

IBM Netfinity®  
High-Availability Cluster Solutions  
using the IBM ServeRAID™ -3H, -3HB, and -3L  
Ultra2 SCSI Controllers

**Installation and User's Guide**





IBM Netfinity®  
High-Availability Cluster Solutions  
using the IBM ServeRAID™ -3H, -3HB, and -3L  
Ultra2 SCSI Controllers

## **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 38.

**Third Edition (December 1999)**

**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, 1999. 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

<b>About this book</b> .....	iv
How this book is organized .....	iv
Related publications .....	v
Downloadable files from the World Wide Web .....	v
<b>Chapter 1. Introduction to IBM High-Availability Cluster Solutions</b> .....	1
Preparing to install or modify a High-Availability Cluster Solution .....	3
Terms and definitions .....	4
ServeRAID considerations .....	4
<b>Chapter 2. Removing a quorum arbitration cable and upgrading an existing High-Availability Cluster Solution</b> .....	7
<b>Chapter 3. Configuring ServeRAID controllers for the IBM High-Availability Cluster Solution</b> ..	8
Configuring the controller for Windows NT 4.0 .....	8
Configuring the controller for Vinca and Novell IntranetWare .....	9
Configuring a controller for clustering .....	10
Configuring ServeRAID arrays and logical drives .....	13
<b>Chapter 4. Installing the Windows NT 4.0 cluster software</b> .....	18
Installing the Microsoft Cluster Software .....	18
Creating a ServeRAID disk resource on a ServeRAID server with MSCS .....	20
Deleting and later recreating a ServeRAID logical disk resource .....	21
Uninstalling MSCS on ServeRAID servers .....	22
Updating ServeRAID software in an MSCS cluster environment .....	22
<b>Chapter 5. Monitoring and updating a High-Availability Cluster Solution using the ServeRAID Manager program</b> .....	25
Checking the ownership of a physical drive .....	25
Adding logical drives to a cluster .....	26
Validating a cluster .....	27
Viewing Merge Group numbers and other cluster information .....	28
<b>Chapter 6. Troubleshooting a High-Availability Cluster</b> .....	29
Rebuilding a defunct drive .....	29
Steps for recovering from defunct drives .....	29
Rebuilding a hot-swap drive .....	30
Recovering from a failed server .....	30
Recovering from a failed ServeRAID controller .....	31
Replacing a ServeRAID controller .....	32
Troubleshooting ServeRAID subsystems in a cluster environment .....	34
<b>Appendix A. Notices</b> .....	38
Trademarks .....	39

---

## About this book

This book provides information and instructions for setting up a high-availability clustering solution using the IBM® Netfinity® ServeRAID™ -3H, ServeRAID-3HB, and ServeRAID 3-L Ultra2 SCSI Controller, herein called the ServeRAID controller.

This book is intended for experienced users who will be involved with setting up clustering and high-availability solutions for their computer installations with Windows NT 4.0 and Novell IntranetWare 4.11.

---

## How this book is organized

Chapter 1, "Introduction to IBM High-Availability Cluster Solutions," provides information you need to know before you start the installation and setup of a high-availability cluster solution or upgrade an existing one. This information includes an overview of high-availability cluster solutions, definitions of important terms used in this manual, and some considerations you need to be familiar with for the ServeRAID controller.

Chapter 2, "Removing a quorum arbitration cable and upgrading an existing High-Availability Cluster Solution," provides the information on removing a quorum arbitration cable and upgrading an existing high availability cluster solution.

Chapter 3, "Configuring ServeRAID controllers for the IBM High-Availability Cluster Solution," provides the information about configuring the ServeRAID controllers in the two servers used in the high-availability cluster solution for both Windows NT and Novell IntranetWare. It also includes information on defining all shared and non-shared logical drives and setting Merge Group numbers.

Chapter 4, "Installing the Windows NT 4.0 cluster software," provides instructions for setting up a high-availability cluster solution using Windows NT cluster software.

Chapter 5, "Monitoring and updating a High-Availability Cluster Solution using the ServeRAID Manager program," provides information about how to use the ServeRAID Manager program to determine the availability of logical drives in a cluster. It also includes procedures to use if you add or remove a logical drive in your high-availability cluster solution, or to uninstall your high-availability cluster solution.

Chapter 6, "Troubleshooting a High-Availability Cluster," provides information on how to recover from some potential problems, such as a defective ServeRAID controller or a defective physical 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* 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.
- The *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide* provides instructions for the installation of the ServeRAID controller, as well as instructions on how to use and maintain RAID configurations.

You can obtain these publications from the IBM Support Web site. See “Downloadable files from the World Wide Web” for additional information.

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.

---

## Downloadable files from the World Wide Web

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

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

To access Clustering and ServeRAID support, follow these steps:

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

- |                            |  |
|----------------------------|--|
| <b>Downloadable files</b>  | Download the latest versions of Clustering software, the ServeRAID Manager program, BIOS/firmware, device driver updates, and other important information. |
| <b>Hints and tips</b>      | Obtain useful information for the IBM Clustering and ServeRAID products and troubleshooting potential problems.  |
| <b>Online publications</b> | Download the Installation and User's Guides, References, Red Books, and other IBM publications.  |

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



---

## 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-server cluster, where both servers, 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.

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

An *adapter* is software that enables different software components or products to interact with one another. A *controller* is a device that coordinates and controls the operation of one or more input and output devices, such as workstations, and synchronizes the operation of such devices with the operation of a system as a whole.

IBM high-availability cluster solutions, as discussed in this manual, use the IBM ServeRAID-3H, ServeRAID-3HB, or ServeRAID-3L Ultra2 SCSI controller. IBM high-availability cluster solutions also use IBM expansion enclosures, such as the IBM Enterprise Expansion Enclosure (Model 3518), the IBM PC Server 3519 Rack Expansion Enclosure, the IBM Netfinity EXP15 Rack Expansion Enclosure, or the IBM Netfinity EXP10 Rack Expansion Enclosure.

You can install IBM high-availability cluster solutions 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
- Novell High Availability Server 1.0

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

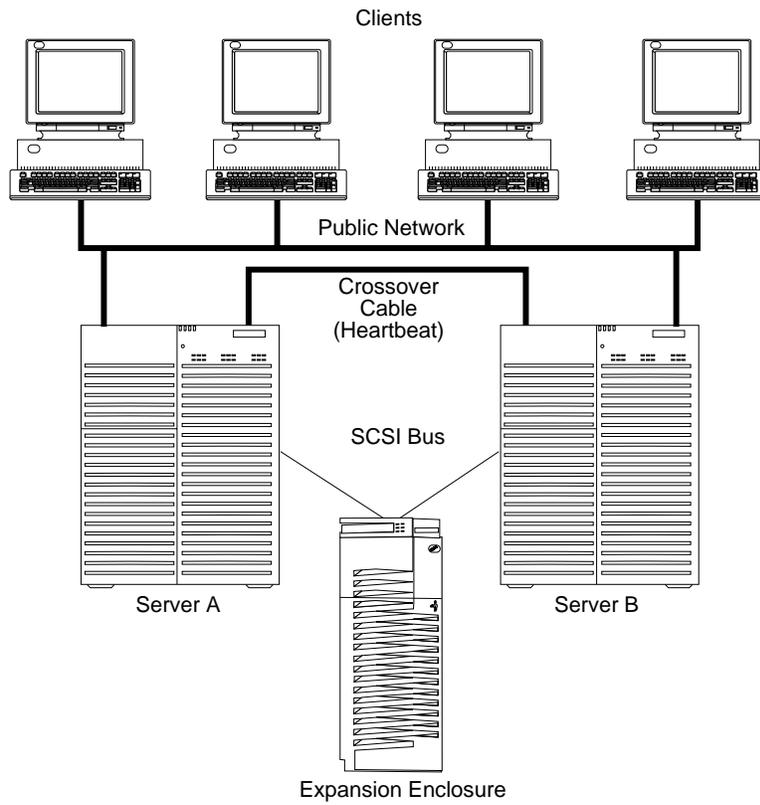


Figure 1. A high-level view of a cluster

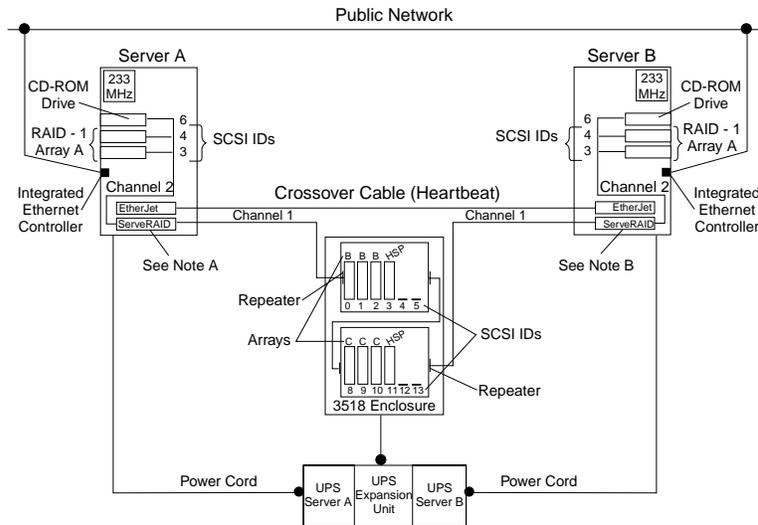


Figure 2. A detailed view of a cluster

**Note A: Server A**

- Shared Channel 1 - SCSI ID 6
- Non-shared Channel 2 - SCSI ID 6
- Non-shared Channel 3 - SCSI ID 7

**Note B: Server B**

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

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

---

## Preparing to install or modify a High-Availability Cluster Solution

**Note:** If you are going to upgrade an existing High-Availability Cluster Solution and are already familiar with the concepts of a cluster, go to Chapter 2, “Removing a quorum arbitration cable and upgrading an existing High-Availability Cluster Solution” on page 7. If you are setting up your first High-Availability Cluster Solution, continue reading this section.

