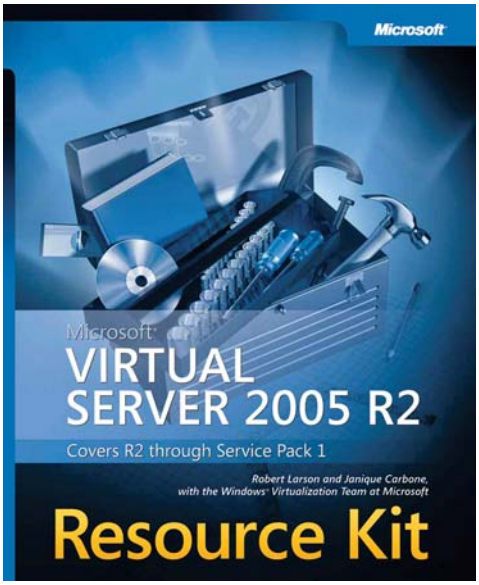

Microsoft® Virtual Server 2005 R2 Resource Kit

Robert Larson and Janique Carbone with the Windows Virtualization team

To learn more about this book, visit Microsoft Learning at http://www.microsoft.com/MSPress/books/10622.aspx

9780735623811
Publication Date: August 2007

Microsoft Press

Table of Contents

Chapter 5

Virtual Server 2005 R2 Advanced Features

In this chapter:

This chapter describes advanced features in Microsoft Virtual Server 2005 Release 2 (R2). You will learn about virtual hard disk, network, and clustering options that you can use to deploy broad virtualization infrastructure solutions. Technical descriptions and configurations are discussed along with common usage scenarios.

Using Virtual Hard Disk Advanced Features

Virtual Server 2005 R2 uses the virtual hard disk (VHD) format to encapsulate virtual machine data into one or more files that are equivalent to physical drives associated with a traditional server. Using the VHD format as a basic building block, Virtual Server 2005 R2 provides advanced virtual hard disk features that enable the creation of virtualized environments that are more functional and flexible than physical equivalents, particularly for disciplines such as development, testing, training, and support. Table 5-1 lists the advanced virtual hard disk features covered in this section.

Table 5-1 Virtual Hard Disk Advanced Features

Feature	Description
Differencing disks	A special type of dynamically expanding virtual hard disk that stores virtual machine data changes while isolating them from the base virtual hard disk.
Undo disks	A special type of dynamically expanding virtual hard disk that stores virtual machine data changes while isolating them from the base virtual hard disk. There are similarities with differencing disks, but differences in options and applicable scenarios.

Table 5-1 Virtual Hard Disk Advanced Features

Feature	Description
Linked disks	A special type of virtual hard disk designed specifically to convert a physical hard disk into a virtual hard disk file. The process associated with the use of linked disks is potentially time consuming depending on the size of the physical disk.
VHDMount command-line tool	This is a new feature provided with Virtual Server 2005 R2 SP1. VHDMount is an essential tool to manipulate virtual hard disk files without booting into a virtual machine.
VHD compaction	This tool is used to regain unused space within a virtual hard disk. The compaction process works only for dynamically expanding virtual hard disks. No other type of virtual hard disk can be compacted.

Differencing Disks

A virtual machine running within Virtual Server 2005 R2 has its data encapsulated in one or more base virtual hard disks. When data changes occur to the guest operating system or the applications running in it, modifications are committed to the virtual hard disks. The changes made to the virtual hard disks are permanent, paralleling the process that would occur with a standard physical system. However, a variety of compelling scenarios are enabled by preserving a base virtual hard disk in an unchanged state, while still capturing and storing ongoing virtual machine changes.

A differencing disk is a special type of dynamic disk that stores changes to virtual machine data in a separate file from a base virtual hard disk. The association of the base virtual hard disk to the differencing disk is defined as a parent-child relationship. In this parent-child relationship, each child differencing disk can derive from only one parent disk, but parent disks can be used as the basis to create multiple, distinct child differencing disks.

Figure 5-1 shows that differencing disks can be created in very simple or very complex parent-child hierarchies. A multilevel differencing disk hierarchy is commonly referred to as a *chain* of differencing disks, reflecting that a child differencing disk can have a parent disk that is also a differencing disk. The chain can consist of several levels, but it always stems from either a standard dynamically expanding or fixed-size virtual hard disk at the top of the hierarchy. This concept is important because data changes in a differencing disk are simply represented as modified blocks in relation to the parent disk. Therefore, a differencing disk is never used independently, but in conjunction with all parent disks in its hierarchy. (See Figure 5-1.)

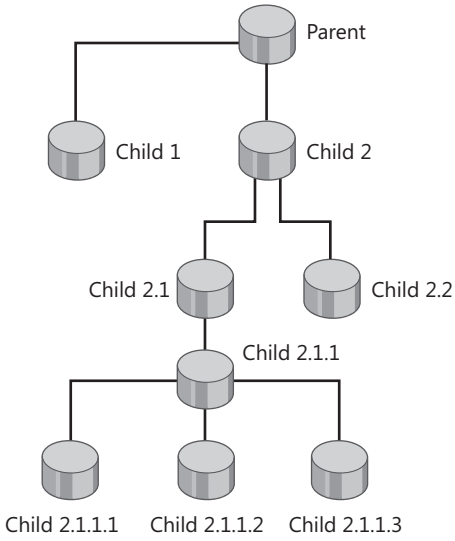

Figure 5-1 Multilevel differencing disk hierarchy

If you examine a Virtual Server 2005 R2 host file system, you will see each differencing disk stored as an individual file. Within the virtual machine file system, you see only a single disk, independent of how many levels of differencing disks are actually associated with a specific virtual hard disk.

Best Practices To quickly and easily identify parent-child differencing disk relationships in a complex chain, adopt a standardized virtual hard disk naming convention. The Virtual Server Administration Website allows you to inspect a differencing disk and discover its parent disk, but it does not report any child differencing disks related to it.

Creating a Differencing Disk

When you create a new differencing disk, the location of the base virtual hard disk that will be the parent for the new differencing disk must be specified. The parent disk can be either a fixed-size or dynamically expanding virtual hard disk. A differencing disk grows as needed, up to the size specified for the parent virtual hard disk.

To create a differencing disk, follow these steps:

1. Open the Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Disks, point to Create and then click Differencing Virtual Hard Disk.
3. In Location, select the folder to store the new virtual hard disk file. If the folder does not appear in the list, type the fully qualified path to the folder as described in the next step.

4. In the Virtual Hard Disk File Name text box, type the fully qualified path to the folder followed by a name for the differencing virtual hard disk. You do not need to include a filename extension.
5. In Known Virtual Hard Disks, select the virtual hard disk file to use as the parent disk.
6. If the virtual hard disk file does not appear in the Known Virtual Hard Disks list, in the Fully Qualified Path To Parent Virtual Hard Disk text box, type the fully qualified path to the parent virtual hard disk file.
7. Click Create.

Note By default, differencing disks use the .vhd file extension, which makes them difficult to distinguish from standard virtual hard disks.

Examining Parent-Child Differencing Disk Relationships

Every dynamic disk contains a standard virtual hard disk header that embeds a specific dynamic disk header. The dynamic disk header format is identical for both standard dynamically expanding and differencing disks. However, several fields in this header are only relevant to differencing disks, as they identify parent disk attributes. A list of the dynamic disk header fields is provided in Table 5-2, with those relating only to differencing disks appearing in boldface type.

Table 5-2 Dynamic Disk Header

Dynamic disk header fields	Description
Cookie	A set field that identifies the header.
Data Offset	Absolute byte offset to next hard disk image structure (*currently unused*).
Table Offset	Absolute byte offset of the block allocation table (BAT) in the file.
Header Version	Dynamic disk header version.
Max Table Entries	Maximum number of entries in the BAT.
Block Size	Size of unit that is used to incrementally expand the dynamic disk.
Checksum	Checksum of the dynamic disk header.
Parent UUID	**128-bit universally unique identifier (UUID) of the parent disk (used only for differencing disks).**
Parent Time Stamp	**Modification time stamp of the parent disk (used only for differencing disks).**
Reserved	Field is set to zero.
Parent Unicode Name	**Unicode string for filename of the parent disk (used only for differencing disks).**

Table 5-2 Dynamic Disk Header

Dynamic disk header fields	Description
Parent Locator Entry 1	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 2	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 3	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 4	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 5	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 6	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 7	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Parent Locator Entry 8	**Platform-specific format containing the absolute byte offset in the file where the parent locator is stored (used only for differencing disks).**
Reserved	Field is set to zero.

