



Virtual SAN

Installation Guide

Table of Contents

Introduction	3
Pre-requisites	3
Installation	4
Preparing for the install	4
Importing the appliance	4
XVS appliance configuration steps	15
Connecting ESX hosts to the XVS appliance volume	22
Planned XVS appliance node shutdown	29
Recovering from an XVS appliance node shutdown or failure.....	29

1 Introduction

This paper offers step-by-step instruction on configuring the Virtual SAN (XVS) appliance for use with VMware ESX3 Server to provide shared VMFS storage without a SAN – the aim being to utilise otherwise unused local storage in the ESX server to facilitate enterprise level features such as vMotion, DRS and HA normally only available through the use of a shared storage device.

Throughout this paper the non-routable range of 192.168.1.0/24 is used, your ESX VMs must have a Service Console and VMkernel Port Group in this range that can ping your iSCSI Server VM. You do not have to use the suggested range (and can use real IPs if you wish) but your ESX hosts must have at least one Service Console and one VMKernel portgroup on the same subnet as the XVS appliances.

1.1 Pre-requisites

Two VMware ESX servers licensed to use iSCSI storage are required as a minimum. No guidance will be provided in this paper on the installation and configuration of ESX server, please see the VMware ESX 3 Server Installation Guide.

http://www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_installation_guide.pdf

To gain all benefits of this appliance, VI3 Enterprise is recommended.

Sufficient local disk space formatted as VMFS must be available in the ESX hosts. 2.5Gb of free space is required for the XVS appliance, any additional space will be used as your Virtual SAN volume so a useful minimum of 20GB is recommended.

The XVS appliance requires 512MB of RAM, this may need to be increased to 1024MB in disk intensive environments

2 Installation

2.1 Preparing for the install

Download the XVS appliance from phdvirtual.com and extract the files to a folder on your desktop.

TIP: WinRAR can be used to extract .tar.gz files. <http://www.win-rar.com/downloadwinrar.html>

Download and install VMware Converter 3.0.2 Update 1. Note: This step can be skipped if you are familiar with uploading Virtual Machines to your servers using SFTP/SCP and registering them manually.

<http://www.vmware.com/products/converter/>

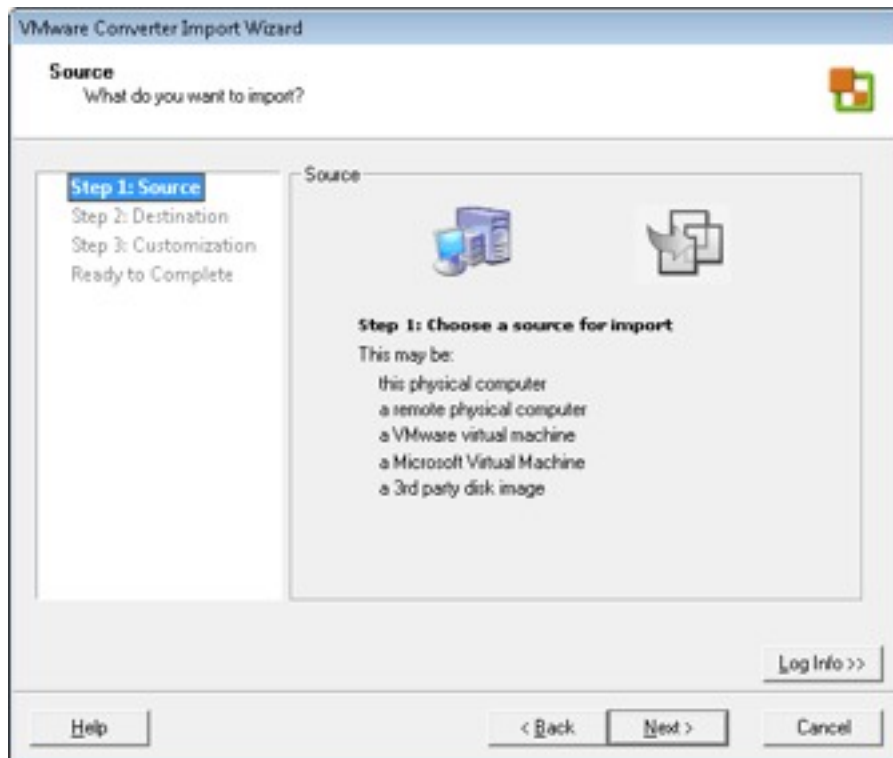
NOTE: Version 3.0.2 Update 1 is required, version 3.0.2 has an error which prevents it from importing the appliance correctly.

2.2 Importing the appliance

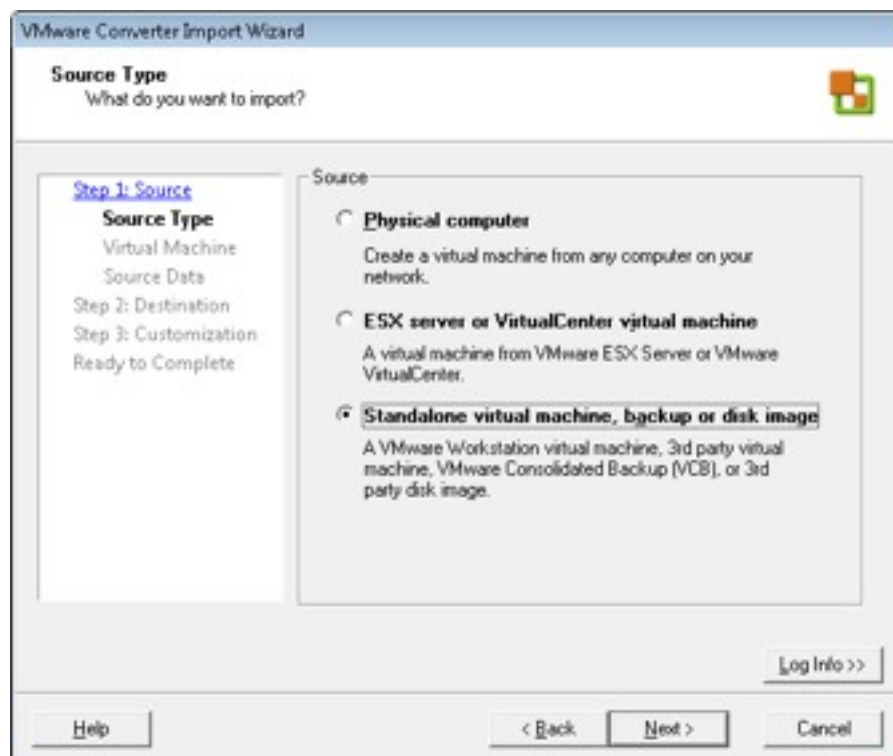
Open VMware Converter and start the Import Wizard.




Click **Next** to select a source for import.



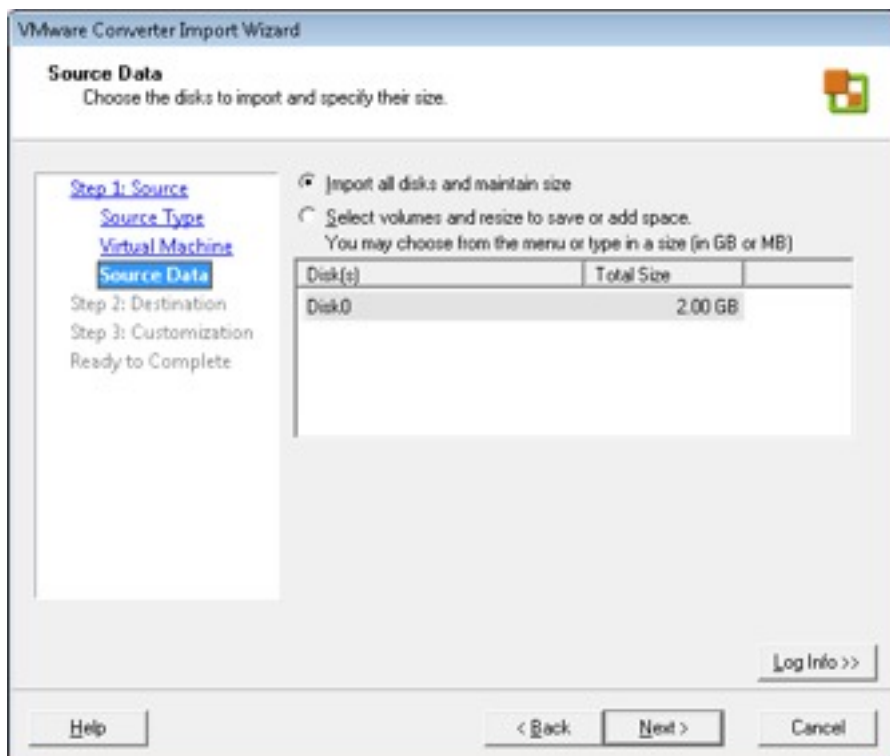
Select **Standalone virtual machine** and click **Next**.



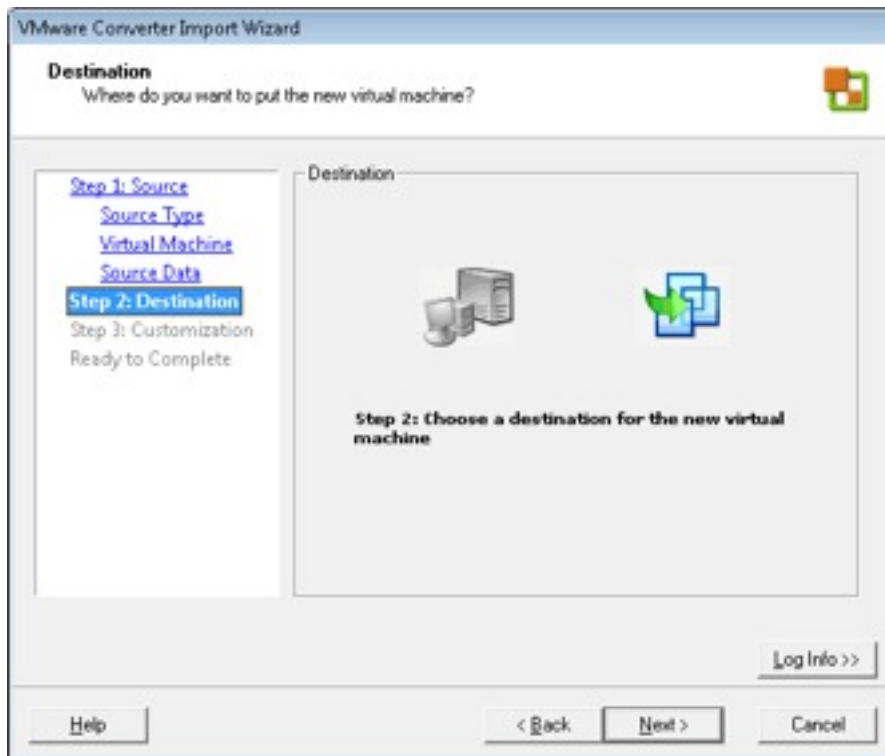
Click **Browse**, browse to the directory you extracted the XVS appliance to, select XVS.vmx. Click **Next**.



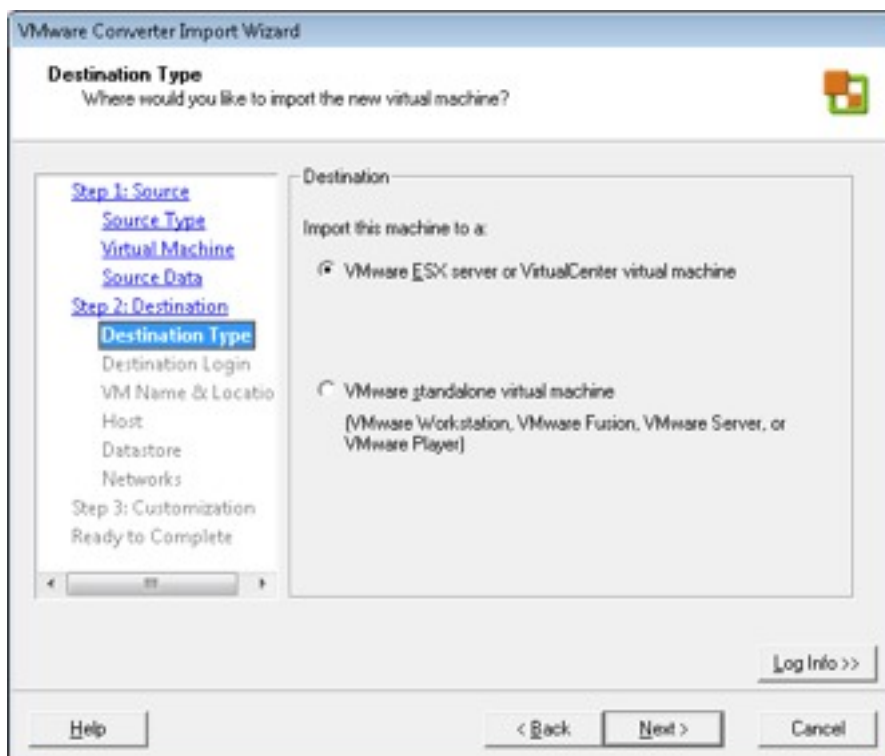
Leave the default **Import all disks and maintain size** selected and click **Next**.



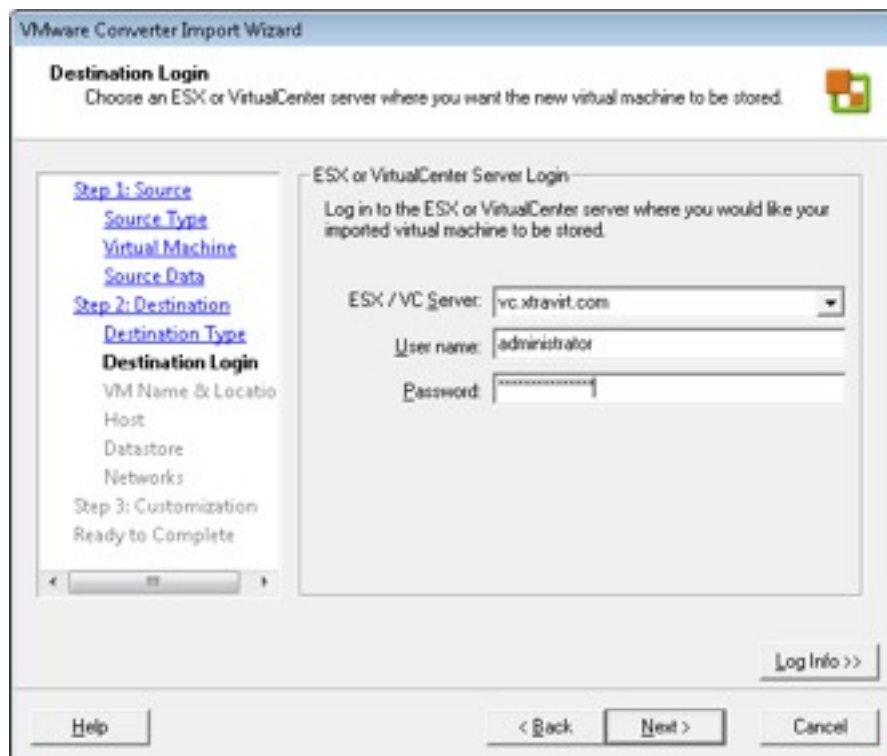
Click **Next** to select a destination host for the appliance.



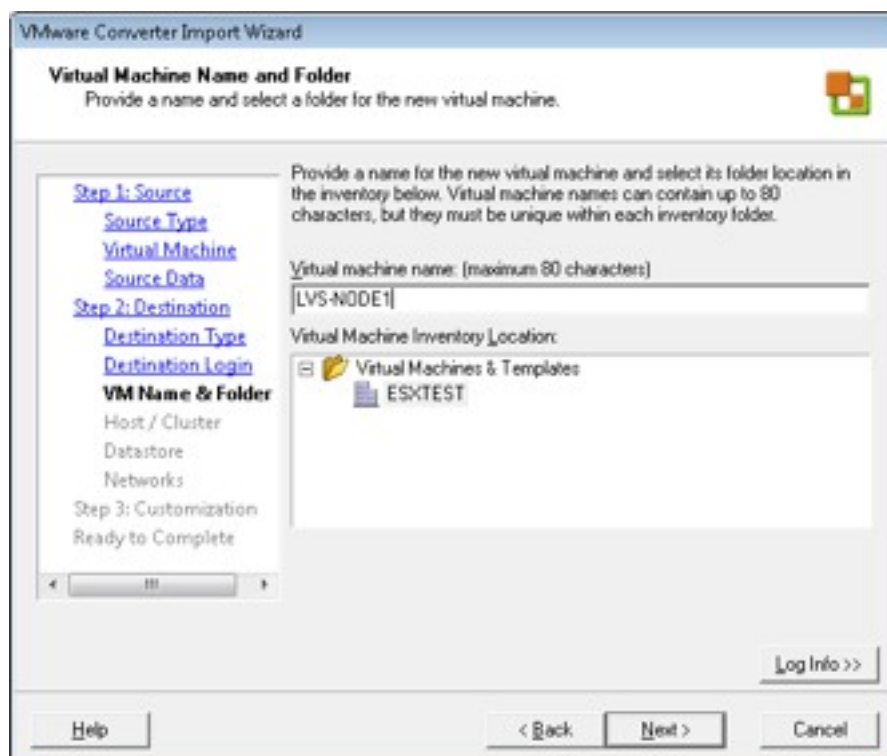
Select **VMware ESX server or VirtualCenter virtual machine** and click **Next**.



