StarWind Virtual SAN® 3-node Compute and Storage Separated Scenario with Windows Server

2024

TECHNICAL PAPERS

Trademarks

Changes

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, StarWind Software assumes no liability resulting from errors or omissions in this document, or from the use of the information contained herein. StarWind Software reserves the right to make changes in the product design without reservation and without notification to its users.

Technical Support and Services

If you have questions about installing or using this software, check this and other documents first - you will find answers to most of your questions on the Technical Papers webpage or in StarWind Forum. If you need further assistance, please contact us .

About StarWind

StarWind is a pioneer in virtualization and a company that participated in the development of this technology from its earliest days. Now the company is among the leading vendors of software and hardware hyper-converged solutions. The company’s core product is the years-proven StarWind Virtual SAN, which allows SMB and ROBO to benefit from cost-efficient hyperconverged IT infrastructure. Having earned a reputation of reliability, StarWind created a hardware product line and is actively tapping into hyperconverged and storage appliances market. In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms” following the success and popularity of StarWind HyperConverged Appliance. StarWind partners with world-known companies: Microsoft, VMware, Veeam, Intel, Dell, Mellanox, Citrix, Western Digital, etc.

Copyright ©2009-2018 StarWind Software Inc.

Applies To: Windows Server 2016, Windows Server 2019, Windows Server 2022

Annotation

Relevant products

This guide applies to StarWind Virtual SAN and StarWind Virtual SAN Free (Version V8 (build 15260) and earlier).

Purpose

This document outlines how to configure a Microsoft Hyper-V Failover Cluster using StarWind Virtual SAN (VSAN), with VSAN running as a Windows application. The guide includes steps to prepare Hyper-V hosts for clustering, configure physical and virtual networking, and set up the StarWind VSAN and devices.

For more information about StarWind VSAN architecture and available installation options, please refer to the
StarWind Virtual (VSAN) Getting Started Guide.

Audience

This technical guide is intended for storage and virtualization architects, system administrators, and partners designing virtualized environments using StarWind Virtual SAN (VSAN).

Expected result

The end result of following this guide will be a fully configured high-availability Windows Failover Cluster that includes virtual machine shared storage provided by StarWind VSAN.

Prerequisites

StarWind Virtual SAN system requirements

Prior to installing StarWind Virtual SAN, please make sure that the system meets the requirements, which are available via the following link:
https://www.starwindsoftware.com/system-requirements

Recommended RAID settings for HDD and SSD disks:

https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/

Please read StarWind Virtual SAN Best Practices document for additional information: https://www.starwindsoftware.com/resource-library/starwind-virtual-san-best-practices

Solution diagram

The diagrams below illustrate the network and storage configuration of the solution:

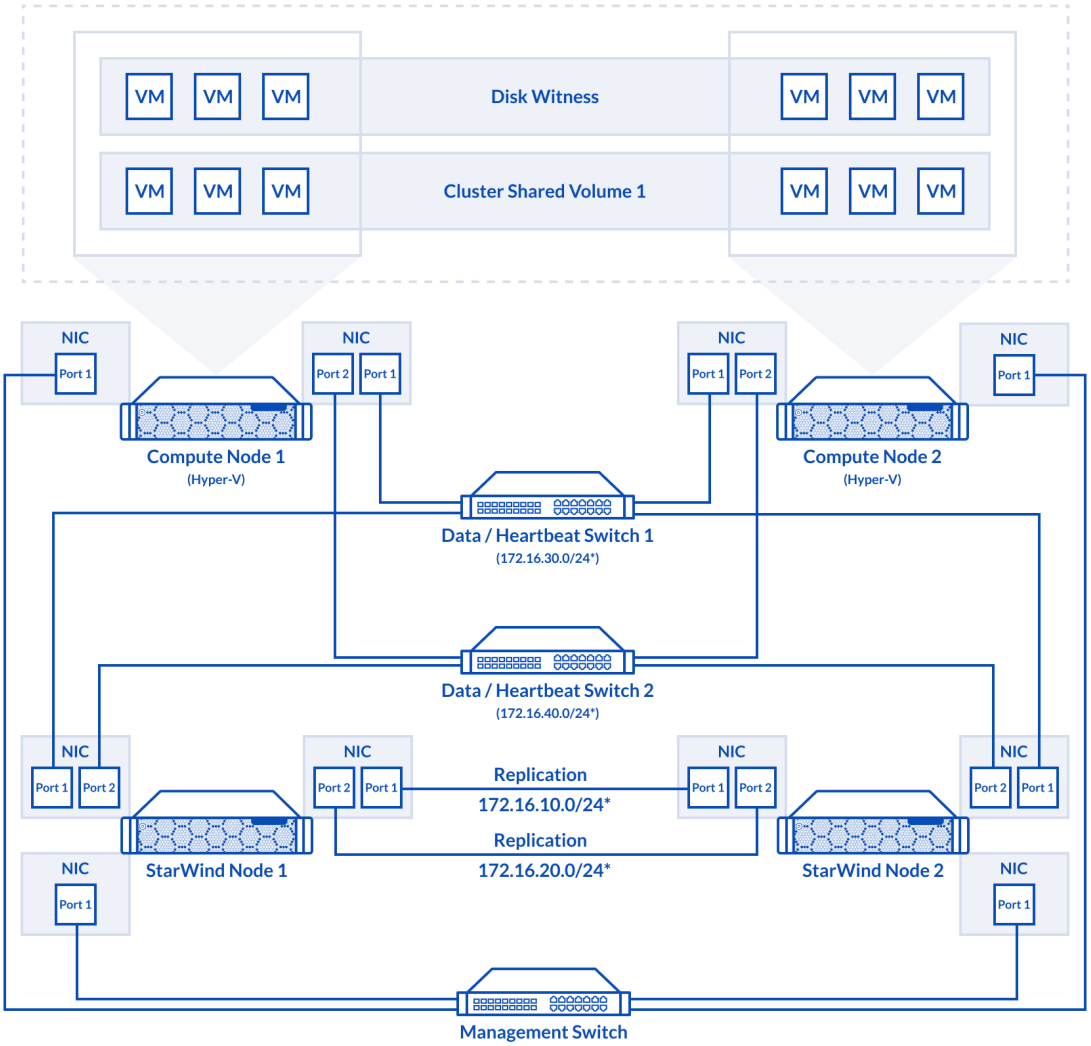

2-node cluster

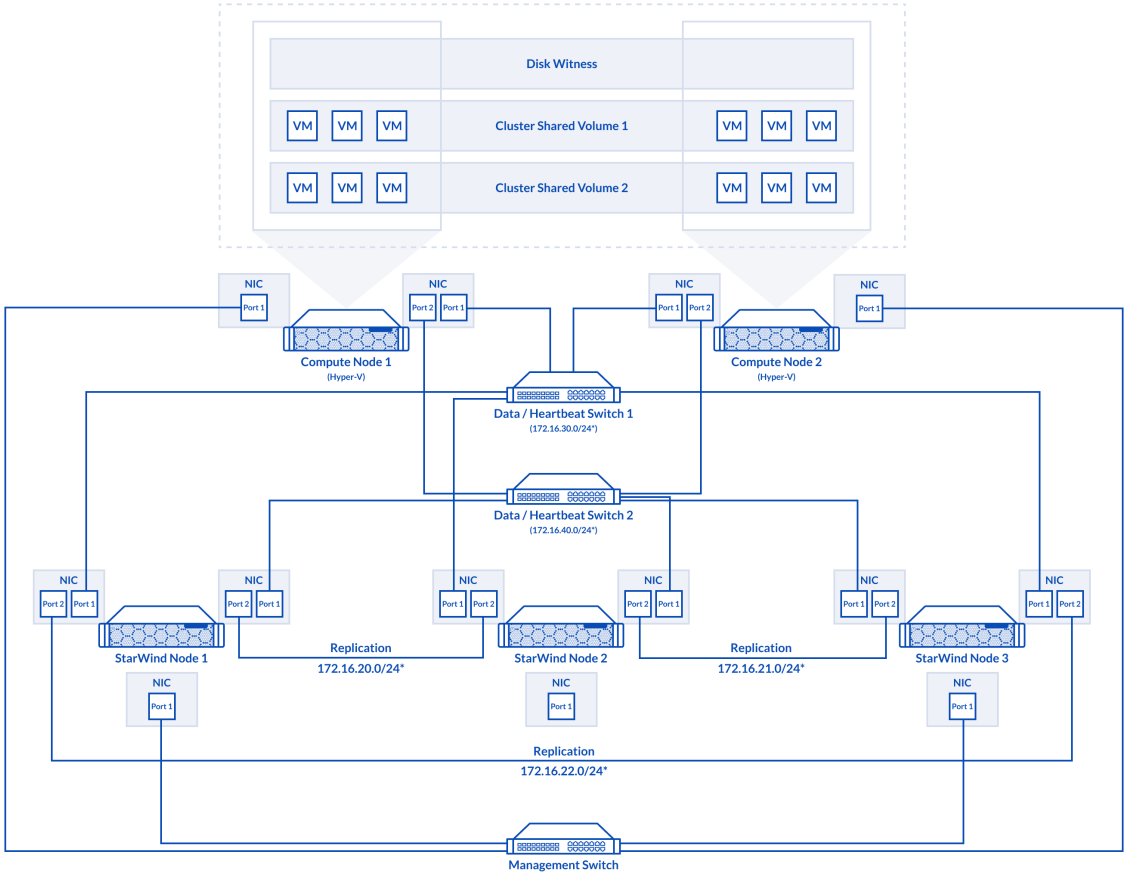

3-node cluster

Preconfiguring cluster nodes

1. Make sure that a domain controller is configured and the servers are added to the domain.
NOTE: Please follow the recommendation in KB article on how to place a DC in case of StarWind Virtual SAN usage.

2. Install Failover Clustering and Multipath I/O features, as well as the Hyper-V role on both cluster nodes. This can be done through the Server Manager (Add Roles and Features) menu item.

3. For a 2-node StarWind setup, configure network interfaces on each node to make sure that Synchronization and iSCSI/StarWind heartbeat interfaces are in different subnets and connected according to the network diagram above. In this document, 172.16.10.x and 172.16.20.x subnets are used for the Synchronization traffic, while

172.16.30.x and 172.16.40.x subnets are used for iSCSI/StarWind heartbeat traffic.

For a 3-node StarWind setup, configure the network interfaces on each node to make sure that the Synchronization and iSCSI/StarWind heartbeat interfaces are in different subnets and connected according to the network diagram above. In this document, 172.16.30.x, 172.16.40.x, subnets are used for the iSCSI/StarWind heartbeat traffic, while 172.16.20.x,172.16.21.x, 172.16.22.x subnets are used for the Synchronization traffic.

4. In order to allow iSCSI Initiators to discover all StarWind Virtual SAN interfaces, the StarWind configuration file (StarWind.cfg) should be changed after stopping the StarWind service on the node where it will be edited. Locate the StarWind Virtual SAN configuration file (the default path is “C:\Program Files\StarWind Software\StarWind\StarWind.cfg”) and open it via WordPad as Administrator. Find the *<iScsiDiscoveryListInterfaces value=”0”/>* string and change the value from 0 to 1 (should look as follows: *<iScsiDiscoveryListInterfaces value=”1”/>*). Save the changes and exit Wordpad. Once StarWind.cfg is changed and saved, the StarWind service can be restarted.

Enabling Multipath Support

5. Open the MPIO Properties manager: Start -> Windows Administrative Tools -> MPIO. Alternatively, run the following PowerShell command:

```
mpiocpl
```

6. In the Discover Multi-Paths tab, choose the Add support for iSCSI devices checkbox and click Add.

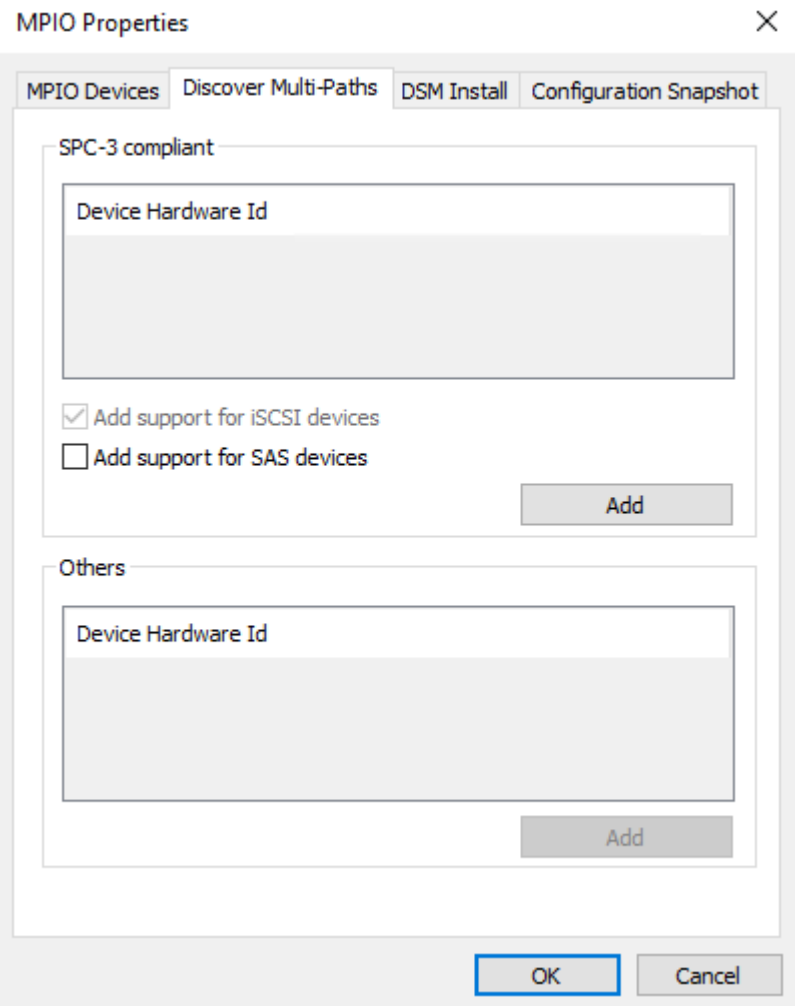

7. When prompted to restart the server, click Yes to proceed.

8. Repeat the same procedure on the other server.

Installing File Server Roles

Please follow the steps below if file shares configuration is required

Scale-Out File Server (Sofs) For Application Data

1. Open Server Manager: Start -> Server Manager.

2. Select: Manage -> Add Roles and Features.

3. Follow the installation wizard steps to install the roles selected in the screenshot below:

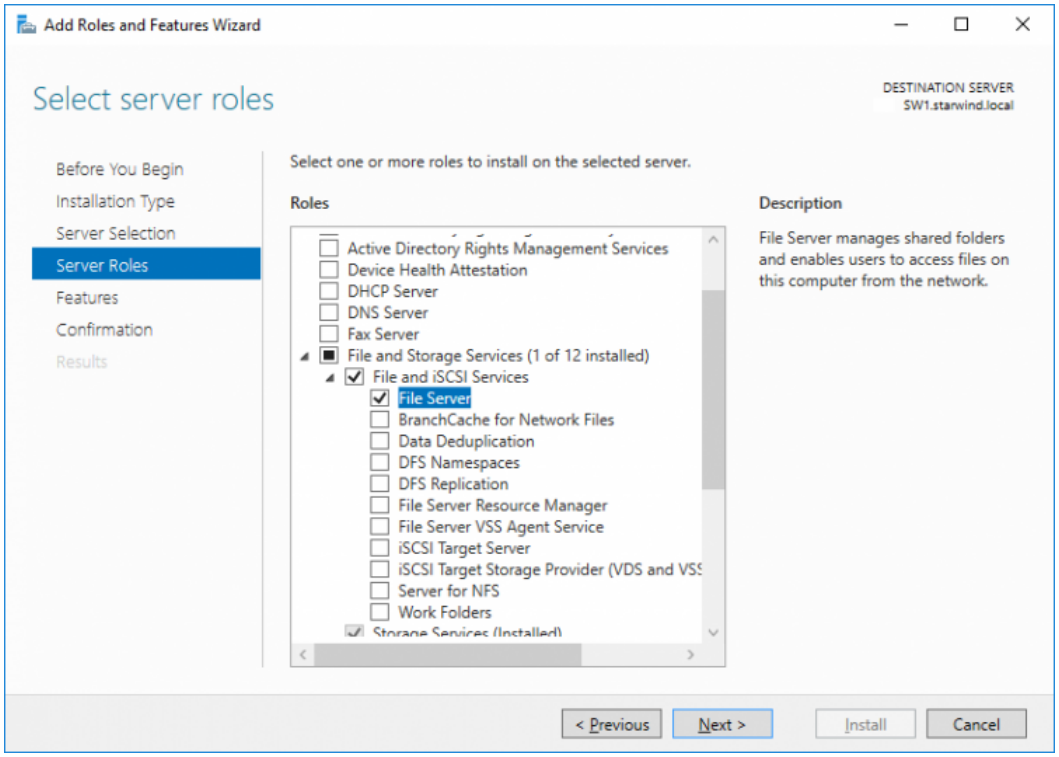

4. Restart the server after installation is completed and perform steps above on the each server.

File Server For General Use With Smb Share

1. Open Server Manager: Start -> Server Manager.

2. Select: Manage -> Add Roles and Features.

3. Follow the installation wizard steps to install the roles selected in the screenshot below:

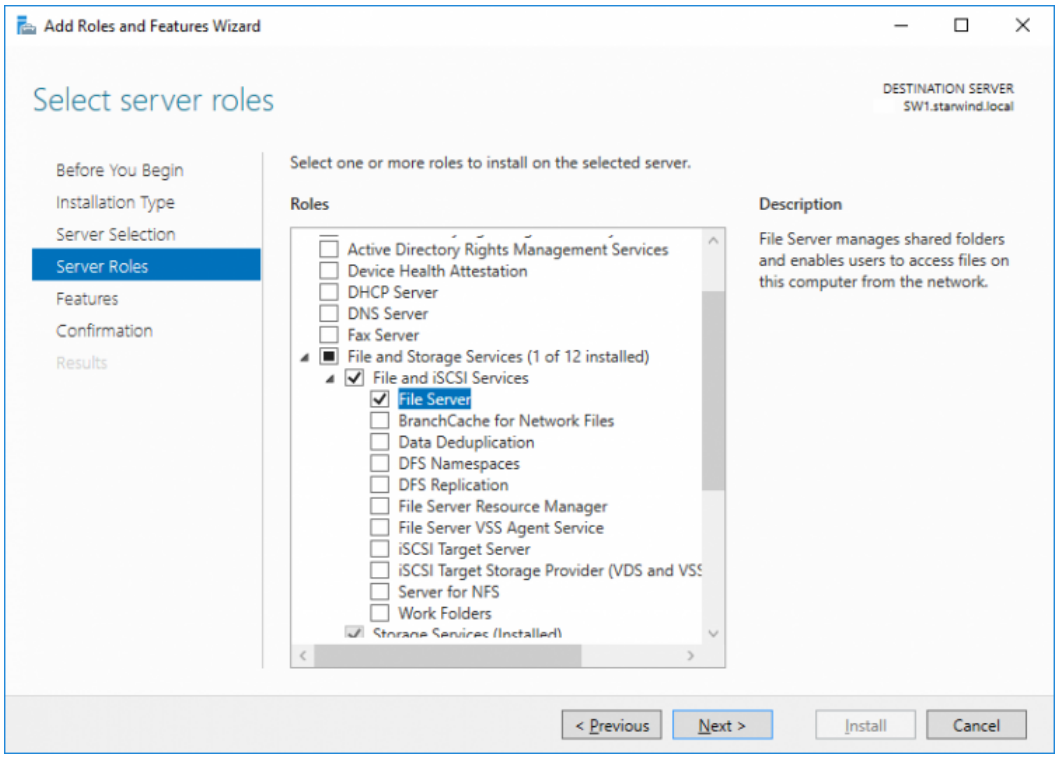

4. Restart the server after installation is completed and perform steps above on each server.

File Server For General Use With Nfs Share

1. Open Server Manager: Start -> Server Manager.

2. Select: Manage -> Add Roles and Features.

3. Follow the installation wizard steps to install the roles selected in the screenshot below:

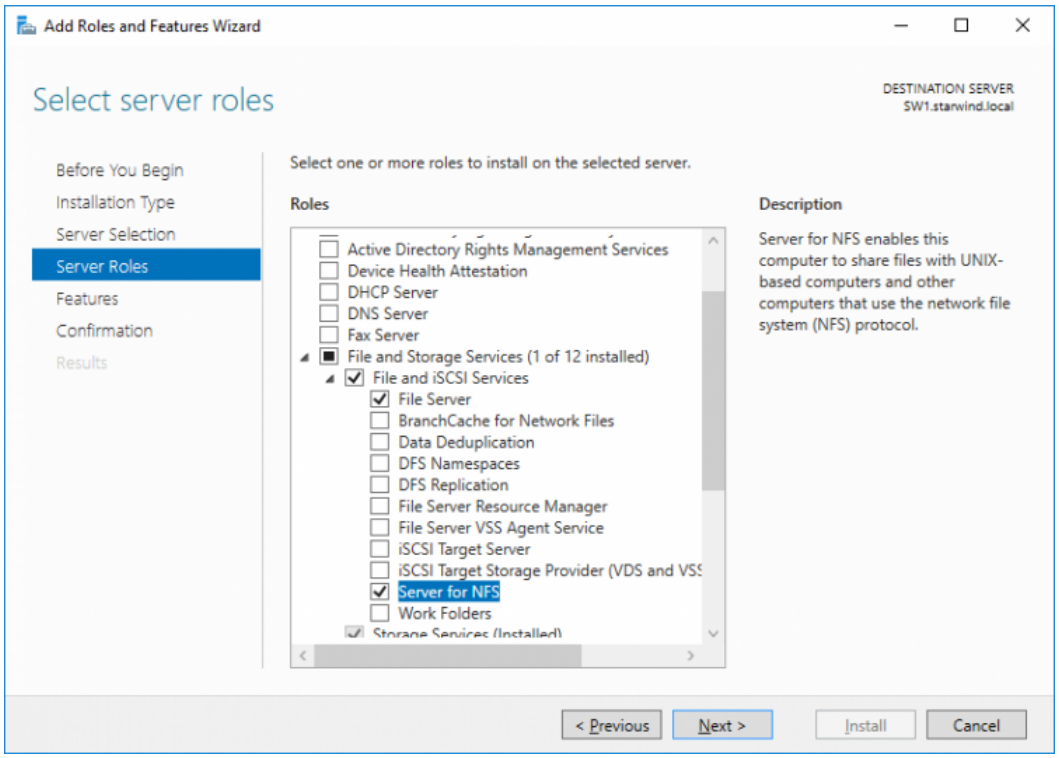

4. Restart the server after installation is completed and perform steps above on each server.

Installing Starwind Vsan For Hyper-V

1. Download the StarWind setup executable file from the StarWind website: https://www.starwind.com/registration-starwind-virtual-san

2. Launch the downloaded setup file on the server to install StarWind Virtual SAN or one of its components. The Setup wizard will appear. Read and accept the License Agreement.

3. Carefully read the information about the new features and improvements. Red text indicates warnings for users that are updating the existing software installations.

4. Select Browse to modify the installation path if necessary. Click on Next to continue.

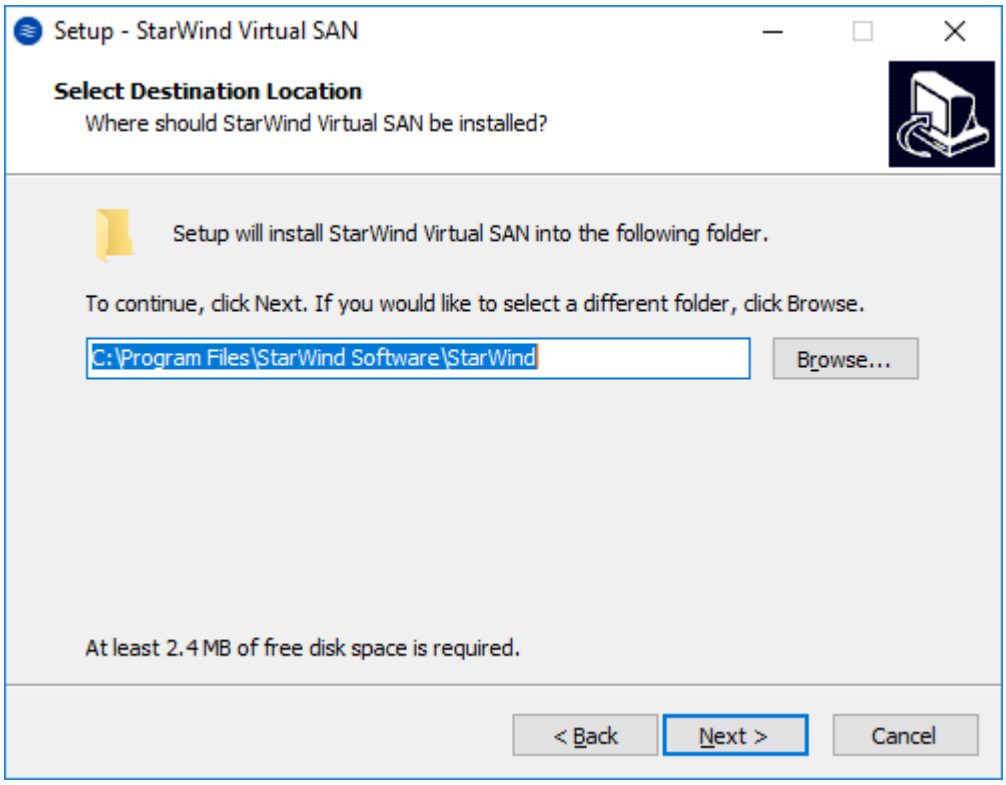

5. Select the following components for the minimum setup:

- StarWind Virtual SAN Service. The StarWind Virtual SAN service is the “core” of the software. It can create iSCSI targets as well as share virtual and physical devices. The service can be managed from StarWind Management Console on any Windows computer that is on the same network. Alternatively, the service can be managed from StarWind Web Console deployed separately.
- StarWind Management Console. Management Console is the Graphic User Interface (GUI) part of the software that controls and monitors all storage-related operations (e.g., allows users to create targets and devices on StarWind Virtual SAN servers connected to the network).
  NOTE: To manage StarWind Virtual SAN installed on a Windows Server Core edition with no GUI, StarWind Management Console should be installed on a different computer running the GUI-enabled Windows edition.

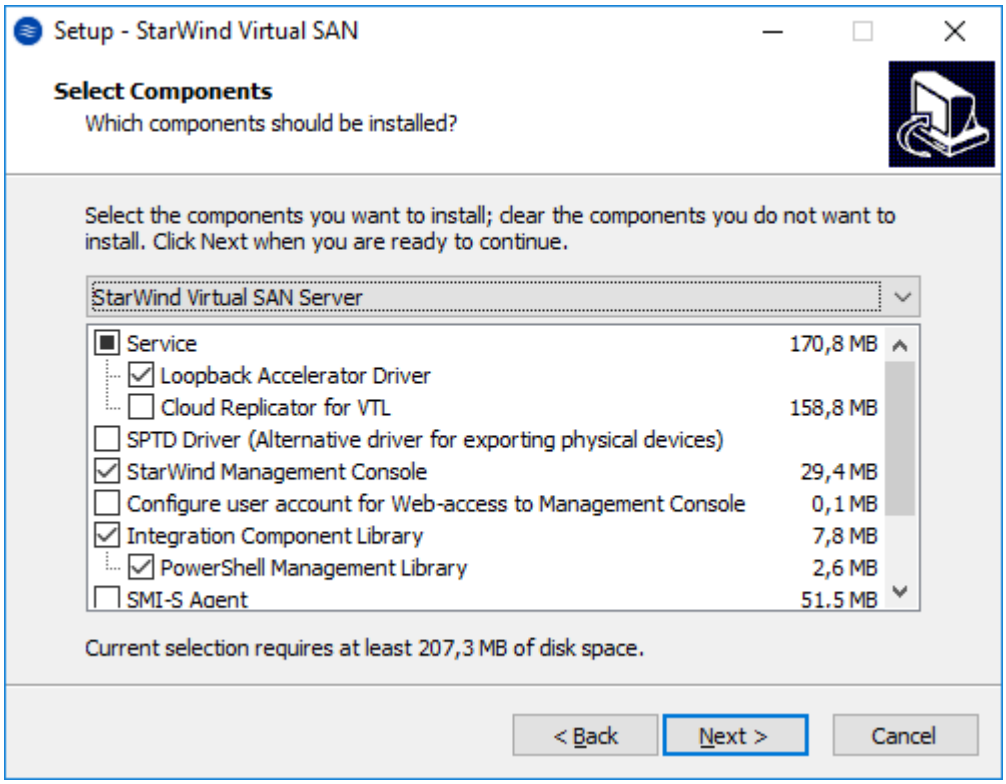

6. Specify Start Menu Folder.

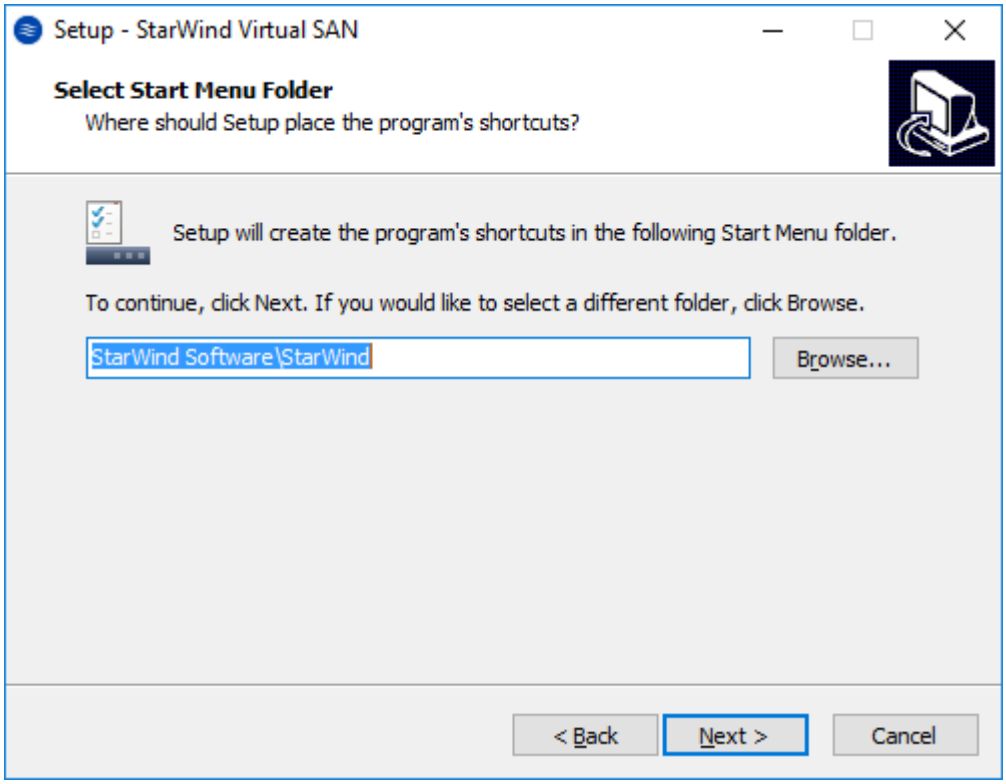

7. Enable the checkbox if a desktop icon needs to be created. Click on Next to continue.

8. When the license key prompt appears, choose the appropriate option:

- request time-limited fully functional evaluation key.
- request FREE version key.
- relect the previously purchased commercial license key.

9. Click on the Browse button to locate the license file.

10. Review the licensing information.

11. Verify the installation settings. Click on Back to make any changes or Install to proceed with installation.

12. Enable the appropriate checkbox to launch StarWind Management Console right after the setup wizard is closed and click on Finish.

13. Repeat the installation steps on the partner node.

Creating Starwind Devices

1. In the StarWind Management Console click to Add Device (advanced) button and open Add Device (advanced) Wizard.

2. Select Hard Disk Device as the type of device to be created.

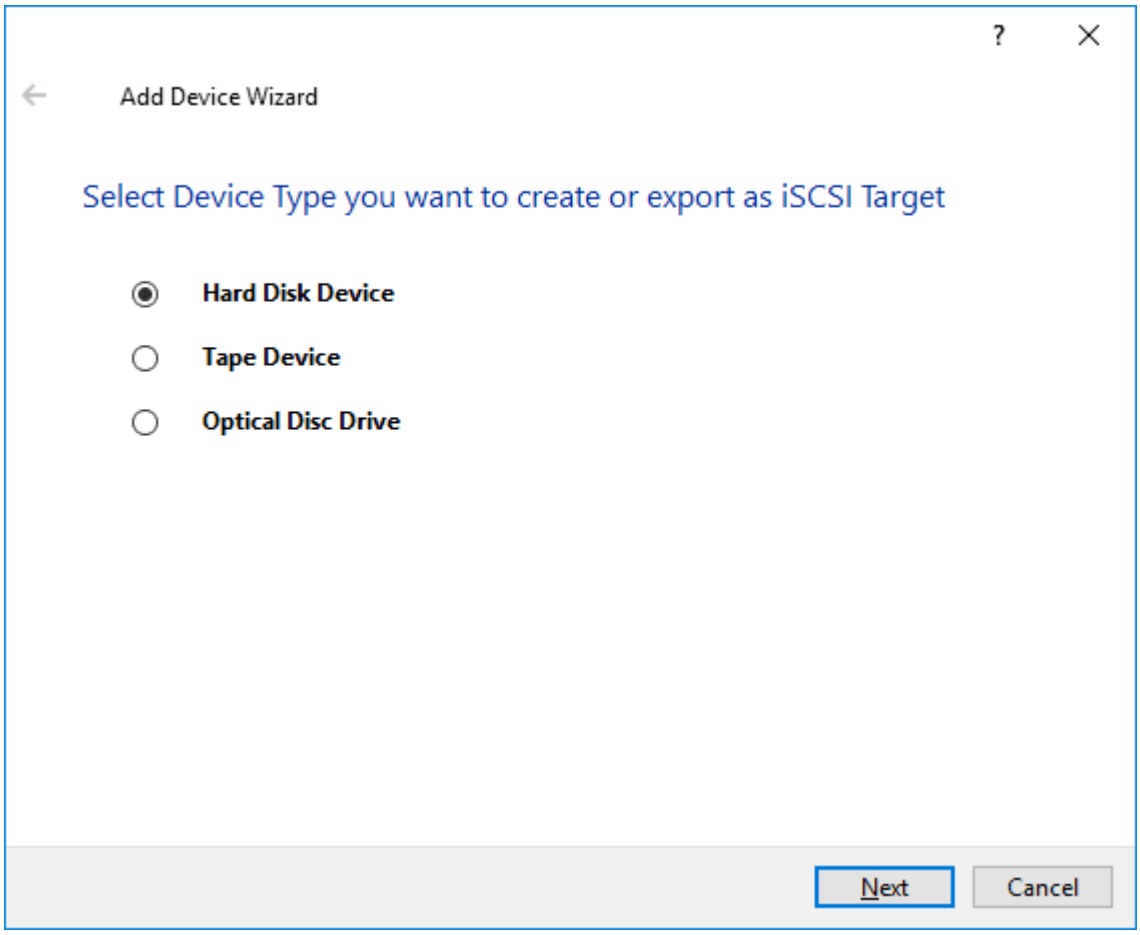

3. Select Virtual Disk.

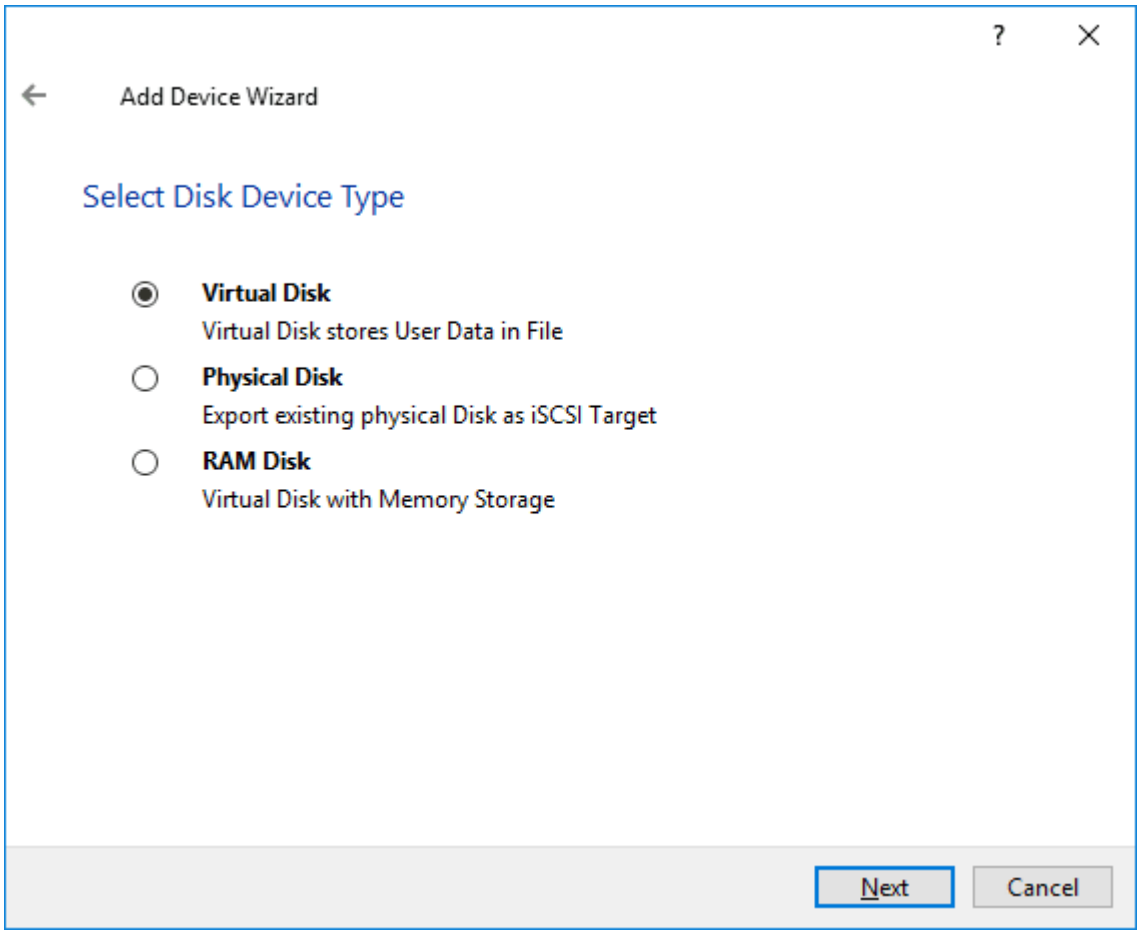

4. Specify a virtual disk Name, Location, and Size.

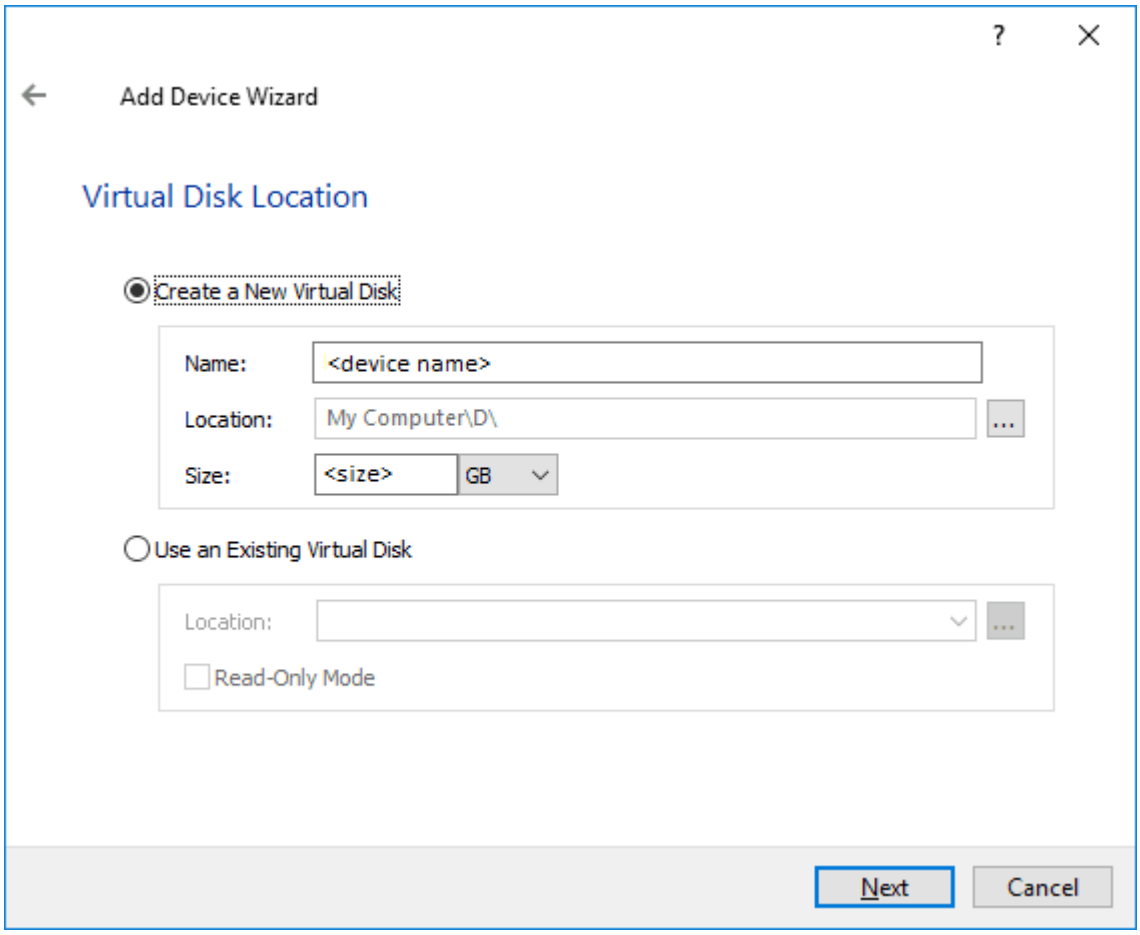

5. Select the Thick provisioned disk type and block size.
NOTE: Use 4096 sector size for targets, connected on Windows-based systems and 512 bytes sector size for targets, connected on Linux-based systems (ESXi/Xen/KVM).

6. Define a caching policy and specify a cache size (in MB). Also, the maximum available cache size can be specified by selecting the appropriate checkbox. Optionally, define the L2 caching policy and cache size.

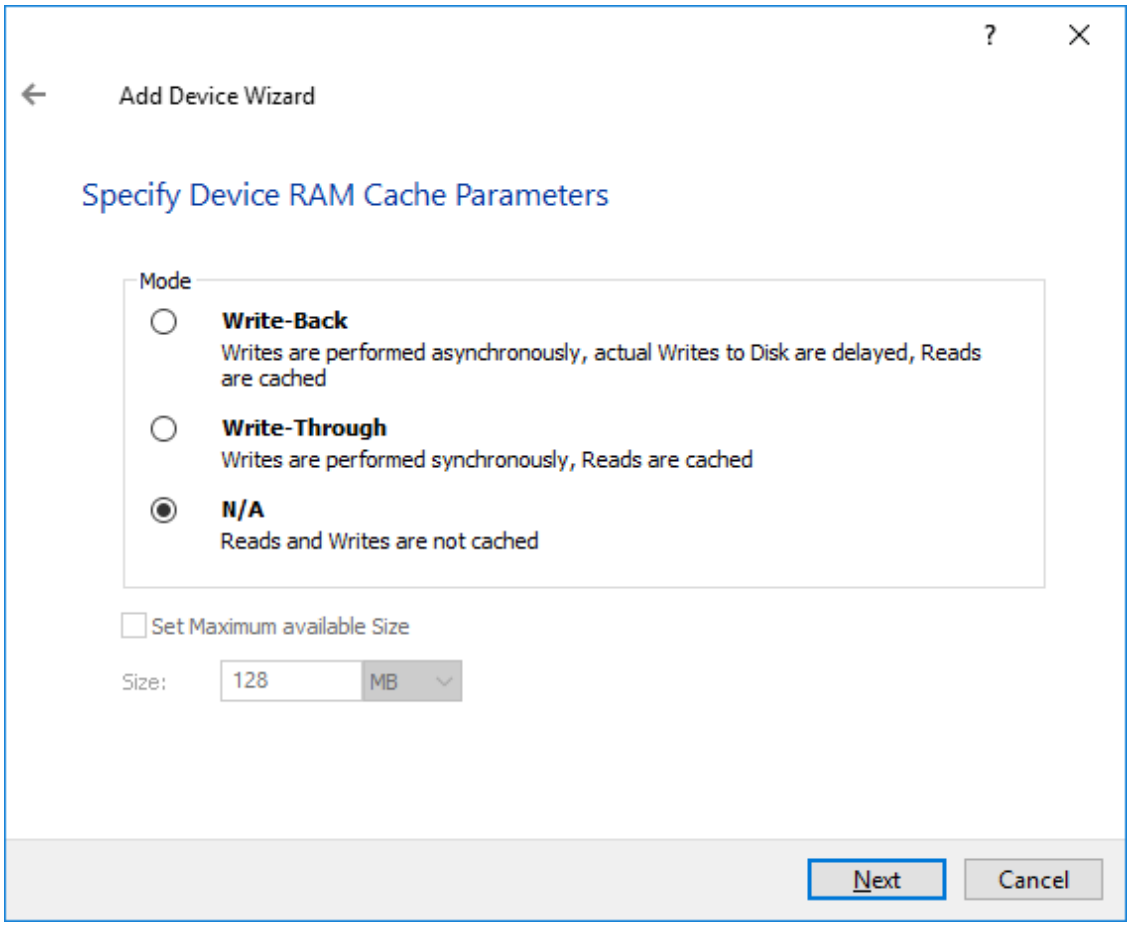

7. Specify Target Parameters. Select the Target Name checkbox to enter a custom target name. Otherwise, the name is generated automatically in accordance with the specified target alias.

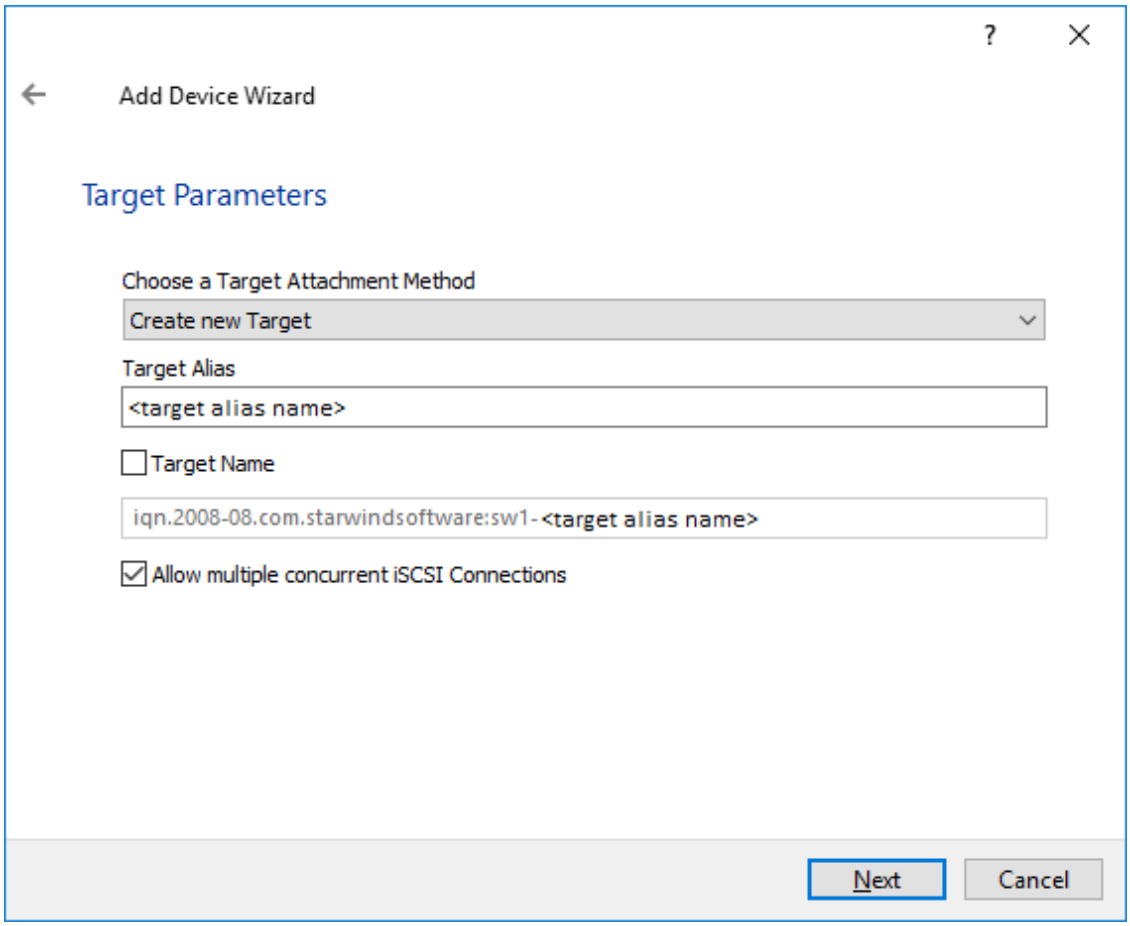

8. Click Create to add a new device and attach it to the target.

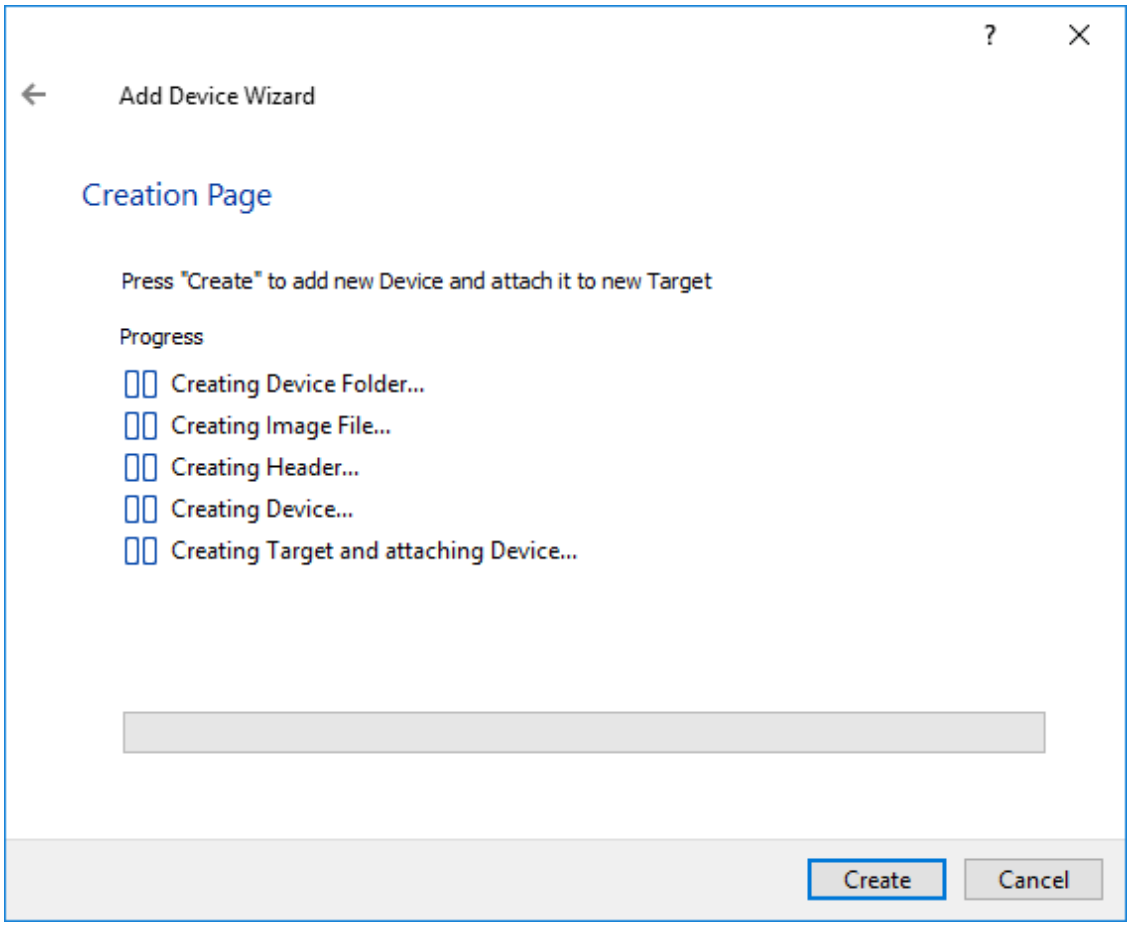

9. Click Close to finish the device creation.

10. The successfully added devices appear in the StarWind Management Console.

Select The Required Replication Mode

The replication can be configured using Synchronous “Two-Way” Replication mode: Synchronous or active-active replication ensures real-time synchronization and load balancing of data between two or three cluster nodes. Such a configuration tolerates the failure of two out of three storage nodes and enables the creation of an effective business continuity plan. With synchronous mirroring, each write operation requires control confirmation from both storage nodes. It guarantees the reliability of data transfers but is demanding in bandwidth since mirroring will not work on high-latency networks.

