VMware for Linux

Version 1.1

We encourage you to refer to our Web site for the latest product and documentation updates.

You will always find the most up to date technical documentation on our Web site at **http://www.vmware.com/support**

VMware, Inc
3145 Porter Drive, Bldg. F
Palo Alto, CA 94304
http://www.vmware.com

Table of Contents

1

VMware for Linux Installation Notes

VMware™ for Linux™ Installation Notes

System Requirements

Hardware requirements for the host machine

- Pentium® II 266MHz and above recommended (Pentium class minimum)
- 64MB RAM minimum (128MB and above recommended)
- Video adapter supported by the XFree86 server (to take advantage of our full-screen option)

Software requirements

VMware host operating system requirements:

- glibc2 to glibc6 (glibc1 does not work)
- Linux kernel 2.0.32 to 36
- Linux kernel 2.2.0 and higher for SMP systems

VMware will not run on systems that do not meet these requirements.

VMware has been tested and is officially supported on the following host operating systems:

- Red Hat™ Linux 5.0 and higher
- Caldera® OpenLinux™ 1.3 (requires optional glibc support installed)
- SuSE™ Linux 6.0 and higher

VMware may run on other Linux distributions. Attempting to do so is recommended for expert Linux users only.

VMware requires an X-Server:

- XFree86-3.3.4 or higher recommended

Installation Steps

Installing VMware for Linux from an X console

1. Mount the VMware for Linux distribution CD-ROM.
2. Go to the vmware-distrib directory on the CD.
3. Become root.

   ```
   su
   ```

4. Install the program.

   ```
   ./vmware-install.pl
   ```

   Answer the questions from the script or hit return to use defaults. See the file INSTALL for more information.

   The installation program should display "Installation successful" on the last line of the install script. If it does not, please run the install program again.

5. Exit from the root shell.

   ```
   exit
   ```

   This completes the installation.

6. Switch into a directory from which you would like to run the virtual machine.

7. Run the command vmware .

   The application starts. You need to configure a virtual machine to run for the first time, as described below, or click Exit to abort the configuration.

   NOTE: VMware for Linux requires glibc versions 2 through 6 to run without error. Red Hat 5.0/5.1/5.2/6.0, SuSE 6.0/6.1 and Caldera OpenLinux 2.2 install glibc as part of their normal installation. However, Caldera OpenLinux 1.3 does not install glibc automatically even though the package is included on the distribution CD. To install the package, mount the Caldera OpenLinux 1.3 distribution CD, switch to the col/contrib/RPMS directory and run

   rpm -i glibc-2.06-1.i386.rpm

   If you are having problems installing, compiling and/or loading the vmmon and vmnet modules, see the note at http://www.vmware.com/support/vmodules.html.

Installing VMware for Linux using Red Hat Package Manager

You may use the Red Hat Package Manager to install VMware for Linux. If you wish to do so, please download the VMware for linux RPM package from our web site at http://www.vmware.com/download/downloadlinux.html. The following three sections document the most common situations that you may face when you install or upgrade using the RPM distribution.

Case 1 - First time installation of VMware using the RPM distribution

1. Become root.

   su

2. To install the RPM package from a Linux terminal window or console, type

   rpm -Uhv VMware-<version>-<build>.i386.rpm

   NOTE: It is possible to install the RPM package using other RPM installers. For example, if you are running Red Hat Linux 6.0 with GNOME you can right-click on the RPM package from the File Manager and select Upgrade to install the new software. You can use the RPM installer of your choice; however, you must be sure that you are root.

3. Run the configuration program

   /usr/bin/vmware-config.pl

   to configure the VMware software for use. Follow the on-screen instructions to finish the VMware for Linux software installation process. Follow the on-screen instructions regarding how to start VMware for Linux.

4. Exit from the root shell.

   exit

   This completes the installation.

   If you decide to uninstall the RPM package in the future, you may do so by typing

   rpm -e vmware

Case 2 - Tar distribution installed, upgrading to RPM distribution

1. Become root.

   su

2. To install the RPM package from a Linux terminal window or console, type

   rpm -Uhv VMware-<version>-<build>.i386.rpm

NOTE: It is possible to install the RPM package using other RPM installers. For example, if you are running Red Hat Linux 6.0 with GNOME you can right-click on the RPM package from the File Manager and select Upgrade to install the new software. You can use the RPM installer of your choice; however, you must be sure that you are root.

3. Run the configuration program

   /usr/bin/vmware-config.pl

   to configure the VMware software for use. Follow the on-screen instructions to finish the VMware for Linux software installation process. Follow the on-screen instructions regarding how to start VMware for Linux.

4. Exit from the root shell.

   exit

   This completes the installation.

   If you decide to uninstall the RPM package in the future, you may do so by typing

   rpm -e vmware

Case 3 - SuSE-provided RPM installed, upgrading to new RPM distribution

1. Become root

   su

2. Remove the existing VMware for Linux RPM package that was provided on the SuSE Linux 6.2 CD. To do this type the following:

   rpm -e vmware

   rpm -e vmmodule

3. To install the RPM package from a Linux terminal window or console, type

   rpm -Uhv VMware-<version>-<build>.i386.rpm

   NOTE: It is possible to install the RPM package using other RPM installers. For example, if you are running Red Hat Linux 6.0 with GNOME you can right-click on the RPM package from the File Manager and select 'Install...' to install the new software. You can use the RPM installer of your choice; however, you must be sure that you are root.

4. Run the configuration program

   /usr/bin/vmware-config.pl

   to configure the VMware software for use. Follow the on-screen instructions to finish the VMware for Linux software installation process. Follow the on-screen instructions regarding how to start VMware for Linux.

5. Exit from the root shell.

   exit

   This completes the installation.

   If you decide to uninstall the RPM package in the future, you may do so by typing

   rpm -e vmware

Installing your license

To install the license, create the directory ~/.vmware , and copy the contents of the file named license from the CD to ~/.vmware/license .

Configuring a Virtual Machine

Run the command vmware. The binary is in the directory where you installed the VMware application. If you use csh/tcsh, you may need to run rehash if you want the system to find the executable automatically.

You will now need to configure your virtual machine. VMware recommends that you select the wizard option.

Comments

Make sure no existing devices use these device numbers:

- /dev/vmmon uses major device number 10, minor device number 165
- /dev/vmnet* uses major device number 119, minor device numbers 0-3

"Real Time Clock" function needs to be compiled into your Linux kernel.

Uninstalling VMware Software

To uninstall the software, run the vmware-install. pl script in the vmware-distrib directory on the VMware CD. When a prompt asks if you want to do an upgrade, answer No. You will then be asked if you want to uninstall the software. Answer Yes.

2

Creating, Configuring, and Operating Virtual Machines Using VMware

Creating, Configuring, and Operating Virtual Machines Using VMware

Before you can run a supported operating system in the virtual machine, you must create a virtual machine configuration file. This is equivalent to building a PC from parts, except there is no risk of losing screws or zapping components with static discharge.

A virtual machine configuration is kept in a text file. This file can be created using the VMware Configuration Wizard or the Configuration Editor from the Settings menu.

A. Creating and Configuring Your Virtual Machine

1. Before starting the configuration you will need to create a directory for your first virtual machine.

   mkdir <name of directory>

2. Go to that directory.

   cd <name of directory>

3.Launch the VMware program.

   vmware

4. If you are creating a virtual machine configuration file for the time first time, VMware recommends using the Run the Configuration Wizard option to help you step through the process.

   If you want to manually create a new configuration file and configure all aspects of the virtual machine, use the Run the Configuration Editor option and continue with the rest of the steps below.

5. Create a virtual disk and associate it to a virtual machine's IDE channel.
   - Click the + next to IDE Devices.
   - Click P-M - Not Installed.
   - Choose Device Type “Hard Disk.”
   - Select an appropriate disk mode for the virtual hard disk.
   - In the Name field, type the filename that you want to use for this virtual disk.
   - Complete the path of the virtual disk by clicking the Browse button and selecting the directory you created for the virtual machine earlier. The virtual disk image will be saved to a file in this directory.
   - Use the default disk size or enter a new value in the Disk Capacity field. The specified size assigns the maximum capacity of the virtual disk. Initially, the space taken by the file will less than 1MB, but the file will grow as software is installed inside the virtual machine. The file will never grow larger than the size specified in the Disk Capacity field.
   - Create a virtual disk file by clicking the Create button.
   - If you selected an Undoable disk mode and want to store the virtual machine's REDO file in a different directory than the actual virtual disk file, click Misc. and specify the new path.
   - Click Install.

NOTE: If you want to configure an additional hard disk, Click "P-S, S-M or S-S - Not Installed" and repeat the above steps.

6. Configure and associate a CD-ROM drive to a virtual machine's IDE channel.
   - Click Not Installed.
   - Select CD-ROM from the Drive Type field.
   - Enter /dev/cdrom or the path to the CD-ROM in the Name field.
   - Click Install.

7. The first floppy drive is configured by default. To add another floppy.
   - Click the + next to Floppy Drives.
   - Click “Not Installed.”
   - Enter /dev/fd1 or the path to your floppy device.
   - Click Install.

8.Optionally configure a network adapter for the virtual machine.
   - Click the + next to Ethernet Adapters.
   - Click Not Installed.
   - Use the default Bridged option (if your real machine does not have an Ethernet card, use Host-only or remove the adapter).
   - Click Install.

9. Add serial port (configure serial mouse under Mouse button).
   - Specify the path your host machine is using to the device.

10. Add parallel port.
   - Specify the path your host machine is using to the device.

11. Change the mouse to a serial mouse (or other).
   - Click on Autodetect and change the setting to the mouse you have.

12. Change default memory settings (32MB).
   - Use the scroll bar to adjust the amount of memory you would like the virtual machine to recognize.

13. Click OK in the bottom of the Configuration Editor and click Save to save the changes. Ensure the path is pointing to the directory where you’re running the virtual machine and enter a file name (any file name you wish).

14. At this point, you can start using the virtual machine you have configured. For additional help on how to run and operate it, refer to Operating Your Virtual Machine below. If you are familiar with using virtual machines and want to install a new operating system for the virtual machine you configured earlier, refer to the Guest Operating System Installation Instructions page on the VMware Web site.

    NOTE: Remember that if you start up a new virtual machine and do not have a bootable CD-ROM or floppy in the drive, the virtual machine will come up with on the display "Operating System not Found."

B. Operating Your Virtual Machine

Your virtual machine is now ready for you to install your choice of the guest operating systems that VMware currently supports. To familiarize yourself with using VMware software, review the section below, then proceed to install an operating system within the virtual machine environment by referring to the Guest OS Installation section.

Mouse and keyboard control

After a virtual machine has been powered up, the mouse and keyboard will continue to work normally on the host. If you click a mouse button while the mouse pointer is on top of the VMware application window, the running virtual machine will grab the mouse and keyboard focus. While the mouse and keyboard are grabbed by the virtual machine, the user will be able to interact with software running inside the virtual machine; the user will not be able to interact with the host OS and applications at this point. To be able to use the mouse/keyboard on the host OS again, press the Ctrl-Alt-Esc keys together. Regardless of whether the virtual machine is in full-screen or window mode, pressing Ctrl-Alt-Esc will release the mouse and keyboard from the virtual machine and return them to the host OS.

Switching between full-screen and window modes

Virtual machines appear in window mode when they are initially powered up. To switch the virtual machine to full-screen, select the Full-screen option from the VMware application's View menu. It is also possible to switch to full-screen by pressing the Ctrl-Alt-F8 keys together. The latter option is particularly useful if you have multiple virtual machines running and want to switch a particular one to full-screen quickly. You can do this by using the Ctrl-Alt-Fx combination keys together; where Fx can be F8, F9, F10, F11, F12, etc. Pressing Ctrl-AltT-F9 will switch the second virtual machine to full-screen, Ctrl-Alt-F10 will switch the third virtual machine to full-screen, and so on.

The description above is true for most Linux systems, where Ctrl-Alt-F7 takes you to the X server. If you have remapped your keys, the first virtual machine will use the key combination following the one that takes you to the X server. For example, if you have remapped your keys so that Ctrl-Alt-F5 takes you to the X server, Ctrl-Alt-F6 will switch the first VMware virtual machine to full-screen mode.

To switch any virtual machine from full-screen back to window mode, press the Ctrl-Alt-Esc keys together.

Accessing the same virtual machine in the future

After the virtual machine is powered down and the VMware application is closed, you can start the virtual machine again by specifying the name of the configuration file when you launch VMware. The following example is for a Windows 98 guest OS.

1. Open a terminal, then change to the virtual machine's directory

   cd /home/username/vmware/win98

   where username is the non-root user you are logged in as and win98 is the directory where the configuration file is stored.

2. Type

   vmware win98.cfg

3. Click the Power On button to power the virtual machine on.

Installing a new operating system in the virtual machine environment

When you power on a virtual machine configured with a virtual disk for the first time, you should have the boot floppy or CD of the operating system you plan to install in the appropriate drive.

After a virtual machine is powered up, its BIOS starts to run and tries to boot from a boot device in the following order: floppy, hard disk, CD-ROM. To change the default boot order, time/date or other settings you can enter the BIOS Setup by pressing the F2 key while the virtual machine's memory is being tested. Data stored in a virtual machine's CMOS is actually saved to a file named nvram in the virtual machine directory.

To begin installing an operating system inside a virtual machine, see the next section, Installing Guest Operating Systems Inside a Virtual Machine. Once an operating system is installed, you may use the virtual machine to do anything the OS supports, given the virtual devices available in the virtual machine environment. You may install applications and work in your new virtual environment.

Configuring a new virtual machine to access a raw disk

Advanced users can configure their virtual machine to use a raw disk instead of a virtual disk. VMware recommends the following for advanced users only. If you have a dual- or multiboot environment, you may configure a virtual machine to boot a previously installed operating system from an existing raw partition; see page 4-8. If you have free space or an unused partition in your real machine, you may install a new operating system to the raw disk through the virtual machine; see page 4-11.

3

Installing Guest Operating Systems inside a Virtual Machine

Installing Guest Operating Systems Inside a Virtual Machine

MICROSOFT® WINDOWS® 2000 INSTALLATION GUIDELINES AND KNOWN PROBLEMS

Problems have been identified with version of Windows 2000 prior to RC1. For example, the Windows 2000 Server Beta 3 installation program may fail to set up Internet Information Server properly. The only recovery to the installation process is to cancel the IIS installation and continue with the Windows 2000 installation.

