



www.chinatag.com

CHINATAG

Microsoft 70-223

Microsoft Windows 2000
Advanced Server Clustering Services

Study Guide

DEMO Version

Copyright (c) 2003 Chinatag LLC. All rights reserved.

Important Note Please Read Carefully

For demonstration purpose only, this free version Chinatag study guide contains **10** full length questions selected from our full version products which have more than **200** questions each.

This Study guide has been carefully written and compiled by Chinatag certification experts. It is designed to help you learn the concepts behind the questions rather than be a strict memorization tool. Repeated readings will increase your comprehension.

For promotion purposes, all PDF files are **not** encrypted. Feel free to distribute copies among your friends and let them know Chinatag website. Our IT certification products start at only **\$7.99**.

Study Tips

This product will provide you questions and answers along with detailed explanations carefully compiled and written by our experts. Try to understand the concepts behind the questions instead of cramming the questions. Go through the entire document at least twice so that you make sure that you are not missing anything.

Latest Version

We are constantly reviewing our products. New material is added and old material is revised. Free updates are available for 180 days after the purchase. You should check the products page on the <http://www.chinatag.com> website for an update 3-4 days before the scheduled exam date.

Please tell us what you think of our products. We appreciate both positive and critical comments as your feedback helps us improve future versions. Feedback on specific questions should be send to feedback@chinatag.com.

Thanks for purchasing our products and look forward to supplying you with all your Certification training needs.

Good studying!

Technical and Support Team
Chinatag LLC.

TABLE OF CONTENTS

List of Tables

List of Figures

Introduction

1. Windows Clustering Technologies

- 1.1 Advantages of Implementing Cluster Service
- 1.2 Cluster Server Units and Terminology
- 1.3 The Methods of Communication between the Nodes
- 1.4 Cluster Service Implementation Models
- 1.5 Cluster Service Configuration Models
- 1.6 Network Configuration and Hardware Cluster Service Implementation Considerations
 - 1.6.1 Drive Technology Selection
 - 1.6.2 Shared Disk and Network Factors
 - 1.6.3 Hardware Factors
- 1.7 Preparing for Cluster Service Implementation
 - 1.7.1 Defining Resource Groups
 - 1.7.2 Determining the Failover Policy
 - 1.7.3 Assessing the Network Infrastructure and Associated Risks
 - 1.7.4 Determining the Suitable Hardware for Implementing Cluster Service

2. Network Adapter, Shared Storage and Cluster Service Installation and Configuration

- 2.1 Network Adapter Implementation and Configuration
 - 2.1.1 Renaming the Network Connections on the First Node
 - 2.1.2 Configuring the Private and Public TCP/IP Parameters
 - 2.1.3 Configuring the Private Network Adapter Settings
- 2.2 Implementing and Configuring the Cluster's Shared Storage Device
 - 2.2.1 Validating Disk Access among the Servers in the Cluster
 - 2.2.2 Using Redundant Array of Independent Disks (RAID) for Fault Tolerance
- 2.3 Cluster Service Installation Options
 - 2.3.1 Installation with Windows 2000 Advanced Server Installation
 - 2.3.2 Installation on an Existing Windows 2000 Advanced Server



- 2.3.3 Unattended Cluster Service Installation
 - 2.3.3.1 Automating Cluster Service Installation for an Initial Windows 2000 Advanced Server
 - 2.3.3.2 Automating Cluster Service Installation on an Existing Windows 2000 Advanced Server

2.4 Using the Cluster Service Configuration Wizard to Install Cluster Service

- 2.5 Performing a Rolling Upgrade from a Windows NT 4 Enterprise Edition Cluster
 - 2.5.1 Steps Executed for a Rolling Upgrade

2.6 Validating and Confirming Cluster Service Installation and Configuration

3. Administering, Managing and Supporting the Cluster

3.1 Managing the Cluster with Cluster.exe

- 3.2 Managing the Cluster with Cluster Administrator
 - 3.2.1 Cluster Administrator's Administrative Functions

