Windows Server 2019 Hyper-V Cluster using PetaSAN

Version 1.0

Revision History

Date	Version	Description
15-07-2019	1.0	Initial version

Contents

1. Purpose

The purpose of this guide is to show how to create a Windows 2019 Hyper-V Cluster using Clustered Shared Volumes (CSV) stored on PetaSAN. In this setup, virtual machines are stored directly by Hyper-V on the CSV volumes. Using shared storage provides high availability for the virtual machines and support s more advanced features such as live migration.

This guide does not cover storing the virtual machines on a Scale Out File Server, which will be covered in a different document.

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 hyperv-1 and hyperv-2 with 4 physical interfaces.
  These will act as our 2 Hyper-V 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 Hyper-V nodes.

3. Network setup

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

	AD	hyperv-1	hyperv-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
VM switch		10.0.6.51	10.0.6.52

- *Note: Subnets 10.0.4.0 & 10.0.5.0 are assigned to the PetaSAN backend networks.*

4. Active Directory Setup

4.1 Setup the AD Server

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

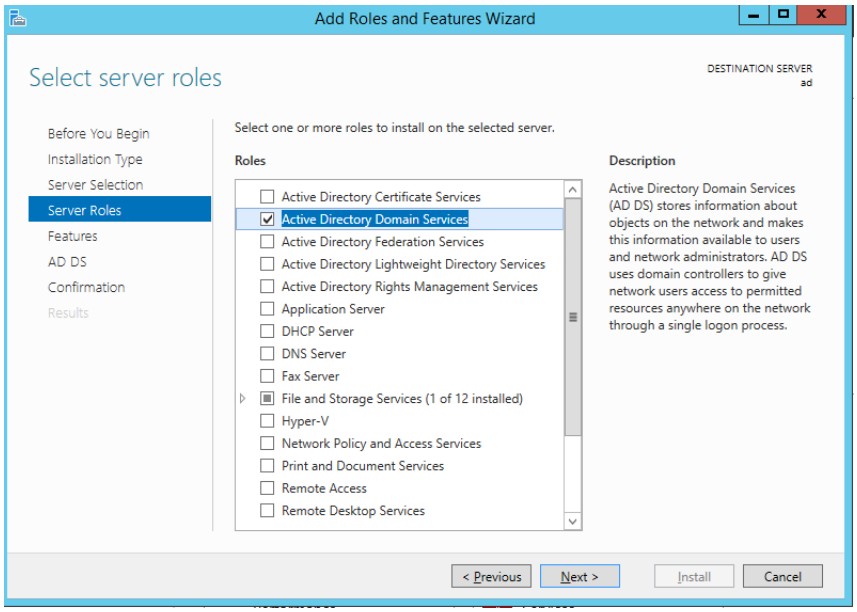

After reboot, select “Promote the server to a domain controller”

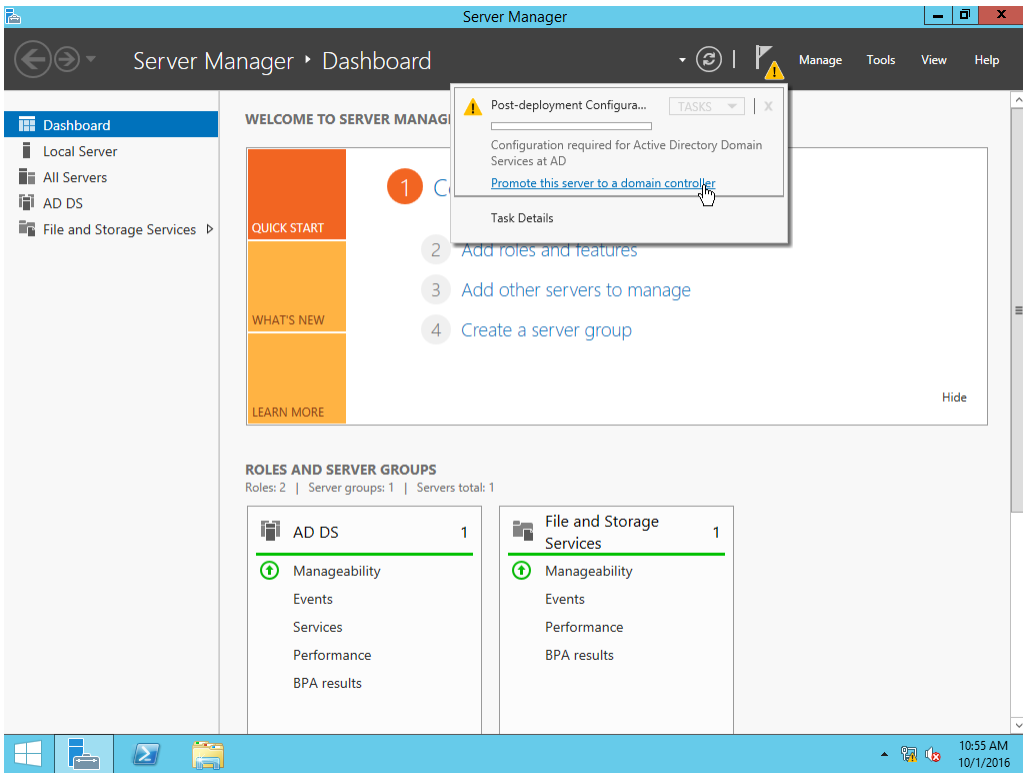

In the configuration wizard choose “Add a new forest” and enter the domain name “demo.local”

