



Technical report:

Microsoft Cluster Service and Microsoft SQL Server 2000 Virtual Server with IBM System Storage N series and SnapDrive

Implementation

• • • • • • • • •

Document NS3248-0

September 24, 2007



Table of contents

Abstract	3
Introduction	3
Assumptions and System Requirements	3
Configuration Overview	4
Network.....	4
Creating the MSCS Quorum and SQL Server Data Disks	4
Microsoft Cluster Server	5
Installing MSCS	5
Testing the MSCS Configuration	8
SQL Server 2000	9
Installing a SQL Server Virtual Server	9
Configuring Node2.....	13
Configuring an Active/Active Cluster	13
Configuring Multiple Disk Resources.....	13
Failover and the Virtual Server	14
Testing Virtual Server Failover	14
Caveats	15
Trademarks and special notices	16



Abstract

This document outlines steps necessary to implement a Microsoft Cluster Server/SQL Server 2000 Virtual Server cluster as one component in a high-availability solution when storing SQL Server 2000 databases on IBM System Storage N series featuring SnapDrive and other storage management tools.

Introduction

This report outlines steps necessary to implement a Microsoft® Cluster Server (MSCS) and SQL Server 2000 Virtual Server cluster as one component in a high-availability solution when storing SQL Server 2000 databases on an IBM® System Storage® N series. Presented below is a description of the steps necessary for both active/passive and active/active SQL Server 2000 Virtual Server configurations in a Microsoft Windows® 2000 Advanced Server two-node MSCS cluster.

Assumptions and System Requirements

IBM System Storage N series with SnapDrive™ 2.0 is a Microsoft Management Console (MMC) plug-in used in the creation and management of Fibre Channel Protocol (FCP) logical unit numbers (LUNs) and their IBM System Storage N series with Snapshot™ copies.

A working installation of Windows 2000 Advanced Server, Service Pack 3 (SP3), and SnapDrive 2.0 as well as familiarity with SnapDrive 2.0 and LUN administration are assumed. All supported IBM System Storage N series with Data ONTAP® versions apply. More information about SnapDrive 2.0, SQL Server 2000, and Data ONTAP administration can be found in related technical reports and documentation listed in the reference section at the bottom of this report.

It is necessary for the host cluster nodes and the filer to reside in the same W2K domain. If you already have a working installation of an MSCS cluster and can successfully failover between nodes, continue to section 4.2, which discusses creating the SQL Server data disk.

Note: It is recommended that you install and test first in a non-production environment using production-like workloads. Once preproduction testing is complete, the configuration can be rolled successfully into production.



Configuration Overview

Network

This configuration used Gigabit Ethernet with separate subnets for public and private communication. FCP is required for use with LUNs. Access to the LUN used to store the SQL Server database was through a Brocade 3800 switch. Section 5.1 describes Microsoft's recommendations regarding the use of public and private networks with MSCS.

Creating the MSCS Quorum and SQL Server Data Disks

The MSCS quorum resource (disk) is used for the recovery of the cluster and stores configuration data consisting of all changes applied to the cluster and the current state of each cluster component. Create a qtree and CIFS share on the filer to be used for the MSCS quorum disk. In this test, drive "Q:" was created within the share to act as the quorum disk for the configuration. A second qtree, CIFS share, and cluster group on the filer and host, respectively, were created for use with SQL Server. Drive "R:" was then created within this cluster group to store the SQL Server 2000 system databases on the filer.

The LUNs were defined as shared resources during their creation.

Microsoft Cluster Server

Installing MSCS

Start > Programs > Administrative Tools > Configure Your Server

If you have not already installed MSCS you will be prompted to configure Cluster Server immediately following installation of Windows 2000 Advanced Server. It is recommended that you apply Windows SP 3 to both nodes before configuring MSCS.

After specifying the cluster name and domain account to be used by the cluster service, be sure the disk to be used for the cluster quorum is listed under "Managed disks."

