



IBM Netfinity
High Availability Clustering Solutions
Using the IBM SSA RAID Cluster Adapter

Installation and User's Guide



International Technical Support Organization

SG24-5387-00

**IBM Netfinity High Availability Clustering Solutions
Using the IBM SSA RAID Cluster Adapter
Installation and User's Guide**

April 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, "Special Notices" on page 77.

First Edition (April 1999)

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 jurisdictions 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 and the United Kingdom. It is possible that this publication may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

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

No part of this publication may be reproduced or distributed in any form or by any means without prior permission in writing from the International Business Machines Corporation.

© Copyright International Business Machines Corporation 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

| | |
|---|----|
| About This Book | v |
| How This Book Is Organized | v |
| Related Publications | v |
| | |
| Chapter 1. Introduction | 1 |
| 1.1 Microsoft Cluster Server | 1 |
| 1.2 SSA RAID Cluster Adapter | 2 |
| 1.2.1 Features | 2 |
| 1.2.2 Supported RAID Levels | 3 |
| 1.2.3 Array Configuration Issues | 4 |
| 1.2.4 Disk Drive Access Modes | 4 |
| 1.2.5 Booting of SSA RAID Cluster Adapter Drives | 4 |
| 1.2.6 Data Access | 5 |
| 1.2.7 Performance Conclusions | 5 |
| 1.2.8 Supported Operating Systems | 5 |
| | |
| Chapter 2. Preparations | 7 |
| 2.1 Terms and Definitions | 7 |
| 2.1.1 Node | 7 |
| 2.1.2 Resources | 7 |
| 2.1.3 Quorum Resource | 11 |
| 2.1.4 Resource Groups | 11 |
| 2.1.5 Failover | 13 |
| 2.1.6 Failback | 14 |
| 2.2 Networking | 15 |
| 2.2.1 Redundancy | 16 |
| 2.2.2 TCP/IP Considerations | 16 |
| 2.3 Domains | 19 |
| 2.4 Clustering Models | 19 |
| 2.5 SSA Considerations | 20 |
| 2.5.1 SSA RAID Cluster Adapter Considerations | 20 |
| 2.5.2 Connection Considerations | 21 |
| 2.5.3 RAID Considerations | 21 |
| 2.5.4 SSA 7133 Enclosure Considerations | 21 |
| 2.5.5 Fiber Optic Considerations | 22 |
| 2.5.6 Disk Types and Disk Numbers | 22 |
| 2.6 Planning for Microsoft Cluster Server | 23 |
| 2.6.1 Identifying Risks | 23 |
| 2.6.2 Cluster Configurations | 25 |
| | |
| Chapter 3. SSA Configurators | 27 |
| 3.1 Obtaining the Software | 27 |
| 3.2 Remote System Management | 27 |
| 3.2.1 Installing RSM | 28 |
| 3.2.2 Using RSM | 28 |
| 3.2.3 Tutorial | 29 |
| 3.3 SSA Text-Mode Configurator | 29 |
| | |
| Chapter 4. Installing Windows NT and SSA Tools | 31 |
| 4.1 Preparation | 31 |
| 4.2 Installing Windows NT Server Enterprise Edition | 31 |

| | | |
|---|--|-----------|
| 4.3 | Installing the Device Drivers | 33 |
| 4.4 | Updating the SSA Firmware | 36 |
| 4.5 | Installing the SSA Event Logger | 37 |
| 4.6 | Installing RSM | 38 |
| 4.6.1 | Installing the Stand Alone Version | 39 |
| 4.6.2 | Installing the Netfinity Manager Version | 41 |
| 4.7 | Next Steps | 42 |
| Chapter 5. Configuring SSA Hardware | | 43 |
| 5.1 | Connect the Enclosure | 43 |
| 5.2 | Formatting the SSA Disks | 44 |
| 5.3 | Defining the Arrays | 48 |
| 5.4 | Defining Hot Spares | 52 |
| 5.5 | Setting the Arrays as System Resources | 55 |
| 5.6 | Creating Partitions in Windows NT | 57 |
| 5.6.1 | Working with Node B | 59 |
| Chapter 6. Installing Microsoft Cluster Server | | 61 |
| 6.1 | The Installation Process | 61 |
| 6.1.1 | Installing MSCS on the First Node | 61 |
| 6.1.2 | Installing MSCS on the Second Node | 69 |
| 6.2 | Testing the Cluster Installation | 71 |
| 6.2.1 | Moving the Cluster Group | 72 |
| 6.2.2 | Shutting Down Node B | 74 |
| 6.3 | Cluster Administrator on a Remote System | 74 |
| Appendix A. Special Notices | | 77 |
| Index | | 79 |

About This Book

This book provides information and instructions for setting up a high availability clustering solution using the IBM SSA RAID Cluster Adapter.

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" provides an overview of high availability cluster solutions.

Chapter 2, "Preparations" 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 SSA RAID Cluster Adapter.

Chapter 3, "SSA Configurators" explains the configuration tools that are available to set up the SSA disk subsystem.

Chapter 4, "Installing Windows NT and SSA Tools" describes how to install Windows NT Server Enterprise Edition, the SSA drivers, SSA adapter firmware, SSA tools and the Remote Systems Management configurator.

Chapter 5, "Configuring SSA Hardware" describes how to configure the SSA disk subsystem in preparation for clustering.

Chapter 6, "Installing Microsoft Cluster Server" describes how to install Microsoft Cluster Server on both nodes.

Related Publications

The following publications may be useful during the preparation and installation of the high-availability cluster solution:

- *IBM SSA RAID Cluster Adapter for PC Servers Installation and User's Guide*, S96H-9839
- *IBM SSA RAID Cluster Adapter for PC Servers Technical Reference*, SA33-3275
- *IBM SSA RAID Cluster Adapter Hardware Maintenance Manual Supplement*, S96H-9840
- *Implementing Netfinity Disk Subsystem: ServeRAID SCSI, Fibre Channel and SSA*, SG24-2098
- *Microsoft Cluster Server Administrator's Guide*, supplied with Windows NT Server Enterprise Edition
- *TCP/IP Tutorial and Technical Overview*, GG24-3376

Chapter 1. Introduction

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.

A *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 the IBM SSA RAID Cluster Adapter and IBM expansion enclosures, such as the IBM 3527 SSA Entry Storage Subsystem and the IBM 7133 Serial Disk System. These IBM high availability cluster solutions can be installed using Microsoft Cluster Server, which is part of Windows NT Server Enterprise Edition.

For the latest information and downloads for Netfinity clustering, go to the following Web page:

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

1.1 Microsoft Cluster Server

Microsoft Cluster Server (MSCS) is Microsoft's first major push into the enterprise computing arena. The first release of Microsoft Cluster Server links two servers together to allow *system* redundancy. Before Microsoft Cluster Server, Vinca's StandbyServer and the like, as well as hardware manufacturers such as IBM, provided redundancy in many of the PC server's *components*, such as power supplies, disks and memory.

System redundancy means that a complete server can fail and client access to server resources will remain largely unaffected. Microsoft Cluster Server extends this theme by also allowing for software failures at an application level as well as an operating system level. If the operating system fails, all applications and services can be restarted on another server. If just one application fails, it can be managed by Microsoft Cluster Server individually. This, in effect, means that a failure can occur but the cluster as a whole remains intact still servicing its customer's requests.

Microsoft Cluster Server achieves this by continually monitoring services and applications. Any program that crashes or hangs can be immediately restarted on the same server or an alternate server.

If a failure does occur, the process of restarting the application on another server is called *failover*. Failover can occur either automatically, such as when an

application or a whole server crashes, or manually. By issuing a manual failover, the administrator is able to move all applications onto one server and bring the other server down for maintenance. Once the server is back online again, applications can, as required, be transferred back to it either manually, or automatically.

Note: Microsoft Cluster Server supports hardware-based RAID solutions such as those offered by IBM. Microsoft Cluster Server does not support software RAID solutions in the shared disk enclosures. Windows NT's software RAID can be used to provide redundancy for local disks, however.

1.2 SSA RAID Cluster Adapter

The SSA RAID Cluster Adapter (option 96H9835) supports 2-way disk configurations. This means that two adapters in one loop are able to share these disk resources. In conjunction with Microsoft Cluster Server, the adapters support additional data protection through host failover.

As shown in Figure 1, the adapter has 1 MB SRAM. It also has 32 MB of DRAM in the form of two DIMMs connected at an angle to the board. DRAM is mostly used for caching of RAID-1 writes, also known as *staging*.




Figure 1. SSA RAID Cluster Adapter

1.2.1 Features

The RAID Cluster Adapter has the following general features:

- Full-length PCI busmaster adapter.
- Supports up to 48 SSA drives on one loop.
- It has two pairs of connectors (or ports) to accommodate two loops for a total of up to 96 devices per adapter. You have to use the corresponding connectors of one pair to connect to a loop, which means A1 together with A2 in a loop and B1 together with B2 for another loop. See Figure 2 on page 3 for an illustration.
- Next to each pair of connectors is an LED as shown in Figure 2 on page 3.
 - The light is on continuously when a closed loop is attached and the ports are functional.
 - The light flashes continuously if one of the ports is not operational which usually means a broken loop.

- The light is off if both ports are not operational.




Figure 2. Serial Storage Architecture Adapter. The adapter has four external connectors. Connectors A1 and A2 are always paired off together and connectors B1 and B2 are always paired off together.

- Each SSA loop has a maximum bandwidth of 40 MBps.
- Booting of SSA drives is supported.

Booting from SSA Drives

If you have problems booting from an SSA disk resource you should check the following first:

- Is the CBIOS of the adapter enabled? (This is by default "disabled".)
- Is the adapter installed in a PCI slot with higher priority than other disk controllers?
- Is the SSA device driver loaded before other (for example, SCSI) disk drivers?
- Is the SSA adapter running the latest firmware?
- Is the resource (array or JBOD) you are trying to boot on top of the system resources-list in the DOS configuration utility? (This is required.)

- Over standard copper cables you can reach distances of up to 25 meters between SSA nodes. Depending on the type of enclosure being used, these distances can be increased even farther, up to 2.4 km using multi-mode fiber optic extenders or 10 km using single-mode fiber optic extenders.

1.2.2 Supported RAID Levels

The SSA RAID Cluster Adapter supports the following disk configurations:

- 2-way JBOD (single disk)
- 2-way RAID-1 arrays

Note: RAID-1 arrays have exactly two members (this is not RAID-1 enhanced as in the ServeRAID II adapter). This means that using the SSA RAID Cluster Adapter, the maximum partition size is limited to the maximum size of the disks used. For example, using 9.1 GB drives will mean the largest partition size is also 9.1 GB.

RAID-1 and More than 2 GB of Memory

RETAIN Tip H165049

Servers with the SSA RAID Cluster Adapter are limited to 2 GB of system memory when any JBOD or RAID-1 arrays are configured. If you have 2 GB or more memory, your system may hang, experience a “blue screen of death” or other unpredictable failure. Review the RETAIN tip for details.

1.2.3 Array Configuration Issues

As with the SSA RAID Adapter, arrays can be spanned over the two loops of the adapter to achieve performance benefits and increased fault tolerance.

- With Microsoft Cluster Server, up to 22 RAID-1 arrays are supported. This number is limited by the maximum number of drive letters available to Windows NT. Typically, A-D are used by the diskette drive, non-SSA boot disk and CD-ROM drive.
- Hot spares are global to the adapter.
- Data scrubbing for the arrays is enabled by default and is done automatically in the disk subsystem idle time.

1.2.4 Disk Drive Access Modes

Since two systems are able to share the same disk resources, disk drive access modes were introduced. According to the function of the resource you can use one of the configuration tools to configure the disk resource as:

- Public: System access is available to this resource through both adapters in the loop. RAID-1 arrays are public by default, so you can't change this mode.
- Local: System access is available to this resource, but only through this adapter.
- Private: System access is not available to this resource through this adapter.
- No access: System access is not available to this resource to either of the adapters in the loop.

1.2.5 Booting of SSA RAID Cluster Adapter Drives

With the SSA RAID Cluster Adapter you can only boot from a JBOD resource which has to be defined as *Local* when viewed from the system you want to boot from. Booting of a RAID-1 array is not supported. The default drive access mode for RAID-1 arrays is public and can't be changed.

If you have problems booting from an SSA drive refer to “Booting from SSA Drives” on page 3.

1.2.6 Data Access

In a 2-way environment, the SSA cluster manages data access in a unique way. One adapter in the loop acts as the *primary* adapter, the other one as the *secondary* adapter. You can manually change this using one of the configuration tools and make the secondary adapter become the primary adapter and vice versa.

When accessing data stored on a single non-array disk (JBOD), the two servers in the cluster access data directly through the SSA RAID Cluster Adapter. However, when RAID-1 data is handled, transactions are executed only by the primary adapter in the cluster. Requests from the system with the secondary adapter are routed to the primary adapter and then executed.

This behavior is unique to the RAID Cluster Adapter and is not to be confused with the function of the *primary initiator*, which is a definition for all multi-initiator loops in SSA generally. This should also not be confused with the possible primary/secondary server function established on the cluster application level (for example, MSCS). If the primary adapter fails, the secondary adapter detects this failure and becomes the primary adapter, which is called a failover.

1.2.7 Performance Conclusions

The behavior described in 1.2.6, “Data Access” on page 5 results in the following recommendations to improve your performance:

- Make the adapter of the system which generates most of the I/O operations to RAID-1 arrays, the primary adapter. This will avoid unnecessary rerouting of I/O requests and improve the overall performance.
- If only one loop is used, we recommend you connect the spare ports of both adapters in a direct loop as shown in Figure 3 on page 5.

If the secondary adapter requests data through the primary adapter (that means from a RAID-1 array), the firmware of the adapters is intelligent enough to find the fastest way for the data transfer. Since the adapter can route data requests over both loops it can use the direct connection as an additional communications loop if it provides faster response.




Figure 3. Connecting the Unused Ports for Additional Loop

1.2.8 Supported Operating Systems

Only Windows NT 4.0 Server with Service Pack 4 or higher and Windows NT Enterprise Edition are supported operating systems for the SSA RAID Cluster Adapter.

Service Pack 4 is required since it resolves a potential problem that exists after multiple reboots of a clustered system using an SSA array as the shared storage

system. After multiple restarts, the SCSI reservations may lose a drive array because of repeated SCSI resets. See Microsoft Knowledge Base entry Q195636 for details.

Chapter 2. Preparations

Before you begin installing and setting up a high availability cluster solution using SSA hardware and Microsoft Cluster Server (MSCS), it is important to familiarize yourself with the following terms and definitions that are used in this manual.

2.1 Terms and Definitions

The following terms are used in clustering:

2.1.1 Node

Each server in a configured cluster is referred to as a *node*, to be connected together to form a cluster. The two nodes are also made available to client workstations via LAN connections. An additional independent network connection is used to perform monitoring within the cluster. One or more disk subsystems are attached to both nodes. Figure 4 shows the basic configuration of Microsoft Cluster Server.

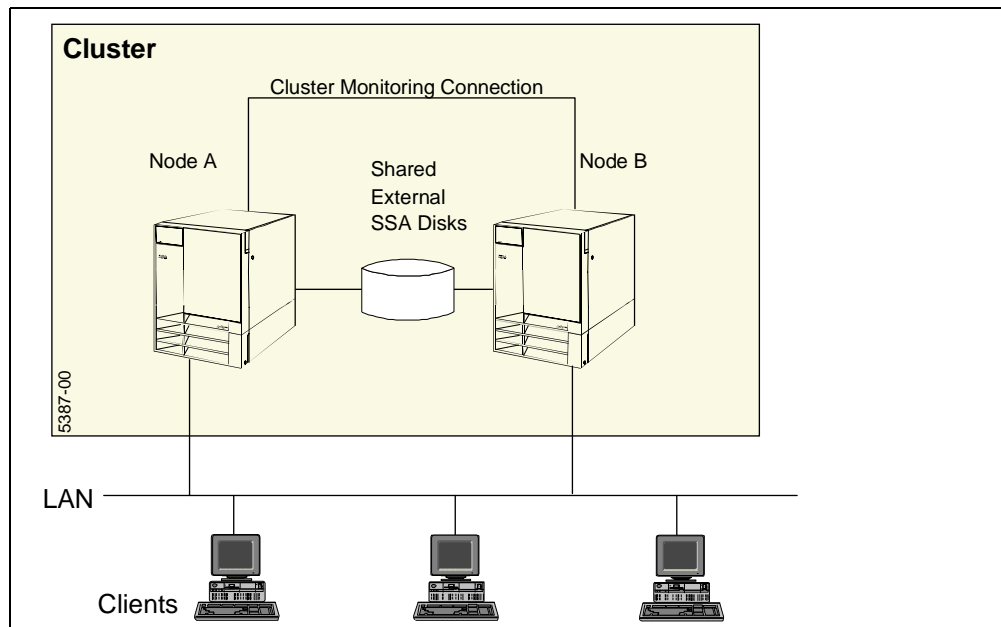


Figure 4. Basic Configuration

2.1.2 Resources

Resources are the applications, services or components running under the control of MSCS. The communication between the Resource Monitor and resources is provided by the resource DLL files. The resource DLLs, or *resource modules*, detect any change to the state of a resource and notify the Resource Monitor, which carries the information to the Cluster Service. Figure 5 shows this communication between the Cluster Service and the resource.

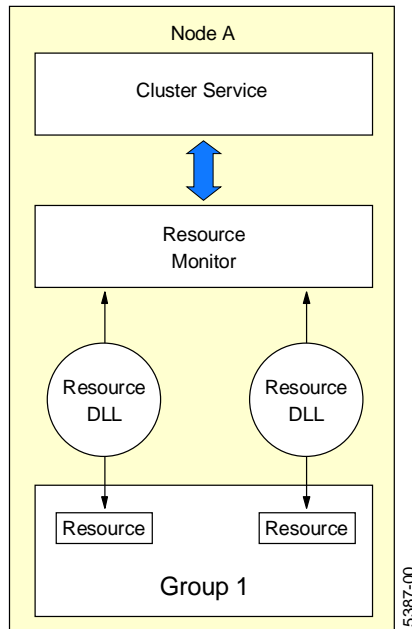


Figure 5. Communication between the Cluster Service and the Resources

2.1.2.1 Resource Monitor

Each node can have one or more resource monitor as shown in Figure 5. The resource monitor watches its assigned resource and notifies the Cluster Service if any change happens to its state. By default, the Cluster Service will start only one resource monitor to service all resource modules in the cluster environment.

This can be changed by choosing to run the resource in a separate resource monitor at the creation of the resource. Resources that will have to be debugged or resource modules that are likely to conflict with other resource modules should have a separate resource monitor.

The resource monitor is separate from the cluster server in order to provide an extra level of security. The resource DLLs are running in the address space of the applications themselves. If the applications fail, the resource DLLs may malfunction causing the resource monitor to fail as well. The Cluster Service, however, should remain available to the cluster.

2.1.2.2 Dependencies

During a failover, MSCS brings each resource online in a specific sequence that is defined when the resources are created. For example, if you want to create a file share resource, then the file share will require a disk drive. As a result, you will need to define a physical disk resource first; then when you define the file share resource, you specify that the physical disk resource is a dependency.

This can be visualized as a tree structure that shows the dependency linkages between resources (for example, Figure 6).

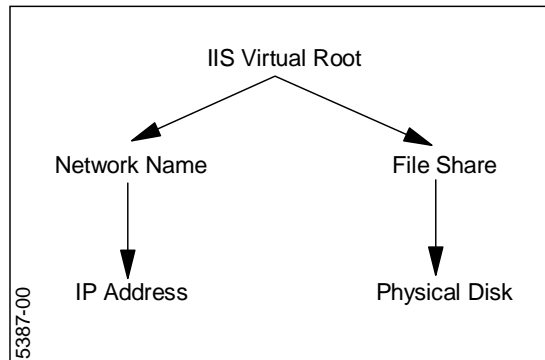


Figure 6. Resource Dependencies

Here the IIS virtual root resource is dependent on a network name resource and a file share resource.

As described in 2.1.4, “Resource Groups” on page 11, all dependent resources must be placed together in a single *resource group*.

2.1.2.3 Resource Types

MSCS defines 12 resource types. Other types of resources can be offered by software vendors using the application programming interface (API) in the Microsoft Platform Software Development Kit (SDK).

- DHCP Server

This resource type allows you to provide Dynamic Host Configuration Protocol (DHCP) services from MSCS. You will need to enter the path to the DHCP database to be able to provide the DHCP service. IP address resource and physical disk resource are the only two required dependencies for the DHCP server resource.

- Distributed Transaction Coordinator

This resource allows you to use Microsoft Distributed Transaction Coordinator (MSDTC) in MSCS. Only two dependencies are required to run this type of resource: physical disk resource and network name resource.

- File Share

This resource type lets you share a directory on one of the shared disks in your configuration and configure the permissions to give access to that directory to clients. You will be asked to enter the name of the share, the network path, a comment and the maximum number of users that can connect to the share at the same time. The configuration of a file share in this resource type is identical to the configuration of a file share in Windows NT Explorer. This resource requires the physical disk resource and the network name resource. For example, you can store files in a shared directory on the server and give access only to a group of clients.

- Generic Application

The generic application resource works with most applications and can be used with general applications that can fail over and be restarted by MSCS. There are no mandatory dependent resources. Examples of generic applications are CLOCK.EXE and NETMON.EXE.

- **Generic Service**

This resource type can be used with a generic service running on Windows NT. You must enter the exact name of the service at the creation of the resource. Like generic applications, the Generic Service resource does not require any dependent resources. An example of a generic service is the logon-authentication service.

- **IIS Virtual Root**

The IIS virtual root resource type provides the failover capabilities for Microsoft Internet Information Service Version 3.0 or later. There are three resource dependencies for an IIS resource: IP address resource, physical disk resource and network name resource. It will let you create an FTP, Gopher or WWW service.

- **IP Address**

The IP address resource can be used to assign an IP address. You need to obtain a static IP address and subnet mask from your network administrator. This new IP address will be bound to the network interface chosen at the *Network to use option* during the creation of the resource. This resource does not have any dependencies.

- **Microsoft Message Queue Server**

This resource supports clustered installations of Microsoft Message Queue Server (MSMQ). This resource is dependent on a distributed transaction coordinator resource, a physical disk resource and a network name resource.

- **Network Name**

This resource gives an identity to the cluster, allowing client workstations to see the cluster as a single server. The only dependency for this resource is an IP address resource. For example, if you create a network name resource WOLF1 and you have a file share resource UTIL, you can access it from a client desktop entering the path \\WOLF1\UTIL. This will give access to the directory on the shared disk regardless of which node actually owns the disk at the time.

- **Physical Disk**

When you first install MSCS on servers, you are asked to select the available disks on the shared storage. Each disk will be configured as a physical disk resource. If you find it necessary to add more disks after the installation, you would use the physical disk resource. This resource does not require any dependencies.

- **Print Spooler**

The print spooler resource provides a spool directory on the shared storage disk where print jobs will be spooled. Two resources are needed to create a print spooler resource: a physical disk resource and a network name resource.

- **Time Service**

This is a special resource that maintains the date and time consistency between two nodes. It does not require dependencies. Only one time service resource is necessary for each node. You should not create one for each group.

2.1.2.4 Resource States

Resources can be in one of five states at any time:

- Offline — The resource is not available for use by any other resource or client.
- Offline Pending — This is a transitional state. The resource is being taken offline. You can specify the amount of time that MSCS allows for a resource to go offline. If this resource cannot be taken offline within this time, the resource will go into the failed state.
- Online — The resource is available.
- Online Pending — This is a transitional state. The resource is being taken online. You can specify the amount of time that MSCS allows for a resource to go online. If this resource cannot be brought online within this time, the resource will go into the failed state.
- Failed — The resource has failed to transition from a pending state (online or offline) and is no longer available.

2.1.3 Quorum Resource

The quorum resource maintains data integrity and cluster unity and plays a critical role in cluster operation. It holds the cluster database which contains recovery logs, configuration and status information.

The quorum resource is used to guarantee that only one set of active, communicating nodes is allowed to operate as a cluster. Only one node has control of the quorum resource at any time. A node can form a new cluster only if it can gain control of the quorum resource and a node can join or remain in an existing cluster only if it can communicate with the node that controls the quorum resource.

The quorum resource must be of resource type Physical Disk and must be formatted as NTFS. The disks used must be installed in the external disk enclosure so that ownership of the resource can be transferred between the nodes in the case of a failover. You specify the location of the quorum resource when you install the first node of MSCS.

It 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. For reliability reasons, the disk should be configured as RAID-1.

The drive containing the quorum resource may also contain other applications and data.

2.1.4 Resource Groups

Resource groups are dependent resources that are grouped together. Some resources need other resources to run successfully. These are known as resource dependencies, as described in 2.1.2.2, “Dependencies” on page 8. When one resource is listed as a dependency for another resource, then the two resources *must* be placed in the same group, known as a *resource group*.

If all resources are ultimately dependent upon one resource (for example, a physical disk resource) then all resources must be in the same group. For this reason, it is possible, for example, that all your cluster resources will need to be in one group.

Any cluster operation on a group is performed on all resources that are present within that group. For example, if a resource within a group needs to be moved from Node A to Node B, all the resources defined in the group will be moved. Figure 7 shows an example of MSCS groups.

