A Comparison of Software and Hardware Techniques for x86 Virtualization

K. Adams and O. Agesen

Presented in *Monday, Jan. 12*

Summered by *Lei Xia (lxia@northwestern.edu)*

The authors compared performance of software VMM with new designed VMM that employs recent hardware extensions support. The experiments result different from the intuition, hardware VMM fails to provide a certain performance enhancement. The reasons cause this situation are analyzed in the paper.

The experiments from paper indicate that, the currently emerged hardware virtualization extension support fails to provide a certain performance enhancement to current VMM software techniques. And this coming from two reasons, first it lacking of MMU virtualization support, the other is this hardware extension can not work co-exist with current software virtualization techniques very well. The paper designed a systematical performance experiments, and thereafter gives an comprehensive comparison result and also a reasonable analysis on the relation between software VMM and hardware VMM. The paper contributes a good suggestions to the future research on VMM design.

In the io performance comparison, the result shows that hardware VMM io performance is far worse than software VMM, in which the authors did not specify explicitly which technique the software VMM used, naive emulated I/O or paravirtualized I/O, with or without any optimizations? Also, recent emerged self-virtualized devices can avoid most of the software virtualization work by providing hardware self-virtualization support, which make the virtualized IOs done in near native performance. Moreover, the paper takes the binary translation techniques as software VMM to compare with hardware. Fully BT is a main technique used in VMM, however, there are still many other techniques widely used to provide software virtualization, like these used in para-virtualization.

As the virtualization techniques gaining more and more attention from industry, more hardware manufactures are heading to this area, such as the hardware MMU support from Intel VT-d and AMD nested paging (although the latter one still not performance very good as expected due to its long walk way of address translation). And also, recently hardware virtualization supports are added to some performance-sensitive devices, like some gigabyte network card, which makes virtualized io work potentially at native performance. The hardware virtualization support is the future trend in this area. Although, as the authors in this paper suggested, the performance of VMM should not rely merely on the hardware. Without co-exist software VMM techniques with this hardware extensions, we can not gain too much from hardware.