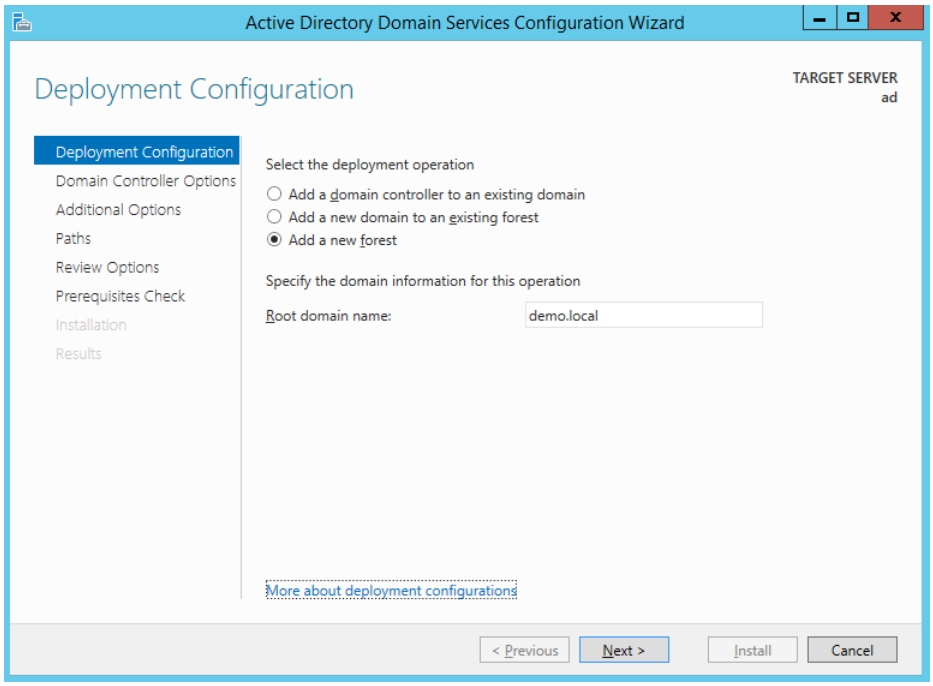

Enter the password

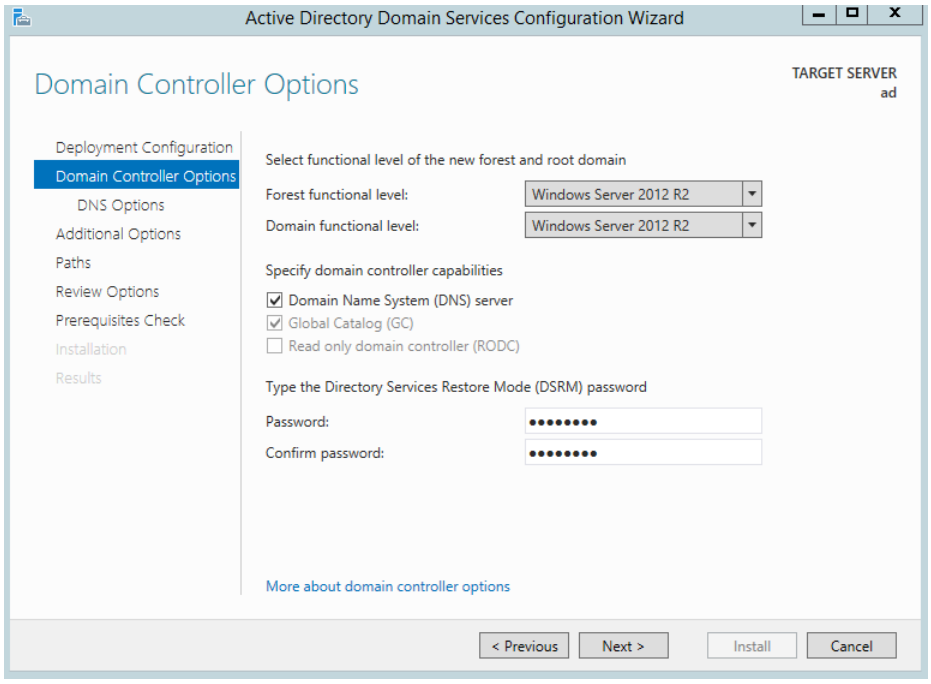

Reboot system when done.

4.2 Joining the AD Server

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

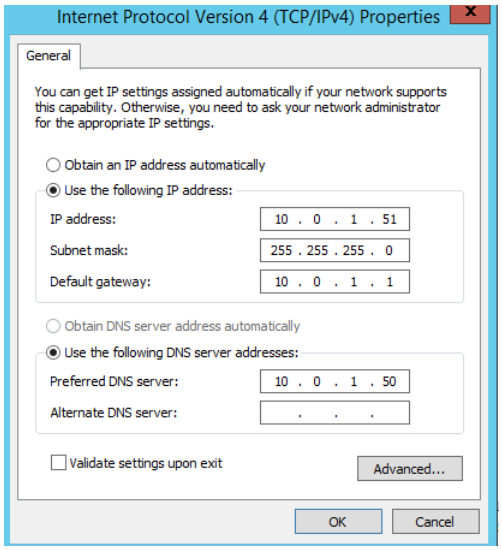

Then in “Server Manager -> Local Server” click on “WORKGROUP” in the “Domain” field.

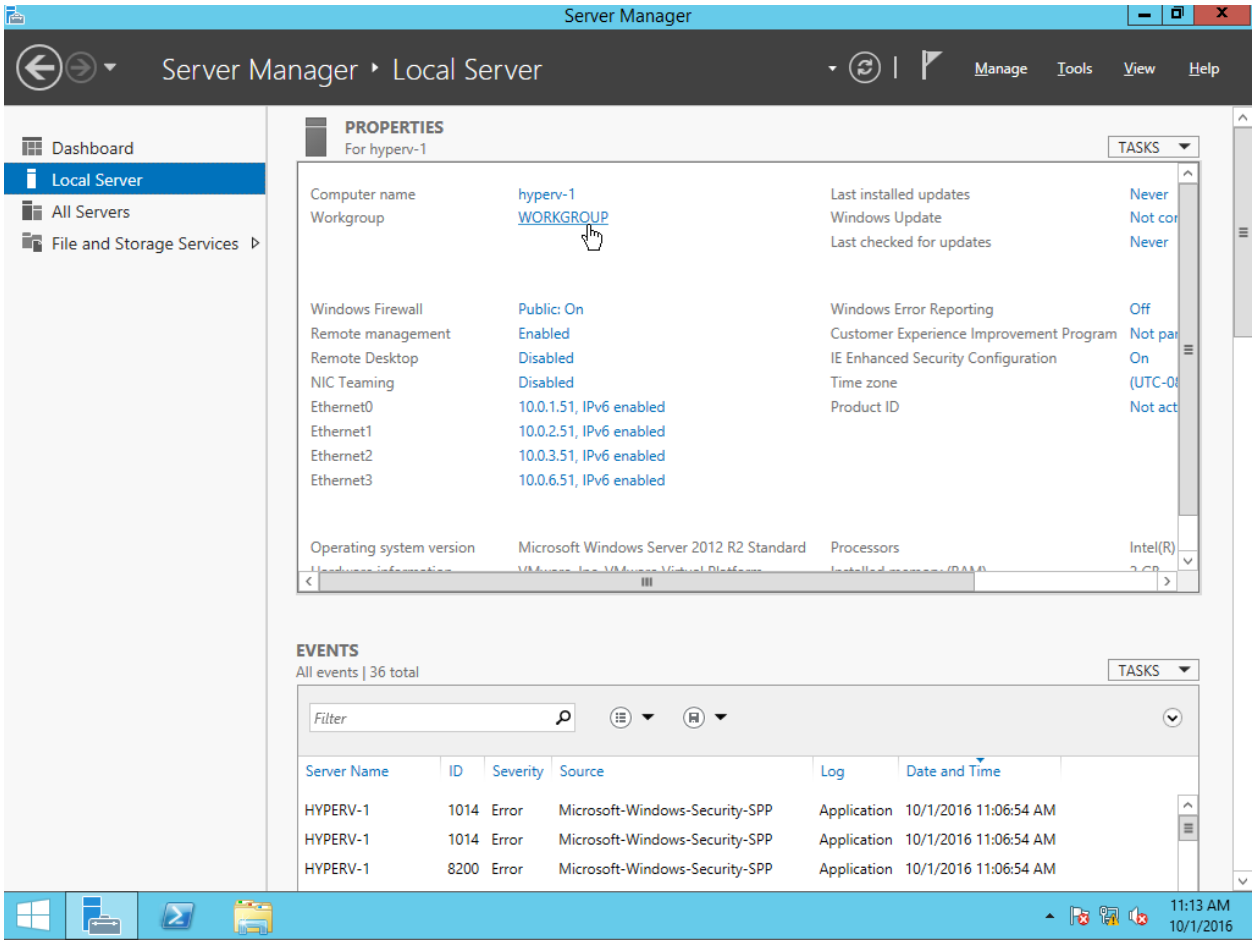

In “System Properties”, click “Change…”

