

Virtualization

Originally prepared by Greg Bosch; last modified November 2014 by B. Davison

- I. Introduction to Virtualization
 - II. Virtual Appliances
 - III. Benefits to Virtualization
 - IV. Example Virtualization Products
-

An Introduction to Virtualization

What is 'Virtualization'?

Technique for hiding the physical characterizes of computing resources from the way other systems, applications or end users interact with them

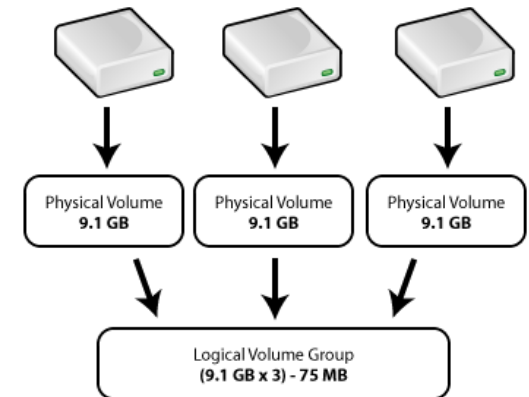
Two common functions:

- Making multiple physical resources appear to function as a single logical resource

We've see this before...

- Making a single physical resource appear to function as multiple logical resources

We'll talk about this today...



An Introduction to Virtualization

What is a 'Virtual Machine (VM)'?

Implementation of a machine that executes programs as if it were a real machine

Separated into two categories:

- **Process Virtual Machine**

Runs as a normal application inside an operating system to abstract away the details of the underlying hardware

- **System Virtual Machine**

Allows multiplexing (time sharing) of the underlying hardware between different operating systems

Virtual Machines

Process Virtual Machines

Designed to provide a platform-independent environment to a single process (i.e., program)

The environment is created when its associated process is started and destroyed when that process exits

Allows program to execute in the same way regardless of the physical platform it is running on

Implemented using an interpreter

The programmer's code is NOT compiled, but the interpreter requires compilation before providing the processing environment



Virtual Machines

System Virtual Machines

Designed to provide a complete platform which can support the execution of multiple, and different, operating systems

Allows for time-sharing of underlying hardware between virtual machines

Think of a scheduler that works on operating systems rather than processes

Operating Systems remain isolated from one another

The Instruction Set Architecture (ISA) provided by the virtual machine can be different from that of the real machine



Virtual Machines

System Virtual Machines

Implemented through the use of a **Virtual Machine Monitor (VMM)** also-known-as a Hypervisor

Two classifications of Hypervisors:

- **Native** (Hardware-Level): software runs directly on top of a given hardware platform as a control program for operating systems
- **Hosted** (OS-Level): software runs within an operating system environment as a control program for other operating systems



Virtual Machines

Native System Virtual Machines

Is where virtualization began... In the time of the mainframe

IBM developed the first Native Hypervisor in the 1960s although 'hypervisor' hadn't entered the lexicon yet...


They called it CP/CMS and it consisted of two main components

- **Control Program (CP)**

Which served to create the virtual machine environment for instances of

- **Cambridge Monitor System (CMS)**


A lightweight single-user OS



Virtual Machines

Hosted System Virtual Machines

- Virtual Machine Monitor layer is moved one level higher as compared to Native VMs
- Runs within a Host operating system environment
- An operating system is installed first; as usual, on top of Hardware
- A Virtual Machine Monitor is then installed within the Host OS
- Guest operating systems can be installed on top of the VMM layer
- Host OS sees the VMM as a process
- VMM controls the allocation of time between Guest OSes
- Guest is segregated from the rest of the environment



Virtual Machines

Emulation or Simulation in Hosted System Virtual Machines

Virtual Machine provides a “guest” OS the (simulated) hardware environment it expects

Advantage: Guest Software need not be modified

Disadvantage: Must pay Performance Penalty

Software is unaware that it is really talking to a virtualized device

Each interaction between Guest device driver with the emulated device hardware requires transaction with VMM

The real hardware does its job as usual, but the VMM must now translate the result for the guest



