

CNT 4603: System Administration Fall 2012

Introduction To Virtualization

Instructor : Dr. Mark Llewellyn
markl@cs.ucf.edu
HEC 236, 4078-823-2790
<http://www.cs.ucf.edu/courses/cnt4603/fall2012>

Department of Electrical Engineering and Computer Science
Computer Science Division
University of Central Florida



Introduction To Virtualization

- In a general sense, **virtualization**, is the creation of a virtual, rather than an actual, version of something.
 - For example, you can take a virtual tour of the White House by going to <http://www.visitingdc.com/white-house/virtual-tour-white-house.htm>
 - In other words, you can take a tour of the White House without actually going to the White House and taking the tour.
- From a computing perspective, you might have already done some virtualization if you've ever partitioned a hard disk drive into more than one “virtual” drive.



Introduction To Virtualization

- More generally, in a computing environment, virtualization will create a virtual version of a device or resource, such as a server, storage device, network or even an operating system where the framework divides the resource into one or more execution environments.
- Virtualization is part of an overall trend in enterprise IT that includes **autonomic computing**, which is a scenario in which the IT environment will be able to manage itself based on perceived activity, and **utility computing**, in which the processor's power is seen as a utility that clients can pay for only as needed.



Introduction To Virtualization

- Despite all the recent buzz about virtualization being a new technology...it isn't. Mainframe computers have offered the ability to host multiple operating systems for more than 40 years.
- However, four trends have converged in recent years that have allowed virtualization to migrate from the exclusive domain of mainframe computers to a front-and-center position in today's computing environment.
- These four trends are:
 - Hardware is under utilized
 - Data centers are running out of space
 - Energy costs are skyrocketing
 - System administration cost are increasing



Introduction To Virtualization

Hardware is under utilized

- Moore's Law demonstrates increasing returns – the amount of improvement in processor speed grows over time because there is an exponential increase in capacity for every generation of processor improvement.
- Many data centers today have machines running at only 10-15% of the total processing capacity. In other words, 85-90% of the machine's power is unused.
- In a way, Moore's Law is no longer relevant to most companies because they aren't able to take advantage of the increased power available to them.
- By applying virtualization, organizations can improve their hardware utilization from 10-15% to 70-80 %.



Introduction To Virtualization

Data Centers Are Running Out Of Space

- 20 years ago the vast majority of business processes were paper based. Computerized systems were confined to so-called backroom automation: payroll, accounting, and the like.
- Again, partly due to Moore's Law, this has all changed with business process after business process being captured in software and automated. The rise of the Internet has exponentially increased this transformation as well.
 - As an example, Boeing's latest airplane, the 787 Dreamliner, is being designed and built using CAD software to design all aspects of the plane. All communication about the project uses the CAD designs, no physical mock-ups are being generated. The data warehouse that houses these CAD designs for this project alone requires 5 petabytes of data!



Introduction To Virtualization

Energy Costs Are Skyrocketing

- In most companies' strategic thinking, budgeting power costs used to rank somewhere below deciding what brand of soda to keep in the vending machines. Electrical power was cheap and endlessly available.
- Several events have changed this line of thinking:
 - Increasing computerization means that companies are using more power than ever.
 - The California power scares and more frequent brown-outs have companies looking for ways to be less power dependent.
 - Escalating energy costs now make electricity one of the top five costs in many company's operation budgets.



Introduction To Virtualization

System Administration Costs Are Increasing

- Computer systems do not operate on their own. Every server requires proper care and feeding by system administrators who, as part of the operations group, ensure that the server runs properly.
- System administration is labor intensive and system administrators don't come cheap.
- Unlike system programmers who may be offsite, system administrators are almost always onsite due to their need to access the physical hardware.
- The steady increase in the number of servers world-wide has meant a very good job market for system administrators.



