



XenServer Virtual Machine Installation Guide

Release 4.0.1

October 8, 2007

XenServer Virtual Machine Installation Guide

Copyright © 2007 XenSource, Inc.

Xen®, XenSource™, XenEnterprise™, XenServer™, XenExpress™, XenCenter™ and logos are either registered trademarks or trademarks of XenSource, Inc. in the United States and/or other countries. Other company or product names are for informational purposes only and may be trademarks of their respective owners.

This product contains an embodiment of the following patent pending intellectual property of XenSource, Inc.:

1. United States Non-Provisional Utility Patent Application Serial Number 11/487,945, filed on July 17, 2006, and entitled 'Using Writeable Page Tables for Memory Address Translation in a Hypervisor Environment'.
2. United States Non-Provisional Utility Patent Application Serial Number 11/879,338, filed on July 17, 2007, and entitled 'Tracking Current Time on Multiprocessor Hosts and Virtual Machines'.

Contents

1	About this document	5
1.1	Overview	5
1.2	How this Guide relates to other documentation	5
2	Creating VMs	7
2.1	Overview	7
2.2	Virtual memory and disk size minimums	7
2.3	XenServer product family virtual device support	8
2.4	Physical to Virtual Conversion (P2V)	8
2.4.1	General Guidelines for Virtualizing Physical Servers	9
2.5	Cloning an existing VM	10
2.6	Importing an exported VM	10
3	Installing Windows VMs	13
3.1	Making the ISO available to XenServer Hosts	13
3.1.1	Copying ISOs to local storage	14
3.2	Windows paravirtualized drivers	14
3.3	Remote Desktop	15
3.4	Preparing to clone a Windows VM	15
3.5	Release Notes	16
3.5.1	General Windows Issues	16
3.5.2	Windows 2003 Server	16
3.5.3	Windows XP SP2	16
3.5.4	Windows 2000 Server	16
4	Installing Linux VMs	17
4.1	Installation of a built-in distribution	17
4.2	Installing Linux from a network installation server to a VM	18
4.3	Physical-to-Virtual Installation of a Linux VM	18
4.3.1	Guest Installation Network	18
4.4	Installing the Linux guest agent	19
4.5	Preparing to clone a Linux VM	20
4.5.1	Machine Name	20
4.5.2	IP address	20
4.5.3	MAC address	20
4.6	Time handling in Linux VMs	20
4.7	Configuring VNC for VMs	21
4.7.1	Setting up Red Hat-based VMs for VNC	21
4.7.1.1	Determining the location of your VNC configuration file	21
4.7.1.2	Configuring GDM to use VNC	21

4.7.1.3	Firewall settings	22
4.7.1.4	VNC screen resolution	23
4.7.2	Setting up SLES-based VMs for VNC	23
4.7.2.1	Checking for a VNCserver	23
4.7.2.2	Enabling Remote Administration	23
4.7.2.3	Modifying the xinetd configuration	24
4.7.2.4	Firewall settings	25
4.7.2.5	VNC screen resolution	25
4.7.3	Setting up Debian-based VMs for VNC	26
4.7.4	Checking runlevels	26
4.8	Release Notes	26
4.8.1	Debian Sarge 3.1 and Etch 4.0	27
4.8.2	Red Hat Enterprise Linux 3	27
4.8.3	Red Hat Enterprise Linux 4	27
4.8.4	Red Hat Enterprise Linux 5	27
4.8.5	CentOS 4	28
4.8.6	CentOS 5	28
4.8.7	SUSE Enterprise Linux 9	28
4.8.8	SUSE Enterprise Linux 10 SP1	28
5	Updating VMs	29
5.1	Updating paravirtualized drivers for Windows VMs	29
5.2	Updating Linux kernels and guest utilities	29
A	Creating ISO images	31
B	Setting Up a Red Hat Installation Server	33
B.1	Copy installation media	33
B.2	Enable remote access	33
B.2.1	NFS	33
B.2.2	FTP	34
B.2.3	HTTP	34
C	Troubleshooting VM problems	35
C.1	VM crashes	35
C.1.1	Linux VMs	35
C.1.2	Windows VMs	36
	Index	37

Chapter 1

About this document

1.1 Overview

This document is a guide to creating Virtual Machines with XenServer™, the platform virtualization solution from XenSource™. It describes the various methods of getting VMs up and running on XenServer Hosts for each of the supported operating systems.

This section summarizes the rest of the guide so that you can find the information you need. The following topics are covered:

- general information about creating VMs
- creating Windows VMs
- creating Linux VMs
- updating VMs
- Creating and using ISO images of vendor media for installing VMs
- Setting up a network repository of vendor media for installing VMs
- Troubleshooting problems with VMs

1.2 How this Guide relates to other documentation

This document is primarily aimed at system administrators who need to set up deployments of XenServer VMs. Other documentation shipped with this release includes:

- *XenServer Installation Guide* provides a high level overview of XenServer, along with step-by-step instructions on installing XenServer Hosts and the XenCenter management console;
- *XenServer Administrator's Guide* describes the tasks involved in configuring a XenServer deployment -- how to set up storage, networking and resource pools, and how to administer XenServer Hosts using the xe command line interface (CLI).
- *Programmer's Guide* presents an overview of the XenServer SDK -- a selection of code samples that demonstrate how to write applications that interface with XenServer Hosts.
- *API Reference* A programmer's reference guide to the XenServer API.
- *Release notes* provide a list of known issues that affect this release.

Chapter 2

Creating VMs

This chapter provides an overview of how VMs are created and lists virtual memory and virtual disk size minimums, describes the differences in virtual device support for the members of the XenServer product family. This chapter also discusses physical to virtual conversion (P2V), cloning Templates, and importing previously-exported VMs.

2.1 Overview

VMs are created from *Templates*. A Template is a "gold image" that contains all the various configuration settings to instantiate a specific VM. XenServer ships with a base set of Templates, which range from generic "raw" VMs that can boot an OS vendor installation CD (Windows) or run an installation from a network repository (Red Hat Enterprise Linux, Suse Linux Enterprise 10) to complete pre-configured OS instances (Debian Etch and Sarge).

There are three basic methods by which VMs are created using Templates:

- using a complete pre-configured Template (Debian Sarge and Etch Linux)
- Installing from a CD or an *ISO image* onto the appropriate Template (Windows 2000 SP4/Windows 2003 Server/Windows XP SP2)
- Installing from vendor media on a network installation server directly onto a Template (Red Hat Enterprise Linux 4.x and 5.0, and SUSE Linux Enterprise Server 10 SP1)

Creating VMs by installing Windows operating systems onto the appropriate Templates is described in [Chapter 3](#).

Creating VMs by installing Linux operating systems onto the appropriate Templates is described in [Chapter 4](#).

Additionally, VMs can be created by

- performing a physical to virtual (P2V) conversion on an existing physical server (Red Hat Enterprise Linux 3.6, 3.8, 4.1-4.4, and SUSE Linux Enterprise Server 9 SP2 and SP3)
- Importing an existing, exported VM
- Converting an existing VM to a Template

These methods are describe in this chapter.

2.2 Virtual memory and disk size minimums

In general, when installing VMs, be sure to follow the memory and disk space guidelines of the operating system and any relevant applications that you want to run when allocating resources such as memory and disk space.

Operating System	RAM	Disk space
Windows 2003	128MB minimum supported; 256MB or more recommended	2GB
Windows XP SP2	128MB minimum supported; 256MB or more recommended	1.5GB
Windows 2000 SP4	128MB minimum supported; 256MB or more recommended	2GB
Red Hat Enterprise Linux 3.6	64MB	1.5GB
Red Hat Enterprise Linux 4.1, 4.4	256M	800MB
Red Hat Enterprise Linux 5	256M	800MB
SUSE Linux Enterprise Server 9 SP2	256MB	1GB
SUSE Linux Enterprise Server 10 SP1	512MB	1.5GB
Debian Sarge, Etch	128MB	4GB

2.3 XenServer product family virtual device support

The current version of the XenServer product family has the following general limitations on virtual devices for VMs. Note that specific guest operating systems may have lower limits for certain features. These limitations are noted in the individual guest installation section.