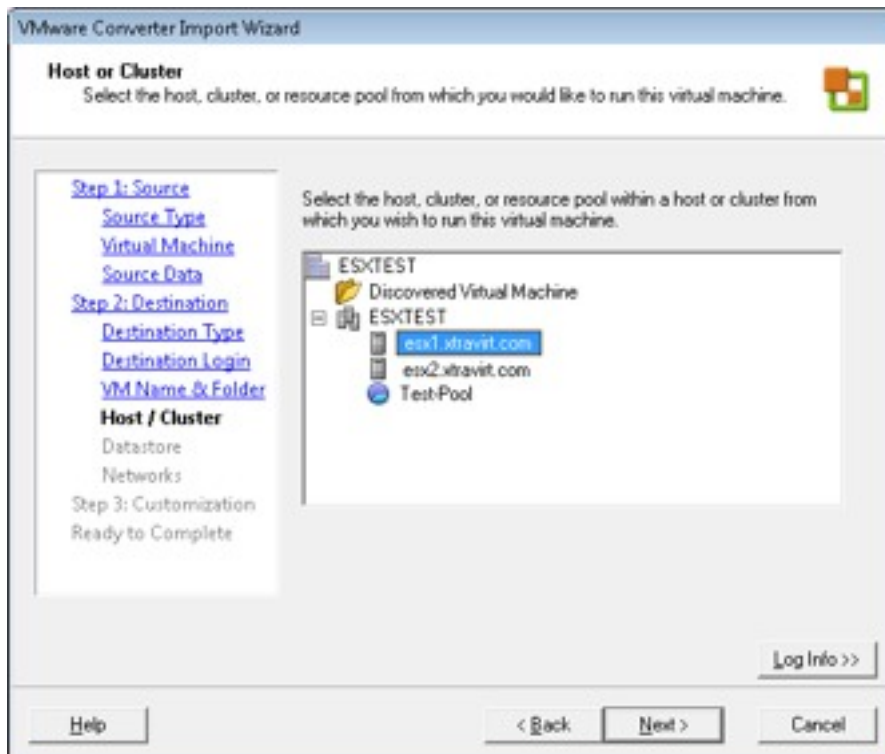
Enter the address and credentials for your VC server or ESX host and click **Next**.



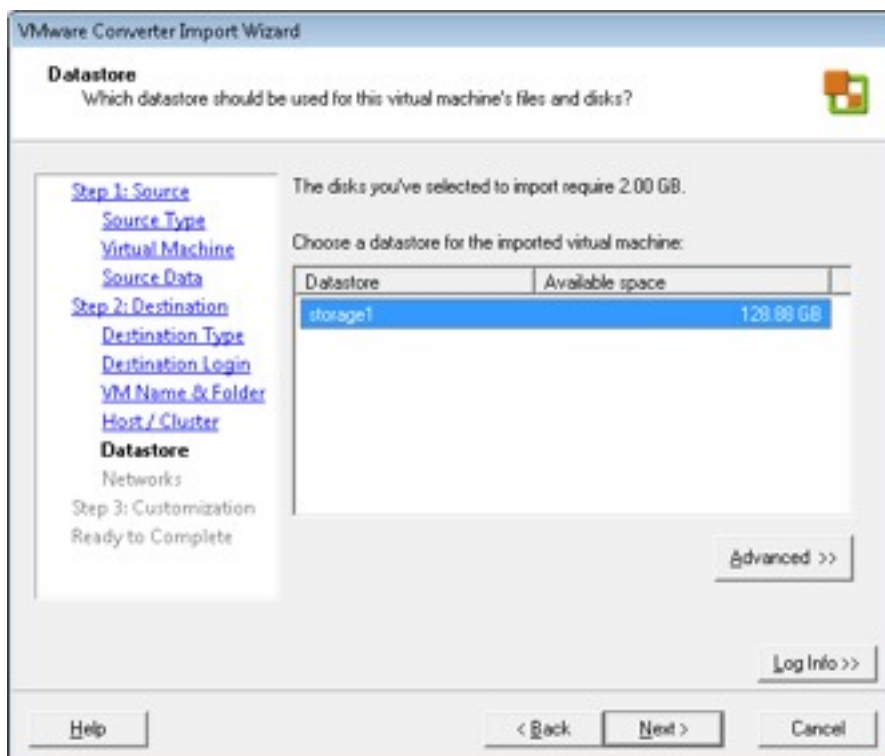
Give the XVS appliance a unique name. In this paper we use **XVS-NODE1** and **XVS-NODE2**. Then select a data centre for the appliance and then click **Next**.



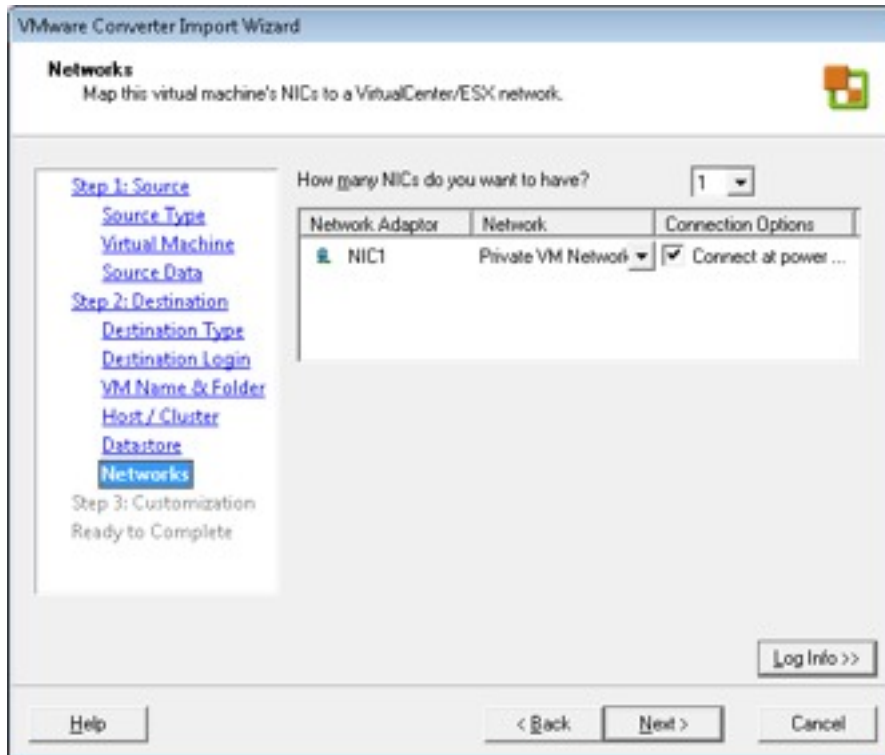
Select an ESX host for the first XVS node and click **Next**.



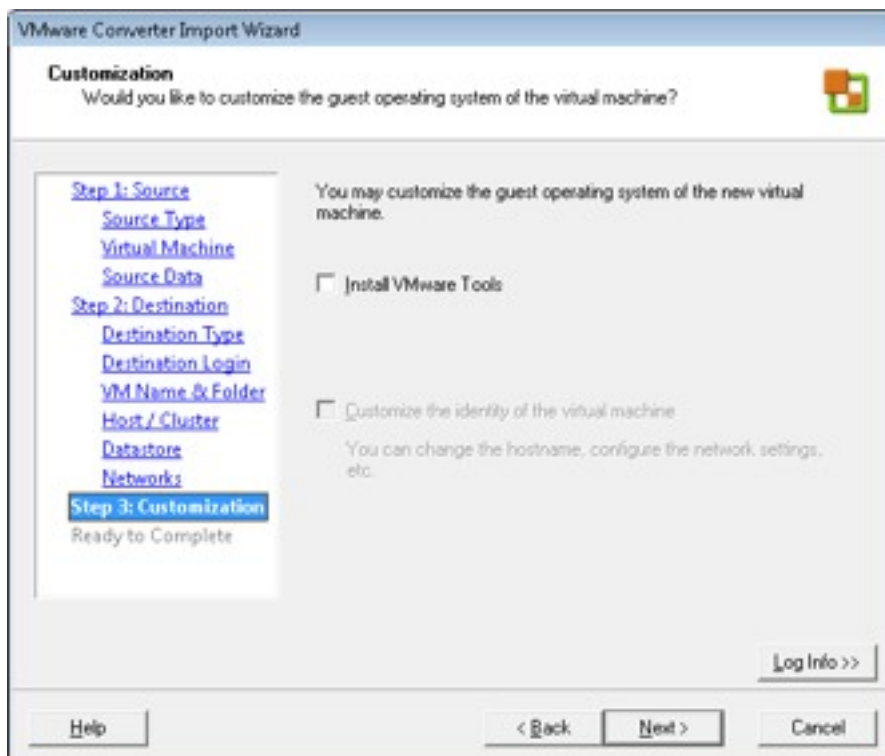
Choose the local storage VMFS volume on the host and click **Next**.



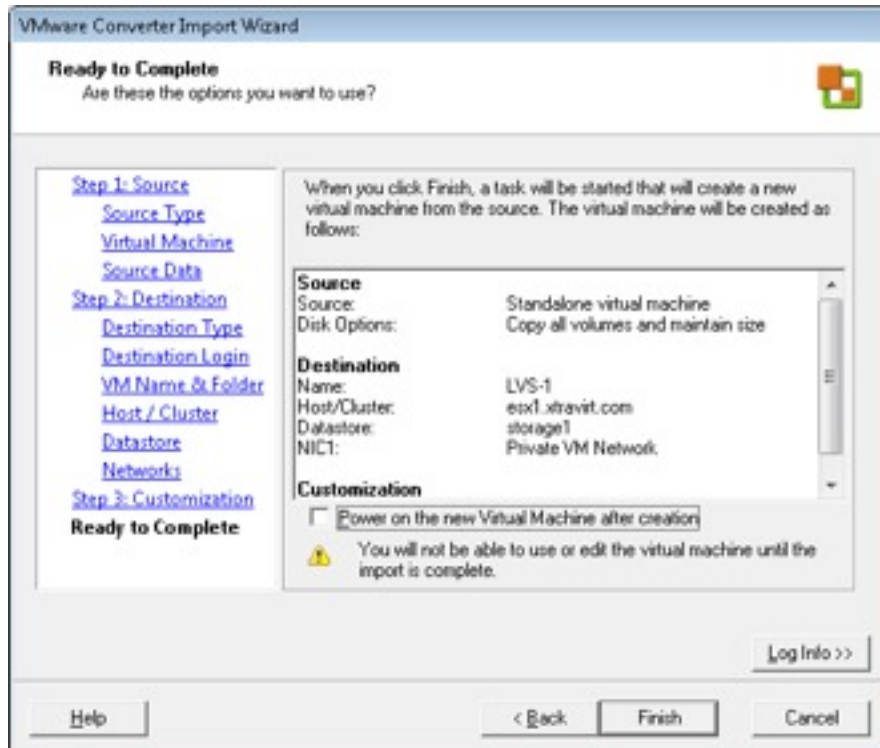
Select a network for the XVS appliance that is connected to a vSwitch with access to the appropriate subnet and click **Next**.



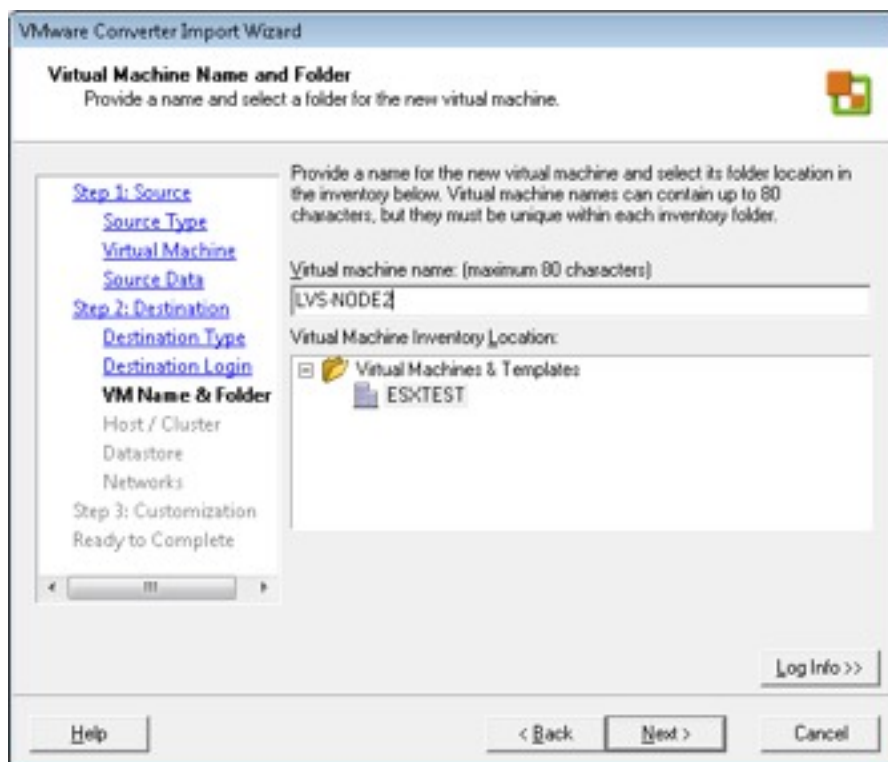
Leave **Install VMware Tools** un-selected, the VMware tools are pre-installed in the appliance.



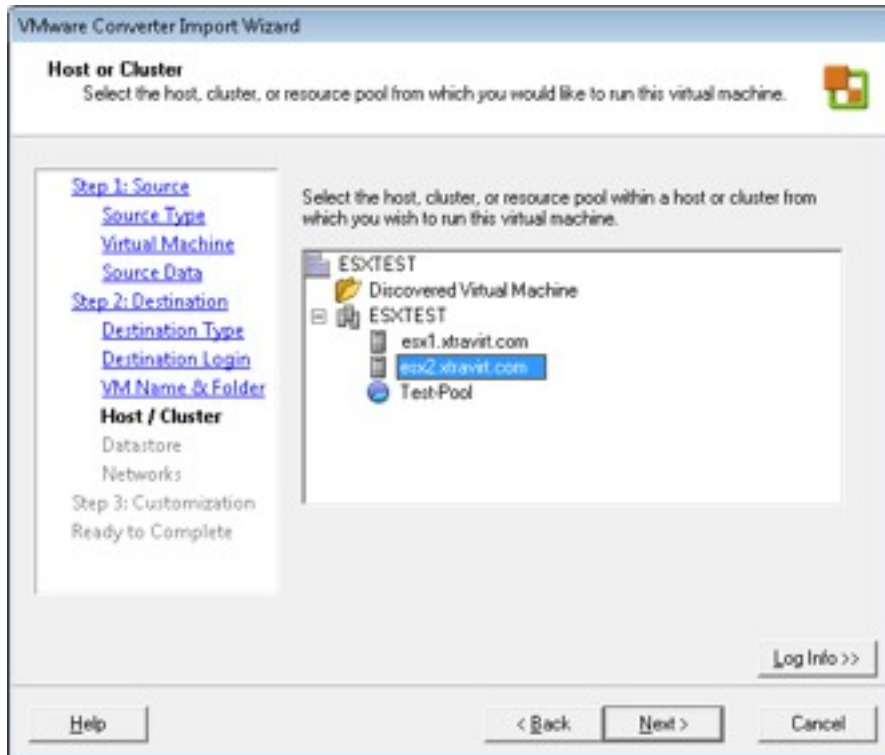
Leave **Power on the new Virtual Machine after creation** un-selected. It is necessary to add a second disk to the appliance before powering it on for configuration. Click **Next**.