Synchronous “Two-Way” Replication

1. Right-click the recently created device and select Replication Manager from the shortcut menu.

2. Select the Add Replica button in the top menu.

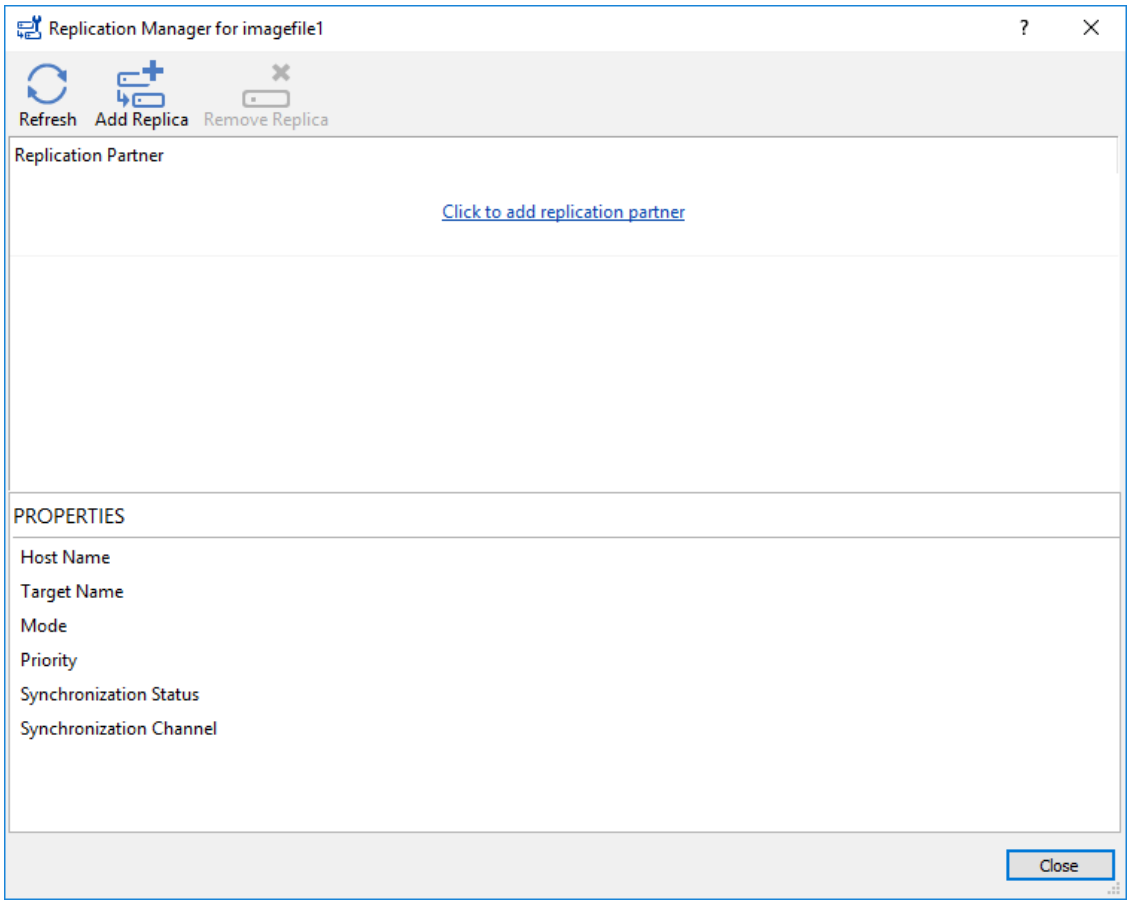

3. Select Synchronous “Two-Way” replication as a replication mode.

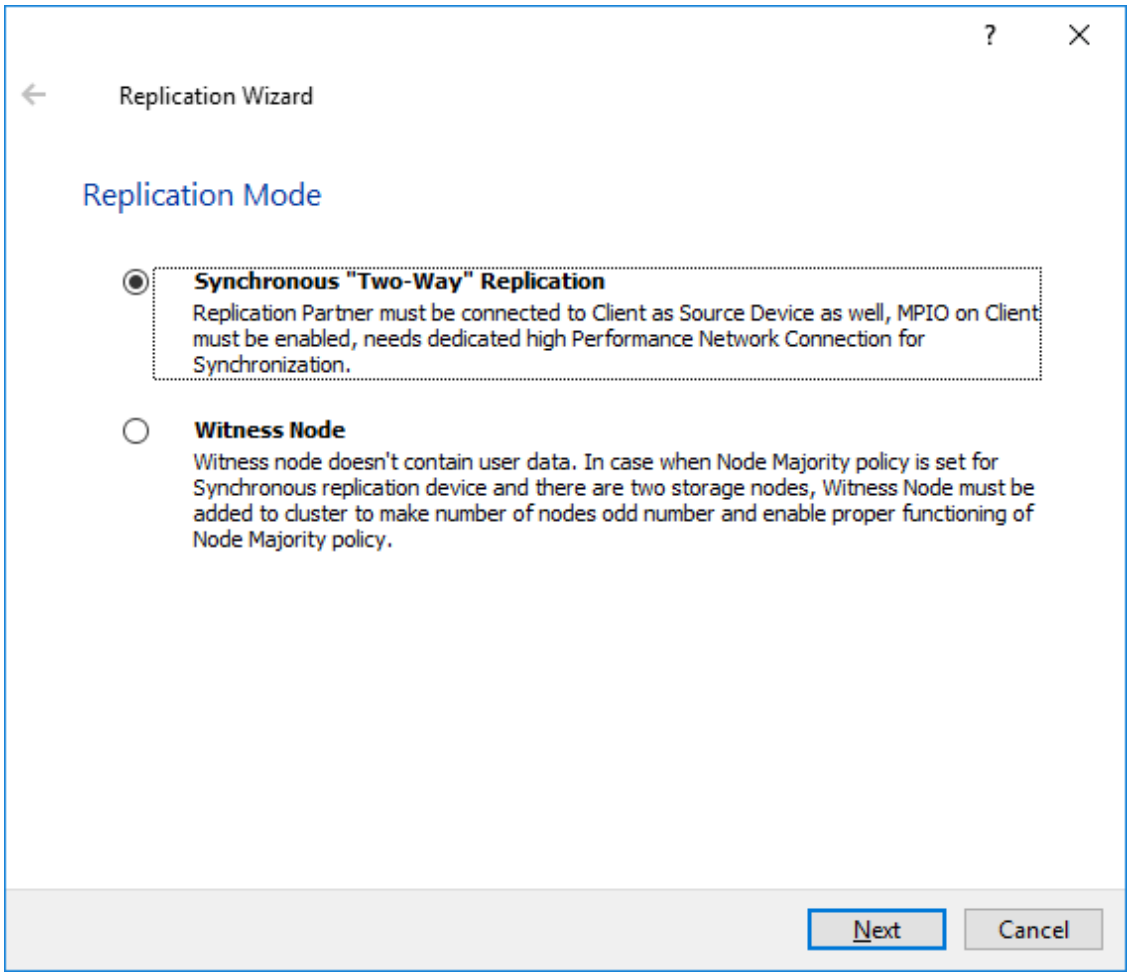

4. Specify a partner Host name or IP address and Port Number.

Selecting The Failover Strategy

StarWind provides 2 options for configuring a failover strategy:

Heartbeat

The Heartbeat failover strategy allows avoiding the “split-brain” scenario when the HA cluster nodes are unable to synchronize but continue to accept write commands from the initiators independently. It can occur when all synchronization and heartbeat channels disconnect simultaneously, and the partner nodes do not respond to the node’s requests. As a result, StarWind service assumes the partner nodes to be offline and continues operations on a single-node mode using data written to it.
If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The device with the lowest priority will be marked as not synchronized and get subsequently blocked for the further read and write operations until the synchronization channel resumption. At the same time, the partner device on the

synchronized node flushes data from the cache to the disk to preserve data integrity in case the node goes down unexpectedly. It is recommended to assign more independent heartbeat channels during the replica creation to improve system stability and avoid the “split-brain” issue.
With the heartbeat failover strategy, the storage cluster will continue working with only one StarWind node available.

Node Majority

The Node Majority failover strategy ensures the synchronization connection without any additional heartbeat links. The failure-handling process occurs when the node has detected the absence of the connection with the partner.
The main requirement for keeping the node operational is an active connection with more than half of the HA device’s nodes. Calculation of the available partners is based on their “votes”.
In case of a two-node HA storage, all nodes will be disconnected if there is a problem on the node itself, or in communication between them. Therefore, the Node Majority failover strategy requires the addition of the third Witness node or file share (SMB) which participates in the nodes count for the majority, but neither contains data on it nor is involved in processing clients’ requests. In case an HA device is replicated between 3 nodes, no Witness node is required.
With Node Majority failover strategy, failure of only one node can be tolerated. If two nodes fail, the third node will also become unavailable to clients’ requests.
Please select the required option:

Heartbeat

1. Select Failover Strategy.

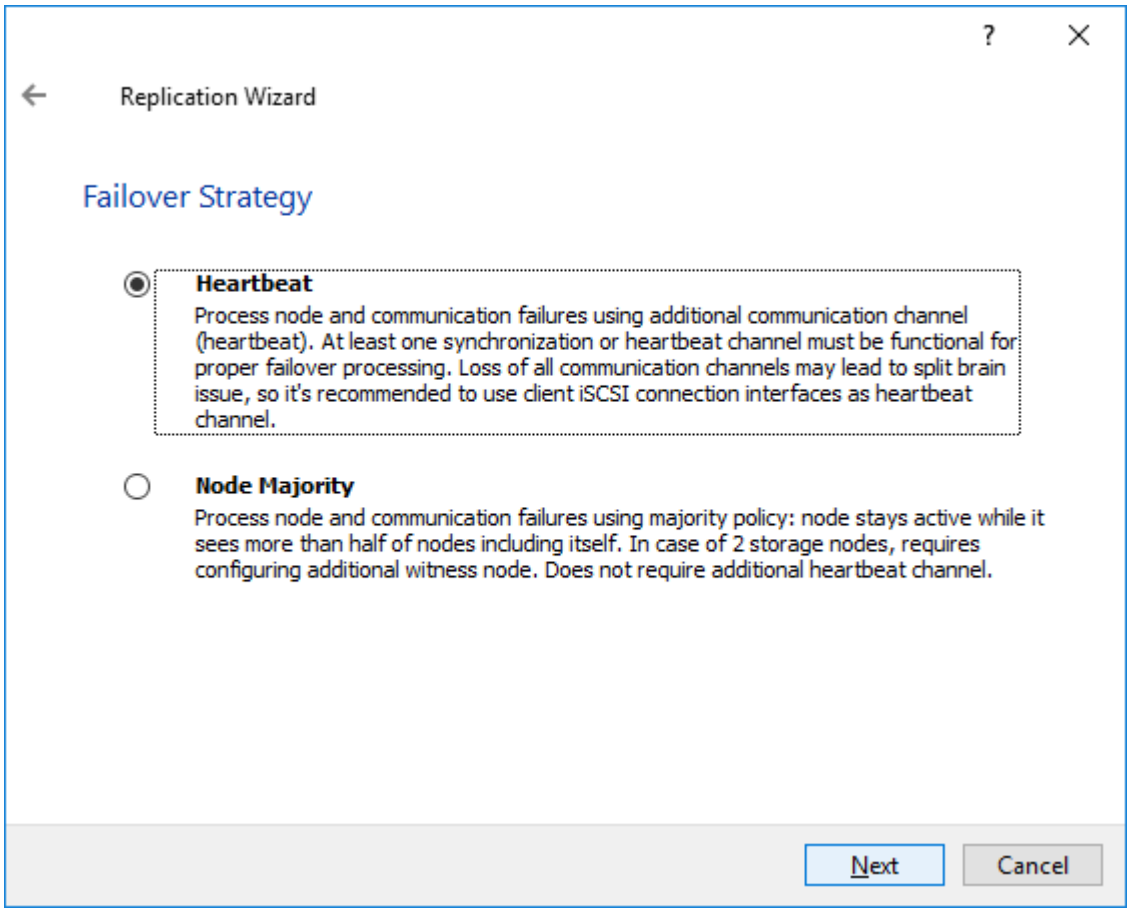

2. Select Create new Partner Device and click Next.

3. Select a partner device Location and click Next.

4. Select Synchronization Journal Strategy and click Next.
NOTE: There are several options – RAM-based journal (default) and Disk-based journal with failure and continuous strategy, that allow to avoid full synchronization cases.

RAM-based (default) synchronization journal is placed in RAM. Synchronization with RAM journal provides good I/O performance in any scenario. Full synchronization could occur in the cases described in this KB:
https://knowledgebase.starwindsoftware.com/explanation/reasons-why-full-synchronization-may-start/

Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes. Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are recommended as the device performance is defined by the disk speed, where the journal is located. For example, it can be placed on the OS boot volume.

It is required to allocate 2 MB of disk space for the synchronization journal per 1 TB of HA device size with a disk-based journal configured with 2-way replication and 4MB per 1 TB of HA device size for 3-way replication.

Failure journal
The strategy provides good I/O performance, as a RAM-based journal, while all device nodes are in a healthy synchronized state. If a device on one node went into a not synchronized state, the disk-based journal activates and a performance drop could occur as the device performance is defined by the disk speed, where the journal is located. Fast synchronization is not guaranteed in all cases. For example, if a simultaneous hard reset of all nodes occurs, full synchronization will occur.

Continuous journal
The strategy guarantees fast synchronization and data consistency in all cases. Although, this strategy has the worst I/O performance, because of frequent write operations to the journal, located on the disk, where the journal is located.

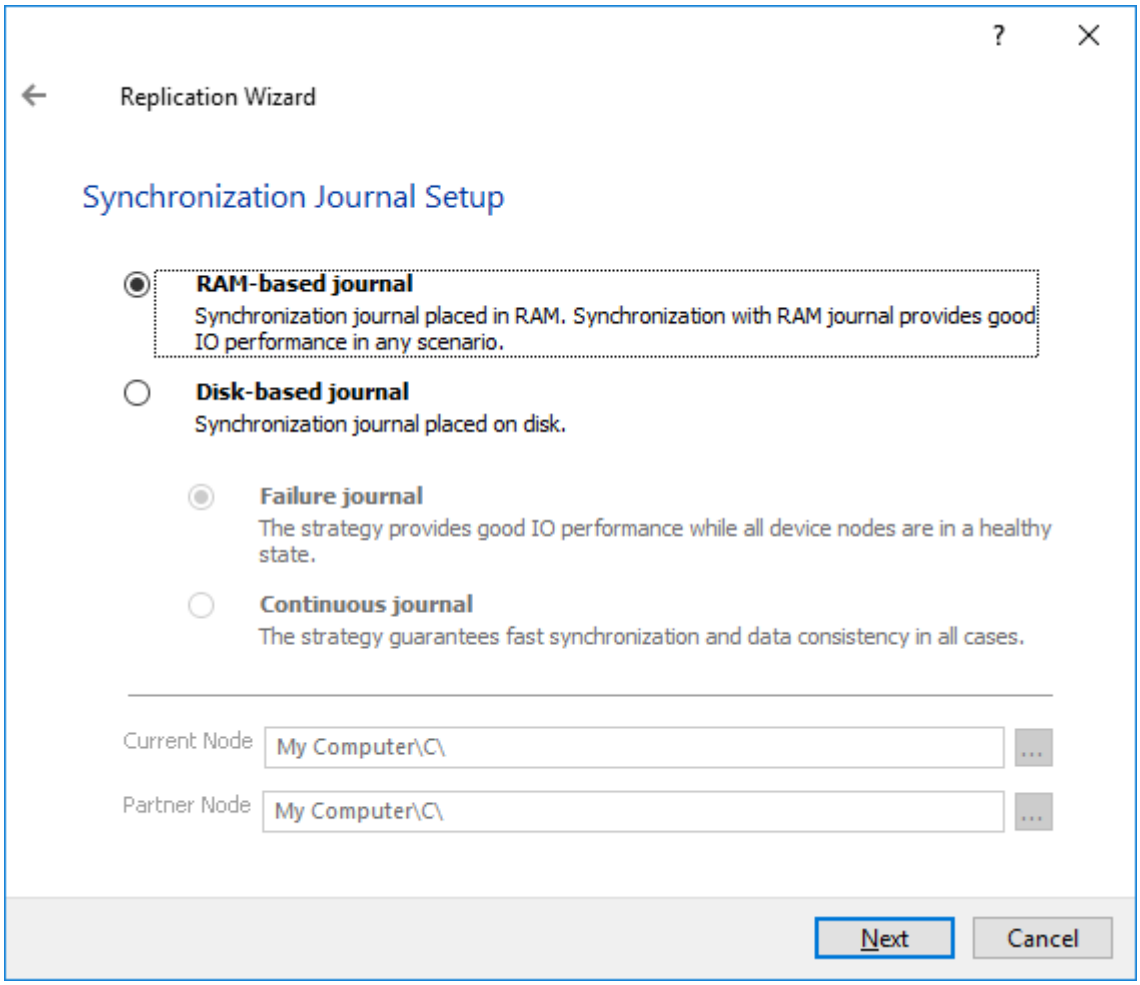

5. Click Change Network Settings and specify the interfaces for Synchronization and

Heartbeat Channels. Click OK and then click Next.

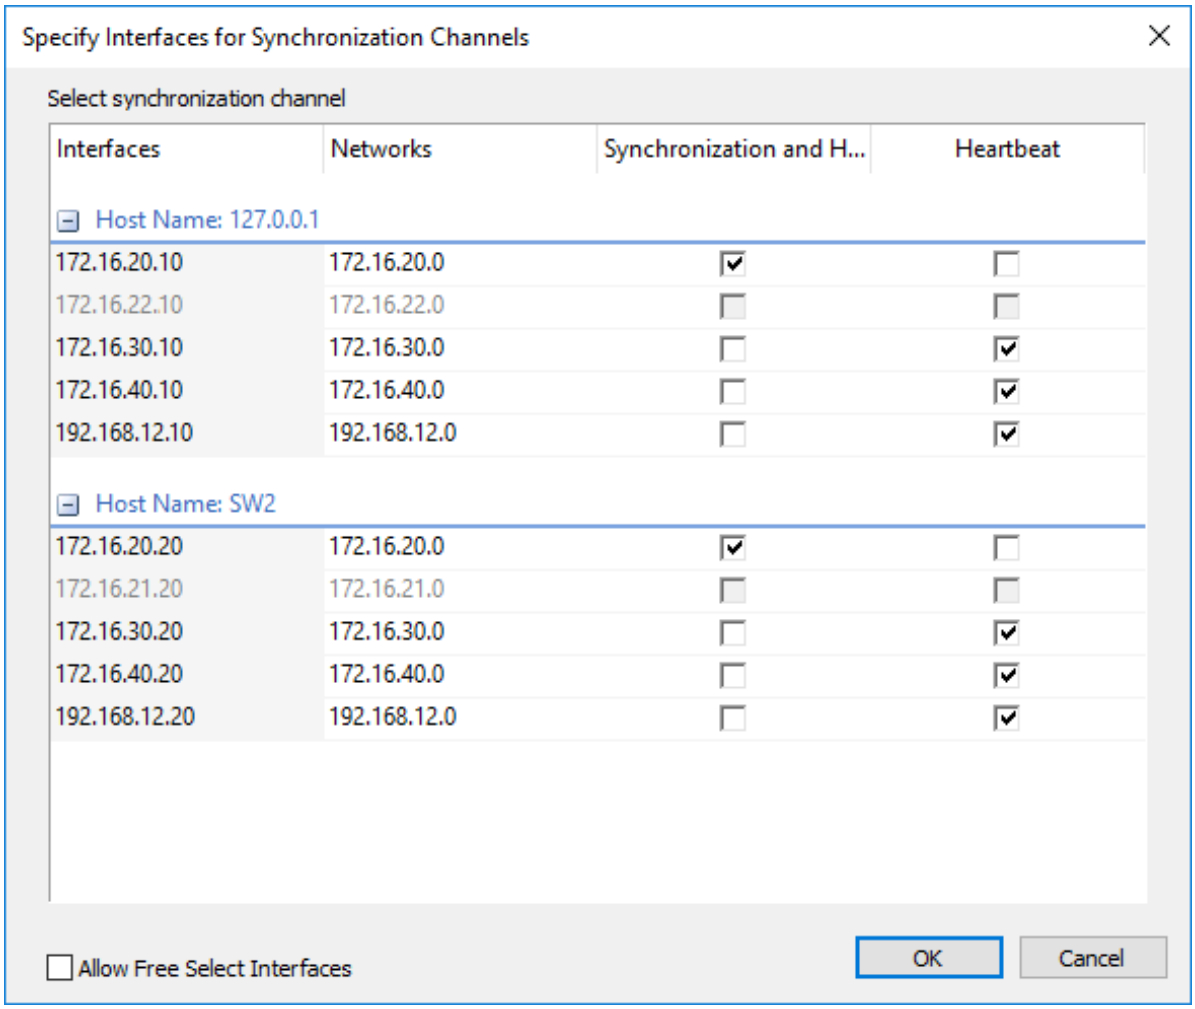

6. In Select Partner Device Initialization Mode, select Synchronize from existing Device and click Next.

7. Click Create Replica. Click Finish to close the wizard.

8. The successfully added device appears in StarWind Management Console.

9. Choose device, open Replication Manager and click Add replica again.

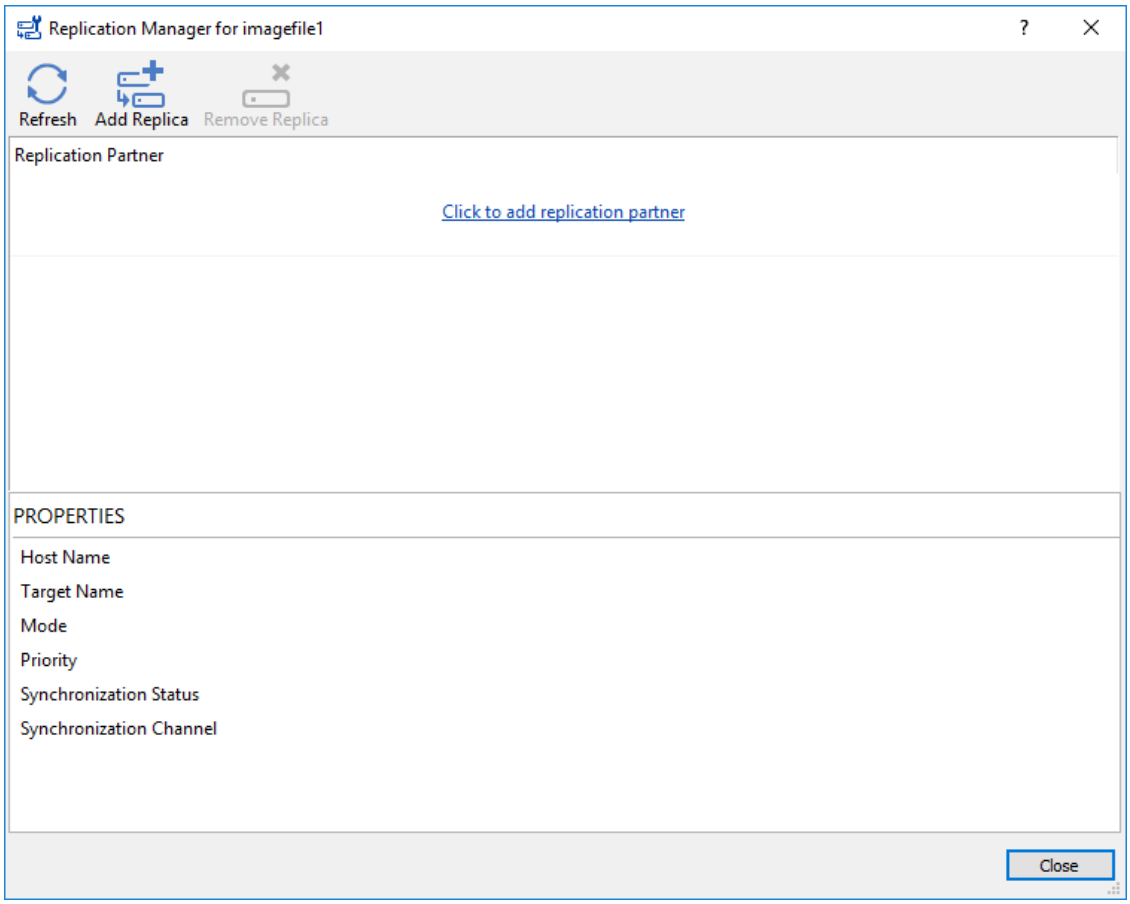

10. Select Synchronous “Two-Way” Replication as a replication mode. Click Next to proceed.

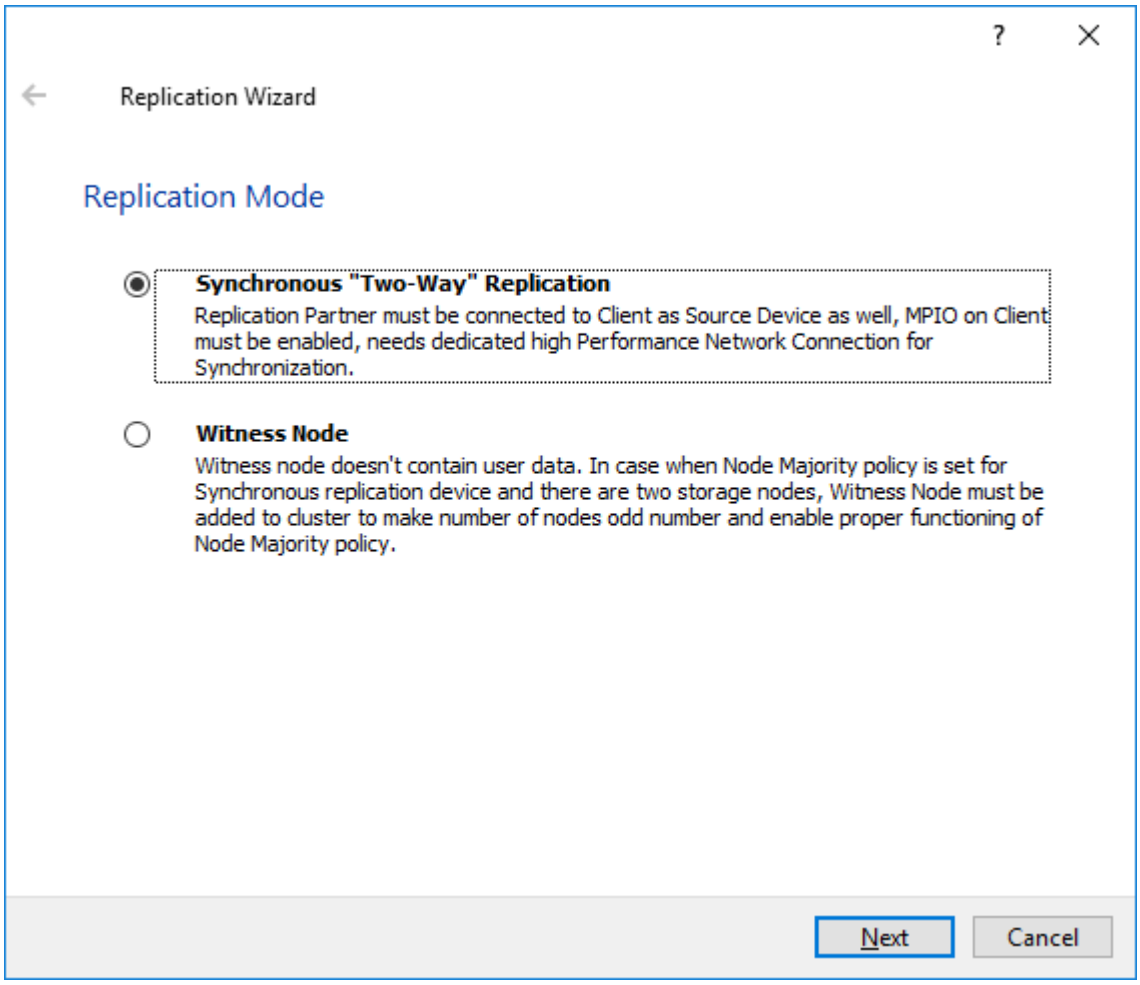

11. Specify a partner Host name or IP address and Port Number.

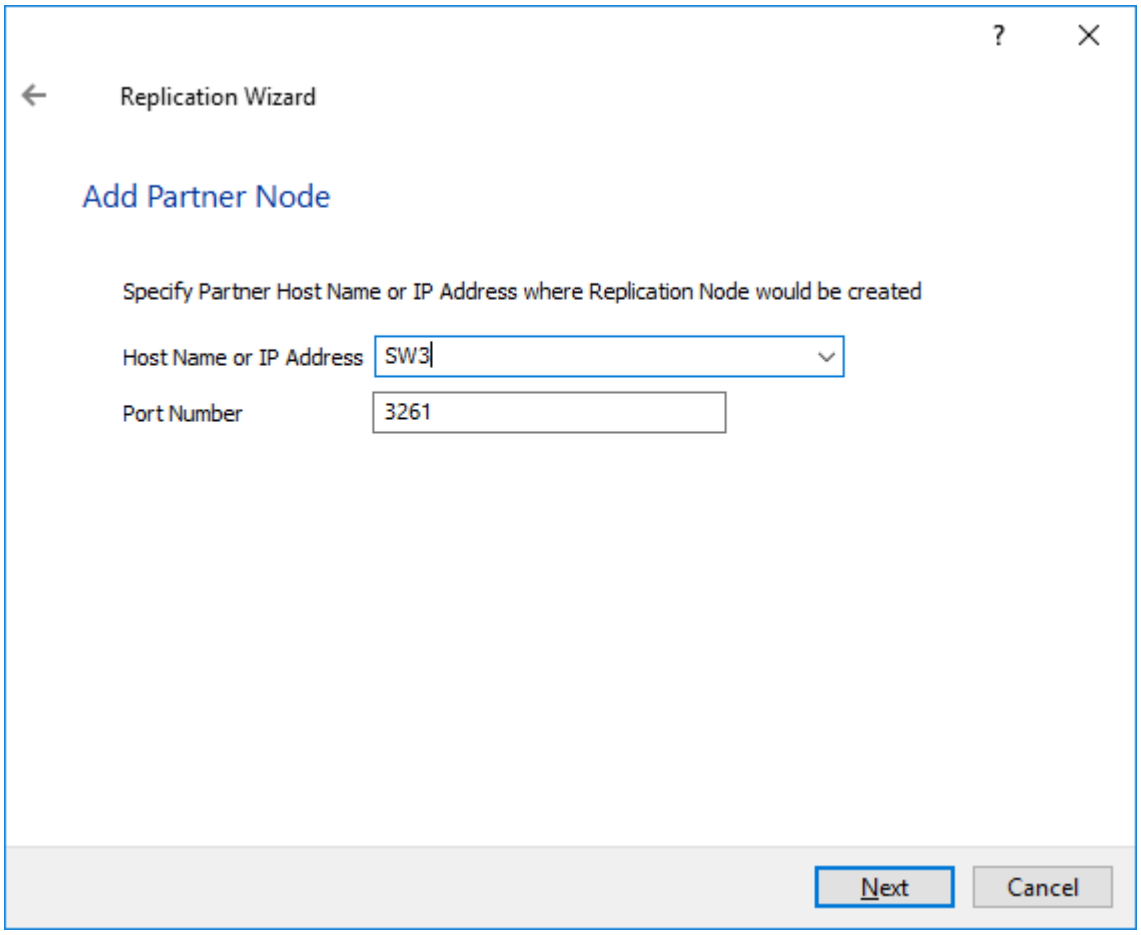

12. Select Failover Strategy.

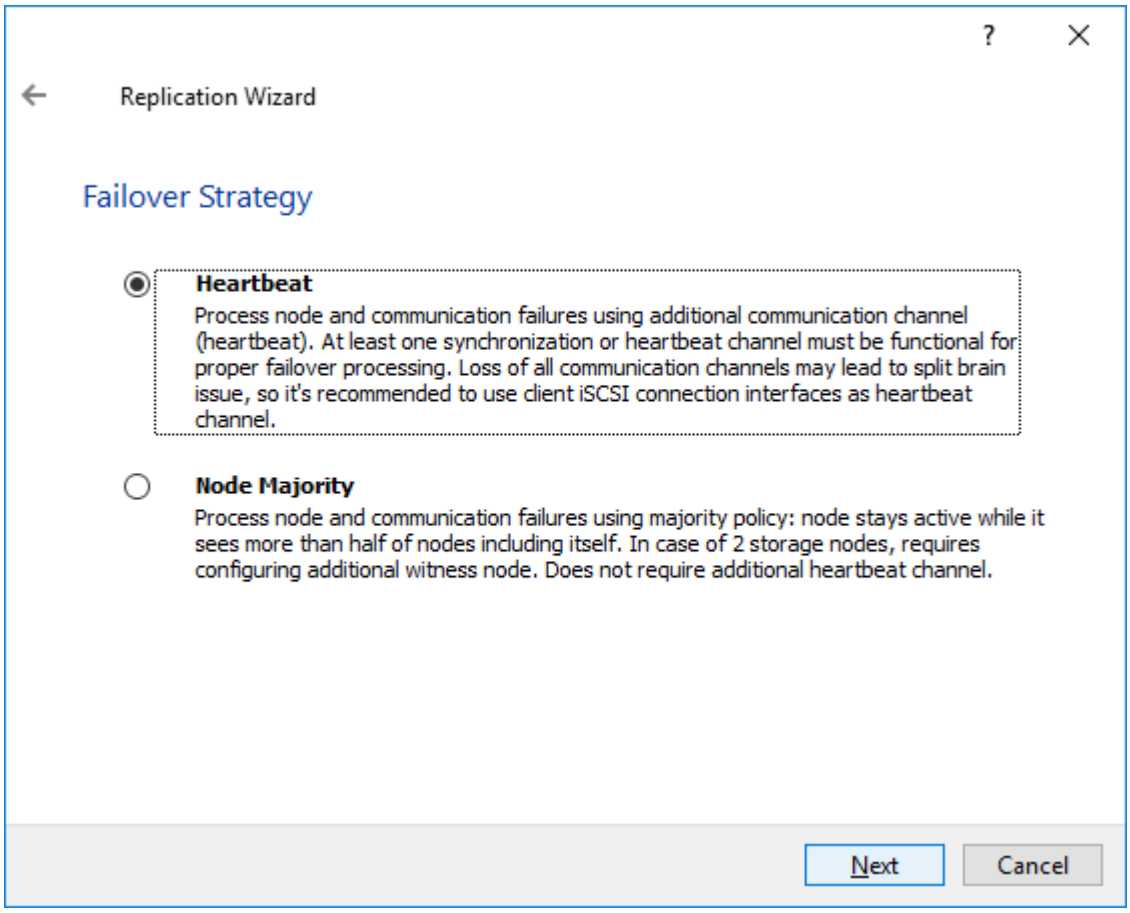

13. Select Create new Partner Device and click Next.

14. Select a partner device Location and Synchronization Journal Strategy and click Next.

15. Click Change Network Settings.

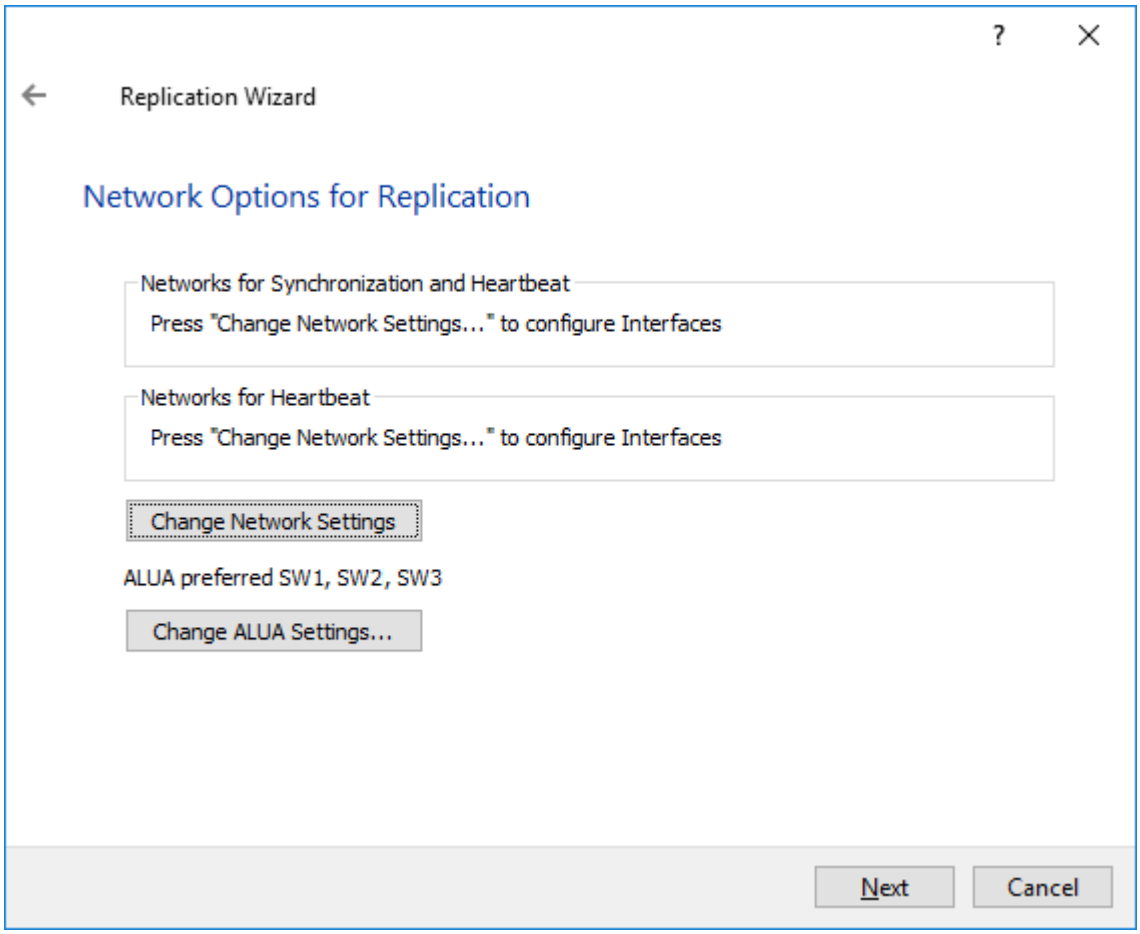

16. Specify the interfaces for Synchronization and Heartbeat Channels. Click OK and then click Next.

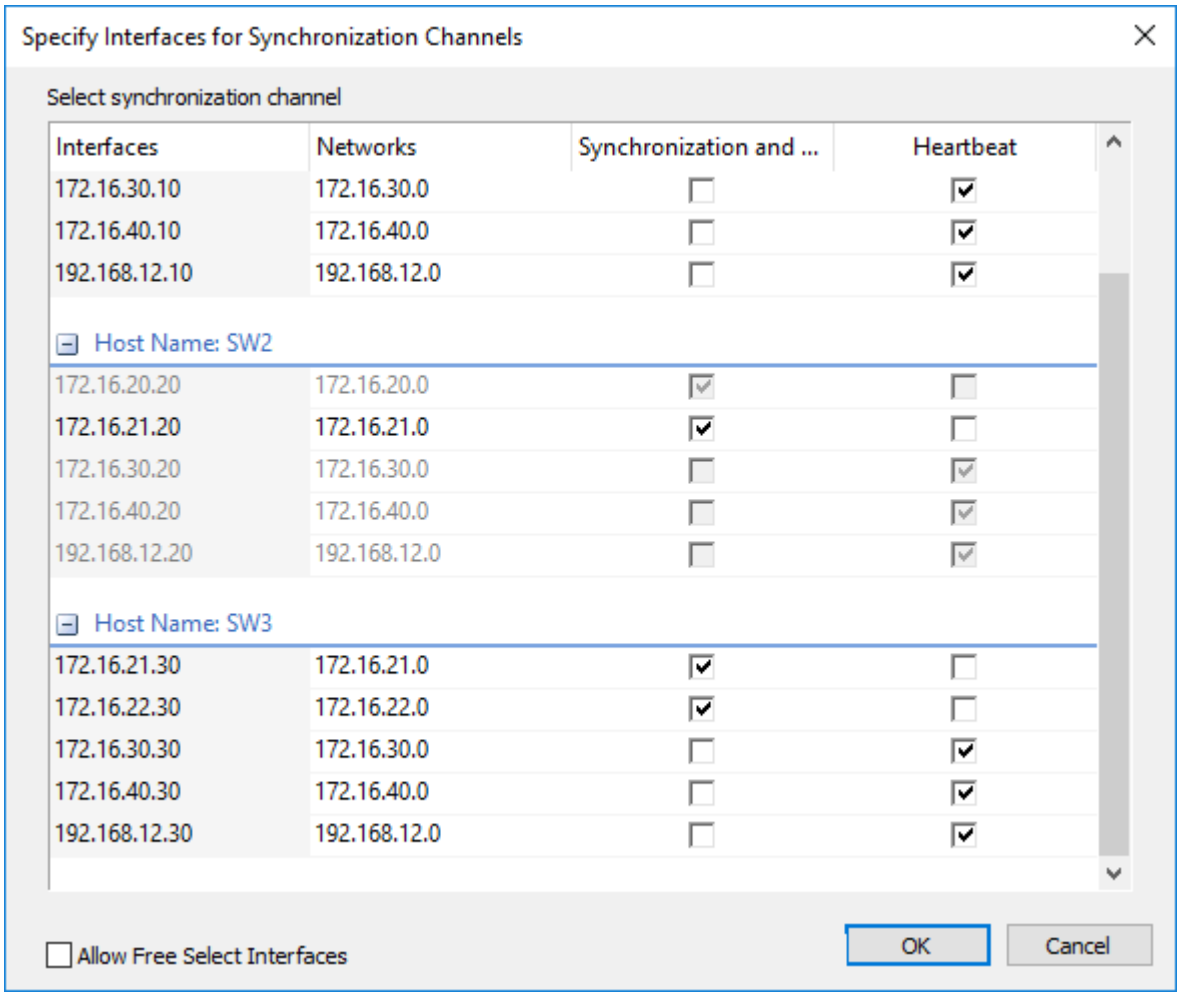

NOTE: It is not recommended to configure the Heartbeat and iSCSI channels on the same interfaces to avoid the split-brain issue. If the Synchronization and Heartbeat interfaces are located on the same network adapter, it is recommended to assign one more Heartbeat interface to a separate adapter.

17. In Select Partner Device Initialization Mode, select Synchronize from existing Device and click Next.

18. Click Create Replica. Click Finish to close the wizard.
The successfully added device appears in StarWind Management Console.

19. Follow the similar procedure for the creation of other virtual disks that will be used as storage repositories.

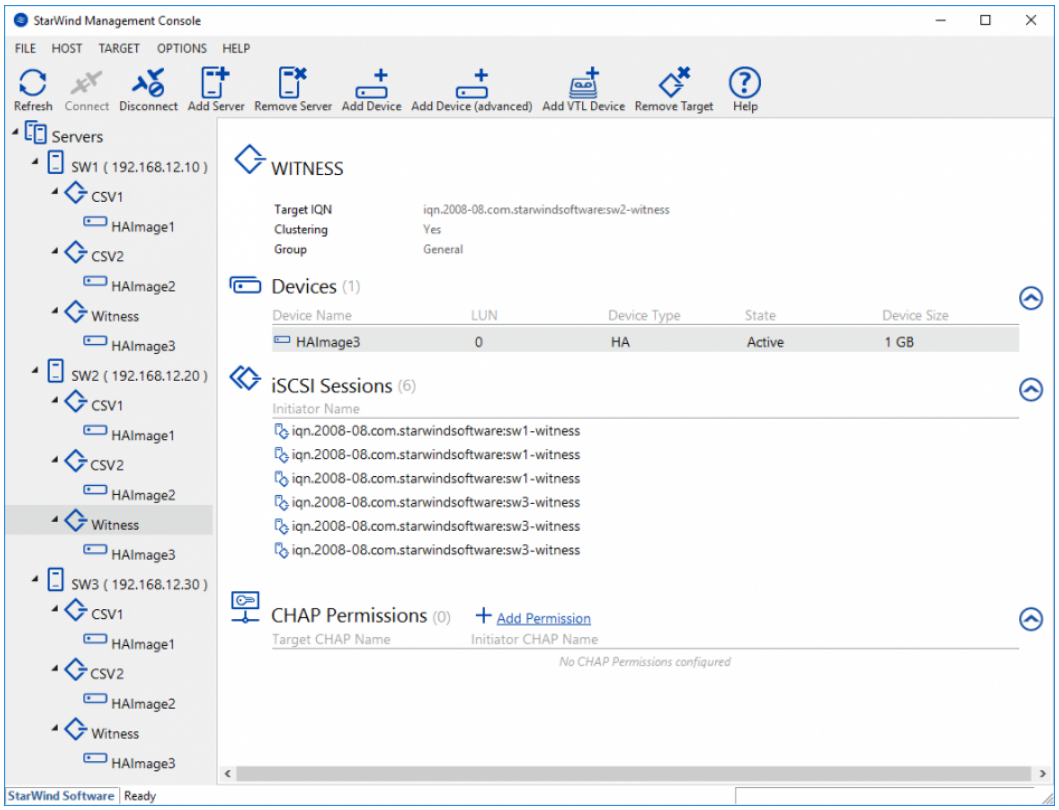

NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:
https://knowledgebase.starwindsoftware.com/maintenance/how-to-extend-image-file-or-high-availability-device/

Node Majority

1. Select the Node Majority failover strategy and click Next.

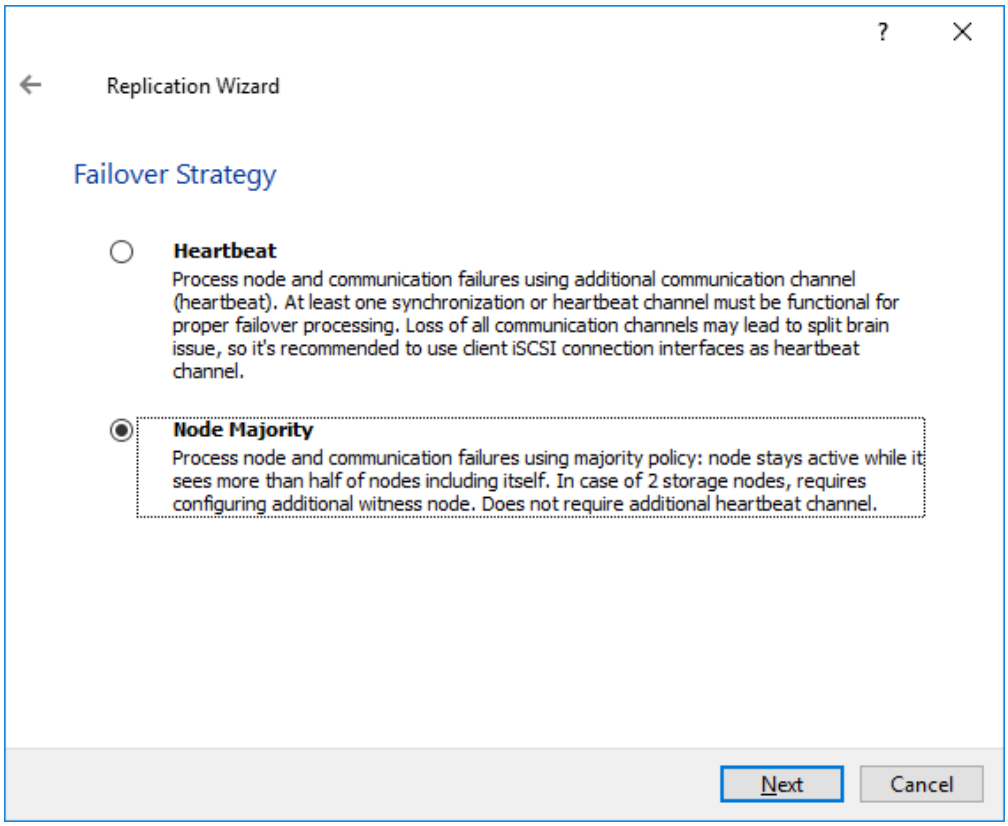

2. Choose Create new Partner Device and click Next.

3. Specify the partner device Location and modify the target name if necessary.
Click Next.

4. Select Synchronization Journal Strategy and click Next.
NOTE: There are several options – RAM-based journal (default) and Disk-based journal with failure and continuous strategy, that allow to avoid full synchronization cases.

