
Installing the IBM[®] ServeRAID[™] Cluster Solution

For IBM Netfinity[®] ServeRAID-4x Ultra160 SCSI Controllers

Chapter 1. Introduction to IBM[®] ServeRAID[™] Cluster Solution

A *cluster* is a type of parallel or distributed system that consists of interconnected computers that are used as a single, unified computing resource. In other words, a cluster is a group of computers linked together in such a way that they share and manage a set of resources that can support a number of users at the same time.

This *high-availability clustering solution* is based on a two-server cluster, where both servers can access the same storage devices, but only one server at a time controls the storage devices shared by both servers. If one of the two running servers fails, the remaining server automatically assumes control of the shared resources managed by the failed server while still controlling its own resources at the same time. The failed server can then be repaired offline without the loss of time or work efficiency because the shared data and applications earlier managed by the failed server are still online to the clients.

When the failed server is operational again, it can be placed back into the cluster; the resources are reallocated between the two servers and the cluster resumes normal operation.

A *controller* is a device that coordinates and controls the operation of one or more input and output devices, such as workstations, and synchronizes the operation of such devices with the operation of a system as a whole.

IBM ServeRAID Cluster Solution, as discussed in this book, uses the IBM Netfinity[®] ServeRAID-4x Ultra160 SCSI controller. IBM ServeRAID Cluster Solution also uses IBM expansion enclosures, such as the following:

- IBM Netfinity EXP15 Rack Expansion Enclosure
- IBM Netfinity EXP200 Rack Expansion Enclosure
- IBM Netfinity EXP300 Rack Expansion Enclosure

You can install IBM ServeRAID Cluster Solution by using one of the following software applications:

- Microsoft[®] Cluster Service, which is part of Windows NT[®] Server Enterprise Edition
- Microsoft Cluster Service using Windows[®] 2000, which is part of Windows 2000 Advanced Server Edition or Windows 2000 Datacenter Server Edition
- Vinca High Availability for Novell IntranetWare, which is part of the IBM Netfinity Cluster Pack by Vinca
- Novell High Availability Server 1.0

The following figures show an example of a simple high-availability clustering solution. The first figure shows a high-level view of a cluster; the second figure shows a more detailed view of a cluster.




Figure 1. A high-level view of a cluster

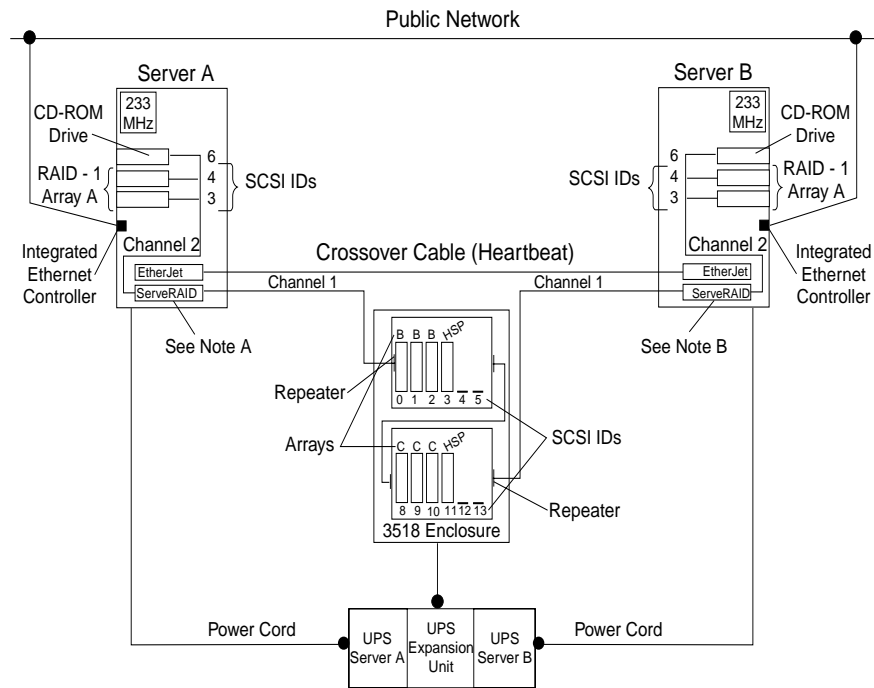


Figure 2. A detailed view of a cluster

Note A: Server A

- Shared Channel 1 - SCSI ID 6
- Non-shared Channel 2 - SCSI ID 6

- Non-shared Channel 3 - SCSI ID 7

Note B: Server B

- Shared Channel 1 - SCSI ID 7
- Non-shared Channel 2 - SCSI ID 7

The solution you use for your installation depends on your specific requirements. For more information about requirements and types of installations, refer to the *IBM Shared Disk Clustering Hardware Reference*. You can obtain this publication from the IBM Support Web site. See "Downloadable files from the World Wide Web" on page 40 for additional information.

Chapter 2. Preparing to install or change a cluster solution

Note: If you are going to upgrade an existing high-availability cluster solution and are already familiar with the concepts of a cluster, go to Chapter 6, "Updating the ServeRAID BIOS, firmware, and software code for clustering," on page 37. If you are setting up a high-availability cluster solution for the first time, continue reading this section.

Before you begin installing and setting up a high-availability cluster solution, familiarize yourself with the following terms and definitions that are used in this manual, as well as some considerations concerning the ServeRAID controller.

Note: The following terms and definitions are for the Microsoft Cluster Service using Windows NT or Windows 2000. For similar terms and definitions for the Novell IntranetWare Cluster Solution using Novell NetWare, refer to the *Vinca High Availability for NetWare* manual.

Terms and definitions

The following terms are used in this manual:

Failover is the process of relocating the cluster resources from a failed server to the remaining server. The resource-monitor program responsible for the resource detects the failure. If a resource failure occurs, the resource-monitor program notifies the Cluster Service, which triggers the actions that are defined in the failover policy for that resource. A failover can occur automatically, such as when an application or a server fails, or it can occur manually, such as when a system administrator moves all applications onto one server and then shuts down the other server for scheduled maintenance. Failover generally includes three tasks:

- Failure detection
- Resource relocation
- Application restart.

Failback is a special case of failover. It is the process of moving back some or all groups of resources to the "preferred owner" after a failover has occurred. The *preferred owner* is the server in the cluster on which you prefer each group of resources to run. If the preferred owner fails, its resources are transferred to the other server. When the preferred owner is repaired and comes back online, and **allow failback** is enabled in the Microsoft Cluster Administration program, the resources are automatically transferred back to that server.

Resource Group is a set of dependent resources that can be grouped together. Some resources need other resources to run successfully and are known as *resource dependencies*. When one resource is listed as a dependency for another resource, the two resources must be placed in the same group. If all resources are ultimately dependent on one resource (for example, one physical disk resource), all resources must be in the same group. Therefore, all resources in a cluster might need to be in the same group.

Furthermore, any cluster operation performed on a group is performed on all resources that are in that group. For example, if a resource must be moved from Server A to Server B, all resources that are in the same group as the resource will be moved also.

ServeRAID considerations

It is important to understand the following ServeRAID considerations before you set up your high-availability cluster solution.

Note: The following ServeRAID considerations are for the Microsoft Cluster Solution using Windows NT and Windows 2000. For similar considerations for the Novell IntranetWare Cluster Solution, refer to the *Vinca High Availability for NetWare* manual.

Hardware considerations

- Do not connect nondisk devices, such as CD-ROM drives, to shared channels. These devices will be accessible to both servers, and conflicts might arise.
- You must cable the SCSI channels in a controller pair so that the controllers share the same channels.
- You must connect all disks that will be shared to the shared channels. This includes hot-spare drives, as well as data drives.
- All shared physical drives must be part of a logical drive or defined as a hot spare to one or the other servers. You must remove all other physical drives from the shared disk enclosure to ensure proper operation.

Controller considerations

- Each ServeRAID controller must have a unique controller name assigned.
- A *controller pair* is a term that refers to two controllers, one in each system, that are connected to each other through the same SCSI bus.
- Do not change the controller name on both ServeRAID controllers in a pair at the same time. Doing so can cause a problem, such as a server *lock-up*.
- Each ServeRAID controller, and controller pair, must have a different SCSI initiator ID assigned to the shared channels. For example, set one server, Server A, to ID 6, and set the other server, Server B, to ID 7. The setup instructions describe this in more detail.
- You must set the stripe-unit size of each ServeRAID controller in a pair to the same value (8 KB, 16 KB, 32 KB, or 64 KB).
- You *cannot* share hot-spare drives between controllers in a cluster pairing. If you want hot-spare protection, *each* ServeRAID controller must have a hot-spare drive defined.

You *must* connect hot-spare drives *only* to shared channels. This is in case the hot-spare drive replaces a failed shared drive. Both servers must be able to access the drive in case a failover occurs *after* the shared drive fails.

- Physical drives that are configured by one controller in the pair must appear as ready drives to the other controller in the pair. Be sure to verify this criterion when you add drives. Restart both servers, or scan for new devices using the ServeRAID Manager program on both servers, to make sure the added physical drives appear as ready drives on both servers before creating logical drives.

Array or logical drive considerations

- The ServeRAID controller does not permit the failover of RAID level-5 logical drives that are in a critical state; that is, a drive in a RAID level-5 array that has failed. For this reason, do not use RAID level-5 for the quorum drive. You must configure the quorum drive as a RAID level-1 logical drive.

Note: Define hot-spare drives in your array to minimize the time that logical drives remain in critical state.

- Every logical drive that is shared by the two servers *must* have its cache policy set to write-through mode to ensure that data integrity is maintained. Logical drives that are not shared between the two servers can be configured for write-back mode for improved performance.
- Create only *one* logical drive for each array.
- You must assign merge group numbers in the range 1– 8 to each logical drive that will be shared. Merge group numbers must be unique for each shared logical drive in the cluster. You must assign merge group numbers 206 or 207 to the non-shared logical drives.
- If you are starting (booting) the operating system from a shared controller, define the first logical drive as the startup drive and assign a merge group number for a non-shared drive, for example, 206 for Server A.
- The total number of logical drives per controller must be eight or less before or after a failover. If you exceed this number, a failover will not complete.
- The following logical drives will not failover:
 - Logical drives currently undergoing Logical Drive Migration (LDM) operations
 - RAID level-5E logical drives undergoing compression or decompression.However, all other logical drives will failover if necessary.
- If a failover occurs while a critical RAID level-1, RAID level-1E, RAID level-10 or RAID level-1E0 logical drive is rebuilding to a spare disk, the Rebuild operation automatically starts.
- The cluster support software will initiate a synchronization of all RAID level logical drives except for the following:
 - RAID level-0 logical drives
 - Any critical or offline logical drives.If a drive fails before this synchronization is complete, logical drive access is blocked and is no longer accessible.
- When a logical drive spans multiple SCSI channels and a failure within the drive subsystem occurs that is unique to a channel (for example, a disconnected cable), the entire physical array will be identified as failed even though access from the surviving server can occur. Therefore, if you have small arrays, consider not spanning multiple channels.

Use the **Validate cluster** action in the ServeRAID Manager program to verify that your cluster is properly configured.

Note: If you must initialize your cluster configuration from the controller, you must manually reset the following settings because backup configurations do not save these settings:

- SCSI initiator IDs
- Host name
- Partner name.

Chapter 3. Configuring ServeRAID controllers for clustering

To install the IBM ServeRAID Cluster Solution, you must first install the BIOS and firmware. Refer to Chapter 6, “Updating the ServeRAID BIOS, firmware, and software code for clustering,” on page 37. Then, you must configure the ServeRAID controller, define logical drives, and set up merge group numbers for your servers.

