



# Introduction to Storage Virtualization

VMware Storage

## Table of contents

Introduction to Storage Virtualization .....	3
Intro to Storage Virtualization .....	3
vSphere 6.5 Storage Features .....	10
NFS 4.1 Multipathing Configuration and Best Practices .....	11
Paths in same Subnet .....	11
Steps to configure .....	11
Best practice configuration in same subnet .....	12
Paths in different Subnets .....	12
Steps to configure .....	13
Best practice configuration for different subnets .....	13
Physical Switches .....	14
Viewing the Multipath connections .....	14
Configuring iSCSI Storage .....	15
vSphere Core Storage .....	31
vSphere 6.7 Storage Features .....	32
vSphere 6.7 Storage .....	32

## Introduction to Storage Virtualization

### Intro to Storage Virtualization



# Getting Started with vSphere with Operations Management

## Introduction to Storage Virtualization

This walkthrough demonstrates the concept of shared storage in a vSphere with Operations Management environment. Use the arrow keys to navigate through the screens.

### Single-Server Environment



Hosting VM Files



Virtual HDD



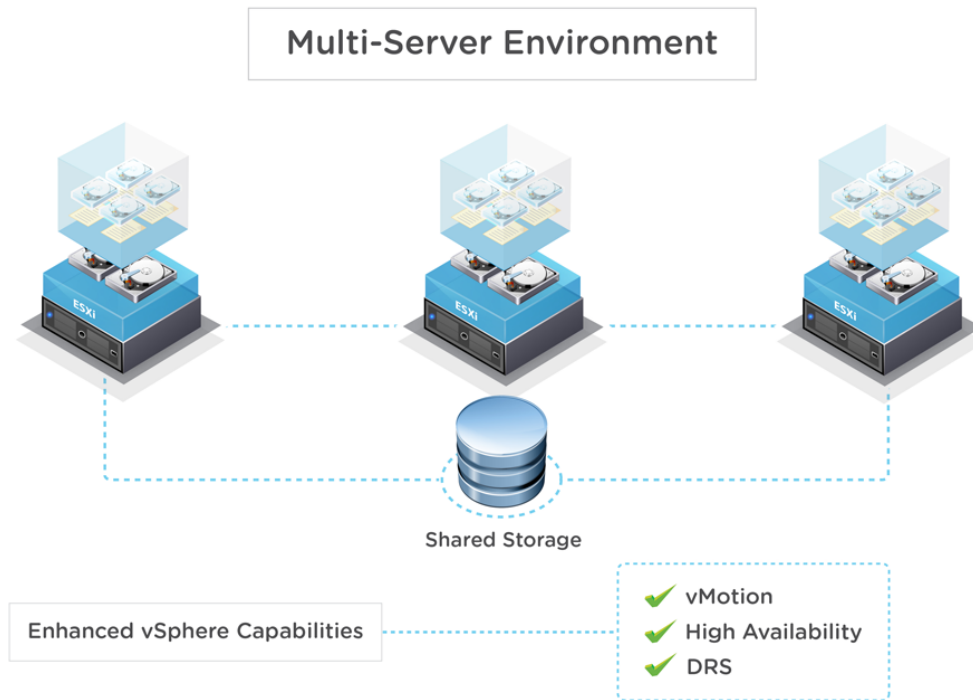
High Availability



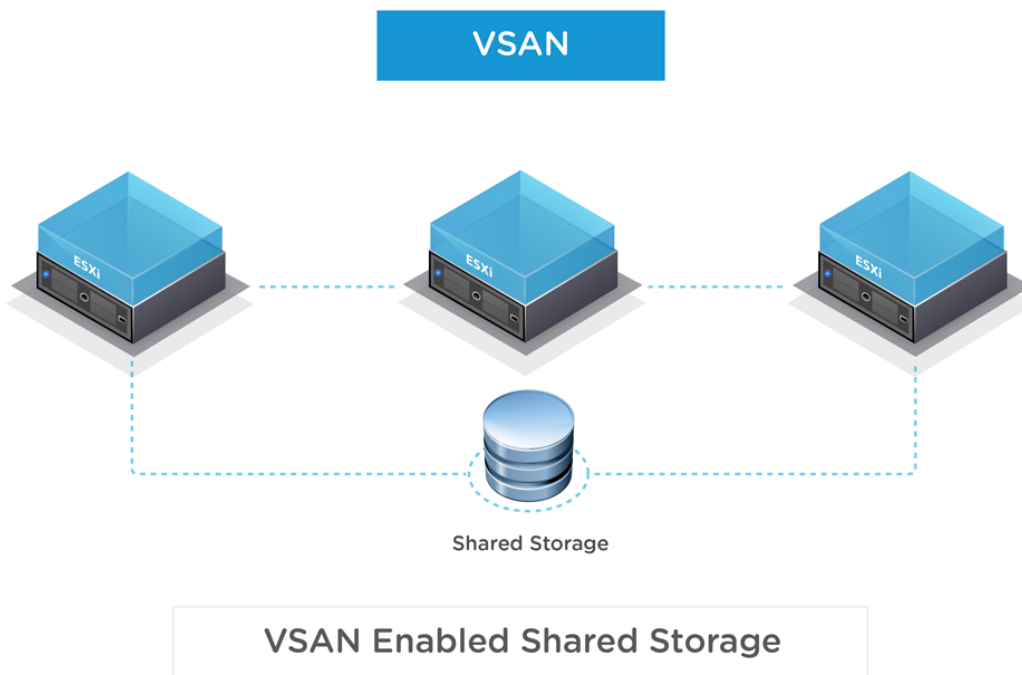
Scale Out




Before we walk through the configuration of iSCSI, let's review some storage basics. Internal hard-drives in a stand-alone ESXi host can be used to host VMware Virtual Machine files including configuration files and virtual hard drives. While this is acceptable for a single-server environment, a single server's internal disks do not provide highly available storage, nor does this solution scale well.




When multiple ESXi servers are clustered together, we can take advantage of VMware vMotion, High Availability (HA) and Distributed Resource Scheduler (DRS) to provide enhanced functionality, availability and manageability. Shared storage provides a common pool of storage for clustered hosts to access, enabling advanced vSphere features.




Shared storage can be provided in several different ways. VMware vSAN aggregates a mix of solid state and magnetic internal hard drives of multiple physical servers into a logical, highly available, high performance datastore.



An external storage array is another option for providing shared storage to multiple ESXi hosts. This storage can be presented to ESXi hosts using either file or block-based storage protocols. In file storage, the storage array (sometimes called a filer) creates and owns a file system, presenting the file system to a host to use. VMware ESXi can mount file-based storage using the NFS protocol.



In block storage, the storage array presents a raw set of hard drive blocks called a Logical Unit (LUN) to the connected hosts. The hosts are responsible for formatting and creating a file system on that space. Block-based storage, plus the storage adapters (HBA) and storage fabric (switches and cabling) are known as a Storage Area Network, or SAN. ESXi supports Fiber Channel (FC), Fiber Channel over Ethernet (FCoE) or iSCSI SANs. We'll focus on iSCSI for the remainder of this walkthrough.



vmware®

File


Block

ESXi hosts connect to iSCSI SANs using either 1Gb or 10Gb Ethernet connections. The SAN presents a raw LUN to ESXi. ESXi formats the LUN with the VMware File System (VMFS). Virtual Machines are stored in the VMFS datastore. Note: Redundancy is important to a reliable VMware vSphere environment. As such, iSCSI should be connected in a highly available manner.




vmware®


In a typical vSphere environment, we combine two or more servers that have VMware ESXi installed into a cluster. Those hosts access the same shared storage datastores.




The physical servers should have two or more storage adapters. For iSCSI, this means that two or more 1GbE or 10GbE NICs should be available in each physical ESXi host.



These NICs are connected to a pair of Ethernet switches. These switches should be capable of handling the expected amount of traffic necessary for the I/O activity of all the virtual machines' storage traffic.