- 3.3 Node Management and Administration
 - 3.3.1 Cluster Security Management
 - 3.3.2 Node Management
 - 3.3.3 Controlling Client Access and Shared Drive Access

3.4 Configuring Resources

3.5 Configuring Resource Groups

3.6 Configuring Virtual Servers

4. Implementing File Shares, Applications, Network Services, and Print Services

- 4.1. Implementing File Shares on a Cluster
 - 4.1.1 Installing and Configuring the Distributed File System (DFS)

- 4.2 Creating a Physical Disk Resource for Managing Storage in the Cluster
 - 4.2.1 The Private Properties of the Physical Disk Resource
 - 4.2.2 The Process for Extending an Existing Cluster Shared Disk's Disk Space
 - 4.2.3 The Process for Adding another Shared SCSI Bus between Nodes
 - 4.2.4 The Process for Adding Disk Devices to the Shared SCSI Bus
 - 4.2.5 The Process for replacing a Failed (online) Cluster Disk
 - 4.2.6 The Process for replacing a Failed (offline) Cluster Disk

4.3 Implementing Applications on a Cluster

4.4 Implementing DHCP on a Cluster



4.5 Implementing WINS on a Cluster

4.6 Implementing Internet Information Services (IIS) on a Cluster

4.7. Implementing Print Services on a Cluster

5. Installing SQL Server 2000 and Exchange Server 2000 on the Cluster

5.1 SQL Server 2000 Installation on the Cluster

5.1.1 Managing SQL Server 2000

5.2 Exchange Server 2000 Installation on the Cluster

5.2.1 Planning Exchange Server 2000 Installation

5.2.2 Setting up for and Installing Exchange Server 2000

5.2.3 Managing Exchange Server 2000

6. Techniques for Troubleshooting Cluster Service

6.1 Disaster Recovery Techniques

6.1.1 Cluster Backup and Restore Mechanisms

6.2 Troubleshooting Cluster Service

6.2.1 Installation

6.2.2 Startup Issues

6.2.3 Cluster SCSI Devices

6.2.4 Troublesome Client Network Connections

6.2.5 Cluster Nodes

6.2.6 Virtual Servers

6.2.7 Resource Groups

6.2.8 Resources

6.2.9 Resource Type

6.3 Troubleshooting using the Cluster Log

7. Network Load Balancing (NLB)

7.1 The Architecture of Network Load Balancing (NLB)

7.1.1 NLB Port Rules

7.1.2 The Convergence Process

7.2 Planning for NLB Installation

7.3 NLB Configuration

7.4 Troubleshooting NLB Configurations



LIST OF TABLES

- TABLE 3.1: Cluster.exe Main Arguments
TABLE 3.2: Cluster.exe Command Line Arguments

LIST OF FIGURES

- FIGURE 2.1: The Create or Join a Cluster Page
FIGURE 2.2: The Select an Account Page

70-223 MCSE Microsoft Windows 2000 Advanced Server Clustering Services

Exam Code: 70-223

Certifications:

Microsoft Certified Systems Engineer (MCSE)

Prerequisites:

This study guide assumes that you are able to install and configure Microsoft Windows 2000 Advanced Server and Active Directory on a Windows 2000 server. You should also be capable of managing users and security on the Windows 2000 domain.

About This Study Guide

This Study Guide is based on the exam questions for the 70-223 MCSE: Microsoft Windows 2000 Advanced Server Clustering Services exam. It presents all the information necessary to pass the 70-223 Microsoft Windows 2000 Advanced Server Clustering Services exam, and is detailed on the particular proficiencies examined in the exam. The Microsoft Windows 2000 Advanced Server Clustering Services exam is designed to test your knowledge on installing, configuring, administering and monitoring the clustering technologies supported by Microsoft. The information provided in this Study Guide is specific to the 70- 223 examination, and does not symbolize a complete reference work on the any other areas covered by the MCSE series of exams.