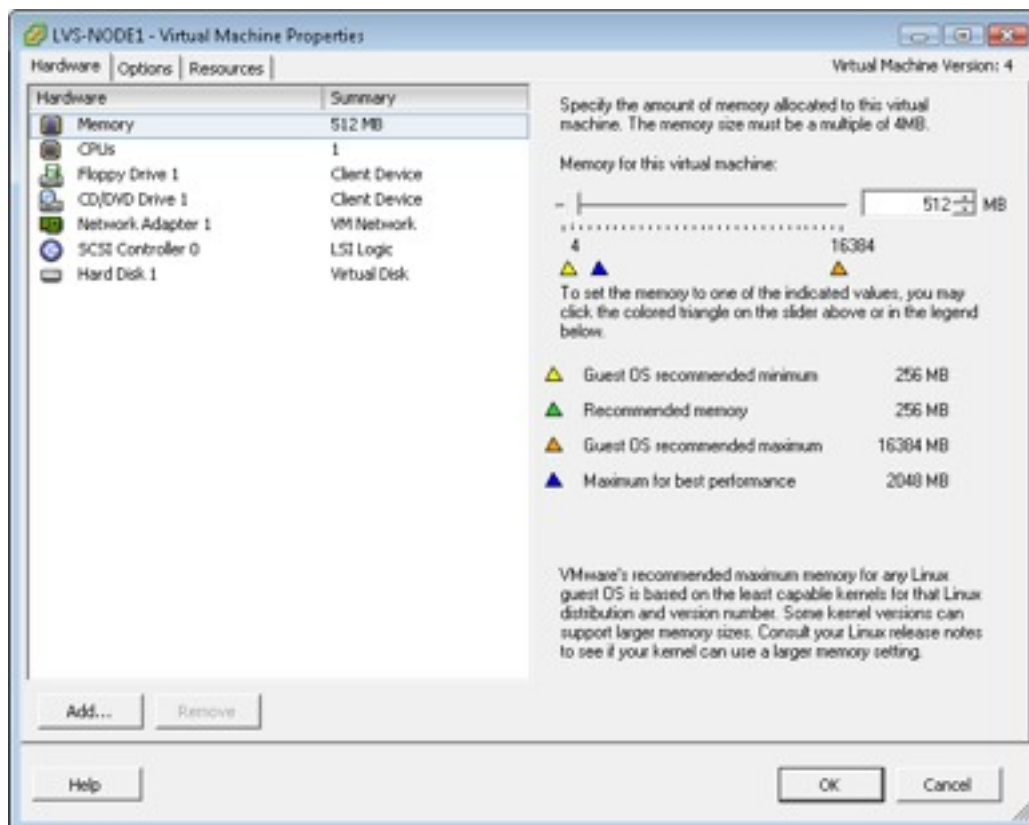
The appliance should now import. Repeat the process above for the second XVS node with the following changes. Give the second XVS appliance a unique name.



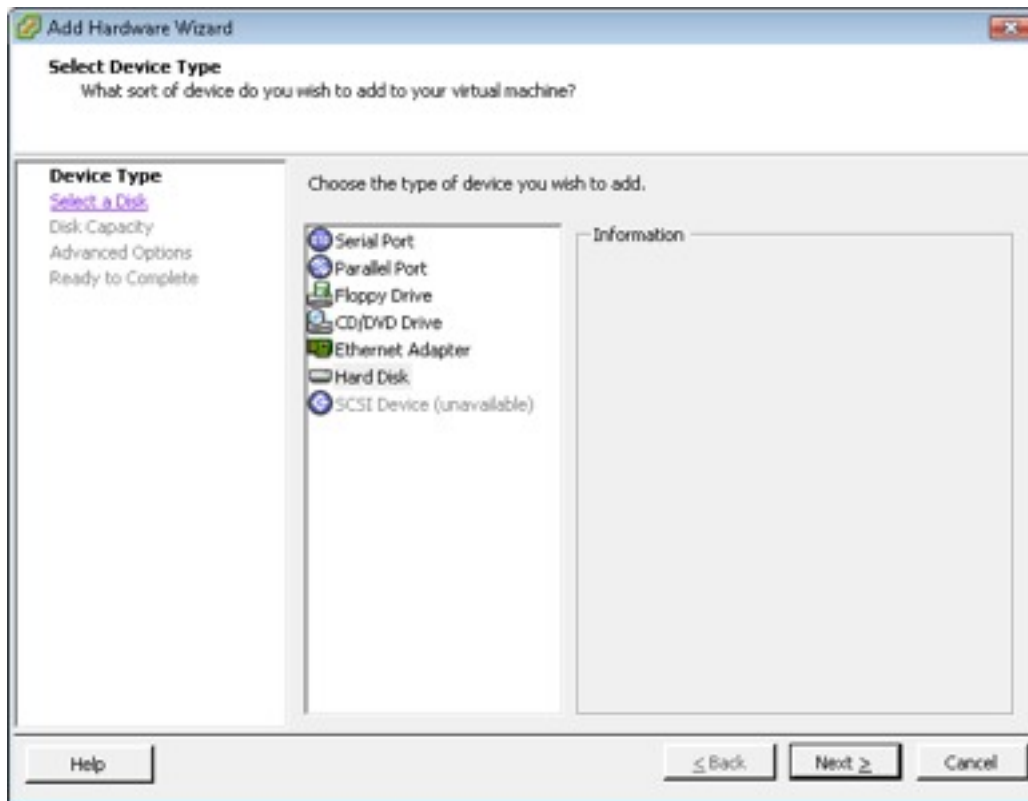
Make sure the second XVS appliance node is placed on a different ESX host to the first node.



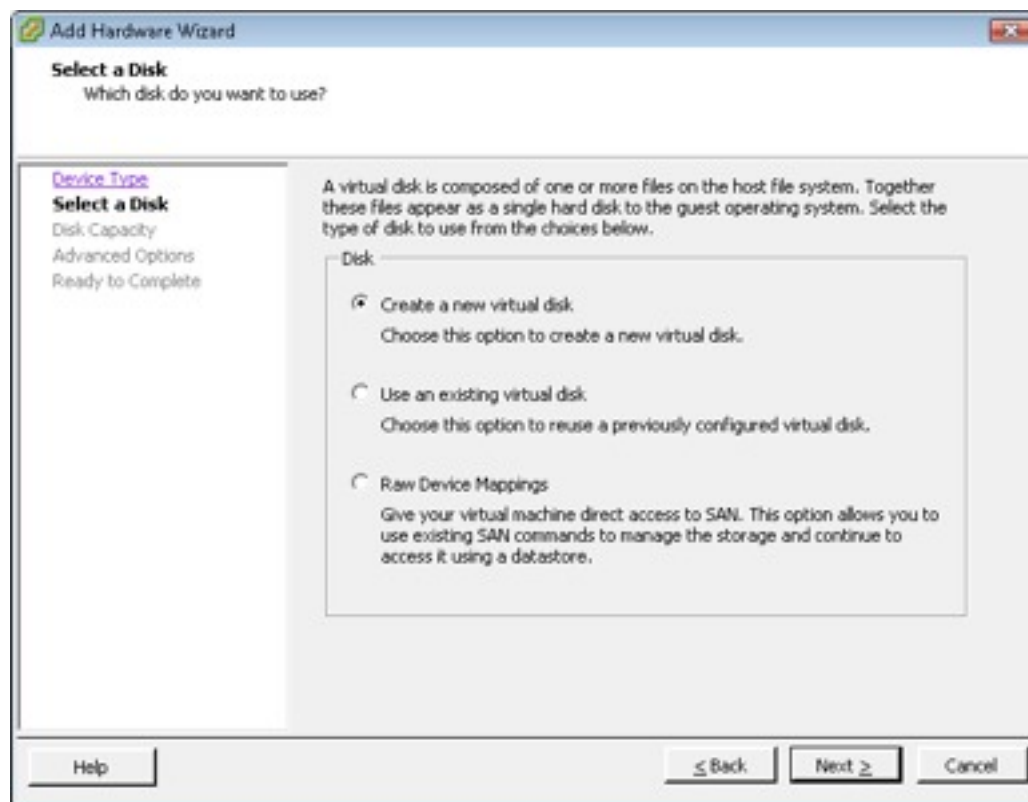
When the appliance is imported, **Edit Settings** on the VMs. Click **Add** to add hardware to the VM.



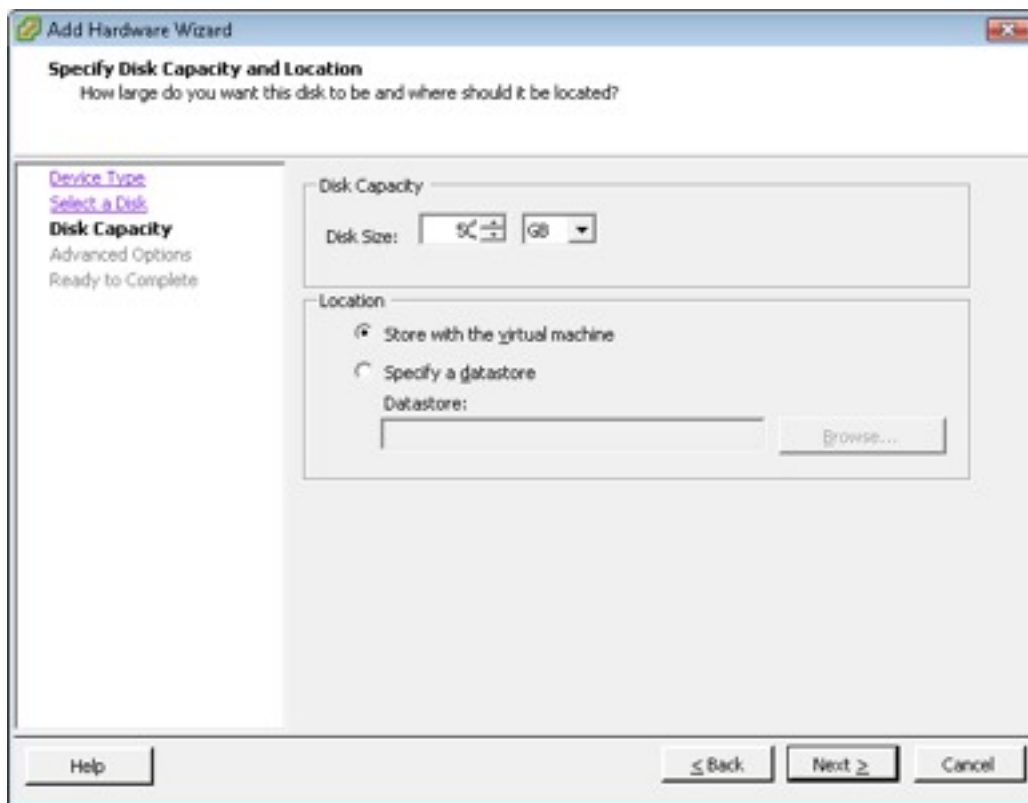
Select **Hard Disk** from the list of devices and click **Next**.



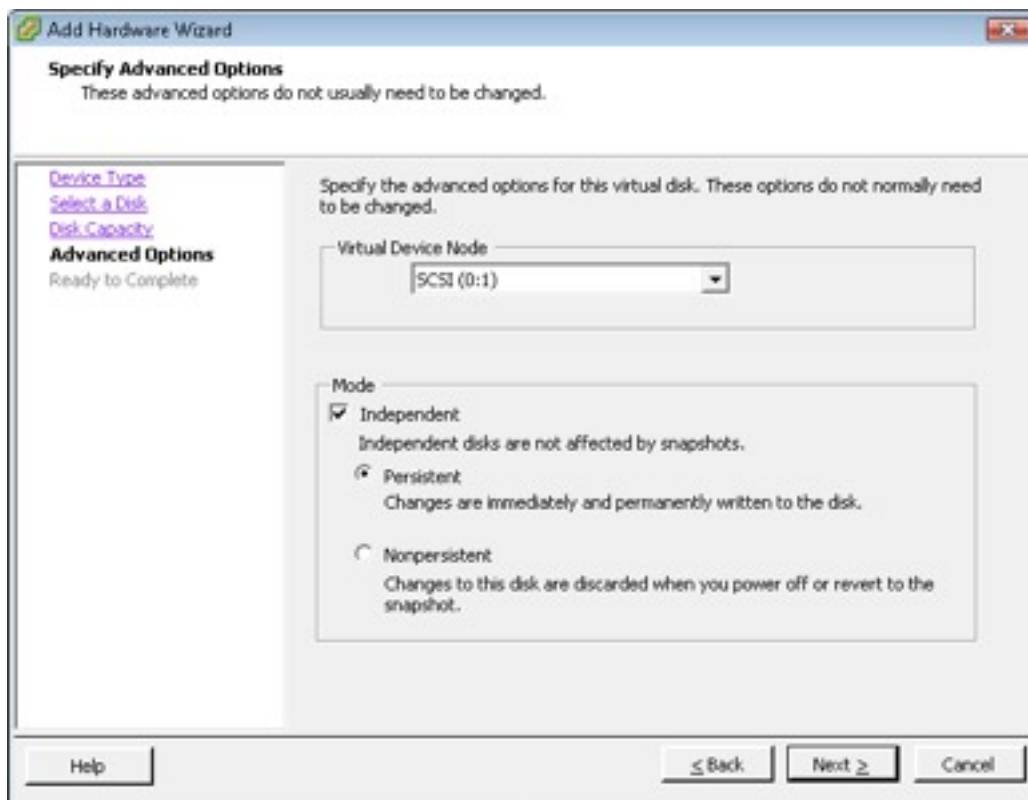
Select **Create a new virtual disk** and click **Next**.



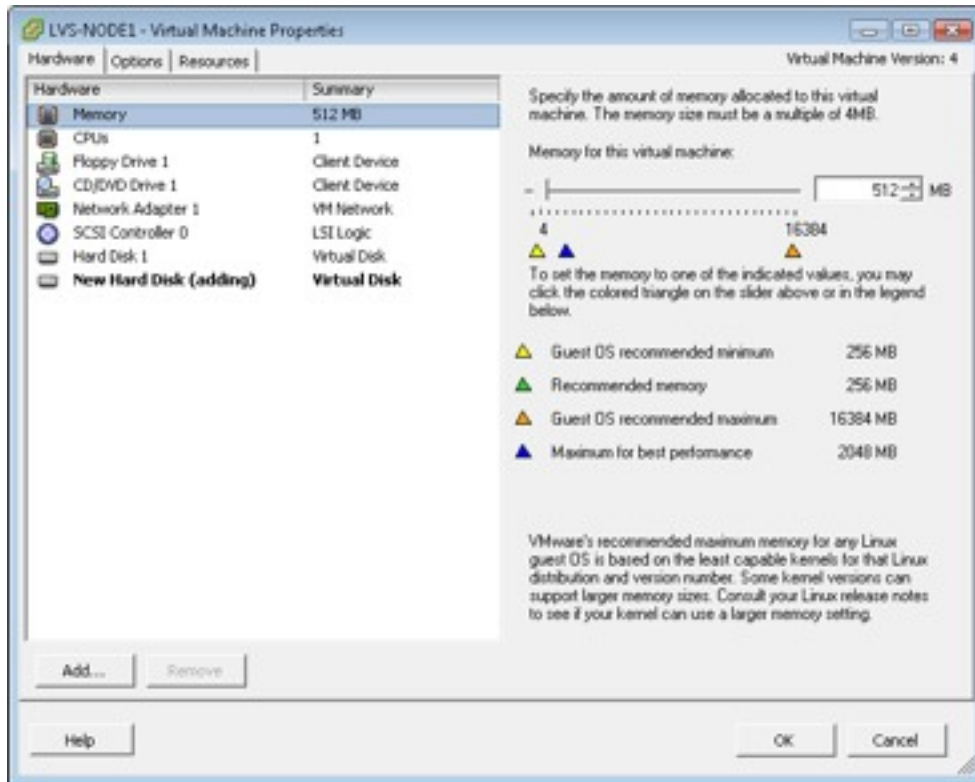
Specify the size of the SAN you would like to create, this can be any size <2Tb depending on available local disk space. Select **Store with the virtual machine** and click **Next**.



Leave the SCSI ID as the default (0:1). Select **Independent** and **Persistent** then click **Next**.



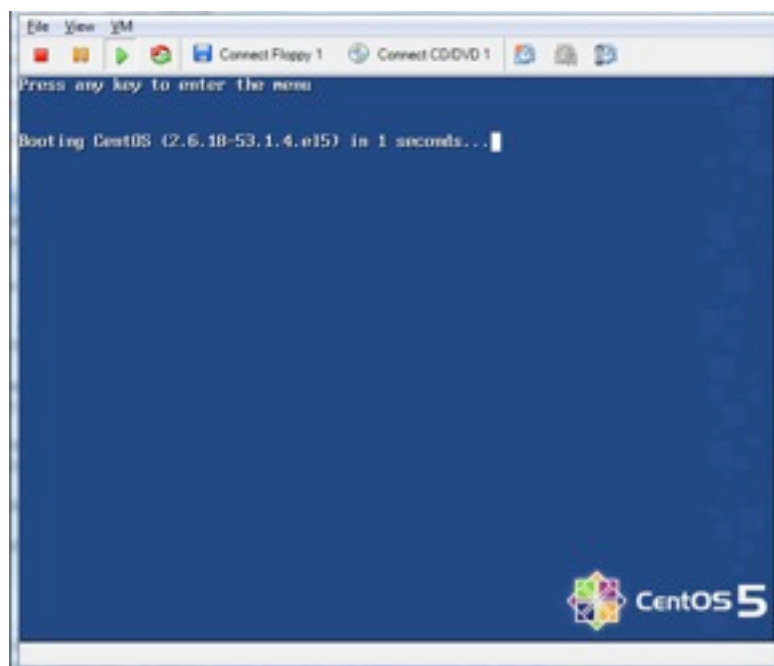
The new SAN disk drive should appear in the hardware list. Click **OK** to confirm the change.