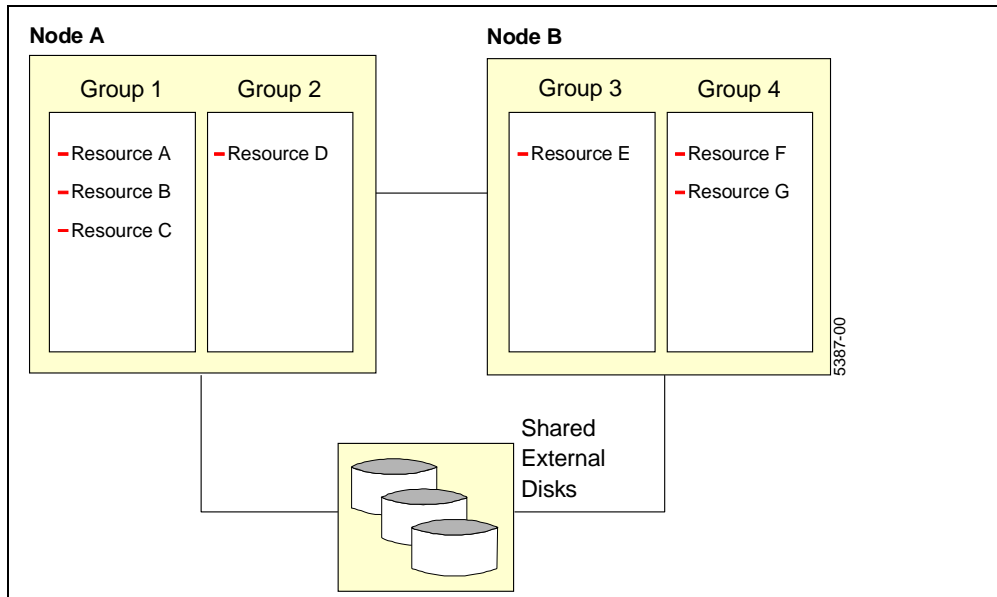


Figure 7. Example of MSCS Groups

2.1.4.1 Group States

A resource group can be in any one of the following states:

- Online – When all resources of the group are online.
- Offline – When all resources of the group are offline.
- Partially Online – When some resources are offline and some are online.

2.1.4.2 Virtual Server

Groups that contain at least an IP address resource and a network name resource appear on the network as *virtual servers*. A virtual server allows clients to connect to the cluster instead of connecting to the individual nodes.

It is very important to always connect to the virtual server to take advantage of the failover policy. For example, if you create a group IIS_Server, then when you browse your network you will see a server called IIS_Server under the same domain of the physical servers. You will also be able to see both physical servers that are under the same domain, but you should not connect to them.

Note: The cluster name is for administrators only. Clients access the applications and services using the virtual network names associated with the individual groups, not through the cluster name.

2.1.5 Failover

Failover is the relocation of resources in a group from a failed node to the surviving node. The detection of a failure is made by the resource monitor responsible for the resource. If a resource failure occurs, the resource monitor 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 steps:

1. Failure detection
2. Resource relocation
3. Application restart

Application restart usually takes the longest time to complete.

You can configure when MSCS performs failover by setting the two variables: Failover Threshold and Failover Period.

Group and Resource Failover Properties

Both groups and resources have failover threshold and period properties with the same name. However, the functions perform different actions as described below.

- Resource Failover Setting

Threshold is the number of times in the specified period that MSCS allows the resource to be restarted on the *same* node. If the threshold count is exceeded, the resource and all other resources in that group will fail over to another node in the cluster.

Period is the time (in seconds) during which the specified number of attempts to restart the resource must occur before the group fails over.

After exceeding the threshold count of restart attempts, MSCS fails over the group that contains the failing resource and every resource will be brought online according to the startup sequence defined by the dependencies.

- Group Failover Settings

Threshold is the maximum number of times that the group is allowed to fail over within the specified period. If the group fails over more often than specified, MSCS will leave it offline or partially online depending on the state of the resources in the group.

Period is the length of time (in hours) in which the group will be allowed to fail over only the number of times specified in Threshold.

For example, consider an application CLOCK.EXE in group CLOCKGROUP. Other resources in the group include a file share resource and a physical disk resource, as shown in Figure 8.

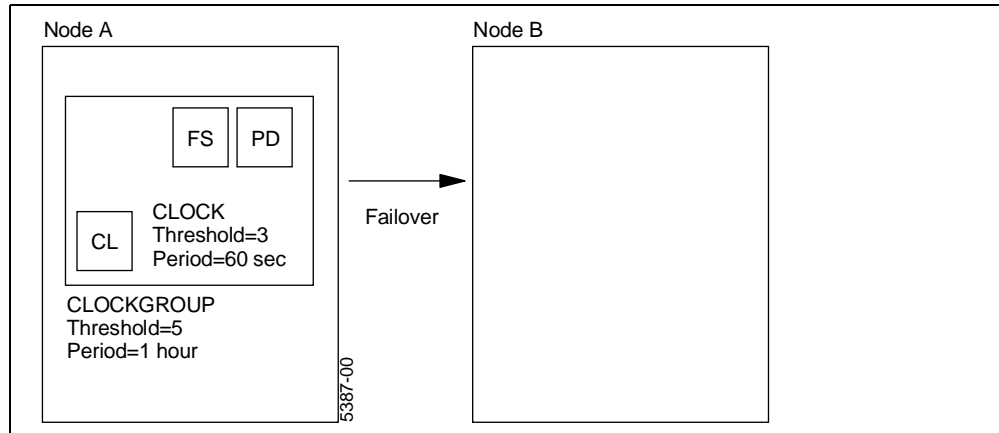


Figure 8. Failover Example

The CLOCK resource has a threshold of 3 and a period of 60 seconds. The CLOCKGROUP group has a threshold of 5 and a period of 1 hour.

Consider the situation when CLOCK continually fails. CLOCK will be restarted on node A three times. On the fourth time within 1 minute, it and its group CLOCKGROUP will fail over to Node B. This counts as one CLOCKGROUP failover. When CLOCK fails four times (that is one more than the threshold) on Node B, it will fail over to node A. This counts as the second CLOCKGROUP failover.

After the fifth CLOCKGROUP failover within 1 hour (Node A > B > A > B > A > B), MSCS will not attempt a restart of CLOCK, nor will it fail over CLOCKGROUP. Instead, it will leave CLOCK in the failed state and CLOCKGROUP will be placed in the partially online state. The other resources will be placed in the failed state if they are dependent on CLOCK or they will remain online if they are not dependent on CLOCK.

2.1.6 Failback

Failback is a special case of failover and is the process of moving back some or all groups to the preferred owner after a failover has occurred.

The *preferred owner* is the node in the cluster on which you prefer each group of resources to run. If the preferred owner fails, its resources will be transferred to another node. When the preferred owner comes back online and *allow failback* is enabled, the resources will automatically transfer back to that node.

You may set the preferred owner to be the more powerful system in your cluster to allow maximum performance. On the other hand, you may choose not to set a preferred owner, as it may not matter where the resources reside (such as when the systems are identical).

You can use the preferred owner settings to set up a simple load-balancing configuration. When both servers are running with failback enabled, the applications and resources will move to their preferred owner thereby balancing out the workload on the cluster as you defined it.

For SSA configurations, you should set the system with the primary SSA adapter to be the preferred owner for all I/O-intensive applications.

When you create a group, the default failback policy is set to prevent failback. In other words, if a failover occurs, the resources will be transferred to the other node and will remain there, regardless of whether the preferred node is online or not. You can allow failback by setting the group to move to the preferred node immediately after the node is available, or you can set the failback to occur between specific hours. The second option is very useful, because you may want the failback to occur only after the working hours or the peak business hours.

2.1.6.1 Looks Alive and Is Alive

You can adjust how often the resource monitor will verify if a resource is available and running. These actions are referred to as polls. There are two levels of polling. Exactly what polling is performed depends on how the resource module was written.

- Looks Alive polling

In Looks Alive polling, the resource monitor makes a superficial verification to determine if the resource is available.

If a resource fails to respond to a Looks Alive poll, then the resource monitor will notify the Cluster Service. When you create a new resource, you define the length of the interval (in milliseconds) between polling attempts.

- Is Alive polling

In Is Alive polling, the resource monitor performs a complete check over the resource to verify if it's fully operational. If it gets a failed response, then the Cluster Service is immediately notified and depending on the configuration defined for the resource, the resource manager can terminate the resource or try to bring it back online on the same node or on the other node.

Consider the Windows NT clock (CLOCK.EXE), and assume that CLOCK's resource is created with the default parameters (Looks Alive=5000 milliseconds and Is Alive=60000 milliseconds). The resource monitor calls the Looks Alive function in the resource module every 5 seconds to check that the clock is available for the cluster environment. Every 60 seconds, the Is Alive function is called to check if the clock is operating.

In both actions, the resource monitor will be looking for any failure of the clock. The only difference is that the Is Alive function is more thorough than the Looks Alive function. Depending on the implementation of the Looks Alive function in the resource DLL, the poll may involve issuing some status command to the resource to verify its successful operation.

2.2 Networking

In each machine, you should have at least two network adapter cards installed. You must define which one to use for cluster communications and the other(s) for client access. During the installation of Microsoft Cluster Server, your choices will be:

- Use for all communications (Cluster and Client)

The adapter will be used for internal communications between the cluster nodes and for client access.

- Use only for internal cluster communications (Cluster only)
Information between the cluster nodes will be interchanged on this adapter. Client access is only possible through the other adapters.
- Use only for client access
Clients are allowed access to the cluster through this adapter. Internal cluster information will be interchanged on another adapter.

An MSCS cluster can only have one network adapter on each node for internal cluster communications. It can, however, have one or more other network adapters for all (internal and client) communications. The “internal communications” adapter is the *primary interconnect*, and the “all communication” adapters serve as *backup interconnect* if the primary ever fails. Internal cluster communication is also called the *heartbeat*.

Only specific adapters are certified for internal cluster communications. Among these are the IBM 10/100 Ethernet adapter and the 3COM Fast Etherlink XL adapter.

If there is one token-ring and one Ethernet, use the token-ring adapter for client access to the cluster and the Ethernet adapter for internal cluster communications.

2.2.1 Redundancy

The reason for setting the client access adapter for all communications is so that in the event the adapter for internal cluster communications fails, MSCS will automatically use the client access adapter for cluster communications.

If there was no redundant path for cluster communications, the cluster nodes would immediately arbitrate for the quorum disk, and the winner would take control of all groups. The loser of the arbitration process would withdraw from the cluster by shutting down its cluster service.

This redundancy is not possible vice versa. Enabling all communications for the dedicated (internal cluster) adapters will not make the client connections redundant. Thus, in a normal configuration with the cluster nodes connected to one public network and by the crossover cable, there is no reason to enable all communications on the private adapter.

The situation changes if the nodes are connected to more than one public LAN. For example, consider the case that both cluster nodes are connected to a small-bandwidth network #1 (used for client access) and also a high-bandwidth network #2 (used for server-to-server transfer such as backup or database communications). Network #2 may act also as the cluster connection. Then you may define some cluster resources served to network #1 and others served to network #2. Now we have the situation where both networks have cluster-private traffic (redundancy) as well as public traffic (some of the cluster resources).

2.2.2 TCP/IP Considerations

MSCS uses TCP/IP to communicate with network applications and all resources. Note that MSCS cannot have an IP address assigned from a Dynamic Host Configuration Protocol (DHCP) server for any IP address resource or for the

cluster administration address (registered at the installation of MSCS). You must use a static IP address for both nodes and for the IP address resources.

Note: IP addresses for the individual nodes can be under the control of DHCP.

In the MSCS solution, each node will have at least two network adapter cards installed. The IP address assigned to each adapter will represent a public address or a private address. If the adapter is assigned a public IP address, the adapter will be used for client access. If the adapter is assigned a private address will be used for internal communications between the cluster nodes. As described in 2.2.1, "Redundancy" on page 16, you can also assign an adapter a public IP address and configure it to be used for both client access and internal communications.

There are IP address ranges that are reserved for private IP networks:

- 10: a single Class A network
- 172.16... 172.31: 16 contiguous Class B networks
- 192.168.0... 192.168.255: 256 contiguous Class C networks

We recommend that Class A addresses be used for internal cluster communications (simply because these are rarely used). For more information refer to *TCP/IP Tutorial and Technical Overview*, GG24-3376-05, "2.1.6 Intranets (Private IP Addresses)" and Chapter 3 of *Microsoft Cluster Server Administrator's Guide*.

Note

You must have TCP/IP installed on both servers in order to use MSCS. Applications that use only NetBEUI or IPX will not work with the failover ability of MSCS. However, NetBIOS over TCP/IP will work.

Additional comments about networking with Microsoft Cluster Server:

- Microsoft SNA Server, Proxy Server and WINS server currently have their own capabilities for high availability and do not use MSCS.
- Clustered servers can be connected to multiple subnets. MSCS additionally supports multi-homing configurations using multiple network cards. If we have two subnets connected by a router, however, there is no way to fail over an IP address resource.
- MSCS cannot use a second network card as a hot backup for the client access. That means that the card may be a critical failure point.

The problem of redundancy for the client network access must be solved in a network layer below the address assignment by NT. We recommend you use redundant NICs. For example, the Adaptec ANA-6911 and ANA-6944 with Dual-Link Option are on the IBM ServerProven list.

There is another interesting network effect which may lead to large cluster problems: the priority of network adapters in a multi-homed Windows NT server. Because each cluster node has adapters to at least two different networks (the cluster-private link and the public LAN), each cluster node is also a multi-homed host. On such a system the question "What's my IP address?" is answered by a whole list of IP addresses assigned to all network cards installed in the machine. The Windows Sockets `gethostbyname()` API is used for that.

Some cluster applications (for example, Oracle FailSafe and SAP R/3) are sensitive about the IP address order in the list returned by `gethostbyname()`. They require that the IP address of the network adapter to which their cluster virtual address will also be bound appears on top. This means that the addresses of the adapter to the public LAN must be shown before the address of the cluster-private link. If not, then it may be impossible to reach the application under the virtual address after a failover.

To avoid such problems, you should check the order in the address list before you install MSCS (to ensure that the network assignments are right from the beginning). The two simplest ways to do this are:

- Ping each node to itself. For example, on Node_A you type in a command window:

```
ping node_a
```

Then the address from which the ping is answered must be the address assigned to the NIC in the public LAN.
- The utility `IPCONFIG` shows all addresses in the order of the `gethostbyname()` list.

If you see addresses are in the wrong order, correct them immediately. When the NICs are of the same type, then it is sufficient to simply exchange their outgoing cable connections. If you have different NIC types (for example, 10/100 EtherJet for the cluster-private link, redundant FDDI for the public LAN), then you need a way to control the internal IP address order.

Under Windows NT 3.51, the IP address list is in the same order as the TCP/IP network card binding order; therefore, altering the TCP/IP network card binding order (**Control Panel > Network > Bindings**) will indirectly change the IP address order returned by `gethostbyname()`.

However, under Windows NT 4.0, that binding order does not influence the IP address order. Using the Move Up or Move Down buttons in the Bindings tab of the network control will not work. This is documented in the Microsoft Knowledge Base article number Q171320. You have to add a registry value, `DependOnService`, to change the IP address order.

Assume that your two network adapter cards have the driver names `Netcard1` and `Netcard2`. (**Note:** The 10/100 EtherJet Adapter has the driver name `IBMFE`.) Ping and `ipconfig` show you `Netcard1` first, but your public LAN is on `Netcard2`. To change the IP address order to list `Netcard2`'s IP addresses first:

1. Start Registry Editor (`REGEDT32.EXE`) and select the following subkey:

```
HKEY_LOCAL_MACHINE\SYSTEM\Current\ControlSet\Service\Netcard1
```

2. From the Edit menu, click **Add Value**, and enter the following data:

```
Value Name: DependOnService
```

```
Value Type: REG_MULTI_SZ
```

```
Data: Netcard2
```

3. Exit from Registry Editor and reboot the machine. Now PING and `IPCONFIG` should show you the right address.

Note that the additional registry value added here is deleted when NT rebuilds the network bindings. Thus after each modification of network parameters you should verify that the order is still right.

2.3 Domains

The following information specifies the criteria for servers in a domain:

- The two servers must be members of the same domain.
- A server can be a member of only one cluster.
- The following configurations of servers in a cluster are possible:
 - A primary domain controller and a backup domain controller
 - Two backup domain controllers
 - Two stand-alone servers

2.4 Clustering Models

There are two main clustering models used in clustering today:

- Shared Disk

In this model, applications running on one node may access any data on disks connected to any other node in the cluster. If two nodes want to read the same data, each node must read the data separately, or one node must transfer the data to the other. If two nodes want to write data to the disk, the writes must be coordinated to ensure data integrity is maintained. The model allows for multiple readers at any one time, but only one write operation can be performed to ensure consistency.

The shared disk model is useful in a load-balancing situation, in that the data on any disk can be obtained by any node at any time. The down side is that the software required to manage this type of environment is significantly more complex than that in the Shared Nothing model.

- Shared Nothing

Here, each node in the cluster owns some of the resources of the cluster. No resource is owned by more than one node at any time. Only one node owns and is allowed to use each disk at any time. If another node needs to access data on a disk, a request is sent to the owning node, which performs the request on its behalf. If a failure occurs, the resource can automatically be transferred to another node and any requests for that resource are automatically re-routed to the new owner. This ensures that all client requests are fulfilled regardless of the status of individual nodes.

This model can be extended so that when a client request is made of the cluster, the request can be divided up and given to each of the nodes in the cluster. The nodes then process their part of the request and the results are then assembled and returned to the client.

Microsoft Cluster Server is based on the Shared Nothing model.

Even though most MSCS hardware implementations are based on a shared SCSI bus, MSCS is considered a Shared Nothing clustering architecture as every server has its own system disk. No concurrent access to data on the physically

shared disk is allowed. If implemented in a specific application, data can be shared via requesting the information from the owner of the disk and then retrieving the data from the owner.

2.5 SSA Considerations

The following should be considered when configuring a cluster configuration using the SSA RAID Cluster Adapter:

2.5.1 SSA RAID Cluster Adapter Considerations

- The SSA RAID Cluster Adapter drivers, BIOS, firmware and utilities should be the latest certified version. Get them from the following Web pages:
 - SSA firmware and device driver diskettes (latest MSCS certified driver)
<http://www.hursley.ibm.com/~ssa/pcserver>
 - System BIOS: <http://www.pc.ibm.com/us/searchfiles>
 - Network adapter device drivers:
<http://www.networking.ibm.com/nes/neshome.html>

Note: The SSA driver should be certified for use with MSCS.

- Only certain configurations are certified for use with Microsoft Cluster Server. This applies to software and firmware as well as hardware. See the following Web pages for the latest information:
<http://www.hursley.ibm.com/~ssa/pcserver>
http://www.pc.ibm.com/us/netfinity/cluster_server.html
<http://www.microsoft.com/hwtest/hcl>
- If there are two adapters in a particular SSA loop, they must both be SSA RAID Cluster adapters. You cannot mix the adapter types.
- Up to three adapters can be installed in a server, but all loops have to be formed such that the cabling, connectors and adapter placement must be mirrored between the two servers. For example, if the SSA cable is connected to connector B1 on the adapter in slot 3 in the first server, then the cable must also be connected to B1 on the adapter in slot 3 of the second server.
- One of the SSA adapters is the primary adapter or *master initiator*; the other is the secondary or remote adapter. By default the adapter with the highest unique ID is the primary adapter. However, this can be user-defined.
- Disk drives that are not in an array may be restricted to local access from one server only. This is a way of configuring an SSA boot disk even in a shared SSA loop. Be aware of setting the disk resource number (described in 2.5.6, “Disk Types and Disk Numbers” on page 22) such that the boot disk appears on a different bus to NT. Arrays are always in public access mode.
- If you want to boot from an SSA disk drive, install the SSA adapter in a PCI slot with a higher priority than the slots used by any SCSI adapters. Check the documentation of your server for the boot sequence.
- If the non-primary adapter detects that the primary fails, and then takes control of the arrays, then it may start to resynchronize some or all arrays. This will happen for all arrays to which the primary adapter made write operations during the last 20 seconds before its failure.

2.5.2 Connection Considerations

- At most two servers can be connected together in one loop. Only one SSA RAID Cluster adapter in each server can be in that loop.
- The adapters must be installed in *identical* PCI slots in each server. For example, if you install the adapter in slot 1 of one server (that is, the highest priority slot), the adapter in the second server must also be installed in slot 1.
- Each adapter has four connectors A1, A2, B1 and B2. When joining the adapters of the two servers together, ensure that they are connected to the same port pair (that is, either the A port or the B port). For example, connector A1 on one adapter must be connected to A1 or A2 on the other adapter and so on. The reason for this is that a single serial interface chip (SIC) controls each port pair and must have control of the entire loop. The SSA RAID Cluster Adapter has two SICs, one for each port pair.
- If two SSA adapters are in a loop and there is a second loop attached to one adapter, it must be attached to the other adapter also. You cannot build a shared and a local loop with different ports of the same adapter. (If you want to have an SSA boot disk, use local access mode and disk numbering as described in 2.5.6, “Disk Types and Disk Numbers” on page 22.)

2.5.3 RAID Considerations

- The SSA RAID Cluster Adapter provides support for RAID-1 and non-RAID disks only. The RAID-1 arrays can only include two disks. Three or more disks per array (that is, RAID 1 enhanced) is not supported.
- The array selected as the quorum resource must be RAID 1. (See 2.1.3, “Quorum Resource” on page 11.)
- All RAID-1 transactions are controlled by the primary adapter. Any transactions that originate from the remote adapter are routed to the primary adapter. There are, therefore, performance benefits in configuring MSCS so that all I/O-intensive applications have their preferred node to be the server with the primary adapter installed in it. (Refer to 2.1.6, “Failback” on page 14 for more information about preferred nodes.)
- If you move a RAID-1 array from a non-cluster SSA RAID Adapter to an SSA RAID Cluster Adapter, you must change the array's SplitConfirm attribute to enabled. If you are going to move a RAID-1 array from an SSA RAID Cluster to a non-cluster SSA RAID Adapter, before you move it change its SplitConfirm attribute to disabled.

2.5.4 SSA 7133 Enclosure Considerations

- When connecting the 7133s to the adapters, it is recommended for performance reasons that you split the drives so that half are connected through the A1 and A2 connectors, and half are through the B1 and B2 connectors of each adapter. If you wish to have all drives connected on the A1 and A2 connectors, we recommend that you join B1-B1 and B2-B2 with SSA cables. This will allow additional I/O traffic to be relayed on the B connectors as well as the A connectors.
- The 7133 can hold 16 disks, divided into groups of four. If you don't have an exact multiple of four disks, you will need to install dummy connectors to fill in the gaps to make up 4, 8, 12 or 16 connections. The 7133 is delivered with dummy connectors (dummy modules) installed in all bays where drives are not

present. For performance reasons, it is recommended that you install the disks in positions 1 and 4 of the group first and put the dummy connectors in positions 2 and 3 of the group.

- The 7133-020 and 7133-600 have bypass circuits used to isolate the groups of four drives from each other. These circuits allow the configuring of one, two or three groups without the use of external jumpering.
- Only up to three dummy modules can be used in sequence. Four or more in a row is not supported.
- Another (but very small) performance difference comes from the position of a disk drive in relation to the adapter. Under heavy load, the performance is a little bit better for disks that are nearer to the adapter.

2.5.5 Fiber Optic Considerations

If you plan a cluster with SSA and fiber extenders between several rooms or buildings, you should consider the risk of cluster partitioning. The term *cluster partitioning* describes a kind of split-brain syndrome, which may happen in distributed systems when all connections are lost. In such a situation, neither node receives a heartbeat from the other, and each resource on both nodes is seen as offline. The cluster service cannot distinguish between connection and node failures. Each node tries to fail over all resources to the other node, which would lead to an inconsistent cluster state.

Normally, the quorum resource acts as a tiebreaker in such situations. But, in a configuration where the quorum disk is also mirrored between two data centers each site would have its own copy of the quorum resource. The tiebreaker algorithm would present incorrect results. Therefore, you run the risk of all mirror sets being broken. Each node assigns all TCP/IP addresses and restarts all applications, but updates only one half of the disks. There is no way to get a consistent copy of data without losing the transactions processed at one node during the partitioning.

In SCSI configurations, this possibility can be ignored because a failure of all SCSI and network connections at the same time (in such a small place such as a rack and in such a way that each machine can access some disks) is very unlikely. With SSA fiber extenders and mirroring between separate machine rooms, we have to consider connection loss as a real danger. The fibers and cables are more volatile when going throughout the building or across the campus. Thus we recommend following these guidelines when using fiber extenders and mirroring over large distances:

- Ensure that all cables (user-traffic network, cluster-private network, SSA) are laid in such a way that there is no possibility to lose all connections at the same time. Carefully check all common cable tubes and building entry points. Always build SSA loops and FDDI rings as real loops with separate paths for each half, not just with cores in the same cable.
- If possible, place the quorum disk set in a third room, distinct from the rooms with the cluster nodes.

2.5.6 Disk Types and Disk Numbers

The SSA RAID Cluster Adapter has a multi-level concept of disks. The lowest levels are new disks and free disks; each disk is initially in one of these states.