RAM-based (default) synchronization journal is placed in RAM. Synchronization with RAM journal provides good I/O performance in any scenario. Full synchronization could occur in the cases described in this KB:
https://knowledgebase.starwindsoftware.com/explanation/reasons-why-full-synchronization-may-start/

Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes. Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are recommended as the device performance is defined by the disk speed, where the journal is located. For example, it can be placed on the OS boot volume.

It is required to allocate 2 MB of disk space for the synchronization journal per 1 TB of HA device size with a disk-based journal configured with 2-way replication and 4MB per 1 TB of HA device size for 3-way replication.

Failure journal
The strategy provides good I/O performance, as a RAM-based journal, while all device nodes are in a healthy synchronized state. If a device on one node went into a not synchronized state, the disk-based journal activates and a performance drop could occur as the device performance is defined by the disk speed, where the journal is located. Fast synchronization is not guaranteed in all cases. For example, if a simultaneous hard reset of all nodes occurs, full synchronization will occur.

Continuous journal
The strategy guarantees fast synchronization and data consistency in all cases. Although, this strategy has the worst I/O performance, because of frequent write operations to the journal, located on the disk, where the journal is located.

5. In Network Options for Replication, press the Change network settings button and select the synchronization channel for the HA device.

6. In Specify Interfaces for Synchronization Channels, select the checkboxes with the appropriate networks and click OK. Then click Next.

7. Select Synchronize from existing Device as the partner device initialization mode.

8. Press the Create Replica button and close the wizard.

9. The added devices will appear in StarWind Management Console.

10. Choose device, open Replication Manager and click Add replica again.

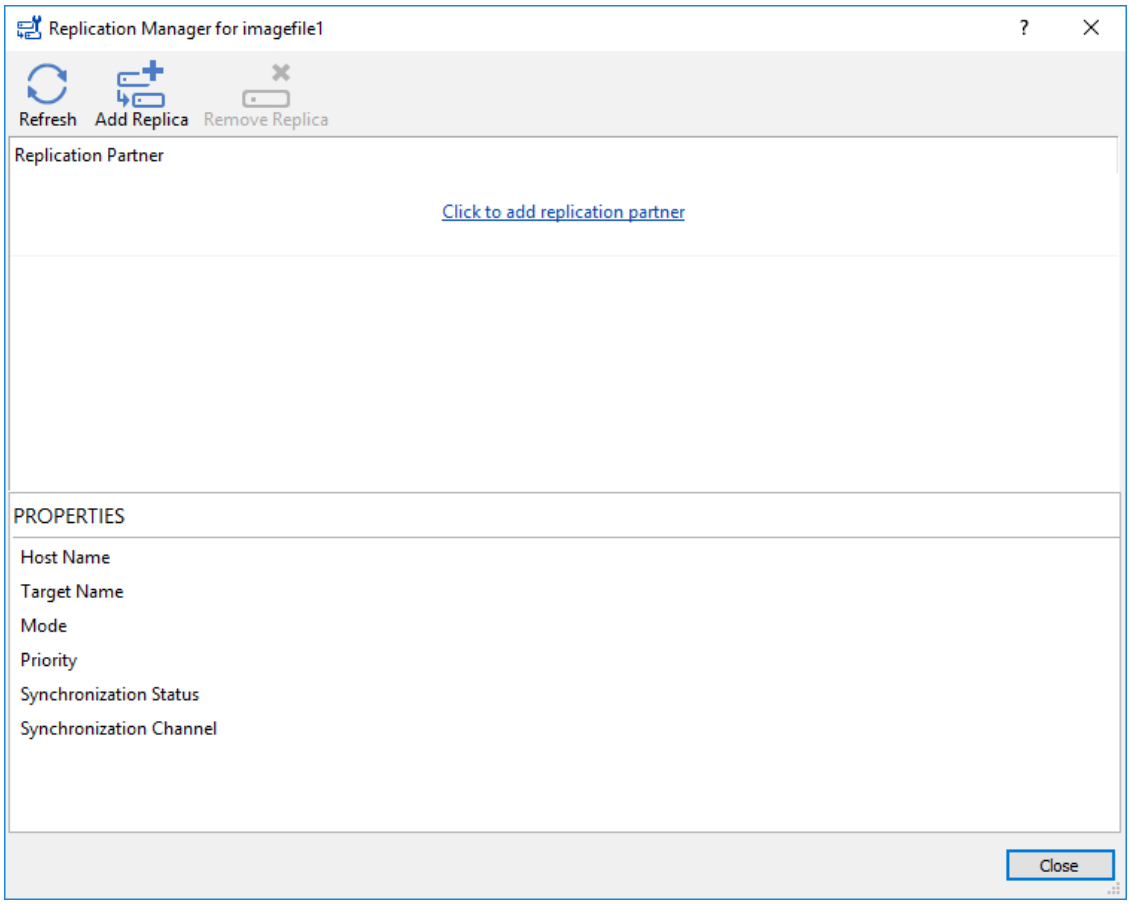

11. Select Synchronous “Two-Way” Replication as a replication mode. Click Next to proceed.

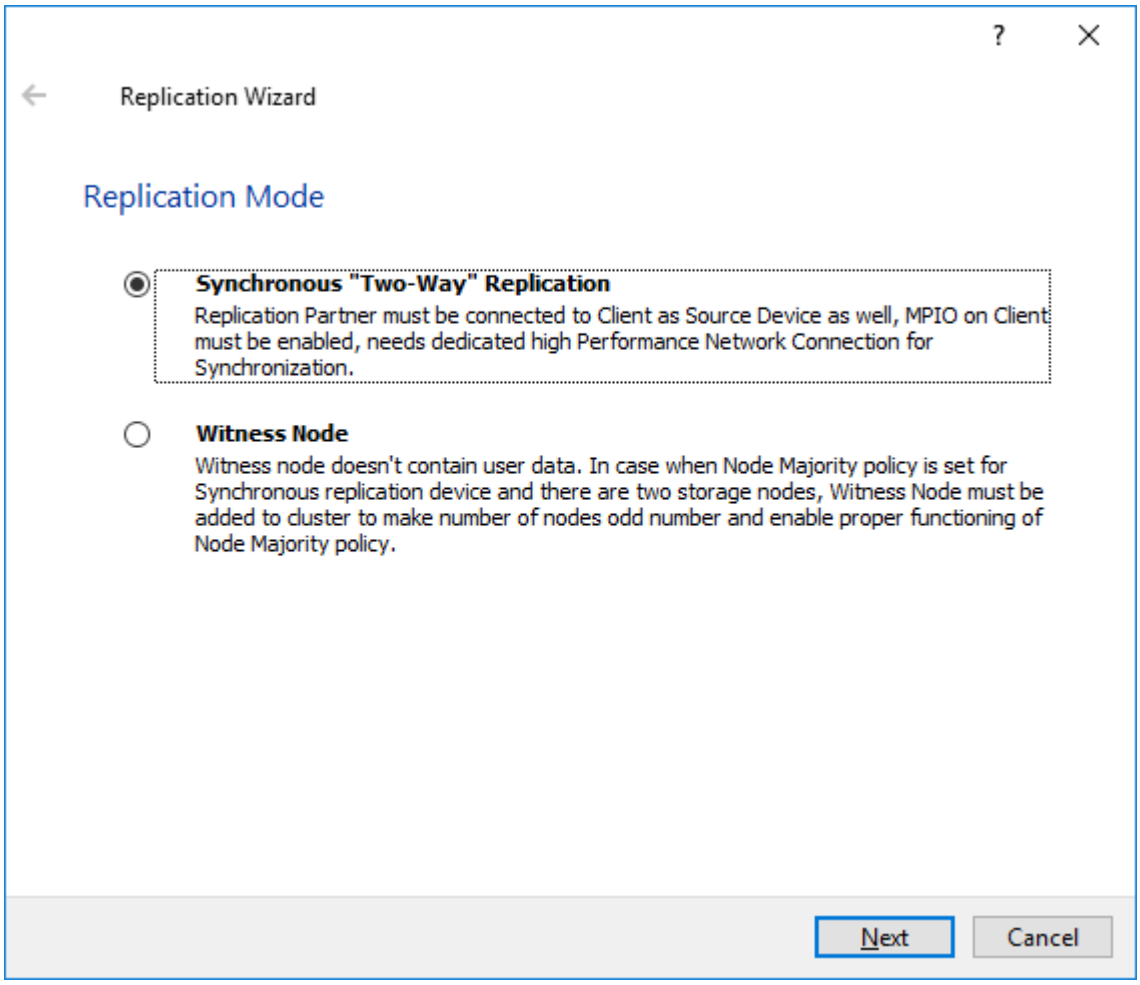

12. Specify a partner Host name or IP address and Port Number.

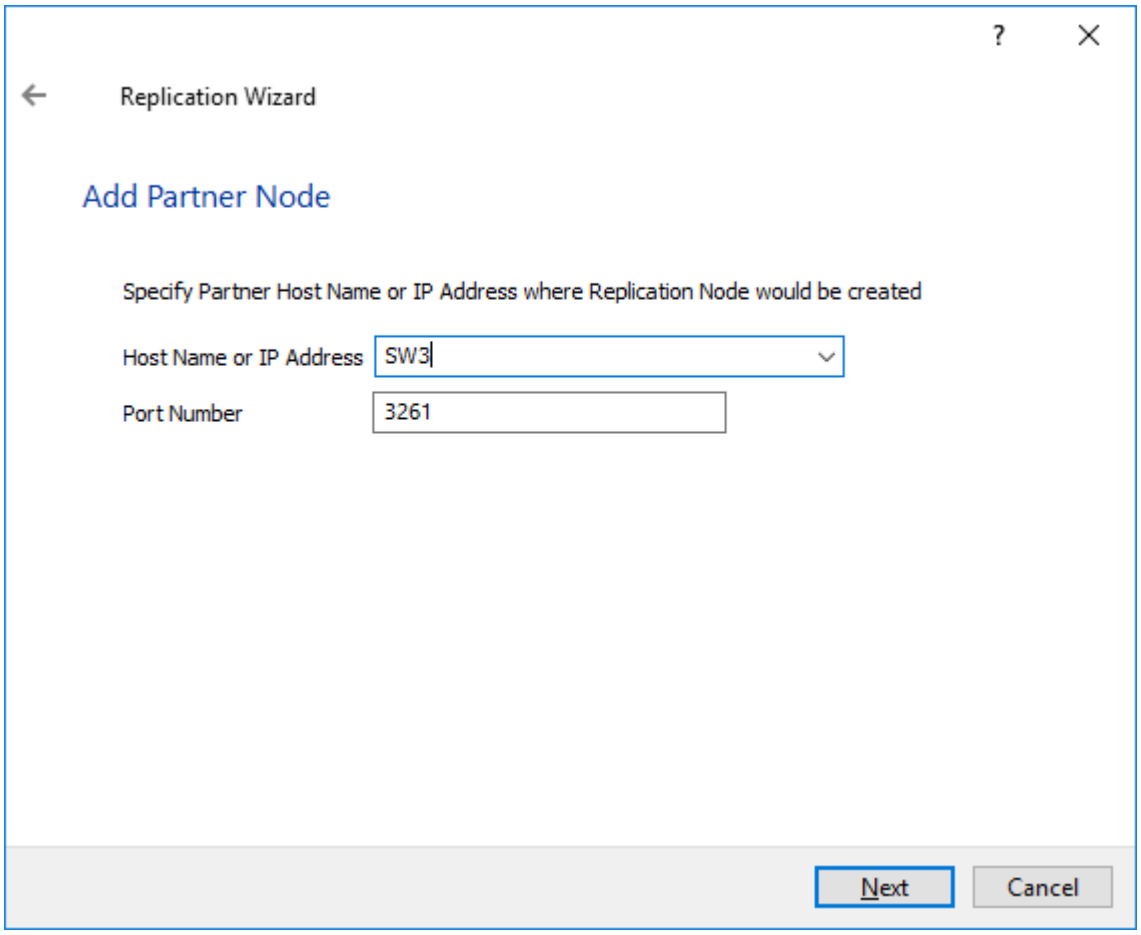

13. Select the Node Majority failover strategy and click Next.

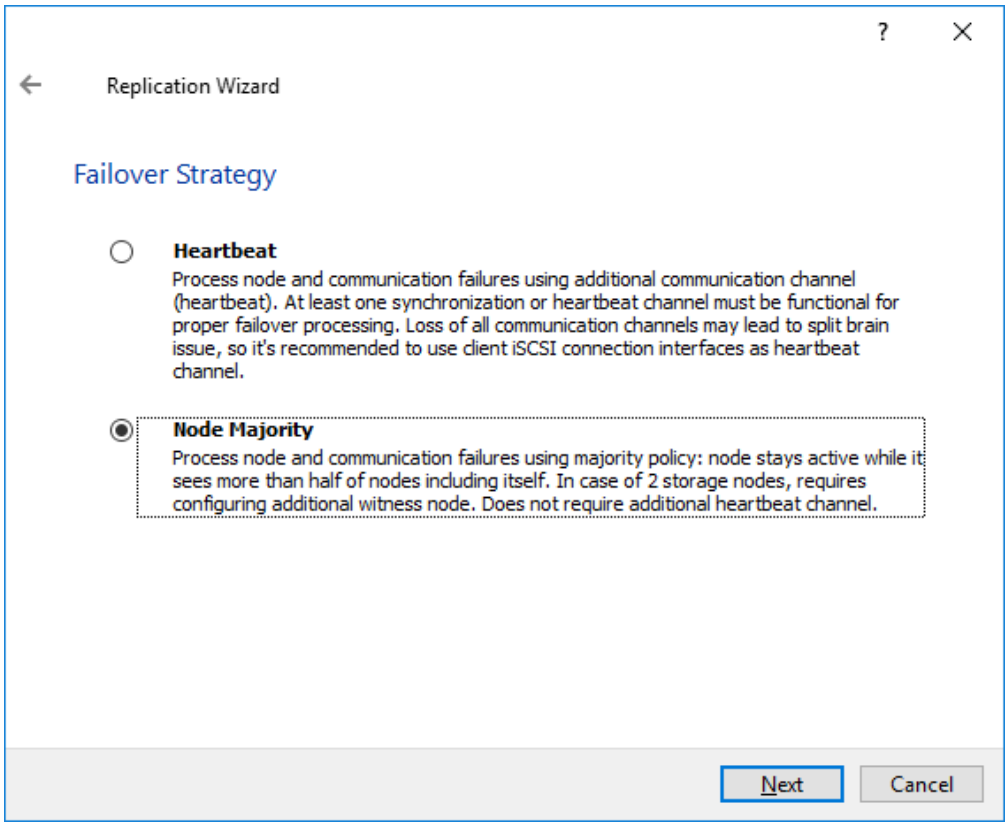

14. Choose Create new Partner Device and click Next.

15. Specify the partner device Location and modify the target name if necessary. Click Next.

16. Select Synchronization Journal Strategy and click Next.

17. In Network Options for Replication, press the Change network settings button and select the synchronization channel for the HA device.

18. In Specify Interfaces for Synchronization Channels, select the checkboxes with the appropriate networks and click OK. Then click Next.

19. Select Synchronize from existing Device as the partner device initialization mode.

20. Press the Create Replica button and close the wizard.

21. The added devices will appear in StarWind Management Console.

Repeat the steps above to create other virtual disks if necessary.
NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:

https://knowledgebase.starwindsoftware.com/maintenance/how-to-extend-image-file-or-high-availability-device/

Provisioning Starwind Ha Storage To Windows Server Hosts

1. Launch Microsoft iSCSI Initiator by executing the following command in the CMD window:

```
iscsicpl
```

2. Navigate to the Discovery tab.

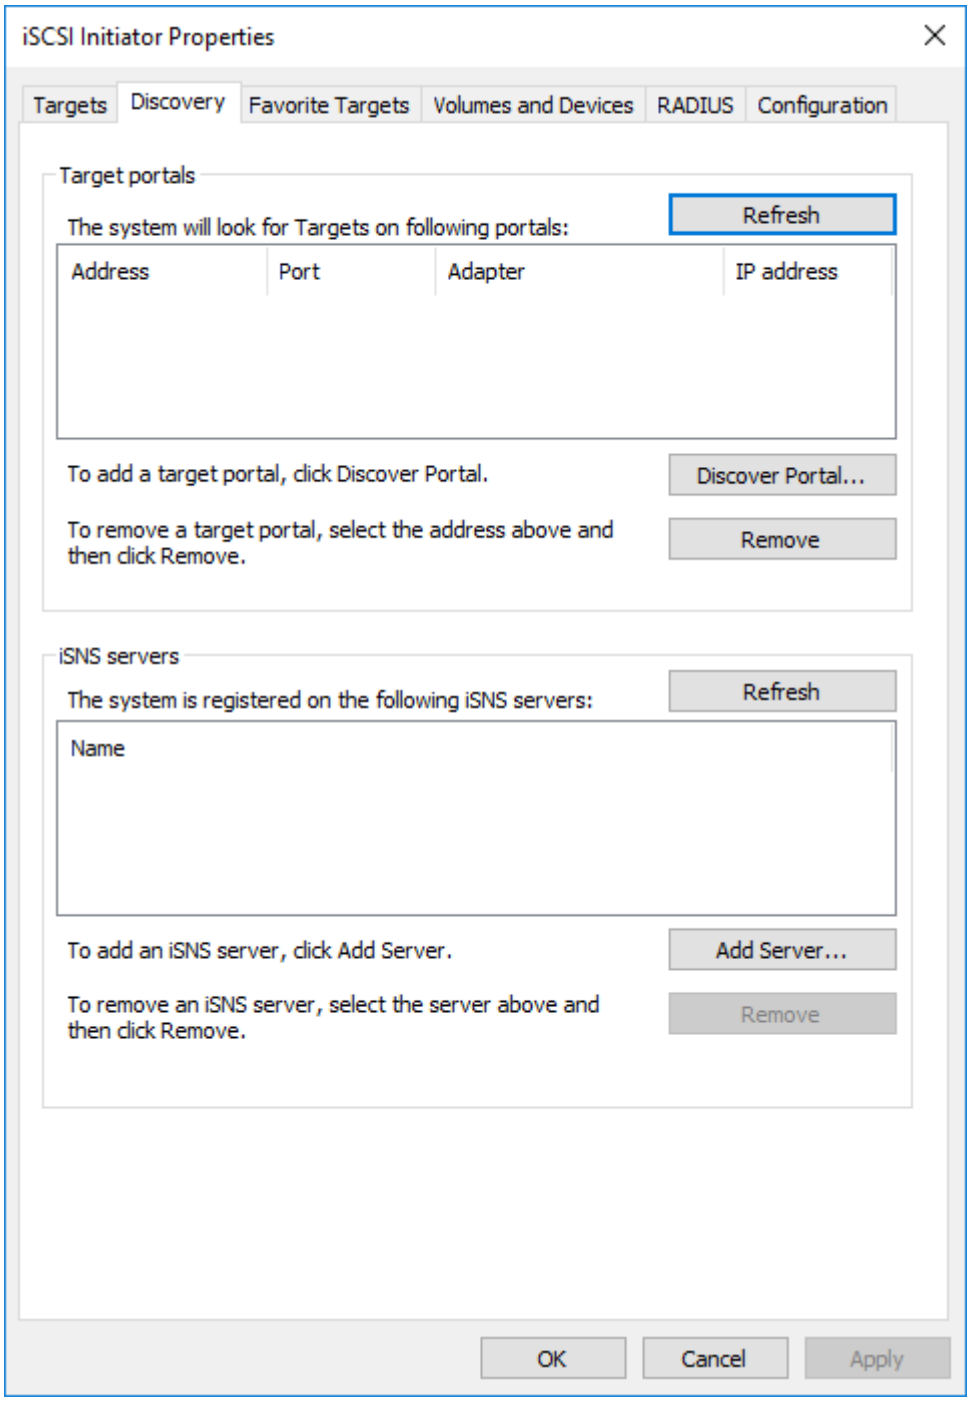

3. Click the Discover Portal button. In the Discover Target Portal dialog, type in the iSCSI interface IP address of the first StarWind node that will be used to connect the StarWind provisioned targets. The steps below provide instructions how to discover targets within 172.16.30.X subnet. The same should be done for 172.16.40.X subnet
Click Advanced...

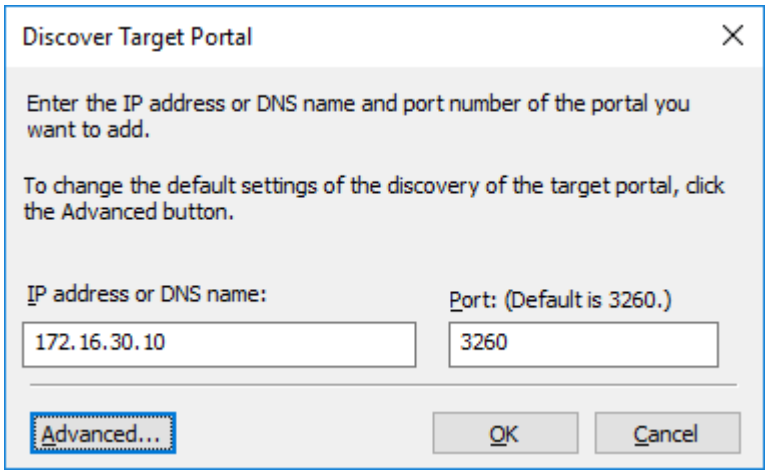

4. Select Microsoft iSCSI Initiator as the Local adapter, select the Initiator IP in the same subnet as the IP address of the partner node from the previous step. Confirm the actions to complete the Target Portal discovery.

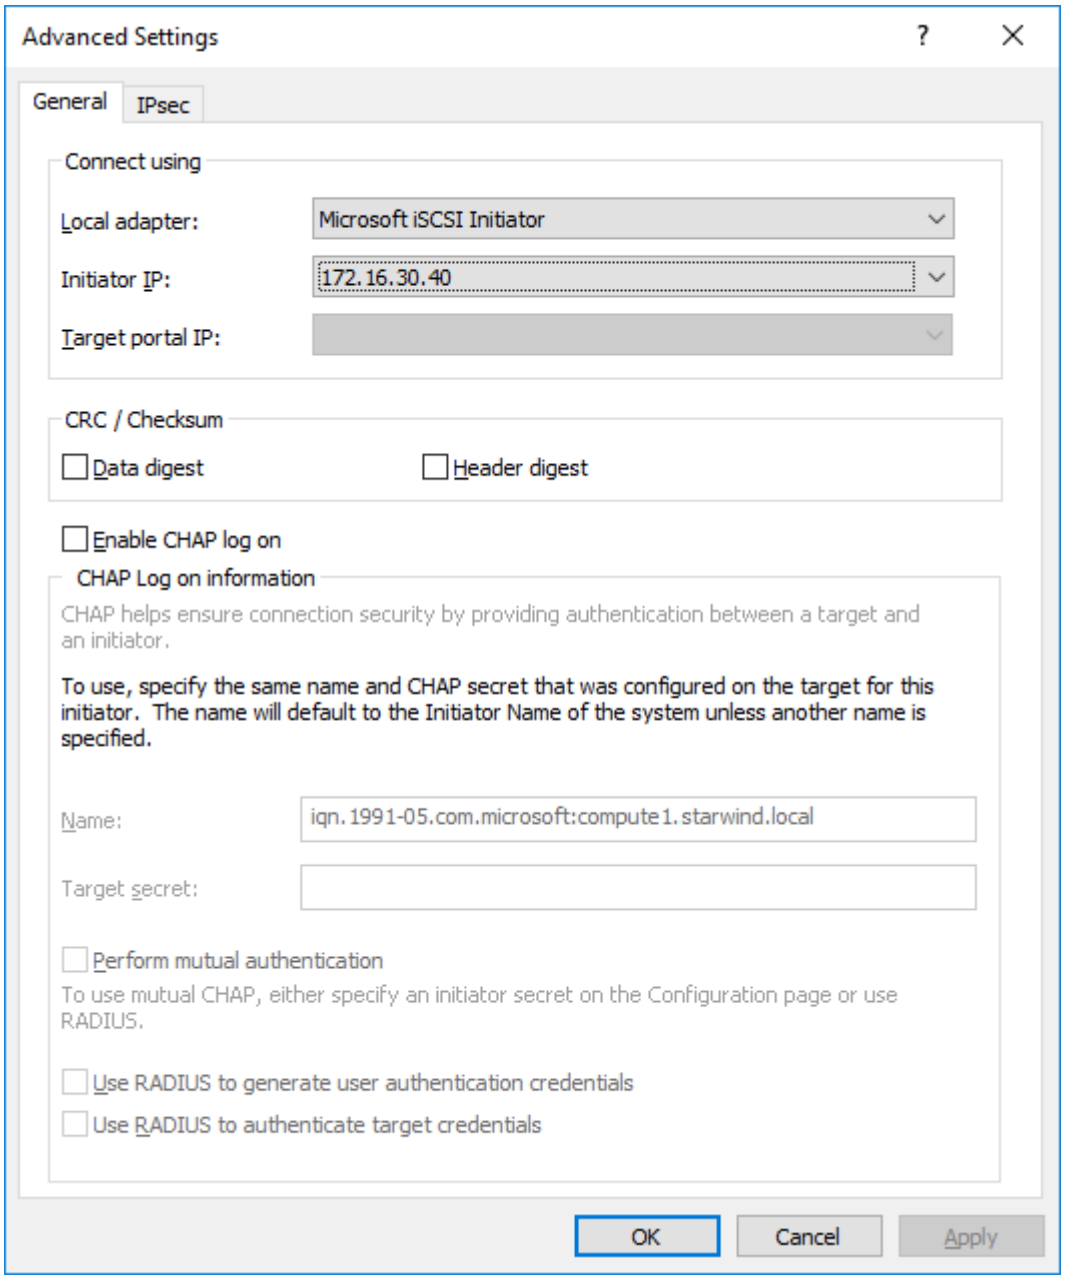

5. Click the Discover Portal button. In the Discover Target Portal dialog, type in the iSCSI interface IP address of the second StarWind node that will be used to connect the StarWind provisioned targets. Click Advanced...

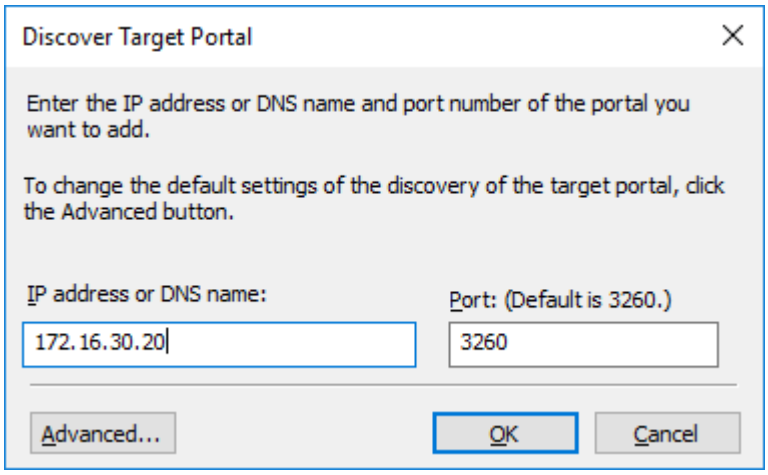

6. Select Microsoft iSCSI Initiator as the Local adapter, select the Initiator IP in the same subnet as the IP address of the partner node from the previous step. Confirm the actions to complete the Target Portal discovery.

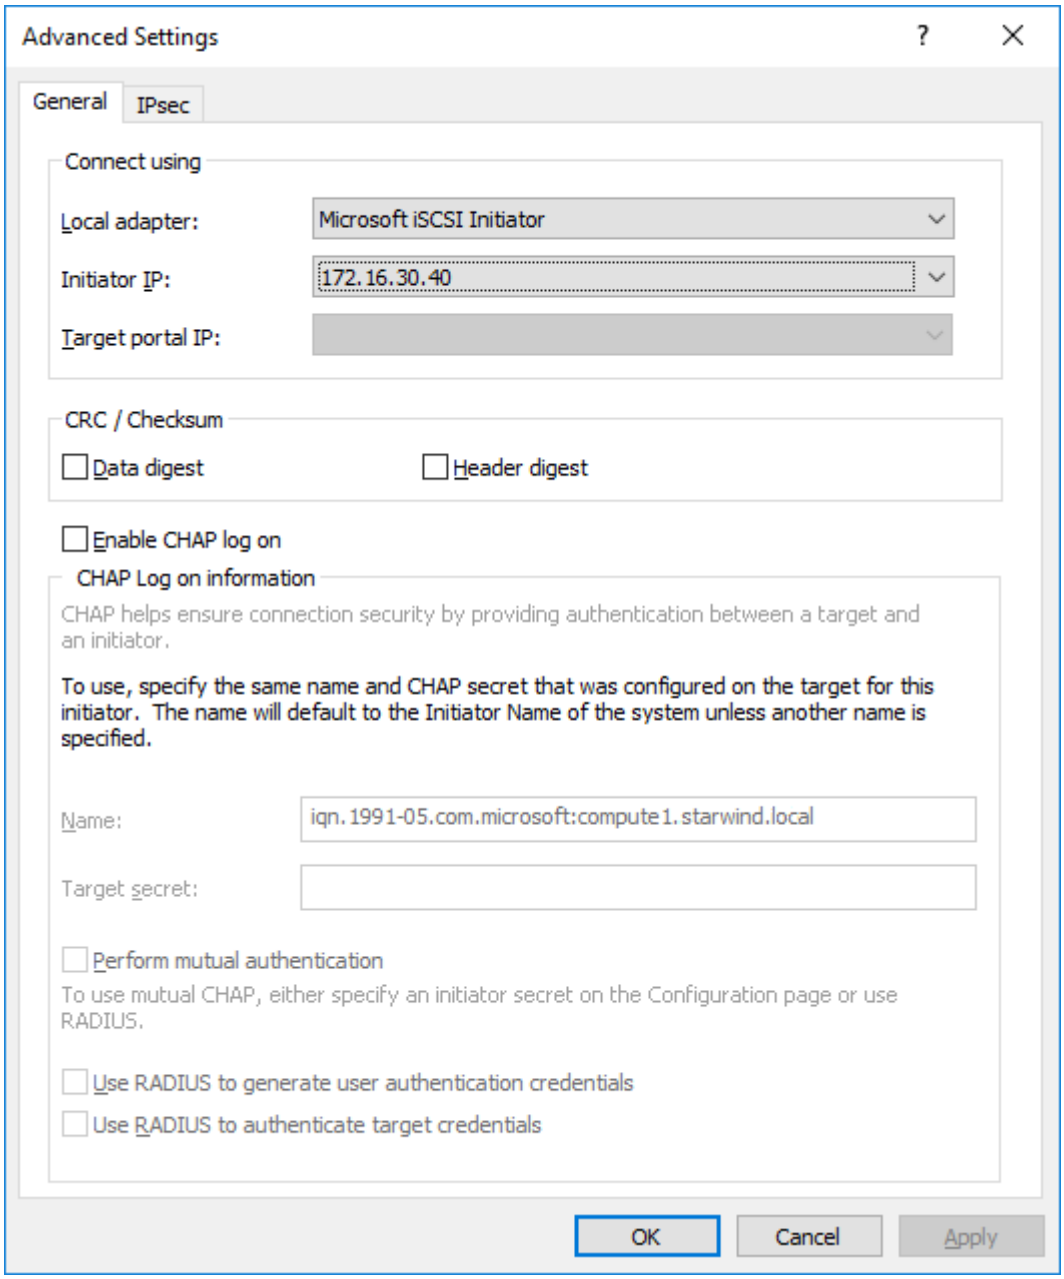

7. Click the Discover Portal button. In the Discover Target Portal dialog, type in the iSCSI interface IP address of the third StarWind node that will be used to connect the StarWind provisioned targets. Click Advanced...

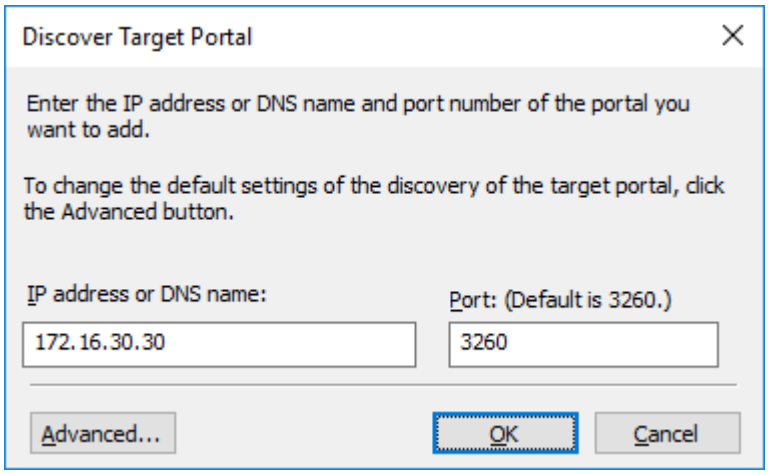

8. Select Microsoft iSCSI Initiator as the Local adapter, select the Initiator IP in the same subnet as the IP address of the partner node from the previous step. Confirm the actions to complete the Target Portal discovery.

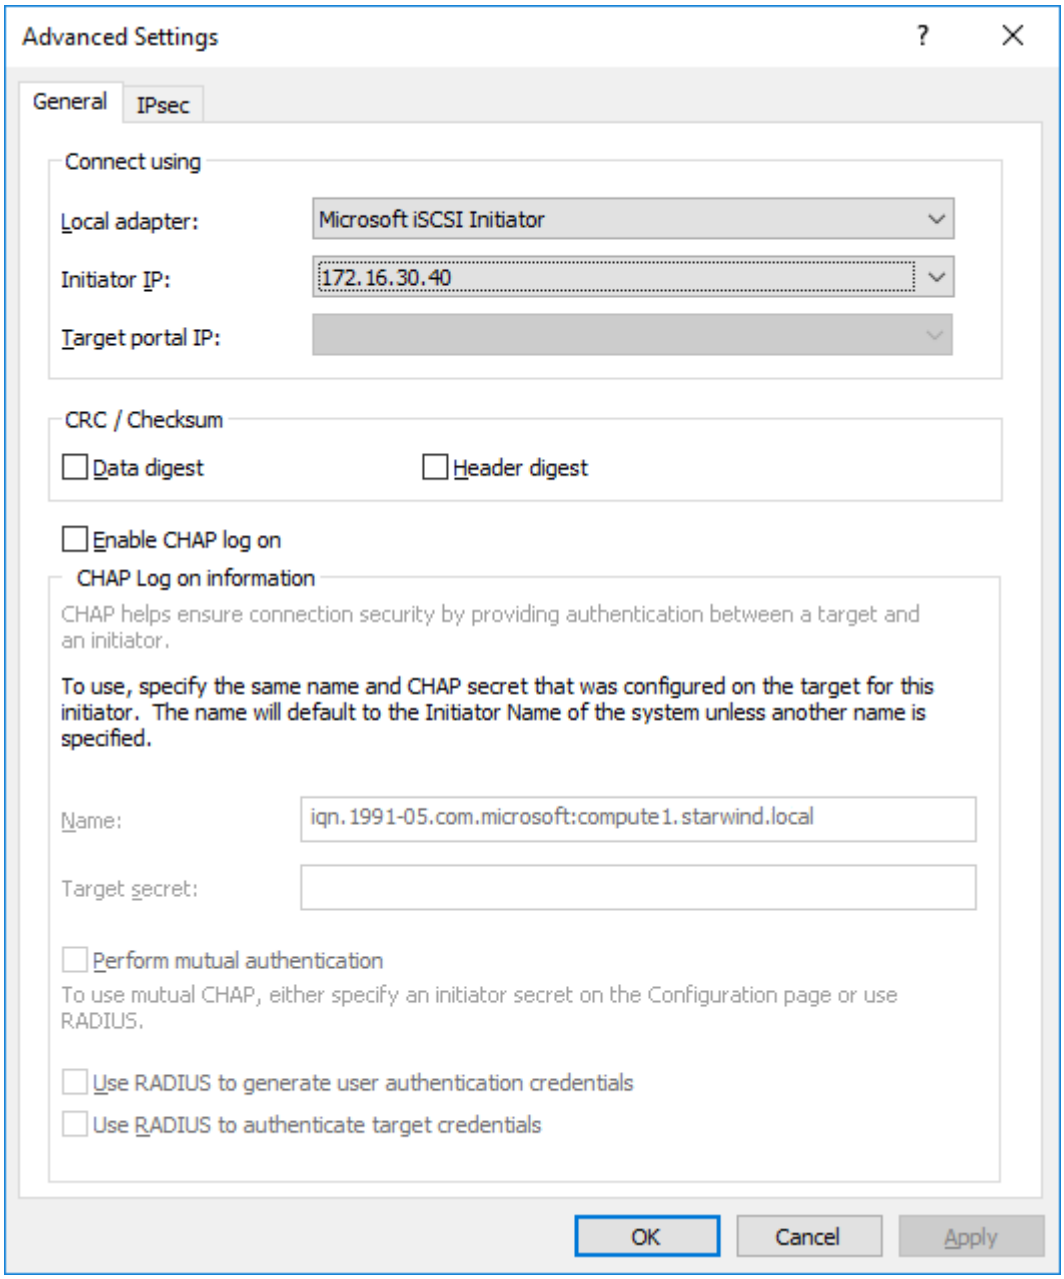

9. Repeat the steps 1-8 on for the 172.16.40.X subnet.

10. Repeat the steps 1-9 on the partner node

11. All the target portals are added on the first Compute node.

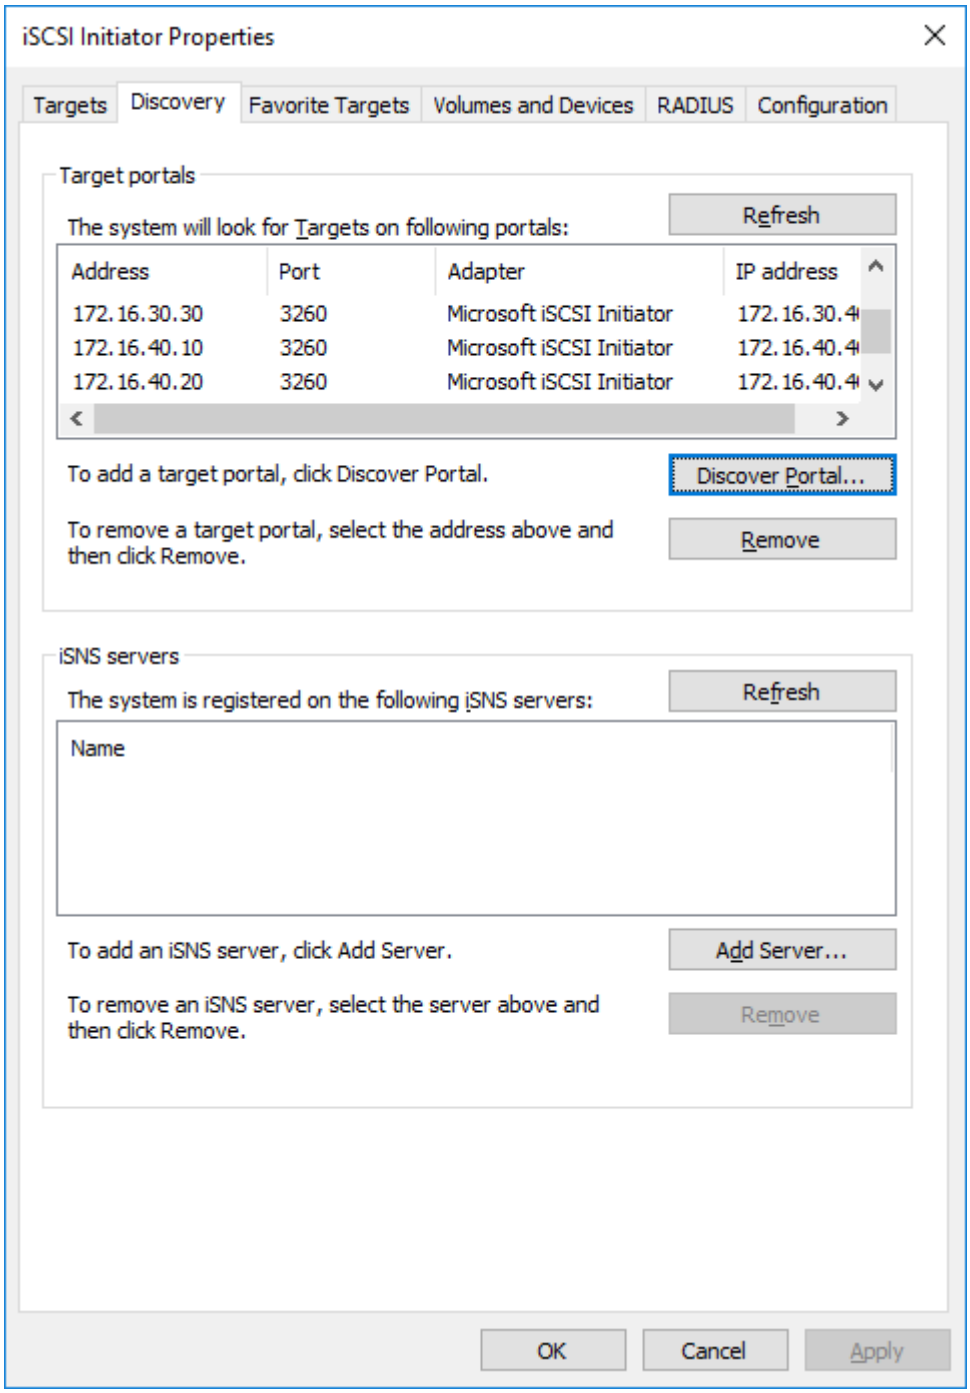

Address	Port	Adapter	IP Adress
172.16.30.10	3260	Microsoft iSCSI initiator	172.16.30.40
172.16.30.20	3260	Microsoft iSCSI initiator	172.16.30.40
172.16.30.30	3260	Microsoft iSCSI initiator	172.16.30.40
172.16.40.10	3260	Microsoft iSCSI initiator	172.16.30.40

172.16.40.20	3260	Microsoft iSCSI initiator	172.16.30.40
172.16.40.30	3260	Microsoft iSCSI initiator	172.16.30.40

12. All the target portals are added on the second Compute node.

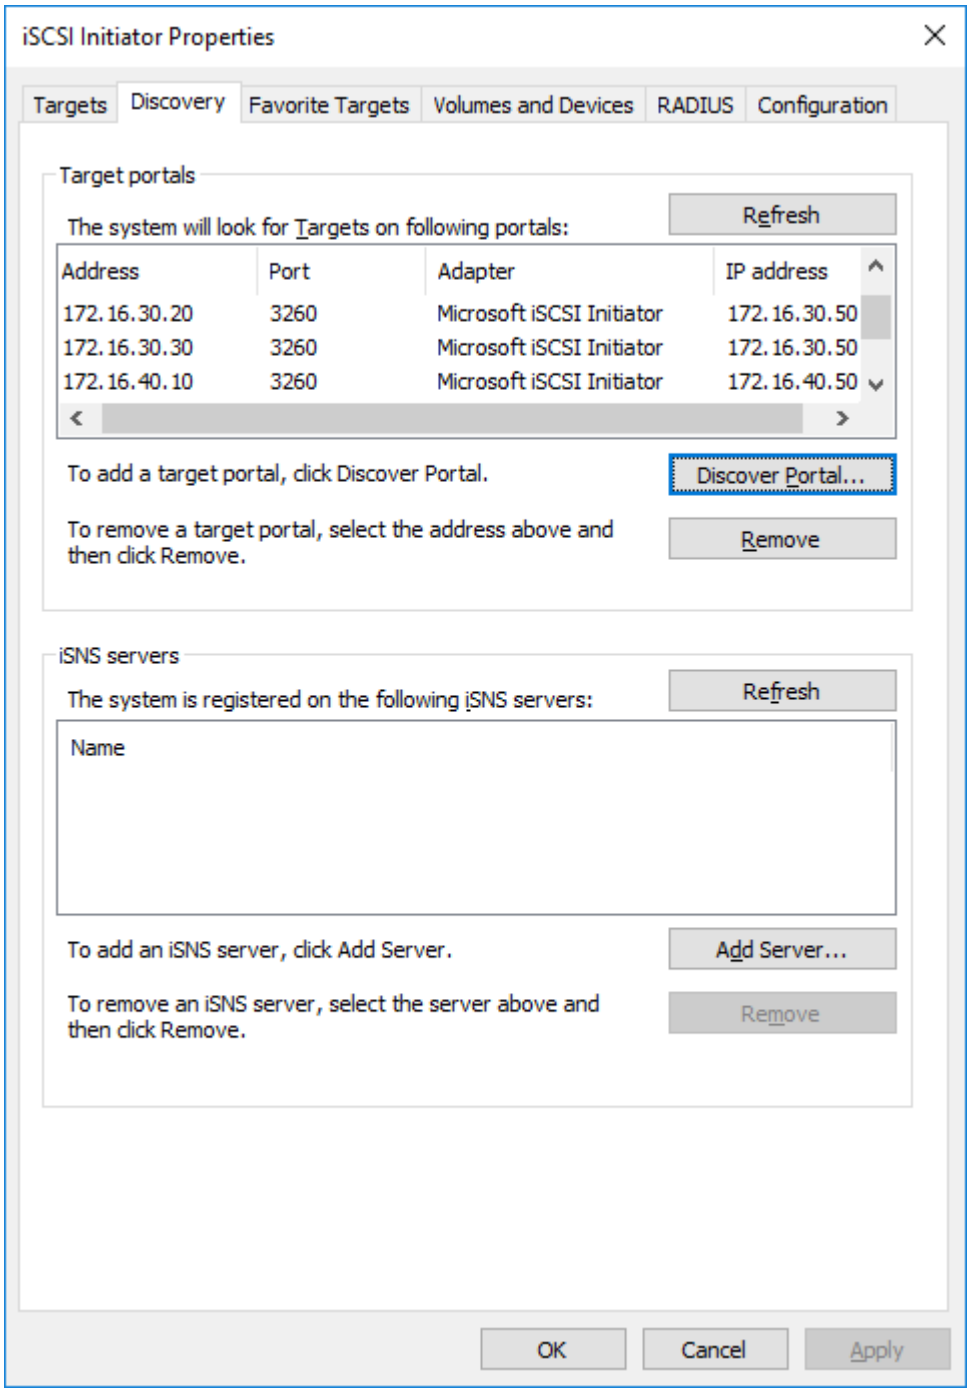

Address	Port	Adapter	IP Adress
172.16.30.10	3260	Microsoft iSCSI initiator	172.16.30.50

172.16.30.20	3260	Microsoft iSCSI initiator	172.16.30.50
172.16.30.30	3260	Microsoft iSCSI initiator	172.16.30.50
172.16.40.10	3260	Microsoft iSCSI initiator	172.16.30.50
172.16.40.20	3260	Microsoft iSCSI initiator	172.16.30.50
172.16.40.30	3260	Microsoft iSCSI initiator	172.16.30.50

Connecting Targets

1. Click the Targets tab. The previously created targets are listed in the Discovered Targets section.
NOTE: If the created targets are not listed, check the firewall settings of the StarWind Node as well as the list of networks served by the StarWind Node (go to StarWind Management Console -> Configuration -> Network). Alternatively, check the Access Rights tab on the corresponding StarWind VSAN server in StarWind Management Console for any restrictions.

