

VMWare^(VCP-310)

Virtual Infrastructure Administrator



**Smarter
Training**

This LearnSmart exam manual covers the most important topics you will encounter on the VMWare Certified Professional exam (VCP-310). By studying this manual, you will become familiar with an array of exam-related topics, including:

- Understanding the Virtual Infrastructure
- Installing and Troubleshooting ESX Servers
- Installing and Using VirtualCenter
- And more!

Give yourself the competitive edge necessary to further your career as an IT professional and purchase this exam manual today!

VMware Virtual Infrastructure Administrator (VCP-310) LearnSmart Exam Manual

Copyright © 2011 by PrepLogic, LLC.
Product ID: 012076
Production Date: July 19, 2011

All rights reserved. No part of this document shall be stored in a retrieval system or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein.

Warning and Disclaimer

Every effort has been made to make this document as complete and as accurate as possible, but no warranty or fitness is implied. The publisher and authors assume no responsibility for errors or omissions. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this document.

LearnSmart Cloud Classroom, LearnSmart Video Training, Printables, Lecture Series, Quiz Me Series, Awdeeo, PrepLogic and other PrepLogic logos are trademarks or registered trademarks of PrepLogic, LLC. All other trademarks not owned by PrepLogic that appear in the software or on the Web Site (s) are the property of their respective owners.

Volume, Corporate, and Educational Sales

Favorable discounts are offered on all products when ordered in quantity. For more information, please contact us directly:

1-800-418-6789
solutions@learnsmartsystems.com

International Contact Information

International: +1 (813) 769-0920

United Kingdom: (0) 20 8816 8036

Table of Contents

| | |
|--|-----------|
| Abstract | 7 |
| What to Know | 7 |
| Tips | 7 |
| Lesson 1: Understanding the VMware Product Line | 8 |
| What is Virtualization? | 8 |
| <i>Virtualization Benefits</i> | 9 |
| VMware Products | 9 |
| <i>Virtualization Products</i> | 9 |
| <i>Management Products</i> | 10 |
| Understanding the ESX 3.5 Editions | 11 |
| Licensing ESX 3.5 | 12 |
| Comparing VMware ESX 3.5 with Microsoft Hyper-V | 13 |
| Lesson 1: Summary | 14 |
| Lesson 2: Installing ESX or ESXi | 14 |
| System Requirements | 14 |
| Installation Process | 15 |
| <i>Disk Partitions</i> | 15 |
| <i>IP Configuration</i> | 16 |
| <i>Post Installation Configuration</i> | 16 |
| Performing the Installation | 16 |
| Acquiring and Installing ESXi | 18 |
| Troubleshooting ESX Installations | 19 |
| <i>Purple Screen of Death (PSoD)</i> | 19 |
| <i>Using Diagnostic Data Collection</i> | 20 |
| <i>Hardware Issues and Configuration Problems</i> | 21 |
| Lesson 2: Summary | 21 |
| Lesson 3: Configuring the Host Machine | 21 |
| Using the Service Console | 21 |
| Using the VMware Infrastructure Client | 23 |
| <i>Modifying the Memory Settings for the Service Console</i> | 26 |
| <i>Viewing System Logs</i> | 27 |
| <i>Using the Event Viewer</i> | 27 |
| <i>Configuring VM Autostart</i> | 28 |

| | |
|---|-----------|
| <i>Creating Users and Groups</i> | 29 |
| Managing Storage | 30 |
| <i>Fibre Channel</i> | 31 |
| <i>iSCSI</i> | 31 |
| <i>NFS</i> | 31 |
| <i>Direct Attached Storage</i> | 31 |
| Configuring Virtual Networking | 32 |
| <i>Creating a New vSwitch</i> | 33 |
| <i>Modifying vSwitch Settings</i> | 33 |
| Lesson 3: Summary | 34 |
| Lesson 4: Creating Virtual Machines | 34 |
| VM Files | 35 |
| <i>Settings Files</i> | 35 |
| <i>Log Files</i> | 35 |
| <i>Storage Files</i> | 35 |
| Virtual Hardware | 36 |
| Creating a VM | 39 |
| <i>Creating a VM: Step-by-Step</i> | 40 |
| Modifying a VM | 41 |
| Deploying a Linux VM | 42 |
| <i>Creating an OpenSUSE Linux 11 VM</i> | 43 |
| <i>Installing OpenSUSE 11</i> | 45 |
| <i>Installing the VMware Tools in OpenSUSE</i> | 48 |
| Deploying a Windows VM | 49 |
| <i>Creating a Windows Server 2003 VM</i> | 49 |
| <i>Installing Windows in the VM</i> | 50 |
| <i>Installing the VMware Tools in a Windows VM</i> | 52 |
| VMware and Hardware Emulation | 53 |
| Lesson 4: Summary | 53 |
| Lesson 5: Understanding and Installing VirtualCenter 2.5 | 54 |
| VirtualCenter Overview | 54 |
| <i>VirtualCenter Components</i> | 54 |
| <i>The VirtualCenter Database</i> | 56 |
| VirtualCenter Requirements | 57 |

| | |
|---|-----------|
| <i>Database Requirements</i> | 58 |
| <i>Using SQL Server Express Edition</i> | 59 |
| <i>Database Size Planning</i> | 59 |
| Installing VirtualCenter | 59 |
| Common Installation Problems | 63 |
| <i>Database Connectivity Problems</i> | 63 |
| <i>Port Conflicts</i> | 63 |
| <i>VI Client Connection Problems</i> | 64 |
| Updating VirtualCenter | 64 |
| Running VirtualCenter in a VM | 64 |
| Configuring the Inventory | 64 |
| <i>Creating a Logical Structure</i> | 65 |
| Clusters | 66 |
| <i>DRS Automation Levels</i> | 67 |
| <i>HA Settings</i> | 68 |
| Adding ESX Hosts | 68 |
| Configuring Settings | 70 |
| <i>Configuring Statistics</i> | 70 |
| <i>Creating Alarms</i> | 71 |
| <i>Scheduling Tasks</i> | 74 |
| Lesson 5: Summary | 77 |
| Lesson 6: Power Administration | 77 |
| Using Templates | 77 |
| Web Access | 79 |
| VMware Security | 81 |
| <i>Creating Roles</i> | 81 |
| <i>Creating Domain Users</i> | 82 |
| <i>Add Permissions</i> | 84 |
| Backup Solutions | 86 |
| <i>Backing Up VMs</i> | 86 |
| <i>Using VMware Consolidated Backup</i> | 87 |
| High Availability | 87 |
| VMware Converter Enterprise | 88 |
| Lesson 6: Summary | 88 |

| | |
|---|------------|
| Lesson 7: Performance Management | 89 |
| Resource Management | 89 |
| <i>Resource Pools</i> | 92 |
| Performance Monitoring..... | 93 |
| <i>Monitoring ESX Host Performance</i> | 94 |
| <i>Monitoring Windows Server VMs Internally</i> | 96 |
| Performance Tuning | 104 |
| Guided Consolidation | 104 |
| Lesson 7: Summary | 108 |
| Appendix: VCP 3.5 Objectives Map | 109 |

Abstract

This Administration Manual is intended to help a candidate prepare for the administration of VMware infrastructures in real-world environments. Furthermore, careful study and application of the material in this administration manual will help prepare VCP-310 test takers for a better score.

What to Know

The VCP-310 on VMware Infrastructure 3 is intended to demonstrate an administrator's ability to set up, administer, support and maintain a Virtual Machine infrastructure using VMware's proprietary software. The exam covers the virtual infrastructure, installing ESX servers, Networking with VMware, installing and using VirtualCenter, creating and managing virtual machines, access control for VMware, resource management, data protection and much more. It is important for someone preparing for this exam to be aware of the objectives for the exam, and to have an understanding of the basics of virtual machines and their applications.

Tips

This Exam Manual is not intended to be an all-inclusive reference for technicians. Rather, this could be considered the "Cliff's Notes" to a very broad area of study. It is advised that a technician with little to no experience in the field preparing for this exam purchase a VMware Reference Guide and, for those candidates seeking the VCP-310 certification, sign up for the boot camp as soon as possible. The Exam Manual should be used for quick reference and study. Also, it is recommended that the technician also purchase the video training that serves as a companion to this Exam Manual.

Lesson 1: Understanding the VMware Product Line

In this lesson, you will first learn what virtualization is. Next, you'll explore the products offered by VMware that provide virtualization services.

What is Virtualization?

Virtualization is a *hardware abstraction* solution that allows *multiple guest* operating systems to run *simultaneously* on a *single physical host*. Each of these qualifying factors will be explored in more detail below:

- Hardware abstraction is a technology that allows different operating systems to run on the same hardware without the need for direct hardware access. The hardware access is provided through a virtualization layer, which is usually called the *hypervisor*.
- While you can run a single operating system within a virtual implementation, the real benefit is realized when multiple operating systems run concurrently. For example, you can run two Windows Server 2003 virtual machines (VMs) alongside a Linux VM – all on the same physical host.
- Virtualization is different than dual-booting. Dual-booting allows you to run different operating systems on the same hardware, but these operating systems do not run simultaneously. With virtualization, the operating systems do run simultaneously, which allows for physical machine consolidation.
- All of the VMs may run on a single physical host machine. Of course, you will need more power in the host machine than would be required to run any of the VMs natively, but the consolidation provides several benefits that will be explained later.

A key term, referenced previously, is the word *hypervisor*. The hypervisor is the layer between the hardware and the VMs. Two basic types of hypervisors exist: bare-metal and hosted. A bare-metal hypervisor (also known as a Type-1 hypervisor) runs directly on the hardware and the VMs run above the hypervisor layer. A hosted hypervisor (also known as a Type-2 hypervisor) runs on top of the host operating system (usually Windows or Linux) and then the VMs run above the hypervisor layer. The bare-metal hypervisors outperform the hosted hypervisors on like hardware. Figure 1.1 provides a visual representation of the two hypervisor classes.

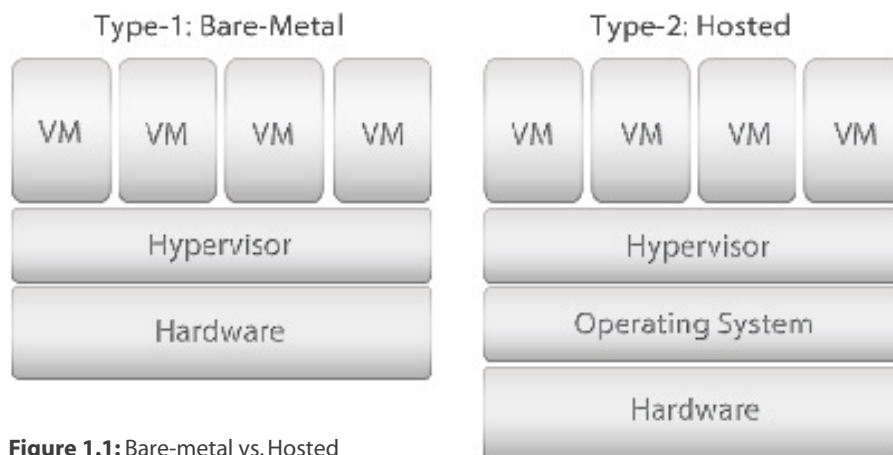


Figure 1.1: Bare-metal vs. Hosted

Examples of bare-metal hypervisors include:

- VMware ESX 3.5 and ESXi
- Microsoft Hyper-V

Examples of hosted hypervisors include:

- VMware Server
- Virtual Server 2005 R2
- VMware Workstation
- Virtual PC 2007

Virtualization Benefits

Virtualization provides several benefits, including:

- **Reduced hardware costs.** You can implement the needed logical servers and services on fewer physical devices.
- **Increased hardware utilization.** Under-utilized servers are no longer an issue. Several currently under-utilized servers can be combined on a single virtual host.
- **Improved energy efficiency.** Energy efficiency is improved from two perspectives. First, fewer physical devices are needed and, therefore, less power is consumed. Second, fewer physical devices are generating heat and, therefore, less cooling is needed.
- **Improved administration.** Virtual servers are easier to administer from a hardware perspective. If the physical host machine should fail, all VMs are simply placed on a new host machine and brought online. Recovery time is usually less than an hour even without the advanced virtual infrastructure (VI) solutions VMware provides.
- **Simplified disaster recovery.** Disaster recovery is simplified as you can implement a hot failover site that is completely driven by virtualization. Since very little is required in the way of hardware, the entire virtual network can be up and running in minutes or hours.

Additional benefits may be realized in unique scenarios, but these are the common benefits for all organizations.

VMware Products

The VMware product line is extensive. As the leader in virtualization solution, VMware provides technologies for both small-scale and large-scale installations. To help you understand the product line best, we'll look at the products from two perspectives: virtualization products and management products.

Virtualization Products

The VMware virtualization product line includes several solutions for different scenarios. Some products are used to provide desktop virtualization and others provide powerful server virtualization solutions.

VMware Workstation

VMware Workstation is a hosted hypervisor solution that is intended for desktop virtualization and test environments. VMware Workstation runs on Windows 2000 through Windows Vista as of version 6.5. Additionally, many distributions of Linux support VMware workstation including Red Hat, SUSE, openSUSE and Ubuntu. VMware Workstation is cheaper than ESX 3.5, but it does not perform as well since it is a hosted virtualization solution.

VMware Server

VMware Server is an earlier virtualization product offered by VMware. It is a hosted product that very closely mirrors the current VMware Workstation product. It has been replaced with ESXi and ESX 3.5 as the recommend virtualization solution for virtual servers. However, it is still available and is a free product. VMware Server 2.0 supports modern operating systems like Windows Server 2008 and Vista.

NOTE: VMware Server is only supported on Windows *server* host operating systems; however, you can install it on several distributions of Linux and it can host Windows client operating systems as guest VMs.

ESX

ESX 3.5 and ESXi are the bare-metal virtualization solutions offered by VMware. With the release of vSphere, ESX moved to version 4.0. The differences between ESX 3.5 and ESXi are important. ESX 3.5 is a licensed product and ESXi is free. ESX 3.5 includes a Service Console (SC) for local management of the host machine and ESXi lacks this SC. ESXi has a smaller footprint of only 32 MB and is managed by the vCenter management console. Don't let the word free fool you, however. ESXi is managed using licensed remote tools. In other words, you really get a smaller footprint VM host with ESXi, but you still pay licenses in order to manage it. VMware Server 2.0, however, is completely free – even from a management perspective.

Management Products

The manager products VMware offers are part of the Virtual Infrastructure (VI) 3 suite or part of the new vSphere suite.

VMware Virtual Infrastructure 3

The Virtual Infrastructure (VI) 3 suite includes several components that work together to form an enterprise-ready virtualization solution. The ESX 3.5 and ESXi components may be implemented independently, with Virtual Center 2.5 as the core of the enterprise management solutions.

The following components are considered part of the VI 3 suite:

- **ESX 3.5 and ESXi:** ESX 3.5 and ESXi, when used within the VI 3 suite, are still the same products; however, they are now managed through the central vCenter 2.5 application.
- **Virtual Center 2.5:** Virtual Center 2.5 has been renamed to vCenter with the release of vSphere, but the 2.5 version is still the same product. With vCenter 2.5, you can use VMotion to move VMs between hosts, use Storage VMotion to move the VM from one storage location to another without downtime, use VMware Converter to convert physical machines to virtual machines (P2V), implement the Distributed Resource Scheduler (DRS) to allow for dynamic VM-to-host assignments based on resource availability, and support Update Manager so that both the host and VMs can be updated automatically.

VI 3 is still the most popular version of VMware's virtualization management products. Of course, as more organizations move to vSphere and as new implementations are performed, this will begin to change.

New in vSphere

vSphere, the new virtual infrastructure management suite released in 2009, includes updates to the following components:

- ESXi and ESX 3.5 have been updated to version 4.0
- vCenter is now on version 4.0
- ESXi embedded is supported

vSphere also introduces several important new features as outlined below:

- Transparent page sharing allows identical memory pages to be stored once for all VMs. For example, if you are running the SQL Server services on two VMs within the same host, it's very likely that several memory pages will be identical between these two VMs. The new versions of ESX and ESXi support storing these pages once and sharing them between the two VMs.
- VMDirectPath allows VMs to directly access the network and storage devices to improve performance in these heavily CPU intensive operations.
- Memory support for VMs has improved from 64 GB to 255 GB and the total number of supported VMs has increased from 170 to 256 on a single host.
- Virtual switches may now span across ESX/ESXi hosts.
- Hot-add for virtual disks is now supported. You can add a new virtual disk in a VM like a hot-add physical disk in physical servers.
- In addition to hot-add for disks, you can hot-add CPUs and memory without user disruption.
- vSphere improves the support for multiple processors in a VM from 4 to 8.

Understanding the ESX 3.5 Editions

VMware ESX 3.5 comes in three different editions. When you purchase and implement ESX 3.5 (ESXi is free), you must understand the features available in each edition. The following table provides a breakdown in the supported features by edition.

| ESX 3.5 Feature | Foundation | Standard | Enterprise |
|------------------------|-------------------|-----------------|-------------------|
| VMFS | Yes | Yes | Yes |
| Virtual SMP | Yes | Yes | Yes |
| VirtualCenter Agent | Yes | Yes | Yes |
| Update Manager | Yes | Yes | Yes |
| Consolidated Backup | Yes | Yes | Yes |
| VMware HA | No | Yes | Yes |
| VMware DRS | No | No | Yes |
| VMware DPM | No | No | Yes |
| VMotion | No | No | Yes |
| Storage VMotion | No | No | Yes |
| VirtualCenter Server | No | No | Yes |

From this table, you can see that the Enterprise edition is the only full featured edition. Foundation is useful in small environments that do not require high availability (HA) in any way. Standard benefits those middle tier organizations that can benefit from basic HA features, but do not require the more advanced features such as dynamic moves from one host to another.

Licensing ESX 3.5

Like any purchased software application, ESX 3.5 comes with its own licensing requirements. Since it is a server-based product, the licensing is a bit more complex than your typical desktop application. Four licensing modes exist:

- Evaluation mode
- Serial number
- License server
- Host license file

Of course, evaluation mode licenses may not be used in a production environment; however, they are beneficial for testing the product and initial learning. Figure 1.2 shows the license management interface within the Virtual Infrastructure client application. Serial number licensing is available in the ESX 3.5 licensing dialog, but it is not supported by ESX 3.5. This option is only listed for communications with ESXi hosts. Using a license server is best in environments with several ESX hosts. In fact, you cannot use features such as VMotion, DRS or HA without a license server. The final option is the host license file. In this case, a license file is uploaded directly to the ESX host for licensing.

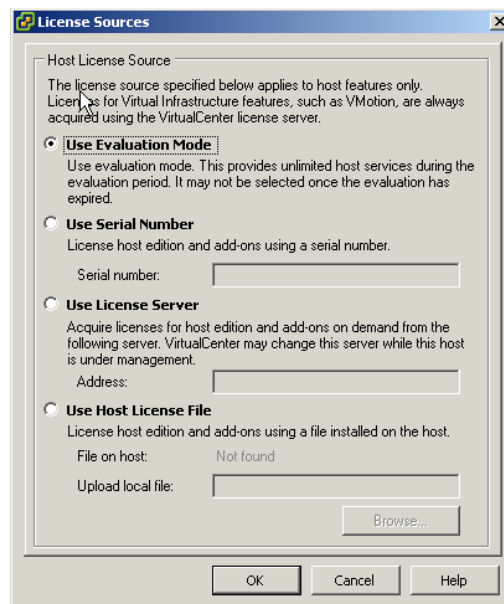


Figure 1.2: ESX Host License Sources

Comparing VMware ESX 3.5 with Microsoft Hyper-V

Microsoft's new virtualization solution, Hyper-V, is also a hypervisor-based technology. However, significant differences exist between ESX and Hyper-V. Administrators must understand these differences in order to make good buying decisions. The following table compares ESX 3.5 with Hyper-V. This comparison is based on the first version of Hyper-V and not Hyper-V R2.

| Feature | ESX 3.5 | Hyper-V |
|---|--------------------------------------|---|
| Hypervisor based virtualization | Yes - Monolithic Hypervisor | Yes - Microkernelized hypervisor |
| Small drive footprint | 32 MB only | No - Hyper-V requires more than 2 GB of drive space |
| Small virtualization kernel | Yes | Yes |
| 64-bit OS support | Yes | Yes |
| Quick Migration (short downtime) | Yes - with added purchase | Yes - with initial license of Windows Server 2008 Hyper-V |
| Live Migration (no downtime) | Yes - with added purchase | No |
| Multiprocessor guests | Yes - 2/4 cores with added licensing | Yes - 4 cores |
| Max host memory | 128 GB | Limited by hardware constraints (up to 1 Terabyte of RAM supported) |
| 64-bit hypervisor | No - but supports 64-bit guests | Yes |
| Maximum guest VMs | | 128 |

One very important difference between ESX 3.5 and Hyper-V is the type of hypervisor used. ESX uses a monolithic hypervisor. This description means that the hypervisor includes the device drivers embedded in the hypervisor. Hyper-V is microkernelized, which means that the hypervisor does not include embedded device drivers for most hardware. Instead, Hyper-V uses a "parent" partition, which is itself a VM running on the hypervisor. This parent partition includes the needed device drivers for guest partitions.

Those who favor the Hyper-V hypervisor suggest that the microkernelized architecture provides for wider hardware support - and indeed it does. Those who favor the ESX hypervisor suggest that it is faster since the drivers are embedded in the hypervisor - and indeed it is. You'll have to decide between these two potential solutions for your production environment. Based on their historic track record and stability of code, VMware ESX still holds grip on the market at the time of this writing.

However, in July of 2009, Microsoft submitted code to the Linux open source project that provides support for Hyper-V within the kernel of Linux. This offering means that, within months, many Linux distributions will run as "aware" operating systems in Hyper-V. The result will be much better performance and stability for Linux VMs running on Hyper-V. The lacking Linux support was a major area of complaint in the virtualization community, but only time will tell whether the complaint was real or not.

Lesson 1: Summary

This lesson has shown the vastness of the VMware product line. If you have a need for virtualization, they probably provide a solution. Virtualization is the process of abstracting hardware to allow multiple operating systems to run concurrently on a single physical host. VMware Workstation and Server are hosted virtualization products. ESX and ESXi are bare-metal virtualization products. The VI 3 product suite includes ESX/ESXi and vCenter Server 2.5. This lesson provided a summary of the new vSphere features; however the remaining lessons will cover the implementation and administration of a VI 3 environment.

Lesson 2: Installing ESX or ESXi

ESX 3.5 and ESXi install as the base operating system on a computer. They do not run on top of another operating system such as Linux or Windows. In fact, neither ESX nor ESXi are based on Linux. The Service Console, used to perform basic management tasks on an ESX installation, is based on Linux, but they, ESX, and ESXi hypervisor (VMKernel) are completely proprietary.

Before you can install ESX or ESXi, you must understand the system requirements. Next, you'll want to preview what you will be doing as you install the product. Both topics will be addressed in this lesson. Finally, a step-by-step installation walkthrough is provided in the section titled *Performing the Installation*.

System Requirements

The system requirements for ESX and ESXi are listed in Table 2.1. The most important requirements include dual processors, sufficient memory and sufficient drive space. Since most servers come with at least one NIC, these requirements will not likely be a problem.

| Component | Requirement |
|-------------------|---|
| Processor | Dual processors (not dual core) running at 1.5 GHz or faster; both Intel and AMD are supported |
| Memory | 1 GB minimum; sufficient extra memory for VMs |
| Network Interface | 1 NIC minimum; 2 NICs recommended – one for VM communications (virtual networking) and one for management |
| Storage | Direct attached, Fibre channel, iSCSI, RAID (ESX can install to and boot from SATA, SCSI, IDE or a SAN) |

Table 2.1: ESX Requirements

While these requirements will meet the minimum demands of ESX, they are not practical for a real world implementation. Particularly in the areas of processor and memory, changes must be made. To determine the real processor and memory requirements for a given ESX host, consider the following questions:

- How many VMs will the host run?
- How much physical memory will each host require?
- What processor demands will each host make?
- What is the total demand for all VMs and the host in the areas of memory and processor?

Consider a scenario like the following. You are planning to install an ESX host that will manage three guest VMs. The first VM will run Windows Server 2003 and act as a dedicated domain controller with no other services. The second VM will run Windows Server 2008 and act as a SQL Server 2008 server. This second VM will manage only small databases that are sporadically accessed throughout the day. The final VM will run Windows Server 2008 and act as an IIS web server. Three websites will be managed within the VM. All websites are intranet sites, but one site is accessed at least once every minute.

In a scenario like this, how do you determine the hardware requirements for the host system? The first step is to start with the minimum requirements and then work up from there. For example, you may decide that a 2 GHz processor with 756 MB of RAM will work very well for the first VM acting as a domain controller. Additionally, the second VM may work well with two 2 GHz processors and 2048 MB RAM. Finally, the last VM should work well with a 2 GHz processor and 1024 MB RAM. This information results in the following calculation:

Memory = 1 GB (minimum requirements) + 756 MB + 2 GB + 1 GB

Memory = 4.75 GB

Processor = 3 GHz (minimum requirements) + 2 GHz + 4 GHz + 2 GHz

Processor = 11 GHz

In this case, I would use 6 GB of RAM in the host. For the processor, I would use two dual-core 3 GHz processors. This would give me 12 GHz of virtual processing. Alternatively, I could use a single 3 GHz quad core processor, but – in my experience – this does not perform as well as the two physical processors.

You may not have to allocate 1 GB of RAM for the ESX host in your calculations. VMware suggests that you allocate 272 MB of RAM for the Service Console in ESX 2.5, but I am a big fan of RAM as are most operating systems. You can never hurt yourself by having an extra 700+ MB of RAM.

The processors introduce another interesting fact. In the real world, you can often get by with less in the virtual implementation than you did in the physical implementation. Because we are implementing three servers on a single host, the inactive time on one server may be the active time on another. You may find that the previous three VMs can work just fine on a single dual-core server running at 3 GHz in each core.

Installation Process

Once you have established the minimum requirements for your host, you can begin the installation process. Before I walk you through the installation, let's consider what's involved and some of the key decisions you'll have to make during the installation.

Disk Partitions

Like most operating systems, ESX requires drive partitioning for installation. In fact, ESX server requires three local partitions: a boot partition, a swap partition and an OS partition. The boot partition will be an ext3 compatible partition as will the OS partition. The swap partition will actually be of type swap. The recommended partitions sizes are as follows:

- Boot partition: 100 MB
- OS partition: 5 GB
- Swap partition: 544 MB or greater

In addition to these three required ESX partitions, you will have to have somewhere to install your virtual machine virtual disks. Virtual disks are installed on a VMFS partition and it will need to be large enough to hold the virtual disk data for your VMs. The VMFS partition may be on the local server or a SAN. Finally, you will need a vmkore partition, which is used to store debugging information for technical support and must be 100 MB. An optional partition is the /var/log partition. It is of type ext3 and should be about 2 GB in size.

IP Configuration

Like all servers that use the TCP/IP protocol suite, the ESX host will use an IP address to communicate on the local network and to receive connections from the network. You should be prepared to provide the following parameters:

- IP address: something like 10.10.12.47
- Subnet mask: something like 255.255.0.0
- Default gateway: something like 10.10.12.1
- Primary DNS: something like 10.10.1.10
- Secondary DNS: something like 192.168.12.40
- Host name: something like esxhost.company.local

You may also have to provide a VLAN ID in some networks. Be sure to use static IP addressing for your servers. The last thing you want is a constantly changing IP address. Many administration tools (mostly third party) connect using the direct IP address instead of host names.

Post Installation Configuration

One of the most important tasks you will perform after this installation completes is licensing configuration. You will need to provide a license file or configure a license server depending on your licensing infrastructure.

In addition to licensing configuration, you may also want to enable SSH so that secure management connections can be made to the ESX server. You should also perform a simple verification task to make sure the server is running properly. At the local console, you can press ALT + F1 to enter the Service Console (ESX only, ESXi does not include the Service Console). Alternatively, you can connect to the server from remote with a web browser and you should see the ESX server web page displayed.

Performing the Installation

The following pages will walk you through the installation of ESX 3.5. If you want to perform the installation on a lab computer, make sure that it is compatible with the ESX software.

Performing a Clean Installation of ESX 3.5

1. Boot the target server with the ESX 3.5 CD in the CD drive.
2. Select to boot from the CD, if necessary.

3. Press **ENTER** if you want to use the graphical installer; otherwise, type **esx text** and press **ENTER** to use the text mode installer.

NOTE: The remaining instructions assume you've chosen the graphical installer.

4. On the media test page, click Test if you want to test the media.

NOTE: The ESX 3.5 installation is not a long one. I usually choose not to test the media as the media is usually good.

5. On the Welcome page, click **Next** to begin the installation.

6. On the Select Keyboard page, select your language and click **Next**.

7. On the Mouse Configuration page, select your mouse settings and click **Next**.

NOTE: Don't fret over the mouse settings. They have no impact on the resulting installation. Remember, you have no GUI on the local ESX 3.5 server with which to use a mouse.

8. You will be presented with the Select Installation Type dialog if a previous version of ESX is detected. From this dialog, you can choose to **Install**, which will overwrite the existing installation, or **Upgrade**, which will attempt to maintain the configuration from the existing installation. Since this is a clean installation, we should not see this dialog. Proceed to step 9.

