }
```

This is because the `hierarchyA` works back up the path two elements at a time because the full name of this package is also two elements (phdf5/1.8.14). The `base` variable now becomes:

```
/apps/gcc-4_8/mpich-3_1/phdf5/1.8.14
```

The last type of modulefile that needs to be discussed is an mpi stack modulefile such as mpich/3.1.2. This modulefile is more complicated because it has to implement the two-digit rule, build the path to the package and build the new entry to the **MODULEPATH**. The modulefile is:

```

local pkgNameVer = myModuleFullName()
local pkgName    = myModuleName()
local fullVersion = myModuleVersion()
local pkgV       = fullVersion:match('( %d+%. %d+)%?.?')

local hierA      = hierarchyA(pkgNameVer,1)
local compilerV  = hierA[1]
local compilerD  = compilerV:gsub("/", "-"):gsub("%.", "_")
local base       = pathJoin("/apps", compilerD, pkgName, fullVersion)
local mpath      = pathJoin("/apps/mfiles/MPI", compilerV, pkgName, pkgV)

prepend_path("MODULEPATH", mpath)
setenv(      "TACC_MPICH_DIR", base)
setenv(      "TACC_MPICH_LIB", pathJoin(base, "lib"))
setenv(      "TACC_MPICH_BIN", pathJoin(base, "bin"))
setenv(      "TACC_MPICH_INC", pathJoin(base, "include"))

whatis("Name: " .. pkgName)
whatis("Version " .. fullVersion)
whatis("Category: mpi")
whatis("Description: High-Performance Portable MPI")
whatis("URL: http://www.mpich.org")

```

The **Two Digit** rule implemented by forming the `pkgV` variable. The `base` and `mpath` are:

```

base = "/apps/gcc-4_8/mpich-3_1/phdf5/1.8.14"
mpath = "/apps/mfiles/MPI/gcc/4.8/mpich/3.1"

```

The `rt` directory contains all the regression test used by Lmod. As such they contain many examples of modulefiles. To complement this description, the `rt/hierarchy/mf` directory from the source tree contains a complete hierarchy.

5.8 The Interaction between Modules, MPI and Parallel Filesystems

Many sites that use modules have hundred to thousands of nodes connected to large parallel filesystems. Some care has to be taken when a user's job scripts start a parallel mpi execution to avoid problems with the parallel filesystem. Module commands in a user's `~/bashrc` or `~/cshrc` can overload or cause timeouts in a parallel filesystem.

At TACC, we do a couple of things to avoid this problem. We have seen this timeout problem for mpi programs that execute more than 1000 nodes, but when this problem occurs will depend on relative speeds of the network and the parallel filesystem.

It is helpful to outline the procedure that we use at TACC to start a job on a compute node for a user

1. The bash user submits a job to the scheduler.
2. The current environment is captured
3. That environment passed to the shell script that starts on a compute node.
4. The user's job script starts an mpi execution.
5. The mpi execution starts an interactive non-login (and not prompt) shell on every node.
6. This non-login interactive shell sources the user's `~/bashrc`
7. Then the environment found at the start of mpi execution is passed to all nodes.

It is the sixth step that causes the problem. If there are any module load commands in `~/bashrc`, these module commands will be started on every node at about the same time. The module command has to walk the directories listed in `$MODULEPATH` to find the modules. This hits the parallel filesystem hard when there are thousands of nodes.

At TACC we define an environment variable call `ENVIRONMENT` (a clever name, no?) to be `BATCH` when a job is started on a compute node. We provide every user with a default `~/bashrc` that has the following section:

```
if [ -z "$__BASHRC_SOURCED__" -a "$ENVIRONMENT" != BATCH ]; then
  export __BASHRC_SOURCED__=1

  #####
  # **** PLACE MODULE COMMANDS HERE and ONLY HERE. ****
  #####

  # module load git

fi
```

This way modules are only loaded once on the initial shell and not reloaded on subshells or on compute nodes during job submission.

Some users won't follow our guidelines so as an extra layer of protection we make the module be a no-op for bash users.

The module command for bash is defined to be:

```
module () {
  eval $( $LMOD_CMD bash "$@" ) ...
}
```

Normally, `$LMOD_CMD` points the `lmod` command but on compute nodes we have the following startup behavior for bash users:


```
# Compute nodes set the ENVIRONMENT var to BATCH for non-interactive shells.
if [ -z "$PS1" ]; then
    export ENVIRONMENT=BATCH

    # If we are in BATCH mode then turn off the module command.
    if [ -z "$TACC_DEBUG" ]; then
        export LMOD_CMD=':'
        export LMOD_SETTARG_CMD=':'
    fi
fi
```

By making `$LMOD_CMD` be a colon, we make the module command silently do nothing. Unfortunately, this only works for bash and not tcsh or zsh users. We want module command to work in a user's job submission script:

```
#!/bin/bash
#SBATCH ...

module load intel mvapich2
ibrun ./my_mpi_program           # start parallel execution.
```

Note that the total environment is passed by our parallel job starter `ibrun`. So there is no need for a user's `~/bashrc` or similar file to load modules. In fact it can lead to problems if the user's current module environment loads modules that don't match the environment that `ibrun` pushes to the compute nodes.

Tcsh and Zsh source the startup scripts sourced in `/etc/profile.d/` for shell script startup. This is how the module command is defined there. If we redefined `$LMOD_CMD` for those shell, module command would not work in tcsh or zsh scripts.

Bash uses a different technique. It only uses the value of `$BASH_ENV`. If that variable points to a file then that file is source. Lmod defines that variable to point to a file that defines the module command. So re-defining `$LMOD_CMD` in the startup scripts won't affect the module commands in the job submission script, just in the parallel execution when `~/bashrc` is sourced.

One drawback to redefining `$LMOD_CMD` is that if bash user tries to invoke a shell script to run in parallel any module command will silently do nothing.

5.9 Lmod on Shared Home File Systems

Many sites have a single Operating System and one set of modules across their cluster. If a site has more than one cluster, they may chose to have a separate home directory for each cluster. Some sites may wish to have multiple clusters share a single home directory. While this strategy has some advantages, it complicates things for your users and administrators. If your site has a single home directory sharing between two or more clusters, you have a shared home file system.

As a further complication, your site may or may not have a shared home file system even if you have two or more clusters. If you have separate login nodes for each cluster then you do have a shared home file system. If you have a single login which can submit jobs to different clusters then you do not have a shared home file system.

The way to think about this is each cluster is going to have at least some modules which are different. Module collections need to be unique to each cluster. The trick described below will make them unique for each cluster.

Sites that use a shared home file system across multiple clusters should take some extra steps to ensure the smooth running of Lmod. Typically each cluster will use different modules.

There are three steps that will make Lmod run smoothly on a shared home filesystem:

1. It is best to have a separate installation of Lmod on each cluster.

2. Define the environment variable “LMOD_SYSTEM_NAME” uniquely for each cluster.
3. If you build a system spider cache, then build a separate cache for each cluster.

A separate installation on each cluster is the safest way to install Lmod. It is possible to have a single installation but since there is some C code build with Lmod, this has to work on all clusters. Also the location of the Lua interpreter must be exactly the same on each cluster.

It is also recommended that you set “LMOD_SYSTEM_NAME” outside of a modulefile. It would be bad if a module purge would clear that value. When you set this variable, it makes the module collections and user spider caches unique for a given cluster.

A separate system spider cache is really the only way to go. Otherwise a “module spider” will report modules that don’t exist on the current cluster. If you have a separate install of Lmod on each cluster then you can specify the location of system cache at configure time. If you don’t, you can use the “LMOD_RC” environment to specify the location of the lmodrc.lua file uniquely on each cluster.

Lmod knows about the system spider cache from the lmodrc.lua file. If you install separate instances of Lmod on each cluster, Lmod builds the scDescriptT table for you. Otherwise you can modify lmodrc.lua to point to the system cache by adding scDescriptT to the end of the file:

```
scDescriptT = {
  {
    ["dir"] = "<location of your system cache directory>",
    ["timestamp"] = "<location of your timestamp file",
  },
}
```

where you have filled in the location of both the system cache directory and timestamp file.

5.10 User Spider Cache

In *System Spider Cache*, we described how to build a system cache. If there is no system cache available then Lmod can produce a user based spider cache. It gets written to `~/.lmod.d/.cache`. It is designed to provide improved speed of performing doing module avail or module spider. But it is not without its problems. The first point is that if Lmod thinks any spider cache is valid, it uses it for the MODULEPATH directories it covers then it uses it instead of walking the tree.

Personal Cache rules:

1. If it can’t find a valid cache then Lmod walks the tree to find all available modules and builds the spider cache in memory.
2. If the time it takes to build the cache is longer than the contents of env var. LMOD_SHORT_TIME (default 10 seconds) then Lmod writes the cache file into the `~/.lmod.d/.cache` directory.
3. A user’s cache is assumed to be valid for the contents of LMOD_ANCEIENT_TIME (default 86400 seconds or 24 hours) based on the date associated with the cache file.

Sites can change these defaults at configure time. Users can set the environment variables to change their personal setting.

To turn off the generation of a user cache, one can set LMOD_SHORT_TIME to some big number of seconds. For example:

```
export LMOD_SHORT_TIME=86400
```

would say that Lmod only write the user cache file if it took longer than 1 day (=86400 seconds). A second way do to this is do make the user cache directory unwritable:

```
$ rm -rf ~/.lmod.d/.cache
$ chmod 500 ~/.lmod.d/.cache
```

5.11 System Spider Cache

It is now very important that sites with large modulefile installations build system spider cache files. There is a shell script called “`update_lmod_system_cache_files`” that builds a system cache file. It also touches a file called “`system.txt`”. Whatever the name of this file is, Lmod uses this file to know that the spider cache is up-to-date.

Lmod uses the spider cache file as a replacement for walking the directory tree to find all modulefiles in your `MODULEPATH`. This means that Lmod only knows about system modules that are found in the spider cache. Lmod won’t know about any system modules that are not in this cache. (Personal module files are always found). It turns out that reading a single file is much faster than walking the directory tree.

The spider cache is used to speed up `module avail` and `module spider` and not `module load`. All the spider cache file(s) provide is a way for Lmod to know what modules exist and any properties that a modulefile might have. It does not save the contents of any modulefiles. Lmod always reads and evaluate the actual modulefile when performing loads, shows and similar commands.

The reason that Lmod does not use the cache with `module load` is that if the spider cache is out-of-date, then Lmod will not be able to load a module. Either Lmod uses the spider cache or it walks the directories in `MODULEPATH`.

A site may choose to use have the spider cache assist the `module load` command by configuring Lmod or setting the environment variable:

```
export LMOD_CACHED_LOADS=yes
```

See *Configuring Lmod for your site* for more details. Just remember that the cache file has to be up-to-date or user’s won’t be able to find system modulefiles! Note too, that a cache file is tied to a particular set of directories in the `MODULEPATH`. Lmod knows which directories in `MODULEPATH` are covered by spider cache file(s) and which are not. So having a system spider cache file and setting `LMOD_CACHED_LOADS=yes` will not hamper modulefiles created by users in personal directories.

While building the spider cache, each modulefile is evaluated for changes to `MODULEPATH`. Any directories added to `MODULEPATH` are also walked. This means if your site uses the software hierarchy then the new directories added by compiler or mpi stack modulefiles will also be searched.

Sites running Lmod have three choices:

1. Do not create a spider cache for system modules. This will work fine as long as the number of modules is not too large. You will know when it is time to start building a cache file when you start getting complains how long it takes to do any module commands.
2. If you have a formal procedure for installing packages on your system, then I recommend you to do the following. Have the install procedure run the `update_lmod_system_cache_files` script. This will create a file called “`system.txt`”, which marks the time that the system was last updated, so that Lmod knows that the cache is still good.
3. Or you can run the `update_lmod_system_cache_files` script say every 30 minutes. This way the cache file is up-to-date. No new module will be unknown for more than 30 minutes.

There are two ways to specify how cache directories and timestamp files are specified. You can use “`--with-spiderCacheDir=dirs`” and “`--with-updateSystemFn=file`” to specify one or more directories with a single timestamp file:

```
./configure --with-spiderCacheDir=/opt/mData/cacheDir --with-updateSystemFn=/opt/
↳mdata/system.txt
```

If you have multiple directories each with their own timestamp file, you can list those in a file that configure will read rather than enumerating them with `--with-spiderCacheDescript=file`. This also enables each cache directory to have its own timestamp. The file is only used at configure time, not when Lmod runs, and is used like:

```
cacheDir1:timestamp1
cacheDir2:timestamp2
```

Lines starting with ‘#’ and blank lines are ignored. It is best if each cache directory has its own timestamp file. This file is used by configure to modify the `$LMOD_DIR/init/lmodrc.lua` file. See the [An Example Setup](#) for a complete example.

5.11.1 How to decide how many system cache directories to have

The answer to this question depends on which machines “owns” which modulefiles. Many sites have a single location where their modulefiles are stored. In this case a single system cache file is all that is required.

At TACC, we need two system cache files because we have two different locations of files: one in the shared location and one on a local disk. So in our case Lmod sees two cache directories. Each node builds a spider cache of the modulefiles it “owns” and a single node (we call it master) builds a cache for the shared location.

5.11.2 What directories to specify?

If your site doesn’t use the software hierarchy, (see [How to use a Software Module hierarchy](#) for more details) then just use all the directory specified in **MODULEPATH**. If you do use the hierarchy, then just specify the “Core” directories, i.e. the directories that are used to initialize Lmod but not the compiler dependent or mpi-compiler dependent directories.

5.11.3 How to test the Spider Cache Generation and Usage

In a couple of steps you can generate a personal spider cache and get the installed copy of Lmod to use it. The first step would be to load the lmod module and then run the **update_lmod_system_cache_files** program and place the cache in the directory `~/moduleData/cacheDir` and the time stamp file in `~/moduleData/system.txt`:

```
$ module load lmod
$ update_lmod_system_cache_files -d ~/moduleData/cacheDir -t ~/moduleData/system.txt
↪ $MODULEPATH
```

If you using Lmod 6 then replace **MODULEPATH** with **LMOD_DEFAULT_MODULEPATH** instead.

Next you need to find your site’s copy of `lmodrc.lua`. This can be found by running:

```
$ module --config
...

Active RC file(s):
-----
/opt/apps/lmod/6.0.14/init/lmodrc.lua
```

It is likely your site will have it in a different location. Please copy that file to `~/lmodrc.lua`. Then change the bottom of the file to be:

```
scDescriptT = {
  {
    ["dir"]      = "/path/to/moduleData/cacheDir",
    ["timestamp"] = "/path/to/moduleData/system.txt",
  },
}
```

where you have changed */path/to* to match your home directory. Now set:

```
$ export LMOD_RC=$HOME/lmodrc.lua
```

Then you can check to see that it works by running:

```
$ module --config
...

Cache Directory          Time Stamp File
-----
$HOME/moduleData/cacheDir  $HOME/moduleData/system.txt
```

Where **\$HOME** is replaced by your real home directory. Now you can test that it works by doing:

```
$ module avail
```

The above command should be much faster than running without the cache:

```
$ module --ignore_cache avail
```

5.11.4 An Example Setup

Suppose that your site has three different modulefile trees. This can be handle in two very different ways. If each tree is on the same computer you can have one spider cache that knows about all three.

Assuming that the tree modulefile trees are named:

```
/sw/ab/modulefiles
/sw/cd/modulefiles
/sw/ef/modulefiles
```

If all tree directory trees are owned by same computer then one can configure Lmod with:

```
$ ./configure --with-spiderCacheDir=/sw/mData/cacheDir --with-updateSystemFn=/sw/
↪mData/cacheTS.txt
```

And build the cache file with:

```
$ export MODULEPATH=/sw/ab/modulefiles:/sw/cd/modulefiles:/sw/ef/modulefiles
$ update_lmod_system_cache_files -d /sw/mData/cacheDir -t /sw/mData/cacheTS.txt
↪$MODULEPATH
```

Now suppose you have the same three module directories but they reside on three different computers or are managed by three different groups. If you have three different groups managing a different module directory tree, you'll obviously want each group to manage each module tree separately.

Many sites place all their module based software on a shared disk across all nodes. Other sites might store some software locally on a node and some in a shared location. It is this scenario which requires some care when generating the spider caches.

So for any number of reasons you might have to have multiple spider cache files. In this case your site would configure Lmod with a spider cache description file (call say: spiderCacheDescript.txt) that contains:

```
/sw/ab/mData/cacheDir:/sw/ab/mData/cacheTS.txt
/sw/cd/mData/cacheDir:/sw/cd/mData/cacheTS.txt
/sw/ef/mData/cacheDir:/sw/ef/mData/cacheTS.txt
```

Next Lmod is configured with this spiderCacheDescript.txt file, which is only used to configure Lmod.:

```
$ ./configure --with-spiderCacheDescript=/path/to/spiderCacheDescript.txt
```

The configure script modifies the \$LMOD_DIR/init/lmodrc.lua file so that the lmod command knows about the caches. The spiderCacheDescript.txt is never used again. Here is what the bottom of the lmodrc.lua would look like:

```
...
scDescriptT = {
  {
    ["dir"]      = "/sw/ab/mData/cacheDir",
    ["timestamp"] = "/sw/ab/mData/cacheTS.txt",
  },
  {
    ["dir"]      = "/sw/cd/mData/cacheDir",
    ["timestamp"] = "/sw/cd/mData/cacheTS.txt",
  },
  {
    ["dir"]      = "/sw/ef/mData/cacheDir",
    ["timestamp"] = "/sw/ef/mData/cacheTS.txt",
  },
}
```

Scenario 1: Three groups managing a separate module tree

Here we are assuming that all software resides on a shared but there are three group each managing a module tree.

So the “ab” group builds their spider cache as follows:

```
$ update_lmod_system_cache_files -d /sw/ab/mData/cacheDir -t /sw/ab/mData/cacheTS.txt
↪ /sw/ab/modulefiles
```

Similar the “cd” group builds their spider cache by:

```
$ update_lmod_system_cache_files -d /sw/cd/mData/cacheDir -t /sw/cd/mData/cacheTS.txt
↪ /sw/cd/modulefiles
```

and so on for each group managing their module tree. Each group has to update their spider cache if they update their module tree. If the “ab” group add new software and new modulefiles. They must update their cache file, but other groups do not have to update their caches if everything has remained the same for their modules

Scenario 2: Different computers owning different module trees

Suppose that the master node controls the directories /sw/ab/... and the /sw/cd/... on a shared disk. Then on the master node, one runs:

```

master$ update_lmod_system_cache_files -d /sw/ab/mData/cacheDir -t /sw/ab/mData/
↪cacheTS.txt /sw/ab/modulefiles
master$ update_lmod_system_cache_files -d /sw/cd/mData/cacheDir -t /sw/cd/mData/
↪cacheTS.txt /sw/cd/modulefiles

```

Then on each local node has a replicated copy of **/sw/ef/...** on a local disk. So each node has to run:

```

$ update_lmod_system_cache_files -d /sw/ef/mData/cacheDir -t /sw/ef/mData/cacheTS.txt ↪
↪ /sw/ef/modulefiles

```

Again if any new modulefiles are added or changed, then the appropriate caches must be updated.

5.12 Using the module spider command

The module spider command reports all the modules that can be loaded on a system. In a flat module layout system, the *module avail* and *module spider* return similar information. In a hierarchical system, *module spider* returns all the modules that are possible where as *module avail* only reports modules that can be loaded directly.

There are three modes to module spider

1. module spider: Report all modules, known as level 0.
2. module spider <name> : report all the versions for the modules that match <name>. This is known as level 1
3. module spider <name/version>: Report detailed information on a particular module version. This is known as level 2

5.12.1 Level 0: module spider

The *module spider* command by itself lists all the modules that can be found. If your site is running a flat layout, then the information displayed will be similar to what *module avail* reports. A typical output looks like:

```

$ module spider
-----
The following is a list of the modules currently available:
-----

autotools: autotools/1.2
  Autoconf, automake, libtool

boost: boost/1.54.0, boost/1.55.0, boost/1.56.0, boost/1.61.0, boost/1.62.0
  Boost provides free peer-reviewed portable C++ source libraries.

fftw2: fftw2/2.1.5
  Numerical library, contains discrete Fourier transformation

gcc: gcc/4.7.3, gcc/4.8.1, gcc/5.3.1, gcc/6.2.0
  The Gnu Compiler Collection

hdf5: hdf5/1.8.12, hdf5/1.8.13, hdf5/1.8.14
  General purpose library and file format for storing scientific data (Serial ↪
↪Version)

lmod: lmod/7.5.9
  Lmod: An Environment Module System

```

(continues on next page)

(continued from previous page)

```
mpich: mpich/3.0.4-dbg, mpich/3.0.4, mpich/3.1.1-dbg, mpich/3.1.1, mpich/3.1.2-dbg, ↵
↵...
    High-Performance Portable MPI

openmpi: openmpi/1.8.2, openmpi/1.10.3, openmpi/2.0.1
    Openmpi Version of the Message Passing Interface Library

phdf5: phdf5/1.8.12, phdf5/1.8.13, phdf5/1.8.14
    General purpose library and file format for storing scientific data (parallel I/O ↵
↵version)
```

This output shows the name of the module followed by a list of the fullnames of the modules. If there are more modules then can fit on one line, the the list is truncated. Below is the description. Lmod looks for a particular whatis command in the modulefile. For example, the autotools module has a whatis function call that looks like the followin in a lua modulefile.:

```
whatis("Description: Autoconf, automake, libtool")
```

In a TCL modulefile it would look like:

```
module-whatis "Description: Autoconf, automake, libtool"
```

If your output of *module spider* doesn't have a description, please ask your site to consider adding the an appropriate whatis line in your modulefiles.

5.12.2 Level 1: module spider name1 name2 ...

You can get more detail about particular modules by adding one or more names to the command to get a level 1 command. For example:

```
$ module spider hdf5

-----
hdf5:
-----

Description:
  General purpose library and file format for storing scientific data
  (Serial Version)