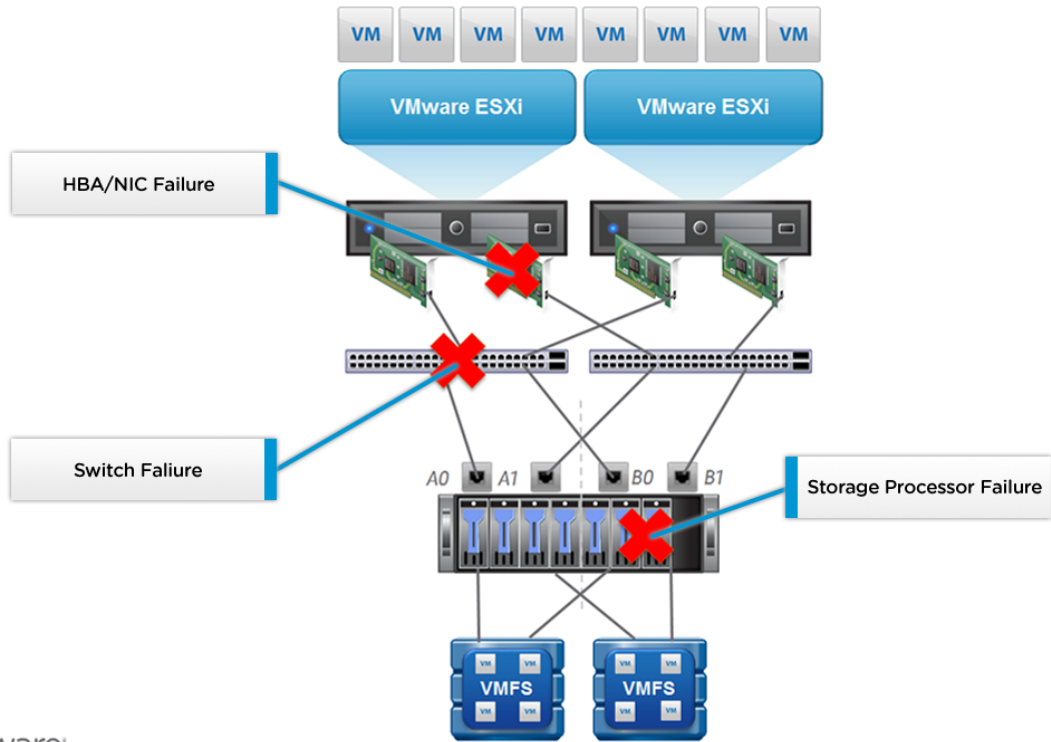


Business-class storage arrays typically contain two or more storage controllers (A and B sides) for high availability. Each storage controller has two or more front-end ports for host connectivity. These ports are cabled to the Ethernet switches in a highly available manner.



LUNs are owned by one or both storage controllers depending on your array architecture. LUNs are formatted with VMFS. Virtual machine files are stored in VMFS. The resulting environment provides multiple paths for storage traffic to flow from an ESXi host to the storage array to a LUN.





If any component in the SAN fails - a cabling failure, a NIC failure, a switch failure, or a storage controller failure - connectivity between the host and the datastore will be maintained on one or more surviving paths. ['The Setting up iSCSI Storage'](#) walkthrough builds on the information presented here to help you configure iSCSI storage.

## vSphere 6.5 Storage Features

This whitepaper describes in detail the various features of the vSphere 6.5 Core Storage. Please visit:

[vSphere 6.5 Storage Features](#)

## NFS 4.1 Multipathing Configuration and Best Practices

With vSphere 6.0, ESXi incorporated the support for NFSv4.1 protocol. Multipathing is one of the feature which NFS4.1 provides. We will talk about configuring and best practices of configuring Multipathing feature of ESXi's NFS4.1 here.

## What is Multipathing?

Multipathing is a method to use multiple paths, each accessible through a different IP address, to access the storage. NFS 3 uses one TCP connection for I/O. As a result, ESXi supports I/O on only one IP address or hostname for the NFS server, and does not support multiple paths. Though we can configure multiple NICs in "NIC teaming" to access a particular IP address through different physical paths, this is not Multipathing

NFS 4.1, by contrast, provides multipathing for servers that support the session trunking. With session trunking, servers maintain the states per session. We can have multiple connections associated with a session. When the session trunking is available, we can use multiple IP addresses to access a single NFS volume.

## Path selection

Currently NFS4.1 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 one wants to add another path, 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 address 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/share1 -v nfs41share1
```

#### vSphere Web Client

The configuration screen using vSphere Web Client is as follows.

Getting Started Summary Monitor **Configure** Permissions VMs Resource Pools Datastores Networks Update Manager

New Datastore

1 Type  
2 Select NFS version  
3 **Name and configuration**  
4 Configure Kerberos authentication  
5 Ready to complete

**Name and configuration**  
Specify name and configuration.

1 If you plan to configure an existing datastore on new hosts in the datacenter, it is recommended to use the "Mount to additional hosts" action instead.

Datastore name:

Folder:   
E.g: /datastore-001

Server(s):  +  
E.g: nas, nas.it.com or 192.168.0.1

1 If the server that will back this datastore has trunking enabled, you can enter additional IPs above, which the ESXi host will use to achieve multipathing to this NFS server mount point.


Servers to be added

0 items Copy

Back Next Finish Cancel

### 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

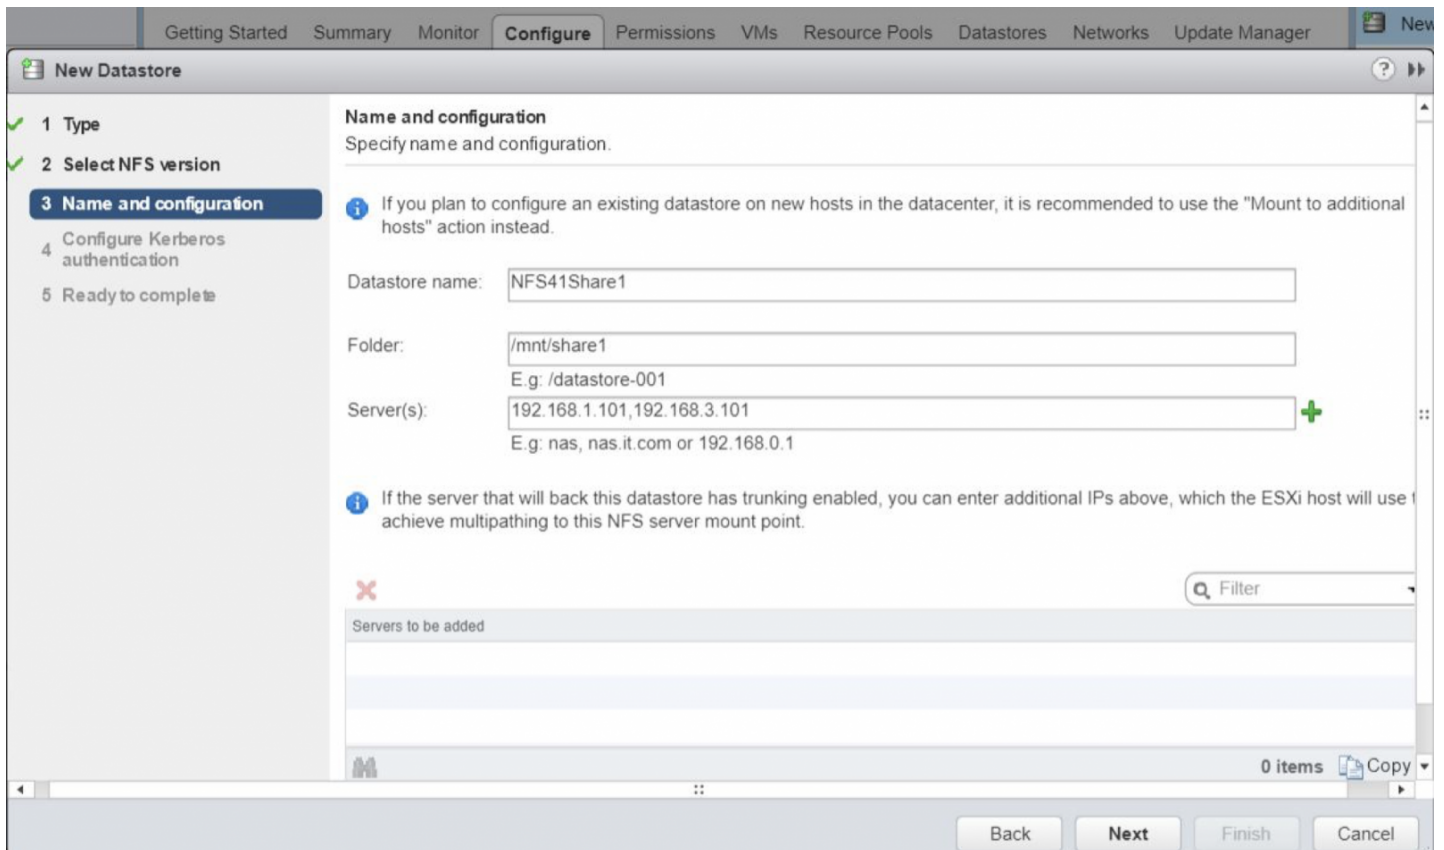
```
[root@w2-nfs-esx106:~] esxcli storage nfs41 -H 192.168.1.101,192.168.3.101 -s /share2 -v nfs41Share2
```

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

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

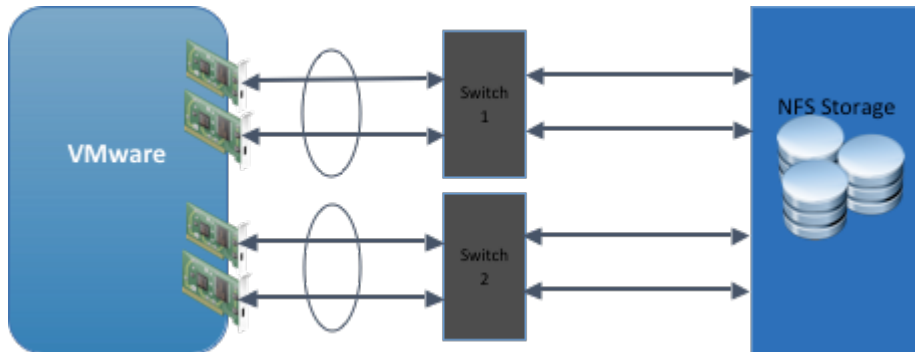
With different subnets we can have independent connections from different vmkernel portgroups. To utilize the network bandwidth and to have better redundancy at physical layer, configure each portgroup to be in different virtual switch and configure each vmkernel portgroup for NIC Teaming.

# 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  isPE  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/000000009A601F597856341200000000/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/000000009A601F597856341200000000/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
}
```