9. The license agreement is displayed. Read and agree to the license agreement and then click **Next**.

10. Choose either Recommended or Advanced for the partitioning method. In this case, we choose **Recommended** and allow the installation routine to build the partitions for us. Click **Next**.

11. You must now select how ESX will boot. We select **From a drive (install on the MBR of the drive)** and click **Next**.

12. Select the appropriate network interface card (NIC) for use by the console.

NOTE: VMs will share the selected NIC unless you create a virtual switch after installation. To prevent VMs from sharing the administrative NIC, a second NIC must be available in the server.

13. Configure the IP settings based on the information you collected before the installation. Click **Next** when you've entered the IP information.

NOTE: During the install, you will usually leave the **Create a default network for virtual machines** box checked. This setting makes the creation of VMs much simpler.

14. You must now select your time zone. Once you've selected the appropriate time zone, click **Next**.

15. Enter a root password. This password is used to logon to the Service Console after installation in root mode. The password must contain at least six characters. After you've entered the password twice, click **Next**.

16. Review the installation configuration settings and, if they are correct, click **Next** to begin installing ESX 3.5 on the server.

17. After the installation is complete, click **Finish**. The system will reboot and you should see a screen similar to the one in Figure 2.1. The initial boot process can take a few minutes.




Figure 2.1: The ESX Server Screen

At this point, the installation is complete. In the next lesson, you will learn how to perform several basic post-installation tasks; however, you should be able to connect to the ESX server from remote by entering the IP address or host name in your web browser on another machine.

Acquiring and Installing ESXi

This lesson focused on the ESX 3.5 installation. You can also install a free version of VMware ESX called ESXi.

Figure 2.2 shows the website after the release of ESXi 4. In the bottom section of the page, you can download older versions. ESXi 3.5 Update 4 is the most widely used version of VMware at the time of this writing in the summer of 2009. Additionally, ESXi 3.5 supports 32-bit hardware and ESXi 4 supports only 64-bit hardware.

Figure 2.2: ESXi Download Page

The ESXi installation is very similar to the ESX 3.5 installation. Note that you will need to acquire a license key for activation even though ESXi is free.

Troubleshooting ESX Installations

Installing any operating system can be problematic. While ESX 3.5 is a virtualization environment, it is still categorized as an operating system. You may experience hardware configuration problems and even crashes. The following sections provide some tips for resolving these issues.

Purple Screen of Death (PSoD)

The PSoD in ESX is comparable to Microsoft's Blue Screen of Death (BSoD) and is shown in Figure 2.3. Like the BSoDs in modern versions of Windows, most PSoDs are related to hardware failures. In fact, you can usually suspect one of two hardware factors: processor and memory.

```

VMware ESX Server [Release build-32039]
Exception type 13 in world 1193:vmware-vmx @ 0x6c4f5a
Frame=0xd Frame=0x36a7e58 eip=0x6c4f5a cr2=0x524ad000 cr3=0x96b3a000 cr4=0x168
eax=0x0 ebx=0xe ecx=0x68000000 edx=0xba0eed es=0xba04028 ds=0x4028
fs=0x0 gs=0x0 ebp=0x36a7ebc esi=0x36a7ed8 edi=0xba00b48 err=0 ef=0x10246
cpu 0 1024 console; cpu 1 1140 mks:OMHQ5; cpu 2 1138 vmm0:OMHQ; cpu 3 1083 vmm0:
OMHQ;
cpu 4 1267 mks:OMHQ5; cpu 5 1107 vmm0:OMHQ; cpu 6 1069 vmm0:OMHQ; cpu 7 1078 mks
:OMHQ5;
cpu 8 1032 idle8; cpu 9 1158 vmm0:OMHQ; cpu 10 1102 vmm0:OMHQ; cpu 11 1165 vmm0:
OMHQ;
cpu 12 1095 vmm0:OMHQ; CPU 13 1193 vmware-vm; cpu 14 1232 mks:OMHQ5; cpu 15 1039
idle15;
0x36a7ebc:[0x6c4f5a]UserObjFDPoll+0x5e(0xba00b48, 0x7, 0x0)
0x36a7ee8:[0x6c4fc7]UserObjPollCleanupWaiters+0x4b(0xba15928, 0xe, 0x0)
0x36a7f40:[0x6c52b8]UserObjPoll+0x1ec(0xba15928, 0xe, 0x12)
0x36a7f74:[0x6daf7d]LinuxFileDesc_Poll+0xad(0xbf5feb84, 0xe, 0x12)
0x36a7fa8:[0x6bf314]User_LinuxSyscallHandler+0x6c(0x36a7fe0, 0x10000023, 0xd2002
8)
0xbf5f4h38:[0x667cfc]CommonTrap+0xc(0x0, 0x0, 0x0)
VMK uptime: 9:10:36:17.943 TSC: 1630888606661150
TSC: 4742056872 cpu0:0)Chipset: 665: Make sure that if 'noapic' is used, it is o
n purpose
0:00:00:31.779 cpu0:1024)PCI: 1650: failed for 000:15.2
0:07:02:41.693 cpu0:1158)APIC: 1265: Lint1 interrupt on popu 0 (port x61 contain
s 0xb0)
Starting coredump to disk using slot 1 of 1... 98766666543210 Disk dump successf
ul.
Debugger is listening on serial port ...
Remote debugger activated. Local debugger no longer available

```

Figure 2.3: PSoD in VMware ESX Server

Like Microsoft's BSoDs, the PSoD information is not the most helpful information, unless of course you speak computer hardware language. However, as you may notice in Figure 2.3, the PSoD does create a disk dump during the crash. These dump files may be helpful in determining the cause of the PSoD.


NOTE: Keep in mind the following two tips. First, the second line of a PSoD tells you what ESX thinks was the cause of the crash. For example, *Exception type 13 in world 1193:vmware-vmx @ 0x6c4f5a* is the problem in Figure 2.3. You can often search for the first few words of line 2 and find discussion forum threads where others have experienced the same problem. The threads often include solutions.

The best solution for PSoD analysis is the diagnostics data and log files gathered during and before the crash. Assuming the ESX host will start up again, you can view the log files from the Virtual Infrastructure client application. You can also perform diagnostic data collection from there.

Using Diagnostic Data Collection

Diagnostic data is used to analyze continuing problems. The data may be sent to VMware for analysis. To gather the diagnostic data using the Virtual Infrastructure client, follow these steps:

1. Launch the Virtual Infrastructure client and connect to the target ESX host.
2. Select File > Export > Export Diagnostic Data.
3. Enter a path for storage of the diagnostic data as seen here:



4. Click OK.

Be patient. The creation of the diagnostic data can take anywhere from 1 minute to 20 minutes. It depends on the speed of your ESX host and the network connections between you and the host. Once the process completes, the Recent Tasks window (in the bottom of the VI client window) will show the status for the Generate Diagnostic Bundles task as Completed.

Hardware Issues and Configuration Problems

In addition to the PSoDs, other hardware issues may be encountered. During installation, the ESX setup program may not be able to detect your storage devices.

You can search for devices by category or by partner name (the vendors partnering with VMware). The results will be a list of compatible hardware. If your hardware is not on the list, you may need to replace the hardware. The best advice is to purchase complete systems that are certified to work with VMware ESX and, specifically, the version you are using.

Lesson 2: Summary

In this lesson, you learned about the hardware requirements of ESX server. You learned that ESX server is its own operating system and that it does not depend on Linux. You also discovered that real-world implementations seldom use the minimum requirements. Finally, you learned to perform a typical installation of ESX 3.5 on a compatible server.

Lesson 3: Configuring the Host Machine

If you've installed ESX server as Lesson 2 described, you should be able to perform all of the tasks covered in this lesson. Lesson 3 will teach you to use the Service Console (SC) and to work with host configuration settings, including storage management and virtual networking.

Using the Service Console

The SC can be used for several basic administration tasks. The SC is based on Red Hat Linux, though it is not a full implementation of Linux. However, the core networking features are included. By default, ESX server does not allow remote SSH (similar to Telnet, but with security) connections to login as root. If you want to root access for remote SSH connections, perform the following steps:

1. Access the local console of the ESX host and press `ALT + F1`.
2. Logon to the local console as root with the password you created during installation.
3. Open the `sshd_config` file with the VI editor by executing the following command:

```
vi /etc/ssh/sshd_config
```
4. Scroll down through the file, using the arrow keys, until you see a line that reads, "PermitRootLogin no"
5. Move your cursor, using the arrow keys, over the "n" in the word "no" and press the delete key twice.
6. Press the "i" key to enter insert mode.
7. Now, type yes so that the line reads, "PermiteRootLogin yes"

8. Press `ESC` to exit edit mode.
9. Type `:wq` and press `ENTER` to write the changes and quit the VI editor.
10. In order for the changes to apply, the service must be restarted. To restart the service, execute `service sshd restart` at the command prompt.

You may prefer to leave root access disabled. In this case, you will need a user account with the ability to access the ESX host through SSH. Follow these instructions:

1. Access the local console of the ESX host and press `ALT + F1`.
2. Logon to the local console as root with the password creating during the installation.
3. Execute the following commands:

```
useradd username  
passwd username
```
4. After executing the `passwd` command, you will be prompted for the user's password. Enter the desired password twice.
5. Connect to the ESX server from a remote machine using an SSH client such as PuTTY.
6. Enter the login name you defined in step 3.
7. Enter the password you defined in step 4.
8. Use the following command to switch to root mode: `su`
9. You will be prompted for the root password. Provide it and you will be in root mode without allowing direct root access.

To validate that you are not accessing the ESX host with the full power of root, perform these steps:

1. Connect to the ESX host from remote with a SSH client such as PuTTY.
2. Use the login created in the preceding exercise.
3. Use the password associated with that login.
4. Execute the command `cat /etc/shadow`; note the access denied message.
5. Execute the command `su` and then enter the root password when prompted.
6. Execute the command `cat /etc/shadow`; note that the information is now displayed.

These six steps show that you can access the ESX host SC as a regular user and then elevate to root level access as needed. This behavior can be compared to that of enable mode in Cisco devices or User Account Control in Vista and Windows 7. To return to regular user mode after entering root mode, simply execute `su username`, where `username` is the name of the original user with which you logged onto the host.

Using the VMware Infrastructure Client

Many administrators will perform all administration from the VMware Infrastructure (VI) client. You can install the VI client by accessing the web server that is built into the ESX host. Follow these steps to install the VI client on a Windows-based machine:

1. Open your web browser and navigate to the IP address of the ESX host machine.
2. If you receive an error message regarding the certificate, choose to continue to the site anyway. This message is normal because the certificate is not yet installed in your local certificate store.
3. Click the “Download VMware Infrastructure Client” link on the page that is displayed. See Figure 3.1.
4. Select to “Save” the download and then choose the Desktop as the save location.
5. When the download completes, close the web browser.
6. Launch the VMware-viclient.exe file on the Desktop.
7. If you see a dialog informing you that the application comes from an unknown source, click Run.
8. Click Next to begin the installation.
9. Agree to the license agreement and then click Next.
10. Enter the user and organization information you desire and click Next.
11. Accept the default installation location and click Next.
12. Click Install to begin the installation.
13. After the installation is complete, click Finish.
14. You can now launch the VI client by double-clicking the new icon on the desktop.

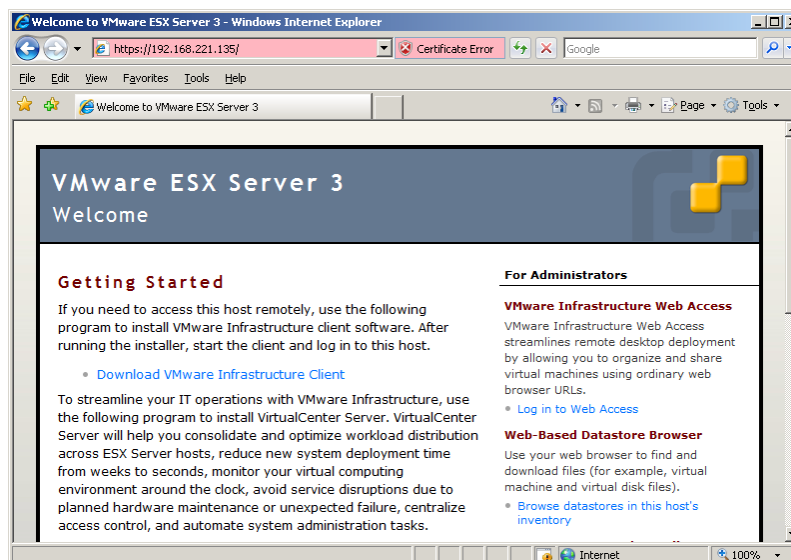


Figure 3.1: The ESX Host Server Web Page

Once the VI client is installed, you must configure its connection to your ESX host. To connect to the ESX host with the VI client, perform these steps:

1. Launch the VI client from the Desktop or Start menu.
2. Enter the IP address, username and password in the dialog as seen here (change the IP address to match your ESX host server):



3. Click Login to connect.
4. If you receive a certificate warning, click Ignore.
5. You should see a screen similar to the one in Figure 3.2.

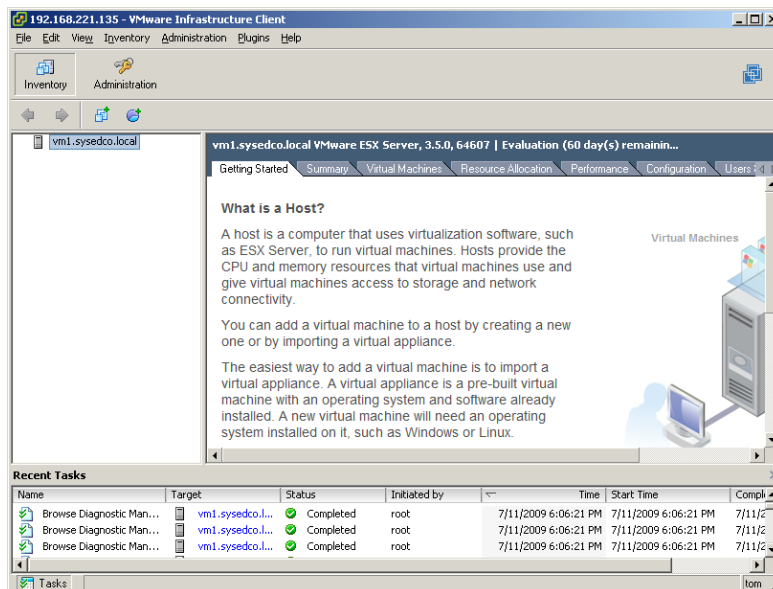
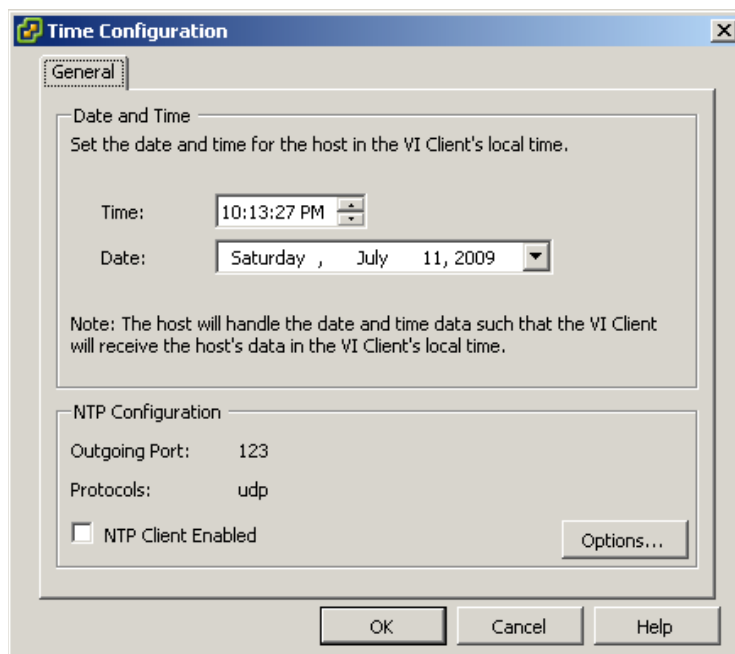


Figure 3.2: Using the VI Client

Once connected to the ESX host, most administration tasks can be performed through the VI client. Be careful not to confuse the VI client with the vCenter or VirtualCenter management tools. VirtualCenter introduces enterprise level features such as templates, VMotion and high availability (HA).

As an example of what you can configure in the VI client, consider the Network Time Protocol (NTP). NTP is used to synchronize or set the time on machines within a networked environment. To configure NTP for an ESC host:

1. Connect to the ESX host using the VI client.
2. Click the *Configuration* tab.
3. Click the *Time Configuration* link in the *Software* pane.
4. Click the *Properties...* link in the upper right corner.
5. Enter the time and date manually, if you do not have access to an NTP server, and then click OK as you are done configuring the time.



6. To use an NTP server, check the box that reads *NTP Client Enabled* and then click the Options button to configure NTP.
7. Click the NTP settings page in the left pane of the NTP daemon Options dialog.
8. Add the proper IP address for your NTP server.
9. Click OK.
10. Click OK.

Remember, this is just one of several administration tasks you can perform from the VI client tool. In addition to the NTP configuration, you may need to perform the following common configuration tasks:


- Modify the Memory Settings for the Service Console
- Viewing System Logs
- Using the Event Viewer
- Configure VM Autostart
- Creating Users and Groups

All of these administrative tasks are covered in the following sections.

Modifying the Memory Settings for the Service Console

The default memory for the SC (remember, this is a modified version of Red Hat Linux) is 272 MB. If you plan to perform all administration from remote and do not intend to install any additional Linux applications in the SC, you may be able to accept the default settings. However, you may need to increase this memory if the opposite is true and you will be performing a lot of administration tasks from the SC or installing extra applications. To increase the memory of the SC, follow these steps:

1. Launch the VI client and connect to the target ESX host.
2. Click the *Configuration* tab.
3. Select *Memory* in the Hardware category.
4. Note the current memory settings.
5. Click the *Properties...* link in the upper right corner.
6. Enter the new memory setting you desire (the minimum allowed is 272 and the maximum is 800) as seen in the following image:



7. Click OK to save your changes.

Note that the changes to SC memory will not take effect until the ESX host is restarted. You can reboot an ESX host by right-clicking on the host in the VI client and selecting Reboot. You'll be informed that the host is not in maintenance mode and given the option to continue or cancel the reboot. You should always ensure that the VMs are properly shut down before rebooting. VMs that support the VMware Tools can be shut down by the ESX host during the reboot process. Those that do not support VMware Tools cannot be shutdown automatically and the ESX host will attempt to save the state of the VM before rebooting.

Viewing System Logs

The system logs provide exhaustive information about the events that take place on your ESX host. To view the logs, follow these steps:

1. Launch the VI client and connect to the target ESX host.
2. Click the Administration button (see the following image).



3. Select the System Logs tab.
4. View the desired log.

Several logs exist including:

- Hostd.log
- Messages
- Vmkernel
- Vmksummary.txt
- Vmkwarning

Using the Event Viewer

The Events tab shows the actions taken and the results within the virtual environment. For example, you will see entries when a VM is started and again when the power on is complete. When you restart an ESX host, the Events are refreshed. Figure 3.3 shows the Events tab for the entire host. Each VM also has its own Events tab for events specific to that VM.

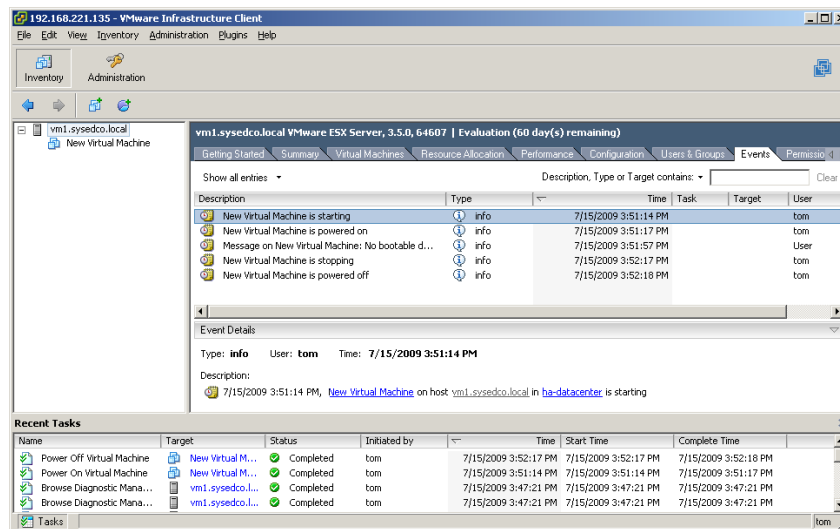


Figure 3.3: Event Viewer in the VI Client

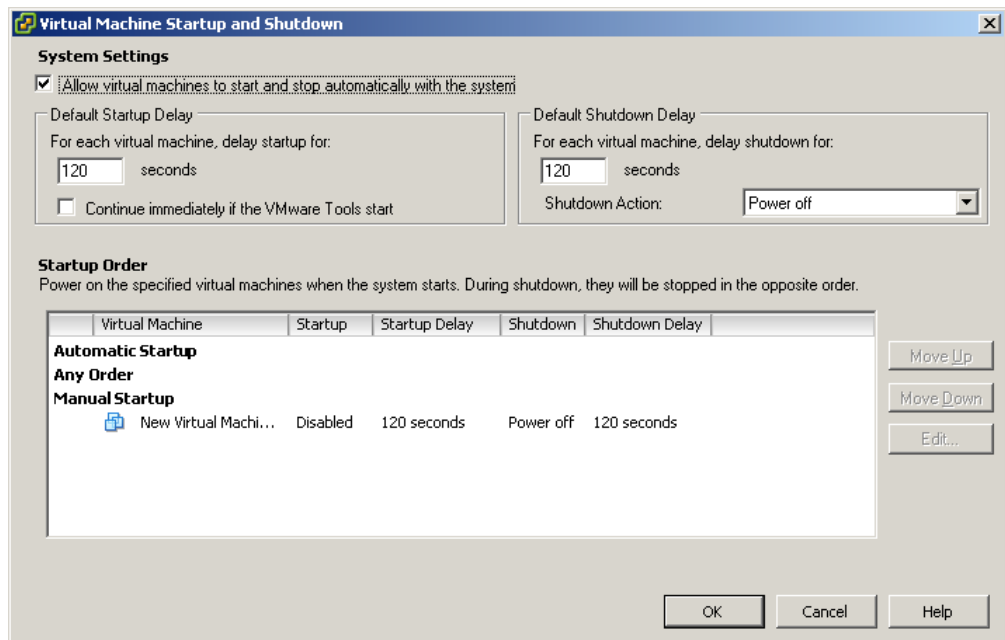
Configuring VM Autostart

A little known fact about ESX hosts is that the VMs are not configured to start automatically when the ESX host starts by default. This makes sense if you consider that many administrators will create several VMs on a given host even if they do not plan to use those VMs continually. Some of the VMs may act as templates, while others are used for testing, while still others are used in production.

The VI client software does provide an interface for configuring autostart parameters for the VMs. You can start the VMs immediately or after some delay. To configure autostart parameters globally, follow these steps:

1. In the VI client, select the ESX host on which the VM that you want to start automatically resides.
2. Click the *Configuration* tab.
3. In the Software section, select *Virtual Machine Startup/Shutdown*.
4. Click the *Properties...* link in the upper right corner.

5. Check the box that reads, "Allow virtual machines to start and stop automatically with the system" as seen in the following image:



6. Configure the startup and shutdown delays as desired.
7. Click the VMs in the Startup Order section and move them up or down to specify the startup order.
8. Leave VMs that should be started manually in the Manual Startup section.
9. Place VMs that can start at any time, but should start automatically, in the Any Order section.
10. Click OK.

When you configure automatic startup for VMs, be thoughtful of the startup order. For example, you would want VMs acting as Windows domain controllers to start before VMs that depend on the domain controllers. Stand alone servers without any dependencies on other servers can be placed in the Any Order section. Test VMs, template VMs and any other VMs that only run part of the time should be placed in the Manual Startup section.

Creating Users and Groups

If you need to allow for additional ESX host administrators and support staff, you may need to create users and groups. Several users exist by default in order to run the various services. To work with users and groups, click the Users & Groups tab in the VI client. You'll see a screen similar to the one in Figure 3.4. From here, you can right-click anywhere in the white space and select Add to add a new user. If you need to add a new group, click the Groups button first and then right-click in the white space and select Add.

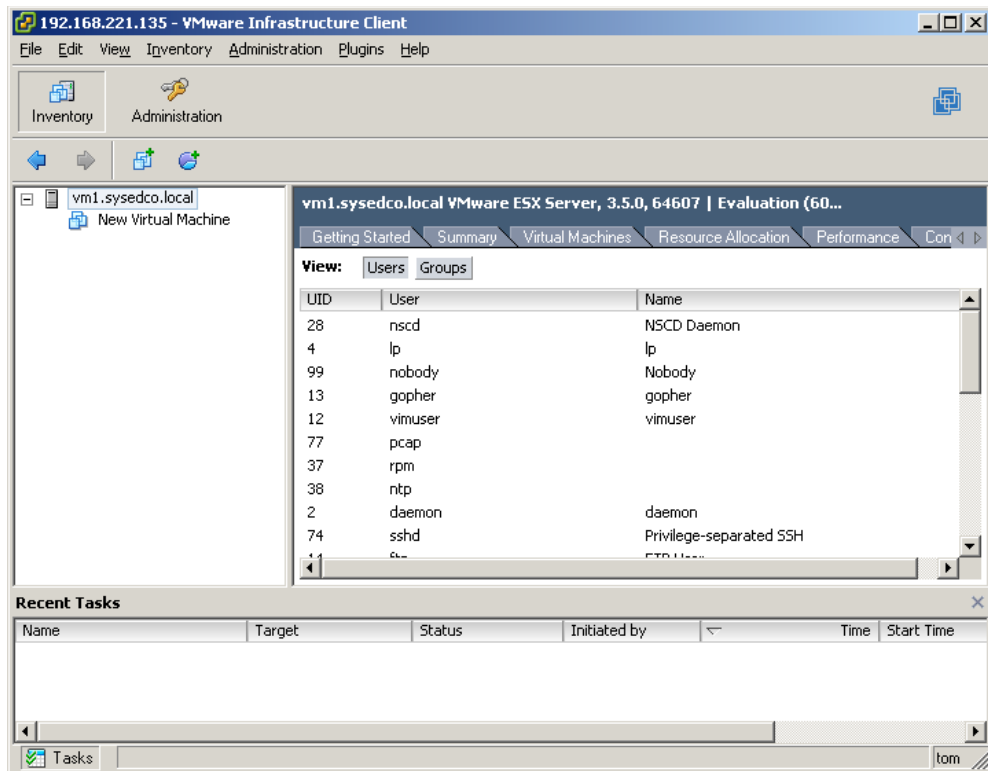


Figure 3.4: Adding and Managing Users and Groups

Managing Storage

One of the key management tasks is storage management. After the installation, storage management is mostly focused on the storage of the virtual disks used by the VM. ESX supports several storage technologies including:

- Fibre Channel
- iSCSI
- NFS
- Direct Attached Storage

Storage Area Network (SAN) protocols such as iSCSI and Fibre Channel are used to access data at the block level using network connections. SANs provide fault tolerance and centralization of data resources for one or more servers. This book will not attempt to cover every aspect of SAN implementation as the topic cannot be covered adequately here; however, a few important notes about the different storage technologies will be addressed in the sections that follow.

Fibre Channel

Fibre Channel (FC) uses SCSI commands over highspeed networks of up to 8 Gbps at this time. ESX server can communicate with FC SANS for both storage and booting from the SAN. FC, or another SAN solution, is used to enable advanced enterprise features such as high availability, VMotion, Storage VMotion and Distributed Resource Scheduling.

FC networks use host bus adapters (HBAs) in the participating servers to communicate with the SAN network. Special devices known as FC switches are used to build the fabric for the SAN. The FC SAN will provide logical unit numbers (LUNs) to the servers so that drives may be added. A LUN may refer to a single drive attached to the SAN or it may refer to a redundant solution such as RAID 0 or RAID 5.

iSCSI

Unlike FC SANs, Internet SCSI (iSCSI) does not use a special dedicated network. Instead, iSCSI uses SCSI commands layered over TCP/IP. iSCSI may communicate using the same network that your users use to access the VMs running on the ESX hosts; however, it is recommended that you build a dedicated gigabit network for the iSCSI communications.

The speeds of iSCSI are the same as Ethernet. As Ethernet moves to 10 Gbps and 100 Gbps, iSCSI performance will increase. For now, FC usually outperforms iSCSI easily.

The performance of iSCSI can be improved by using dedicated iSCSI controllers communicate on a dedicated SAN network. The iSCSI controllers are really special network cards that offload the bottom four layers of the OSI stack from the operating system. This means that everything from TCP down to the physical layer is handled through special optimized device drivers.

NFS

Network Attached Storage (NAS) devices use many different file systems and communications protocols. ESX servers only support the network file system (NFS) for communications with NAS devices. Since NFS uses packet level access to the NAS device (much like accessing a network share), the performance is inferior to either Fibre Channel or iSCSI. NFS and NAS are recommended for test use only and not for production implementations.

