

# SUSE Linux Enterprise Server

10 SP2

[www.novell.com](http://www.novell.com)

April 30, 2008

Virtualization with Xen



# Virtualization with Xen

All content is copyright © Novell, Inc.

## Legal Notice

This manual is protected under Novell intellectual property rights. By reproducing, duplicating or distributing this manual you explicitly agree to conform to the terms and conditions of this license agreement.

This manual may be freely reproduced, duplicated and distributed either as such or as part of a bundled package in electronic and/or printed format, provided however that the following conditions are fulfilled:

That this copyright notice and the names of authors and contributors appear clearly and distinctively on all reproduced, duplicated and distributed copies. That this manual, specifically for the printed format, is reproduced and/or distributed for noncommercial use only. The express authorization of Novell, Inc must be obtained prior to any other use of any manual or part thereof.

For Novell trademarks, see the Novell Trademark and Service Mark list <http://www.novell.com/company/legal/trademarks/tmlist.html>. \* Linux is a registered trademark of Linus Torvalds. All other third party trademarks are the property of their respective owners. A trademark symbol (®, ™ etc.) denotes a Novell trademark; an asterisk (\*) denotes a third party trademark.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither Novell, Inc., SUSE LINUX Products GmbH, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

# Contents

<b>Preface</b>	<b>v</b>
<b>1 Introduction to Xen Virtualization</b>	<b>1</b>
1.1 Basic Components . . . . .	1
1.2 Understanding Virtualization Modes . . . . .	2
1.3 Xen Virtualization Architecture . . . . .	3
1.4 The Virtual Machine Host . . . . .	4
<b>2 Setting Up a Virtual Machine Host</b>	<b>7</b>
2.1 Best Practices and Suggestions . . . . .	8
2.2 Setting Up a 32-Bit Domain 0 on a 64-Bit Hypervisor . . . . .	9
2.3 Managing Domain 0 Memory . . . . .	11
2.4 Starting the Virtual Machine Host . . . . .	12
<b>3 Setting Up Virtual Machines</b>	<b>15</b>
3.1 Creating a Virtual Machine . . . . .	15
3.2 Installing an Operating System . . . . .	17
3.3 Including Add-On Products in the Installation . . . . .	21
3.4 Using the Command Line to Create Virtual Machines . . . . .	22
3.5 Deleting Virtual Machines . . . . .	23
3.6 Troubleshooting . . . . .	23
<b>4 Managing a Virtualization Environment</b>	<b>25</b>
4.1 Virtual Machine Manager . . . . .	26
4.2 Controlling the Host by Modifying xend Settings . . . . .	27
4.3 Configuring a Virtual Machine by Modifying its xend Settings . . . . .	28

4.4	The <code>xm</code> Command . . . . .	29
<b>5</b>	<b>Virtualization: Configuration Options and Settings</b>	<b>31</b>
5.1	Virtual CD Readers . . . . .	31
5.2	Remote Access Methods . . . . .	36
5.3	VNC Viewer . . . . .	37
5.4	The Boot Loader Program . . . . .	40
5.5	Changing the Host's Desktop from Text Mode to GUI Mode . . . . .	43
5.6	Paravirtual Mode and Journaling File Systems . . . . .	44
5.7	Virtual Machine Initial Startup Files . . . . .	45
5.8	Sparse Image Files and Disk Space . . . . .	47
5.9	Virtual Keyboards . . . . .	47
5.10	Mapping Physical Storage to Virtual Disks . . . . .	50
5.11	Virtual Disks . . . . .	51
5.12	Migrating Virtual Machines . . . . .	52
5.13	Passing Key Combinations to Virtual Machines . . . . .	52
5.14	<code>xm block-list</code> Device Numbers . . . . .	53
5.15	File-Backed Virtual Disks and Loopback Devices . . . . .	54
5.16	Saving Virtual Machines . . . . .	54
5.17	Restoring Virtual Machines . . . . .	55
5.18	Virtual Machine States . . . . .	57
<b>6</b>	<b>Xen Guest Systems</b>	<b>59</b>
6.1	NetWare Virtual Machines . . . . .	59
6.2	OES Linux Virtual Machines . . . . .	59
6.3	SUSE Linux Virtual Machines . . . . .	59
6.4	Windows Virtual Machines . . . . .	68
<b>7</b>	<b>Virtual Machine Drivers</b>	<b>69</b>
7.1	Buying a Support Subscription . . . . .	70
7.2	Loading the Driver Pack . . . . .	70
7.3	Using the Installation Wizard to Install the Driver Pack . . . . .	71
7.4	SUSE Linux Enterprise Server 9 . . . . .	79
7.5	SUSE Linux Enterprise Server 10 . . . . .	83

# Preface

This manual offers an introduction to virtualization technology of your SUSE Linux Enterprise Server. It features an overview of the various fields of application and installation types of each of the platforms supported by SUSE Linux Enterprise Server as well as a short description of the installation procedure.

Quality service is also available. Experts can answer questions about installation or configuration, make reliable security updates available, and support development projects.

## 1 Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation and enter your comments there.

## 2 Additional Documentation

For additional documentation on this product, refer to <http://www.novell.com/documentation/sles10/index.html>:

### *Architecture-Specific Information*

Architecture-specific information needed to prepare a SUSE Linux Enterprise Server target for installation.

### *Installation and Administration*

In-depth installation and administration for SUSE Linux Enterprise Server.

For a documentation overview on the SUSE® Linux Enterprise Desktop product, refer to <http://www.novell.com/documentation/sled10/index.html>.

Documentation regarding the Open Enterprise Server 2 is found at <http://www.novell.com/documentation/oes2/index.html>.

# 3 Documentation Conventions

The following typographical conventions are used in this manual:

- `/etc/passwd`: filenames and directory names
  - *placeholder*: replace *placeholder* with the actual value
  - `PATH`: the environment variable `PATH`
  - `ls, --help`: commands, options, and parameters
  - `user`: users or groups
  - `Alt, Alt + F1`: a key to press or a key combination; keys are shown in uppercase as on a keyboard
  - *File, File > Save As*: menu items, buttons
  - This paragraph is only relevant for the specified architectures. The arrows mark the beginning and the end of the text block.
- This paragraph is only relevant for the specified architectures. The arrows mark the beginning and the end of the text block.
- *Dancing Penguins* (Chapter *Penguins*, ↑Another Manual): This is a reference to a chapter in another manual.

# Introduction to Xen Virtualization

# 1

Virtualization of operating systems is used in many different computing areas. It finds its applications in server consolidation, energy saving efforts, or the ability to run older software on new hardware, for example. This chapter introduces and explains the components and technologies you need to understand to set up and manage a Xen-based virtualization environment.

## 1.1 Basic Components

The basic components of a Xen-based virtualization environment are the *Xen hypervisor*, the *Domain0*, any number of other *VM Guests*, and the tools, commands, and configuration files that let you manage virtualization. Collectively, the physical computer running all these components is referred to as a *virtual machine host* because together these components form a platform for hosting virtual machines.

### The Xen Hypervisor

The Xen hypervisor, sometimes referred to generically as a virtual machine monitor, is an open-source software program that coordinates the low-level interaction between virtual machines and physical hardware.

### The Domain0

The virtual machine host environment, also referred to as *domain0* or controlling domain, is comprised of several components, such as:

- The SUSE Linux operating system, which gives the administrator a graphical and command line environment to manage the virtual machine host components and its virtual machines.

---

## NOTE

The term “Domain0” refers to a special domain that provides the management environment. This may be run either in graphical or in command line mode.

---

- The xend daemon (xend), which stores configuration information about each virtual machine and controls how virtual machines are created and managed.
- A modified version of QEMU, which is an open-source software program that emulates a full computer system, including a processor and various peripherals. It provides the ability to host operating systems in full virtualization mode.

### Xen-Based Virtual Machines

A Xen-based virtual machine, also referred to as a *domain*, consists of the following components:

- At least one virtual disk that contains a bootable operating system. The virtual disk can be based on a file, partition, volume, or other type of block device.
- Virtual machine configuration information, which can be modified by exporting a text-based configuration file from xend or through Virtual Machine Manager.

### Management Tools, Commands, and Configuration Files

There is a combination of GUI tools, commands, and configuration files to help you manage and customize your virtualization environment.

## 1.2 Understanding Virtualization Modes

Guest operating systems are hosted on virtual machines in either full virtualization mode or paravirtual mode. Each virtualization mode has advantages and disadvantages.



- Full virtualization mode lets virtual machines run unmodified operating systems, such as Windows\* Server 2003 and Windows XP, but requires the computer running as the virtual machine host to support hardware-assisted virtualization technology, such as AMD\* Virtualization or Intel\* Virtualization Technology.

Some guest operating systems hosted in full virtualization mode, can be configured to run SUSE Linux Enterprise Virtual Machine Drivers instead of operating-system-specific drivers. Running virtual machine drivers improves performance dramatically on guest operating systems, such as Windows XP and Windows Server 2003. For more information, see [Chapter 7, \*Virtual Machine Drivers\*](#) (page 69).

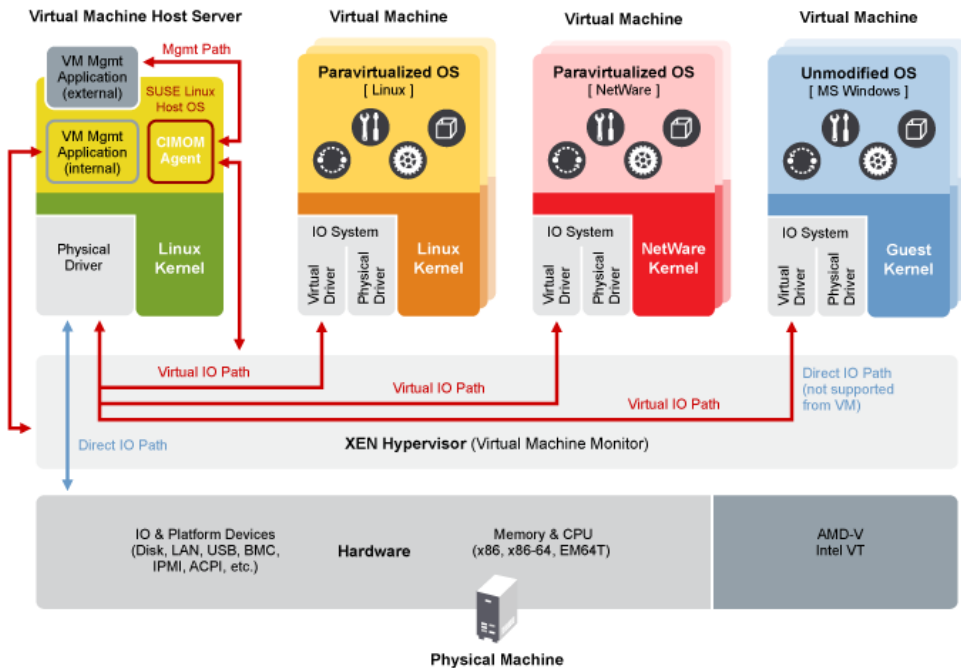
- Paravirtual mode does not require the host computer to support hardware-assisted virtualization technology, but does require the guest operating system to be modified for the virtualization environment. Typically, operating systems running in paravirtual mode enjoy better performance than those requiring full virtualization mode.

Operating systems currently modified to run in paravirtual mode are referred to as *paravirtualized operating systems* and include SUSE Linux Enterprise Server 10 and NetWare® 6.5 SP7.

## 1.3 Xen Virtualization Architecture

The following graphic depicts a virtual machine host with four virtual machines. The Xen hypervisor is shown as running directly on the physical hardware platform. Note, that the controlling domain is also just a virtual machine, although it has several additional management tasks compared to all other virtual machines.

**Figure 1.1** *Virtualization Architecture*



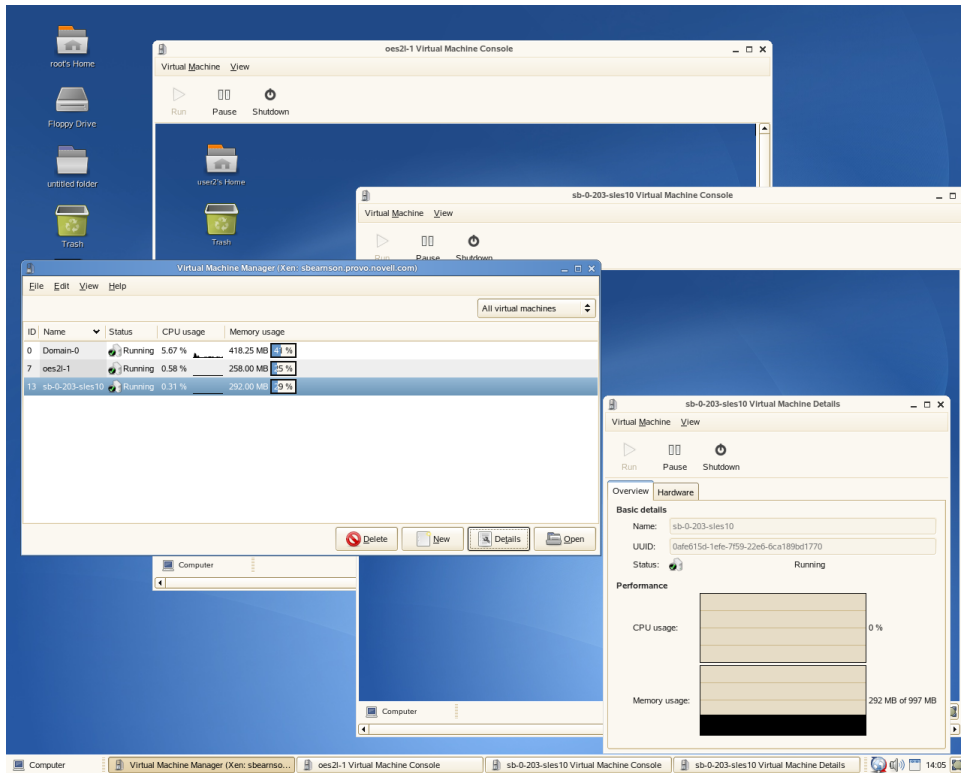
On the left, the virtual machine host's Domain0 is shown running the SUSE Linux operating system. The two virtual machines shown in the middle are running paravirtualized operating systems. The virtual machine on the right shows a fully virtual machine running an unmodified operating system, such as Windows Server 2003 or Windows XP.