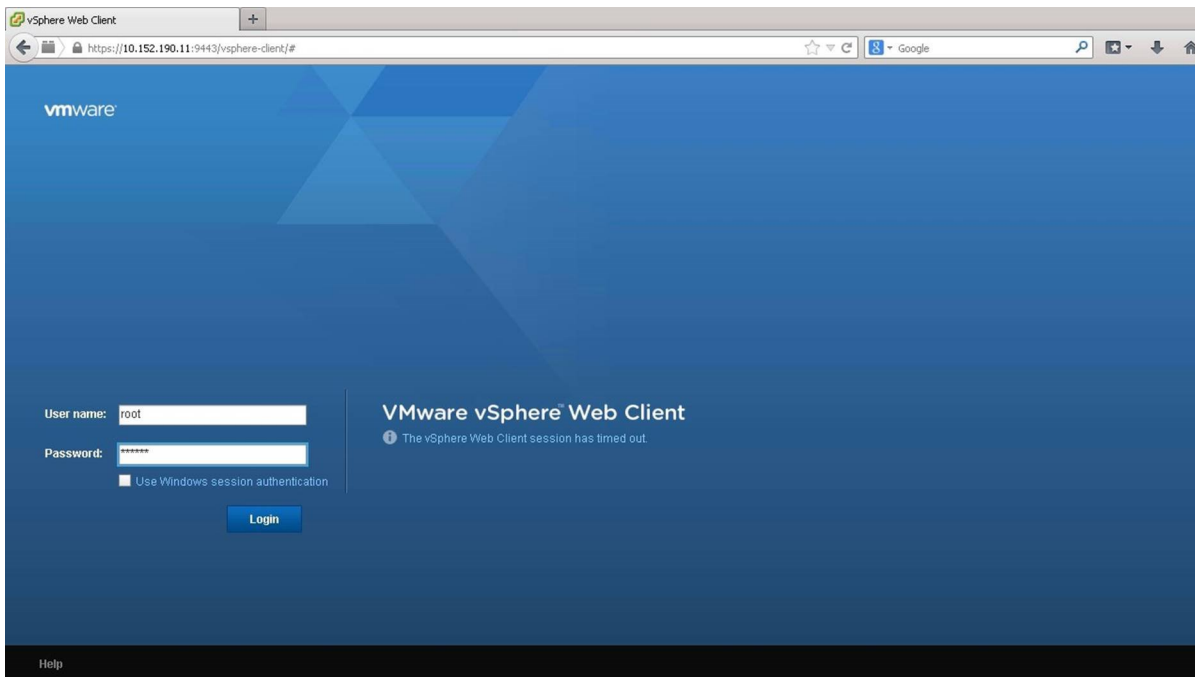
## Configuring iSCSI Storage




# 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 log on to the vSphere Web Client.

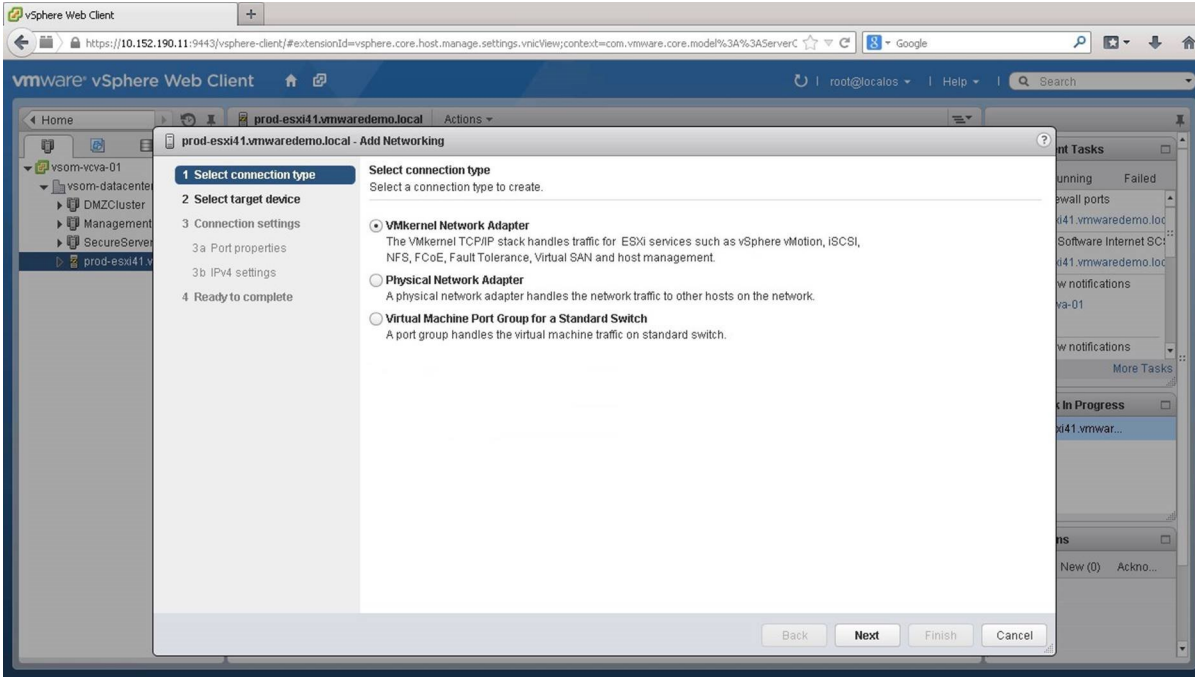


Go to the [Hosts and Clusters] view.




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






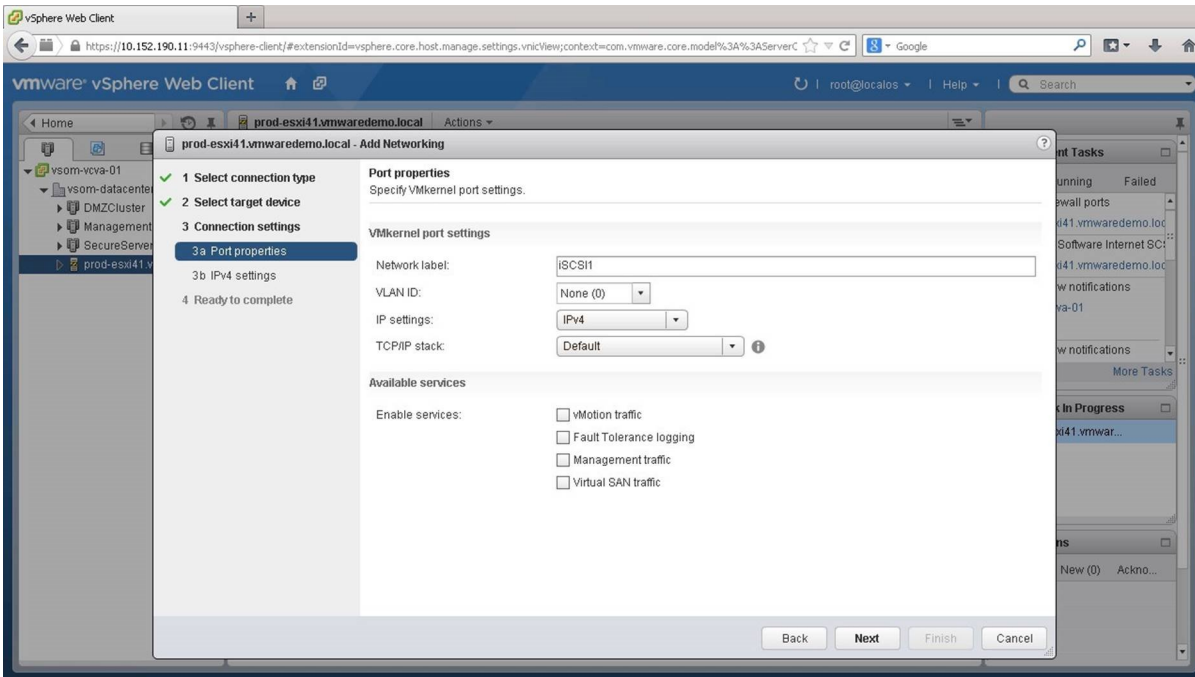
Select [VMkernel Network Adapters] and click on [Next].