Then you can define arrays or hot spares with members from the free disk list. These arrays or single disks are not visible to the operating system. They must be attached to the system by defining them as system resources. When the SSA configuration tool attaches the array or the single disk drive to the system, then a disk number is assigned to that disk drive or array.

The number that is assigned is related to the position in the System Resource List at which the resource (disk drive or array) is attached. The SSA configurator provides an option to change the disk number to another number, if required. Valid numbers are 1 through 255. They are shown always as hexadecimal numbers in the configuration tool (see Figure 9).

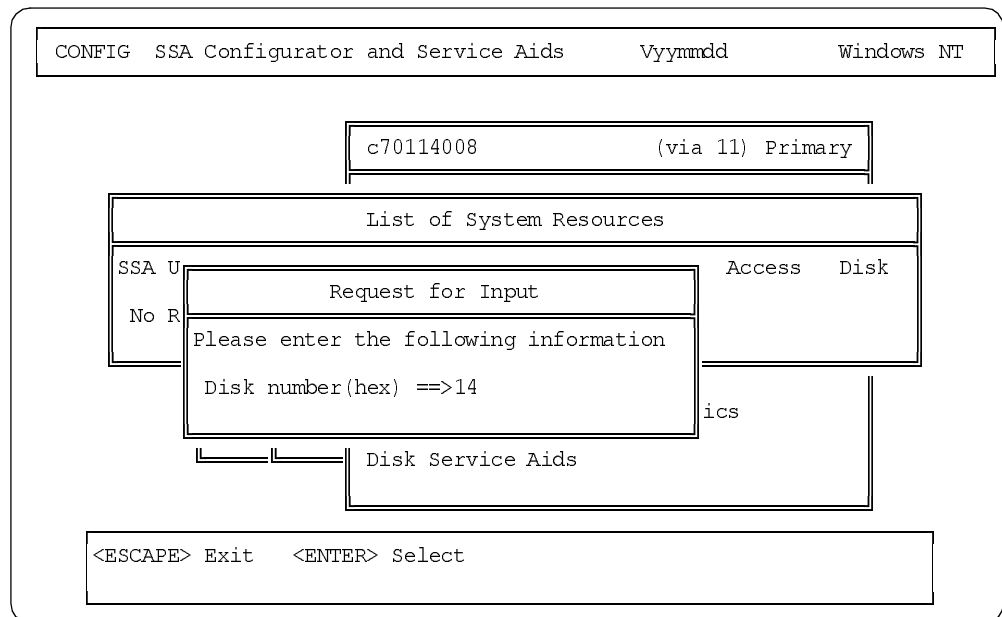


Figure 9. Disk Number

Disk numbers are similar to the SCSI ID settings on SCSI disk drives. The system software uses the disk numbers to map the disk drives to a logical bus, target, and LUN. The logical bus number to which a system resource is mapped is:

$$\text{Logical bus number} = \text{The disk number divided by } 32$$

where 32 is the number of targets on each bus. Therefore, a disk number that is lower than 32 is mapped to logical bus 0. The target address is the remainder that results from dividing the disk number by 32. In this way, you can manipulate which drives or arrays are seen by NT on the same or on different buses. MSCS requires that the boot device be on a different logical bus from any shared disks.

2.6 Planning for Microsoft Cluster Server

In this section, we discuss topics that you should consider when planning to implement a Microsoft Cluster Server environment in your organization.

2.6.1 Identifying Risks

Microsoft Cluster Server is a component of Windows NT Server 4.0 Enterprise Edition that provides improved availability for your data and applications. It is

designed to allow applications to be restarted on a second server if either the first server fails or the application fails.

MSCS is *not* designed to act as backup software and it can't be used to protect your data from all types of problems. MSCS protects the availability of data to the user community, but can't protect the data itself. You should plan to use MSCS with other high availability products and techniques such as redundant RAID disk arrays, backup schemes, UPS units and disaster recovery strategies.

Failures that MSCS and Windows NT can address include (by failing over from one server to another):

- Server connection failure
- Server hardware such as CPU or memory failure
- Operating system or application failure
- Network hub failure using redundant network connections

Failures that MSCS cannot address include:

- Power supply – use a UPS
- Disk – use RAID and backup solutions
- Data – use a backup solution
- Network cards – use fault-tolerant NICs
- Routers – use redundant links
- Dial-up – use multiple modems/lines
- User workstations – have spare workstations
- Major disaster – use off-site disaster recovery facilities

2.6.1.1 Clusterable Applications

Most applications can be used in an MSCS environment. The only restrictions are that the application must communicate via TCP/IP and that it can place data on any disk in the server.

Applications can be aware of the cluster and can therefore communicate with MSCS via APIs, or they can be unaware or *generic* applications. Cluster-aware applications respond intelligently to *Looks Alive* and *Is Alive* polling and can be managed using MSCS's administration tools or the IBM Cluster Systems Management software. (See 2.1.6.1, “Looks Alive and Is Alive” on page 15.)

A failover can occur when there is a hardware malfunction. This can mean loss of data stored in the server's memory. Consequently those applications that store information about their current state on disk rather than memory are more suitable for use in MSCS. Similarly, applications that can restart where they left off based on that saved status information are also more suitable. Those applications that simply “start over” are more likely to lose data during a failover.

2.6.1.2 Failover Options

As described in 2.1.5, “Failover” on page 13 and 2.1.6, “Failback” on page 14, if an application fails rather than the operating system or the server itself, MSCS can be configured to restart the application on the *same* server a specific number of times before transferring it to the second server. You can also specify how many times this failover will occur before MSCS gives up and leaves the application in a non-operating state.

In configuring MSCS, you can also specify which server or node is the preferred node for a particular application. If that node fails, the application is restarted on the second node. If the preferred node comes back online, MSCS can be configured to automatically return the application to its preferred node. You can specify whether this *failback* occurs immediately or between certain hours in the day.

2.6.2 Cluster Configurations

As per Chapter 2 of the *Microsoft Cluster Server Administrator's Guide*, MSCS can be set in one of a number of configurations depending on your availability requirements. Here are some examples:

- Active-Active

This is typically how an MSCS configuration is set up. Applications run in production on both servers in the cluster. Should either server fail, the applications on that server are restarted on the surviving server. The surviving server may run in a degraded state due to the extra workload placed upon it.

- Active-Hot Spare

In this cluster configuration, the second node does not normally perform any useful work, but should the first server fail, all applications will restart on the second node with no performance degradation. The downside of this configuration is the extra hardware required to implement the hot-spare.

- Mixed Failover

While you may have applications that are suitable candidates for MSCS, you may also have applications (such as those which are based on NetBEUI and cannot use TCP/IP) that cannot be clustered.

In this situation, you can configure MSCS to allow failover of some applications but not others. If a hardware failure occurs, those NetBEUI applications will not fail over, but when the server comes back online, they can be automatically restarted.

2.6.2.1 Determining Groups

After deciding what form your cluster will take, you need to group your applications and resources together. Review Chapter 2 of the *Microsoft Cluster Server Administrator's Guide*.

Things you should consider when setting up your groups:

- All the resources that are dependent upon one another must be in the same group.
- Resources can only be in one group.
- Whole groups fail over, not individual resources. When a resource fails over, the entire group fails over with it.
- It can be helpful to draw a dependency tree depicting all the resources and how they interact with each other.
- Refer to 2.1.2, "Resources" on page 7 for details about resources, resource types and dependencies between them.

Chapter 3. SSA Configurators

With the SSA RAID Cluster Adapter, you can configure your SSA disk subsystem with any of the following configurators:

- SSA Remote System Management and Netfinity Web Extension, a Web browser-based SSA Configurator
- SSA Remote System Management Configurator, a stand-alone Windows NT Service Web browser-based SSA Configurator
- Offline DOS SSA Configurator and Utilities
- Online Windows NT Configurator

If you want to boot from an SSA disk you need to use the DOS configurator to configure your server with the SSA enclosure, prior to installing the operating system there. This DOS Configurator and Utility diskette is a bootable diskette.

To use the online utilities, you need a running operating system with the SSA device driver already installed. You should also copy the utilities to the ISSA directory on your boot device. You can then run the ISSACFG command to configure your SSA environment.

3.1 Obtaining the Software

To download the latest DOS SSA Configurator, Device Driver and Utilities diskettes, BIOS/Firmware upgrade or the latest SSA Remote System Management configurator, go to:

<http://www.hursley.ibm.com/~ssa/pcserver>

The software is available for download as a self-extracting EXE.

3.2 Remote System Management

RSM is a browser-based tool for configuration of SSA adapters and devices. It enables you to configure SSA devices attached to any IBM SSA RAID adapter that is running the SSA Remote Systems Management (RSM) Service, locally or remotely, via any of the supported browsers. RSM was designed as a replacement for the text-based SSA Configurator.

RSM runs on Windows NT Server 4.0 only.

Two versions are currently available:

- A Netfinity Web Extension for servers running Netfinity Manager V5.0, V5.1 or V5.2, with Web Manager installed and activated.
- A stand-alone Windows NT Server 4.0 service.

Both versions have the same look and operability. Both versions can be installed on the same server and run at the same time; however, the Server Service version is only single-threaded so two concurrent users will notice delays in response. The Netfinity Manager version is multi-threaded and supports multiple user accesses. For this reason, if you have Netfinity Manager installed, we recommend you install and use that version.

3.2.1 Installing RSM

See 4.6, “Installing RSM” on page 38 on how to install and run RSM.

3.2.2 Using RSM

This user ID has full administrator access to your SSA subsystem, so you should create other, more secure user IDs and delete this one as soon as practical. User IDs and passwords are case-sensitive and can only be from 1 to 10 characters. Blank passwords cannot be used. We recommend that at least two users be created, one to administer the SSA devices and one to simply view the current configuration

After entering the user ID and password, the RSM main window appears, similar to Figure 10:

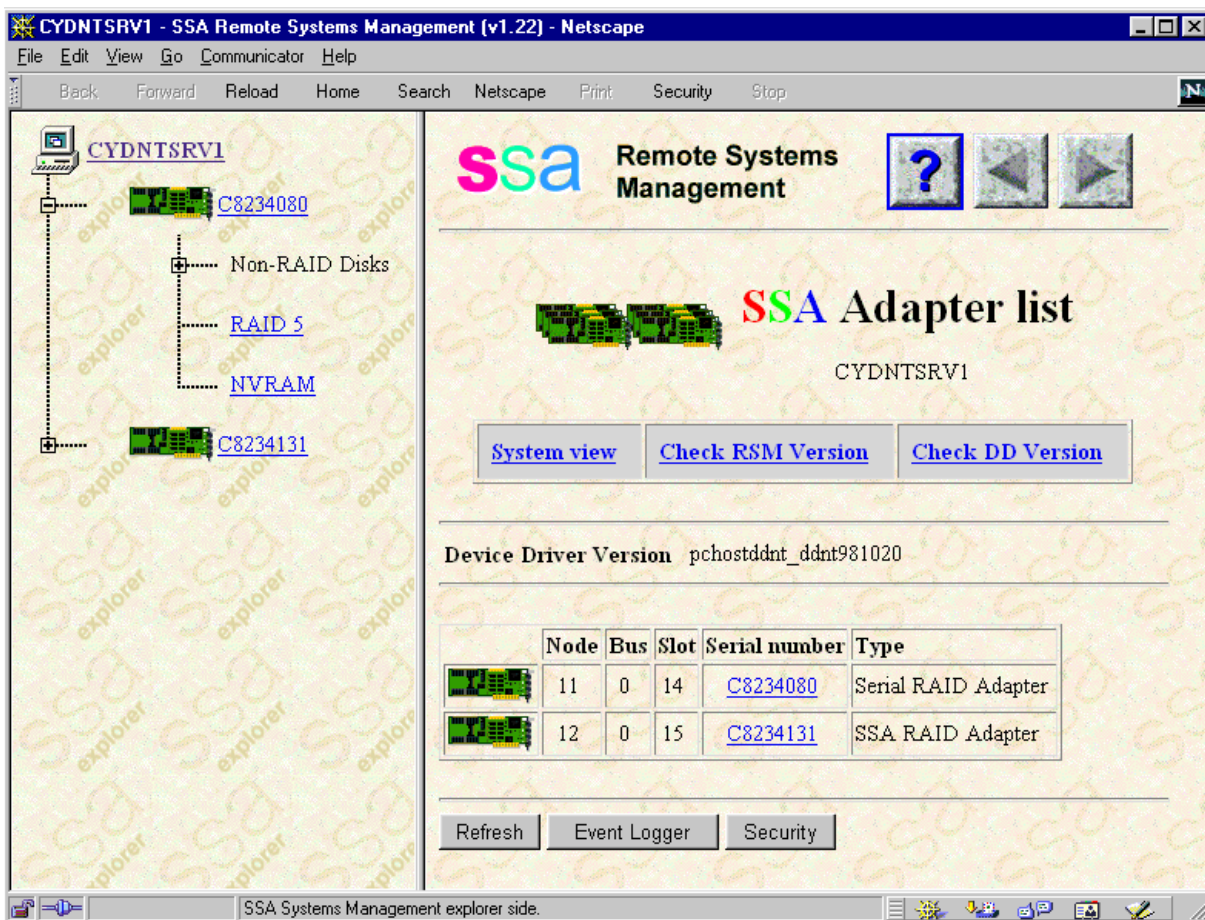


Figure 10. RSM Main Window — Stand-Alone Version

Proxy Settings

If you use a proxy gateway, add the TCP/IP address 127.0.0.1 into the “do not use proxy server for addresses beginning with” proxy configuration field of your Web browser.

From the main RSM screen, you can access any of the following:

- The configurator itself
- The RSM tutorial

- HTML versions of all SSA readme files
- Other reference material

3.2.3 Tutorial

The RSM tutorial is very comprehensive and we recommend you follow the tutorial to familiarize yourself with the product. You will find the following topics there:

- SSA RSM Tutorial Overview
 - What is RSM
 - SSA Terminology
 - Design Concepts
 - Local and Remote Configuration
 - Help
 - Code Updates
- SSA RSM Tutorial Using RSM
 - Views (Explorer, Adapter, Logical, Physical, Enclosure, System)
 - Security
 - Event Log
- SSA RSM Tutorial Configuration Examples
 - System RAID 5 Resource
 - Making This Array a Fast Write Resource
 - System Non-RAID Disk
 - Hot Spare Non-RAID Disk

3.3 SSA Text-Mode Configurator

The DOS configurator (and its online counterpart in Windows NT) is a text-mode configurator that allows you to configure the SSA disk subsystem in much the same way that the RSM software does.

You would use the DOS configurator, in particular, to configure the SSA disk subsystem if you planned to boot from SSA disks. The DOS configurator runs offline from a bootable diskette.

Note: If you install Version 1.10 or later of the SSA Event Logger, you will also need to obtain the latest version of the text-mode configurator. You can download this from:

<http://www.hursley.ibm.com/~ssa/pcserver>

Figure 11 shows the SSA RAID Configurator for DOS. However, the configurator for Windows NT looks almost identical.

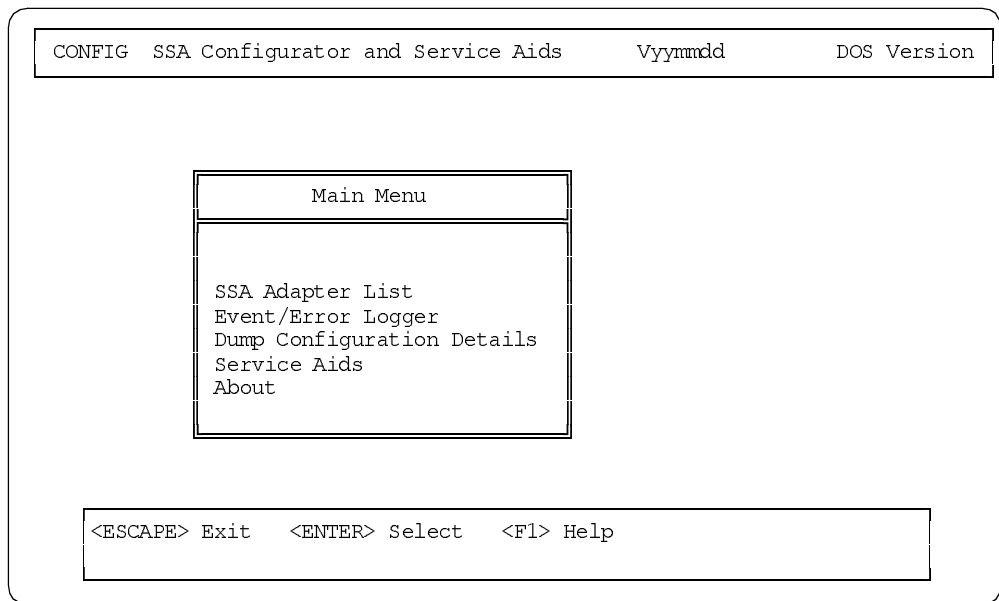


Figure 11. SSA DOS Configurator Main Menu

See the *SSA RAID Cluster Adapter User's Guide* for more information about how to use the configurator.

Chapter 4. Installing Windows NT and SSA Tools

This chapter explains, step by step, how to install Windows NT Server Enterprise Edition and the latest SSA drivers and firmware for use in a two-node SSA shared-disk environment. We cover MSCS installation in Chapter 6, "Installing Microsoft Cluster Server" on page 61.

The steps covered here are:

1. Install Windows NT Server Enterprise Edition
2. Install SSA RAID Cluster Adapter device drivers
3. Update SSA RAID Cluster Adapter firmware
4. Install SSA Remote Systems Management (RSM)

Note: The operating system is to be installed on a non-SSA disk.

4.1 Preparation

Prior to the installation of Windows NT Server Enterprise Edition, ensure the following activities have been performed:

- Upgrade the server's system BIOS and the SCSI adapter's firmware to the latest level. Obtain the latest level from <http://www.pc.ibm.com/support>.
- You have the latest drivers and firmware for the SSA RAID Cluster Adapter. Obtain them from <http://www.hursley.ibm.com/~ssa/pcserver>.
- Install the SSA RAID Cluster Adapter in both servers.
- Windows NT is not installed on either node.
- We recommend you perform a low-level format on your local disks. If they are ServeRAID-attached, you can use the format option in the Advanced Features menu of the ServeRAID Configuration Utility diskette. If they are attached to an Adaptec controller, press Ctrl-A during boot to access the function.
- If you plan to install the operating system on ServeRAID-attached devices, you will need to configure the ServeRAID adapter so as to create arrays and a logical drive on the local disks on which to install the operating system.
- You do not have either node connected to the SSA external disk enclosures.
- If a Windows NT domain environment has already been set up, MSCS nodes can be members of that domain. If you plan to have Node A and Node B perform the roles of PDC and BDC, you can do so, but you will have to leave Node A up and running as you install Node B. You should not, however, log on to the domain on Node A. See 2.3, "Domains" on page 19 for other information.

You should also consult Chapter 2, "Preparations" on page 7.

4.2 Installing Windows NT Server Enterprise Edition

To install Windows NT Server Enterprise Edition on Node A and Node B, do the following:

1. Insert the Windows NT Server Enterprise Edition CD-ROM 1 and power on Node A to begin the installation

2. If your local disks are attached to a ServeRAID controller, you will need to press F6 when you first see Setup is Inspecting your hardware to indicate you wish to supply a device driver. When requested, insert the ServeRAID device driver diskette.
3. Do *not* install the SSA RAID Cluster Adapter drivers during the installation process. You will install the drivers and update the adapter firmware later.
4. Format your primary partition as NTFS.
5. When prompted for the computer name, use NODE_A or something suitable to show which node in the cluster it is. Consult with your network administrator about computer naming conventions at your site.
6. When prompted for the administrator password, ensure it is not left blank. Microsoft Cluster Server requires the administrator to have a password.
7. If you plan to install Internet Information Server (IIS) as a cluster-aware application, you should install it later after the SSA drives have been brought online. Deselect the checkbox to install IIS when prompted.
8. You will need to install the network adapter device drivers for all network adapters, including the interconnect adapters. Be aware that TCP/IP is the only network protocol supported for failover by Microsoft Cluster Server.

In our configuration we used the following TCP/IP addresses:

- Public network: 9.9.9.1 and 9.9.9.2, subnet mask 255.0.0.0
- Private interconnect network: 10.10.10.1 and 10.10.10.2, subnet mask 255.255.0.0

Note: Although the 255.0.0.0 subnet mask for the public network was used in our example, your network may require the use of Class C addressing and subnetting.

9. When prompted, install Service Pack 3. You will install Service Pack 4 after you have installed Microsoft Cluster Server, as described in Chapter 6, “Installing Microsoft Cluster Server” on page 61.
10. Reboot the server and log on to Windows NT as an administrator.
11. When the system boots, you will see the Windows NT Server Enterprise Edition Installer window, Figure 12:

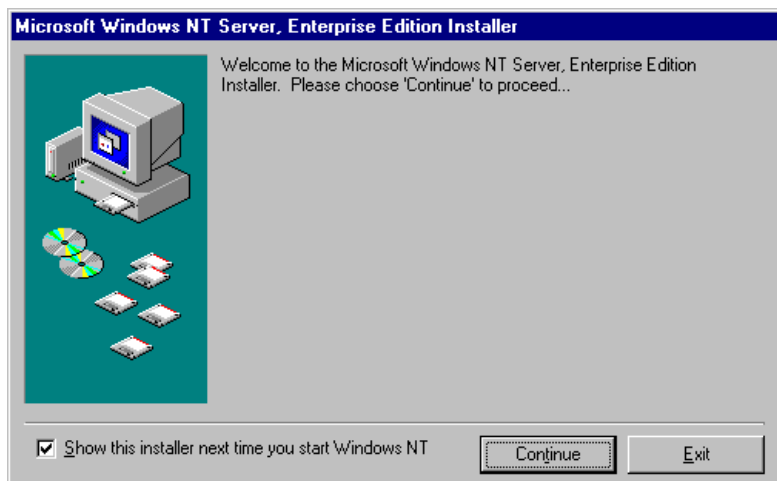


Figure 12. Enterprise Edition Installer

12. Do not use the installer at this time. Uncheck the box **Show this installer next time you start Windows NT** and click on **Continue**. (Do not exit from this screen.)

13. On the next screen, click on **Exit** and confirm you wish to exit Setup.

14. If you haven't done so already, power on Node B and repeat steps 1 to 13.

Note: If you plan to have Node A and Node B perform the roles of PDC and BDC, you will have to leave Node A up and running but not logged on to the domain. If your clustered servers are simply being installed as stand-alone servers in a pre-existing domain, you should power off Node A once Windows NT is installed.

At this point, Windows NT Server Enterprise Edition and Service Pack 3 are now installed on local disks in both servers. Both servers are still not connected to any SSA disk enclosures.

4.3 Installing the Device Drivers

You are now going to load the device drivers for the SSA adapter on each of the two server nodes. The process involves installing two drivers: the driver for the SSA adapter and the SSA-SCSI driver which replaces the standard DISK.SYS driver. The process also installs the SSA Event Logger and the text-mode SSA Configurator tools.

1. Open the Control Panel and double-click on the **SCSI Adapters** icon.
2. In the SCSI Adapters window, go to the Drivers tab, click on **Add** and then on the **Have disk** button.
3. Insert the driver diskette you downloaded from <http://www.hursley.ibm.com/~ssa/pcserver>.
4. Enter A:\ if required and click on **OK**.
5. When prompted, choose **96H9835 IBM SSA RAID Cluster Adapter** within the list of drivers and click **OK**.
6. Enter A:\ when prompted for the location of the SSA adapter files and click **Continue**.
7. The files are now copied from diskette. You will then see Figure 13:



Figure 13. Unique Identifier Required

8. Click on **OK** on the pop-up window. You will now see Figure 14.

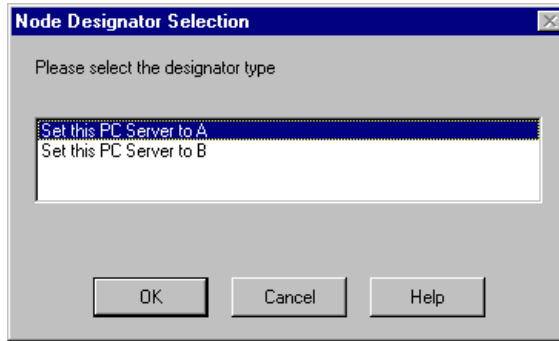


Figure 14. Node Designator Selection Window

9. The SSA RAID Cluster Adapter must be assigned a unique identifier, in the Node Designator Selection window. Select **Set this PC Server to A** if you are installing on Node A or **Set this PC Server to B** if you are installing on Node B and then click on **OK**.

10. A pop-up window, Figure 15, informs you that you need to install the IBM SSA-SCSI driver. Click on **OK**.




Figure 15. SSA-SCSI Driver Required

11. When you are prompted to restart your computer you must click on **No**. Do not restart your computer at this time.

12. In the drivers window, Figure 16, you'll see the SSA RAID Cluster Adapter listed as well as any other SCSI drivers that are loaded (on-board Adaptec SCSI driver, IDE CD-ROM, and so on).