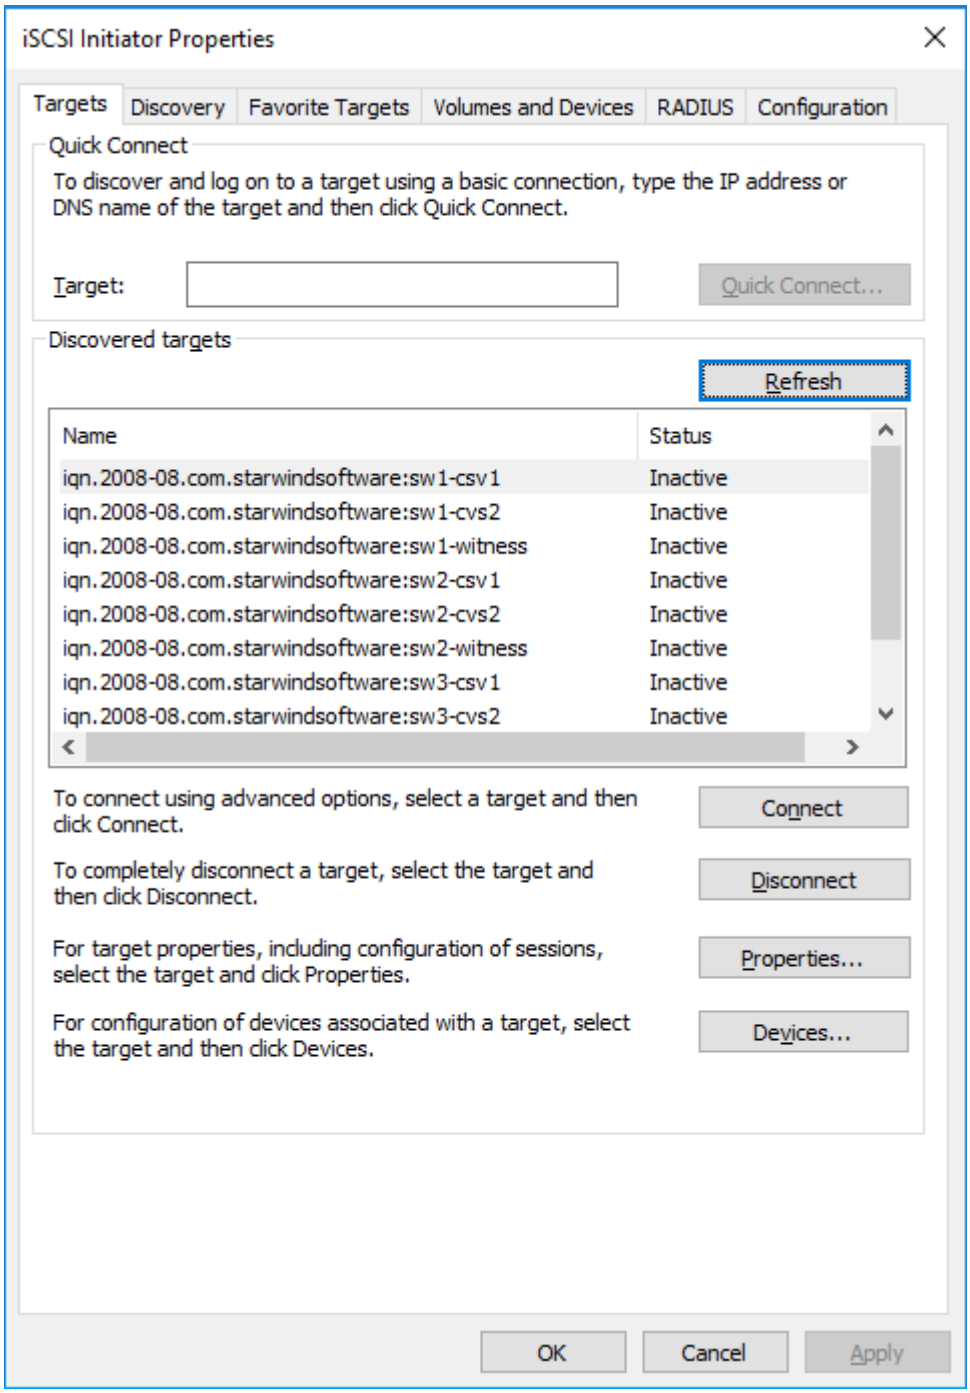

2. Select the Witness target from the first StarWind Node and click Connect.

3. Enable checkboxes as shown in the image below. Click Advanced...

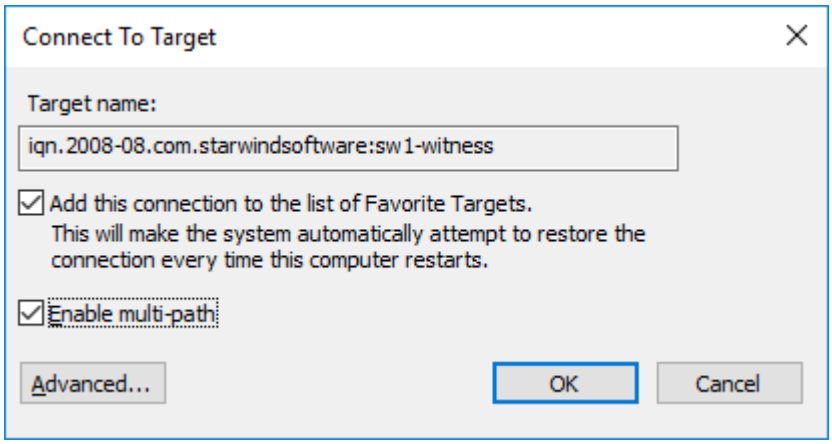

4. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Initiator IP, select 172.16.30.40, in Target portal IP, select 172.16.30.10. Confirm the action

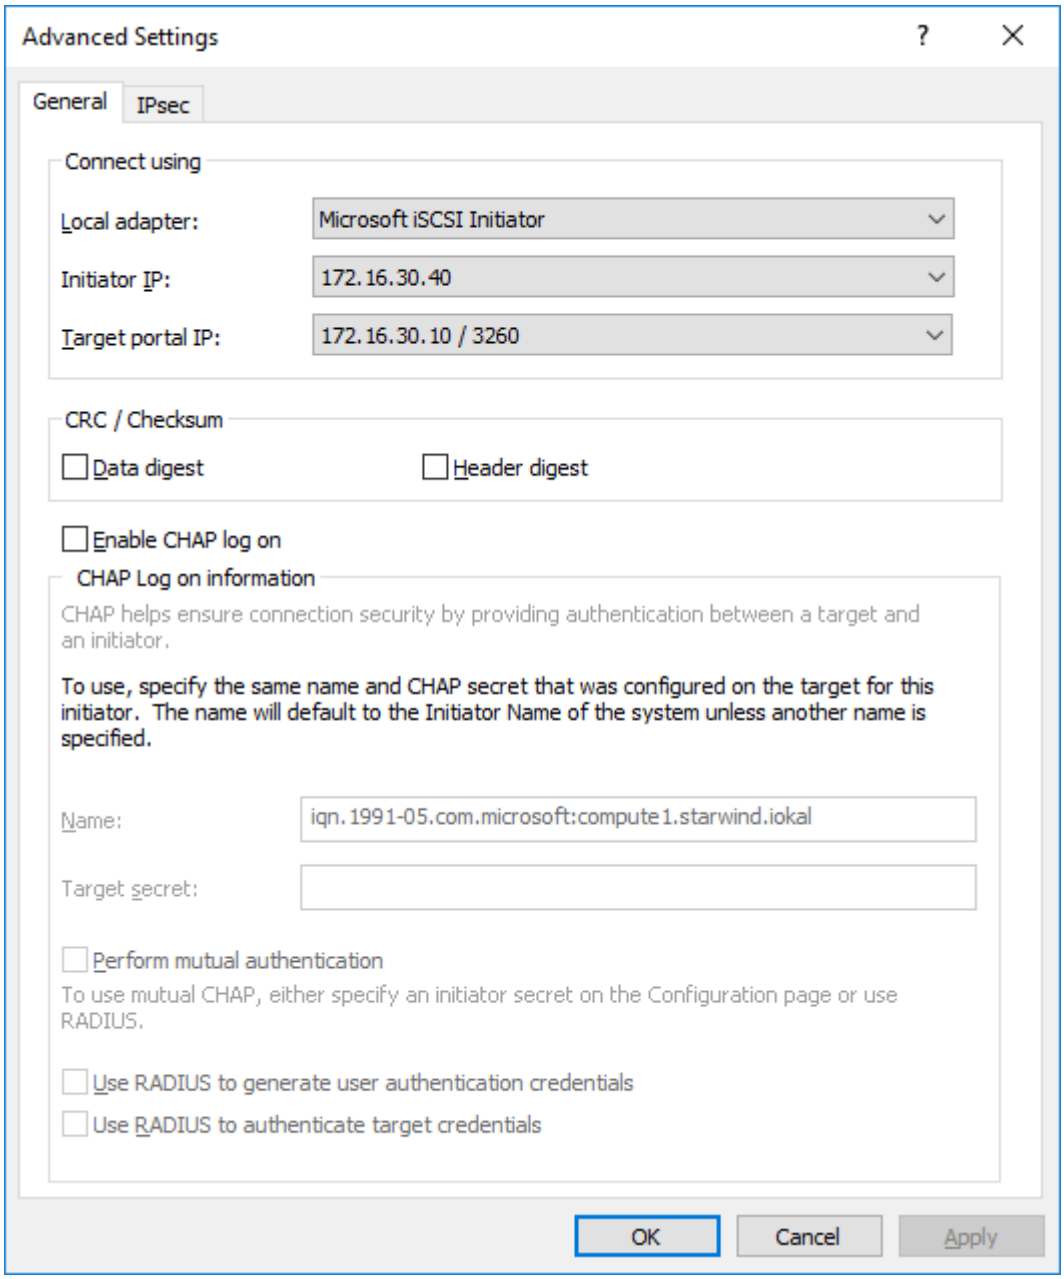

5. Select the Witness target from the first StarWind Node again and click Connect.

6. Enable checkboxes as shown in the image below. Click Advanced...

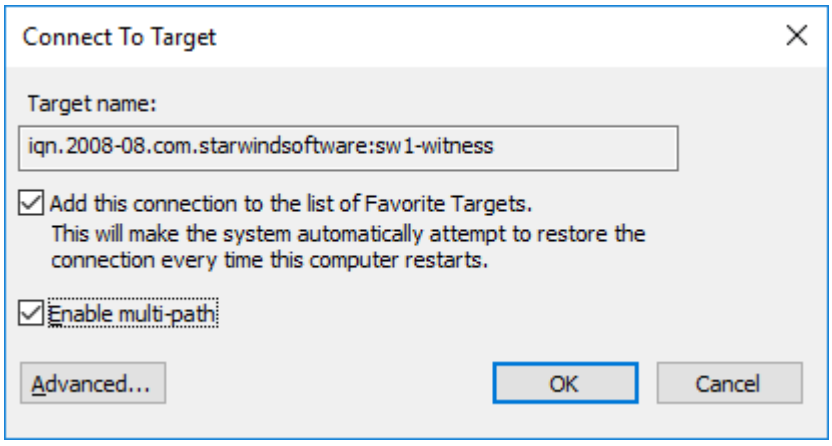

7. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Initiator IP, select 172.16.40.40, in Target portal IP, select 172.16.40.10. Confirm the action

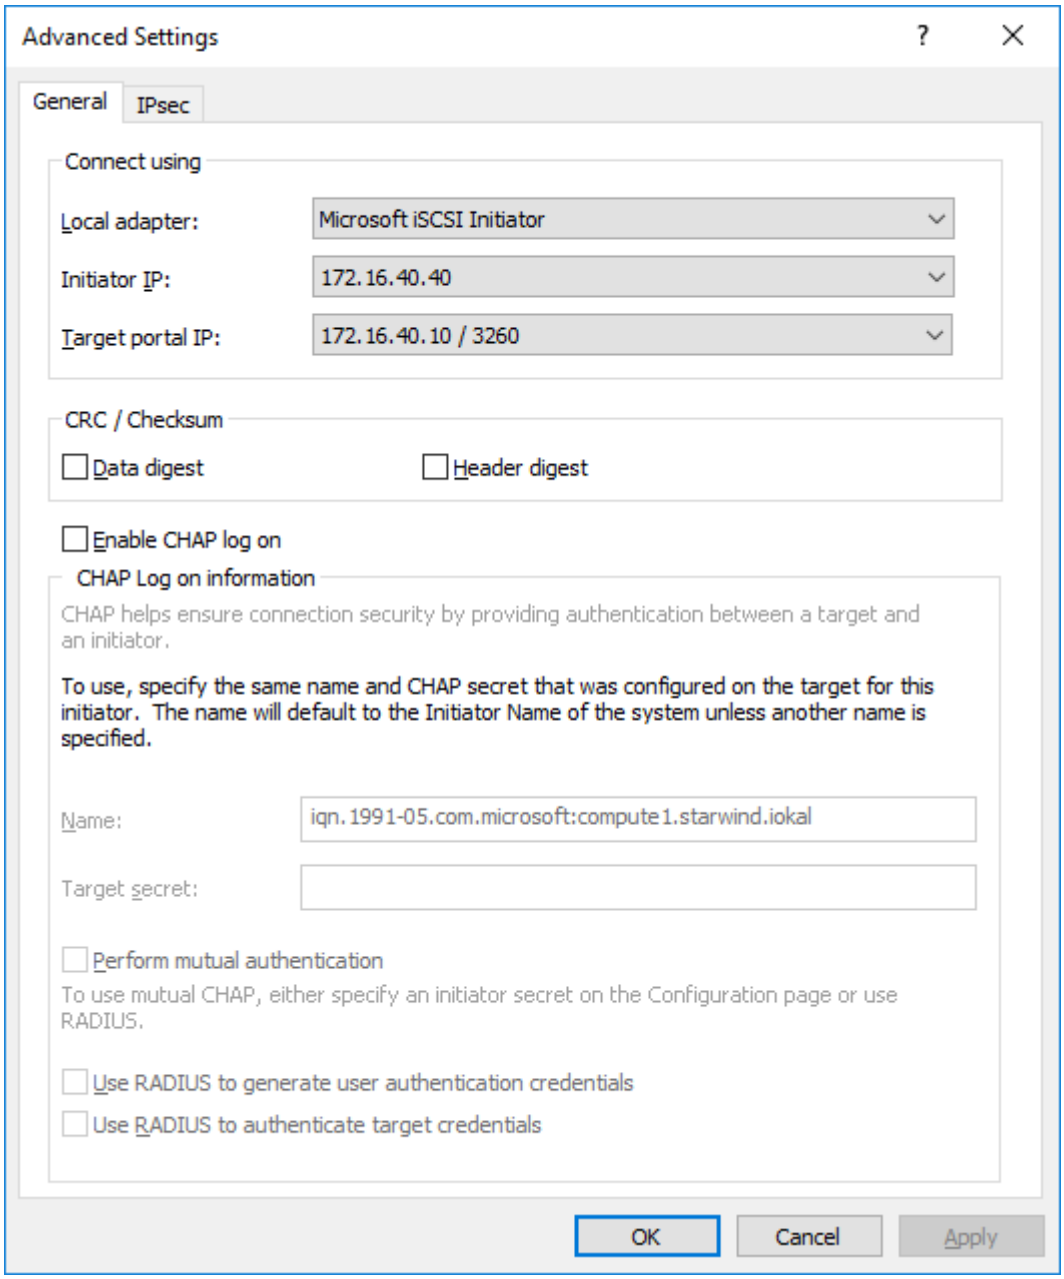

8. Select the Witness target from the second StarWind Node and click Connect.

9. Enable checkboxes as shown in the image below. Click Advanced...

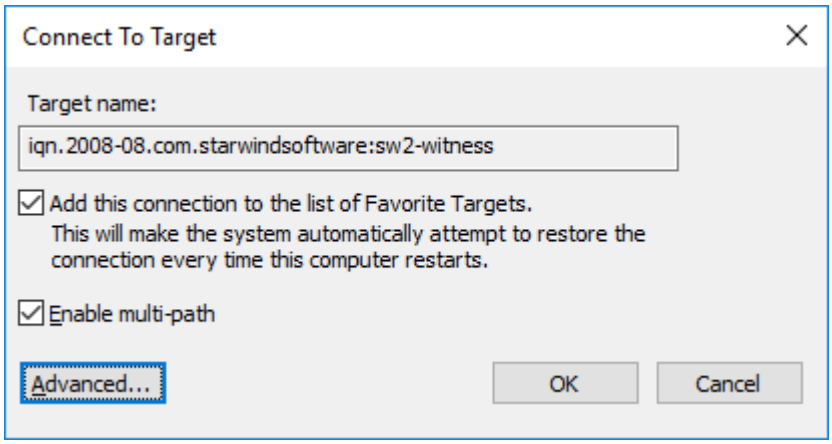

10. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Initiator IP, select 172.16.30.40, in Target portal IP, select 172.16.30.20. Confirm the action

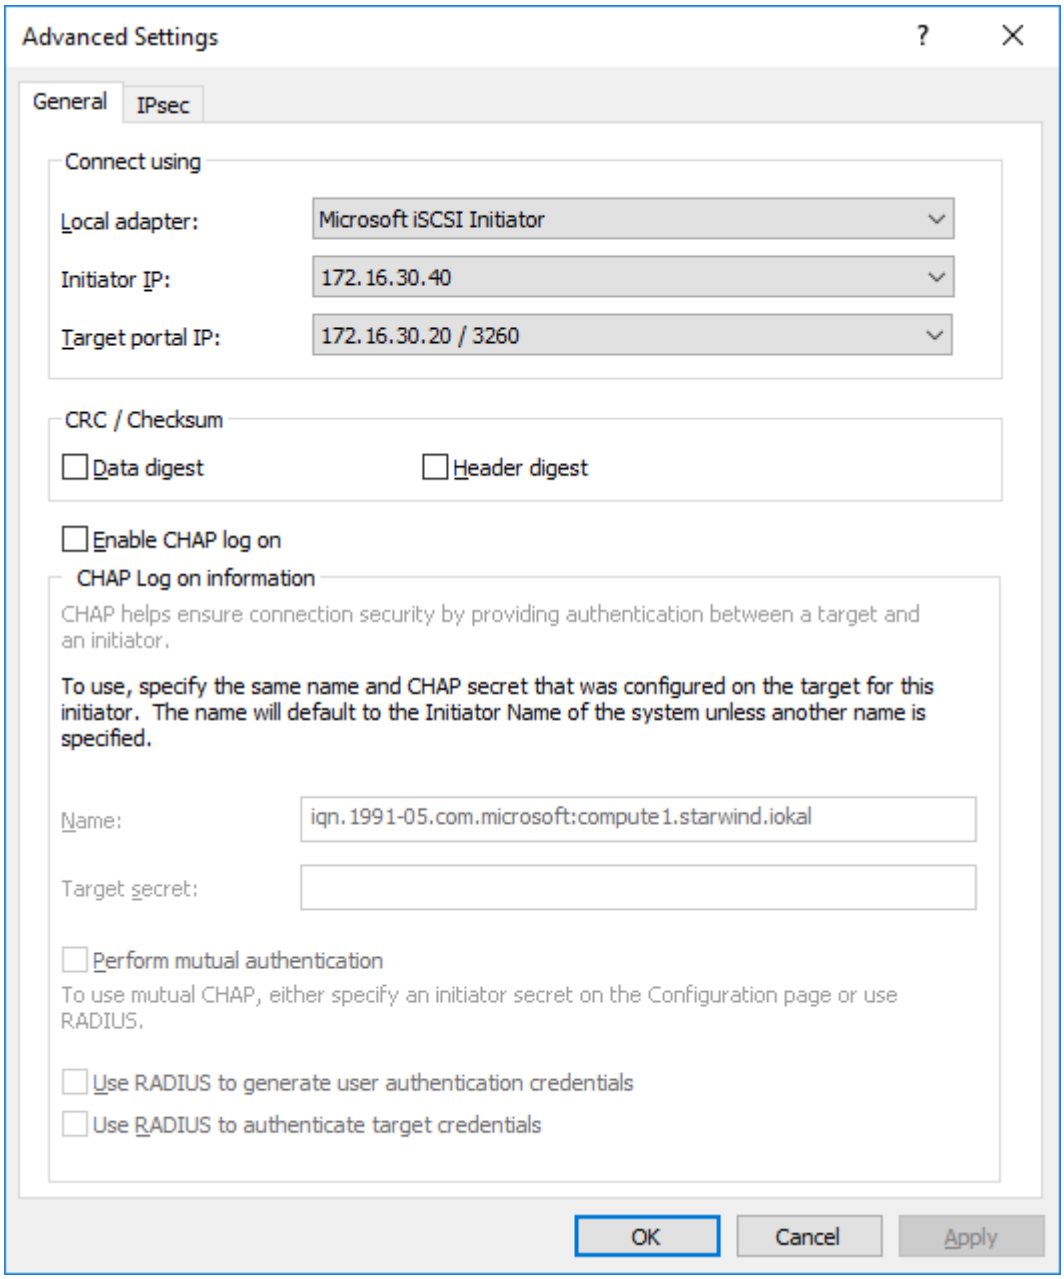

11. Select the Witness target from the second StarWind Node once again and click Connect.

12. Enable checkboxes as shown in the image below. Click Advanced...

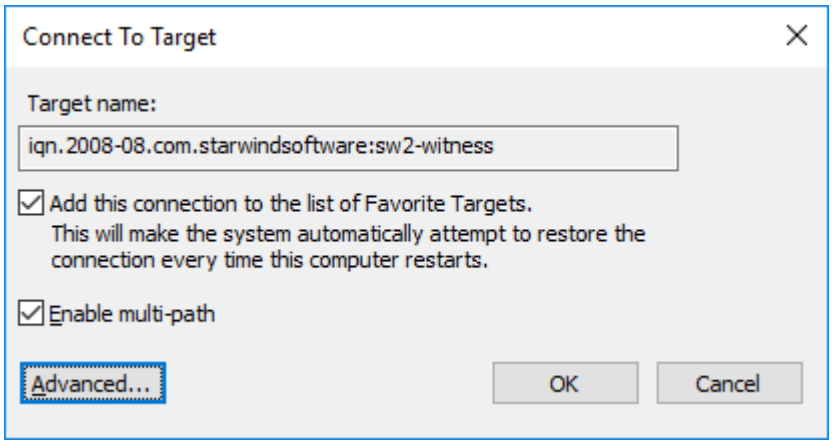

13. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Initiator IP, select 172.16.40.40, in Target portal IP, select 172.16.40.20. Confirm the action.

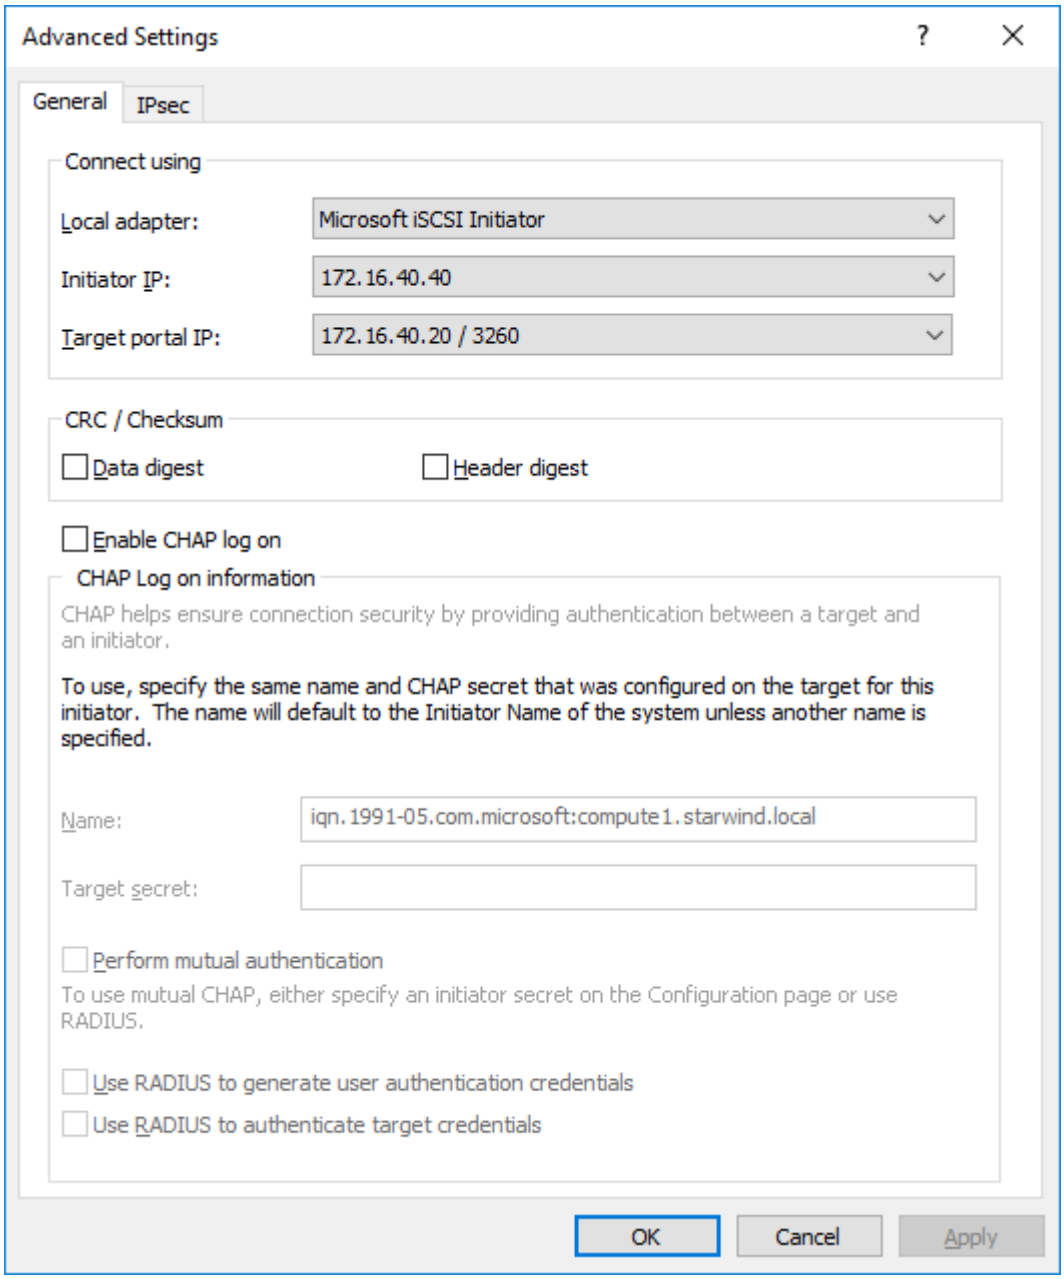

14. Select the Witness target from the third StarWind Node and click Connect.

15. Enable checkboxes as shown in the image below. Click Advanced...

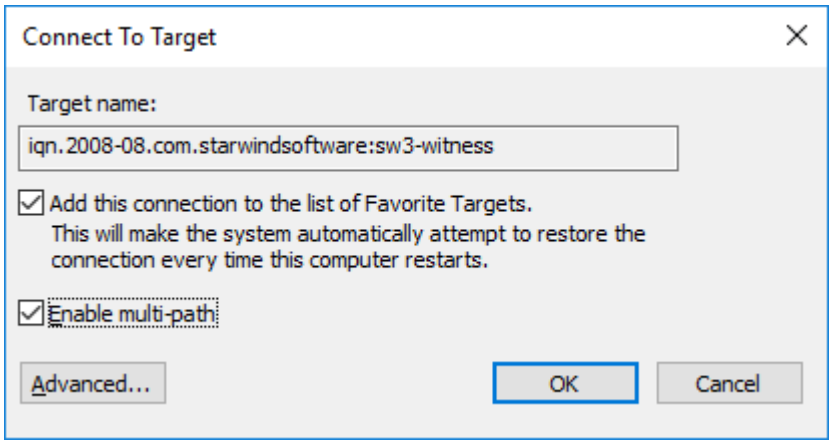

16. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Initiator IP, select 172.16.30.40, in Target portal IP, select 172.16.30.30. Confirm the action.

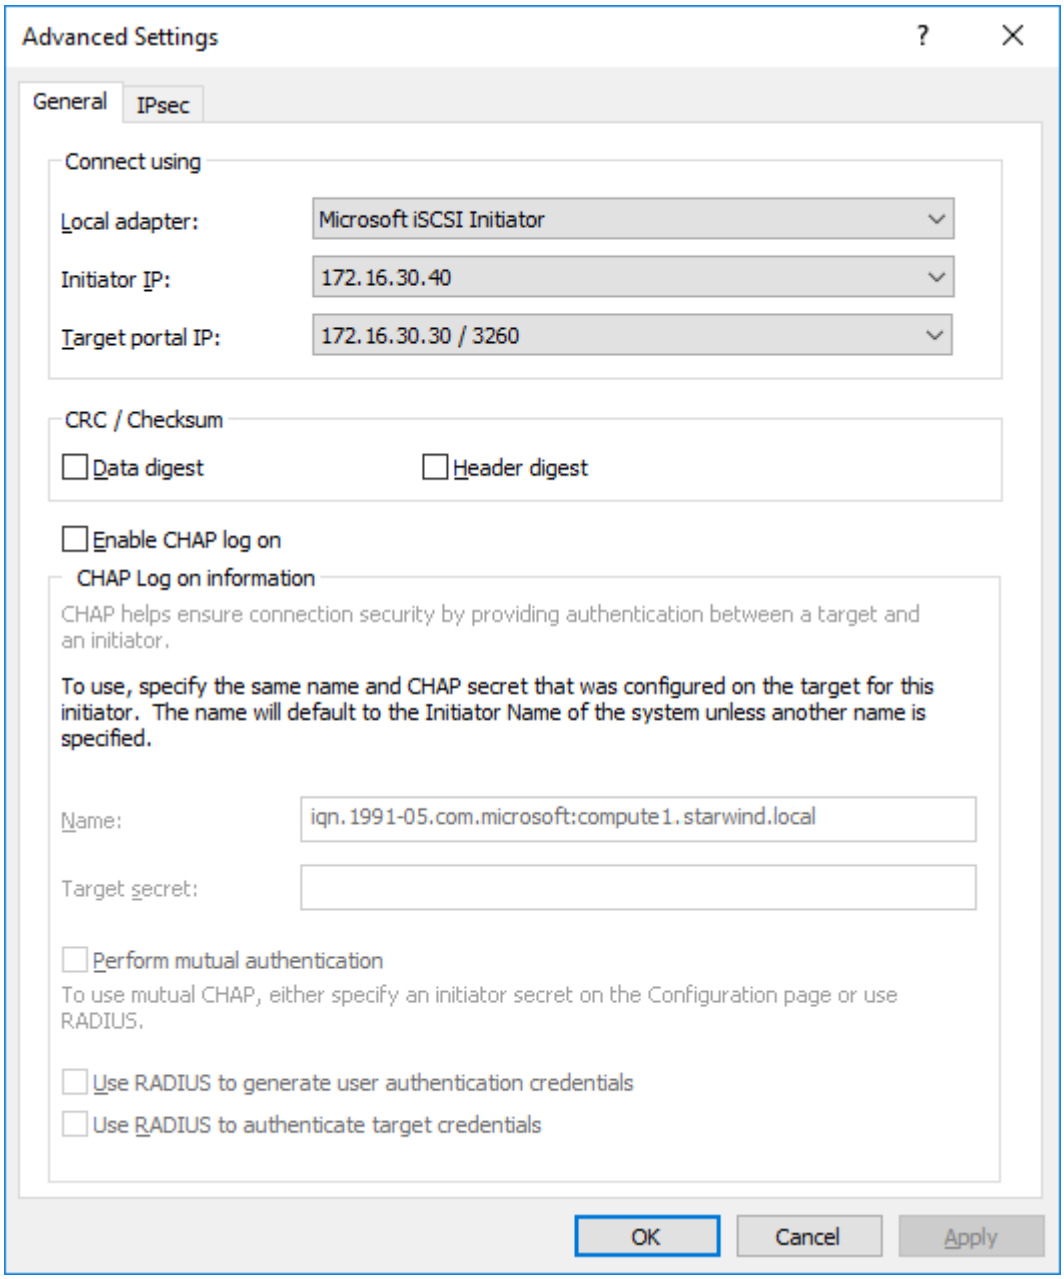

17. Select the Witness target from the third StarWind Node one more and click Connect.

18. Enable checkboxes as shown in the image below. Click Advanced...

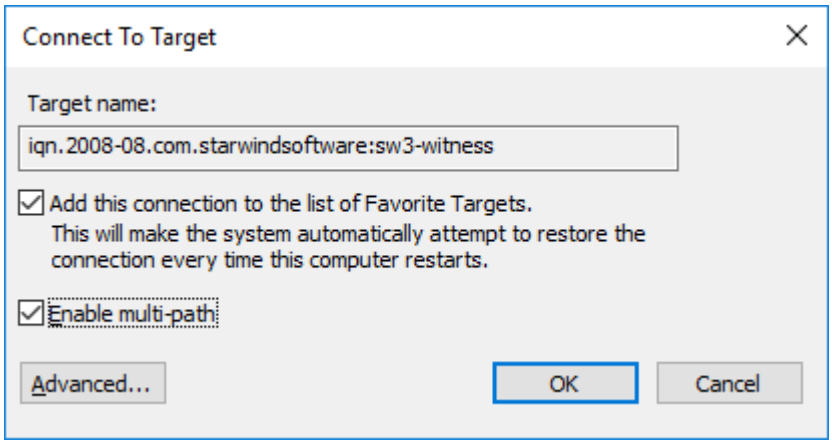

19. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Initiator IP, select 172.16.40.40, in Target portal IP, select 172.16.40.30. Confirm the action.

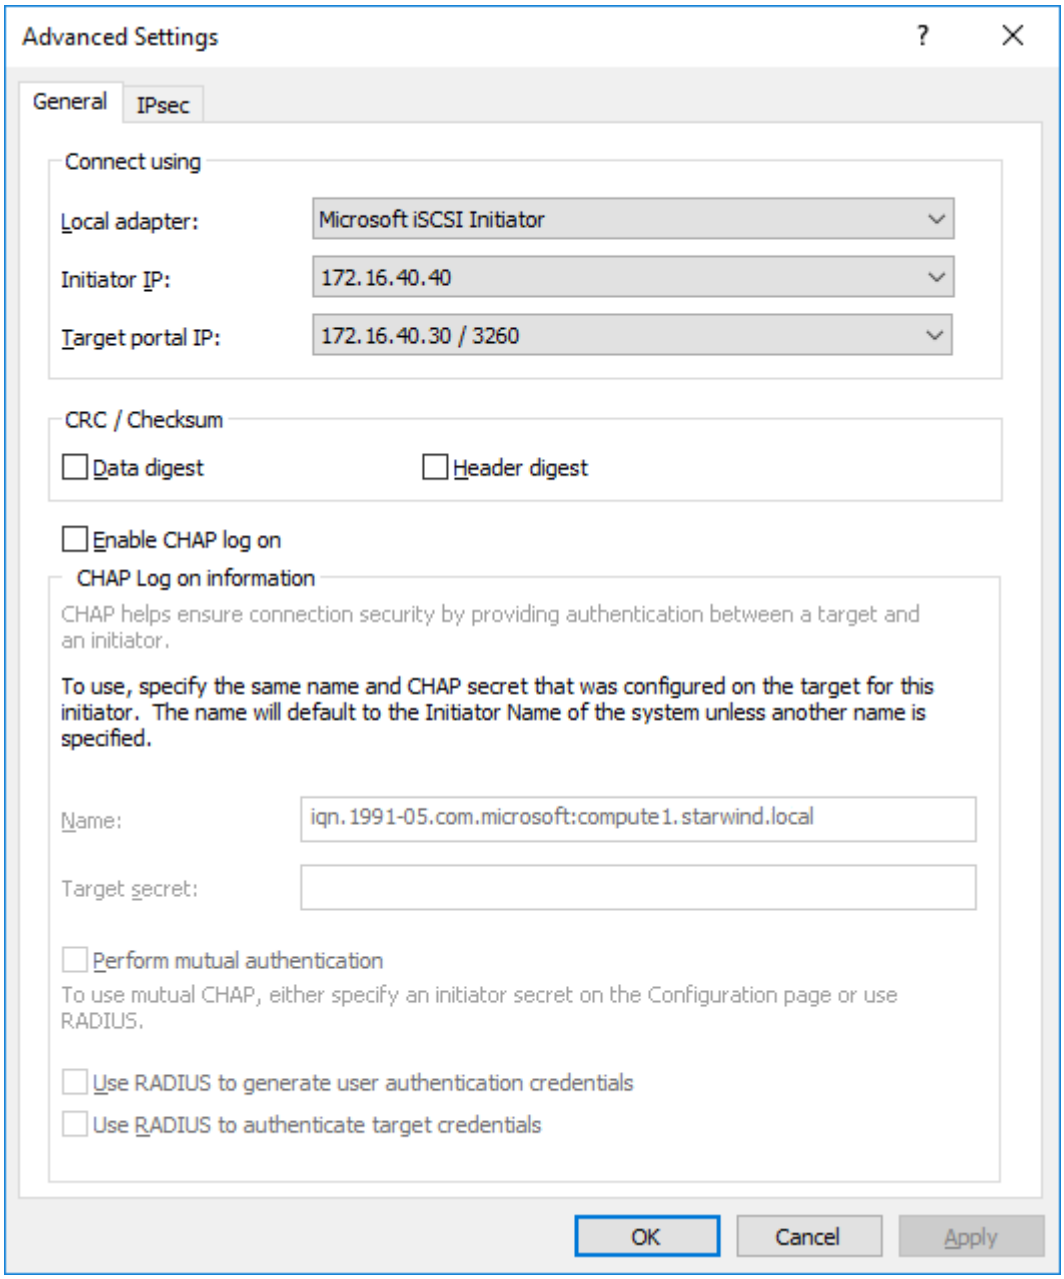

20. Repeat the steps 1-19 for all remaining HA device targets.

21. Repeat the steps 1-19 on the other Compute node, specifying the corresponding local and data channel IP addresses. The result should look like in the screenshot below.

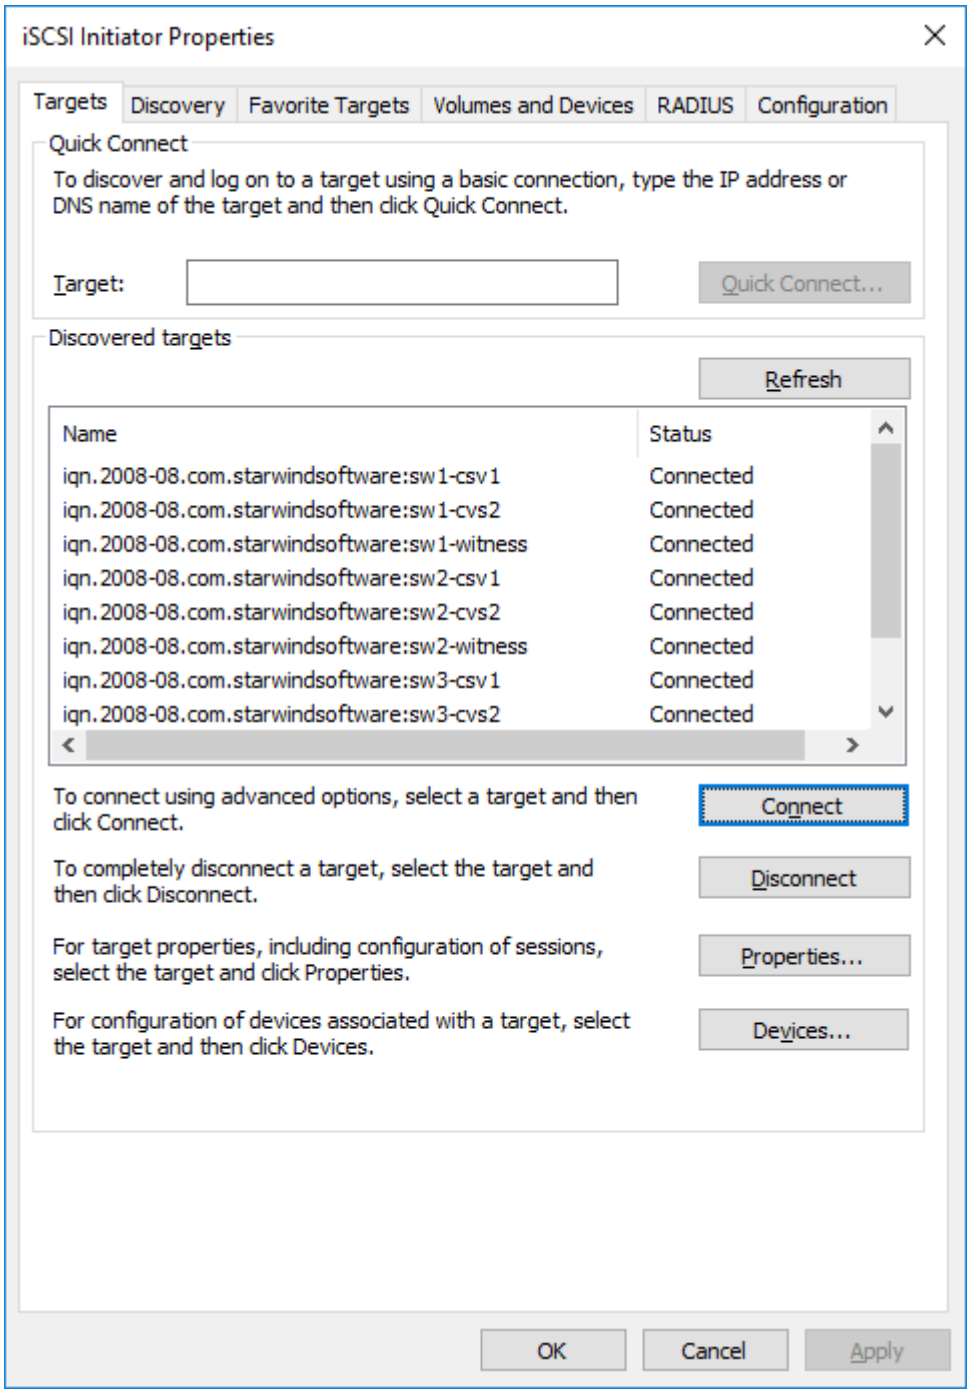

Configuring Multipath

NOTE: It is recommended to set the Round Robin or Least Queue Depth MPIO load balancing policy.

1. Configure the MPIO policy for every target leaving it with the load balance policy of choice. Select the Target located on the local server and click Devices.

2. In the Devices dialog, click MPIO

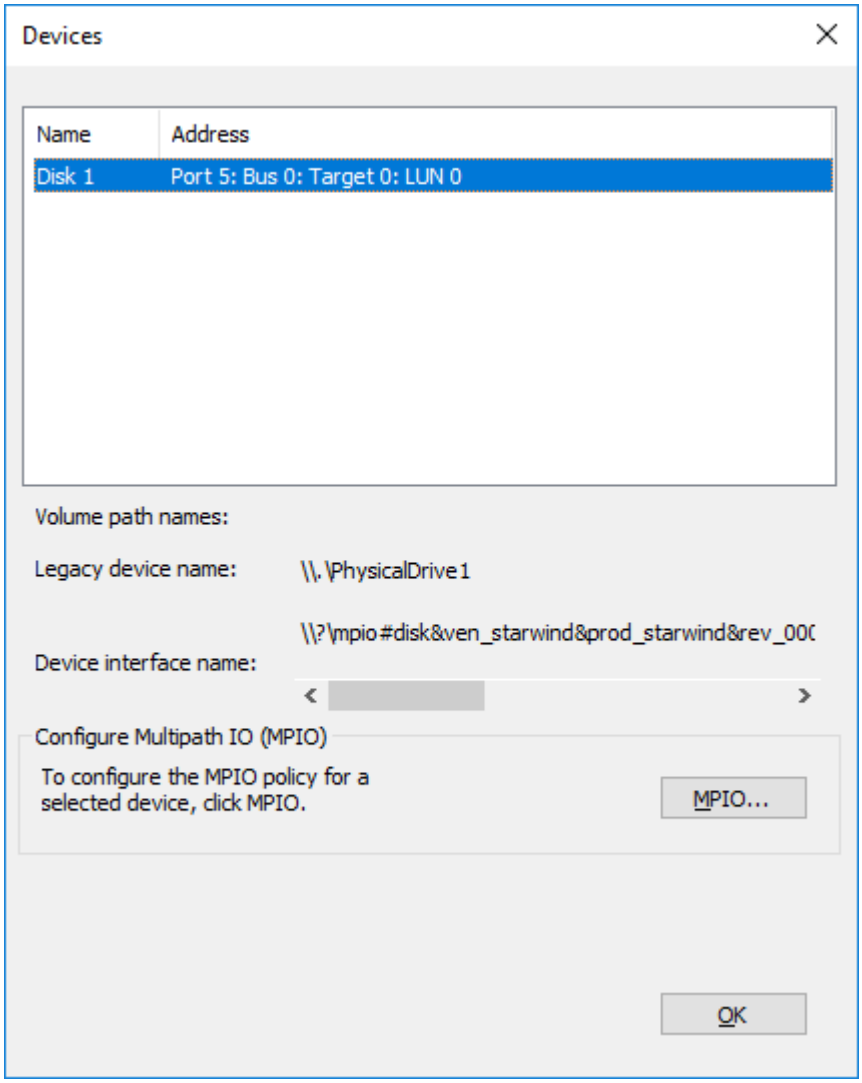

3. Select the appropriate load balancing policy.

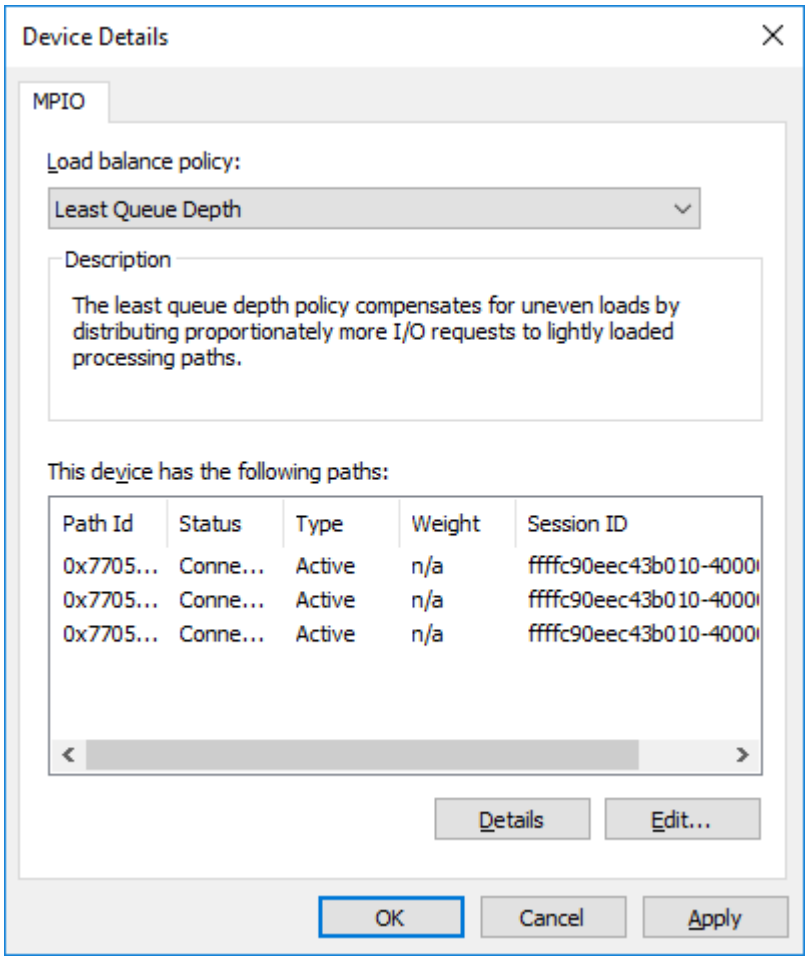

Configuring Disks to Servers

1. Open the Disk Management snap-in. The StarWind disks will appear as unallocated and offline.

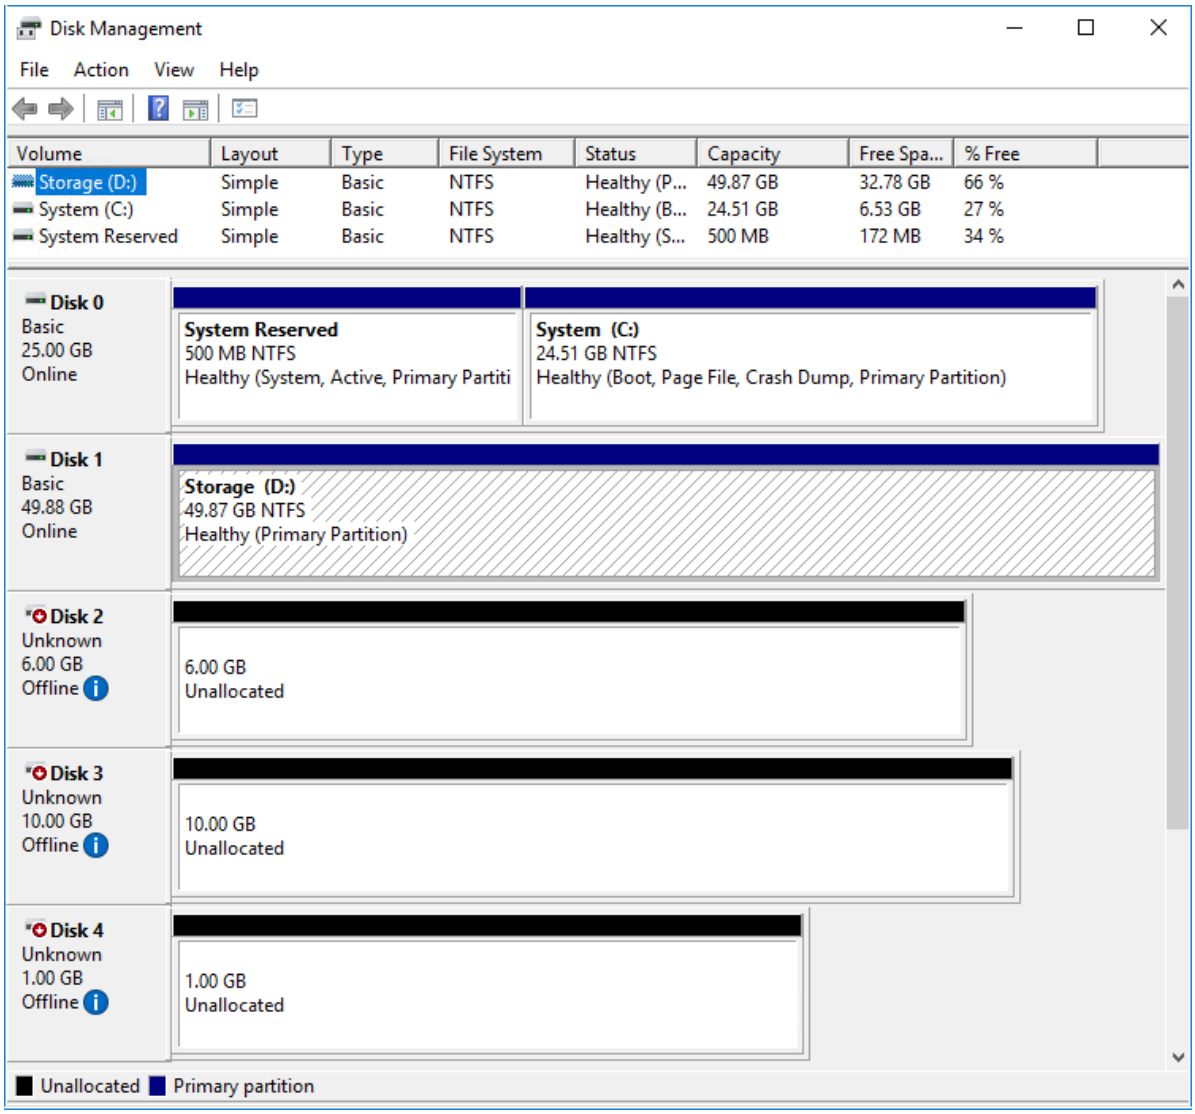

2. Bring the disks online by right-clicking on them and selecting the Online menu option.

3. Select the CSV disk (check the disk size to be sure) and right-click on it to initialize.

4. By default, the system will offer to initialize all non-initialized disks. Use the Select Disks area to choose the disks. Select GPT (GUID Partition Style) for the partition style to be applied to the disks. Press OK to confirm.