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 Controller.

**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 the *Vinca High Availability for NetWare manual* and the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide*.

## Terms and definitions

The following terms will be used in this manual.

**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 about failover, refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide*.

**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.

## ServeRAID considerations

It is important to understand the following ServeRAID controller 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 the *Vinca High Availability for NetWare* manual and the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide*.

- The ServeRAID controller 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 level-1 logical drive.

**Note:** You should define hot-spare drives in your array to minimize the time that logical drives remain in critical state.

- Every logical drive that is shared by the two servers *must* have its cache policy set to write-through mode to ensure 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, -3HB, and -3L Ultra2 SCSI Controllers 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 controller must have a unique Controller Name assigned.
- Do not attempt to change the Controller Name on both ServeRAID controllers in a pair at the exact same time. Doing so can cause a problem, such as a server *lock-up*.
- Each ServeRAID controller, and controller pair, must have a different SCSI channel initiator ID assigned to the shared channels. For example, one server, Server A, could be set to ID 6, and the other server, Server B could be set to ID 7. This is described during the setup instructions.
- The stripe-unit size of each ServeRAID controller in a pair must be set to the same value (8 KB, 16 KB, 32 KB, or 64 KB).
- The SCSI channels in a controller pair must be cabled so that the same channels are shared between controllers.
- All disks that will be shared must be connected to the shared channels. This includes hot-spare drives, as well as data drives.
- Only one logical drive must be created for each array.
- Merge Group numbers in the range 1-8 must be assigned to each logical drive that will be shared. Merge Group numbers must be unique for each shared logical drive in the cluster. Merge Group numbers 206 or 207 must be assigned to the non-shared logical drives.
- If you are starting (booting) the operating system from a shared controller, define the first logical drive as the startup drive and assign a non-shared Merge Group number, for example, 206 for Server A.
- Do not connect nondisk devices, such as CD-ROMs to shared channels. They will be accessible to both servers and conflicts can arise.
- Hot-spare drives are **not** shared between controllers in a cluster pairing. If you want hot-spare protection, **each** ServeRAID controller must have a hot-spare drive defined.

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

- The total number of logical drives per controller 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 logical drives will be able to failover if necessary.
- Physical drives that are configured by one controller in the pair must appear as Ready drives to the other controller in the pair. Be sure to verify this consideration when drives are added after the servers are started. Restart both servers, or scan for new devices using the ServeRAID Manager program on both servers, to make sure the added physical drives appear as ready drives on both servers before creating logical drives.
- All shared physical drives must be part of a logical drive or defined as a hot spare to one or the other servers. All other physical drives should be removed from the shared disk enclosure to ensure proper operation.

- 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 and is no longer assessable.
- When a logical drive spans across multiple SCSI channels and a failure within the drive subsystem occurs that is unique to a channel (for example, a disconnected cable), the entire physical array will be identified as unsuccessful even though access from the surviving server can occur. Therefore, you might want to consider, if you have small arrays, not spanning across multiple channels.

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

---

## Chapter 2. Removing a quorum arbitration cable and upgrading an existing High-Availability Cluster Solution

If you already have an existing High-Availability Cluster Solution set up and running using a quorum arbitration cable (SCSI Heartbeat Cable), and you want to remove it from your cluster to free a SCSI channel for future use with more storage devices, do the following:

### Important

Before starting this procedure, you must obtain the new software for the IBM ServeRAID controllers. If you need to update any RAID software, you must update all RAID software at the same time to ensure compatibility with all of the levels of software. This includes:

- BIOS/firmware
- Device drivers
- ServeRAID Manager
- Clustering software

You can obtain this software from the IBM Support Web site. See “Downloadable files from the World Wide Web” on page v for additional information.

1. Shut down both servers in the cluster.
2. Follow the upgrade procedure “Updating ServeRAID software in an MSCS cluster environment” on page 22.
3. Remove the Quorum Arbitration cable (SCSI heartbeat cable) on channel 3 of both servers.

Once you have removed the quorum arbitration cable, you can now use that SCSI channel to add more physical drives. For information on how to add a hard disk drive to a cluster, go to “Adding logical drives to a cluster” on page 26.

---

## Chapter 3. Configuring ServeRAID controllers for the IBM High-Availability Cluster Solution

To install a high-availability cluster solution, you must first configure the ServeRAID controller, define logical drives, and set up Merge Group numbers for your servers.

- If you are installing high-availability using Windows NT 4.0, refer to “Configuring the controller for Windows NT 4.0.”
- If you are installing high-availability using Vinca and Novell IntranetWare, refer to “Configuring the controller for Vinca and Novell IntranetWare” on page 9.

### Important

Before beginning this installation, consider the following:

- Read “ServeRAID considerations” on page 4.
- Use the following CD and diskette, together as a set, with the high-availability cluster solution:
  - *IBM ServeRAID Support CD*
  - *IBM ServeRAID Device Driver Diskette*

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

- BIOS/firmware
- Device drivers
- ServeRAID Manager
- Clustering software

Ensure that the ServeRAID controllers have the latest software installed. You must have version 3.5, or higher, as a minimum. You can obtain this software from the IBM Support Web site. See “Downloadable files from the World Wide Web” on page v for additional information. If you are using Microsoft Cluster Server, be sure that you use a version that has been certified by Microsoft as noted on the IBM Support Web site.

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

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

---

### Configuring the controller for Windows NT 4.0

If you are installing high-availability using Windows NT 4.0, follow these steps:

1. If necessary, perform any operating system shut down procedure on **both** servers that you are going to use in the cluster.
2. Turn off **both** servers.
3. If you have not done so, connect any shared external drive enclosures between both servers; then, turn on the enclosures.
4. Follow the procedure “Configuring a controller for clustering” on page 10 for each shared controller in **Server A**.

5. Continue with the procedure “Configuring ServeRAID arrays and logical drives” on page 13 for **Server A**.
6. Remove the *IBM ServeRAID Support CD* from the CD-ROM drive.
7. Shut down **Server A**.
8. Follow the procedure “Configuring a controller for clustering” on page 10 for each shared controller in **Server B**.
9. Remove the *IBM ServeRAID Support CD* from the CD-ROM drive.
10. Shut down **Server B**.
11. Restart **Server A** and **Server B**.
12. Continue to Chapter 4, “Installing the Windows NT 4.0 cluster software” on page 18.

---

## Configuring the controller for Vinca and Novell IntranetWare

If you are installing high-availability using Vinca and Novell IntranetWare, following these steps:

1. If necessary, perform any operating system shut down procedure on **both** servers that you are going to use in the cluster.
2. Turn off **both** servers.
3. If you have not done so, connect any shared external drive enclosures between both servers; then, turn on the enclosures.
4. Follow the procedure “Configuring a controller for clustering” on page 10 for each shared controller in **Server A**.
5. Continue with the procedure “Configuring ServeRAID arrays and logical drives” on page 13 for **Server A**.  
**Note:** In a typical Novell environment, create the logical drives on the server that will maintain primary ownership. For example, create one RAID level-5 logical drive on Server A and one RAID level-5 logical drive on Server B.
6. Remove the *IBM ServeRAID Support CD* from the CD-ROM drive.
7. Shut down **Server A**.
8. Follow the procedure “Configuring a controller for clustering” on page 10 for each shared controller in **Server B**.
9. Continue with the procedure “Configuring ServeRAID arrays and logical drives” on page 13 for **Server B**.  
**Note:** If you define an online drive or hot-spare drive for Server A, you cannot use that drive when creating arrays or hot-spare drives on Server B.
10. Remove the *IBM ServeRAID Support CD* from the CD-ROM drive.
11. Shut down **Server B**.
12. Restart **Server A** and **Server B**.
13. Install Novell NetWare on your servers. For information, refer to the documentation that comes with your Novell NetWare software package.
14. 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 Controller*.

## Configuring a controller for clustering

**Note:** For complete information on using the ServeRAID Manager program, refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide*.