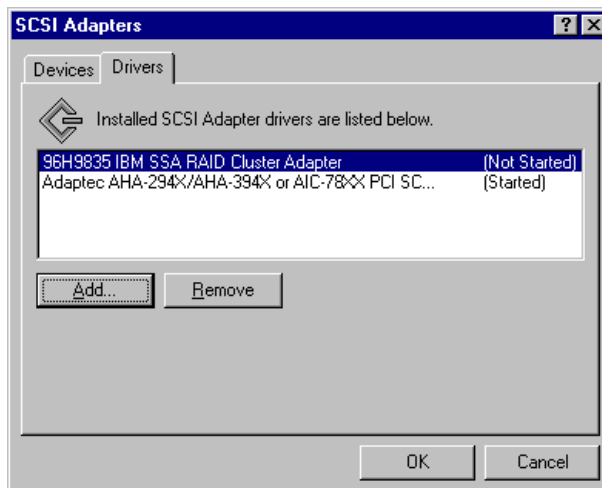


Figure 16. SSA RAID Cluster Adapter Driver Installed

13. Click on the **Add** button again to add the SSA-SCSI driver.

14. Figure 17 appears, displaying the list of drivers available to install (from the previous installation of drivers).

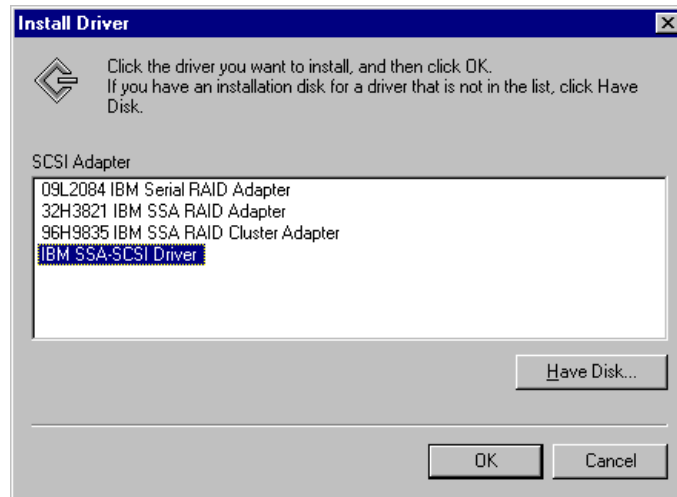


Figure 17. List of Drivers to Install

15. Highlight **IBM SSA-SCSI** and click on **OK**. Type in A:\ to specify the path to the driver files and click **Continue**. Files are now copied from diskette.

16. Figure 18 appears indicating that the IBM SSA-SCSI driver has been installed replacing DISK.SYS. Read the message for the information regarding removal and click on **OK**.




Figure 18. SSA-SCSI Driver Installed, Replacing DISK.SYS

17. Figure 19 appears.



Figure 19. Installation Completed

18. Click on **OK**. Remote Systems Management will be installed later as described in 4.6, "Installing RSM" on page 38.

19. When prompted to restart the system, click on **Yes**.

20. After reboot, you should see windows similar to those in Figure 20 in **Control Panel > SCSI**:

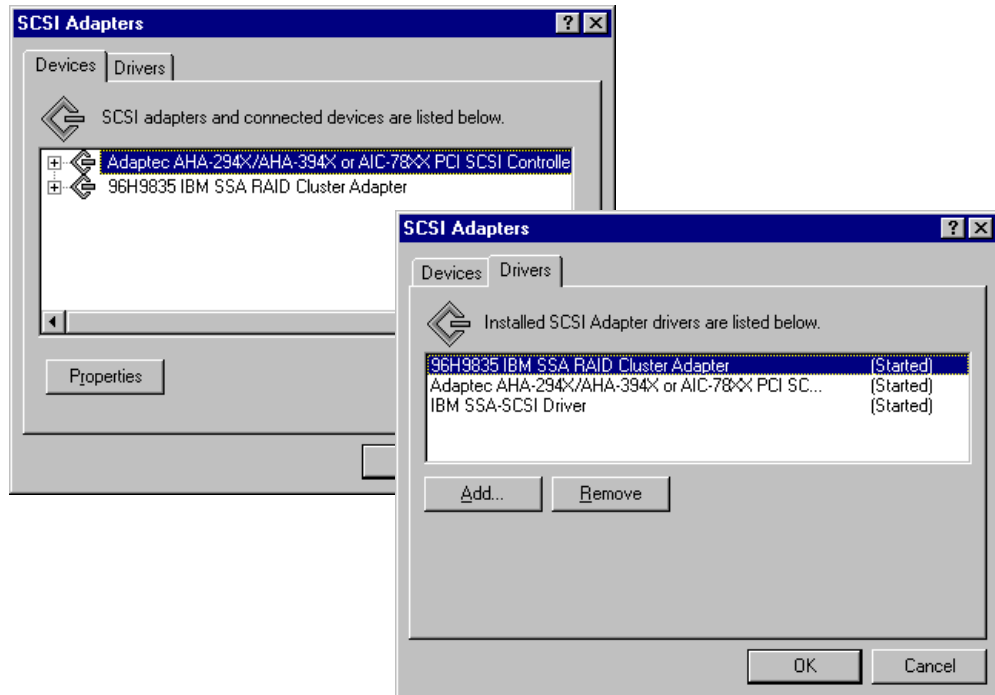


Figure 20. Drivers Successfully Installed

21. Repeat steps 1 to 20 for Node B. In step 9, ensure you select **Set this PC Server to B.**

4.4 Updating the SSA Firmware

Now you need to update the SSA adapter in each server with the firmware you obtained from <http://www.hursley.ibm.com/~ssa/pserver>.

Proceed as follows:

1. Insert the firmware diskette.
2. From within Windows NT, open a command prompt window and change to the C:\ISSA directory where the driver installation placed its files.
3. From the C:\ISSA subdirectory enter the command:

```
ISSAADLD A:\xxxxxxx.L10
```

Where xxxxxxx.L10 (for example, ADAP0200.L10) is the name of the firmware file to be loaded. ISSAADLD is the IBM SSA Adapter firmware downloader.

Note: The file type of the firmware must be L10. If it is anything else, then you have downloaded the wrong firmware.

If the download is successful, you will see messages similar to the following:

```
C:\issa>issaadld a:\adap0200.110
```

```
IBM SSA Adapter Firmware Downloader for Windows, version 1.5
```

```
Loading firmware file: a:\adap0200.110
```

```
Loaded firmware file, CodeID=0200 , length=355172 (0x56b64), LRC=0xcdcf1c56e
```

```
Copyright message :-
```

```
---
```

```
Licensed Internal code - Property of IBM.SSA xxxxxx Adapter Microcode (C) Copyri  
ght IBM Corp, 19xx. All rights reserved US Government Users Restricted Rights -  
Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM  
Corp
```

```
---
```

```
Downloading 0200 to adapter, node=11 ... done
```

```
1 successful download performed
```

Check the messages to ensure the procedure was successful. For example, if you attempted to download SSA RAID Cluster Adapter firmware (LL10) onto an SSA RAID adapter, (LL11) you would get the following messages:

```
Skipping adapter with LL=11, node=11
```

```
0 successful downloads performed
```

4. Shut down Windows NT, power off the server and power it back on again.
5. Repeat the process for Node B.

4.5 Installing the SSA Event Logger

You must install Version 1.10 or later of the SSA Event Logger. You can download the latest version from:

<http://www.hursley.ibm.com/~ssa/pcserver>

To install the Event Logger, follow these steps:

1. Stop the currently installed version. Depending on the version, this involves either closing the application from the task bar or stopping the NT service **PC's SSA Event Logger**. To stop the service, use **Control Panel > Services**, highlight the service and click on **Stop**.
2. Run SETUP.EXE from the directory where you unpacked the ZIP file.
3. Follow the initial installation screens. You will see Figure 21.

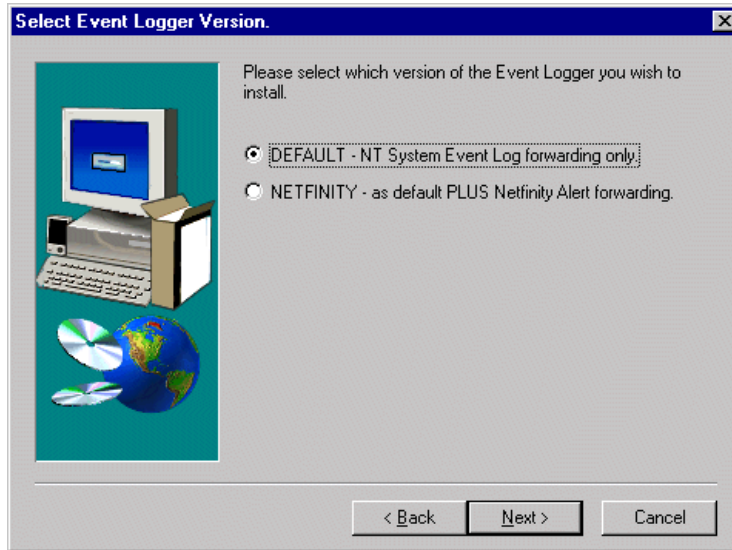


Figure 21. Selecting the Alerting Method

4. From here you can specify how you want the SSA events to be distributed. You can either choose to have them recorded in the Windows NT Event Viewer (the default), or you can have them stored in Event Viewer *and* sent to Netfinity Manager as alerts.

If you have Netfinity Manager installed or plan to do so, select **NETFINITY** and click on **Next**.

5. Follow the remainder of the installation until it is completed. The newly installed Event Logger service will be automatically started.

Note: The event control file C:\ISSA\EVNCTRLF.TXT controls the forwarding of SSA events as alerts to Netfinity Manager. When you select **NETFINITY**, this is automatically enabled. If you select **DEFAULT**, you can enable this later by editing the file. To enable event forwarding to Netfinity Manager, remove the “#” from the start of the line “#Netfinity”. The POPUP section in the file controls which SSA events are raised as Netfinity pop-up alerts. You may want to modify this section as well. See the comments in the file for more information.

4.6 Installing RSM

The recommended way to configure the SSA adapters is to use Remote Systems Manager, RSM. You can download it free of charge from:

<http://www.hursley.ibm.com/~ssa/pcserver>

RSM has the following prerequisites:

- Netscape Navigator 3.x, 4.x or Microsoft Internet Explorer 4.x.
- IBM SSA Event Logger V1.10 or later (downloadable from <http://www.hursley.ibm.com/~ssa/pcserver>). See 4.5, “Installing the SSA Event Logger” on page 37.
- For the Netfinity Manager version, you also need Netfinity Manager V5.0, V5.1 or V5.2 (check <http://www.hursley.ibm.com/~ssa/pcserver> for the latest information about other versions).

Note: Installing Version 1.10 or later of the SSA Event Logger also requires an update of the text-mode configurator. You can download this from <http://www.hursley.ibm.com/~ssa/pcserver> as well.

To install RSM, run SETUP.EXE from the RSM CD-ROM or from the directory where you unpacked the downloaded version. Follow the prompts to install the software. After the initial setup windows, you will see Figure 22:

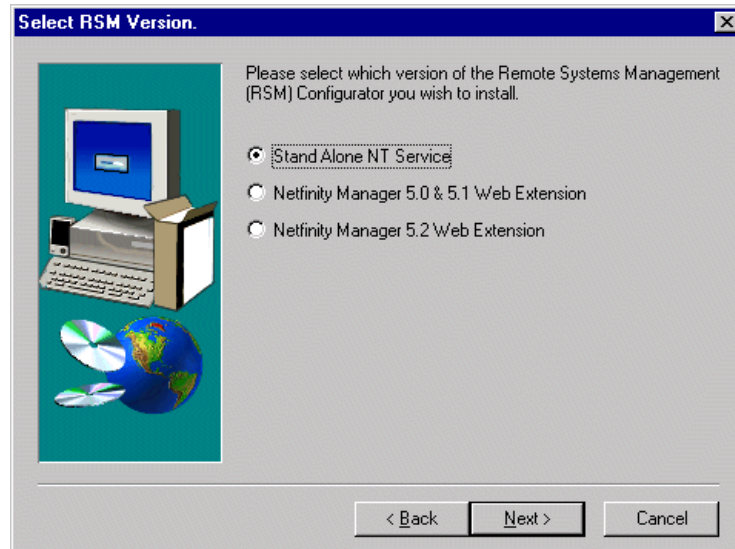


Figure 22. RSM Version Selection

Select the version of RSM you wish to install. See 3.2, “Remote System Management” on page 27 for more information about the versions.

4.6.1 Installing the Stand Alone Version

If you select **Stand Alone NT Service** from Figure 22, the software will be installed and the service will be automatically started. You will see Figure 23 towards the end of the installation. Ensure you set the appropriate levels of access to RSM as per the READMESA.TXT file.




Figure 23. Security Settings

To verify installation, click on **Start > Programs > SSA Tools** which will start a browser session to URL <http://127.0.0.1:511/ssaindex.htm> as shown in Figure 24.

Note: 127.0.0.1 is the loop-back address for the local server. If you are accessing the server from another system in your network, substitute the server's TCP/IP address. If you use a proxy gateway, add the address 127.0.0.1 into the "do not use proxy server for addresses beginning with" proxy configuration field of your Web browser.



Figure 24. RSM Main Window

You can then access RSM by clicking on **SSA RSM Configurator**. You will then be prompted for a user ID and password. These are initially:

- User ID: SSA
- Password: SSA

You should change these as per the guidelines in the READMESA.TXT file. Once you have installed RSM on each server, shut it down and power it off.

Removing the Stand Alone RSM Version

If you plan to uninstall the stand-alone version of RSM, you *must* first run the following command:

```
c:\issa\servutil -rm rsm
```

This will remove the SSA RSM Service from the Registry. You can then run **Control Panel > Add/Remove Programs** to remove the software.

If you do not run the SERVUTIL command, you will receive an error message each time NT reboots, as it attempts to start the RSM service, and removal will fail as the service will be running. See the READMESA.TXT file for more information.

4.6.2 Installing the Netfinity Manager Version

If you wish to integrate RSM with Netfinity Manager, select one of the Netfinity Manager choices from Figure 22 on page 39.

During the process, Netfinity Manager is stopped and restarted. Once installation is complete, start Netfinity Manager by clicking **Start > Programs > Netfinity > Netfinity Service Manager**. Open Web Manager and enable Web access as shown in Figure 25.

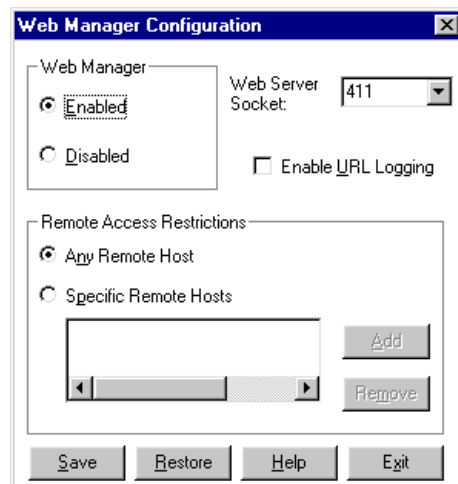


Figure 25. Netfinity Manager — Web Manager Configuration

You can then access Netfinity Manager via the URL <http://ssa-server:411/main>. The Netfinity Manager main window appears as in Figure 26 with the additional SSA icon:

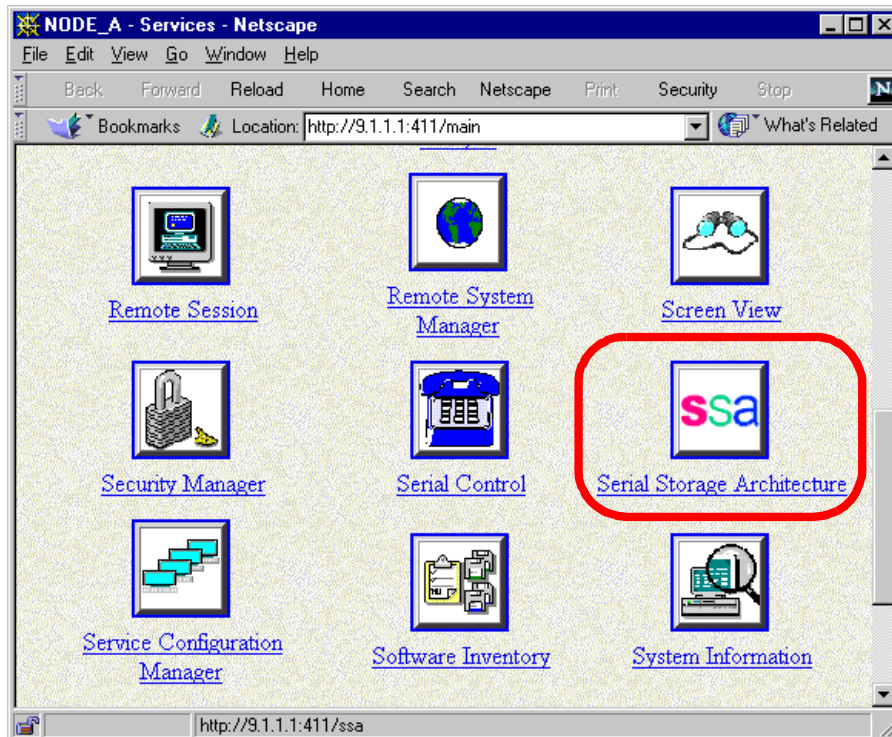


Figure 26. RSM Icon in Netfinity Manager

To run RSM, click on the **Serial Storage Architecture** icon. Alternatively, you can access RSM directly via URL <http://ssa-server:411/ssa>.

Once you have installed RSM on each server, shut it down and power it off.

4.7 Next Steps

Both servers should now be powered off. The next steps are to connect the SSA enclosures to the server and configure the arrays. These are described in Chapter 5, “Configuring SSA Hardware” on page 43.

Chapter 5. Configuring SSA Hardware

Now that Windows NT Server Enterprise Edition and the SSA driver and tools are installed, the next step is to configure the disks in the SSA external storage enclosures.

We will do the following:

1. Connect the SSA enclosures.
2. Format the SSA disks.
3. Define the RAID arrays.
4. Define the hot spares.
5. Prepare the arrays for use with Windows NT.
6. Format the partitions in Windows NT.

In Chapter 6, "Installing Microsoft Cluster Server" on page 61, we cover MSCS installation.

5.1 Connect the Enclosure

At this point, both servers have the following installed:

- The latest BIOS for your server
- Windows NT Server Enterprise Edition
- Service Pack 3
- SSA device driver
- SSA Remote Systems Management (RSM)
- The latest firmware on each SSA RAID Cluster Adapter

Both servers are powered off.

Prior to defining the RAID arrays, you now need to connect the SSA disk enclosures to the two servers and power on Node A:

1. Both Node A and Node B are powered off.
2. Connect the SSA storage devices to the servers via SSA cables.
Refer to the SSA external enclosure documentation for information about how to connect the SSA cables.
3. With both nodes connected to the enclosures, power on the enclosures.
4. Power on Node A.
5. Log on to Windows NT on Node A when you are prompted to do so.
6. Power on Node B. Press the Spacebar when you see the message:

OS Loader - Please select the operating system to start

This will halt the loading of Windows NT on Node B.

Warning: Do not let Node B progress beyond this screen. Otherwise, you will have to reboot.

5.2 Formatting the SSA Disks

At this point, Node A is running and logged on as an administrator. Node B is powered on, but it remains at the initial OS Loader screen. Proceed to work from Node A.

For each of the disk drives in the SSA enclosures, you should mark all disks as Free Resources and perform low-level formats on each one. Follow the steps below. Your configuration may be different to the one here (eight hard disks in a single 7133 unit) but the process is still the same.

Using RSM

These instructions use the text-mode configurator. If you prefer, you may use the browser-based RSM configurator. The steps remain the same, however.

1. Start the SSA Configurator by running C:\ISSA\ISSACFG.EXE.
2. You will then see the SSA Configurator program main menu as shown in Figure 27.

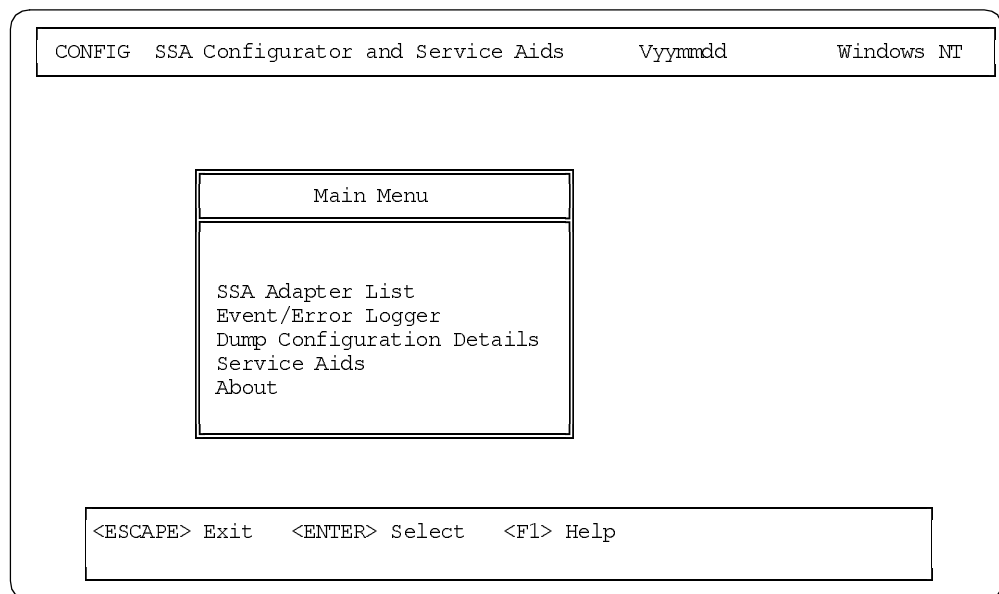


Figure 27. SSA Configurator Main Menu

3. Select **SSA Adapter List** and press Enter. You will see a screen similar to Figure 28.

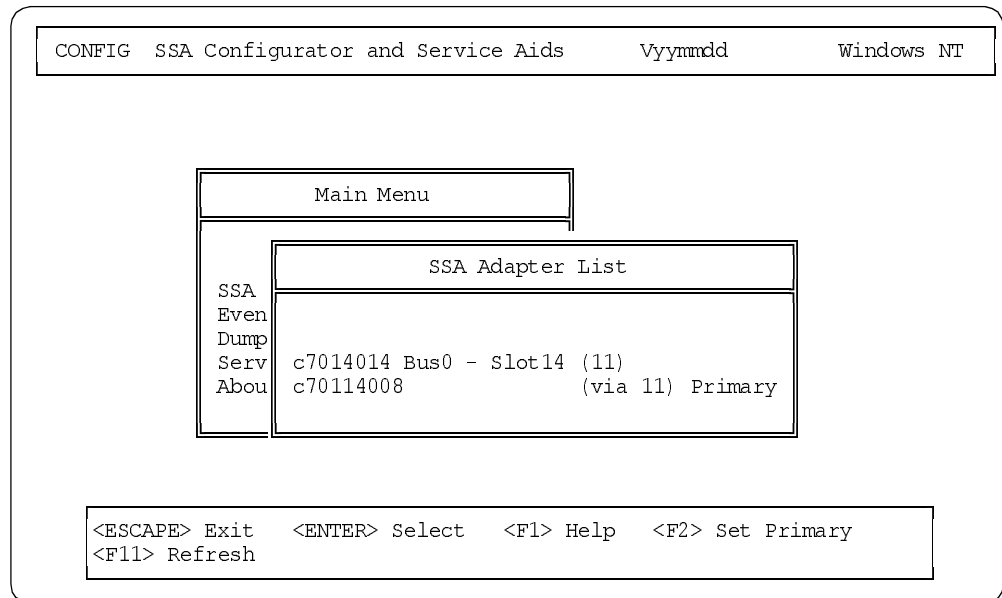


Figure 28. SSA Adapter List Window

4. You can see that both adapters (one on each node) have been discovered. Only the primary adapter (the second one in Figure 28) has the ability to configure the RAID-1 arrays.

Tip: If the local adapter is not the primary (which is the case in Figure 28), you can highlight the local adapter (the first adapter in the list in our case) and press F2 to set it as the primary.

Highlight the primary adapter and press Enter. Figure 29 appears.

