

# Creating a 2 node virtual SQL Server 2008 Cluster Configuration Using Windows 2003 MSCS

Reference No:  
Date: 10/08/2010  
File Name: Creating a 2 node SQL Server 2008 Cluster Configuration using Windows  
2003 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).....	29
3.4	ATTACHING THE ISCSI LUNS .....	37
3.5	INSTALLING THE WINDOWS 2003 CLUSTER .....	41
3.6	CREATE THE MICROSOFT DISTRIBUTED TRANSACTION COORDINATOR RESOURCE.....	53
<b>4</b>	<b>INSTALLING THE SQL SERVER INSTANCE.....</b>	<b>54</b>
4.1	CREATE THE SQL SERVER FAILOVER INSTANCE. ....	54
4.2	ADD A CLUSTER NODE TO AN EXISTING CLUSTERED INSTANCE .....	65
<b>5</b>	<b>APPENDIX A NAMED INSTANCES &amp; NETWORK NAMES.....</b>	<b>70</b>
5.1	NON CLUSTERED.....	70
5.2	CLUSTERED .....	70

# 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 2003 Enterprise 32bit and SQL Server 2008 Enterprise 32 bit (you may use 64 bit if your hardware supports it). 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 virtual network name 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

To re create storage scenarios you will use FREENAS to deploy a virtual SAN. The first steps in this article, install and configure the VMWare Server hypervisor and then the FreeNAS SAN virtual machine.

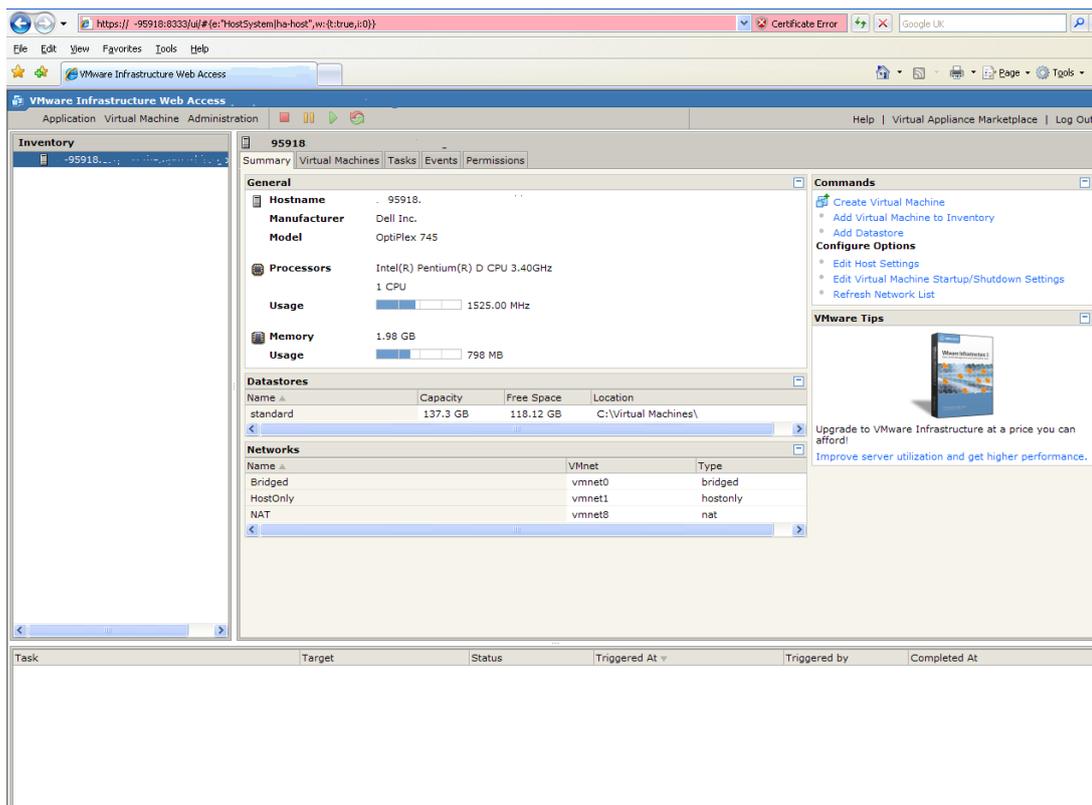
#### 3.1 INSTALL & CONFIGURE VMWARE SERVER

**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 Windows\domain account to the following local group on your pc

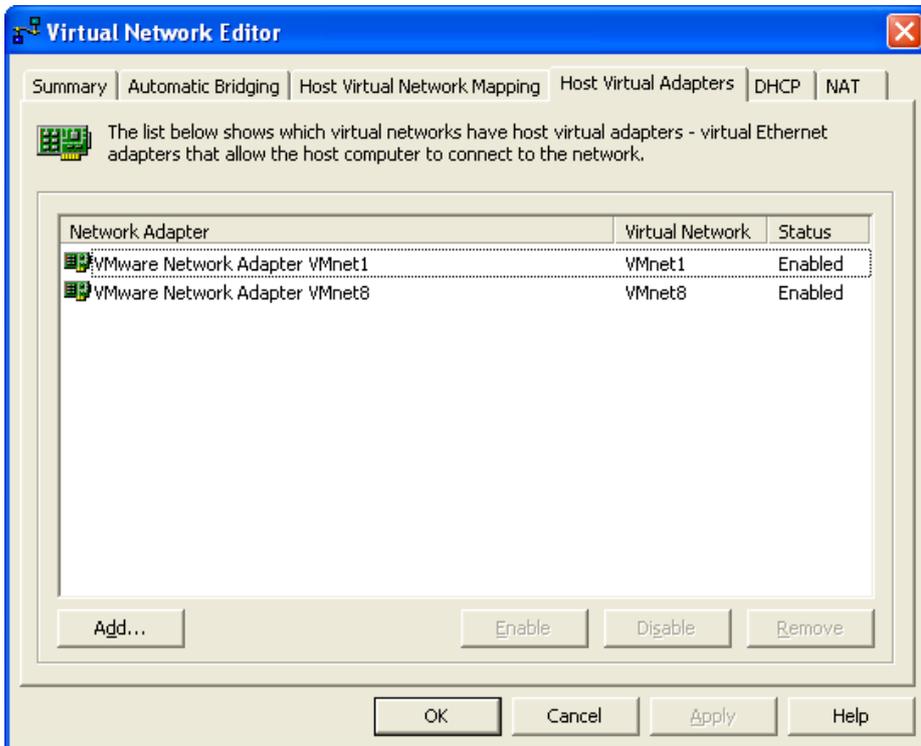
\_\_vmware\_\_

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

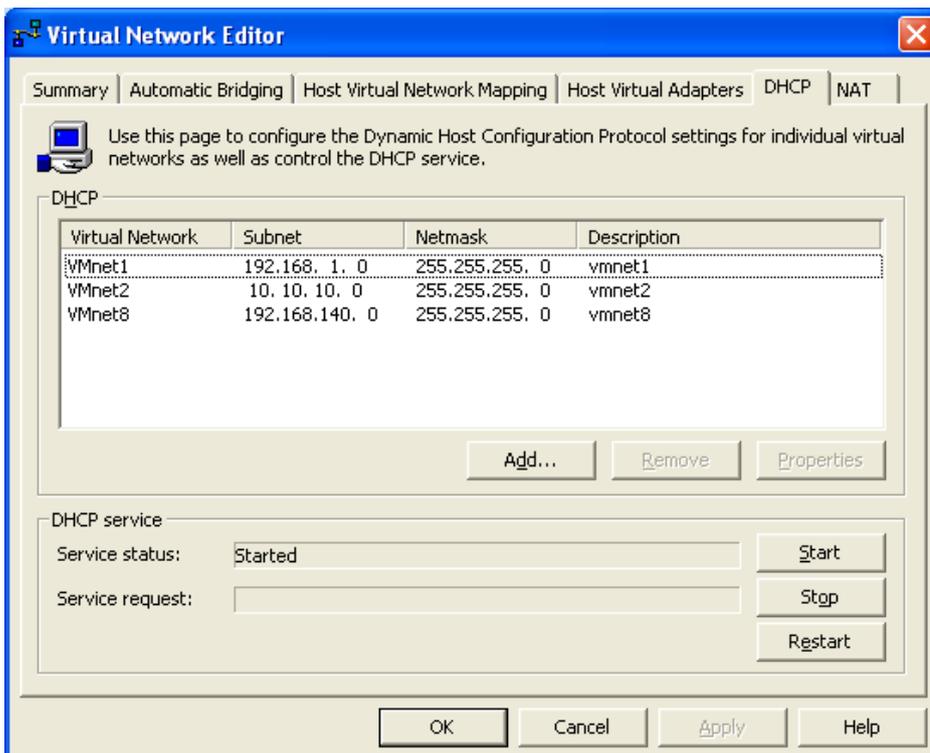


Open the VMWare “Manage Virtual Networks” option from the Start menu and go straight to the “Host virtual adapters” tab as shown below.

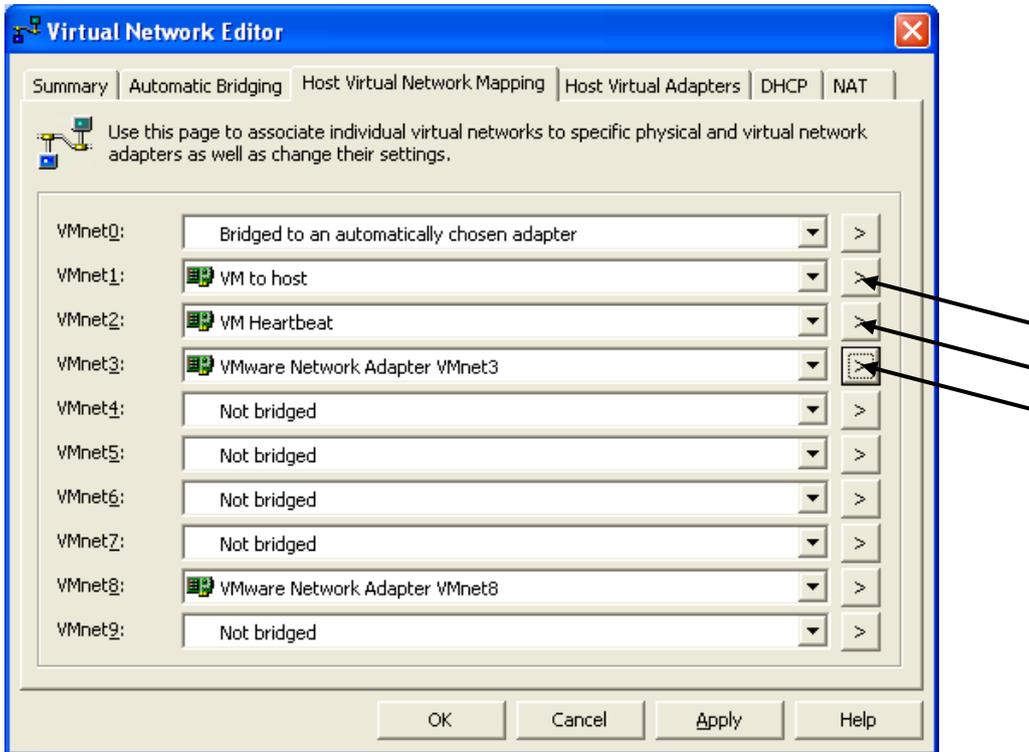
Add a new virtual adapter for VMnet2 and VMnet3 then click “Apply”



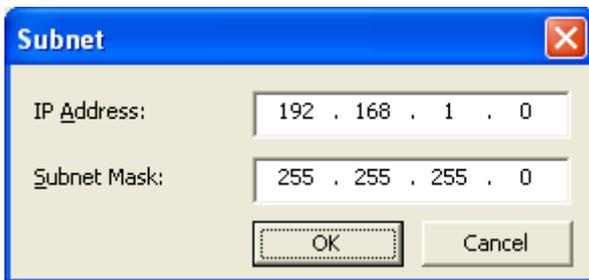
Now go to the DHCP tab and remove any DHCP assignments (click each item and remove) the click “Apply”.



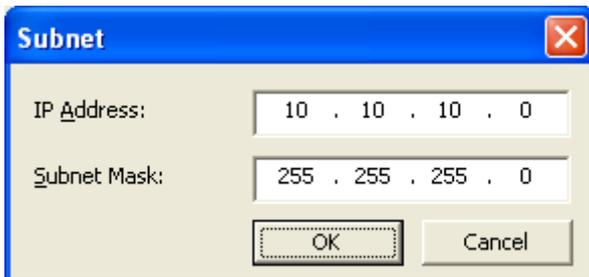
Go to the “Host virtual network mapping” option and change the subnets to be used for VMnet 1, 2 & 3. Do this by clicking the arrows (indicated) and selecting “Subnet” from the pop-up menu



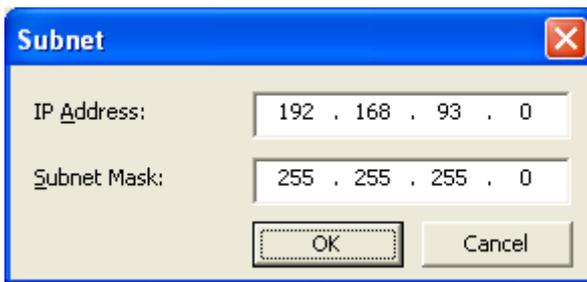
For VMnet1 set the IP address to 192.168.1.0



For VMnet2 set the IP address to 10.10.10.0

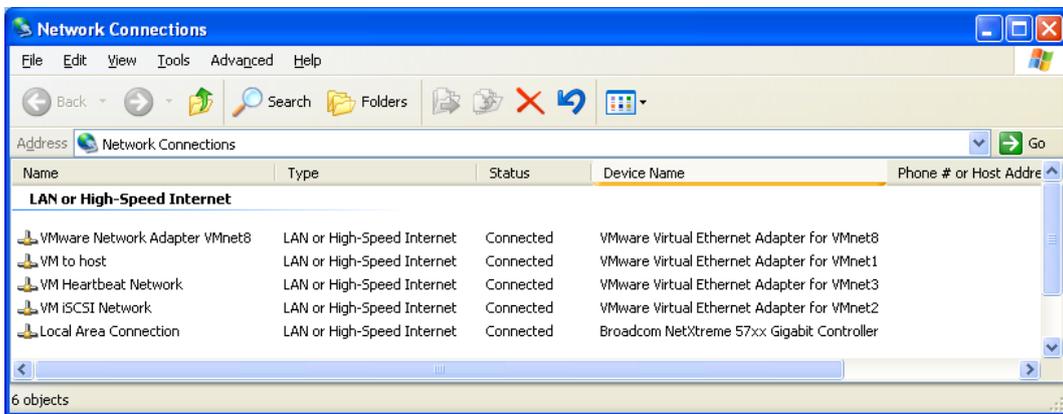


For VMnet3 set the IP address range to 192.168.93.0



Click "OK" to exit the virtual network editor

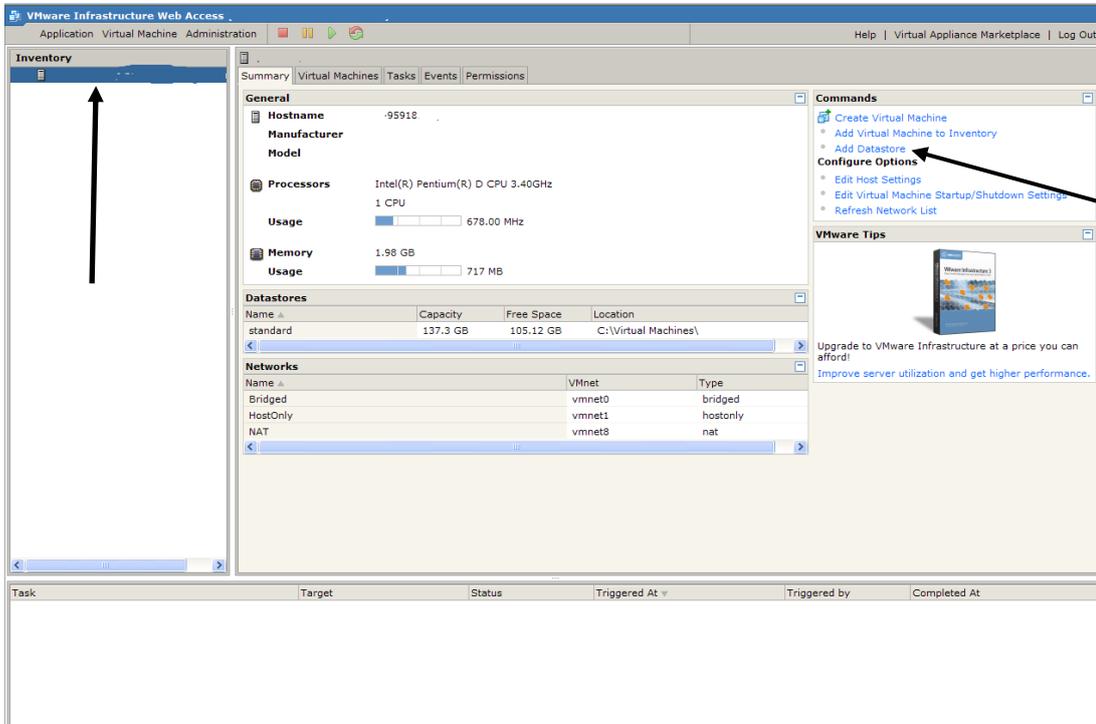
It's a good idea to open your host machine network connections and rename the virtual LAN adapters to something a little more meaningful, as shown below;



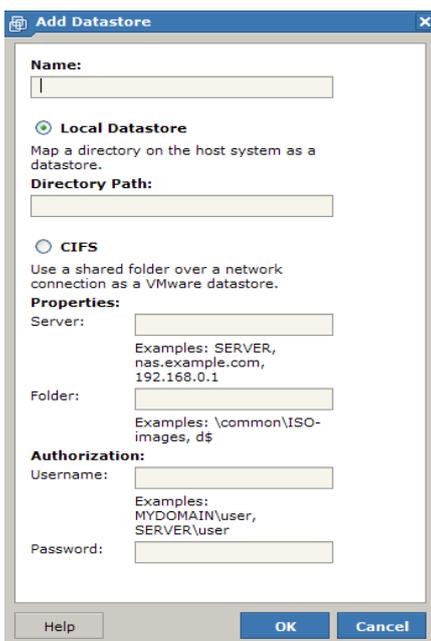
Click the "refresh networks list option" within the VMWare Server console to refresh the virtual networks.

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”;

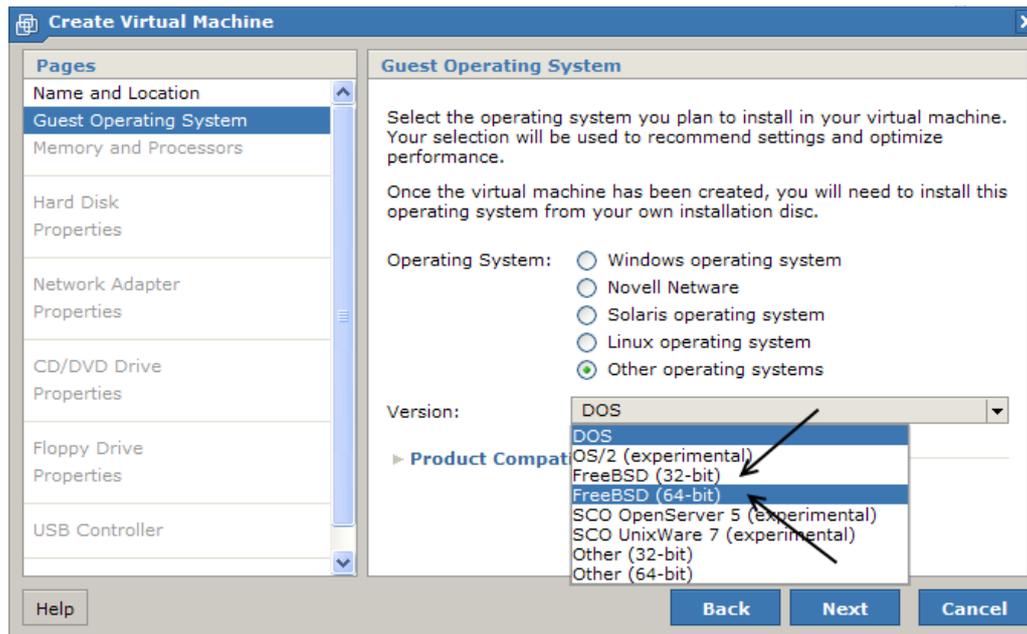


At the dialog supply a datastore name (ISOs) and a local path (C:\ISOs) then click “OK”;



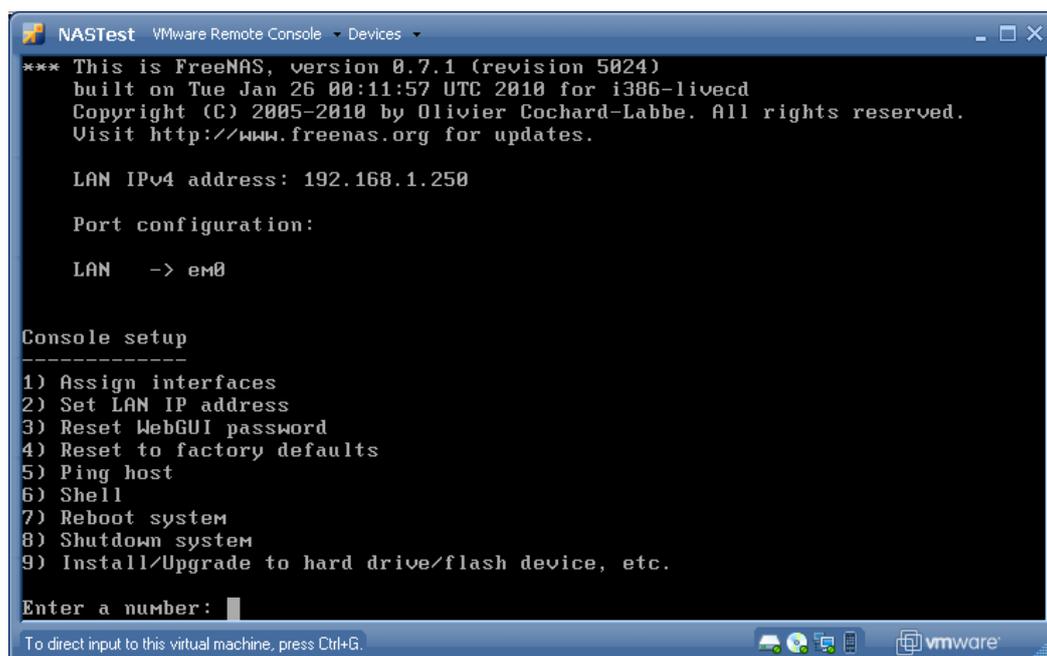
### 3.2 CREATE & CONFIGURE THE SANNAS VM

The first VM we need to create is the NAS server. Create a new VM with the following properties (select either 32 bit or 64 bit as your host system supports);

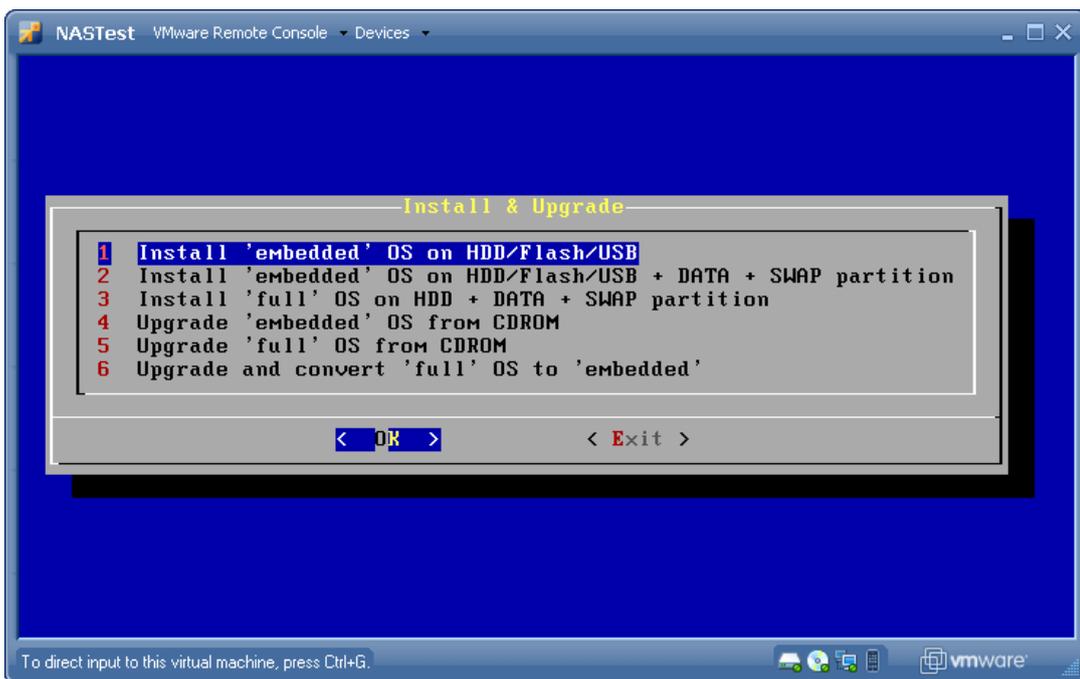


Add a 2GB virtual hard disk and a network adapter. Bind the virtual NIC to whichever virtual switch you want to use for the Public network (VMNET1 in my case). Add a virtual CD/DVD drive and bind to the FreeNAS ISO image. Do not add a USB controller or a Floppy disk drive. The VM will need 400MB of RAM.

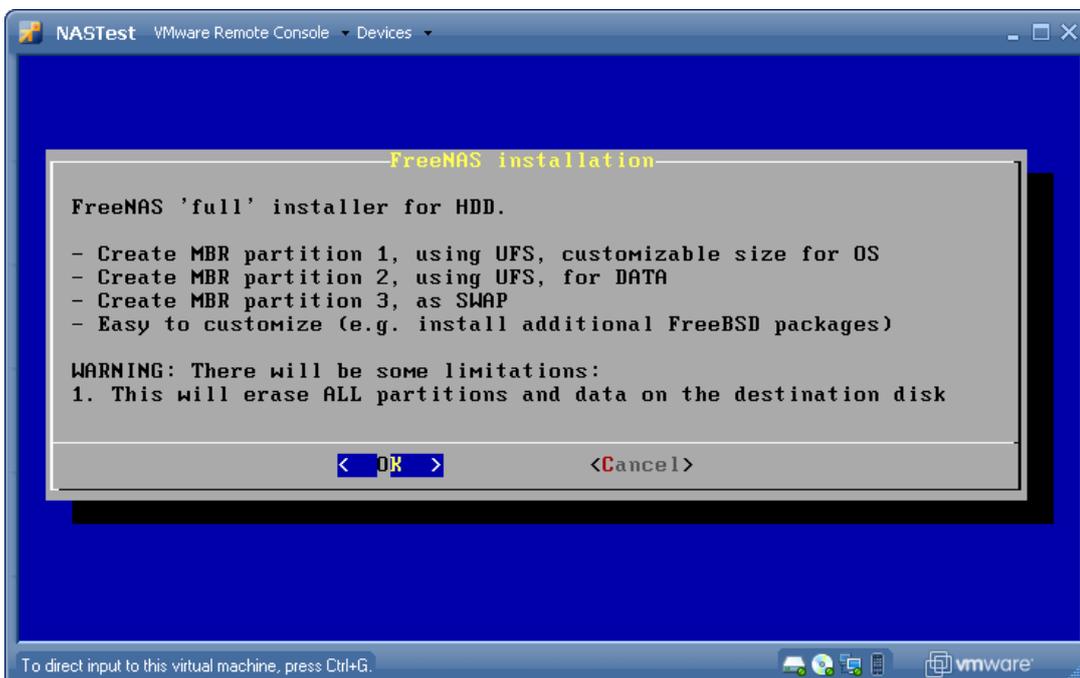
Boot the ISO image accepting all defaults. Once the NOS has booted (shown below), select option 9 from the menu.



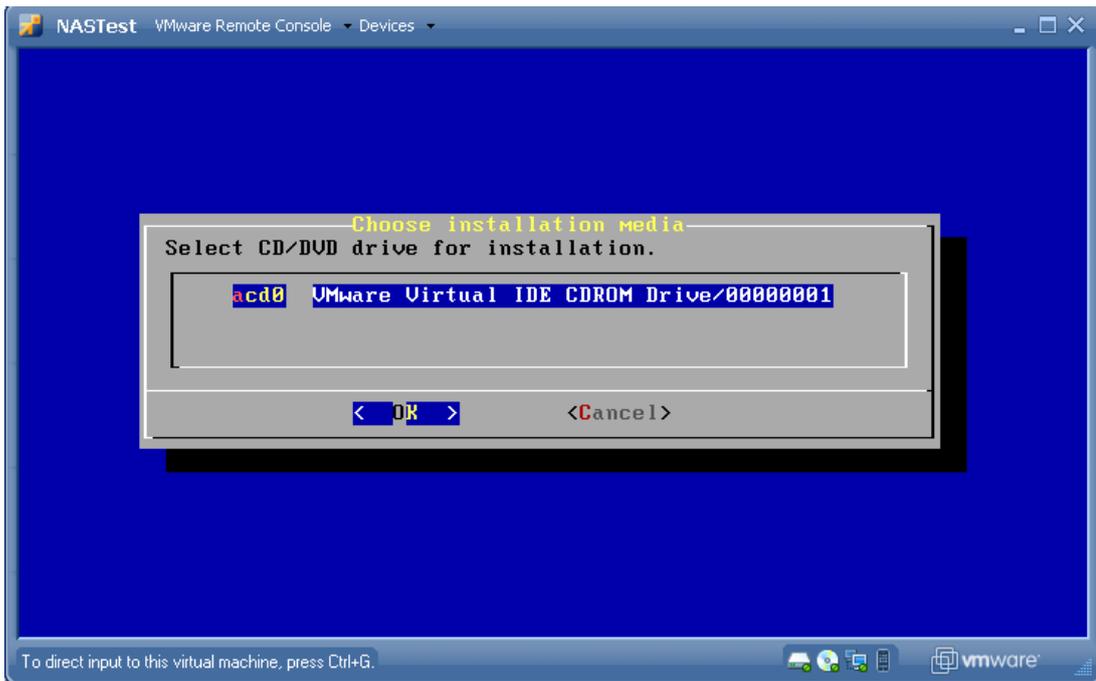
You want to install the Full OS using the 3rd option.