- If you are installing the IBM ServeRAID Cluster Solution using Microsoft Cluster Service, refer to “Configuring the controller for Microsoft Cluster Service using Windows NT 4.0 or Windows 2000”.
- If you are installing the IBM ServeRAID Cluster Solution using Vinca and Novell IntranetWare, refer to “Configuring the controller for Vinca High Availability for Novell IntranetWare or Novell High Availability Server” on page 12.

Configuring the controller for Microsoft Cluster Service using Windows NT 4.0 or Windows 2000

If you are installing the IBM ServeRAID Cluster Solution by using Microsoft Cluster Service on Windows NT 4.0 or Windows 2000, follow these steps:

1. If necessary, perform any operating system shutdown procedure on *both* servers that you are going to use in the cluster.
2. Turn off *both* servers.
3. If you have not done so, connect any shared external drive enclosures to both servers; then, turn on the enclosures.
4. Follow the procedure “Configuring a controller for clustering” on page 12 for each shared controller in Server A, then return here.
5. Continue with the procedure “Configuring ServeRAID arrays and a logical drive” on page 14 for Server A, then return here.
6. Before you configure Server B, for complicated clusters you might want to use **View shared drives** to identify Online drives in Server B as Reserved.
 - a. From Server A, right-click on the controller, select **View shared drives**.
 - b. Click the checkbox **Enable view shared drives**.
 - c. Set the channel and SCSI ID for the drives that are identified as Ready on Server A and Online on Server B.

Drives used in the cluster will display as **Reserved** on Server B.

7. Follow the procedure “Configuring a controller for clustering” on page 12 for each shared controller in Server B, then return here.
8. Remove the *IBM ServeRAID Support* CD from the CD-ROM drive.
9. Restart Server A and Server B.
10. If you are installing the IBM ServeRAID Cluster Solution using Microsoft Cluster Service using Windows NT 4.0, continue to “Installing the Microsoft Cluster Service software on Windows NT 4.0” on page 17.

If you are installing the IBM ServeRAID Cluster Solution using Microsoft Cluster Service using Windows 2000, continue to “Installing the Microsoft Cluster Service software on Windows 2000” on page 24.

Note: After you have completed installing the IBM ServeRAID Cluster Solution, use the **Validate cluster** action in the ServeRAID Manager program to verify the cluster configuration. See “Validating a cluster” on page 33 for more information.

Configuring the controller for Vinca High Availability for Novell IntranetWare or Novell High Availability Server

If you are installing the IBM ServeRAID Cluster Solution by using Vinca High Availability for Novell IntranetWare or Novell High Availability Server, following these steps:

1. If necessary, perform any operating system shutdown procedure on *both* servers that you are going to use in the cluster.
2. Turn off *both* servers.
3. If you have not done so, connect any shared external drive enclosures to both servers; then, turn on the enclosures.
4. Follow the procedure “Configuring a controller for clustering” for each shared controller in Server A.
5. Continue with the procedure “Configuring ServeRAID arrays and a logical drive” on page 14 for Server A.

Note: In a typical Novell environment, create the logical drives on the server that will maintain primary ownership. For example, create one RAID level-5 logical drive on Server A and one RAID level-5 logical drive on Server B.

6. Follow the procedure “Configuring a controller for clustering” for each shared controller in Server B.
7. Continue with the procedure “Configuring ServeRAID arrays and a logical drive” on page 14 for Server B.

Note: If you define an online drive or hot-spare drive for Server A, you cannot use that drive when creating arrays or hot-spare drives on Server B.

8. Restart Server A and Server B.
9. Install Novell NetWare on your servers. For installation information, refer to the documentation that comes with your Novell NetWare software package.
10. Install the Vinca IntranetWare software on your servers. For installation information, refer to the documentation that comes with Vinca High Availability for NetWare and IBM ServeRAID II SCSI Controller.

Configuring a controller for clustering

To configure a controller, do the following.

Note: For complete information on using the ServeRAID Manager program, refer to the ServeRAID Manager online help.

1. Insert the *IBM ServeRAID Support* CD (or the configuration CD that was included with your server) into the CD-ROM drive; then, turn on the server. The ServeRAID Manager program starts.

Note: If the ServeRAID Manager program opens in Configuration mode, a window similar to the following opens. Click **Cancel** to go to Information mode and continue to the next step.

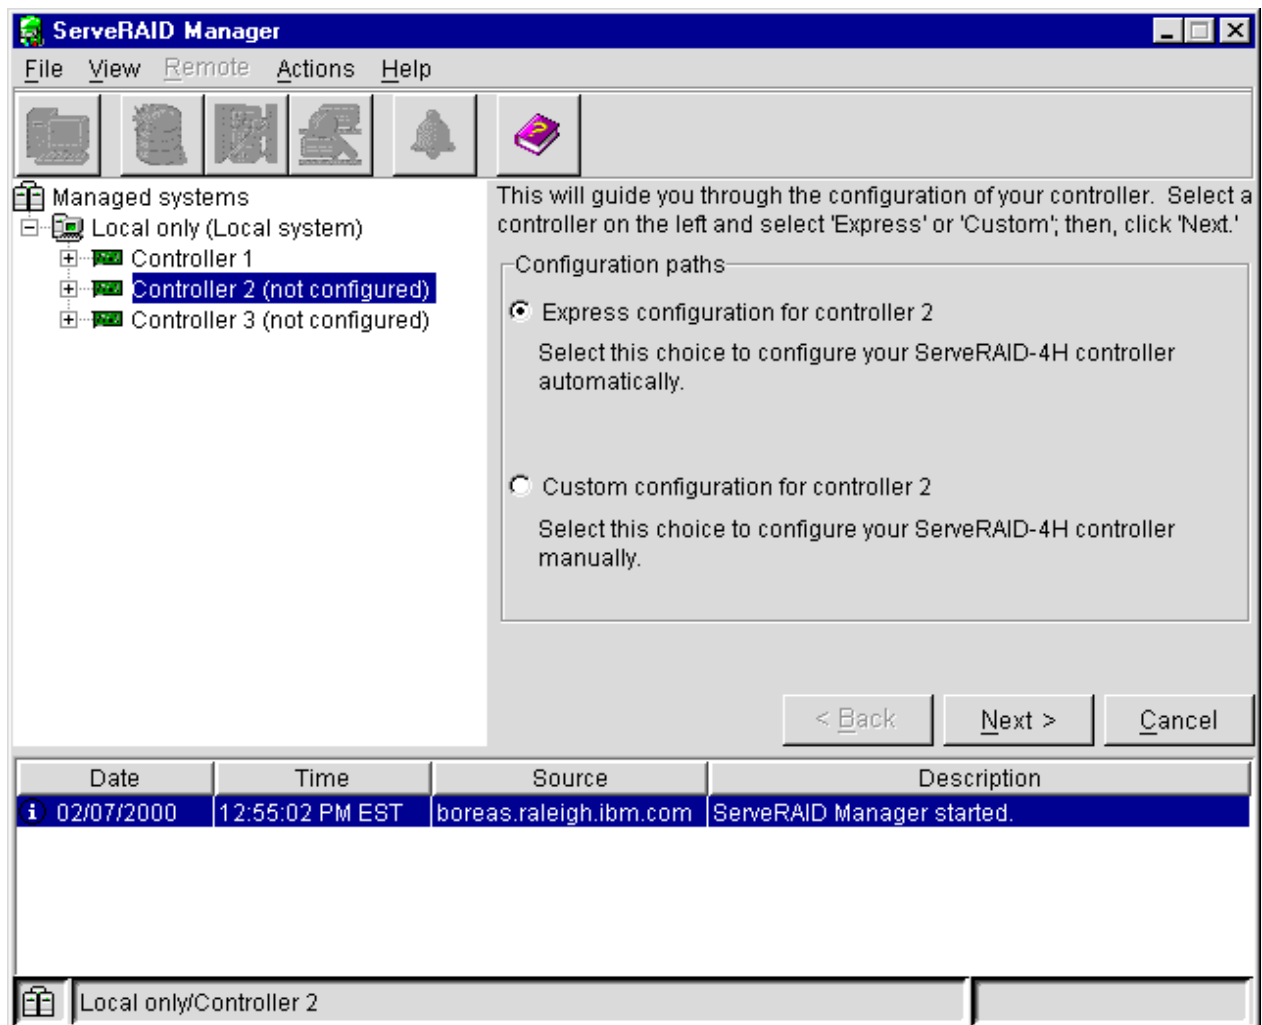


Figure 3. Configuration mode window

2. In the Main Tree, click the ServeRAID controller that you want to configure.

Attention

The ServeRAID controller should be initialized only one time, when the controller is first installed. Initializing the controller removes all configuration information from NVRAM. This information includes the SCSI IDs, the Controller host ID and the Cluster partner host ID, and the array information. This point is especially important if the Windows NT or Windows 2000 system startup drive is attached to the IBM ServeRAID controller. If the initialization occurs after the startup drive array has been created, the array information will be removed and the system will no longer start in Windows NT or Windows 2000.

If you already have logical drives defined that you want to keep, go to step 6 on page 14. Completing step 1 on page 12 through step 5 on page 14 will cause loss of data by deleting all logical drives already defined.

3. Click **Actions** → **Restore to factory default settings**.
4. If a confirmation window opens, click **Yes**. When the operation is complete, a message displays at the bottom of the window indicating that the restore to factory default settings was successful.

5. In the Main Tree, click the ServeRAID controller that you are configuring.

6. Click  on the tool bar. The Configure for clustering window opens.

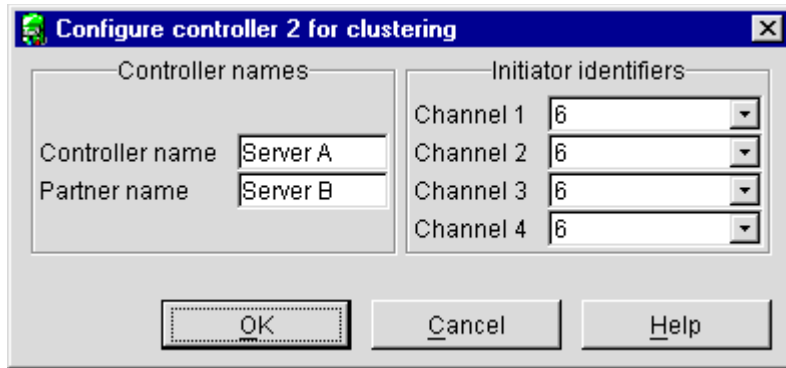


Figure 4. Configure for clustering window

7. Provide the following information.

Notes:

- a. The controller name and partner name are case sensitive.
- b. If your configuration will contain multiple ServeRAID controllers pairs, each controller name must be unique.

Controller name Enter a name to identify the first controller, for example, Server A.

Partner name Enter a name that you will use to identify the second controller, for example, Server B.

SCSI initiator IDs The system presets these IDs to 7. For Server B, do *not* change these values. For Server A, set each ID to 6 for shared channels.

Note: The ServeRAID Manager program does not permit you to enter numbers other than 6 and 7.


8. Click **OK**. The controller is now configured for clustering.

Configuring ServeRAID arrays and a logical drive

Important

- **Create only one logical drive for each array when defining shared logical drives for use in a cluster.**
- Hot-spare and standby hot-spare drives must be defined only on a shared channel, and each controller must have its own hot-spare or standby hot-spare drive defined. This means that you must not define the same hot-spare or standby hot-spare drive on both controllers in the pair.

To configure arrays and a logical drive, do the following.

