Server Virtualization

Raj Jain
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@cse.wustl.edu

These slides and audio/video recordings of this class lecture are at:

http://www.cse.wustl.edu/~jain/cse570-13/

1. Why Virtualize?
2. Server Virtualization Concepts
3. Virtualization Products
4. A Sample of Virtualization Products
5. Open Virtualization Format (OVF)

Note: Storage Virtualization and Network Virtualization will be discussed in subsequent lectures

Virtualization

"*Virtualization means that Applications can use a resource without any concern for where it resides, what the technical interface is, how it has been implemented, which platform it uses, and how much of it is available.*"

-Rick F. Van der Lans
in Data Virtualization for Business Intelligence Systems

5 Reasons to Virtualize

1. Sharing: Break up a large resource
   Large Capacity or high-speed
   E.g., Servers
2. Isolation: Protection from other tenants
   E.g., Virtual Private Network
3. Aggregating: Combine many resources in to one, e.g., storage
4. Dynamics: Fast allocation, Change/Mobility, load balancing, e.g., virtual machines
5. Ease of Management ⇒ Easy distribution, deployment, testing

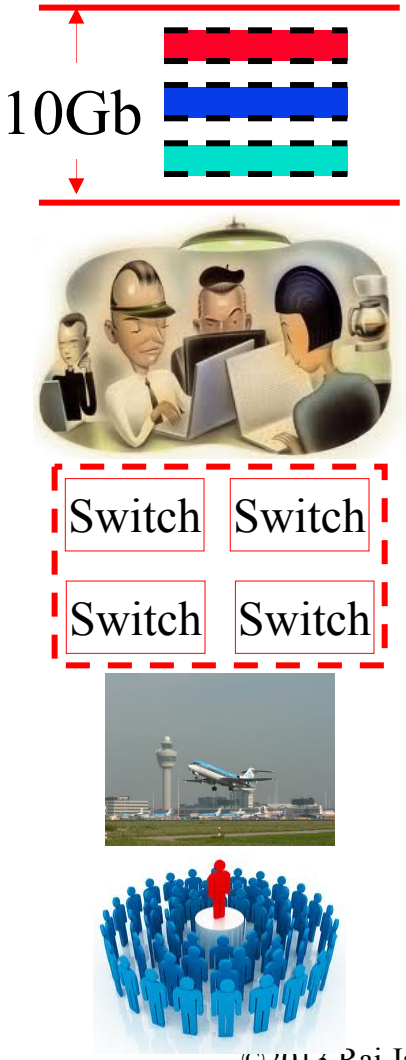

Advantages of Virtualization

- Minimize hardware costs (CapEx)
  Multiple virtual servers on one physical hardware
- Easily move VMs to other data centers
  - Provide disaster recovery. Hardware maintenance.
  - Follow the sun (active users) or follow the moon (cheap power)
- Consolidate idle workloads. Usage is bursty and asynchronous.
  Increase device utilization
- Conserve power
  Free up unused physical resources
- Easier automation (Lower OpEx)
  Simplified provisioning/administration of hardware and software
- Scalability and Flexibility: Multiple operating systems

Ref: http://en.wikipedia.org/wiki/Platform_virtualization
Ref: K. Hess, A. Newman, "Practical Virtualization Solutions: Virtualization from the Trenches," Prentice Hall, 2009, ISBN:0137142978

Virtualization in Computing

- **Storage**:
  - Virtual Memory ⇒ L1, L2, L3, ... ⇒ Recursive
  - Virtual CDs, Virtual Disks (RAID), Cloud storage
- **Computing**:
  - Virtual Desktop ⇒ Virtual Server ⇒ Virtual Datacenter

    Thin Client ⇒ VMs ⇒ Cloud
- **Networking**: Plumbing of computing
  - Virtual Channels, Virtual LANs, Virtual Private Networks