A differencing disk uses the parent UUID and Unicode file name information stored in its dynamic disk header to locate and open the parent disk. Because a parent disk can also be a differencing disk, it is possible that the entire hierarchy of parent disks will be opened, up to the base virtual hard disk.

Portability of parent and child differencing disks across server platforms is provided by the Parent Locator entries listed in Table 5-2. Parent locator entries store platform-specific information to locate the parent differencing disk on the physical drive.

Important For the Microsoft Windows platform, both the absolute (for example, c:\parent\parent.vhd) and relative (for example, .\parent\parent.vhd) paths of the parent disk are stored in the Parent Locator entry of a differencing disk. As long as you copy the virtual hard disks to the same relative directory hierarchy on a new host, you will be able to add the virtual machine to Virtual Server and turn it on without having to make any additional changes.

When a virtual machine using differencing disks issues a write operation, the data is written only to the child differencing disk. As part of the process, an internal virtual hard disk data structure is updated to reflect changes that supersede data in the parent disk. During read operations, the same internal virtual hard disk data structure is checked to determine which data to read from the child differencing disk. Unchanged data is read from the parent disk.

Direct from the Source: Configure Parent Disks as "Read-Only"

A child differencing disk stores the parent disk modification time stamp when it is created. Any modifications made to the parent disk after creation of the child differencing disk will be detected and will invalidate the child differencing disk. To ensure that nothing can be written to the parent disk that will corrupt the parent-child disk relationship, configure the parent disk as "read-only."

Bryon Surace
Program Manager, Windows Virtualization

Merging Differencing Disks

Although a differencing disk can be used to permanently store virtual machine data changes, you might need to combine the child differencing disk with the parent disk. Virtual Server 2005 R2 provides two ways to accomplish this. You can either merge the differencing disk into the parent disk or merge the differencing disk and the parent disk into a new virtual hard disk. If you merge a differencing disk into the parent disk, the differencing disk is deleted upon completion of the process and any other differencing disk that pointed to the original parent disk is invalidated. If you need to retain the differencing disk, you should choose to merge the differencing disk and parent disk into a new virtual hard disk. This approach is recommended to lower the risk of data loss. You can verify that the merge operation is successful prior to deleting the original files.

To merge differencing disks, follow these steps:

1. Open the Virtual Server R2 Administration Website.
2. In the navigation pane, under Virtual Disks, click Inspect.
3. In the Inspect Virtual Hard Disk pane, do one of the following, and then click Inspect:
   - ❑ In Known Virtual Hard Disks, select the virtual hard disk that you want to merge.
   - ❑ In the Fully Qualified Path To File text box, type the fully qualified path to the virtual hard disk file that you want to merge.
4. In the Actions pane, click Merge Virtual Hard Disk.
5. Proceed with one of the following two choices:
   - ❑ Select the Merge With Parent Virtual Hard Disk option.

- Select the Merge To New Virtual Hard Disk option, and then select a folder in which to store the new virtual hard disk. If the folder is not listed, type a fully qualified path and filename for the new virtual hard disk. You do not need to include a filename extension.

6. In Merged Virtual Hard Disk Type, select a type for the new virtual hard disk.
7. Click Merge.

Important Prior to merging a differencing disk and parent disk into a new virtual hard disk, make sure there is enough space on the physical disk to perform the operation.

Using Differencing Disks

Functionality gains from using differencing disks become evident when considering a typical support scenario. A support engineer often needs to troubleshoot server configurations for different operating system update levels or with different applications. Using one or more physical test servers, even with preconfigured build images, the setup and testing of multiple server configurations is a lengthy, complex process that results in protracted problem response time. Using Virtual Server 2005 R2 with differencing disks, a support engineer can quickly create a virtual machine for each unique server configuration. Starting with a common parent virtual hard disk that contains the base operating system, each individual server configuration is created as a new virtual machine with one or more differencing disks to capture incremental operating system patches and application stacks.

Important Differencing disks should not be used with cluster configurations.

As shown in Figure 5-2, implementing a virtualized support environment using differencing disks can help significantly reduce the setup and test cycle associated with problem resolution response time. Even with a single physical server constraint, a Virtual Server 2005 R2 host can run multiple virtual machines (VMs) concurrently, allowing parallel testing of distinctive server configurations. In addition to creating an environment that can lead to faster support response time, this solution also has the additional benefit of saving significant amounts of physical disk space for any scenario that requires multiple complex configurations sharing a large common software base.

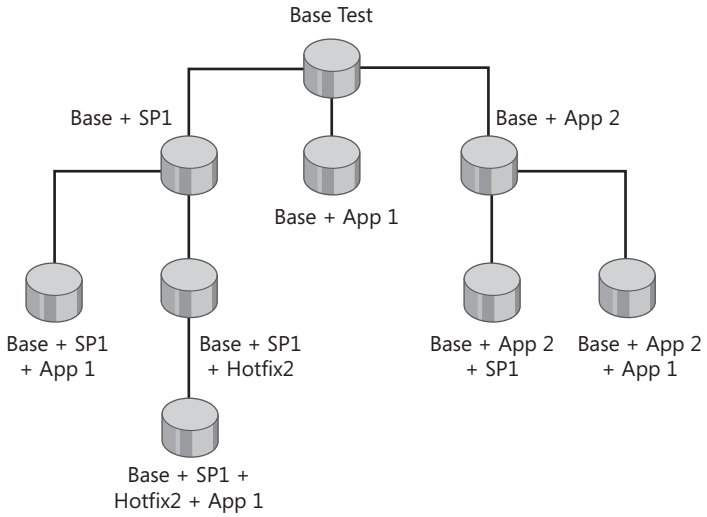

Figure 5-2 Using differencing disks to create guest VMs for concurrent testing

Undo Disks

Undo disks are quite similar to differencing disks. Like a differencing disk, an undo disk is used to isolate virtual machine data changes from a base virtual hard disk. Undo disks also share the special dynamic disk characteristics previously defined for differencing disks. However, in an environment where virtual machine data changes need to be quickly discarded or a rapid rollback to the base virtual machine state is required, undo disks are a better solution than differencing disks. There are other environments that require the use of a shared common software base and rapid rollbacks to a baseline state. In these cases, differencing disks can be used in combination with undo disks to implement the virtualization solution.

Note Unlike a differencing disk, which has a .vhd filename extension, an undo disk uses a .vud filename extension. Also, undo disks are stored in the same directory as the virtual machine configuration file (which uses a .vmc filename extension).

Configuring Undo Disks

One major distinction between differencing disks and undo disks is in the configuration process. A differencing disk is created at an individual virtual hard disk level and usually associated with the creation of a new virtual machine. In contrast, undo disks are either enabled or disabled for an existing virtual machine and created for every virtual hard disk associated with the virtual machine. In other words, you do not have the ability to individually choose the virtual hard disks for which undo disks are generated.

Important If you need to move a virtual machine from one Virtual Server 2005 R2 host to another, don't forget to move parent disks and virtual machine configuration files (.vmc) along with child differencing disks and undo disks.

To configure undo disks for a virtual machine, follow these steps:

1. Open the Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Machines, point to Configure and then click the desired virtual machine.
3. In the Configuration section, select Hard Disks.
4. In the Virtual Hard Disk Properties section, select the Enable Undo Disks check box and then click OK.

Important Undo disks can be enabled or disabled only when a virtual machine is in a powered-off state. The option to enable undo disks is not available if the virtual machine is in a saved state.

Managing Undo Disks

Another major distinction between differencing disks and undo disks is that you are required to decide what to do with the changes saved in undo disks every time a virtual machine is shut down or placed in a saved state. Virtual Server 2005 R2 provides three options to manage undo disks:

- **Keep Undo Disks** This option saves the changes stored in the undo disk and preserves the state of the base virtual hard disk.
- **Commit Undo Disks** This option saves the changes stored in the undo disk to the base virtual hard disk.
- **Discard Undo Disks** This option deletes the undo disk without saving any changes to the base virtual hard disk.