1. Click  on the tool bar. The Configuration mode window opens.
2. In the Main Tree, click the ServeRAID controller that you want to configure.

3. Click the **Custom configuration** button.
4. Click **Next**. The Create arrays window opens.
5. Using the right mouse button, click the ready drive or SCSI channel icons in the Main Tree to select the drives that you want to add to your arrays or delete from your arrays; then, select a choice from the pop-up list.

Notes:

- a. You must use all the ready drives when creating your arrays (that is, all ready drives in the Main Tree must be online).
 - b. Within the same array, you must have two or more physical drives for RAID level-1 and three or more physical drives for RAID level-5.
 - c. For Windows NT and Windows 2000, you must create at least one logical drive, the quorum drive, as RAID level-1 to store the quorum resource information. This drive can be used to store programs and data.
 - d. For improved reliability, do not create arrays that span channels.
 - e. Hot-spare drives are *not* shared between controllers in a cluster pairing. If you want hot-spare protection, *each* ServeRAID controller must have a hot-spare drive defined.
 - f. You *must* connect hot-spare drives *only* to shared channels. This is in case the hot-spare drive replaces a failed shared drive. Both servers must be able to access their uniquely-defined hot-spare drive in case a failover occurs *after* the shared drive failure.
6. After you select the ready drives for your arrays, click **Next**. The Create logical drives window opens.

Note: If you change your mind, you can remove a specific drive from a newly defined array or delete an entire newly defined array. To do this:

- a. Click **Back**.
 - b. Using the right mouse button, click the specific drive or Array icon in the Main Panel on the right.
 - c. Click **Remove from new array** or **Delete new array**.
7. Select a RAID level for the logical drive from the RAID pull-down menu in the Main Panel.

Notes:

- a. If you have four or more ready drives in an array, the recommended RAID level is RAID level-5E.
 - b. Only one logical drive can be defined for each array.
8. Click **Next**. The Configuration summary window opens.
 9. Review the information that displays in the Configuration summary window. Click **Apply** to accept and save the configuration.

Notes:

- a. To change the configuration, click **Modify arrays** or **Modify logical drives**.
 - b. For Windows NT and Windows 2000 clusters, define all shared logical drives on Server A.
10. Click **Yes** to continue working with your ServeRAID configuration. After creating the arrays and logical drives, you must assign merge group numbers to the shared logical drives.

Note: A merge group number uniquely identifies each logical drive, even when the logical drive moves between servers.

11. In the Main Tree, click the ServeRAID controller that contains the logical drives.

Notes:

- a. Valid merge group numbers for non-shared logical drives are 206 and 207. For example, set all non-shared drives on Server A to merge group number 206 and all non-shared drives on Server B to merge group number 207.
- b. The merge group number must be unique for each logical drive that is shared by the controller pair. For example, if you have a logical drive with a merge group number 1 on Server A, you cannot assign the same group number to a logical drive on Server B. Valid merge group numbers for shared logical drives are in the range 1-8.

12. Click  on the tool bar. The Configure for clustering window opens.

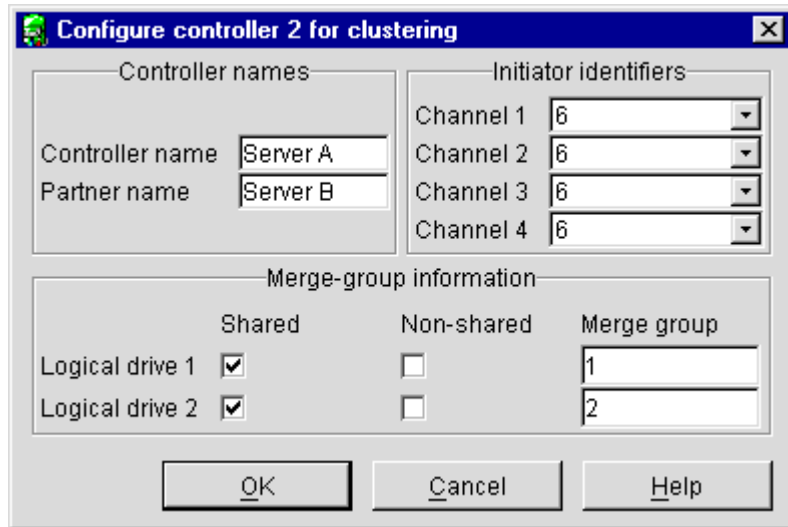


Figure 5. Configure for clustering window

13. Click the **Shared** check boxes for all the logical drives.
14. Do not change the merge group number from the default that is provided.
15. Record the controller name, partner name, and merge group numbers and store them in a safe place. You will need this information to recover the configuration if the ServeRAID controller fails.

Note: You can save this information to a file by using the ServeRAID Manager program:

- a. In the Main Tree, click the server.
- b. Click **Actions** → **Save printable configuration**.

16. Click **OK**.
17. To exit the ServeRAID Manager program, click **File** → **Exit**.
18. Remove the *IBM ServeRAID Support* CD from the CD-ROM drive.
19. Shut down the server.

Chapter 4. Installing the Microsoft Cluster Service software

After configuring the ServeRAID controller, defining logical drives, and setting up merge group numbers, you can install Microsoft Cluster Service (MSCS) on Windows NT 4.0 or Windows 2000. Refer to the following instructions:

- “Installing the Microsoft Cluster Service software on Windows NT 4.0”
- “Installing the Microsoft Cluster Service software on Windows 2000” on page 24

Important

Before beginning the procedure, consider the following:

- Read “ServeRAID considerations” on page 8, which contains important information on ServeRAID considerations.
- Read “Configuring the controller for Microsoft Cluster Service using Windows NT 4.0 or Windows 2000” on page 11 for configuring Server A and Server B for the high-availability cluster solution.

Installing the Microsoft Cluster Service software on Windows NT 4.0

Important

Before beginning the procedure, ensure that you are using the *IBM ServeRAID Device Driver* diskette (version 4.0, or later) during the installation of the Microsoft Windows NT Server Enterprise Edition.

To install the Microsoft Cluster Service, complete the following procedure:

1. If you do not have Microsoft Windows NT Server Enterprise Edition installed already, install it with Service Pack 3 (or later) on both servers, but do not install the Microsoft Cluster Service software. When the system prompts you, deselect the option to install MSCS; it will be installed later in this procedure.

Note: The *Microsoft Windows NT Server Enterprise Edition* CD contains Service Pack 3.

For information on installing Windows NT Server Enterprise Edition, refer to documentation that is included with it. When you are finished, return here and continue with this procedure.

2. Use the ServeRAID Manager program to scan for new drives.
3. On Server A, run the utility IPSHAHTO.EXE from the *IBM ServeRAID Windows NT Cluster Solutions Diskette* or the *IBM ServeRAID Support CD*.
 - If you are using the diskette, type the following command from a Windows NT command line:
`a:\support\ipshahto`
where *a* is the diskette drive letter.
 - If you are using the CD, type the following command from a Windows NT command line:
`d:\programs\winnt\cluster\support\ipshahto`
where *d* is the CD-ROM drive letter.

Running the program makes all the shared logical drives available for access by Server A.

4. Using the Windows NT Disk Administrator, format all shared logical drives as NTFS (NT File System), assign volume labels, and specify drive letters for each.

Important

Before creating the ServeRAID logical cluster resource, you must assign a permanent (sticky) drive letter to the shared logical drive on both servers.

When configuring a cluster, you must assign a permanent drive letter because the two servers must have the same drive letters on each shared logical drive. If you do not assign permanent drive letters, Windows NT might assign different drive letters on each server. You must assign a permanent drive letter to ensure Windows NT assigns the same drive letter to the partition on both servers.

Notes:

- a. It is helpful if part of the volume label that you assign contains the drive letter that you assigned to the drive, such as DRIVE E.
 - b. If you want to use the disk in a cluster, create only one partition on the disk. Make the partition the full size of the disk and format the partition to NTFS.
5. On Server B, run the same utility, IPSHAHTO.EXE, to make all shared logical drives available for access by Server B.

Note: Use the volume label as a guide to make sure you assign the same letters on Server B as you assigned on Server A.

6. Using the Windows NT Disk Administrator, assign the same drive letters that you assigned to each drive on Server A.
7. Run IPSHAHTO.EXE from Server A again, so that control of all shared logical drives returns to Server A.
8. Install the Microsoft Cluster Service (MSCS) software.

Note: The directory name and path where the MSCS software is installed must be the same on both servers.

To start the installation, do the following.

Important

You must use the *localquorum* option when installing the Microsoft Cluster Service (MSCS) software with ServeRAID controllers.

- a. Insert the *Windows NT Server Enterprise Edition Component CD* into the CD-ROM drive on Server A, type the following command at the command line and press Enter.

```
d:\MSCS\CLUSTER\I386\SETUP /localquorum
```

where *d* is the CD-ROM drive letter.
Note: You *must* type the parameter "localquorum" in lowercase.
 - b. Install the Microsoft Cluster Service software following the instructions in the documentation that is included with the MSCS software.
Server A will restart when the installation is completed.
9. To make Server B *join* the cluster, do the following.

- Note:** The directory name and path where the MSCS software is installed must be the same on both servers.
- a. Insert the *Windows NT Server Enterprise Edition Component CD* into the CD-ROM drive on Server B. Type the following command (from either a command line or Add/Remove Programs in the Control panel) and press Enter:
`d:\MSCS\CLUSTER\I386\SETUP`
where *d* is the CD-ROM drive letter.
 - b. Install the Microsoft Cluster Service software following the instructions in the documentation that is included with the MSCS software.
Server B will restart when the installation is completed.
10. Install the IBM ServeRAID Windows NT Cluster Solution program by doing the following:
- a. On Server A, run the utility SETUP.EXE program from the *IBM ServeRAID Windows NT Cluster Solution* diskette or *IBM ServeRAID Support CD* and follow the instructions that appear on the screen.
 - If you are using the diskette, type the following command from a Windows NT command line:
`a:\setup`
where *a* is the diskette drive letter.
 - If you are using the CD, type the following command from a Windows NT command line:
`d:\programs\winnt\cluster\setup`
where *d* is the CD-ROM drive letter.
 - b. When the program has completed, run the SETUP.EXE program on Server B.
After the utility has been installed on both servers, all shared ServeRAID groups and disk resources will be automatically created by the IBM ServeRAID Cluster Solution setup program and then brought online. Server A is the current owner of these groups.
11. Move the quorum resource from *localquorum* to a ServeRAID disk resource by doing the following:
- a. On Server A, run the Microsoft Cluster Administrator program and click the cluster name to select the cluster you are using.
 - b. Verify that the same server (for example, Server A) owns the ServeRAID disk resources and the localquorum resource and that the quorum resource is moved to a RAID level-1 logical drive.

Note: It is recommended that you set the quorum resources on a RAID level-1 logical drive. Do *not* set the quorum resources on a RAID level-5 logical drive.
 - c. Using the right mouse button, click the cluster name to open the cluster.
 - d. Select **Properties**; then, select the **Quorum tab**.
 - e. In the **Quorum Resource** field, select a ServeRAID shared disk.
 - f. Click **OK**.
 - g. Delete the Local Quorum cluster resource and cluster group.
Refer to the *Microsoft Cluster Solution Administrator's Guide* for more information.
12. If required, install Service Pack 4 or later.
- Note:** If you installed Service Pack 4 or later previously, you must reinstall it now.

You can now move the shared disk resources manually using the Microsoft Cluster Administrator program or automatically during failover.