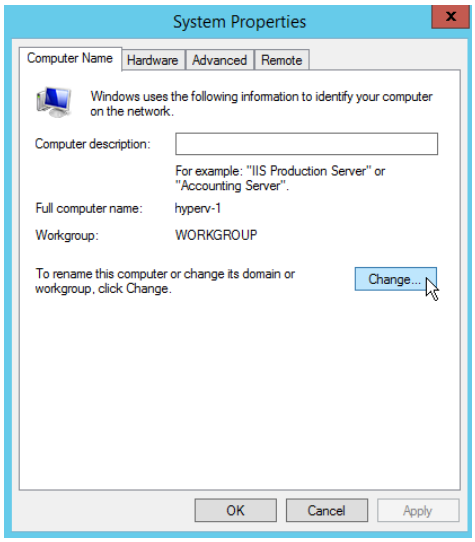

Enter “demo.local” in the domain field

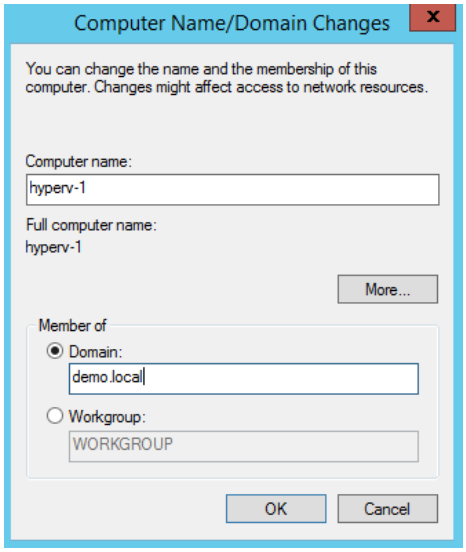

Enter the AD password

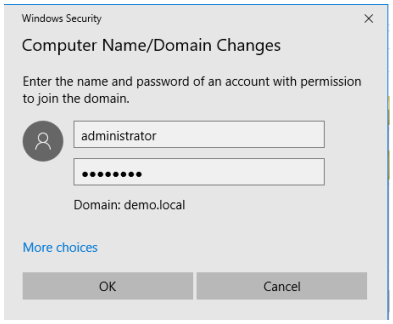

This should be all for joining the domain, please repeat the same steps for hyperv-2.

5. Configuring node roles

On all three nodes we need to add the “Hyper-V” role and the “Failover Clustering” feature. We are adding these to the AD server since we are using it as a central location to manage the Hyper-V nodes. However there is a slight difference when setting up the AD server, we need not specify a network interface for use by the VM network since we will not be running any VMs on that node.

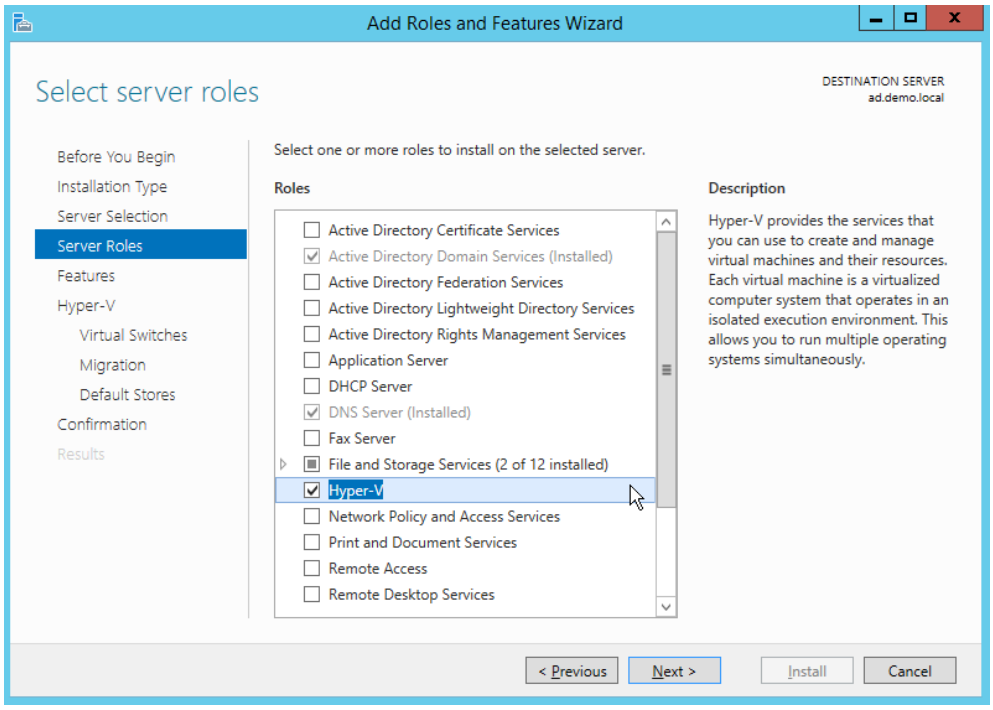

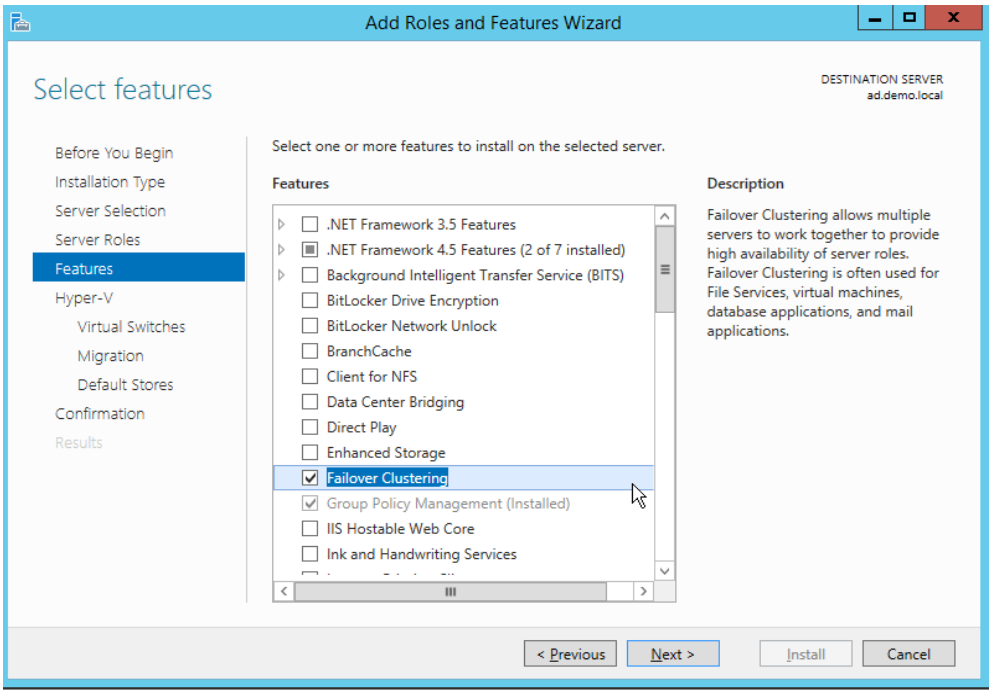

On hyperv-1 and hyper-v2 nodes specify the interface to use for the VM traffic. As discussed earlier, we reserved our fourth interface (Ethernet3) specifically for that purpose.