To configure a controller, do the following:

1. Start the server with the *IBM ServeRAID Support CD* (or the configuration CD that was included with your server) in the CD-ROM drive. The ServeRAID Manager program starts. A screen similar to the following appears:

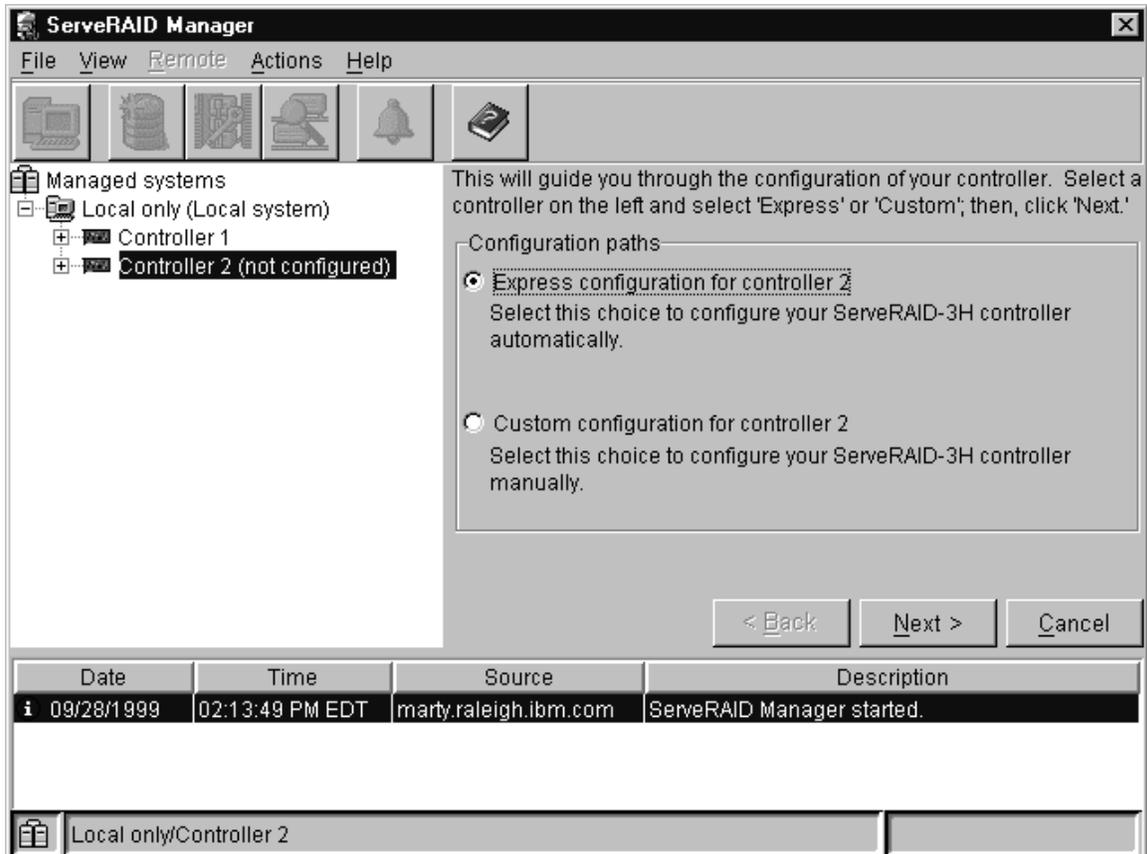


Figure 3. Configuration mode screen

2. Click **Cancel**.
3. Click the ServeRAID controller in the Main Tree that you are configuring.

**Important**

If you already have logical drives defined that you wish to keep, go to step 7 on page 11. The steps prior to step 7 will cause loss of data by deleting all logical drives already defined.

The ServeRAID controller should be restored to factory default settings only one time, when the controller is first installed. Restoring to factory default settings removes all configuration information from the controller. This information includes the SCSI initiator IDs, the Controller Name and the Partner Name, and the logical drive information. This point is especially important if the Windows NT system boot drive is attached to the IBM ServeRAID controller. If the restore to factory default settings occurs after the boot logical drive has been created, this logical drive will be removed and the server will not be able to start up into Windows NT.

4. Select **Restore to factory default settings** from the Actions menu.
5. If a confirmation window appears, click **Yes**. When complete, a message appears at the bottom of the screen indicating the restore to factory default settings was successful.
6. Click the ServeRAID controller in the Main Tree that you are configuring.
7. Click  on the tool bar. The Configure for clustering window appears.

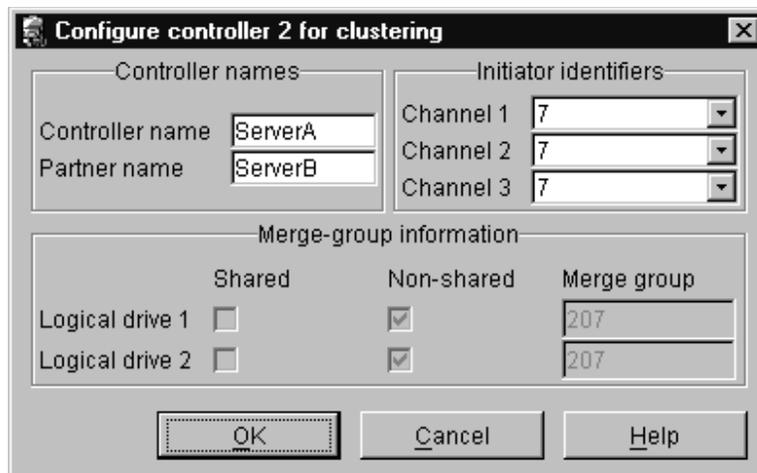


Figure 4. Configure for clustering window

8. Fill in the following information:

**Important**

Keep the following in mind:

- The Controller Name and Partner Name are case-sensitive.
- Both servers cannot have the same SCSI channel initiator IDs on shared channels.
- If your configuration will contain multiple ServeRAID controller pairs, each Controller Name must be unique.

**Controller Name**

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

**Partner Name**

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

**SCSI channel initiator IDs** The system presets these to seven (7). For System A, set each ID to six (6) for shared channels. For System B, set each ID to seven (7) for shared channels.

