Texas Advanced Computing Center

Allinea Distributed Debugging Tool Helps Researchers Scale Code for Massive Parallel Processing

Customer Success Story

Client

- The Texas Advanced Computing Center in Austin, TX is home to Ranger, one of the world's largest supercomputers for open academic research, rated at 579.4 trillions of floating-point operations per second (TeraFLOPS).

Challenge

- Few people know how to write software that can scale to run across thousands of cores. Traditional debugging methods like printf() cannot help on such a large system. But TACC needs to support hundreds of users preparing their applications for the system.

Solution

- TACC now includes the Allinea Distributed Debugging Tool (DDT) software as part of the standard toolbox for all developers.

Results

- Researchers on Ranger can now use Allinea DDT to debug software written in Fortran, C, or C++ and help their applications scale and run better across many processors

URL References

- sun.com/hpc
- sun.com/sunconstellationsystem
- allinea.com
- tacc.utexas.edu

The Texas Advanced Computing Center (TACC) at The University of Texas in Austin hosts the Ranger system delivered by Sun, one of the most powerful supercomputers for open science. To help developers scale code to run across hundreds or thousands of CPUs, TACC provides the Allinea Distributed Debugging Tool.

Success at a glance

The Texas Advanced Computing Center (TACC) in Austin is home to Ranger, one of the most powerful supercomputers in the world for open science research. By any measure, the Ranger system is big — and fast. Designed, built, and delivered by Sun with a grant from the National Science Foundation (NSF), Ranger's peak performance is rated at a sizzling 579.4 TeraFLOPS.

The Ranger system includes 82 racks of Sun Blade™ 6048 Modular Systems, containing 3,936 blades and 15,744 Third Generation Quad-Core AMD Opteron™ processors. The system also has 125 terabytes of RAM and 1.7 petabytes of raw disk storage. While these specifications are mind-boggling, one simple fact sums it up — Ranger delivers more peak computing power than all the other big systems on the NSF TeraGrid for American academics put together.

Who needs all that number-crunching anyway? "There are more and more things that we just can't afford to simulate in an actual physical lab," says Karl Schulz, Associate Director of High Performance Cmputing at TACC. "Much more research is being done computationally, numerically. That's why the demand has been very high for these big systems."

In a few months since it went online, Ranger has attracted more than 1,000 academic users, who are busy tackling some of the biggest questions of all. For example, one researcher is modeling the birth of the universe in more detail than ever before. Some are studying hurricanes or earthquakes, while others are modeling all the earth's weather systems, or looking for new drugs. In every case, high-speed processing is vital for breakthrough results.

Debugging is critical

Few researchers have experience writing software that can scale to run in parallel across hundreds or thousands of processors. "This is the first time a lot of academics have had access to a system anything like this," says Schulz. "Ranger is so big, we have to make sure the application development process is relatively rapid. That means having a good toolbox available for our users and developers."

This toolbox includes a carefully chosen set of software products that researchers can use to recompile their code, visualize their results, and troubleshoot any problems. Researchers apply for computer time on Ranger. If successful, they upload their own code, most often written in Fortran, C or C++. Then every program must be recompiled, tested, and debugged. "It's non-trivial to write a good parallel application," says Schulz.

"Let's say a user has some code, and the code works fine at 512 processors, so then they run it at 1,024 processors and it gives them a different answer, or the code blows up. This happens all the time. That's why debugging is a very critical need on these big systems."

Debugging cannot be done efficiently using traditional methods, such as using print() statements and stepping through the code to find an error. "Imagine you have a problem with a particular variable in a program running on 2,000 processors. If you put in a print statement, all of a sudden you've got 2,000 prints," says Schulz. "So the classic printf() debugging gets very nasty, very quick."

Allinea Distributed Debugging Tool to the rescue

To help developers scale code to run in parallel across many processors, TACC now includes the Allinea Distributed Debugging Tool (DDT) in the baseline environment for all Ranger users. Allinea DDT is a graphical debugger for scalar, multithreaded, large-scale parallel applications that is surprisingly easy to learn and use.

The intuitive interface makes debugging even the most demanding parallel applications faster and easier. DDT users can debug applications either by processes or threads, with the software automatically analyzing values across processors to find problems fast.

In addition, the unique ability of DDT to control thousands of processes in its parallel stack view shows at a glance the state of every process in a parallel job, along with exactly what line of code is being executed — all in a single view. This design lets developers control thousands of processes from a single window.

"There is value in a graphical debugger that lets you step through a distributed memory program all at once, and gives you one master view," says Schulz. "That's why we wanted to make a graphical debugger available to our users. "DDT provides added debugging for application researchers, and we are pretty satisfied with it," he says. "We've definitely been able to use DDT to find user bugs.

"Time is incredibly precious in our environment," notes Schulz. "If we can save our users time debugging their problem, that's a good return on investment. When people have a tool that can help them solve a problem quicker, they're very happy."

A partnership that works for everyone

The TACC personnel deal directly with the key developers of Allinea DDT to get answers and fixes to any issues. This direct interaction, scientist to scientist, is valuable for both sides. "We view our vendor relationships more as partnerships, because the technology transfer really does go in both directions," says Schulz.

"It's a good opportunity to work with Allinea, so we can expose them to all these different stacks, to lots of different applications, and to a machine of very large scale, so they can develop for the next platform as well. They're very responsive and it's been working pretty well."

> "Debugging is a critical need on these big systems, and we've definitely been able to use Allinea DDT to find user bugs."
>
> **Karl Schulz**
> Associate Director of High Performance Computing, TACC

This teamwork is giving everyone fresh insights into parallel computing. "We're pleased with this partnership," says Dave Maples, VP of Sales for North America for Allinea. "We're happy to collaborate with the researchers and developers at TACC, and to bring them a powerful and cost-effective tool that helps them scale up their applications to run smoothly on Ranger."

"The success of a system like Ranger depends on a number of technologies," says Bjorn Andersson, Director of HPC at Sun Microsystems. "Although Sun can provide state-of-the-art blades, networking, and storage products, developers must be able to write new or modify existing applications to take advantage of such a large system. That is where we depend on partners like Allinea to work with customers such as TACC."

The massive computer system, and the software that scales to run on it, bring new capabilities to researchers at TACC. Using thousands of individual computers to work on a problem cuts down time to results and increases the amount of scientific analysis that can be done. The close partnership between TACC, Sun, and Allinea Software is helping researchers do groundbreaking science that can reshape the world in a big way.

allinea
SCALE TO NEW HEIGHTS

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

Sun microsystems

SunWIN #542877 Lit. #CRSS14498-0 09/08