

Installing and Configuring Linux Guest Operating Systems

VMware® Infrastructure 3

Datacenters today have a mix of Windows and Linux workloads. IDC estimates in 2008 that 68 percent of all physical servers shipped are Windows-based, compared to 23 percent that are Linux-based. However, the proliferation of Linux environments is steadily increasing. From 2006 to 2011, IDC forecasts the compounded annual growth rate (CAGR) of physical server units running Linux at 28.1 percent, with Windows trailing at 25.0 percent. As more datacenters are virtualized with VMware Infrastructure 3, it makes sense that these virtualized environments are also trending towards increased use of Linux. The CAGR of virtual server units running Linux is forecasted by IDC at 44.1 percent, with Windows behind at 39.0 percent. Linux operating systems now host applications from databases to Web servers to application servers, much as their Windows counterparts do. Linux guest operating systems are here, and VMware is dedicated to supporting them.

This technical note describes installing, configuring, updating, and administering Linux guest operating systems in virtual machines running on VMware Infrastructure 3 version 3.5 (specifically VMware VirtualCenter 2.5 Update 2 and VMware ESX 3.5 Update 2). In addition, this note includes a collection of useful tips and tricks in fine-tuning your Linux virtual machines, which may or may not apply to all your Linux usage scenarios. Although the recommendations in this paper apply to most Linux distributions, they are tailored specifically to Red Hat Enterprise Linux 5. IDC observed in 2007 that of all paid Linux subscriptions, Red Hat Enterprise Linux came in at 62.1 percent and Novell SUSE Linux Enterprise Server placed second at 29 percent. Linux administrators can use this paper as a source for guidelines when building and maintaining Linux virtual machines in their VMware Infrastructure environments. Some working knowledge of VirtualCenter 2.5 Update 2, ESX 3.5 Update 2, and Linux operating systems is required.

This technical note covers the following topics:

- [“Linux Support on VMware ESX”](#) on page 2
- [“Installing Linux in a Virtual Machine”](#) on page 2
- [“Installing and Upgrading VMware Tools”](#) on page 4
- [“Cloning a Linux Guest Operating System”](#) on page 9
- [“VMware Update Manager”](#) on page 11
- [“Linux Time Synchronization Recommendations”](#) on page 12
- [“Additional Notes”](#) on page 15
- [“Resources”](#) on page 19
- [“Appendix A: Linux Versions Supported on ESX Server”](#) on page 20
- [“Appendix B: Command-Line Options for VMware Tools Upgrades”](#) on page 21
- [“Appendix C: Enabling VMI in a Linux Kernel and in ESX 3.5”](#) on page 22

Linux Support on VMware ESX

VMware ESX supports the widest range of Linux guest operating systems of any virtualization product. ESX supports Red Hat Enterprise Linux 2.1, 3, 4, and 5, SUSE Linux Enterprise Server 8, 9, and 10, and Ubuntu Linux 7.04, 8.04, and 8.10. In addition, ESX supports almost all updates to these releases as well as specialized variants of them. See [“Appendix A: Linux Versions Supported on ESX Server”](#) on page 20 for a listing of Linux versions supported at the time this paper was written. For a complete, up-to-date listing of guest operating systems supported by VMware products, see the *Guest Operating System Installation Guide*. For a link, see [“Resources”](#) on page 19. Choosing a Linux distribution from this list offers performance benefits over nonsupported Linux distributions because VMware products optimize hypervisor settings based on guest operating system types.

Installing Linux in a Virtual Machine

When you create the virtual machine in which you plan to install your Linux guest operating system, be sure that its devices are set up as you expect. This section describes the creation of a Linux virtual machine from installable media. Once you have created a virtual machine, you can create templates and clones from the base virtual machine. This enables you to provision future virtual machines quickly. See [“Cloning a Linux Guest Operating System”](#) on page 9 for details. For general installation guidelines for all supported Linux distributions, see the *Guest Operating System Installation Guide*. For a link, see [“Resources”](#) on page 19.

Memory Recommendations

Be sure the virtual machine is configured with at least 512MB of memory for Red Hat Enterprise Linux 5 or with 256MB of memory for Red Hat Enterprise Linux 3 or Red Hat Enterprise Linux 4. If the memory in the virtual machine is lower than the recommended values, Red Hat Enterprise Linux presents an error message as it loads certain VMware drivers.

Network Adapter Recommendations

Be sure to select the correct network adapter. For most 32-bit guest operating systems, you can select **Flexible** or **Enhanced vmxnet**. For most 64-bit guest operating systems, you can select **E1000** or **Enhanced vmxnet**. Enhanced vmxnet is not supported on every 32- and 64-bit Linux distribution, but if the option exists, we recommend you select Enhanced vmxnet as your network adapter. See [“vmxnet”](#) on page 8 for details.

SCSI Adapter Recommendations

When creating the virtual machine, be sure to select the LSI Logic SCSI adapter. Red Hat Enterprise Linux 5 does not include a driver for the BusLogic SCSI adapter. Many Linux guest operating systems encounter problems in a virtual machine configured to use the BusLogic virtual SCSI adapter. In most cases, VMware recommends that you use the LSI Logic virtual SCSI adapter with all Red Hat guest operating systems. However, ESX Server 2.5.2, 2.5.3, 2.5.4, and 2.5.5 support only the BusLogic SCSI adapter. VMware provides a separate BusLogic driver for Red Hat Enterprise Linux 4 Upgrades 1, 2, 3, 4, and 5. For instructions on downloading and installing a driver for the BusLogic adapter, see Archives for VMware ESX Server 2.x. For a link, see [“Resources”](#) on page 19.

You can install Red Hat Enterprise Linux 5 in a virtual machine using the standard Red Hat distribution CD, via the boot floppy/network method, or from a PXE server. If you plan to use a PXE server to install the guest operating system over a network connection, you do not need the operating system installation media. When you power on the virtual machine, the virtual machine detects the PXE server. For more details on installation through PXE, see the VMware knowledge base article [“Using PXE \(Preboot Execution Environment\) to Install Guest Operating Systems over a Network.”](#) For a link, see [“Resources”](#) on page 19.

Rather than installing from a physical CD-ROM, you can create an ISO image file from the installation CD-ROM. Using an ISO image file in this way can be particularly convenient if you need to install the same operating system in multiple virtual machines. You can store the ISO file on the host machine or on a network drive accessible from the host machine. Then in the VI Client:

- 1 Right-click the virtual machine in which you want to install the new guest operating system and click **Edit Settings**.
- 2 Select **CD/DVD Drive 1** under the Hardware tab, select **Connect at power on**, and browse to the ISO file under **Datastore ISO file**.
- 3 Power on your virtual machine by clicking the **Power On** button.
- 4 Follow the instructions in [“Installation Steps”](#) on page 3.
- 5 Change back to their normal settings for CD/DVD Drive 1 if you do not want the ISO file to remain connected during subsequent reboots.

Installation Steps

- 1 Follow the installation steps as you would for a physical machine. Be sure to make the choices outlined in the following steps.
- 2 Allow automatic partitioning of the disk to occur in the Automatic Partitioning screen or partition the virtual disk manually if you do not want to use the Red Hat defaults. You might see a warning that begins “The partition table on device <devicename> was unreadable. To create new partitions it must be initialized, causing the loss of ALL DATA on the drive.” This does not mean that anything is wrong with the hard drive on your physical computer. It simply means that the virtual hard drive in your virtual machine needs to be partitioned and formatted.
- 3 Click **Yes** to partition and format the virtual hard drive.
- 4 Do not select the **Virtualization software** option during the installation. Uncheck the **Virtualization** box. If you select this option, Red Hat Enterprise Linux 5 or Fedora Core 7 installs the Xen hypervisor and a XenLinux kernel. You might experience a number of performance and functionality issues. For more information, see the VMware knowledge base article “RHEL 5 and FC 7 Guests Installed with Red Hat Virtualization Affects Performance of Virtual Machine.” For a link, see [“Resources”](#) on page 19.