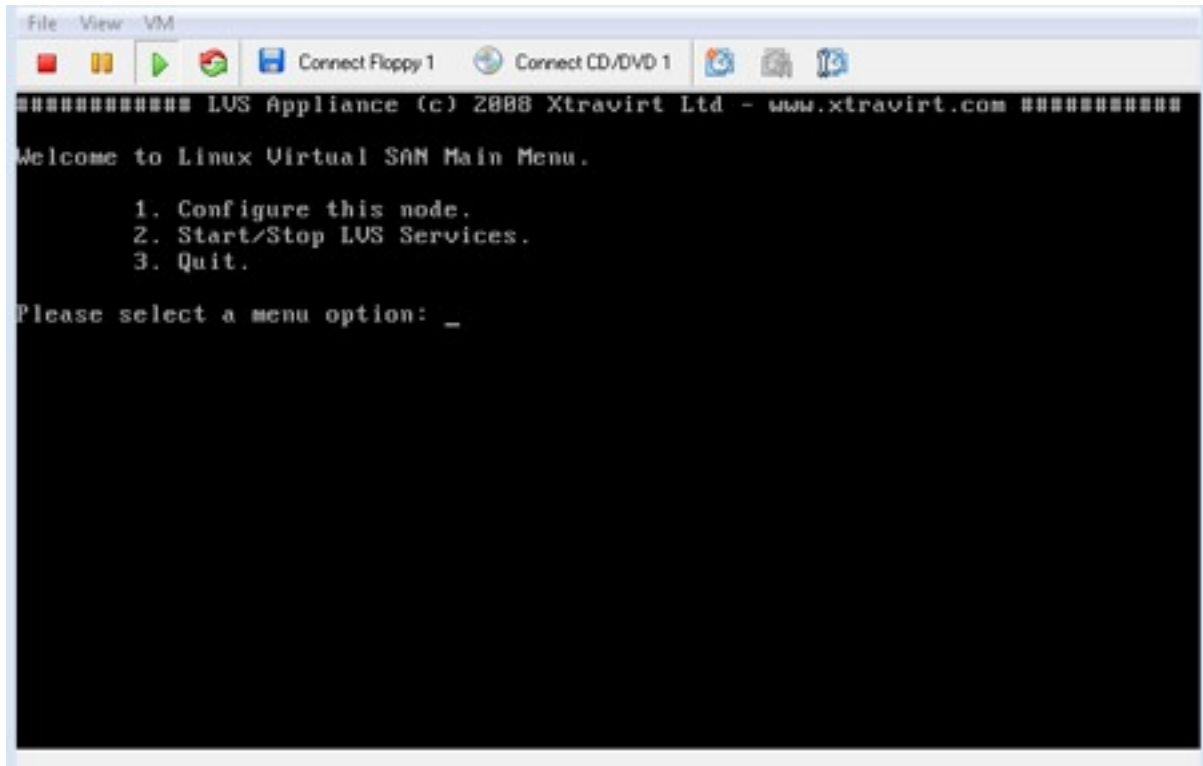
Repeat these steps for both appliance nodes. The import is now complete.

2.3 XVS appliance configuration steps

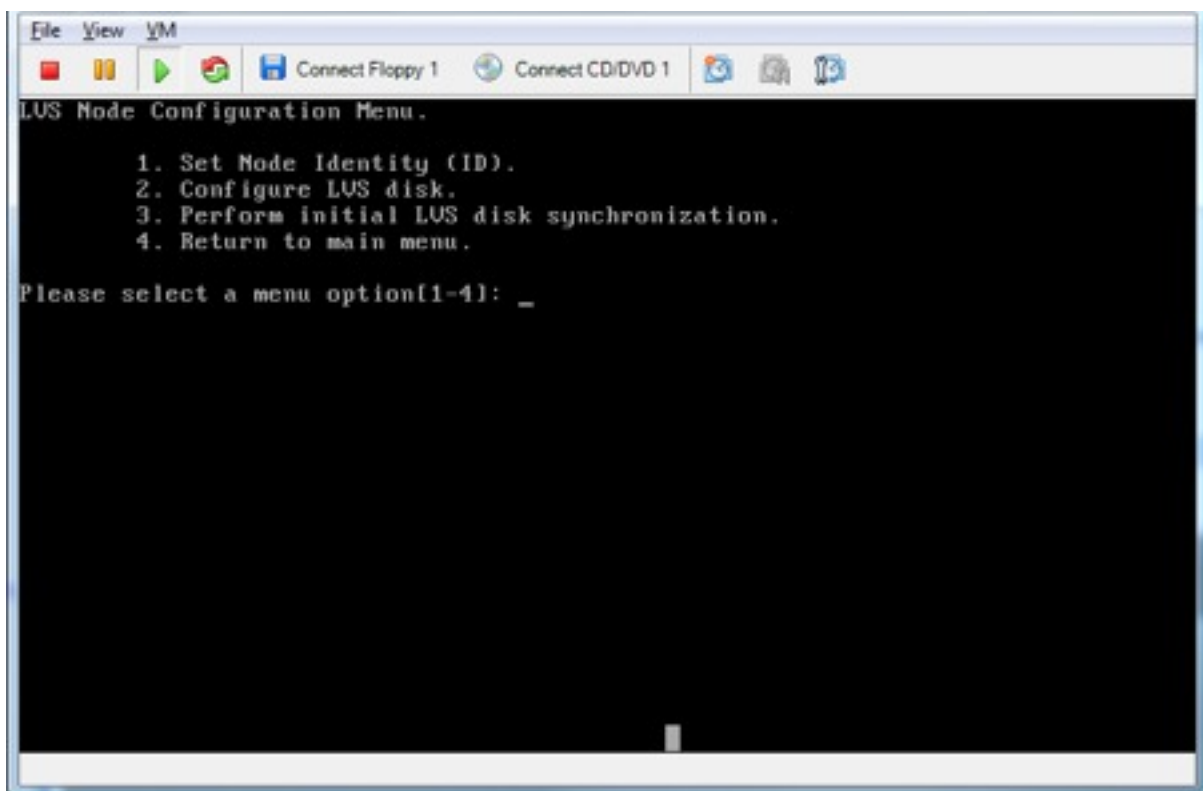
Power on both of the XVS appliance nodes and allow them to boot. Open a console to each. **NOTE:** Don't press a key at this point, just allow the default option to boot.



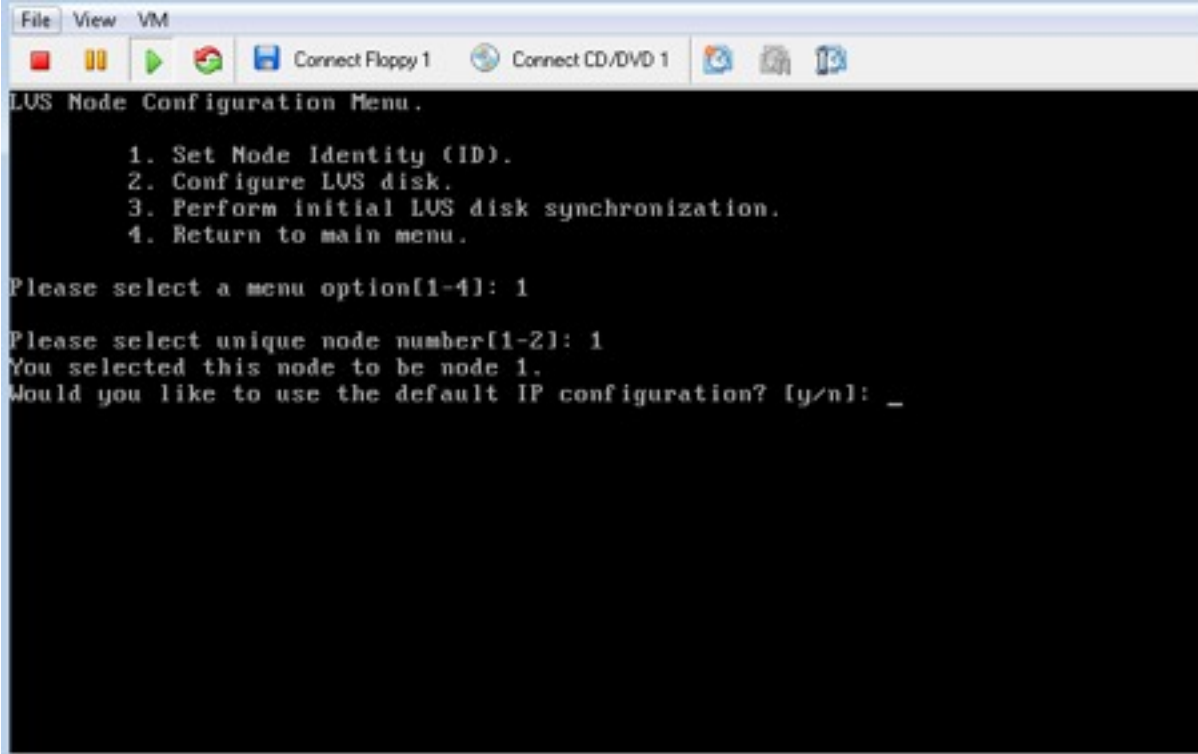
Once the boot process is complete, you will be presented with the XVS Main Menu at the console.



Select **Option 1** to enter the **Node Configuration Menu**.



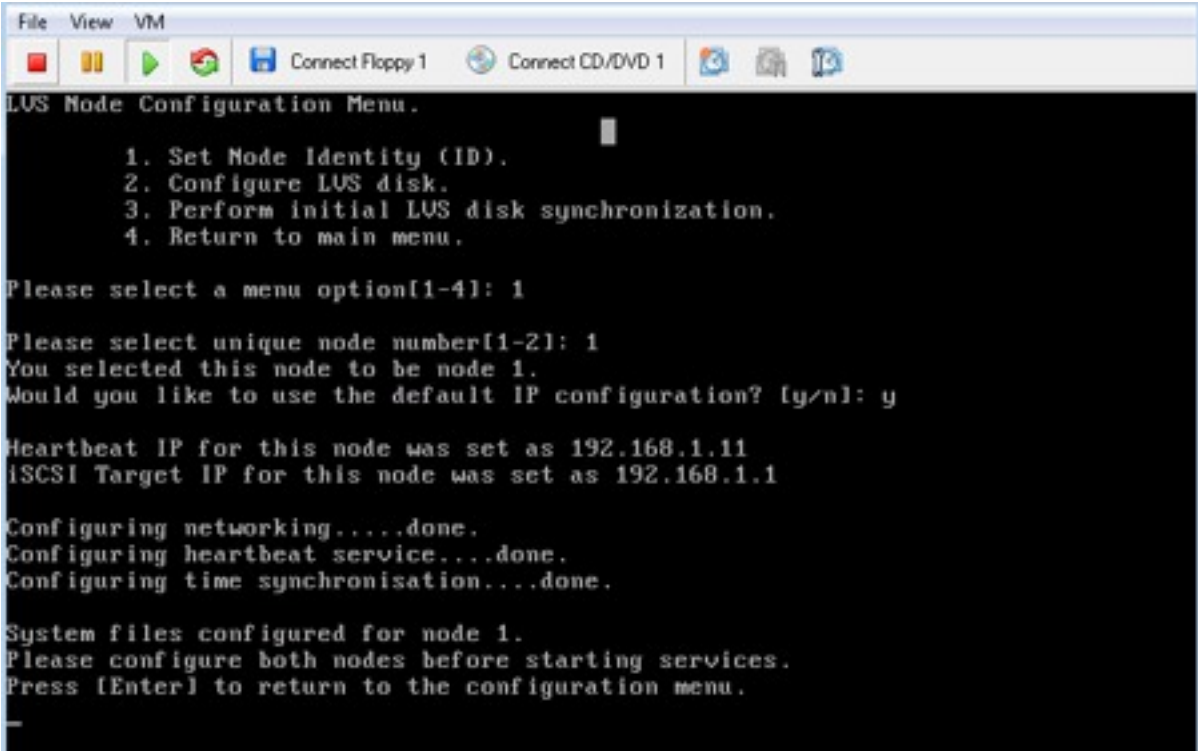
Enter the unique node number **1 or 2** for the node.



The screenshot shows a VMware window titled 'File View VM' with a toolbar containing icons for 'Connect Floppy 1' and 'Connect CD/DVD 1'. The main content is a terminal window with the following text:

```
LUS Node Configuration Menu.  
  
1. Set Node Identity (ID).  
2. Configure LUS disk.  
3. Perform initial LUS disk synchronization.  
4. Return to main menu.  
  
Please select a menu option[1-4]: 1  
  
Please select unique node number[1-2]: 1  
You selected this node to be node 1.  
Would you like to use the default IP configuration? [y/n]: _
```

Press **"y"** if you want to accept the default IP address configuration (details below).



The screenshot shows the same VMware window as above, but with the following text in the terminal window:

```
LUS Node Configuration Menu.  
  
1. Set Node Identity (ID).  
2. Configure LUS disk.  
3. Perform initial LUS disk synchronization.  
4. Return to main menu.  
  
Please select a menu option[1-4]: 1  
  
Please select unique node number[1-2]: 1  
You selected this node to be node 1.  
Would you like to use the default IP configuration? [y/n]: y  
  
Heartbeat IP for this node was set as 192.168.1.11  
iSCSI Target IP for this node was set as 192.168.1.1  
  
Configuring networking....done.  
Configuring heartbeat service....done.  
Configuring time synchronisation....done.  
  
System files configured for node 1.  
Please configure both nodes before starting services.  
Press [Enter] to return to the configuration menu.  
_
```

Default IP address configuration is as follows:

Heartbeat/replication IP addresses:

Node1 - 192.168.1.11

Node2 - 192.168.1.12

These IP addresses only need to be able to ping each other, they require no external connectivity.

iSCSI Target IP address:

Node1 – 192.168.1.1

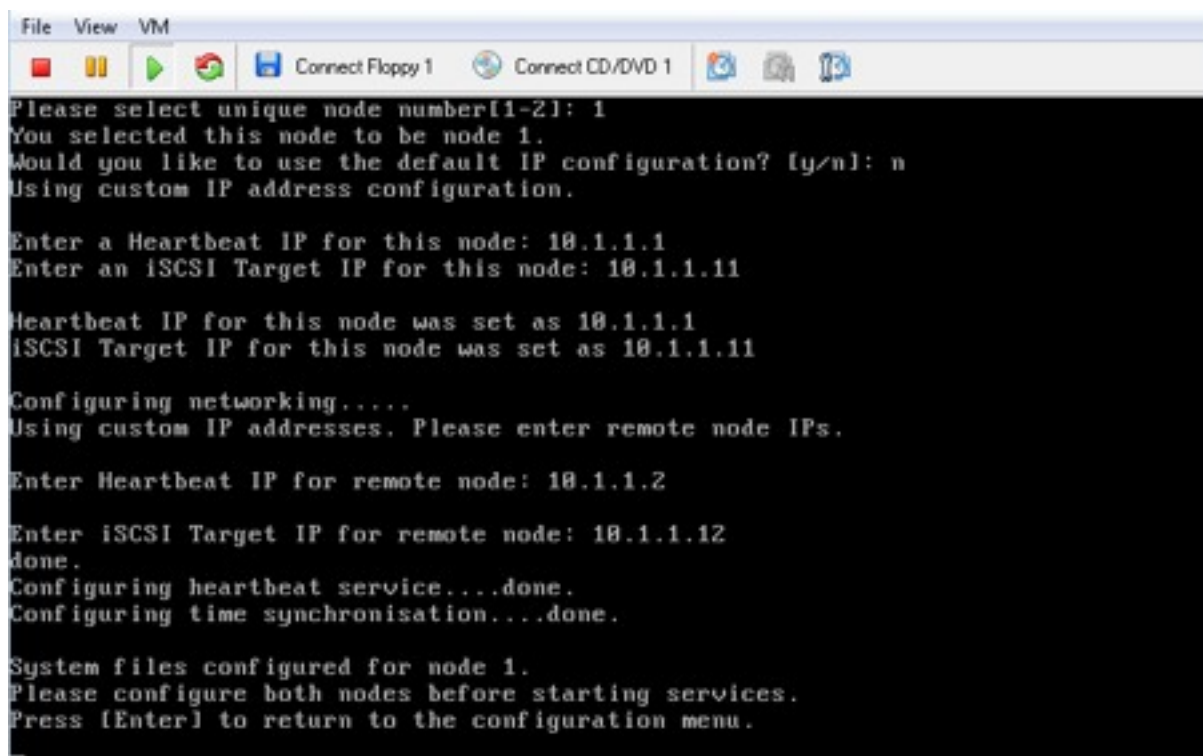
Node2 – 192.168.1.2

These IP addresses must be on the same subnet as at least one Service Console and one VMkernel Port Group on each ESX host. The shared storage will not be accessible without this.

