



Technical Report

High-Availability Support for DataFabric Manager Server

Jagatheesan Krishnamoorthy and Prasad S, NetApp
July 2009 | TR-3767

EXECUTIVE SUMMARY

This technical report explains how to set up high-availability solution for NetApp® DataFabric® Manager Server. The high-availability solution for DataFabric Manager Server on Windows® is implemented using Microsoft® Cluster Server (MSCS), while a high-availability solution for DataFabric Manager Server on Linux® and Solaris™ server is implemented using Veritas™ Cluster Server (VCS).

Customers who want a high-availability solution for DataFabric Manager Server can refer to this document. The solution described uses Operations Manager, SnapDrive®, MSCS (on Windows server), and VCS (on Linux and Solaris servers).

TABLE OF CONTENTS

1. INTRODUCTION	3
1.1 ABOUT MICROSOFT CLUSTER SERVER	3
1.2 ABOUT VERITAS CLUSTER SERVER	3
1.3 ABOUT SNAPDRIVE	3
2. HIGH-AVAILABILITY SUPPORT FOR DATAFABRIC MANAGER SERVER ON WINDOWS SERVER 4	
2.1 DEPLOYMENT SCENARIO	4
2.2 REQUIREMENTS AND LIMITATIONS FOR MSCS	5
2.3 SETTING UP HIGH AVAILABILITY FOR DATAFABRIC MANAGER SERVER USING MSCS	5
3 HIGH-AVAILABILITY SUPPORT FOR DATAFABRIC MANAGER SERVER ON LINUX OR SOLARIS SERVER	51
3.1 DEPLOYMENT SCENARIO	51
3.2 REQUIREMENTS AND LIMITATIONS FOR VCS	52
3.3 SETTING UP HIGH AVAILABILITY FOR DATAFABRIC MANAGER USING VCS	52
4 GLOSSARY	62

1. INTRODUCTION

As DataFabric Manager Server is used for day-to-day storage management activities, the availability of DataFabric Manager Server services is crucial. In case of DataFabric Manager Server failure, customers should be able to access the DataFabric Manager services quickly from another host with minimal interruption. This document explains how to set up a high-availability solution for DataFabric Manager Server.

This document is meant for the customers, who want to set up a high-availability solution for DataFabric Manager Server. This document describes how to set up:

- High-availability solution DataFabric Manager Server on Windows Server using MSCS:
This solution is available from DataFabric Manager Server 3.3 release.
- High-availability solution for DataFabric Manager Server on Linux and Solaris Servers using VCS:
This solution is available from DataFabric Manager Server 3.7 release.

To understand this technical report, you should understand the concepts and working of Operations Manager, Data ONTAP®, SnapDrive, MSCS, and VCS.

Note: DataFabric Manager Server is the management server that hosts the services and data required for Operations Manager, Protection Manager, Performance Advisor, and Provisioning Manager.

1.1 ABOUT MICROSOFT CLUSTER SERVER

Microsoft Cluster Server provides failover support for applications and services that require high availability, scalability and reliability. With clustering, organizations can make applications and data available on multiple servers linked together in a cluster configuration. Back-end applications and services, such as those provided by database servers, are ideal candidates for Microsoft Cluster Server.

Note: Microsoft Cluster Server is called as failover cluster in Windows 2008 Server

1.2 ABOUT VERITAS CLUSTER SERVER

Veritas Cluster Server is the industry's leading cross-platform clustering solution for minimizing application downtime. Through central management tools, automated failover features to test disaster recovery plans without disruption, and advanced failover management based on server capacity, Cluster Server allows IT managers to maximize resources by moving beyond reactive recovery to proactive management of application availability.

1.3 ABOUT SNAPDRIVE

SnapDrive software is an optional management package for Microsoft Windows and some UNIX® hosts. SnapDrive simplifies the management and data protection tasks associated with iSCSI and FCP storage.

SnapDrive for Windows software integrates with the Windows Volume Manager so that storage systems can serve as storage devices for application data in Windows environments. SnapDrive manages LUNs on a storage system, making the storage available as local disks on Windows hosts. This allows Windows hosts to interact with the LUNs as if they belonged to a directly attached disk array.

SnapDrive for UNIX is a tool that simplifies data backup management so that you can recover should data be accidentally deleted or modified. SnapDrive for UNIX uses Snapshot™ technology to create an image of the data stored on a storage system attached to a UNIX host. You can then restore that data at a later time. In addition, SnapDrive for UNIX lets you provision storage on the storage system. SnapDrive for UNIX provides a number of storage features that enable you to manage the entire storage hierarchy, from the

host-side application-visible file down through the volume manager to the storage-system-side LUNs providing the actual repository.

2. HIGH-AVAILABILITY SUPPORT FOR DATAFABRIC MANAGER SERVER ON WINDOWS SERVER

This section describes how to set up and configure a high-availability solution for DataFabric Manager Server on Windows Server. The section also lists the hardware and software required and the steps to be followed to configure DataFabric Manager Server for high-availability.

2.1 DEPLOYMENT SCENARIO

In the example below, two independent Windows DataFabric Manager Servers are connected to a shared LUN on a NetApp storage system through iSCSI or FCP protocol. The DataFabric Manager Server data resides on the shared LUN. Microsoft Cluster Service has been installed and configured on both the DataFabric Manager Servers. A virtual IP is configured for the cluster using which a client can access the DataFabric Manager Server data.

In case of a system failure or repeated DataFabric Manager Server service failure on the active HA node, cluster server will automatically failover to the other node.

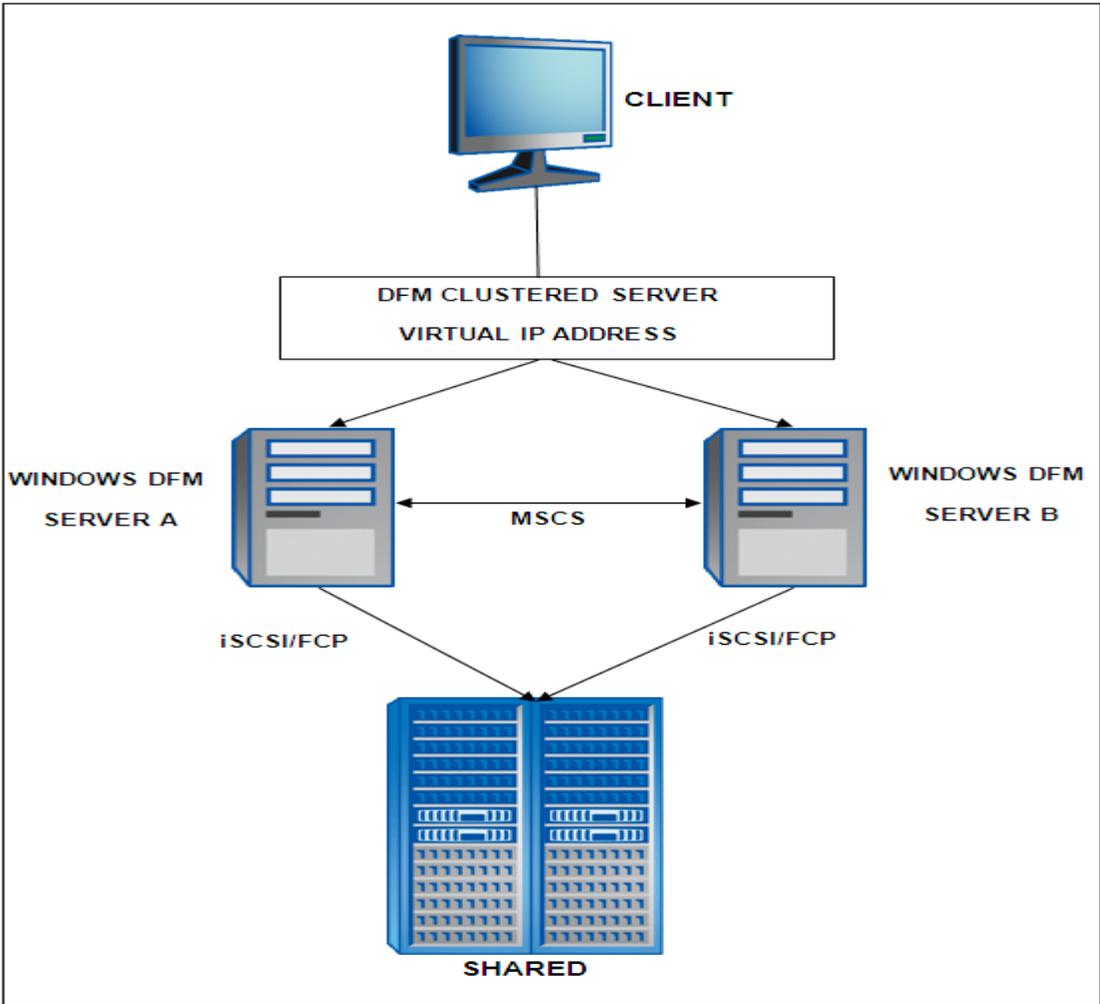


Figure 1 - MSCS deployment scenario

2.2 REQUIREMENTS AND LIMITATIONS FOR MSCS

This section describes the hardware and software requirements and limitations for MSCS

- Microsoft Windows Servers running Windows 2003 or Windows 2008 Enterprise edition or Data Center edition, running the same patch version of identical hardware.
- The DataFabric Manager Servers should be SAN capable, with either iSCSI or FCP connectivity to a storage system.
- Only two node clustering is supported.
- The members of the cluster should not be a domain controller. They Should be only member servers.
- The same version of DataFabric Manager Server should be running on both the servers. Also the DataFabric Manager Server should be installed at the same location on both the servers (say default location C:\Program Files\NetApp\DataFabric\DFM).
- The DataFabric Manager Server administrator should be domain user rather than local system user, so that the user login works even when the DataFabric Manager Server services fail over to the partner node.
- Fibre Channel switched fabric or iSCSI-based storage for shared data disks with a NetApp storage system as the storage back end.
- Only a single quorum disk cluster set up is supported on Windows 2003.
- Only “node and disk majority” quorum type is supported in failover cluster(Windows 2008 server).
- Ipv6 virtual IP address is supported only on failover clustering(Windows 2008).
- SnapDrive for Windows though an optional requirement for high-availability solution, helps you to manage the LUN easily. It also helps in creating host-consistent backups. LUN management can also be done using storage system command-line interface or FilerView®. Refer to the SDW compatibility matrix at <http://now.netapp.com/matrix/mtx/login.do>.
- In Windows 2008(64bit) server 64bit perl to be installed for running DataFabric Manager Server configuration scripts.

2.3 SETTING UP HIGH AVAILABILITY FOR DATAFABRIC MANAGER SERVER USING MSCS

To set up high-availability for DataFabric Manager Server using MSCS, the user needs to do the following high level steps:

- Install same version of SnapDrive for Windows (optional) on both the cluster nodes.
- Create a shared LUN on NetApp storage system for quorum disk and another shared LUN for holding DataFabric Manager Server data.
- Create a new cluster group and configure the group with cluster name, virtual IP address, quorum disk, and data disk.
- Add the data disk to the cluster group.
- Install DataFabric Manager Server on both the cluster nodes to the default location or same location.
- Configure MSCS using DataFabric Manager Server configuration scripts.

The section below explains the steps to set up high-availability for DataFabric Manager Server using MSCS.

2.3.1 SNAPDRIVE FOR WINDOWS INSTALLATION

This section explains the steps to install SnapDrive for Windows. This section can be skipped if faster database backup of DataFabric Manager Server data solution and LUN management through SnapDrive for Windows are not required.

- Install SnapDrive for Windows as mentioned in “Install Upgrade Administration Guide” available at <http://now.netapp.com/NOW/knowledge/docs/snapdrive/relsnap60/pdfs/admin.pdf> .

Note: Even though LUN creation and management can be done without using SnapDrive, it is recommended to use SnapDrive as it simplifies LUN management.

2.3.2 QUORUM AND DATA DISK CREATION

This section explains how to create shared LUNs (quorum disk, data disk) on the NetApp storage system and how to add the data disk in cluster administrator. The section explains the disk creation with and without SnapDrive for Windows.

Note: Before proceeding to below steps, make sure that the cluster nodes having either software/hardware iSCSI (for iSCSI LUNs) installed or HBA card (for FCP).

Creating Quorum, Data Disk Using SnapDrive

- Launch Computer Management wizard on one of the cluster nodes. As shown in Figure 2, right-click Disks to launch Create Disk Wizard.

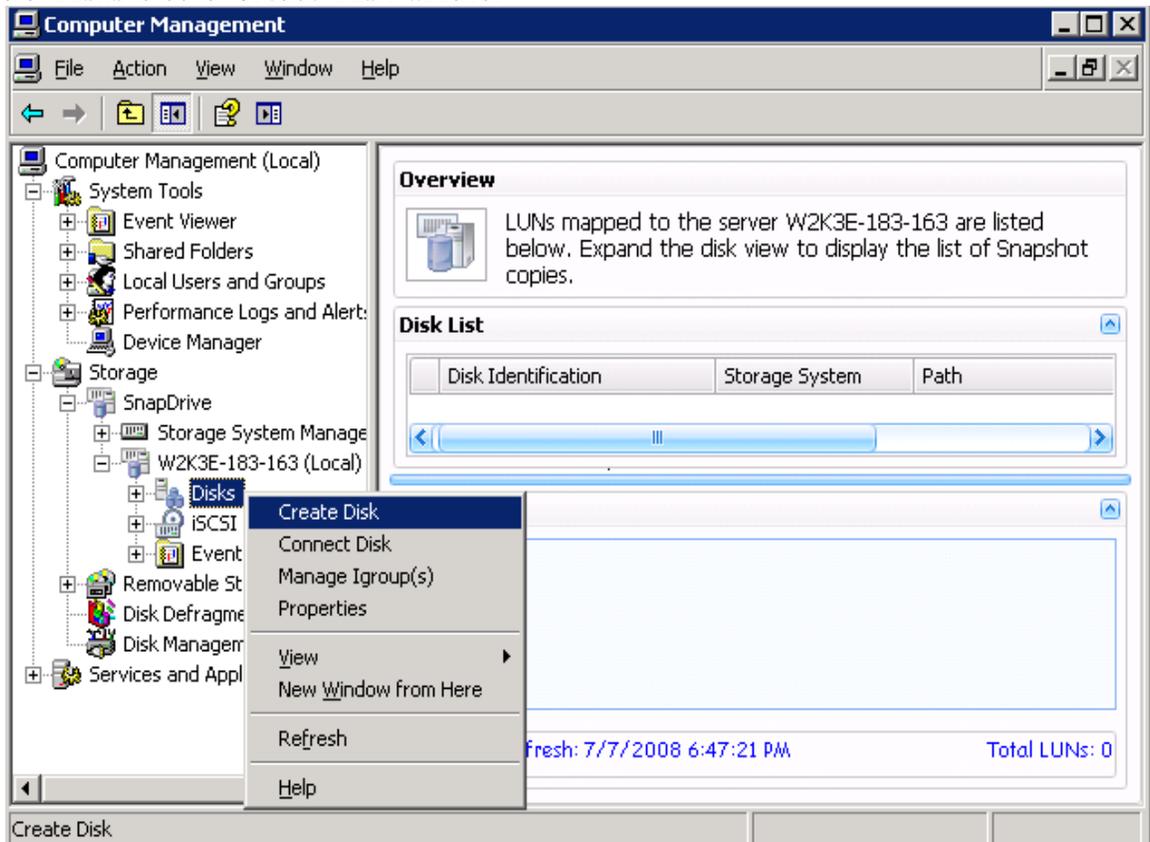


Figure 2 - Launch create disk wizard

- In the Create Disk Wizard, enter the storage system IP address and select a volume or qtree under which the shared disk is to be created. Enter the LUN name as shown in Figure 3.

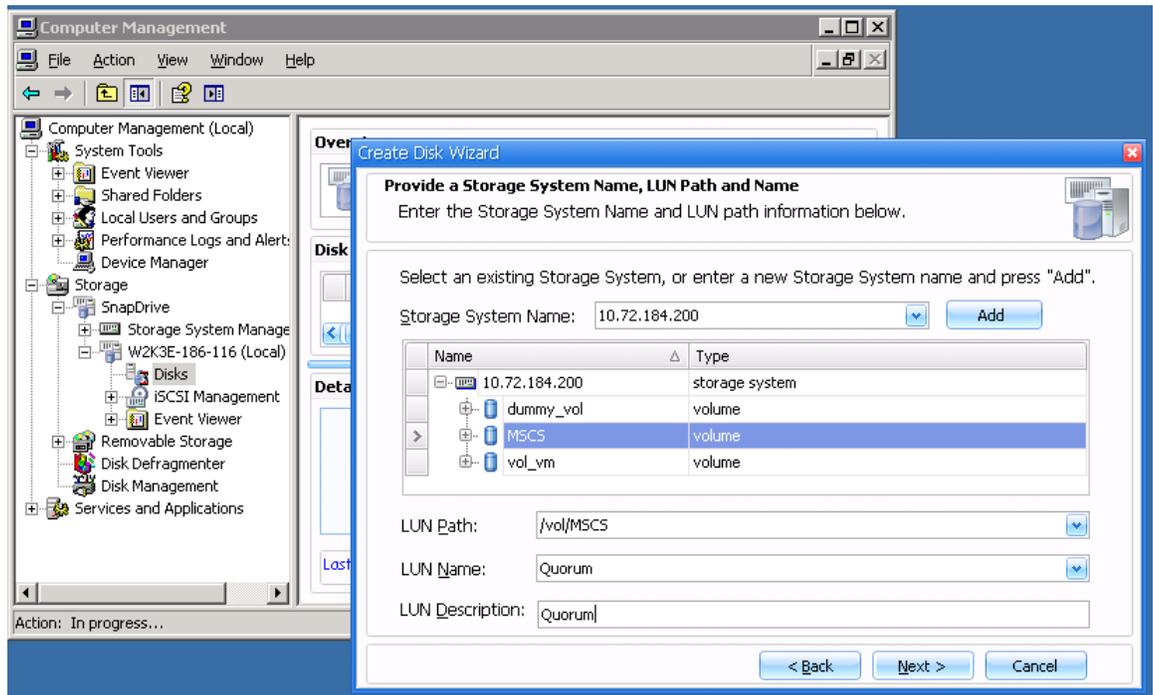


Figure 3 – Create disk wizard

- Select the LUN type as Shared and press Next.
- Select the drive letter and LUN size as shown in Figure 4.

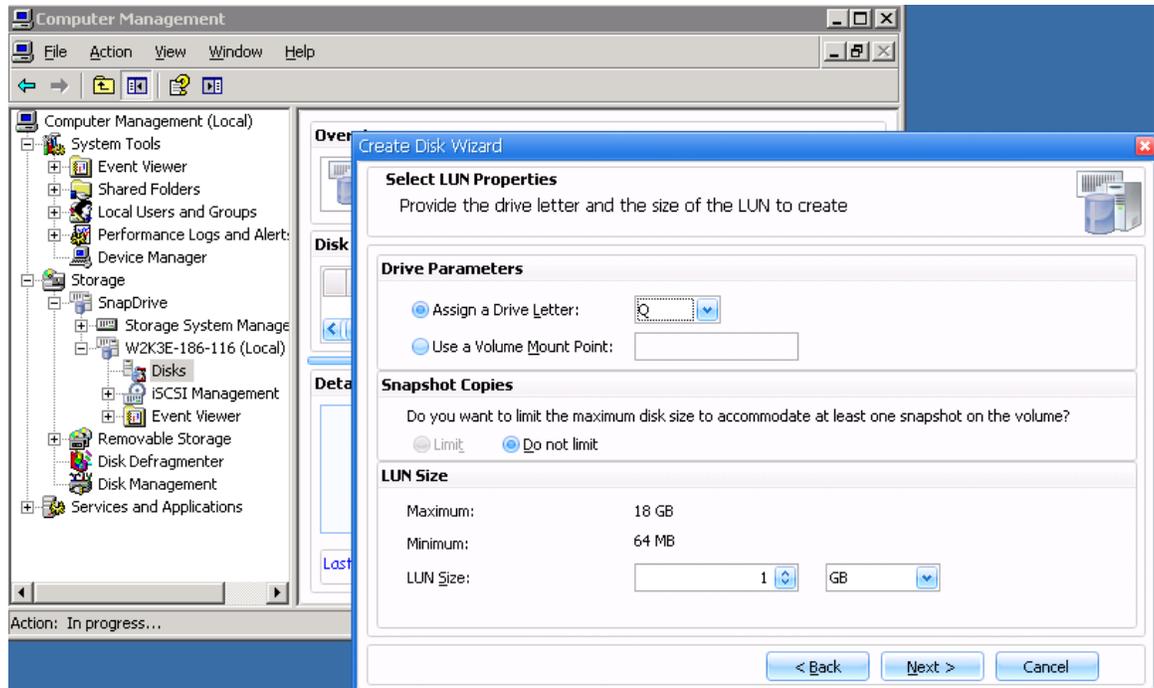


Figure 4 – Enter disk parameters

Note: The recommended quorum disk size should be a minimum of 1GB.

- Select the initiators in the list and press Next. Continue till the end to create a shared disk.
Note: Select the IQN for iSCSI LUN and WWPN for the FC LUN
- Repeat the same steps and create another shared disk. It will act as a data disk to hold DataFabric Manager Server data

Creating Quorum, Data Disk Without SnapDrive

- If you are not using SnapDrive, refer to the steps mentioned in the below link to create a shared data disk. <https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb17762>

Note: The disks should be mapped to the same drive letters on both the cluster nodes.

2.3.3 CLUSTER CONFIGURATION

2.3.3.1 CLUSTER CONFIGURATION ON WINDOWS 2003

This section explains how to create a cluster group with cluster name, virtual IP, quorum disk, and data disk.

Note: Make sure that the cluster nodes are added to an Active Directory domain

- Log in as a domain user to the cluster node1 and launch Cluster Administrator.
- Select Create new cluster in the Open Connection to Cluster window as shown in Figure 5.

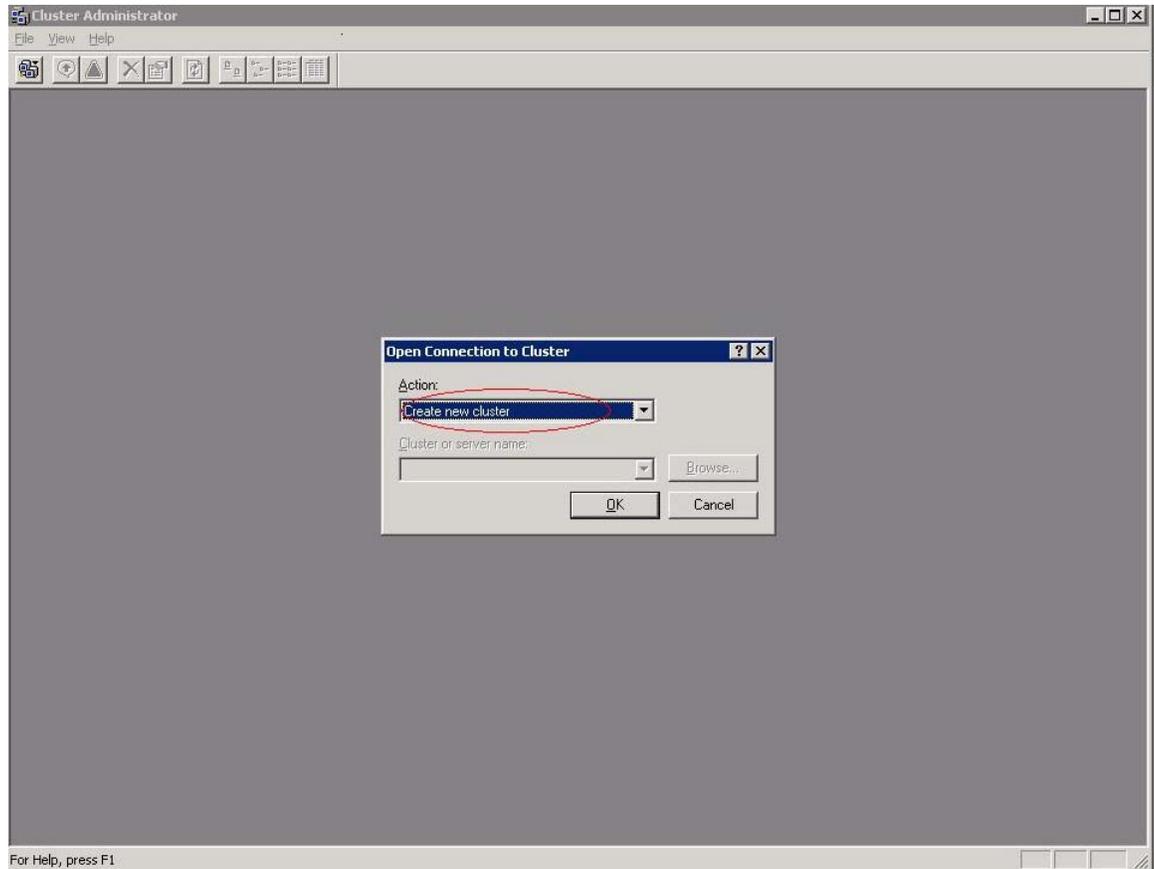


Figure 5 – Create new cluster wizard

- Select the domain name (both the cluster nodes should be in this domain) and enter a cluster name as shown in Figure 6.

