Xen and the Art of Virtualization

Paul Barham, Boris Dragovic,
Keir Fraser, Steven Hand,
Tim Harris, Alex Ho,
Rolf Neugebauer,
Ian Pratt, Andrew Warfield

Outline

- Motivation
- Overview of Xen
- CPU virtualization
- MMU virtualization
- Experimental results
- Recent Developments

Motivation

- Stronger isolation between applications
  - Using separate machines is too expensive
  - Separate processes is not sufficient
- Excess computing power
- Different OSs on the same machine

Types of Virtualization

- Hardware-level virtualization
  - Vmware, Xen
- Operating system-level virtualization
  - Jails
- High-level language virtual machines
  - Java VM

Overview of Xen

- Requires the guest OS to be ported
- Applications run without modifications
- Does not use a host OS

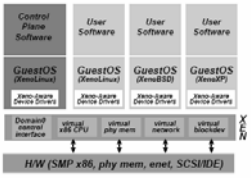

Ideal VM CPU

- Sensitive instructions cause exceptions
  - Instructions that change the machine state
  - Instructions that read or write sensitive registers/memory

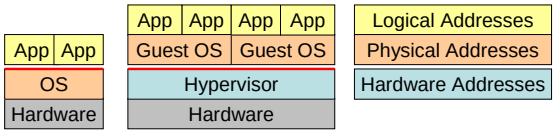

x86 CPU

- Privileged instructions can only be successfully executed from below the red line
- Some sensitive instructions are not privileged

7

VMware CPU virtualization

- Checks for sensitive instructions before execution

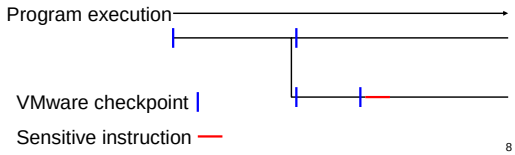

8

Xen CPU virtualization

- When the guest OS executes privileged instructions, the x86 raises exceptions
- Xen catches these exceptions
- Guest OSs directly call Xen code instead of using sensitive, unprivileged instructions

9

Ideal VM MMU

- Page translation occurs in software
- OSs provide a TLB miss handler
- Hypervisor executes guest mapping routine

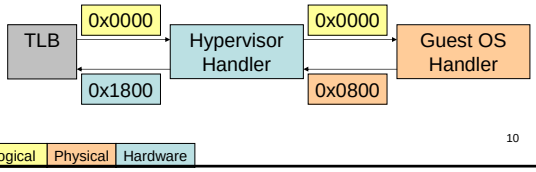

10

x86 MMU

- TLB misses are handled directly by the MMU
- OSs must create a page table that maps logical to physical addresses
  - The table must be laid out as specified by the MMU
  - The OS sets a register to point to the table

11

VMware MMU virtualization

- Maintains shadow page tables

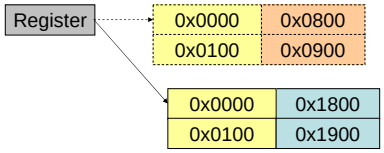

12

Xen MMU virtualization

- Xen exposes the hardware addresses to the guest OS
- The guest OS constructs a page table that maps from logical to hardware addresses
- Updates to the page table must pass through Xen

13

Experiments

- Minimal performance degradation over plain Linux

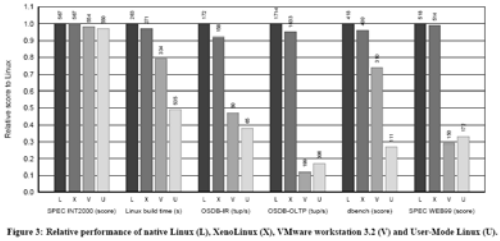

Figure 3: Relative performance of native Linux (L), XenoLinux (X), VMware workstation 3.2 (V) and User-Mode Linux (U).

14

Performance Isolation

- Prevented misbehaving guests from interfering with other guests

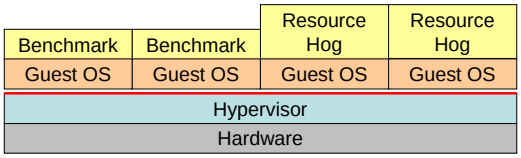

15

Recent Developments

- Many Linux distros have Xen support
- Unmodified Windows XP ran on Xen with Intel VT-enabled processors
- Blazingfast provides virtual servers using Xen

16