

Virtualization

An overview of the hardware features that give rise to world-class virtualization and cloud computing

Robin Maffeo

AMD Alliance Manager for Microsoft DevDiv



What is Virtualization?

Creating an abstraction around real hardware to present a (transparent) virtual CPU and platform to Guest OSes

ie, a Hypervisor

Hypervisors must virtualize CPU, memory mapping, devices, time, etc.

What does virtualizing mean? Making a resource that participates properly in a virtual system, by some combination of hypervisor control, through state-swapping or emulation



Why Virtualize?

Efficiency (server consolidation) – pool underutilized machines without sacrificing performance

Power and cooling

Heterogeneous machine abstraction

Manageability, security

Development and test environments



Hardware Support for Virtualization (AMD-V)

Provides context switch mechanism (VMRUN and VMEXIT) to jump from Hypervisor to guest and back. VMRUN saves state and loads guest state, including segments and control registers

Hypervisor timeslices between guests using VMRUN/VMEXIT, but these are large context swaps

Provides interception for instructions such as CPUID and RDTSC

Provides CR3 interception for Shadow Page Table support

In shipping Barcelona/Family 10h processors, AMD-V includes Nested Page Table support



Software Virtualization (Shadow Page Tables)

OSes use page tables to map virtual addresses to physical

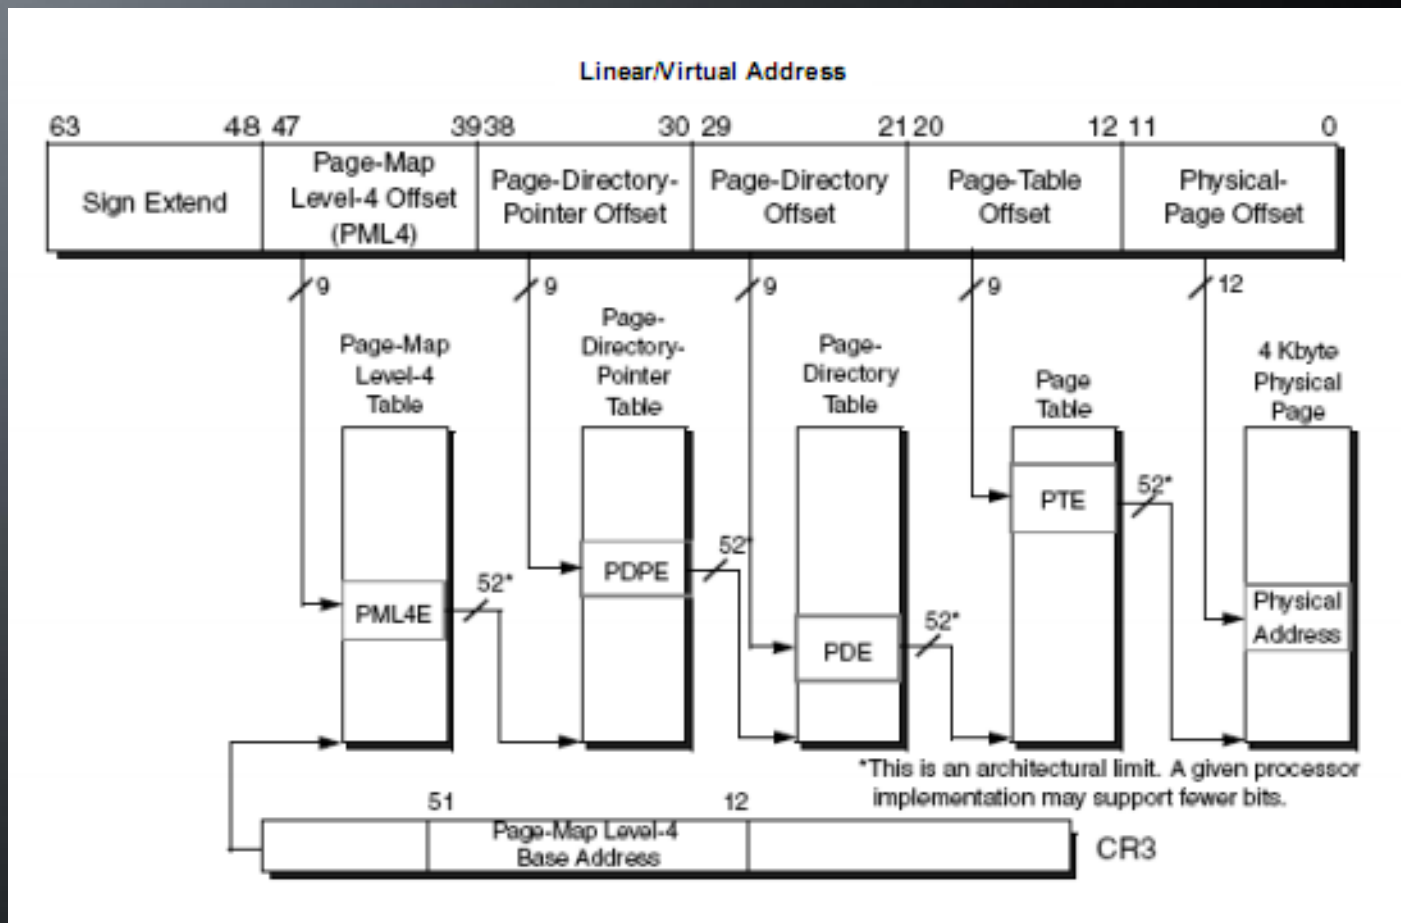
When guest is active, Hypervisor forces processor to use shadow tables to perform address translation

