

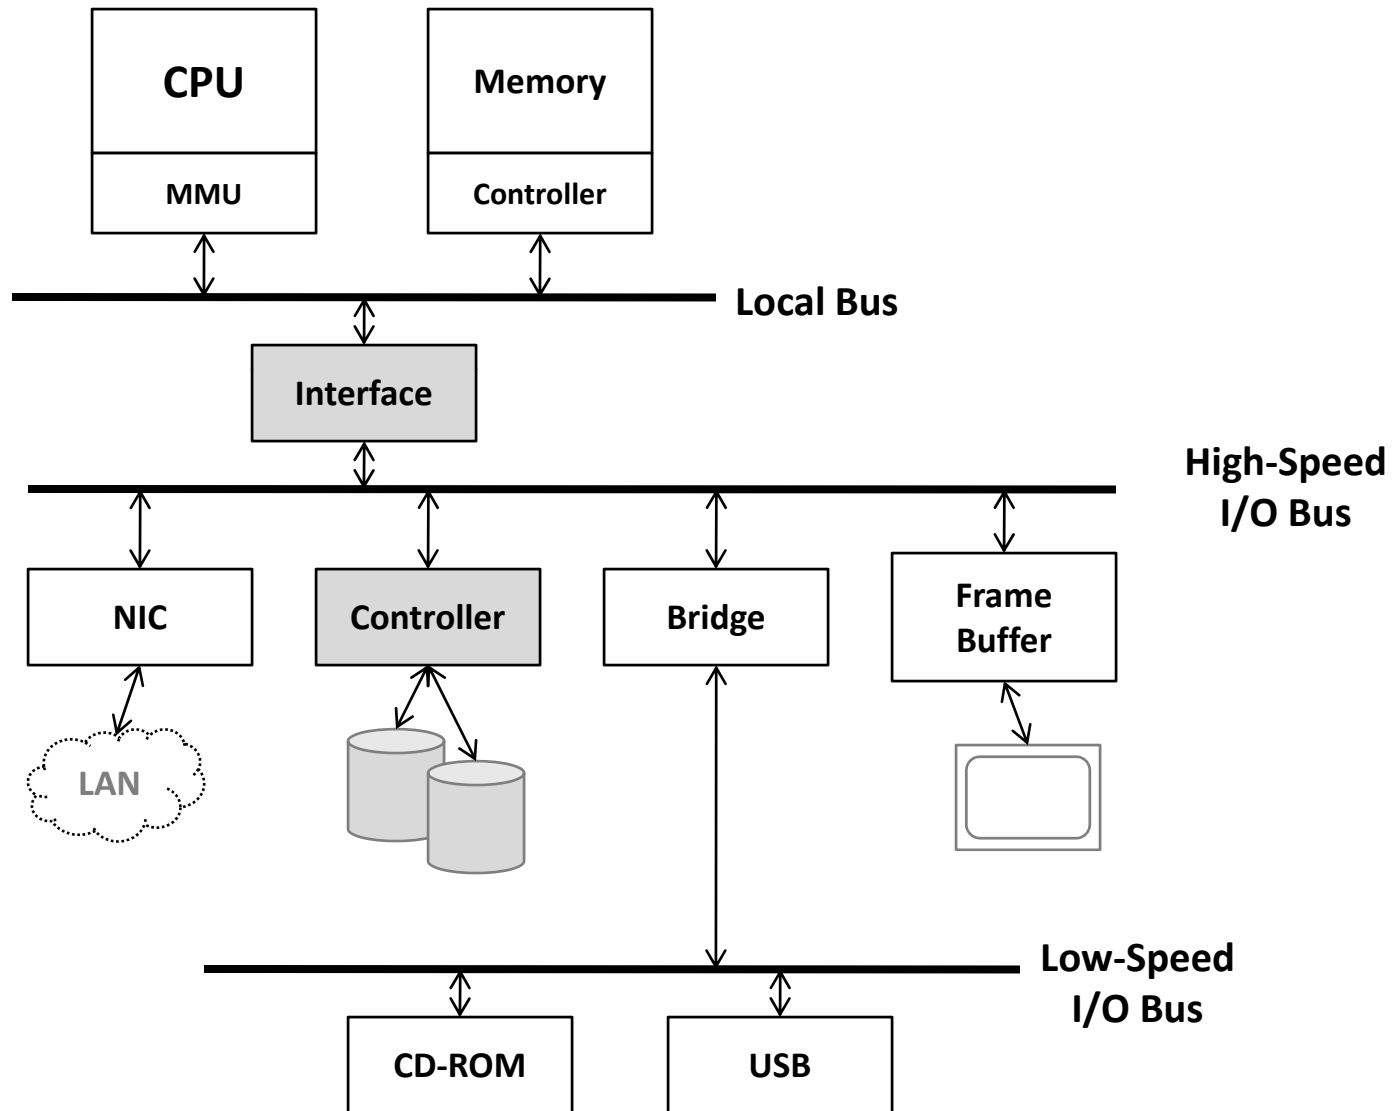
E6998 - Virtual Machines
Lecture 4
Device Virtualization

Scott Devine
VMware, Inc.