NOTE: it is recommended to have the Heartbeat and iSCSI Target IP addresses on the same subnet.

If you wish to use custom IP addresses, enter **n** and then enter your desired IP addresses.

NOTE: The remote node IP addresses must match what you set on the other node or the services will not start.



```
File View VM
Please select unique node number[1-2]: 1
You selected this node to be node 1.
Would you like to use the default IP configuration? [y/n]: n
Using custom IP address configuration.

Enter a Heartbeat IP for this node: 10.1.1.1
Enter an iSCSI Target IP for this node: 10.1.1.11

Heartbeat IP for this node was set as 10.1.1.1
iSCSI Target IP for this node was set as 10.1.1.11

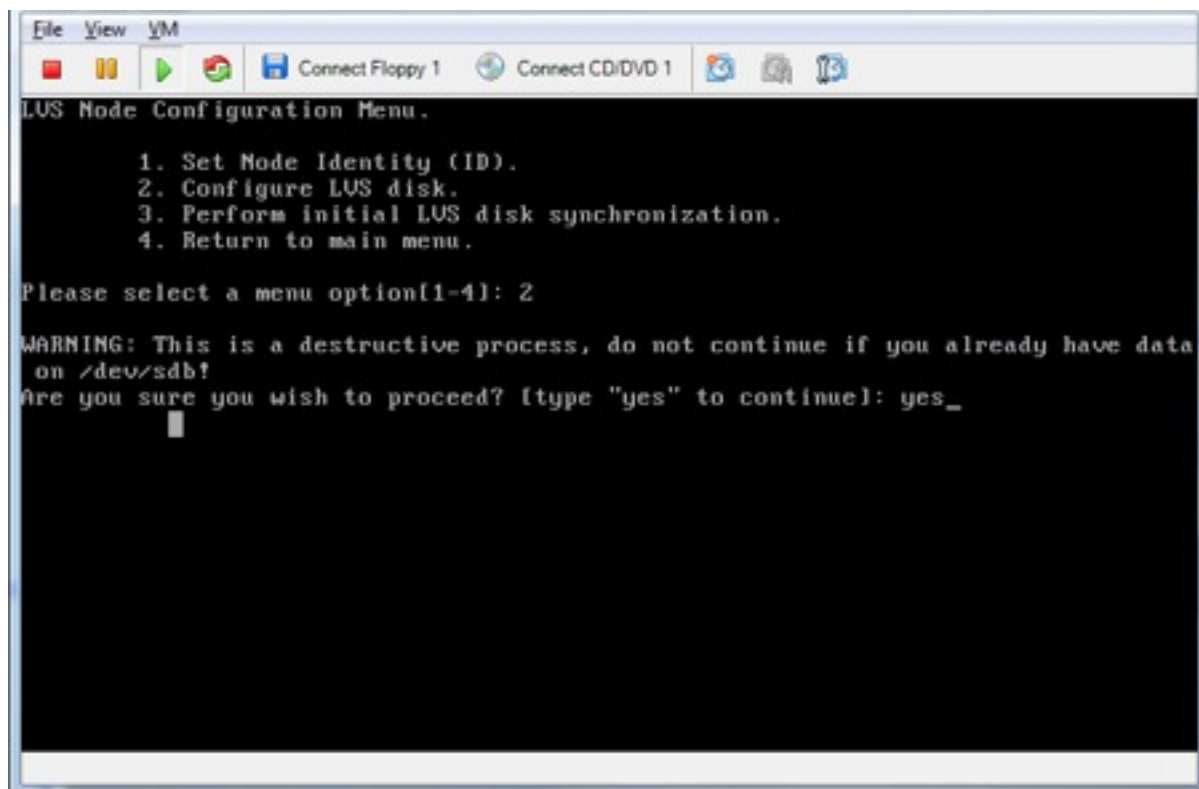
Configuring networking....
Using custom IP addresses. Please enter remote node IPs.

Enter Heartbeat IP for remote node: 10.1.1.2

Enter iSCSI Target IP for remote node: 10.1.1.12
done.
Configuring heartbeat service....done.
Configuring time synchronisation....done.

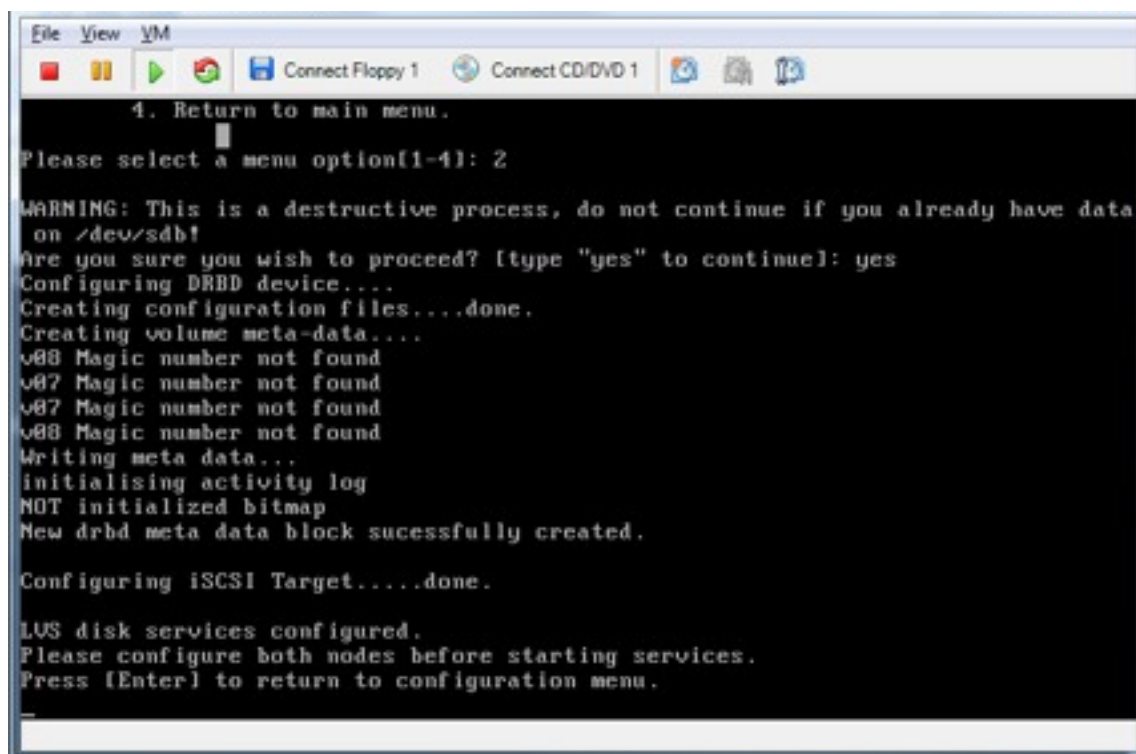
System files configured for node 1.
Please configure both nodes before starting services.
Press [Enter] to return to the configuration menu.
```

When the IP address configuration is complete, select **Option 2** from the configuration menu.



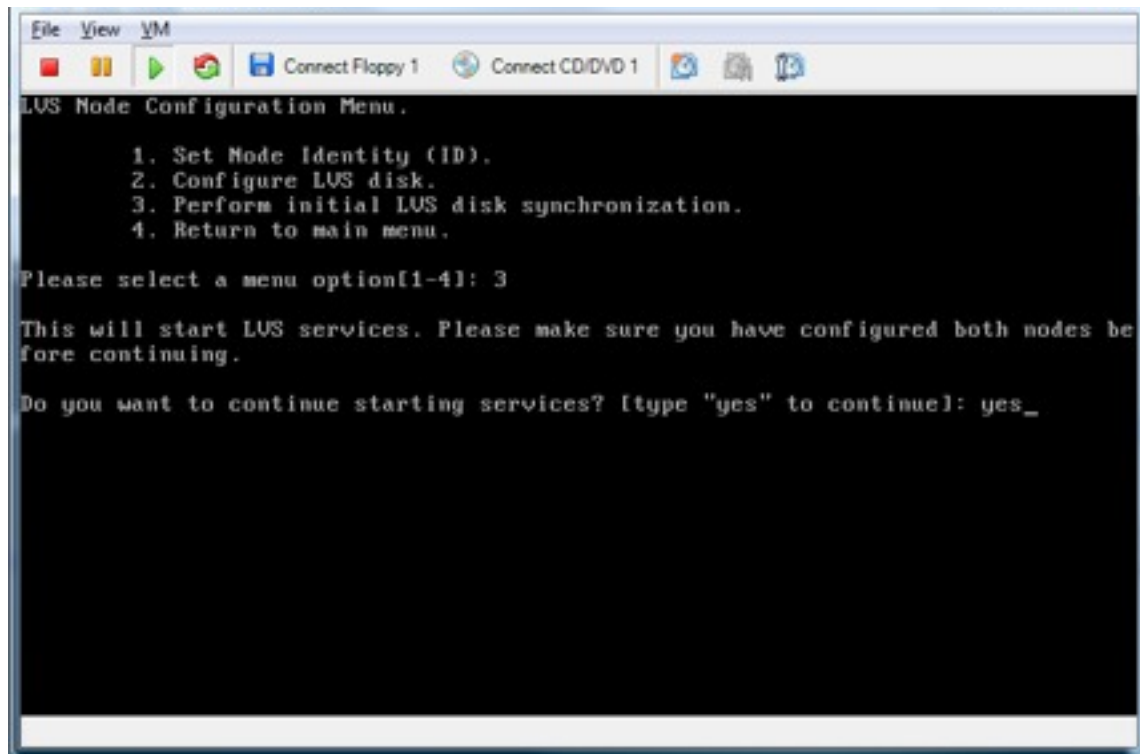
Type **"yes"** to proceed with the XVS volume configuration.

NOTE: Do not perform this step if you already have live data on the appliance, it will delete all existing data from the appliance.

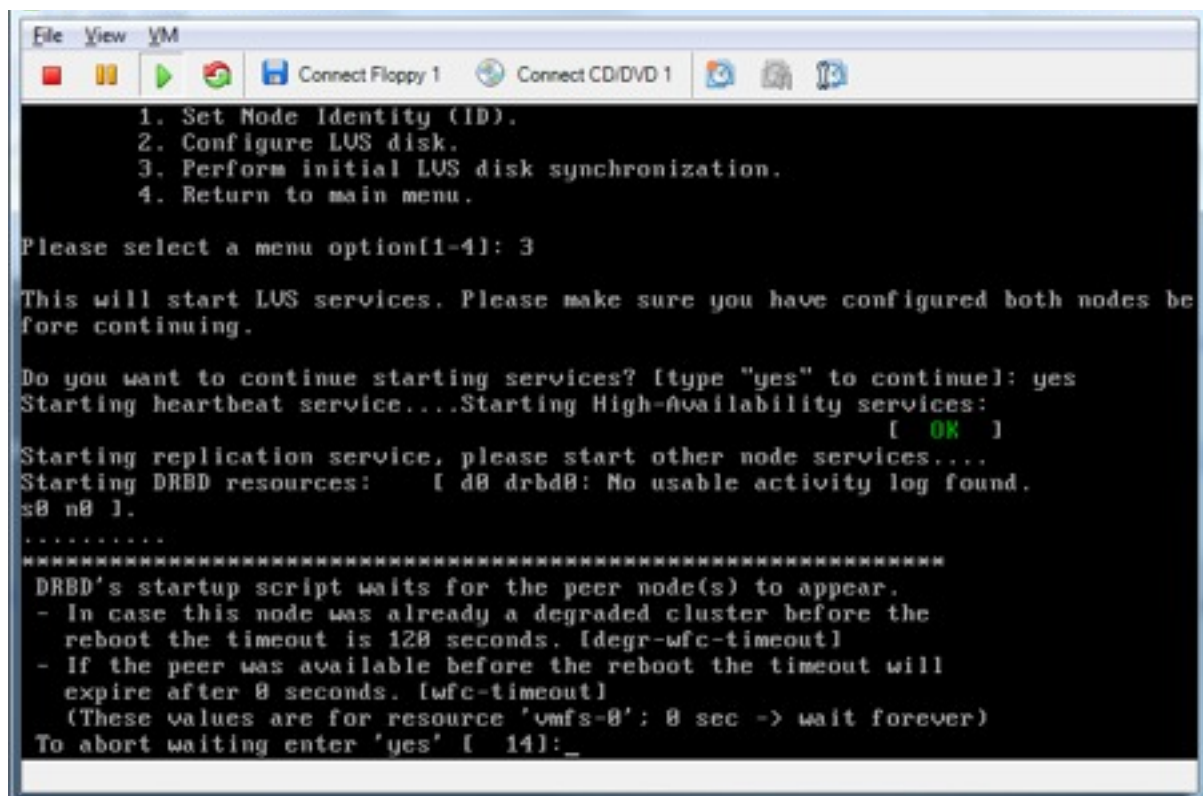


NOTE: Make sure you have completed the above steps on **BOTH** nodes before continuing.

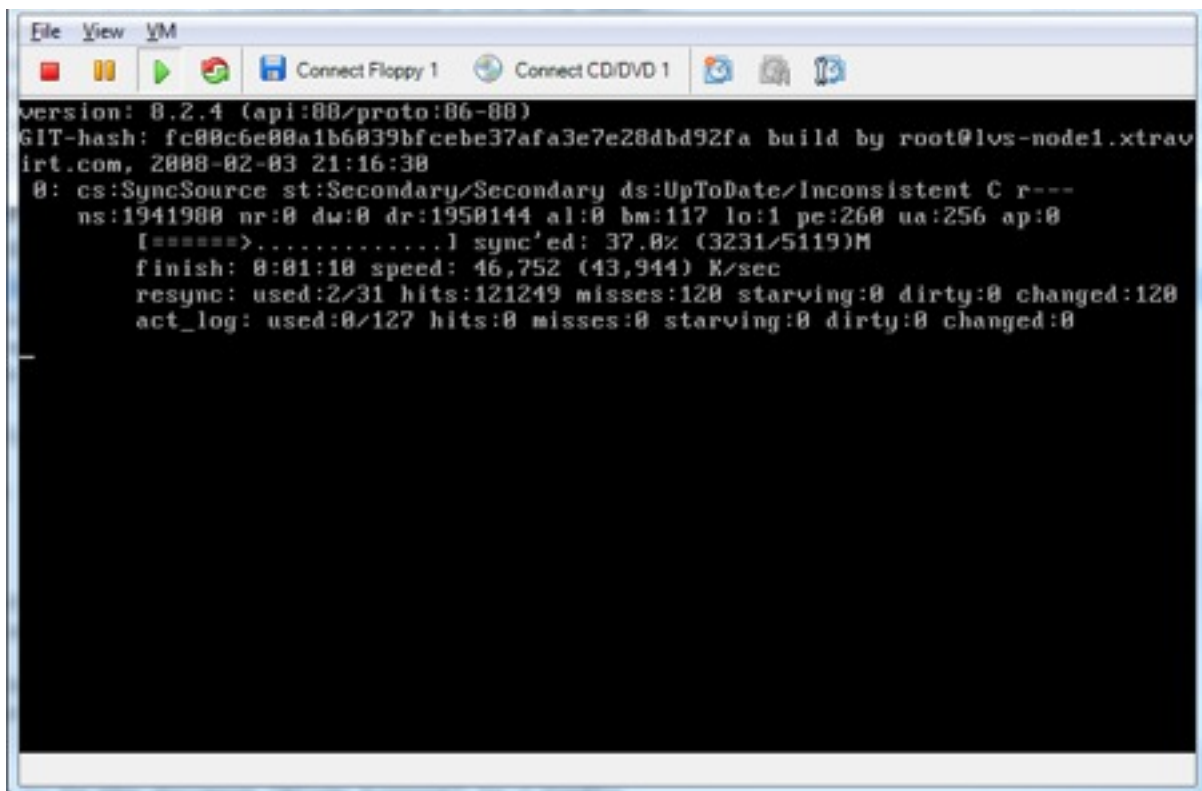
Select **Option 3** to perform the initial XVS disk synchronization. Do this on **node1** first.



When you have started the synchronisation on **node1** and it is waiting for a peer, repeat the process on **node2**.

