


Introducing Microsoft Virtual Server 2005 on IBM **@server** xSeries Servers

Describes the major features of this new product

Explains the architecture and best practices

Introduces the management tools that are available



Kenny Bain
Martin Dippold
David Watts



International Technical Support Organization

**Introducing Microsoft Virtual Server 2005 on
IBM @server xSeries Servers**

November 2004

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (November 2004)

This edition applies to Microsoft Virtual Server 2005 running on IBM @server xSeries servers.

© Copyright International Business Machines Corporation 2004. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team that wrote this Redpaper	ix
Become a published author	xi
Comments welcome	xi
Chapter 1. Product overview	1
1.1 Microsoft Virtual Server 2005 architecture	2
1.1.1 Physical system	3
1.1.2 Host operating system	3
1.1.3 Virtual machine environment	4
1.1.4 Guest operating systems	5
1.2 Virtual Server management	6
1.2.1 Virtual Server Administration Web site	6
1.2.2 Active Directory integration	9
1.2.3 Microsoft Operations Manager 2005 Management Pack for Virtual Server	9
1.2.4 IBM Director 4.2 with Virtual Machine Manager	9
1.2.5 Virtual Server Migration Toolkit	9
1.3 Resource management	10
1.3.1 Processor resource allocation	10
1.3.2 Memory resource allocation	11
1.3.3 Virtual hard disks	11
1.3.4 Virtual networking	12
1.4 Sizing	14
1.4.1 Calculating requirements	15
1.5 Performance	15
1.6 Best practices	16
1.6.1 Ensure that the computer has adequate memory and disk space	16
1.6.2 Enable large memory support where appropriate	16
1.6.3 Disable Hyper-Threading	17
1.6.4 Correctly configure physical network adapters	17
1.6.5 Use NTFS file system on computers running the Administration Web site	18
1.6.6 Keep virtual machine components in a single folder	18
1.6.7 Install Virtual Machine Additions	18
1.6.8 Disable Host time synchronization when operating in multiple domains	18
1.6.9 Run Sysprep on your virtual hard disks	19
1.6.10 Do not use the VMRC client over a Terminal Services connection	19
1.6.11 Shut down virtual machines when making changes	19
1.6.12 Ensure adequate storage space when using Undo Disks	20
1.6.13 Write-protect the parent disk before using a differencing disk	20
1.6.14 Use Relative Weight	20
1.6.15 Ensure proper configuration for high disk-use scenarios	20
1.6.16 The implications of using Virtual Server with Virtual PC 2004	21
1.6.17 Secure Virtual Server and its associated components	21
1.6.18 Stop the Virtual Server service before host operating system shutdown	22
1.6.19 Improve Virtual Server application performance	22
1.6.20 Use appropriate high-availability techniques	22

1.6.21 Use the event log when troubleshooting	23
Chapter 2. Virtual Server architecture	25
2.1 Virtual machine technology	26
2.2 Virtual Server application structure	26
2.2.1 Virtual Machine Monitor kernel	27
2.2.2 VMM driver	27
2.2.3 NDIS driver	28
2.2.4 Virtual Server service	28
2.3 Virtual networking	28
2.4 Virtual hard disks	29
2.4.1 VHD architecture	29
2.4.2 Clustering support	29
2.4.3 VHD types	30
2.5 Virtual Server Additions	31
2.6 Virtual machine security	31
2.6.1 Virtual machine isolation and encapsulation	32
2.6.2 User authentication	32
2.6.3 Administrative Network connections	33
Chapter 3. Customer scenarios	35
3.1 Consolidating test and development environments	36
3.2 Migrating legacy applications	36
3.3 Consolidate multiple server workloads	37
3.3.1 Disaster Recovery	38
3.4 Simulate distributed applications on one physical server	39
Chapter 4. Management and deployment	41
4.1 IBM Director 4.2	42
4.2 Virtual Machine Manager 1.0	42
4.2.1 VMM integration with IBM Director	43
4.2.2 Virtual Machine Manager tasks	44
4.2.3 Scheduler tasks for use with VMM objects	45
4.2.4 Power operations for all virtual machines on a host	48
4.2.5 VMM event filters and actions	49
4.3 Microsoft Operations Manager 2005	53
4.3.1 Management consoles	54
4.3.2 Management packs	54
4.3.3 MOM Connector Framework	55
4.4 MOM Management Pack for Virtual Server	55
4.4.1 Feature overview	56
4.5 Virtual Server Migration Toolkit	60
4.5.1 VSMT infrastructure and system prerequisites	60
4.5.2 Using VSMT for image capture and deployment	63
4.5.3 Configuring virtual machines	69
Chapter 5. Automation	73
5.1 The Virtual Server 2005 COM API	74
5.2 Scripting scenarios	75
5.2.1 Backup	75
5.2.2 Load balancing	75
5.2.3 Disaster recovery	75
5.2.4 Automation of repeating tasks	76

Related publications	77
IBM Redbooks	77
Product publications	77
Referenced URLs	77
Online resources	78
How to get IBM Redbooks	78
Help from IBM	79
 Abbreviations and acronyms	 81

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:


This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®
server®
ibm.com®
xSeries®
BladeCenter™

DB2®
IBM®
Lotus®
NetView®
OS/2®

Redbooks™
Redbooks (logo) ™
ServerProven®
Tivoli Enterprise™
Tivoli®

The following terms are trademarks of other companies:

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Preface

Businesses continually seek ways to reduce cost and risk while increasing quality and agility in their IT infrastructure. Virtualization is a key enabling technology that businesses can use to achieve these business benefits. With virtualization technology, clients run multiple operating systems concurrently on a single physical server, improving hardware efficiency, reducing IT costs, and increasing administrator productivity.

Microsoft Virtual Server 2005 hosted on Windows® Server 2003 and IBM @server® xSeries® servers uses virtualization technology to deliver the performance necessary for completing time-saving and cost-saving tasks. It provides businesses with an enterprise-ready computing environment with advanced levels of scalability, manageability, and reliability.

This IBM® Redpaper introduces Virtual Server and describes its main features and functions, architecture, and typical uses. It also introduces the management tools available for Virtual Server, including IBM Director with Virtual Machine Manager and Microsoft® Operations Manager.

This IBM Redpaper is intended for IT specialists who would like to learn about this new product and how it can be used in their environment.


The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Kenny Bain is an xSeries server specialist with IBM in the Advanced Technical Support (ATS) organization of EMEA and is based in Greenock, Scotland. He holds a degree in Electrical and Electronic Engineering from Strathclyde University. He has worked with IBM for eight years, previously in development testing, but now in pre-sales. He co-wrote the *IBM @server xSeries 440 and 445 Solution Assurance Product Review Guide*. His area of expertise is xSeries hardware, with particular interest in high-end systems and virtualization software solutions. His responsibilities are EMEA-wide and include product compatibility and acting as a focal point for high-end performance-related issues.

Martin Dippold is an IT specialist with IBM Mannheim, Germany. He has ten years of experience in networking with a wide variety of devices and operating systems. He holds a degree in Computer Science and has worked at IBM for 12 years. He co-wrote *IBM WorkSpace On-Demand 3.01*, SG24-6006. His areas of expertise include Intel® server operating systems, software distribution, and server consolidation.

David Watts is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces IBM Redbooks™ on hardware and software topics related to IBM @server xSeries systems and associated client platforms. He has authored over 30 IBM Redbooks and Redpapers. He has a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for more than 15 years. He is an IBM @server Certified Specialist for xSeries and an IBM Certified IT Specialist.



The residency team (left to right): Martin Dippold, Kenny Bain, David Watts

Thanks to the following people for their contributions to this project:

From IBM:

David Archer, Advisory Software Engineer
Paul Branch, xSeries Technical Alliance Manager for Microsoft
Donn Bullock, Senior Product Manager for xSeries
Vinod Menon, Enterprise System Management
Joakim Hansson, Project Manager for Virtual Machine Manager
Massimo Re Ferre', xSeries IT Specialist
Margaret Ticknor, ITSO Raleigh Center IT Support

From Microsoft:

Chris Hallum, Program Manager for MOM Pack for Virtual Server
Jim Katsandres, Presales Engineer for IBM Relationship
Eric Keyser, Marketing Manager for IBM Relationship
James Ni, Senior Technical Product Manager for Virtual Server
Eric Winner, Program Manager for Virtual Server Migration Toolkit
Jeff Woolsey, Development Program Manager for Virtual Server

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- Send your comments in an email to:

redbook@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 662
P.O. Box 12195
Research Triangle Park, NC 27709-2195



Product overview

Microsoft Virtual Server 2005 is a server application for the Windows operating system that enables users to run a broad range of operating systems simultaneously on an industry-standard x86-based single physical server. The operating systems include Windows Server 2003, Windows 2000 Server, and Windows NT® 4.0 Server,

A typical server configuration runs a single operating system on one physical hardware platform. The hardware includes the mouse and keyboard, processor, memory, disk drives and drive controllers, video and network adapters, and other physical devices. The operating system runs on the hardware and controls it, and applications run on the operating system.

By contrast, the virtual machine technologies built into Virtual Server enable one physical server to run several operating systems. Each operating system is installed in an encapsulated environment and runs specific applications. Each virtual machine uses software and hardware devices to create an emulated operating environment.

Microsoft Virtual Server 2005 is a product optimized to provide this capability in a Windows Server 2003 operating system environment. In its first release, Virtual Server targets four key developer and server administrator scenarios:

- ▶ Software testing and development
- ▶ Legacy application re-hosting
- ▶ Testing of distributed server applications on a single server
- ▶ Server consolidation

Virtual Server is available in two separate editions: Microsoft Virtual Server 2005, Enterprise Edition and Microsoft Virtual Server 2005, Standard Edition. The features of both are the same; scalability is the only difference. Standard Edition supports as many as four processors and Enterprise Edition supports as many as 32 physical processors.

This chapter provides a basic overview of Virtual Server's features, management, and best practices.

1.1 Microsoft Virtual Server 2005 architecture

Figure 1-1 illustrates the basic structure of the Virtual Server architecture. For a more detailed discussion, see Chapter 2, “Virtual Server architecture” on page 25.

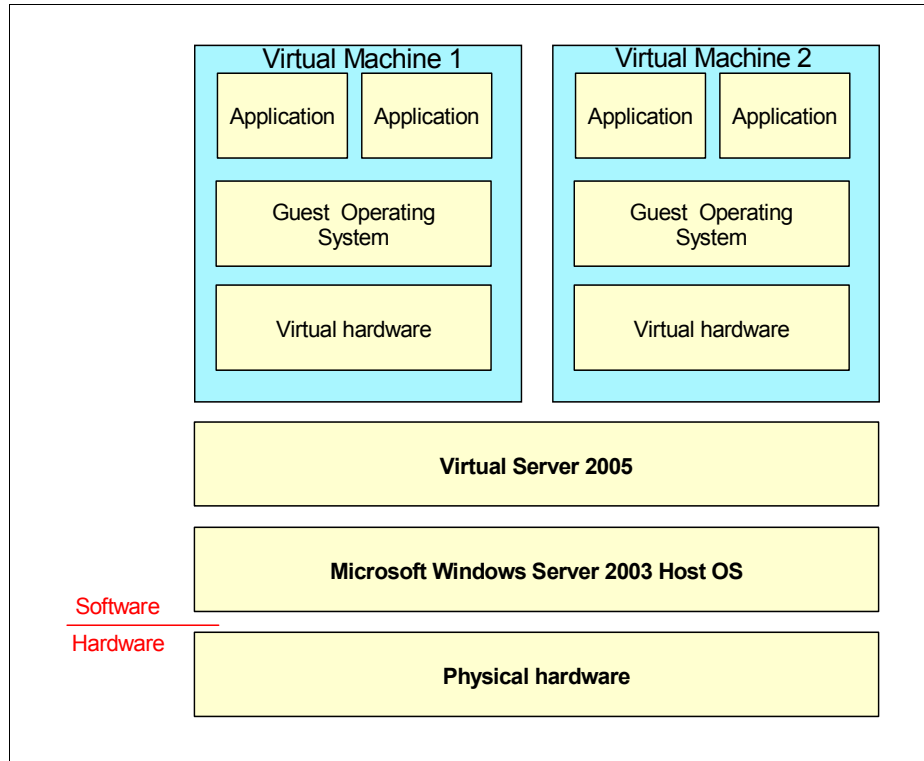


Figure 1-1 Microsoft Virtual Server 2005 architecture

Figure 1-1 shows the four building blocks that are required for the Virtual Server architecture:

- ▶ **Physical system**
This is where necessary hardware is installed.
- ▶ **Host operating system**
This provides an interface to the underlying hardware.
- ▶ **Virtual Server application with a Virtual Machine Monitor (VMM) virtualization layer**
VMM manages the virtual machines. This provides the software infrastructure for hardware emulation and the necessary links to the underlying host operating system. Each virtual machine consists of a set of virtualized devices that are presented to the guest operating system as virtual hardware.

The virtual machine can be described by two files stored on the host operating system. As such, it is fully portable between host systems running Microsoft Virtual Server 2005.
- ▶ **Guest operating systems running in the virtual machines**
These use the virtual devices that are manually assigned by the virtual machine creator. When a guest operating system is running, the special-purpose VMM kernel takes mediated control over the processor and hardware during virtual machine operations. This creates an isolated and encapsulated environment, whereby the guest operating system and applications run close to the hardware at the highest possible performance level.

These four distinct parts are described in the sections that follow in more detail.

1.1.1 Physical system

The two versions of Microsoft Virtual Server 2005 provide two different scalability capabilities:

- ▶ Enterprise Edition supports up to 32 processors and 64 GB of memory
- ▶ Standard Edition supports up to four processors and 64 GB of memory

The physical system can be any of the supported xSeries or IBM @server BladeCenter™ systems. The supported systems can be found at the following Web site:

<http://www.pc.ibm.com/us/compat/nos/microsoft.html>

All systems must have the following devices:

- ▶ Processors
These can be any 32-bit Intel or AMD processor running at 550 MHz or faster (1 GHz recommended). This includes the following specific processors:
 - Intel Xeon or Pentium® families
 - Intel Xeon with EM64T extensions (with 32-bit operating systems only)
 - AMD AMD64 (with 32-bit operating systems only) or Athlon families
- ▶ CD-ROM or DVD drive
- ▶ Super VGA (800 × 600) or higher resolution monitor recommended
- ▶ Storage and network adapters should also be included for most guest operating systems

Note: Virtual Server does not support Itanium or Itanium 2 processors.

1.1.2 Host operating system

You can run Microsoft Virtual Server 2005 on a server with multiple processors; in fact, this is a normal configuration. The maximum number of processors for Microsoft Virtual Server 2005, Enterprise Edition is determined by the host operating system. The maximum number is 32.

Important: The host operating system must be the 32-bit version of the operating systems listed in Table 1-1. Virtual Server does not support 64-bit operating systems.

Table 1-1 Maximum processor and memory configurations for each supported host operating system

Operating System	Edition	Maximum CPU	Maximum memory
Microsoft Windows Server 2003	Datacenter	32	64 GB
	Enterprise	8	32 GB
	Standard	4	4 GB
Microsoft Windows Small Business Server 2003	Premium	2	4 GB
	Standard	2	4 GB
Microsoft Windows XP	Professional (see note)	2	4 GB

Note: Microsoft does not support Windows XP Professional as a host operating system in a production environment.

1.1.3 Virtual machine environment

Virtual Server creates a virtualization layer called the Virtual Machine Monitor (VMM). VMM provides the software infrastructure for creating virtual machines, managing instances, and interacting with guest operating systems. Virtual Server is a multi-threaded application that runs as a system service, with each virtual machine running in its own thread of execution. Virtual machine input/output (I/O) occurs in child threads.

Virtual Server derives two core functions from the host operating system:

- ▶ The underlying host kernel schedules CPU resources.
- ▶ Host device drivers enable access to system devices.

Virtual machine isolation ensures that if one virtual machine crashes or hangs, it cannot impact any other virtual machine or the host system. Maximum application compatibility is achieved through isolation. Clients, therefore, can leverage existing storage, network, and security infrastructures.

Virtual Server can support up to a maximum of 64 virtual machines. The practical limit of how many virtual machines you can run simultaneously depends on system resources, the amount of memory assigned to each virtual machine, and the total memory available on the physical server.

Each virtual machine provides a set of virtualized devices that are shown in Table 1-2. These devices are the virtual hardware for use within the guest operating system.

Table 1-2 Virtualized hardware presented to the guest operating system

Device	Description
Processor	In the current version of Virtual Server, the guest operating systems see a single processor per virtual machine (with its ID based on that of the real processor).
Memory	Virtual Server supports up to 3.6 GB per virtual machine. Note: Host memory cannot be overcommitted to virtual machines. See 1.3.2, “Memory resource allocation” on page 11.
Small Computer System Interface (SCSI) virtual storage	Virtual Server emulates a multi-port Adaptec 7870 SCSI controller with 4 virtual SCSI buses. With virtual SCSI storage enabled, each virtual machine can connect to more than 56.5 TB. Virtual Server supports two-node failover clustering between virtual machines by means of shared SCSI storage.
Graphics	Virtual Server emulates the S3 Trio64 graphics adapter with 4 MB of VRAM for VESA 2.0-compliant VGA and SVGA resolutions, 2-D acceleration, hardware cursor, and Microsoft DirectX support.
integrated drive electronics (IDE) and ATA Packet Interface (ATAPI) virtual storage	Virtual Server emulates up to 4 IDE hard disk drives or CD/DVD-ROM drives (or ISO images). Virtual hard disk drives can provide up to 128 GB of storage per IDE channel.
Virtual diskette drive	One 1.44-MB virtual diskette drives can be mapped to a host drive or image.
Virtual Serial (COM) ports	Two virtual serial ports can be mapped to host serial ports
Virtual Parallel (LPT) ports	One virtual printer port can be mapped to the host parallel port

Device	Description
Virtual Ethernet controller	Virtual Server emulates a multi-port Intel 21141 10/100TX Ethernet controller with 4 virtual network adapters: <ul style="list-style-type: none"> ▶ A local virtual network connection does not need to be associated with a device ▶ Virtual machines can also be configured to have no virtual network connection
Virtual Machine BIOS	AMI BIOS
Sound	Virtual Server does not currently include an emulated sound card
Virtual motherboard	Virtual Server emulates an Intel 440BX chip set with PIIX4 and the following components: <ul style="list-style-type: none"> ▶ CMOS ▶ Real-time clock ▶ RAM and VRAM ▶ Memory Controller ▶ DMA Controller ▶ PCI Bus ▶ ISA Bus ▶ SM Bus ▶ Power Management ▶ 8259 PIC ▶ PIT

1.1.4 Guest operating systems

Virtual Server supports the following operating systems as guests:

- ▶ Windows Server 2003, Standard Edition
- ▶ Windows Server 2003, Enterprise Edition
- ▶ Windows Server 2003, Web Edition
- ▶ Windows Small Business Server 2003
- ▶ Windows 2000 Server
- ▶ Windows 2000 Advanced Server
- ▶ Windows NT Server 4.0 with Service Pack 6a

Important: Currently Virtual Server only provides support for up to 1 processor and up to 3.6GB of memory in each virtual machine.

Although Virtual Server can run most variants of Linux®, as well as NetWare, OS/2®, and BSD, as guest operating systems, Microsoft does not support them. Only the above Windows operating systems are supported as guests.

At the time of publication, IBM had no plans to offer support for these operating systems.

Virtual Machine Additions

For complete integration and optimization of the guest operating systems, Virtual Server provides Virtual Machine Additions, a component of Virtual Server with these features:

- ▶ Mouse pointer integration
- ▶ Video driver optimization
- ▶ Host time synchronization

Virtual Machine Additions is an architectural component for integrating host and guest operating system for optimized performance and enhanced user experience. Virtual Machine Additions must be manually installed by the user after installation of the guest operating system.

Note: Full integration and optimization features are available only for Windows guests.

Physical machine equivalency

Virtual Server hides the fact that guest operating systems are not running on a physical server where appropriate. This is called physical machine equivalency.

Physical machine equivalency provides interoperability with the following system management products:

- ▶ Automated Deployment Services
- ▶ Windows Update Services
- ▶ System Management Services
- ▶ Microsoft Operations Manager

Physical machine equivalency has additional attributes:

- ▶ Event logging has been integrated with the host Windows event log with integrated management solutions.
- ▶ Performance monitoring is integrated with the host Perfmon with CPU, RAM, and heartbeat counters, which integrate with host management solutions.

1.2 Virtual Server management

Microsoft Virtual Server 2005 virtualized management can be accomplished in many ways with many standard tools. This section describes the tools and services available from both Microsoft and IBM.

1.2.1 Virtual Server Administration Web site

The Virtual Server service is installed as a headless service (that is, a service with no direct display, keyboard, or mouse access) on the host operating system. The Common Object Model (COM) application programming interface (API) is used to create the content in the Administration Web site.

All actions performed through the Web site can be automated through the API. As such, the creation, configuration, and control of virtual machines becomes automatic. The Administration Web site controls a single instance of Virtual Server and enables authenticated administration and remote access.

The Virtual Server Administration Web site is accessible using the URL:

`http://systemname:1024/VirtualServer/VSWebApp.exe?view=1`

In this URL, systemname is the address of the server that has Microsoft Internet Information Server (IIS) installed on it (that is, if IIS is installed on a separate server).

The main page of the Virtual Server Administration Web site is illustrated in Figure 1-2.