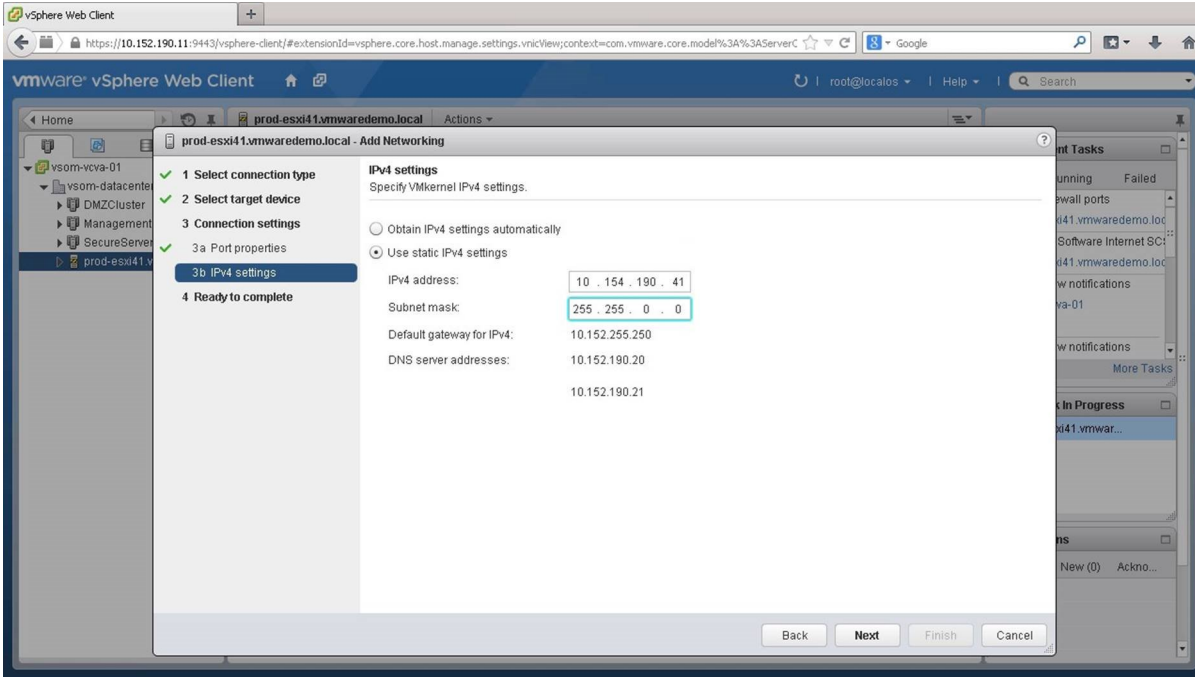
We click on [Browse] to select an existing standard switch. You can also choose to create a new standard switch.




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



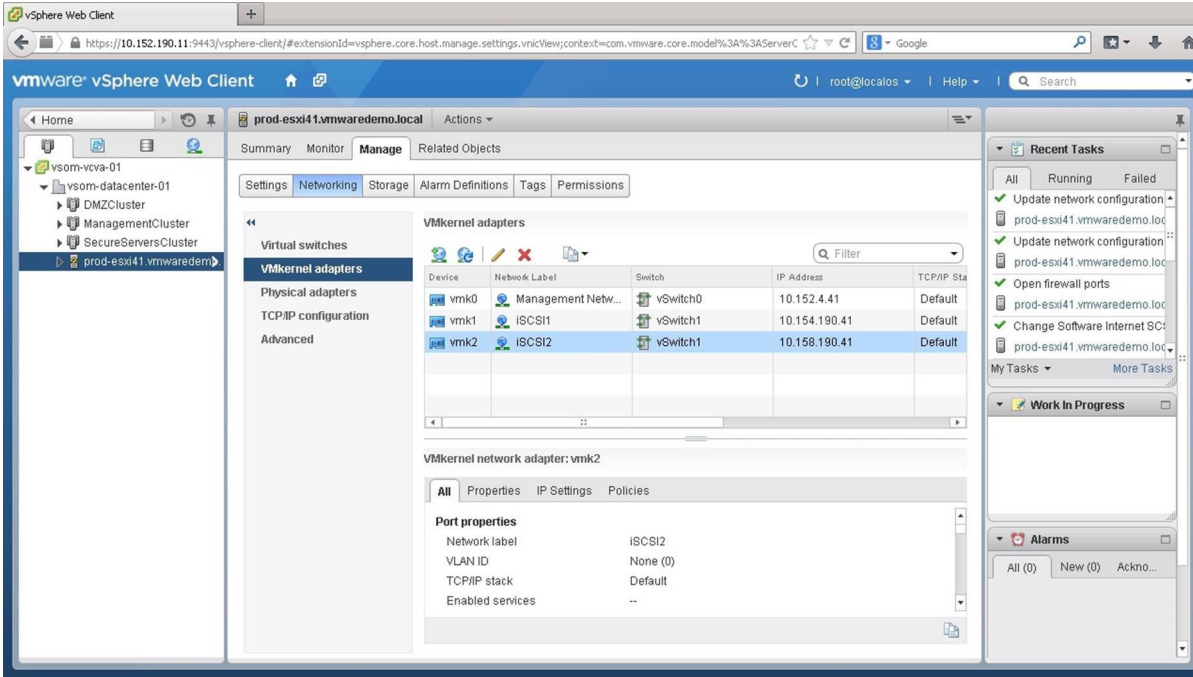
Assign a name to the network and click on [Next].



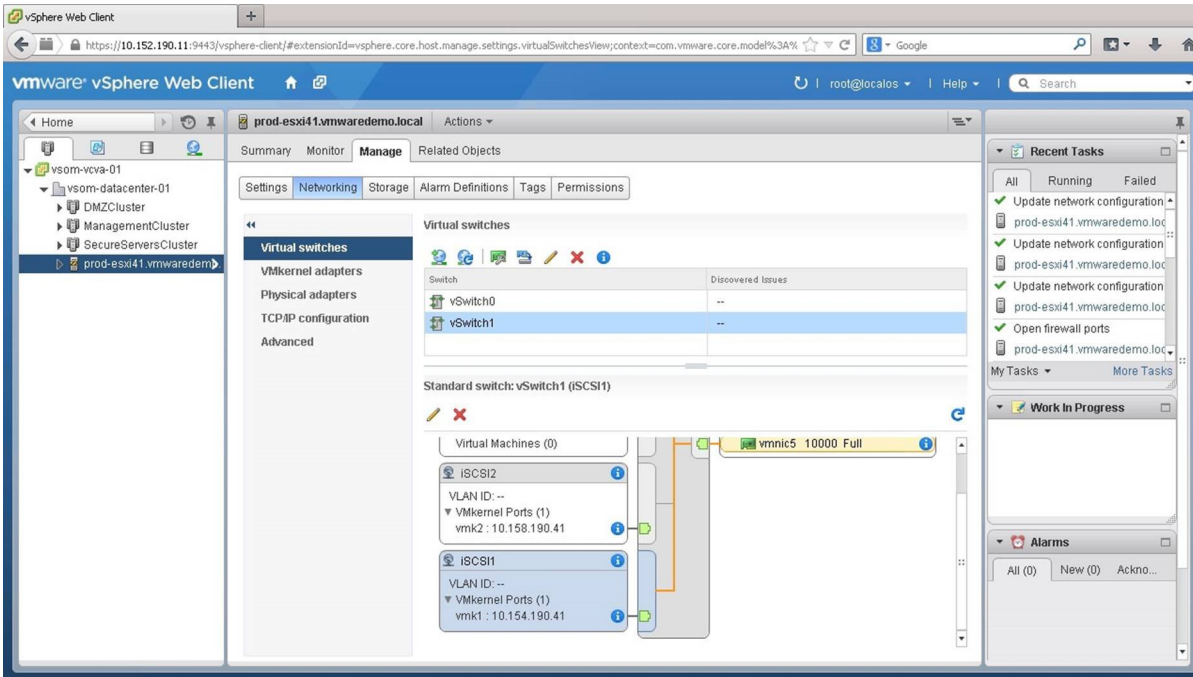
Configure the IP settings. We choose to use static IP addresses and click on [Next].



