

DEPLOYMENT GUIDE

INDUSTRY: ALL

ESRI ArcGIS Server 9.3 for VMware Infrastructure

Deployment and Technical Considerations Guide

December 2009



Table of Contents

Introduction	1
ESRI ArcGIS Server 9.3 Overview	1
VMware Infrastructure Overview.....	1
ArcGIS Server Architecture and Deployment Strategy	3
Architecture	3
Deployment Strategies.....	5
Single-machine deployment scenario.....	5
Multiple-machine deployment scenario	6
Testing Process and Results.....	7
Hardware Configuration.....	7
Physical Server Configuration	7
Virtual Machine Configuration	7
Workload Used	8
Results Observed.....	8
Deployment Best Practices	11
Summary.....	12
Resources	13

Introduction

This document describes the best practices for running a typical ESRI ArcGIS® Server deployment on VMware® Infrastructure. It provides basic guidance on the architectural design of an ArcGIS deployment and the value of utilizing the VMware platform to provide virtual infrastructure. It then discusses results of recent testing conducted jointly by VMware and ESRI characterizing the performance and functionality of ArcGIS Server running on VMware infrastructure version 3.5i update 3. Finally, it outlines some best practices for utilizing the two products together in your datacenter.

ESRI ArcGIS Server 9.3 Overview

ESRI's ArcGIS Server connects people with the geographic information they need. Organizations use ArcGIS Server to distribute maps and other GIS capabilities provided by Web mapping applications and services to improve internal workflows, communicate vital issues, and engage stakeholders. With ArcGIS Server, you can

- **Publish fast, intuitive Web maps tailored to your audience:** dramatically strengthen business and resource decisions with real-time geointelligence.
- **Geographically enable your IT investments:** shrink data and application redundancy, optimizing system configurations and consolidating enterprise systems.
- **Centrally manage your geodata:** provide better data security and integrity for your most important information assets.
- **Simplify access to large volumes of imagery resources:** significantly reduce storage costs and data processing overhead.
- **Extend GIS to your mobile workforce:** increase the accuracy and value of field data collection projects and asset monitoring, as well as provide resource and event management.

VMware Infrastructure Overview

VMware infrastructure is a leading virtualization solution that provides multiple benefits to IT administrators and users. VMware infrastructure provides a layer of abstraction between the resources required by an application and operating system, and the underlying hardware that provides those resources. Benefits offered by VMware infrastructure include the following:

- **Consolidation:** VMware virtualization technology allows multiple physical servers to be consolidated into one server, with little or no decrease in overall performance. Consolidation usually leads to a reduced total cost of ownership (TCO) for running ESRI ArcGIS Server in VMware virtual machines.

- **Ease of Provisioning:** VMware infrastructure encapsulates an application into a self-contained image. With the addition of VMware vCenter™ management software, it becomes possible to leverage templates to create a golden master of a virtual machine environment that can be duplicated or moved. This process can greatly reduce the cost and time of application provisioning and deployment as well as make it easier to scale ArcGIS Server deployments.
- **Manageability:** Virtual machines may be moved from server to server with no down time using VMware VMotion™ technology. VMotion also simplifies common operations such as planned downtime for hardware maintenance.
- **Availability:** VMware High Availability (HA) can shorten unplanned downtime and provide higher service levels to an application. In the case of an unplanned hardware failure, VMware HA restarts any affected virtual machines on another host in a VMware cluster.