Virtual device	Linux VMs	Windows VMs
Number of virtual CPUs	32 ¹	8
Number of virtual disks	8 (including virtual CD-ROM)	8 (including virtual CD-ROM)
Number of virtual CD-ROM drives	1	1
Number of virtual NICs	7	7
Hotplugging virtual disks	add/remove	add only
Hotplugging virtual NICs	add/remove	add/remove

XenExpress, XenServer, and XenEnterprise also differ in the following ways that are relevant for creating VMs:

	XenEnterprise	XenServer	XenExpress
Amount of physical RAM on XenServer Host	up to 128GB	up to 128GB	up to 4GB
Number of concurrent VMs	50	50	4
Support for VLANs	yes	no	no
Support for shared storage	yes	no	no
Support for server pools	yes	no	no
Support for additional QoS control	yes	no	no

If you attempt to create a fifth VM with a XenExpress license, for example, an error message will appear, which suggests that you can upgrade your license and gives a web address for the XenSource website.

2.4 Physical to Virtual Conversion (P2V)

Physical to Virtual Conversion (P2V) is the process by which an existing operating system on a physical server — its filesystem, configuration, etc. — is cast into a virtualized instance of the same operating system and filesystem, transferred, instantiated, and started as a VM on the XenServer Host. For existing physical instances of Windows servers, a third-party tool is required. Windows P2V software and documentation is available for download from the XenSource website at <http://www.xensource.com/partners/offers>. For existing physical instances of Linux servers, this is accomplished by booting from the XenServer installation CD and choosing the P2V option. The

¹ A maximum of 8 vCPUs are supported via XenCenter.

filesystem is copied across the network onto a XenServer Host, where it appears as a normal VM. We recommend that you perform P2V operations during off-peak hours since the process involves transferring a large amount of data, which could impact other Virtual Machines running on the XenServer Host.

The P2V tool requires a 64-bit capable CPU by default. If you have an existing Linux instance on an older machine that you want to transfer via P2V, you can boot the CD via the *p2v-legacy* option at the initial prompt. This does require at least a PAE-enabled machine, so for very old machines you can physically move the hard drive to a PAE-enabled machine and perform the operation from there.

To P2V an existing Linux server directly to a XenServer Host:

1. Reboot the physical server that you want to convert and boot from the XenServer installation CD. If the boot fails, start again and use the *p2v-legacy* option.
2. After the initial boot messages, you will see the "Welcome to XenServer" screen. (In this and the screens that follow, use **Tab** or **Alt+Tab** to move between elements, **Space** to select, and **F12** to move to the next screen.)
Select OK to proceed.
3. The installer does some hardware detection and initialization, then presents a screen with four choices.
Select Convert an existing OS on this machine to a VM (P2V) and choose OK to proceed.
4. A "Welcome to XenServer P2V" screen with a descriptive message is displayed next. Click OK to proceed and follow the on-screen prompts.

When the P2V process is complete and the new VM is created, you will need to create and attach a VIF for it to have external network connectivity. Similarly, extra disks may also be added to take advantage of additional storage capacity available to the XenServer host.

Since the VM has new virtual network hardware, the MAC addresses it sees will also be different. Follow the Linux cloning guidelines (see Section 4.5) for customizing the configuration files to make the VM re-run any hardware detection scripts at startup.

2.4.1 General Guidelines for Virtualizing Physical Servers

When considering how to best begin virtualizing a collection of physical servers, it is best to gain some comfort level and experience with virtualizing servers that are more simply configured, moving later to servers with more complex configurations.

Good candidates typically include servers that are used for test and development environments, and servers used for in-house IT infrastructure (intranet web servers, DNS, NIS, and other network services, etc.). Typically servers that are doing heavily CPU-intensive tasks (sophisticated mathematical modeling, video rendering) or are I/O-intensive (high-traffic commercial web sites, highly-used database servers, streaming audio/video servers) are not the best candidates for virtualization at the start.

Once you've identified some physical servers that seem reasonable to work on first, you should take a close look at how you are currently using them. What applications are they hosting? How I/O intensive are they? How CPU-intensive are they?

To make a reasonable assessment, you should gather a reasonable amount of data on the current physical servers that you are thinking about virtualizing. Look at system monitoring data for disk usage, CPU usage, memory usage, and network traffic, and consider both peak and average values.

Good candidates for virtualization are:

- servers whose CPU and memory usage and NIC and disk throughput are low will be more likely to coexist as a VM on a XenServer Host with a few other VMs without unduly constraining its performance.
- servers that are a few years old - so their performance as VMs hosted on a newer server would be comparable to their existing state.

- servers that don't use any incompatible hardware which cannot be virtualized, such as dongles, serial or parallel ports, or other unsupported PCI cards (serial cards, crypto accelerators, etc.).

Once you have identified a set of machines that you want to virtualize, you should plan the process to accomplish the task. First, provision the physical servers that will serve as your XenServer Hosts. The chief constraint on the number of VMs you can run per XenServer Host is system memory.

Next, plan how you will create the VMs. Your choices are to P2V an existing server, install a fresh server from network-mounted vendor media, or install a base operating system using a pre-existing template.

If you P2V an existing server, it's best to P2V a test instance of the server, and run it in parallel with the existing physical server until you are satisfied that everything works properly in the virtual environment before re-purposing the existing physical machine.

Next, plan how to arrange the desired VMs on the XenServer Hosts. Don't "mix up" servers - assign VMs to specific XenServer Hosts, giving consideration to complementary resource consumption (mixing CPU-intensive and I/O-intensive workloads) and complementary peak usage patterns (for instance, assigning overnight batch processing and daytime interactive workloads to the same XenServer Host).

For configuring individual VMs themselves, keep these guidelines in mind:

- create single-processor VMs unless you are serving a multi-threaded application that will perform demonstrably better with a second virtual CPU.
- when you configure the memory settings for a VM, consult the documentation for the guest operating system you plan to run in that VM and for the applications you plan to run on them.

2.5 Cloning an existing VM

You can make a copy of an existing VM by *cloning* from a template. Templates are just ordinary VMs which are intended to be used as master copies to instantiate copies from. A VM can be customized and converted into a template, but be sure to follow the appropriate preparation procedure for the VM (see Section 3.4 for Windows and Section 4.5 for Linux). Templates cannot be used as normal VMs without first cloning them.

XenServer has two mechanisms for cloning VMs: a full copy, or a faster "Copy-on-Write" (CoW) mode which only writes modified blocks to disk. The CoW mode is only supported for file-backed VMs. CoW is designed to save disk space and allow fast clones, but will slightly slow down normal disk performance. A template can be fast-cloned multiple times without slowdown, but if a template is cloned into a VM and the clone converted back into a template, disk performance can linearly decrease depending on the number of times this has happened. In this event, the **vm-copy** CLI command can be used to perform a full copy of the disks and restore expected levels of disk performance.

2.6 Importing an exported VM

You can make a VM by *importing* an existing exported VM. Like cloning, exporting and importing a VM is a means for creating additional VMs of a certain configuration. You might, for example, have a special-purpose server configuration that you use many times. Once you have set up a VM the way you want it, you can export it, and later import it, creating another copy of your specially-configured VM. This also provides a simple way to move a VM to another XenServer Host.

When importing a VM, you can choose to preserve the MAC address on any virtual network interfaces associated with it. If you do choose to generate a new MAC address, be sure to follow the appropriate preparation procedure for the imported VM (see Section 3.4 for Windows and Section 4.5 for Linux).

Importing an exported VM will take some time, and depends on the size of the VM and the speed and bandwidth of the network connection between the XenServer Host and XenCenter.



Note that an exported VM that originated on one XenServer Host might or might not be able to be resumed on a different XenServer Host. For example, if a Windows VM created on a XenServer Host with an Intel VT-enabled CPU is exported, then imported to a XenServer Host with an AMD-V CPU, it will not start.