Outline

- **Types of Device Virtualization**
 - Direct Access
 - Emulated
 - Para-virtualized
- **Storage Virtualization**
 - Copy-on-Write Disks

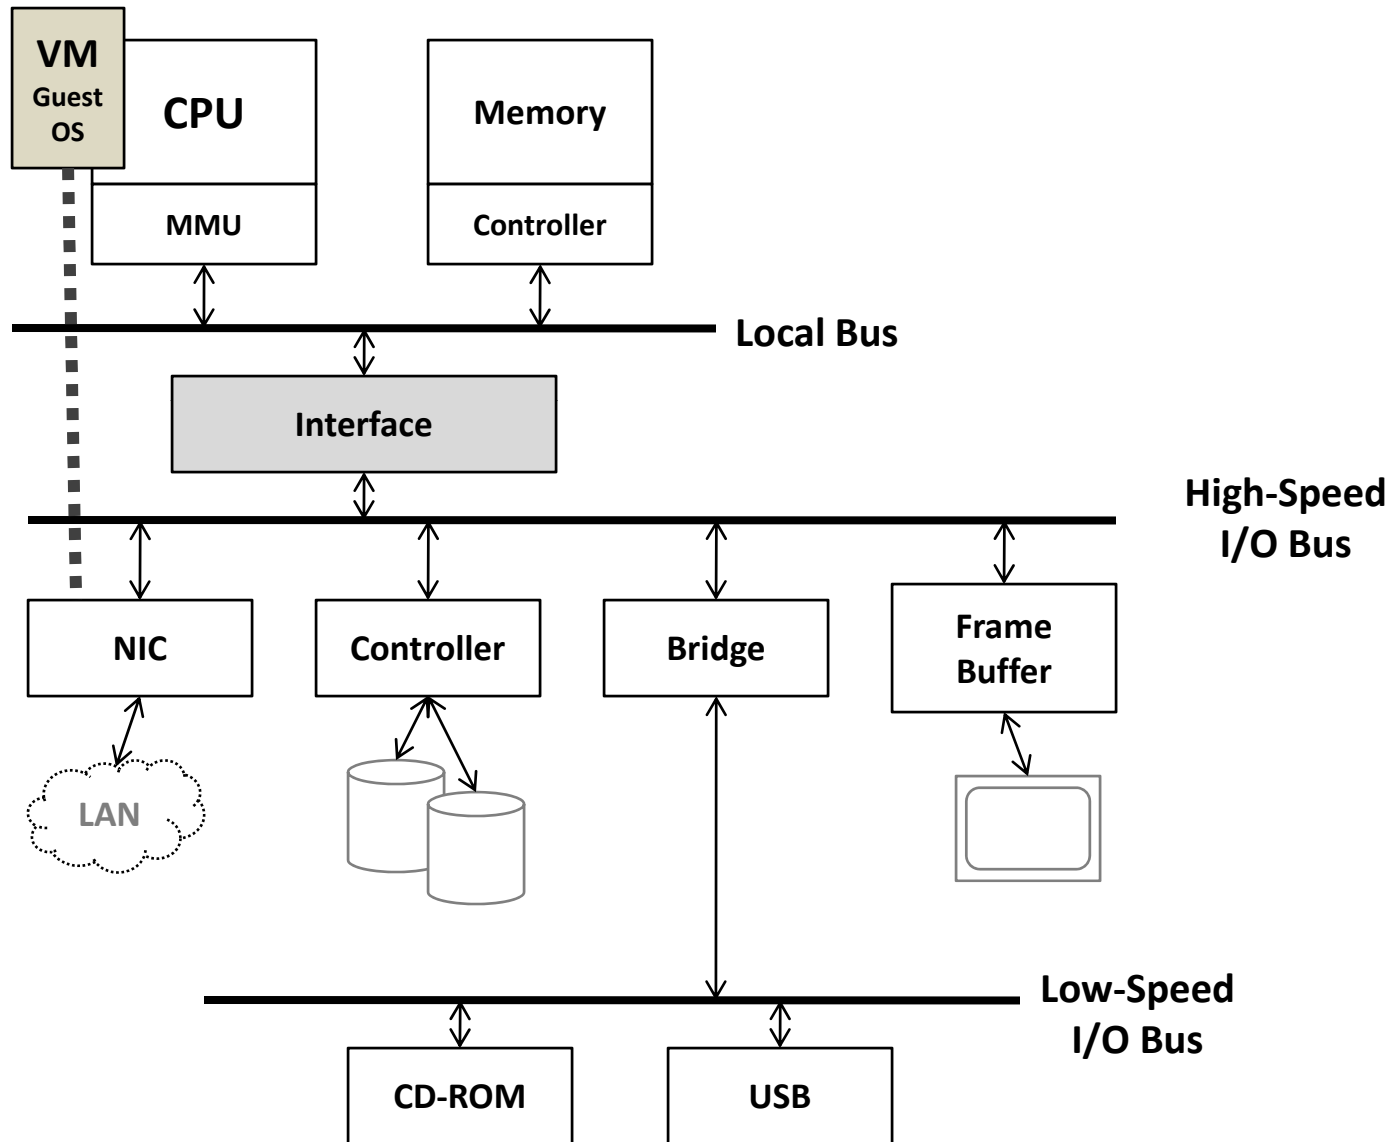
Computer System Organization



