

# Creating a 2 node virtual SQL Server 2008 Cluster Configuration Using Windows 2008 MSCS

Reference No:  
Date: 16/02/2010  
File Name: Creating a 2 node SQL Server 2008 Cluster Configuration using Windows  
2008 MSCS.doc  
Version: 1.0

**Author:** Perry Whittle

### Modification Log

Work Order	Date	Author	Comments

## Contents

<b>1</b>	<b>PURPOSE OF THE DOCUMENT .....</b>	<b>4</b>
1.1	AUDIENCE .....	4
<b>2</b>	<b>CLUSTERING BRIEF OVERVIEW .....</b>	<b>4</b>
<b>3</b>	<b>CREATING THE WINDOWS CLUSTER .....</b>	<b>5</b>
3.1	INSTALL & CONFIGURE VMWARE SERVER .....	5
3.2	CREATE & CONFIGURE THE NAS VM.....	10
3.3	CREATE THE CLUSTER VM'S (DC AND CLUSTER NODES).....	27
3.4	ATTACHING THE ISCSI LUNS .....	35
3.5	INSTALLING THE WINDOWS 2008 CLUSTER .....	40
3.6	CREATE THE MICROSOFT DISTRIBUTED TRANSACTION COORDINATOR RESOURCE.....	54
<b>4</b>	<b>INSTALLING THE SQL SERVER INSTANCE.....</b>	<b>58</b>
4.1	INSTALLING THE FIRST NODE.....	58
4.2	ADD A CLUSTER NODE.....	68
4.3	SO HOW WOULD I HAVE AN ACTIVE\ACTIVE CONFIGURATION?.....	71
<b>5</b>	<b>APPENDIX A NAMED INSTANCES &amp; NETWORK NAMES.....</b>	<b>72</b>
5.1	NON CLUSTERED.....	72
5.2	CLUSTERED .....	72

# 1 PURPOSE OF THE DOCUMENT

This document explains how to create a virtual 2 node SQL Server cluster using VMWare Server 2.0.1, Windows 2008 Enterprise 64bit and SQL Server 2005 Enterprise 64 bit. For the purposes of this document the following apply;

Fail-Over	A Microsoft Cluster implementation method
SSMS	SQL Server Management Studio
T-SQL	Transact_SQL (the native SQL Server command language)
MSCS	Microsoft Cluster Services. The underlying technology for SQL Server Fail-Over clustering
NIC	Network interface card
Virtual Server Name	A unique computer name in the domain. During fail-over this computername is passed from one node to another
Virtual IP Address	A unique IP Address in the network. During fail-over this IP Address is passed from one node to another
NOS	Operating System
Node	A host which participates in a cluster
Quorum	The centralised repository used by MSCS
Heartbeat	A segregated private network for communication detection between nodes
Active\Passive	A cluster configuration which involves an active node and a passive node. The passive node becomes active on Fail-Over.

## 1.1 AUDIENCE

The document is intended to be accessible by Support representatives expressing a wish to learn more about Windows\SQL Server clustering with a view to supporting the SQL Server application under this platform. It is not expected that the reader is familiar with the Windows operating system and MSCS.

# 2 CLUSTERING BRIEF OVERVIEW

MSCS involves 2 or more computers (they don't have to be physical you can use virtual machines too) configured into a cluster relationship, however they do all have to use the same NOS (i.e. Enterprise or Datacentre). This technology requires a central, shared storage (it cannot exist on the machine itself). Clusters use Virtual Server Names and Virtual IP Addresses to create a reference for the network connection to the clustered application. For example SQL Node 1 has a computername of S-DBA-SQL-P01 and IP Address of 10.20.0.120. SQL Server instance, INST1 has a computername of S-DBA-SQL-C03 and IP Address of 10.20.0.126. All network calls to the SQL Server instance 1 are made through the virtual server name and IP address, not the Nodes actual name or IP Address. During fail-over this virtual name is de registered and re registered and along with the IP address passed to the partner Node like a ticket, re directing network calls to the new Node.

### 3 CREATING THE WINDOWS CLUSTER

MSCS for Windows 2008 has changed and is now more stringent upon the supported resources for clustering, especially storage. To emulate storage that supports SCSI-3 persistent reservations, we are using FreeNAS as a virtual SAN device. The first steps in this article, install and configure the VMWare Server hypervisor and then the FreeNAS virtual machine.

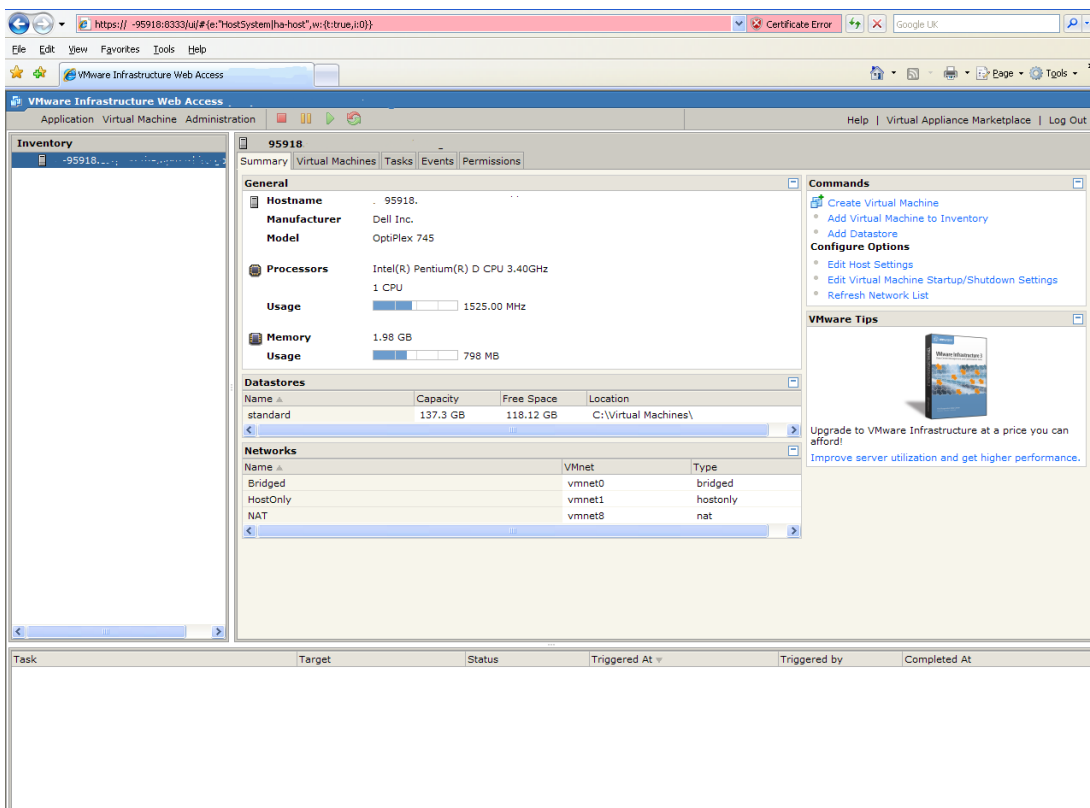
#### 3.1 INSTALL & CONFIGURE VMWARE SERVER

**Note: For this to be workable you must have sensible hardware available. You will need a multi socket, multi core machine and plenty of RAM and disk space.**

Install VMware Server 2.0.1 and then ensure you add your domain account to the following local group on your pc

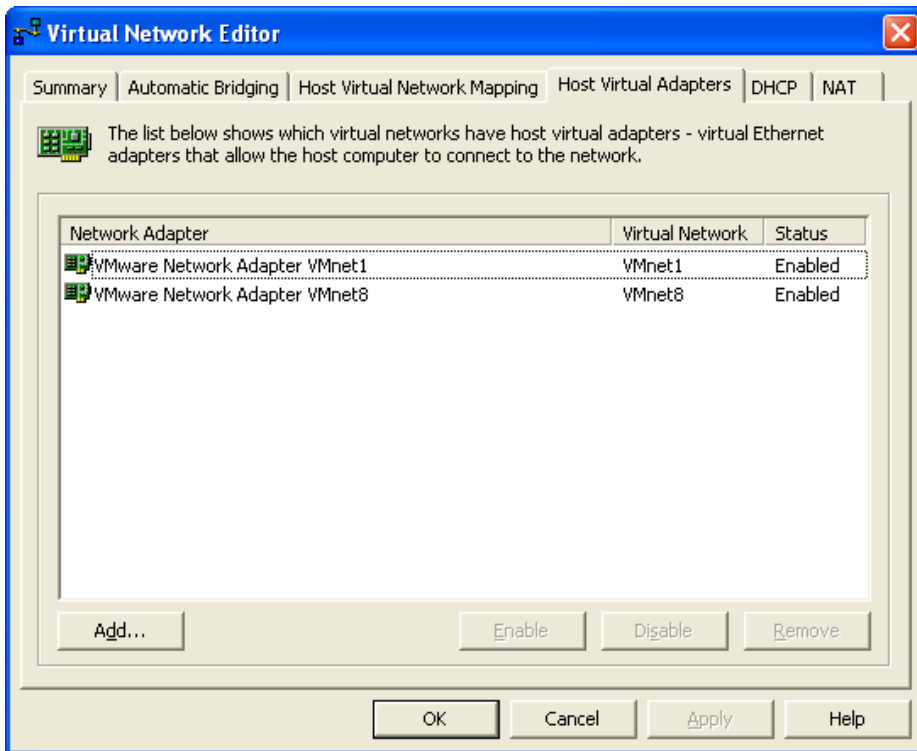
\_\_vmware\_\_

Login to the Vmware Server console supplying your domain account and password, you see the following;

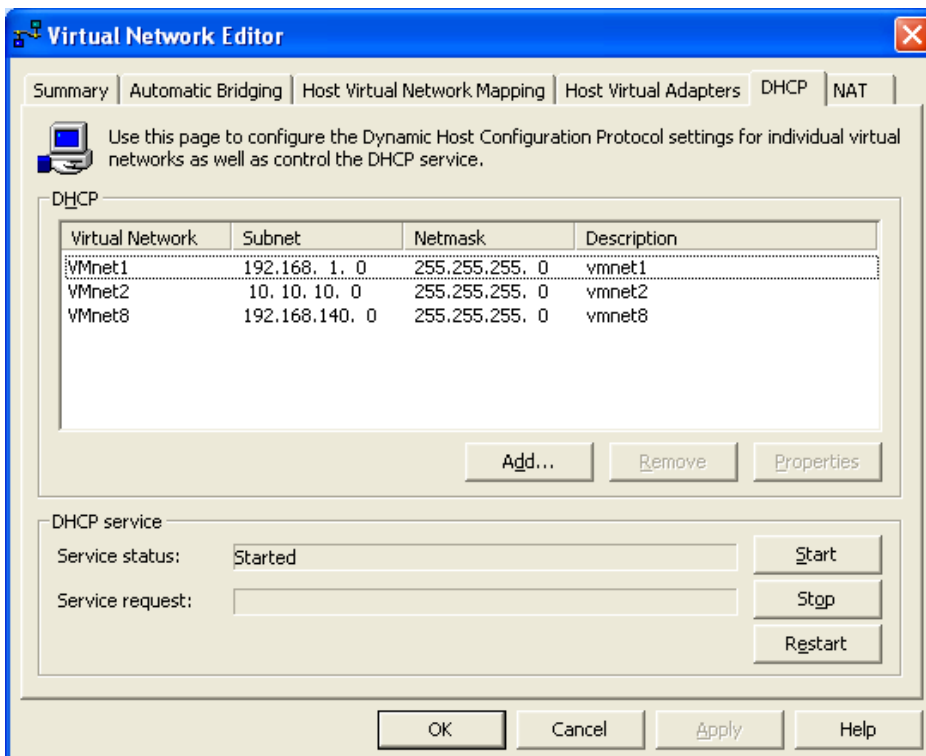


Open the VMWare “Manage Virtual Networks” option from the Start menu and go straight to the “Host virtual adapters” tab as shown below.

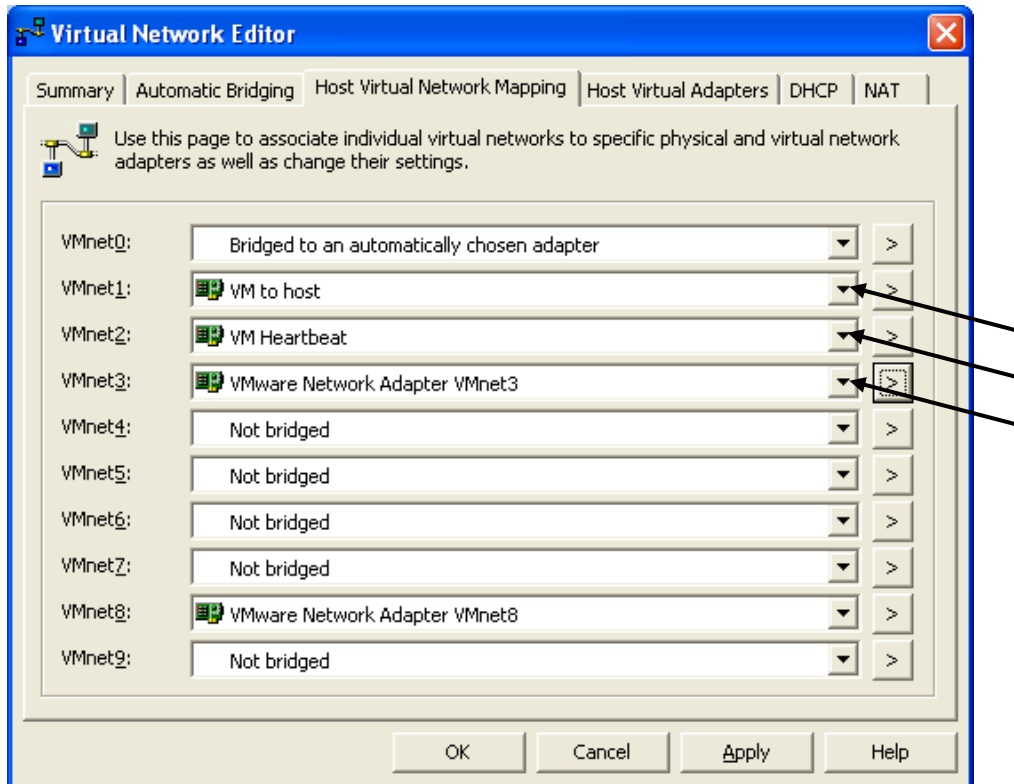
Add a new virtual adapter for VMnet2 and VMnet3 then click “Apply”




Now go to the DHCP tab and remove any DHCP assignments (click each item and remove) the click “Apply”.




Go to the “Host virtual network mapping” option and change the subnets to be used for VMnet 1, 2 & 3. Do this by clicking the arrows (indicated) and selecting “Subnet” from the pop-up menu




For VMnet1 set the IP address to 192.168.1.0



For VMnet2 set the IP address to 10.10.10.0

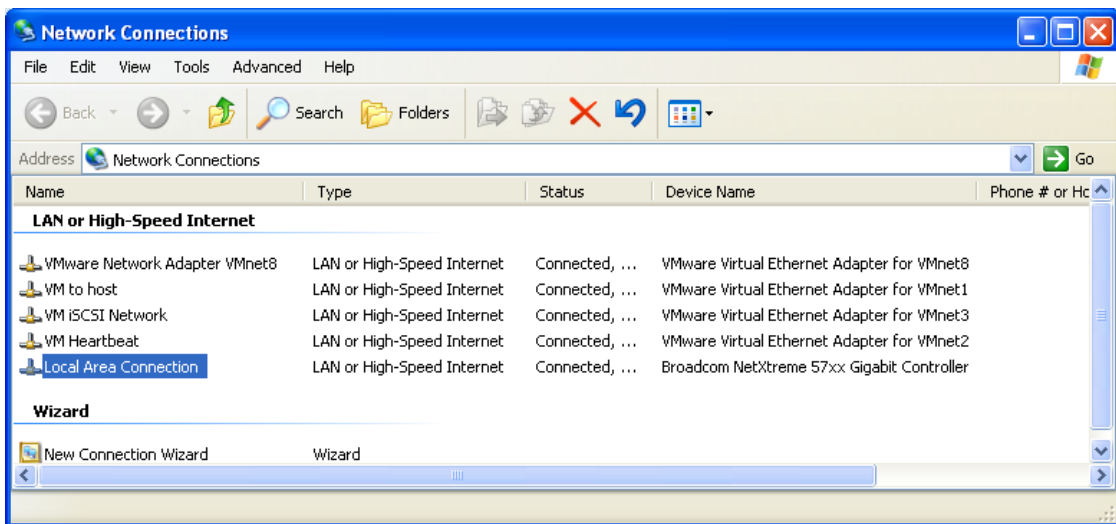


For VMnet3 set the IP address range to 192.168.93.0



Click “OK” to exit the virtual network editor

It’s a good idea to open your host machine network connections and rename the virtual LAN adapters to something a little more meaningful, as shown below;

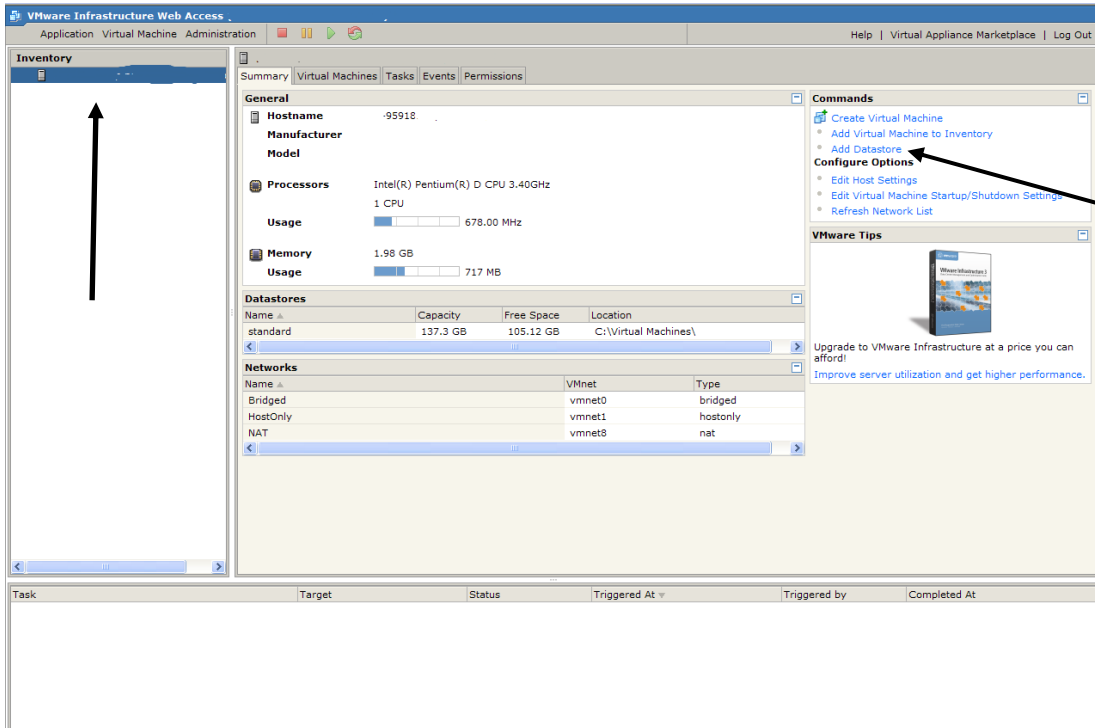


Click the “refresh networks list option” within the VMWare Server console to refresh the virtual networks.

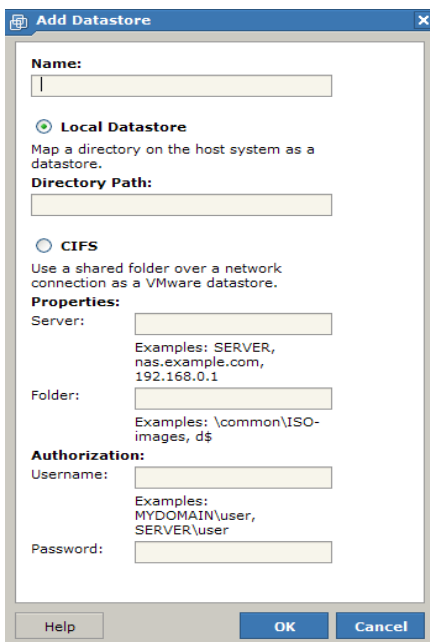


With the VMware Server console installed and configured and the virtual networks created, it's now time to create a datastore to hold all the ISO images we wish to install from.

Select the host node in the server console as shown below and under “Commands” click “Add datastore”;

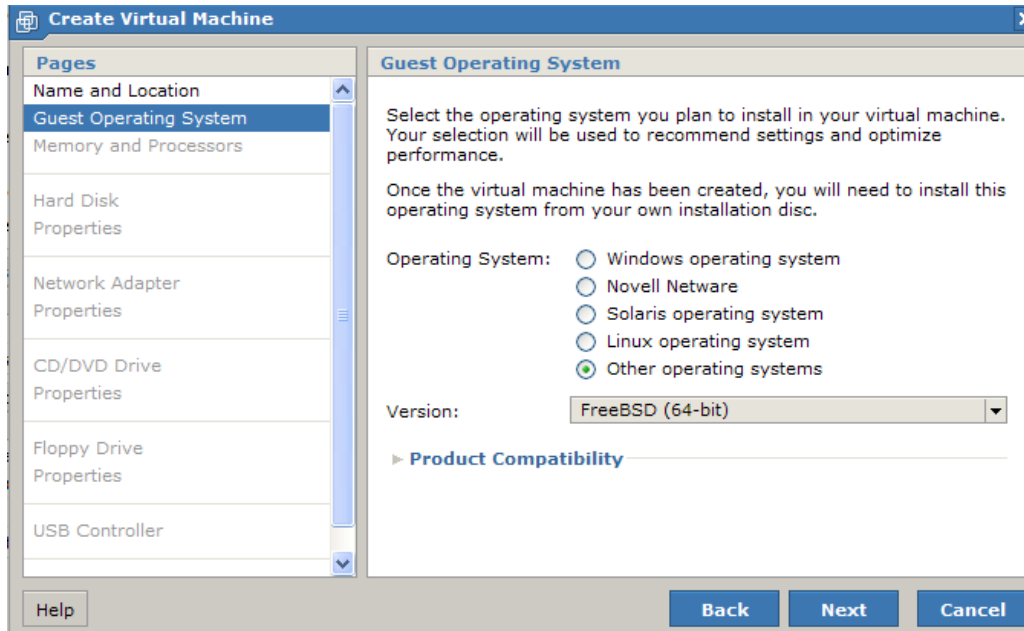


At the dialog supply a datastore name (ISOs) and a local path (C:\ISOs) then click “OK”;



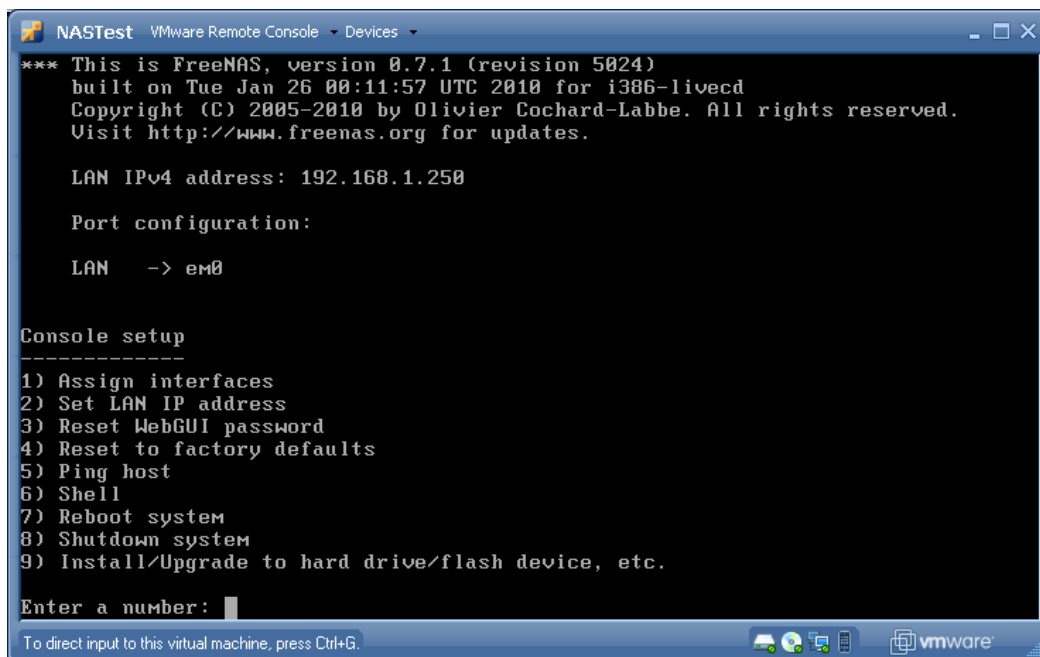
### 3.2 CREATE & CONFIGURE THE NAS VM

The first VM we need to create is the NAS server. Create a new VM with the following properties;

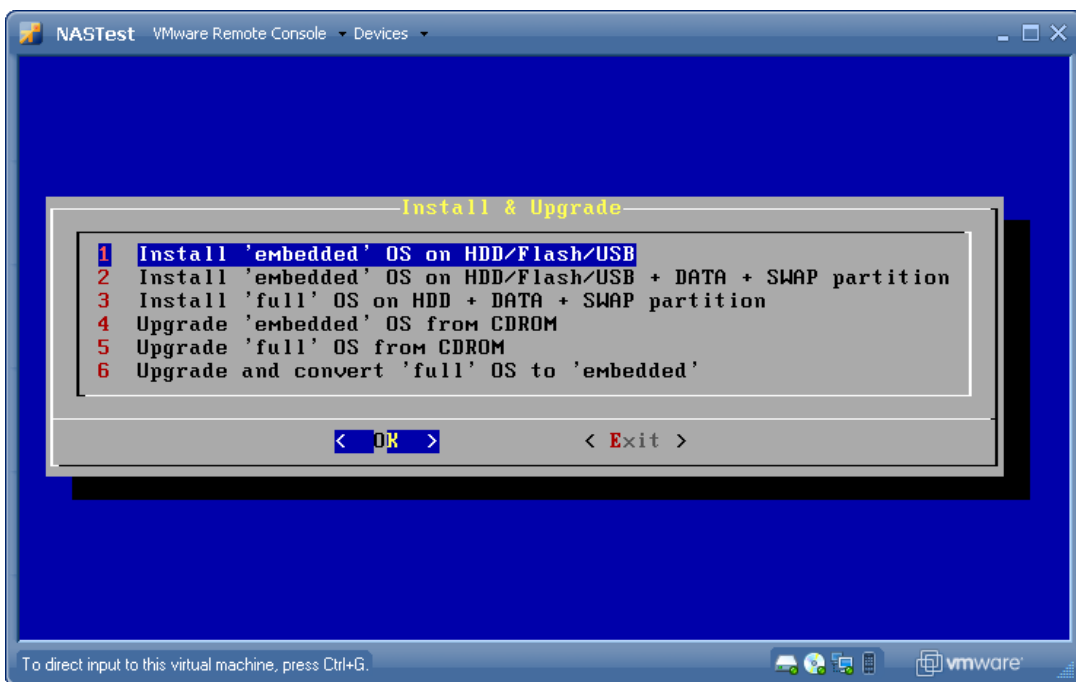


Add a 2GB virtual hard disk and a network adapter. Bind the virtual NIC to whichever virtual switch you want to use for the iSCSI network (VMNET2 in my case).

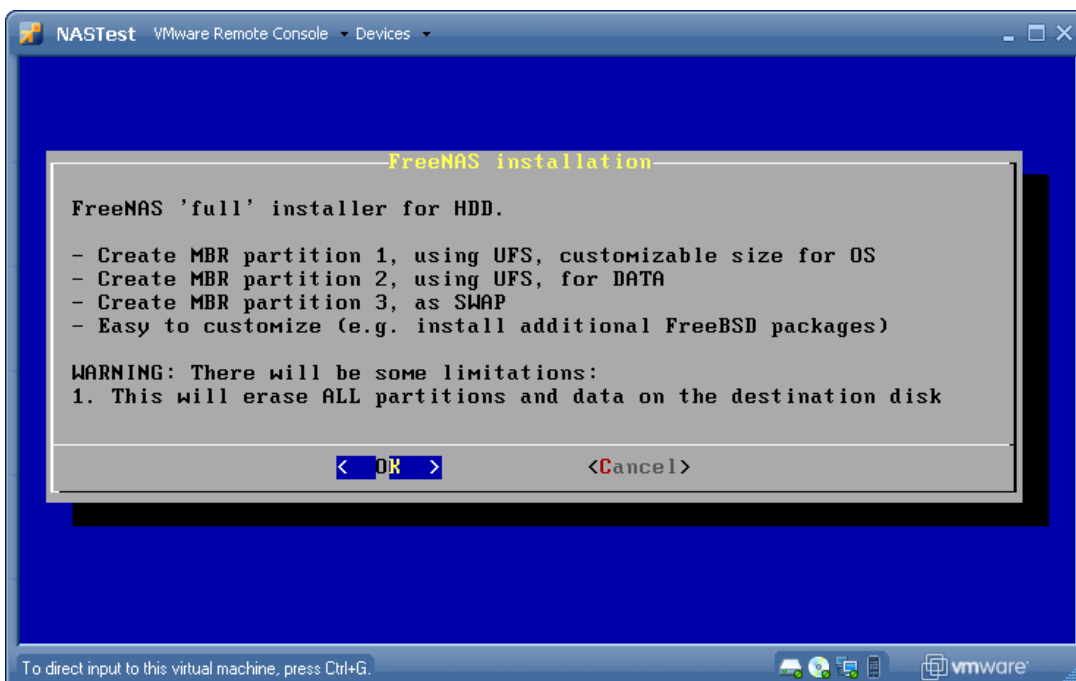
Boot the ISO image accepting all defaults. Once the NOS has booted (shown below), select option 9 from the menu.




You want to install the Full OS using the 3rd option.




Answer "OK" to continue




Select "OK" for the CD drive,




Select "OK" for the virtual disk,



Supply the OS partition size of 512MB,



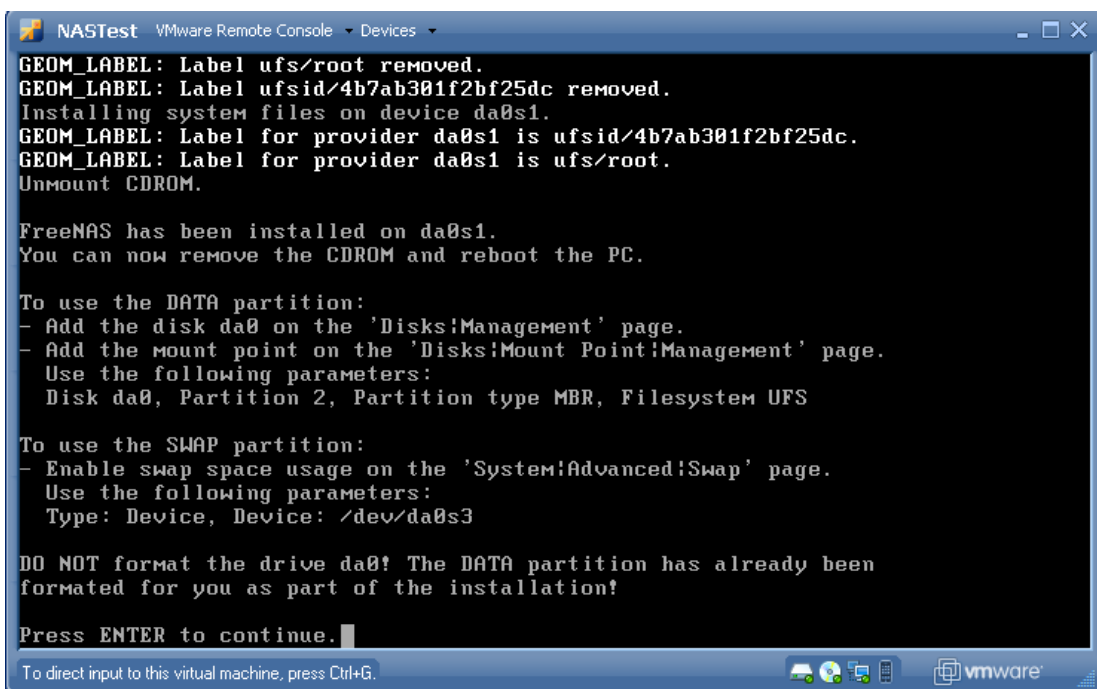
Select "yes" to add a swap partition,




Supply the swap partition size of 1024MB,



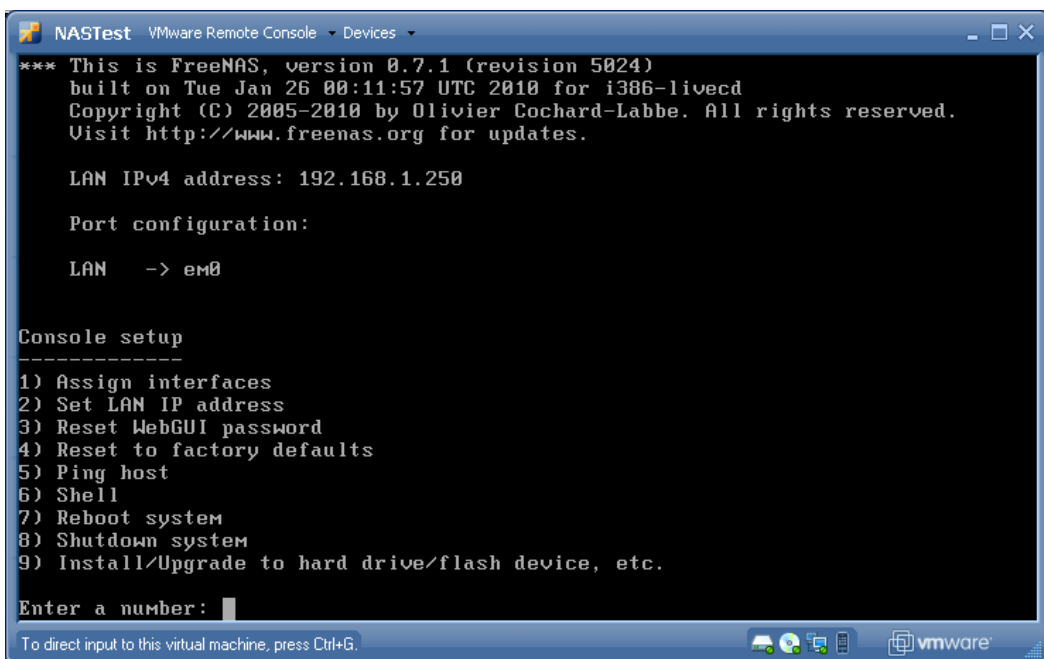
Installation completed, press "Enter" to continue.