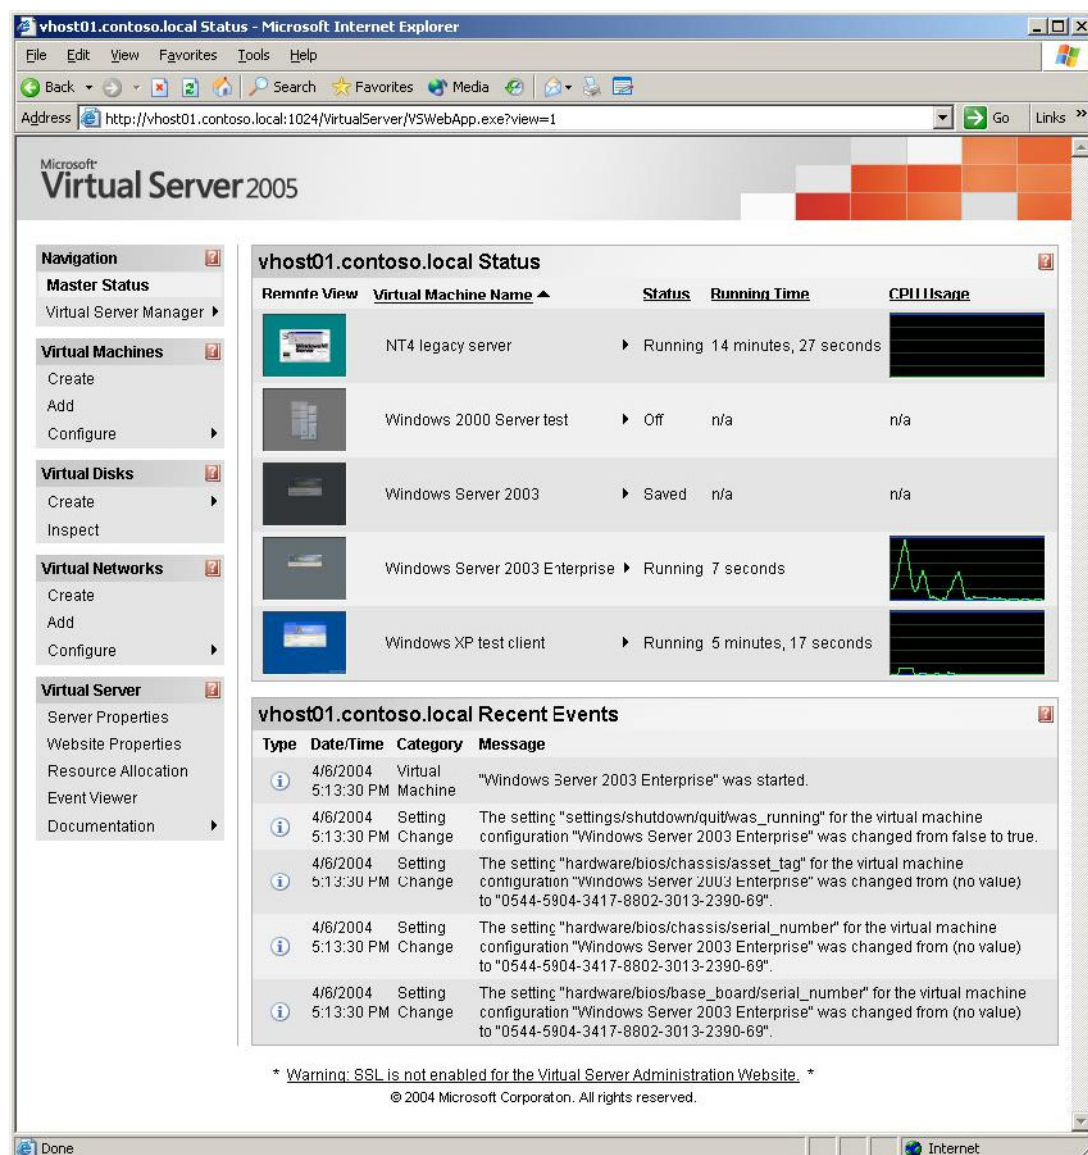


Figure 1-2 The Virtual Server Administration Web site

The Administration Web site is the user interface for administrating and controlling virtual machine configuration securely anywhere in the network. It allows remote control of the virtual machine user interfaces through Virtual Machine Remote Control (VMRC), and a VMRC link provides interface access for each virtual machine.

VMRC uses Secure Sockets Layer (SSL) data encryption and Active Directory and Kerberos user authentication. Remote access can also be accomplished with a Remote Desktop Protocol (RDP) or Terminal Services, although this may require the installation of an appropriate application and the configuration of the relevant permissions.

SSL configuration is provided through the Administration Web site. You can use the Administration Web site in conjunction with standard server management tools to administer virtual machines.

Virtual Server provides three levels of administration access to virtual machines based on host access control lists (ACLs):

- ▶ Creation of new virtual machines
- ▶ Reconfiguration of a specific virtual machine
- ▶ User access to a specific virtual machine

Administrators of the host operating system having full access to all three levels by default.

The Administration Web site makes it possible to control the state of each virtual machine with the controls shown in Table 1-3.

Table 1-3 Virtual Server Administration Website controls

Control	Description
Remote Control	Presents the user interface for the virtual machine selected
Turn On	Starts the virtual machine
Pause	Pauses the virtual machine without discarding memory
Resume	Resumes a paused virtual machine
Save State	Saves the current state of the virtual machine, and stops the virtual machine from running
Restore Saved State	Turns on the virtual machine and restores it to the state it was in when its state was saved
Discard Saved State	Discards the Saved State file for the virtual machine and leaves the virtual machine turned off
Shut Down Guest Operating System	Requires the additions to be installed
Turn Off	Turns off the virtual machine without saving any state information. This action has the same effect on the virtual machine as does pulling the plug on a physical computer
Reset	Resets the virtual machine. This action has the same effect on the virtual machine as does pressing a server's reset button
Remove	Removes the virtual machine from the Virtual Server Administration Web site without altering the virtual machine files

Through the Administration Web site, users can also control the configuration of each virtual machine as follows:

- ▶ Set the default action of the virtual machine on host startup and shutdown.
- ▶ Set the access controls for the virtual machine.
- ▶ Change the hardware configuration, including memory, network, storage, floppy drive, CD drive, and port configurations.
- ▶ Change the virtual machine processor resource management settings.
- ▶ Install additions in the virtual machine.
- ▶ Attach a script to a virtual machine event.

Important: The World Wide Web Service component of IIS is a requirement for the computer running the Virtual Server Administration Web site. This does not have to be the system on which Virtual Server is installed. For more details, see the online User's Guide (installed when Virtual Server is installed).

1.2.2 Active Directory integration

Integration with Active Directory enables delegated administration and authenticated guest access. Virtual Server enables fine-grained administrative control over virtual machines with per-virtual machine ACLs that can be managed from within the Active Directory Group Policy Management Console. Event logs are integrated with Active Directory and Microsoft Management Consoles.

1.2.3 Microsoft Operations Manager 2005 Management Pack for Virtual Server

This management pack developed specifically for Virtual Server enables advanced management features within virtual machines. For more details, see 4.3, "Microsoft Operations Manager 2005" on page 53.

Note: At the time of publication, the management pack was not available. Once it is available, you can download it from the Management Pack and Product Connector Catalog Web site:

<http://www.microsoft.com/management/mma/catalog.aspx>

1.2.4 IBM Director 4.2 with Virtual Machine Manager

Virtual Machine Manager is an extension available for IBM Director 4.2 that provides management capability for virtualized environments, including Microsoft Virtual Server 2005. For more details see 4.2, "Virtual Machine Manager 1.0" on page 42.

1.2.5 Virtual Server Migration Toolkit

Virtual Server Migration Toolkit (VSMT) is a free, downloadable tool for Microsoft Virtual Server 2005 that automates the migration of a supported operating system and installed applications from a physical server to a server running in a virtual machine provided and managed by Virtual Server. The physical server can be running any of the following operating systems:

- ▶ Windows NT 4.0 Standard or Enterprise editions with SP6A
- ▶ Windows 2000 Server SP4 or later
- ▶ Windows 2000 Advanced Server SP4 or later
- ▶ Windows Server 2003 Standard or Enterprise editions

To use VSMT, Automated Deployment Services (ADS) 1.0 and Virtual Server must be deployed in the environment where the migration is to be performed.

Note: You can download the Virtual Server Migration Toolkit from the following Web site:

<http://www.microsoft.com/windowsserversystem/virtualserver/evaluation/vsmt.mspx>

For more information, see 4.5, "Virtual Server Migration Toolkit" on page 60.

Note: VSMT supports both editions of Microsoft Virtual Server 2005, but does not support Virtual PC.

1.3 Resource management

Microsoft Virtual Server 2005 provides policy-based control for balanced workload management. Administrators can also use the Administration Web site or scripting tools to adjust some resource allocations dynamically.

1.3.1 Processor resource allocation

Virtual Server supports both weighting and constraint methods for fine-grained control and weight-based and constraint-based resource allocation for balanced workload management as follows:

- ▶ Multi-threaded capability provides for highly scalable performance on systems with up to 32 processors and 64 GB RAM.
- ▶ Each virtual machine can use up to 100% of a single host processor; however, there is no facility to ensure that each virtual machine runs on its own dedicated physical processor.
- ▶ We do not recommend Hyper-Threading. For more information, see 1.6.3, “Disable Hyper-Threading” on page 17. If you do enable Hyper-Threading, the single host processor is a logical processor.
- ▶ Multiple virtual machines can execute concurrently to make use of multiple host processors. The Virtual Server process creates a thread of execution for each virtual machine. Other threads are created to perform I/O, service the COM API, and so forth.
- ▶ The number of virtual machines that can be hosted on any server depends on the combined processor, memory, and I/O load that the virtual machines put on the host. This number also depends on the processor, memory, and I/O capacity available on the host system.
- ▶ Host processor resources can be changed dynamically using the Administration Web site or the COM API.
- ▶ Processor capacity is guaranteed and cannot be overcommitted. Virtual Server provides the allocated amount of processor resources to the virtual machine.
- ▶ Maximum capacity is the highest percentage of the total resources of a single processor that can be consumed by a virtual machine at any given time.
- ▶ Relative weight is a number between 1 and 10,000 that identifies the relative priority of each virtual machine. The relative weight is allotted to the resource needs of a virtual machine compared to those needed by all other virtual machines.

A virtual machine with a higher relative weight is dynamically allocated resources as needed from other virtual machines that have lower relative weights. By default, all virtual machines have a relative weight of 100, so that their processing resource requirements are equal and none is given preference.

- ▶ Using capacity settings lets you assign minimum and maximum percentages of the resources that can be consumed by the virtual machine. You can use the reserved setting to ensure that a minimum amount of resources is always available to the virtual machine. You can use the maximum setting to ensure that the virtual machine does not consume more than the specified percentage of the processor's resources.

1.3.2 Memory resource allocation

Virtual Server supports memory resizing at virtual machine start time. Virtual Server enables flexible memory configuration on a per-virtual machine basis. Since the virtual RAM of a virtual machine cannot exceed the physical RAM of the host machine, the over-commitment of memory (page sharing) is not supported.

Also, memory size cannot currently be reconfigured while a virtual machine is running. In addition, support is included for non-uniform memory access (NUMA) aware scheduling and memory allocation. On non-NUMA systems, Virtual Server relies on the host operating system scheduler.

1.3.3 Virtual hard disks

Virtual Server encapsulates virtual machines in portable virtual hard disks (VHDs) for flexible configuration, versioning, and deployment. VHDs are used to assign storage to the virtual machine, and they contain the data for a virtual machine. The data are packaged as a single file on the host file system. To the guest operating system, the VHD appears as a single storage volume.

Each VHD is created with the Virtual Server Virtual Disk Manager (VDM). When you create a virtual machine, you specify the size and type of the drive. The VHD is then created automatically.

With Virtual Server, you can:

- ▶ Create a new VHD
- ▶ Use an existing VHD
- ▶ Connect no VHD to the virtual machine.

The Virtual Disk Manager at the Administration Web site can create a VHD in any storage to which the host file system has access. This includes IDE, SCSI, redundant array of independent disks (RAID), storage area network (SAN), and network-attached storage (NAS).

When any given virtual machine accesses a virtualized IDE/SCSI device, the Virtual Machine Monitor (VMM) maps that to the physical device access. Virtual machines can access VHDs on any storage topology that is supported by the underlying Windows host server. This includes SAN and NAS configurations.

VHDs support the following formats and functionality:

- ▶ Dynamically expanding virtual hard disks
This VHD format starts with an NT File System (NTFS) Sparse File and grows with the addition of data. Virtual Server sends low disk warnings if a disk image on the host file system grows too large.
- ▶ Fixed-size virtual hard disks
This VHD format consists of a fixed-extent file that resides on a host system's hard drive. Fixed-size drives cannot be resized.
- ▶ Linked virtual hard disks
A linked virtual hard disk is a VHD that points to and uses an entire physical disk for the purpose of converting physical disk to a virtual hard disk.
- ▶ Differencing virtual hard disks

A differencing virtual hard disk is a VHD associated with another VHD in a parent-child relationship. It enables disk hierarchy, with one or more child disks and a parent disk. The differencing disk is the child, and the associated virtual hard disk is the parent. The differencing disk stores a record of all changes made to the parent disk, which saves changes without altering the parent disk.

We recommend that you write-protect the parent disk before using the differencing disk. Otherwise, if the parent disk is modified by another process, all differencing disks related to it become invalid. On the host file system, differencing disks appear as two files, but from within the virtual machine, they appear as a single VHD.

- ▶ Undo disks

Undo disks is a feature that saves changes to a virtual machine's data in a separate undo disk file in case you want to reverse the changes. The feature provides you with a way to decide whether to permanently modify a virtual machine and its disks each time you end a virtual machine session. When you enable undo disks, it applies to all VHDs installed on the virtual machine. When you run a virtual machine that is using undo disks, any changes to a virtual hard disk are temporarily stored in an undo disk (.vud) file, while reads occur from the original virtual hard disk.

Note: The Virtual Disk Manager supports conversion between drive types.

Once created, virtual machines in Virtual Server consist of two components:

- ▶ An XML configuration file that contains metadata describing the virtual machine
- ▶ A VHD file containing the virtual machine data

Virtual Server allows for dynamic disks, fixed-size disks, differencing disks, and undo disks. This flexibility enables fast and economical deployment and leverages existing storage, networking, security, and management infrastructure. Each virtual machine can connect up to 32 VHDs using virtualized IDE or SCSI controllers as follows:

- ▶ Up to four VHDs can be connected to the virtual IDE controller. Each VHD can be up to 128 GB.
- ▶ Up to four virtual SCSI controllers, where:
 - The guest operating system sees a virtualized multi-function Adaptec 2940 PCI adapter.
 - Each controller creates a single SCSI bus.
 - Up to seven VHDs can be attached to each SCSI bus.
 - Each VHD can be up to 2 terabytes.

Note: The maximum storage per virtual machine is over 56 TB.

1.3.4 Virtual networking

Virtual Server provides secure and flexible networking with the following connectivity options:

- ▶ Guest-to-guest
- ▶ Guest-to-host
- ▶ Guest-to-network

Virtual networks are used to connect a virtual machine to internal or external networks. Virtual Server supports up to four virtual network adapters per virtual machine.

By default, one internal virtual network is created to enable virtual machine-to-virtual machine networking. Another virtual network is created for each host Ethernet interface, and this virtual network connects the virtual machines to external networks.

Additional virtual networks can be created with the network settings Web page on the Virtual Server Administration Web site. A unique media access control (MAC) address is created for each virtual network connection. The virtual networking architecture is illustrated in Figure 1-3.

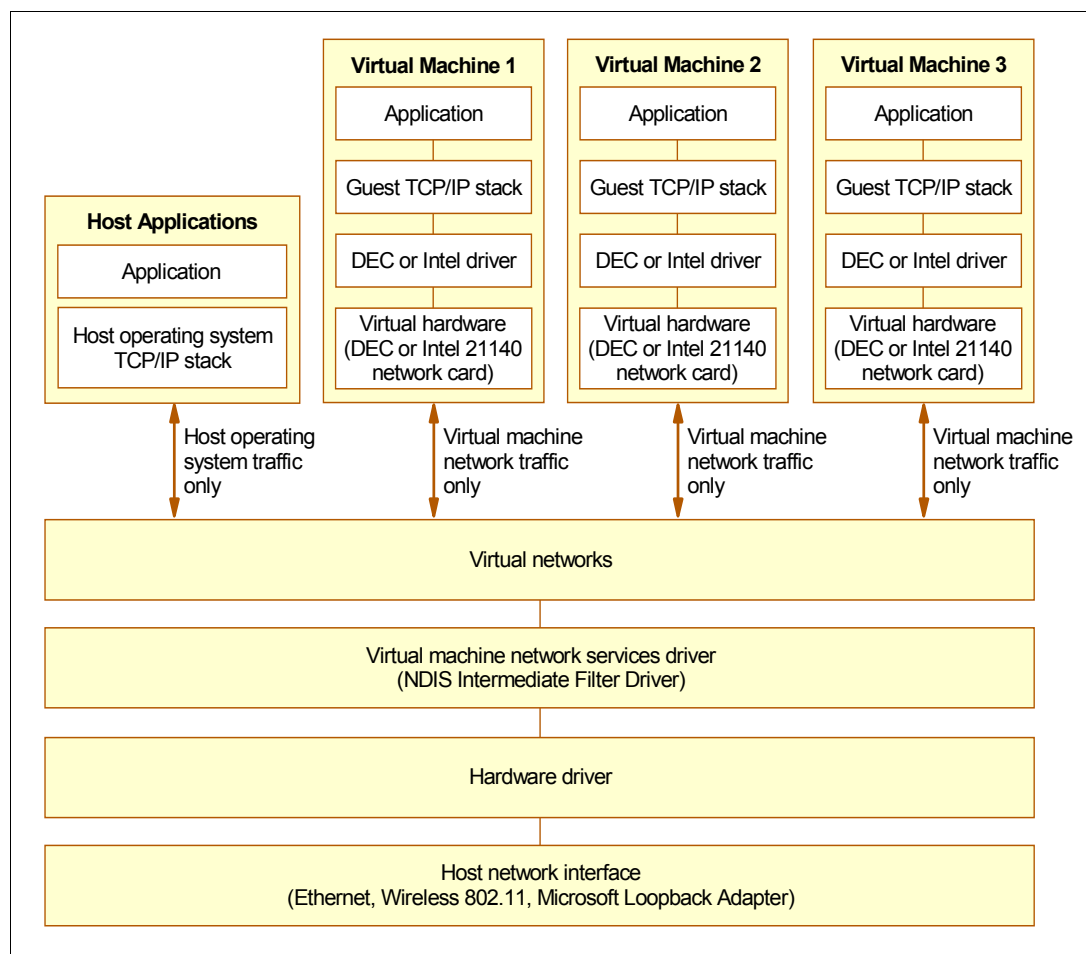


Figure 1-3 Virtual networking architecture

For each virtual network without a connected network adapter, all networking to Virtual Server is performed locally using memory copies. A virtual Dynamic Host Configuration Protocol (DHCP) server is available to provide IP addresses to virtual machines. (Gateways, DNS/WINS servers, and IP leases are configurable.)

Virtual Server supports any Ethernet interface on a host system, including teamed network adapters, enabling powerful, security-enhanced configurations for optimized network operations. With the Virtual Networking Manager, users can create an unlimited number of virtual networks with the following topologies:

- External networking

Virtual machines interoperate with physical machine equivalency as stand-alone servers in a network and can communicate with other nodes in the network, including other virtual machines configured for external networking.

- ▶ Virtual networking

Virtual machines communicate by means of simulated Ethernet and a virtual DHCP server with other virtual machines also running on the same host computer. Virtual networking does not transmit packets over the physical network; instead, it uses host system memory copies.

- ▶ Virtual machine-to-host networking

Virtual machines communicate with the host system by means of a loopback adapter.

1.4 Sizing

You should refer to the product documentation included with each operating system to determine the minimum system requirements for the guest operating system. Be aware that the minimum required disk space you need is the sum of the required disk space of each guest operating system and the host operating system. This sum is in addition to disk space for any applications to be installed on the operating systems.

If you plan to run more than one guest operating system simultaneously, you need to add the memory requirements for those operating systems. Again, these requirements are in addition to the memory required by the host operating system. Also, Virtual Server requires an additional 32 MB of memory per virtual machine.

To enhance performance, consider increasing the amount of memory beyond the minimum requirements described here and in the product documentation. Note that memory allocated to virtual machines is not available for paging by the host operating system.

Table 1-4 provides the minimum disk space and memory required by the host operating system. These requirements are general guidelines only. You should consult the product documentation provided with each operating system for specific requirements.

Table 1-4 Minimum memory and disk requirements for the host operating system

Host operating system	Minimum RAM	Minimum disk space
Windows Small Business Server 2003, Standard Edition	256 MB	4 GB
Windows Small Business Server 2003, Premium Edition	512 MB	4 GB
Windows Server 2003, Standard Edition	256 MB	2 GB
Windows Server 2003, Enterprise Edition	256 MB	2 GB
Windows Server 2003, Datacenter Edition	512 MB	2 GB

It is essential that you allow for the additional space that will be required by all guest operating systems that you plan to deploy. Not only do you need enough disk space for the host operating system and each guest operating system, but you must also take into account the extra space needed for the following:

- ▶ Each virtual machine's paging file
- ▶ All dynamically expanding virtual hard disks
- ▶ The contents of each virtual machine's RAM

1.4.1 Calculating requirements

This example illustrates how to calculate memory and storage required by the host operating system and guest operating systems you want to run. Assume that you want to run the following operating systems simultaneously on three virtual machines:

- ▶ Host operating system for the physical computer: Windows Server 2003, Standard Edition
- ▶ Guest operating system for virtual machine 1: Windows NT Server 4.0 with SP6a
- ▶ Guest operating system for virtual machine 2: Windows NT Server 4.0 with SP6a
- ▶ Guest operating system for virtual machine 3: Windows 2000 Server

To calculate disk space and memory overhead required by all the operating systems, add them as shown in Table 1-5. Our recommendations are in bold; the documentation recommendations are presented as normal text for comparison.