Creating additional ServeRAID disk resources on a Windows NT 4.0 ServeRAID cluster

You can use a ServeRAID logical disk that has not already been designated as a cluster resource, or you can create a new array and logical disk on Server A, using the ServeRAID Manager program. For more information about creating new arrays and logical drives, refer to "Configuring ServeRAID controllers" in the *IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference*.

Note: Be sure to assign a unique shared merge group number to the new shared logical drive.

Important

Before creating the ServeRAID logical cluster resource, you must assign a permanent (sticky) drive letter to the shared logical drive on both servers.

When configuring a cluster, you must assign a permanent drive letter because the two servers must have the same drive letters on each shared logical drive. If you do not assign permanent drive letters, Windows NT might assign different drive letters on each server. You must assign a permanent drive letter to ensure Windows NT assigns the same drive letter to the partition on both servers.

1. On Server A, do the following:
 - a. Click **Start** in the lower-left corner of the screen; then, click **Programs** → **Tools** → **Administrative Tools** → **Disk Administrator**. The Disk Administrator window opens.
 - b. Click the partition that you want to assign a permanent drive letter.
 - c. Click **Tools** → **Assign Drive Letter**. The Assign drive letter window opens.
 - d. In the Assign drive letter pull-down menu, click the drive letter for the partition.
 - e. Click **OK**. Windows NT permanently assigns the selected drive letter to that partition, even if you add new partitions or rearrange your controllers.
 - f. To exit the Disk Administrator, click **Partition** → **Exit**.
2. On Server B, do the following:
 - a. Using the ServeRAID Manager program, scan for new or removed ready drives.
 - b. Using the IPSHAHTO.EXE program, move the newly defined logical drive.
 - c. Using the Windows NT Disk Administrator program, assign the same permanent (sticky) drive letter as in step 1.

Perform the following operation on the cluster server that currently owns the shared logical disk drive. The cluster server that owns the logical disk drives is the server on which you last ran the IPSHAHTO.EXE program.

1. On the server owning the newly created logical drive, run Cluster Administration program and open a connection to the installed cluster.
2. To create the ServeRAID disk resource in an existing group, ensure that the cluster server that owns the shared logical disk drive is the owner of that group. If the cluster server is not the owner of the group, use **Move Group** to make this cluster server the owner of the group where you will create the ServeRAID disk resource. Then, create the ServeRAID disk resource.

3. Bring the new ServeRAID disk group online. If you create a new ServeRAID group, the cluster server that owns the shared logical disk drive will be the owner of this group by default.
4. Create a ServeRAID disk resource as part of the new group.
5. Bring the ServeRAID disk group online.

Deleting and re-creating a ServeRAID logical drive resource

This section describes how to delete and re-create a ServeRAID logical drive resource without deleting a ServeRAID logical drive.

To delete a resource or resource group:

1. If the resource or resource group is not currently owned by the cluster server from which you will delete the group, move the resource or resource group to that cluster server.

Note: The cluster service must be running on both systems.

2. To delete the resource, run the Microsoft Cluster Administration program from the cluster server that currently owns the resource or resource group.

Note: You can also run the Microsoft Cluster Administration program from a remote administration system.

Use the normal MSCS procedure for deleting a resource or a group. Refer to the documentation that is included with the MSCS package for instructions.

To create a resource:

1. If the group that will own the resource you are creating is not currently owned by the cluster server on which you want to create the resource, move the group to the cluster server.
2. To create the resource, run the Microsoft Cluster Administration program from the cluster server that currently owns the resource group.

Create a new ServeRAID logical disk. Refer to "Creating additional ServeRAID disk resources on a Windows NT 4.0 ServeRAID cluster" on page 20.

Uninstalling MSCS on ServeRAID servers

Note: *Do not* use the uninstallation procedure if you are upgrading MSCS on ServeRAID servers. To upgrade MSCS on ServeRAID servers, follow the instructions in "Updating ServeRAID software in an MSCS cluster environment" on page 22.

To uninstall MSCS on ServeRAID servers, you must perform the following procedure on all servers in your cluster:

1. Select one of the cluster servers from which to uninstall MSCS.
2. Click **Start** in the lower-left corner of the screen; then, click **Settings** → **Control Panel**. The Windows NT control panel window opens.
3. Double-click **Services**. The Services window opens.
4. Select **Cluster Service** from the list of services; then, click **Stop**.
5. Click **Close** to exit the Services window. For more information, refer to the documentation that is included with the cluster services.
6. In the Windows NT control panel, double-click **Add/Remove Programs**.
7. Select **Microsoft Cluster Service** from the list of installed components; then, click **Add/Remove**.

Note: Ignore the warnings regarding the removal of IBM ServeRAID Logical Disks.

8. Follow the on-screen prompts by clicking **OK** or **Yes**.
9. After the uninstallation of MSCS is complete, restart the server when prompted.
10. After the server has restarted, click **Start**; then, click **Settings** → **Control Panel**.
11. Double-click **Add/Remove Programs**.
12. Select **IBM ServeRAID Windows NT Cluster Solution** from the list of installed components; then, click **Add/Remove**.
13. Repeat step 1 on page 21 through step 12 for the other server in the cluster.

Updating ServeRAID software in an MSCS cluster environment

Note: If you are upgrading your ServeRAID software after downloading it from the IBM Support Web site, perform this procedure. If you are installing your software using the *IBM ServeRAID Support* CD, refer to "Configuring ServeRAID controllers" in the *IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference* for instructions.

Important

This upgrade procedure enables the cluster to remain online at all times. To upgrade the cluster, perform the following steps on one server at a time. When all steps are completed for Server A and both servers are reported in the Microsoft Cluster Administrator program as *up* and all resources are *online*, repeat the steps for Server B.

Note: Use these instructions on a cluster running Microsoft Cluster Service (MSCS) Release 1.00 and IBM ServeRAID Cluster Solution.

Step 1: Upgrading the IBM ServeRAID Manager program

To upgrade the ServeRAID Manager program, insert the *IBM ServeRAID Support* CD into the CD-ROM drive. The Setup program starts automatically.

Step 2: Upgrading the IBM ServeRAID Microsoft Windows NT Cluster Solution

To upgrade the IBM ServeRAID Microsoft Windows NT Cluster Solution, do the following on the server you want to upgrade:

Run the utility SETUP.EXE program from the *IBM ServeRAID Windows NT Cluster Solution* diskette or *IBM ServeRAID Support* CD and follow the instructions that appear on the screen.

- If you are using the diskette, type the following command from a Windows NT command line:

```
a:\setup
```

where *a* is the diskette drive letter.

- If you are using the CD, type the following command from a Windows NT command line:

```
d:\programs\winnt\cluster\setup
```

where *d* is the CD-ROM drive letter.

Step 3: Upgrading the IBM ServeRAID Windows NT device driver

To upgrade the IBM ServeRAID Windows NT device driver on a server in the cluster, do the following:

1. Insert the *IBM ServeRAID Support* CD, version 4.0 or later into the CD-ROM drive.
2. Click **Start** in the lower-left corner of the screen; then, click **Settings** → **Control Panel**. The Windows NT control panel window opens.
3. Double-click **SCSI-Adapters**.
4. Click the **Drivers** tab.
5. Click **Add**.
6. Click **Have Disk**.
7. For the path, type:
`d:\programs\winnt\driver\nt4_0`
where *d* is the CD-ROM drive.
8. Click **OK**.
9. Click **OK** when the IBM ServeRAID controller is highlighted.
10. Shut down and restart the server from the CD.

Step 4: Upgrading the IBM ServeRAID BIOS and firmware

To upgrade the IBM ServeRAID BIOS and firmware on a server in the cluster, do the following:

1. If you have not already done so, start the server with the *IBM ServeRAID Support* CD, version 4.0 or later.
2. Flash the BIOS and download the firmware by clicking **Update** when prompted. If no update is required, the ServeRAID Manager program starts.
3. When the download is complete, remove the CD from the CD-ROM drive and press Ctrl+Alt+Del to restart (reboot) the server.
4. Allow the server to rejoin the cluster and when the server is reported in the Microsoft Cluster Administrator program as “Up” and all resources are online, the upgrade for this server is complete.

Step 5: Upgrading the software for the other server in the cluster

To upgrade the software for the *other* server in the cluster, do the following:

1. Ensure that Server A is online in the cluster.
2. For Server B, repeat “Step 1: Upgrading the IBM ServeRAID Manager program” on page 22 through “Step 4: Upgrading the IBM ServeRAID BIOS and firmware”.

You have completed the steps necessary for upgrading the IBM ServeRAID Cluster Solution.

Installing the Microsoft Cluster Service software on Windows 2000

Important

Before beginning the procedure, consider the following:

- Ensure that you are using the *IBM ServeRAID Device Drivers* diskette (version 4.0 or later) during the installation of Windows 2000 Advanced Server Edition or Windows 2000 Datacenter Server Edition.
- To create a ServeRAID cluster, you must install the Windows 2000 operating system on an NTFS partition, and any shared cluster disks must be formatted as NTFS.
- While installing Windows 2000, do not install Microsoft Cluster Service. You will install MSCS later.
- To install the ServeRAID device driver version 4.0 (or later), refer to "Installing device drivers after installing Windows 2000" in the *IBM ServeRAID-4x Ultra 160 SCSI User's Reference*.

Installing the ServeRAID Manager program

Before installing MSCS, you must install the ServeRAID Manager program:

1. If the server is not running, turn on the server.
2. Log on to Windows 2000.
3. Insert the *IBM ServeRAID Support* CD (or the configuration CD that was included with your server) into the CD-ROM drive of Server A. The ServeRAID Manager setup program starts.
4. Click **Next**.
5. Click **Yes** to the Software License Agreement.
6. When prompted to select an operating systems, select Windows 95/98/NT; then, click **Next**.
7. When prompted to install Windows NT Service, select **Yes**; then, click **Next**.
8. Select the **Destination Folder**.
9. Click **Next**.
10. After setup completes copying the files, click **Finish**.

Installing the Microsoft Cluster Service

To install the Microsoft Cluster Service, complete the following procedure:

1. Use the ServeRAID Manager program to scan for new drives on all ServeRAID controllers.
2. On Server A, run the utility IPSHAHTO.EXE from the *IBM ServeRAID Windows NT Cluster Solutions Diskette* or the *IBM ServeRAID Support* CD.
 - If you are using the diskette, type the following command from a Windows 2000 command line:
`a:\support\ipshahto`
where *a* is the diskette drive letter.
 - If you are using the CD, type the following command from a Windows 2000 command line:
`d:\programs\winnt\cluster\support\ipshahto`
where *d* is the CD-ROM drive letter.

Running the program makes all the shared logical drives available for access by Server A.

3. Using the right-mouse button, click **My Computer** on the desktop; then, click **Manage**. The Computer Management window opens.
4. On the left side of the window, click to expand the **Storage** icon; then, click **Disk Management**. The Disk Signature and Upgrade Wizard window opens.
5. To assign disk signatures for the shared disks, click **Next**.
6. Microsoft Cluster Service using Windows 2000 requires you to set shared logical drives to Basic Disk. To do this, click **Cancel** in the Upgrade Disks panel of the window.
7. Format all shared logical drives to NTFS (NT File System), assign volume labels, and specify drive letters for each.

Important

Before creating the ServeRAID logical cluster resource, you must assign a permanent (sticky) drive letter to the shared logical drive on both servers.

When configuring a cluster, you must assign a permanent drive letter because the two servers must have the same drive letters on each shared logical drive. If you do not assign permanent drive letters, Windows 2000 might assign different drive letters on each server. You must assign a permanent drive letter to ensure Windows 2000 assigns the same drive letter to the partition on both servers.