Virtual Machines

Paravirtualization

Application Programming Interface (API) is provided to the Guest OS by the VMM so the guest may utilize the hardware

Advantage: Better Performance

Disadvantage: Guest OS must be modified to use API

Guest interacts with VMM at a higher level of abstraction

Instead of supplying the specifics of how to use the hardware, software provides general requests to the VMM

Decreases the number of interactions between Guest and VMM for a specific operation



Virtual Machines

Containers (OS virtualization)

Instead of virtualizing the hardware, run multiple virtual instances of same OS on single hardware

Advantage: Best Performance / Scalability, Ease of Admin.

Disadvantage: Only virtualizes copies of same OS

Single kernel means very low overhead (1..3%) compared to standalone server.

Containers provide isolation between processes, appear as separate OS.

Examples: Solaris Containers/Zones,
FreeBSD Jails, OpenVZ,
Linux-VServer



Outline

- I. Introduction to Virtualization
- II. Virtual Appliances
- III. Benefits to Virtualization
- IV. Example Virtualization Products

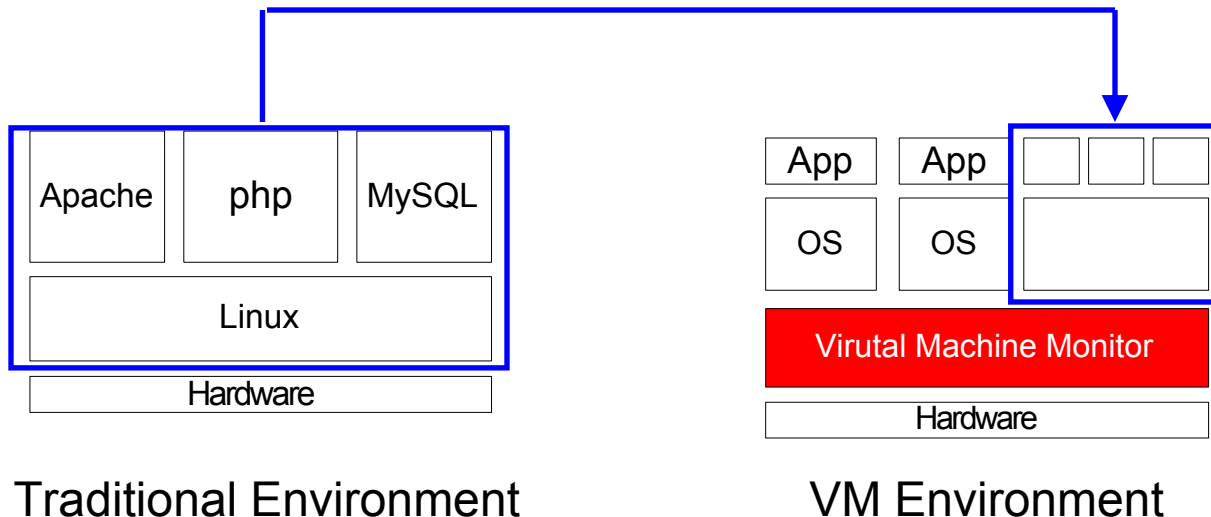
Virtual Appliance

Virtual Appliance is a virtual machine prepackaged with the necessary components to serve its intended purpose

Aimed to eliminate the installation, configuration and maintenance costs associated with complex stacks of software

For instance

LAMP Appliances (Linux + Apache + MySQL + PHP)



Virtual Appliance

Media Wiki



Software that runs wikipedia.org

Freely available to organizations

Packaged as a Virtual Appliance

Uses a minimal OS that installs within supported VMM

And includes all other necessary software packages

TRULY a Turn-Key system

Lets take a quick look if we aren't behind schedule:

<http://www.vmware.com/appliances/>

<http://bitnami.org/>

<http://www.turnkeylinux.org/all>

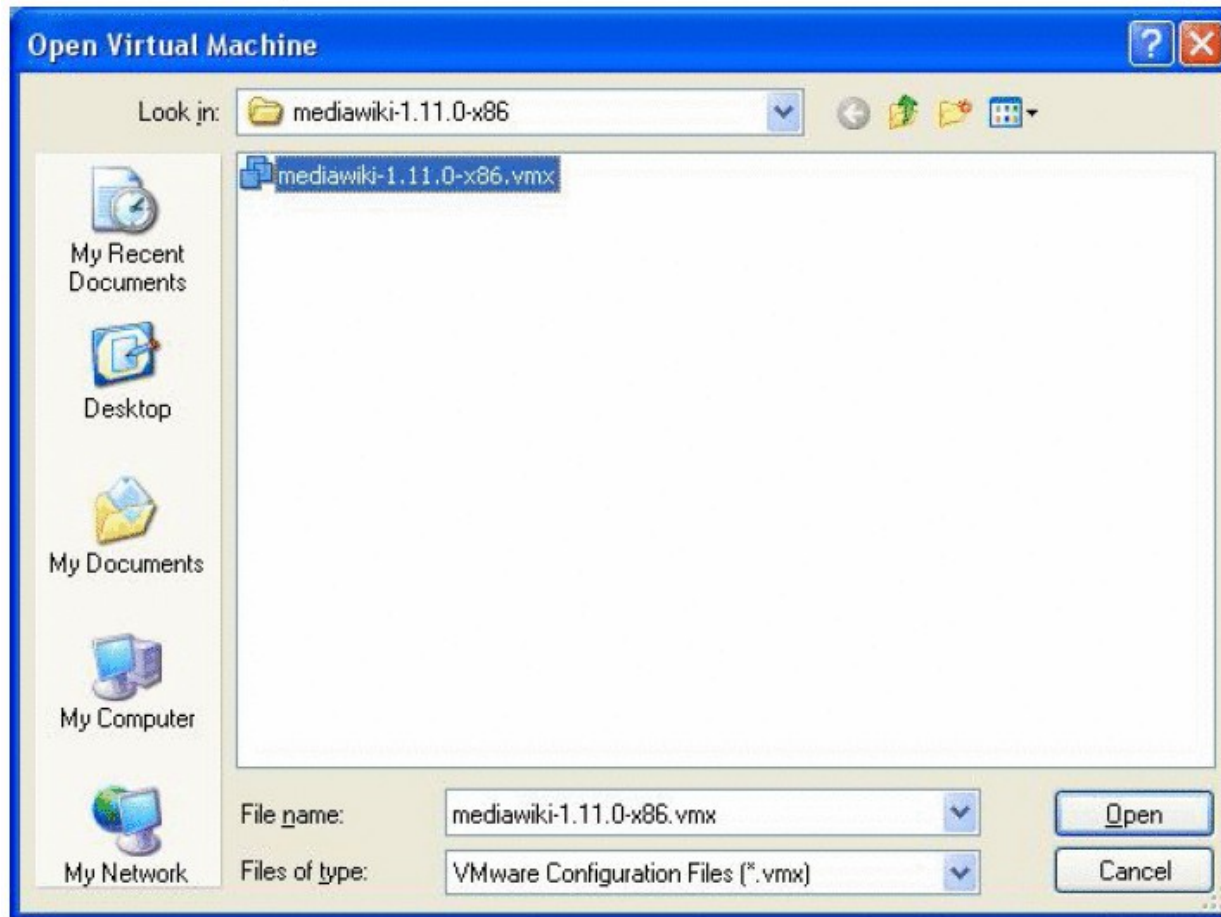
Virtual Appliance

Step 1: Install VMware's Player available from vmware.com and open the MediaWiki Virtual Machine (originally from rpath.com)