## 1.4 The Virtual Machine Host

After you install the virtualization components and reboot the computer, the GRUB boot loader menu should display a Xen menu option. Selecting the Xen menu option loads the Xen hypervisor and starts the Domain0 running the SUSE Linux operating system.

Running on domain 0, the SUSE Linux operating system displays the installed text console or desktop environment, such as GNOME or KDE. When opened, each virtual machine is displayed in its own window inside the controlling domain 0.

**Figure 1.2** Desktop Showing Virtual Machine Manager and Virtual Machines





# Setting Up a Virtual Machine Host

# 2

This section documents how to set up and use SUSE Linux Enterprise 10 SP2 as a virtual machine host.

In most cases, the hardware requirements for the Domain0 are the same as those for the SUSE Linux Enterprise Server operating system, but additional CPU, disk, memory, and network resources should be added to accommodate the resource demands of all planned VM Guest systems.

---

## TIP

Remember that VM Guest systems, just like physical machines, perform better when they run on faster processors and have access to more system memory.

---

The following table lists the minimum hardware requirements for running a typical virtualized environment. Additional requirements have to be added for the number and type of the respective guest systems.

**Table 2.1** *Hardware Requirements*

System Component	Minimum Requirements
Computer	Computer with Pentium II or AMD K7 450 MHz processor
Memory	512 MB of RAM for the host

System Component	Minimum Requirements
Free Disk Space	7 GB of available disk space for the host.
Optical Drive	DVD-ROM Drive
Hard Drive	20 GB
Network Board	Ethernet 100 Mbps
IP Address	<ul style="list-style-type: none"> <li>• One IP address on a subnet for the host.</li> <li>• One IP address on a subnet for each VM Guest.</li> </ul>

Xen virtualization technology is available in SUSE Linux Enterprise Server products based on code path 10 and later. Code path 10 products include Open Enterprise Server 2 Linux, SUSE Linux Enterprise Server 10, SUSE Linux Enterprise Desktop 10, and openSUSE 10.x.

The virtual machine host requires a number of software packages and their dependencies to be installed. To install all necessary packages, run YaST *Software Management*, select the filter *Patterns* and choose *Xen Virtual Machine Host Server* for installation.

After the Xen software is installed, restart the computer.

Updates are available through your update channel. To be sure to have the latest updates installed, run YaST *Online Update* after the installation has finished.

## 2.1 Best Practices and Suggestions

When installing and configuring the SUSE Linux Enterprise operating system on the host, be aware of the following best practices and suggestions:

- If the host should always run as Xen host, run YaST *System > Boot Loader* and activate the Xen boot entry as default boot section.
- In YaST, click *System > Boot Loader*.

- Change the default boot to the *Xen* label, then click *Set as Default*.
- Click *Finish*.
- Disable powersave functionality on the host and all guest operating systems. On the host computer and other Linux computers, you can use the `chkconfig powersaved off` command.
- Close Virtual Machine Manager if you are not actively using it and restart it when needed. Closing Virtual Machine Manager does not affect the state of virtual machines.
- For best performance, only the applications and processes required for virtualization should be installed on the virtual machine host.
- When using both, iSCSI and OCFS2 to host Xen images, the latency required for OCFS2 default timeouts in SP2 may not be met. To reconfigure this timeout, run `/etc/init.d/o2cb configure` or edit `O2CB_HEARTBEAT_THRESHOLD` in the system configuration.

## 2.2 Setting Up a 32-Bit Domain 0 on a 64-Bit Hypervisor

A virtual machine host running on a 64-bit platform can access much more memory than hosts running on a 32-bit platform. However, there might be circumstances where you want to run the 32-bit version of the SUSE® Linux operating system (as domain 0) on the 64-bit Xen hypervisor.

- 1 On 64-bit hardware, start the 32-bit SUSE Linux operating system installation program.
- 2 Select *Installation Settings > Software* to include the *Xen Virtual Machine Host Server* software package selection.
- 3 Select *Software > Software Selection and System Tasks* and click *Details* to view all packages.

- 4 In the *Filter* drop-down list, select *Search*, then search for `kernel-xen`. Make sure that `kernel-xen-pae` is selected for installation and that `kernel-xen` is not selected.

- 5 Complete the SUSE Linux installation program.

The installation program copies files, reboots, and prompts you for additional setup information.

- 6 After entering the additional setup information, log in to the computer.

- 7 Access the SUSE Linux installation source media.

The 64-bit Xen hypervisor RPM is included in both 32-bit and 64-bit versions of the SUSE Linux installation media.

- 8 On the SUSE Linux installation source media, find the `x86_64` directory that contains `xen.rpm`.

- 9 Use the following command to install the RPM:

```
rpm -U xen.rpm --ignorearch --force
```

- 10 Edit the GRUB boot loader settings by running *YaST* > *System* > *Boot Loader* or editing the file at `/boot/grub/menu.lst`.

- 11 Change Hypervisor to `/boot/xen.gz`.

The `xen.gz` hypervisor is the 64-bit hypervisor.

- 12 Reboot the computer.

- 13 During the boot process, select the Xen option from the GRUB boot loader.

- 14 Verify that the computer is running the 32-bit Domain 0 by entering:

```
uname -m
```

It should report a 32-bit machine name, such as `i686`.

- 15 Verify that the computer is running the 64-bit hypervisor by entering:



```
xm info | grep xen_caps
```

It should report a 64-bit Xen hypervisor, such as xen-3.0-x86\_64.

## 2.3 Managing Domain 0 Memory

When the host is set up, a percentage of system memory is reserved for the hypervisor, and all remaining memory is automatically allocated to Domain0.

A better solution is to set a minimum and maximum amount of memory for domain 0, so the memory can be allocated appropriately to the hypervisor. An adequate maximum amount would be 20 percent of the total system memory up to 2 GB. An appropriate minimum amount would at least be 512 MB.

### 2.3.1 Setting a Maximum Amount of Memory

- 1 Determine the amount of memory to set for domain 0.
- 2 At Domain0, type `xm list` to view the currently allocated memory.
- 3 Run *YaST* > *Boot Loader*.
- 4 Select the Xen section.
- 5 In *Additional Xen Hypervisor Parameters*, add `dom0_mem= mem_amount` where *mem\_amount* is the maximum amount of memory to allocate to Domain0. Add K, M, or G, to specify the size, for example, `dom0_mem=768M`.
- 6 Restart the computer to apply the changes.

## 2.3.2 Setting a Minimum Amount of Memory

To set a minimum amount of memory for Domain0, edit the `dom0-min-mem` parameter in the `/etc/xen/xend-config.sxp` file and restart `xend`. For more information, see [Section 4.2, “Controlling the Host by Modifying xend Settings”](#) (page 27).

## 2.4 Starting the Virtual Machine Host

If virtualization software is correctly installed, the computer boots to display the GRUB boot loader with a *Xen* option on the menu. Select this option to start the virtual machine host.

---

### NOTE: Xen and Kdump

In Xen, the hypervisor manages the memory resource. If you need to reserve system memory for a recovery kernel in Domain0, this memory has to be reserved by the hypervisor. Thus, it is necessary to add the parameter `crashkernel=size@offset` to the `kernel` line instead of using the line with the other boot options. For more about Kdump, see also [http://www.novell.com/rc/docrepository/public/37/basedocument.2007-09-13.2960758124/4622069\\_f\\_en.pdf](http://www.novell.com/rc/docrepository/public/37/basedocument.2007-09-13.2960758124/4622069_f_en.pdf).

---

If the *Xen* option is not on the GRUB menu, review the steps for installation and verify that the GRUB boot loader has been updated. If the installation has been done without selecting the Xen pattern, run the *YaST Software Management*, select the filter *Patterns* and choose *Xen Virtual Machine Host Server* for installation.

After booting the hypervisor, the Domain0 virtual machine starts and displays its graphical desktop environment. If you did not install a graphical desktop, the command line environment appears.

Before starting to install virtual guests, make sure that the system time is correct. To do this, configure NTP (Network Time Protocol) on the controlling domain:

- 1 In YaST select *Network Services > NTP Configuration*.
- 2 Select the option to automatically start the NTP daemon during boot. Provide the IP address of an existing NTP time server, then click *Finish*.

---

**NOTE: Time Services on Virtual Guests**

Hardware clocks commonly are not very precise. All modern operating systems try to correct the system time compared to the hardware time by means of an additional time source. To get the correct time on all VM Guest systems, also activate the network time services on each respective guest or make sure that the guest uses the system time of the host. For more about `Independent Wallclocks` in SUSE Linux Enterprise Server see [Section 6.3.4, “Virtual Machine Clock Settings”](#) (page 65).

---

For more information about managing virtual machines, see [Chapter 4, \*Managing a Virtualization Environment\*](#) (page 25).



# Setting Up Virtual Machines

A virtual machine is comprised of data and operating system files that define the virtual environment. Virtual machines are hosted and controlled by the virtual machine host. This section provides generalized instructions for installing virtual machines.

Virtual machines have few if any requirements above those required to run the operating system. If the operating system has not been optimized for the virtual machine host environment, the unmodified OS can run only on hardware-assisted virtualization computer hardware, in full virtualization mode, and requires specific device drivers to be loaded.

You should be aware of any licensing issues related to running a single licensed copy of an operating system on multiple virtual machines. Consult the operating system license agreement for more information.

---

**NOTE: Virtual Machine Architectures**

The virtual machine host runs only on x86, AMD64 and Intel 64 hardware. It does not run on other system architectures such as Itanium, or POWER. A 64-bit virtual machine host can also run a 32-bit operating system, while a 32-bit host cannot run a 64-bit client. This holds true for both, the controlling domain and all further guests.

---

## 3.1 Creating a Virtual Machine

Before creating a virtual machine, you need the following:

- Install a host server as described in *Chapter 2, Setting Up a Virtual Machine Host* (page 7).
- If you want to use an automated installation file (AutoYaST, NetWare® Response File, or RedHat Kickstart), you should create and download it to a directory on the host machine server or make it available on the network.
- For NetWare and OES Linux virtual machines, you need a static IP address for each virtual machine you create.
- If you are installing Open Enterprise Server (OES) 2 Linux, you need a network installation source for OES 2 Linux software including the SUSE Linux Enterprise Server 10 SP2 media and the OES 2 Linux add-on CD. For procedures to create the installation sources, see "Setting Up the Server Holding the Installation Sources" in the *SUSE Linux Enterprise Server 10 Installation and Administration Guide* [[http://www.novell.com/documentation/sles10/sles\\_admin/data/sec\\_deployment\\_remoteinst\\_instserver.html](http://www.novell.com/documentation/sles10/sles_admin/data/sec_deployment_remoteinst_instserver.html)]

For further prerequisites, consult the manuals of the respective operating system to install.

The Create Virtual Machine Wizard helps you through the steps required to create a virtual machine and install its operating system. The information that follows is generalized for installing any operating system. In addition to these instructions, you should consult *Chapter 6, Xen Guest Systems* (page 59) for information specific to the operating system you are installing.

The actual configuration files for the Xen guests are stored at `/etc/xen/vm/`. The default location for image files is `/var/lib/xen/images`.

Launch the *Create Virtual Machine Wizard* by using one of the following methods:

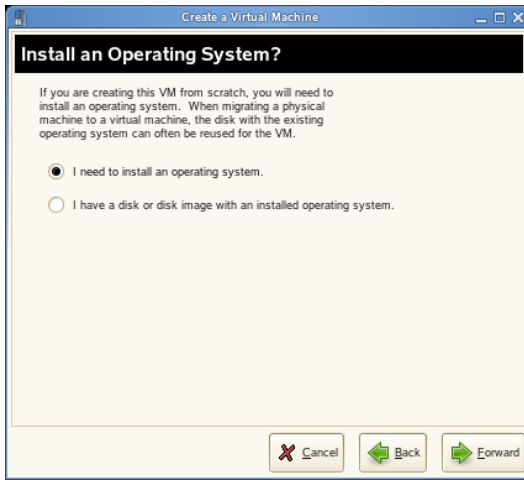
- From the virtualization host server desktop, click *YaST > Virtualization > Create Virtual Machine*
- From within Virtual Machine Manager, click *New*.
- At the command line, enter `vm-install`.

