



# NFS & iSCSI Multipathing in vSphere

VMware Storage

## Table of contents

NFS & iSCSI Multipathing in vSphere .....	3
Path selection .....	3
Multipathing configuration .....	4
Paths in same Subnet .....	4
Paths in different Subnets .....	5
Avoid Single Point of Failures .....	7
Physical Switches .....	7
Viewing the Multipath connections .....	7
vSphere 6.7 Storage .....	24

## NFS & iSCSI Multipathing in vSphere

### Path selection

Currently, NFS 4.1 client selects the paths in a Round-Robin fashion. It selects the path only from a list of active paths. If a path goes down then it is removed from the list of active paths till the connectivity is restored.

## Multipathing configuration

Before configuring Multipathing, check whether the NFS server has support for it. The IP addresses of the interfaces exposed by the server for Multipathing can be in the same subnet or in a different subnet.


The IP addresses for multiple paths need to be specified during volume mount.

If you want to add another path, the volume needs to be unmounted and remounted. Although there can be more than two paths, the following sections are based on having two paths.

### Paths in same Subnet

When the two paths have IP addresses in the same subnet, the ESXi host will make an independent connection to each of the servers IP through its vmkernel portgroup IP which is in same subnet.

Note: Even if you configure multiple IP addresses in same subnet at ESXi, ESXi will chose only one IP adress as source to establish connections with server IPs.



### Steps to configure

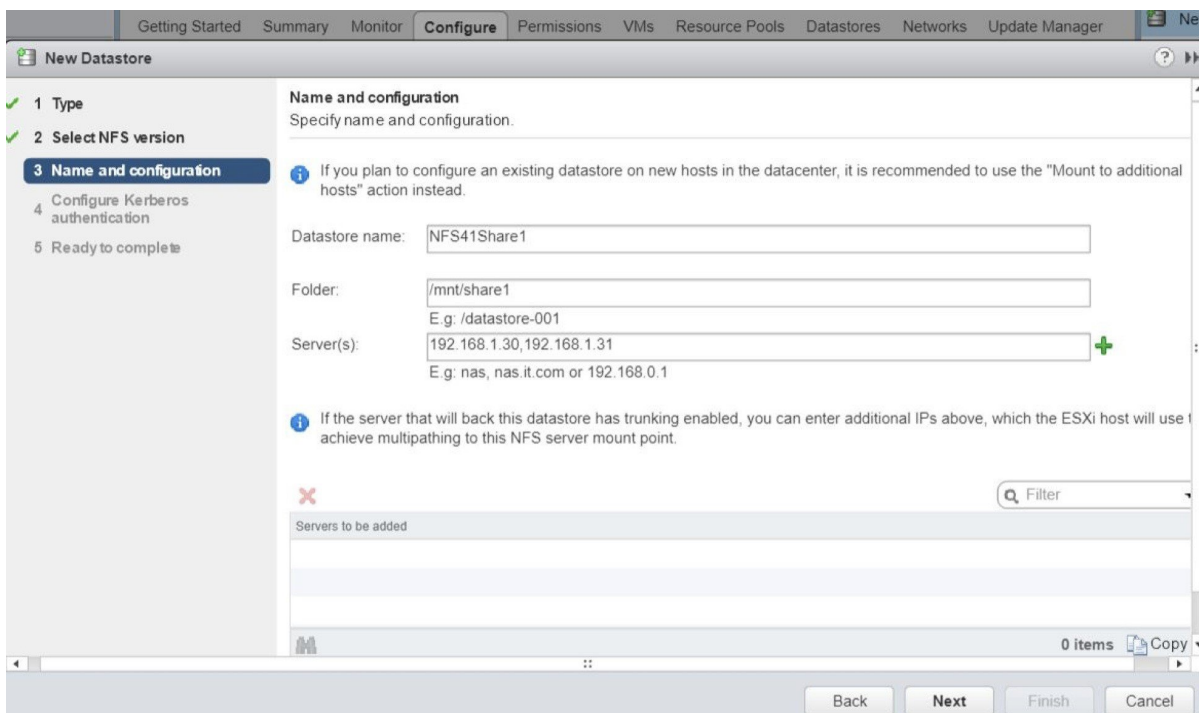
Multipathing Configuration can be done through command line or through vSphere Web Client.

Command line

```
#esxcli storage nfs41 -H 192.168.1.30, 192.168.1.31 -s /mnt/sh are1 -v nfs41share1
```

vSphere Web Client


The configuration screen using vSphere Web Client is as follows.



### Best practice configuration in same subnet

It is recommended to configure NIC Teaming to better utilize the bandwidth of physical NICs and avoid single point of failure of NICS. To configure NIC Teaming, attach multiple adapters to the vmkernel portgroup of NFS.

Configure NIC Teaming with the IP hash load-balancing policy.




### Paths in different Subnets

With different subnets we can configure multiple vmkernel portgroups to communicate with IP address in the respective subnets. This configuration provides independent connections from each vmkernel portgroup and provides better redundancy and bandwidth utilization than paths in the same subnet.

In the example below

- 192.168.1.y can be reached by vmk1, which is in same subnet as vmk1.
- 192.168.2.y can be reached by vmk2, which is in same subnet as vmk2.



### Steps to configure

Command line

```
[root@w2-nfs-esx106:~] esxcli storage nfs41 -H 192.168.1.101,1
```

```
92.168.3.101 -s /share2 -v nfs41Share2
```

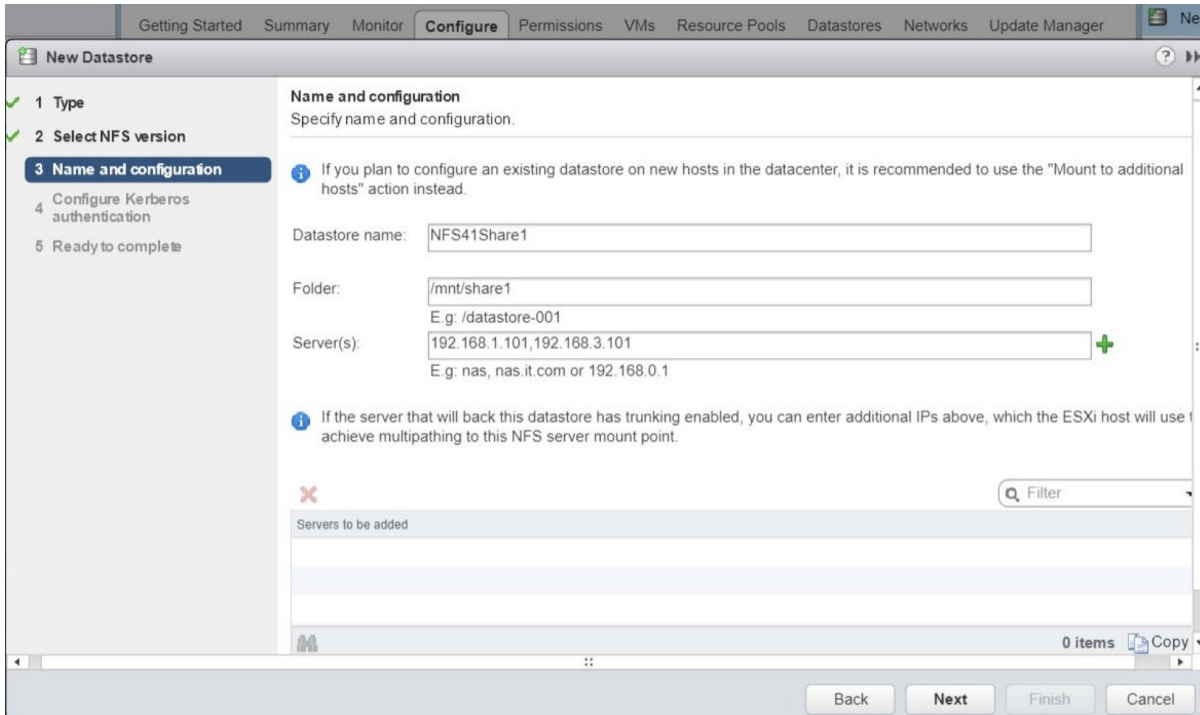
```
[root@w2-nfs-esx106:~] esxcli storage nfs41 list
```

Volume Name Host(s) Share Accessible Mounted Read-Only Security sPE Hardware Acc eleration

```
-----
krb5i_200g 192.168.1.101,192.168.3.101 /krb5i_200g true true false AUTH_SYS false Not Supported
```

vSphere Web Client

The configuration screen using vSphere Web Client is as follows.



## Best practice configuration for different subnets


With different subnets, we can have independent connections from different vmkernel portgroups. To utilize the network bandwidth and to have better redundancy at the physical layer, configure each portgroup to be in different virtual switches and configure each vmkernel portgroup for NIC Teaming.