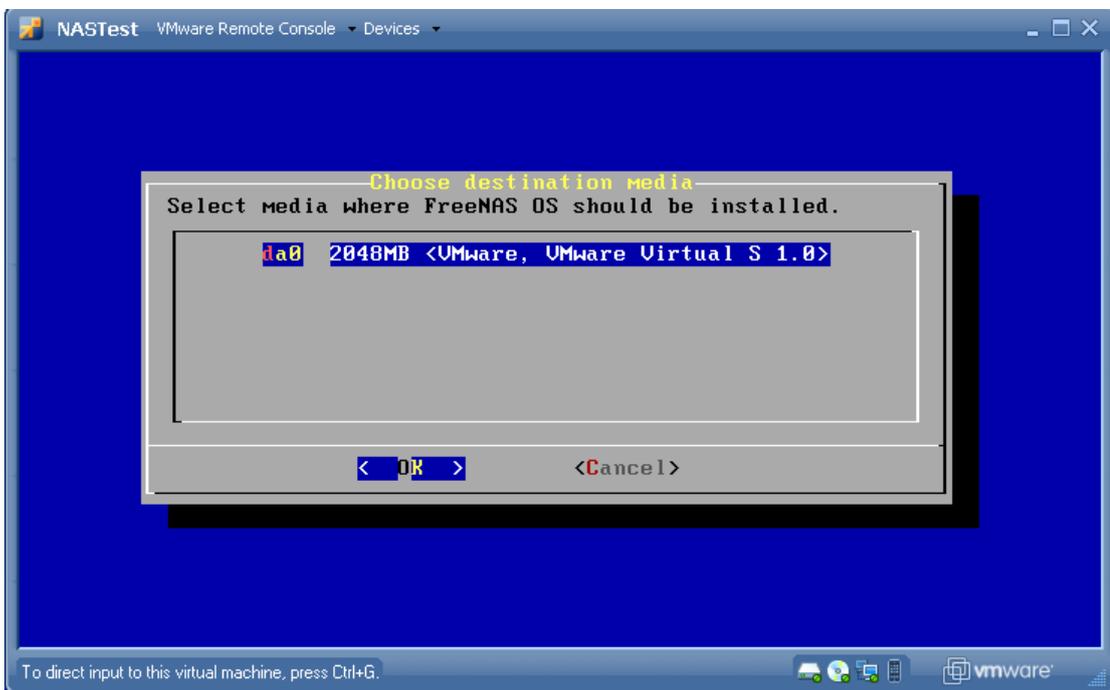
Answer "OK" to continue



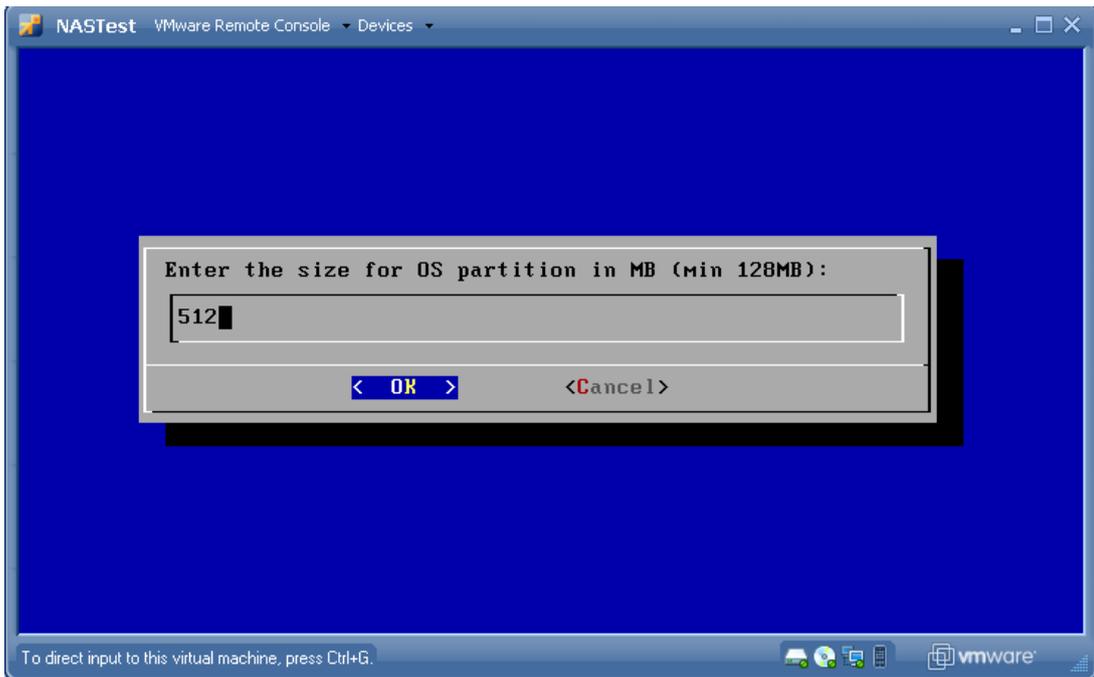
Select "OK" for the CD drive,



Select "OK" for the virtual disk,



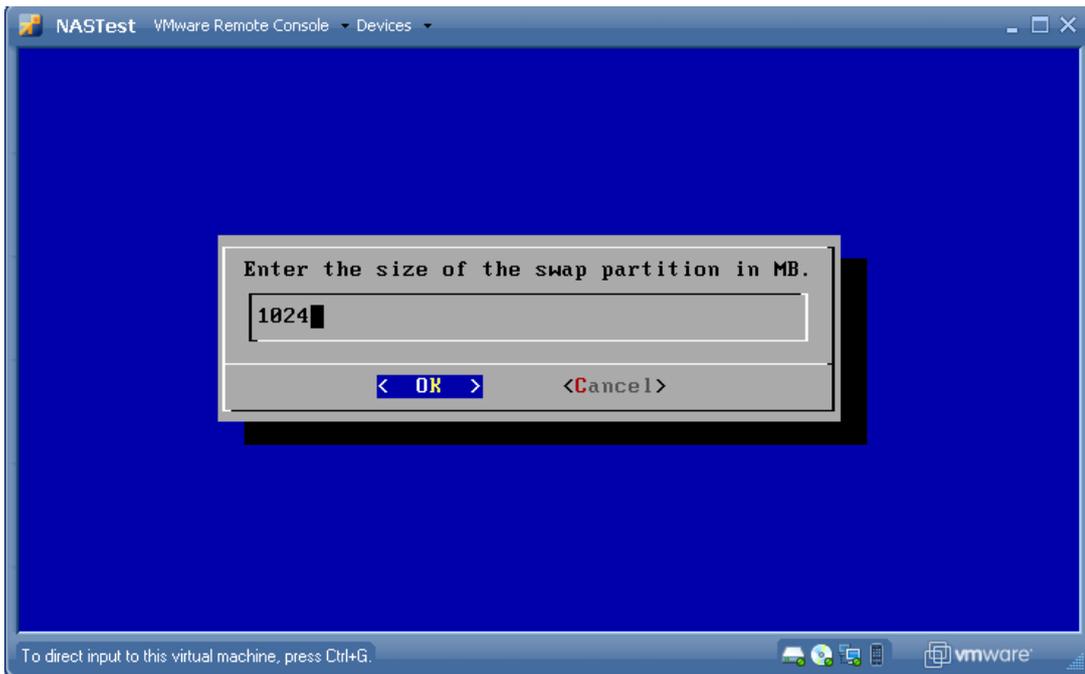
Supply the OS partition size of 512MB,



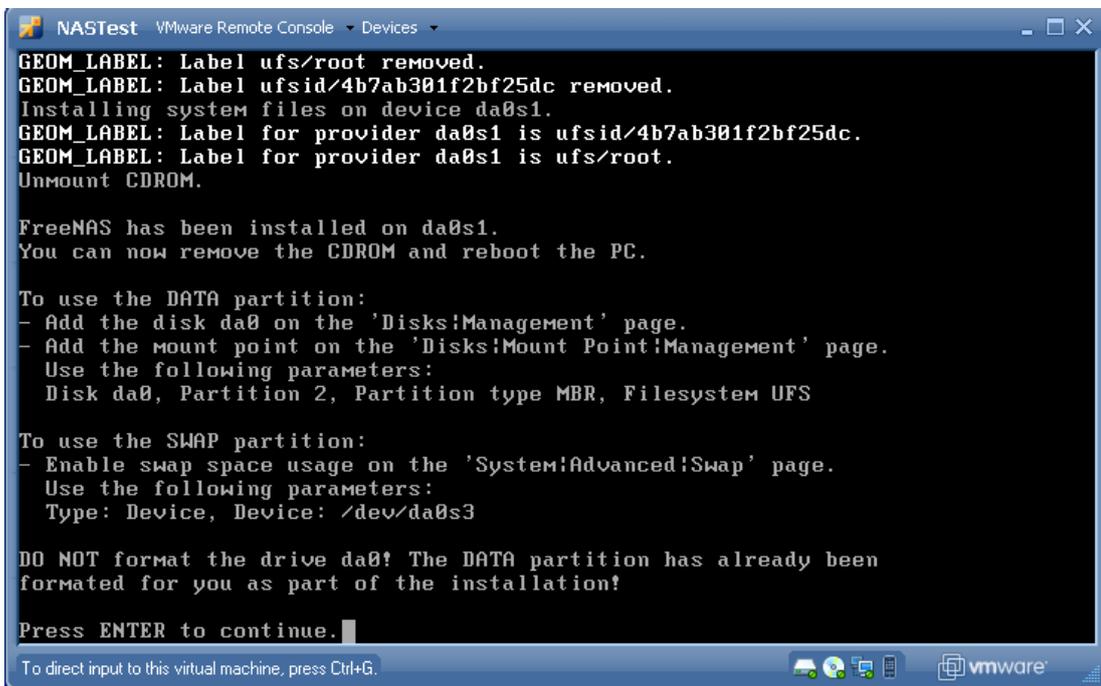
Select "yes" to add a swap partition,



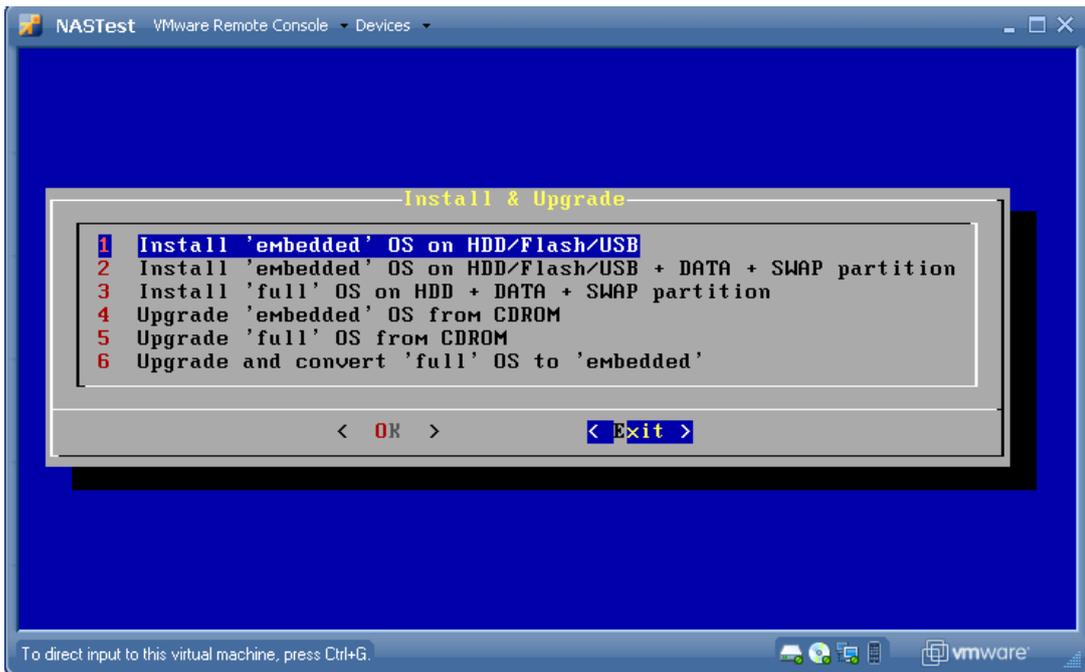
Supply the swap partition size of 1024MB,



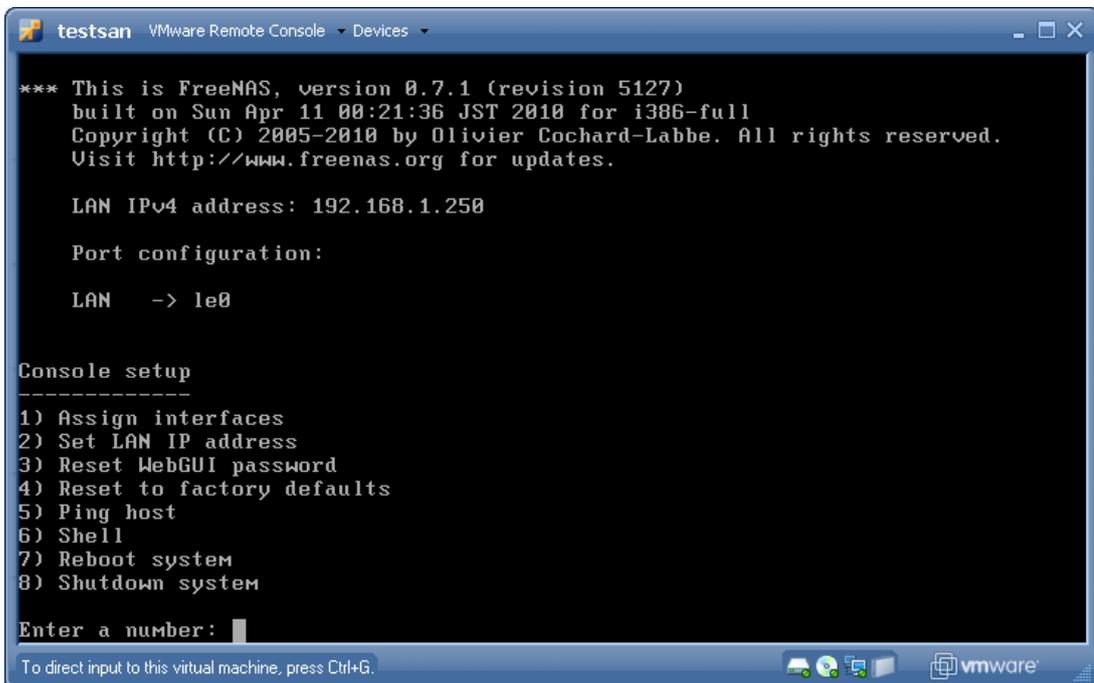
Installation completed, press "Enter" to continue.



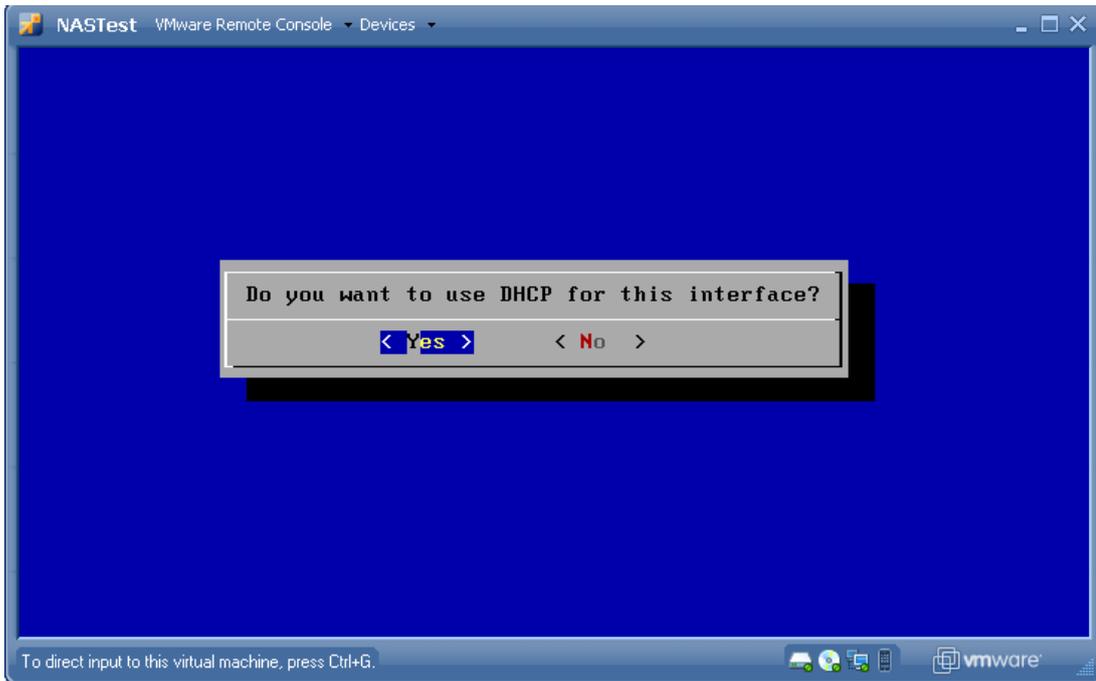
Now select "Exit" to finish.



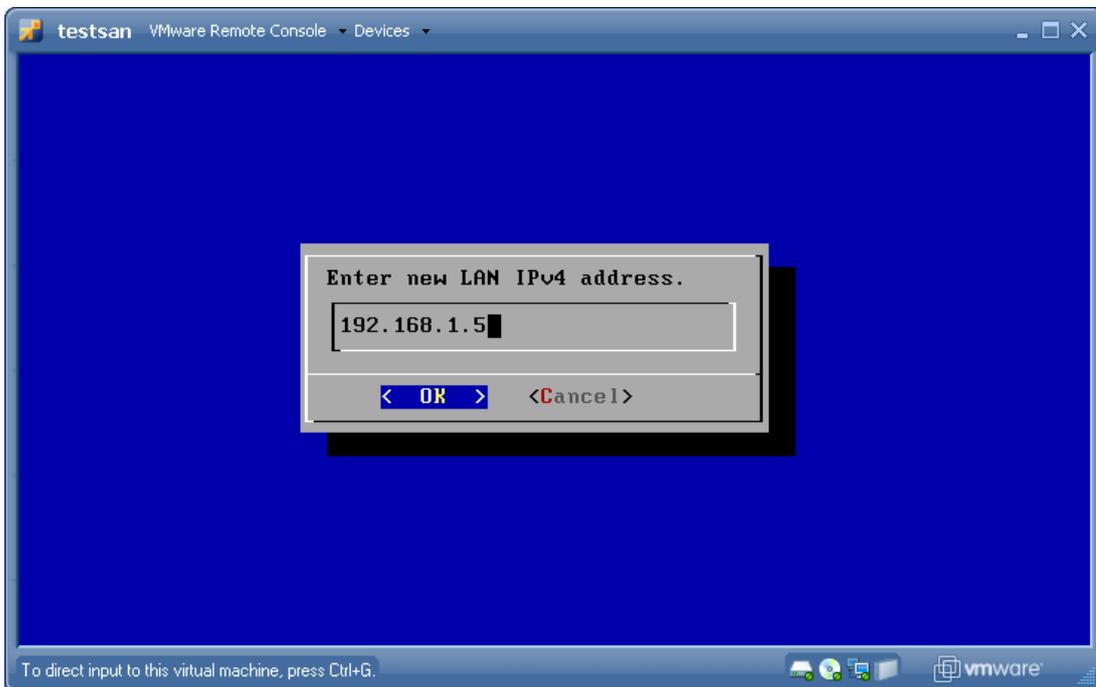
Now select option 7 and reboot the SAN! Once this has been done we need to set the LAN IP address. Select option 2 from the menu



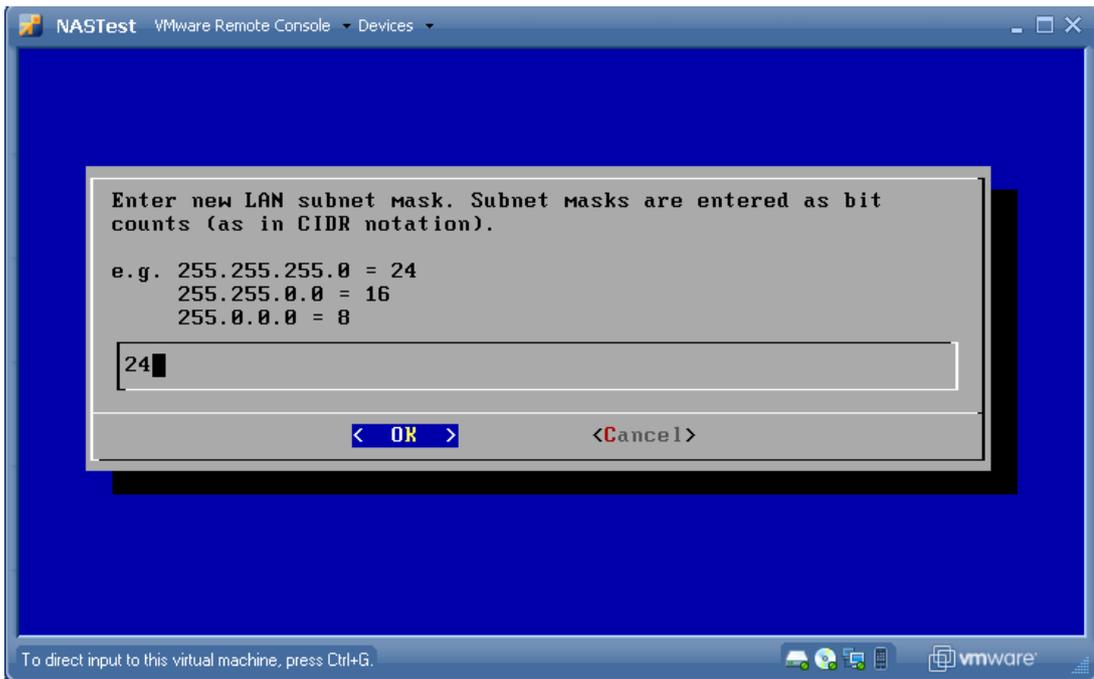
The LAN configuration starts as follows. Select whether to use DHCP (no for our scenario)



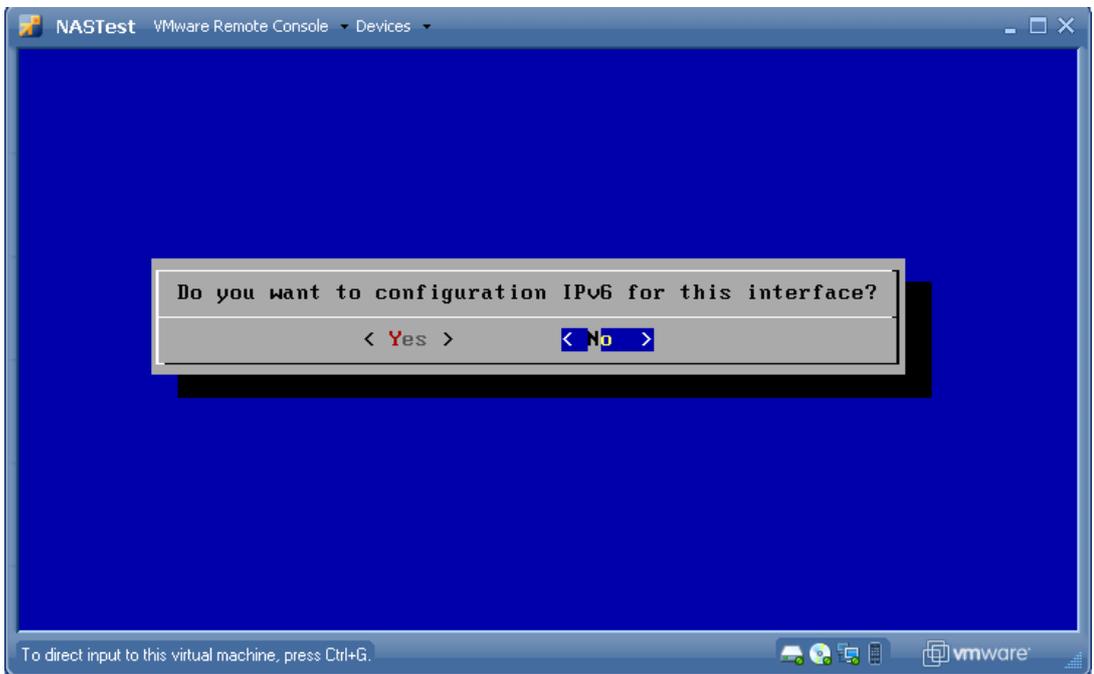
Enter the new LAN IP address



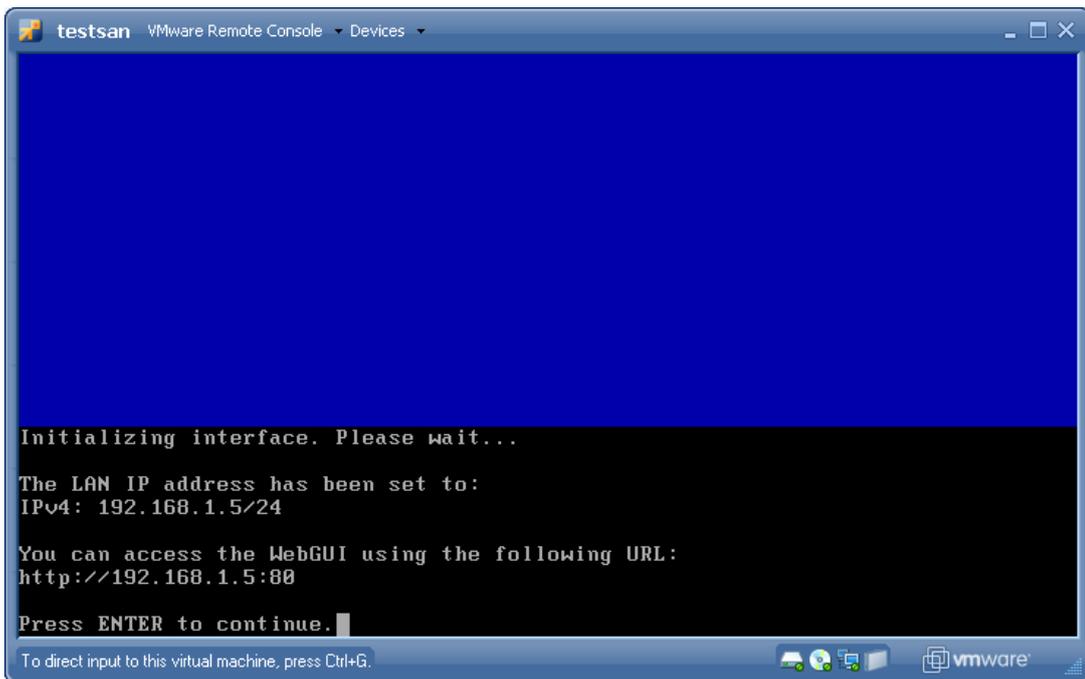
Supply the Network mask (using CIDR notation),



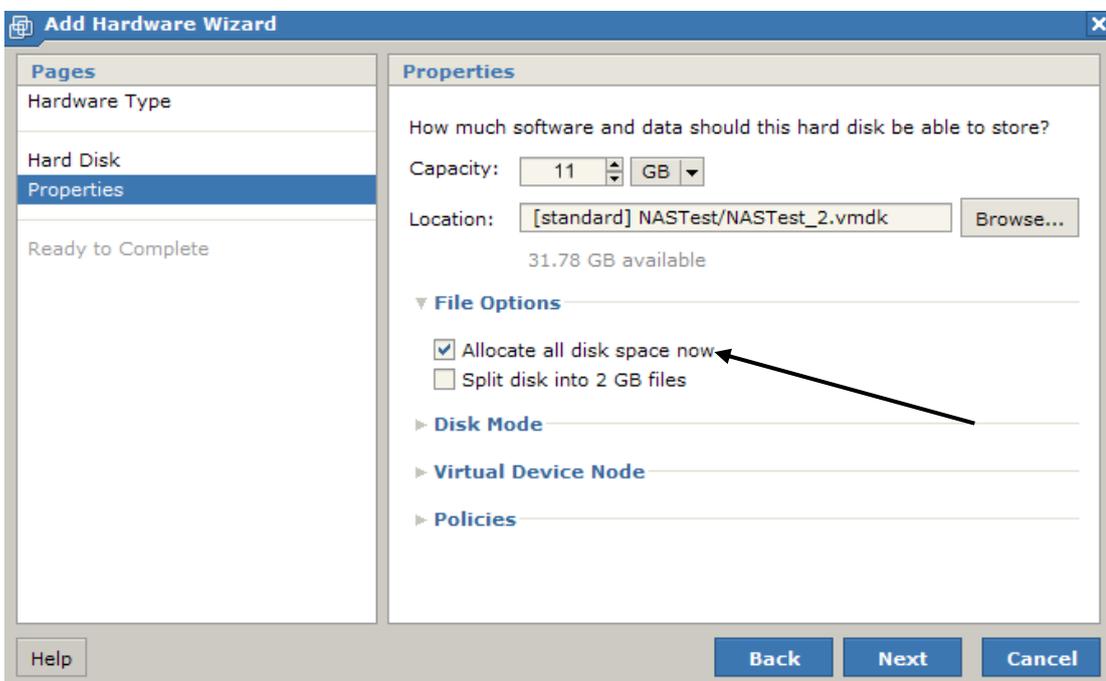
Select "no" for the IPv6 configuration,



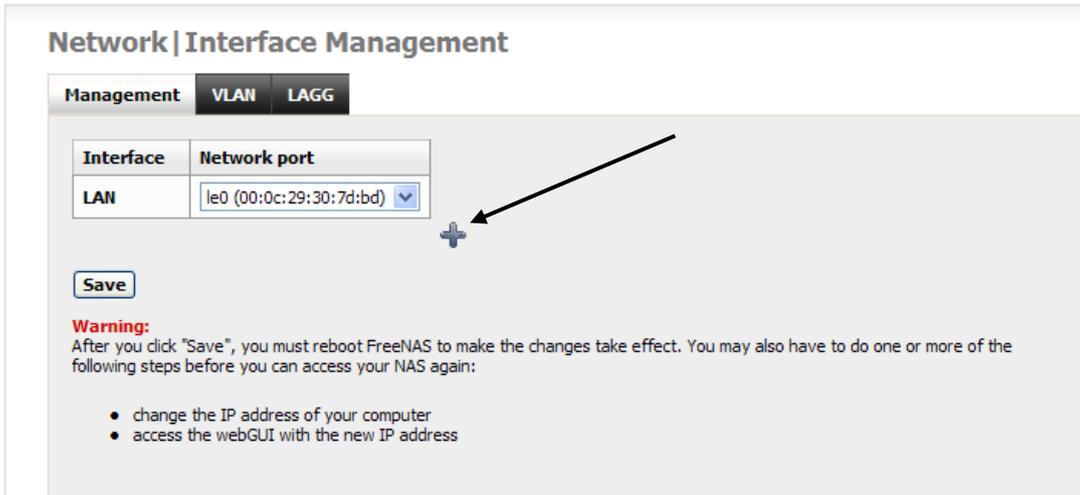
Press "Enter" to continue,



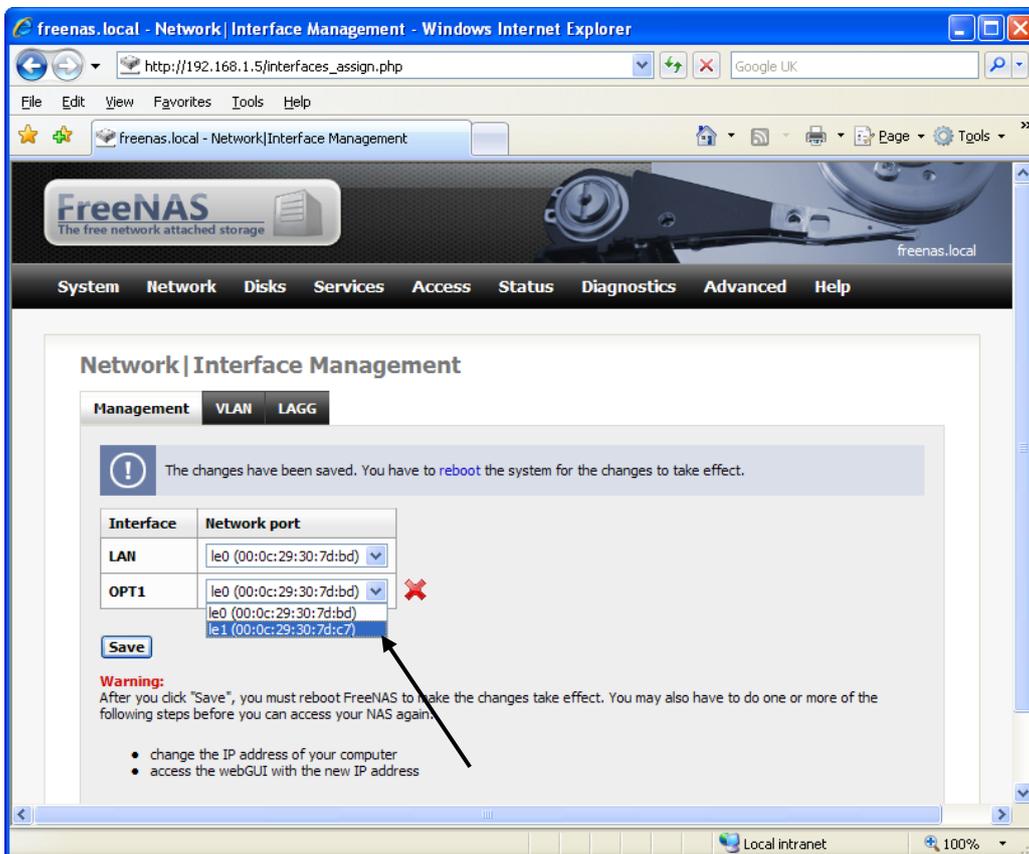
Shutdown the server using option 8 from the menu as we now have to create\attach the virtual hard disk to the NAS VM that will hold our iSCSI LUNs. Click the "Add hardware" option and add a new disk 11GB in size, pre allocating all disk space (shown below).