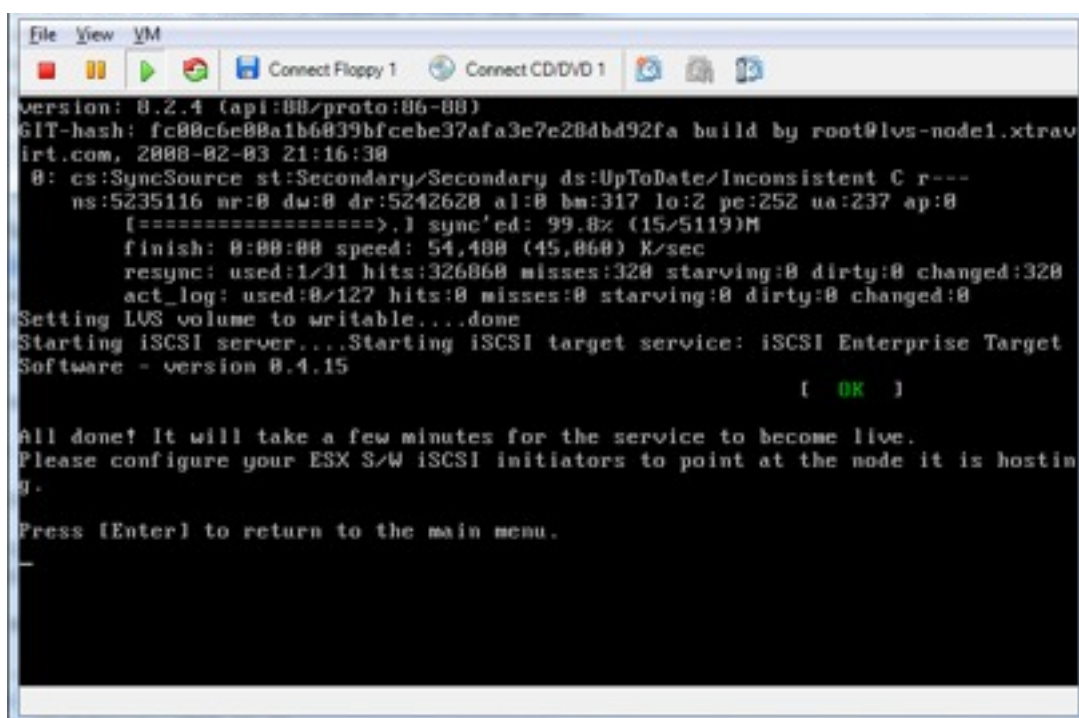


Assuming correct IP configuration the initial synchronisation will start and a progress bar will be shown at the console. **NOTE:** for very large disks the initial synchronisation may take some time.



```
File View VM
Connect Floppy 1 Connect CD/DVD 1
version: 0.2.4 (api:00/proto:06-00)
GIT-hash: fc00c6e00a1b6039bfcebe37afa3e7e20dbd92fa build by root@lvs-node1.xtrav
irt.com, 2008-02-03 21:16:38
0: cs:SyncSource st:Secondary/Secondary ds:UpToDate/Inconsistent C r---
ns:1941988 nr:0 du:0 dr:1950144 al:0 bm:117 lo:1 pe:260 ua:256 ap:0
[=====>.....] sync'ed: 37.0% (3231/5119)M
finish: 0:01:10 speed: 46,752 (43,944) K/sec
resync: used:2/31 hits:121249 misses:120 starving:0 dirty:0 changed:120
act_log: used:0/127 hits:0 misses:0 starving:0 dirty:0 changed:0
```

Once the synchronisation is complete, the iSCSI services are started and configuration is complete. Please allow a couple of minutes for the virtual IP addresses to establish themselves before attempting to connect to the storage. Your Virtual SAN is now ready for use.

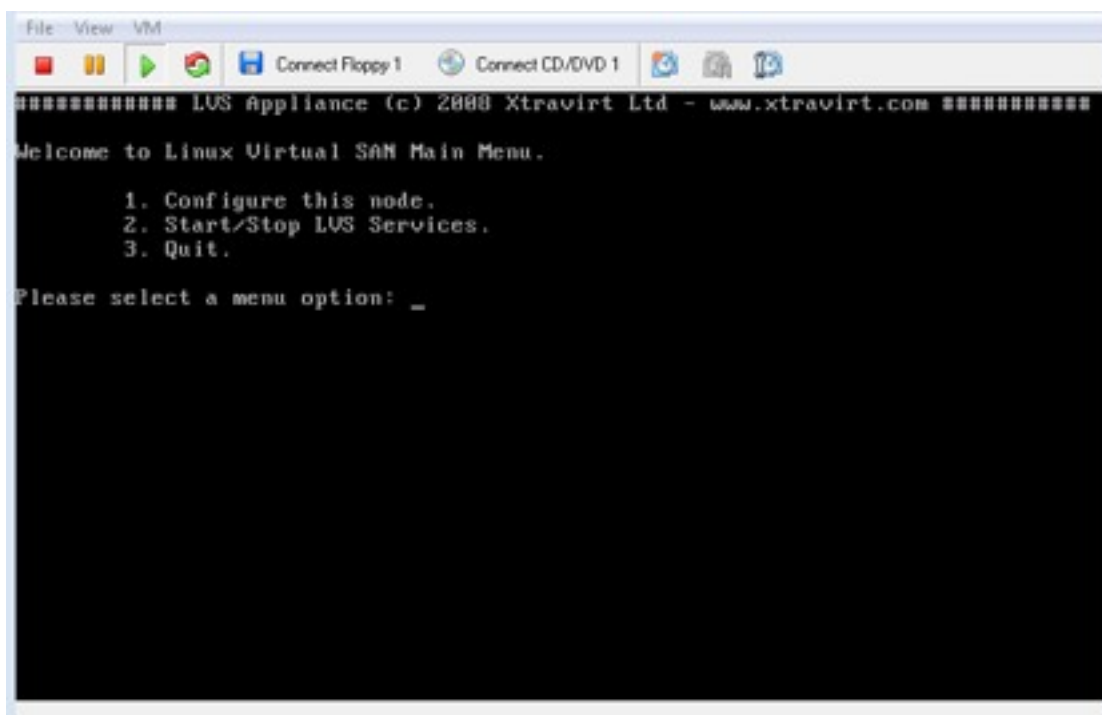


```
File View VM
Connect Floppy 1 Connect CD/DVD 1
version: 0.2.4 (api:00/proto:06-00)
GIT-hash: fc00c6e00a1b6039bfcebe37afa3e7e20dbd92fa build by root@lvs-node1.xtrav
irt.com, 2008-02-03 21:16:38
0: cs:SyncSource st:Secondary/Secondary ds:UpToDate/Inconsistent C r---
ns:5235116 nr:0 du:0 dr:5242620 al:0 bm:317 lo:2 pe:252 ua:237 ap:0
[=====>.....] sync'ed: 99.8% (15/5119)M
finish: 0:00:00 speed: 54,400 (45,060) K/sec
resync: used:1/31 hits:326060 misses:320 starving:0 dirty:0 changed:320
act_log: used:0/127 hits:0 misses:0 starving:0 dirty:0 changed:0
Setting LVS volume to writable...done
Starting iSCSI server...Starting iSCSI target service: iSCSI Enterprise Target
Software - version 0.4.15
[ OK ]

All done! It will take a few minutes for the service to become live.
Please configure your ESX S/W iSCSI initiators to point at the node it is hostin
g.

Press [Enter] to return to the main menu.
```

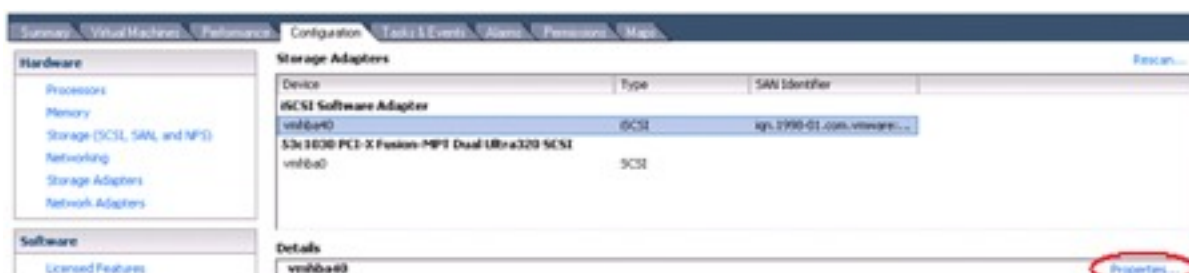
Configuration is now complete. It is recommended to leave the appliance at the main menu and simply close the console. If you wish to enter the troubleshooting console select **Option 3** and then log in with username **root** and password **xtravirt**.



NOTE: To return to the XVS Main Menu type **"configure_node"** at the troubleshooting console.

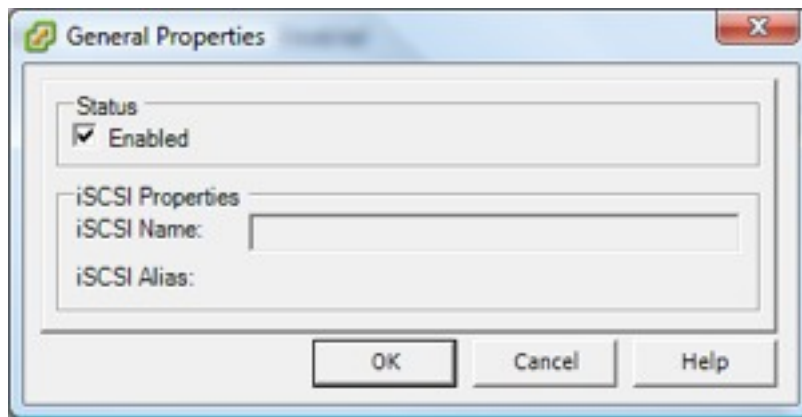
2.4 Connecting ESX hosts to the XVS appliance volume

First we need to enable the ESX software iSCSI initiator. In the VI client, click the ESX host you want to enable iSCSI on, and select the **configuration** tab. Select the **iSCSI Software Adapter** and then click **Properties**.

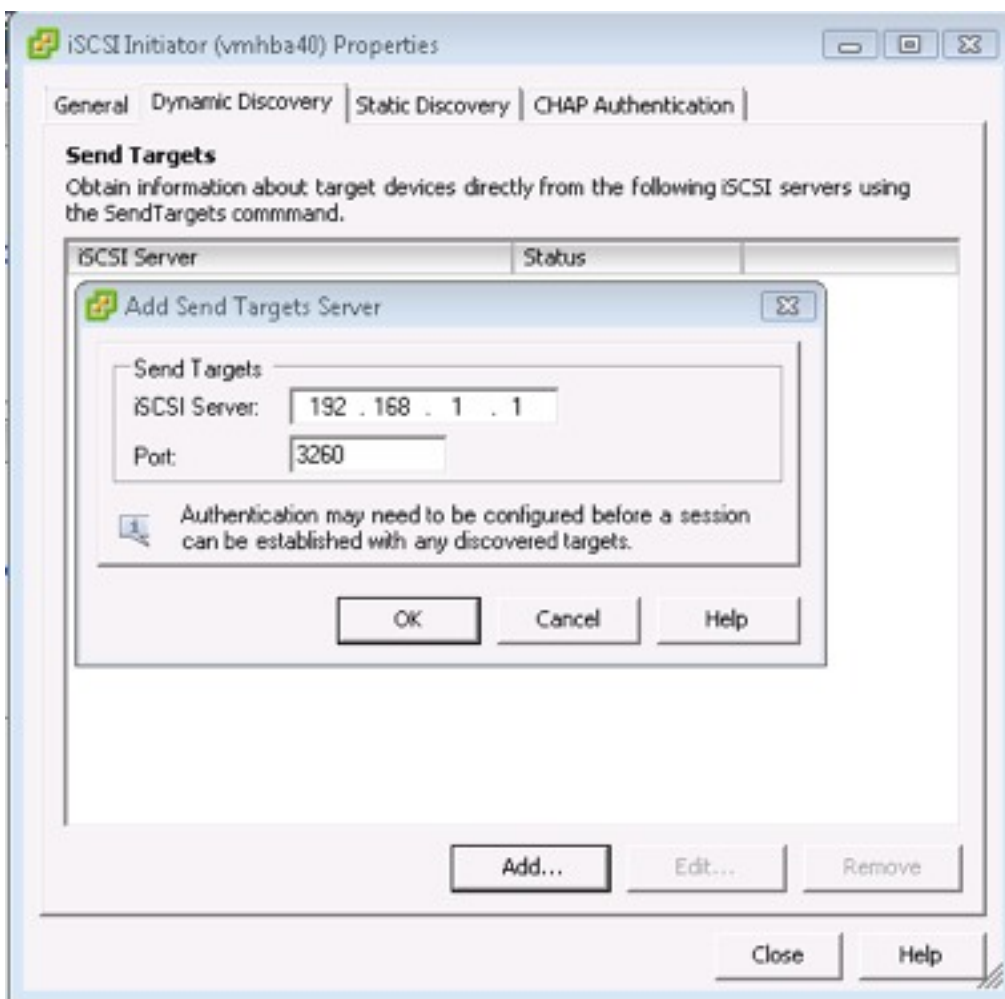


Click **"Configure"** and the following window will pop up. Tick the **"Enabled"** box and press **OK**.

NOTE: In some circumstances it is necessary to reboot the ESX host, however this is usually not necessary and the message can be safely ignored.



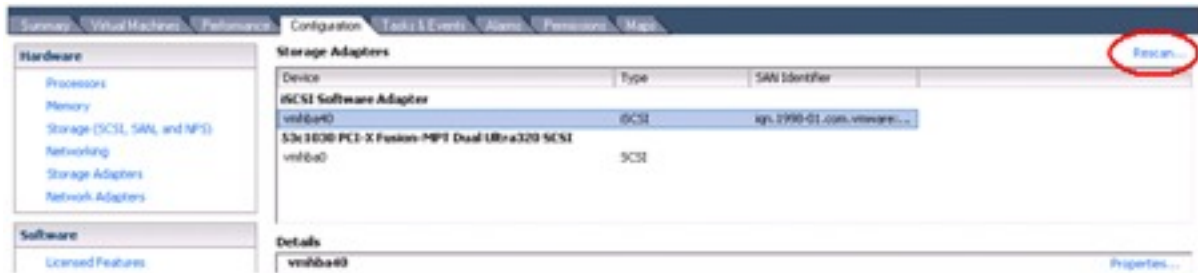
Click "**Properties**" again, and select the "**Dynamic Discovery**" tab. Click "**Add**" and enter the IP of your locally stored XVS appliance iSCSI Target IP – in this case **192.168.1.1**. Click **OK** and then **Close**.



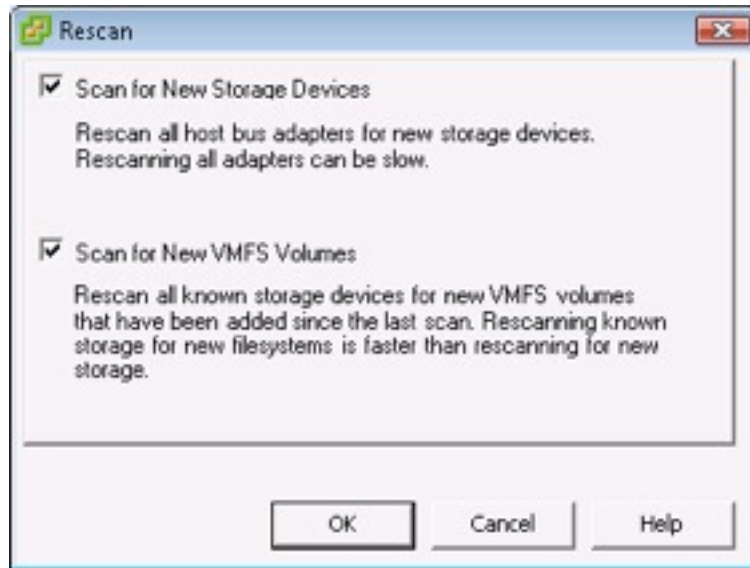
NOTE: to improve read performance and reduce network traffic, it is recommended that you add only the iSCSI Target IP address of the XVS appliance that the ESX host is hosting on its local disks.

For example, in the above screen shot – node1 is hosted on ESX1 so we have configured ESX1’s software iSCSI initiator to point at node1’s default iSCSI Target IP address 192.168.1.1. Assuming we are using the defaults, ESX2’s software iSCSI initiator would be configured to point at node2 – 192.168.1.2.

When the task is complete, click **Rescan**.




Leave both boxes checked and click **OK**.