Versions:
  hdf5/1.8.12
  hdf5/1.8.13
  hdf5/1.8.14
Other possible modules matches:
  phdf5
```

Since the name “hdf5” matches a module name, Lmod only report on the hdf5 module and not the phdf5 module. It does report that other matches are possible (such as phdf5). The reason for this is some sites name the R stat package as R. This rule is to prevent getting every module that has an ‘r’ in it. Note that the searching for modules is case insensitive. So *module spider openmpi* would match a module named *OpenMPI*.

If you search a name that only partially matches a module name then Lmod reports all matches:

```
$ module spider df5
```

(continues on next page)

(continued from previous page)

```
-----
hdf5:
```

```
-----
Description:
```

```
General purpose library and file format for storing scientific data
(Serial Version)
```

```
-----
Versions:
```

```
hdf5/1.8.12
hdf5/1.8.13
hdf5/1.8.14
-----
```

```
-----
phdf5:
```

```
-----
Description:
```

```
General purpose library and file format for storing scientific data
(parallel I/O version)
```

```
-----
Versions:
```

```
phdf5/1.8.12
phdf5/1.8.13
phdf5/1.8.14
-----
```

Finally, you can perform regular expression matches with:

```
$ module -r spider '.*hdf5.*'
```

5.12.3 Level 2: module spider name/version

The level 2 output provides a detailed report on a module:

```
$ module spider phdf5/1.8.14
```

```
-----
phdf5: phdf5/1.8.14
-----
```

```
Description:
```

```
General purpose library and file format for storing scientific
data (parallel I/O version)
```

```
You will need to load all module(s) on any one of the lines below
before the "phdf5/1.8.14" module is available to load.
```

```
gcc/4.8.1 mpich/3.1.1
gcc/4.8.1 mpich/3.1.2
gcc/4.8.1 openmpi/1.8.2
```

```
Help:
```

```
The HDF5 module defines the following environment variables:
...
```

5.13 The spider tool

Lmod provides a tool called *spider* to build the spider cache and other files that can help a site manage their modulefiles. The shell script *update_lmod_system_cache_files* described in [System Spider Cache](#) uses the spider command to build the spiderT.lua file which is the filename of the spider cache. The command *spider* is different from the *module spider* command. The first one will be mainly used by system administrators where as the second command is typically for users.

The *spiderT.lua* file contains information about all the modulefiles on your system that can be found from \$MODULEPATH and is built by *spider*. The way this file is built is by evaluating each module and remembering the help message, the *whatis* calls and the location of the modulefile. It also remembers any properties set by the modulefile. It doesn't remember the module commands inside the modulefile.

When evaluating the modulefiles, Lmod looks for changes in \$MODULEPATH. Each new directory added to MODULEPATH is also searched for modulefiles. This is why this command is called "spider".

There are some important points about Lmod walking the module tree. Each modulefile is evaluated separately. This means that any changes in MODULEPATH need to be static. That is new path must be defined in the modulefile. If the change in MODULEPATH is dynamic or depends on some environment variable that is dynamic, it is very likely that Lmod's *module spider* won't know about these dynamic changes to MODULEPATH.

The spiderT.lua file is used by Lmod to know any properties set in a modulefile. This information is used by *module avail*, *module spider* and *module keyword*. By default Lmod doesn't use the spider cache (aka spiderT.lua) to load modulefiles, however a site can configure Lmod to use cache based loads. This just means that a site *must* keep the spider cache up-to-date.

Lmod looks for special *whatis* calls to know what the description for a module is. See [Using the module spider command](#) for details.

5.13.1 Software page generation

Since the package modulefiles are the gold standard of the packages your sites offers, it should be that information which should be used to generate the list of software packages. Lmod provides two kind of output depending on your sites needs. Suppose you have the simple module tree:

```
foo
├── .2.0.lua
├── 1.0.lua
├── 1.1.lua
└── default -> 1.0.lua
```

Then Lmod can produce the following information in json format:

```
$ MODULEPATH=.
$ spider -o jsonSoftwarePage $MODULEPATH | python -mjson.tool
[
  {
    "defaultVersionName": "1.0",
    "description": "foo description",
    "package": "foo",
    "versions": [
      {
        "canonicalVersionString": "000000001.*zfinal",
        "full": "foo/.2.0",
        "help": "foo v.2.0",
        "hidden": true,
```

(continues on next page)

(continued from previous page)

```

        "markedDefault": false,
        "path": "foo/.2.0.lua",
        "versionName": "1.0",
        "wV": "000000000.000000002.*zfinal"
    },
    {
        "canonicalVersionString": "000000001.000000001.*zfinal",
        "full": "foo/1.1",
        "help": "foo v1.1",
        "markedDefault": false,
        "path": "foo/1.1.lua",
        "versionName": "1.1",
        "wV": "000000001.000000001.*zfinal"
    },
    {
        "canonicalVersionString": "000000001.*zfinal",
        "full": "foo/1.0",
        "help": "foo v1.0",
        "markedDefault": true,
        "path": "foo/1.0.lua",
        "versionName": "1.0",
        "wV": "^000000001.*zfinal"
    }
]
}
]

```

The versions array block is sorted by the “wV” fields. This a weighted version of the canonicalVersionString, where the only difference is that the first character in the string is modified to know that it is marked default. Also if a module is hidden then the “hidden” field will be set to true.

The last entry in the versions array is used to set the description.

Another json output may be of interest. There is more information but it will be up to the site build the summarization that jsonSoftwarePage provides:

```

$ MODULEPATH=.
$ spider -o spider-json $MODULEPATH | python -mjson.tool
{
  "foo": {
    "foo/.2.0.lua": {
      "Description": "foo description",
      "Version": "1.0",
      "fullName": "foo/.2.0",
      "help": "foo v.2.0",
      "hidden": true,
      "pV": "000000000.000000002.*zfinal",
      "wV": "000000000.000000002.*zfinal",
      "whatis": [
        "Description: foo description",
        "Version: .2.0",
        "Categories: foo"
      ]
    },
    "foo/1.0.lua": {
      "Description": "foo description",
      "Version": "1.0",

```

(continues on next page)

(continued from previous page)

```

    "fullName": "foo/1.0",
    "help": "foo v1.0",
    "hidden": false,
    "pV": "000000001.*zfinal",
    "wV": "^000000001.*zfinal",
    "whatis": [
        "Description: foo description",
        "Version: 1.0",
        "Categories: foo"
    ]
},
"foo/1.1.lua": {
    "Description": "foo description",
    "Version": "1.1",
    "fullName": "foo/1.1",
    "help": "foo v1.1",
    "hidden": false,
    "pV": "000000001.000000001.*zfinal",
    "wV": "000000001.000000001.*zfinal",
    "whatis": [
        "Description: foo description",
        "Version: 1.1",
        "Categories: foo"
    ]
}
}
}

```

5.14 Deprecating Modules

There may come a time when your site might want to mark a module for deprecation. If you track module usage, you can find the modules that are rarely used, and you can find out which users are using the modules. Once you have decided which modules are marked for removal, you can make a message be printed when the module is loaded.

Suppose you want to mark a module for later removal. You can setup a message that is reported to the user *Every time the module is loaded*.

For example:

```

$ module load abc/1.2.3

abc/1.2.3:
  This module is deprecated and will be removed from the system
  on June 19.  Please load abc/2.3.4 instead.

```

Note that this message is just text and in no way controls user access to the module. Your site will have to remove the module. This nag message is a way to let your users know that removal will happen ahead of time.

Also note that a user only gets this message when loading the module. There is no special marking in `module avail` or `module spider`.

You can create a file called “admin.list” and place it in “/path/to/lmod/etc/admin.list”. Note that typically the lmod script will be in “/path/to/lmod/lmod/libexec/lmod”. The etc directory is independent to the version of Lmod. You can see the location that Lmod is looking for by executing:

```
$ module --config
```

Look for “Admin File”. You can also set the “LMOD_ADMIN_FILE” to point to the admin.list file.

The admin file consists of key-value pairs. For example:

```
moduleName/version:  message
<blank line>
```

Or:

```
Full_PATH_to_Modulefile: message
<blank line>
```

Or:

```
foo/1.1 | bar/1.2 | Full_Path_to_Modulefile:
message
<blank line>
```

Lmod treats assume that if a pattern starts with a ‘/’ then it is a full path to a modulefile. If it doesn’t then it is a moduleName/version string. Also, you can have several modulefiles use the same message by separating them with |

You can use Lua regular expression to also match one or modules for both the full path to a module or the module fullname. Lua regular expressions are much less powerful (see <http://lua-users.org/wiki/PatternsTutorial> for more details). In particular they do not support patterns like:

```
(foo|bar)
```

Lmod searches the list of modules and/or paths from top to bottom and it uses the first match it finds. So you might want to place the module patterns from specific to general in the admin file to avoid incorrect matching.

Also some care may be necessary for ‘.’ as it matches every character so you might have to change the ‘.’ to ‘%.’ to match the actual period. Also ‘-’ is a regular expression character so matching module names that have a ‘-’ in their name requires % quoting.

The message can be as many lines as you like. The message ends with a blank line. Below is an example:

```
gcc/2.95:    This module is deprecated and will be removed from the system on Jan 1.↵
↪1999.
           Please switch to a newer compiler.

boost/1%.54%.0:
We are having issues

/opt/apps/modulefiles/Compiler/gcc/4.7.2/boost/1.55.0:
We are having issues

boost/1%.[5-7].*:
    We are having more issues.

boost%-mpi/1%.[5-7].*:
    This module will be removed.
```

Note that you don’t include the .lua part when specifying the version number.

5.15 Module Properties

Lmod support giving modules properties. For modules written Lua, the `add_property()` function looks like:

```
add_property("key", "value")
```

In TCL, it is written as:

```
add-property key value
```

The *key* and *value* are controlled by a file called `lmodrc.lua`.

5.15.1 The Properties File: `lmodrc.lua`

Lmod provides a standard `lmodrc.lua` which is copied to the installation directory. For example Lmod version X.Y.Z is installed in `/apps/lmod/X.Y.Z` then `lmodrc.lua` would be installed in `/apps/lmod/X.Y.Z/init/lmodrc.lua`. During the install process this file is modified to include the location of the system spider cache.

Lmod searches the properties in several location in order given below. Assuming again that Lmod is installed in `/apps/lmod/X.Y.Z` then Lmod searches for the property information in the following order:

1. `/apps/lmod/X.Y.Z/init/lmodrc.lua`
2. `/apps/lmod/etc/lmodrc.lua`
3. `/etc/lmodrc.lua`
4. `$HOME/.lmodrc.lua`
5. `$LMOD_RC`

Where `$LMOD_RC` is an environment variable that can be set to point to any file location. If there are more than one of these files exist then they are merged and not a replacement. So a site can (and should) leave the first file as is and create another one to specify site properties and Lmod will merge the information into one.

The format of this file looks like:

```
local i18n = require("i18n")
propT = {
  arch = {
    validT = { mic = 1, offload = 1, gpu = 1, },
    displayT = {
      ["mic:offload"]      = { short = "(*)", color = "blue", full_color = false,
↪doc = "built for host, native MIC and offload to the MIC", },
      ["mic"]              = { short = "(m)", color = "blue", full_color = false,
↪doc = "built for host and native MIC", },
      ["offload"]          = { short = "(o)", color = "blue", full_color = false,
↪doc = "built for offload to the MIC only", },
      ["gpu"]              = { short = "(g)", color = "red", full_color = false,
↪doc = "built for GPU", },
      ["gpu:mic"]          = { short = "(gm)", color = "red", full_color = false,
↪doc = "built natively for MIC and GPU", },
      ["gpu:mic:offload"] = { short = "(@)", color = "red", full_color = false,
↪doc = "built natively for MIC and GPU and offload to the MIC", },
    },
  },
  state = {
    validT = { experimental = 1, testing = 1, obsolete = 1 },
```

(continues on next page)

(continued from previous page)

```

    displayT = {
      experimental = { short = "(E)", full_color = false, color = "blue", doc =
↪ i18n("ExplM"), },
      testing      = { short = "(T)", full_color = false, color = "green", doc =
↪ i18n("TstM"), },
      obsolete     = { short = "(O)", full_color = false, color = "red", doc =
↪ i18n("ObsM"), },
    },
  },
  lmod = {
    validT = { sticky = 1 },
    displayT = {
      sticky = { short = "(S)", color = "red", doc = i18n("StickyM"), }
    },
  },
  status = {
    validT = { active = 1, },
    displayT = {
      active = { short = "(L)", color = "yellow", doc = i18n("LoadedM") },
    },
  },
}

```

This file defines a table called `propT`. A table is a generic name for a hash table or dictionary or associative array. That is, it stores key value pairs. It is an Lmod convention that a table is named with a trailing `T` to remind us that it is a table.

In this case `propT` defines the valid keys and values that are possible for a modulefile to use with `add_property()`. In the case of the above table, the only valid keys in a modulefile would be `arch`, `state`, `lmod` and `status`.

The value for `state` controls the valid values. In particular, the only valid values for `state` are `experimental`, `testing` or `obsolete`. Please note that a modulefile can have multiple properties but each property key can have only one value. So:

```

add_property("state", "testing")
add_property("state", "obsolete")

```

would make the `state` property have a value of `obsolete`. On the other hand a modulefile could have two or more properties.:

```

add_property("state", "testing")
add_property("lmod", "sticky")

```

Lmod itself depends on the keys `lmod` and `status`. So as a site, it is expected that any `lmodrc.lua` file will contain these properties.

5.15.2 The tables *validT* and *displayT*

The function `add_property()` expects a key and value. So for the `state` key, possible values are `experimental`, `testing` or `obsolete`. Those strings must appear in two tables: the `validT` and the `displayT` tables. For example, we can see that `testing` appears both in the `validT` and `displayT` tables. This exists for checking for valid values when the `add_property()` function is called from modulefiles.

The `displayT` table controls how the property is displayed. The fields in the table control how a property is displayed. For example:

```
testing = { short = "(T)", full_color = false, color = "green", doc = i18n("TstM"), }  
↪,
```

says that a module with this property will have a ‘(T)’ next to its name when printed by module avail. If the terminal display has “xterm” as part of the environment variable TERM. then the ‘T’ will be in green. If the field *full_color* is set to *true* then the name and ‘T’ will be in green.

The possible color values are: *black*, *red*, *green*, *yellow*, *blue*, *magenta*, *cyan*, and *white*. In practice, since users can use light letters on dark backgrounds or dark letters on light backgrounds, sites may wish to avoid *black*, *white* and possibly *yellow*.

The arch key shows that the values can be combined. If the value is colon separated then each string between the colons have to be valid keys.

5.16 Debugging Modulefiles

Most modulefiles are simple combination of a `help()` message, a couple of `setenv()` and a `prepend_path()` or two and don’t require much in the way of debugging. However modulefiles are a program and might need debugging.

5.16.1 TCL modulefiles

As was discuss in *How does Lmod convert TCL modulefile into Lua*, Lmod converts TCL modulefiles into a lua modulefile by executing normal tcl commands and translates TCL module commands into lua functions. To see what Lmod does with your TCL modulefile, you can run `tcl2lua.tcl` to see the translation:

```
$ $LMOD_DIR/tcl2lua.tcl <path_to_modulefile>
```

For example, suppose you have a TCL modulefile in `~/my_modules/foo/1.0`:

```
#!/Module  
  
global env  
set home $env(HOME)  
set pkg "$home/foo"  
prepend-path PATH $pkg/bin  
setenv FOO_DIR $pkg
```

Then running the command produces:

```
$ $LMOD_DIR/tcl2lua.tcl ~/my_modules/foo/1.0  
  
prepend_path{"PATH","/home/user/foo/bin",delim=":",priority="0"}  
setenv("FOO_DIR","/home/user/foo")
```

5.16.2 Lua Modulefiles

It is important to remember that Lmod uses a two part process to change the your environment. The lmod program produces text that is appropriate for the shell choice: bash commands for bash shell; csh commands for C-shell and so on. Then that text is evaluated by the shell to change your environment.

We can take advantage of this two part process to debug modulefiles by getting Lmod to produce the commands but not evaluate them. So starting with a simple lua modulefile called `~/my_modules/foo/1.0.lua`:


```

local home = os.getenv("HOME")
local pkg  = pathJoin(home, "foo")
io.stderr:write("home: ", home, "\n")
io.stderr:write("pkg:   ", pkg, "\n")

prepend_path("PATH", pathJoin(pkg, "bin"))
setenv("FOO", pathJoin(pkg, "bin"))

```

You can see that the above modulefile two extra print debugging statements that you'll want to remove after debugging. Running the Lmod command produces:

```

$ module use ~/my_modules
$ $LMOD_CMD bash load foo/1.0

home: /home/user
pkg:   /home/user/foo
FOO="/home/user/foo/bin";
export FOO;
PATH="/home/user/foo/bin:..."
export PATH;
...

```

Actually the lmod command will produce much more text. It contains other environment variables such as:

```

LOADEDMODULES="...";
export LOADEDMODULES;
__LMFILES__="...";
export __LMFILES__;
MODULEPATH="...";
export MODULEPATH;
__LMOD_REF_COUNT_PATH="/home/user/foo/bin:1;..."
export __LMOD_REF_COUNT_PATH;
PATH="/home/user/foo/bin:..."
export PATH;
__ModuleTable001__="...";
export __ModuleTable001__;
__ModuleTable_Sz__="6";
export __ModuleTable_Sz__;

```

LOADEDMODULES and __LMFILES__ are the list of modules loaded and their locations. These variables is made available to be compatible with Tmod and can be used by modulefiles. Variables like __LMOD_REF_COUNT_PATH are used to support reference counting for path like variables. Finally, Lmod uses a lua table called the __ModuleTable__ which contains the information used between Lmod invocations. This table is base64 encoded and split into 256 character blocks and stored in \$_ModuleTable001, \$_ModuleTable002, ...

5.17 SitePackage.lua and hooks

Sites may wish to alter the behavior of Lmod to suit their needs. A good place to do this is the `SitePackage.lua`. Anything in this file will automatically be loaded every time the Lmod command is run. This file can be used to provide common functions that can be used in a sites modulefiles. It is also a place where a site can implement their hook functions.

Hook functions are normal functions that if implemented and registered with Lmod will be called when certain operations happen inside Lmod. For example, there is a load hook. A site can register a function that is called every time a module is loaded. There are several hook functions that are discussed in [Hook functions](#).

5.17.1 How to set up SitePackage.lua

Here are two suggestions on how to use your SitePackage.lua file:

1. Install Lmod normally and then overwrite your SitePackage.lua file over this one in the install directory. This is good for testing but it is recommended that sites use the second method for permanent changes. Otherwise changes may be lost when upgrading Lmod.
2. Create a file named “SitePackage.lua” in a different directory separate from the Lmod installed directory and it will override the one in the Lmod install directory. Then you should modify your z00_lmod.sh and z00_lmod.csh (or however you initialize the “module” command) with:

```
(for bash, zsh)
export LMOD_PACKAGE_PATH=/path/to/the/Site/Directory

(for csh)
setenv LMOD_PACKAGE_PATH /path/to/the/Site/Directory
```

5.17.2 Implementing functions in SitePackage.lua

For example your site might wish to provide the following function to set MODULEPATH inside your SitePackage:

```
function prependModulePath(subdir)
    local mroot = os.getenv("MODULEPATH_ROOT")
    local mdir = pathJoin(mroot, subdir)
    prepend_path("MODULEPATH", mdir)
end
```

This function must be registered with the sandbox so that Lmod modulefiles can call it:

```
sandbox_registration{ prependModulePath    = prependModulePath }
```

5.17.3 Hook tips

If you want to know which modules are being loaded, use the FrameStack:

```
local frameStk = FrameStk:singleton()
-- check if anything else gets loaded
if not frameStk:empty() then
    print("This modules has no dependencies loaded.")
end

-- the module at the top of the FrameStack is the one the user requested
local userload = (frameStk:atTop()) and "yes" or "no"
```

If you want to know the short name or path of a loaded module, you can use the ModuleTable:

```
local mname    = MName:new("mt", FullModuleName)
local sn       = mname:sn()
local version  = mname:version()
```

5.17.4 Hook functions

load (...): This function is called after a modulefile is loaded in “load” mode.

unload (...): This function is called after a modulefile is unloaded in “unload” mode.

parse_updateFn (...): This hook returns the time on the timestamp file.

writeCache (...): This hook return whether a cache should be written.

SiteName (...): This hook is used to specify Site Name. It is used to generate family prefix:
`site_FAMILY_COMPILER site_FAMILY_MPI ...`

msgHook (...): Hook to print messages after avail, list, spider.

errWarnMsgHook (...): Hook to print messages after LmodError, LmodWarning, LmodMessage.

groupName (...): This hook adds the arch and os name to moduleT.lua to make it safe on shared filesystems.

avail (...): Map directory names to labels

restore (...): This hook is run after restore operation

startup (UsrCmd): This hook is run when Lmod is called

packagebasename (s_patDir, s_patLib): This hook gives you a table with the current patterns that spider uses to construct the reverse map.

load_spider (...): This hook is called when spider is evaluating a modulefile.

isVisibleHook (modT): This hook is called when evaluating whether a module is visible or hidden. It’s argument is a table containing: fullName, sn (short name), fn (file path) and isVisible (boolean) of the module.

5.17.5 Example Hook: msgHook

A site might like to control the output of list, avail, and spider commands by adding text to the beginning or end of the generated text.

Here is an example of how to use the msgHook. So inside a site’s SitePackage.lua file one would do:

```
local hook = require("Hook")

function myMsgHook(kind,a)
  if (kind == "avail") then

    -- Here is text that would go at the top of avail:
    table.insert(a,1,"This system has ...\n")
    table.insert(a,2,"blah blah blah ...\n")
    table.insert(a,3,"more blah blah blah ...\n")