Table 1-5 How to calculate disk space and memory overhead for multiple virtual machines

Operating system	Recommended disk space	Memory (including overhead for virtual machines)
Host operating system for the physical computer	8 GB 2 GB	2000 MB 256 MB
Windows NT Server 4.0 with SP6a in virtual machine 1	4 GB 500 MB	512 MB + 32 MB = 608 MB 96 MB
Windows NT Server 4.0 with SP6a in virtual machine 2	2 GB 500 MB	256 MB + 32 MB = 288 MB 96 MB
Windows 2000 Server in virtual machine 3	8 GB 1 GB	2 GB + 32 MB = 2032 MB 160 MB
Minimum total required	22 GB 4 GB	4928 MB 608 MB

Note: As with a physical computer, adding memory to a virtual machine improves performance.

1.5 Performance

Performance of Microsoft Virtual Server 2005 can be improved by ensuring the following general recommendations are adhered to:

- ▶ As with a physical computer, the overall objective in improving Virtual Server performance is to eliminate any performance bottlenecks. In general, the more memory, processors, and hard disk space that exist, the better Virtual Server's performance is.
- ▶ We recommend that you overestimate the amount of RAM required to run the host operating system, Virtual Server, and the guest operating systems. Performance degrades when a virtual machine runs out of available RAM and starts paging.

In addition, poor host operating system performance translates directly into poor virtual machine performance; therefore, you should make sure you have sufficient RAM to run all host processes. We also recommend that you use fast hard disks and SANs to enhance virtual hard disk performance.

- ▶ Add network adapters to improve network performance. You should add up to one adapter per virtual machine.

- Consider load-balancing your applications. Deploy applications that have similar resource requirements on different instances of Virtual Server. Mix network-intensive applications with those that are processor intensive and hard disk intensive. This prevents excessive demands on a single system resource, such as processors.

If, after taking steps to correct the situation, you find that Virtual Server performance still does not meet your requirements, you can easily migrate some of the virtual machines to a different instance of Virtual Server to balance the load.

1.6 Best practices

The configuration tips described in this section help ensure that your installation is successful.

1.6.1 Ensure that the computer has adequate memory and disk space

The host operating system and all virtual machines that are running concurrently require adequate memory. To calculate the total memory needed, you should allocate enough memory:

- To each virtual machine to run its operating system
- To the memory required by the host operating system
- For the associated overhead per virtual machine

You must make sure that you have adequate disk space for each guest operating system that you plan to deploy along with enough disk space for the host operating system. You must take into account the extra space needed for each virtual machine's paging file and all dynamically expanding virtual hard disks. Finally, you must also consider the space needed to save the contents of each virtual machine's RAM when putting the virtual machine into a saved state.

For information about calculating memory and disk space requirements for the host and guest operating systems, see 1.4, "Sizing" on page 14.

1.6.2 Enable large memory support where appropriate

You should enable Physical Address Extension (PAE) on your host operating system, if it is appropriate to do so. PAE allows the host operating system to access more than 4 GB of physical memory. To enable PAE add the /pae switch in the boot.ini file as shown:

```
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Windows Server 2003, Enterprise"  
/fastdetect /pae
```

Since virtual machines can only access up to 3.6 GB of memory, you should not enable PAE on guest operating systems.

If the host operating system is running on a server with more than 16 GB of memory, you should not enable PAE in conjunction with the /3GB parameter on physical systems.

Note: For more information about PAE, see:

http://www.microsoft.com/resources/documentation/windowsserv/2003/enterprise/proddocs/en-us/vlm_1.asp

1.6.3 Disable Hyper-Threading

If workload is heavy, Hyper-Threading, which allows a single processor to appear as two processors, may cause poor server performance. Because workload is dynamic and can vary quickly, Microsoft recommends disabling Hyper-Threading on the physical server to prevent this potential problem.

1.6.4 Correctly configure physical network adapters

Virtual machines that require a high level of network activity can burden the physical network adapters on the physical computer. This can make it difficult to connect to the host operating system or administer the host operating system with the network adapters.

To alleviate this problem, we recommend configuring the physical server with two or more network adapters. One of the network adapters should be dedicated for use by only the host operating system. To do this, unbind the Virtual Machine Network Services from the dedicated network adapter as shown in Figure 1-4.

You should unbind all network protocols, such as Transmission Control Protocol/Internet Protocol (TCP/IP) and Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX), except for the Virtual Machine Network service from all the network adapters that you will use for virtual machines. This ensures that these network adapters are dedicated to Virtual Server networking.

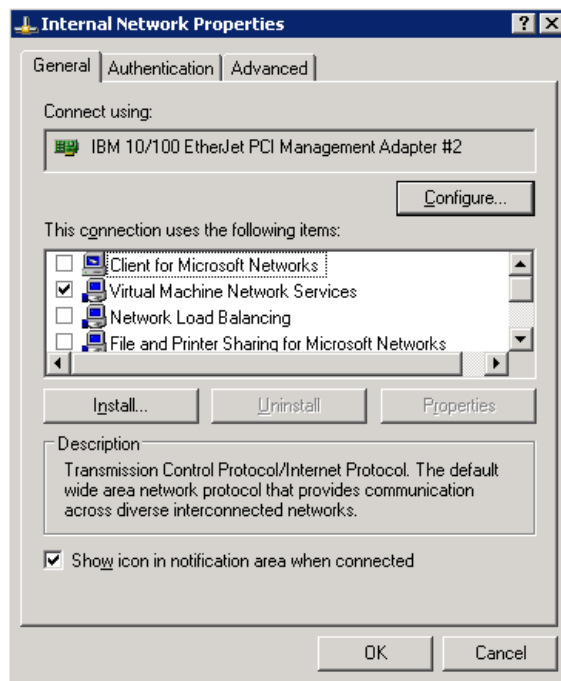


Figure 1-4 Network settings window, showing virtual machine network setting selected

Therefore, to ensure the best possible network configuration, you should:

- ▶ Distribute the networking load. If you are running multiple instances of Virtual Server, you can distribute the networking load between them as though they were physical servers. To do this, run a mix of network-intensive and non-network-intensive applications on a single instance of Virtual Server.
- ▶ Add physical network adapters. For best performance, dedicate at least one physical network adapter to each virtual machine.

- ▶ Virtual machines cannot take advantage of Network Load Balancing (NLB) because the Virtual Server network driver runs below the NLB driver on the host operating system network stack. This ensures that guest operating systems and the host operating system are isolated from one another and cannot read, monitor, and capture one another's network traffic. However, guest operating systems that share the same virtual network can read, monitor, and capture one another's network traffic.
- ▶ The location in the network of the Virtual Server service and the Web service (when they are running on separate servers) with respect to one another is not a major performance consideration.

1.6.5 Use NTFS file system on computers running the Administration Web site

Using the NTFS file system access permissions on the server running IIS (and thereby hosting the Administrator Web site) restricts access to the Web site to only those user accounts with Execute permissions for the folder that contains Web application.

Furthermore, the security of the virtual machine configuration (.vmc) files, the virtual hard disk (.vhd) files, and the virtual network configuration (.vnc) files all rely on the NTFS file system.

1.6.6 Keep virtual machine components in a single folder

Virtual machines consist of two primary components:

- ▶ Virtual machine configuration file
- ▶ VHD file

These can be named whatever you deem appropriate for your installation. We recommend you give them names that are easily identified. The default location for these files is C:\Documents and Settings\All Users\Documents\Shared Virtual Machines.

For ease of management, we recommend that you keep each of these files in a single folder that represents the virtual machine. Other files associated with this virtual machine, such as undo disk and saved-state files, are automatically placed in this folder because these files are created in the same location as the virtual machine configuration (.vmc) file.

By default, all virtual machine configuration files and the related files (for example, VHD files), are created in the Shared Virtual Machines folder. As this might not be appropriate for your installation, you can create these files in a folder structure that more closely resembles your organization. This might allow you to configure security for these files more appropriately.

1.6.7 Install Virtual Machine Additions

Virtual Machine Additions is a set of features that improve the integration of the host and guest operating systems. It also improves the performance and manageability of the guest operating system.

1.6.8 Disable Host time synchronization when operating in multiple domains

The Host time synchronization feature of Virtual Machine Additions allows guest operating systems to synchronize their system clocks with the system clock of the host operating system. This can lead to unpredictable and unwanted results if the host operating system and guest operating system are in different domains or are operating in different time zones.

Therefore, if your guest operating system and host operating system are members of different domains or are operating in different time zones, you should disable Host time synchronization.

Attention: Host time synchronization should not be enabled on a virtual machine that is configured as a domain controller. Instead, you should use an external source for time synchronization. Virtual machines use the processor frequency to measure time; however, because they do not get complete access to the processor, they lose timing clicks from the processor and end up losing time.

1.6.9 Run Sysprep on your virtual hard disks

A significant feature of Virtual Server is the portability of the VHD (.vhd) files. However, to use a.vhd file that contains an operating system installed on multiple virtual machines (for example, a template system held in a system repository for rollout purposes), you must first run Sysprep on the VHD.

You should run Sysprep on the original VHD before you create additional copies. Sysprep configures various operating system settings on the VHD to ensure that every copy of the operating system is unique when you distribute it to additional virtual machines.

Sysprep can be found on any Microsoft Windows Server 2003 or Microsoft Windows XP Professional operating system CD. To obtain Sysprep, open the Support\Tools folder on any Windows XP Professional or Windows Server 2003 operating system CD, and then open Deploy.cab.

For more information about using Sysprep, see the Microsoft Windows Corporate Deployment Tools User's Guide (Deploy.chm). Deploy.chm is included in the Deploy.cab file in the Support folder on the Windows Server 2003 operating system CD.

1.6.10 Do not use the VMRC client over a Terminal Services connection

You may experience poor mouse control if you attempt to use Terminal Services to access the host operating system and then use the Virtual Machine Remote Control (VMRC) client to access a guest operating system. You can, however, use Terminal Services to access the guest operating system directly if the guest operating system supports Terminal Services.

Important: Virtual Server is not a replacement for Terminal Services. The VMRC client does not operate in the same manner as Terminal Services or Remote Desktop because multiple users can use the VMRC client to connect to the same virtual machine, and each user can access the guest operating system without the knowledge of the other users. This is designed for training and lab scenarios where one user wants to demonstrate a task to other users who are connected to the same remote session.

If you want exclusive access to a virtual machine, then you should connect to it with Terminal Services or Remote Desktop and not the VMRC client.

1.6.11 Shut down virtual machines when making changes

When moving or copying virtual machines or any virtual machine components such as virtual hard disks, we recommend that you first shut down the virtual machine.

1.6.12 Ensure adequate storage space when using Undo Disks

The Undo Disks feature saves changes to a virtual disk in a separate file, which is stored in the same location as the .vmc file. Changes to the VHD are saved in the undo disk file until you either delete the changes, which erases the undo disk file, or commit the changes, which saves the changes to the original VHD.

Depending on the number of changes saved in the undo disk file, it can grow quite large, up to the maximum size allocated to the parent disk. If you use Undo Disks, you should first make sure that there is enough storage space for the undo disk file. By default, the .vmc file and undo disk file are stored in a folder with the same name as the virtual machine.

1.6.13 Write-protect the parent disk before using a differencing disk

You should always write-protect or lock the parent disk before using a related differencing disk (see 1.3.3, “Virtual hard disks” on page 11 for a discussion of these terms). Otherwise, if the parent disk is modified by another process, all differencing disks related to it become invalid, and all data written to the differencing disks is lost. You must also reconfigure the virtual machine to use the differencing disk instead of the parent disk. If you write-protect the disk and do not reconfigure the virtual machine, you will receive an error when you try to start the virtual machine because it cannot use a read-only disk.

1.6.14 Use Relative Weight

When you use the Relative Weight option on the CPU Resource Allocation page, Virtual Server adjusts the processor usage of each virtual machine dynamically, based on the existing workload. On the other hand, if you use the Reserved Capacity and Maximum Capacity options, you may need to actively monitor and adjust the settings to ensure you have the optimum configuration. To use Relative Weight, set Reserved Capacity to 0% and Maximum Capacity to 100%.

1.6.15 Ensure proper configuration for high disk-use scenarios

If workloads are extreme, the guest operating system may report a disk time-out. This can occur when a physical disk experiences extremely high usage due to a particular virtual machine configuration.

There are several ways to avoid this problem:

- ▶ Use a hard disk solution that allows fast access, such as a RAID array, or a Fibre Channel-based SAN.
- ▶ Put each virtual hard disk on a dedicated volume, SCSI hard disk, redundant array of independent disks (RAID), or in a SAN. It is easiest to put virtual hard disks together with their associated virtual machine configuration files on a RAID or SAN because this keeps everything in one place.
- ▶ Put VHDs on a different physical disk than the one used for the host operating system. In particular, you want to put VHDs on a different physical disk than the one used for the host page file.
- ▶ Reduce disk fragmentation. Defragment the physical disk regularly, especially if you are using a dynamically expanding VHD. The data stored on a dynamically expanding VHD

grows increasingly fragmented as the capacity of the disk increases because storage space is used only as it is needed.

As the capacity grows, it is less likely that the space will be contiguous. By contrast, a fixed-size virtual hard disk uses a reserved block of storage space, which means that data is less likely to be fragmented as it is stored.

- Compact virtual hard disks to free more physical disk space.

If none of these options is feasible, you can modify a registry key in the affected guest operating system as described in Knowledge Base article 818877:

<http://support.microsoft.com/?kbid=818877>

You should change the default value of the TimeOutValue registry key described in this article from 45 to 90.

Important: Incorrectly editing the registry may severely damage your system. Before making changes to the registry, you should back up any valued data on your computer.

1.6.16 The implications of using Virtual Server with Virtual PC 2004

When using Virtual Server and Virtual PC together, consider the following points:

- Sound cards

Virtual Server does not include an emulated sound card in its virtual machines, but Virtual PC does. If you are using a virtual machine in both Virtual Server and Virtual PC, you should disable the emulated sound card in Virtual PC. This prevents the sound card's Plug and Play capability from causing errors in the virtual machine that you created with Virtual Server.

- SCSI support

Virtual Server provides SCSI support, but Virtual PC does not. If you move a virtual machine with virtual SCSI disks that you created on Virtual Server to Virtual PC, the SCSI disks will be ignored. This can lead to negative consequences in many situation (for example, if the virtual machine page file is on the SCSI disk or if you are trying to use a SCSI disk as the startup disk).

If you plan to move VHDs between machines created with Virtual PC and Virtual Server, we recommend that you attach the VHDs only to a virtual IDE bus in Virtual Server.

- Configuration files

The .vmc files are typically compatible between Virtual Server and Virtual PC. However, if a virtual machine is in a saved state, do not move the .vmc file between Virtual Server and Virtual PC. The saved-state (.vsv) files are incompatible between the two products.

- CD-ROM drives

Although Virtual Server allows for virtual machines with multiple CD-ROM drives, Virtual PC supports virtual machines with only one CD-ROM drive. If you are moving virtual machines between the two products, you should configure only a single CD-ROM drive on the virtual machine and attach it to secondary channel 0, which is the default setting.

1.6.17 Secure Virtual Server and its associated components

You should take steps to secure Virtual Server and its associated components, including the physical computer, the Administration Website, the virtual machines, and the virtual machine components, such as the VHDs and virtual networks.

In addition, to the steps provided in the Securing Virtual Server feature, you should secure your virtual machines in the following ways:

- ▶ Enable a firewall on each virtual machine.
- ▶ Apply the latest security patches and updates to the operating system and to the applications that are running on each virtual machine.
- ▶ Install antivirus software on each virtual machine, as appropriate.
- ▶ Implement additional security lockdown procedures on each virtual machine, as appropriate.

1.6.18 Stop the Virtual Server service before host operating system shutdown

Before you shut down the host operating system, be sure that you stop the Virtual Server service. This ensures that the virtual machines are turned off in the manner that you have specified for them. Otherwise, the virtual machines may be effectively turned off rather than shutdown, which can result in a loss of data.

Similarly, if you are using an uninterruptible power supply, you should ensure that when you stop it, you also run a script to accomplish the following tasks in sequence:

1. Save the state of every virtual machine.
2. Stop the Virtual Server service.
3. Shut down the host operating system.

You should ensure that your UPS can provide power long enough to accomplish all of these tasks.

1.6.19 Improve Virtual Server application performance

You can improve the performance of Virtual Server by making the following configuration changes on the host operating system:

- ▶ Increase the acceleration of your graphics hardware to full. On the Display control panel, click the **Settings** tab, and then click **Advanced**. Click the **Troubleshoot** tab, and then move the Hardware acceleration slider to **Full**.
- ▶ Disable all unnecessary pointer options, such as pointer trails and shadow cursors. On the Mouse control panel, click the **Pointer Options** tab, and clear the check boxes of any options that you do not need.
- ▶ Optimize system settings for performance. On the System control panel, click the Advanced tab. Under Performance, click **Settings**, and then click **Adjust** for best performance. Click the Advanced tab, select **Background services**, and then click **OK** twice.
- ▶ Improve the speed with which the Administration Web site Master Status page displays and refreshes by hiding the virtual machine thumbnails that appear under Remote View. Clearing the **Remote View** check box on the Administration Website Properties page

1.6.20 Use appropriate high-availability techniques

You can use server clusters and Virtual Server together for training purposes or to test failover between two virtual machines on a single physical computer. You cannot, however, use server clusters with Virtual Server to fail over between host operating systems or virtual machines running on separate physical computers.

Server clusters are typically used with Virtual Server in a test or training environment and are not intended as a high-availability solution in a production environment. If you require high availability from your physical computer, you should implement other methods to achieve this, such as the use of fault-tolerant hardware.

1.6.21 Use the event log when troubleshooting

Virtual Server uses Windows event logging to record information about important events. You can view the event log with Event Viewer. We recommend that you review the event log when troubleshooting Virtual Server errors.



Virtual Server architecture

Modern, industry-standard servers consist of multiple layers of hardware and software that operate together as a system. Hardware resources typically include a processor, display, storage, networking, peripheral devices, and so on. The device drivers control hardware resources, translating operating system instructions into a specific device control language. Drivers are developed with an assumption of exclusive device ownership.

For example, it is assumed that a video driver owns a video adapter exclusively. Any software application that calls the video adapter must interact with the hardware through the video driver. When assumptions about exclusive device ownership are broken, systems typically fail to function properly.

The concept of exclusive device ownership typically precludes the possibility of running more than one operating system concurrently on a server. One approach to overcoming this limitation is virtual machine technology or virtualization. Virtualization involves redirecting interactions with device resources at lower levels so that higher-level application layers are unaffected. With Microsoft Virtual Server 2005, customers can run multiple operating systems concurrently on a single physical server.

In this chapter, we discuss the architecture of Microsoft Virtual Server 2005 in detail. We also describe networking, Virtual Server Additions, and virtual machine security.

2.1 Virtual machine technology

Figure 2-1 illustrates the basic architecture of the Microsoft virtual machine technology.

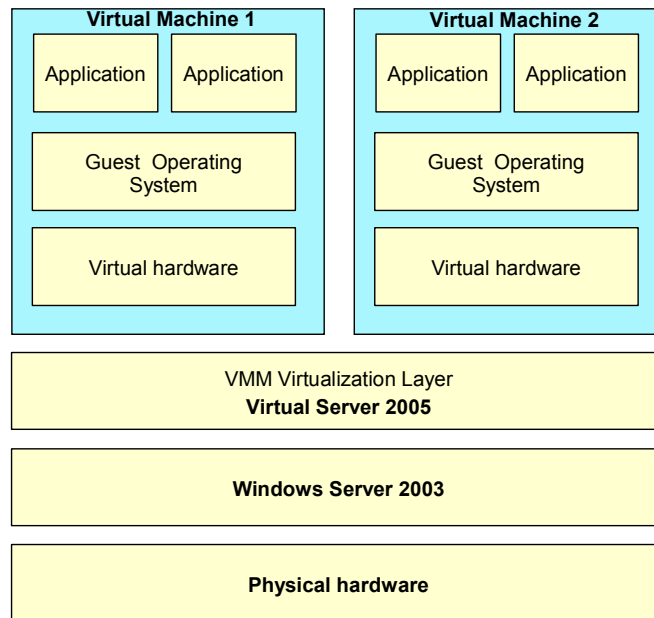


Figure 2-1 Virtual Server architecture

Starting from the bottom of the logical stack, Figure 2-1 shows the following activities:

- ▶ The host operating system, Windows Server 2003, manages the host system.
- ▶ Microsoft Virtual Server 2005 provides a Virtual Machine Monitor (VMM) virtualization layer that manages virtual machines. This provides the software infrastructure for hardware emulation.
- ▶ Each virtual machine consists of a set of virtualized devices. These devices comprise the virtual hardware for each virtual machine.
- ▶ A guest operating system and applications run in the virtual machine, unaware, for example, that the network adapter they interact with through Virtual Server is only a software simulation of a physical Ethernet device.

2.2 Virtual Server application structure

Microsoft Virtual Server 2005 is a multi-threaded application that runs as a system service. Each virtual machine runs in its own thread of execution. Input/output (I/O) occurs in child threads.

The host operating system provides the following core functions to Virtual Server:

- ▶ The underlying host operating system kernel schedules processor resources.
- ▶ The device drivers of the host operating system provide access to system devices.

The Virtual Server VMM provides the software infrastructure (illustrated in Figure 2-2) to create virtual machines, manage instances, and interact with guest operating systems.