There have also been problems in trying to add IIS after completing a Windows 2000 Server Beta 3 installation. You may also see the following VMware error message: "The cpu has been disabled by the guest operating system. You will need to power off or reset the virtual machine at this point." This is most likely because the Windows 2000 guest operating system has panicked or "blue screened." If you encounter any such problems, please report them to VMware using the form on the VMware Web site at http://www.vmware.com/incident.

Windows 2000 Professional or Server (RC1 or later versions) can be installed in a virtual machine using the corresponding Windows 2000 distribution CD. Before installing the operating system, be sure that you have already created a new virtual machine and configured it using the VMware Configuration Wizard (or Editor).

Windows 2000 installation steps

1. Before starting the installation, use the VMware Configuration Editor to verify the virtual machine's devices are set up as you expect. For example, if you would like networking software to be installed during the Windows 2000 installation, make sure the virtual machine's Ethernet adapter is configured and enabled.
2. Insert the Windows 2000 CD in the CD-ROM drive.
3. Power on the virtual machine to start installing Windows 2000.
4. If you enabled the virtual machine's Ethernet adapter, then an "AMD PCNET Family Ethernet Adapter" will be detected and set up automatically.
5. Finish the Windows 2000 installation. After Windows has been installed, install the VMware Tools – using the file vmware-toolsxxx.exe on the VMware CD-ROM – for improved video performance and added functionality.

Enabling Sound After Installing Windows 2000

If sound was disabled during the Windows 2000 installation, it can be enabled after the OS has been installed. To set up the virtual machine to play sound, please read the technical note about VMware and Sound, page 4-15.

Enabling Networking After Installing Windows 2000

If networking was disabled during the Windows 2000 installation, it can be enabled after the OS has been installed. To set up networking for a virtual machine, follow the instructions below:

1. Shut down Windows 2000 and power off the virtual machine.
2. From the VMware application window, select Configuration Editor from the Settings menu and click the + beside Ethernet Adapters, then click Ethernet.
3. Select a network connection type for the virtual machine and click the Install button.
4. Save the updated configuration and power on the virtual machine.

5. When Windows 2000 boots, it will automatically detect a new network adapter and load drivers for an AMD PCNET Family PCI Ethernet Adapter.
6. You should be able to access the network after logging on to the Windows 2000 guest OS.

WINDOWS NT® INSTALLATION GUIDELINES AND KNOWN PROBLEMS

Windows NT 4.0 can be installed in a virtual machine using the standard Windows NT CD-ROM. Before installing the OS, be sure that you have already created a new virtual machine and configured it using the VMware Configuration Wizard (or Editor).

Windows NT installation steps

1. Use the VMware Configuration Editor to verify the virtual machine's devices are set up as you expect before starting the installation. For example, if you would like networking software to be installed during the Windows NT installation, be sure the virtual machine's Ethernet adapter is configured and enabled. VMware also recommends that you disable the screen saver on the host system before starting the installation process.
2. Insert the Windows NT CD in the CD-ROM drive.
3. Power on the virtual machine to start installing Windows NT.
4. If you enabled the virtual machine's Ethernet adapter, then an AMD PCNET Family Ethernet Adapter will be detected and set up automatically. The default settings should work fine and do not need to be changed.
5. Finish the Windows NT installation.
6. VMware's virtual disks support DMA transfers for better performance. The feature can be enabled after Windows NT has been successfully installed. You will need the NT Service Pack 3 or 4 CD to enable this option. Once the virtual machine is running Windows NT, insert the SP3 or SP4 CD in the drive, run DMACHECK.EXE from the \SUPPORT\UTILS\I386 directory on the CD and click the Enabled option for both IDE channels.

After Windows has been installed, be sure you install the VMware Tools for improved video performance and added functionality.

Enabling sound after installing Windows NT

If sound was disabled during the Windows NT installation, it can be enabled after the OS has been installed. To set up the virtual machine to play sound, please read the technical note VMware and Sound.

Enabling networking after installing Windows NT

If networking was disabled during the Windows NT installation, it can be enabled after the OS has been installed. To set up networking for a virtual machine, follow the instructions below:

1. Shut down Windows NT and power off the virtual machine.
2. From the main program window, select Configuration Editor from the Settings menu and click the + next to Ethernet Adapters.
3. Click Ethernet – Not Installed.
4. Select a network connection type for the virtual machine and click the Install button.
5. Save the updated configuration and power on the virtual machine.
6. While Windows NT is booting, insert the Windows NT 4.0 CD in the CD-ROM drive.

7. Log in to Windows NT and install the AMD PCNET driver:
   a. Open the Network properties page by double-clicking the Network icon in Control Panel. Change to the Network Adapters screen by clicking the Adapters tab.
   b. Use the Add button and select the AMD PCNET Family Ethernet Adapter from the list.
   c. A message will pop up prompting you to enter a path for the Windows NT files; specify the \I386 directory on the CD in the path you enter – for example, type d:\i386 if the CD is in drive D: – and click Continue.
   d. Windows NT setup will prompt you for the Windows NT files again; simply click Continue.
   e. Use the default adapter settings; they do not need to be changed. Windows NT setup will prompt you again for a path to the Windows NT files; simply click Continue to finish installing the driver.

Windows 98 Installation Guidelines and Known Problems

Windows 98 can be installed in a virtual machine using the standard Windows 98 CD. Before installing the OS, please be sure that you have already created a new virtual machine and configured it using the VMware Configuration Wizard (or Editor).

Windows 98 Installation Steps

1. Use the VMware Configuration Editor to verify the virtual machine's devices are set up as you expect before starting the installation. For example, if you would like Windows 98's Setup program to install a sound driver, be sure that Sound is enabled in the virtual machine's configuration. VMware also recommends that you disable the screen saver on the host system before starting the installation process.
2. Insert the Windows 98 CD in the CD-ROM drive.
3. Power on the virtual machine to start installing Windows 98.
4. Choose Boot from CD-ROM, then select the Start Windows 98 Setup from CD-ROM option. The setup program will run FDISK and reboot.
5. Choose Boot from CD-ROM, then select the Start Windows 98 Setup from CD-ROM option. The setup program will continue installing 98.
6. Follow the Windows 98 installation steps as you would for a real PC.

After Windows has been installed, install the VMware Tools for improved video performance and added functionality.

Enabling Sound After Installing Windows 98

If sound was disabled during the Windows 98 installation, it can be enabled after the OS has been installed. To set up the virtual machine to play sound, please read the technical note VMware and Sound.

Enabling Networking After Installing Windows 98

If networking was disabled during the Windows 98 installation, it can beenabled after the OS has been installed. To set up networking for a virtual machine, follow the instructions below:

1. Shut down Windows 98 and power off the virtual machine.
2. From the main program window, select Configuration Editor from the Settings menu and click the + next to Ethernet Adapters.
3. Click Ethernet – Not Installed.

4. Select a network connection type for the virtual machine and click the Install button.
5. Save the updated configuration and power on the virtual machine.
6. When Windows 98 reboots, it will auto-detect an AMD PCNET Family Ethernet Adapter (PCI-ISA) PCI Ethernet controller and prompt for the Windows 98 CD-ROM to install drivers. The default Ethernet adapter settings should work fine and do not need to be changed.
7. Use the Network icon from Control Panel to view or change network settings. For example, you may want to add the TCP/IP protocol since Windows 98 does not install it by default.

Windows 95 Installation Guidelines and Known Problems

Windows 95 can be installed in a virtual machine using a standard Windows 95 boot disk and CD-ROM.

NOTE: Some Microsoft Windows 95 OEM disks included with new computers are customized for those computers and include device drivers and other utilities specific to the hardware system. Even if you can install this Windows 95 operating system on your actual computer, you may not be able to install it within a VMware virtual machine. You may need to purchase a new copy of Windows to install within a virtual machine.

NOTE: The setup instructions for some Windows 95 distributions do not include the steps to fdisk and format a C: drive. You must fdisk and format the VMware virtual IDE hard disk drives before running Windows 95 setup.

The instructions below are for the simplest case of one virtual IDE hard disk drive and one virtual IDE CD-ROM drive. If you have configured the virtual machine with more than one IDE hard drive, you should also fdisk and format these drives before installing Windows 95. If you have configured the virtual machine with more than one virtual hard drive or one virtual CD-ROM, you may need to use different device letters than those in the instructions below.

Before installing the OS, please make sure that you have already created a directory for the new virtual machine and configured the virtual machine using the VMware Configuration Wizard (or Editor).

Windows 95 Installation Steps

1. Use the VMware Configuration Editor to verify the virtual machine's devices are set up as you expect before starting the installation. For example, if you would like Windows 95's Setup program to install a sound driver, be sure that sound is enabled in the virtual machine's configuration. VMware also recommends that you disable the screen saver on the host system before starting the installation process.
2. Insert the Window 95 CD-ROM Setup Boot Disk in floppy drive A: and insert the Windows 95 CD in the CD-ROM drive.
3. Power on the virtual machine.
4. After the virtual machine boots, if you are presented with a choice of CD-ROM drivers, select the first IDE driver option available (even if your computer has a SCSI CD-ROM drive).
5. Partition the virtual disk. Type:

   A:\> FDISK

   and answer the questions.

   NOTE: If you create a primary partition that is smaller than the size of the hard disk, then be sure the partition is marked active.

6. Reboot Windows 95: If the cursor is not already within the VMware window click in the window, then press Ctrl-Alt-Ins. If prompted on reboot to select a CD-ROM driver, select the first IDE CD-ROM driver from the list.

7. Format the C: drive with

   A:\> FORMAT C: /S

8. Now start the Windows 95 installation. Type:

   A:\> D:\WIN95\SETUP /IS

   NOTE: An intermittent problem can occur during Windows 95 installations in a virtual machine. Shortly after the Windows 95 Setup program is started, scandisk runs to completion, and when the Windows 95 Setup program should start its graphical user interface, the virtual machine returns to an MS-DOS prompt. VMware recommends you reboot the computer and rerun Windows 95 Setup. You will not need to fdisk and format the drive again. If this problem occurs reproducibly, please contact VMware customer support (http://www.vmware.com/support).

9. If the virtual machine's Ethernet adapter is enabled, you will have to manually add an Ethernet driver because Windows 95 will not detect it during the Analyzing Computer phase (even if you selected the Network Adapter detection option). Do the following to enable networking:
   a. Continue with the Windows 95 installation, until you get to the screen titled Windows 95 Setup Wizard/Setup Options. Change the default setting from Typical to Custom and click the Next button to continue.
   b. From the screen titled Network Configuration (which appears after the Analyzing Computer phase), click the Add button, select the Adapter component, select Advanced Micro Devices from the manufacturer window and AMD PCNET Family Ethernet Adapter (PCI & ISA) from the network adapter window.
   c. If you need TCP/IP networking, add it from the Network Configuration screen (Windows 95 Setup does not enable TCP/IP by default). If you don't do this, the first phase of the Windows 95 installation will not copy some of the files it will need later, and the entire installation will fail.

10. Finish the Windows 95 installation.

11. VMware's virtual disks support DMA transfers for better performance. The feature can be enabled after Windows 95 has been successfully installed. To enable the feature: right-click My Computer and select Properties. From the System Properties dialog box click the Device Manager tab, double-click the Disk Drives device category, double-click the GENERIC IDE DISK TYPE01 device, Click the Settings tab and enable the DMA check box.

After Windows has been installed, install the VMware Tools for improved video performance and added functionality.

Enabling Sound After Installing Windows 95

If sound was disabled during the Windows 95 installation, it can be enabled after the OS has been installed. To set up the virtual machine to play sound, please read the technical note VMware and Sound.

Enabling Networking After Installing Windows 95

If networking was disabled during the Windows 95 installation, it can be enabled after the OS has been installed. To set up networking for a virtual machine, follow the instructions below:

1. Shut down Windows 95 and power off the virtual machine.

2. From the main program window, select Configuration Editor from the Settings menu and click the + next to Ethernet Adapters.

3. Click Ethernet – Not Installed.

4. Select a network connection type for the virtual machine and click the Install button.
5. Save the updated configuration and power on the virtual machine.
6. When Windows 95 reboots, it will auto-detect an AMD PCNET Family Ethernet Adapter (PCI & ISA) PCI Ethernet controller and prompt for the Windows 95 CD-ROM to install drivers. The default Ethernet adapter settings should work fine and do not need to be changed.
7. Use the Network icon from Control Panel to view or change network settings. For example, you may want to add the TCP/IP protocol since Windows 95 does not install it by default.

Red Hat Linux 6.0 Installation Guidelines and Known Problems

The easiest method of installing Red Hat Linux 6.0 in a virtual machine is to use the standard Red Hat distribution CD. The notes below describe an installation using the standard distribution CD, however, installing Red Hat Linux 6.0 via the boot floppy/network method is supported as well. Before installing the OS, be sure that you have already configured a new virtual machine and configured it using the VMware Configuration Wizard (or Editor).

NOTE: During the Red Hat Linux 6.0 installation, a standard VGA16 X server (without support for VMware's X server) will be installed. To get an accelerated SVGA X server running inside the virtual machine, you should install the VMware Tools for Linux package immediately after installing Red Hat Linux 6.0.

Red Hat Linux 6.0 installation steps

1. Use the VMware Configuration Editor to verify the virtual machine's devices are set up as you expect before starting the installation. For example, if you would like networking software to be installed during the Red Hat Linux 6.0 install process, be sure the virtual machine's Ethernet adapter is enabled and configured.
2. Insert the Red Hat Linux 6.0 CD in the CD-ROM drive and click the Power On button. The virtual machine should start booting from the CD and the installation process will begin.
3. Follow the installation steps as you would for a real PC.

   NOTE: If the virtual machine's Ethernet adapter was enabled, the installation program will auto-detect and load the AMD PC/Net 32 driver (no command line parameter will be necessary to load the driver).
4. During the Linux installation select the standard VGA16 X server: Select the Generic VGA compatible/Generic VGA card from the list in the Choose a Card screen. Select the Generic Monitor entry from the list in the Monitor Setup screen. Select the Probe button from the Screen Configuration dialog and select OK from the Starting X dialog. After Linux is installed, the generic X server will be replaced with the accelerated X server included in the VMware Tools for Linux package.
5. Finish installing Red Hat Linux 6.0 as you would on a real PC.

At this point Red Hat 6.0 will boot and present a login screen. Follow the steps below to install VMware Tools for Linux inside the virtual machine.