Server Virtualization Concepts

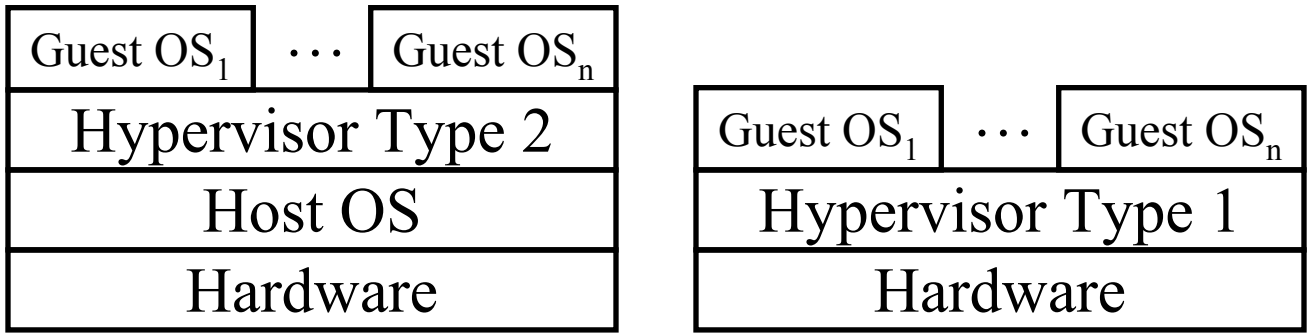

- Host OS: Runs on the bare metal
- Guest OS: Runs on the host OS,
  e.g., Windows XP Mode on Win 7
- Hypervisor: Software to support multiple virtual machines
  - Type 1: Runs on bare metal, e.g., Xen, VMware ESXi
  - Type 2: Runs on a host OS, e.g., MS Virtual PC
  - Type 0: Both 1 and 2, e.g., Linux KVM

Ref: http://en.wikipedia.org/wiki/Hypervisor

Levels of Virtualization

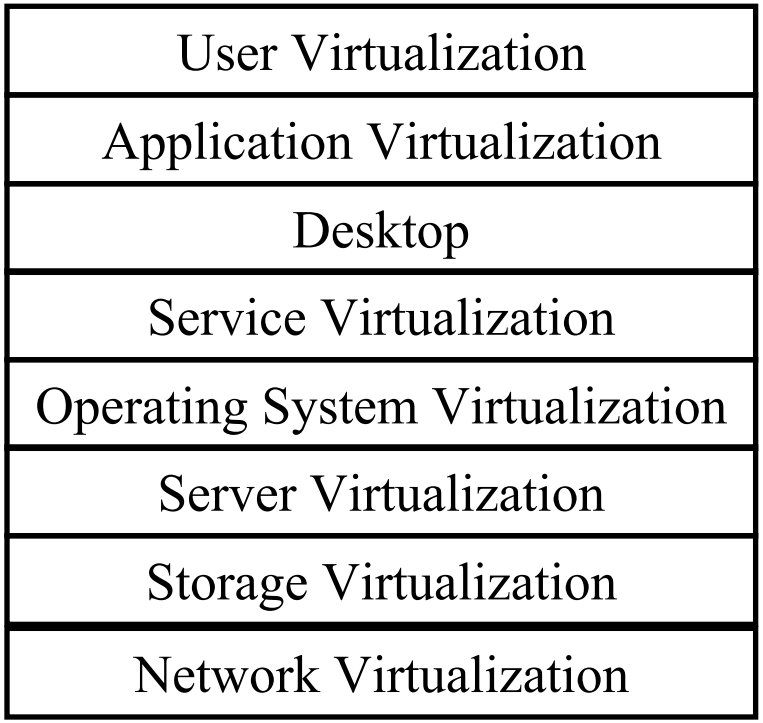

- Network and Storage virtualization will be discussed in later modules of this course.

Operating System-Level Virtualization

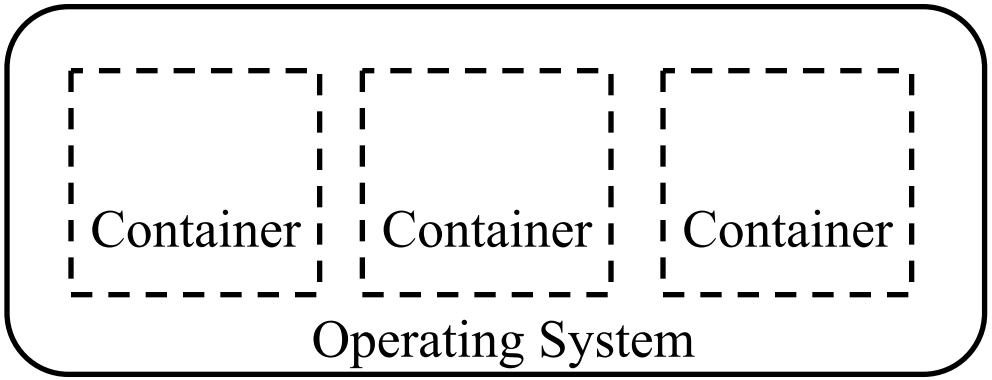

- Multiple isolated user spaces on the same operating system
- Example: Windows Server 2003, 2008, 2012, …
- Multiple users can remotely login and use the system. Only one operating system instance
- Guest operating system is similar or same as the host, E.g., Windows on Windows
- Applications of one user cannot affect other users aka Jails => Can be used to run suspect software E.g., iCore Virtual Accounts and Sandboxie