The following topics are covered in this Study Guide: The Advantages of Implementing Cluster Service, Cluster Server Units and Terminology, Cluster Service Implementation Models, Cluster Service Configuration Models, Network Configuration and Hardware Cluster Service Implementation Considerations, Determining the Failover Policy for Cluster Service, Assessing the Network Infrastructure and Associated Risks, Using Redundant Array of Independent Disks (RAID) for Fault Tolerance, Cluster Service Installation with Windows 2000 Advanced Server Installation, Cluster Service Installation on an Existing Windows 2000 Advanced Server, Unattended Cluster Service Installation, Using the Cluster Service Configuration Wizard to Install Cluster Service, Administering, Managing and Supporting the Cluster, Implementing File Shares on a Cluster, Installing and Configuring the Distributed File System (DFS), Creating a Physical Disk Resource for Managing Storage in the Cluster, Extending an Existing Cluster Shared Disk's Disk Space, Adding another Shared SCSI Bus between Nodes, Replacing a Failed (online) Cluster Disk, Implementing Applications on a Cluster, Implementing a DHCP Server on a Cluster, Implementing a WINS Server on a Cluster, Implementing Internet Information Services (IIS) on a Cluster, Implementing Print Services on a Cluster, Installing SQL Server 2000 and Exchange Server 2000 on the Cluster, Troubleshooting Cluster Service, Troubleshooting using the Cluster Log, Network Load Balancing (NLB), Architecture of Network Load Balancing (NLB) , Planning for NLB Installation, NLB Configuration, and Troubleshooting NLB Configurations

Intended Audience

This Study Guide is targeted specifically for people that have knowledge on configuring Microsoft Windows 2000 Advanced Server and Active Directory on a Windows 2000 who now desire to complete the 70-223 Microsoft Windows 2000 Advanced Server Clustering Services exam. The information in this Study Guide is specific to that exam. It is not a complete reference work on the other exam topics relating to the MCSE qualification. Although our Study Guides are aimed at new comers to the world of IT, the concepts dealt with in this Study Guide are complex. Knowledge on configuring Microsoft Windows 2000 Advanced Server and Active Directory on a Windows 2000 would be advantageous.

How To Use This Study Guide

To benefit from this Study Guide we recommend that you:

- Study each chapter carefully until you fully understand the information. This will require regular and disciplined work. Where possible, attempt to implement the information in a lab setup.
- Be sure that you have studied and understand the entire Study Guide before you take the exam.

Good luck!

1. Windows Clustering Technologies

Microsoft Clustering Server (MSCS) was initially launched in the Windows NT Server Enterprise Edition. MSCS made it possible for organizations and enterprises to increase server **availability** for key business applications by grouping multiple physical servers that had network access to a common drive, into a cluster. What this means is that when a server in the cluster has a failure, another server in the cluster would resume the services of, or those applications hosted by the failed server. Servers in the cluster are typically referred to as **nodes**, while services and applications are referred to as **resources**. The physical servers in the cluster operate as one network server and users would continue to use one computer name or server name to access those key applications hosted by the servers in cluster. This removes the need of reconfiguring user rights when a server has a failure, and its applications are hosted by another server in the cluster.

A cluster can therefore be defined as the grouping of two or multiple physical servers that are portrayed as, and operate as one network server. These servers provide redundancy to the enterprise network by resuming operations of a failed server within the cluster. This procedure is known as **failover (FO)**. The process of **failback** occurs when a failed server automatically recommences performing its former operations once it is available again or online. The cluster can also be configured to provide load balancing features.

With the introduction of Windows 2000 this technology became known as **Microsoft Cluster Service**. Microsoft Cluster Service introduced further enhancements to its predecessor. The functions associated with installing and configuring the technology has also been simplified. Microsoft Cluster Service is best suited for network services that require a high degree of availability. The technology is therefore ideal for use with file and print servers, and messaging applications.