Choosing a Kernel

Some older Linux distributions such as Red Hat Enterprise Linux 3 and Red Hat Advanced Server 2.1 install kernels optimized for AMD processors when one is detected at install time. If a virtual machine with such an AMD Linux kernel is ever run on an Intel-based host, it can encounter problems, as described in the VMware knowledge base article “Linux Guest Moved to a System with Different Processor Type Panics During Boot.” For a link, see [“Resources”](#) on page 19. Kernels that are optimized for Intel processors are also compatible with AMD processors and do not have this issue, however, it is better to match the kernel to the hardware platform. Currently for 64-bit processors, Linux vendors have a single kernel for AMD and Intel platforms, so this is not an issue for 64-bit guest operating systems.

It is also important to select a Linux kernel that is appropriate for your needs. Linux distributions typically offer a choice of precompiled kernels optimized for various system memory and SMP configurations. As a rough rule of thumb, kernels optimized for the following systems are listed in order of performance:

- Uniprocessor (UP).
- Symmetric multiprocessor (SMP).
- Symmetric multiprocessor with physical address extensions (SMP-PAE).
- Symmetric multiprocessor with physical address extensions and separation of kernel and user space entirely so that each can make full use of the 4GB virtual address space on 32-bit systems (SMP-PAE, 4G/4G split). These SMP-PAE, 4G/4G split kernels are not supported by VMware products.

On Intel platforms, 64-bit guests can be run only with hardware assist (Intel VT). On AMD platforms 64-bit guests can be run either with hardware assist (AMD-V) or using binary translation. A VT virtual machine monitor incurs less overhead on the newer Intel Core 2 processors than on the older Pentium 4 processors. If you have a choice between a 32-bit or 64-bit guest, select the 32-bit guest on older Intel platforms (Pentium 4). A 64-bit guest is a better choice for newer Intel platforms and for guest operating systems that access a large amount of memory (more than 4GB), because 64-bit kernels can address the entire memory space without complex memory management overhead. Also, the VT monitor for 64-bit guests works faster on newer Intel processors than the older Pentium 4 processors.

The tables below highlight the main Red Hat Enterprise Linux kernel packages and what they contain. See the respective URL links for full details on all available kernel packages for each version of Red Hat Enterprise Linux.

Table 1. Red Hat Enterprise Linux 5.2 Kernel RPM Descriptions

Information link	https://www.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5.2/html/Deployment_Guide/s1-kernel-packages.html
Kernel RPMs	<ul style="list-style-type: none"> ■ kernel = UP, SMP, non-PAE (4G) ■ kernel-PAE = SMP, PAE (up to 64G)

Table 2. Red Hat Enterprise Linux 4 Kernel RPM Descriptions

Information link	https://www.redhat.com/docs/manuals/enterprise/RHEL-4-Manual/en-US/System_Administration_Guide/Manually_Upgrading_the_Kernel-Overview_of_Kernel_Packages.html
Kernel RPMs	<ul style="list-style-type: none"> ■ kernel = UP, SMP, non-PAE (4G) ■ kernel-hugemem (not supported) = SMP, PAE (up to 64G), 4G/4G split ■ kernel-smp = SMP, PAE (up to 16G)

Table 3. Red Hat Enterprise Linux 3 Kernel RPM Descriptions

Information link	https://www.redhat.com/docs/manuals/enterprise/RHEL-3-Manual/sysadmin-guide/ch-kernel.htm#S1-KERNEL-PACKAGES
Kernel RPMs	<ul style="list-style-type: none"> ■ kernel = UP, SMP, non-PAE (4G) ■ kernel-hugemem (not supported) = SMP, PAE (up to 64G), 4G/4G ■ kernel-smp = SMP, PAE (up to 16G)

The 64-bit kernels can address the entire memory space in the virtual machine directly and do not need a special memory management kernel.

Installing and Upgrading VMware Tools

VMware Tools is a suite of utilities that enhances the performance of a virtual machine's guest operating system and improves management of the virtual machine. Installing VMware Tools in the guest operating system is vital. Although the guest operating system can run without VMware Tools, you lose important functionality and convenience. See [“Open Virtual Machine Tools”](#) on page 19 for information on an open source project that allows the use, modification, and redistribution of most of the components of VMware Tools.

The following are installed with VMware Tools:

- The VMware Tools service (`vmware-guestd` on Linux guests). This service synchronizes the time in the guest operating system with the time in ESX.
- A set of VMware device drivers, including an SVGA display driver, the `vmxnet` accelerated networking driver (as described in [“Major Components Installed by VMware Tools”](#) on page 8), the BusLogic SCSI driver (as described in [“Installing Linux in a Virtual Machine”](#) on page 2), the memory control driver for efficient memory allocation between virtual machines, the sync driver to quiesce I/O for VMware Consolidated Backup, and the VMware mouse driver.

- The VMware Tools control panel, which lets you modify settings, shrink virtual disks, and connect and disconnect virtual devices.
- A set of scripts that helps you automate guest operating system operations. The scripts run when the virtual machine's power state changes if you configure them to do so.
- The VMware user process (`vmware-user` on Linux guests), which enables you to copy and paste text between the guest and managed host operating systems. In Linux guests, this process controls grabbing and releasing the mouse cursor when the SVGA driver is not installed.

Make sure you configure the guest operating system to include the development packages before installing or reinstalling VMware Tools. This enables VMware Tools to determine the correct mouse configuration and module configuration. An improperly configured guest operating system can cause problems, including guest operating system crashes, when you install VMware Tools.

VMware Tools has the following limitations in Linux virtual machines:

- Shrink disk is not supported.
- The mouse driver installation fails in X Window System versions earlier than 4.2.0.

NOTE If you do not have VMware Tools installed in your virtual machine, you cannot use the shutdown or restart options in VirtualCenter. You can use only the power options. If you want to shut down the guest operating system, shut it down from within the virtual machine console before you power off the virtual machine.

The installers for VMware Tools are built into ESX as ISO image files. An ISO image file looks like a CD-ROM to your guest operating system. You do not use an actual CD-ROM disc to install VMware Tools, nor do you need to download the CD-ROM image or burn a physical CD-ROM of this image file.

When you select to install VMware Tools, VirtualCenter temporarily connects the virtual machine's first virtual CD-ROM disk drive to the ISO image file that contains the VMware Tools installer for your guest operating system. You are ready to begin the installation process.

To install or upgrade VMware Tools on a Linux guest from X with the RPM installer

You can perform RPM installations only in certain Linux distributions such as Red Hat and SUSE Linux Enterprise Server, among others. For other Linux distributions, follow the RMP installation guidelines provided with the distribution or see the next section.

- 1 Open a console to the virtual machine.
- 2 Power on the virtual machine.
- 3 After the guest operating system starts, right-click the virtual machine and select **Install/Upgrade VMware Tools**. The remaining steps take place inside the virtual machine.
- 4 Do one of the following:
 - If you see a VMware Tools CD icon on the desktop, double-click it, and after it opens, double-click the RPM installer in the root of the CD-ROM.
 - If you see a file manager window, double-click the RPM installer file.

In some Linux distributions, the VMware Tools CD icon might fail to appear on the desktop. In this case, double-click **Computer** and double-click the CD-ROM drive and you should see a `VMwareTools-3.5.0-xxxxx.i386.rpm` file as well as a `VMwareTools-3.5.0-xxxxx.tar.gz` file (where `<xxxxx>` is the build number of the ESX release—for example, build number 82663 for ESX 3.5 Update 1). Or you may install VMware Tools from the command line, as described in the next section.

- 5 If prompted, enter the root password and click OK. The installer prepares the packages.
- 6 Click **Continue** when the installer presents a dialog box that shows Completed System Preparation. A dialog box appears with a progress bar. When the installer is done, VMware Tools is installed. There is no confirmation or finish button, however you can verify installation in VI Client in the Summary tab of the Linux virtual machine. You can delete the `vmware-tools-distrib` folder.

- 7 In an X terminal, as root (`su -`), run the following file to configure VMware Tools:

```
vmware-config-tools.pl
```

- 8 When done, exit from the root account:

```
exit
```

- 9 In an X terminal, open the VMware Tools Properties dialog box:

```
vmware-toolbox &
```

To install or upgrade VMware Tools on a Linux guest with the tar installer or RPM installer

For enhanced performance, you should install VMware Tools even if you are not running the X Window System.

- 1 Open a console to the virtual machine.
- 2 Power on the virtual machine.
- 3 After the guest operating system starts, right-click the virtual machine and select **Install/Upgrade VMware Tools**. The remaining steps take place inside the virtual machine.
- 4 As root (`su -`), mount the VMware Tools virtual CD-ROM image and change to a working directory (for example, `/tmp`), as follows.

