



Spring 2010 CIS Colloquium

Architecting a Symbiotic Virtual Machine Monitor for Scalable High Performance Computing

John R. Lange
(Northwestern University)

11 AM - 12 PM, Tuesday, March 2
Wachman Hall 447
(4th Floor Conference Room)

ABSTRACT:

While virtual machines (VMs) have become commonplace in datacenter and cloud computing environments, they have yet to make inroads into high performance computing (HPC). Despite this, there is considerable interest in bringing the well established benefits of virtualization to that domain. However, there are many concerns about parallel application performance and scalability. Currently, there are currently no virtual machine monitor (VMM) architectures designed with HPC in mind.

In response, we have developed the Palacios VMM, an OS-independent VMM designed specifically for research and use in HPC, as well as for systems and architecture research. Palacios is part of a collaborative research effort between Northwestern University, the University of New Mexico, and Sandia National Laboratories. As of December, Palacios has successfully been tested at scales exceeding 4096 nodes on the RedStorm Cray XT supercomputer at Sandia, the largest scale performance study of virtualized high performance computing ever conducted. Overheads were extremely low, and Palacios-based virtualization had minimal (less than 5%) impact on scalability. These results demonstrate the practicality of virtualization in high performance computing. Palacios is an open source project, and has been downloaded over 1,000 times as of July 2009. This talk will focus on the Palacios VMM and its use in HPC environments.

I will also describe symbiotic virtualization and its integration inside the Palacios VMM. Symbiotic virtualization is a new approach to the interface between the guest OS and the VMM in which the OS makes a high-level semantically rich interface available to a VMM. The VMM can use the interface to asynchronously or synchronously acquire information from the guest. This interface is not required for correct functionality, nor does it prevent an OS from running on physical hardware. When the symbiotic interface is present, a VMM can leverage the information it provides to enable optimizations and services that are difficult or impossible to implement given current VMM/guest structures. I will show an example optimization, SwapBypass, that is enabled by symbiotic virtualization, and describe my current and future efforts in applying symbiotic virtualization to HPC and other environments.