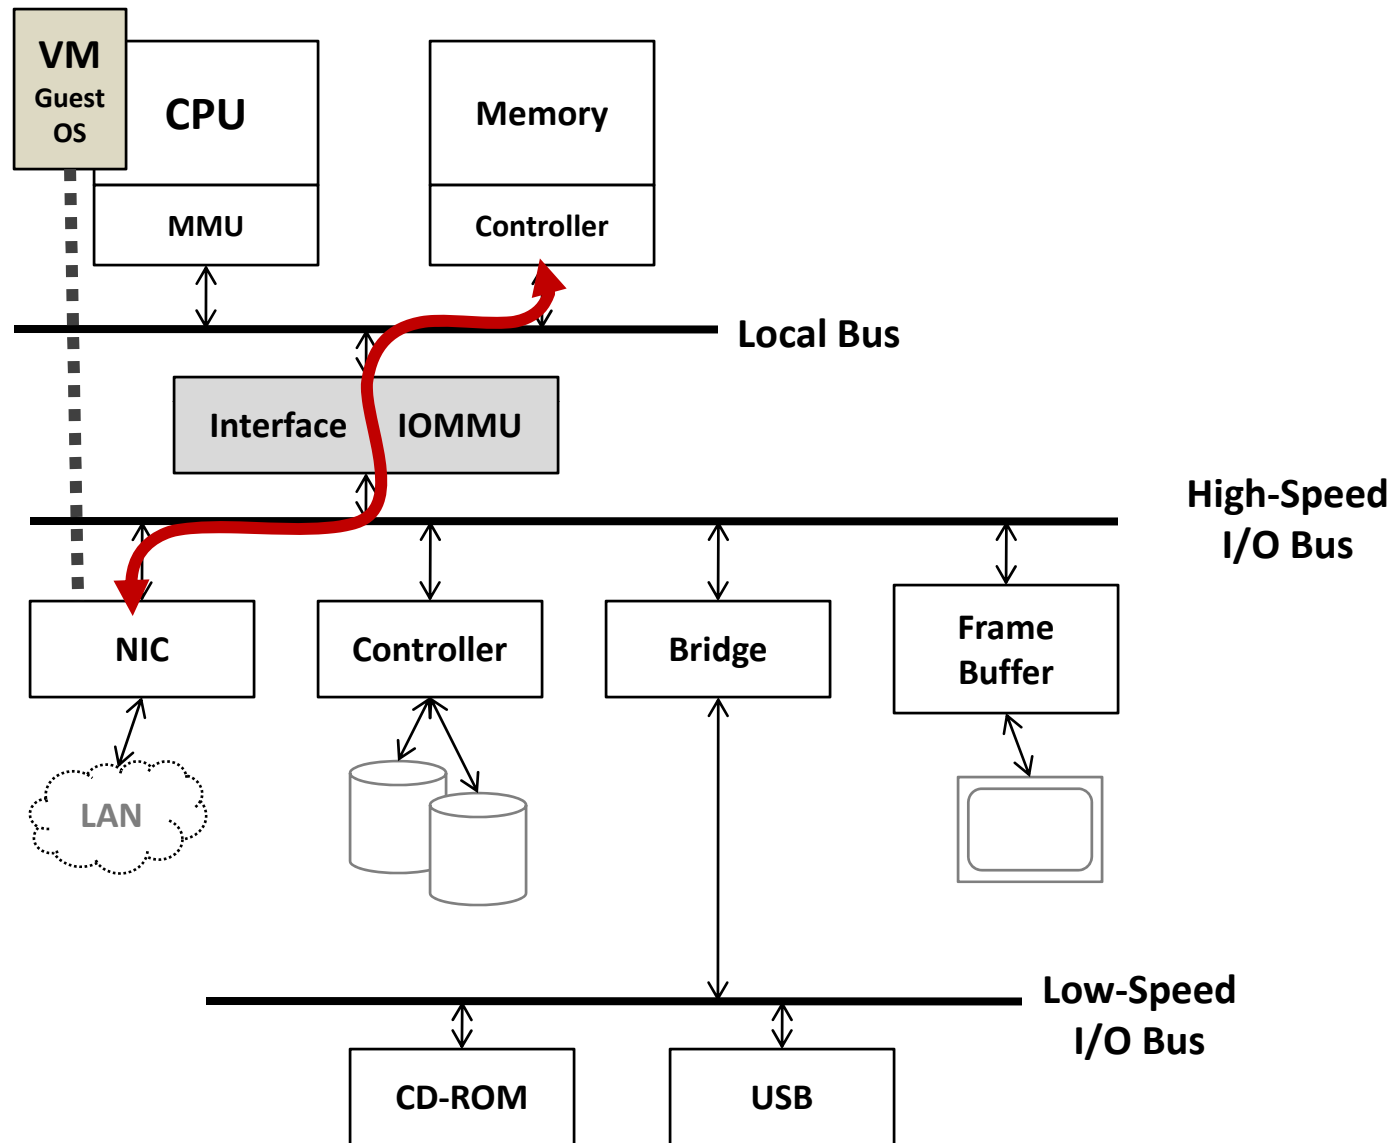
Device Virtualization

- **Goals**
 - Isolation
 - Multiplexing
 - Speed
 - Mobility
 - Interposition
- **Device Virtualization Strategies**
 - Direct Access
 - Emulation
 - Para-virtualization