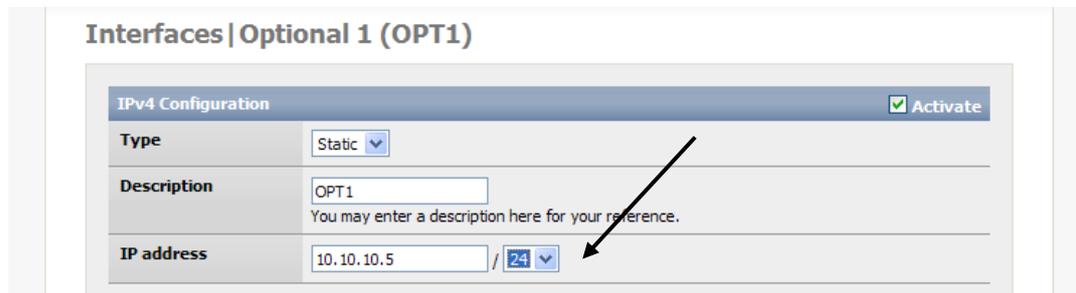
At this point, you now need to complete the configuration for the networks the SAN will be using. Shutdown the SAN VM and add a second vNIC binding to the virtual switch you will use for iSCSI traffic (VMNET2 in my case). Start the VM and login to the console, from the main menu select 'Network' > 'Interface management'. You will only see the 1 vNIC. Click the 'Add Interface' icon shown below;



Select the second vNIC to bind to the iSCSI network (OPT1 interface), as shown below;



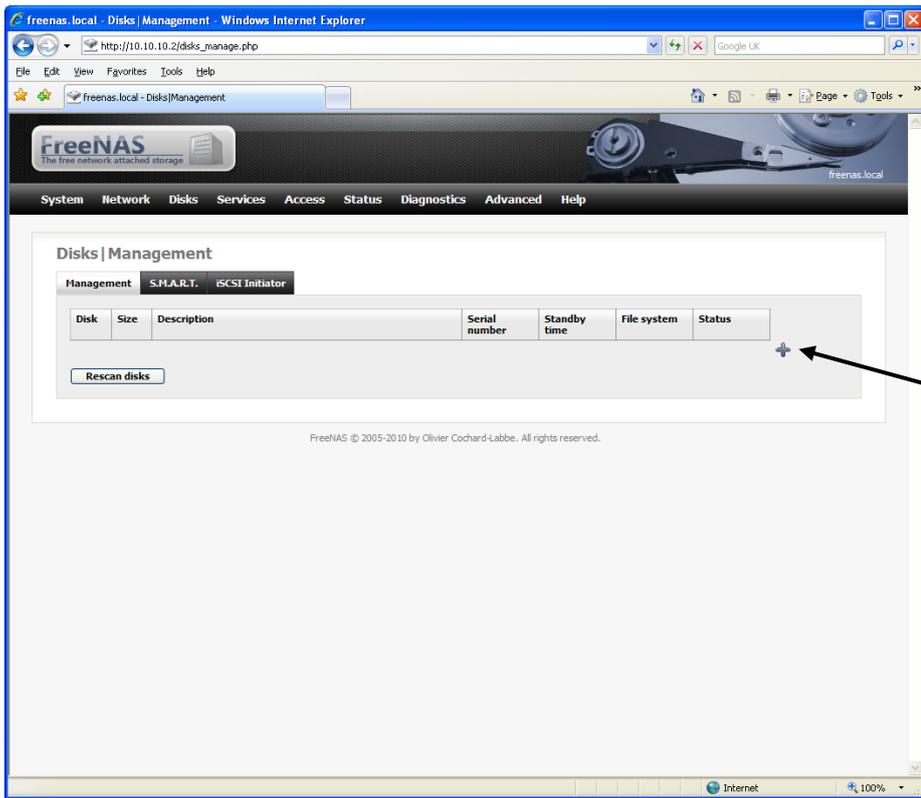
Click 'save' and then reboot the system! This will create a new network interface named 'OPT1'. Once you have saved and rebooted, login to the console and go to 'Network' > 'OPT1'. Configure the iSCSI network information as shown below, (this IP address will later become the iSCSI portal address). Save and reboot again!



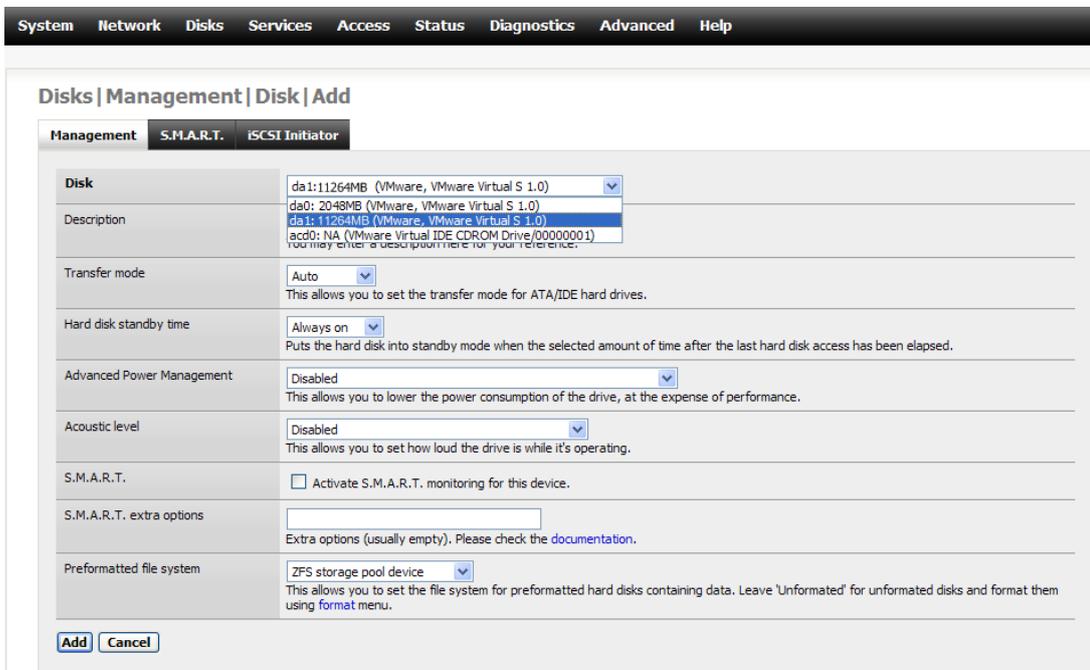
Once the disk has been added and the OPT1 interface configured, boot the NAS VM and login to the NAS web management page using the username "admin" and password "FreeNAS" (without quotes). You may have to bypass the proxy for local addresses in your IE settings!



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



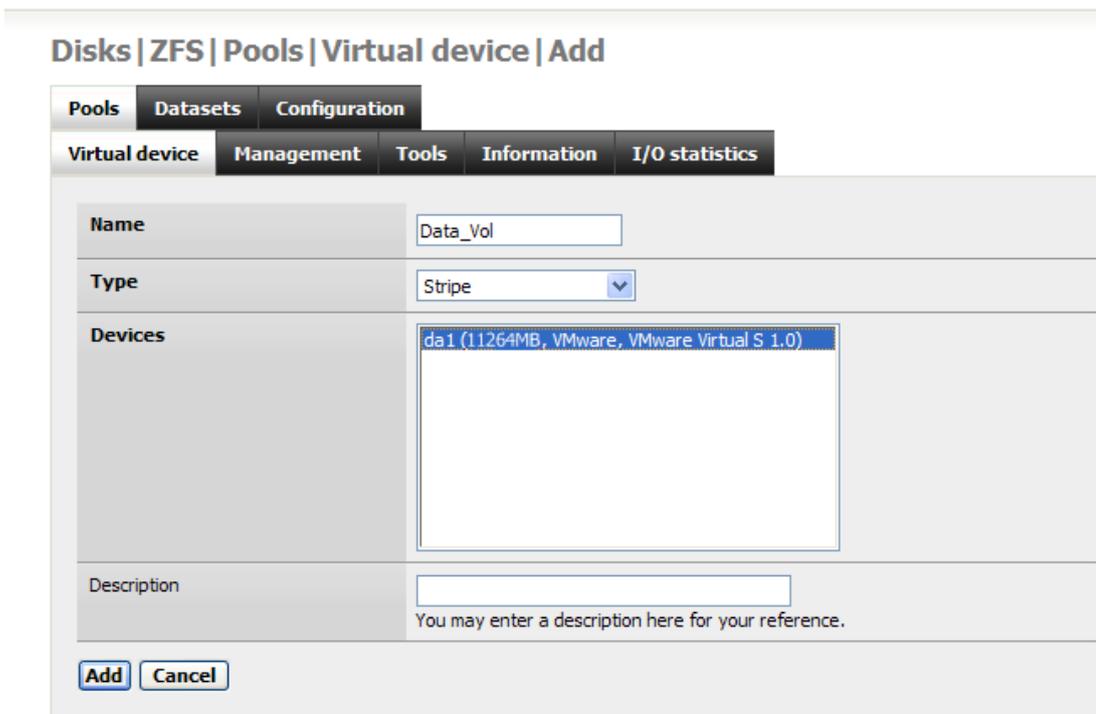
Select the new 11GB disk and use ZFS pre formatted file system, then click “add”. After adding the disk you must click the “Apply changes” button.



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.



Now click "Management" as shown below and then click the + sign



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

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”.

The screenshot shows the 'Services | iSCSI Target' configuration page. At the top, there are tabs for 'Settings', 'Targets', 'Portals', 'Initiators', 'Auths', and 'Media'. The 'Settings' tab is active. The 'iSCSI Target' section has an 'Enable' checkbox checked. Below this, there are fields for 'Base Name' (iqn.2007-09.jp.ne.peach.istgt), 'Discovery Auth Method' (Auto), and 'Discovery Auth Group' (None). An 'Advanced settings' section contains fields for 'I/O Timeout' (30), 'NOPIN Interval' (20), 'Max. sessions' (32), 'Max. connections' (8), 'FirstBurstLength' (65536), 'MaxBurstLength' (262144), and 'MaxRecvDataSegmentLength' (262144). At the bottom, there is an 'iSCSI Target Logical Unit Controller' section with an 'Enable' checkbox checked and a 'Save and Restart' button.

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

The screenshot shows the 'Services | iSCSI Target | Initiator Group' configuration page. At the top, there are tabs for 'Settings', 'Targets', 'Portals', 'Initiators', 'Auths', and 'Media'. The 'Initiators' tab is active. The 'Initiator Groups' section contains a table with columns 'Tag', 'Initiators', and 'Networks'. Below the table, there is a plus sign (+) and a description: 'A Initiator Group contains authorised initiator names and networks to access the target.'

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

The screenshot shows the 'Initiator Group' configuration page. At the top, there are tabs for 'Settings', 'Targets', 'Portals', 'Initiators', 'Auths', and 'Media'. The 'Initiators' tab is selected. Below the tabs, there are several fields: 'Tag number' with a value of '1', 'Initiators' with a dropdown menu set to 'ALL', 'Authorised network' with a dropdown menu set to '10.10.10.0/24', and a 'Comment' field. At the bottom, there are 'Add' and 'Cancel' buttons.

Click “Portal group” and then click the + sign

The screenshot shows the 'Portal Group' configuration page. At the top, there are tabs for 'Settings', 'Targets', 'Portals', 'Initiators', 'Auths', and 'Media'. The 'Portals' tab is selected. Below the tabs, there is a 'Portal Groups' section with a table. The table has two columns: 'Tag' and 'Portals'. Below the table, there is a description: 'A Portal Group contains IP addresses and listening TCP ports to connect the target from the initiator.' To the right of the table, there is a blue plus sign (+) button, which is pointed to by a black arrow.

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	512 MIB	
	SQLBak	/mnt/Data_Vol/SQLBak	1024MIB	
	SQLData	/mnt/Data_Vol/SQLData	2048MIB	
	SQLLog	/mnt/Data_Vol/SQLLog	1024MIB	

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 'qn.'.

**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	512MiB
SQLBak	/mnt/Data_Vol/SQLBak	1024MiB
SQLData	/mnt/Data_Vol/SQLData	2048MiB
SQLLog	/mnt/Data_Vol/SQLLog	1024MiB

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

Target

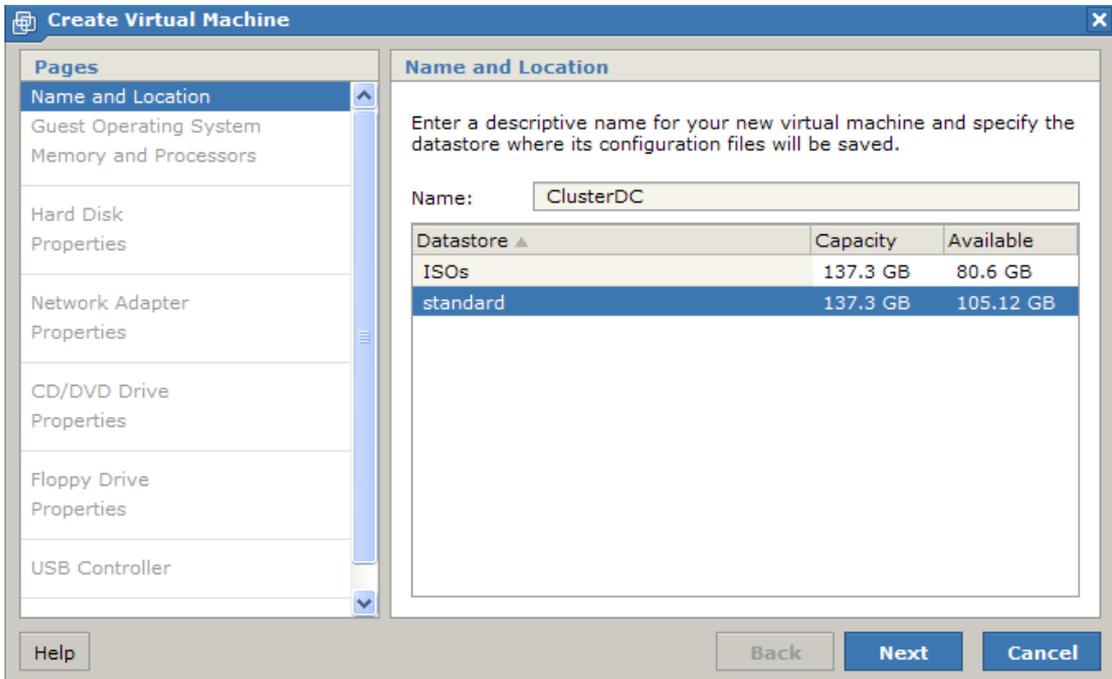
Name	Flags	LUNs	PG	IG	AG
iqn.2007-09.jp.ne.peach.istgt:Quorum	rw	LUN0=/mnt/Data_Vol/Quorum	1	1	none
iqn.2007-09.jp.ne.peach.istgt:SQLBak	rw	LUN0=/mnt/Data_Vol/SQLBak	1	1	none
iqn.2007-09.jp.ne.peach.istgt:SQLData	rw	LUN0=/mnt/Data_Vol/SQLData	1	1	none
iqn.2007-09.jp.ne.peach.istgt:SQLLog	rw	LUN0=/mnt/Data_Vol/SQLLog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

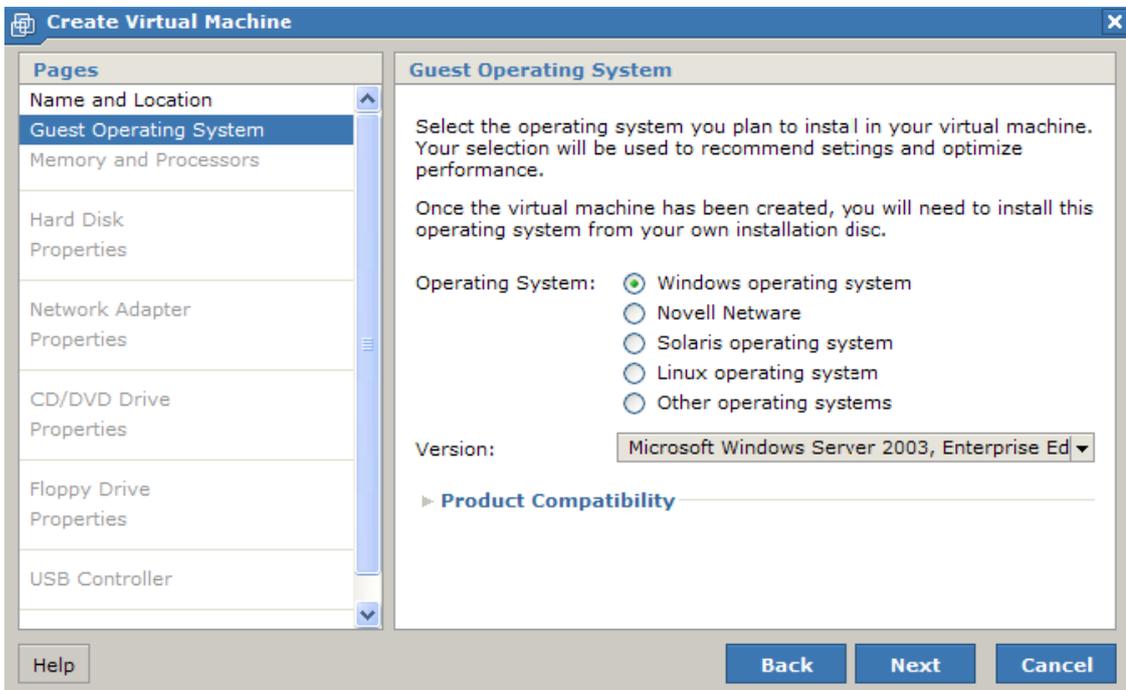
**Note:**  
 To configure the target, you must add at least Portal Group and Initiator Group and Extent.  
 Portal Group which is identified by tag number defines IP addresses and listening TCP ports.  
 Initiator Group which is identified by tag number defines authorised initiator names and networks.  
 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)

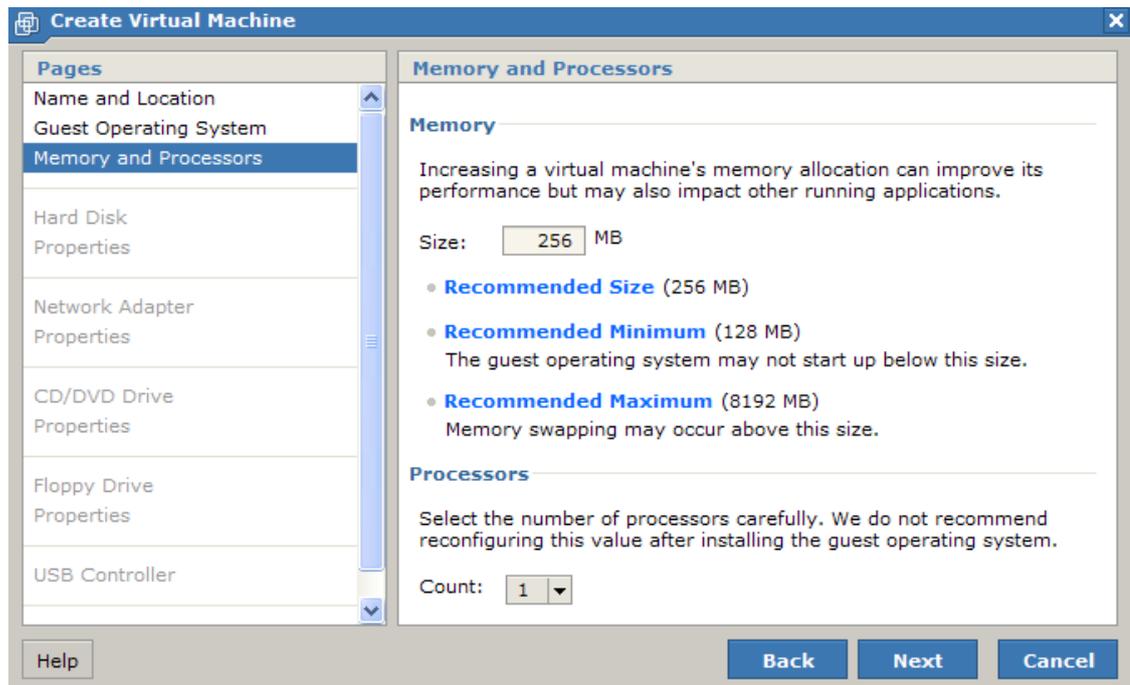
Now we are ready to start creating the VM's. Under the "Commands" section, select "Create virtual machine". Enter a Virtual machine name and select a datastore (this datastore maps to a folder on your local drive) then click "Next";



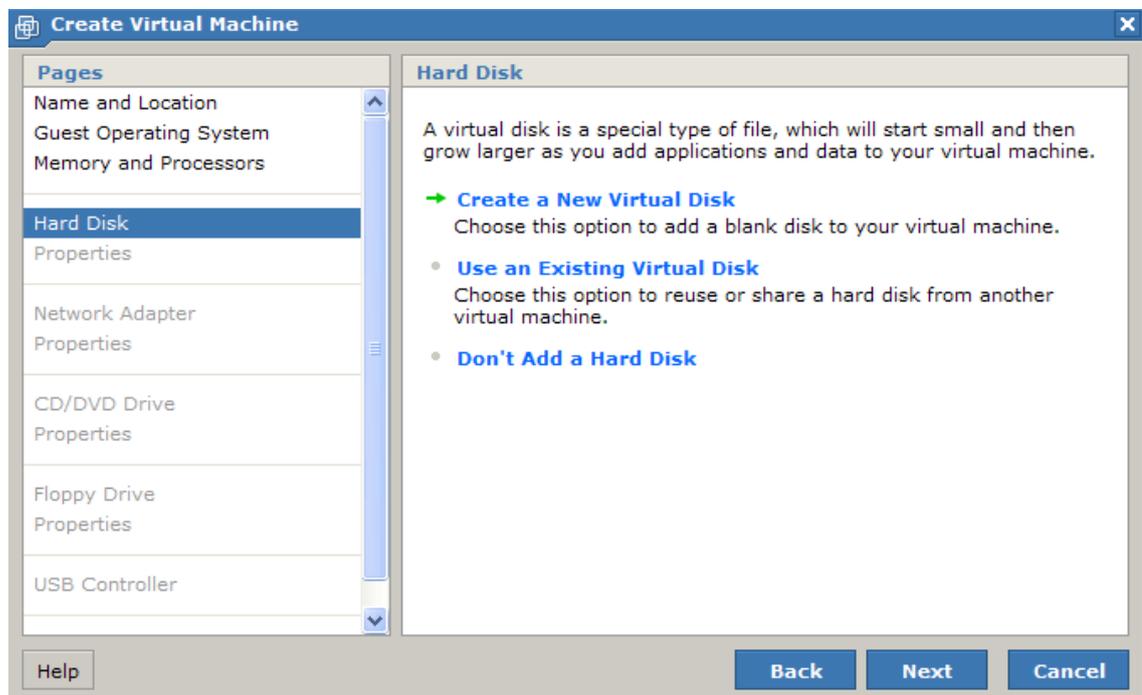
Select the Operating System type (Enterprise Edition for clustering) and click "Next";



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

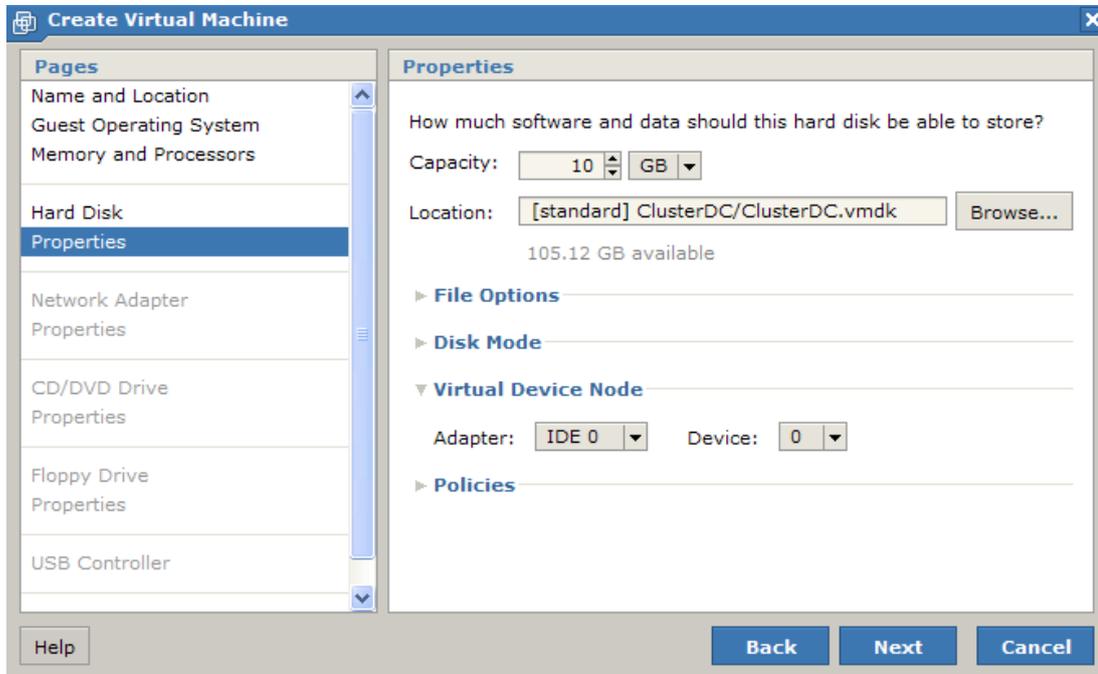


Select to create a new virtual disk;

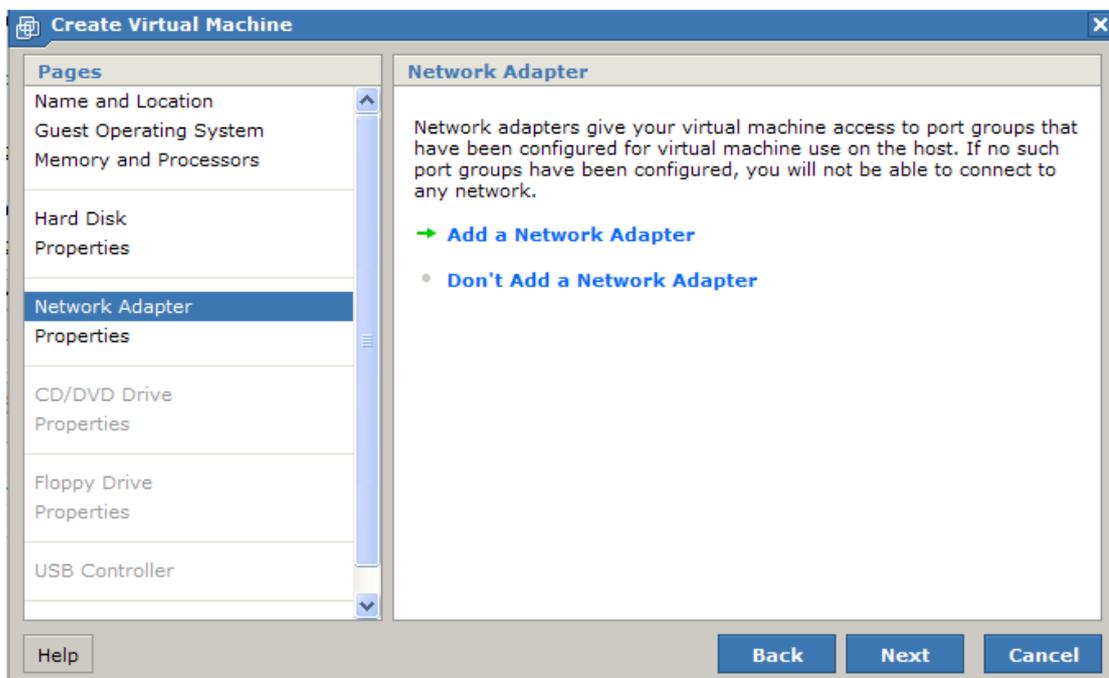


Supply a size for the virtual disk and select SCSI bus (SCSI ID 0 for the VM boot drives in our example), make sure to pre allocate the disk space for the virtual disks. Click “Next”;

