

Los Alamos National Laboratory breaks petaflop barrier with hybrid supercomputer from IBM

Overview

■ Challenge

To allow scientists at Los Alamos National Laboratory to work on compute-intensive national security science projects that require speed and reliability for large system runs

■ Solution

Building a hybrid supercomputer using approximately 12,000 IBM PowerXCell™ 8i Cell Broadband Engine™ processors and more than 6,500 AMD Opteron dual-core processors

■ Key Benefits

Offers the world's fastest computational performance, advancing scientific research in areas such as climate change and plasma physics while enabling low power consumption and a small data center footprint



Los Alamos National Laboratory (LANL) has been at the forefront of national security science since its founding in 1943. The organization now finds itself at the forefront of supercomputing technology, having shattered the petaflop computing barrier with a new hybrid high-performance system from IBM.

LANL, approximately 100 miles north of Albuquerque, New Mexico, is home to over 9,000 scientists and supporting staff, many of whom need a high-performance computing system that can support large-scale full-system runs in the areas of gene sequencing, virus research, plasma physics, and other scientific projects.

“The IBM proposal was deemed to have the best overall value in terms of technical capabilities, price, and IBM’s business capabilities.”

— John Morrison, division leader of the High-Performance Computing Division, Los Alamos National Laboratory

Going for the brass ring

“It’s been a while since Los Alamos has had a large system,” says John Morrison, division leader of the High-Performance Computing Division at LANL. “So we wanted to go for the brass ring.”

Working with LANL scientists, IBM built a hybrid system that combines IBM PowerXCell 8i Cell Broadband Engine processors and AMD Opteron dual-core processors to achieve incredible speed in a small and efficient package. Dubbed “Roadrunner” by LANL after the state bird of New Mexico, the system smashed the petaflop barrier (that’s a quadrillion floating point calculations per second) by running at a sustained 1.026 petaflops, putting the system in the #1 position on the June 2008 TOP500 list of the world’s fastest supercomputers.¹

“The main characteristic of Roadrunner is simple: It’s the first and largest hybrid computer that mixes Cell processors as accelerators to commodity Opteron processors,” explains Ken Koch, technical manager of the Roadrunner project. The project was funded by the U.S. Department of Energy’s National Nuclear Security Administration primarily for use in the Stockpile Stewardship program, which maintains the nation’s nuclear arsenal without the need for nuclear testing.

Partnering with IBM

Los Alamos National Laboratory had a definite need for all that power. Their research projects routinely involve very large scientific simulations and data sets, so they started thinking about building a new high-performance system to handle the workload.

“We didn’t want a cluster of Cells,” says Koch. “We wanted Cells attached to Opterons as acceleration devices. IBM was able to leverage existing products to give us the performance we needed.”

“We had several responders to this procurement, and the IBM proposal was deemed to have the best overall value in terms of technical capabilities, price, and IBM’s business capabilities,” says Morrison.

IBM’s proposal was to build a system with over 12,000 IBM PowerXCell 8i processors—derived from the same chips that power today’s most popular video game consoles. An additional 6,500 AMD Opteron dual-core processors perform basic compute functions, freeing the IBM PowerXCell 8i chips for math-intensive calculations.

A phased approach

The implementation process was broken out into phases, starting with deployment of a base system. “Our phase one deliverable was the original four-socket Opteron system,” says Koch. “It really helped get the IBM system software and Linux® environment ironed out.”

“It also provided much-needed computing resources for our user community,” adds Morrison. “It was a welcome addition to our supercomputing assets.”

This first phase also included some reworking of the architecture that provided better connectivity to the Cells using PCI Express x8. This, along with decreasing the node size and increasing cluster connectivity, provided a significant boost in performance. The current Roadrunner system is based on a tri-blade configuration with IBM BladeCenter® LS21 and IBM BladeCenter QS22 components.

Toward the top of the TOP500

For Roadrunner, it was a quick and easy trip to the top of the TOP500 list. “The system was fully assembled and ready for a full scale run on Friday, May 23, 2008, with a full system impact run the next day,” says Morrison. “That’s a fairly significant accomplishment. One of the world’s foremost LINPACK experts predicted the test would take a month.”

The hybrid architecture inherent in the system design gives the LANL team a chance to experiment with mixed processors and mixed cores. “These types of mixed environments are likely to appear in future systems or microprocessor designs,” says Koch. “The Cells pack a huge punch in terms of computing performance.”

“It’s just very difficult to design any kind of node that is that small, that compact, and with that kind of performance—and we have it” he says.

A green supercomputing machine

Having so much compute power in a small form factor is important, but Roadrunner is also energy efficient and therefore also environmentally friendly, delivering 437 million calculations per watt. This slender energy profile landed Roadrunner at spot number 3 on the Green500 list.²

“Having a system with the computational horsepower of Roadrunner, and yet being very power efficient, really is an advantage for us,” notes Koch.

Solution Components

- IBM BladeCenter® LS21
 - IBM BladeCenter QS22
 - IBM PowerXCell™ 8i Cell Broadband Engine™ processors
 - IBM SDK for Multicore Acceleration
-



Exploring new science

For its first year at LANL, Roadrunner is dedicated to non-classified research projects. “We issued a call for proposals for science research and got a significant response—more than we could accommodate—for full system runs that they could not achieve on any other machine in the world,” recalls Koch. “We’ve got some biological gene information technology projects going on. We have other efforts looking at astrophysics and plasma physics.”

At the conclusion of those projects in 2009, Roadrunner will be dedicated to addressing national nuclear security problems. A smaller additional Roadrunner system consisting of two connected-units, or approximately 12 percent of the main machine, was recently purchased and will be used to continue much of the non-classified research after the main machine goes classified.

“We will have built this enthusiasm among our open science community for Roadrunner resources, so we felt like we needed to have some resources for them to compute on,” says Morrison. “That was the motivation for expanding the footprint of the Roadrunner technology here at the laboratory.”

Creating a computing buzz

What does it mean to Los Alamos National Laboratory to have the world’s fastest supercomputer? “Having a computational resource of this power gives people an exciting opportunity to explore new science,” Morrison says. “So we’ve seen quite a buzz amongst the staff with the arrival of Roadrunner.”

“And not only was this machine a petaflop machine, but its unique architecture gave quite a stir in the industry—wondering how it would work, if it would work,” adds Koch. “There were lots of people questioning how well it might work and I think there are lots of people that are still going to follow it and see how well it plays out.”

For more information

Contact your IBM sales representative or IBM Business Partner. Visit us at:

ibm.com/systems

For more information about Los Alamos National Laboratory, visit: www.lanl.gov

¹ TOP500 List, June 2008.
<http://www.top500.org/list/2008/06/100>

² Green500 List, June 2008.
www.green500.org/lists/2008/06/ranks1-100.php

© Copyright IBM Corporation 2009

IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.

Produced in the United States of America
May 2009

All Rights Reserved

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to imply that only IBM’s product, program or service may be used. Any functionally equivalent product, program or service may be used instead. Offerings are subject to change, extension or withdrawal without notice.

All client examples cited represent how some clients have used IBM products and the results they may have achieved. Performance data for IBM and non-IBM products and services contained in this document was derived under specific operating and environmental conditions. The actual results obtained by any party implementing such products or services will depend on a large number of factors specific to such party’s operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS-IS” WITHOUT ANY WARRANTY, EITHER EXPRESSED OR IMPLIED.



Recyclable, please recycle.

BLC03024-USEN-00