## Avoid Single Point of Failures

NIC teaming provide first level of redundancy at NIC level. As a generic NFS best practice and to avoid single point of failure further at physical switches and NAS level, configure them for redundancy at different levels.

### Physical Switches

This configuration provides second level of redundancy at physical switches. If switch1 goes down, traffic can be routed through switch2. With this solution, one has four potential NIC cards in the ESX server configured with IP hash failover and two pairs going to separate LAN switches - with each pair configured as teamed at the respective LAN switches.



For NAS level redundancy and more details, refer to the [NFS Best Practices guide](#).

### Viewing the Multipath connections

“vsish” utility can be used to view the connection details for each mounted NFS41 share. The vSphere Client currently doesn’t provide this information.

Following is the output for a share with two paths in different subnets.

```
[root@w2-nfs-esx106:~] esxcli storage nfs41 list
```

```
Volume Name Host(s) Share Accessible Mounted Read-Only Security i sPE Hardware Acceleration
```

```
-----
share2 192.168.1.101,192.168.3.101 /share2 true true false AUTH_SYS false Not Supported share2 Datastore is associated with cluster "1". There are two connections associated with the session of this cluster.
```

```
/>get /vmkModules/nfs41client/clusters/1/servers/1/sessions/00
```

```
0000009A601F597856341200000000/connections/1/info NFSv4.1 Connection Info {
network Address:tcp:192.168.1.101.8.1 state:NFS41_S_CONN_UP
sched:1 orphaned:0
}
```

```
/>get /vmkModules/nfs41client/clusters/1/servers/1/sessions/00
```

```
0000009A601F597856341200000000/connections/2/info NFSv4.1 Connection Info {
network Address:tcp:192.168.3.101.8.1 state:NFS41_S_CONN_UP sched:1 orphaned:0 }
```


4 Configuring iSCSI Storage



# Getting Started with vSphere with Operations Management

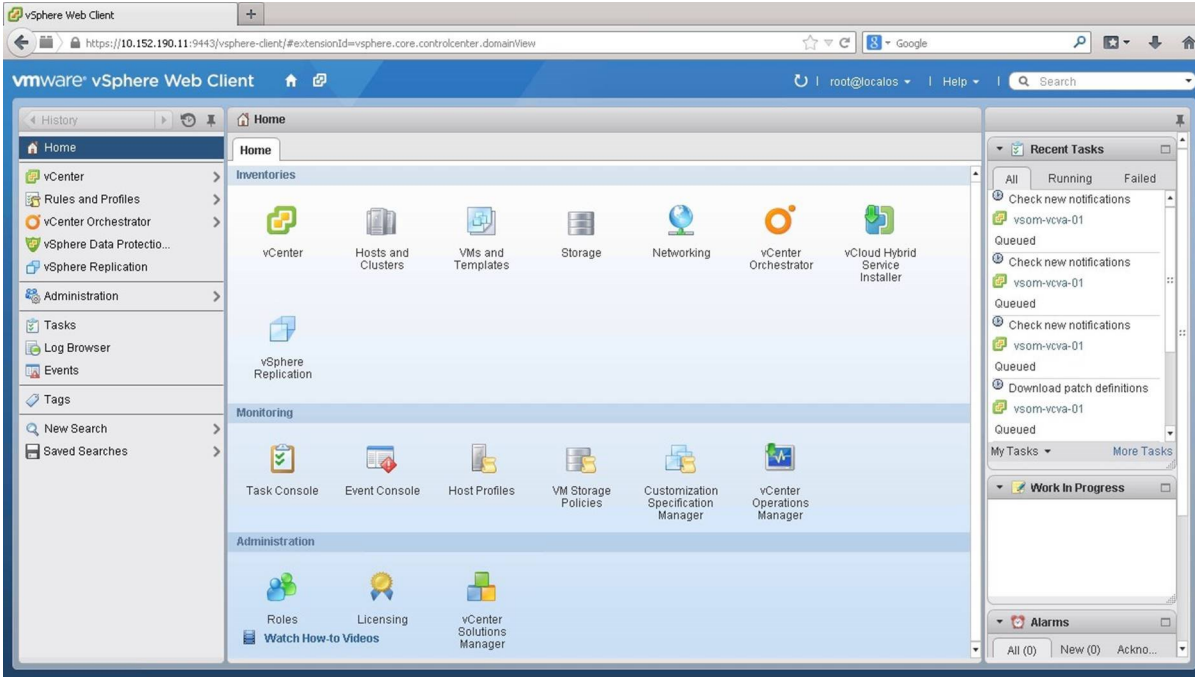
## Configuring iSCSI Storage

This walkthrough demonstrates how to connect to iSCSI storage on an ESXi host managed by vCenter with network connectivity provided by vSphere Standard Switches. Use the arrow keys to navigate through the screens.

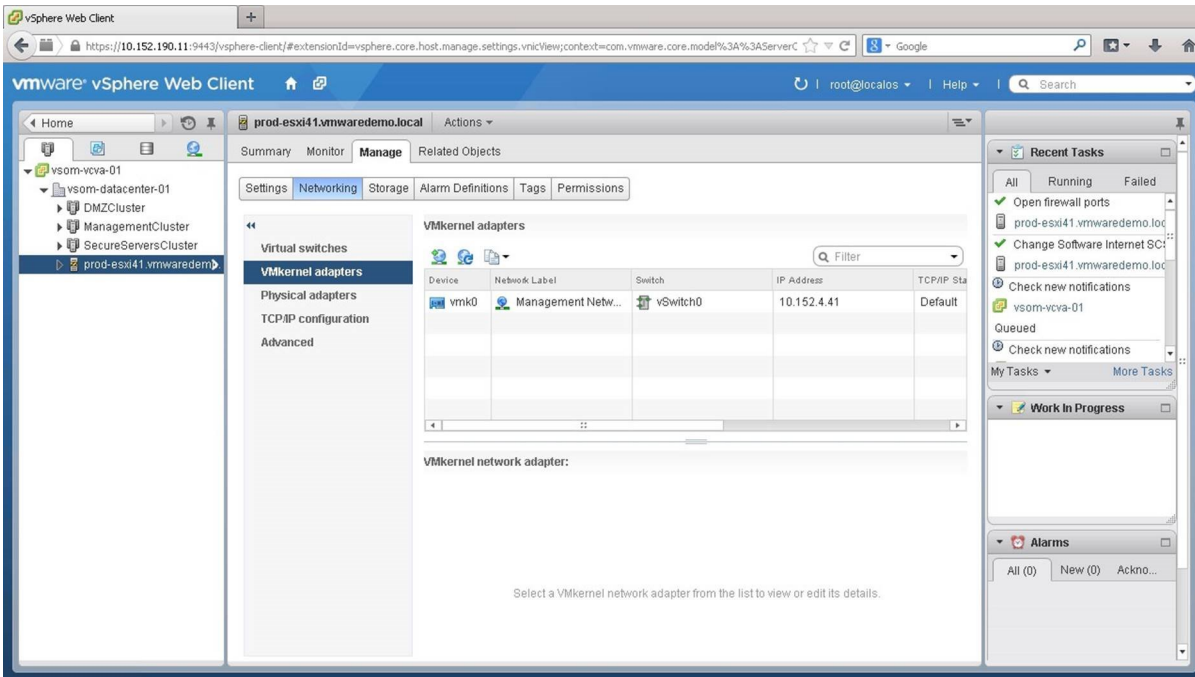


Now that you understand how iSCSI is presented and connected, let's look at how to configure iSCSI in ESXi. We logon on to the vSphere Web Client.