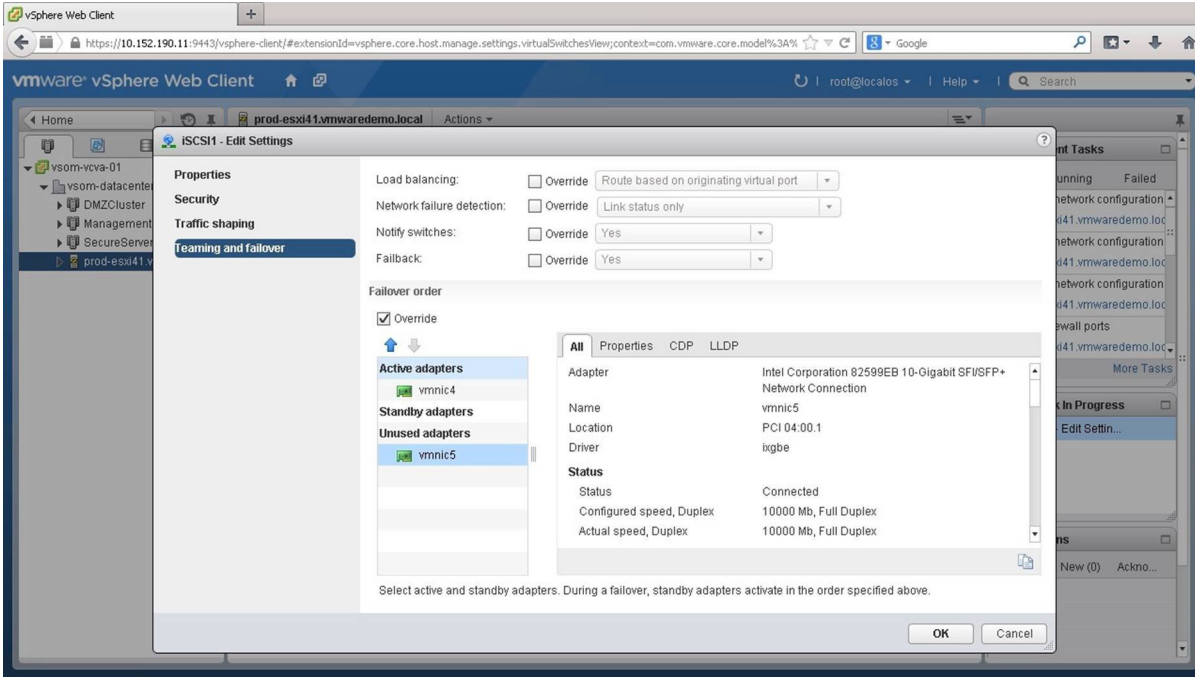
Review the settings and click on [Finish].



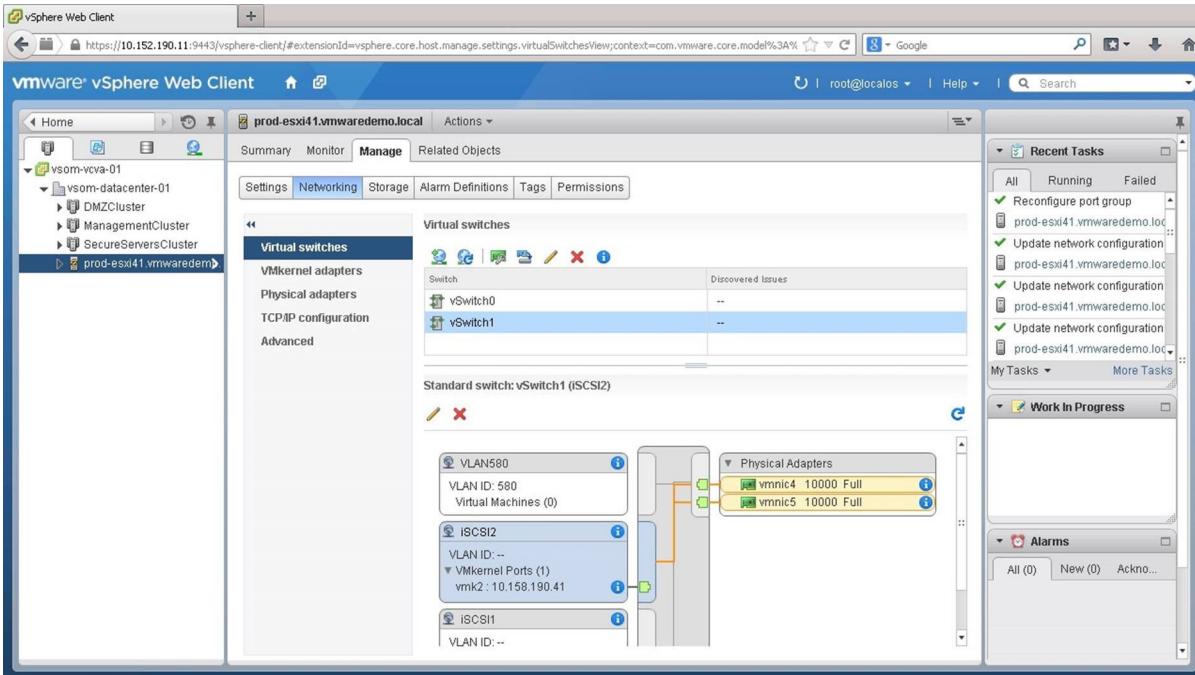
As covered in the introduction, we need to have multiple iSCSI adapters configured across multiple physical NICs and physical switches. So we have also created another VMkernel Adapter here. Note that the second iSCSI network will need to be configured with a different IP address. The ESXI software iSCSI 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 on [OK].



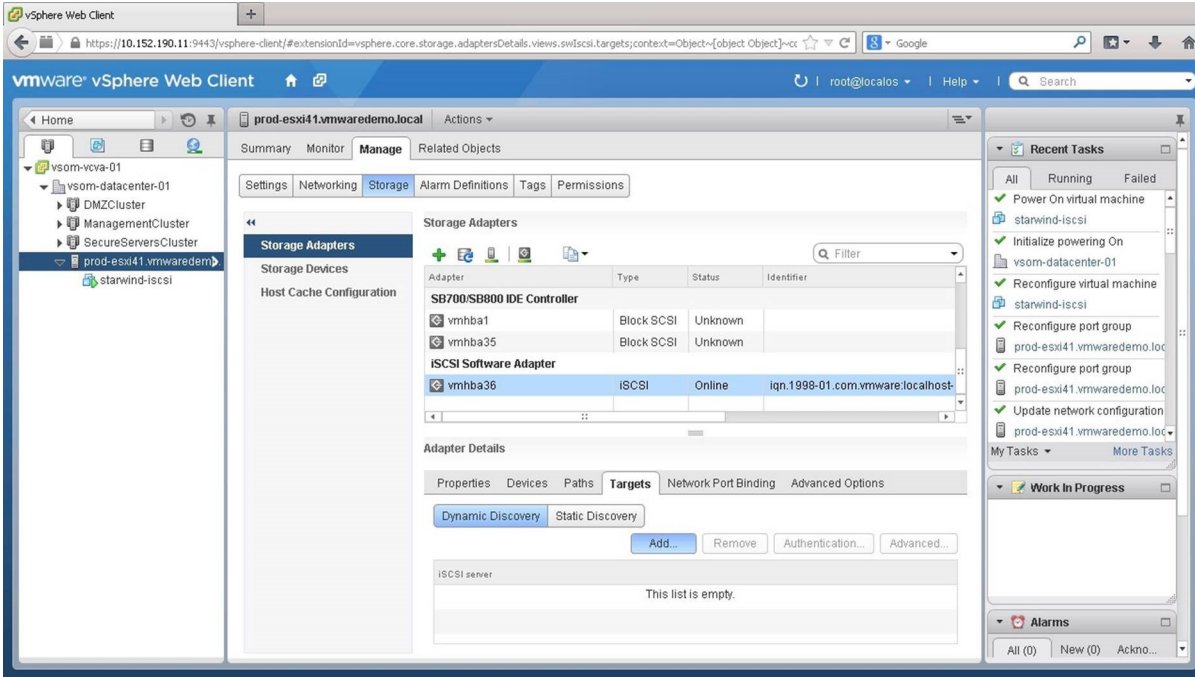
We repeated the same process on the iSCSI2 and have changed the failover order to use "vmnic5" as active, and "vmnic4" as unused.