Notes:

- a. It is helpful if part of the volume label that you assign contains the drive letter that you assigned to the drive, such as DRIVE E.
- b. If you want to use the disk in a cluster, create only one partition on the disk. Make the partition the full size of disk and format the partition to NTFS.
- c. Make sure that all shared logical drives are the type Basic Disk. If they are marked Dynamic Disk, change them to Basic Disk by doing the following:
 - (1) With the right-mouse button, click the Disk icon at the bottom of the Disk Management window (see Figure 6 on page 26).
 - (2) Click **Revert to Basic Disk**.

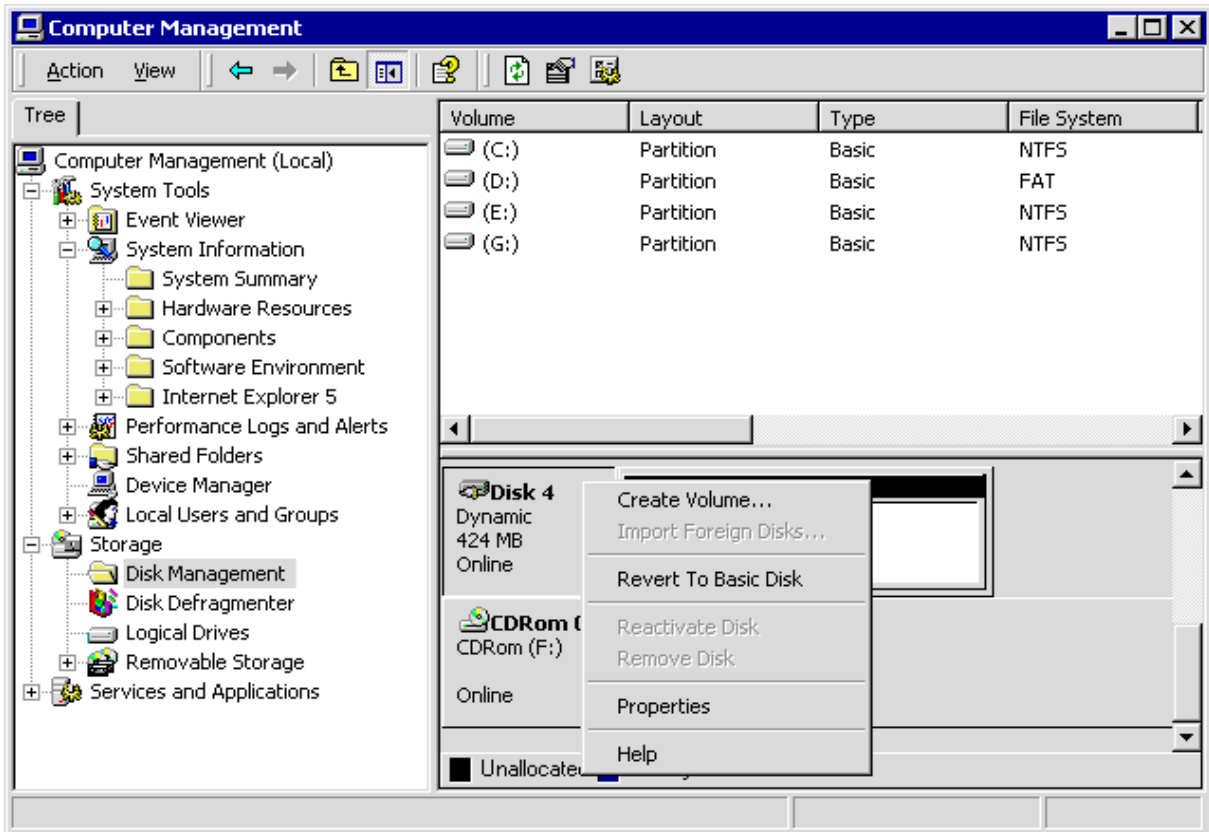


Figure 6. Changing a shared logical drive to Basic Disk

8. On Server B, use the ServeRAID Manager program to scan for new drives on all ServeRAID controllers.
9. On Server B, run the same utility, IPSHAHTO.EXE, to make all shared logical drives available for access by Server B.

Note: If all logical drives do not display as Okay, there was a problem bringing drives over from the other server. To diagnose the problem, use the log files or ServeRAID Manager to determine if the drives are Critical or Offline.

10. Using the Windows 2000 Disk Management Snap-in Computer Management program, scan for new drives and assign them the same drive letters that you assigned to each drive on Server A.

Notes:

- a. Use the disk volume label as a guide to make sure you assign the same drive letters on Server B as you assigned on Server A.
- b. You might need to press the F5 key multiple times to update the displayed information. Wait until all logical drive letters appear before attempting to change the property of any drives.
- c. If the Disk Signature Wizard prompts you to write signatures on drives, the drives have not been correctly recognized. Do the following:
 - (1) On Server A, run the IPSHAHTO utility.
 - (2) After Server A has successfully brought over all drives, run the IPSHAHTO utility on Server B.
 - (3) Repeat step 10.

11. Run IPSHAHTO.EXE from Server A again, so that control of all shared logical drives returns to Server A.
12. Using the Windows 2000 Disk Management Snap-in Computer Management program on Server A, verify that the shared logical drives display with the correct size, volume label, and drive letter.

Note: The drive letters for the shared disks must be the same on each server in the cluster.
13. Ensure that both servers are turned on.

Note: If you have a two-node cluster that has a private network consisting of the two servers directly connected to each other using an Ethernet cross-over cable, make sure that *both* servers are turned on before you proceed.

Failure to have both servers turned on will cause Windows 2000 to mark the network card as failed and you will not be able to select it as the private network interface during cluster setup.
14. Insert the *IBM ServeRAID Windows NT Cluster Solution* diskette or *IBM ServeRAID Support* CD into the Server A diskette or CD-ROM drive; then, type the following command and press Enter:
 - If you are installing from diskette, type the following command from a Windows 2000 command line and press Enter:


```
a:\setup
```

where *a* is the diskette drive letter.
 - If you are installing from CD, type the following command from a Windows 2000 command line and press Enter:


```
d:\programs\winnt\cluster\setup
```

where *d* is the CD-ROM drive letter.

The SETUP.EXE program copies some files and launches Windows Component Wizard.
15. In the Windows Component Wizard, click the check box to install the Cluster Service component; then, click **Next**.
16. When prompted, insert the CD labeled *Windows 2000 Advanced Server CD-ROM* into the CD-ROM drive; then, click **OK**. The Microsoft Cluster Service Configuration Wizard opens.

Note: To configure the Microsoft Cluster Service, follow the instructions in the Microsoft Cluster Service documentation.
17. After you install and configure the cluster, IBM ServeRAID Cluster Configuration Wizard opens and applies ServeRAID-specific configurations to the cluster. When prompted, select a shared ServeRAID logical drive on which to set quorum.

Note: Set the quorum resources on a RAID level-1 logical drive. Do *not* set the quorum resources on a RAID level-5 logical drive.
18. Repeat step 13 through step 16 for Server B. The cluster installation is complete and the cluster is ready for use.

Creating additional ServeRAID disk resources on a Windows 2000 ServeRAID cluster

You can use a ServeRAID logical disk that has not already been designated as a cluster resource, or you can create a new array and logical disk on Server A, using the ServeRAID Manager program. For more information about creating new arrays and logical drives, refer to "Configuring ServeRAID controllers" in the *IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference*.

Note: Be sure to assign a unique shared merge group number to the new shared logical drive.

Important

Before creating the ServeRAID logical cluster resource, you must assign a permanent (sticky) drive letter to the shared logical drive on both servers.

When configuring a cluster, you must assign a permanent drive letter because the two servers must have the same drive letters on each shared logical drive. If you do not assign permanent drive letters, Windows 2000 might assign different drive letters on each server. You must assign a permanent drive letter to ensure Windows 2000 assigns the same drive letter to the partition on both servers.

1. On Server A, do the following:
 - a. Using the right-mouse button on the desktop, click **My Computer**; then, click **Manage**. The Computer Management window opens.
 - b. Double-click **Storage**.
 - c. Click **Disk Management**. The Write Signature and Upgrade Disk Wizard opens.
 - d. Using the right-mouse button, click **Disk Management** → **Rescan Disks**.
 - e. On each of the new disks, do the following:
 - (1) Using the right-mouse button, click on a new disk; then, click **Write Signature**. The Write Signature Wizard opens. Follow and complete the wizard.
 - (2) Make sure that all shared logical drives are the type Basic Disk. If they are marked Dynamic Disk, change them to Basic Disk by doing the following:
 - (a) With the right-mouse button, click the Disk icon at the bottom of the Disk Management window.
 - (b) Click **Revert to Basic Disk**.
 - (3) Create one partition on the disk (the full size of disk), assign a drive letter, and format the partition to NTFS.

Note: To help recognize the drive on the cluster partner server, assign a descriptive volume label that includes the drive letter.

2. On Server B, do the following:
 - a. Using the ServeRAID Manager program, scan for new or removed ready drives on all ServeRAID controllers.
 - b. To move the newly defined logical drive to Server B, start the IPSHAHTO.EXE program by typing:
`%systemroot%\cluster\ipshahto.exe`
where `%systemroot%` is expanded by the operating system to the root directory of the operating system on Server B.
 - c. Using the Windows 2000 Disk Management Snap-in Computer Management program, assign the same drive letters that you assigned to the new drives on Server A.
Note: Use the disk volume label as a guide to make sure you assign the same drive letters on Server B as you assigned on Server A.
 - d. To create a new ServeRAID disk resource and group for each newly defined logical drive on Server B, start the MSCSCFG.EXE program by typing:
`%systemroot%\cluster\mscscfg.exe`

where `%systemroot%` is expanded by the operating system to the root directory of the operating system on Server B.

Deleting and re-creating a ServeRAID logical drive resource

This section describes how to delete and re-create a ServeRAID logical drive resource without deleting a ServeRAID logical drive.

To delete a resource or resource group:

1. To delete the resource, run the Microsoft Cluster Administration program from the cluster server that currently owns the ServeRAID resource or resource group.
 - a. The cluster service must be running on both servers.
 - b. You can also run the Microsoft Cluster Administration program from a remote administration system.
2. Use the normal MSCS procedure for deleting a resource or a group. Refer to the documentation that is included with the MSCS package for instructions.
3. After deleting the ServeRAID resource, shut down and restart the cluster server from which the resource was deleted.

Note: Restarting one cluster server will *not* bring down the entire cluster.

To create the new ServeRAID disk resource and group for the ServeRAID logical drive, start the MSCSCFG.EXE program by typing:

```
%systemroot%\cluster\mscscfg.exe
```

where `%systemroot%` is expanded by the operating system to the root directory of the operating system on Server B.

Upgrading the IBM ServeRAID Cluster Solution from Microsoft Windows NT 4.0 to Windows 2000

To upgrade from Windows NT 4.0 with Service Pack 5 (or later) to Windows 2000 Advanced Server, do the following.

Note: Before upgrading from Windows NT 4.0 Enterprise Edition, make sure you have updated Windows NT 4.0 Enterprise Edition with the latest Service Pack on both Server A and Server B.

1. On the server you want to upgrade, click **Start** → **Programs** → **Administrative Tool** → **Cluster Administrator**.
2. Move all cluster resources and groups to the other cluster node (that is, the server that is *not* running the Cluster Administrator program).
3. Insert the *IBM ServeRAID Support* CD (version 4.0 or later) into the CD-ROM drive; then, type the following command and press Enter:

```
d:\programs\winnt\cluster\setup.exe
```

where *d* is the CD-ROM drive.