NOTE Some Linux distributions automatically mount CD-ROMs. If your distribution uses auto-mounting, do not use the `mount` and `umount` commands described in this procedure. For example, your CD-ROM may already be mounted to `/mnt/cdrom` or `/media`. However, you still must untar the VMware Tools installer to `/tmp`.

Some Linux distributions use different device names or organize the `/dev` directory differently. Modify the following commands to reflect the conventions used by your distribution:

```
mount /dev/cdrom /mnt/cdrom
cd /tmp
```

NOTE If you have a previous installation, delete the previous `vmware-tools-distrib` directory before installing. The default location of this directory is:
`/tmp/vmware-tools-distrib`

- 5 Uncompress the installer and unmount the CD-ROM image.

Depending on whether you are using the tar installer or the RPM installer, do one of the following:

- For the tar installer, at the command prompt, enter:

```
tar xzpf /mnt/cdrom/VMwareTools-3.5.0-.tar.gz
umount /dev/cdrom
```

Where `<xxxxx>` is the build number of the ESX release. For example, build number 82663 for ESX 3.5 Update 1.

- For the RPM installer, at the command prompt, enter:

```
rpm -Uhv /mnt/cdrom/VMwareTools-3.5.0-.i386.rpm
umount /dev/cdrom
```

Where `<xxxxx>` is the build number of the ESX release. For example, build number 82663 for ESX 3.5 Update 1.

NOTE If you attempt to install an RPM installation over a tar installation—or the reverse—the installer detects the previous installation and must convert the installer database format before continuing.

- 6 Depending on whether you are using the tar installer or the RPM installer, do one of the following:

- For the tar installer, run the VMware Tools tar installer:

```
cd vmware-tools-distrib
```

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

Respond to the configuration questions on the screen. Press Enter to accept the default value.

The `vmware-config-tools.pl` script should run automatically at the end of this installation.

- For the RPM installer, you need to configure VMware Tools by running `vmware-config-tools.pl` manually:

```
vmware-config-tools.pl
```

VMware Tools is now installed and can be verified in VI Client in the Summary tab of the Linux virtual machine. You can delete the `vmware-tools-distrib` folder.

- 7 Log off the root account.

```
exit
```

- 8 Start your graphical environment.

- 9 In an X terminal, open the VMware Tools Properties dialog box:

```
vmware-toolbox &
```

Configuring VMware Tools with `vmware-config-tools.pl`

The VMware Tools installation is incomplete until you run `vmware-config-tools.pl`, as described in the installation procedures above. The `vmware-config-tools.pl` script configures VMware Tools before the package runs for the first time. In addition, you must run this configuration script again if you upgrade or modify the Linux kernel of this virtual machine. Running `vmware-config-tools.pl` also gives you the opportunity to set the display resolution of the X Window System. Typically the resolutions range from 640 × 480 to 2364 × 1773 with the default selection as 1024 × 768.

Verifying VMware Tools Setup

To verify that the components of VMware Tools are installed, check whether the specific modules you want to confirm are installed. For example, listing the modules to check for `vmxnet` is the best way to verify that the `vmxnet` module is available in memory. Use the following `lsmod` command, which lists information about all loaded modules:

```
lsmod | grep vmxnet
```

This should return a line showing statistics on the `vmxnet` driver.

Also check `/etc/modprobe.conf` and make sure that the `pcnet32` or `vmxnet` driver is named next to the appropriate network devices. If `vlan` is included in one of these lines, the `vmxnet` driver is not used even if it has been loaded into memory. The `modprobe.conf` file specifies which options are to be used with which modules when there are conflicting modules loaded. For descriptions of `vmxnet` and `vlan`, see “[vmxnet](#)” on page 8.

Displaying the VMware Tools Properties Dialog Box

Use the VMware Tools Properties dialog box to configure VMware Tools inside your virtual machine. Use this dialog box to configure such things as time synchronization between host and guest, notifications of VMware Tools updates, and specifying which scripts to run when the virtual machine’s power state changes. For instructions on using the VMware Tools Properties dialog box, click the Help button inside the dialog box.

To display the VMware Tools Properties dialog box, open a console to the virtual machine and then open a terminal window and enter the command:

```
/usr/bin/vmware-toolbox &
```

Major Components Installed by VMware Tools

The major components of VMware Tools when installed include `vmblock`, `vmdesched` (to be deprecated in upcoming releases of VMware Tools), `vmhgfs`, `vmmemctl`, `vmsync`, and `vmxnet`. This section takes a closer look at `vmxnet`, the component of most interest to Linux administrators.

vmxnet

`vmxnet` is a Linux kernel device driver for the VMware high-speed virtual networking device. The network devices in a virtual machine are based on real hardware. For example, `vlance` is a virtual device that provides strict emulation of the AMD Lance PCNet32 Ethernet adapter for 32-bit guests, and `e1000` is a virtual device that provides strict emulation of the Intel E1000 Ethernet adapter for 32-bit and 64-bit guests. When creating a 32-bit virtual machine, if you select **Flexible** as the network adapter, the virtual machine defaults to `vlance` before VMware Tools is installed but switches to `vmxnet` after VMware Tools is installed. When creating a 64-bit virtual machine, if you select **E1000** as the network adapter, the virtual machine uses `e1000` as the network adapter regardless of whether VMware Tools is installed.

Alternatively, for some 32- and 64-bit guests, you can select **Enhanced vmxnet** as the network adapter. `vmxnet` is specifically designed for virtual machines to improve performance. Enhanced VMXNET, introduced in ESX 3.5, provides a new version of the `vmxnet` virtual device (the VMware paravirtualized virtual networking device for guest operating systems). Enhanced VMXNET includes several new networking I/O performance improvements, such as support for TCP/IP segmentation offload (TSO) and jumbo frames. All other networking features, such as teaming and VLANs, are fully supported. To enable TSO and jumbo frames, see the *ESX Server 3 Configuration Guide*. For a link, see “[Resources](#)” on page 19. You can also use the command line interface to configure MTU (maximum packet size) and to enable or disable TSO. Enhanced VMXNET is not supported on every 32- and 64-bit Linux distribution, but if the option exists, we recommend you select Enhanced VMXNET as your network adapter for better performance. This is especially true on Intel VT-x systems where VMEXIT costs are high. `vmxnet` has a lower rate of VMEXITs compared with `e1000`.

Enhanced VMXNET is supported for only the following Linux guest operating systems:

- Red Hat Enterprise Linux 5 (32 and 64 bit)
- Red Hat Enterprise Linux 4 (64 bit, no jumbo frame support)
- SUSE Linux Enterprise Server 10 (32 and 64 bit)

Experimental support is provided for:

- Red Hat Enterprise Linux 3 (64 bit, no TSO, no jumbo frame support)
- SUSE Linux Enterprise Server 9 (64 bit, no jumbo frame support)
- Ubuntu 7.04 (64 bit, no jumbo frame support)

Because it is backed by actual virtual hardware, Enhanced VMXNET should be loaded automatically by `hotplug` or `udev` as needed. For best performance, we recommend that you enable TSO on all interfaces driven by `vmxnet` using `ethtool`.

To enable TSO, use shell code similar to the following:

```
if which ethtool >/dev/null 2>&1; then
    for ethif in `ifconfig -a | grep ^eth
                | cut -d' ' -f1`; do
        ethtool -K $ethif tso on >/dev/null 2>&1
    done
fi
```

VMware Tools Upgrades

You can upgrade VMware Tools manually, or you can configure virtual machines to check for and install newer versions of VMware Tools automatically. The following are required for automatic upgrades:

- Virtual machines must have a version of VMware Tools shipped with ESX Server 3.0.1 or greater installed.

- Virtual machines must be hosted on ESX Server 3.0.1 or greater, and the VirtualCenter server must be version 2.0.1 or greater.
- Virtual machines must be running a guest operating system that is supported by ESX Server 3.0.1 or greater and VirtualCenter 2.0.1 or greater.
- Virtual machines must be powered on.