Introduction To Virtualization

- The typical goal of virtualization is to centralize administrative tasks while improving scalability and work loads across a system or network.
- Virtualization in a computing environment can be present in many different forms, some of which are:
 - Hardware virtualization
 - Storage and data virtualization
 - Software virtualization
 - Network virtualization



Hardware Virtualization

- **Hardware virtualization** is the execution of software in an environment separated from the underlying hardware resources. There are four primary categories of hardware virtualization:
 - **Hardware-assisted** – where some support for virtualization is built into the underlying hardware. For example, Intel VT-x or AMD-V processors.
 - **Full virtualization** – provides a sufficiently complete simulation of the underlying hardware to allow software, typically a guest OS, to run unmodified.
 - **Partial virtualization** – some, but not all of the target environment is simulated. This may require that some guests be modified in order to run in this environment.
 - **Paravirtualization** – a virtualization environment in which the guest OS requires modification to run in exchange for higher efficiency.



Storage And Data Virtualization

- The amount of data that organizations are creating and storing is exploding.
 - Many applications generate more data than can be stored on a single server.
 - Many applications have multiple machines that need to access the same data. Having all the data sitting on one physical machine creates a bottleneck.
 - Creating backups can be a huge task when hundreds of machines are involved.
- These trends have led to data virtualization. Moving data from many physical servers to a central location enables backups to be performed more efficiently. Furthermore, the central data repository can be configured with multiple physical storage devices to ensure that no hardware crash will ever make the organization's data unavailable.



Storage And Data Virtualization

- The three basic approaches to data storage are:

Direct-Attached Storage (DAS)

- DAS is the traditional mode of data storage: hard drives attached to whatever physical server is running the application.

Network-Attached Storage (NAS)

- NAS is a machine that sits on your network and offers storage to other machines. If you've ever had a remote drive in Windows, you've experienced NAS.

Storage-Area Network (SAN)

- SAN uses highly specialized hardware and software to transform simple disk drives into a data storage solution that transfers data on its own high-speed network. SANs typically use their own network protocol (Fibre Channel) for the SAN network communications.



Software Virtualization

- **Software virtualization** is most often viewed from the perspective of OS virtualization. This allow for the hosting of multiple virtualized environments within a single OS instance.
- **Application virtualization** allows the hosting of individual applications in an environment separated from the underlying OS.
- **Virtual machine (VM)** is a software (possibly hardware assisted) implementation of a computer system that executes applications like an actual computer.



The Benefits Of Virtualization

- **Server consolidation:**
 - One of the primary benefits of virtualization software is that it allows you to increase the scale of your server infrastructure without purchasing additional pieces of hardware.
- **Energy conservation:**
 - In addition to savings in hardware costs, virtualization software may also save you money on your energy bill. According to [Energy Star](#), the energy costs for running a server for a year will soon exceed the price of acquiring it.



The Benefits Of Virtualization

- Improving ease of management:
 - Managing virtual machines is a lot easier than managing “real” machines, since hardware upgrades, for example, can be done with the click of several buttons, rather than having to power down the machine, install the hardware, verify the change, then power up again. Moreover, managing virtual machines can often be done via a console server, thereby reducing the time needed to deploy them.



The Benefits Of Virtualization

- Reducing backup and recovery time:
 - Since virtual machines are essentially files, backing up and restoring them is a lot less time-consuming. And while the files can be huge, a directory of many 2-GB files is still easier to restore than a real machine of the same specifications. Moreover, hardware failures — such as a failed hard drive — will not affect virtual machines in the same way they would a physical machine. (Of course, the real hard drive on which the virtual machine reside needs to be backed up as usual.)



The Benefits Of Virtualization

- Testing software configurations:
 - Another way you can use virtualization software is for testing software configurations before deploying them on a live system. If you needed to verify whether a program is incompatible with your existing setup, for example, you may try testing in on a virtual machine first. This can be immensely useful for organizations that have legacy systems or applications and must test out systems before deploying them. Virtual machines can also interact with one another in virtual networks, allowing you to test server-client applications virtually.



The Benefits Of Virtualization

- **Maintaining legacy applications:**
 - If you do have old applications that have compatibility issues with newer software or that must run on a certain version of an operating system, you can dedicate a virtual machine just for those tasks. That way, your IT architecture and planning won't be constrained because of a few applications.
- **Maintaining a cross-platform office:**
 - It is not uncommon for offices that run mostly Macs to need to run one or two Windows-only programs; in this case, virtual software can be an affordable, easy way to do this. Note, however, that the reverse is not applicable; many virtualization applications for PCs allow you to run Linux, but not Mac operating systems.