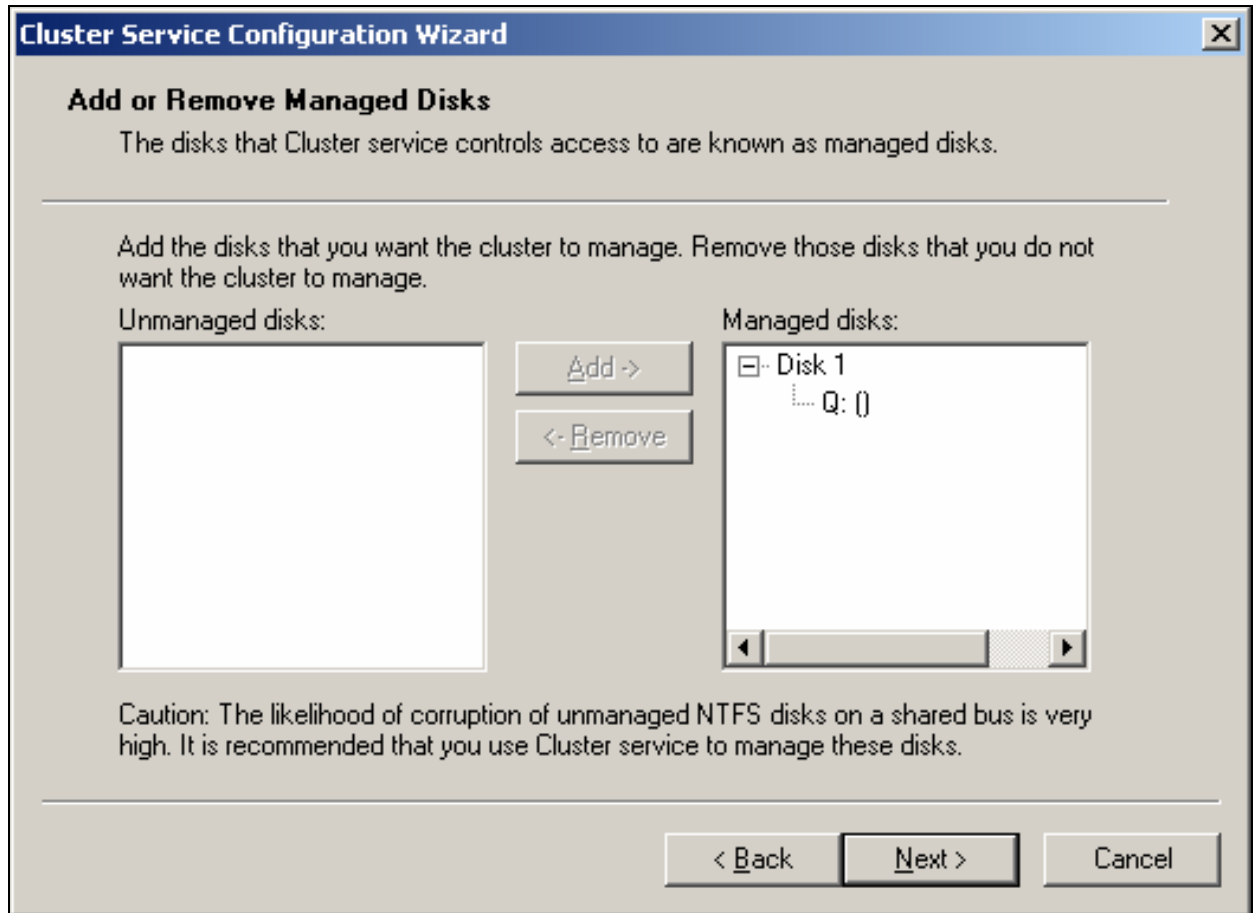


Figure 1, Cluster Administrator: Managed Disks

Only a single disk resource is required in the primary MSCS cluster group for configuration of the cluster quorum disk.

Then, specify the disk be used for the quorum resource:

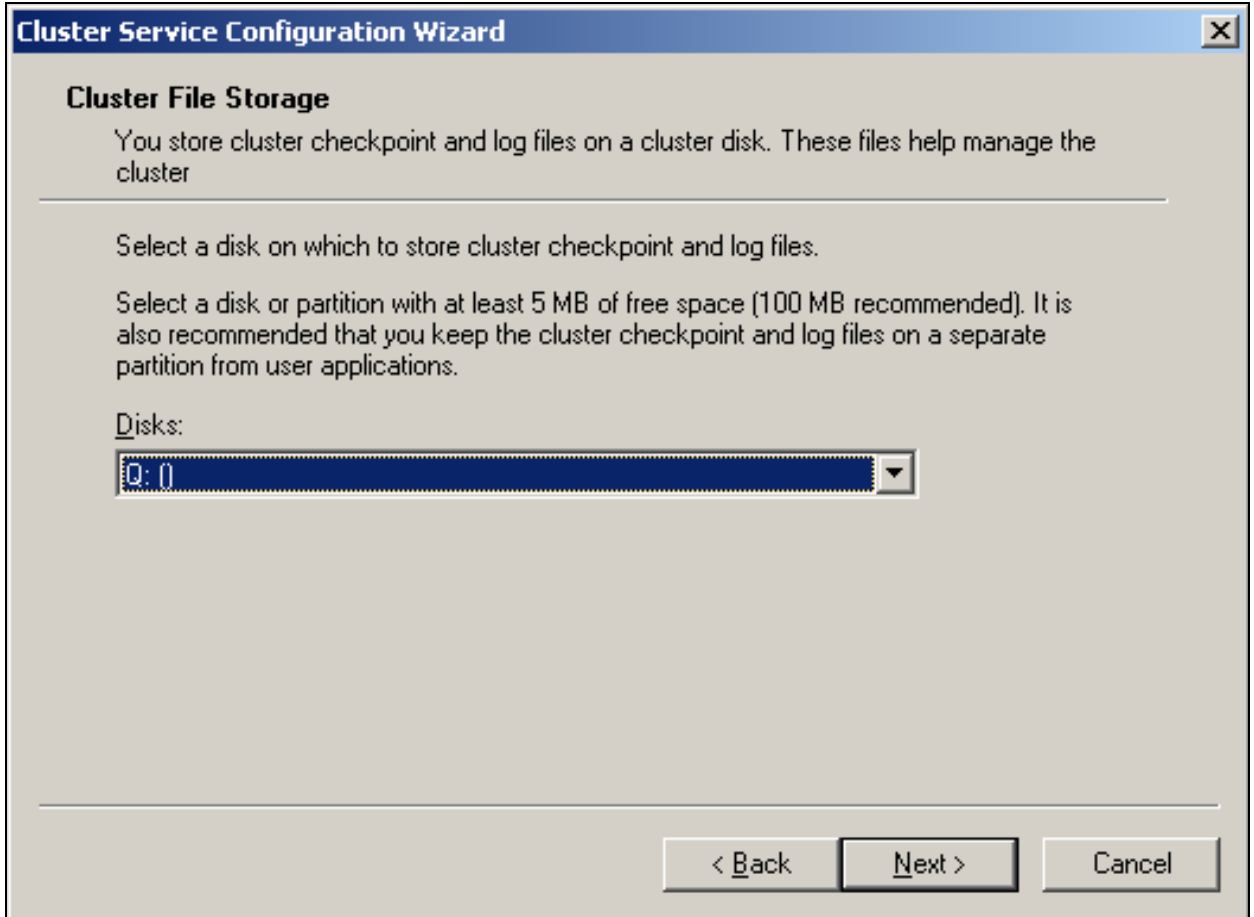


Figure 2. Cluster Administrator: Quorum Disk Selection

Note: Microsoft recommends using a private network for internal cluster communication between nodes.

The cluster IP address specifies a new IP address, not already in use on the network, to manage the cluster over the public network interface.

After configuration of node1 is complete, repeat the process on node2 by specifying the installation is for a second node in the cluster:

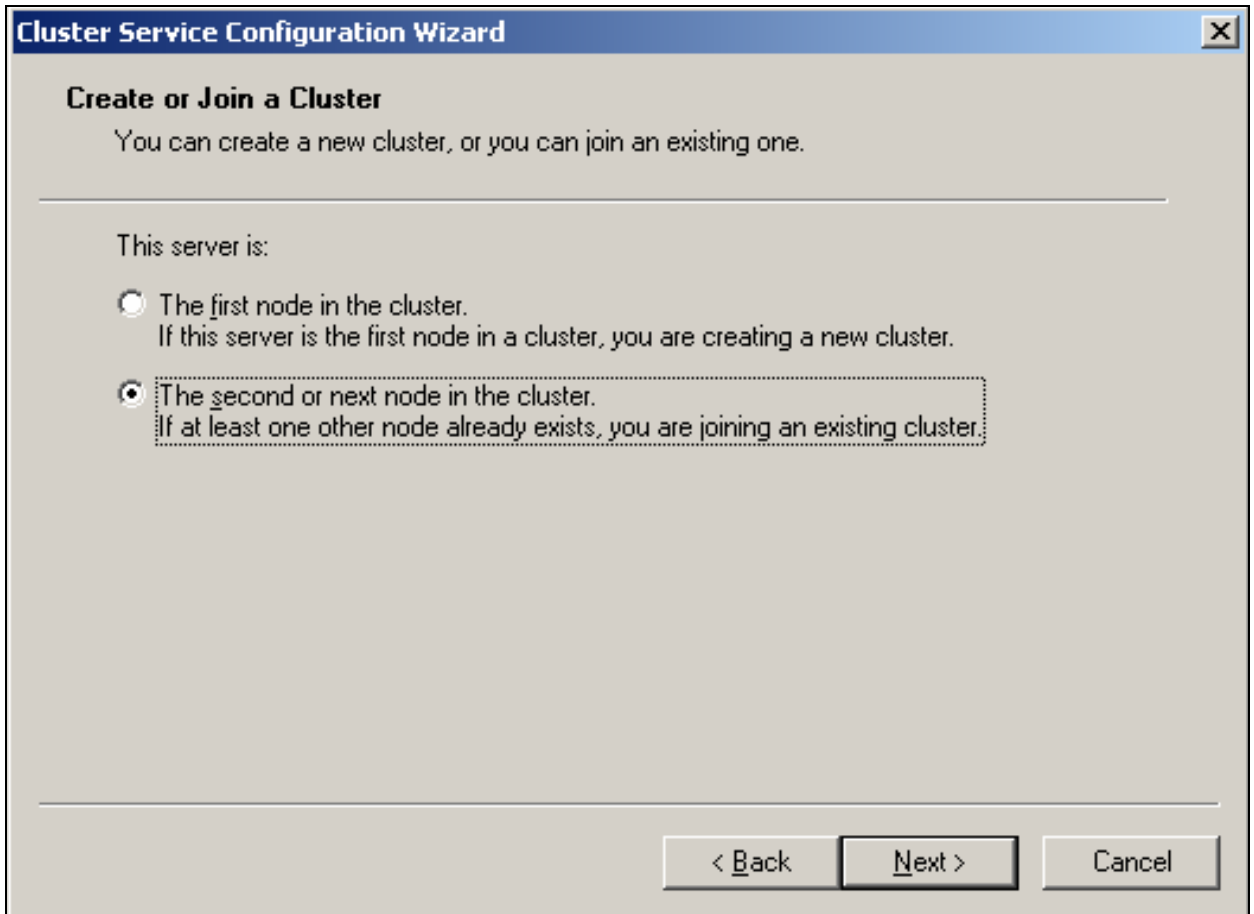


Figure 3. Cluster Administrator: Node2 Configuration

Specify the same cluster and “Connect to cluster as” names used during configuration of node1. Again, use the same domain account for the cluster service as that used in the configuration of node1.

When configuration of node2 is complete, both nodes should be displayed in Cluster Administrator and the quorum disk resource online:




Figure 4. Cluster Administrator: Cluster Group Resources

A second cluster group was created and, within it, a LUN (drive R:) to be used to store the SQL Server database. Refer to section 6.2 below for information about using multiple LUNs with your SQL Server Virtual Server installation.

Testing the MSCS Configuration

Several methods exist for testing successful configuration of the cluster. If a SQL Server 2000 database is not presently installed and stored on the filer, test failover of the MSCS cluster by right-clicking the cluster group and selecting “move group.” The cluster group and its resources should failover to the other node. In a more dramatic test, one could unplug power to the node presently owning the cluster resources. Again, the cluster group and its resources should failover to the surviving node.

SQL Server 2000

Installing a SQL Server Virtual Server

Before beginning installation, please read the installation documentation accompanying the SQL Server 2000 software. It is recommended in the documentation that some services (example. IIS) be temporarily stopped before installation of SQL Server 2000. Only 16 virtual servers can be defined for a failover cluster. For more information, see the references section at the bottom of this report.

Also, prior to installation it is recommended that the latest MDAC version be downloaded and installed on both nodes. As specified in its description on Microsoft's download page, there is an MDAC version containing specific updates for SQL Server/ODBC. Following reboot of both nodes, retest MSCS failover. When failover is successful between nodes, ensure node1 owns the cluster resources.

Install SQL Server 2000 as a Virtual Server on node1, specifying the name of a new virtual server. You may name the virtual server anything you would like:

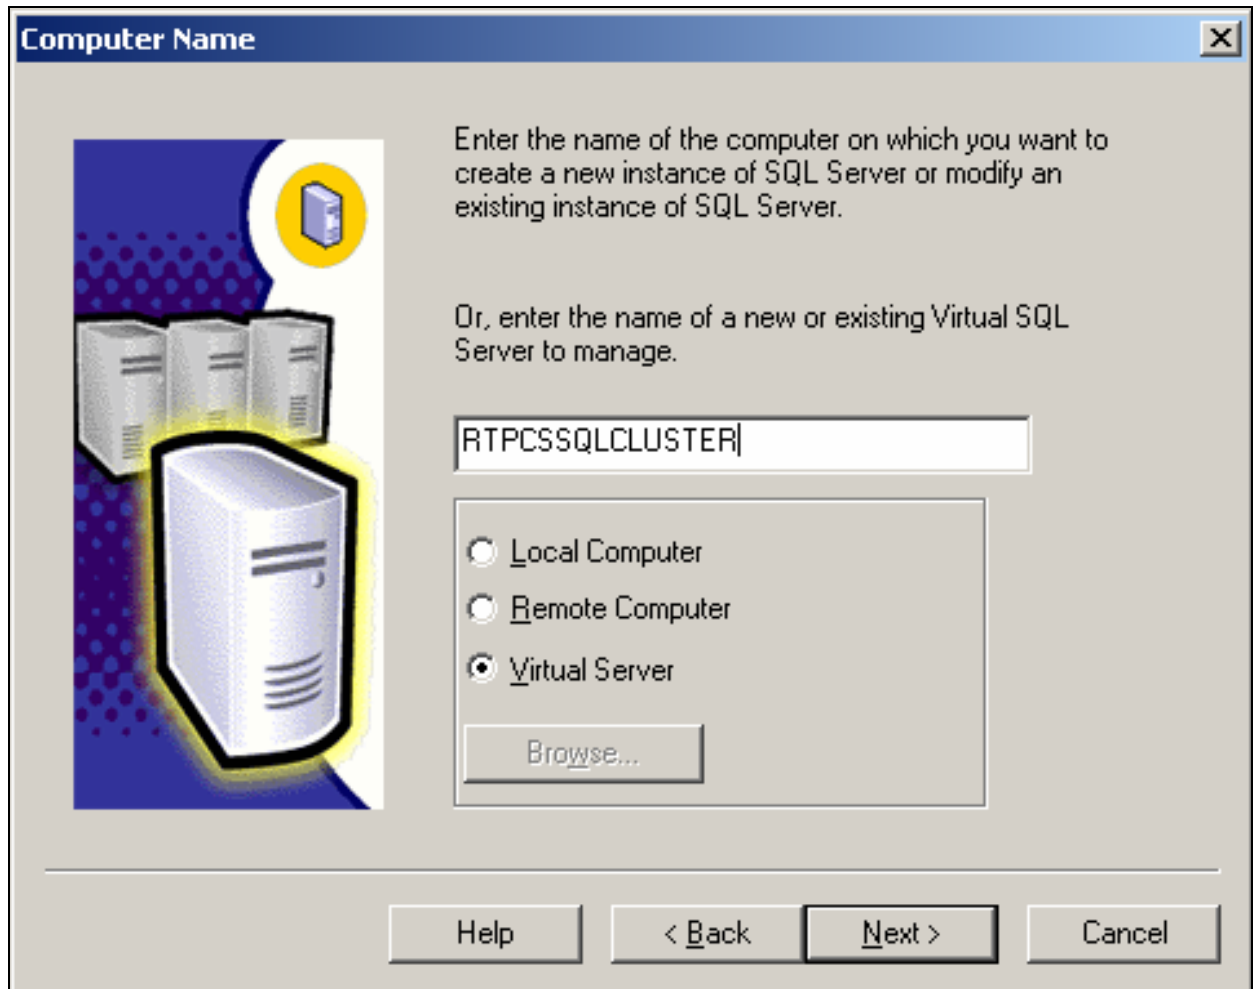


Figure 5. SQL Server: Installing a Virtual Server

The name entered here is not the name of the MSCS cluster group but a new name chosen to represent the SQL Server Virtual Server.

Enter an IP address not presently in use on the network to be used by applications/hosts to connect over the *public* network to connect to the virtual server:

The screenshot shows a 'Failover Clustering' dialog box with the following fields and controls:

- Virtual server name: RTPCSSQLCLUSTER
- IP address: [Empty field with dots for decimal separators]
- Subnet: [Empty field with dots for decimal separators]
- Network to Use: [Dropdown menu]
- Buttons: Add, Remove
- Table with columns: IPAddress, Subnet, Network
- Table content: 10.60.192.198, 255.255.255.0, PUBLIC
- Bottom buttons: Help, < Back, Next >, Cancel

IPAddress	Subnet	Network
10.60.192.198	255.255.255.0	PUBLIC

Figure 6. SQL Server: Virtual Server IP Address

Specify the LUN to be used to store the system databases. See section 6.2 for recommendations regarding the use of separate LUNs for storage of system and user database:

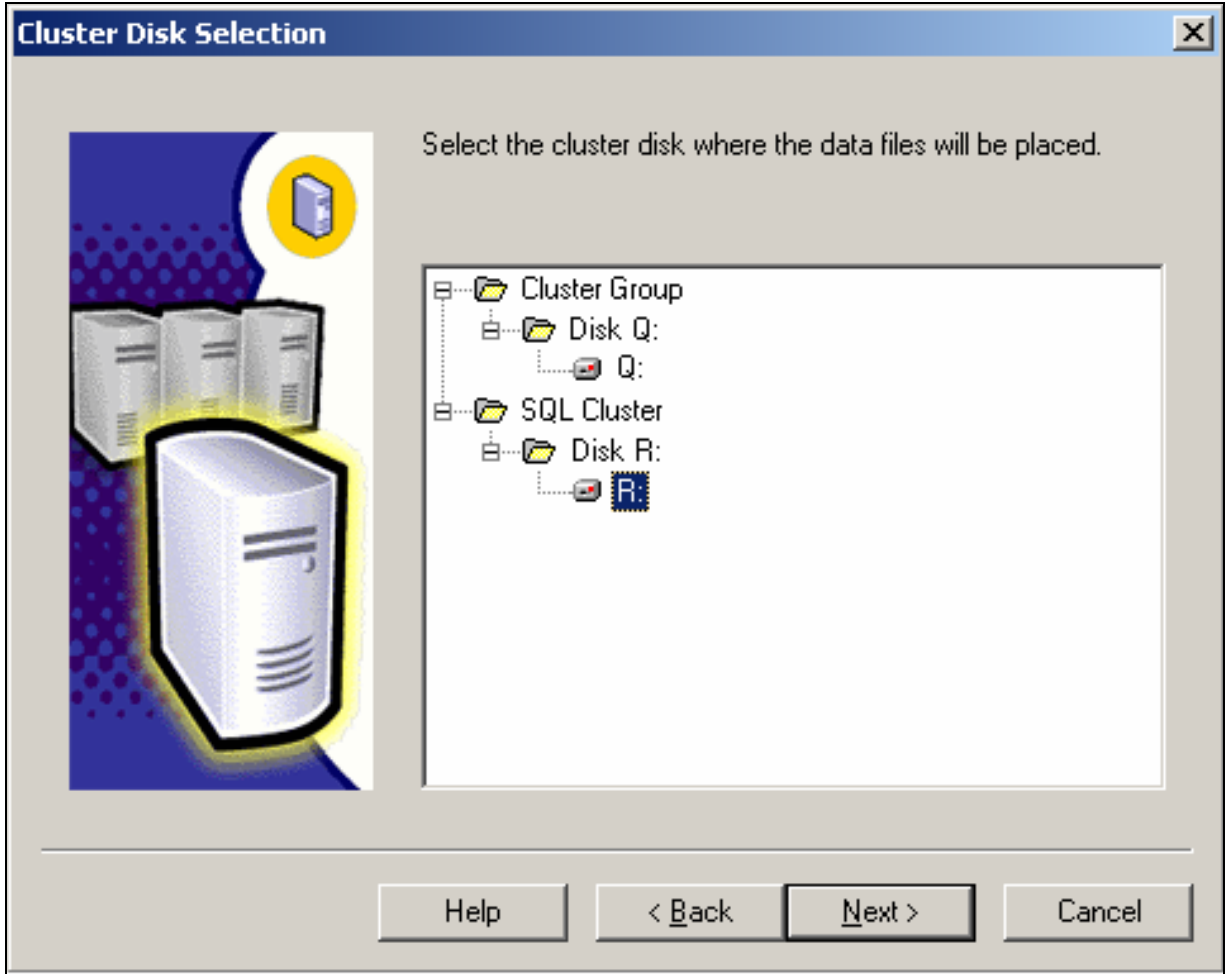


Figure 7. SQL Server: System Databases Disk

Install as a single-node installation by removing node2 from the list of configured nodes:

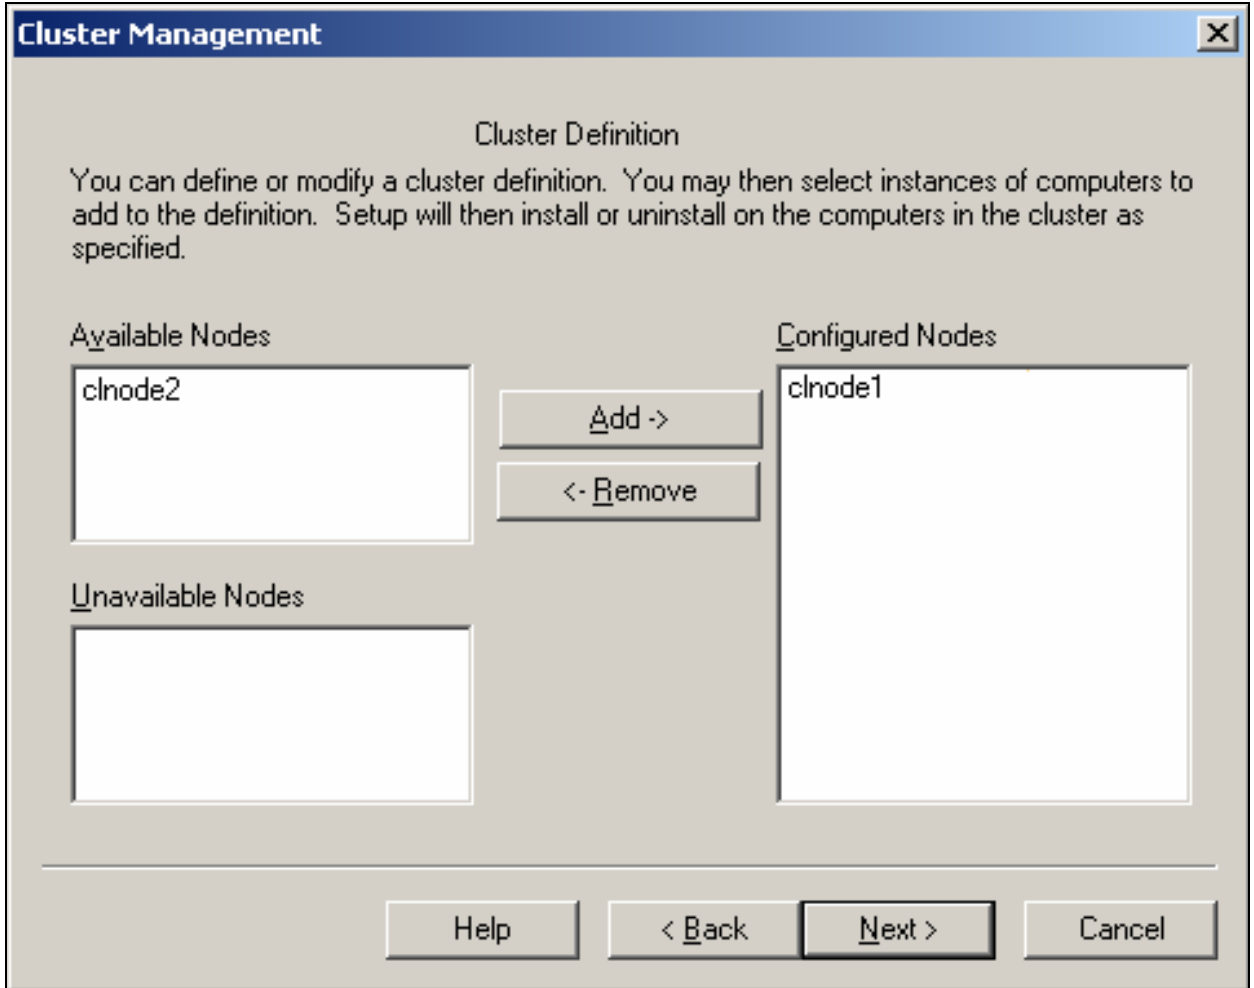


Figure 8. SQL Server: Single-Node Installation

For simplicity, specify the domain administrator account used during installation and configuration of MSCS. Next, you are presented with the option of a default installation or a named instance installation. Either is fine, but if you plan to install SQL Server in an active/active configuration, only one instance can be installed as a default instance. All other instances in the configuration must be installed as named instances.

Note: A SQL Server 2000 failover cluster cannot exist on the same server as a SQL Server 6.5 or 7.0 failover cluster. Also, you must specify a domain account with administrator privileges on both nodes for the SQL Server services (SQL Server, SQL Server Agent, and all cluster group services).

After installation on node1 is complete, you will see SQL Server Virtual Server resources added to the cluster group containing the disk specified during installation to store the database. At this point, you may choose to apply the latest SQL Server Service Pack.