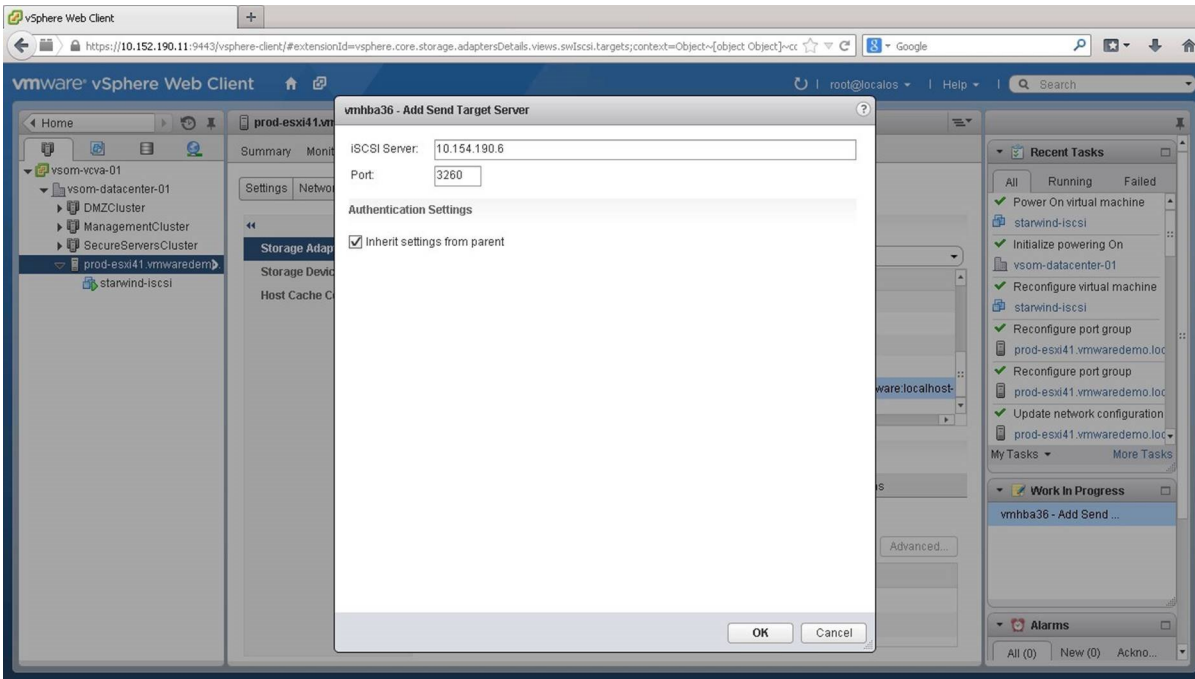
Here we see that we have a list of local storage adapters and no iSCSI adapter, so we will add a new iSCSI adapter by clicking on the add [+] icon and selecting [Software iSCSI Adapter].




Click on [OK] on this prompt.



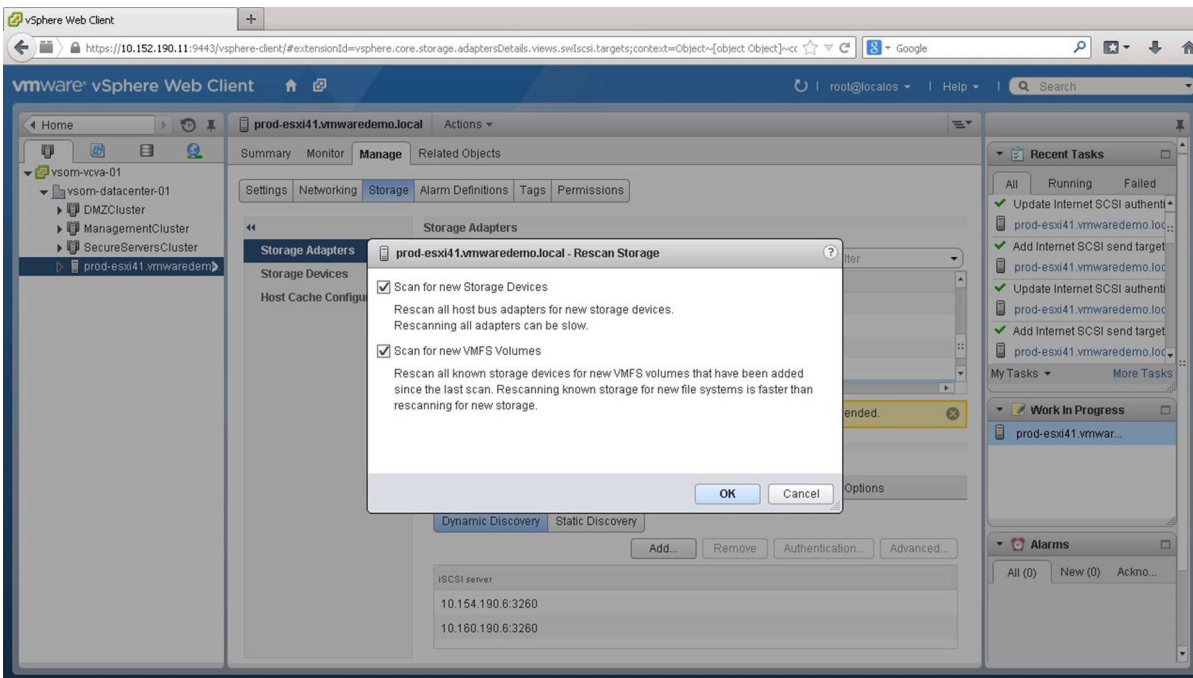
Here we will add the IP addresses of the iSCSI storage array. Click on [Add].



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

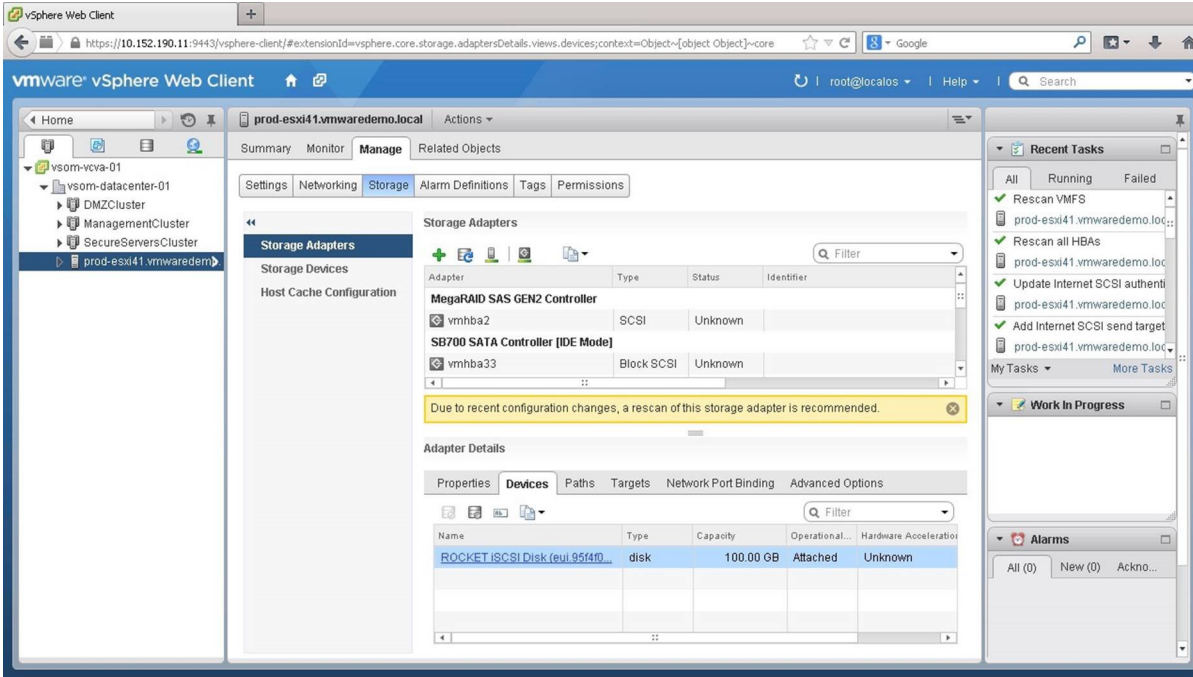


In our demo environment, the iSCSI array has two controllers with a single NIC each. So we have also added a second target address here. With both the targets added, we will now rescan the iSCSI software adapters to find the LUNs that are available. Click on the [Rescan] icon.