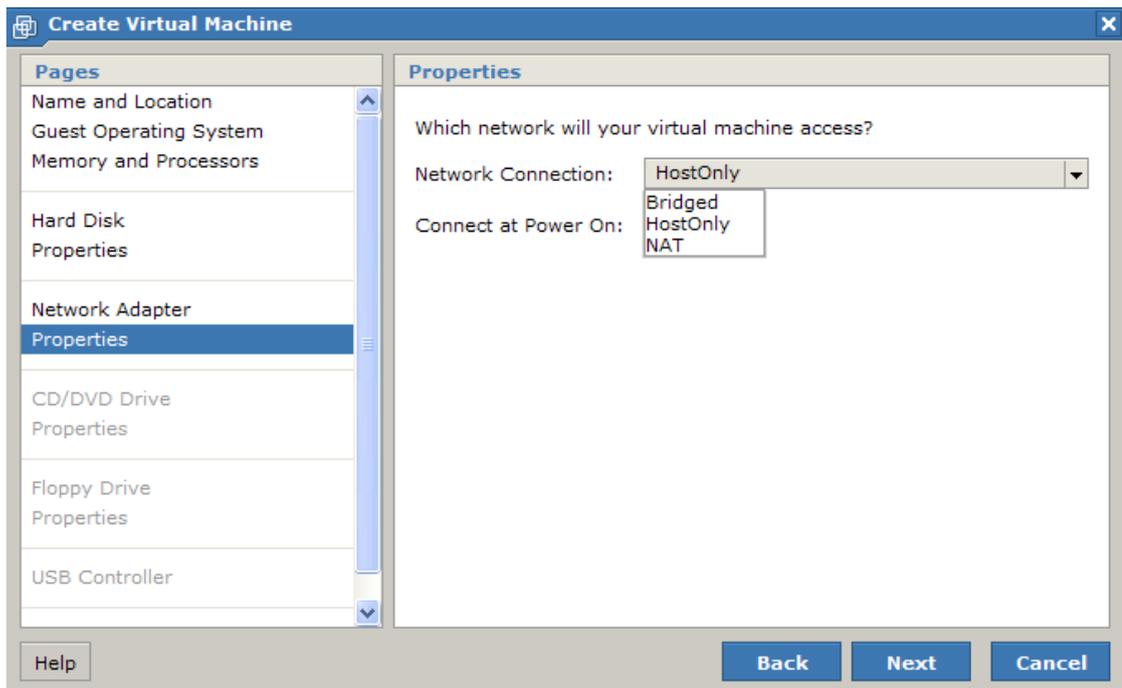
*The Windows 2003 cluster nodes need a min 15GB disk drive and 10GB for the DC.*



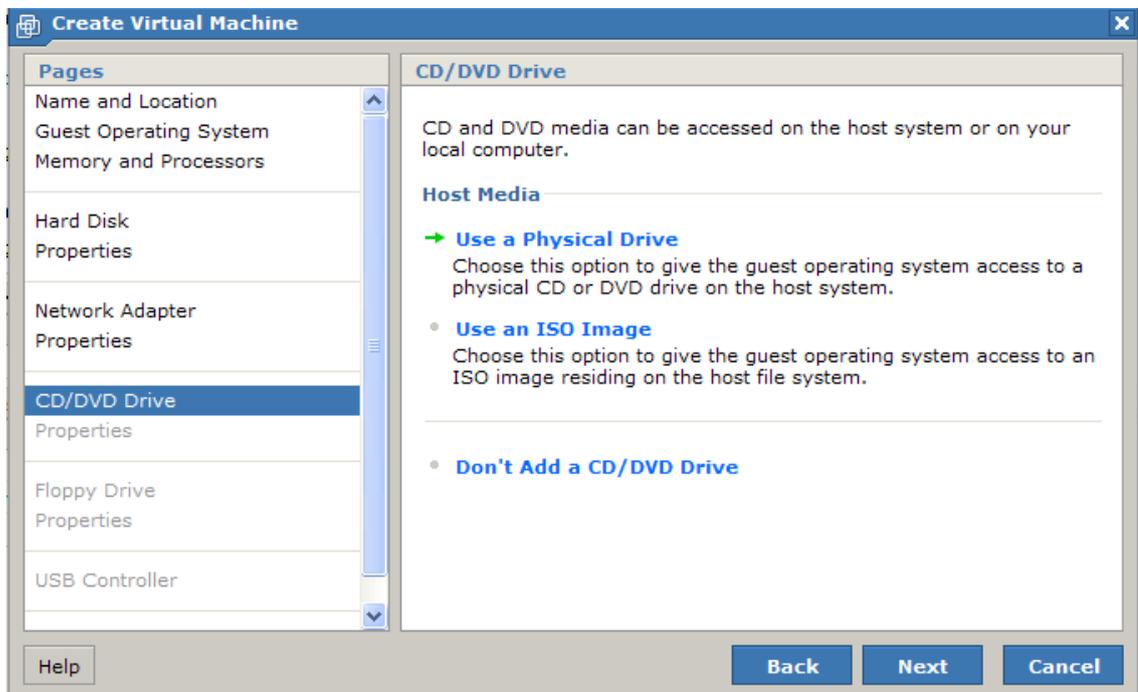
Select “Add a network adapter” and the network selection browse opens



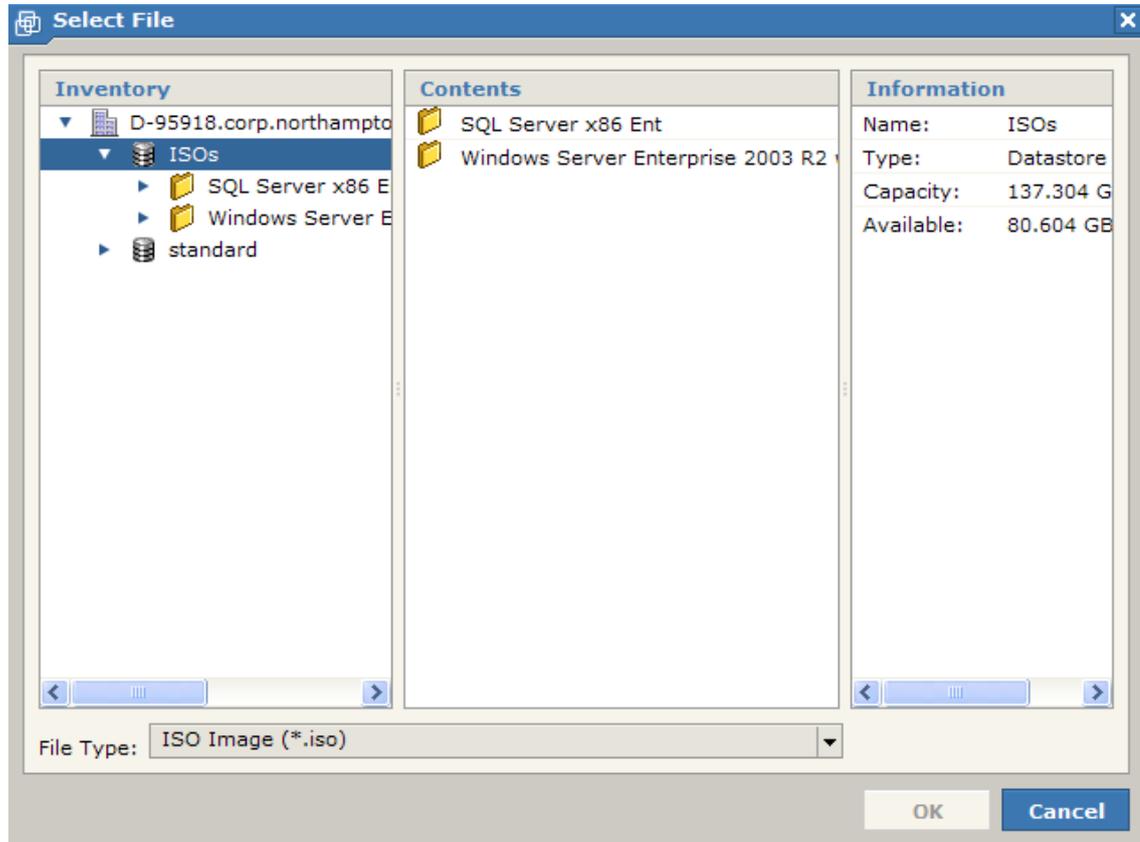
Select "Host only" for this vNIC and click "Next";



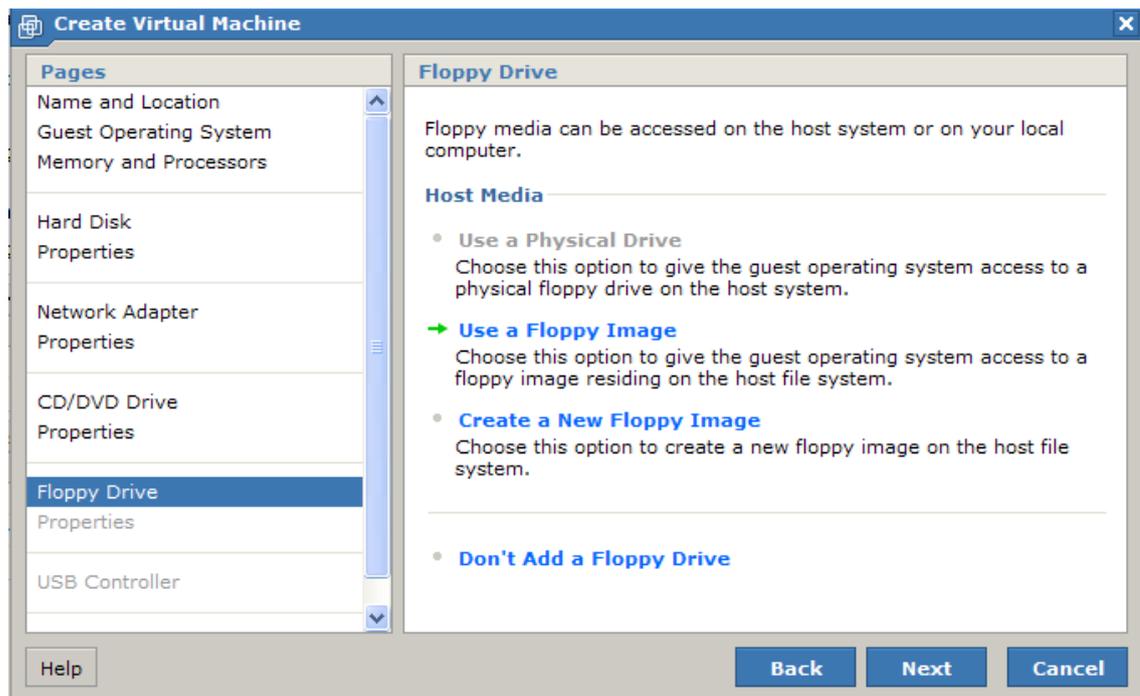
Select to use an ISO image for the vCD drive and click the browse button,



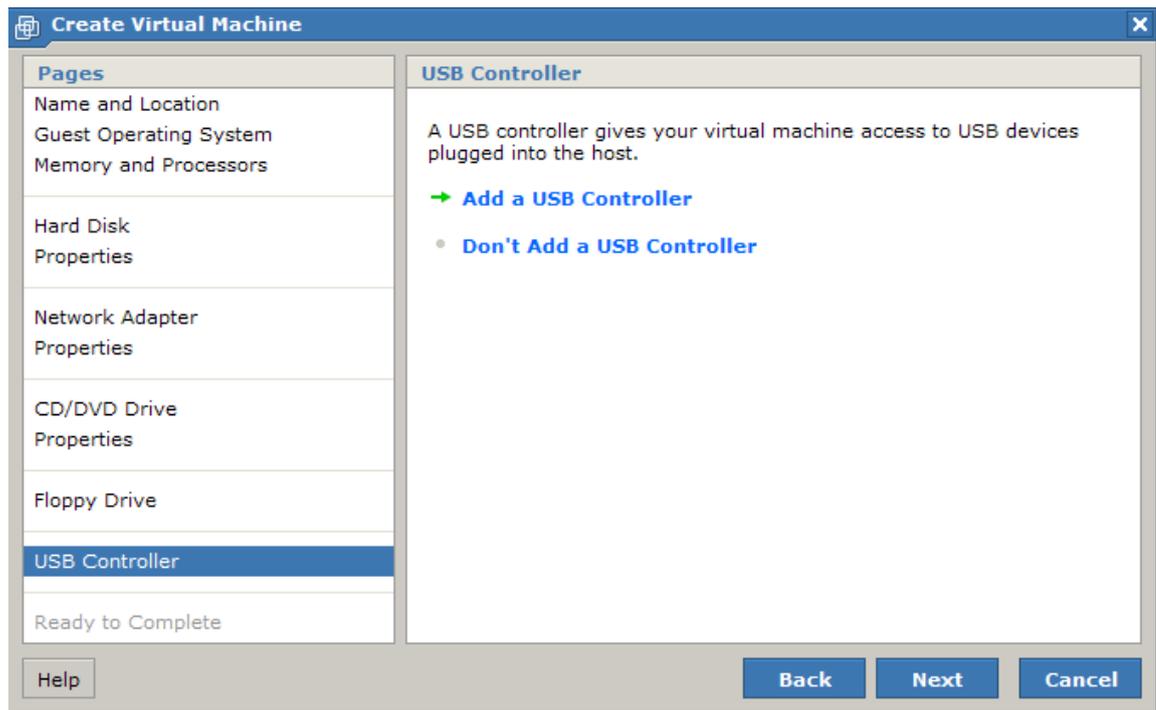
The datastore browser opens, drill down and select the Windows 2003 R2 Enterprise ISO (or Windows 2003 Enterprise for a cluster node) and click “OK”. At the “create virtual machine window” click “Next”;



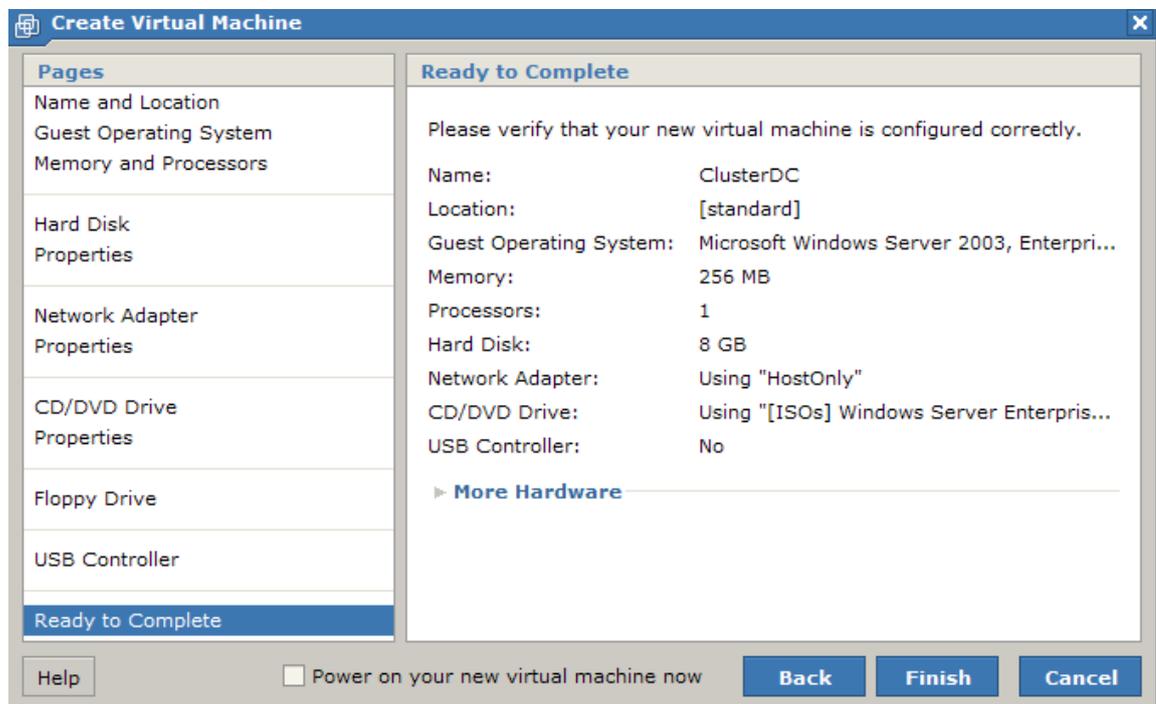
At the next screen do not add a floppy drive and click “Next”;



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

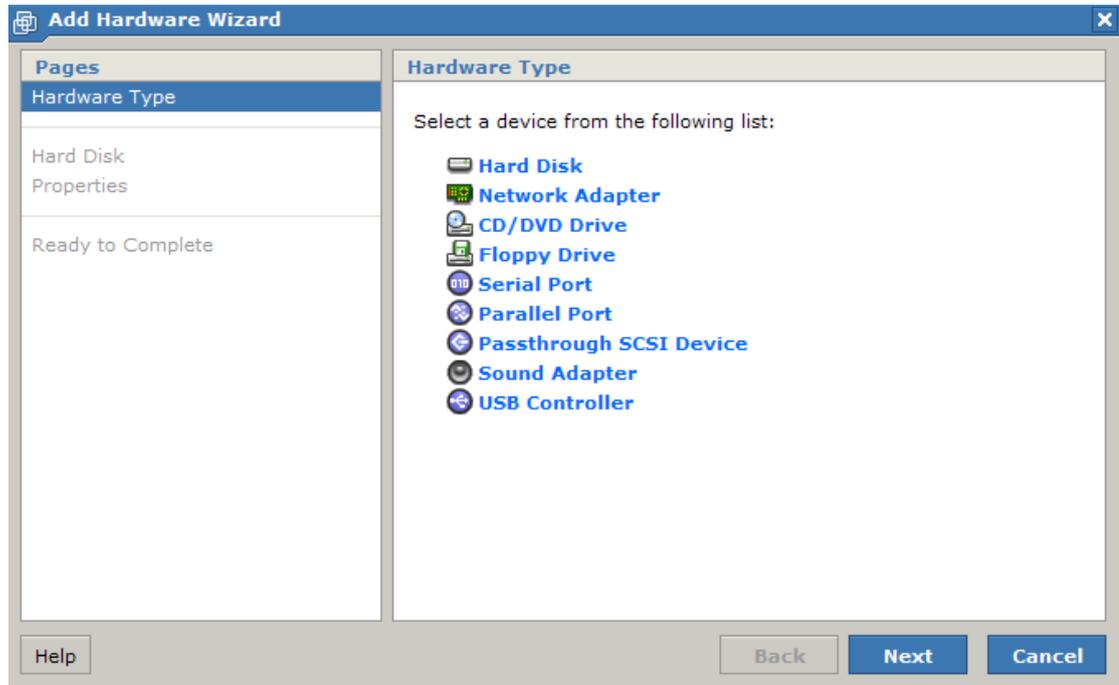


At the last screen click "Finish" to complete the VM

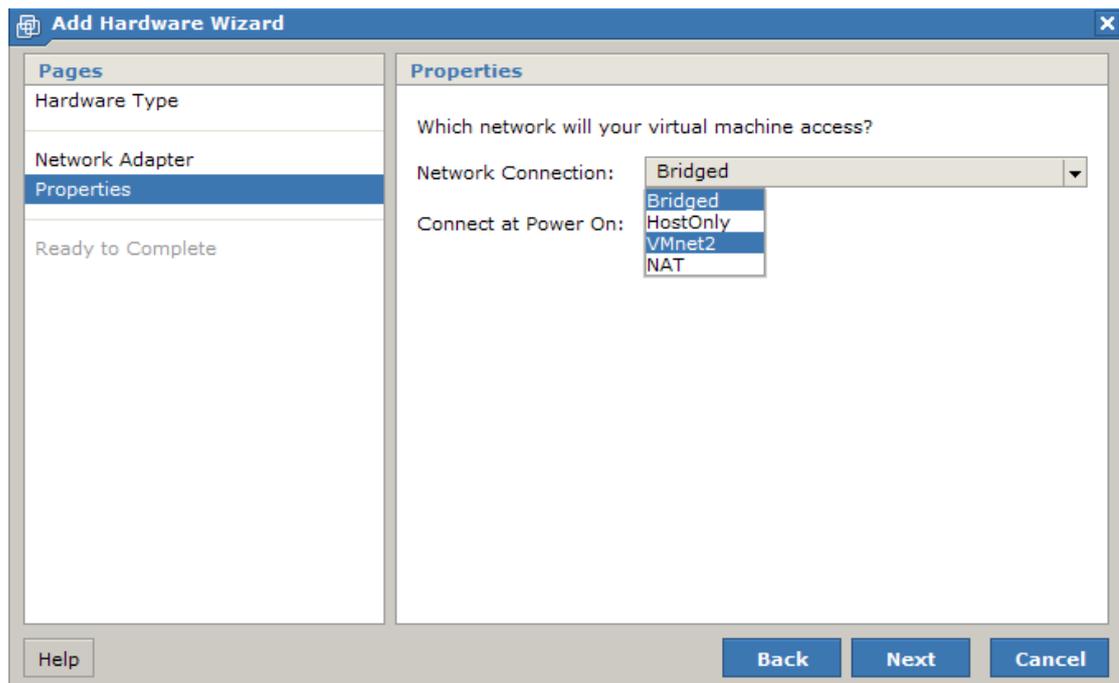


Create the remaining VMs (node 1 and node 2) using Windows 2003 Enterprise software. Use a sparse (not pre allocated) 15 - 20GB virtual disk for each node, also add 2 more vNICs, these will be used for the Heartbeat and iSCSI networks. This is done as follows;

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.



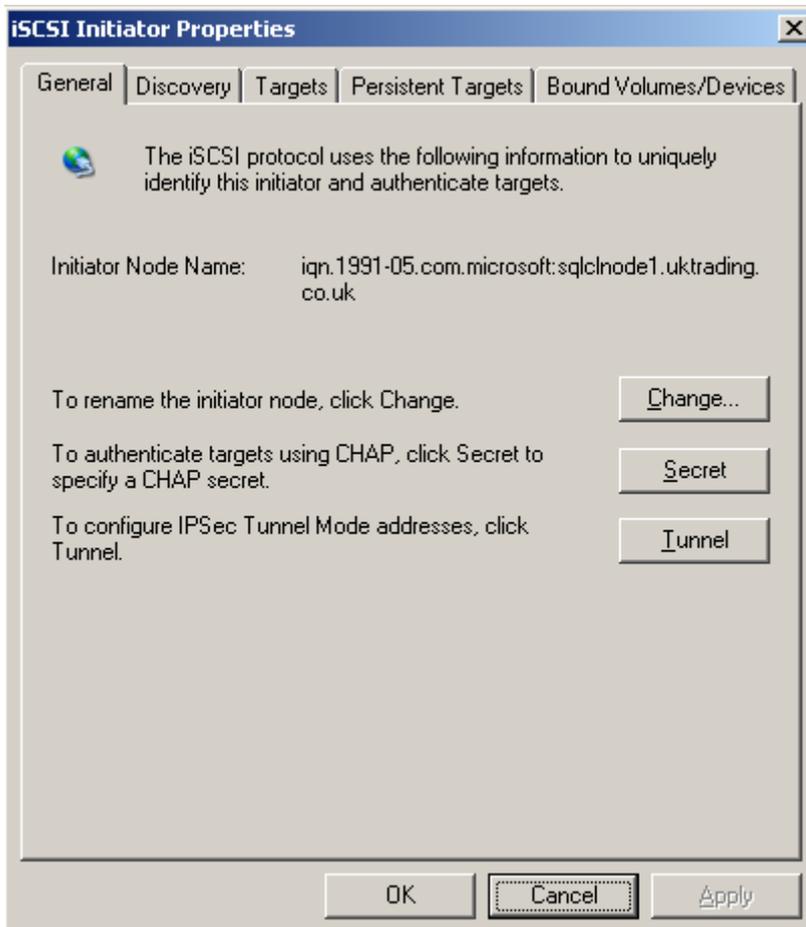
Now complete the following;

- 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.
- Mount the Windows 2003 server iSCSI initiator software installer and install the iSCSI extensions, this will require a reboot.

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.*

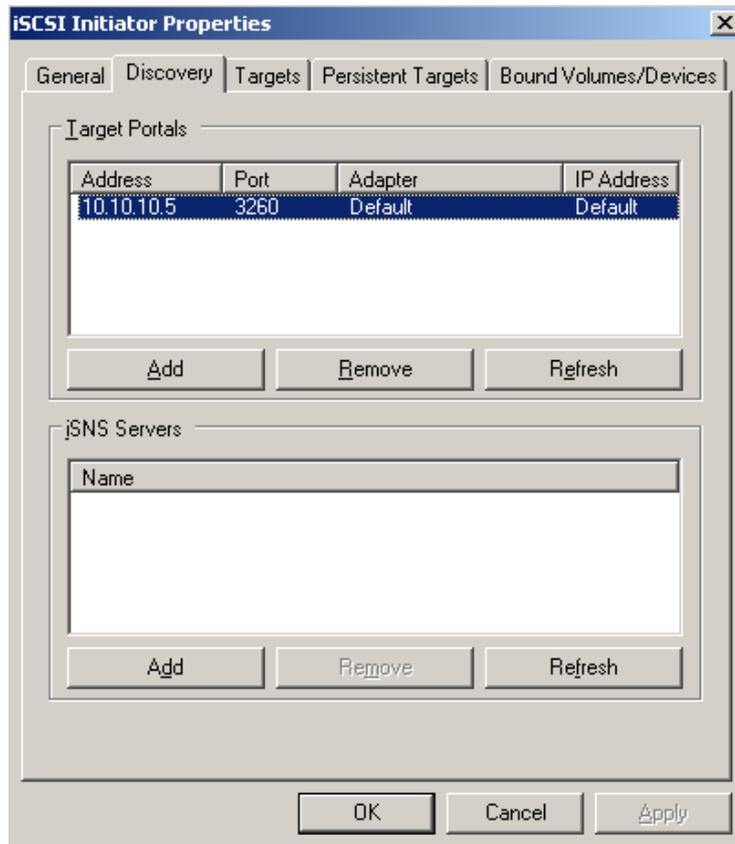
### 3.4 ATTACHING THE iSCSI LUNS

Once the cluster nodes have been created and networking configured we need to start the Windows iSCSI initiator and discover the iSCSI LUNs. Open the Windows Control Panel and click the iSCSI Initiator. The following dialog appears! Click the 'Discovery' tab.

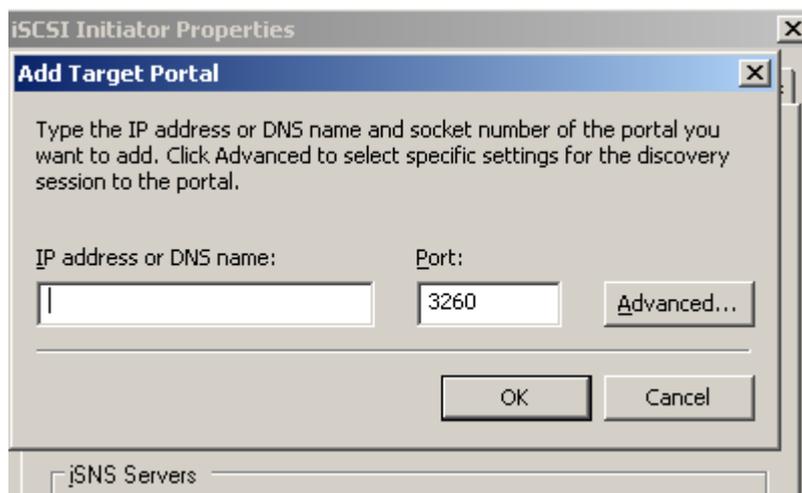


*You may use IPsec for secure communications. Most importantly your iSCSI traffic should pass over a private, segregated network.*

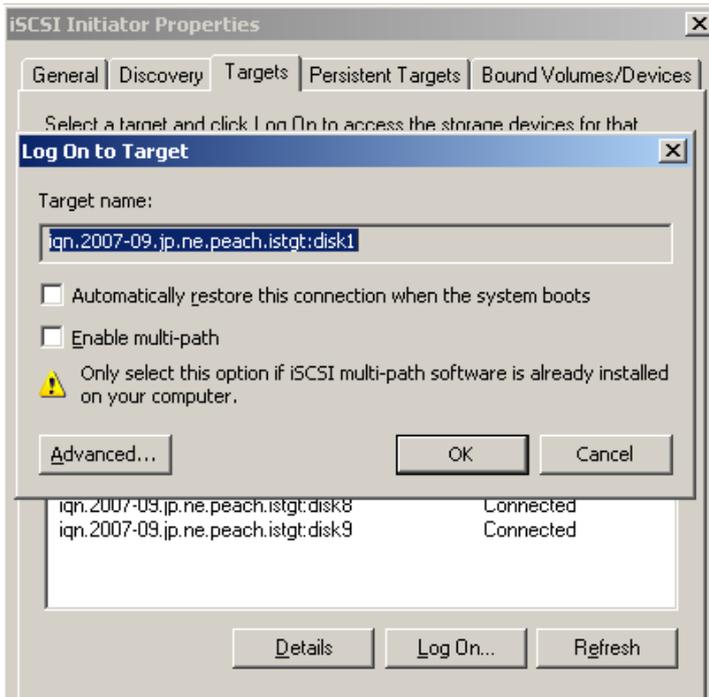
Click 'Add' and enter the NAS VM IP address from the iSCSI network not the Public network and click "OK". Leave the default port 3260.



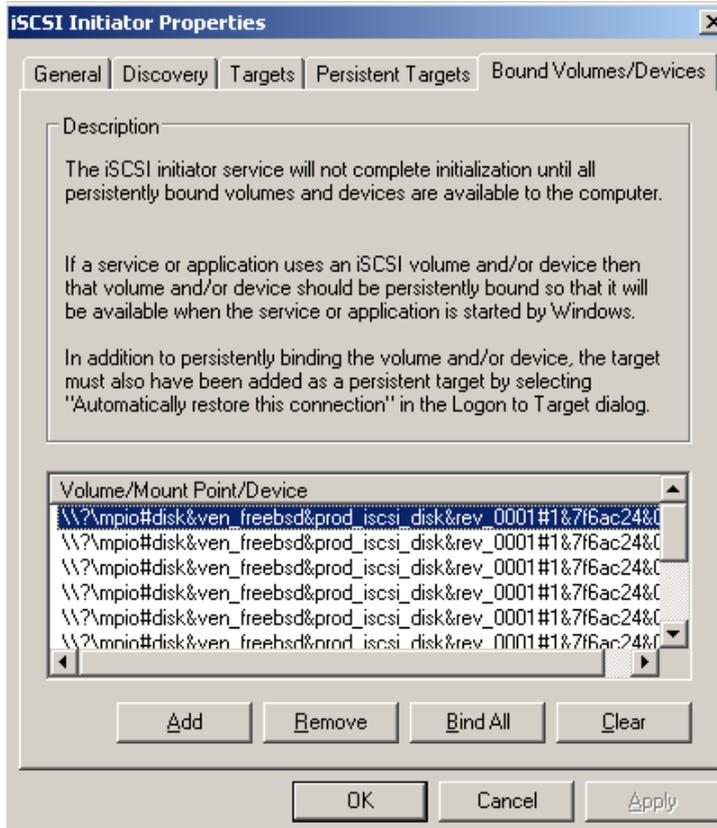
The 'Add' portal dialog



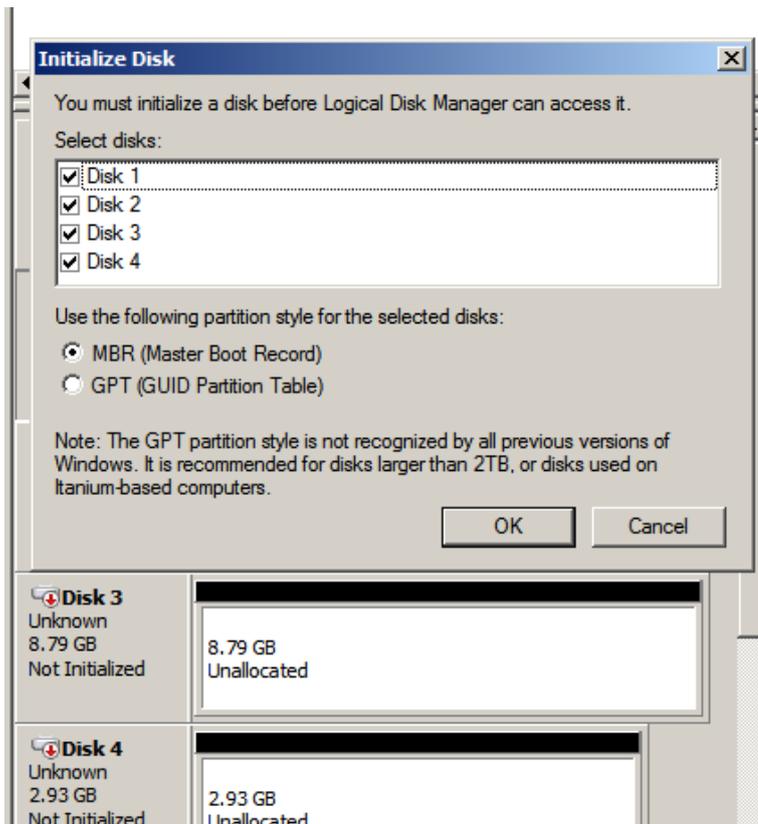
On the 'Targets' tab click each disk in turn and select 'Logon'. Check the 'Automatically restore.....' checkbox and select 'OK'.