Ref: http://en.wikipedia.org/wiki/Operating_system-level_virtualization

Desktop Virtualization

- A Desktop system with multiple operating systems
- Example: Mac OS X and Windows at the same time
  Parallels Desktop for Mac
- Hypervisor type 1 similar to server virtualization
- Useful for testing software on multiple OS
- Reduced hardware cost
- This is local desktop virtualization

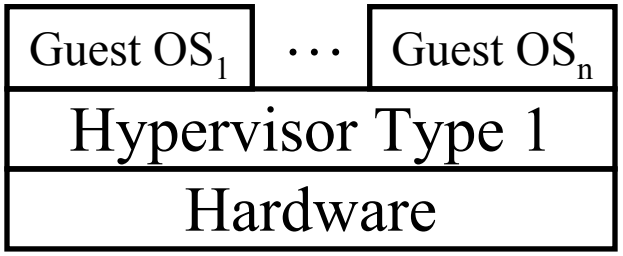

Thin Client

- Also called Remote Desktop Virtualization or Virtual Desktop Infrastructure (VDI) or Access Virtualization
- Remote server provides most of the computing and storage resources
- All programs and data is stored on remote servers
- Local computer is simple with limited resources
- Easy maintenance and upgrades
- Users can log in on any client
- Significant energy savings
- Example: Chrome books
- Zero Client: Power cable only. USB and communication via WiFi

Ref: http://en.wikipedia.org/wiki/Desktop_virtualization
Ref: http://en.wikipedia.org/wiki/Thin_client

Application Virtualization

- Allows an application to run on many different operating systems and hardware
- Application byte code can run on different processors and operating systems usually using an interpreter or just-in-time (run-time) compilation
- Examples: Java Virtual Machine (JVM)

Ref: http://en.wikipedia.org/wiki/Comparison_of_application_virtual_machines

Service Virtualization

- Service is similar to but different from applications in that they are hidden from the end user
  Example services: Firewall, load balancer, Proxy Server
- Service virtualization: Virtual appliances implemented in software
- A load balancer's IP address is advertised outside.
  It distributes the load to multiple servers.

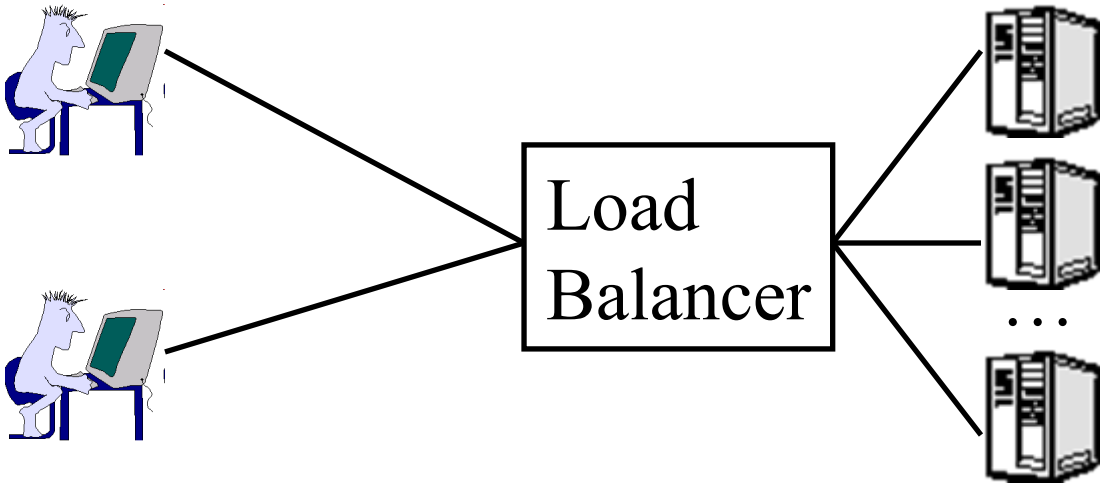

User Experience Virtualization

- Aka User Virtualization
- Users see the same application interface regardless of the device: laptop, tablet, Smart Phone
- Users can roam from device to device
- User profiles and application settings are stored in a central data center or cloud
- Example: Microsoft UE-V