Windows Load Balancing (WLB) was initially introduced in Windows NT Server Enterprise Edition. The WLB technology has been modified as well with Windows 2000, and became **Network Load Balancing (NLB)**. NLB makes it possible for multiple servers to be grouped to respond to users' requests. Network Load Balancing (NLB) is an example of a technology that provides **dynamic load balancing**. NLB is well suited for a Web server farm because of its fast expansion capabilities. NLB can also be used for e-commerce sites that need its fast expansion capabilities.

These technologies are integrated into the Microsoft Windows 2000 Operating System. They supply administrators with a sophisticated ability for managing and administering network services, both enterprise and network.

1.1 The Advantages of Implementing Cluster Service

The benefits associated with implementing Cluster Service are:

- Windows 2000 Advanced Server computers can be **integrated** with current network resources, that is, they can reside in the exact cluster as Windows NT 4 Server computers.
- The technology provides **redundancy** for network failures because another node in the cluster resumes the services of the failed server. This increases server availability for key business applications and network services.
- There is no manual **configuration** associated with failback because the failed server automatically takes on its former operations.



- Application **response time** can be improved by dispersing applications across multiple servers.
- Cluster Service also reduces downtime associated with scheduled **maintenance** downtime. When a server in the cluster is scheduled for an upgrade, its services / applications can be manually shifted to another node in the cluster. Cluster Service therefore provides high availability by not disrupting client or user access when upgrading nodes, application and resources.
- A network using Cluster Service also enjoys improved **scalability** because servers can be expanded while client access is still ensured.
- The nodes, services and applications in the cluster can be managed, controlled and administered **remotely**, and in the same manner as though they were all hosted on one server. Applications and services can also be taken off the network for upgrades.

1.2 Cluster Server Units and Terminology

Before delving into the modules of the Cluster Service, it is important to study the following expressions that are generally used in the remainder of the discussion on clustering:

- **Active/Active cluster implementation:** In this implementation, every node is able to administer the cluster's resource groups, and each node is able to automatically resume the services / applications of another node.
- **Active/Passive cluster implementation:** In this implementation, a particular node is responsible for the cluster's resource groups explicitly assigned to it.
- A **common resource** refers to a resource, such as the private cluster network, that each node in the cluster has access to.
- **Dependency** is the terminology used to describe the relationship amid two resources that function in the identical resource group. In this instance, each of the two resources has to operate in the identical group.
- **Domainlet** is an alternative domain that is typically used for the cluster nodes instead of the conventional domain. The domainlet has restricted qualities for groups, authentication and policies. The use of a domainlet for the cluster nodes saves on server overhead that would have been increased had the conventional domain been used.
- **IsAlive Check** is used by the Resource Monitors to thoroughly examine the condition of a resource. An IsAlive Check failure initiates the failover procedure.
- The **LooksAlive Check** is simpler than the IsAlive Check, and is also used by the Resource Monitors to examine the condition of a resource. When this check is returned as inconclusive, or fails, the more thorough IsAlive Check is initiated.
- **Offline** refers to a cluster's resource that **is not** capable of providing the related service.
- **Online** refers to cluster's resource that **is** capable of providing the related service.
- A **quorum resource** has to exist in order for a node in the cluster to carry out its responsibilities. This common resource holds the cluster database's synchronized version that stores management data for the cluster. The quorum resource is situated on the physical disk of the shared drive of the cluster.

- **Resources**, as mentioned earlier, refers to components like applications and services controlled by Cluster Service. These components run on one node at any given time.
- A **resource group** is set of resources like the IP address or disk space needed for a specific application to run. A resource group has a unique IP address and a network name associated to it. Where the resources are reliant on each other, they have to be in the identical resource group, and on the identical node. A resource group is administered as one logical unit.
- A **virtual server** enables the means for clients to access a resource via the identical IP address and NetBIOS name. With virtual servers, clients are not dependent on the actual node that is currently managing the resource. A cluster can have multiple virtual servers that each has their own individual IP address and NetBIOS name.