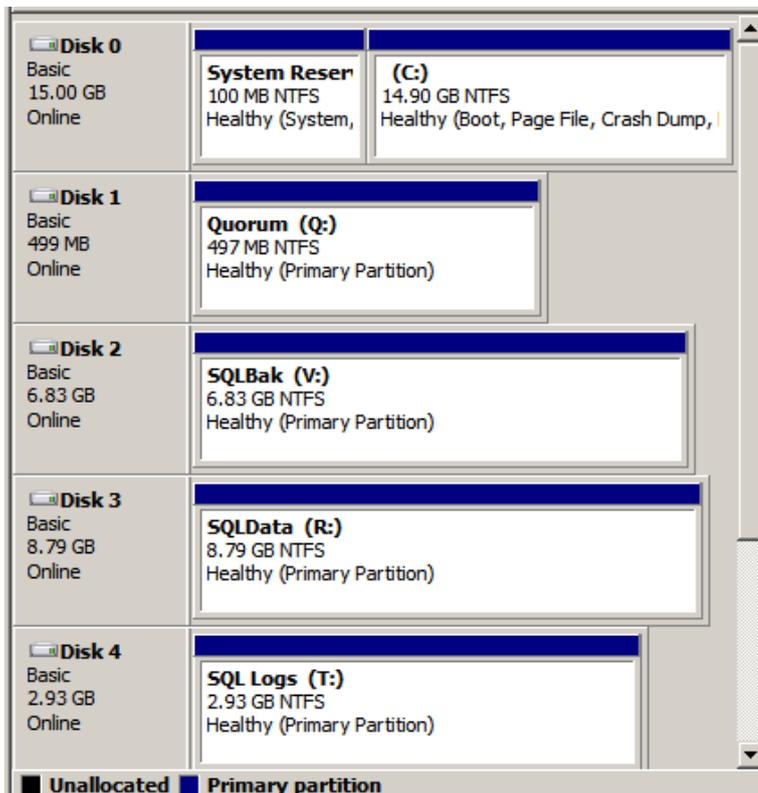
Go to the 'Bound volumes\devices' tab and select 'Bind All'. Finally, click 'OK'.



If you now open disk management you will be asked to initialise the disks. Do this, ensuring you do not select the option to convert the disks to dynamic disks.

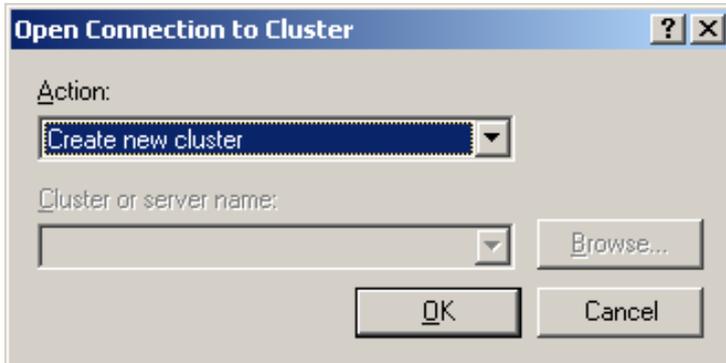


Create your disk partitions in the usual manner. All disks online and formatted!



### 3.5 INSTALLING THE WINDOWS 2003 CLUSTER

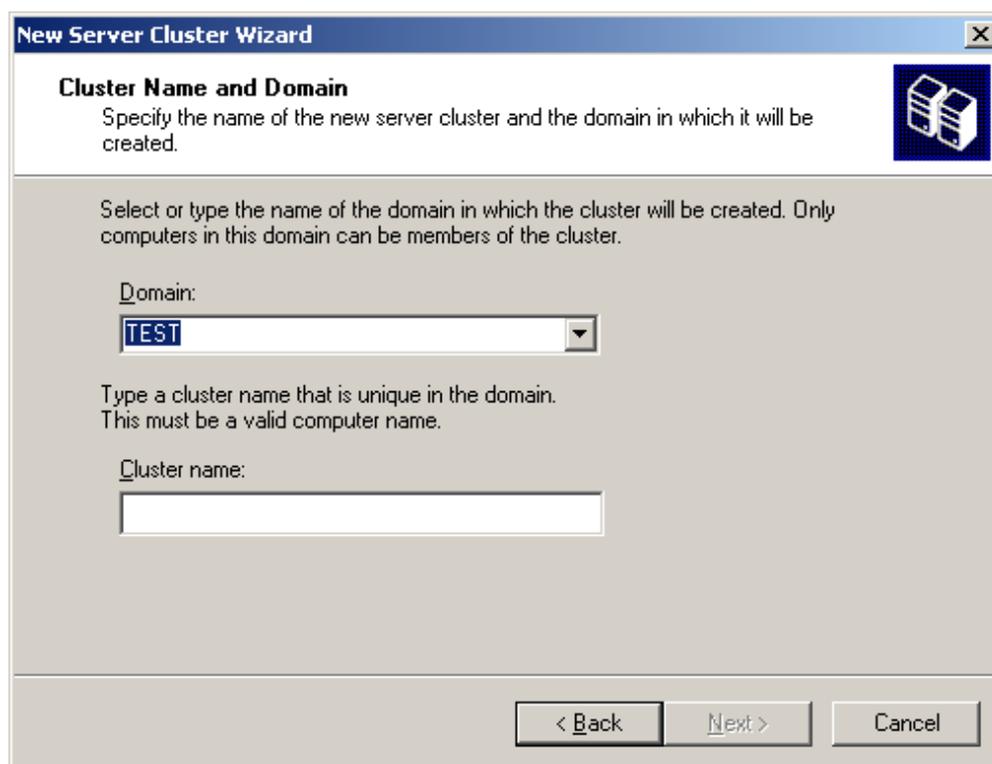
Now we need to install and configure the first node into the new cluster, shutdown Node2. Boot the first Node and open cluster administrator and select the option to create a new cluster;



Click "Next" at the welcome screen;

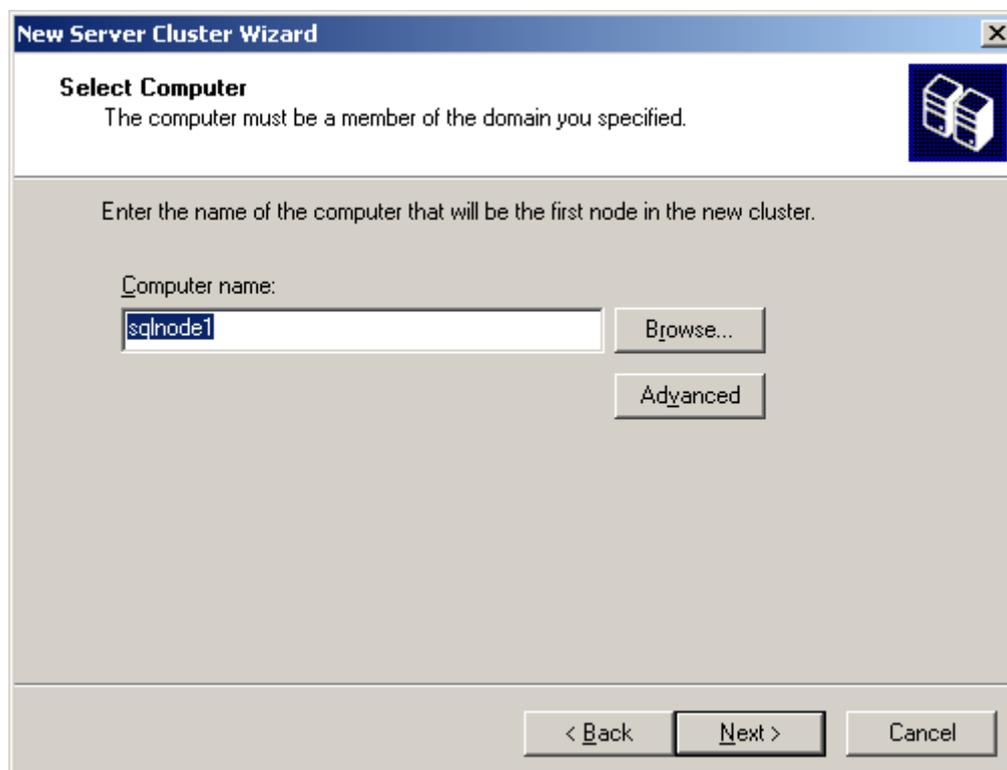


Select the domain and enter a unique cluster name, then click “Next”;



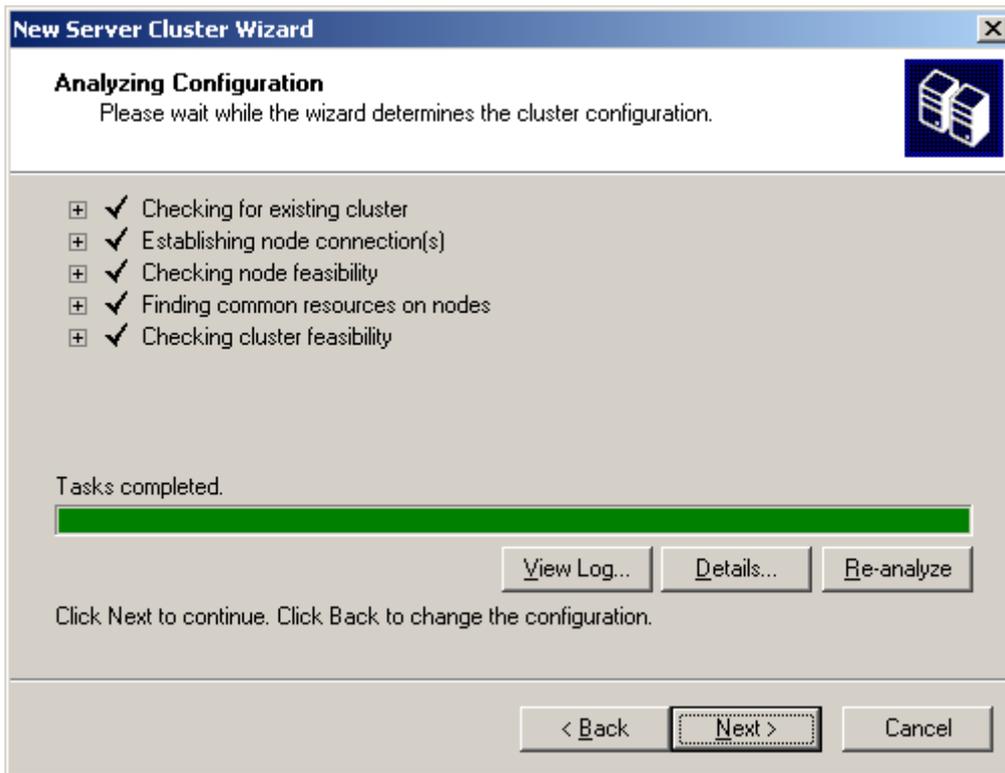
The screenshot shows the 'New Server Cluster Wizard' dialog box. The title bar reads 'New Server Cluster Wizard'. The main heading is 'Cluster Name and Domain'. Below the heading, it says 'Specify the name of the new server cluster and the domain in which it will be created.' There is a small icon of a server rack in the top right corner. The main area contains the following text: 'Select or type the name of the domain in which the cluster will be created. Only computers in this domain can be members of the cluster.' Below this is a 'Domain:' label followed by a dropdown menu with 'TEST' selected. Then, it says 'Type a cluster name that is unique in the domain. This must be a valid computer name.' Below this is a 'Cluster name:' label followed by an empty text box. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Confirm the node to add to the cluster and click “Next”;

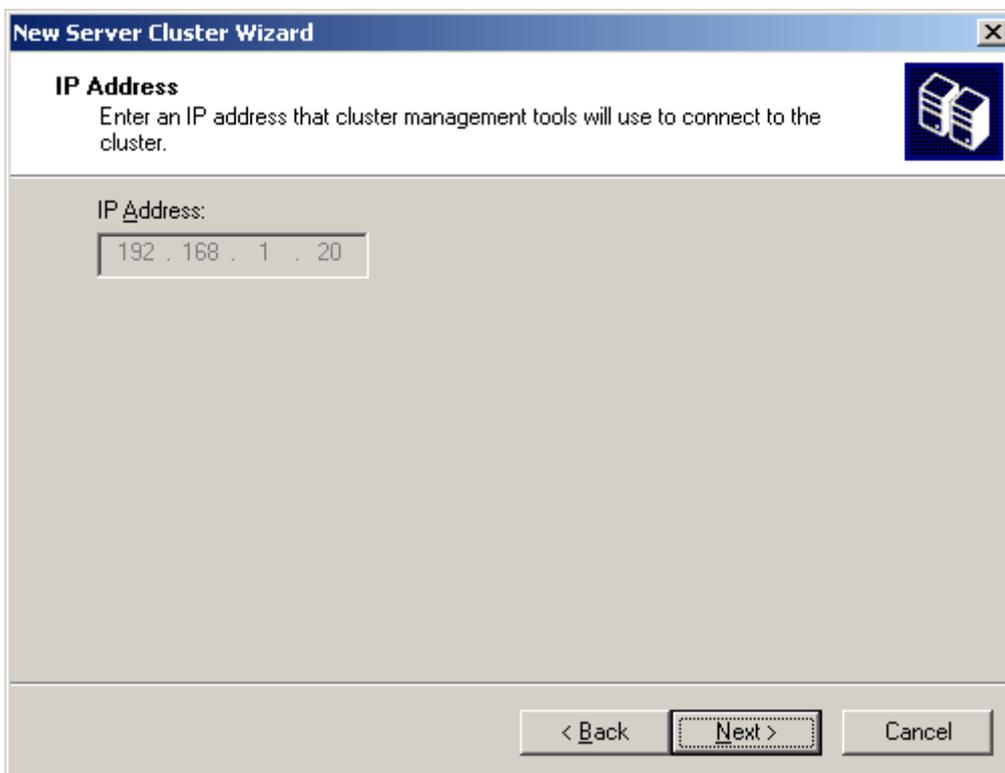


The screenshot shows the 'New Server Cluster Wizard' dialog box. The title bar reads 'New Server Cluster Wizard'. The main heading is 'Select Computer'. Below the heading, it says 'The computer must be a member of the domain you specified.' There is a small icon of a server rack in the top right corner. The main area contains the following text: 'Enter the name of the computer that will be the first node in the new cluster.' Below this is a 'Computer name:' label followed by a text box containing 'sqlnode1'. To the right of the text box is a 'Browse...' button. Below the text box and 'Browse...' button is an 'Advanced' button. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

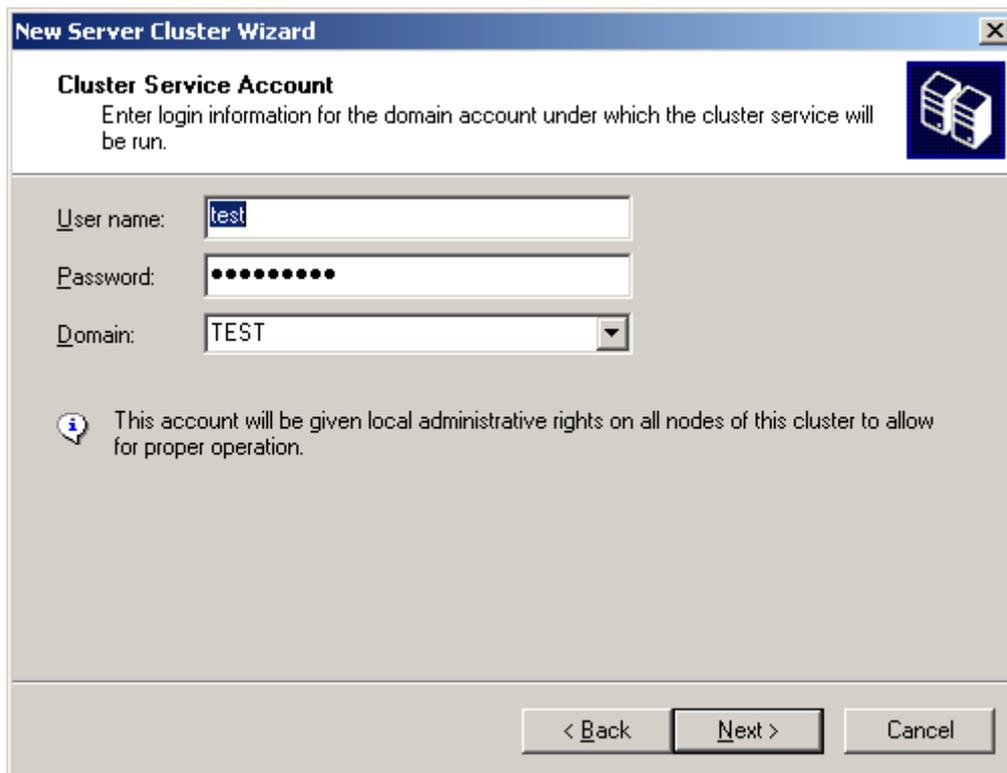
Review any errors or warnings from the analysis stage and click “Next” to proceed;



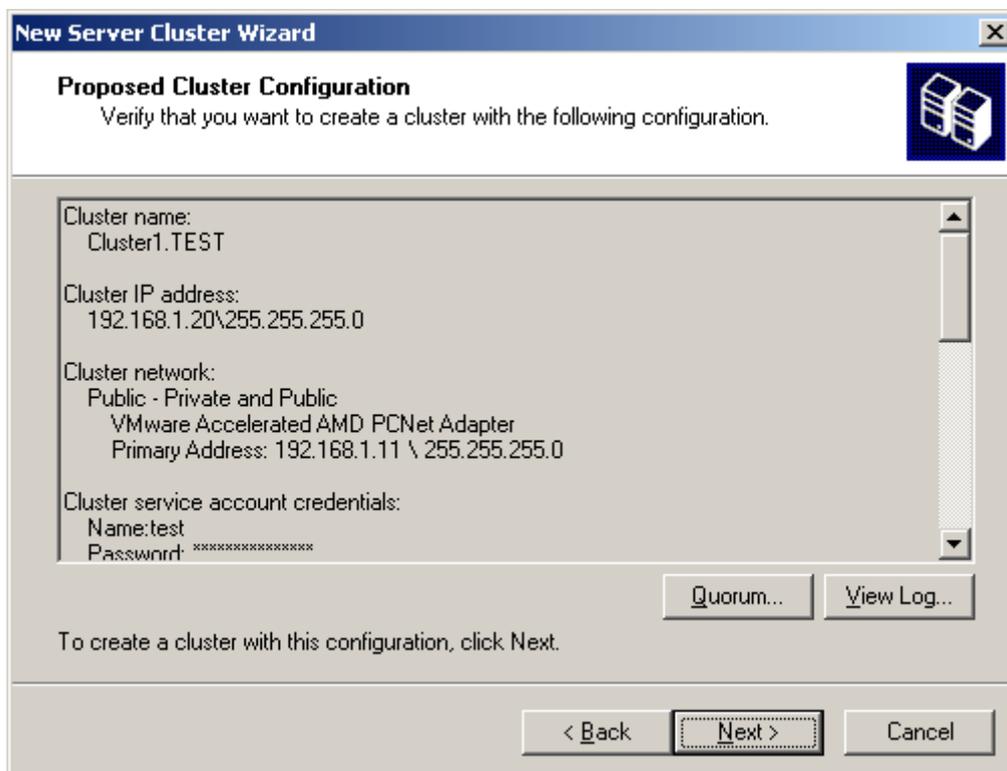
Enter a unique IP Address for the cluster and click “Next”;



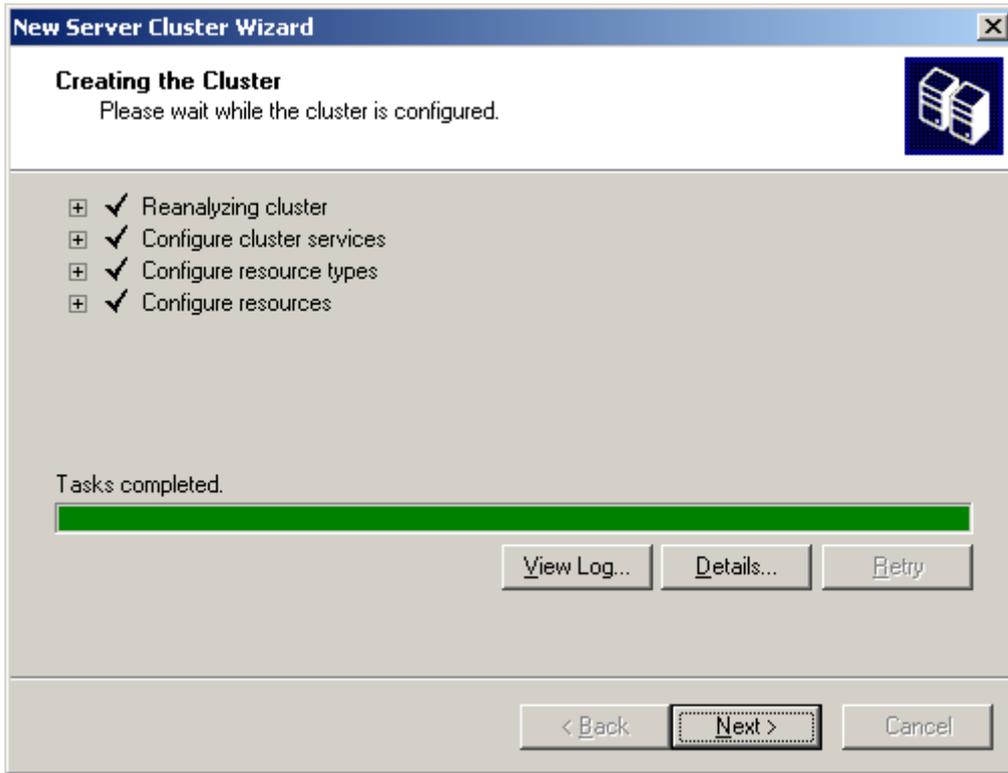
Enter the cluster service account credentials and click “Next”;



Review the proposed configuration and click “Next”;



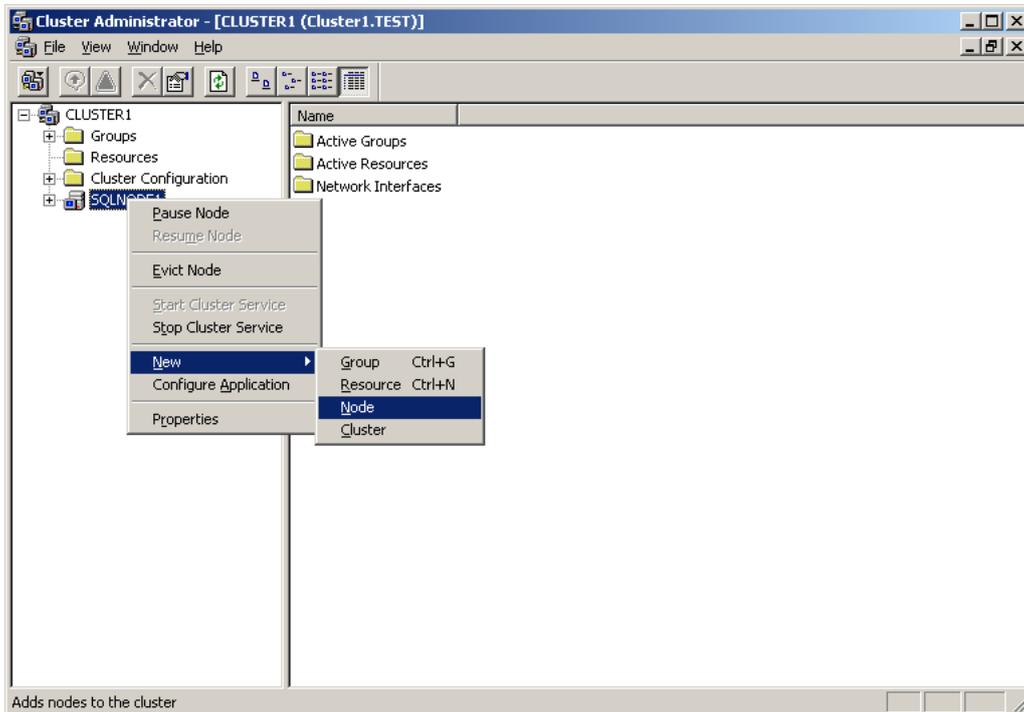
Once the cluster has been configured review any errors or warnings and click “Next” to proceed;



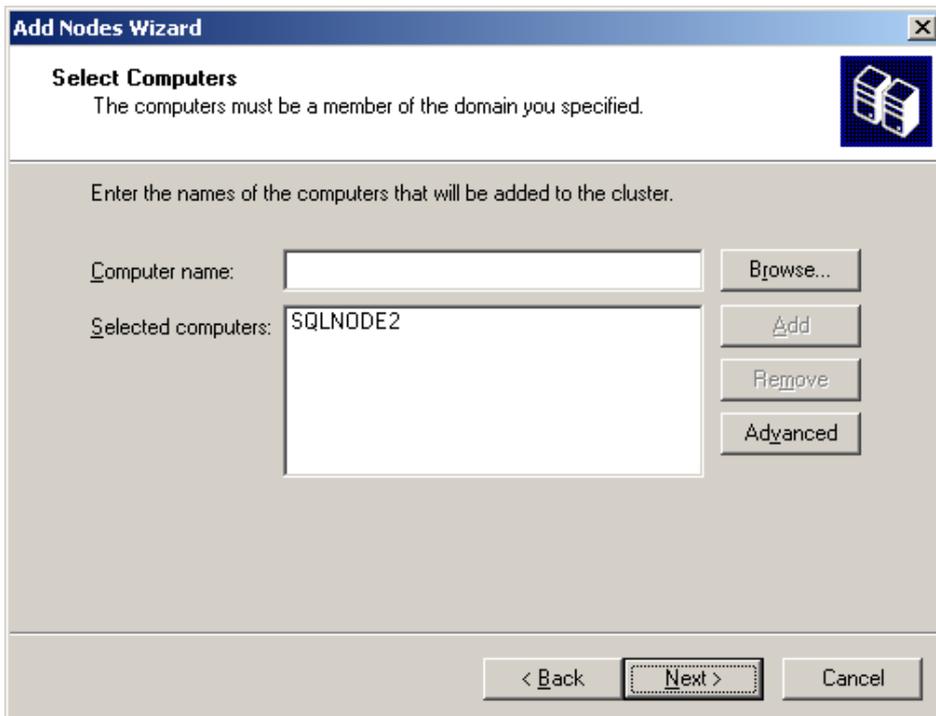
Finally click “Finish” to exit;



Now the first cluster node is active and managing the resources you may now boot SQLNode2. From within inside cluster administrator on SQLNode1, right click the cluster and select new > node and the add node wizard will start.

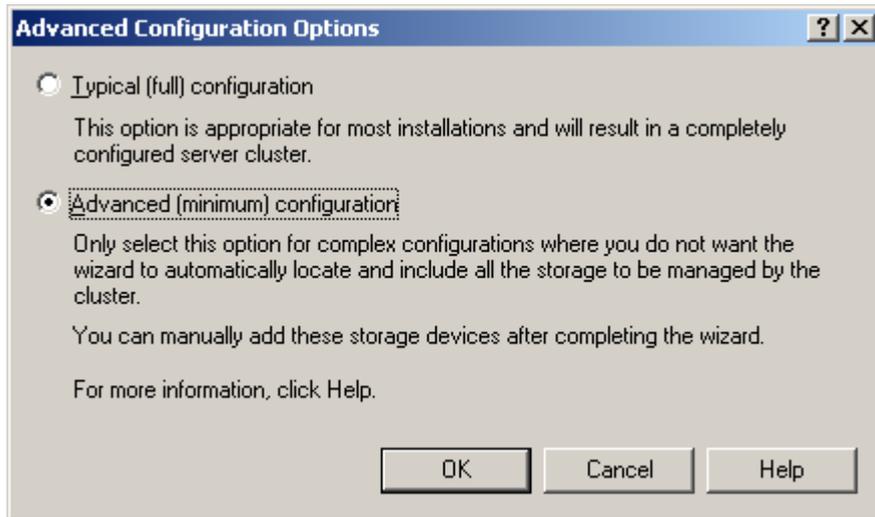


Click “Next” through the welcome screen and you will be asked to provide the computer name of the node you wish to add. Browse for the computer name and select from the list returned. Click the “Add” button to confirm the node;

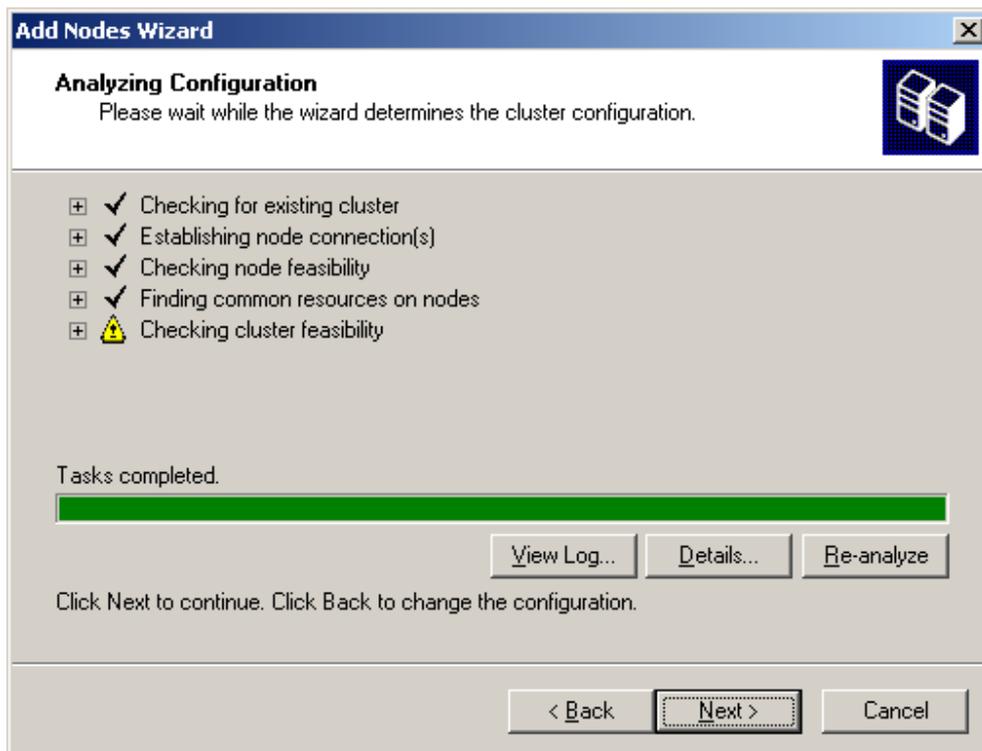


For the add node action use a minimal configuration by clicking the advanced button. Click "Next" to proceed.