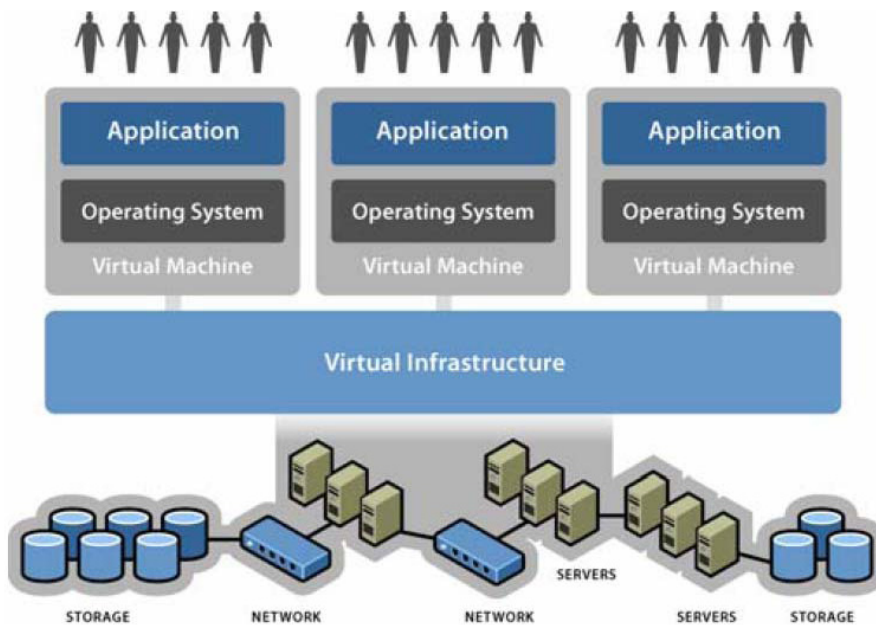


Figure 1. VMware Infrastructure Solution

ArcGIS Server Architecture and Deployment Strategy

Architecture

An ArcGIS Server system includes the following components:

- **GIS server** – The GIS server hosts GIS resources, such as maps, globes, and address locators, and exposes them as services to client applications.

The GIS server itself is composed of two distinct parts: the Server Object Manager (SOM) and Server Object Containers (SOCs). As the name implies, the SOM manages the services running on the server. When a client application requests the use of a particular service, it is the SOM that actually provides the service for the client to use.

The SOM connects to one or more SOCs. The SOC machines host the services that the SOM manages. Depending on your configuration, you can run the SOM and SOCs on different machines and also have multiple machines to run SOCs. Figure 2 below shows a SOM machine connected to two SOC machines.

- **Web server** – The Web server hosts web applications and services that use the resources running on the GIS server.
- **Clients** – Clients are Web, mobile, and desktop applications that connect to ArcGIS Server Internet services or ArcGIS Server local services.
- **Data server** – The data server contains the GIS resources that have been published as services on the GIS server. These resources can be map documents, address locators, globe documents, geodatabases, and toolboxes.
- **Manager and ArcCatalog administrators** – ArcGIS Server administrators can use either Manager or ArcCatalog to publish their GIS resources as services.
 - Manager is a Web application that supports publishing services, administering the GIS server, creating Web applications, and publishing ArcGIS Explorer maps on the server.
 - ArcCatalog includes a GIS Server node that can be used to add connections to GIS servers for either general server access or administration of a server's properties and services.
- **ArcGIS Desktop content authors** – To author GIS resources such as maps, geoprocessing tools, and globes that will be published to your server, you need to use ArcGIS Desktop applications such as ArcMap, ArcCatalog, and ArcGlobe. Additionally, if you are creating a cached map service, you will need to use ArcCatalog to create the cache.

The following illustration shows the ArcGIS Server system architecture.