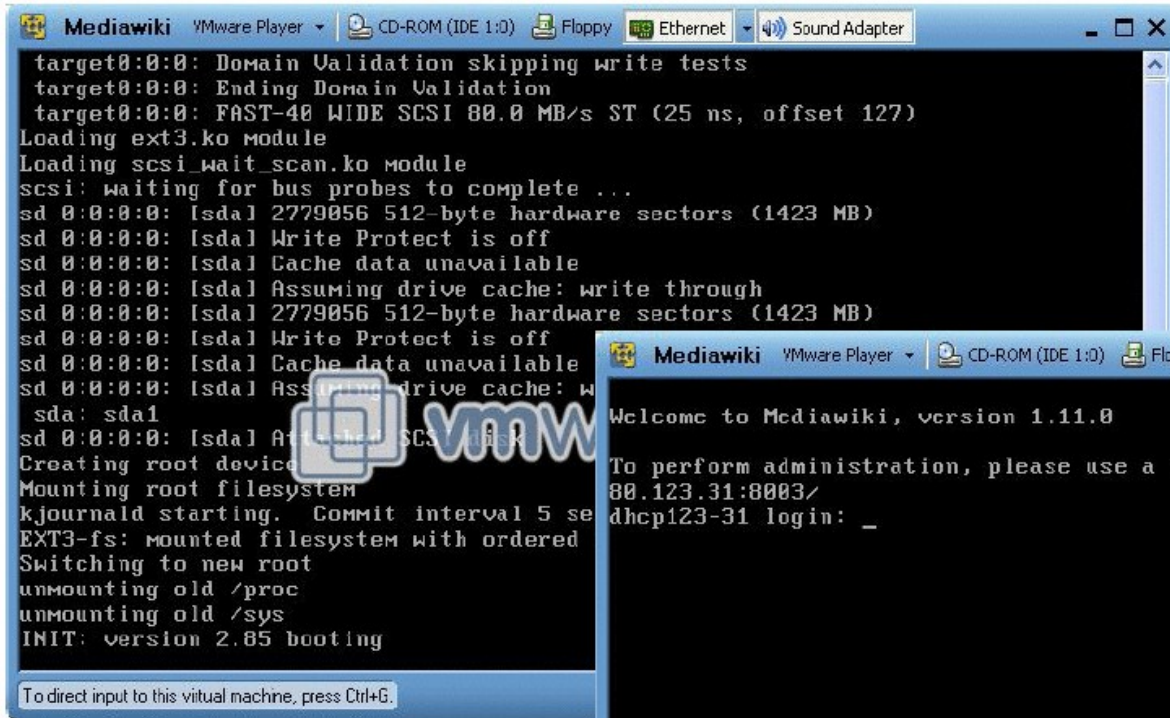
Virtual Appliance

Step 2: Opening the Virtual Appliance



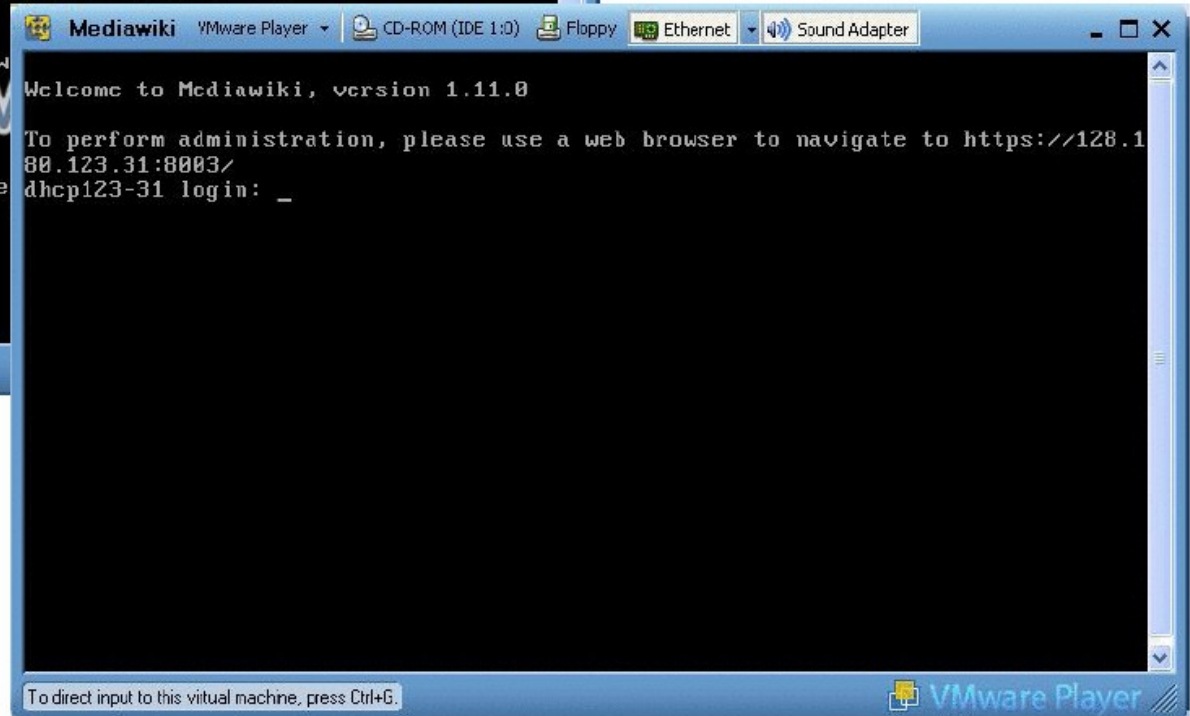
Virtual Appliance

Step 3: Just enough Operating System loading (Red Hat variant) and booting



```
target0:0:0: Domain Validation skipping write tests
target0:0:0: Ending Domain Validation
target0:0:0: FAST-40 WIDE SCSI 80.0 MB/s ST (25 ns, offset 127)
Loading ext3.ko module
Loading scsi_wait_scan.ko module
scsi: waiting for bus probes to complete ...
sd 0:0:0:0: [sd] 2779056 512-byte hardware sectors (1423 MB)
sd 0:0:0:0: [sd] Write Protect is off
sd 0:0:0:0: [sd] Cache data unavailable
sd 0:0:0:0: [sd] Assuming drive cache: write through
sd 0:0:0:0: [sd] 2779056 512-byte hardware sectors (1423 MB)
sd 0:0:0:0: [sd] Write Protect is off
sd 0:0:0:0: [sd] Cache data unavailable
sd 0:0:0:0: [sd] Assuming drive cache: write through
sda: sda1
sd 0:0:0:0: [sd] Attached SCSI disk
Creating root device
Mounting root filesystem
kjournald starting. Commit interval 5 seconds
EXT3-fs: mounted filesystem with ordered data mode
Switching to new root
unmounting old /proc
unmounting old /sys
INIT: version 2.85 booting
```

To direct input to this virtual machine, press Ctrl+G.



```
Welcome to Mediawiki, version 1.11.0
To perform administration, please use a web browser to navigate to https://128.1
80.123.31:8003/
dhcp123-31 login: _
```

To direct input to this virtual machine, press Ctrl+G.

VMware Player