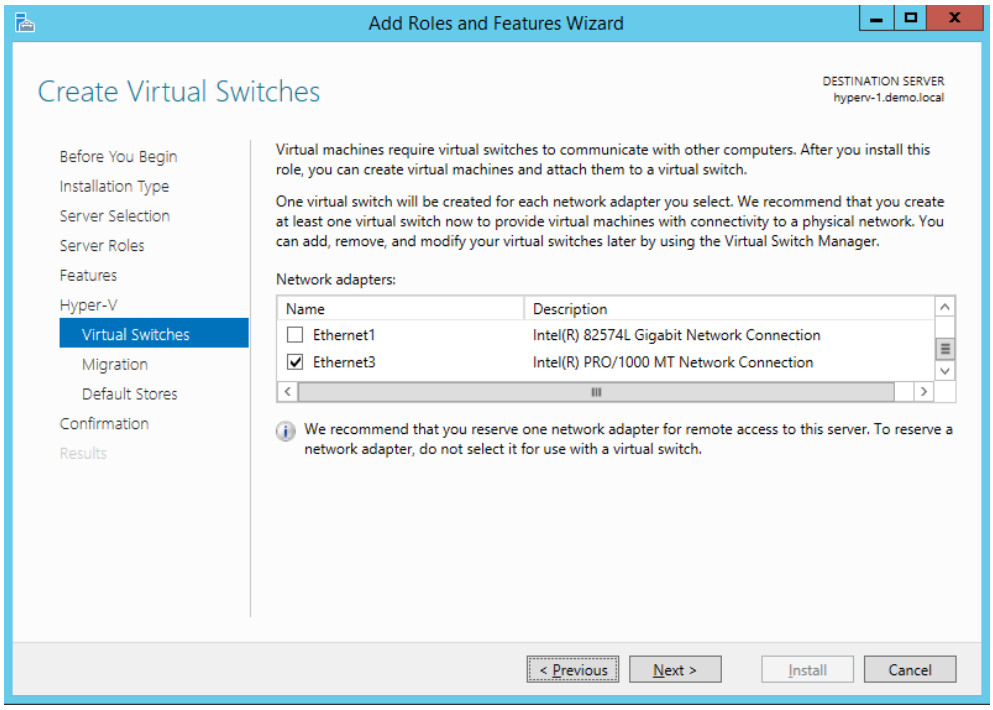

In contrast, our AD server has one interface, keep it unchecked.

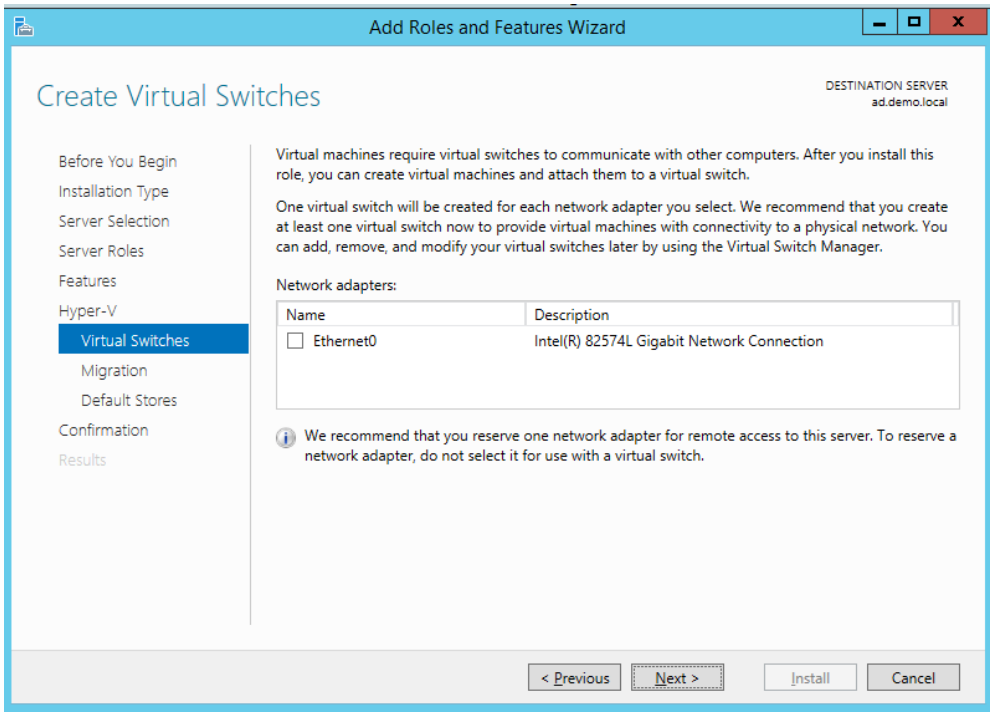

Reboot each node when done.

6. 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 Hyper-V virtual machines.
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 hyperv-1 and hyperv-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 hyperv-1.

7. 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”

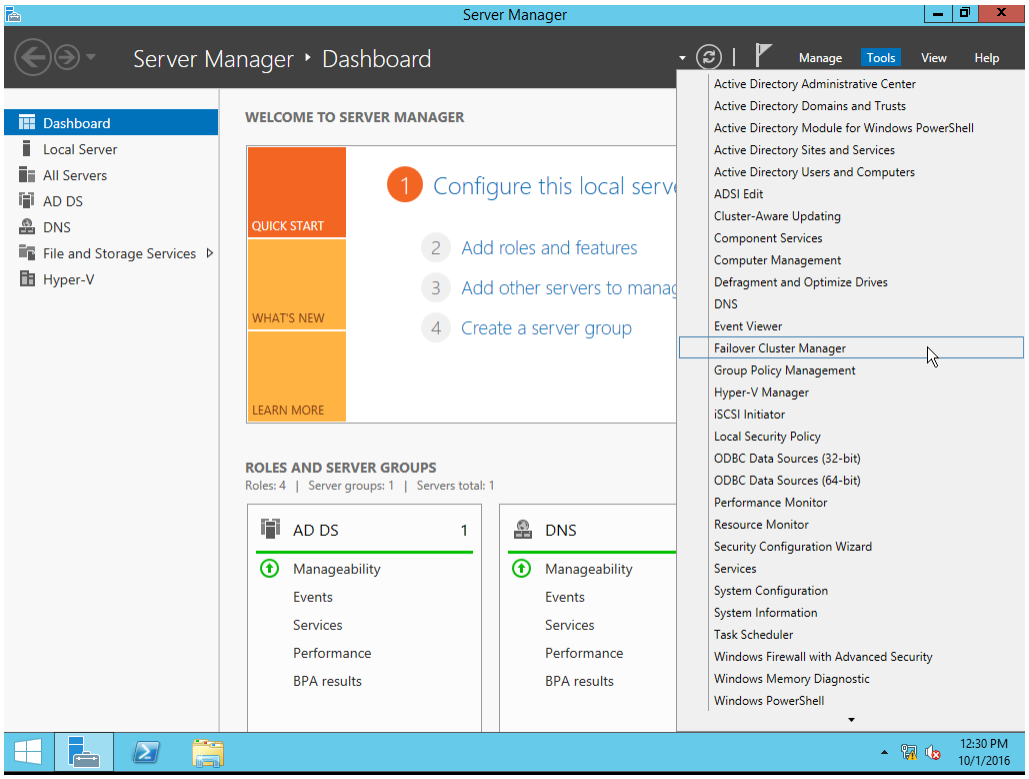

Next click on “Validate Configuration”