4. When prompted "You are currently running Windows NT 4.0. Is this upgrade to "IBM ServeRAID NT Cluster Solution" being done as procedure to upgrade cluster to run Windows 2000?", click **Yes**.
5. The IBM ServeRAID Windows NT Solution Setup program copies the necessary Windows 2000 files; then, prompts you to upgrade the server to Windows 2000.

Note: Completing this step means you can not use this server as a Windows cluster node anymore.

6. Insert the *Microsoft Windows 2000 Advanced Server CD* or the *Microsoft Windows 2000 Datacenter Server CD* into the CD-ROM drive.

7. Run the Windows 2000 Advanced Server Setup program or the Windows 2000 Datacenter Server Setup program to upgrade this node to Windows 2000. After successfully upgrading the server, it joins the cluster.
8. Repeat step 1 on page 29 through step 7 for the other server in the cluster.

Uninstalling MSCS on ServeRAID servers

To uninstall MSCS on a ServeRAID server using Windows 2000 Advanced Server, do the following:

1. On the server you want to remove from the cluster, click **Start** → **Programs** → **Administrative Tools** → **Cluster Administrator**.
2. Move all cluster resources and groups to the other cluster server (that is, the server that is *not* running the Cluster Administrator program).
3. In the tree view, click the server from which you want to remove MSCS.
4. Using the right mouse button, click **Stop Cluster Service**; then, click **Evict Node**.
5. Click **Start** → **Settings** → **Control Panel**. The Control Panel window opens.
6. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
7. In the left panel of the window, click **Add/Remove Windows Components**. The Windows Components Wizard opens.
8. Deselect **Cluster Service** in the Components list; then, click **Next**.
9. After uninstalling the Microsoft Cluster Service completes, the Windows Component Wizard prompts you to restart the server.
10. After the server restarts, click **Start** → **Settings** → **Control Panel**. The Control Panel window opens.
11. Double-click **Add/Remove Programs**. The Add/Remove Programs window opens.
12. Select IBM ServeRAID NT Cluster Solution from the list of installed components.
13. To uninstall the IBM ServeRAID Windows NT Cluster Solution, click **Change/Remove**.

Chapter 5. Monitoring and updating an IBM ServeRAID Cluster Solution

You can use the ServeRAID Manager program to check the ownership of a physical drive on another server or to add or remove a logical drive in the cluster. The ServeRAID Manager program supports three features for cluster functions:

- View shared drives
- Validate cluster
- Configure for clustering

Note: Make sure that you installed the ServeRAID Manager program on your server. If you need to install it, refer to "Installing the ServeRAID Manager program" in the *IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference* for instructions.

Checking the ownership of a physical drive

A physical drive that is owned (and configured) on Server B will appear in the ready or defunct state on Server A. To determine if a ready drive is owned by the other server, do the following:

1. In the Main Tree, click the controller that contains the ready drives. A window similar to the following opens.

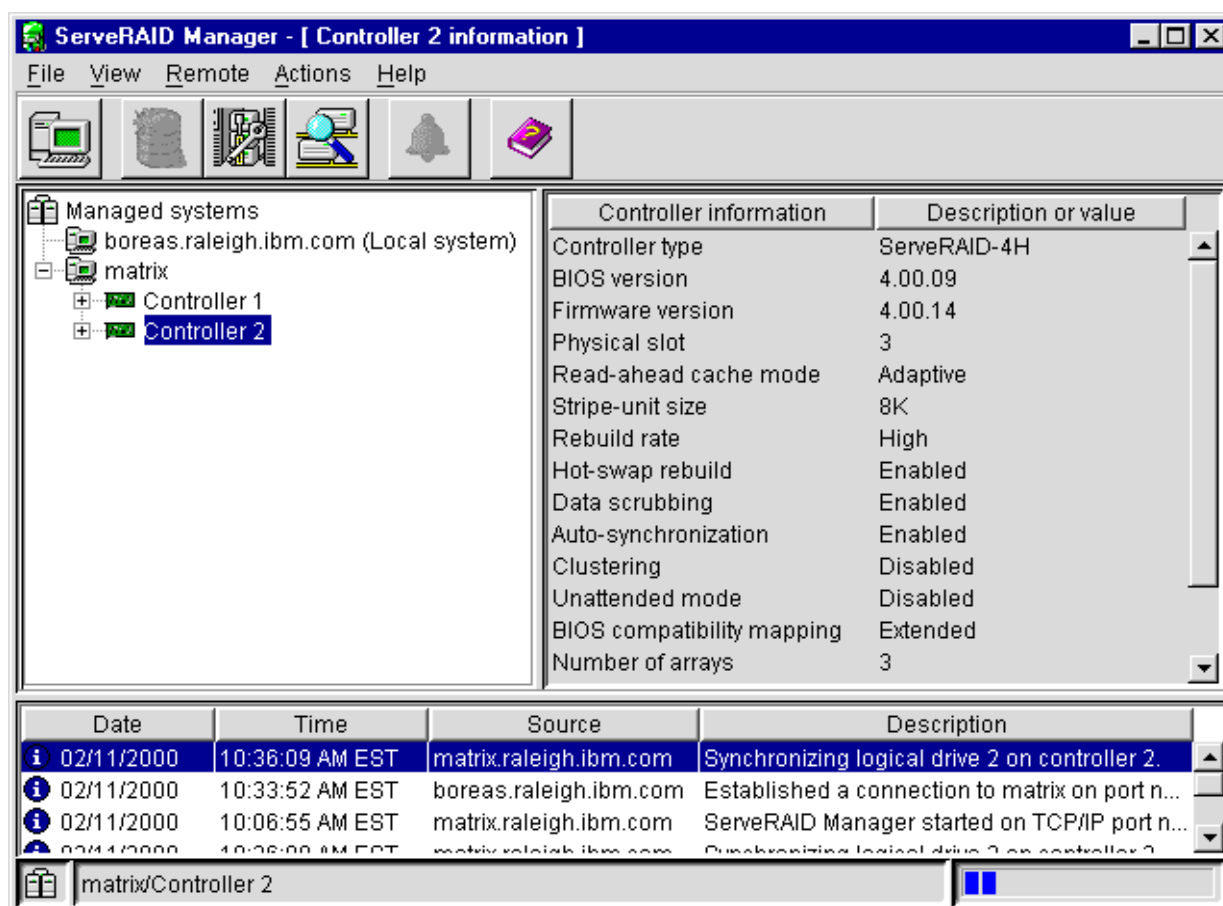


Figure 7. Information mode window

2. Click **Actions** → **Clustering actions**.

3. Click **View shared drives**. The View shared drives window opens.
4. Click **Enable view shared drives**.
5. Select the SCSI channel and SCSI ID of a physical drive on a shared channel.
6. Click **OK**. Physical drives that are available for use are labeled **Ready** in the Main Tree. If the physical drive is not available for use (that is, it is owned by the other server), the drive is labeled **Reserved** in the Main Tree.

Important

Hot-spare drives owned by the other server are not labeled **Reserved**. Be sure you do not inadvertently use a hot-spare drive in *both* servers.

Adding logical drives to a cluster

If you want to add logical drives to a cluster, do the following:

1. Install the hard disk drives in the server or shared expansion enclosure. For installation information, refer to the documentation that is included with the hard disk drive and the server.


Attention

Physical drives that are configured by one controller in the pair must appear as ready drives to the other controller in the pair. Be sure to verify this status when you add drives *after* the servers are started. Before creating logical drives, restart both servers, or scan for new drives using the ServeRAID Manager program on both servers, to make sure the added physical drives appear as ready drives on both servers. For more information, see "ServeRAID considerations" on page 8.

2. Using the ServeRAID Manager program, configure the physical drives for the ServeRAID controller. Refer to "Configuring ServeRAID controllers" in the *IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference* for instructions.

Note: If the drives will be shared, define only one logical drive per array.

3. In the Main Tree, click the controller for the cluster.

4. Click  on the tool bar. A window similar to the following opens.

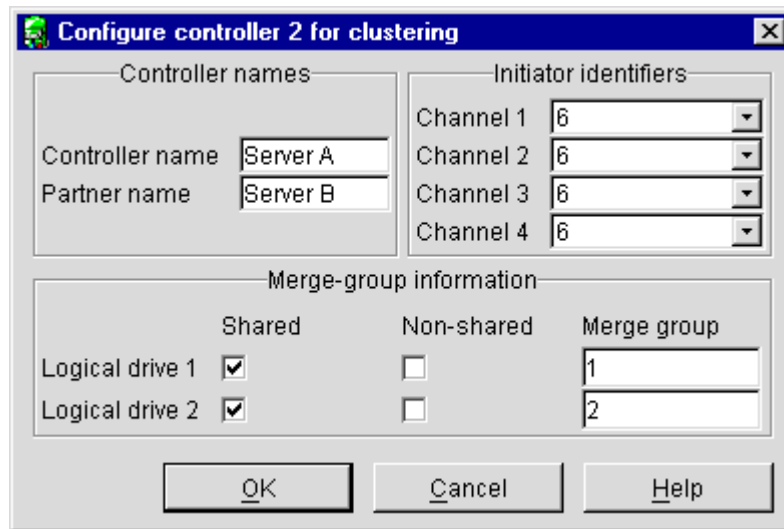


Figure 8. Configure for clustering window

5. Click **Shared** or **Non-shared** for the logical drive you have just defined.
6. Select the merge group number you want to use for the logical drive you just defined.
7. Click **OK**.
8. If you are using Windows NT 4.0, continue with “Creating additional ServeRAID disk resources on a Windows NT 4.0 ServeRAID cluster” on page 20.
If you are using Windows 2000, continue with “Creating additional ServeRAID disk resources on a Windows 2000 ServeRAID cluster” on page 27.

Validating a cluster

Use this feature to determine if your cluster is configured correctly.

Note: You must have the ServeRAID Manager program installed on all servers in the cluster for this feature to work.

To validate a cluster, do the following:

1. In the Main Tree, click a server in the cluster.
2. Click **Actions** → **Validate cluster**. A window similar to the following opens.




Figure 9. ServeRAID cluster validation window

The server you selected in step 1 on page 33 is listed in the Node A box.

Note: You cannot change the Node A system in the Cluster validation window. To select a different system, click **Cancel** to return to the Main Tree.

3. From the Node B pull-down menu, select your second server.


Note: The Node B system defaults to the first system in the Main Tree other than the Node A system. If there are no remote systems in the Main Tree, click **Cancel** to return to the Main Tree and add a remote system. You can also type the Node B name.

4. Click **Start**. The ServeRAID Manager program scans both systems to verify the ServeRAID cluster configuration. The progress indicator reports the progress of the validation.
5. If you want to stop the validation before it completes, click **Cancel**. Otherwise, if the ServeRAID Manager program finds problems with the ServeRAID cluster configuration, it reports the problems in the **Problems found** message field. If there are no problems with the ServeRAID cluster configuration, the message field reports only the cluster pairings found.
6. To view the problem report events in an easy-to-read window, double-click an event in the **Problems found** message field.
7. To save the contents of the **Problems found** message field, click **Save**. A Save as window opens. Specify the file name for the message field log and click **OK**.
8. Click **Done** to exit the Cluster validation window. If the ServeRAID cluster configuration has problems, fix the problems; then, use **Validate cluster** again.