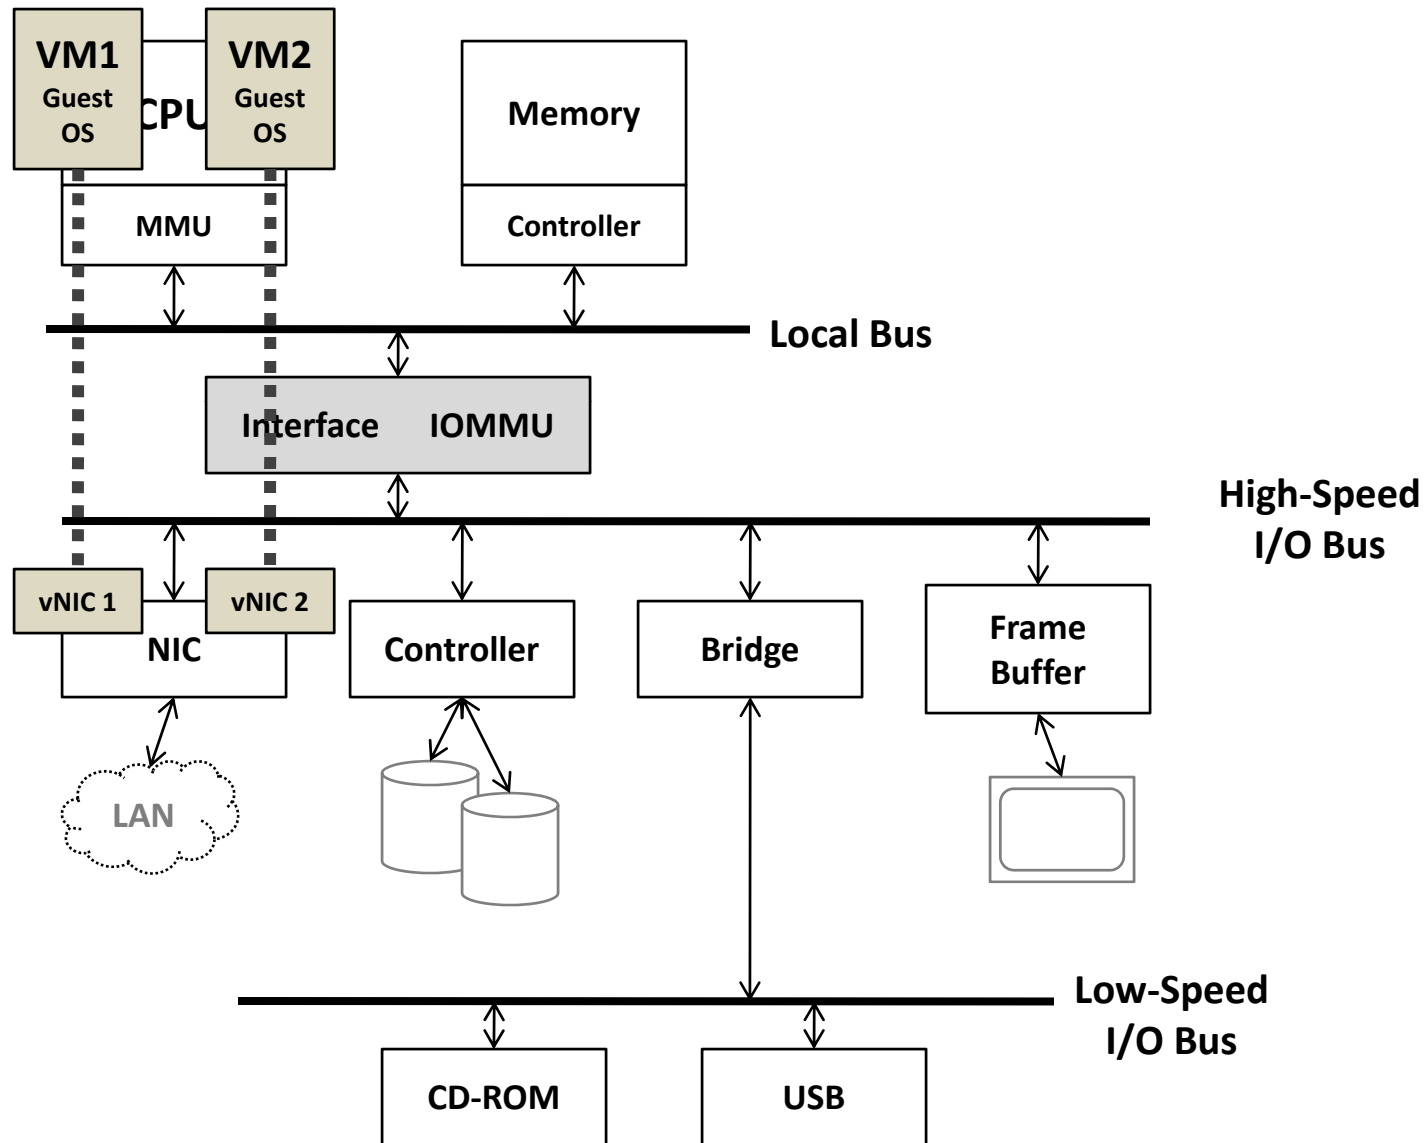
Direct Access Device



Memory Isolation w/ Direct Access Device



Virtualization Enabled Device



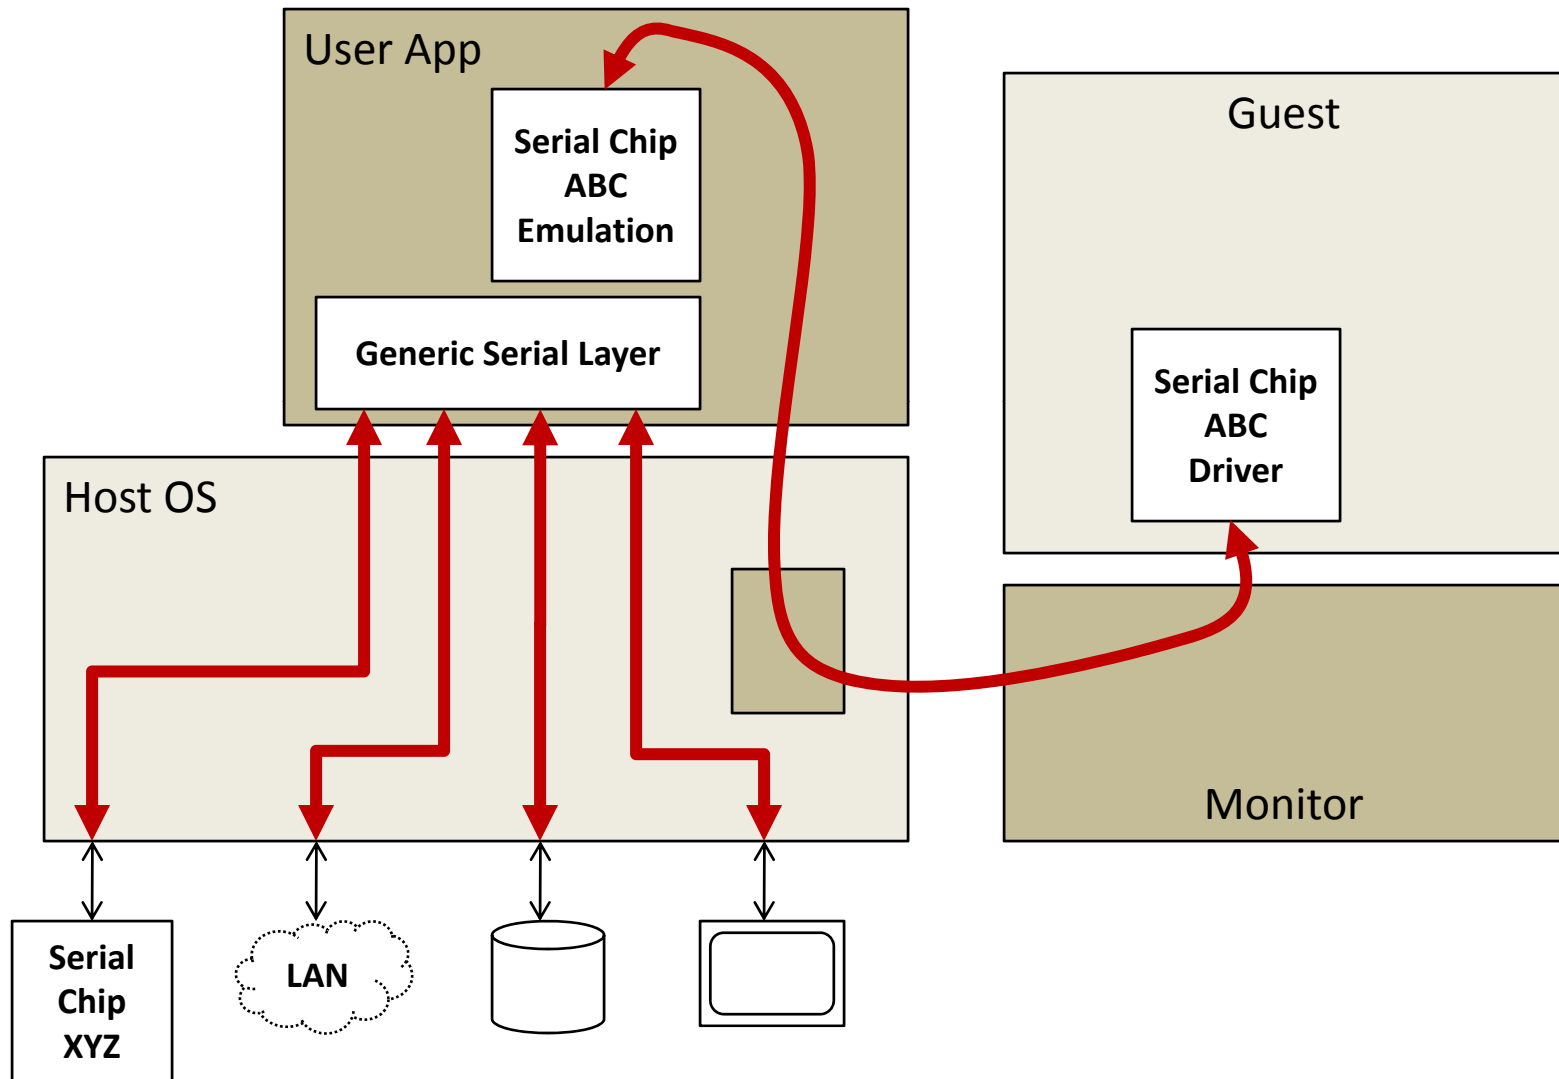
Direct Access Device Virtualization

- **Allow Guest OS direct access to underlying device**
- **Positives**
 - Fast
 - Simplify monitor
 - Limited device drivers needed
- **Negatives**
 - Need hardware support for safety (IOMMU)
 - Need hardware support for multiplexing
 - Hardware interface visible to guest
 - Limits mobility of VM
 - Interposition hard by definition

Emulated Devices

- **Emulate a device in class**
 - Emulated registers
 - Memory mapped I/O or programmed I/O
- **Convert**
 - Intermediate representation
- **Back-ends per real device**

Serial Port Example



Emulated Devices

- **Positives**

- Platform stability
- Allows interposition
- No special hardware support needed
 - Isolation, multiplexing implemented by monitor


- **Negatives**

- Can be slow
- Drivers needed in monitor or host


Para-Virtualized Devices

- **Guest passes requests to Monitor at a higher abstraction level**
 - Monitor calls made to initiate requests
 - Buffers shared between guest / monitor
- **Positives**
 - Simplify monitor
 - Fast
- **Negatives**
 - Monitor needs to supply guest-specific drivers
 - Bootstrapping issues