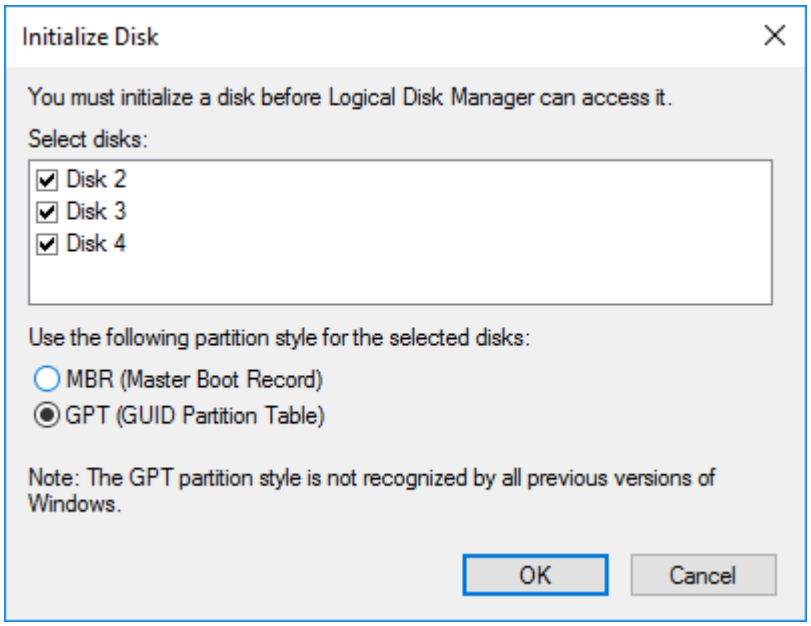

5. Right-click on the selected disk and choose New Simple Volume.

6. In New Simple Volume Wizard, indicate the volume size. Click Next.

7. Assign a drive letter to the disk. Click Next.

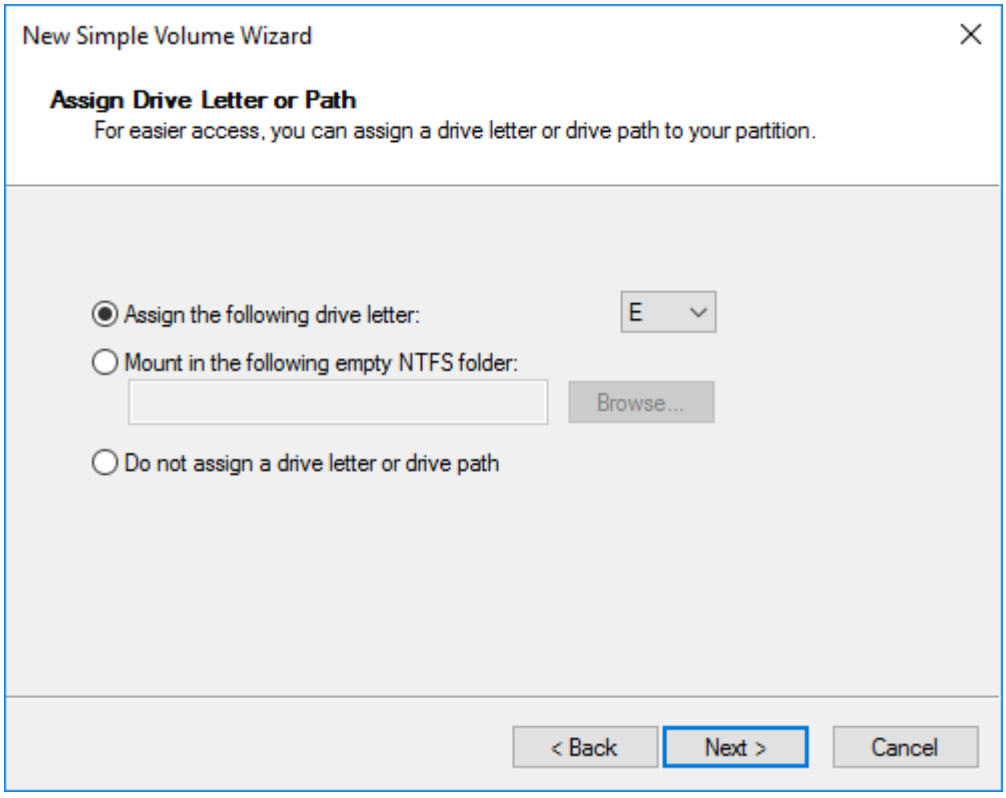

8. Select NTFS in the File System dropdown menu. Keep Allocation unit size as Default. Set the Volume Label of choice. Click Next.

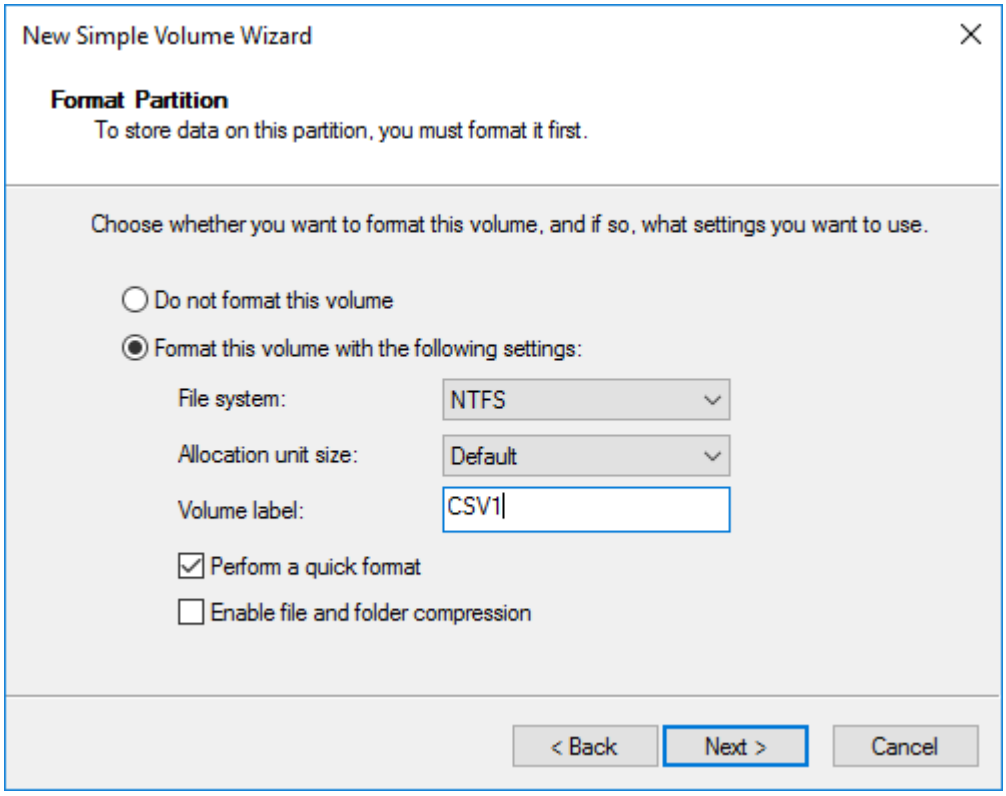

9. Press Finish to complete.

10. Complete the steps 1-9 for the Witness disk. Do not assign any drive letter or drive path for it.

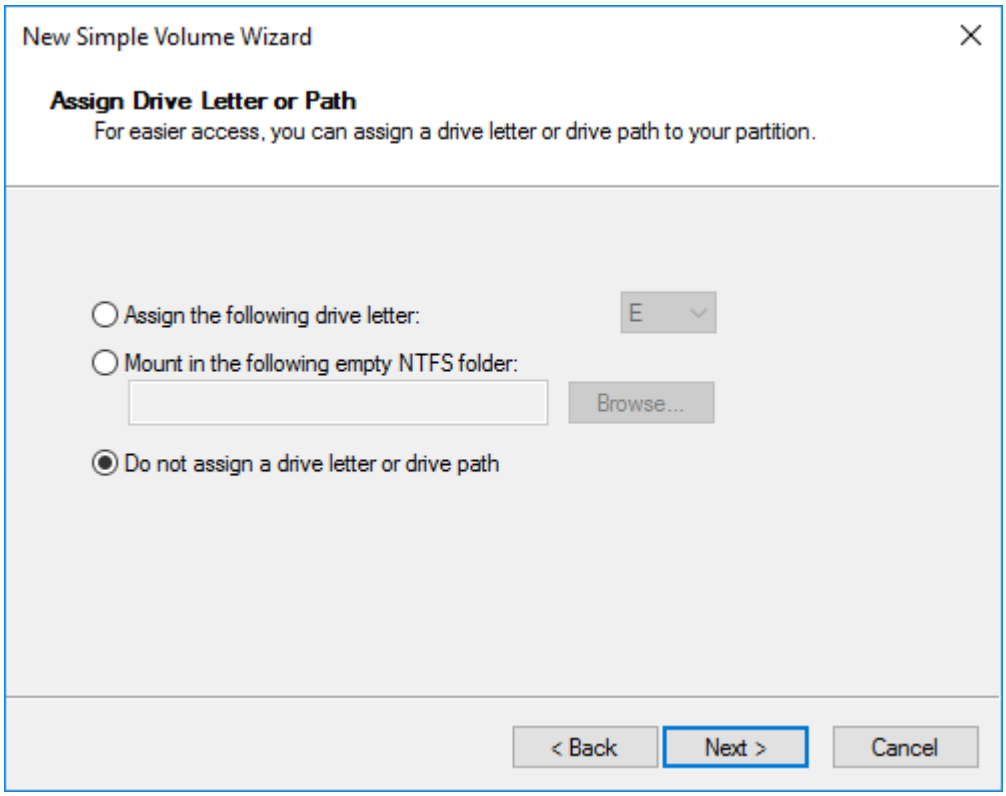

11. On the partner node, open the Disk Management snap-in. All StarWind disks will appear offline. If the status is different from the one shown below,
click Action->Refresh in the top menu to update the information about the disks.

12. Repeat step 2 to bring all the remaining StarWind disks online.

Creating A Failover Cluster In Windows Server

NOTE: To avoid issues during the cluster validation configuration, it is recommended to install the latest Microsoft updates on each node.

1. Open Server Manager. Select the Failover Cluster Manager item from the Tools menu.

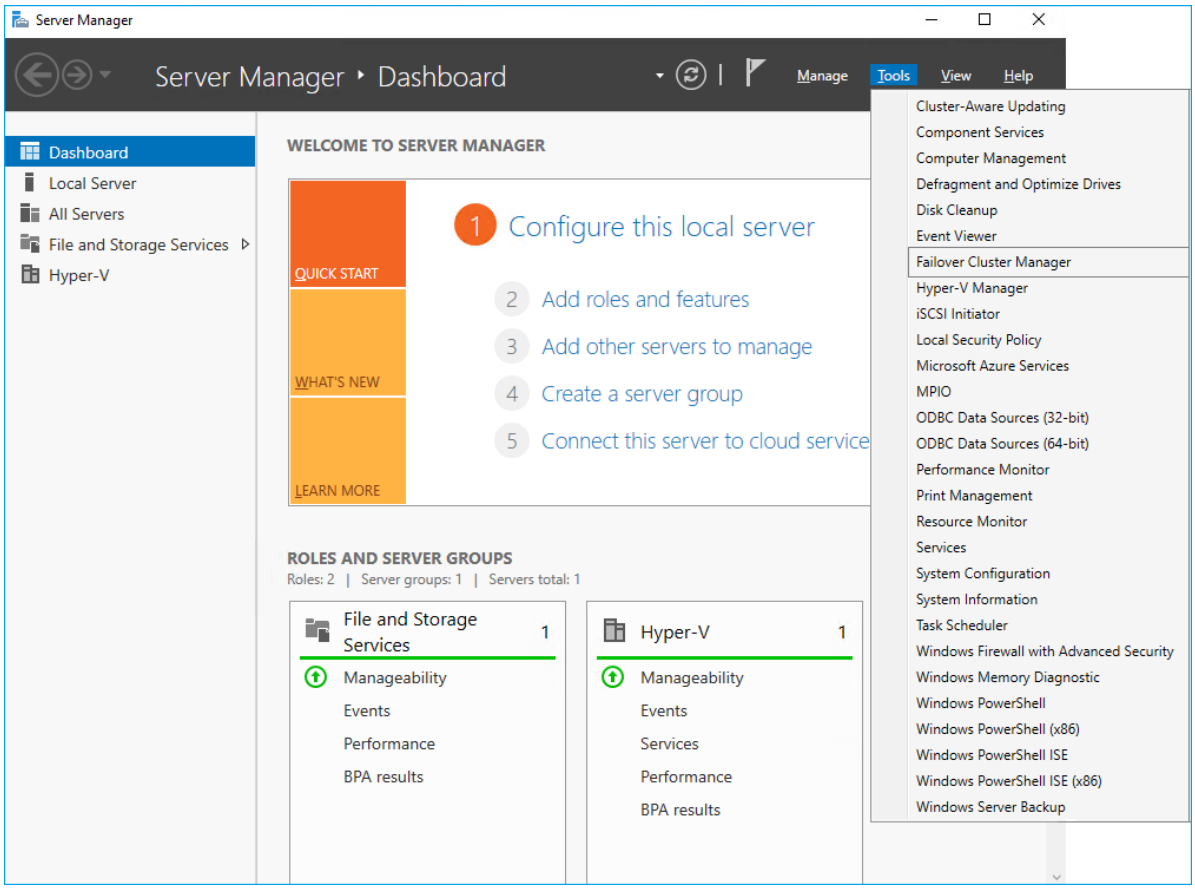

2. Click the Create Cluster link in the Actions section of Failover Cluster Manager.

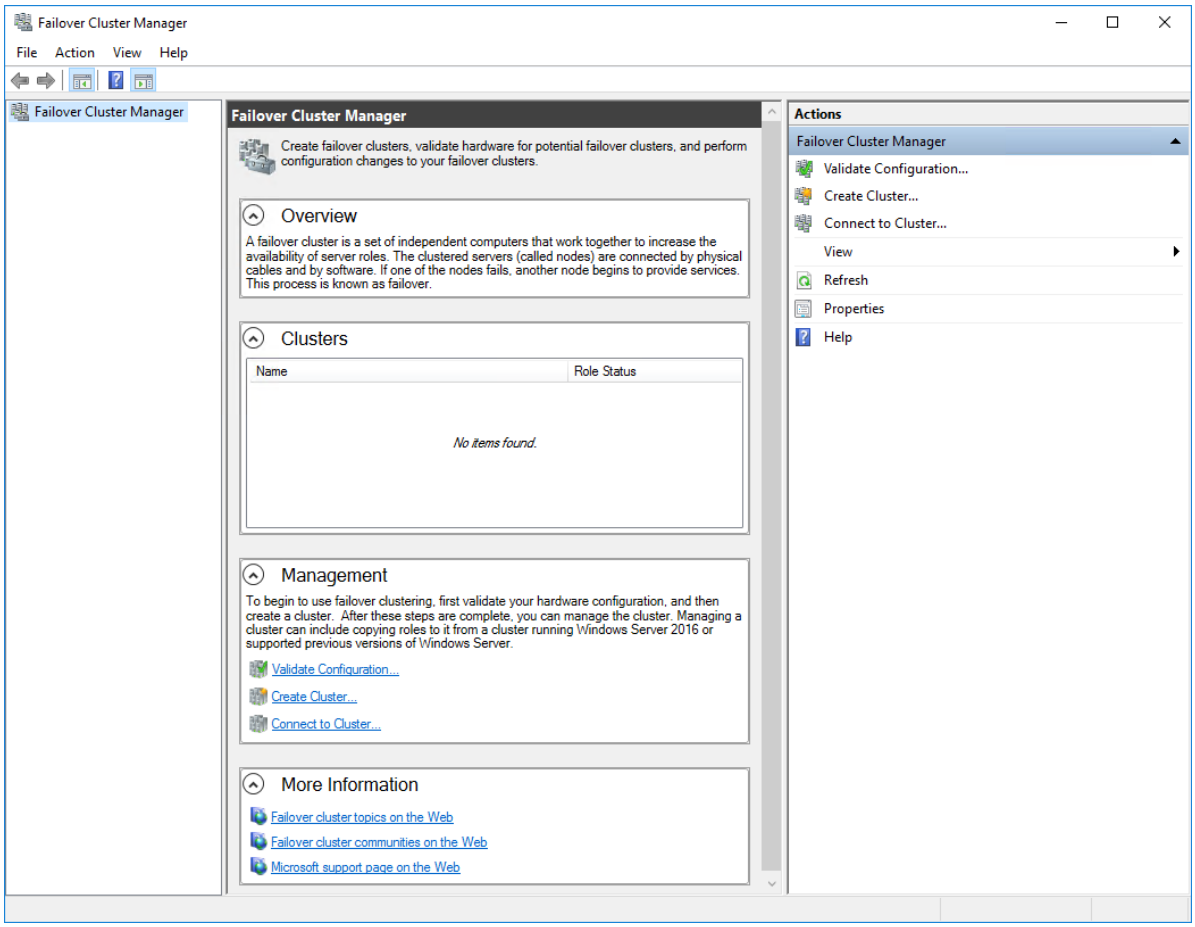

3. Specify the servers to be added to the cluster. Click Next to continue.

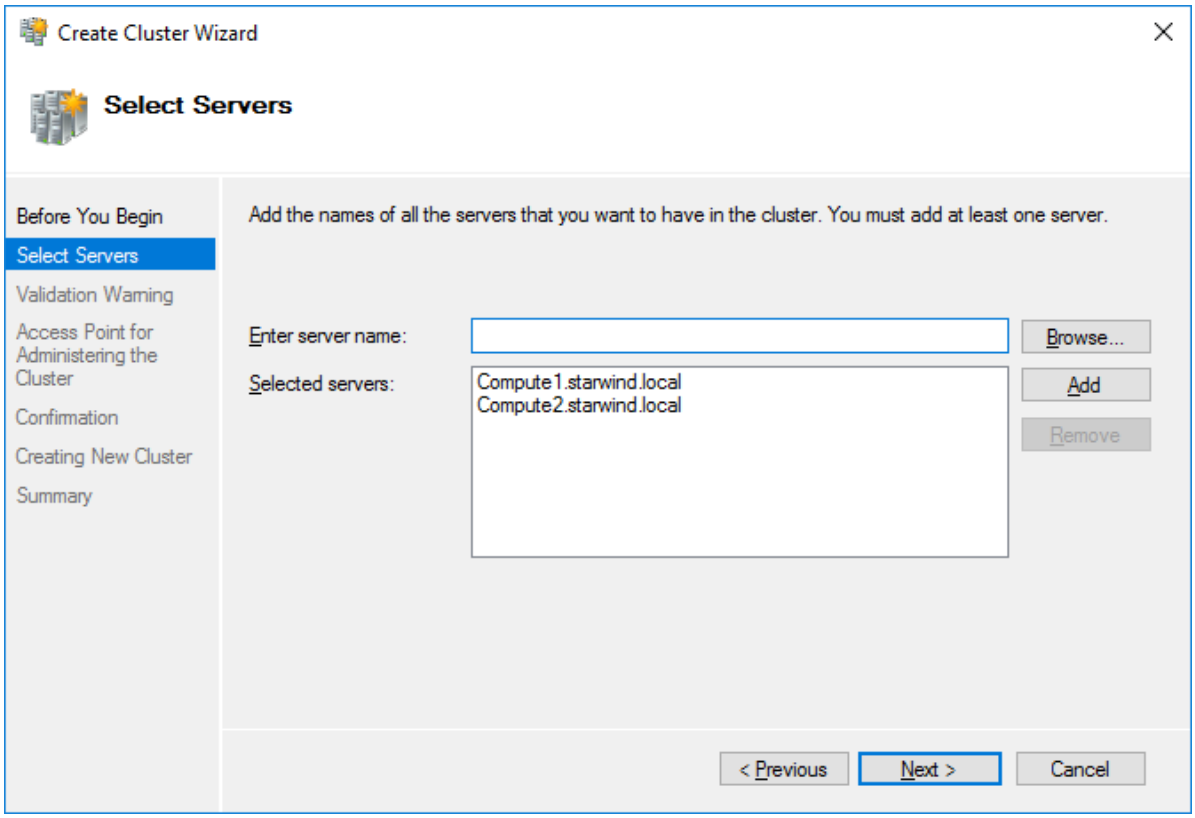

4. Validate the configuration by running the cluster validation tests: select Yes... and click Next to continue.

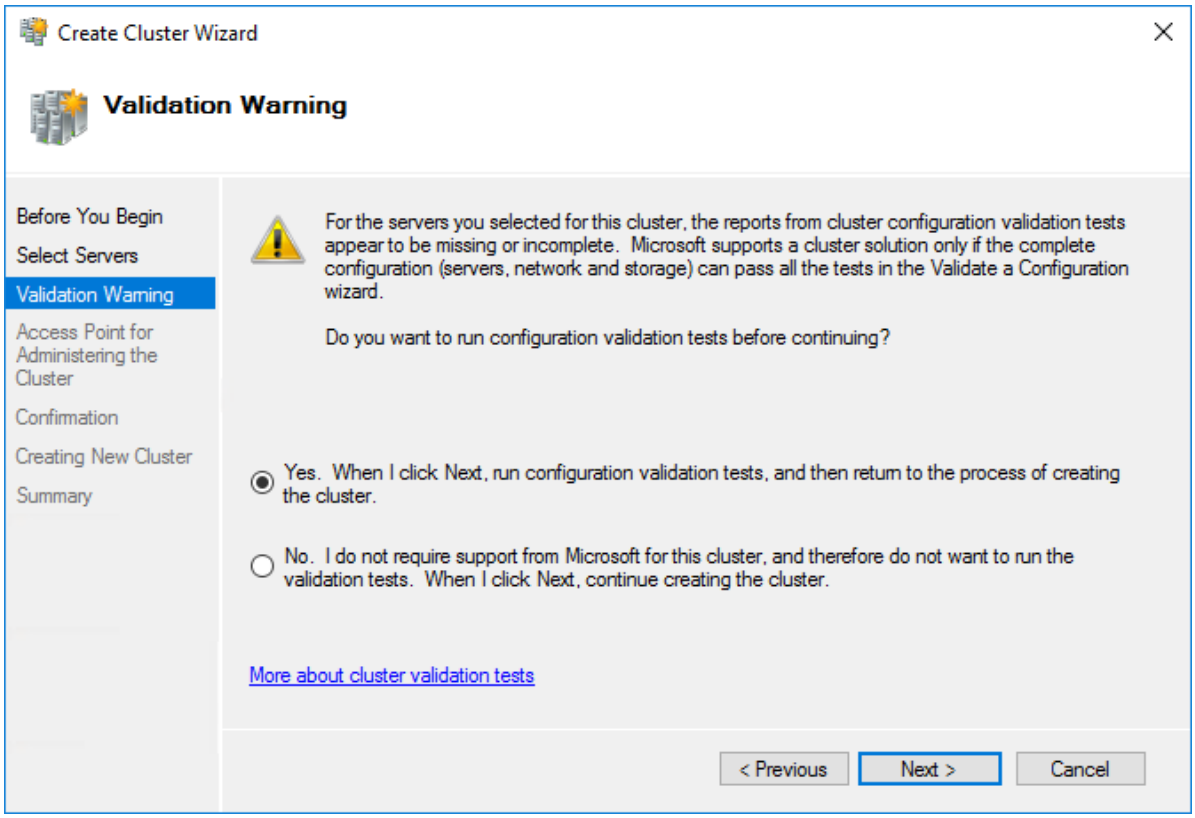

5. Specify Cluster Name.
NOTE: If the cluster servers get IP addresses over DHCP, the cluster also gets its IP address over DHCP. If the IP addresses are set statically, set the cluster IP address manually.

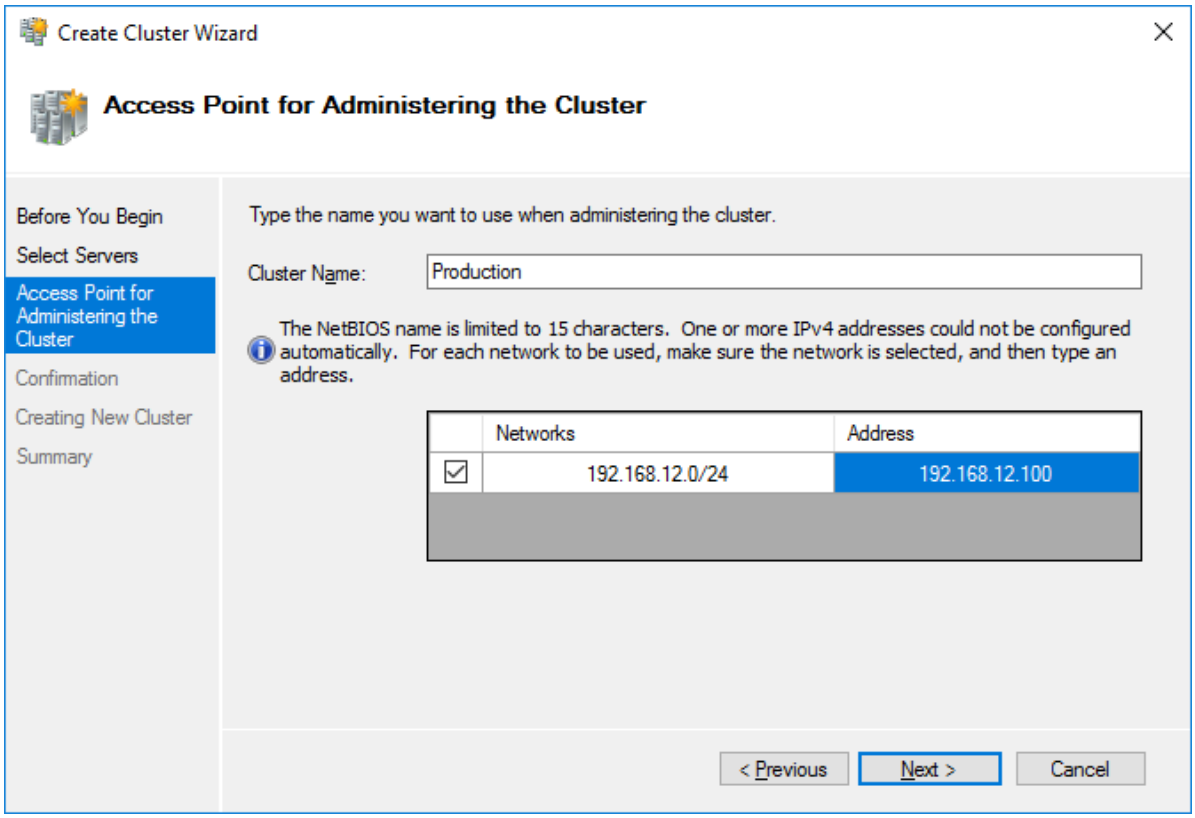

6. Make sure that all settings are correct. Click Previous to make any changes or Next to proceed.

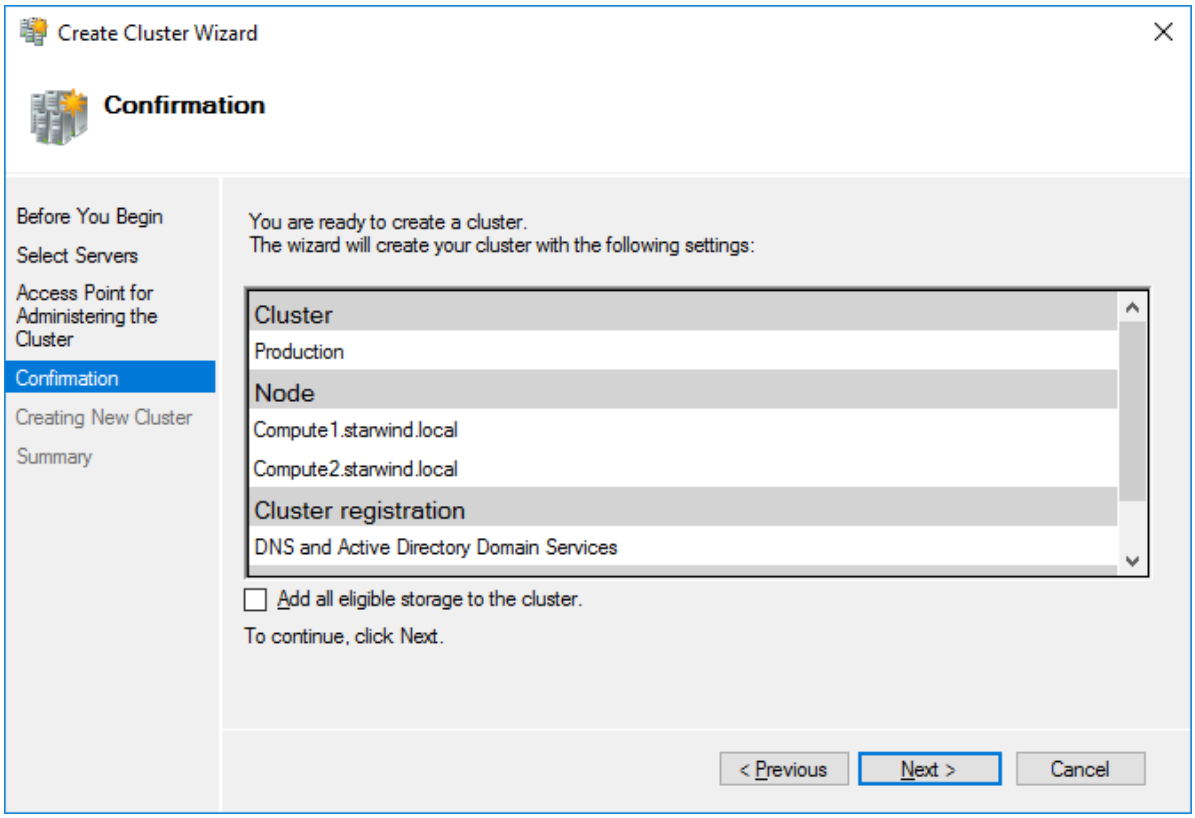

NOTE: If checkbox Add all eligible storage to the cluster is selected, the wizard will add all disks to the cluster automatically. The device with the smallest storage volume will be assigned as a Witness. It is recommended to uncheck this option before clicking Next and add cluster disks and the Witness drive manually.

7. The process of the cluster creation starts. Upon the completion, the system displays the summary with the detailed information. Click Finish to close the wizard.

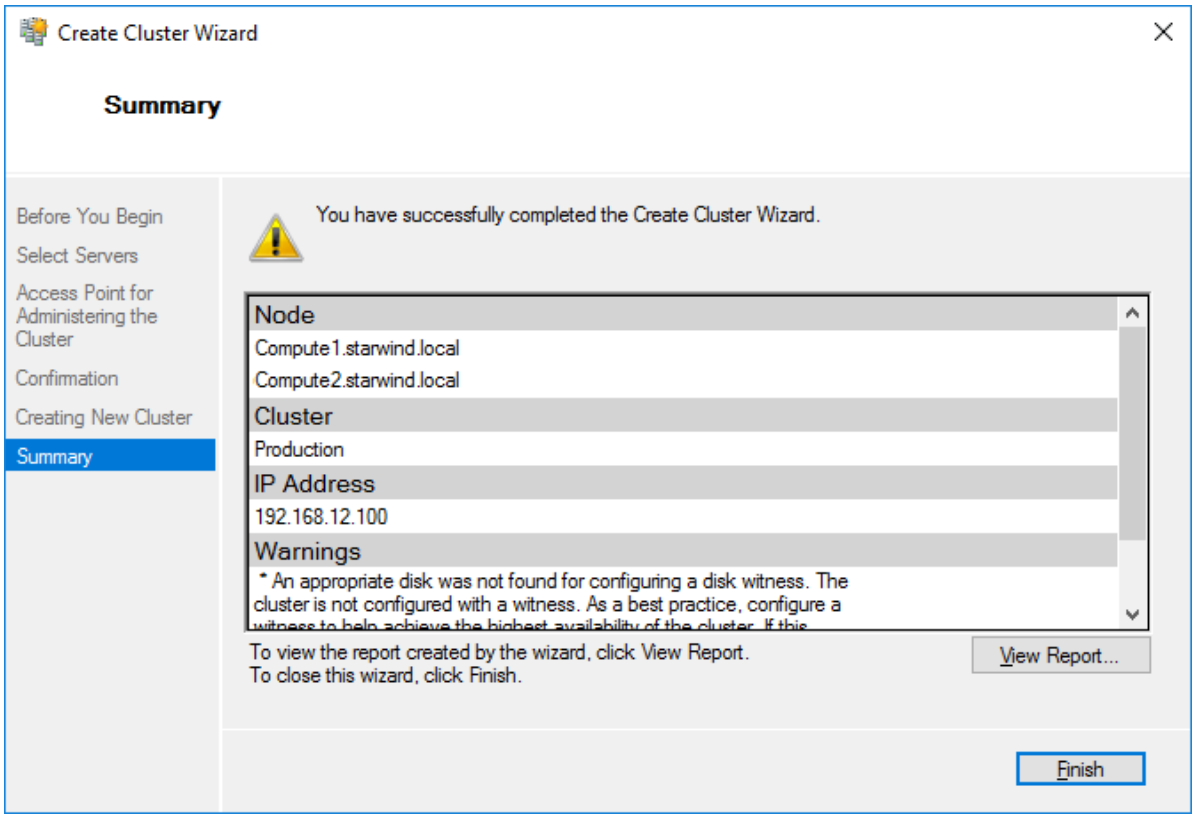

Adding Storage to the Cluster

1. In Failover Cluster Manager, navigate to Cluster -> Storage -> Disks. Click Add Disk in the Actions panel, choose StarWind disks from the list and confirm the selection.

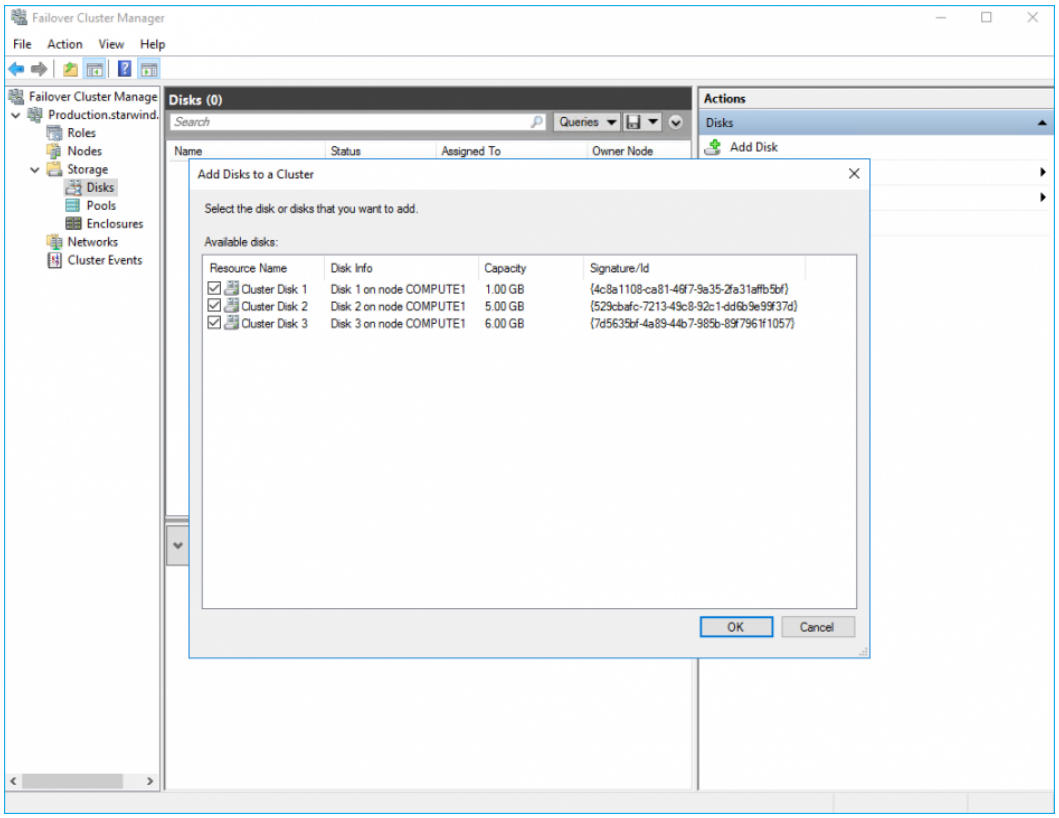

2. To configure the cluster witness disk, right-click on Cluster and proceed to More Actions -> Configure Cluster Quorum Settings.

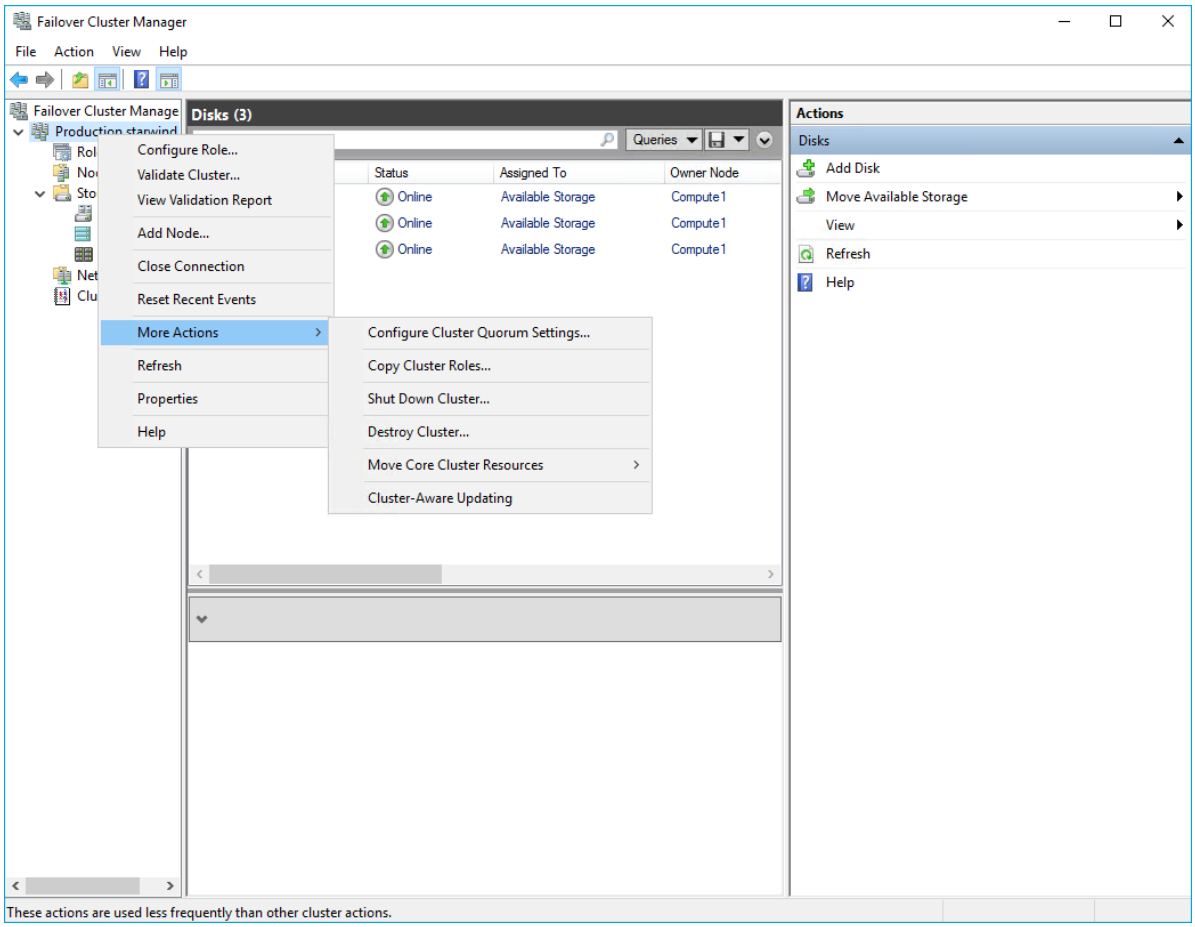

3. Follow the wizard and choose the Select the quorum witness option. Click Next.

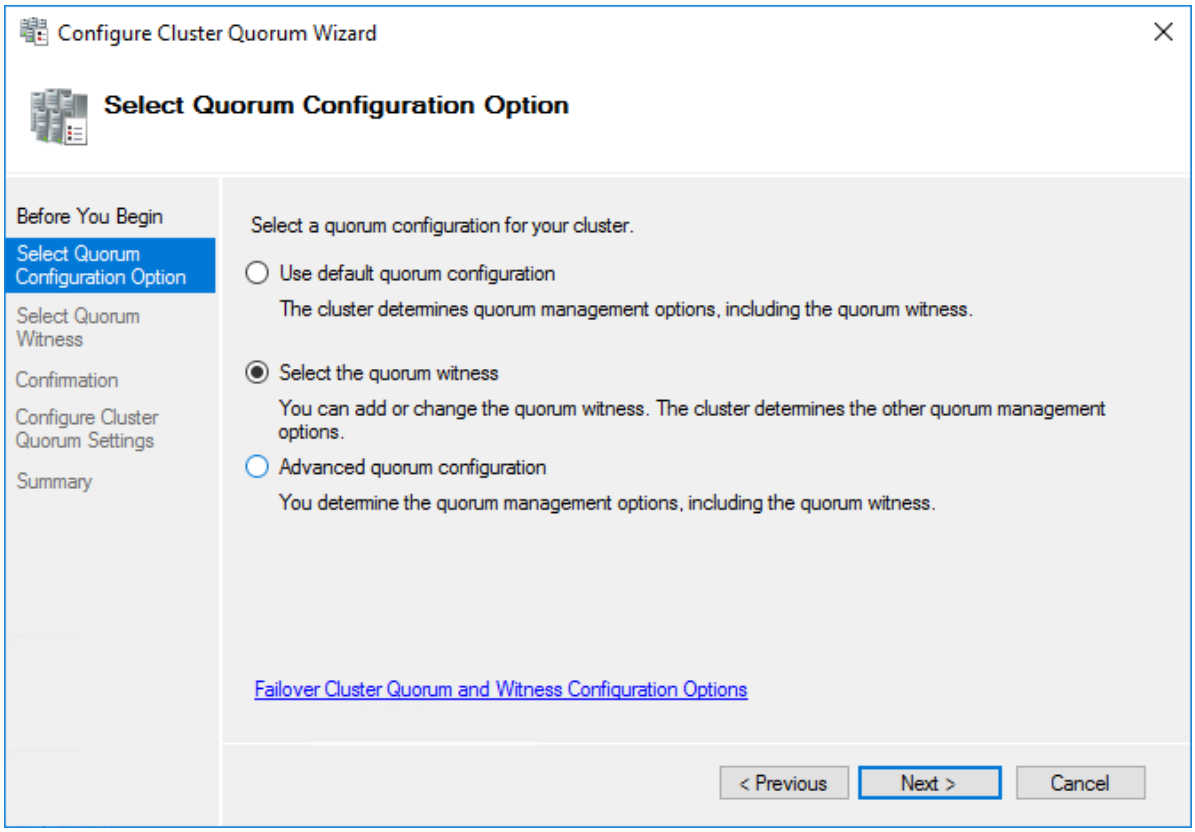

4. Select Configure a disk witness. Click Next.

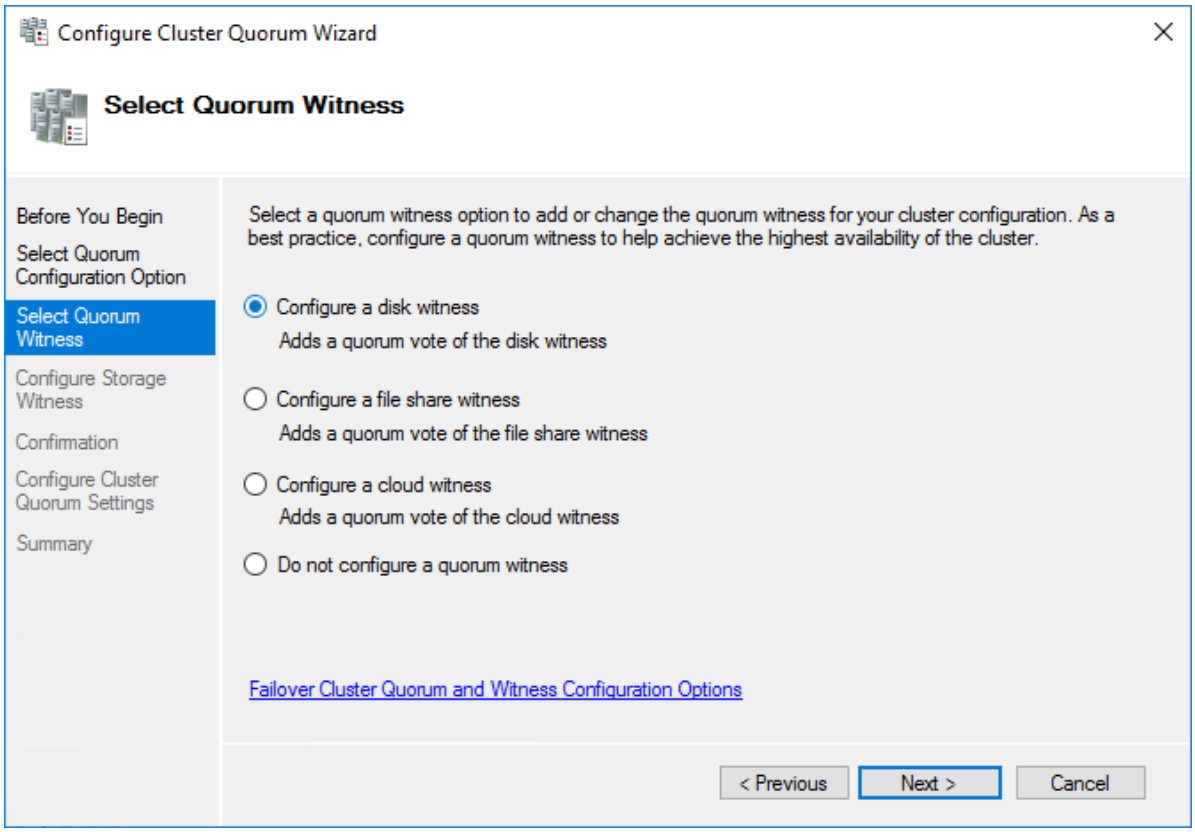

5. Select the Witness disk to be assigned as the cluster witness disk. Click Next and press Finish to complete the operation.