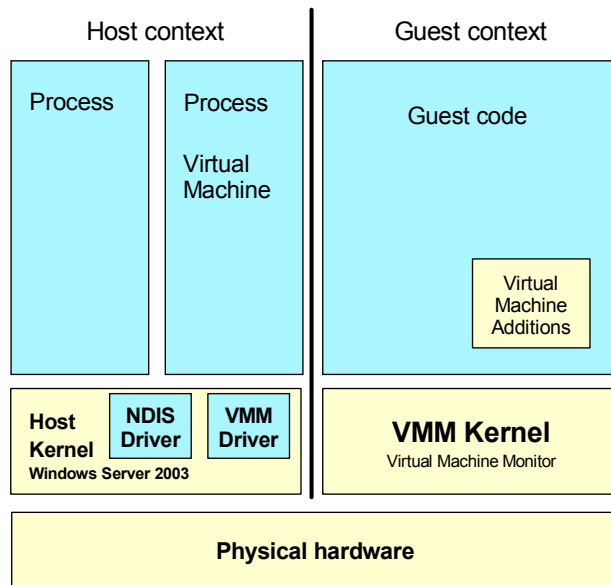


Figure 2-2 Virtual Server application structure

2.2.1 Virtual Machine Monitor kernel

The VMM virtualization layer manages virtual machines, providing the software infrastructure for hardware emulation. Each virtual machine consists of a set of virtualized devices.

All virtual machines run within a single process. Each virtual processor runs on its own thread. Additional threads are created to perform input/output (I/O), service the Component Object Model (COM) application programming interface (API), and so forth.

All software code in the virtual machine runs in a separate VMM context. This context consists of an address space that is completely separate from any Windows process, or from any other virtual machine. In this sense, the architecture is stronger than a simple per-process design.

The guest operating system and applications run on the virtual machine as though they were running on physical hardware, rather than emulated hardware. When a guest operating system is running, the special-purpose VMM kernel manages the processor and hardware during virtual machine operations. This creates an isolated environment on which the guest operating system and applications run close to the hardware at the highest possible performance level.

2.2.2 VMM driver

The VMM driver component of Virtual Server is a kernel-mode driver running on the host operating system. It has limited and specific functionality, acting as a firewall between the host operating system and virtual machines.

The VMM drivers has access to the physical processors and manages resources between the two environments. This prevents malicious or poorly designed applications that are running in a guest operating system from requesting excessive hardware resources from the host operating system.

2.2.3 NDIS driver

The virtual machine Network Driver Interface Specification (NDIS) driver is installed on the host operating system at a low level, just above the hardware network driver. It handles communication with the host operating system and devices in the network. It determines the routing of network packets, sending them to the host operating system or a virtual network adapter assigned to a virtual machine. This driver also spoofs unique media access control (MAC) addresses for the virtual network interface cards (NICs) and inserts their outgoing packets into the host Ethernet stack.

2.2.4 Virtual Server service

The Virtual Server service (vssrv.exe) creates virtual machines and provides all virtual machine functionality. It projects the emulated hardware into the virtual machine environment. It provides each virtual machine with its own 32-bit address space.

The Virtual Server service also includes the Virtual Machine Remote Control (VMRC) server. This server uses the VMRC protocol to provide a mechanism by which remote VMRC clients can interact with virtual machines.

A COM interface is implemented in the Virtual Server service that permits the user to monitor and control the virtual machine environment. All Web-based user interfaces in Microsoft Virtual Server 2005 use this COM interface through a scripting language.

2.3 Virtual networking

The virtual network architecture of Virtual Server allows the traffic in each virtual network to be isolated from that of other virtual networks. Communication with the host operating system and devices in the network is handled by the virtual machine NDIS driver. For more information about the NDIS driver, see 2.2.3, “NDIS driver” on page 28.

The degree to which the network traffic of virtual machines and the host operating system is isolated depends on the configuration of the virtual networks and virtual machines. Configuration options are as follows:

- ▶ Virtual network is not attached to a physical network adapter
In this scenario, the virtual network is a self-contained private network with its own optional virtual DHCP server. The network traffic for the virtual machines attached to this network. This completely isolates the host operating system.
Because it is isolated, the host operating system cannot read, monitor, or capture the network traffic of the virtual machines. By the same token, the virtual machines cannot read, monitor, or capture the network traffic of the host operating system. In addition, all network traffic is confined to the physical computer—in other words, isolated from the physical network. This scenario can be used to simulate a complete networking environment inside Virtual Server.
- ▶ Virtual network is attached to a dedicated physical network adapter
In this scenario, when no other virtual networks are attached to this physical network adapter, the virtual machines attached to this network cannot read, monitor, or capture the host operating system's network traffic. The host operating system, however, can read, monitor, or capture network traffic between a virtual machine and another device on the physical network, although not the network traffic between two virtual machines.
With this configuration option, a group of virtual machines can have their own isolated connection to a specific network.

- Two or more virtual networks are attached to the same physical network adapter.

When two virtual networks are attached to the same physical network adapter, the network traffic is only partly isolated. Virtual machines attached to such virtual networks can read, monitor, and capture one another's inbound network traffic, although they cannot read, monitor, and capture one another's outbound traffic.

This network configuration is normally used when there are more virtual networks than network adapters or network ports available.

2.4 Virtual hard disks

A virtual hard disk (VHD) provides storage for a virtual machine. Within the virtual machine, the VHD is represented as a physical disk and is used by the virtual machine as though it were a physical disk.

2.4.1 VHD architecture

Technically, the VHD is a file that resides on a physical disk that the host operating system can access. On the physical disk, the VHD file is stored as a .vhd file. As a general rule, a .vhd file can be on any type of storage device, as long as the host operating system can access that device. If the device is available to the host operating system, it is also available to Virtual Server 2005, and as a result, to any of the virtual machines.

For example, any of the following types of storage can be used:

- Integrated drive electronics (IDE) drive
- Small Computer System Interface (SCSI) drive
- Redundant array of independent disks (RAID)
- Storage area network (SAN)

The maximum size for a VHD is 2040 GB. However, any VHD attached to the IDE controller cannot exceed 127 GB. To support a larger VHD size, attach the VHD to a SCSI adapter.

Recommendation: We recommend attaching one or more virtual SCSI adapters, instead of using virtual IDE adapters, as this improves disk I/O performance. IDE is limited to one transaction at a time, regardless of whether the bus is physical or virtual. This means that a virtual machine with two VHDs attached to the IDE adapter is limited to a single transaction for both disks.

By contrast, a SCSI adapter allows for multiple simultaneous transactions, which provides better performance than disks attached to the IDE adapter. However, IDE is the default to maintain compatibility with the Virtual PC product.

In a virtual machine, there is no direct access to the physical disk that stores the .vhd file. This means no information about the physical disk can be accessed from the virtual machine. For example, format and label information about a physical disk cannot be obtained from a virtual machine.

2.4.2 Clustering support

Virtual Server clustering uses a virtual shared SCSI bus to implement the quorum device. The feature provides software failover support between the clustered virtual machines.

2.4.3 VHD types

Virtual Server provides different types of VHDs

- ▶ Fixed-size
- ▶ Dynamically expanding
- ▶ Linked
- ▶ Undoable
- ▶ Differencing

Fixed-size disks

A fixed-size VHD is a .vhd file, the size of which is determined when the file is created. Even when the amount of data being stored on the file changes, the size of the .vhd file remains fixed. For example, if a fixed-size VHD of 1 GB is created, Virtual Server creates a 1 GB .vhd file.

The size of a fixed-size VHD does not change because all of the storage space that is available on a fixed-size VHD is reserved when the VHD is created. The file utilizes as much contiguous space as is available on the physical disk that stores the .vhd file. The reserved space is filled as needed when data is written to the disk. The space on a fixed-size disk is more likely to be contiguous than on a dynamically expanding disk, so fixed-size disks generally provide better performance. Also, the file size of a fixed-size VHD does not need to be expanded before data is written to the file, which also helps provide better performance.

Dynamically expanding disks

A dynamically expanding VHD is one in which the size of the .vhd file grows as data is written to it. This is the default type of VHD created by Virtual Server.

When a dynamically expanding VHD is created, the maximum file size is specified. This size restricts how large the disk can become. However, the initial size of the .vhd file is only about 3 MB. For example, if a 1 GB, dynamically expanding VHD is created, the initial size of the .vhd file will be about 3 MB. As a virtual machine uses the VHD, the size of the .vhd file grows to accommodate the new data. The size of any dynamically expanding disk only grows; it does not shrink, even when data is deleted. The size of a dynamically expanding disk may be reduced by compacting it.

Occasionally, a Virtual Server may be unable to expand the VHD to the maximum size. This problem occurs when there is not enough free space on the physical disk that stores the .vhd file. Virtual Server monitors the free space on the physical disk. If the dynamically expanding disk starts to approach the limits of available space left on the volume in which the VHD file is stored, Virtual Server pauses the virtual machine and records an error in the Virtual Server event log.

Linked disks

A linked VHD is a VHD that points to and uses an entire physical disk to convert a physical disk to a VHD. A linked disk can be associated only to a drive; it cannot be associated to a volume. Because linked disks are intended only for conversion, a virtual machine cannot be turned on if a linked disk is attached to the virtual machine.

Undo and differencing disks

Undo disks and differencing disks store all state changes to a virtual machine or VHD in a separate file. This isolates changes to a virtual machine and maintains a VHD in an unchanged state. The main difference between an undo disk and a differencing disk is that undo disks apply to all VHDs associated with a virtual machine and a differencing disk applies to one VHD only.

A size cannot be specified for either type of disk. Both types of disks can be considered a special type of dynamically expanding disk. The size of any dynamically expanding disk only grows; it does not shrink, even when data is deleted. Undo disks and differencing disks can grow as large as the parent disks to which they are associated. However, unlike dynamically expanding disks, undo disks and differencing disks cannot be compacted directly. The parent disk can be updated with the changes stored in the undo or differencing disk. Then the parent disk can be compacted, if it is a dynamically expanding disk. If the parent disk is a fixed-size disk, it can be compacted by converting it to a dynamically expanding disk and then compacting the converted disk.

Undo and differencing disks can be very useful in test and development environments. They provide an easy way to start tests many times from the same clean state by discarding all changes afterwards. Service packs or new application levels can be tested for some time before committing the changes.

2.5 Virtual Server Additions

The Virtual Machine Additions component is very important to the process of running a virtual machine. Virtual Machine Additions should be installed on all virtual machines.

Virtual Machine Additions adds the following enhancements to a guest operating system:

- ▶ Improved mouse cursor tracking and control
- ▶ Greatly improved overall performance
- ▶ Virtual machine heartbeat generator
- ▶ Optional time synchronization with the clock of the physical computer

Virtual Machine Additions also adds keys to the registry of the guest operating system. The registry then can be queried for information about the virtual machine and the physical computer running Virtual Server.

Virtual Machine Additions is included for the following supported server operating systems:

- ▶ Microsoft Windows Server 2003 (all versions)
- ▶ Microsoft Windows 2000 Server
- ▶ Microsoft Windows NT Server 4.0 with Service Pack 6a (SP6a)

Virtual Machine Additions is provided as a convenience for the following client operating systems:

- ▶ Microsoft Windows XP (all versions)
- ▶ Microsoft Windows 2000 Professional
- ▶ Microsoft Windows Millennium Edition
- ▶ Microsoft Windows 98
- ▶ Microsoft Windows 95

Note: Virtual Machine Additions are the same for Virtual PC and Virtual Server. Additions are also included for IBM OS/2, but these are not supported in Virtual Server by Microsoft.

2.6 Virtual machine security

Microsoft Virtual Server 2005 offers the following security features:

- ▶ Isolation of virtual machines running on the same server
- ▶ User authentication
- ▶ Administrative network connections

2.6.1 Virtual machine isolation and encapsulation

The architecture of Virtual Server enables two key features: isolation and encapsulation.

- ▶ Isolation

Isolation is the separation or insulation of one entity from other entities. Virtual Server isolates virtual machines to prevent them from accessing resources or data owned by other virtual machines or the host operating system. The VMM ensures that each virtual machine has its own dedicated 32-bit address space that is fully isolated from the address space of other virtual machines and the host operating system.

If a virtual machine encounters a software failure, isolation enables other virtual machines and the host operating system to continue running. Isolation results in a robust and resilient architecture, which means that you can run even badly behaved applications in a stable, secure environment.

- ▶ Encapsulation

Encapsulation is the packaging of data and processing within a single object. In the case of virtual machines, this means that virtual machines are packaged into a single .vhd file that you can easily move among host operating systems running Virtual Server. Using encapsulation to decouple applications from hardware simplifies capacity planning, deployment, and management tasks.

2.6.2 User authentication

Virtual Server uses standard Windows features for user authentication. Authentication for VMRC connections can be configured on the administration Web site. The administration Web site uses the authentication method configured in Microsoft Internet Information Services (IIS).

VMRC connections

Virtual Server can be configured to use one of three different user authentication methods for virtual machine remote control connections:

- ▶ NTLM/Active Directory

With this authentication method, all Windows users with appropriate permissions can administer the virtual machine.

- ▶ Kerberos

When selected and used in conjunction with an Active Directory domain controller, all users of Windows Server 2003, Windows XP, and Windows 2000 operating systems with appropriate permissions can administer the virtual machine.

- ▶ Automatic

With this method, the VMRC client automatically negotiates the authentication protocol, either NT LAN Manager (NTLM) or Kerberos. This allows all Windows users with appropriate permissions to administer the virtual machine.

Virtual Server administration

Virtual Server does not require any additional IIS security configuration beyond the security model determined to be appropriate for the target environment. IIS should be configured to allow access to the Administration Web site as is appropriate for that environment.

By default, Virtual Server uses Integrated Windows authentication with anonymous access disabled. This is the preferred method of authentication and typically should not be changed. However, if a computer running Virtual Server is in an untrusted domain, the Administration

Web site cannot be accessed with Integrated Windows authentication. Instead, basic authentication must be used. More information about these authentication methods can be found in the documentation for IIS.

Virtual Server provides three levels of administration access to virtual machines based on host access control lists (ACLs). By default, administrators of the host operating system have access to all three:

- ▶ Creation of new virtual machines
- ▶ Reconfiguration of a specific virtual machine
- ▶ User access to a specific virtual machine

2.6.3 Administrative Network connections

Network transport for VMRC connections can be encrypted using Secure Sockets Layer (SSL). SSL configuration is provided through the Virtual Server administration Web site.

SSL encryption for the administration Web site is recommended when using basic authentication. This is because with Basic authentication, passwords are transmitted in plain text. The use of SSL encryption is configured from within IIS.



Customer scenarios

In this chapter we discuss the various benefits and pitfalls of four typical implementations of Microsoft Virtual Server 2005.

The benefits offered by implementing a virtualized environment provided by Virtual Server are twofold:

- ▶ Improved hardware efficiency

The combination of Virtual Server and Windows Server 2003 provides a virtualization platform that can improve hardware efficiency across a wide range of xSeries and BladeCenter servers. This platform can also run different Windows operating systems in the guest environment. Policy-based management features offer both weighting and constraint methods for fine-grained control of individual virtual machines.

- ▶ Increased administrator productivity

Virtual Server's comprehensive Component Object Model (COM) application programming interface (API), in combination with the Virtual Hard Disk (VHD) format and support for virtual networking, help make administrators more productive. This is because this combination allows scripted control of portable, connected virtual machines. These features offer deployment automation and ongoing change configuration.

The Virtual Server Administration Web site offers users remote access to and administration of individual virtual machines, along with Secure Sockets Layer (SSL) authentication. They can also manage a Virtual Server environment with their existing physical server management tools because Virtual Server integrates with a wide variety of IBM and Microsoft tools. These tools include:

- IBM Director 4.2
- IBM Virtual Machine Manager
- Active Directory
- Microsoft Operations Manager
- Windows Server 2003 Automated Deployment Services

3.1 Consolidating test and development environments

Developers and system managers are continually seeking ways to lower costs while accelerating application and infrastructure installations and upgrades and delivering a comprehensive level of quality assurance to prevent costly system failure or business application downtime. To achieve testing coverage goals prior to going into production, multiple challenges must be overcome:

- ▶ **Server operational and capital costs:** High-quality application test coverage requires replicating production environments in development sandboxes, adding risk to budgets and schedules.
- ▶ **Network operations:** Incorrect configuration of a test network could impact production networks.
- ▶ **Developer productivity:** Developer time is wasted on time-consuming administrative tasks.

Virtual machine technology was developed to address these challenges with side-by-side testing and production partitions on the same physical system. With Virtual Server, Businesses can consolidate their test and development server farm and automate the provisioning of virtual machines. This can improve hardware utilization and operational flexibility.

3.2 Migrating legacy applications

A recurring challenge for many businesses is the management and maintenance of existing server-based applications. Business applications often outlive their original operating system or hardware. As support for these primary infrastructure elements diminishes over time, cost of ownership steadily increases.

Under ideal conditions, customers would prefer to continue running business applications unchanged, but three factors increase the urgency of legacy application re-hosting:

- ▶ **Diminishing hardware support for legacy operating systems**, such as Microsoft Windows NT 4.0
- ▶ **High time and cost of administering server-based legacy applications**
- ▶ **High cost and risk to upgrade or rewrite legacy applications for increased quality and agility**

Customers need support for legacy or orphaned applications, such as those written for Windows NT 4.0, so that they may run them on new hardware and software platforms. Many small and medium-sized organizations find, for example, that the independent software vendor (ISV) that installed their application years ago is no longer in business.

Customers have called for a solution that assures smooth application re-hosting while delivering solid application compatibility and support. Virtual Server 2005 delivers the best of both worlds. It offers application compatibility with legacy environments, plus the advantage of the reliability, manageability, and security-enhanced features of Windows Server 2003.

Virtual Server allows customers to run legacy applications in their native software environment on virtual machines, without rewriting application logic, reconfiguring networks, or retraining end users. Customers can refresh out-of-warranty infrastructure systems first. Then they can either upgrade or rewrite out-of-service applications with a timetable that best fits their business needs.

3.3 Consolidate multiple server workloads

Server consolidation reduces the amount of server hardware under management for an optimized IT infrastructure. Virtual Server can be used for server consolidation of infrastructure services, disaster recovery environments, and department or branch office services, especially for server-based applications on xSeries hardware.

Virtual Server accomplishes this because it allows multiple applications and services on disparate operating systems to coexist on the same hardware, thus increasing hardware utilization and manageability while reducing capital and recurring costs related to hardware and hosting. Ultimately, fewer physical systems are required, which reduces the enterprise hardware footprint.

Specific workloads demand specific approaches, and not all workloads are suitable for consolidation, as shown in Table 3-1.

Table 3-1 Recommended approaches for differing target workloads

Target Workload	Recommended Approach	Consolidation Scenario	Customer Benefit
Branch office and department applications	Windows Server 2003 and Virtual Server 2005	Enterprise migrates Windows NT 4.0 applications from 1,000 stand-alone servers to 50 centralized, 4-way, rack-mount systems	Legacy applications run unchanged in virtual machine partitions with improved performance, reliability, and manageability
Enterprise applications	Windows Server 2003 and Windows System Resource Manager (WSRM)	Medium-sized business moves enterprise resource planning suite onto 8-way systems running Windows Server 2003 and WSRM	Scalability enhancements are coupled with resource management for outstanding application performance
File and print	Microsoft Windows Storage Server 2004	Small business consolidates file and print servers by using network attached storage (NAS)	Migration wizards enable fast, simple upgrades for improved performance and reliability
Databases	Windows Server 2003 and Microsoft SQL Server 2000	Enterprise consolidates databases on highly scalable IA-64 systems	Consolidation tools such as SQL multi-instance improve scalability and reliability
E-mail	Windows Server 2003 and Microsoft Exchange Server 2003	Medium-sized business consolidates e-mail servers using Exchange Server on scalability cluster	Robust consolidation tools enable move from legacy systems
Web	Windows Server 2003 and Internet Information Services (IIS) 6.0	Hosting service consolidates extranet applications on blades	IIS 6.0 application pools better manage Web applications

Workload characteristics determine which approach brings the highest return:

- ▶ Departmental applications

For departmental-level and branch office applications, Virtual Server 2005 is optimal, because it allows applications to coexist on the same hardware. This increases hardware utilization and manageability while reducing capital costs.

- ▶ Enterprise applications

Hardware isolation and scalability provides I/O performance and system resiliency; scalable servers can also be clustered for high availability and disaster recovery. Windows System Resource Manager (WSRM) provides mainframe-style resource management for large, scalable applications.

- ▶ File and print servers

File and print performance under Windows Server 2003 is higher than under Windows NT 4.0. New features such as the Volume Shadow Copy Service and enhanced collaboration capabilities provide better user experiences. Upgrading and consolidating servers to a single operating system results in substantial savings in operational costs.

- ▶ Database servers

Homogeneous consolidation enables fewer instances of a database and operating system. The SQL Server 2000 multi-instance capability enables organizations to consolidate many databases on a single server for maximum scalability, reliability, and manageability.

- ▶ E-mail servers

As with database servers, homogeneous consolidation of e-mail servers is best achieved with scalable clustered hardware. Exchange Server 2003 removes the scalability limits that once required the division of workload across multiple servers by off loading storage for multiple databases. Exchange Server 2003 can compress communication over the network, ensuring that servers can be centralized and fully utilize even low-bandwidth connections.

- ▶ Web servers and terminal servers

Improvements in Windows Server 2003 make it possible to consolidate Web sites, Web applications, and terminal servers through workload management tools such as WSRM. Virtual machine technology also helps older implementations of terminal services achieve higher scalability. Windows Server 2003 and IIS 6.0 provide significant cost savings opportunities.

Virtual Server is the recommended approach for server consolidation at the department or branch level, especially for custom server-based applications running on industry-standard hardware.

Virtual Server enables another form of server consolidation through disaster recovery.

3.3.1 Disaster Recovery

Rather than maintaining redundancy with costly physical servers, customers can use Virtual Server 2005 to back up their mission-critical functionality cost effectively with virtual machines. The Virtual Machine Monitor (VMM) and VHD technologies in Virtual Server 2005, combined with its comprehensive COM API, can be used to create similar failover functionality for standard, hardware-driven disaster recovery solutions.