    -- Here is text that would go at the end of avail:
    a[#a+1] = "More blah blah ...\n"
    a[#a+1] = "yet more blah blah ...\n"
  elseif (kind == "list") then
    ...
  elseif (kind == "spider") then
    ...
  end
  return a
end

hook.register("msgHook", myMsgHook)
```

As you can see you can add text to the beginning and/or the end of the text that is generated by avail, spider and list.

5.18 Lmod Localization

Lmod places all messages, warning and error messages into a file. It also uses a slightly modified version of the `i18n` package from:

<https://github.com/kikito/i18n.lua>

to allow for translate into multiple languages. The standard message file is in `messageDir/en.lua`. Other languages are such as French, Spanish, German, and Mandarin Chinese are available. You can submit pull requests if you wish to have other languages. Lmod uses the `LANG` variable to pick which language to use. Note that Lmod only uses the first part of the `LANG` variable to choose which language to use. So there is only one version of English and one version of Spanish and so on. So a typical value of `LANG` is:

```
LANG=en_US.UTF-8
```

Lmod removes the underscore and just uses “en”. A site can also configure Lmod or set `LMOD_LANG` to override `$LANG`.

5.18.1 Site Tailoring

It is also possible to modify the standard messages to be tailored for a site. You can leave the Language to be English but change the value of a particular message to better match your site. For example you might change message `e118` (as seen in `messageDir/en.lua`). Please note that the message string should be on one line.:

```
e118 = "User module collection: \"${collection}\" does not exist.\n Try \"module_\nsavelist\" for possible choices.\n",
```

by creating a file called: `/path/to/site_msgs.lua`:

```
return {  
  site = {  
    e118 = "User module collection: \"${collection}\" does not exist.\n Try_\n\"module savelist\" for possible choices. For questions see https://...\n";  
  }  
}
```

Then configure Lmod (or set env var. `LMOD_SITE_MSG_FILE`) to use `/path/to/site_msg.lua`.

5.19 Integration of EasyBuild and Lmod

EB & Lmod work great together.

There are variables you should look into for tuning your system optimally, such as `EASYBUILD_MODULES_TOOL=Lmod`, `LMOD_PACKAGE_PATH` (keep track of loads), `LMOD_SYSTEM_DEFAULT_MODULES` (choose buildsets/stages etc)

Here is an example of how to ensure that your users can choose to have (or not) the default directory of EasyBuild under home (`$HOME/.local/easybuild/modules/all`) in their `$MODULEPATH`. Conveniently, by going this way users can save/restore environments via Lmod, at will.

Features obtained:

1. user modulefiles generated via EasyBuild are easy to enable - less typing, too
2. `$MODULEPATH` is managed via modulefiles in itself
3. technique is compatible with `$LMOD_SYSTEM_DEFAULT_MODULES` feature

4. `ml save/restore` work as desired

Caveats:

- This feature is super-picky on broken module environments, fi: https://bugzilla.redhat.com/show_bug.cgi?id=1326075 (testable via <https://github.com/hpcugent/easybuild-framework/issues/1756>) . Fix that asap, or your modulefiles experience may not be as good as it can!

5.19.1 How to put into use

The two examples are visible below and are ready to be copied in respective filenames `use.own.eb/append` & `use.own.eb/prepend`. Simply place both files in a directory of the existing `$MODULEPATH` and decide if you with to *append* or *prepend*:

```
$ ml av use.own

----- /etc/site/modules -----
use.own.eb/append      use.own.eb/prepend (D)

----- /cm/local/modulefiles -----
use.own

Where:
D:  Default Module

Use "module spider" to find all possible modules.
Use "module keyword key1 key2 ..." to search for all
possible modules matching any of the "keys".

$ ml use.own.eb/append
$ echo $MODULEPATH
/etc/site/modules:/cm/local/modulefiles:/home/user/.local/easybuild/modules/all
```

5.19.2 Example use.own.eb/append

```
$ cat /etc/site/modules/use.own.eb/append
##Module1.0#####
##
## use.own.eb modulefile
##
## modulefiles/use.own.eb.  Generated by fgeorgatos
##
proc ModulesHelp { } {

    puts stderr "\tThis module file will add \${HOME}/.local/easybuild/modules/all to_
↳the"
    puts stderr "\tlist of directories that the module command will search"
    puts stderr "\tfor modules.  EasyBuild places your own modulefiles here."
    puts stderr "\tThis module, when loaded, will create this directory if_
↳necessary."
    puts stderr "\n\n"
}

module-whatis    "adds your EasyBuild modulefiles directory to MODULEPATH"
```

(continues on next page)

(continued from previous page)

```
eval set  [ array get env HOME ]
set      ownmoddir      $HOME/.local/easybuild/modules/all

# create directory if necessary
if [ module-info mode load ] {
    if { ! [ file exists $ownmoddir ] } {
        file mkdir $ownmoddir
    }
}

module use --append $ownmoddir
```

5.19.3 Example use.own.eb/prepend

```
$ cat /etc/site/modules/use.own.eb/prepend
##Module1.0#####
##
## use.own.eb modulefile
##
## modulefiles/use.own.eb.  Generated by fgeorgatos
##
proc ModulesHelp { } {
    puts stderr "\tThis module file will add \${HOME}/.local/easybuild/modules/all to
↳the"
    puts stderr "\tlist of directories that the module command will search"
    puts stderr "\tfor modules.  EasyBuild places your own modulefiles here."
    puts stderr "\tThis module, when loaded, will create this directory if
↳necessary."
}

module-whatism  "adds your EasyBuild modulefiles directory to MODULEPATH"

eval set  [ array get env HOME ]
set      ownmoddir      $HOME/.local/easybuild/modules/all

# create directory if necessary
if [ module-info mode load ] {
    if { ! [ file exists $ownmoddir ] } {
        file mkdir $ownmoddir
    }
}

module use --prepend $ownmoddir
```

5.20 Providing Custom Labels for Avail

Lmod writes out the modules in alphabetical order for each directory in MODULEPATH in order:

```
$ module avail

----- /opt/apps/modulefiles/Common -----
```

(continues on next page)

(continued from previous page)

```

abc/8.1    def/11.1    ghi/2.3

----- /opt/apps/modulefiles/Core -----
xyz/8.1    xyz/11.1 (D)

----- /opt/apps/modulefiles/Compilers -----
gcc/6.3    intel/17.0

```

Sites can replace the directory paths with any label they like. This is implemented by adding a SitePackage.lua file and calling the avail hook. See [SitePackage.lua and hooks](#) for how to create SitePackage.lua.

Suppose you wish to merge the Common and Core sections above into a single group named “Core Modules” and change the directory to “Compiler Modules”. The result would be:

```

$ module avail

----- Core Modules -----
abc/8.1    def/11.1    ghi/2.3    xyz/8.1    xyz/11.1 (D)

----- Compiler Modules -----
gcc/6.3    intel/17.0

```

To make this happen you need to do the following. Create a SitePackage.lua file containing:

```

require("strict")
local hook = require("Hook")

local mapT =
{
  en_grouped = {
    ['/Compilers$'] = "Compilers",
    ['/Core$']      = "Core Modules",
    ['/Common$']    = "Core Modules",
  },
  fr_grouped = {
    ['/Compilers$'] = "Compilateurs",
    ['/Core$']      = "Modules de base",
    ['/Common$']    = "Modules de base",
  },
}

function avail_hook(t)
  local availStyle = masterTbl().availStyle
  local styleT     = mapT[availStyle]
  if (not availStyle or availStyle == "system" or styleT == nil) then
    return
  end

  for k,v in pairs(t) do
    for pat,label in pairs(styleT) do
      if (k:find(pat)) then
        t[k] = label
        break
      end
    end
  end
end

```

(continues on next page)

(continued from previous page)

```
end

hook.register("avail", avail_hook)
```

And then set the LMOD_AVAIL_STYLE variable to be:

```
export LMOD_AVAIL_STYLE="system:<en_grouped>:fr_grouped"
```

The angle brackets define the default which in this case is en_grouped. A user can set:

```
export LMOD_AVAIL_STYLE="fr_grouped"
```

to change to the french labels.

5.21 Load Storms: Long load times or Fails to Load

Sometimes it can take a long time for a module to load. Or Lmod will produce an error that it can't load a module it should. One likely possibility is that you have a `load storm`. That is Lmod is loading and reloading the same modulefiles over and over. This can happen when a module tries to load other modules and other modules do the same. For example A/1.0.lua is:

```
load("B/2.0")
load("C/2.0")
```

The module B/2.0.lua is:

```
load("C/2.0")
load("D/2.0")
```

And module C/2.0.lua is:

```
load("D/2.0")
```

The `load()` function always loads the requested module file even if that modulefile is already loaded.

Lmod can report what is happening. Using the `-D` debug flag it is possible to track what gets loaded:

```
$ module purge
$ module -D load A          2> ~/load_storm.log
$ grep 'MasterControl:.*load(' ~/load_storm.log
```

The results from the `grep` is:

```
MasterControl:load(mA={A}) {
  MasterControl:load(mA={B/2.0}) {
    MasterControl:load(mA={C/2.0}) {
      MasterControl:load(mA={D/2.0}) {
        MasterControl:load(mA={D/2.0}) {
          MasterControl:unload(mA={D}) {
            MasterControl:load(mA={D/2.0}) {
              MasterControl:load(mA={C/2.0}) {
                MasterControl:unload(mA={C}) {
                  MasterControl:unload(mA={D/2.0}) {
```

(continues on next page)

(continued from previous page)

```
MasterControl:load(mA={C/2.0}) {
    MasterControl:load(mA={D/2.0})
}
```

We can see that the `D/2.0.lua` module is loaded 4 times in this example. To avoid this problem, one can reduce the number of loads. In this case, the `B/2.0.lua` module can only load the `C` module as `D` is already being loaded by the `C` module. If this is not practical then placing guards around the load statement can also reduce the number of loads. For example changing `A/1.0.lua` to:

```
if (not isloaded("B/2.0")) then
    load("B/2.0")
end
if (not isloaded("C/2.0")) then
    load("C/2.0")
end
```

and similarly for the other modules will reduce loading of each module to one time. This can be seen by executing the debug module load and greping the results as before:

```
MasterControl:load(mA={A}) {
    MasterControl:load(mA={B/2.0}) {
        MasterControl:load(mA={C/2.0}) {
            MasterControl:load(mA={D/2.0}) {
            }
        }
    }
}
```

The above guard statements won't unload the dependent module, unloading the `A` won't unload `B`, `C`, or `D`. Changing the guard statements to the following will allow for loading and unloading:

```
if (not isloaded("B/2.0") or mode() == "unload") then
    load("B/2.0")
end
if (not isloaded("C/2.0") or mode() == "unload") then
    load("C/2.0")
end
```

5.22 Tracing Lmod

Lmod supports a tracing option to report the major steps that Lmod performs. It can be set by either setting the `LMOD_TRACING` to `yes` or using the `-T` or `--trace` option on the command line. This can be useful to understand what happens a shell startup. By setting `LMOD_TRACING` to `yes` for a particular user you can see something like the following for a user which as a default collection:

```
$ module -T restore
running: module -T restore
Using collection:      /home/user/.lmod.d/default
Setting MODULEPATH to :/opt/apps/modulefiles/Linux:/opt/apps/modulefiles/Core
Loading: unix (fn: /opt/apps/modulefiles/Core/unix/unix.lua)
Loading: gcc (fn: /opt/apps/modulefiles/Darwin/gcc/5.2.lua)
Loading: noweb (fn: /opt/apps/modulefiles/Core/noweb/2.11b.lua)
Loading: StdEnv (fn: /opt/apps/modulefiles/Core/StdEnv.lua)
```

5.22.1 Using Shell Startup Debug with Tracing

Sites may wish to install the shell startup debug package from sourceforce:

```
https://sourceforge.net/projects/shellstartupdebug
```

It tracks the startup actions during shell startup. It can tell you what files in `/etc/profile.d/*` are started. For example:

```
% cat > ~/.init.sh
export SHELL_STARTUP_DEBUG=1
^D
% zsh -l
/etc/zshenv{
  /etc/profile.d/lmod.sh{
    } Time = 0.1875
  /etc/profile.d/z00_lmod.sh{
    } Time = 0.2294
  /etc/profile.d/z01_StdEnv.sh{
    } Time = 0.2889
  } Time = 0.2929
```

One of the tricks of this package is that it guarantees that a `~/.init.sh` is read before any files in `/etc/profile.d/*.sh` are sourced for bash and zsh. The file `~/.init.csh` is used for csh/tcsh. This means that you can also easily track Lmod startup actions:

```
% cat > ~/.init.sh
export LMOD_TRACING=yes
^D

% ssh localhost
...
running: module --initial_load --no_redirect restore
Using collection:      /home/user/.lmod.d/default
Setting MODULEPATH to: /opt/apps/modulefiles/Darwin:/opt/apps/modulefiles/Core
Loading: unix (fn: /opt/apps/modulefiles/Core/unix/unix.lua)
Loading: gcc (fn: /opt/apps/modulefiles/Darwin/gcc/5.2.lua)
Loading: noweb (fn: /opt/apps/modulefiles/Core/noweb/2.11b.lua)
Loading: StdEnv (fn: /opt/apps/modulefiles/Core/StdEnv.lua)
```

This way you can trace Lmod startup without having to edit any files in `/etc/profile.d/*` or the shell startup files.

5.23 Sticky Modules

Sites may wish to mark a module as *sticky*. This means that if the module is loaded then it won't be removed with a normal unload or purge. One possible use of sticky modules is where a site wants to define some environment variables that define what the architecture or operating system that allows users to use their system. The reason to make these values sticky is that the system may be difficult to use without these variables set.

Suppose you have a module named “site” that will be sticky. A lua module would look like:

```
setenv("ARCH","abc")
add_property("lmod","sticky")
```

A TCL module would look like:

```
##Module
setenv ARCH abc
add-property lmod sticky
```

It is the `add_property()` function in Lua or `add-property` command in TCL which makes the module sticky. If the “site” module is loaded then it can be unloaded by either command:

```
$ module --force unload site
```

or:

```
$ module --force purge
```

Since a user can unload a sticky module if they really want to. You may wish to the startup scripts (i.e. `/etc/profile.d/*`) instead of modules to define environment variables that you don’t want users to easily change.

5.23.1 Lmod Implementation

Lmod unloads all requested modules, including the sticky modules. As each module is unloaded it remembers any modules which are sticky. After all modules have been unloaded, Lmod tries to load any sticky modules found from the previous step.

5.24 Modify Lmod behavior with SitePackage.lua

Lmod provides a standard way for sites to modify its behavior. This file is called `SitePackage.lua`. Anything in this file will automatically be loaded every time the Lmod command is run. Here are two suggestions on how to use your `SitePackage.lua` file

1. Install Lmod normally and then overwrite your `SitePackage.lua` file over this one in the install directory.
2. Create a file named “`SitePackage.lua`” in a different directory separate from the Lmod installed directory. Then you should modify your `z01_lmod.sh` and `z01_lmod.csh` (or however you initialize the “module” command) with:

```
(for bash, zsh, etc) export LMOD_PACKAGE_PATH=/path/to/the/Site/Directory
```

```
(for csh) setenv LMOD_PACKAGE_PATH /path/to/the/Site/Directory
```

A “`SitePackage.lua`” in that directory will override the one in the Lmod install directory.

You should check to see that Lmod finds your `SitePackage.lua`. If you do:

```
$ module --config
```

and it reports:

```
Modules based on Lua: Version X.Y.Z 3027-02-05 16:31
  by Robert McLay mclay@tacc.utexas.edu

Description          Value
-----
...
Site Pkg location    standard
```

Then you haven’t set things up correctly.

5.24.1 Possible functions for your SitePackage.lua file

There are two common reason a site might setup this file. The main reason is to specify hook functions. If your site wishes to track module usage, you can use the load hook to have a message sent to syslog. The details are in `contrib/track_module_usage/README`.

The second reason is to provide common functions that all your modulefiles can use. These functions must be registered. To prevent modulefiles from calling arbitrary functions all modulefiles are evaluated in a “sandbox” which limits the functions that are called.

The following SitePackage.lua file is a simple example of how to implement a function. In this case, we are going to provide a simple prepend to the MODULEPATH that can be called by any modulefile. Suppose your site is wants to use MODULEPATH_ROOT as the top level directory and all modulepath entry are subdirectories of it. You could create a `prependModulePath()` function to simplify your modulefiles.

Here is the SitePackage.lua file:

```
require("sandbox")

function prependModulePath(subdir)
    local mroot = os.getenv("MODULEPATH_ROOT")
    local mdir = pathJoin(mroot, subdir)
    prepend_path("MODULEPATH", mdir)
end

sandbox_registration{ prependModulePath    = prependModulePath, }
```

Then in your lua modulefiles you can use it to extend MODULEPATH:

```
-- gcc modulefile:
local version = "7.1"
prependModulePath(pathJoin("Compiler/gcc/",version))
```

Note that if you wish to use any function defined in SitePackage.lua in a TCL modulefile, you will need to modify `tcl2lua.tcl` to know about these functions. You’ll also have to merge your changes to `tcl2lua.tcl` into any new version of Lmod.

There is no need to modify `tcl2lua.tcl` if you only add hook functions.

5.25 Converting shell scripts to modulefiles

As it is stated in *Frequently Asked Questions*, site should convert a shell script into a modulefile. This way the environment initialization can be reversed when the module unloaded. Lmod provides a script called `sh_to_modulefile` which will convert a script to a modulefile. An example is:

```
% $LMOD_DIR/sh_to_modulefile ./foo.sh > foo_1.0.lua
```

or:

```
% $LMOD_DIR/sh_to_modulefile --output foo_1.0.lua ./foo.sh
```

This program defaults to generating a lua based modulefile. It is possible to generate a TCL modulefile with:

```
% $LMOD_DIR/sh_to_modulefile --to TCL --output foo_1.0 ./foo.sh
```

See:

```
% $LMOD_DIR/sh_to_modulefile --help
```

for all the options.

The way it works is that remembers the initial environment and runs the script. The program then compares the initial environment and generate environment. The output is a report of the environment changes.

Note that it doesn't know about any aliases or shell functions that the script creates. They will have to added by hand.

5.26 Tracking Module Usage

Once you have a module system, it can be important to know what modules your users are using or not. This ability collect in some fashion has existed for a long time. What is new here is a complete solution: Using syslog to track module usage. Then collect this data into a database.

There are a number of steps but all are easy. The following is an overview. I am going to explain our setup and point out where you site may be different. In this setup I will be assuming that you are using a MySQL database and Rsyslog (version 5.8)

For a cluster, it is common to have each node in that cluster to send its syslog messages to a single central machine (called master in this discussion). I will be sending the module usage message through syslog to a separate machine. That separate machine will collect the module tracking data into a log file that just contains tracking data. Finally this data is written into the database.

Also provided is a script to analyze the data as shown below:

```
$ analyzeLmodDB --sqlPattern '%fftw%' counts --start '2015-01-01 --end '2015-02-01'
```

Module path	Distinct Users
-----	-----
/apps/intel13/mvapich2_1_9/modulefiles/fftw3/3.3.2	151
/apps/intel13/mvapich2_1_9/modulefiles/fftw2/2.1.5	62
/apps/intel13/impi_4_1/modulefiles/fftw3/3.3.2	45
/apps/intel13/impi_4_1/modulefiles/fftw2/2.1.5	19

This shows number of users of any fftw module for the first month of 2015.

5.26.1 Detailed Steps

Step 1

Use SitePackage.lua to send a message to syslog.:

```
-----
-- load_hook(): Here we record the any modules loaded.

local hook      = require("Hook")
local uname     = require("posix").uname
local cosmic    = require("Cosmic"):singleton()
local syshost   = cosmic:value("LMOD_SYSHOST")

local s_msgA = {}

local function load_hook(t)
```

(continues on next page)

(continued from previous page)

```

-- the arg t is a table:
--      t.modFullName:  the module full name: (i.e: gcc/4.7.2)
--      t.fn:           The file name: (i.e /apps/modulefiles/Core/gcc/4.7.2.lua)

-- use syshost from configuration if set
-- otherwise extract 2nd name from hostname: i.e. login1.stamped2.tacc.utexas.edu
local host      = syshost
if (not host) then
    local i,j, first
    i,j, first, host = uname("%n"):find("([^.]*)%."([^.]*)%.")
end

if (mode() ~= "load") then return end
local msg      = string.format("user=%s module=%s path=%s host=%s time=%f",
                               os.getenv("USER"), t.modFullName, t.fn, uname("%n
↪"),
                               epoch())

    local a      = s_msgA
    a[#a+1]      = msg
end

hook.register("load", load_hook)

local function report_loads()
    local a = s_msgA
    for i = 1, #a do
        local msg = a[i]
        lmod_system_execute("logger -t ModuleUsageTracking -p local0.info " .. msg)
    end
end

ExitHookA.register(report_loads)

```

This code uses two “hook” functions. The first is `load_hook`. This means that every load will be saved. The second hook is called at exit. If there were no errors then any module loads are reported by sending a syslog message with the tag “ModuleUsageTracking”

Please read the file `src/SitePackage.lua` to see how to use the environment variable `LMOD_PACKAGE_PATH` to point to your own `SitePackage.lua`.

You should check to see that Lmod finds your `SitePackage.lua`. If you do:

```
$ module --config
```

and it reports:

```

Modules based on Lua: Version X.Y ...
    by Robert McLay mclay@tacc.utexas.edu

Description                               Value
-----
...
Site Pkg location                         standard

```

Then you haven’t set things up correctly.

Step 2

Have “master” send the tracking messages to a separate computer.

You can add the following to master’s /etc/rsyslog.conf file:

```
if $programname contains 'ModuleUsageTracking' then @module_usage_tracking
&~
```

Where you change “module_usage_tracking” into a real machine name. Adding this to rsyslog.conf will direct all syslog messages to be sent to the “module_usage_tracking” machine.

Remember to restart the rsyslog daemon on master.

Step 3

On the “module_usage_tracking” machine you add to /etc/rsyslog.conf the following:

```
# read in include files
$IncludeConfig /etc/rsyslog.d/*.conf...
```

Then in /etc/rsyslog.d/moduleTracking.conf:

```
$Ruleset remote
if $programname contains 'ModuleUsageTracking' then /var/log/moduleUsage.log
$Ruleset RSYSLOG_DefaultRuleset

# provides UDP syslog reception
$ModLoad imudp
$InputUDPServerBindRuleset remote
$UDPServerRun 514
```

The above commands are in the language of rsyslog version 5.8. What this says is accept outside syslog messages on port 514 and if any are tagged with “ModuleUsageTracking” then write them to /var/log/moduleUsage.log

Remember to restart the rsyslog daemon on the “module_usage_tracking” machine.

Step 4

Create the file /etc/logrotate.d/moduleUsage:

```
/var/log/moduleUsage.log{
    missingok
    copytruncate
    rotate 4
    daily
    create 644 root root
    notifempty
}
```

This will log rotate the moduleUsage.log. Remember to restart the logrotate daemon. Note that it will be the second day before the log is rotated. On Centos machines, it seems that the log rotate happens at about 3am.

Step 5

I found the following site helpful in getting the MySQL database setup:

```
http://zetcode.com/db/mysqlpython/
```

1. Install MySQL db. Create a mysql root password. Then create an account in the database like this:

```
$ mysql -u root -p
Enter password:

mysql> CREATE DATABASE lmod;

mysql> CREATE USER 'lmod'@'localhost' IDENTIFIED BY 'test623';

mysql> USE lmod;

mysql> GRANT ALL ON lmod.* TO 'lmod'@'localhost';

mysql> quit;
```

You will want to change ‘test623’ to some other password. You’ll also probably want to allow access to this database from outside machines as well.

2. Use the “conf_create” program from the contrib/tracking_module_usage directory to create a file containing the access information for the db:

```
$ ./conf_create
Database host:
Database user: lmod
Database pass:
Database name: lmod
```

Where you’ll have to fill in the correct name for the database host and password. This creates a file named lmod_db.conf which is used by createDB.py, analyzeLmodDB and other programs to access the database.

3. Make sure your python knows about the MySQLdb module. Please use pip or something similar if it is unavailable.
4. Create the database by running the createDB.py program.:

```
$ ./createDB.py
```

Step 6

1. If you have more than one cluster and you want to store them in the same database you might want to modify the store_module_data program found in the contrib/tracking_module_usage directory. It assumes that host names are of the form: node_name.cluster_name.something.something and the current store_module_data program picks off the second field in the hostname. If your site names things differently you should modify that routine to match your needs.
2. I use a cron job to load the moduleUsage.log-* files. This is the script I use:

```
#!/bin/bash

PATH=<path_to_python>:$PATH
cd ~mclay/load_module_usage

for i in /var/log/moduleUsage.log-*; do
    ./store_module_data $i
```

(continues on next page)

(continued from previous page)

```

if [ "$?" -eq 0 ]; then
    rm -f $i
fi
done

```

Where <path_to_python> has a python that can also import MySQLdb python module. Also you'll probably want to change ~mclay/load_module_usage to where ever you have the store_module_data program and lmod_db.conf files.

I am running this cron job on the "module_usage_tracking" machine at 5am every morning. This is after the log rotation has been done.

Step 7

Once data is being written to the database you can now start analyzing the data. You can use SQL commands directly into the MySQL data base or you can use the supplied script found in the contrib/tracking_module_usage directory: analyseLmodDB:

```

% ./analyzeLmodDB --help
usage: analyzeLmodDB [-h] [--dbname DBNAME] [--syshost SYSHOST]
                  [--start STARTDATE] [--end ENDDATE]
                  [--sqlPattern SQLPATTERN]
                  cmdA [cmdA ...]

positional arguments:
  cmdA                  commands: counts, usernames, modules_used_by

optional arguments:
  -h, --help            show this help message and exit
  --dbname DBNAME       lmod db name
  --syshost SYSHOST     system host name
  --start STARTDATE     start date
  --end ENDDATE         end date
  --sqlPattern SQLPATTERN sql pattern for matching

```

There are three kinds of reports this program will report. Only one command at a time.

1. counts: Report the number of distinct users of a particular module:

```

$ analyzeLmodDB --sqlPattern '%fftw%' --start '2015-01-01' --end '2015-02-01' counts

```

Module path	Distinct Users
/apps/intel13/mvapich2_1_9/modulefiles/fftw3/3.3.2	151
/apps/intel13/mvapich2_1_9/modulefiles/fftw2/2.1.5	62
/apps/intel13/impi_4_1/modulefiles/fftw3/3.3.2	45
/apps/intel13/impi_4_1/modulefiles/fftw2/2.1.5	19

To get all modules loaded in a date range do:

```

$ analyzeLmodDB --sqlPattern '%' --start '2015-01-01' --end '2015-02-01' counts

```

2. usernames: Report users of a particular pattern:

```

$ ./analyzeLmodDB --sqlPattern '%/apps/modulefiles/settarg%' usernames

```

(continues on next page)

(continued from previous page)

Module path	User Name
-----	-----
/opt/apps/modulefiles/settarg/5.8	user1
/opt/apps/modulefiles/settarg/5.8	user2
/opt/apps/modulefiles/settarg/5.8	user3
/opt/apps/modulefiles/settarg/5.8.1	mclay
/opt/apps/modulefiles/settarg/5.9.1	user5

3. modules_used_by: Report the modules used by a particular user:

```
$ ./analyzeLmodDB --start '2015-01-01' --end '2015-02-01' --sqlPattern 'mclay' ↵
↵modules_used_by
```

Module path	User Name
-----	-----
/opt/apps/gcc4_9/modulefiles/boost/1.55.0.lua	mclay
/opt/apps/gcc4_9/modulefiles/mvapich2/2.1	mclay
/opt/apps/gcc4_9/mvapich2_2_1/modulefiles/phdf5/1.8.16.lua	mclay
/opt/apps/gcc4_9/mvapich2_2_1/modulefiles/pmetis/4.0.2.lua	mclay
/opt/apps/intel13/modulefiles/boost/1.55.0.lua	mclay
/opt/apps/intel13/modulefiles/mvapich2/1.9a2	mclay

5.27 Combining modules and build systems with settarg

5.27.1 Introduction

The settarg module works with Lmod to help developers manage their compiled software projects. It does so by making it easy to switch between optimized or debug builds; change compiler or other modules; and let the build system know about the changes. The secret of settarg is that it consolidates the state of a build into one environment variable called \$TARG.

One of the big advantages of developing software with a module system is the ability to easily switch between compilers and mpi stacks and other modules. Sometimes one gets a strange compiler error or runtime error with one set of tools and it is convenient try a different compiler to see if it tells you something different.

Another useful trick to switch between the level of optimization with one set of builds that have minimal optimization and the “-g” flag to include the symbol table for ease of debugging and other builds with full optimization.

Finally there are some sites which have “shared-home” filesystems. That is where two or more clusters share the same “home” directory tree. Unless they are exactly the same hardware and running the same version of the operating system, software built on one system might not work on another.

For all these reasons, it is convenient to have the built software reside in different directories. Settarg makes it easy to place all the objects, libraries and executables in a separate directory for each kind of build so that they never mix and this avoids hard to resolve bugs.

Settarg manages these environment variables but it is up to the software developer to integrate these variables in to their build system. More on how to modify a Makefile to know about all the TARG variables later.

The keys to settarg

The keys to settarg are:

1. It manages a set of environment variables including \$TARG

2. It is integrated with Lmod so that changes in modules automatically changes \$TARG
3. Your PATH is dynamically adjusted with \$TARG
4. The title bar is dynamically managed.
5. It is easy to configure with site, user and project control.

All you need to is load the settarg module and it will control these variables including PATH that you can use to control your build process. The most important one of these is TARG short for target. It contains a string for the machine architecture, the build scenario, (opt, dbg), the compiler and the mpi module. For example, TARG might be:

```
TARG = OBJ/_x86_64_06_2d_dbg_gcc-7.1_openmpi-2.2
```

Settarg also generates other environment variables to be used to control your Makefile or other build tool. So for the above \$TARG, the following variables are also set:

```
TARG_MACH           = x86_64_06_2d
TARG_BUILD_SCENARIO = dbg
TARG_COMPILER       = gcc/7.1
TARG_COMPILER_FAMILY = gcc
TARG_MPI            = openmpi/2.2
TARG_MPI_FAMILY     = openmpi
```

Knowing that TARG_COMPILER_FAMILY is gcc or intel can mean it is easy to set the compiler flags appropriate for each compiler.

But the most important feature of settarg is that if you change your compiler from gcc to intel/18.0.1 then TARG and the other environment variable change automatically:

```
TARG           = OBJ/_x86_64_06_2d_dbg_intel-18.0.1_openmpi-2.2
TARG_COMPILER  = intel/18.0.1
TARG_COMPILER_FAMILY = intel
```

and all the other variables remain the same. If you change the build scenario from dbg to opt then the following would be the result:

```
TARG           = OBJ/_x86_64_06_2d_opt_intel-18.0.1_openmpi-2.2
TARG_BUILD_SCENARIO = opt
```

Integration with your build system

Once we have these environment variables, we can use them to control where our software is built. It is possible to place all the objects, libraries and executable stored in the \$TARG directory. If all the generated files are in the \$TARG directory, then changing the compiler will result in a different TARG directory. So each \$TARG directory is independent and we won't require a `make clean` between changing compilers or build scenarios.

Integration with PATH

It is useful to have your PATH point to the new \$TARG directory, so settarg changes your path to include \$TARG by removing the old value of \$TARG and replacing it with the new value of \$TARG. This way you can set \$TARG, build, then run the new executable.

Xterm title bar support

If the environment variable `LMOD_SETTARG_TITLE_BAR=yes` and `$TERM` has “xterm” in the string, then loading the module `settag` will turn on title bar support. A typical string in the title bar might be:

```
(D G/5.2 M/3.2)user@host: ~/bin
```

where `user` is your user name and `host` is your hostname followed by your current directory. The string in the parentheses are what `settag` are providing. The “D” is `dbg` build scenario, the “G/5.2” is an abbreviation for the `gcc/5.2` compiler module and “M/3.2” is an abbreviation for the `mpich/3.2 mpi` module. The abbreviations are controlled by configuration files. This string is `TARG_TITLE_BAR_PAREN`.

`Settag` uses the following rules to define what the host is. If the environment variable `$SHOST` is defined then that value is used for the hostname. Otherwise the first name is used from the hostname. In other words, your host name is “`login1.stampede2.tacc.utexas.edu`” then “`login1`” will be used as the host. Obviously, if you’d rather have “`login1.stampede2`” be in the title bar then define `SHOST` to do so.

Settag configuration

Lmod provides a default configuration for `settag` in the file `settag_rc.lua`. Sites may have to tailor this file to match the names of their compilers and mpi modules and other module names. Then users may wish to set their own preferences. Finally a project may wish to have specialized settings. All files are merged together in an intelligent fashion into a single configuration. They do not overwrite the previous setting. More on this in [Settag configuration files](#)

Commands

The environment `TARG`’s value is typically used as a name of the build directory. So the `settag` module provides some helpful aliases to take advantage of this.

1. `gettargdir`: it is simply an alias for “`echo $TARG`”
2. `cdt`: Another alias: “`cd $TARG`”
3. `settag`: How to set the build scenario and to access other features.

By default `settag` has an “empty” build scenario. This can be changed by:

```
$ settag dbg
$ settag opt
```

Which will change `TARG_BUILD_SCENARIO` to “`dbg`” or “`opt`”. Also:

```
$ settag --report
```

report the state of the `.settag` table after combining all the possible `.settag.lua` files.

For those of you who like short commands, please configure Lmod with `--with_settag=full` or set the environment variable `LMOD_SETTARG_FULL_SUPPORT=full` before loading the `settag` module. One useful command is:

```
$ targ
```

which is a short for “`gettargdir`”. Also if you switch between build scenarios frequently may wish to define the following shortcuts for setting the build scenario:

```
dbg() { settarg "$@" dbg; }
opt() { settarg "$@" opt; }
mdbg() { settarg "$@" mdbg; }
empty() { settarg "$@" empty; }
```

What environment variables are defined by settarg

Below are a typical list of variables:

```
TARG_SUMMARY=x86_64_06_2d_dbg_gcc-7.1_openmpi-2.2
TARG=OBJ/_x86_64_06_2d_dbg_gcc-7.1_openmpi-2.2

TARG_TITLE_BAR=D G/7.1 O/2.2
TARG_TITLE_BAR_PAREN=(D G/7.1 O/2.2)

TARG_BUILD_SCENARIO=dbg

TARG_MACH=x86_64_06_2d

TARG_COMPILER=gcc/7.1
TARG_COMPILER_FAMILY=gcc

TARG_MPI=openmpi/2.2
TARG_MPI_FAMILY=openmpi

TARG_OS=Linux-2.6.32-279
TARG_OS_Family=Linux
TARG_HOST=stampede
```

Here is a glossary of what each of these variables mean:

TARG_SUMMARY: The dynamic combination of items like the machine architecture, build scenario, etc. See below for how this gets built.

TARG: This variable contains all the “interesting” items. How is put together is described later.

TARG_TITLE_BAR: This contains everything in TARG_SUMMARY but it is abbreviated to fit the space available. This string is provided in case the user wishes to use this variable as part of their own title bar string.

TARG_TITLE_BAR_PAREN: This is \$TARG_TITLE_BAR with parentheses around the string. This variable is typically used in the xterm title bar.

TARG_BUILD_SCENARIO: This can be used to control compiler flags so that “dbg” might mean to create a debuggable executable. Where as “opt” might mean to build a fully optimized build. To clear this field use the command `settarg empty`.

TARG_MACH: This is the machine architecture along with the cpu family and model number in two hex numbers when on Linux system that has the pseudo file `/proc/cpuinfo`. The architecture is what is reported by “`uname -m`”

TARG_COMPILER: The name of the compiler and version written as `<compiler>/<version>`

TARG_COMPILER_FAMILY: The name of the compiler without the version.

TARG_MPI: The name of the mpi module and version written as `<mpi>/<version>`

TARG_MPI_FAMILY: The name of the mpi module without the version.

TARG_OS, TARG_OS_FAMILY: These are the OS name and family. These variables are always defined even if there are not part of TARG_SUMMARY.

TARG_HOST: See below on how this is extracted from *hostname -f*

5.27.2 Settarg configuration files

Below is a typical configuration file. This file contains several tables written in Lua. If you don't know Lua, it still should be easy to modify this table. Just remember the comma's.

The BuildScenarioTbl table maps host name to initial Build Scenario state. So the default is “empty” which means that the TARG_BUILD_SCENARIO is undefined. If you are on “login1.stampedede.tacc.utexas.edu” your default TARG_BUILD_SCENARIO will be “opt”. Similarly, any host with “foo.bar.edu” will have a default scenario of “dbg”.

```
BuildScenarioTbl = {
  default          = "empty",
  ["tacc.utexas.edu"] = "opt",
  ["foo.bar.edu"]   = "dbg",
}

ModuleTbl = {
  build_scenario = { "dbg", "opt", "empty"},
  compiler       = { "intel", "pgi", "gcc", "sun", },
  mpi            = { "mpich", "mpich2", "openmpi", "mvapich2", "impi"},
  solver         = { "petsc", "trilinos"},
  profiling      = { "mpiP", "tau"},
  file_io        = { "hdf5", "netcdf", },
}

TargetList = { "mach", "build_scenario", "compiler", "mpi"}

SettargDirTemplate = { "$SETTARG_TAG1", "/", "$SETTARG_TAG2", "$TARG_SUMMARY" }

NoFamilyList = {"mach", "build_scenario"}

TitleTbl = {
  dbg      = 'D',
  opt      = 'O',
  impi     = "IM",
  mvapich2 = 'M',
  openmpi  = "O",
  mpich    = "M",
  mpich2   = "M2",
  intel    = "I",
  gcc      = "G",
  phdf5    = "H5",
  hdf5     = "H5"
}

TargPathLoc = "first"

HostnameTbl = { 2}
```

ModuleTbl connects module names with a category. It is also used to define “build_scenario” which is just words to declare a build state. In other words, in the above table “dbg” and “opt” could be anything. The only hard-wired name is “empty”. The category “build_scenario” is also hard-wired. The names of all other categories are not fixed and you are free to add other categories.

This table is also how settarg knows what the names of the compiler and mpi stacks are. If your site uses the name “ompi” for openmpi then the above table will have to be modified to match.

TargetList defines how TARG_SUMMARY is assembled. It is an array of categories. The category “mach” is special it is always defined to be *uname -m* plus on Linux systems it contains the cpu family and model from */proc/cpuinfo*. Each piece is concatenated together with “_”. If an item is undefined then the extra “_” is removed.

Settarg ships with the order given above, but sites and users can change the order to be anything they like. Also notice that there are many more categories than are listed in TargetList. More on this aspect in the “Custom Configuration” section.

SettargDirTemplate specifies how TARG is assembled. In the case shown above then env. var SETTARG_TAG1 is combined with “/” and SETTARG_TAG2 followed by TARG_SUMMARY. Both “TAG” variables have to be set in the environment. Here we have assumed that SETTARG_TAG1 is “OBJ” and SETTARG_TAG2 is “_”. This leads to TARG being:

```
TARG=OBJ/_x86_64_06_2d_dbg_gcc-7.1_openmpi-2.2
```

The NoFamilyList is an array of categories that do not get the FAMILY version. All categories do. For example, if TARG_COMPILER is “gcc/7.1” then TARG_COMPILER_FAMILY is “gcc”.

The TARG_TITLE_BAR and TARG_TITLE_BAR_PAREN are strings that could be used in a terminal title bar. Every item in the TARG_SUMMARY is in the TITLE bar variables (except for TARG_MACH). Because the title bar space is limited, TitleTbl is a way to map each item into an abbreviation. The order in which categories appear on the title bar is the same as TargetList. So a title bar with “O G/7.1 O/2.2” would mean that you are in “opt” mode with gcc/7.1 and openmpi/2.2 loaded.

TargPathLoc controls where (or if) \$TARG. Note that the environment variable LMOD_SETTARG_TARG_PATH_LOCATION is use to control TargPathLoc. Normally the value of TARG is placed in the PATH at the beginning of your PATH. You can place it at the end of your PATH when TargPathLoc = “last”. If TargPathLoc is “empty” then TARG is removed from your path. Actually the rules controlling where TARG goes in your path are slightly more complicated. TargPathLoc controls where \$TARG is placed in your path when TARG was not there before. After the first time TARG is added to your path, TARG maintains its relative location.

Finally, HostnameTbl tells settarg how to extract an entry from the full hostname to be used as TARG_HOST. If your host has multiple components then a “2” would say to use the second component as TARG_HOST. So if your hostname is “login1.stampedede.tacc.utexas.edu” then TARG_HOST would be “stampedede”. If HostnameTbl was “{ 3,2}” then TARG_HOST would be “tacc.stampedede”. If your hostname has a single component then that is used for TARG_HOST.

Custom configuration

Settarg will read up to three separate copies of settarg configuration files. The first one is in the same directory as the settarg command is and is called settarg_rc.lua. The second place is in the user’s home directory (if ~/.settarg.lua exists). Then from the current directory up to “/” it looks for another .settarg.lua (if it exists). It will not re-read the ~/.settarg.lua. Typically a user should copy the system settarg_rc.lua to their home directory (as ~/.settarg.lua) and specify the generally desired behavior. Then in top directory of a project place a simple .settarg.lua that specifies how the target list should be put together for that project:

Suppose that TargetList ~/.settarg.lua is:

```
TargetList = { "mach", "build_scenario", "compiler", "mpi", }
```

Then in ~/project/a there is another .settarg.lua that just has:

```
TargetList = { "mach", "build_scenario", "compiler", "mpi", "file_io" }
```

Normally in any directory your TARG will be the default, but in any directory below ~/project/a TARG will have hdf5 or netcdf if either are loaded.

To see the state of the configuration execute:

```
$ settarg --report
```

5.27.3 Makefile integration

See the contrib/settarg/make_example directory and the README.txt inside. That directory contains a simple Makefile and a more complicated one to a way to use \$TARG in a Makefile so that all generated files (*.o and the executable) are in the \$TARG directory.

There are four main points to converting a Makefile to know about settarg. The first is to set the compiler based on TARG_COMPILER_FAMILY:

```
CC := gcc
#####
# Use TARG_COMPILER_FAMILY to set the C compiler name

ifeq ($(TARG_COMPILER_FAMILY),gcc)
    CC := gcc
endif

ifeq ($(TARG_COMPILER_FAMILY),intel)
    CC := icc
endif
```

The second is to set the optimization based on TARG_BUILD_SCENARIO:

```
CF := -O2
#####
# Use TARG_BUILD_SCENARIO to set the compiler options for either
# debug or optimize.

ifeq ($(TARG_BUILD_SCENARIO),dbg)
    CF := -g -O0
endif

ifeq ($(TARG_BUILD_SCENARIO),opt)
    CF := -O3
endif
override CFLAGS := $(CFLAGS) $(CF)
```

The third point is to force the make file to use the \$TARG directory if defined and change the compilation rules:

```
#####
# Use O_DIR as equal to $(TARG)/ so that if TARG is empty then O_DIR
# will be empty. But if $(TARG) as a value then O_DIR will have a
# trailing slash.

ifneq ($(TARG),)
    override O_DIR := $(TARG)/
endif

##### compilation rules #####

$(O_DIR)%.o : %.c
    $(COMPILE.c) -o $@ -c $<
```


The four point is that the dependencies have to change to use `$(O_DIR)`:

```
##### Dependencies #####  
  
$(O_DIR)main.o : main.c hello.h  
  
$(O_DIR)hello.o: hello.c hello.h
```

For small projects, generating the dependencies by hand is manageable. But for larger projects it can get unwieldy. The `Makefile` shows how to generate the dependencies automatically.

TOPICS YET TO BE WRITTEN

1. Optional Software layout, two digit rule
2. Advanced Topics: priority path, .modulerc tricks
3. inherit
4. internal structure of lmod.

INDICES AND TABLES

- `genindex`
- `modindex`
- `search`