Traditional Architecture



Hosted Monitor Architecture



VMware ESX 2.0

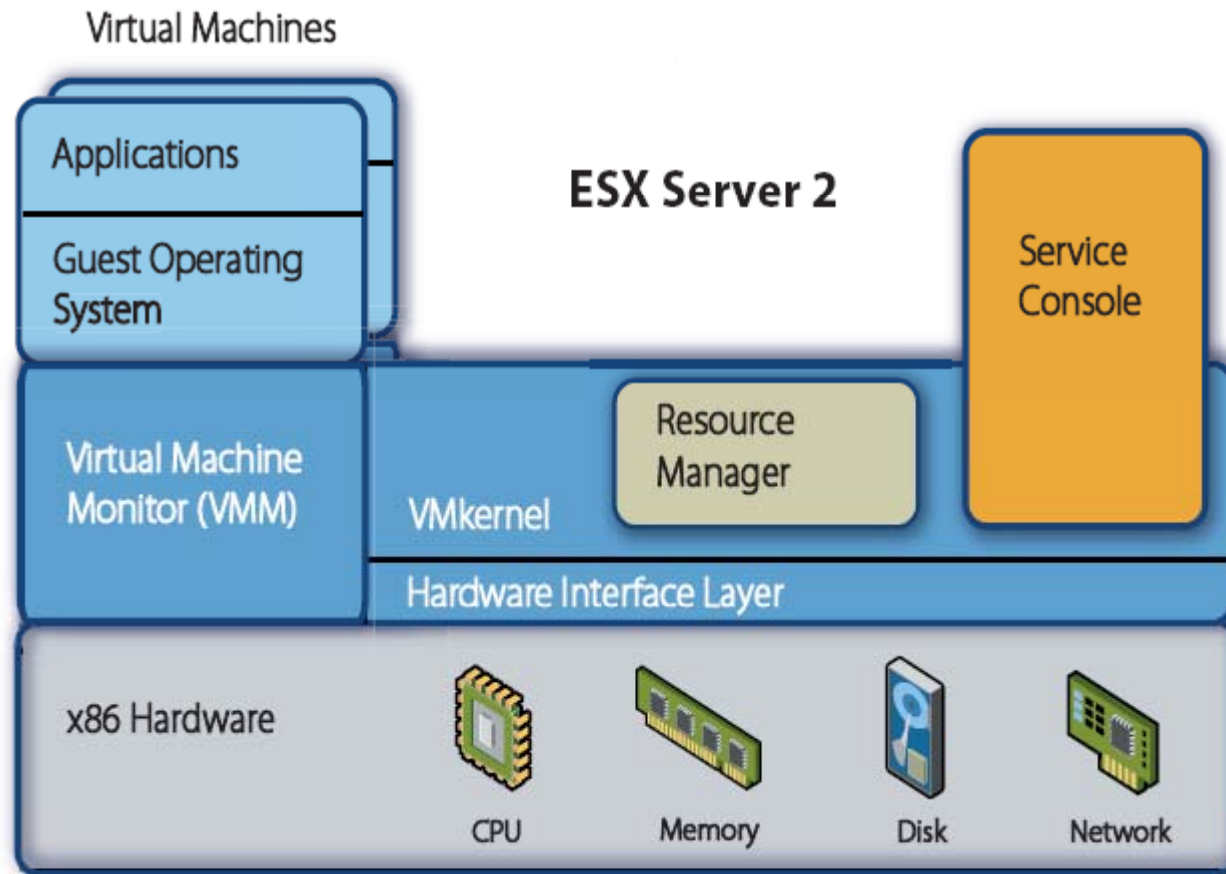
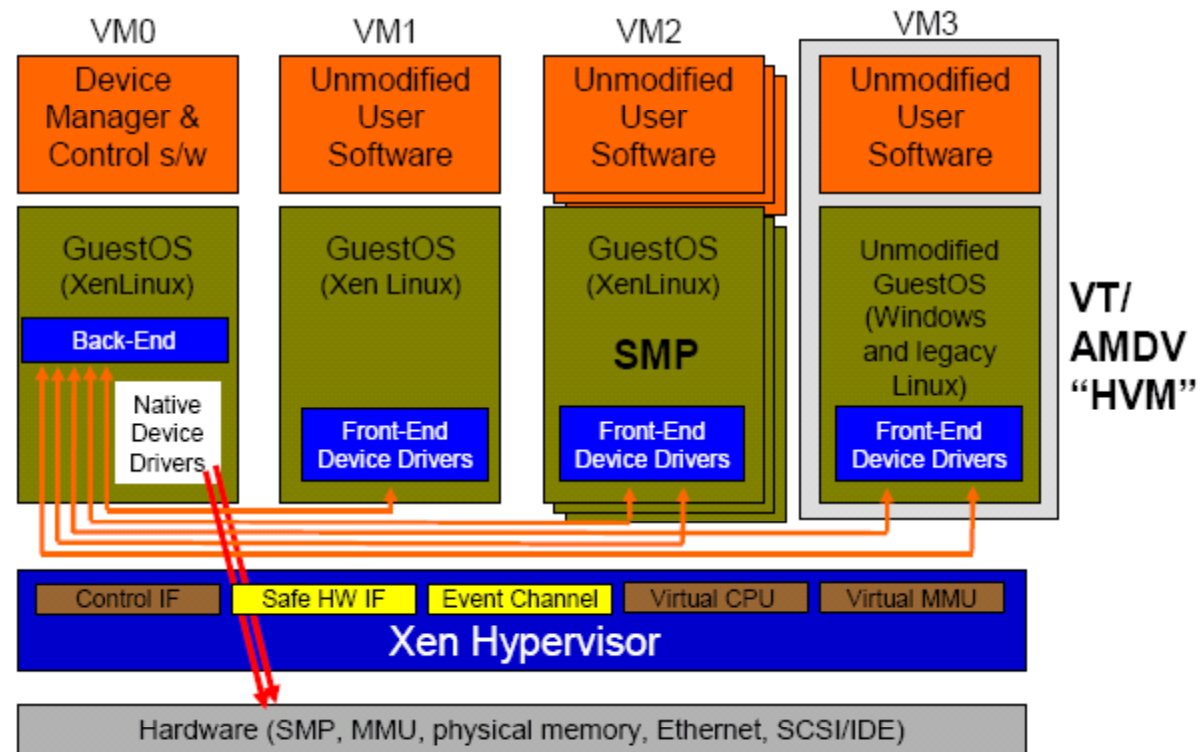