Customers can use the Virtual Server COM API to script periodic duplication of physical hard disks containing vital business applications to virtual machine VHDs. Additional scripts can

switch to the virtual machine backup in the case of catastrophic failure. A failing device can be stopped for troubleshooting, or the application or database can be moved to another physical or virtual machine. Moreover, because VHDs are a core Virtual Server technology, they can be used as a disaster recovery agent, wherein business functionality and data can be easily archived, duplicated, or moved to other physical machines.

3.4 Simulate distributed applications on one physical server

It typically takes significant hardware resources and considerable amounts of time to deploy and test distributed server applications. The hardware and software systems must be configured in a lab environment to simulate a desired scenario, as described in “Consolidating test and development environments” on page 36

Virtual Server helps minimize these requirements in distributed server application testing scenarios. With Virtual Server, individual developers can easily deploy and test a distributed server application by using multiple virtual machines on one physical server. Virtual Server features, such as disk hierarchy and networking, gives developers an efficient way to simulate complex network environments on one physical server. The result is a time-saving and cost-effective solution because less hardware and less time are required for build-out.

For example, consider a large business with many distributed sites such as a bank. There would be a large headquarters with many branches, and in these branches, there would typically be a file server, print server, application server, a mail server, and so on. In the smaller branches, these systems would be required but not necessarily highly stressed; therefore, virtualization could be considered in these cases.

Currently, to deploy one of these sites would typically require 1-2 days of installation: cabling, power, hardware, and then software installation. However, with Microsoft Virtual Server 2005 and Virtual Server COM API, it could be possible to script the installation of the complete site, write it to a DVD, and supply it to the business complete with its one physical unit. The client then would only have plug in his server and install the DVD, reducing installation time, costs, and complexity considerably.



Management and deployment

When a server is running a virtualization application such as Microsoft Virtual Server 2005, managing that server becomes increasingly important and increasingly complex. Not only must you manage the physical hardware, the host operating system, and the applications that run on the host operating system, but you must also manage the guest operating systems and the applications that run in each of them.

In this chapter we describe some of the ways you can manage and deploy Virtual Server.

4.1 IBM Director 4.2

IBM Director Version 4.2 (IBM Director) is a workgroup management application for Intel processor-compatible systems. IBM Director consists of a suite of systems management tools that help users manage hardware to achieve maximum system availability and lower IT costs.

IBM Director's industry-standard foundation enables heterogeneous hardware support and works with a variety of operating systems and network protocols. It is a comprehensive hardware management solution that provides:

- ▶ Inventory of hardware features and settings
- ▶ System health notification
- ▶ Proactive and automated systems management capabilities
- ▶ Built-in subscription services for proactive upgrade protection
- ▶ Upward Integration into many other packages including:
 - Tivoli® Enterprise™ Framework
 - Tivoli NetView® NT
 - Computer Associates Unicenter TNG
 - Microsoft SMS and Microsoft Operations Manager
 - HP OpenView Network Node Manager
 - BMC Patrol
 - NetIQ

For more information about IBM Director 4.2 see the product announcement letter at:

<http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=an&subtype=ca&appname=redbooks&htmlfid=897/ENUS104-281>

For more detailed technical information about the product see:

http://www.ibm.com/servers/eserver/xseries/systems_management/director_4.html

For information about IBM Director 4.1 solutions, see *Implementing Systems Management Solutions using IBM Director*, SG24-6188.

4.2 Virtual Machine Manager 1.0

IBM Virtual Machine Manager (VMM) Version 1.0 is a free plug-in for IBM Director. You can download it from the Web at:

http://www.ibm.com/servers/eserver/xseries/systems_management/xseries_sm/vmm.html

With VMM, you can use Microsoft Virtual Server 2005 (and VMware VirtualCenter) in an IBM Director environment. For integration with IBM Director, VMM provides an agent extension that supports Virtual Server.

VMM uses the application programming interface (API) provided by Virtual Server that allows you to perform the following tasks from IBM Director Console:

- ▶ Correlate relationships between physical platforms and virtual components
- ▶ Report status of physical platforms and their corresponding virtual components
- ▶ Log in to the management interface of the virtualization application
- ▶ Discover virtual components
- ▶ Perform power operations on virtual machines
- ▶ Create event action plans that involve virtual components

With VMM, IBM Director can recognize systems that contain virtual components and create the appropriate managed objects for these systems. These managed objects are coordinators, hosts, and guest operating systems.

Without VMM, IBM Director can recognize them only as managed systems. Together, IBM Director and VMM can also create new IBM Director managed objects that represent the logical concepts of farms and virtual machines.

4.2.1 VMM integration with IBM Director

VMM supports integrated physical and virtual management by providing the following extensions to IBM Director components:

- ▶ VMM Server extension for IBM Director Server
- ▶ VMM Console extension for IBM Director Console
- ▶ VMM Agent extensions for IBM Director Agent

These are illustrated in Figure 4-1.

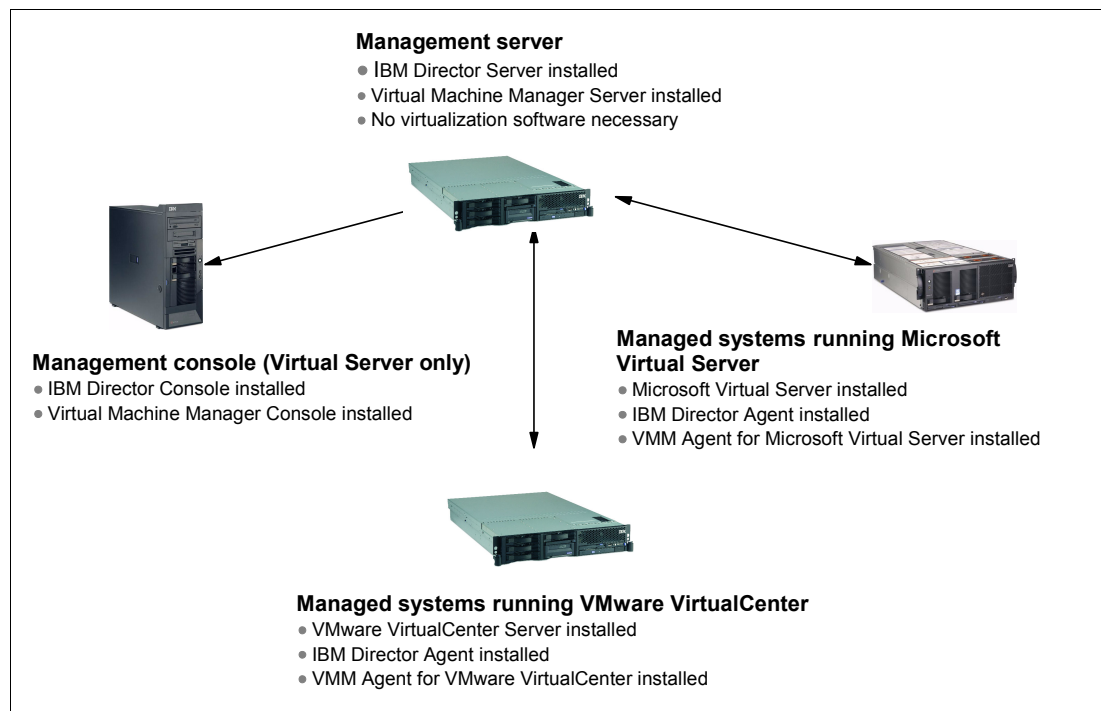


Figure 4-1 Overview of Virtual Machine Manager integration with IBM Director

VMM Server extension

The VMM Server extension provides the following features:

- ▶ VMM objects that are required for managing virtual components from IBM Director:
 - Coordinators
 - Farms (VMware only)
 - Hosts
 - Virtual machines
 - Guest operating systems
- ▶ Correlation of VMM objects with physical objects

- ▶ IBM Director groups and associations that are used when viewing VMM objects in the IBM Director Console Group Contents
- ▶ Event filtering and actions to support IBM Director Event Action Plans that involve VMM objects
- ▶ Event logging

VMM Console extension

The VMM Console extension does not provide its own distinct user interface. Instead, it adds functionality directly to IBM Director Console.

The VMM Console extension provides the following features:

- ▶ Context-sensitive menu options in the Group Contents pane for VMM objects and the windows that support these functions
- ▶ Virtual Machine Manager tasks in the Tasks pane of IBM Director Console
- ▶ Icons that depict the states of VMM objects

In addition, when you perform an operation on a VMM object in IBM Director Console, VMM uses the Virtual Server API to perform the operation on the virtual component itself. Both IBM Director Console and Virtual Server are updated with the results of the operation that was performed.

VMM Console operations, such as powering on a virtual machine, are executed as background processes. Therefore, IBM Director Console can be used for other operations while VMM performs its operations.

VMM Agent extensions

VMM provides a VMM Agent extension for Virtual Server (and VMware VirtualCenter).

The VMM Agent extension for Virtual Server provides the following features:

- ▶ Communication between the VMM Server extension and Virtual Server
- ▶ Discovery of Virtual Server and the creation of the corresponding VMM objects to enable IBM Director management of virtual components
- ▶ Translation from events provided by Virtual Server into events that can be used in IBM Director Event Action Plans

Note: The VMM Agent extension for Virtual Server only provides support for those features provided by Virtual Server itself. For example, the VMM Agent for Virtual Server does not provide this support for the migration of virtual machines from one host to another. The VMM Console component controls which features you can access for VMM objects.

4.2.2 Virtual Machine Manager tasks

The console component of Virtual Machine Manager is an extension of IBM Director. It provides the following console subtasks:

- ▶ Help subtask

Use the Help subtask to access and view the help pages that are provided with Virtual Machine Manager. The Help subtask is highlighted in Figure 4-2 on page 45.

► Start Vendor Software subtask

Use the Start Vendor Software subtask to start Virtual Server (see Figure 4-2 on page 45). If Virtual machine Manager cannot locate the vendor software, it displays a prompt that asks if you want to browse for the location.

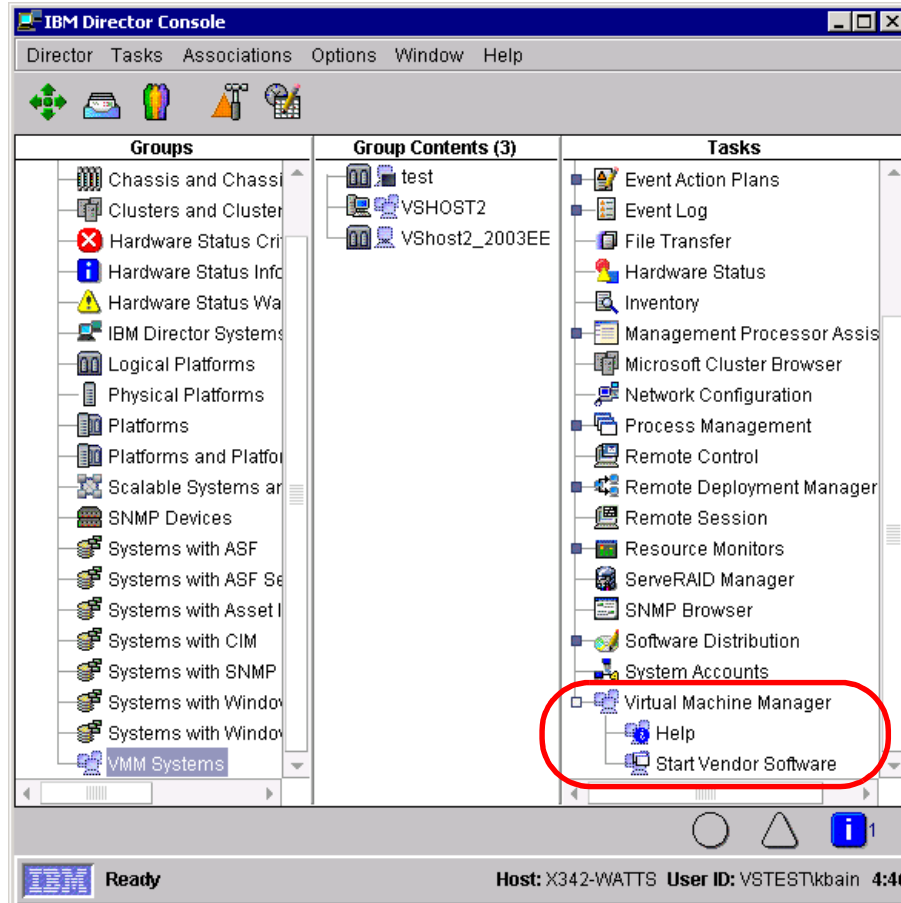


Figure 4-2 Displaying VMM tasks in IBM Director Console

4.2.3 Scheduler tasks for use with VMM objects

You can use IBM Director Scheduler to create jobs that perform tasks on one host or on one virtual machine.

Tasks for one host

You can use the Scheduler in IBM Director Console to create scheduled jobs that perform operations on a specific host.

You can perform operations that affect a host directly or affect all virtual machines on a host.

Important: Enabling a virtual-machine-related power operation on the Virtual Server host causes all virtual machines defined for that host, even those virtual machines that are not represented in the IBM Director Console, to be affected.

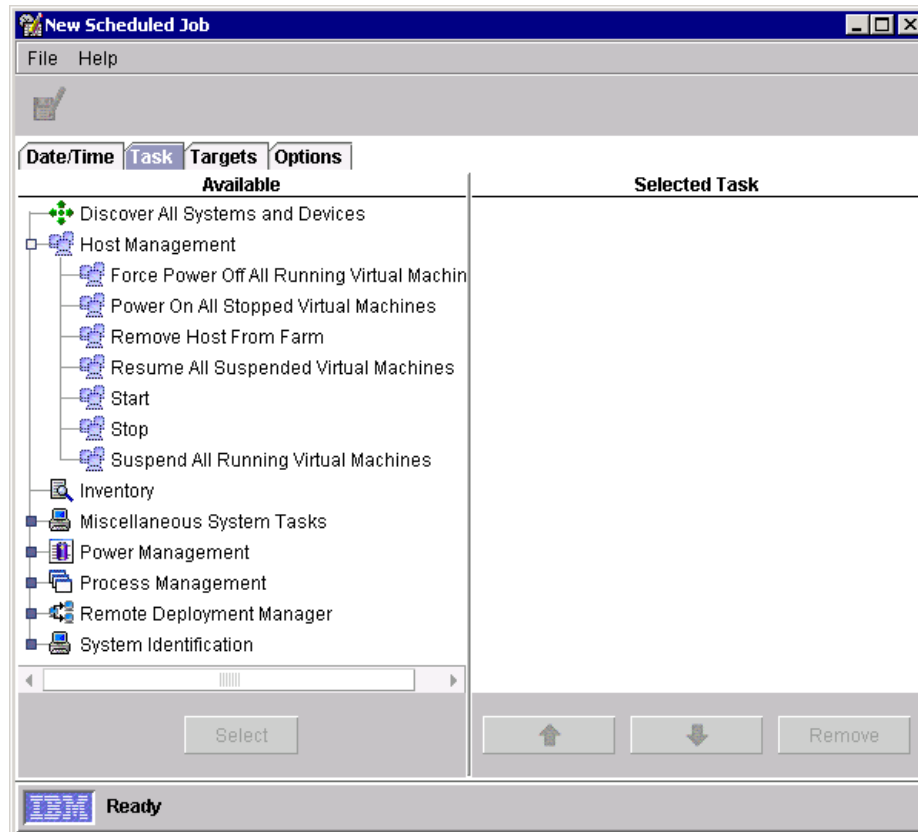


Figure 4-3 Scheduler tasks for host management

The possible tasks that may be scheduled on a Virtual Server host are shown in Figure 4-3.

Table 4-1 describes each of the Virtual Server host's scheduled tasks.

Table 4-1 Definition of scheduled power operations available on the Virtual Server host

Scheduler Task	Description
Force Power Off All Running Virtual Machines	Turns off all running virtual machines defined for the host without gracefully shutting down any guest operating systems
Power On All Stopped Virtual Machines	Turns on all stopped virtual machines defined for the host
Resume All Suspended Virtual Machines	Resumes all suspended virtual machines for the host
Start	Applies only to hosts that are running Virtual Server. Starts the host represented by the managed object. You can only create scheduled jobs that use this task for hosts that are currently stopped
Stop	Applies only to hosts that are running Virtual Server. Stops the host represented by the managed object. You can create scheduled jobs that use this task only for hosts that are currently started
Suspend All Running Virtual Machines	Suspends all running virtual machines defined for the host.

Tasks for one virtual machine

You can also use the Scheduler in IBM Director Console to create scheduled jobs that perform operations on a specific virtual machine. The possible options for all systems are shown in Figure 4-4.

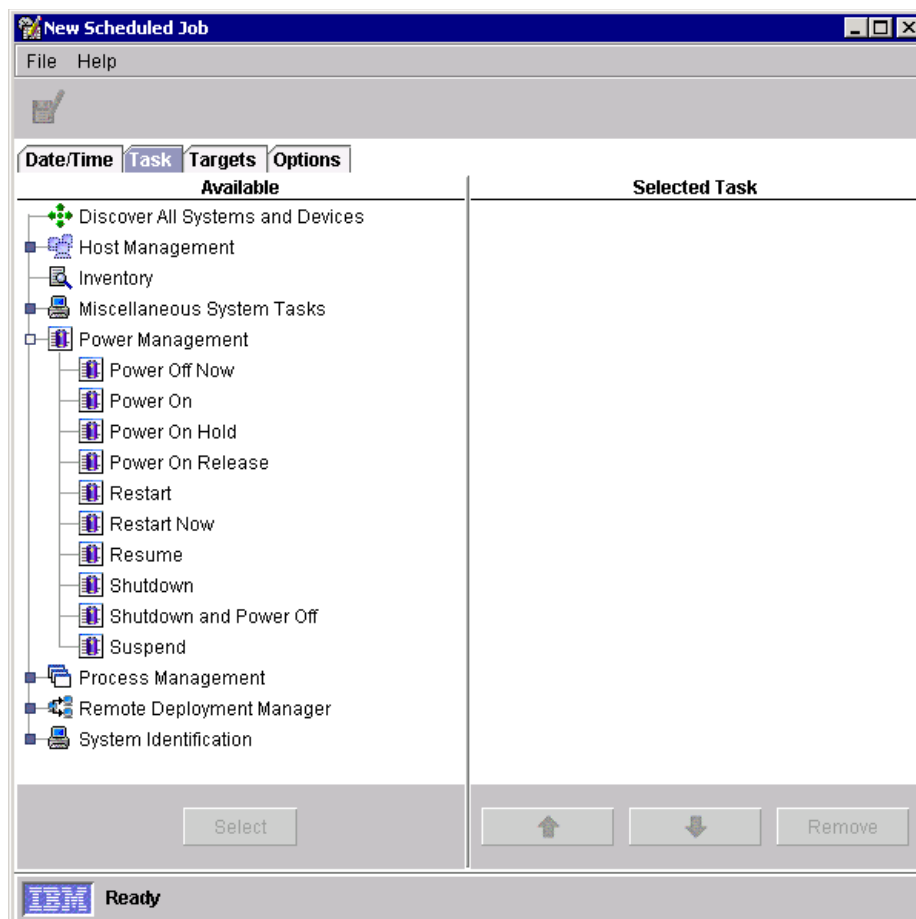


Figure 4-4 Scheduler tasks for power management on any system

Only a subset of these options apply and only relevant options apply to specific virtual machines in specific states. For example, if a virtual machine is suspended, the Power Management menu only contains the Resume menu option.

Table 4-2 lists the possibilities and their meanings.

Table 4-2 Descriptions of possible power options for virtual machines

Menu option	Description
Power Off Now	The virtual machine is turned off without gracefully shutting down its guest operating system. This abruptly stops all applications that are in use in that guest operating system.
Power On	The virtual machine is turned on.
Restart Now	The virtual machine is turned off without gracefully shutting down its guest operating system and then turned on again. This abruptly stops all applications that are in use in that guest operating system.
Resume	The virtual machine resumes operation and is no longer suspended.

Menu option	Description
Shutdown and Power Off	The guest operating system on the virtual machine is gracefully shut down and then the virtual machine is turned off.
Suspend	The virtual machine remains turned on, but is suspended from use.

4.2.4 Power operations for all virtual machines on a host

Use the Host Management menu to perform power operations for all virtual machines on a selected host. The Host Management menu is available when you right-click a virtual machine in the Group Contents pane of IBM Director Console as shown in Figure 4-5.

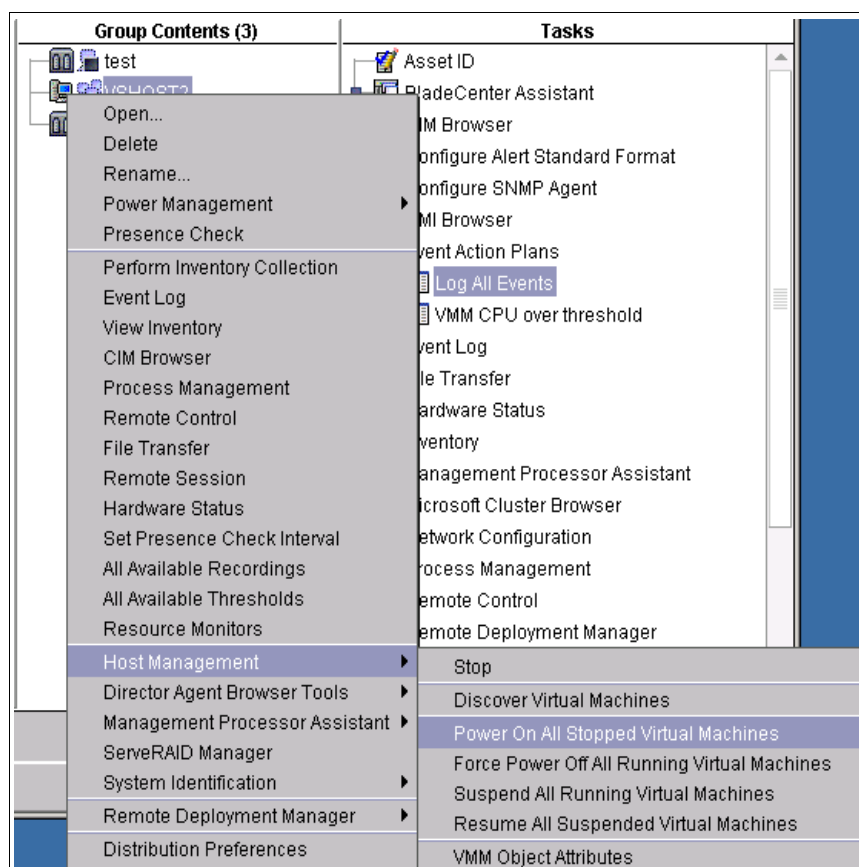


Figure 4-5 Host power operations available directly from the console

You can use the Host Management menu to perform the following power operations:

- ▶ Turn on all stopped virtual machines
- ▶ Suspend all running virtual machines
- ▶ Resume all suspended virtual machines
- ▶ Turn off all running virtual machines

It is important to note that the power operations in the Host Management menu apply to all virtual machines defined for Virtual Server, even those virtual machines that are not represented in IBM Director Console. Host management power operations are applied to all virtual machines defined for the selected host.