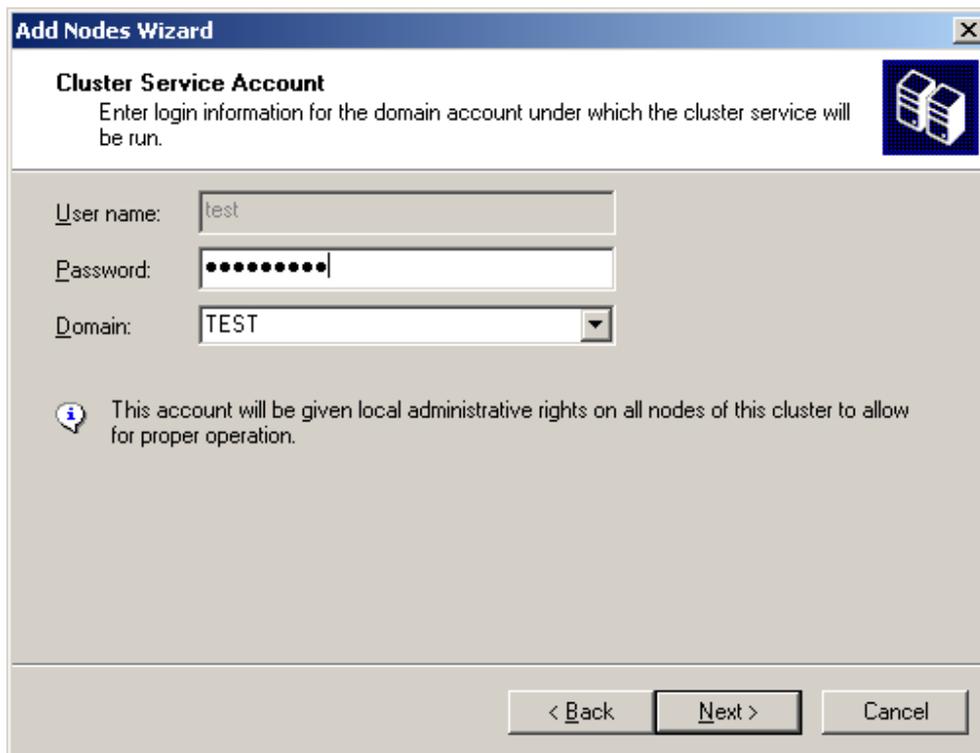
*Note: This step is not necessary when deploying a Windows 2008 cluster!*



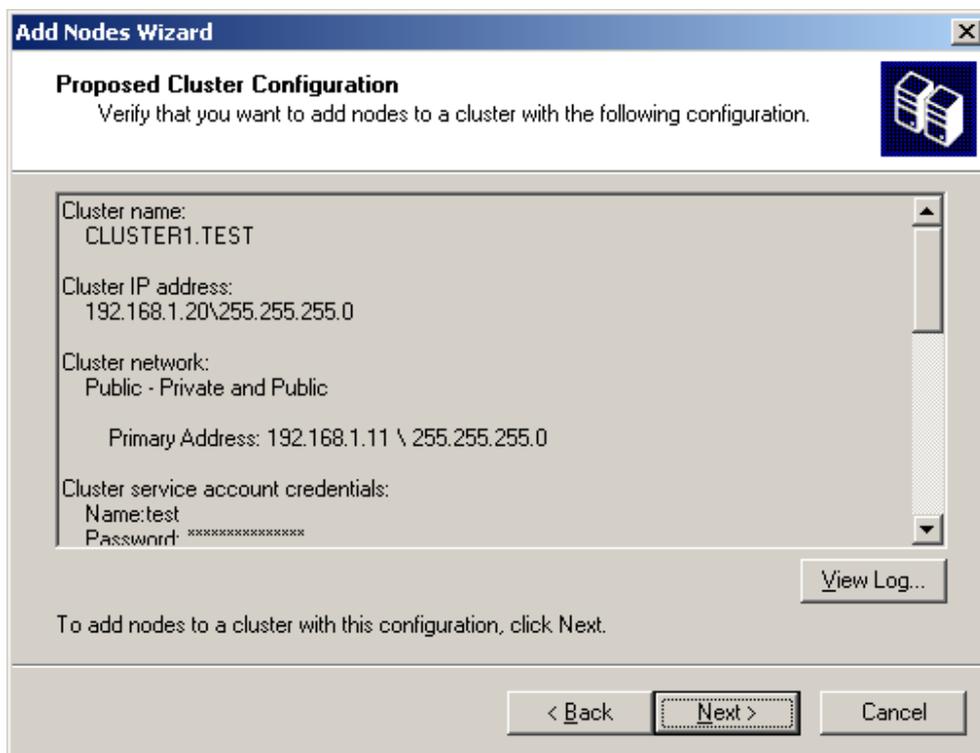
Cluster administrator analyses the configuration, review any errors or warnings.



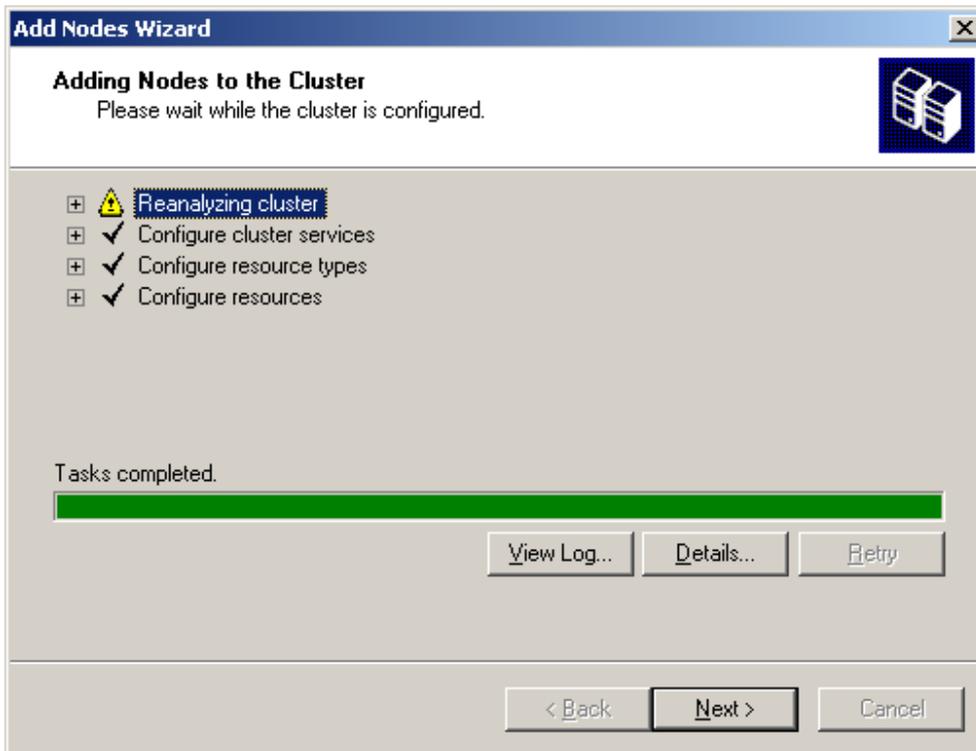
Supply the cluster service account credentials



Review the proposed configuration and click "Next";



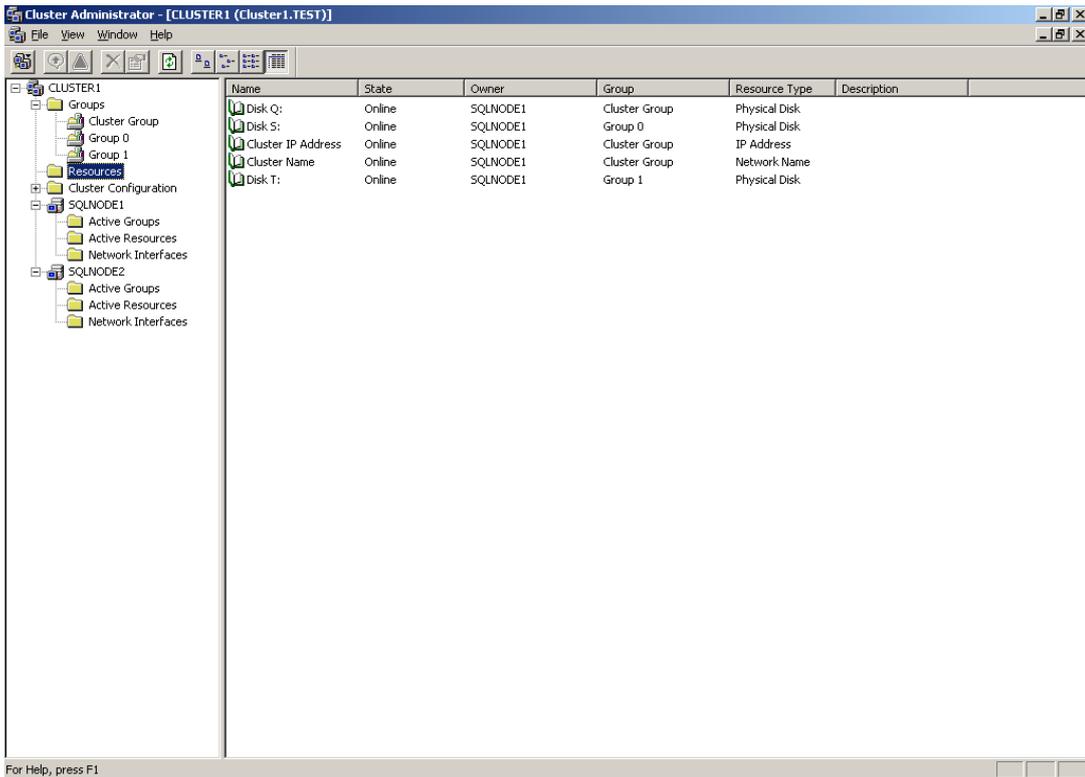
The node(s) is/are added to the cluster, again review any errors or warnings



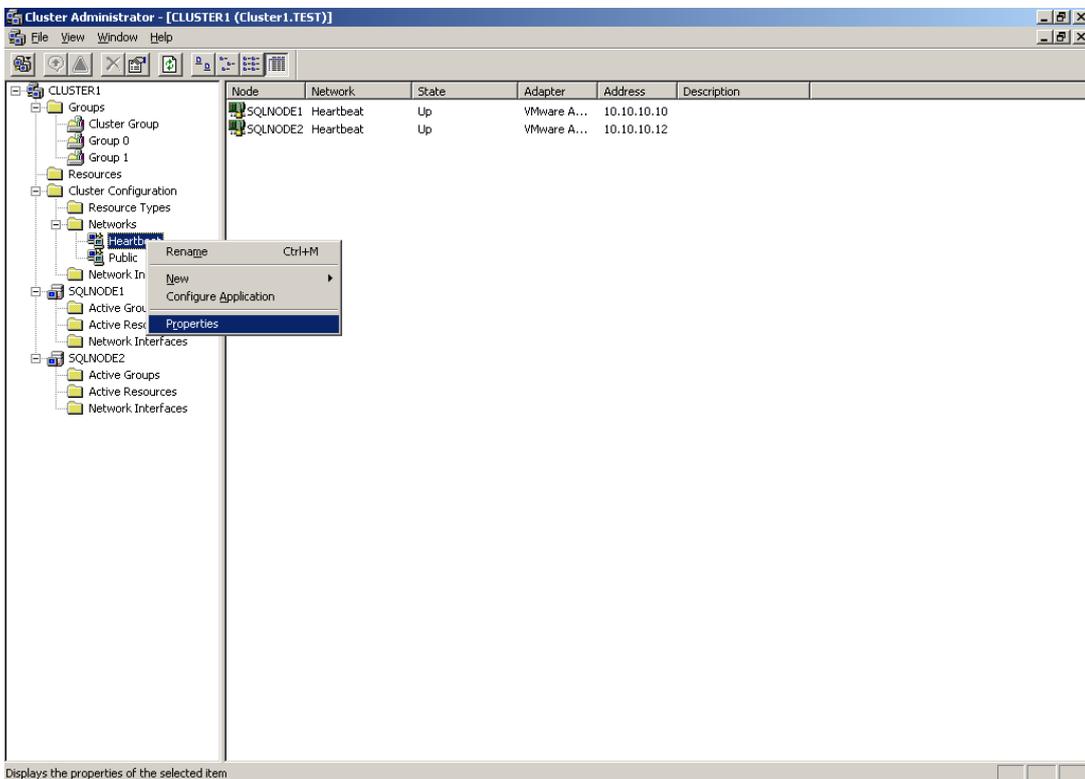
Finally click "Finish" when the wizard has completed;



Now the cluster has both nodes configured you should see the following in cluster administrator. Notice all resources owned by SQLNode1

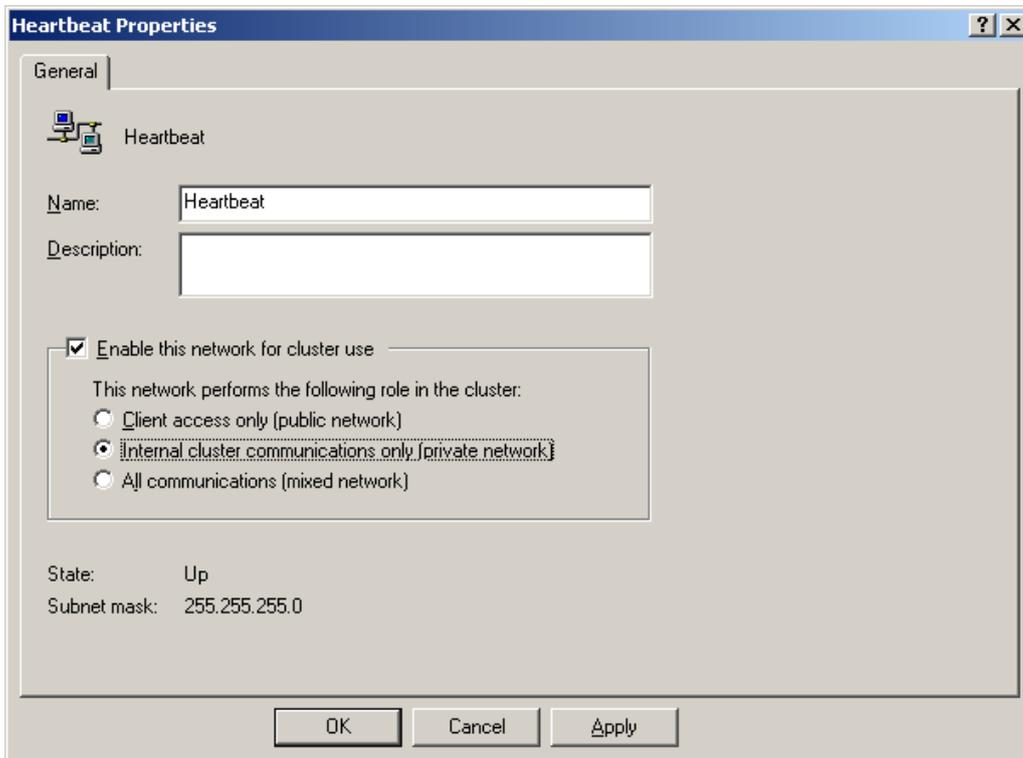


Open the networks information and set the properties for each item;

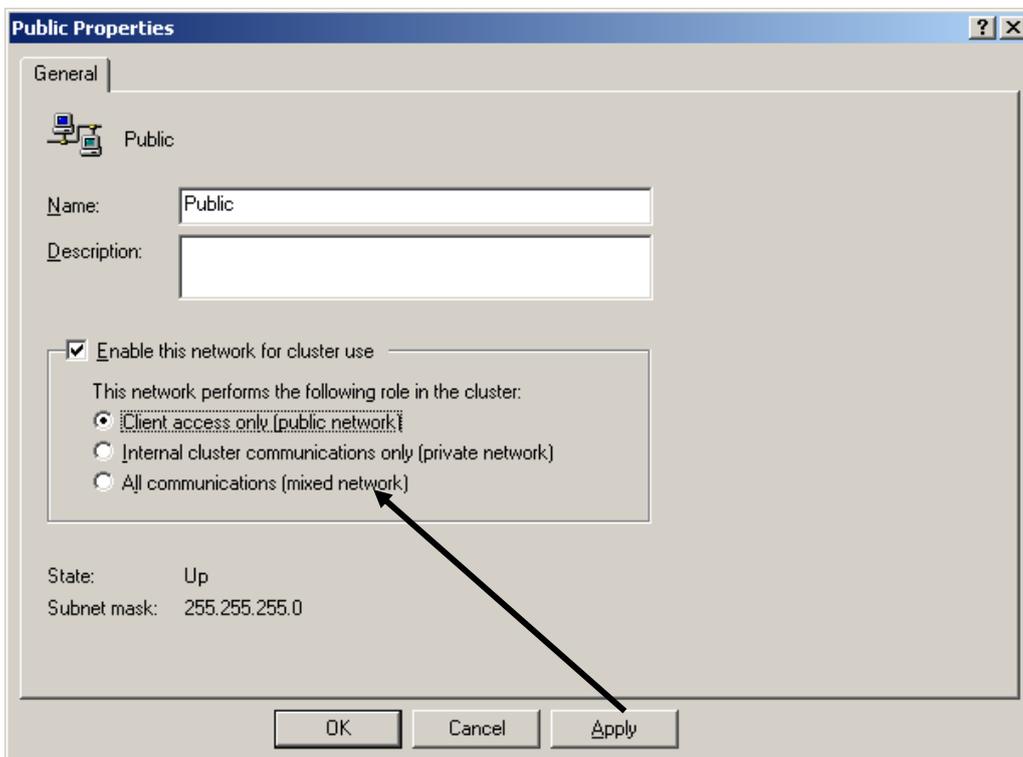


**IMPORTANT:** Ensure that on the iSCSI network connection in Cluster Administrator, you uncheck the 'Enable this network for cluster use' checkbox.

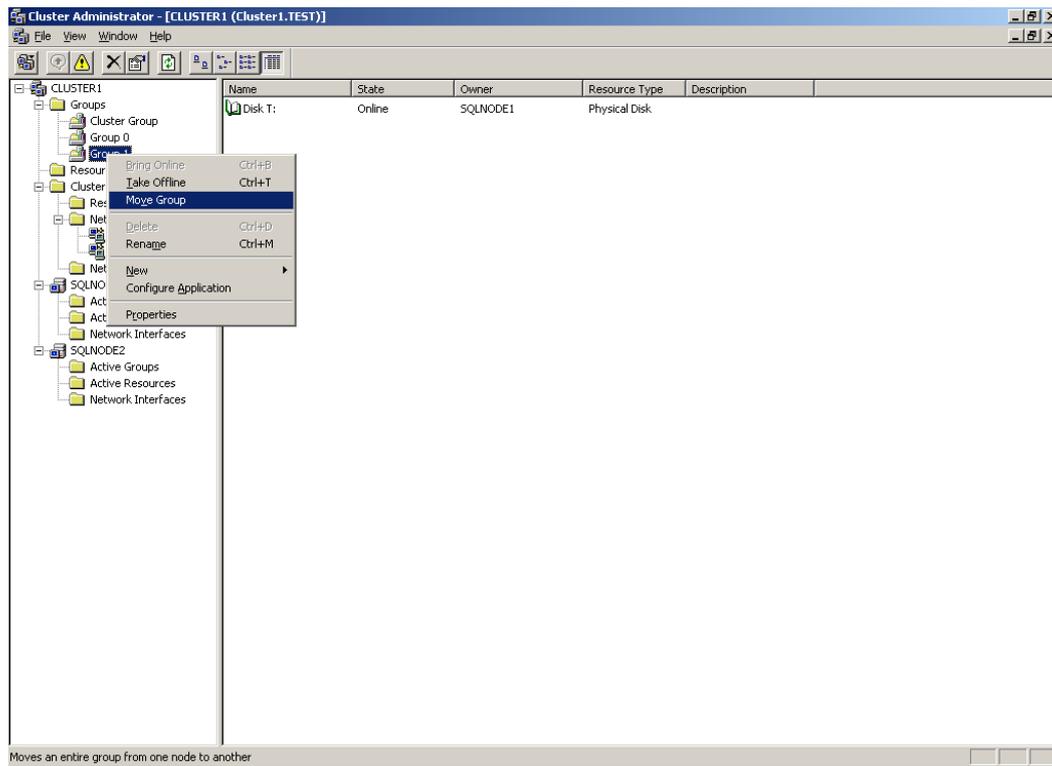
The Heartbeat needs only internal access;



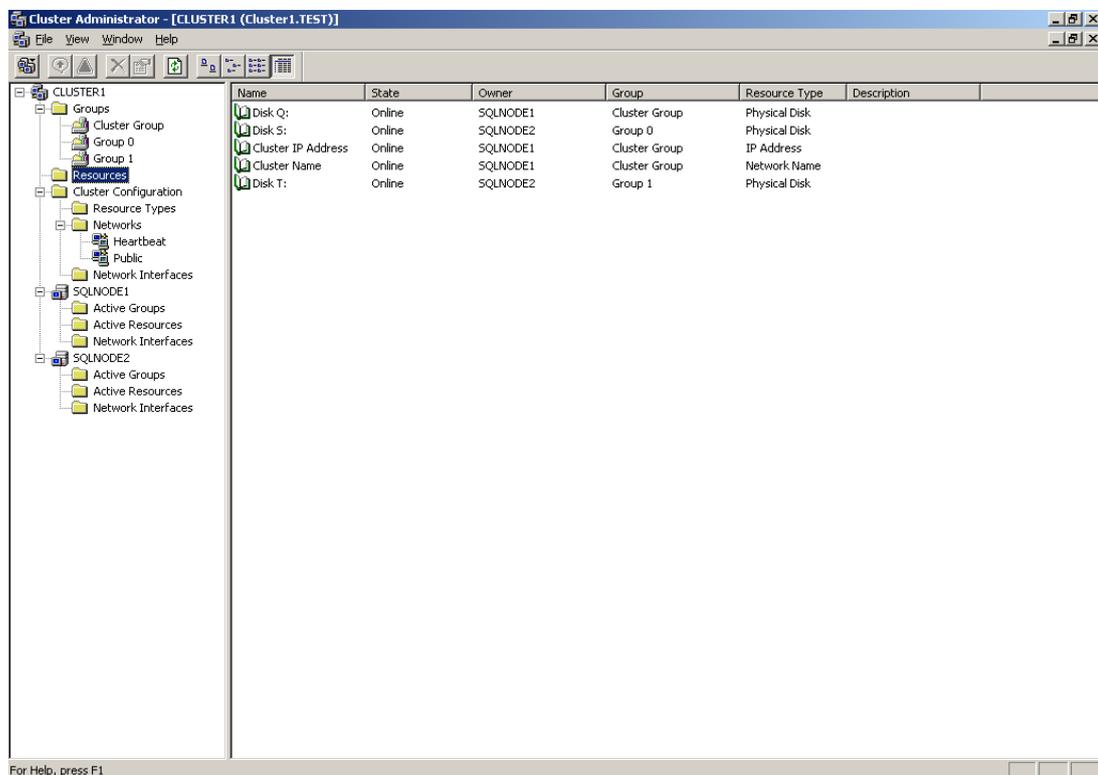
The Public should be set to Mixed Communication access;



Test your cluster configuration by moving a cluster group from one node to another. Right click a group and select "Move group". The resource will transfer to the partner node;



Groups 0 and 1 transferred to SQLNode2;



### 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.

In Cluster Administrator right click 'Groups' and select 'New' > 'Group'. Give the group a name and click 'Next' then add available nodes and click 'Finish'.

Right click the newly created group and select 'New' > 'Resource'. Give the resource a name 'INST1DTC IP', from 'Resource Type' drop down list select 'IP Address' and click 'Next'. Select available nodes and click 'Next'. Click 'Next' through dependencies, enter an IP Address and mask for the public network (192.168.0.30 and 255.255.255.0) and click 'Finish'.

Right click the group and select 'New' > 'Resource'. Give the resource a name 'INST1DTC Name'. From the drop down list select 'Network name' as the resource type and click 'Next'. Select available nodes and click 'Next'. Add IP Address resource as a dependency and click 'Next'. Enter the unique network name and uncheck the 'DNS registration must succeed' checkbox then click 'Finish'.

If the DTC disk drive has already been discovered by Cluster Administrator you may skip this paragraph. Otherwise, right click the group and select 'New' > 'Resource'. Give the resource a name 'DTC Data' and from the drop down list select 'Physical Disk' as the resource type, then click 'Next'. Select available nodes and click 'Next'. Click 'Next' through the dependencies. From the drop down list on the parameters dialog select the disk drive to use (P:) then click 'Finish'.

Lastly create the DTC resource by right clicking the group and selecting 'New' > 'resource'. Give the resource a name 'INST1DTC SVC' and select 'Distributed Transaction Coordinator' from the 'Resource Type' drop down list, then click 'Next'. Select available nodes and click 'Next'. At the dependency dialog select the resources for Network Name and Physical Disk only (These resources must be online for the service to start) then click 'Finish'. Now right click the group and bring it 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 these services. With that done, it's now time to start the installation.

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

Due to the issues with the RTM version of the SQL Server 2008 installation media it is necessary to 'Slipstream' the Service Pack 1 media to create a trouble free installation.

Installing SQL Server 2008 onto a Windows 2003 cluster requires the following prerequisites;

**Windows 2003 SP2  
Hotfix 937444 (Filestream Hotfix)  
Windows Installer 4.5  
.NET Framework 3.51**

Due to the issues with the RTM version of the SQL Server 2008 installation media it is also necessary to 'Slipstream' the Service Pack 1 media to create a trouble free installation.

### 4.1 CREATE THE SQL SERVER FAILOVER INSTANCE.

Firstly extract the service pack executable to a temporary folder using the following syntax at the command prompt.