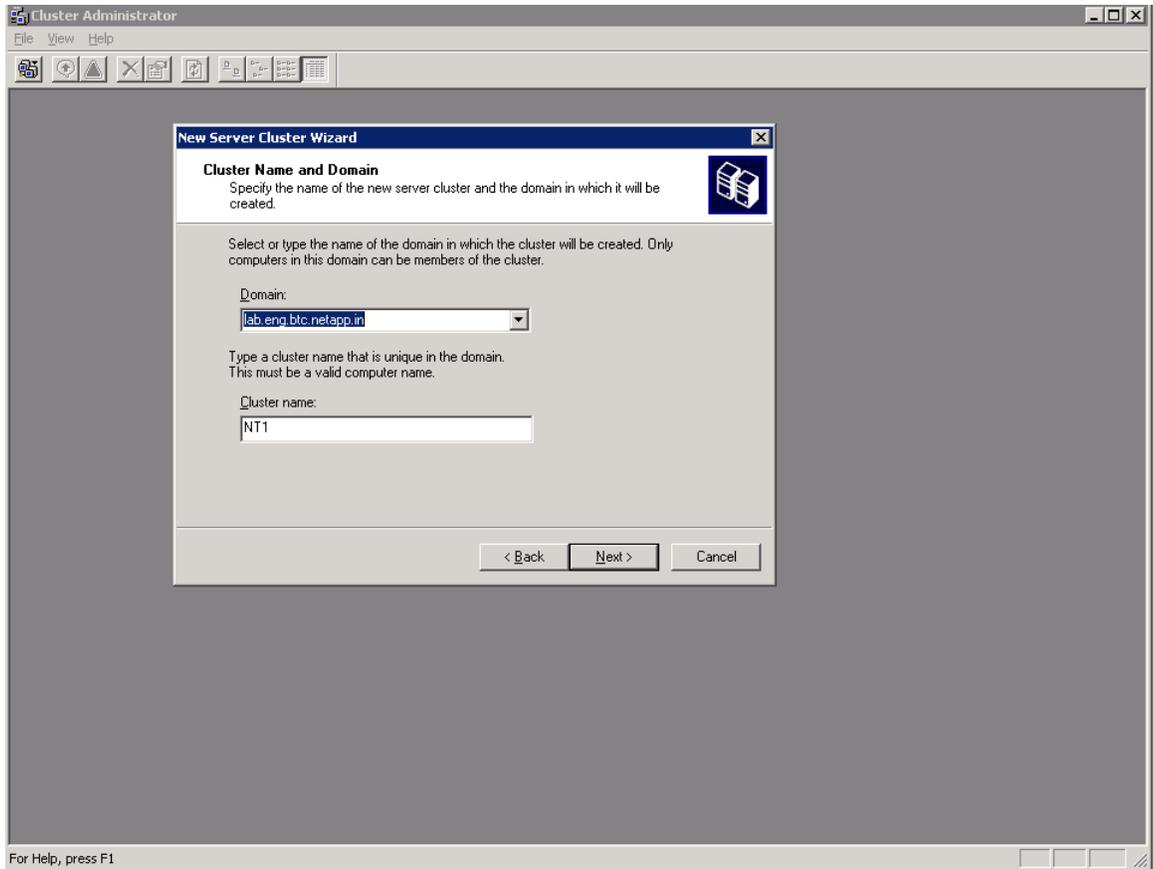


Figure 6 – Enter cluster name and domain

- Enter the host name of any of the cluster nodes as shown in Figure 7.

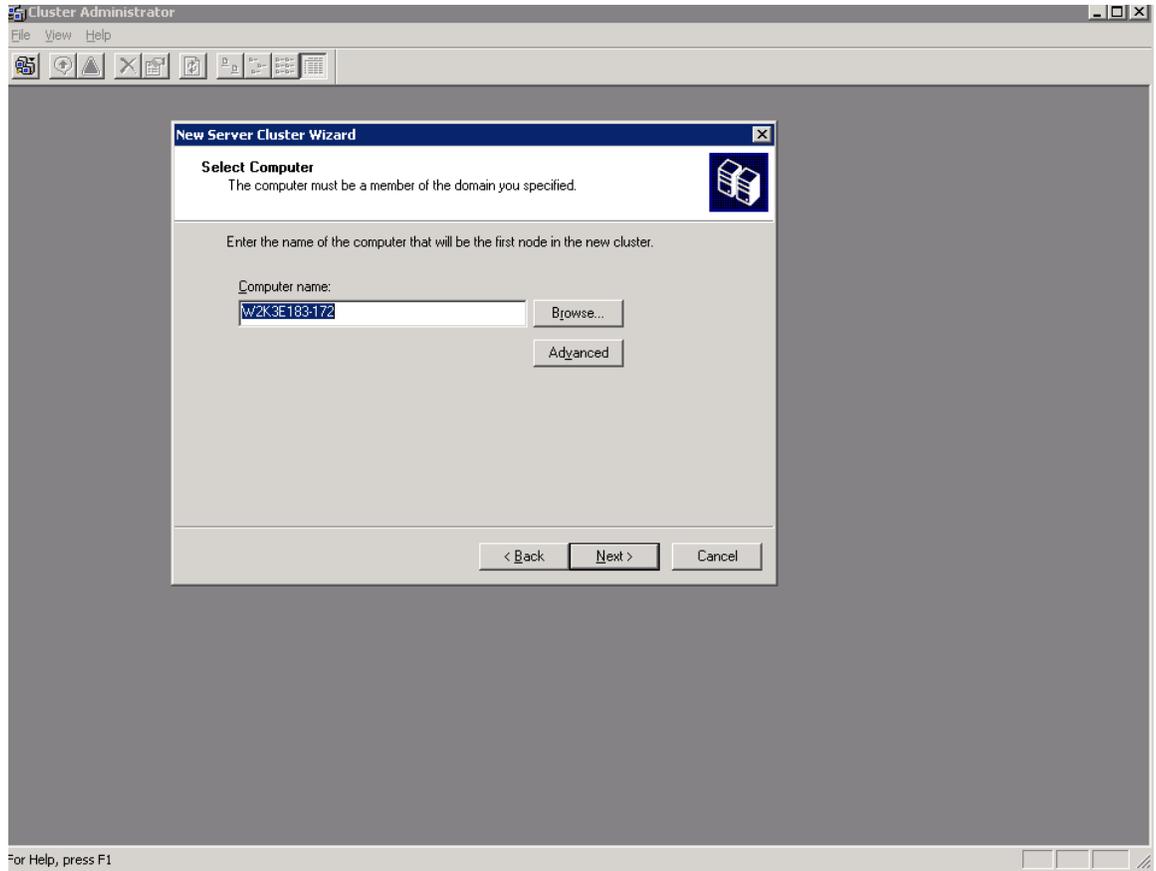


Figure 7 – Add cluster node

- Check that the "Analyzing Configuration" succeeds as shown in Figure 8.

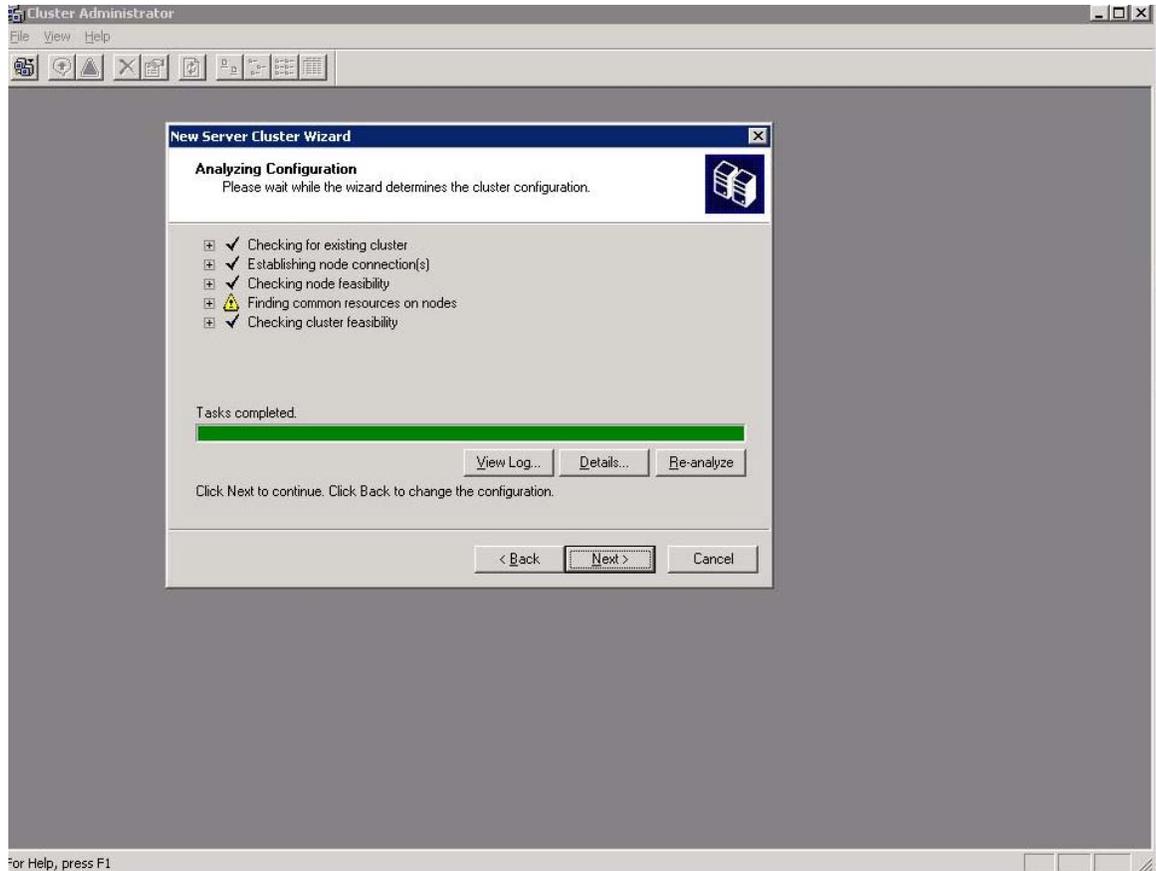


Figure 8 – Analyze cluster configuration

Note: If there is some error in the cluster configuration steps, it will be indicated in red. Expand the failed operation to know about the reason for failure. Once the problem is fixed, reanalyze the configuration by clicking the “Re-analyze” button.

- Proceed to next page after the “Analyzing Configuration” test is successful. Enter the virtual IP as shown in Figure 9.

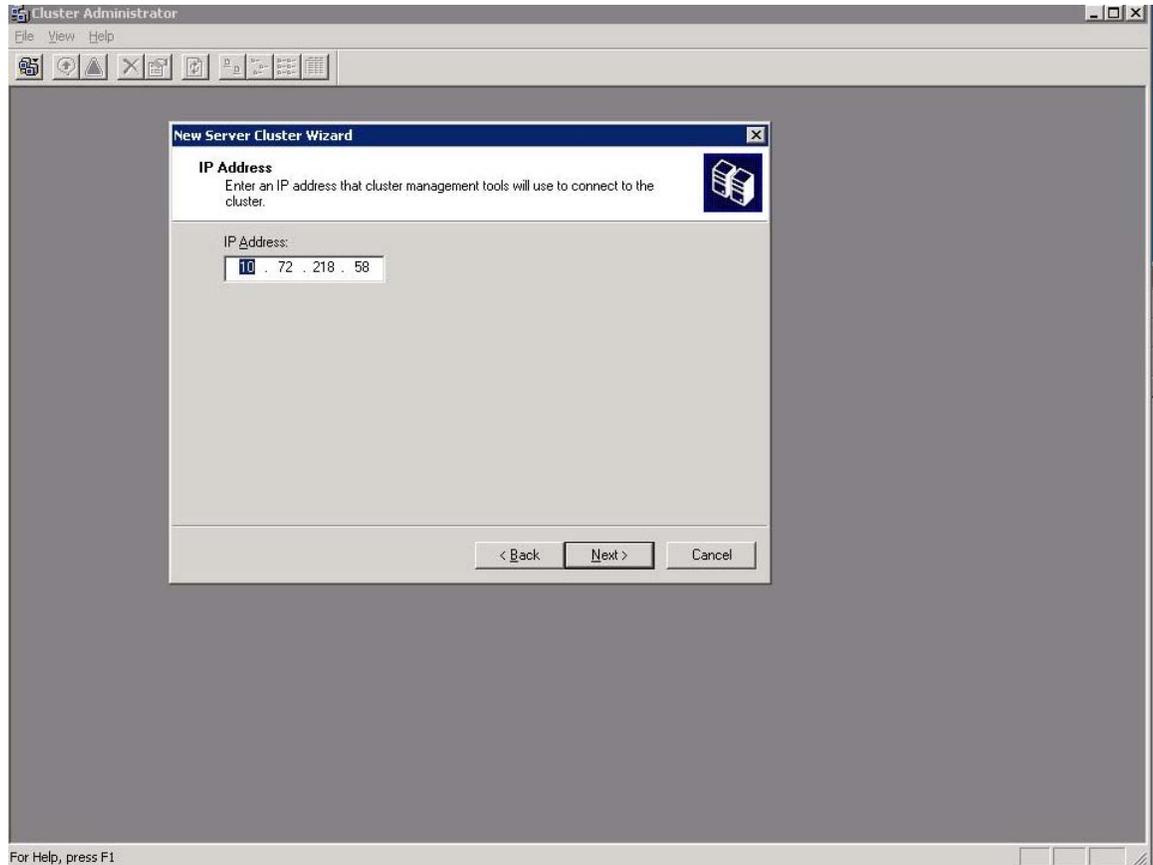


Figure 9 – Enter cluster IP address

Note: Make sure that the virtual IP is in the same subnet as the cluster nodes. Operations Manager and NetApp Management Console use the virtual IP to access the DataFabric Manager server.

- Enter the domain user name and password and select the domain in the next page as shown in Figure 10.

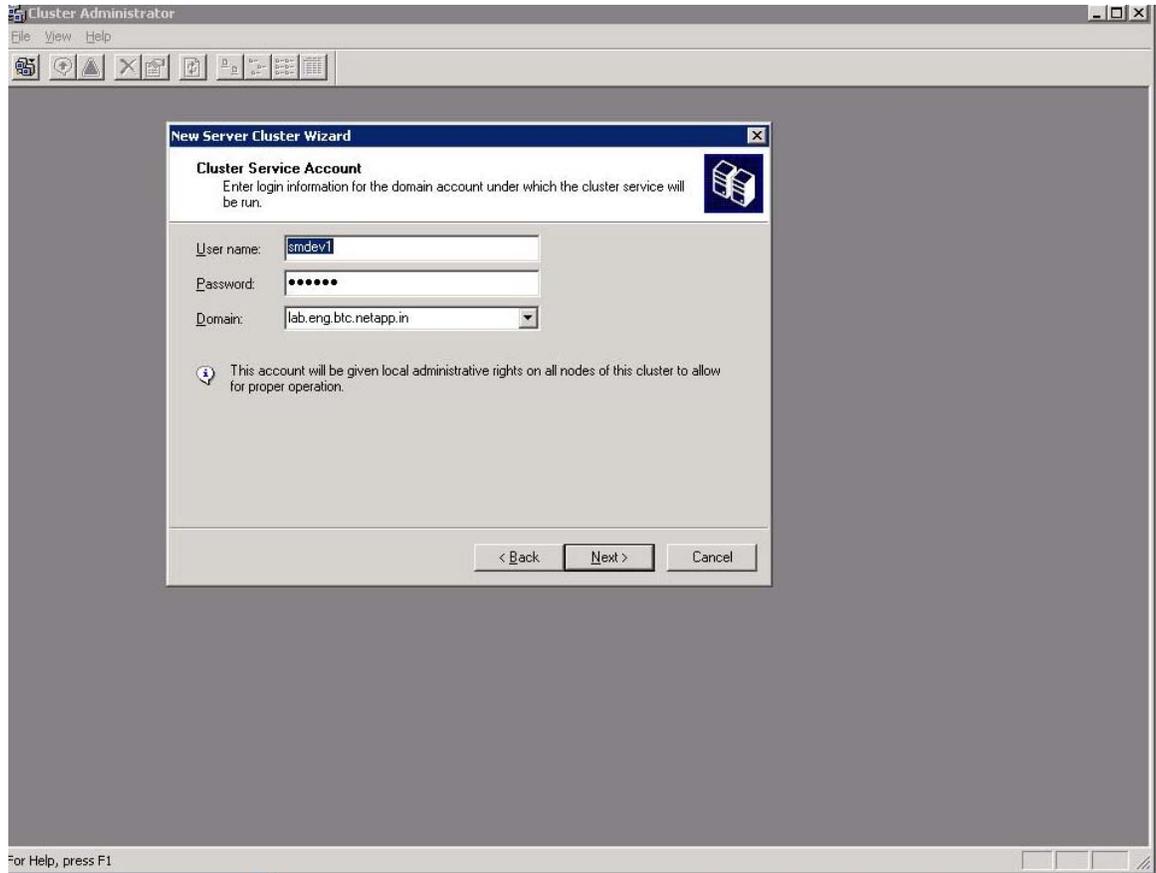


Figure 10 – Enter domain account details

- Select a shared disk for quorum as shown in Figure 11.

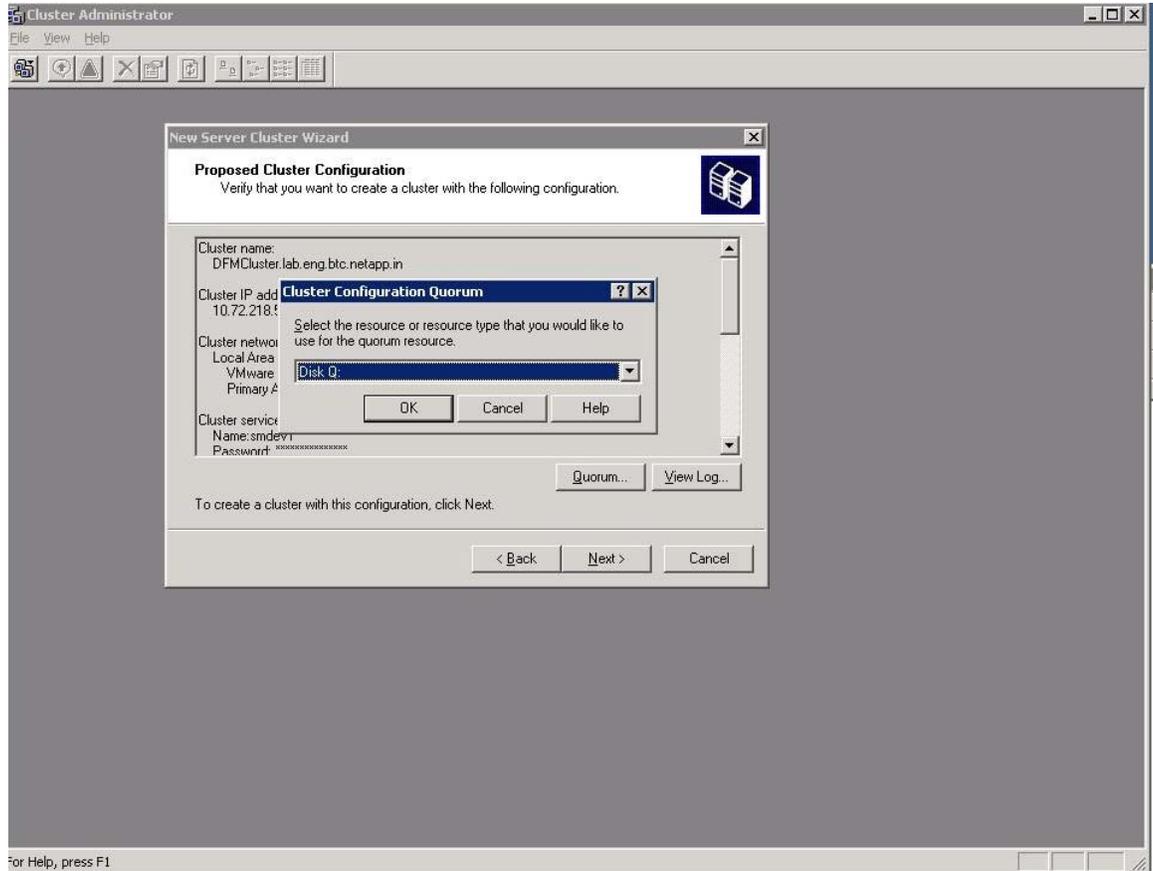


Figure 11 – Select quorum disk

Note: The Cluster Configuration Quorum dialog box lists both the quorum as well as data disk. Make sure that you select the quorum disk.

- Press Next to complete the wizard. The cluster administrator will show the cluster details of the newly created cluster as shown in Figure 12.