Ref: http://en.wikipedia.org/wiki/User_virtualization

Related Concepts

- Hardware Emulation
- Hardware Assisted Virtualization
- Parallel Virtual Machines
- Paravirtualization

Hardware Emulation

- Mimic a different type of hardware, e.g., SPARC on PC
- Examples:
  - Microsoft Virtual PC for Mac: x86 on PowerPC
  - Bochs (pronounced Box): emulates x86 on UNIX, Linux, Windows, and Mac OS X
  - QEMU: Emulates x86, $x86_{64}$, ARM, SPARC, PowerPC, MIPS, and m68k on x86, $x86_{64}$, and PowerPC

Ref: http://en.wikipedia.org/wiki/Hardware_emulation

Hardware Assisted Virtualization

- Processors are designed to help virtualization
- For example, Intel-VT and AMD-V provide additional instructions that help virtualization
- Hypervisors can use these instructions to improve the performance
- Linux KVM, Microsoft Hyper-V, Microsoft Virtual PC, Xen, etc. use these features

Ref: http://en.wikipedia.org/wiki/Hardware-assisted_virtualization

Parallel Virtual Machines

- One computing environment running on multiple computers
- PVM is also the name of an open source software that allows running a program in parallel on multiple machines
- Now grid computing and cluster computing is generally used

Ref: http://en.wikipedia.org/wiki/Parallel_Virtual_Machine

Paravirtualization

- Hypervisor offers special APIs for operations that are difficult to run in a virtual environment
- Requires modifying the guest OS.
  Requires source code for the guest OS.
  Change and recompile for the hypervisor
- Example: paravirt-ops code in Linux provides a hypervisor agnostic interface between the hypervisor and guest kernels.

Ref: http://en.wikipedia.org/wiki/Paravirtualization

Virtualization Products

- **Microsoft Virtual PC**: Included with Windows
- **Xen**: Popular open source hypervisor from University of Cambridge. Endorsed by many manufacturers. Allows running multiple Linux-like operating systems at nearly native speeds. Citrix provides Xen extensions and support.
- Oracle VM **VirtualBox**: runs on Linux, Mac OS X, Windows XP/Vista/7/8, Solaris, OpenSolaris, FreeBSD. Free.
- **Microsoft Hyper-V**: Part of Windows 2008+ Server
- Linux **KVM** (Kernel Virtual Machine): Uses Linux kernel as a hypervisor. Each VM uses its own unique kernel. Uses processor virtualization extensions (Intel-VT and AMD-V).
- User-Mode Linux (**UML**): Uses an executable kernel and a root file system to create a VM.

Ref: http://en.wikipedia.org/wiki/Comparison_of_platform_virtual_machines
Ref: http://en.wikipedia.org/wiki/Virtualbox
Ref: http://en.wikipedia.org/wiki/ESXi

VMware ESX/ESXi

- ESX: Bare metal embedded hypervisor for server virtualization in an enterprise data center. Creates VM0 for a Linux based service consol for management of hypervisor and other VMs
  - vMotion feature in ESX allows a live VM to move to another host. A copy of memory is sent to the new location and updated until the transfer.
- ESXi: Smaller footprint version of ESX. No service console. Free download. Many features require vCenter license.
  - **vCenter** allows managing multiple servers and to run **vMotion** and **svMotion** (Storage sync on the remote location)

Ref: http://en.wikipedia.org/wiki/VMware_ESX

VMware vMotion

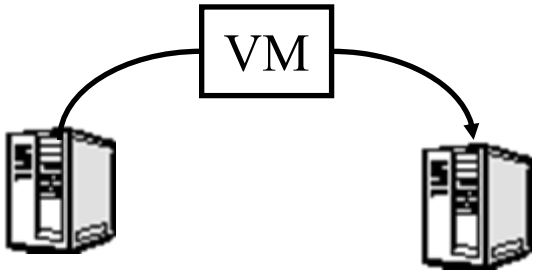

- Live migration of running VMs from one datacenter to another
  Live= <1 second interruption. Warm = few tens of seconds
- Allows VMs to move in/between datacenters
- VMware ESX servers
- VMs keep the same IP address
  ⇒ Same IP subnet/broadcast domain
- Requires LAN extension and Storage Extension (access storage from previous data center)
- Storage vMotion: Storage is moved before the VM
- Storage Caching can also be used