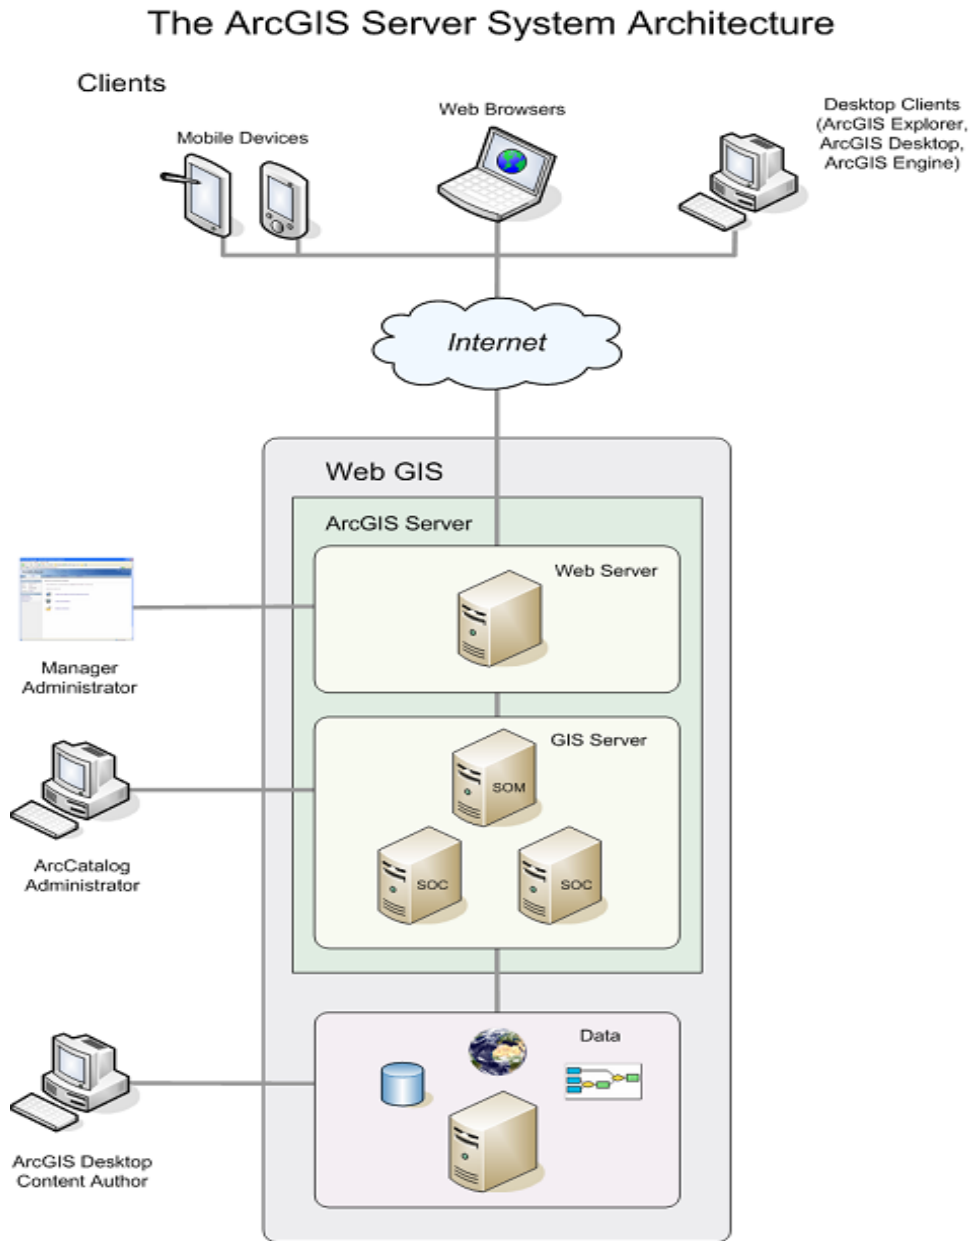


Figure 2. ArcGIS Server System Architecture

For more information, see the *ArcGIS Server Installation Guide* available at:

http://resources.esri.com/help/9.3/arcgisserver/install_gds/dotnet/Step_2_Server_site_config.htm

Deployment Strategies

The scalable architecture of ArcGIS Server provides numerous deployment options. Large deployments may require multiple Web servers, Server Object Managers (SOMs), and Server Object Containers (SOCs), while smaller organizations may want to consolidate these resources on only one or two machines. How you deploy ArcGIS Server depends on what you want to do with it and the workload demands it needs to handle. If you are using the product for development or testing purposes, you probably don't need a very large deployment configuration, but if you are publishing GIS services to be accessed by a large community of users, you need to provide extra consideration and resources to deal with factors such as increased processing loads, eliminating single points of failure, and added security.

The following deployment scenarios are presented as guides for you to consider as you prepare to build your own ArcGIS Server system. Although you could deploy your system exactly as presented in each of these scenarios, you may want to use the scenarios just to get ideas of what is possible to implement with ArcGIS Server and then adjust your own deployment configuration to fit your specialized needs and hardware resources.

Single-machine deployment scenario

If you're using ArcGIS Server exclusively for development, testing, or small deployments, and you don't have to accommodate large numbers of requests for services, a simple configuration is probably sufficient. This first scenario shows how you can install all the ArcGIS Server components on a single machine, along with the Web server, which also resides on the same machine. The machine must also have access to an ArcGIS Server administrative interface, such as Manager.

Using this first deployment scenario, the data needed by ArcGIS Server resides on the same machine as all the other components. If ArcSDE is used to access the data, then the data is most likely stored in a Microsoft SQL Server Express database. ArcGIS Server Workgroup uses the configuration shown in Figure 3.