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

## Configuring ServeRAID arrays and logical drives

### Important

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

1. Click the ServeRAID controller that you want to configure in the Main Tree.

2. Click  on the tool bar. A screen similar to the following appears:

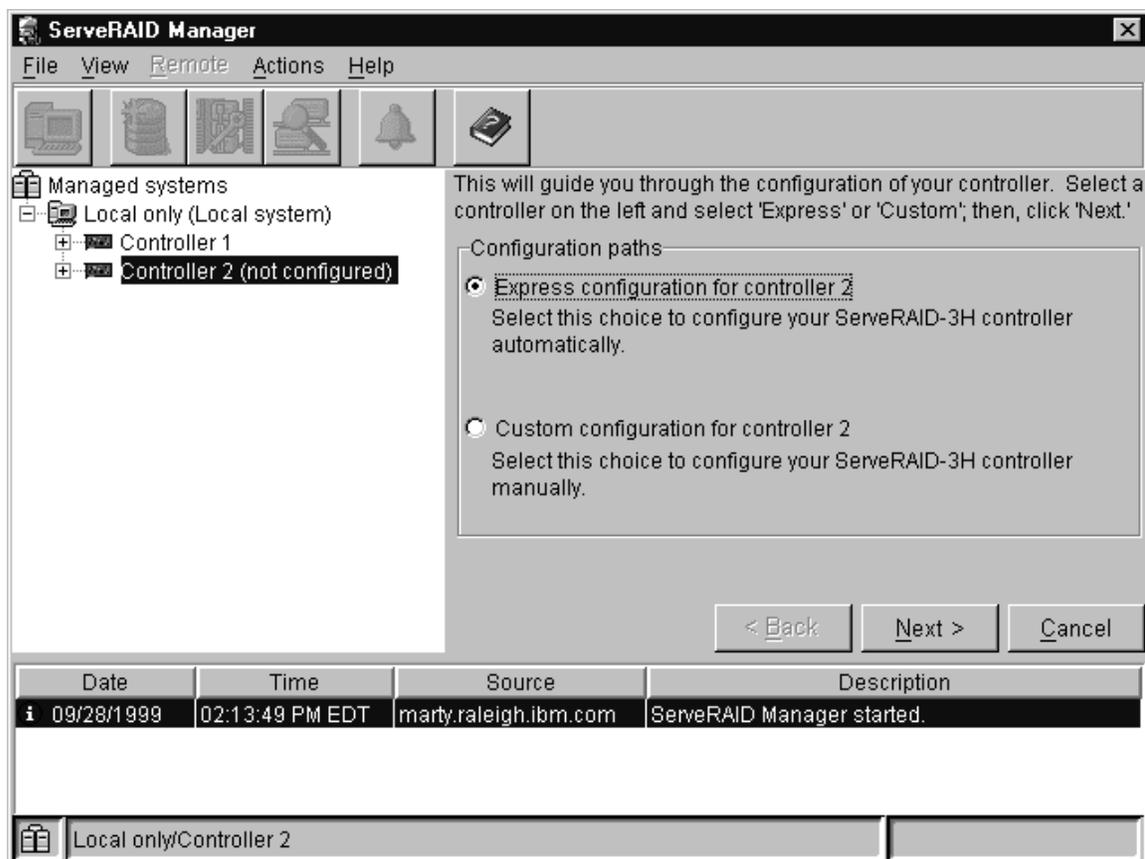


Figure 5. Configuration mode screen

3. Click the **Custom configuration** radio button.

4. Click **Next**. The Create arrays screen similar to the following appears.

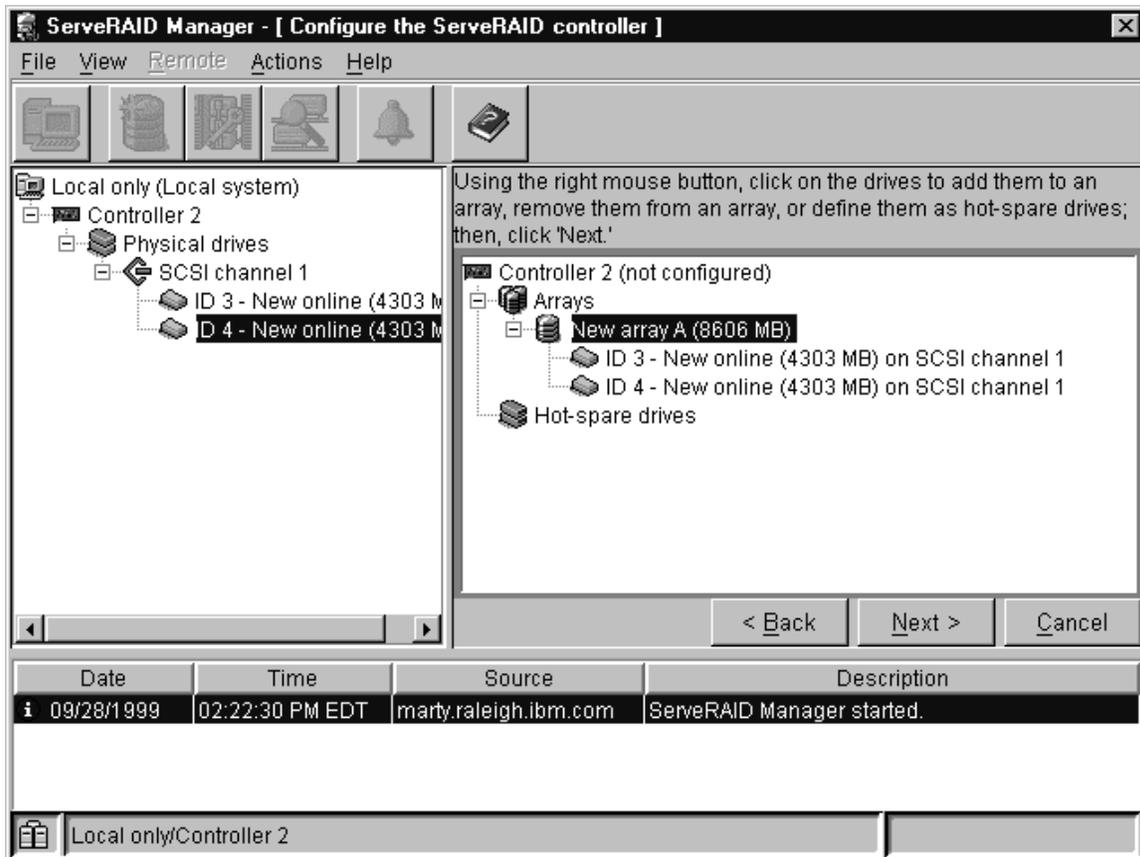


Figure 6. Create arrays screen

**Notes:**

- a. You must have two or more physical drives for RAID level-1 and three or more physical drives for RAID level-5 within the same array.
  - b. 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.
5. Right-mouse click the ready drive or SCSI channel icons in the Main Tree to select the drives that you want to add to your arrays, or define as hot-spare drives; then, select a choice from the pop-up list.

or

Drag the ready 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.

**Important**

Hot-spare drives are **not** shared between controllers in a cluster pairing. If you want hot-spare protection, **each** ServeRAID controller must have a hot-spare drive defined.

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

6. After you select the ready drives for your arrays and define your hot-spare drives, click **Next**. The Create logical drives screen appears.

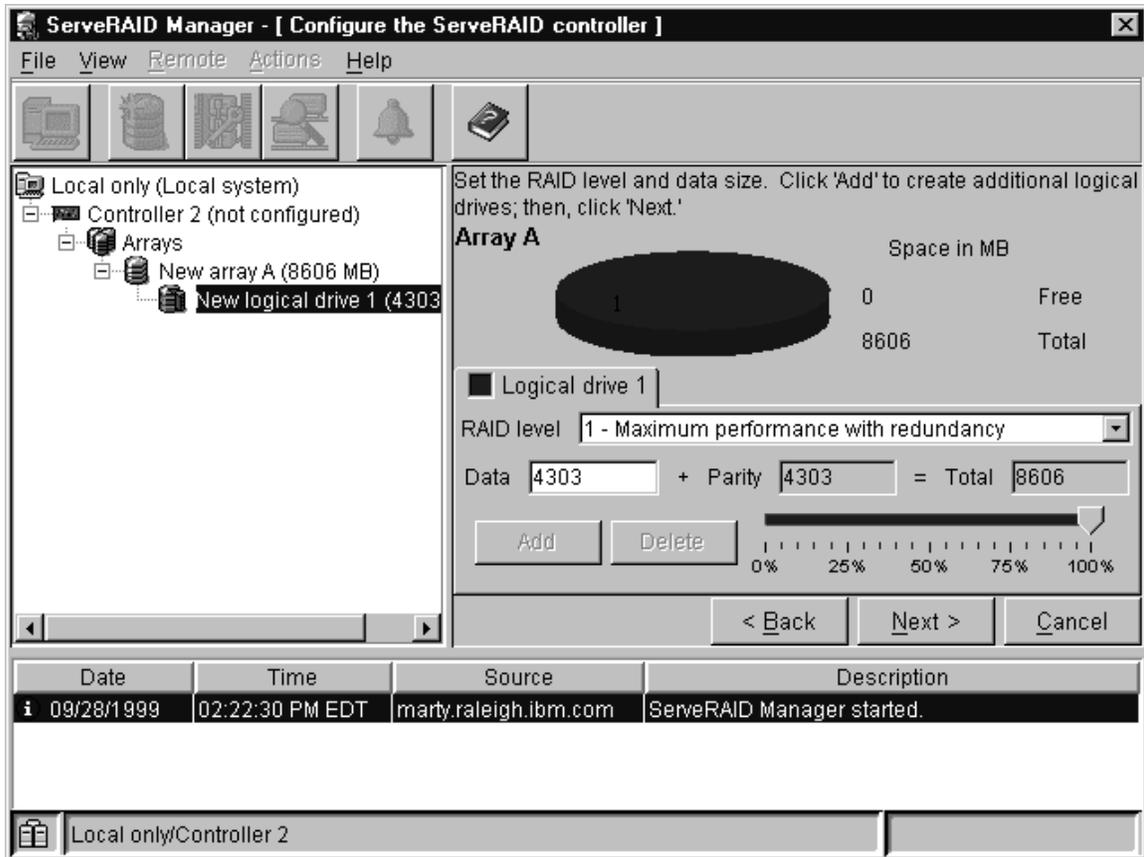


Figure 7. Create logical drives screen

If you change your mind, you can:

- Remove a specific ready drive from a newly defined array or delete an entire newly defined array. To do this, click **Back**; then, right-mouse click the specific drive or Array icon in the Main Panel on the right. 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 **Back**; then, right-mouse click the Hot-spare drive icon in the Main Panel on the right. Select **Delete new hot-spare drive** or **Delete all new hot-spare drives**.
7. Select a RAID level for the logical drive from the RAID drop-down list box in the Main Panel. (Refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide* for descriptions of the supported levels.)
  8. Click **Next**. The Configuration summary screen appears.

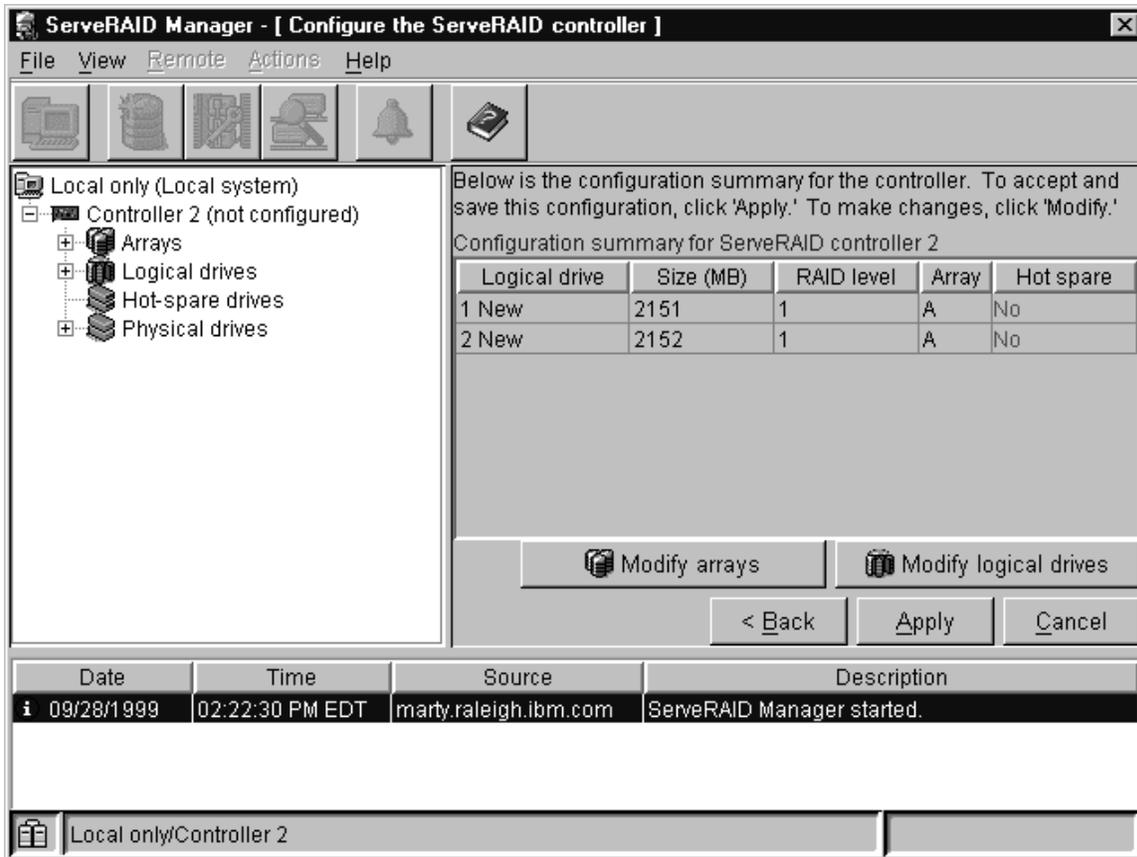


Figure 8. Configuration summary screen

9. Review the information that appears on the Configuration summary screen. Click **Apply** to accept and save the configuration.