The Benefits Of Virtualization

- Despite its many benefits, keep in mind that virtualization software is not for everyone. There is a learning curve in both conceptualizing how virtual machines will function in your network and organization, as well as managing them reliably and cost-effectively.
- If the employees of your organization have trouble with “real” computers, you may need to consider making which machine is the guest and which is the host extremely transparent to your users, or explaining to them in simple terms how this will affect their day-to-day work, if at all.



Virtualization Technology

- In the IT world, most virtualization refers to the use of virtualization software that allows the physical hardware of a single PC (or server) to run multiple operating systems simultaneously in virtual machines (VMs).
- The virtualization software simulates enough hardware to create an environment that allows an unmodified guest operating system (the one running inside a VM) to be run in isolation on a host operating system (the one running on the physical computer system).
- We'll introduce you to virtualization using Microsoft Virtual PC 2007 in this course.

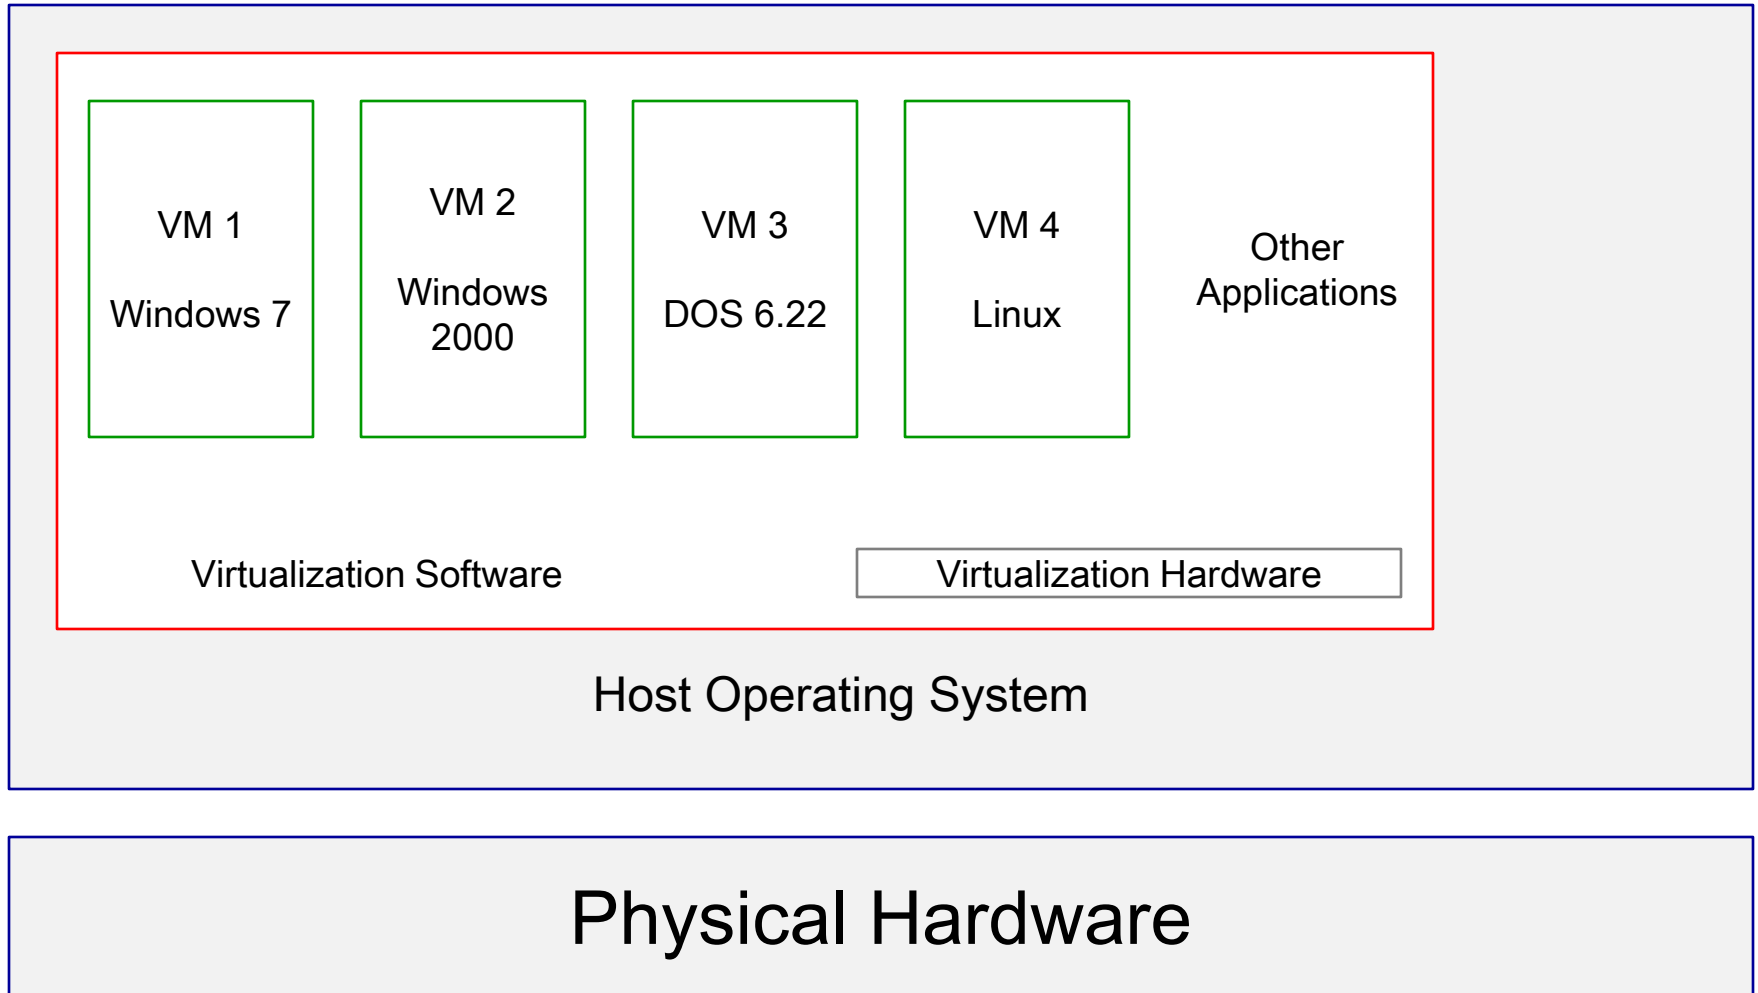


Virtualization Technology

- The diagram on the next page illustrates, at a high-level, how virtualization works:
 - At the base layer you have the hardware of the physical computer system.