Remember that each node in the cluster has to be part of the **same** domain, and Cluster Service must be installed on **every node** within the cluster. Cluster Service is divided into units. Each unit or component in the Cluster Service has specific functions related to it. The units and their related responsibilities are outlined below:

- The **Checkpoint Manager** is responsible for **verifying and recording** the registry data of a resource. Check point data (resource alterations) is saved to and stored in the quorum recovery log. When a resource is offline, the unit records a checkpoint to the quorum disk. The Checkpoint Manager unit also writes to a new resource's registry data prior to it being online. This unit therefore makes certain that failover can take place in the cluster by carrying out registry checkpointing.
- The **Communications Manager / Cluster Network Driver** unit is responsible for **communication** among the cluster nodes. The unit also informs the nodes of a node failure. The Communications Manager uses **Remote Procedure Calls (RPCs)**.
- The **Configuration Database Manager / Database Manager** is responsible for the management of the cluster **configuration database** kept in the registry of every node. Each node in the cluster has a Configuration Database Manager that makes certain that any updates to the configuration database are precise, constant and reliable. The configuration database holds cluster configuration information on the resource groups, resources and the cluster.
- For the failover process, the **Failover Manager** determines the node that should manage a resource. When a cluster has multiple nodes, each node's Failover Manager takes part in the process of determining or negotiating which node would resume the functions of the node that had a failure. When communication in the cluster is down, the Failover Manager proceeds with failover. Because the nodes are not able to communicate among each other, they communicate with the quorum resource that in turn ensures that a single node has an online resource. This node then returns the remainder of the nodes' resources to the online state.
- The **Event Processor** is the unit that starts the Cluster Service. This unit transmits event signals to the cluster's nodes. Event signals are messages that contain vital activity information like status changes, which are also distributed to different units in Cluster Service. This unit also provides support for the cluster API event component.
- The **Event Log Manager** is responsible for making sure that the nodes each have identical **event log entries**. Event log information of a specific node is replicated to the remainder of nodes.
- The **Log Manager** records any changes to the quorum resource's **recovery logs / quorum logs**.

- The **Global Update Manager** supplies the interface for the remainder of the Cluster Service units to manage **state changes**. State changes are broadcasted to all active nodes in the cluster.
- The **Node Manager** determines the nodes which should perform resource group management. This assignment is dependent on the availability of the node, and group preference lists. The Node Manager informs the Membership Manager to send regroup events.
- The **Membership Manager** is responsible for managing the membership to the cluster. The Membership Manager initiates a regroup event when a node in the cluster has a failure. This results in the remainder of the online nodes updating their membership lists accordingly. When the previously offline node is online again, another regroup event is initiated and these lists are updated again.
- The **Resource Manager** controls the resources and dependencies, and sets off resource group failover. The Node Manager and Resource Monitors communicate resource and state information to the Resource Manager
- **Resource Monitors** supply the communication vehicle for the cluster and resource DLLs. These **Resource Monitors** use Remote Procedure Calls (RPCs) to communicate with Cluster Service. The default configuration is that one Resource Monitor is running on each node. This can however be overridden. This unit makes certain that every resource is running correctly by using callbacks to the resource DLLs. The IsAlive and LookAlive cluster APIs are used for confirming resource availability.
- The **Object Manager** manages a database on the cluster's mechanisms or objects including the nodes, resource groups and resources.
- **Resource DLLs** supply the mechanism for Cluster Service to uphold communications with its supported applications. Applications supported by Cluster Service could include **cluster aware** applications and **cluster unaware** applications, and **management applications**. Cluster management applications like the Cluster Administrator, are applications that use cluster APIs to communicate with Cluster Service. Cluster aware applications reside on cluster nodes. These applications use their own DLLs that are specific for the actual application. Cluster aware applications can be divided into applications that are controlled as resources, like Exchange 2000, and applications that interrelate with the cluster. Applications that interrelate with the cluster are not regarded as resources of the cluster. Cluster unaware applications use the default resource DLL of Cluster Service. These applications are unaware of the cluster. These cluster unaware applications therefore do not use the cluster API for communication.