To manually upgrade VMware Tools:

- 1 Launch the VI Client and log in to the VirtualCenter server.
- 2 Select the **Inventory > Hosts and Clusters** view.
- 3 Select the host or cluster that contains the virtual machines you want to upgrade.
- 4 Select the **Virtual Machines** tab.
- 5 Select the virtual machines you want to upgrade and power them on.
- 6 Right-click your selections and select **Install/Upgrade VMware Tools**.
- 7 (Optional) Enter command-line options in the **Advanced** field (see [“Appendix B: Command-Line Options for VMware Tools Upgrades”](#) on page 21).
- 8 Click **OK**.

To configure virtual machines to upgrade VMware Tools automatically:

- 1 Make sure your virtual machine is powered off and open the Virtual Machine Properties dialog box for the virtual machine you want to upgrade.
- 2 Select **Options** tab > **VMware Tools**.
- 3 Select **Check and upgrade Tools before each power-on** under **Advanced**.
- 4 Click **OK**.

The next time the virtual machine is powered on, it checks the ESX host for a newer version of VMware Tools. If one is available, it is installed and, if required, the guest operating system is restarted.

Cloning a Linux Guest Operating System

After you have created and configured a Linux virtual machine, you have the option of converting it into a template. You can then use tools in VirtualCenter to clone additional virtual machines from this template. During the cloning process a technique called customization allows you to customize the identity and network settings of your virtual machine’s guest operating system so it is ready to begin work immediately in your target environment. Cloning and customization allow you to provision new virtual machines much more quickly than you can by building them manually. You can save your virtual machine settings in a specification that you can recall later and reuse. You do this using the Guest Customization wizard. There are several ways to access the customization wizard, as described below.

Linux Requirements for Guest Customization

Guest customization of a Linux guest operating system can occur if:

- The clone or template has one of the following Linux versions installed:
 - Red Hat Enterprise Linux AS versions 2 through 5 (including 64-bit versions)
 - Red Hat Application Server versions 2 through 5 (including 64-bit versions)
 - SUSE LINUX Enterprise Server 8, 9, or 10

NOTE Customization for Red Hat Linux 4 and greater and SUSE Linux Enterprise Server 9 and greater is supported only on hosts running ESX 3.5 and greater or ESXi 3.5 and greater.

- The most recent version of VMware Tools is installed in the guest operating system.

- The clone or template has a root volume formatted with an ext2, ext3, or ReiserFS file system.

NOTE SUSE Linux Enterprise Server distributions register the Ethernet MAC address information in the guest operating system's configuration files. This causes the process of obtaining IP addresses after cloning to fail. To avoid this issue edit the `/etc/sysconfig/network/config` file, adding the following:

```
FORCE_PERSISTENT_NAMES=no
MANDATORY_DEVICES=no
```

Customizing a Linux Guest Operating System

You can start the Guest Customization wizard indirectly from:

- The Deploy Template wizard
- The Clone Virtual Machine wizard
- The Customization Specification Manager

by selecting a virtual machine, then selecting **Edit > Customization Specifications**, and then clicking **New**.

- 1 Start the Guest Customization wizard using a method described above.
- 2 If you opened the Guest Customization Wizard from the Customization Specification Manager, select Linux as the **Target Virtual Machine OS** and give it a **Name** and (optional) **Description**. Click **Next**.
- 3 On the Computer Name page, specify the **Computer Name** and the **Domain Name** for the virtual machine. Click **Next**.

The computer name is the name given to the particular instance of a guest operating system. The operating system uses this name to identify itself on the network. On Linux systems, it is called the host name. This is not the same as the VMware Infrastructure 3 virtual machine name that was declared earlier in the Deploy Template wizard or Clone Virtual Machine wizard. You can set the computer name using the following options:

- **Use a specific name:** The name can contain alphanumeric characters and the underscore (`_`) and hyphen (`-`) characters. It cannot contain periods (`.`) or blank spaces and cannot be made up of digits only. If you want to ensure that the name is unique so that you do not incur conflicts, select **Append a numeric value to ensure uniqueness**. Names are case-insensitive: the name `my_vm` is identical to `My_Vm`.
 - **Use the virtual machine name:** The computer name that VirtualCenter creates is identical to the name of the virtual machine on which the guest operating system is running.
 - **Prompt the user for a name in the Deploy Wizard:** VI Client populates the Deploy Virtual Machine wizard with a prompt for the computer name after you complete all the steps in the wizard.
 - **Use a custom application configured with the VirtualCenter Server to generate a name:** Enter a parameter that can be passed to the custom application.
- 4 If you want VirtualCenter to configure all network interfaces automatically from a DHCP server, select **Typical settings** on the Network page. Click **Next**.

If using a DHCP server is not appropriate for your environment, select **Custom settings**, and click **Next**.

- 5 Select the network interface card (NIC) to customize and click **Customize** to make additional specifications.

Use the Network Properties dialog box to perform these steps:

- a On the General tab, select whether you want to use DHCP to obtain an IP address automatically or enter the IP addresses manually. Also select whether you want to use DHCP to obtain a DNS server address automatically or enter the DNS server addresses manually.
- b Click **OK** to return to the previous dialog box. Click **Next** on the Network Guest Customizations page.

- 6 On the DNS and Domain settings page, enter the IP addresses for the DNS servers. Specify the DNS connections by entering DNS suffixes. For each DNS suffix you enter, click **Add**. If you are adding multiple DNS connections, use **Move Up** and **Move Down** to specify the order in which a virtual machine is to use the connections.
- 7 If you opened the Guest Customization Wizard from the Deploy Template wizard or the Clone Virtual Machine wizard, you see a Save Specification page where you can optionally save the customized options as an .xml file by completing these steps:
 - a Select **Save this customization specification for later use**.
 - b Specify the filename for the specification, and click **Next**.

VirtualCenter saves the customized configuration parameters in the VirtualCenter database. If the customization settings are saved, the administrator and domain administrator passwords are stored in encrypted format in the database. Because the certificate used to encrypt the passwords is unique to each VirtualCenter Server, reinstalling the VirtualCenter Server, or attaching a new instance of the server to the database, invalidates the encrypted passwords. The passwords must be re-entered before they can be used.

- 8 Click **Finish**.

VirtualCenter closes the Guest Customization wizard and returns you to the Deploy Template wizard, Clone Virtual Machine wizard, or Customization Specification Manager.

Completing Linux Guest Operating System Customization

The customization process is not complete until the guest operating system boots, runs the finalization scripts, and reaches the log in page. A customized Linux virtual machine does not need any additional rebooting and is operational as soon as the log in page appears after the first boot. If configuration errors occur, they are displayed in the virtual machine's console window while the guest operating system is booting. These errors are also logged to `/var/log/vmware/customization.log`.

VMware Update Manager

VMware Update Manager enables administrators to apply updates and patches across ESX hosts and all managed virtual machines. Update Manager provides the ability to create user-defined security baselines which represent a set of security standards. Security administrators can compare hosts and virtual machines against these baselines to identify systems that are not in compliance. See the VMware Update Manager Release Notes for the latest list of supported guest operating systems. For a link, see "[Resources](#)" on page 19.

Currently, VMware Update Manager 1.0 Update 2 supports scanning of the following powered off and powered on Linux guest operating systems:

- Red Hat Enterprise Linux AS 3.0 (Update 5 or later)
- Red Hat Enterprise Linux ES 3.0 (Update 5 or later)
- Red Hat Enterprise Linux AS 4.0 (Update 2 or later)
- Red Hat Enterprise Linux ES 4.0 (Update 2 or later)

Installing the Update Manager Guest Agent

VMware Update Manager Guest Agent facilitates Update Manager processes. The Guest Agent is installed at different times depending on the operating system the virtual machine is running. For Linux guest operating systems, using VirtualCenter 2.5 Update 2, the Guest Agent is installed when a scan operation is performed by VMware Update Manager for the first time.

For best results, ensure that the latest version of the Guest Agent is installed. If, for some reason, Guest Agent installation does not complete successfully, scanning operations will fail. In such a case, manually install the Guest Agent. The Guest Agent installation packages for Linux guests are in the Windows location you specified when installing Update Manager server. In that directory, Guest Agent installation packages are at

\docroot\vmci\guestAgent\. For example, if Update Manager was installed in C:\Program Files\VMware\Infrastructure\Update Manager, the Guest Agent installers are at C:\Program Files\VMware\Infrastructure\Update Manager\docroot\vmci\guestAgent\. The Guest Agent requires no user input so that the installation completes silently.

