Chapter 5.13

4.

The operation sequence for shadow page table and nested page table are

1. Shadow page table:

(1) Virtual Machine creates page table and hypervisor updates shadow table
(2) Nothing happens
(3) Hypervisor intercepts page fault, creates new mapping and invalidates the old mapping in TLB
(4) Virtual machine notifies the hypervisor to invalidate the process's TLB entries

2. Nested page table

(1) Virtual Machine creates new page table and hypervisor adds new mappings in Physical address to machine address table
(2) Hardware walks both page tables to translate Virtual address to machine address
(3) virtual machine and hypervisor update their page tables, hypervisor invalidates stale TLB entries
(4) Virtual machine notifies the hypervisor to indicate the process's TLB entries

2 of

Q. 13

step 1

Given that the number of page table levels = 4

The number of memory references is needed to serve a TLB miss for nested page table =

$L \times (L+2)$ (where L is the levels of page table

step 2

The number of memory references are needed to serve a TLB miss for native page table =

4 (since the four-level page table)

No. of memory references are needed to serve a TLB miss for nested page table

$= 4 \times (4+2)$

$= 24$

(3) of S. 13

The page fault rate is more important metric for shadow page table

TLB miss rate is the more important metric for nested page table

(5)

Combining multiple page table modifications can be used to reduce page table shadowing induced overhead

(6)

NPT caching is used to reduce NPT induced overhead. It is similar to TLB caching