Virtual Appliance

Step 4: Configuration of MediaWiki through web interface

The image displays two overlapping screenshots of the MediaWiki 1.11.0 installation web interface in Internet Explorer.

The left window shows the "MediaWiki 1.11.0 Installation" page. It features a sunflower logo and a "Checking environment..." section. Below this, there is a list of system requirements and a "Site config" section. The "Site config" section includes fields for "Wiki name", "Contact e-mail", and "Language".

The right window shows the "MediaWiki 1.11.0 Installation" page with the "Site config" section expanded. It includes the following fields and options:

- Wiki name:** myWiki (Must not be blank or "MediaWiki")
- Contact e-mail:** root@localhost
- Language:** en - English
- Copyright/license:** No license metadata, GNU Free Documentation License 1.2 (Wikipedia-compatible), A Creative Commons license - choose
- Admin username:** WikiSysop
- Password:** (Cannot be blank)
- Password confirm:** (Cannot be blank)

The bottom of the right window shows the Windows taskbar with the "Done" button and the "Internet" icon.

Virtual Appliance

Step 4: Configuration of MediaWiki through web interface

The image displays two overlapping screenshots of the MediaWiki 1.11.0 installation web interface, accessed via a Windows Internet Explorer browser. The browser address bar shows the URL: `http://dhcp123-31.eecs.lehigh.edu/wik/config/index.php`.