Now select "Exit" to finish.



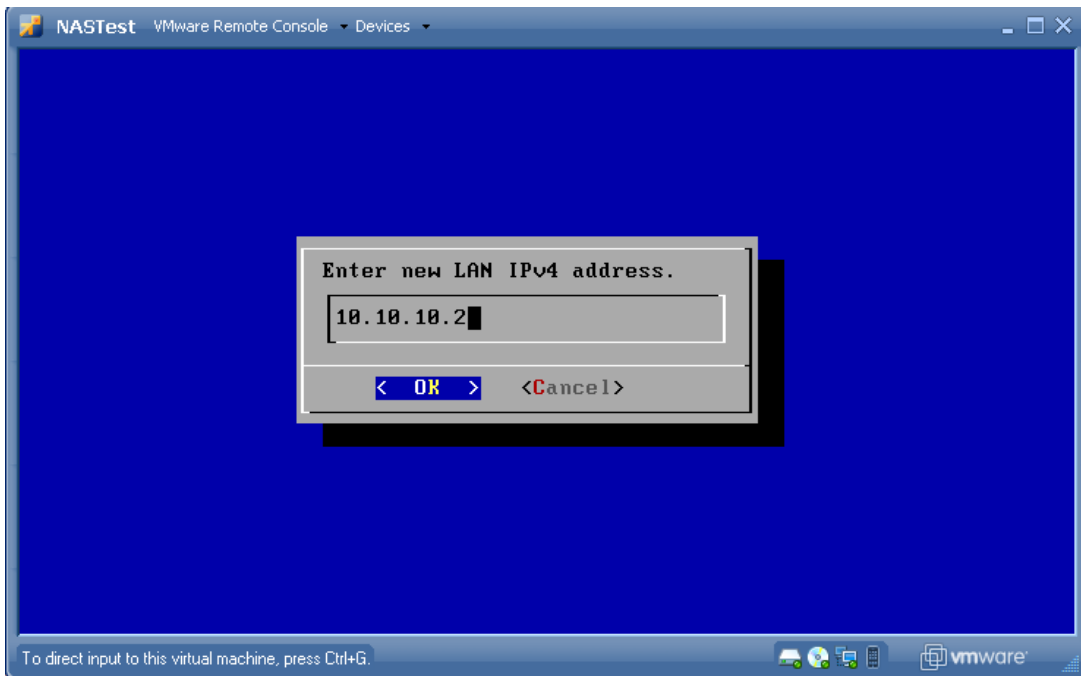
Now we need to set the LAN IP address. Select option 2 from the menu



The LAN configuration starts as follows. Select whether to use DHCP (no for our scenario)

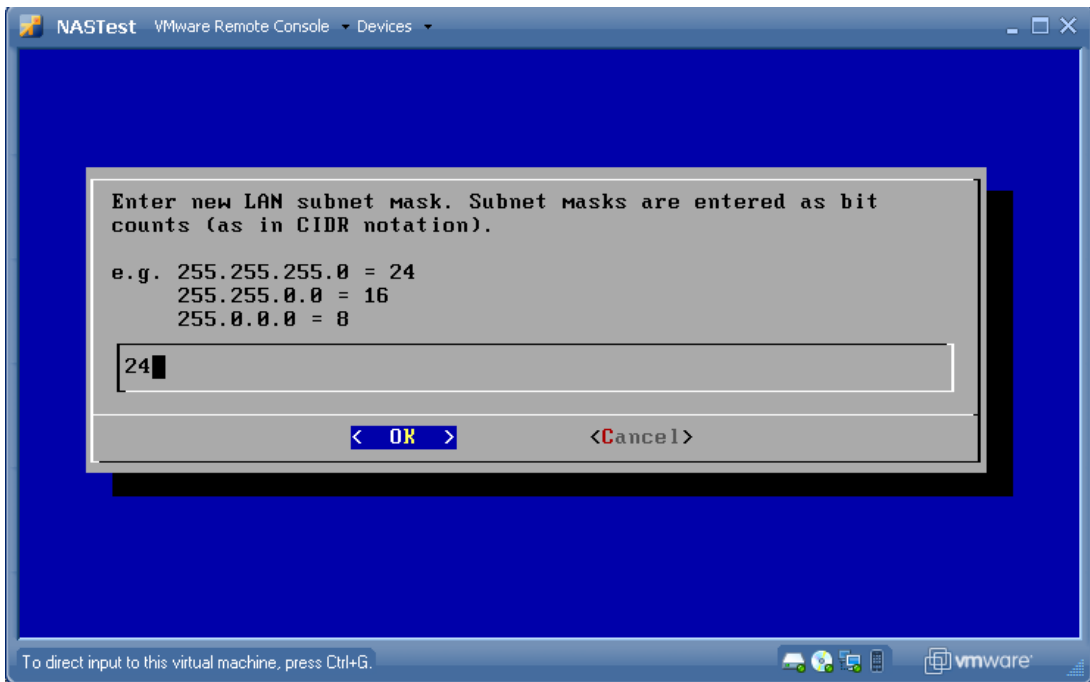


Enter the new LAN IP address






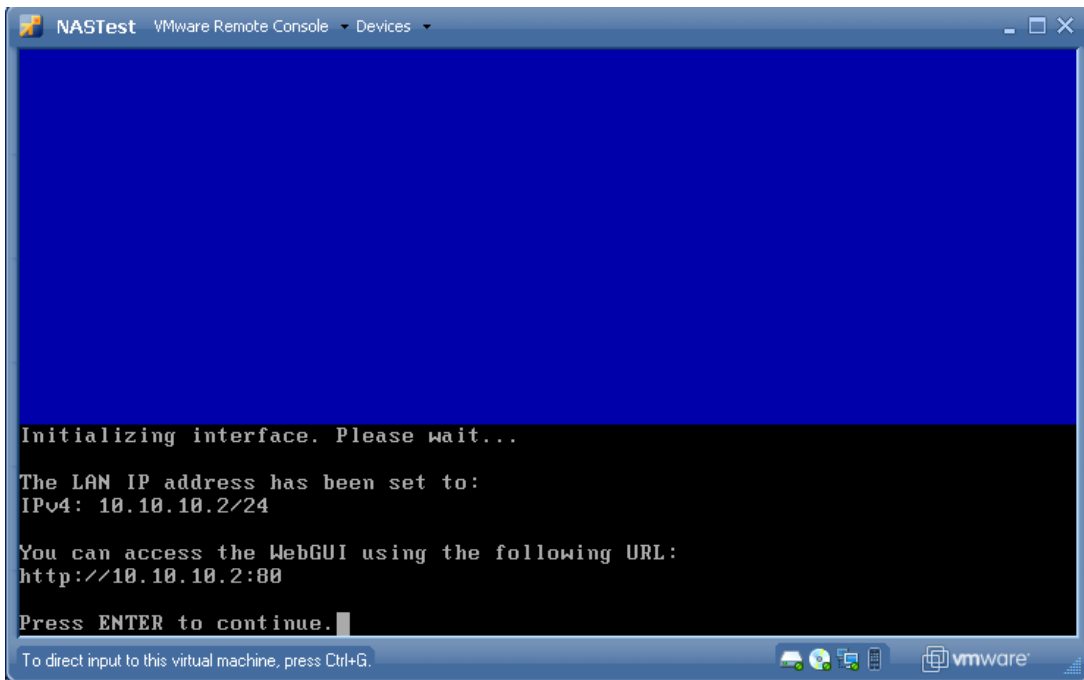
Supply the Network mask (using CIDR notation),



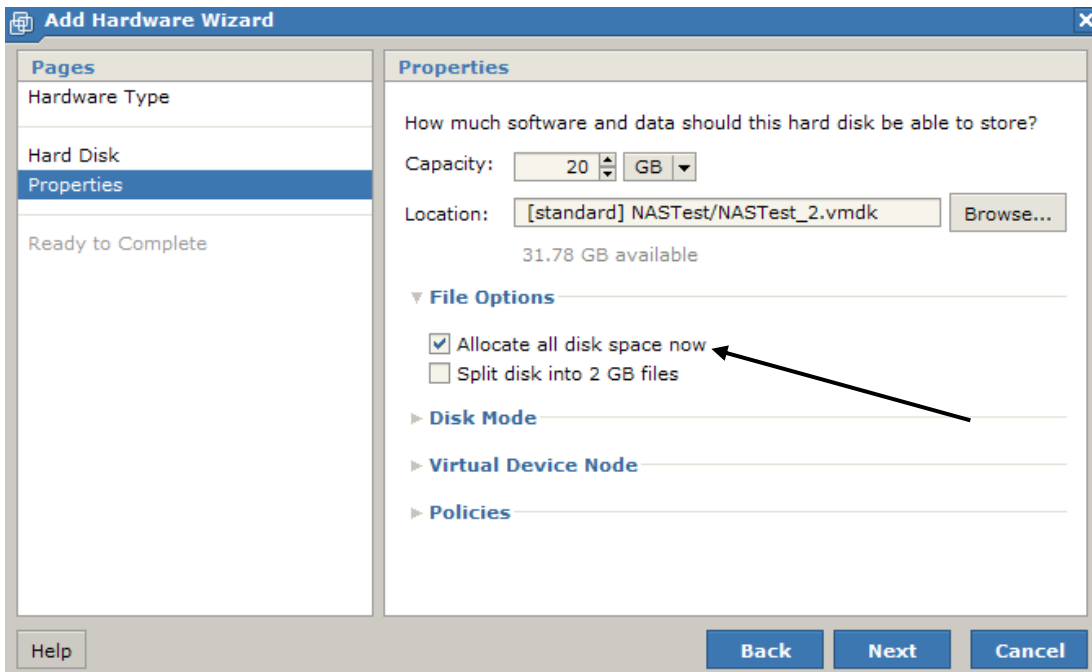
Select "no" for the IPv6 configuration,



Press "Enter" to continue,



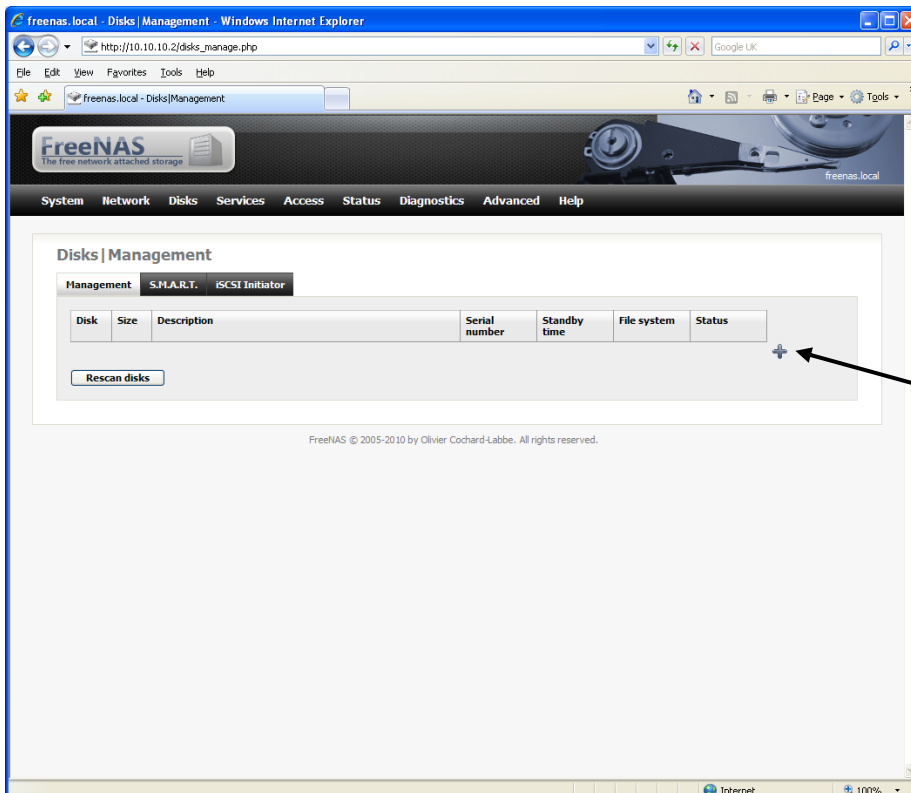
Shutdown the server using option 8 from the menu as we now have to create\attach the virtual hard disk to the NAS VM that will hold our iSCSI LUNs. Click the "Add hardware" option and add a new disk 20GB in size, pre allocating all disk space (shown below).



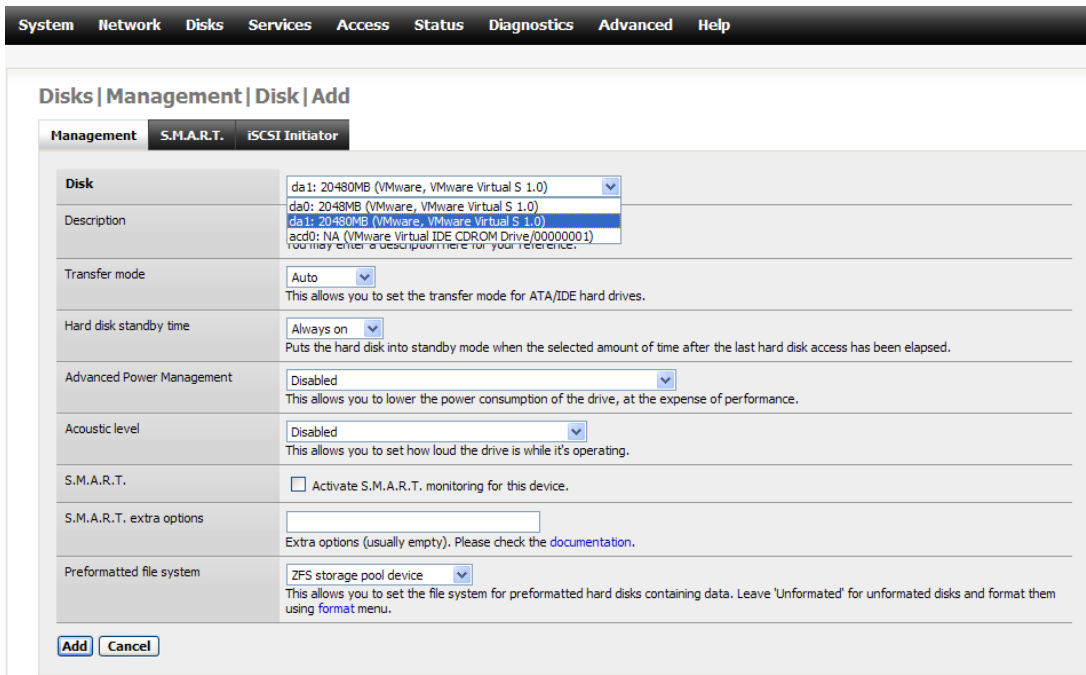
Once the disk has been added, boot the NAS VM and login to the NAS web management page using the username admin and password "freenas" (without quotes).



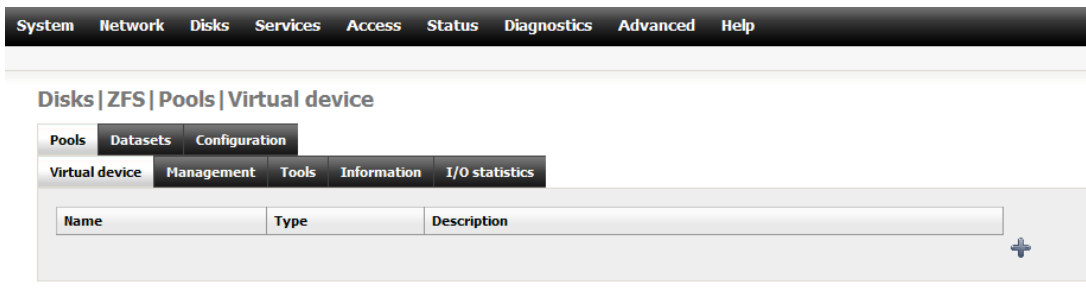
Go to "Disks" > "Management". Click the + sign as indicated,



Select the disk and use ZFS pre formatted file system, then click “add”. After adding the disk you must click the “Apply changes” button.



After adding the disk click “Disks” > “ZFS”. Select “Pools” > “Virtual device”. Click the + sign



Supply a device name and select the disk then click “Add”, you must click the “Apply changes” button afterwards.

**Disks | ZFS | Pools | Virtual device | Add**

**Pools** | **Datasets** | **Configuration**

**Virtual device** | **Management** | **Tools** | **Information** | **I/O statistics**

**Name**

**Type**

**Devices**

da1 (20480MB, VMware, VMware Virtual S 1.0)

**Description**

You may enter a description here for your reference.

Now click “Management” as shown below and then click the + sign

**Disks | ZFS | Pools | Management**

**Pools** | **Datasets** | **Configuration**

**Virtual device** | **Management** | **Tools** | **Information** | **I/O statistics**

Name	Size	Used	Free	Capacity	Health	AltRoot
+						

FreeNAS © 2005-2010 by Olivier Cochard-Labbe. All rights reserved.

Supply a name and select the virtual device then click “Add”, you must click the “Apply changes” button afterwards.

**Disks | ZFS | Pools | Management | Add**

**Pools | Datasets | Configuration**

**Virtual device | Management | Tools | Information | I/O statistics**

**Name:** Data\_Vol

**Virtual devices:** Data\_Vol (stripe)

**Root:**   
Creates the pool with an alternate root.

**Mount point:**   
Sets an alternate mount point for the root dataset. Default is /mnt.

**Description:**   
You may enter a description here for your reference.

**Add** **Cancel**

Once this is done you may then go to “Services” > “iSCSI target”. The first task is to enable the iSCSI Target by checking the box indicated and then clicking “Save and restart”.

**Services | iSCSI Target**

**Settings | Targets | Portals | Initiators | Auths | Media**

**iSCSI Target**  Enable

**Base Name:** iqn.2007-09.jp.ne.peach.istgt  
The base name (e.g. iqn.2007-09.jp.ne.peach.istgt) will append the target name that is not starting with 'iqn.'.

**Discovery Auth Method:** Auto  
The method can be accepted in discovery session. Auto means both none and authentication.

**Discovery Auth Group:** None  
The initiator can discover the targets with correct user and secret in specific Auth Group.

**Advanced settings**

**I/O Timeout:** 30  
I/O timeout in seconds (30 by default).

**NOPIN Interval:** 20  
NOPIN sending interval in seconds (20 by default).

**Max. sessions:** 32  
Maximum number of sessions holding at same time (32 by default).

**Max. connections:** 8  
Maximum number of connections in each session (8 by default).

**FirstBurstLength:** 65536  
iSCSI initial parameter (65536 by default).

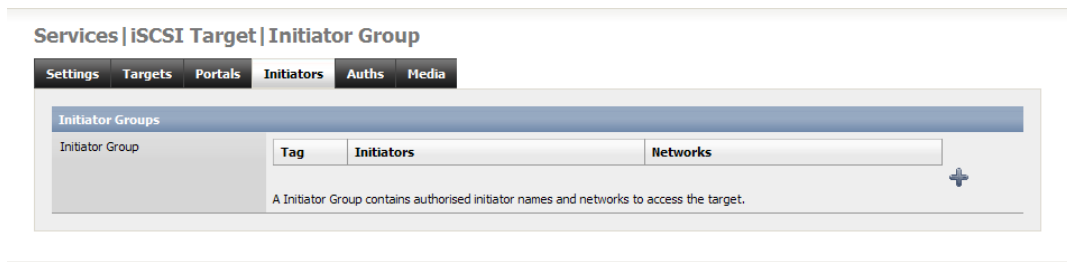
**MaxBurstLength:** 262144  
iSCSI initial parameter (262144 by default).

**MaxRecvDataSegmentLength:** 262144  
iSCSI initial parameter (262144 by default).

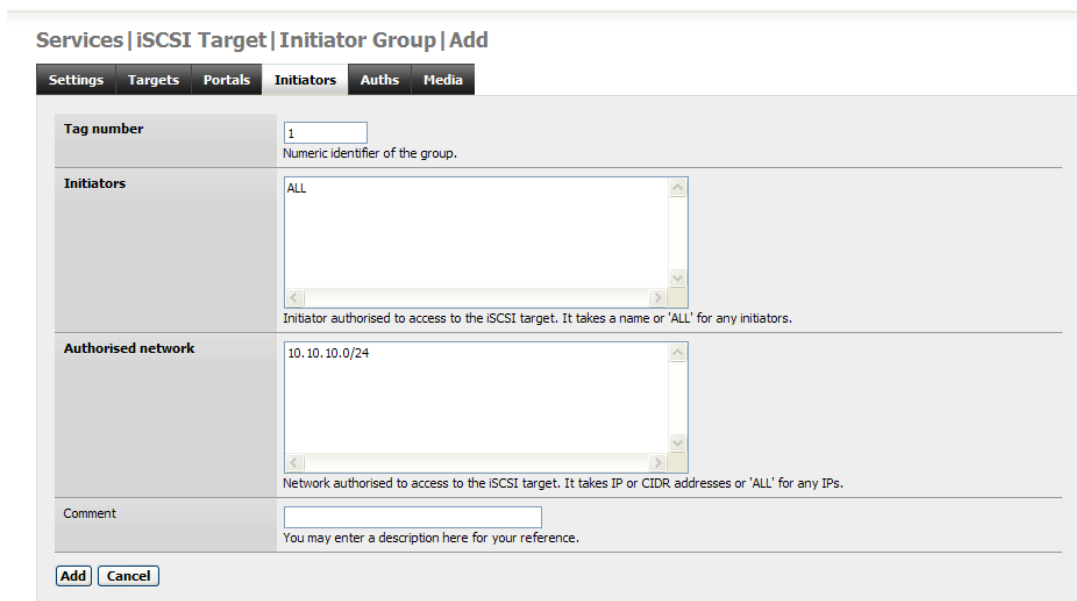
**iSCSI Target Logical Unit Controller**  Enable

**Save and Restart**


Go to the “Initiators” section and click the + sign.



If you want all initiators to connect click “Add” (you could lock down which machines connect by specifying their IDs here), you must click the “Apply changes” button afterwards.



Click “Portal group” and then click the + sign



Specify the portal (accept the default in our case) then click “Add”. You must click the “Apply changes” button afterwards.

Services | iSCSI Target | Portal Group | Add

Settings Targets Portals Initiators Auths Media

**Tag number**  
  
 Numeric identifier of the group.

**Portals**  
  
The portal takes the form of 'address:port'. For example '192.168.1.1:3260' for IPv4, '[2001:db8:1:1::1]:3260' for IPv6. the port 3260 is standard iSCSI port number. For any IPs (wildcard address), use '0.0.0.0:3260' and/or '[:]:3260'. Do not mix wildcard and other IPs at same address family.

**Comment**  
  
 You may enter a description here for your reference.

Now click “Targets”. You may now set up the LUNs you wish to present to the cluster nodes. Click the + sign to add an extent.

Services | iSCSI Target | Target

Settings Targets Portals Initiators Auths Media

**Targets**

Extent	Name	Path	Size
+			

Extents must be defined before they can be used, and extents cannot be used more than once.

Target	Name	Flags	LUNs	PG	IG	AG
+						

At the highest level, a target is what is presented to the initiator, and is made up of one or more extents.

**Note:**  
 To configure the target, you must add at least Portal Group and Initiator Group and Extent.  
 Portal Group which is identified by tag number defines IP addresses and listening TCP ports.  
 Initiator Group which is identified by tag number defines authorised initiator names and networks.  
 Auth Group which is identified by tag number and is optional if the target does not use CHAP authentication defines authorised users and secrets for additional security.  
 Extent defines the storage area of the target.

Supply the extent name, type, path and size. Remember to click “Apply changes” button when prompted.

Services | iSCSI Target | Extent | Add

Settings Targets Portals Initiators Auths Media

**Extent Name**  
  
 String identifier of the extent.

**Type**  
  
 Type used as extent.

**Path**  
   
 File path (e.g. /mnt/sharename/extent/extent0) used as extent.

**File size**  
   
 Size offered to the initiator. (up to 8EiB=8388608TiB. actual size is depend on your disks.)

**Comment**  
  
 You may enter a description here for your reference.



Extents added! Now click the + sign next to targets and configure those.

Services | iSCSI Target | Target

Settings Targets Portals Initiators Auths Media

The changes have been applied successfully.

**Targets**

Extent	Name	Path	Size	
	Quorum	/mnt/Data_Vol/Quorum	500MIB	
	SQLBak	/mnt/Data_Vol/SQLBak	7000MIB	
	SQLData	/mnt/Data_Vol/SQLData	9000MIB	
	SQLLog	/mnt/Data_Vol/SQLLog	3000MIB	

Extents must be defined before they can be used, and extents cannot be used more than once.

**Target**

Name	Flags	LUNs	PG	IG	AG

At the highest level, a target is what is presented to the initiator, and is made up of one or more extents.

**Note:**  
 To configure the target, you must add at least Portal Group and Initiator Group and Extent.  
 Portal Group which is identified by tag number defines IP addresses and listening TCP ports.  
 Initiator Group which is identified by tag number defines authorised initiator names and networks.  
 Auth Group which is identified by tag number and is optional if the target does not use CHAP authentication defines authorised users and secrets for additional security.  
 Extent defines the storage area of the target.

Configure the following options for each target (Quorum, SQLData, SQLLog, SQLBak).

Services | iSCSI Target | Target | Add

Settings Targets Portals Initiators Auths Media

**Target Name**   
Base Name will be appended automatically when starting without 'iqn.'

**Target Alias**   
Optional user-friendly string of the target.

**Type**   
Logical Unit Type mapped to LUN.

**Flags**

**Portal Group**   
The initiator can connect to the portals in specific Portal Group.

**Initiator Group**   
The initiator can access to the target via the portals by authorised initiator names and networks in specific Initiator Group.

**Comment**   
You may enter a description here for your reference.

---

**LUN0**

**Storage**   
The storage area mapped to LUN0.

Targets added! The target name is the label presented to the computer node over the iSCSI network.

**Services | iSCSI Target | Target**

Settings **Targets** Portals Initiators Auths Media

The changes have been applied successfully.

**Targets**

Extent	Name	Path	Size
	Quorum	/mnt/Data_Vol/Quorum	500MB
	SQLBak	/mnt/Data_Vol/SQLBak	7000MB
	SQLData	/mnt/Data_Vol/SQLData	9000MB
	SQLLog	/mnt/Data_Vol/SQLLog	3000MB

Extents must be defined before they can be used, and extents cannot be used more than once.

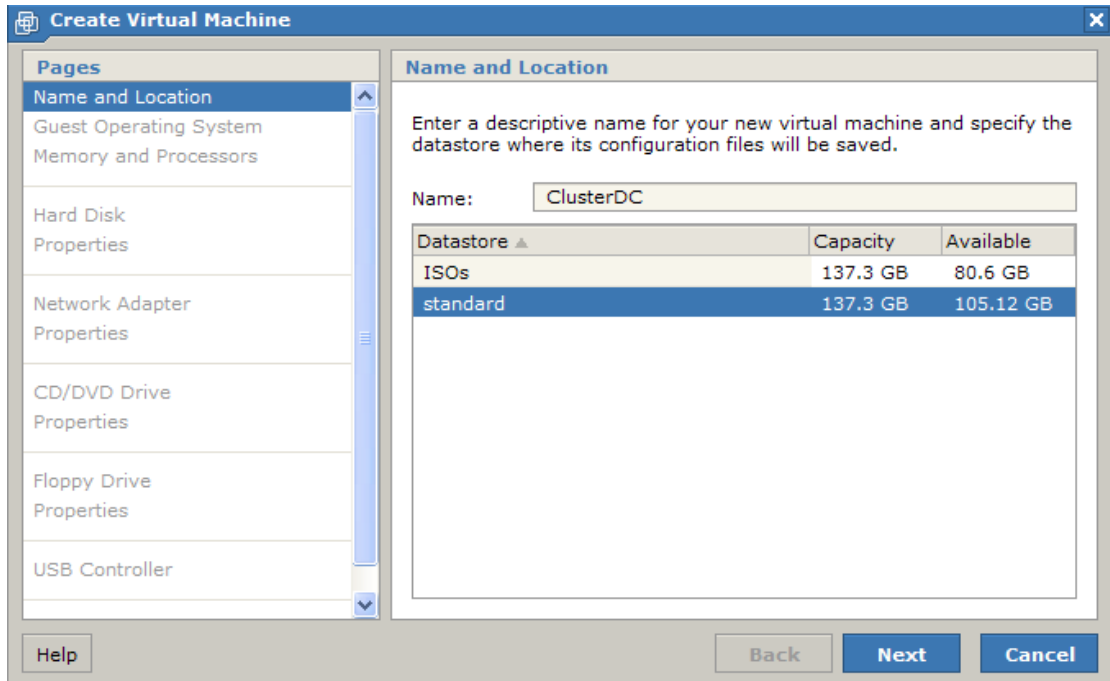
Target	Name	Flags	LUNs	PG	IG	AG
	iqn.2007-09.jp.ne.peach.istgt:Quorum	rw	LUN0=/mnt/Data_Vol/Quorum	1	1	none
	iqn.2007-09.jp.ne.peach.istgt:SQLBak	rw	LUN0=/mnt/Data_Vol/SQLBak	1	1	none
	iqn.2007-09.jp.ne.peach.istgt:SQLData	rw	LUN0=/mnt/Data_Vol/SQLData	1	1	none
	iqn.2007-09.jp.ne.peach.istgt:SQLLog	rw	LUN0=/mnt/Data_Vol/SQLLog	1	1	none

At the highest level, a target is what is presented to the initiator, and is made up of one or more extents.

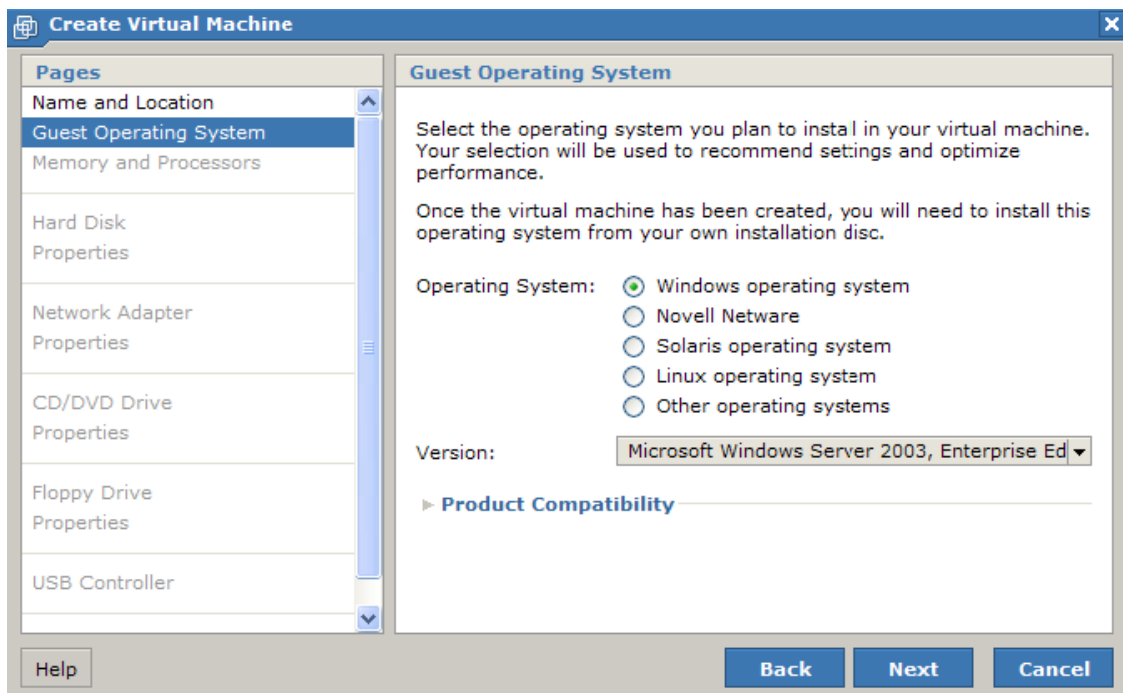
**Note:**  
 To configure the target, you must add at least Portal Group and Initiator Group and Extent.  
 Portal Group which is identified by tag number defines IP addresses and listening TCP ports.  
 Initiator Group which is identified by tag number defines authorised initiator names and networks.  
 Auth Group which is identified by tag number and is optional if the target does not use CHAP authentication defines authorised users and secrets for additional security.  
 Extent defines the storage area of the target.