VMware Tools for Linux installation steps

1. Log in as root and copy the VMware Tools for Linux package (tools-for-linux.tar.gz) to a directory in the virtual machine. If the virtual machine can access the Internet, then use Lynx, HTTP or FTP to download the package from http://www.vmware.com/support/vmwaretools.html directly to the virtual machine's disk. If you have the software on a floppy disk, mount the floppy drive under the Linux guest and copy it to a directory on the virtual machine's disk.

2. Unpack the software by running tar zxf tools-for-linux.tar.gz , then switch into the vmware-linux-tools directory.
3. Install the VMware X server by running ./install.p l. If you prefer some variation on the standard installation, look at what the Perl script does and install the files by hand.
4. Since the X server installed by Red Hat will be running at this time, startx will not work. To start the VMware X server, either kill the running X process and then run startx , or do a shutdown -r now (in which case Linux will reboot and automatically start the new X server when it comes up).
5. Once the VMware-provided X server starts, run the vmware-linux-toolbox program to take advantage of virtual machine mouse and device enhancements that make interacting with the host and other virtual machines more convenient.

SuSE Linux 6.1 and Above Installation Guidelines and Known Problems

The easiest method of installing SuSE Linux 6.1 and above in a virtual machine is to use the standard SuSE distribution CD. The notes below describe an installation using the standard distribution CD, however, installing SuSE Linux via the boot floppy/network method is supported as well. Before installing the OS, be sure that you have already created a new virtual machine and configured it using the VMware Configuration Wizard (or Editor).

NOTE: During the SuSE Linux installation, a standard VGA16 X server (without support for VMware's X server) should be installed. To get an accelerated SVGA X server running inside the virtual machine, you should install the VMware Tools for Linux package immediately after installing SuSE Linux.

SuSE Linux installation steps

1. Use the VMware Configuration Editor to verify the virtual machine's devices are setup as you expect before starting the installation. For example, if you would like networking software to be installed during the SuSE Linux installation process, be sure the virtual machine's Ethernet adapter is enabled and configured.
2. Insert the SuSE Linux CD in the CD-ROM drive and click the Power On button. The virtual machine should start booting from the CD and the installation process will begin.
3. If the virtual machine's Ethernet adapter was enabled, you must configure Linux to load the appropriate driver. Select Kernel Modules (Hardware Drivers) from the install program's main menu, then select Autoload of Modules. A pop-up message will appear stating that the pcnet32 driver will be loaded automatically. The driver's long name will be displayed as AMD PCI PCnet32 (PCI bus NE2100).
4. Follow the installation steps as you would for a real PC, until you get to the selection screens described in the next step.
5. From the Installation - YaST screen select the packages that you want to install and select the Start Installation option to continue.

   NOTE: If you want the virtual machine to use DHCP to get its IP address, verify that the dhclient package is selected before selecting Start Installation. To check, choose the Change/Create Configuration option, press the F4 key and select All Packages from the pop-up dialog, then select All Packages (Excluding Sources) and verify that the dhclient package is selected in the Package Selection screen.
6. Finish installing SuSE Linux as you would on a real PC.

At this point SuSE Linux will boot and present a login screen. Follow the steps below to install VMware Tools for Linux inside the VM.

VMware Tools for Linux installation steps

1. Log in to SuSE as root and copy the VMware Tools for Linux package (tools-for-linux.tar.gz) to the virtual machine.
2. Unpack the software by running tar zxf tools-for-linux.tar.gz , then switch into the vmware-linux-tools directory.
3. Install the VMware X server by running ./install.pl . If you prefer some variation on the standard installation, look at what the Perl script does and install the files by hand.
4. Since the X server installed by SuSE Linux will be running at this time, startx will not work. To start the VMware X server, either kill the running X process and then run startx , or do a shutdown -r now (in which case Linux will reboot and automatically start the new X server).
5. Once the VMware-provided X server starts, you can run the vmware-linux-toolbox program to take advantage of virtual machine mouse and device enhancements that make interacting with the host more convenient.

CALDERA OPENLINUX 2.2 INSTALLATION GUIDELINES AND KNOWN PROBLEMS

The easiest method of installing OpenLinux 2.2 in a virtual machine is to use the standard Caldera distribution CD. The notes below describe an installation using the standard distribution CD, however, installing OpenLinux 2.2 via the boot floppy/network method is supported as well. Before installing the OS, be sure that you have already creatred a new virtual machine and configured it using the VMware Configuration Wizard (or Editor).

NOTE: During the OpenLinux 2.2 installation, a generic VGA/SVGA X server will be installed. After the installation is complete, the X server will be started automatically but the VM screen will be cropped and unuseable. To get a working X server running inside the VM, you should install the VMware Tools for Linux package immediately after installing OpenLinux 2.2.

OpenLinux 2.2 installation steps

1. Use the VMware Configuration Editor to verify the virtual machine's devices are set up as you expect before starting the installation. For example, if you would like networking software to be installed during the OpenLinux 2.2 installation process, be sure the virtual machine's Ethernet adapter is enabled and configured.
2. Before clicking the Power On button, add Shift to the Ctrl-Alt-Esc hot-key sequence for the virtual machine. Enable this option from the Configuration Editor's Hot Key dialog screen.

   NOTE: After OpenLinux 2.2 is installed with VMware for Linux, it will be necessary to use the Ctrl-Alt-Fx key combination to switch to a Linux virtual terminal within the virtual machine (where x represents the virtual terminal number). VMware Tools for Linux cannot be installed unless the Shift option is enabled.

   ALSO NOTE: When the Shift option is enabled, to release the mouse/keyboard from the virtual machine, all four buttons must be pressed at the same time (Shift-Ctrl-Alt-Esc).
3. Insert the OpenLinux 2.2 CD in the CD-ROM drive and click the Power On button. The virtual machine should start booting from the CD and the install process will begin.
4. Follow the installation steps as you would for a real PC, until you get to the selection screens described in the next step.
5. At the Mouse, Keyboard, Video Card, Monitor and Video Mode selection screens, choose the installation program's default settings without making changes.

NOTE: It is not necessary to use the Probe or Test Mode button at the Video Card/Video Mode selection screens; if you use the buttons, ignore the results and click the Next button to continue with the installation.

6. If the virtual machine's Ethernet adapter was enabled, the OpenLinux installation program will auto-load the correct driver. You will need to choose between DHCP and static IP addressing. The hostname parameter must be provided in either case.
7. Finish installing OpenLinux 2.2 as you would on a real PC.

At this point OpenLinux 2.2 will boot and present a login screen. As you may notice, the dialog box will be cropped inside the virtual machine window. Follow the steps below to install VMware Tools for Linux inside the virtual machine.

VMware Tools for Linux installation steps

1. From the OpenLinux 2.2 login screen, press the Ctrl-Alt-F2 keys together to switch the virtual machine to OpenLinux's second vritual terminal.
2. Log in as root and copy the VMware Tools for Linux package (tools-for-linux.tar.gz) into a directory on the virtual machine.
3. Unpack the software by running tar zxf tools-for-linux.tar.gz , then switch into the vmware-linux-tools directory.
4. Install the VMware X server by running ./install.pl . If you prefer some variation on the standard installation, look at what the Perl script does and install the files by hand.
5. Since the X server installed by OpenLinux will be running at this time, startx will not work. To start the VMware X server, either kill the running X process and then run startx , or do a shutdown -r now (in which case Linux will reboot and automatically start the new X server).
6. Once the VMware-provided X server starts, run the vmware-linux-toolbox program to take advantage of virtual machine mouse and device enhancements that make interacting with the host and other virtual machines more convenient.

INSTALLING OTHER GUEST OPERATING SYSTEMS

You will find detailed installation instructions on VMware's Web site for additional guest operating systems. Please visit http://www.vmware.com/support/guestnotes.html .

THE VMWARE TOOLS

VMware installs unmodified operating systems directly from floppies and/or CD-ROMs. This installation process is the first and only necessary step in building a virtual machine. VMware, however, highly recommends that you install the VMware Tools suite within each virtual machine as soon as it is installed. The suite consists of two components:

- A graphics driver that is optimized for the VMware virtual graphics card. For Linux guest operating systems, this consists of our own version of the XFree86 X server.The graphics driver is installed by the wizard. Once installed, it will be used for subsequent reboots.
- A small background application that allows you to change configuration settings within the virtual machine. The VMware application controls the cursor settings and the connection state of removable devices (such as floppy and CD-ROM).The cursor settings allow users to smoothly move the mouse cursor

between the virtual machine and the host graphical user interface. The cursor settings also allow users to copy and paste text buffers between virtual machines and the host. Note that these enhancements are available only when the application is running.

An installation wizard will automatically install all the VMware software that needs to run within a virtual machine.

The VMware Tools suite is specific to each guest operating system. It is currently available only for specific guest operating systems.

The installation files are on the distribution CD-ROM. You can also download them from the VMware Web site at http://www.vmware.com/download/downloadtools.html.

If your virtual machine is already connected to the Internet, you can download the VMware Tools directly from within the virtual machine.

If your virtual machine is not connected to the Internet but you can access the Internet from the host machine (or any PC), follow the steps below to transfer the VMware Tools package into the virtual machine:

1. From a PC that has access to the internet, download the appropriate VMware Tools suite.
2. Format a floppy disk so it can be read by the guest OS inside the virtual machine.
3. Copy the downloaded VMware Tools package to the disk you prepared in the previous step.
4. Insert the disk in the host machine's floppy drive (if it's not already there) and copy the VMware Tools package to the virtual machine's disk.

NOTE: If you are downloading the tools from the host OS, be sure the guest OS does not have the floppy allocated. You can verify this by selecting Settings > Removable Device in the guest OS menu; be sure the floppy is not checked. When the guest OS needs to access the floppy, be sure to reattach the floppy by selecting Settings > Removable Device and ensure the floppy is checked.

HINT: If you attempt to access the floppy from the host OS and that results in "A:\ is not accessible. The parameter is incorrect," this may be a sign that a guest OS has the floppy allocated.

NOTE: Once the VMware Tools suite is downloaded to the virtual machine's disk, complete the installation by skipping to the appropriate section below:

Installing VMware Tools for Windows NT and Windows 2000

Go to the directory where you downloaded tools-for-windows.exe , and double-click to install VMwareTools. This gives you the option of installing VMware Toolbox and the VMware SVGA driver.

The VMware Tools will automatically create a directory path under Program Files called \VMware\Drivers . After the installation, change the display type by selecting Have Disk ..., and typing this path into the browse box. Change to the desired color depth for the new SVGA driver. Then reboot Windows NT. At this point you can change the display resolution to the desired size.

VMware Toolbox is started when you log on to Windows NT and is accessible as an item in the system tray. Right-click on the VMware logo to see the menu. The VMware Toolbox lets you control mouse behavior and removable devices from within the VM.

The VMware SVGA driver gives you improved graphics performance and resolution in the virtual machine.

If you've tried to upgrade the SVGA driver under a Windows 2000 virtual machine and had problems, and if you've already installed the VMware Tools (including the SVGA display driver), follow these steps to upgrade the Windows 2000 SVGA drivers after installing the current version of VMware Tools.

1. Double-click the Display icon in the control panel. This will start the Display Properties control panel.
2. Click the Settings tab.
3. Click the button labeled Advanced...
4. Click the Adapter tab.
5. Click the button labeled Properties.
6. Click the Driver tab.
7. Click the button labeled "Update Driver..." This will start the Upgrade Device Driver wizard. Click Next.
8. Click the option labeled "Display a list of the known drivers for this device so that I can choose a specific driver." Click Next.
9. Click the button labeled Have Disk.
10. Browse to the drivers subdirectory under your VMware Tools directory. For example:
    C:\Program Files\VMware\Drivers
11. Select the "VMware, Inc. svga" display adapter and click Next. Click Next again. Click the Yes button to indicate that you want to continue with the installation.
12. Click the Finish button.
13. Click Yes to restart Windows 2000 and start using the new video driver.

Installing VMware Tools for Windows 95 and 98

Go to the directory where you downloaded tools-for-windwos.exe , and double-click to install VMware Tools. This gives you the option of installing VMware Toolbox and the VMware SVGA driver.

After the installation, change to the desired color depth for the new SVGA driver. Then reboot Windows 98. At this point you can change the display resolution to the desired size.

VMware Toolbox is started when you log on to Windows 98 and is accessible as an item in the system tray. Right-click on the VMware logo to see the menu. The VMware Toolbox lets you control mouse behavior and removable devices from within the virtual machine.

The VMware SVGA driver gives you improved graphics performance in the virtual machine.

Installing VMware Tools for Linux

The program vmware-toolbox is installed at the same time you install the VMware X server from the file tools-for-linux.tar.gz . To run the toolbox, type vmware-toolbox . For more information, see the section (above) on installing your Linux distribution as a guest OS.

Using the XFree86 X Servers

VMware for Linux should run with generic X servers, but technical support is available only for customers running XFree86 Project X servers. The XFree86 servers will provide significant improvements in full-screen graphics mode. VMware highly recommends customers install XFree86 version 3.3.4 or later. For information on versions of XFree86 servers packaged with your particular Linux distribution, see your distribution vendor's Web site.

The X servers starting with XFree86 version 3.3.4 contain VMware extensions to Direct Graphics Access (DGA) 1.1 that allow virtual machines to take advantage of hardware-based 2-D graphics acceleration. To utilize DGA, the guest operating system must also be running the VMware SVGA graphics driver included in VMware Tools for Windows or Linux (depending on the guest operating system).

VMware plans to support DGA 2.0 when it becomes available.

The latest XFree86 X servers are available at http://www.xfree86.org/.

If you are unable to use the XFree86 X servers, older versions of the X servers, modified by VMware, are provided on the VMware Web site at http://www.vmware.com/download/downloadxserver_pre.html.

4

VMware Technical Notes

VMware Technical Notes

VMware Disk Modes - Persistent, Nonpersistent and Undoable

In the VMware configuration menu, disks can be configured in one of three modes: persistent, nonpersistent and undoable.

Persistent

Persistent mode is the simplest disk mode that VMware supports. Persistent disks behave like conventional disk drives on your computer. All writes done to a persistent disk are written out permanently to the disk. This happens for disks being emulated within a file on the Linux host operating system or for disks that are actual raw disk devices.

Nonpersistent

Any changes made to nonpersistent disks during a VMware session are lost after that session is powered down. Nonpersistent disks are convenient for people who always want to start with a virtual machine in the same state.

Example uses include providing known environments for software test and technical support users as well as doing demonstrations of software. A nonpersistent disk can be either a file on the Windows host operating system or an actual raw disk device. VMware only reads the original disk; any writes to the disk during the session are saved to a log file that is deleted at the end of the session.