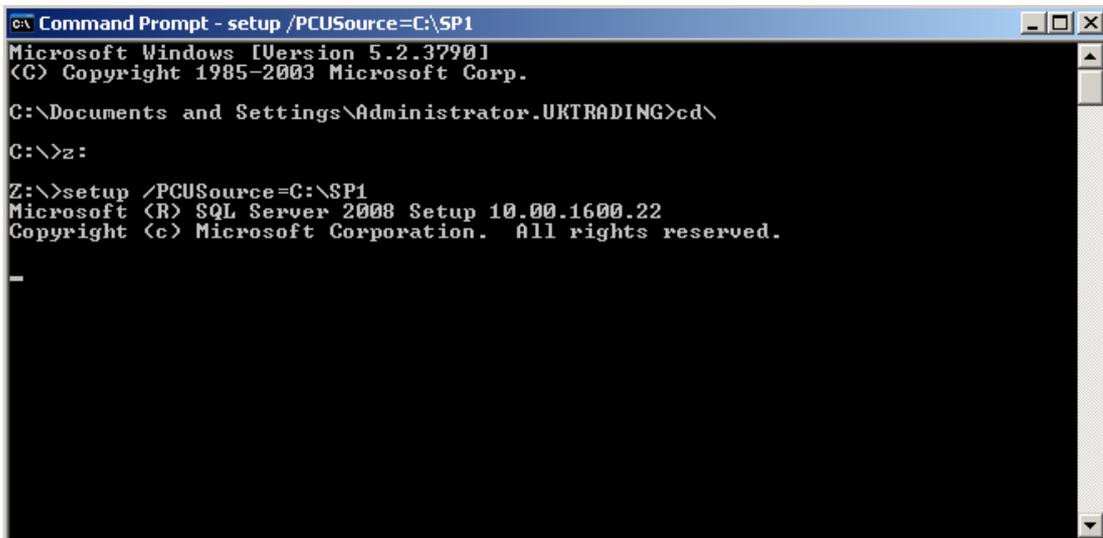
```
C:\en_sql_server_2008_sp1_x64.exe /x:C:\SP1
```

Now browse the SQL Server media directories and locate\install the Windows installer and .NET Framework, also install the Filestream hotfix (listed above). Once you have all the prerequisites installed, you may launch the SP1 support files installer. The installer is located at;

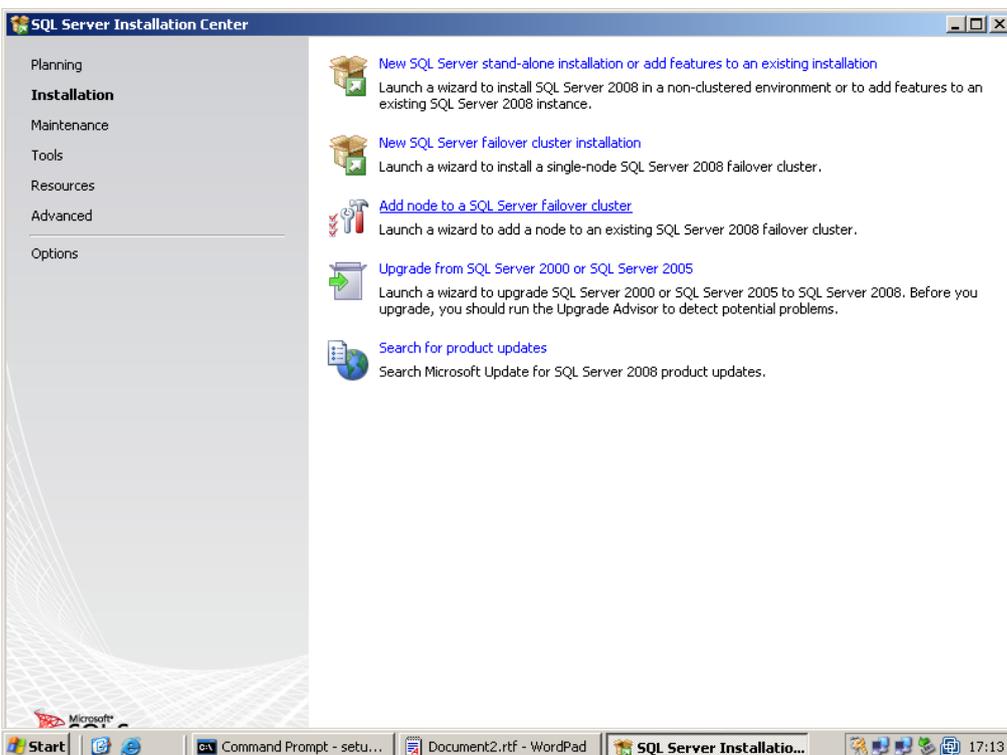
```
C:\SP1\x64\setup\1033\sqlsupport.msi
```

After the support files have been installed, launch the SQL Server 2008 setup from the command prompt using;

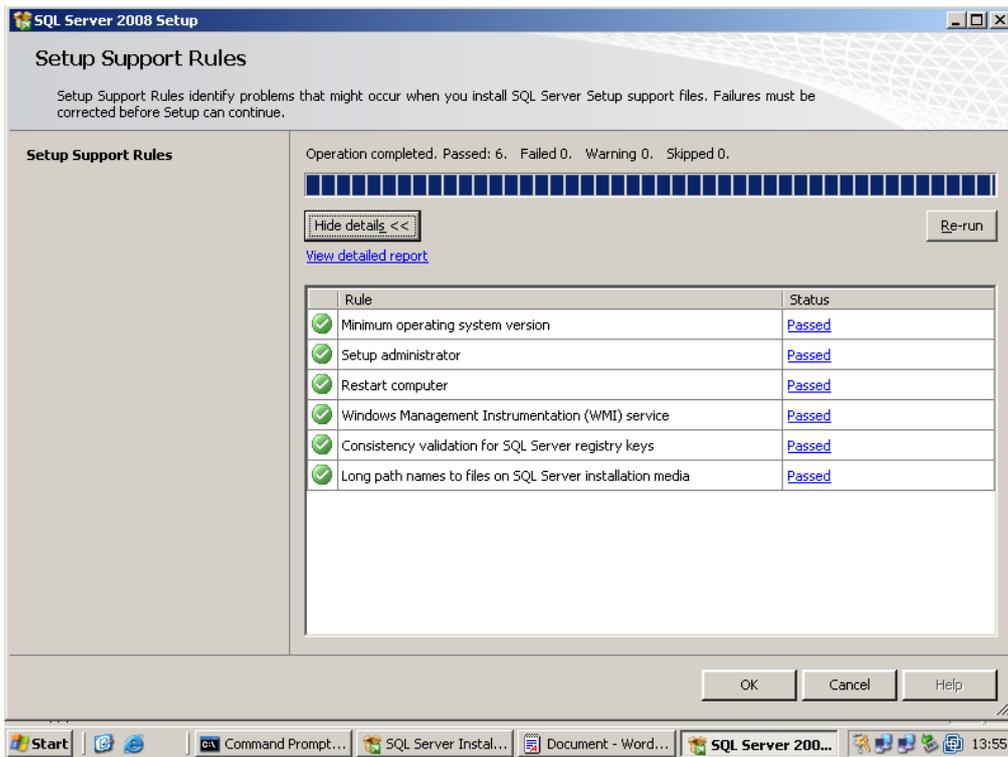
Z:\>setup.exe /PCUSource=C:\SP1



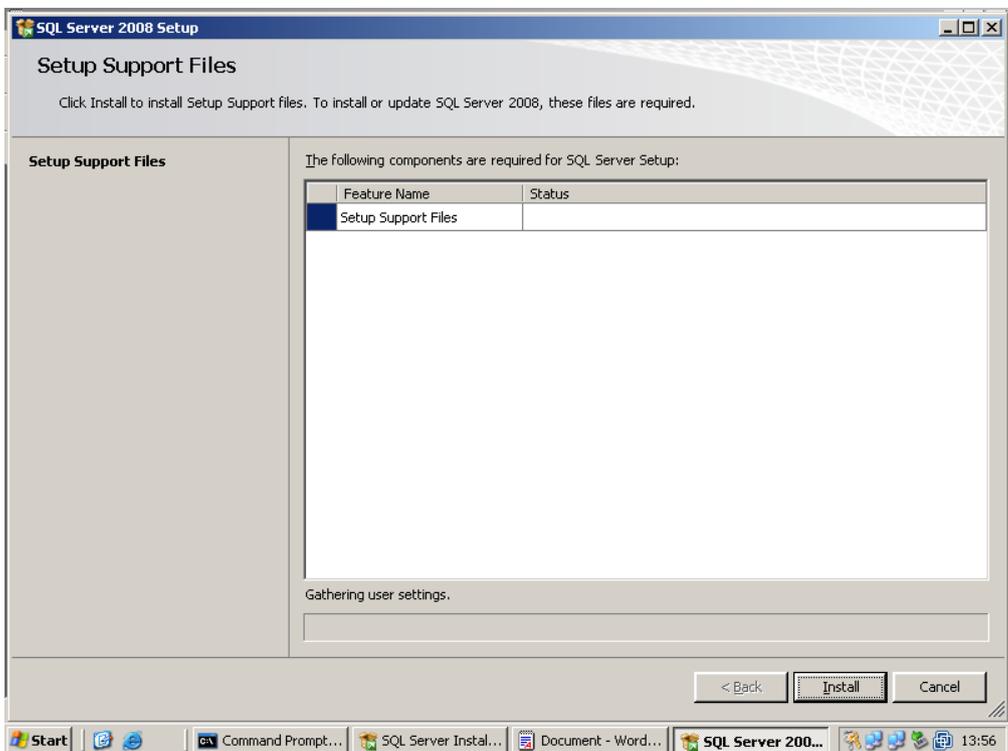
At the SQL Server 2008 splash screen select "Installation" from the left menu, then select "New SQL Server failover cluster installation".



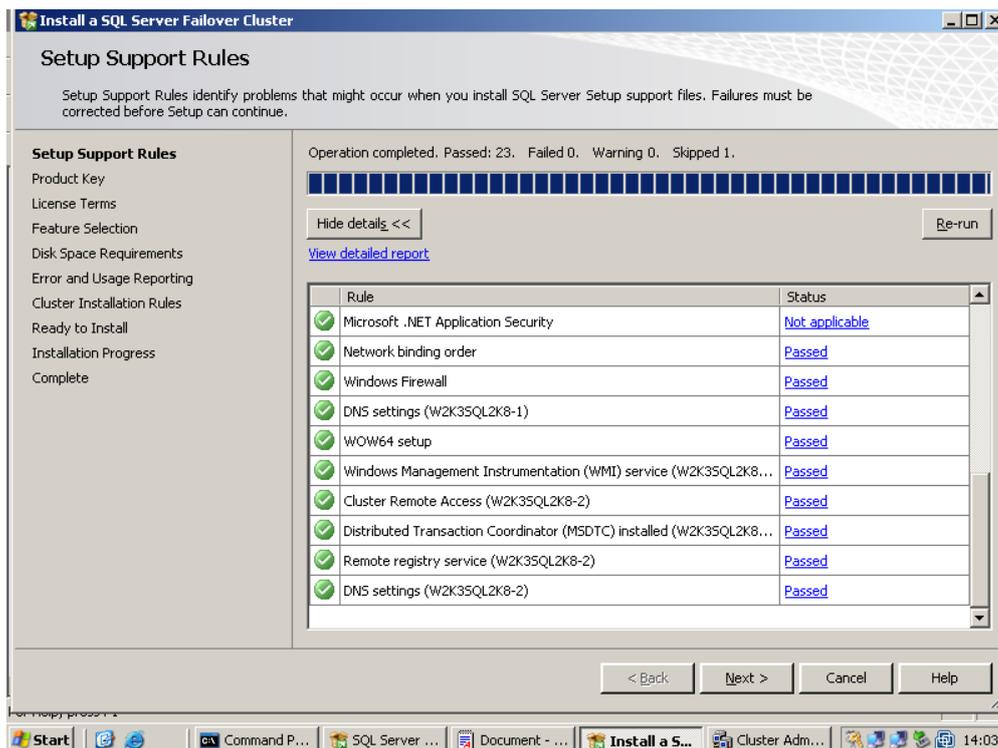
The installation will start and the following screen will shortly appear, click “OK” when prompted.



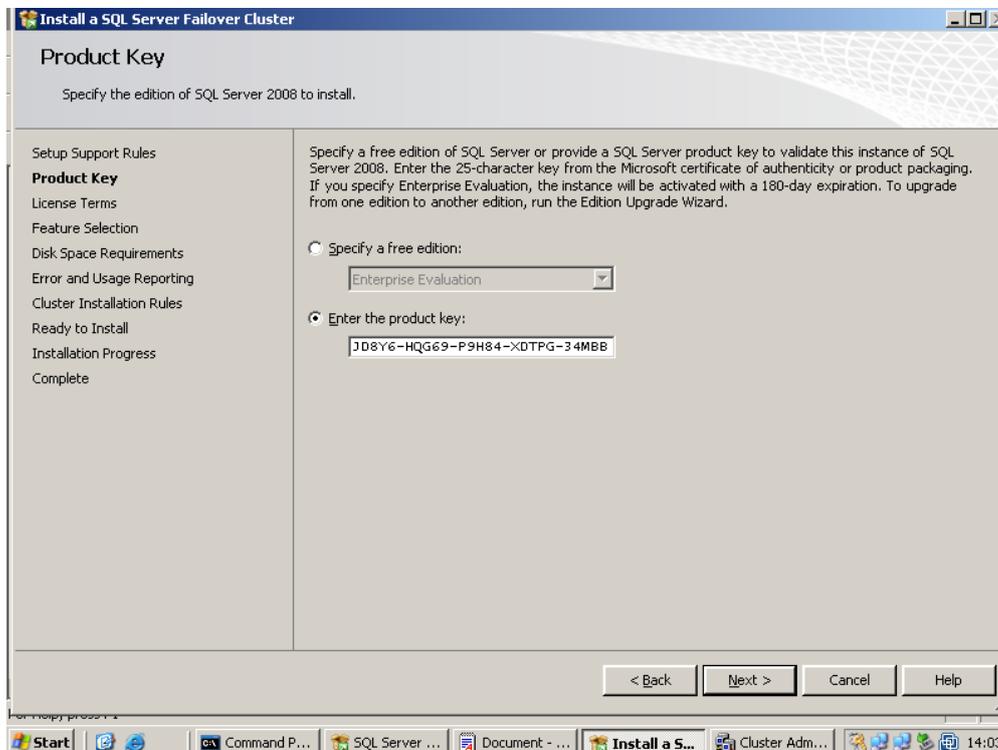
Click “Install” to install setup support files.



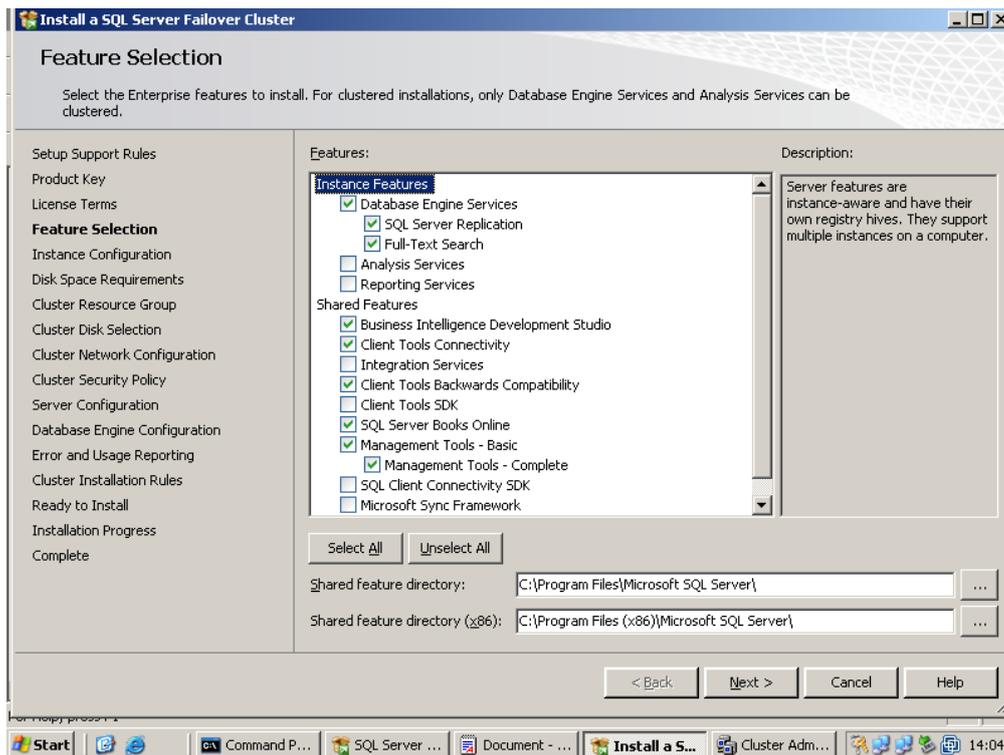
After the support rules have been checked, click "Next" to continue



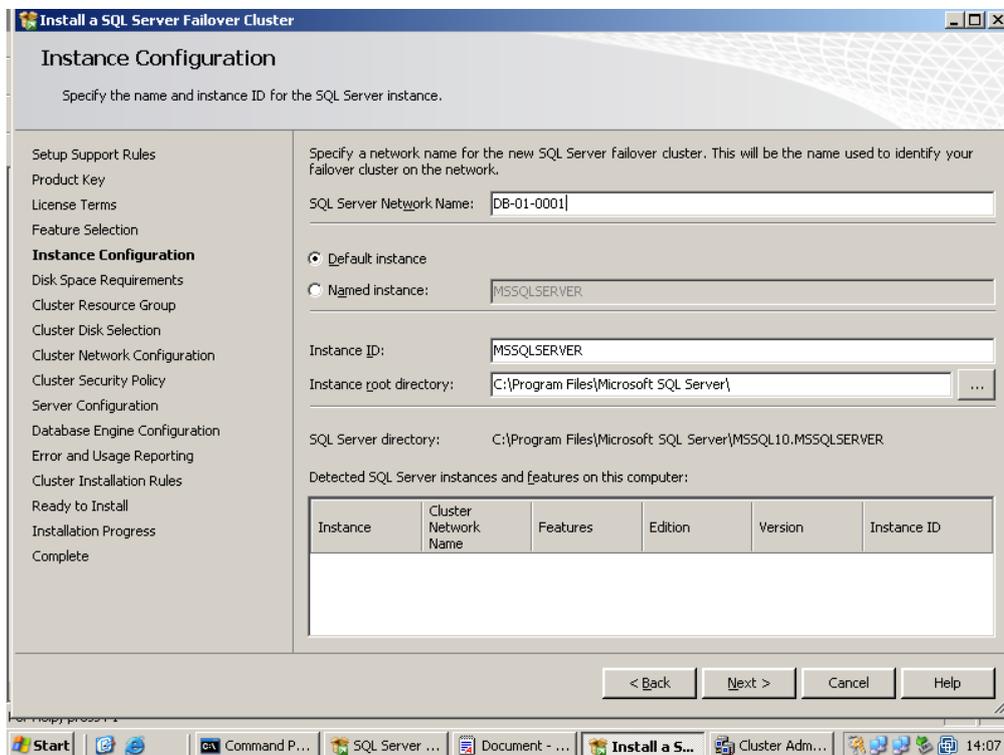
Enter your licence key or select a free edition to install!



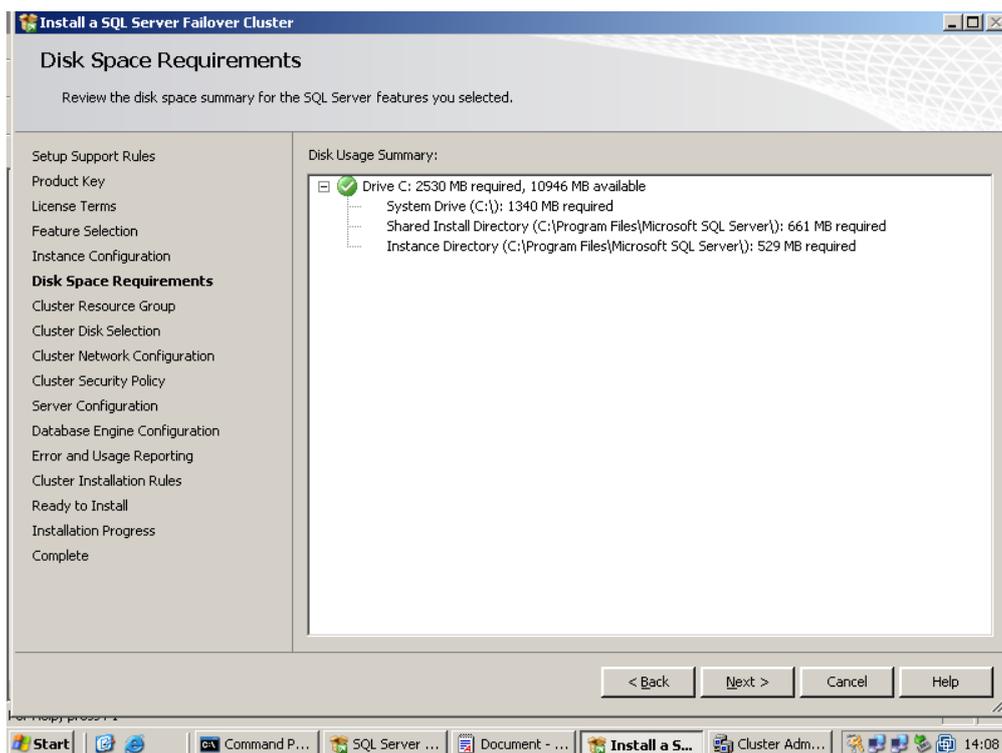
Select the features you wish to install and supply the path for the shared files.



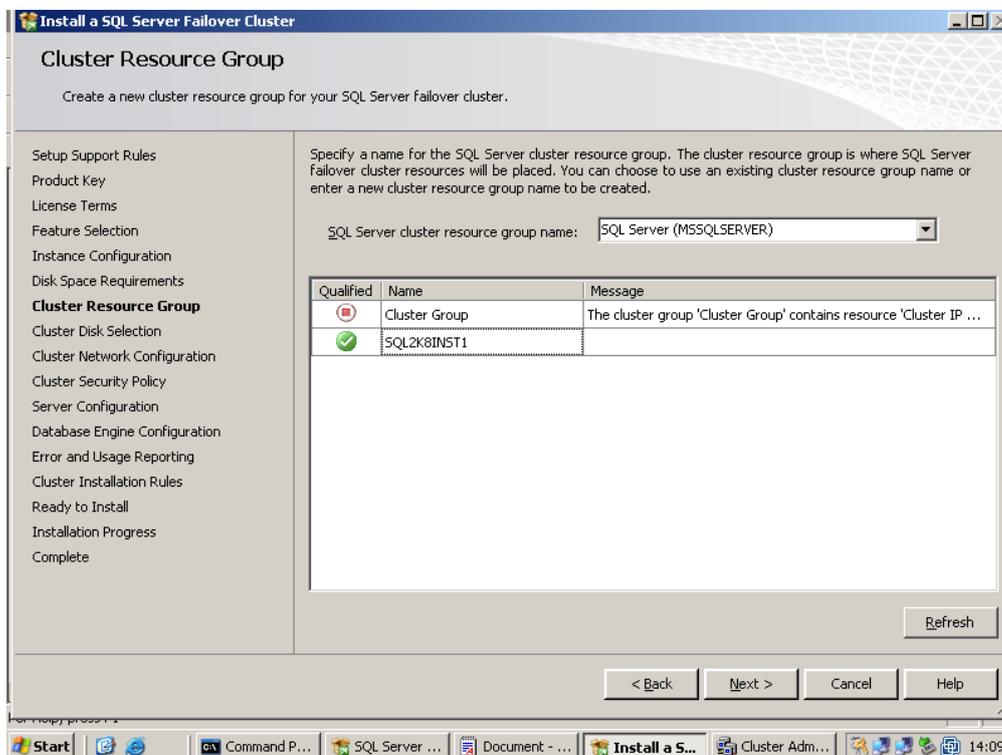
Supply a unique virtual network name and an instance name



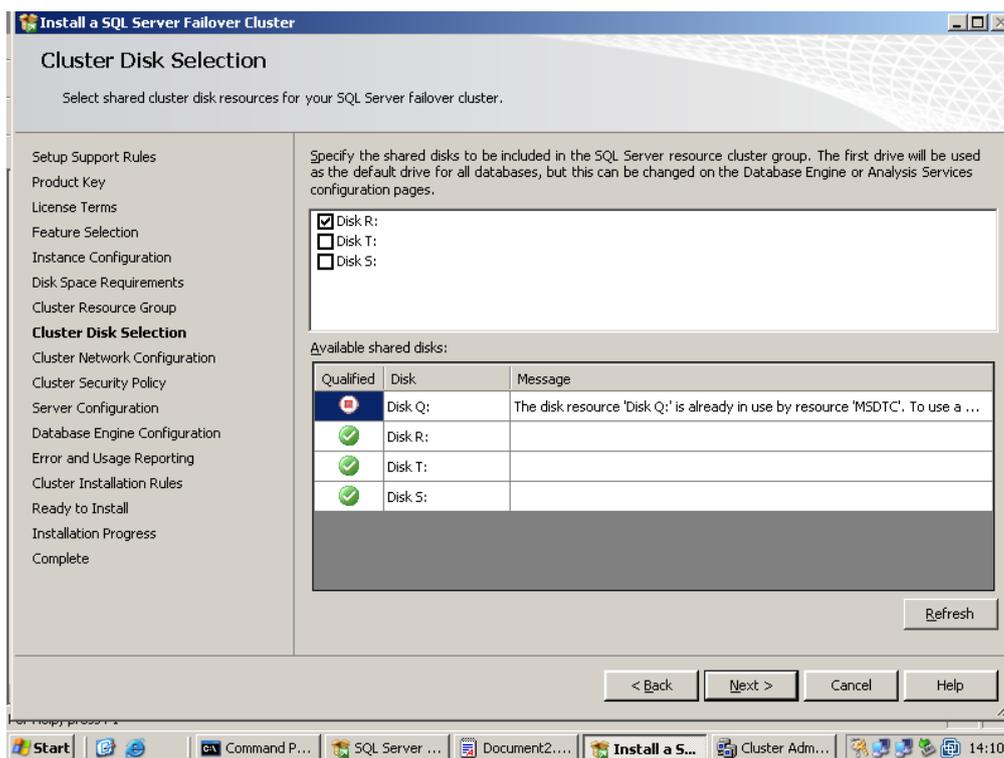
Review the space requirements and click “Next”,



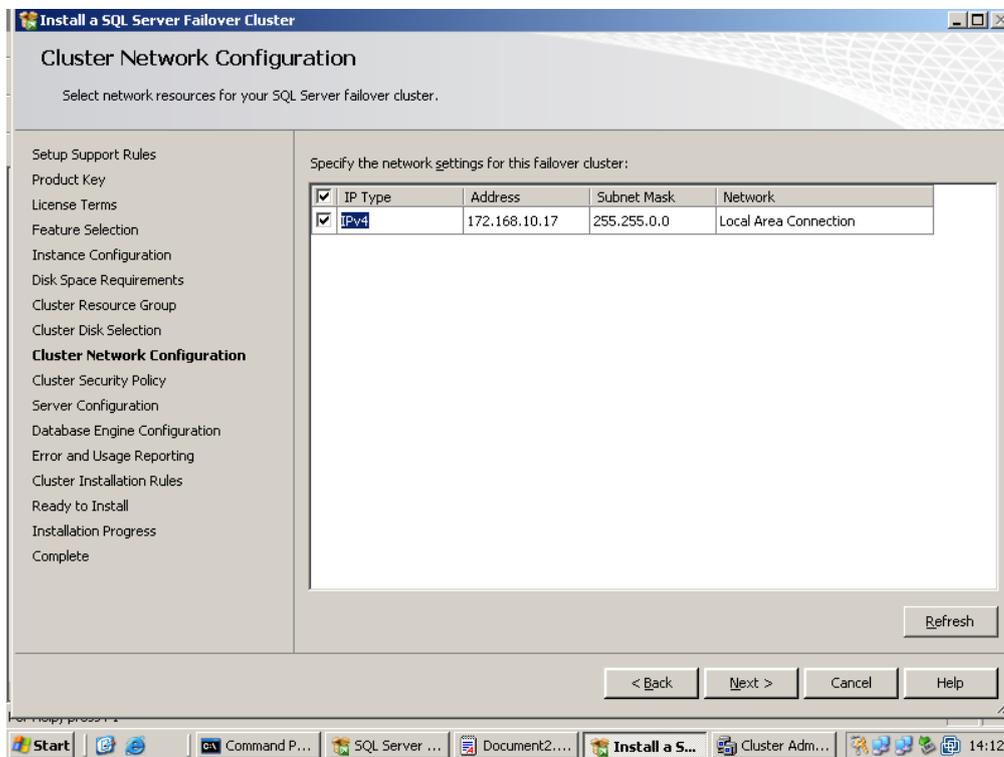
Select the default name for the SQL Server cluster resource group or select a free group that already exists (displays a green tick).



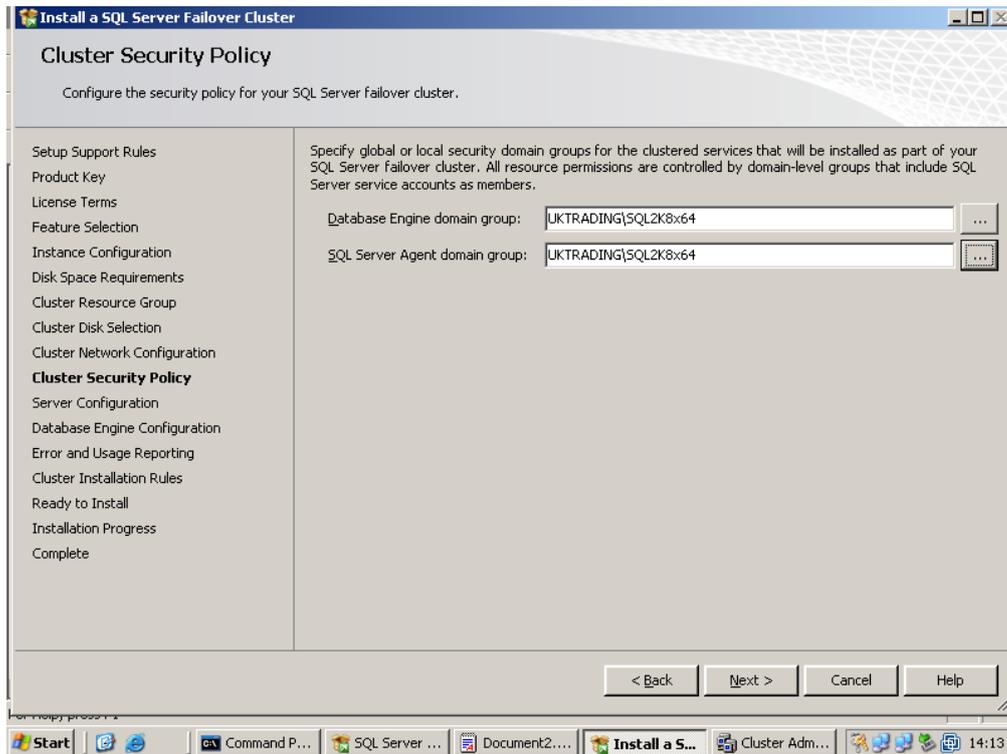
Select disk resources and click “Next”,



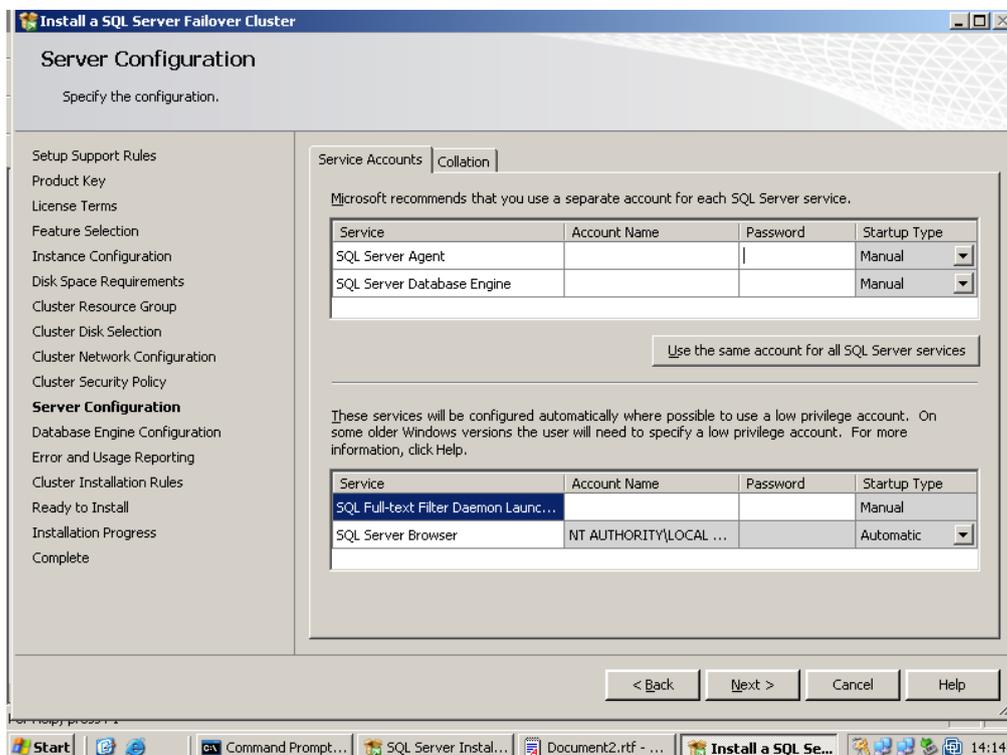
Supply an IP address and click “Next”,



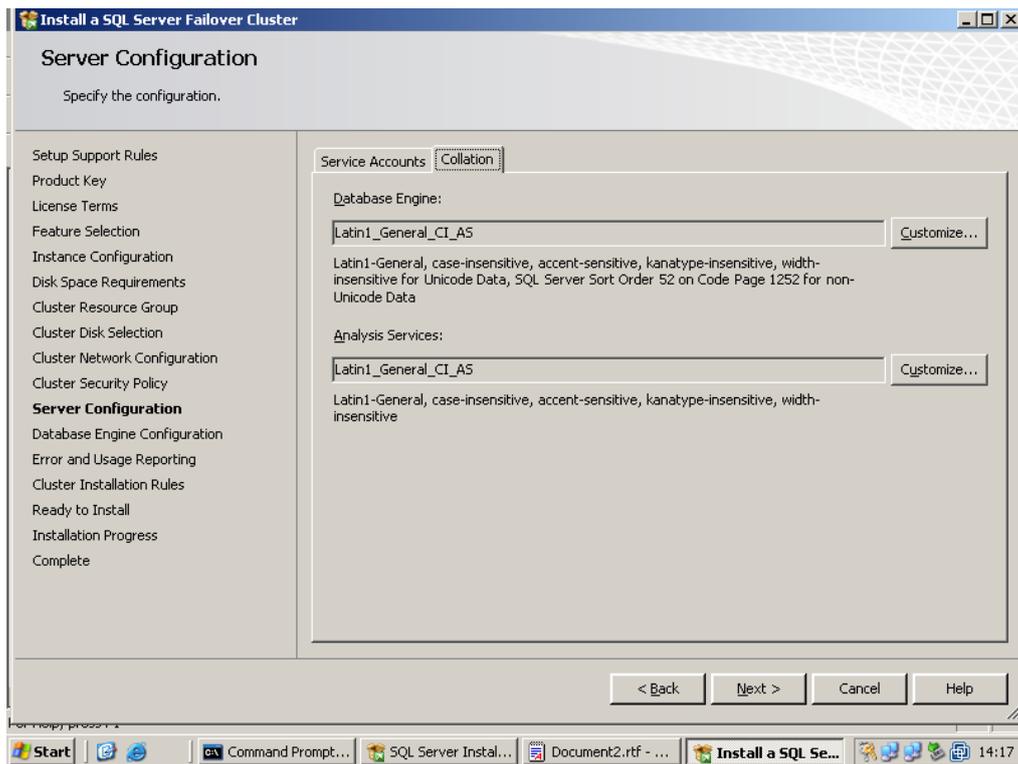
Supply the security group resources for the clustered SQL Server services and click “Next”,