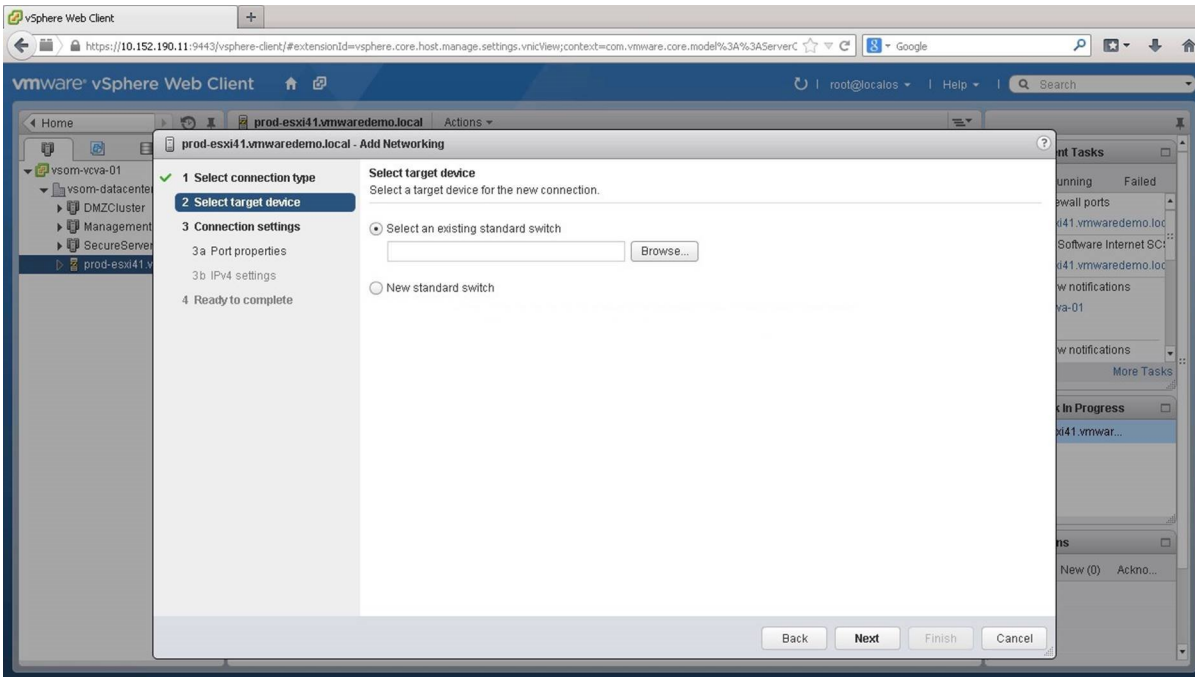
Go to the [Hosts and Clusters] view.



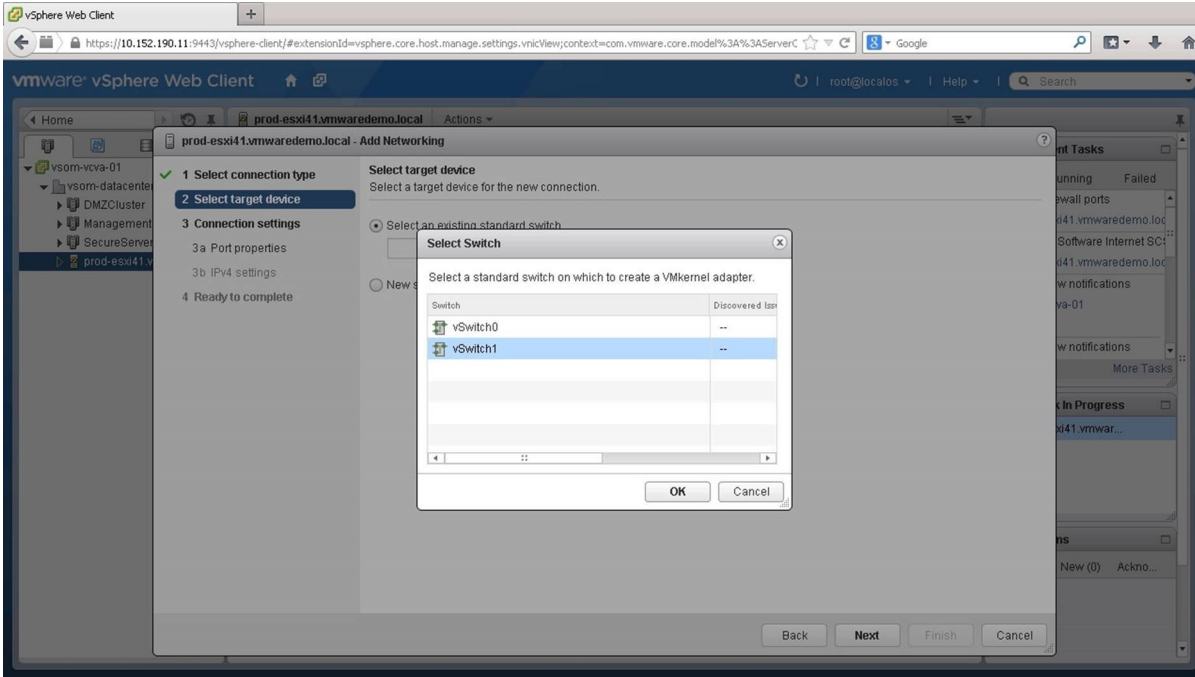
Select the [Host], go to the [Manage] tab, and click on [Networking]. Under [VMkernel Adapters], click on [Add Host Networking].




Select [VMkernel Network Adapters] and click on [Next].




We click on [Browse] to select an existing standard switch. You can also choose to create a new standard switch.



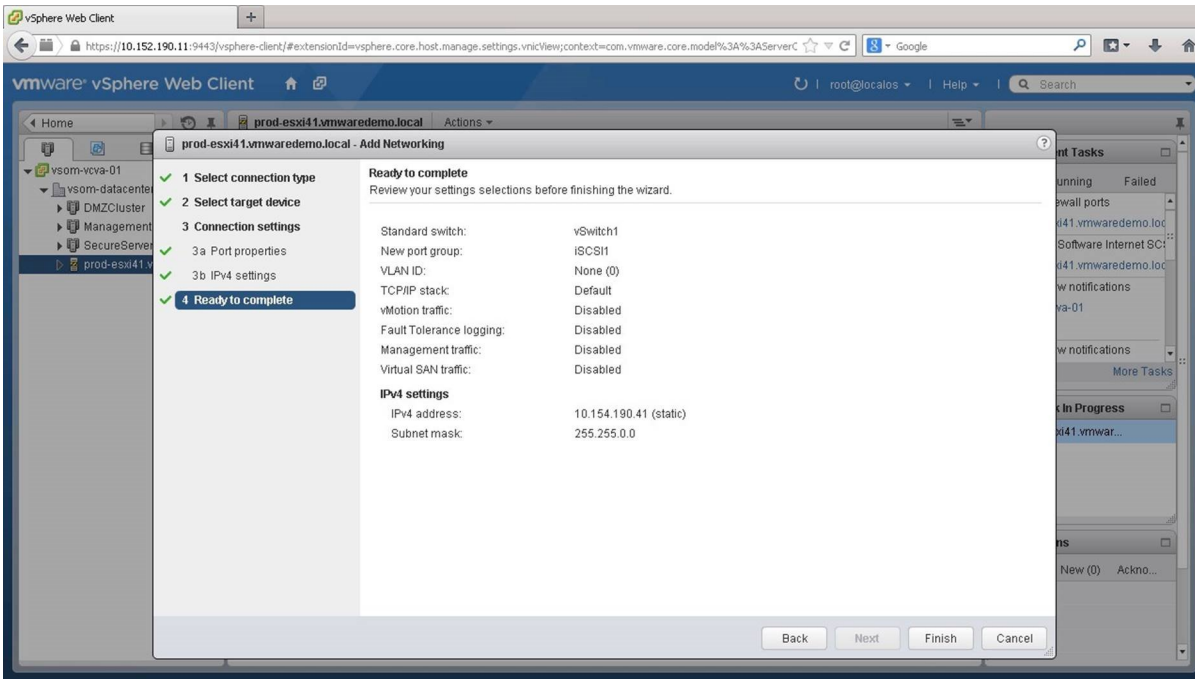
We select the standard switch [vSwitch1] and click on [OK] and click on [Next].



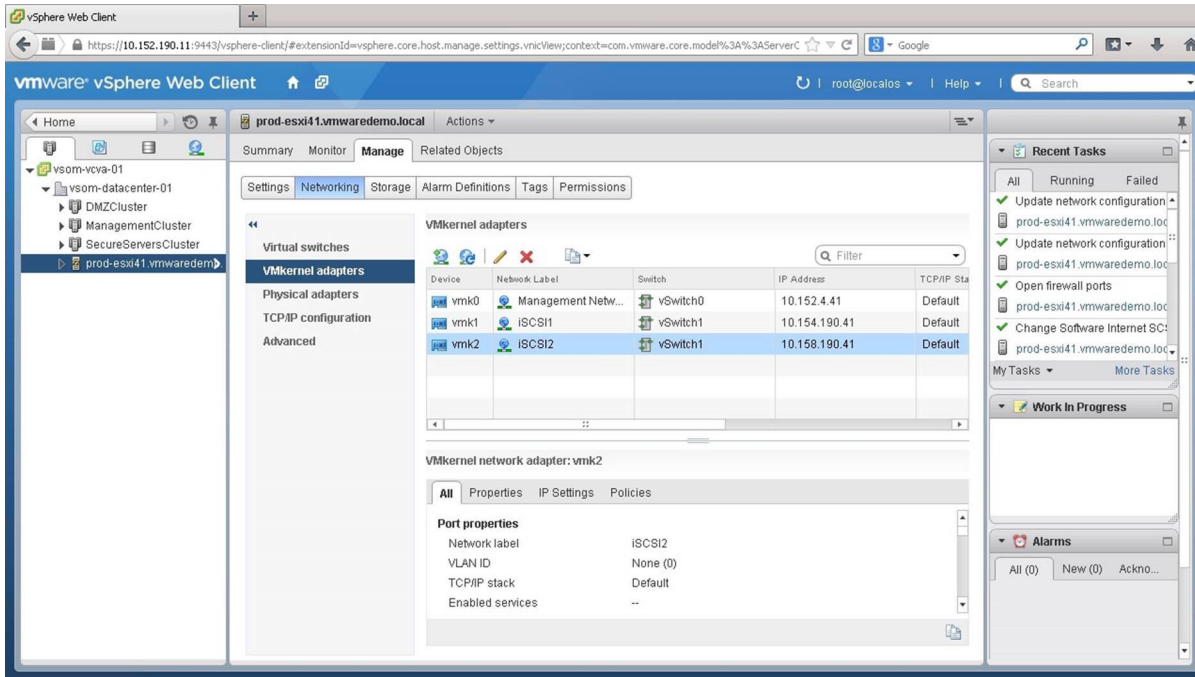
Assign a name to the network and click on [Next].