Chapter 3

Installing Windows VMs

XenServer allows you to install Windows 2000 SP4, Windows Server 2003 (32-/64- bit), or Windows XP SP2 into a VM. Installing Windows VMs on XenServer Host requires hardware virtualization support (Intel VT or AMD-V). Installing a Windows VM can be broken down into two main steps:

- installing the Windows operating system
- installing the *paravirtualized device drivers*

Windows VMs are installed by cloning an appropriate Template from either XenCenter or the CLI. The Templates for individual guests have predefined platform flags set which define the configuration of the virtual hardware. For example, all Windows VMs are installed with the ACPI Hardware Abstraction Layer (HAL). If you subsequently change one of these VMs to have multiple virtual CPUs, Windows automatically switches the HAL to multi-processor mode.

The list of pre-defined Windows templates are:

- **Windows Server 2003:** can be used to install Windows Server 2003 32-bit SP0, SP1, and SP2. The Server, Enterprise, Data Centre, and SBS editions are supported.
- **Windows Server 2003 x64:** can be used to install Windows Server 2003 64-bit. The Server, Enterprise, Data Centre, and SBS editions are supported.
- **Windows 2000 SP4:** can be used to install Windows 2000 Server Service Pack 4. Earlier service packs are not supported.
- **Windows XP SP2:** can be used to install Windows XP Service Pack 2. Earlier service packs are not supported.

The Windows VM can be installed either from a install CD in a physical CD-ROM on the XenServer Host, or from an ISO image of your Windows media (see Appendix A for information on how to make an ISO image from a Windows install CD and make it available for use).

3.1 Making the ISO available to XenServer Hosts

To make an ISO library available to XenServer Hosts, create an external NFS or SMB/CIFS share directory. The NFS or SMB/CIFS server must be set to allow root access to the share. For NFS shares, this is accomplished by setting the *no_root_squash* flag when you create the share entry in `/etc/exports` on the NFS server.

Then either use XenCenter to attach the ISO library, or connect to the host console and type in:

```
xe-mount-iso-sr host:/volume
```

Additional arguments to the mount command may be passed in, for advanced use. If making a Windows SMB/CIFS share available to the XenServer host, either use XenCenter to make it available, or connect to the host console and type in:

```
xe-mount-iso-sr unc_path -t smbfs -o username=myname/myworkgroup
```

The `unc_path` argument should have back-slashes replaced by forward-slashes. **-t cifs** can be used for CIFS instead of SMB. Examples:

```
xe-mount-iso-sr //server1/myisos -t cifs -o username=johndoe/mydomain
xe-mount-iso-sr //server2/iso_share -t smbfs -o username=alice
```

After mounting the share, any ISOs in it should be available by name from the CD pulldown list in XenCenter, or as CD images from the CLI commands. The ISO should be attached to an appropriate Windows template:

- Windows Server 2003
- Windows Server 2003 x64
- Windows 2000 SP4
- Windows XP SP2

3.1.1 Copying ISOs to local storage

In XenServer 3.2 and earlier, ISOs could be copied directly to the control domain into the `/opt/xen/source/packages/iso` directory. In XenServer 4.0.1 hosts, this directory is reserved for use of the built-in ISO images, and is *not intended for general use*. This directory is considered to be identical across hosts in a Resource Pool, and CD images may fail to attach if the contents are modified.

To use local ISO storage from the control domain:

1. Log onto the host console.
2. Create a directory to copy the local ISOs into: **mkdir -p /var/opt/xen/iso_import**
3. Create an ISO storage repository by: **xe-mount-iso-sr /var/opt/xen/iso_import -o bind**
4. Copy the ISO images into this directory, taking care not to fill up the control domain filesystem.
5. Verify that the ISO image is available for use by **xe vdi-list**, or checking the CD drop-down box in XenCenter.



Be extremely careful with copying ISOs directly onto the control domain filesystem, as it has limited space available. A network share is a much safer mechanism for storing large numbers of ISO images. If the control domain does fill up, unpredictable behavior will result.

3.2 Windows paravirtualized drivers

The XenSource paravirtualized network and SCSI drivers provide high performance I/O services without the overhead of traditional device emulation found in first-generation virtualization products. During the installation of a Windows operating system, Xen will use traditional device emulation to present a standard IDE controller and a standard network card to the Virtual Machine. This allows Windows to complete its installation using built-in drivers, but with reduced performance due to the overhead inherent in emulation of the controller drivers.

After Windows is installed, you install the XenSource high-speed paravirtualized drivers. These are on an ISO available to the virtual CD-ROM drive of the Virtual Machine. These drivers replace the emulated devices and provide high-speed transport between Windows and the XenServer product family software.



While a Windows VM will function without them, performance is significantly hampered unless these drivers are installed. Running Windows VMs without these drivers is *not* supported. Some features, such as live relocation across physical hosts, will only work with the paravirtual drivers installed and active.

The Windows paravirtualized drivers ISO can be attached to the VM by using the "Install Tools" menu in XenCenter, or by directly inserting the built-in `xs-tools.iso` ISO image to the VM using the CLI. Once the ISO is attached, double-click on the `xensetup.exe` installer executable and follow the on-screen prompts.

3.3 Remote Desktop

The graphical console for Windows can be either a standard console via emulated graphics card, or an RDP connection. For Windows VMs, there is a Switch to Remote Desktop button on the Console tab. Clicking it disables the standard graphical console, and switches to using Remote Desktop instead.

The button will be greyed out if you do not have Remote Desktop enabled in the VM, and the paravirtualized drivers must be installed.

To enable Remote Desktop on a Windows VM:

1. From the Start menu, select Control Panel.
2. From the Control Panel window, select System.
3. In the System Properties dialog box, select the Remote tab.
4. In the Remote Desktop section of this dialog box, check the checkbox labeled Allow users to connect remotely to this computer (Windows XP) or Enable Remote Desktop on this computer (Windows 2003 Server).
5. If you want to select any non-administrator users that can connect to this Windows VM, click the Select Remote Users... button and provide the usernames. (Users with Administrator privileges on the Windows domain can connect by default.)

3.4 Preparing to clone a Windows VM

You need to use the Windows utility **sysprep** to prepare a Windows VM for cloning. This is the only supported way to properly clone a Windows VM.

Computers running Windows operating systems use a Security ID (SID) to uniquely identify themselves. When cloning a Windows VM, it is important to take steps to ensure the uniqueness of these Security IDs. Cloning an installation without taking the recommended system preparation steps can lead to duplicate SIDs and other problems. Because the SID identifies the computer or domain as well as the user, it is critical that it is unique. Refer to the [Microsoft KnowledgeBase article 162001, "Do not disk duplicate installed versions of Windows,"](#) for more information.

Sysprep modifies the local computer Security ID (SID) to make it unique to each computer. The Sysprep binaries are on the Windows product CDs in the `\support\tools\deploy.cab` file.

Here are the overall steps you need to follow to clone Windows VMs:

1. Create, install, and configure the Windows VM as desired.
2. Apply all relevant Service Packs and updates.

3. Install the XenSource PV drivers.
4. Install any applications and perform any other tailoring that is desired.
5. Copy the contents of `\support\tools\deploy.cab` from the Windows product CD to a new `\sysprep` folder in the VM.
6. Run `sysprep` (this will shut down the VM when it completes).
7. In XenCenter, convert the VM into a template. From the CLI, this can be done by setting the **is-a-template** parameter to true on the VM.
8. Clone the newly created template into new VMs as required.
9. When the cloned VM starts, it will get a new system ID and name, then run a mini-setup to prompt for configuration values as necessary, and finally restart, before being available for use.



The original, sysprepped VM (the "source" VM) should *not* be restarted again after the sysprep stage, and should be converted to a template immediately afterwards to prevent this. If the source VM is restarted, sysprep must be run on it again before it can be safely used to make additional clones.

For more information on using sysprep, refer to the Microsoft TechNet page "[Windows System Preparation Tool.](#)"

3.5 Release Notes

There are many versions and variations of Windows with different levels of support for the features provided by XenServer. This section lists notes and errata for the known differences.