### 3.3 CREATE THE CLUSTER VM'S (DC AND CLUSTER NODES)

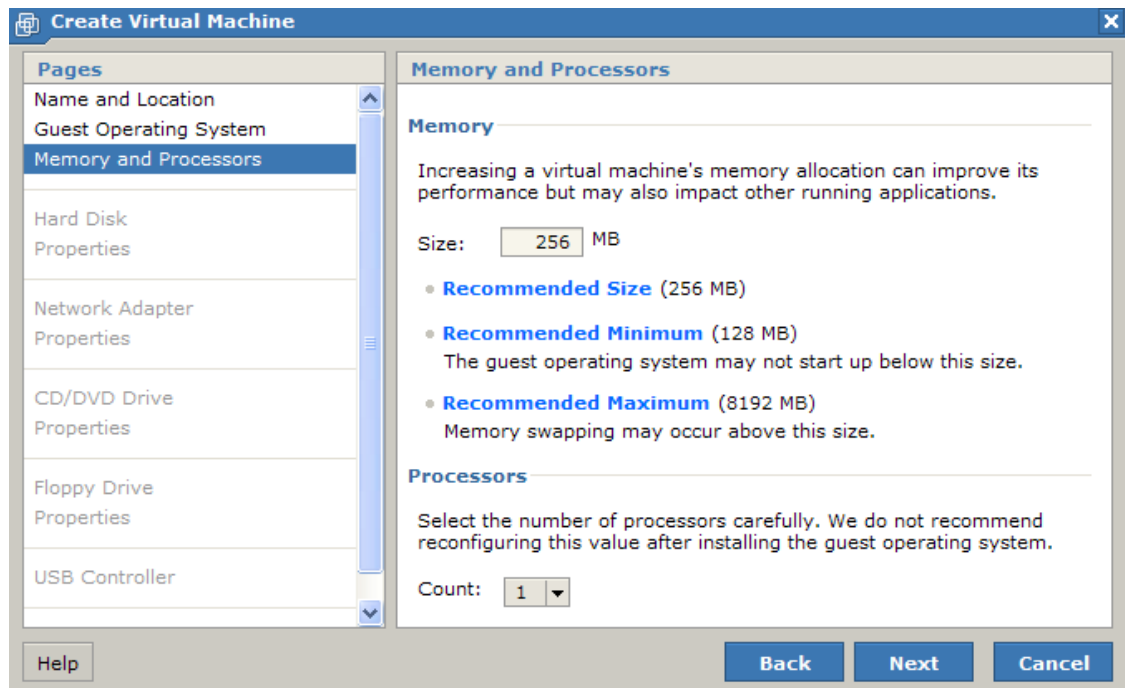
Now we are ready to start creating the VM's. Under the "Commands" section, select "Create virtual machine". Enter a Virtual machine name and select a datastore (this datastore maps to a folder on your local drive) then click "Next";



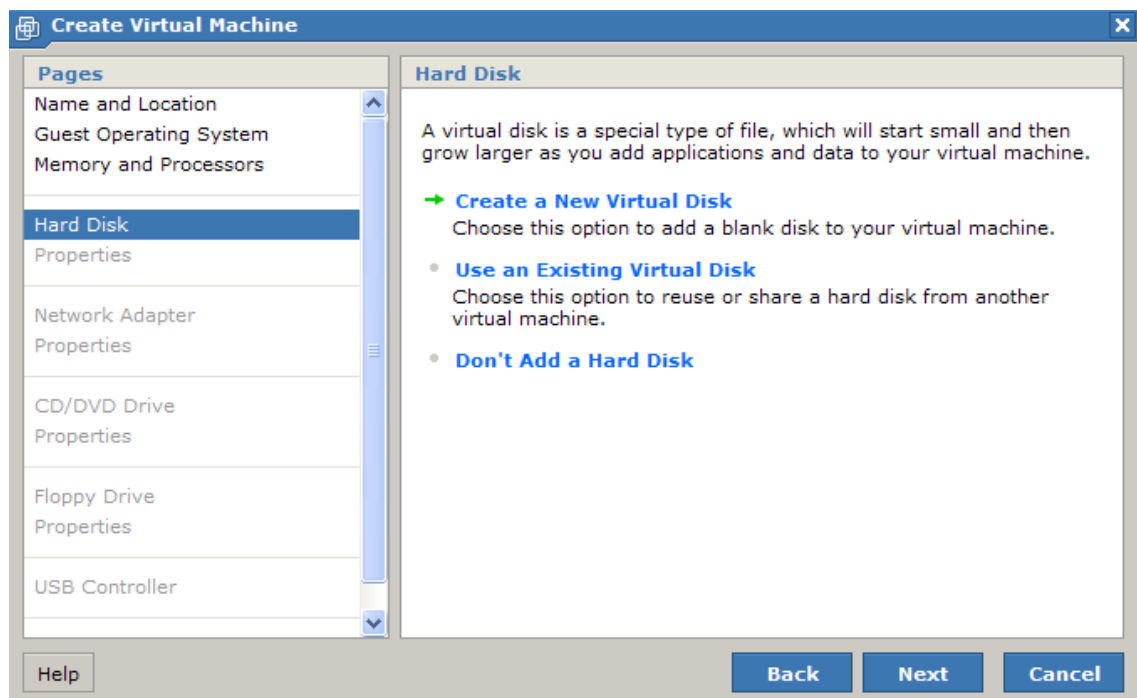
Select the Operating System type (Enterprise Edition for clustering) and click "Next";



Select the RAM (256Mb for DC and 512MB for each cluster node) and CPU (1 each) then click “Next”;

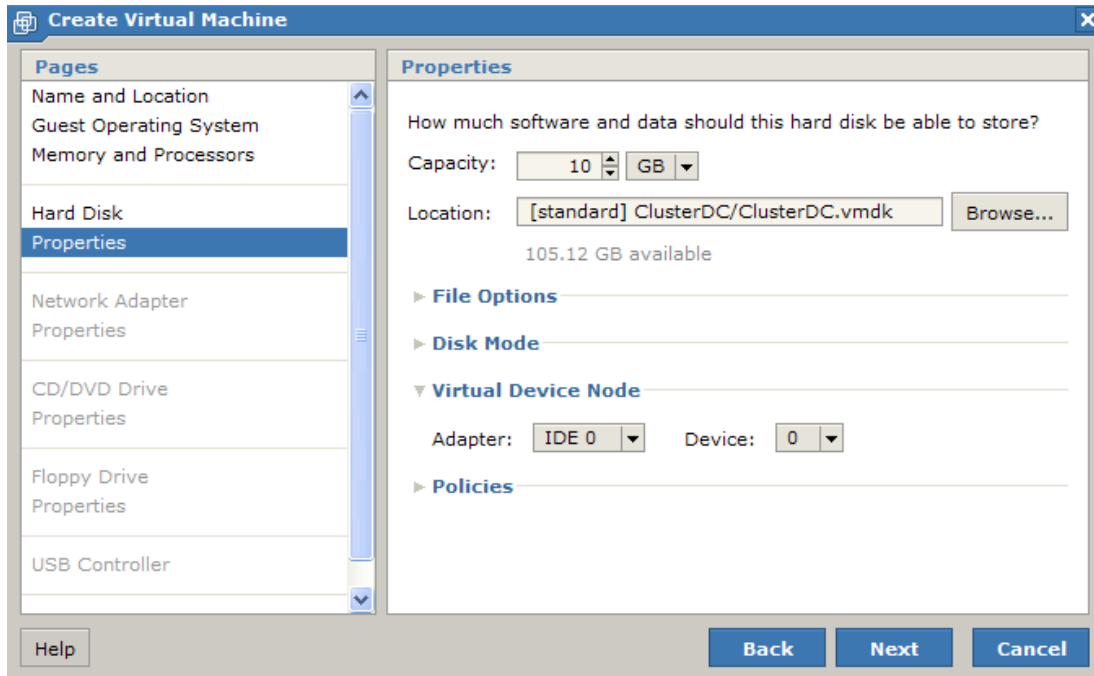


Select to create a new virtual disk;

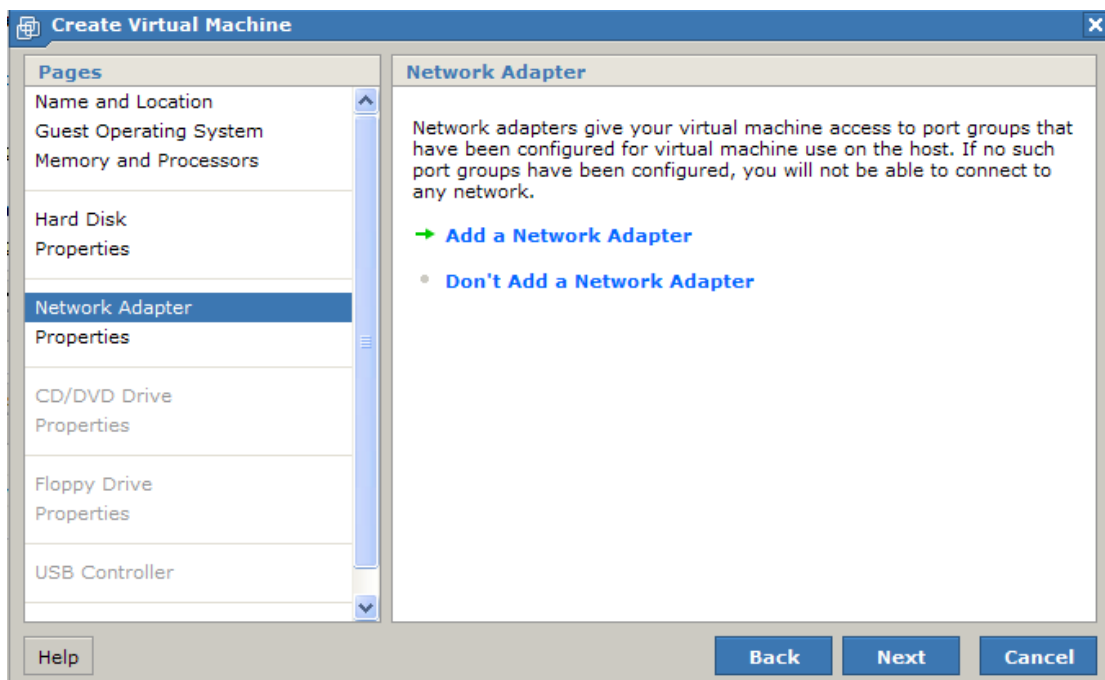


Supply a size for the virtual disk and select SCSI bus (SCSI ID 0 for the VM boot drives in our example), make sure to pre allocate the disk space for the virtual disks. Click “Next”;

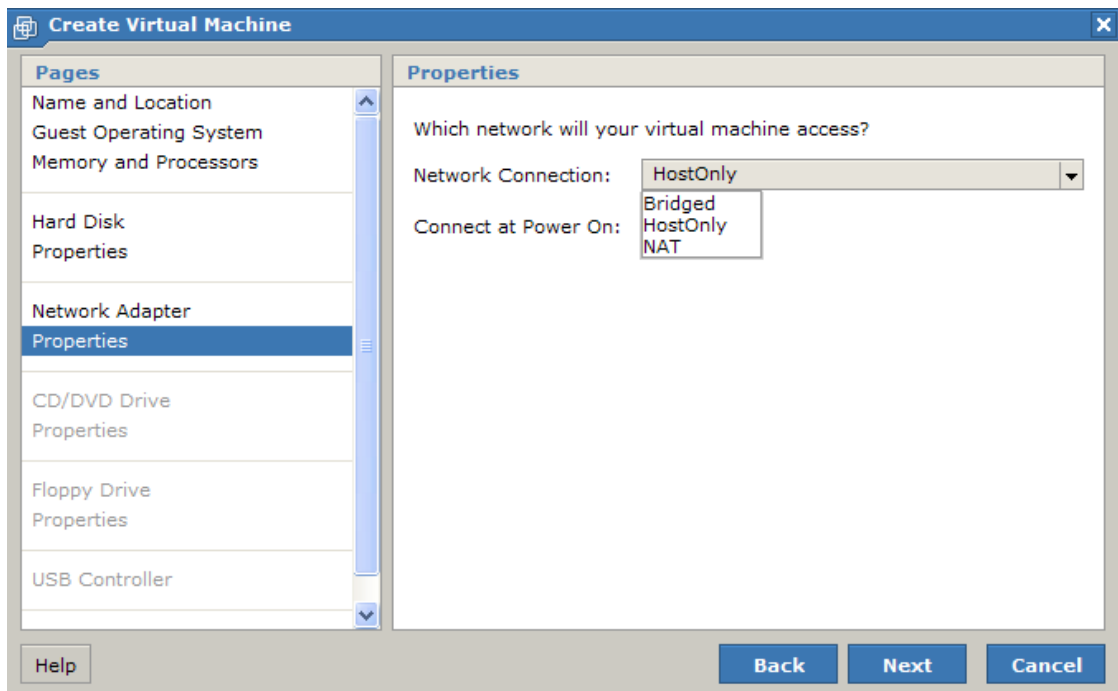
*The Windows 2008 nodes need a 20GB disk drive and 10GB for the DC.*



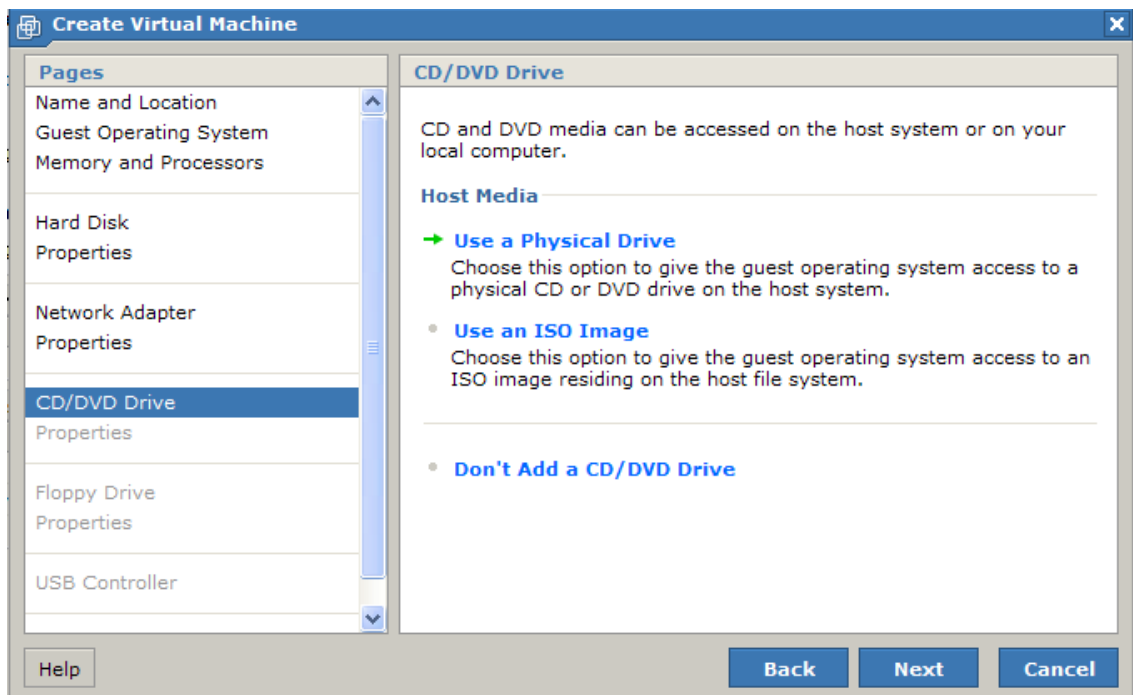
Select “Add a network adapter” and the network selection browse opens




Select "Host only" for this vNIC and click "Next";



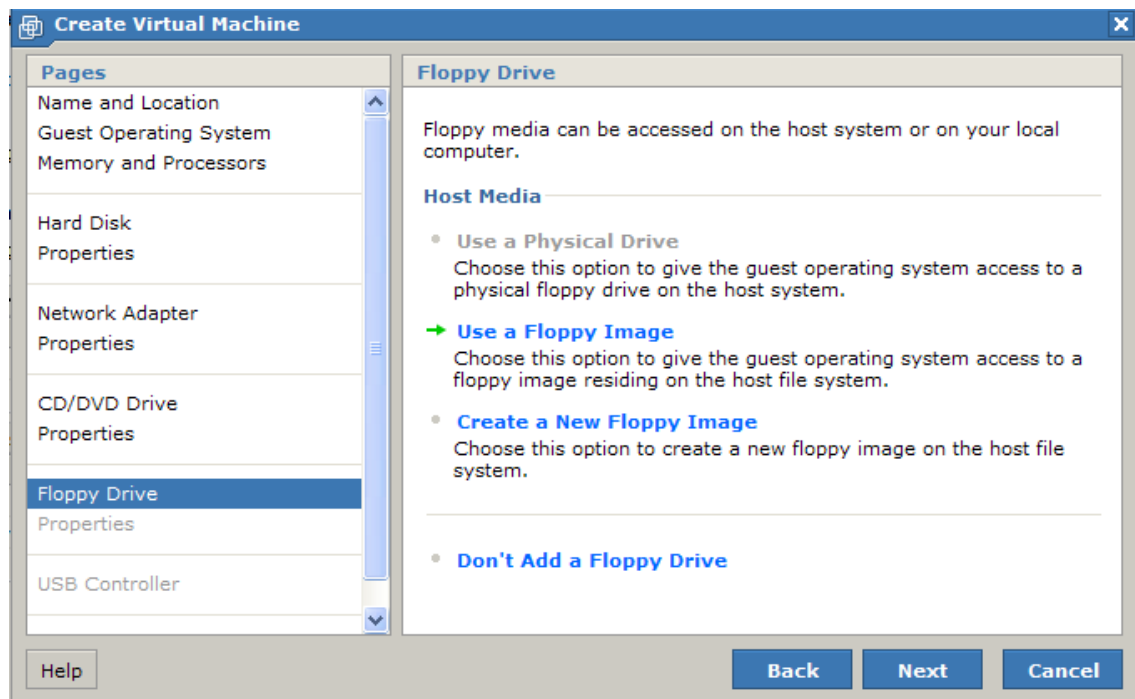
Select to use an ISO image for the vCD drive and click the browse button,



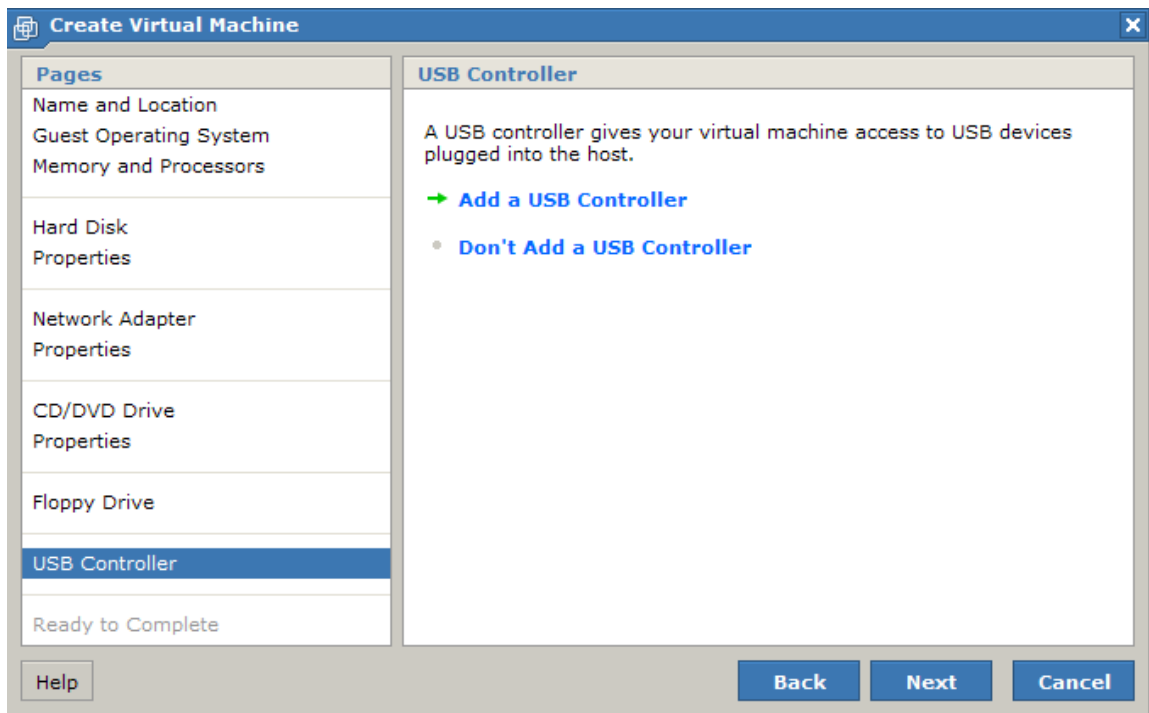
The datastore browser opens, drill down and select the Windows 2003 R2 Enterprise ISO (or Windows 2008 for a cluster node) and click “OK”. At the “create virtual machine window” click “Next”;



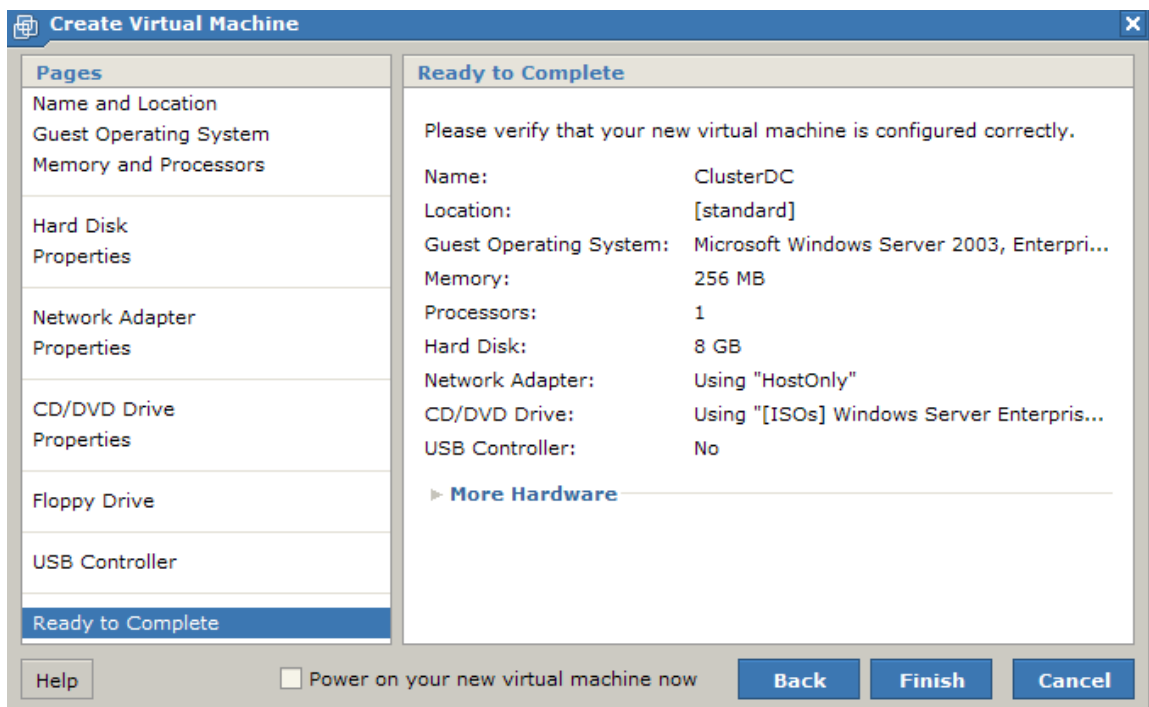
At the next screen do not add a floppy drive and click “Next”;



Do not add a USB controller and click "Next";



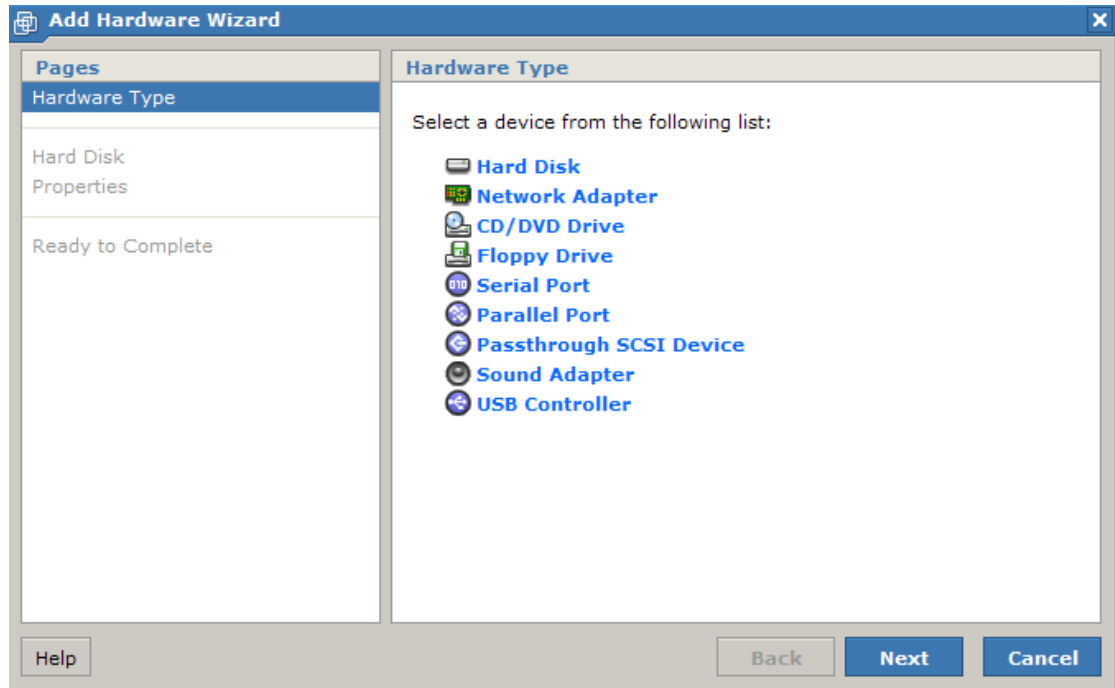
At the last screen click "Finish" to complete the VM



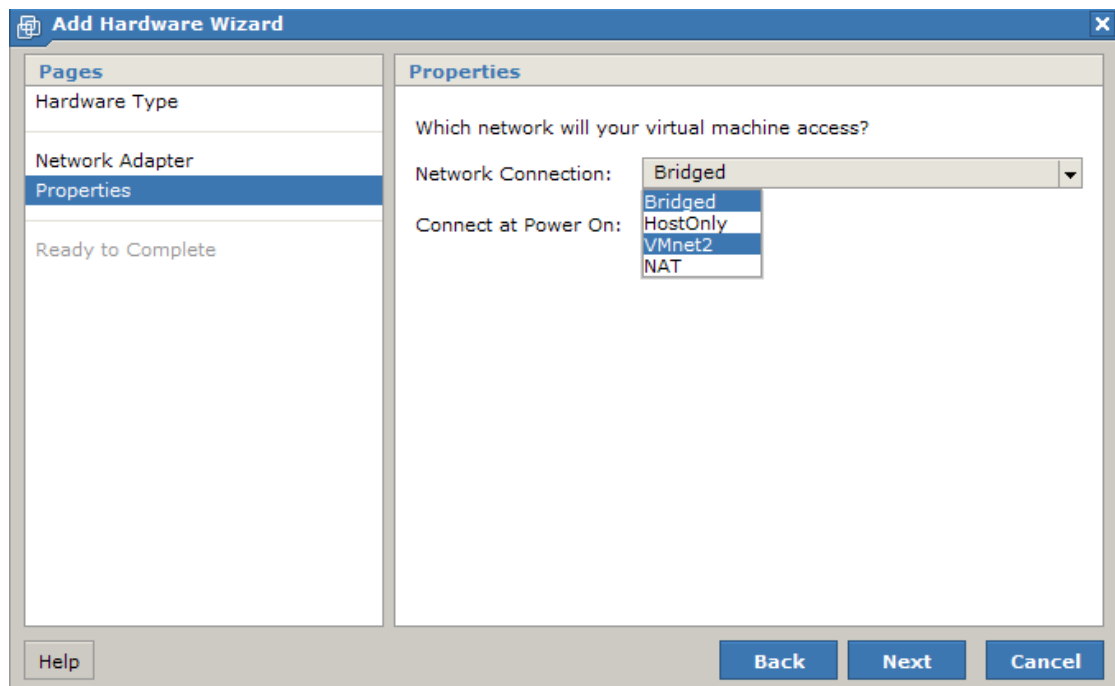


Create the remaining VMs (node 1 and node 2) using Windows 2008 Enterprise software. Use a pre allocated 20GB virtual disk for each node, also add 2 more vNICs, these will be used for the Heartbeat and iSCSI networks. This is done as follows;

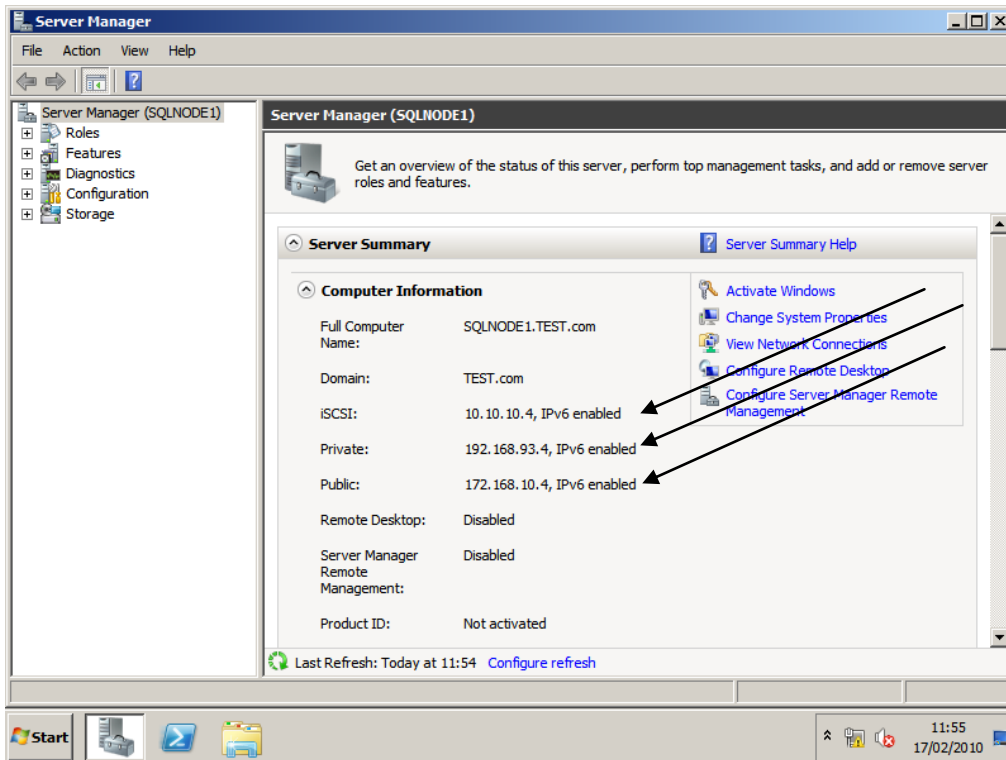
With the base VM created, select the first node and under the “Command” section click “Add hardware”, the following screen appears. Click “Network adapter” and the vNIC properties appear;



Select the “VMnet2” option from the drop down list and click “Next”. Add in a 3<sup>rd</sup> vNIC for VMnet3, then click “Finish” to complete. Do this for the second cluster node too.

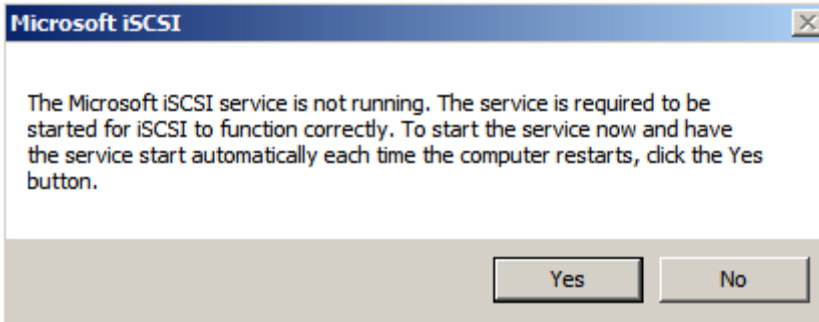


Install the windows operating systems on each VM and create a domain controller with a test domain on ClusterDC. Configure Public networking between the 3 VM's and join the 2 nodes to the domain. Configure the iSCSI network between the VM's and the NAS VM. Configure the heartbeat network between the cluster node VMs. Once this is done you may proceed with the rest of the instructions in this document. If you are unsure with any aspect of the NOS build or configuration consult your Windows administrator for help with this. The screenshot below shows the deployed Windows 2008 server with network adapters.