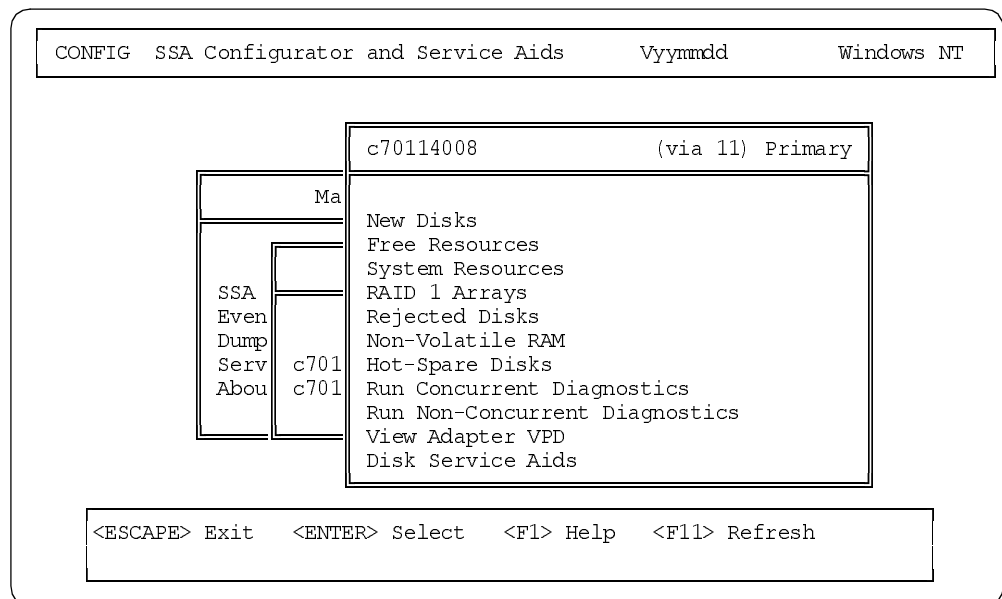


Figure 29. SSA Primary Adapter Options Window

5. Before configuring the new RAID-1 arrays, you must first set all of the disk drives to Free Resources. From Figure 29, select each of the following menu entries and delete all drives that appear in them:

- New Disks
- System Resources
- RAID 1 Arrays
- Rejected Disks
- Non-Volatile RAM
- Hot-spare Disks

This will mark all those disks as Free Resources.

6. Once you have selected each item and deleted all of the disks, then select **Free Resources** in Figure 29 and press Enter.
7. You need to make sure that all of your SSA disks (eight drives in this example configuration) are listed as shown in Figure 30.

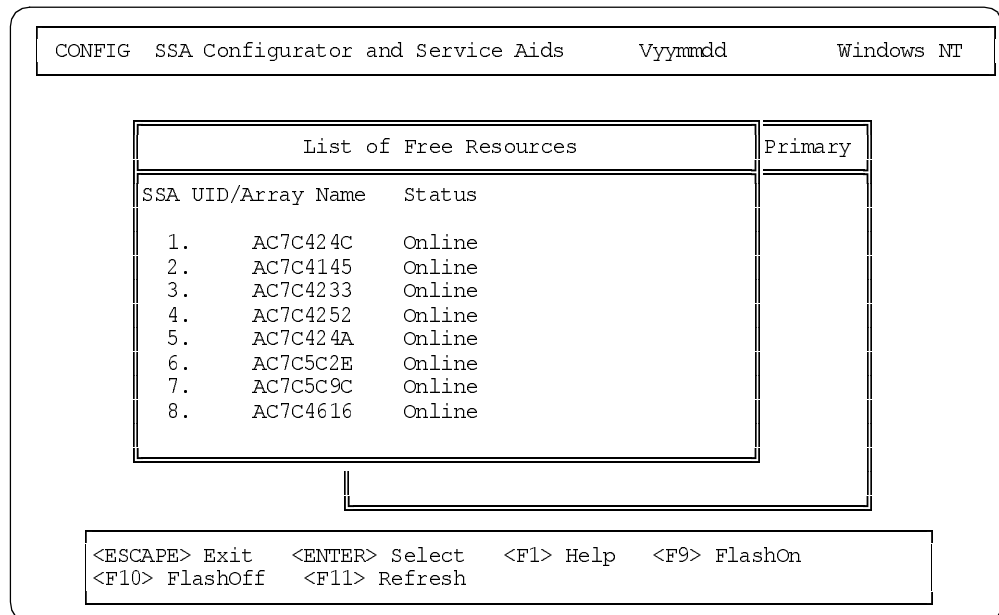


Figure 30. SSA Free Resources

8. Press Esc to return to primary adapter list options. You will see the adapter options window again (Figure 29 on page 45).
9. Perform a low-level format on each disk drive. Although this step is not mandatory, it is recommended so as to build your cluster on disks in a known state (that is, formatted). Select **Disk Service Aids** and press Enter. You will see a screen similar to Figure 31.

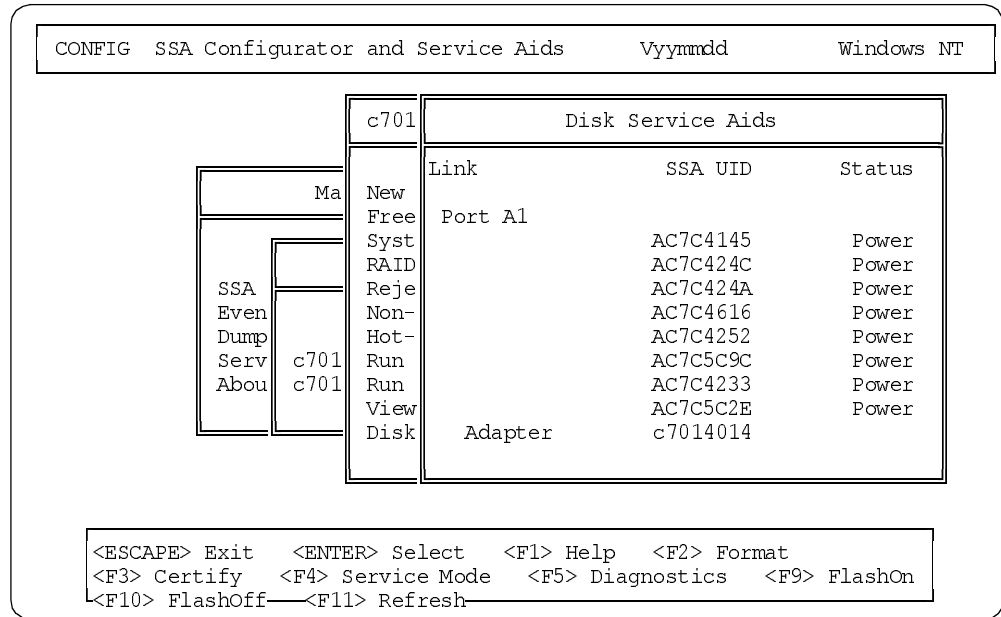


Figure 31. Disk Service Aids Window

10. Select one drive at a time and press the F2 key. You will be asked to confirm the format operation. Highlight **Yes** and press Enter.

11. After all drives have been selected to be formatted, you will see Figure 32.

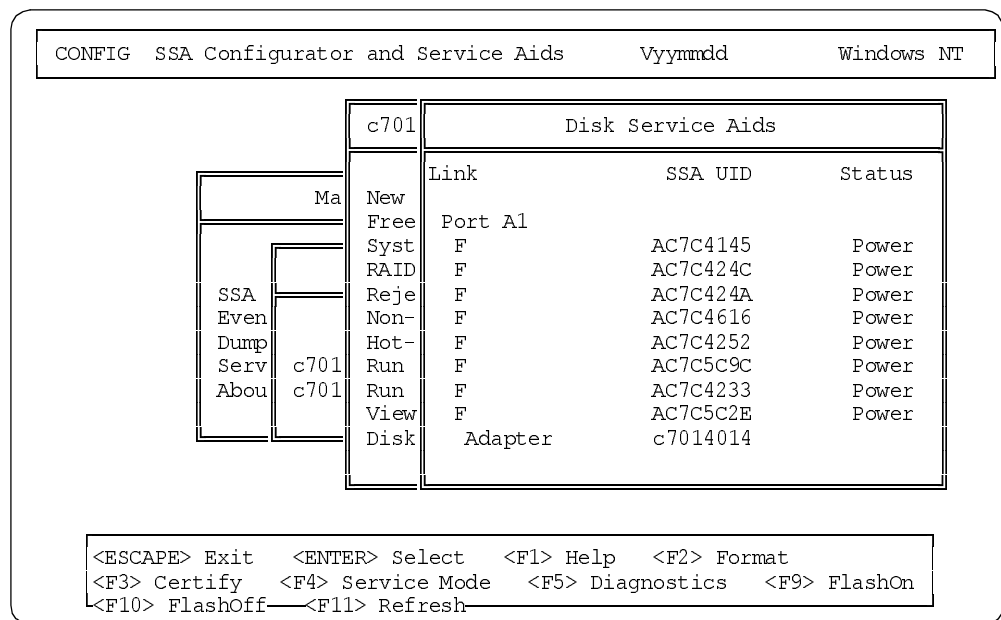


Figure 32. Disk Service Aids Window

The letter F displayed in front of each drive indicates that it is being formatted. To see the status of the formatting, press Esc to get back to the main menu. When you are at the main menu, scroll up, select **Free Resources** and press Enter. You will see a screen similar to Figure 33.

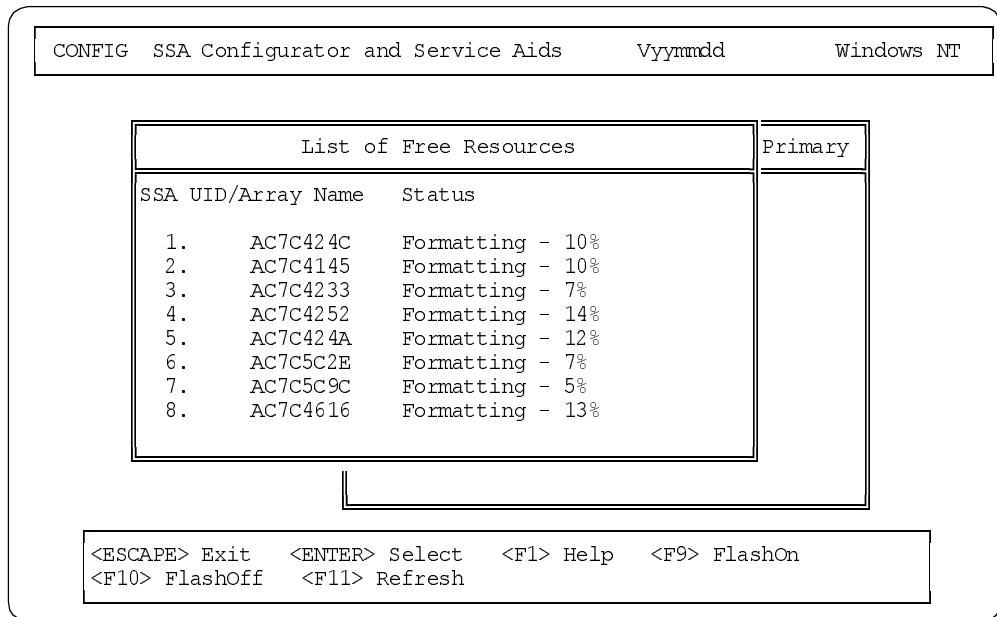


Figure 33. Low-Level Format in Progress

12. You can check the progress of the format process from this panel. Press F11 to refresh the screen. Once the formatting is complete the status of each drive will change to Online.
13. Press Esc to return to the primary adapter option window (Figure 29 on page 45).

5.3 Defining the Arrays

You are still working on Node A. Node B is powered on but still at the OS Loader screen. Now you need to create the RAID-1 arrays you require for your installation.

1. From the primary adapter option window (Figure 29 on page 45) highlight **RAID 1 Array Resource** and press Enter. A screen similar to Figure 34 appears.

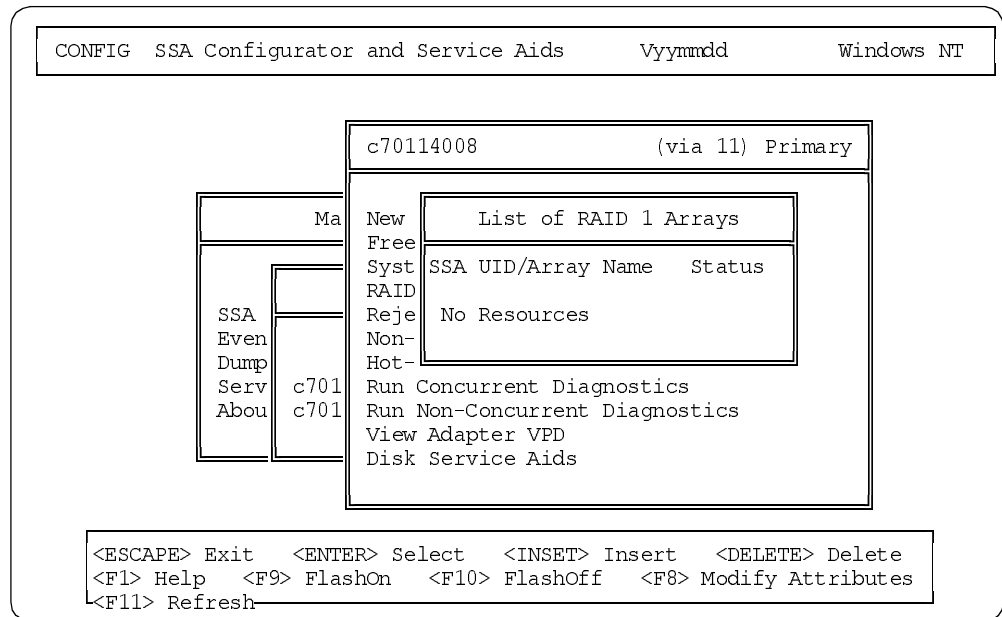


Figure 34. List of RAID 1 Arrays

2. There are no resources listed in this window yet. Press the Insert key to create a new array. You will see a screen similar to Figure 35 prompting you to specify the attributes of the new array:

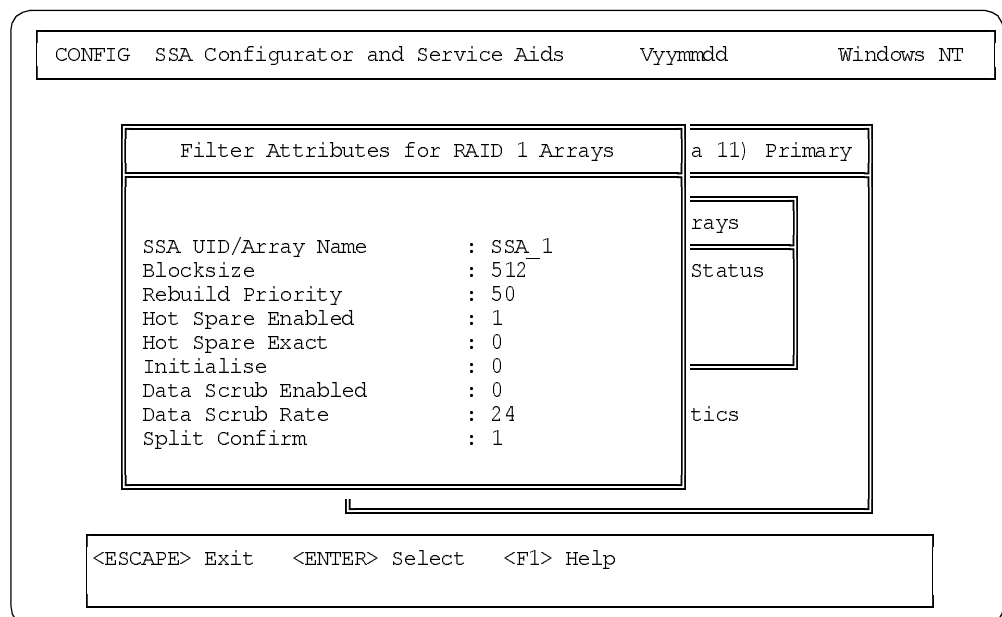


Figure 35. Specifying Attributes of the New Array

3. Type an array name in the SSA UID/Array Name field. It is wise to use the name of the Windows NT volume which will be placed on this array as the array name. This makes it easier for you to compare the physical SSA layer with the operating system's view. We used SSA_1.
4. Press Enter to scroll down through the options and make the changes needed regarding your production environment. When you press Enter at the last field,

you will be prompted to select which disks you want to be part of this array (Figure 36):

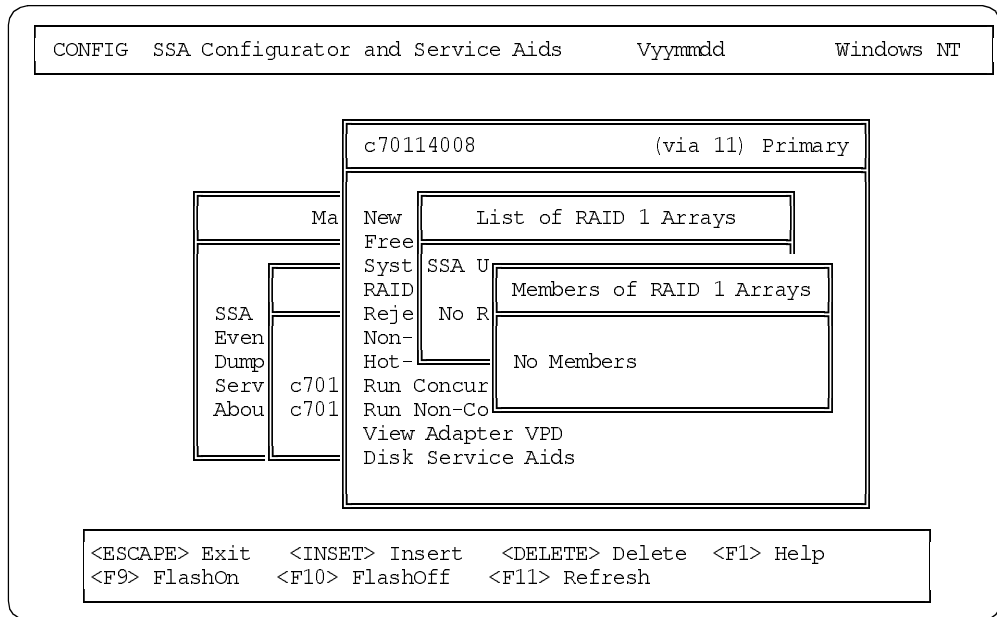


Figure 36. Members of RAID 1 Array

5. This array is new and empty so the drives must be added. Press the Insert key and select from the free resources the drives you want for your array. Figure 37 shows the disks that are candidates for inclusion in the array.

Note: If your drives are not listed, ensure they are in the Free Resources list as explained in step 5 on page 45.

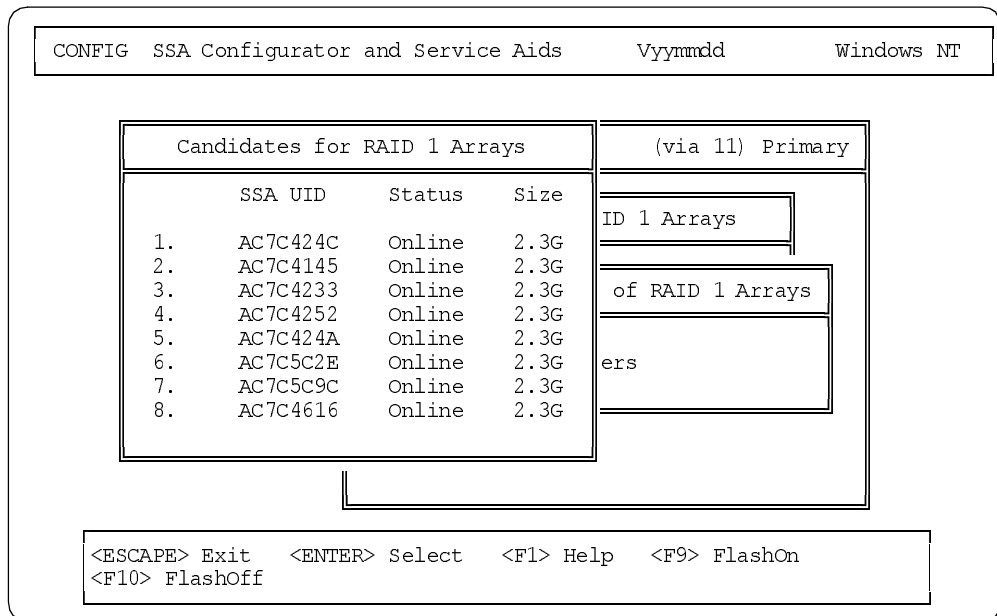


Figure 37. Disks Eligible for Inclusion in the New Array

6. Highlight the drive you want and press Enter to add it to the array. You will see a screen similar to Figure 38.

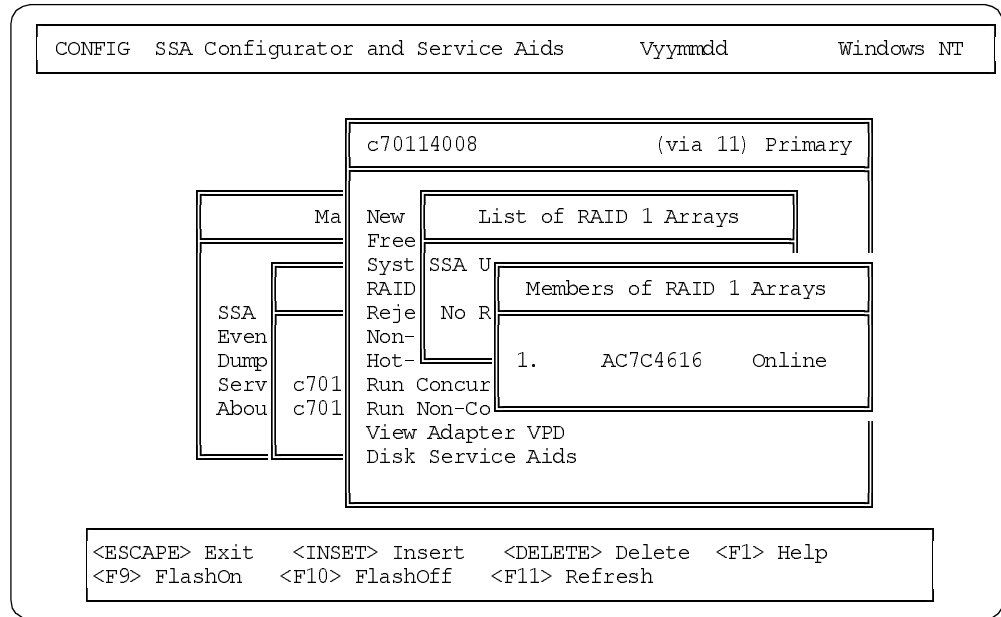


Figure 38. Members of the Array

7. You can check in this window that the drive you chose has been included in the array and is online. Press Insert again to add the second disk.

8. Press Esc and click **Yes** at the confirmation window. You will see a screen similar to Figure 39.

Tip: You can verify which drives are in the array by pressing the F9 key. This will cause the LEDs on all array members to blink (FlashOn); thus, you can easily see them in the 7133. Pressing the F10 key stops the flashing (FlashOff).

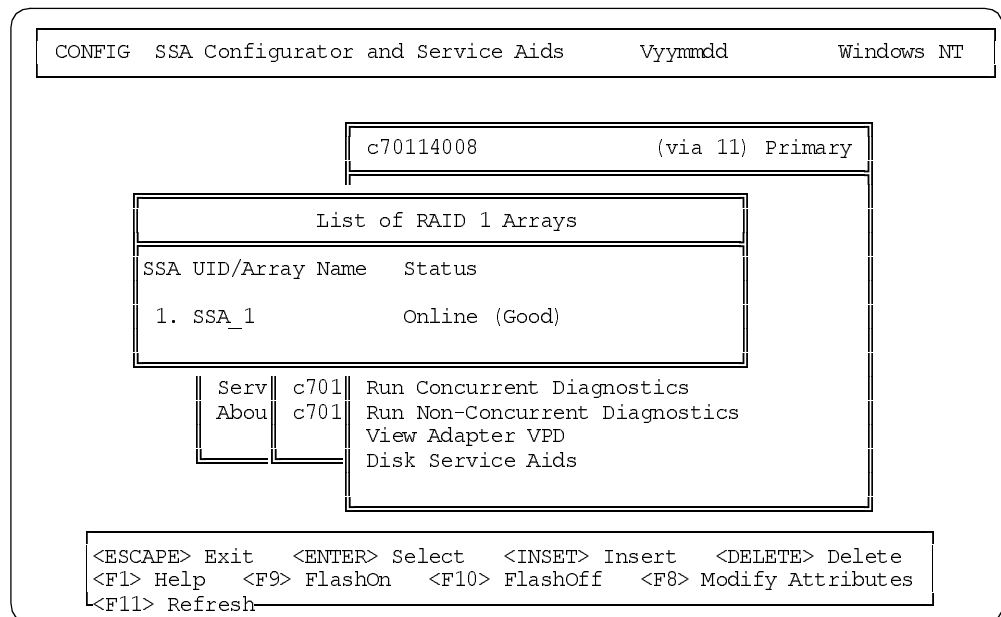


Figure 39. List of RAID 1 Arrays

- The two drives have been added to the array and the array is now online. If you need to define other RAID-1 arrays, press Insert and repeat the operations described above. We will create three RAID-1 arrays, which results in Figure 40.