Left Screenshot: E-mail, e-mail notification and authentication setup

E-mail features (global): Enabled
 Disabled

Use this to disable e-mail notifications if sending e-mail is not desired.

User-to-user e-mail: Enabled
 Disabled

The user-to-user e-mail exchange e-mail will be used for user-to-user e-mail.

E-mail notification about changes: Disabled
 Enabled for changes
 Enabled for changes and recommended for latest changes

For this feature to work, you must have the e-mail options in the user's preferences enabled. For testing the feature, test it yourself.

There are additional options in `LocalSettings.php` for e-mail notification.

E-mail address authentication: Disabled
 Enabled

If this option is enabled, users will be asked for their e-mail address whenever they set their password and/or change their e-mail address. This is to prevent potential abuse of the e-mail address field.

Right Screenshot: Database config

Database type: MySQL
 PostgreSQL

Database host:

If your database server isn't on your web server, enter the name or IP address here.

Database name:

DB username:

DB password:

DB password confirm:

Must not be blank

If you only have a single user account and database available, enter those here. If you have database root access (see below) you can specify new accounts/databases to be created. This account will not be created if it pre-exists. If this is the case, ensure that it has SELECT, INSERT, UPDATE, and DELETE permissions on the MediaWiki database.

Superuser account: Use superuser account

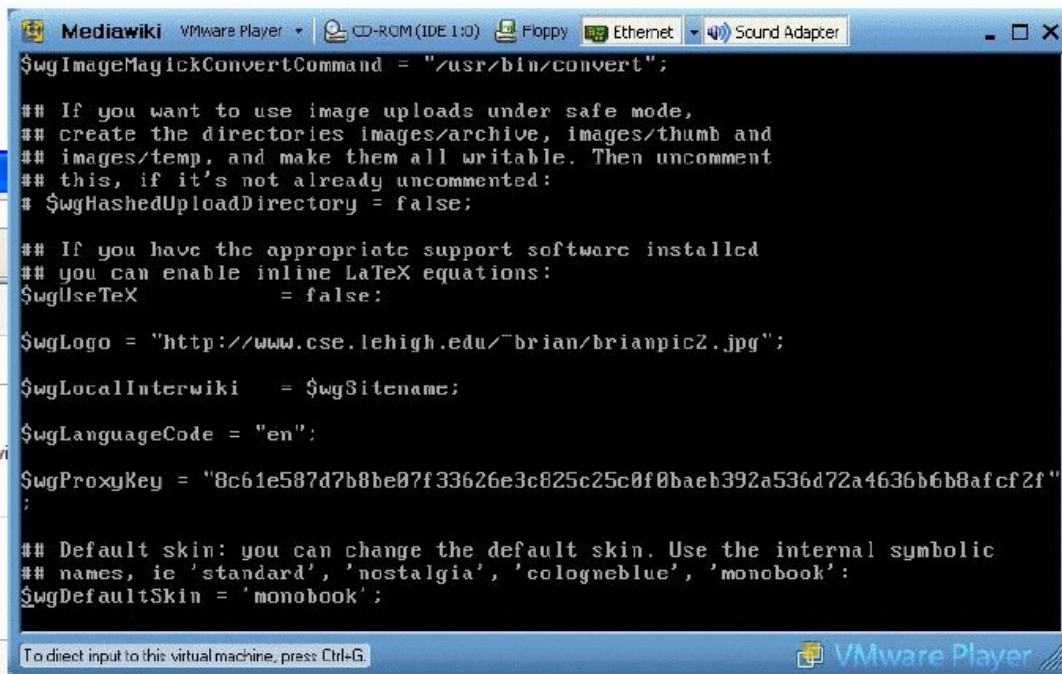
Superuser name:

Superuser password:

If the database user specified above does not exist, or does not have access to create the database (if needed) or tables within it, please check the box and provide details of a superuser account, such as root, which does.