During the VMware session any blocks that have been modified and written to the log file are read from there instead of the disk. At the end of the VMware session the log file is discarded. The guest operating system is entirely unaware that the disk is nonpersistent.

Normal guest operating system file buffering works on top of this mechanism, providing efficient buffered I/O. Some disk operations may even be faster to a VMware nonpersistent disk than to an actual disk.

The log file for a nonpersistent disk is implemented by opening and then unlinking a file in the same directory as the disk file. This has the advantage that if a VMware session fails or something else goes wrong, the file is automatically deleted; however, it has the disadvantage that the filesystem can fill up without it being obvious what is consuming the space. The log file is placed in the same directory as the disk file, by default. However, the location of the log file can be changed in the configuration GUI.

Undoable

Undoable mode is similar to nonpersistent disk mode in that writes to the disk are logged elsewhere in a redo-log file. The difference is that with the disk configured to undoable mode, the user has the option later of permanently applying the changes saved in the redo-log file to the disk. One thing that makes undoable mode disks especially convenient is that errors or problems can be undone by simply discarding the redo log.

Examples of uses for undoable disks include installing software or doing administration tasks that may need to be undone if there are problems.

While the VMware session is running, disk blocks that have been modified and written to the redo log are read from there instead of the main disk. Undoable disks can be either disks emulated within a file on the Linux file system or actual raw disk devices.

When you power off of a VMware session with an undoable disk, you are given the option of either committing the modifications in the redo log to the disk, discarding the changes or saving

the redo log. If you choose to save the redo log, you are prompted again when you power on the next VMware session to see if you want to commit the redo-log changes from the previous session, discard the redo log, continue appending changes to the redo log or cancel the power on. You could also manually remove the redo-log file between sessions, although this is not recommended.

The guest operating system is entirely unaware that the disk is undoable. Normal guest operating system file buffering works on top of the redo-log mechanism, proving efficient buffered I/O. Some disk operations, including file system creation, repair and other operations, may even be faster on a VMware undoable disk than on an actual disk.

For disks being emulated within files on a Linux filesystem, the redo log for a disk called filename is created by default in the same directory and is called filename.REDO. This location can be changed using the Configuration Editor. When you are running a VMware guest operating system from raw devices, VMware recommends creating a symbolic link to the actual master disk device. The redo log for the undoable disk is then created with the name symlink.REDO in the same directory as the symbolic link. This way you can control the location of the redo-log file.

MEMORY USAGE NOTES

VMware allows users to set the memory size of each virtual machine and the amount of physical host memory that is reserved for virtual machines. By adjusting the memory sizes of each virtual machine and the amount of reserved memory, users can affect both virtual machine and overall system performance. In this note we describe how VMware uses the memory configuration parameters to properly manage virtual machine and reserved memory.

Virtual machine memory size

The first configuration parameter that users can set is the size of the virtual machine's physical memory. This configuration parameter can be set via the Configuration Editor. The minimum size of the memory for the virtual machine should be set based on the recommendations of the operating system provider. The Configuration Wizard sets what VMware belives are reasonable defaults for the memory size of a virtual machine based on the type of the guest OS. The actual size that should be given to a virtual machine depends on a few parameters:

- What kinds of applications are to be run in the virtual machine
- What other virtual machines will be contending with this virtual machine for memory resources
- What other applications are going to be running on the host at the same time as the virtual machine

Linux does not behave well when it runs low on free memory. For this reason users should not run virtual machines whose memory requirements exceed those of the host and other applications.
To help guard against virtual machines causing the host to thrash, VMware enforces a limit on the total amount of memory that may be consumed by virtual machines: The sum of the memories of all currently running virtual machines cannot exceed the amount of physical memory on the host minus some memory that must be kept available for the host.

Some memory must be kept available on the host to ensure that the host will be able to operate properly while virtual machines are running. The amount of memory reserved for the host depends on the host OS and the amount of memory installed on the host computer.

Reserved memory

The second configuration parameter that users can set is the amount of memory reserved for all running virtual machines. This parameter can be set via a slider in the Total Reserved Memory panel under the Settings menu.

In general, the memory that a virtual machine uses comes out of the same pool of memory used by all other applications and the host. However, in order to improve virtual machine performance, VMware will reserve up to a certain amount of memory for virtual machines. The user can set this limit. When VMware is using this memory, the memory is not available to other applications or the host. When VMware is not using this memory, it is available to other applications and the host. By reserving memory, VMware can allow virtual machines to execute more efficiently.

The memory used by VMware includes the memory made available to the guest operating systems as well as overhead memory associated with running a virtual machine. The amount of reserved memory used by a virtual machine will depend on the working set of the guest operating system (the working set is how much memory the guest needs to run without experiencing poor performance) plus a portion of the overhead memory. The amount of overhead memory used by a virtual machine depends on several factors and can vary from a few megabytes to over 10 megabytes. The actual amount that comes out of the reserved memory pool is always well under 10 megabytes.

The amount of reserved memory actually used for a particular virtual machine varies dynamically as a virtual machine runs. If multiple virtual machines run simultaneously, they will work together to manage the reserved memory. If all this reserved memory is in use by one or more virtual machines, the Linux host will not be able to use this memory for any other purpose. VMware uses the reserved memory only if it determines that a virtual machine needs reserved memory to have reasonable performance. Even if multiple virtual machines are running at the same time, VMware may be using only a fraction of the reserved memory, thus allowing any unused reserved memory to be used by other applications.

The recommended amount of memory to reserve for all running virtual machines is 50 percent of the host's physical memory. If you determine you want VMware to reserve more or less physical memory, you can change this amount by using the slider in the Total Reserved Memory panel under Settings. Changing the amount of reserved memory is recommended only for advanced users because it can have an adverse impact on host or virtual machine performance. Selecting too much physical memory to reserve can cause the host to thrash, or even hang, if other applications are run on the host. Selecting too little physical memory to reserve can cause virtual machines to perform very poorly and also limit the number of virtual machines that can be run.

VMware limits how many virtual machines can run at once based on the amount of reserved memory. This is done to prevent virtual machines from causing each other to perform very poorly. If a virtual machine is powered on and there is not enough reserved memory available, then the power-on will fail.

NETWORKING DOCUMENTATION

Overview

Each virtual machine can have its own distinct network configuration. There are four choices for configuring networking:

- No networking
- Host-only networking
- Bridged networking
- Custom networking

No networking simply means a virtual machine is run in isolation; it will not be able to communicate with the host operating system or any other virtual machine running on the host. This option is useful if you want complete isolation for testing or security purposes. To set up your virtual machine in this way, simply do not install a network interface adapter when you configure the virtual machine.

Host-only networking means a virtual machine can communicate with the host operating system and any other virtual machines set up to use host-only networking, but the virtual machine cannot communicate with any systems off the host machine without the use of a proxy server. This facility is most useful when the host is itself isolated or when you want to isolate your virtual machines from systems outside the host computer. This configuration is analogous to the way corporations connect their internal networks to the Internet with a firewall and proxy services. To set up a virtual machine in this way, you need to install a network interface adapter and mark it Host-only. Once the guest operating system is installed, you may need to do some additional configuration, described below.

Bridged networking means a virtual machine runs on a virtual network that is "connected" to an existing physical network. This permits a virtual machine to appear as a full-fledged host on an existing physical network.

A bridged virtual machine may transparently use any of the services available on the network that it is bridged to: printers, file servers, gateways, etc. Likewise, when a virtual machine is bridged, any physical host – or other virtual machine configured with bridged networking – can use resources on that virtual machine. This is the most commonly used networking configuration. To configure bridged networking manually you need to install a network interface adapter and mark it Bridged. Once the guest operating system is installed, you may need to do some additional configuration, described below.

Custom networking refers to any network configuration other than those described above. For example, a collection of virtual machines, possibly on multiple physical hosts, might be configured on a private virtual network. This might be done to set up a private file-sharing environment, or for testing a group of virtual machines in an isolated network environment. Configuration of custom networking requires a thorough understanding of networking concepts and potentially the implementation of some simple user-level applications. Setting up custom networking is not detailed in this document. If you want to set up your own custom network environment and have trouble doing so, please file an incident at http://www.vmware.com/incident.

What you will see on the host

VMware networking support is done on the host machine through a virtual network device driver that implements four network interfaces: vmnet0, vmnet1, vmnet2, and vmnet3. Each interface is associated with a virtual Ethernet hub through which any number of virtual machines and the host may communicate. By convention vmnet0 is used for bridged networking, vmnet1 is used for host-only networking, and the other two interfaces are available for custom network configurations. In addition to the network interfaces there are two applications: vmnet-bridge and vmnet-dhcpd. The vmnet-bridge application is used by the bridged networking support to effect transparent communication between vmnet0 and another network interface, typically eth0. vmnet-dhcpd is an optional process that runs only when host-only networking is configured; it implements the DHCP protocol for virtual machines running on vmnet1.

What you will see on the guest operating system

Network support on the guest operating system appears through the virtual Ethernet adapter(s) that are configured for the virtual machine. Each device appears to the operating system as an AMD PCNET PCI adapter. Most operating systems will recognize this virtual hardware and automatically configure use of the appropriate device driver. The main issue in completing network configuration in the guest operating system is assigning a network address for the virtual machine.

More details about host-only networking

Setting up host-only networking on the guest operating system

Host-only networking means a virtual machine can communicate with the host operating system and any other virtual machines set up to use host-only networking, but the virtual machine cannot communicate with any systems off the host machine without the use of a proxy server. This is done by creating a private virtual network on which the host and all host-only configured virtual machines reside. Typically all the parties on this private network use the TCP/IP protocol suite, although there is no requirement for this.

Regardless of the communication protocols used, each virtual machine and the host must be assigned addresses on the private network. This can be done „statically‰ (that is, by consulting a fixed database) or "dynamically" using protocols such as the Dynamic Host Configuration Protocol (DHCP).

When host-only networking is enabled at the time VMware is installed, a custom DHCP server application is set up to run on the host machine. This server implements the DHCP protocol only for virtual machines running on the host-only network associated with the virtual network interface vmnet1. Guest operating systems that are set up to use DHCP at boot time to obtain an IP address will then work without any additional configuration (but note the discussion on setting up names in the following section). Guest operating systems that do not use DHCP to obtain an IP address must be setup with a static IP address.

Selecting IP addresses for virtual machines on a host-only network

You have two choices for setting up IP addresses for virtual machines on a host-only network: dynamic assignment using DHCP or static assignment. Using DHCP to assign IP addresses is simpler and more automatic than statically assigning them. Most Windows operating systems, for example, come preconfigured to use DHCP at boot time so they will work the first time they are booted, without additional configuration. (Using this option requires you to have a DHCP server installed on your host.) If, however, you want your virtual machines to communicate with each other using names instead of IP addresses, then you need to set up a naming convention and/or a name server on the host machine. In this case it may be simpler to use static IP addresses.

VMware recommends that if you have virtual machines you intend to use frequently or for extended periods of time, you should assign them static IP addresses or configure the host-only DHCP server to assign the same IP address to the virtual machine every time. For virtual machines that you do not expect to keep for long, use DHCP and let it allocate an IP address. A useful convention is to split up the available IP addresses on a host-only network (VMware suggests that you use a Class C address for host-only networks):

Range	Address Use	Example
<net>.1	host machine	192.168.0.1
<net>.2-<net>.127	static addresses	192.168.0.1-192.168.0.127
<net>.128-<net>.254	DHCP-assigned	192.168.0.128-192.169.0.254
<net>.255	Broadcasting	192.168.0.255

(where <net> is the network number assigned to your host-only network)

Avoiding IP packet leakage in a host-only network

Each host-only network is intended to be confined to the host machine on which it is set up. That is, no packets sent by virtual machines on this network should "leak out" to a physical network attached to the host. Packet leakage can occur only if a machine actively forwards packets. Note that this can be true of the host machine or any virtual machine running on the host-only network.

Systems that support the TCP/IP protocols are usually capable of forwarding IP packets they receive but which are not addressed to them. By default, however, these systems come with IP packet forwarding disabled. If you find packets leaking out of a host-only network, check to see if forwarding has been mistakenly enabled on the host machine, and if it is enabled, disable it. For Linux systems, this is done by writing a "0" to the special file /proc/sys/net/ipv4/ip_forward:

```
machinename# echo 0 >/proc/sys/net/ipv4/ip_forwar    d
```

NOTE: This must be done as the super-user.

For other systems there is a system configuration option that can be set somehow: through a control panel, at compile time, or possibly at boot time; consult your system documentation.

If the host has multiple network adapters then it is likely intentionally configured to do IP forwarding and you do not want to disable it. In this case the only way to avoid packet leakage is to enable a "packet filtering" facility and specify that packets from the host-only network should not be sent off-machine. An explanation of how to do this is beyond the scope of this document; consult your system documentation.

Finally, be aware that virtual machines may leak packets as well. For example, if you use Dial-Up Networking support on a virtual machine, then if packet forwarding is enabled, host-only network traffic may leak out through the dial-up connection.

Controlling routing information for a host-only network

A host-only network is a full-fledged network. It has a network interface associated with it (vmnet1) that is marked "up" at the time the host operating system is booted. Consequently routing server processes that operate on the host operating system, such as routed and gated, will automatically discover it and propagate information on how to reach it unless you explicitly configure them not to. If either of these programs is being run only to receive routing information, the easiest solution is to run them with a -q option so that they do not supply routing information, only receive it. If, however, they are running because they are to supply routing information, you need to configure them to not advertise routes to the host-only network.

Unfortunately, the version of routed that comes with many versions of Linux has no support for specifying an interface should not be advertised. Consult the routed(8) manual page for your system in case you have a more contemporary version of the software.

For gated, configuration is involved. You need to explicitly exclude the vmnet1 interface from any protocol activity. If you need to run virtual machines on a host-only network on a multi-homed system where gated is used and have problems doing so, please file an incident at http://www.vmware.com/incident.

