

CASE STUDY

U.S. Dept. of Energy ASC Chooses Moab Cluster Suite for Department of Energy TriLabs Program



KEY FACTS

Overview

ASC integrates the high-performance computing work of Los Alamos, Lawrence Livermore, and Sandia National Labs, as well as academic researchers, into a nationally coordinated program administered by NNSA.

Challenge

Find a resource- and workload-management solution that would be common to all ASC computing systems, improve usability and manageability, improve ROI across the board, offer enhanced reporting functions, and optimize utilization.

Solution

Moab Cluster Suite® and Moab Grid Suite® from Adaptive Computing (formerly Cluster Resources)

Results

Moab® provides optimization of NNSA's compute, network, and storage resources; automated failure detection and recovery; flexible policies; true visualization of cluster activity; detailed accounting; and reduced costs.



ABOUT ASC

Established in 1995, the Advanced Simulation and Computing (ASC) Program supports the Department of Energy's National Nuclear Security Administration (NNSA) Defense Programs' shift in emphasis from test-based confidence to simulation-based confidence.

Under ASC, computer simulation capabilities are developed to analyze and predict the safety, security, and reliability of nuclear weapons and to certify their functionality.

ASC integrates the high-performance computing work of three national security laboratories (Los Alamos National Laboratory, Lawrence Livermore National Laboratory, and Sandia National Laboratories) and university researchers nationally into a coordinated program administered by NNSA.

CHALLENGE

In 2006, ASC initiated a search for a common resource- and workload-management solution to improve usability and manageability of its diverse resources and to improve the

return on its significant computing investment. In addition, the program sought enhanced reporting for managed resources and the ability to optimize resource utilization while still maintaining the flexibility required to meet the individual needs of each site and project.

ASC has a highly heterogeneous environment, with resources that range from large-scale Intel and AMD Opteron-based systems provided by IBM, HP, Dell, and others to more exotic and powerful systems such as Cray's XT4 and IBM's Blue Gene. Going into the assessment, ASC had a high degree of knowledge in the resource-management space, as a result of their own development of advanced resource-management and scheduling tools such as BProc, SLURM, and LCRM.

Roadrunner ushers in a new era for science at Los Alamos National Laboratory. Just a week after formal introduction of the machine to the world, we are already doing computational tasks that existed only in the realm of imagination a year ago.

Terry Wallace, Associate Director for Science, Technology, and Engineering at LANL

"The fundamental needs of ASC are not all that different from today's IT centers," said David Jackson, CTO of Adaptive Computing (formerly Cluster Resources). "They need total optimization of compute, network, and storage resources; scalability; automated failure detection and recovery; more flexible policies; true visualization of cluster activity; detailed accounting; and reduced costs. When you're dealing with more than 100,000 processors, the approaches used to deliver this optimization must become more efficient and manageable."

SOLUTION

At the core of the chosen solution are Moab Cluster Suite® and Moab Grid Suite®, professional cluster-management solutions that leverage the power of Moab Workload Manager®, a policy-based workload-management and scheduling tool, as well as a graphical cluster-administration interface and a Web-based end-user job-submission and management portal.

The contract grants the three national security laboratories ongoing use of the Moab family of software products for workload management, system accounting, capacity planning, automated failure recovery, virtualization, and a host of other capabilities. In addition, Adaptive Computing also works with each of the labs in the areas of collaborative research and development, consulting, 24/7 support, and other professional services. Adaptive Computing will fully support ASC throughout the usage lifecycle, providing onsite and online

training, best-practices consulting, and other enabling services.

RESULTS

The Moab solution adds significant manageability and optimization to HPC resources while providing deployment methods that effectively minimize the risk and cost of adoption.

Unique Moab capabilities allow it to be transparently deployed with little or no impact on the end user. These capabilities include system workload, resource, and policy simulation; batch language translation; capacity-planning diagnostics; nonintrusive test facilities; and infrastructure stress testing.

WORTH NOTING

In a test run in May 2008, Roadrunner (built by IBM with funding from NNSA for Los Alamos National Laboratory) achieved a long-sought supercomputing goal: performing at the petaflop/s level—more than a thousand trillion floating-point operations per second.

Roadrunner is the first supercomputer to use a hybrid processor architecture based on both x64 processors from Advanced Micro Devices (AMD) and the IBM Cell Broadband Engine (Cell BE) processing elements. It will be used by NNSA's laboratories to perform calculations that will improve the labs' ability to certify that the U.S. nuclear weapons stockpile is safe, secure, and reliable without underground nuclear tests. It will also be used for research into astronomy, energy, human genome science, and climate change.

"Roadrunner ushers in a new era of science at LANL," stated Terry Wallace, associate director for science, technology, and engineering. "Just a week after formal introduction of the machine to the world, we are doing computational tasks that existed only in the realm of imagination a year ago."

In breaking the petaflop/s barrier, Los Alamos National Laboratory relied on Moab Workload Manager and TORQUE resource manager from Adaptive Computing when running the LINPACK benchmark on Roadrunner.

In total, Roadrunner connects 6,562 dual-core AMD Opteron chips and 12,240 Cell chips (on IBM Model QS22 blade servers). The Roadrunner system has 98 terabytes of memory and is housed in 278 refrigerator-sized IBM BladeCenter racks occupying 5,200 square feet. Its 10,000 connections—both Infiniband and Gigabit Ethernet—require 55 miles of fiber-optic cable. Roadrunner weighs 500,000 pounds.

Roadrunner's hybrid format sips power (2.35 megawatts) and delivers world-leading efficiency—437 million calculations per watt.

In June 2008, Roadrunner ranked as the No. 1 supercomputer in the Top500 list of the world's most powerful supercomputers and as No. 3 on the Green500 list for energy efficiency.

To learn more about how Moab technology can improve HPC, data center, or cloud computing, contact us—

North America, Latin America, Asia-Pacific +1 (801) 717-3700
U. S. toll-free. +1 (888) 221-2008
EMEA (Europe, Middle East, Africa) +44 (0) 1483 243578
Email info@adaptivecomputing.com
Web site www.adaptivecomputing.com