**Notes:**

- a. To change the configuration, click either **Modify arrays** or **Modify logical drives**.
- b. For Microsoft Windows NT clusters, you should define all shared logical drives on Server A.

10. Click **Yes** to continue working with your ServeRAID configuration.
11. After creating the arrays and logical drives, you must assign a Merge Group number to the shared logical drives. Click the controller in the Main Tree that contains the logical drives.

**Important**

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

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

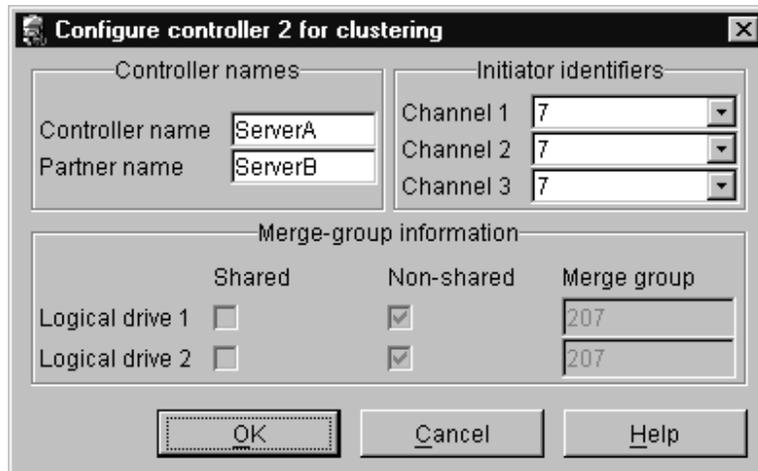


Figure 9. Configure for clustering window

13. Click the **Shared** check box for the logical drives that you want to share between the two controllers.
14. Set the Merge Group number.
 

**Note:** Although you can change the Merge Group number, **use** the default provided.
15. Record the Controller Name, Partner Name, and Merge Group numbers and store in a safe place.
 

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

  - a. Click the system from the Main Tree.
  - b. Select **Save printable configuration** from the Actions menu.
16. Click **OK**.
17. Select **Exit** from the File menu to exit the ServeRAID Manager program.

---

## Chapter 4. Installing the Windows NT 4.0 cluster software

After configuring the ServeRAID controller, defining logical drives, and setting up Merge Group numbers, you can install the Windows NT 4.0 cluster software.

---

### Installing the Microsoft Cluster Software

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

#### Important

Before beginning the procedure, consider the following:

- Read "Preparing to install or modify a High-Availability Cluster Solution" on page 3, which contains important information on ServeRAID considerations.
- Read "Configuring the controller for Windows NT 4.0" on page 8 for configuring Server A and Server B for the high-availability cluster solution.
- Ensure that you are using the *IBM ServeRAID Device Driver Diskette* (version 3.50, or higher) during the installation of the Microsoft Windows NT Server Enterprise Edition.

1. If you do not have Microsoft Windows NT Server Enterprise Edition, install it with Service Pack 3 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 user's guide and the documentation that is included 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 ServeRAID Manager program to scan for new drives.
- Check to see if all synchronization of RAID level-5 logical drives is completed. To check if any synchronization is running, use the ServeRAID Manager program. If logical drives are currently being synchronized, a progress indicator appears in the lower right corner showing the progress of the synchronization.

- a. On Server A, run the utility IPSHAHTO.EXE from the ServeRAID Windows NT Cluster Solutions Diskette or the IBM ServeRAID Support CD: If you are installing from diskette, type the following command from a Windows NT command line:

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

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

```
d:\programs\winnt\cluster\support\ipshahto
```

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

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

- 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.

**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.

- 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.EXE 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* into the CD-ROM drive on Server A and issue the following command from the command line:

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

**Note:** The parameter, *localquorum*, must be typed in lowercase.

- b. Install the Microsoft Cluster Server software following the instructions in the documentation that is included 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* 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 is included 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 is included with the MSCS software. Refer to the documentation that is included with the MSCS software for instructions.
- b. Select the **Cluster Group**.
- c. Select **Properties** from the **File** menu to display the properties window.
- d. Click **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 **OK**. Server A appears in the Preferred Owner text box.
  - g. Make sure that the cluster group has failover enabled. For instructions on how to do this, refer to the documentation that is included 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 Windows NT Cluster Solution utility by doing the following:
- a. From Server A, run the SETUP.EXE program from the *IBM ServeRAID Windows NT Cluster Solution Diskette* or IBM ServeRAID Support CD and follow the instructions that appear on the screen.
  - b. When complete, run the SETUP.EXE program on 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 *localquorum* to a ServeRAID disk resource by doing the following:
- a. From the Cluster Administrator, click the cluster name to select the cluster you are using.
  - b. Right-mouse click the cluster name to open the cluster.
  - c. Select **Properties**; then, select the **Quorum tab**.
  - d. In the **Quorum Resource** field, select a ServeRAID shared disk.
  - e. Click **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 localquorum resource and that the quorum resource is moved to a RAID level-1 logical drive.
8. Install Service Pack 4 or higher.

**Note:** If you installed Service Pack 4 or higher previously, you must re-install it now.

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

---

## 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 on Server A, using the ServeRAID Manager program. For more information about creating new arrays and logical drives, refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide*.

**Note:** Be sure to assign a unique shared Merge Group number to the new shared logical drive.

Before creating the ServeRAID logical cluster resource, assign a sticky drive letter to the shared logical disk drive on both servers using the Windows NT Disk Administrator Program by following the next procedure. (Use the IPSHAHTO.EXE utility to move the logical drive to the other server so an identical drive letter can be assigned.)

1. On Server A, use the Windows NT Disk Administrator Program to assign a previously unused drive letter to the new logical disk. For information on how to assign sticky drive letters, see the MSCS documentation.

2. On Server B, use the ServeRAID Manager program to scan for new or removed ready drives.
3. Run the IPSHAHTO.EXE utility on Server B to move the newly defined logical disk drive to Server B.
4. Using the Windows NT Disk Administrator Program on Server B, assign the same sticky drive letter.

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. 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.
5. Create a ServeRAID disk resource as part of the new group.
6. 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. Doing so ensures 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.

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

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 is included with the MSCS package for instructions.

3. When 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.

Create a new ServeRAID logical disk using the normal MSCS procedure for creating cluster resources. Refer to the documentation that is included 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 “Updating ServeRAID software in an MSCS cluster environment.”

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

1. Select one of the cluster servers from which to uninstall MSCS.
2. Click **Start** in the lower-left corner of the screen; then, select **Control Panel** from the **Settings** menu. The Windows NT control panel window appears.
3. Double-click **Services**. The Services window appears.
4. Select **Cluster Service** from the list of services; then, click **Stop**.
5. Click **Close** to exit the Services window. For more information, refer to the documentation that is included with the cluster services.
6. In the Windows NT control panel, double-click **Add/Remove Programs**.
7. Select **Microsoft Cluster Service** from the list of installed components; then, click **Add/Remove**.  
**Note:** Ignore the warnings regarding the removal of IBM ServeRAID Logical Disks.
8. Follow the on-screen prompts by clicking **OK** or **Yes**.
9. Restart the system when prompted to complete uninstallation on MSCS.
10. Click **Start**; then, select **Control Panel** from the **Settings** menu.
11. Double-click **Add/Remove Programs**.
12. Select **IBM ServeRAID Windows NT Cluster Solution** from the list of installed components; then, click **Add/Remove**.

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

13. Repeat step 1 to step 12 for the other server in the cluster.

---

## Updating ServeRAID software in an MSCS cluster environment

**Note:** If you are upgrading your RAID software after downloading it from the IBM Support Web site, perform this procedure. If you are installing your software using the *IBM ServeRAID Support CD*, refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide* for instructions.

This upgrade procedure allows the cluster to remain online at all times. 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 Manager program

To upgrade the ServeRAID Manager program, run the SETUP.EXE program on the *IBM ServeRAID Support CD*.

## Upgrading the IBM ServeRAID Microsoft Windows NT Cluster Solution

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

1. Click **Start** in the lower-left corner of the screen; then, select **Control Panel** from the **Settings** menu. The Windows NT control panel window appears.
2. Double-click **Services**. The Services window appears.
3. Select **Cluster Service** from the list of services; then, click **Stop**.
4. Click **Close** to exit the Services window. For more information, refer to the documentation that is included with the cluster services.
5. Insert the *IBM ServeRAID Support CD* into the CD-ROM drive and type:  
d:\programs\winnt\cluster\setup.exe  
where *d* is the CD-ROM drive.
6. Follow the instructions that appear on the screen.
7. Repeat this procedure for the other server in the cluster.