Direct Attached Storage

Direct Attached Storage (DAS) is the fancy name for internal or directly connected external drives. For example, an external SCSI or eSATA drive would be considered a DAS device as would an internal SCSI or SATA drive. The main benefit of DAS is the high performance levels that may be achieved. At this time, SATA drives may commonly reach speeds of 3 Gbps and can be configured in RAID arrays to achieve more than that. However, modern SCSI RAID controllers can usually outperform SATA RAID controllers.

The real question, however, is a matter of price. You must meet your budget constraints with your performance demands at some happy medium. For many ESX hosts, a SATA RAID controller with four drives providing internal storage will be just fine. For others, you'll need the power of external 20 plus drive SCSI RAID cabinets. The decision will be based on your budget, in the end, more than your performance desires – if my experience proves true.

Configuring Virtual Networking

An additional configuration concern is the virtual networking in ESX server. ESX server uses the concept of a virtual switch. In fact, you can purchase Cisco virtual switches so that the switches in your VMware servers mirror your physical switches in behavior. However, at this time, the Cisco virtual switches only work with the new VSphere product line.

A virtual switch is very similar to a physical switch. It allows several VMs to connect to it and it links to the physical NIC to bridge data onto the real network. The virtual switches support:

- Avoiding broadcasts to unintended target ports by using VLANs.
- Frame forwarding from one port to another.
- Medium Access Control (MAC) port table management.
- MAC address checking and routing.

However, it is important to note that virtual switches are isolated. This isolation means that the switches cannot share MAC table information, but they also do not require the Spanning Tree Protocol (STP) since loops cannot occur on isolated switches. Additionally, you cannot configure virtual switches to uplink with each other. They can, however, connect to the physical network and then connect, through physical switch ports, to other virtual switches.

VMware supports three types of virtual switches:


- **Internal virtual switches.** The internal virtual switch only allows communications among the VMs on a single ESX host. Think of it like an isolated physical switch that is not connected to any other physical switches or routers. To communicate with a node on that switch, you would have to be connected to a port on the same switch. This behavior is also true for internal virtual switches and only VMs can connect to the switch.
- **Single adapter virtual switch.** The single adapter virtual switch is the normal switch type used to allow VMs to communicate on the physical network and to allow clients from the physical network to communicate with the VMs. One physical NIC port is used to service the single adapter switch.
- **Multiple adapter virtual switches.** The multiple adapter virtual switch, also called NIC teaming, provides fault tolerance because two NICs provide connectivity to the physical network. Think of it like a physical switch with two trunk ports connected to another physical switch. If one trunk port fails, the second can still be used for communications between the switches. The same is true for multiple adapter virtual switches.

During the installation of ESX 3.5, the Service Console is connected to vSwitch0 on the port named vswif0. This connection allows remote SSH and Telnet connections (with modifications to the default configuration) into the Service Console.

Creating a New vSwitch

The VI client can be used to create a new vSwitch on an ESX host. To create a new vSwitch, follow these instructions (NOTE: You will need an available and unused NIC in the ESX host to perform this operation.):

1. Launch the VI client on the Windows machine you use for ESX host administration.
2. Connect to the target ESX host (the one on which you wish to create the new vSwitch).
3. Click the **Configuration** tab and select the **Networking** page. You should see your current vSwitch configuration. It will look something like the following.



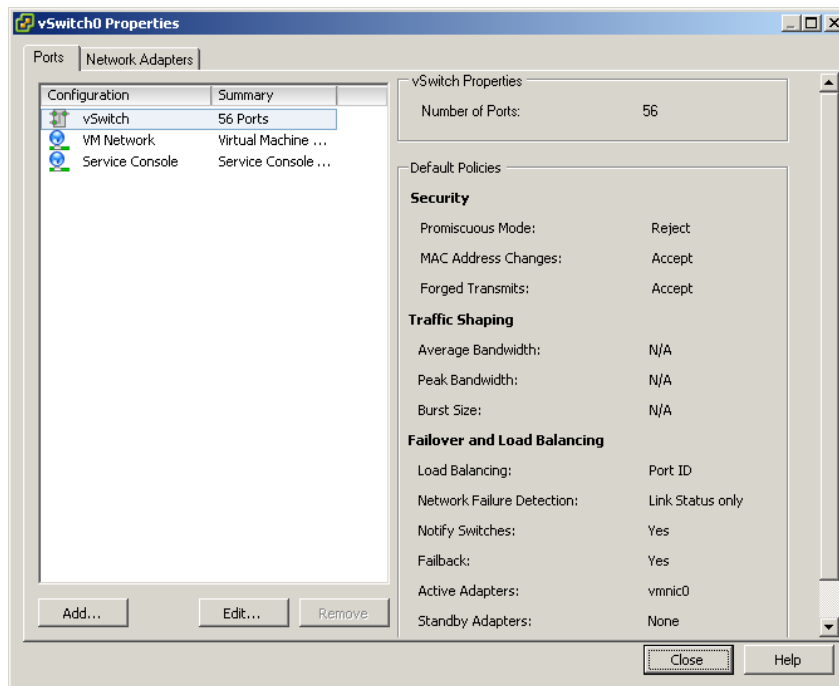
4. Click the **Add Networking** link in the upper right corner. From here, you can add a Virtual Machine, VMKernel or Service Console vSwitch. We want a new vSwitch for VMs, so select **Virtual Machine** and click Next.
5. Leave the default setting, "Create a virtual switch" and click Next.
6. Select the appropriate connection parameters and click Next.

Modifying vSwitch Settings

You may also need to change the settings for a vSwitch for security or performance reasons. To access the vSwitch settings, follow these steps:

1. Launch the VI client on the Windows machine you use for ESX host administration.
2. Connect to the target ESX host (the one on which you wish to create the new vSwitch).
3. Click the **Configuration** tab and select the **Networking** page.

4. Click the **Properties** link for the vSwitch you wish to manage. The following dialog is presented for vSwitch0:



5. From here, you can configure several items including:
 - a. The number of available ports on the switch
 - b. The security parameters for the switch
 - c. Traffic shaping
 - d. NIC teaming
 - e. Speed and duplex of the network

Lesson 3: Summary

In this lesson, you learned how to work with the Service Console to perform basic administrative tasks including enabling remote root access, creating user accounts and restarting services. You also learned how to install the VI client and manage virtual switches and NTP with this client. Additionally, you reviewed the different storage technologies supported by ESX servers.

Lesson 4: Creating Virtual Machines

Without the creation of virtual machines (VMs), ESX hosts serve little purpose. In this lesson, you will learn about the files that are used to create a VM as well as the hardware that may be virtualized for a VM. In addition, you will learn the basic processes used to create and modify VMs and then deploy a Linux VM and a Windows VM.

VM Files

VMs are comprised of several files including:

- VMX
- NVRAM
- LOG
- #.LOG
- VMDK
- FLAT.VMDK
- VSWP
- VMDS

This list may seem long, but each file serves an important purpose. The files can be categorized into three categories: settings, logs and storage. Each category and the associated files are covered in the following sections.

Settings Files

The settings files consist of the VMX and NVRAM files. The *.vmx file contains the VM configuration settings. These settings include the amount of RAM, number of CPUs and so on. The NVRAM file contains the BIOS of the VM. You may have noticed that, when you power on a VM, you have a few moments where you can press F2 to enter the BIOS of the machine. VMware VMs utilize a Phoenix BIOS.

Log Files

The *.log and *-#.log files contain logs for troubleshooting and reference. The *-#.log files are the old logs, which are numbered starting with 1 - for example *-1.log.

Storage Files

The storage files can really be broken into two further categories: VM storage and feature storage. The VM storage files are the VMDK files. The *.vmdk files are the description files used to describe the virtual hard drives. The *-flat.vmdk files are the actual storage containers for the virtual hard drives.

In addition to the VMDK files, ESX servers support swap files - *.vswp - for each VM and the VMs may have snapshots. Snapshots are stored in *.vmsd files.

You will need to understand these basic files and their uses when planning for backups and VM moves and add-ons. Using the tools provided by VMware (such as VirtualCenter), you will not have to worry about the individual files as you move a VM from one host to another; however, you may encounter situations where you have to move VMs manually and, in such scenarios, understanding these files is essential.

Virtual Hardware

Just as physical machines have hardware options, VMs have hardware options. You can choose the number of processors, the amount of RAM, the hard drives and their sizes, and several other hardware settings within the limits imposed by the ESX host. ESX 3.5 hosts impose the following virtual hardware limitations (in addition to the real limits of the hardware on which ESX is installed):

- VMs max out at 64 GB of RAM in ESX 3.5; however, you cannot add more RAM to a VM than that which is available in the ESX host.
- The CPU limit in ESX 3.5 is 4; you may assign 4 CPUs to each VM as long as 4 or more total CPUs are available in the ESX host.
- Two ports are available for both serial and parallel devices.
- Two floppy drives may also be connected as physical hardware or floppy images.
- Up to four IDE CD or DVD drives may be assigned and they may connect to physical hardware or ISO images.
- Six PCI devices are supported within the following constraints:
 - ▶ One of the devices is used for the video card, which leaves five total devices available for NICs or SCSI controllers.
 - ▶ You can have a maximum of four NICs (network interface cards) or SCSI controllers, but you only have five available devices. If you have four NICs in your VM, you can only have one SCSI controller. If you have four SCSI controllers, you can only have one NIC.
 - ▶ The five available PCI devices can be divided among NICs and SCSI controllers in any way you like within the maximum limit of four per device type.

The SCSI adapters supported within VMs may have fifteen hard drives connected to each. The result is a maximum of 60 hard drives on the maximum of four SCSI adapters.

Figure 4.1 shows the main screen used to edit VM settings. The interface is similar to that used in VMware Workstation and Microsoft's Virtual PC product. To view this interface, simply run the VI client and right-click on any VM and select Edit Settings. From here you can adjust the memory settings, number of CPUs and several other parameters.

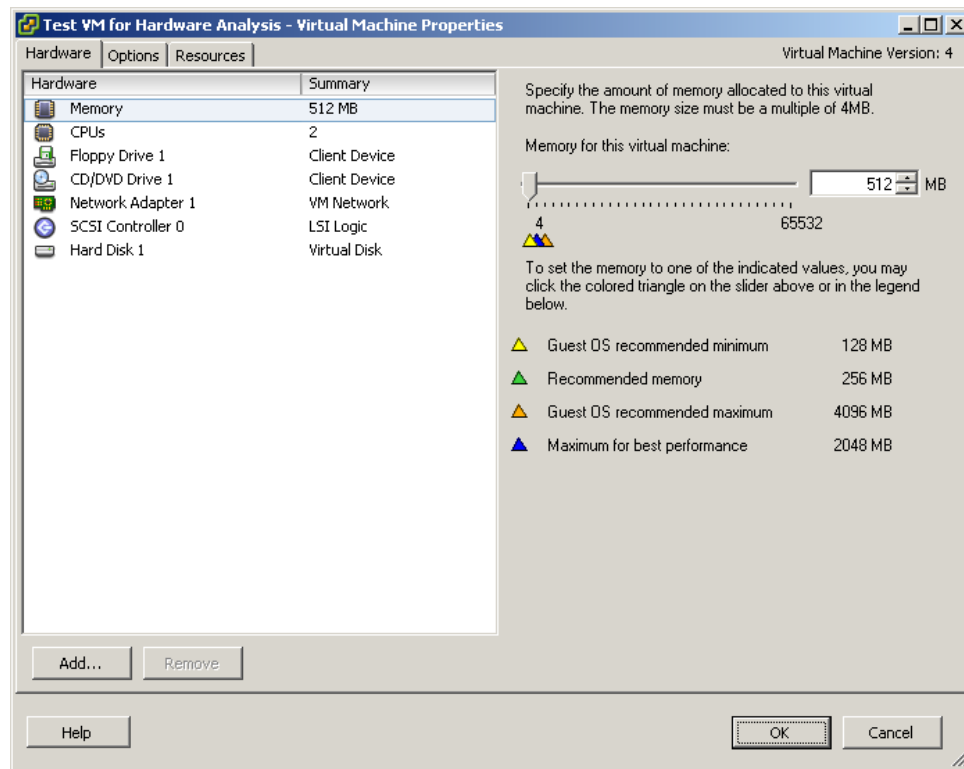


Figure 4.1: VM Hardware Settings

Since several VM configuration settings are available in the BIOS of the VM, it is important to be able to access the BIOS. You can do this in one of two ways. First, you can launch a console connection to the VM and then click inside the VM and press F2. This method can be a little tricky, because of timing issues, but the instructions are as follows:

1. Right-click the VM and select Open Console.
2. In the newly opened console, click the VM menu and select Power On (if the VM is already started, you can restart the guest operating system).
3. Click inside the VM window so that the mouse is captured.
4. When you see the prompt stating, "Press F2 to enter Setup," press F2 quickly. You only have a moment.

If you perform the previous steps successfully, you should see a console similar to that shown in Figure 4.2. You can interact with this BIOS in the same way you would interact with the normal BIOS on a physical machine.

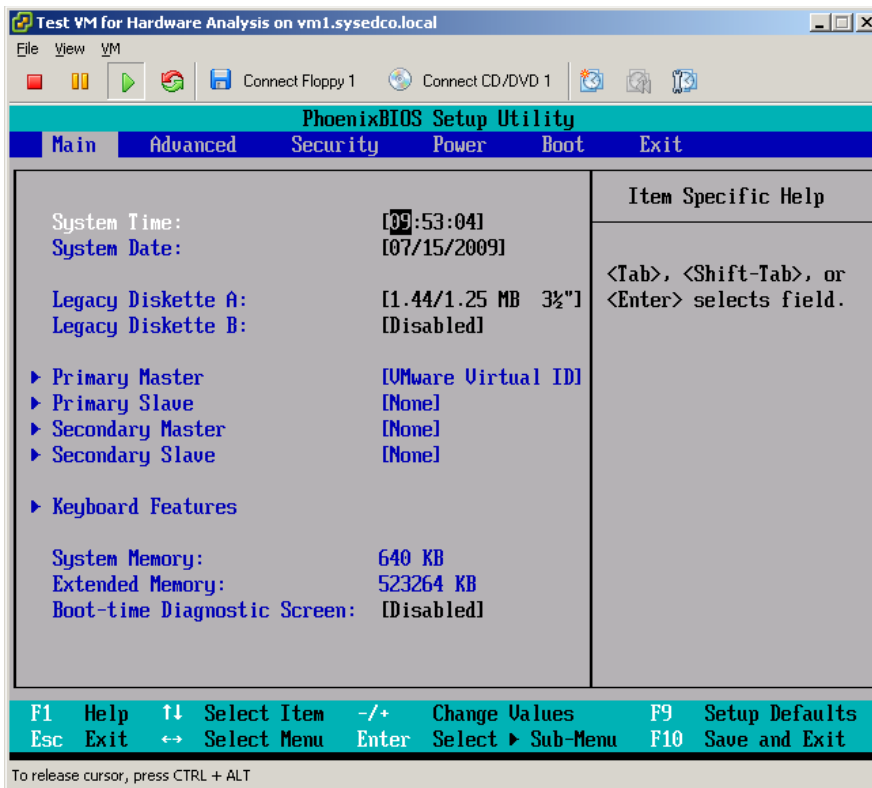


Figure 4.2: VM BIOS in an ESX 3.5 VM

You can also enter the BIOS of a VM by following these steps:

1. Right-click the VM and select Edit Settings.
2. Click the Options tab.
3. Select the Boot Options page.
4. Check the checkbox that reads, "The next time the virtual machine boots, force entry into the BIOS setup screen."

This latter method is more accurate as it is not dependent on you pressing F2 in time. Once in the BIOS, you can configure important hardware settings such as:

- Cache memory settings
- Parallel port modes
- Serial port modes

Among the other BIOS settings, few will impact the operations of the VM. However, you may be required to adjust memory settings for some operating systems.

Creating a VM

Creating a VM begins with planning rather than actions. It is better to create the VM you need from the start rather than to continuously tweak a VM to support the features you need. The following questions should help you better prepare to create a VM:

- What operating system will the VM run?
- How much memory will be required?
 - Remember to consider the operating system and the installed applications or services.
- How many processors will be required?
- How much hard drive space is needed and how fast must it be?
 - Remember that DAS is faster, but does not offer high availability like Fibre Channel.
 - Fibre Channel is usually faster than iSCSI.
- Do you need floppy drives, serial ports or parallel ports?
- How many CD/DVD drives will be needed?

Do these questions sound familiar? If you've ever had to select a physical server, they should be recognizable. You must plan for a VM in the same way that you plan for a physical server.

If anything, you must do more planning because you also have to consider the answer to the following important question: What other VMs are already running on the target ESX host? You would never have to worry about the "other" running operating systems on a traditional physical server; but you must consider the impact of multiple VMs on a single host now.

Interestingly, one of the most common mistakes made when creating a VM is to oversize the VM. For example, the administrator may create a VM and assign it 2 GB of RAM when it really only needs 1 or 1.5 GB of RAM. Oversizing is a big problem because you cannot get those resources back without reconfiguring the VM. If three VMs are running and all three VMs use 2 GB of RAM instead of the 1.5 GB they require, they are using 6 GB of RAM total. If the host machine has only 8 GB of RAM and you now need a new VM requiring 3 GB, you cannot create this VM. If the first three VMs were using the 1.5 GB they require, you could indeed create this fourth VM.

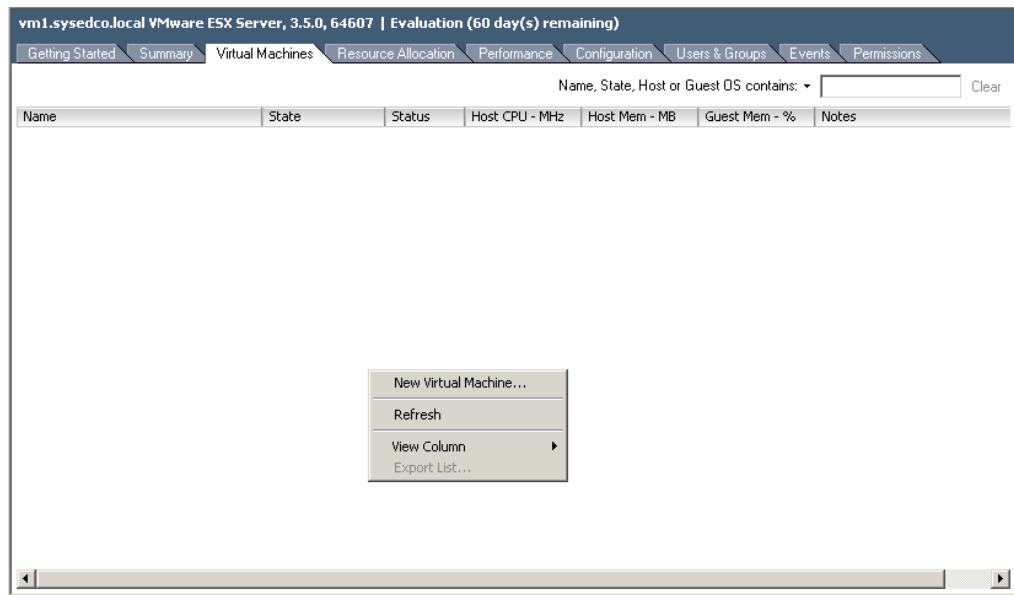
One of the primary benefits of virtualization is server consolidation. You can consolidate multiple physical servers that are underutilized onto a single ESX host. If you oversize the ESX VMs, by using the same settings that were in the physical servers (for example, memory and CPU settings), you are still underutilizing your resources. Take advantage of the virtual environment and scale the resources back enough so that quality demands are met, but resource waste does not occur.

Finally, just as with physical servers, remember to plan for the future. If you are implementing ESX servers without VirtualCenter, this is even more important. With VirtualCenter, you can move VMs from one host to another easily. Without it, the process is more complex. Make sure you are planning for future growth in VM utilization by leaving some resources available on the ESX host. Again, this recommendation only applies to non-centrally managed deployments that do not utilize VirtualCenter.

Creating a VM: Step-by-Step

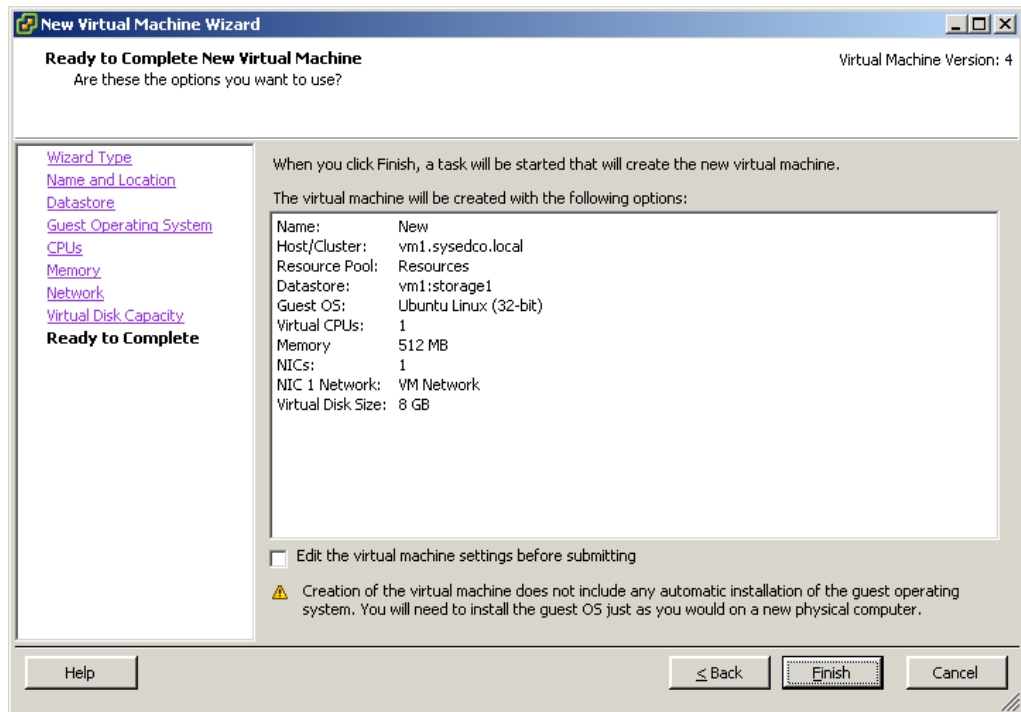
To create a VM for any operating system, follow these general steps:

1. Launch the VI client on a Windows machine.
2. Click the ESX host on which you wish to create the VM in the left pane of the VI client.
3. Click the Virtual Machines tab.
4. Right-click in the white space of the Virtual machines tab and select New Virtual Machine as shown here:



5. A New Virtual Machine Wizard dialog is displayed. Select Custom to configure I/O adapters or Typical to accept defaults for I/O adapters and click Next.
6. Provide a name for the VM and click Next.
7. Choose the data store location. Only locations previously configured on the Configuration tab of the VI client will be available. Click Next.
8. Choose the appropriate guest operating system. If your version does not show up, choose one that is close in version number to increase the likelihood of compatibility; however, only displayed operating systems and versions are supported by VMware. Click Next.
9. Choose the number of processors and click Next.
10. Configure the available RAM by either dragging the slider or manually entering the amount of RAM desired. Click Next.
11. Choose the number of NICs to make available, remembering that the virtual switch does have a limited number of ports, and click Next.

12. If you chose custom in step 5, you must now select the I/O controller type and click Next. You will then be prompted to configure new disks, use existing disks or map to a raw device, if you're using a SAN.
13. If you chose typical in step 5, you will choose the amount of space to provide on the virtual disk and click Next.
14. Review the settings as seen in the following image and, if they are correct, click Finish.



As you can see, the process of creating a VM is very straightforward. In the next section, we'll look at how you can modify an existing VM configuration set.

Modifying a VM

Modifying a VM is a much simpler process. When you right-click on any VM in the VI client, you can select Edit settings. The settings available for edit will depend on a number of factors:

- If the VM is powered on and does have the VMware Tools installed, many hardware settings can be changed. Most settings will not apply until the next VM reboot.
- If the VM is powered on and does not have the VMware Tools installed, most settings cannot be changed.
- If the VM is powered off, all hardware settings can be modified.

The most common settings you'll need to modify for a VM include:

- **Memory:** It is common to need more RAM later in the life of a server. For example, you may run a virtual Microsoft SQL Server that works fine for the first six months; however, at the six month point, several departments may ask you to add more databases to the server. Changes like these can often require additional memory.
- **Disks:** Like memory, disks may eventually fill up and need to be enlarged. You may also choose to add new disks instead of growing the existing ones. Adding more disks may require adding more physical storage space to the server.
- **CPUs:** VMs can support from 1 to 4 virtual CPUs (vCPUs). For many of the same reasons that you would upgrade memory and disks, you may need to upgrade to multiple CPUs. Additionally, it is not uncommon to deploy an application that only supports single processor machines only to upgrade to a newer version later on that supports multiple processors.
- **CD/DVD:** While this hardware is added infrequently, it is modified very often. It is not uncommon to attach a new ISO image to a VM monthly, weekly or even daily.

Deploying a Linux VM

VMware ESX 3.5 supports several variants of Linux. The supported versions include, but are not limited to:

- Ubuntu 7.04 Desktop and Server (32-bit and 64-bit)
- SUSE Linux Enterprise Desktop and Server 10 (32-bit and 64-bit)
- SUSE Linux Enterprise Desktop and Server 9 (32-bit and 64-bit)
- Red Hat Enterprise Linux Desktop and Server 2.1 through 5.2 (32-bit and 64-bit)

In order to gain support for newer operating systems, you may need to apply updates. For example, ESX Server 3.5 update 4 is required to support Ubuntu 8.04 or SUSE Enterprise Linux 11.

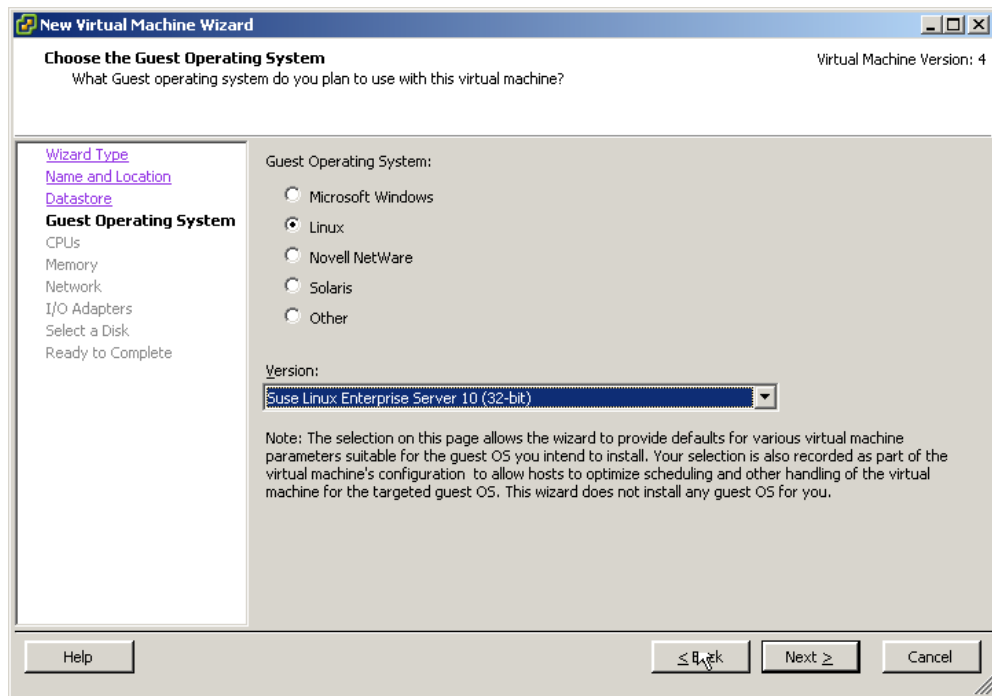
NOTE: If you decide to run a 64-bit version of Linux as a guest OS on VMware ESX servers running on EM64T systems, be sure to enable the Execute Disable bit in the BIOS. Otherwise, the Linux OS may fail to operate or become unstable.

It is important to remember that VMware ESX virtualizes a computer and a set of hardware. For this reason, most operating systems can be installed. For example, in addition to the Linux distributions mentioned previously, you may install CentOS, Debian, Mandriva, Mandrake and Novell Linux - just to name a few.

In the following exercise, you will install OpenSUSE 11 in a VM on the ESX 3.5 host. OpenSUSE 11 is the open source version of SUSE Linux Enterprise and it works very well on VMware servers. In order to perform this exercise, you will need a running ESX host and the OpenSUSE 11 distribution DVD.