Figure 3. Single Machine ArcGIS Server Deployment

Multiple-machine deployment scenario

The multiple-machine deployment scenario is ideal for many internal deployments and Web deployments. In this scenario, the SOM and Web server reside on the same machine. Since the SOM uses relatively little memory, it can usually coexist with the Web server without conflict. Web Application Development Frameworks (ADF) is also installed on the same machine.



Figure 4. Multiple-Machine ArcGIS Server Deployment

The multiple-machine scenario includes one or more SOC machines that perform GIS tasks and can be scaled out depending on the number of users your system needs to accommodate. The number of SOC machines you should add depends on how many users will make requests to the system at any one time and the processing intensity of the operations they will request. Under average conditions, a CPU in a SOC machine can support about four concurrently active service instances. Figure 4 shows two SOC machines. If each machine is a dual-CPU system, this configuration can accommodate about 16 users simultaneously performing operations on services. Using that benchmark, adjust the number of SOC machines to accommodate the number of concurrent users you anticipate in your deployment environment.

The data tier of the multiple-machine configuration consists of a separate data server machine running Enterprise ArcSDE and a DBMS. The SOC machines have permissions to access the data on this machine based on the username and password you save when you make the ArcSDE connection in ArcCatalog.

For more information on ArcGIS Server configuration scenarios, go to:

http://webhelp.esri.com/arcgisserver/9.3/DotNet/deploy_scenarios.htm

Testing Process and Results

VMware and ESRI jointly conducted performance tests to characterize the performance of ArcGIS Server on VMware Infrastructure. The following sections summarize the configurations tested, workloads used in testing, and the results.

Hardware Configuration

Host Server:

- Dell 2950 Server running VMware ESXi 3.5u3
- 2 x Quad E5450 Core CPUs and 16 GB RAM

Storage:

- Local attached SAS disks with 2 disk RAID 1; ESXi and VMware VMFS were installed on same disk.
- NFS with 6 SAS disk RAID 5; ESXi was installed on local disk and NFS was used for virtual machine storage.

Physical Server Configuration

- Dell 2950 Server running Windows 2003 x64
- 2 x Quad E5450 Core CPUs and 16 GB of RAM

Note: During the testing of the physical server, Microsoft Windows 2003 boot.ini parameters were passed to the operating system, to limit the number of CPU and memory resources available, to be able to make appropriate comparisons. You can see an overview of these parameters at:

<http://support.microsoft.com/kb/833721>

Virtual Machine Configuration

The following ArcGIS Server configurations were used in testing:

- 4 virtual machines: 1 vCPU / 2GB RAM: Equivalent to native 4 cores with 8 GB RAM each.
- 4 virtual machines: 1 vCPU / 4GB RAM: Equivalent to native 4 cores with 16GB RAM each.
- 4 virtual machines: 2 vCPU / 2 GB RAM: Equivalent to native 8 cores with 8 GB RAM each.
- 4 virtual machines: 2 vCPU / 4 GB RAM: Equivalent to native 8 cores with 16 GB each.
- 8 virtual machines: 1 vCPU / 2GB RAM: Equivalent to native 8 cores with 16 GB each.

Figure 5 shows the configuration of four ArcGIS Server virtual machines running on one VMware ESXi version 3.5u4 host.