If you shut down the guest operating system from within the virtual machine, undo disks are saved. If you choose to discard undo disk changes, new undo disks are created when the virtual machine is turned back on.

Caution If you disable undo disks while a virtual machine is turned off, the undo disks are immediately deleted.

Using Undo Disks

Undo disks are most useful in scenarios where frequent rollbacks to a base configuration are required. Two mainstream examples are software testing and end-user training. Working in these scenarios with only physical components, one of the most time consuming and tedious tasks is rebuilding the baseline environment—whether it is to re-create the steps to isolate a software bug or to prepare the system for the next user of a training lab. This is even more of a burden if the environment consists of several, incrementally different workloads, although the process can again be somewhat simplified by using imaging tools to more quickly reset each system. A better solution for working in these scenarios is to use Virtual Server 2005 R2 virtual machines that enable undo disks. As illustrated in Figure 5-3, the more complex software testing scenario—which requires multiple, incrementally different virtual machine configurations—is optimized by using undo disks in conjunction with differencing disks. The simple end-user training configuration only requires the implementation of undo disks. At the end of each training session, the system only needs to be reset to the base configuration.

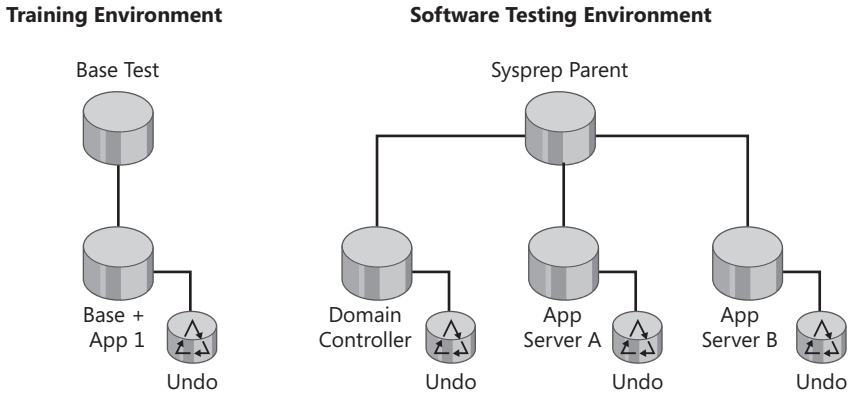

Figure 5-3 Using undo disks with and without differencing disks to achieve a quick rollback

In either case, a rollback to a baseline configuration is accomplished by simply choosing to discard the changes in the undo disks. This process takes just a few moments to complete before the system software is reset to the original configuration.

Best Practices To restrict the ability to commit undo disk changes and ensure the integrity of the virtual machine baseline configuration, you should set the base virtual hard disk files to read-only.

Linked Disks

A linked disk is a virtual hard disk that points to a physical drive with the single purpose of duplicating the contents into a new virtual hard disk. There are several requirements to con-

sider when using this method to migrate the contents of the physical disk into a virtual hard disk. The limitations are as follows:

- A linked disk can be associated only with a physical disk, not a volume.
- A linked disk must be used only to migrate a data disk; startup disks are not supported.
- A linked disk must be stored on a physical disk that is separate from the drive being converted.
- The physical disk must not be accessed by the host operating system or applications during the conversion process.
- If the physical disk that is being converted is larger than 127 GB, you must attach the virtual hard disk into which the disk contents will be copied to a virtual SCSI adapter.

Best Practices Prior to creating the linked disk, you should use the Disk Management Microsoft Management Console (MMC) or other similar tool to remove the drive letter for the target drive. This will make the drive inaccessible to the host operating system, preventing disk corruption during the conversion process.

Creating a Linked Disk

The creation of a linked disk is simple, but it is only the first step in converting a physical disk into a new virtual hard disk. Follow these steps to create a linked disk:

1. Open the Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Disks, point to Create, and then click Linked Virtual Hard Disk.
3. In Location, select the folder in which to store the virtual hard disk file. If the folder you want does not appear in the list, you must type the fully qualified path to the folder in the following step.
4. In the Virtual Hard Disk File Name text box, after the path to the folder, type a name for the virtual hard disk. You do not need to include a filename extension.
5. In the Physical Computer Drive section, select the physical hard disk to which you want to link the virtual hard disk and then click Create.

At this point, you have only created a virtual hard disk that is essentially a pointer to the physical drive.

Using the Linked Disk to Convert the Physical Disk

To complete the process and copy the physical drive content to a new virtual hard disk, follow these additional steps:

1. Open the Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Disks, click Inspect.
3. In the Inspect Virtual Hard Disk pane, select the virtual hard disk to convert in Known Virtual Hard Disks. If the file does not appear in the list, in the Fully Qualified Path To File text box, type the fully qualified path to the virtual hard disk file to convert.
4. Click Inspect.
5. Under Actions, click Convert Virtual Hard Disk.
6. In Location, select the folder in which to store the converted virtual hard disk file. If the folder you want does not appear in the list, in the Converted Virtual Hard Disk Name text box, type the fully qualified path including the filename.
7. In Converted Virtual Hard Disk Type, select the type of virtual hard disk that you want to create.
8. Click Convert.

Once the conversion process completes, you can attach the new virtual hard disk to a virtual machine. You should delete the linked disk that you created prior to the physical disk conversion.

Note Virtual Server 2005 R2 will prevent you from attaching a linked disk to a virtual machine.

VHDMount Command-Line Tool

The VHDMount command-line tool is a new feature delivered with Virtual Server 2005 R2 Service Pack 1 (SP1). This tool allows you to mount a virtual hard disk file as a virtual disk device on a host machine. Using this method, you can inspect, inject, or delete files in the virtual hard disk without having to boot into a virtual machine.

Note By default, the VHDMount command-line tool is located in %systemdrive%\Program Files\Microsoft Virtual Server\Vhdmount.

VHDMount leverages the Virtual Disk Service (VDS), which is a set of application programming interfaces (APIs) that permit management of disks and volumes at the operating system level. Although VDS is available only with Windows Server 2003 and later operating systems, you can still run VHDMount in Windows XP.

When VHDMount is used to mount a virtual hard disk file, VDS interacts with the Plug and Play Manager to discover the virtual hard disk as a new disk and mount it (assigning a drive letter) in the host operating system. Once the virtual hard disk is successfully mounted, a new entry is listed in Device Manager | Disk Drives and is registered as *MS Virtual Server SCSI Disk*

Device. There is also a new entry listed in Device Manager | Microsoft Server Virtual Storage Devices that is registered as *Microsoft Server Virtual Storage DeviceXX*, where *XX* is a unique number that increases sequentially with each mounted device. At this stage, the virtual hard disk file contents can be accessed using standard file system browsing tools such as Windows Explorer.

Important By default, all changes made by VHDMount to a mounted disk are written to an undo disk that is created in the temporary folder for the current user. You can use the /f option to mount a VHD without an undo disk. You can also use the /c option to commit or /d option to discard changes when unplugging a mounted disk.

Defining VHDMount Command-Line Options

VHDMount is a simple utility to use, with only a few options needed to mount and dismount virtual hard disks. Table 5-3 lists the VHDMount command-line options.

Table 5-3 VHDMount Command-Line Options

Command-line option	Description
/p	Plugs in a virtual hard disk file as a virtual disk device without mounting the volume.
/m	Plugs in a virtual hard disk file as a virtual disk device and mounts the volume.
/u	Unplugs a virtual disk device.
/q	Returns the disk name of a mounted virtual disk device.

Important Even though VDS is not available in Windows XP, the virtual disk device should be automatically detected and mounted. However, because you cannot use the /m option with VHDMount, you are unable to specify a starting drive letter to mount a virtual hard disk in Windows XP.

Using VHDMount to Plug In a Virtual Hard Disk File

The following command-line shows the VHDMount option and parameter needed to plug in a virtual hard disk file without mounting the volume:

```
VHDMOUNT.EXE /p VHDFileName
```

VHDFileName indicates the fully qualified path to the virtual hard disk file. For example, to plug in a virtual hard disk file named test.vhd located in folder c:\virtual machines, you type the following:

```
VHDMOUNT.EXE /p "c:\virtual machines\test.vhd"
```

Important When you specify any path values in the command line and those paths contain spaces, you must enclose the entire path in quotes.