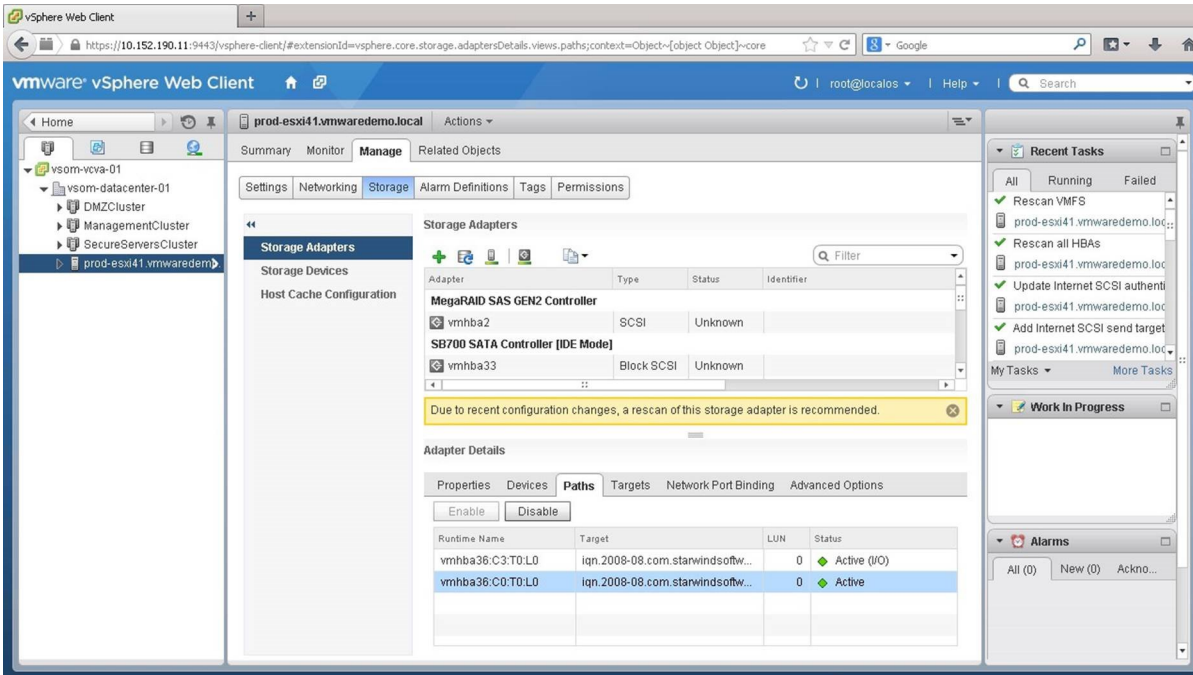


Select the areas that need to be scanned and click on [OK].







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.



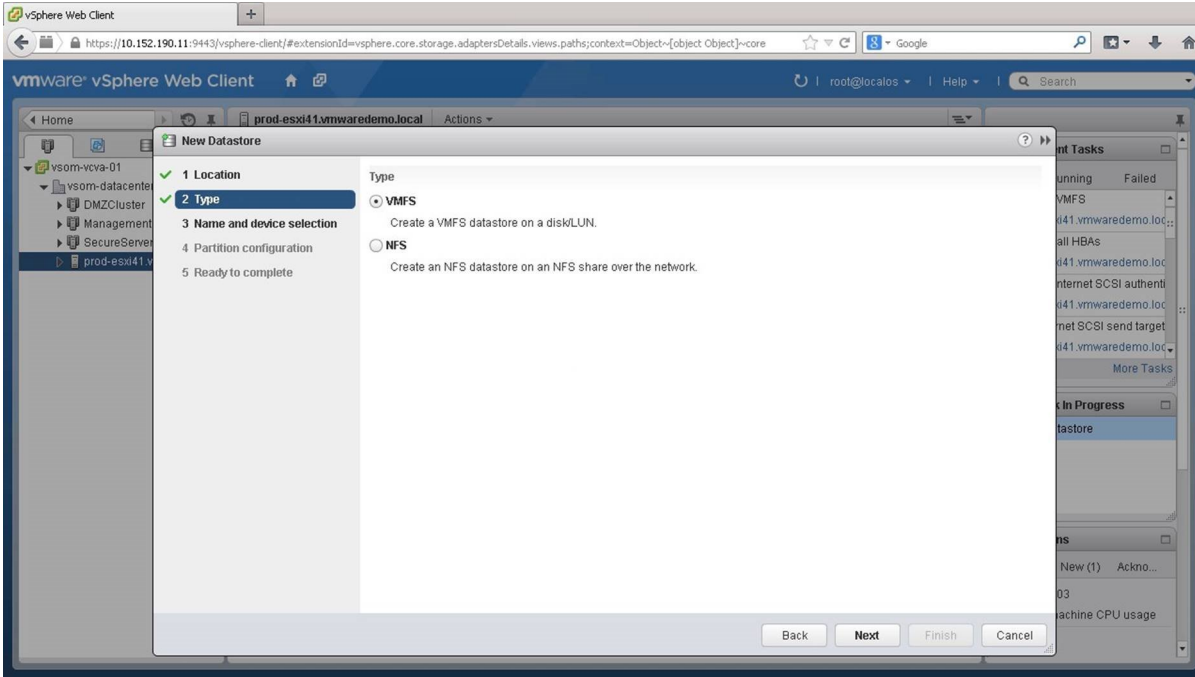
We see that there is one path for each VMkernel adapter that was configured. This enables high availability to our iSCSI storage array. With the new LUN found, we will go ahead and configure a new VMware file system on the newly found LUN. We click on [Actions].




Click on [New Datastore].



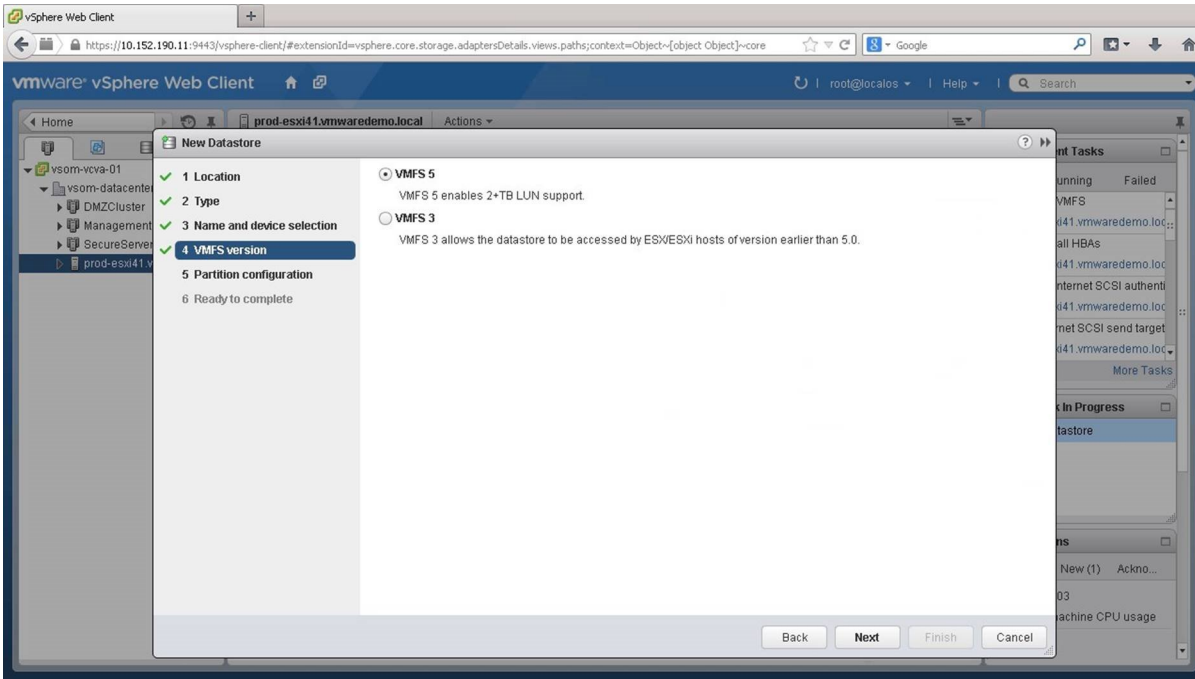
Click on [Next].



