

Lmod: The 'secret sauce' behind module management at TACC

November 14 2012



Lmod protects users so they can't load mismatched compilers, libraries, and other parts of the software stack. Credit: Texas Advanced Computing Center

Modules "set the table" for users of high performance computing system by loading the packages each user needs. Lmod, developed at TACC, protects users so they can't load mismatched compilers, libraries, and other parts of the software stack. This capability is part of our "secret sauce" that makes TACC systems easy to use.

It's likely that fellow users on your favorite [supercomputer](#) have widely varying needs. The applications, compilers, and libraries that you need are probably different from the ones other users need. That is where modern environment module systems come in. A good module system "sets the table " for users by loading the packages each user needs.

The Texas Advanced Computing Center (TACC) has developed an innovative module system that addresses some unique challenges facing modern computing centers.

Life in the world of computing has become much more complicated since John Furlani first introduced the concept of modules in 1991. A few years ago, [system administrators](#) at TACC noticed a recurring problem. The existing module system was not preventing researchers from making mistakes. These mistakes were leading to extra work and frustrations for both the scientists and the TACC staff that supported them.

Incompatible packages, the work of manually disabling or re-enabling modules, and having to choose from multiple versions of tools was the norm for users of TACC's Ranger or Lonestar supercomputers. In addition, a parallel application or library usually depends on both the compiler and MPI (Message Passing Interface) stack. At TACC, with three major compiler families and three MPI families, this means that as many as nine versions of a given package might be available. Only one version will run properly and give the correct answer when paired with complementary tools.

In 2009, Robert McLay, research associate at TACC, decided to address some of the recurrent challenges and remedy this situation by creating Lmod, a complete rewrite of the Environment Module system.

"Lmod protects our users so they can't load mismatched compilers, libraries, and other parts of the software stack," McLay said.

McLay presented the new tool at Supercomputing '11 in Seattle, as well as the 2011 IEEE Conference in Austin. He has worked with staff at many leading universities and computing centers to install the tool broadly. Lmod is freely available via Sourceforge and is already being

used at computing centers across the country, including those at the National Center for Atmospheric Research, the Ohio Supercomputing Center, the University of Florida and the University of Tromsø in Norway.

"I advocated the use of Lmod over the alternative implementations because of its flexibility and full access to the power provided by the module files," said Oleksandr Moskalenko, Bioinformatics Specialist at the University of Florida and an early adopter of Lmod. "As we upgraded our cluster from CentOS 5 to RHEL 6 last Summer we made Lmod the default and the only supported way for making the software we provide available to the users"

While Lmod helps prevent users from making mistakes in building their programs and setting up their computations, it also simplifies the work of systems administrators, who previously received several requests for assistance each week from users unsure of which version of a tool to use.

Lmod allows administrators to encode dependencies into the modules themselves.

Only valid options are available to users. Lmod disables the "wrong" choices, and informs users about precisely which versions of software are the right ones to use. But it goes even further: suppose you have gcc (a popular compiler system that supports various programming languages) loaded, along with a compatible MPI stack and a parallel solver. If you change the compiler, Lmod will automatically replace the MPI stack and other parallel libraries with compatible versions. The user cannot load a module that is not present and the system ensures loaded modules remain compatible with each other.

"This capability is part of our secret sauce," McLay said. "It's part of what makes TACC systems easy to use. It has reduced software selection

errors to nearly zero."

The tool is helpful for beginners and expert users. Every aspect of the tool can be optimized by experienced command-line users to suit their needs. It also provides some behind-the-scenes advantages to administrators. For one, it allows them to track the usage of modules and mine this data to determine which packages are being used, by whom, and how.

This information is useful to centers like TACC as they determine which software to install, upgrade or optimize.

"The metrics built into Lmod have allowed us to track adoption of new biocomputing-oriented software modules and to monitor the effectiveness of our outreach and training strategy," said Matthew Vaughn, director of TACC's Life Sciences group. "As we begin to transition to the next generation of supercomputers powered by extremely parallel microprocessors, this has allowed us to focus our efforts in software optimization for those chips on the most popular pieces of software."

In January 2013, TACC will deploy Stampede, its newest supercomputer and the first large-scale deployment of the Intel Xeon Phi coprocessors. McLay is already developing modules for Lmod that are Xeon Phi-aware, with color coded properties that show a user which packages can be run on the coprocessor.

As high performance computing systems grow larger and more heterogeneous, the need for simplifying tools grows in importance.

Said Moskalenko, "We hope to see TACC continue to develop and improve Lmod for years to come as the computing environments become more complex on the back-end and the need for making it easier

for our clients to access and use the provided software grows."

For McLay, the reason to use the tool is obvious: "If you're not using Lmod, you're working too hard."

Provided by University of Texas at Austin

Citation: Lmod: The 'secret sauce' behind module management at TACC (2012, November 14)
retrieved 3 May 2024 from <https://phys.org/news/2012-11-lmod-secret-sauce-module-tacc.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.