3.5.1 General Windows Issues

Multiple VCPUs are exposed as CPU sockets to Windows guests, and are subject to the licensing limitations present in the guest. The number of CPUs present in the guest can be confirmed by checking the Device Manager. The number of CPUs actually being used by Windows can be seen in the Task Manager.

The disk enumeration order in a Windows guest may differ from the order in which they were initially added. This is a behavioral artifact between the paravirtualized drivers and the PnP subsystem in Windows. For example, the first disk may show up as "Disk 1", the next disk hotplugged as "Disk 0", a subsequent disk as "Disk 2", and then upwards in the expected fashion.

3.5.2 Windows 2003 Server

No known issues.

3.5.3 Windows XP SP2

Windows XP does not support disks larger than 2TB (terabytes) in size. See [this article in the Windows Hardware Developer Central website.](#)

3.5.4 Windows 2000 Server

No known issues.

Chapter 4

Installing Linux VMs

XenServer supports the installation of many Linux distributions into paravirtualized VMs. There are three installation mechanisms at present: complete distributions provided as built-in templates, Physical-to-Virtual (P2V) of an existing native install (see Section 2.4), and using the vendor media to perform a network installation. All use of Linux VMs requires the Linux Pack to be installed onto the XenServer Host.

The supported Linux distributions are:

Distribution	Built-in	P2V	Vendor Install
Debian Sarge 3.1	yes	no	no
Debian Etch 4.0	yes	no	no
Red Hat Enterprise Linux 3.6-3.8	no	yes	no
Red Hat Enterprise Linux 4.1-4.5	no	yes	yes
Red Hat Enterprise Linux 5.0	no	no	yes
SUSE Linux Enterprise Server 9	no	yes	no
SUSE Linux Enterprise Server 10 SP1	no	no	yes
CentOS 4.5	no	no	yes
CentOS 5.0	no	no	yes

Distributions which use the same installation mechanism as Red Hat Enterprise Linux 5 (e.g. Fedora Core 6) may successfully install using the same template. However, use of distributions not present in the above list is *unsupported*.

4.1 Installation of a built-in distribution

This is the simplest way of installing a VM. The template provided with XenServer can be used to directly create a VM running version 3.1 (Sarge) or 4.0 (Etch) of the Debian Linux distribution without need for vendor installation media and without performing a P2V conversion of an existing physical server.

The VMs are instantiated by using the **vm-install** from the CLI, or by cloning the template using XenCenter. For example, using the CLI on Linux:

```
# xe vm-install template=Debian\ Etch\ 4.0 new-name-label=ExampleVM  
f21cd819-5b7d-002d-7a1e-861a954e770
```

When the VM is first booted, it will prompt you for a root password, a VNC password (for graphical use), and a hostname. After values are entered for these, it will finish at a standard login prompt, ready for use. You will need to add a network interface if installed via the CLI.

4.2 Installing Linux from a network installation server to a VM

The XenServer guest installer allows you to install an operating system from a network-accessible copy of vendor media into a VM. In preparation for installing from vendor media, you need to make an exploded network repository of your vendor media (*not* ISO images), exported via NFS, HTTP or FTP accessible to the XenServer Host administration interface. See Appendix B for information on how to copy a set of installation CDs to a network drive. For further information, see the section "Preparing for a Network Installation" in the *Red Hat Enterprise Linux 4 Installation Guide* for details.

The network repository must be accessible from the Control Domain of the XenServer host, normally via the management interface. The URL should point to the base of the CD/DVD image on the network server, and be of the form:

- **HTTP:** `http://<server>/<path>`
- **FTP:** `ftp://<server>/<path>`
- **NFS:** `nfs:<server>:/<path>`

The XenServer "New VM Wizard" provides an additional step for vendor-installable templates which prompts for the repository URL. When using the CLI, install the template as normal using **vm-install** and then set the **other-config-install-repository** key to the value of the URL. When the VM is subsequently started, it will begin the network installation process.



When installing a new Linux-based VM, it is important to fully finish the installation and reboot it before performing any other operations on it. This is analogous to not interrupting a Windows installation, which would leave you with a non-functional VM.

4.3 Physical-to-Virtual Installation of a Linux VM

Older Linux distributions such as Red Hat Linux Enterprise 3.6 do not support Xen directly, and are typically legacy installations which benefit from virtualization for the purposes of server consolidation or hardware upgrades. The XenServer P2V feature analyzes existing installations and converts them into VMs.

When an installation is converted into a VM using P2V (see Section 2.4), the kernel used is also automatically switched to a Xen paravirtualized kernel. XenServer contains ports of the Red Hat Enterprise Linux 3/4 and SUSE Enterprise Linux 9 kernels to support the native Xen hypervisor interface directly. These kernels are present in the built-in `xs-tools.iso` image in the default CD list, or via the Install XenSource Tools command in the VM menu in XenCenter.

4.3.1 Guest Installation Network

During the installation of a VM via P2V, a special network is used to assign a temporary IP address to the VM to enable the installation to proceed. It is possible that the range of IP addresses used might conflict with real IP addresses already in use in your network. The default range of IP addresses is 192.168.128.1 to 192.168.128.254, and the default netmask is 255.255.255.0.

To change the guest installer network values:

1. Open a text console on the XenServer Host or install the CLI for remote use.
2. Find the guest installer network:

```
xe network-list
```

The command will return the list of networks available to the XenServer Host. The one you want has the name-label *Guest installer network*.

3. Examine the other-config parameters of the guest installer network:

```
xe network-param-list uuid=<UUID of the guest installer network>
```

The command will a subset of the guest installer network's parameters, including the *other-config* parameter. If the values are set to the default described above, you will see the line:

```
other-config (MRW): is_guest_installer_network: true; ip_begin: 169.254.0.1; ip_end: ←
169.254.255.254; netmask: 255.255.0.0
```

4. To change the IP address range the guest installer network will use, edit the *ip_begin*, *ip_end*, and *netmask* values as follows:

```
xe network-param-set uuid=<UUID of the guest installer network> other-config:ip_begin=< ←
desired IP range beginning> other-config:ip_end=<desired IP range end> other-config: ←
netmask=<desired netmask>
```

Do *not* change the value of the parameter *is_guest_installer_network*.

4.4 Installing the Linux guest agent

Although all the supported Linux distributions are natively paravirtualized (and thus do not need special drivers for full performance), XenServer includes a guest agent which provides additional information about the VM to the host. This additional information includes:

- Linux distribution name and version (major, minor revision).
- Kernel version (uname).
- IP address of each Ethernet interface.
- Total and free memory within the VM.

It is important to install this agent and keep it up-to-date (see Chapter 5) as you upgrade your XenServer host.

To install the guest agent:

1. The files required are present on the built-in `xs-tools.iso` CD image, or alternatively by using the 'Install Tools' option in XenCenter.
2. Mount the image into the guest via:

```
mount /dev/xvdd /mnt
```

3. Execute the installation script as the root user:

```
/mnt/Linux/install.sh
```

4. If the VM was upgraded from XenServer 3.1, you will also need to run:

```
chkconfig xengmond off
```

5. If the kernel has been upgraded, or the VM was upgraded from a previous version, reboot the VM now.



CD-ROM drives and ISOs attached to Linux Virtual Machines appear as `/dev/xvdd` instead of as `/dev/cdrom` as you might reasonably expect. This is because they are not "true" CD-ROM devices, but normal devices. When the CD is ejected by either XenCenter or the CLI, it hot-unplugs the device from the VM and the device disappears. This is different from Windows Virtual Machines, where the CD remains in the VM in an empty state.

4.5 Preparing to clone a Linux VM

When a Linux VM is cloned, some virtual hardware parameters are changed in the new VM. The VM may need to be customized in order to be aware of these changes. For instructions for specific supported Linux distributions, please see Section [4.8](#).

4.5.1 Machine Name

Of course, a cloned VM is another computer, and like any new computer in a network, it must have a unique name within the network domain it is part of.

4.5.2 IP address

A cloned VM must have a unique IP address within the network domain it is part of. This is not a problem in general if DHCP is used to assign addresses; when the VM boots the DHCP server will assign it an IP address. If the cloned VM had a static IP address, the clone must be given an unused IP address before being booted.