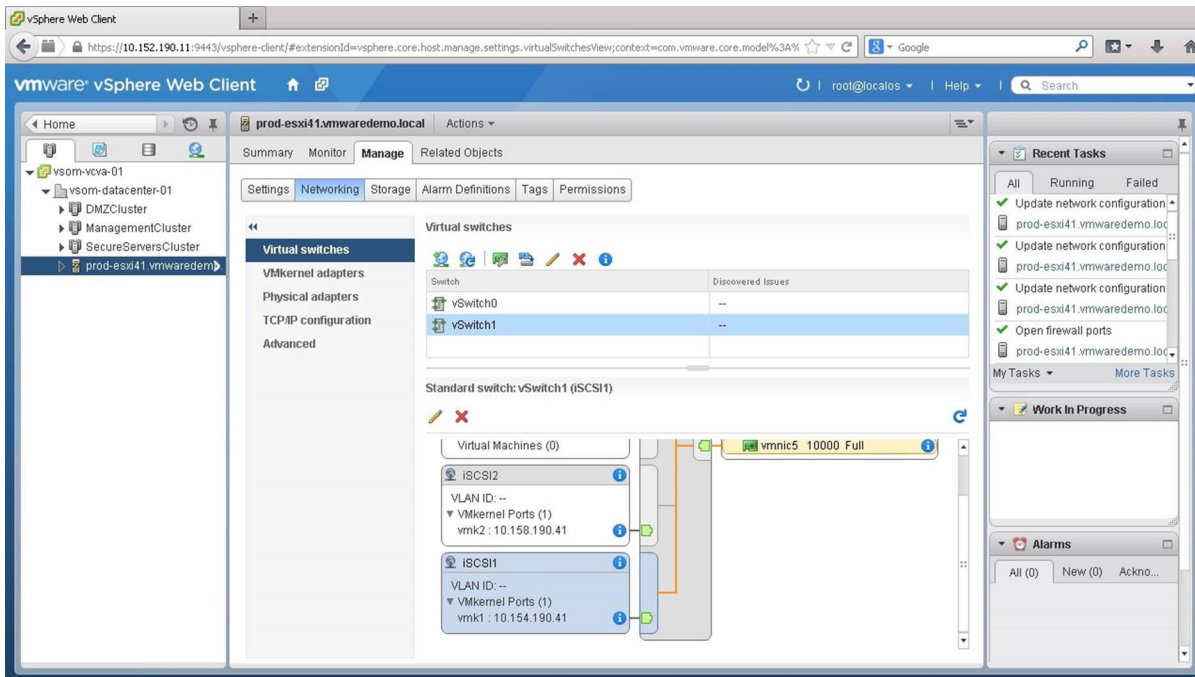
Configure the IP settings. We choose to use static IP addresses and click on [Next].



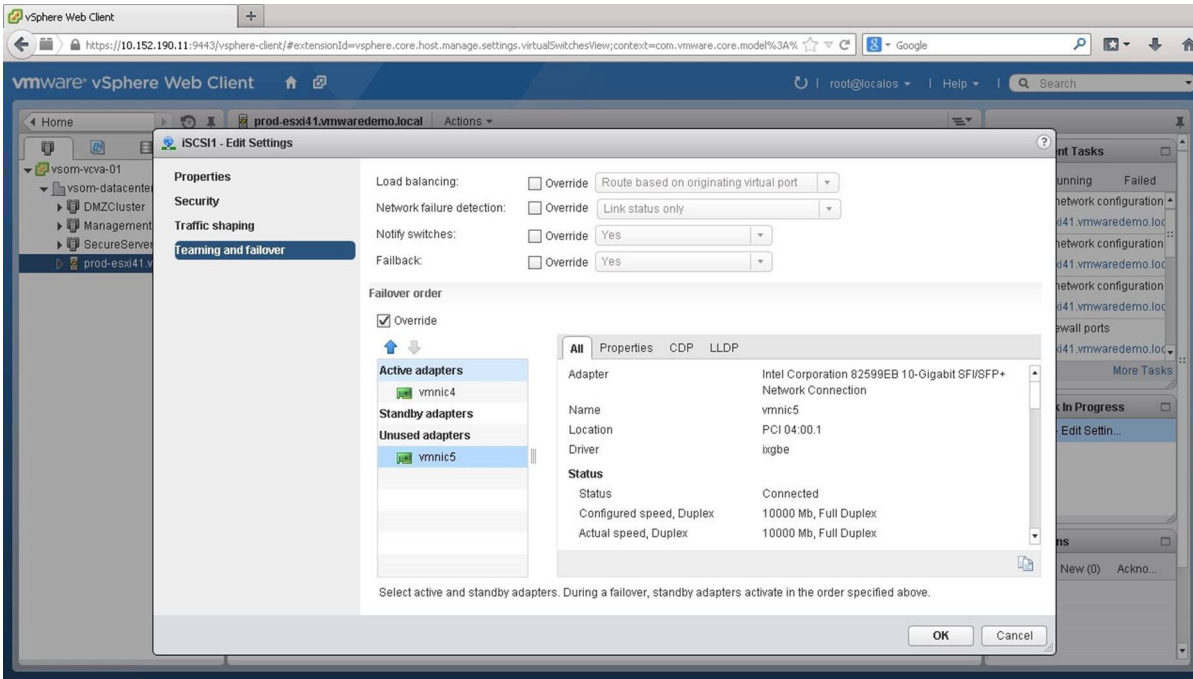
Review the settings and click on [Finish].



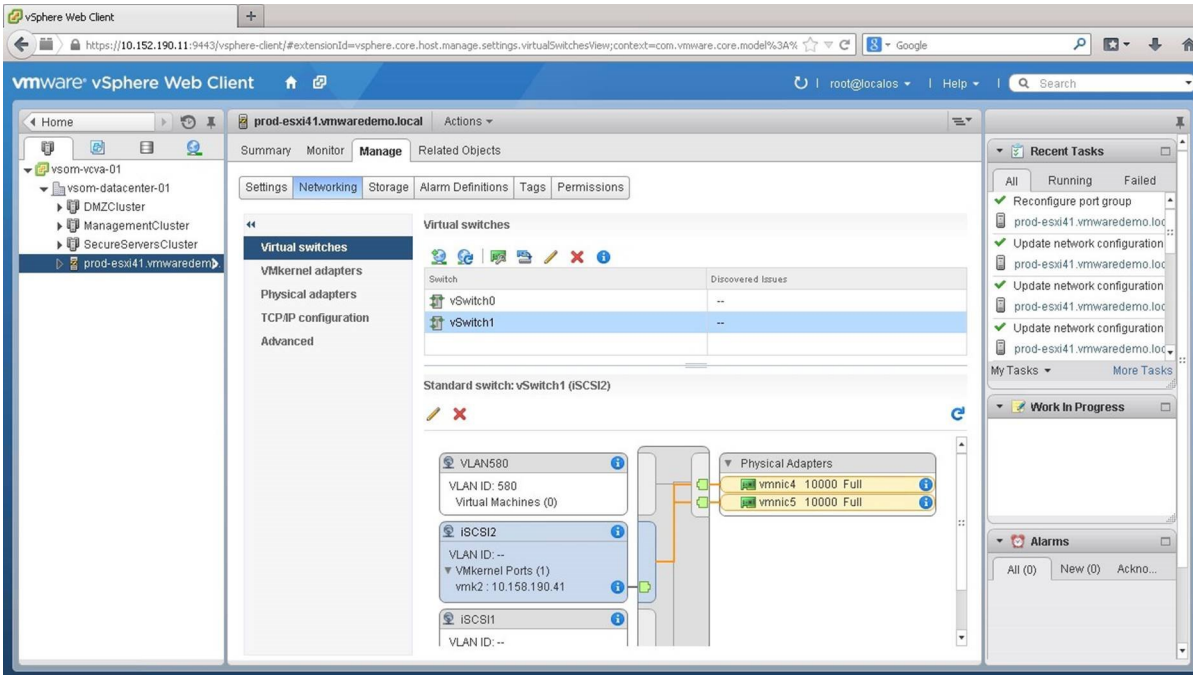
As covered in the introduction, we need to have multiple VMkernel adapters configured across multiple physical NICs and physical switches. So we have also created another VMkernel Adapter here. Note that the second iSCSI network will need to be configured with a different IP address. The ESXi software iSCSI initiator supports a single iSCSI session with a single TCP connection for each iSCSI target. Next we go into [Virtual Switches].