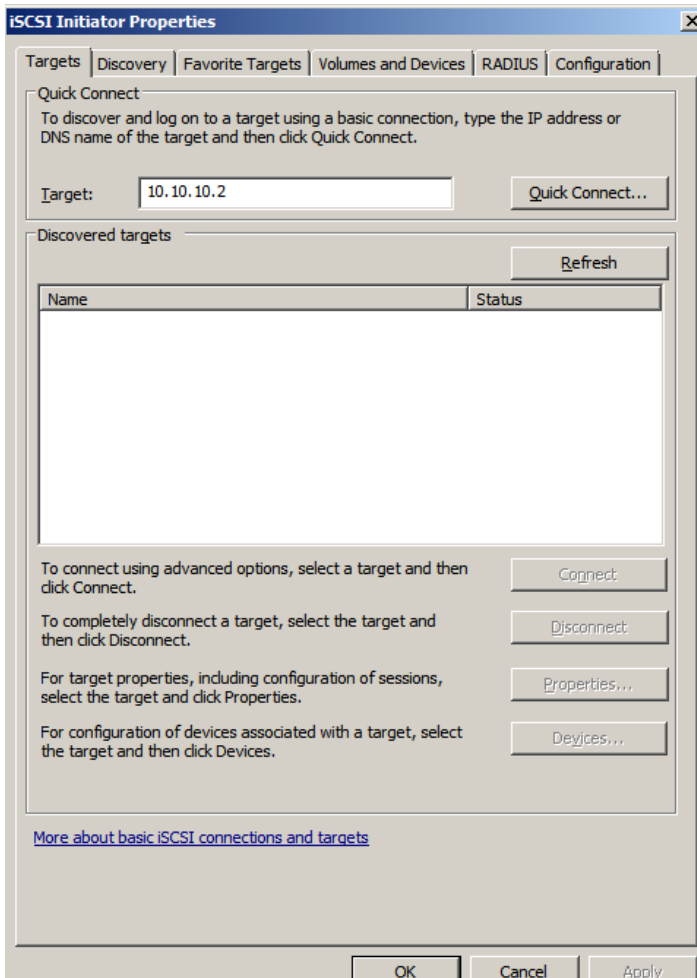


### 3.4 ATTACHING THE ISCSI LUNS

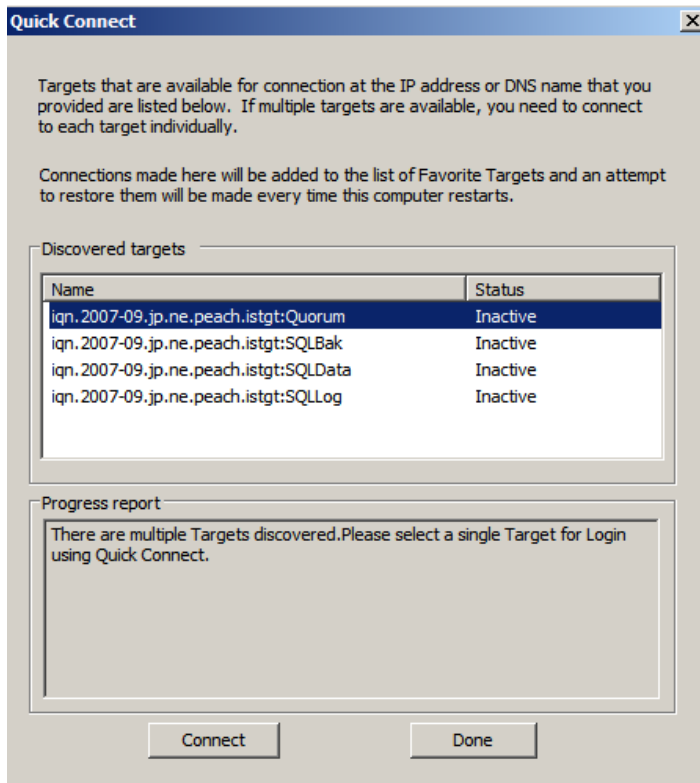
Once the cluster nodes have been created and networking configured we need to start the Windows iSCSI initiator and discover the iSCSI LUNs. Open administrative tools and double click the iSCSI Initiator. You may receive a message indicating the service is not running and needs to be started, this is shown below. Accept this message to continue.



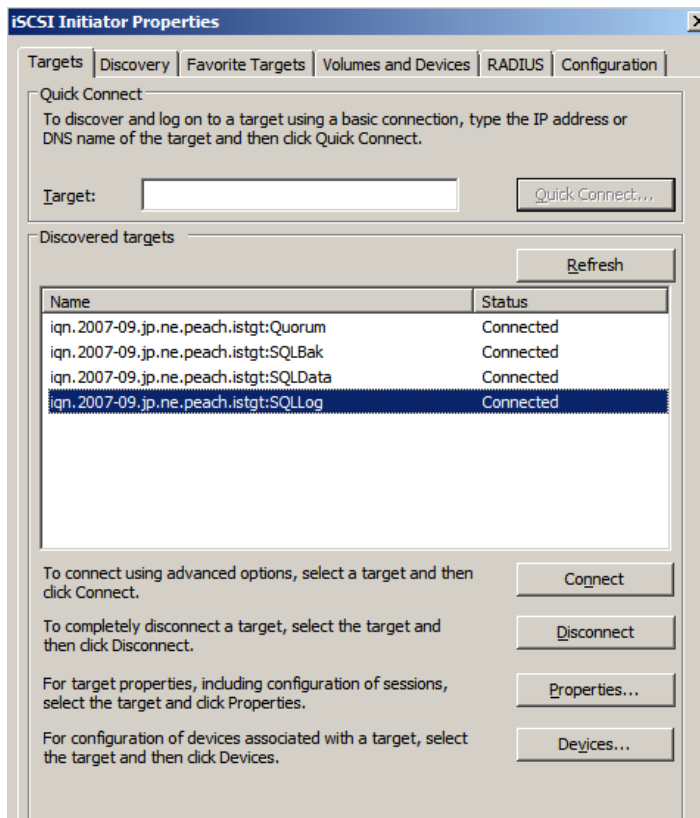
With the service now started the following dialog should appear. Enter the NAS VM IP address and click "Quick connect". *You may even use IPsec for secure communications. Most importantly your iSCSI traffic should pass over a private, segregated network (much like the VMWare Vmotion network).*



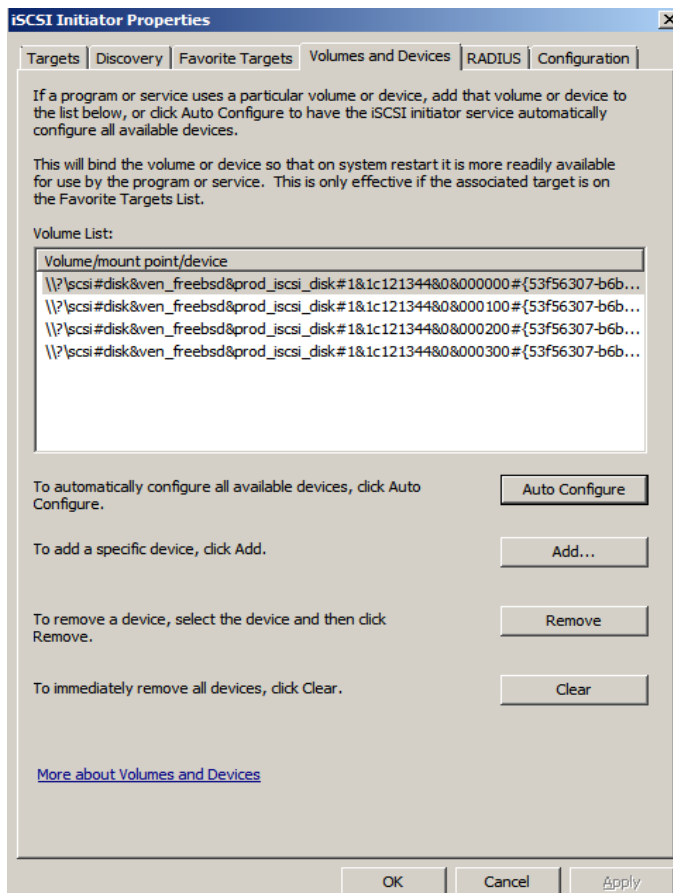
The quick connect dialog opens as shown below. Click each target and then click connect. When all targets are connected, click "Done".



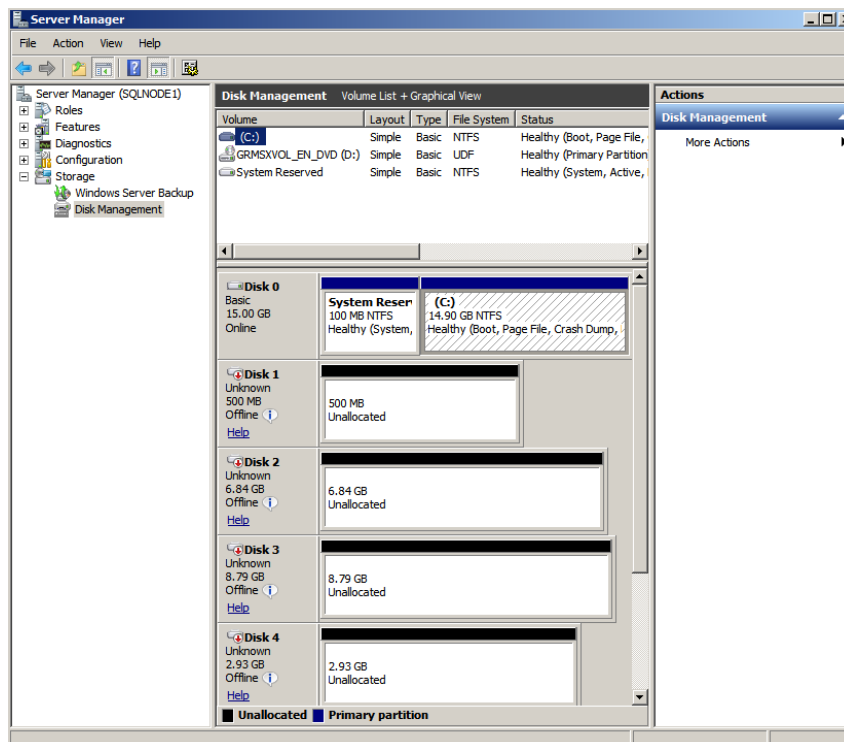
All targets connected!



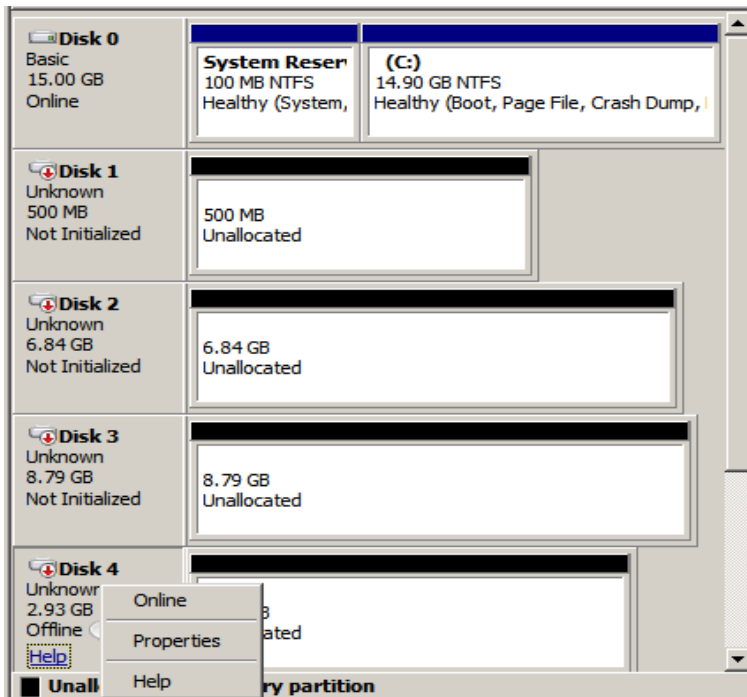
Go to the “Volumes and devices” tab and click “Auto configure” then click “OK”



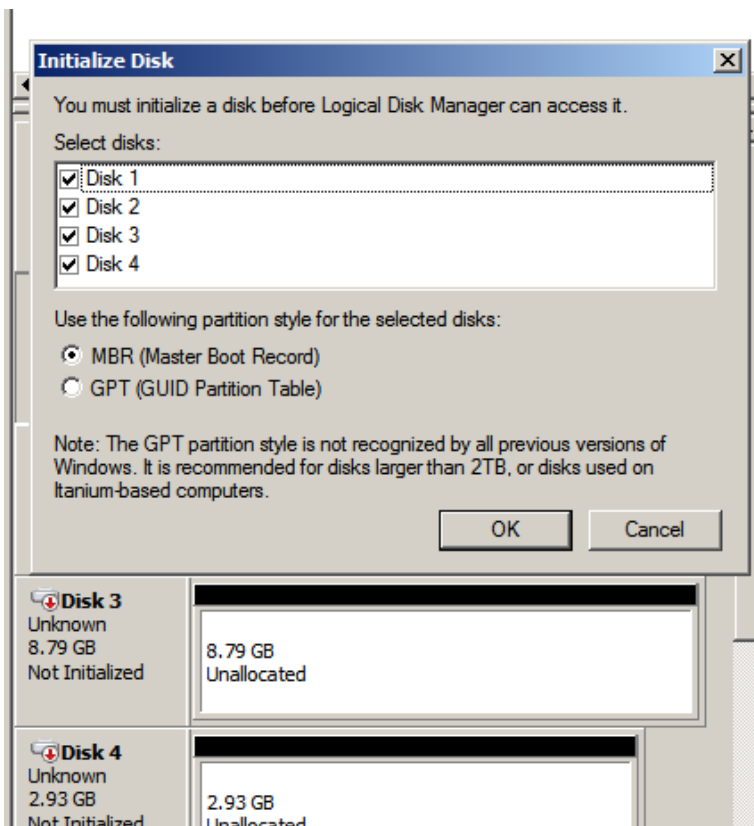
If you now open disk management you will see the following, notice the disks are offline and unknown!



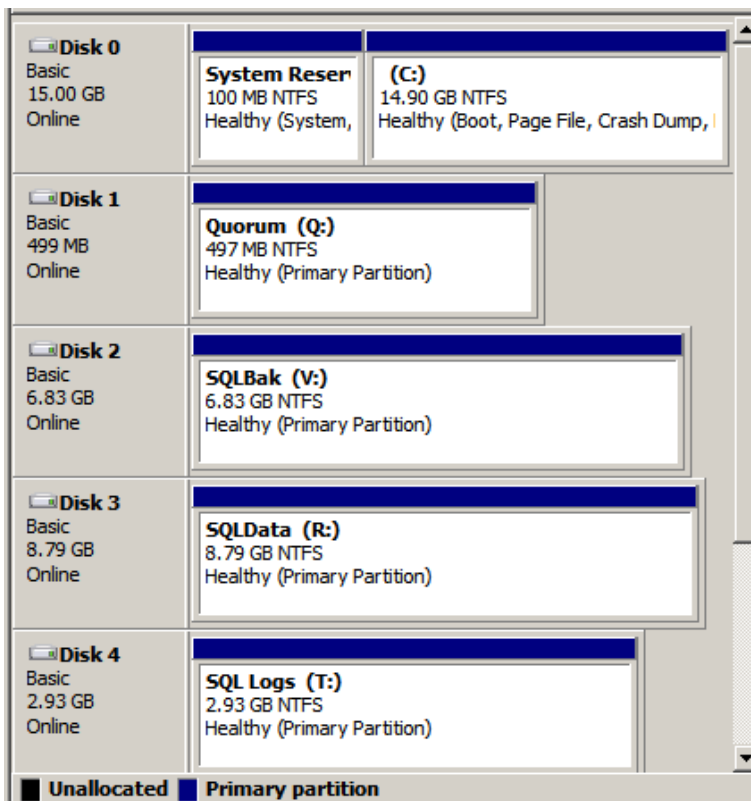
Right click each disk and select "Online". Once all disks are online, right click any disk again and select "Initialise". The wizard will initialise all drives found online.



And Initialise

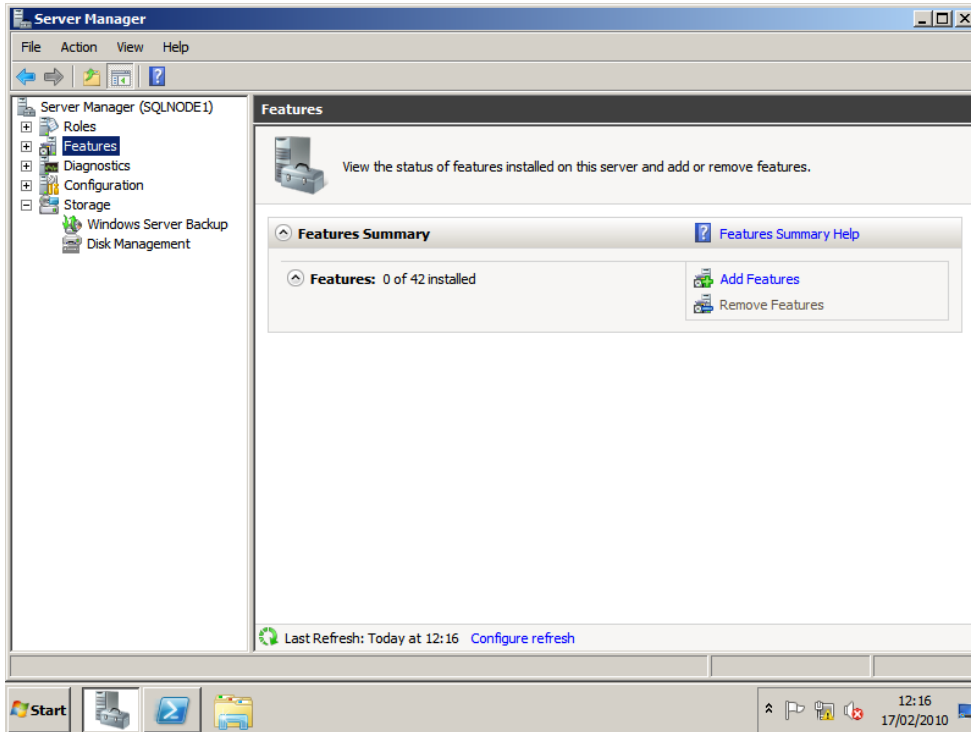


Create your disk partitions in the usual manner. All disks online and formatted!

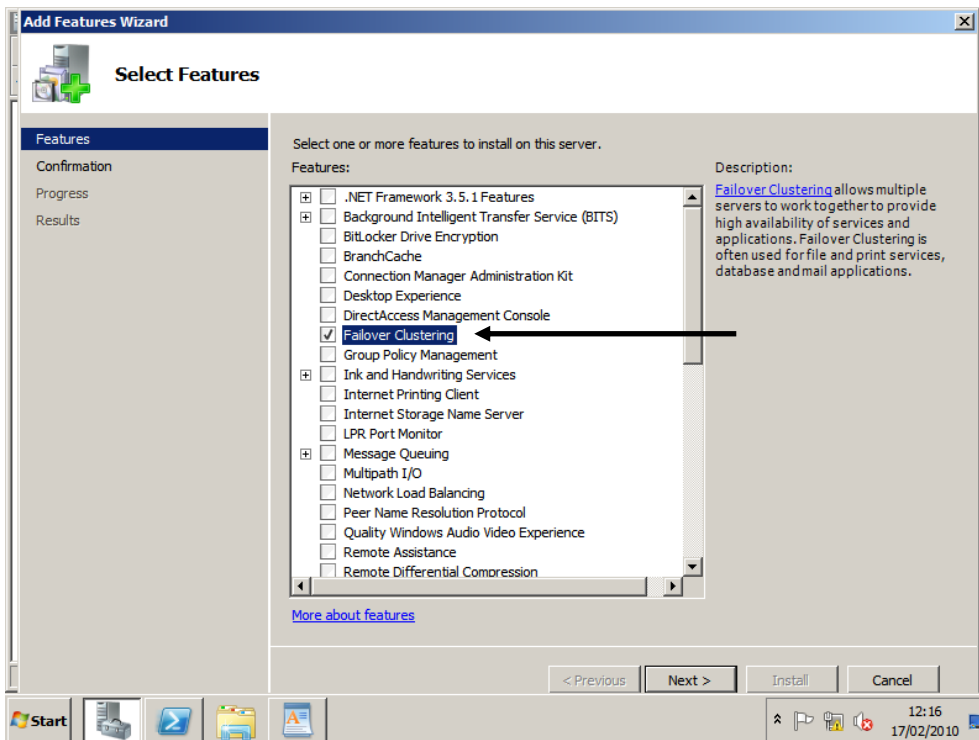


### 3.5 INSTALLING THE WINDOWS 2008 CLUSTER

The Windows 2008 cluster may now be verified\created, but first we need to enable this feature in Windows 2008 feature manager. Click “Add features” to continue,

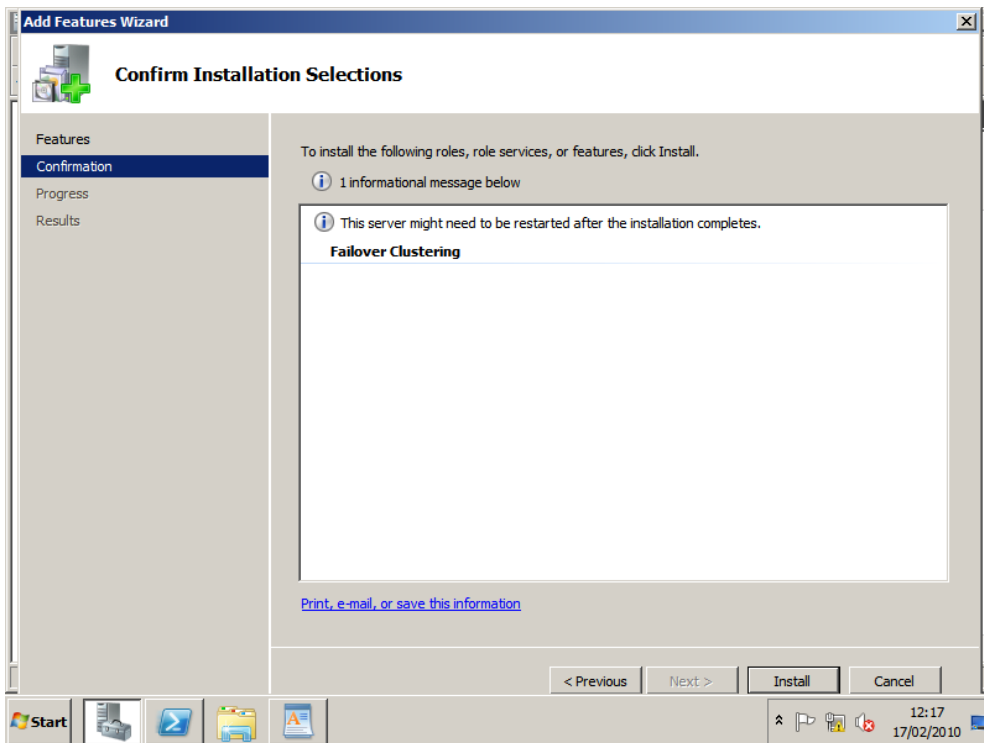


Select the Failover Clustering checkbox and click “Next”.

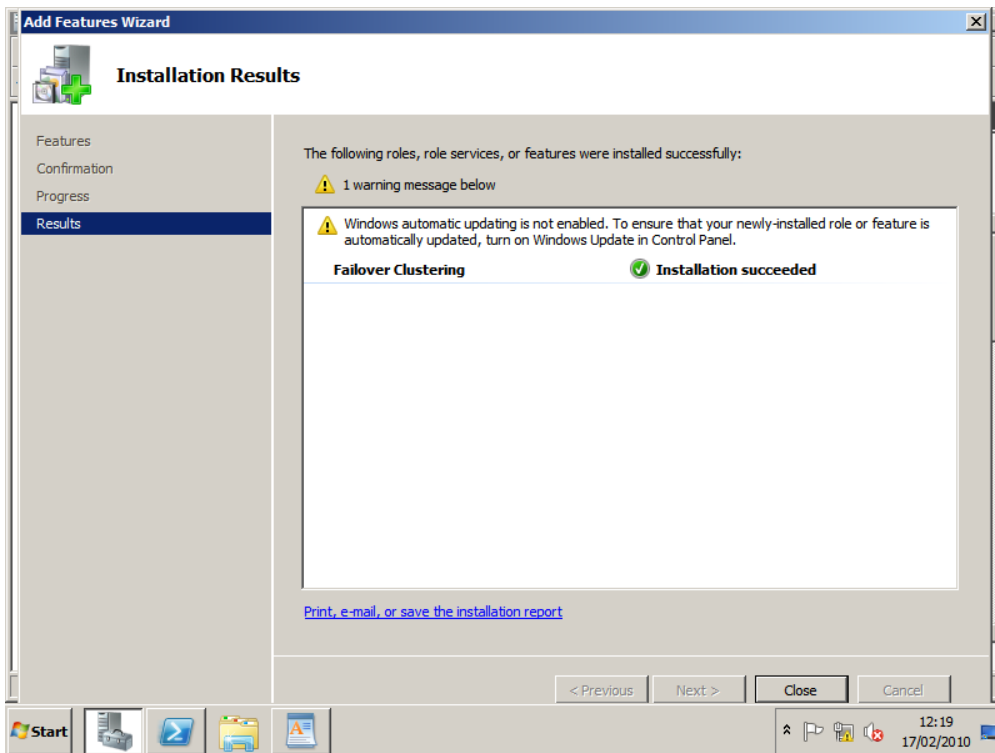




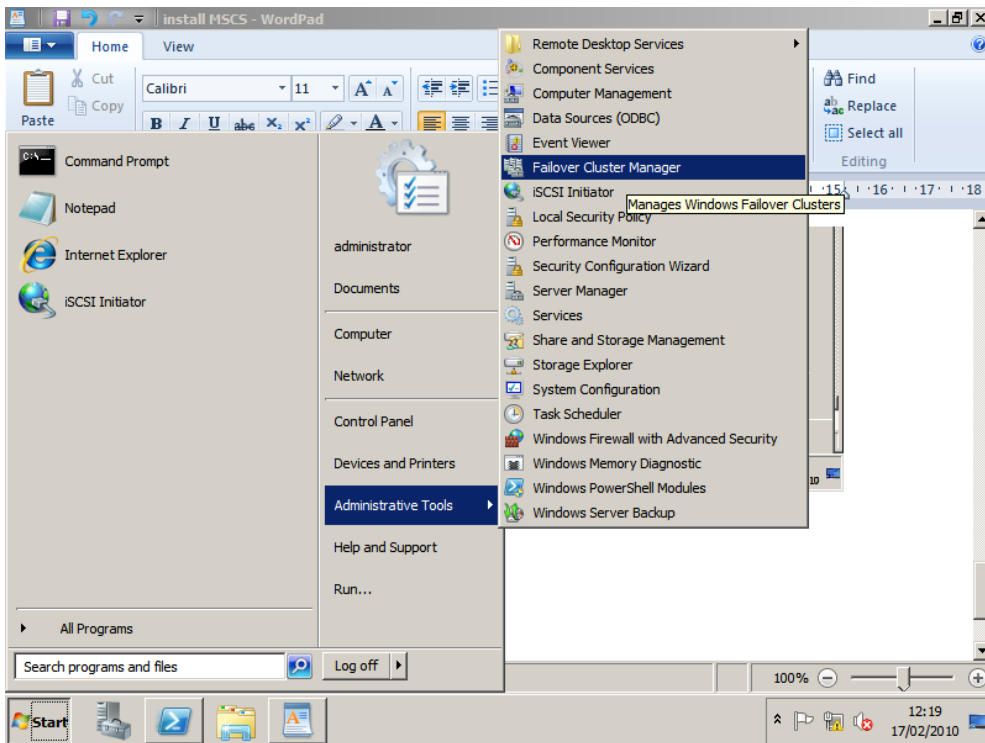
Click "Install" to proceed.



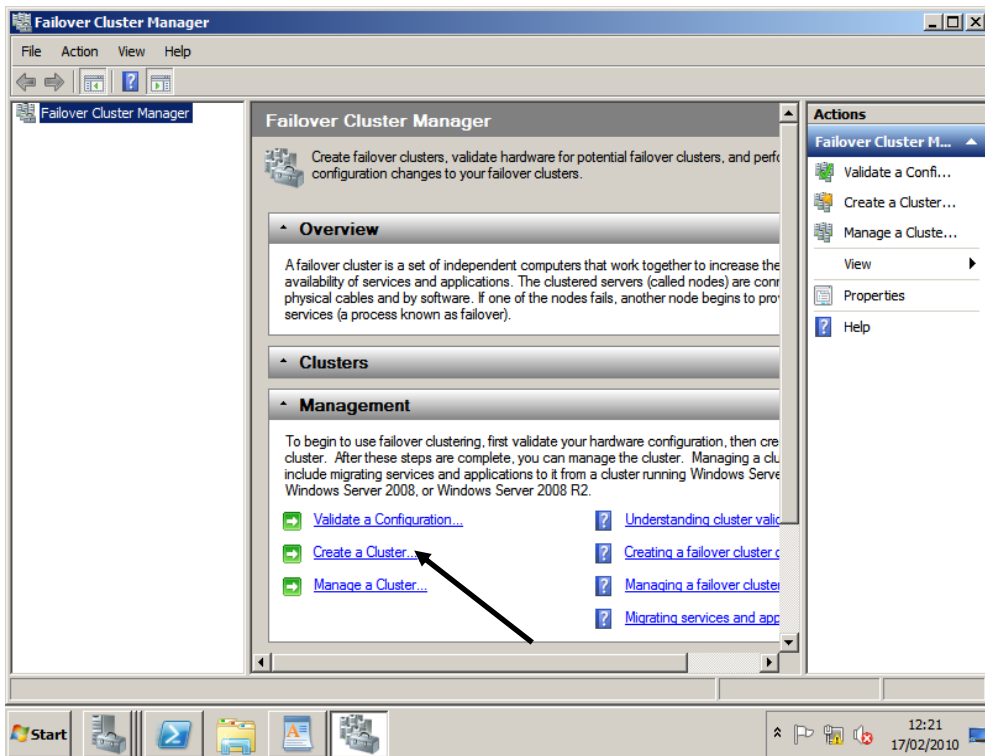
Review any errors or warnings and click "Close" once the installation finishes.



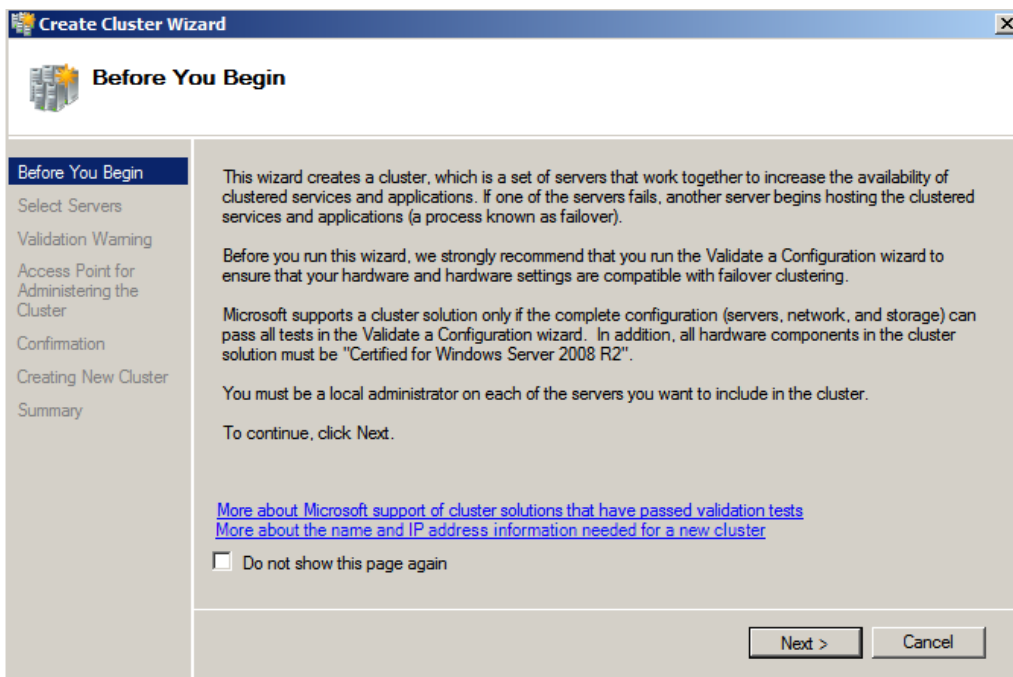
Open Failover Cluster manager from the Administrative Tools menu.