4.5.3 MAC address

In some cases, the MAC address of a cloned VM's virtual network interface is recorded in the network configuration files. After the VM is cloned, the new cloned VM has a different MAC address. Therefore, when started, the network does not come up automatically.

Some Linux distributions use udev rules to remember the MAC address of each network interface, and persist a name for that interface. This is intended so that the same physical NIC always maps to the same `ethn` interface, which is particularly useful with removable NICs (like laptops). But this behavior is problematic in the context of Virtual Machines. For example, if you configure two virtual NICs when you install a VM, and then shut it down and remove the first NIC, on reboot XenCenter shows just one NIC, but calls it `eth0`. Meanwhile the VM is deliberately forcing this to be `eth1`. The result is that networking doesn't work.

If the VM uses persistent names, the best thing to do is to turn these rules off. If for some reason you do not want to turn persistent names off, be aware that you will need to reconfigure networking inside the VM in the usual way, and the information shown in XenCenter will be out of synch with reality.

4.6 Time handling in Linux VMs

By default, the clocks in a Linux VM are synchronized to the clock running on the control domain, and cannot be independently changed. This mode is a convenient default, since only the control domain needs to be running the NTP service to keep accurate time across all VMs. Upon installation of a new Linux VM, make sure you change the time-zone from the default UTC to your local value (see Section [4.8](#) for specific distribution instructions).

Individual Linux VMs can be set to maintain independent times:

1. From a root prompt on the VM, type the command: **echo 1 > /proc/sys/xen/independent_wallclock**
2. This can be persisted across reboots by changing the `/etc/sysctl.conf` configuration file and adding:

```
# Set independent wall clock time
xen.independent_wallclock=1
```

3. As a third alternative, the `independent_wallclock=1` may also be passed as a boot parameter to the VM.

4.7 Configuring VNC for VMs

With the exception of VMs based on the Debian templates, VMs might not be set up to support VNC by default. For example, if you P2V a server that does not have a VNC server installed, the resulting VM won't have VNC installed either. Before you can connect with the XenCenter graphical console, you need to ensure that the VNC server and an X display manager are installed on the VM and properly configured. This section describes the procedures for configuring VNC on each of the supported Linux operating system distributions to allow proper interactions with the XenCenter graphical console.

CentOS-based VMs should use the instructions for the Red Hat-based VMs below, as they use the same base code to provide graphical VNC access. CentOS 4 is based on Red Hat Enterprise Linux 4, and CentOS 5 is based on Red Hat Enterprise Linux 5.

4.7.1 Setting up Red Hat-based VMs for VNC



Before setting up your Red Hat VMs for VNC, be sure that you have installed the Linux guest agent. See Section 4.4 for details.

In order to configure VNC on Red Hat VMs, you need to modify the GDM configuration. The GDM configuration is held in a file whose location varies depending on the version of Red Hat Linux you are using. Before modifying it, we first must determine the location of this configuration file; this file will then be modified in a number of subsequent procedures in this section.

4.7.1.1 Determining the location of your VNC configuration file

If you are using Red Hat Linux version 3 or 4 the GDM configuration file is `/etc/X11/gdm/gdm.conf`. This is a unified configuration file that contains both default values as specified by the provider of your version of GDM in addition to your own customized configuration. This type of file is used by default in older versions of GDM, as included in these version of Red Hat Linux.

If you are using Red Hat Linux version 5 the GDM configuration file is `/etc/gdm/custom.conf`. This is a split configuration file that contains only user-specified values that override the default configuration. This type of file is used by default in newer versions of GDM, as included in these versions of Red Hat Linux.

4.7.1.2 Configuring GDM to use VNC

1. As root at the prompt in the VM's text console, type in **rpm -q vnc-server gdm**. The package names `vnc-server` and `gdm` should appear, with their version numbers specified.

If these package names are displayed, the appropriate packages are already installed. If you see a message saying that one of the packages is not installed, then you may not have selected the graphical desktop options during installation. You will need to install these packages before you can continue. See the appropriate *Red Hat Linux x86 Installation Guide* for details regarding installing additional software on your VM.

2. Open the GDM configuration file with your preferred text editor and add the following lines to the file:

```
[server-VNC]
name=VNC Server
command=/usr/bin/Xvnc -SecurityTypes None -geometry 1024x768 -depth 16 -BlacklistTimeout 0
flexible=true
```

- With configuration files as found on Red Hat Linux 3 and 4, this should be added above the `[server-Standard]` section.
- With configuration files as found on Red Hat Linux 5, this should be added into the empty `[servers]` section.

3. Modify the configuration so that the Xvnc server is used instead of the standard X server:

- If you are using Red Hat Linux 3 or 4, there will be a line just above that says:

```
0=Standard
```

Modify it to read:

```
0=VNC
```

- If you are using Red Hat Linux 5 or greater, you will need to add the above line just below the `[servers]` section and before the `[server-VNC]` section.

4. Save and close the file.

Restart GDM for your change in configuration to take effect, by running `/usr/sbin/gdm-restart`.

Note that, for Red Hat Linux, runlevel 5 is used for graphical startup. If your installation is configured to start up in runlevel 3, you will need to change this in order for the display manager to be started (and therefore to get access to a graphical console). Please refer to Section 4.7.4 for further details.

4.7.1.3 Firewall settings

The firewall configuration by default does not allow VNC to traffic to go through. If you have a firewall between the VM and XenCenter, you need to allow traffic over the port that the VNC connection uses. By default, a VNC server listens for connections from a VNC viewer on TCP port 5900 + N, where N is the display number (usually just zero). So a VNC server setup for Display-0 will listen on TCP port 5900, Display-1 is TCP-5901, etc. Consult your firewall documentation to make sure these ports are open.

You might want to further customize your firewall configuration if you want to use IP connection tracking or limit the initiation of connections to be from one side only.

To customize Red Hat-based VMs firewall to open the VNC port:

1. For Red Hat Linux 3, use **redhat-config-securitylevel-tui**.
For Red Hat Linux 4 and 5, use **system-config-securitylevel-tui**.
2. Select 'Customize' and add 5900 to the other ports list.

Alternatively, you can disable the firewall until the next reboot by using **service iptables stop**, or permanently by using **chkconfig iptables off**. This can of course expose additional services to the outside world and reduce the overall security of your VM.

4.7.1.4 VNC screen resolution

If, after connecting to a Virtual Machine with the graphical console, the screen resolution is mismatched (for example, the VM's display is too big to comfortably fit in the Graphical Console pane), you can control it by setting the VNC server's `-geometry` parameter as follows:

1. Open the GDM configuration file with your preferred text editor. Please refer to Section 4.7.1.1 for information about determining the location of this file.
2. Find the `[server-VNC]` section you added above.
3. Edit the command line to read, for example,

```
command=/usr/bin/Xvnc -SecurityTypes None -geometry 800x600
```

where the value of the `-geometry` parameter can be any valid screen width and height.

4. Save and close the file.

4.7.2 Setting up SLES-based VMs for VNC



Before setting up your SUSE Linux Enterprise Server VMs for VNC, be sure that you have installed the Linux guest agent. See Section 4.4 for details.

SLES has support for enabling 'Remote Administration' as a configuration option in YaST. You can select to enable Remote Administration at install time, available on the Network Services screen of the SLES installer. This will allow you to connect an external VNC viewer to your guest to view the graphical console; the methodology for using the SLES remote administration feature is slightly different than that provided by XenCenter, but it is possible to modify the configuration files in your SUSE Linux VM such that it is integrated with the graphical console feature.

4.7.2.1 Checking for a VNCserver

Before making configuration changes, you should verify that you have a VNC server installed. SUSE ships the `tightvnc` server by default; this is a suitable VNC server, but you can also use the standard RealVNC distribution if you prefer.

You can check that you have the `tightvnc` software installed by running the command:

```
rpm -q tightvnc
```

4.7.2.2 Enabling Remote Administration

If Remote Administration was not enabled during installation of the SLES software, you can enable it as follows:

1. Open a text console on the VM and run the YaST utility:

```
# yast
```



Due to the complex control characters used to draw the YaST configuration screens, your screen usually becomes corrupted while using the text-mode YaST configuration tools. For example, you will be unable to see portions of the display, for example. In the steps that follow, use the key combination **Ctrl+L** to redraw the display and remove artifacts when necessary.