Additionally, the selected power operation applies only to the virtual machines in the appropriate state. For example, a power-on operation applies only to those virtual machines

that are turned off. It does not affect those virtual machines that are already turned on or suspended. This aspect is important to consider when you use IBM Director to create scheduled jobs for host management power operations that involve virtual machines.

4.2.5 VMM event filters and actions

VMM provides several events and event actions for use with virtual objects.

Important: Before using event filters and actions with virtual objects, use the Management Processor Configuration subtask of the IBM Director Management Processor Assistant (MPA) task.

In the alert-forwarding profile, the connection type must be set to *IBM Director Comprehensive* and the IP address must be set to the IP address of the management server that is being used to manage virtual objects.

The events that are provided by VMM are for VMM objects. These events are in the VMM event type in the Event Filter Builder and are divided into two categories:

- ▶ Virtual machine events
- ▶ Host events

Note: Make sure that any event action plans that use these events are actually directed to the applicable VMM object; otherwise, the intended actions will not occur. Use the IBM Director Event Log task to view details about all VMM event types that have been received and logged by IBM Director Server.

Simple event filters can be created with or without specifying extended attributes for the event filter. The target object for an event filter differs, depending on which method was used to create the event filter. If an event filter is created without specifying any extended attributes, the event action plan that contains this event filter should be applied directly to the VMM object identified as the target object in Table 4-3 on page 50 and Table 4-4 on page 52.

If an event filter is created, and it specifies values for extended attributes, the event action plan that contains this event filter can be applied directly to the VMM object identified as the target object or to any higher-level VMM objects that are associated with the target object. For example, if you create an event filter with extended attributes that identifies a specific virtual machine, you can apply the event action plan that contains this event filter directly to the virtual machine itself or to its associated host, farm, or coordinator.

In this scenario, all VMM objects that are associated with the target object listed in the following tables receive the event notification. This means that the IBM Director Event Log contains one event for each associated VMM object, each with the same detail. To continue the example, there would be four log entries, one each for the coordinator, farm, host, and virtual machine.

Virtual machine events

Virtual machine events are in the Virtual Machine subcategory of the VMM event type category. Figure 4-6 on page 50 shows the available virtual machine events that can be used in conjunction with an event action to produce specific task.

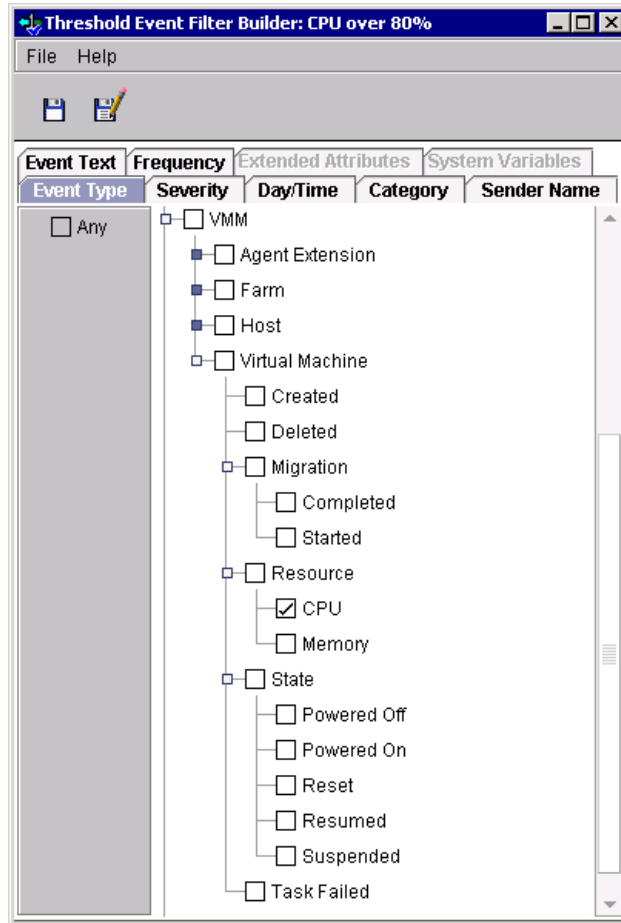


Figure 4-6 Available virtual machine events

In this case, we have created a customized event based on processor resources.

The definition of the resources shown above are shown in Table 4-3.

Table 4-3 Description of virtual machine events and triggers

Event subcategory	Additional event subcategory	Event trigger	Description
Created		A virtual machine object is created	IBM Director Console displays a virtual machine object for a virtual machine
Deleted		A virtual machine object is deleted	IBM Director Console no longer displays a virtual machine object for a virtual machine
Migration	Started	A virtual machine has started migration from one host to another	VMM has started the migration of a virtual machine from one host to another
	Completed	A virtual machine has completed migration from one host to another	VMM has completed the migration of a virtual machine from one host to another

Event subcategory	Additional event subcategory	Event trigger	Description
Resource	CPU	The values for the processor resource for a virtual machine have changed	One of the following events has occurred: <ul style="list-style-type: none"> ▶ The reserved (minimum) or maximum capacity of CPU utilization changed ▶ The relative weight of the CPU changed
	Memory	The values for the memory resource for a virtual machine have changed	The memory size has changed
State	Powered off	The state of a virtual machine has changed to off	A virtual machine is turned off
	Powered on	The state of a virtual machine has changed to on	A virtual machine is turned on
	Reset	The state of a virtual machine has changed to restarted	A virtual machine is restarted
	Resume	The state of a virtual machine has changed from suspended to on	A virtual machine resumes operations after being in the suspended state
	Suspended	The state of a virtual machine has changed to suspended	A virtual machine is suspended
Task Failed		An operation on a virtual machine has failed to be completed successfully	A power or migration operation failed for a virtual machine

Host events

VMM also provides events for hosts. These events are in the VMM event type category. Figure 4-7 shows the available event types.

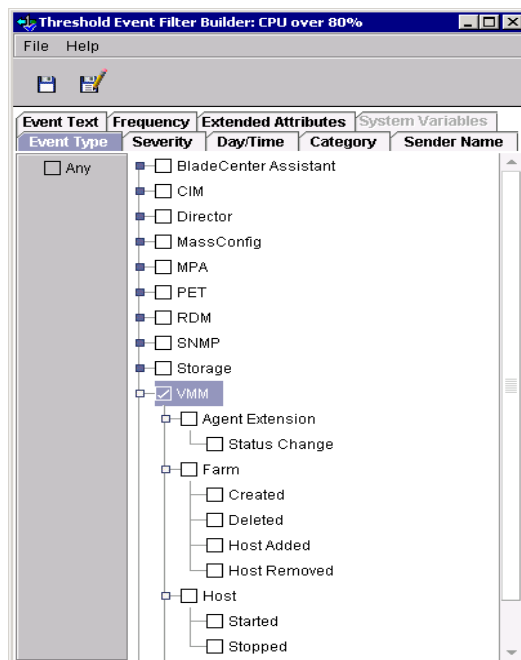


Figure 4-7 Host events

The definitions that relate to Virtual Server are described in Table 4-4.

Table 4-4 Description of virtual machine events and triggers

Event subcategory	Additional event subcategory	Target object for event action plan	Event trigger	Description
Agent Extension	Status Change	Host	VMM Agent detects a status change for the host	The following are examples of such events: <ul style="list-style-type: none"> ▶ VMM Agent is not running or the Virtual Server is not installed. ▶ VMM Agent and the Virtual Server are installed, but they are not communicating properly with each other. ▶ The host is ready for use. This means that Microsoft Virtual Server services are started.
Host	Started	Host	VMM agent detects that Virtual Server has started on a host	Virtual Server services are started on a host.
	Stopped	Host	VMM agent detects that Virtual Server has been stopped on a host	Virtual Server services are stopped on a host.

Note: Events of type Farm and Coordinator are VMware VirtualCenter only.

Event actions provided by Virtual Machine Manager

Figure 4-8 shows the options available in the event action plan builder.

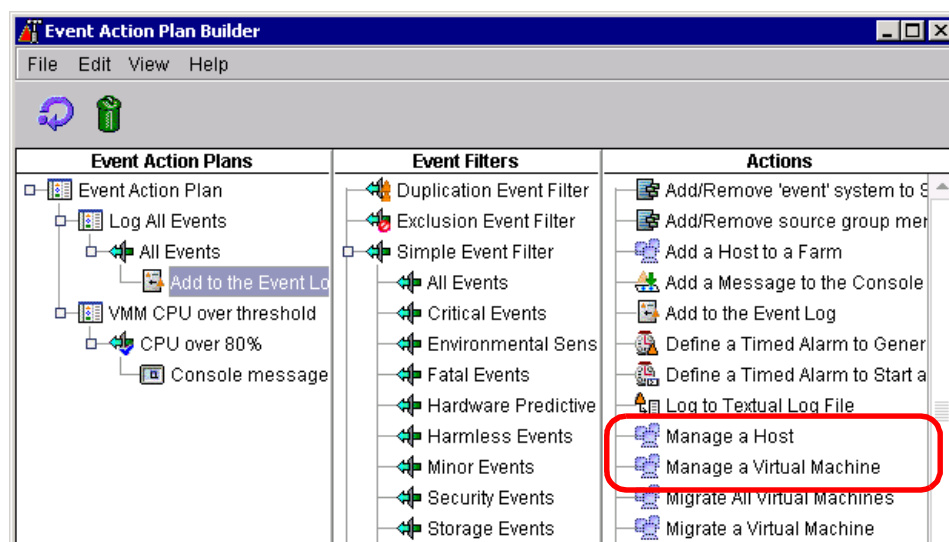


Figure 4-8 Actions added by Virtual Machine Manager for Virtual Server

Note: Other actions not listed (such as adding or removing a host from a farm) are for VMware VirtualCenter only

VMM provides the following event actions for performing power operations on hosts and virtual machines:

- ▶ Manage a Host
 - Start
 - Stop
 - Turn off all virtual machines
 - Resume all virtual machines
 - Suspend all virtual machines
- ▶ Manage a virtual machine
 - Turn on
 - Turn off
 - Shutdown and turn off
 - Turn off now
 - Suspend
 - Resume
 - Restart now

Note: VMM power operations do not apply to virtual machines that use the Undo Disks feature.

By creating custom event actions, you can specify which action you want IBM Director to take as a result of an event occurrence that is triggered by an event filter.

To create event filters for VMM objects, you can use VMM events. After you have created custom event actions and event filters, you can create an event action plan that contains specific filters and their associated actions. For example, as we demonstrated, you can create an event action plan that sends an alert when a specified resource is changed.

All event action plans must target the appropriate VMM object, or the action defined in the plan will not occur. For example, if you have an event filter for the Virtual Machine, Task Failed event, it needs to be included in an event action plan that targets a virtual machine for it to trigger an action. To use a custom event action, you must add it to an event filter that is already in an event action plan.

For more information about Virtual Machine Manager event filters and actions, please see the *Virtual Machine Manager Installation and User's Guide* available in PDF form from:

http://www.ibm.com/servers/eserver/xseries/systems_management/xseries_sm/vmm.html

4.3 Microsoft Operations Manager 2005

Microsoft Operations Manager 2005 (MOM) delivers extensible and scalable enterprise-class operational management with the following features:

- ▶ Comprehensive event management
- ▶ Proactive monitoring and alerting
- ▶ Reporting and trend analysis

MOM also has system-specific and application-specific tasks to improve the manageability of Windows Server system environments:

- ▶ User interfaces optimized for the current task, whether it be administration, operations, or reporting
- ▶ Operational knowledge obtained directly from developers and operational staff to help identify, understand, and resolve IT issues
- ▶ Easily extensible architecture to manage custom applications and integrate with third party solutions

MOM interfaces with Microsoft Virtual Server 2005 using add-ons known as management packs, as described in 4.4, “MOM Management Pack for Virtual Server” on page 55.

Note: For more information about MOM, go to the following Web site:

<http://www.microsoft.com/mom>

4.3.1 Management consoles

MOM consists of four management consoles for managing and configuring incoming data:

- ▶ Administration console

This console is a Microsoft Management Console (MMC) snap-in. It provides MOM administrators with the tools for configuring and optimizing MOM, installing management packs or creating rules, configuring and deploying agents, and identifying user views and security. The MOM Administrator can configure views that would be used in the Operator Console for each operator so that an Exchange administrator only works with Exchange Servers.

- ▶ Operator console

The Operator Console is an application built on .NET. It was designed to assist the administrators or operators that provide day-to-day monitoring and administration of IT services. It provides multiple views into the systems, helping administrators identify alerts, take actions, and resolve those alerts.

- ▶ Web console

The Web Console allows operators to access to health information from any computer in the network without requiring a full client. The console provides the basic functionality for viewing and resolving alerts and tracking system states.

- ▶ Reporting console

From a Web page, the Reporting Console provides access to all published reports. These reports can be made available to executives or systems operators, and can provide insight into capacity planning information, performance against service level agreements, trends in the environment, and summaries of open alerts.

4.3.2 Management packs

Management packs are MOM add-ons that allow interaction with applications such as Virtual Server. Managements packs consist of:

- ▶ Rules

Rules can be simple, such as detecting the occurrence or absence of a given event or determining whether a performance counter has exceeded a threshold. They can also be complex, such as running a series of scripts to perform synthetic transactions that verify

server response times and availability. Rules can create and prioritize alerts. These alerts identify conditions that require administrator intervention to prevent potential outages or quickly address issues that affect server performance or availability.

- Alerts

First level support can receive real-time alerts of problems as they occur.

- State monitoring

State monitoring provides a real-time view of the condition of line-of-business servers and applications, verifying that services are available and that they are providing end users with good performance. This information helps operations staff prioritize issues and provide higher overall availability for systems.

- Tasks

Tasks can take specific diagnostic or management actions based on the role of the server. They are initiated by the administrator and can include sequences that run on the MOM Console, MOM Server, or the server experiencing a problem. Examples include flushing the Domain Name System (DNS) cache on a specific DNS server or pinging the server to check connectivity.

- Reports

Reports are based on data from the MOM database. Reporting Console is a Web console that can be extended with custom reports using the report writing capabilities of SQL Server 2000 Reporting Services. Visual Studio .NET provides the development environment to create custom reports.

Management Packs also provide information to the discovery capability of a MOM server to automatically identify specific roles of servers (e.g. SQL Server, Exchange Server, Domain Controllers, and so on) and deploy the corresponding rules to those servers without any administrative action required.

Administrators can create their own customized Management Packs for internally developed applications. This gives an enterprise the ability to leverage its investment in MOM across all of its mission-critical applications. Likewise, third-parties can also create Management Packs for the solutions they produce.

4.3.3 MOM Connector Framework

Many organizations have invested in enterprise framework tools that provide a single, end-to-end view of the health of their IT infrastructure. As a result, it is important for monitoring tools to have bidirectional communications with those frameworks. This is so the appropriate administrator can learn about events and resolve the issues satisfactorily, while still allowing the event and performance information to aggregate into corporate-wide reports.

The MOM Connector Framework (MCF) facilitates this bidirectional communication by providing the ability to forward alerts and other data to an enterprise framework or other tools in a heterogeneous environment and receive status messages from those frameworks and tools to resolve alerts as necessary. MCF is based on Web services, so it is simple to create connections to various third-party systems.

4.4 MOM Management Pack for Virtual Server

MOM Management Pack for Virtual Server provides extensible guest-host mapping for event and performance management. Microsoft's manageability strategy for Virtual Server is to enhance the current manageability product set with virtual machine awareness. This is

evident in Virtual Server's ability to distinguish host/guest hierarchical relationships, represent and manage the topology, and expose it to other management solutions.

Virtual Server provides Windows Management Instrumentation (WMI) counters on the host system. Therefore, it can be integrated into a third party management solution or passed to MOM for health monitoring and alerting.

Events are logged and displayed in the Virtual Server event log. This log is available on the Virtual Server Administration Web site from the Master Status Web page. It is also displayed on the host system event log, which can be passed to MOM for alerting, or it can be integrated into a third-party management solution.

Virtual Server uses an extensible XML file format for virtual machine configuration management from external management software. This can be leveraged through the COM API for asset management, automated provisioning and deployment and more.

4.4.1 Feature overview

MOM collects and processes information from the virtual server environment using Virtual Server Management Pack. Virtual Server Management Pack provides pre-packaged rules, thresholds, knowledge, tasks reports. Operators can use these packages to better understand the event and the recommended response.

The Virtual Server Management Pack contains the following components:

- ▶ Six views (include State)
- ▶ 10 tasks
- ▶ 30 rules (include State Monitoring)
- ▶ Two reports

These components can be used to monitor the following key scenarios:

- ▶ Detect Virtual Server and virtual machine installation and configuration issues
- ▶ Detect virtual machine availability
- ▶ Detect server conditions that affect Virtual Server or virtual machine health, including disk space and system resources
- ▶ Detect virtual machine-related network issues
- ▶ Detect virtual machine image-related issues

Note: At the time of publication, the management pack was not available. Once it is available, you can download it from the Management Pack and Product Connector Catalog Web site:

<http://www.microsoft.com/management/mma/catalog.aspx>

Views

MOM information can be displayed in as many as six views. One of these is a state view that allows the user to monitor the current state of a virtual machine as shown in Figure 4-9.

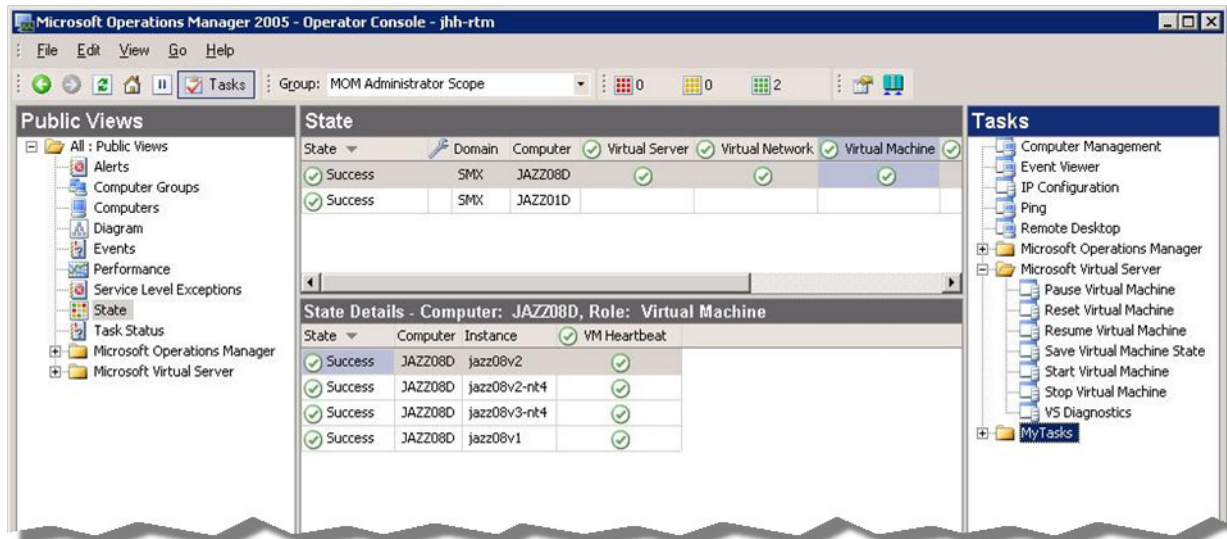


Figure 4-9 State view with associated tasks

Tasks

The tasks that can be performed are divided into three categories:

- ▶ Virtual Server health diagnostics
 - Virtual Server diagnostics
- ▶ Virtual Server service management
 - Start Virtual Server
 - Stop Virtual Server
 - Start virtual machine helper
 - Stop virtual machine helper
- ▶ Virtual machine management
 - Start virtual machine
 - Stop virtual machine
 - Save virtual machine state
 - Resume virtual machine
 - Reset virtual machine
 - Pause virtual machine

These tasks allow the user to have some measure of control over the Virtual Server service and the virtual machines themselves.