We need to ensure that the TCP connections always traverse the same network as the initiator connects to the iSCSI. We will configure teaming and failover settings. We first select [iSCSI1] and click on the [Edit] icon.

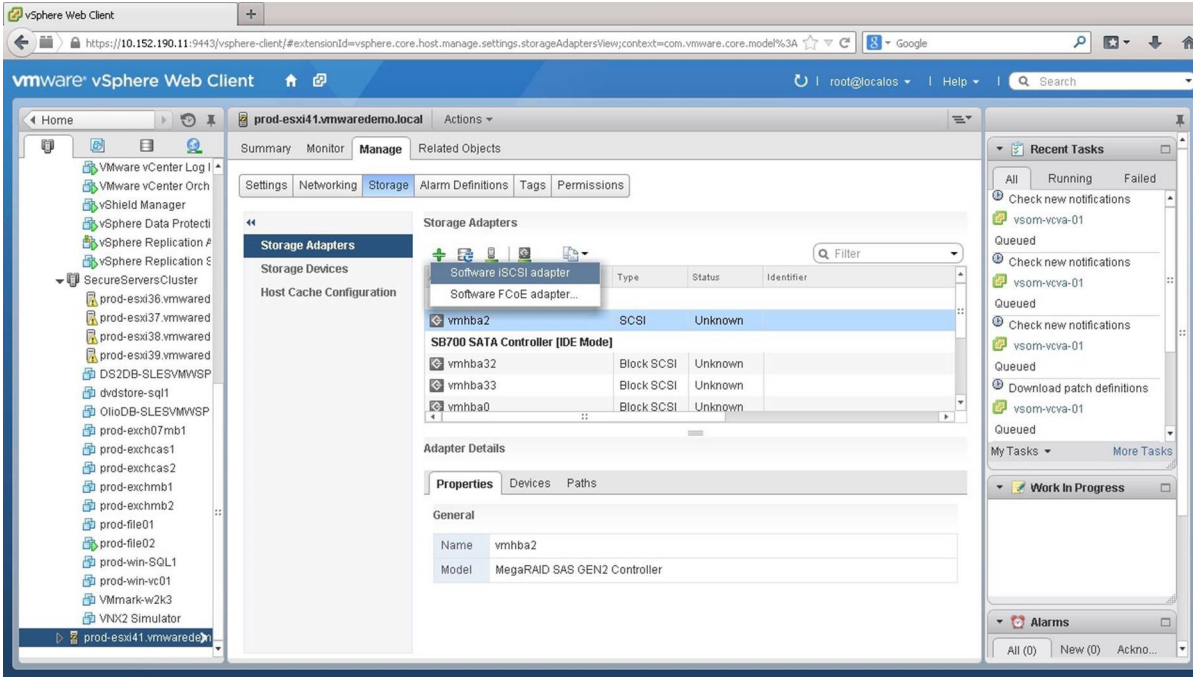


Go to [Teaming and Failover], click on the check box to [Override] failover border and make "vmnic4" as the only active adapter on iSCSI1 and click [OK].

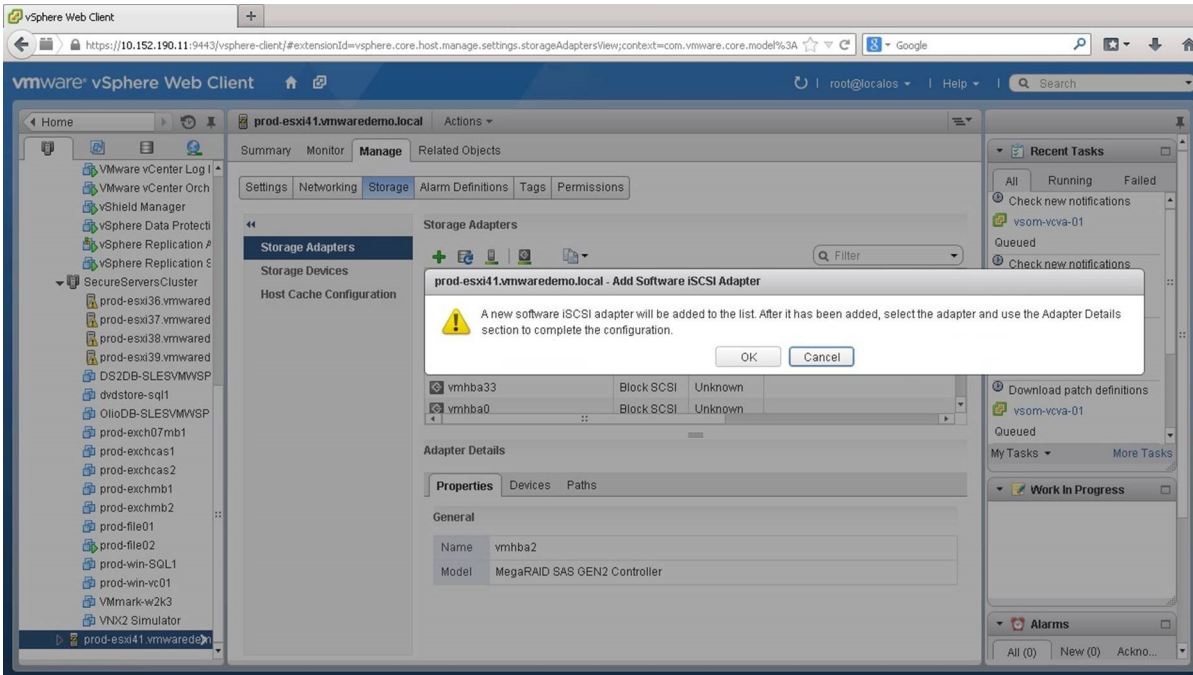


We repeated the same process on the iSCSI2 and have changed the failover order to use "vmnic5" as active, and "vmnic4" as unused.