VMware ThinApp, View, Fusion

- **ThinApp**: Portable application creator. Virtualizes resources such as environment variables, files, and registry keys. Allows programs to run from USB keys.
- **View**: Allows thin clients.
  Multiple clients share a remote server.
  A Linux or Mac OS X client can connect to remote server.
- **Fusion**: Hypervisor for Intel-based Macs
  - Allows running Windows, Linux, Solaris, along with OS X
  - Requires Intel-VT hardware virtualization support

Ref: http://en.wikipedia.org/wiki/VMware_ThinApp
Ref: http://en.wikipedia.org/wiki/Vmware_view
Ref: http://en.wikipedia.org/wiki/Vmware_fusion

VMware Workstation and Player

- **Workstation**: Hypervisor for virtual computing.
  - Virtualizes network adapters, disk drives, USB devices, CD/DVD drives.
  - Useful for testing client-server environments.
  - Requires hardware virtualization support.
  - Many ready-made VMs are available.
- **Player**: A freeware version of VMware workstation with limited
- VMware High Availability (**HA**)
- Virtual Machine File System (**VMFS**): Cluster file system for ESX cluster

VMware vSphere

vSphere is the VMware's product suite consisting of:

- **vCenter**: Provisioning, management, and monitoring console for the cloud
- **ESX** or **ESXi**: Type 1 hypervisor that runs on bare metal
- Virtual SMP: VM running on two or more physical processors
- Cluster = collection of servers that pool CPU/memory/storage for VMs
- Pod = Collection of clusters grouped as a management unit
- Distributed Resource Scheduler (**DRS**): System for provisioning VMs and load balancing across multiple physical servers
- vNetwork Distributed Switch (**DVS**): Maintains network runtime state as VMs move.

Ref: http://en.wikipedia.org/wiki/VMware

File System for VMs

Using VMware's example:

- All physical disks are clustered in to Pools
- Each VM has a virtual disk in some pool with extension .vmdk
- The disk is formatted using Virtual Machine File System (VMFS)
- Thin Provisioning: Dynamic disk size. Occupies the minimum space required and grows as the data grows
- Thick Provisioning: Static. Full size is allocated at creation
- Linked Clone: All VMs use the same disk (generally read-only), e.g., operating system disk
- Non-Linked Disk: Separate disks for each VM (writeable)
- Persistent/Non-Persistent: Saved/Deleted on VM release. Permanent/Temporary files

Ref: J. Langone and A. leibovici, "VMware View 5 Desktop Virtualization Solutions," Packt Publishing, 2012, 288 pp. ISBN: 1849681120

Open Virtualization Format (OVF)

- Standalone software can be distributed as a virtual machine image, called, virtual appliance
- Independent of hypervisor or processor architecture
- OVF is the standard format for virtual appliances
  Standardized by DTMF (Distributed Management Task Force)
  Now ISO/IEC standard.
- OVF package consists of several files in a directory.
  An XML file with extension .ovf or a compliant format, e.g., .vmdk in the directory contains all the meta data required to run the package, e.g., hardware requirements, descriptions, security certificates, etc.
- VMware, Microsoft, Oracle, Citrix, IBM and many others support OVF

Ref: http://en.wikipedia.org/wiki/Virtual_appliance
Ref: http://en.wikipedia.org/wiki/Open_Virtualization_Format

OVF (Cont)

- OVF V1.1.0 supports single VM packages and packages containing multiple VMs constituting a multitier service
- Other popular format are Microsoft's Virtual Hard Disk (VHD), and VMware's Virtual Machine Disk (VMDK)

Ref: http://en.wikipedia.org/wiki/VHD_%28file_format%29
Ref: http://en.wikipedia.org/wiki/VMDK

PCoIP Protocol

- VMware proprietary protocol for virtual desktops
- Content-Aware: Text and graphics are treated differently
  Graphics is compressed
- Server Rendered: All pixels are rendered on the server
  No codec in the client
- Delivery optimization based on real-time network conditions
- Layered rendering: Graphics becomes clearer slowly
- Can be used with/without a hardware PCoIP card
- Alternatives:
  - Microsoft's Remote Display Protocol (RDP)
  - Virtual Network Computing (VNC)
  - X Window System (X11)

Ref: http://en.wikipedia.org/wiki/PCoIP#PCoIP_Protocol
Ref: http://en.wikipedia.org/wiki/Remote_Desktop_Protocol