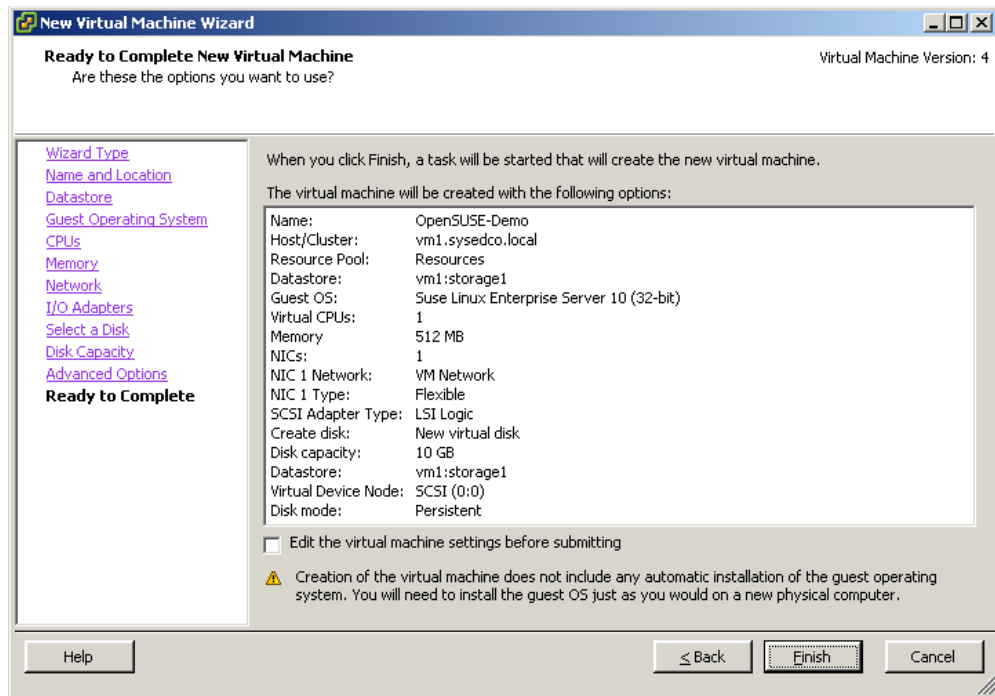
Creating an OpenSUSE Linux 11 VM

1. On a Windows machine running the VI client, launch the VI client and connect to the target ESX host.
2. Click the target ESX host in the left pane to ensure it is selected.
3. Click the Virtual Machines tab.
4. Right-click in the open white space and select *New Virtual Machine*.
5. Select Custom and click Next.
6. Name the VM OpenSUSE-Demo and click Next.
7. Select the storage location you wish to use (the location should have at least 2.5 to 3 GB free for an optimal installation) and click Next.
8. Select the Linux operating system and a version close to the OpenSUSE 11 distribution. Note that SUSE 10 has been chosen in the following image:



9. Click Next to continue.
10. Choose a single processor and click Next.
11. Accept 512 MB of memory and click Next.
12. Accept the default for the network interface and click Next.


13. Choose an LSI Logic controller for the drives and click Next.
14. Select Create a new virtual disk and click Next.
15. Set the disk size to 10 GB and click Next.
16. Accept the defaults for virtual device node and click Next.
17. Ensure your summary screen looks similar to the following:



18. Click Finish to create the VM.

Now that the VM is created, you can begin the installation process. Before you power on the VM, you should ensure that the OpenSUSE DVD is in the appropriate drive or attached as an ISO image. You may choose to upload the ISO image to the datastore on the ESX host so that you can mount the image. To upload the ISO image from your Windows client running the VI client to the ESX host, follow these instructions:

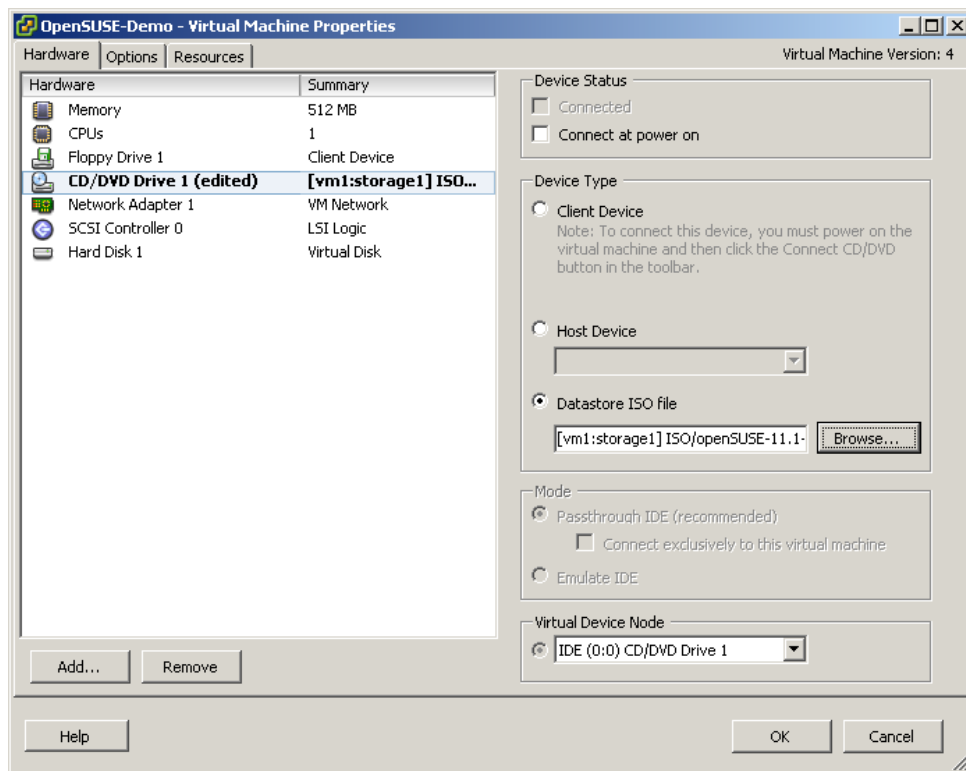
1. In the VI client, select the appropriate ESX host in the left pane.
2. Click the Configuration tab and select Storage in the Hardware section.
3. Double-click on the storage location to which you wish to copy the ISO image.
4. In the Datastore Browser, create a new folder called ISO.
5. Open the ISO folder by double-clicking it.

6. Click the Upload Files button () to upload the ISO image.
7. Browse for and select the ISO file to upload and click Open.

With the ISO image uploaded to the ESX host, you can begin the installation of OpenSUSE. The following instructions assume you are mounting an ISO image. If you are booting from a DVD instead, you can skip to step 8 to begin.

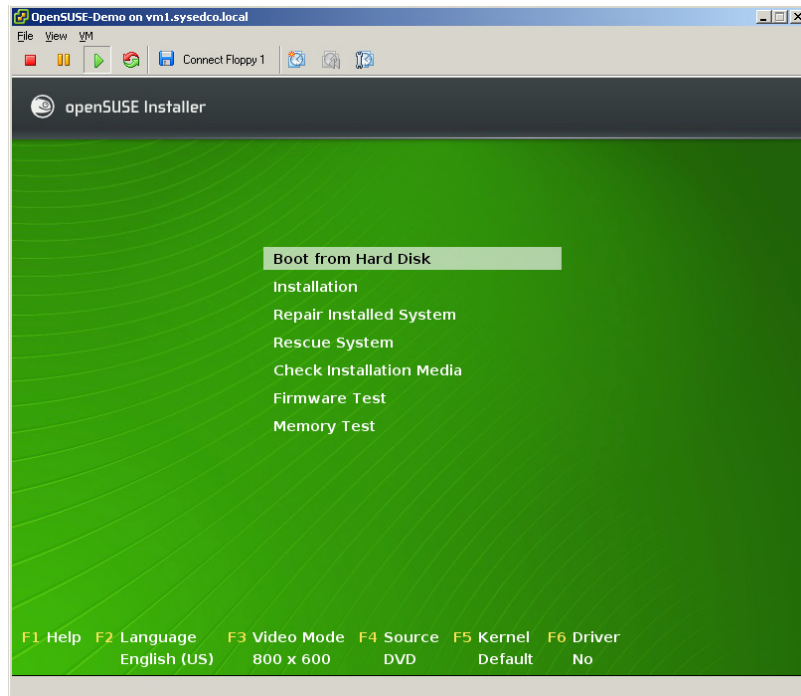
Installing OpenSUSE 11

1. Right-click the OpenSUSE-Demo VM and select Edit Settings.
2. Click the CD/DVD drive hardware item.
3. Choose Datastore ISO file under Device Type.
4. Click the Browse button to select the ISO image.
5. Open the datastore to which you've copied the ISO file and then open the ISO folder.
6. Select the OpenSUSE ISO file and click OK. You should see a screen similar to the following:



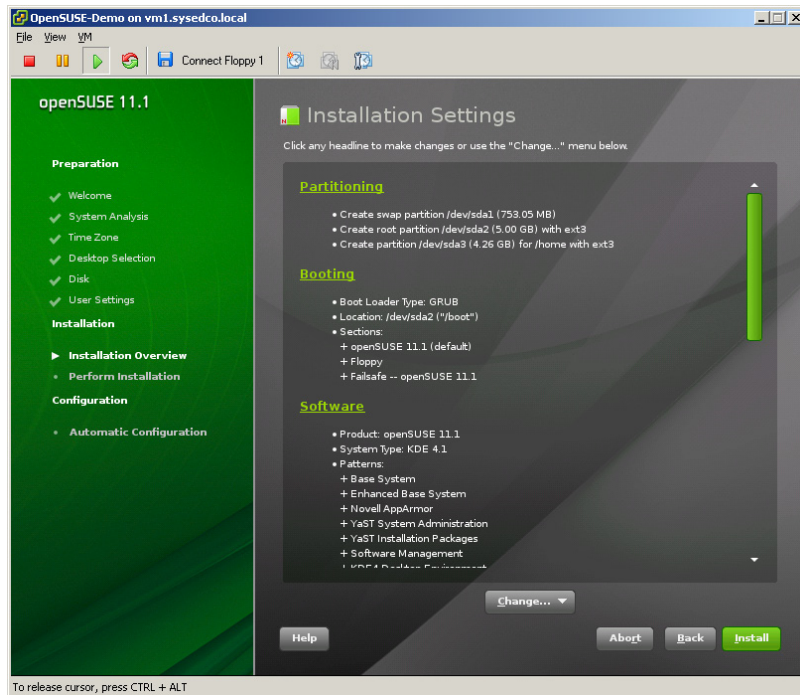
7. Choose the Connect at power on option and then click OK to save the settings and exit the dialog.
8. To access the VM for installation, right-click the OpenSUSE-Demo VM and select Open Console.

9. When the console is displayed, click the VM menu and select Power On (or press CTRL + B).
10. The VM should boot directly into the OpenSUSE installation process. You will see a screen similar to the following:



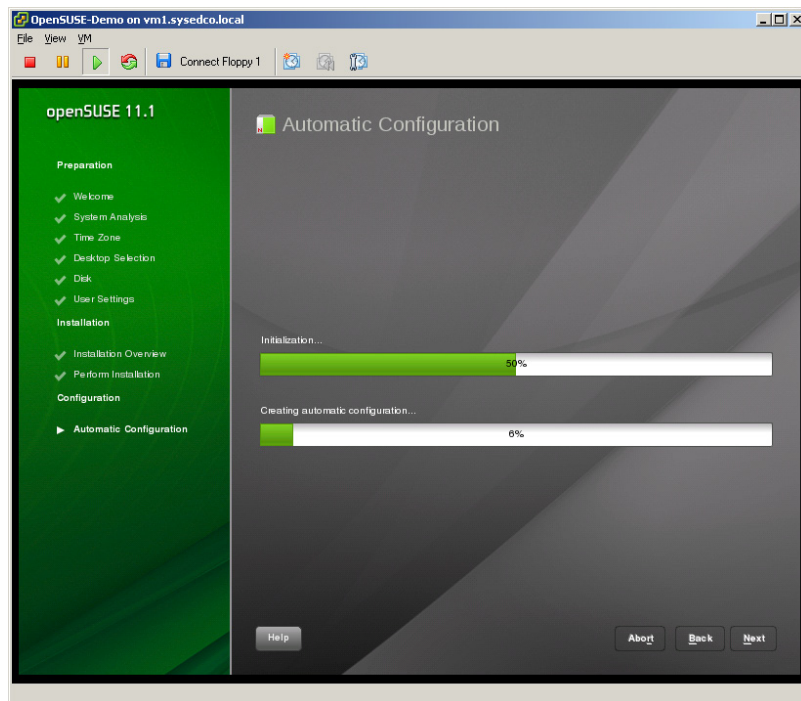
11. Choose Installation, by using the arrow keys, and press ENTER.
12. After the Linux kernel is loaded, the installation wizard will begin.
13. Accept the defaults for the language and keyboard layouts (unless you are installing a non-English version) and click Next.
14. The system analysis will begin, which evaluates your hardware and attempts to locate the appropriate drivers and configuration set.
15. When the system analysis completes, choose a New installation and click Next accepting all defaults.
16. Next, you'll need to configure the appropriate Time Zone. Choose your Time Zone and click Next.
17. On the next screen you must choose the desktop environment. For this installation, choose Gnome and click Next.
18. Accept the default for partitioning and click Next.
19. Enter a user name of Demo with a simple password of *password* and click Next, accepting all other defaults on the Create New User screen.

20. You will receive two prompts asking if you really want to use the simple password. Click Yes each time since this is only a demo installation. You should use secure passwords in production installations.
21. You will see an installation review screen similar to the following:



22. Once you're sure the installation plan is correct, click Install to begin the installation.
23. Another installation confirmation screen will appear, click Install again. The installation process will take from 10 to 45 minutes.

24. When the installation is complete, the system will reboot and launch an Automatic Configuration process as shown in the following image:



25. When the configuration is complete, the installation process will move to the Writing Configuration phase. Upon completion, the VM will restart automatically.
26. When you see the login screen, click the Demo user and enter the password of *password*, which was set during the installation. Click Login to complete the login process.

At this point, the basic installation of an OpenSUSE Linux distribution is complete; however, since this is a VMware environment, one more set of tasks remains. We must install the VMware tools to get the best performance out of the new Linux VM.

Installing the VMware Tools in OpenSUSE

1. With the OpenSUSE-Demo VM console in view, click VM and select Install VMware tools.
2. If you're not logged on, logon to the OpenSUSE VM as the Demo user created during installation.
3. If you were already logged on, the VMware tools CD window should be displayed. If not, double-click the VMware Tools icon on the desktop.
4. Right-click in the open white space within the VMware Tools window and select Open in Terminal.
5. Copy the VMware Tools compressed file to the /tmp directory with the following command:

```
cp VMwareTools-3.5.0-64607.tar.gz /tmp
```


6. Change to the /tmp directory by executing `cd /tmp`
7. Extract the contents of the VMware Tools compressed file with the following command:

```
tar zxvf VMwareTools-3.5.0-64607.tar.gz
```
8. Change to the /vmware-tools-distrib directory by executing `cd vmware-tools-distrib`
Change to root user mode by executing `su root` and entering the password or *password*.
9. Run the PERL installation script by executing `perl vmware-install.pl`
10. Simply press ENTER to accept the defaults for all prompts.
11. If you receive prompts asking for a GCC compiler, you may have to rerun the installation after installing the compiler.
12. When you are returned to the prompt, the installation is complete. Restart the VM to enable the VMware Tools.

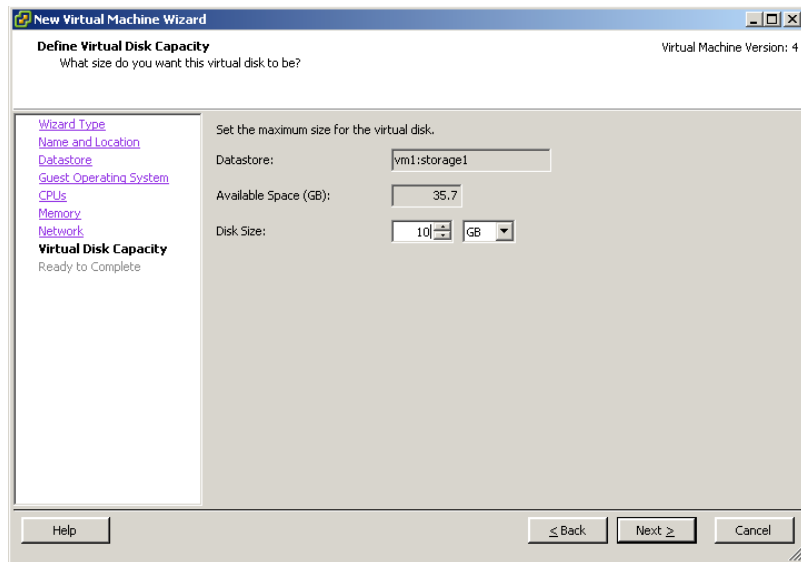
Deploying a Windows VM

Installing a Windows-based VM guest is simpler than the Linux guests for one reason - the VMware Tools are installed in a GUI interface. When you select to install the VMware tools, a graphical installation process is launched. You'll see this at the end of the installation process in the next exercise.

You can download an evaluation edition of Windows Server to work through this exercise. Visit Microsoft's website and search for *download evaluation edition*.

Creating a Windows Server 2003 VM

1. In the VI client, click the target host for the new VM in the left pane.
2. Click the Virtual Machines tab.
3. Right-click in the open white space and select New Virtual Machine.
4. Choose Typical for the VM type and click Next.
5. Name the VM Windows2003-Demo and click Next.
6. Choose the datastore for VM storage and click Next.
7. On the Choose the Guest Operating System screen, choose Microsoft Windows and your version of Windows Server 2003 (Standard, Enterprise, etc.) and click Next.
8. Choose 1 for the number of virtual processors and click Next.
9. Increase the memory setting to 512 MB and click Next.
10. Accept the defaults on the Choose Networks screen and click Next.
11. On the Define Virtual Disk Capacity screen, set the drive space to 10 GB as seen in the following image and click Next.

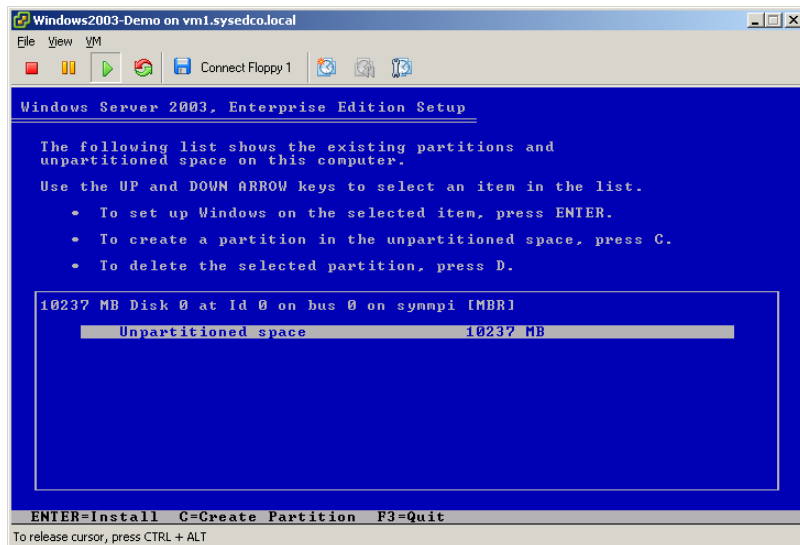


12. Click Finish to complete the creation of the VM.

Now that the VM is in place, you can attach the Windows Server 2003 ISO and begin the installation process. To upload the ISO to the datacenter, follow the instructions provided earlier for the upload of the OpenSUSE ISO image. To attach the ISO to the Windows Server 2003 VM, follow the instructions provided earlier for attaching the ISO to the OpenSUSE VM. With the ISO attached to the VM, you can power it on to begin the installation.

Installing Windows in the VM

1. Right-click on the Windows2003-Demo VM and select Open Console.
2. Click the Power On button in the console window to start the VM.
3. The Windows Server 2003 installation process should begin automatically.
4. When you reach the Welcome to Setup screen, press ENTER to perform an installation.
5. Press F8 to agree to the license agreement.
6. The next screen shows the available drive space as seen below. Simply press ENTER to use all of the available space for the Windows installation.




7. On the next screen, choose to *Format the partition using the NTFS file system (Quick)* in order to shorten the installation time. Press ENTER once you've highlighted the appropriate option.
8. The setup process will prepare the drive and then copy several files to the drive in order to launch the GUI portion of the installation. When the GUI portion begins, it will continue for several minutes without interaction. The entire process will take between 10 and 30 minutes.
9. The next interactive screen asks for your Regional Settings. Configure them as desired and click Next.
10. Enter a value for the organization and user and click Next.
11. Leave the Product Key blank to install an evaluation copy of Windows and click Next.
12. Accept the default licensing setting and click Next.
13. Enter the computer name of *WIN2K3DEMO* and a password of *Password1* to comply with minimum password rules and click Next.
14. Choose your Time Zone and click Next.
15. On the Network Settings screen, either accept the defaults for a dynamic IP configuration or choose Custom to configure the static settings you require. Click Next.
16. Accept the default workgroup name of *WORKGROUP* and click Next.
17. The installation will proceed to copy files and may take as long as 30 minutes at this point.
18. When the installation completes, the VM will automatically reboot.
19. Press **CTRL+ALT+I** (to send the CTRL-ALT-DEL keystroke to the guest) and logon as the Administrator with the password of *Password1* that was configured during the installation.

20. Because this is the first logon, it may take a few moments to show the desktop. Eventually, the Manage Your Server window will appear.


At this point the installation is complete; however, we still need to install the VMware tools.

Installing the VMware Tools in a Windows VM

1. Close any open windows in the guest VM.
2. If you need to unlock the mouse press the right CTRL + ALT keys together.
3. Click the VM menu in the guest console window and choose Install/Upgrade VMware Tools.
4. Click OK if an information dialog like the following appears.

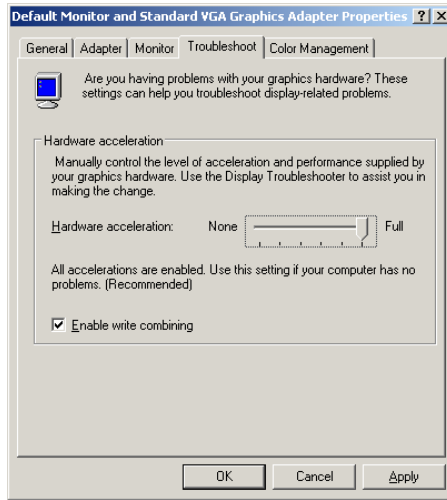


5. The VMware Tools installation process starts automatically. Choose Next to begin the installation.
6. Choose a Complete install and click Next.
7. Click Install to begin the installation.
8. If prompted with a message about hardware acceleration, as in the following image, choose Yes to configure the hardware acceleration.



9. Close the Notepad window that appears.
10. Click the Advanced button in the Display dialog and select the Troubleshoot tab.

11. Drag the hardware acceleration slider all the way to the right as in the following image:



12. Click OK to save your changes. Click OK again to exit the Display dialog.
13. In the VMware Tools installation window, click Finish to complete the installation.
14. Click Yes to reboot the server.

When the server is finished rebooting, the VMware tools will be installed and you should notice that the performance is much improved.

VMware and Hardware Emulation

At this point, you may be wondering just what hardware is emulated by the VMware ESX hosts for the guests. This is not only an interesting question, but the answer can certainly help if you ever have to install a very old OS - like OS/2 Warp 3 or Windows NT 3.51. Though the likelihood is low, the following information may prove helpful.

| Hardware Device | Emulated Device |
|-----------------|---|
| Network adapter | Generic AMD PCNet |
| Video adapter | Standard VGA drivers from most OSes work |
| Sound adapter | Not supported (ESX is intended for server virtualization) |
| SCSI adapter | BusLogic Model BT-958 PCI Wide Ultra SCSI Host Adapter |

Lesson 4: Summary

In this lesson, you learned to create virtual machines using the free VI client. You actually installed OpenSUSE Linux and Windows Server 2003 in virtual machines. With the information provided, you should be able to install most any supported guest OS in the ESX environment.

Lesson 5: Understanding and Installing VirtualCenter 2.5

In large enterprise environments - and even smaller organizations with many ESX hosts, a centralized management solution is needed. VirtualCenter 2.5 provides that centralized management for your VMware environment. In addition, features such as high availability, VMotion and most of the other advanced features VMware offers require the implementation of VirtualCenter.

In this lesson, you will learn about the VirtualCenter product and the features it provides. You will understand the requirements of the product and learn to install it successfully on Windows machines. Finally, you will learn to add ESX hosts and manage VirtualCenter settings and inventories.

VirtualCenter Overview

VirtualCenter is the center or hub of your virtualization management infrastructure (usually called the Virtual Infrastructure or VI). Figure 5.1 shows the components needed for a VI with VirtualCenter at the core. As you can see, multiple components are involved and VMware does not provide all of the components. For example, the database server may be provided by Microsoft or Oracle, but VMware does not use a database system that is developed in house.




Figure 5.1: VirtualCenter Components and Communications

VirtualCenter Components

The individual components are reviewed in the following pages:

VirtualCenter Management Server: The VirtualCenter server is an application that runs on the Windows operating system. VirtualCenter is the core of the VI and it is responsible for monitoring ESX hosts. It also provides the advanced features offered by VMware including high availability, VMotion, Storage VMotion, distributed resource scheduler, update manager, and VMware consolidated backup.

Database Server: The database server is required to use VirtualCenter. VirtualCenter connects to the database during startup and the database can be run on SQL Server or Oracle. SQL Server Express is supported for smaller implementations. VirtualCenter reads configuration information from the database as well as information about the managed ESX hosts.

ESX Servers (Hosts): After startup, VirtualCenter connects to the ESX hosts and ensures that the proper licensing is in place. Additionally, heartbeats are sent from the ESX hosts to the VirtualCenter server using UDP (user datagram protocol) port 902. It is important that you allow UDP port 902 through the network to the VirtualCenter server; if you do not allow this communication, the high availability features will not work properly.

Active Directory: VirtualCenter performs authentication for support staff using Active Directory (AD). If you want to deploy an enterprise implementation of VMware's VI3 (virtual infrastructure 3, which includes VirtualCenter 2.5 and ESX 3.5), you must have an existing AD infrastructure. This issue is often addressed in passing by VMware technical support and webcasts, but it is a very important dependency. You can authenticate using a local Windows account database in a small implementation.

SMTP (Simple Mail Transfer Protocol): An optional feature of VirtualCenter is alert notification. An SMTP server is required to send the email alerts.

VI Clients or WebAccess: The VI client, which we used in Lessons 2-4 to work directly with ESX hosts, is also used to manage VirtualCenter. Using the VI client you can manage ESX hosts directly in a VI3 environment as well. The WebAccess interface provides a web browser interface to the VirtualCenter, but the complete feature set is only exposed through the VI client.

As you can see, VMware's VI3 is not a self-contained or self-sustained solution. Instead, they depend on Oracle or Microsoft for the database and Microsoft for the authentication (Active Directory or Windows local databases). This is a key difference between VMware VI3 and Microsoft's Hyper-V. In an enterprise Hyper-V implementation, Microsoft provides everything from the virtualization (Hyper-V) through to the database (SQL Server) and enterprise management (System Center). While this makes the product appealing to Microsoft shops, the existing large deployment base of VMware products certainly provides an uphill battle for the Hyper-V product line.

NOTE: When you connect to VirtualCenter and then open a console connection to an ESX VM, that connection is not proxied or routed through the VirtualCenter server. Instead, a direct connection is made with the ESX host.

VirtualCenter Features

Several features, which are marketed as the most important features of ESX and VI3, are only available if you also run VirtualCenter. The following list outlines these features:

VMotion: VMotion allows the administrator to move a VM from one ESX host to another without interrupting user sessions. Of course, the VMs must be stored on a shared storage location that both ESX hosts can access, but the ability to move the actual processing from one ESX host to another means that maintenance of the physical host machines, among other tasks, does not have to result in down time for the users.

Storage VMotion: Storage VMotion takes VMotion to the next level and was first introduced in VirtualCenter 2.5. Storage VMotion allows the VMs to be moved from one storage location to another without impacting the users' access to the VM. This is accomplished by copying the VM files to the alternate storage location and then redirecting the ESX host to the new location before finally deleting the VM files from the old location.

High Availability: High Availability or HA is a feature that allows the administrator to quickly bring a failed VM online on a different ESX host. Technically, VirtualCenter is only required to initially configure HA, but it would be a rare situation where it is used to configure it and then is uninstalled.

Distributed Resource Scheduler (DRS): DRS is a service that monitors the ESX hosts to measure utilization. If a given ESX host is over-utilized, DRS can move an allowed VM to another ESX host automatically. VMs may be moved (using VMotion) among any of the ESX hosts in the DRS cluster (not to be confused with “real” clusters, this cluster is simply a collection of ESX hosts participating in the DRS function).