2. Use the arrow keys to select Network Services in the left menu, then **Tab** to the right menu and use the arrow keys to select Remote Administration. Press **Enter**.
3. In the Remote Administration screen, **Tab** to the Remote Administration Settings section. Use the arrow keys to select Allow Remote Administration and press **Enter** to place an X in the checkbox.
4. **Tab** to the Firewall Settings section. Use the arrow keys to select Open Port in Firewall and press **Enter** to place an X in the checkbox.
5. **Tab** to the Finish button and press **Enter**.
6. A message box appears telling you that you will need to restart the display manager for your settings to take effect. Press **Enter** to acknowledge the message.
7. The original top-level menu of YaST appears. **Tab** to the Quit button and press **Enter**.

4.7.2.3 Modifying the xinetd configuration

After enabling Remote Administration, you need to modify a configuration file if you want to allow XenCenter to connect, or else use a third party VNC client.

1. Open the file `/etc/xinetd.d/vnc` in your preferred text editor.

The file contains sections like the following:

```
service vnc1
{
  socket_type = stream
  protocol   = tcp
  wait       = no
  user       = nobody
  server     = /usr/X11R6/bin/Xvnc
  server_args = :42 -inetd -once -query localhost -geometry 1024x768 -depth 16
  type      = UNLISTED
  port      = 5901
}
```

2. Make the following changes:
 - edit the `port` line to read `port = 5900`
 - edit the `server_args` line to include the argument `-BlacklistTimeout 0` at the end
3. Save and close the file.
4. Restart the display manager and xinetd service with the following commands:

```
/etc/init.d/xinetd restart
rcxdm restart
```

SUSE Linux uses runlevel 5 for graphical startup. If your remote desktop does not appear, verify that your VM is configured to start up in runlevel 5. Refer to Section 4.7.4 for details.

4.7.2.4 Firewall settings

The firewall configuration by default does not allow VNC to traffic to go through. If you have a firewall between the VM and XenCenter, you need to allow traffic over the port that the VNC connection uses. By default, a VNC server listens for connections from a VNC viewer on TCP port 5900 + N, where N is the display number (usually just zero). So a VNC server setup for Display-0 will listen on TCP port 5900, Display-1 is TCP-5901, etc. Consult your firewall documentation to make sure these ports are open.

You might want to further customize your firewall configuration if you want to use IP connection tracking or limit the initiation of connections to be from one side only.

To customize SLES-based VMs firewall to open the VNC port:

1. Open a text console on the VM and run the YaST utility:

```
# yast
```



Due to the complex control characters used to draw the YaST configuration screens, your screen usually becomes corrupted while using the text-mode YaST configuration tools. For example, you will be unable to see portions of the display, for example. In the steps that follow, use the key combination **Ctrl+L** to redraw the display and remove artifacts when necessary.

2. Use the arrow keys to select Security and Users in the left menu, then **Tab** to the right menu and use the arrow keys to select Firewall. Press **Enter**.
3. In the Firewall screen, **Tab** to the Firewall Configuration: Settings section. Use the arrow keys to select the Allowed Services in the left menu.
4. **Tab** to the Firewall Configuration: Allowed Services fields on the right. Use the arrow keys to select the Advanced... button (near the bottom right, just above the Next button) and press **Enter**.
5. In the Additional Allowed Ports screen, type *5900* in the TCP Ports field. **Tab** to the OK button and press **Enter**.
6. **Tab** back to the list of screens on the left side and use the arrow keys to select Start-Up. **Tab** back to the right and **Tab** to the Save Settings and Restart Firewall Now button and press **Enter**.
7. **Tab** to the Next button and press **Enter**, then in the Summary screen **Tab** to the Accept button and press **Enter**, and finally on the top-level YaST screen **Tab** to the Quit button and press **Enter**.
8. Restart the display manager and xinetd service with the following commands:

```
/etc/init.d/xinetd restart
rcxdm restart
```

Alternatively, you can disable the firewall until the next reboot by using the **rcSuSEfirewall2 stop** command, or permanently by using YaST. This can of course expose additional services to the outside world and reduce the overall security of your VM.

4.7.2.5 VNC screen resolution

If, after connecting to a Virtual Machine with the Graphical Console, the screen resolution is mismatched (for example, the VM's display is too big to comfortably fit in the Graphical Console pane), you can control it by setting the VNC server's `-geometry` parameter as follows:

1. Open the `/etc/xinetd.d/vnc` file with your preferred text editor and find the `service_vnc1` section (corresponding to displayID 1).

2. Edit the **-geometry** argument in the **server-args** line to the desired display resolution. For example,

```
server_args = :42 -inetd -once -query localhost -geometry 800x600 -depth 16
```

where the value of the `-geometry` parameter can be any valid screen width and height.

3. Save and close the file.

4. Restart the vnc server:

```
/etc/init.d/xinetd restart  
rcxdm restart
```

4.7.3 Setting up Debian-based VMs for VNC

The built-in Debian Sarge and Etch templates come pre-configured with VNC setup and ready use. However, the default VNC configuration in Debian does not permit the root administration user to log in by default. To log in by VNC, you can either:

- Log in to the text console and create a new, unprivileged user via the **adduser** command. This is the recommended course of action.
- At the graphical console login prompt, select Actions, Configure the Login Manager, type in your root password, then select Security, Allow local system administrator login, and finally select Close.

If you need to reset the VNC password, use the command

```
vnc4passwd /etc/vncpass
```

4.7.4 Checking runlevels

Red Hat and SUSE Linux VMs use runlevel 5 for graphical startup. This section describes how to verify that your VM is configured to start up in runlevel 5 and how to change it if it is not.

1. Check `/etc/inittab` to see what the default runlevel is set to. Look for the line that reads:

```
id:n:initdefault:
```

If *n* is not 5, edit the file to make it so.

2. You can run the command **telinit q ; telinit 5** after this change to avoid having to actually reboot to switch runlevels.

4.8 Release Notes

Most modern Linux distributions support Xen paravirtualization directly, but have different installation mechanisms and some kernel limitations.

4.8.1 Debian Sarge 3.1 and Etch 4.0

XenServer includes a custom Xen kernel for Debian guests installed via the built-in template. This kernel is a cut-down version of the dom0 kernel used in XenServer, and as such is the most heavily tested and reliable guest kernel available. After installation, the time-zone in a Debian VM defaults to UTC (see Section 4.6). It can be changed to your local value by using the `tzconfig` command.

To prepare a Debian guest for cloning (see Section 4.5.3), Ethernet name persistence must be disabled. For Debian Sarge VMs, name persistence is controlled through `/etc/udev/persistent-net-generator.rules` in the `udev` package. It can be disabled by removing the following symlink to that file by:

```
rm -f /etc/udev/rules.d/z45_persistent-net-generator.rules
```

4.8.2 Red Hat Enterprise Linux 3

XenServer includes a custom port of the RHEL3.8 kernel with native Xen paravirtualized guest support. This kernel is installed during the P2V process for RHEL3.6-3.8 guests. Since the kernel is based on Linux 2.4, the following limitations apply:

- Only 3 virtual network interfaces and 3 virtual block devices are supported.
- VMs with multiple VCPUs cannot be suspended. If you wish to suspend these guests, you must reduce the number of VCPUs to 1 while the guest is halted.

4.8.3 Red Hat Enterprise Linux 4

XenServer includes a custom port of the RHEL 4.5 kernel with additional bug fixes and expanded Xen support. This kernel is used in the vendor installation templates of RHEL 4.1 and 4.4, but not in the RHEL 4.5 template (since RHEL4.5 is the first release with native Xen support).

