

Helping Developers at Lawrence Livermore National Laboratory Create Grand Challenge Applications on the IBM Blue Gene Supercomputer

End-User

Lawrence Livermore National Laboratory (LLNL) is a premier applied science laboratory, part of the National Nuclear Security Administration (NNSA) within the Department of Energy (DOE). LLNL has been managed since its inception by the University of California for the U.S. government.

Challenge

Researchers at the LLNL develop mission-critical Grand Challenge applications using the IBM Blue Gene/L, the fastest supercomputer in the world. Applications written for the IBM Blue Gene/L are highly complex, using thousands of processors and consuming gigabytes of memory, and developing efficient code for such an advanced supercomputer presents great challenges for developers.

Solution

Developers at LLNL use TotalView, the industry's most advanced multithreaded debugging tool, to help develop and debug their mission-critical applications on the IBM Blue Gene/L. TotalView was designed from the ground up to handle complex parallel applications that scale, by offering many advanced features that understand parallel programming models like MPI and OpenMP. TotalView has made it significantly easier and quicker for LLNL developers to create their complex applications.

Background

As a national security laboratory, LLNL is responsible for ensuring that the nation's nuclear weapons remain safe, secure and reliable by applying the latest advances in science and engineering. With its special capabilities, the laboratory also meets other pressing national security needs, which include countering the proliferation of weapons of mass destruction and strengthening homeland security.

Given the critical role that LLNL plays in national security, it is no surprise that they utilize the world's fastest supercomputer, the IBM Blue Gene/L, to develop their mission-critical applications.

The Blue Gene/L installation at LLNL is an Accelerated Strategic Computing Initiative (ASCI)/LLNL-only capability resource that is used for very large, highly parallel (512 – 65,000 compute nodes) jobs. The system has 121,072 CPUs (2 CPUs/node) with a theoretical system peak performance of 367 teraflops per second.

At LLNL, Blue Gene/L is optimized to run molecular dynamics applications at extreme speeds to address materials aging issues confronting the Stockpile Stewardship Program. Blue Gene/L is also used to explore the potential of system-on-a-chip technologies to achieve extreme speed while minimizing floor space and electrical power consumption.

Application Development Challenges

Applications being developed on the Blue Gene/L range from "simple" scalable linear solvers to large hydrodynamic and simulation codes that use multiple languages and network communication patterns.

One developer at LLNL described his application: "The code is a large, highly portable hydrodynamics code. It is a mixture of C, C++, Fortran and Fortran 90. It compiles to a 37 MB executable on Blue Gene/L when optimized. It has a variety of network communication patterns that are dynamic over time. It incorporates many different third-party libraries and thus must embrace a large number of coding styles with different language features used. It will soon be able to run on all 12,800 processors of the Blue Gene/L machine."

For this developer, the biggest challenge is debugging his code that crashes when scaled on a very large number of processors. "We don't always have the luxury of scaling back to 2048 processors," he said. "On the other hand, the debugger needs to work fairly quickly at this scale to be of real use."

Another programmer is developing scalable linear solvers, mainly algebraic multi-grid, written in C and consisting of short programs designed to be used by bigger applications. The goal is scalability across a large number of processors, which is difficult since

the algorithms require a large amount of communication across many processors.

"I like the fact that I can look at all jobs in parallel and really see what is going on all processors at the same time. Also, being able to set conditional breaks has been helpful."

TOTALVIEW® Case Study

Other applications developed on Blue Gene/L at LLNL include a large multiphysics code written in C, which runs on a variety of platforms and has been used for scaling studies up to 12,000 processors.

How TotalView Helps

LLNL developers use the TotalView debugger from TotalView Technologies (formerly Etnus) to understand and reduce the complexity of developing applications on Blue Gene/L. TotalView is the most proven and scalable debugging product of its kind, able to handle from one to thousands of processes. The advanced debugging capabilities of TotalView, including independent thread control, multi-platform support, register and instruction level debugging, and a built-in memory debugger, have been proven to reduce development time by more than 20 percent.

For the programmer developing scalable linear solvers, using TotalView has yielded great benefits. "TotalView has been extremely helpful as part of my development process in finding bugs... I like the fact that I can look at all jobs in parallel and really see what is going on all processors at the same time. Also, being able to set conditional breaks has been helpful."

The developer working on large hydrodynamics code lauds TotalView's breakpoint management capabilities, as well as its ability to scale transparently up to thousands of processors or processes while remaining easy to use. "The ability to step through individual processes in addition to aggregations of processors is useful, as is the fact that breakpoints are saved when a job exits. The fact the TotalView understands C++ method calls is also useful," he noted.

"TotalView has been extremely helpful as part of my development process in finding bugs."

According to the third developer, "Typically, we debug on no more than 64 – 128 processors. It is easy to use TotalView on 4,096 processors and this has helped get the code scaled up. The nicest feature is the ability to trap on memory writes to specific locations. That's irreplaceable and can save time on nasty bugs."

About TotalView Technologies

TotalView Technologies (formerly Etnus) is the world's leading provider of debugging and analysis software solutions for the multi-core age. TotalView Technologies products enable software developers to quickly, easily and effectively debug UNIX, Linux, and Mac OS X applications running on development machines with single, dual-core, multi-core, or multiple processors. For more than 20 years, TotalView Technologies products have been at work in research institutions, government laboratories, and technical computing centers, as well as commercial enterprises in the financial services, telecommunications, biotech, aerospace, weather prediction, film special effects and animation, oil and gas exploration, and computer-aided engineering markets. Recognized worldwide as the gold standard for debugging in high-performance, distributed or cluster computing environments, TotalView's award-winning technology is used by some of the world's largest supercomputers. For more information, visit www.totalviewtech.com.

"The nicest feature is the ability to trap on memory writes to specific locations. That's irreplaceable and can save time on nasty bugs."

About IBM Blue Gene®

The IBM® System Blue Gene® solution is the result of an IBM supercomputing project begun over five years ago, dedicated to building a new family of supercomputers optimized for bandwidth, scalability and the ability to handle large amounts of data while consuming a fraction of the power and floor space required by today's fastest systems. Today Blue Gene ranks as the number one and number three fastest supercomputers on the TOP500 list along with 10 other entries in the top 50, and 28 entries in the top 150.

TotalView Technologies

24 Prime Park Way
Natick, MA 01760
P.508.652.7700
F.508.652.7701