

Introduction to Computer Architecture

Assignment 3 - Solution

Due May 13, 2014

1. [100 = 20 + 20 + 30 + 30] Virtual machines can lose performance from a number of events, such as the execution of privileged instructions, TLB misses, traps, and I/O. These events are usually handled in system code. Thus, one way of estimating the slowdown when running under a VM is the percentage of application execution time in a system versus user mode. For example, an application spending 10% of its execution in a system mode might slow down by 60% when running on a VM. Figure 2.32 lists the early performance of various system calls under native execution, pure virtualization, and paravirtualization for LMBench using Xen on an Itanium system with times measured in microseconds (courtesy of Matthew Chapman of the University of New South Wales).

- What types of programs would be expected to have smaller slowdowns when running under VMs?
- If slowdowns were linear as a function of system time, given the slowdown above, how much slower would a program spending 20% of its execution in system time be expected to run?
- What is the median slowdown of the system calls in the table above under pure virtualization and paravirtualization/
- Which functions in the table above have the largest slowdowns? What do you think the cause of this could be?

Benchmark	Native	Pure	Para
Null call	0.04	0.96	0.50
Null I/O	0.27	6.32	2.91
Stat	1.10	10.69	4.14
Open/close	1.99	20.43	7.71
Install sighandler	0.33	7.34	2.89
Handle signal	1.69	19.26	2.36
Fork	56.00	513.00	164.00
Exec	316.00	2084.00	578.00
Fork + exec sh	1451.00	7790.00	2360.00

Solution

- a. Programs that do a lot of computation but have small memory working sets and do little I/O or other system calls.
- b. The slowdown above was 60% for 10%, so 20% system time would run 120% slower.
- c. The median slowdown using pure virtualization is 10.3, while for para virtualization the median slowdown is 3.76.
- d. The null call and null I/O call have the largest slowdown. These have no real work to outweigh the virtualization overhead of changing protection levels, so they have the largest slowdowns.