Summary

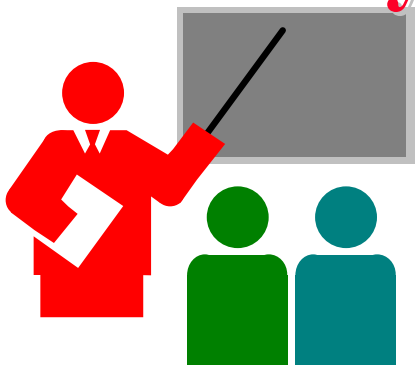

1. Virtualization allows computation to be done anywhere anytime on any infrastructure ⇒Easy and efficient resource scheduling and management
2. Servers, storage, and network all need to be virtualized
3. Hypervisors of type 1 run on bare metal. Type 2 require a host OS.
4. MS Virtual PC, Xen, VMware ESX, and Virtual Box are examples of popular virtualization products.
5. OVF is the standard format for virtual images

Acronyms

ACE	Application Control Engine (Cisco)
API	Application Programming Interface
AVM	Application Virtual Machines
CD	Compact Disk
DOM	Domain
DRS	Distributed Resource Scheduler
DTMF	Distributed Management Task Force
DVD	Digital Video Disk
ESX	VMware Product Name
FreeBSD	Free Berkeley System Distribution
GSX	VMware Product Name
IEC	International Electro technical Commission
ISO	International Standards Organization
JVM	Java Virtual Machine
KVM	Kernel Virtual Machine

Acronyms (Cont)

MS	Microsoft
OS	Operating System
OVF	Open Virtualization Format
PC	Personal Computer
SMP	Symmetric Multiprocessing
SPARC	Scalable Processor Architecture (Sun/Oracle)
svMotion	Storage virtual motion
UML	User-Mode Linux
USB	Universal Serial Bus
VHD	Virtual Hard Disk
VM	Virtual Machine
VMFS	Virtual Machine File System
vMotion	Virtual Motion
XML	eXtensible Markup Language

Reading List

- K. Hess, A. Newman, "Practical Virtualization Solutions: Virtualization from the Trenches," Prentice Hall, 2009, ISBN:0137142978 (Safari Book)
- L. C. Miller, "Server Virtualization for Dummies," Wiley, 2012, Oracle Special Edition, http://www.oracle.com/oms/hardware/extremeperformance/assets/ept-eb-dummies-server-1641465.pdf
- C. Scheffy, "Virtualization for Dummies," Wiley 2007, AMD Special Edition, http://www.amd.com/us/Documents/Virt_for_Dummies.pdf
- B. Golden, "Virtualization for Dummies," Wiley, 2011, HP special edition, https://ssl.www8.hp.com/de/de/pdf/virtuallisation_tcm_144_1147500.pdf

Wikipedia Links

- http://en.wikipedia.org/wiki/Application_virtualization
- http://en.wikipedia.org/wiki/Comparison_of_application_virtual_machines
- http://en.wikipedia.org/wiki/Comparison_of_platform_virtual_machines
- http://en.wikipedia.org/wiki/Desktop_virtualization
- http://en.wikipedia.org/wiki/ESXi
- http://en.wikipedia.org/wiki/Full_virtualization
- http://en.wikipedia.org/wiki/Hardware-assisted_virtualization
- http://en.wikipedia.org/wiki/Hardware_emulation
- http://en.wikipedia.org/wiki/Hardware_virtualization
- http://en.wikipedia.org/wiki/Hypervisor
- http://en.wikipedia.org/wiki/Open_Virtualization_Format
- http://en.wikipedia.org/wiki/Operating_system-level_virtualization

Wikipedia Links (Cont)

- http://en.wikipedia.org/wiki/Parallel_Virtual_Machine
- http://en.wikipedia.org/wiki/Paravirtualization
- http://en.wikipedia.org/wiki/Platform_virtualization
- http://en.wikipedia.org/wiki/VHD_%28file_format%29
- http://en.wikipedia.org/wiki/Virtual_appliance
- http://en.wikipedia.org/wiki/Virtual_machine
- http://en.wikipedia.org/wiki/Virtualbox
- http://en.wikipedia.org/wiki/Virtualization
- http://en.wikipedia.org/wiki/VMDK
- http://en.wikipedia.org/wiki/VMware_ESX
- http://en.wikipedia.org/wiki/VMware_ThinApp
- http://en.wikipedia.org/wiki/Windows_Virtual_PC
- http://en.wikipedia.org/wiki/Xen

Related Web Sites

- Open Virtualization Format (OVF), http://www.dmtf.org/standards/ovf (OVF standard documents)