Figure 1: ESX Server architecture

Source: http://www.vmware.com/pdf/esx2_performance_implications.pdf

Hybrid Ex 2 - Xen 3.0

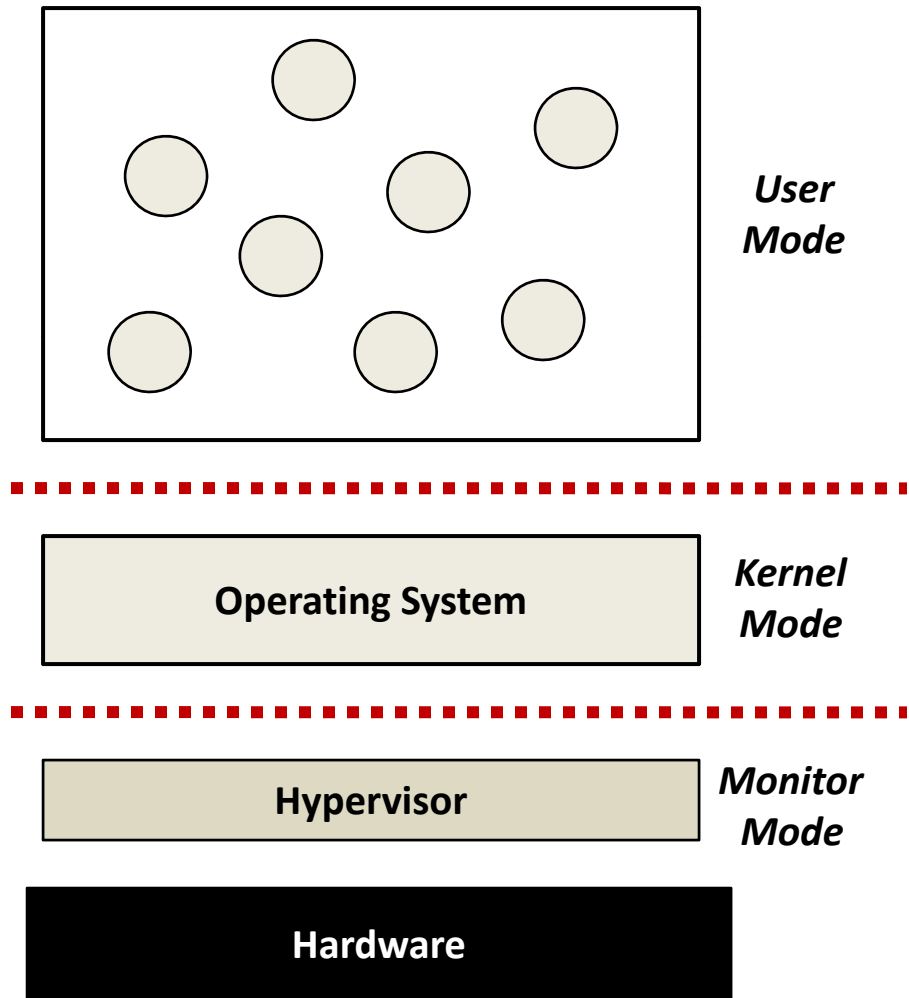
- **Para-virtualization**
 - Linux Guest
- **Hardware-supported virtualization**
 - Unmodified Windows
- **Isolated Device Drivers**



[Source: Ottawa Linux Symposium 2006 presentation.](http://www.cl.cam.ac.uk/netos/papers/)
<http://www.cl.cam.ac.uk/netos/papers/>

Hypervisor


- **Hardware-supported single-use monitor**
- **Characteristics**
 - Small size
 - Runs in a special hardware mode
 - Guest OS runs in normal privilege level
- **Uses**
 - Security
 - System management
 - Fault tolerance




Virtual Disks

- **File on file system**
 - Host file system
 - Hypervisor file system
- **Meta-data in file**
 - Maps disk blocks to file offsets
 - Flat file
 - Indexed file
 - Allows disk to grow on demand

Copy-on-Write Disks



Copy-on-Write Disks



Non-persistent Execution