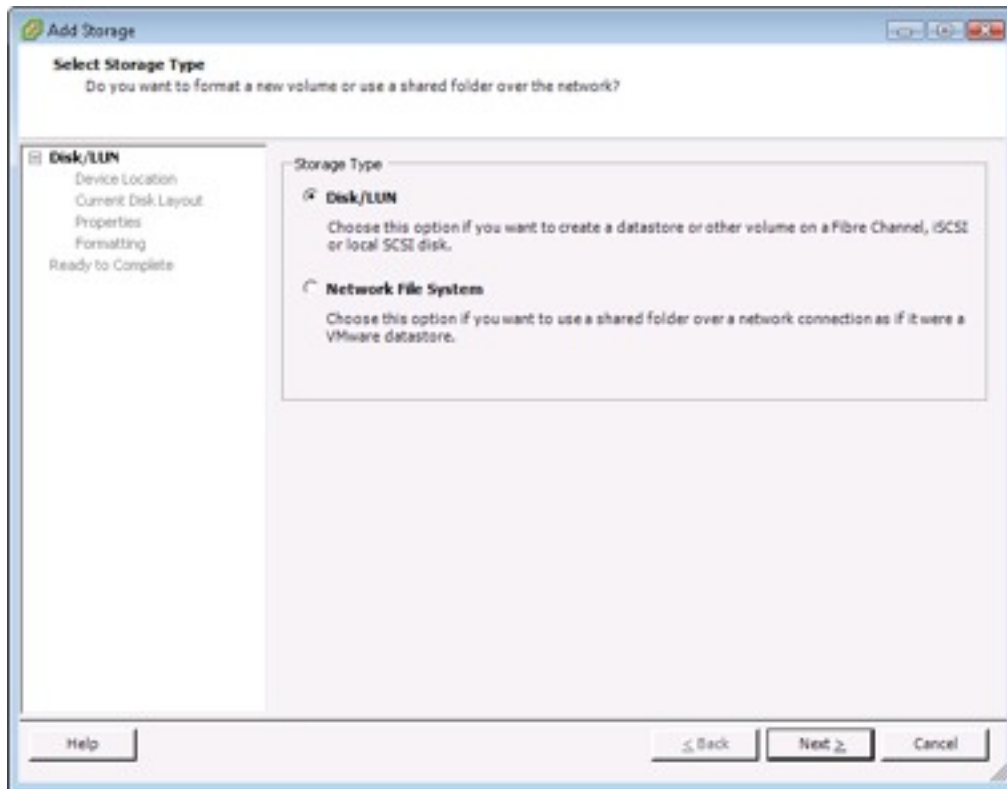


When the rescan is complete you should see your XVS iSCSI target listed in the bottom pane when you click the Software iSCSI Adapter.

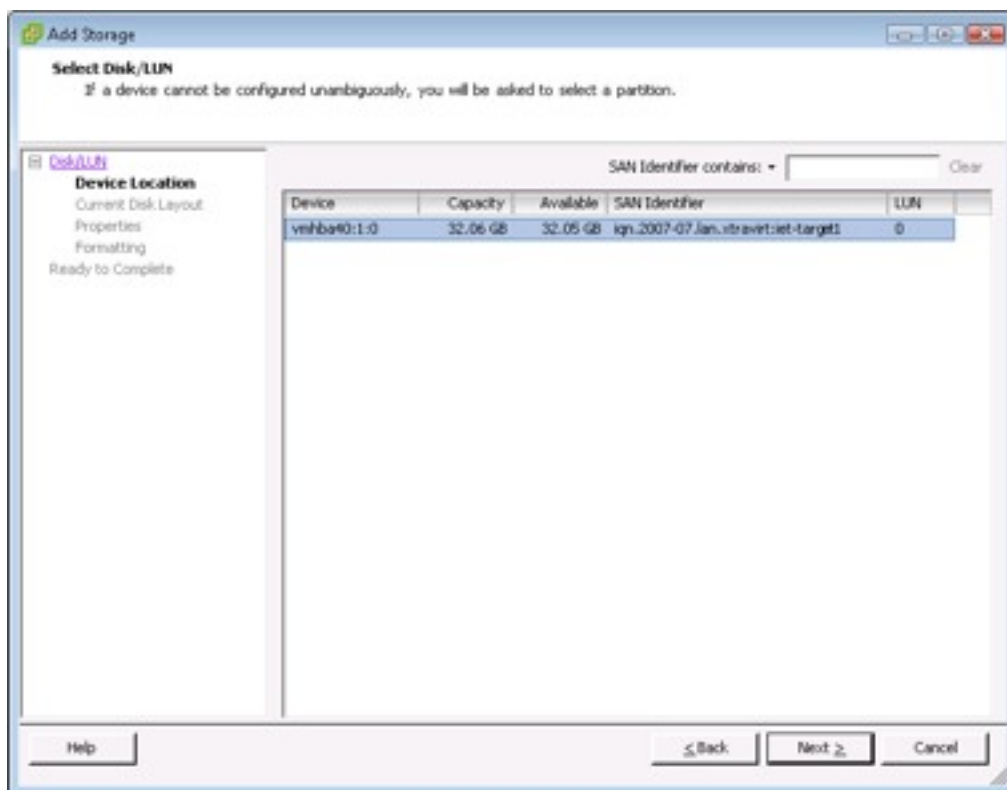
The storage needs to be formatted as VMFS. Select the **Storage (SCSI, SAN and NFS)** link under the host configuration tab and then click **Add Storage**.



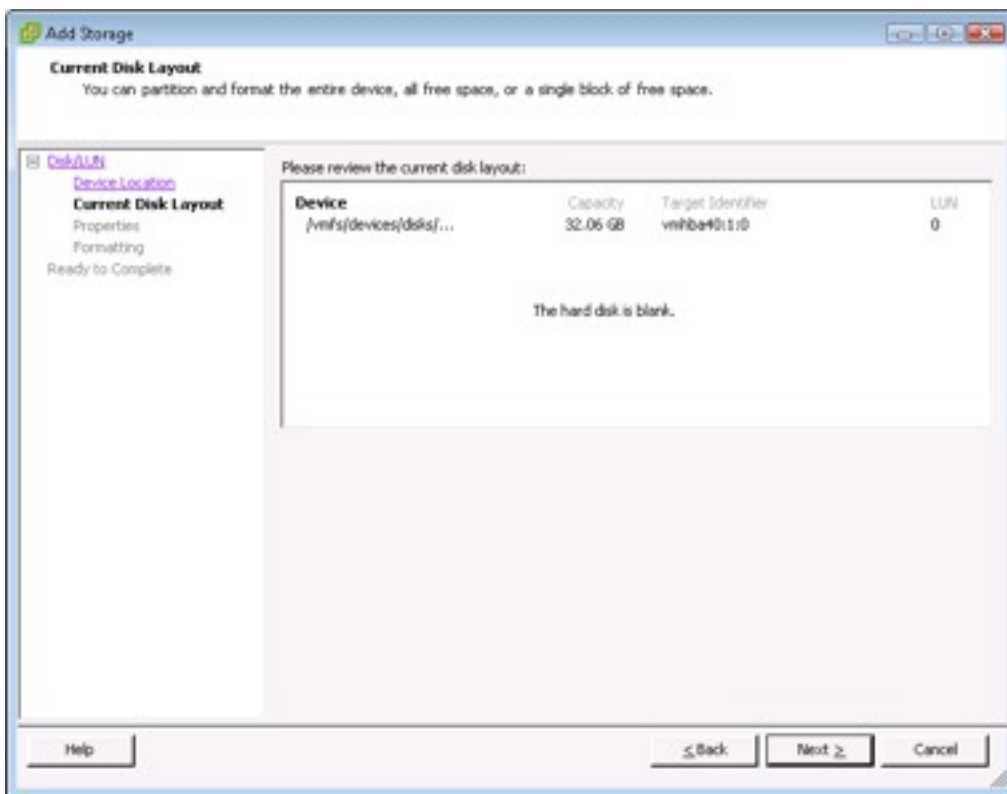
Select **"Disk/LUN"** and click **Next**.



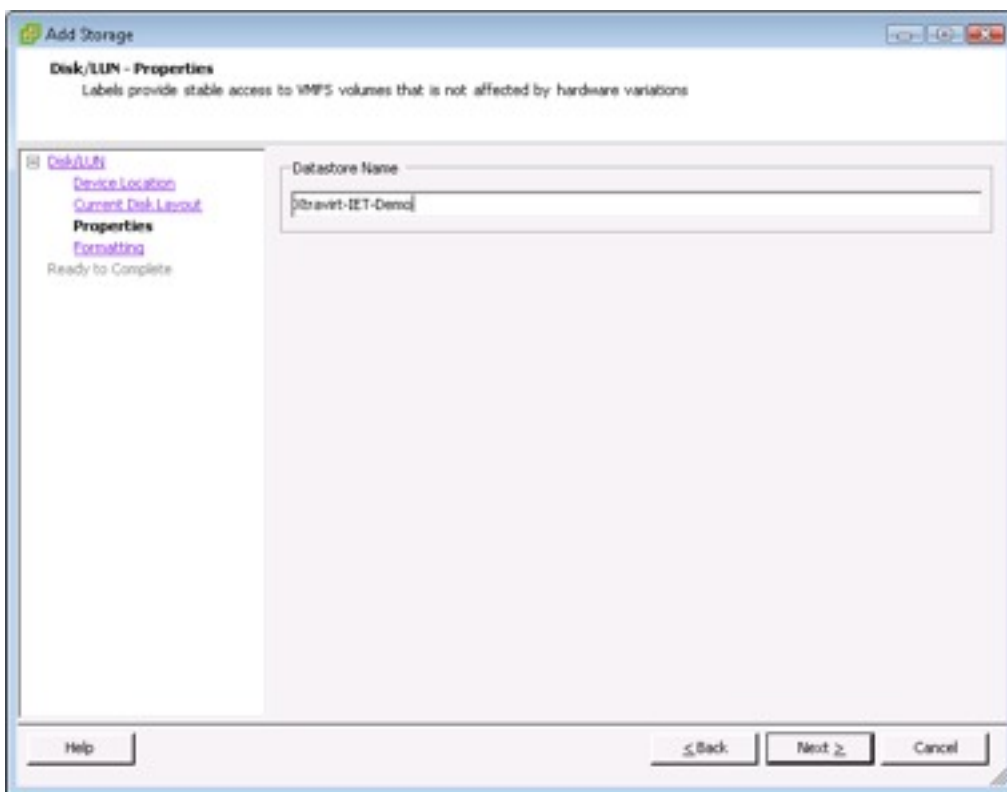
Select the LUN and click **Next**.



The volume should be blank. Click **Next**.

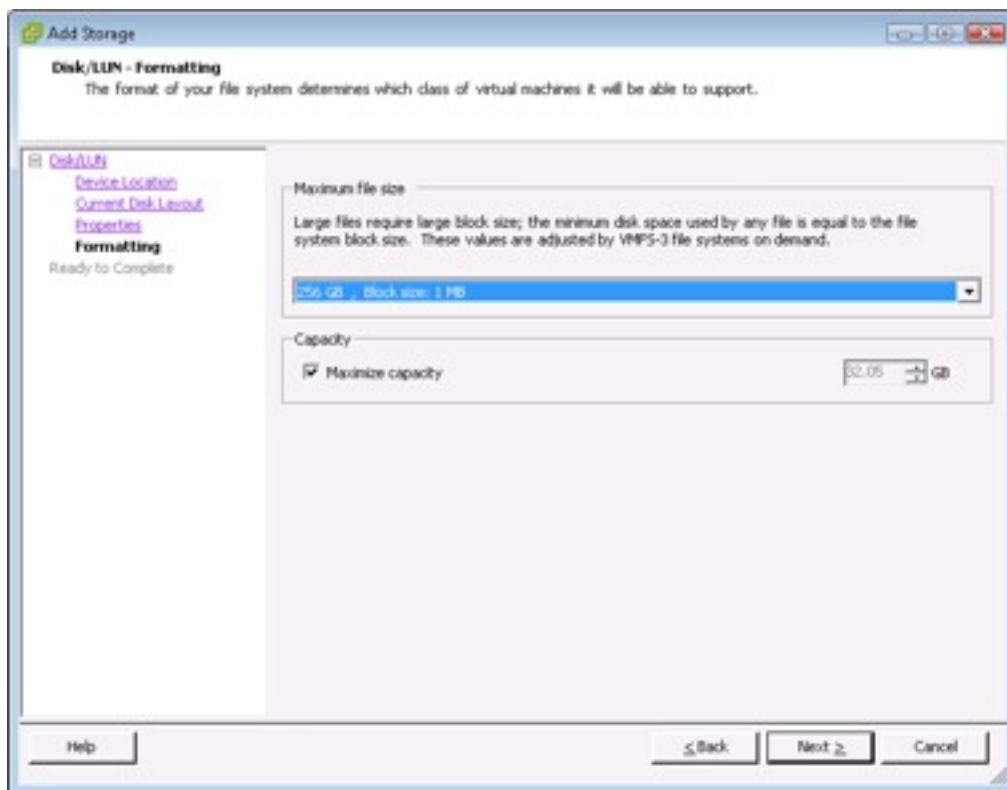


Give the new VMFS datastore a name and click **Next**.

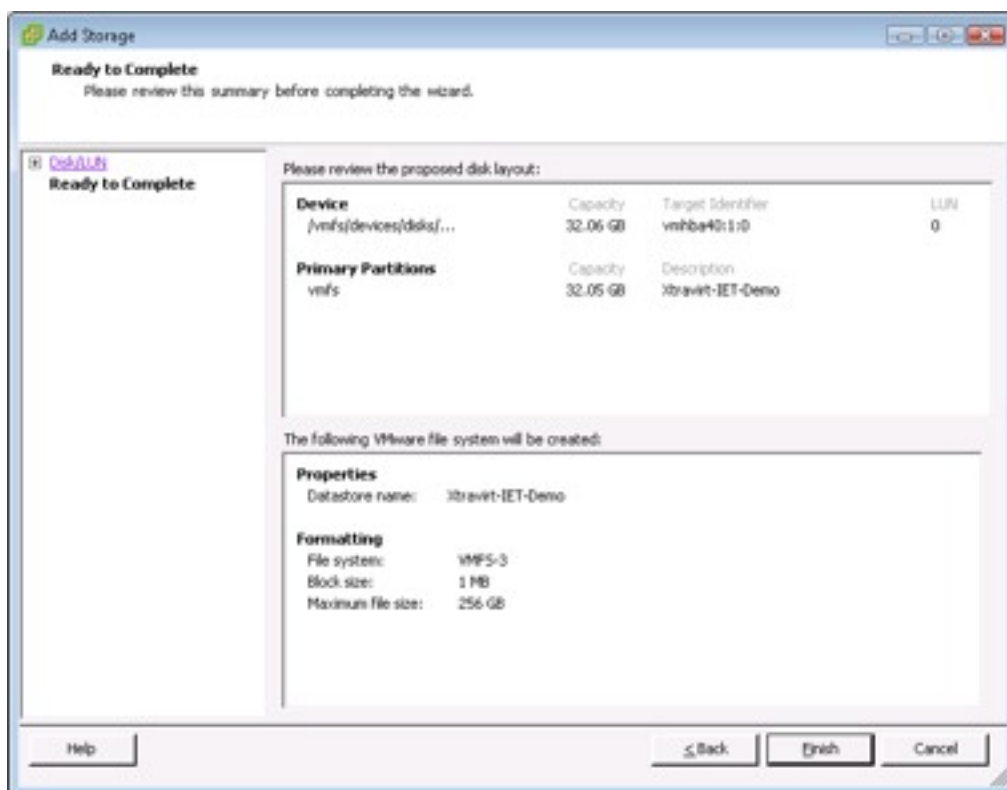


Accept the default block size, click **Next**.

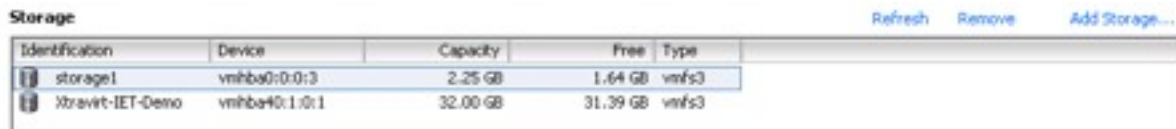
NOTE: if you are using a very large volume, you may wish to increase the block size. Please refer to VMware's best practice guide for VMFS volumes.



Your summary should look like the below screenshot. Click **Finish** to format the Volume as VMFS.

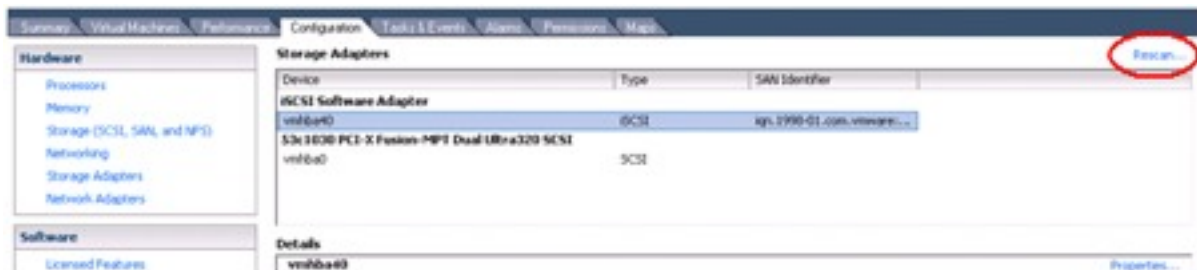


The XVS appliance volume should now be listed as an available VMFS volume.