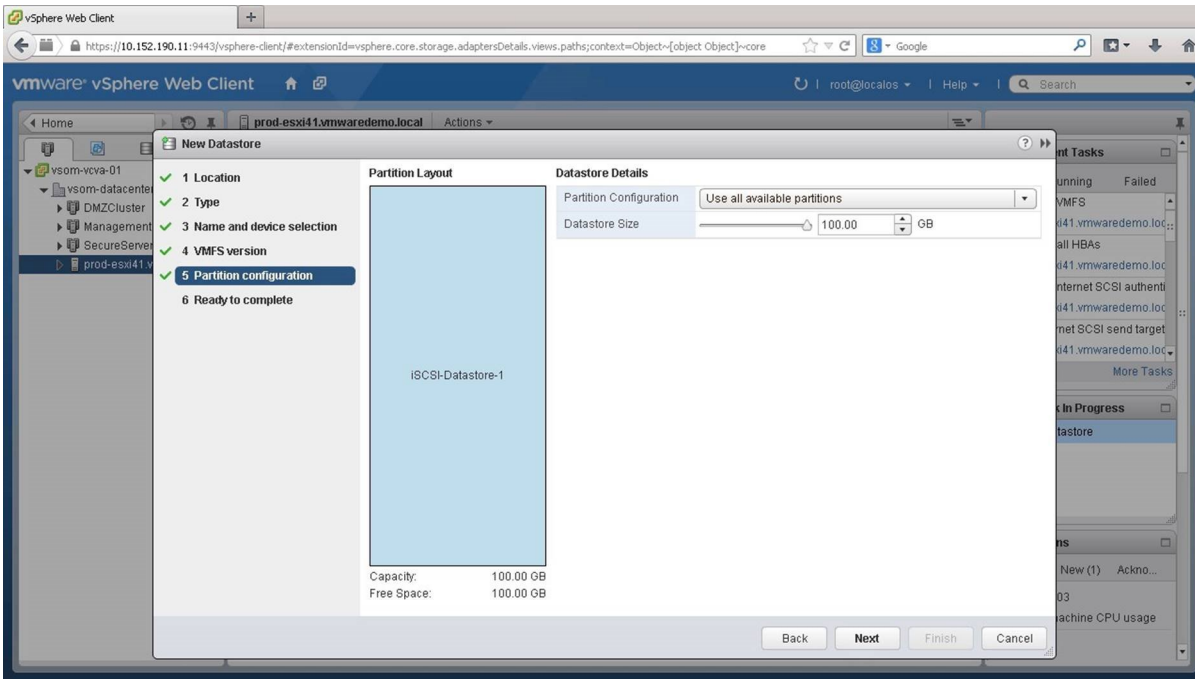
Select VMFS and click on [Next].




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



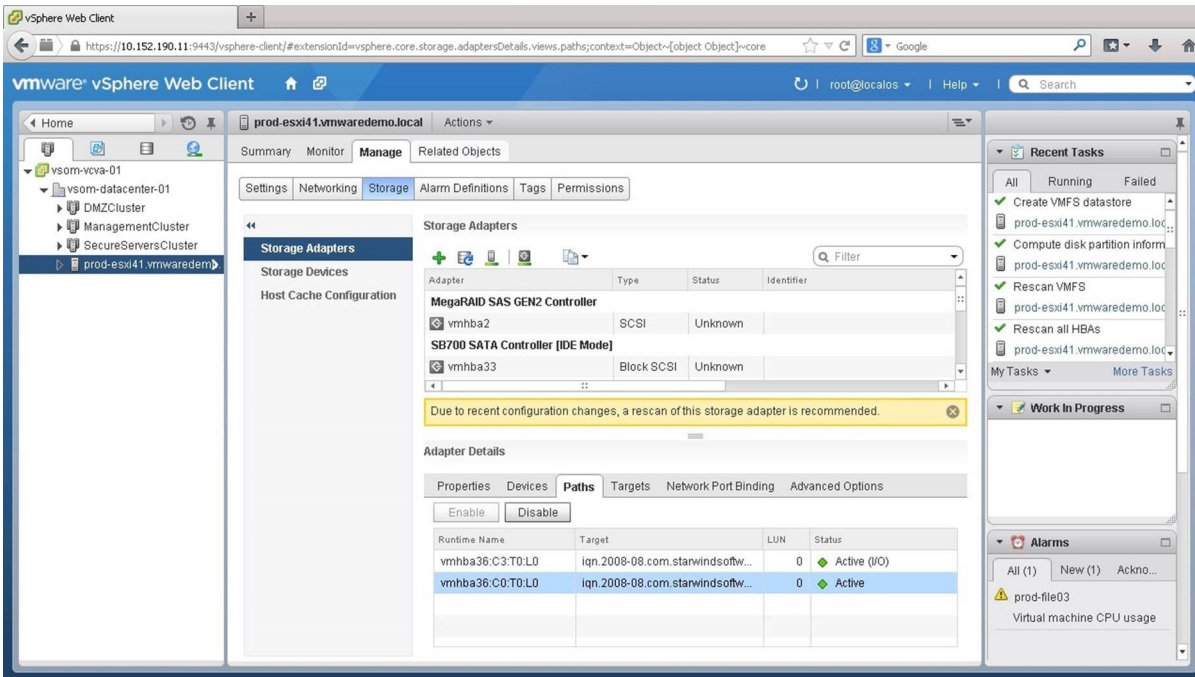
Select the VMFS version and click on [Next].



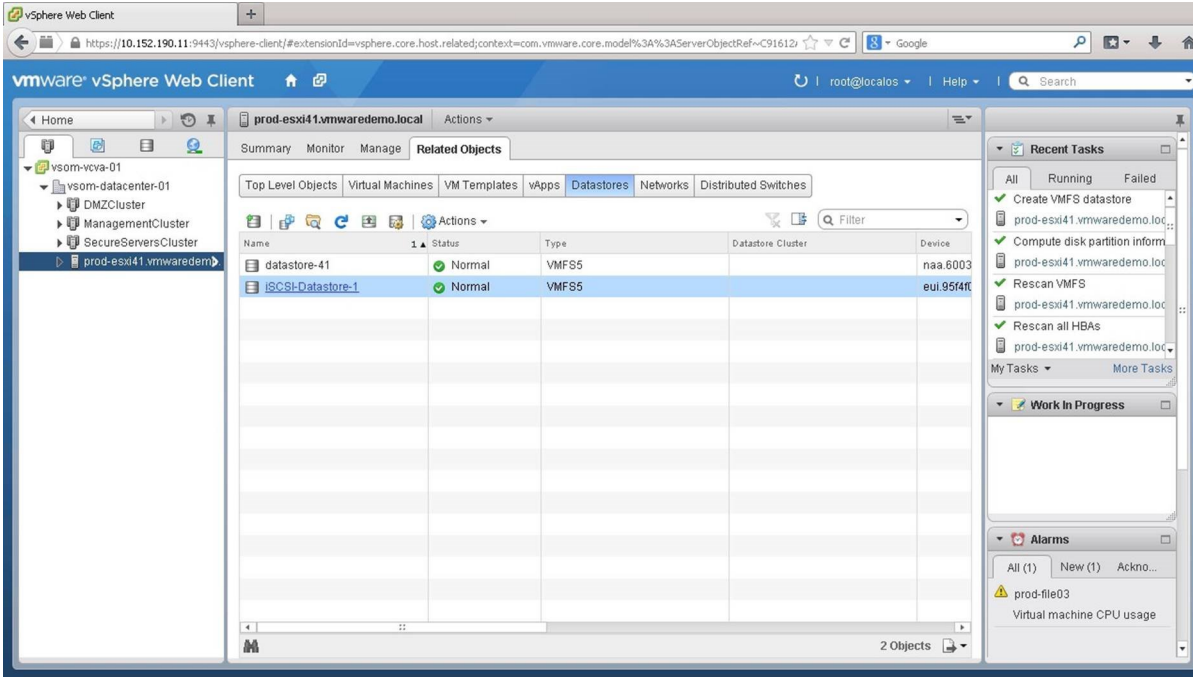
Choose the partition layout and click on [Next].



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.

## vSphere Core Storage

vSphere Core Storage on [core.vmware.com](https://core.vmware.com)

Visit [core.vmware.com](https://core.vmware.com) to learn about core storage features and capabilities in vSphere.

[vSphere 6.7 Storage Features](#)

## [vSphere 6.7 Storage](#)

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

[vSphere 6.7 Storage](#)