Virtual Appliance

Having some fun



While it may seem off-topic, but it is **the point** of a virtual appliance: using your time to manage the application and not the system software

Outline

- I. Introduction to Virtualization
- II. Virtual Appliances
- III. Benefits to Virtualization**
- IV. Example Virtualization Products

Benefits to Virtualization

Cost Benefit

Traditional production servers (physical box) run a single application

Utilization of hardware by that appliance estimated at 5-15%

Reduction of physical assets reduces substantial expenses

- Cost of hardware,
- Data center footprint,
- Electricity, and
- others

Virtualization allows for consolidation of appliances at ratios between 10:1 and 20:1 (perhaps closer to 2:1 for general purpose servers)

Benefits to Virtualization

Cost Benefit continued...

Man Hour Reduction too:

Instantiation of new virtual machine requires just minutes

As compared to the “old fashioned way” involving

- Sourcing of new hardware (purchasing and installation)
 - Installation of Operating System, Patching
 - Installation of relevant Applications
 - Testing
-

Benefits to Virtualization

Operational Benefits

Automation of Installation process from OS to Applications
(i.e., Appliances)

Isolation (Sandboxing)

- Program Development and Testing
- Beneficial for running un-trusted Operating Systems or
- Un-trusted Applications

Provide legacy system support without allocating physical resources

Teaching environment for classes such as these

Check-pointing

The state of the machine can be saved, paused, restarted, even migrated to another machine

Outline

- I. Introduction to Virtualization
- II. Virtual Appliances
- III. Benefits to Virtualization
- IV. Example Virtualization Products**

Xen

Originated as a University of Cambridge Project

Project led to founding of XenSource, Inc.

Acquired by Citrix in October 2007 for \$500 Million



Products include:

- Citrix XenServer (Free)

 - A free starter package for bringing virtualization to every server

- Citrix XenDesktop

 - On-demand Windows desktop anywhere

and more...

Xenproject.org open source Xen hypervisor (behind Amazon Web Services, RackSpace Public Cloud, Verizon Cloud, etc)

Xen hypervisor

Requires modification to Kernel of pre-installed OS

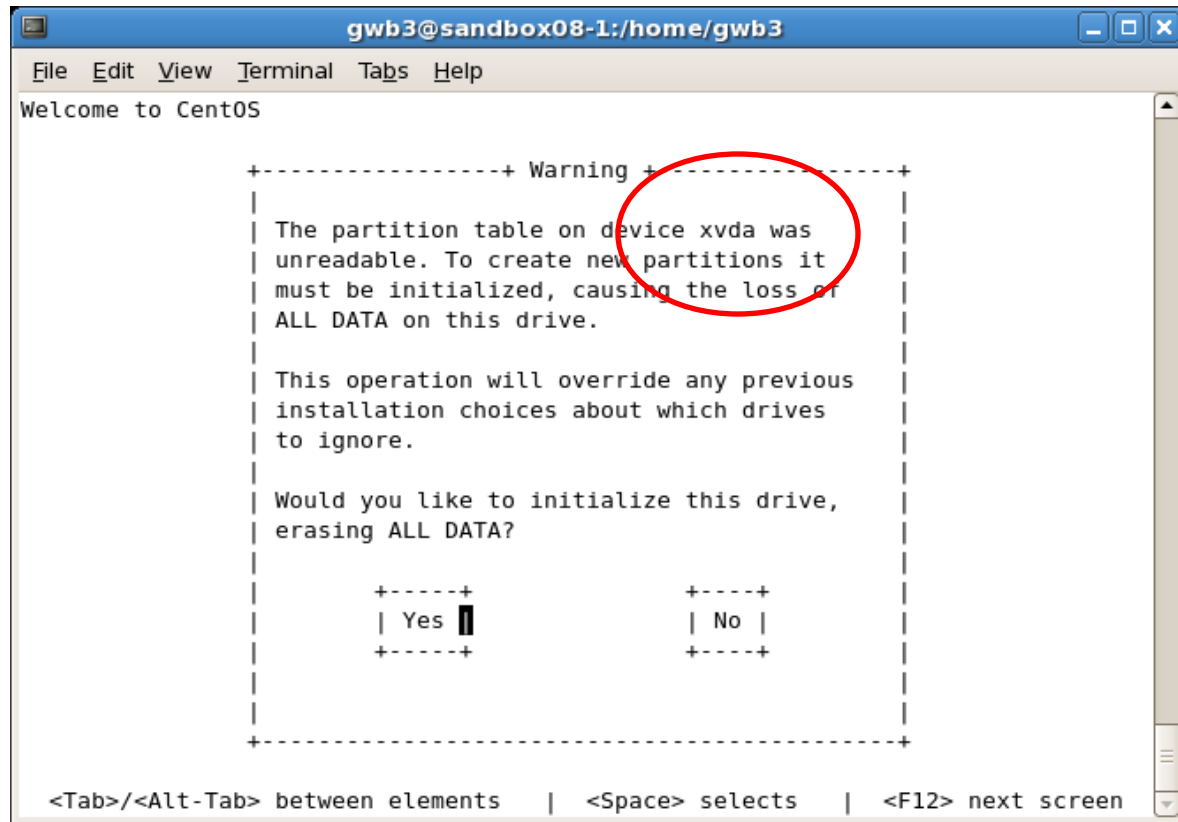
Modification installs Xen VMM just above the hardware (Native System Virtual Machine)