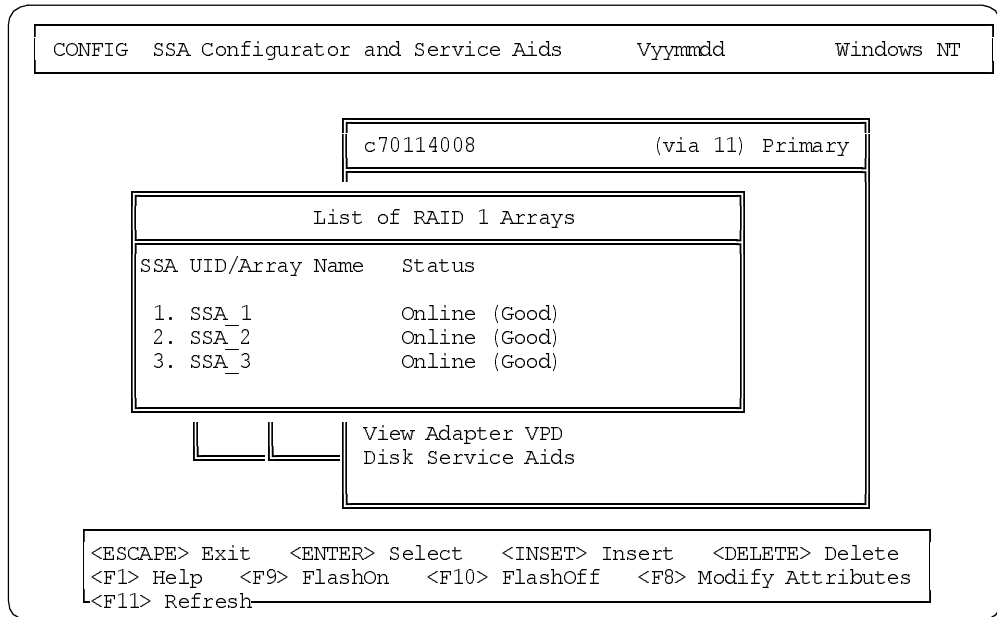


Figure 40. List of RAID 1 Arrays

- When all arrays are defined and they appear as Online <Good>, you can press Esc to return to the primary adapter option window and define hot spares if needed.

5.4 Defining Hot Spares

- Select the item **Hot-spare Disks** and press Enter. You will see a screen similar to Figure 41.

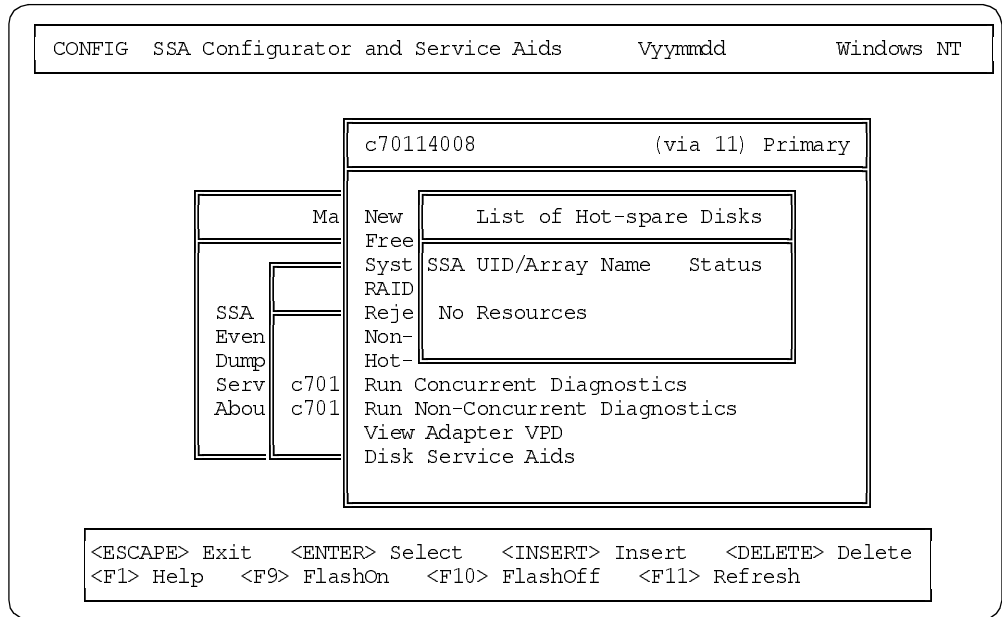


Figure 41. List of Hot Spare Disks

- The list is empty. Press the Insert key to add hot-spare drives. You will see a screen similar to Figure 42 prompting you to enter the block size for the hot spare disk.

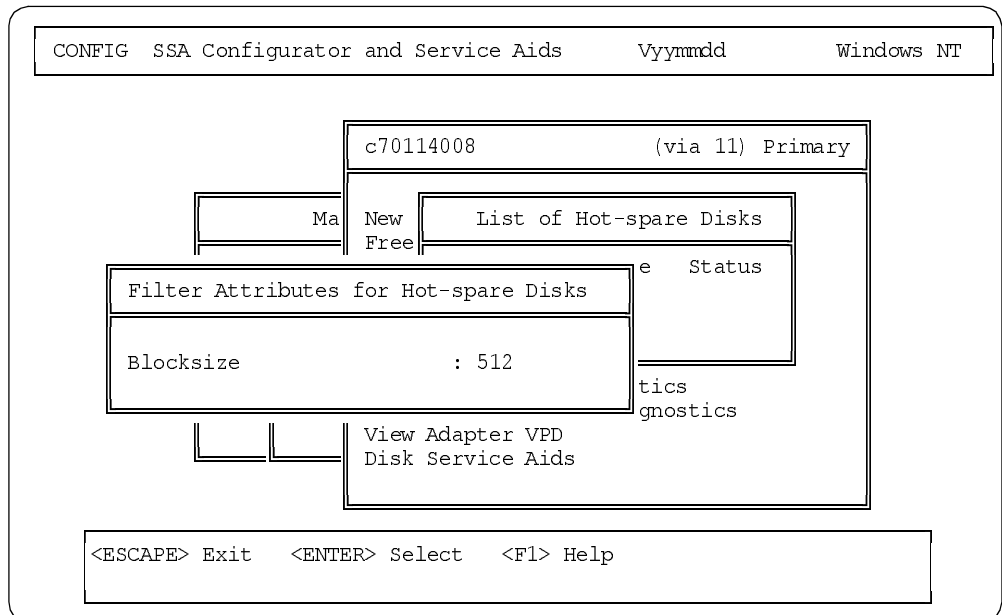


Figure 42. Blocksize for the Hot Spare Disks

- You are prompted to specify the blocksize you want to use for the hot-spare drive. Use the same value you specified in the definition of the arrays (Figure 35 on page 49). We shall keep the default and press Enter. Figure 43 appears:

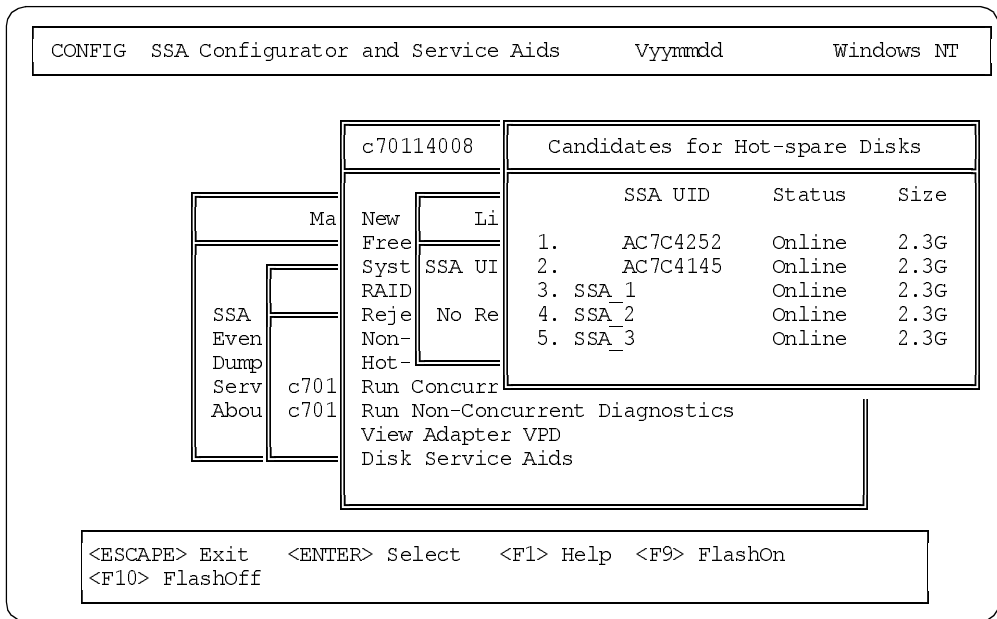


Figure 43. Candidates for Hot-Spare Disks

4. Highlight the disk you want as hot-spare drive and press Enter.

Note: Do not select a RAID-1 array. The arrays are not valid selections for hot spares. If you need more than one hot-spare drive, repeat this operation.

We selected two hot-spare drives, as shown in Figure 44:

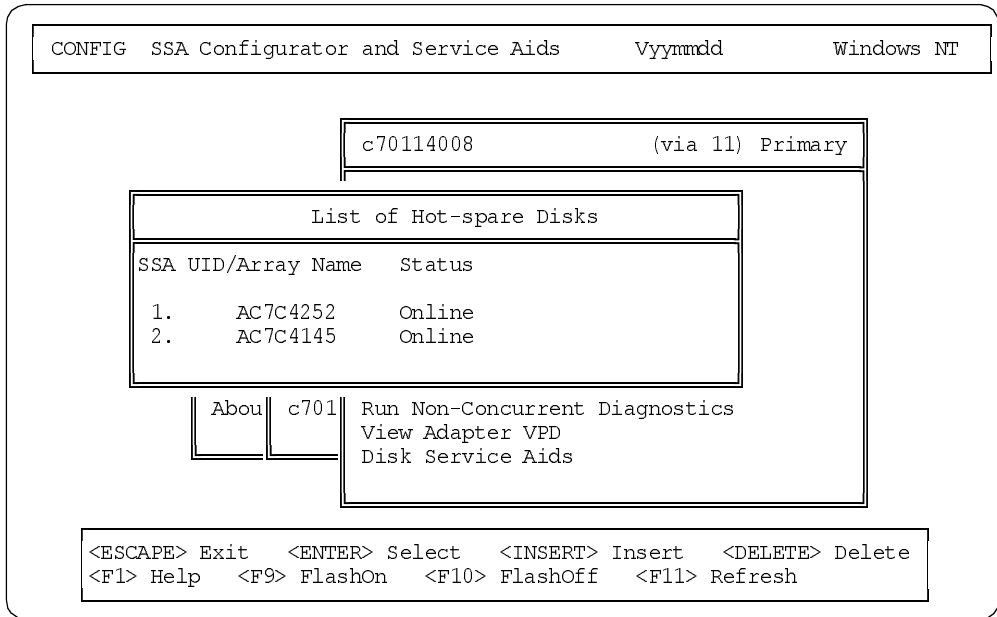


Figure 44. List of Defined Hot-Spare Disks

5. Press Esc to return to the primary disk main menu.

5.5 Setting the Arrays as System Resources

The last step is to mark the newly defined RAID-1 arrays as system resources. Before the arrays are recognized by Windows NT as “physical disks”, they must be marked as system resources.

1. From the primary adapter options window (Figure 29 on page 45), highlight **System Resource** and press Enter. You will see a screen similar to Figure 45.

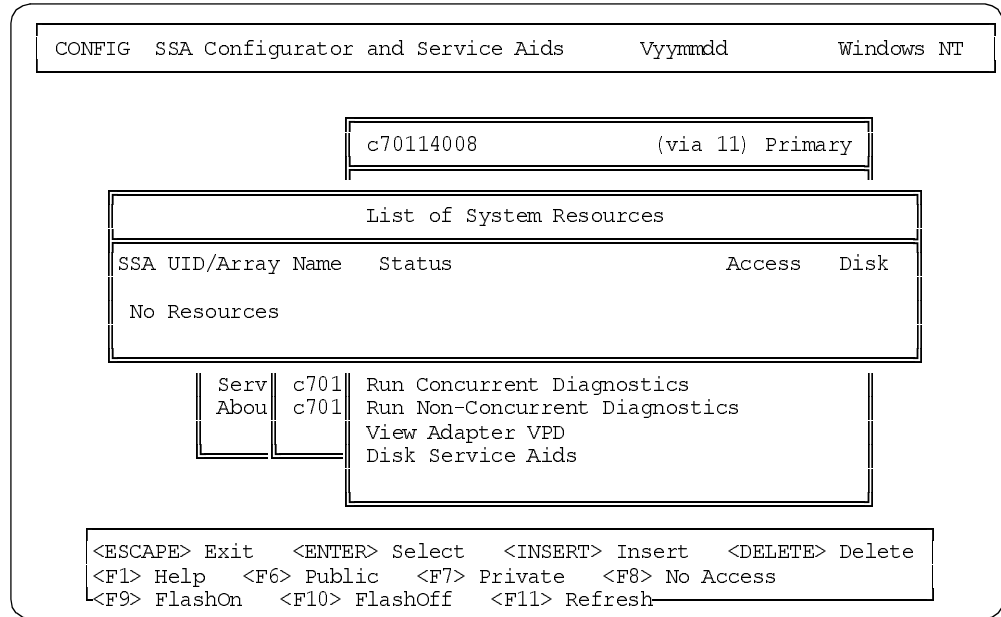


Figure 45. List of System Resources

2. At this time there are no resources defined as system resources. Press the Insert key to add a new resource. You will see a screen similar to Figure 46.

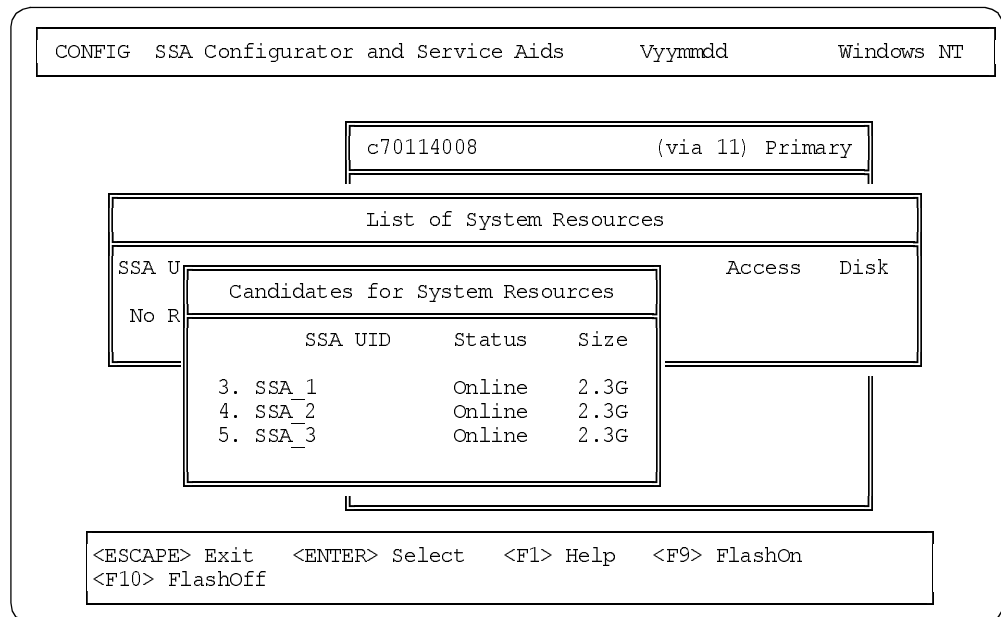


Figure 46. Candidates for System Resources

- The window Candidates for System Resources displays the RAID-1 arrays defined in 5.3, “Defining the Arrays” on page 48, which can be used as system resources. Select an array and press Enter. You will see a screen similar to Figure 47.

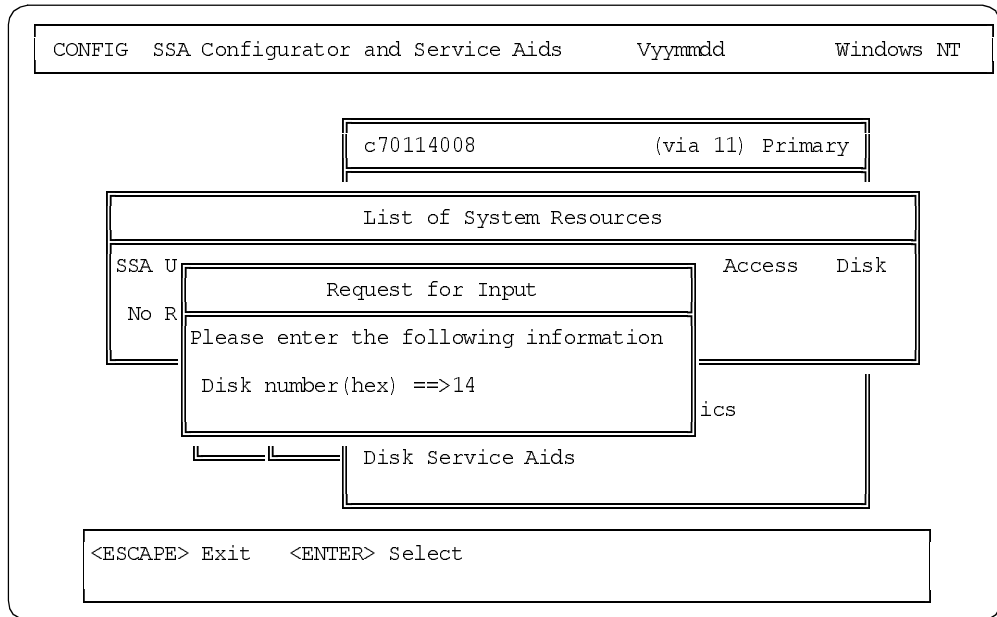


Figure 47. Disk Number

- You are prompted to give a disk number in hexadecimal for this new logical drive. Accept the default and press Enter.

For more information on the disk number, see 2.5.6, “Disk Types and Disk Numbers” on page 22.

- Select each of the other two arrays and also add them to the list of system resources. This results in Figure 48:

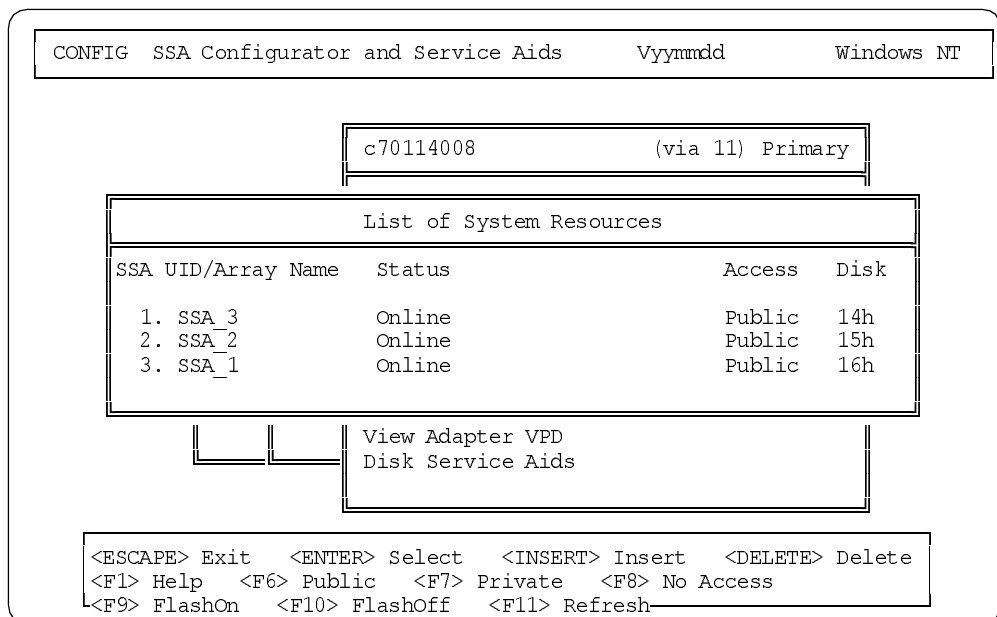


Figure 48. List of System Resources after Resource Inclusion

6. The three RAID-1 arrays have now been included in the system resources. The SSA drives configuration is now complete. You can exit the SSA Configurator program by pressing Esc four times and answering **Yes** on the confirmation panel.

The next step is to configure the drives in Windows NT's Disk Administrator.

5.6 Creating Partitions in Windows NT

You are still working on Node A. Node B is powered on but held at the OS Loader window. Now that the SSA arrays are defined, the next step is to create operating system partitions and assign drive letters with Windows NT Disk Administrator.

1. Start Windows NT Disk Administrator on Node A. Figure 49 appears.




Figure 49. Disk Administrator: New Disks Found

2. Click **OK** to continue. You will then see Figure 50:

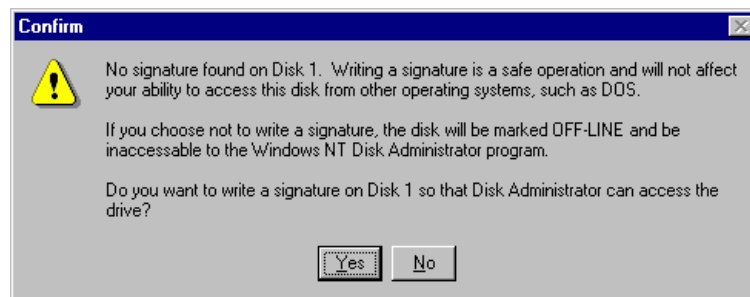


Figure 50. Disk Administrator: No Signature Found

3. Click on **Yes**. You will see one of these windows for each of the system resources you added in 5.5, "Setting the Arrays as System Resources" on page 55.

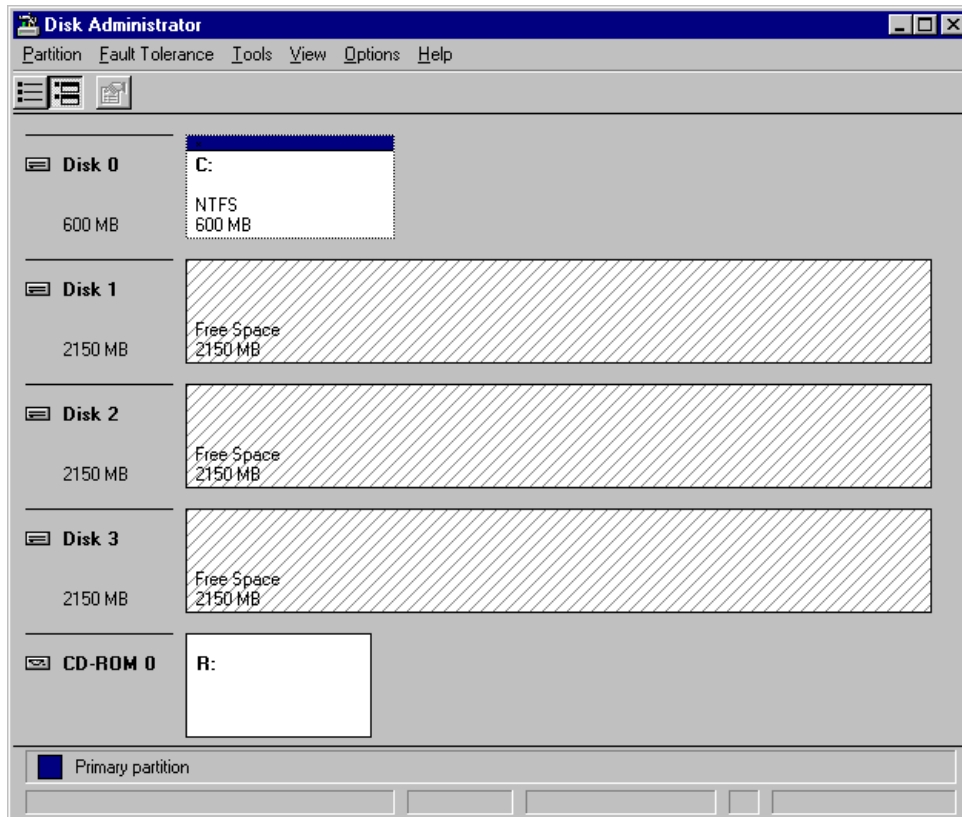


Figure 51. Windows NT Disk Administrator - New Drives

4. Change the letter of your CD-ROM drive (R, for example).

Note: If you are sure more drive letters will be required in your configuration, you should assign the letter Z to your CD-ROM.
5. Right-click on Disk 1 drive space and select **Create**.
6. Accept the default size proposed in the resulting window and click **OK** to continue.
7. Repeat the steps above for all new disks (each SSA system resource).
8. Right-click on the D: drive and click **Commit Changes Now** in the resulting pop-up.
9. Choose **Yes** on the pop-up Confirm menu. The partitions have now been created.
10. Right-click on each new SSA partition and format it as NTFS. Microsoft Cluster Server only supports NTFS partitions.
11. Enter a disk label. We labeled D as Quorum, E and F as Disk_E and Disk_F.

We highly recommend you use labels containing the drive letters. With MSCS, you must always ensure that the drive letters of shared disks are the same on both cluster nodes. The label is the only NTFS volume structure information that is written to the partition itself and can be easily read. (All other information kept in the registry or the disk signature can be obtained only with special tools.) This will be important if the drive letters become exchanged for some reason.

12. Check the **Quick Format** box (a quick format is sufficient) and select **Start**. Your format settings will look similar to Figure 52.