 - Above the hardware layer is the host operating system. The virtualization software runs within the host OS (as an application) and provides the virtual platform for the VMs.
 - An emulated hardware system (which imitates the function of the computer system) is provided within each VM for each guest OS. The OS and software within each VM are unaware of the other VMs and have full access to the virtual platform.
- Page 23 illustrates an actual setting; the one on my office computer, in which two VMs are running alongside Vista and DOS.



Virtualization Technology



Windows Vista - real

Windows Vista - virtual

Windows Server 2008 - virtual

DOS 6.0 - real

The screenshot displays a virtual machine environment. At the top, a Windows Vista desktop is visible with a Recycle Bin icon. Below the desktop, the Server Manager console for 'WIN-4EVID7P6TAF' is open, showing a 'Server Summary' and 'Computer Information' section. The 'Computer Information' section lists details such as Full Computer Name (WIN-4EVID7P6TAF), Workgroup (WORKGROUP), Local Area Connection (IPv4 address assigned by DHCP, IPv6 enabled), Remote Desktop (Disabled), and Product ID (92573-082-2500 115-76594). In the foreground, an 'Administrator: Command Prompt' window is open, displaying a directory listing of files and folders on the C: drive. The listing includes dates, times, file names, and sizes. A clock in the background of the desktop shows the time as approximately 10:10.