Hypervisor must keep track of guest page tables in order to maintain shadow page table

Since Hypervisors virtualize physical addresses, shadow copies must be kept by the Hypervisor which adds complexity, pathlength overhead due to page faults when entry is not in SPT



Normal VA to PA mapping

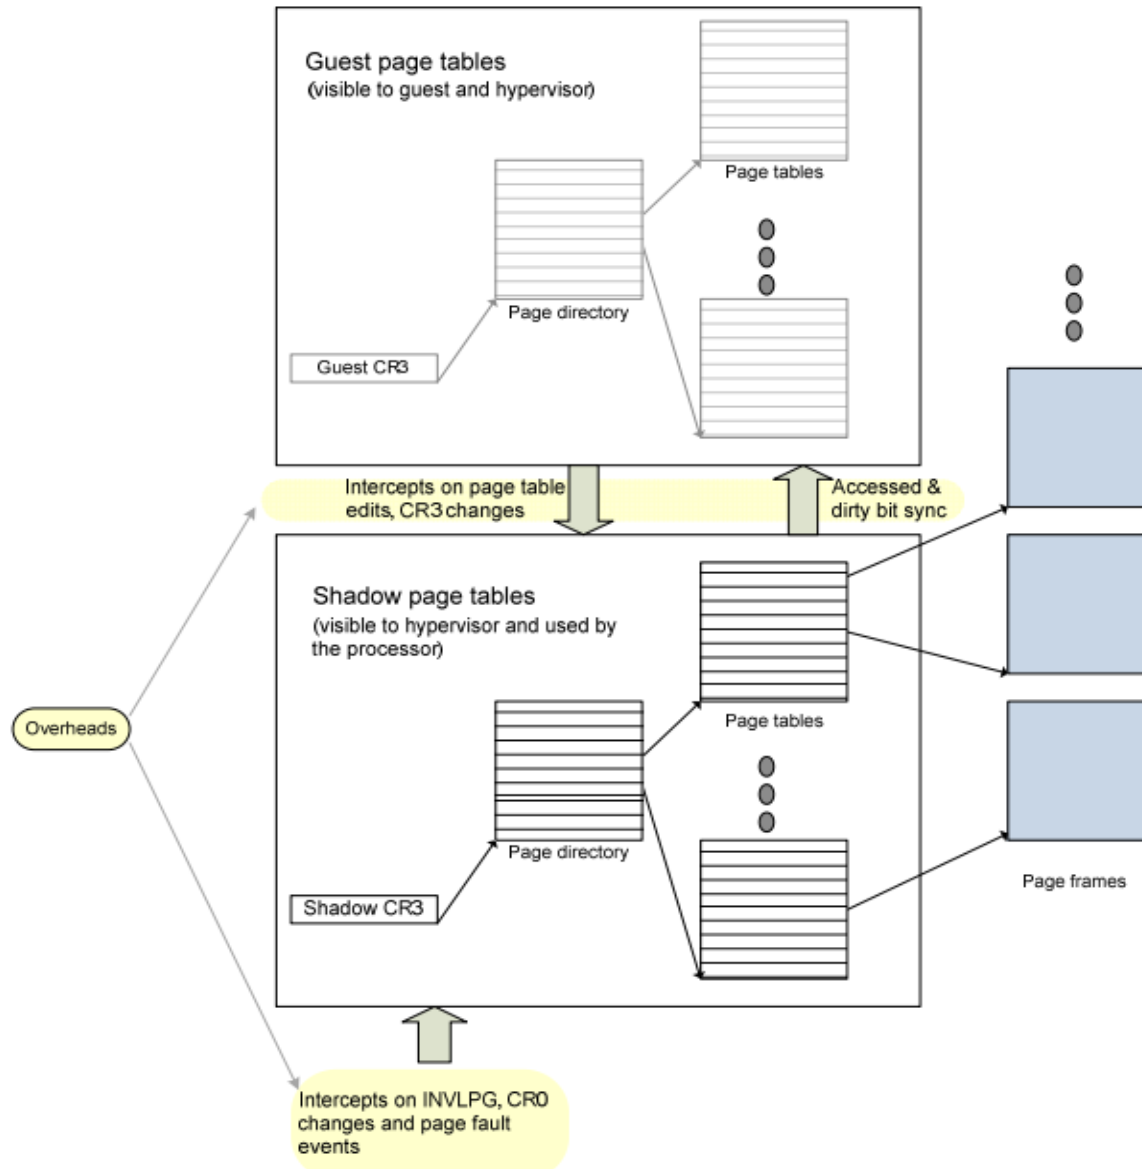


Processor contains hardware page walker to perform address translation

6 | Intro to AMD Virtualization | PDC October 2008



Guest and shadow page tables