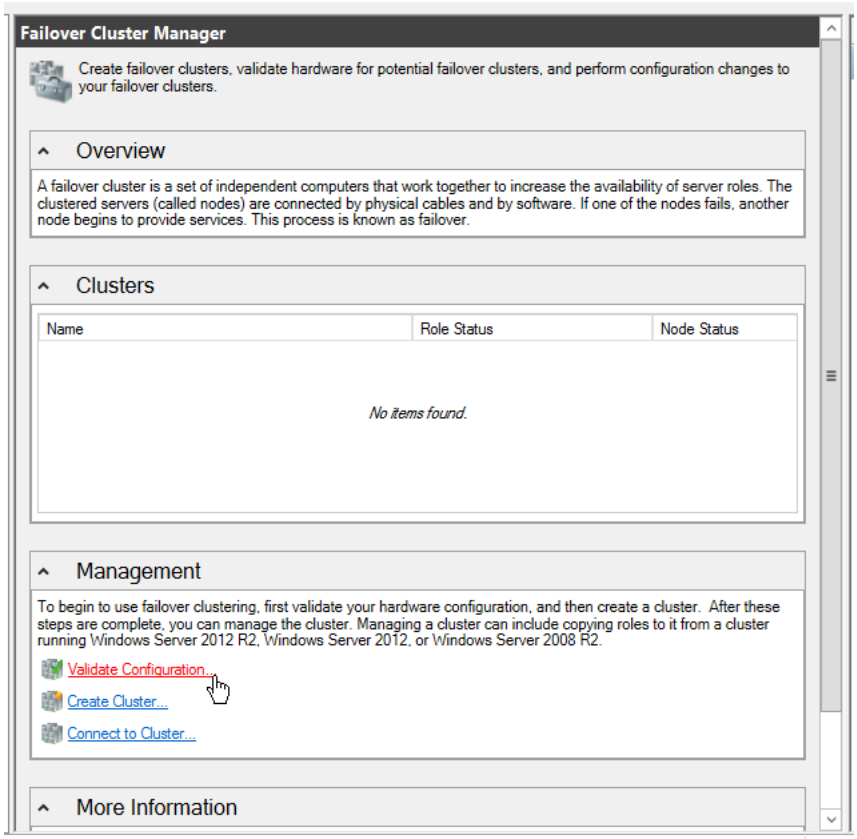

In the validation wizard, add both hyperv-1 and hyperv-2

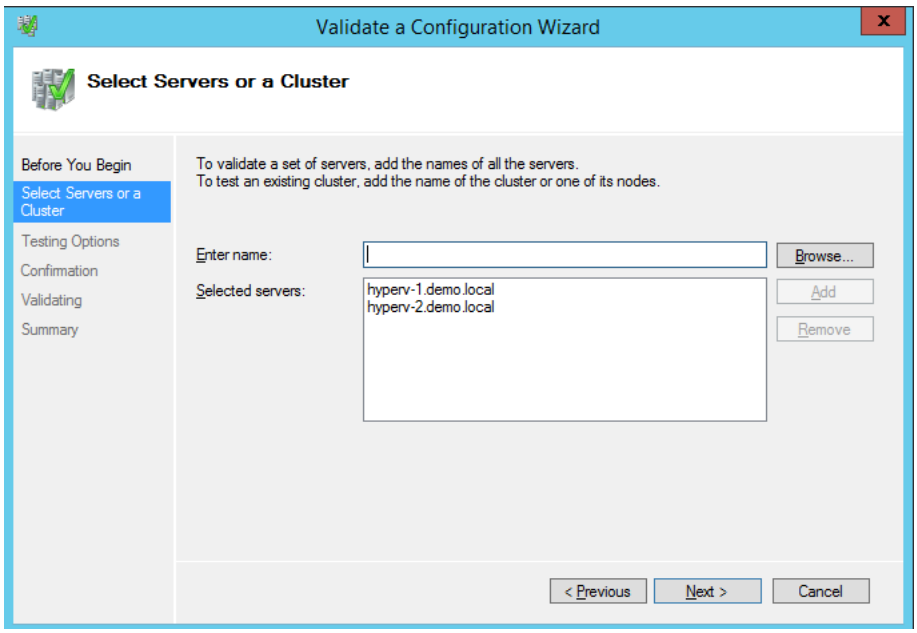

Next choose to run all tests

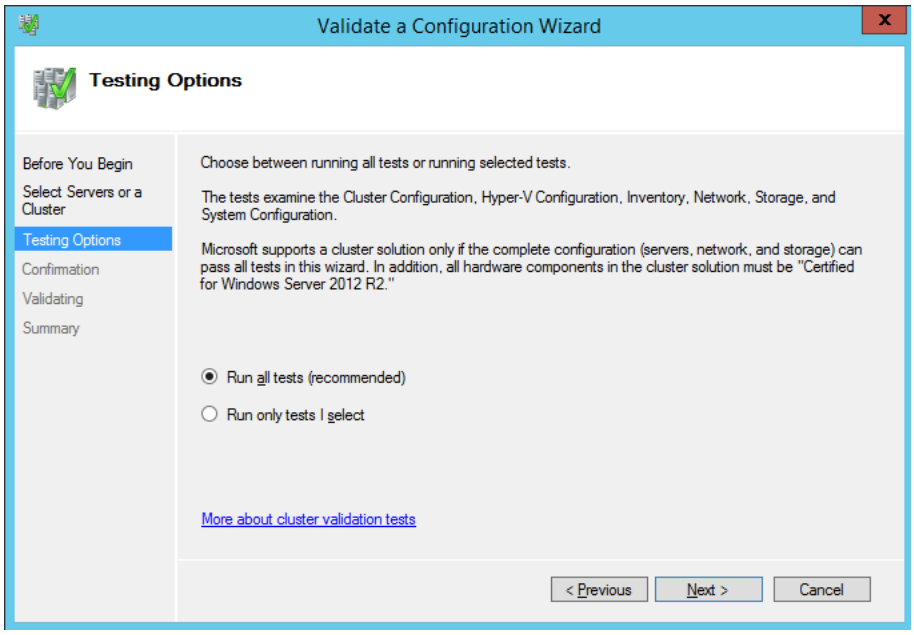

The wizard will take a couple of minutes to run various tests; these include many important storage failover and validation tests for our PetaSAN disks.

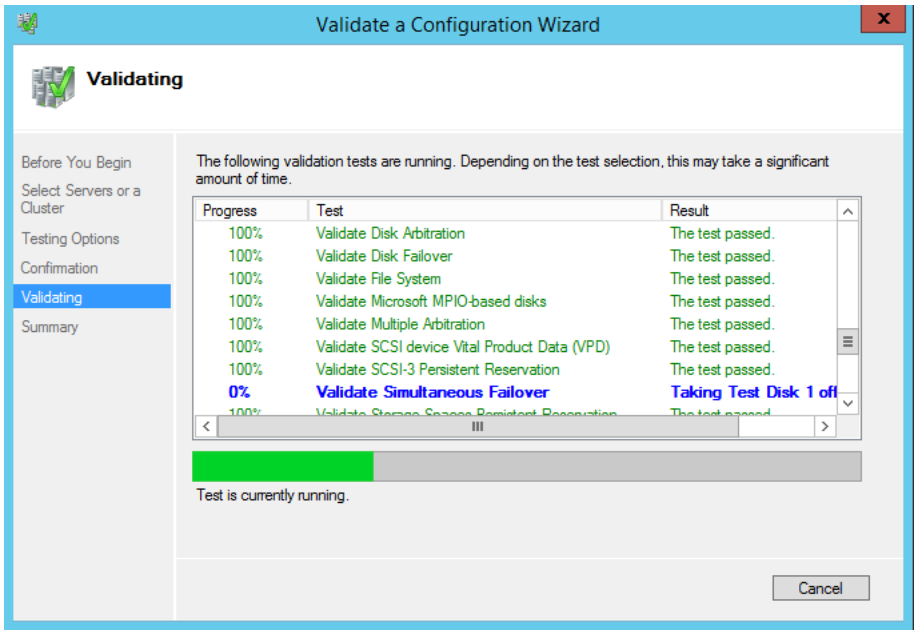

Once completed, the wizard displays a cluster validation report.