If the wizard does not appear or the `vm-install` command does not work, review the process of installing and starting the virtualization host server. The virtualization software might not be installed properly.

## 3.2 Installing an Operating System

You can choose to run an installation program or choose a disk or disk image that already has an installed and bootable operating system.

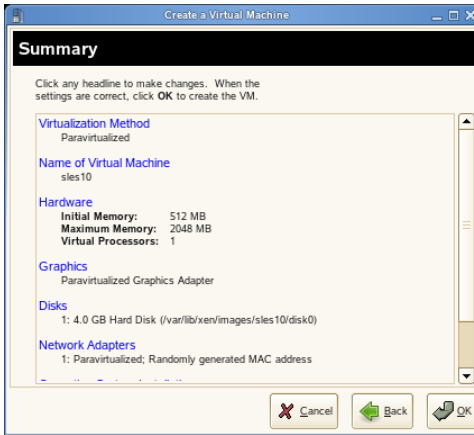
**Figure 3.1** *Installing an Operating System*



If you choose to run an installation program, you are presented with a list of operating systems. Select the one you want to install.

The Summary page shows you a summary of the virtual machine you are creating. You can click on any of the headings to edit the information. As you edit the information in the Summary, consult [Chapter 6, \*Xen Guest Systems\*](#) (page 59) for instructions specific to the operating system you are installing.

**Figure 3.2** *Summary*



Information about the following operating systems is included:

- OES NetWare Virtual Machines, see [Section 6.1, “NetWare Virtual Machines”](#) (page 59).
- OES Linux Virtual Machines, see [Section 6.2, “OES Linux Virtual Machines”](#) (page 59).
- SUSE Linux Virtual Machines, see [Section 6.3, “SUSE Linux Virtual Machines”](#) (page 59).
- Windows Virtual Machines, see [Section 6.4, “Windows Virtual Machines”](#) (page 68).

When running through the creation of a VM Guest, the following steps have to be accomplished:

- 1 Select if the VM Guest should run as full or paravirtualized guest.

If your computer supports hardware-assisted virtualization, you can create a virtual machine that runs in fully virtual mode. If you are installing an operating system that is modified for virtualization, you can create a virtual machine that runs in paravirtual mode. For more information about virtualization modes, see [Section 1.2, “Understanding Virtualization Modes”](#) (page 2).



- 2 Each virtual machine must have a unique name. The name entered on this page is used to create and name the virtual machine's configuration file. The configuration file contains parameters that define the virtual machine and is saved to the `/etc/xen/vm/` directory.
- 3 The Hardware page allows you to specify the amount of memory and number of virtual processors for your virtual machine.

#### Initial Memory

The amount of memory initially allocated to the virtual machine (specified in megabytes).

#### Maximum Memory

The largest amount of memory the virtual machine will ever need.

#### Virtual Processors

If desired, you can specify that the virtual machine has more virtual CPUs than the number of physical CPUs. You can specify up to 32 virtual CPUs; however, for best performance, the number of virtual processors should be less than or equal to the number of physical processors.

- 4 Select the graphics mode to use:

#### No Graphics Support

The virtual machine operates like a server without a monitor. You can access the operating system through operating system supported services, such as SSH or VNC.

#### Paravirtualized Graphics Adapter

Requires that an appropriate graphics driver is installed in the operating system.

- 5 A virtual machine must have at least one virtual disk. Virtual disks can be:
  - File backed, which means that the virtual disk is a single image file on a larger physical disk.
  - A sparse image file, which means that the virtual disk is a single image file, but the space is not preallocated.
  - Configured from a block device, such as an entire disk, partition, or volume.

For best performance, create each virtual disk from an entire disk or a partition. For the next best performance, create an image file but do not create it as a sparse image file. A virtual disk based on a sparse image file delivers the most disk-space flexibility but slows installation and disk access speeds.

By default, a single, file-backed virtual disk is created as a sparse image file in `/var/lib/xen/images/vm_name` where `vm_name` is the name of the virtual machine. You can change this configuration to meet your specific requirements.

- 6 If you want to install from DVD or CD-Rom, add the drive to the list of available hard disks. To learn about device names of the available optical drives, run `hwinfo --cdrom` and search for the line starting with `Device File:`. Add this device file to the available hard disks of the VM Guest.
- 7 By default, a single virtual network card is created for the virtual machine. It has a randomly generated MAC address that you can change to fit your desired configuration. The virtual network card will be attached to a default bridge configured in the host. You can also create additional virtual network cards in the Network Adapters page of `vm-install`.

---

**NOTE: Using Arbitrary Bridge Names**

If installing a fully virtualized guest and you are using a bridge name that is different than the default names, explicitly specify the bridge by selecting the bridge name from the *Source* menu on the Virtual Network Adapter page. Paravirtual guests by definition are aware they are running on a virtualization platform and therefore, do not need to have the bridge explicitly specified, thus leaving *Source* as *Default* will suffice.

---

- 8 The operating system can be installed from a CD/DVD device or an ISO image file. In addition, if you are installing a SUSE Linux operating system, you can install the operating system from a network installation source.

If you are installing a paravirtual machine's operating system from CD, you should remove the virtual CD reader from the virtual machine after completing the installation, because the virtual machine assumes that the original CD is still in the CD reader, even if it is ejected. If it is ejected, the virtual machine cannot access the CD (or any other newly inserted CD) and receives I/O errors.

---

## WARNING: Changing CD-Roms

On paravirtual guests, it is currently not possible to go on with the installation of a system after changing the CD-Rom. Instead, you should use a DVD or a network based installation. For a workaround to this problem, see [Section 5.1.1, “Virtual CD Readers on Paravirtual Machines”](#) (page 32).

---

If the installation program is capable of recognizing an installation profile, response file, or script, you can automate the installation settings by specifying the location of the profile, response file, or script you want to use. For example, SUSE Linux uses an AutoYaST profile, NetWare uses a NetWare Response File, and Red Hat uses a Kickstart file to move through the installation screens with no interaction.

You can also pass instructions to the kernel at install time by entering parameters for the *Additional Arguments* field. These arguments may either be kernel options, or options for `linuxrc`. More information about `linuxrc` can be found at [http://www.novell.com/documentation/sles10/sles\\_admin/data/sec\\_deployment\\_remoteinst\\_bootinst.html](http://www.novell.com/documentation/sles10/sles_admin/data/sec_deployment_remoteinst_bootinst.html)

If all the information on the *Summary* screen is correct, click *OK* to create the virtual machine. A TightVNC screen appears and at this point you begin the installation of your OS. From this point on, follow the regular installation instructions for installing your OS. More information for any instructions that might be different for installing your OS in a virtualized environment may be found in [Chapter 6, \*Xen Guest Systems\*](#) (page 59).

## 3.3 Including Add-On Products in the Installation

In order to include an Add-On product in the installation process of a VM Guest, it is necessary to provide the installation system with both, the standard installation images and the image for the Add-On product.

First, add the system disk, the SUSE Linux Enterprise Server 10 SP2 installation image and the physical CD-Rom or Add-On image as disks to the VM Guest. For example, you may have:

xvda

Main system disk.

xvdb

ISO image of the installation medium.

xvdc

ISO image of the Add-On product.

During the installation, add the Add-On product to the installation by entering the device path. Commonly, this path looks like `hd:///?device=/dev/xvd<letter>`. In the special example with “xvdc” as Add-On product, this would look like:


```
hd:///?device=/dev/xvdc
```

## 3.4 Using the Command Line to Create Virtual Machines

From the command line, you can enter `vm-install` to run a text version of the Create Virtual Machine Wizard. The text version of the wizard is helpful in environments without a graphical user interface. This command defaults to using a graphical user interface if available and if no options were given on the command line.

For information on scripting a virtual machine installation, see the man pages of `vm-install` and `vm-install-jobs`.

**Figure 3.3** *Command Line Interface*



```
Terminal
File Edit View Terminal Tabs Help
sbearnson:~ # vm-install
Gathering settings...

Please specify the type of operating system that will run within the virtual
machine. This defines many defaults, and helps decide how to start
paravirtualized operating systems.
 1: Microsoft Windows (other)
 2: Microsoft Windows 2000, 2003
 3: Microsoft Windows 9x, NT, XP
 4: Novell NetWare 6.5 SP6 and older
 5: Novell Open Enterprise Server 2 (Linux)
 6: Novell Open Enterprise Server 2 (NetWare)
 7: Other operating system
 8: RedHat (other)
 9: RedHat Enterprise Linux 3
10: RedHat Enterprise Linux 4
11: RedHat Enterprise Linux 5
12: SUSE (other)
13: SUSE Linux Enterprise Desktop 10
14: SUSE Linux Enterprise Server 10
15: SUSE Linux Enterprise Server 8
16: SUSE Linux Enterprise Server 9
17: Solaris 10
18: Solaris 9 and older
[14] >
```

## 3.5 Deleting Virtual Machines

When you use Virtual Machine Manager or the `xm` command to delete a virtual machine, it no longer appears as a virtual machine, but its initial startup file and virtual disks are not automatically deleted.

To delete all components of a virtual machine configured with a file-backed virtual disk, you must manually delete its virtual disk image file (`/var/lib/xen/images/`) and its initial startup file (`/etc/xen/vm`).

## 3.6 Troubleshooting

In some circumstances, problems may occur during the installation of the VM Guest. This section describes some known problems and their solutions.

During boot, the system hangs

The software I/O translation buffer allocates a large chunk of low memory early in the bootstrap process. If the requests for memory exceed the size of the buffer it usually results in a hung boot process. To check if this is the case, switch to console 10 and check the output there for a message similar to


```
kernel: PCI-DMA: Out of SW-IOMMU space for 32768 bytes at device  
000:01:02.0
```

In this case you need to increase the size of the `swiotlb`. Add “`swiotlb=128`” on the `Domain0` cmdline. Note that the number can be adjusted up or down to find the optimal size for the machine.

# Managing a Virtualization Environment

Graphical utilities, text-based commands, and modified configuration files are methods you can choose from to manage your virtualization environment. Virtual Machine Manager is a graphical utility available in YaST that can be launched from the virtual machine Domain0.

**Figure 4.1** Desktop Showing Virtual Machine Manager and Virtual Machines



From a command line interface on the virtual machine host, you can use the `vm-install` program and `xm` commands to create and manage virtual machines. You can also edit configuration files to change the settings of the virtual machine host or a virtual machine.

## 4.1 Virtual Machine Manager

From the desktop of the management virtual machine, the YaST Virtual Machine Manager provides a graphical user interface you can use to create and manage virtual machines.


---

### NOTE

Close Virtual Machine Manager if you are not actively using it and restart it when needed. Closing Virtual Machine Manager does not affect the state of virtual machines.

---

**Figure 4.2** *Virtual Machine Manager Main Console*



- Selecting a virtual machine and clicking *Open* displays the virtual machine window showing the virtual machine's current state.
- Clicking *Run* on the virtual machine window boots the virtual machine and displays the user interface or text console running on the virtual machine.
- Selecting a virtual machine and clicking *Details* lets you view performance and configure hardware details associated with the virtual machine.



- Clicking *New* in Virtual Machine Manager launches the *Create Virtual Machine Wizard*, which walks you through the steps required to set up a virtual machine. See also [Section 3.1, “Creating a Virtual Machine”](#) (page 15).

## 4.2 Controlling the Host by Modifying xend Settings

The xend is a key component of Xen virtualization. It performs management functions and stores settings that relate to the host environment and each virtual machine. You can customize xend to meet your specific configuration requirements.

- To specify xend operating parameters, edit the `/etc/xen/xend-config.sxp` file. The settings take effect the next time xend starts.

```
# -*- sh -*-

#
# Xend configuration file.
#

# This example configuration is appropriate for an installation that
# utilizes a bridged network configuration. Access to xend via http
# is disabled.

# Commented out entries show the default for that entry, unless otherwise
# specified.

#(logfile /var/log/xen/xend.log)
#(loglevel DEBUG)

# The Xen-API server configuration.
#
# This value configures the ports, interfaces, and access controls for
# the
# Xen-API server. Each entry in the list starts with either unix, or a
# port
```

- To start the xend daemon, enter `rcxend start`.
- To stop the xend daemon, enter `rcxend stop`.
- To restart the xend daemon, enter `rcxend restart`.

- To check the status of the xend daemon, enter `rcxend status`.

The parameters in the `xend-config.sxp` file can be customized to meet your requirements for virtualization. For a full list of all available options, read the manual page of `xend-config.sxp`.

## 4.3 Configuring a Virtual Machine by Modifying its xend Settings

The machine settings of each virtual guest are stored in an internal database managed by `xend`. You can change a virtual machine's settings by modifying the settings stored in `xend`. This process requires you to export a virtual machine's settings from the `xend` database to a text file, edit the settings in the file to meet your configuration requirements, import the file back into `xend`, and restart the virtual machine.

---

### NOTE

It is no longer recommended that you edit the initial startup files stored in `/etc/xen/vm`, because they are used only during the creation of a new virtual machine.

---

To modify a virtual machine's settings that is administrated with the virtual machine manager:

- 1 At `Domain0`, enter

```
xm list -l vm_name > filename
```

where `vm_name` is the name of the virtual machine you want to modify and `filename` is whatever you want to name the text file.

- 2 Use a text editor to make and save any desired changes.

```
(domain
  (domid 1)
  (on_crash destroy)
  (memory 384)
  (uuid 4fbc943-871c-9a51-3a48-3ad99d933841)
  (bootloader_args '- -entry=xvda2:/boot/vmlinuz-xen,/boot/initrd-xen')
```

```

(name SLES10SP1withOES2)      (maxmem 512)
(on_reboot restart)
(on_poweroff destroy)
(vcpus 1)
(bootloader /usr/lib/xen/boot/domUloader.py)
(shadow_memory 0)
(cpu_weight 256)
(cpu_cap 0)
(features )
(on_xend_start ignore)
(on_xend_stop ignore)
(start_time 1178219902.47)
(cpu_time 4574.26779201)
(online_vcpus 1)
(image
  (linux
    (kernel /var/lib/xen/tmp/kernel.Bg0o6h)
    (ramdisk /var/lib/xen/tmp/ramdisk.Bch8YM)
    (args 'TERM=xterm ')
  )
)
(status 2)
....

```

- 3 Delete the existing configuration from xenstore with the command `xm del vm_name`
- 4 Enter `xm new -F filename` to import the virtual machine's new settings into xend.
- 5 Enter `xm start vm_name` to start the virtual machine with its new settings.

You should repeat the entire process of exporting the file each time you want to make changes to a virtual machine's settings.

## 4.4 The `xm` Command

The `xm` command provides a command line interface for managing virtual machines. It can be used to create, pause, and shut down virtual machines. It can also be used to list the current domains, enable or pin virtual CPUs, and attach or detach block devices. For example, the `xm list` command displays the status of all virtual machines.

```
# xm list
```

Name	ID	Mem	VCPUs	State	Time(s)
Domain-0	0	457	2	r-----	2712.9
OES	7	512	1	-b-----	16.3
SLES10		512	1		12.9

The syntax of the `xm` command usually follows the format:

```
xm <subcommand> [domain-id] [OPTIONS]
```

where `subcommand` is the `xm` command to run, `domain-id` is the ID number assigned to a domain or the name of the virtual machine, and `OPTIONS` indicates subcommand-specific options.

Other useful `xm` commands include:

- `xm start` starts a virtual machine
- `xm reboot` reboots a virtual machine
- `xm destroy` immediately terminates a virtual machine
- `xm block-list` displays all virtual block devices attached to a virtual machine
- All `xm` operations require that the Xen control daemon, `xend`, be running. For this reason, you should make sure `xend` starts whenever the host boots.
- Most `xm` commands require root privileges to allow interaction with the Xen hypervisor. Entering the `xm` command when you are not logged in as root returns an error.
- Some `xm` commands return no information even though the action is completed. In some instances, for example, when shutting down a virtual machine, the action can take several seconds to complete. To verify that the action has completed, you might need to view its status another way, such as, using the `xm list` command.

For a complete list of `xm` command parameters, enter `xm help` at the command line or read the manual page of `xm`.

# Virtualization: Configuration Options and Settings

# 5

The documentation in this section, describes advanced management tasks and configuration options that might help technology innovators implement leading-edge virtualization solutions. It is provided as a courtesy and does not imply that all documented options and tasks are supported by Novell, Inc.

Virtualization technology is being rapidly developed. In an effort to keep you up to date on the most current information, this documentation will be updated more often than traditional Novell product documentation. To create a more dynamic environment, it will quickly include any feedback and new virtualization discoveries made by you and other users through the User Comments feature at the bottom of each page.

Although a more dynamic documentation environment can be beneficial, there will also be some drawbacks. The information might not have gone through a formal editing cycle before being published. The information might be less authoritative, might not completely cover obvious topics, and could have more typos than traditional product documentation. It might also include untested instructions for virtualization procedures you might want to try.

## 5.1 Virtual CD Readers

Virtual CD readers can be set up when a virtual machine is created or added to an existing virtual machine. A virtual CD reader can be based on a physical CD/DVD or based on an ISO image. Virtual CD readers work differently depending on whether they are paravirtual or fully virtual.

---

## NOTE

Only limited support for virtual machine removable media is available in the original version of SUSE Linux Enterprise Server. Much of the following information pertains to functionality available in maintenance updates.

If you are installing drivers from the SUSE Linux Enterprise Virtual Machine Driver Pack, proceed by completing one of the following:

- If you have updated to virt-manager 0.3.1-0.24 or later, complete the instructions in [Section 5.1.3, “Adding Virtual CD Readers”](#) (page 33)--> to set up the ISO image file as a virtual CD reader.

The most recent virt-manager package is available through the SUSE Linux Enterprise 10 SP1 update channel.

- If you are running an earlier version of virt-manager, you must set up the ISO image file as a virtual CD reader when you create the virtual machine, or add a virtual CD reader based on the ISO image file following the instructions in [Section 5.1.4, “Adding Virtual CD Readers \(Command Line Method\)”](#) (page 35)-->.

---

## 5.1.1 Virtual CD Readers on Paravirtual Machines

A paravirtual machine can have up to 16 block devices comprised of virtual CD readers and virtual disks. On paravirtual machines, virtual CD readers present the CD as a virtual disk with read-only access. Unless you remove the virtual CD reader, the virtual machine assumes that the originally-inserted CD is still in the CD reader, even if you have ejected it. If it has been ejected, the virtual machine will not be able to access the CD (or any other newly-inserted CD) and will receive I/O errors. Virtual CD readers cannot be used to write data to a CD. They are configured as read-only devices.

After you have finished accessing a CD on a paravirtual machine, it is recommended that you remove the virtual CD reader from the virtual machine.

If several CDs are needed for the installation of an operating system, make all of them available to the paravirtual machine and change the path to the data on request. If you

do not have enough CD readers, create image files from the CDs and make those available. To create an image file, the following command may be used:

```
dd if=/dev/cdrom of=<path to image>/cdimage.iso
```

On request of the installation system, press *Details* and manually point to the correct disk by changing `xvdb` to `xvdc`, `xvdc` to `xvdd` and so on.

## 5.1.2 Virtual CD Readers on Fully Virtual Machines

A fully virtual machine can have up to four block devices comprised of virtual CD readers and virtual disks. A virtual CD reader on a fully virtual machine interacts with the inserted CD in the way you expect a physical CD reader to interact. For example, in a Windows\* XP\* virtual machine, the inserted CD appears in the `Devices with Removable Storage` section of `My Computer`.

When a CD is inserted in the physical CD reader on the host computer, all virtual machines with virtual CD readers based on the physical CD reader, such as `/dev/cdrom/`, are able to read the inserted CD. Assuming the operating system has automount functionality, the CD should automatically appear in the file system. Virtual CD readers cannot be used to write data to a CD. They are configured as read-only devices.

## 5.1.3 Adding Virtual CD Readers

Virtual CD readers can be based on a CD inserted into the CD reader or on an ISO image file.

- 1 Make sure that the virtual machine is running and the operating system has finished booting.
- 2 Insert the desired CD into the physical CD reader or copy the desired ISO image to a location available to the virtual machine's file system.
- 3 Run Virtual Machine Manager.
- 4 Select the virtual machine, then click *Details*.

**5** Click *Hardware > Disk*.

**6** Click *CD-ROM*.

**7** Specify the path to the physical CD reader, such as `/dev/cdrom`. If you are adding a virtual CD reader based on an ISO, specify the path to the ISO image file.

**8** Click *OK* to apply the changes.

A new block device, such as `/dev/xvdb`, is added to the virtual machine.

**9** If the virtual machine is running Linux, complete the following:

**9a** Open a terminal in the virtual machine and enter `fdisk -l` to verify that the device was properly added. You could also enter `ls /sys/block` to see all disks available to the virtual machine.

The CD is recognized by the virtual machine as a virtual disk with a drive designation, for example,

```
/dev/xvdb
```

**9b** Enter the command to mount the CD or ISO image using its drive designation. For example, enter

```
mount -o ro /dev/xvdb /cd1
```

mounts the CD to a mount point named `/cd1`.

The CD or ISO image file should be available to the virtual machine at the specified mount point.

**10** If the virtual machine is running Windows, reboot the virtual machine.

Verify that the virtual CD reader appears in its `My Computer` section



## 5.1.4 Adding Virtual CD Readers (Command Line Method)

- 1 Make sure that the virtual machine is running and that the operating system has finished booting.
- 2 Insert the CD into the physical CD reader or copy the ISO image to the host.
- 3 In a terminal on the host, enter the appropriate command

- To set up an ISO image file, enter:

```
xm block-attach vm_id file:/mycd.iso hdc:cdrom r
```

where `vm_id` is the virtual machine ID, `mycd.iso` is the path to the ISO image file, and `hdc` is the drive designation presented to the virtual machine.

If the virtual machine is paravirtual, replace the device designation with a `vdb` `x` device designation, such as

```
xm block-attach vm_id file:/mycd.iso xvdb:cdrom r
```

- To set up a physical CD reader, enter:

```
xm block-attach vm_id phy:/dev/cdrom hdc:cdrom r
```

where `vm_id` is the virtual machine ID, `dev/cdrom` is the path to the physical CD reader, `hdc` is the physical drive designation to be presented to the virtual machine, and `cdrom` is the name you specify for the device.

If the virtual machine is paravirtual, replace the `hdc` drive designation with a `vdb` `x` drive designation, such as

```
xm block-attach vm_id phy:/dev/cdrom xvdb:cdrom r
```

- 4 Conditional: If running full virtualization mode, reboot the virtual machine.

The CD or ISO image file should appear in the virtual machine's operating system.

## 5.1.5 Removing Virtual CD Readers

- 1 Make sure that the virtual machine is running and the operating system has finished booting.
- 2 If the virtual CD reader is mounted, unmount it from within the virtual machine.

---

### TIP

You can enter `fdisk -l` in the virtual machine's terminal to view its block devices.

---

- 3 Run Virtual Machine Manager.
- 4 Select the virtual machine, then click *Details*.
- 5 Click *Hardware > Disk*.
- 6 Select the virtual CD-ROM device to remove.
- 7 Click *Remove* to remove the virtual CD-ROM device.
- 8 Press the hardware eject button to eject the CD.

## 5.2 Remote Access Methods

Some configurations, such as those that include rack-mounted servers, require a computer to run without a video monitor, keyboard, or mouse. This type of configuration is often referred to as `headless` and requires the use of remote administration technologies.

Typical configuration scenarios and technologies include:

### Graphical Desktop with X Window Server

If a graphical desktop, such as GNOME or KDE, is installed on the virtual machine host you can use a remote viewer, such as a VNC viewer. On a remote computer, log in and manage the host environment by using graphical tools, such as Virtual Machine Manager.

### Text and Graphical Applications

If neither a graphical desktop nor the X Window Server, but the X Windows libraries are installed on the virtual machine host, you can use the `ssh -X` command from the remote computer to log in and manage the virtualization host environment. You can then use Virtual Machine Manager and the `xm` command to manage virtual machines and the `vm-install` command to create them.

### Text Only

You can use the `ssh` command from a remote computer to log in to a virtual machine host and access its text-based console. You can then use the `xm` command to manage virtual machines and the `vm-install` command to create new virtual machines.

## 5.3 VNC Viewer

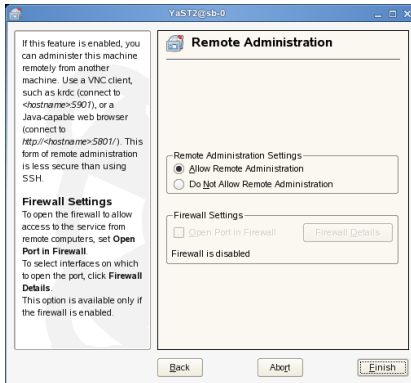
By default, Virtual Machine Manager uses the VNC viewer to show the display of a virtual machine. You can also use VNC viewer from Domain0 (known as local access or on-box access) or from a remote computer.

You can use the IP address of a VM Guest and a VNC viewer to view the display of this VM Guest. When a virtual machine is running, the VNC server on the host assigns the virtual machine a port number to be used for VNC viewer connections. The assigned port number is the lowest port number available when the virtual machine starts. The number is only available for the virtual machine while it is running. After shutting down, the port number might be assigned to other virtual machines.

For example, if ports 1 and 2 and 4 and 5 are assigned to the running virtual machines, the VNC viewer assigns the lowest available port number, 3. If port number 3 is still in use the next time the virtual machine starts, the VNC server assigns a different port number to the virtual machine.

To use the VNC viewer from a remote computer, remote administration must be enabled on the host. You can use the YaST Remote Administration module.

**Figure 5.1** *YaST Remote Administration*



In addition to this, change `vnc-listen` in `/etc/xen/xend-config.sxp` to open the access to the VM Guest. This is necessary, if you encounter messages like `vncviewer: ConnectToTcpAddr: connect: Connection refused` from a remote host. For more information about modifying `xend-config.sxp` see [Section 4.2, “Controlling the Host by Modifying xend Settings”](#) (page 27).

To access the virtual machine from the local console or a remote computer running a VNC viewer client, enter one of the following commands:


- `vncviewer host_ip::590#`
- `vncviewer host_ip:#`

where `host_ip` is the IP address of the virtual machine host and `#` is the VNC viewer port number assigned to the virtual machine.

```
# vncviewer 192.168.3.100::5901
```

If you are using the built-in VNC viewer of an Internet browser, such as Internet Explorer\* or Mozilla\* Firefox\*, use `580#` as the port number, where `#` is the VNC viewer port number assigned to the virtual machine.

**Figure 5.2** *Mozilla Firefox VNC Viewer*



When successfully connected, the VNC viewer shows the display of the running virtual machine.

## 5.3.1 Assigning VNC Viewer Port Numbers to Virtual Machines

Although the default behavior of VNC viewer is to assign the first available port number, you might want to assign a specific VNC viewer port number to a specific virtual machine.

To assign a specific port number on a fully virtualized guest, edit the virtual machine's setting (`/etc/xen/vm/name`) as follows: include the line (`vncdisplay #`) where

```
vnc=1
vncdisplay=#
vncunused=0
```

`vnc` must be set to one, `#` is the assigned port number and `vncunused` must exist and be set to 0 in order to make this work.

On a paravirtualized guest, edit the line `vfb=[ 'type=vnc,vncdisplay=# ' ]` to set the viewer port number to the desired value.

---

### TIP

Assign higher port numbers to avoid conflict with port numbers assigned by the VNC viewer, which uses the lowest available port number.

---

## 5.3.2 Using SDL instead of a VNC Viewer

If you access a virtual machine's display from the virtual machine host console (known as local or on-box access), you might want to use SDL instead of VNC viewer. VNC viewer is faster for viewing desktops over a network, but SDL is faster for viewing desktops from the same computer.

To set the default to use SDL instead of VNC, change the virtual machine's configuration information to the following. For instructions, see [Section 4.3, “Configuring a Virtual Machine by Modifying its xend Settings”](#) (page 28).

- If it is a fully virtual machine, use `vnc=0` and `sdl=1`.
- If it is a paravirtual virtual machine, use `vfb=["type=sdl"]`.


Remember that, unlike a VNC viewer window, closing an SDL window terminates the virtual machine.

## 5.4 The Boot Loader Program

The boot loader controls how the virtualization software boots and runs. You can modify the boot loader properties by using YaST or by directly editing the boot loader configuration file.

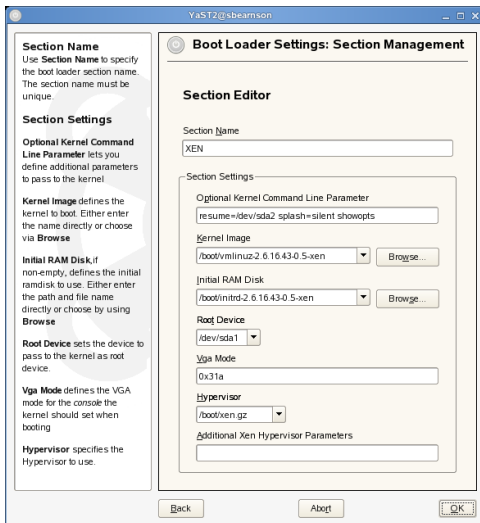
The YaST boot loader program is located at *YaST > System > Boot Loader*. The Boot Loader Settings screen lists the sections that appear as options on the boot menu. From this screen, you can change the boot loader so it auto-selects the virtual machine host option when booting.

**Figure 5.3** *Boot Loader Settings*



Select the *Xen* section, then click *Edit* to manage the way the boot loader and Xen function.

**Figure 5.4** *Boot Loader Settings: Section Management*



You can use the Boot Loader program to specify functionality, such as:

- Pass kernel command line parameters
- Specify the kernel image and initial RAM disk

- Select a specific hypervisor
- Pass additional parameters to the hypervisor (see `/usr/share/doc/packages/xen/pdf/user.pdf` section “Xen Boot Options” after installing the package `xen-doc-pdf`).

You can customize your virtualization environment by editing the `/boot/grub/menu.lst` file.

If the Xen option does not appear on the GRUB boot menu, you can compare your updated GRUB boot loader file with the examples below to confirm that it was updated correctly.

The first example shows a typical GRUB boot loader file updated to load the kernel that supports virtualization software. The second example shows a GRUB boot loader file that loads the PAE-enabled virtualization kernel.

**Example 5.1** *Xen Section in the menu.lst File (Typical)*

```
title XEN
root (hd0,5)
kernel /boot/xen.gz hyper_parameters
module /boot/vmlinuz-xen kernel_parameters
module /boot/initrd-xen
```

**Example 5.2** *Xen Section in the menu.lst File (PAE)*

```
title XEN
root (hd0,5)
kernel /boot/xen-pae.gz hyper_parameters
module /boot/vmlinuz-xenpae kernel_parameters
module /boot/initrd-xenpae
```

The `title` line defines sections in the boot loader file. Do not change this line, because YaST looks for the word *XEN* to verify that packages are installed.

The `root` line specifies which partition holds the boot partition and `/boot` directory. Replace `hd0, 5` with the correct partition. For example, if the drive designated as `hda1` holds the `/boot` directory, the entry would be `hd0, 0`.

The `kernel` line specifies the directory and filename of the hypervisor. Replace *hyper\_parameters* with the parameters to pass to the hypervisor. A common parameter is `dom0_mem=<amount_of_memory>`, which specifies how much memory



to allocate to Domain0. The amount of memory is specified in KB, or you can specify the units with a K, M, or G suffix, for example 128M. If the amount is not specified, the Domain0 takes the maximum possible memory for its operations.

For more information about hypervisor parameters, see `/usr/share/doc/packages/xen/pdf/user.pdf` section “Xen Boot Options” after installing the package `xen-doc-pdf`.

The first `module` line specifies the directory and filename of the Linux kernel to load. Replace `kernel_parameters` with the parameters to pass to the kernel. These parameters are the same parameters as those that can be passed to a standard Linux kernel on physical computer hardware.

The second `module` line specifies the directory and filename of the RAM disk used to boot the virtual machine host.

To set the GRUB boot loader to automatically boot the Xen virtualization software, change the `default` entry from 0, which means the first `title` entry, to the number that corresponds to the `title XEN` entry. In the example file, Xen is the second `title` line, so to specify it, you would change the value of `default` from 0 to 1.

## 5.5 Changing the Host’s Desktop from Text Mode to GUI Mode

If selecting the Xen option from the boot loader starts SUSE Linux desktop environment (domain 0) in text mode, the graphics card is probably not configured correctly. To properly configure the graphics card so you can switch to graphical mode, complete the following steps:

- 1 When the computer boots and loads the boot loader menu, select the option that includes the word *Xen*.

The computer boots and displays a command prompt.

- 2 At the command prompt, enter `telinit 3` to make sure you are in text mode.
- 3 Enter `SaX2` to run the SaX2 graphical configuration module.

4 Configure the GUI as desired.

5 Enter `telinit 5` to switch to GUI mode.

## 5.6 Paravirtual Mode and Journaling File Systems

It is recommended that operating systems running in paravirtual mode set up their kernel on a separate partition that uses a non-journaling file system, such as ext2.

Before a paravirtualized operating system can boot, the management domain must construct a virtual machine and place the paravirtualized kernel in it. Then, the paravirtualized operating system boots. To retrieve the kernel during the bootstrapping process, the virtual machine's boot disk is mounted in read-only mode, the kernel is copied to the virtual machine's memory, and then the boot disk is unmounted.

When a virtual machine's operating system crashes, its disks are not shut down in an orderly manner. This should not pose a problem to a virtual machine running in full virtualization mode because the pending disk entries are checked and corrected the next time the operating system starts. If the disk is using a journaling file system, the journal is replayed to update and coordinate any pending disk entries.

This type of system crash poses a potential problem for paravirtualized operating systems. If a paravirtualized operating system using a journaled file system crashes, any pending disk entries cannot be updated and coordinated because the file system is initially mounted in read-only mode.

Therefore, it is recommended that you set virtual machine boot files, such as the kernel and ramdisk, on a separate partition that is formatted with a non-journaling file system, such as ext2.

## 5.7 Virtual Machine Initial Startup Files

During the process of creating a new virtual machine, initial startup settings are written to a file created at `/etc/xen/vm/`. During the creation process, the virtual machine starts according to settings in this file, but the settings are then transferred and stored in xend for ongoing operations.

---

### IMPORTANT

Modifying the initial startup file to create or make changes to a virtual machine is not recommended. The preferred method for changing a virtual machine's settings is to use Virtual Machine Manager as described in [Section 4.3, “Configuring a Virtual Machine by Modifying its xend Settings”](#) (page 28) or to follow the instructions in [Section 5.4, “The Boot Loader Program”](#) (page 40).

---

When a virtual machine's settings are stored in xend, it is referred to as a `xen-managed domain` or `xen-managed virtual machine`. Whenever the `xen-managed` virtual machine starts, it takes its settings from information stored in the xend database—not from settings in the initial startup file.

Although it is not recommended, you might need to start an existing virtual machine based on settings in the initial startup file. If you do this, any xend settings stored for the virtual machine are overwritten by the startup file settings. Initial startup files are save to `/etc/xen/vm/vm_name`. Values must be enclosed in single quotes, such as `localtime = '0'`.

**Table 5.1** *Initial Startup File Entries and Descriptions*

Entry	Description
<code>disk =</code>	Virtual disks for the virtual machine.  For example:  <pre>disk = [ 'file:/var/lib/xen/images/VM1_SLES10/hda,xvda,w'</pre>

Entry	Description
	<p>This entry specifies a virtual disk based on a file ( <code>file:</code> ) named <code>hda</code> located at <code>/var/lib/xen/images/VM1_SLES10/</code>. It presents itself as the first drive (<code>xvda</code>) and has read/write access (<code>w</code>).</p> <p>Disks can also be based on a block device.</p>
<code>memory =</code>	Virtual memory in Mb.
<code>vcpus =</code>	Number of virtual CPUs.
<code>builder =</code>	Specifies paravirtual mode (Linux) or full virtualization mode ( <code>hvm</code> ).
<code>name =</code>	Name of the virtual machine.
<code>vif =</code>	Randomly-assigned MAC addresses and bridges assigned to use the virtual machine's network addresses.
<code>localtime =</code>	Specifies a localtime (0) or UTC (1) time setting.
<code>on_poweroff =</code>	Specifies the action that the virtual machine performs when the operating system is powered off.
<code>on_reboot =</code>	Specifies the action that the virtual machine performs when the operating system reboots.
<code>on_crash =</code>	Specifies the action that the virtual machine performs when the operating system crashes.
<code>extra =</code>	Parameters passed to the kernel.
<code>bootloader =</code>	Location and filename of the domU boot loader.
<code>bootentry =</code>	Location of the kernel and initial ramdisk.
<code>ostype =</code>	Type of operating system.

Entry	Description
<code>uuid =</code>	Identification number for a virtual drive.

## 5.8 Sparse Image Files and Disk Space

If the host's physical disk reaches a state where it has no available space, a virtual machine using a virtual disk based on a sparse image file is unable to write to its disk. Consequently, it reports I/O errors.

The Reiser file system, perceiving a corrupt disk environment, automatically sets the file system to read-only. If this situation happens, you should free up available space on the physical disk, remount the virtual machine's file system, and set the file system back to read-write.

To check the actual disk requirements of a sparse image file, use the command `du -h <image file>`.

## 5.9 Virtual Keyboards

When a virtual machine is started, the host creates a virtual keyboard that matches the `keymap` entry according to the virtual machine's settings. If there is no `keymap` entry in the virtual machine's settings, the host uses the `keymap` entry specified in host's `xend` file ( `xend-config.sxp`). If there is no `keymap` entry in either the host's `xend` file or the virtual machine's settings the virtual machine's keyboard defaults to English (US).

Unless you manually specify it, a `keymap` entry is not specified in the host's `xend` file or for any virtual machine. Therefore, by default, all virtual machine settings use the English (US) virtual keyboard. It is recommended that you specify a `keymap` setting for `xend` and for each virtual machine, especially if you want to migrate virtual machines to different hosts

To view a virtual machine's current `keymap` entry, enter the following command on the Domain0:

```
xm list -l vm_name | grep keymap
```

You can specify a keymap entry to be used for all virtual machines and keymap entries for specific machines.

- To specify a global keymap entry for virtual machines on the host, edit the host's `xend-config.sxp` file.
- To specify a keymap entry for a specific virtual machine, edit the virtual machine's settings by following instructions in [Section 4.3, “Configuring a Virtual Machine by Modifying its xend Settings”](#) (page 28).

In the `device > vfb` section, add the desired `keymap` entry to the file. For example, you can specify a German keyboard. Make sure the virtual machine's operating system is set to use the specified keyboard. After you specify the host's `keymap` setting, all virtual machines created by using the Create Virtual Machine Wizard on the host add the host's `keymap` entry to their virtual machine settings.

Virtual machines created before a host's `keymap` entry is specified are not automatically updated. These virtual machines start with the keyboard specified by the host, but the `keymap` entry is not a permanent part of the virtual machine's settings. For the entry to be permanent, it must be explicitly stated in the virtual machine's settings.

**Table 5.2** *Language and Keymap Settings*

Language	Keymap Setting
Danish	da
German	de
Swiss-German	de-ch
English (UK)	en-gb
English (US)	en-us
Spanish	es
Finnish	fi

<b>Language</b>	<b>Keymap Setting</b>
French	fr
French - Belgium	fr-be
French - Canada	fr-ca
French - Switzerland	fr-ch
Hungarian	hu
Icelandic	is
Italian	it
Japanese	ja
Dutch	nl
Dutch - Belgium	nl-be
Norwiegen	no
Polish	pl
Portuguese	pt
Portuguese - Brazil	pt-br
Russian	ru
Swedish	sv

## 5.10 Mapping Physical Storage to Virtual Disks

To specify a mapping between physical storage and the virtual disk, you might need to edit the virtual machine’s disk information. Follow the instructions in [Section 4.3](#), “Configuring a Virtual Machine by Modifying its xend Settings” (page 28), to change the respective device entry to the desired setting.

**Example 5.3** *Example: Virtual Machine Output from Xend*

```
(vbd
  (dev xvda:disk)
  (uname file:/var/lib/xen/images/oes21/disk0)
  (mode w)
  (type disk)
  (backend 0)
)
```

**Table 5.3** *uname Settings*

Protocol	Description	Example
phy:	Block devices, such as a physical disk, in domain 0	phy:/dev/xvda
file:	Raw disk images accessed by using loopback	file:/path/file
tap:aio:	Raw disk images accessed by using blktap. Similar to loopback but with better performance	tap:aio:/path/file
tap:vmdk:	VMware disk images accessed by using blktap	tap:vmdk:/path/file
nbd:	Raw disk images accessed by using NBD	nbd: ip_port
tap:qcow:	QEMU disk images accessed by using blktap	tap:qcow/path/file



Protocol	Description	Example
iscsi:	iSCSI targets using connections initiated from domain 0	<code>iscsi: IQN, LUN</code>
npiv:	Fibre Channel connections initiated from domain 0	<code>npiv: NPIV, LUN</code>

## 5.11 Virtual Disks

Virtual disks can be based on the following types of physical devices and files. Each type includes an example statement.

- A physical disk device, such as a DVD, that is accessible as a device to the host.

```
phy:/dev/cdrom
```

- A file that contains a disk image accessible from the file system of the host. Disk images can be encoded as raw, QEMU, or VMware.

```
file:/mnt/disks/sles10sp1.iso
```

`tap:aio:/mnt/disks/sles10sp1.iso` specifies a raw disk that might be taken from a different virtualization platform.

```
tap:qcow:/mnt/disks/sles10sp1.iso.qcow
tap:vmdk:/mnt/disks/sles10sp1.iso.vmdk
```

- A remote storage device specified using the Internet SCSI (iSCSI) protocol.

```
iscsi:iqn.2001-04.com.acme@0ac47ee2-216e-452a-a341-a12624cd0225
```

- A remote storage device specified using a Fibre Channel (NPIV) protocol.

```
npiv:210400e08b80c40f
```

## 5.12 Migrating Virtual Machines

A running virtual machine can be migrated from its source virtual machine host to another virtual machine host. This functionality is referred to as *live migration*. For live migration to work, the virtual machine being migrated must have access to its storage in exactly the same location on both source and destination host platforms.

Live migration only works when every entity involved is the same architecture. For example, a 64-bit paravirtualized guest running on a 64-bit hypervisor can be migrated a host running a 64-bit hypervisor. If any of the pieces do not match exactly, migration will fail.

Another requirement is, that the involved filesystems are available on both machines. The options to accomplish this task include Network Block Devices (NBD), iSCSI, NFS, drbd and fiber channel devices. Furthermore, the routing of the network connection to the virtual network device must be correct.

The following xend options, which are located in the `/etc/xen/xend-config.sxp` file, need to be set on both hosts in order for live migration to work.

```
(xend-relocation-server yes)
(xend-relocation-port 8002)
(xend-relocation-address "")
(xend-relocation-hosts-allow "")
```

For information on modifying xend settings, see [Section 4.2, “Controlling the Host by Modifying xend Settings”](#) (page 27).

## 5.13 Passing Key Combinations to Virtual Machines

When in a virtual machine window, some key combinations, such as Ctrl + Alt + F1, are recognized by the virtual machine host and are not passed to the virtual machine. To bypass the virtual machine host, Virtual Machine Manager provides sticky key functionality. Pressing Ctrl, Alt, or Shift three times makes the key sticky, then you can press the remaining keys to pass the combination to the virtual machine.

For example, to pass Ctrl + Alt + F2 to a Linux virtual machine, press Ctrl three times, then press Alt + F". You could also press Alt three times, then press Ctrl + F2.

The sticky key functionality is available in Virtual Machine Manager during and after installing a virtual machine.

## 5.14 xm block-list Device Numbers

Virtual devices are named internally in Xen using the Linux-style major and minor numbers. This is true even for non-Linux virtual machines. The number returned from `xm block-list` is a decimal representation of the combined major and minor numbers.

For example, if a virtual machine's virtual disk is designated as `hda`, the number returned from `xm block-list` is 768. `hda` has major number 3 and minor number 0. The major number is stored as a high-order byte; the minor is the lower byte. A decimal representation is  $(3*256)+0 = 768$ . Another example, `sda3` has major number 8 and minor number 3, so its decimal representation is  $(8*256)+3 = 2051$ .

Of course, it is possible to manually work backwards from a number to discover the human-readable device name. For your reference, some common mappings listed in the following table.

<code>/dev/hda</code>	768
<code>/dev/hdb</code>	832
<code>/dev/hdc</code>	5632
<code>/dev/hdd</code>	5696
<code>/dev/sda</code>	2048
<code>/dev/sdb</code>	2064
<code>/dev/sdc</code>	2080
<code>/dev/sdd</code>	2096
<code>/dev/xvda</code>	51712
<code>/dev/xvdb</code>	51728
<code>/dev/xvdc</code>	51744
<code>/dev/xvdd</code>	51760

## 5.15 File-Backed Virtual Disks and Loopback Devices

When a virtual machine is running, each of its file-backed virtual disks consumes a loopback device on the host. By default, the host allows up to eight loopback devices to be consumed.

To simultaneously run more file-backed virtual disks on a host, you can increase the number of available loopback devices by adding the following option to the host's `/etc/modprobe.conf.local` file.

```
options loop max_loop=x
```

where `x` is the maximum number of loopback devices to create.

Changes take effect after the module is reloaded.

---

### TIP

Enter `rmmod loop` and `modprobe loop` to unload and reload the module. In case `rmmod` does not work, unmount all existing loop devices or reboot the computer.

---

## 5.16 Saving Virtual Machines

The save operation preserves the exact state of the virtual machine's memory. The operation is somewhat similar to *hibernating* a computer. The virtual machine is off, but it can be quickly restored to its previously-saved running condition. The operation does not make a copy of any portion of the virtual machine's virtual disk.

When saved, the virtual machine is paused, its current memory state saved to a location you specify, and then the virtual machine is stopped. The amount of time to save the virtual machine depends on the amount of memory allocated. When saved, a virtual machine's memory is returned to the pool of memory available on the host.

The restore operation is used to return a saved virtual machine to its original running state.

---

## IMPORTANT

After using the save operation, do not boot, start, or run a virtual machine that you intend to restore. If the virtual machine is at any time restarted before it is restored, the saved memory-state file becomes invalid and should not be used to restore.

---

### **Procedure 5.1** *Save a Virtual Machine's Current State (Virtual Machine Manager)*

- 1 Run Virtual Machine Manager.
- 2 Make sure the virtual machine to be saved is running.
- 3 Select the virtual machine.
- 4 Click *Open* to view the virtual machine console or *Details* to view virtual machine information.
- 5 Select *Virtual Machine > Save* from the menu.
- 6 Name and save the file.

### **Procedure 5.2** *Save a Virtual Machine's Current State (xm Command)*

- 1 Make sure the virtual machine to be saved is running.
- 2 In the host environment, enter `xm save ID state-file` where *ID* is the virtual machine ID you want to save and *state-file* is the name you specify for the memory-state file.

## 5.17 Restoring Virtual Machines

The restore operation loads a virtual machine's previously saved memory-state file and starts the virtual machine. The virtual machine does not boot the operating system but resumes at the point that it was previously saved. The operation is somewhat similar to coming out of hibernation.

The restore operation deletes the previously-saved memory-state file and assigns the virtual machine a new ID. The virtual machine name and UUID remain the same as previously saved.

---

## IMPORTANT

After using the save operation, do not boot, start, or run the virtual machine you intend to restore. If the virtual machine is at any time restarted before it is restored, the saved memory-state file becomes invalid and should not be used to restore.

---

### **Procedure 5.3** *Restore a Virtual Machine's Current State (Virtual Machine Manager)*

- 1 Make sure the virtual machine to be restored has not been started since you ran the save operation.
- 2 Run Virtual Machine Manager.
- 3 In Virtual Machine Manager, click *File > Restore Saved Machine* from the drop-down menu.
- 4 Specify the previously-saved file.
- 5 Click Open.

The virtual machine and the guest operating system are restored to the previously-saved state. The memory-state file is automatically deleted.

### **Procedure 5.4** *Restore a Virtual Machine's Current State (xm Command)*

- 1 Make sure the virtual machine to be restored has not been started since you ran the save operation.
- 2 In the host environment, enter `xm restore state-file` where `state-file` is the previously-saved memory-state file.

## 5.18 Virtual Machine States

A virtual machine's state can be displayed in Virtual Machine Manager or by viewing the results of the `xm list` command, which abbreviates the state using a single character.

- `r` - running - The virtual machine is currently running and consuming allocated resources.
- `b` - blocked - The virtual machine's processor is not running and not able to run. It could be waiting for I/O or has stopped working.
- `p` - paused - The virtual machine is paused. It does not interact with the hypervisor but still maintains its allocated resources, such as memory.
- `s` - shutdown - The guest operating system is in the process of being shutdown, rebooted, or suspended, and the virtual machine is being stopped.
- `c` - crashed - The virtual machine has crashed and is not running.
- `d` - dying - The virtual machine is in the process of shutting down or crashing.





# Xen Guest Systems

## 6.1 NetWare Virtual Machines

In Novell® Open Enterprise Server (OES) 2, the NetWare® operating system has been modified to run in paravirtual mode on a Xen virtual machine.

The documentation of OES Netware virtual machines is discussed in a separate document. For more information on this topic see [http://www.novell.com/documentation/oes2/inst\\_oes\\_nw/data/b8m3eeb.html](http://www.novell.com/documentation/oes2/inst_oes_nw/data/b8m3eeb.html).

## 6.2 OES Linux Virtual Machines

The documentation of OES Linux virtual machines is discussed in a separate document. For more information on this topic see [http://www.novell.com/documentation/oes2/inst\\_oes\\_lx/data/b9kmg9x.html](http://www.novell.com/documentation/oes2/inst_oes_lx/data/b9kmg9x.html).

## 6.3 SUSE Linux Virtual Machines

On host computers running SUSE Linux code 10 with Service Pack 1, the following SUSE Linux operating systems are supported as guests.

Operating System	Fully virtual	Paravirtual
SUSE Linux Enterprise Server 10 SP1	Yes	Yes
SUSE Linux Enterprise Server 10		Yes
Open Enterprise Server 2 Linux*	Yes	Yes
Open Enterprise Server 1 Linux*		Yes
SUSE Linux Enterprise Desktop 10 SP1	Yes	Yes
SUSE Linux Enterprise Server 9 SP4	Yes	Yes
SUSE Linux Enterprise Server 9 SP3	Yes	

## 6.3.1 SLE Virtual Machines on Host Architectures

SUSE Linux operating systems supported as guest operating systems, such as SUSE Linux 10 SP1 or later and SUSE Linux 9 SP3 or later, can be installed as guests on a host that is running any of the following combinations of domain 0 and hypervisor architecture.

**Table 6.1** *Virtual Machines and Host Architectures*

Operating System and Virtualization Mode	Hypervisor and Domain 0
SLE (32-bit) - Fully virtual	<ul style="list-style-type: none"> <li>• 32-bit hypervisor / 32-bit domain 0</li> <li>• 32-bit PAE hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 64-bit domain 0</li> </ul>

<b>Operating System and Virtualization Mode</b>	<b>Hypervisor and Domain 0</b>
SLE (32-bit) - Paravirtual	<ul style="list-style-type: none"> <li>• 32-bit hypervisor / 32-bit domain 0</li> </ul>
SLE (32-bit PAE) - Fully virtual	<ul style="list-style-type: none"> <li>• 32-bit PAE hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 64-bit domain 0</li> </ul>
SLE (32-bit PAE) - Paravirtual	<ul style="list-style-type: none"> <li>• 32-bit PAE hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 64-bit domain 0</li> </ul>
SLE (64-bit) - Fully virtual	<ul style="list-style-type: none"> <li>• 64-bit hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 64-bit domain 0</li> </ul>
SLE (64-bit) - Paravirtual	<ul style="list-style-type: none"> <li>• 64-bit hypervisor / 32-bit PAE domain 0</li> <li>• 64-bit hypervisor / 64-bit domain 0</li> </ul>

## 6.3.2 Updating SLE 10 Virtual Machines to SLE 10 SP2

Typically, the guest operating system on a virtual machine can be updated using the same procedures as if it were running on a physical machine; however, SLE 10 virtual machines created on a SLE 10 host need to be updated to SLE 10 SP1 by re-creating the virtual machine from the settings stored in the virtual machine's configuration file.