The issues below have been reported upstream to Red Hat and are already fixed in our kernel (which can be installed by using the `/mnt/Linux/install.sh` script in the built-in `xs-tools.iso` CD image):

- Only 3 virtual network interfaces and 3 virtual block devices are supported. The XenServer-supplied kernel supports up to 7.
- Occasional kernel crash on boot in `queue_work()` (Red Hat Bugzilla [246586](#))
- Disks sometimes do not attach correctly on boot (Red Hat Bugzilla [247265](#))
- Live migration can occasionally crash the kernel under low memory conditions (Red Hat Bugzilla [249867](#))
- Guest kernel can occasionally hang due to other xenstore activity (Red Hat Bugzilla [250381](#))

To prepare a RHEL4 guest for cloning (see Section 4.5.3), edit `/etc/sysconfig/network-scripts/ifcfg-eth0` before converting the VM into a template and remove the `HWADDR` line. Note that Red Hat recommend the use of Kickstart to perform automated installations, instead of directly cloning disk images (see [Red Hat KB Article 2415](#)).

4.8.4 Red Hat Enterprise Linux 5

XenServer uses the standard Red Hat kernel supplied with RHEL5 as the guest kernel. Any bugs found in this kernel are reported upstream to Red Hat, and are listed below:

- Only 3 virtual network interfaces and 3 virtual block devices are supported.
- Random segmentation faults on loading ELF binaries (Red Hat Bugzilla [247261](#))
- Disks sometimes do not attach correctly on boot (Red Hat Bugzilla [247265](#))

- Soft lockup messages after suspend/resume or live migration (Red Hat Bugzilla [250994](#)). These messages are harmless, but there may be a period of inactivity in the guest during live migration as a result of the lockup.
- Network blackout during live relocation for up to a minute (Red Hat Bugzilla [251527](#)). After migration has completed, the kernel sends a gratuitous ARP to cause ARP caches to get refreshed and minimise network downtime. However, carrier detect is delayed in the kernel and so there is a network blackout until the ARP caches expire or the guest generates an ARP for some other reason.

When you install the XenServer `xe-guest-utilities` RPM, it adds an entry to the `yum` configuration, allowing you to pick up kernel updates provided by XenSource as they become available.

4.8.5 CentOS 4

Please refer to Section [4.8.3](#) for the list of CentOS 4 release notes.

Unlike RHEL4, CentOS includes a third-party updates mechanism known as `yum`. The `xe-guest-utilities` RPM will install a XenServer entry for `yum`, allowing you to pick up kernel updates provided by XenSource via the standard update mechanism as they become available.

4.8.6 CentOS 5

Please refer to Section [4.8.4](#) for the list of CentOS 5 release notes.

4.8.7 SUSE Enterprise Linux 9

XenServer includes a custom port of the SLES9 SP3 kernel with additional bug fixes and expanded support. This kernel is automatically installed during the P2V process for SLES9 guests.

To prepare a SUSE Linux guest for cloning (see Section [4.5.3](#)), edit `/etc/sysconfig/network/config` and edit the line:

```
FORCE_PERSISTENT_NAMES=yes
```

to

```
FORCE_PERSISTENT_NAMES=no
```

When you P2V a SLES 9 server, the networking configuration files that were present on the physical server will remain on the VM. You may wish to move these aside, or update them accordingly, when you add virtual interfaces to the VM.

4.8.8 SUSE Enterprise Linux 10 SP1

XenServer uses the standard Novell kernel supplied with SLES10 SP1 as the guest kernel. Any bugs found in this kernel are reported upstream to Novell and listed below:

- Only 3 virtual network interfaces and 3 virtual block devices are supported.
- Disks sometimes do not attach correctly on boot. (Novell Bugzilla [290346](#)).

When installing SLES10 VMs, you may experience some display corruption in the text-mode install. This is more pronounced when using SLES10, and is partially fixed in SLES10 Service Pack 1. In either case, pressing ‘Control-L’ after the corruption happens will force a screen redraw, and improve the installation experience.

Chapter 5

Updating VMs

This chapter discusses updating VMs with new Linux kernel revisions, applying Windows Service Packs, and updates to XenSource paravirtualized drivers and VM utilities.

Upgrades to VMs are typically required when moving to a new version of XenServer. The following are current issues involving upgrading VMs running on XenServer to this version:

- XenMotion of Windows VMs is not supported until the paravirtualized drivers are upgraded.
- Suspend/Resume of Windows VMs is not supported until the paravirtualized drivers are upgraded.
- The use of certain anti-virus and firewall applications may crash the Windows VM unless the paravirtualized drivers are upgraded.

5.1 Updating paravirtualized drivers for Windows VMs

The paravirtualized drivers are present on the built-in `xs-tools.iso` (if using the CLI), or in XenCenter via the Install XenSource Tools command from the VM, which will attach a CD image containing the drivers to the VM. Either wait for the auto-run facility, or manually click on the `xensetup.exe` program. Follow the on-screen prompts to install the new drivers, which will automatically deactivate and upgrade the old drivers.

When updating Windows 2000 SP2 VMs from XenServer 3.2 to the current version, you may occasionally see a dialog popup containing 'the xennet.sys device driver could not locate the entry point `__XenTrace` in driver `xvtchn.sys`'. This warning message is harmless and installation will continue normally after clicking past the dialog.

5.2 Updating Linux kernels and guest utilities

The Linux guest utilities can be updated by re-running the `Linux/install.sh` script from the built-in `xs-tools.iso` CD image (see Section 4.4). From time to time, XenSource also supplies updated Linux kernels for supported distributions. Supported distributions are:

- Red Hat Enterprise Linux 5.x
- CentOS 5.x
- Red Hat Enterprise Linux 4.x
- CentOS 4.x
- Red Hat Enterprise Linux 3.x
- Debian Sarge and Etch

The updates are posted online at: <http://updates.xensource.com/XenServer/4.0.1/>.

For example, the RHEL 3.x kernel would be at: <http://updates.xensource.com/XenServer/4.0.1/rhel3x/>.

This is of particular importance for RHEL 4.5 and CentOS 4.5, where you will get the upstream kernel by default, which has certain limitations (see Section 4.8).

For yum-enabled distributions (CentOS 4 and 5, RHEL 5), `xe-guest-utilities` installs a yum configuration file to enable subsequent updates to be done via yum in the standard manner. Note that RHEL 4 in particular does not use yum.

For Debian, `/etc/apt/sources.list` is populated to enable updates via apt by default.

Appendix A

Creating ISO images

XenServer can use ISO images of CD-ROM or DVD-ROM disks as installation media and data sources for Windows or Linux VMs. This section describes how to make ISOs from CD/DVD media.

On a Linux computer:

1. Put the CD- or DVD-ROM disk into the drive. The disk should not be mounted. To check, type the command:
mount
If the disk is mounted, unmount the disk. Refer to your operating system documentation for assistance if required.

2. As root, type the command
dd if=/dev/cdrom of=/path/cding_filename.iso

This will take some time. When the operation is completed successfully, you should see something like

```
1187972+0 records in
1187972+0 records out
```

Your ISO file is ready.

On a Windows computer:

1. Windows computers do not have an equivalent operating system command to create an ISO. Most CD-burning tools have a means of saving a CD as an ISO file.

One simple and free utility is [ISO Recorder](#). It works on Windows XP SP2, Windows 2000, and Windows Server 2003. Once installed, you simply right-click on a CD/DVD drive and select Create image from CD from the context menu.

Appendix B

Setting Up a Red Hat Installation Server

This chapter explains how to set up a server as an installation server for Red Hat Linux.

For a server to act as a Red Hat Linux network installation server, you need space on your server to copy the entire contents of each CD onto your server. This is typically the number of CDs or ISO images times 650MB.

Ensure that the space you intend to use is formatted with your chosen filesystem and is mounted. You can check this space with the command:

```
df -h
```

B.1 Copy installation media

1. First create a directory to contain the installation files, for example `/install`
2. Mount your CD. Refer to your operating system documentation for assistance if needed. This example assumes that it is mounted at `/mnt/cdrom`:

```
mount /mnt/cdrom
```

3. Copy the data from the CD to the installation directory:

```
cp -var /mnt/cdrom/RedHat /install
```

4. Unmount the CD:

```
umount /mnt/cdrom
```

5. Remove the first CD, put in the next one, and repeat for each of your CDs you have.

B.2 Enable remote access

Next, you need to make your installation data available to other machines on the network. You can use NFS, HTTP, or FTP protocols. You can enable all three services on your server or any subset of the three.

