

Virtual Supercomputer for HPC and HTC

Georg Birkenheuer, André Brinkmann, Hubert Dömer, Sascha Effert,
Christoph Konersmann, Oliver Niehörster, and Jens Simon

Paderborn Center for Parallel Computing PC²
University of Paderborn, Germany
{birke, brinkman, homerj, fermat, c_k, nieh, jens}@upb.de

Abstract Virtualization is a well-known approach in modern data centers to share computing power of servers and therewith to increase energy efficiency and ecological and economical performance. The reason for the great success of virtualization is the high computational capacity of current hardware architectures, which typically have a low workload of 15 percent for a single application and mostly consume idle cycles and wait for queries. Nevertheless, physical servers with a workload of 70 percent still have acceptable response times and a number of virtual machines can be executed on a single physical server.

Based on a different load profile, the use of virtualization technology in high performance computing (HPC) and high throughput computing (HTC) is still very uncommon. HPC and HTC applications are usually not spending idle cycles, but have computing intensive applications to calculate. Thus the computing nodes have a very high workload even with a single application.

Nevertheless, in both scenarios, the usage of virtualization might be very beneficial. Users of HPC and HTC applications often suffer from a huge management burden to transfer their applications between different cluster environments with different operating systems versions and libraries. Based on the provision of virtualized images, this management burden can be significantly reduced. Furthermore, at least in HTC, it might be beneficial to use virtualization technology to improve overall throughput and to interleave IO intensive schedules with compute intensive schedules.

Inside this paper, we evaluate the influence of the virtualization solutions VMWare, Xen, KVM, and VirtualBox on HPC and HTP computing scenarios. Therefore, we have measured and analyzed their impact on computing and IO performance for different processors. In contrast to standard evaluations, we focus on typical HPC and HTP settings, where only a single virtual server is running on each hosting physical server.