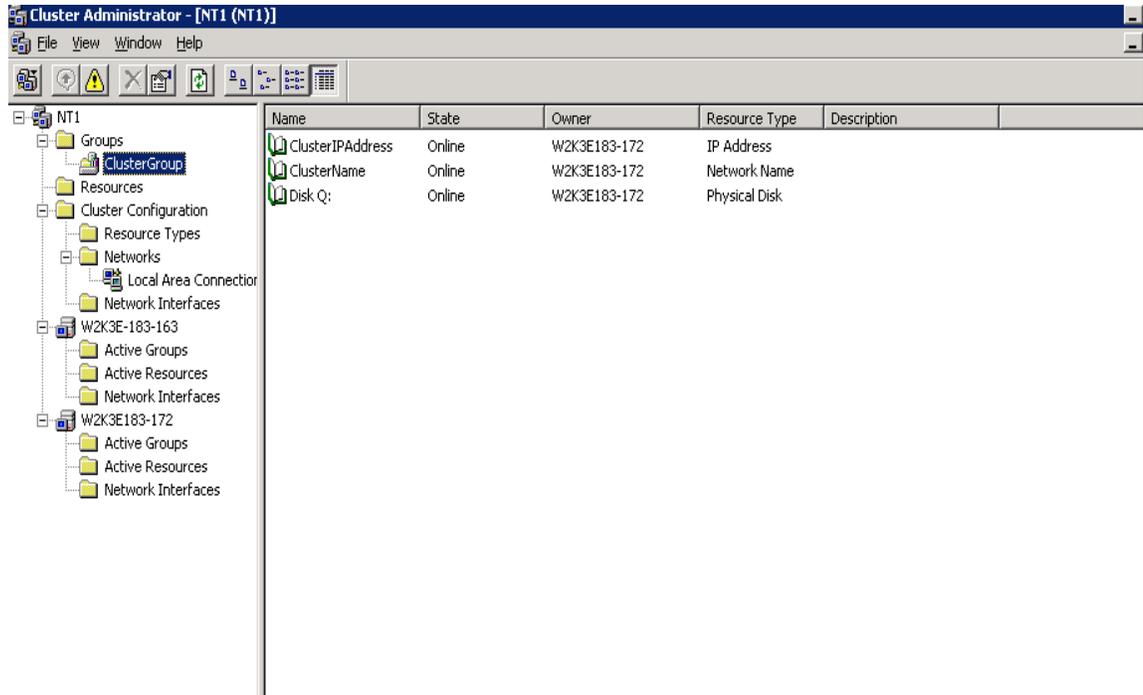


Figure 12 – Cluster administrator window

- Select the cluster group and check if cluster name, cluster IP address, and quorum disk are reported correctly.

As of now only one node is part of the cluster. To add the second node to the cluster, follow the below steps.

- Click File → Open Connection to launch the “Open connection to cluster” wizard. Select the action add nodes to the cluster as shown in Figure 13. Specify the name of the cluster created earlier.

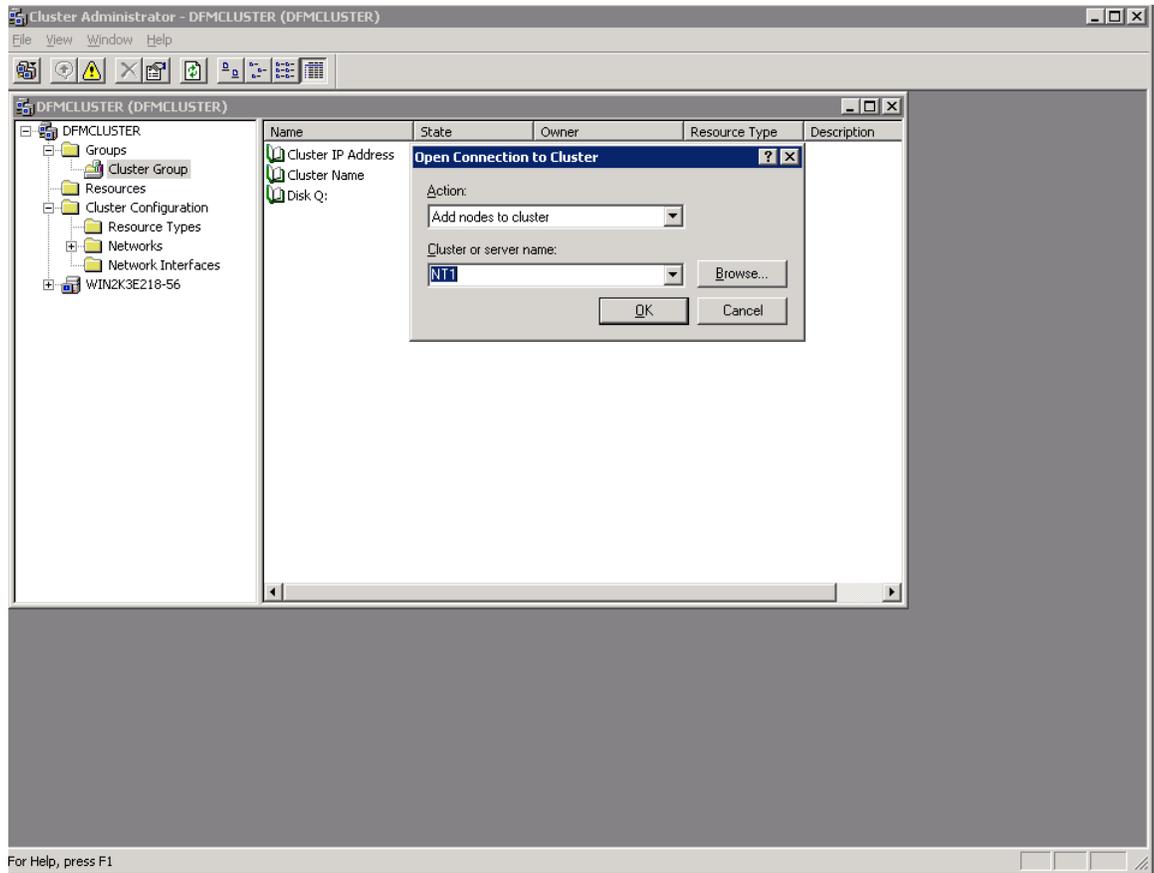


Figure 13 – Add second node to cluster

- Enter the name of the second cluster node as shown in Figure 14.

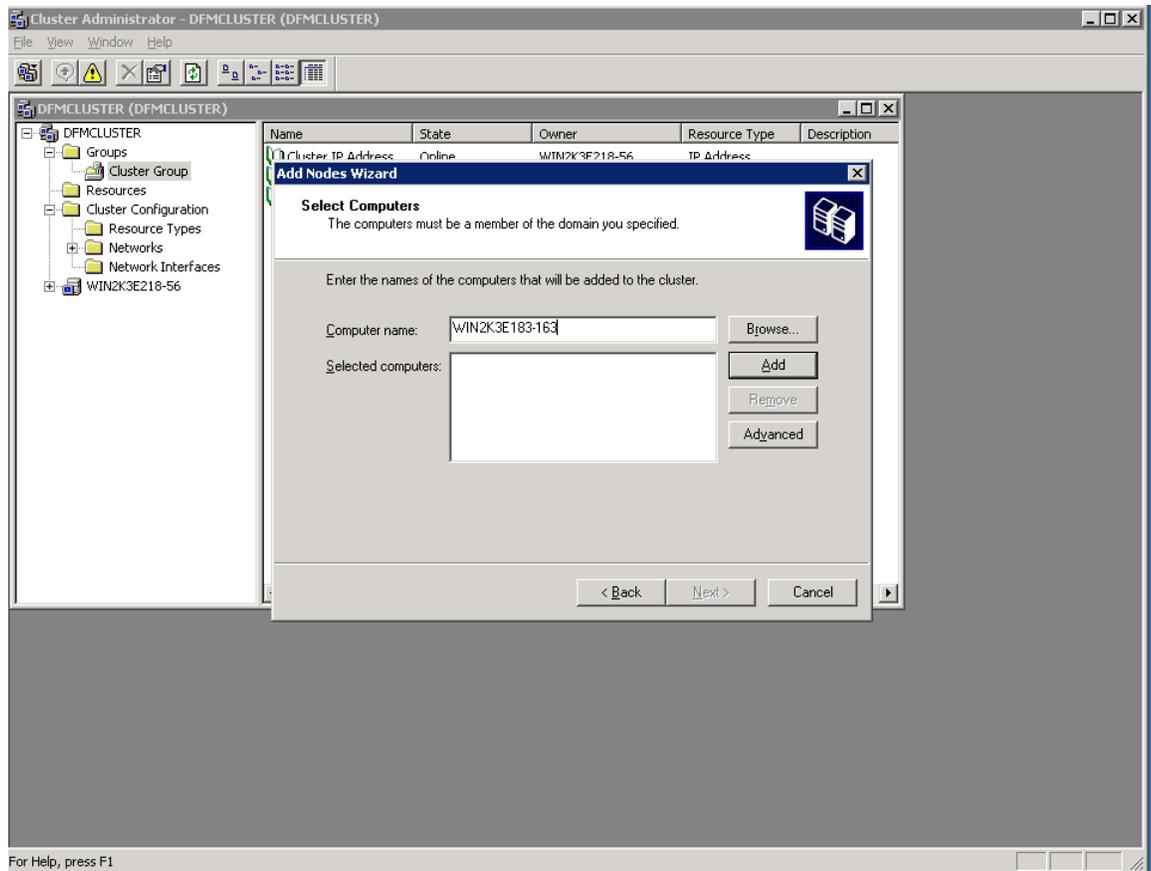


Figure 14 – Enter second node

- Clicking the Next button will take you to the Analyze Configuration wizard as shown in Figure 6. If there are no errors, proceed to complete the wizard.

Once the cluster is configured with the nodes, you need to add the data disk as a resource to the cluster.

- Add the data disk as a new resource in the cluster as shown in Figure 15. This disk acts as a datastore for DataFabric Manager Server data.

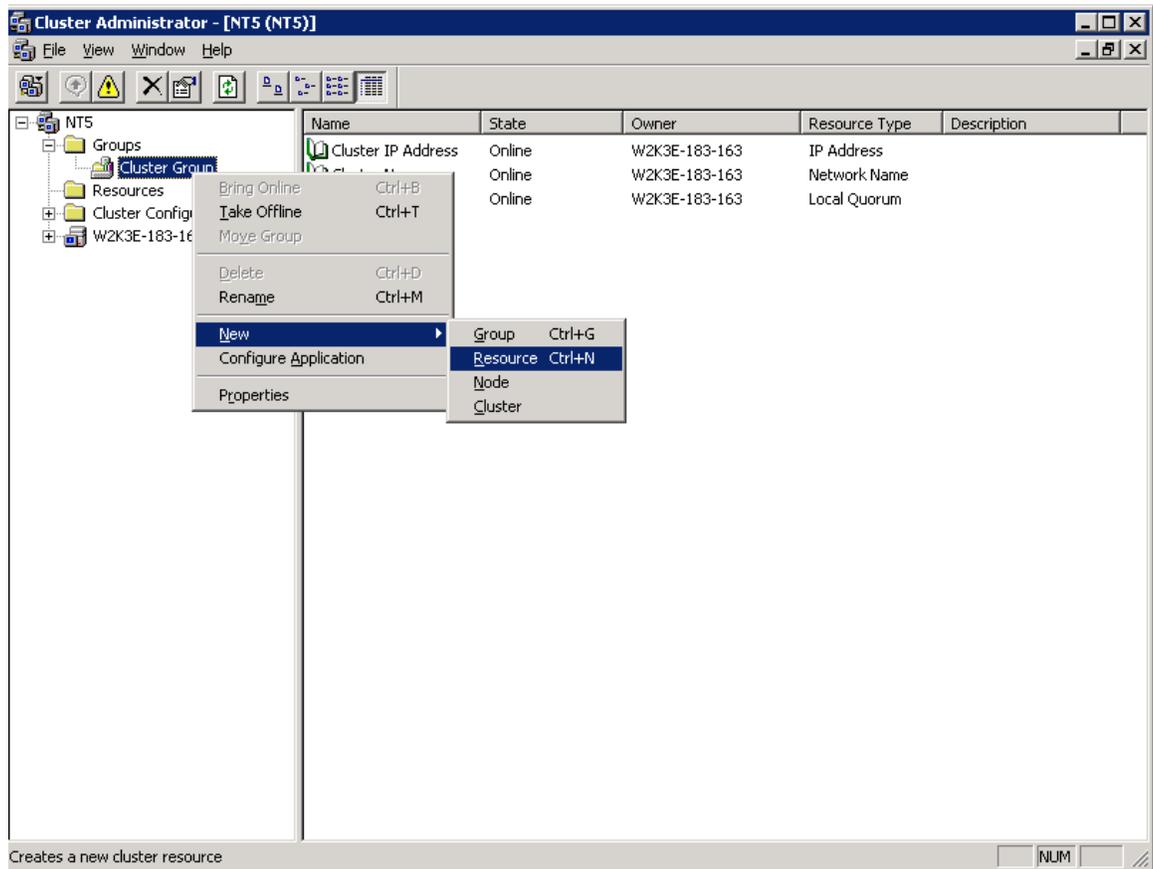


Figure 15 – Add new resource

- Select the resource type as “physical disk” and name it as shown in Figure 16. Press Next to complete the wizard.

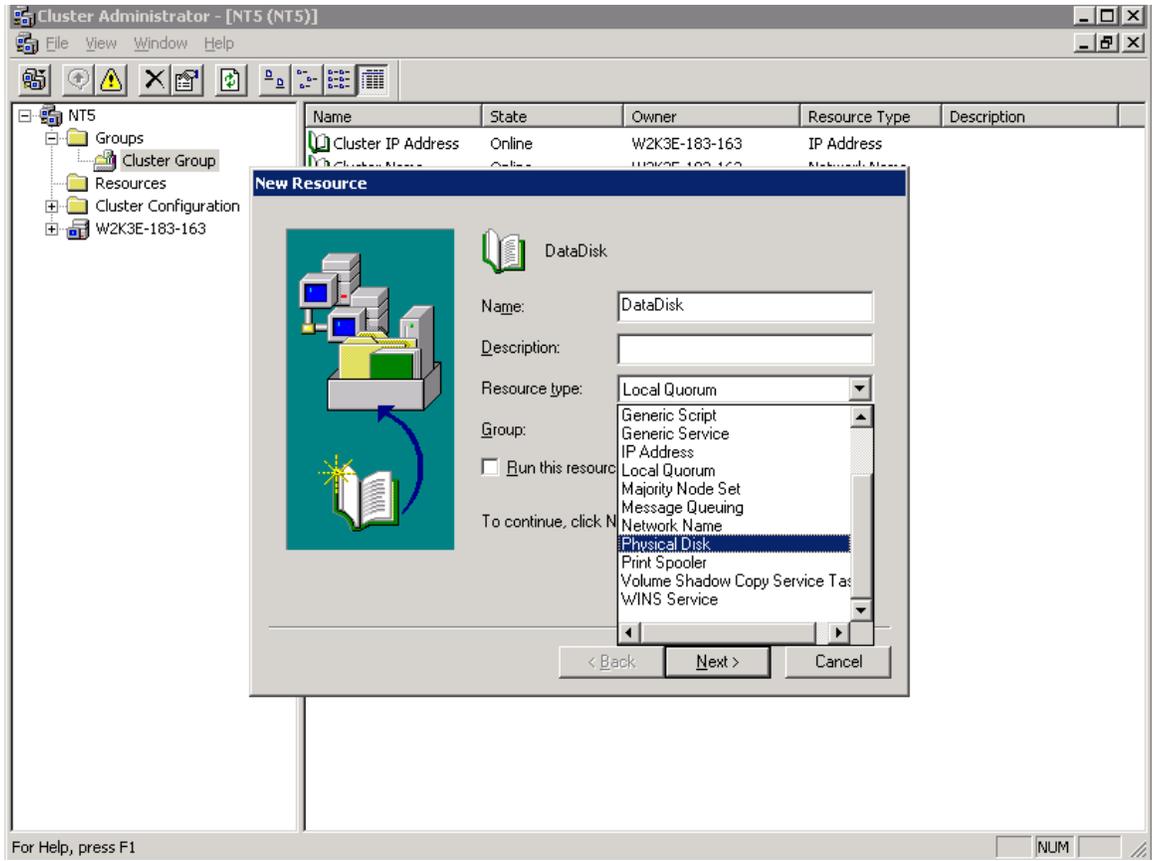


Figure 16 – Select physical disk

- The disk will be in offline state as shown in Figure 17. Right-click and bring the disk online.

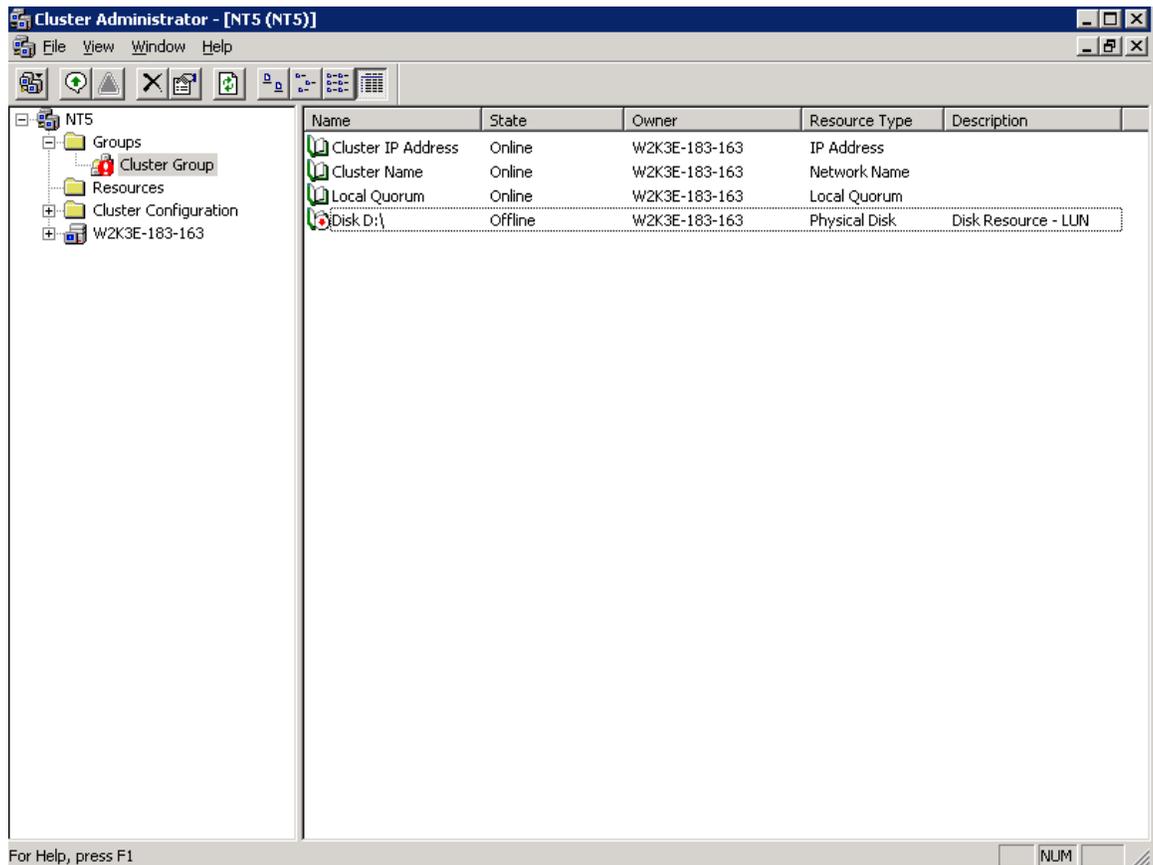


Figure 17 – Cluster administrator with offline data disk

- Manually move the cluster to the secondary node and check if cluster failover is working as shown in Figure18.

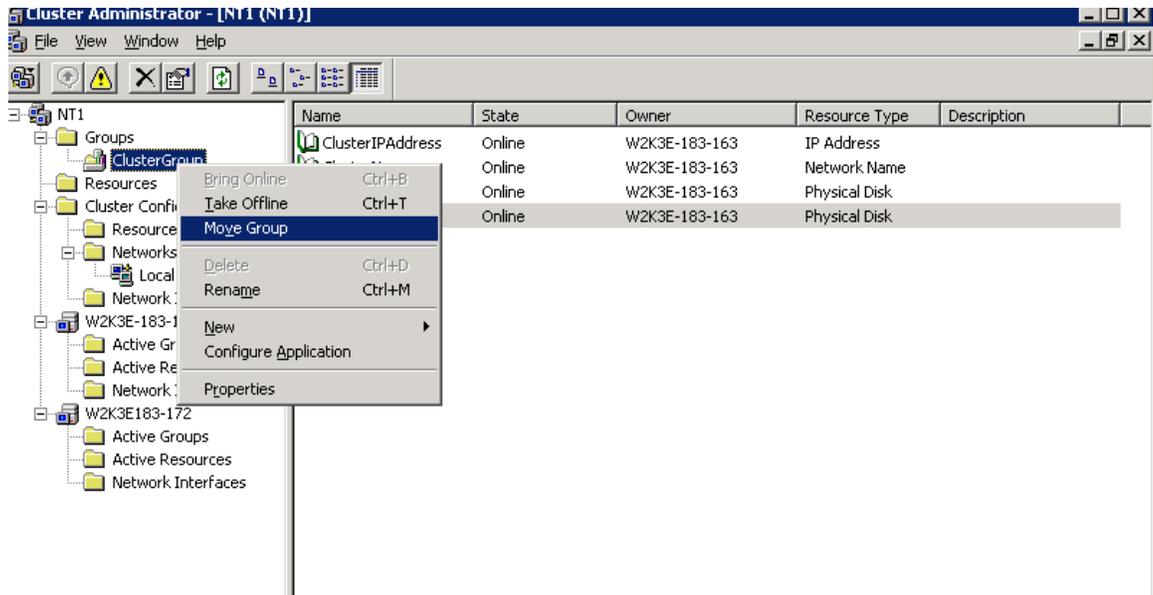


Figure 18 – Cluster manual failover

Note: Also check the below items to make sure that the cluster failover is working fine.

1. Check if the cluster IP address is migrated to the second cluster node. This can be checked by running the `ipconfig` command on both the nodes. The cluster IP address must be available only on the second node.

```

C:\Documents and Settings\smdev1>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection 2:

    Media State . . . . . : Media disconnected

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . :
    IP Address. . . . . : 10.72.186.61
    Subnet Mask . . . . . : 255.255.255.224
    IP Address. . . . . : 10.72.186.36
    Subnet Mask . . . . . : 255.255.255.224
    Default Gateway . . . . . : 10.72.186.33
  
```

2. Check the data disk is migrated to the second cluster node. This can be checked by accessing the data disk on the second node. The data disk must be accessible only on the second node.

2.3.3.2 CLUSTER CONFIGURATION ON WINDOWS 2008

This section explains how to create a cluster group with cluster name, virtual IP, quorum disk and data disk on a Windows 2008 server.

Note: Make sure that the cluster nodes are added to an Active Directory domain and the domain user should have administrator privileges on the cluster nodes.

- Install Failover Clustering from Server Manager on both the cluster nodes.

Note: Failover clustering is not installed by default on windows 2008.

- Log in as a domain user to the cluster node1 and launch Failover Cluster Management.
- Launch Create cluster Wizard window as shown in Figure 19.