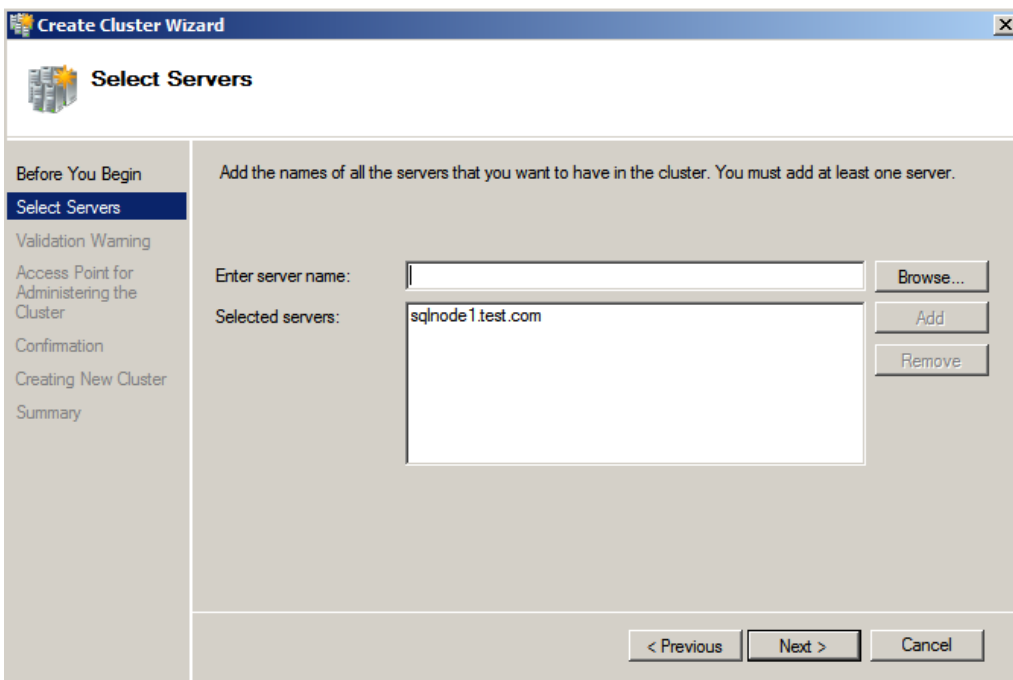
Select the option to create a cluster.



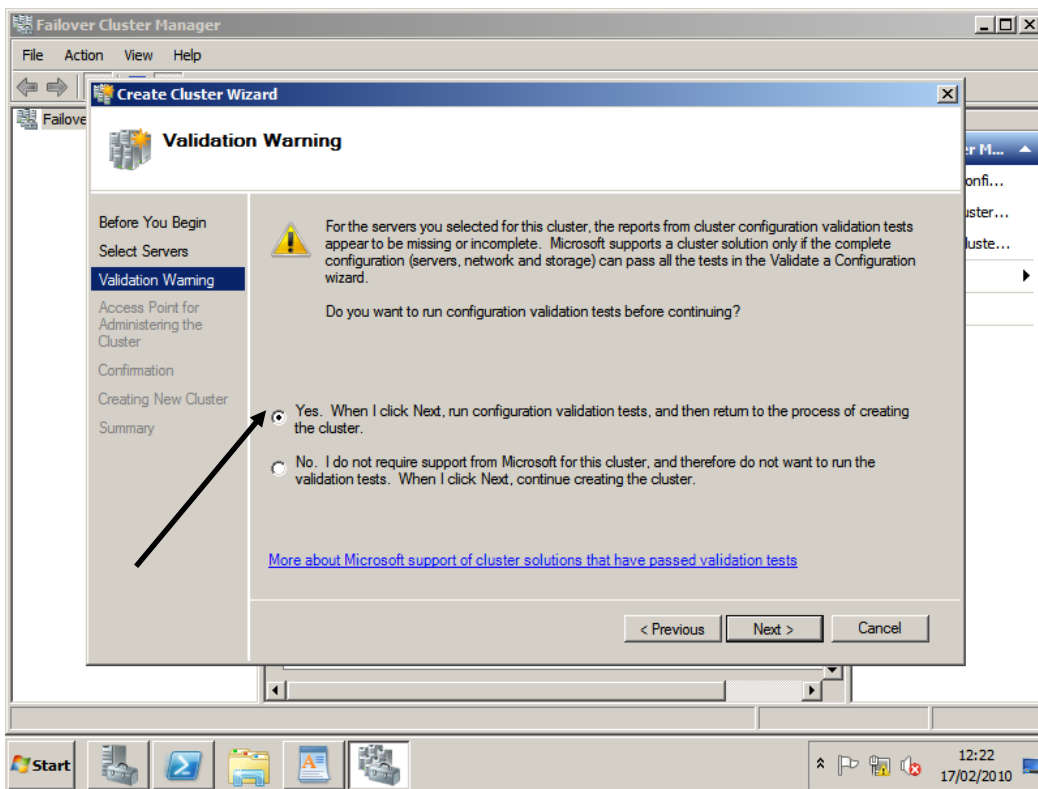
Review the beginning notes and click “Next”.



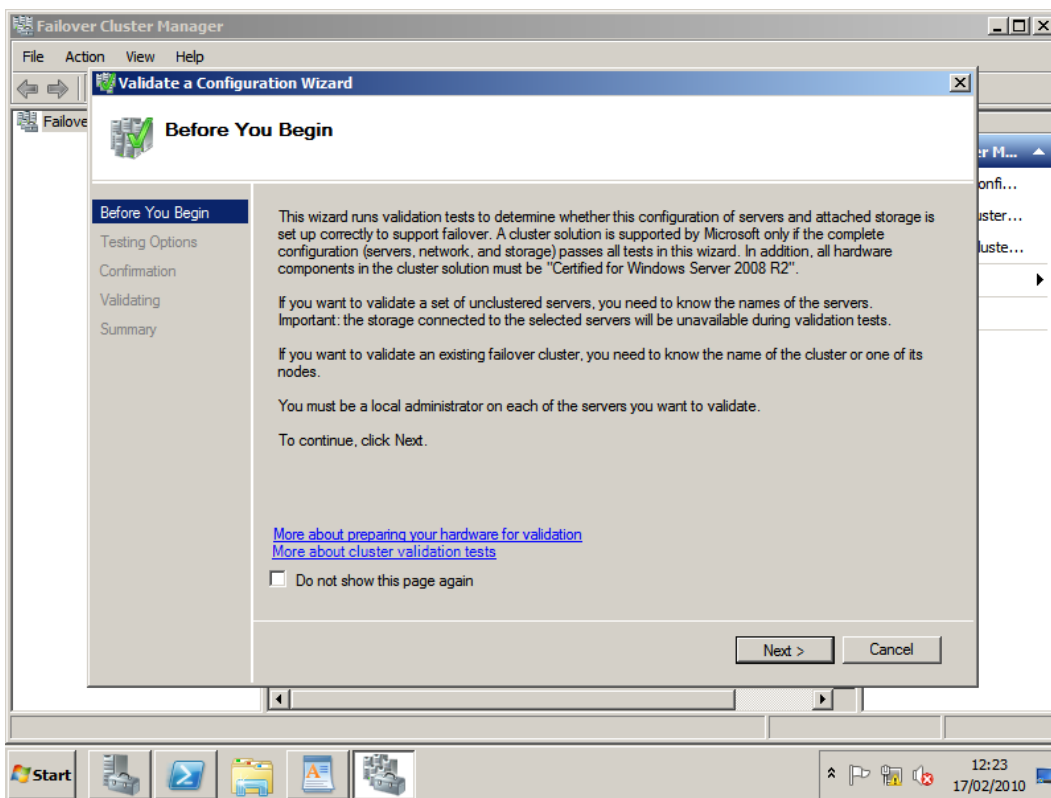
Select the server(s) to include in the cluster



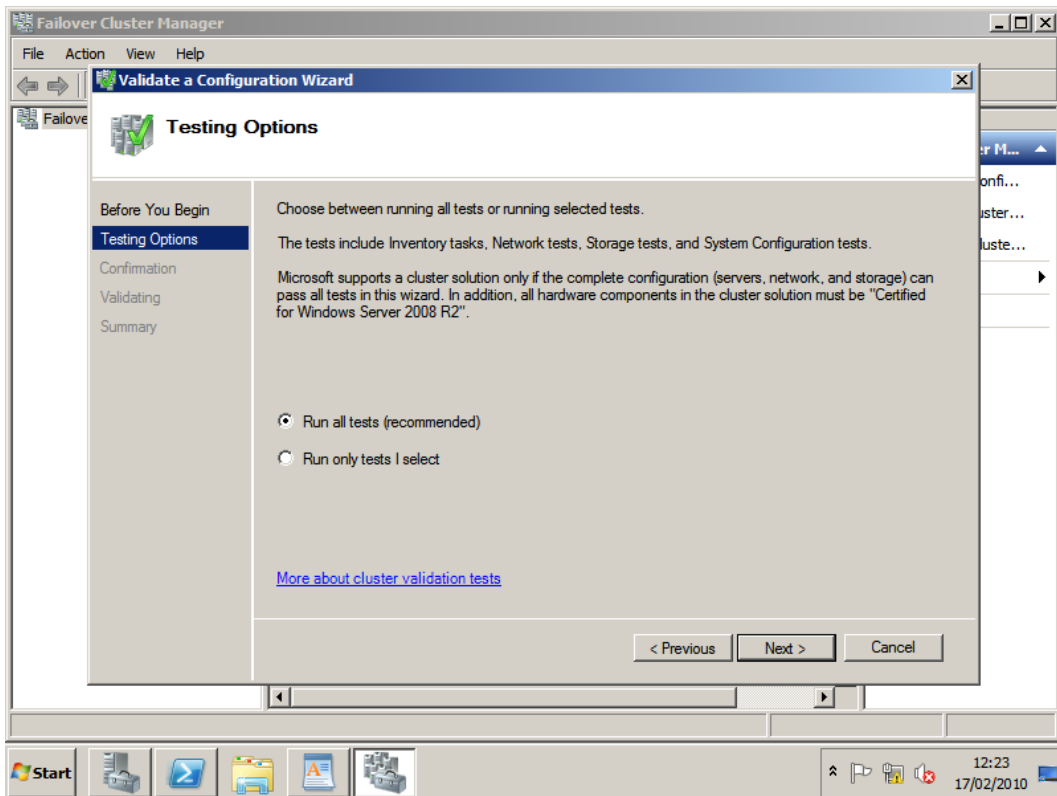
Review the validation warning. Select the radio button shown below and click “Next”.



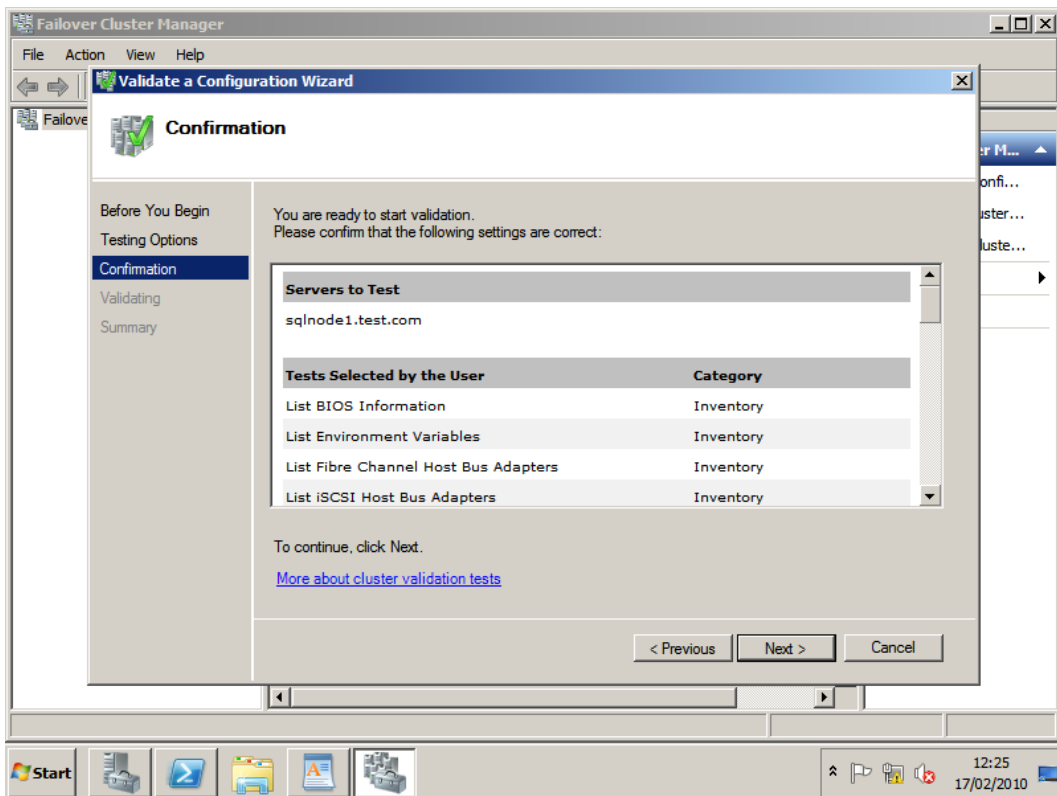
Click “Next” to start the cluster validation wizard. This is required to be eligible for Microsoft support if it is ever needed.



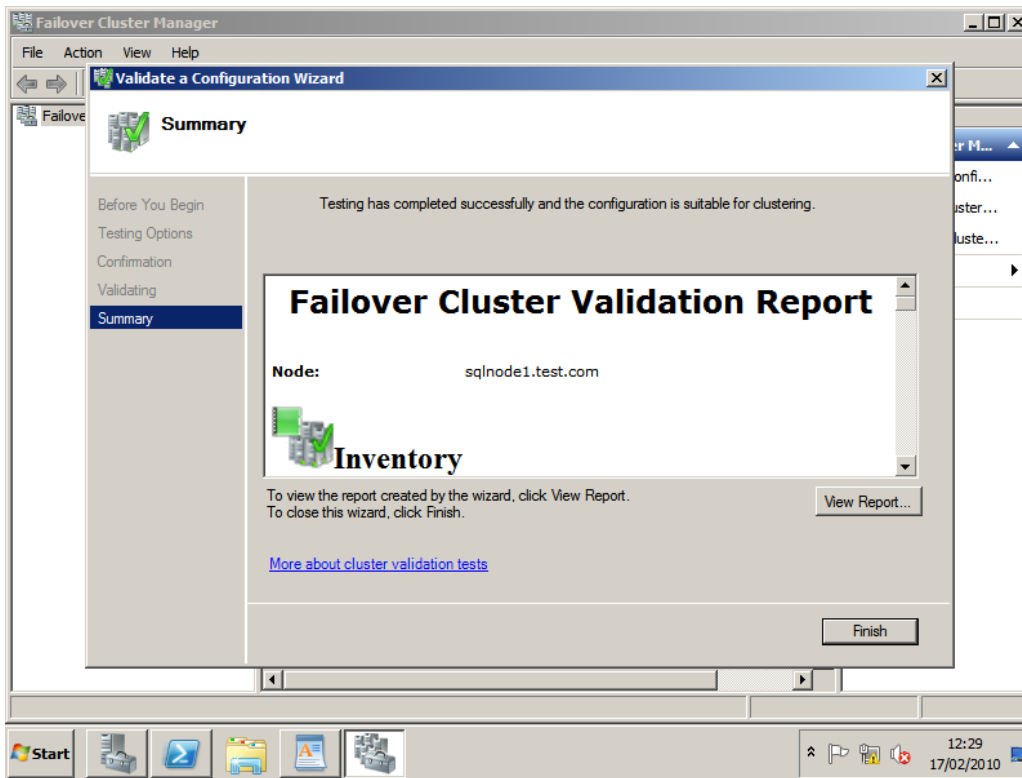
Select the option to “Run all tests”. This will thoroughly test all storage, networks, operating system, etc.



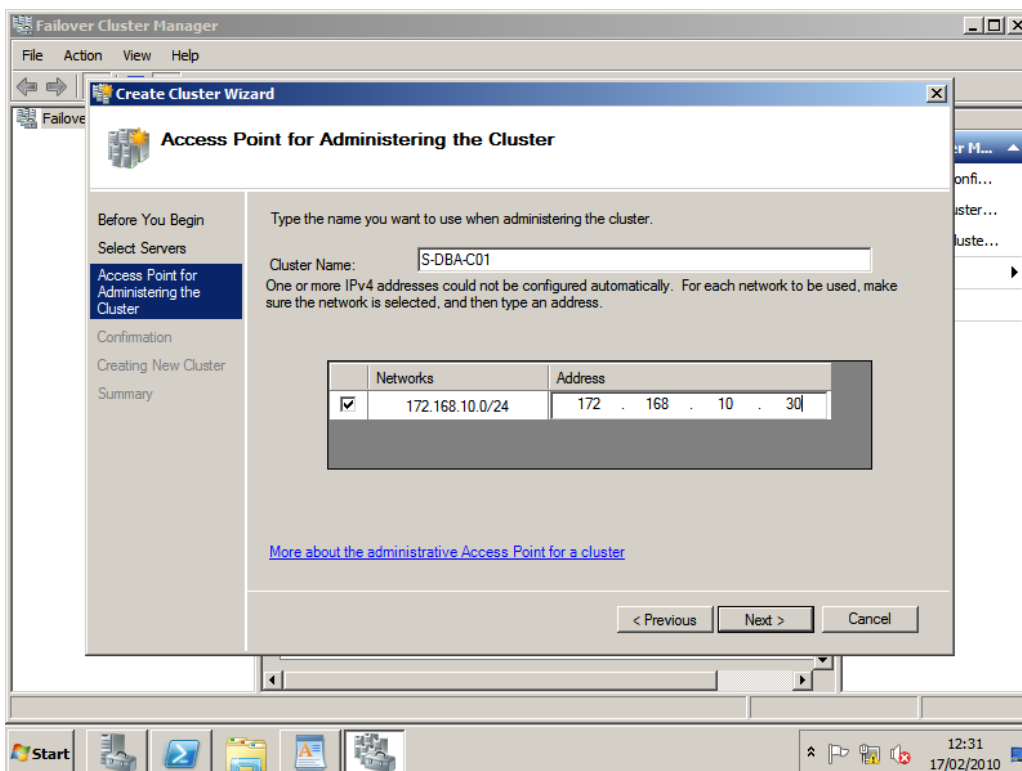
Click “Next” to start validation



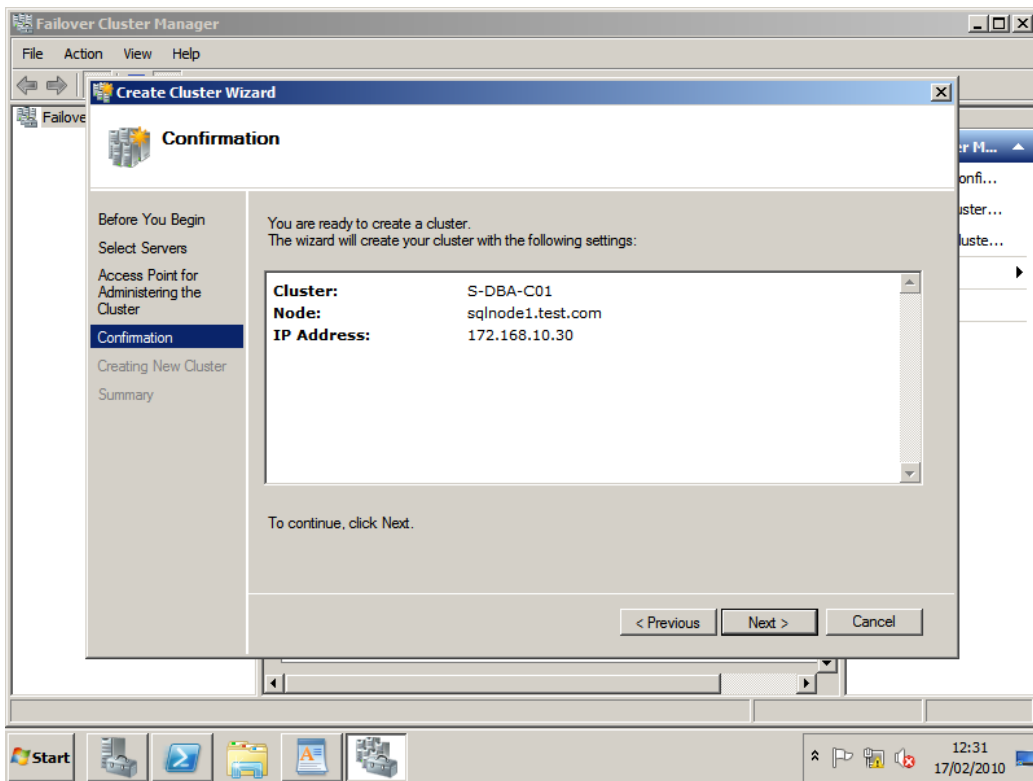
Once validation has completed you may review the report. When you have finished click “Finish”.



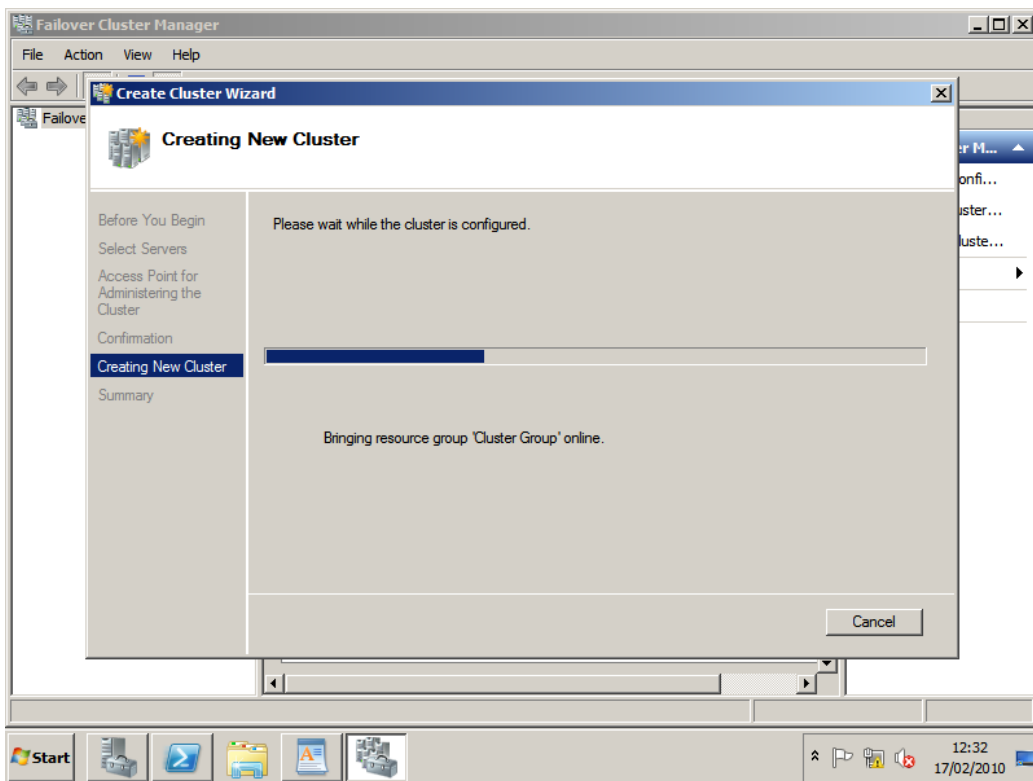
If you have any failures you must remediate these before you may continue. The cluster installation will now continue if the tests were successful. Supply a unique cluster name and IP address and then click “Next”.



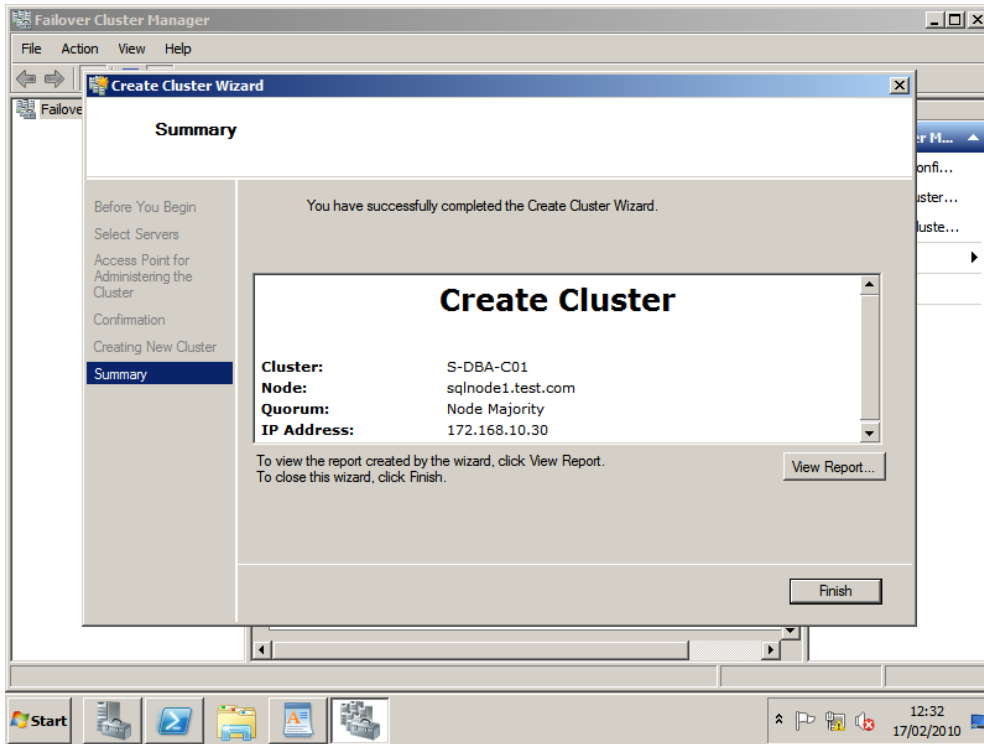
Click "Next" to confirm the details.



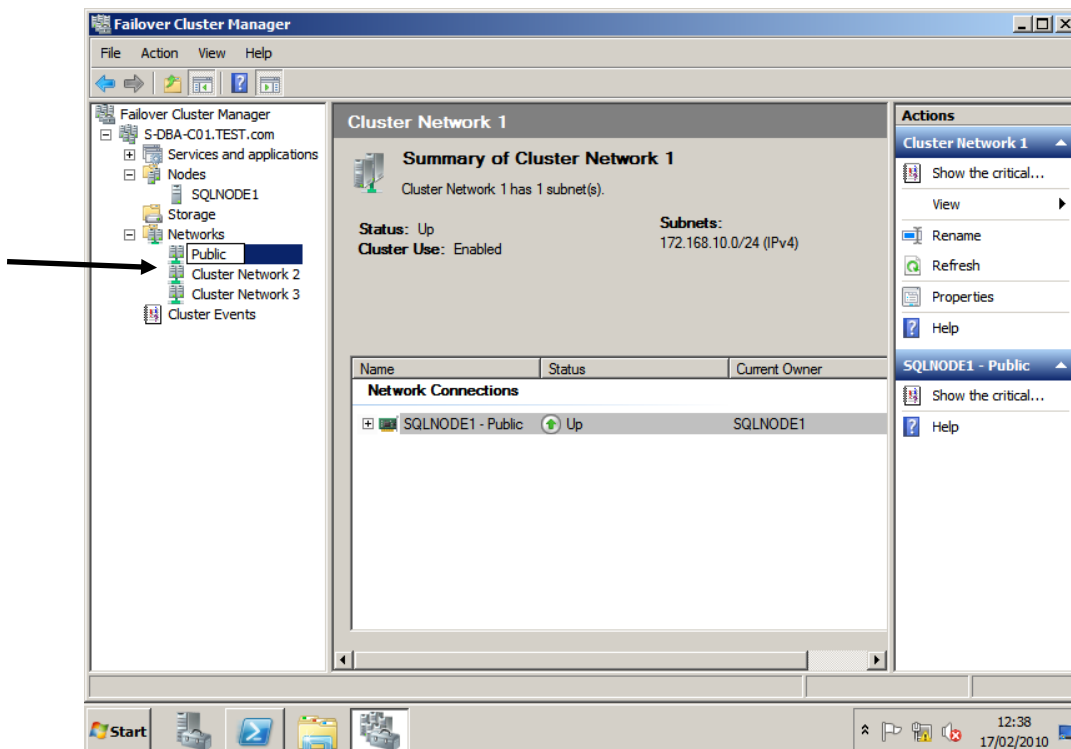
The cluster configuration starts.



Once the wizard completes you may review the report and finally click “Finish”.  
*Note the Quorum type used by Windows 2008 MSCS. We will change this later for our 2 node cluster.*

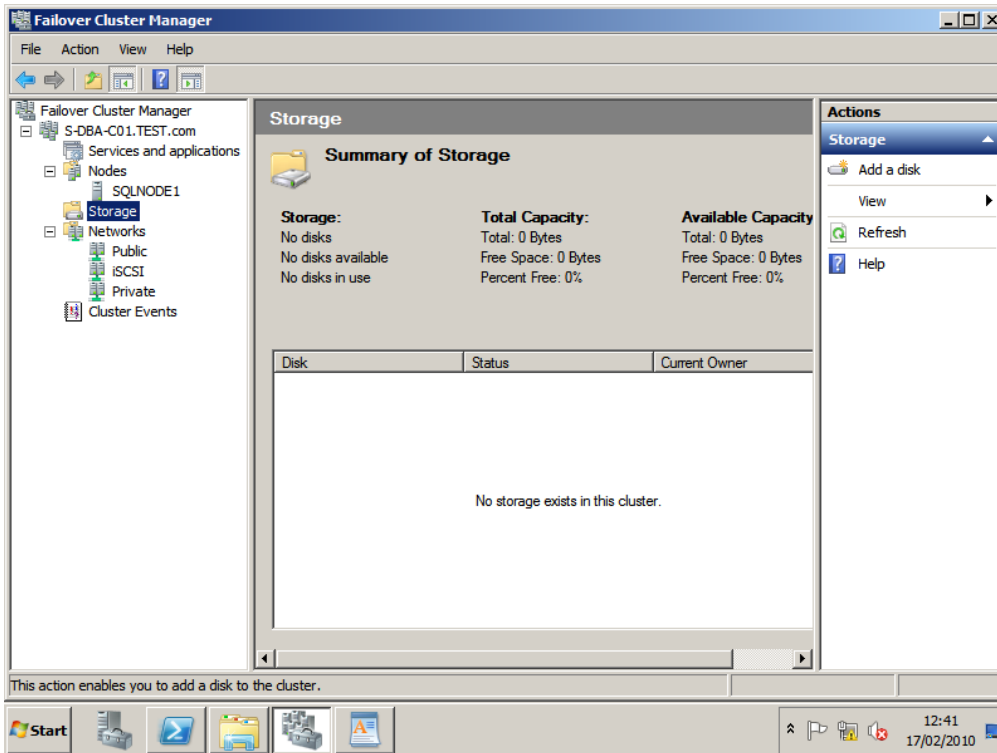


Change the networks names to match their intended use. Check the properties of each and set the network usage. The Public should allow cluster and public communications. The Private should allow cluster communication only and the iSCSI should allow no cluster or public communications.






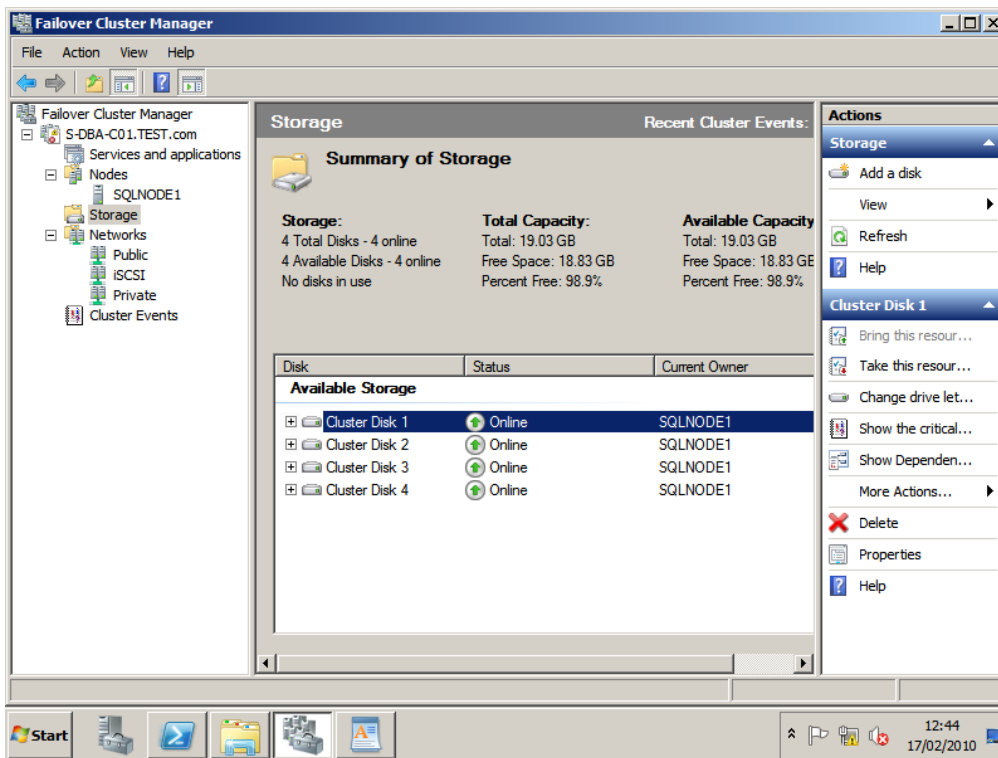
We have no Storage so add the disks now by clicking “Add a disk”.