Rules

The Virtual Server Management Pack contains 30 rules, based on the health modelling process:

- ▶ Event
 - VHD operation failure: Compacting, converting, creating, or merging failed VHD
 - VHD operation failure: Disk operations may fail due to critically low disk space
 - Virtual machine/Server: Configuration error
 - Virtual machine: Cannot commit changes or merge a VHD
 - Virtual machine: Cannot mount virtual floppy disk
 - Virtual machine: Fatal error: the parent differencing disk is invalid
 - Virtual machine: Floppy disk image errors

- Virtual machine: General errors.
 - Virtual machine: VHD failure with performing action on differencing or undo drive
 - Virtual machine: VHD failure with performing action on differencing or undo drive
 - Virtual machine: Virtual Machine Additions are out of date
 - Virtual machine: Virtual machine cannot initialize its virtual network adapter
 - Virtual machine: Virtual machine has experienced a critical error
 - Virtual machine: Virtual machine is unable to start or write to disk
 - Virtual machine: Virtual machine: device is not available
 - Virtual machine: Virtual network could not be found
 - Virtual machine: Virtual machine could not be started
 - Virtual machine: Virtual machine restore failed
 - Virtual machine: Virtual machine save failed.
 - Virtual network error
 - Virtual network error: Virtual Server cannot access the Network Services driver
 - Virtual Server: Fatal error due to low system resources.
 - Virtual Server: General errors.
 - Virtual Server: Installation has become corrupt
- Performance
- VM CPU Utilization
 - VM Disk Space Used
 - VM RAM Used

These event and performance rules can be used to create and prioritize alerts. Alerts call attention to conditions that require administrator intervention. The administrator can use them to prevent potential outages or to quickly address issues affecting server performance or availability.

Reports

Virtual Server Management Pack provides 2 types of report:

- ▶ All Virtual Servers (illustrated in Figure 4-10)
 - Shows all Virtual Server host servers and the virtual machines they are hosting
 - Adds the free space on the hosted drive
- ▶ All Virtual Machine Details (illustrated in Figure 4-11)
 - Helps administrators get the detail view of their Virtual Servers
 - Helps to plan the resource utilization of each Server

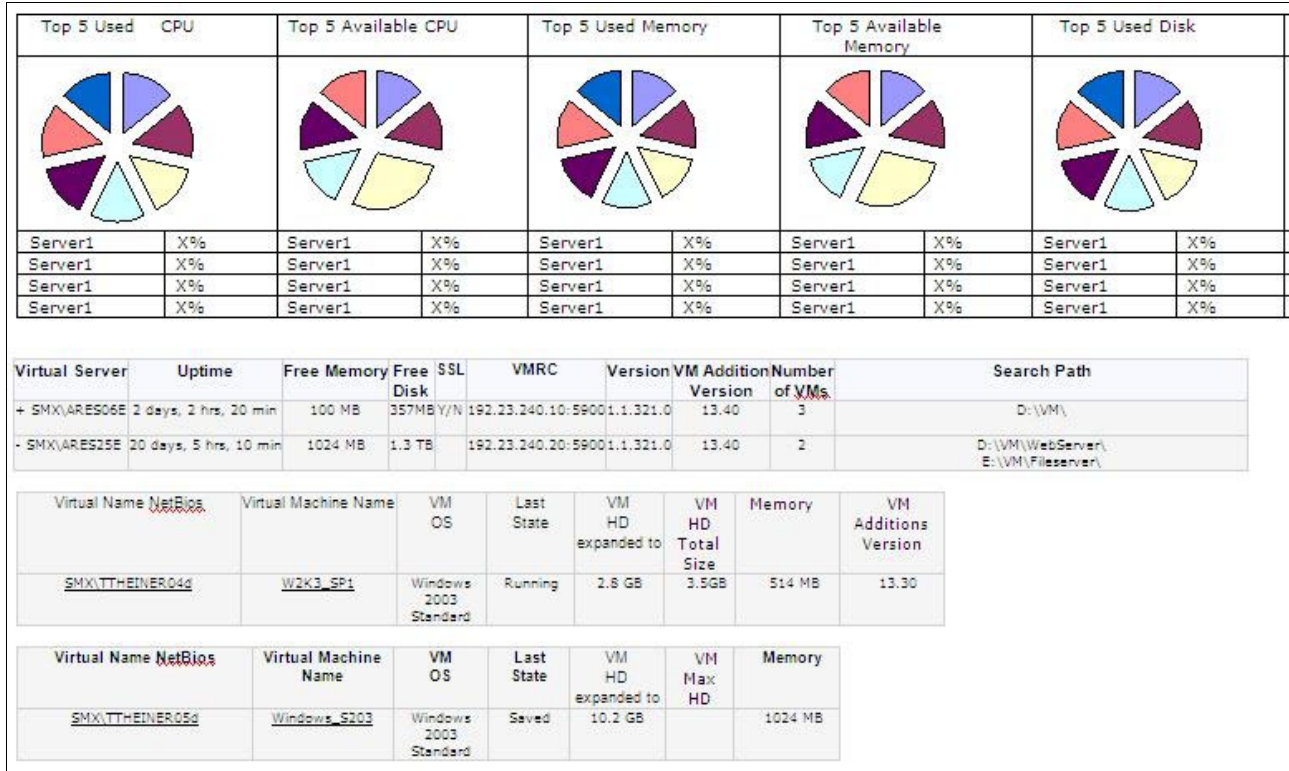


Figure 4-10 All Virtual Servers Report

Virtual Machine	Virtual Machine Name	Hosted on Virtual Server	Save State size	Undo size	VM Max HD	Total Size	Max CPU%	Network	Location
NetBios	Windows_S203	SMX\TTHEINER08D	4GB	2GB	16 GB	10.2GB	100	IP: 192.23.225.114 Subnet: 255.255.0.0 DNS: ttheiner04d.smx.netD:	D:\vm\mom2004beta3.vmd D:\vm\mom2004beta3.vsv D:\vm\pcundomom2004beta3.vud

VM Virtual Server Drive E: 1 Terabyte free Space				
Location	VM OS	Last State	VM HD expanded to	VM HD Total Size
D:\vm\mom2004beta3.vmd	Windows 2003 Standard	Running	2.8 GB	3.5GB

Figure 4-11 All Virtual machine details

4.5 Virtual Server Migration Toolkit

Virtual Server Migration Toolkit (VSMT) is a free set of tools that lets you migrate operating systems and applications that are installed on physical machines to virtual machines on a host server that is running Virtual Server. This is called “P2V”. Although not directly a systems management tool, VSMT can be useful to a system administrator in the appropriate environment.

VSMT is available from the following Web site:

<http://www.microsoft.com/windowsserversystem/virtualserver/evaluation/vsmt.mspx>

VSMT works with Automated Deployment Services (ADS) 1.0 to automate the migration of physical servers to virtual machines. ADS 1.0 is a free add-on for Windows Server 2003, Enterprise Edition and Datacenter Edition. It is a deployment provisioning and management tool for Windows Server operating systems.

You can download ADS from the following Web site:

<http://www.microsoft.com/windowsserver2003/techinfo/overview/adsbenefits.mspx>

4.5.1 VSMT infrastructure and system prerequisites

This section discusses the system and infrastructure prerequisites for installing and running VSMT.

Infrastructure requirements

ADS and Virtual Server must be deployed in the environment where the migration will be performed. The computers running the ADS Controller and Virtual Server must be either members of the same workgroup or joined to the same Active Directory directory service domain.

The ADS Controller can be installed on computers running Windows Server 2003, Enterprise Edition or Windows Server 2003, Datacenter Edition. The ADS Agent can be installed on computers running all editions of Windows 2000 Server and Windows Server 2003.

You can use an existing ADS infrastructure that is deployed in the environment where you perform the migrations. If an ADS infrastructure does not exist, you must create one.

All ADS services (Image Distribution, Network Boot Services, and Controller) must be available. In addition, you must be able to PXE-boot the physical source computers and Pre-Boot Execution Environment (PXE) broadcast must be enabled across all routers or IP Helpers must be configured.

You should verify that there is adequate disk space on the computer hosting the store for the Image Distribution service. This may also require that DHCP be made available in the environment and that network functionality be changed.

The ADS Controller and Virtual Server 2005 can be installed on the same computer or on different computers. In addition, the environment on which the migration must be performed must be configured as follows:

- ▶ Dynamic Host Configuration Protocol (DHCP) addressing must be available on the subnets hosting the source and destination computers.
- ▶ Device systems must be in the same broadcast domain as the ADS PXE and DHCP services or IP helpers must be configured on router interfaces where these broadcasts are blocked

- ▶ Device systems must be in the same multicast domain as the ADS Image Deployment Service
- ▶ All device systems to be imaged, deployed, or managed must have logical, bidirectional connectivity to the computers hosting the ADS services
- ▶ If the ADS services are installed on separate servers, they must communicate with one another over the adapter they use to access the device systems.

If there are existing PXE services in your environment, such as those that support Remote Installation Services (RIS) or stateless devices, these services must be isolated from the ADS implementation.

Communication ports required by ADS must be able to traverse the network between ADS servers and managed systems. For details on the ports, protocols, and services required see:

http://www.microsoft.com/resources/documentation/WindowsServ/2003/all/ADS/en-us/ads_device_and_service_ports.asp

You can use an existing public key infrastructure (PKI) to issue the certificates necessary for system and device authentication. If one does not exist, ADS can issue certificates for this purpose.

System capacity requirements

Sizing physical computers that will host virtual machines is critical to successful deployment.

Before performing a migration, you should verify that the destination computer (for example, the one running Virtual Server) has adequate processor, memory, and hard disk resources to host the migrated virtual machines. We recommend that you add a reasonable buffer of at least 10% to your estimated resource requirements, or make sure that there is a simple and quick way to extend these resources if the need arises.

We also recommend allowing at least 384 MB RAM for the host operating system. In addition, make sure that the disk space you have allocated to store ADS images is sufficient for the images you plan to migrate.

Source computer requirements

You can use VSMT to migrate source computers running the following operating systems to virtual machines in Virtual Server:

- ▶ Windows NT 4.0 Server SP6a, Standard and Enterprise Editions
- ▶ Windows 2000 Server SP4 or later
- ▶ Windows 2000 Advanced Server SP4 or later
- ▶ Windows Server 2003, Standard Edition and Enterprise Edition

In addition, source computers must meet the following requirements:

- ▶ Windows Management Instrumentation (WMI) must be installed and functional. Unless you have done this already, you must install WMI on Windows NT 4.0 Server. It is available from:
<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=C174CFB1-EF67-471D-9277-4C2B1014A31E>
- ▶ The server must have a minimum of 96 MB of physical memory. If it has any file allocation table (FAT) partitions, then 160 MB is required. We recommend that you convert FAT partitions to NTFS prior to migration.
- ▶ The server must have a network adapter that is enabled for PXE version 0.99c or above.

Other technology requirements

VSMT has the following requirements for specific technologies:

► Hardware devices

Certain hardware devices cannot be represented in or accessed from the virtual environment because Virtual Server does not emulate them. The following devices are examples of those Virtual Server does not emulate:

- USB devices other than keyboard and mouse
- IEEE 1394 (High Performance Serial Bus) devices
- Non-Ethernet network adapters
- Devices that support modem, fax, and TTY functionality
- Specialized add-in boards for video, audio, and co-processing
- SCSI controllers
- Backup devices

► Dynamic physical disks

You cannot migrate dynamic physical disks. GatherHW.exe will gather information about any dynamic disks, but VMScript.exe will skip any disk that is formatted as a dynamic disk, and it will not be migrated.

There are two options for handling dynamic physical disks. You can convert a dynamic disk back to a basic disk and then perform the migration; however, once you do this, you cannot change the dynamic volume back to a partition. Alternatively, you can exclude the dynamic physical disk when running VMScript.exe. You can use the -excludeDrives option to specify the drive to exclude.

► Storage area network (SAN) connections

SAN connections are not migrated. You must reconfigure these connections in the guest operating system after the migration. The SAN is presented as a logical unit number (LUN) to the host operating system and mounted to a drive letter. You can give the virtual machine access to the SAN with a linked virtual hard disk.

► Extended partitions

VSMT supports extended partitions, but it converts them to primary (basic) partitions. The maximum number of partitions supported by primary partitions is four. Therefore, you can only migrate up to four partitions per disk in any combination of primary partitions and logical drives.

Figure 4-5 provides some examples of how VSMT migrates partitions.

Table 4-5 How VSMT migrates partitions

Source configuration	Migrated configuration
4 primary partitions	4 primary partitions
1 primary partition and 1 extended partition containing 1 logical drive	2 primary partitions
2 primary partitions and 1 extended partitions containing 2 logical drives	4 primary partitions
1 primary partition and 1 extended partition containing 4 logical drives	Cannot be migrated because the number of partitions in the migrated configuration would exceed 4

Security accounts

Before performing a migration, we recommend creating appropriate security accounts for administering Virtual Server and each virtual machine. You may want to create security groups for the following roles:

- ▶ Virtual Server administrators

This group has full control over Virtual Server and virtual machine configuration. You can also create administrator groups with a subset of permissions. You can configure permissions for these groups by using the Virtual Server Administration Web site.

- ▶ Virtual machine administrators

This group has the permissions necessary for configuring a specified virtual machine or set of virtual machines, and it can control the virtual machine. You can grant permissions to this group on the file system that hosts the virtual machine configuration and resource files.

- ▶ Virtual machine operators

This group can control a specified virtual machine or set of virtual machines. You can grant permissions to this group on the file system that hosts the virtual machine configuration and resource files.

In addition, if you want to configure a virtual machine to automatically turn on when the Virtual Server service starts, or if you want to run scripts for a virtual machine, you must specify a user account under which to run the virtual machine. We recommend creating a special user account with a low level of privileges for this purpose.

4.5.2 Using VSMT for image capture and deployment

This section provides technical information about image capture and deployment using VSMT. To create a successful image of a physical system and deploy it in a virtual machine, you must take the following steps:

1. Prepare the source servers
2. Generate the migration scripts
3. Capture the source image
4. Create the virtual machine
5. Deploy the image - DeployVM.cmd
6. Complete the deployment
7. Configure the virtual machines

Preparing the source servers

Before initiating the scripts, determine whether your source server is configured to allow PXE booting. On newer servers, this may require reviewing the system BIOS settings. On older computers, this may require verifying manufacturer specifications for the network adapters. The minimum support required is PXE 0.99c.

A Remote Boot Floppy Generator (RBFPG) or RIS (Remote Installation Services) startup disk may be required. For example, it may be needed for computers that do not support PXE as an automatic boot selection or the option to press F12 (a functionality available via BIOS) to select a device to boot from at boot time.

VSMT includes a virtual floppy disk named Ris2003.vfd. This file can be written to a physical 1.44 MB diskette with a diskette image creation and extraction tool (such as WinImage available from <http://www.winimage.com>).

You can also create the PXE boot diskette from a Windows Server 2003 installation by running **RBF.G.EXE**, which is located in the `system32\dlcache` folder.

If you have multiple network cards in the source server that support PXE, enable only the first network card. Enabling multiple network cards may prevent the server from connecting to the ADS Network Boot Service, in which case the capture process will fail.

Generating the migration scripts

Generating the migration scripts involves three tasks:

- ▶ Executing **GatherHW.exe** on the source computer
This generates an XML file of configuration information. This information can be examined by **VMScript.exe** for issues that could prevent migration. **VMScript.exe** also uses this information to generate the scripts that carry out the migration process.
- ▶ Executing **VMScript.exe**
This task helps identify any configuration issues that could prevent a successful migration.
- ▶ Executing **VMScript.exe** to generate the actual migration scripts
From the XML file that **GatherHW.exe** generated, **VMScript.exe** creates the CMD scripts and XML files that perform the remainder of the migration.

Further details about these tasks are discussed in the sections that follow.

Gathering configuration information

This step uses the **GatherHW.exe** command. **GatherHW.exe** uses WMI to gather information. If you are migrating Windows NT 4.0 Server 6a, you must install the WMI Core on the source computer, if it has not already been installed. You may download WMI from this Web site:

<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=C174CFB1-EF67-471D-9277-4C2B1014A31E>

Before running **GatherHW.exe** on Windows NT 4.0 Server, you must first run Disk Administrator on the source computer. Disk Administrator is available from the Administrative Tools menu.

You must specify full paths for **GatherHW.exe**, as follows:

- ▶ Local paths, such as `drive:\folder\subfolder`. Unlike **VMScript.exe**, you can use environment variables with **GatherHW.exe**
- ▶ Universal Naming Convention (UNC) names in the following form:
`\\full computer name\folder\subfolder`
`\\IP address\folder\subfolder`

Run **GatherHW.exe** as follows:

1. Log on to the source server with an account that is a member of the local Administrators group.
2. Copy **GatherHW.exe** to the source server.
3. From the local system, execute **GatherHW.exe** by using the following syntax:
`GatherHW.exe /f: path\source_configuration_file.xml`

The `/f:` parameter lets you specify a path and file name for the generated XML file. We recommend that you use the name of the source server. The XML file will be used in the generation of the scripts.

Validating configuration information

This step uses the VMscript.exe command.

1. Log on to the server running the ADS Controller with an account that is a member of the local Administrators group.
2. Copy the XML file that was generated when you ran GatherHW.exe to the ADS Controller.
3. From the VSMT installation folder (by default, this is %systemdrive%\Program Files\Microsoft VSMT), execute the following command:

```
VMScript -hwvalidate -hwinfile:source_configuration.xml
```

After executing this command, you should receive an output file. Examine any incompatibilities listed in the output. If you cannot correct the incompatibilities, it may not be possible to migrate this server.

Generating scripts

This step uses the VMscript.exe command.

Log on to the server running the ADS Controller with an account that is a member of the local Administrators group, and execute the command in Example 4-1 as a single command on one line (we have divided into multiple lines for readability).

Example 4-1 VMscript.exe command

VMScript

```
-hwgeneratep2v  
-hwinfile:source_configuration_file.xml  
-name:virtual machine name  
-hwtaskseqpath:full path to VSMT installation folder\p2v  
-hwpatchdir:full path to VSMT installation folder\patches  
-vmconfigpath:full path to the .vmc file  
-virtualDiskPath:full path to the .vhd file  
-hwdestvs:FQDN of Virtual Server host
```

Important: The path that you specify for **vmconfigpath** or **virtualDiskPath** must already exist on the server that is running Virtual Server before you run the generated scripts for migration. If not, the operation to create the virtual machine will fail.

VSMT Setup creates a VirtualMachines folder under the VSMT installation folder that you can use for this purpose (by default, %systemdrive%\Program Files\Microsoft VSMT).

Paths that you specify when using VMScript must be fully qualified. In other words, they cannot contain environment variables. This is because these paths are also used in the generated scripts, and paths are resolved at script runtime. You should specify paths as local paths in the following form: drive:\folder\subfolder.

The following optional switches may also be useful:

► `-virtualdiskdynamic`

By default, VHD (.vhd) files created in the migration process are of a fixed size. Using this switch creates them as dynamically expanding VHDs. These VHDs are the size of the actual data and can grow to the size of the original physical disk.

► `-adminmac:MAC-address`

If the source computer has more than one network adapter, you can use this switch to specify which media access control (MAC) address to use for PXE booting the device into the ADS Deployment Agent.

► `-vmmemory:size` (in MB)

By default, VSMT allocates an amount of memory to the virtual machine that is equal to the amount of physical memory on the source computer. This switch allows you to manually allocate memory to the virtual machine during the migration process. Enter an amount in MB.

An output file is generated giving details of success or failure.

The `virtual_machine_P2V_readme.txt` file contains information about the network configuration of the source computer. This file is copied to the virtual machine and displays upon the first logon following migration as a reminder that you may need to take additional configuration steps for the virtual machine.

If `vmscript.exe` is successful, three scripts are created:

- `virtual_machine_CaptureVM.cmd`
- `CreateVM.cmd`
- `DeployVM.cmd`

Capturing the source image

This step uses the `virtual_machine_capture.cmd` script.

Important: If the ADS Controller and Image Deployment services are not on the same computer, and the startup drive uses a FAT partition, the last actions taken by the script (mounting of the captured boot image and `boot.ini` fix-ups) will fail.

You can work around this issue and ignore the error by removing the System, Hidden, and Read-only attributes from `boot.ini` on the source server before you run the capture scripts. This can be accomplished by executing the following at a command line:

```
attrib -s -h -r %systemdrive%\boot.ini
```

The following steps and processes are involved in capturing the image of the source server:

1. Log on to the server running the ADS Controller with an account that is a member of the local Administrators group.
2. From the folder containing the generated scripts, execute the script
`virtual_machine_CaptureVM.cmd`
where `virtual_machine` is the name of the virtual machine that you specified for the virtual machine in “Generating scripts” on page 65.
3. When prompted, restart the source computer and allow it to PXE boot.

Note: If the device is not PXE enabled, try inserting the RIS floppy disk into the source machine at this time. In this case, the RIS floppy disk will use one of the available network adapters for booting. If this is different than the network adapter that VSMT selected, then the process will not continue. You must generate new scripts by using the -adminMac option, specifying the MAC address that the RIS floppy disk should use.

4. In the ADS management console, you will see a new device object representing the source server. ADS takes control of this device. The sequence necessary for the capture is created in the form of a job to run the appropriate virtual_machine_CaptureVM.xml template.
5. The source computer boots to the ADS Deployment Agent, and the capture begins.

The progress for each disk capture is displayed in the ADS job that is running the Capture Image task sequence. Capture occurs for each directly attached logical disk. ADS compresses the image file created for each disk, although the XML definition of the computer retains information about the original disk size.

A typical startup or system disk with 1.5 GB of operating system data will take 10-15 minutes to capture, assuming a 100 Mbps network and no other undue load on either system. Disks with more data will extend the time required for this capture phase.
6. When the capture is complete, ADS initiates a task to shut down the source server, remove control, and delete its device object.
7. The command script then attempts to mount the image as a volume on the ADS Image Deployment server to perform the configuration work defined in the boot.ini file. When this is complete, the script exits.

Note: While the script is running, you can monitor progress from ADS Management as well as from the Event Viewer in the Virtual Server Administration Web site. Errors might include network issues during the capture process or a driver missing for a particular SCSI controller. The latter will be reported as a bad device path.

Creating the virtual machine

This step uses the CreateVM.cmd script.

To create the virtual machine, log on to the server running the ADS Controller with an account that is a member of the local Administrators group, and execute the command

```
virtual_machine_CreateVM.cmd
```

Virtual_machine is the name of the virtual machine that you specified in “Generating scripts” on page 65.

Important: This step can fail if there is not enough physical memory or disk space on the destination computer to allocate to the virtual machine.