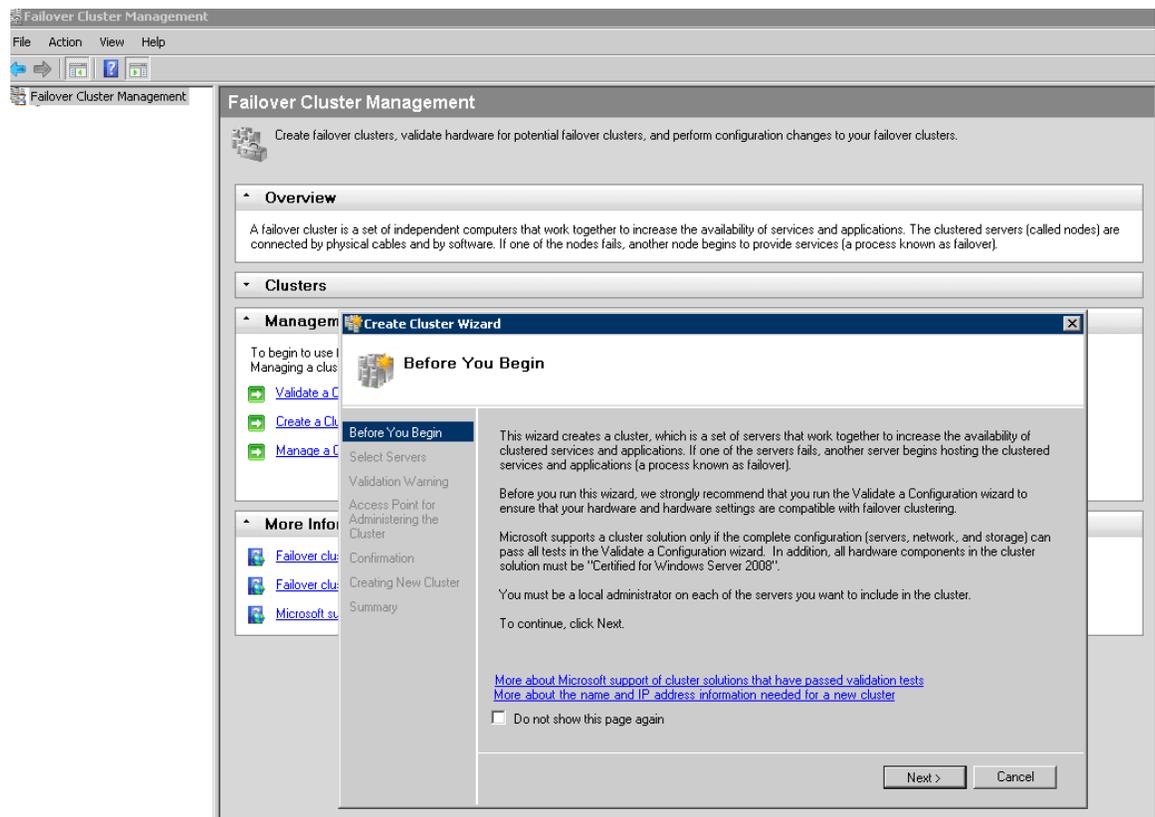


Figure 19 – Launch create cluster wizard

- Enter the name of the cluster node1 as shown in Figure 20.

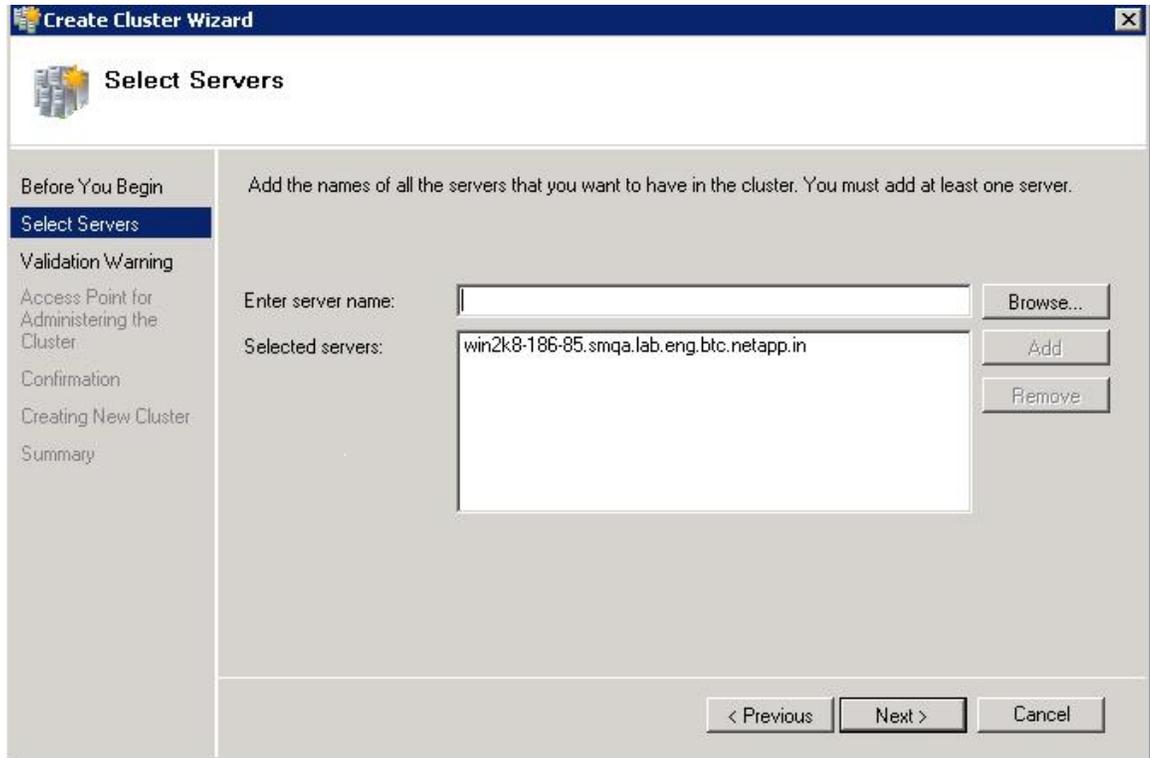


Figure 20 – Add node1

- Continue the wizard and make sure that all the tests are completed and successful.

Note: Make sure all the validation tests are run in the above create cluster wizard.

- Enter a unique cluster name and a virtual IPv4 address in the same sub network accessible by both of the clustered nodes as shown in Figure 21.

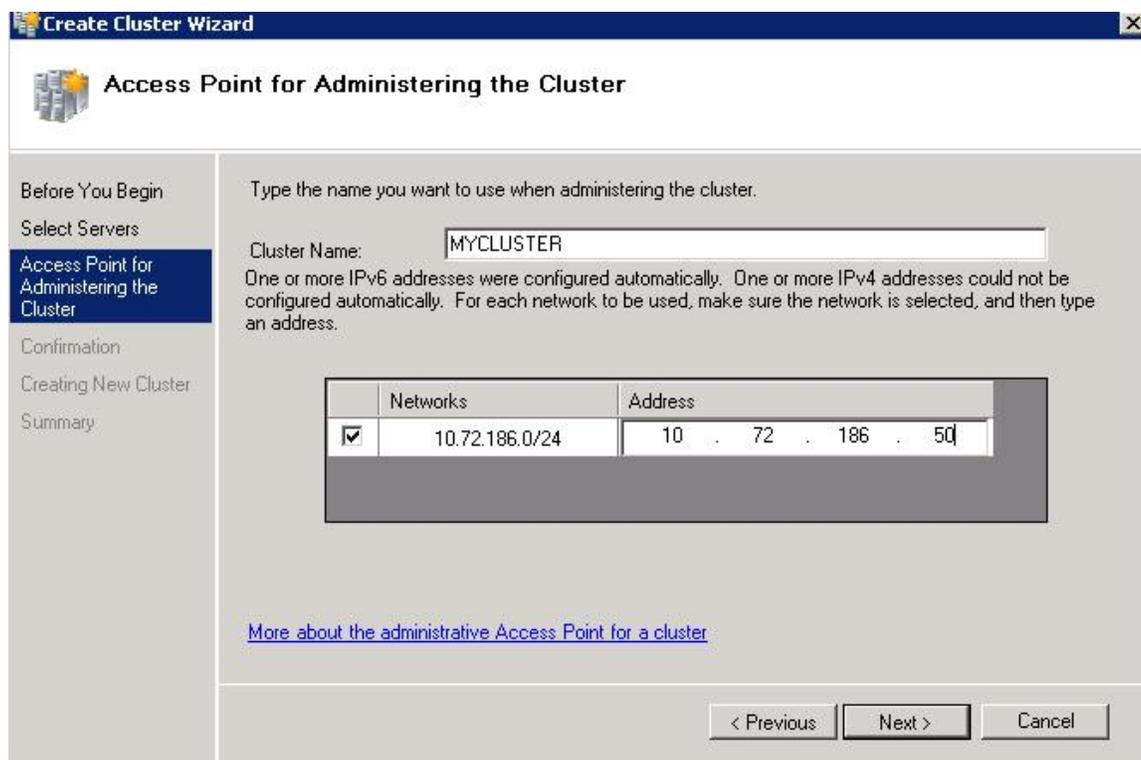


Figure 21 – Enter cluster name and IP address

- An ipv6 virtual IP address will be configured automatically, if the cluster nodes are enabled with IPv6 feature as shown in Figure 22.

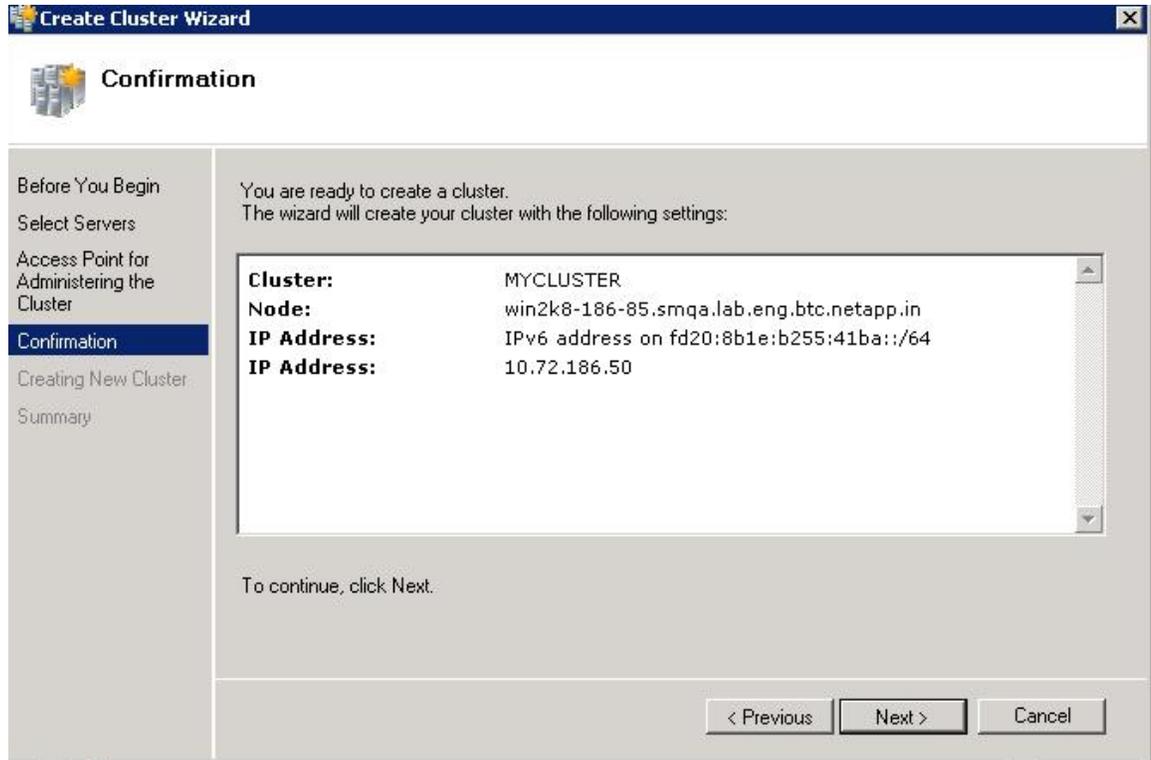


Figure 22 – Confirmation window

- Continue the wizard and till the end.

Note:

- The IPv4 and IPv6 virtual IP addresses to be used to invoke Operations Manager and the Management Console.
- Addition of quorum disk in failover clustering will be taken care separately.
- If there is some error in the cluster configuration steps, it will be indicated in red. Expand the failed operation to know about the reason for failure. Once the problem is fixed, reanalyze the configuration by clicking the Re-analyze button.

- . The created cluster should be listed with one host as mentioned in Figure 23.

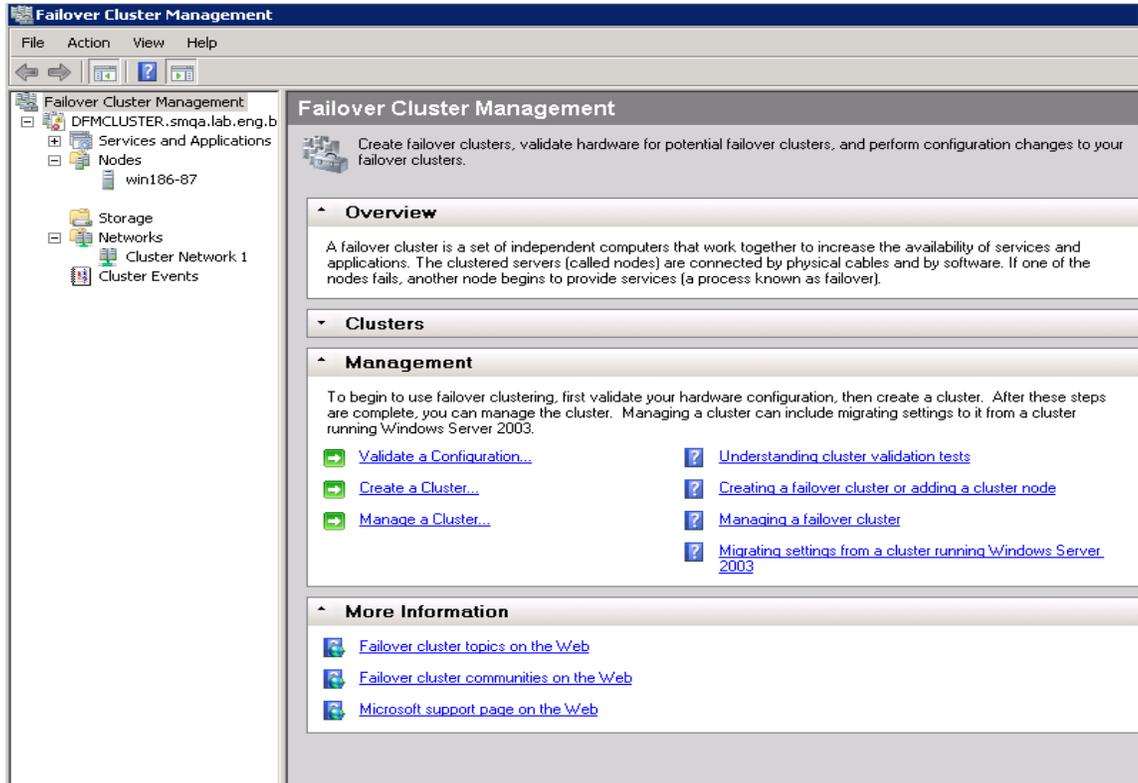


Figure 23 – Failover cluster management window

- Select the created cluster and press Add Node to launch Add node wizard as shown in Figure 24.

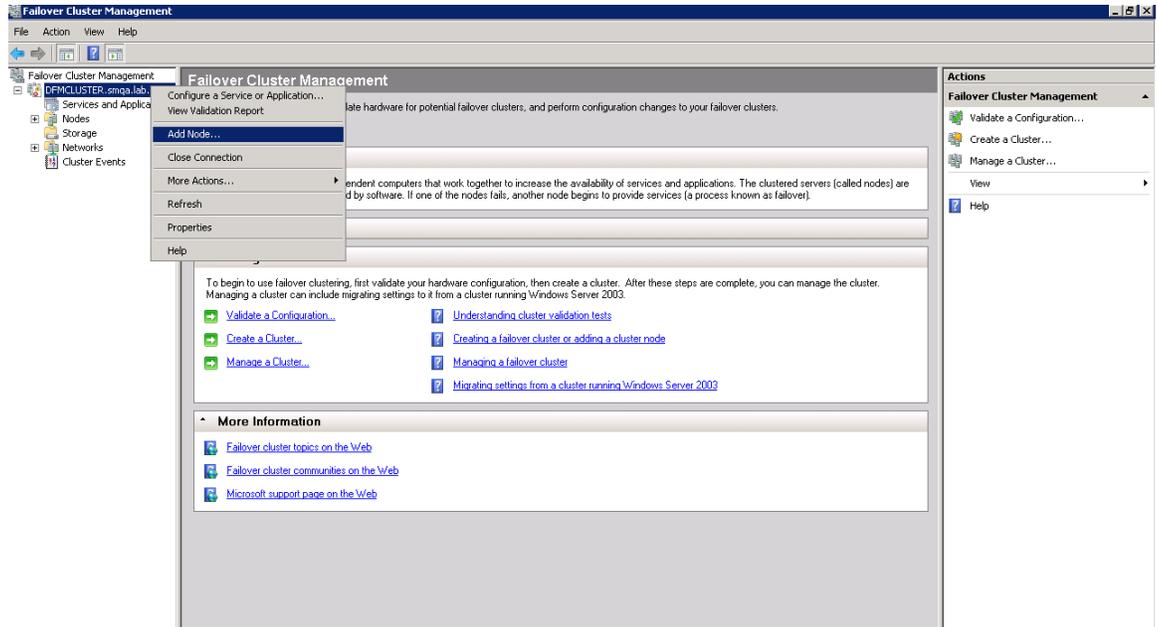


Figure 24 – Add second node into cluster

- Enter the name of the cluster node2 as shown in Figure 25 and continue the wizard till the end.

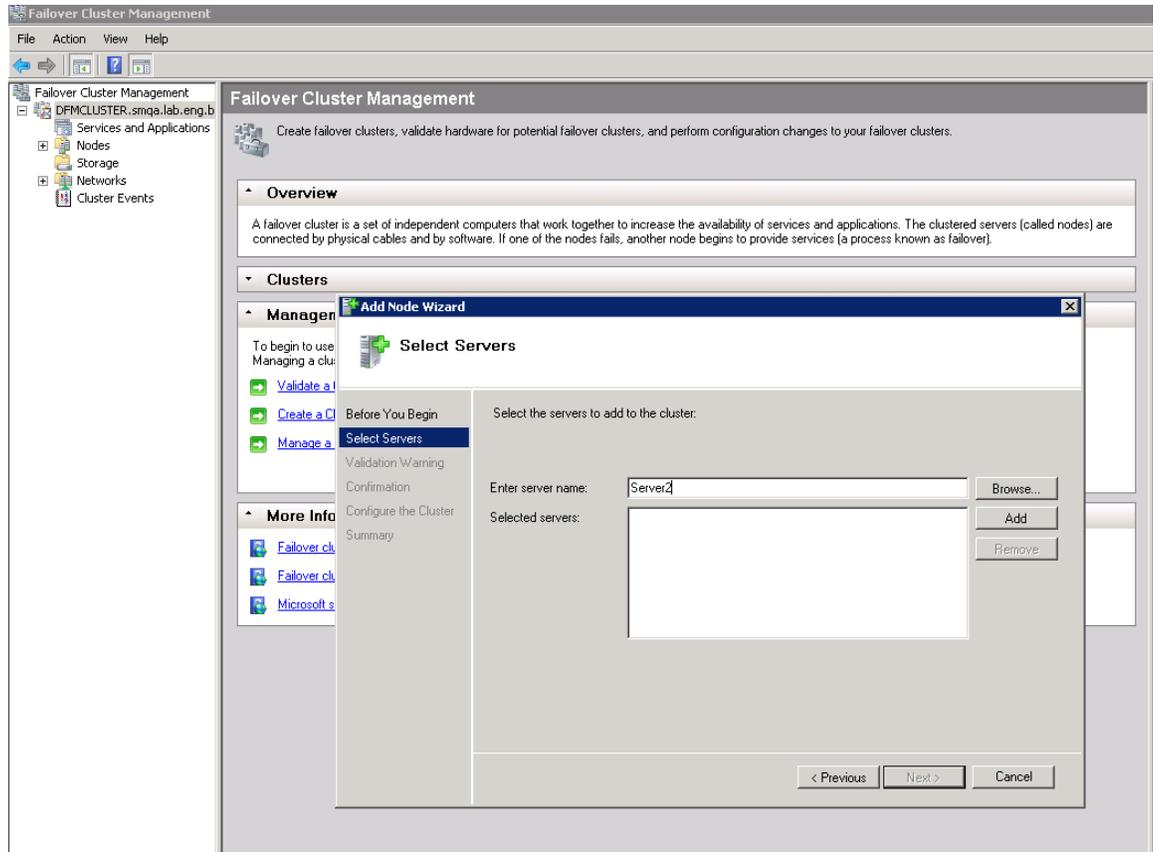


Figure 25 – Launch add node wizard

Note: Make sure all the specified tests have been run in the Add Node wizard.

- The failover cluster should show both the nodes as shown in Figure 26.

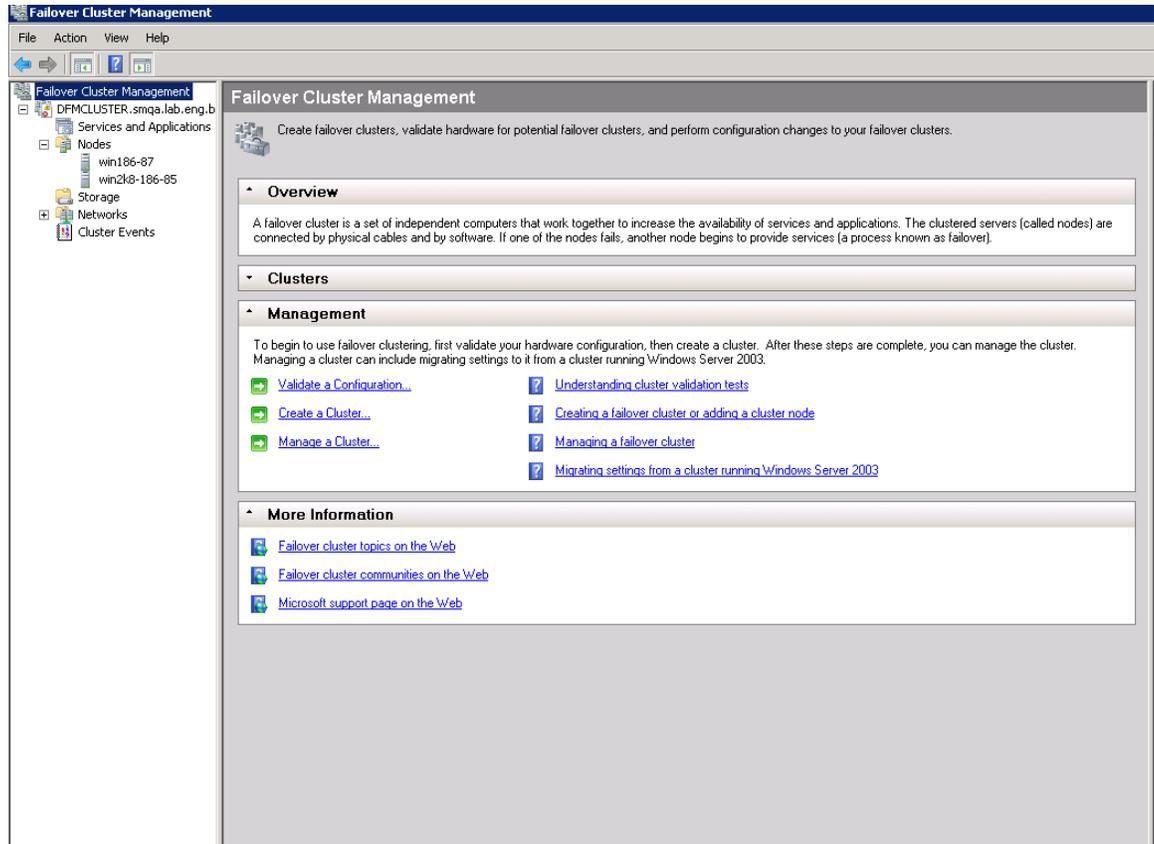


Figure 26 – Failover cluster with two nodes

- Select Cluster-> More actions->Configure Cluster Settings to open the quorum configuration wizard as shown in Figure 27.