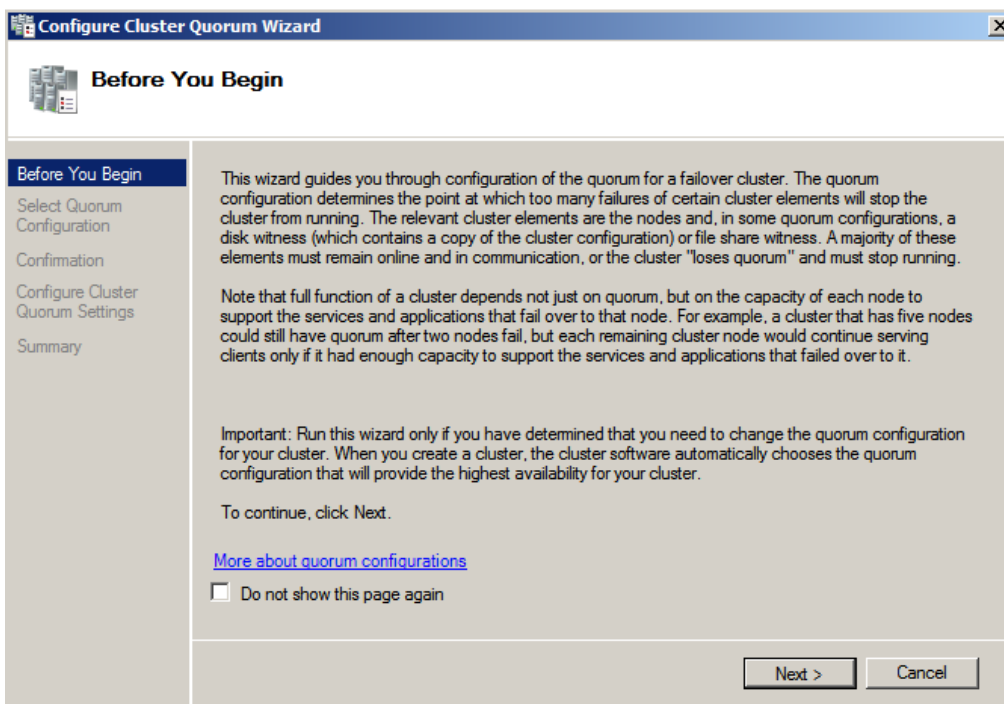
Our disks appear in the selection screen, ensure they are all checked and click “OK”.



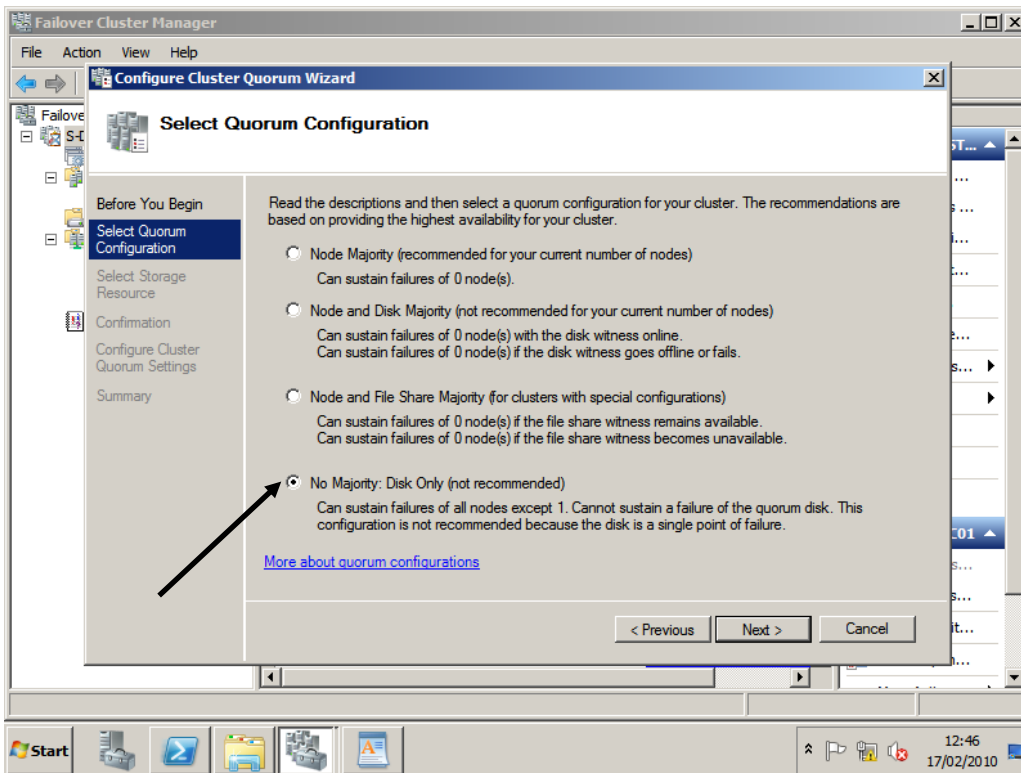
Now that the disks are added to the cluster storage we can start to re configure the Quorum settings to use a disk instead of a majority node set share. Select/highlight the cluster in Failover Cluster Manager then select “Action” > “More Actions” > “Configure Cluster Quorum Settings”.



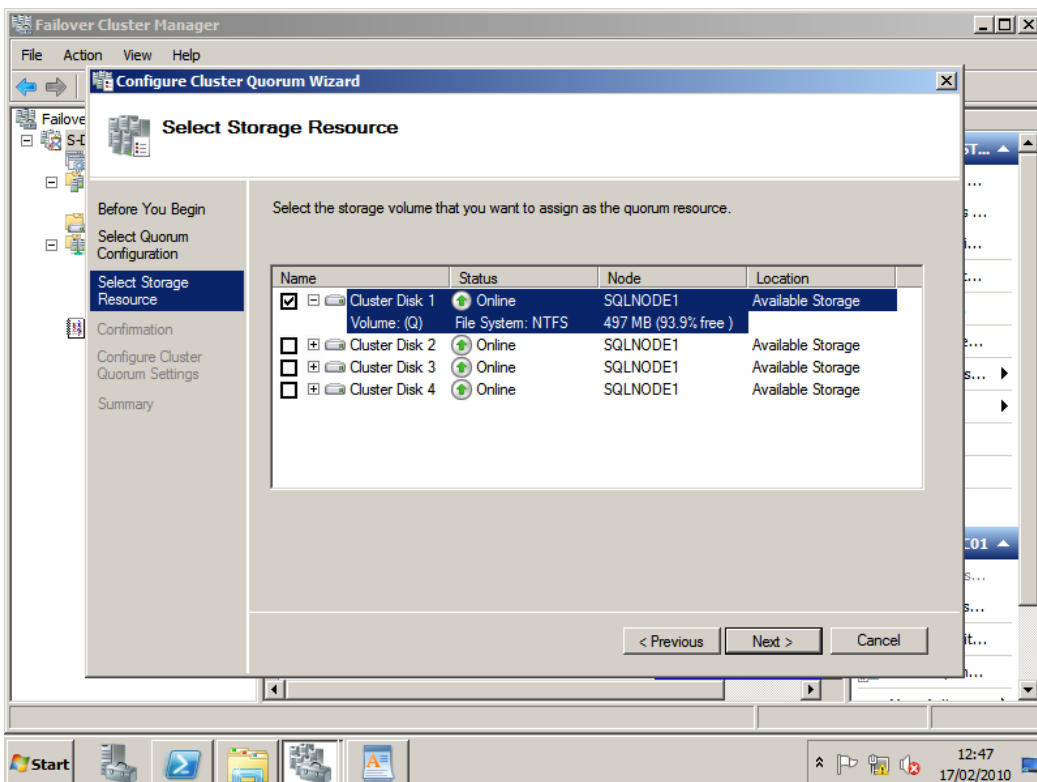
The wizard starts as shown below, review the screen and click “Next” to continue.



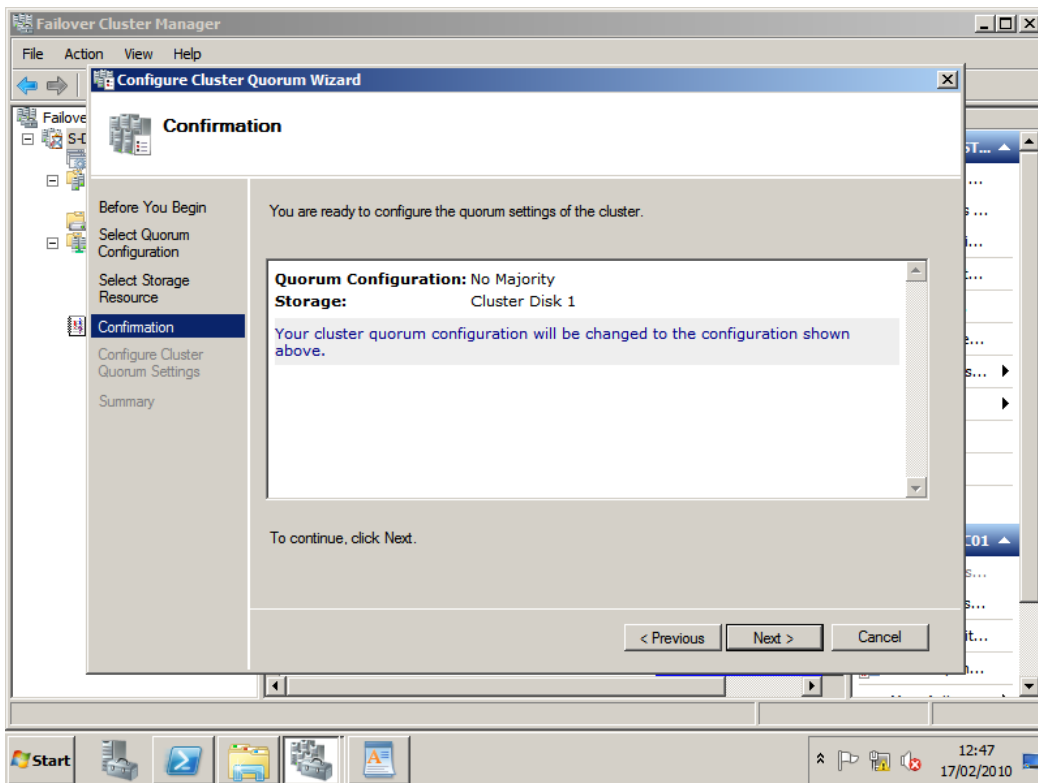
Select the option below for our virtual 2 node cluster and click “Next”.  
 For clusters of 3 nodes or more you should use the majority node set. For more information on majority node sets check the online *Technet* documentation.



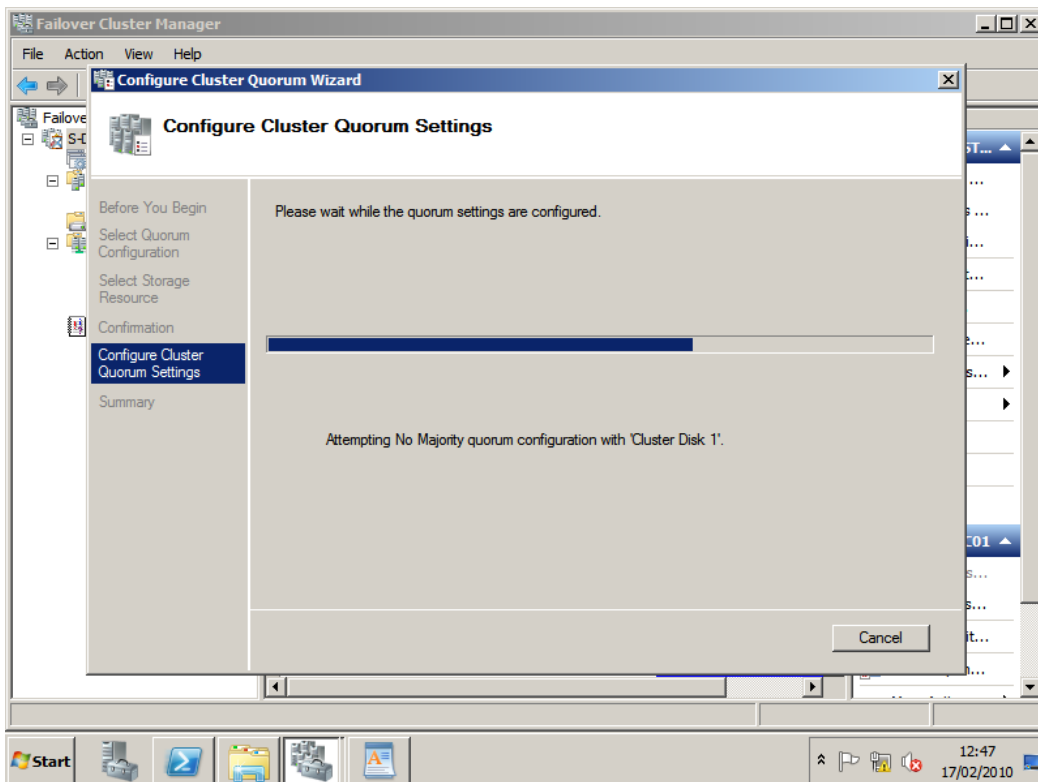
Select the storage resource to use for the Quorum drive and click “Next”.



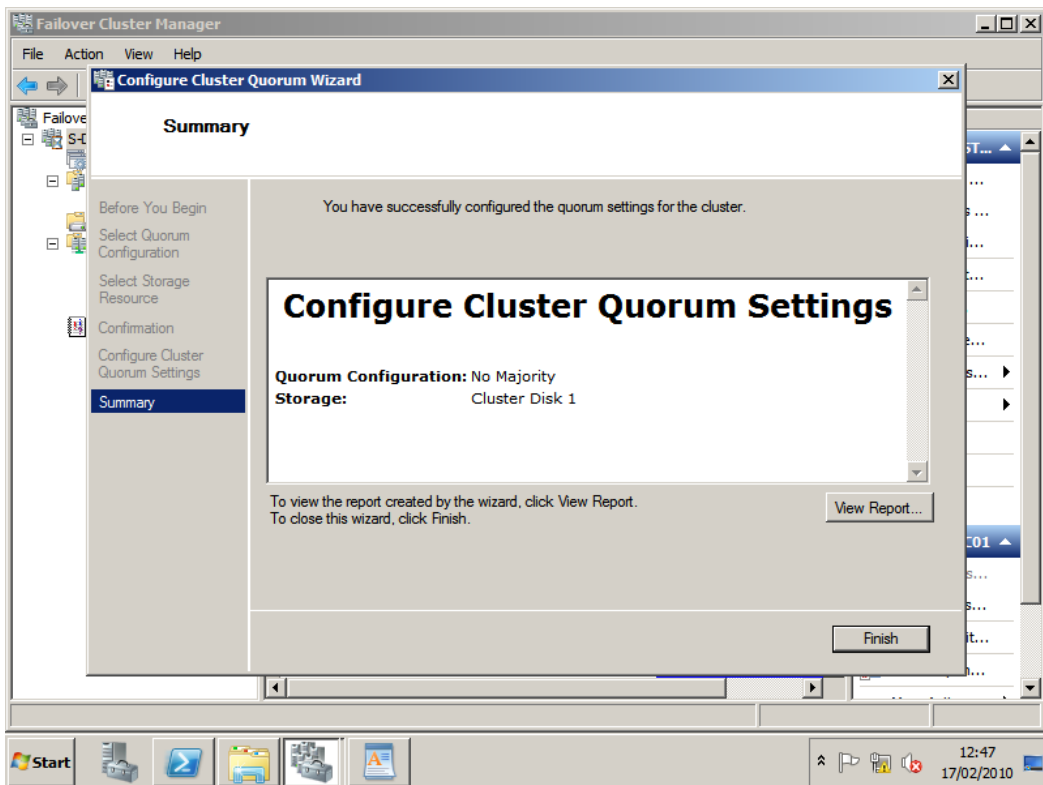
Review the confirmation screen and click "Next".



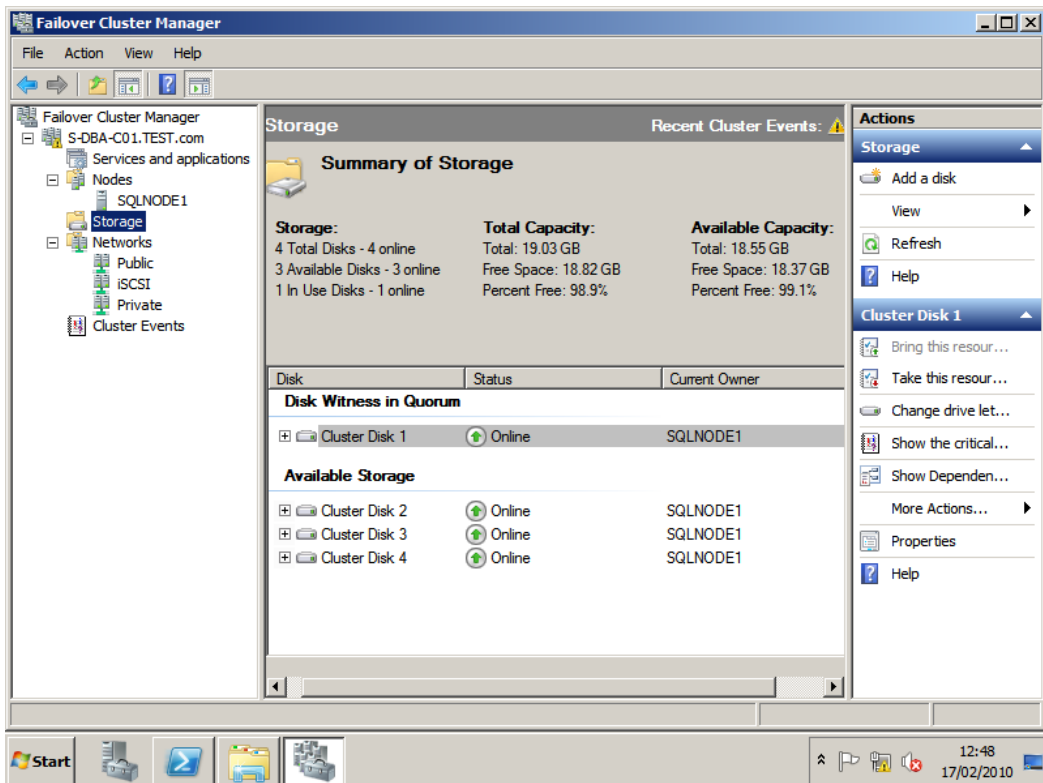
Cluster reconfiguration progress!



Finally click “Finish” to complete the re configuration.

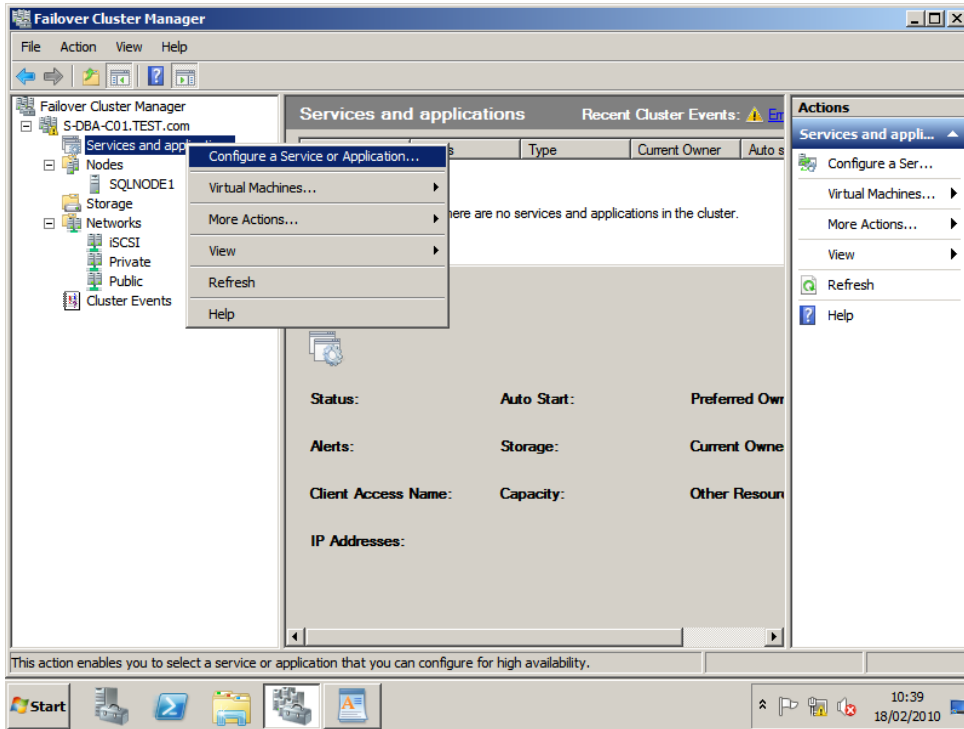


Disk 1 (Q drive) now set as the Quorum witness resource.

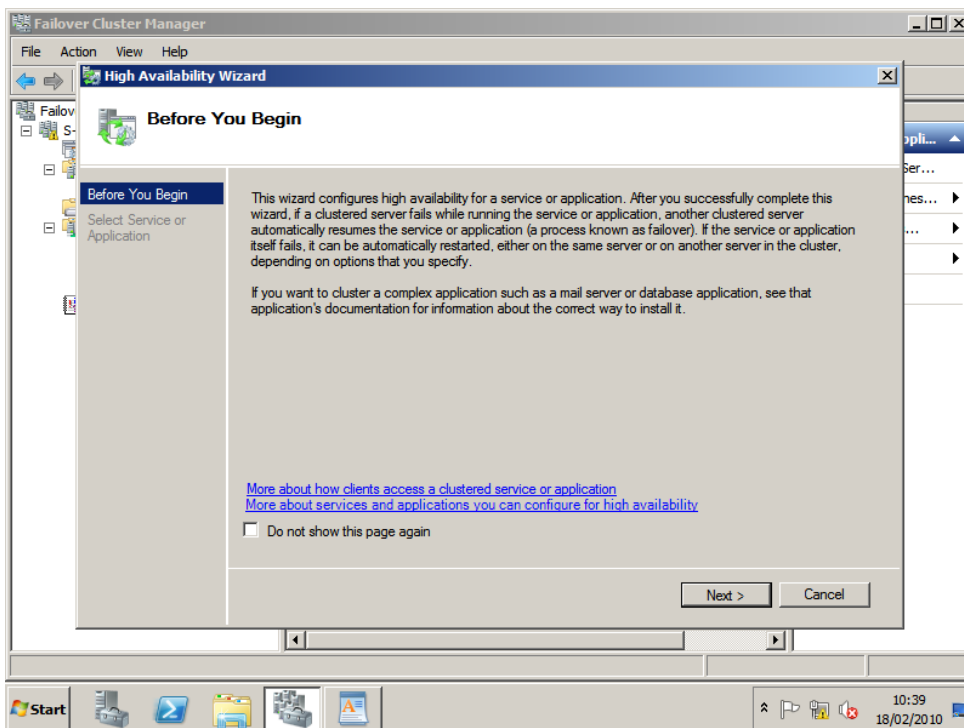


### 3.6 CREATE THE MICROSOFT DISTRIBUTED TRANSACTION COORDINATOR RESOURCE

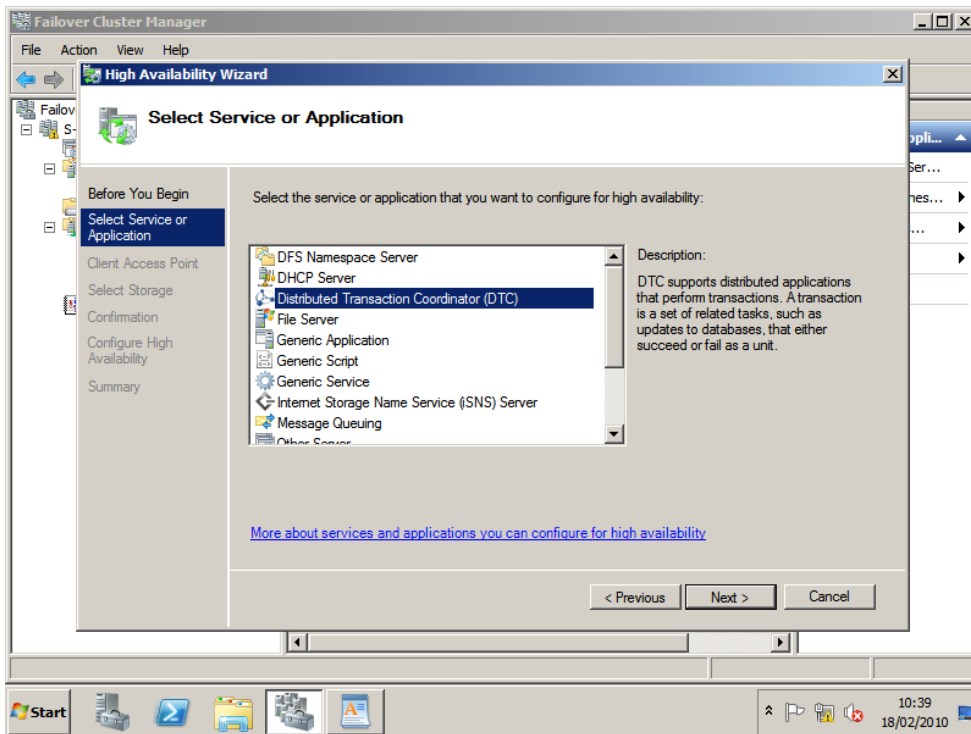
Next we need to create a cluster resource and group for the Distributed Transaction Coordinator service. This action is now wizard driven in Windows 2008 MSCS. In Failover Cluster Manager right click 'Services and applications'



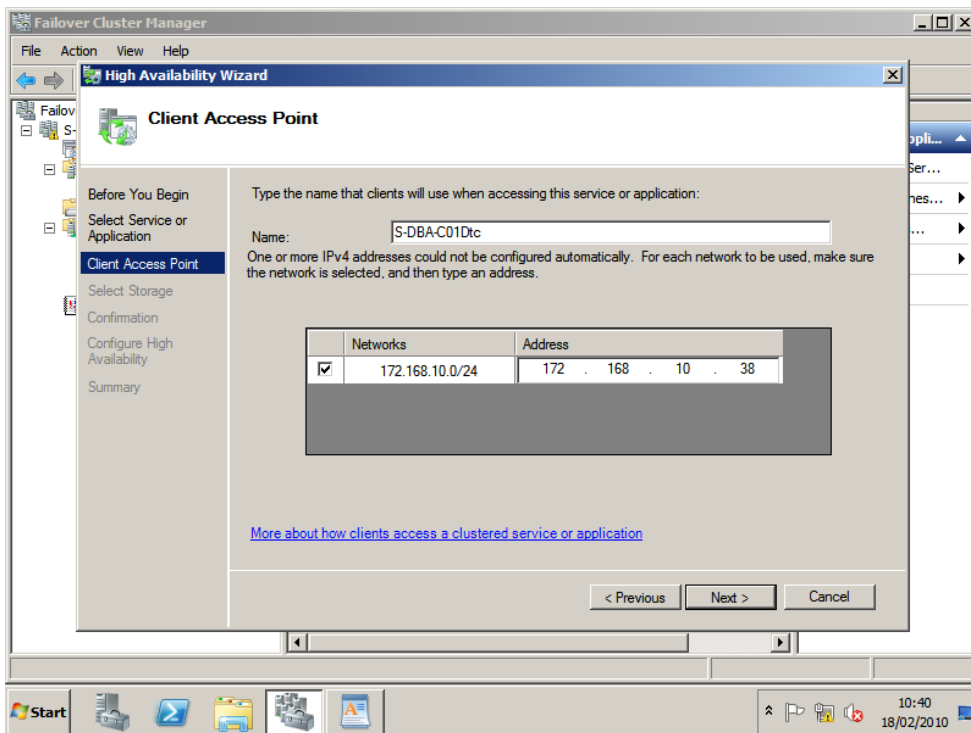
Review the following screen and click "Next".



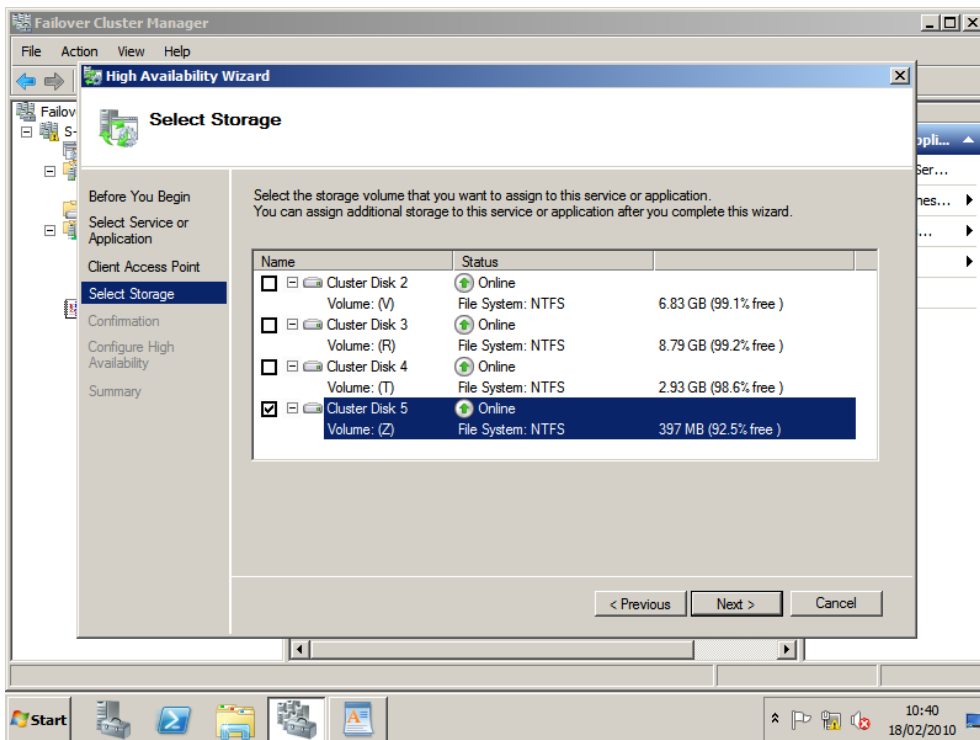
Select the “Distributed Transaction Coordinator” and click “Next”.



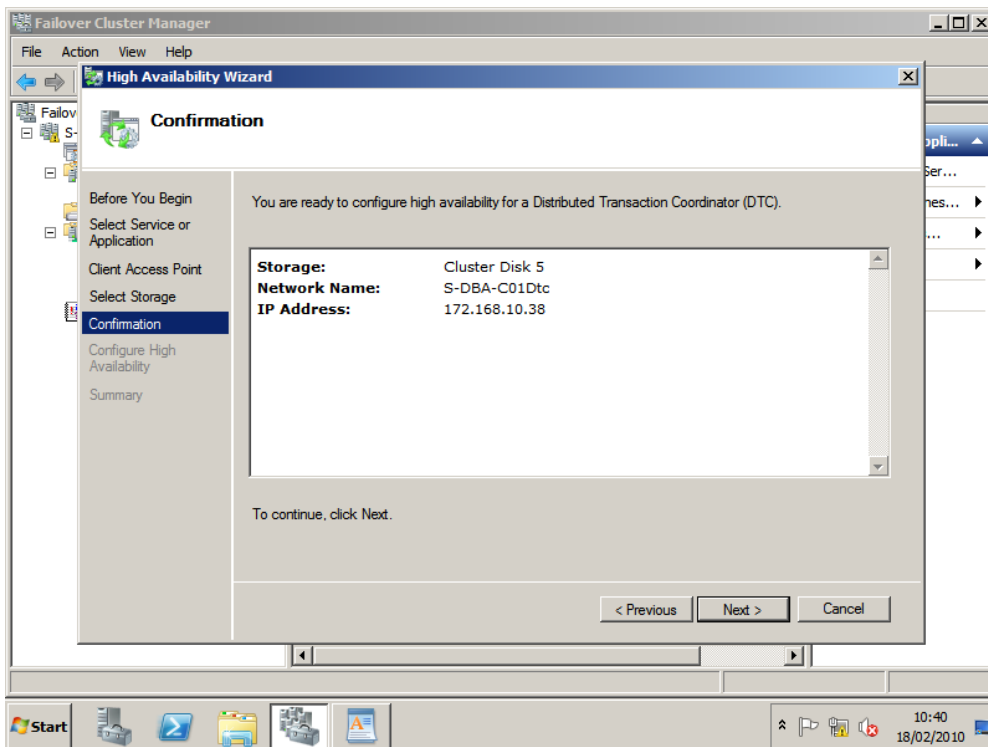
Supply a unique network name and IP address and click “Next”.



Select the storage drive to use!