Configuring Node2

To add node2 to the virtual server cluster, begin installation again on *node1*. No reboot is required. Specify Virtual Server and enter the name of the virtual server created during the first installation process. Select Advanced Options and then Maintain a Virtual Server for Failover Clustering.

Note: It is not necessary to specify the IP address of the virtual server since it is already configured on node1.

Add node2 to the list of configured nodes. Enter the domain administrator account used in other parts of the installation and configuration of the MSCS and SQL Server Virtual Server cluster. Again, a domain administrator account is required for the SQL Server services.

The installation process will then copy the SQL Server 2000 binaries to the C:\ drive of node2 and configure it for use in the cluster. At this point, installation of an active/passive configuration is complete.

If Windows Service Pack 3 is reapplied on both nodes following installation, reinstalling the latest MDAC version on both nodes is recommended.

Configuring an Active/Active Cluster

Installation of an active/active configuration uses the same process, beginning on node2. Before installing on node2, create a second cluster group with the necessary disk resources (LUNs) to be used by node2 for database storage. During installation of SQL Server on node2, you will need to specify a new Virtual Server, IP, disk for system database storage, and database instance name.

See section 6.2 for more information on active/active disk resource requirements and how to configure multiple disk resources for use by the Virtual Server.

Configuring Multiple Disk Resources

During SQL Server installation you will specify a default disk within a given cluster group to be used to store the system databases. For ease of maintenance and recovery, use separate LUNs for user-defined databases.

For example: Install the SQL Server system databases in a LUN. After installation, create your own databases in one or more LUNs separate from the LUN used to install the system databases.

Using multiple LUNs in a SQL Server Virtual Server cluster can be accomplished easily by creating additional LUNs in the cluster group containing the default SQL Server Virtual Server data disk.

If an active/active installation is planned, you must create two cluster groups, each containing the disk resources to be associated with a given node's instance.

By default, SQL Server Virtual Server installation makes only the disk specified to store the system databases a dependency of the SQL Server service. To enable Virtual Server use/failover of additional disk resources, use Cluster Administrator to make the SQL Server Service dependent on the other disks intended for use with a given node's instance.



For example: Create a cluster group containing multiple LUNs. Install SQL Server and specify this cluster group for the Virtual Server. Select one of the LUNs to be used for installation of the system databases. Following installation, go into Cluster Administrator and for each additional LUN to be used by the node for storage of user-defined databases:

Right-click the disk resource > offline (and/or move the disk resource to an appropriate cluster group).

Within a given cluster group, right-click the SQL Server resource > properties > dependencies tab > modify > add the disk resource.

In this way, SQL Server will not start the Virtual Server until the disks upon which it depends are available. When complete, bring the new disk resource online and retest for successful failover between nodes.

Failover and the Virtual Server

In an active/passive cluster configuration, each node is configured to be a member of the Virtual Server, but only the resource-owning node is active in its connection to the database. Node2 stands waiting to take over in the event node1 becomes unavailable.

If node1 fails, MSCS transfers ownership of the SQL Server Virtual Server resources to node2. Node2 then restarts the default instance. Applications continue to connect via the name of the virtual server.

In an active/active cluster configuration, each node has its own SQL Server Virtual Server and active instance. When one node fails, ownership of the failed node's SQL Server Virtual Server resources transfers to the surviving node. The surviving node then restarts the default instance of the failed node and applications continue access via their respective Virtual Server names.

At this point, the surviving node owns and is running two SQL Server instances: its own instance and the instance of the failed node.

Testing Virtual Server Failover

With its database stored on the filer, initiate an INSERT query from node1. While running the query, fail node1 either by:

1. Selecting Move Group from Cluster Administrator
2. Removing power from node1

After node2 acquires the Virtual Server resources of node1, use Enterprise Manager to verify the database running the INSERT query at the time of failover is not in suspect mode. Query the table to view the inserted rows.

Note: Remember to conduct these tests in a preproduction environment using production like workloads. Tests of this nature should never be done against production databases.

Use of filer option *nvfail* is recommended for all database environments. Please refer to the *Data ONTAP Administration Guide* or the NOW™ site for more information on the nvfail option.



Caveats

Material presented here has not been tested for all hardware and software platforms. Your environment may require modifications to attain a working configuration. If you experience problems making information presented in this report work, please contact IBM Support for assistance.



Trademarks and special notices

© International Business Machines 1994-2007. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Network Appliance, the Network Appliance logo, Data ONTAP, SnapDrive and Snapshot are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

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

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.