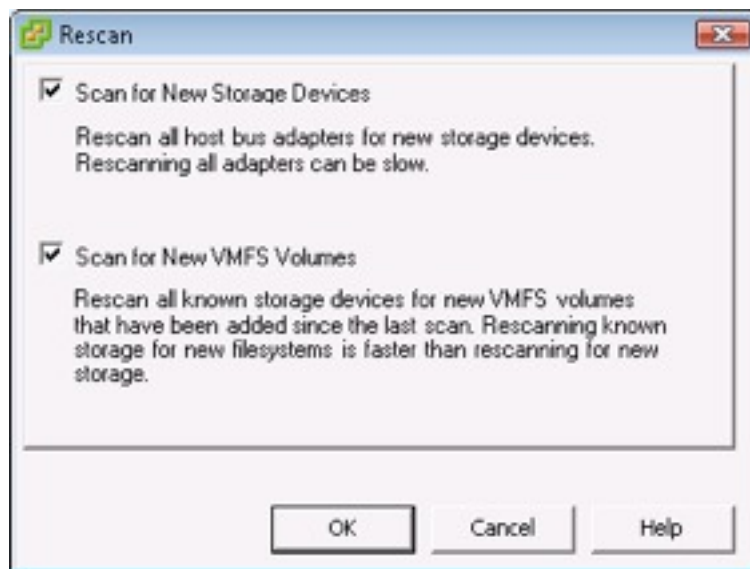


Identification	Device	Capacity	Free	Type
storage1	vmhba0:0:0:3	2.25 GB	1.64 GB	vmfs3
Xtravirt-IET-Demo	vmhba0:1:0:1	32.00 GB	31.39 GB	vmfs3

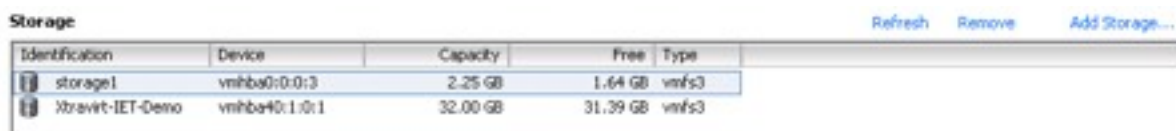
Click your second ESX host, select the **Configuration** tab and click **Storage Adapters**. Click **Rescan**.



Leave both boxes checked and click **OK**.



The XVS appliance volume should now be listed as an available VMFS volume on the second ESX host.



Identification	Device	Capacity	Free	Type
storage1	vmhba0:0:0:3	2.25 GB	1.64 GB	vmfs3
Xtravirt-IET-Demo	vmhba0:1:0:1	32.00 GB	31.39 GB	vmfs3

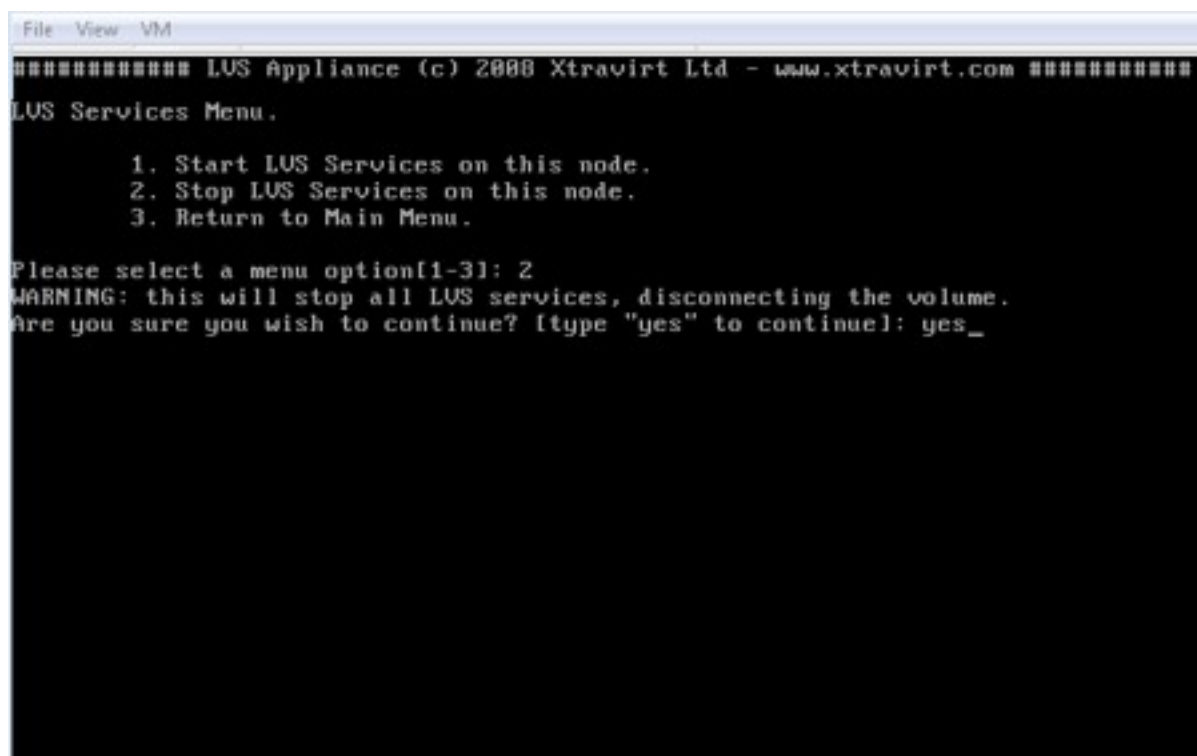
You now have a virtual SAN VMFS volume that resides on both local disks of your ESX hosts, you can vMotion and perform all of the other functions that are available with standard shared storage devices using this volume and all data is synchronously replicated between the disks.

2.5 Planned XVS appliance node shutdown

If one ESX host needs to be taken down for a reboot or maintenance it may be necessary to shutdown one of the XVS appliance nodes. All services can continue to function from the other node, but a clean shutdown greatly reduces replication times on re-synchronisation and is recommended where possible.

On the node that is to be shutdown select **Option 2** from the main menu to enter the **Services Menu**.

At the **Services Menu** select **Option 2** to shutdown the services on this XVS node. Type **"yes"** to confirm you wish to continue with the shutdown.

A screenshot of a terminal window titled 'File View VM'. The terminal displays the 'LUS Services Menu' with three options: 1. Start LUS Services on this node, 2. Stop LUS Services on this node, and 3. Return to Main Menu. The user has selected option 2, and the terminal shows a warning: 'WARNING: this will stop all LUS services, disconnecting the volume. Are you sure you wish to continue? [type "yes" to continue]: yes_'. The terminal background is black with white text.

```
File View VM
##### LUS Appliance (c) 2008 Xtravirt Ltd - www.xtravirt.com #####
LUS Services Menu.

    1. Start LUS Services on this node.
    2. Stop LUS Services on this node.
    3. Return to Main Menu.

Please select a menu option[1-3]: 2
WARNING: this will stop all LUS services, disconnecting the volume.
Are you sure you wish to continue? [type "yes" to continue]: yes_
```

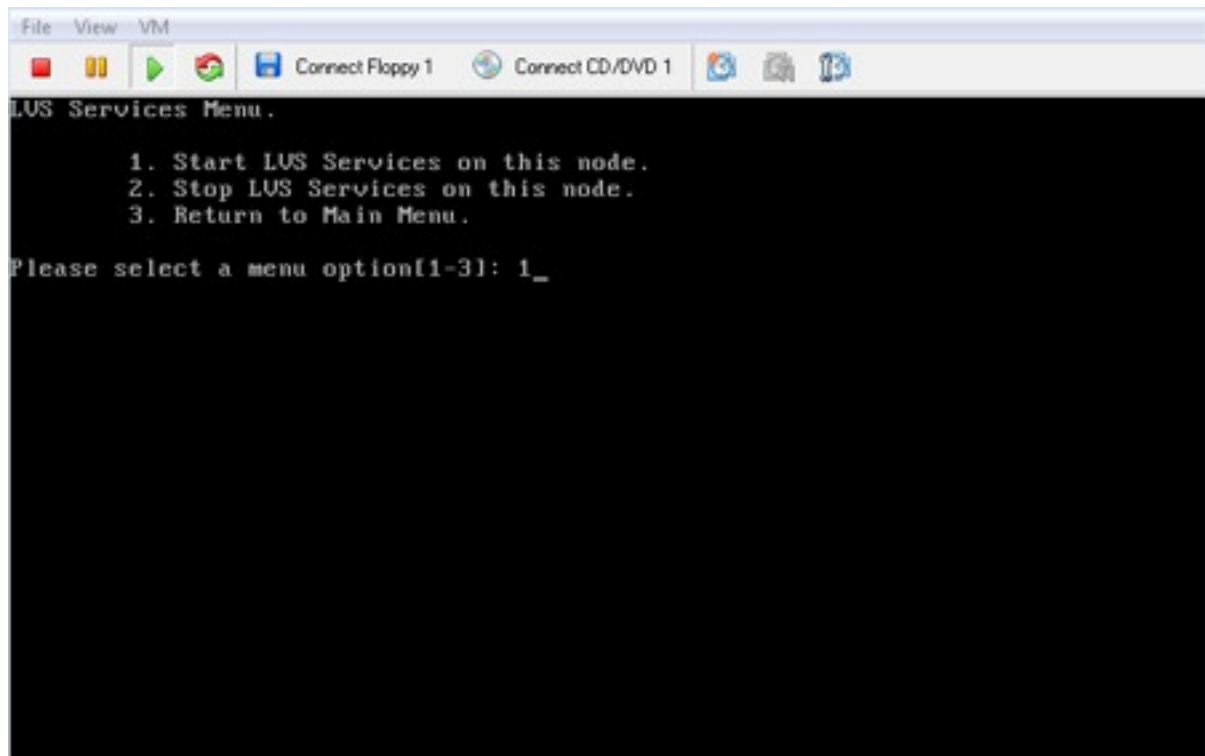
The services will then shutdown and the node can safely be powered off.

2.6 Recovering from an XVS appliance node shutdown or failure

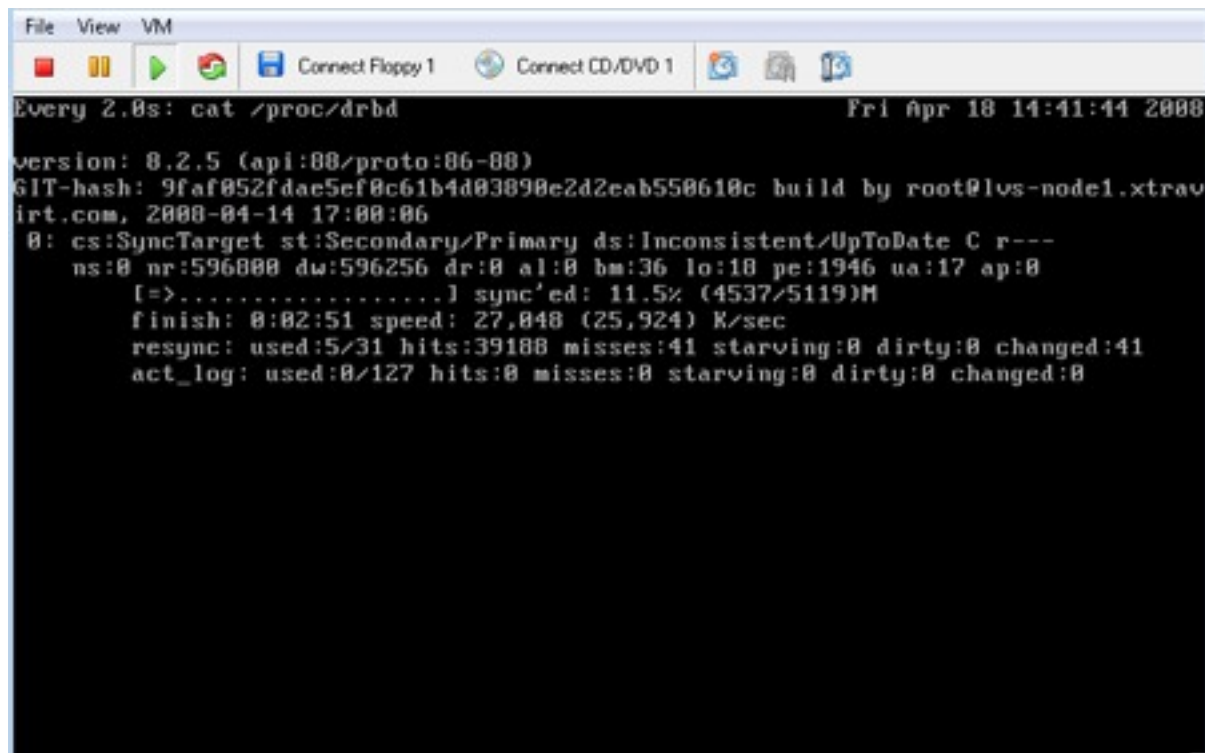
If one of the nodes is lost for any reason, including ESX host failure, the other node will take over all storage operations and the volume will continue to function. When the node is restored to operation it is necessary to re-sync the volume to resume full failover capability of the SAN.

Power on the XVS node that was disconnected and allow it to boot. Select **Option 2** at the **Main Menu** to enter the **Services Menu**.

At the **Services Menu** select **Option 1** to start the XVS services.

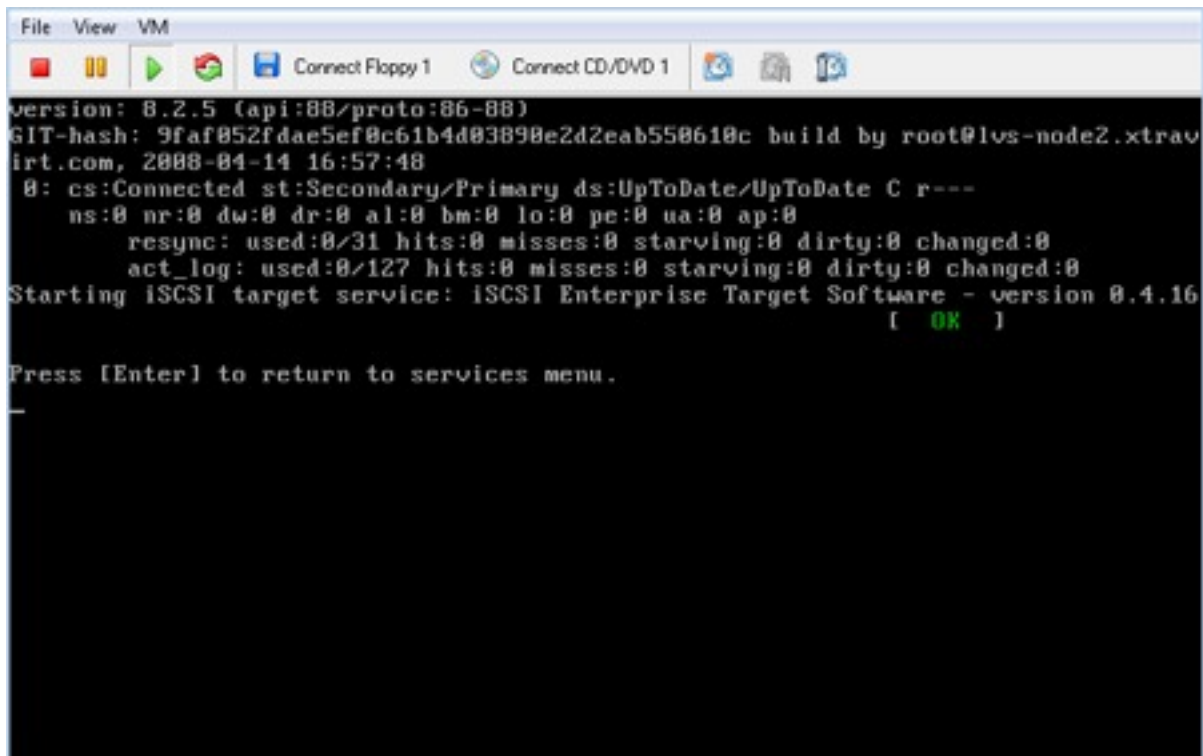


Assuming network connectivity between the nodes is functional, you should be presented with a resynchronisation progress screen. This will automatically exit when resynchronisation is complete.



NOTE: The length of time required to resynchronise the volumes varies dependant on how much data has changes since the node was disconnected and whether the node was cleanly shut down or not.

When the resynchronisation is complete, press **Enter** to return to the main menu. The XVS SAN is now restored to full redundancy and the iSCSI Target IP address will automatically move back to the recovered node.



```
File View VM
Connect Floppy 1 Connect CD/DVD 1
Version: 0.2.5 (api:88/proto:86-88)
GIT-hash: 9faf052fdae5ef0c61b4d03890e2d2eab550610c build by root@lvs-node2.xtrav
irt.com, 2008-04-14 16:57:48
0: cs:Connected st:Secondary/Primary ds:UpToDate/UpToDate C r---
ns:0 nr:0 dw:0 dr:0 al:0 bm:0 lo:0 pe:0 ua:0 ap:0
resync: used:0/31 hits:0 misses:0 starving:0 dirty:0 changed:0
act_log: used:0/127 hits:0 misses:0 starving:0 dirty:0 changed:0
Starting iSCSI target service: iSCSI Enterprise Target Software - version 0.4.16
[ OK ]

Press [Enter] to return to services menu.
```

We hope you found this guide and appliance useful.