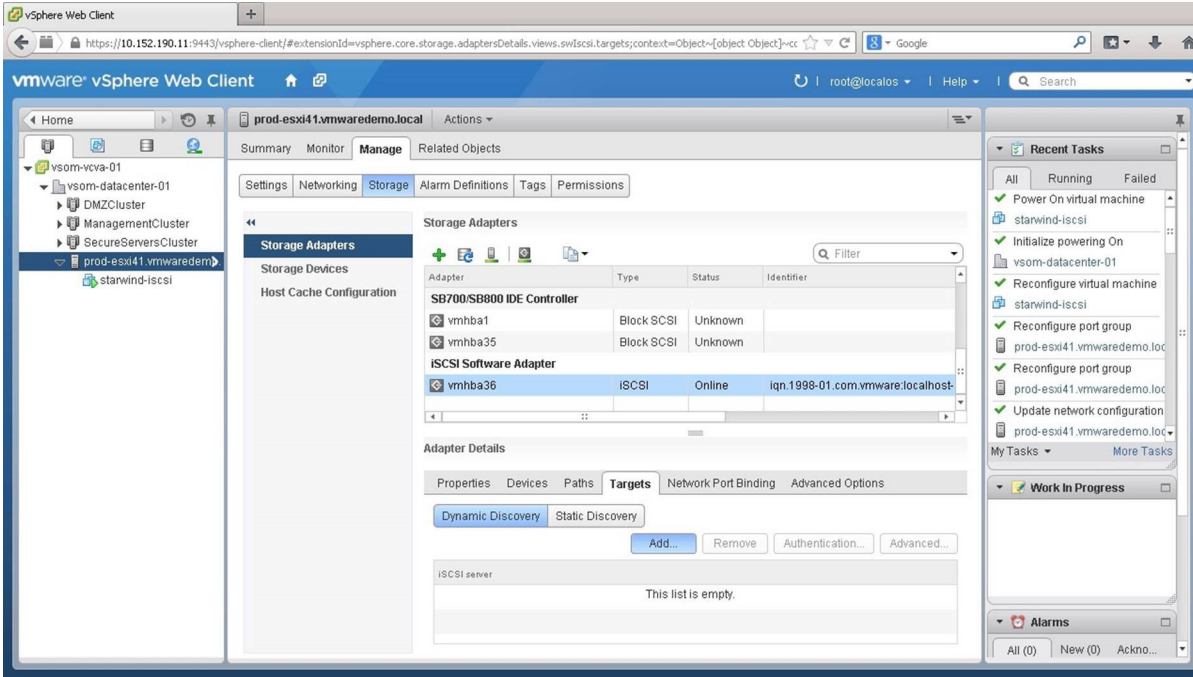
# NFS & iSCSI Multipathing in vSphere



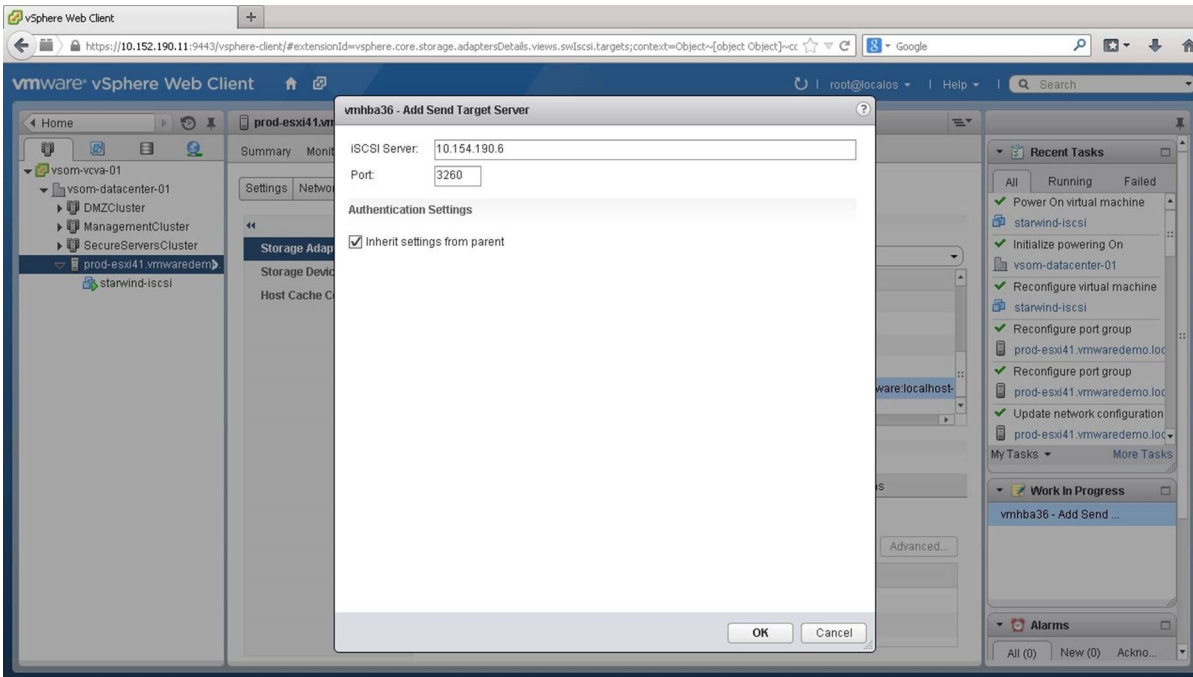
Here we see that we have a list of local storage adapters and no iSCSI adapter, so we will add a new iSCSI adapter by clicking on the add [+] icon and selecting [Software iSCSI Adapter].



Click [OK] on this prompt.

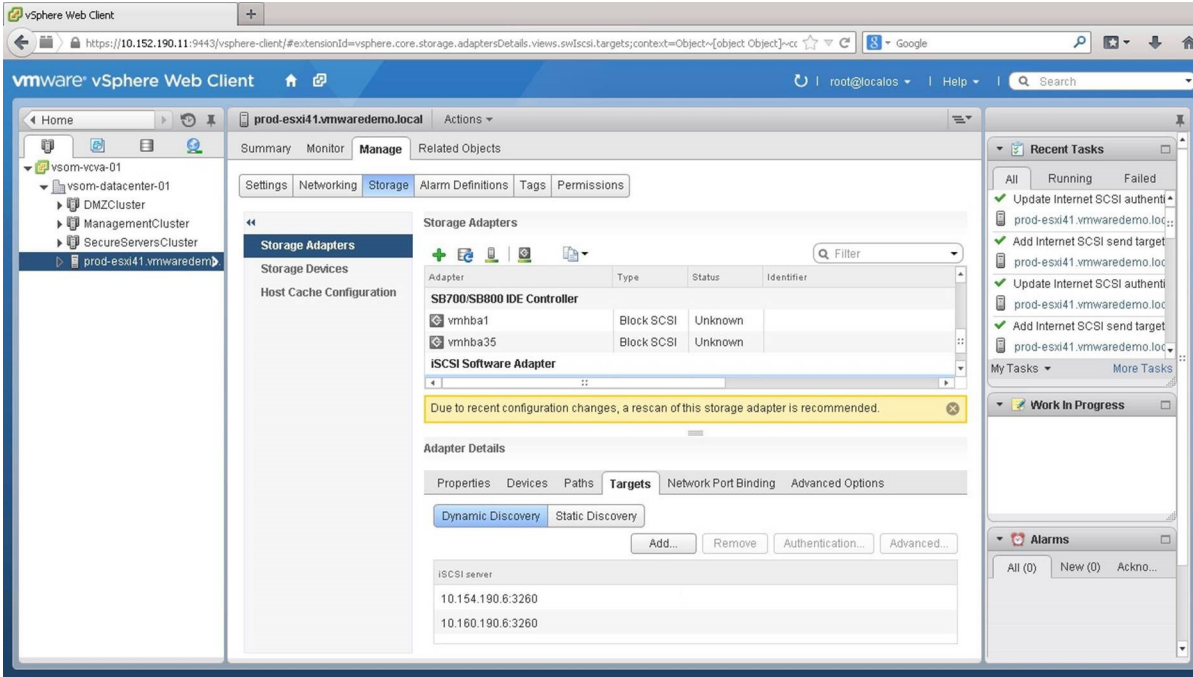


Here we will add the IP addresses of the iSCSI storage array. Click on [Add].

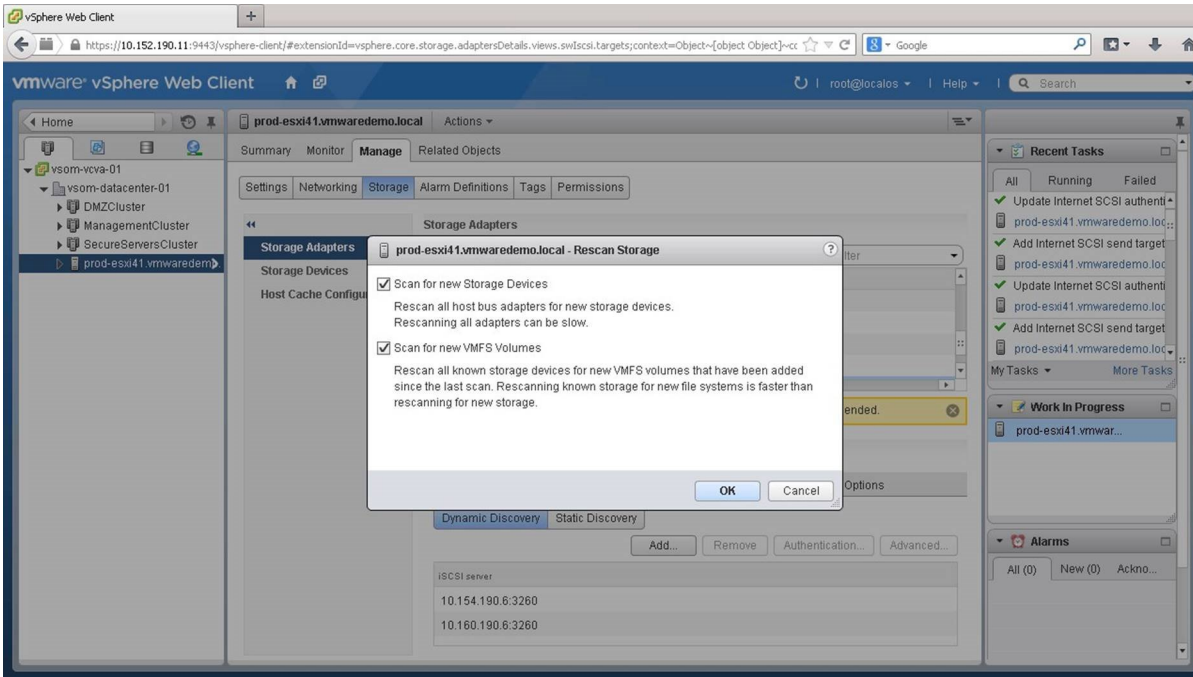


We specify the IP and the port and click [OK].







