

IBM Netfinity
High Availability Cluster Solutions
using the IBM ServeRAID-3H Ultra2 SCSI Adapter

Installation and User's Guide

IBM Netfinity
High Availability Cluster Solutions
using the IBM ServeRAID-3H Ultra2 SCSI Adapter

Installation and User's Guide



Note

Before using this information and the product it supports, be sure to read the general information under Appendix A, "Notices" on page 70.

Fourth Edition (December 1998)

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

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

This publication was developed for products and services offered in the United States of America. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local IBM representative for information on the products, services, and features available in your area.

Requests for technical information about IBM products should be made to your IBM reseller or IBM marketing representative.

© Copyright International Business Machines Corporation 1997, 1998. All rights reserved.

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

About This Book	v
How This Book Is Organized	v
Related Publications	vii
Chapter 1. Introduction to IBM High-Availability Cluster Solutions	1
Chapter 2. Preparing to Install a High-Availability Cluster Solution	4
Terms and Definitions	4
ServeRAID Considerations	6
Chapter 3. Configuring the ServeRAID Adapter for the High-Availability Cluster Solution	10
Preparing the Servers	11
Configuring Server A	12
Configuring Server B	20
Chapter 4. Installing a High-Availability Cluster Solution Using Windows NT	28
Special Installation Instructions	34
Creating a ServeRAID Disk Resource on a ServeRAID Server with MSCS	34
Deleting and Later Recreating a ServeRAID Logical Disk Resource	36
Uninstalling MSCS on ServeRAID Servers	38
Upgrading MSCS on ServeRAID Servers	39
Upgrading the IBM ServeRAID Administration and Monitoring Program	39
Upgrading the IBM ServeRAID Microsoft Windows NT Cluster Solution	39
Upgrading the IBM ServeRAID Windows NT Miniport Driver	40
Upgrading the IBM ServeRAID BIOS/Firmware Diskette	40
Upgrade the Software for the Other Node in the Cluster	41
Chapter 5. Installing a High-Availability Cluster Solution Using Vinca and Novell IntranetWare	42

Chapter 6. Monitoring and Updating a High-Availability Cluster Solution	43
Checking the Availability of a Drive	43
Adding Hard Disk Drives to a Cluster	44
Viewing Merge Group Numbers and Other Cluster Information	46
Chapter 7. Troubleshooting a High Availability Cluster	47
Recovering From a Hard Disk Drive Failure	47
Recovering From a Failed ServeRAID Adapter	48
Recovering From a Failed Server	57
Troubleshooting ServeRAID Subsystems in a Cluster Environment	65
Appendix A. Notices	70
Trademarks	71

About This Book

This book provides information and instructions for setting up a high-availability clustering solution using the IBM ServeRAID-3H Ultra2 SCSI Adapter, herein called ServeRAID adapter.

Note: The ServeRAID-3L Ultra2 SCSI adapters do not support clustering.

This book is intended for experienced users who will be involved with setting up clustering and high-availability solutions for their computer installations.

How This Book Is Organized

Chapter 1, "Introduction to IBM High-Availability Cluster Solutions," provides an overview of high-availability cluster solutions.

Chapter 2, "Preparing to Install a High-Availability Cluster Solution," provides the information you need to know before you start the installation and setup of a high-availability cluster solution. It provides definitions of important terms used in this manual and some considerations you need to be familiar with for the ServeRAID adapter.

Chapter 3, "Configuring the ServeRAID Adapter for the High-Availability Cluster Solution," provides the information necessary to configure the ServeRAID adapters in the two servers used in the high-availability cluster solution. Information also is provided for defining all shared and non-shared logical drives and setting Merge IDs.

Chapter 4, "Installing a High-Availability Cluster Solution Using Windows NT," provides instructions for setting up a high-availability cluster solution with Windows NT.

Chapter 5, "Installing a High-Availability Cluster Solution Using Vinca and Novell IntranetWare," directs you to the information needed to set up a high-availability cluster solution with Novell IntranetWare.

Chapter 6, “Monitoring and Updating a High-Availability Cluster Solution,” provides information on how to use the ServeRAID Administration and Monitoring Utility to determine the availability of logical drives in a cluster. It also includes the procedures to use if you add or remove a hard disk drive to your high-availability cluster solution, or to uninstall your high-availability cluster solution.

Chapter 7, “Troubleshooting a High Availability Cluster,” provides information on how to recover from some potential problems, such as a defective ServeRAID adapter or a defective hard disk drive in a logical drive.

Appendix A, “Notices,” includes legal notices and trademarks.

Related Publications

During the installation of the High-Availability Cluster Solution, you might need to refer to one or more of the following manuals:

- The *IBM Shared Disk Clustering Hardware Reference*, Part Number 33L3945, provides general information about planning and configuring a shared-disk cluster using IBM Netfinity and PC Server products. It contains illustrations, descriptions, and parts listings for various high-availability, shared-disk cluster examples. This publication is currently available on the World-Wide Web at

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

Once you are on the web page, perform a search for the clustering information.

- The *Microsoft Cluster Server User's Guide*, provides instructions for setting up and using the Microsoft Cluster Server Software.
- The *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*, Part Number 05L2054, provides instructions for the installation of the ServeRAID adapter, as well as instructions on how to use and maintain RAID configurations.

For more information about these publications, call 1-800-879-2755 in the United States and Puerto Rico, or call 1-800-426-4968 in Canada. In all other countries, contact your IBM reseller or marketing representative.

Chapter 1. Introduction to IBM High-Availability Cluster Solutions

A *cluster* is a type of parallel or distributed system that consists of interconnected computers used as a single, unified computing resource. In other words, 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-node cluster, where both nodes, or servers, can access the same storage devices, but only one server at a time controls the storage devices shared by both servers. If one server fails, the remaining server automatically assumes control of the resources that the failed server was using, 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 access to that server's data and applications is still available.

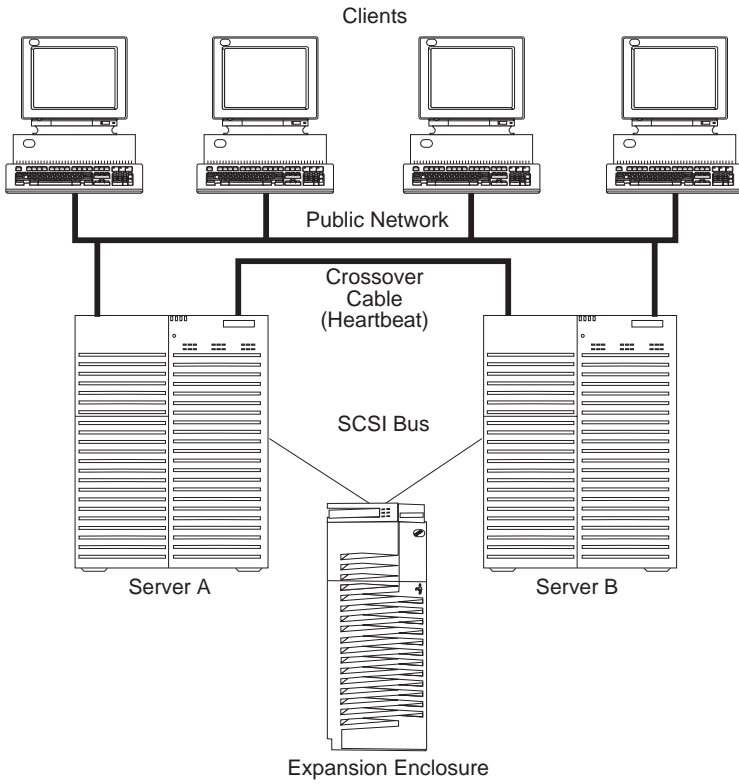
Once the failed server is operational again, it can be placed back into the cluster, the resources are reallocated between the two servers and the cluster can then resume normal operations.

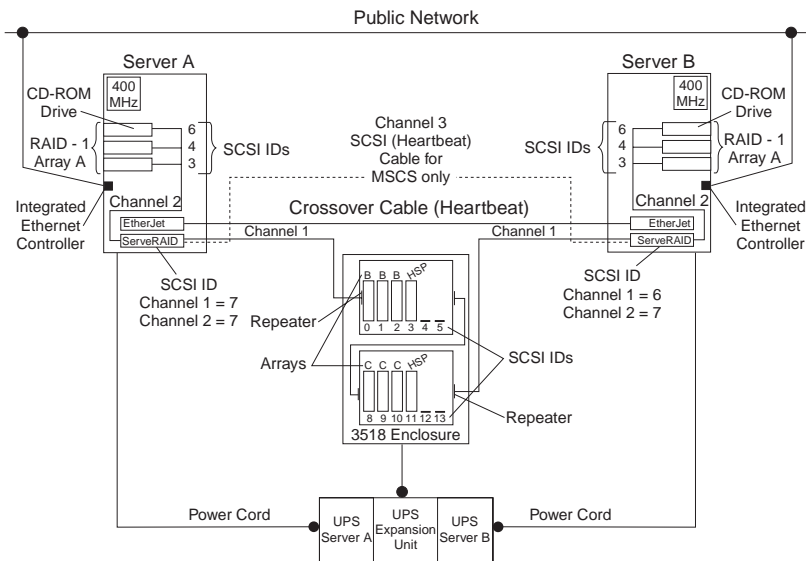
IBM high-availability cluster solutions, as discussed in this manual, use IBM ServeRAID-3H Ultra2 SCSI Adapters (ServeRAID-3L Ultra2 SCSI adapters do not support clustering) and IBM expansion enclosures, such as the IBM Enterprise Expansion Enclosure (Model 3518), the IBM PC Server 3519 Rack Expansion Enclosure, the IBM Netfinity EXP10 Rack Expansion Enclosure or the IBM Netfinity EXP15 Rack Expansion Enclosure.

IBM high-availability cluster solutions can be installed using one of the following:

- Microsoft Cluster Server, which is part of Windows NT Server Enterprise Edition
- Vinca High Availability for Novell IntranetWare, which is part of the IBM Netfinity Cluster Pack by Vinca

The following figures show an example of a simple high-availability clustering solution. Figure 1 shows a high-level view of a cluster. Figure 2 shows a more detailed view of a cluster.





The solution you use for your installation depends on your own specific requirements. For more information about requirements and types of installations, refer to the *IBM Shared Disk Clustering Hardware Reference*, Part Number 33L3945. This publication can be found on the IBM World Wide Web page at

<http://www.pc.ibm.com/us/netfinity>

Once you are on the web page, perform a search for the clustering information.

Chapter 2. Preparing to Install a High-Availability Cluster Solution

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

Note: The following terms and definitions are for the Microsoft Cluster Solution using Windows NT. For similar terms and definitions for the Novell IntranetWare Cluster Solution using Novell NetWare, refer to *Vinca High Availability for NetWare and IBM ServeRAID-3H and ServeRAID-3R Ultra2 SCSI Adapters Installation and User's Guide*.

Terms and Definitions

The following terms will be used in this manual.

Quorum Resource: A dedicated drive that both cluster servers share to store and retrieve quorum information. This resource maintains logged data that is critical to recovery when servers fail to communicate with each other. This data can be owned by one cluster server at a time only and is used to determine which server will take ownership of the resources if the two servers cannot communicate with each other. When the quorum owner fails, the surviving server takes over ownership of that quorum drive and all the other resources in the cluster. This drive is a critical component of the cluster. If the quorum drive becomes defective, recovery is impossible. A quorum drive must be configured as RAID level 1.

Failback: This is a special case of failover and 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.

Failover: This is the process of relocating the cluster resources from a failed server to the surviving server. The detection of a failure is made by the resource-monitor program responsible for the resource. If a resource failure occurs, the resource-monitor program notifies the Cluster Service, which triggers the actions 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 brings the other server down for scheduled maintenance.

Failover generally consists of three components:

- Failure detection
- Resource relocation
- Application restart

For more information on failover, refer to the *Clustering and High Availability Guide for IBM Netfinity and IBM PC Servers*.

Resource Group: Dependent resources that are 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. It is possible that all resources in a cluster will need to be in one group.

Also, any cluster operation performed on a group is performed on all resources that are present within that group. For example, if a resource needs to be moved from server A to server B, all the resources defined in the group that the resource is in will be moved.

SCSI Heartbeat Connection: This is a SCSI cable connection that joins two IBM ServeRAID adapters, one in each server, together.

ServeRAID Considerations

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

Note: The following ServeRAID considerations are for the Microsoft Cluster Solution using Windows NT. For similar considerations for the Novell IntranetWare Cluster Solution, refer to *Vinca High Availability for NetWare and IBM ServeRAID II SCSI Adapter User's Guide*.

- Microsoft Cluster Server (MSCS) Configurations using the ServeRAID II adapters in each of the servers require a SCSI connection joining the two ServeRAID II adapters together. This connection is known as a *SCSI heartbeat*. The SCSI heartbeat connection is required to manage the Microsoft Cluster Server quorum disk and, in the event of a server failure, to determine which is the surviving server. The connection is a combination of any length 0.8-mm to 68-pin cable and a pigtail 68-pin to 0.8 mm cable.
- The SCSI heartbeat connection *must* be connected to the third channel of the ServeRAID adapter pair that has the quorum drive connected to it. No disks can be installed on this heartbeat channel. If you choose to move the quorum drive to another ServeRAID II adapter, you must also move the SCSI heartbeat cable on both servers to the new quorum ServeRAID adapter pair.
- The quorum disk can be placed on any ServeRAID channel shared by the two servers, except for the channel used by the SCSI heartbeat cable. The quorum logical drive can be used to store customer programs and data.
- The ServeRAID adapter does not allow the failover of RAID level-5 logical drives that are in a critical state; that is, a drive in a RAID level-5 array has failed. For this reason, do not use RAID level-5 for the quorum disk. The quorum drive must be configured as a RAID 1 logical drive.

Note: You should define hot-spare drives (HSP) 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 data integrity is maintained. Logical drives that are not shared between the two servers can be configured for write-back mode for improved performance. For more information about write policy, refer to the *IBM Netfinity ServeRAID-3H and ServeRAID-3R Ultra2 SCSI Adapters Installation and User's Guide*.
- The SCSI-2 Fast/Wide Enhanced Repeater Card is required for configurations using the IBM Enterprise Expansion Enclosure (Model 3518) or the IBM Model 3519 Rack Expansion Enclosure. The previous version of the SCSI-2 Fast/Wide Repeater is not supported in these configurations. For information on the latest version of the SCSI Fast/Wide Enhanced Repeater Card, contact your IBM remarketer or dealer.
- Each ServeRAID adapter must have a unique host ID assigned.
- Do not attempt to change the host ID on both ServeRAID II adapters in a pair at the exact same time. Doing so can cause a problem, such as a server “lock-up.”
- Each ServeRAID adapter, and adapter pair, must have a different SCSI Initiator ID assigned to the shared channels. For example, one server, Node A, could be set to ID 6, and the other server, Node B, could be set to ID 7. This will be described during the setup instructions.
- The stripe unit size of each ServeRAID adapter in a pair must be set to the same value (8 KB, 16 KB, 32 KB, or 64 KB).
- The SCSI channels in an adapter pair must be cabled so that the same channels are shared between adapters.
- All disks that will be shared must be connected to the shared channels. This includes hot-spare (HSP) drives, as well as data drives.
- Only one logical drive must be created for each physical array because only one server at a time can own a physical drive. For

example, only logical drives A0, B0, C0, and so on are supported, not A1, B1, and so on.

- Merge ID numbers in the range of 1 to 8 must be assigned to each logical drive that will be shared. Merge ID numbers in the range of 200–215, where the last two numbers are the SCSI initiator of a shared channel, must be assigned to the non-shared logical drives.
- If you are starting (booting) the operating system from a shared adapter, define the first logical drive as the startup drive and assign a non-shared Merge ID, for example, 206 for Node A.
- Do not connect nondisk devices, such as tape drives, to shared channels. They will be accessible to both servers and conflicts can arise.
- Hot-spare drives are not shared between the adapters in a pair. Each ServerRAID adapter must have a hot-spare drive defined if hot-spare protection is desired, no matter which node is controlling a logical drive. All hot-spare drives must be connected to shared channels so that if the hot-spare drive is used to replace a failed shared disk, the disk can be accessed from either node in the case a failover occurs later.
- RAID level-0 and RAID level-1 logical drives can failover in normal or critical state. RAID level-5 logical drives cannot failover in critical state. Use hot-spare drives to minimize the time that logical drives are in critical state and exposed to a second disk or system failure.
- The total number of logical drives per adapter is eight before or after a failover. A failover will not complete if this number is exceeded.
- Logical drives that are currently undergoing Logical Drive Migration (LDM) operations, will not be allowed to failover. However, all other drives will be able to failover if necessary.

- Drives that are configured by one adapter in the pair must appear as RDY (ready) drives to the other adapter in the pair. The same must also be true in reverse. Be sure to verify this consideration when drives are added after the nodes are started. Restart the servers, or scan for new devices, or set the drives to RDY state using the RAID Administration and Monitoring Utility after adding new drives to a server.
- If a failover occurs while a critical RAID level-1 logical drive is rebuilding to a spare disk, the rebuild automatically starts a few seconds after the failover is completed.
- The cluster support software will initiate a synchronization of RAID level-1 and RAID level-5 logical drives immediately after a failover. If a drive fails before this synchronization is complete, logical drive access is placed in the blocked state.
- When a logical drive spans across multiple channels and a failure within the drive subsystem occurs that is unique to a channel (for example, a disconnected cable), then the entire physical array will be marked as “bad” even though access from the surviving server can occur. Therefore, make sure that drives do not span across multiple channels.
- A ServeRAID adapter pair can have only eight shared logical drives.

Chapter 3. Configuring the ServeRAID Adapter for the High-Availability Cluster Solution

To begin the installation of a high-availability cluster solution, you must first configure the ServeRAID adapter, define logical drives, and set up Merge IDs. To do this, perform the following procedure.

Important

Before beginning this procedure, read the following information.

- Read "ServeRAID Considerations" on page 6.
- Ensure that the ServeRAID adapters have the latest BIOS/firmware (microcode) installed. You must have Version 3.0 or higher as a minimum. Refer to the following website to download the latest level

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

For instructions on how to determine the versions of the BIOS/firmware (microcode) on your ServeRAID adapter, refer to the documentation that comes with your server or your adapter.

Preparing the Servers

Before you configure Server A and Server B, do the following:

1. If necessary, perform any operating system shutdown procedure on the two servers, herein called Server A and Server B, you are going to use in the cluster, then turn off the servers.

Note: If you are using the Microsoft Cluster Server (MSCS), determine which ServeRAID logical drive will be used as the quorum resource and the ServeRAID adapter pair that it is connected to. Then, connect the channel 3 connectors of that ServeRAID adapter pair together, using the appropriate cable options, to create the SCSI heartbeat connection. See the *IBM Shared Disk Clustering Hardware Reference* for more information and part numbers.

2. If you have not done so, connect any shared external drive enclosures between both servers; then, turn on the enclosures.

Configuring Server A

Note: For complete information on using the ServeRAID Configuration Program, refer to the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*.

To configure Server A, do the following:

Important

If you already have logical drives defined that you wish to keep, go to step 6 on page 14. The steps prior to step 6 will cause loss of data by deleting all logical drives already defined.

1. Start server A using the *IBM ServeRAID Configuration CD* (or the configuration CD that came with your server). At this time, leave server B turned off.

A screen similar to the following appears:

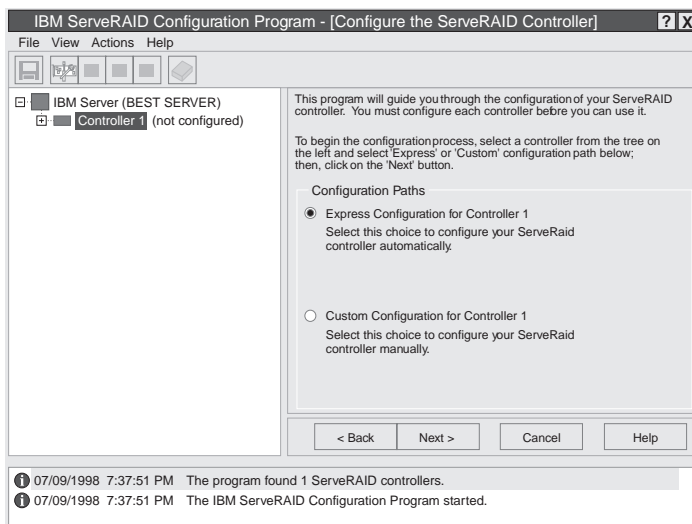


Figure 1. Configuration Mode Screen

2. Click on **Cancel**.

Important

The ServeRAID adapter only should be initialized one time when the adapter is first installed. Initializing the adapter removes all configuration information from NVRAM. This information includes the SCSI IDs, the cluster and the cluster partner name, and the array information. This point is especially important if the Windows NT system boot drive is attached to the IBM ServeRAID adapter. If the initialization occurs after the boot drive array has been created, this array information will be removed and the server will not be able to startup into Windows NT.

3. Using the right mouse button, click on the ServeRAID controller in the main tree that you are configuring and a pop-up list appears.
4. Click on **Initialize the Controller Configuration**
If a confirmation window appears, click on **Yes**.
5. When complete, a message appears at the bottom of the screen saying that the initialization was successful.

Note: Do not attempt to change the host ID on both ServeRAID adapters in a pair at the exact same time. Doing so can cause a problem, such as a server “lock-up.”

6. Using the right mouse button, click on the ServeRAID controller in the main tree that you are configuring and the pop-up list appears.
7. Click on **Configure for Clustering**.
8. A window appears. Fill in the following information:

Important

If your configuration will contain multiple ServeRAID adapter pairs, each adapter Host Configuration ID must be unique.

The name of the Host ID and the Cluster Partner ID are case sensitive.

- Host ID - Enter a name for the Host ID for the first controller, for example, Server_A.
- Cluster Partner ID - Enter a name for the Cluster Partner ID for the first controller, for example, Server_B.

Note: Servers cannot have the same initiator_identifier ids.

- Initiator Identifiers (Channels) - The system presets these to an unused value, for example, 6.

Note: Be sure to record the Host IDs and Cluster Partner IDs in a safe place for future reference. You will need this information to recover the configuration if the ServeRAID adapter should fail.

9. Record the Host ID and Cluster Partner ID.
10. Click **OK**

Notes:

- a. When creating logical drives, be aware of the following way in which the information is displayed on the ServeRAID configuration program screens. **You only should create one logical drive for each array when defining shared logical drives for use in a cluster.**

When logical drives are created, they appear on the screen in the order they are created. For ease of understanding, create an array and then all logical drives for that array *before* you create another array. If they are created in that order, they will appear on the screen as shown in the first figure. (***** is other information that will appear on the screen, such as RAID level, Write Policy, and so on, but is not needed for this explanation).

```
A ***** A0 ***** ***** ***** ***** *****
B ***** B0 ***** ***** ***** ***** *****
C ***** C0 ***** ***** ***** ***** *****
```

If logical drives are created out of order, meaning you do not follow the order mentioned above, then they are listed on the screen out of order, which might cause some confusion. For example, say you created two arrays, A and B, but did not define any logical drives at that time. Then, afterwards, you go back and create a logical drive on array B. That logical drive will then appear in the first line of the logical drive box. Now, if you go back and create another logical drive on array A, that logical drive will show up on the second line of the logical drive box, as shown in the second figure.

```
A ***** B0 ***** ***** ***** ***** *****
B ***** A0 ***** ***** ***** ***** *****
```

- b. Hot-Spare/Standby Hot-Spare (HSP/SHS) devices must be defined only on a shared channel and each adapter must have its own HSP/SHS device defined. This means that the same HSP/SHS device must not be defined to both adapters in the adapter pair. For example: if node A has a particular drive set for hot-spare, this same drive on node B should remain in a Ready state.

11. Create RAID arrays and logical drives on server A as follows:

- a. Click on the ServeRAID controller that you want to configure in the Main Tree.
- b. Click on the **Custom Configuration** radio button.
- c. Click on the **Next** button.

Notes:

- 1) You must have two or more drives for RAID level 1 and three or more drives for RAID level 5 within the same array.
 - 2) In the Windows NT environment, 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 customer programs and data.
- d. Using the right mouse button, click on the drive or SCSI Channel icons in the Main Tree to select the drives that you want to add to your arrays, delete from your arrays, or define as hot-spare drives; then, select a choice from the pop-up list.

or

Drag the drive or SCSI Channel icon from the Main Tree and drop it on the Array or Hot-Spare Drive icon in the Main Panel on the right. If you change your mind, you can drag the icons back to the Main Tree to remove them from the configuration.

- e. After you select the data drives for your arrays and define your hot-spare drives, click on the **Next** button.

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, click on the **Back** button, use the right mouse button to click on the specific drive or Array icon in the Main Panel on the right, and then select **Remove from New Array** or **Delete New Array**.

- Remove a specific hot-spare drive or all newly defined hot-spare drives. To do this, click on the **Back** button, use the right mouse button to click on the Hot-Spare Drive icon in the Main Panel on the right, and then select **Remove New Hot-Spare Drive** or **Remove All New Hot-Spare Drives**.
- f. Select a RAID level for the logical drive from the RAID pull-down menu in the Main Panel. (Refer to the ServeRAID documentation for descriptions of the supported levels.)
 - g. Move the sliding bar in the Main Panel from right to left to allot data and parity space for the logical drive.

Notes:

- 1) Some operating systems have size limitations for logical drives. Before you save the configuration, you might want to verify that the size of the logical drive is appropriate for your operating system. See your operating-system documentation for more detailed information.
 - 2) Typically, the first logical drive defined on the first ServeRAID adapter found by system BIOS during startup will be your startup (boot) drive.
- h. Click on the **Next** button.
 - i. Review the information that appears on the Configuration Summary screen.

Note: To change the configuration, click on **Modify** icon button.
 - j. Click on the **Apply** to accept and save the configuration.

Note: For Microsoft Windows NT clusters, you should define all shared logical drives on server A.

Then, continue with step 12.

12. After creating the arrays and logical drives, you must assign a Merge ID to the shared drives. To set the Merge IDs, do the following: :

Note: Shared and Non-Shared Merge ID considerations:

- Shared logical drives must have unique Merge IDs.
 - Valid Merge IDs for non-shared logical drives are 200–215. An example would be to set all non-shared drives on server A to Merge ID 206 and all non-shared drives on server B to Merge ID 207.
 - The Merge ID must be unique for each shared logical drive that will be shared by the adapter pair. For example, if you have a logical drive with a merge ID of 1 on server A, you cannot assign the same merge ID to a logical drive on server B. Valid IDs for shared logical drives are 1 to 8.
13. Using the right mouse button, click on the ServeRAID controller in the main tree that contains the logical drives and a pop-up list appears.
 14. Click on **Configure for Clustering**.
 15. A window appears.
 16. Click on the **Shared** checkbox for the logical drives that will be shared between the two ServeRAID adapters.
 17. Set the Merge ID. You can change the Merge Group value, but use the default provided is ok.
 18. Click on **Ok**
 19. Exit the Configuration Program by either clicking on the box with the "X" in the upper right corner, or by clicking on **File**, and then **Exit**.
 20. Shutdown Server A.
 21. Restart the server.
 22. When the ServeRAID Mini-Configuration prompt appears, press **Cntl+i** to go to the mini-configuration program.
 23. Select the ServeRAID adapter you are working on by using the Up Arrow (↑) or Down Arrow (↓) and then press **Enter**
 24. Select **Advanced Functions**.
 25. Select **Configure BIOS Settings**.
 26. Select BIOS Unattended Mode and press **Enter** to set it On.

27. Select **Save Configuration to ServeRAID Controller** and press **Enter**.

Configuring Server B

To configure Server B, do the following:

Note: For complete information on using the ServeRAID Configuration Program, refer to the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*.

Important

If you already have logical drives defined that you wish to keep, go to step 6 on page 22. The steps prior to step 6 will cause loss of data by deleting all logical drives already defined.

1. Start server B using the *IBM ServeRAID Configuration CD* (or the configuration CD that came with your server). At this time, leave server A turned off.

A screen similar to the following appears:

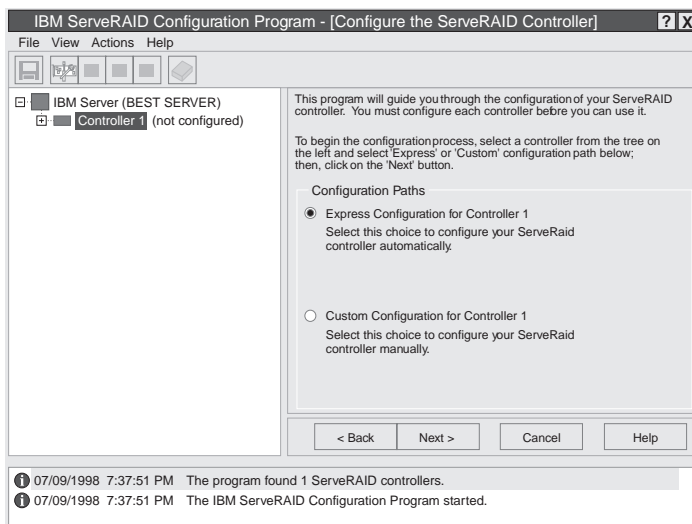


Figure 2. Configuration Mode Screen

2. Click on **Cancel**.

Important

The ServeRAID adapter only should be initialized one time when the adapter is first installed. Initializing the adapter removes all configuration information from NVRAM. This information includes the SCSI IDs, the cluster and the cluster partner name, and the array information. This point is especially important if the Windows NT system boot drive is attached to the IBM ServeRAID adapter. If the initialization occurs after the boot drive array has been created, this array information will be removed and the server will not be able to startup into Windows NT.

3. Using the right mouse button, click on the ServeRAID controller in the main tree that you are configuring and a pop-up list appears.
4. Click on **Initialize the Controller Configuration**
If a confirmation window appears, click on **Yes**.
5. When complete, a message appears at the bottom of the screen saying that the initialization was successful.

Note: Do not attempt to change the host ID on both ServeRAID 3 adapters in a pair at the exact same time. Doing so can cause a problem, such as a server “lock-up.”

6. Using the right mouse button, click on the ServeRAID controller in the main tree that you are configuring and the pop-up list appears.
7. Click on **Configure for Clustering**.
8. A window appears. Fill in the following information:

Important

If your configuration will contain multiple ServeRAID adapter pairs, each adapter Host Configuration ID must be unique.

The name of the Host ID and the Cluster Partner ID are case sensitive.

- Host ID - Enter a name for the Host ID for the first controller, for example, Server_B.
- Cluster Partner ID - Enter a name for the Cluster Partner ID for the first controller, for example, Server_A.
- Initiator Identifiers (Channels) - The system presets these to an unused value, for example, 7.

Note: Be sure to record the Host IDs and Cluster Partner IDs in a safe place for future reference. You will need this information to recover the configuration if the ServeRAID adapter should fail.

9. Record the Host ID and Cluster Partner ID.
10. Click **OK**

Important

- If you are using Microsoft Windows NT Server Enterprise Edition, you have now completed the configuration for your ServeRAID adapters and hard disk drives, which are the initial installation steps for a high-availability cluster solution. The rest of the procedure you need to do depends on the operating system you are using. Go to step 20 on page 26.
- If you are using Vinca and Novell IntranetWare, continue with the following steps.

Notes:

- a. When creating logical drives, be aware of the following way in which the information is displayed on the ServeRAID configuration program screens. **You only should create one logical drive for each array when defining shared logical drives for use in a cluster.**

When logical drives are created, they appear on the screen in the order they are created. For ease of understanding, create an array and then all logical drives for that array *before* you create another array. If they are created in that order, they will appear on the screen as shown in the first figure. (***** is other information that will appear on the screen, such as RAID level, Write Policy, and so on, but is not needed for this explanation).

```
A ***** A0 ***** ***** ***** ***** *****
B ***** B0 ***** ***** ***** ***** *****
C ***** C0 ***** ***** ***** ***** *****
```

If logical drives are created out of order, meaning you do not follow the order mentioned above, then they are listed on the screen out of order, which might cause some confusion. For example, say you created two arrays, A and B, but did not define any logical drives at that time. Then, afterwards, you go back and create a logical drive on array B. That logical drive will then appear in the first line of the logical drive box. Now, if you go back and create another logical drive on array

A, that logical drive will show up on the second line of the logical drive box, as shown in the second figure.

```
A ***** B0 ***** ***** ***** ***** *****
B ***** A0 ***** ***** ***** ***** *****
```

- b. Hot-Spare/Standby Hot-Spare (HSP/SHS) devices must be defined only on a shared channel and each adapter must have its own HSP/SHS device defined. This means that the same HSP/SHS device must not be defined to both adapters in the adapter pair. For example: if node A has a particular drive set for hot-spare, this same drive on node B should remain in a Ready state.

11. Create RAID arrays and logical drives on server B as follows: .

- a. Click on the ServeRAID controller that you want to configure in the Main Tree.

Notes:

- 1) You must have two or more drives for RAID level 1 and three or more drives for RAID level 5 within the same array.
 - 2) In the Windows NT environment, 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 customer programs and data.
- b. Using the right mouse button, click on the drive or SCSI Channel icons in the Main Tree to select the drives that you want to add to your arrays, delete from your arrays, or define as hot-spare drives; then, select a choice from the pop-up list.

or

Drag the drive or SCSI Channel icon from the Main Tree and drop it on the Array or Hot-Spare Drive icon in the Main Panel on the right. If you change your mind, you can drag the icons back to the Main Tree to remove them from the configuration.

- c. After you select the data drives for your arrays and define your hot-spare drives, click on the **Next** button.

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, click on the **Back** button, use the right mouse button to click on the specific drive or Array icon in the Main Panel on the right, and then select **Remove from New Array** or **Delete New Array**.
 - Remove a specific hot-spare drive or all newly defined hot-spare drives. To do this, click on the **Back** button, use the right mouse button to click on the Hot-Spare Drive icon in the Main Panel on the right, and then select **Remove New Hot-Spare Drive** or **Remove All New Hot-Spare Drives**.
- d. Select a RAID level for the logical drive from the RAID pull-down menu in the Main Panel. (Refer to the ServeRAID documentation for descriptions of the supported levels.)
- e. Move the sliding bar in the Main Panel from right to left to allot data and parity space for the logical drive.

Notes:

- 1) Some operating systems have size limitations for logical drives. Before you save the configuration, you might want to verify that the size of the logical drive is appropriate for your operating system. See your operating-system documentation for more detailed information.
 - 2) Typically, the first logical drive defined on the first ServeRAID adapter found by system BIOS during startup will be your startup (boot) drive.
- f. Click on the **Next** button.
- g. Review the information that appears on the Configuration Summary screen.
- Note:** To change the configuration, click on **Modify** icon button.
- h. Click on the **Apply** to accept and save the configuration.

Note: For Microsoft Windows NT clusters, you should define all shared logical drives on server A.

Then, continue with step 12.

12. After creating the arrays and logical drives, you must assign a Merge ID to the shared drives. To set the Merge IDs, do the following:

Note: Shared and Non-Shared Merge ID considerations:

- Shared logical drives must have unique Merge IDs; however, non-shared logical drives can share the same Merge ID.
- Valid Merge IDs for non-shared logical drives are 200–215. An example would be to set all non-shared drives on server A to Merge ID 206 and all non-shared drives on server B to Merge ID 207.
- The Merge ID must be unique for each shared logical drive that will be shared by server A and server B. For example, if you have a logical drive with a merge ID of 1 on server A, you cannot assign the same merge ID to a logical drive on server B. Valid IDs for shared logical drives are 1 to 8.

13. Using the right mouse button, click on the ServeRAID controller in the main tree that contains the logical drives and a pop-up list appears.
14. Click on **Configure for Clustering**.
15. A window appears.
16. Click on the **Shared** checkbox for the logical drives that will be shared between the two ServeRAID adapters.
17. Set the Merge ID. You can change the Merge Group value, but use the default provided is ok.
18. Click on **Ok**
19. Exit the Configuration Program by either clicking on the box with the "X" in the upper right corner, or by clicking on **File**, and then **Exit**.
20. Shutdown Server B.

21. Restart the server.
22. When the ServeRAID Mini-Configuration prompt appears, press **Cntl+i** to go to the mini-configuration program.
23. Select the ServeRAID adapter you are working on by using the Up Arrow (↑) or Down Arrow (↓) and then press **Enter**
24. Select **Advanced Functions**.
25. Select **Configure BIOS Settings**.
26. Select BIOS Unattended Mode and press **Enter** to set it On.
27. Select **Save Configuration to ServeRAID Controller** and press **Enter**.

You have now completed the configuration of your ServeRAID 3 adapters and hard disk drives, which are the initial installation steps for a high-availability cluster solution. The rest of the procedure you need to do depends on the operating system you are using.

- If you are using Vinca and Novell IntranetWare, go to Chapter 5, “Installing a High-Availability Cluster Solution Using Vinca and Novell IntranetWare” on page 42.

Chapter 4. Installing a High-Availability Cluster Solution Using Windows NT

To install a high-availability cluster solution using Windows NT, complete the following procedure.

Important

Before beginning this procedure, make sure you have:

- Read Chapter 2, "Preparing to Install a High-Availability Cluster Solution" on page 4, which contains important information on ServeRAID considerations.
- Read Chapter 3, "Configuring the ServeRAID Adapter for the High-Availability Cluster Solution" on page 10, which steps you through the installation and setup procedure for configuring the ServeRAID adapters in the two servers you will be using for your high-availability cluster solution.
- Ensure that you are using the *ServeRAID Device Driver Diskette*, Version 2.80 or higher during the installation of the Microsoft Windows NT Server Enterprise Edition. Refer to the README file on the ServeRAID device drivers diskette for more information.

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

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

2. Format all shared logical drives and assign drive letters to them by doing the following:

Important

Before using the IPSHAHTO.EXE program:

- Use the Administration and Monitoring Utility to scan for new devices.
- Check to see if all synchronization of RAID level-5 logical drives is completed. To check if any synchronization is running, use the ServeRAID Administration and Monitoring Utility. If arrays are currently being synchronized, a popup window appears showing the progress of the synchronization.

- a. On server A, run the utility IPSHAHTO.EXE from the ServeRAID NT Cluster Solutions Diskette by typing the following command from a Windows NT command line or from the Add/Remove Programs control panel:

```
a:\support\ipshahto
```

This utility will make all the shared logical drives available for access by server A.

Note: It is helpful if part of the volume label that you assign contains the drive letter that is assigned to a drive, such as drive_E.

- b. Using the Windows NT Disk Administrator, format all shared logical drives as NTFS (NT File System), assign volume labels, and specify drive letters to each.
- c. On server B, run the same utility, IPSHAHTO.EXE, to make all shared logical drives available for access by server B.

Note: The volume label is used as a guide to make sure the same letters are assigned on server B as were assigned on server A.

- d. Using the Windows NT Disk Administrator, assign the same drive letters that you assigned to each drive on server A.

e. Run IPSHAHTO from server A again to regain control of all shared logical drives to that server.

3. Install the Microsoft Cluster Server (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 Server (MSCS) software with ServeRAID.

a. Insert the *Windows NT Server Enterprise Edition Component CD ROM* into the CD-ROM drive on Server A and issue the following command, either from a command line or from Add/Remove Programs in the Control panel:

Note: The "localquorum" parameter must be typed in lowercase

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

Note: A local quorum drive will be installed on the server's local/boot drive. Later in this procedure, the local quorum drive must be moved to a ServeRAID logical array resource. Instructions are given when this needs to be done.

b. Install the Microsoft Cluster Server software following the instructions in the documentation that comes with the MSCS software.

Server A will restart when the installation is completed.

4. 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 ROM* into the CD-ROM drive on Server B and issue the

following command, either from a command line or from
Add/Remove Programs in the Control panel:

```
\MSCS\CLUSTER\I386\SETUP
```

- b. Install the Microsoft Cluster Server software following the instructions in the documentation that comes with the MSCS software to have Server B join the newly created high-availability cluster solution.

Server B will restart when the installation is completed.

5. Make server A the preferred owner of the cluster group by doing the following:
 - a. Run the Cluster Administrator program that comes with the MSCS software. Refer to the documentation that comes with the MSCS software for instructions.
 - b. Select the Cluster Group.
 - c. Display its properties window by selecting **File**; then, selecting **Properties**.
 - d. Click on **Modify** next to the Preferred Owner text box.
 - e. Select **Server A** and move it to the right side, into the Preferred Owners panel.
 - f. Click on **OK**.

Server A appears in the Preferred Owners text box.

- g. Make sure that the Cluster Group has failover enabled. For instructions on how to do this, refer to the documentation that comes with the Microsoft Cluster Services package.
- h. Close the connection to the cluster and open it again to ensure that the connection is through server A.

6. Install the IBM ServeRAID NT Cluster Solution utility by doing the following:

- a. From server A, run the SETUP.EXE program from the ServeRAID NT Cluster Solution Diskette and follow the instructions that appear on the screen.
- b. When complete, repeat the procedure from server B.

Upon completion of the installation of the utility on both servers, all shared ServeRAID groups and disk resources will be automatically created by the IBM Cluster Solution setup program and then brought online. Server A will be the current owner of these groups.

7. Move the quorum resource from Local Quorum to a ServeRAID disk resource by doing the following:

- a. From the Cluster Administrator, select the cluster you are using by clicking the left mouse button on the cluster name.
- b. Click the right mouse button to open the cluster.
- c. Select **Properties**; then, select the **Quorum tab**.
- d. In the **Quorum Resource** field, select a ServeRAID shared disk and click on **OK**.

Refer to the *Microsoft Cluster Solution Administrator's Guide* for more information.

You will need to ensure that the same server (for example, server A) owns the ServeRAID disk resources and the local quorum resource and that the quorum resource is moved to a RAID level-1 logical drive.

Shared disk resources can now be moved manually using the Microsoft Cluster Administrator utility or automatically during failover.

Special Installation Instructions

You must update the MSCS supplied by Microsoft for each cluster server. To obtain and install this update, do the following:

1. Download the file CLUSFIXI.EXE currently at the following web site:

`ftp://ftp.microsoft.com/bussys/winnt/winnt-public/fixes/
usa/nt40/hotfixes-postSP3/roll-up/CLUSTER/`

2. Copy CLUSFIXI.EXE file to the boot hard disk drive on each cluster server.
3. Run CLUSFIXI.EXE on server A. This will automatically update server A and restart (reboot) the system.
4. Allow server A to rejoin the cluster and be reported in the MSCS cluster administrator as "Up" and all resources are "online".
5. Run CLUSFIXI.EXE on server B. This will automatically update server B and restart (reboot) the system. Allow server B to rejoin the cluster and be reported in the MSCS cluster administrator as "Up" and all resources are "online".
6. You have completed the steps necessary for installing the MSCS update.

Creating a ServeRAID Disk Resource on a ServeRAID Server with MSCS

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 using the ServeRAID Configuration Program or the ServeRAID Administration and Monitoring Utility.

Before creating the ServeRAID logical disk cluster resource, assign a sticky drive letter to the shared logical disk drive on both servers using the Windows NT Disk Administrator Program. For information on how to assign sticky drive letters, see the MSCS documentation. Use the IPSHAHTO.EXE utility to move the logical drive to the other server so an identical drive letter can be assigned.

Perform this 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 system that owns the logical disk drive, run the Cluster Administration Utility program and make this server the preferred owner of **Cluster Group**.
2. Using the Cluster Administration Utility program, close the connection to the MSCS cluster and then reopen the connection. This ensures that your connection to the cluster is using the cluster server that owns the shared logical disk drive.
3. 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.
4. Bring the new ServeRAID disk group ONLINE.
5. 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.
6. Create a ServeRAID disk resource as part of the new group.
7. Bring the ServeRAID disk group ONLINE.

Important

When adding logical drives to an already existing configuration after you assign them as disk resources from within Cluster Administrator you must restart your computer. This restart will ensure that the MSCS software properly initializes the new drives.

Deleting and Later Recreating a ServeRAID Logical Disk Resource

This section describes how to delete and later, recreate, a ServeRAID logical disk 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.
2. To delete the resource, run the Cluster Administration Utility program from the cluster server that currently owns the resource or resource group.

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

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

3. Once the ServeRAID resource is deleted, 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 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 Cluster Administration Utility program from the cluster server that currently owns the resource group.

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

Create a new ServeRAID logical disk using the normal MSCS procedure for creating cluster resources. Refer to the documentation that comes with the MSCS package for instructions.

Uninstalling MSCS on ServeRAID Servers

Note: *Do not* use the uninstall procedure if you are upgrading MSCS on ServeRAID servers. To upgrade MSCS on ServeRAID servers, use “Upgrading MSCS on ServeRAID Servers” on page 39.

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 to uninstall MSCS from.
2. Stop the cluster services. To do this, select the **Services** icon in the NT control panel. Highlight the **Cluster Service** entry and then select the **Stop** button. For more information, refer to the documentation that came with the cluster services.
3. Select the Start button, then select **Control Panel** from the **Settings** menu.
4. Double-click **Add/Remove Programs**.
5. Select **Microsoft Cluster Service** from the list of installed components and then select **Add/Remove**.
Note: Ignore the warnings regarding the removal of IBM ServeRAID Logical Disks.
6. Follow the on-screen prompts by selecting **OK** or **Yes**.
7. Restart the system when prompted to complete uninstallation on MSCS.
8. Select the Start button, then select **Control Panel** from the **Settings** menu.
9. Double-click **Add/Remove Programs**.
10. Select **IBM ServeRAID NT Cluster Solution** from the list of installed components and then select **Add/Remove**.

You can also remove the IBM ServeRAID NT Cluster Solution by running the IBM ServeRAID NT Cluster Solution setup program and selecting the uninstall option.

11. Repeat step 1 to step 10 for the other server in the cluster.

Upgrading MSCS on ServeRAID Servers

This upgrade procedure allows the cluster to remain online at all time. To upgrade the cluster, perform the following steps on one server at a time. When all of the steps are completed for server A and both servers are reported in the MSCS cluster administrator as "Up" and all resources are "online", repeat the steps for server B.

Note: These instructions should be used on a cluster running Microsoft Cluster Server (MSCS) Release 1.00 and IBM ServeRAID Cluster Solution

Upgrading the IBM ServeRAID Administration and Monitoring Program

To upgrade the IBM ServeRAID Administration and Monitoring program, do the following:

Run SETUP.EXE program from the IBM ServeRAID Administration and Monitoring diskette and follow the instructions that appear on the screen.

Upgrading the IBM ServeRAID Microsoft Windows NT Cluster Solution

To upgrade the IBM ServeRAID Microsoft Windows NT Cluster Solution, do the following:

1. Select the Start button, then select **Control Panel** from the **Settings** menu.
2. Double-click **Services**.
3. Select **Microsoft Cluster Service** from the list of installed components and then select **Stop**.
4. Run SETUP.EXE program from *IBM ServeRAID Windows NT Cluster Solution* diskette and follow the instructions that appear on the screen.

Upgrading the IBM ServeRAID Windows NT Miniport Driver

To upgrade the IBM ServeRAID Windows NT Miniport Driver, do the following:

1. Insert the IBM ServeRAID Device Driver diskette in drive A:
2. Select the Start button, then select **Control Panel** from the **Settings** menu.
3. Double-click on **SCSI-Adapters**.
4. Select the **Drivers** tab.
5. Click on **Add**.
6. Click on **Have Disk**.
7. Use **A:\NT** for the path and then click on **OK**.
8. Click on **OK** when the IBM PC ServeRAID Adapter is highlighted
9. Click on **New**.
10. Use **A:\NT** for the path to install the driver.
11. Insert the IBM ServeRAID BIOS/Firmware diskette, version 3.00 or higher, in drive **A:** and then, shutdown Windows NT and restart the server from the diskette.

Upgrading the IBM ServeRAID BIOS/Firmware Diskette

To upgrade the IBM ServeRAID BIOS/Firmware, do the following:

1. Start the system from the *IBM ServeRAID BIOS/Firmware Diskette* Version 2.41 or later, if it was not previously started in step 11.
2. Flash the BIOS and download the firmware by selecting item 1 from the main menu.
3. When the download is complete, remove the diskette from the diskette drive and press **control-alt-delete** to restart (reboot) the system.

4. Allow the server to rejoin the cluster and when the server is reported in the MSCS cluster administrator as "Up" and all resources are "online", the upgrade for this server is complete.

Upgrade the Software for the Other Node in the Cluster

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

1. Ensure that server A is online in the cluster.
2. Repeat all of the steps between "Upgrading the IBM ServeRAID Windows NT Miniport Driver" on page 40 through "Upgrading the IBM ServeRAID BIOS/Firmware Diskette" on page 40 on the other server (server B) in the cluster.

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

Chapter 5. Installing a High-Availability Cluster Solution Using Vinca and Novell IntranetWare

To install a high-availability cluster solution using Novell IntranetWare, do the following:

1. Setup your ServeRAID adapter configuration as described in Chapter 3, "Configuring the ServeRAID Adapter for the High-Availability Cluster Solution" on page 10.
2. Install Novell NetWare on your servers. For information, refer to the documentation that comes with your Novell NetWare software package.
3. Install the Vinca IntranetWare software on your servers. For information, refer to the documentation that comes with *Vinca High Availability for NetWare and IBM ServeRAID II SCSI Adapter*.

Chapter 6. Monitoring and Updating a High-Availability Cluster Solution

You can use the ServeRAID Administration and Monitoring Utility to check the availability of a hard disk drive on another server or to add or remove a hard disk drive to the cluster.


There are two selections for high availability functions supported by the Administration and Monitoring Utility.

- Viewed Shared Devices
- Manage Merge Groups

Checking the Availability of a Drive

Note: Make sure the ServeRAID Administration and Monitoring Utility is installed on your server. If you need to install it, refer to the *IBM ServeRAID Adapter Installation and User's Guide* for instructions.

To check the availability of a hard disk drive on another server, do the following:

1. From the Administration and Monitoring Utility screen, click on the  icon on the tool bar or select **High Availability** from the **Advanced** pull-down menu and then select **View Shared Devices**.
2. Click on the box to **Enable View Shared Devices**.
3. Select the channel that the hard disk drive you want to check is connected to.
Note: This procedure can determine only if ONL Devices are available.
4. Select the bay of the hard disk drive that you are checking is connected to.

Click on **Refresh Now**. If the hard disk drive is available for use, it will appear as RDY (Ready) on the Administration and Monitoring Utility screen. If the hard disk drive is not available

for use, meaning that it is owned by the other server, it will appear as RSV (Reserved) on the Administration and Monitoring Utility screen.

All of the other hard disk drives owned by the other server also will appear as RSV at this time.

Note: Hot-spare drives owned by the other server will not display as RSV. Care must be used to make sure that the other servers hot-spare drives are not used inadvertently by both servers.


Adding Hard Disk Drives to a Cluster

If you add hard disk drives to a cluster, do the following:

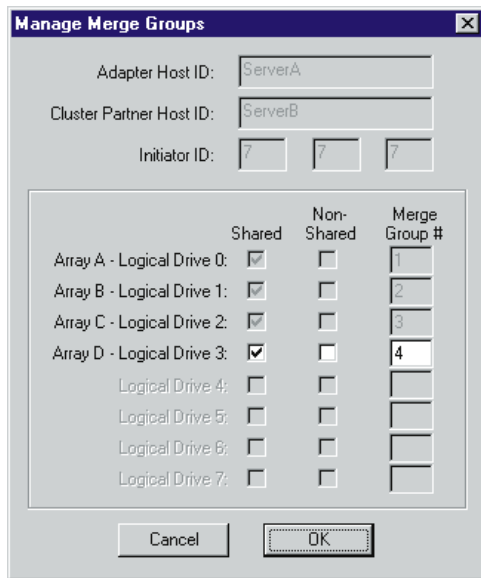
1. Install the hard disk drives in the server or shared expansion enclosure. For information, refer to the documentation that comes with the hard disk drive and the server.
2. Configure and setup the hard disk drives for the ServeRAID adapter. Refer to the *IBM ServeRAID Adapter Installation and User's Guide* for instructions.

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

3. After creating a logical drive, a screen similar to the following appears.



Select **Yes** and a screen similar to the following appears.




4. Select shared or non-shared for the logical drive you have just defined.
5. Select the Merge Group # you want to use for the drive you just defined.
6. Click on **OK** when you are done for the changes to take effect.

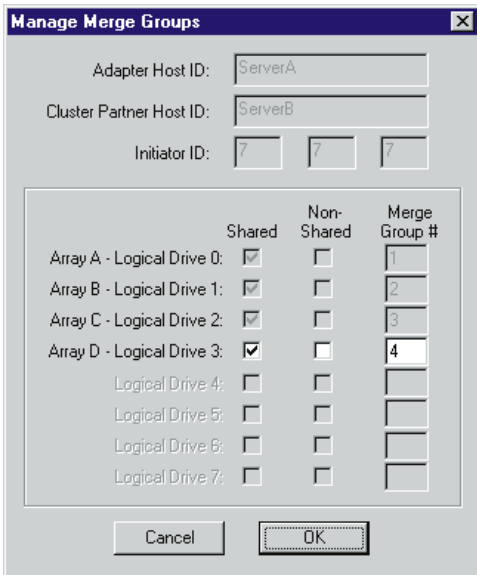
Note: Depending on the operating system, additional steps such as assigning drive letters, partitioning, formatting, and creating cluster resources will be required in order to use the new logical drive to store programs and data. For information, see “Creating a ServeRAID Disk Resource on a ServeRAID Server with MSCS” on page 34.

Attention: Drives that are configured by one adapter in the pair must appear as RDY (ready) drives to the other adapter in the pair. The same must also be true in reverse. Be sure to verify this consideration when drives are added after the servers are started. Restart the servers, or scan for new devices, or set the drives to RDY state using the RAID Administration and Monitoring Utility after adding new drives to a server. For more information, see “ServeRAID Considerations” on page 6.

Viewing Merge Group Numbers and Other Cluster Information

To view the merge group numbers and other cluster information,

click on the  icon on the tool bar or select **High Availability** from the **Advanced** pull-down menu and then select **Manage Merge Groups**. A screen similar to the following appears.



	Shared	Non-Shared	Merge Group #
Array A - Logical Drive 0:	<input checked="" type="checkbox"/>	<input type="checkbox"/>	1
Array B - Logical Drive 1:	<input checked="" type="checkbox"/>	<input type="checkbox"/>	2
Array C - Logical Drive 2:	<input checked="" type="checkbox"/>	<input type="checkbox"/>	3
Array D - Logical Drive 3:	<input checked="" type="checkbox"/>	<input type="checkbox"/>	4
Logical Drive 4:	<input type="checkbox"/>	<input type="checkbox"/>	
Logical Drive 5:	<input type="checkbox"/>	<input type="checkbox"/>	
Logical Drive 6:	<input type="checkbox"/>	<input type="checkbox"/>	
Logical Drive 7:	<input type="checkbox"/>	<input type="checkbox"/>	

In this case, this screen will be for viewing only. Information can be read, but not accessed.

Chapter 7. Troubleshooting a High Availability Cluster

There is the possibility that one of the servers, one of the ServeRAID adapters, or one of the hard disk drives in a cluster fails. If this occur, there are steps you can take to recover from the problem.

In this chapter are some steps 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.

Important

To perform the procedures in this chapter, you must use the *IBM ServeRAID Configuration Diskette*. If you do not have one, you must get one from the IBM web page at:

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

Recovering From a Hard Disk Drive Failure

Use the following procedure to recover from a hard disk drive failure:

1. Remove the failed hard disk drive from your server.

For information on how to remove the hard disk drive from your server, see the documentation that came with your server and the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*.

2. Install a new hard disk drive.

For information on how to install the new hard disk drive, see the documentation that came with your server and the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*.

3. Restore logical drives.

If the failed hard disk drive is part of a critical RAID level 1 or RAID level 5 logical drive, you must perform a rebuild operation to restore the logical drive to the OKY state. If the failed hard disk drive is part of a RAID level 0, 1, or 5, logical drive in the OFFLINE state the data may have been lost. You must set each DDD drive to the ONL state except the last DDD drive in RAID level 1 and 5 must be set to the RBL state; then, restore the logical drive from a backup copy.

Recovering From a Failed ServeRAID Adapter

The following configuration information is required to configure your new ServeRAID adapter:

- SCSI Bus Initiator_IDs
- Adapter Host_ID
- Cluster Partner's Host_ID
- Stripe Size
- Unattended Mode

If the ServeRAID adapter you are replacing *is* functional, you can obtain the SCSI Bus Initiator_IDs, the Adapter Host_ID, the Cluster Partner's Host_ID, the stripe size, and the unattended mode by starting your server with the *IBM ServeRAID Configuration Diskette* and selecting the **Display/Change Adapter Params** item from the Advanced Functions menu.

If the ServeRAID adapter is *not* functional, you will need to refer to a record of the settings that was made when the adapter was previously configured. If you do not have a record of the configuration information, the following hints **might** help you to assign the proper values.

- The SCSI Bus Initiator_IDs for a non-shared SCSI channel are set to 7 by default. For shared SCSI channels, the IDs will usually be 7 or 6 and must be different than the SCSI Bus Initiator_IDs for the corresponding SCSI channels of the cluster partner adapter. Therefore, if you can obtain the SCSI Bus Initiator_IDs for the corresponding cluster partner system, you can infer the correct SCSI Bus Initiator_IDs for this system. For

example, if the cluster partner system shared SCSI Bus Initiator_IDs were set to 7, then usually the adapter you are replacing would need to have its IDs set to 6.

To obtain the SCSI Bus Initiator_IDs from the corresponding cluster partner adapter, start the cluster partner system with the *IBM ServeRAID Configuration Diskette* and select the **Display/Change Adapter Params** option from the Advanced Functions menu.

- The proper settings for the Host_ID and cluster partner Host_ID of the adapter being replaced can also be determined by reading the settings from the cluster partner system by using the **Display/Change Adapter Params** option. The Host_ID for the new adapter must be identical to the cluster partner Host_ID value defined on the corresponding adapter in the cluster partner system, and the cluster partner Host_ID value for the new adapter must be identical to the Host_ID value defined on the corresponding adapter in the cluster partner system.

Use the following procedure to replace a failed ServeRAID Adapter from your high-availability configuration.

To replace the ServeRAID adapter, do the following:

1. Remove the failing ServeRAID adapter.

For instructions on how to remove the ServeRAID adapter, see the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*. Before you remove the adapter, be sure to:

- a. Turn off the server.
- b. Note which SCSI cables are connected to the SCSI channel connectors on the adapter.
- c. Note which PCI slot has the adapter installed.

2. Install the new ServeRAID adapter.

For instructions on how to install the ServeRAID adapter, see the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*. Be sure to install the adapter to

the same PCI slot from which you removed the failed ServeRAID adapter.

Important

Do **not** reconnect the SCSI channel cables to the adapter at this time.

3. Start the system from the *IBM ServeRAID Configuration Diskette* Version 3.00 or later.

4. Initialize the adapter configuration.

To initialize the adapter:

- a. Select **Advanced Functions** from the Main Menu.
- b. Select **Init/View/Synchronize Config**.
- c. Select **Initialize Config**.

5. Ensure that the adapter is at the latest BIOS level.

The BIOS level of the adapter is displayed after system POST when the adapter BIOS loads. The latest BIOS levels are available from the IBM website at

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

Once you have connected to this URL, search for *RAID BIOS*. Download and read the text file to determine the latest levels available. If your adapter BIOS is downlevel, download and apply the BIOS update.

6. Update the configuration parameters.

To update the configuration parameters:

- a. Start the system from the *IBM ServeRAID Configuration Diskette* Version 3.00 or later.
- b. Select **Advanced Functions** from the main menu.
- c. Select **Display/Change Adapter Params**.
- d. Using the settings that were assigned to the ServeRAID adapter you are replacing, select and configure each of the following parameters:
 - SCSI Bus Initiator_IDs

- Adapter Host_ID
- Cluster Partner's Host_ID

e. Select **Change RAID Parameters** from the Advanced Functions menu and enable unattended mode.

7. Shut down and turn off the server.
8. Reconnect the SCSI channel cables to the adapter.

Note: Be sure to connect the cables to the correct SCSI channels as noted in step 1 on page 49.

Important

If the ServeRAID adapter being replaced *is* the adapter that attaches to the server startup disk array, other non-shared disk arrays, or shared arrays which have not failed over to the other server, continue with the following steps; otherwise, restart the system normally.

9. If the adapter that was replaced attaches the operating system startup disk array for the system or if other non-shared disk arrays are attached to this adapter, start the system using the *IBM ServeRAID Configuration Diskette* Version 3.00 or later and then:
 - a. Select **Advanced Functions** from the Main Menu.
 - b. Select **Merge Group Management**.

c. Restore the adapter disk array configuration.

- To restore non-shared disk array configurations:

- 1) Select **Merge Group Management** from the Advanced Functions menu; then, press **Enter**.

A screen similar to the following appears.

IBM PC ServeRAID Adapter Disk Array Configuration Ver. X.XX			
Adapter Number: 1 Bus Number: 0 Host ID = Null Config			
Merge Group Management	Ch 1 Array Bay	Ch 2 Array Bay	Ch 3 Array Bay
1. Help	1 ONL A	1 HSP B	1
2. Create Shared Logical Drive ID	2 ONL A	2 ONL B	2
3. Create Non-Shared Logical Drive ID	3 ONL A	3	3
4. Delete Logical Drive Merge ID	4	4	4
5. Merge/Unmerge Logical Drive	5	5 ONL B	5
6. Exit	6	6 RDY	6
	7 INI	7 CDR	7 INI
	8	8 INI	8
	9	9	9
	10	10	10
	11	11	11
	12	12	12
	13	13	13
	14	14	14
	15	15	15
	16	16	16

Select a choice using the Up Arrow (↑) and Down Arrow (↓) keys and press Enter. Press Esc to return to previous menu.

- 2) Select **Merge/Unmerge Logical Drive** and then press **Enter**.

IBM PC ServeRAID Adapter Disk Array Configuration Ver. X.XX			
Adapter Number: 1 Bus Number: 0 Host ID = Null Config			
Merge/Unmerge Logical Drive	Ch 1 Array Bay	Ch 2 Array Bay	Ch 3 Array Bay
1. Help	1	1	1 ONL B
2. Merge Partner Logical Drive	2 ONL A	2	2 ONL B
3. Unmerge Partner Logical Drive	3	3	3
4. Merge Own Shared Logical Drive	4	4	4
5. Unmerge Own Shared Logical Drive	5	5	5
6. Merge Group ID Logical Drive	6	6	6
7. Unmerge Logical Drive	7 INI	7	7 INI
8. Exit	8	8 INI	8
	9	9	9
	10	10	10
	11	11	11
	12	12	12
	13	13	13
	14	14	14
	15	15	15
	16	16	16

Enter a unique number, 1 thru 8. Press the F10 key to make the changes. Press the Esc key to abort.

- 3) Select **Merge Own Non-Shared Logical Drive**.

- 4) Select from the **Group ID** field the Merge Group ID for the array

2xx

where xx is the shared SCSI Bus Initiator_ID, and then press **Enter**. The Merge Group ID value is typically 206 or 207. A message appears at the bottom of the screen saying

Merging own shared logical drive(s). Please wait...

- 5) When the process completes, a message appears saying

Merge/Unmerge operation completed successfully.
Press any key to continue.

- 6) Press **Esc** to return to the previous menu. Continue pressing **Esc** to return to the Main Menu. If the adapter you are replacing is the startup adapter, the system will now be able to startup the operating system properly.

- To restore shared disk array configurations:

Note: Usually all shared arrays will have failed over and will not need to be merged.

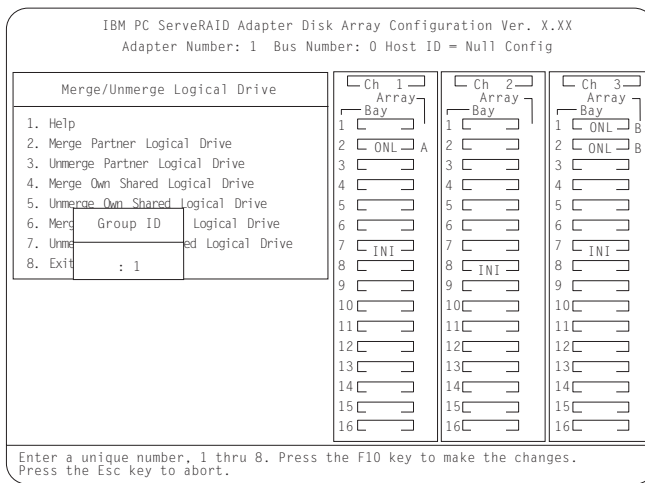
- 1) Select **Merge Group Management** from the Advanced Functions menu; then, press **Enter**.

A screen similar to the following appears.

IBM PC ServeRAID Adapter Disk Array Configuration Ver. X.XX Adapter Number: 1 Bus Number: 0 Host ID = Null Config			
Merge Group Management	Ch 1 Array Bay	Ch 2 Array Bay	Ch 3 Array Bay
1. Help	1 ONL A	1 HSP B	1
2. Create Shared Logical Drive ID	2 ONL A	2 ONL B	2
3. Create Non-Shared Logical Drive ID	3 ONL A	3	3
4. Delete Logical Drive Merge ID	4	4	4
5. Merge/Unmerge Logical Drive	5	5 ONL B	5
6. Exit	6	6 RDY	6
	7 INI	7 CDR	7 INI
	8	8 INI	8
	9	9	9
	10	10	10
	11	11	11
	12	12	12
	13	13	13
	14	14	14
	15	15	15
	16	16	16

Select a choice using the Up Arrow (↑) and Down Arrow (↓) keys and press Enter.
Press Esc to return to previous menu.

- 2) Select **Merge/Unmerge Logical Drive** and then press **Enter**.



- 3) Select **Merge Own Shared Logical Drive** for each shared array (Merge Group IDs in the range from 1 to 8) that has not failed over to the cluster partner system (for example, RAID level 5 arrays in critical or degraded state) to restore the configuration of these shared arrays.
- 4) Press **Esc** to return to the previous menu. Continue pressing **Esc** to return to the Main Menu.

Repeat this process for each shared array (Merge Group IDs in the range from 1 to 8) that has not failed over to the cluster partner system (for example, RAID level 5 arrays in critical or degraded state) to restore the configuration of these shared arrays.

Important

The *IBM ServeRAID Configuration Diskette* **must not** be used to perform failover and failback to merge or unmerge drives belonging to the other server. Failover and failback to merge or unmerge drives belonging to the other server is normally handled by the operating system software and cluster support software.

10. Restart your server.
11. When the ServeRAID Mini-Configuration prompt appears, press **Cntl+i** to go to the mini-configuration program.
12. Select the ServeRAID adapter you are working on by using the Up Arrow (↑) or Down Arrow (↓) and then press **Enter**
13. Select **Advanced Functions**.
14. Select **Configure BIOS Settings**.
15. Select **BIOS Unattended Mode** and press **Enter** to set it On.
16. Select **Save Configuration to ServeRAID Controller** and press **Enter**.

Once all array configurations have been restored, the server can be restarted normally.

Recovering From a Failed Server

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

Important

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

Use the following procedure to recover from a failed server:

1. Remove the failed server from your high-availability configuration.
2. Remove all hard disk drives from the disk drive 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 adapter is functional, remove it from the failed server.

As you remove the adapter from the failed server, be sure to:

- Note which SCSI cables are connected to the SCSI channel connectors on the adapter.

- Note which PCI slot has the adapter installed.

If you are replacing your failed server with an identical server, you can reinstall the ServeRAID adapter in an identical configuration and get your server up and running quickly.

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. Install the new ServeRAID adapter.

For instructions on how to install the ServeRAID adapter, see the *IBM ServeRAID-3H and ServeRAID-3L Ultra2 SCSI Adapters Installation and User's Guide*. If you are installing a ServeRAID adapter that was previously installed in the failed server, install the adapter in the same PCI slot as it was installed in the failed server.

Important

Do **not** reconnect the SCSI channel cables to the adapter at this time.

6. Configure the ServeRAID adapter in the new server.

Note: If you have installed the adapter in a server that is identical to the failed server and have installed it in the same PCI slot that it was installed in the failed system, you might not need to configure the adapter.

You will need the following information to configure your new ServeRAID adapter:

- SCSI Bus Initiator_IDs
- Adapter Host_ID
- Cluster Partner's Host_ID

If the ServeRAID adapter and server that you are replacing *are functional*, you can obtain this information by starting the server with the *IBM ServeRAID Configuration Diskette* and selecting the

Display/Change Adapter Params item from the Advanced Functions menu.

If the ServeRAID adapter or server is *not functional*, you will need to refer to the record of the settings that you made when the adapter was previously configured. If you do not have a record of the configuration information, the following hints **might** help you to assign the proper values.

- The SCSI Bus Initiator_IDs for a non-shared SCSI channel are set to 7 by default. For shared SCSI channels, the IDs will usually be 7 or 6 and must be different than the SCSI Bus Initiator_IDs for the corresponding SCSI channels of the cluster partner adapter. Therefore, if you can obtain the SCSI Bus Initiator_IDs for the corresponding cluster partner system, you can infer the correct SCSI Bus Initiator_IDs for this system. For example, if the cluster partner system's shared SCSI Bus Initiator_IDs were set to 7, then usually the adapter you are replacing would need to have its IDs set to 6.

To obtain the SCSI Bus Initiator_IDs from the corresponding cluster partner adapter, start the cluster partner system with the ServeRAID Configuration Diskette and select the **Display/Change Adapter Parameters** option from the Advanced Functions menu.

- The proper settings for the Host_ID and Cluster Partner's Host_ID of the adapter being replaced can also be determined by reading the settings from the cluster partner system by using the **Display/Change Adapter Params** option. Ensure that the Host_ID for the new adapter is identical to the Cluster Partner's Host_ID value defined on the corresponding adapter in the cluster partner system, and that the Cluster Partner's Host_ID value for the new adapter is identical to the Host_ID value defined on the corresponding adapter in the cluster partner system.

7. Start the system from the *IBM ServeRAID Configuration Diskette* Version 3.00 or later.

8. Initialize the adapter configuration.

To initialize the adapter:

- a. Select **Advanced Functions** from the main menu.
- b. Select **Init/View/Synchronize Config.**
- c. Select **Initialize Config.**

9. Ensure that the adapter is at the latest BIOS level.

The BIOS level of the adapter is displayed after system POST when the adapter BIOS loads. The latest BIOS levels are available from the IBM website at

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

Once you have connected to this URL, search for *RAID BIOS*. Download and read the text file to determine the latest levels available. If your adapter BIOS is downlevel, download and apply the BIOS update.

10. Update the configuration parameters.

To update the configuration parameters:

- a. Start the system from the *IBM ServeRAID Configuration Diskette* Version 2.40 or later.
- b. Select **Advanced Functions** from the main menu.
- c. Select **Display/Change Adapter Params.**
- d. Using the settings that were assigned to the ServeRAID adapter you are replacing, select and configure each of the following parameters:
 - SCSI Bus Initiator_IDs
 - Adapter Host_ID
 - Cluster Partner's Host_ID
- e. Select **Change RAID Parameters** from the Advanced Functions menu and enable unattended mode.

11. Shut down the system and reconnect the SCSI channel cables to the adapter. Be sure to connect the cables to the correct SCSI channels as noted in step 3 on page 57.

Important

If the ServeRAID adapter being replaced is not the adapter that attaches to the server startup disk array or other non-shared disk arrays, you do not need to perform any of the following steps. The system can now be restarted normally.

12. If the adapter that was replaced attaches to the operating system startup disk array for the system or if other non-shared disk arrays are attached to this adapter, start the system using the *IBM ServeRAID Configuration Diskette* Version 2.40 or later and then:
 - a. Select **Advanced Functions** from the main menu.
 - b. Select **Merge Group Management**.
 - c. Restore the adapter disk array configuration.
 - To restore non-shared disk array configurations:
 - 1) Select **Merge Group Management** from the Advanced Functions menu; then, press **Enter**.
 - 2) Select **Merge/Unmerge Logical Drive** and then press **Enter**.
 - 3) Select **Merge Own Non-Shared Logical Drive**.
 - 4) Type in the **Group ID** field the Merge Group ID for the array
2xx
where xx is the shared SCSI Bus Initiator_ID, and then press **Enter**. The Merge Group ID value is typically 206 or 207. A message appears at the bottom of the screen saying
Merging own shared logical drive(s). Please wait...

5) When the process completes, a message appears saying

Merge/Unmerge operation completed successfully.
Press any key to continue.

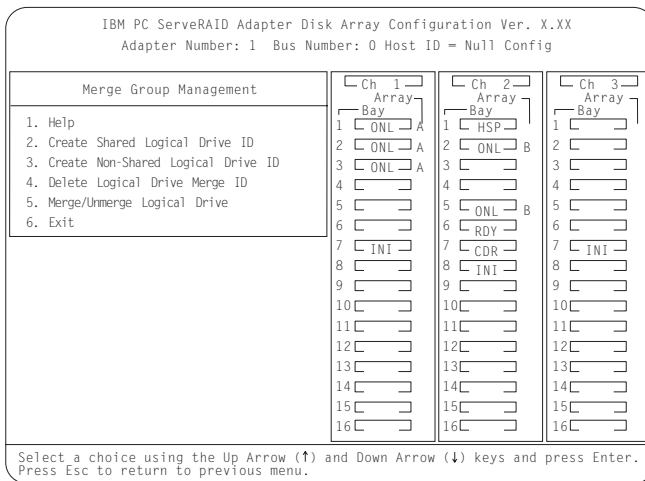
6) Press **Esc** to return to the previous menu. Continue pressing **Esc** to return to the Main Menu. If the adapter you are replacing is the boot adapter, the system should now be able to startup the operating system properly.

- To restore shared disk array configurations:

Note: Usually all shared arrays will have failed over and will not need to be merged.

1) Select **Merge Group Management** from the Advanced Functions menu; then, press **Enter**.

A screen similar to the following appears.



2) Select **Merge/Unmerge Logical Drive** and then press **Enter**.

Merge/Unmerge Logical Drive			Ch 1 Array Bay	Ch 2 Array Bay	Ch 3 Array Bay
1. Help			1	1	1
2. Merge Partner Logical Drive			2 ONL A	2	2 ONL B
3. Unmerge Partner Logical Drive			3	3	3
4. Merge Own Shared Logical Drive			4	4	4
5. Unmerge Own Shared Logical Drive			5	5	5
6. Merge Group ID Logical Drive			6	6	6
7. Unmerge Group ID Logical Drive			7 INI	7	7 INI
8. Exit : 1			8	8 INI	8
			9	9	9
			10	10	10
			11	11	11
			12	12	12
			13	13	13
			14	14	14
			15	15	15
			16	16	16

Enter a unique number, 1 thru 8. Press the F10 key to make the changes.
Press the Esc key to abort.

- 3) Select **Merge Own Shared Logical Drive** for each shared array (Merge Group IDs in the range from 1 to 8) that has not failed over to the cluster partner system (for example, RAID level 5 arrays in critical or degraded state) to restore the configuration of these shared arrays.
- 4) Press **Esc** to return to the previous menu. Continue pressing **Esc** to return to the Main Menu.

Repeat this process for each shared array (merge group IDs in the range from 1 to 8) that has not failed over to the cluster partner system (for example, RAID level 5 arrays in critical/degraded state) to restore the configuration of these shared arrays.

Important

The *IBM ServeRAID Configuration Diskette* **must not** be used to perform failover and failback to merge or unmerge drives belonging to the other server. Failover and failback to merge or unmerge drives belonging to the other server is normally handled by the operating system software and cluster support software.

13. Restart your server.

Once all array configurations have been restored, the server can be restarted normally.

Troubleshooting ServeRAID Subsystems in a Cluster Environment

Following is a series of common problems and solutions that can help you troubleshoot your high-availability solution.

Problem: ServeRAID Administration Utility program shows physical devices as DDD state.

Action:

1. Enable the shared disk display function to allow disks that have been moved or have failed over to the other node in the cluster to be displayed as RSV state instead of DDD state.

Note: It is normal for disks that have been moved or have failed over to be displayed in the DDD state if the shared disk display feature has not been enabled. In this case the disks shown in the DDD state are not really defective.

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 disk and perform a rebuild operation.

Problem: ServeRAID shared logical drives do not failover properly.

Action:

1. Ensure that the resource type of each ServeRAID adapter shared disk resource is "IBM ServeRAID logical disk".

If the resource type is shown as "physical disk", the localquorum option was not specified properly when MSCS was installed.

To correct this problem, you must reinstall a high-availability cluster solution using Microsoft Windows NT. Refer to Chapter 4, "Installing a High-Availability Cluster Solution Using Windows NT" on page 28 for instructions.

2. Ensure that shared SCSI buses on the ServeRAID adapter pair are connected such that corresponding SCSI channels are connected (for example SCSI channel 1 on the adapter in the first cluster node server is connected to SCSI channel 1 on the adapter in the second cluster node server, channel 2 is connected to channel 2, and so forth).
3. Ensure that physical SCSI disks that contain logical drives are all connected to shared SCSI channels.

4. Ensure that there are no more than 8 shared logical disk drives defined per pair of ServeRAID adapters for use in your cluster.
5. For Windows NT Server clusters, ensure that SCSI channel 3 of the adapter pair that attaches to the ServeRAID logical drive that has been designated as the NT Cluster Quorum Resource is used for arbitration. Also, ensure that it is connected from the first cluster node to SCSI channel 3 in the second cluster node and that there are no SCSI devices connected to that channel.

The SCSI heartbeat connection *must* be connected to the third channel of the ServeRAID adapter pair that has the quorum drive connected to it. No disks can be installed on this heartbeat channel. If you choose to move the quorum drive to another ServeRAID adapter, ID level 5 you must also move the SCSI heartbeat cable on both servers to the new quorum ServeRAID adapter pair. For more information, see "ServeRAID Considerations" on page 6.

6. The quorum disk can be placed on any ServeRAID channel shared by cluster node server.
7. Make sure each shared logical drive has a Merge ID assigned. Merge IDs **must** be in a range of 1 to 8
8. Make sure each ServeRAID adapter has been assigned a unique Host ID and that each ServeRAID adapter has its cluster partner Host ID assigned properly to correspond to the ServeRAID adapter in the other cluster node server that is attached to the shared SCSI buses.
9. Check for loose shared SCSI bus cables.
10. Ensure that SCSI repeater cards in Model 3518 or 3519 disk expansion enclosures are at the latest revision level.
11. Ensure that physical disks that are expected to be moved or failover show up as RDY or RSV state on the node that is attempting to take over control of these disks.

Problem: RAID level 5 logical disks cannot be accessed by the operating system after a failover.

Action: Use the ServeRAID Administration Utility program to check the state of the logical disk drive to ensure that it is not blocked. Using the utility program, select the logical disk drive and look for Blocked state **Yes**. If the logical disk drive is blocked, make sure all physical disks that are part of the logical drive are in the ONL state. If all physical disks are not in the ONL state, then a disk might have gone bad during a failover or during the resynchronization process after a failover. Data integrity cannot be guaranteed in this case and the array has been blocked to prevent the possibility of incorrect data being read from the logical drive.

Reinitialize and synchronize the logical drive and restore the data from a back up source. Depending on the type of data contained on the logical drive and the availability of a recent back up copy, you can unblock the drive and continue normal operation or replace/rebuild one or more DDD disks. However, if you do not reinitialize, synchronize and restore the drive, be aware that some data on the disk drive could be lost or corrupted.

Problem: If one of the cluster nodes fails and the surviving node takes over the cluster resources, occasionally one or more of the IP address resources will stay in the ONL pending state for several minutes after moving over to the surviving node. After this, the resource will go to the failed state and the following error message will be displayed in the surviving node's system log (as viewed with the Event Viewer).

For Example: NT Event Log Message:

Date: ???	Event ID: 1069
Time: ???	Source: ClusSvc
User: N/A	Type: Error
Computer: ???	Category: (4)

Description:

Cluster resource 'ip address resource name' failed

Action: Do the following:

1. Using the right mouse button, click on the IP resource in the Cluster Administrator.
2. Select **Properties** from the **General** tab.
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: After one of the cluster nodes (servers) has been shutdown normally and the surviving node (server) takes over the cluster resources, occasionally one or more of the IBM ServeRAID logical disk resources will stay in the 'online pending' state for several minutes, after moving over to the surviving node (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 node's (server's) system log (as viewed with the Event Viewer).

For Example: NT Event Log Message:

Date: ???	Event ID: 1069
Time: ???	Source: ClusSvc
User: N/A	Type: Error
Computer: ???	Category: (4)

Description:

Cluster resource 'IBM ServeRAID Logical Disk name' failed.

Action: No action is necessary to bring the resource online after the failover. MSCS will successfully re-attempt to bring this resource online on the surviving node within about four minutes.

Problem: 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 asking if you want to perform an upgrade.

Action: You must delete *C3E76E53-F841-11D0-BFA1-08005AB8ED05* registry key. To delete the registry key, do the following:

1. Select RUN
2. Type: REGEDIT and click OK. The Registry Editor screen appears.
3. Select "HKEY_CLASSES_ROOT\CLSID" and delete *C3E76E53-F841-11D0-BFA1-08005AB8ED05*
4. Reinstall the ServeRAID Windows NT Cluster Solution. Refer to Chapter 4, "Installing a High-Availability Cluster Solution Using Windows NT" on page 28 for instructions.

Problem: The following error message is displayed when running the IPSHAHTO program on a node. The warning message is: *"Warning: CONFIG_SYNC with 0xA0 command FAILED on Adapter #"* and one or more HSP/SHS devices are defined on adapter pairs or READY (RDY) devices are not part of any logical array configuration on an adapter pair in a cluster.

Action: If all shared disk resources moved successfully when running the IPSHAHTO program, it is safe to ignore the error message and no further action is required.

If shared disk resources fail to move when running the IPSHAHTO program, then perform a low-level format on all HSP/SHS and RDY devices that are not part of any logical array configuration on an adapter pair in a cluster. Refer to the instructions on low-level formatting HSP/SHS drives in *ServeRAID Adapter Installation and User's Guide (P/N4227022)* for further details.

Problem: Array identifiers and logical drive numbers might change during a failover condition.

Action: By design, the array identifiers and logical drive numbers may change during a failover condition. Consistency between the merge identifiers and Windows NT sticky drive letters is maintained, while the ordering process during a failover condition is controlled by the Microsoft Cluster Management Software and the available array identifiers on the surviving node.

Appendix A. Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Department 80D, P.O. Box 12195, 3039 Cornwallis, Research Triangle Park, NC 27709, U.S.A. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries:

Netfinity IBM OS/2

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of the Microsoft Corporation.

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



Part Number: 33L3890

Printed in U.S.A.

33L3890

