

Performance Improvement on Memory Virtualization

Switching between SPT and TDP& Simplified TDP with Large Tables

Yu Zhang

October 20, 2015

Abstract

Among the performance bottlenecks for virtual machine, memory comes next to the I/O as the second major source of overhead to be addressed. Although the current two widely adopted mechanisms - SPT (shadow page table) and TDP (two dimensional paging) have proved to be quite effective and mature in memory virtualization, each of them have their drawbacks. SPT is software-based, more portable, but suffers performance overhead due to frequent `vmexit` upon guest page fault. TDP eliminated `vmexit` caused by guest page fault, but incurs overhead due to the multilevel page table walking in the second dimension of paging process. Even with TLB the performance gain may still be offset by this overhead in a certain case. The fact is, each method is only capable of handling a subset of the workloads, and neither is guaranteed to deliver the best performance for all types of the workloads, especially considering that the performance of HPC workloads is more sensitive to the virtual than to the native execution environment. For this, we propose two solutions: 1. Dynamic Switching between SPT and TDP, and 2. A Simplified TDP with Large Tables. The first solution is able to combine the best qualities of the two mechanisms by dynamically switching between the two as a response to the ever-changing memory access behavior of the workload. The second tries to avoid the negative aspects of both SPT and TDP by modifying the operation of TDP only in the software side of the hypervisor.