



Virtualizing Microsoft Compute Cluster Server 2003 on Scientific Linux using XEN

1. INTRODUCTION	2
2. REQUIREMENTS	2
2.1. REQUIRED HARDWARE	2
2.2. REQUIRED SOFTWARE	3
3. INSTALLATION	3
3.1. INSTALLING SCIENTIFIC LINUX 4.4	3
3.2. XEN 3.0.4	3
3.2.1. INSTALLING XEN 3.04 SOURCE	3
3.2.2. CREATING XEN KERNEL AND STARTING IT UP	4
3.3. INSTALLING WINDOWS COMPUTE CLUSTER SERVER 2003	4
4. SECURITY	6
5. CONCLUSION	7
5.1. HYPERLINKS	7
6. ATTACHMENTS	7

1. Introduction

This paper describes how to set up a virtual machine of a Windows Compute Cluster Server 2003 (Link 1) compute node on top of an installation of Scientific Linux 4.4 by using XEN.

XEN¹ is an Open Source virtualization Software. It uses a form of virtualization which is called para-virtualization (Link 2). This means, that the guest system has to be explicitly adapted to the underlying operating system. Unlike Linux Windows source code is not publicly available so modifying it is almost impossible. But with the use of processors that supports para-virtualization in this case (Virtualization Technology, VT) it is possible to run a Windows virtual machine (VM).

The physical server will be installed with an almost standard installation of Scientific Linux. After that XEN virtualization software will be installed. On top of this it is then possible to install virtual Windows 2003 Compute Cluster Servers.

By using VMs the resources can be requested when needed. Meanwhile their computing power is not needed they can be deactivated, so that other VMs on the same host can use these resources. Figure 1 gives an overview of the virtualization. The number of VMs is only limited by the resources of the host system (which is also called domain 0).

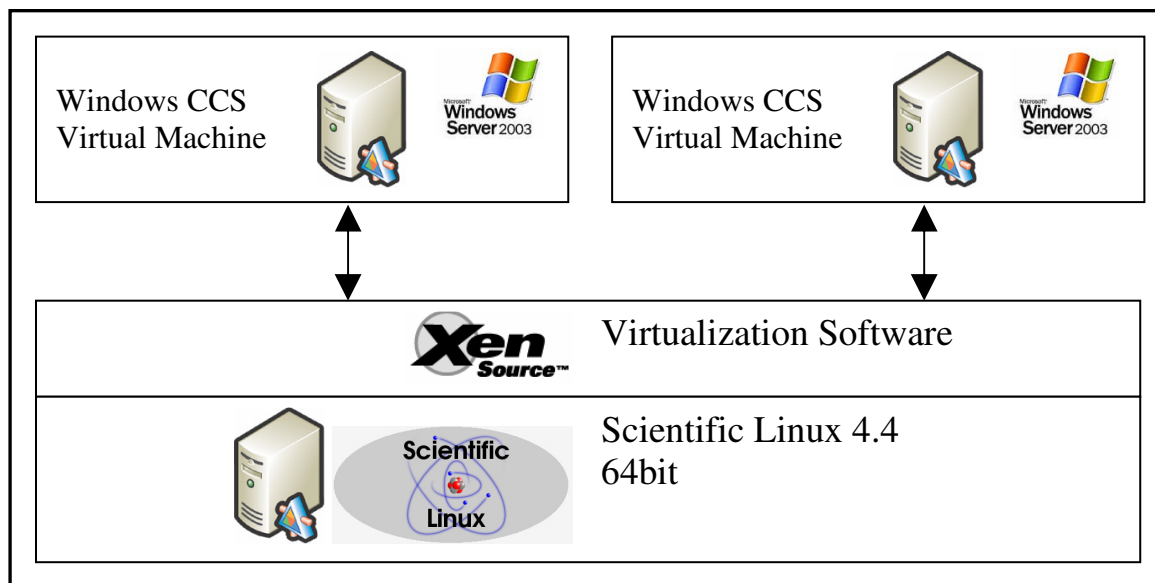


Figure 1: Overview of virtualizing Windows Compute Cluster Server 2003 on one server

2. Requirements

2.1. Required Hardware

The Platform of the host system must be 64bit architecture and in addition to that the Intel processor it must support Virtualization Technology. Alternatively it should also be possible to use AMD processors which support AMD virtualization (Codename: Pacifica).

Note that on some systems the support for VT is deactivated in the BIOS by default so it has to be activated first.