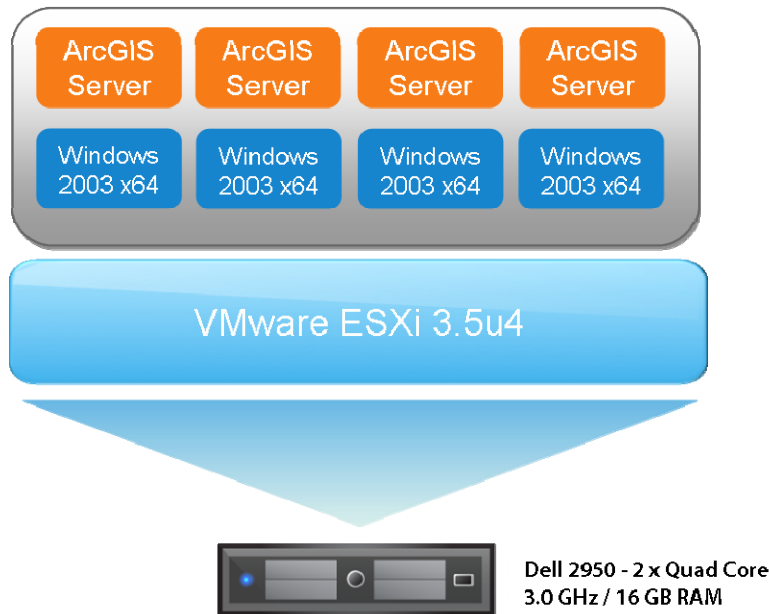


Figure 5. Configuration of Four ArcGIS Server Virtual Machines on Single Box

Workload Used

ESRI used Microsoft’s Visual Studio Team edition to generate a simulated workload to the ArcGIS Server instances. The tests consisted of either 4 or 8 clients simultaneously submitting export map requests to ArcGIS Server 9.3 optimized REST MapServices. Each client continuously submits a request, waits for a response and submits another request with no think time for a period of 60 minutes. Throughput is calculated as the number of transactions per hour that are supported on the machine with a minimally acceptable transaction time of three seconds.

Results Observed

ESRI ArcGIS Server version 9.3 was tested running on both a single physical server and a set of virtual machines using two types of storage architectures, locally attached disks (2 x SAS RAID 1) as well as NFS mounted disks (6 x SAS RAID 5). The tables below provide comparisons of throughput and transaction time performance for the physical versus virtual machine configurations.

Table 1. Physical Servers configured with 2 x SAS RAID-1 direct-attached disks

Configuration	ESRI ArcGIS Server Configuration	Throughput (Transactions/ Hour)	Average Transaction Time
1 Physical: 4 cores, 8GB RAM	All ArcGIS Server Components	41154	0.7
1 Physical: 4 cores, 16GB RAM	All ArcGIS Server Components	42264	0.68
1 Physical: 8 cores, 8GB RAM	All ArcGIS Server Components	56814	0.51
1 Physical: 8 cores, 16GB RAM	All ArcGIS Server Components	56574	0.51

Table 2. Virtual Machines configured with 2 x SAS RAID-1 direct attached disks

Configuration	ESRI ArcGIS Server Configuration	Throughput (Transactions/ Hour)	Average Transaction Time
1 Virtual: 4 vCPU, 16 GB RAM (equivalent to native 4 cores with 16GB RAM) running Windows 2003 x64	All ArcGIS Server Components	27342	1.08
2 Virtual: 2 vCPU, 8 GB RAM (equivalent to native 4 cores with 8GB RAM) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM & Data Server 1 x Data Server (ArcSOC)	32790	0.88
2 Virtual: 2 vCPU, 4 GB RAM (equivalent to native 4 cores with 8GB RAM) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM & Data Server 1 x Data Server (ArcSOC)	33000	0.88
2 Virtual: 2 vCPU, 4 GB RAM (equivalent to native 4 cores with 8GB RAM) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM & Data Server 1 x Data Server (ArcSOC)	33624	0.86

Tables 1 and 2 provided a comparison of physical versus virtual machine performance using the same storage architecture, locally-attached disks.

Table 3 provides a performance comparison for virtual machines using NFS network-attached storage.

Table 3. Virtual Machines located on 6 x SAS RAID-5 network attached storage (NFS)

Configuration	ESRI ArcGIS Server Configuration	Throughput (Transactions/ Hour)	Average Transaction Time
4 Virtual: 1 vCPU, 2GB RAM (equivalent to native 4 cores with 8 GB RAM) running Windows 2003 x64	All ArcGIS Server Components	37476	0.77
4 Virtual: 1 vCPU, 4GB RAM (equivalent to native 4 cores with 16GB RAM) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM), Data Server, and 3 x Data Server (ArcSOC)	35994	0.8
4 Virtual: 2 vCPU, 2 GB RAM (equivalent to native 8 cores with 8 GB RAM) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM), Data Server, and 3 x Data Server (ArcSOC)	49524	0.58
4 Virtual: 2 vCPU, 4 GB RAM (equivalent to native 8 cores with 16 GB) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM), Data Server, and 3 x Data Server (ArcSOC)	50562	0.57
8 Virtual: 1 vCPU, 2GB RAM (equivalent to native 8 cores with 16 GB) running Windows 2003 x64	1 x Controller (IIS, Object Manager ArcSOM), Data Server, and 7 x Data Server (ArcSOC)	52344	0.55

The virtual configuration that provided the highest throughput of 52344 transactions per minute used eight separate Microsoft Windows 2003 x64 virtual machines, each with 1 x vCPUs and 2 GB of RAM.