Viewing merge group numbers and other cluster information

To view the merge group numbers and other cluster information, do the following:

1. In the Main Tree, click the controller you are using for the cluster.

2. Click  on the tool bar. A window similar to the following opens.

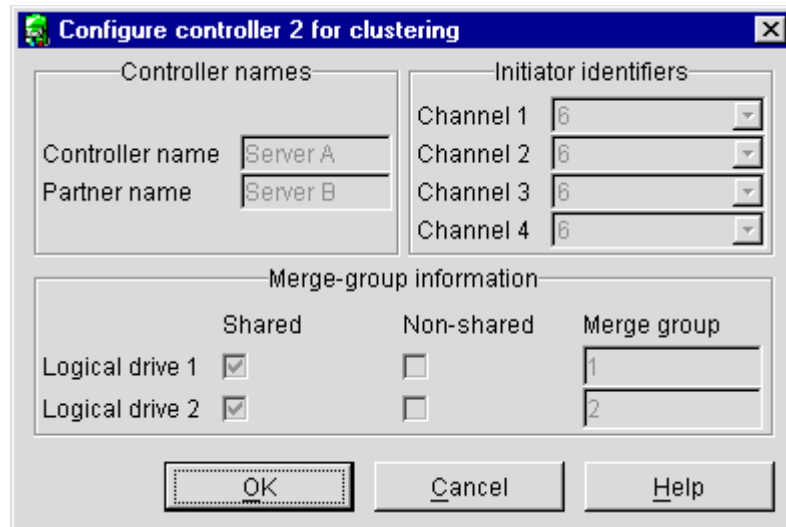


Figure 10. Configure controller for clustering window

Note: In this case, you can view this information only. You cannot change the information.

Chapter 6. Updating the ServeRAID BIOS, firmware, and software code for clustering

Use the following CD and diskette (version 4.0 or later), together as a set, with the IBM ServeRAID Cluster Solution:

- *IBM ServeRAID Support CD*
- *IBM ServeRAID Device Drivers* diskette

Note: If you are using the latest version of the *IBM ServeRAID Support CD*, the BIOS/firmware code is automatically verified when you start the server with the CD in the CD-ROM drive. If the BIOS/firmware code needs updating, you will be prompted to do so.

When you update any RAID software, you must update *all* RAID software at the same time to ensure compatibility with all of the levels of the software. This includes:

- BIOS/firmware
- Device drivers
- ServeRAID Manager
- IBM ServeRAID Cluster Solution software

Ensure that the ServeRAID controllers have the latest IBM ServeRAID software installed; you must have version 4.0, or later, as a minimum. You can obtain this software from the IBM Support Web site. See “Downloadable files from the World Wide Web” on page 40 for additional information. If you are using Microsoft Cluster Service, be sure that you use a version of the ServeRAID software that has been certified by Microsoft for Windows NT and Windows 2000 compatibility, as noted on the IBM Support Web site.

For instructions on how to determine the versions of the software on your ServeRAID controller, refer to the documentation that was included with your server or your controller.

Note: If you are using Windows 2000, check the Windows 2000 Hardware Compatibility List (HCL) included on the *Windows 2000 Advanced Server* CD or the IBM Server Proven List on the IBM Support Web site to determine if your server devices are supported by the Windows 2000 operating system. If not, contact the device manufacturer to receive a device driver for Windows 2000.

Chapter 7. Obtaining ServeRAID updates

IBM periodically makes updated versions of the ServeRAID firmware, device drivers, and utility programs available from the IBM Support page on the World Wide Web. In addition to the updated device drivers and utility programs, there is a command-line program available from the World Wide Web.

Note: If you need to update any ServeRAID software, you must update *all* ServeRAID software at the same time to ensure compatibility with all of the levels of the software. This includes:

- BIOS/firmware
- Device drivers
- ServeRAID Manager program
- Command-line programs

You can download the most current versions of the ServeRAID device drivers and utility programs or a copy of the *IBM ServeRAID Command-Line Programs* diskette from the World Wide Web (see “Downloadable files from the World Wide Web” on page 40).

Refer to the README file on your diskette or CD for installation instructions.

If you do not have access to the World Wide Web, contact your place of purchase, your IBM reseller, or your IBM marketing representative for replacement CDs or diskettes.

The *IBM ServeRAID Support* CD contains the following:

- ServeRAID Manager program. Use this program to configure arrays and logical drives using ServeRAID controllers. This program is available in the following directory:

`d:/PROGRAMS/operatingsystem/MANAGER`

where *d* is the CD-ROM drive and *operatingsystem* is the specific operating system you are using in the ServeRAID installation.

- Device drivers. These files are available in the following directory:

`d:/PROGRAMS/operatingsystem/DRIVER`

where *d* is the CD-ROM drive and *operatingsystem* is the specific operating system you are using in the ServeRAID installation.

- Command-line programs. These command-line programs are available in the following directory:

`d:/PROGRAMS/operatingsystem/CMDLINE`

where *d* is the CD-ROM drive and *operatingsystem* is the specific operating system you are using in the ServeRAID installation.

- ServeRAID files needed to install Microsoft Cluster Service (MSCS) when using Microsoft® Windows NT® or Windows® 2000 and an IBM ServeRAID-4x Ultra160 SCSI controller. These files are available in the following directory:

`d:/PROGRAMS/WINNT/CLUSTER`

where *d* is the CD-ROM drive.

The *IBM ServeRAID Device Drivers* diskette contains the following:

- Device drivers for Microsoft Windows NT, Windows 2000, IBM OS/2®, Novell NetWare, and Red Hat Linux 6.1
- DOS CD-ROM drive device drivers

The *IBM ServeRAID Device Drivers for SCO Operating Systems - UnixWare HBA and OpenServer BTL* diskette contains the following:

- Device drivers for SCO OpenServer 5.0.5
- Device drivers for SCO UnixWare 7.1

Downloadable files from the World Wide Web

You can download new and updated files for the IBM ServeRAID and Clustering products from the IBM Support Web site:

<http://www.ibm.com/pc/support/>

To access ServeRAID and Clustering support, follow these steps:

1. Click **Servers** on the IBM Support page. The IBM Netfinity and PC Server Support page opens.
2. Select **ServeRAID** or **Clustering** from the **Family** list box. The ServeRAID or Clustering Support page opens.
3. On this page, you can select from the following categories in the left margin:

Downloadable files Download the latest versions of Clustering software, the ServeRAID Manager program, BIOS/firmware, device driver updates, and other important information.

Hints and tips Obtain useful information for the IBM Clustering and for ServeRAID products and troubleshooting potential problems.

Online publications Download the installation and user's guides, references, redbooks, and other IBM publications.

You can also review information on parts, products, and software.

Chapter 8. Troubleshooting an IBM ServeRAID Cluster Solution

There is the possibility that one of the servers, one of the ServeRAID controllers, or one of the hard disk drives in a cluster might fail. This section presents procedures you can take to recover from these potential problems. If you still have problems after following these procedures, contact your IBM service representative for further information. Also, you can find hints and tips to help you solve problems on the IBM Support Web site. See “Downloadable files from the World Wide Web” on page 40 for additional information.

Recovering from a failed ServeRAID server

When replacing a failed server in a high-availability configuration, you might not need to replace the ServeRAID controller. However, if you replace your ServeRAID controller, you must reconfigure the controller after you have installed your new server.

Important

The following procedure requires specific configuration settings for the ServeRAID controller. If the server and controller that are being replaced are functional, you can obtain these settings from the controller. However, if the controller or the server is not functional, you will need a record of these settings, such as one that was created when the controller was previously configured. If you are replacing your ServeRAID controller with your server, you *must* have correct configuration information to complete this procedure.

Use the following procedure to recover from a failed server:

1. Shut down the failed server.
2. Remove all hard disk drives from the array in the failed server.

As you remove your hard disk drives, be sure to note the bay in which each drive was installed. If you are replacing your failed server with an identical server, you can reinstall the drives in an identical configuration and get your server up and running quickly.
3. If the ServeRAID controller is functional, remove it from the failed server. As you remove the controller from the failed server, be sure to note the following:
 - Which SCSI cables are connected to the SCSI channel connectors on the controller.
 - Which PCI slot has the controller installed.
4. Install the hard disk drives in the new server.

For information on how to install a hard disk drive, see the documentation that comes with your server. If you are replacing the failed server with an identical server, install each hard disk drive in the same bay as the one it was removed from in the failed server.
5. If the ServeRAID controller in the failed server is *not functional*, or you are not sure, continue with “Recovering from a failed ServeRAID controller” on page 42. Otherwise, perform the following steps with the functional ServeRAID controller:
 - a. Install the controller in the new server.
 - b. Reinstall the cables as they were in the failed server.
 - c. Start the server.

Recovering from a failed ServeRAID controller

You must have the following configuration information to configure your new ServeRAID controller:

- SCSI initiator IDs

The SCSI initiator IDs for a non-shared SCSI channel are set to 7 by default. For shared SCSI channels, the initiator IDs must be 7 or 6 and must be different from the SCSI initiator IDs for the corresponding SCSI channels of the cluster partner controller. Therefore, if you can obtain the SCSI initiator IDs for the corresponding cluster partner system, you know the correct SCSI initiator IDs for this system.

For example, if the cluster partner system-shared SCSI initiator IDs were set to 7, the controller you are replacing would need to have its SCSI initiator IDs set to 6.

- Controller and partner name

You can also determine the proper settings for the controller name and partner name of the controller being replaced by reading the settings from the cluster partner system.

- Stripe-unit size: The stripe-unit size *must* be the same on both controllers in a high-availability cluster.
- Unattended mode setting: This setting is configured automatically as Enabled when a controller is set up for clustering.

If the ServeRAID controller you are replacing *is* functional, you can obtain the above configuration information by starting your server with the *IBM ServeRAID Support* CD in the CD-ROM drive and using the ServeRAID Manager program. For additional information, refer to the ServeRAID Manager online help.


Note: The configuration information is also available in the printable configuration file you might have saved when configuring the ServeRAID controllers for clustering.

If the ServeRAID controller is *not* functional, you will need to refer to a record of the settings that was made when the controller was previously configured. If you do not have a record of the configuration information, refer to “Obtaining the current configuration information” for information that might help you to assign the proper values.

Obtaining the current configuration information

To obtain the SCSI initiator IDs, controller name, partner name, and stripe-unit size from the corresponding cluster partner controller, do the following:

1. Start the cluster partner system with the *IBM ServeRAID Support* CD in the CD-ROM drive. The ServeRAID Manager program starts.
2. In the Main Tree, click the cluster partner controller.

3. Click  on the tool bar. The Configure controller for clustering window opens.

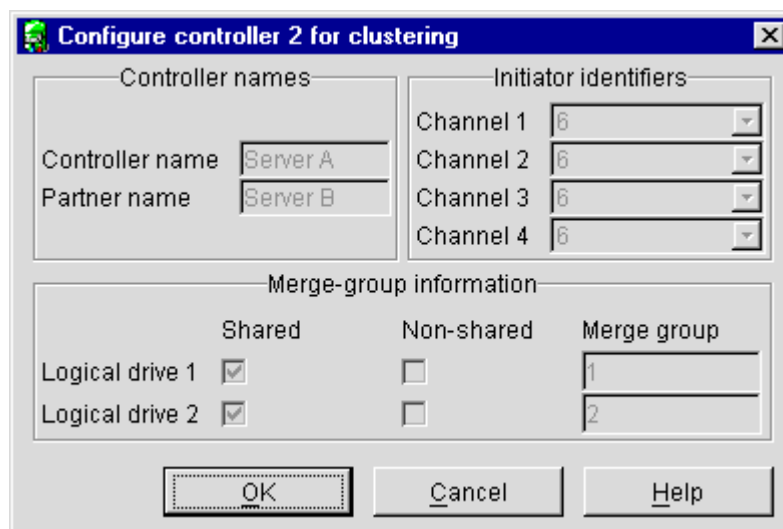


Figure 11. Configure controller for clustering window

4. Note the settings in the fields.
5. To determine stripe-unit size, refer to the Information panel in Figure 7 on page 31.

Replacing a ServeRAID controller

Use the following procedure to replace a failed ServeRAID controller in your high-availability configuration:

1. Shut down the server.
2. Note which SCSI cables are connected to the SCSI channel connectors on the controller.
3. Note which PCI slot has the controller installed.
4. Remove the failing ServeRAID controller.

For instructions on how to remove and install the ServeRAID controller, see "Installing and cabling a ServeRAID controller" in the *IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference*.

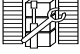
5. Install the new ServeRAID controller. Be sure to install the controller in the same PCI slot from which you removed the failed ServeRAID controller.

Important

- Do *not* reconnect the SCSI channel cables to the controller at this time.
- You must ensure that you have the same level of ServeRAID BIOS and firmware on both controllers in the cluster.

The software is available on the IBM Support Web site. See "Downloadable files from the World Wide Web" on page 40 for additional information.

6. Start the system with the *IBM ServeRAID Support CD* (version 4.0, or later) in the CD-ROM drive. The ServeRAID Manager program starts.
7. In the Main Tree, click the new controller.
8. Click **Actions** → **Restore to factory default settings**. The ServeRAID Manager program initializes the controller configuration.

9. With the new controller still selected in the Main Tree, click  on the tool bar. The Configure controller for clustering window opens.

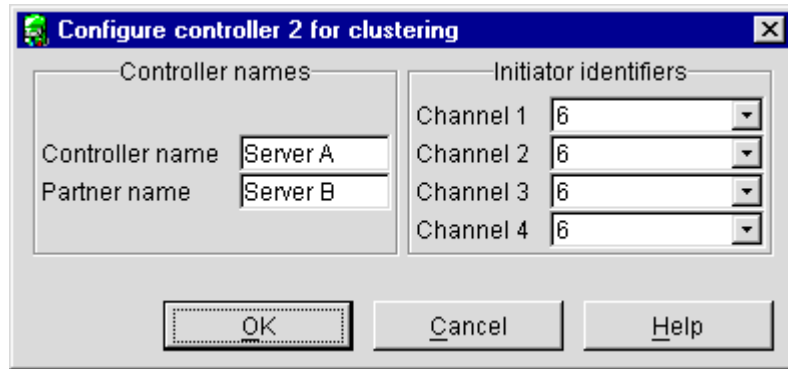


Figure 12. Configure controller for clustering window

10. Using the settings that were assigned to the previous ServeRAID controller (that is, the controller you are replacing), fill in each of the following fields:
- SCSI initiator IDs
 - Controller name
 - Partner name
11. Click **OK**. The new controller stripe-unit size setting defaults to 8 KB. If you need to change this setting to match the partner controller, go to step 12. Otherwise, go to step 15.
12. In the Main Tree, click the new controller.
13. Click **Actions** → **Change stripe-unit size**.
14. Click the new stripe-unit size for your installation.
15. Shut down and turn off the server.
16. Reconnect the SCSI channel cables to the controller.
- Note:** Be sure to connect the cables to the correct SCSI channels as noted in step 2 on page 43.
17. If the ServeRAID controller being replaced attaches to any of the following:
- The server startup disk array
 - Other non-shared disk arrays,
- start the system with the *IBM ServeRAID Support* CD (version 4.0, or later) in the CD-ROM drive. The ServeRAID Manager program starts.
- If the ServeRAID controller being replaced does not attach to any of these items, go to step 21.
18. In the Main Tree, click the new controller.
19. Click **Actions** → **Copy configuration from drives**.
20. Click the merge group number for the non-shared logical drives (that is, 206 or 207).
- Note:** The last number in the merge group number is the shared SCSI initiator ID. That is, if the merge group number is 206, the SCSI initiator ID is 6.
21. Remove the *IBM ServeRAID Support* CD from the CD-ROM drive.
22. Restart your server.
- Note:** If the controller you replaced is the startup controller, the system will now start the operating system properly.

Recovering from a failover in a Microsoft Cluster Service server

After a failover occurs in a Microsoft Cluster Services server, the red drive lights in the IBM Netfinity EXP15 enclosure might turn on for physical drives that are online. There is no hardware problem; the lights turn off within an hour.

Recovering from a failed server in a ServeRAID cluster

Contact IBM technical support for assistance.

Troubleshooting ServeRAID subsystems in a cluster environment

Following is a series of common problems and solutions that can help you troubleshoot your IBM ServeRAID Cluster Solution.

Problem	Action
<p>The ServeRAID Manager program shows physical drives in defunct state.</p>	<ol style="list-style-type: none"> 1. Enable the View shared drives feature in the ServeRAID Manager program. Doing so displays a drive as reserved, instead of defunct, in the following situations: <ul style="list-style-type: none"> • A drive has been moved. • A drive has failed-over to the other server in the cluster. <p>Note: It is normal for drives in the above situations to display as defunct if you have not enabled the View shared drives feature. In this case, the drives shown in the defunct state are not really defective.</p> 2. Check RAID level-1 and RAID level-5 arrays to make sure they are not in critical state. If they are in critical state, replace the failed drive and perform a rebuild operation.
<p>ServeRAID shared logical drives do not failover properly.</p>	<ol style="list-style-type: none"> 1. Ensure that the resource type of each ServeRAID controller shared disk resource is <i>IBM ServeRAID logical disk</i>. If the resource type is shown as <i>physical drive</i>, the localquorum option was not specified properly when MSCS was installed. To correct this problem, you must reinstall the high-availability cluster solution using Microsoft Windows NT. Refer to "Configuring ServeRAID controllers" in the <i>IBM ServeRAID-4x Ultra160 SCSI Controller User's Reference</i> for instructions. 2. Ensure that shared SCSI buses on the ServeRAID controller pair are connected in a way that corresponding SCSI channels are connected. For example, SCSI channel 1 on the controller in the first cluster server is connected to SCSI channel 1 on the controller in the second cluster server, channel 2 is connected to channel 2, and so forth. 3. Ensure that physical SCSI drives that contain logical drives are all connected to shared SCSI channels. 4. Ensure that there are no more than eight shared logical drives defined per pair of ServeRAID controllers. 5. Make sure each shared logical drive has a unique merge group number assigned. Shared merge group numbers <i>must</i> be in the range 1–8. 6. Make sure each ServeRAID controller has been assigned a unique controller name and that each ServeRAID controller has its partner name assigned properly to correspond to the ServeRAID controller in the other cluster server that is attached to the shared SCSI buses. 7. Check for loose shared SCSI bus cables. 8. Ensure that physical drives that are expected to be moved or to failover show up as ready or reserved state on the server that is attempting to take over control of these drives.

Problem	Action
<p>RAID level-5 logical drives cannot be accessed by the operating system after a failover.</p>	<p>Use the ServeRAID Manager program to check the state of the logical drive to ensure that it is not blocked. Using this program, select the logical drive and look for Blocked state Yes. If the logical drive is blocked, make sure all physical drives that are part of the logical drive are in the online state. If all physical drives are <i>not</i> in the online state, a drive might have failed during one of the following situations:</p> <ul style="list-style-type: none"> • A failover • A resynchronization process after a failover <p>Data integrity cannot be guaranteed in this case and the logical drive has been blocked to prevent the possibility of incorrect data being read from the logical drive.</p> <p>Unblock, reinitialize, and synchronize the logical drive and restore the data from a backup source. Depending on the type of data contained on the logical drive and the availability of a recent backup copy, you can do one of the following:</p> <ul style="list-style-type: none"> • Unblock the drive and continue normal operation. • Replace and rebuild one or more defunct drives. <p>However, if you do not reinitialize, synchronize, and restore the drive, be aware that some data could be lost or damaged.</p>
<p>If one of the cluster servers fails and the surviving server takes over the cluster resources, occasionally one or more of the IP address resources will stay in the online pending state for several minutes after moving over to the surviving server. After this, the resource will go to the Failed state and the following error message will be displayed in the surviving server system log (as viewed with the Event Viewer).</p> <p>For example, the Windows NT Event Log Message:</p> <pre> Date: ??? Event ID: 1069 Time: ??? Source: ClusSvc User: N/A Type: Error Computer: ??? Category: (4) </pre> <p>Description:</p> <p>Cluster resource 'ip address resource name' failed</p>	<p>Do the following:</p> <ol style="list-style-type: none"> 1. Using the right mouse button, click the IP resource in the Cluster Administrator. 2. From the General tab, select Properties. 3. Check the box labeled Run this resource in a separate Resource Monitor. A message appears stating that the resource must be restarted for the change to take effect.

Problem	Action
<p>After one of the cluster servers has been shut down normally and the surviving server takes over the cluster resources, occasionally one or more of the IBM ServeRAID logical drive resources will stay in the online-pending state for several minutes, after moving over to the surviving server (when viewed with the Cluster Administrator). After this, the resource will go to the Failed state and the following error message will be displayed in the surviving server system log (as viewed with the Event Viewer).</p> <p>For example, the Windows NT Event Log Message:</p> <pre>Date: ??? Event ID: 1069 Time: ??? Source: ClusSvc User: N/A Type: Error Computer: ??? Category: (4)</pre> <p>Description:</p> <p>Cluster resource 'IBM ServeRAID Logical Disk name' failed.</p>	<p>No action is necessary to bring the resource online after the failover. MSCS will bring this resource online on the surviving server within about four minutes.</p>
<p>You cannot reinstall the ServeRAID Windows NT Cluster Solution. If a previous version of IBM ServeRAID Cluster Solution has been uninstalled, when attempting to reinstall the IBM ServeRAID Windows NT Cluster solution, a message incorrectly appears prompting you to perform an upgrade.</p>	<p>You must delete the <i>C3E76E53-F841-11D0-BFA1-08005AB8ED05</i> registry key. To delete the registry key, do the following:</p> <ol style="list-style-type: none"> 1. From the Start menu, click Run. 2. Type: REGEDIT and click OK. The Registry Editor window opens. 3. Select HKEY_CLASSES_ROOT\CLSID and delete <i>C3E76E53-F841-11D0-BFA1-08005AB8ED05</i>. 4. Reinstall the ServeRAID Windows NT Cluster Solution. Refer to Chapter 3, "Configuring ServeRAID controllers for clustering," on page 11 for instructions.
<p>When running the IPSHAHTO program on a server, the following error message appears:</p> <pre>Warning: CONFIG_SYNC with 0xA0 command FAILED on COntröller #</pre> <p>and either of the following occurs:</p> <ul style="list-style-type: none"> • One or more hot-spare or standby hot-spare drives are defined on controller pairs. • Ready drives are not part of any logical drive configuration on a controller pair in a cluster. 	<p>If all shared drive resources moved successfully when you ran the IPSHAHTO program, it is safe to ignore the error message and no further action is required.</p> <p>However, if shared drive resources <i>failed</i> to move when you ran the IPSHAHTO program, remove all ready drives that are not part of any logical drive configuration on a controller pair in a cluster.</p>
<p>Array identifiers and logical drive numbers might change during a failover condition.</p>	<p>By design, the array identifiers and logical drive numbers might change during a failover condition. Consistency between the merge group numbers and Windows NT or Windows 2000 permanent (sticky) drive letters is maintained, while the ordering process during a failover condition is controlled by the Microsoft Cluster Service software and the available array identifiers on the surviving server.</p>