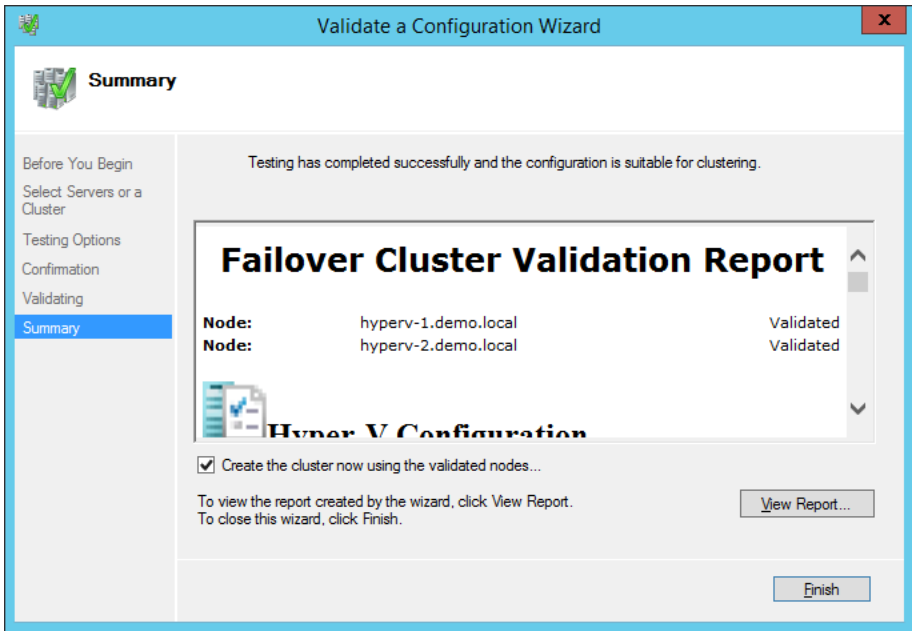

If all is good, leave the “Create the cluster now using the validated nodes” checked and click on the “Finish” button.

8. Cluster Creation

After successful validation, the create cluster wizard is displayed

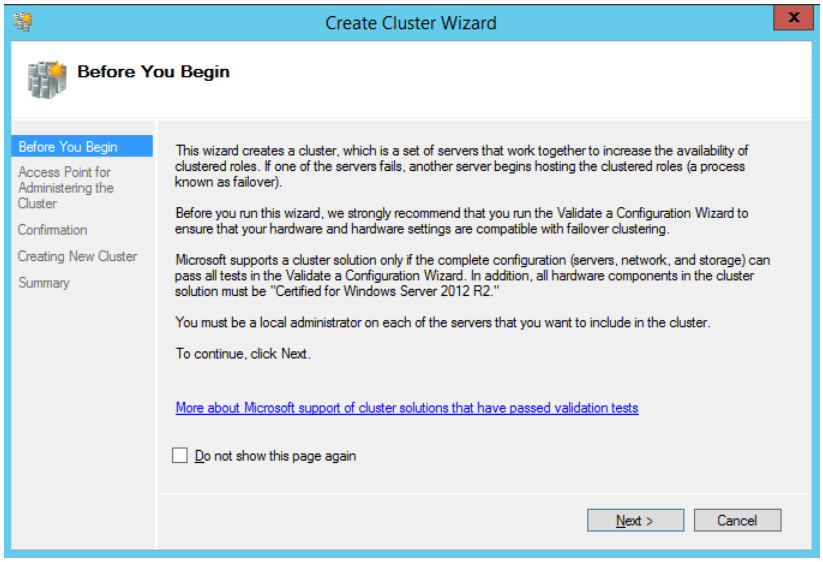

Enter the cluster name, for example “Hyper-V Cluster”

We need to assign an IP address for the cluster, in our example enter IP 10.0.1.100

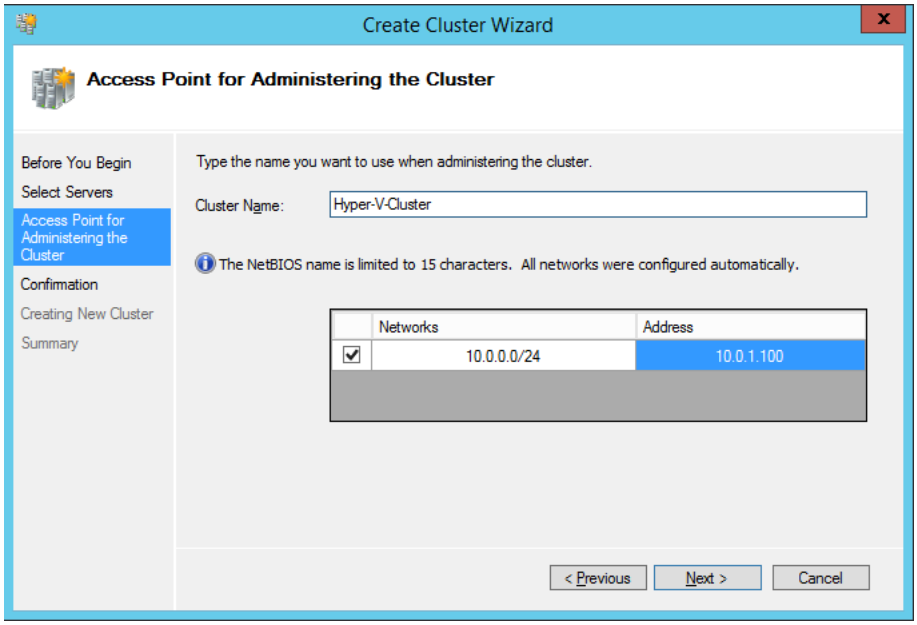

Click “Next”

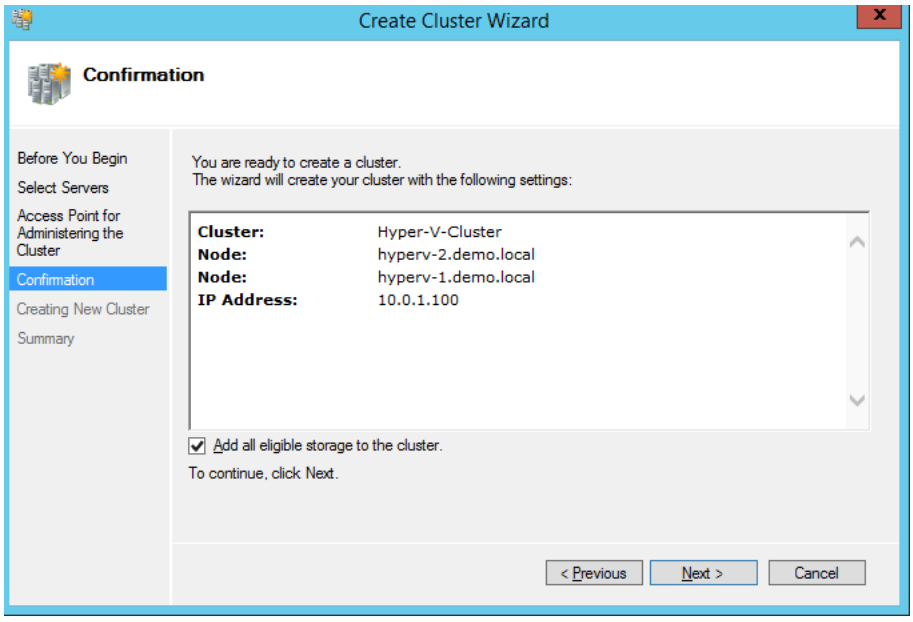

Click “Next”, then “Finish”