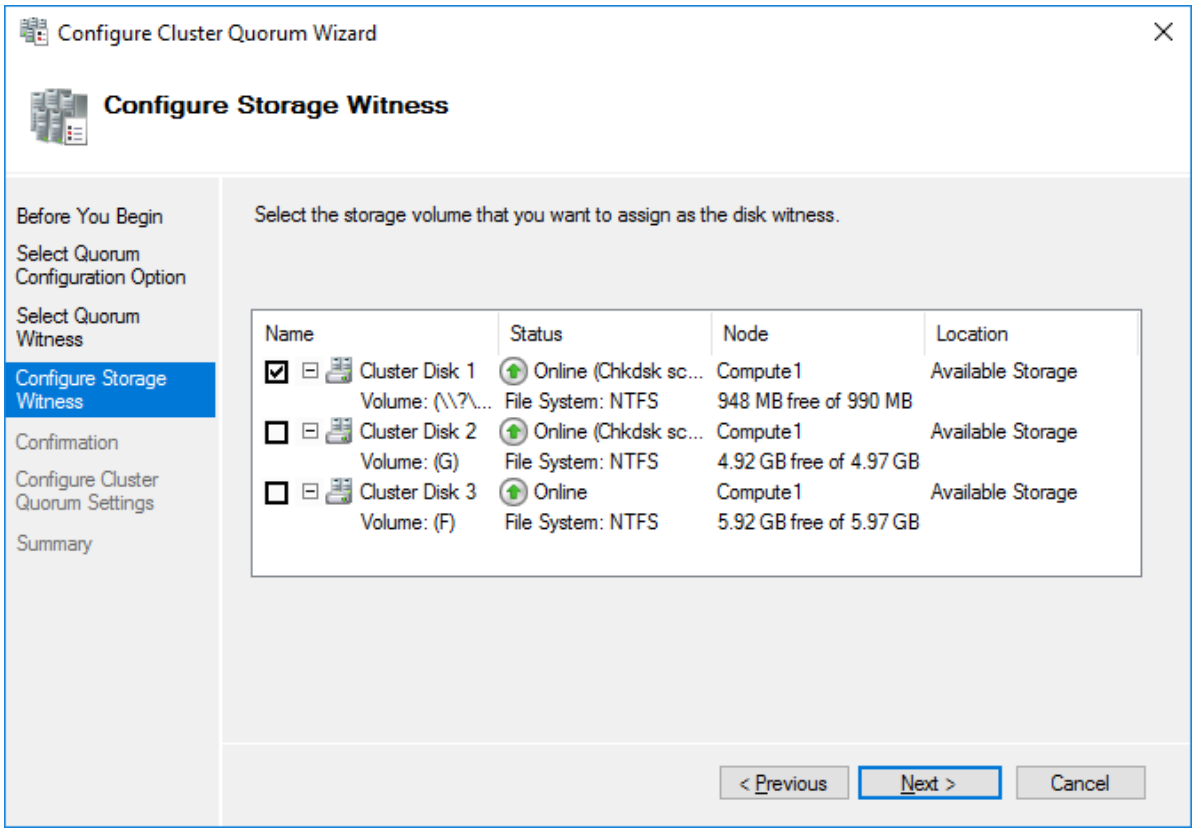

6. In Failover Cluster Manager, right-click the disk and select Add to Cluster Shared Volumes.

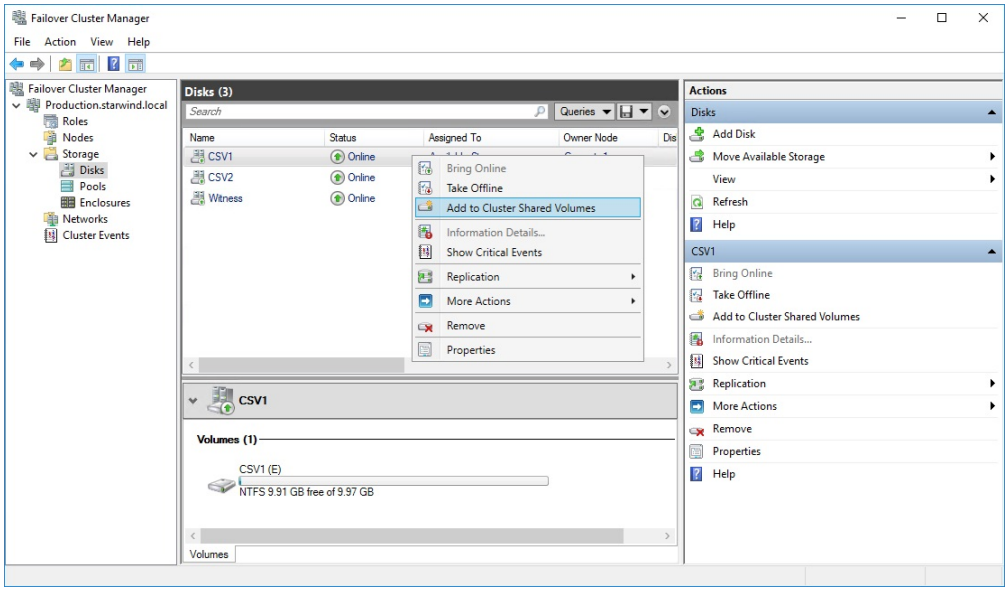

7. If renaming of the cluster shared volume is required, right-click on the disk and select Properties. Type the new name for the disk and click Apply followed by OK.

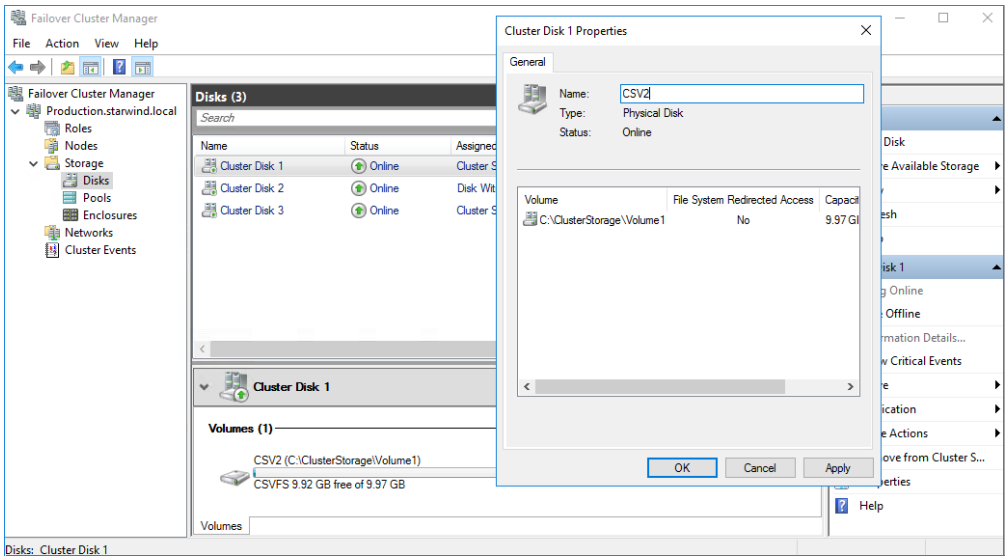

8. Perform the steps 6-7 for any other disk in Failover Cluster Manager. The resulting list of disks will look similar to the screenshot below.

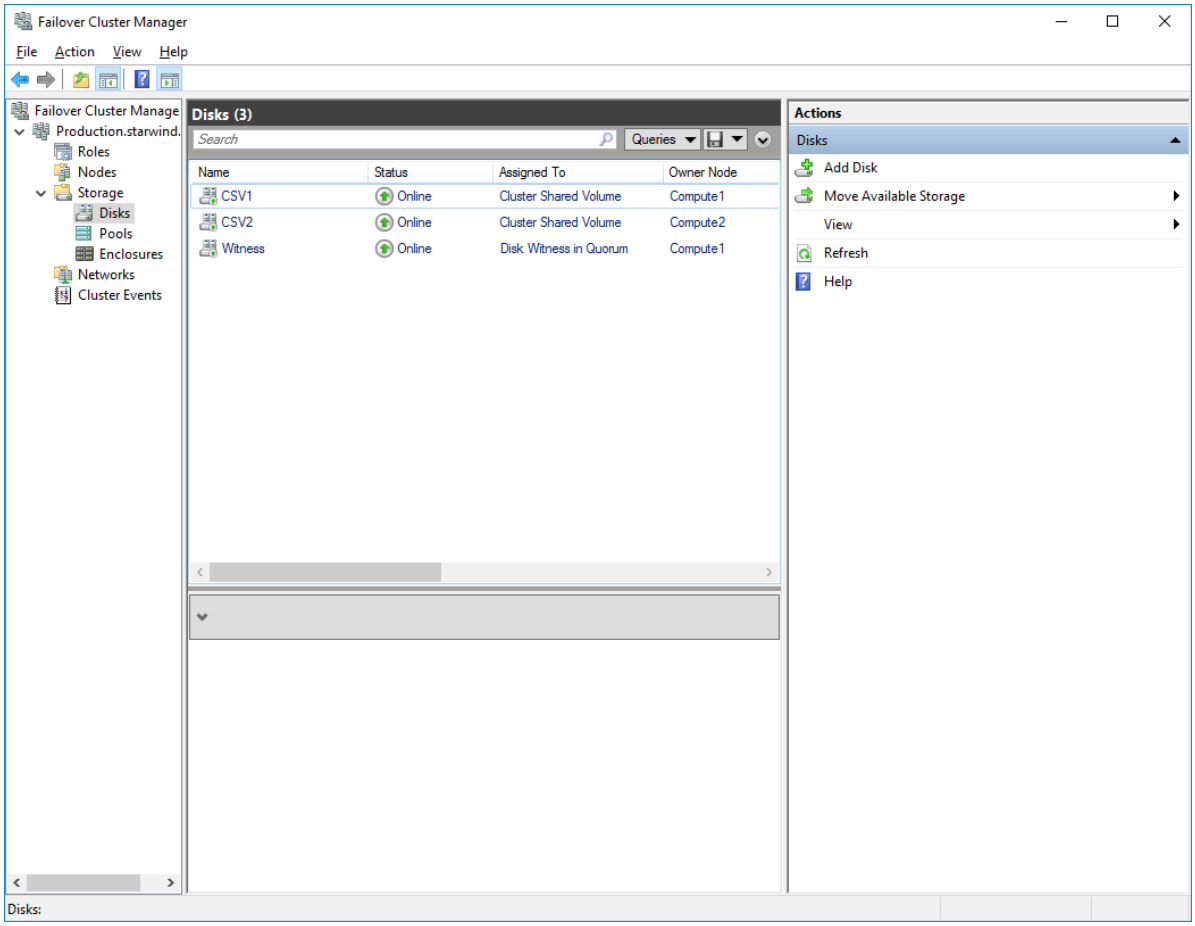

Configuring Cluster Network Preferences

1. In the Networks section of the Failover Cluster Manager, right-click on the network from the list. If required, set its new name to identify the network by its subnet. Apply the change and press OK.
NOTE: Please double-check that cluster communication is configured with redundant networks:
https://docs.microsoft.com/en-us/windows-server/failover-clustering/smb-multichannel

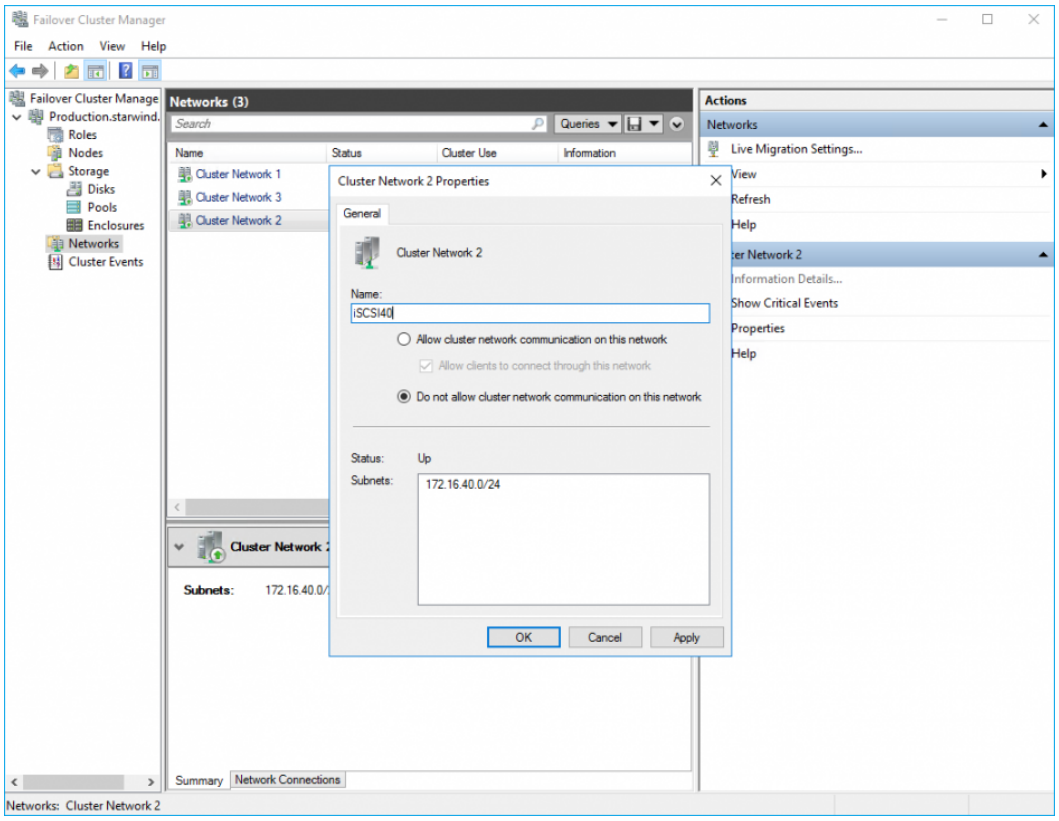

2. Rename other networks as described above, if required.

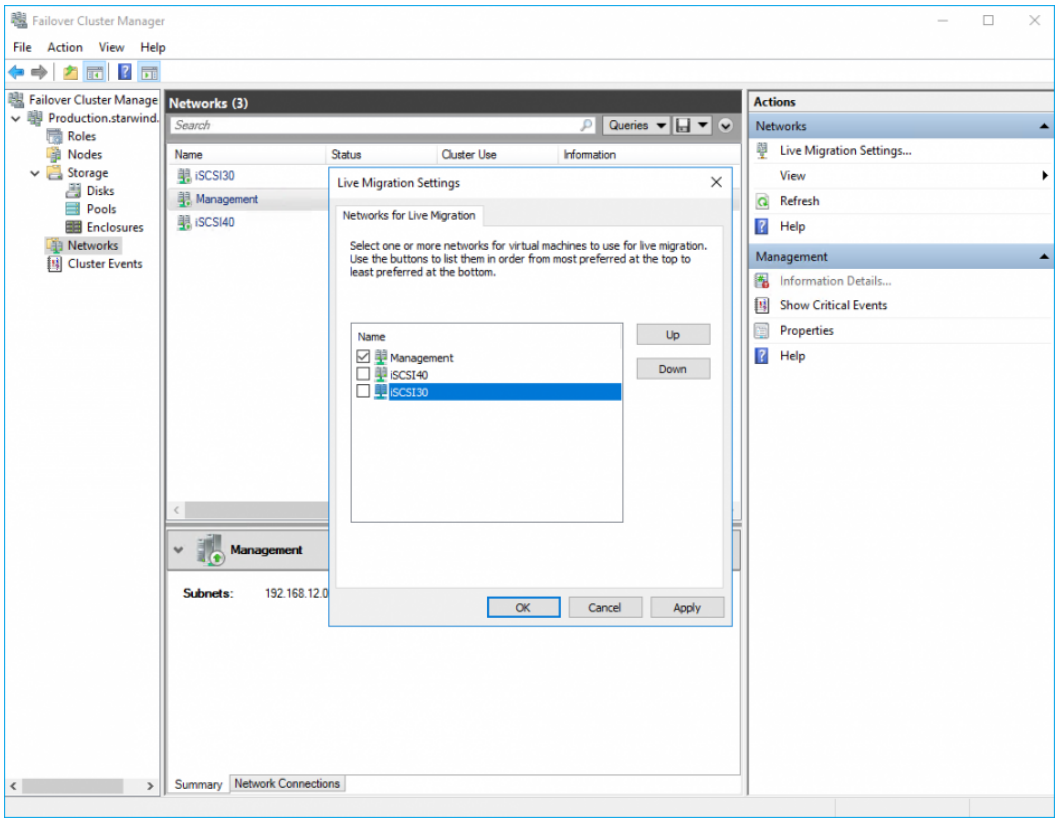

3. In the Actions tab, click Live Migration Settings. Uncheck the synchronization network, while the iSCSI network can be used if it is 10+ Gbps. Apply the changes and click OK.

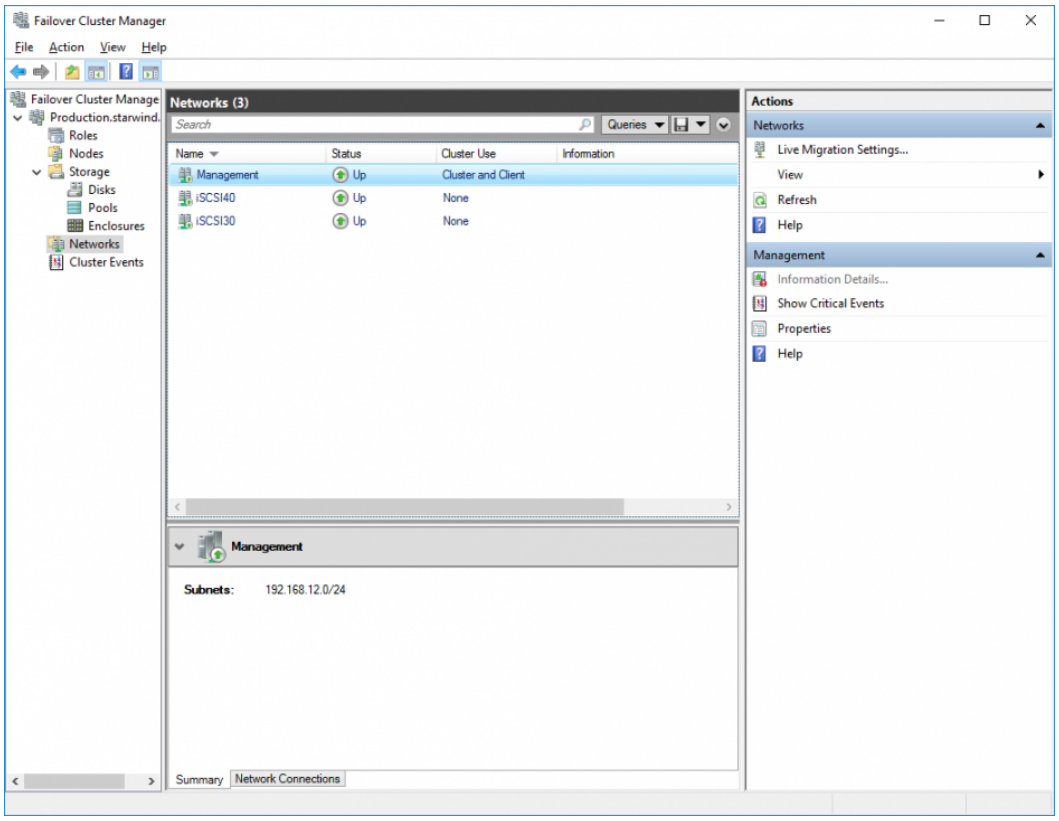

The cluster configuration is completed and it is ready for virtual machines deployment. Select Roles and in the Action tab, click Virtual Machines -> New Virtual Machine. Complete the wizard.

Configuring File Shares

Please follow the steps below if file shares should be configured on cluster nodes.

Configuring The Scale-Out File Server Role

1. To configure the Scale-Out File Server Role, open Failover Cluster Manager.

2. Right-click the cluster name, then click Configure Role and click Next to continue.

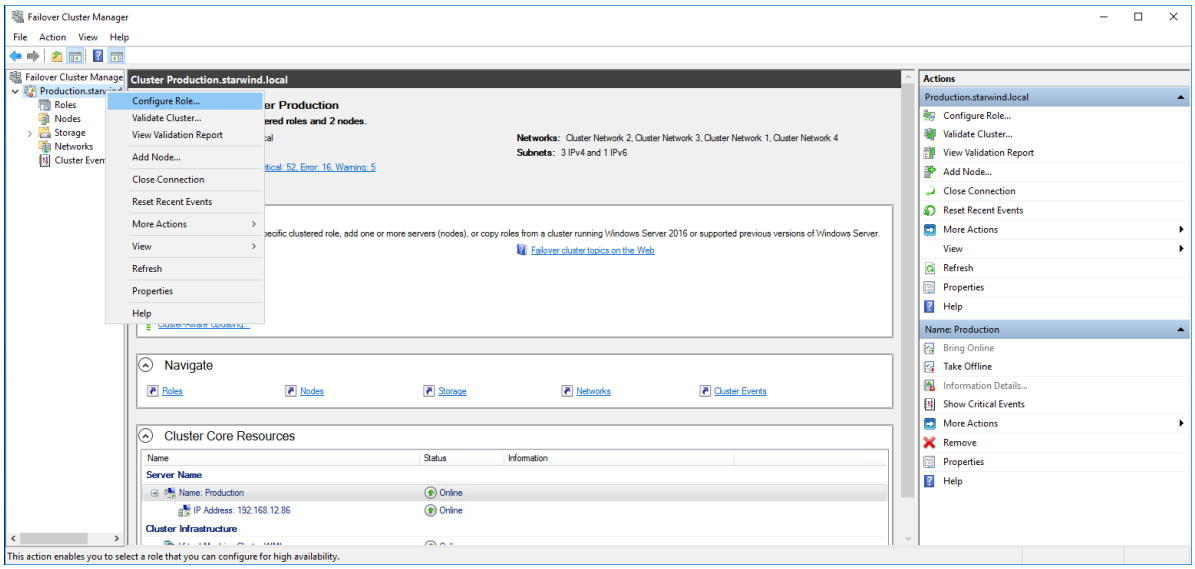

3. Select the File Server item from the list in High Availability Wizard and click Next to continue.

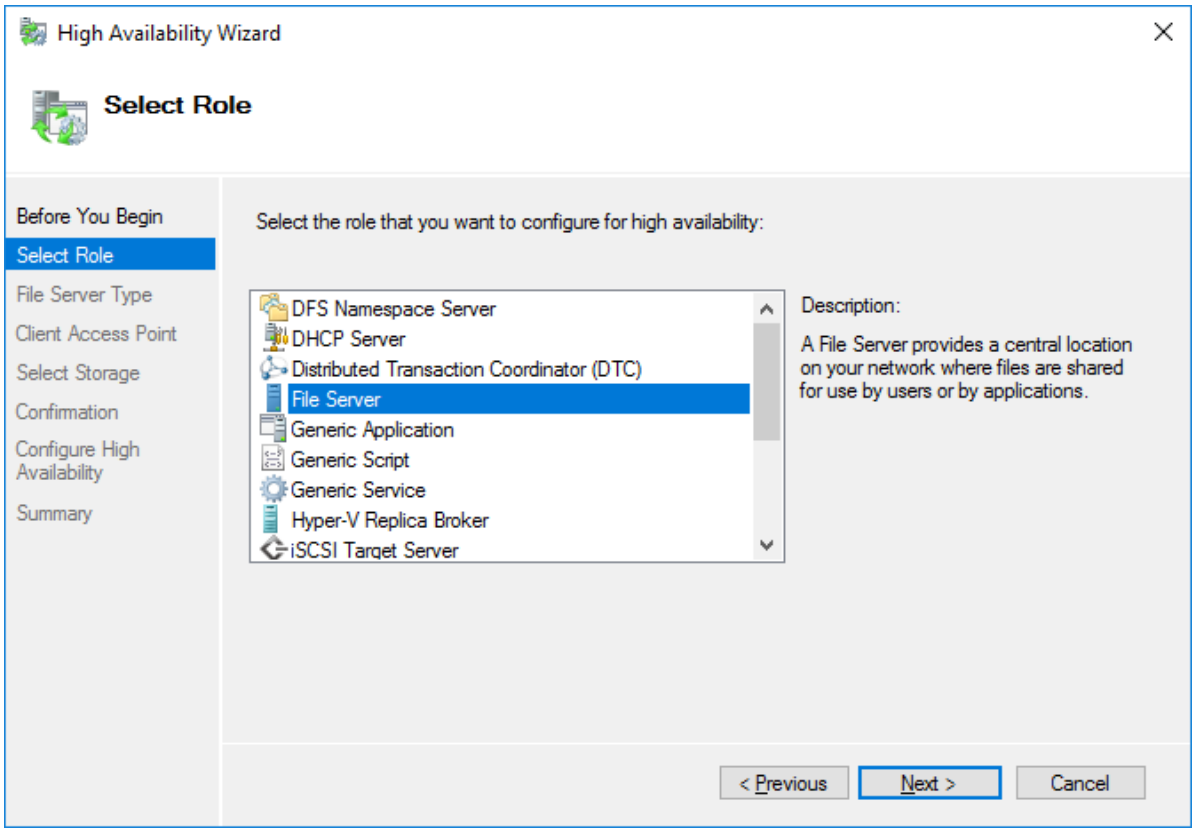

4. Select Scale-Out File Server for application data and click Next.

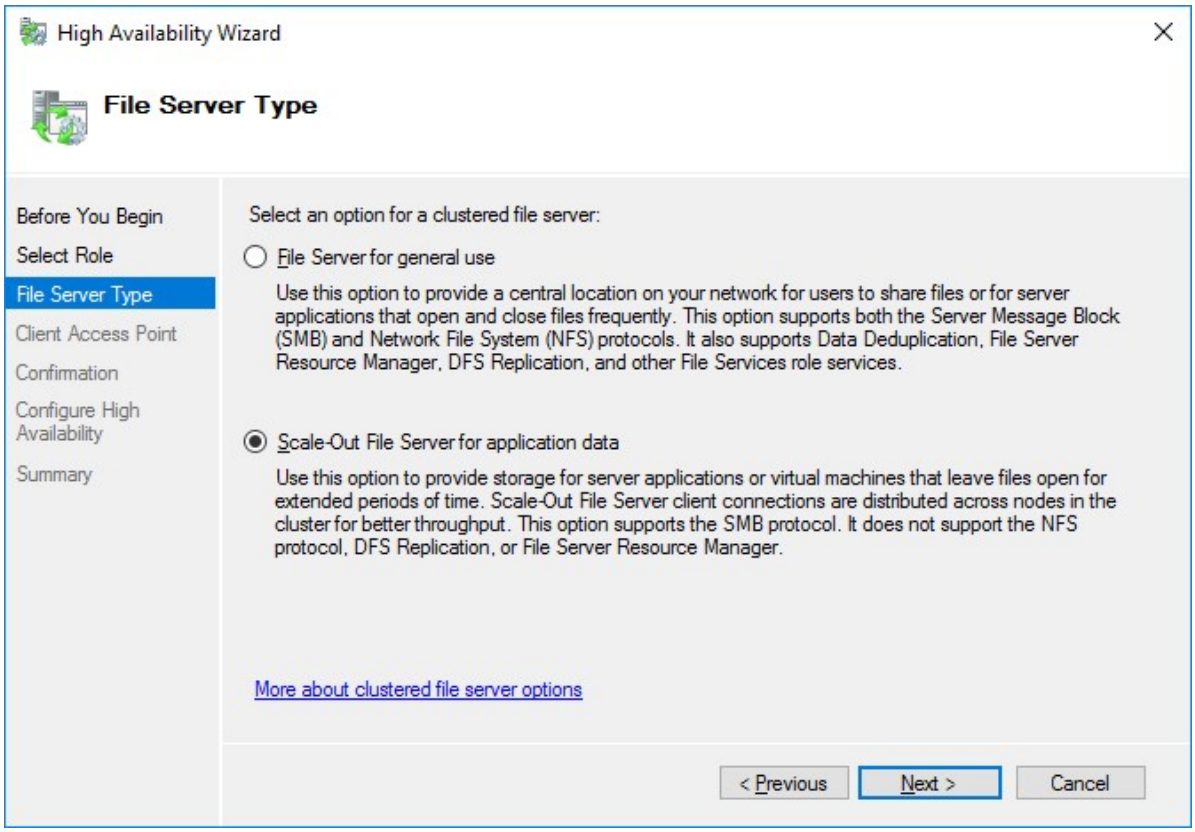

5. On the Client Access Point page, in the Name text field, type the NetBIOS name that will be used to access a Scale-Out File Server.

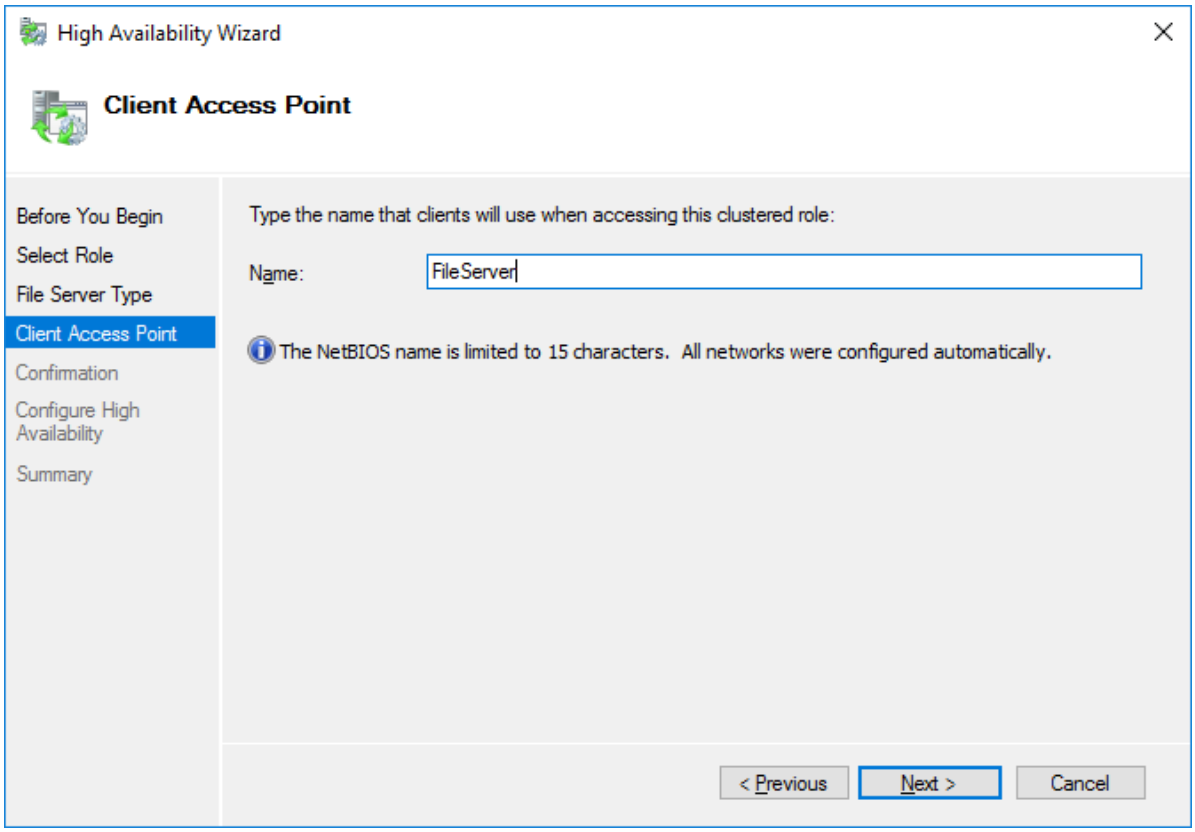

Click Next to continue.

6. Check whether the specified information is correct. Click Next to continue or Previous to change the settings.

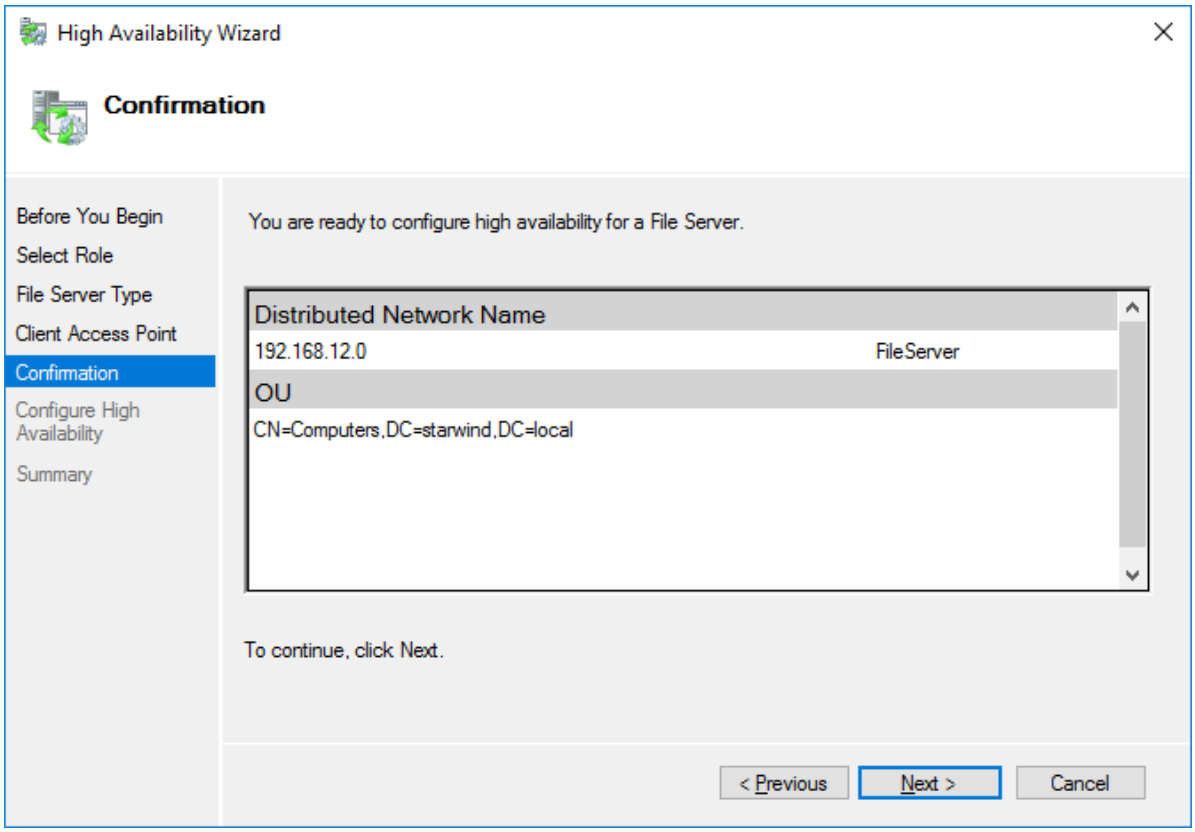

7. Once the installation is finished successfully, the Wizard should now look like the screenshot below.
Click Finish to close the Wizard.

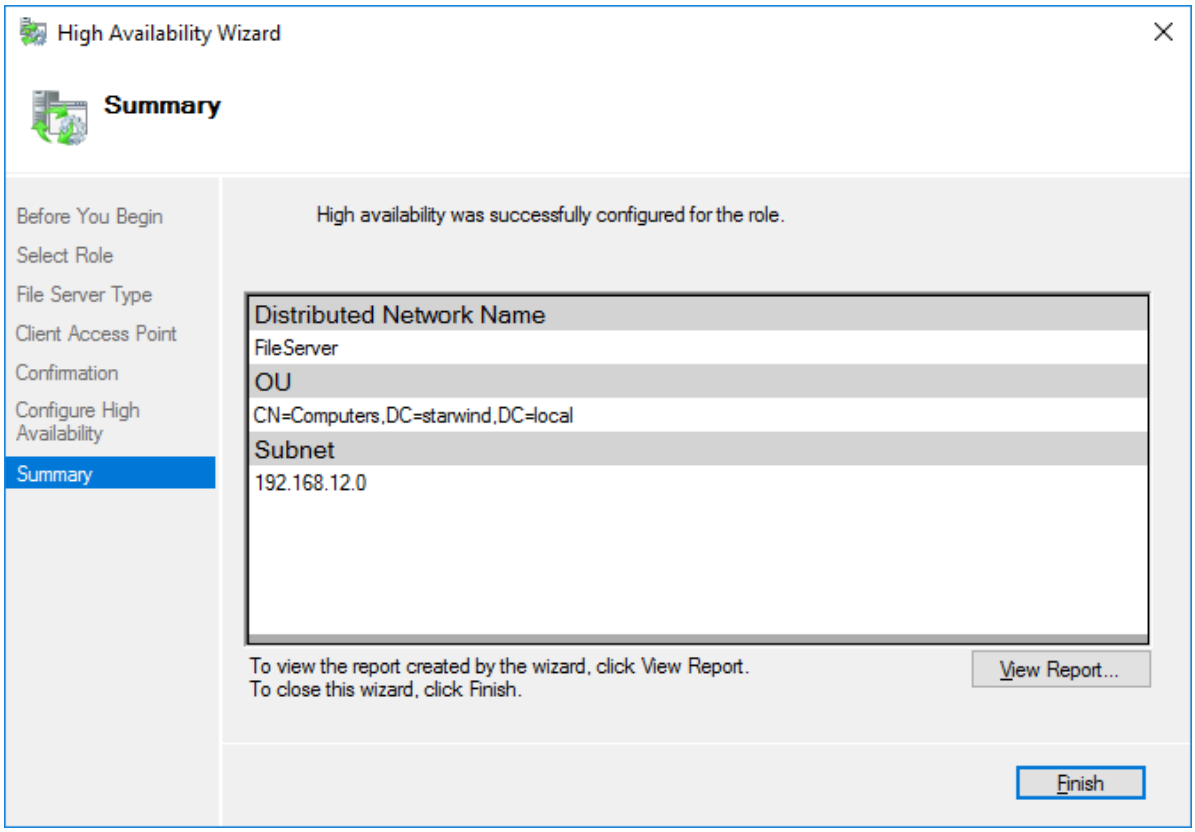

8. The newly created role should now look like the screenshot below.

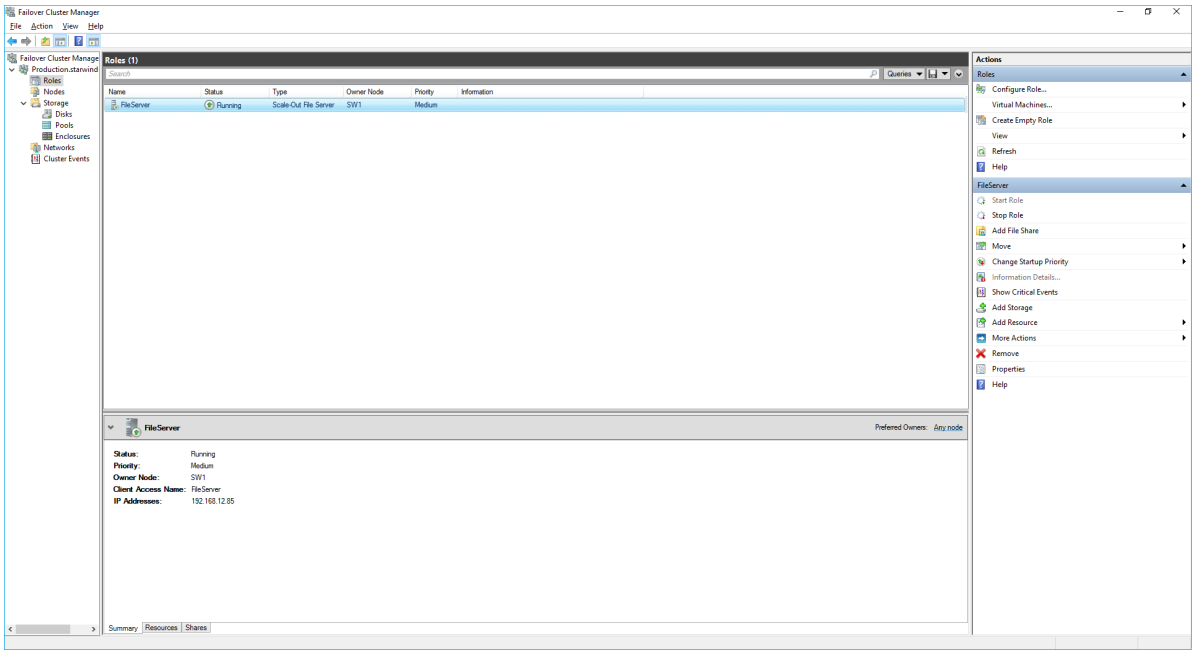

NOTE: If the role status is Failed and it is unable to Start, please, follow the next steps:

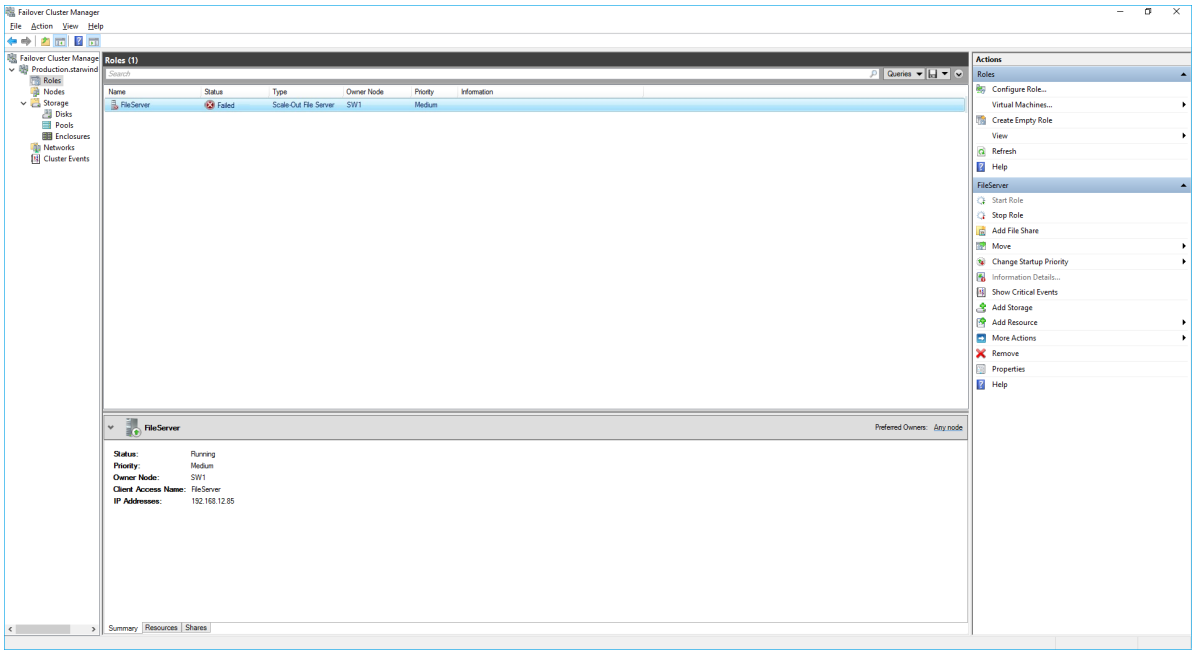

- open Active Directory Users and Computers
- enable the Advanced view if it is not enabled
- edit the properties of the OU containing the cluster computer object (in this case – Production)
- open the Security tab and click Advanced
- in the appeared window, press Add (the Permission Entry dialog box opens), click Select a principal
- in the appeared window, click Object Types, select Computers, and click OK
- enter the name of the cluster computer object (in this case – Production)

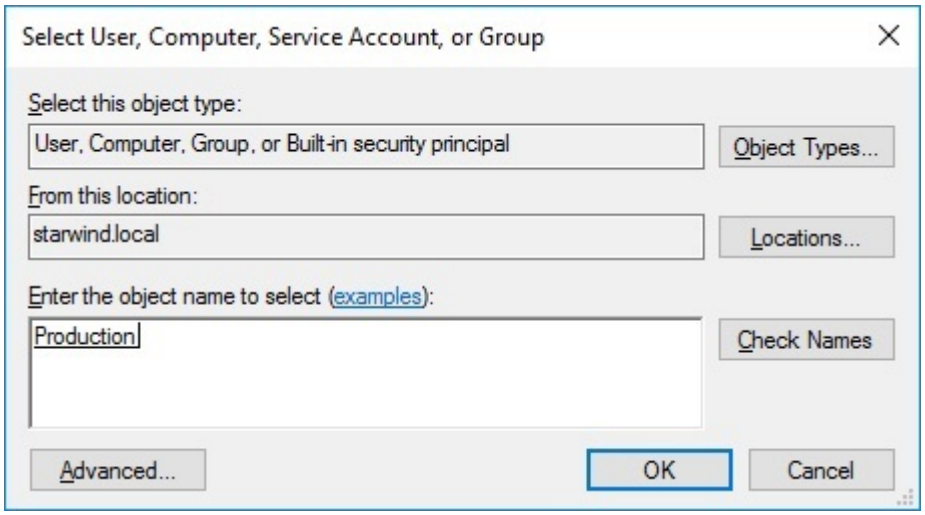

- go back to Permission Entry dialog, scroll down, and select Create Computer Objects,

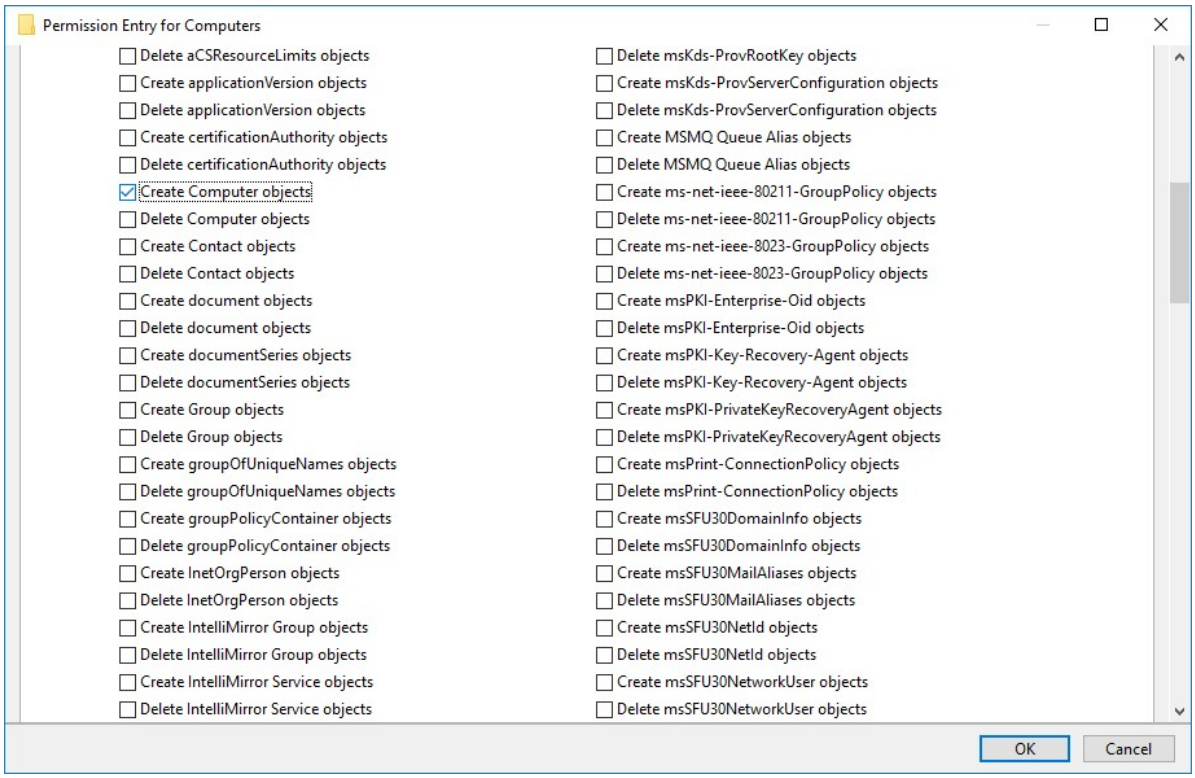

- click OK on all opened windows to confirm the changes
- open Failover Cluster Manager, right-click SOFS role and click Start Role

Configuring File Share

To Add File Share:

- open Failover Cluster Manager
- expand the cluster and then click Roles
- right-click the file server role and then press Add File Share
- on the Select the profile for this share page, click SMB Share – Applications and then click Next