In Linux, install the VMware-VCIGuestAgent-Linux.rpm file by issuing the following command:

```
rpm -ivh VMware-VCIGuestAgent-Linux.rpm
```

The Guest Agent monitors changes to the guest operating system's RPM package database. Whenever that database is changed, the Guest Agent reports the RPM package database contents to Update Manager server. If the virtual machine is online at the time of the scan, the Guest Agent reports the current package information during the scan.

Scanning a Linux guest operating system

During a scanning operation, Update Manager compares a guest operating system's package database contents against a predefined baseline of patches. Specifically, the Update Manager Guest Agent monitors changes to the guest operating system's RPM package database. Whenever that database is changed, the Guest Agent reports the RPM package database contents to the Update Manager server. If the virtual machine is powered on at the time of the scan, the Guest Agent immediately reports the current package information during the scan. The scan result indicates whether the Linux guest operating system is compliant or noncompliant to the baseline. For noncompliant operating systems, the scan indicates the specific patches that are missing. You can remediate these missing patches in the Linux guest operating system by invoking either the software updater utility or the `update` utility found in most Linux distributions. These utilities enable you to download and install specific patch packages.

Linux Time Synchronization Recommendations

The recommendations below include best practices on the particular kernel command line options to use for specific Linux operating systems. This section also provides a description of the recommended settings and usage for NTP time synchronization, configuration of VMware Tools time synchronization, and virtual hardware clock configuration to achieve best timekeeping results. For more details, see VMware knowledge base article 1006427 "Timekeeping best practices for Linux" and the paper "Timekeeping in VMware Virtual Machines." For links, see "[Resources](#)" on page 19.

Tables 4 and 5 show a listing of common Linux guest operating systems and gives kernel parameters required for better timekeeping. The Notes column indicates whether that guest operating system is recommended or if you should avoid using it for time-sensitive virtual machines. In all other cases, the guest operating system has acceptable timekeeping performance.

When both SMP and UP kernels are available, they must be used in virtual machines configured with the corresponding virtual hardware—multiple processors for an SMP kernel or a single processor for a UP kernel—otherwise the mismatch may cause time to drift.

Editing Kernel Configuration

Kernel command line parameters are specified in the `/etc/lilo.conf` or `/boot/grub/grub.conf` file, depending on your choice of boot loader.

For LILO, put the kernel command line parameters at the end of the `append` line. For example, if the `append` line looks like:

```
append="resume=/dev/hda6 splash=silent"
```

and you want to add `clock=pmtmr divider=10`, the updated text is:

```
append="resume=/dev/hda6 splash=silent clock=pmtmr divider=10"
```

Remember to run `/sbin/lilo` after editing `lilo.conf`, so that your edits take effect.

For GRUB, put the kernel command line parameters at the end of the `kernel` line. For example if the kernel line looks like:

```
kernel /vmlinuz-2.6.18 ro root=/dev/hda2
```

and you want to add `clock=pmtmr divider=10`, the updated text is:

```
kernel /vmlinuz-2.6.18 ro root=/dev/hda2 clock=pmtmr divider=10
```

For additional information about working with boot loaders, see your Linux distribution's documentation.

Table 4. Settings for 32-bit Kernels

Linux Version	Kernel Parameters	Notes
Red Hat Enterprise Linux 5.2	<code>divider=10 clocksource=acpi_pm</code>	
Red Hat Enterprise Linux 5.1	<code>divider=10 clocksource=acpi_pm</code>	
Red Hat Enterprise Linux 5.0	<code>clocksource=acpi_pm</code>	
Red Hat Enterprise Linux 4.7	<code>clock=pmtmr divider=10</code>	
Red Hat Enterprise Linux 4.0, 4.1, 4.2, 4.3, 4.4, 4.5, 4.6	<code>clock=pmtmr</code>	
Red Hat Enterprise Linux 3 (all updates)		No additional kernel parameters required.
SUSE Linux Enterprise Server 10 SP2 on ESX 3.5 and later		Recommended. Use a VMI-enabled kernel.
SUSE Linux Enterprise Server 10 SP2 on ESX Server 3.0.x and earlier	<code>clock=pmtmr</code>	
SUSE Linux Enterprise Server 10 SP1	<code>clock=pmtmr</code>	
SUSE Linux Enterprise Server 10	<code>clock=pmtmr</code>	
SUSE Linux Enterprise Server 9 (all updates)	<code>clock=pmtmr</code>	
Ubuntu 8.04 on ESX 3.5 and later		Recommended. Use a VMI-enabled kernel.
Ubuntu 8.04 on ESX Server 3.0.x and earlier	<code>clocksource=acpi_pm</code>	
Ubuntu 7.04, 7.10	<code>clocksource=acpi_pm</code>	
Ubuntu 5.04, 5.10, 6.06, 6.10	<code>clock=pmtmr</code>	

Among Red Hat Enterprise Linux 4 versions, Red Hat Enterprise Linux 4.7 has the best timekeeping performance because of the availability of the `divider=10` option.

Among Red Hat Enterprise Linux 5 version, Red Hat Enterprise Linux 5.1 and 5.2 have the best timekeeping performance because of the availability of the `divider=10` option.

Table 5. Settings for 64-bit Kernels

Linux Version	Kernel Parameters	Notes
Red Hat Enterprise Linux 5.2	<code>notsc divider=10</code>	
Red Hat Enterprise Linux 5.1 with RHSA-2007:0993-13	<code>notsc divider=10</code>	
Red Hat Enterprise Linux 5.1 without RHSA-2007:0993-13	<code>notsc</code>	
Red Hat Enterprise Linux 5.0		No additional kernel parameters required.
Red Hat Enterprise Linux 4.7	<code>notsc divider=10</code>	
Red Hat Enterprise Linux 4.2, 4.3, 4.4, 4.5, 4.6	<code>notsc</code>	
Red Hat Enterprise Linux 4.0, 4.1		Does not support <code>notsc</code> . Avoid using if possible.

Table 5. Settings for 64-bit Kernels

Linux Version	Kernel Parameters	Notes
Red Hat Enterprise Linux 3 (all updates)		Has no workaround for lost tick overcompensation. Avoid using if possible.
SUSE Linux Enterprise Server 10 SP2 on ESX 3.5 and later		Recommended.
SUSE Linux Enterprise Server 10 SP2 on ESX Server 3.0.x and earlier	notsc	
SUSE Linux Enterprise Server 10 SP1	notsc	
SUSE Linux Enterprise Server 10	notsc	
SUSE Linux Enterprise Server 9 with kernel version 2.6.5-7.312 or later	ignore_lost_ticks	
SUSE Linux Enterprise Server 9 with kernel version 2.6.5-7.311 or earlier		Has no workaround for lost tick overcompensation. Avoid using if possible.
Ubuntu 7.10, 8.04	clocksource=acpi_pm	
Ubuntu 5.10, 6.06, 6.10, 7.04	notsc	
Ubuntu 5.04		Does not support notsc. Avoid using if possible.

Recommended Configurations

Based on vendor support in their Linux kernels, we expect the following configurations to have the best timekeeping behavior in a Linux virtual machine:

- SUSE Linux Enterprise Server 10 SP2 32-bit or 64-bit running on ESX 3.5 Update 2 or later
- Ubuntu 8.04 32-bit running on ESX 3.5 Update 2 or later

Among Red Hat Enterprise Linux 4 and 5, versions with the `divider=10` option have better timekeeping behavior than those without.

For more information on VMI-enabled kernels, see the following VMware knowledge base articles:

- “How to enable a virtual machine interface in a Linux kernel and in ESX Server 3.5”
- “Enabling VMI with SLES10 SP2 32bit virtual machines on ESX”

For links, see “[Resources](#)” on page 19.

NTP Recommendations

Whenever possible, use NTP instead of VMware Tools periodic time synchronization. Also, you may need to open the firewall (UPD 123) to allow NTP traffic.

To enable NTP, you must modify `/etc/ntp.conf`. The following is a sample `/etc/ntp.conf`:

```
tinker panic 0
restrict 127.0.0.1
restrict default kod nomodify notrap
server 0.vmware.pool.ntp.org
server 1.vmware.pool.ntp.org
server 2.vmware.pool.ntp.org
driftfile /var/lib/ntp/drift
```

The following is a sample `/etc/ntp/step-tickers`:

```
0.vmware.pool.ntp.org
1.vmware.pool.ntp.org
```

The configuration directive `tinker panic 0` instructs NTP not to give up if it sees a large jump in time. This is important for coping with large time drifts and also resuming virtual machines from their suspended state.