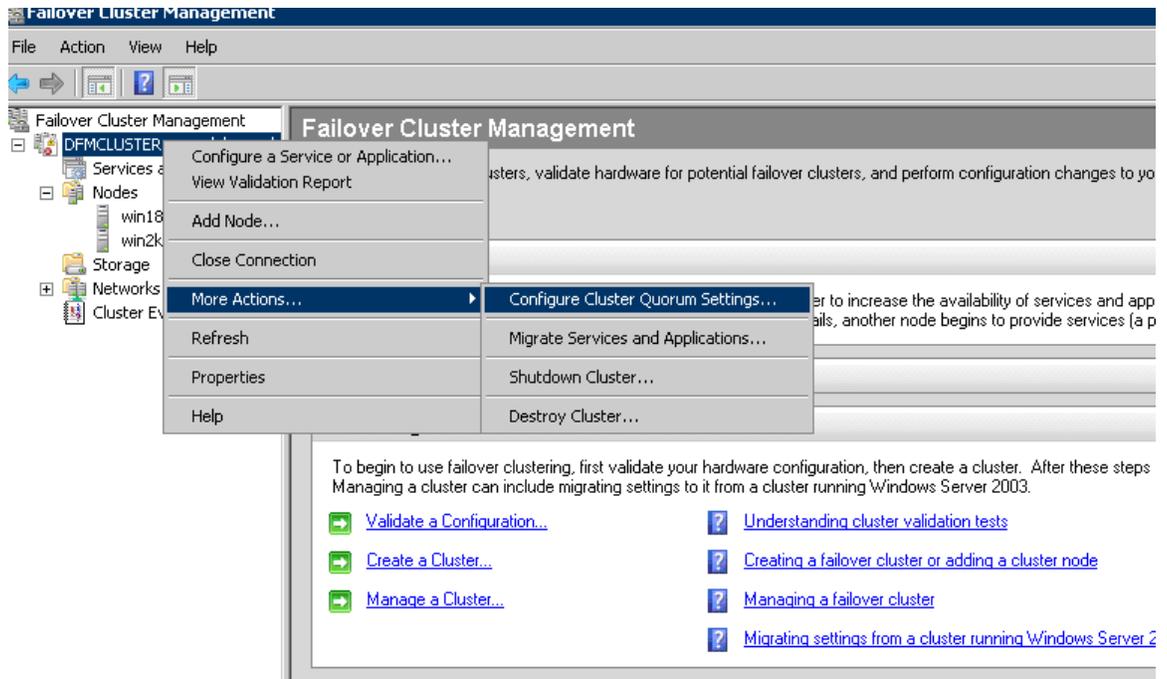


Figure 27 – Launch configure quorum settings wizard

- Select Node and Disk Majority in the quorum configuration wizard as shown in Figure 28.

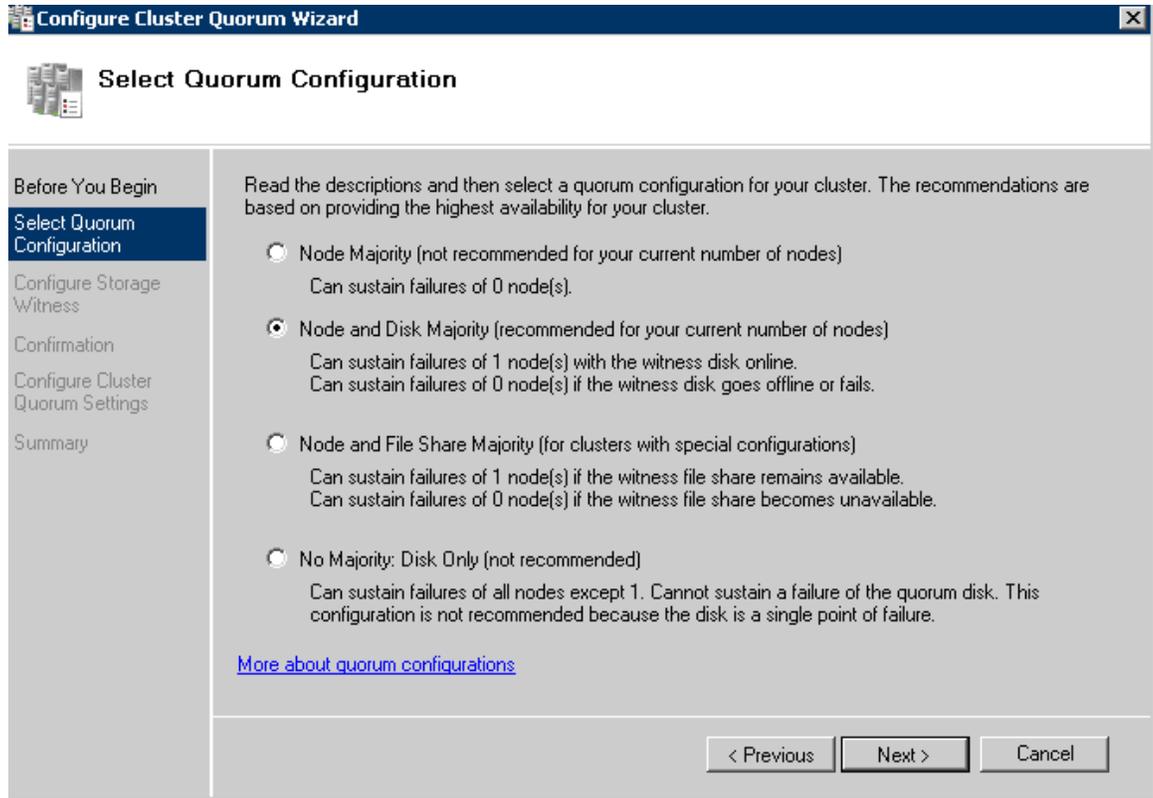


Figure 28 – Select quorum configuration

- Select and add the created quorum disk as shown in Figure 29 and continue the wizard till the last page.

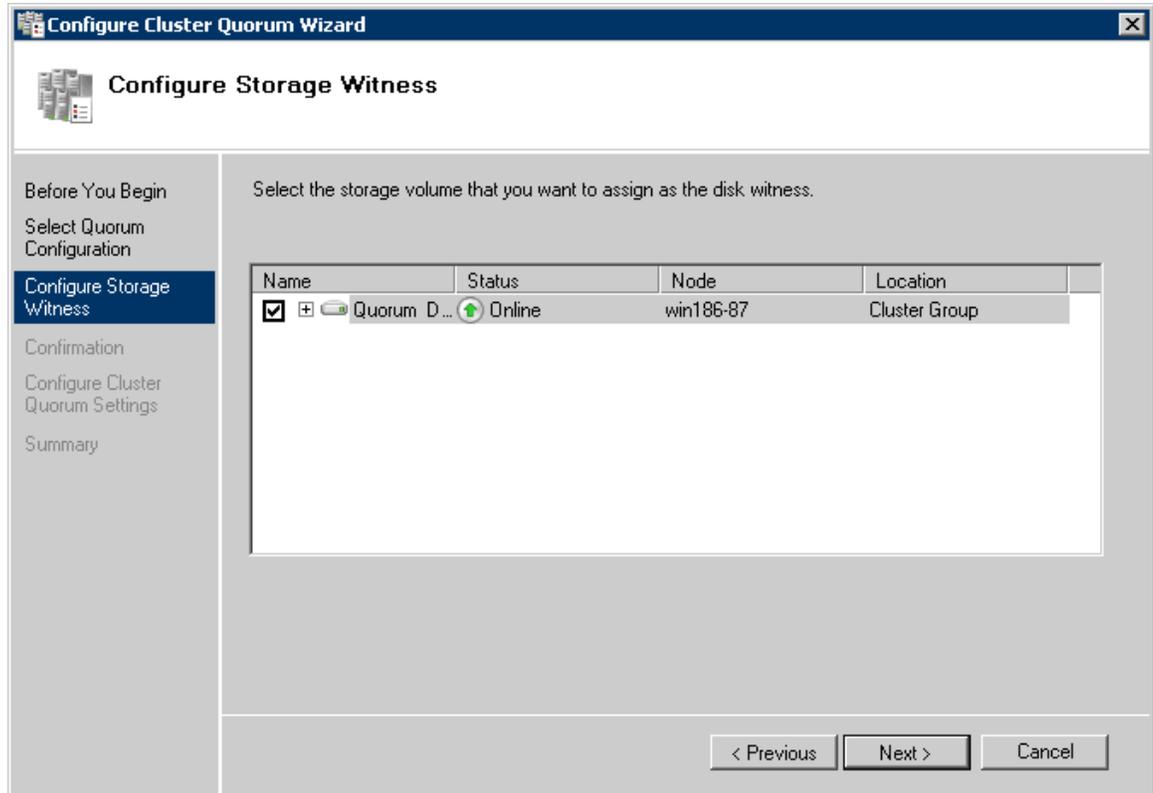


Figure 29 – Select quorum disk

- Select the Cluster -> Storage -> Add storage and add the data disk as shown in Figure 30.

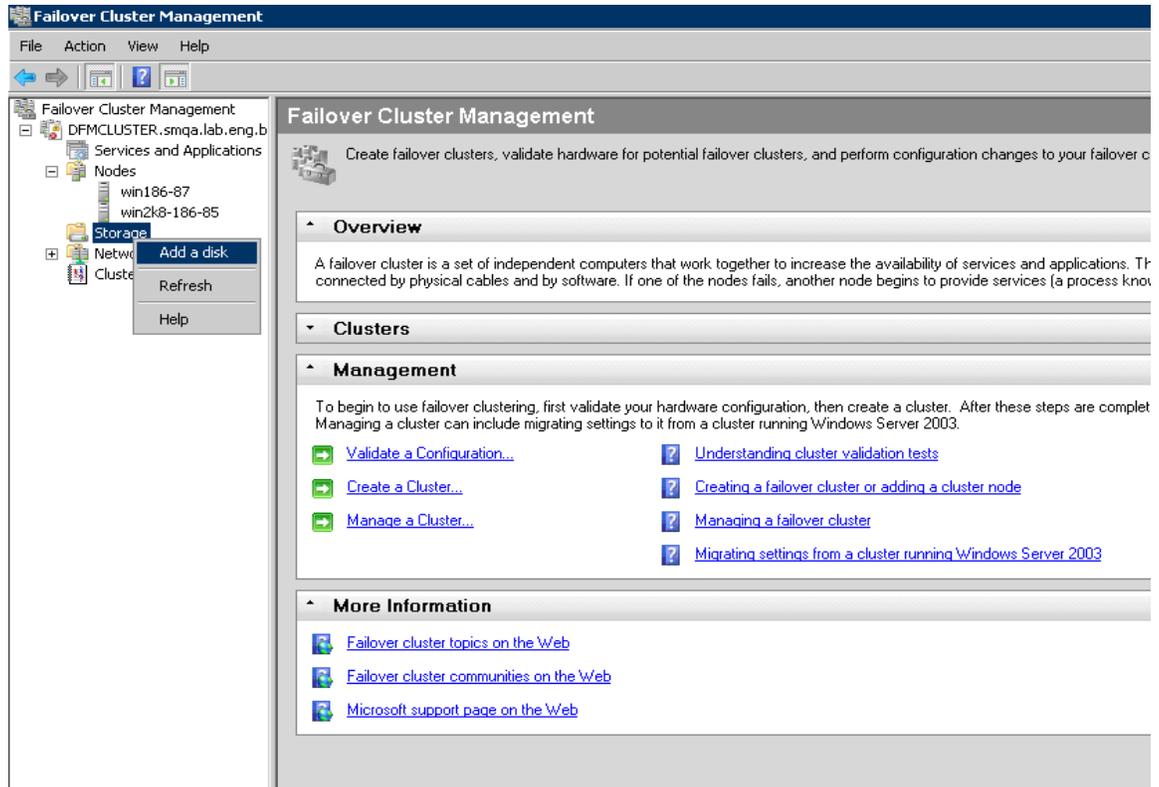


Figure 30 – Launch add data disk wizard

- Select the storage in the failover cluster management tree as shown in Figure 31. It should show the quorum disk and data disk.

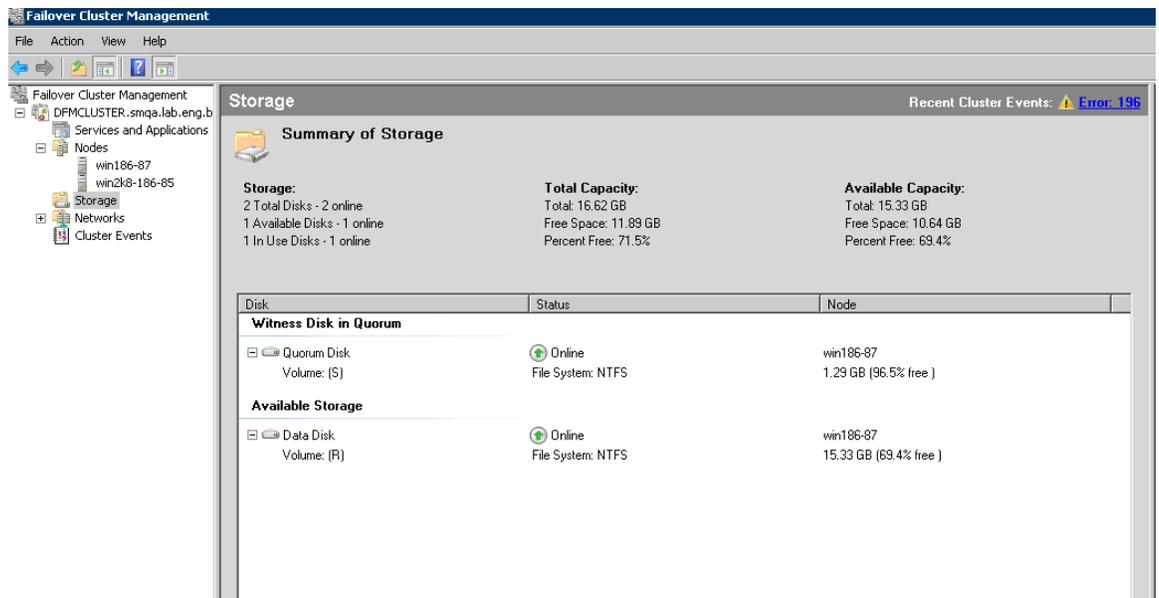


Figure 31 - View storage

Note: By default the name of the disks will be shown as Available Storage and Witness Disk in Quorum. We can rename the same into Data Disk and Quorum Disk for ease of reference. This can be done by selecting the disk -> Configuration and rename the Resource Name.

- Select the cluster node having access to the virtual IP and storage and click more actions->Stop to stop the cluster server as shown in Figure 32.

Note: This is to make sure that the virtual IP address and storage are migrated to another node of the cluster.

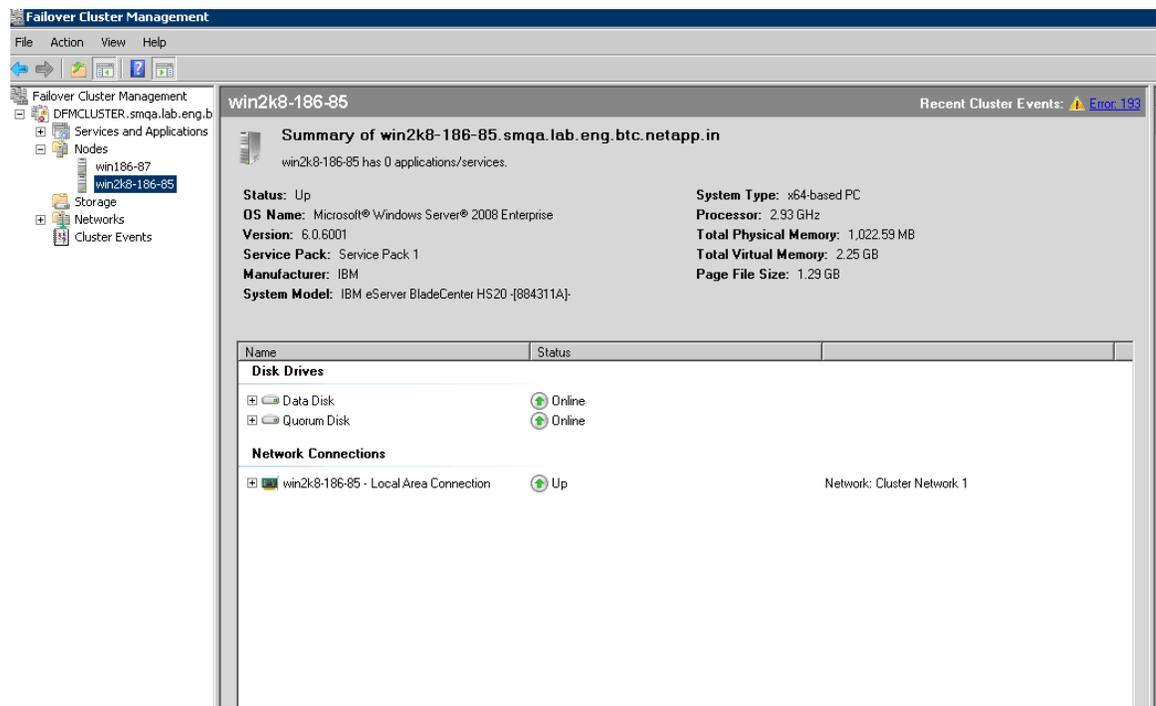


Figure 32 – Migrate cluster resources to second node

Note: Check if the cluster IP address and storage are migrated to the second cluster node. This can be checked by selecting the cluster node on the Failover Cluster Management tree. The right side shows the disk drives and network connections as shown Figure 32.

- Select the node where the cluster service is stopped and start the cluster service. Repeat the above step to move back the cluster resources to Node1.

2.3.4 INSTALLING DATAFABRIC MANAGER SERVER

This section explains how to install DataFabric Manager on the cluster server nodes.

- Log in as a domain user to both the nodes of the cluster.
- Install DataFabric Manager Server on both the nodes.
Note: DataFabric Manager Server should be installed on the same path in both the cluster nodes (for example, C:\Program Files\NetApp\DataFabric).
- Make the DataFabric Manager Server Service start up in manual mode by running the command `dfm service enable -m`.

Note:

- This is very important as the services should not be started automatically on node reboot. Service start and stop should be managed only by cluster service.
- `dfm service enable` command is not available for the DataFabric Manager Server releases less than 3.7. Follow the below steps to set the service start to manual mode.
 - From the Start menu, select Control Panel. In the Control Panel, select Administrative Tools > Services.
 - Double-click the DFM Sybase ASA icon to launch its properties window.
 - In the General tab, change the startup type option to Manual and click Apply.
 - Click the Log On tab.
 - Enter the name of the user account you want to use to access the DataFabric Manager database on the shared disk and click OK.
 - Repeat the above steps for DFM Apache, DFM Server, DFM Monitor, DFM Scheduler, DFM Event, and DFM Watchdog services.

2.3.5 SETTING UP DATAFABRIC MANAGER SERVER WITH MSCS USING CONFIGURATION SCRIPTS AND DFM SERVER COMMANDS

2.3.5.1 SETTING UP DATAFABRIC MANAGER SERVER ON WINDOWS 2003

This section explains how to add DataFabric Manager Server services as the cluster resources

- Log in to the node owning cluster resources. The owner of the cluster resources can be obtained by launching the cluster administrator as shown in Figure 33.

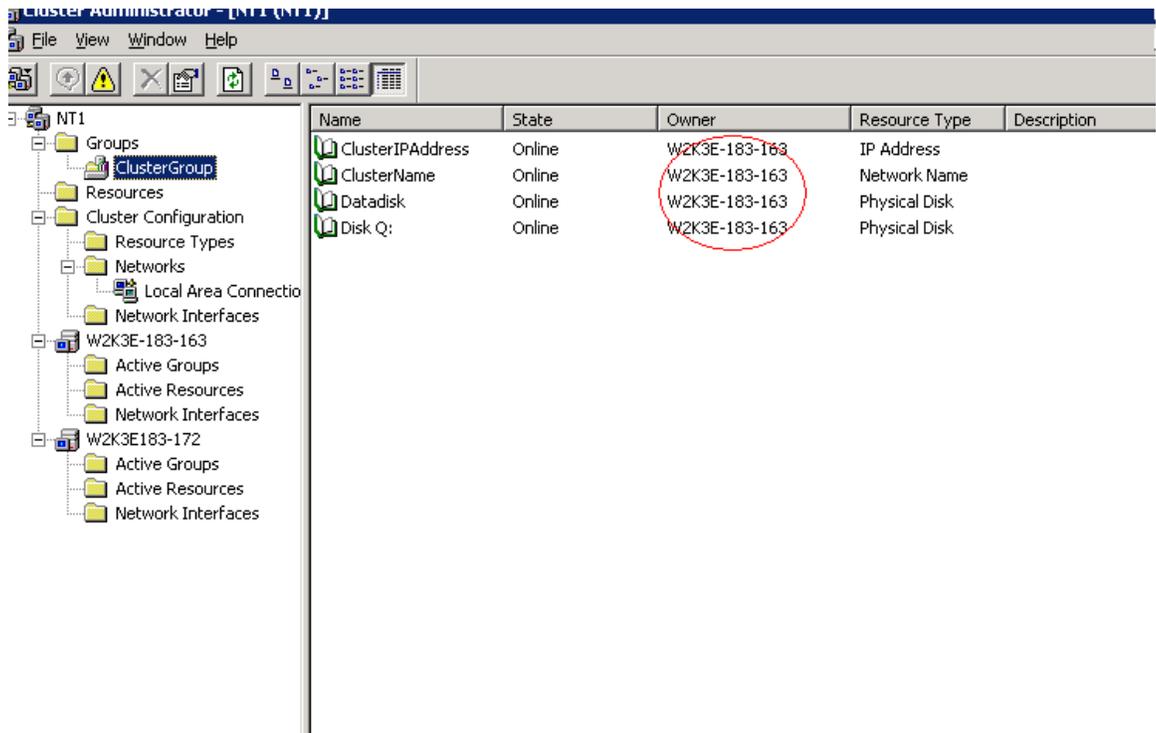


Figure 33 – Find the cluster owner

- Run `dfm datastore setup <Drive Name>` command to move the DataFabric Manager Server data onto the shared data disk as shown below.

```

C:\>dfm datastore setup E:\
Required space for data:13.0 MB, Available space:22.9 GB
Stopping all services and jobs...
Copying database files to E:\data.
2 File(s) copied
Copying perf data files to E:\perfddata.
Copying script output files to E:\script-plugins.
Changing data configuration settings ...
Changed dbDir to E:\data.
Updated dbLogDir to E:\data.
Starting sql service...
Changed perfArchiveDir to E:\perfddata.
Changed scriptDir to E:\script-plugins.
Changed databaseBackupDir to E:\data

```

Note: The above command `dfm datastore setup` is not available for the DataFabric Manager Server release less than 3.6. If you are using lower versions of DFM Server, follow the steps mentioned below to move the `dfm` data onto the shared LUN.