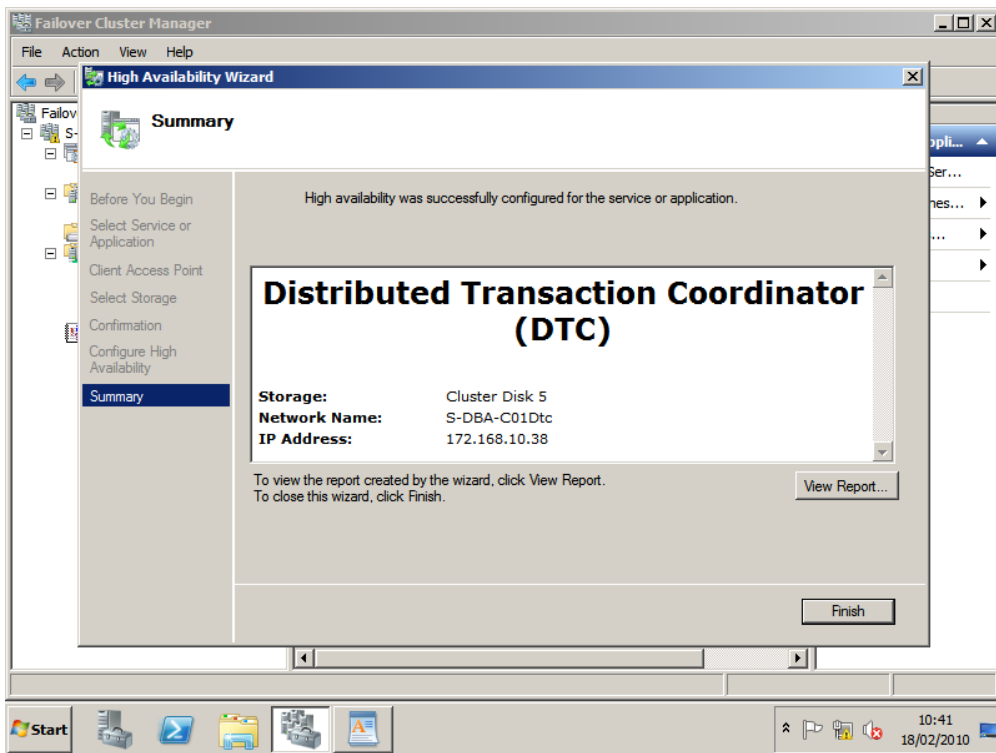


Click "Next" to create the clustered service.

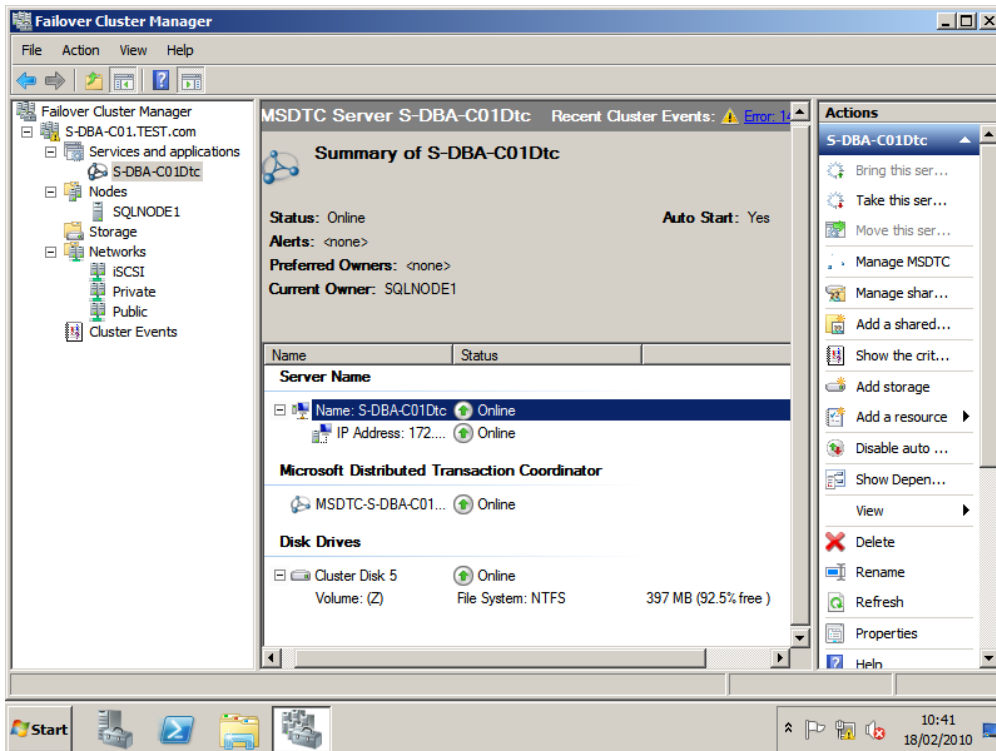




Click "Finish" to complete the wizard



The DTC clustered service is now configured and online.



## 4 INSTALLING THE SQL SERVER INSTANCE.

For clarification the following terms are used;

VM is a virtual machine.

NIC is a network Interface Card.

vNIC is a virtual Network Interface Card.

DC is a Windows Domain Controller.

NOS refers to the Windows operating system.

HA is the VMWare high availability technology used by VMWare Virtual Infrastructure.

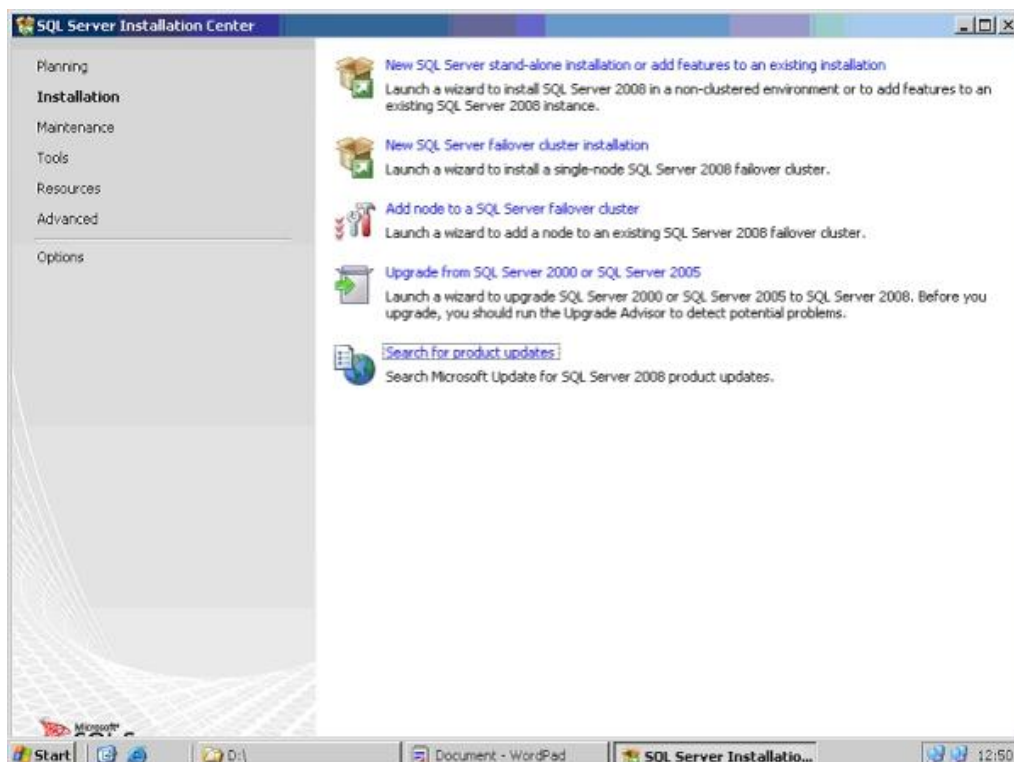
ESX is the server operating system used by host machines in the VMWare Virtual Infrastructure.

You should have already created the user account(s) for the SQL Server services and also the cluster Windows groups for the services. With that done it's now time to start the installation. Launching the Setup on Node 1, Installation proceeds as shown below.

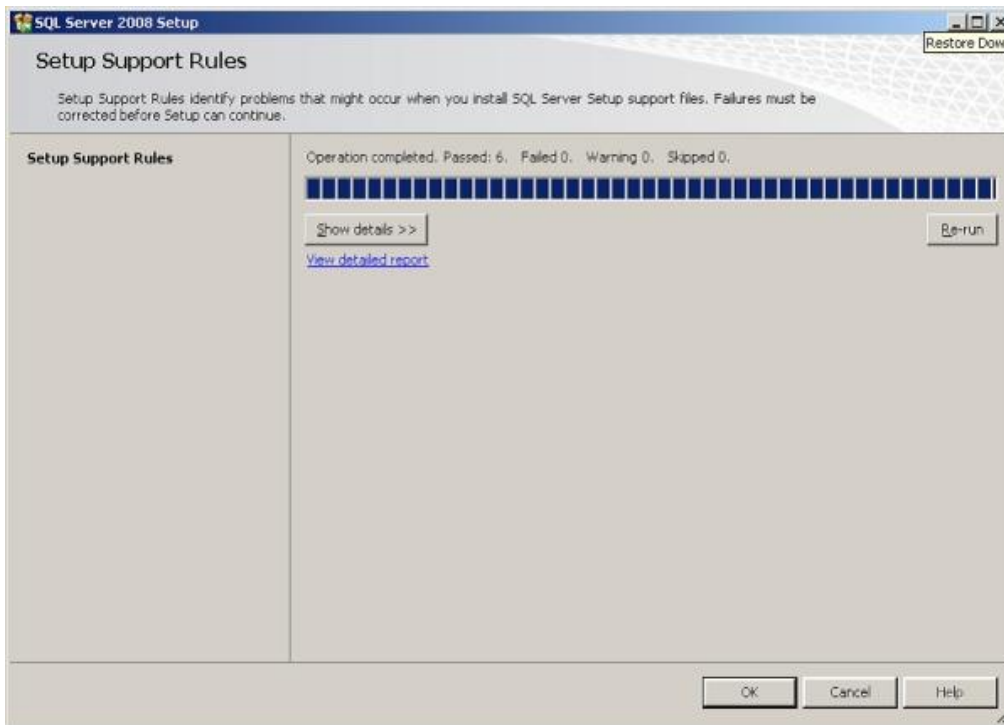
*Note: The installation process between SQL Server 2005 and SQL Server 2008 has changed extensively.*

### 4.1 INSTALLING THE FIRST NODE

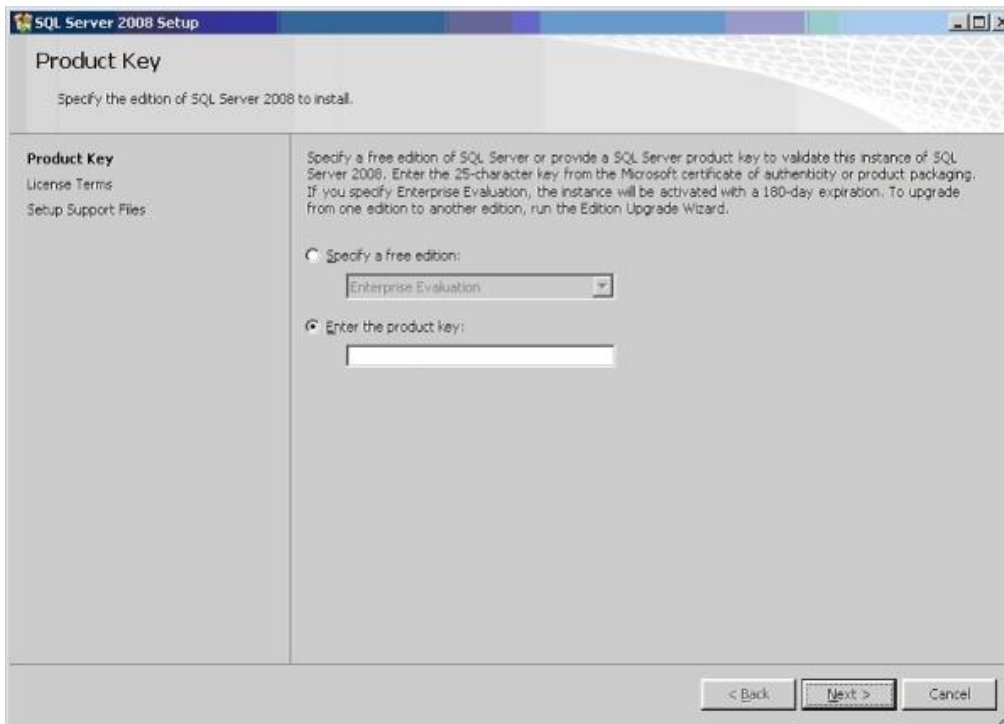
Click Installation and then select 'New SQL Server Failover Cluster'.




You will eventually see the following. Click 'OK' if the checks have been successful.



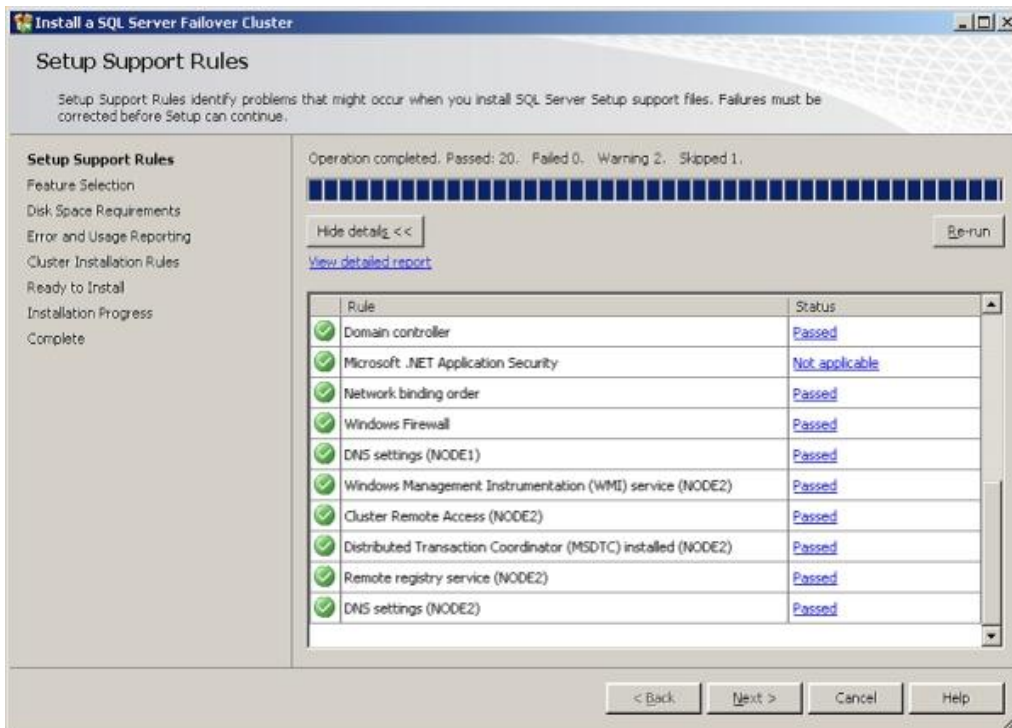
The 'Product Key' screen will appear. Either select an edition or enter your valid licence key to continue and click "Next".



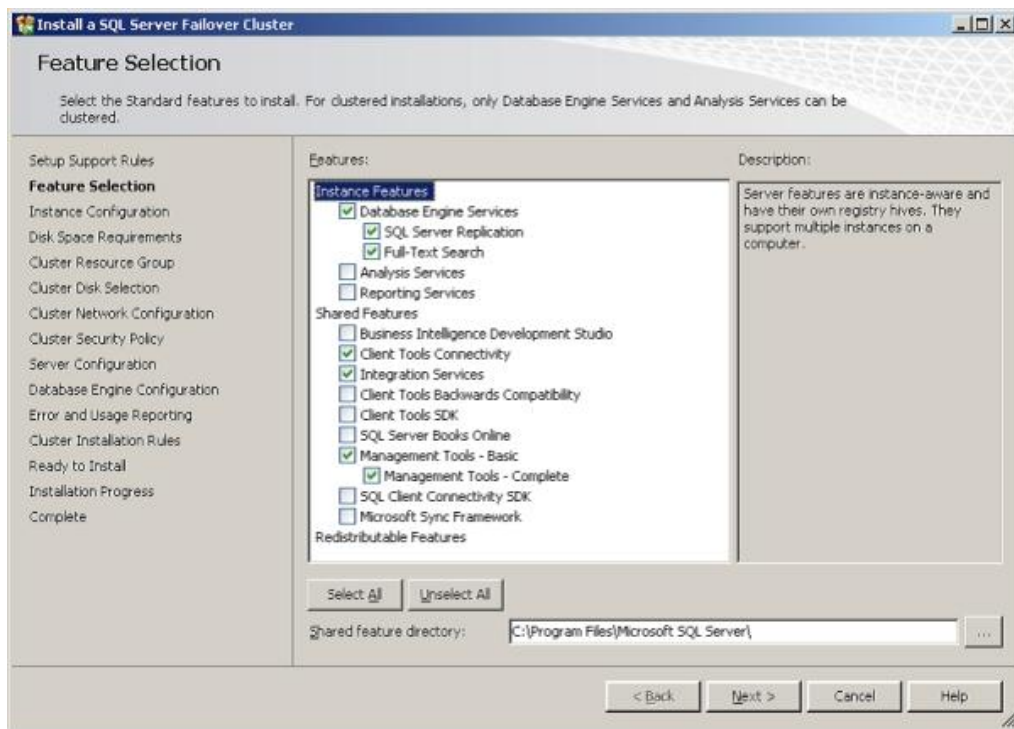
Click 'Next' through the Licence Terms (assuming you wish to accept them). You will be presented with the 'setup support files' install dialog as shown below.



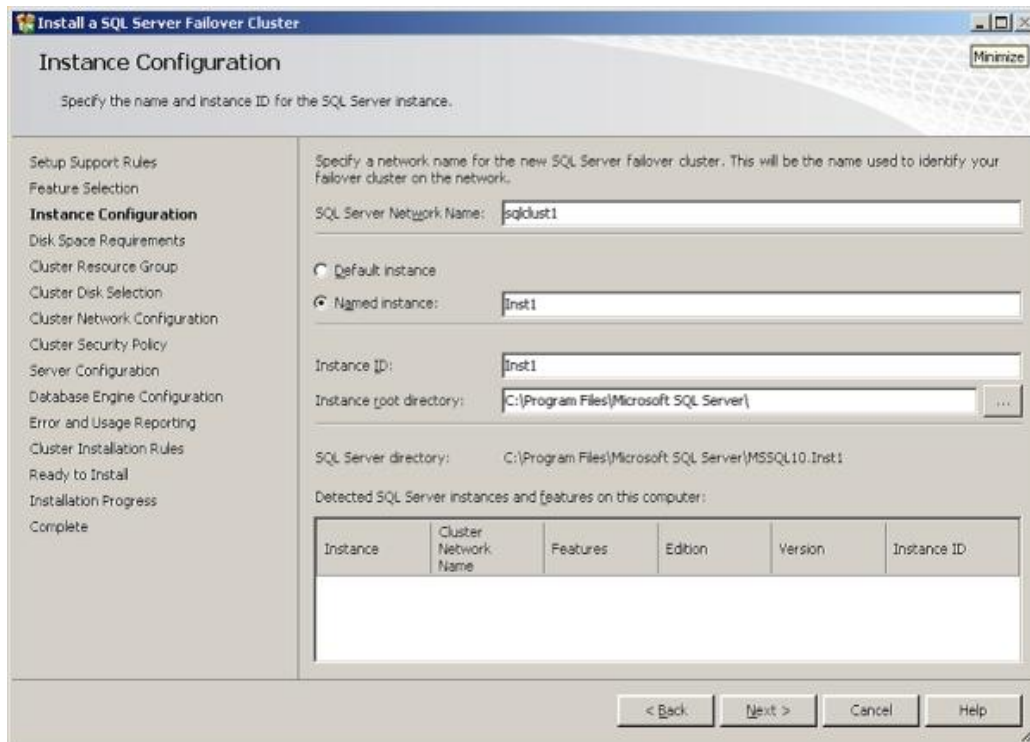
Click 'Install' through this and ensure the 'Setup Support Rules' checks complete as below.




If they all display the green check, click 'Next' and select the features to install.



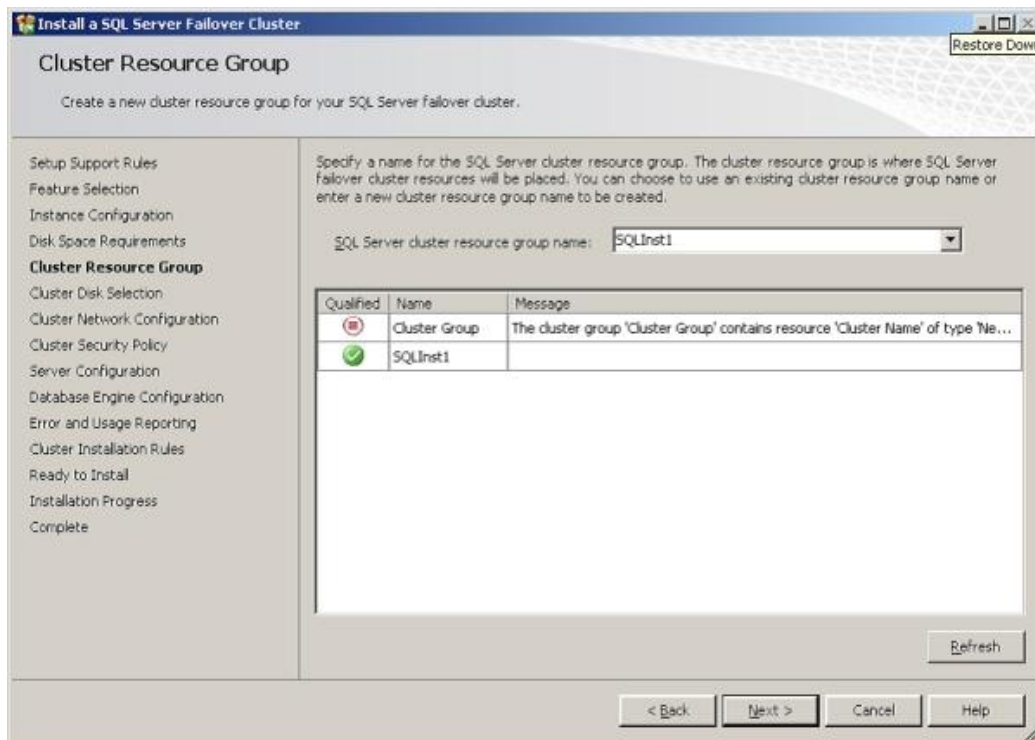
After making your selection, click 'Next' and the 'Instance Configuration' dialog appears as shown below.



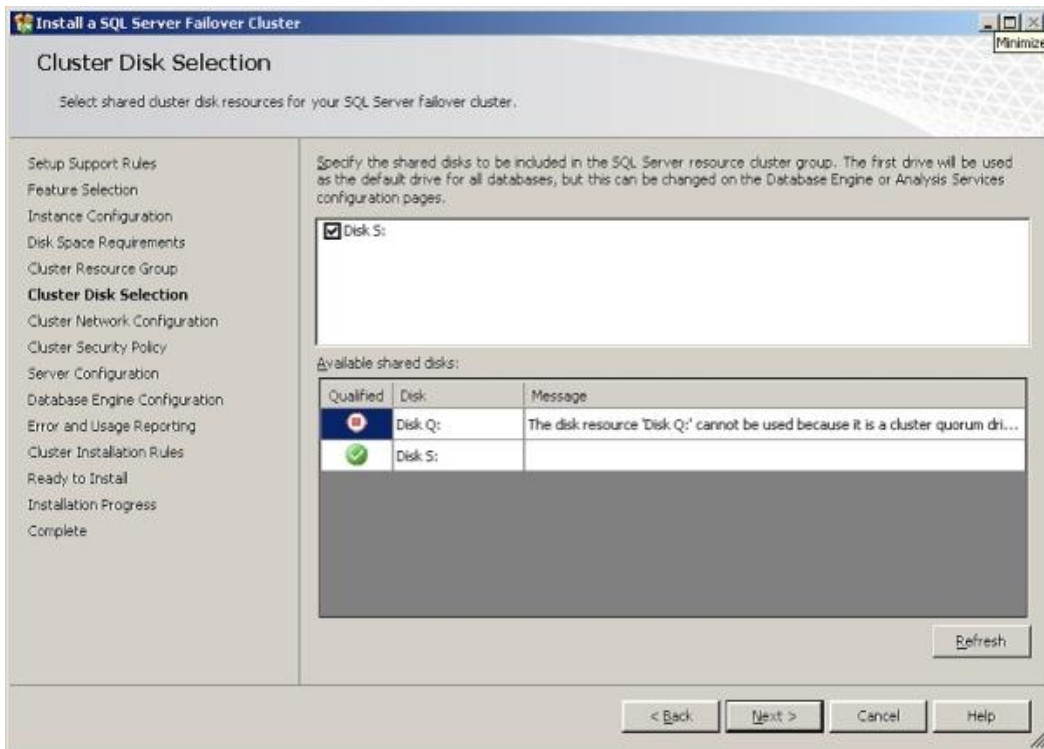
Supply the Instance configuration and click 'Next' to proceed to the Disk Space Requirements.



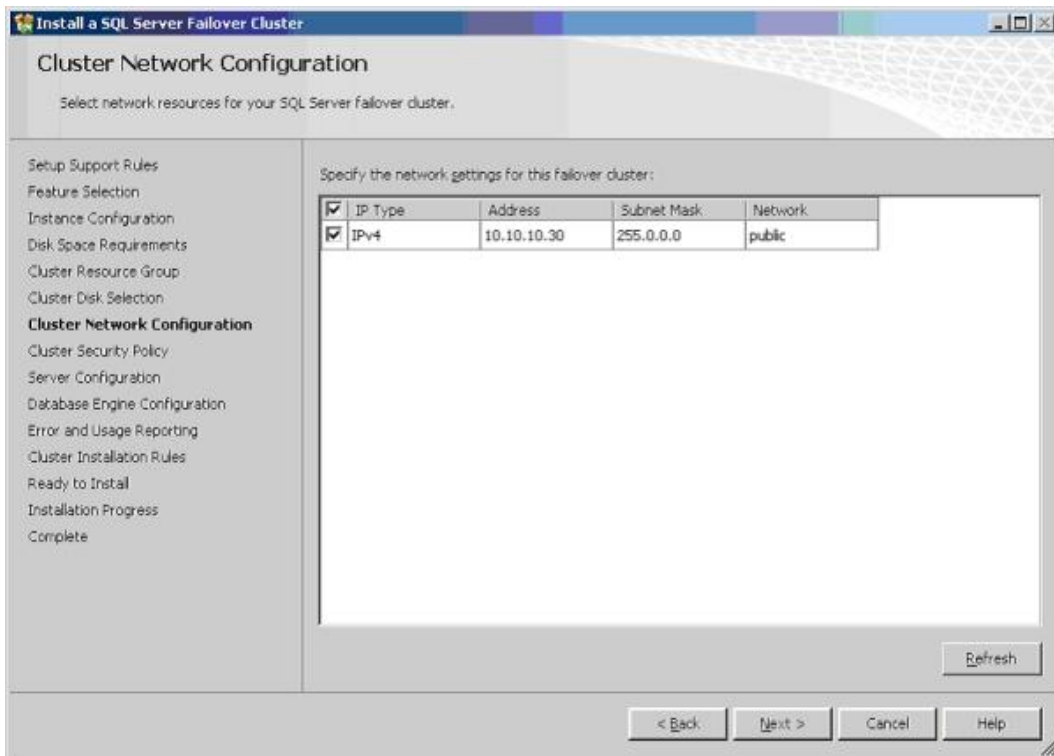
Click 'Next' and you will be prompted for the Cluster Resource Group to use for your SQL Server 2008 Cluster.



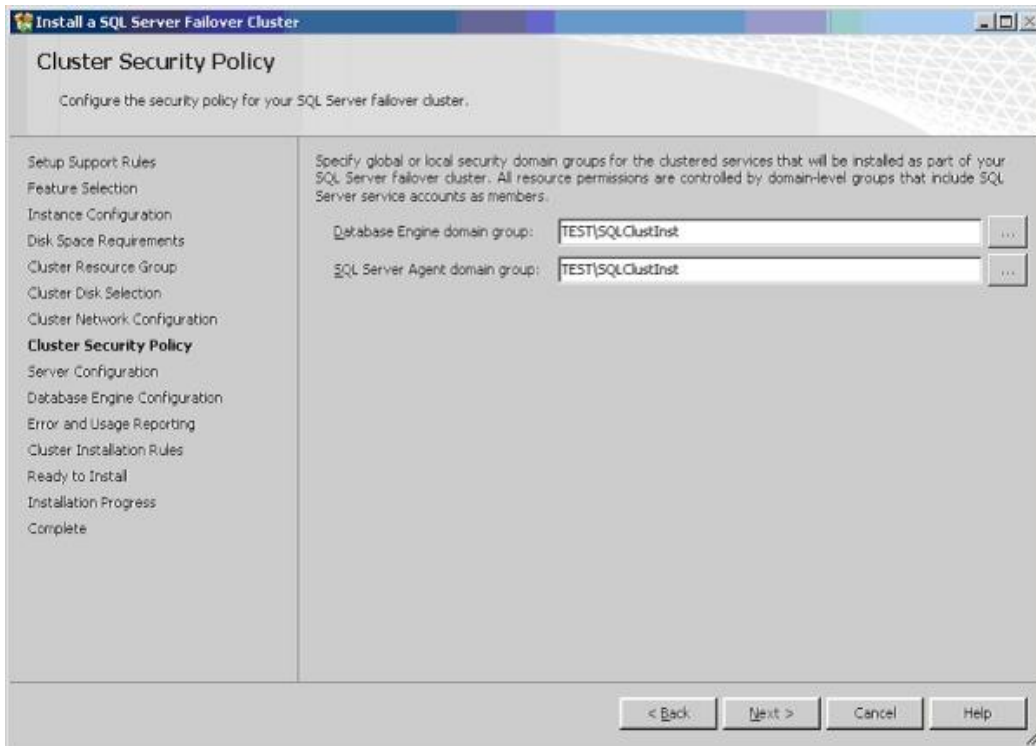
Click 'Next' again and you will be prompted for the Cluster Disk Resource to use.



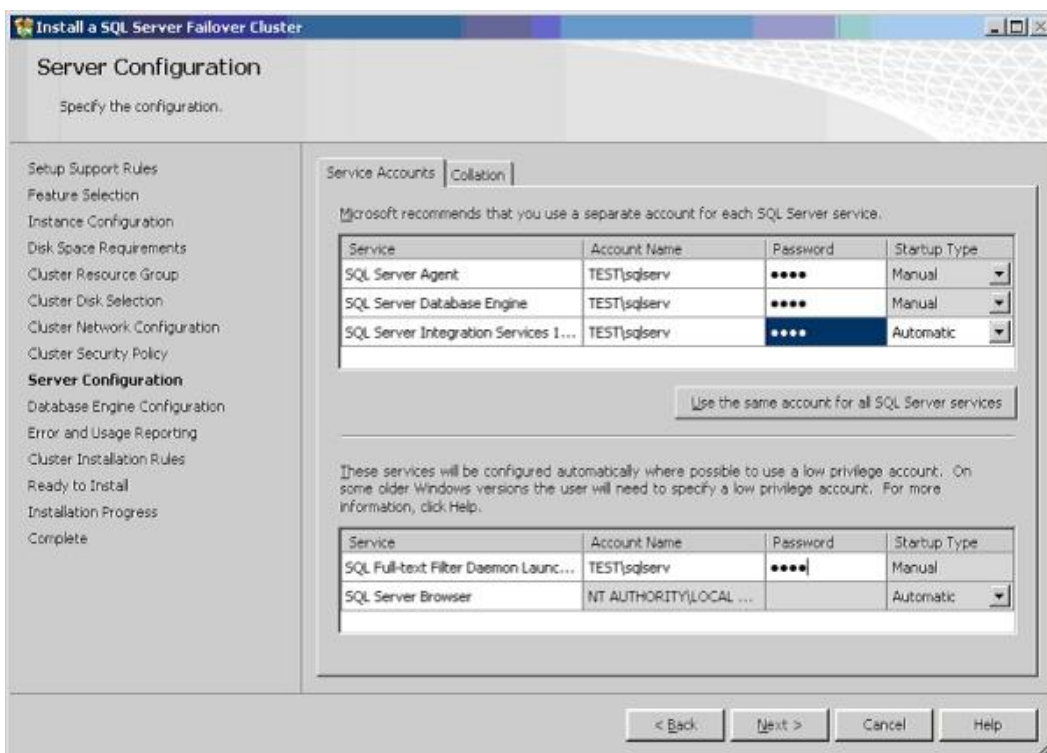
Click 'Next' and supply the network configuration.



Click 'Next' and provide the details of the previously created domain security groups. The service account the SQL Server Services run under, must be a member of these domain groups.