NOTE The directive `tinker panic 0` must be at the top of the `ntp.conf` file.

It is also important not to use the local clock, often referred to as the undisciplined local clock, as a time source. NTP has a tendency to fall back to this source in preference to the remote servers when there is a large amount of time drift.

An example of such a configuration is:

```
server 127.127.1.0
fudge 127.127.1.0 stratum 10
```

Comment out both of these lines from `/etc/ntp.conf`.

After making changes to the NTP configuration, you must restart the NTP daemon. See your operating system vendor's documentation for details.

VMware Tools Time Synchronization Configuration

When using NTP in the guest operating system, disable VMware Tools periodic time synchronization using one of the following options:

- 1 Set `tools.syncTime = "0"` in the configuration file (`.vmx` file) of the virtual machine.
- 2 Deselect **Time synchronization between the virtual machine and the host operating system** in the VMware Tools toolbox GUI in the guest operating system.
- 3 Run the command `vmware-guestd --cmd "vmx.set_option synctime 1 0"` in the guest operating system.

These options do not disable one-time synchronizations done by VMware Tools for events such as tools startup, taking a snapshot, reverting to a snapshot, resuming from suspend, or VMotion. These events synchronize time in the guest operating system with time in the host operating system, so it is important to make sure that the host operating system's time is correct.

To do this for VMware ACE, VMware Fusion, VMware Player, VMware Server, and VMware Workstation, run time synchronization software such as NTP or `w32time` on the host. For VMware ESX, run NTP in the service console. For VMware ESXi, run NTP in the VMkernel.

Virtual Hardware Clock Configuration

When configuring a Linux guest operating system, if you are given a choice between keeping the "hardware" clock (that is, the virtual CMOS time of day clock) in UTC or local time, choose UTC. This avoids any confusion when your local time changes between standard and daylight saving time (in England, summer time).

Additional Notes

This section provides additional useful notes for Linux administrators who want to optimize their Linux guest operating systems. These notes do not detail step-by-step instructions for fine-tuning your Linux virtual machines, but they offer further insight into understanding the behavior of your virtual machines. This section covers the following topics:

- ["Asynchronous I/O"](#) on page 16
- ["Linux Timer Rates for Virtual Machines on VMware ESX"](#) on page 16
- ["Linux Swappiness"](#) on page 17
- ["Paravirtualization and Virtual Machine Interface"](#) on page 17
- ["Open Virtual Machine Tools"](#) on page 19

Asynchronous I/O

Support for asynchronous I/O in the kernel and corresponding system calls are included in Linux kernel version 2.6. But the API for the appropriate system calls differs from the original functions, so a simple recompilation of existing applications might be insufficient to use the new change. The change requires source code modification that has occurred in different applications at different rates.

Be aware of this issue when considering applications that use asynchronous I/O. In benchmarking, for example, Iometer is a tool that can be used for storage system performance measurements on both Windows and Linux. However, because its source code was never updated to use the new, corrected asynchronous I/O system calls on Linux, the numbers it produces on Linux (using raw devices) are artificially and unfairly low. Consider this when doing Linux-to-Windows comparisons and use asynchronous I/O-aware tools such as `aiostress`, when needed.

Linux Timer Rates for Virtual Machines on VMware ESX

Modern systems use a hardware timer for a variety of fine-grained operations at the operating system level. VMware's virtualization platforms virtualize this timer in the ESX kernel. Because the virtual timer provided to the virtual machine is actually software, it is subject to the same resource restrictions as other processes. The busier the system, the more the timer execution must contend with other hypervisor activities. There are two implications of this:

- When the system is very busy, the software timer might not execute as regularly as it does when the system is less busy and virtual time can fall behind.
- Depending on how frequently the operating system is interrupted by the timer, the hypervisor must do different amounts of work.

The amount of work required to manage the virtual timer is greatest with Red Hat Enterprise Linux 5 SMP systems, which use a clock frequency of 1000Hz and suffer from a multiplicative amount of work because of SMP support. For instance, the following table shows the number of timer interrupts on a 1000Hz Red Hat Enterprise Linux 5 virtual machine:

Table 6. Timer Interrupts for Multiple Virtual CPUs in Red Hat Enterprise Linux 5

Virtual CPU Count	Timer Interrupts per Second per Virtual Machine
1	1000
2	6000
4	20000
8	72000

Table 7. Timer Interrupts for Selected Linux Distributions

Linux Version	Timer Interrupts per Second per Virtual CPU
Red Hat Enterprise Linux 5	1000
Red Hat Enterprise Linux 4	100
SUSE Linux Enterprise Server 10 SP1	250
SUSE Linux Enterprise Server 9	1000
SUSE Linux Enterprise Server 8	100

The amount of work that needs to be done by the hypervisor increases dramatically with the addition of virtual CPUs. In addition, decreasing the timer interrupt rate greatly decreases the work that needs to be done by the VMkernel to virtualize the timer. Red Hat Enterprise Linux 5.1 and Red Hat Enterprise Linux 4.7 offer a Linux kernel that enables reducing the timer rate. By adding the parameter `divider=10` to the boot parameters as described in [“Linux Time Synchronization Recommendations”](#) on page 12, you can reduce the amount of work required of the VMkernel to virtualize the timer by an order of magnitude.

The mainline Linux kernel has moved to a tickless timekeeping model in which, instead of counting periodic interrupts to tell the time, the Linux kernel queries a hardware time source when it needs to know the time. These changes make Linux a much better-behaved operating system in physical as well as virtual environments. Tickless kernels keep better time and reduce power utilization on laptops and other mobile devices where battery life is important. In virtual machines, tickless kernels are ideal because they impose less overhead on the hypervisor and keep very accurate time. To serve our customers better, VMware has worked directly with the mainline Linux kernel as well as Canonical, the publisher of the Ubuntu distribution, to implement operating system changes to support tickless timekeeping. Because the migration to a tickless kernel was completed recently, only recent Ubuntu Linux releases incorporate this functionality.

Accurate tickless timekeeping can also be achieved through paravirtualization, in which the operating system is modified to be aware that it is running on a hypervisor. The Virtual Machine Interface (VMI), developed by VMware in conjunction with the Linux community and now part of the mainline Linux kernel, is a paravirtualization solution that, in addition to providing higher performance, ensures accurate timekeeping in virtual machines. VMware worked with Novell and Canonical to support VMI-enabled guest operating systems, such as 32-bit SUSE Linux Enterprise Server 10 SP2 and Ubuntu 8.04 LTS, for ESX 3.5 Update 2 and later. The 32-bit versions of Ubuntu 7.10 and 7.04 support VMI on ESX 3.5 and later. VMware has encouraged Red Hat to incorporate support for VMI into its releases. Outside of VMI, VMware also worked with Novell to implement tickless timekeeping through paravirtualizing 64-bit SUSE Linux Enterprise Server 10 SP2.

Linux Swappiness

In a virtualized environment, the frequency of swapping applications in and out of memory can have adverse effects on a Linux virtual machine's responsiveness after long idle periods. You can adjust this frequency, or "swappiness," by setting `/proc/sys/vm/swappiness` to a range between 0 and 100. Tune Linux swappiness down to make Linux virtual machines more responsive using a command inside the Linux virtual machine similar to the following:

```
echo 0 > /proc/sys/vm/swappiness
```

If you want this setting to persist, add the following line to `/etc/sysctl.conf` in the Linux virtual machine:

```
vm.swappiness=0
```

When you lower the swappiness setting, idle processes in the virtual machine are not penalized as Linux, by default, tries to swap them out. Swapping and page faults in a virtual machine are more expensive than in a native environment. For example, a large, idle application with its own large cache would not benefit from being swapped out of memory. If it is swapped out, it takes a long time for its pages to swap back into memory when the application is accessed. The impact of changes to the swappiness setting is more pronounced in Linux virtual machines than in Linux distributions running natively.

Paravirtualization and Virtual Machine Interface

ESX 3.5 introduces support for guest operating systems that use VMware's paravirtualization standard, Virtual Machine Interface.