## 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. Click **Start** in the lower-left corner of the screen; then, select **Control Panel** from the **Settings** menu. The Windows NT control panel window appears.
3. Double-click **SCSI-Adapters**.
4. Click the **Drivers** tab.
5. Click **Add**.
6. Click **Have Disk**.
7. Use **A:\NT** for the path; then, click **OK**.
8. Click **OK** when the IBM ServeRAID controller is highlighted.
9. Click **New**.
10. Use **A:\NT** for the path to install the driver.
11. Insert the *IBM ServeRAID Support CD*, version 3.50 or higher.
12. Shut down and restart the server from the CD.

## Upgrading the IBM ServeRAID BIOS/firmware

To upgrading the IBM ServeRAID BIOS/firmware, do the following:

1. Start the server from the *IBM ServeRAID Support CD*, version 3.50 or higher, if the server was not previously started in step 11 of the procedure “Upgrading the IBM ServeRAID Windows NT miniport driver.”
2. Flash the BIOS and download the firmware by clicking **Update** when prompted. If no update is required, the ServeRAID Manager program starts.
3. When the download is complete, remove the CD from the CD-ROM drive and press **Ctrl+Alt+Del** to restart (reboot) the server.
4. Allow the server to rejoin the cluster and when the server is reported in the MSCS cluster administrator as "Up" and all resources are online, the upgrade for this server is complete.

## Upgrading the software for the other server in the cluster

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

1. Ensure that Server A is online in the cluster.
2. Repeat all the steps between “Upgrading the IBM ServeRAID Windows NT miniport driver” on page 23 through “Upgrading the IBM ServeRAID BIOS/firmware” on the other server (Server B) in the cluster.

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

---

## Chapter 5. Monitoring and updating a High-Availability Cluster Solution using the ServeRAID Manager program

You can use the ServeRAID Manager program to check the ownership of a physical drive on another server or to add or remove a logical drive in the cluster. There are three features for high-availability functions supported by the ServeRAID Manager program:

- Validate cluster
- Configure for clustering
- View shared drives

**Note:** Make sure the ServeRAID Manager program is installed on your server. If you need to install it, refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide* for instructions. See "Downloadable files from the World Wide Web" on page v for additional information.

---

### Checking the ownership of a physical drive

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

1. Click the controller that contains the ready drives from the Main Tree. A screen similar to the following appears:

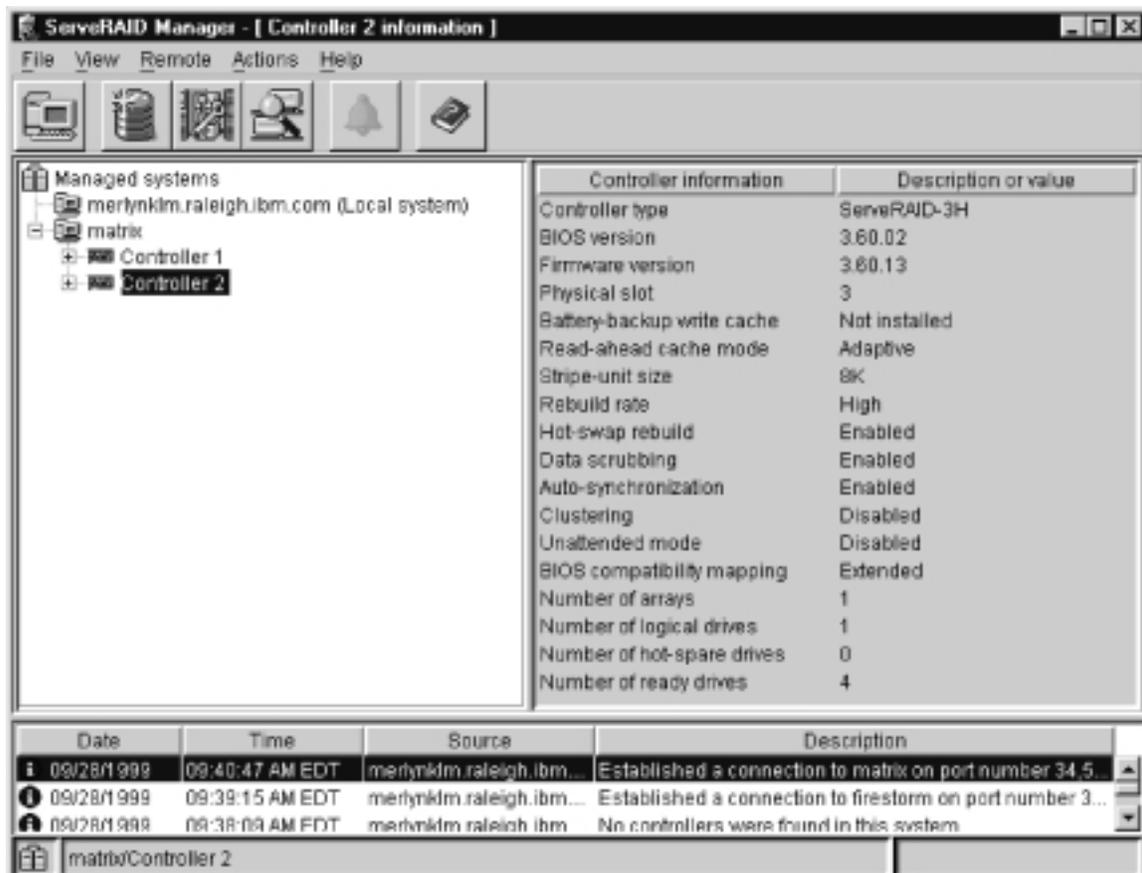


Figure 10. Information mode screen

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

**Important**

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

---

## Adding logical drives to a cluster

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

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

**Attention**

Physical drives that are configured by one controller in the pair must appear as Ready drives to the other controller in the pair. Be sure to verify this consideration when you add drives *after* the servers are started. Restart both servers, or scan for new drives using the ServeRAID Manager program on both servers, to make sure the added physical drives appear as ready drives on both servers before creating logical drives. For more information, see "ServeRAID considerations" on page 4.

2. Using the ServeRAID Manager program, configure the hard disk drives for the ServeRAID controller. Refer to the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide* for instructions.

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

3. Click the controller for the cluster from the Main Tree.

4. Click  on the tool bar. A screen similar to the following appears:

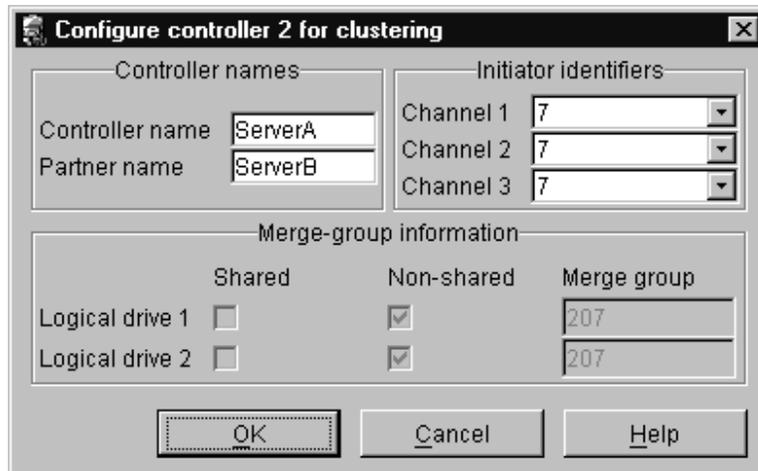


Figure 11. Configure for clustering window

5. Click **Shared** or **Non-shared** for the logical drive you have just defined.
6. Select the Merge Group Number you want to use for the logical drive you just defined.
7. Click **OK**.
8. Continue with “Creating a ServeRAID disk resource on a ServeRAID server with MSCS” on page 20.

## Validating a cluster

Use this feature to determine if your cluster configuration is correctly configured. To validate a cluster, do the following:

1. Click a system in the cluster from the Main Tree.
2. Select **Validate cluster** from the Actions menu. A screen similar to the following appears:

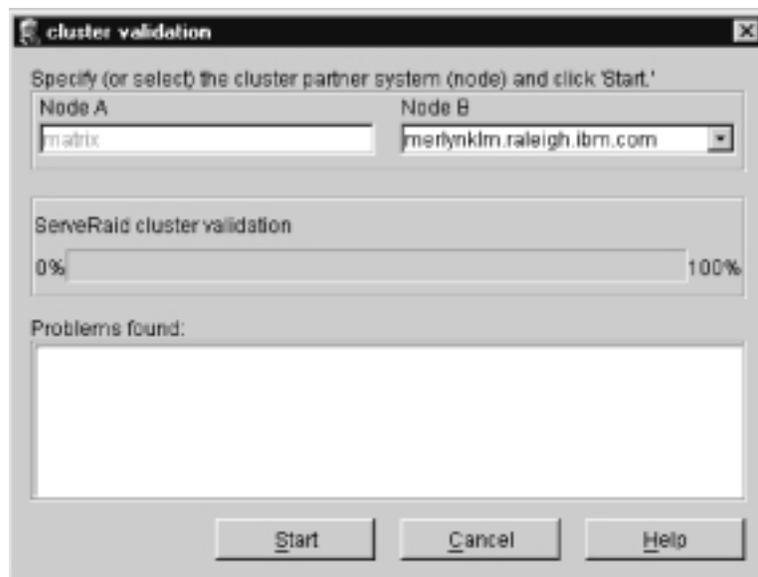


Figure 12. Validate cluster window

The system you selected in step 1 is listed in the Server A box.