Setting up Samba for virtual machines on a host-only network

Samba is a free software package that implements the Microsoft network file system protocols for UNIX®. It is installed by default on most Linux systems. Samba is very useful if you have virtual machines that run Microsoft operating systems because it allows you to share files among all your operating systems. The Samba server process normally operates on only one network interface. If you enable host-only networking on the host machine, Samba will not, by default, serve requests from virtual machines connected to this network. You must explicitly configure the Samba server, smbd, to listen on both the primary network interface (or all network interfaces that exist) and vmnet1, where the host-only network is present. To do this you edit the smb.conf file that is usually located at /etc/smb.conf. Add an "interfaces" line of the form:

```
interfaces = <physical networks> <host-only-net>.1/24
```

where <physical networks> is a list of the physical networks that are to be served and <host-only-net> is the network number that was assigned to the host-only network. For example, a typical configuration for a machine is to be connected to a network. Say this network is numbered 209.220.166 with a netmask of 255.255.255.0 (that is, it is a Class C network) and the

host address is 209.220.166.34, then the interfaces line would read

interfaces = 209.220.166.34/24 192.168.0.1/24

if the host-only network is assigned 192.168.0.0. If you are not familiar with the /24 notation you can also give the netmask directly:

interfaces = 209.220.166.34/255.255.255.0 192.168.0.1/255.255.255.0

Consult the smb.conf(5) manual page for more information.

Other potential issues with host-only networking

The following are common questions and issues that may arise when you are configuring a host-only network.

Q: DHCPD on the host machine does not work after I installed VMware.

A: If you were running the DHCP server program dhcpd on your machine before installing VMware, it probably was configured to respond to DHCP requests from clients on any network interface present on the machine. When host-only networking is configured, an additional network interface, vmnet1, is marked "up" and available for use, and dhcpd may notice this. In this case some dhcpd implementations abort if their configuration files do not include a "subnet" specification for the interface, even if dhcpd is not to respond to messages that arrive through the interface. The best solution to this problem is to add a line to the dhcpd configuration file of the form:

subnet <net>.0 netmask 255.255.255.0 {}

(where <net> is the network number assigned to your host-only network; for example, 192.168.0). This informs dhcpd about the host-only network and tells it explicitly not to respond to any DHCP requests it sees coming from it.

An alternative solution is to explicitly state the set of network interfaces you want dhcpd to listen to each time you start the program. For example, if your machine has one Ethernet interface, eth0, then each time you start dhcpd you would list it on the command line: machinename# dhcpd eth0 rather than have it probe for all available network interfaces. If the above solutions do not work for your DHCP server program, it likely is old. You can try upgrading to a more current version such the Version 2 DHCP software available from the ISC (see http://www.isc.org).

Q: Is there any way to use DHCP and Dynamic Domain Name Service (DDNS) on a host-only network?

A: DHCP can be used to hand out IP addresses as well as other information such as the identity of a host running a name server and the nearest router or gateway. But it does not currently provide a means to dynamically establish a relationship between the IP address it assigns and a client's "name" (that is, to update a DNS server using DDNS). This facility is scheduled to be part of the Version 3 DHCP server available from the Internet Software Consortium (ISC). When that is available we will update our software to use that server.

In the meantime, this means that if you want to use names to communicate with other virtual machines you will need to either edit the DHCP configuration file for vmnet1 (/etc/vmware/vmnet1.conf) or use IP addresses that are statically bound to a host name. Editing the DHCP server configuration file requires information that is best obtained directly from the DHCP server documentation; consult the UNIX manual pages dhcpd(8) and dhcpd.conf(8).

More details about bridged networking

Setting up bridged networking on the guest operating system

Bridged networking means a virtual machine appears just like any other host on the physical

network. You need to configure operating system support for the virtual Ethernet adapter, then either assign a fixed network address or enable use of DHCP for dynamic address assignment. Assigning a network address is done according to local conventions. If your site runs DHCP, you may choose to enable DHCP. Otherwise you will need to consult a network administrator to obtain a network address. Be aware that if the host machine is set up to boot multiple operating systems and you run one or more of them in virtual machines, you will need to configure each operating system with a unique network address. Many people assign all systems the same address because they assume only one will be running at a time.

Changing the MAC address of a virtual machine

When a virtual machine is powered on, VMware software automatically assigns it a MAC address. The software guarantees that virtual machines will be assigned unique MAC addresses within a given host system. However, the software does not guarantee that a given virtual machine will be assigned the same MAC address every time it is powered on. In addition, VMware does its best, but cannot guarantee, to automatically assign unique MAC addresses for virtual machines running across multiple host systems.

If you want to guarantee that the same MAC address is assigned to a given virtual machine every time, or want to guarantee a unique MAC address for each virtual machine within a networked environment, you can assign it manually instead of allowing VMware to assign it automatically. It is possible to manually assign the same, unique MAC address to any virtual machine by adding the following line to its configuration file:

ethernet0.address = 00:50:56:XX:YY:ZZ

where 'XX' is a valid hex number between 00h and 3Fh, and 'YY' and 'ZZ' are valid hex numbers between 00h and FFh. Because VMware virtual machines do not support arbitrary MAC addresses, the above format must be used.

NOTE: as long as you choose 'XX:YY:ZZ' unique among your hard-coded addresses (where 'XX' is a valid hex number between 00h and 3Fh, and 'YY' and 'ZZ' are valid hex numbers between 00h and FFh), conflicts between the automatically assigned MAC addresses and the manually assigned ones should never occur.

Configuring Dual/Multiboot Systems to Run With VMware

It is possible to install VMware for Linux to run one or more guest operating systems from raw disk partitions. This is useful for people who already have a dual- or multiboot system and who want to run those operating systems under VMware. If you don't have a dual-boot system, but want to install a guest operating system into an unused partition through a virtual machine please refer to the technical note Installing an OS onto a Raw Partition from a Virtual Machine Using VMware. For some applications raw disk may also provide better disk I/O performance than using a virtual disk. Just as with virtual disks, VMware enables a raw disk to be used in persistent, undoable, and nonpersistent modes. See the explanation of these modes in the technical note VMware Disk Modes.

> **Caution:** Raw disk support is an advanced feature of VMware software and should be enabled only by users who are already familiar with the product. To familiarize yourself, you should, at minimum, create and configure a virtual machine with a virtual disk and install an OS. In addition, booting a previously installed OS within a virtual machine may not work on some existing installations. Your experience may vary depending on your hardware configuration and guest operating system installation.

VMware supports using raw disk partitions only on IDE drives. Booting guest operating systems on raw SCSI drives is not currently supported. However, if a virtual machine is configured with a virtual disk, instead of a raw disk partition, then its disk (file) can be stored on the Linux filesystem, regardless of whether the underlying drive(s) containing the file system are IDE or SCSI.

VMware uses meta configuration files to access each raw IDE device on the system. These meta files contain access privilege information that controls a virtual machine's access to certain partitions on the disks. This mechanism is used to prevent users from accidentally trying to run the host operating system again as a guest, or another guest operating system that the virtual machine was not configured for. The meta configuration file also prevents accidental writes to raw disk partitions from badly behaved operating systems or applications.

The VMware Configuration Wizard with the -rawdisk command line option is used to configure VMware to use existing raw disk partitions. The Wizard will step you though creating a configuration for a new virtual machine including configuring the raw disk meta configuration files. The Wizard is typically rerun to create a separate configuration for each guest operating system installed on a raw partition.

VMware works with existing boot managers installed on the computer system. The boot manager will run inside VMware and present the user with the choice of guest operating systems to run. The user has to manually choose the guest operating system that this configuration was intended to run.

Use the following steps to run a guest operating system from a raw device:

Create a separate configuration for each guest operating system. Allow read/write access to the partitions used by that operating system only.

1. Before starting, if you are running a Windows or Windows NT guest operating system, you should read the technical note Setting Up Hardware Profiles in Virtual Machines. VMware recommends booting the guest operating system natively on the computer and creating a hardware profile for the virtual machine before proceeding.

2. Check operating system partition mounts:
   Be sure the existing raw disk partition(s) that you plan to configure the virtual machine to use are not mounted by Linux.

3. Set the device group membership or device ownership:
   The master raw disk device(s) needs to be readable and writeable by the user who runs VMware. On most distributions, the raw devices (such as /dev/hda, /dev/hdb) belong to group-id named "disk." If this is the case, you can add VMware users to the disk group. Another option is to change the owner of the device. Think carefully about security in exploring different options here. It is typically a good idea to grant VMware users access to all /dev/hd[abcd] raw devices that contain operating systems or boot managers, then rely on VMware's raw disk configuration files to guard access. This helps provide boot managers access to configuration and other files they may need to boot the operating systems. For example lilo needs to read /boot on a Linux partition to boot a non-Linux operating system that may be on another drive.

4. If you will be running a second Linux installation from an existing partition as a guest OS, and your real machine's /etc/lilo.conf has a memory register statement such as Append= "mem=...." You may want to adjust the Append memory parameter or create a new entry in lilo for running Linux in a virtual machine. Many newer Linux distributions recognize all memory in the real machine, whereas many older Linux distributions see only the first 64MB of memory by default. Machines with more than 64MB of memory that run the older distributions may have the Append= "mem=...." parameter added under the Image=....' section of lilo.conf to tell Linux to look for more memory than seen by default. If the amount of memory configured in lilo.conf exceeds the amount of memory assigned to the virtual machine, then when the virtual machine tries to boot the second Linux installation, the guest OS will most likely panic. You can simply create another entry in lilo.conf for running Linux in a virtual machine by specifying a different amount of memory than what should be normally recognized when Linux boots directly on the real machine.

5. Run the VMware Configuration Wizard with the rawdisk option:
   Run the VMware Configuration Wizard program and specify the raw disk option. Type vmware-wizard -rawdisk . The default location for vmware-wizar d is /usr/local/bin .

6. In the Configuration Wizard
   - Read the introductory text.
   - Click Next.
   - Check the box for the operating system you have on raw disk. This is used to pick some simple defaults for VMware configuration settings and give default names to configuration files. The default settings can be changed later with the Configuration Editor.
   - Click Next.
   - Specify the virtual machine directory This is where the configuration files are placed. Leaving the default is often a good choice.
   - Click Next.
   - Chose Existing Partition from Virtual Disk Type Settings.
   - Click Next.
   - Select the read/write option for the disk partition(s) that contain the guest operating system being configured.

   NOTE: Corruption is possible if you allow the virtual machine to modify a partition that is simultaneously mounted under Linux. Since the virtual machine and guest OS will be accessing an existing partition while the host continues to run Linux, it is critical that the virtual machine not be allowed to modify any partition mounted under Linux or in use by another virtual machine. To safeguard against this problem, be sure the partition you mark read/write for the virtual machine is not mounted under the Linux host.
   - You need to leave the MBR at least read only.
   - Leaving the other partitions read only is recommended. The lilo boot manager will often have to read files from /boot (on a Linux partition) to boot a guest operating system.
   - Click Next.
   - Configure CD-ROM Device Settings as required.
   - Click Next.
   - Configure Floppy Device Settings as required.
   - Click Next.
   - Configure Networking Settings as required.
   - Click Next.
   - Look through the Confirmation page to see what is about to be done.
   - Click Back if you need to change settings. Note the location of the configuration file (.cfg). This is the file you will need to specify when you run VMware.
   - Click Done.

7. Start VMware and check the configuration
   - Type vmware <config-file> . Where <config-file> is the path of the configuration file created by the Wizard. These files end in .cfg.
   - Open Settings > Configuration Editor and check that your IDE configuration specifies at least one raw disk meta configuration file, these files are named <configuration-name>.hda, <configuration-name>.hdb , etc.
   - Also modify any configuration options you want to change from the Wizard default. For example, you may change the amount of memory allocated to the guest operating system.

8. If you have multiple IDE drives configured on a system, the VMware BIOS will normally attempt to boot them in this sequence:
    - Primary Master
    - Primary Slave
    - Secondary Master
    - Secondary Slave

    If the system has multiple bootable IDE drives and you want to boot from an arbitrary drive in this list, the boot sequence can be changed in the virtual machine's Phoenix BIOS boot menu. After powering on VMware, press F2 during the BIOS boot in VMware to enter the BIOS setup menu.

9. Power on the virtual machine.
    - Click the Power On button. VMware should start, run the Phoenix BIOS, then boot from the master boot record (MBR).
    - Choose the target operating system from the list of options offered by the boot manager.

10. Remember that your virtual machine hardware environment that the guest operating system is about to run in for the first time probably differs significantly from the real hardware of your machine. For Windows guest operating systems, Plug and Play will start reconfiguring Windows. For Windows and Windows NT systems, you should set up your virtual hardware profile. See the technical note Setting Up Hardware Profiles in Virtual Machines for more information.

Warning: If you designate your safe raw disk as an undoable disk, you will need to either commit or discard the changes to the undoable disk before you reboot your guest OS natively. This is because any changes to sectors on the real disk that have been modified on the undoable disk will invalidate the redo file corresponding to the undoable disk. See the technical note VMware Disk Modes for more information on undoable disks and their corresponding redo files.

NOTE: It is possible to specify using a raw device directly in the VMware Configuration Editor. This is potentially hazardous, and VMware does not support users doing this. In future releases VMware will detect this situation and refuse to power on.

Installing an OS onto a Raw Partition from a Virtual Machine Using VMware

In some situations raw disks may provide better disk I/O performance than using a virtual disk. Hence, it is possible to install an operating system onto an extra, unused disk (or partition in some cases) from a virtual machine and access it while the Linux host runs simultaneously. However, with this configuration, the real computer will may not be able to boot the freshly installed operating system; the new OS will most likely only be accessible from a virtual machine. If you have a dual-boot system and want to configure a virtual machine to boot from an existing partition, please see the technical note Configuring Dual/Multiboot Systems to Run With VMware. Just as with virtual disks, VMware enables a raw disk to be used in persistent, undoable, or nonpersistent modes; see the explanation of these modes in the technical note VMware Disk Modes.

NOTE: VMware virtual machines support raw disk partitions only on IDE drives. Booting guest operating systems from raw SCSI drives is not currently supported. Of course, a virtual machine configured with a virtual disk can have its disk file stored on the Linux file system, regardless of whether the underlying drive(s) containing the file system are IDE or SCSI.

VMware uses meta configuration files to access each raw IDE disk on the host system. These meta files contain access privilege information that controls a virtual machine's access to certain

partitions on the disks. This mechanism is used to prevent users from accidentally trying to run the host operating system again as a guest, or another guest operating system that the virtual machine was not configured for. The meta configuration file also prevents accidental writes to raw disk partitions from badly behaved operating systems or applications.

NOTE: It is possible to specify using a raw device directly in the VMware Configuration Editor. This is potentially hazardous and VMware does not support users doing this. In future releases VMware will detect this situation and refuse to power on.