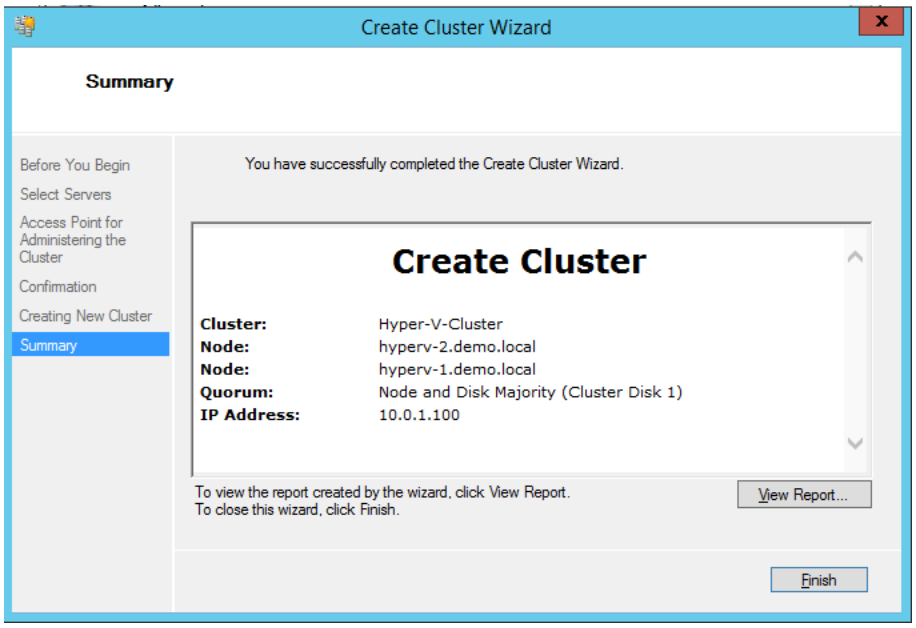

Once the cluster is created, go to Storage -> Disks

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

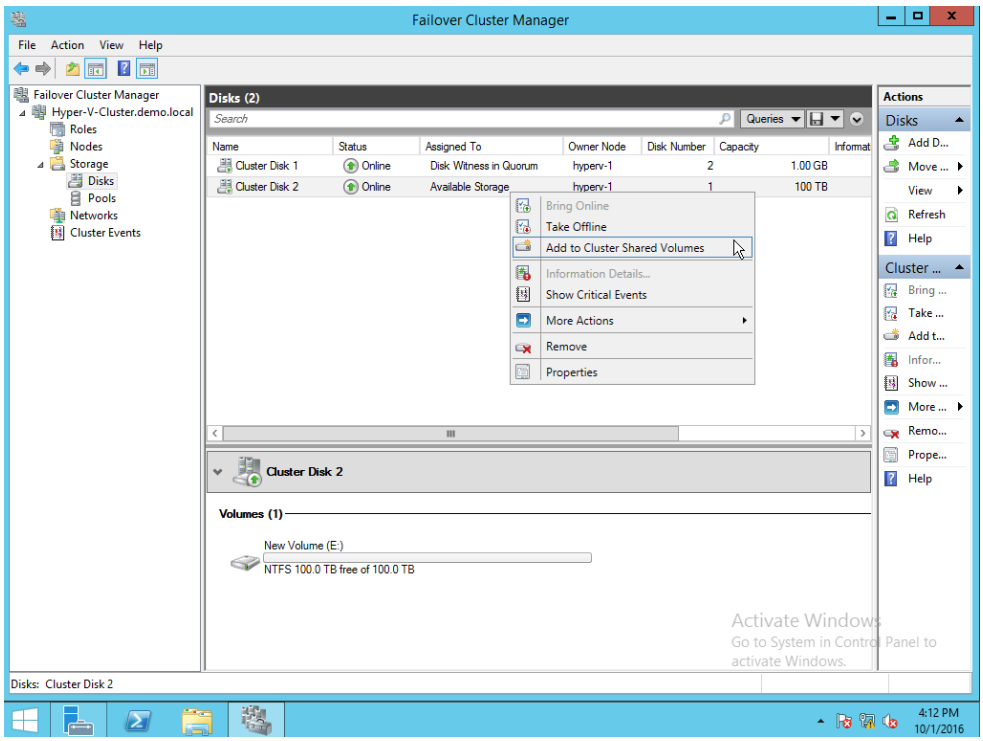

On the bottom pane, the 100 TB disk volume will change from NTFS to CSVFS (Clustered Shared Volume File System), this allows the volume to be used by many Hyper-V nodes concurrently. Notice too that it is now accessible as “C:\ClusterStorage\Volume1”.

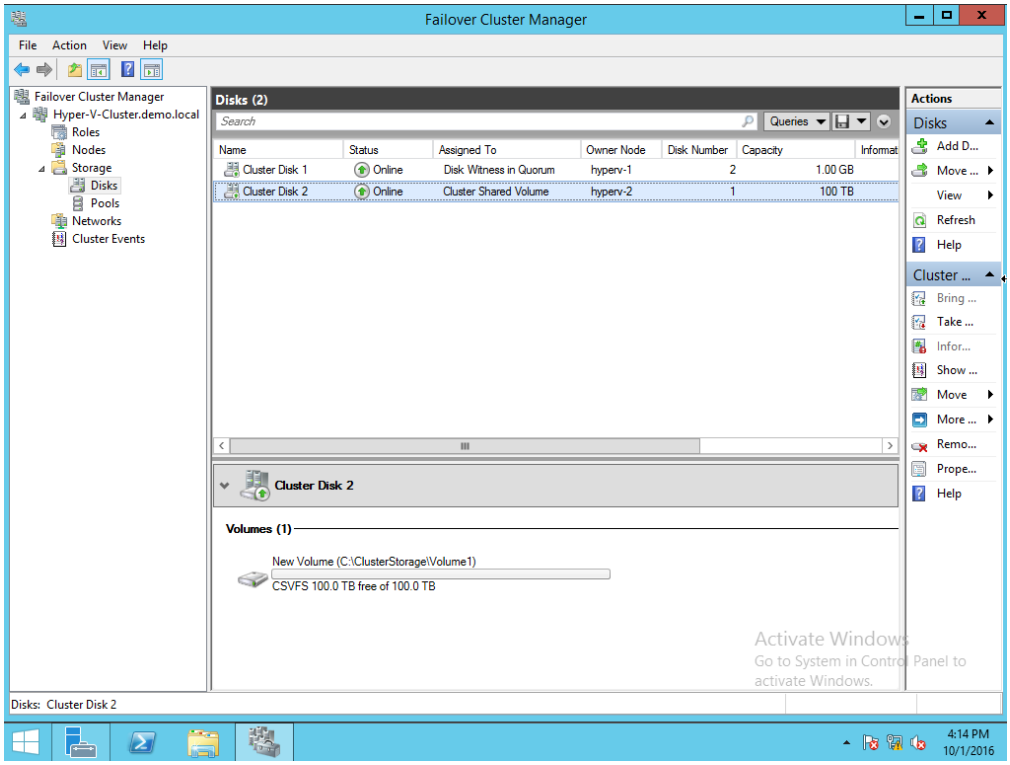

Our next step is to create virtual machines, storing them on “C:\ClusterStorage\Volume1”