B.2.1 NFS

To install over NFS you need to meet certain conditions on the server:

- The installation directory must be exported

To export your installation directory, edit the `/etc/exports` file and add an entry for `/install` to it:

```
/install *(ro)
```

Save the edited exports file and tell the NFS daemon to reread its configuration file:

exportfs -r

This configures the most basic read-only export to all hosts on our network. If you want to include more advanced options in your export, such as exporting to certain hosts only, or on a certain subnet only, see the man page for the exports file at exports (5).

- NFS needs to be installed and running

To check, type the command:

showmount -e hostname

Entering the showmount command without the hostname parameter will check the local system.

If NFS is not active, you will see a message similar to

showmount: ServerA: RPC: Program not registered

- Portmap needs to be running, and can be checked with the command:

service portmap status

B.2.2 FTP

To enable installing over FTP, you need to allow FTP access to the installation directory on the server. This can be either anonymous FTP access or access through a named account with a password.

If you want anonymous FTP to point to a different directory, you can use sym links to point to the installation directory on the server.

B.2.3 HTTP

If you have a web server running and want to enable HTTP access to your installation server, then add sym links from your document root to the installation server directory to grant access.

The installation server is now ready to use. Make sure you note the server name or IP address and the directory path to the installation directory you created.

Appendix C

Troubleshooting VM problems

If you experience odd behavior, application crashes, or have other issues, this chapter is meant to help you solve the problem if possible and, failing that, describes where the application logs are located and other information that can help your XenSource Solution Provider and XenSource track and resolve the issue.

Troubleshooting of installation issues is covered in the *XenServer Installation Guide*. Troubleshooting of XenServer Host issues is covered in the *XenServer Administrator's Guide*.



We recommend that you follow the troubleshooting information in this chapter solely under the guidance of your XenSource Solution Provider or XenSource Support.

XenSource provides two forms of support: you can receive free self-help support via the [Support site](#), or you may [purchase our Support Services](#) and directly submit requests by filing an online Support Case. Our free web-based resources include product documentation, a Knowledge Base, and discussion forums.

C.1 VM crashes

If you are experiencing VM crashes, it's possible that a kernel crash dump might help identify the problem. If the crash is reproducible, follow this procedure to send the dumps to XenSource.

C.1.1 Linux VMs

For Linux VMs, the crashdump behavior can be controlled through the `actions-after-crash` parameter. The following are the possible values:

Value	Description
preserve	leave the VM in a paused state (for analysis)
coredump_and_restart	record a core dump, then reboot the VM
coredump_and_destroy	record a core dump, leave VM halted
restart	no core dump, just reboot VM (this is the default)
destroy	no coredump, leave VM halted

To enable saving of Linux VM crash dumps:

1. On the XenServer Host, determine the UUID of the desired VM using the command:

```
xe vm-list name-label=<name> params=uuid --minimal
```

2. Change the *actions-after-crash* value using the **xe vm-param-set**; for example:

```
xe vm-param-set uuid=<vm_uuid> actions-after-crash=coredump_and_restart
```

C.1.2 Windows VMs

For Windows VMs, the core dump behavior cannot be controlled through the *actions-after-crash* parameter. By default Windows crash dumps are put into %SystemRoot%\Minidump in the Windows VM itself.

You can configure the VMs dump level by going into My Computer --> Properties --> Advanced --> Startup and Recovery. You can then use lomount to retrieve it.

Index

A

AMD-V (AMD hardware virtualization), [11](#), [13](#)

C

Cloning VMs, [10](#), [20](#)

Configuring VNC

firewall settings, RHEL, [22](#)

firewall settings, SLES, [25](#)

for Debian VMs, [26](#)

for Red Hat VMs, [21](#)

for SUSE VMs, [23](#)

Converting a VM to a Template, [7](#)

Creating an ISO image, [31](#)

Creating VMs

converting VM to a Template, [7](#)

from pre-configured Template, [7](#)

importing an exported VM, [7](#)

installing OS from a CD or ISO, [7](#)

installing OS from a network repository, [7](#)

Linux, [17](#)

overview, [7](#)

physical to virtual conversion (P2V), [7](#)

Windows, [7](#), [13](#)

D

Drivers, Linux paravirtualized, [19](#)

Drivers, Windows paravirtualized, [14](#)

F

Firewall settings and VNC, [22](#), [25](#)

G

General guidelines for virtualizing physical servers, [9](#)

Guest agent, Linux, [19](#)

H

Hardware abstraction layer (HAL), [13](#)

I

Importing VMs, [7](#), [10](#)

Installation server, for installing Red Hat VMs, [33](#)

Installing VMs

by P2V, [8](#)

installing from vendor media, [17](#), [18](#)

Linux, [17](#)

Windows, [13](#)

Intel VT (Intel hardware virtualization), [11](#), [13](#)

ISO image

creating, [31](#)

Making available to XenServer Hosts, [13](#)

L

Limitations on virtual devices, [8](#)

Linux

Creating VMs, [17](#)

guest agent, [19](#)

runlevels, [26](#)

time handling, in VMs, [20](#)

M

Making an ISO image available to XenServer Hosts, [13](#)

Minimums, virtual disk space, [7](#)

Minimums, virtual memory, [7](#)

N

Network installation server, [17](#), [18](#)

NFS server, mounting ISO from, [13](#)

P

P2V, [7](#), [8](#)

general guidelines for virtualizing physical servers, [9](#)

guest installation network, [18](#)

Linux, [8](#), [9](#), [17](#), [18](#)

p2v-legacy option, [9](#)

Windows, [8](#)

Paravirtualized drivers

Windows, [14](#)

Physical to virtual conversion, *see* P2V

R

Red Hat installation server, for installing VMs, [33](#)

Release notes

Linux VMs, [26](#)

Windows VMs, [16](#)

Remote Administration, SUSE Linux, [23](#)

Remote Desktop, [15](#)

Runlevels, Linux, [26](#)

S

Setting up a Red Hat installation server, for installing VMs, [33](#)

SMB/CIFS share, mounting ISO from, [14](#)

Sysprep, for preparing Windows VM for cloning
sysprep, [15](#)

T

Template

definition of, [7](#)

Linux VMs, [7](#), [17](#)

pre-configured (Debian), [7](#)

Windows, [13](#)

Windows VMs, [7](#), [13](#)

Time handling, in Linux VMs

time handling, in VMs, [20](#)

Troubleshooting

Linux VM problems, [35](#)

Windows VM problems, [36](#)

V

- Virtual devices, limitations on, [8](#)
- Virtual disk size minimums, [7](#)
- Virtual memory minimums, [7](#)
- Virtualizing physical servers, general guidelines for, [9](#)

VMs

- Cloning, [10](#)
- general guidelines for virtualizing physical servers, [9](#)
- importing, [10](#)
- installing by P2V, [8](#)
- installing operating system from network installation server, [17](#), [18](#)
- installing Windows, [13](#)
- non-paravirtualized (Windows), [13](#)
- Paravirtualized, [18](#)
- paravirtualized, [17](#)
- paravirtualized (Linux), [9](#)
- Remote Desktop, [15](#)

VNC

- checking runlevels, [26](#)
- firewall settings, RHEL, [22](#)
- firewall settings, SLES, [25](#)
- for Debian VMs, [26](#)
- for Red Hat VMs, [21](#)
- for SUSE VMs, [23](#)

- VT (Intel hardware virtualization), [11](#), [13](#)

W

Windows

- ACPI HAL, [13](#)
- Creating VMs, [13](#)
- hardware abstraction layer (HAL), [13](#)
- list of VM Templates, [13](#)
- multi-processor HAL, [13](#)
- paravirtualized drivers, [14](#)
- Remote Desktop, [15](#)
- SMB/CIFS share, mounting ISO from, [14](#)
- sysprep, [15](#)

X

- XenSource product family, differences, [8](#)