The Configuration Wizard, used with the -rawdisk command line option, will create the raw disk meta files and set up the raw disks to appear in the virtual machine on the same controller/channel as the real machine. Matching the IDE controller/channel assignments of the host system inside a virtual machine is essential for users who have a dual-boot system and want the virtual machine to boot a previously installed OS from an existing partition. However, to actually install a new OS onto an unused disk, you must create a virtual machine with the Configuration Wizard and then additionally configure it through the Configuration Editor to recognize that the extra, unused raw disk is attached to its primary master IDE channel. If the latter step is not done, the guest operating system installation program will not be able to install its files nor will the virtual machine be able to boot from the raw disk.

Use the following steps to prepare a virtual machine to install a new OS onto an unused raw disk:

CAUTION: Raw disk support is an advanced feature of VMware software and should be enabled only by users who are already familiar with the product. To familiarize yourself, you should at a minimum create and configure a virtual machine with a virtual disk, then install and use an OS from the virtual machine.

1. Identify the raw partition where the guest operating system will be installed:
   Check the guest operating system documentation regarding the type of partition on which the OS can be installed. For example, operating systems like DOS, Windows 95 and Windows 98 must be installed on the first primary partition while others, including Linux, can be installed to a primary or extended partition on any part of the drive.

   Depending on the guest operating system you will run, identify an appropriate raw partition or disk for it to use. Be sure that the raw partition is not mounted by the Linux host and not in use by others. Also, be sure the raw partition or disk does not have data you will need in the future; if it does, back it up.

2. Set the device group membership or device ownership:
   The master raw disk device(s) needs to be readable and writeable by the user who runs VMware. On most distributions, the raw devices (for example, /dev/hda, /dev/hdb) belong to the group-id named “disk.” If this is the case, you can add VMware users to the disk group or change the owner of the device.

3. Launch the VMware Configuration Wizard with the -rawdisk option.
   Type vmware-wizard -rawdisk . The default location for vmware-wizard is /usr/local/bin . If you are not familiar with the VMware Configuration Wizard, see the section of this guide titled Create and Configure Your Virtual Machine.

4. In the Configuration Wizard:
   - Read the introductory text and click Next.
   - Check the box for the operating system you have on raw disk and click Next.
   - Specify the Virtual Machine Directory and click Next.
   - Choose Existing Partition from Virtual Disk Type Settings and click Next.
   - Select the read/write option only for the raw partition/disk (and its MBR)to which you want to install the guest operating system. If the raw disk you plan to use has multiple partitions already on it, be aware that certain OSes (specifically, DOS, Windows 95, Windows 98) must be installed on the first primary partition.

CAUTION: Corruption is possible if you allow the virtual machine to modify a partition that is simultaneously mounted under Linux. Since the virtual machine and guest OS will be accessing a raw disk partition while the host continues to run Linux, it is critical that you not allow the virtual machine to modify any partition mounted under Linux or in use by another virtual machine. To safeguard against this problem, be sure the raw disk partition you mark read/write for the virtual machine is not in use.

Select No Access for the remaining raw partitions and click Next.

- Configure CD-ROM Device Settings as required and click Next.
- Configure Floppy Device Settings as required and click Next.
- Configure Networking Settings as required and click Next.
- Look through the Confirmation page to see what is about to be done. Click Back if you need to change settings. Note the location of the configuration file (.cfg). This is the file you will need to specify when you run VMware. Click Done

5. Start VMware and manually change the controller/channel assignment selected by the wizard:
   - Type vmware <config-file> . Where <config-file> is the path of the configuration (.cfg) file created by the Wizard in the previous step.
   - Open the Configuration Editor from the Settings menu and check that your IDE configuration specifies at least two raw disk meta configuration files. These files are named <configuration-name>.hda , <configuration-name>.hdb , etc.
   - Identify the meta configuration file for the raw disk to which you will install the new guest operating system. For example, if your real machine has an unused disk on the secondary master IDE channel and you want to use this device for the virtual machine, you should see a file called <configuration-name>.hdc next to the virtual machine's IDE 1:0 configuration entry.
   - Replace the meta configuration (.hda) file next to the virtual machine's IDE 0:0 channel with the meta file you identified in the previous step.
   - Remove the other raw disk meta file(s) from the virtual machine's IDE configuration dialog and click OK.
   - Save the changes to the virtual machine's configuration file.

6. At this point you are ready to begin installing the guest operating system into the raw disk you configured for the virtual machine. Installation notes for various guest operating systems are included above.

Setting Up Hardware Profiles in Virtual Machines

If you have a dual-boot system and want to use a virtual machine to boot a previously installed operating system (such as Windows 95, Windows 98, Windows NT or Windows 2000) from an existing partition, you must set up "real" and "virtual" hardware profiles. Certain operating systems use hardware profiles to load the appropriate drivers for a given set of hardware devices. VMware recommends using them only if you are familiar with VMware virtual machines and the Windows hardware profiles concept. Also, if you haven't already done so, look at the technical note Configuring Dual/Multiboot Systems to Run With VMware before proceeding.

Each virtual machine provides a platform that consists of the following set of virtual devices:

- Virtual CD-ROM
- Virtual IDE hard disk drives
- Standard PCI graphics adapter
- Standard floppy disk drive

- Intel 82371 PCI Bus Master IDE Controller (includes primary and secondary IDE controllers)
- Standard 101/102-key keyboard
- PS/2-compatible mouse
- AMD PCNET Family Ethernet Adapter (PCI-ISA)
- Serial ports (COM1-COM4)
- Parallel ports (LPT1-LPT2)
- SoundBlaster 16-compatible sound card

This set of virtual devices is different from the set of real hardware devices and is independent of the underlying hardware with a few exceptions (the processor itself is such an exception). This feature provides a stable platform and allows OS images installed within a virtual machine to be migrated to other virtual machines, regardless of the configuration of the real machine.

If an operating system is installed directly into a VMware virtual machine, the operating system will properly detect all the virtual devices by scanning the hardware. However, if an operating system is already preinstalled on the real machine (for example, in a dual-boot configuration), the operating system will already be configured for the real hardware devices. In order to boot such preinstalled operating systems in a virtual machine, you will need to create separate hardware profiles in order to simplify the boot process.

Microsoft operating systems (including Windows95, Windows98, Windows NT 4.0) have the notion of hardware profiles. Each hardware profile is associated with a set of known devices. If more than one hardware profile exists, the user is prompted to choose between different hardware profiles at boot time.

Windows 95, Windows 98 and Windows 2000 use Plug and Play at boot time to confirm that the actual devices match the chosen hardware profile. Mismatches lead to the automatic detection of new devices. Although this operation succeeds, it can be a fairly slow process.

Windows NT does not have Plug and Play support and uses the hardware profiles to initialize its devices. Mismatches will lead to errors reported by the device drivers and the devices are disabled.

In order to set up hardware profiles for your real and virtual machines, we recommend that you follow these steps:

1. Before running VMware to boot a preinstalled operating system on a disk partition, boot it natively and create two hardware profiles that you can call Real Machine and Virtual Machine. To do this, open Control Panel > System then select the Hardware Profiles tab. Click the Copy button and name the copies appropriately.
2. **NT only:** While still running natively, use the Device Manager to disable some devices from the Virtual Machine hardware profile. To do this, open Control Panel > Devices, then select the individual devices to disable. Things to disable in the Virtual Machine hardware profile include audio, MIDI and joystick devices, Ethernet and other network devices and USB devices. Remember to disable them from the Virtual Machine hardware profile only. Skip this step if you are running Windows 95 or Windows 98. The initial Plug and Play phase will detect device mismatches.
3. Reboot the machine into Linux.
4. Use the VMware Configuration Wizard to configure your virtual machine. See the technical note Configuring Dual/Multiboot Systems to Run With VMware for more information on your partition configuration.
5. Boot the virtual machine and select the guest OS via your existing boot manager. Choose Virtual Machine at the hardware profile menu prompt. You will encounter device failure messages and delays during this initial boot.

6. **Windows 2000 only:** After you log in to Windows 2000 (now running as a guest operating system), you should see a Found New Hardware dialog box for the video controller as Plug and Play runs and discovers the virtual hardware. Do not install drivers at this time; click Cancel to close the Found New Hardware dialog boxes. Do not reboot the virtual machine; instead, click No when you see the System Settings Change/Reboot pop-up message. Windows 2000 will automatically detect and load the driver for the AMD PCnet PCI Ethernet card. At this point, you should install the VMware Tools for Windows add-on package inside the virtual machine. Allow the virtual machine to reboot after the VMware, Inc. SVGA video driver (included in VMware Tools for Windows) has been installed. Once Windows 2000 reboots inside the virtual machine, select a new SVGA resolution from the Display Properies > Settings tab dialog to increase the size of the virtual machine's display window. If you want to enable the virtual machine's sound adapter to work inside the Windows 2000 guest operating system, finish the remaining steps on this page and then refer to the technical note VMware and Sound.

   Windows 95/98 only: You should see New Hardware Detected dialog boxes as Plug and Play runs and discovers the virtual hardware. Windows will prompt you for locations to search for device drivers. Most of the device drivers will be available in the existing OS installation, but you may need the installation CD-ROM for some networking device drivers. Windows will also ask you to reboot your system several times as it installs the device drivers. In some instances, Windows may not recognize drivers during the initial hardware detection. In such cases, you can cancel the installation of the particular device, or try pointing to C:\windows\system\ to search for device drivers on the hard disk. Any failed device installations may be performed at a later time after the CD-ROM drive is recognized. After Windows has installed the virtual hardware and its drivers, you can remove the failed devices corresponding to the real hardware using the Device Manager Control Panel > System > Device Manager tab). Select the device then click the Remove button. If a device appears in multiple hardware profiles, you can select the hardware profile(s) from which to remove the device.

7. **Windows NT only:** After the OS has finished booting, view the event log to see which real devices have failed to start properly. You can disable them from the Virtual Machine hardware profile using the Device Manager (Control Panel > Devices).

8. Confirm that your virtual devices are working properly, specifically the network adapter. **Windows 95/98 only:** If any virtual devices are missing, you can detect them by running Control Panel > Add New Hardware.

9. Install the VMware Tools. (**Windows 2000 only:** You already did this in step 6.) The VMware tools will appear and run in both hardware configurations but will have an effect only in the virtual hardware configuration.

   NOTE: The next time you reboot Windows natively using the Real Machine hardware profile, some virtual devices may appear in the device list. You can disable or remove these virtual devices from Real Machine hardware profile in the same way that you removed real devices from the Virtual Machine hardware profile in steps 6 and 7 above.

VMware and Sound

VMware provides a Creative Technology Sound Blaster® 16 compatible audio device and supports sound in Windows 95, Windows 98, Windows NT, Windows 2000 and Linux guest operating systems. The VMware sound device is disabled by default and must be enabled in the VMware Configuration Editor. Sound support is currently limited to PCM (pulse code modulation) output (that is, any application that produces sound without using MIDI).

Configuring sound with VMware

1. Configure sound on the Linux host operating system. Please refer to your documentation for your particular Linux distribution. You may need to install additional software

packages on your system to support sound. VMware cannot provide support assistance in configuring sound on your host operating system. Please contact your host operating system support provider or PC manufacturer for help.

2. Enable sound within the virtual machine. The sound virtual device is not installed in the virtual machine by default. In the VMware Configuration Editor open the Sound Adapter panel and Click the Install button. Save the configuration and power on the virtual machine.
3. Configure the guest operating system to use the VMware virtual sound device. This device is compatible with a Creative Technology Sound Blaster 16.
4. For Windows 2000 guest operating systems
    - Double-click the Add/Remove icon from the Windows 2000 Control Panel.
    - From the Add/Remove Hardware Wizard dialog, select “Add a new device” and click Next.
    - From the Find New Hardware screen, select “No, I want to select the hardware from a list” and click Next.
    - From the Hardware Type screen, select "Sound, video and game controllers" from the list and click Next.
    - From the Select a Device Driver screen, select "Creative" from the Manufacturers list and select "Sound Blaster 16 or AWE32 or compatible (WDM)" from the Models list, then click Next.
    - From the Start Hardware Installation screen, click Next to install the Sound Blaster 16 drivers.
    - From the Completing the Add/Remove Hardware Wizard screen, click Finish and reboot the virtual machine. Sound should be working the next time the virtual machine boots Windows 2000.
5. For Linux guest operating systems

    Please refer to your documentation for your particular Linux distribution. You may need to install additional software packages on your system to support sound. When configuring the sound, please use the following parameters:

IO PORT	IRQ	8-bit DMA	16-bit DMA
0x220	5	1	7

6. For Windows 95 and Windows 98 guest operating systems

    If you have never installed a Sound Blaster 16 card in this Windows system, you will need a Windows 95 or Windows 98 installation CD-ROM.
    - Launch Add New Hardware from the Windows Control Panel
    - Click Next
    - Select Yes for "Do you want Windows to search for new hardware?"
    - Click Next
    - Click Next again
    - Windows should run the autodetection and say it is ready to finish
    - If prompted, insert the Windows CD-ROM into the drive and click OK
    - Click Finish

    If you have problems with Windows auto detection, add the device manually.
    - Launch Add New Hardware from the Windows Control Panel
    - Click Next
    - Select No for "Do you want Windows to search for new hardware?"
    - Click Next

- Select "Sound, video and games controllers"
- Click Next
- Select "Creative Labs Sound Blaster 16 or AWE-32"
- Click Next
- Click Finish

7. For Windows NT guest operating systems

   If you have never installed a Sound Blaster 16 Card in this Windows NT system you will need a Windows NT 4.0 installation CD-ROM.

   - Launch Multimedia from the Windows NT Control Panel
   - Click the Devices tab
   - Click the Add button
   - Select the "Creative Labs Sound Blaster 1.X, Pro, 16"
   - Click OK
   - Insert the Windows NT 4.0 CD-ROM in the CD-ROM drive when prompted
   - Specify D:\I386 (if D: is your CD-ROM drive; otherwise, substitute the correct drive letter)
   - Click OK
   - Configure the Sound Blaster base I/O Address
     - IO Address: 220
   - Click OK
   - Complete the Sound Blaster 16 Configuration
     - IRQ: 5
     - 8-bit DMA: 1
     - 16-bit DMA: 7
     - MPU-401 I/O Address: Disable (MPU-401 MIDI device is not supported)
   - Click OK
   - When prompted to restart, select Restart Now

VMware for Linux Parallel Port Behavior

This document explains the current behavior and limitations of parallel ports under VMware for Linux.

VMware for Linux supports two types of virtual parallel port devices: the unidirectional ports (SPP) supported in VMware for Linux versions prior to 1.0.8 and a partial emulation of bidirectional PS/2-style ports. Unidirectional ports are supported in all Linux host versions. Bidirectional ports require Linux kernel versions 2.2.5 or later.