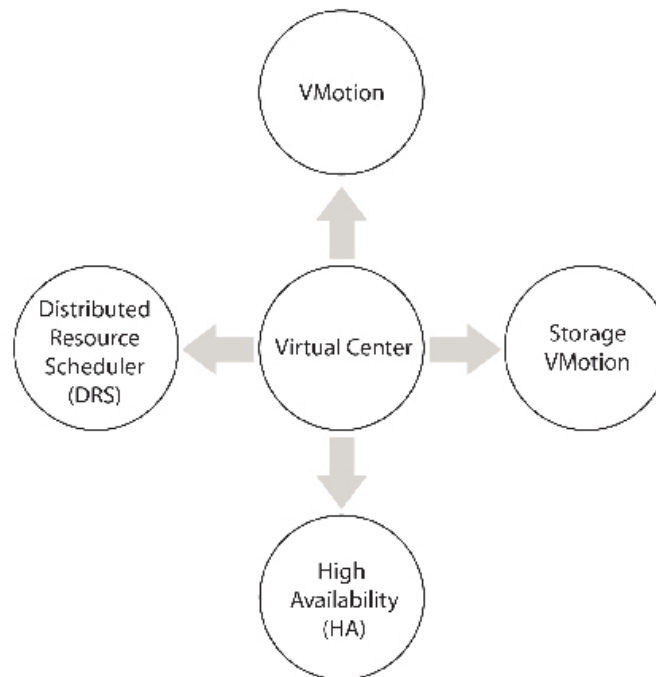


Figure 5.2: Features Requiring VirtualCenter

The VirtualCenter Database

The VirtualCenter database contains the configuration settings for your VI. It is very important. Several databases are supported, including:

- SQL Server 2005 (Express or Enterprise)
- SQL Server 2000 with manual tweaks
- Oracle

A little known fact is that the VirtualCenter depends so much on the database that, if the database must be restarted, VirtualCenter must be restarted as well. For this reason, it's important that you implement database monitoring for the database used by VirtualCenter. If the database goes down, you need to be notified immediately. This notification can be accomplished using jobs and alerts in SQL Server or with third party database monitoring solutions.

An additional concern related to the database is the collection of statistics. By default, VirtualCenter collects statistics from the ESX hosts every 5 minutes. By default, the statistics are collected at level 1, which is the lowest level and the smallest data set. Even at this level, however, with dozens of ESX hosts, the database can grow very rapidly. In these larger environments, you may need to implement the database on a larger server with more memory and hard drive space.

Most environments are implementing VirtualCenter using Microsoft SQL Server. This being the case, the following recommendations should help to improve the database performance on these SQL Servers:

Create a maintenance plan. A database maintenance plan includes integrity checking, rebuilding or reorganizing of indexes, and backing up the database. These tasks are essential to a well-performing and recoverable VMware database.

Dedicate the SQL Server instance to VirtualCenter. By dedicating the SQL Server instance to VirtualCenter, you can use Windows System Resource Manager (WSRM) to prioritize the VirtualCenter database. WSRM only works on the Enterprise edition and higher for Windows Server 2003; however, it is available on all editions of Windows Server 2008 and SQL Server 2005 can be installed on this newer Windows version. You must run SQL Server 2005 SP2 when installing on Windows Server 2008.

Monitor the database server to ensure continued performance. Over time, the VirtualCenter database will grow and the performance of the server may suffer. Using the Performance tool in Windows, you can monitor the performance of the server and compare it to a baseline. If the performance is lowered by a significant amount, you can consider adding memory, removing other instances, or upgrading the physical server.

VirtualCenter Requirements

Like most applications, VirtualCenter does have minimum requirements. The following table outlines the hardware and software requirements for VirtualCenter 2.5.

| Factor | Requirement |
|-----------------------------|--|
| Operating System | Windows 2000 SP4 with update rollup 1 Windows 2003 Server SP1 Windows 2003 Server R2 Windows XP Pro SP2 |
| Memory | 2 GB or higher |
| CPU | 2 GHz or higher |
| Hard Drive | 560 MB minimum, 2 GB recommended |
| Database Requirement | SQL Server 2000 Standard or Enterprise with SP4 SQL Server 2005 Enterprise with SP1 SQL Server 2005 Express with SP1 Oracle 9i release 2 Standard or Enterprise Oracle 10g release 1 or 2 Standard or Enterprise |

In addition to the requirements of VirtualCenter, the client from which VirtualCenter will be administered must meet the minimum requirements of the Virtual Infrastructure (VI) Client application. The VI Client requires a 266 MHz or faster processor with 256 MB or more of memory. Additionally, the client must have 150 MB of free space for the installation. The VI Client can run on Windows 2000 with SP4 through Windows Vista Enterprise.

Database Requirements

In addition to requiring a supported database system (SQL Server or Oracle), the database itself must be properly prepared. If you plan to perform a small installation, you can simply use SQL Server Express and let the VirtualCenter 2.5 installation process configure the database for you. Otherwise, you will need to create and prep the database.

The VirtualCenter database can have any name you desire. A best practice would be to name the database something that indicates its purpose. A name like VirtualCenter will work well. The following steps assume you are using SQL Server 2005 to prep the database:

1. Launch the SQL Server Management Studio.
2. Right-click the Databases container in Object Explorer and select New Database.
3. Enter a database name (such as *VirtualCenter*) and accept all other defaults.
4. Click OK to create the database.
5. Now, expand the Security container in Object Explorer and then expand Logins.
6. Right-click Logins and select New Login.
7. Enter a name like *VCuser*.
8. Select SQL Server authentication and provide a password.
9. Select the Server Roles page and add the *sysadmin* role.
10. Click OK to create the user.

Remember, these steps are only required when using SQL Server Enterprise edition. When using SQL Server Express Edition, you can simply let the installation process do all the work. If you use SQL Server Enterprise Edition and have followed the preceding steps, you will also need to create an ODBC DSN that maps to the SQL Server. The SQL Server Native Client is supported in VirtualCenter 2.5 and it is the fastest communications method. During the installation of VirtualCenter, if you do not use Express Edition, you will need the following three items:

- The DSN name.
- The sysadmin user name (VCuser)
- The user password

Using SQL Server Express Edition

You should be aware of a few limitations when using SQL Server 2005 Express Edition for your database. The following list provides a breakdown of the limitations.

- **5 hosts:** VMware recommends running no more than 5 ESX hosts with a SQL Server 2005 Express database. Remember, statistics are gathered from each host and the database can grow large very quickly.
- **50 VMs:** Each VM added to a host adds to the data in the database. A collection of 5 hosts with 10 VMs is estimated to require 1 GB of database space after one year. If the same 5 hosts include 50 VMs, the database will be 2.1 GB.

Database Size Planning

Since the database can grow rather large, it's a good idea to plan for it. The VMware website provides an Excel spreadsheet that can be used for these estimations. The calculator can be downloaded here: http://www.vmware.com/support/vi3/doc/vc_db_calculator.xls. Figure 5.3 shows the calculator for 10 ESX hosts with 53 VMs running on them.



VirtualCenter Database Sizing Calculator for Microsoft SQL Server
Valid for VirtualCenter 2.5

| | | | | | |
|--|--|---|------|------|------------------------|
| Number of hosts | 10 | | | | |
| Number of virtual machines | 53 | | | | |
| Number of clusters | 0 | | | | |
| Number of resource pools | 25 | | | | |
| Average number of CPUs per host | 4 | | | | |
| Average number of virtual CPUs (VCPUs) per VM | 1.5 | | | | |
| Average number of network interfaces per host | 2 | | | | |
| Average number of virtual network devices per VM | 1 | | | | |
| Average number of disk devices per host | 10 | | | | |
| Average number of disk devices per VM | 2 | | | | |
| | Number of samples collected every 5 mins | Potential DB size in gigabytes at the end of 1 Year | +15% | -15% | Space need for temp DB |
| Statistics Collection Level 1 | 481 | 0.1 | 0.2 | 0.1 | 0.1 |
| Statistics Collection Level 2 | 1,610 | 0.5 | 0.5 | 0.4 | 0.5 |
| Statistics Collection Level 3 | 4,733 | 1.3 | 1.5 | 1.1 | 1.3 |
| Statistics Collection Level 4 | 7,032 | 2.0 | 2.3 | 1.7 | 2.0 |


Figure 5.3: The VMware Database Sizing Calculator

Installing VirtualCenter

The installation of VirtualCenter is a standard installation process; however, a few important factors must be considered along the way. The following instructions walk you through the installation procedure. The instructions assume you are installing on a Windows 2003 Server machine with SP2 or on a Windows 2003 Server R2 machine. Additionally, the instructions are for an installation using SQL Server 2005 Express Edition.

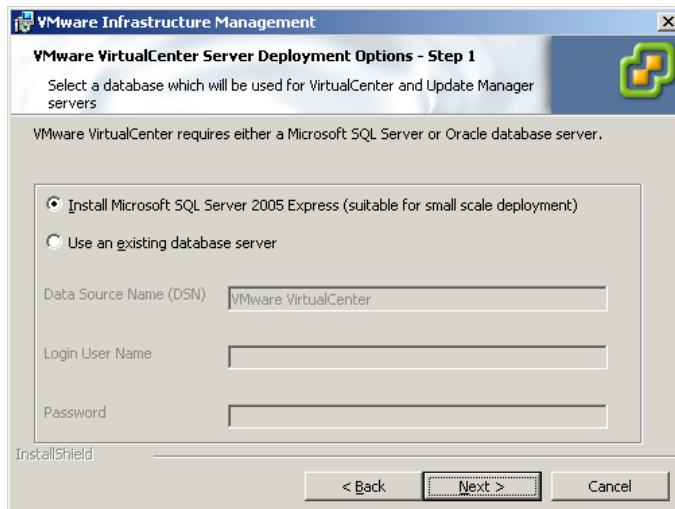
1. Insert the Virtual Infrastructure CD.
2. If the installation program does not launch automatically, open My Computer and double-click on the CD-ROM drive.
3. On the Welcome screen, click Next to begin the installation.
4. Read the helpful reminders on the Introduction screen and click Next.

5. Read and accept the license agreement and click Next.
6. Provide your user name and organization name and click Next.
7. You are presented with three choices in the Installation Type dialog as seen below:



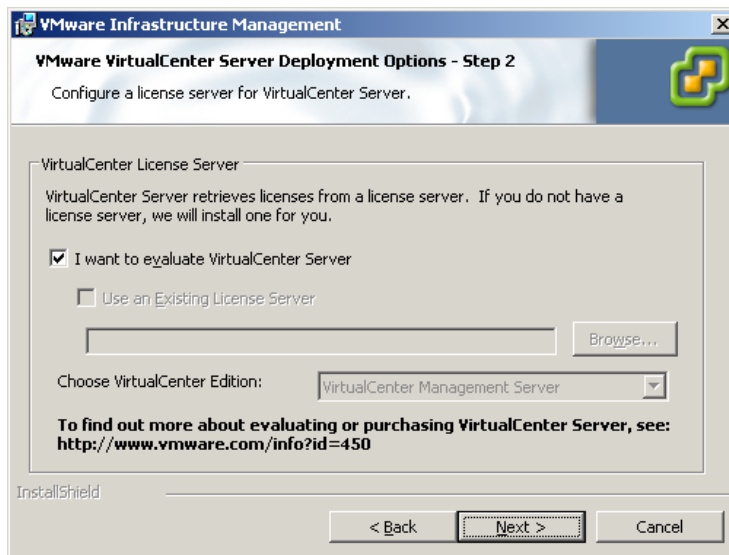
You may choose only to install the VI Client. This option is useful on administrator's computers. You may also choose Custom to determine exactly which components you need. In this case, we will choose VMware VirtualCenter Server and click Next.

8. If you already have the VI Client installed, you'll see a notification box that it is already installed and that only the extra needed components will be added. If this notification is displayed, click OK to continue.
9. You are presented with a database selection screen similar to the following:



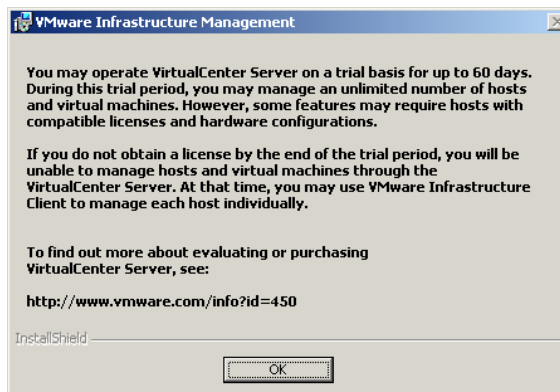
This is where you can select the full SQL Server Enterprise Edition database if needed. We will choose *Install Microsoft SQL Server 2005 Express* and click Next.

10. Since this is a non-production server, on the next screen, choose *I want to evaluate VirtualCenter Server* as seen below and click Next:

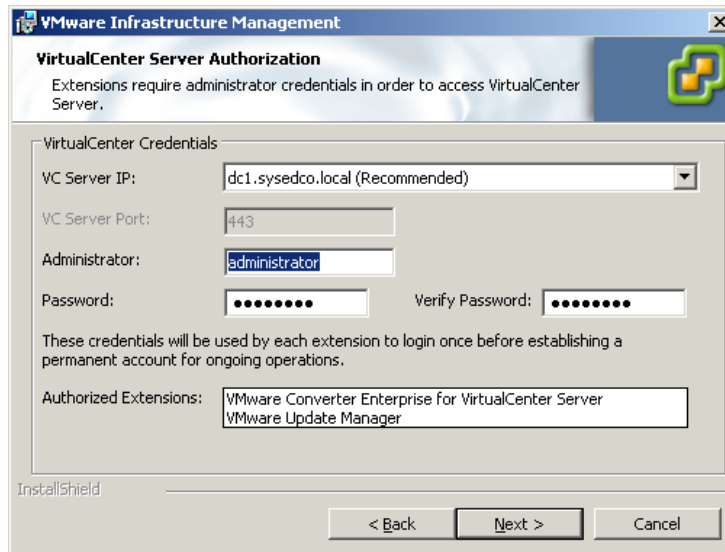


NOTE: You can also configure VirtualCenter to use an existing licensing server here.

11. Since we've selected evaluation mode, a warning will be shown that indicates the time limit for evaluation mode as seen below, click OK:



12. On the next screen, enter the administrative account and password (see image below). These credentials are used by supporting services the first time they connect to the VirtualCenter, click Next.



13. Click Install to begin the actual installation process. The process can take from 10 minutes to 20 minutes depending on the speed of the machine and whether SQL Server 2005 Express is already installed on the machine or not.
14. When the installation complete, click Finish.

Now that the installation is complete, you can check the services list to see if the VI services are installed. Right-click on My Computer and select Manage. Expand the Services and Applications node and then select Services. In the right pan, you should see several VMware services. Figure 5.4 shows the VMware services in the Services management tool.

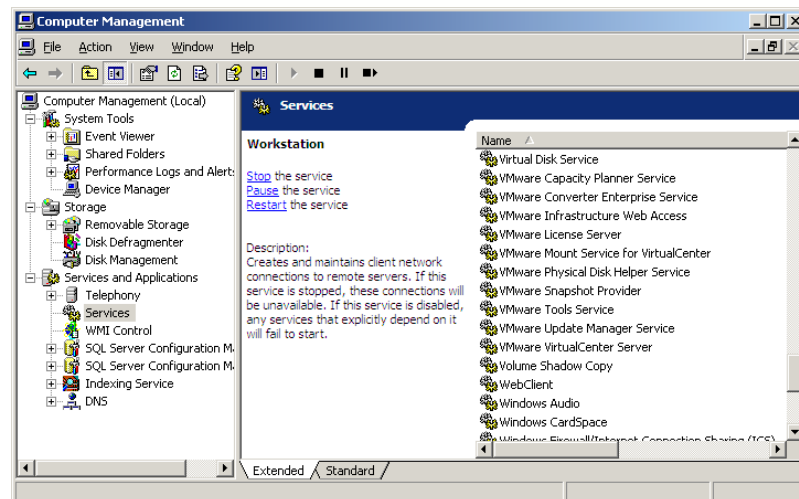


Figure 5.4: VMware Services

Common Installation Problems

Aside from the normal damaged media and corrupt operating systems issues, which can occur with any application, VirtualCenter faces some common problems. These problems are usually related to database connectivity issues, port conflicts or connection problems from the VI client. All three are addressed in the following sections.

Database Connectivity Problems

If the VirtualCenter application cannot connect to the database, verify the following items:

- Make sure the database service is running.
- Ensure that the ODBC DSN has been properly configured.
- Ensure that the user name provided is a sysadmin user.

If all of these factors are true, make sure that TCP port 1433 is allowed (for the default instance of SQL Server) for communications.

Port Conflicts

Port conflicts may also be an issue. Ports should be considered from two perspectives: the in use ports on the local machine and the network ports needed for communications. From the first perspective, if you are running VirtualCenter on a machine that is already running IIS, you may run into some problems. VirtualCenter uses an embedded version of the Apache TomCat web server to provide WebAccess. Since IIS is already using ports 80 and 443, you will have a port conflict.

The `netstat -nab` command will show all of the ports on which services are listening. You can use this command to find out what ports are available in order to use an appropriate port. The best practice is to install VirtualCenter on a server that is not running IIS.

The following table lists the important ports that are used in a VI3 environment.

| Service | Ports |
|-----------------------------|------------------------------------|
| License Server | 27000 and 27010 |
| Oracle | 1521 |
| SQL Server | 1433 (assuming a default instance) |
| WebAccess | 80 (HTTP) and 443 (HTTPS) |
| ESX hosts | 902 |
| Connections from VI Clients | 443 |
| Remote Consoles | 903 |
| NFS | 2049 |
| iSCSI | 3260 |
| VMotion | 8000 |

VI Client Connection Problems

In most cases, VI Client connection problems are related to port filtering within the network or proxy settings. The VI client uses HTTP for the initial connection, so the proxy settings must be configured appropriately on a client machine that requires a proxy for HTTP communications.

Updating VirtualCenter

When you must update VirtualCenter, keep the following tips in mind:

- All VirtualCenter updates are complete installs. The installation completely reinstalls VirtualCenter including all patches up to the point of the update.
- Make sure your database is backed up before performing an update.
- Backup your certificates (used for SSL/SSH connections) before performing an update.

As long as you keep these factors in mind, your updates should go well. If you have a problem, you can restore the database and certificates and try again.

Once VirtualCenter is updated - at the next startup, it will push new agents out to all of the ESX hosts. The update of VirtualCenter is not complete until both the VirtualCenter application and the ESX hosts are updated.

Finally, the VI client should be updated on all administrative workstations. This must be performed manually.

Running VirtualCenter in a VM

VMware supports the option of running VirtualCenter in a VM. This configuration results in the use of a VM on an ESX host for the management of the VI. The VI is truly managed from within the VI. In order to implement a VM-based VirtualCenter deployment, you will need to consider several factors:

- At least one ESX host must use host-based licensing and that ESX host must run the VirtualCenter server.
- If your infrastructure uses DRS, you must disable it for the VirtualCenter VM. Otherwise, should the VirtualCenter VM be unreachable, you'll have to connect to each ESX host in order to find the VirtualCenter VM (since DRS may have moved it).
- VMware recommends running the database on a separate VM or physical server from the VirtualCenter VM.
- You cannot clone the VirtualCenter VM and you cannot perform a cold migration.

Configuring the Inventory

Now that VirtualCenter is installed, it's time to configure the inventory. The inventory includes *datacenters* and *folders*. These two elements provide a hierarchy for management of ESX hosts and VMs.

Creating a Logical Structure

Folders are used for functional and logical organization. For example, if you have three physical network sites in Ohio and two such sites in Texas, you may create a folder named Ohio and another folder named Texas. Additionally, folders can contain datacenters and datacenters can contain folders. For example, if the three locations in Ohio are Columbus, Springfield and Athens, you may choose to create a datacenter for each of these locations. Furthermore, in each datacenter, you may choose to create folders representing server functions, such as database servers, email servers and so on. Figure 5.5 shows a typical structure. Notice the use of both folders and datacenters.

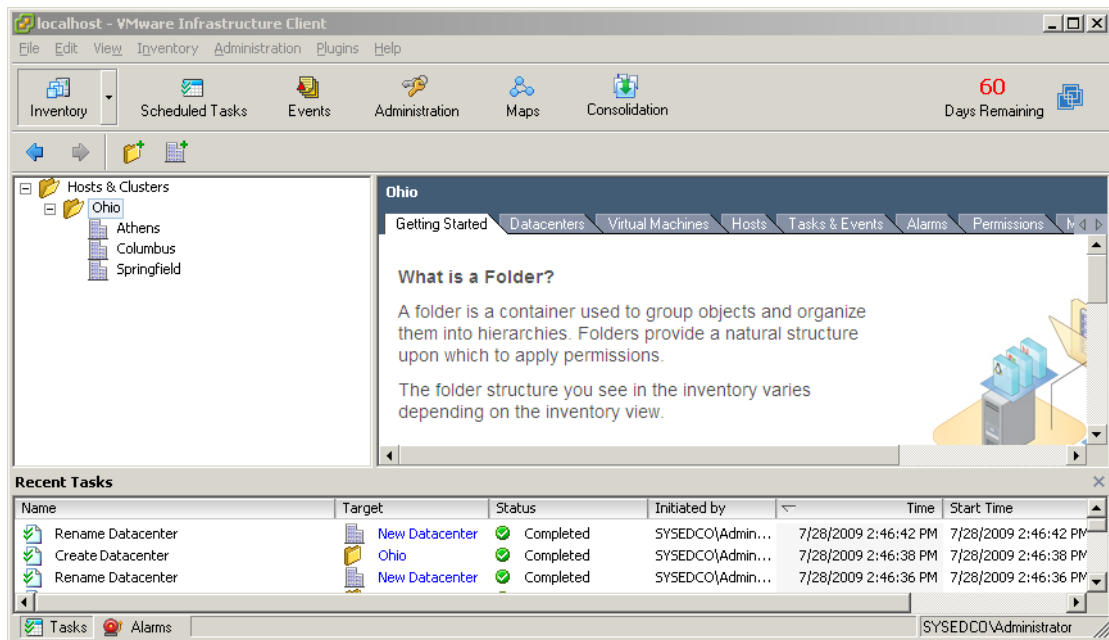


Figure 5.5: Using Logical Structures with Folders and Datacenters

A datacenter, however, must be considered with more care than a folder. VMs cannot be VMotioned from one datacenter to another. The datacenter is considered to be a local area in which all physical hosts and storage reside.

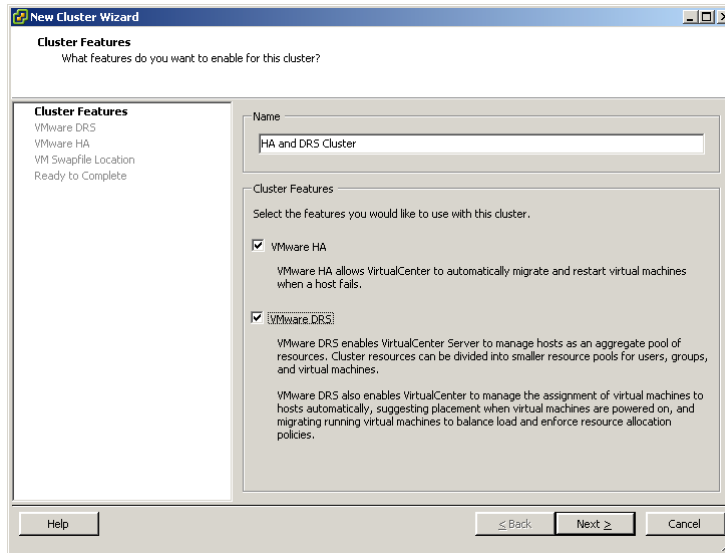
To create the structure represented in Figure 5.5, follow these steps:

1. Right-click on the Hosts & Clusters node in the VI Client that is connected to the VirtualCenter server and select New Folder.
2. Name the folder Ohio.
3. Right-click on the Ohio folder and select New Datacenter.
4. Name the Datacenter Athens.
5. Repeat steps 3 and 4 two more times while changing the name to Columbus and Springfield respectively.

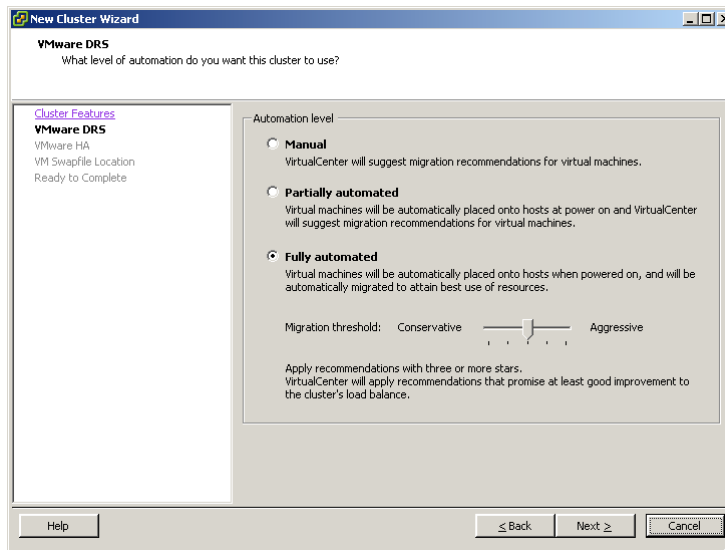
Clusters

Clusters are used to group multiple ESX hosts together within a shared resource pool. Once clustered, the ESX hosts can participate HA and DRS functions. When you create the cluster, you determine the functions available (HA only, DRS only or both). To create a cluster, follow these steps:

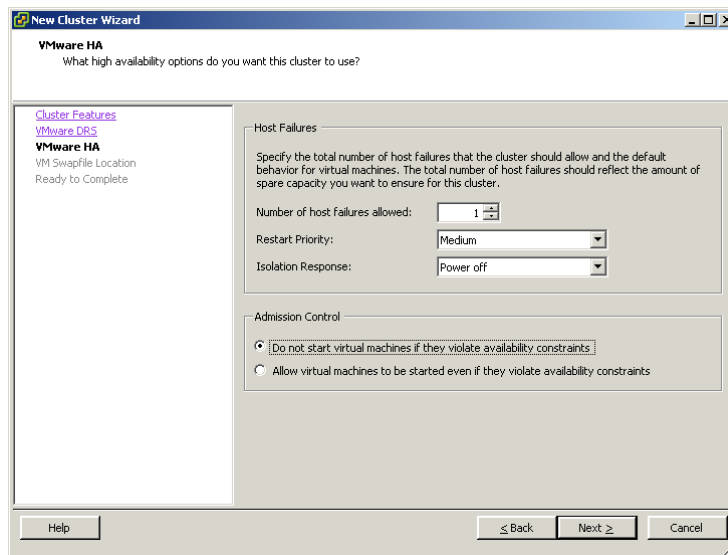
1. Right-click on the datacenter in which you wish to place the cluster and select New Cluster.
2. Name the cluster and choose the desired functions as seen in the following image, click next.



3. If you chose DRS, you will be presented with a screen used to configure VMware DRS as seen below. Choose the migration method you desire for DRS functions and click Next.



4. Choose the HA settings (see image below) you desire (described in the HA Settings section later) and click Next.



5. Choose to store the swapfile in the same directory as the VM and click Next.
6. Click Finish to create the cluster.

Once the cluster is created, you can add ESX hosts to it just like a datacenter.

DRS Automation Levels

The DRS automation levels are used to determine how and if DRS can move VMs from one ESX host to another. Three levels are available:

- **Manual:** When set to manual, DRS suggests where the various VMs should be moved and you must move them manually using VMotion. This is a good option for static environments. You do not want the VMs moving around a lot, but you also don't want to make the decision about where each VM should be located yourself.
- **Partially Automated:** In this mode, VMs are initially placed based on the DRS function; however, they are never moved after power on. DRS will make recommendations after power on, but it will not move the VMs automatically.
- **Fully Automated:** When fully automated, DRS initially places the VMs on the ESX host it determines best. During continued operations, DRS will automatically relocate VMs for you in order to maintain a consistently well-performing environment.

HA Settings

When creating a cluster, you are presented with options for HA operations. The following options can be configured as seen in the image presented during the preceding step-by-step procedure for creating a cluster.

- **Number of host failures allowed:** This setting can be from 1 to 4. You are configuring the number of host failures you can tolerate without taking down the entire cluster.
- **Restart Priority:** This setting determines the priority for VM restarts within the cluster. You can change this priority on a VM-basis, which allows you to determine which VMs should be allowed to run if enough hosts fail so that all VMs can no longer be supported. VMs are started in the order of high, then medium, then low. Disabled VMs are never started unless you are within the tolerance set for the number of host failures allowed.
- **Isolation Response:** The isolation response setting determines what happens when an ESX host is no longer receiving heartbeats from other hosts in the cluster. The options are Power off or Leave powered on.
- **Admission Control:** This section provides two options. The default option, *do not power on virtual machines if the violate availability constraints*, will force a VM to remain powered off if resources are not available to meet the VMs configuration settings. The alternate option, *allow virtual machines to be powered on even if they violate availability constraints*, allows the VM to be powered on anyway.