9. Creating Virtual Machines

Right click on “Roles” -> “Virtual Machines…” -> “New Virtual Machine…”

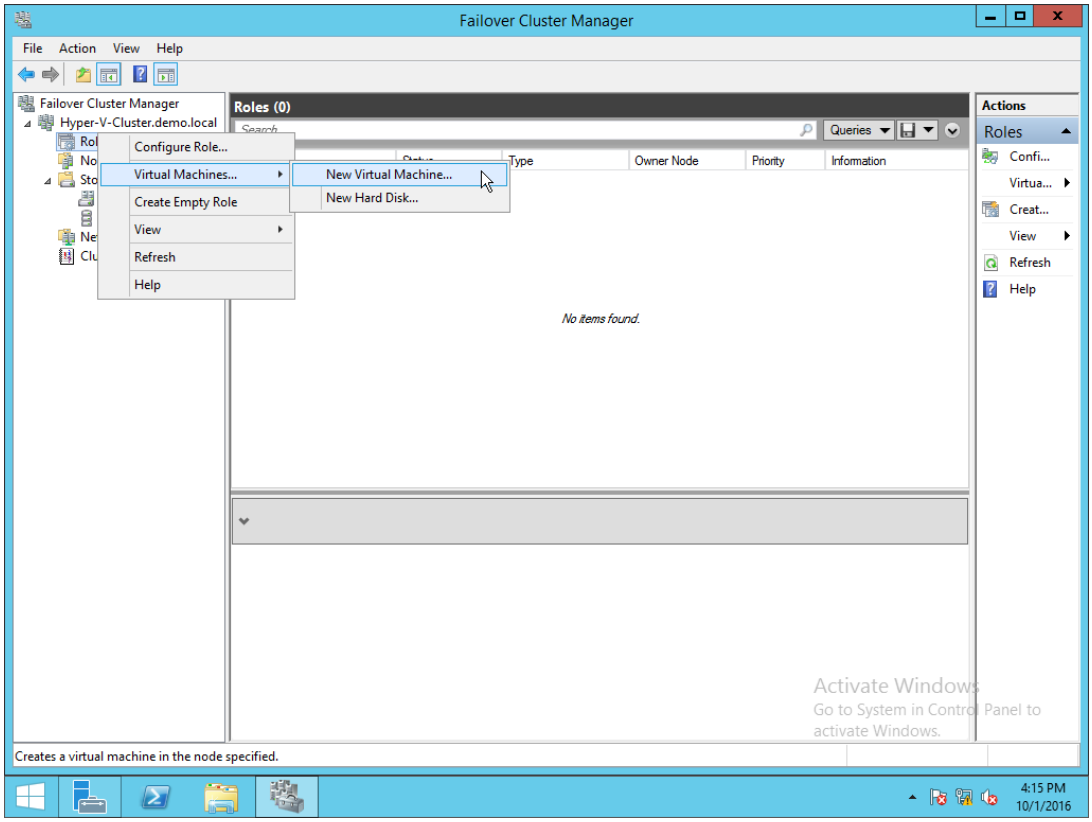

Select which hyper-v node that will (initially) house the new VM. Choose hyperv-1.

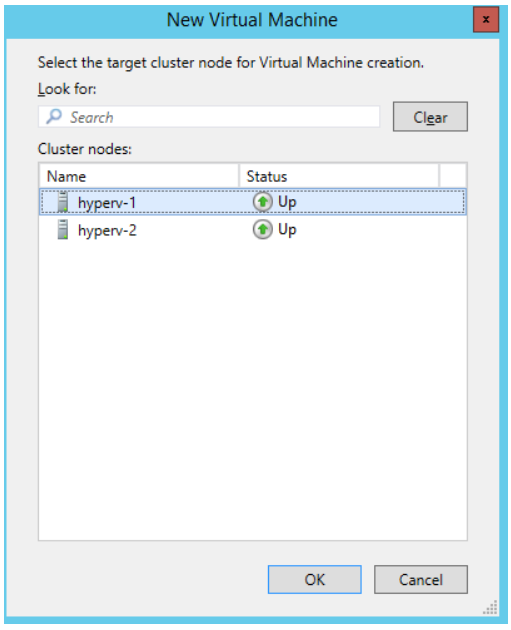

Check the “Store the virtual machine in a different location” and specify our “C:\ClusterStorage\Volume1” clustered volume.

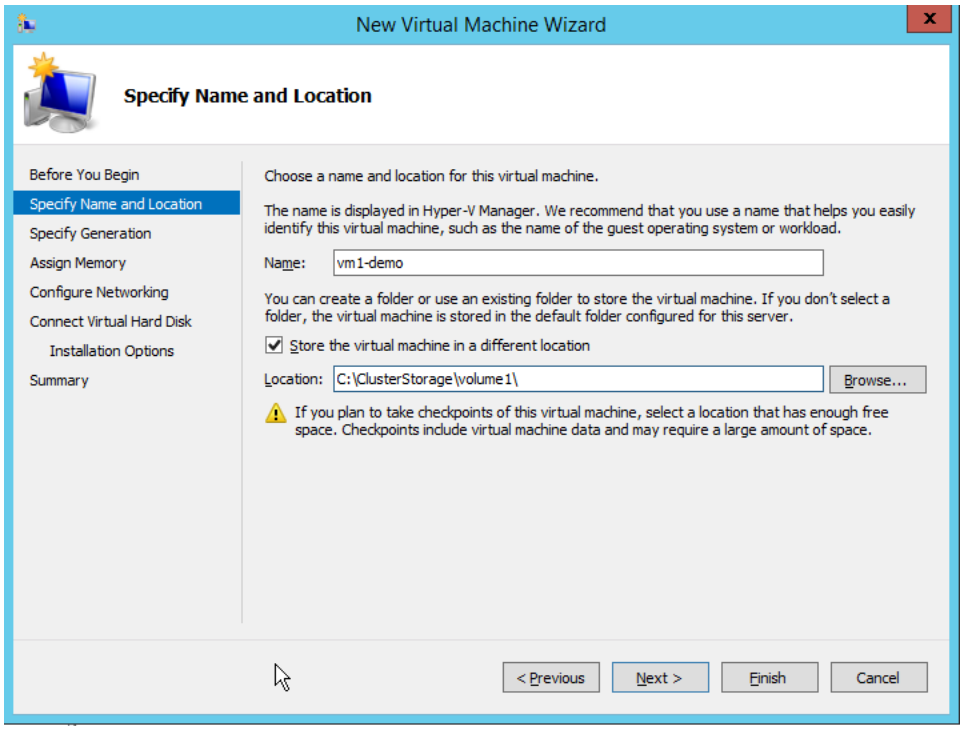

Follow the wizard to complete the creation process. The new virtual machine will initially run on hyperv-1 as we had specified but in case of node failure it will be picked up by other nodes in the cluster (in our case hyperv-2). We can also perform live migration to transfer the virtual machine from one Hyper-V node to another.