By default, the physical memory allocated to the virtual machine matches what was present on the source server. If necessary, you can change the memory allocation as described in “Generating the migration scripts” on page 64.

In addition, the amount of disk space available must equal the space required for the .vhd file plus at least a 10% buffer. In any case, be sure there is sufficient system capacity on the destination computer before you perform the migration.

This part of the process uses most of the information collected by GatherHW.exe to create a virtual representation of the source server. Executing `virtual_machine_CreateVM.cmd` calls the `virtual_machine_CommonInit.cmd` and `VSMT_initenv.cmd` files to set environment variables that are passed to the `VMClient.exe` commands to create the virtual machine. (Note that these environment variables only exist in the open command session. They no longer exist once the window is closed.)

By default, the virtual machine is created with the same amount of RAM, same number of network adapters (up to four) with the same MAC addresses, and the same disks (partitions, sizes, and so forth) as the source server unless you specify different settings as described in “Generating the migration scripts” on page 64.

In addition, the virtual machine is allocated 100% of a single CPU, less the amount that the host operating system is using and the amount allocated to other virtual machines.

The script displays the actions being performed on Virtual Server by `VMClient.exe`. You can also monitor the creation of the new virtual machine and assignment of properties from the Event Viewer page of the Virtual Server Administration Web site and ADS Management.

The `createvm.cmd` script performs the following tasks:

1. Creates a virtual machine configuration (.vmc) file and adds it to the Virtual Server configuration
2. Creates virtual networks on Virtual Server
3. Creates SCSI controllers (if necessary)
4. Creates VHDs that represent the source computer logical disks

Note: If fixed-size virtual hard disks are being created, a .vhd file equal to the size of each original disk is created and zeroed. This can take some time to complete because it happens in sequence.

5. In Virtual Server, adds the virtual hard disks to the virtual machine
6. In Virtual Server, attaches RIS2003.vfd to the virtual floppy drive for the virtual machine
7. In Virtual Server, attaches the virtual network adapter to the virtual network specified during creation
8. In ADS Management, creates a new ADS device object for the virtual machine and populates it with the necessary variables to perform the image deployment and post-imaging fix-ups or tasks
9. In ADS Management, takes control of the device object and assigns the job that corresponds to the imaging and post-imaging tasks
10. Opens the Virtual Server Administration Web site, if it is not already open

Deploying the image - DeployVM.cmd

To deploy the image, log on to the server running the ADS Controller with an account that is a member of the local Administrators group, and execute the script

```
virtual_machine_DeployVM.cmd
```

where `virtual_machine` is the name of the virtual machine that was created. This is the name specified in the `-name:virtual machine` argument of `vmscript.exe`.

Attention: This step may fail if there are not enough physical memory resources on the destination server to start the virtual machine. Before you begin, make sure that there is sufficient memory available.

The virtual machine created in the previous step turns on. Because the RIS2003 floppy disk is attached to the virtual floppy drive, the virtual machine PXE boots into the ADS Deployment Agent.

Image deployment to each virtual hard disk occurs in sequence. As with the capture process, the length of time required depends on the amount of data being moved and the number of virtual hard disks being imaged. Generally, image deployment to the virtual machine takes longer than image capture from the physical source computer. This can be due to the overhead of the virtualized environment, but virtual hard disk type is also an important factor. Dynamically expanding virtual hard disks take more time.

At the end of the imaging sequences, operating system patches and other configuration settings are applied. The script exits when this is complete.

The ADS device is left in a controlled state and the virtual machine is left in the ADS Deployment Agent by default. To specify the state, you can use the option `-postDeployAction` to leave the device in one of the following states:

- ▶ Leave the VM in the Deployment Agent: `-postDeployAction:0`
- ▶ Shutdown the VM: `-postDeployAction:1`
- ▶ Reboot the VM: `-postDeployAction:2`

Completing the deployment

To complete the deployment, perform the following tasks:

1. In the Virtual Server Event Viewer and the ADS History log, verify that there are no errors.
2. Send the `ResetVM` job to the virtual machine.
3. In ADS Management, remove control from the device.
4. Allow the new system to start up. The system should start automatically once drivers are installed and other configuration steps are performed.
5. Log on to the guest operating system with an account that is a member of the local Administrators group, access the CD image, and run setup for Virtual Machine Additions.

In the Virtual Server Administration Web site, navigate to the CD/DVD page for the virtual machine, and capture `VMAdditions.iso` from Known image files. See 1.6.7, “Install Virtual Machine Additions” on page 18 for more information.
6. Take any additional steps required to configure the virtual machine so that it runs properly, as described in the next heading.
7. After testing the virtual machine to make sure that it is operating correctly, remove the captured image from the ADS image store.

4.5.3 Configuring virtual machines

This section describes the steps to take after a migration to configure the virtual machines so that they run properly.

Allocating system resources to virtual machines

You can allocate physical memory and CPU resources to each virtual machine running on Virtual Server. By default, when VSMT migrates a server, it gives the migrated system on the

destination computer the same amount of memory that existed on the physical source computer. It also allocates 100% of the available processor resources on the destination computer. Available resources include any processor capacity not currently being used by the host operating system or other virtual machines.

When you migrate multiple source computers to a single destination computer, there may not be enough memory or processor capacity on the destination computer to support their combined requirements. Symptoms may include an inability to turn on the virtual machine or boot the virtual machine to the ADS Deployment Agent. In this case, you will need to change the resource allocations.

You can change the amount of memory allocated to a virtual machine by editing the **-vmmemory** parameter when running VMScript, as described in “Generating scripts” on page 65. You can also change the memory or CPU capacity allocated to a virtual machine following migration from within Virtual Server.

Installing and registering device drivers

When a server is migrated from a physical computer to a virtual machine, certain drivers must be installed because of differences between the physical hardware on the source computer and the hardware emulated by Virtual Server.

For example, for the guest operating system to start, certain drivers must be present. These drivers include the hardware abstraction layer (HAL) and kernel that is compatible with the destination device for the hardware emulated by Virtual Server and drivers for the storage controllers used in the guest operating system. If these drivers are not present, the operating system cannot start, and you will not be able to install any additional drivers.

After deployment, VSMT automatically installs drivers. VSMT setup installs the binary files that are needed to install drivers for the operating systems listed in Figure 4-6.

These drivers are organized in a tree under the patch folder based on the build number and service pack number. By default, the patch folder is c:\Program Files\Microsoft VSMT\Patches. Build numbers (and folder names) for each supported operating system are listed in Table 4-6.

Table 4-6 Driver levels for each supported operating system (folder names)

Driver build number	Operating system
4.0.1381	Windows NT 4.0 Server SP6a
5.0.2195	Windows 2000 Server, SP2, SP3, and SP4, all editions
5.2.3790	Windows Server 2003, all editions

In the same folder are several task sequences for uploading the driver binary and setting appropriate registry settings to the destination device. These task sequences are:

- ▶ HAL.XML: Fixes the HAL on the guest operating system.
- ▶ STORAGE.XML: Adds support for ATAPI and PCI IDE storage controllers.
- ▶ NETWORK.XML: Uploads the network configuration script.
- ▶ OTHER1.XML, OTHER2.XML, ..., OTHER_n.XML: User-specified task sequences to be executed on the guest operating system.
- ▶ FINISH.XML: Shuts down the guest operating system.

During the script generation phase, VSMT examines the patch folder for the existence of the appropriate task sequences. If, for example, you are migrating a computer running Windows

2000 Server SP4 (corresponding to driver version 5.0.2195), then VSMT searches (by default) the folder c:\Program Files\Microsoft VSMT\Patches\5.0.2195\sp4 for the appropriate task sequence and includes it in the script that deploys the virtual machine.

If you want to change the installed drivers, you can edit the location from which VSMT picks up the task sequence and make the drivers available in that location. To edit the location, you specify an alternative folder for the VMScript -patch option. This folder should have the same structure and the same task sequences. You can update the binary files with the versions you want to use. You should not need to modify any task sequence.

If VSMT does not find the appropriate task sequence, the process cannot continue.

The rest of the drivers needed for the hardware device to function, such as those for networking, must be installed once the device has been booted in the guest operating system. In Windows 2000 Server and Windows Server 2003, Plug-and-Play functionality installs new hardware devices as long as their drivers are present. In Windows NT 4.0 Server there is no Plug-and-Play functionality, and you must manually install new drivers. For instructions, see “Configuring Windows NT 4.0 Server” on page 71.

Configuring IP addresses

If you migrate an application that is dependent on a specific IP address, and the guest operating system to which it is migrated has been assigned a different IP address, the application will not function. ADS is an example of an application that has this type of dependency.

If you encounter this issue, you can either change the IP address used by the application, or change the IP address assigned to the guest operating system.

Changing virtual machine startup options

By default, virtual machines do not turn on when Virtual Server starts. You can change this setting on the General Properties Web page for the machine at the Virtual Server Administration Web site. For this feature to work, you must specify a user account for running the virtual machine. We strongly recommend a special account with a low level of privileges.

Configuring Windows NT 4.0 Server

Because Windows NT 4.0 Server lacks Plug-and-Play functionality, some manual steps are required to install hardware that corresponds to new virtualized devices. These steps are straightforward but must be performed from within the migrated operating system of the virtual machine:

1. Configure the network
 - a. From the Virtual Server Administration Web site, on the Hard disks Web page for the virtual machine, capture NT4 Network Driver.vfd from Known floppy disks.

Note: Virtual Server allows only one virtual machine to capture the same virtual floppy disk at the same time. It will allow you to capture the virtual floppy disk even if another virtual machine has it, but it will not function.

- b. Log on to the guest operating system with an account that is a member of the local Administrators group.
 - c. On the Adapters tab of the Network Control Panel, delete the network adapters.
 - d. Add a new adapter, select **Have Disk**, select **A:**, and install the DEC 21140 adapter.

- e. Remove NT4 Network Driver.vfd from the virtual machine, and restart the guest operating system.

Note: If the source server had multiple network adapters, in the virtual machine, they will all be assigned to the same virtual network, and the guest operating system will report a conflict in the network name. Either remove all except one network adapter or assign each to a different physical network adapter on the destination computer.

2. Configure SCSI adapters

- a. Log on to the guest operating system with an account that is a member of the local Administrators group.
- b. Open the SCSI Adapters Control Panel.
- c. In the list of devices, identify those that are no longer used.
- d. On the Drivers tab, delete unused adapters.
- e. If they are not already installed, install Virtual Machine Additions in the guest operating system.
- f. Follow the instructions in “Install the SCSI driver after installing Virtual Machine Additions” in the *Virtual Server 2005 Administrator's Guide*.
- g. Follow the instructions for adding a SCSI adapter in “Add or remove a SCSI adapter” in the *Virtual Server 2005 Administrator's Guide*.
- h. Restart the guest operating system.



Automation

Microsoft Virtual Server 2005 provides support for automation through Microsoft Component Object Model (COM) technology. This means that any scripting language that can connect to COM automation objects will work with Virtual Server.

In this chapter, we discuss the Virtual Server COM application programming interface (API) and present scripting scenarios.

5.1 The Virtual Server 2005 COM API

A full-featured Component Object Model (COM) scripting model ensures that every aspect of Virtual Server functionality can be controlled by scripts. Because the scripting model is based on COM, users are not tied to a specific scripting language. They can choose from Microsoft Visual Basic, C#, Perl, and many other modern development languages. Furthermore, scripts can be triggered by certain events within Virtual Server.

Scripts are executed by Windows Script Host, the standard Windows operating system scripting environment. You can use any script program that Windows Script Host is capable of executing. Windows includes built-in support for Microsoft Visual Basic Scripting Edition (VBScript) and Microsoft JScript development software. You can use other scripting languages as well.

If the script host application is not specified, scripts are executed by **CScript**, the command line version of Windows Script Host. You may also execute scripts with **WScript**, a version of Windows Script Host that provides a graphical user interface, by using the **WScript** command.

Automation does not need to be in the form of scripts, however. You can manage Virtual Server with other programming languages, such as C, C++, and C#.

Virtual Server exposes a rich COM interface that permits the user to monitor and control the virtual machine environment. All of the Virtual Server Web-based user interfaces use this COM interface by way of a scripting language. This allows customization of the virtual machines. The COM application programming interface (API) contains 42 interfaces and hundreds of calls, allowing scripts to control nearly every aspect of the product.

The scripting API makes Virtual Server a powerful platform for virtualization solutions by:

- ▶ Enabling programmatic control over the configuration and administration of virtual machines
- ▶ Automating virtual machine deployment and operations
- ▶ Helping customers integrate virtual machines into their existing IT infrastructure and operations for reduced cost and enhanced manageability

For example, the Visual Basic script in Figure 5-1 save states all running virtual machines on a single host.

```
Option Explicit
Dim vpcApp, vmCollection
Call main()
sub main()
Dim VirtualMachine
Set vpcApp = CreateObject( "VirtualServer.Application","localhost" )
Set vmCollection = vpcApp.VirtualMachines
for each VirtualMachine in vmCollection
    If VirtualMachine.State = 5 Then ' Guest is running
        wScript.Echo "Suspending: " + VirtualMachine.Name
        VirtualMachine.save 'Save state (keeps undo drives if present)
    End If
next
end sub
```

Figure 5-1 sample script myscript.vbs

To run this script with CScript, the command-line version of Windows Script Host, the following command can be used: **cscript myscript.vbs**

If a script is attached to an event that is associated with a virtual machine, the quoted name of the virtual machine will be passed to the script as the first parameter.

Restriction: Scripts are executed in a new process in the logon session of the network service account, under which the Virtual Server service is running. As a result, any user interface that the script may try to present will not be visible on the local computer. In general, displaying UI from these scripts should not be attempted.

5.2 Scripting scenarios

Scripting Virtual Server can be useful in many scenarios. This section lists a few of them.

5.2.1 Backup

A sample script for nightly backups of virtual machines could automate the following steps:

1. Shutdown the virtual machines.
2. Start backup application to backup the .vhd files.
3. Wait for backup completion.
4. Write logfile entries.
5. Restart the virtual machines.

5.2.2 Load balancing

Simple load balancing can be implemented on a server farm with shared SAN storage by scripting the following steps:

1. Monitor server load (CPU, network utilization).
2. If load on one server reaches a defined limit
 - a. Shut down selected virtual machines on that server.
 - b. Select target server based on lowest load.
 - c. Restart the virtual machines on that server.

Using these scripts, automatic load balancing can be achieved with only brief downtime for the virtual machine shutdown and restart.

5.2.3 Disaster recovery

Scripts can be used for disaster recovery for a whole Virtual Server:

1. Monitor Virtual Server availability.
2. If the Virtual Server goes down:
 - a. Turn off that server.
 - b. Restart the virtual machines on a backup/standby server or distribute them to other servers in the server farm.
 - c. Create alert to management console or to inform administrator.

Similar scripts can be used to monitor single virtual machines on a server and restart them in there are problems.

5.2.4 Automation of repeating tasks

Test and development environments often need repeating tasks like setting up machines for installation and application tests. Here scripts can be used for automation such as:

1. Define new undo or differencing disk for the test virtual machine.
2. Start the virtual machine.
3. Wait for completion of the test.
4. Discard the undo/differencing to return the test virtual machine to a clean state.

For more information about scripting and the Virtual Server object model, see the *Virtual Server 2005 Programmer's Guide*.

Related publications

The publications listed in here are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks

The following IBM Redbook is relevant to this redpaper:

- *Implementing Systems Management Solutions using IBM Director*, SG24-6188

Product publications

These publications are relevant to this redpaper:

- *Virtual Server 2005 Product Overview*, available from:
<http://www.microsoft.com/windowsserversystem/virtualserver/overview/vs2005prod.msp>
- *Virtual Server 2005 Technical Overview*, available from:
<http://www.microsoft.com/windowsserversystem/virtualserver/overview/vs2005tech.msp>
- *Virtual Machine Manager Installation and User's Guide*, available from:
http://www.ibm.com/servers/eserver/xseries/systems_management/xseries_sm/vmm.html
- *Microsoft Windows Corporate Deployment Tools User's Guide*, an online help file, in the deploy.cab file in the Support folder of the Windows Server 2003 installation CD-ROM

The following guides are available as part of the Virtual Server installation:

Virtual Server 2005 User's Guide
Virtual Server 2005 Administrator's Guide
Virtual Server 2005 Programmer's Guide

Referenced URLs

The following URLs were referenced in this redpaper:

IBM Web pages:

- IBM ServerProven®: BladeCenter servers that support Windows operating systems:
<http://www.pc.ibm.com/us/compat/nos/microsoft.html>
- IBM ServerProven: xSeries servers that support Windows operating systems:
<http://www.pc.ibm.com/us/compat/nos/microsoft.html>
- Virtual Machine Manager home page:
http://www.ibm.com/servers/eserver/xseries/systems_management/xseries_sm/vmm.html
- IBM Director home page:
http://www.ibm.com/servers/eserver/xseries/systems_management/director_4.html

- ▶ IBM Director 4.2 announcement:

<http://www.ibm.com/common/ssi/fcgi-bin/ssialias?infotype=an&subtype=ca&appname=redbooks&htmlfid=897/ENUS104-281>

Microsoft Web pages:

- ▶ Enterprise Memory Architecture overview:

http://www.microsoft.com/resources/documentation/windowsserv/2003/enterprise/proddocs/en-us/vlm_1.asp

- ▶ Knowledge base entry, Cluster Servers May Experience Connection Timeouts to Drives After You Install or Upgrade to the Windows Clustering Feature in Windows Server 2003:

<http://support.microsoft.com/?kbid=818877>

- ▶ Automated Deployment Services technical overview:

<http://www.microsoft.com/windowsserver2003/techinfo/overview/adsbenefits.mspx>

- ▶ Windows Management Instrumentation for Windows NT 4.0:

<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=C174CFB1-EF67-471D-9277-4C2B1014A31E>

- ▶ Virtual Server Migration Toolkit Beta page:

<http://www.microsoft.com/windowsserversystem/virtualserver/evaluation/vsmtbeta.mspx>

- ▶ Downloads for Microsoft Virtual Server 2005:

<http://www.microsoft.com/windowsserversystem/virtualserver/downloads> –

Other Web sites

- ▶ Diskette imaging tool:

<http://www.winimage.com>

Online resources

These Web sites are also relevant as further information sources:

- ▶ Microsoft Virtual Server 2005 home page

<http://www.microsoft.com/windowsserversystem/virtualserver>

- ▶ IBM Director home page

http://www.ibm.com/servers/eserver/xseries/systems_management/director_4.html

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Abbreviations and acronyms

ACL	access control list	OS	operating system
ADS	Automated Deployment Services	PAE	Physical Address Extension
AMD	Advanced Micro Devices, Inc	PC	personal computer
API	application programming interface	PCI	Peripheral Component Interconnect
ATAPI	ATA Packet Interface	PKI	public key infrastructure
ATS	Advanced Technical Support	PXE	Preboot eXecution Environment
BIOS	basic input/output system	RAID	redundant array of independent disks
BSD	Berkeley software distribution	RAM	random access memory
CD-ROM	compact disk-read only memory	RBFG	Remote Boot Floppy Generator
CMD	command	RDP	Remote Desktop Protocol
COM	Component Object Model	RIS	Remote Installation Services
CPU	central processing unit	SAN	Storage Area Network
DHCP	Dynamic Host Configuration Protocol	SCSI	Small Computer System Interface
DNS	Domain Name System	SMS	System Management Server
EMEA	Europe/Middle East/Africa	SQL	structured query language
FAT	file allocation table	SSL	Secure Sockets Layer
GB	gigabyte	TCP/IP	Transmission Control Protocol/Internet Protocol
HAL	hardware abstraction layer	TTY	teletypewriter
I/O	input/output	UNC	Universal Naming Convention
IBM	International Business Machines Corporation	UPS	uninterruptible power supply
IDE	integrated drive electronics	URL	Uniform Resource Locator
IEEE	Institute of Electrical and Electronics Engineers	USB	universal serial bus
IIS	Internet Information Services	VDM	Virtual Disk Manager
IP	internet protocol	VGA	video graphics array
IPX	Internetwork Packet eXchange	VHD	virtual hard disks
ISV	independent software vendor	VM	virtual machine
ITSO	International Technical Support Organization	VMC	virtual machine configuration
LUN	logical unit number	VMM	Virtual Machine Manager
MAC	media access control	VMRC	Virtual Machine Remote Control
MB	megabyte	VS	Virtual Server
MCF	MOM Connector Framework	VSMT	Virtual Server Migration Toolkit
MMC	Microsoft Management Console	WMI	Windows Management Instrumentation
MOM	Microsoft Operations Manager	WSRM	Windows System Resource Manager
NAS	Network Attached Storage	XML	eXtensible Markup Language
NDIS	network driver interface specification		
NIC	network interface controller		
NLB	Network Load Balancing		
NTFS	NT File System		
NTLM	NT LAN Manager		



Introducing Microsoft Virtual Server 2005 on IBM *@*server xSeries Servers



Describes the major features of this new product

Explains the architecture and best practices

Introduces the management tools that are available

Businesses continually seek ways to reduce cost and risk while increasing quality and agility in their IT infrastructure. Virtualization is a key enabling technology that businesses can use to achieve these business benefits. With virtualization technology, clients run multiple operating systems concurrently on a single physical server, improving hardware efficiency, reducing IT costs, and increasing administrator productivity.

Microsoft Virtual Server 2005 hosted on Windows Server 2003 and xSeries servers uses virtualization technology to deliver the performance necessary for completing time-saving and cost-saving tasks. It provides businesses with an enterprise-ready computing environment with advanced levels of scalability, manageability, and reliability.

This IBM Redpaper introduces Virtual Server and describes its main features and functions, architecture, and typical uses. It also introduces the management tools available for Virtual Server, including IBM Director with Virtual Machine Manager and Microsoft Operations Manager.

This IBM Redpaper is intended for IT specialists who would like to learn about this new product and how it can be used in their environment.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks