Windows Server 2019 Scale Out File Server Cluster using PetaSAN

Version 1.0

Revision History

Date	Version	Description
15-7-2019	1.0	Initial version

Contents

1. Purpose

The purpose of this guide is to show how to create a Windows 2019 Scale Out File Server Cluster using PetaSAN for its underlying scale-out storage.

2. Pre-requisites

This guide assumes the reader has followed the Quick Start guide and has deployed a working PetaSAN cluster. We will be using the same subnet assignments as given in the Quick Start example.

Additionally this guide requires:

- 2 x Windows Server 2019 named sofs-1 and sofs-2 with 3 physical interfaces.
  These will act as our 2 SOFS servers
- 1 x Windows Server 2019 named AD with 1 physical interface
  This will act as our Active Directory server. We will also use it for central cluster management of the SOFS nodes.

3- Why PetaSAN and SOFS ?

SOFS has become the most popular solution for storing data for applications such as Hyper-V and SQL Servers. PetaSAN is ideal for use as the underlying storage for SOFS due to the following:

- PetaSAN provides scale-out storage at the disk block level, whereas SOFS provides scale-out at the file system and file share levels.
- A single PetaSAN disk can serve as the main storage volume for SOFS, making administration much simpler. In PetaSAN, a single disk is served by all storage nodes in parallel.
- PetaSAN storage is symmetric allowing all SOFS servers to run in Direct IO mode, concurrently writing to the same shares, for details on Direct IO, please refer to https://blogs.technet.microsoft.com/josebda/2013/10/30/automatic-smb-scale-out-rebalancing-in-windows-server-2012-r2/
- PetaSAN uses cloud based technology which supports storage over-commitment, it is possible to create a very large flat disk whose initial size exceeds physical storage available and then add physical storage as needed.

4. Network setup

The Windows servers used in this guide are configured with the following IP addresses

	AD	sofs-1	sofs-2
Management	10.0.1.50	10.0.1.51	10.0.1.52
Gateway	10.0.1.1	10.0.1.1	10.0.1.1
iSCSI 1		10.0.2.51	10.0.2.52
iSCSI 2		10.0.3.51	10.0.3.52

5. Active Directory Setup

5.1 Setup the AD Server

On the designated node for AD, add the role “Active Directory Domain Services” and reboot

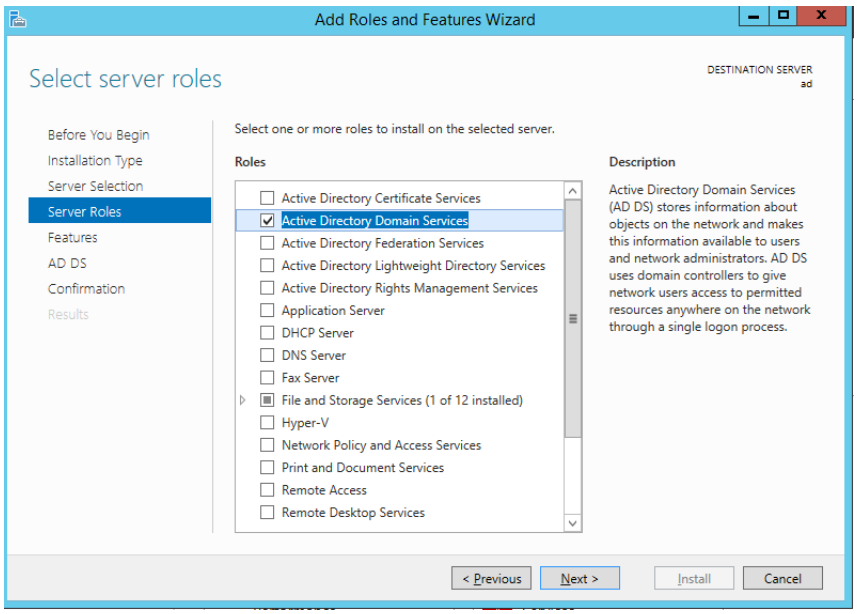

After reboot, select “Promote the server to a domain controller”

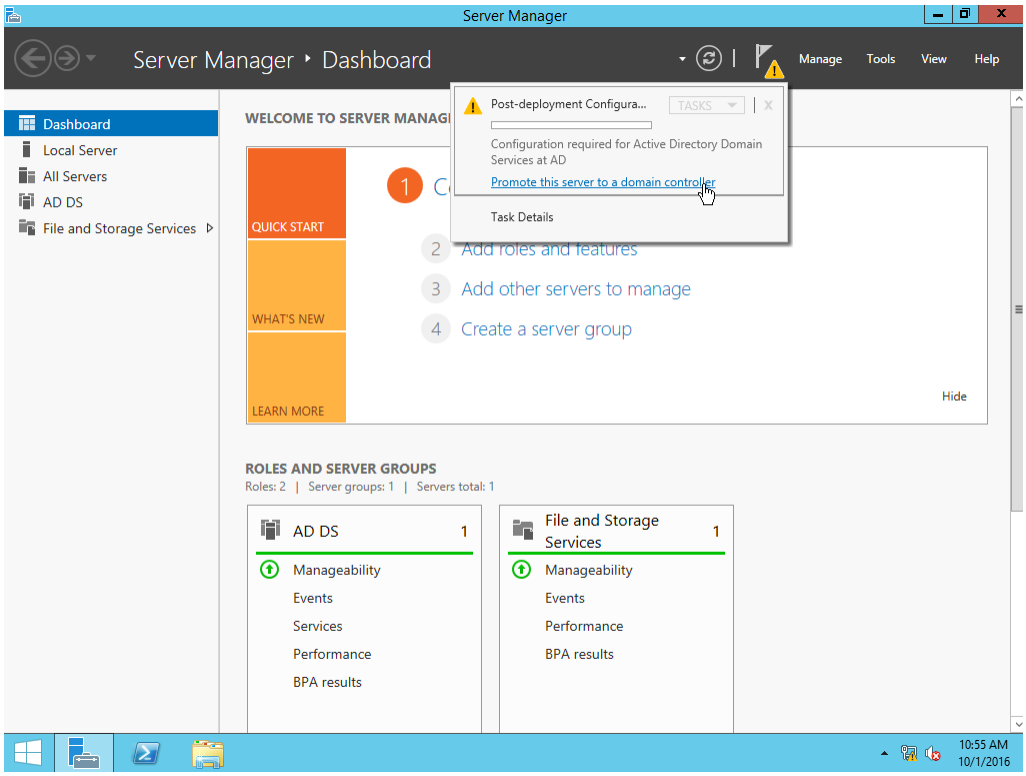

In the configuration wizard choose “Add a new forest” and enter the domain name “demo.local”

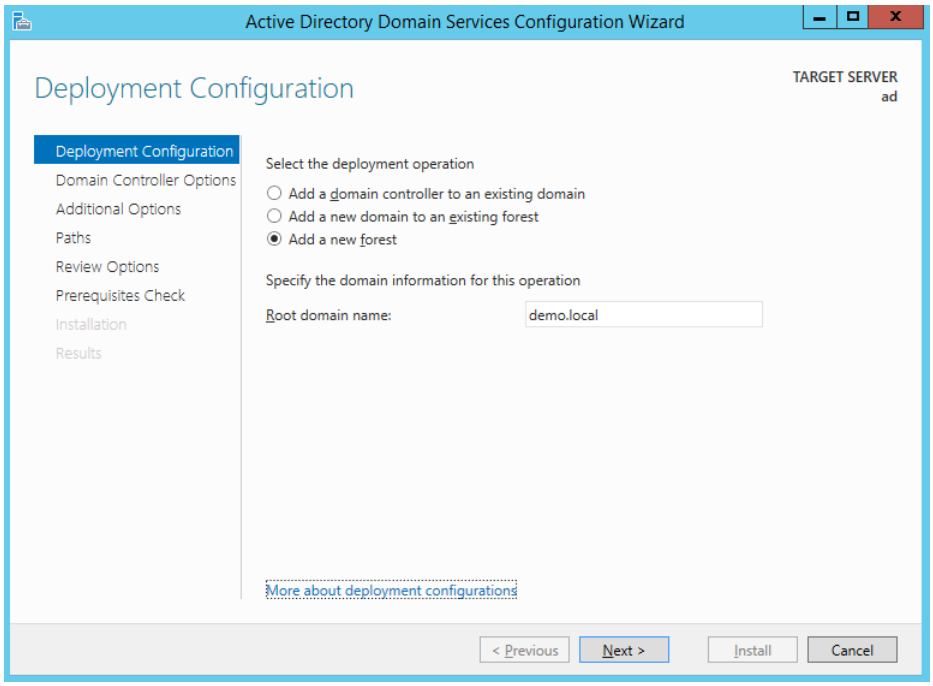

Enter the password

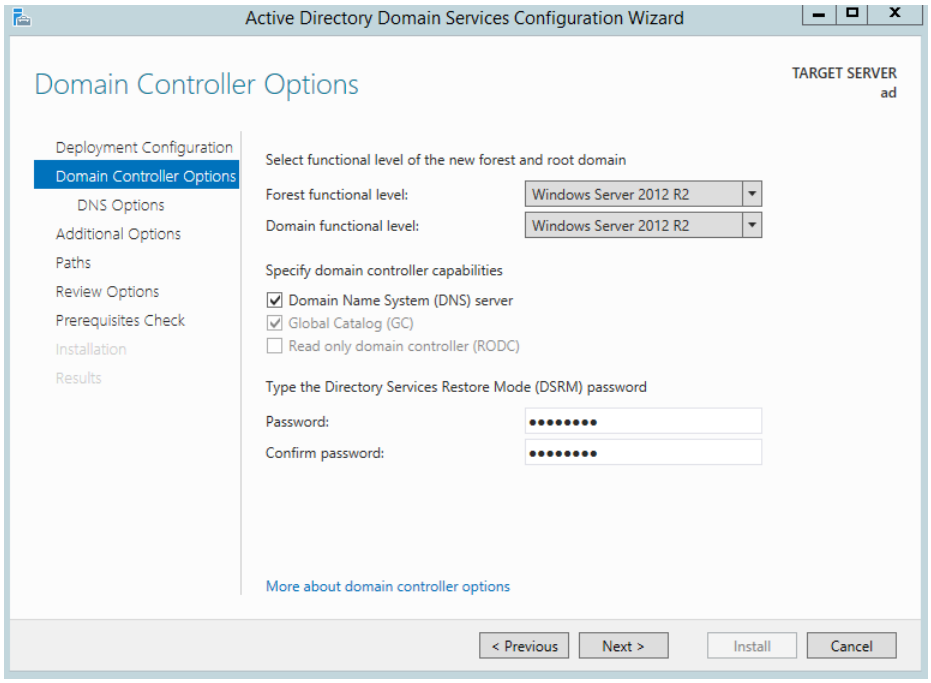

Reboot system when done.

5.2 Joining the AD Server

On both sofs-1 and sofs-2 nodes, edit the DNS setting to point to the AD server

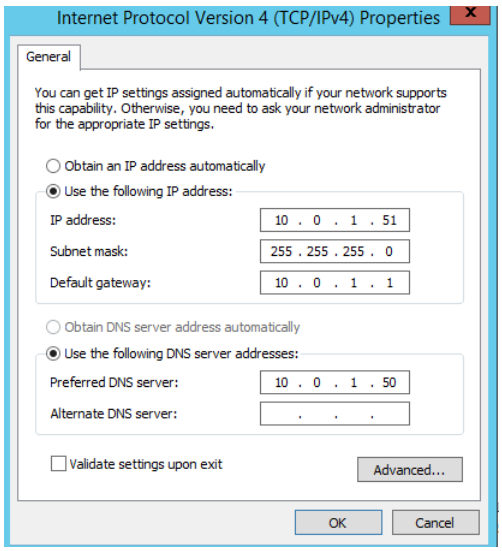

Then in “Server Manager -> Local Server” click on “WORKGROUP” in the “Domain” field.

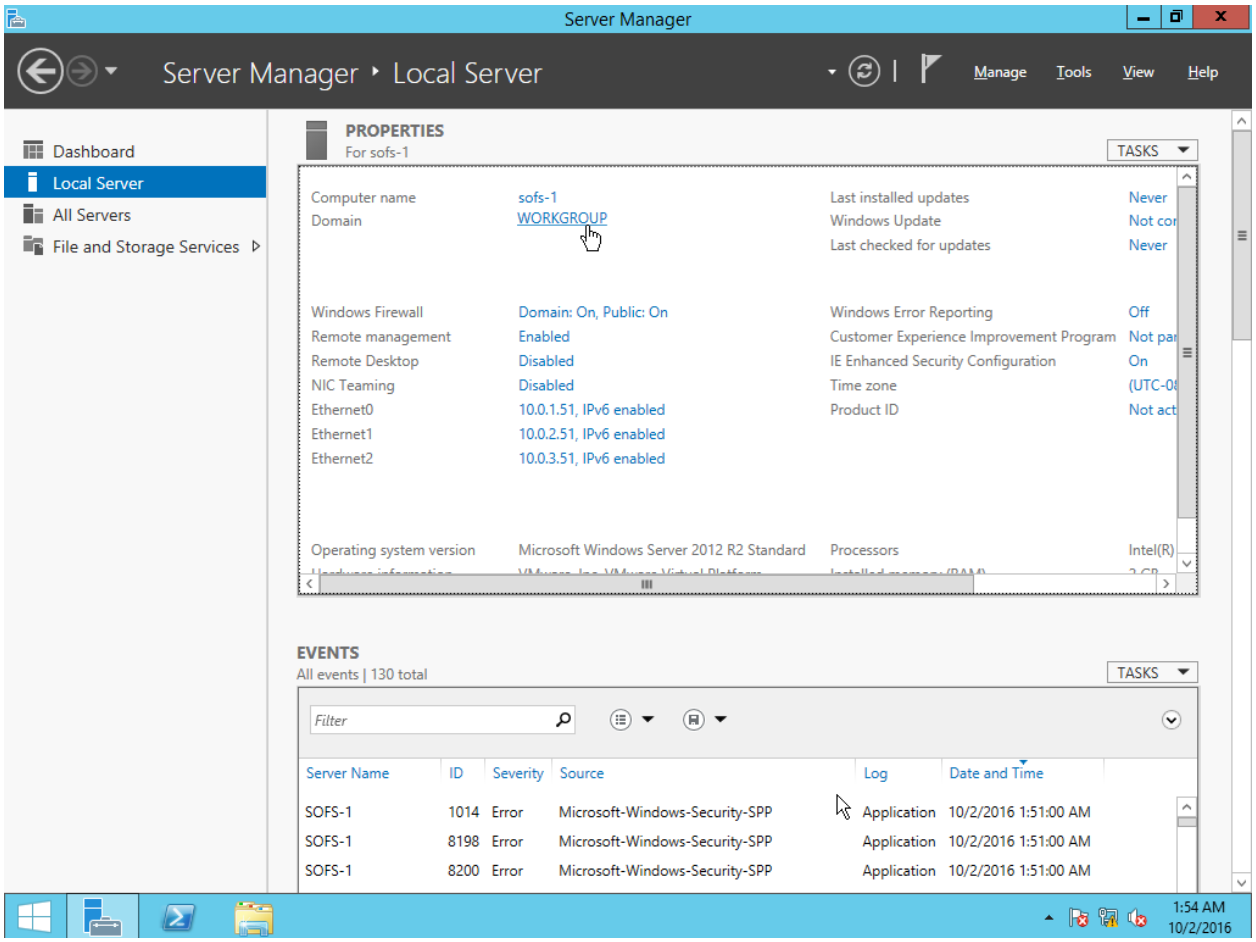

In “System Properties”, click “Change…”

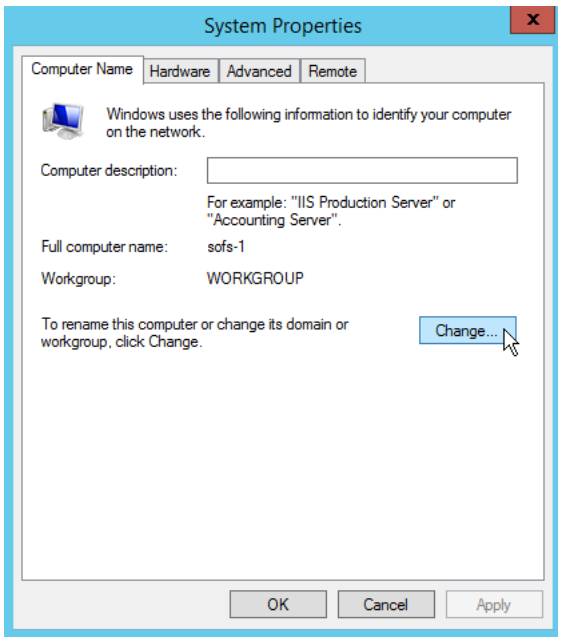

Enter “demo.local” in the domain field

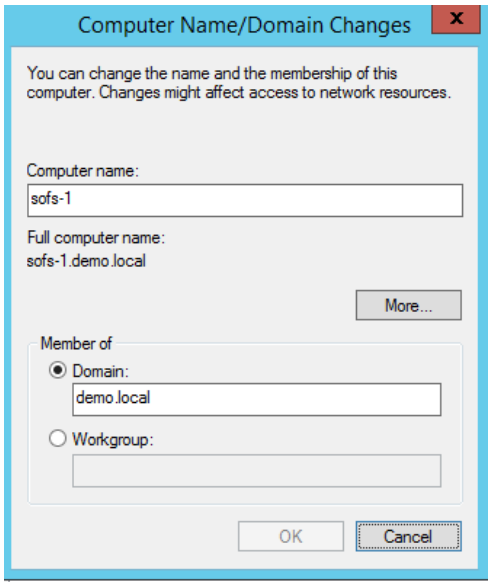

Enter the AD password

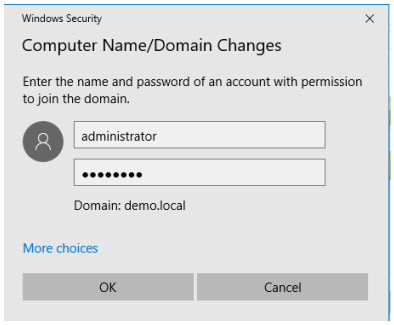

This should be all for joining the domain, please repeat the same steps for sofs-2.

6. Configuring node roles

On both sofs-1 and sofs-2 add the “File Server” Role.

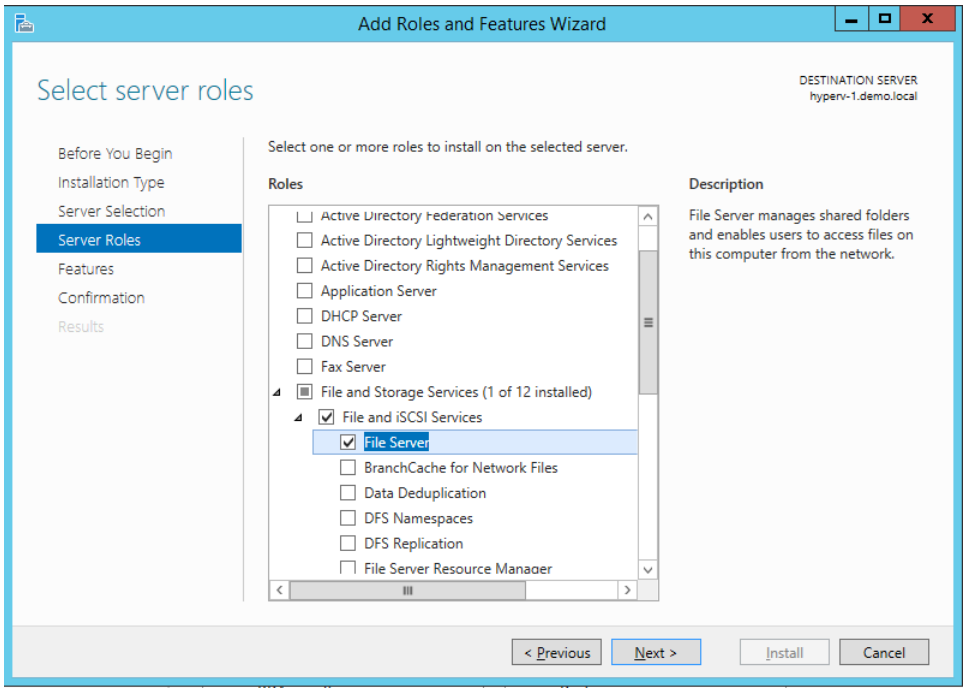

On the AD server as well on sofs-1 and sofs-2 add the “Failover Clustering” feature.

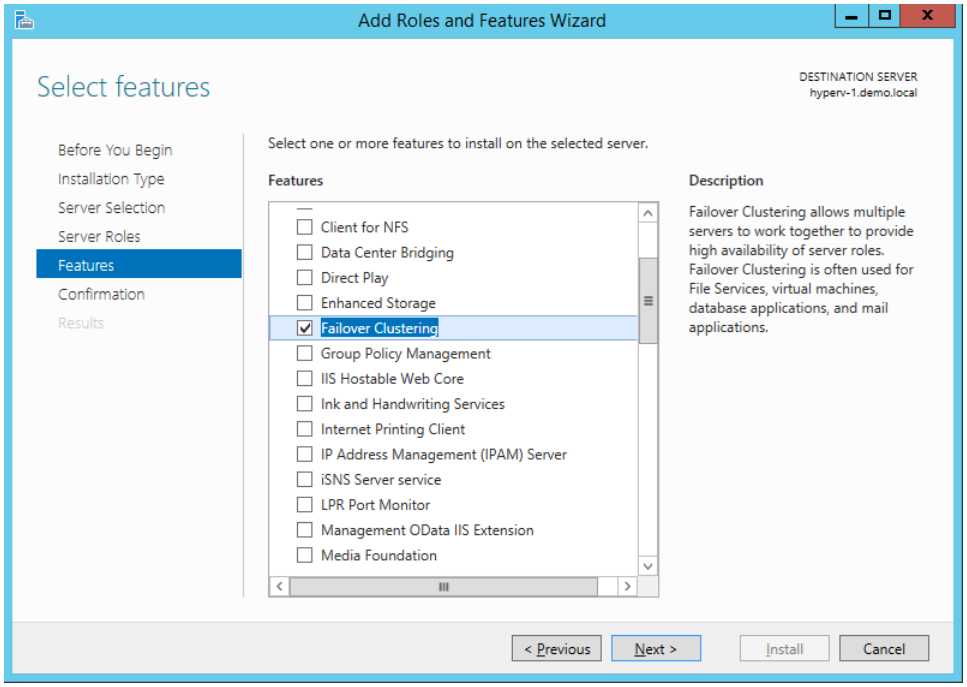

Reboot each node when done.

7. Connecting the PetaSAN disks

We need to create 2 disks in PetaSAN:

1. Disk1: 100TB x 4 paths with CHAP authentication. This will serve as the main data store for the SOFS shares.
2. Disk2: 1G x 4 paths with CHAP authentication. This will serve as a quorum disk, this is used by Windows Clustered Shared Volumes (CSV) to control concurrent access to the first disk from multiple machines.

We need to connect to these 2 disks from both sofs-1 and sofs-2 servers. Please refer to *Connecting to PetaSAN from Windows 2019 using MPIO* guide for step by step instructions.

Note that initializing and formatting the disks should be done once from one node only, for example from sofs-1.

8. Validating the cluster

As discussed earlier, we selected to use our AD server as the machine we use for cluster management. Before we create our cluster, we should let Windows validate it first by running a couple of tests.

On the AD server open the “Failover Cluster Manager”

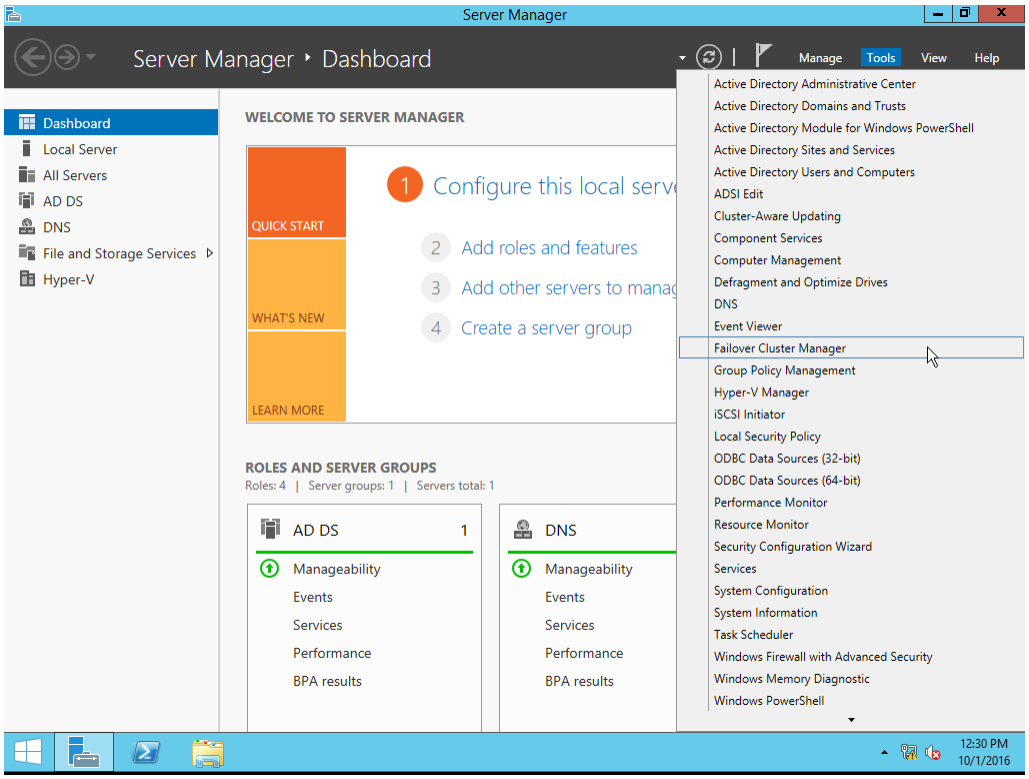

Next click on “Validate Configuration”

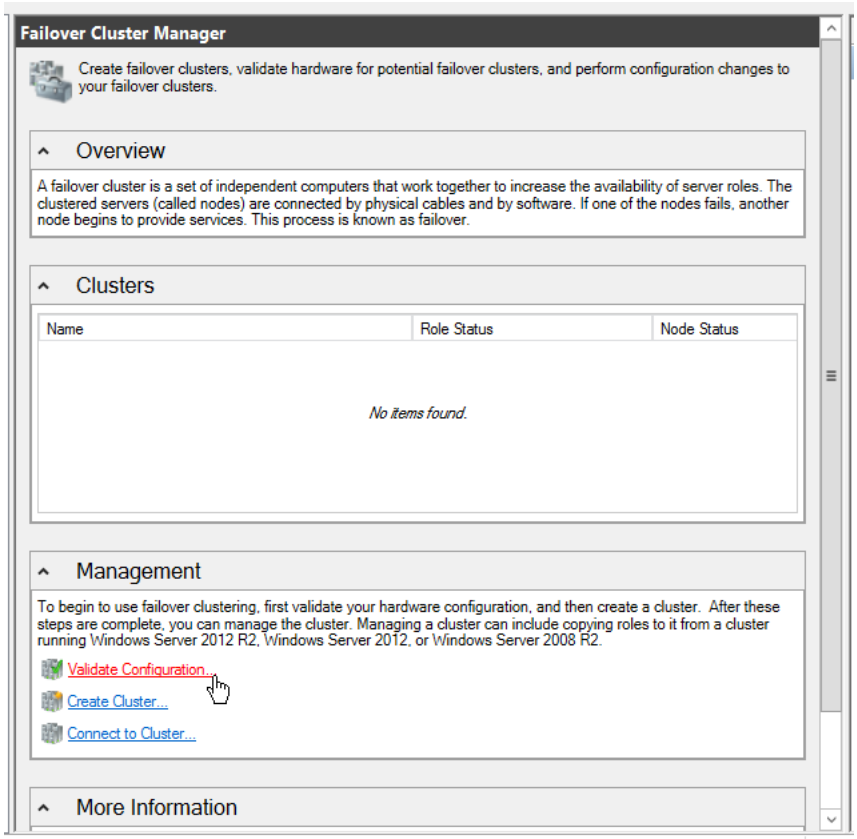

In the validation wizard, add both sofs-1 and sofs-2

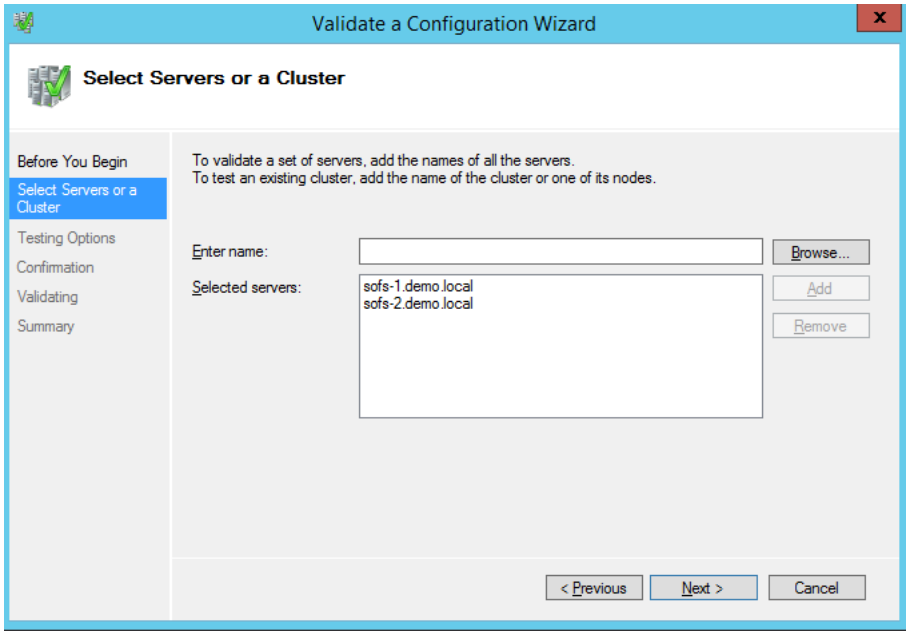

Next choose to run all tests

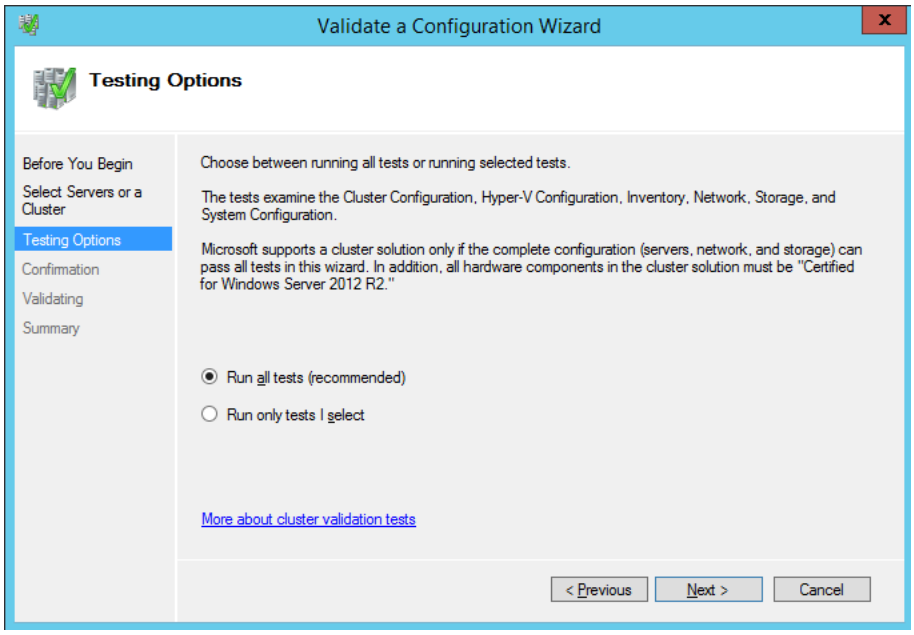

The wizard will take a couple of minutes to run various tests; these include many important storage failover and validation tests for our PetaSAN disks.

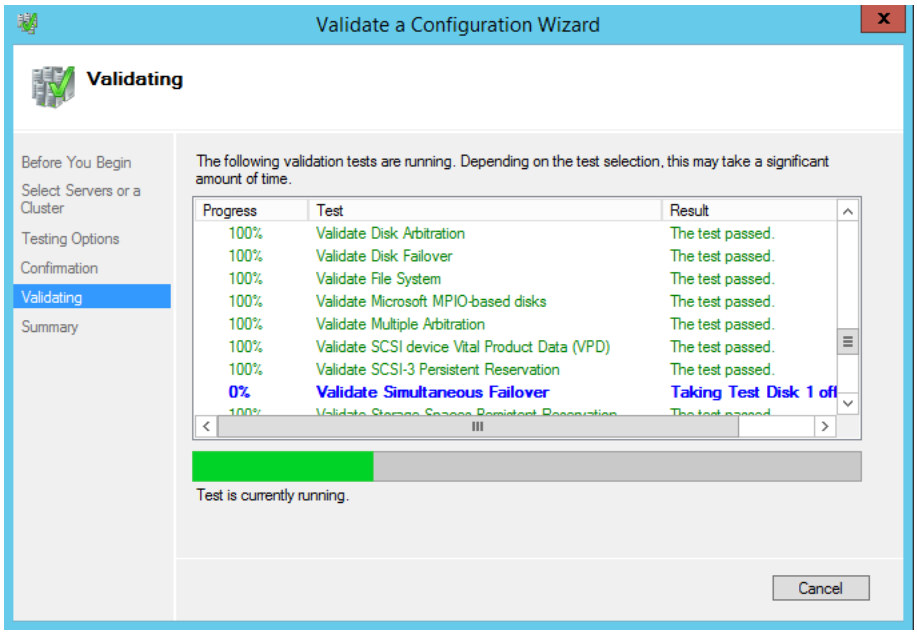

Once completed, the wizard displays a cluster validation report.

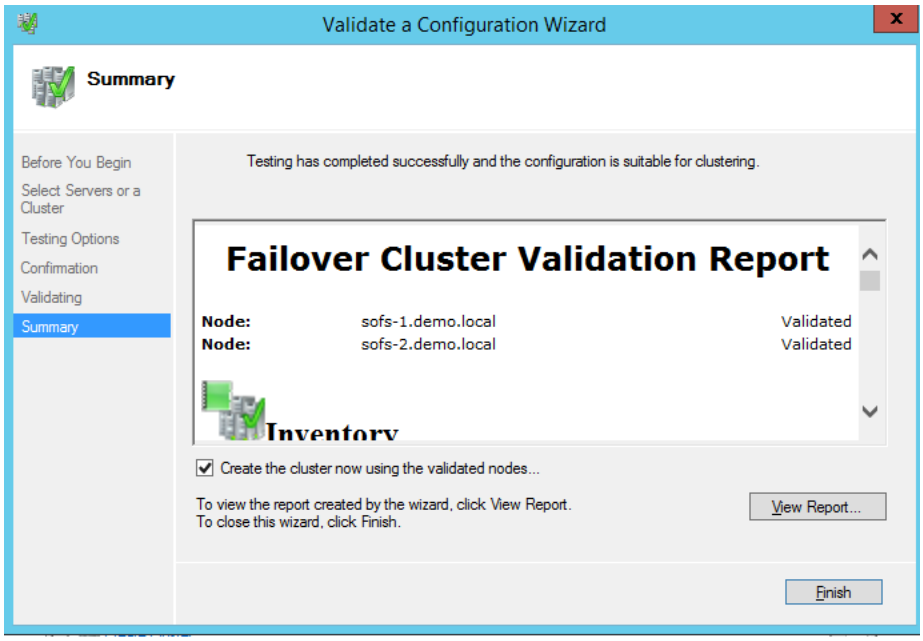

If all is good, leave the “Create the cluster now using the validated nodes” checked and click on the “Finish” button.

9. Cluster Creation

After successful validation, the create cluster wizard is displayed

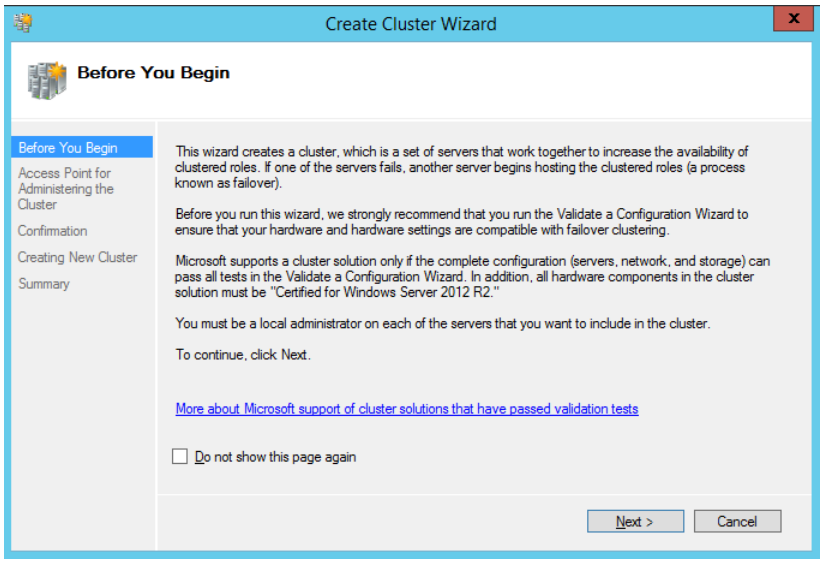

Enter the cluster name, for example “SOFS Cluster”

We need to assign an IP address for the cluster and specify which network interface it will use.

Windows will display a list of candidate interfaces to choose from, it will not list interfaces it thinks are not appropriate, such as those used for iSCSI storage.

In our example we chose IP 10.0.1.100 for our cluster. When clients connect to our SOFS, they connect using this IP.

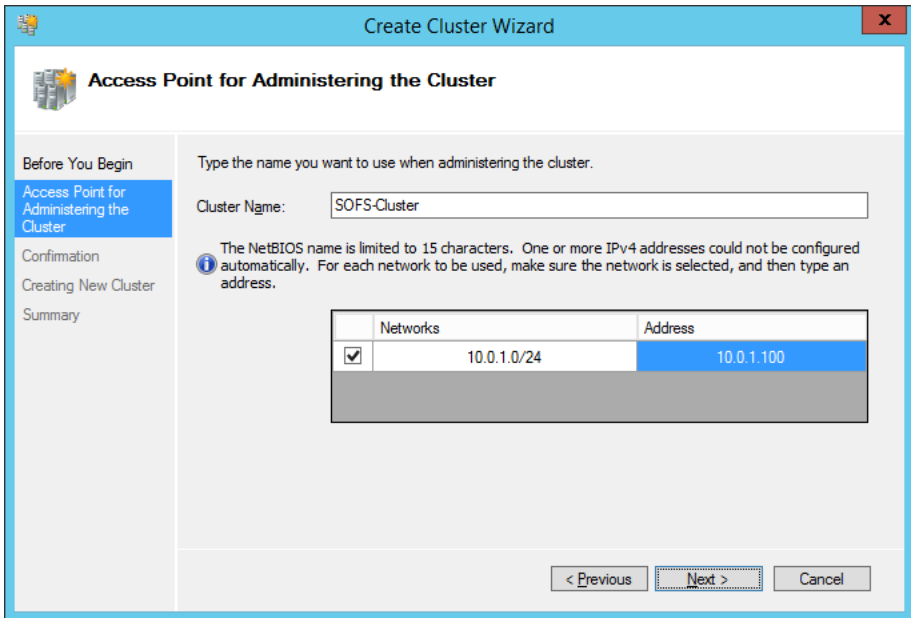

- *Note: In our simple setup, the SOFS cluster IO resides on the same network as our Management traffic. For setups requiring better isolation, we could have added a fourth interface card and created a subnet specifically for SOFS IO traffic.*

Click “Next”

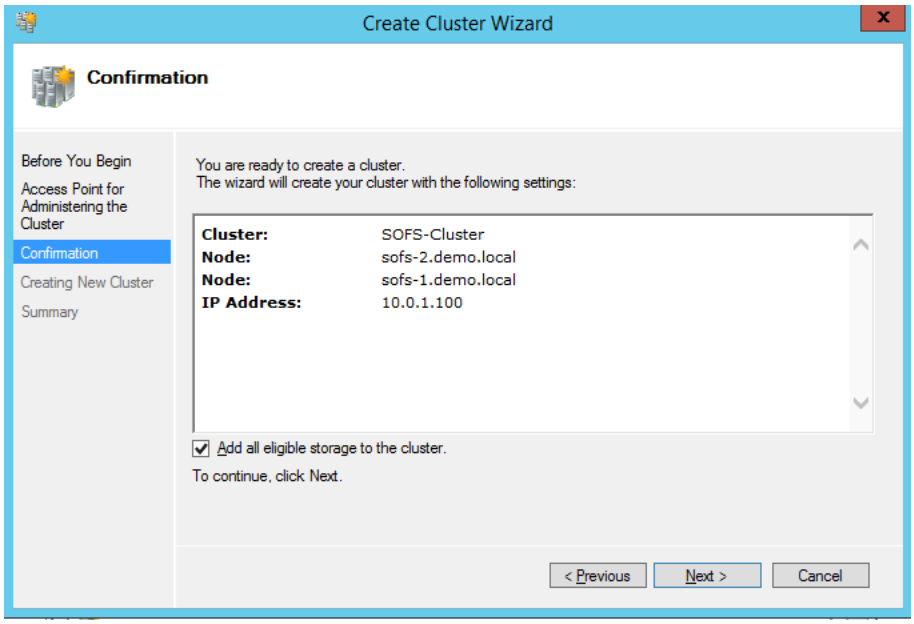

Click “Next”, then “Finish”

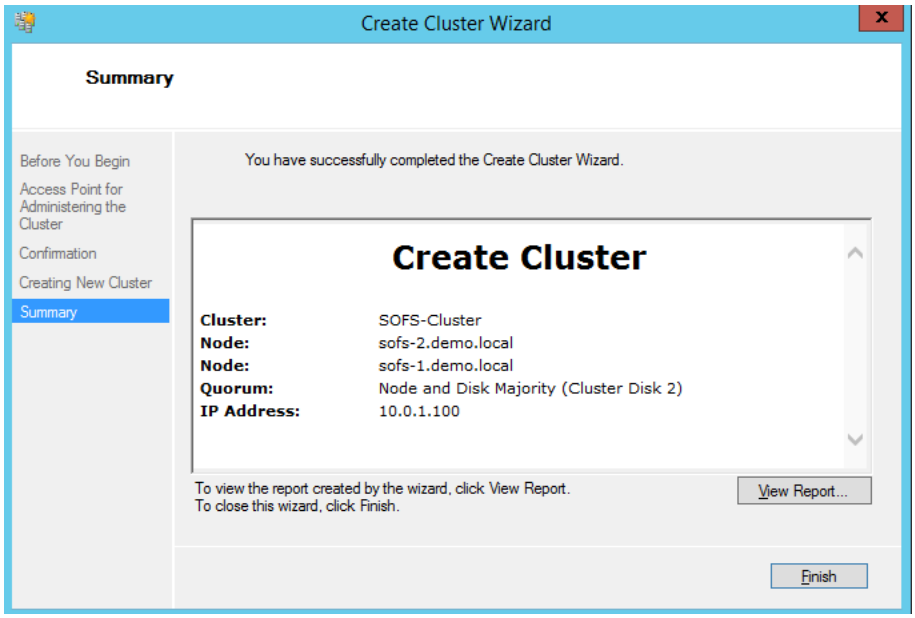

Once the cluster is created, go to Storage -> Disks

Right click on the 100 TB disk add select “Add to Cluster Shared Volumes”

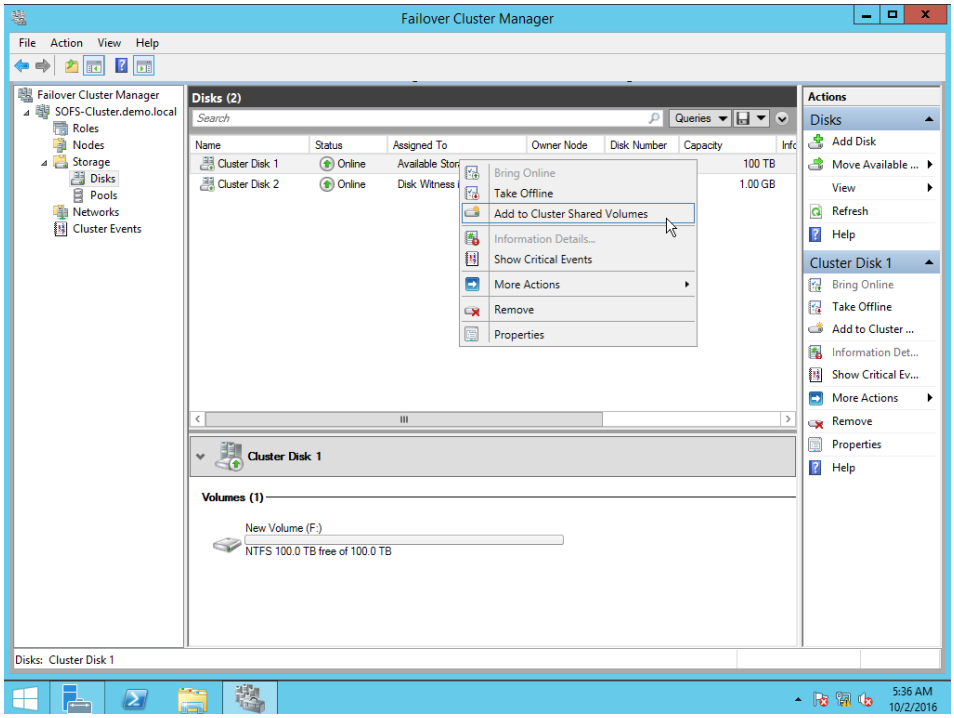

On the bottom pane, the 100 TB disk volume will change from NTFS to CSVFS (Clustered Shared Volume File System), this allows the volume to be used by many SOFS nodes concurrently. Notice too that it is now accessible as “C:\ClusterStorage\Volume1”

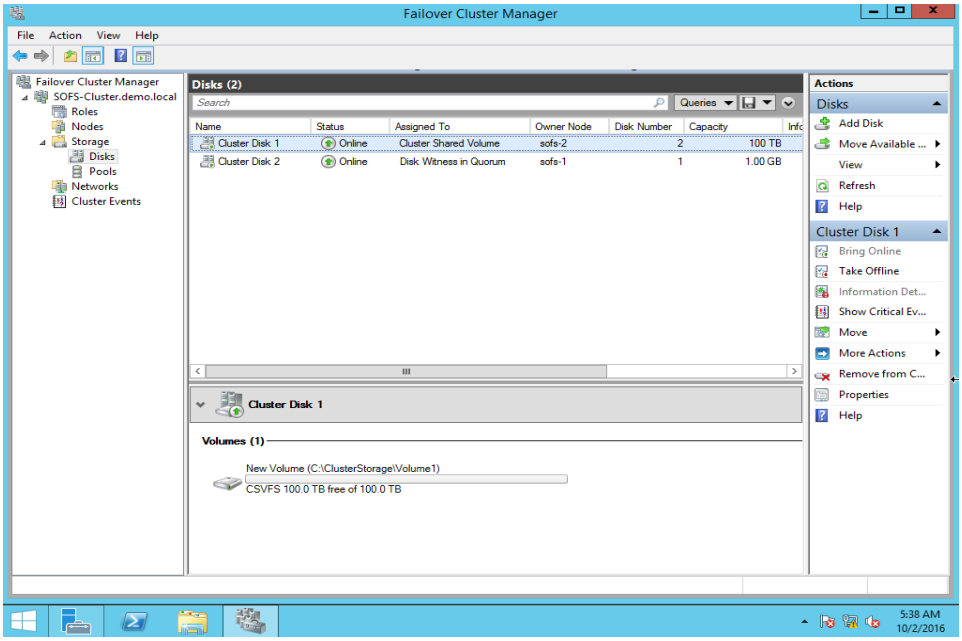

Right click on “Roles” and select “Configure Role…”

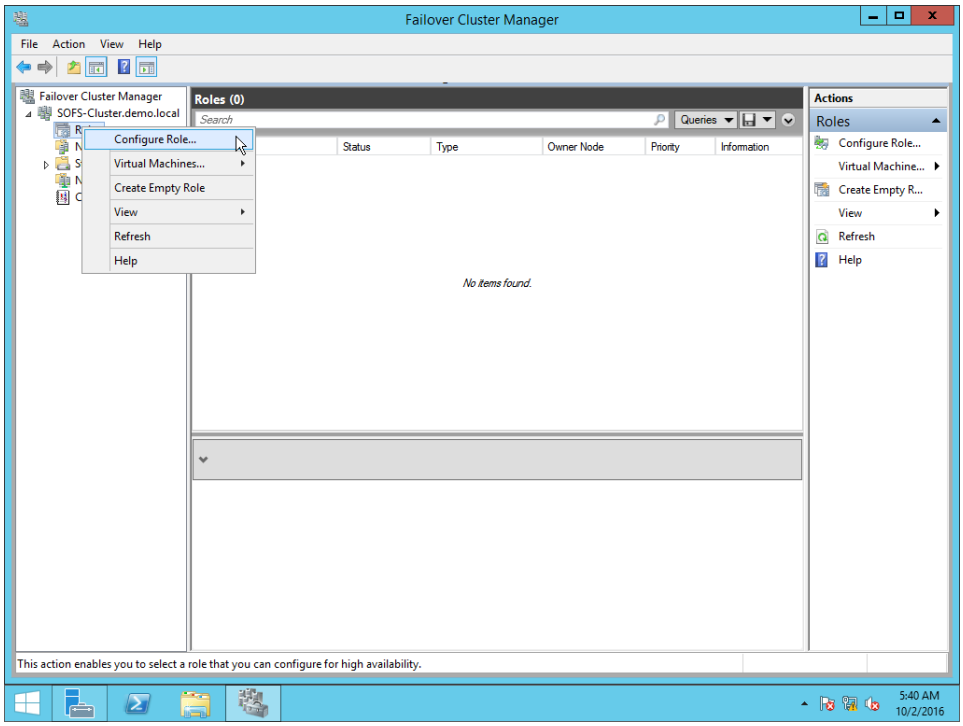

In the “Select Role” choose “File Server”

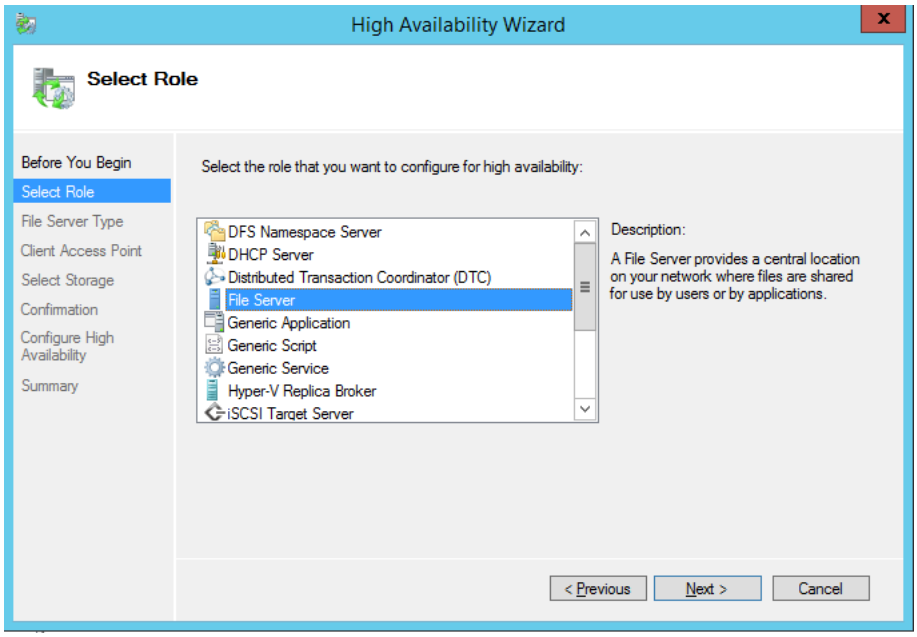

Select “Scale Out File Server for application data”

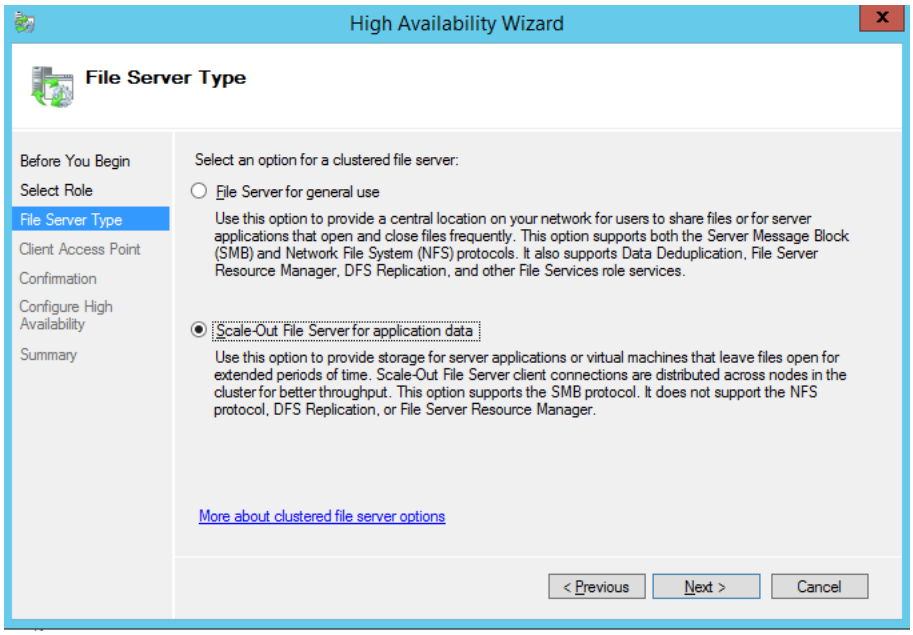

We then enter the cluster network (NetBIOS) name; in our example it is “SOFS”

Clients will access shares on the cluster using this name

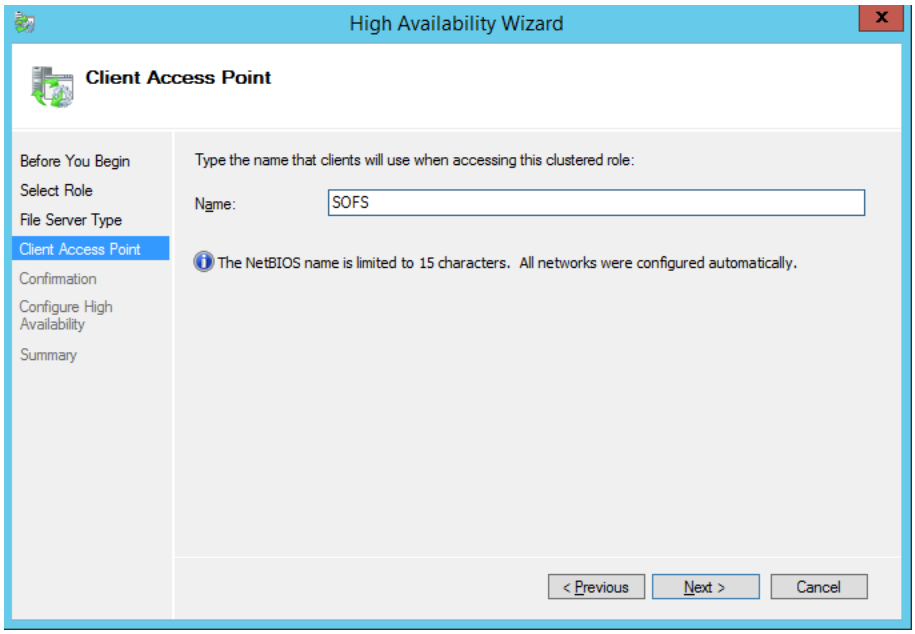

Click “Finish” to complete.

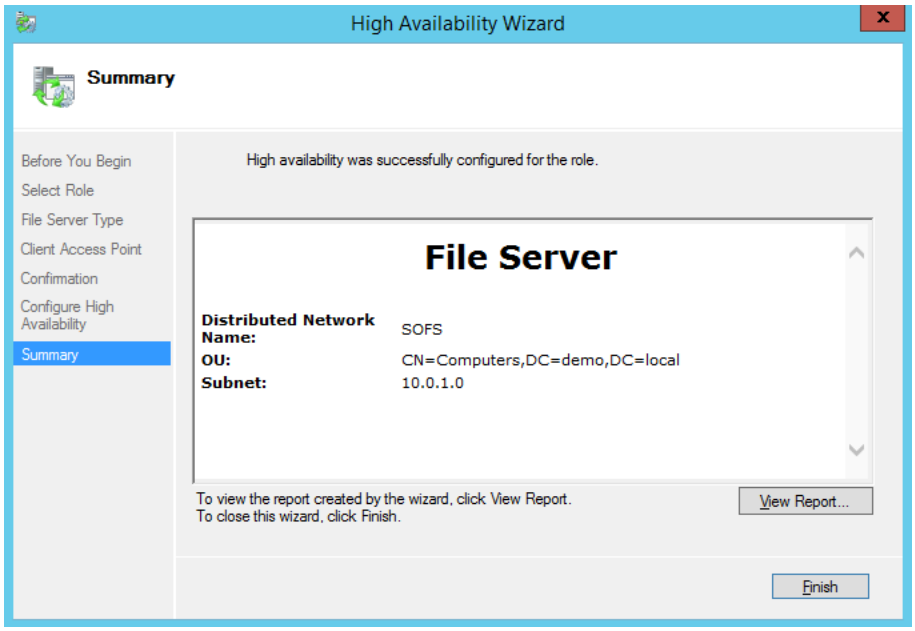

We have successfully created our SOFS cluster, our next step is to add shares to it.

10. Adding Shares

To add a share, right click on the “SOFS” cluster role we added and click “Add File Share”

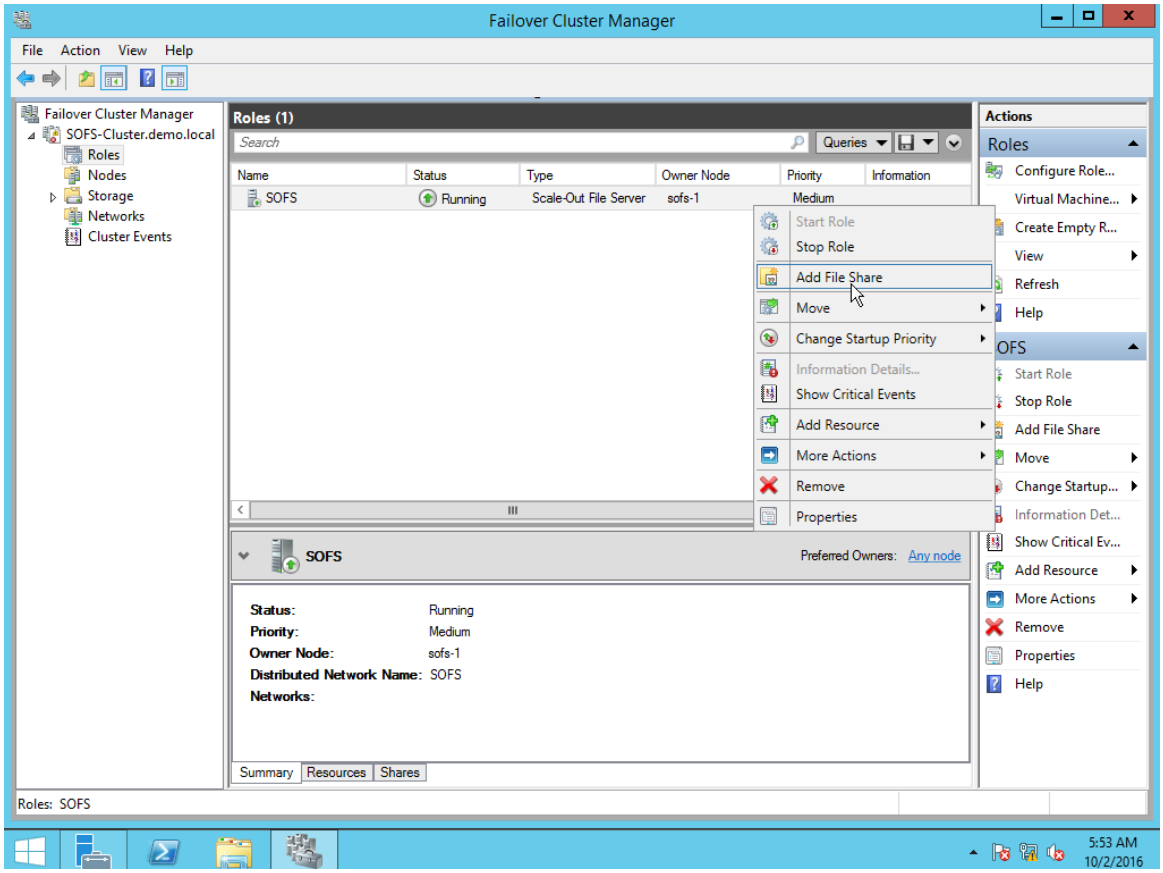

In the wizard select “SMB Share – Applications” profile.

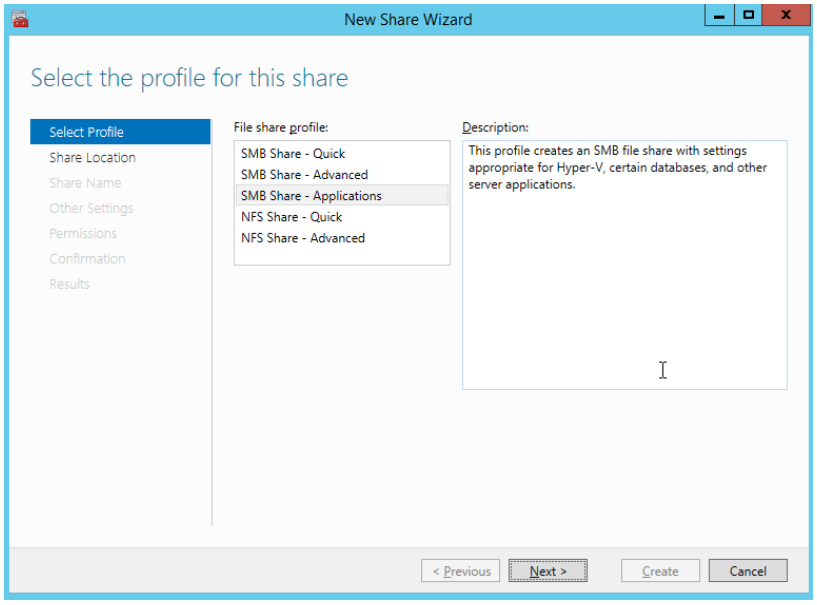

In the “Select the server and path for this share” accept the default values.

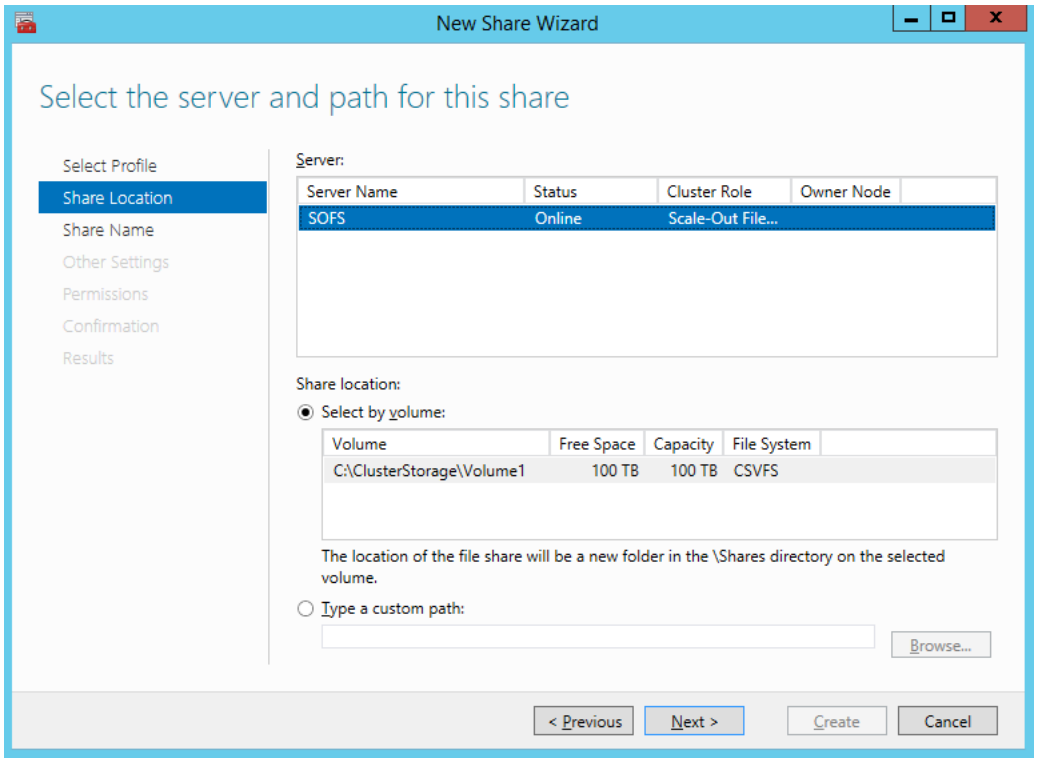

Now enter the name and description for the new share, for example “vm-disks”

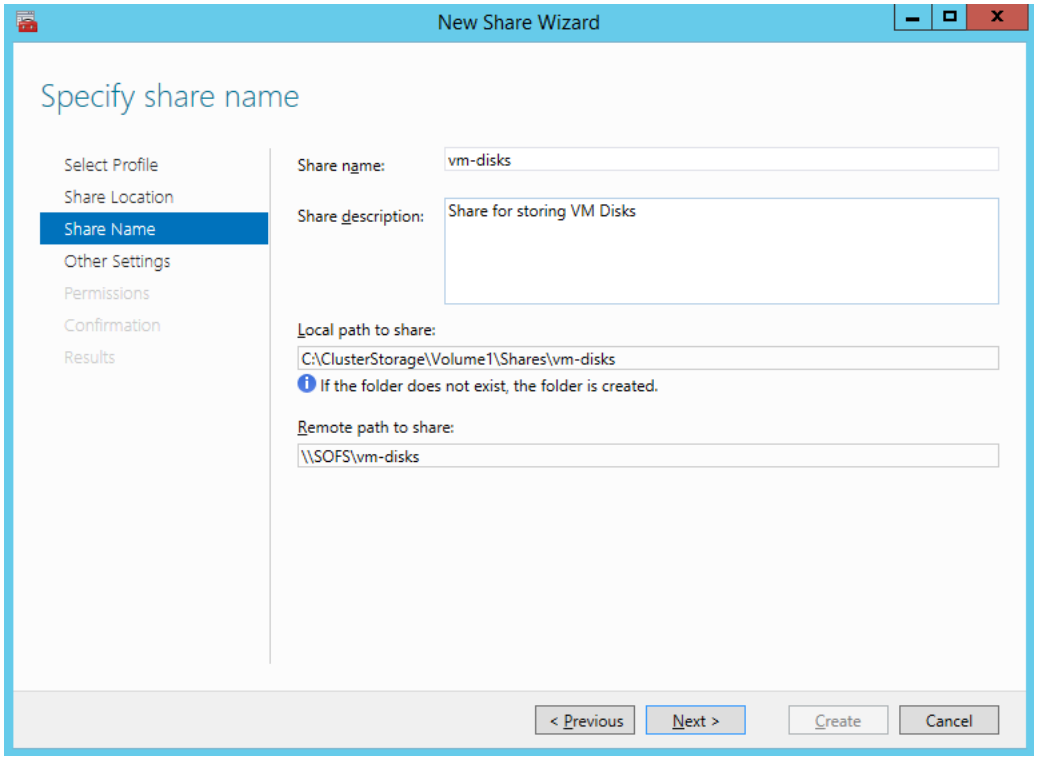

For the share settings, accept the defaults and click “Next”

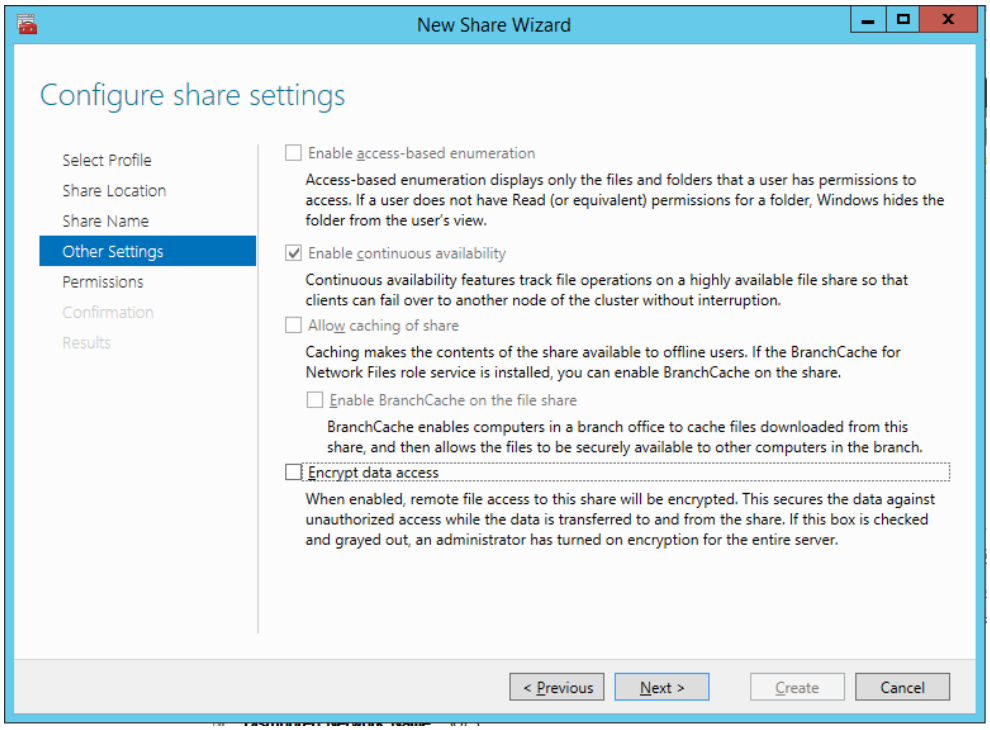

For Permissions, accept the defaults and click “Next”

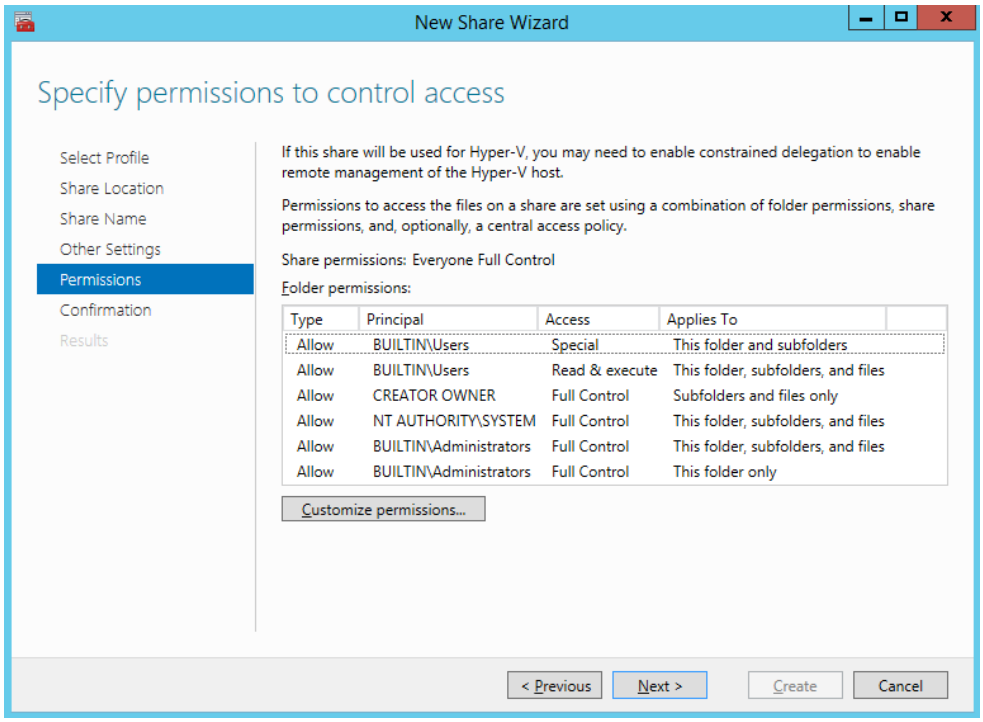

Once the new share wizard is done, we should see our new share added on the lower section of the window as per below:

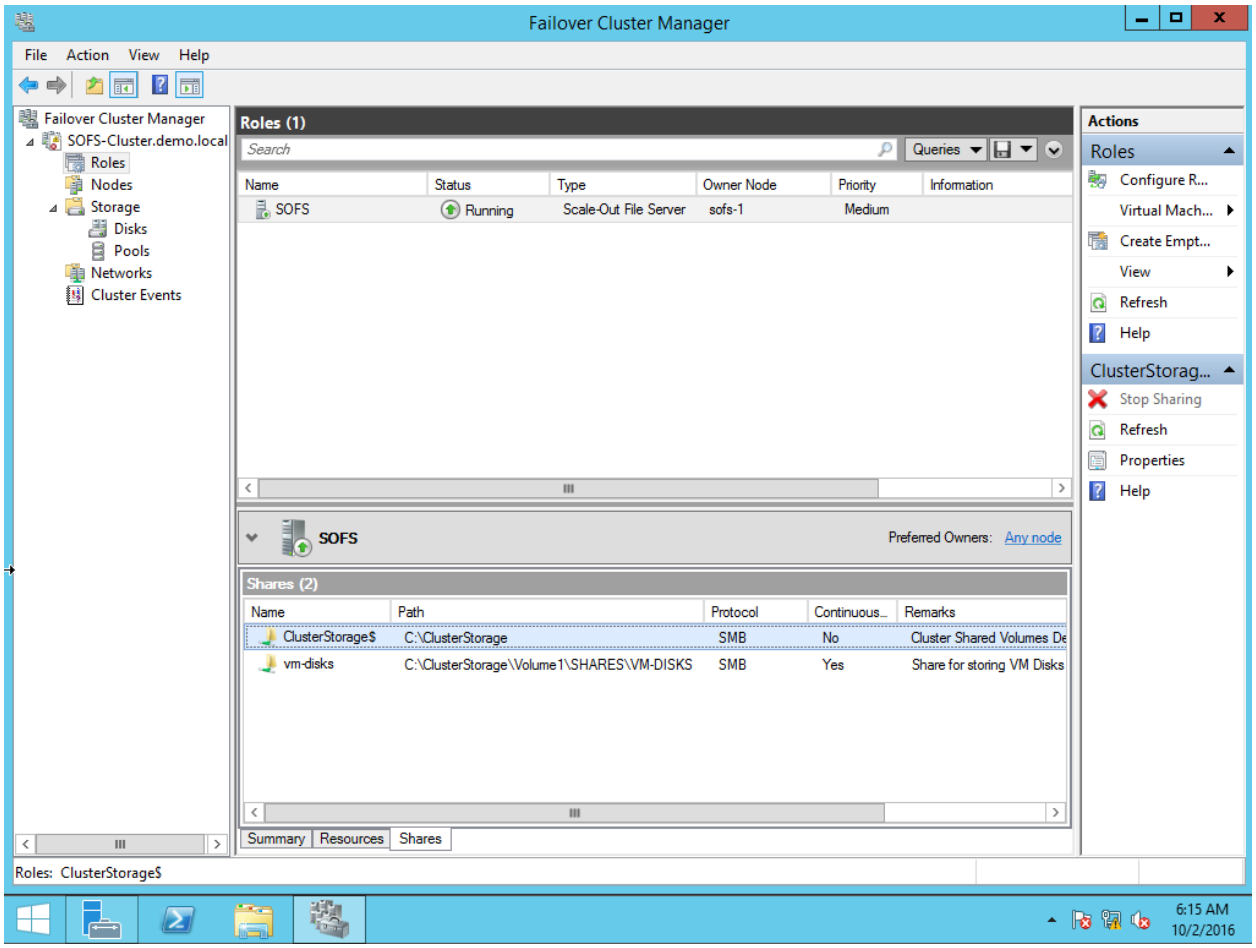

That is it, our share is up and running. In our example setup, clients access the share using the name \\SOFS\vm-disks

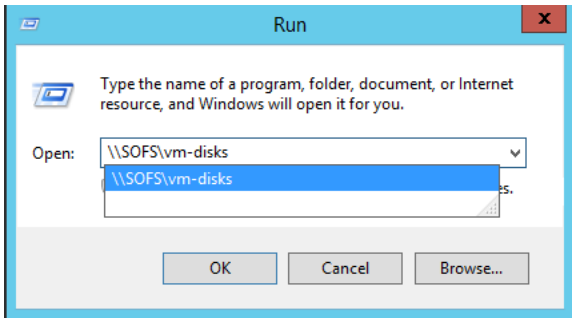

Finally we can add more nodes to our SOFS cluster to increase throughput. All SOFS nodes will have simultaneous concurrent IO on the share.