Date	Time	File/Dir	Size
01/20/2008	09:33 PM	<DIR>	
10/05/2009	08:24 AM	5,763 prog2dbscript.sql	
11/18/2010	03:00 PM	<DIR>	
01/14/2011	02:55 PM	<DIR>	
01/14/2011	02:37 PM	<DIR>	
01/13/2011	10:17 AM	<DIR>	
05/04/2009	12:59 PM	1,754 script.txt	
10/27/2008	03:22 PM	<DIR>	
11/06/2008	10:15 AM	86 sprint kart.txt	
09/17/2009	06:19 PM	616 state script.sql	
09/17/2009	06:24 PM	149 states.txt	
09/21/2009	08:14 AM	222 states2.sql	
09/21/2009	08:21 AM	172 states2.txt	
09/21/2009	08:23 AM	215 states3.txt	
09/13/2010	02:49 PM	<DIR>	
11/03/2010	08:06 AM	5,576,450 UCFUGRDCatalog1011.pdf	
10/22/2008	11:17 AM	<DIR>	
11/12/2009	08:48 PM	<DIR>	
01/19/2011	10:40 AM	<DIR>	
04/14/2010	12:29 PM	<DIR>	
01/11/2011	12:15 PM	<DIR>	
		14 File(s)	5,940,454 bytes
		19 Dir(s)	120,029,212,672 bytes free



Drawbacks Of Virtualization

- Unfortunately, there are also some disadvantages to virtualization. These include a drain on performance, limited scalability, and common hardware.
- In most cases the benefits of virtualization outweigh the drawbacks, however, you need to be aware of the disadvantages when considering virtualization so that you can make an informed decision as to the use of virtualization.



Drawbacks Of Virtualization

Performance Issues:

- Virtualization software adds a certain amount of overhead to the computer system on which it is running.
- Determining the overall impact that virtualization software imposes is a difficult task, but generally speaking, virtualization software has a minimal impact on processor and memory performance. A more significant impact is imposed on the disk and networking performance of the system.

Scalability:

- Most virtualization software is scalable to multiple VMs in the same physical system, however, almost all are uni-processor systems and limited to the amount of physical memory of the host system and can only run 32-bit OS.



Drawbacks Of Virtualization

Common Hardware:

- Since the only physical hardware on which the VMs run is the same for all VMs, this represents both benefits and drawbacks when it comes to virtualization. Rapid provisioning and system flexibility are benefits of a common hardware arrangement. However, if the OS and application that you want to virtualize are not capable of running on the system hardware or require other hardware, you're out of luck with virtualization.



Common Uses Of Virtualization

- Virtualization does not lend itself to all situations. As an IT professional you will need to know how to recognize situations in which virtualization can be beneficial to an organization.
- In addition to the areas we've already mentioned, the most common areas in which virtualization can be readily applied are: application development and testing, training, demonstrations, and help desk environments.



Common Uses Of Virtualization

Application Development and Testing:

- Virtualization allows software development to run multiple VMs rather than multiple physical machines.
- When developing and testing an application, the developer might need a network of machines: a database server, a transaction server, and a client. Virtualization could be used to run this network on one physical machine.
- By using rapid provisioning, the developer can quickly and economically create the necessary network infrastructure.
- The developer can repeatedly test the software inside of VMs with the ability to quickly return to a clean, known state after each test.



Common Uses Of Virtualization

Training:

- Virtual machines are great tools for implementing training and learning. Training employees about new OSs and computer applications involves extensive amounts of time and hardware.
 - For example, if you were training students to be network administrators, you would need to provide students with access to demonstration networks. You would need to acquire enough hardware to supply isolated computer systems and the network infrastructure to provide an authentic training environment for each student. Either you or your students would need to install the OSs and configure the computer systems appropriately.



Common Uses Of Virtualization

Training (continued):

- In many training environments, the trainees will be changing the system configuration. What happens if they do something accidentally wrong? You don't want this happening on the physical system.
- Virtualization can reduce the amount of hardware needed for the necessary training environments.
- Allows for almost instant restoration of the training environment by using undo disks and differencing disks to discard changes made by a trainee.
- Multiple VMs allow for different learning environments to be running simultaneously on the same hardware.



Common Uses Of Virtualization

Help Desks:

- Help desks must support several varieties of hardware and software configurations. The combinations of multiple Windows versions and various applications becomes unwieldy.
- Help desk personnel must maintain multiple computers or restart their computers to support people who use various configurations. They might not even have access to a configuration that a customer is using, and so they must guess or fly blind with their assistance. Customer satisfaction can suffer under these scenarios.
- Using virtualization, help desks are better able to duplicate users' environments, including the OS and applications. For example, help desk personnel using Vista can support previous Windows OS by simply starting a prebuilt VM from a saved state, which takes only seconds.