Adding ESX Hosts

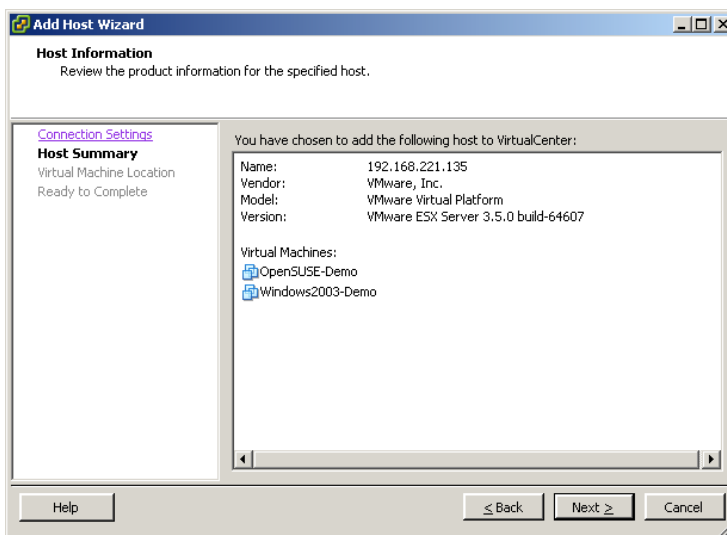
Once you have your inventory in place you can add ESX hosts to the VirtualCenter inventory. Adding an ESX host is a simple task. You can add hosts to datacenters or to clusters. You cannot add ESX hosts to top level folders, but you may add hosts to sub-folders within datacenters or clusters.

To add a host, follow these general steps:

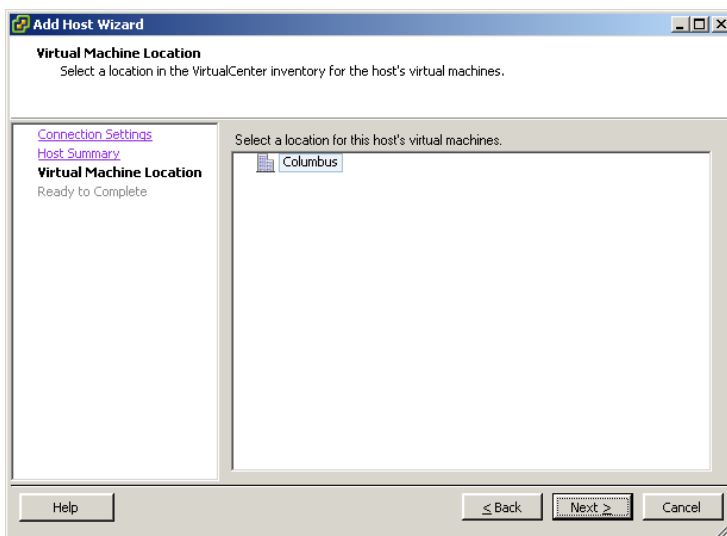
1. Right-click the supporting container and select Add Host.
2. Provide the host name or IP address and a valid user name and password, click Next.

The screenshot shows the 'Add Host Wizard' dialog box with the 'Specify Connection Settings' step selected. The dialog is titled 'Add Host Wizard' and has a subtitle 'Specify Connection Settings'. Below the subtitle, it says 'Type in the information used to connect to this host.' On the left side, there is a 'Connection Settings' panel with a tree view showing 'Host Summary', 'Virtual Machine Location', and 'Ready to Complete'. The main area of the dialog is divided into two sections: 'Connection' and 'Authorization'. The 'Connection' section has a text box for 'Host name' containing '192.168.221.135'. The 'Authorization' section has text boxes for 'Username' containing 'tom' and 'Password' which is masked with asterisks. At the bottom of the dialog, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

- Verify the proper host has been selected by reviewing the current VMs on the host, click Next.



- Choose the location or accept the default, click Next.



- Click Finish to add the host.

Configuring Settings

At this point, you should have a fully functioning VMware virtual infrastructure. At the very least, the software is in place. You will not actually be able to use features like VMotion and DRS unless you have a SAN solution in place. Several universal settings should be configured to ensure proper operations in your environment. These settings include:

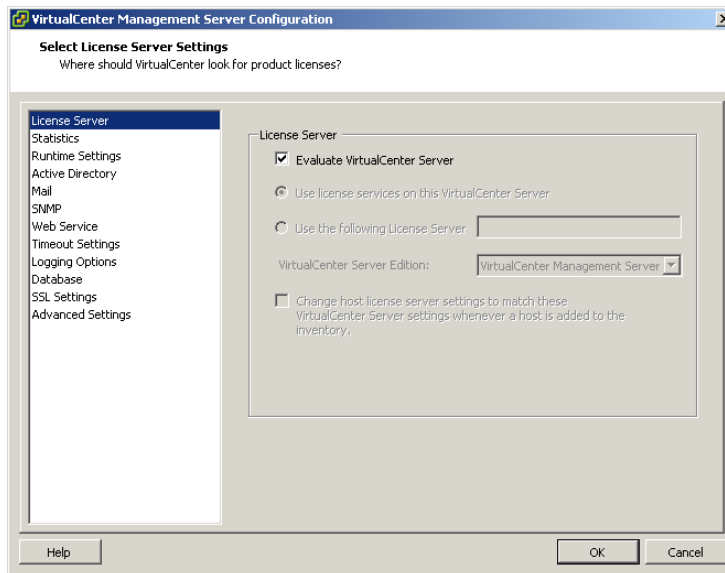
- Statistics
- Alarms
- Scheduled tasks

Configuring Statistics

If statistics are not configured properly, your database can grow too large too fast. VMware recommends setting the statistics to level 1 or 2. A best practice may be to set statistics at level 1 during normal operations and then use level 2 only during performance analysis windows. Levels 3 and 4 are usually used only during troubleshooting times.

To configure statistics:

1. Connect to the VirtualCenter server using the VI client.
2. Click the Administration menu and select VirtualCenter Management Server Configuration.
3. You will see the dialog shown here:



4. Select the Statistics page. Here you can configure intervals of 5 minutes, 30 minutes, 2 hours and 1 day. For each, you can choose to keep the statistics for a specific window of time and you can choose the verbosity (level) of the statistics.
5. Choose the statistics settings you desire and click OK.

Consider Figure 5.6. Notice that the estimated space required is nearly 400 GB. Why is this? The first statistics setting has been configured to gather statistics every minute, keep them for 5 months, and use statistics at level 4. This screenshot should make it very clear: be careful when configuring statistics.

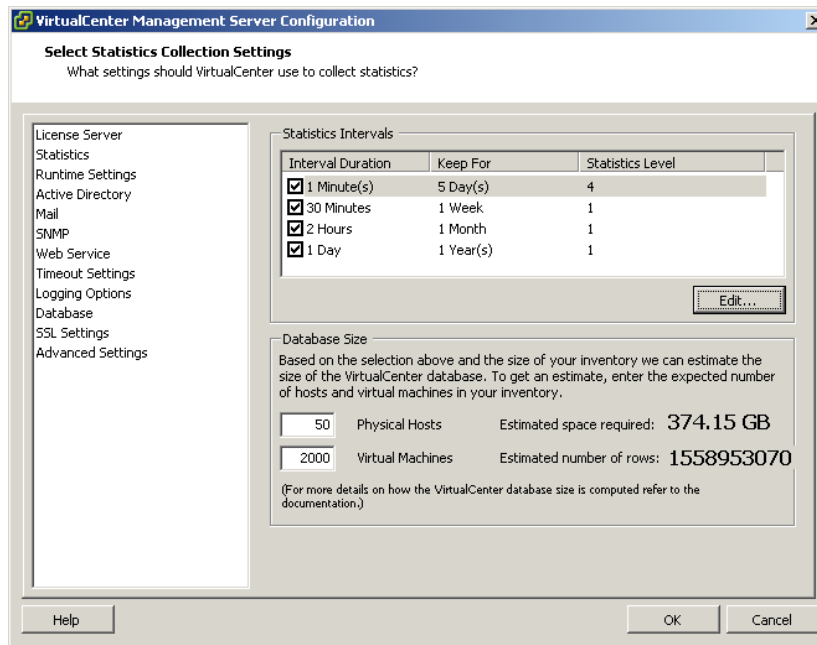


Figure 5.6: Statistics Configured Improperly for Long-Term Use

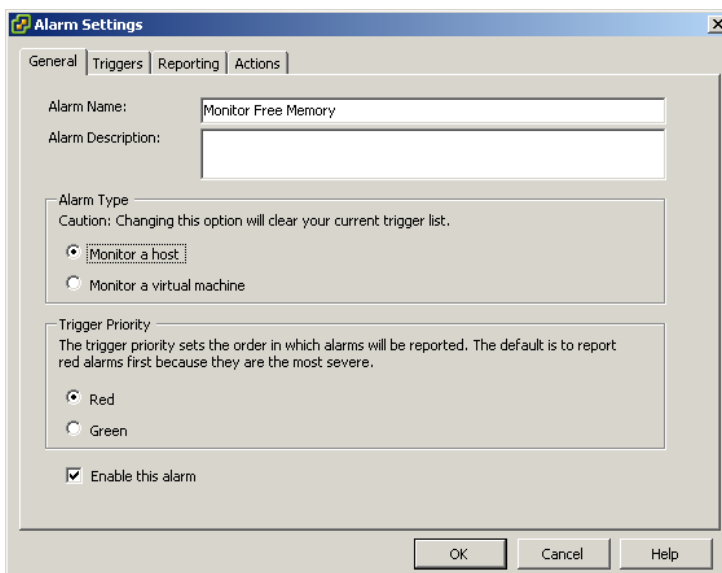
Creating Alarms

Alarms are useful in that they can notify administrators of problematic events. Several hundred combinations of settings can be achieved. The simplest way to understand alarms is to create one. While performing the following steps, take the time to look around at all of the options.

To create an alarm that monitors free memory on a specific host:


1. Launch the VI Client and connect to the VirtualCenter server.
2. Select the Alarms tab and click the Definitions button.
3. Right-click in the white space where existing definitions are listed and select New Alarm.

4. Configuring the General tab as seen here:



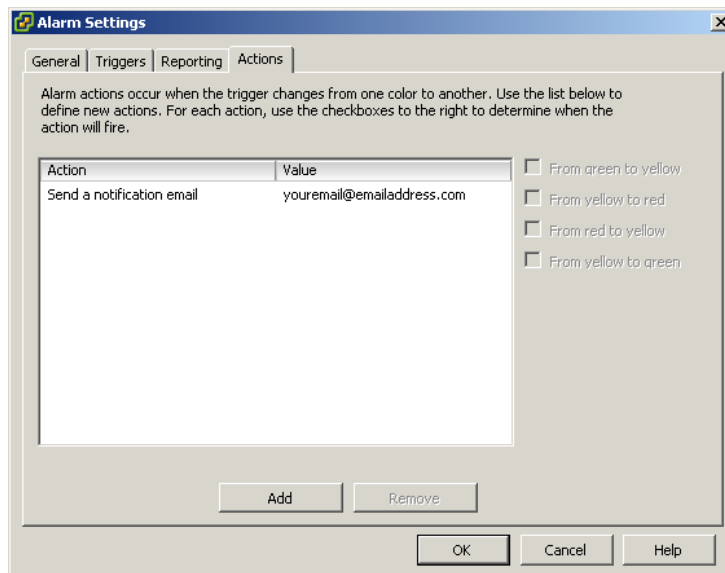
Notice that you can monitor a host or a single VM on a host. You can also indicate whether this is a high (red) or low (green) priority alert.

5. Click the Triggers tab.
6. Click the Add button to add a new trigger.
7. Configure the trigger as seen in the following image:



NOTE: You can click in the Warning or Alert fields and type any value you like instead of choosing a predefined value in the drop down lists.

8. Click the Reporting tab. Determine whether you want the alert to fire multiple times or not. Leave the values at 0 for a single trigger.
9. Click the Actions tab.
10. Click the Add button to add an action for the alert to take when fired.
11. Configure the action as seen in the following image:

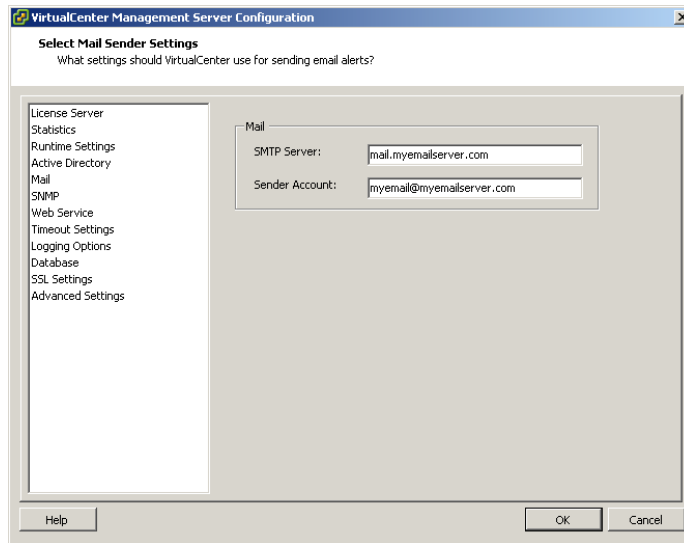


12. Click OK to create the alert.

At this point, you may be wondering how VirtualCenter knows to which hosts or VMs the alert applies. The answer lies in the hierarchy. If you create an alert at a datacenter level, it will only apply to objects in that data center. If you create an alert at the folder level, it will apply to objects in that folder.

In addition to configuring the alert, if you intend for the alert to send an email, you must configure SMTP settings. To configure SMTP settings follow these steps:

1. Connect to the VirtualCenter server using the VI client.
2. Click the Administration menu and select VirtualCenter Management Server Configuration.
3. Select the Mail page.
4. Enter your SMTP server information as seen in the following image:



5. Click OK.

Scheduling Tasks


VirtualCenter 2.5 supports 8 different tasks that may be scheduled. The tasks are:

- Change the power state of a virtual machine
- Clone a virtual machine
- Deploy a virtual machine
- Move a virtual machine with VMotion
- Relocate a virtual machine
- Create a virtual machine
- Make a snapshot of a virtual machine
- Add a host

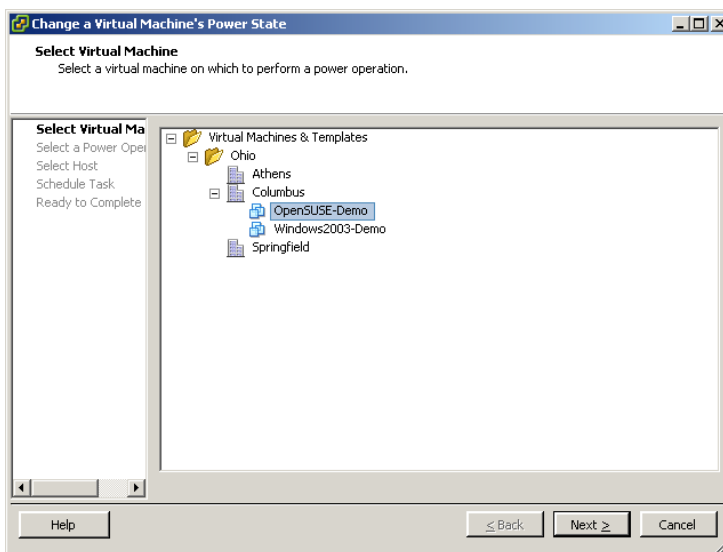
Of these, the least commonly scheduled tasks will likely be creating a VM and adding a host. The following instructions provide an example of scheduled task creation.

To schedule the powering on of a VM:

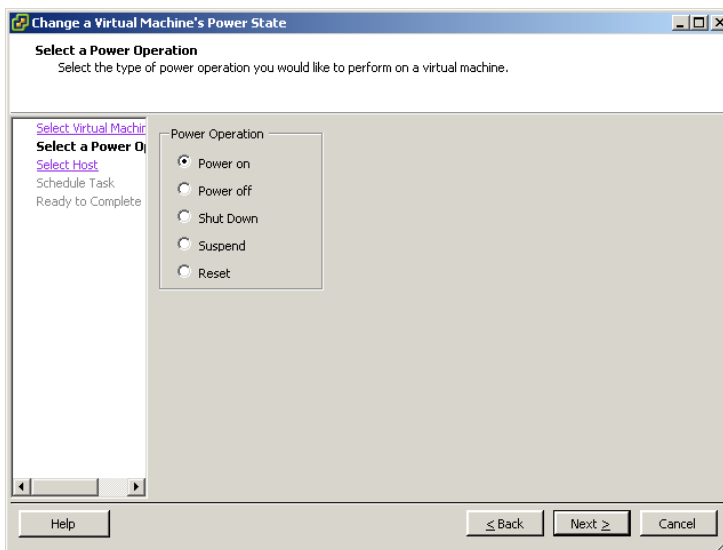
1. In the VI Client, while connected to the VirtualCenter server, click the Scheduled Tasks button.
2. Click the New Task button.
3. Choose *Change the power state of a virtual machine* and click OK.



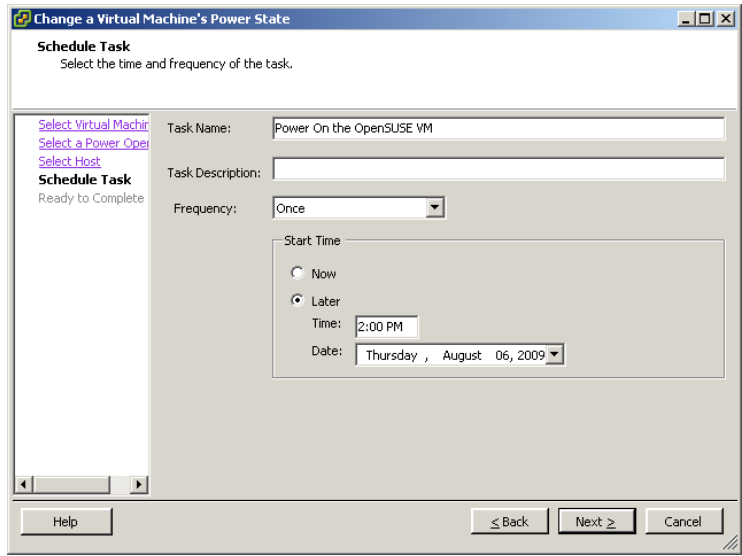
4. Choose the virtual machine you want to power on and click Next.



5. Choose the Power on option and click Next.



6. Choose the host onto which you want to power the VM and click Next.
7. Configure the task properties and the scheduled time and click Next.



8. Click Finish to add the scheduled task.

Once the task is added, it will show up in the list as seen in Figure 5.7. If the task is a one-off task, it will be removed from the list automatically after it runs.




Figure 5.7: Viewing Scheduled Tasks in the VI Client

Lesson 5: Summary

In the lesson, you learned about the features of VirtualCenter and the VMware virtual infrastructure. You installed VirtualCenter and learned to prepare a SQL Server database for an enterprise deployment. Once installed, you learned to configure inventories and add ESX hosts. Finally, you learned to configure several important settings including statistics and alerts.

Lesson 6: Power Administration

When you implement VirtualCenter, as discussed in Lesson 5, you open a whole new world of administrative possibilities. In this lesson, you will explore those possibilities, including:

- Using Templates
- Web Access
- VMware Security Management
- Backup Solutions
- High Availability
- VMware Converter Enterprise

Using Templates

One of the most important features added by VirtualCenter is the availability of templates. Templates are very similar to the traditional imaging solutions (Ghost, TruImage, etc.) that we've used for physical machines for years. Templates can be updated and new VMs can be provisioned from the templates. Templates cannot, however, be powered on.

To create a template, you will follow these high-level steps:

1. Create a VM on which you wish to base the template. Fully configure it according to the needs of the template.
2. Create a template from the prepared VM.

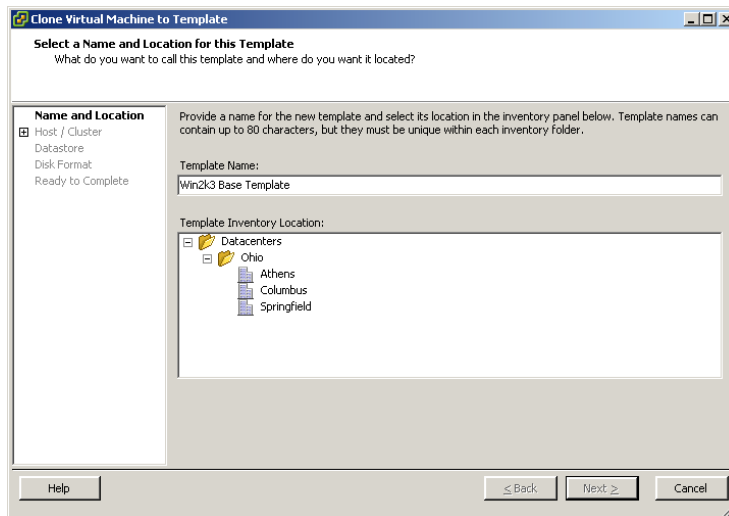
It's really that simple. If you've worked with physical imaging solutions, the process is very similar. With templates, you can create them in one of two ways: cloning or converting. When you clone to a template, the original VM still remains. When you convert to a template, the VM is converted and you can only use it as a template. A VM can only be cloned to a template if it is powered off.

In order to update templates, you convert them back to VMs and then apply patches and updates. Once the updates are complete, you can convert them back to templates. Just right-click on any template and you can select to convert it to a VM.

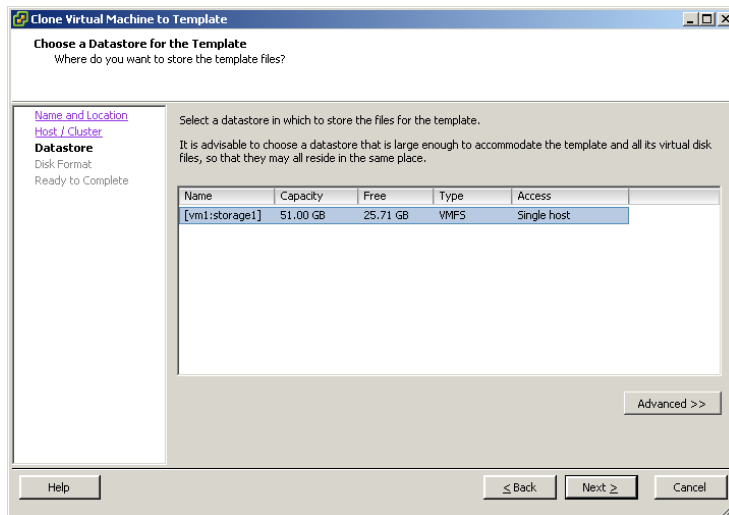
Templates can be stored in two modes: normal and compact. If you store the template in normal mode, provisioning is fast, but the template takes up more space. If you store the template in compact mode, provisioning is slower, but space is conserved. If you have plenty of storage space, it's usually best to use normal mode so that provisioning is less intensive.

To create a template from an existing VM, follow these steps:

1. Using the VI Client, connect to the VirtualCenter server.
2. Locate the VM in the inventory.
3. Right-click the VM and select Clone to Template.
4. Provide a name for the template, select the inventory location and click Next.



5. Choose the host on which to store the template and click Next.
6. Choose the appropriate datastore and click Next.



7. Choose the Normal disk format for fast provisioning and click Next.
8. Click Finish to create the template.

The template creation process can take several minutes depending on the size of the VM on which the template is based and the performance of the datastore. When the template creation is completed, you should see a notification like the following:

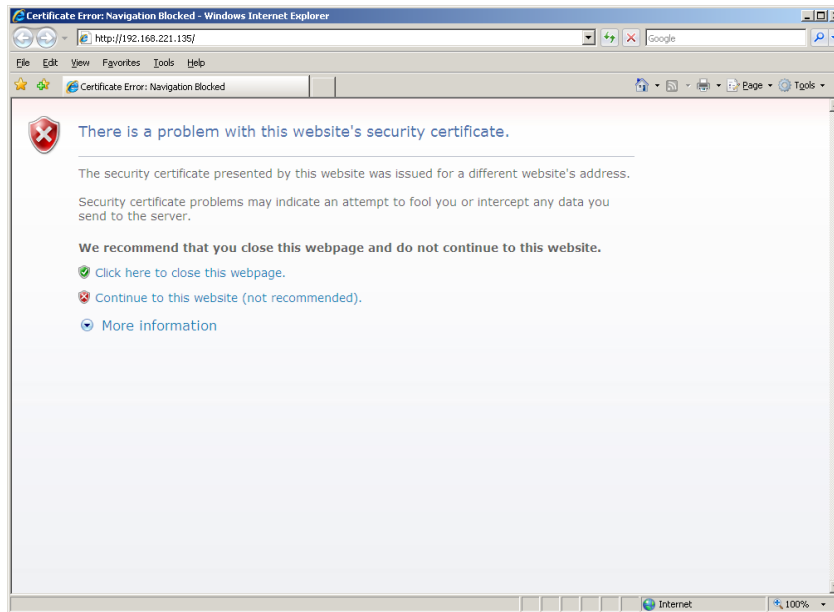
| Recent Tasks | | | | | | | |
|-----------------------|------------------|-----------|------------------|------|----------------------|------------------------|------------------------|
| Name | Target | Status | Initiated by | Time | Start Time | Complete Time | |
| Clone Virtual Machine | Windows2003-Demo | Completed | SYSEDCO\Admin... | | 7/29/2009 11:24:08 A | 7/29/2009 11:24:08 ... | 7/29/2009 11:26:54 ... |

Web Access


In Lesson 2, you learned that you can connect to an ESX host using a web browser. After you install VirtualCenter, you can also connect to the VirtualCenter server with a web browser. With Web Access, you cannot perform all of the same functions that you can with the VI Client; however, it is still a beneficial feature.

To connect to the VirtualCenter server via Web Access, perform the following steps:

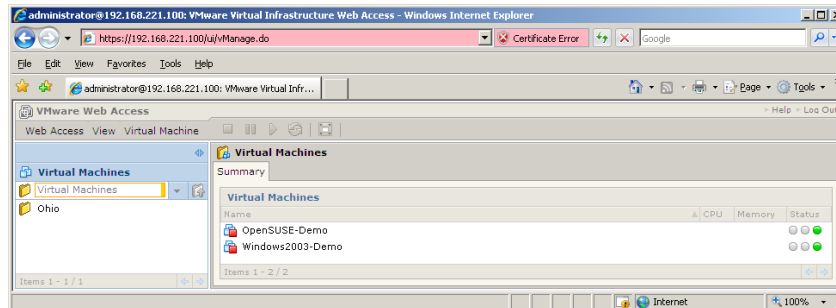
1. Open a web browser on a computer that can connect to the VirtualCenter server.
2. Enter the IP address or host name of the VirtualCenter server in the URL field and append `/ui` to the end of it and press ENTER. (For example: `http://192.168.221.100/ui`)
3. If you see a certificate error message, continue to the website anyway.



4. Enter the credentials for a valid user of Web Access and click the Log In button.



5. You will be presented with an interface similar to the following:



From the Web Access console, you can perform any of the following tasks:

- Manage the power state of VMs.
- Shutdown or restart guests.
- Open and work in a console for any VM.
- View the summary information for a VM.
- Modify settings of existing VMs.
- View the Events, Alarms and Tasks related to a VM.
- Attach a CD or other disk to a VM.
- Configure virtual networking for a VM.

You cannot create VMs from within Web Access. The VI Client must be used for advanced administration tasks like creating VMs.

Web Access requires the following browser versions or newer on Windows:

- Internet Explorer 6.0
- Firefox 1.0.7
- Netscape Navigator 7.0
- Mozilla 1.x

The following minimum requirements hold true for Linux machines:

- Firefox 1.0.7
- Mozilla 1.x
- Netscape Navigator 7.0

VMware Security

With so many ways into the administration of the VMware VI3 environment, it is important to consider security options. You can limit the number of users who have access to the administration tools and you can provide different users with different levels of access.

VMware uses a roles-based model. Several roles exist out-of-the-box and you can create additional roles as needed. Users are assigned to roles in order to grant them permissions. Figure 6.1 represents this process.




Figure 6.1: The Security Process in VI3

Creating Roles

Roles are assigned to users at an object level. For example, you can assign a user the Administrator role at the datacenter level and the users will not be an Administrator in different datacenters.

Three system roles exist in VirtualCenter by default:

- **No Access:** This role provides exactly what it sounds like: no access.
- **Read Only:** The Read Only role is used for support staff such as the help desk. They can view configuration parameters for verification, but they cannot change anything.
- **Administrator:** The Administrator role is all powerful. Any administrative task can be performed by a member of this role.

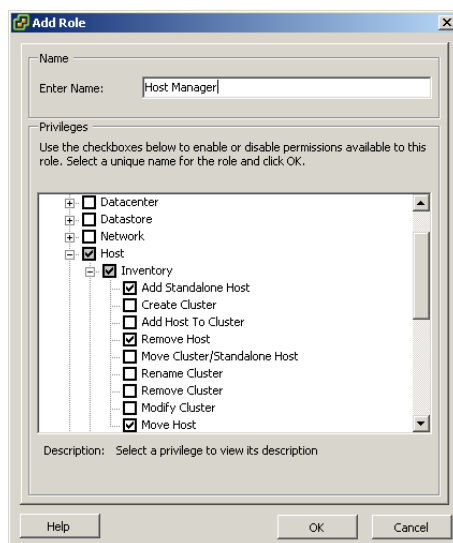
The remaining default roles are sample roles, which exist so that you can see the various ways in which roles may be utilized. The sample roles include:

- Virtual Machine Administrator
- Datacenter Administrator
- Virtual Machine Power Users

- Virtual Machine User
- Resource Pool Administrator
- VMware Consolidated Backup User

To create a new role, following these steps:

1. Click the Administration button in the VI Client while connected to the VirtualCenter server.
2. Select the Roles tab.
3. Right-click on any role and select Add. This will add a new blank role.
4. Provide a name for the role and choose the appropriate privileges.