Unidirectional ports

Unidirectional ports are supported for backward compatibility. They are used typically to connect to printers or to send the printer output to a file. The speed is usually adequate for printing text, but expect long delays when printing images.

The pathnames of the host devices for unidirectional ports are typically /dev/lp0, /dev/lp1 , etc.

Bidirectional ports

Bi-directional ports are used by a variety of devices, such as printers, scanners, dongles, and disk drives.

Currently, VMware provides only partial emulation of PS/2 hardware. Specifically, interrupts requested by a device connected to the physical port are not passed to the virtual machine. Also, the guest operating system cannot DMA (direct memory access) data to or from the port. This will be resolved in the future.

For this reason, and for the time being, not all devices that attach to the parallel port are guaranteed to work correctly. Below is a partial list of devices, which we will update on the VMware Web site as we gain further information. If you try out a device that is not in the list, we would like to hear about it.

Bidirectional emulation is slower than native access but faster than unidirectional emulation, so this is the recommended mode, when possible, even when the device is unidirectional – printers, for example.

The pathnames of the host device for bidirectional ports are usually /dev/parport0 , /dev/parport1 , etc.

Default configuration

Parallel ports by default are bidirectional on Linux hosts 2.2 or later; on earlier versions they are unidirectional. Their default base addresses are, in order, 0x3bc, 0x378 and 0x278. None of the ports have an assigned IRQ or DMA channel. The ports are not present by default.

Installation on guest operating systems

Most guest operating systems automatically detect the parallel port(s) at installation time and install the required drivers. Some operating systems, including Linux, Windows NT, and Windows 2000, auto detect the port(s) at boot time. Others, such as Windows 9x, do not. On Windows 9x, when a port is changed from unidirectional to bidirectional or vice versa, it is necessary to remove the device driver for that port (from the System icon in the Control Panel), and add a new one. Adding a new driver is also required when a new port is added. In both cases, use the Add New Hardware icon in the Control Panel and let Windows detect the new device(s). Manually selecting the devices from a list may result in an incorrect configuration.

Troubleshooting

If an error message is displayed when you power on the virtual machine stating the parallel port on the host does not have an ECR (extended control register), it is possible the hardware supports it but it has been disabled in the BIOS. In this case, reboot your host, enter the BIOS configuration editor (typically by holding down the Delete key during early execution of the BIOS), find the parallel port field, and enable ECP mode (or another combination of modes that includes ECP). Most modern computers should support ECP mode.

Devices known to work

Adobe dongle	Windows 95 guest
RIO MP3 player	Windows 95 guest
UMAX Astra 1220 P scanner	Windows 95 guest
Hewlett-Packard LaserJet 5MP printer	Windows 9X and Windows NT/2000 guests
Canon Bubble Jet BJ-200e printer	Windows 9X and Windows NT/2000 guests
Iomega ZIP drive	Linux and Windows NT/2000 guests only (see note below)

Devices that probably work

Dongles	Most dongles are likely to work.
Printers	Most printers are likely to work.
HP Deskjet 722C	Reported by customer
CARDport Swift Smart Media Digital ImageReader/Writer from Chase AdvancedTechnologies	Reported by customer

Special notes for the Iomega ZIP drive

The Iomega Zip drive currently works reliably on Linux and NT guests only. On Windows 95 or 98 guests, it intermittently locks up the guest at boot time or during installation. However, there is a workaround. The lockups happen only when using the newer drivers provided directly by Iomega. If, instead, the older drivers available from the Microsoft FTP site are used, the drive works reliably. The drivers are available at:

ftp://ftp.microsoft.com/Softlib/MSLFILES/PPA3.EXE

Using Raw Disks with VMware for Linux

This technical note explains issues involved in using raw disks with VMware for Linux.

Raw disks

A raw disk is a physical IDE disk that your host operating system knows about. For example, with Linux as the host operating system, the Raw Disks are /dev/hda through /dev/hdl.

Safe raw disks

A safe raw disk is a small file that describes how the different parts of a raw disk should be accessed by a virtual machine. For example, here is the content of my_first_safe_raw_disk, a typical safe raw disk for a Windows NT virtual machine running inside VMware for Linux:

```
===
DEVICE /dev/hda
# Partition type: MBR
RDONLY 0 62
# Partition type: HPFS/NTFS
ACCESS 63 8193149
# Partition type: Linux swap
NO_ACCESS 8193150 8466254
===
```

The virtual machine will be able to access the /dev/hda raw disk. The access information for sectors on the Raw Disk is as follows:

Partition Type	Sectors	Access Rights
Linux boot information	0 through 62 inclusive	Read-Only
NTFS or FAT	63 through 8193149 inclusive	Read-Write
Linux swap	8193150 through 8466254 inclusive	Forbidden

If the guest operating system inside the virtual machine attempts a forbidden read or write operation to the safe raw disk, VMware will display a pop-up window asking the user to authorize or deny the access.

Installing a safe raw disk in a virtual machine

To install a safe raw disk in your virtual machine:

1. Determine which raw disk you would like to access from a virtual machine.
2. Open the Configuration Editor in the virtual machine for which you want to create a safe raw disk.
3. Expand the IDE Drives node.
4. Look for "Not installed" drives. If there are none, your virtual machine is already configured with four IDE drives, and you will not be able to configure another one. You must remove one of the configured IDE Drives. To do this, click on the one you want to remove, then click the "Remove" button.
5. Click on a "Not installed" drive. For example, if you choose "P-S Not Installed," it means that your raw disk will be seen by the virtual machine as the slave IDE device of the primary IDE controller.
6. Set the Device Type field to Raw Disk.
7. Set the Name field to the name of your safe raw disk (for example, my_first_safe_raw_disk).
8. Set the Device to your raw disk device (for example, /dev/hda).
9. Click the Edit Raw Disk... button. A new window appears, displaying a list of partitions present on your raw disk.
10. For each partition, select the access rights the virtual machine will have. At this point, you need to know what access rights your guest operating system will need. Here are a few tips to help you with this step:

    The following options are possible:

    - No Access - The virtual machine will not be able to read or write on the partition at all. Use it only if you suspect that your guest operating system is buggy, and if you want to track random off-range read accesses.
    - Read/Write - The virtual machine will be able to read from and write to the partition. Use it for partitions that your guest operating system knows about natively.
    - Read-Only - The virtual machine will only be able to read data from the partition. Use it everywhere else.
11. Click the Save button. Sometimes, a pop-up window will warn you that two partitions overlap (they have a range of sectors in common) and that consequently they should have the same access rights. If this is the case, change the access rights of one of these two partitions, and retry clicking the Save button.
12. The safe raw disk file is now written in your virtual machine directory. Click the Install button to actually connect the raw disk to your virtual machine.

You are now ready to have your raw disk used by your virtual machine.

As with a virtual disk, you should decide in which mode (persistent, nonpersistent, or undoable) the raw disk will be used. Before booting your virtual machine, please carefully read the section on risks in using raw disks, below.

If you need to install an operating system on the raw disk, you may want to look at the technical note Installing an OS onto a Raw Partition from a Virtual Machine Using VMware. If there is already an operating system on the raw disk and you want to use it sometimes as a host operating system and sometimes as a guest operating system, you may want to look at the technical note Configuring Dual/Multiboot Systems to Run With VMware.

Removing a safe raw disk from a virtual machine

1. Open the Configuration Editor in the virtual machine from which you want to remove the safe raw disk.
2. Expand the "IDE Drives" node.
3. Look for raw disk drives. Click on the one you want to remove.
4. Click the Remove button.

Modifying a safe raw disk

Follow this procedure if a safe raw disk file has been generated previously (either because you used the Configuration Wizard or because you used the Configuration Editor as described above in this note) and you would like to modify it.

If you want to use another raw disk or if you have modified the layout of the partitions on the raw disk, first follow the instructions above to remove (or rename) the safe raw disk, then go to the installation section of this note and follow the instructions there to create a new safe raw disk file corresponding to the new raw disk.

If, on the other hand, the raw disk hasn't changed, but you want to change the access rights of the virtual machine to the raw disk (which means that you want to modify the safe raw disk), just click the Edit Raw Disk... button, follow the directions in the Edit Raw Disk... paragraph of the installation section of this technical note, then click the Save button.

Risks in using raw disks

It is perfectly safe to use raw disks if it is done right. The only danger is that the raw disk may be accessed simultaneously from the host OS and a guest operating system. Because each operating system is unaware of the other, data corruption can occur if both operating systems read or write to the same sectors at the same time. The point of using a safe raw disk file is to regulate disk operations of the guest operating system. Currently, VMware does nothing to regulate disk operations of the host operating system.

Consequently, you should ensure that your host operating system doesn't "see" the partitions with which your guest operating system works. The safety of raw risks depends on this one requirement.

If you need to exchange data between a host and a guest operating system, either perform disk accesses sequentially (for example, on Linux, mount the raw disk, put the data on it, unmount the disk, start VMware, read the data, stop VMware), or use network protocols such as SMB (Windows Networking, implemented by Samba under Linux) or NFS.

Frequently Asked Questions

1. What does this message "The partition information in file doesn't match that of device." mean?

   It means the configuration editor has detected that the access rights in your safe raw disk file cannot be applied to the list of partitions present on the device you specified in the "Device" field.

   It typically happens in these cases:

   - The raw disk used to create the safe raw disk file is different from the one you typed in the "Device" field.

- You have modified the layout of the partitions on the raw disk. This can happen, when:
  - You have replaced a drive in the physical machine.
  - You have moved the safe raw disk to another physical machine.
- To get rid of this message, the safest thing to do is to remove the safe raw disk (see instructions above).

2. Can I move a safe raw disk from one physical machine to another?

   No, unless both machines have disks with identical content. Otherwise, VMware will detect the situation and issue the warning message described above in question 1.

3. Can I mix raw disks and virtual disks in the same virtual machine?

   Yes. This feature is particularly useful when you want to transfer a lot of data to a virtual disk.

VMware for Linux Keyboard Mapping

This technical note addresses these issues (and some others too):

- My (language-specific) keyboard is not supported by VMware.
- Some of the keys on my keyboard don't work right in the virtual machine.
- My keyboard works fine when I run a virtual machine locally but not when I run the same virtual machine with a remote X server.

Quick answers

- If your keyboard works correctly with a local X server, and you just want the same behavior with a remote X server (which is also an XFree86 server running on a PC), just add the line

  xkeymap.useKey codeMapIfXFree86 = true

  to the virtual machine configuration file or ~/.vmware/config (on the host machine, where you run the virtual machine, not on the machine with the X server).
- If you are using an XFree86-based server that VMware doesn't recognize as an XFree server, use this instead:

  xkeymap.useKey codeMap = true
- If you are using an XFree86 server running locally, and the keyboard does not work correctly, please file an incident report using the form at http://www.vmware.com/incident.

The complete story

Unfortunately, keyboard support for the PC (virtual or otherwise) is a complex affair. To do it justice, we have to start with some (greatly simplified and not strictly correct) background information.

Pressing a key on the PC keyboard generates a scan code based roughly on the position of the key. For example, the Z key on a German keyboard generates the same code as the Y key on an English keyboard, because they are the same key. Most keys have one-byte scan codes, other keys (with one exception) have two-byte scan codes with prefix 0xe0. Internally, VMware uses a simplified version of the PC scan code that is a single 9-bit numeric value, called v-scan code. V-scan codes are written as a three-digit hexadecimal number. The first digit is 0 or 1. For example, the left-hand control key has a one-byte scan code (0x1d); its v-scan code is 0x01d. The right-hand control key scan code is two bytes (0xe0, 0x1d); its v-scan code is 0x11d.

An X server uses a two-level encoding of keys. An X key code is a one-byte value. The assignment of key codes to keys depends on the X server implementation and the physical keyboard. As a result, an X application normally cannot use key codes directly. Instead, the key codes are mapped into keysyms that have names like "space," "escape," "x," "2." The mapping can be controlled by an X application via XChangeKeyboardMapping() or by the program xmodmap. Also, xev is a handy tool that shows the key codes and keysyms for the keys typed into its window.

To recap, a key code corresponds roughly to a physical key, and a keysym corresponds to the symbol on the key top. For example, with an XFree86 server running on a PC, the Z key on the German keyboard has the same key code as the Y key on an English keyboard. The German Z keysym, however, is the same as the English Z keysym, and different from the Y keysym.

For an XFree86 server on a PC, there is a one-to-one mapping from X key codes to PC scan codes (or v-scan codes, which is what we really use). VMware takes advantage of this fact. When it is using an XFree86 server on the local host (therefore it must be on a PC), then it uses the built-in mapping from X key codes to v-scan codes. This mapping is keyboard independent and should be correct for most (if not all) languages. In other cases (not an XFree86 server or not a local server), VMware must map keysyms to v-scan codes, using a set of keyboard-specific tables.

Key code mapping is simple, automatic, and foolproof. (Keysym mapping is more complex and described later.) However, because the program cannot tell whether a remote server is running on a PC, it errs on the safe side and uses key code mapping only with local X servers. This is often too conservative and undesirable. Luckily, this and other behavior related to key code mapping can be controlled by configuration settings:

- xkeymap.useKey codeMapIfXFree86 = true
  Use key code mapping if using an XFree86 server, even if it is remote.
- xkeymap.useKey codeMap = true
  Always use key code mapping regardless of server type.
- xkeymap.noKey codeMap = true
  Never use key code mapping.
- xkeymap.key code.<code> = <v-scan code>
  If using key code mapping, map key code <code> to <v-scan code>. <code> must be a decimal number. <v-scan code> should be a C-syntax hexadecimal number (for example, 0x001).

The easiest way to find the X key code for a key is to run xev or xmodmap -pk. For v-scan codes, try the (incomplete) table at the end of this document. The keysym mapping tables described below are also helpful.

Use this feature to make small modifications to the mapping. For example, do this to swap left control and caps lock:

```
xkeymap.key code.64 = 0x01d  # X Caps_Lock -> VM left ctrl
xkeymap.key code.37 = 0x03a  # X Control_L -> VM caps lock
```

These configuration lines can be added to the individual VM configuration, to your personal VMware configuration (~/.vmware/config), or even to the host-wide (/etc/vmware/config) or installation-wide (usually /usr/local/lib/vmware/config) configuration.

When key code mapping cannot be used (or is disabled), VMware maps keysyms to v-scan codes. This is done using one of the tables in the "xkeymap" directory in the VMware installation (usually /usr/local/lib/vmware). Which table to use depends on the keyboard layout. The normal distribution includes tables for PC keyboards for the United States and a number of European countries and languages. And for most of these, there are both the 101-key (or 102-key) and the 104-key (or 105-key) variants.

