Case Study: Allinea Software Helps Launch Discoveries On World-Class Supercomputer

Snapshot

Client: The Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Laboratory (ORNL), home to one of the fastest supercomputers in the world running at 3.3 petaflops on 299,000 cores.

Challenge: As researchers scale up their code to work on the colossal OLCF machines, traditional debugging methods, such as printfs, are too time-consuming to be feasible.

Solution: Allinea DDT, which can quickly find problems, even on code running on more than 200,000 cores.

Results: Software fixes now take hours instead of weeks, so developers can focus on their research, not on debugging.

Summary quote: "From one computational scientist to another, I don't think anybody does debugging any better than Allinea Software. Not even close." – Joshua Ladd, Tools Project Technical Officer during the OLCF3 Project.

Case study

In the race to build the world's fastest supercomputer, researchers need the best tools they can find. That's why developers at the Oak Ridge Leadership Computing Facility (OLCF) are working with Allinea DDT.

The OLCF's Jaguar is already among the world's most powerful supercomputers, a Cray XK6. Developers are now adding graphics processing units (GPUs) to transform it into a speedier machine, re-named Titan. This system is expected to have a peak theoretical performance of more than 20 petaflops, making it one of the world's fastest supercomputers.

The goal is to attract top computational researchers to perform breakthrough research at Oak Ridge National Laboratory – the home of the OLCF. Investigators who use OLCF's supercomputer span the universe of scientific inquiry, including alternative energy, astrophysics, climate, new materials, nuclear physics, and combustion.

"Part of the mission of the Titan project is to provide a comprehensive programming ecosystem that allows researchers to be as productive as possible," says Joshua S. Ladd, Tools Project Technical Officer during the OLCF3 Project. "A major component of that ecosystem is the debugger."

A supercomputer needs a super debugger. Across thousands of cores, the traditional method of doing printfs to locate problem code becomes intractable. With Allinea DDT, developers can pinpoint any failures quickly because it gives them a single view of every process in a parallel job, along with exactly what line of code is being executed.

"Titan is really cutting-edge technology, and it's even more exciting because it's not immediately clear what kind of issues users are going to run into when porting their code to the GPUs," says Ladd. "To help encourage researchers to use the GPU accelerators, they must have the most powerful and effective tools at their disposal, tools like Allinea DDT," he says. "We're excited for users to run into bugs on the GPUs to see this tool in action."

Allinea DDT Propels Nuclear Research To The Next Level

To gain access to OLCF's facilities, researchers compete for computing time through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program. INCITE projects tackle grand-challenge investigations in science and technology with scientific applications that scale to roughly 60,000 cores up to the full machine. Since most researchers have never scaled their code up to this size, problems are inevitable.

"You need high-grade software tools that can scale along with your code – a debugger in particular – because when problems arise at scale, you are in a totally different universe," says Ladd.

Computational Scientist, Hai Ah Nam, frequently proves Allinea DDT's value in her role as a scientific computing liaison for the INCITE program. Her most recent project involved an application, Bigstick, intended to describe the properties of various atomic nuclei of different substances. This basic science research contributes to many fields, including energy and medical research.

Although Nam had a good understanding of the algorithm, she had only a few days to work with the researcher.

"He was having trouble with this 'Heisenbug' (Heisenbugs are bugs that mysteriously vanish whenever you try to "observe" them, typically with a printf, because you've altered latencies between interprocessor communications) that only showed up when he scaled to a certain number of processors," she says. "And it happened sporadically."

By the time the researcher got through his first print statement and looked at one part of his code, Nam had figured out the problem with Allinea DDT.

"I stunned him by finding his problem so quickly," she says. "I was able to do it in one sitting, about an hour. I suspect it would have taken him at least a couple of weeks."

Nam can empathize with scientists who want to focus on their research rather than learning yet another software application.

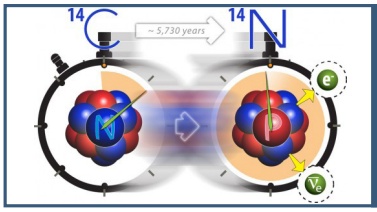

Research In Action – Physicists Explore What Makes Carbon-14 Tick

Carbon-14 decays through beta decay, in which the nucleus emits an electron and an anti-neutrino and becomes a nitrogen-14 nucleus.

Visualization by Hai Ah Nam and Andrew Sproles, ORNL.

"At small scale, I could printf myself out of a problem," she admits. "But now, with codes running on tens to hundreds of thousands of cores, I have to use Allinea DDT if I want to solve the problem quickly."

Speed And Scalability Out Of This Universe

With Allinea DDT's graphical interface and simple point-and-click operation, OLCF reports a high adoption rate and positive feedback from its researchers.

"Allinea DDT is tightly integrated into the Cray programming environment. We worked with Allinea Software to ensure that," says Ladd. "All you really need to do is load the Allinea DDT module and type 'ddt' on the command line to fire up the GUI, and you're ready to go. And the GUI is just point and click with a mouse."

And it's fast.

"I was able to fire up Allinea DDT on 130,000-plus cores in less than 30 seconds," says Ladd. That's faster than starting some laptops with only two to four cores.

While researchers use Allinea DDT to debug across tens of thousands of cores, Ladd and his team place much larger demands on the tool.

"It was a big deal when Allinea Software came here two years ago, and they were able to start Allinea DDT on all of Jaguar's 225,000 cores," he said.

Plus, Allinea DDT has solved some unusual challenges. Ladd and his team used the program to debug an open source implementation of the Message Passing Interface (MPI) middleware. The work was at a very large scale, a half-million lines of code running on 100,000 to 225,000 cores.

"Even your typical nuclear supernova application is not that size," says Ladd. "Debugging inside MPI is a vast universe of complexity that touches all aspects of a supercomputer – the network, the CPU, and the memory. All of these factors can conspire to cause problems at scale."

"By having the ability to step through the code, we could identify and resolve issues that I don't think we would have been able to without Allinea DDT."

Debugging also gets tricky when code has errors but still runs. To address this problem, Allinea Software is collaborating with VisIt – open source software used to visualize large scientific data sets. A visual inspection enables researchers to look at a picture of the data, click on different cells, and inspect the process generating the data.

"So let's say the output is a video of a star exploding," says Ladd. "As that star explodes, if there are all kinds of weird asymmetries, you probably have some bug in your math. With a visualized debugging tool, if it doesn't look like you expected, you go through the process to determine if you've got a bug in your code, or if you've discovered something new."

Match Made In Heaven

Overall, Ladd describes the working relationship with Allinea Software as a gratifying partnership.

"I think it has been rewarding for the Allinea Software folks to see their baby running at this scale, and it's been rewarding for us to have it as a productive contributor in our tools suite."

The OLCF has an Allinea Software staff member onsite and Ladd usually does biweekly conference calls with Allinea Software's CTO, David Lecomber.

"He has highly respectable technical chops and can talk shop with the best of them," says Ladd.

By creating the fastest supercomputer with the best super-tools to support it, OLCF and Allinea Software are creating a solid launch pad for breakthrough discoveries.