Using VHDMount to Plug In and Mount a Virtual Hard Disk File

The next command line shows the VHDMount option and parameter needed to plug in and mount a virtual hard disk file:

```
VHDMOUNT.EXE /m VHDFileName [DriveLetter]
```

DriveLetter is an optional parameter that defines the starting drive letter used to mount virtual hard disk files.

For example, to plug in and mount the virtual hard disk file used in the previous example as drive E, you type the following:

```
VHDMOUNT.EXE /m "c:\virtual machines\test.vhd" E
```

Important When you specify a drive letter in your VHDMount command, do not type a colon after the drive letter. If you do, the VHDMount help screen will display and your command will be ignored.

Using VHDMount to Unmount a Virtual Hard Disk File

The following command line represents the VHDMount option and parameters needed to unmount a virtual hard disk file:

```
VHDMOUNT.EXE /u VHDFileName | All
```

All is an optional parameter that applies the operation to all mounted virtual disk devices.

For example, to dismount all virtual hard disk files, you type the following:

```
VHDMOUNT.EXE /u All
```

Using VHDMount to Determine a Virtual Hard Disk Name

The next command line shows the VHDMount option and parameter needed to determine the disk name associated with the mounted virtual hard disk file:

```
VHDMOUNT.EXE /q VHDFileName | All
```

All is an optional parameter that applies the operation to all mounted virtual disk devices.

For example, to get a listing of disk names associated with all mounted virtual hard disk files, you type the following:

```
VHDMOUNT.EXE /q All
```

On the Companion Media On the companion media, you will find a directory called \Chapter Materials\Files\VHDMount. Inside the directory there is a registry file named Vhdmenu.reg. This file will make registry modifications that add mount and dismount selections to the context menu that appears when you right-click a virtual hard disk file.

VHD Compaction

VHD compaction is a process that reduces the size of a virtual hard disk file on the physical disk. Virtual Server 2005 R2 provides a compaction tool that achieves minor reductions in a virtual hard disk file size if used solely on its own. A better approach is to use a three-step process that includes defragmentation, precompaction, and compaction. Defragmentation and precompaction prepare the virtual hard disk file for the compaction process, resulting in greater reductions in virtual hard disk file size.

Note Prior to virtual hard disk file defragmentation, remove temporary files and folders, delete any other unwanted data, and empty the recycle bin.

VHD compaction can be performed only on dynamically expanding disks. Fixed-size virtual hard disks have to be converted to a dynamically expanding disk prior to being compacted. Special dynamically expanding virtual hard disks, such as differencing or undo disks, cannot be directly compacted. Differencing disks and undo disk changes must be merged into their parent disk, and the parent disk can be compacted if it is a dynamically expanding disk.

Best Practices Because of processor and disk resource requirements, you should use a non-production server, when possible, to perform the virtual hard disk compaction process. In Virtual Server 2005 R2, you can perform the defragmentation step within the virtual machine or while the virtual hard disk is offline. However, it is best to perform defragmentation, precompaction, and compaction with the virtual hard disk file offline.

Defragmenting the Virtual Hard Disk File

The first step in the process to reduce the size of a virtual hard disk file is defragmentation. As new information is written to disk, data might not be saved in contiguous disk blocks. In time, as you delete data on the disk, empty blocks will be randomly filled with file fragments. Performance is adversely affected when the disk fragmentation is excessive because it takes longer to retrieve related data spread across a disk than if it were located in a contiguous set of blocks. Defragmentation reduces or eliminates the number of fragmented files on a disk, resulting in larger areas of empty contiguous blocks.

To defragment a virtual hard disk offline, you first have to use the VHDMount command-line tool to mount the virtual hard disk file. You can find the VHDMount command syntax in the

"Using VHDMount to Plug In and Mount a Virtual Hard Disk File" section earlier in this chapter. Once the virtual hard disk file is mounted, use the Windows Defrag utility on the host system to defragment the virtual hard disk file. Table 5-4 lists the defrag command lines for Windows XP, Windows Server 2003, and Windows Vista. The time required to defragment the virtual hard disk file depends on several factors, including the degree of fragmentation, file size, and disk characteristics.

Table 5-4 Platform-Specific Defragmentation Command Lines

Command line	Operating system
Defrag *DriveLetter* ■ *DriveLetter* is the drive letter associated with the mounted virtual hard disk.	Windows XP Windows Server 2003
Defrag *DriveLetter* –w ■ *DriveLetter* is the drive letter associated with the mounted virtual hard disk. ■ –*w* specifies that all file fragments should be consolidated, regardless of size.	Windows Vista

Precompacting the Virtual Hard Disk File

The second step in the process is precompaction. Virtual Server 2005 R2 includes the Virtual Disk Precompactor tool, which is designed to overwrite any unallocated disk blocks in a virtual hard disk file with zeros. This step is crucial to ensure that the compaction tool can make the virtual hard disk file as small as possible.

The Virtual Disk Precompactor tool is contained in the Precompact.iso disk image located in the %systemdrive%\Program Files\Microsoft Virtual Server\Virtual Machine Additions folder. Use your favorite virtual CD tool to mount the Precompact.iso image on your Virtual Server 2005 R2 host and retrieve the Precompact.exe tool. Table 5-5 lists the options that are available when you invoke the Virtual Disk Precompactor tool from the command line.

Table 5-5 Virtual Disk Precompactor Command-Line Options

Command-line option	Description
-Help	Displays the help dialog box that lists the command-line options, product version, and syntax examples.
-Version	Displays the help dialog box that lists the command-line options, product version, and syntax examples.
-Silent	Executes the precompactor in unattended mode, and suppresses all dialog boxes.

Table 5-5 Virtual Disk Precompactor Command-Line Options

Command-line option	Description
-SetDisks:<list> <*list*> is an optional parameter that represents one or more drive letters.	Defines the list of virtual hard disks to precompact. If this option is not specified, all virtual hard disks attached to a virtual machine are compacted.

For example, the following command precompacts virtual hard disks mounted to drive letters F and G, in unattended mode:

```
Precompact –Silent –SetDisks:FG
```

More Info Virtual Server 2005 R2 allows precompacting virtual hard disk files from within a virtual machine. Once you capture the Precompact.iso image on the virtual machine CD or DVD drive, you can double-click the drive to launch Virtual Disk Precompactor. Using this process, you cannot specify which virtual hard disk to precompact. Instead, Virtual Disk Precompactor precompacts all virtual hard disks attached to the virtual machine.

Compacting the Virtual Hard Disk File

The third and final step in the process to reduce the virtual hard disk size is disk compaction. After running the Virtual Disk Precompactor tool, empty disk blocks in the virtual hard disk file contain zeros. The Virtual Server compaction tool finds the disk blocks that contain zeros and removes them, reducing the virtual hard disk file size.

Caution The Virtual Server compaction tool requires that you have enough disk space to concurrently store the original virtual hard disk file and an additional temporary file that contains the compacted virtual hard disk. The original virtual hard disk file will be deleted at the end of the compaction process and replaced with the compacted virtual hard disk file. If the disk runs out of space before completing the compaction process, an event will be recorded in the Virtual Server event log and no changes will be made to the disk.

To use the Virtual Server compaction tool, follow these steps:

1. Open the Virtual Server 2005 R2 Administration Website.
2. Turn off the virtual machine associated with the dynamically expanding virtual hard disk that you want to compact.
3. In the navigation pane, under Virtual Disks, click Inspect.
4. In the Inspect Virtual Hard Disk pane, select the virtual hard disk to compact in Known Virtual Hard Disks. If the virtual hard disk file does not appear in the list, type the fully qualified path to the virtual hard disk in the Fully Qualified Path To File text box.

5. Click Inspect.
6. Under Actions, click Compact Virtual Hard Disk.
7. In the Compact Virtual Hard Disk pane, click Compact.

The VHD compaction process can also be scripted using the Virtual Server 2005 R2 COM API. This API allows you to create scripts and compact the virtual hard disk files outside of the Virtual Server Administration Website.

