

VIRTUALIZING SUPERCOMPUTERS



IN A SANDIA NATIONAL LABORATORIES FACILITY PROTECTED BY ARMED GUARDS IN NEW MEXICO SITS A 3,500-SQUARE-FOOT MASS OF BLACK CABINETS THAT HOLD THE 38,000 PROCESSORS OF THE OMINOUSLY NAMED RED STORM, **THE 13TH FASTEST COMPUTER IN THE WORLD.**

Over the past several years Red Storm has modeled and simulated complex problems in nuclear weapons stockpile stewardship, determined how much explosive power it would take to destroy an asteroid headed for Earth, and gauged how changes in the composition of the atmosphere affect our planet. It's a formidable tool with computing power that could provide solutions to complex problems in a wide range of scientific areas. But it's a tool that isn't used as often as it could be because Red Storm's highly specialized hardware doesn't mesh with most scientific software.

That's where Peter Dinda comes in.

Dinda, associate professor of electrical engineering and computer science, and graduate student Jack Lange have created a virtual machine monitor (VMM) called Palacios that acts as a translator between a user's software and the specialized hardware of the Red Storm system and other supercomputers. While Palacios was made specifically for supercomputers like Red Storm, researchers say that, thanks to the changing nature of computers, VMMs could solve problems that will eventually show up in desktop computers. (Palacios can also run on standard PCs, and Dinda uses it for teaching on such platforms.)

A VMM works by separating a computer's operating system from its hardware — by “virtualizing” it so an operating system from one machine can run on another, or by allowing one machine to run multiple operating systems simultaneously. In essence, a VMM is software that behaves as if it is hardware.

Dinda and his group created Palacios as part of a National Science Foundation and Department of Energy-funded project to build an open-source, publicly available VMM. The need for VMMs stems from the changing world of computers. For about the last 20 years, the dominant processing language (instruction set) for computer processors has been Intel x86, and most programs have been created using that language. But now computer hardware is changing. Researchers are finding it difficult to make individual processing cores faster, so instead they are increasing the number of processing cores per chip and supporting different kinds of cores. Typical operating systems and applications built for Intel x86 assume that only one or a very small number of processing cores are available,

however. VMMs can act as mediators between the new hardware that employs an exponentially increasing number of homogeneous or heterogeneous cores and the old operating systems and applications designed for a few homogeneous cores. “Because of this disconnect between the language and the architecture, the

problems that happen now in high-performance computers will eventually show up on desktop computers,” Dinda says.

“One way computer architects can use VMMs is by exploring different extensions of the Intel x86 language and different ways of implementing that language,” Lange says. “We're taking a step back and asking, Given that we have VMMs, how can we move beyond the constrictions of that language? How can we redesign everything?”

Dinda's group has been developing Palacios since 2007. Lange wrote most of the program, and an undergraduate student, Andy Gocke (computer science, '12), spent last summer importing it into the Intel VT architecture. This, Dinda says, “is impressive for anybody, but it's especially impressive for a freshman.” Gocke spent his childhood learning how to program computers and was up to the challenge. “As an undergrad you get exposed to a lot of problems that have basically been solved for 30 or 40 years,” says Gocke. “It was really fun to work on something that didn't really have a solution.”

Researchers hope that Palacios can run on Red Storm all the time. That would allow more people to use the system without rebooting, a process that can take an hour on such a large machine. Dinda's collaborators at Sandia National Labs recently tested Palacios on Red Storm — a task that required the tester to sit in a guarded room for several hours as the program ran. The results were successful: the VMM had minimal impact on performance and was tested across thousands of “nodes” on the supercomputer's network. The test indicated that it is possible to bring the benefits of virtualization to even the largest computers in the world without performance compromises.

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

The Palacios VMM developed by Dinda's group and his collaborators at the University of New Mexico and Sandia National Labs is available for anyone to use and can be found at <http://v3vee.org>. — Emily Ayshford **M**