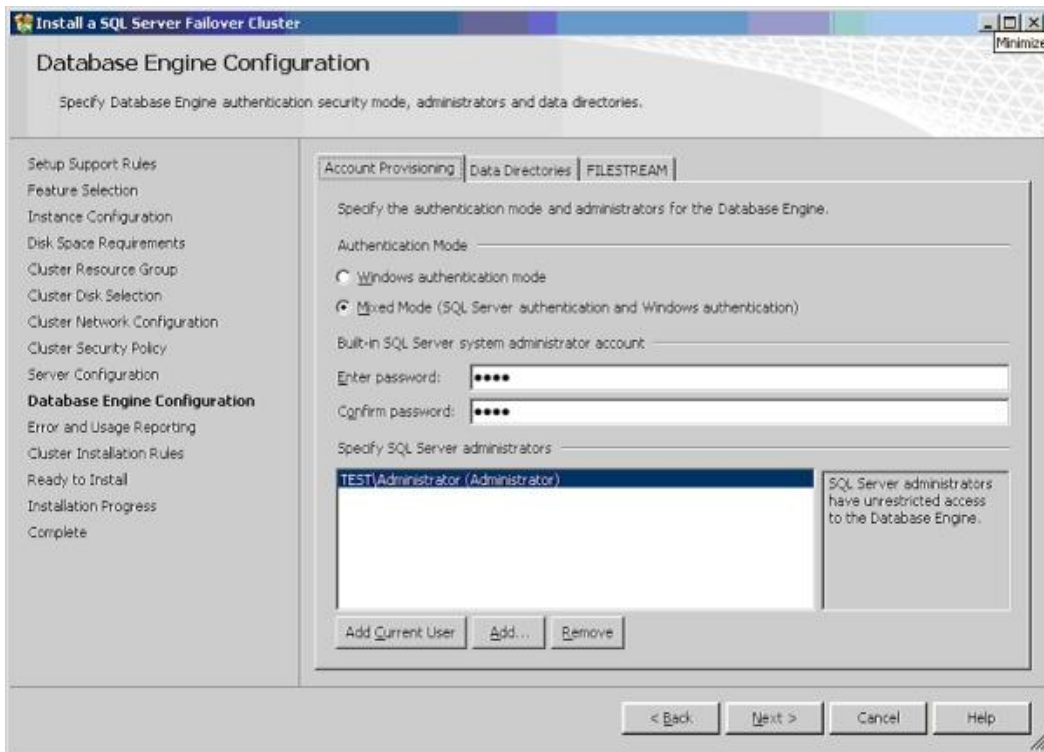


Click 'Next' through the wizard and you will see the Server Configuration dialog. Here you provide details of the SQL Server services security context. Also set the collation type on the 'Collation' tab.

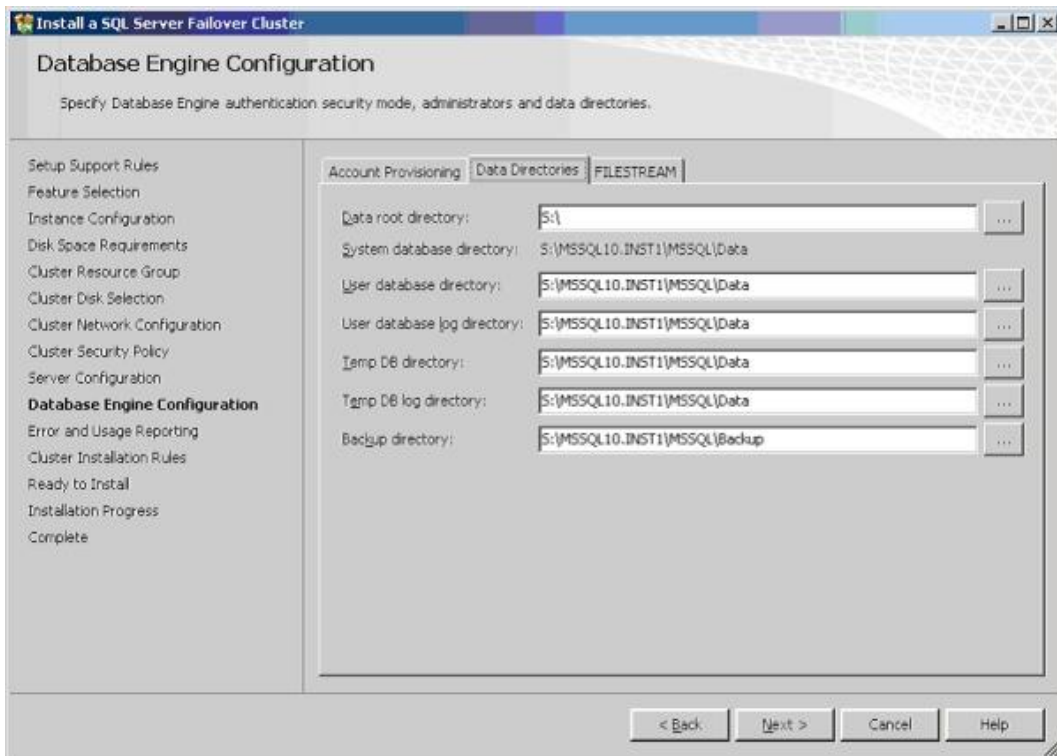




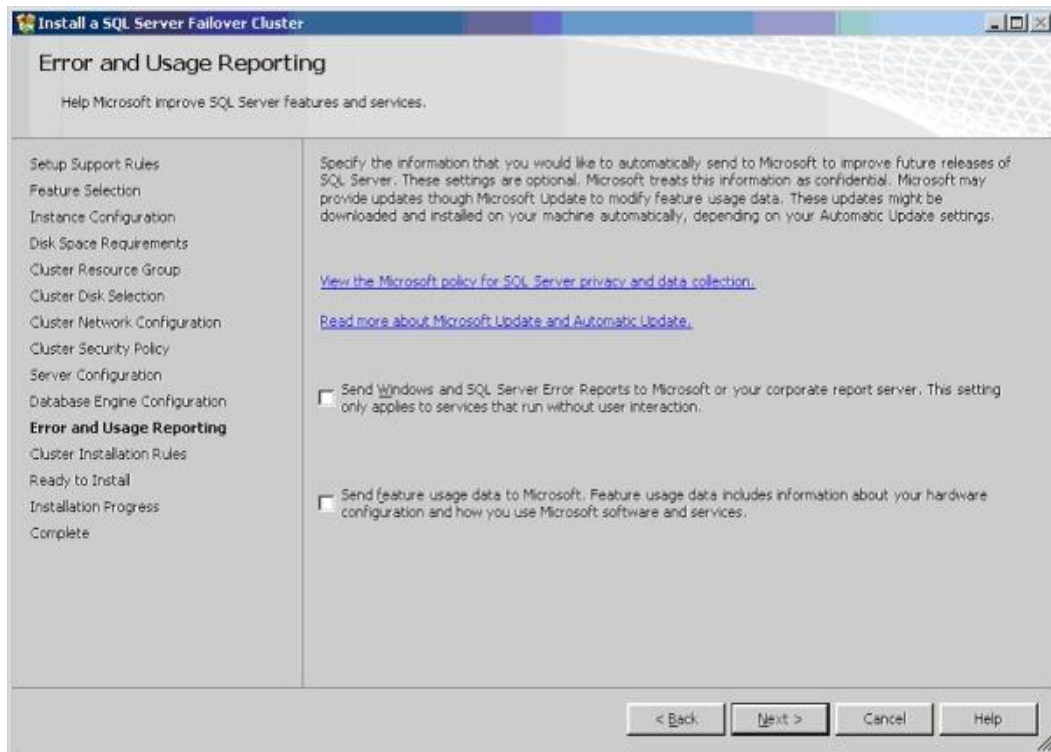
Click 'Next' and supply the Database Engine Configuration, Account Provisioning...



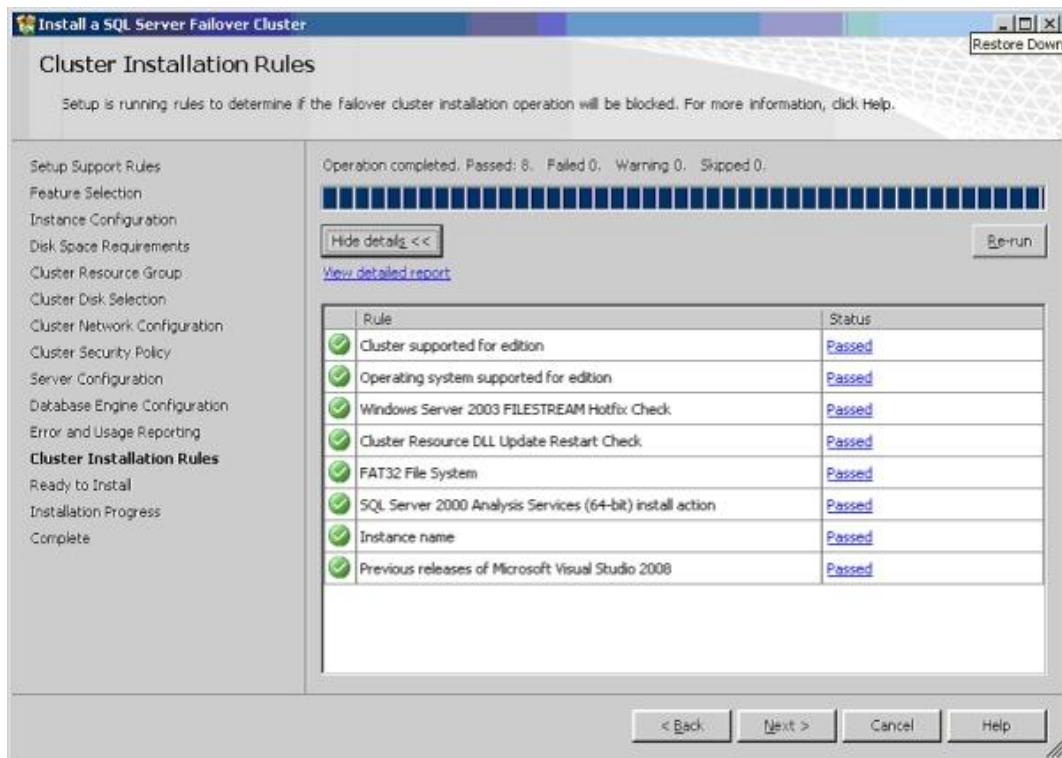
And Data Directory configurations. Enable 'Filestream' if required



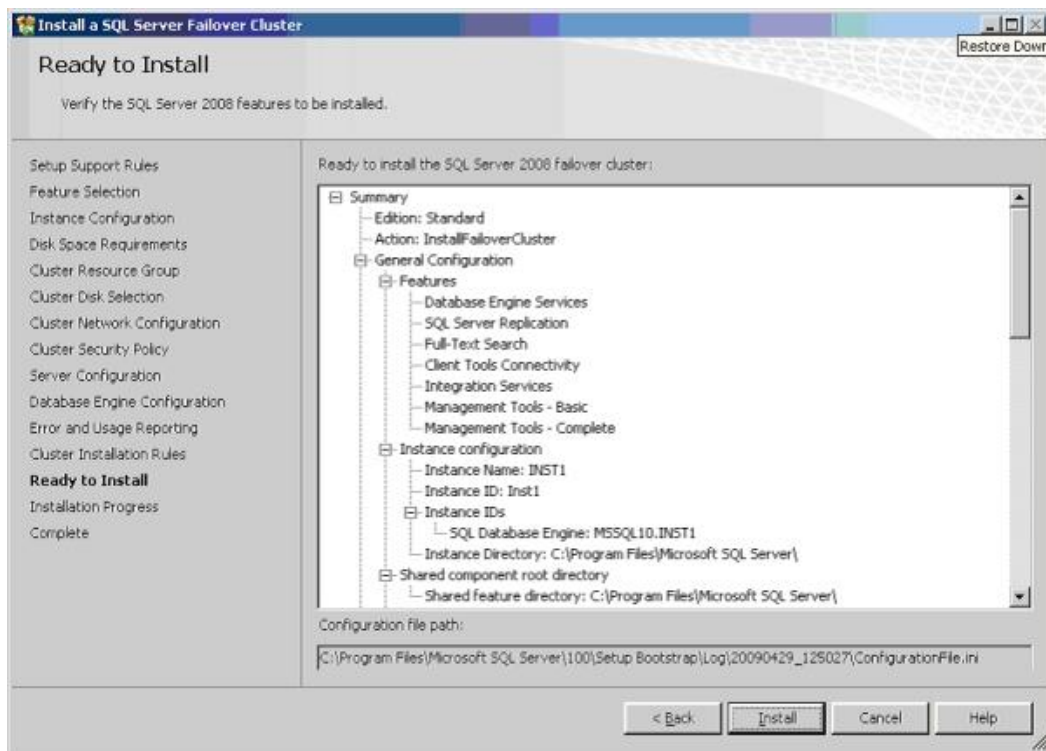
Click 'Next' and proceed to the 'Error and Usage Reporting' dialog.



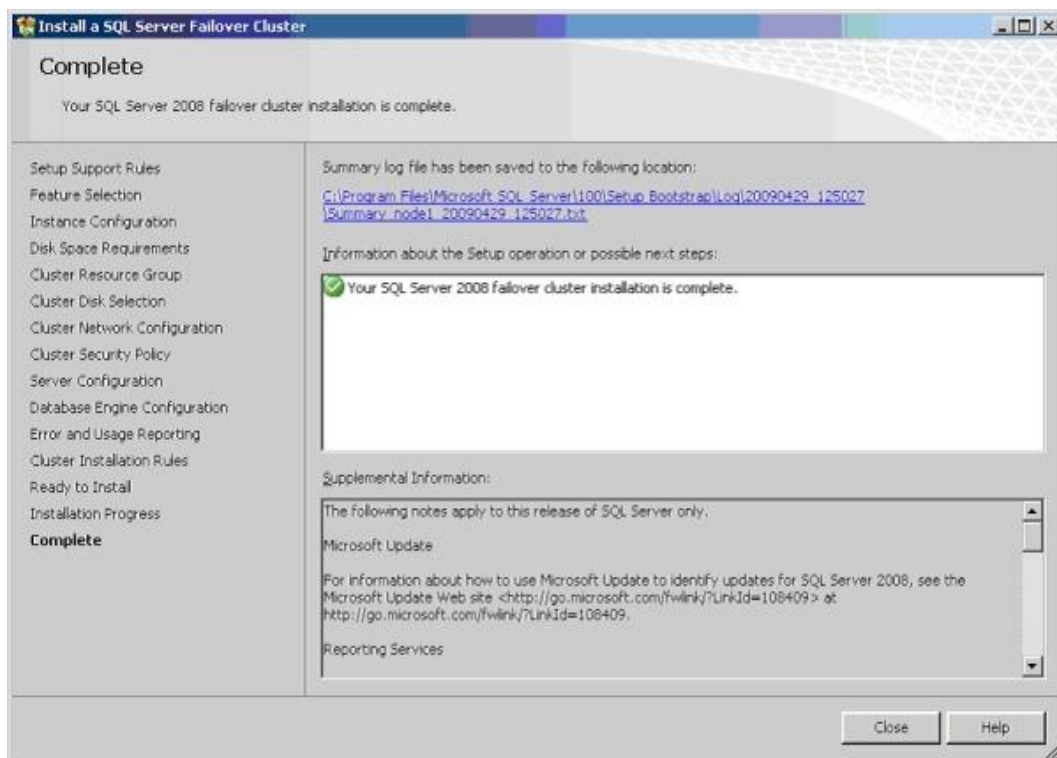
Click 'Next' and proceed to the Cluster Installation Rules dialog.



Ensure all checks complete and click 'Next' to review your summary.



Click 'Next' and follow the installation progress. Once the installation has completed you should see the following...



This installs SQL Server to the first\active node, you now have to launch setup on each cluster node you wish to participate in the SQL Server cluster.

*Previously SQL Server 2005 would deploy to all nodes in the configured cluster group and the installation progress for all nodes was viewable on the progress dialog via a drop down list.*

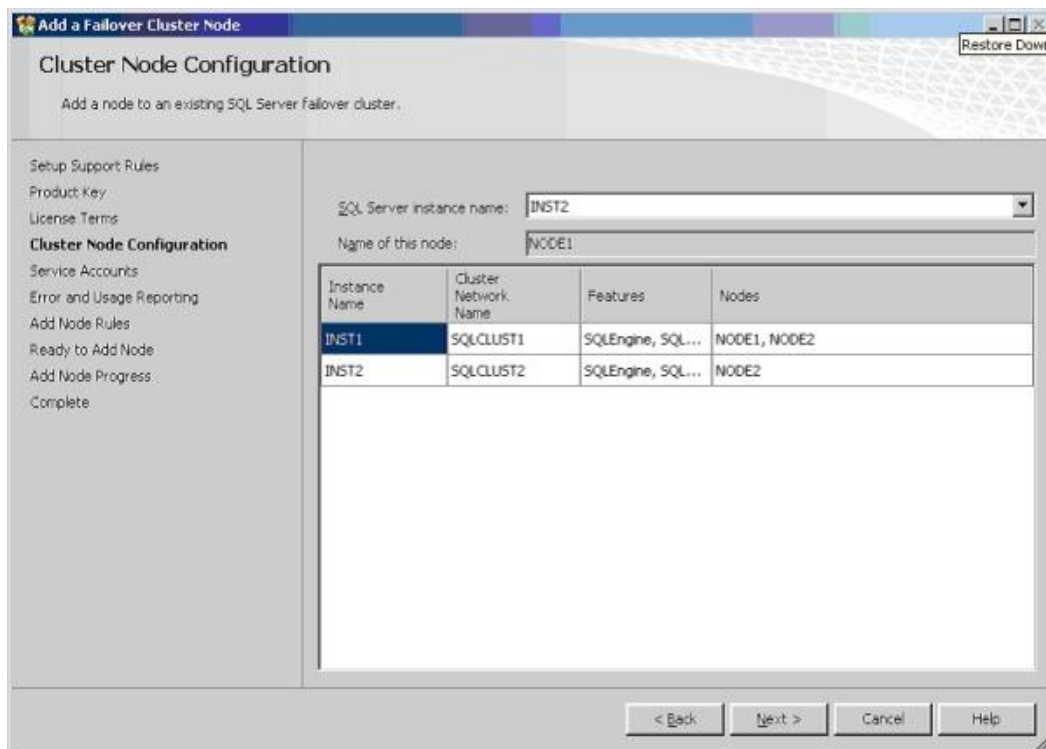
## 4.2 ADD A CLUSTER NODE

To add a node, launch setup on the node to be installed and from the Installation menu select 'Add Node to a SQL Server Failover Cluster'. Installation proceeds in a similar way to the new node installation. You will see the dialogs for 'Setup Support Rules', 'Setup Support Files'. At some point after this you are required to provide the product key, unfortunately there is a bug around this portion of the installer and the following error may be encountered...

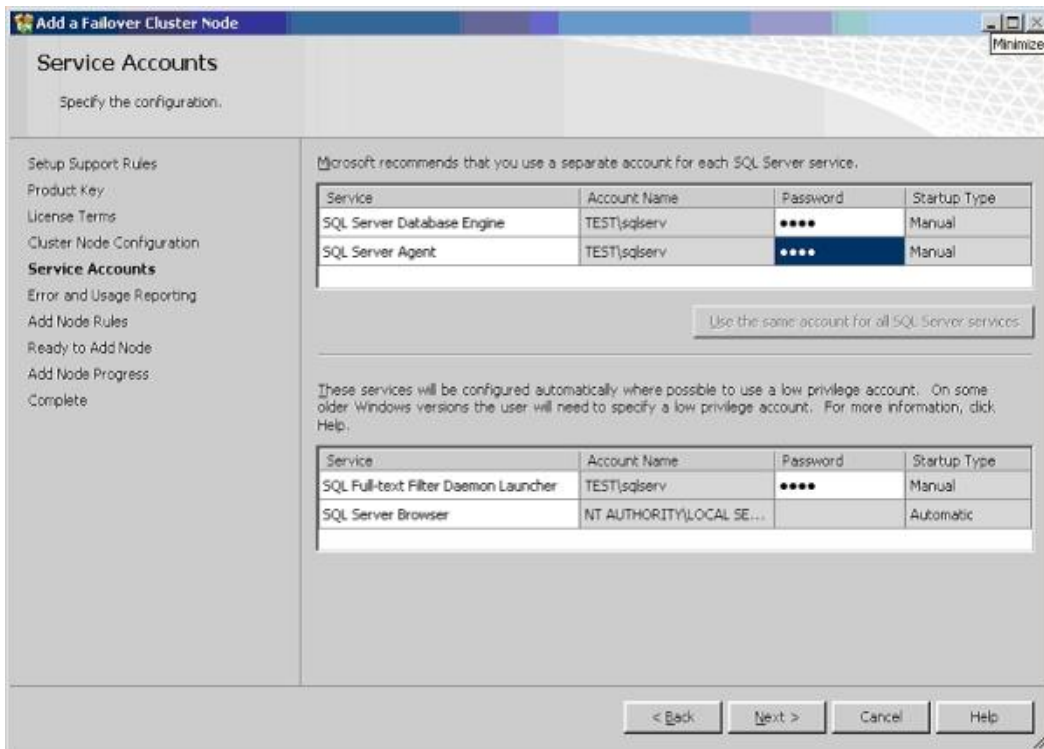
'The current SKU is invalid'

Although a hotfix is available, to work around this remove the default licence key and click 'Next'. Now click 'Back' and manually type the key then click 'Next' again. Installation should proceed without error. The node configuration should show as follows

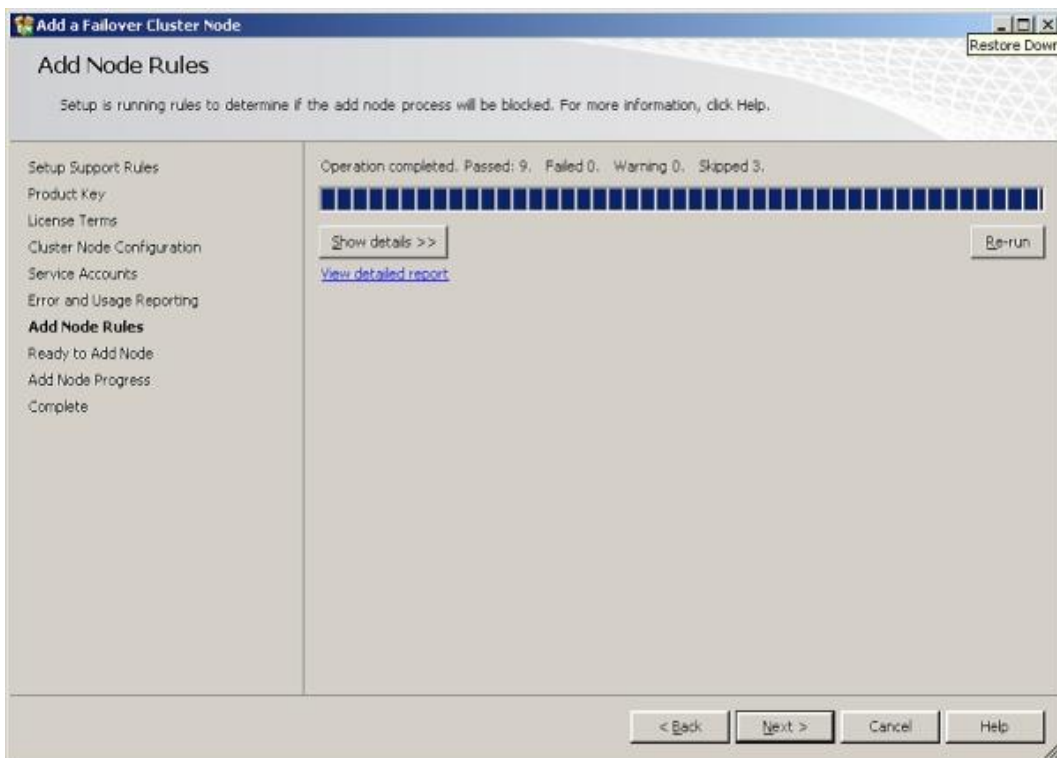
*Note the extra SQL Instance that is previously configured.*



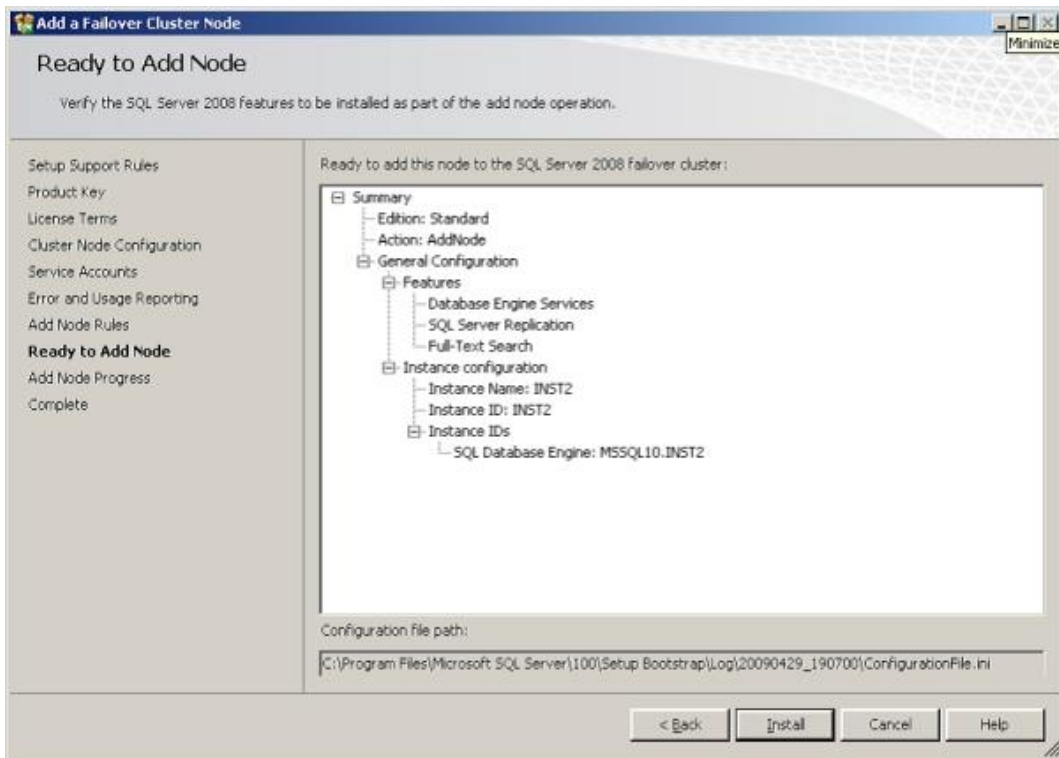
Select the InstanceName for which you are adding the node to and click 'Next' to proceed to the Service Accounts dialog.



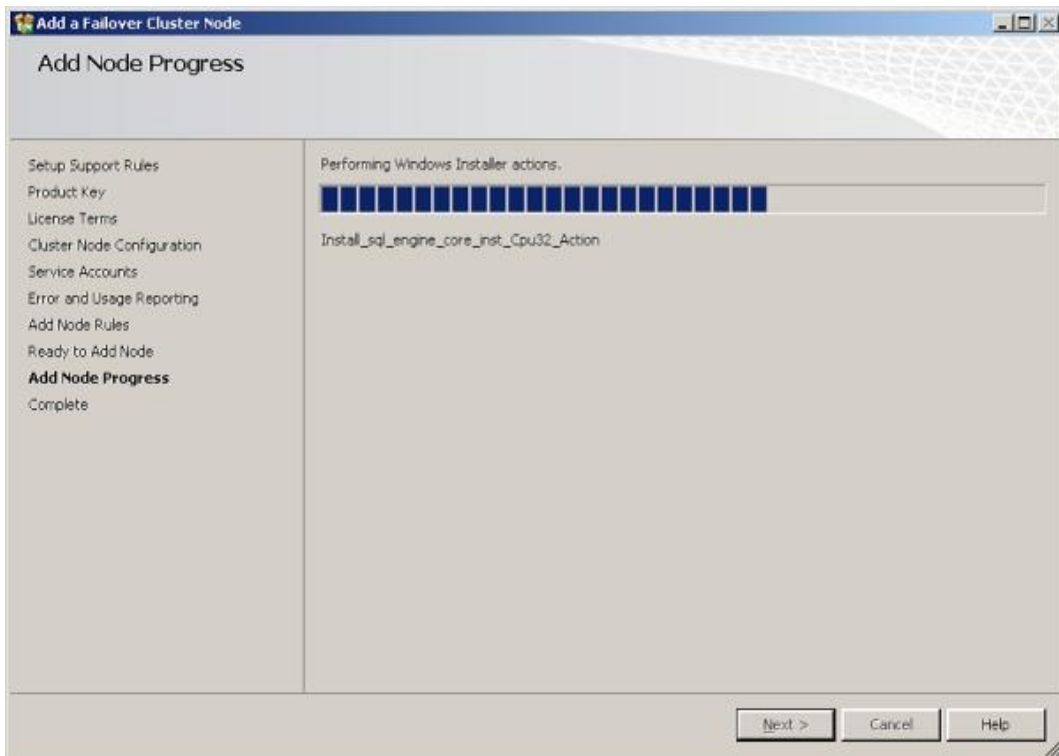
Configure the service account details and click 'Next' through the wizard past the error and usage reporting dialog. You should see the Add Node Rules dialog as shown below.



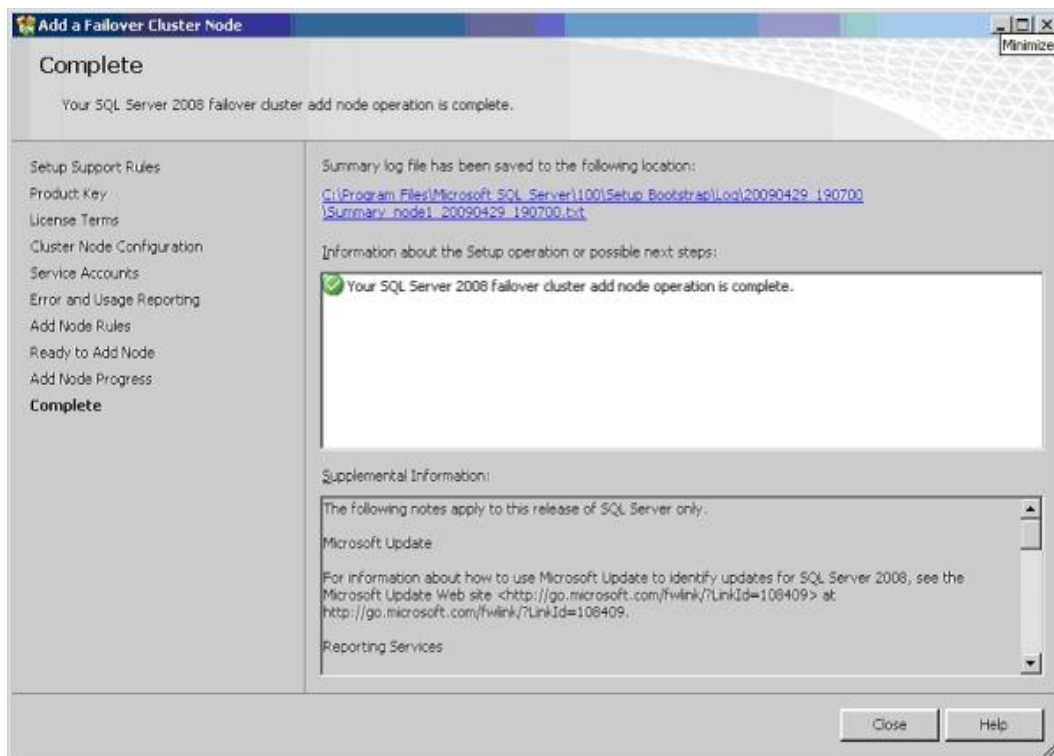
Click 'Next' and view the Add Node Rules. Click 'Install' to continue.



You will see the progress dialog below.



Once installation has completed you should see the following screenshot.



Should you receive any errors, address these issues and if necessary re launch Setup.

### 4.3 So how would I have an active\active configuration?

The more astute among you would have noticed the already installed SQL instance in the Cluster Node Configuration screenshot above. To create an active\active cluster simply create the extra resources required on the FreeNAS VM and present them over the iSCSI transport to the cluster nodes and add them as disk resources to Failover Cluster Manager. Once created, launch the installation on a passive node. Install the SQL Instance as shown above and then add nodes to the new instance. Keep the quorum drive in the cluster group (move the group to the required node before installing your secondary SQL instance).

## 5 APPENDIX A NAMED INSTANCES & NETWORK NAMES

During SQL Server installation setup requires you to supply an instance name for the SQL Server instance you are installing. On a non clustered system, the machine name is the instance name prefix. A virtual network name is essentially the same as a machine name in that it must be unique on the network.

Installations of SQL Server allow only one default instance, after that the rest must be named.

Take the following scenarios

### 5.1 NON CLUSTERED

A non clustered server named [\MYSERVER](#) has 2 SQL Server instances installed to it. When the administrator installed SQL server they created a named instance using the name "Instance1" and a default instance.

To log on to each instance you would use,

Default

MYSERVER

Named

MYSERVER\Instance1

### 5.2 CLUSTERED

A 2 node clustered system exists using the following details,

Node1

Name = CL-01-001  
IP = 10.200.1.17

Windows Cluster

Name = CLU-01-001  
IP = 10.200.1.12

Default SQL Instance

Network Name = DB-01-011  
Instance Name =

Node 2

Name = CL-01-002  
IP = 10.200.1.18

Named SQL Instance

Network Name = DB-01-013  
Instance Name = MSUAT



You can see the complexity added to the system as there are now more names and IP addresses used for the same computers. In the clustered environment it is important to remember to disregard the node names and IP addresses when connecting to SQL Server. Everything is referenced by the Network name. In the clustered environment the Network Name and note the machine name forms the instance prefix.

To logon to the default instance you use

DB-01-011

To logon to the named instance you use

DB-01-013\MSUAT

Practice these by using names of your own to substitute the items above and ensure you understand default and named instances in both clustered and non clustered environments.