On the Companion Media On the companion media, you will find a directory called \Chapter Materials\Scripts\Compact. Inside the directory there are two files, Vhdprep.bat and Compaction.vbs. The Vhdprep.bat file mounts the virtual hard disk file and runs the defragmenter and Virtual Disk Precompactor before calling the Compaction.vbs script. The Compaction.vbs script invokes the Virtual Server compaction tool to compact the virtual hard disk offline.

Using Virtual Network Advanced Features

The Virtual Server 2005 R2 network architecture allows virtual machine network traffic to be isolated from other virtual machines, the Virtual Server 2005 R2 host, and external networks. It also allows virtual machines to be connected to each other, the Virtual Server 2005 R2 host, corporate networks, and the Internet. Many configuration options are available and some depend on the implementation of advanced network settings. Table 5-6 lists Virtual Server 2005 R2 advanced network features covered in this section.

Table 5-6 Virtual Network Advanced Features

Configuration	Description
Microsoft Loopback Adapter	A software-based network adapter that is used to connect virtual machines to internal networks.
Host-to-Guest Networking	Uses the Microsoft Loopback Adapter to enable network connectivity between a Virtual Server 2005 R2 host and virtual machines.
Internet Connection Sharing with Network Address Translation	Uses the Microsoft Loopback Adapter to enable virtual machines to share the Virtual Server 2005 R2 server network access to the Internet.

Using the Microsoft Loopback Adapter

The Microsoft Loopback Adapter is a built-in, software-based network interface that can be attached to virtual networks to provide connectivity between virtual machines. The Microsoft Loopback Adapter can also be used to attach to internal virtual networks linking virtual machines to the Virtual Server 2005 R2 host. Network traffic between virtual machines and the Virtual Server 2005 R2 host is constrained to the internal virtual networks and isolated from external, physical networks.

Installing the Microsoft Loopback Adaptor

The Microsoft Loopback Adaptor is installed on the Virtual Server 2005 R2 host just like a physical network adapter. Here are the steps to install the Microsoft Loopback Adaptor on Windows Server 2003 R2:

1. On the Virtual Server 2005 R2 host, click Start and then click Control Panel.
2. In Control Panel, click Add Hardware and then click Next.
3. In the Is The Hardware Connected dialog box, choose Yes (I Have Already Connected The Hardware) and then click Next.
4. In the Installed Hardware list, choose Add A New Hardware Device and then click Next.
5. In the What Do You Want The Wizard To Do check list, choose Install The Hardware That I Manually Select From A List (Advanced) and then click Next.
6. In the Common Hardware Types list, choose Network Adapters and then click Next.
7. In the Manufacturer list, click Microsoft.
8. In the Network Adapter list, choose Microsoft Loopback Adapter and then click Next.
9. In the Hardware To Install dialog box, click Next.
10. In the Completing The Add Hardware Wizard dialog box, click Finish.

Important You must be a member of the administrators group to install a new network adapter in the Virtual Server host operating system.

Configuring the Microsoft Loopback Adaptor

Before you can use the Microsoft Loopback Adaptor, you must ensure that it is properly configured on your Virtual Server 2005 R2 host. The Microsoft Loopback Adaptor must be bound to Virtual Machine Network Services to allow communications through a virtual network. Once the configuration is complete, you can create virtual networks in the Virtual Server 2005 R2 Administration Website to enable virtual machine network connectivity. Follow these steps to configure the Microsoft Loopback Adaptor bindings on the Virtual Server 2005 R2 host:

1. On the Virtual Server 2005 R2 host, click Start and select Control Panel.
2. Select Network Connections, right-click the local area connection associated with the Microsoft Loopback Adapter and then click Properties.
3. In This Connection Uses The Following Items, ensure that the Virtual Machine Network Services check box is selected.
4. Click Internet Protocol (TCP/IP), and then click Properties.

5. On the General tab, select Use The Following IP Address and then type the IP address and subnet mask, but do not enter a gateway address.
6. Click OK, and then click Close.

Note Use one of the reserved ranges of nonroutable TCP/IP addresses when you configure the Microsoft Loopback Adaptor network address properties. The network address and network mask must be the same on the Virtual Server 2005 R2 host as on the virtual machines that you want to connect to the virtual network.

Implementing Host-to-Guest Networking

Virtual PC 2007 has a Shared Folders feature that allows file sharing between the Virtual PC host and virtual machines. Although no similar feature exists in Virtual Server 2005 R2, you can use the Microsoft Loopback Adapter and virtual networks to enable network connectivity between a Virtual Server 2005 R2 host and virtual machines. Once you have configured this arrangement, you can use standard Windows file sharing features between the physical server and virtual machines.

Creating a Virtual Network for Host-to-Guest Networking

After the Microsoft Loopback Adapter has been installed and configured on the Virtual Server 2005 R2 host, you can create a new virtual network to which you connect the virtual machines. To accomplish this, perform the following steps:

1. Open the Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Networks, click Create.
3. In the Virtual Network Name text box, type a name for the virtual network.
4. In Network Adapter On Physical Computer, select the Microsoft Loopback Adapter.
5. In Disconnected Virtual Network Adapters, select the Connected check box for any virtual machine network adapter that you want to attach to the new virtual network.
6. In the Virtual Network Notes text box, type in a description for the new virtual network and then click OK.

You can now boot the virtual machines, configure the network address for the new local connection, and configure firewall settings to enable resource sharing, as required.

Enabling a Virtual DHCP Server on a Virtual Network

If you intend to connect several virtual machines to the host-to-guest virtual network, you should configure the Virtual DHCP Server option on the virtual network. The Virtual DHCP

server will manage and provide network configuration options to connecting virtual machines. These are the steps to enable the Virtual DHCP Server option:

1. Open the Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Networks, select Configure and then click the appropriate virtual network.
3. In the Virtual Network Properties pane, click DHCP server.
4. Choose the Enabled check box, and configure the DHCP server options as needed.
5. Click OK.

Note In the DHCP Server options, you can see that the first 16 IP addresses from the start of the specified range are reserved. These 16 IP addresses are never assigned; use one in that range to configure the Virtual Server host adapter.

Configuring Internet Connection Sharing and Network Address Translation

Using the Microsoft Loopback Adapter, you can also configure Internet Connection Sharing (ICS) on the Virtual Server 2005 R2 host to provide virtual machine connectivity to external networks using Network Address Translation (NAT). This configuration provides external network access without the provisioning of official network addresses or direct virtual machine connection to the physical network. The major steps to implement this scenario are as follows:

1. Install the Microsoft Loopback Adapter on the Virtual Server 2005 R2 host.
2. Configure Internet Connection Sharing on the Microsoft Loopback Adapter.
3. Create a virtual network using the Microsoft Loopback Adapter.
4. Connect virtual machines to the virtual network.

All steps are covered in previous examples, with the exception of the Internet Connection Sharing configuration on the Virtual Server 2005 R2 host. Here are the steps to complete the Internet Connection Sharing configuration on Windows Server 2003 R2:

1. On the Virtual Server 2005 R2 host, click Start and select Control Panel.
2. Select Network Connections, and click on the connection that provides Internet connectivity.
3. In the Local Area Connection Status dialog box, on the General tab, click Properties.
4. Click the Advanced tab.

5. In Internet Connection Sharing, select the Allow Other Network Users To Connect Through This Computer's Internet Connection check box.
6. Click OK.

You can use the network connection Repair option in the virtual machines to force connections to refresh the IP address configuration from the Internet Connection Sharing host.

Caution If IPSec is configured on the Virtual Server 2005 R2 host, you cannot use Internet Connection Sharing to provide external network access to virtual machines.

Using Clustering Advanced Features

A common issue that arises when considering the deployment of a virtualized infrastructure is that a single physical server running multiple workloads becomes a more critical point of failure, with an impact on a larger user and business base than a single physical server running a single workload. Clustering addresses this risk by providing high-availability solutions that are as applicable in the virtualization space as in the physical server space. In this section, you will learn how to configure virtual machines and Virtual Server 2005 R2 hosts to implement the clustering scenarios listed in Table 5-7.

Table 5-7 Virtual Server 2005 R2 Advanced Cluster Configurations

Feature	Description
Virtual Machine Cluster Using iSCSI	A cluster based on Microsoft Cluster Server (MSCS) that consists of two or more virtual machine cluster nodes supporting a cluster-aware application. Virtual machine cluster nodes can be located across Virtual Server 2005 R2 hosts, but they require iSCSI-based disks.
Virtual Server Host Cluster	A cluster based on Microsoft Cluster Server that consists of two or more Virtual Server 2005 R2 host cluster nodes.