SLE 10 virtual machines created on a SLE 10 SP1 host do not need to follow these same procedures to be updated to SP1. Additionally, SLE 10 did not support SLE 10 fully virtual machines.

---

## NOTE

In SLE 10, changing the configuration of a virtual machine by editing its configuration file (`etc/xen/vm/vm_name`) was acceptable. In SLE 10 SP1 and later, editing the configuration file is no longer recommended. Instead, virtual machine configuration changes should be made in Virtual Machine Manager and by editing virtual machine settings stored in the xend database.

---

To update a SLE 10 virtual machine to SLE 10 SP1, complete the following procedure.

- 1 Make sure the host computer is running the most recent SLE updates. The host computer must be running software that is more recent than the software planned for the virtual machine update.
- 2 Prepare the virtual machine's operating system environment for the update by completing any prerequisite tasks. Making a copy of the entire virtual disk is recommended.
- 3 Shut down the virtual machine you want to update.
- 4 View or print the virtual machine's configuration file located at `etc/xen/vm/vm_name`.
- 5 On Domain0, use Virtual Machine Manager to create a new virtual machine.
- 6 Choose the selection to install an operating system.
- 7 Specify the settings for the new virtual machine as the same settings specified in the VM's configuration file.

Make sure that the following settings in Virtual Machine Manager match those in the VM's configuration file.

- Network MAC address
- Memory size

- Virtual disk size

- 8 Make sure to select or specify the virtual machine's already-created disk or disk image, for example, `hda`, as the virtual disk.

For example, if the SLE 10 virtual machine was using `/var/lib/xen/images/sles10/disk0` as its disk, the updated virtual machine should specify the same disk.

- 9 Specify the virtual machine operating system installation source as the SLE 10 SP1 installation media.
- 10 In the Operating System Installation section, add the following argument in Additional Arguments.

```
root=/dev/xvda2
```

where `/dev/xvda2` is the path and drive designation for the root partition that contains the virtual machine's boot files.

- 11 Click OK to create the virtual machine and start the SLE 10 SP1 installation program.

A new window displaying the installation program opens on the Domain0.

If the new windows does not display, select the virtual machine in Virtual Machine Manager, and click Open.

- 12 During the installation program, select Update from the Installation Mode screen.
- 13 If you receive a message stating that an error occurred while installing GRUB, click *OK*.
- 14 If you receive a message stating that an error occurred during boot loader installation, click *No*.
- 15 Continue the installation (update) by following the on-screen instructions.

The installation program continues by booting the virtual machine and displaying the rest of the installation in text mode.

After completing the installation program, the virtual machine should now be running SLE 10 SP1 and registered with xend.

**16** Log in to the SLE 10 SP1 virtual machine.

**17** (Conditional) If you want the SLE 10 SP1 virtual machine to run in GUI mode, complete the following from its command line:

**17a** Enter `init 3`.

**17b** Enter `sax2` to configure the GUI environment.

**17c** Enter `init 5` to restart the GUI.

**18** Open the virtual machine's `/etc/inittab` file with a text editor.

**19** Find the section titled

```
# getty-programs for the normal runlevels
```

**20** After the line that begins with `1:2345:respawn:sbin/mingetty`, add the following lines:

```
2:2345:respawn:/sbin/mingetty tty2
3:2345:respawn:/sbin/mingetty tty3
4:2345:respawn:/sbin/mingetty tty4
5:2345:respawn:/sbin/mingetty tty5
6:2345:respawn:/sbin/mingetty tty6
```

**21** Save the file and restart the virtual machine's operating system.

If the mouse pointer stops working within the virtual machine, you might need to restart the GUI on the virtual machine host by entering `init 3` then `init 5` on the host's command line.

After completing these instructions, the updated SLE 10 SP1 virtual machine should perform correctly. Configuration changes should be made using Virtual Machine Manager or by modifying xend settings.

## 6.3.3 Using the Add-On Products Program

The Add-On Products program is available during the SLE operating system installation and after installation at *YaST > Software > Add-On Products*. It lets you install additional products that may reside on a separate CD, ISO image file, or installation source.

Because paravirtual machines present removable media, such as a CD inserted in the CD reader, as a non-removable disk device, the Add-On Product program does not recognize inserted CD as valid add-on product media.

To use the Add-On Products program on a paravirtual machine, you must set up the add-on product media as a network installation source or copy the ISO image file to the virtual machine's filesystem.

On fully virtual machines, you can use the Add-On Products program to specify add-on product media as a network installation source, an ISO image file, or as a CD inserted in the host's CD reader.

## 6.3.4 Virtual Machine Clock Settings

When booting, virtual machines get their initial clock time from their host. After getting their initial clock time, fully virtual machines manage their time independently from the host. Paravirtual machines manage clock time according to their independent wall-clock setting. If the independent wallclock is enabled, the virtual machine manages its time independently and does not synchronize with the host. If the independent wallclock is disabled, the virtual machine periodically synchronizes its time with the host clock.

---

### NOTE

OES 2 NetWare virtual machines manage clock time independently after boot up. They do not synchronize with the host clock time.

---

If a guest operating system is configured for NTP and the virtual machine's independent wallclock setting is disabled, it will still periodically synchronize its time with the host time. This dual type of configuration can result in time drift between virtual machines that need to be synchronized. To effectively use an external time source, such as NTP, for time synchronization on a virtual machine, the virtual machine's independent wall-

clock setting must be enabled (set to 1). Otherwise, it will continue to synchronize its time with its host.

### **Procedure 6.1** *Viewing the Independent Wallclock Setting*

- 1 Log in to the virtual machine's operating system as `root`.
- 2 In the virtual machine environment, enter

```
cat /proc/sys/xen/independent_wallclock
```

- 0 means that the virtual machine is getting time from the host and not using independent wallclock.
- 1 means that the virtual machine is using independent wallclock and managing its time independent from the host.

### **Procedure 6.2** *Permanently Changing the Independent Wallclock Setting*

- 1 Log in to the virtual machine environment as `root`.
- 2 Edit the virtual machine's `/etc/sysctl.conf` file.
- 3 Add or change the following entry:

```
xen.independent_wallclock=1
```

Enter 1 to enable or 0 to disable the wall clock setting.

- 4 Save the file and reboot the virtual machine operating system.

When it boots, a virtual machine gets its initial clock time from the host. Then, if the wallclock setting is set to 1 in the `sysctl.conf` file, it manages its clock time independently and does not synchronize with the host clock time.

### **Procedure 6.3** *Temporarily Changing the Independent Wallclock Setting*

- 1 Log in to the virtual machine environment as `root`.
- 2 Enter the following command:



```
echo "1" > /proc/sys/xen/independent_wallclock
```

Enter 1 to enable or 0 to disable the wall clock setting.

**3** Add or change the following entry:

```
xen.independent_wallclock=1
```

Enter 1 to enable or 0 to disable the wall clock setting.

Although the current status of the independent wallclock changes immediately, its clock time might not be immediately synchronized. It might take a long time to sync. The setting persists until the virtual machine reboots. Then, it gets its initial clock time from the host and uses the independent wallclock according to setting specified in the `sysctl.conf` file.

## 6.3.5 Updating a Network Installation Source

The installation of SUSE Linux Enterprise Server 9 is only supported from a network installation source. To have the right device names supported in Xen, you must update the `kernel` and `initrd` that are used to install the system. Furthermore, the updated kernel must be available in the installation source. In the following example, the network installation source is found at `/srv/ftp`. Create this directory manually, if it does not yet exist.

- 1** Get the latest kernel package from the Novell Customer Center. Depending on your system,
- 2** Create a directory for executables in your home directory: `mkdir -p $HOME/bin`
- 3** Copy the script `create_update_source.sh` from `http://www.suse.de/~ug/tools/create_update_source.sh` to the `bin/` directory and make it executable.

```
cd $HOME/bin
wget http://www.suse.de/~ug/tools/create_update_source.sh
chmod 755 create_update_source.sh
```

- 4 Install the package `inst-source-utils`. Then, change your working directory to your network installation source.
  - 5 Run the command `$HOME/bin/create_update_source.sh /srv/ftp`.
  - 6 Copy all updated packages to the directory `/srv/ftp/updates/suse/<arch>/.`
  - 7 Run the following commands to make all the new packages known to the installation source:
- 8 Create the checksums needed for the installation process with the commands:

```
cd /srv/ftp/updates/suse;  
perl /usr/bin/create_package_descr -x setup/descr/EXTRA_PROV  
  
cd /srv/ftp/updates/suse/setup/descr  
for i in *; do echo -n "META SHA1 "; \  
shasum $i|awk '{ORS=" "; print $1}'; \  
echo -n " "; basename $i; done >> /srv/ftp/updates/content
```

After this procedure, the packages that are copied to the updates directory are available during the installation. However, they will only be used if they are newer than the packages provided by the installation itself.

Note, that in order to use a new kernel during the installation, you must also create an appropriate `installation initrd` as it is found in `/srv/ftp/boot/`.

## 6.4 Windows Virtual Machines

The installation of Windows virtual machines is supported by the virtual machine manager. To install a Windows VM Guest, select the respective system from the list provided there.

To get the best available support for Windows systems, install the driver pack. For more information about this, see [Section 7.3, “Using the Installation Wizard to Install the Driver Pack”](#) (page 71).

# Virtual Machine Drivers

SUSE® Linux Enterprise Virtual Machine Driver Pack contains disk and network device drivers for a number of third-party operating systems in order to enable the high-performance hosting of these unmodified guest operating systems on top of SUSE Linux Enterprise Server 10 (SLES 10) SP1. The guest operating systems supported by this driver pack are:

- Red Hat\* Enterprise Linux\* 4
- Red Hat Enterprise Linux 5
- Windows Server\* 2003
- Windows XP
- Windows 2000 Server
- SUSE Linux Enterprise Server 9
- SUSE Linux Enterprise Server 10

Virtualization allows the consolidation of Linux workloads on newer, more powerful, energy-efficient hardware. Paravirtualized operating systems such as SUSE Linux Enterprise Server are aware of the underlying virtualization platform, and can therefore interact efficiently with it. Unmodified operating systems are unaware of the virtualization platform and expect to interact directly with the hardware. Because this is not possible when consolidating servers, the hardware must be emulated for the operating system. Emulation can be slow, but it is especially troubling for high-throughput disk and network subsystems. Most performance loss occurs in this area.

The device drivers in SUSE Linux Enterprise Virtual Machine Driver Pack bring many of the performance advantages of paravirtualized operating systems to unmodified operating systems because only the paravirtualized device driver (not the rest of the operating system) is aware of the virtualization platform. For example, a paravirtualized disk device driver appears as a normal, physical disk to the operating system. However, the device driver interacts directly with the virtualization platform (with no emulation) to efficiently deliver disk access, allowing the disk and network subsystems to operate at near native speeds in a virtualized environment, without requiring changes to existing operating systems.

---

**TIP: Best Practices**

Before installing SUSE Linux Enterprise Virtual Machine Driver Pack in your production environment, we strongly recommend that you run it in a test environment to ensure that it functions properly with your system.

---

## 7.1 Buying a Support Subscription

The customer support you receive for the SUSE Linux Enterprise Virtual Machine Driver Pack is at the same level as your SUSE Linux Enterprise support subscription [[http://support.novell.com/linux/sles\\_support.html](http://support.novell.com/linux/sles_support.html)]. You can choose the following levels of support:


- Basic
- Standard
- Priority

Descriptions of these levels of support and how to buy them are found at SUSE Linux Enterprise Server 10 How to Buy [<http://www.novell.com/products/server/howtobuy.html>].

## 7.2 Loading the Driver Pack

- 1 Download the SLES-Virtual-Machine-Driver-Pack-10.iso file to your SLES 10 SP1 virtual machine host server.

- 2 Open YaST.
- 3 Select *Software*, then click *Add-on Product*.
- 4 Click *Local Directory*, then click *Next*.



- 5 Select the *ISO Image* box, then browse to the location of the driver pack ISO image you downloaded.
- 6 Select the driver pack ISO, then click *Open*. Click *OK*.
- 7 Select *Patterns* from the *Filter* drop-down list. The drivers are located under Virtual Machine Driver Pack.
- 8 Select the drivers you want to install, then click *Accept*.
- 9 When asked if you want to install more packages, click *No*.

The drivers are now on your system in the `/opt/novell/vm-driver-pack` directory.