Two techniques for virtualizing unmodified guest operating system kernels are binary translation and hardware virtualization. A comparison of these "full virtualization" techniques can be found in the paper "A Comparison of Software and Hardware Techniques for x86 Virtualization." A different technique is paravirtualization, which requires modifications to the guest operating system and can be used in conjunction with the first two techniques or on its own. For additional details on this technique, see the paper "Performance of VMware VMI." For links to these papers, see "[Resources](#)" on page 19.

Paravirtualization

Paravirtualization is a technique in which a modified guest operating system kernel communicates to the hypervisor its intent to perform privileged CPU and memory operations. This technique reduces the work required of the hypervisor, thus making it simpler than a binary translation hypervisor. Although paravirtualization does not eliminate virtualization overhead, it can improve guest operating system performance.

The idea of guest-host interaction is not a new concept. It has long been a part of VMware products in the form of VMware Tools. For example:

- The VMware SVGA driver shares data structures with the hypervisor to allow faster screen updates.
- The VMware high-performance virtual Ethernet driver, `vmxnet`, shares data structures with the hypervisor to reduce CPU overhead.
- The VMware “balloon driver” is used by the hypervisor to control the guest operating system’s memory usage.
- The VMware Tools service enables time synchronization between host and guest.

While such guest-host communication provides improved performance and can be classified as paravirtualization, none of these examples involve changes to the underlying guest operating system kernel. However, in order to paravirtualize the CPU and the memory management unit, changes to the guest operating system kernel are required. Open-source operating systems, such as Linux, allow us to make such changes.

Virtual Machine Interface

Early paravirtualization required operating systems and applications to be certified on many different kernels—some for native hardware and others to run on a hypervisor. Also, the lack of a standard guest-host interface led to frequent interface changes, which in turn caused version dependencies between the guest operating system kernel and the hypervisor. In order to address these issues, VMware proposed a new guest-host interface, called Virtual Machine Interface, which defines a set of hypercalls an operating system can use to communicate with the hypervisor. For details, see the Paravirtualization API Version 2.5 specification. For a link, see [“Resources”](#) on page 19.

The standardized interface provided by VMI allows the guest operating system kernel and the hypervisor to evolve independently. The VMI specification also makes it possible for other vendors to enable their hypervisors to support guest operating systems that use VMI.

VMI was designed to abstract native hardware. This feature, called transparent paravirtualization, allows a VMI-enabled kernel to run both on native hardware and on hypervisors that support the interface with no additional modification to the operating system kernel.

VMI code is included in 32-bit Linux mainline kernels 2.6.22 and above. The Ubuntu Linux distribution includes VMI support in version 7.04 (Feisty Fawn) and later. Novell includes VMI support in SUSE Linux Enterprise Server 10 SP2. To enable VMI, see [“Appendix C: Enabling VMI in a Linux Kernel and in ESX 3.5”](#) on page 22.

VMI Performance Benefits

The VMware implementation of VMI offers a number of performance and resource-utilization benefits:

- The syscall entry and exit path is faster. This speeds up syscall-dominated workloads.
- VMI-enabled Linux kernels by default use an alternate timer interrupt mechanism that results in reduced physical CPU consumption, especially when the virtual machine is idle, and in more accurate timekeeping, even when running many virtual machines.
- Because the guest kernel communicates to the hypervisor its intent to perform memory management unit (MMU)-related operations, MMU virtualization overhead is reduced. Depending on the workload, this can have varying performance benefits.
- SMP virtual machines running VMI-enabled operating systems use shared shadow page tables. As a result they have less memory space overhead than those running non-VMI-enabled operating systems. For more information on shadow page tables, see the paper “ESX Server Architecture and Performance Implications.” For a link, see [“Resources”](#) on page 19.

When a workload runs in user mode, the VMware virtual machine in which it is running is in direct-execution mode. Because directly executed code already runs at native speeds in both binary translation and VMI-style paravirtualization, workloads that spend the majority of their time in user mode gain only modest performance improvements from paravirtualization.

Open Virtual Machine Tools

As part of the Open Virtual Machine Tools project, VMware has open sourced portions of VMware Tools for Linux, FreeBSD and Solaris and moved to a collaborative development model with the open source community. This move allows the use, modification, and redistribution of the code being open sourced. Third-party developers outside VMware can contribute through community development and collaborate with VMware engineers for rapid innovation, development of ports to additional guest operating systems for which VMware Tools packages do not yet exist, and overall ease of maintenance. Customers and partners can also simplify the incorporation of VMware Tools installation and upgrade as part of their guest operating system lifecycle management and redistribution as open source software components.

Open Virtual Machine Tools also will aid Linux distributions in bundling Open Virtual Machine Tools for a better “out of the box” guest operating system experience on VMware platforms by allowing the distributions to provide Open Virtual Machine Tools that match the guest operating system’s kernel version. This is especially useful for Linux releases that ship between releases of VMware platform software.

In addition, the virtual appliance market is driving needs for support for newer, nontraditional guest operating systems. This support requires a rapid development model made possible through community involvement and development. In general, now that virtualization has become mainstream, opening VMware Tools facilitates collaborative development involving the wider community.

Currently, only the following components that are part of VMware Tools are not part of Open Virtual Machine Tools:

- VMware Tools upgrader
- (experimental) VMware Descheduled Timer Accounting (VMDesched)

The central activity center of the Open Virtual Machine Tools project is at Sourceforge. For a link, see “Resources” on page 19. The Sourceforge page provides links to announcements, technical discussions, administrative processes, documents, development processes, bugs, user groups, contributions, projects, merges, and more. It also allows you to browse, download, and contribute to the source code.

Resources

- “A Comparison of Software and Hardware Techniques for x86 Virtualization”
<http://www.vmware.com/resources/techresources/528>
- Archives for VMware ESX Server 2.x
http://www.vmware.com/download/esx/drivers_tools.html
- *Basic System Administration*
http://www.vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_admin_guide.pdf
- “Enabling VMI with SLES10 SP2 32bit virtual machines on ESX”
<http://kb.vmware.com/kb/1005701>
- *ESX Server 3 Configuration Guide*
http://www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_3_server_config.pdf
- “ESX Server Architecture and Performance Implications”
<http://www.vmware.com/vmtn/resources/433>
- *Guest Operating System Installation Guide*
http://www.vmware.com/pdf/GuestOS_guide.pdf
- “How to enable a virtual machine interface in a Linux kernel and in ESX Server 3.5”
<http://kb.vmware.com/kb/1003644>

- “Improving Guest Operating System Accounting for Descheduled Virtual Machines in ESX Server 3.x Systems”
<http://www.vmware.com/resources/techresources/526>
- Instructions on downloading and installing a driver for the BusLogic virtual SCSI adapter
http://www.vmware.com/download/esx/drivers_tools.html
- “Linux Guest Moved to a System with Different Processor Type Panics During Boot”
<http://kb.vmware.com/kb/1572>
- Open Virtual Machine Tools main page at Sourceforge
<http://open-vm-tools.sourceforge.net/>
- Paravirtualization API Version 2.5 specification
http://www.vmware.com/pdf/vmi_specs.pdf
- “Performance of VMware VMI”
<http://www.vmware.com/resources/techresources/1038>
- “RHEL 5 and FC 7 Guests Installed with Red Hat Virtualization Affects Performance of Virtual Machine”
<http://kb.vmware.com/kb/9134325>
- SUSE Linux Enterprise Server information
<http://www.novell.com/products/server/>
- “Timekeeping best practices for Linux”
<http://kb.vmware.com/kb/1006427>
- “Timekeeping in VMware Virtual Machines”
<http://www.vmware.com/resources/techresources/1066>
- Ubuntu Linux download page
<http://www.ubuntu.com/getubuntu/download>
- VMware Update Manager Release Notes
http://www.vmware.com/support/pubs/vi_pages/vi_pubs_35u2.html
- “Using PXE (Preboot Execution Environment) to Install Guest Operating Systems over a Network”
<http://kb.vmware.com/kb/1162>

Appendix A: Linux Versions Supported on ESX Server

Table 8 lists the Linux and FreeBSD versions supported on specific versions of ESX Server. See the *Guest Operating System Installation Guide* for the most recent compatibility information. For a link, see “Resources” on page 19.