Nested Page Tables

AKA Rapid Virtualization Indexing

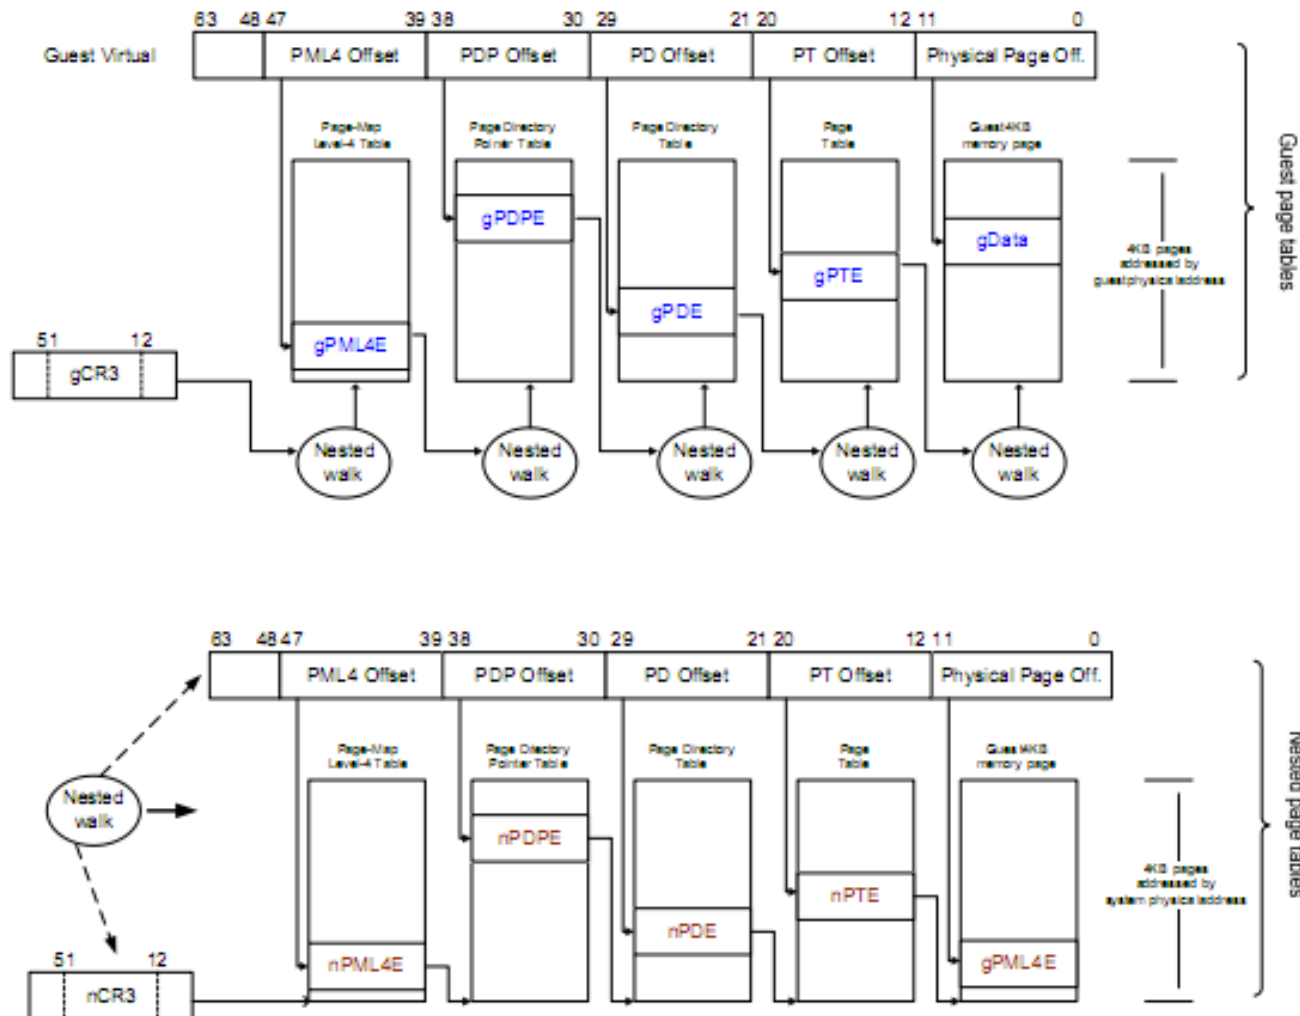
Improves performance, scaling, reliability, and memory usage

Hypervisor no longer required to maintain shadow page tables, decreasing complexity and overhead. Leaves guest in complete control of its page tables.

Hardware walks table to translate guest PA to system PA directly



Guest VA to system PA translation



Virtualization now and in the future

Azure's back-end Hypervisor supports NPT

Hyper-V in Windows Server 2008 R2 (Win7 Server) supports NPT

Based on close AMD and Microsoft collaboration

Further improvements in AMD-V

Enhanced NPT

Faster world switch times (VMEXIT/VMRUN)

Additional investments being made by AMD in hardware support for improved Hypervisor performance, reliability, and flexibility (make device access faster, for example)



Trademark Attribution

AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.

©2008 Advanced Micro Devices, Inc. All rights reserved.