Xen boots from boot-loader (GRUB) and then loads the modified Host OS into the privileged domain (Dom0)

Administrator can use Host OS to install and then start guest OSs in the unprivileged domain (DomU)

Some commercial implementations (including Citrix, Oracle, etc.)

Scary moment during installation of Guest OS



```
gwb3@sandbox08-1:/home/gwb3
File Edit View Terminal Tabs Help
Welcome to CentOS

+-----+ Warning +-----+
|
| The partition table on device xvda was
| unreadable. To create new partitions it
| must be initialized, causing the loss of
| ALL DATA on this drive.
|
| This operation will override any previous
| installation choices about which drives
| to ignore.
|
| Would you like to initialize this drive,
| erasing ALL DATA?
|
| +-----+           +-----+
| | Yes |           | No |
| +-----+           +-----+
|
+-----+

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen
```

However, taking a closer look Xen is doing it's job

Xen

Sandbox Machine
running Xen Linux
Kernel

```
gwb3@sandbox08-1:/  
File Edit View Terminal Tabs Help  
[root@sandbox08-1 /]# uname -a  
Linux sandbox08-1 2.6.18-53.1.14.el5xen #1 SMP Wed Mar 5 12:39:19 EST 2008 i686  
i686 i386 GNU/Linux  
[root@sandbox08-1 /]# ls  
afs  dev      home      media  net    root    smb  tmp  vm  
bin  etc      lib       misc   opt    sbin    srv  usr  
boot gmon.out lost+found mnt    proc   selinux sys  var  
[root@sandbox08-1 /]#
```

Virtualized CentOS
environment within
original Operating
System

```
gwb3@sandbox08-1:/home/gwb3  
File Edit View Terminal Tabs Help  
[root@dhcp123-62 /]# uname -a  
Linux dhcp123-62.eecs.lehigh.edu 2.6.18-53.el5xen #1 SMP Mon Nov 12 03:26:12 EST  
2007 i686 i686 i386 GNU/Linux  
[root@dhcp123-62 /]# ls  
bin  dev  home  lost+found  misc  net  proc  sbin  srv  tmp  var  
boot  etc  lib  media      mnt  opt  root  selinux  sys  usr  
[root@dhcp123-62 /]#
```

VMware

Proprietary virtualization software developer

x86-compatible architectures only



Both Desktop and Server space software packages:

Desktop:

VMware Workstation (orig. 1999) allowed users to run multiple x86 operating systems

VMware Fusion is the Mac-Intel platform product



VMware Player allows users to run but not create VMs

Server:

VMware ESXi is an Enterprise marketed product



VMware Server is the less optimized, freeware version

VMware Player



Vmware Server



Other Technologies