In Virtual Server 2005, you could create only a two-node virtual machine cluster based on virtual SCSI adapters. This required the cluster nodes to be located on the same Virtual Server 2005 host. A two-node virtual machine cluster could be useful in test environments using cluster-aware applications, but it was not a solution that could be deployed and supported in a production environment. In effect, the Virtual Server 2005 host represented a single point of failure, so the solution could not meet high-availability production requirements.

Virtual Server 2005 R2 removed the two-node and single-host virtual machine cluster limitations by adding support for the iSCSI protocol. Using iSCSI shared disks, multinode clusters can be created using virtual machines hosted on separate Virtual Server 2005 R2 hosts. This type of cluster is still recommended for virtual machines running cluster-aware applications.

Virtual Server 2005 R2 also introduced support for Virtual Server host clusters. Virtual Server 2005 R2 host clusters allow failing over individual or all virtual machine workloads to other Virtual Server 2005 R2 host cluster member nodes. For virtual machines running non–cluster aware applications, Virtual Server 2005 R2 host clusters are a basic building block for the implementation of high-availability solutions.

Virtual Server 2005 R2 SP1 includes all the clustering features found in Virtual Server 2005 R2. Once you have installed Virtual Server 2005 R2 SP1 on a host, you will have access to a whitepaper with detailed information concerning Virtual Server 2005 R2 host clusters. Previously provided as a download from the Microsoft Web site, the whitepaper is now packaged in the Virtual Server 2005 R2 SP1 distribution media. You can find the whitepaper in the %systemdrive%\Program Files\Microsoft Virtual Server\Host Clustering directory on your Virtual Server 2005 R2 SP1 host.

Implementing a Virtual Machine Cluster Using iSCSI

With Virtual Server 2005 R2 SP1, virtual machine clusters are now supported for production workloads when used in conjunction with iSCSI-based shared disk systems. Using iSCSI to deploy a cluster eliminates the need for the specialized hardware that was previously required to configure clustering. The requirements for an iSCSI-based solution are network adapters to connect the storage to the cluster nodes, and a storage unit that uses iSCSI. The iSCSI protocol defines the rules and processes for transmitting and receiving block storage data over TCP/IP networks. iSCSI-based implementations consist of an iSCSI initiator and an iSCSI target with an interconnecting network.

Virtual machine clusters implemented with iSCSI require each cluster node to be located on separate Virtual Server 2005 R2 hosts. Virtual machine clusters can range from two-node to eight-node active clusters. Physical distance between cluster nodes is restricted by the iSCSI protocol and the maximum latency that a cluster heartbeat signal can support.

Table 5-8 lists implementation requirements prior to creating a two-node virtual machine cluster based on an iSCSI storage device.

Table 5-8 Requirements for an ISCSI-Based Virtual Machine Cluster

Requirement	Description
Operating System	Windows Server 2003 R2 Enterprise Edition must be installed on each virtual machine cluster node.
Virtual Machine Additions	Virtual Machine Additions must be installed on each virtual machine node.
iSCSI Quorum and Shared Disks	iSCSI Quorum and Shared Disks targets must be created prior to configuring the cluster nodes. The Quorum disk must be at least 50 MB in size to satisfy Microsoft Cluster Server requirements.

Table 5-8 Requirements for an ISCSI-Based Virtual Machine Cluster

Requirement	Description
Network Adapters	Three network adapters must be added and configured for the Public, Private, and iSCSI networks on each virtual machine cluster node.
Virtual Networks	Virtual networks must be created for non-cluster traffic and iSCSI traffic (Public, Private, and iSCSI).
Active Directory	Virtual machine cluster nodes must be members of an Active Directory domain.
Cluster Service Account	A cluster service account must be created in Active Directory.

To deploy a two-node virtual machine cluster using iSCSI, you must perform the following major steps:

1. Create a shared drive for quorum and data storage using the iSCSI Initiator.
2. Configure virtual networks on each of the Virtual Server 2005 R2 hosts.
3. Configure shared drives on each virtual machine cluster node.
4. Install Microsoft Cluster Server on the first virtual machine cluster node and assign the shared drive.
5. Install Microsoft Cluster Server on the second virtual machine cluster node, join it to the cluster, and assign the shared drive.

Note The Microsoft iSCSI Initiator service is included in the Microsoft iSCSI Software Initiator package, which you can download from the Microsoft Web site at *http://go.microsoft.com/fwlink/?linkid=44352*.

Configuring the iSCSI Shared Disks

After you build your base virtual machines, you can configure the cluster shared disks. Follow these steps to configure virtual machine cluster node access to iSCSI shared disks:

1. Install the Microsoft iSCSI Initiator software in the first virtual machine.
2. Click Start, click All Programs, click Microsoft iSCSI Initiator, and then click Microsoft iSCSI Initiator again.
3. Click the Discovery tab, and in Target Portals, click Add.
4. Enter the name or IP address of the server where the target iSCSI drive is defined.
5. Click the Targets tab to display a list of disk targets.
6. Select Quorum and click Log On.

7. Select Automatically Restore This Connection When The System Boots And Enable Multipath, if you have multipath software installed.
8. Repeat steps 6 and 7 for the Shared target, and then click OK.
9. In the Disk Management MMC, format each disk with a single partition, using drive letter Q for the Quorum disk and drive letter S for the Shared disk.
10. Shut down the virtual machine.
11. Repeat steps 1 to 8 for the second virtual machine.
12. In the Disk Management MMC, set the Quorum drive letter to Q and the Shared drive letter to S.

Configuring Microsoft Cluster Server on the First Virtual Machine

When you create the first node in a cluster, you specify all parameters that define the cluster configuration. The Cluster Configuration Wizard guides you through the installation and completes the cluster setup when you have entered all the required information.

Caution During the configuration of Microsoft Cluster Server on the first cluster node, you must power-off all other nodes. This is to avoid data corruption on the shared disks. Ensure that the first cluster node can successfully access all volumes before attempting to join additional cluster nodes.

Follow these steps to configure Microsoft Cluster Server on the first virtual machine cluster node:

1. Log in to the virtual machine with Domain Administrator credentials.
2. Click Start, click All Programs, click Administrative Tools, and then click Cluster Administrator.
3. When prompted with the Open Connection To Cluster dialog box, select Create New Cluster in the Action drop-down list.
4. Review the information list in the New Server Cluster Wizard, and then click Next.
5. In the Cluster Name text box, type a name for the cluster and then click Next.
6. In the Computer Name text box, type the computer name of the virtual machine that is the first node in the cluster.
7. Click Next.
8. Remedy any errors found in the Analyzing Configuration step, and then re-analyze. If there are no further errors, click Next.
9. In the IP Address text box, type an IP address on the public network that will be used to manage the cluster and click Next.

10. In the User Name text box, type the name of the cluster service account that you created in Active Directory.
11. In the Password text box, type the password for the cluster service account.
12. In Domain, select your domain name from the drop-down list and then click Next.
13. Review the Summary page to verify that all information used to create the cluster is correct.
14. Click Quorum, select Disk Q: from the drop-down list, and then click OK.
15. Click Next.
16. Once the cluster creation is complete, click Next.
17. Click Finish to complete the installation.

Configuring Microsoft Cluster Server on the Second Virtual Machine

Installing Microsoft Cluster Server on the second virtual machine is much quicker because the cluster configuration already exists. Additional cluster nodes are simply joined to the defined cluster.

When adding subsequent nodes, leave the first cluster node and all shared disks turned on, and power-up additional nodes. The cluster service will control access to the shared disks to eliminate any chance of corruption. Follow these steps to configure the second node (and any subsequent node) in the cluster:

1. Open Cluster Administrator on the first cluster node.
2. Click File, click New, and then click Node.
3. On the Add Cluster Computers Wizard Welcome page, click Next.
4. In the Computer Name text box, type the computer name for the second cluster node and then click Add.
5. Click Next.
6. Remedy any errors found in the Analyzing Configuration step, and then re-analyze. If there are no further errors, click Next.
7. Type the password for the cluster service account, and then click Next.
8. Review the summary information that is displayed for accuracy, and then click Next.
9. Review any warnings or errors encountered during cluster creation, and then click Next.
10. Click Finish to complete the installation.