In our demo environment, the iSCSI array has two controllers with a single NIC each. So we have also added a second target address here. With both the targets added, we will now rescan the iSCSI software adapters to find the LUNs that are available. Click on the [Rescan] icon.



Select the areas that need to be scanned and click on [OK].

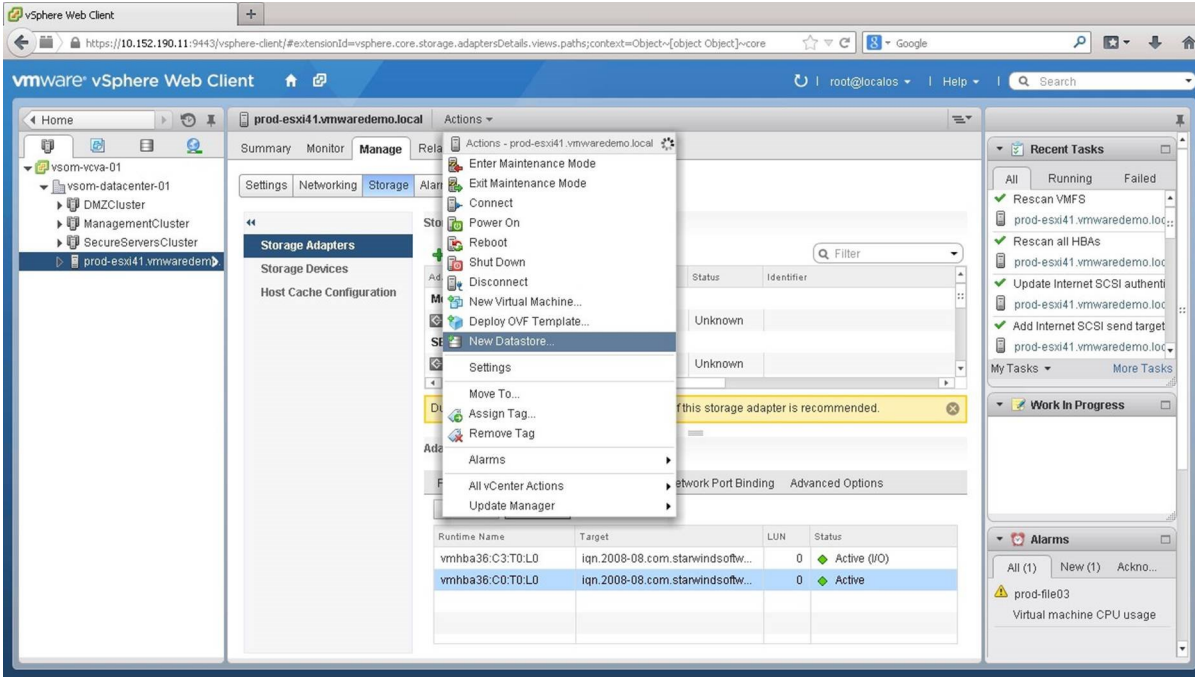


Once the scan is complete; we go into the [Devices] tab and see that an iSCSI disk has been found. We then switch to the [Paths] tab.

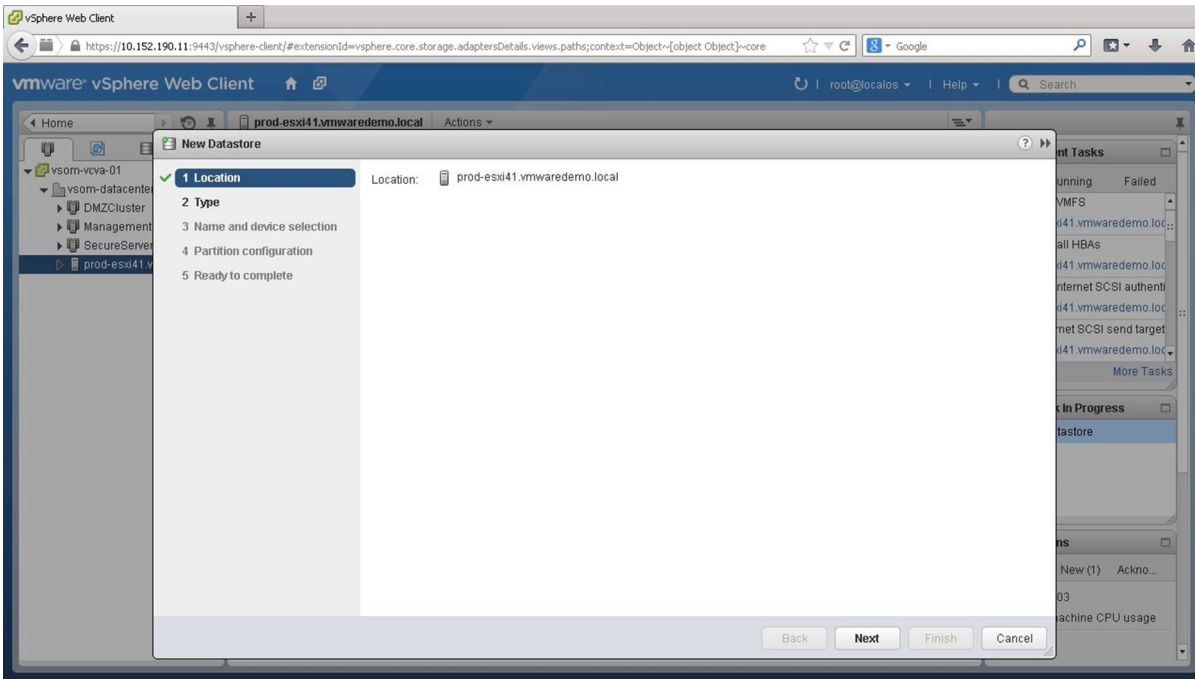


We see that there is one path for each VMkernel adapter that was configured. This enables high availability to our iSCSI storage array. With the new LUN found, we will go ahead and configure a new VMware file system on the newly found LUN. We click on [Actions].


# NFS & iSCSI Multipathing in vSphere



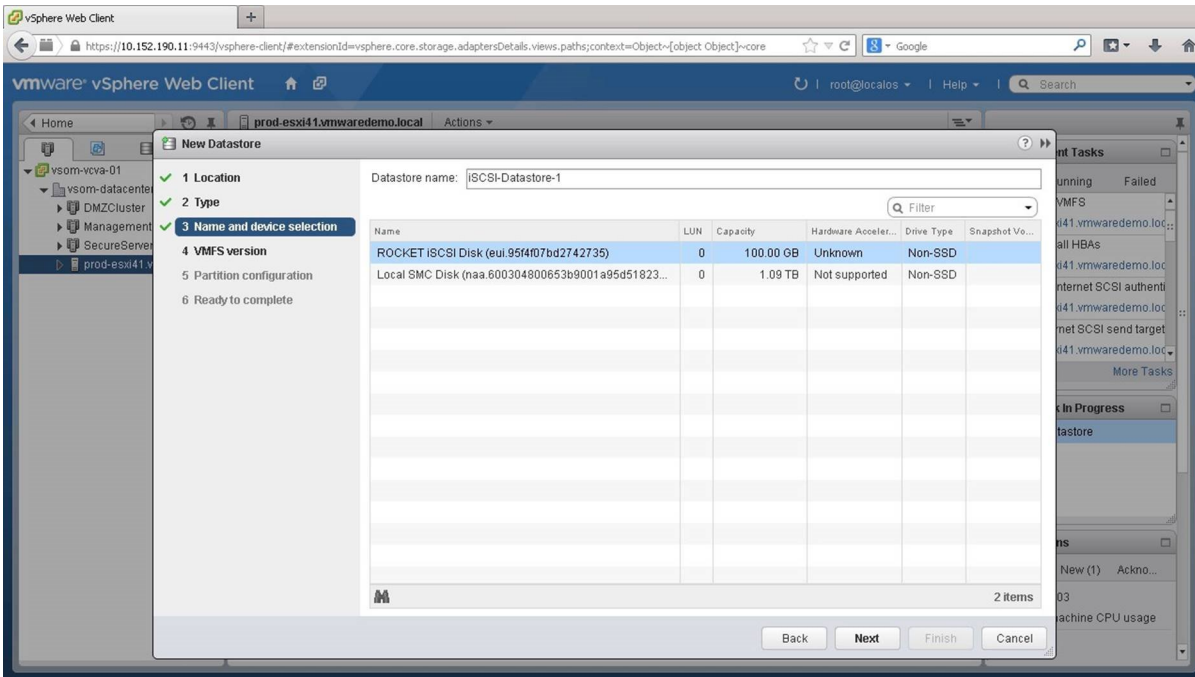
Click on [New Datastore].