**Note:** You cannot change the Server A system in the Validate cluster window. To select a different system, click **Cancel** to return to the Main Tree.

3. Select your second system from the Server B drop-down list box.

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

4. Click **Start**. The ServeRAID Manager program scans both systems to verify the ServeRAID cluster configuration. The progress indicator reports the progress of the validation.

5. If you want to stop the validation before it completes, click **Cancel**. Otherwise, if the ServeRAID Manager program found problems with the ServeRAID cluster configuration, it reports the problems in the Problems found message field. If there are no problems with the ServeRAID cluster configuration, the message field reports only the cluster pairings found.

6. To view the problem report events in an easy-to-read window, double-click an event in the Problems found message field.

7. To save the contents of the Problems found message field, click **Save**. A Save as window appears. Specify the file name for the message field log and click **OK**.

8. Click **Done** to exit the Validate cluster window. If the ServeRAID cluster configuration has problems, fix the problems; then, use **Validate cluster** again.

---

## Viewing Merge Group numbers and other cluster information

To view the Merge Group numbers and other cluster information, do the following:

1. Click the controller you are using for the cluster from the Main Tree.

2. Click  on the tool bar. A screen similar to the following appears:

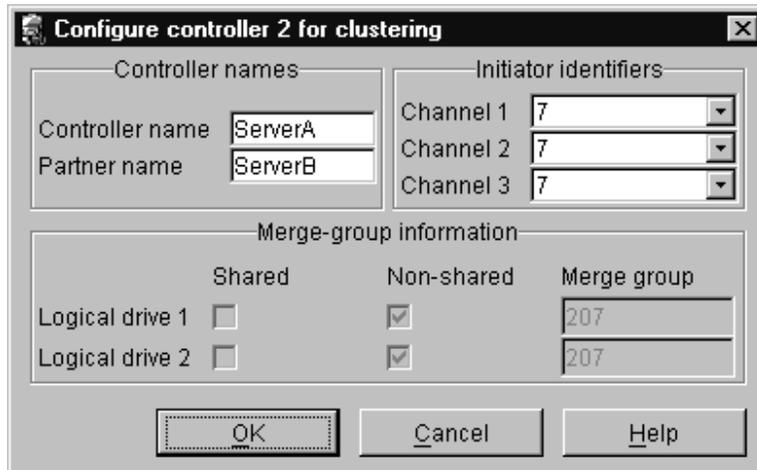


Figure 13. Configure controller for clustering window

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

---

## Chapter 6. Troubleshooting a High-Availability Cluster

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

---

### Rebuilding a defunct drive

A hard disk drive goes defunct when there is a loss of communication between the controller and the hard disk drive. This can be caused by any of the following:

- An improperly connected cable, hard disk drive, or controller
- A loss of power to a drive
- A defective cable, backplane, hard disk drive or controller

In each case, the communication problem needs to be resolved, and then a Rebuild operation is required to reconstruct the data for the device in its disk array. The ServeRAID controllers can reconstruct RAID level-1 and RAID level-5 logical drives, but they cannot reconstruct data stored in RAID level-0 logical drives.

To prevent data-integrity problems, the ServeRAID controllers set the RAID level-0 logical drives to Blocked during a Rebuild operation. After the Rebuild operation completes, you can unblock the RAID level-0 logical drives, and access them once again. Remember, however, that the logical drive might contain damaged data.

### Steps for recovering from defunct drives

If the defunct drives are not part of an array, contact your IBM service representative.

If the defunct drives are part of an array, do the following:

- If more than one hard disk drive in an array is defunct, contact your IBM service representative.
- If a Rebuild operation is in progress, wait until the Rebuild is complete.
- If a Rebuild is not in progress and only one hard disk drive in the array is defunct, do the following:
  1. Verify the cables, hard disk drives, and controllers are installed properly.
  2. Attempt to rebuild the defunct hard disk drive by performing a hot-swap rebuild. Refer to “Rebuilding a hot-swap drive” on page 30 for instructions.
  3. If the hot-swap rebuild fails, contact your IBM service representative.

## Rebuilding a hot-swap drive

A hot-swap rebuild refers to a rebuild operation that is started by the ServeRAID controller when it detects that a drive that is part of an array and in the defunct state, has been removed and reinserted on the SCSI cable or backplane. The reinsertion of the hard disk drive, whether it is the same drive or a new drive, will trigger the ServeRAID controller to start the Rebuild operation. During the Rebuild operation, the drive being rebuilt is in the rebuild state, and the logical drive remains critical until the Rebuild operation has been successfully completed.

On IBM servers, when a hot-spare drive is available, the Rebuild operation will begin automatically without the requirement to replace the failed drive.

To start a hot-swap rebuild, do the following:

1. Without removing the drive completely, gently remove the hard disk drive from the server, using the handle of the hot-swap tray. If necessary, refer to the documentation that comes with your server on removing a hard disk drive.
2. Wait 20 seconds to allow the disk drive to completely spin down.

**Note:** When power is removed from a hot-swap drive, the drive immediately parks the heads, locks the actuator in the “landing zone,” and begins spinning down. However, the spinning down of the disk might require up to 20 seconds after power is removed. Do not move the drive while it is spinning down. Moving the drive while it is spinning down may damage the drive.

3. Gently reinstall into the server the hard disk drive that you removed. Make sure the drive is completely installed in the backplane connector.

---

## Recovering from a failed server

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

### Important

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

Use the following procedure to recover from a failed server:

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

As you remove your hard disk drives, be sure to note the bay in which each drive was installed. If you are replacing your failed server with an identical server you can reinstall the drives in an identical configuration and get your server up and running quickly.

3. If the ServeRAID controller is functional, remove it from the failed server. As you remove the controller from the failed server, be sure to note the following:
  - Which SCSI cables are connected to the SCSI channel connectors on the controller.

- Which PCI slot has the controller installed.
4. Install the hard disk drives in the new server.

For information on how to install a hard disk drive, see the documentation that comes with your server. If you are replacing the failed server with an identical server, install each hard disk drive in the same bay as the one it was removed from in the failed server.

5. If the ServeRAID controller in the failed server is *not functional*, or you are not sure, continue with “Recovering from a failed ServeRAID controller.” Otherwise, perform the following steps with the functional ServeRAID controller:
  - a. Install the controller in the new server.
  - b. Reinstall the cables as they were in the failed server.
  - c. Start the server.

---

## Recovering from a failed ServeRAID controller

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

- SCSI channel initiator IDs
- Controller Name
- Partner Name
- Stripe-unit size
- Unattended mode

If the ServeRAID controller you are replacing *is* functional, you can obtain the above configuration information by starting your server with the *IBM ServeRAID Support CD* in the CD-ROM drive and using the ServeRAID Manager program. Refer to *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide* for additional information on using the ServeRAID Manager program.

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

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

**Obtaining the current SCSI channel initiator IDs:** The SCSI channel initiator IDs for a non-shared SCSI channel are set to 7 by default. For shared SCSI channels, the initiator IDs must be 7 or 6 and must be different from the SCSI channel initiator IDs for the corresponding SCSI channels of the cluster partner controller. Therefore, if you can obtain the SCSI channel initiator IDs for the corresponding cluster partner system, you know the correct SCSI channel initiator IDs for this system. For example, if the cluster partner system-shared SCSI bus initiator IDs were set to 7, the controller you are replacing would need to have its IDs set to 6.

To obtain the SCSI channel initiator IDs from the corresponding cluster partner controller, do the following:

1. Start the cluster partner system with the *IBM ServeRAID Support CD* in the CD-ROM drive. The ServeRAID Manager program starts.
2. Click the cluster partner controller from the Main Tree.
3. Click  on the tool bar. The Configure for clustering window appears.

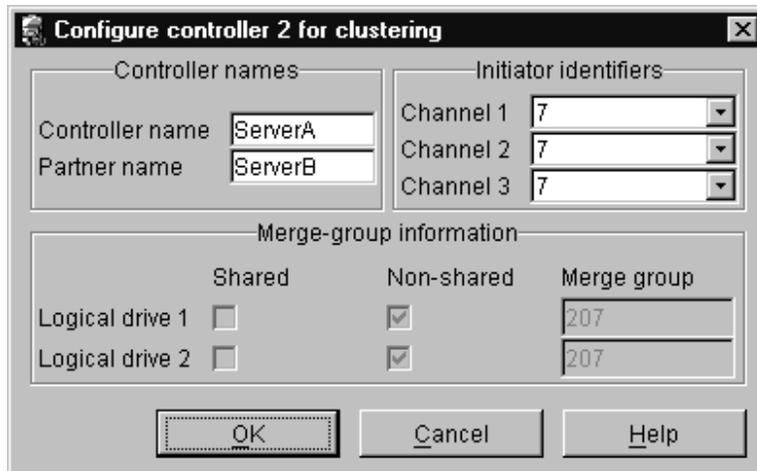


Figure 14. Configure for clustering window

- Note the settings in the fields.

**Obtaining the current Controller Name and Partner Name:** You can also determine the proper settings for the Controller Name and Partner Name of the controller being replaced by reading the settings from the cluster partner system.

- Start the cluster partner system with the *IBM ServeRAID Support CD* in the CD-ROM drive. The ServeRAID Manager program starts.
- Click the cluster partner controller from the Main Tree.