1.3 The Methods of Communication between the Nodes

There are three methods in which the nodes in the cluster can communicate:

- The **quorum resource** can be used as a method of communication. A node may be offline when a configuration modification is performed. When this occurs, the configuration modifications are held in the quorum resource's quorum log. Once the node is online again, the configuration modifications are accessible to the node.
- **Remote Procedure Calls (RPCs)** are used in conjunction with User Data Protocol (UDP) packets. RPCs communicate information among online nodes when resource changes take place.



- A node in the cluster uses **cluster heartbeats** to confirm whether the remainder of the nodes' network interface is active and online. A cluster heartbeat is a UDP packet / datagram that is sent periodically by the Node Manager of the nodes in the cluster. The datagram is 48 bytes. A node is regarded as failed when it does not reply to a cluster heartbeat. The initial node in the cluster starts sending the cluster heartbeat at 0.50 second intervals once another node receives membership to the cluster. The second node in the cluster replies to the cluster heartbeat in 0.20 seconds. When the initial node does not receive a reply, it commences transmitting a total of 18 cluster heartbeats to the second node. The 18 cluster heartbeats take up a total of 5.3 seconds. The node is regarded as failed when it fails to reply to these consecutive 18 cluster heartbeats.

1.4 Cluster Service Implementation Models

Shared Device and Shared Nothing are the two implementation models that are based on industry standards. In the **Shared Device implementation model**, the node's software applications can access the hardware available in the cluster. Applications have a Distributed Lock Manager (DLM) for data synchronization management. This is essential when more than one application use the same processor. The DLM also solves hardware resource conflicts. This implementation model can lead to traffic overhead.

With the **Shared Nothing implementation model**, a single device can access a hardware resource at any given point in time. This is the default Cluster Service implementation model. Every node basically manages its local disk. Devices, such as disk devices and applications that are common are managed by one node. When another node needs to access a hardware resource managed by a different node, it has to submit a request to that specific node. Another node can only claim management of a hardware resource when the particular hardware resource's node has a failure.

1.5 Cluster Service Configuration Models

Cluster implementations offer a choice between five configuration models. This is an important consideration because the configuration model chosen would have an impact on cluster performance, and the degree of availability ensured during a failure. The different configuration models are outlined below:

- With the **High Availability with Static Load Balancing Configuration Model**, the nodes each have particular resources that they are accountable for. To ensure availability during failover, each node has to be sufficiently capable of supporting another node's resources. Because a node resumes the responsibilities of another node during failover, this configuration model leads to decreased performance for the duration of the failover.
- With the **Hot Spare Node with Maximum Availability Configuration Model**, a single primary node manages the resources. The hot spare node is not used at the same time as the primary node, and it only manages the resources when the primary node has a failure. This model therefore ensures very high availability and very high performance during failover. Unfortunately, the costs associated with the configuration model are also very high.
- The **Partial Cluster Service Configuration Model** builds on the principles of the former model. The only difference being that the node that is responsible for resources also supports cluster unaware applications. The data of these cluster unaware applications are held on the local disks of the node. When failover occurs, the cluster unaware applications stay unavailable for the duration of the failover. This configuration model provides high availability for resources that are included in the

failover process. Cluster unaware applications are not part of this process and performance for these applications is greatly reduced at times of failover.

- With the **Virtual Server Configuration Model**, a single node exists in the cluster. No failover process exists in the cluster. Virtual servers can be implemented to react to user's requests. At a later stage, when additional nodes are implemented for the cluster, resources can be grouped into the virtual servers without needing to reconfigure any clients.
- The **Hybrid Configuration Model** can be regarded as a grouping of the above configuration models. In this configuration model, every node in the cluster manages their individual resources. Because this model is a grouping of the other models, availability during failover is ensured for those resources specified for failover.