To quickly verify that cluster failover is successful, you can shut down the first cluster node. When you open Cluster Administrator on the second cluster node, you will see that it now

owns all cluster resources. Once you have tested that cluster failover is successful, you can proceed with the installation of the cluster-aware application.

Implementing a Virtual Server Host Cluster Using iSCSI

To achieve high availability for non–cluster aware applications running in virtual machines, you must implement a Virtual Server 2005 R2 host cluster. Virtual Server 2005 R2 host clusters can be deployed using SCSI, SAN, or iSCSI-based shared storage. Like virtual machine clusters, Virtual Server 2005 R2 host clusters can range from two-node to eight-node active clusters. It is important to understand that in this configuration, you are clustering the Virtual Server 2005 R2 hosts, not the applications running in the virtual machines. If one of the Virtual Server 2005 R2 host cluster nodes fails, virtual machines defined as resource groups in the cluster configuration are restarted on other Virtual Server 2005 R2 host cluster member nodes. In contrast, failure of an application running within a virtual machine will not result in a failover event.

Important The complete set of hardware used to implement a Virtual Server Host cluster must be listed in the Windows Server Catalog as a qualified cluster solution for Windows Server 2003.

There are many scenarios to which you can apply a Virtual Server 2005 R2 host cluster solution. Table 5-9 lists the most common scenarios that benefit from a Virtual Server 2005 R2 host cluster implementation.

Table 5-9 Virtual Server Host Cluster Scenarios

Scenario	Virtual Server host cluster benefits
Host hardware scheduled maintenance	Prior to performing hardware maintenance on a Virtual Server cluster node, hosted virtual machines can move groups over to other nodes in the cluster with minimal impact on application availability.
Host software updates	Before applying potentially disruptive software updates to the host, hosted virtual machines can fail over to other nodes in the cluster with minimal impact on application availability.
Non–cluster aware applications	Non–cluster aware applications running in virtual machines on a Virtual Server 2005 R2 host cluster node are protected from unexpected downtime caused by a host failure. If the Virtual Server 2005 R2 host cluster node fails, the virtual machine can fail over to other nodes in the cluster with minimal impact on application availability.

Table 5-9 Virtual Server Host Cluster Scenarios

Scenario	Virtual Server host cluster benefits
Workload rebalancing	Virtual machine performance might dictate a need to rebalance the workload on a Virtual Server 2005 R2 host cluster node. If there is another cluster node with the required resources available, the virtual machine can be quickly failed over with minimal impact on application availability.

During an unplanned cluster failover event, there is always some short period of time during which the cluster-defined resources are unavailable as they are restarted on a different cluster node. Microsoft Cluster Server ensures that the applications experience minimal service disruptions. If an administrator performs a normal shutdown on a cluster node or moves a guest from one host to another for planned maintenance, Virtual Server 2005 R2 can save the virtual machine state before it is moved.

Because virtual machines running in Virtual Server 2005 R2 are not cluster-aware, Microsoft created a script that ensures that virtual machines function correctly during cluster failover events. Each virtual machine is configured as a cluster resource group. Inside each cluster resource group, the script is configured as a Generic Script resource that has the effect of turning a virtual machine into a cluster-aware-like application. The script can also restart a virtual machine when it stops running. Underlying this whole process is the Microsoft Cluster Server, which provides the health monitoring and automatic recovery for the virtual machine.

On the Companion Media On the companion media, you will find a directory called \Chapter Materials\Scripts\Cluster. Inside the directory there are two files: Stop_clussvc_script.cmd and Havm.vbs. These files are needed during the configuration of Virtual Server 2005 R2 host cluster nodes. A listing of the script is also included in the Virtual Server Host Clustering Step-by-Step Guide for Virtual Server 2005 R2," located at %systemdrive%\Program Files\Microsoft Virtual Server\Host Clustering.

Table 5-10 lists implementation requirements prior to creating a Virtual Server 2005 R2 host cluster based on iSCSI shared storage that is supported in a production environment.

Table 5-10 Requirements for iSCSI-Based Virtual Server Host Cluster

Requirement	Description
Physical Hardware	Creation of a Virtual Server 2005 R2 host cluster supported in production requires two or more identical physical servers that are listed in the Windows Server Catalog.

Table 5-10 Requirements for iSCSI-Based Virtual Server Host Cluster

Requirement	Description
Operating System	Windows Server 2003 Enterprise Edition (SP1 or R2). Windows Server 2003 Datacenter Edition (SP1 or R2).
iSCSI	Microsoft iSCSI Software Initiator 2.0 or later version.
iSCSI Quorum and Shared Disks	iSCSI Quorum and Shared Disks targets must be created prior to configuring the cluster nodes. The Quorum disk must be at least 50 MB to satisfy Microsoft Cluster Server requirements. The Shared disk must be sized to contain virtual machine VHD files.
Network Adapters	Three network adapters must be added and configured for the Public, Private, and iSCSI networks on each Virtual Server 2005 R2 host cluster node.
Active Directory	Virtual Server 2005 R2 host cluster nodes must be members of an Active Directory domain.
Cluster Service Account	A cluster service account must be created in Active Directory.
Virtual Machine Additions	Virtual Machine Additions must be installed on each virtual machine.
Support Files	Havm.vbs and Stop_clussvc_script.cmd, located on the companion media.

To deploy a two-node Virtual Server 2005 R2 host using iSCSI, you must perform the following major steps:

1. Create a shared drive for quorum and data storage using the iSCSI Initiator.
2. Configure Microsoft Cluster Server on each Virtual Server 2005 R2 host.
3. Configure both Havm.vbs and Stop_clussvc_script.cmd on each Virtual Server 2005 R2 host.
4. Configure a cluster disk resource, resource group, and resource script.
5. Configure a virtual machine on one of the Virtual Server 2005 R2 hosts.

Important For a more detailed list of limitations and requirements, refer to the Virtual Server Host Clustering Step-by-Step Guide for Virtual Server 2005 R2 SP1.

Configuring the iSCSI Shared Disks

Follow these steps to configure virtual machine cluster node access to iSCSI shared disks:

1. Install the Microsoft iSCSI Initiator software on the first Virtual Server 2005 R2 host.
2. Click Start, click All Programs, click Microsoft iSCSI Initiator, and then click Microsoft iSCSI Initiator again.
3. Click the Discovery tab, and in Target Portals, click Add.
4. Enter the name or IP address of the server where the target iSCSI drive is defined.
5. Click the Targets tab to display a list of disk targets.
6. Select Quorum and click Log On.
7. Select Automatically Restore This Connection When The System Boots And Enable Multipath if you have multipath software installed.
8. Repeat steps 6 and 7 for the Shared target, and then click OK.
9. In the Disk Management MMC, format each disk with a single partition, using drive letter Q for the quorum disk and drive letter S for the Shared disk.
10. Shut down the Virtual Server 2005 R2 host.
11. Repeat steps 1 through 8 for the second Virtual Server 2005 R2 host.
12. In the Disk Management MMC, set the Quorum drive letter to Q and the Shared drive letter to S.

Configuring Microsoft Cluster Server on the First Virtual Server Host

Follow these steps to configure Microsoft Cluster Server on the first virtual server host:

1. Log in to the first Virtual Server 2005 R2 host with Domain Administrator credentials.
2. Click Start, click All Programs, click Administrative Tools, and then click Cluster Administrator.
3. When prompted with the Open Connection To Cluster dialog box, select Create New Cluster in the Action drop-down list.
4. Review the information list on the New Server Cluster Wizard Welcome page, and then click Next.
5. In the Cluster Name text box, type a name for the cluster and then click Next.
6. In the Computer Name text box, type the computer name of the virtual machine that is the first node in the cluster.
7. Click Next.
8. Remedy any errors found in the Analyzing Configuration step and then re-analyze. If there are no further errors, click Next.