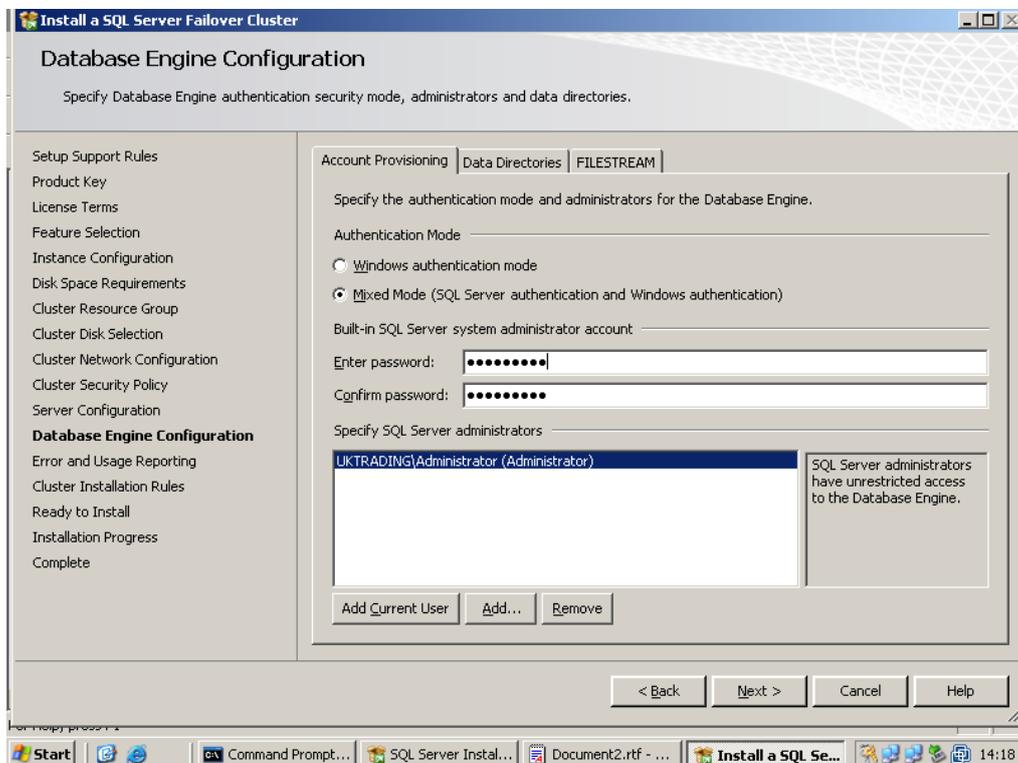
Supply your service account details and passwords and click the “Collation” tab,

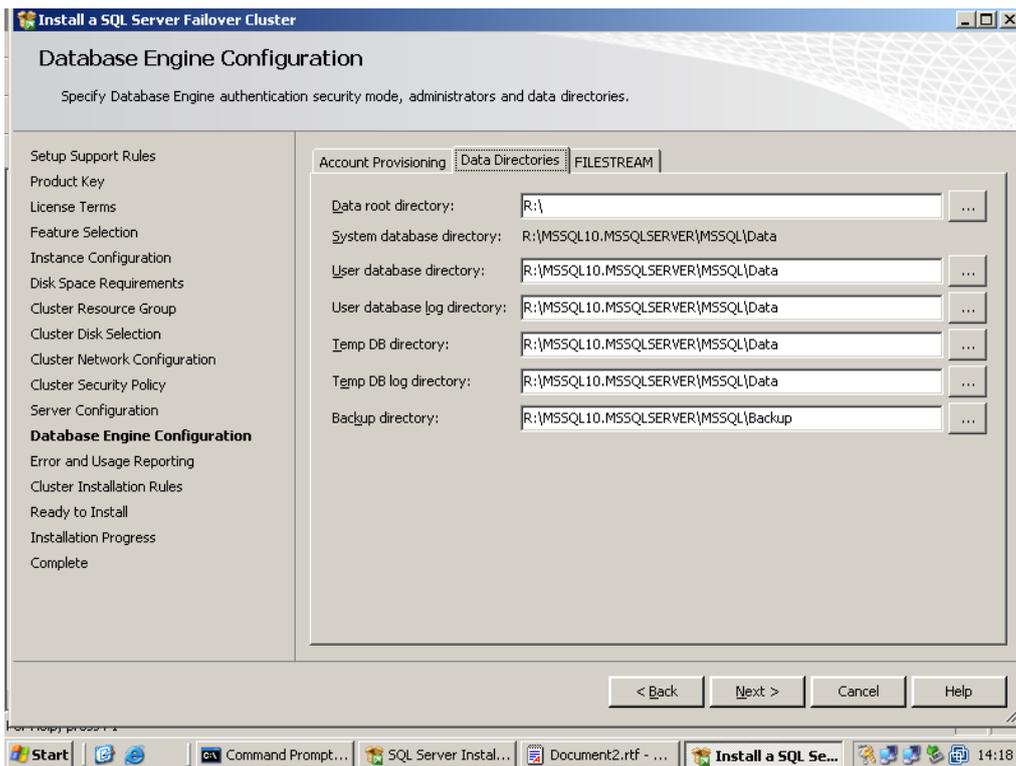


Select the collation and click “Next”,

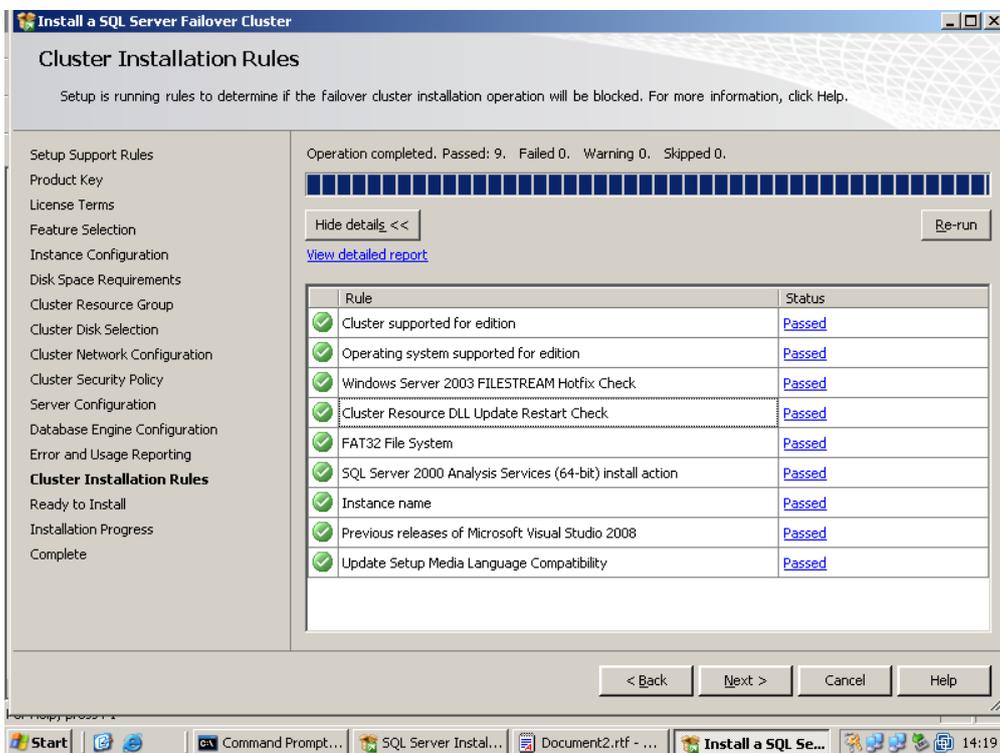


Provision administrator accounts and any data directories then click “Next”,

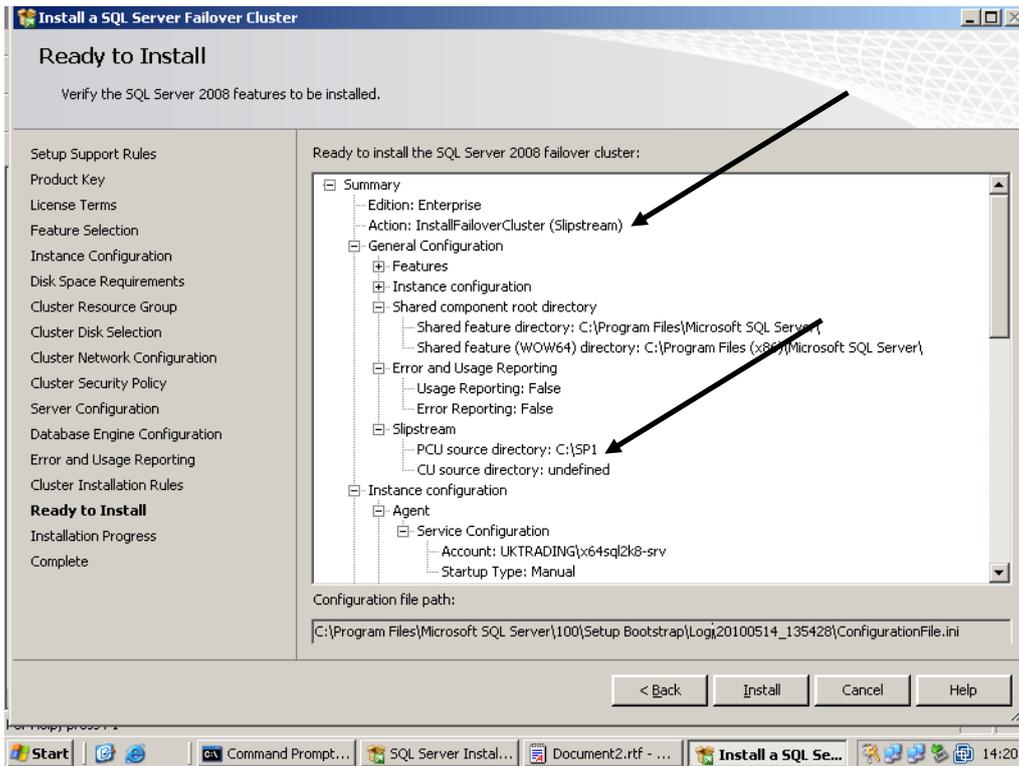




Review the cluster installation rules and click “Next”,



Click “Install” to create the clustered instance of SQL Server 2008. The arrows below indicate that this installation is slipstreaming Service Pack 1.



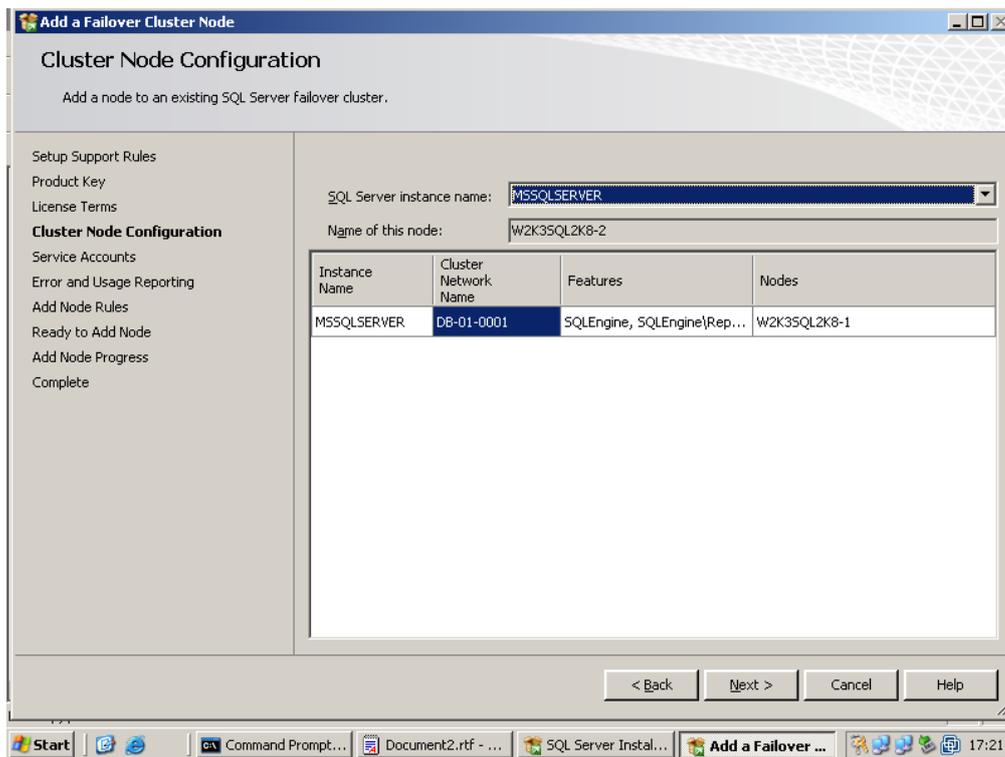
When the installation completes click “Close” to finish and you will be returned to the SQL Server 2008 splash screen. Close the splash screen and the active setup in the command window will complete, you may now close the command window.

## 4.2 ADD A CLUSTER NODE TO AN EXISTING CLUSTERED INSTANCE

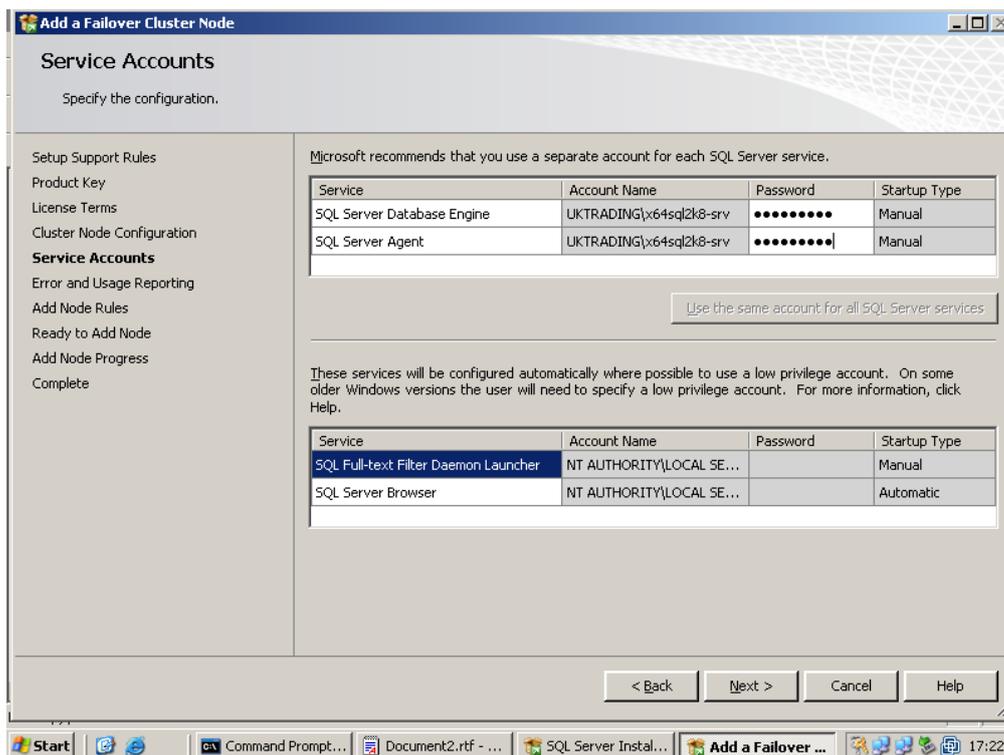
Before launching the installer you will need to perform exactly the same steps as you carried out for the installation of the first node (pre reqs and slipstream initiation). The setup will slipstream the service pack to ensure a smooth installation.

To add a node to an existing instance, select “Add node to a SQL Server failover cluster”. The Add cluster node wizard is exactly the same as the initial install up to the point where you specify licence key details. From here the installation continues as shown below.

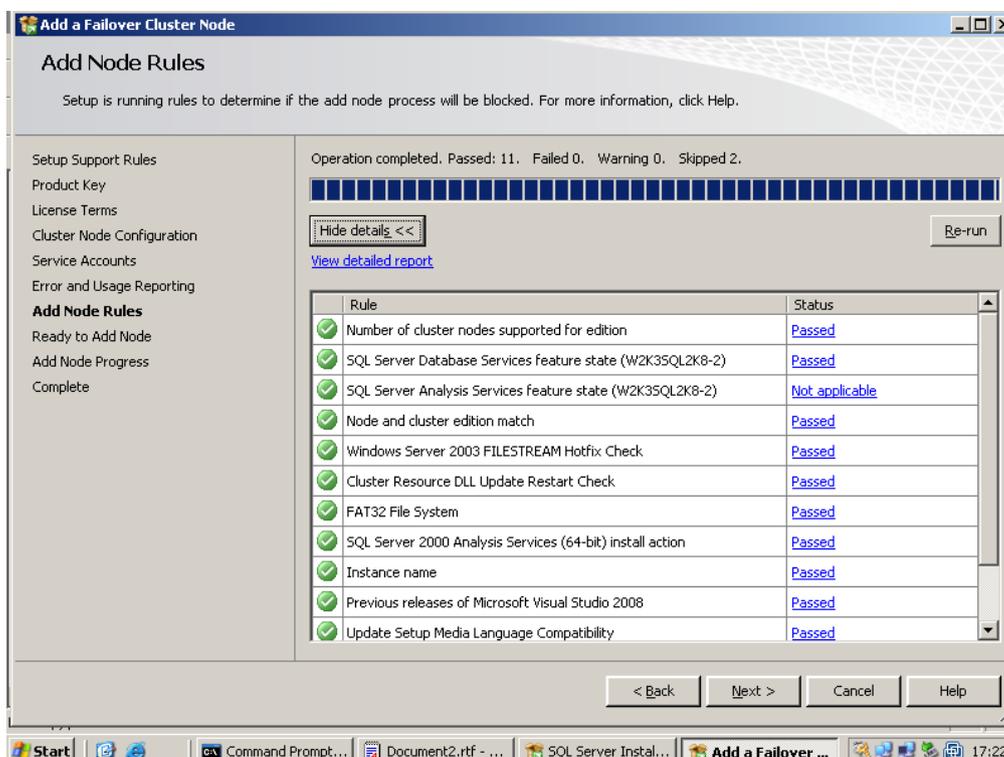
Confirm the instance you want to join and click “Next”,



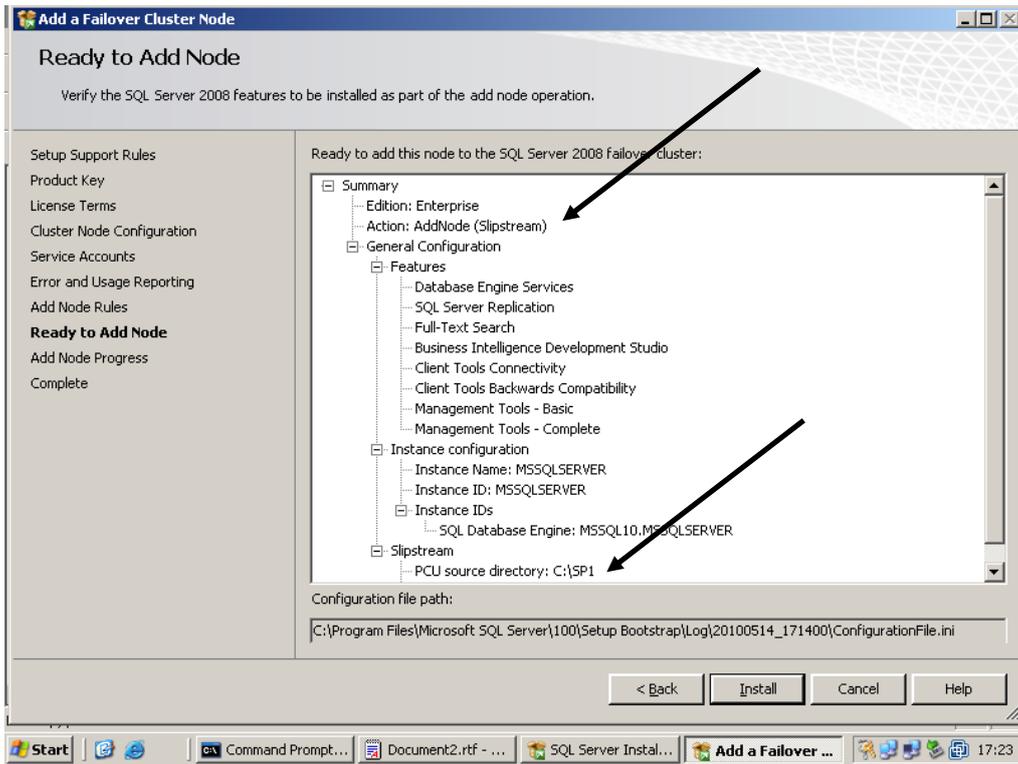
Supply any passwords and click “Next” to continue,



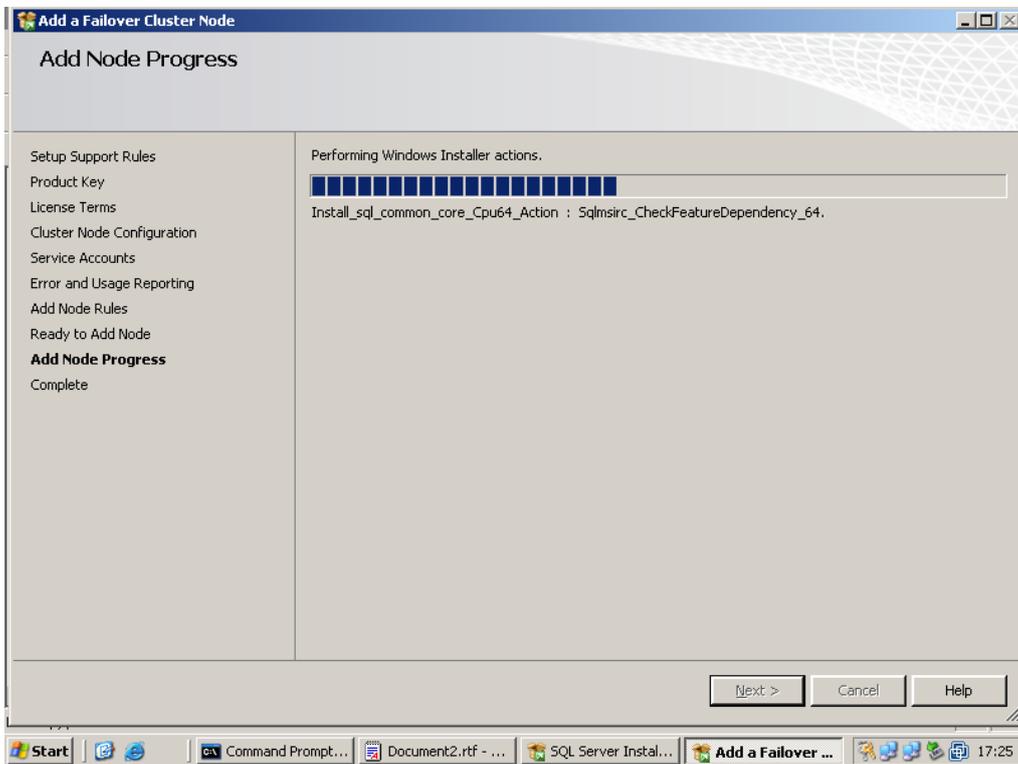
Review the install rules and click “Next”,



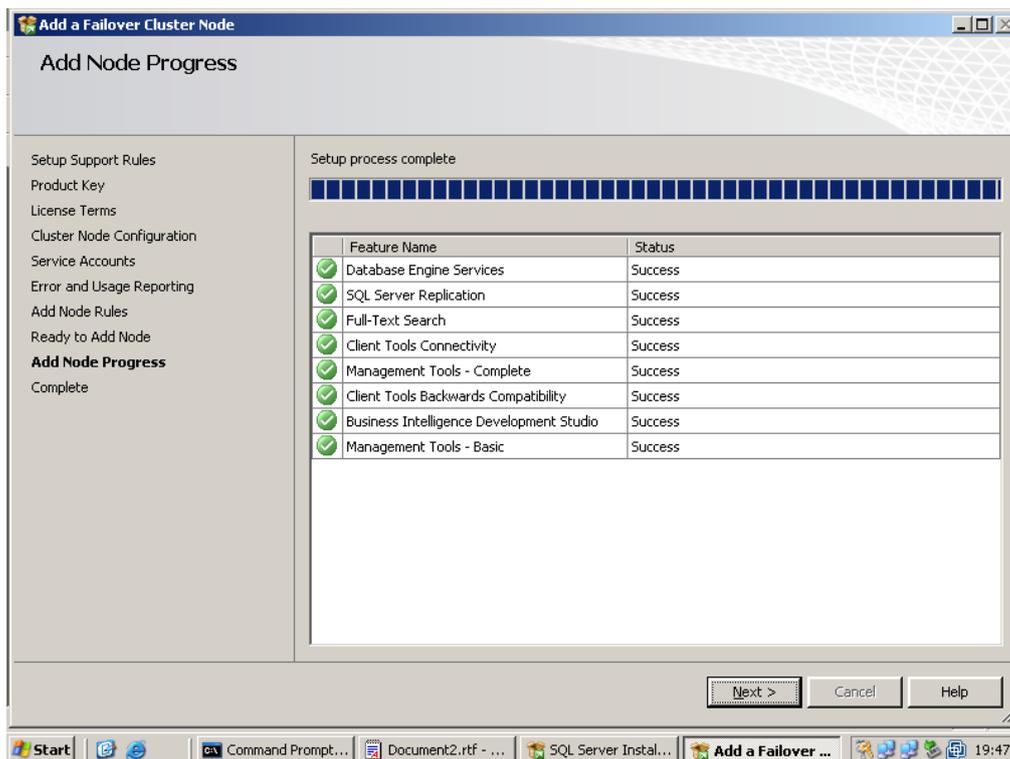
To add the cluster node click "Install". The arrows indicate that this installation is slipstreaming the service pack.



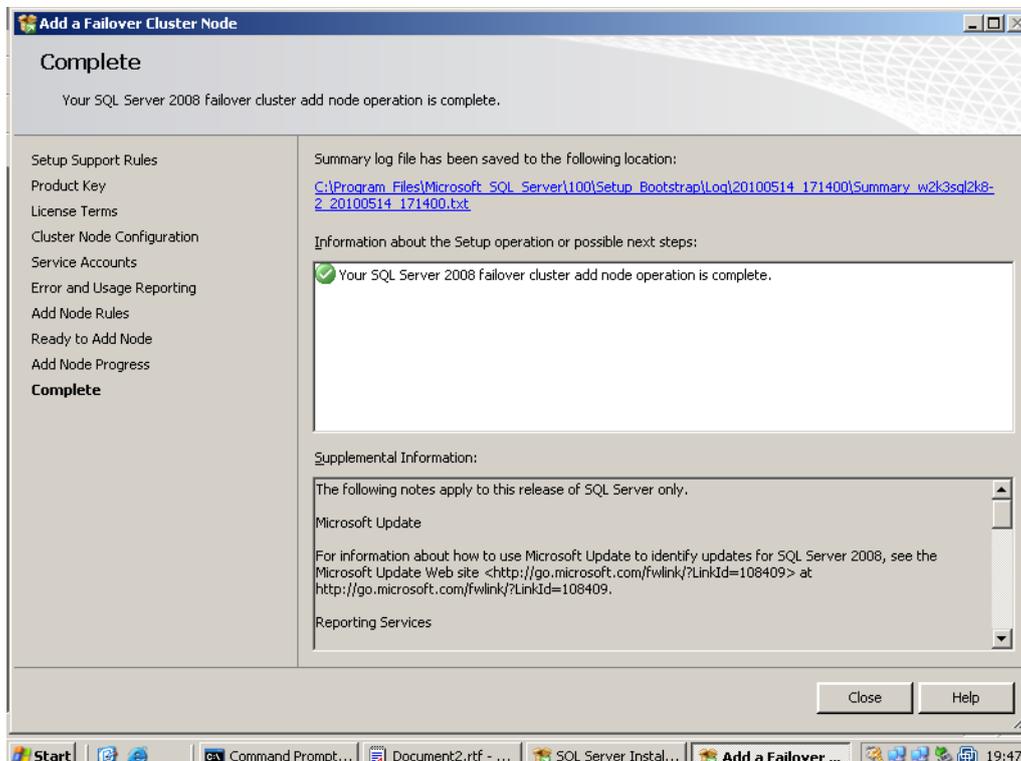
Setup progress!



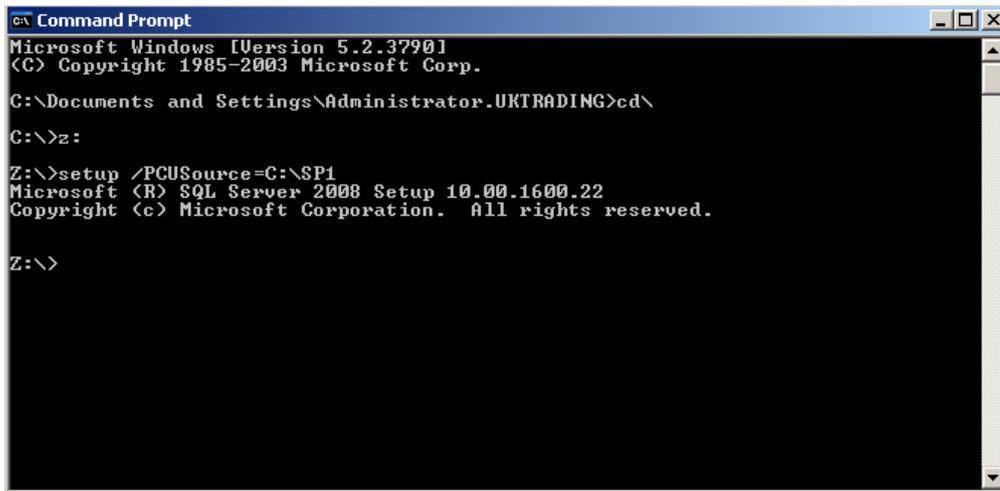
Click "Next" to complete setup.



Click "Close" to finish.



The command window completes. The node has now been added to the failover cluster.



```
Command Prompt
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

C:\Documents and Settings\Administrator.UKTRADING>cd\
C:\>z:
Z:\>setup /PCUSource=C:\SP1
Microsoft (R) SQL Server 2008 Setup 10.00.1600.22
Copyright (c) Microsoft Corporation. All rights reserved.

Z:\>
```

## 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.