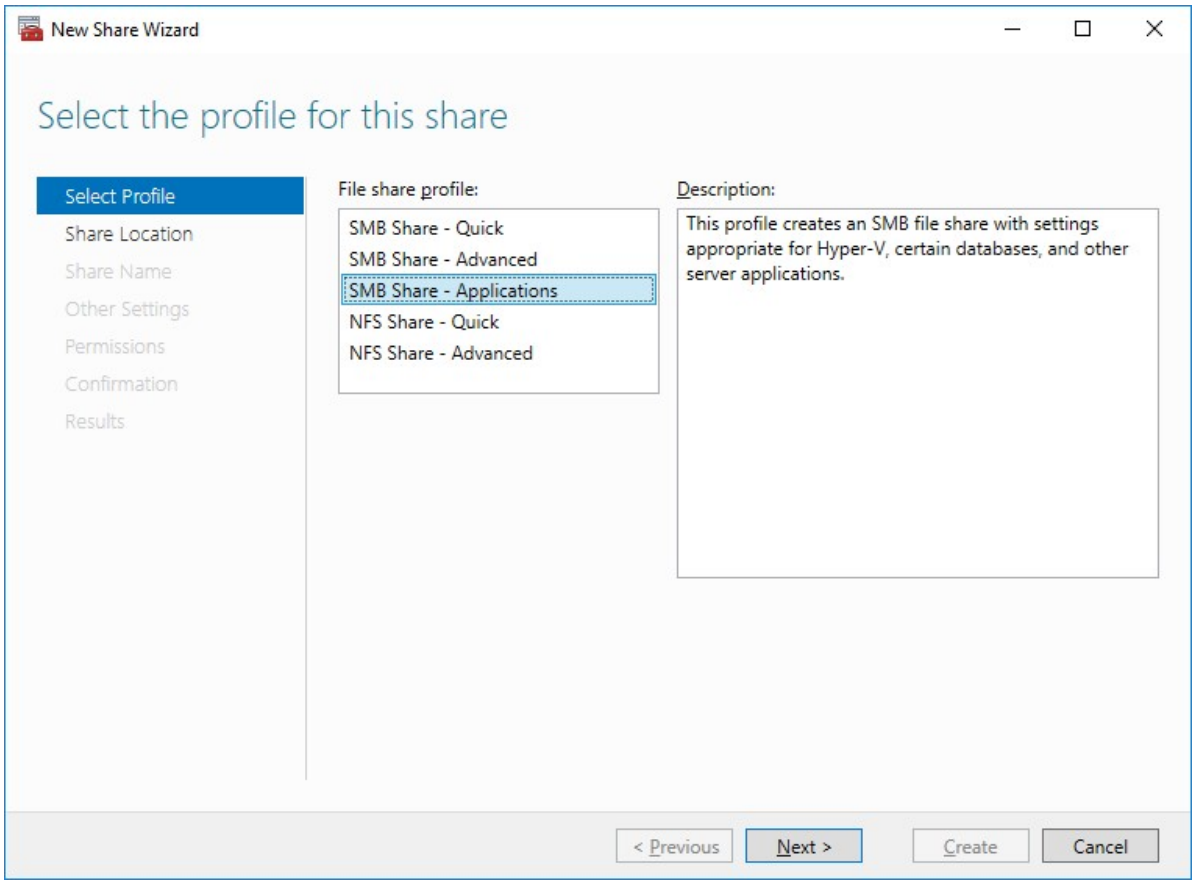

5. Select a CSV to host the share. Click Next to proceed.

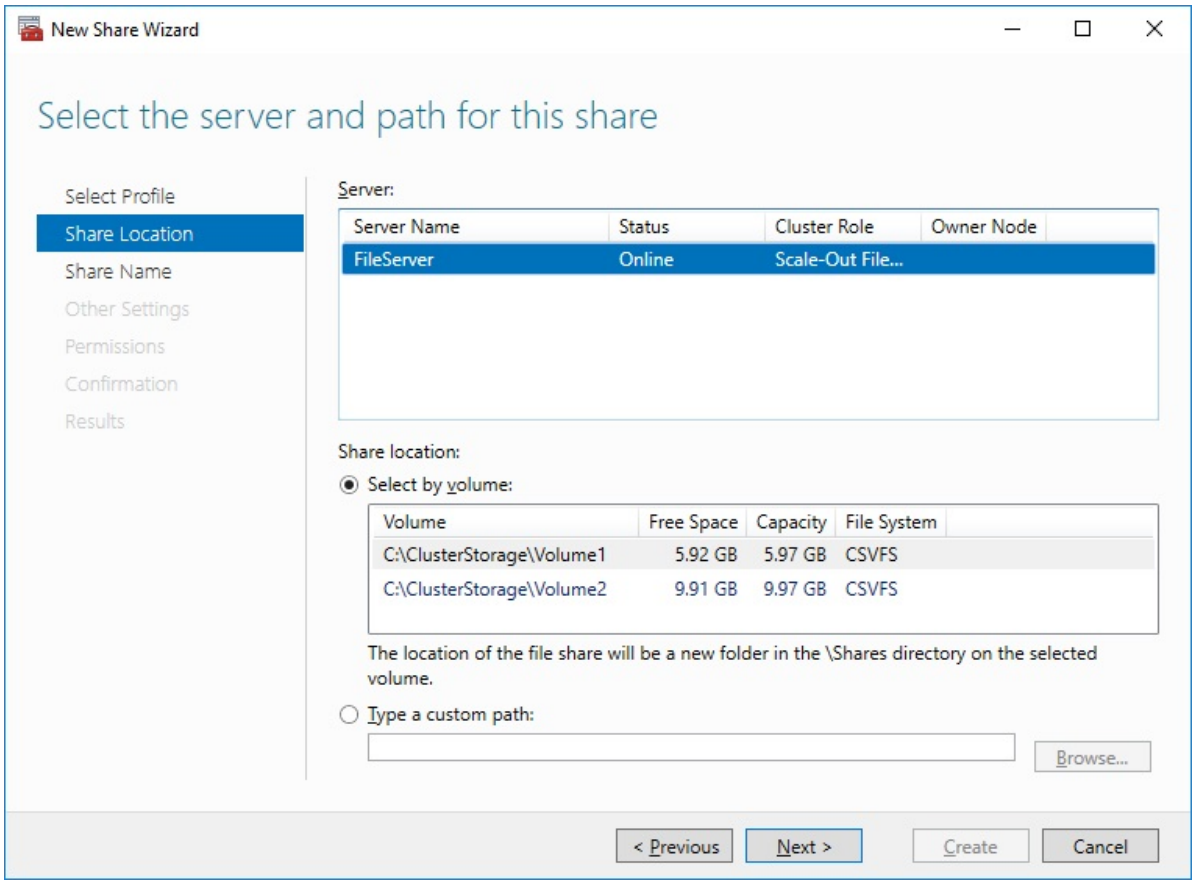

6. Type in the file share name and click Next.

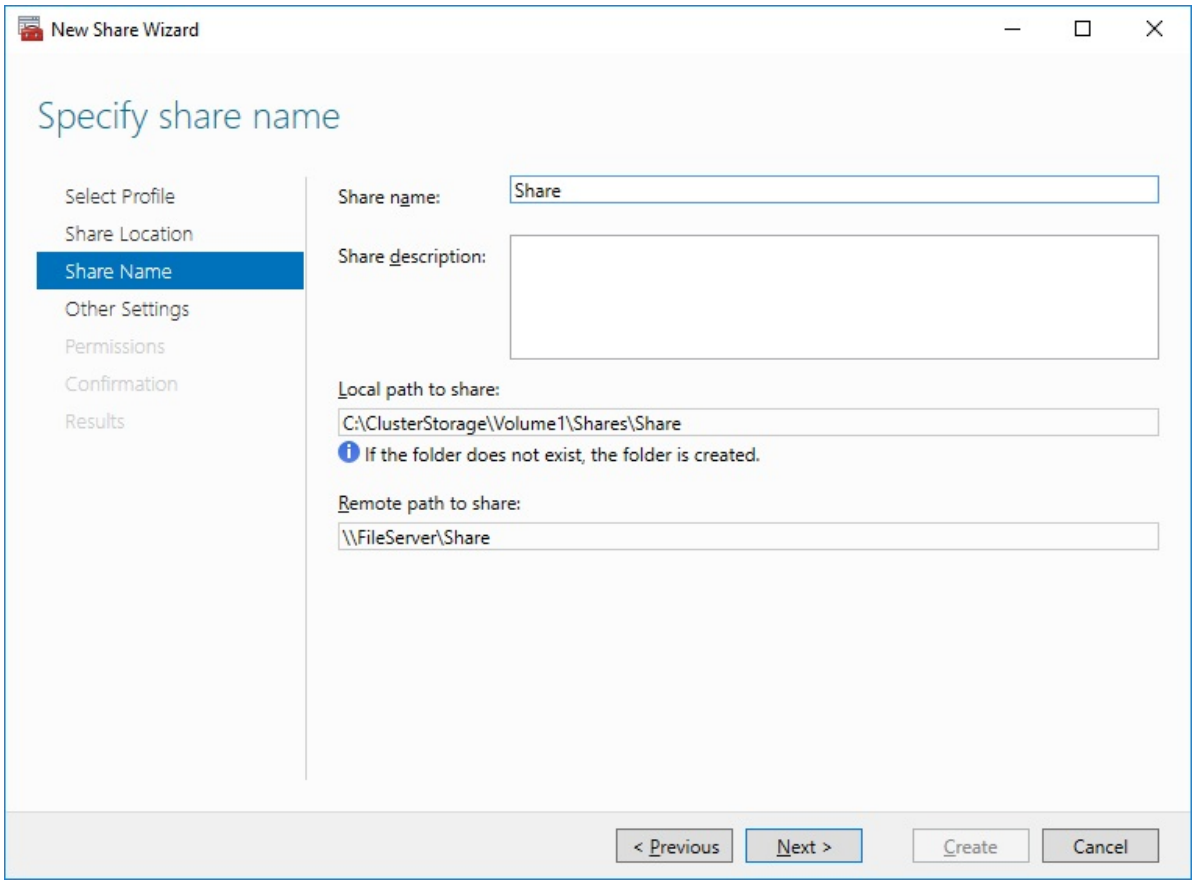

7. Make sure that the Enable Continuous Availability box is checked. Click Next to proceed.

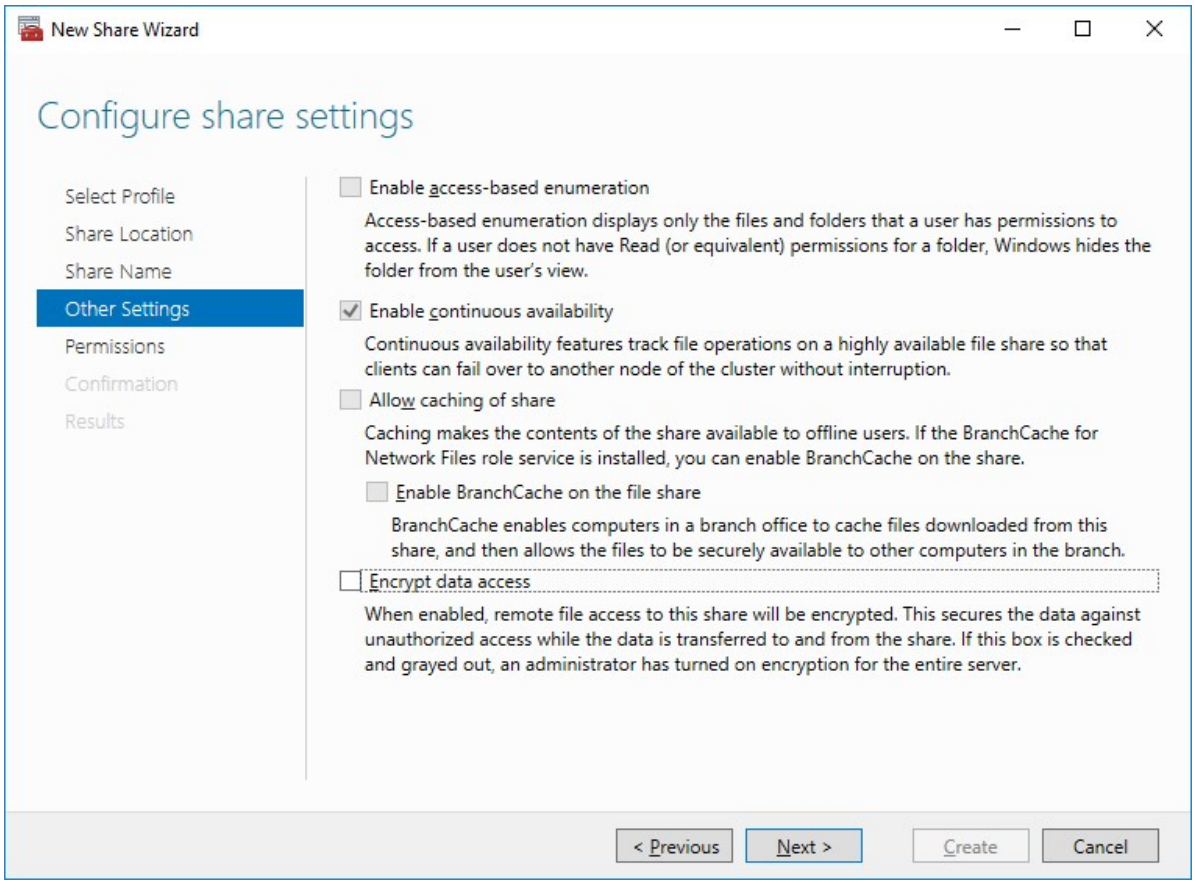

8. Specify the access permissions for the file share.

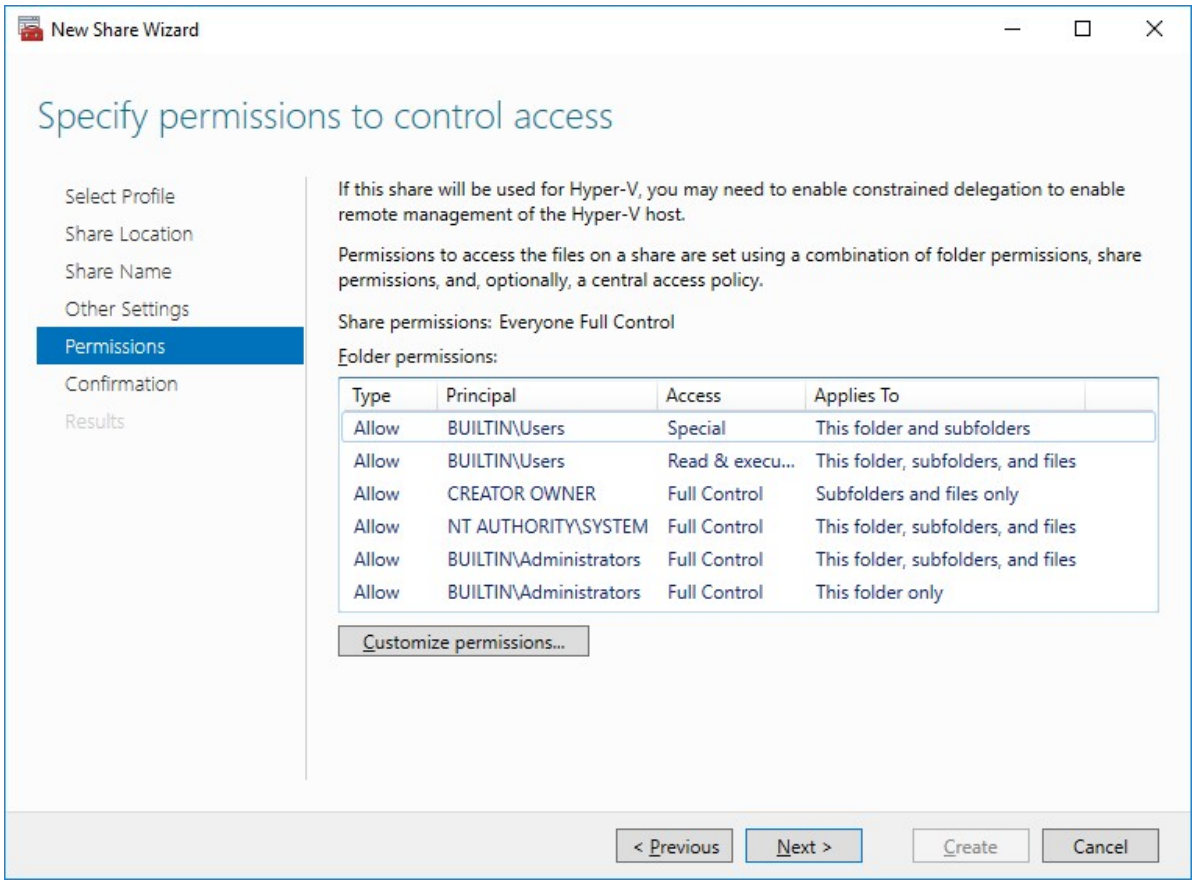

NOTE:

- for the Scale-Out File Server for Hyper-V, all Hyper-V computer accounts, the SYSTEM account, and all Hyper-V administrators must be provided with the full control on the share and file system
- for the Scale-Out File Server on Microsoft SQL Server, the SQL Server service account must be granted full control on the share and the file system

9. Check whether specified settings are correct. Click Previous to make any changes or click Create to proceed.

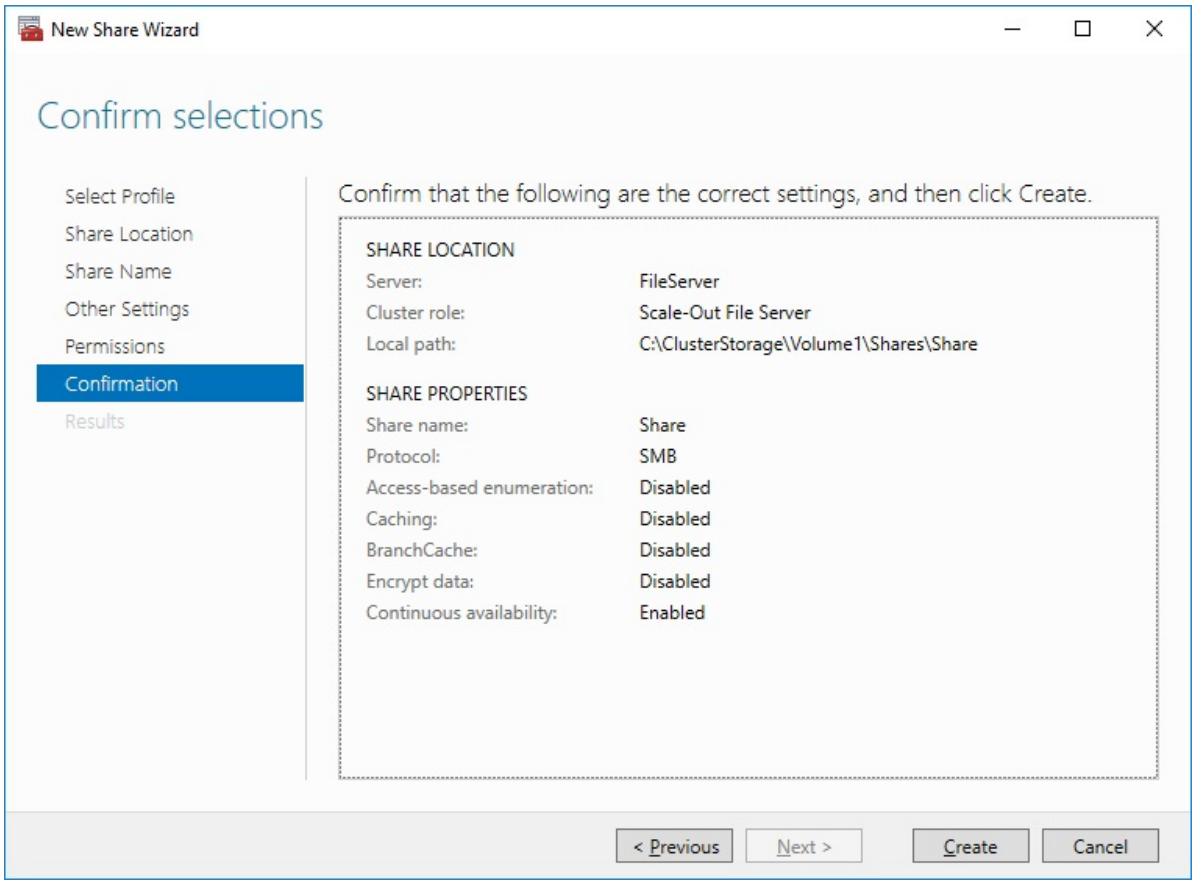

10. Check the summary and click Close to close the Wizard.

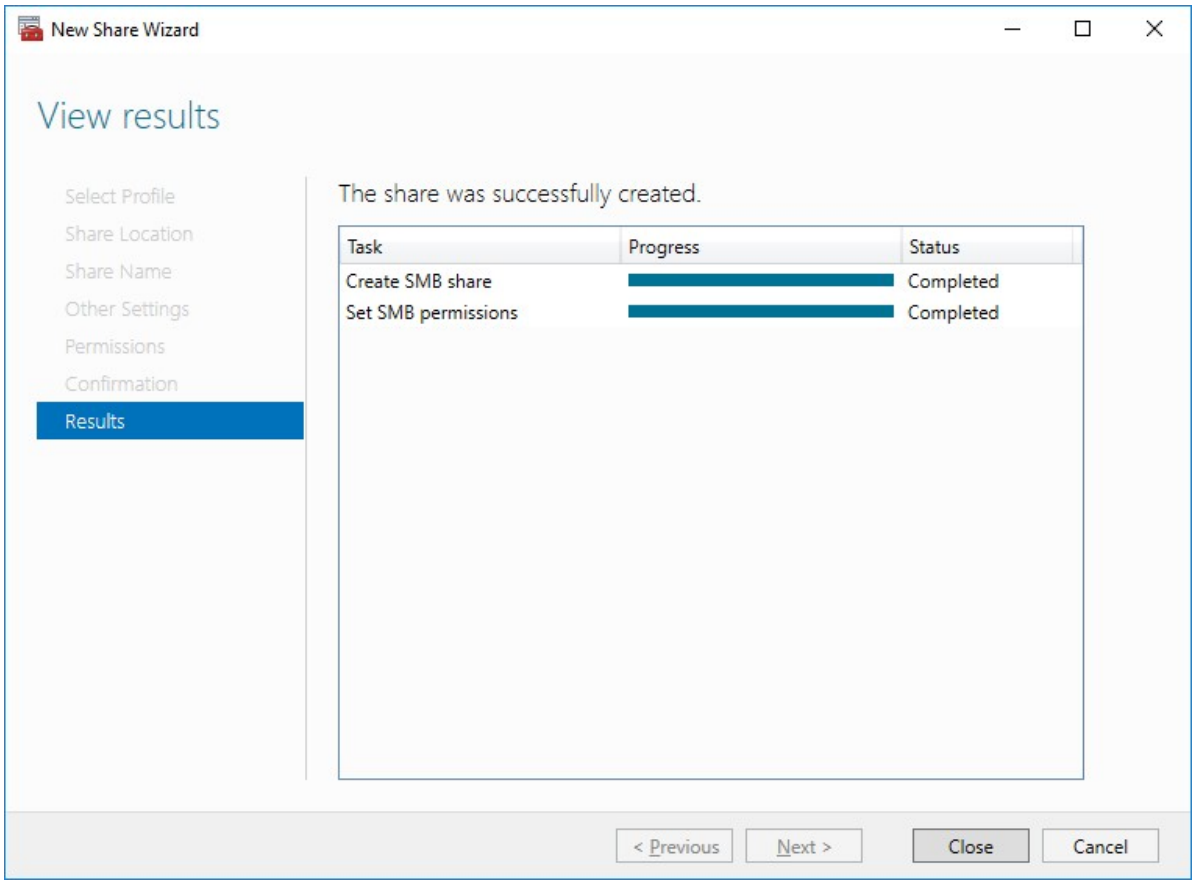

To Manage Created File Shares:

- open Failover Cluster Manager
- expand the cluster and click Roles
- choose the file share role, select the Shares tab, right-click the created file share, and select Properties:

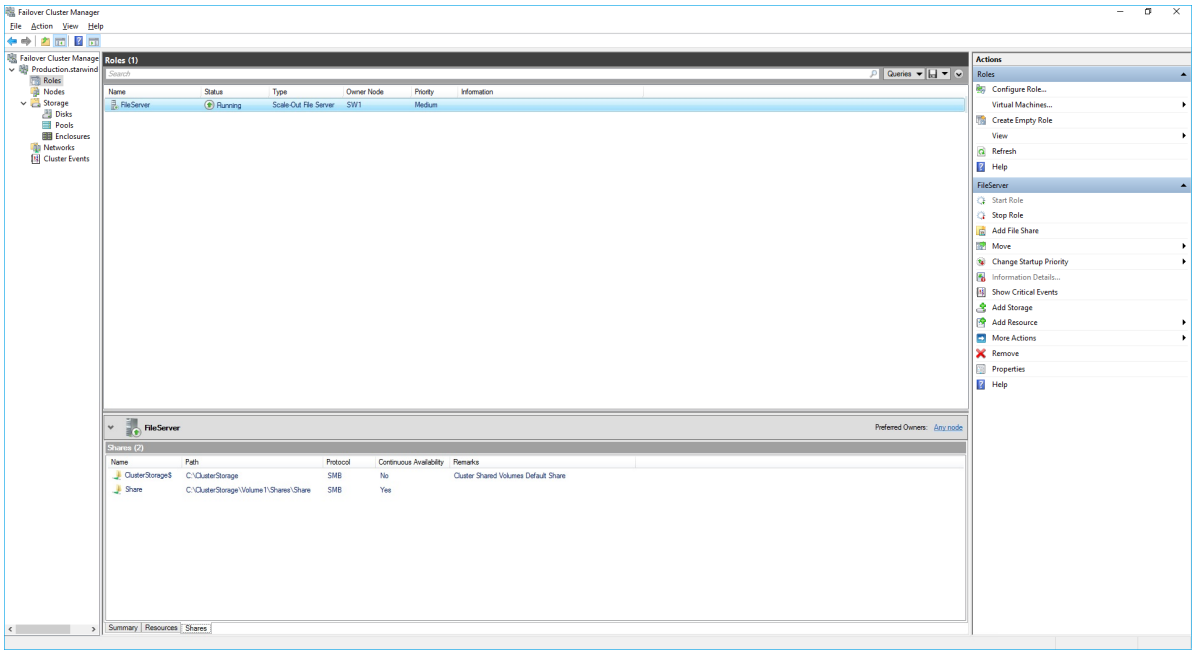

Configuring The File Server For General Use Role

NOTE: To configure File Server for General Use, the cluster should have available storage

1. To configure the File Server for General Use role, open Failover Cluster Manager.

2. Right-click on the cluster name, then click Configure Role and click Next to continue.

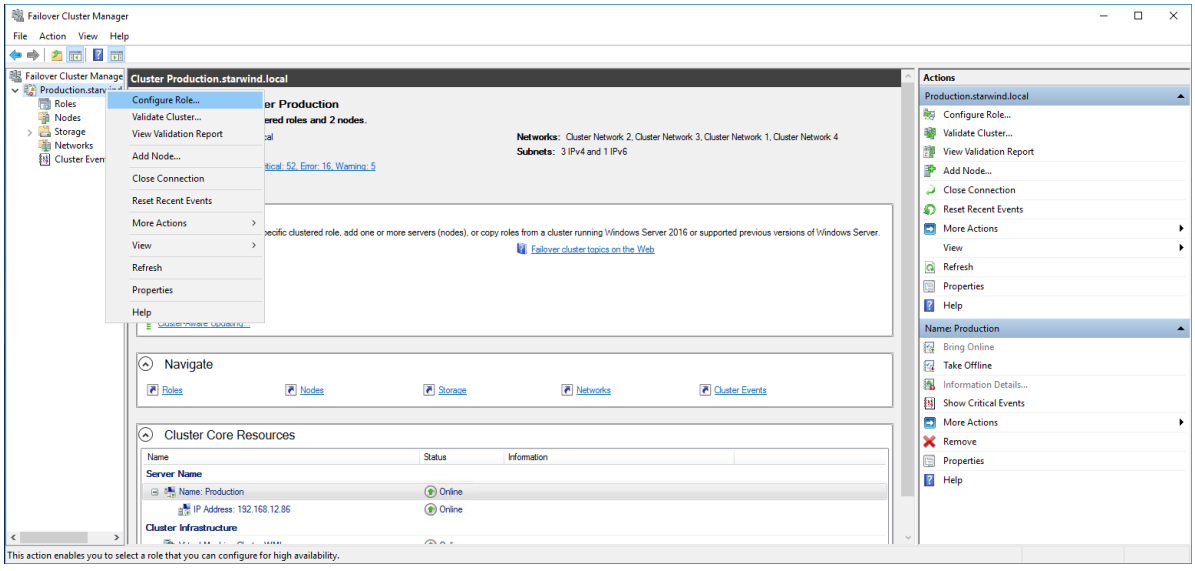

3. Select the File Server item from the list in High Availability Wizard and click Next to continue.

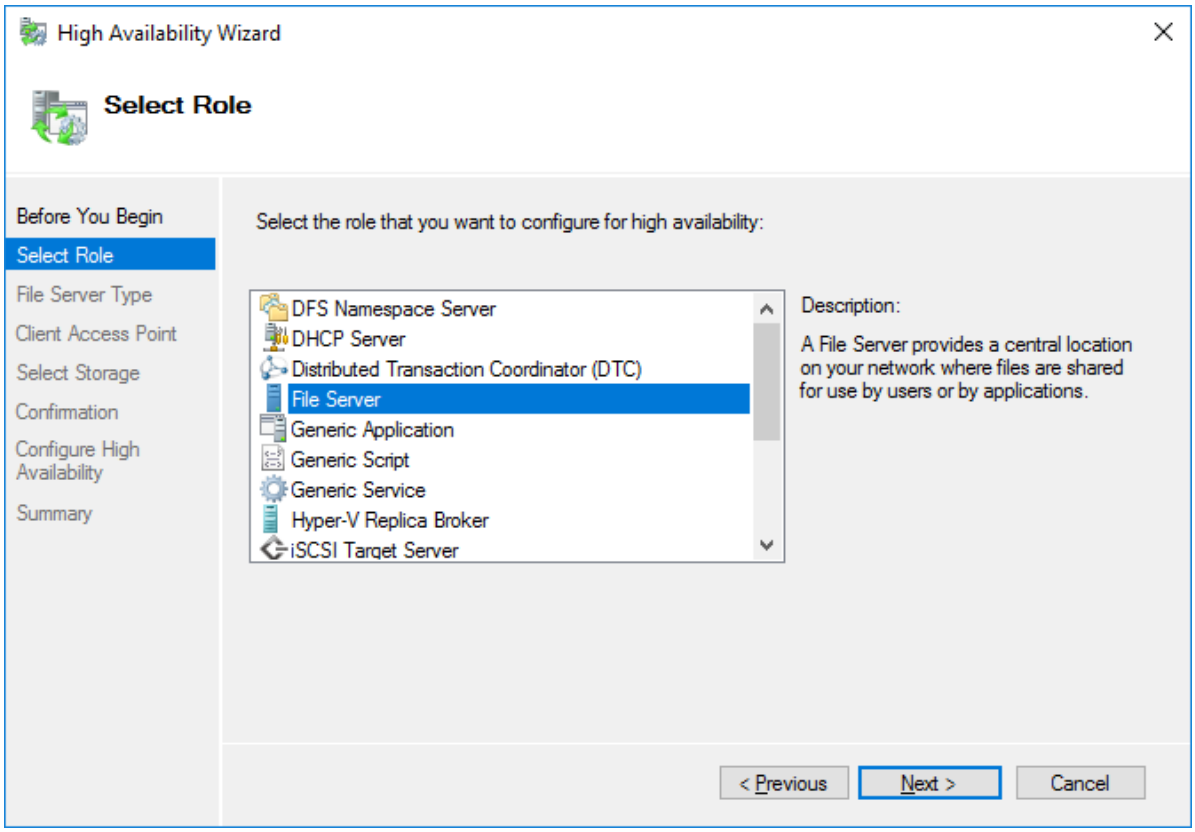

4. Select File Server for general use and click Next.

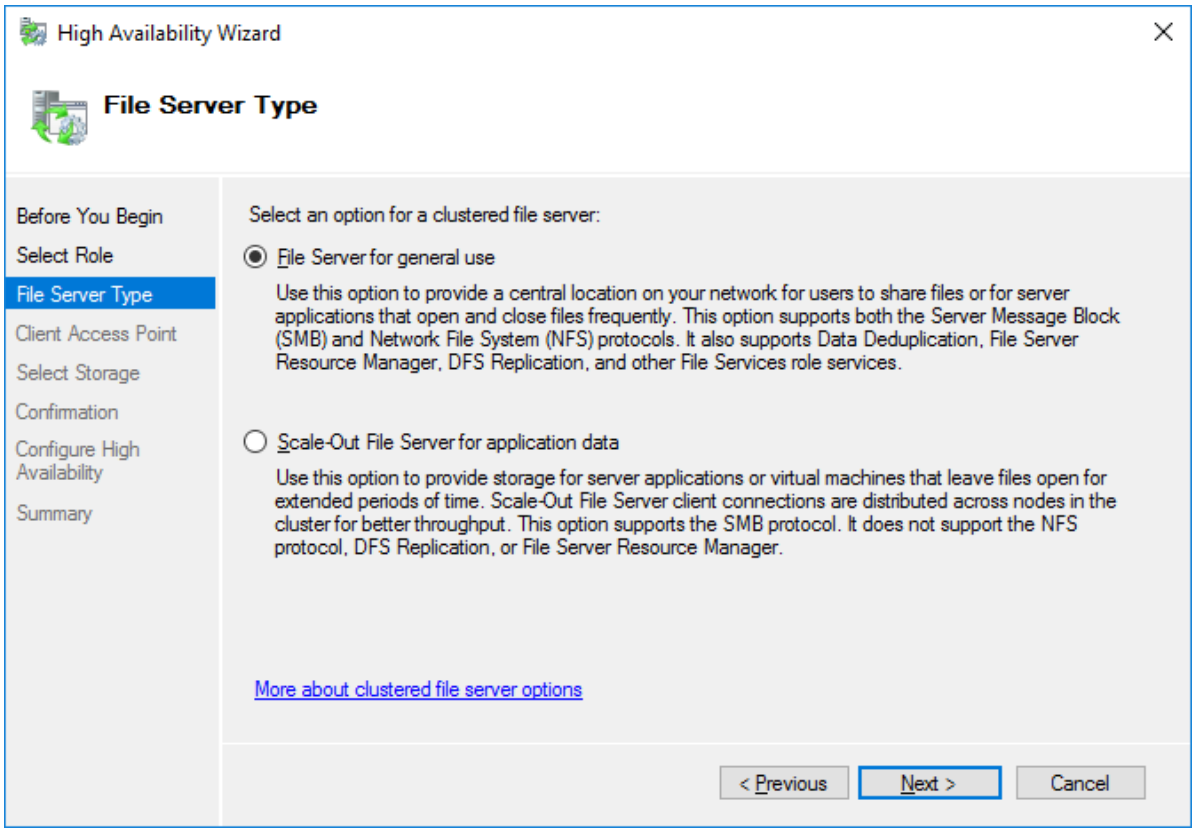

5. On the Client Access Point page, in the Name text field, type the NETBIOS name that will be used to access the File Server and IP for it.

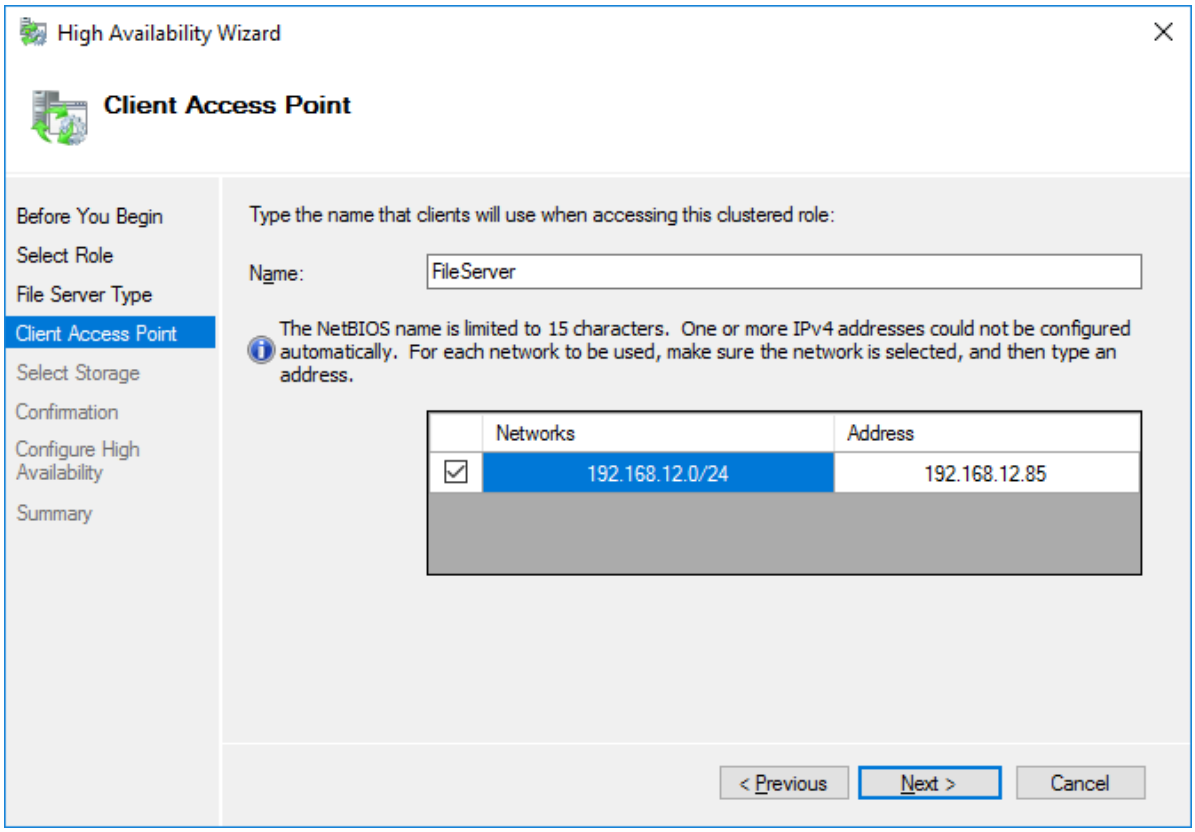

Click Next to continue.

6. Select the Cluster disk and click Next.

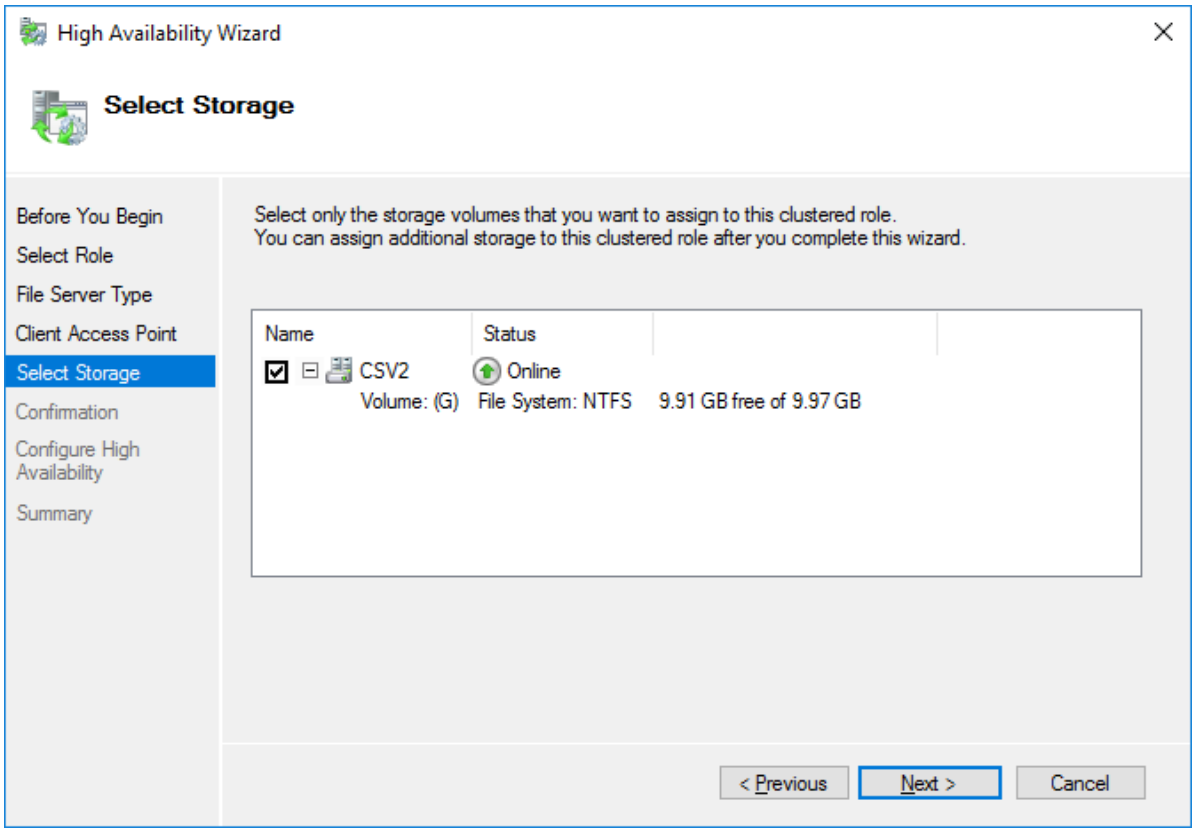

7. Check whether the specified information is correct. Click Next to proceed or Previous to change the settings.

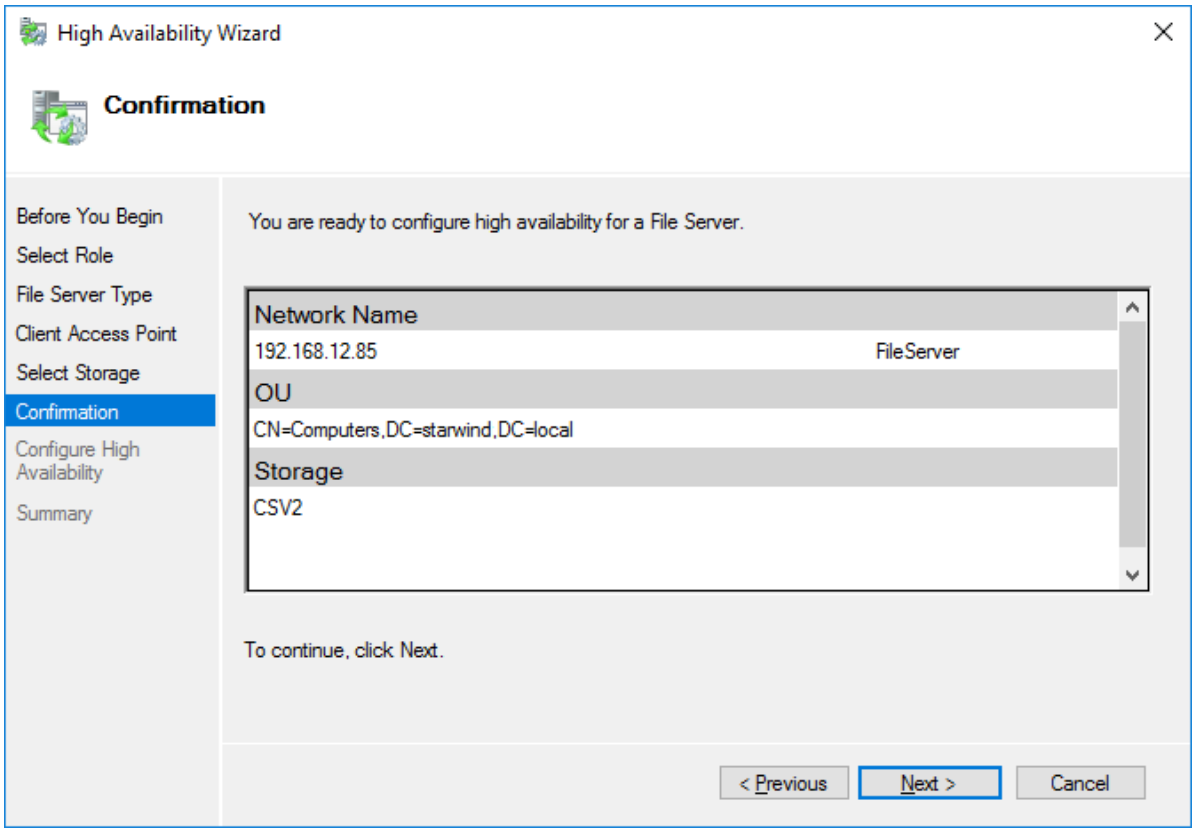

8. Once the installation has been finished successfully, the Wizard should now look like the screenshot below.

Click Finish to close the Wizard.