5. Click OK.

You may also clone an existing role. To do this, right-click the role and select Clone. Finally, you may modify the privileges assigned to a role. To do this, right click the role and select Edit.

Creating Domain Users

Simply creating roles does not provide access to any users. You must assign the users to roles; this action creates a permission. Before you can add users to roles, they must exist in the Windows domain database or in the user database on the local machine.

To add a user to the local machine database:

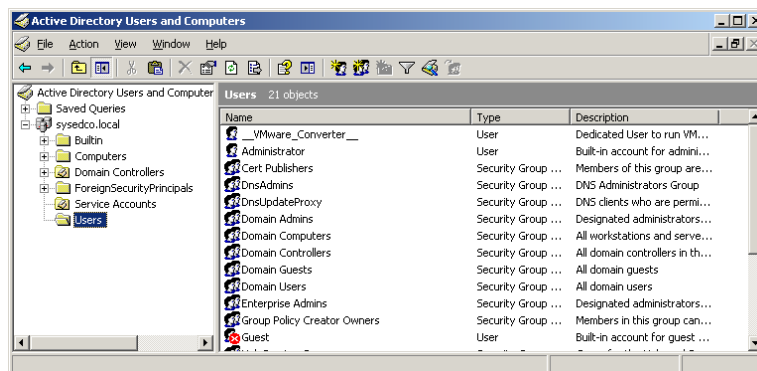
1. Right-click on My Computer and select Manage.
2. Expand Local Users and Groups.

3. Right-click the Users container and select New User.
4. Configure the desired settings for the user.

5. Click Create.

To add a user to the domain database:

1. Logon to a domain controller as a domain Administrator (a member of the Domain Admins group).
2. Launch Active Directory Users and Computers.
3. Select the appropriate container in the left pane.



4. Right-click the container and select New > User.
5. Enter the user name information and click Next.

New Object - User

Create in: sysedco.local/Users

First name: Tom Initials:

Last name: Carpenter

Full name: Tom Carpenter

User logon name: TomCa @sysedco.local

User logon name (pre-Windows 2000): SYSEDCCO\ TomCa

< Back Next > Cancel

6. Enter a valid password and select the appropriate password expiration options and click Next.

New Object - User

Create in: sysedco.local/Users

Password:

Confirm password:

User must change password at next logon

User cannot change password

Password never expires

Account is disabled

< Back Next > Cancel

7. Click Finish to create the user account.

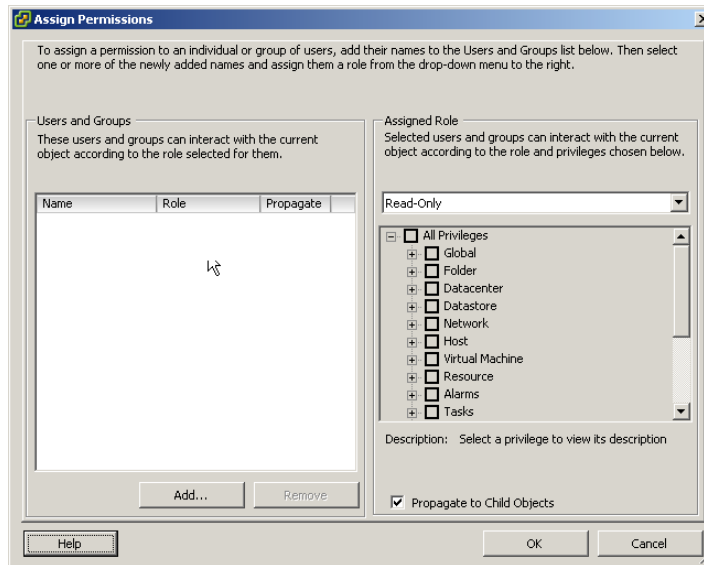
Add Permissions

Now that you've created a role and a valid user account, you can add the user to the role in order to provide that user with the needed privileges.

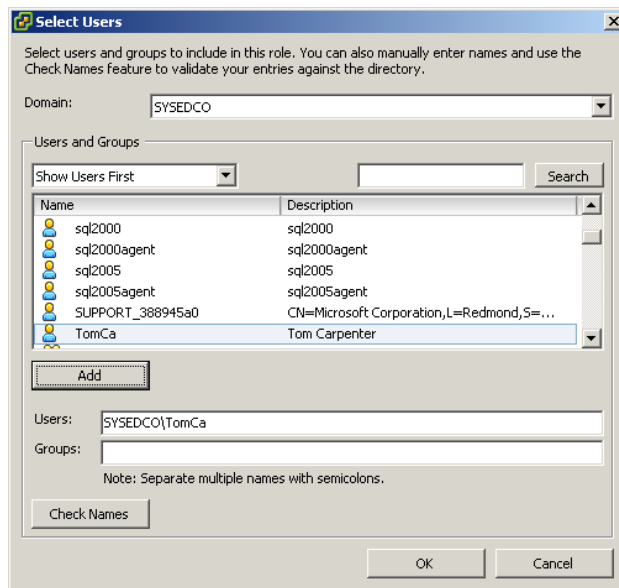
To add a user to a role:

1. In the VI Client, choose the level within the inventory where you want the user to have the role privileges.
2. Click the Permissions tab.

- Right-click in the white space and select Add Permission.

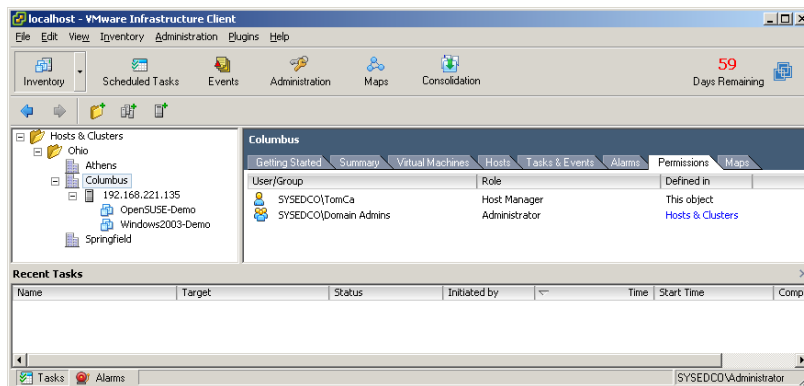


- Click the Add... button to add a Windows user.
- Select the user you created previously and click Add and then OK.



- Choose the role to which you wish to assign the user and click OK.

Once the addition is complete, your Permissions tab should reflect the changes as seen in the following image:



Backup Solutions

Just as you must backup physical servers, the VMs require a stable backup solution. In addition to the VMs, you will need to backup the following components of your VI:

- ESX hosts
- VirtualCenter database
- Active Directory

It's probably safe to assume that some backup solution is already in place for Active Directory; however, you should verify that it is working correctly. The VirtualCenter database will most likely be a SQL Server database, so you can use the built-in SQL Server backup engine for that. The ESX hosts do not change drastically over time - particularly in an environment with VirtualCenter. In such cases, you may choose not to backup the ESX hosts since they can be reinstalled in less than 20 minutes.

Backing Up VMs

VM backups are a top priority. These are the really working servers on your network that provide services and resources to your users and network components. Two traditional options are available:

- Internal backups: With internal backups, you are installing a backup software solution within the VM. This implementation is no different than a physical server. This method results in slower recovery windows than the next method.
- VM file backups: With VM file backups, you are backing up the virtual disks and configuration files that make up the VM. This method may take longer and may require that the VM be powered off, but it results in a fast recovery window.

Using VMware Consolidated Backup

VMware Consolidated Backup (VCB) allows backup and restoration for VMs at the file level or at the image level. VCB uses snapshots in order to perform the backups. VCB provides several benefits:

- Reduced backup processing: VCB creates a snapshot of the VM and then the backup software simply backs up the snapshot files. This backup software is usually running on a different machine than the VM.
- Reduced backup windows: The VM does not have to be taken offline since a snapshot is used. This means that the backup window is non-existent from the viewpoint of availability. However, performance may still be degraded during the creation of the snapshot.

High Availability

High Availability (HA) provides failover solutions for ESX hosts. If you have several VMs running on a single standalone ESX host and that host fails, all of your VMs become unavailable with no method to bring them online again quickly. HA provides a solution to this dilemma. HA does not provide automatic failover (without down time), but it does provide a quick way to restart the VMs on a different ESX host.

In order for HA to operate, all of the ESX hosts must participate in the same cluster. Every 15 seconds, the ESX hosts send out a heartbeat. If an ESX host does not send out its heartbeat, the cluster assumes that host has failed and begins the restart process for the VMs located on the failed host.

In addition to host monitoring, HA can support VM failure monitoring. Only VMs with the VMware Tools installed may participate. VM failure monitoring works from heartbeats like ESX host monitoring. If the VM does not provide a heartbeat, HA attempts to restart the machine. The assumption is that if the heartbeat is not being transmitted, the machine must have crashed in some way.

The following requirements exist in order to implement HA:

- VirtualCenter must be installed and operating properly.
- ESX hosts must use shared storage in order to move VMs during a failure.
- ESX hosts must have the same networks configured so that clients can continue to connect after the VMs are moved.
- All ESX hosts must be able to resolve all other hosts' names to IP addresses using DNS.

While it is not technically required for operation, you should have two NICs (NIC or adapter teaming) dedicated to the Service Console in each ESX host. This means you would need three NICs at a minimum. Two for the Service Console and one for the VMs. Again, this is not required, but you will see error messages from HA if NIC teaming is not used for the Service Console.

In addition to the HA features of VMware, support is also provided for the Microsoft Clustering Service.

VMware Converter Enterprise

VMware Converter Enterprise is a tool that allows for the conversion of multiple sources to VMs. With this tool, you can convert the following existing machines to ESX-compatible VMs:

- Physical machines
- VMware Workstation and GSX Server machines
- Microsoft Virtual Server 2005 machines
- Convert disk images to VMs
- Restore VCB backups to VMs

In addition to these import and convert functions, you may also export ESX VMs to the following formats:

- VMware ESX Server 3.x
- VMware Workstation 4.5.x, 5.x, 6.x
- VMware Player 1.x, 2.x
- VMware Server 1.x
- VMware ACE 1.x, 2.x
- VMware Fusion 1.x
- VMware GSX Server 3.x standalone

VMware Converter Enterprise supports both hot and cold cloning. Hot cloning is used to convert a machine to a VMware VM while the machine is running. Cold cloning converts the machine without the existing OS activated. To perform cold cloning, you will boot from a CD that includes the converter.

Lesson 6: Summary

In this lesson, you learned to use enhanced management features provided by VirtualCenter. These features include templates, Web Access, security management, backups, HA and VMware Converter Enterprise. You learned to implement and configure these features. In the next and final lesson, you will learn to monitor and optimize the performance of your VMware infrastructure.

Lesson 7: Performance Management

If you've worked your way through Lessons 1 through 6, you should have a functioning virtual infrastructure (VI) based on ESX hosts and the VirtualCenter management toolset. Over time, this environment will be more heavily utilized and one of the responsibilities of a VI administrator is managing the performance of his or her environment. VMware provides several tools and features for performance management, including:

- Resource Management
- Resource Pools
- Performance Monitoring
- Performance Settings
- Guided Consolidation

While Guide Consolidation is not marketed as a performance feature, it is such a feature because it helps you consolidate multiple physical servers onto a single ESX host as virtual servers. Even more important, Guided Consolidation makes recommendations based on how well the physical server will perform as VM running alongside other servers. You might say that it does automated performance optimization. Even with a feature like Guided Consolidation, you still have to know the nuts-and-bolts of performance management. Guided Consolidation is great for the initial implementation, but what about six months down the road when 15 percent more users are using the VMs? At that point you have to understand all of the other topics in this lesson.

Resource Management

ESX hosts must be managed effectively in order to achieve the best performance and loading of the servers. Server loading is the process used to ensure that you are getting the most out of your hardware. The last thing you want to do is to use a virtual environment that requires nearly as many physical servers as a physical implementation would have.

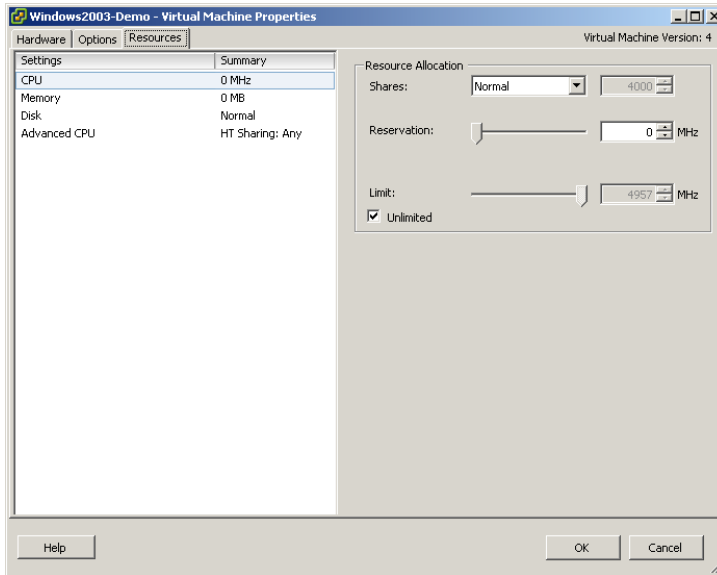
The following technologies are used to help with resource management:

- **vCPU Load Balancing:** VMware virtual CPU (vCPU) technology allows for a physical collection of CPUs to be shared among many VMs. The process used is called *vCPU load balancing*. The architecture works like this (in summary):
 - ▶ VMware implements hardware execution contexts (HECs), which are threads running on the physical processors and vCPUs are given time on these HECs.
 - ▶ The VMkernel decides which VM should get time on an HEC at any given moment.
 - ▶ Every 20 milliseconds, the VMkernel makes a new decision about vCPU to HEC assignments. In order for a VM to gain an HEC assignment (time on the physical processors), one HEC must be available for each vCPU in the VM.
 - ▶ If a VM has two vCPUs and only one HEC is currently available, that VM cannot receive any CPU time.
- **Hyperthreading:** *Hyperthreading* is a feature of certain Intel processes that allows multiple threads to be scheduled on the same processor simultaneously. Hyperthreading is not as good as multiple core (dual- or quad-core) technologies in a VMware implementation. VMs that are CPU intensive will perform far better on multiple-socket (more than one physical processor) or multiple-core servers than on hyperthreading servers.

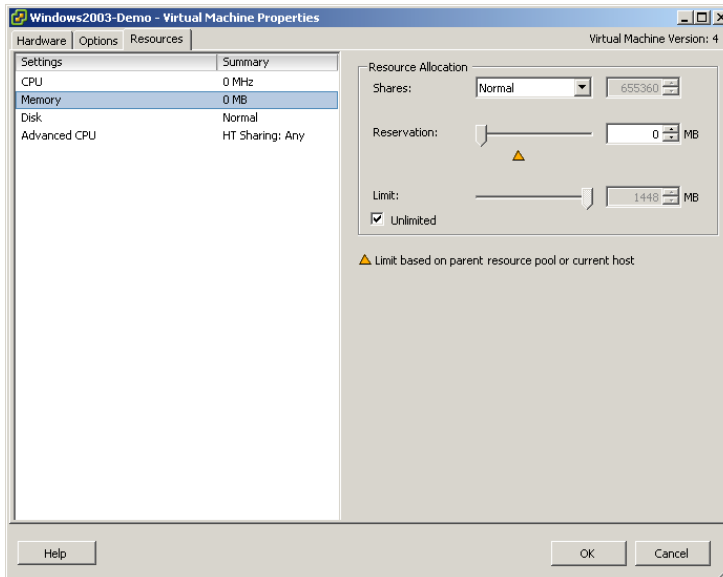
- **CPU affinity:** If you want to constrain a VM so that it only runs on a specified physical processor, you can use CPU affinity. *CPU affinity* is used to affiliate a VM with a specific CPU. This should not be confused with share allocation.
- **Share Allocation:** *Share allocation* is a VMware feature that allows you to prioritize different VMs based on their processor needs. The system is relative. For example, you can assign one VM 1000 allocations and another 500 allocations. Between these two VMs, the first will receive 66 percent of the CPU time. The exact same results could be achieved by assigning the first VM 500 allocations and the second 250.
- **CPU reservations and limits:** In addition to the relative share allocations, you can also set reservations and limits for VMs. For example, you could indicate that VM1 will have a minimum of 25 percent CPU time (reservation) whenever it wants it and, further, that VM2 can never have more than 20 percent CPU time (limit).
- **NUMA:** The Non-Uniform Memory Architecture (NUMA) system divides the processors in a multi-processor system into multiple collections. Each collection has local memory and has access to total system memory as needed; however, access to local memory is much faster. Operating systems that understand how to use the NUMA structure will be able to take advantage of local memory as much as possible, thus improving the performance of the system overall. Other operating systems may run on NUMA structures, but they will not be able to take advantage of local memory. ESX 3.5 can take advantage of NUMA structures if the hardware on which it is running supports them.
- **General memory management:** Several memory management features improve the performance of ESX server as a virtualization solution.
 - *Transparent memory page sharing* is a feature that takes advantage of the fact that multiple VMs may need to access the exact same memory pages. Rather than allocate a separate memory page for each VM, the data is stored once and shared among the different VMs.
 - The *balloon-driver* or *vmmemctl* is installed in the guest operating systems with VMware Tools. This component acts as a memory collector for VMs. For example, if one VM needs more memory, ESX will enlarge the memory “consumed” by the balloon-driver in another VM and then provide the newly available memory to the strained VM. The balloon-driver may consume up to 75 percent of the memory assigned to a VM.

To set CPU and Memory reservations and limits:

1. Right-click on the VM you wish to configure in the VI Client and choose Edit Settings.
2. Click the Resources tab.
3. Choose the CPU page to configure CPU settings.



4. Choose the Memory page to configure memory settings.



5. Click OK to apply your settings.

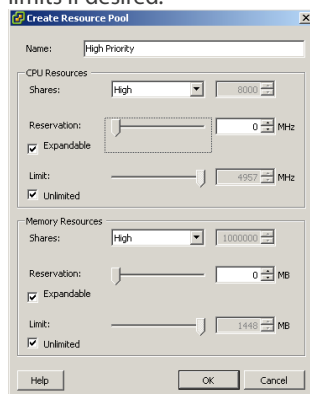
In both the CPU page and the Memory page shown in the preceding instructions, two important settings may be configured. The first is the share allocations (called shares) and the second is the reservation and limit configuration. It is up to the administrator to plan the configuration of these settings. For example, you cannot set four VMs to require a reservation of 2000 MHz on a system with only two 1.6 GHz single core processors.

Resource Pools

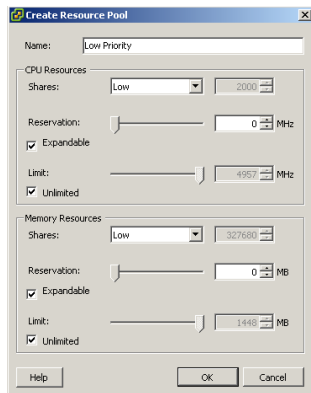
An important alternative to direct CPU and memory resource configuration is the utilization of resource pools. A resource pool is both a configuration set for resource utilization and a container for VMs. You can create multiple resource pools and place multiple VMs in each pool. For example, you may choose to create a Low Priority pool and a High Priority pool. Next, you can place mission critical VMs in the High Priority pool and less critical VMs in the Low Priority pool.

To create High Priority and Low Priority pools, follow these instructions:

1. In the VI Client, right-click on the ESX host where you wish to create the resource pool and select New Resource Pool.
2. Name the pool High Priority.
3. Set the Shares value to High for both CPU and Memory and adjust reservations and limits if desired.



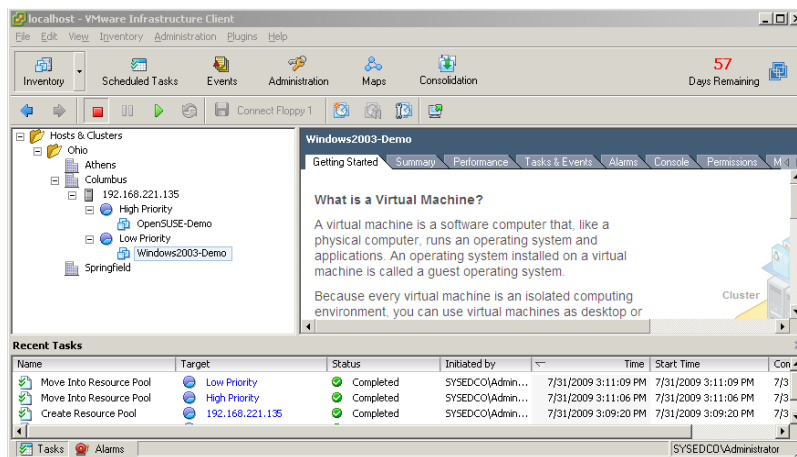
4. Click OK to create the resource pool.
5. Right-click on the same ESX host again and select New Resource Pool.
6. Name the pool Low Priority.
7. Set the Shares value to low for both CPU and Memory and adjust reservations and limits if desired.



- Click OK to create the resource pool.

To place a VM in a resource pool, follow these instructions:

- In the VI Client, click-and-drag the VM into the target resource pool.
- You should see results similar to the following screenshot:



You can also move a VM into a resource pool by migrating it to the resource pool; however, this latter option requires several clicks through a migration wizard. It is much easier to drag-and-drop the VM. Be careful during the drag-and-drop operation. It is easy to inadvertently drop the VM in the wrong container.

Performance Monitoring

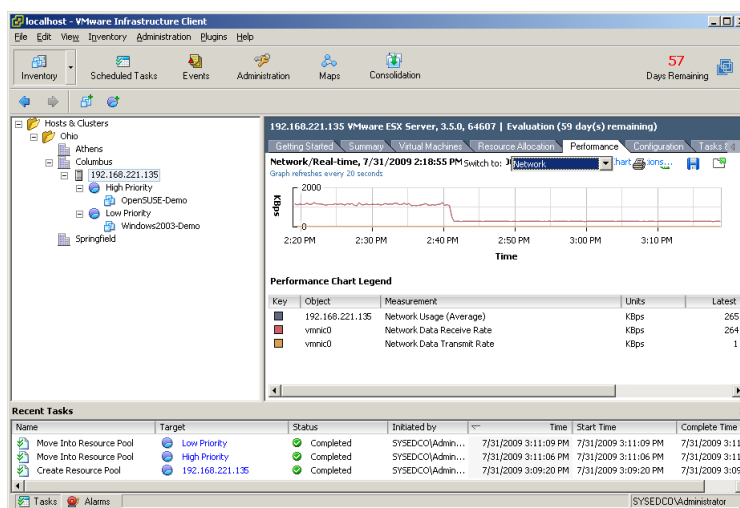
Performance monitoring in a VMware environment must be considered from two perspectives. The first is the performance of the ESX host machines. The second is the performance of the virtual machines. The *Monitoring ESX Host Performance* section will present the various ways in which VirtualCenter exposes the performance of the ESX hosts. The *Monitoring Windows Server VMs* section will provide an overview of the Performance tool provided in Windows.

Monitoring ESX Host Performance

The VI Client provides built-in performance reports for the ESX hosts in your VI. Several performance reports are available from within the VI Client interface.

To view internal VI Client reports:

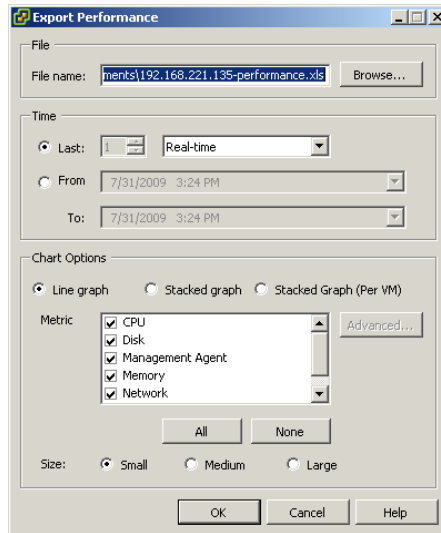
1. Launch the VI Client and connect to the ESX host or the VirtualCenter server.
2. In the VI Client, select the ESX host on which you wish to report.
3. Click the Performance tab.
4. Choose the report you desire from the drop down list. The following image shows the network performance report.



5. If you need a printed report, click the print button.
6. If you need a simple digital report to send to someone, click the save button and provide a name for the JPEG image.

The VI Client may also be used to generate Microsoft Excel reports for any ESX host or VM. To generate an Excel report for an ESX host, follow these steps:

1. Right-click on the host for which you wish to generate a report and select Report Performance. (NOTE: You may also select Report Summary to generate an HTML overview report.)
2. In the Export Performance dialog, provide a name for the report file.
3. Choose the reporting window and select the items on which to report.



4. Click OK to generate the report.

The report generated by the VI Client will look similar to that shown in Figure 7.1 Notice that a tab is provided for each metric measures: CPU, Disk, Memory, Network and System. In addition to the reported data, graphs are automatically generated.

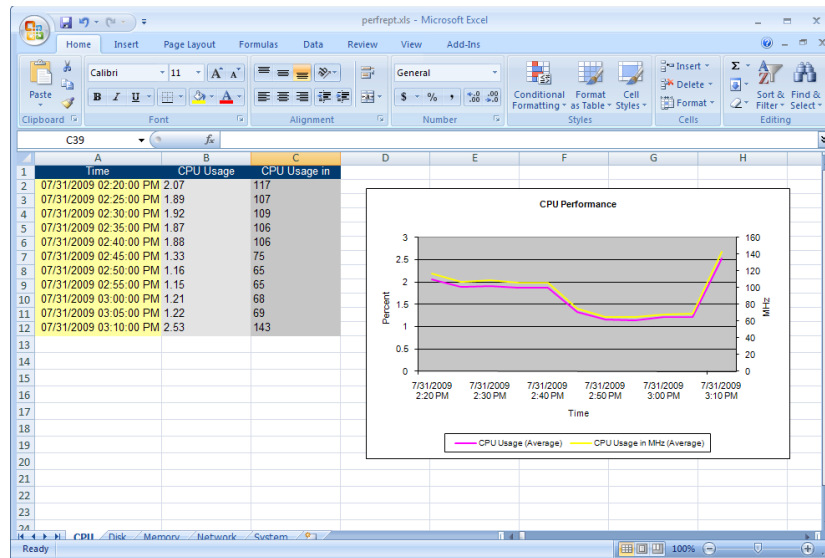


Figure 7.1: VI Client Generated Performance Reports

Monitoring Windows Server VMs Internally

When you want your Windows VMs to perform at their best, the built-in performance monitoring and analysis tools provide insight into potential areas for enhancement. The tools provide you with current performance information and the ability to log this information over time. However, in order to use the tools effectively, you must understand the core performance factors of any server - whether virtual or physical. The performance of Windows servers, Linux servers and any other servers are equally impacted by these factors; this universal impact is based on the truth that these core factors are hardware-related.

The four key server hardware components that can be altered to impact performance are the CPU, memory, hard disks and the network interface card (NIC). Three are internal components (CPU, memory and hard disks) and the fourth is the gateway to the network. Internal server performance determines whether the full NIC capabilities can be utilized. NIC performance determines whether a well-performing internal system matters. As you can see, all four are important and they depend on each other.

While the system bus does impact performance, this component is much harder to replace. The video subsystem is usually not as important in servers as it is in desktop-class machines. Terminal Servers may be an exception to this general guideline. Of course, advanced server hardware – such as RAID controllers, Fibre Channel and iSCSI adapters, can also impact performance, but we will focus on the components that exist in every server as we explore performance analysis in this guide.

Systems Thinking

As you monitor and analyze Windows Server performance, it is essential to employ systems thinking. Systems thinking requires the consideration of the relationships among the hardware components. For example, if CPU utilization is high, the CPU is not automatically seen as the problem. Instead, memory and hard disk utilization is considered. Is the system using an excessive amount of virtual memory? Is the system delayed while waiting on the drive queue to be processed? The CPU utilization may be a symptom of a memory or drive problem rather than evidence of an insufficient CPU speed.

Having performed analysis on hundreds of Windows Servers I've learned one important general guideline: faster CPUs do not always solve performance problems. It's tempting to throw more speed at the problem, but remember the old saying: If a man is lost in a city and he drives faster, he just gets more lost, faster. You could rephrase this for virtual server performance tuning and say that a faster processor just loops faster while it's waiting for the true bottleneck to finish working.

The end result is simple. When analyzing the performance of a Windows virtual server, you should analyze all four core components at the same time. Using this systems thinking process will allow you to locate the true performance bottleneck faster.

CPU Counters

The Performance tool in Windows Server 2003 R2 and earlier, or the Reliability and Performance Monitor in Windows Server 2008 provide several important counters related to the four core components. The key CPU counters are listed under the Processor and Process objects. Here are my favorites:

- % Processor Time
- % User Time
- % Privileged Time

These three counters are available in the Processor object and may be monitored for all CPUs or specific CPUs (see Figure 7.2). They are also available in the Process object and may be monitored for all processes or individual processes. If you notice that the % Processor Time counter is high in the Processor object, it may be helpful to monitor it in the Process object for each individual process. This action will provide you with insight into which processes are monopolizing the processor's time. You may decide to offload some of the processes to a different server or you may even be able to stop running some processes. It's amazing how many unused processes often run on Windows servers.