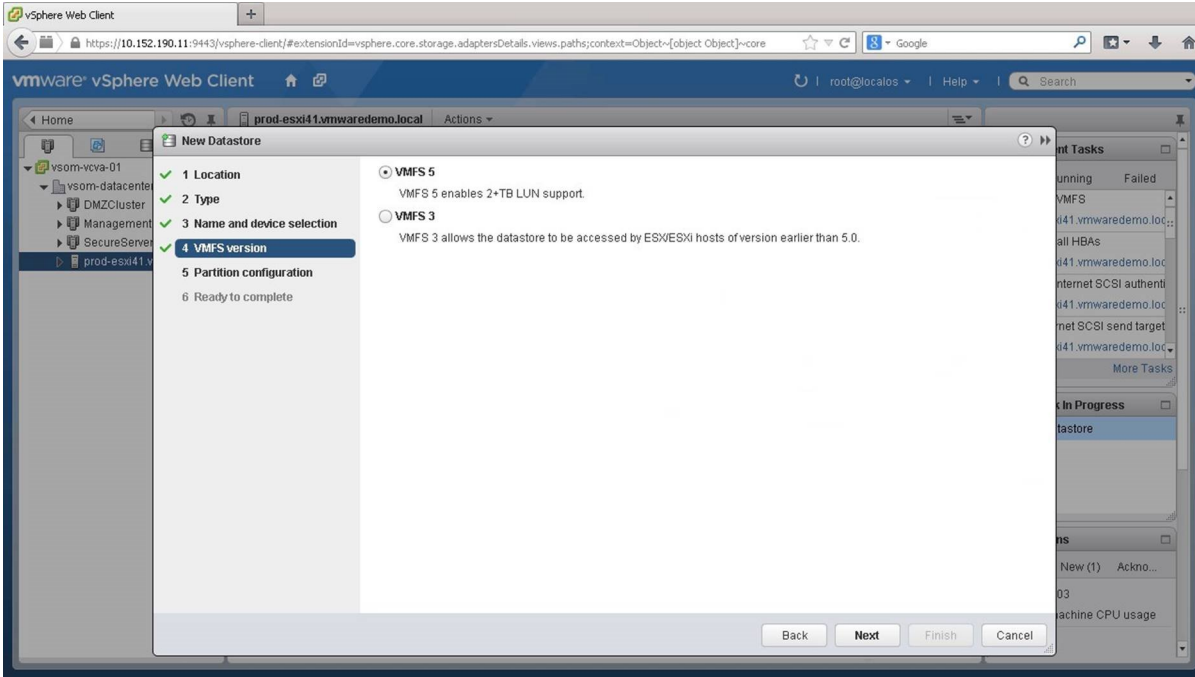
Click [Next].




Select VMFS and click on [Next].



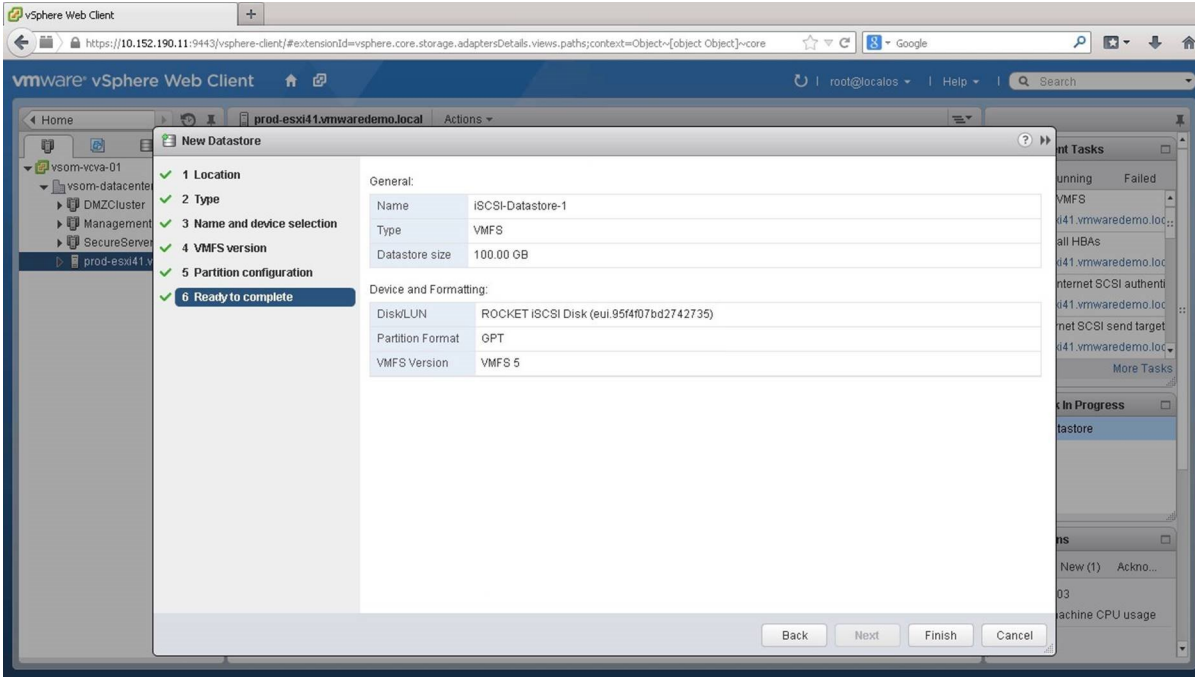
Assign a name and select the newly found LUN. Click on [Next].




Select the VMFS version and click on [Next].



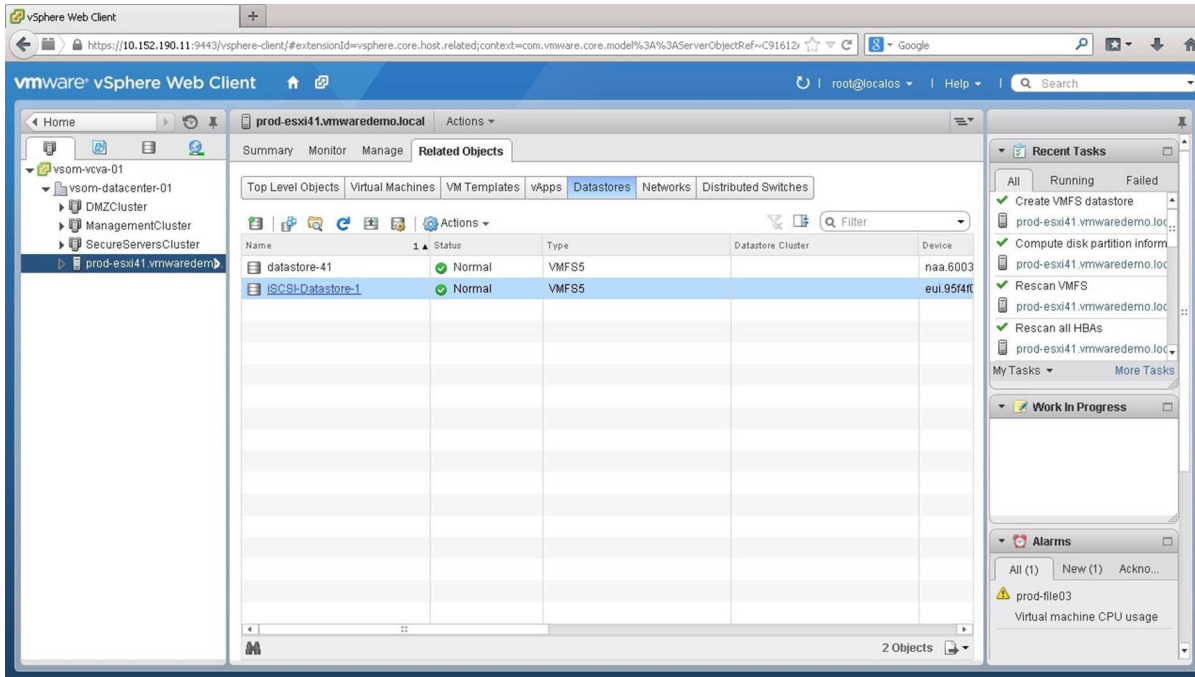
Choose the partition layout and click on [Next].



Review the settings and click on [Finish].



The VMFS datastore has been created on the iSCSI LUN. We go into the [Related Objects] tab.



Under [Datastores], we see that the new iSCSI datastore is now available. Here we can create virtual machines and virtual machine templates. This concludes the walkthrough on iSCSI storage. Select the next walkthrough of your choice using the navigation panel.

## 5 Storage DRS FAQ

## vSphere 6.7 Storage

This whitepaper describes in detail the various features of vSphere 6.7 Storage.

Please visit: [vSphere 6.7 Storage](#)