The second best throughput of 50562 transactions per minute, and a transaction time of 0.58 seconds, was provided by the configuration with four virtual machines, each with 2 vCPUs and 4 GB of RAM.

Deployment Best Practices

VMware best practices are outlined in the VMware documentation available at <http://www.vmware.com/support/pubs/>. VMware also offers specific documents on VMware Infrastructure 3 performance enhancements, networking and storage performance, best practices for VMware VMFS and resource management with VMware Distributed Resource Scheduler (DRS). Additional specific recommendations for VMware infrastructure deployments are:

- Make sure that the host server and storage that you will be using to deploy VMware ESX are listed on the VMware *Systems and Storage Hardware Compatibility List (HCL)* available at:
<http://www.vmware.com/resources/compatibility/>
- For best performance, it is recommended that you run ESRI ArcGIS Server on VMware ESX 3.5, update 3, or later versions.
- Disconnect unused, unnecessary devices on both the guest and on the host. These include COM ports, LPT ports, Floppy drives, CD-ROM drives, and USB adapters. Disabling devices on the host frees IRQ resources and eliminates IRQ sharing conflicts that can cause performance problems.
- Make sure that you are running the latest version of VMware Tools in the guest operating system of virtual machines.
- Do not set resource reservations and limits unless required. Set the limit as “unlimited,” which is the default specified by VMware.
- To establish a network between two virtual machines that reside on the same ESX host, connect both virtual machines to the same virtual switch. If the virtual machines are connected to different virtual switches, traffic will go through “wire” and incur unnecessary CPU and network overhead.
- Using a SAN device for storage is recommended to best meet I/O requirements for applications and to leverage all VMware Infrastructure features and capabilities. Using iSCSI or NFS, respectively, provide the next best storage performance. If ArcGIS Server is deployed onto an iSCSI array or NFS Server, VMware recommends you have at least a 1 Gbps connection.

If local disks must be used, VMware recommends placing the VMware ESX operating system on disks separate from the VMFS file system where the virtual machines reside.

Summary

You can virtualize ESRI ArcGIS Server successfully using VMware Infrastructure. The majority of ArcGIS Server multi-server deployments are good candidates for virtualization and can benefit from advantages offered by a virtualized infrastructure — such as improved management, availability, and scalability — thus reducing TCO.

VMware Infrastructure makes it simpler and less expensive to provide higher levels of availability for ArcGIS Server. Taking advantage of key VMware Infrastructure features as VMotion, DRS, and HA, you can eliminate planned downtime, reduce unplanned downtime, and recover rapidly from component or system outages.

To deploy ArcGIS Server successfully in VMware Infrastructure, you should first clearly understand your organization's business and technical needs, availability requirements, and other operational requirements for implementing ArcGIS Server.

It is critical that you follow best practice guidelines specific to ArcGIS Server as well as those applicable for VMware Infrastructure. In general, best practices in physical environments also apply to deployments on VMware Infrastructure.

Storage configuration is critical to any successful database deployment, especially in virtual environments in which you may consolidate many different SQL Server workloads on a single ESX host. Often, performance issues in virtual environments can be traced to improper storage configuration.

Resources

Customers can find more information about VMware and ESRI products via the links listed below:

- VMware Web site:
<http://www.vmware.com/>
- ESRI Web site:
<http://www.esri.com/>
- VMware Infrastructure 3:
http://www.vmware.com/products/data_center.html
- VMware Infrastructure 3 download page:
<http://www.vmware.com/download/vi/eval.html>
- VMware support Web site:
<http://www.vmware.com/vmtn/>
- VMware Performance Tuning Paper:
http://www.vmware.com/pdf/vi_performance_tuning.pdf
- System Compatibility Guide for a complete list of compatible hardware:
http://www.vmware.com/pdf/vi35_systems_guide.pdf
- Storage/SAN Compatibility Guide for a complete list of compatible storage devices:
http://www.vmware.com/pdf/vi35_san_guide.pdf
- I/O Compatibility Guide for a complete list of compatible networking devices:
http://www.vmware.com/pdf/vi35_io_guide.pdf