¹ <http://www.xensource.com>

Concerning RAM, at least 512 MB should be installed. 256 MB will be reserved for the Linux host operating system and at least 256 are set aside for each Windows Server operating system. However, at least 512MB of RAM for every virtual machine are recommended. Otherwise the system will be very slow.

The need for hard disk space depends mainly on the number of virtual machines. Here, 4GB were used for the host system running Scientific Linux with an adapted XEN-kernel (and a swap partition with 512MB size was created as well). Furthermore 10 GB were reserved for the first Windows 2003 CCS Cluster server and another partition of 10 GB for backup data and saved ISO images.

The hardware which was used here for building this prototype was a HP Proliant ML 110 G4 with a dual core Intel XEON 3040 processor (1.86 GHz) and 512 MB of RAM.

2.2. Required Software

Scientific Linux 4.4 64bit was chosen to be the operating system. Scientific Linux is a widely spread Linux distribution in the science sector. Previous installation trials with 32bit systems showed that the installation a virtual Windows Server 2003 CCS is not possible (Error message: *Attempting to load an x64 operating system, however this CPU is not compatible with x64 mode. Please install a 32-bit operating system. Setup can not continue. Press any key to exit.*). Because Windows Server 2003 CCS is only available as 64bit operating system the host system had to be 64bit as well.

3. Installation

3.1. Installing Scientific Linux 4.4

The installation of Scientific Linux 4.4 (Link 3) was done using standard installation type "server". Additionally, the X-Window system, the Editors and graphical internet were installed for easier configuration. Not needed packages can of course later be uninstalled. Also during the installation the option for Security-Enhanced Linux (SELinux) should be set to warning. If not, network connection will be very restricted.

With these installation settings only the first two CDs out of five of the Scientific Linux package are needed.

Using the integrated partition tool, both the root and swap partition were created. Later Logical Volume Manager (LVM) was used to create the partitions for the virtual machines. This was used instead of standard partition use because it has a better disk throughput. For further information take a look at 4).

3.2. XEN 3.0.4

3.2.1. Installing XEN 3.04 Source

In order to be able to run a virtualized Windows 2003 CCS, the old version of XEN (version 3.0.4) has to be used. The reason why the new version 3.1 cannot be used is because it doesn't (yet) allow virtual machines with 64bit operating system.

XEN 3.0.4 is Open Source and can be downloaded on the XEN Website (Link 5) but before the it can be set up, additional components have to be downloaded and installed:

- Bridge-utils
- gcc
- openssl-devel

- `apt-get install xorg-x11-devel`
- `ncurses-devel`
- `python-devel`
- optionally `SDL-devel` (see *Installing Windows Compute Cluster Server 2003*)

These packages can be installed with `apt-get install <package name>`². After that XEN is installed by executing `make world` and `make install` in the XEN folder. Please note that you have to disable iptables temporarily in order to connect to the internet. Now the system is ready for installing the first virtual machine.

More information and help about installing XEN Source can be found on the following CERN website (Link 6) and of course in the XEN 3.0 user manual (Link 7).

3.2.2. Creating XEN Kernel and starting it up

It is required to use a XEN modified XEN kernel in order to use virtualization. The necessary steps for doing this will be explained in the following paragraph:

1. Executing `/sbin/mkinitrd /boot/initrd-2.6.16.33-xen.img 2.6.16.33-xen` will create the necessary initial ram disk to boot XEN
2. Create an additional entry in `/etc/boot/grub/menu.lst`³

```
title XEN
  root(hd0,0)
  kernel /boot/xen-3.0.gz
  module /boot/vmlinuz-2.6.16.33-xen ro root=/dev/sda1 rhgb quiet
  module /boot/initrd-2.6.16.33-xen.img
```

3. The XEN demon `xend` should be set to start automatically in the runlevels 3, 4 and 5. Therefore the command `/sbin/chkconfig --level 345 xend on` needs to be run.
4. To start VMs automatically the configuration file needs to be in `/etc/xen/auto` and the service `xendomains` must be started automatically:
 - a. `ln -s <path_of_configfile> /etc/xen/auto/someName`
 - b. `/sbin/chkconfig --level 345 xendomains on`
5. After that the system should be rebooted it will start using the XEN-kernel.

3.3. Installing Windows Compute Cluster Server 2003

After that the configuration file for starting virtual Windows Server 2003 CCS can be created. Example files can be found in the folder `/etc/xen`. The most important parameters will be explained in the following paragraph and the whole configuration file which was used here will be available in the attachments.