- Click  on the tool bar. The Configure for clustering window appears.
- Note the settings in the fields.

**Obtaining the current stripe-unit size:** The stripe-unit size **must** be the same on both controllers in a high-availability cluster. To obtain the current stripe-unit size, do the following:

- Start the cluster partner system with the *IBM ServeRAID Support CD* in the CD-ROM drive. The ServeRAID Manager program starts.
- Click the cluster partner controller from the Main Tree.
- In the Main Panel (that is, the right screen), look for the stripe-unit size setting.

**Status of unattended mode:** This setting is configured automatically as Enabled when a controller is set up for clustering.

## Replacing a ServeRAID controller

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

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

For instructions on how to remove and install the ServeRAID controller, see the *IBM Netfinity ServeRAID-3H, -3HB, and -3L Ultra2 SCSI Controllers Installation and User's Guide*.

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

**Important**

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

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

6. Start the system with the *IBM ServeRAID Support CD* (version 3.50, or higher) in the CD-ROM drive. The ServeRAID Manager program starts.
7. Click the new controller in the Main Tree.
8. Select **Restore to factory default settings** on the Actions menu. The ServeRAID Manager program initializes the controller configuration.

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

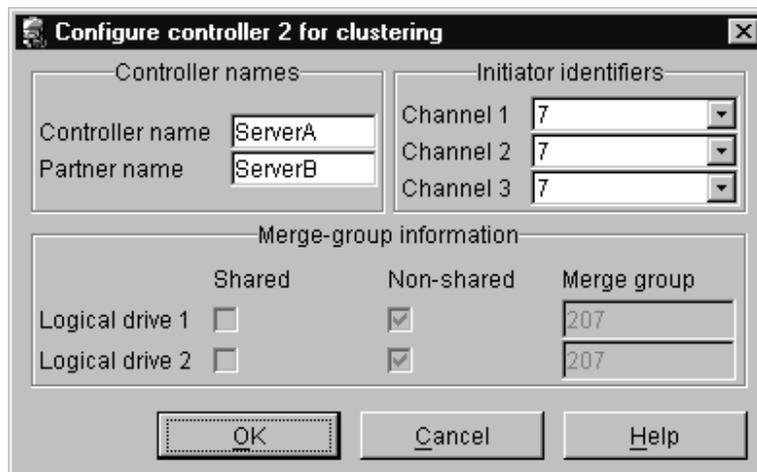


Figure 15. Configure for clustering window

10. Using the settings that were assigned to the previous ServeRAID controller (that is, the controller you are replacing), fill in each of the following fields:
  - SCSI channel initiator IDs
  - Controller Name
  - Partner Name
11. Click **OK**. The new controller stripe-unit size setting defaults to 8 KB. If you need to change this setting to match the partner controller, go to step 12. Otherwise, go to step 15.
12. Click the new controller in the Main Tree.
13. Select **Change stripe-unit size** on the Actions menu.
14. Select the new stripe-unit size for your installation on the menu.
15. Shut down and turn off the server.
16. Reconnect the SCSI channel cables to the controller.

**Note:** Be sure to connect the cables to the correct SCSI channels as noted in step 2 on page 32.

17. If the ServeRAID controller being replaced attaches to any of the following:

- The server startup disk array
- Other non-shared disk arrays
- Shared arrays which have not failed-over to the other server,

start the system with the *IBM ServeRAID Support CD* (version 3.50, or higher) in the CD-ROM drive. The ServeRAID Manager program starts.

**Otherwise**, go to step 21.

18. Click the new controller in the Main Tree.

19. Select **Copy configuration from drives** on the Actions menu.

20. Select the **Merge Group number** for the non-shared logical drives (that is, 206 or 207) on the menu.

**Note:** The last number in the Merge Group number is the shared SCSI channel initiator ID. That is, if the Merge Group number is 206, the SCSI channel initiator ID is 6.

21. Remove the *IBM ServeRAID Support CD* from the CD-ROM drive.

22. Restart your server.

**Note:** If the controller you replaced is the startup controller, the system will now start the operating system properly.

---

## 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.

---

### The ServeRAID Manager program shows physical drives in defunct state.

**Action:**

1. Enable the **View shared drives** feature in the ServeRAID Manager program. Doing so displays a drive as reserved, instead of defunct, in the following situations:

- A drive has been moved.
- A drive has failed-over to the other server in the cluster.

**Note:** It is normal for drives in the above situations to display as defunct if you have not enabled the **View shared drives** feature. In this case, the drives shown in the defunct state are not really defective.

2. Check RAID level-1 and RAID level-5 arrays to make sure they are not in critical state. If they are in critical state, replace the failed drive and perform a rebuild operation.

---

### ServeRAID shared logical drives do not failover properly.

**Action:**

1. Ensure that the resource type of each ServeRAID controller 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 the high-availability cluster solution using Microsoft Windows NT. Refer to Chapter 3, "Configuring ServeRAID controllers for the IBM High-Availability Cluster Solution" on page 8 for instructions.

2. Ensure that shared SCSI buses on the ServeRAID controller pair are connected in a way that corresponding SCSI channels are connected. For example, SCSI channel 1 on the controller in the first cluster server is connected to SCSI channel 1 on the controller in the second cluster server, channel 2 is connected to channel 2, and so forth.
3. Ensure that physical SCSI disks that contain logical drives are all connected to shared SCSI channels.
4. Ensure that there are no more than eight shared logical disk drives defined per pair of ServeRAID controllers for use in your cluster.
5. If you are using a ServeRAID-II with a quorum drive, then with Windows NT Server clusters, ensure that SCSI channel 3 of the pair that attaches to the ServeRAID logical drive that has been designated as the Windows NT Cluster Quorum Resource is used for arbitration. Also, ensure that it is connected from the first cluster server to SCSI channel 3 in the second cluster server and that there are no SCSI devices connected to that channel.  
  
If you are using a ServeRAID-II with a quorum drive, then the SCSI heartbeat connection *must* be connected to the third channel of the ServeRAID controller 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 controller, ID level 5, you must also move the SCSI heartbeat cable on both servers to the new quorum ServeRAID controller pair. For more information, see "ServeRAID considerations" on page 4.
6. Make sure each shared logical drive has a Merge Group number assigned. Shared Merge Group numbers **must** be in the range 1–8.
7. Make sure each ServeRAID controller has been assigned a unique Controller Name and that each ServeRAID controller has its Partner Name assigned properly to correspond to the ServeRAID controller in the other cluster server that is attached to the shared SCSI buses.
8. Check for loose shared SCSI bus cables.
9. Ensure that SCSI repeater cards in Model 3518 or 3519 disk expansion enclosures are at the latest revision level.
10. Ensure that physical drives that are expected to be moved or failover show up as ready or reserved state on the server that is attempting to take over control of these disks.

---

**RAID level-5 logical drives cannot be accessed by the operating system after a failover.**

**Action:** Use the ServeRAID Manager program to check the state of the logical drive to ensure that it is not blocked. Using this program, select the logical drive and look for Blocked state **Yes**. If the logical drive is blocked, make sure all physical drives that are part of the logical drive are in the online state. If all physical drives are *not* in the online state, a drive might have gone bad during one of the following situations:

- A failover.
- A resynchronization process after a failover.

Data integrity cannot be guaranteed in this case and the logical drive has been blocked to prevent the possibility of incorrect data being read from the logical drive.

Unblock, reinitialize, and synchronize the logical drive and restore the data from a backup source. Depending on the type of data contained on the logical drive and the availability of a recent backup copy, you can do one of the following:

- Unblock the drive and continue normal operation.
- Replace and rebuild one or more defunct drives.

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.

---

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

**For example, the Windows 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. Right-mouse click 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.

---

**After one of the cluster servers has been shut down normally and the surviving server takes over the cluster resources, occasionally one or more of the IBM ServeRAID logical drive resources will stay in the online pending state for several minutes, after moving over to the surviving server (when viewed with the Cluster Administrator). After this, the resource will go to the failed state and the following error message will be displayed in the surviving server system log (as viewed with the Event Viewer).**

**For example, the Windows 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 reattempt to bring this resource online on the surviving server within about four minutes.

---

**You cannot reinstall the ServeRAID Windows NT Cluster Solution. If a previous version of IBM ServeRAID Cluster Solution has been uninstalled, when attempting to reinstall the IBM ServeRAID Windows NT Cluster solution, a message incorrectly appears prompting you to perform an upgrade.**

**Action:** You must delete the *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 3, "Configuring ServeRAID controllers for the IBM High-Availability Cluster Solution" on page 8 for instructions.

---

**When running the IPSHAHTO program on a server, the following error message appears:** *Warning: CONFIG\_SYNC with 0xA0 command FAILED on Controller # and either of the following occurs:*

- **One or more hot-spare or standby hot-spare drives are defined on controller pairs.**
- **Ready drives are not part of any logical drive configuration on a controller pair in a cluster.**

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

However, if shared drive resources *fail* to move when running the IPSHAHTO program, remove all ready drives that are not part of any logical drive configuration on a controller pair in a cluster.

---

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

**Action:** By design, the array identifiers and logical drive numbers might change during a failover condition. Consistency between the Merge Group numbers 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 server.

---

## 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 or both:

IBM

Netfinity

OS/2

ServeRAID

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

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



Part Number: 09N7385

Printed in U.S.A.

09N7385