1.6 Network Configuration and Hardware Cluster Service Implementation Considerations

1.6.1 Drive Technology Selection

The drive technologies available for utilization in Service Cluster implementation are Small Computer System Interface (SCSI) and Fibre Channel.

Small Computer System Interface (SCSI), the more frequently deployed drive technology, is the main drive technology referred to in this study guide. This drive technology is relatively cheaper than the Fibre Channel drive technology. The devices can either be connected internal to the server or it can be connected externally. When connected externally, a chain of devices are established. SCSI can support a maximum of 15 devices for each adapter. It can also support many adapters for a server.

Fibre Channel drive technology on the other hand is more costly than the previously described drive technology. It also requires high expertise for its intricate implementation. Fibre Channel does however provide a faster physical connection for the servers, and an enhanced connection for the serial disk interfaces. With Fibre Channel, a single interface is offered with a transmission rate of 100 Mbps in either direction. Fibre Channel also supports a vast quantity of protocols such as Internet Protocol (IP), Asynchronous Transfer Mode (ATM), SCSI, and the IEEE 802.2 standard. Fibre Channel is compatible with the SCSI commands.

1.6.2 Shared Disk and Network Factors

The **shared disk hardware requirement considerations** for Service Cluster are listed below:

- The primary requirement is to ensure that the shared drives are attached (physically) to the nodes that are part of the cluster.
- Shared disks should be configured as a Basic disk, and the NTFS file system has to be used when formatting each partition of a shared disk.
- SCSI drives and adapters should each have its own unique SCSI Identifier (ID).

The **network requirement considerations** for Service Cluster are:

- The cluster should be defined with its own unique NetBIOS name.
- As mentioned earlier, nodes in the cluster have to belong to the same domain. A domain user account has to be defined for this domain that would be used by Cluster Service. The only instance that

allows the domain account to be configured in a different domain is where a trust relationship exists between the concerned domains.

- Five unique IP addresses have to be defined for a Service Cluster implementation that has two nodes:
 - Two private addresses are used for communication among the nodes
 - Two public addresses are used by the network interface cards (NICs) connected to the LAN. Each node must have a NIC for the private network, and a NIC for the public network.
 - One public address is used by the cluster.

1.6.3 Hardware Factors

The **hardware requirement considerations** for Service Cluster are:

- Windows 2000 Advanced Server / Windows 2000 Datacenter Server has to be installed on the boot partition of each server. The servers are not allowed to share these boot partitions.
- Each server has to have two PCI NICs.
- The storage host adapter for SCSI / Fibre Channel has to be separate.
- At least one external drive that contains multiple RAID configured drives, a SCSI drive, must be connected to the servers. More than one drive can be connected though. These are the drives that are shared among the servers. The suitable cabling should be used to connect these external drives to the servers.
- It is recommended to implement the identical RAM processor, as well as drive configurations in every node.

1.7 Preparing for Cluster Service Implementation

Detailed planning for Service Cluster implementation is vital in ensuring that the actual implementation has the correct defined and configured resource groups, failover policies, protocols and hardware requirements. The following section outlines some key components that should be deliberated during the planning phase of Cluster Service implementation.

1.7.1 Defining Resource Groups

Recall from earlier discussions that resource groups are defined as a group of resources like the IP address or disk space needed for a specific application to run. Where the resources are reliant on each other, they have to be in the identical resource group, and on the identical node. It is therefore important to group resources according to their function and dependencies. A resource group typically has an IP address, unique network name, the applications / services that are part of a resource group, and a physical disk. Clients access the applications / services via a virtual server that has a NETBIOS name and IP address on the public / external network. During failover, the node resuming another's applications / services receives this address. The applications and their dependencies, and all other resources that are not considered applications, to be hosted by the cluster have to be determined. Cluster Service includes a feature that makes certain that all related dependencies are started prior to the applications / services initiating. Applications have to be classed as cluster aware and cluster unaware applications. This determines which resource groups would be part of the failover process. Each resource group that is created from these resources is virtual servers. A dependency

tree would be helpful in this planning component. A dependency tree is a diagram that holds the resource groups and their related dependencies.