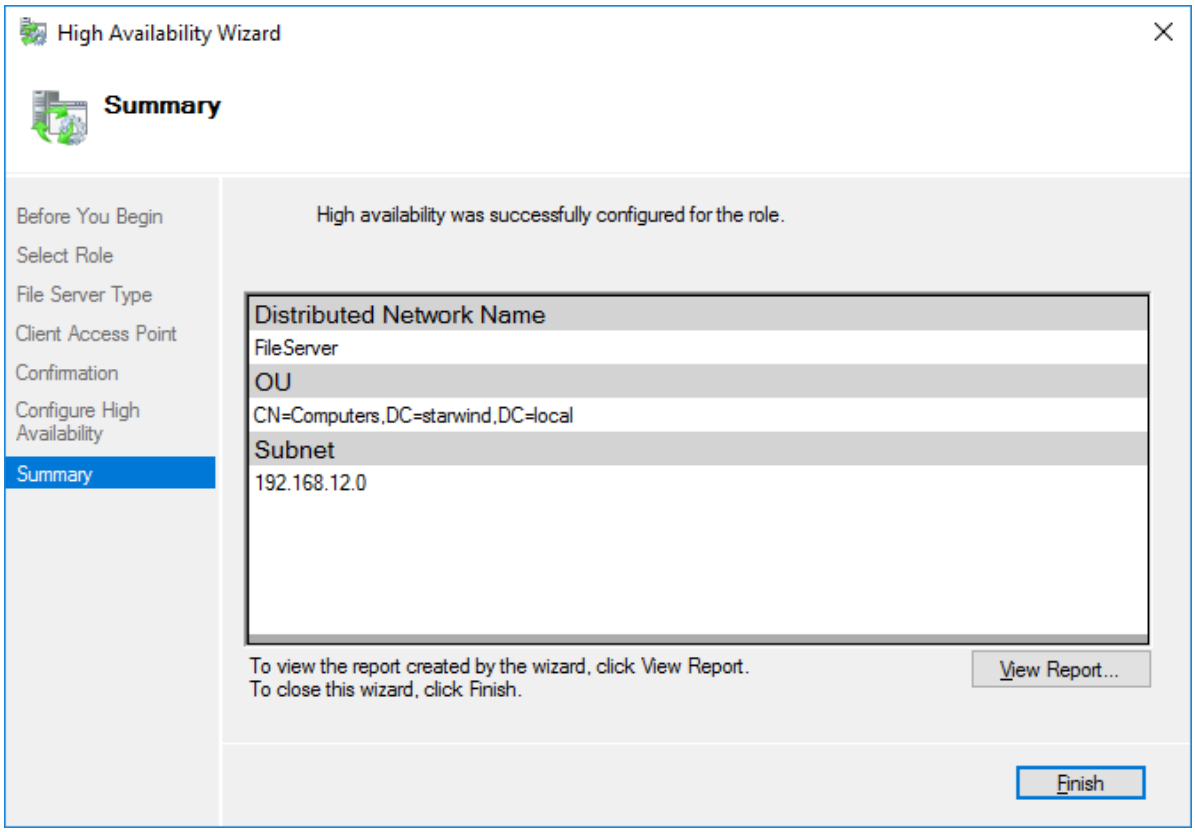

9. The newly created role should now look like the screenshot below.

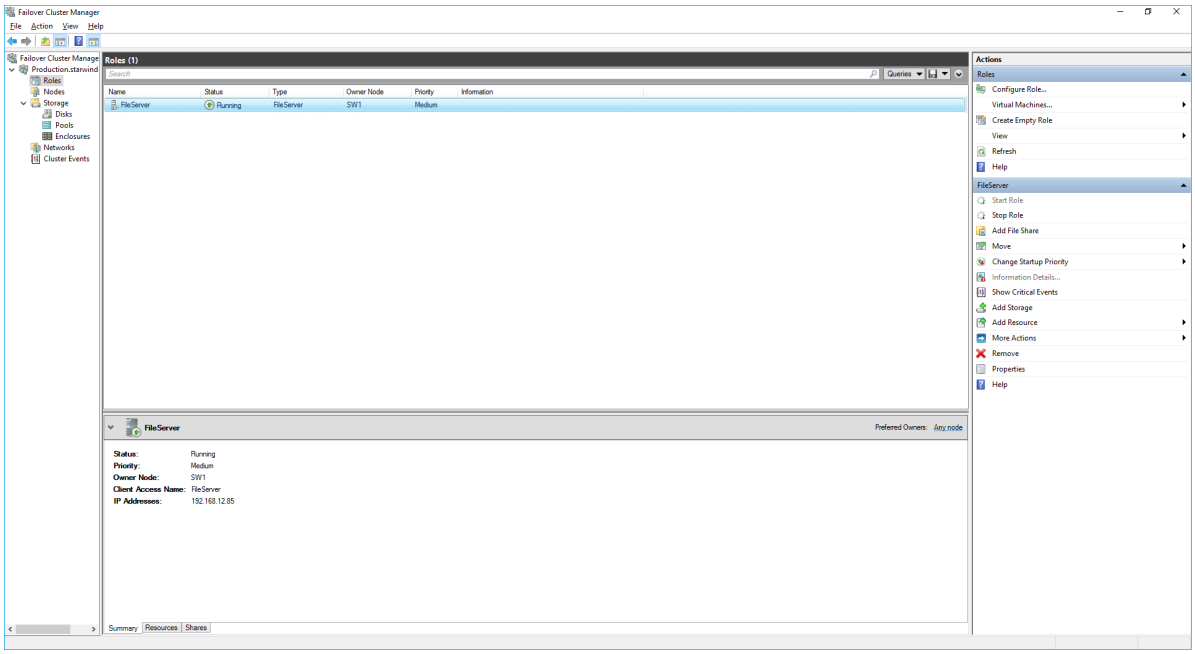

NOTE: If the role status is Failed and it is unable to Start, please, follow the next steps:

- open Active Directory Users and Computers

- enable the Advanced view if it is not enabled
- edit the properties of the OU containing the cluster computer object (in this case – Production)
- open the Security tab and click Advanced
- in the appeared window, press Add (the Permission Entry dialog box opens), click Select a principal
- in the appeared window, click Object Types, select Computers, and click OK
- enter the name of the cluster computer object (in this case – Production)

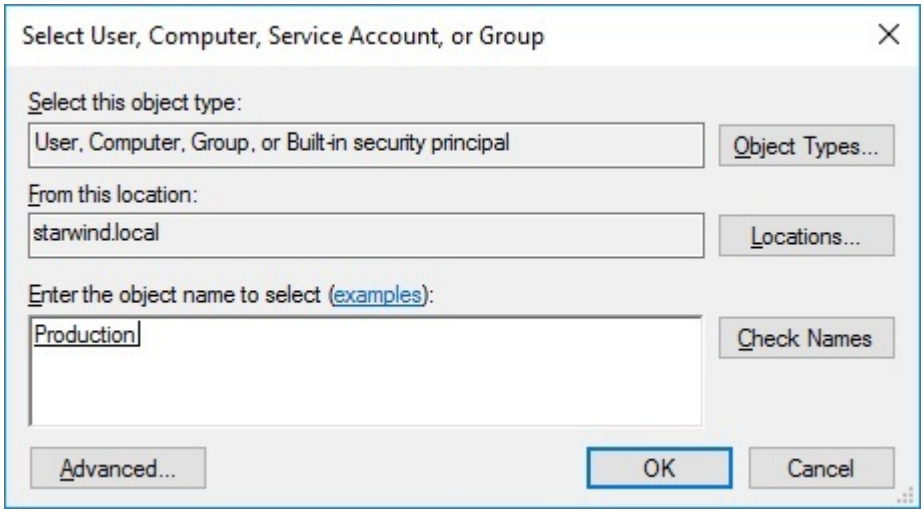

- go back to Permission Entry dialog, scroll down, and select Create Computer Objects

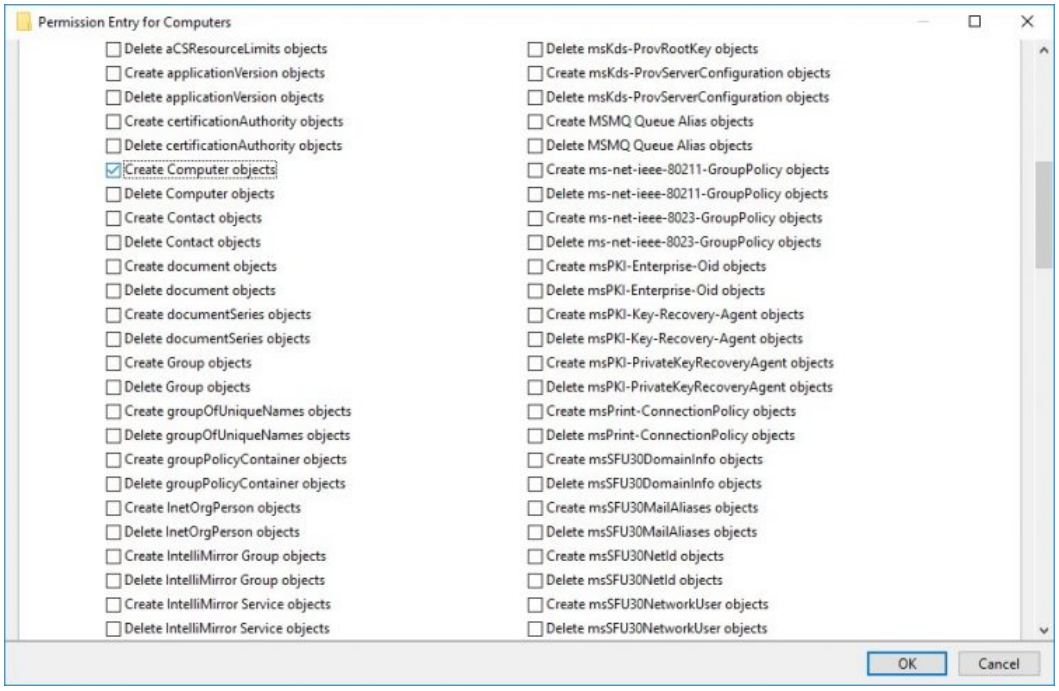

- click OK on all opened windows to confirm the changes

- open Failover Cluster Manager, right-click File Share role and click Start Role

Configuring Smb File Share

To Add SMB File Share

1. Open Failover Cluster Manager.

2. Expand the cluster and then click Roles.

3. Right-click the File Server role and then press Add File Share.

4. On the Select the profile for this share page, click SMB Share – Quick and then click Next.

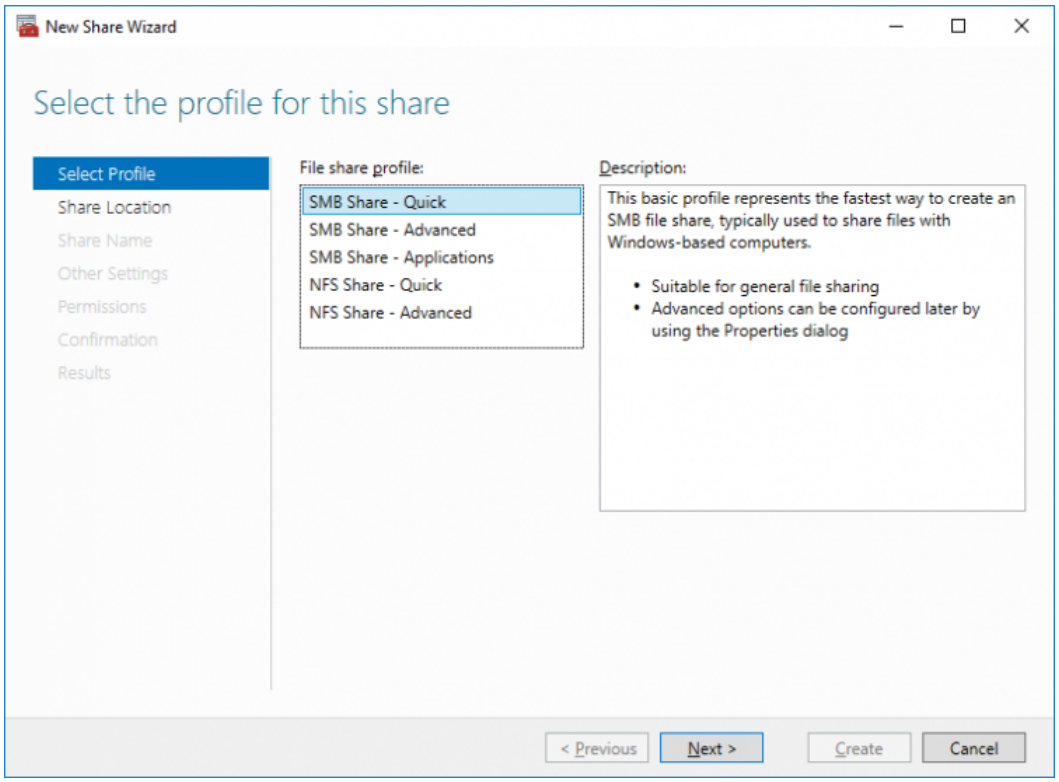

5. Select available storage to host the share. Click Next to continue.

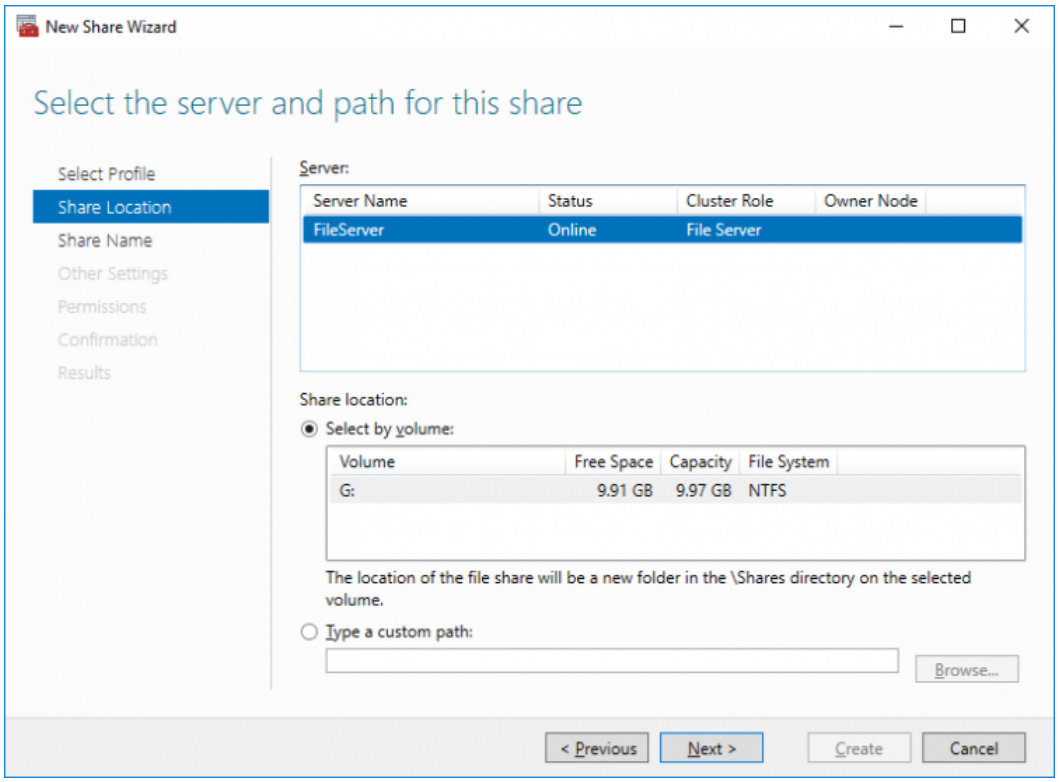

6. Type in the file share name and click Next.

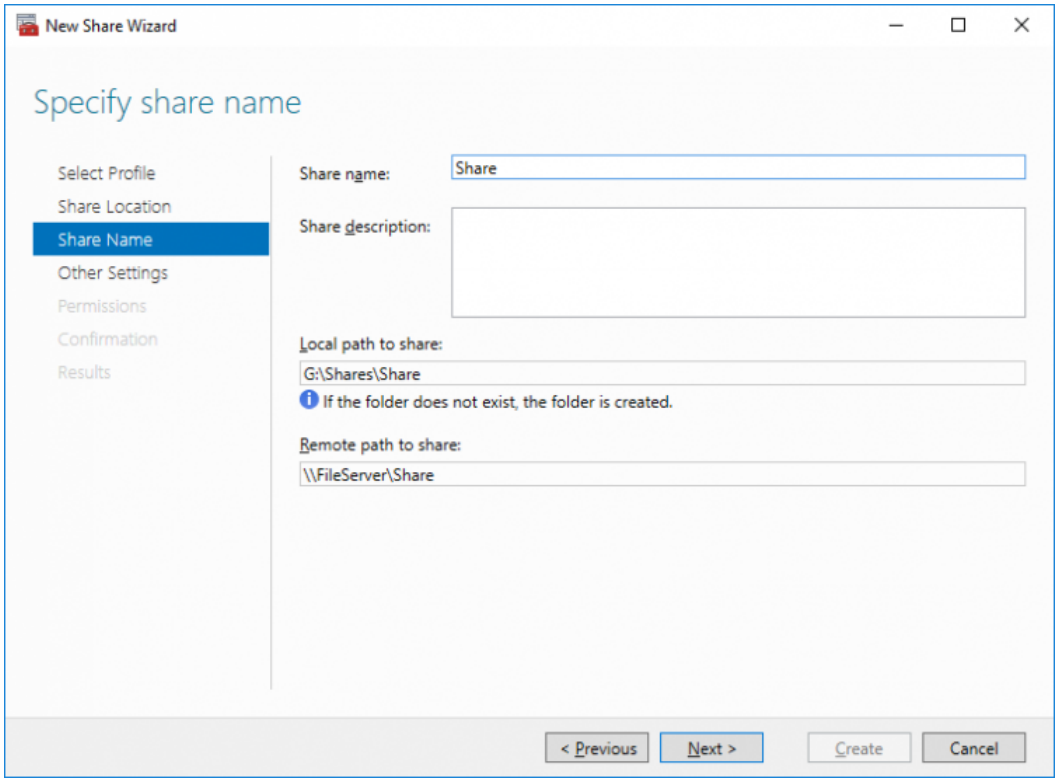

7. Make sure that the Enable Continuous Availability box is checked. Click Next to

continue.

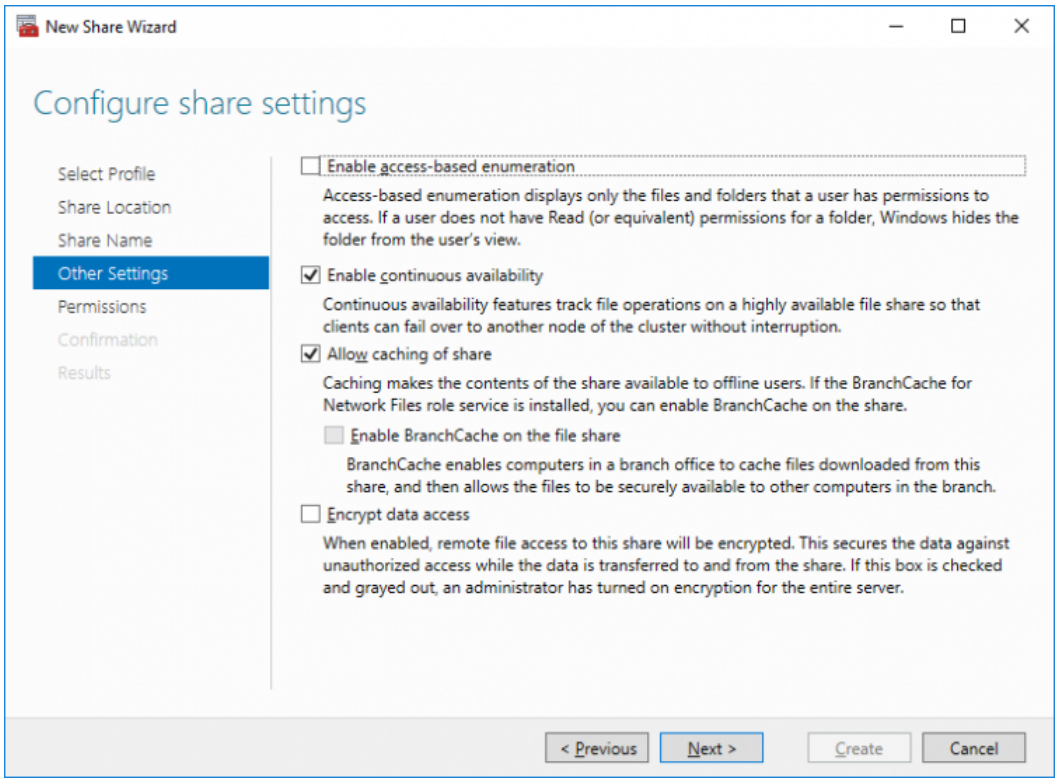

8.Specify the access permissions for the file share.

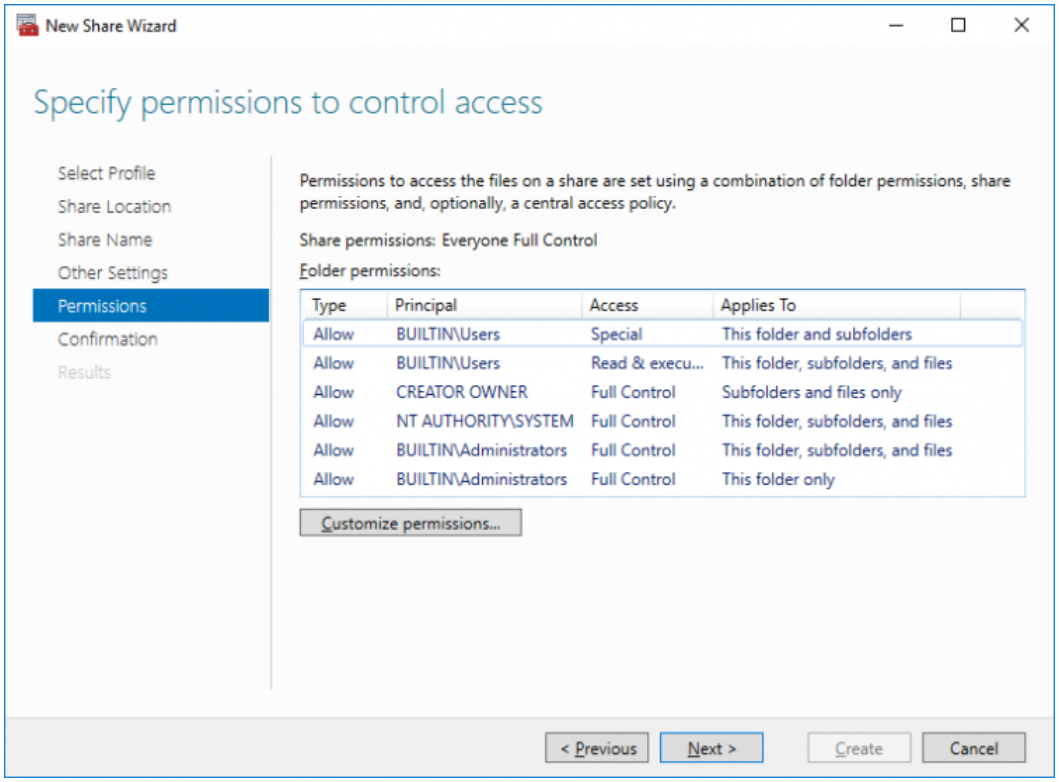

9. Check whether specified settings are correct. Click Previous to make any changes or Next/Create to continue.

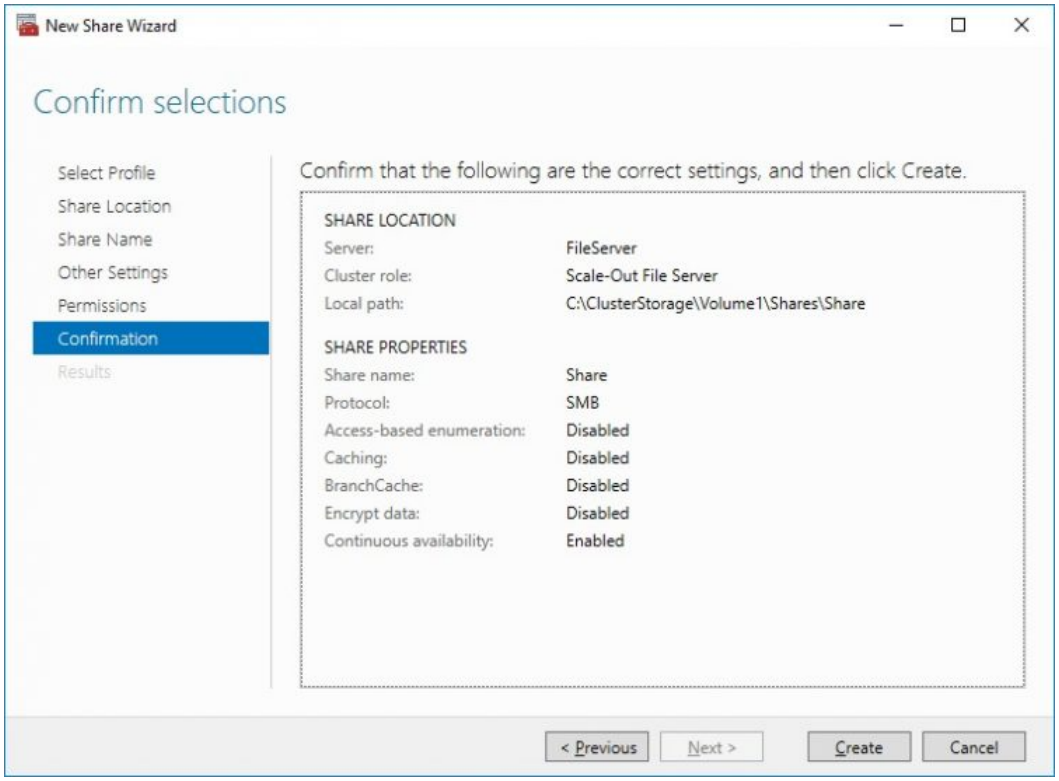

10. Check the summary and click Close.

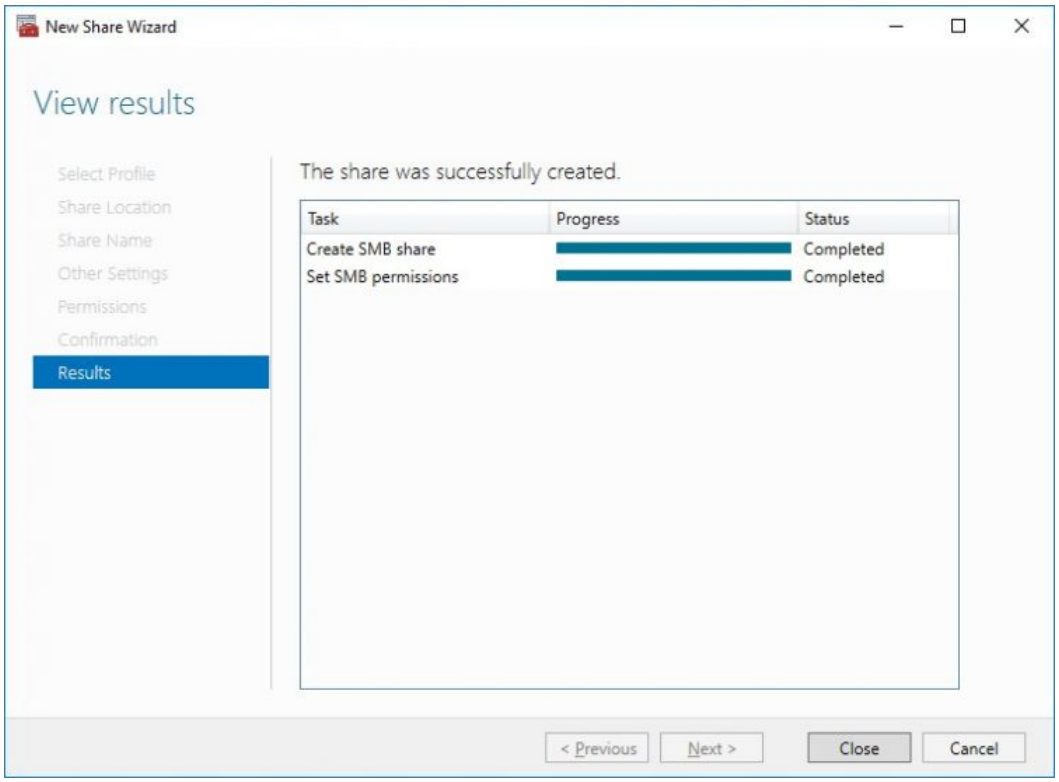

To manage created SMB File Shares

11. Open Failover Cluster Manager.

12. Expand the cluster and click Roles.

13. Choose the File Share role, select the Shares tab, right-click the created file share, and select Properties.

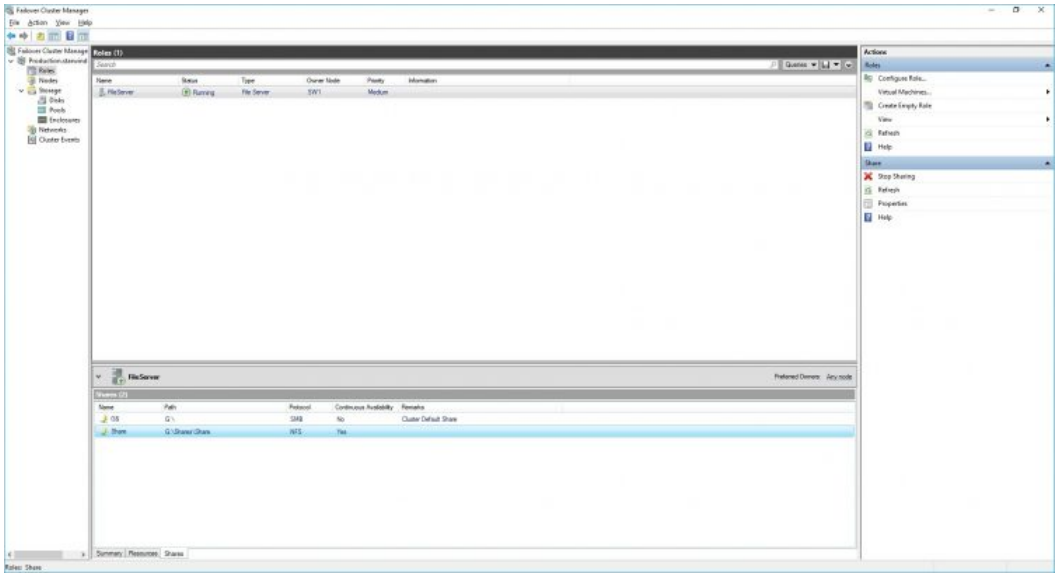

Configuring Nfs File Share

To Add NFS File Share

1. Open Failover Cluster Manager.

2. Expand the cluster and then click Roles.

3. Right-click the File Server role and then press Add File Share.

4. On the Select the profile for this share page, click NFS Share – Quick and then click Next.

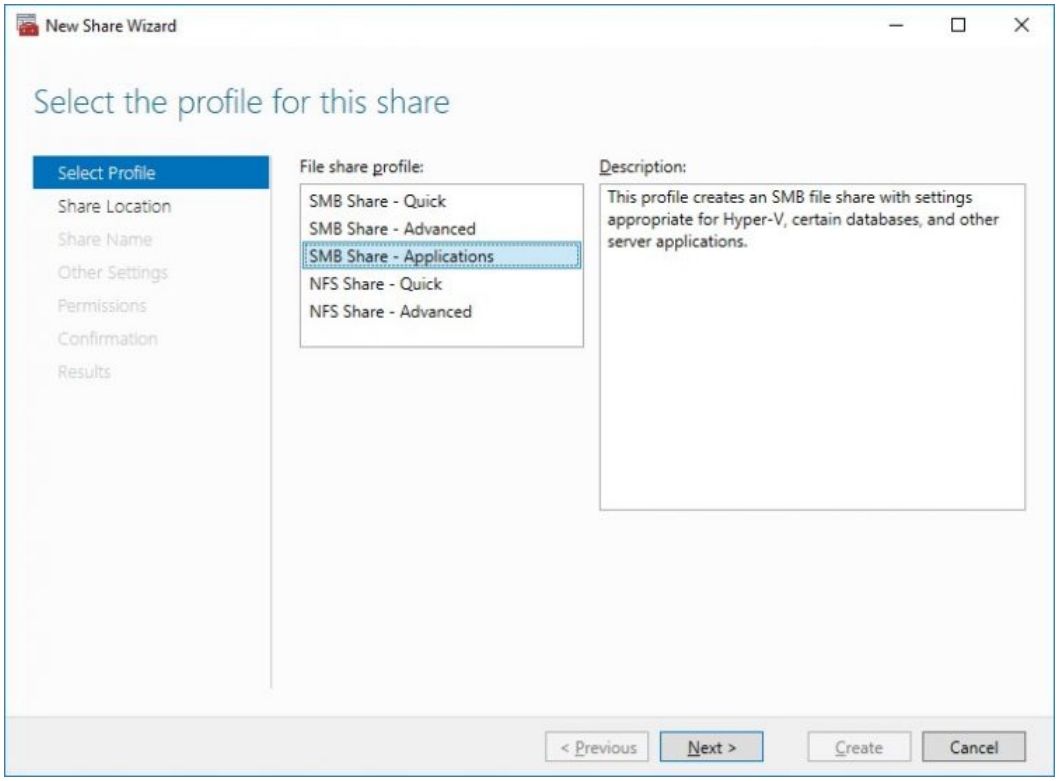

5. Select available storage to host the share. Click Next to continue.

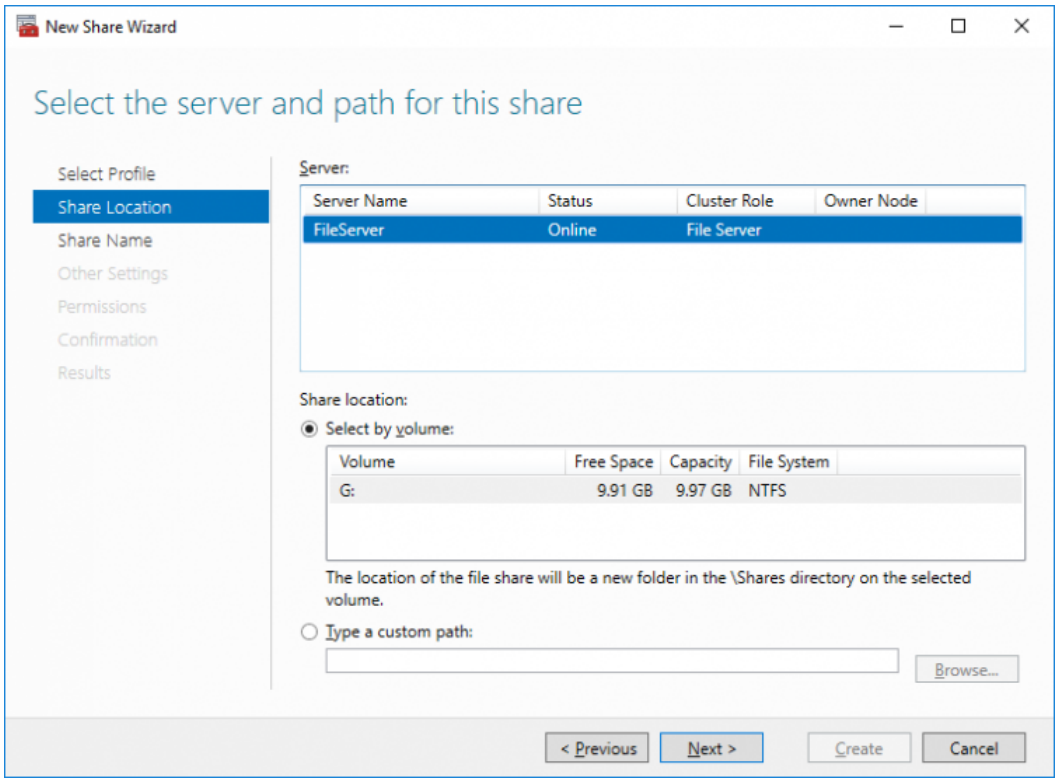

6. Type in the file share name and click Next.

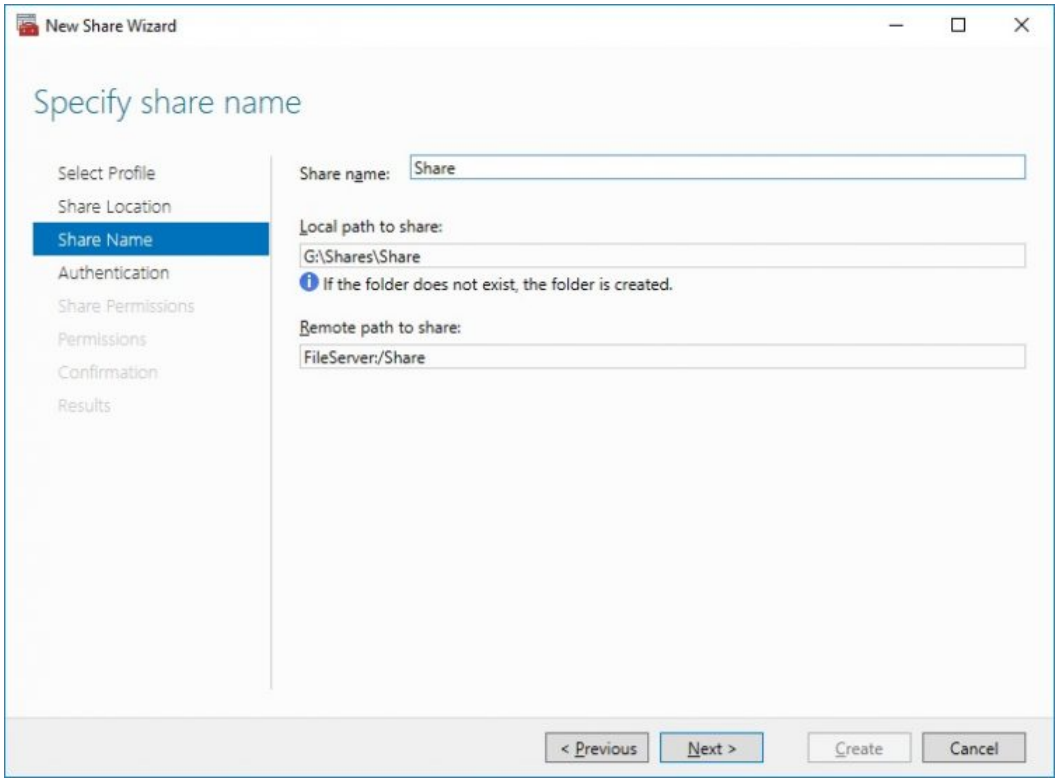

7. Specify the Authentication. Click Next and confirm the message in pop-up window to

continue.

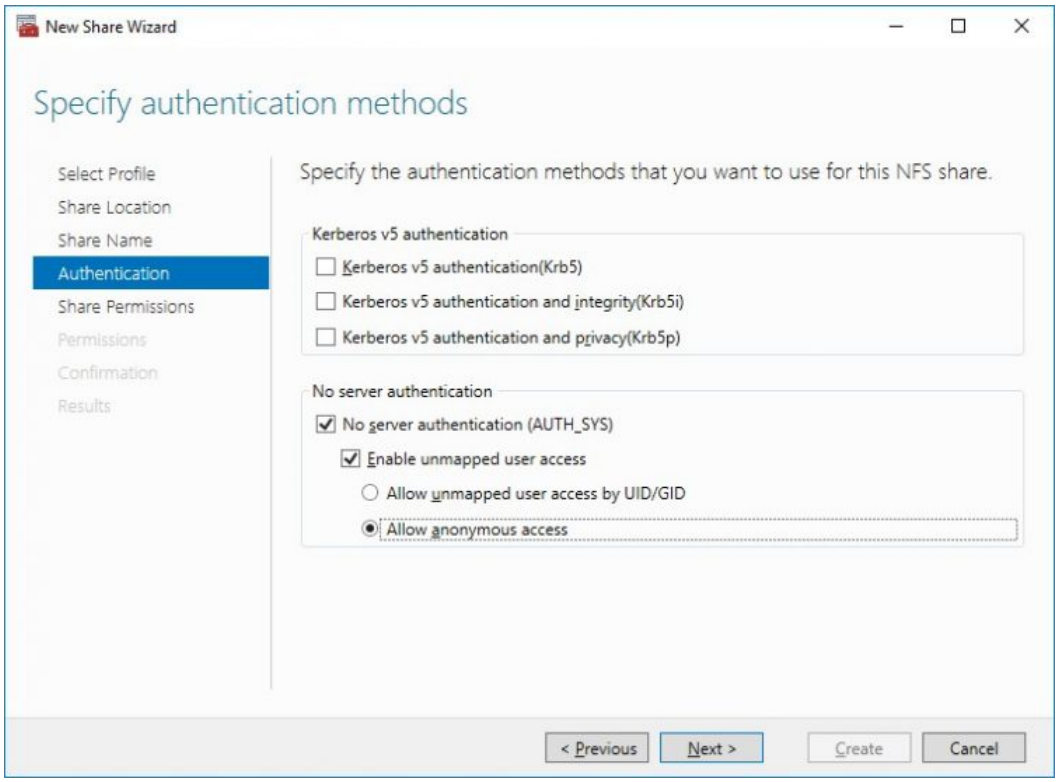

8. Click Add and specify Share Permissions.

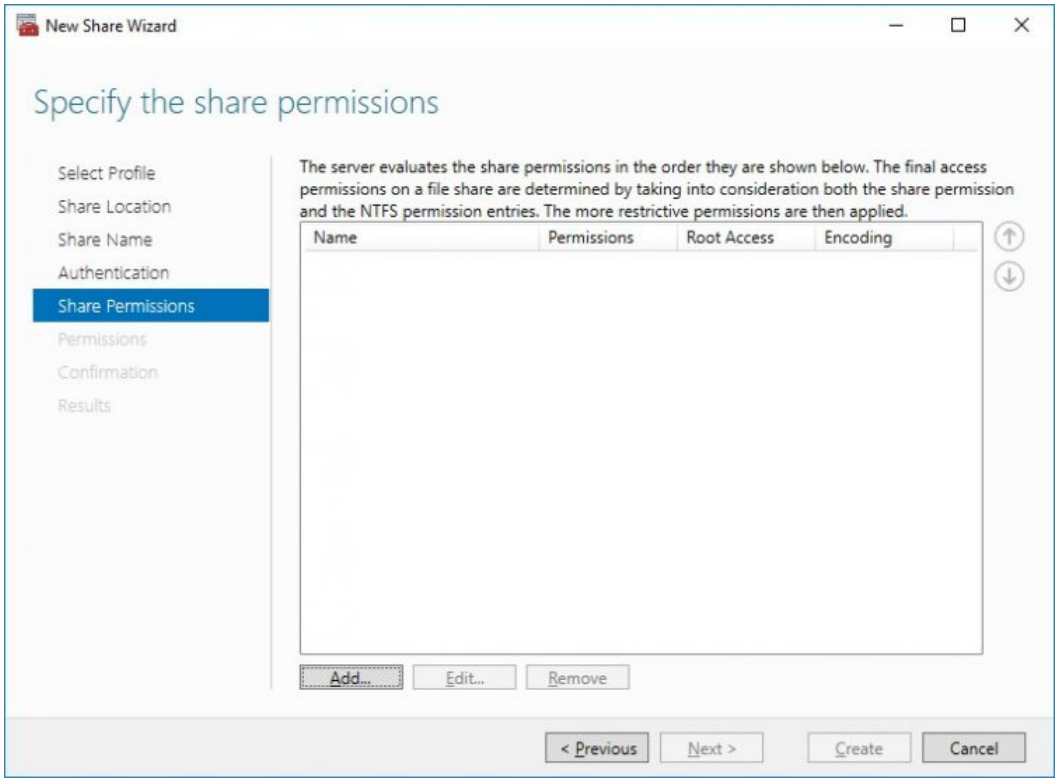

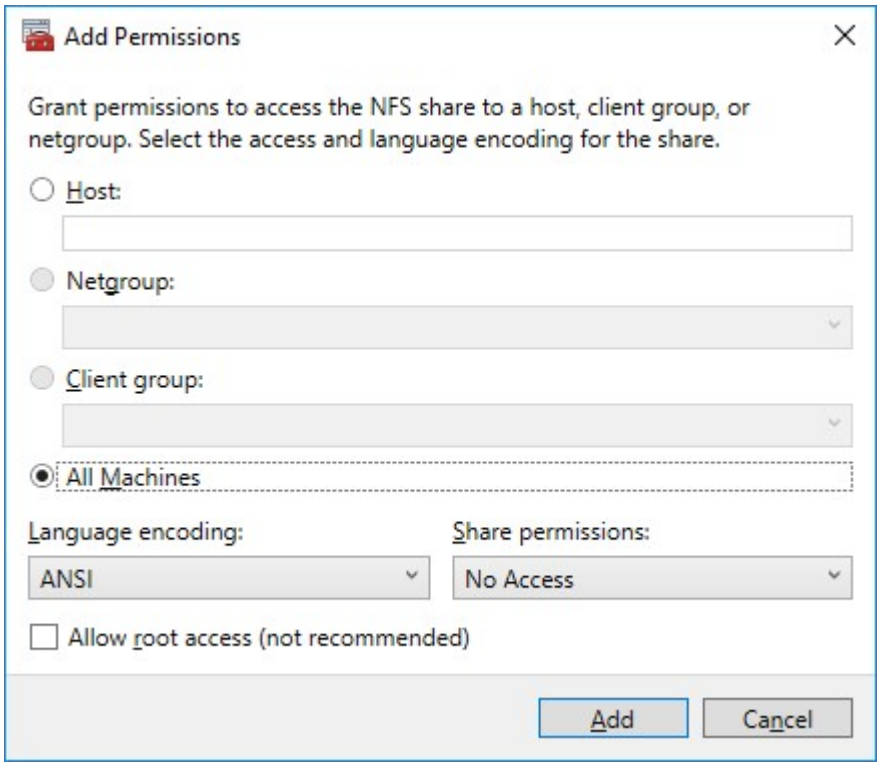

9. Specify the access permissions for the file share.

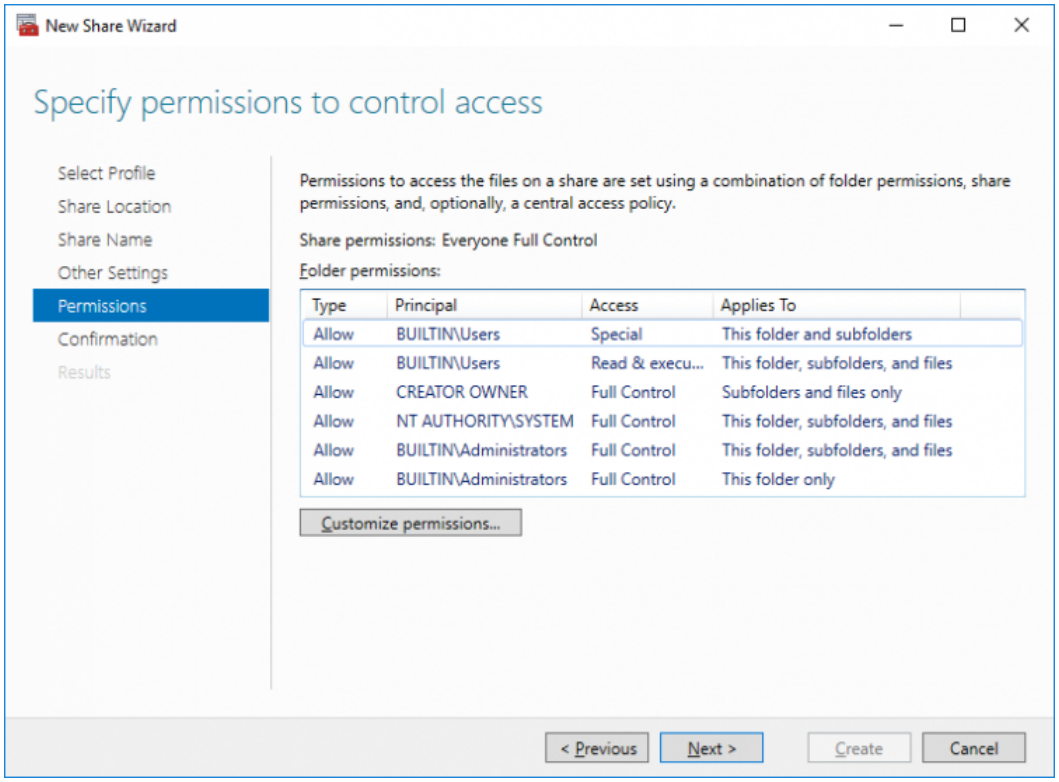

10. Check whether specified settings are correct. Click Previous to make any changes or click Create to continue.

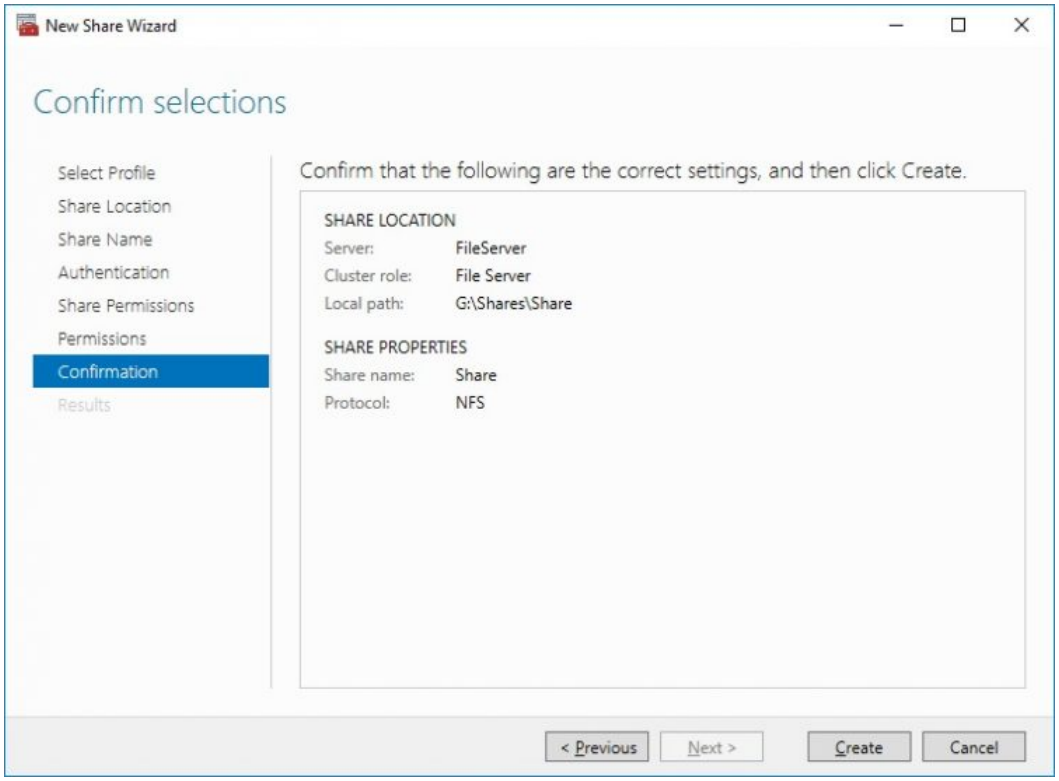

11. Check a summary and click Close to close the Wizard.

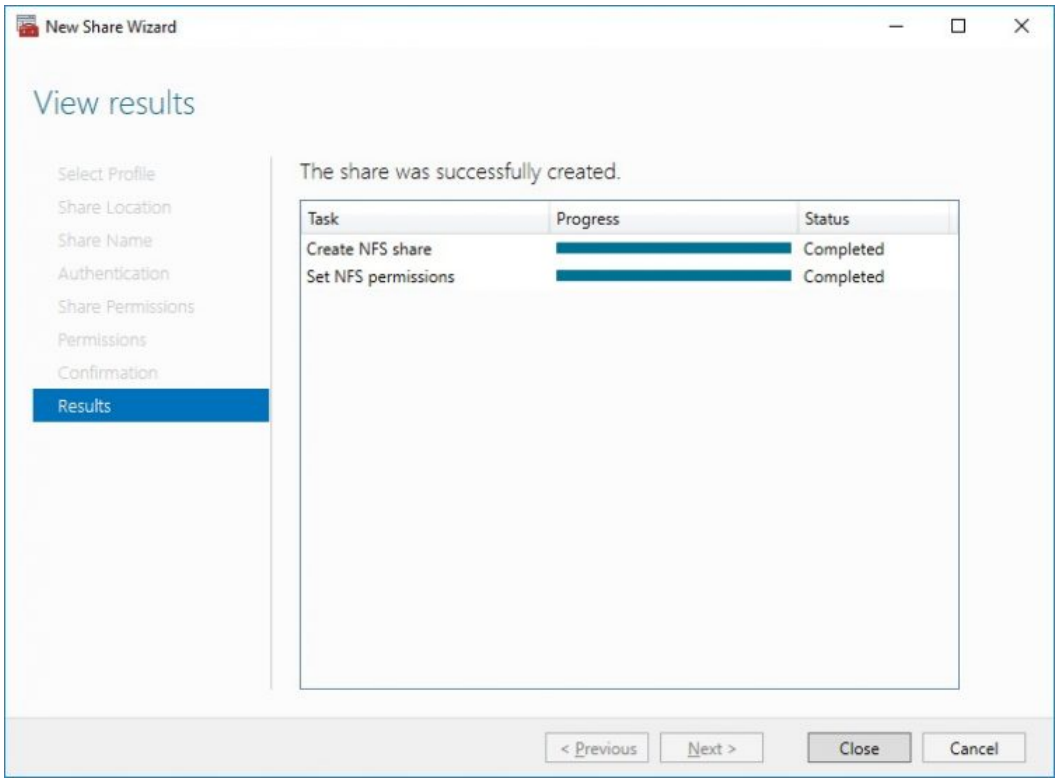

To manage created NFS File Shares:

- open Failover Cluster Manager
- expand the cluster and click Roles
- choose the File Share role, select the Shares tab, right-click the created file share, and select Properties

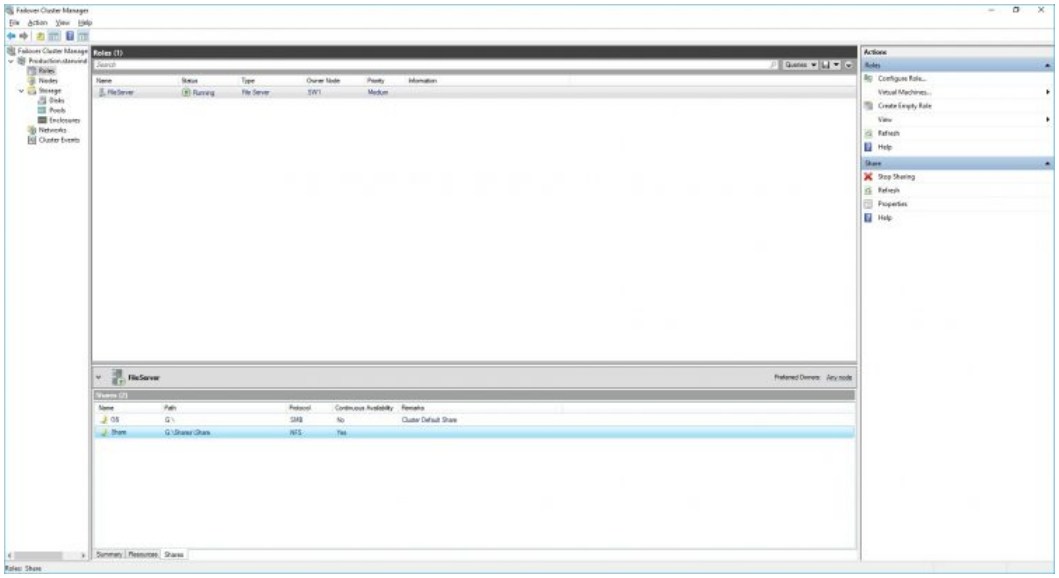

Conclusion

Following this guide, the Failover Cluster was deployed and configured with StarWind Virtual SAN (VSAN) running in Windows application on each host. As a result, a virtual shared storage “pool” accessible by all cluster nodes was created for storing highly available virtual machines.

Contacts

US Headquarters	EMEA and APAC
+1 617 829 44 95	+44 2037 691 857 (United Kingdom)
+1 617 507 58 45	+49 800 100 68 26 (Germany)
+1 866 790 26 46	+34 629 03 07 17 (Spain and Portugal)
	+33 788 60 30 06 (France)

Customer Support Portal: https://www.starwind.com/support
Support Forum: https://www.starwind.com/forums
Sales: sales@starwind.com
General Information: info@starwind.com

StarWind