VMware automatically determines which table to use, by examining the current X keymap. However, its heuristics may sometimes fail. In addition, each mapping is fixed and may not be completely right for any given keyboard and X key code-to-keysym mapping. For example, a

user may have swapped control and caps lock using xmodmap. This means the keys will be swapped in the virtual machine when using a remote server (keysym mapping) but unswapped when using a local server (key code mapping).

Therefore, keysym mapping is necessarily imperfect. To make up for this defect, most of the behavior can be changed with configuration settings:

- xkeymap.language = <keyboard-type>
  Use this if VMware has a table in xkeymap for your keyboard but can't detect it. <keyboard-type> must be one of the tables in the xkeymap directory. (See above for location.) However, the failure to detect the keyboard probably means the table isn't completely correct for you.
- xkeymap.keysym.<sym> = <v-scan code>
  If using keysym mapping, map keysym <sym> to <v-scan code> , <sym> must be an X keysym name. <v-scan code> should be a C-syntax hexadecimal number (for example, 0x001). The easiest way to find the keysym name for a key is to run xev or xmodmap -pk . The X header file /usr/X11R6/include/X11/keysymdef.h has a complete list of keysyms. (The name of a keysym is the same as its C constant without the XK_ prefix.) For v-scan codes, try the (incomplete) table at the end of this document. The xkeymap tables themselves are also helpful. Use this feature to fix small errors in an existing mapping.
- xkeymap.fileName = <file-path>
  Use the keysym mapping table in <file-path> . A table is a sequence of configuration lines of the form

  <sym> = <v-scan code>

  where <sym> is an X keysym name, and <v-scan code> should be a C-syntax hexadecimal number (for example, 0x001). (See xkeymap.keysym above for how to find the keysyms and v-scan codes for your keyboard.)

  Compiling a complete keysym mapping is hard. It is best to start with an existing table and make small changes.

V-scan code table

These are the v-scan codes for the 104-key US keyboard:

Symbol	Shifted symbol	Location	v-scan code
Esc			0x001
1	!		0x002
2	@		0x003
3	#		0x004
4	$		0x005
5	%		0x006
6	^		0x007
7	&		0x008
8	*		0x009
9	)		0x00a
0	(		0x00b
-	_		0x00c

=	+		0x00d
Backspace			0x00e
Tab			0x00f
Q			0x010
W			0x011
E			0x012
R			0x013
T			0x014
Y			0x015
U			0x016
I			0x017
O			0x018
P			0x019
[	{		0x01a
]	}		0x01b
Enter			0x01c
Ctrl		(left)	0x01d
A			0x01e
S			0x01f
D			0x020
F			0x021
G			0x022
H			0x023
J			0x024
K			0x025
L			0x026
;	:		0x027
'	"		0x028
`	~		0x029
Shift		(left)	0x02a
\	|		0x02b
Z			0x02c
X			0x02d
C			0x02e
V			0x02f
B			0x030

N			0x031
M			0x032
,	<		0x033
.	>		0x034
/	?		0x035
Shift		(right)	0x036
*		(numeric pad)	0x037
Alt		(left)	0x038
Space bar			0x039
Caps Lock			0x03a
F1			0x03b
F2			0x03c
F3			0x03d
F4			0x03e
F5			0x03f
F6			0x040
F7			0x041
F8			0x042
F9			0x043
F10			0x044
Num Lock		(numeric pad)	0x045
Scroll Lock			0x046
Home	7	(numeric pad)	0x047
Up arrow	8	(numeric pad)	0x048
PgUp	9	(numeric pad)	0x049
-		(numeric pad)	0x04a
Left arrow	4	(numeric pad)	0x04b
5		(numeric pad)	0x04c
Right arrow	6	(numeric pad)	0x04d
+		(numeric pad)	0x04e
End	1	(numeric pad)	0x04f
Down arrow	2	(numeric pad)	0x050
PgDn	3	(numeric pad)	0x051
Ins	0	(numeric pad)	0x052
Del		(numeric pad)	0x053
F11			0x057

F12			0x058
Break	Pause		0x100
Enter		(numeric pad)	0x11c
Ctrl		(right)	0x11d
/		(numeric pad)	0x135
SysRq	Print Scrn		0x137
Alt		(right)	0x138
Home		(function pad)	0x147
Up arrow		(function pad)	0x148
Page Up		(function pad)	0x149
Left arrow		(function pad)	0x14b
Right arrow		(function pad)	0x14d
End		(function pad)	0x14f
Down arrow		(function pad)	0x150
Page Down		(function pad)	0x151
Insert		(function pad)	0x152
Delete		(function pad)	0x153
Windows		(left)	0x15b
Windows		(right)	0x15c
Menu			0x15d

The 84-key keyboard has a Sys Req on the numeric pad:

Symbol	Shifted symbol	Location	v-scan code
Sys Req		(numeric pad)	0x054

Keyboards outside the US usually have an extra key (often < > or < > |) next to the left shift key:

Symbol	Shifted symbol	Location	v-scan code
<	>		0x056

Installing and Compiling the VMware for Linux modules vmmon and vmnet

Scope

This technical note discusses problems related to the installation of VMware for Linux, specifically with respect to the installation of its two kernel modules. You should normally have to refer to this technical note in only two circumstances:

1. You are unable to install VMware successfully, that is, the vmware-install.pl script fails.

2. You have changed your linux kernel, either by upgrading your distribution or by compiling your own kernel. Any kernel upgrade requires you to recompile vmmon and vmnet . The easiest way to do this is simply to upgrade VMware by running vmware-install.pl in the distribution directory and answering Yes to the question "Would you like to upgrade VMware for Linux."

Background

VMware for Linux's installation script (vmware-install.pl) uses a two-step process to install both vmmon and vmnet. First, it tries to find a matching precompiled module in the vmware-distrib directory. If such a module exists, it will try to load it into the kernel using the /sbin/insmod command.

If no precompiled modules exist, or if the precompiled module does not load, the second step is to automatically compile the module from its sources, which are included in the distribution.

The second step will use the system header files to compile the module. The success of the second step requires a perfect match between the header files installed on the system and the header files previously used to compile the kernel.

The majority of the installation problems are due to this discrepancy. Fixing this system discrepancy is the easiest way to install VMware.

Fixing /usr/include/linux and /usr/include/asm

If your installation fails, the first step is to check that you system include directories /usr/include/linux and /usr/include/asm. These are soft links that point to the correct directories within the Linux source tree.

On many popular distributions:

/usr/include/linux is a soft link to /usr/src/linux/include/linux

/usr/include/asm is a soft link to /usr/src/linux/include/asm

/usr/src/linux is a soft link to /usr/src/linux-<kernel version>

If the links are set up differently – for example, /usr/include/linux does not point within the Linux source – a version mismatch is likely.

We recommend that you change your setup so that the two directories /usr/include/linux and /usr/include/asm are in fact soft links that point within the Linux source tree.

This change is required only to compile the modules as part of VMware's installation process. Once VMware is installed, you may reset your system configuration to its original state if you prefer.

Common problems

1. Script fails because <linux/version.h> is inconsistent
   The script fails with this error message:

   ### Something is wrong with the system include files on
   ### your machine! The file <linux/version.h> is for a
   ### [<ACTUAL>] Linux system but you are running a [<EXPECTING>]
   ### kernel. This will not work for building the VMware device
   ### drivers; you must have include files that match the version
   ### of your operating system.

In your case, [<ACTUAL>] will be replaced with the kernel version that is specified by <linux/version.h> and [<EXPECTING>] will be replaced with the kernel version that is returned by uname -r . The C compiler will likely find <linux/version.h> in/usr/include/linux/version.h. On your system, you are also likely to find another version.h file in /usr/src/linux/include/linux/version.h. The latter one should contain the version that corresponds to your kernel. This problem is usually fixed by following the guidelines described above in "Fixing /usr/include/linux and /usr/include/asm. "

2. Script fails with unresolved symbols
   The script fails with error messages similar to:

   ```
   [...]
   unresolved symbol __pollwait_Rsmp_0b89dd34
   unresolved symbol free_pages_Rsmp_234535e0
   unresolved symbol misc_deregister_Rsmp_632a685b
   [...]
   ```

   The Linux build does symbol mangling to guarantee the consistency of configuration options between the kernel and its modules. When the kernel is built, the kernel configuration options are written in to <linux/autoconf.h> , and the kernel mangled symbols into <linux/modversions.h> and <linux/modules/*.ver> .

   In this failure mode, vmmon or vmnet refuse to load because of unresolved symbols. Since they were just built, this again indicates an inconsistency between the system header files installed on your machine and the ones used to build the kernel that you are running.

   First, be sure that your /usr/include/linux and /usr/include/asm point to the right place, as discussed above in Fixing /usr/include/linux and /usr/include/asm. This might solve your problem.

   Often, rebuilding your kernel using the same configuration will fix the problem. Note that you don't need to reinstall the newly built kernel. Just recompiling it will regenerate the <linux/autoconf.h> and the <linux/modules/*.ver> files.

   To rebuild your kernel, first find the /usr/src/linux/.config that should correspond to you kernel version. It is often, but not always, already installed in /usr/src/linux/.config.

   If you have a .config file, do the following:

   ```
   rm /usr/src/linux/include/linux/modules/*.ver
   make oldconfig
   make dep clean zImage
   ```

   and try to reinstall VMware.

   If you still can't install it because of unresolved symbols, it is likely because of a difference between the configuration options of the kernel that you are running and the one that you have just built.

   If this is the case and you cannot find the matching configuration, installing the newly built kernel will allow you to install VMware. Of course, you'll be running a different kernel now.

3. Script fails because CONFIG_UMISC is not defined
   The script fails with the following message:

   ```
   ### You appear to have a [<VERSION>] Linux kernel that was not
   ### built with the CONFIG_UMISC configuration parameter set.
   ### The VMware software will not operate without this facility .
   ### Please update your kernel to include this facility and then
   ### redo the installation procedure.
   ```

/dev/vmmon registers itself as a "miscellaneous character device." Your kernel is not configured with this option, and VMware therefore cannot be installed.

If this is the case, you must recompile your Linux kernel with that option defined. VMware will not run otherwise.

5

Glossary

Glossary

Configuration Editor
A dialog window that can be launched from the VMware Settings menu and allows users to view and modify the configuration of a virtual machine.

Configuration Wizard
A mini-applet that is launched automatically when VMware is started without specifying a configuration file; the applet can also be launched from the VMware File menu. When a new virtual machine is to be set up, the VMware Configuration Wizard will prompt the user for information and create a virtual machine configuration file. To access the same virtual machine in the future, the user can simply run the configuration file that was already created by the wizard.

Disk modes
Setting that defines how disk write commands issued by software running inside a virtual machine should be handled by the VMware software layer. The disk mode setting is transparent to the software that runs inside the virtual machine. There is a mode setting for each virtual machine disk; the setting can be viewed/modified from the Configuration Editor's IDE Drives panel. The disk mode settings are persistent, non-persistent and undoable.

Existing partition
A partition on a real disk in the host machine; also see raw disk.

Guest operating systems
Operating systems that run inside a virtual machine (for example, when Windows 95 is run inside a virtual machine, it can be said that the virtual machine runs Windows 95 as a guest OS).

Host machine (or host system)
A real, physical computer that runs an operating system and application software, such as Linux and VMware for Linux.

Host operating system
An operating system that runs on the host machine (for example, Linux, Windows NT, etc.)

Meta configuration file
In the context of VMware virtual machines configured with raw disks, a meta configuration file defines the level of access rights (read/write) a virtual machine will have to certain partitions on a raw disk device. Same as safe raw disk file.

Non-persistent disk mode
From a virtual machine's point of view, all disk writes issued by software running inside the VM appear to be written to disk, but they are in fact discarded after the session is powered down. As a result, a virtual disk or raw disk set to a nonpersistent disk mode is not modified by VMware software.

Persistent disk mode
All disk writes issued by software running inside a virtual machine are immediately and permanently written to the virtual machine's disk. As a result, a virtual disk or raw disk set to a persistent disk mode behaves like conventional disk drives on your computer.

Raw disk
A physical disk drive in the host machine. From VMware software's point of view, a virtual machine's disk can be stored as a file on the host filesystem (see virtual disk) or on a local IDE raw disk device. When a virtual machine is configured to use a raw disk VMware software will directly access the local disk/partition as a raw device (not as a file on a filesystem). It is possible to boot a previously installed operating system on an existing partition within a virtual machine environment. The only limitation is that the existing partition must reside on a local IDE drive; VMware virtual machines cannot access SCSI drives as a raw device.

Safe raw disk file
A file containing user-specified access privilege information that controls a virtual machine's read/write access to certain partitions on the raw disk. VMware software uses this file to prevent dual-boot users from accidentally trying to run the host operating system again as a guest, or another guest operating system that the virtual machine was not configured for. Safe raw disk files can also prevent accidental writes to raw disk partitions from badly behaved operating systems or applications. Safe raw disk files are created for each IDE raw disk device when the VMware Configuration Wizard is launched with a -rawdisk command line parameter.

Undoable disk mode
From a virtual machine's point of view, all disk writes issued by software running inside the virtual machine appear to be written to disk, but they are in fact stored in a temporary (.REDO) file on the host filesystem for the duration of the VM session. When the virtual machine is powered down, VMware software will prompt the user to decide how the disk changes for the session should be handled. The user will have three choices: (1) save changes to disk permanently, (2) delete changes to prevent modifying the virtual machine's original disk image, or (3) keep the changes in the REDO file and be able to continue running the virtual machine with what appears to be the modified disk image.

Virtual disk
A virtual disk is a file on the host filesystem that appears as a physical disk drive to a guest operating system. The filesystem can be on the host machine or on a remote computer. Configuring a virtual machine with a virtual disk makes it possible to inscall a new operating system onto the disk file without the need to repartition a physical disk or reboot the host. VMware virtual machines can be configured with virtual disk(s), raw disk(s) or a combination of both.

Virtual machine
A virtualized x86 PC environment on which a guest operating system and associated application software can run. Multiple virtual machines can operate on the same host machine concurrently.

Virtual machine configuration file
A file where configuration information about a specific virtual machine (such as disk size, memory amount, etc.) is stored. Virtual machine configuration files are created using the Configuration Wizard or Configuration Editor. Once a configuration file has been created, a user can access the same virtual machine sometime in the future by running the configuration file as an executable.