1.7.2 Determining the Failover Policy

The failover policy assigned to a particular resource group stipulates the resource group's actions during the failover process. A failover policy has to be determined for each resource group in the cluster. The following failover policies can be stipulated for a resource group:

- **Failover Timing:** In this failover policy, Cluster Service initiates the failover process when a resource group fails. The failed resource group is shifted to another node in the cluster. An optional configuration option available is to allow Cluster Service to restart the resources in the failed resource group for a specific period of time, prior to shifting the resource group to the other node.
- **Failback Timing:** The default configuration is that the failed resource group is failed back to its primary node or preferred node as soon as it is online again. This can be overridden to a better option which is to allow fail back to only occur during off peak network hours.
- **Preferred Node:** This option is useful in ensuring that the node specified for, or preferred for a resource group remains the preferred node for the resource group. Specifying a preferred node for a resource group eliminates the need of manually moving a resource back to the node preferred for the resource group.

After defining the failover policies for the resource groups, it is equally important to ensure that the nodes have the capacity to run and manage the applications / services in each resource group. Future expansion possibilities for the cluster have to be included here as well. CPU and memory requirements for the applications have to be determined, as well as hard drive storage capacity.

1.7.3 Assessing the Network Infrastructure and Associated Risks

Certain network services currently running in the network can have an affect on Cluster Service implementation. The network adaptor installed in a cluster's node has to use the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol. This could be an important planning consideration if the current network environment does not already use the protocol. The current IP configurations used by the existing network environment have to be established and confirmed. The public network adapters need to have TCP/IP and NetBIOS for the networking protocol. NetBIOS is used by clients to access the virtual server name. It is also used to access the Cluster Administrator remotely. The private network adapters must have the NetBIOS over TCP/IP option disabled during configuration. The private adapters and virtual clusters have to be statically configured addresses. The public adapters can have static addresses or the address information can be acquired from a Dynamic Host Configuration Protocol (DHCP) server. If a DHCP server is used instead of static addressing, it is important that permanent leases are used. The only downfall to the DHCP option is that the cluster would be unavailable when the lease expires while the DHCP server is offline.

Recall from an earlier discussion that the nodes in the cluster have to belong to the same domain. The domain should have a unique user account. The only exception to having the user account outside the nodes' domain is where a trust relationship exists between the particular domains. This account has to be the same for all the nodes hosted in the cluster. The nodes should also have access to a domain controller. A domainlet can be implemented when the nodes are domain controllers within the domain. Once the domain

model is clarified, it is important to ensure that the nodes can access a DNS server. The DNS server is used by the nodes for name resolution.

Cluster Service ensures availability during the failover process that can be enhanced by eliminating network points of failure as far as possible. It is good practice to establish failure **backup plans** that could include items like planning for and using Uninterruptible Power Supplies (UPSs).

1.7.4 Determining the Suitable Hardware for Implementing Cluster Service

This planning component involves identifying the node's **hardware requirements** for Cluster Service to function. The actual applications / services implemented on the node needs to be kept in mind when planning these hardware requirements. Some factors to consider could include verifying that the network adapters supply high throughput, and that hardware configurations supply fast disk access if it is a requirement. When implementing a file sharing cluster, it is important to ensure that the network adapters supply high throughput and do not restrict performance. When determining these hardware requirements, the support offered by the Original Equipment Manufacturer (OEM) should also be considered.

Windows 2000 Advanced Server also supports certain applications not used by Cluster Service that can be disabled. Resources normally used by these services could also be freed up. Services like the Server Service and the Microsoft Computer Browser Service could then be optimized.

The paging file and physical RAM make up the virtual memory of the server. The **paging file** must not be placed on the shared disk. It is good practice to install a different local drive for this file on every node in the cluster. This in turn would augment virtual memory performance of the node. The paging file should be 2.5 times the amount of physical RAM in the node for those applications that consume vast amounts of memory resources.