

Research team 'virtualizes' supercomputer

January 20 2010

A collaboration between researchers at Northwestern University, Sandia National Labs and the University of New Mexico has resulted in the largest-scale study ever done on what many consider an important part of the future of computing -- the virtualization of parallel supercomputing systems.

As part of this collaboration, Peter A. Dinda, associate professor of electrical engineering and computer science at Northwestern's McCormick School of Engineering and Applied Science, and his graduate student Jack Lange led the development of a [virtual machine](#) monitor called Palacios specifically for supercomputers. The system was tested Dec. 3 on Sandia's world-class Red Storm supercomputer. Sandia researchers, led by Kevin Pedretti, assisted in adapting and optimizing Palacios for the Red Storm environment and directed the testing effort.

Results show that the team successfully virtualized Red Storm using the Palacios virtual machine monitor and ran communication intensive, fine-grain parallel benchmarks of critical interest to Sandia with extremely high performance. Testing went up to 4,096 nodes, making this the largest-scale study by at least two orders of magnitude.

"Virtualizing a parallel supercomputer is particularly challenging because of the need to support extremely low latency, high-bandwidth communication among thousands of virtual machines," Dinda says.

"Supercomputing users and the owners of supercomputers will not tolerate any performance compromises because the machines are so expensive to acquire and maintain, but, on the other hand, they also want

access to the benefits of virtualization."

A virtual machine monitor (VMM) works by separating a computer's [operating system](#) from its hardware. This indirection exposes a range of benefits. For example, a VMM allows an operating system from one machine to be run on another. (If it needs more memory, for example.) It can also allow one machine to simultaneously run multiple operating systems, and it is possible to migrate running operating systems from one computer to another.

In the case of supercomputing, the VMM also acts as a translator between a user's software and the highly specialized hardware and software environments of the system, which could potentially allow more researchers to use supercomputers to solve complex problems.

With more than 38,000 processors, Red Storm is a massively parallel processing supercomputer that was uniquely designed to support modeling and simulation of complex problems in nuclear weapons stockpile stewardship. It is currently the 17th fastest computer in the world, with a theoretical peak performance of 284 trillion floating point operations per second in a relatively compact 3,500-square foot footprint.

Virtualization on such a machine is important because it will allow more researchers to run scientific computing and simulation programs without reconfiguring their software to the machine's specific hardware and software environments. In this context, thousands of virtual machines must cooperate in order to solve large problems. But because the system is extremely expensive to run, any VMM must have low overhead, which is magnified through the fine-grain interactions among the virtual machines.

At these massive scales, the Palacios virtual machine monitor had a

measured overhead of less than 5 percent. The results clearly indicate that it is possible to bring the benefits of virtualization to even the largest computers in the world without performance compromises.

Virtualization is big business, with the market research and analysis firm IDC forecasting annual revenues to grow from \$5.5 billion in 2007 to \$11.7 billion in 2011.

"If we can virtualize supercomputers without performance compromises, we will make them easier to use and easier to manage, generally increasing the utility of these very large national infrastructure investments," Dinda says.

"The end goal is to provide a more flexible supercomputer environment to end users without sacrificing performance," Pedretti says. "The successful experiments with Palacios on Red Storm demonstrate the feasibility of our approach, and we hope to incorporate this technology in future capability supercomputer platforms."

More information: Researchers can learn more about Palacios and download the latest version of Palacios at v3vee.org .

Provided by Northwestern University

Citation: Research team 'virtualizes' supercomputer (2010, January 20) retrieved 25 April 2024 from <https://phys.org/news/2010-01-team-virtualizes-supercomputer.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.