² At first `apt-get update` to update the package information and after that for example “`apt-get install bridge-utils`” to install the bridge utilities package. Note also that the iptables service should be stopped in order to be able to download from the internet.

³ it makes sense to set XEN as default boot option (see *default* parameter)

Parameter *boot*:

At first the parameter *boot* should be set do *d* because the installation will be from a CD-Rom (ISO-image). After finishing the installation of Windows *boot* should be changed to *c*, because from then it will be able to boot from the hard disk, just like any other Windows.

Parameter *kernel*:

The parameter *kernel* points to the generated *hvmloader* file which is located in `/user/lib/xen/boot/hvmloader`. If the *hvmloader* is not present in this directory, the tools have to be installed again: Execute *make* and *make install* in the folder `.../xen/tools`.

Parameter *disk*:

Copying the installation CD on the disk will save some time because the installation will be faster. The command `dd if=/def/cdrom of=/PathOnDisc` will do this. The parameter *disk* tells the system where its root partition is located and sets the path for the cdrom iso. Also, we won't have problems accessing the CD-drive from our guest system.

```
disk = [ 'phy:/dev/vg/w2k3,ioemu:had,w', 'file: /images/winccs2003.iso,hdc:cdrom,r' ]
```

Parameter *SDL* and *VNC*:

These parameters decide how the window of the virtual machine is accessed. Either you can use *SDL* or *VNC*. This is only necessary for the initial configuration. However, once the Windows operating system is installed it is recommended to activate Remote Desktop Connection to administrate the Windows VM through RDP (Link 8).

If you use *SDL*, a window will pop up once the VM is started. But this only works if the VM is started while logged on locally. Usually it should be enough just to configure the VM this way and later connect remotely through RDP.

Alternatively you can activate *VNC* instead. With *VNC* configured you would also be able to connect from other machines to this VM. Note that when later configuring the VMs to start automatically you have to set *SDL* to 0 and activate *VNC* instead. Otherwise the VM is not able to start correctly.

The virtual machine is now ready to be started via `xm create <config_file>`. By running `xm list` all running VMs will be listed.

Once started, the operating System can be installed just like any other Windows Server (see Figure 2 Running Windows 2003 Compute Cluster Server inside Scientific Linux (using *SDL*). For integrating it into the existing cluster basically three things have to be done:

- join the existing Active Directory Domain
- configure the server as Compute Node (or new Head Node)
- if declared as Compute Node it has to be approved on the head node