- Run the script `dfmcluster_firstnode.pl` located at
C:\Program Files\NetApp\DataFabric\DFM\examples\.
- Before running the script, make sure that :
 - You are logged in as domain user with administrator privileges on the local system.
 - The domain user should exist in DataFabric Manager Server as an administrator with full control privileges.
 - First node is the current owner of the cluster resources.
 - Make sure Perl is installed on the server.
 - In Windows 2008(64bit) server, 64bit Perl to be installed for running DataFabric Manager Server configuration scripts.
- The script copies the DataFabric Manager Server data onto the shared disk and configures DataFabric Manager Server Services to log in using specified domain user name.
 - Pass the following parameter to the script:
 - `-u [domain-user-name]`: domain user who has administrator privileges on the local system
 - `-p [password]`: password for the domain user
 - `-s [dfm-install-dir]`: DataFabric Manager Server installation path. Source path for the data to be moved (same as in `dfm about output`)
 - `-d [data-path]`: Path to a destination folder in shared data disk where to move DataFabric Manager Server database and other files

```
C:\Program Files\NetApp\DataFabric\DFM\examples>C:\Perl\bin\perl dfmcluster_firstnode.pl
-u btclab\smdev1 -p xxxxx -s "C:/Program Files/NetApp/DataFabric/DFM" -d "E:\DFM"
```

- Stop the DataFabric Manager Server services by `dfm service stop` command.
Note: This make sure that the `dfm` services do not try to access the data disk, which is moved to the secondary cluster node.
- Using the cluster administrator now manually `move Cluster group` into the second node (that is, make the second node the owner). The procedure to move is shown in Figure 18. Once the operation is successful, verify that the secondary node is the owner of the cluster resources.
- Run `dfm datastore setup -n <Drive Name>` command on the secondary node. Make sure that the same drive letter is used on both the nodes.

```
C:\Documents and Settings\smdev1>dfm datastore setup -n E:\
Stopping all services and jobs...
Changing data configuration settings...
Changed dbDir to E:\data.
```

```
Updated dbLogDir to E:\data.
Starting sql service...
Changed perfArchiveDir to E:\perfddata.
Changed scriptDir to E:\script-plugins.
Changed databaseBackupDir to E:\data.
NOTE: You must move all archive database backups from 'C:\Program
Files\NetApp\DataFabric\DFM\data\' to new directory
No database upgrade is necessary.
```

Note: The option “-n” in `dfm datastore setup` command make sure that data is not copied again to the shared data disk but the dfm is pointed to the shared disk as it location for db and other stuff. The above command `dfm datastore setup` is not available for the DataFabric Manager Server version less than 3.6. If you are using lower versions of DataFabric Manager Server, follow the steps mentioned below to move the DataFabric Manager Server data onto the shared LUN.

1. Run the script `dfmcluster_secondnode.pl` located at
`C:\Program Files\NetApp\DataFabric\DFM\examples\.`
2. The script configures DataFabric Manager Server to access the data from the shared disk and configures DataFabric Manager Server services to log in using specified domain username.

```
C:\Program Files\NetApp\DataFabric\DFM\examples>C:\Perl\bin\perl dfmcluster_secondnode.pl
-u btclab\smdev1 -p smdev1 -d E:\
```

- Stop all the DataFabric Manager Server services by running `dfm service stop` command
- To add the DataFabric Manager Server Service as resources to the MSCS go to the install directory `C:\Program Files\NetApp\DataFabric\DFM\examples`
- Run `dfmcluster_add_resources.pl -g <cluster Group Name> -I <cluster IP resource Name> -n <Cluster Resource Name> -K <data disk Name>` script to configure the dfm services in MSCS. Please refer to Figure 34 to get values of the options to be specified.

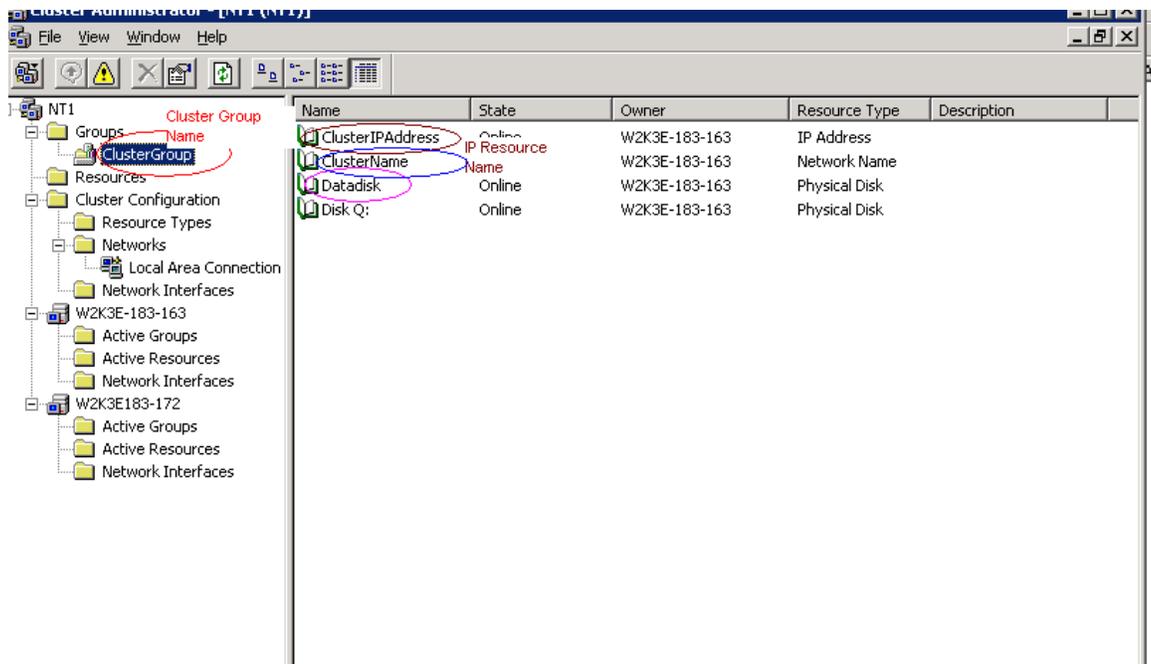


Figure 34 – Cluster entities

```

C:\Program Files\NetApp\DataFabric\DFM\examples>perl dfmcluster_add_resources.pl -g ClusterGroup -i
ClusterIPAddress -n ClusterName -k Data_disk

=== DataFabric Manager Cluster Configuration Script Stopping services ...
Configuring cluster resources ...
=== DataFabric Manager Node Configured Successfully for Clustering ==

```

- Check for successful addition of resources in 'Cluster Administrator'. Please refer Figure 35.

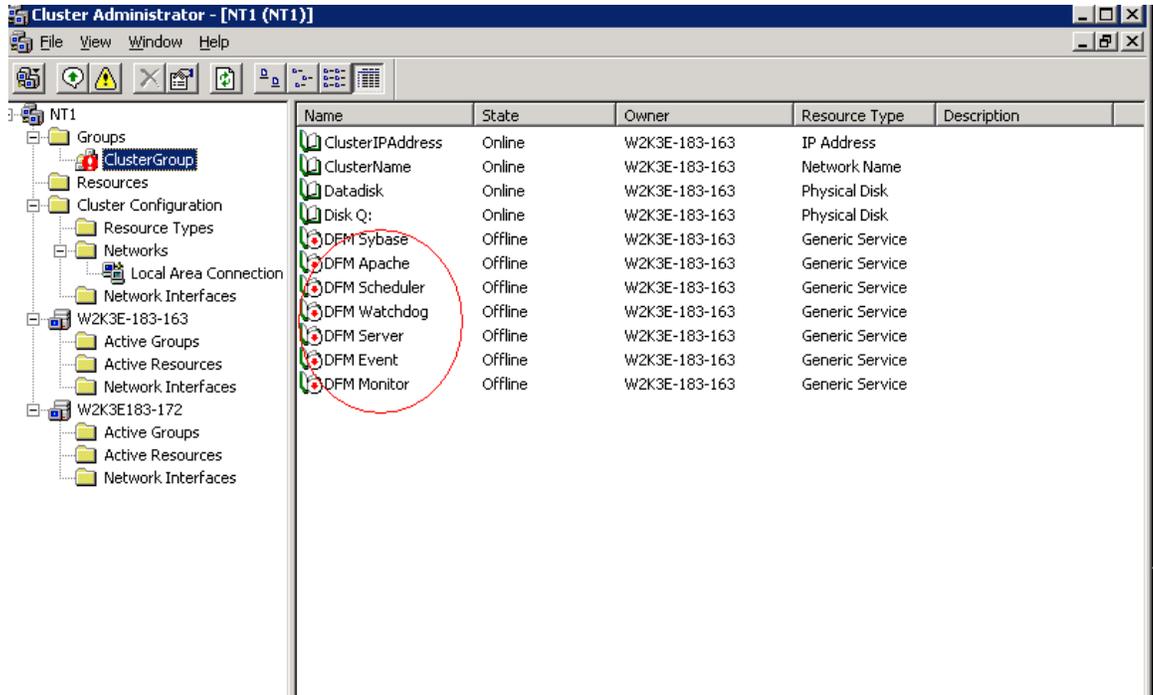


Figure 35 – Operations Manager services in Cluster Administrator

- Right-click on the cluster group and select **Bring Online** to online all the DataFabric Manager Server services. Please refer to Figure 36.

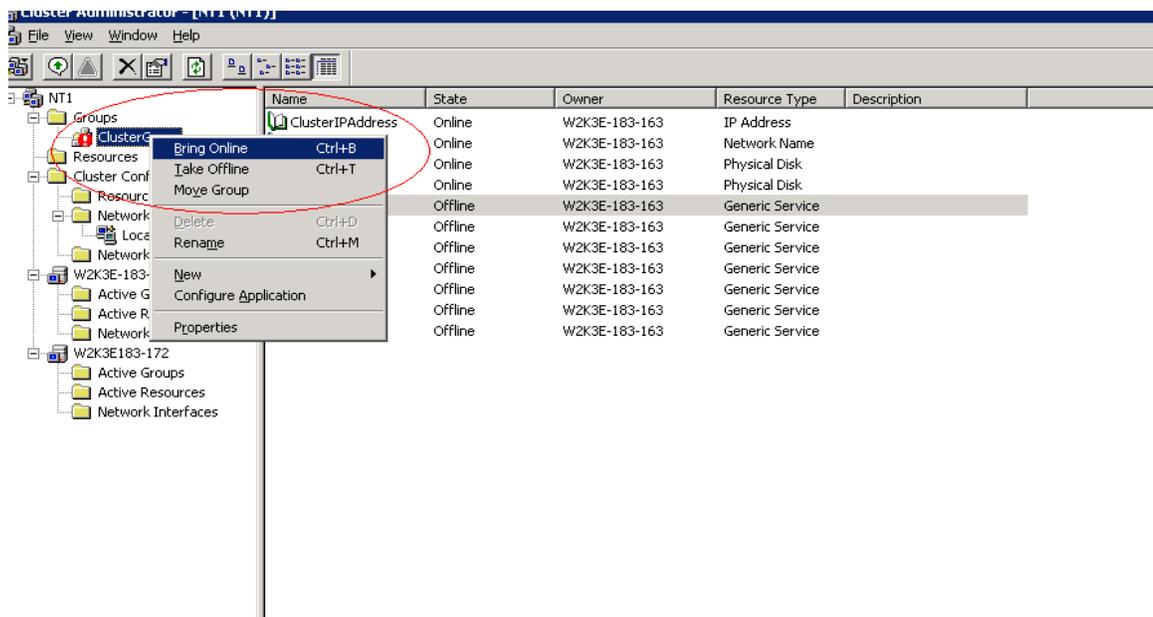


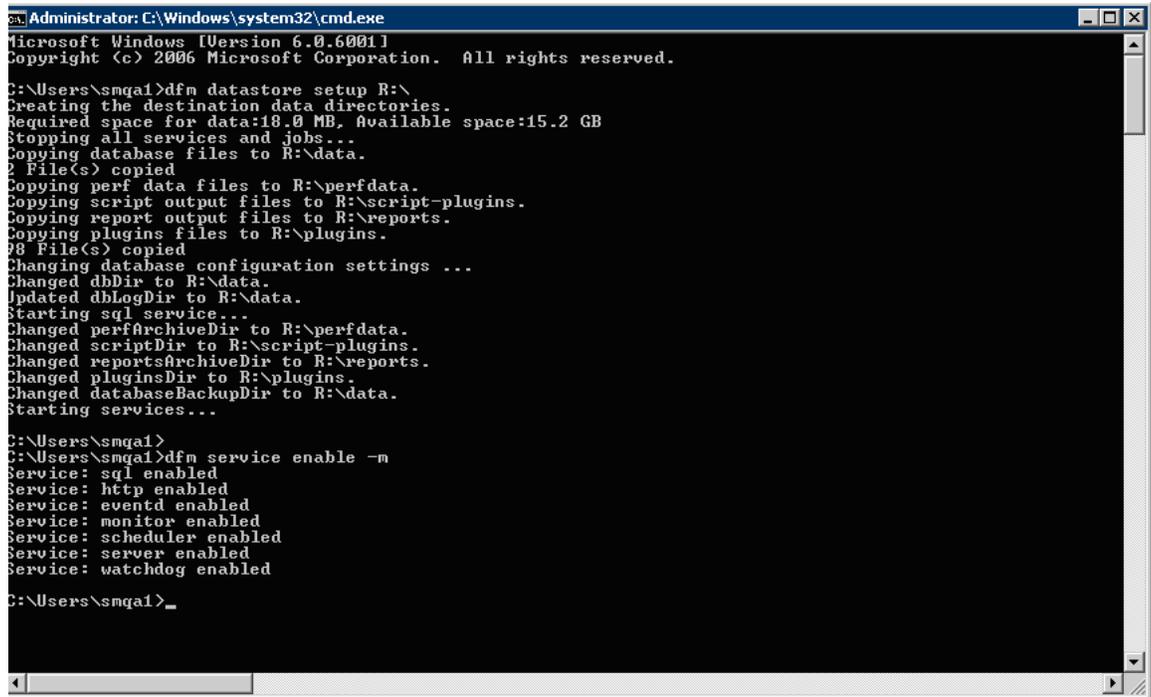
Figure 36 – Online Operations Manager services

- Once the services are online, move the group to the second node and check that the cluster is now up on the second node.

2.3.5.2 SETTING UP DATAFABRIC MANAGER SERVER ON WINDOWS 2008

This section explains how to add DataFabric Manager Server Services as the cluster resources in Windows 2008 Server

- Log in to the node owning cluster resources. The owner of the cluster resources can be obtained by selecting the cluster node in failover cluster management as shown in Figure 32.
- Run `dfm datastore setup <Drive Name>` command to move the DataFabric Manager Server data onto the shared data disk as shown Figure 37.



```
Administrator: C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.0.6001]
Copyright (c) 2006 Microsoft Corporation. All rights reserved.

C:\Users\sngal>dfm datastore setup R:\
Creating the destination data directories.
Required space for data:18.0 MB. Available space:15.2 GB
Stopping all services and jobs...
Copying database files to R:\data.
2 File(s) copied
Copying perf data files to R:\perfdata.
Copying script output files to R:\script-plugins.
Copying report output files to R:\reports.
Copying plugins files to R:\plugins.
98 File(s) copied
Changing database configuration settings ...
Changed dbDir to R:\data.
Updated dbLogDir to R:\data.
Starting sql service...
Changed perfArchiveDir to R:\perfdata.
Changed scriptDir to R:\script-plugins.
Changed reportsArchiveDir to R:\reports.
Changed pluginsDir to R:\plugins.
Changed databaseBackupDir to R:\data.
Starting services...

C:\Users\sngal>
C:\Users\sngal>dfm service enable -m
Service: sql enabled
Service: http enabled
Service: eventd enabled
Service: monitor enabled
Service: scheduler enabled
Service: server enabled
Service: watchdog enabled

C:\Users\sngal>_
```

Figure 37 – Moving dfm data into data disk on node1

- Move the cluster resources to cluster node2 by stopping the cluster service on node1 and run `dfm datastore setup -n` command to make DataFabric Manager Server Database refer to the shared disk as shown in Figure 38.

```

Microsoft Windows [Version 6.0.6001]
Copyright (c) 2006 Microsoft Corporation. All rights reserved.

C:\Users\smqa1>dfm datastore setup -n R:\
Stopping all services and jobs...
Changing database configuration settings ...
Changed dbDir to R:\data.
Updated dbLogDir to R:\data.
Starting sql service...
Changed perfArchiveDir to R:\perfdata.
Changed scriptDir to R:\script-plugins.
Changed reportsArchiveDir to R:\reports.
Changed pluginsDir to R:\plugins.
Changed databaseBackupDir to R:\data.
NOTE: You must move all archive database backups from 'C:\Program Files (x86)\NetApp\DataFabric Manager\DFM\data\' to new directory.
No database upgrade is necessary.
Redefining SQL Views ...
Defining SQL debug procedures ...
Starting services...
Error: SnapDrive version: . Supported version is 4.1 or higher.

C:\Users\smqa1>

```

Figure 38 – Refer DFM DB in the data disk in node 2

- Create a new application group by selecting Services and Applications -> More Actions -> Create Empty Service or Application from failover cluster management as shown in Figure 39.

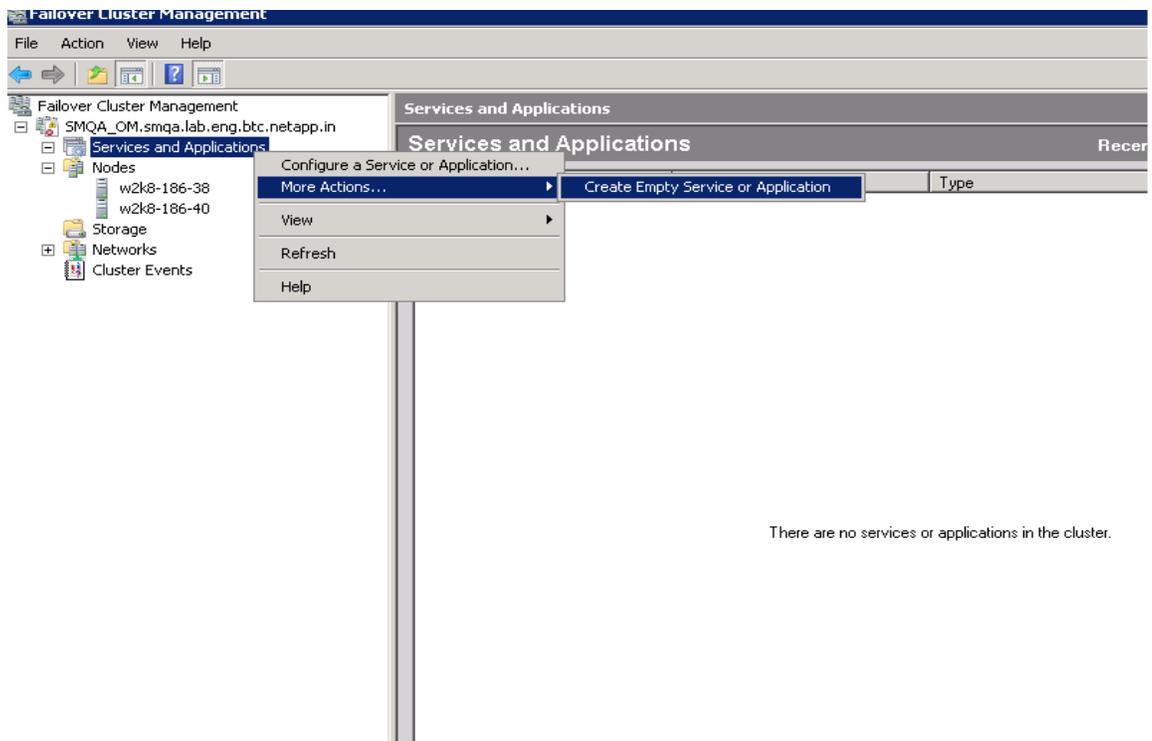


Figure 39 – Create Application group for DFM services

- Rename the application group to DFM group for naming convention.
- Run the `cluster group` command from the command prompt and check the above created application group is listed as shown in Figure 40.

```

Administrator: C:\Windows\system32\cmd.exe
C:\Users\smqa1>cluster group
Listing status for all available resource groups:
Group                Node                Status
-----
Available Storage    w2k8-186-38         Offline
Cluster Group        w2k8-186-38         Online
DFM Group             w2k8-186-40         Online
C:\Users\smqa1>

```

Figure 40 – Run cluster group command

- Offline all the cluster resources as shown in Figure 41. This step is to be performed for moving all the cluster resources into DFM group cluster group.

```

Administrator: C:\Windows\system32\cmd.exe
C:\Users\smqa1>cluster res "Cluster IP Address" /off
Taking resource 'Cluster IP Address' offline...
Resource            Group                Node                Status
-----
Cluster IP Address  Cluster Group        w2k8-186-38         Offline

C:\Users\smqa1>cluster res "Cluster Name" /off
Taking resource 'Cluster Name' offline...
Resource            Group                Node                Status
-----
Cluster Name        Cluster Group        w2k8-186-38         Offline

C:\Users\smqa1>cluster res "Data Disk" /off
Taking resource 'Data Disk' offline...
Resource            Group                Node                Status
-----
Data Disk           Cluster Group        w2k8-186-38         Offline

C:\Users\smqa1>cluster res "Quorum Disk" /off
Taking resource 'Quorum Disk' offline...
Resource            Group                Node                Status
-----
Quorum Disk         Cluster Group        w2k8-186-38         Offline

C:\Users\smqa1>

```

Figure 41 – Offline cluster resources

- Move all the cluster resources into DFM Group application group as shown in Figure 42.

```

ca. Administrator: C:\Windows\system32\cmd.exe

C:\Users\smqa1>cluster res "Cluster IP Address" /move:"DFM Group"
Moving resource 'Cluster IP Address' to group 'DFM Group'...
Resource          Group              Node              Status
-----
Cluster IP Address  DFM Group         w2k8-186-40      Offline

C:\Users\smqa1>cluster res "Data Disk" /move:"DFM Group"
Moving resource 'Data Disk' to group 'DFM Group'...
Resource          Group              Node              Status
-----
Data Disk         DFM Group         w2k8-186-40      Offline

C:\Users\smqa1>cluster res "Quorum Disk" /move:"DFM Group"
Moving resource 'Quorum Disk' to group 'DFM Group'...
Resource          Group              Node              Status
-----
Quorum Disk       DFM Group         w2k8-186-40      Offline

C:\Users\smqa1>cluster res
Listing status for all available resources:
Resource          Group              Node              Status
-----
Cluster IP Address  DFM Group         w2k8-186-40      Offline
Cluster IP Address 10.72.186.39 DFM Group         w2k8-186-40      Offline
Cluster Name        DFM Group         w2k8-186-40      Offline
Data Disk           DFM Group         w2k8-186-40      Offline
Quorum Disk         DFM Group         w2k8-186-40      Offline

C:\Users\smqa1>

```

Figure 42 – Move cluster resources into DFM group

- Online all the offline cluster resources as shown in Figure 43.

```

Administrator: C:\Windows\system32\cmd.exe
C:\Program Files\NetApp\DataFabric Manager\DFM\examples>cluster res "Cluster IP Address 10.72.186.39" /on
Bringing resource 'Cluster IP Address 10.72.186.39' online...
Resource          Group              Node              Status
-----
Cluster IP Address 10.72.186.39 DFM Group         w2k8-186-40      Online

C:\Program Files\NetApp\DataFabric Manager\DFM\examples>cluster res "Data Disk" /on
Bringing resource 'Data Disk' online...
Resource          Group              Node              Status
-----
Data Disk         DFM Group         w2k8-186-40      Online

C:\Program Files\NetApp\DataFabric Manager\DFM\examples>cluster res "Quorum Disk" /on
Bringing resource 'Quorum Disk' online...
Resource          Group              Node              Status
-----
Quorum Disk      DFM Group         w2k8-186-40      Online

C:\Program Files\NetApp\DataFabric Manager\DFM\examples>cluster res "Cluster Name" /on
Bringing resource 'Cluster Name' online...
Resource          Group              Node              Status
-----
Cluster Name     DFM Group         w2k8-186-40      Online

C:\Program Files\NetApp\DataFabric Manager\DFM\examples>

```

Figure 43 – Online cluster resources

- Run `dfmcluster_add_resources.pl -g <cluster Group Name> -I <cluster IP resource Name> -n <Cluster Resource Name> -K <data disk Name>` script to configure the dfm services in failover cluster. Please refer to Figure 44 to get values of the options to be specified.

```

Select Administrator: C:\Windows\system32\cmd.exe
C:\>cluster res
Listing status for all available resources:
Resource          Group              Node              Status
-----
Cluster IP Address DFM Group         w2k8-186-38      Online
Cluster IP Address 10.72.186.39 DFM Group         w2k8-186-38      Online
Cluster Name      DFM Group         w2k8-186-38      Failed
Data Disk         DFM Group         w2k8-186-38      Online
DFM Apache        DFM Group         w2k8-186-38      Offline
DFM Event         DFM Group         w2k8-186-38      Offline
DFM Monitor       DFM Group         w2k8-186-38      Offline
DFM Scheduler     DFM Group         w2k8-186-38      Offline
DFM Server        DFM Group         w2k8-186-38      Offline
DFM Sybase        DFM Group         w2k8-186-38      Offline
DFM Watchdog      DFM Group         w2k8-186-38      Offline
Quorum Disk       DFM Group         w2k8-186-38      Online

C:\>_

```

Figure 44 – List cluster resources

```

C:\Program Files\NetApp\DataFabric Manager\DFM\examples>perl dfmcluster_add_resources.pl
-g "DFM Group" -i "Cluster IP Address" -n "Cluster Name" -k "Data Disk"
=== DataFabric Manager Cluster Configuration Script ===
Stopping services ...