Table 8. Supported Linux and FreeBSD Operating Systems

Guest Operating System	ESX Server
CentOS 5.0	3.0.3 - 3.5 U2
Red Hat Enterprise Linux 5	3.0.2 - 3.5 U2
Red Hat Enterprise Linux 4	2.5.2 - 3.5 U2
Red Hat Enterprise Linux 3	2.0.1 - 3.5 U2
Red Hat Enterprise Linux 2.1	2.0 - 3.5 U2
Red Hat Linux 9.0	2.0 - 2.5.5
Red Hat Linux 8.0	2.0 - 2.5.5
Red Hat Linux 7.3	2.0 - 2.5.5
Red Hat Linux 7.2	2.0 - 2.5.5
SUSE Linux Enterprise Server 10	3.0.1 - 3.5 U2
SUSE Linux Enterprise Server 9	2.5 - 3.5 U2

Table 8. Supported Linux and FreeBSD Operating Systems

Guest Operating System	ESX Server
SUSE Linux Enterprise Server 8	2.0–3.5 U2
SUSE Linux 9.3	2.5.2–2.5.5
SUSE Linux 9.2	2.5.1–2.5.5
SUSE Linux 9.1	2.5–2.5.5
SUSE Linux 9.0	2.1–2.5.5
SUSE Linux 8.2	2.0–2.5.5
Ubuntu Linux 8.04	3.5 U2
Ubuntu Linux 7.10	3.5 U1–U2
Ubuntu Linux 7.04	3.0.2–3.5 U2
FreeBSD 4.11	2.5.4–2.5.5
FreeBSD 4.10	2.5–2.5.5
FreeBSD 4.9	2.5

Appendix B: Command-Line Options for VMware Tools Upgrades

When you install or upgrade VMware Tools, you can enter several command-line options, listed in Table 9. Right-click a virtual machine and select **Install/Upgrade VMware Tools**. Then enter the command-line options in the Advanced field. See *Basic System Administration* for more details. For a link, see [“Resources”](#) on page 19.

Table 9.

Option	Description
-u user	Specifies a user with sufficient privileges on the target virtual machine, including VirtualMachine.Config.*, VirtualMachine.Interact.*, and VirtualMachine.Provisioning.*
-p password	Specifies a password on the command line. If this is omitted, the tool immediately prompts for a password.
-n vmname	The name of the virtual machine to upgrade. This name corresponds to the display name of a virtual machine. Specify multiple virtual machines using multiple -n parameters. The -n option is ignored if -h is specified.
-h host	Attempts to upgrade all the virtual machines on a particular host. Fails if the specified host is not ESX version 3.0 or greater.
-m maxpowerons	On a particular host, power on only this number of virtual machines at a time.
-o port	Specifies the VirtualCenter Server port, if one other than the default port 902 has been configured.
-t maxpowerontime	After the tools upgrade is scheduled on a virtual machine, the virtual machine is powered on and allowed to run through the tools installation process. In most cases, the guest powers off the machine when the process completes. This parameter allows a user to set the maximum amount of time for a virtual machine to be powered on in case the guest is unable to shut down the machine itself.
-s	Skips the tools and does only the virtual hardware upgrade.
-q	Works quietly. Does not produce status or completion messages on shutdown.

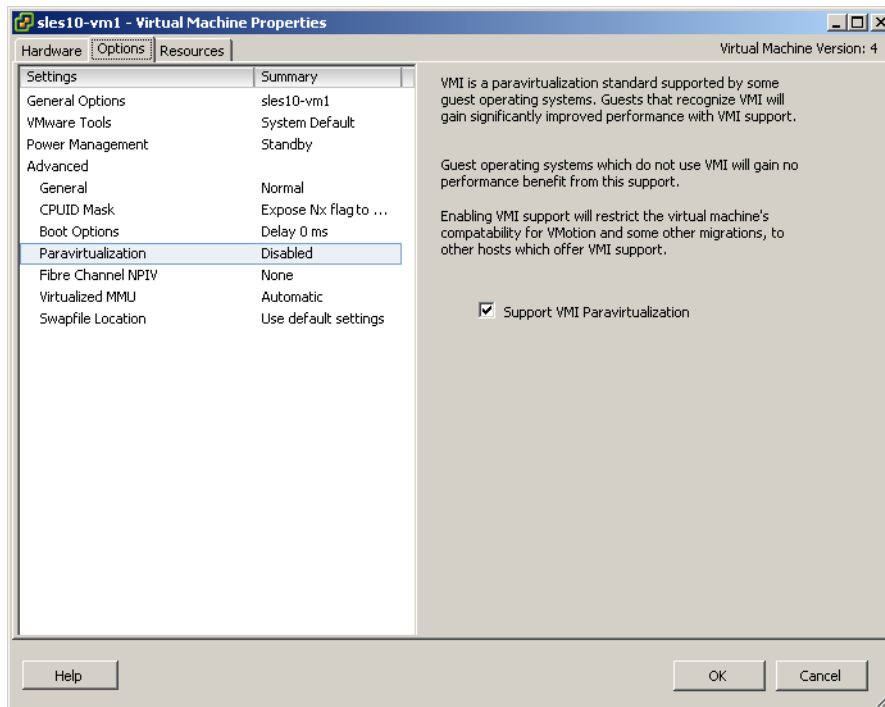
Appendix C: Enabling VMI in a Linux Kernel and in ESX 3.5

To use VMI, you must enable it in your Linux kernel and for the ESX 3.5 virtual machines in which that kernel is running. The following instructions can also be found in the VMware knowledge base article “How to enable a virtual machine interface in a Linux kernel and in ESX Server 3.5.” For a link, see “Resources” on page 19.

Enabling VMI in ESX 3.5

In ESX 3.5, each virtual machine can have VMI either enabled or disabled. To enable VMI for a particular virtual machine, on the Summary tab for that virtual machine, click **Edit Settings**, select the **Options** tab, click **Paravirtualization**, and make sure **Support VMI Paravirtualization** is checked, as illustrated in Figure 1.

Figure 1. VMI Paravirtualization option in ESX 3.5



When VMI is enabled in the virtual machine settings, the `lspci` output in the guest operating system includes a new PCI device (identified as a memory controller), as shown below.

```
linux-iwvp:~ # lspci
00:00.0 Host bridge: Intel Corporation 440BX/ZX/DX - 82443BX/ZX/DX Host bridge (rev 01)
00:01.0 PCI bridge: Intel Corporation 440BX/ZX/DX - 82443BX/ZX/DX AGP bridge (rev 01)
00:07.0 ISA bridge: Intel Corporation 82371AB/EB/MB PIIX4 ISA (rev 08)
00:07.1 IDE interface: Intel Corporation 82371AB/EB/MB PIIX4 IDE (rev 01)
00:07.3 Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI (rev 08)
00:0f.0 VGA compatible controller: VMware Inc [VMware SVGA II] PCI Display Adapter
00:10.0 SCSI storage controller: LSI Logic / Symbios Logic 53c1030 PCI-X Fusion-MPT Dual Ultra320
        SCSI (rev 01)
00:11.0 Ethernet controller: Advanced Micro Devices [AMD] 79c970 [PCnet32 LANCE] (rev 10)
00:12.0 Memory controller: VMware Inc Unknown device 0801 (rev 01)
```

NOTE The VMI PCI device, like other virtual PCI devices, takes up a virtual PCI slot.

The presence of the VMI device in the `lspci` output, however, does not confirm that the guest operating system is running in VMI mode. The following `dmesg` output confirms VMI mode, indicating that VMI is enabled in the virtual machine settings as well as in the Linux kernel.

```
Detected VMI ROM version 3.0  
VMI Timer active.
```

NOTE VMI-enabled kernels are included in some Linux distributions. Check with your Linux vendor to see if a VMI-enabled kernel is available.

If you have comments about this documentation, submit your feedback to: docfeedback@vmware.com

VMware, Inc. 3401 Hillview Ave., Palo Alto, CA 94304 www.vmware.com

Copyright © 2008 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,944,699, 6,961,806, 6,961,941, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, 7,269,683, 7,275,136, 7,277,998, 7,277,999, 7,278,030, 7,281,102, 7,290,253, and 7,356,679; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Revision 20081120 Item: TN-070-PRD-01-01