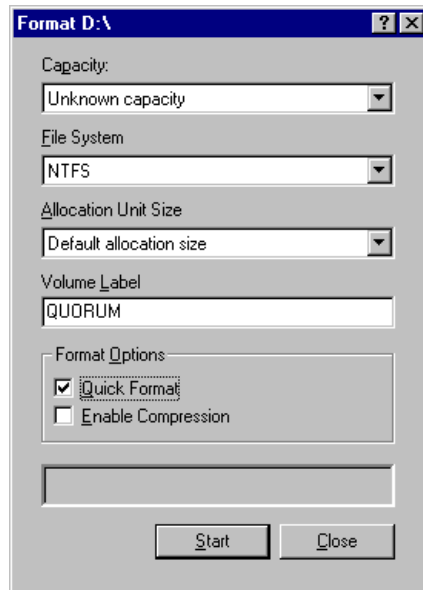


Figure 52. Windows NT Disk Administrator - Drives Format

At this point, the drives have been formatted, lettered and labeled on Node A.

5.6.1 Working with Node B

Now that the SSA RAID arrays have been defined and partitions formatted from Node A, you need to format the partitions on Node B as well. You do not need to define the RAID arrays or allocate SSA system resources again since these were done on Node A and do not need to be redone on Node B.

Keep Node A running (and logged on as an administrator). Work from Node B:

1. On Node B, press Enter to let the boot process resume.
2. Log on to Windows NT as an administrator.
3. Start Windows NT Disk Administrator.
4. Set the CD-ROM to R: or whatever drive letter you set the CD-ROM to on Node A.
5. Make sure that all of the settings for all of the drives (size, letters, labels) match what you have just defined on Node A in 5.6, "Creating Partitions in Windows NT" on page 57.

This concludes the installation of Windows NT Server Enterprise Edition and the configuration of the SSA subsystem. Leave both nodes running and logged on.

You can now install Microsoft Cluster Server and Service Pack 4. Refer to Chapter 6, "Installing Microsoft Cluster Server" on page 61 for the installation of Microsoft Cluster Server.

Chapter 6. Installing Microsoft Cluster Server

Now that Windows NT Server Enterprise Edition is installed and the SSA disk subsystem is configured, you should now install Microsoft Cluster Server.

We assume that the network connection with the TCP/IP protocol between the two servers has been established successfully.

Note: You should also refer to the *MSCS Administrator's Guide* for information about installing MSCS.

The installation process is as follows:

1. Install MSCS software on the first node.
2. Install MSCS software on the second node.
3. Start Cluster Administrator to see if installation was successful.
4. Install Cluster Administrator on a remote system.
5. Update the MSCS software.

6.1 The Installation Process

You have just assigned drive letters to the external shared disks and are ready to install the MSCS software on the machines. Both servers are running Windows NT Server Enterprise Edition and are logged on as an administrator.

6.1.1 Installing MSCS on the First Node

This procedure shows how to install MSCS on the first node.

If you are installing MSCS in a configuration where one node will be the primary domain controller (PDC), install MSCS on the PDC first. MSCS requires the PDC for authentication during the installation process.

Domains

1. The two servers must be members of the same domain.
2. A server can be a member of only one cluster.
3. The following configurations of servers in a cluster are possible:
 - A primary domain controller and a backup domain controller
 - Two backup domain controllers
 - Two stand-alone servers

You should also ensure that the configured gateway and domain name server (or WINS server) are available. If not, some steps in the installation may take a number of minutes to complete while MSCS completes name resolution.

The steps to install MSCS on the first node are as follows:

1. Insert the Windows NT 4.0 Enterprise Edition CD 2 into the CD-ROM drive on the Node A.
2. Run **Start > Programs > Administration Tools (Common) > Enterprise Edition Installer** or execute NHLOADER.EXE from the Run window or from the command line.

3. Click on **Continue**. Figure 53 appears.

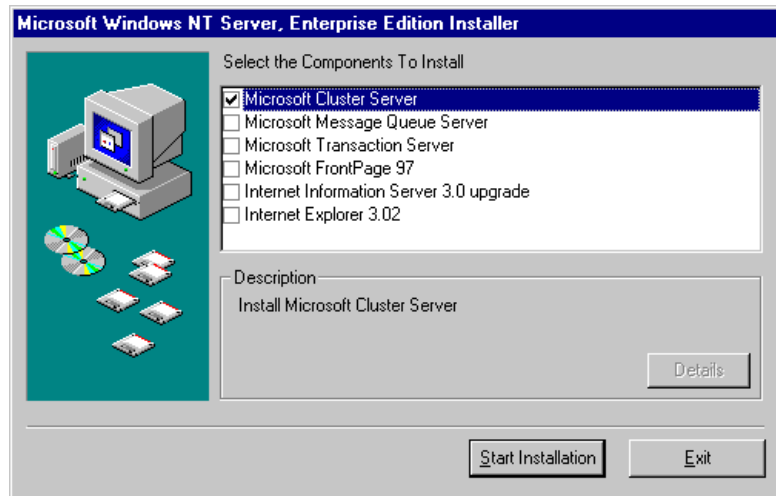


Figure 53. Enterprise Edition Installer

4. From the Enterprise Edition Installer window, check **Microsoft Cluster Server** and then click on the **Start Installation** button.
5. You are then prompted for the location of the files on Enterprise Edition CD 2. Adjust the drive letter of the CD-ROM drive then click on **OK**.
6. The Microsoft Cluster Server Setup window is displayed. Select **Next**. Figure 54 appears.

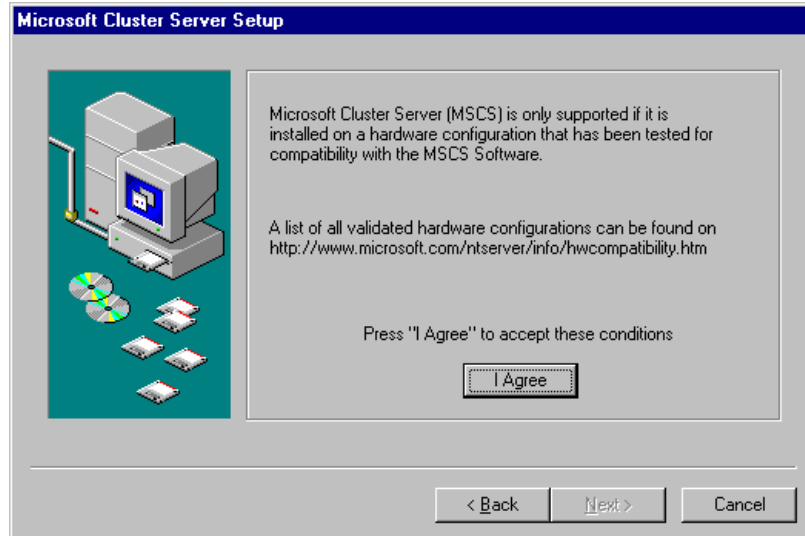


Figure 54. Only Certified Hardware is Supported

7. Click on **I Agree** and then **Next**. Figure 55 appears.

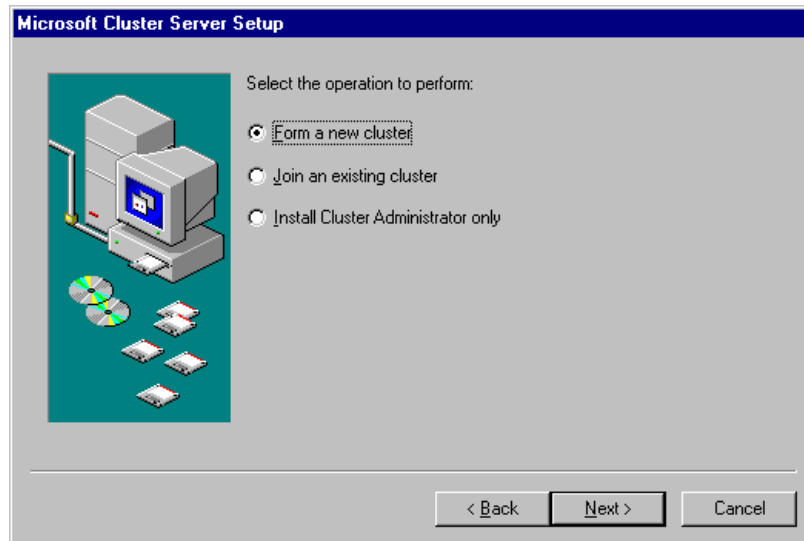


Figure 55. Form a New Cluster

8. Select **Form a new cluster**, then click on **Next**. Figure 56 appears.

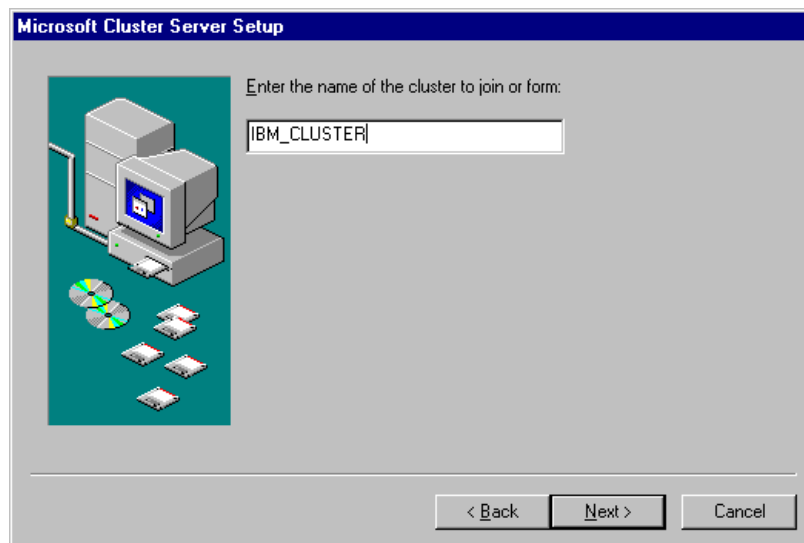


Figure 56. Specify the Cluster Name

9. Type any name you choose for your new cluster in the text box. In this example, we use **IBM_CLUSTER**. Click on **Next**.

10. You will see a panel where you have to define the directory in which the cluster files will be stored. This directory must be on a local drive. Accept the default: **C:\WINNT\CLUSTER**. Select **Next**. Figure 57 appears.

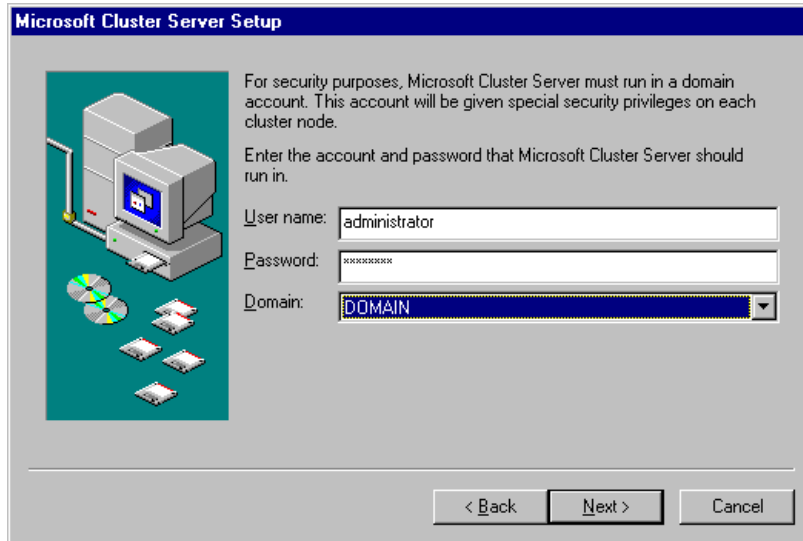


Figure 57. MSCS Logon Screen

11. You are asked to fill in a User name and Password. The User Name you choose does not have to be Administrator, but should have administrator privileges, and the domain is the domain to which you assigned the server when you installed Windows NT. Type in the user ID and password then click on **Next**.

Passwords

The password is only requested once. Be sure you enter the right characters. If you type a character other than what you intended, you will need to cancel the installation and start again.

Figure 58 appears.

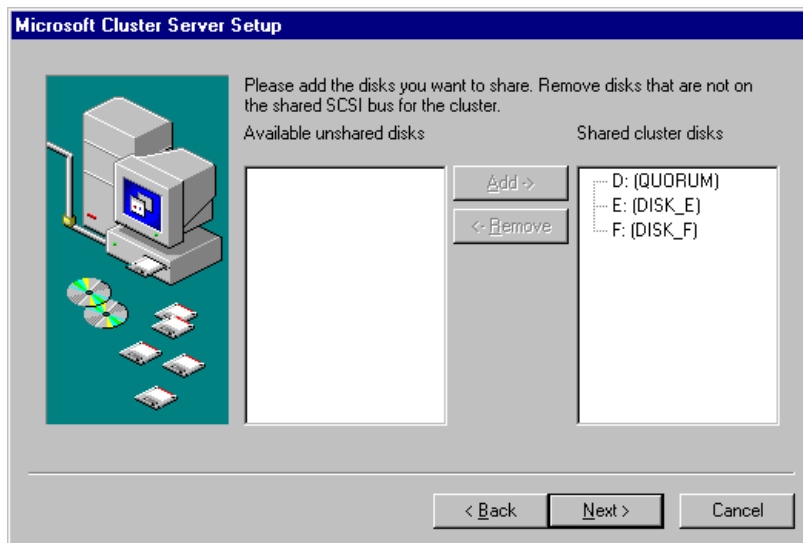


Figure 58. Shared Cluster Disks for SSA

12. Here you specify which drives will be shared between the cluster nodes. Click on **Next** to proceed when the selections are what you want in your configuration. Figure 59 appears.




Figure 59. Quorum Disk for SSA

13. Select the shared disk where the permanent cluster files will be stored. We chose to store these quorum files on D: (with the label QUORUM). Select the appropriate drive and click on **Next** to proceed. Figure 60 appears.

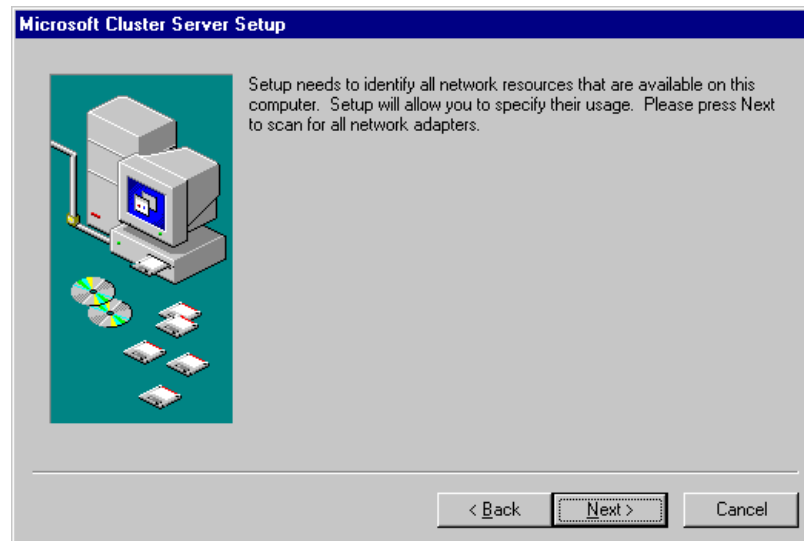


Figure 60. About to Scan for Network Adapters

14. Click on **Next** to scan for all network adapters. Figure 61 appears.

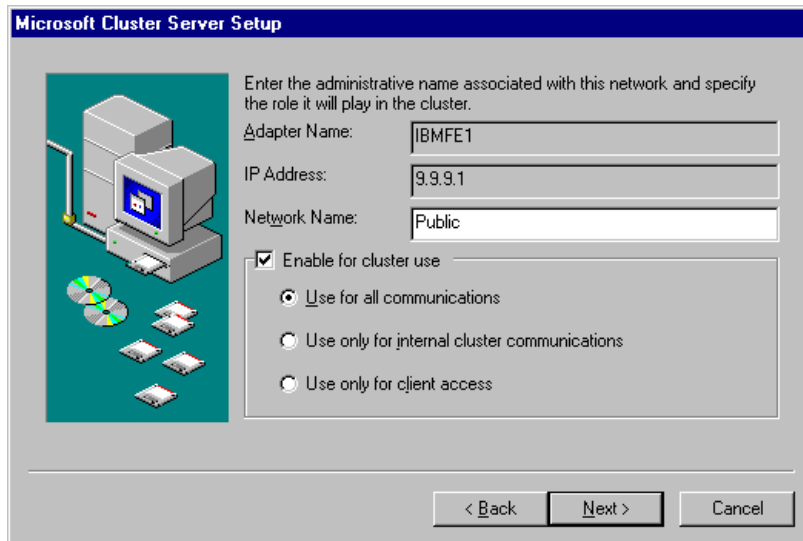


Figure 61. Public Network Adapter Definition

As discussed in 2.2, “Networking” on page 15, in each machine, you should have at least two network adapter cards installed. You must define which one to use for cluster communications and the other(s) for client access. Your choices are:

- Use for all communications (Cluster and Client) -- The adapter will be used for internal communications between the cluster nodes and for client access.
- Use only for internal cluster communications (Cluster only) -- Information between the cluster nodes will be interchanged on this adapter. Client access is achieved through the other network adapters.
- Use only for client access -- Clients are allowed access to the cluster through this adapter. Internal cluster information will be interchanged on another adapter.

In this example, two 10/100 Ethernet adapters were used on each node. The adapter used for client access was configured for “all communications”. The adapter used for the cluster heartbeat was configured for “internal cluster communications only”.

15. Figure 61 shows the first adapter found on the machine. The Ethernet adapter in our configuration is shown as IBMFE1. You may enter the network name you wish to use to identify the adapter/IP address combination.

Type the network name you want, for example, Public, mark the check box **Enable for cluster use**, then select one of the options describing how the adapter should be used. In our example, select **Use for all communications**. Select **Next**. Figure 62 appears.

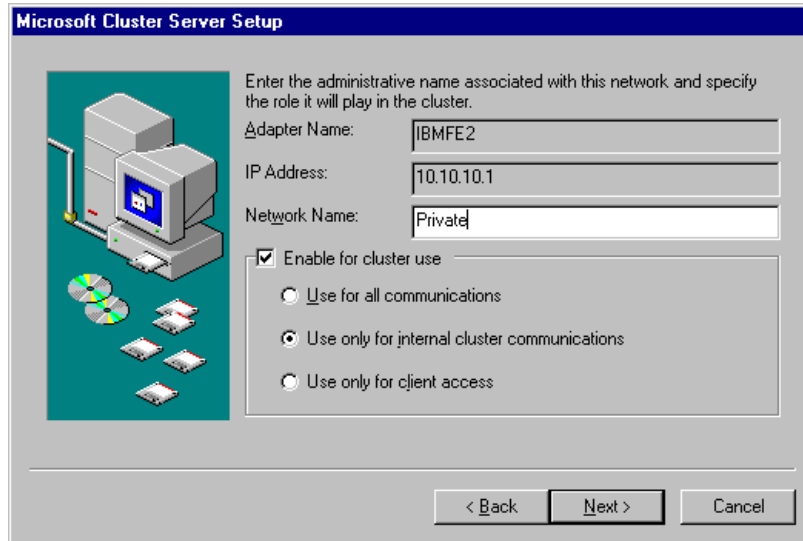


Figure 62. Private Network Adapter Definition

16. When more than one network card is configured as with this example configuration, additional windows for the remaining NICs will require configuration. Enter the Network Name for the second adapter, such as Private. Mark the check box **Enable for cluster use** and select **Use only for internal cluster communications**.

Note: If you only have one adapter selected for internal cluster communications and no adapters for all communications, then you will see a warning message:

Only a single adapter is configured for internal use. If you have multiple adapters, you may reconfigure them to avoid a single point of failure.

Select **Next**. Figure 63 appears.

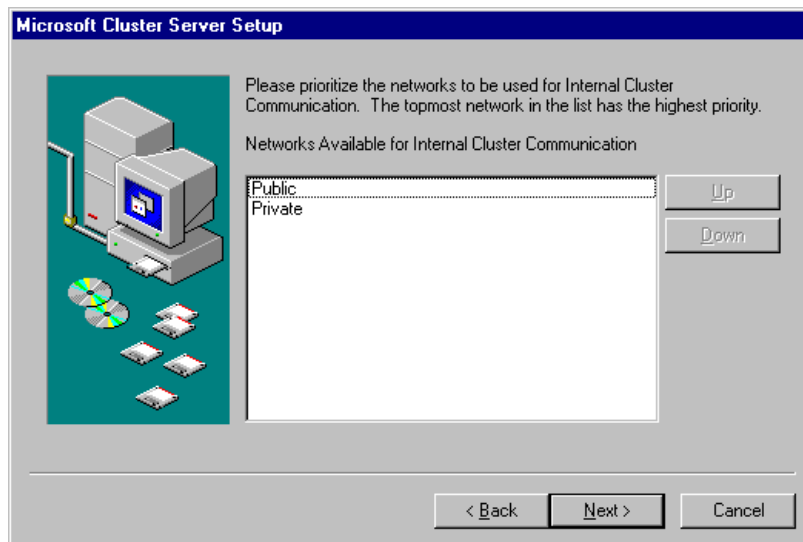


Figure 63. Network Adapter Priority

17. Here you prioritize the network to be used for internal cluster communications so that it has the highest priority. This is the private network. Select **Private**,

and then select **Up** to make the private network the highest priority. Click on **Next**. Figure 64 appears.

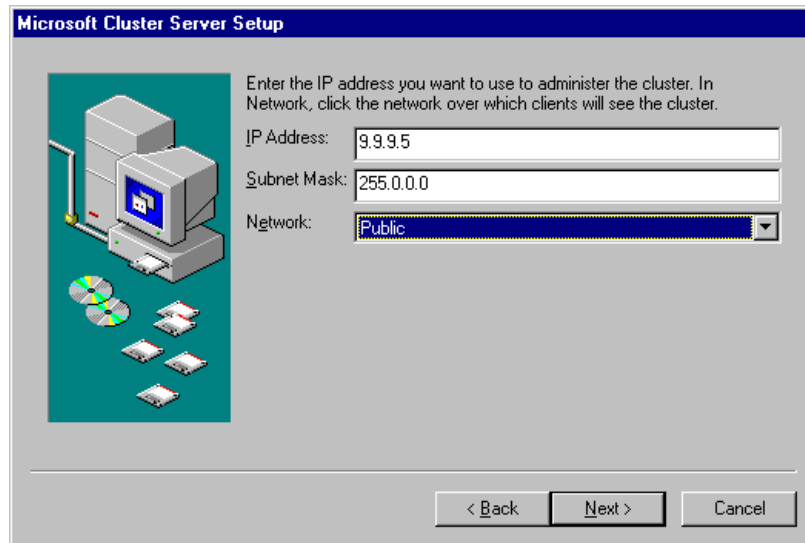


Figure 64. IP Address to Administer the Cluster With

18. You are now prompted to type the IP address and subnet mask to be used to administer the cluster. Select the network name the clients will use from the pull-down list. For example:

- IP Address: 9.9.9.5
- Subnet Mask: 255.0.0.0
- Network: Public

Note: Although the 255.0.0.0 subnet mask for the public network was used in our example, your network may require the use of Class C addressing and subnetting.

This is the IP address with which the cluster can be accessed by other machines in the network. It is primarily used for administration purposes. For normal client access you would define other virtual IP addresses.

Note that this is the IP address for the whole cluster, not just for one machine. If you connect through the server-specific IP address, you will lose your connection when that server fails. However, if you connect through the cluster IP address, you will continue to have access to the defined cluster resources, regardless of which node they may be running on. See Figure 65.

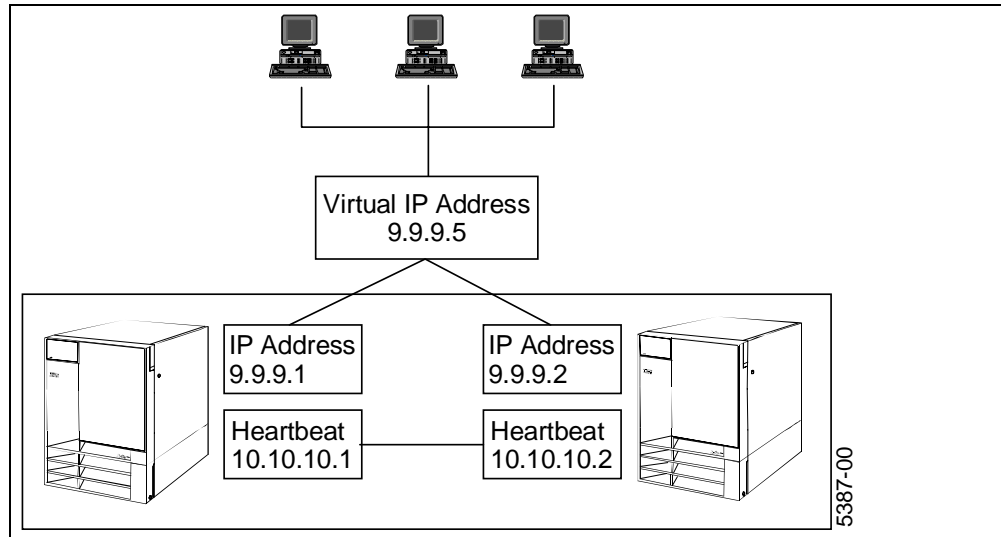


Figure 65. MSCS Virtual IP Address

This IP address is in addition to the addresses of the single machines and has to be a valid address within your network. Fill in the subnet mask as provided by your LAN administrator.