```

Configuring cluster resources ...

=== DataFabric Manager Node Configured Successfully for Clustering ===

- Select the DFM Group from Failover Cluster management and online the services as shown in Figure 44.

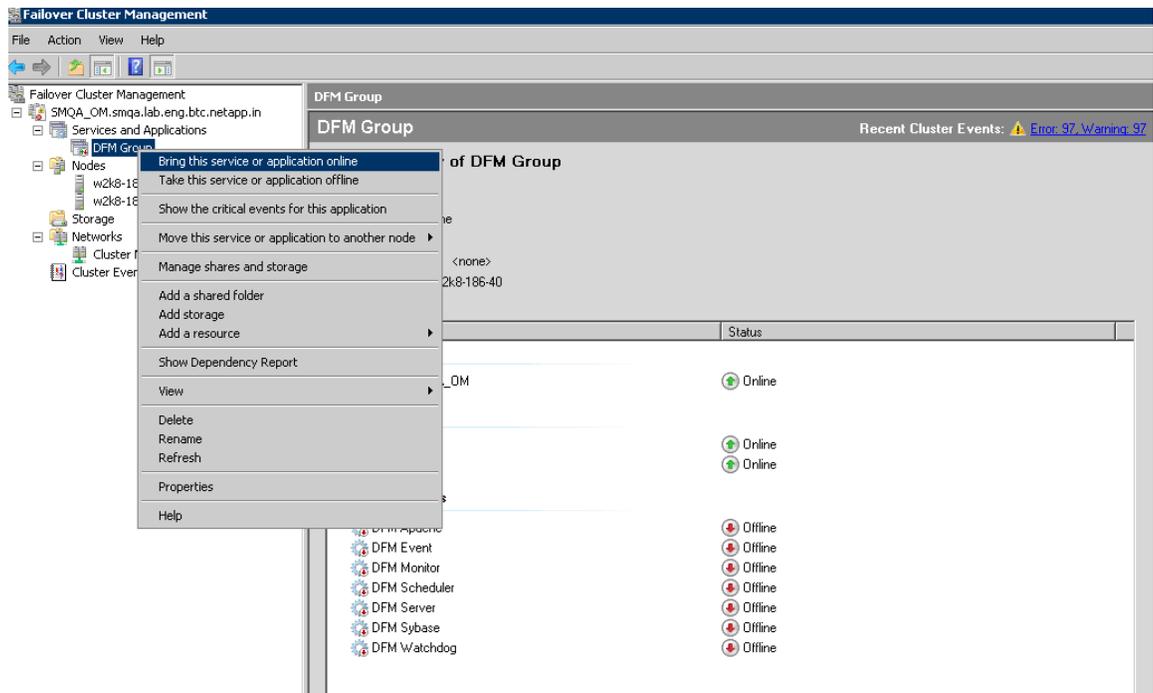


Figure 45 – Online DFM services

- Stop the DataFabric Manager services by `dfm service stop` command.
Note: This make sure that the dfm services do not try to access the data disk, which is moved to the secondary cluster node.
- Using the cluster administrator now manually `move Cluster group` into the second node (that is, make the second node as the owner) as shown in Figure 45. Once the operation is successful, verify that the secondary node is the owner of the cluster resources.

Once the services are online move the resources to other node to check failover is working fine with DFM Server services.

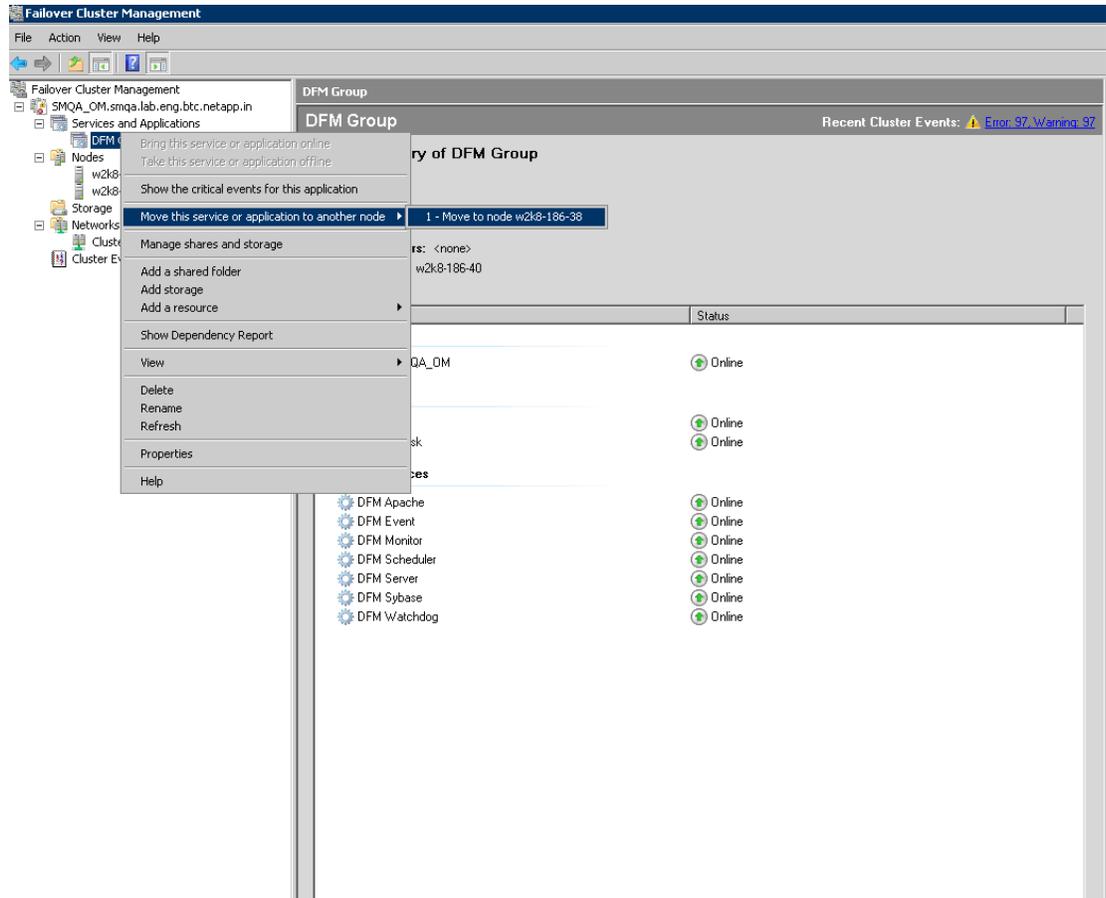


Figure 46 – Manual failover the cluster

2.3.6 CONFIGURING DATAFABRIC MANAGER SERVER TO USE THE CLUSTER NAME

By default, DataFabric Manager Server uses the local system name in the Web links present in alert e-mail messages sent. To make sure that the URLs are always accessible if failover happens, it is necessary to configure DataFabric Manager Server to use the cluster name instead of the local system name. To achieve this it is recommended to set the following options.

```
$ dfm option set localHostName=<FQDN-of-cluster>
```

2.3.7 MANAGING DATA FABRIC MANAGER SERVER WITH MSCS

This section explains how to manage DataFabric Manager Server with MSCS configured.

STOP AND START DATAFABRIC MANAGER SERVER SERVICES

Always use the cluster administrator to stop or start the services. Starting and stopping the services manually, will interfere with the cluster working.

RESTORE THE DATAFABRIC MANAGER SERVER DATA BASE

This section will explain how to restore DataFabric Manager Database with MSCS.

1. Log in to the cluster node currently owning the resources.
2. Bring down the DataFabric Manager Services offline using cluster administrator.
3. Restore the DataFabric Manager backup using `dfm backup restore <backup name>` command.

```
$ dfm backup restore backup_28_may.ndb
```

4. Once the restore is completed, online the DataFabric Manager services from the cluster administrator.

UPGRADING THE DATAFABRIC MANAGER SERVER

This section explains how to upgrade the DataFabric Manager Server with MSCS.

1. Log in to the cluster node currently owning the resources.
2. Offline the DataFabric Manager Server services using cluster administrator.
3. Upgrade the DataFabric Manager Server on this node.
4. Manually move cluster group to the second cluster node.
5. Upgrade the DataFabric Manager Server on the second cluster node.

CONFIGURING DATAFABRIC MANAGER SERVER WEB UI TO USE HTTPS

1. Log in to the first node of the cluster.
2. Take the DataFabric Manager Server Services offline using cluster administrator.
3. Start the DataFabric Manager Server sql service using `dfm service start sql` command.
4. Enter the `dfm ssl server setup` command to set up HTTPS, by creating an SSL certificate. This command creates two files, `server.crt` and `server.key`, in the `C:\Program Files\NetApp\DataFabric\DFM\conf` folder. Copy these files into the `C:\Program Files\NetApp\DataFabric\DFM\conf` folder at second cluster node.
5. Enable the https by `dfm option set httpsenabled=yes` command.
6. Start the DataFabric Manager Server http service using `dfm service start http` command.

Note: Starting the service by using `dfm service start` recreates the `httpd.conf` file with the changed options.

7. Stop the DataFabric Manager Server services using `dfm service stop` command.
8. Move the cluster group to the second node using cluster administrator.
9. Log in to the second cluster node and start the DataFabric Manager Server services. This will recreate `httpd.conf` with changed options.
10. Stop the DataFabric Manager Server service using `dfm service stop` command.
11. Online the DataFabric Manager Server services in the cluster administrator.

CHANGING HTTP OPTIONS

1. Log in to the first cluster node.
2. Take the DataFabric Manager Server services offline using cluster administrator.
3. Start the DataFabric Manager Server sql service using `dfm service start sql` command.
4. Enter the `dfm option set <option-name>=<option-value>` command to set the required HTTP.

```
$ dfm option set httpsPort=9443
```

5. Start the DataFabric Manager Server http service using `dfm service start http` command.
6. Stop the DataFabric Manager Server services using `dfm service stop` command.
7. Move the cluster group to the second node using cluster administrator.
8. Log in to the second cluster node and start the DataFabric Manager Server services. This will recreate `httpd.conf` with changed options.
9. Stop the DataFabric Manager Server service using `dfm service stop` command.
10. Online the DataFabric Manager Server services in the cluster administrator.

2.3.8 HIGH-AVAILABILITY USE CASES

This feature can be used to address the following use cases

SCENARIO 1

The active cluster node needs to be shut down for maintenance purpose or to test the failover is working fine. So a manual failover to the other node has to be performed to access the DataFabric Manager Server services.

- a. Launch cluster administrator from any of the cluster nodes.
- b. Select the cluster group and press move group to manually move to the other cluster node.

SCENARIO 2:

The cluster will automatically failover to the other node for the below reasons in the primary node.

1. Host down.
2. If any of the DataFabric Manager services stopped.
3. Data disk or quorum disk is not accessible.
4. Network failure.

3 HIGH-AVAILABILITY SUPPORT FOR DATAFABRIC MANAGER SERVER ON LINUX OR SOLARIS SERVER

This section describes how to set up and configure a high-availability solution for DataFabric Manager Server on Linux or Solaris servers. The section also lists the hardware and software required and the steps to be followed to provide high-availability.

3.1 DEPLOYMENT SCENARIO

In the example below, two independent Linux DataFabric Manager Servers are connected to a shared LUN on a NetApp storage system through FC Protocol. The DataFabric Manager Server data resides on the shared LUN. Veritas Cluster Server has been installed and configured on both the DataFabric Manager Servers. The two servers should have a redundant network connection called as heart-beat link. A virtual IP is configured on both the servers using which Operations Manager can be accessed.

For a system failure or any of the DataFabric Manager Server service failures on the active HA node, the Veritas Cluster Server will make an automatic failover to the other node.

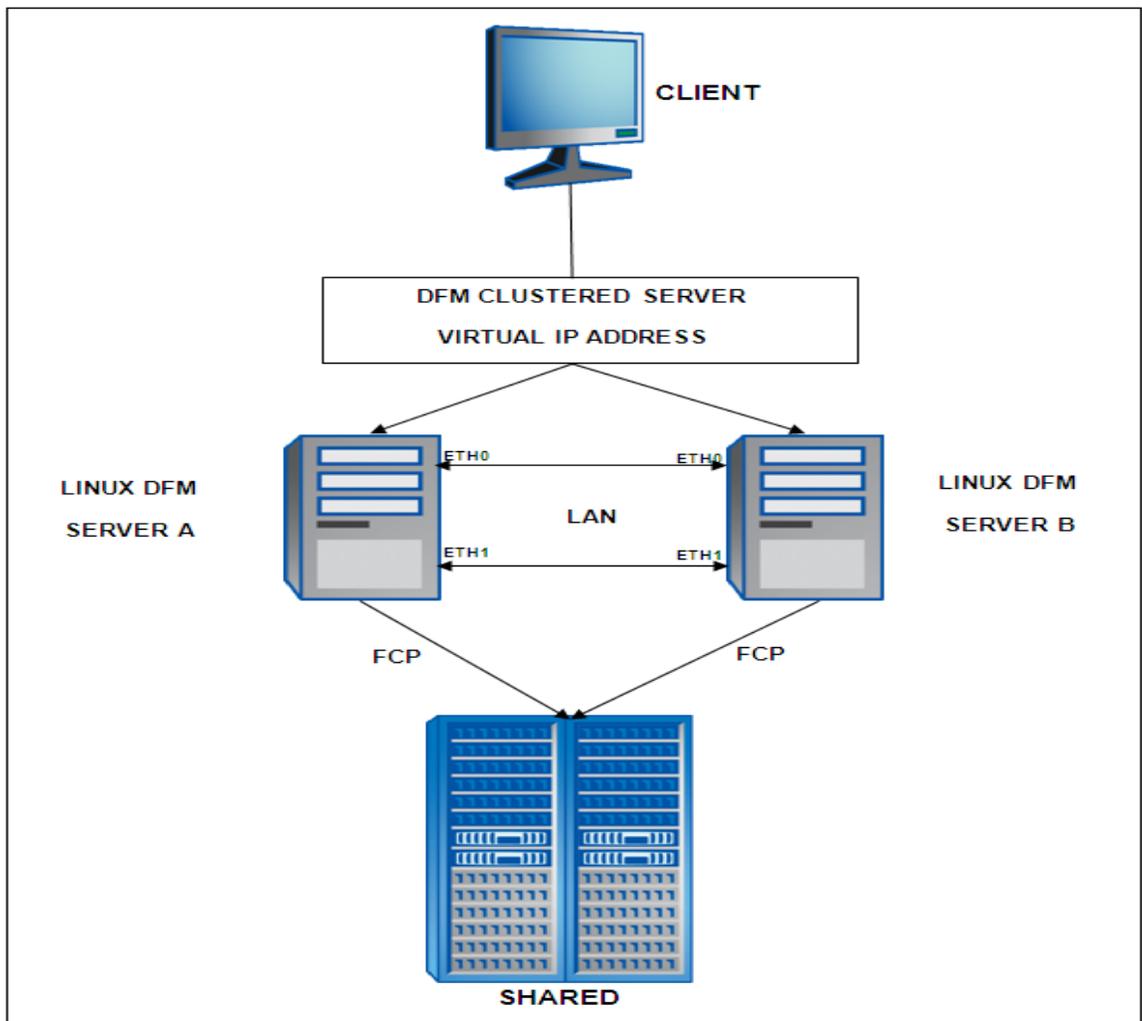


Figure 47 – VCS deployment scenario

3.2 REQUIREMENTS AND LIMITATIONS FOR VCS

This section describes the hardware and software requirements, limitations for VCS

- Minimum supported Linux Version is RHEL 4 update 3, SLES 9 with SP3. Veritas storage foundation and high-availability solutions 5.0 with MP1 (Maintenance Pack 1) should be installed.
- Supported on both Solaris 9 and 10.
- The same version of DataFabric Manager Server should be running on both the servers. Also the DataFabric Manager Server should be installed at the same location on both the servers (say default location /opt/NTAPdfm).
- On Solaris DataFabric Manager Server uses Veritas File System and Volume Manager, whereas on Linux it uses native ext3 file system and LVM.
- The DataFabric Manager servers need to have FCP connectivity to the storage system. iSCSI storage is not supported.
- Fibre Channel switched fabric based storage for shared data disks with a NetApp storage system as the storage back end
- At a minimum, two network interfaces should be set up on each system: one for node-to-node communication and the other for node-to-client communication. The network interface name used for node-to-client communication should be the same on both the systems.
- Only two nodes in the clustered set up; currently only two-node clustering is supported.
- SnapDrive for UNIX is an optional requirement for high-availability solution. LUN management can also be done using storage system command-line interface or FilerView and host commands. Refer to SnapDrive for UNIX compatibility matrix at <http://now.netapp.com/matrix/mtx/login.do>

3.3 SETTING UP HIGH AVAILABILITY FOR DATAFABRIC MANAGER USING VCS

To set up high-availability for DataFabric Manager Server using VCS, user needs to do the following:

- Install VCS software on both the cluster server nodes.
- Install SnapDrive for UNIX in both the cluster nodes.
- Create a shared LUN on NetApp storage system for holding DataFabric Manager data.
- Install DataFabric Manager on both the cluster nodes.
- Configure VCS using DataFabric Manager Server configuration scripts.

The below topics explain how to set up high-availability for DataFabric Manager Server using VCS.

3.3.1 VERITAS CLUSTER SERVER INSTALLATION

Install Veritas Cluster Server according to the instructions given in the Veritas Cluster Server 5.0 installation guide.

3.3.2 SNAPDRIVE INSTALLATION AND CONFIGURATION

This section explains the steps to install “SnapDrive for Unix”. Before proceeding to the below steps, make sure that the cluster nodes have HBA card. Cluster nodes and the storage system should be connected to the Fibre Channel switch and zoning configured correctly on the Fibre Channel switch.

- Install “SnapDrive for UNIX” as mentioned in “Install Upgrade Administration Guide” available:
<http://now.netapp.com/NOW/knowledge/docs/snapdrive/relunix30/html/software/install/install6.htm> for Linux cluster nodes
<http://now.netapp.com/NOW/knowledge/docs/snapdrive/relunix30/html/software/install/install8.htm> for Solaris cluster nodes

Note: "SnapDrive for UNIX" helps to simplify the SAN management. SAN management can be done even without using "SnapDrive for UNIX". It involves a series of steps. SnapDrive will make the SAN management easier. This document explains SAN management using "SnapDrive for UNIX".

- Edit the `snapdrive.conf` file in the SnapDrive installation directory (`/opt/NetApp/snapdrive`) and modify the below field. This will set the default transport as `fc`.

```
default-transport="fc"
```

- Run the SnapDrive command `snapdrive config prepare luns -count <no>` to prepare the LUNs.

```
#snapdrive config prepare luns -count 10
LUN IDs prepared adjusted to be: 18
now available (prepared, unreserved and not in use): 10
now reserved via available-lun-reserve: 8
```

- Run the SnapDrive command `snapdrive config set <user name> <filer name or ip>` to set the password for the storage system that will act as the datastore for DataFabric Manager data.

```
# snapdrive config set root 10.72.184.100
Password for root:
Retype password:
```

3.3.3 DATA DISK CREATION

This section explains how to create a shared disk on NetApp storage system to hold DataFabric Manager Server data:

- Log in to the first node of the cluster.
- Create a disk using the `snapdrive storage create` SnapDrive command.

```
# snapdrive storage create -lun 10.72.184.100:/vol/simulator/dfm_disk -fs /mnt/dfm -lunsize 1g
LUN f6070-184-100:/vol/simulator/dfm_disk ... created
mapping new lun(s) ... done
discovering new lun(s) ... done
LUN to device file mappings:
- f6070-184-100:/vol/simulator/dfm_disk => /dev/sdc
disk group dfm_SdDg created
host volume dfm_SdHv created
file system /mnt/dfm created
```

- Login to the second node of the cluster and connect to the LUN using `snapdrive storage connect` SnapDrive command.

```
#snapdrive storage connect -lun 10.72.184.100:/vol/simulator/dfm_disk -fs /mnt/dfm -hostvol
dfm_SdDg/dfm_SdHv
```

```
mapping lun(s) ... done
discovering lun(s) ... done
LUN 10.72.184.100:/vol/simulator/dfm_disk connected
- device filename(s): /dev/sdb
Importing dfm_SdDg
Connected fs /mnt/dfm
```

3.3.4 INSTALLING DATAFABRIC MANAGER SERVER WITH VCS

This section explains how to install DataFabric Manager Server on the cluster server nodes.

- Log in to the both the nodes of the cluster.
- Install DataFabric Manager Server on both the nodes.

Note: DataFabric Manager Server should be installed on the same path in both the cluster nodes (for example `/opt/NTAPdfm`).

- Make the DataFabric Manager Server service startup in manual mode by running the command `dfm service enable -m`.

Note:

This is very important as the services should not be started automatically on node reboot. Service starting and stopping should be done only by cluster service.

- Stop the DataFabric Manager Server services by `dfm service stop` command.

Note: This make sure that the DataFabric Manager Server services don't try to access the data disk, which is moved to the secondary cluster node.

3.3.5 SETTING UP DATAFABRIC MANAGER SERVER WITH VCS USING CONFIGURATION SCRIPTS

This section explains how to add DataFabric Manager Services, NIC card, volume group and virtual IP address into the Cluster Manager.

- Log in to the first node of the cluster.
- Run `dfm datastore setup <mount point>` to move the DataFabric Manager Server data into the LUN.

```
#dfm datastore setup /mnt/dfm/
Creating the destination data directories.
Required space for data:2.64 GB, Available space:18.6 GB
Stopping all services and jobs...
Adaptive Server Anywhere Stop Engine Utility Version 9.0.2.3397
Copying database files to /mnt/dfm/data.
Copying perf data files to /mnt/ dfm/perfdata.
Copying script output files to /mnt/ dfm/script-plugins.
Copying report output files to /mnt/simulator/reports.
```

```
Copying plugins files to /mnt/ dfm/plugins.
Copying cms files to /mnt/dfm/web/cms.
Copying software files to /mnt/dfm/web/software.
Changing database configuration settings ...
Changed dbDir to /mnt/dfm/data.
Updated dbLogDir to /mnt/dfm/data.
Starting sql service...
Using service restart timeout = 180 seconds.
Changed perfArchiveDir to /mnt/dfm/perfdata.
Changed scriptDir to /mnt/dfm/script-plugins.
Changed reportsArchiveDir to /mnt/dfm/reports.
Changed pluginsDir to /mnt/dfm/plugins.
Changed cmsDir to /mnt/dfm/web/cms.
Changed softwareDir to /mnt/dfm/web/software.
Changed databaseBackupDir to /mnt/dfm/data.
Starting services...
Completed datastore setup
```

- Stop the DataFabric Manager services by `dfm service stop` command.
- Set the path `/opt/VRTSvcs/bin`.

```
# PATH=/opt/VRTSvcs/bin:$PATH
```

- Run the script `dfmcluster_add_resources.pl` available at `/opt/NTAPdfm/examples`.

```
#cd /opt/NTAPdfm/examples
# [root@Inx186-117 examples]# perl dfmcluster_add_resources.pl -h Inx186-118 Inx186-117 -g VCS-186-125 -e eth0 -i 10.72.186.125 -n DFM -f /mnt/dfm -v dfm_SdHv -d dfm_SdDg -m 255.255.255.224
=== DataFabric Manager Cluster Configuration Script ===
Stopping services ...
Configuring cluster resources ...
Creating the service group VCS-186-125.
Configuring NIC and Virtual IP resource groups.
Configuring the data resource groups.
Configuring application resource groups.
Creating the dependency tree.
Starting sql service.
Setting the global options.
Stopping sql service.
```

=== DataFabric Manager Node Configured Successfully for Clustering ===

[root@lnx186-117 examples]#

Note: 1. The disk group and volume group information can be taken from `snapdrive storage show -all` command.