## 7.3 Using the Installation Wizard to Install the Driver Pack

---

### TIP

Do not give the server a fixed IP address before installing the driver pack. Else you have to switch to a different intermediate IP address before finishing the installation.

---

- 1 Launch the Virtual Machine Manager and start your Windows virtual machine.

- 2 Set up the ISO image or CD of the drivers you want to install as a virtual CD reader by following the instructions at [Section 5.1.3, “Adding Virtual CD Readers”](#) (page 33). The ISO images are found in the SLES 10 virtual machine host `opt/novell/vm-driver-pack/xen` directory .
- 3 Right-click the virtual CD-ROM you just added, then click *Autoplay*.

If you see a security alert that indicates the Bus driver has not been properly signed, click *Yes* to install the Bus driver software. Important: Do not click *No*.




The Novell Virtual Device Driver Installer launches.



- 4 Click *Next*.

The End User License Agreement window appears.




- 5 Click *I accept this agreement*, then click *Next*.

A software installation note appears, indicating that the software has not passed Windows logo testing.

If you are installing Windows XP or Windows Server 2003, you see the following:



If you are installing Windows 2000, you see the following:



**6** Click *Continue Anyway* or click *Yes* for each occurrence of this window.

A Congratulations window indicates that the drivers were installed.




## NOTE

Clicking *Cancel* on this screen does not cancel the installation. If you want to cancel, you will have to uninstall (see [Section 7.3.1, “Uninstalling the Driver Pack”](#) (page 77)).

**7** Click *Finish*.

A message appears, indicating that you must restart the system.






8 Click *Yes*.

The system restarts. The Found New Hardware Wizard appears, indicating that it has found new hardware.



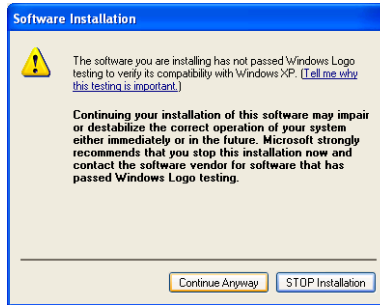
9 Click *Yes this time only*, then click *Next*.




10 Click *Install the software automatically*, then click *Next*.

A software installation note appears, indicating that the software has not passed Windows logo testing.

If you are installing Windows XP or Windows Server 2003, you see the following:




If you are installing Windows 2000, you see the following:



**11** Click *Continue Anyway* or click *Yes*.

A message indicates the Found New Hardware Wizard is done.



**12** Click *Finish*.

**13** If you are installing the Windows 2000 Server, you must reboot the system again.

## 7.3.1 Uninstalling the Driver Pack

**1** Make sure the installation CD is detached from the virtual machine.

**2** Browse to `c:\Program Files\Novell\XenDrv`.

**3** Double-click `uninstall.exe`.

You will be prompted to reboot the system.

**4** Close all applications that are running and click *OK*.

The system restarts. The Found New Hardware Wizard appears, indicating that it has found new hardware.

**5** Click *Yes, this time only*, then click *Next*.

The wizard asks to install software for the PCI Device.



- 6 Click *Cancel*.

The driver pack is now uninstalled from your system.

## 7.3.2 Upgrading the Driver Pack

To upgrade to a newer version of the driver pack:

- 1 Uninstall the current driver pack (see [Section 7.3.1, “Uninstalling the Driver Pack”](#) (page 77)).
- 2 Download and install the new driver pack.

## 7.3.3 Recovering from a Catastrophic Failure

In the event of a catastrophic failure caused by the driver pack, you can safely remove the driver pack and return to your previous system configuration.

- 1 Shut down the virtual machine and reboot in safe mode.
- 2 Uninstall the driver pack (see [Section 7.3.1, “Uninstalling the Driver Pack”](#) (page 77)).

## 7.3.4 Avoiding Problems with the Drivers

To avoid problems and potential failure of the device drivers in the driver pack, you should avoid the following actions:

- Using the Device Manager to manage (Update ..., Disable, Uninstall, etc.) the device drivers in the driver pack.
- Deleting driver files manually. Always use the uninstall utility (see [Section 7.3.1, “Uninstalling the Driver Pack”](#) (page 77)).
- Uninstalling the driver pack with the installation CD attached to the virtual machine.
- Canceling the install when you see security alert that indicates the Bus driver has not been properly signed. You must click *Yes* to continue the installation.
- Upgrading the Linux<sup>\*</sup> kernel of the virtual machine host without upgrading the driver pack at the same time.
- Installing or uninstalling the driver pack by any other process than what is documented in this guide (see [Section 7.3, “Using the Installation Wizard to Install the Driver Pack ”](#) (page 71) and [Section 7.3.1, “Uninstalling the Driver Pack”](#) (page 77)).

## 7.4 SUSE Linux Enterprise Server 9

---

### NOTE

The drivers for SLES 9 are not included in the driver pack. They are available through the standard update channels.

---

This section provides instructions for running SLES 9 SP4 as a fully virtual guest with Xen drivers. Running SLES 9 SP4 requires a VM Host Server based on SUSE Linux Enterprise Server 10 SP2 or later.

To use this software, you must have SUSE Linux Enterprise Server 9 SP4 or later.

## 7.4.1 Setting Up a new Virtual Machine

- 1 Create a new fully virtual machine or start an existing fully virtual machine. Use *YaST > Virtualization > Virtual Machine*.
- 2 Get the latest updates for the operating system from your update channel. For information on updating SLES 9, see the SUSE Linux Enterprise Server 9 Documentation [<http://www.novell.com/documentation/sles9/index.html>].

---

### IMPORTANT

Make sure you perform this step.

---

- 3 If you want to back up the driver files, open a terminal and enter `cp initrd-xxx.img to initrd- xxx.img.bak` , where `xxx` is the kernel version.
- 4 Verify that the Xen drivers are installed.
  - 4a Launch YaST by entering `yast2` at a terminal.
  - 4b Select *Software > Software Management*.
  - 4c Search for `xen-kmp` and make sure that the appropriate drivers are installed for your system.

If they are not selected, select them and click *Accept* to complete the installation.

## 7.4.2 Using an Existing SLES 9 Virtual Machine

In SUSE Linux Enterprise Server 10, the device naming is different than the device naming of SUSE Linux Enterprise Server 9. Therefore, a SUSE Linux Enterprise Server 9 VM Guest will not be able to find its root file system when running on a SUSE Linux Enterprise Server 10 VM Host Server.

To be able to still boot the system, you must know which device is used for the root partition of your virtual system. For example, `hdaxx` will be changed to `xvdaxx` where `xx` is the partition number.

When booting the system, you have to append an extra root option to the kernel command line, that tells the system about its root file system. If your VM Guest used to live on `/dev/hda2`, append the string `root=/dev/xvda2` to the kernel command line. This option should enable you to boot the system, although additional filesystems still will not be available to the system.

To make all the needed file systems available to the VM Guest, edit the file `/etc/fstab` and correct all device entries.

In order to have a valid initial ramdisk that knows about the new location of the root filesystem, run the command `mkinitrd`.

- 1 Edit the virtual machine's `/boot/grub/menu.lst` file. At the kernel line, append an `xvd x=noprobe` parameter where `xvdx` is the drive designation specified in the virtual machine configuration file.

For example, if the virtual machine configuration file specifies a virtual disk as `hda`, such as:  
`disk=[ 'file:/var/lib/xen/images/vml/hda,ioemu:xvda,w' ],`  
you should see something like `kernel vmlinuz-kernel_ver`  
`xvda=noprobe` appended to the kernel line.

- 2 Reboot the virtual machine.

## 7.4.3 Editing the Configuration File to Recognize the Virtual Network Card

- 1 On the virtualization host server, open a terminal and enter `xm delete vm_name` to delete the virtual machine's reference in the Xen store so it is no longer managed.

---

### TIP

Enter `xm list` to see the name of the virtual machine (`vm_name`).

---

- 2 Open the virtual machine configuration file located on virtualization host server (`/etc/xen/vm/vm_cfgfile`, where `vm_cfgfile` is usually the same as the name of the virtual machine).
- 3 Remove all parameters from the `vif=` line except the `mac=mac_address` parameter. Make sure to leave the single quotes.

For example, change the line

```
vif=[ 'mac=00:16:3e:55:24:21,model=pcnet,type=ioemu', ]
```

to

```
vif=[ 'mac=00:16:3e:55:24:21', ]
```

- 4 At the terminal, enter `xm new vm_cfgfile` to reference the virtual machine in the Xen store so it is managed again.
- 5 Restart the virtual machine by using `YaST > Virtualization > Virtual Machine Manager`.

You can also restart the virtual machine by using `xm start vm_cfgfile`, but make sure to connect to it with a viewer.

The first time the virtual machine boots, you might see a message indicating that new hardware was found. This message can be safely canceled.

## 7.4.4 Configuring the Virtual Network Card

- 1 Identify the number of each virtual network card.

Virtual network cards are identified by `vif-x` files located in the `/sys/class/net/ethx/driver` directories. For example, network card 0 is designated as `/sys/class/net/eth0/driver/vif-0`. Look in each `ethx/driver` directory for a `vif-x` file, such as `vif-0`, `vif-1`, and so on.

- 2 Run YaST.

**2a** Make sure you are logged in as root.



**2b** Open a terminal.

**2c** Enter `YaST2`.

**3** In YaST, click *Network Devices > Network Card*.

**4** Click *Change*, then select and delete each network card that is already configured.

**5** Click *Add*, then change the `ethx` parameter to a network card number identified above. Repeat the process to add all virtual network cards.

## 7.5 SUSE Linux Enterprise Server 10

---

### NOTE

The drivers for SLES 10 are not included in the driver pack. They are available through the standard update channels.

---

This section provides instructions for running SLES 10 SP1 as a fully virtual guest with Xen drivers.

To use this software, you must have SUSE Linux Enterprise Server 10 SP1 or later.

### 7.5.1 Setting Up the Virtual Machine

**1** Create a new fully virtual machine or start an existing fully virtual machine. Use *YaST > Virtualization > Virtual Machine Manager*.

**2** Get the latest updates for the operating system from your update channel. For information on updating SLES 10, see the SUSE Linux Enterprise Server 10 SP1 documentation [<http://www.novell.com/documentation/sles10/index.html>].

**3** Verify that the Xen drivers are installed.

**3a** Launch YaST by entering `yast2` at a terminal.

**3b** Select *Software > Software Management*.

**3c** Search for `xen-kmp` and make sure that the appropriate drivers are installed for your system.

If they are not selected, select them and click *Accept* to complete the installation.

## 7.5.2 Setting Up the Virtual Machine's Block Device

**1** Edit the virtual machine's `/boot/grub/menu.lst` file. At the kernel line, append an `hdx=noprobe` parameter where *hdx* is the drive designation specified in the virtual machine configuration file.

For example, if the virtual machine configuration file specifies a virtual disk as `xvda`, such as:

```
disk=['file:/var/lib/xen/images/vml/xvda,ioemu:xvda,w']
```

the following kernel line should be appended:

```
kernel vmlinuz-kernel_ver xvda=noprobe
```

Also at the kernel line, change `root=/dev/disk/by-id/xxx` to `root=/dev/hdxx`, where *hdxx* is the disk and partition of your root partition. For example, `root=/dev/xvda2`.

**2** Shut down the virtual machine.

## 7.5.3 Editing the Configuration File to Recognize the Virtual Network Card

**1** On the virtualization host server, open a terminal and enter `xm delete vm_name` to delete the virtual machine's reference in the Xen store so it is no longer managed.

---

**TIP**

Enter `xm list` to see the name of the virtual machine (*vm\_name*).

---

- 2 Open the virtual machine configuration file located on virtualization host server (`/etc/xen/vm/vm_cfgfile`, where *vm\_cfgfile* is usually the same as the name of the virtual machine).
- 3 Remove all parameters from the `vif=` line except the `mac= mac_address` parameter. Make sure to leave the single quotes.

For example, change the line

```
vif=[ 'mac=00:16:3e:55:24:21,model=pcnet,type=ioemu', ]
```

to

```
vif=[ 'mac=00:16:3e:55:24:21', ]
```

- 4 At the terminal, enter `xm new vm_cfgfile` to reference the virtual machine in the Xen store so it is managed again.
- 5 Restart the virtual machine by using *YaST* > *Virtualization* > *Virtual Machine Manger*.

You can also restart the virtual machine by using `xm start vm_cfgfile`, but make sure to connect to it with a viewer.

The first time the virtual machine boots, you might see a message indicating that new hardware was found. This message is to be expected and can be safely cancelled.

## 7.5.4 Configuring the Virtual Network Card

- 1 Identify the number of each virtual network card.

Virtual network cards are identified by `vif-x` files located in the `/sys/class/net/ethx/device/driver` directories. For example, network card 0 is

designated as `/sys/class/net/eth0/driver/vif-0`. Look in each `ethx/driver` directory for a `vif-x` file, such as `vif-0`, `vif-1`, and so on.

- 2** In YaST, click *Network Devices* > *Network Card*. Select and delete each network card that is already configured.
- 3** Click *Add*, then change the `ethx` parameter to a network card number identified above and enter `xen-vnif` as the module name. Repeat the process to add all virtual network cards.