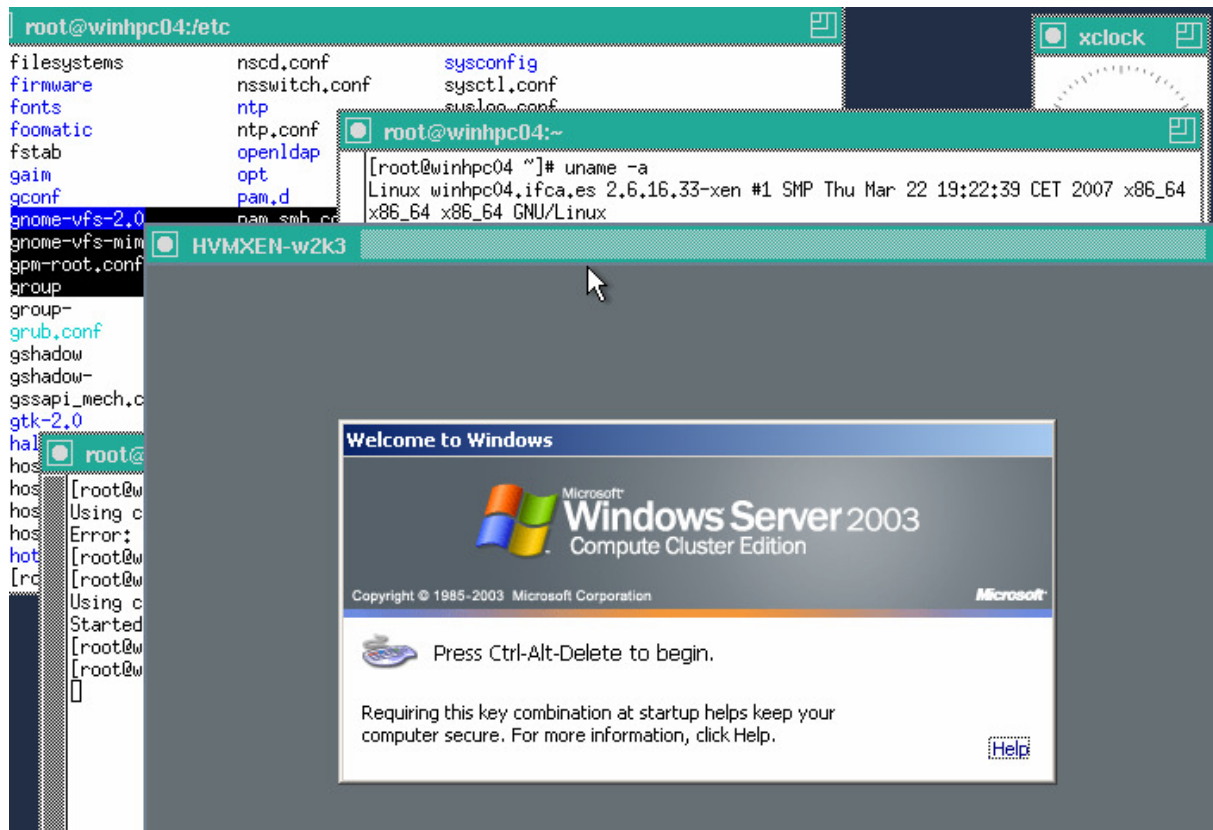


Figure 2 Running Windows 2003 Compute Cluster Server inside Scientific Linux (using SDL)

4. Security

For security reasons it is strongly recommended to use a firewall or a packet filter software. This prototype uses *iptables*. The following example shows some *iptables* settings for administrative purposes:

- Open RDP port (TCP:3389) for remote administration through RDP protocol (Remote administration has to be activated on the Windows Server as well)
- The *forward* entry in *iptables* should be set to ACCEPT because XEN creates several virtual network adapters and they need to forward the traffic.
- Open SSH ports to administrate the host system

For more information see the *iptables* config file named *firewall.sh* in the attachments⁴.

⁴ To execute: `sh firewall.sh` and after that you the settings are saved *with* `iptables-save>/etc/sysconfig/iptables`. Then *iptables* should be restarted with `service iptables restart`.

5. Conclusion

Setting up VMs of Windows Server 2003 CCS is very good way to start with Windows Computing. At the beginning, when the load is not very high, these VMs can be shutdown to give the resources to other servers. When needed, it is very easy to start the machines and use them to compute jobs. The advantage of using XEN is its ability to virtualize Windows operating systems while losing not very much of the performance.

Certainly this new virtual machine has to be tested thoroughly for its performance. However, this should be confirmed by testing all kinds of performance counters like CPU speed, disk throughput, network speed, etc. and compared to the non-virtualized system.

The virtual machine does not know that it is being virtualized so it can be used just like any other physical server. By adding some more RAM many more Virtual Machines could be set up, Windows or Linux. If there are a lot of machines to be installed you could use RIS (Remote Installation Services). It might also be possible to copy the existing installation, remove all server specific data (e.g. with sysprep).

Hopefully future versions of XenExpress (Version 3.1) will support 64bit VMs because installing this host system is much easier, faster and less error-prone. You just have to insert the CD and install the server (a lean Linux system). It is not necessary to install sources, create boot images and so on. Administration of XEN 3.1 can be done through an administration console which can even be installed on Windows systems.

5.1. Hyperlinks

1. More information about Windows Compute Cluster Server:
<http://www.microsoft.com/windowsserver2003/ccs/default.msp>
2. Para-virtualization: <http://en.wikipedia.org/wiki/virtualization>
3. Download of Scientific Linux 4.4 64bit:
ftp://ftp.scientificlinux.org/linux/scientific/44/iso/x86_64/
4. VM integration in INT.EU.GRID
<http://savannah.fzk.de/autobuild/artifacts/Xen/Xen.pdf>
5. XEN download sources: http://bits.xensource.com/oss-xen/release/3.0.4-1/src.tgz/xen-3.0.4_1-src.tgz
6. CERN Website – installing XEN Scientific Linux: http://project-xen.web.cern.ch/project-xen/xen/howto_slcXen.html
7. XEN 3.0 user manual <http://tx.downloads.xensource.com/downloads/docs/user/>
8. RDP – Remote Desktop Protocol:
http://en.wikipedia.org/wiki/Remote_Desktop_Protocol

6. Attachments

- Config file for Windows Virtual Compute Cluster Server 2003: *w2k3.hvm*
- Config file for iptables rules: *firewall.sh*