2. The script will add DataFabric Manager Server services, NIC card, virtual IP resource, mount point and volume, and disk groups into the Veritas Cluster Manager.

- Launch Cluster Manager window by running `hagui` command available at `/opt/VRTSvcs/bin`.

Note: 1. Before running the above command, make sure that the path variable set to `/opt/VRTSvcs/bin`.

2. Check that the DataFabric Manager Server services, mount point, virtual IP, NIC card and volume group have been added into the cluster group. Refer to Figure 47.

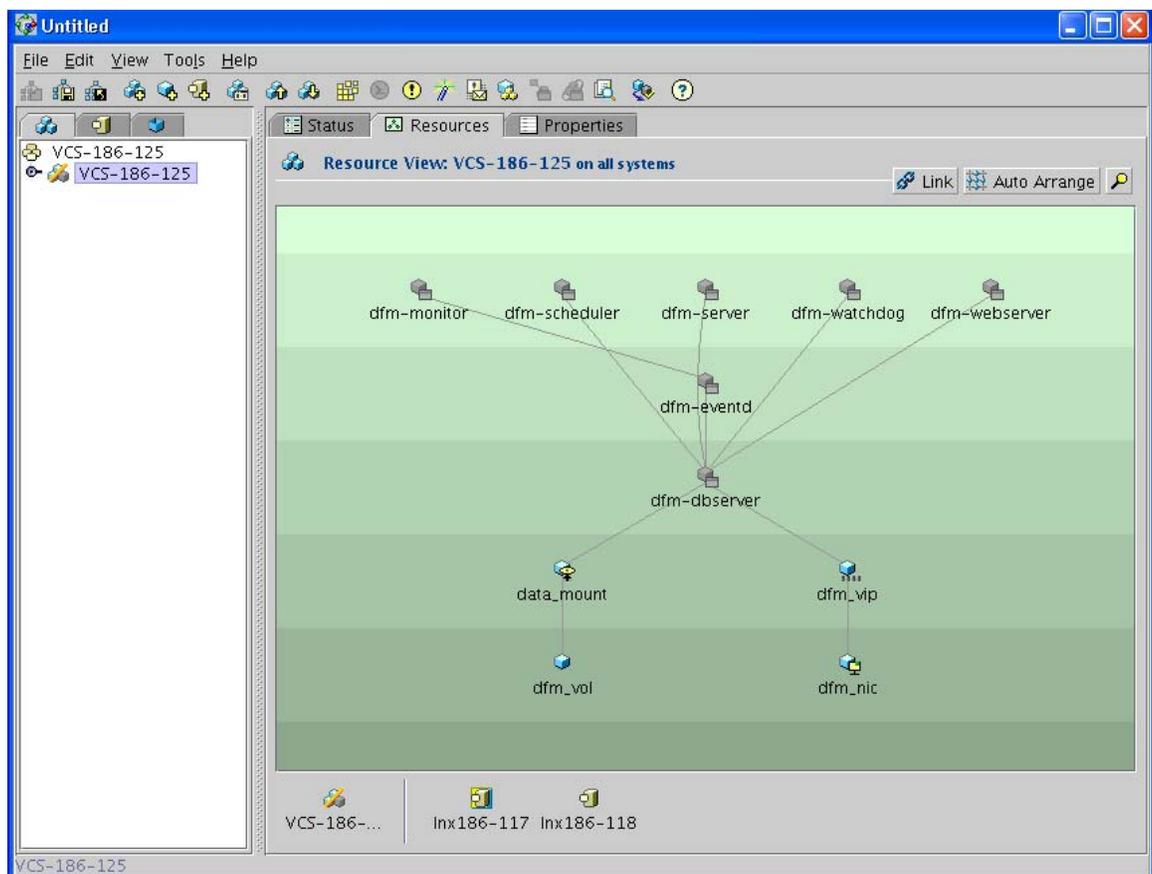


Figure 48 – View VCS resources

Note: We had configured the cluster manager with DataFabric Manager Server services, but we did not bring them online. So only the mount point, virtual IP, volume group and NIC card will be shown up in the cluster manager.

- Manually move the cluster to the secondary node and check if cluster failover is working as shown Figure 48.

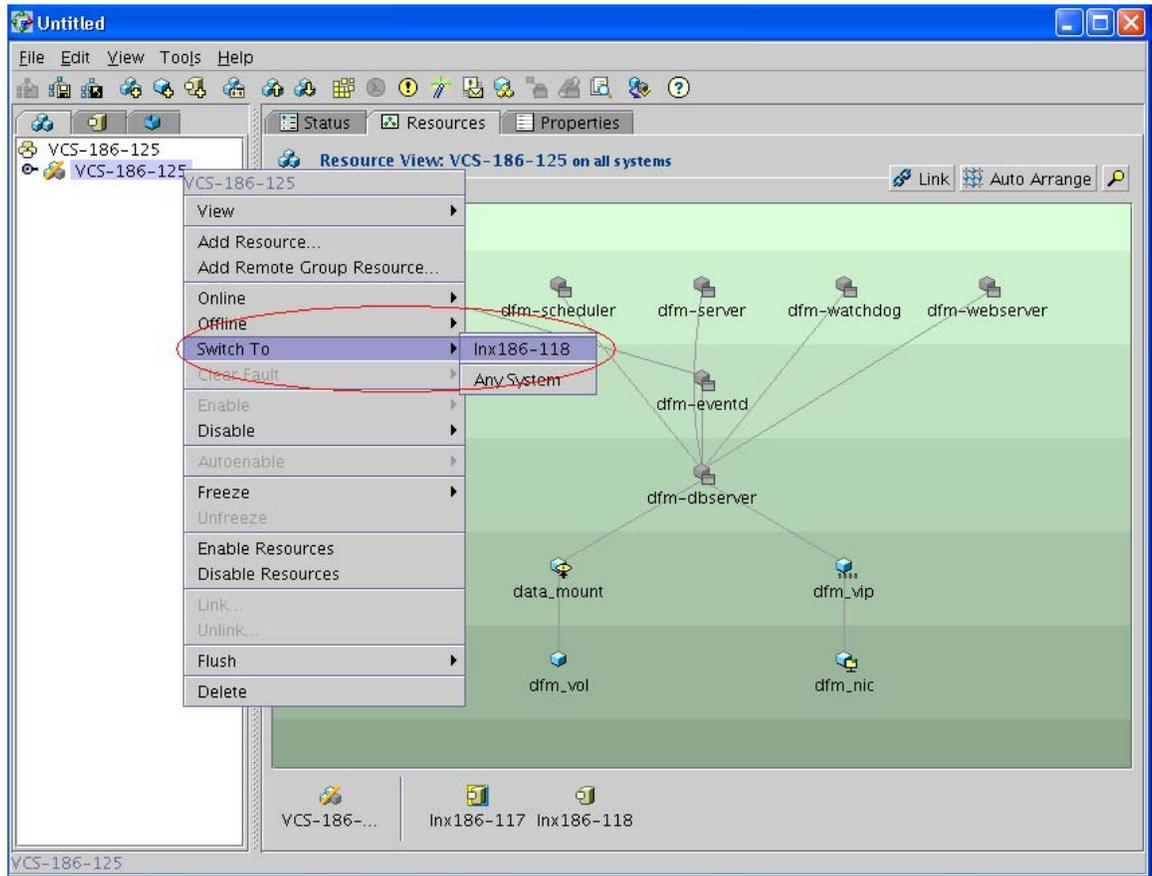


Figure 49 – Manual failover

- Check the Veritas Cluster Server has failed over to the second node of the cluster as shown in Figure 49.

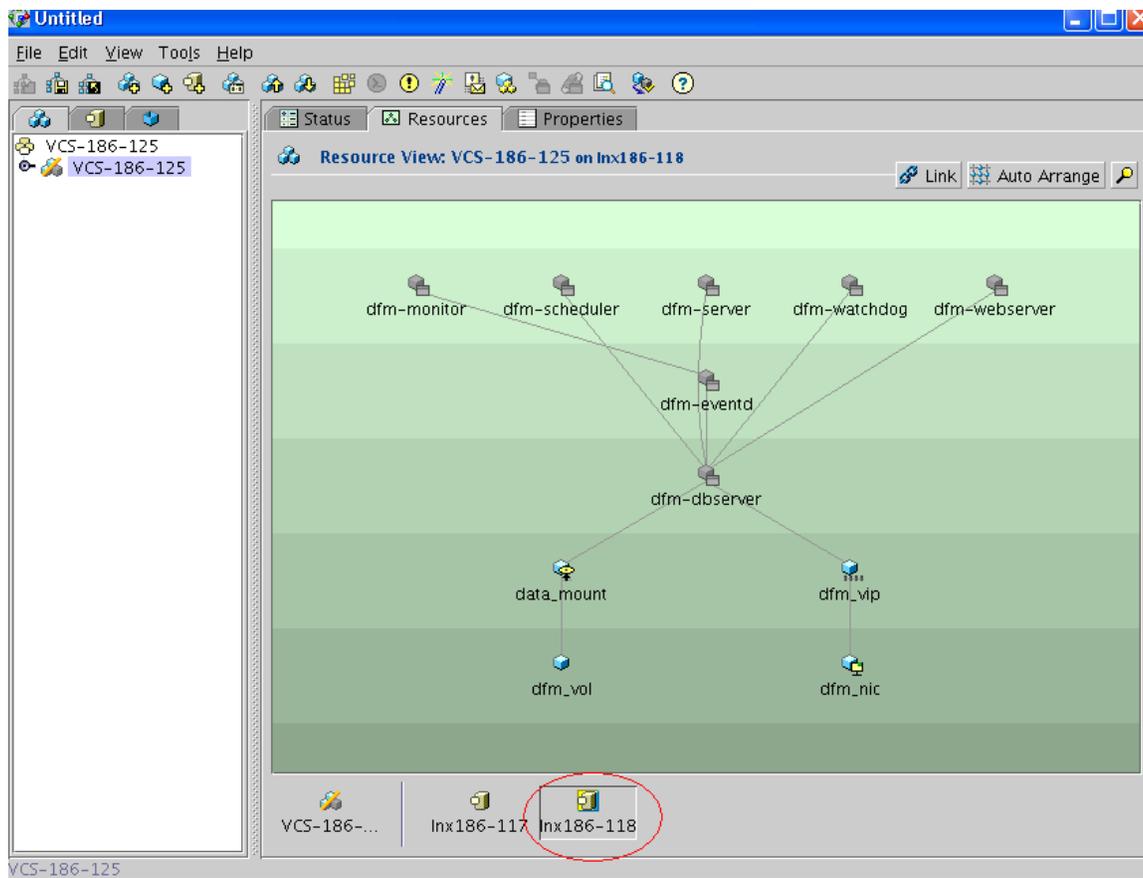


Figure 50 – VCS resources after failover

Note: Make sure that the data mount, virtual IP, volume group and NIC cards are up on the second node of the cluster.

- Log in to the second node of the cluster and run `dfm datastore setup -n <Mount Point>` command. Make sure that the same mount point is as used in first node of the cluster.

```
[root@Inx186-118 /]# dfm datastore setup -n /mnt/dfm
Stopping all services and jobs...
Changing database configuration settings ...
Updated dbLogDir to /mnt/kjag/data.
Starting sql service...
Using service restart timeout = 180 seconds.
Changed perfArchiveDir to /mnt/kjag/perfdata.
Changed scriptDir to /mnt/kjag/script-plugins.
Changed reportsArchiveDir to /mnt/kjag/reports.
Changed pluginsDir to /mnt/kjag/plugins.
Changed cmsDir to /mnt/kjag/web/cms.
Changed softwareDir to /mnt/kjag/web/software.
```

No database upgrade is necessary.

Redefining SQL Views ...

Completed data setup.

- Stop the DataFabric Manager Server services by `dfm service stop` command.
- Make the DataFabric Manager Server service startup in manual mode by running the command `dfm service enable -m`.

Note: This is very important as the services should not be started automatically on node reboot. Service starting and stopping should be done only by cluster service.

- Launch Veritas cluster manager. Bring the services online by selecting the cluster group and `Online - > <Server>` as shown in Figure 50.

Note: The Veritas Cluster Manager can be launched from any of the cluster nodes.

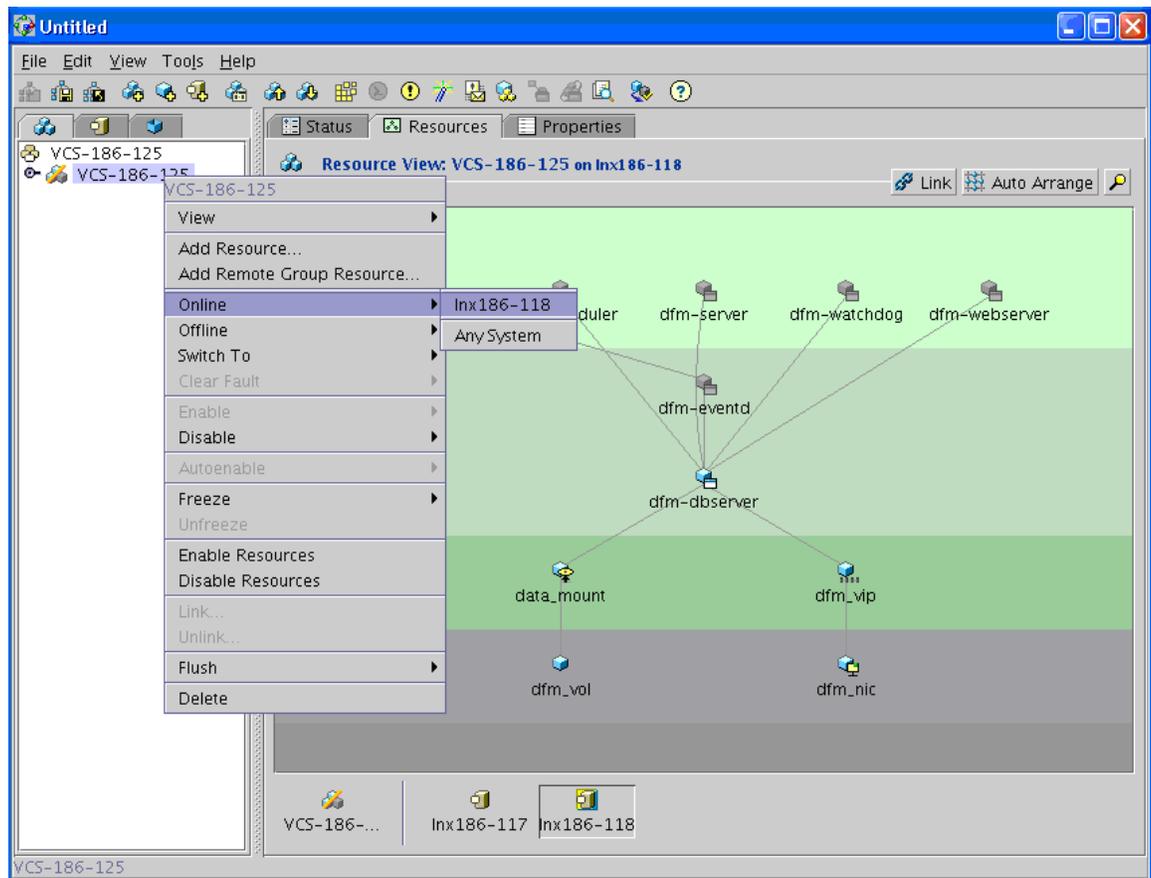


Figure 51 – Online dfm services on cluster node2

- Manually switch to the other node of the Veritas cluster as shown in Figure 27. Make sure that the DataFabric Manager Server services are starting properly on the other node of the cluster.

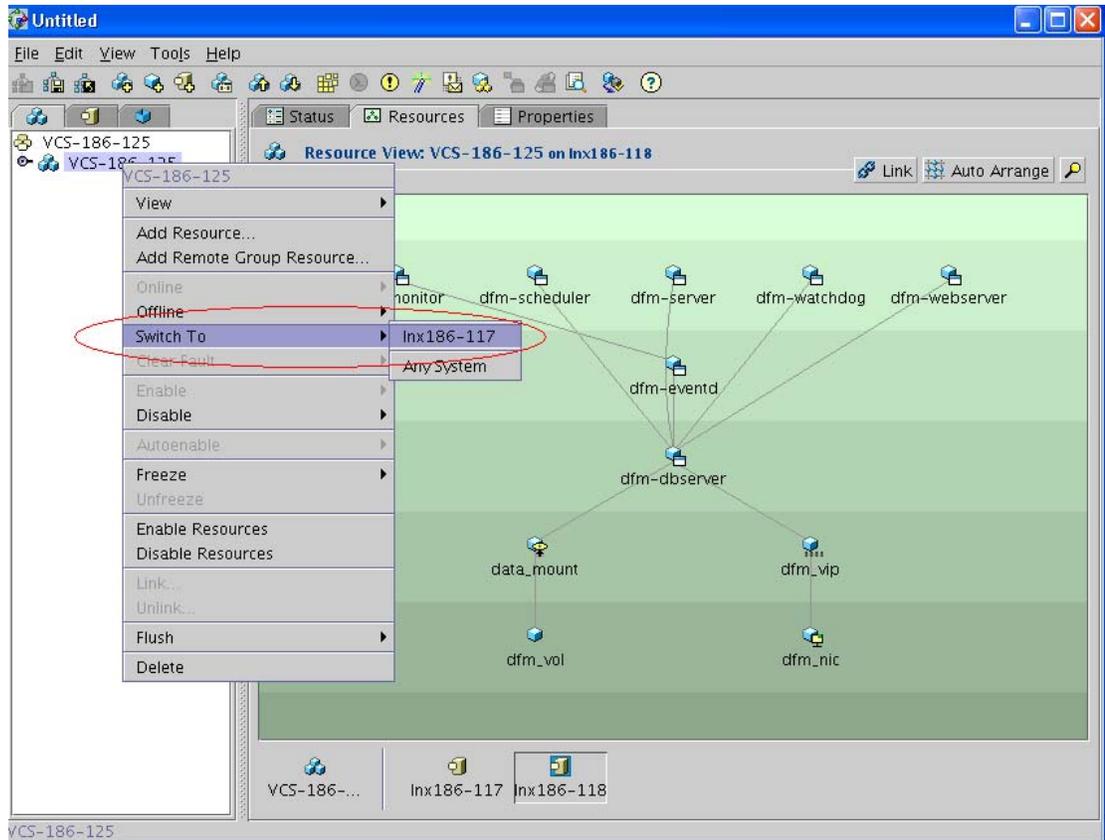


Figure 52 - Move cluster resources to second node

3.3.6 CONFIGURING DATAFABRIC MANAGER SERVER TO USE THE CLUSTER NAME

By default, DataFabric Manager Server uses the local system name in the Web links present in alert e-mail messages sent to administrators. To make sure that the URLs are always accessible if failover happens, it is necessary to configure DataFabric Manager Server to use the cluster name instead of the local system name. To achieve this it is recommended to set the following options.

```
$ dfm option set localHostName=<fqdn-of-cluster>
```

3.3.7 MANAGING DATAFABRIC MANAGER SERVER WITH VCS

This section explains how to manage DataFabric Manager Server with VCS configured.

STOP AND START DATAFABRIC MANAGER SERVER SERVICES

Always use the Veritas Cluster Manager to stop or start the services. Starting and stopping the services manually, will interfere with the cluster working.

RESTORE THE DATAFABRIC MANAGER SERVER DATABASE

This section will explain how to restore DataFabric Manager Server Database with VCS.

1. Log in to the cluster node currently owning the resources.
2. Bring down the DataFabric Manager Server services offline using `Cluster Manager`.
3. Restore the DataFabric Manager Server backup using `dfm backup restore <backup name>` command.

```
$ dfm backup restore backup_28_may_2008.ndb
```

4. Once the restore is completed, online the DataFabric Manager Server services from the `Cluster Manager`.

UPGRADING THE DATAFABRIC MANAGER

This section explains how to upgrade the DataFabric Manager Server with VCS.

1. Log in to the cluster node currently owning the resources.
2. Upgrade the DataFabric Manager Server.
3. Manually move cluster group to the second cluster node.
4. Upgrade the DataFabric Manager Server.

CONFIGURING DATAFABRIC MANAGER Server WEB UI TO USE HTTPS

1. Log in to the first node of the cluster.
2. Take the DataFabric Manager Server services offline using `cluster manager`.
3. Start the DataFabric Manager Server `sql` service using `dfm service start sql` command.
4. Enter the `dfm ssl server setup` command to set up HTTPS, by creating an SSL certificate. This command creates two files, `server.crt` and `server.key`, in the `/opt/NTAPdfm/conf` folder. Copy these files into the `/opt/NTAPdfm/conf` folder at second cluster node.
5. Enable the `https` by `dfm option set httpsenabled=yes` command.
6. Start the DataFabric Manager Server `http` service using `dfm service start http` command

Note: Starting the service by using `dfm service start` recreates the `httpd.conf` file with the changed options.

7. Stop the DataFabric Manager Server services using `dfm service stop` command.
8. Move the cluster group to the second node using `cluster manager`.
9. Log in to the second cluster node and start the DataFabric Manager Server services. This will recreate `httpd.conf` with changed options.
10. Stop the DataFabric Manager Server service using `dfm service stop` command.
11. Online the DataFabric Manager Server services in the cluster manager.

CHANGING HTTP OPTIONS

1. Log in to the first cluster node.
2. Take the DataFabric Manager Server services offline using `cluster manager`.
3. Start the DataFabric Manager Server `sql` service using `dfm service start sql` command.

4. Enter the dfm option set <option-name>=<option-value> command to set the required HTTP.

```
$ dfm option set httpsPort=9443
```

5. Start the DataFabric Manager Server http service using `dfm service start http` command.

6. Stop the DataFabric Manager Server services using `dfm service stop` command.

7. Move the cluster group to the second node using cluster administrator.

8. Log in to the second cluster node and start the DataFabric Manager Server services. This will recreate `httpd.conf` with changed options.

9. Stop the DataFabric Manager Server service using `dfm service stop` command.

10. Online the DataFabric Manager Server services in the cluster administrator.

4 GLOSSARY

- DFM: Data Fabric Manager
- MSCS: Microsoft Cluster Service
- VCS: Veritas Cluster Service
- FCP: Fibre Channel Protocol
- iSCSI: Internet Small Computer System Interface
- LUN: Logical unit number
- SDU: SnapDrive for UNIX
- SDW: SnapDrive for Windows



www.netapp.com

© Copyright 2009 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. NetApp, the NetApp logo, Go further, faster, DataFabric, Data ONTAP, FilerView, SnapDrive, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft and Windows are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. Solaris is a trademark of Sun Microsystems, Inc. Veritas is a trademark of Symantec Corporation or its affiliates in the U.S. and other countries. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3767