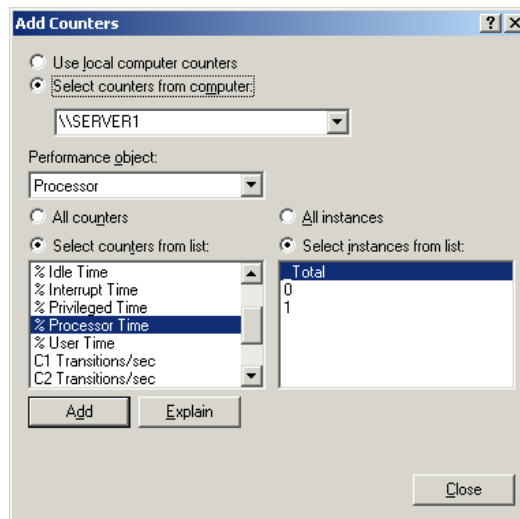


Figure 7.2: Processor Counters

The **% Processor Time** counter is inclusive of both user mode and kernel mode operating system functions. It is technically a measurement of the time in which the System Idle process is not running. The System Idle process runs only when no other process is seeking processor time. It's best to look for average % Processor Time values greater than 65-70 percent before becoming concerned about the processor.

The **% User Time** and **% Privileged Time** counters allow you to monitor user mode and kernel mode activities independently. These counters may help you to determine whether a bottleneck is occurring within an application or within the operating system. However, it is important to remember the architecture of the Windows operating system. Most actions are performed in kernel mode so it is not uncommon to see 70 percent or more of the activity occurring within kernel or privileged mode.

Memory Counters

The most valuable memory counters for general server analysis are located in the Memory object (see Figure 7.3). The memory counters that are most useful include:

- Available KBytes
- Pages/sec

The Available KBytes counter sits happily in the middle between Available Bytes and Available MBytes. The level of detail provided by tracking kilobytes is better than the limited detail of megabytes and the overwhelming detail of bytes.

Pages/sec is used to track the number of virtual memory pages read or written each second. On most systems, a 4 kilobyte memory page is used and so you can multiply the Pages/sec value by 4 to calculate the kilobytes passing to or from the virtual memory file each second.

Both memory counters and processor counters are best measured against a baseline. A baseline provides a picture of the server's performance when it is perceived as performing acceptably. You can create a baseline by logging the key performance counters covered in this article for an acceptable period of time. Some administrators capture the log data for approximately an hour while others prefer to capture it for at least a day.

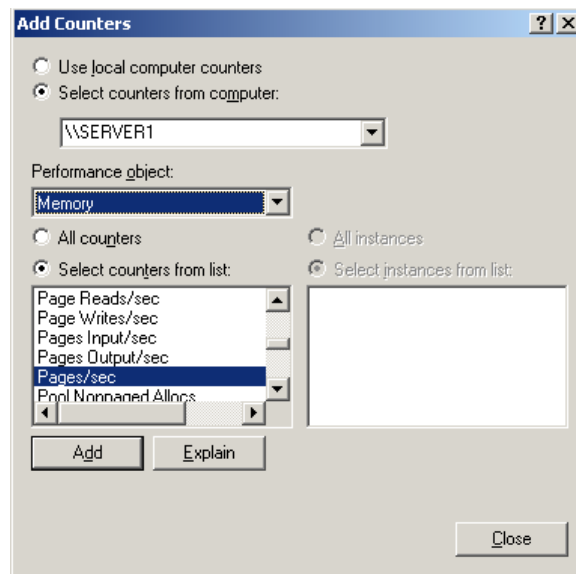


Figure 7.3: Memory Counters

Hard Disk Counters

The hard disk counters are divided into two objects: LogicalDisk and PhysicalDisk (see Figure 7.4). The counters are very similar and the difference is in the way the disks are referenced. LogicalDisk references the disk by the drive letter and PhysicalDisk references the disk by the drive number (for example, drive 0). The key counters to watch include:

- Avg. Disk Queue Length
- Disk bytes/sec
- Free Megabytes

The Average Disk Queue Length counters can reveal whether the drive is keeping up with the demand of running processes. The most frequently cited threshold is 2 items in the queue. If the average is greater than 2, a drive bottleneck may be occurring. This counter should also be compared to the baseline. If the baseline shows an average of 2.3 items in the disk queue and performance was perceived as acceptable, there is no reason to suggest that performance is unacceptable – at a later time – if the average is the same or lower. Remember, performance is measurable with statistics, but whether performance is “good” or “bad” is a relative perspective issue.

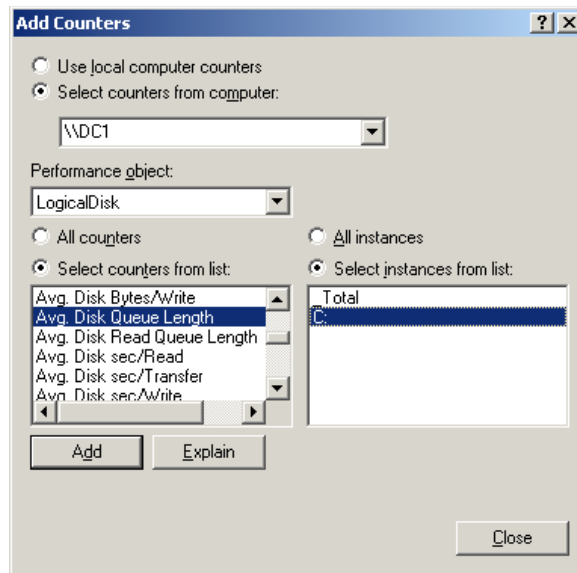


Figure 7.4: Hard Disk Counters

The Disk bytes/sec counter can reveal whether the drive is living up to expectations or not. Many drives are rated at a certain speed, but they perform at lower speeds. This counter can reveal such behavior. In many cases, updating drive controller drivers may resolve such performance issues.

Free Megabytes is not really a performance counter, but it is very useful in predicting future needs. For example, if you measure Free Megabytes for each volume once per month, you can determine consumption rates. With these consumption rates documented, you can predict when you will need to archive old data or upgrade to larger hard drives.

Network Counters

The final set of counters presented in this article is the network counters. These counters are found in the Network Interface object. The two key network counters are:

- Bytes Total/sec
- Output Queue Length

The Bytes Total/sec counter should be compared to the baseline. If this has increased dramatically, it is only certain that the server is more heavily utilized than it was when the baseline was captured; however, it could also be a sign of a network attack or the need for the offloading of some processes. The Output Queue Length may help you decide. If this counter is averaging more than 2, it indicates that the network card (or the data rate of the infrastructure) is not able to handle the capabilities provided by the server. Stated differently, the server is throwing data at the NIC faster than it can transmit it out on the wire.

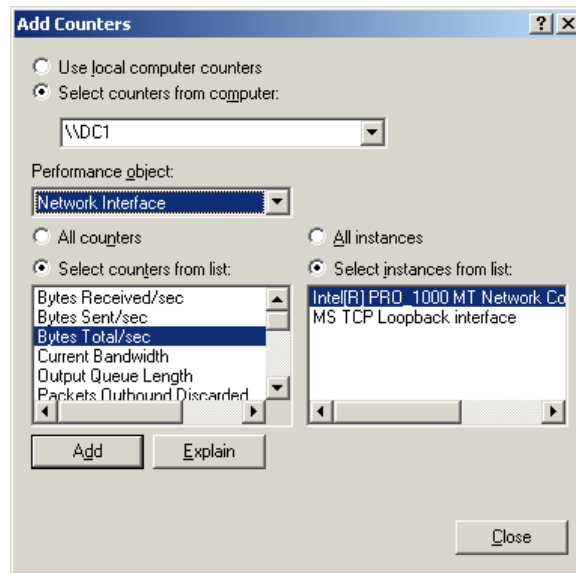


Figure 7.5: Network Counters

Capturing Counters

Now that you have discovered the ten most important counters that help you track the core performance factors in your virtual server, let's look at the process used to capture these counters. The following instructions assume you are using Windows Server 2003 or Windows Server 2003 R2.

1. Launch the Performance tool by click Start > All Programs > Administrative Tools > Performance. You will see that the Pages/sec, Avg. Disk Queue Length and % Processor Time counters are already loaded.
2. Click the button with the plus sign icon (or press CTRL + I) to add more counters.
3. The Processor object is selected by default. With this object select, click the % Privileged Time counter in the Select counters from list section and then click the Add button.
4. Next, scroll down, if necessary, and select the % User Time counter and then click the Add button.
5. Now, select the LogicalDisk object and add the Disk bytes/sec and Free Megabytes counters using the same methods that you used for the % Privileged Time and % User Time counters.
6. Select the Memory object and add Available KBytes.
7. Finally, select the Network Interface object and add the Bytes Total/sec and Output Queue Length counters.

After performing this operation, you should see something similar to Figure 7.6. You may have more or less activity on your server depending on current operations. This exercise results in the ability to monitor live activity. Monitoring live activity is just one way to use this powerful performance tool. In addition, you may want to create a performance log.

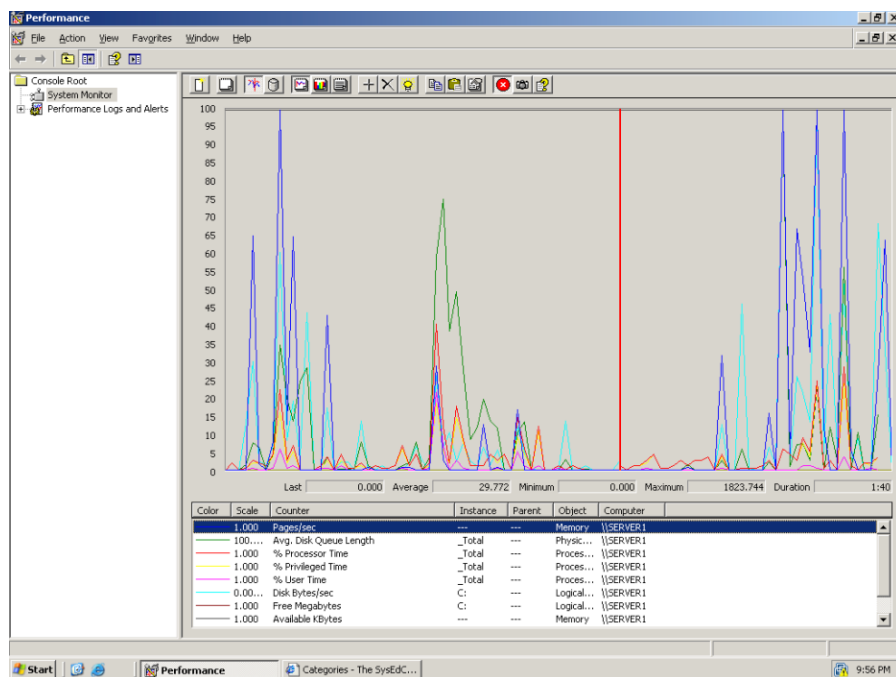
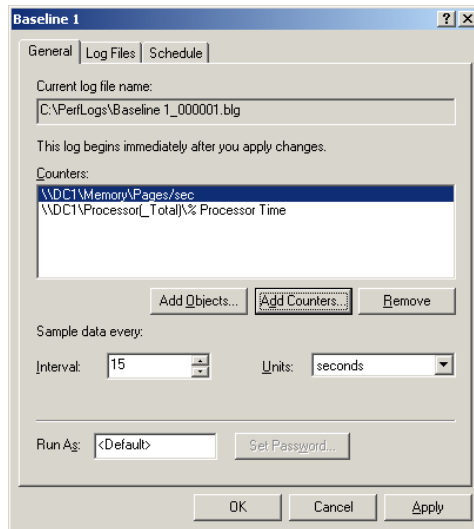


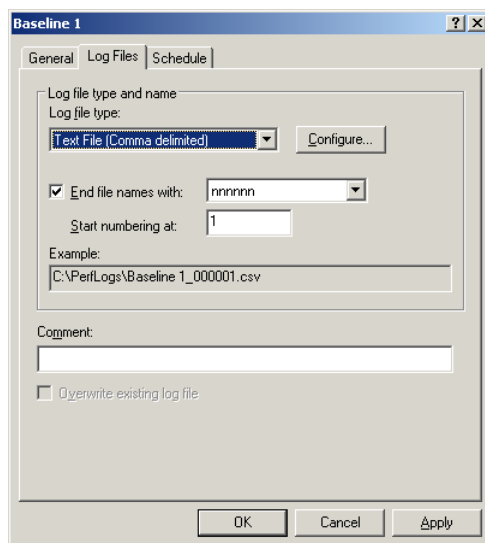
Figure 7.6: Live Performance Activity

Follow these instructions to create a log that will capture performance data for any desirable length of time:

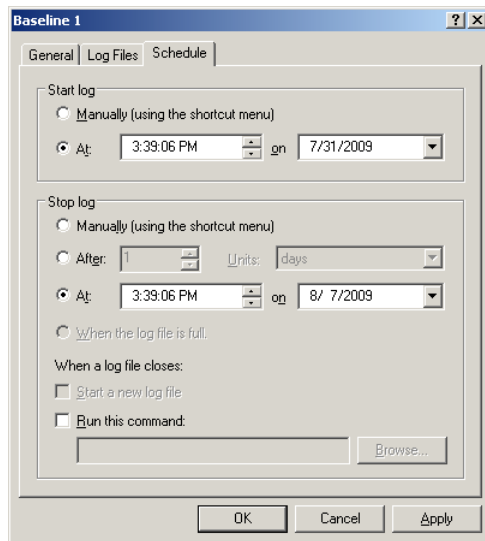
1. Launch the Performance tool as previously described.
2. Expand the Performance Logs and Alerts node in the left pane.
3. Right-click Counter Logs and select New Log Settings.
4. Enter a name for the log such as Baseline 1.
5. Click the Add Counters button on the General tab and add the counters you wish to log.



6. After adding the counters, select the Log Files tab and select the log file format you desire (I prefer the comma delimited text file so that I can analyze the data easily in Excel).



7. On the Schedule tab, either schedule a start and stop time for the log or set it to start manually. Note that you can run a command after the log is created.



8. Click OK.
9. If prompted to create the log directory, choose Yes.

That's really all there is to it. You now have a performance log configuration. If you created the log configuration with the ten counters covered in this guide, you have an excellent configuration for creating baselines. Use this log to capture a baseline of your server's performance when it is perceived as performing well. Then, when users inform you that it is not performing well, you can run the log again and compare the two log files. Figure 7.7 shows line graphs generated in Excel 2007 from CSV log files created in the Performance tool.




Figure 7.7: Performance Measurements shown in Excel

Performance Tuning

A major question must still be answered. After all of the performance analysis of the ESX hosts and the virtual machines, how do you resolve performance problems? The answer: you tweak what needs tweaking or replace what needs to be replaced. Ultimately you have four potential high-level solutions to performance problems in your VI:

- Reconfigure the resource settings for the problematic ESX host or hosts.
- Reassign VMs across your ESX hosts so that performance is improved.
- Upgrade hardware to improve performance.
- Convert a poor performing VM back to a physical machine.

The last option should be considered a last resort. You do not want to go from VM to physical if you do not have to. This action would undermine the very value that virtualization brings. However, we do work in a real world and, sometimes, servers simply cannot be virtualized. For example, if you have a physical server that is using 8 physical CPUs, each running at more than 2 GHz, and it contains 16 GB of RAM and the server is 75 percent utilized, you probably will not be able to virtualize that server - at least not today. However, in just a few years, it may be a prime candidate for consolidation.

Just like physical servers, the primary candidates for replacement are the core four:

- CPU
- Memory
- Hard Disks
- Network Interface

Your performance analysis should reveal which of these core performance factors seem to be the problem. You may need both faster processors and faster hard disks in order to achieve the performance gain you desire. Remember this formula for internal system performance:

$$\text{CPU} + \text{Memory} + \text{Hard Disks} = \text{Performance}$$

Notice that I did not include the system bus because this is not something we usually replace without replacing the entire server. Another important performance formula is as follows:

$$\text{Local Network Interface} + \text{Network Infrastructure} + \text{Remote Network Interface} = \text{Network Performance}$$

The point of this second formula is to remind us that network performance problems are not always localized. We may need to replace or upgrade infrastructure devices or even change some factor on the remote node.

Guided Consolidation

The Guided Consolidation or simply Consolidation feature of VirtualCenter helps you pick the physical computers that would be good candidates for virtualization. Consolidation uses a three phase process represented in Figure 7.8.

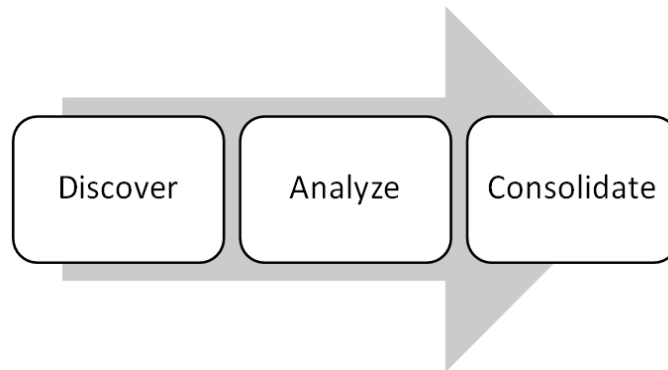


Figure 7.8: The Three Phases of Guided Consolidation

In the Discover phase, the consolidation tools locate physical machines that may be candidates for eventual consolidation. The candidates are limited to the scope of the discovery. You can choose the specific systems you want analyzed and the Guided Consolidation process will only evaluate those systems.

In the Analyze phase, information is gathered from the target systems. Memory, CPU information and performance statistics are gathered in order to provide the best possible automated recommendations.

Finally, in the Consolidate phase, the systems are placed on VMs according to your preferences.

NOTE: If you receive errors when attempting to work with the Consolidation interface in the VI Client, make sure the VMware Capacity Planner Service is running. Go to Start and right-click My Computer and select Manage. Expand Services and ensure the status for the service is Started.

Figure 7.9 shows the interface used to select the scope for consolidation. You first select the domain and then choose the computers from that domain that you think are potential candidates for consolidation. When you click “Add to Analysis,” you may be prompted for credentials that can be used to connect to the machine with administrative capabilities. You can access the dialog shown in Figure 7.9 in two ways:

- Click the Consolidation button in the VI Client. On the Getting Started tab, click Analyze physical computers for consolidation (don't let the phrasing fool you; you can analyze other virtual machines too if you like).
- Click the Consolidation button. On the Analysis tab, click the Add to Analysis button.

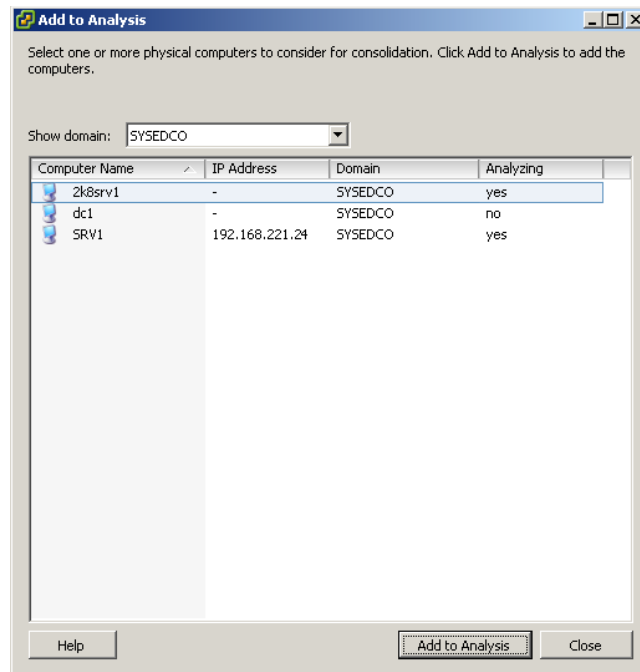


Figure 7.9: Selecting the Scope or Computers for Consolidation

On the Analysis tab, you can view that status of the analysis for each machine. If you are analyzing several machines at the same time, the entire process can take several hours. Even when you are analyzing a few machines, the process can take more than thirty minutes. Be patient and wait for the process to complete. In Figure 7.10, you can see that one machine's analysis is complete and another is still processing.

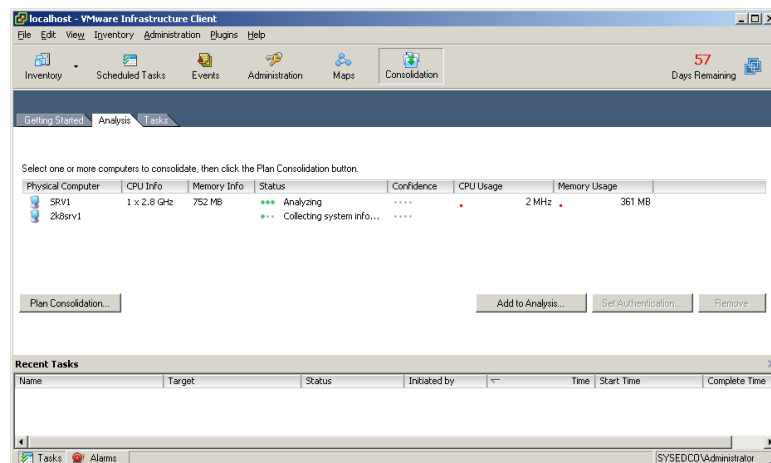


Figure 7.10: Viewing the Analysis Process

Once the analysis is completed, you can click the Plan Consolidation button and see the recommendations made by the Guided Consolidation process. The confidence in the recommendation is based on the amount of time VMware had to analyze the candidates. If a full 24 hours was provided - meaning that a typical day of work was analyzed, the confidence will be very high.

To see the recommended consolidation plan, click the machine for which you wish to see the recommendation and then click the Plan Consolidation button. You'll see a screen similar to the one in Figure 7.11. In this screen, you can change the destination if you disagree with the recommendation. When a destination has five stars, it means that the Guided Consolidation tool is very confident that the destination can handle the new VM that will be created.




Figure 7.11: Planning the Consolidation

If you click the Consolidate button, know that the process begins immediately (this is one of the reasons that these tools are only recommended for smaller environments). The process may take several minutes or even hours depending on the size of the data set on the physical machine. Figure 7.12 shows the import process at 2 percent completion.

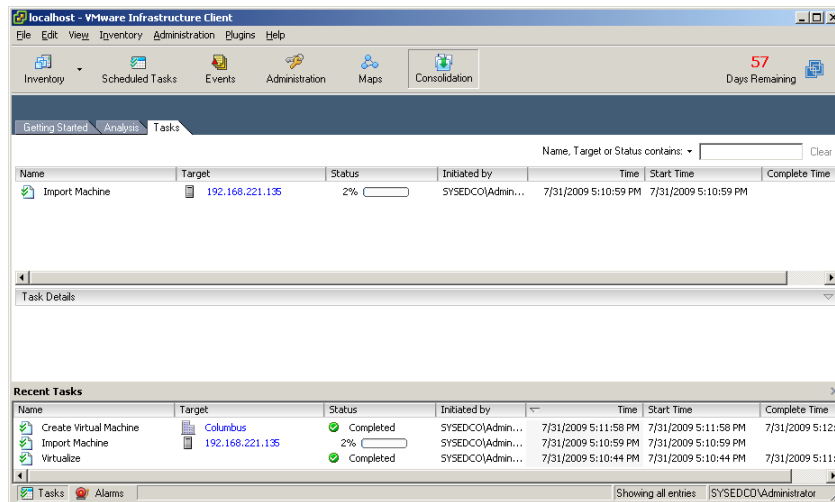


Figure 7.12: The Import Process of Guided Consolidation

Finally, Figure 7.13 shows the VM that was created by the Guided Consolidation process. In Figure 7.11, you saw that the original machine's name was SRV1 and in Figure 7.13 you can see that the Guided Consolidation tools named the VM SRV1 as well. Interestingly, if you look at the recent Tasks pane in figure 7.13, you can see that the import is still not complete, but the VM already shows up. In fact, you can right-click the VM at this point and select Edit Settings to see exactly what settings the consolidation tools assigned to the VM.

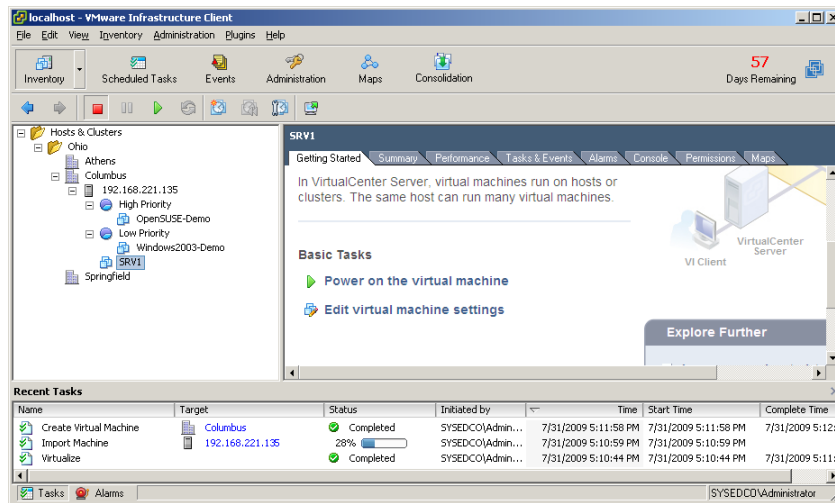


Figure 7.13: The VM Created by the Guided Consolidation Tools

Once the VM has been created completely (the Import Machine task is completed with a green circle and a white check mark), you can modify the settings. For example, you may want to assign more processors, add more memory or include extra drives.

Lesson 7: Summary

In this lesson, you learned about the resources you can manage that impact performance. You learned to create and configure resource pools. The Windows Performance tool was explored because it can be used in VMs just as it can be used in physical machines. You also learned about the different ways in which you can tweak the performance of your VMs and your ESX hosts. Finally, you learned how to use the Guided Consolidation feature to locate and migrate physical machines to virtual machines.

Appendix: VCP 3.5 Objectives Map

While this VMware administrator's guide is not intended to be solely a VCP study guide, it does cover the objectives sufficiently for an experienced administrator to prepare for and pass the VCP 3.5 exam. This appendix provides a mapping of the objectives to the lessons that cover those objectives.

| Objective | Lessons |
|---|-----------------------|
| 1.1 - Identify the characteristics of the products in the VMware product line and when to use each product | Lesson 1 |
| 2.1 - Install VMware ESX Server in different environments including local storage and boot from SAN | Lesson 2 |
| 2.2 - Troubleshoot ESX Server problems | Lesson 6 and Lesson 7 |
| 3.1 - Given a set of requirements and a particular hardware environment, create and configure a virtual switch | Lesson 3 |
| 3.2 - Modify the virtual switch for advanced switch capabilities including VLAN, NIC team, traffic shaping, and layer 2 security policies | Lesson 3 |
| 4.1 - Connect a VMware ESX Server To Fibre Channel SAN storage | Lesson 3 |
| 4.2 - Given a VMware ESX Server attached to the appropriate storage, create and modify VMFS datastores | Lesson 3 and Lesson 5 |
| 4.3 - Connect a VMware ESX Server to NFS storage | Lesson 3 |
| 4.4 - Connect a VMware ESX Server to iSCSI storage | Lesson 3 |
| 5.1 - Given a set of requirements, deploy a VirtualCenter implementation | Lesson 5 |
| 5.2 - Describe VirtualCenter features | Lesson 5 |
| 6.1 - Given a set of requirements, deploy a virtual machine | Lesson 4 |
| 6.2 - Understand applications for VMware Converter Enterprise | Lesson 4 |
| 6.3 - Given a business or technical need, migrate a Virtual Machine using either a standard migration or migration with VMotion | Lesson 5 |
| 6.4 - Utilize Guided Consolidation capabilities | Lesson 7 |
| 7.1 - Create and modify user permissions in VirtualCenter | Lesson 6 |
| 7.2 - Create and modify user permission in ESX Server | Lesson 6 |

| Objective | Lessons |
|---|-----------------------|
| 7.3 - Restrict access to virtual machines using Web Access | Lesson 6 |
| 8.1 - Given an ESX Server, describe the technologies used for resource management | Lesson 7 |
| 8.2 - Given a resource constrained system, make adjustments to improve performance | Lesson 7 |
| 8.3 - Given an ESX environment, allocate resources using resource pools | Lesson 7 |
| 8.4 - Given an ESX Server environment, create and configure DRS clusters | Lesson 1 |
| 9.1 - Given a virtual machine, monitor performance of that virtual machine including CPU, memory, disk and network performance | Lesson 7 |
| 10.1 - Given a VMware Infrastructure environment, develop virtual machine back-up and restore strategies | Lesson 6 |
| 10.2 - Describe the basic operation of VMware HA and other high availability technologies | Lesson 6 |
| 11.1 - Given an environment, determine whether VMware ESX Server could be installed | Lesson 2 |
| 11.2 - Given a set of requirements, plan a VirtualCenter implementation | Lesson 5 |
| 11.3 - Given a set of requirements, plan VI storage deployment | Lesson 2 and Lesson 3 |