9. In the IP Address text box, type an IP address on the public network that will be used to manage the cluster and click Next.
10. In the User Name text box, type the name of the cluster service account that you created in Active Directory.
11. In the Password text box, type the password for the cluster service account.
12. In Domain, select your domain name from the drop-down list and then click Next.
13. Review the Summary page to verify that all information used to create the cluster is correct.
14. Click Quorum, select Disk Q: from the drop-down list, and then click OK.
15. Click Next.
16. Once the cluster creation is complete, click Next.
17. Click Finish to complete the installation.

Configuring Microsoft Cluster Server on the Second Virtual Server Host

Installing Microsoft Cluster on the second Virtual Server 2005 R2 host is again a quick process because the cluster configuration already exists. Additional cluster nodes just have to be added to the existing cluster.

When adding subsequent nodes, leave the first cluster node and all shared disks turned on, and power-up additional nodes. The cluster service will control access to the shared disks to eliminate any chance of corruption. Follow these steps to configure the second node (and any subsequent node) in the cluster:

1. Open Cluster Administrator on the first cluster node.
2. Click File, click New, and then click Node.
3. On the Add Cluster Computers Wizard Welcome page, click Next.
4. In the Computer Name text box, type the computer name for the second cluster node and then click Add.
5. Click Next.
6. Remedy any errors found in the Analyzing Configuration step and then re-analyze. If there are no further errors, click Next.
7. Type the password for the cluster service account, and then click Next.
8. Review the summary information that is displayed for accuracy, and then click Next.
9. Review any warnings or errors encountered during cluster creation, and then click Next.
10. Click Finish to complete the installation.

Configuring the Shutdown Script for Virtual Server Host Cluster Nodes

Because Virtual Server 2005 R2 is not a cluster-aware application, you have to ensure that the cluster service shuts down and all virtual machines are failed over prior to a Virtual Server Host shutdown. Follow these steps to configure the shutdown script for the Virtual Server 2005 R2 host cluster nodes:

1. In the root directory of the local hard disk on each Virtual Server 2005 R2 host, copy the Stop_clussvc_script.cmd file from the companion media.
2. Click Start, click Run, and then type **gpedit.msc**.
3. Click Enter.
4. Navigate to Local Computer Policy, click Computer Configuration, click Windows Settings, and then click Scripts.
5. In the right-hand pane, double-click Shutdown, and click Add.
6. In the Script Name text box, type the fully qualified path name of the batch file, and then click OK twice.

Configuring the Disk Resource, Resource Group, and Havm.vbs Script

Follow these steps to configure the cluster disk resource, resource group, and cluster control script:

1. On the first Virtual Server 2005 R2 host, click Start, click Control Panel, click Administrative Tools, and then click Cluster Administrator.
2. In Cluster Administrator, create a new resource group and name it **Group0**. If you want to specify a Preferred Owner for the group, specify the node on which you want the guest to run most of the time.
3. In Cluster Administrator, create a new disk resource, or use the appropriate disk resource if it has already been created. Verify that it is the Shared disk configured as a Physical Disk Resource with no dependencies, assigned to resource group Group0, and both cluster nodes are listed as Possible Owners.
4. With Group0 online, create a folder on the Shared disk called **GuestVM1**.
5. On each Virtual Server 2005 R2 host cluster node, create a folder on the local disk in %systemroot%\Cluster and copy the Havm.vbs script into it from the companion media.

Important If you want to create and fail over multiple virtual machines independently, you have to configure each guest in its own resource group. If you want to fail over certain virtual machines together, you need to configure them in the same resource group.

Creating a Virtual Machine on the First Virtual Server Host

Follow these steps to configure the virtual network and virtual machine on the first Virtual Server 2005 R2 host cluster node:

1. Click Start, click All Programs, click Microsoft Virtual Server, and then click Virtual Server 2005 R2 Administration Website.
2. In the navigation pane, under Virtual Networks, click Create.
3. In the Virtual Network Name text box, type in a name for the cluster network.
4. In Network Adapter On Physical Computer, select the network adapter associated with the public network, and then click OK.
5. In the navigation pane, under Virtual Networks, click Configure, and then click View All.
6. In Virtual Networks, click on the virtual network you created, and then click Edit Configuration.
7. Copy the fully qualified path of the .vnc file.
8. Open Explorer, and paste the fully qualified path of the .vnc file (without the filename) into the address bar.
9. Right-click the cluster network name you just created, and then click Cut.
10. In Explorer, navigate to the GuestVM1 folder on the Shared disk and paste the .vnc file.
11. Open the Virtual Server 2005 R2 Administration Website.
12. In the navigation pane, under Virtual Networks, click Add.
13. In the Existing Configuration (.vnc) File text box, type the fully qualified path to the new .vnc file that you created in the Shared disk GuestVM1 folder and then click Add.
14. Copy an existing virtual machine into the Shared disk GuestVM1 folder.
15. In the navigation pane, under Virtual Machines, click Add.
16. In the Fully Qualified Path To File text box, type the fully qualified path to the virtual machine .vmc file and then click Add.
17. In the virtual machine Configuration pane, click Network Adapters.
18. In the Virtual Machine Network Adapter Properties pane, in the Connected To drop-down box, select the cluster network that you created and then click OK.

Completing the Virtual Machine Configuration on the Second Virtual Server Host

Follow these steps to complete the configuration of the virtual machine on the second Virtual Server 2005 R2 host:

1. On the second Virtual Server 2005 R2 host, click Start, click All Programs, click Administrative Tools, and then click Cluster Administrator.
2. Move Group0 to the second Virtual Server 2005 R2 host cluster node.
3. Open the Virtual Server 2005 R2 Administration Website.
4. In the navigation pane, under Virtual Networks, click Add.
5. In the Existing Configuration (.vnc) File text box, type the fully qualified path to the .vnc file located on the Shared disk in the GuestVM1 folder and then click Add.
6. In the navigation pane, under Virtual Machines, click Add.
7. In the Fully Qualified Path To File text box, type the fully qualified path to the virtual machine .vmc file and then click Add.
8. Open Cluster Administrator, and create a new script resource called **GuestVM1Script**.
9. Configure the resource as a Generic Script resource, assign it to Group0 with Possible Owners listing both cluster nodes, and add a Shared disk as a resource dependency.
10. In the Script Filepath text box, type **%windir%\Cluster\Havm.vbs**.
11. Click Start, and then click Run.
12. Type '**cluster res "Guest1Script" /priv VirtualMachineName=*GuestVM1***', replacing *GuestVM1* with the name of the virtual machine that you added, and then press Enter.
13. Open Cluster Administrator and bring Group0 online.
14. Open the Virtual Server 2005 R2 Administration Website.
15. Verify that the virtual machine is in the Running state.

You can now verify that the virtual machine fails over to the first Virtual Server 2005 R2 host cluster node. To do this, open the Cluster Administrator and choose the Move Group option for the Group0 resource group. You should see the Owner field change when the virtual machine has failed over.

Summary

There are many advanced features in Virtual Server 2005 R2 that you can leverage to optimize virtualization infrastructure deployments. If you are going to create complex testing, support desk, or user-training scenarios, use differencing disks and undo disks to enable quick provisioning of new virtual machine configurations with the ability to roll back to the baseline state. When you need to reduce the size of dynamically expanding disks, use defragmentation and precompaction prior to the VHD compaction tool to minimize the size of compacted virtual hard disks. Configure the Microsoft Loopback Adapter to create isolated network connections between a Virtual Server 2005 R2 host and its hosted virtual machines. For cluster-aware applications running within virtual machines, use a virtual machine cluster to minimize downtime from virtual machine failures. In the case of non–cluster aware applications, deploy high-availability Virtual Server 2005 R2 host clusters to reduce planned and unplanned downtime.

Additional Resources

The following resources contain additional information related to the topics in this chapter:

- Knowledge Base article 311272, "The DevCon command-line utility functions as an alternative to Device Manager," at *http://support.microsoft.com/kb/311272*
- White paper, "Virtual Hard Disk Image Format Specification," at *http://www.microsoft.com/windowsserversystem/virtualserver/techinfo/vhdspec.mspx*
- White paper, "Using iSCSI with Virtual Server 2005 R2," at *http://go.microsoft.com/fwlink/?LinkId=55646*
- White paper, "Virtual Server Host Clustering Step-by-Step Guide for Virtual Server 2005 R2," in %systemdrive%\Program Files\Microsoft Virtual Server\Host Clustering