Contact your TCP/IP administrator for an IP address that will be used for the whole cluster.

Click on **Next** to continue.

19. At this point, all information necessary to install MSCS has been collected. Click on **Finish**. The installation program will then copy all the necessary files from the CD-ROM. It will also install the Cluster Administrator program, which will be used to administer the cluster later.

When installation has finished, you will see a window telling you that MSCS has been successfully installed and asking you to reboot the machine for all changes to take effect. Click on **OK**.

20. The Cluster Service will be automatically restarted upon startup.

21. Install Service Pack 4 and reboot the machine when prompted.

Note: You must have Node A running before you start installing MSCS on Node B.

6.1.2 Installing MSCS on the Second Node

The installation of MSCS on Node B is similar to that on Node A as described in 6.1.1, "Installing MSCS on the First Node" on page 61. Perform the following instructions to install MSCS on Node B:

1. Run **Start > Programs > Administration Tools (Common) > Enterprise Edition Installer**, select Microsoft Cluster Server to install and click on the **Start Installation** button.
2. You are then prompted for the location of the files on Enterprise Edition CD 2. Adjust the drive letter of the CD-ROM drive then click on **OK**.
3. Click on **Next** at the welcome screen. Click on **Accept** and **Next** when reminded that only certified configurations are supported. Figure 66 appears.

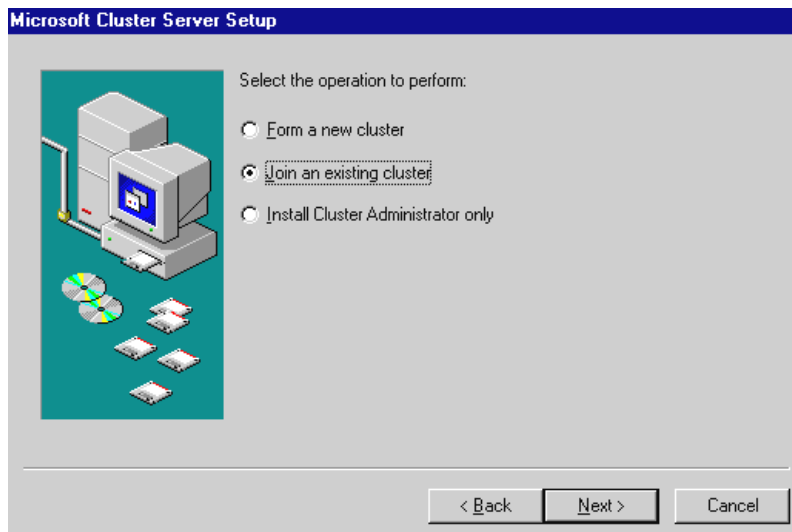


Figure 66. Join an Existing Cluster

4. Select **Join an existing cluster**, then click on **Next**. Figure 67 appears.

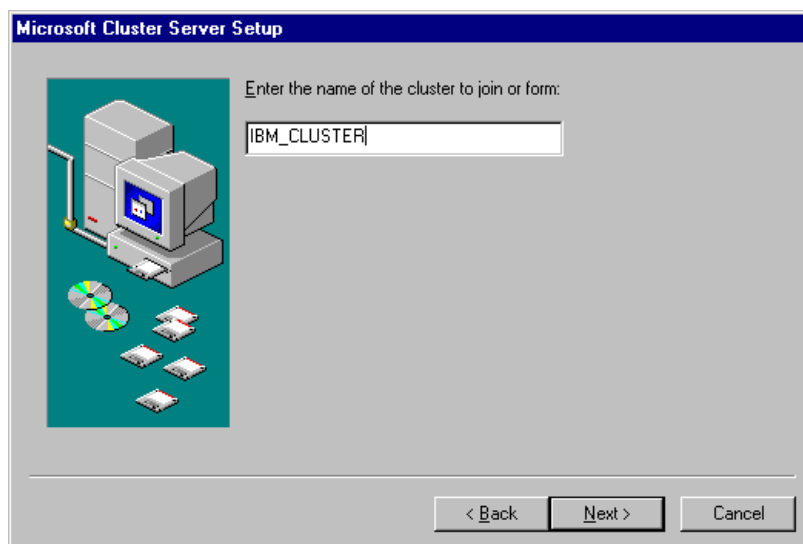


Figure 67. Specify the Cluster Name

5. Type the name of the cluster you specified in step 9 on page 63. In this example, we used IBM_CLUSTER. Click on **Next**.
6. You will see a panel where you have to define the directory in which the cluster files will be stored. This directory must be on a local drive. Accept the default: C:\WINNT\CLUSTER. Click on **Next**. Figure 57 appears.

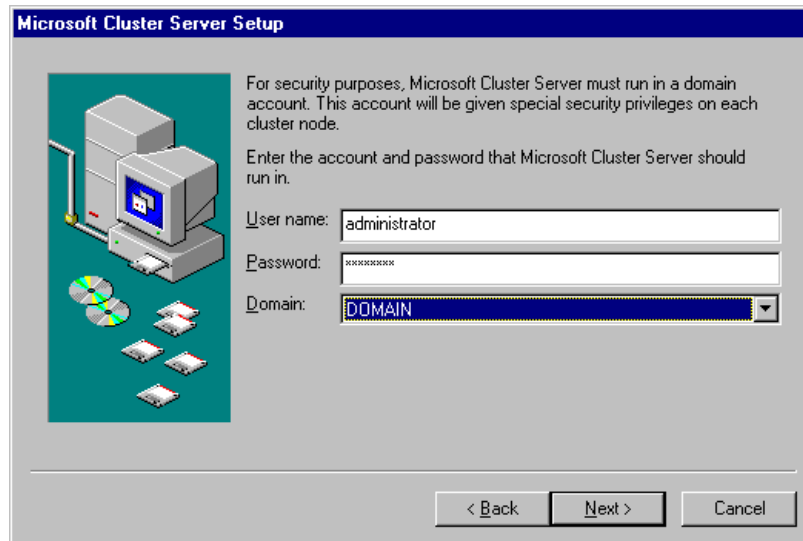


Figure 68. MSCS Logon Screen

7. Enter the password for the administrator user ID. The User name and the Domain name are already specified, based on the information you specified during the installation of Node A. Click on **Next**.

Note: You are not asked to verify the password, so ensure you type it correctly.

8. You will see a window telling you that all information has been collected. Click on the **Finish** button and the installation program will begin to copy all cluster and quorum files. It will also install the Cluster Administrator program.
9. When the installation has finished, a pop-up window tells you that MSCS has been successfully installed and asks you to reboot the machine for all changes to take effect. Click on **OK**.
10. The Cluster Service will be started automatically upon startup.
11. Install Windows NT Service Pack 4 and reboot.

Uninstall

If you have to uninstall MSCS, do the following before reinstalling to avoid problems:

- Use the Add/Remove program from the Control Panel folder to remove MSCS.
- On both cluster nodes, remove the remaining files from the directory where the program files were installed (for example, C:\WINNT\CLUSTER). Check also for hidden files.
- Delete the quorum resource (log files) on the shared drive or format the disk.

6.2 Testing the Cluster Installation

In this section, we've provided two test procedures to test your cluster installation. The first test includes moving a cluster group from one node to another. The second test involves shutting down Node B.

6.2.1 Moving the Cluster Group

To test the cluster installation, you are going to move the cluster group from NODE_A to NODE_B. NODE_B will then take over ownership of the group. Perform the following steps to test the installation:

1. Make sure that you are logged on to Windows NT on both servers.
2. From NODE_A, select **Start > Programs > Administrative Tools > Cluster Administrator**. Figure 69 appears.




Figure 69. Connecting to the Cluster

3. Type in either the cluster name (IBM_CLUSTER in this example), or one of the server names of the nodes (NODE_A or NODE_B here). Click on **Open**. Figure 70 appears.

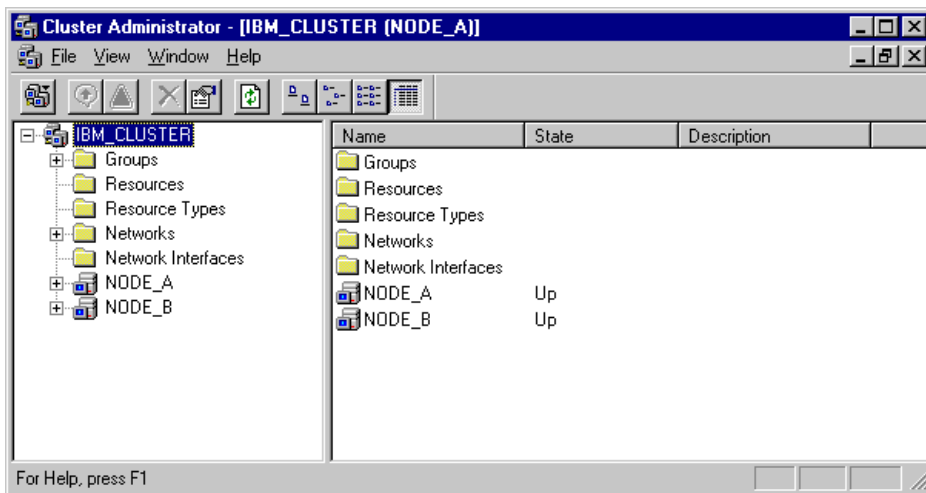


Figure 70. Cluster Administrator

Notice that both Nodes are listed in the window. If only one node is displayed, then it is likely you have a problem. Check the Windows NT Event Viewer and the status of the Cluster Server service on both servers. Check all hardware connections.

4. In Figure 70, expand the **Groups** subtree and click on **Cluster Group** in the left pane. Figure 71 appears.

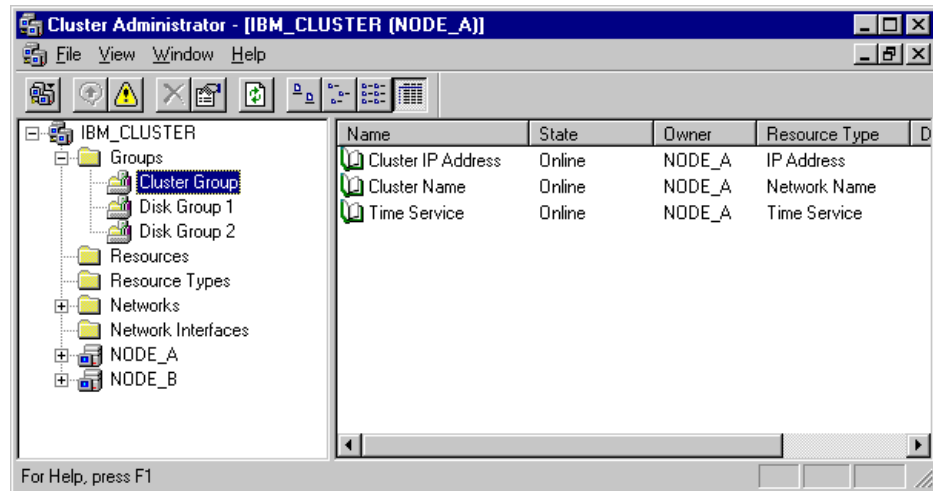


Figure 71. Cluster Group

5. Right click **Cluster Group** and select **Properties**. Figure 72 appears. Notice that the Cluster Group owner is currently NODE_A.

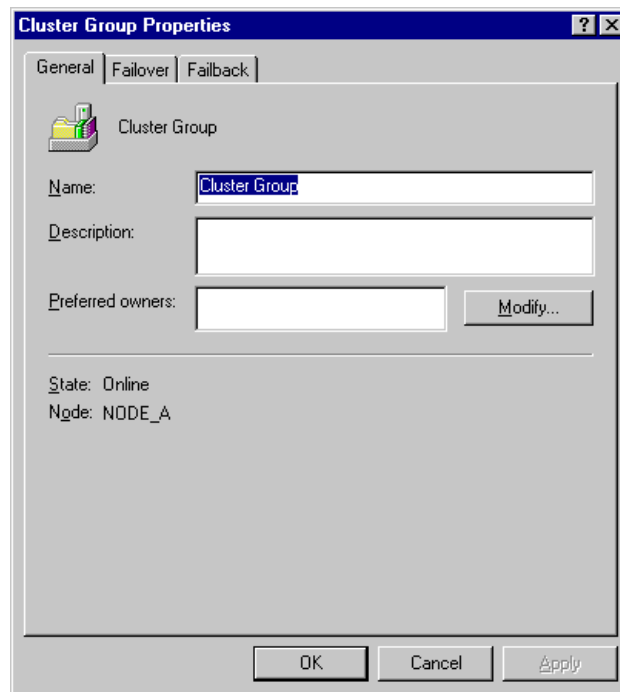


Figure 72. Cluster Group Properties

6. Right click on **Cluster Group** again, but this time select **Move Group**. The state of the group will go offline then online pending then online as the resource ownership is shifted from NODE_A to NODE_B.
7. Now NODE_B owns the group and the state is Online as shown in Figure 73.

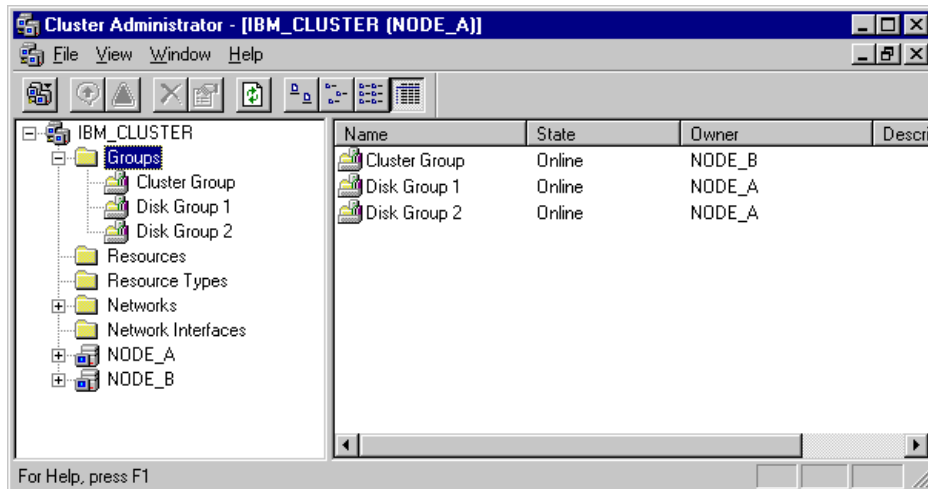


Figure 73. Cluster Group Now Owned by Node B

6.2.2 Shutting Down Node B

Shutting down Node B (or Node A) will transfer all resources and resource groups to the surviving node, Node A (or Node B). The same should occur if the server was simply powered off.

1. Leaving Cluster Administrator running on Node A, shut down Node B.
2. During the shutdown process, you should see the resource groups that were on Node B transferred to Node A, as shown in Figure 74:

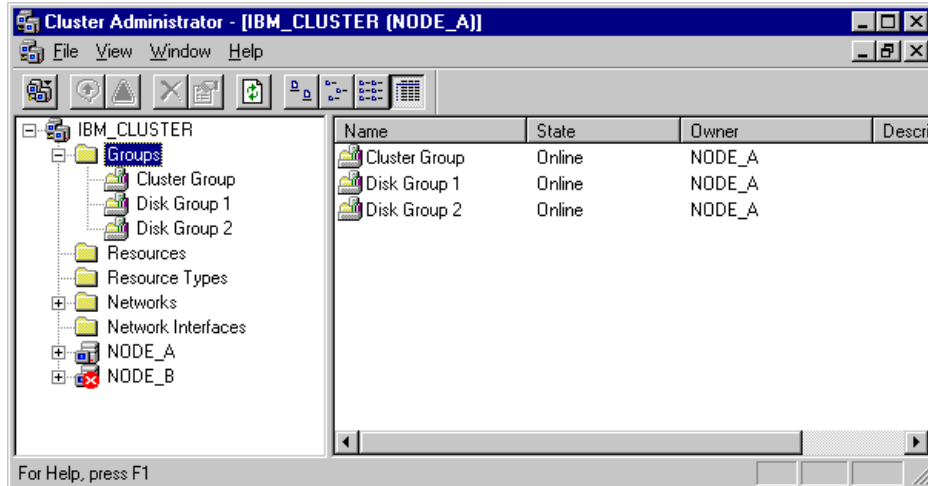


Figure 74. Node B Offline

6.3 Cluster Administrator on a Remote System

If you want to administer the cluster from another system in your network, you can install the Cluster Administrator on that workstation.

This workstation must have Windows NT Server 4.0 (or above) or Windows NT Workstation 4.0 (or above) installed.

Follow these steps to install the MSCS Cluster Administrator:

1. Run the following program from the Windows NT Server Enterprise Edition Disk 2 to start the installation process:
`\MSCS\CLUSTER\I386\SETUP.EXE`
2. Continue the installation by clicking on **Next**. Figure 75 appears.

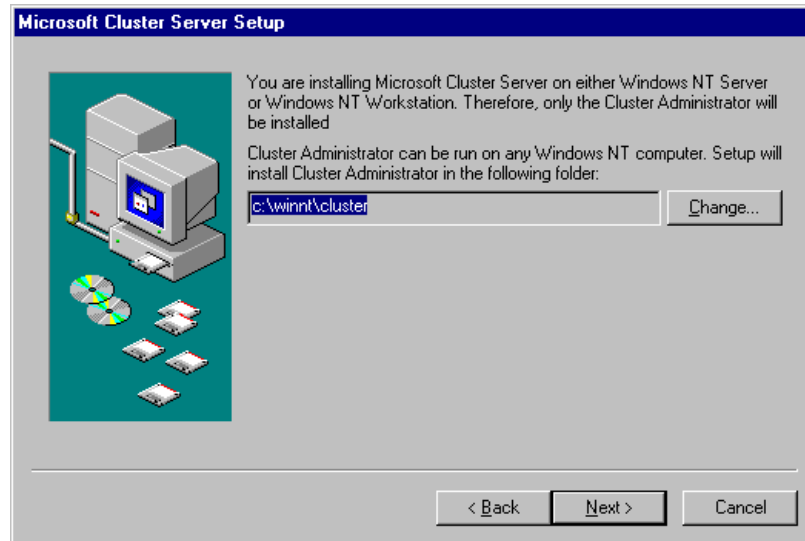


Figure 75. Installing the Cluster Administration Client

3. The installation program has detected that this is not a cluster node so it assumes you want to install the client only. Specify a directory or click on **Next** to continue.
4. You will then see a window telling you that all information has been collected. Click the **Finish** button here. The installation program will install the Cluster Administrator program.
5. You should now reapply Service Pack 4.

Appendix A. Special 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 IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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 the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

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, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

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

| | |
|--------------|-------------------|
| EtherJet | IBM ® |
| Netfinity | Netfinity Manager |
| RETAIN | ServeRAID |
| ServerProven | |

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

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

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries. (For a complete list of Intel trademarks see www.intel.com/dradmarx.htm)

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

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

Index

Numerics

10/100 Ethernet adapter 16
2 GB of server memory 4
3COM Fast Etherlink XL adapter 16
7133 enclosure 21
96H9835 2

A

access modes 4
active-active 25
active-hot spare 25
activity LEDs 2
applications, clusterable 24
arrays per adapter 4
arrays, creating 48

B

backup domain controller 19
backup interconnect 16
basic configuration 7
binding order 18
booting from SSA 3, 4, 20
bypass circuits 22

C

cables 3
cache, adapter 2
CBIOS 3
client access 16
CLOCK.EXE 9, 15
cluster 1
cluster partitioning 22
clustering models 19
configurations, cluster 7, 25
configurators 27
connectors 2, 21
copper cables 3

D

data access 5
data scrubbing 4
definitions 7
 backup interconnect 16
 cluster 1
 cluster partitioning 22
 dependencies 8
 failback 14
 failover 1, 13
 group states 12
 heartbeat 16
 is alive 15
 looks alive 15
 nodes 7
 polling 15

 preferred owner 14
 primary interconnect 16
 quorum resource 11
 resource group 9, 11
 resource monitor 8
 resource states 11
 resource types 9
 resources 7
 system redundancy 1
 virtual server 12

dependencies 8
DependOnService, registry value 18
device drivers, installing 33
DHCP 9, 16, 17
Disk Administrator 57
disk labels 58
disk number 22, 56
disk types 22
distributed transaction coordinator 9
domain controller 19, 33, 61
domains 19
DOS configurator 29
DRAM 2
dummy modules 21

E

ethernet adapters supported 16
Event Logger 29
 installing 37
Event Viewer 38
EVNCTRLF.TXT 38

F

failback 14, 25
 failback policy 15
 timing 15
failed state 11
failover 1, 13
 failover period 13
 failover settings 13, 24
 failover threshold 13
failures MSCS can address 24
failures MSCS cannot address 24
FDDI 18
features of the adapter 2
fiber cables 3
fiber optic considerations 22
file share 9
firmware 36
FlashOn 51
form a new cluster 63
formatting disks 44
free resources 45

G

generic application 9, 24

- generic service 10
- group states 12

H

- heartbeat 16
- hot spares 4, 52

I

- IIS 32
- IIS virtual root 9, 10
- initiator, primary 5
- installing
 - device drivers 33
 - firmware 36
 - IIS 32
 - MSCS 61
 - RSM 38
 - service packs 32, 69
 - SSA device drivers 33
 - SSA Event Logger 37
 - Windows NT 31, 31
- internal communications 15
- IP address 10
- IPX 17
- is alive 15, 24
- ISSAADLD 36
- ISSACFG 44

J

- join an existing cluster 70

L

- LEDs 2
- LEDs, disk 51
- load balancing 14
- local access 20
- local arrays 4
- looks alive 15, 24
- loop 2, 21
- low-level format 31, 46

M

- master initiator 20
- memory
 - 2 GB maximum 4
- memory, adapter 2
- Message Queue Server 10
- Microsoft Cluster Administrator 72
 - installing on a workstation 74
- Microsoft Cluster Server 1
 - installing 61
 - removing 71
 - testing 71
- mixing adapters 20
- models, clustering 19
- multi-homing configurations 17
- multi-mode fiber 3

- multiple network adapters 17
- multiple SSA adapters 20
- multiple subnets 17

N

- NetBEUI 17, 25
- Netfinity Manager 27, 38, 41
- network name 10
- networking 15
- nodes 7
- NTFS 11, 32

O

- offline state 11, 12
- online state 11, 12
- operating systems supported 5
- Oracle FailSafe 18

P

- partially online state 12
- partitions 57
- passwords 32, 64
- PCI slot 20, 21
- pending state 11
- performance 5, 21, 22
- physical disk 10
- planning for MSCS 23
- polling 15
- ports on the adapter 5
- preferred owner 14, 25
- preparations 7
- primary adapter 5, 14, 20, 21, 45
- primary initiator 5
- primary interconnect 16
- print spooler 10
- private arrays 4
- private IP networks 17
- proxy gateway 40
- Proxy Server 17
- public access 20
- public arrays 4
- publications v

Q

- quorum resource 11, 21, 22

R

- RAID
 - software-based 2
- RAID arrays per adapter 4
- RAID levels 3, 21
- redundancy, networking 16
- redundant network cards 17
- related publications v
- remote adapter 20, 21
- resource DLLs 7, 8
- resource group 9, 11, 25

- resource monitor 8
- resource states 11
- resource types 9
 - DHCP 9
 - distributed transaction coordinator 9
 - file share 9
 - generic application 9
 - generic service 10
 - IIS virtual root 10
 - IP address 10
 - Message Queue Server 10
 - network name 10
 - physical disk 10
 - print spooler 10
 - time service 10
- resources 7
- resynchronizing arrays 20
- RETAIN tip H165049 4
- risks 23
- RSM 27, 27
 - installing 38
 - Netfinity Manager 41
 - obtaining 27
 - password 28, 40
 - removing 41
 - tutorial 29
 - versions 27

S

- SAP R/3 18
- secondary adapter 5, 20
- serial interface chip 21
- server memory
 - 2 GB maximum 4
- ServeRAID 31, 32
- service packs 5, 32, 69
- shared disk model 19
- shared nothing model 19
- single-mode fiber 3
- SNA Server 17
- software RAID 2
- SplitConfirm 21, 49
- SRAM 2
- SSA
 - configurators 27
 - connectors 21
 - dummy drive modules 22
 - firmware 36
 - location of adapters 21
 - performance 21
 - RAID levels 21
 - RSM 27
 - staging 2
- SSA Event Logger 29
 - installing 37
- SSA RAID Cluster Adapter 2
 - staging 2
- stand-alone servers 19
- states
 - group states 12

- resource states 11
- static TCP/IP address 17
- subnets 17
- system redundancy 1
- system resources 55

T

- TCP/IP 16
- TCP/IP address order 18
- TCP/IP address ranges 17
- text-mode configurator 29
- time service 10

U

- user guides v

V

- virtual server 12

W

- Web resources
 - certifications 20
 - configurators 27
 - Event Logger 37
 - Netfinity clustering 1
 - network drivers 20
 - SSA 20
 - system BIOS 20
- Windows NT
 - installing 31
- Windows NT Event Viewer 38
- WINS server 17

SG24-5387-00

Printed in the U.S.A.

