

# StarWind Virtual SAN: Configuration Guide for Microsoft Azure Stack HCI [Hyper-V], VSAN Deployed as a Controller VM using PowerShell CLI

2026

TECHNICAL PAPERS



## Trademarks

“StarWind”, “StarWind Software” and the StarWind and the StarWind Software logos are registered trademarks of StarWind Software. “StarWind LSFS” is a trademark of StarWind Software which may be registered in some jurisdictions. All other trademarks are owned by their respective owners.

## 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.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of StarWind Software.

## Annotation

### Relevant products

This guide applies to StarWind Virtual SAN and StarWind Virtual SAN Free (Version V8 (Build 15260, OVF Version 20230901) and earlier).

For newer versions of StarWind Virtual SAN (Version V8 (Build 15260, CVM Version 20231016) and later), please refer to this configuration guide: [StarWind Virtual SAN: Configuration Guide for Microsoft Azure Stack HCI \[Hyper-V\], VSAN Deployed as a Controller Virtual Machine using PowerShell CLI](#)

### Purpose

This guide is to assist users in the deployment and configuration of StarWind VSAN CVM on Microsoft Azure Stack HCI platforms, emphasizing best practices and system requirements.

### Audience

This document is designed for IT professionals, system administrators, and technical users who are looking to deploy and configure StarWind VSAN CVM on Microsoft Azure Stack HCI platforms using PowerShell CLI.

### Expected Results

Following this guide, users will get a fully configured Failover Cluster on Microsoft Azure Stack HCI with shared storage provided by StarWind VSAN.

## Introduction To Starwind Virtual San Cvm

StarWind Virtual SAN Controller Virtual Machine (CVM) comes as a prepackaged Linux Virtual Machine (VM) to be deployed on any industry-standard hypervisor. It creates a VM-centric and high-performing storage pool for a VM cluster.

This guide describes the deployment and configuration process of the StarWind Virtual SAN CVM.

## Starwind Vsan 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:

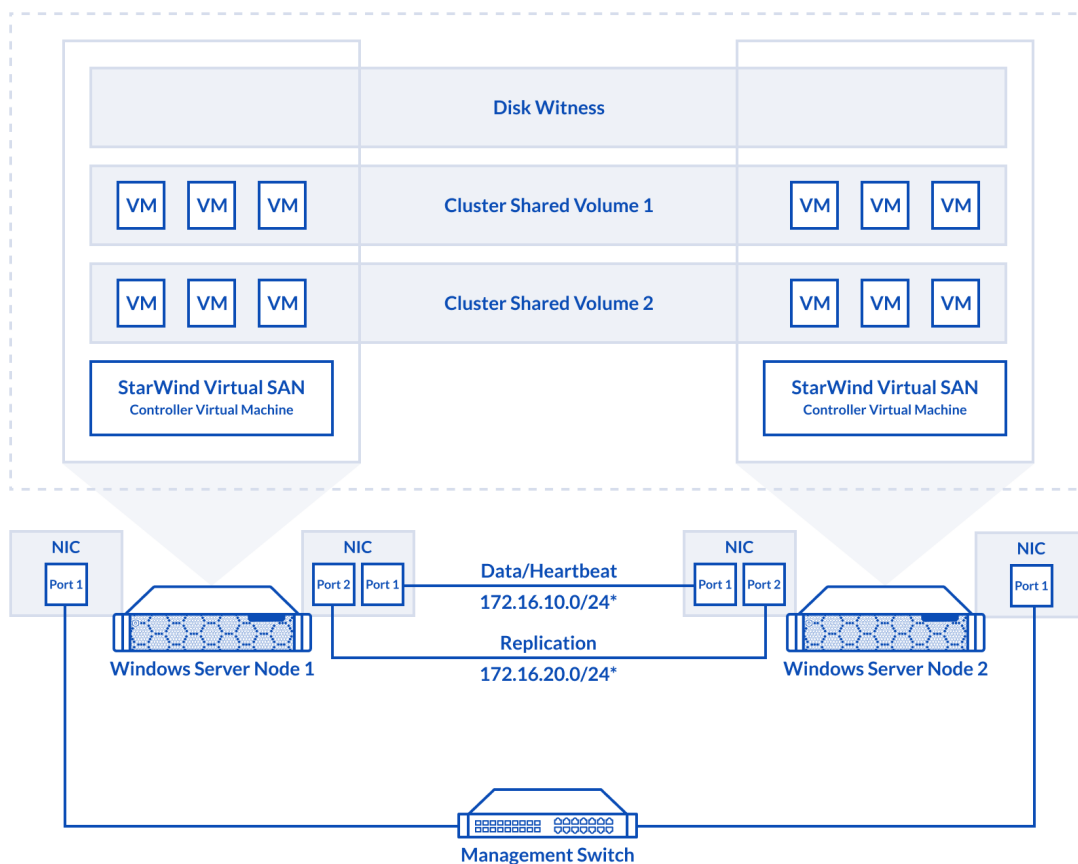
<https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/>

Please read StarWind Virtual SAN Best Practices document for additional information:

<https://www.starwindsoftware.com/resource-library/starwind-virtual-san-best-practices>

## Pre-Configuring The Windows Server Hosts

The diagram below illustrates the network and storage configuration of the solution:



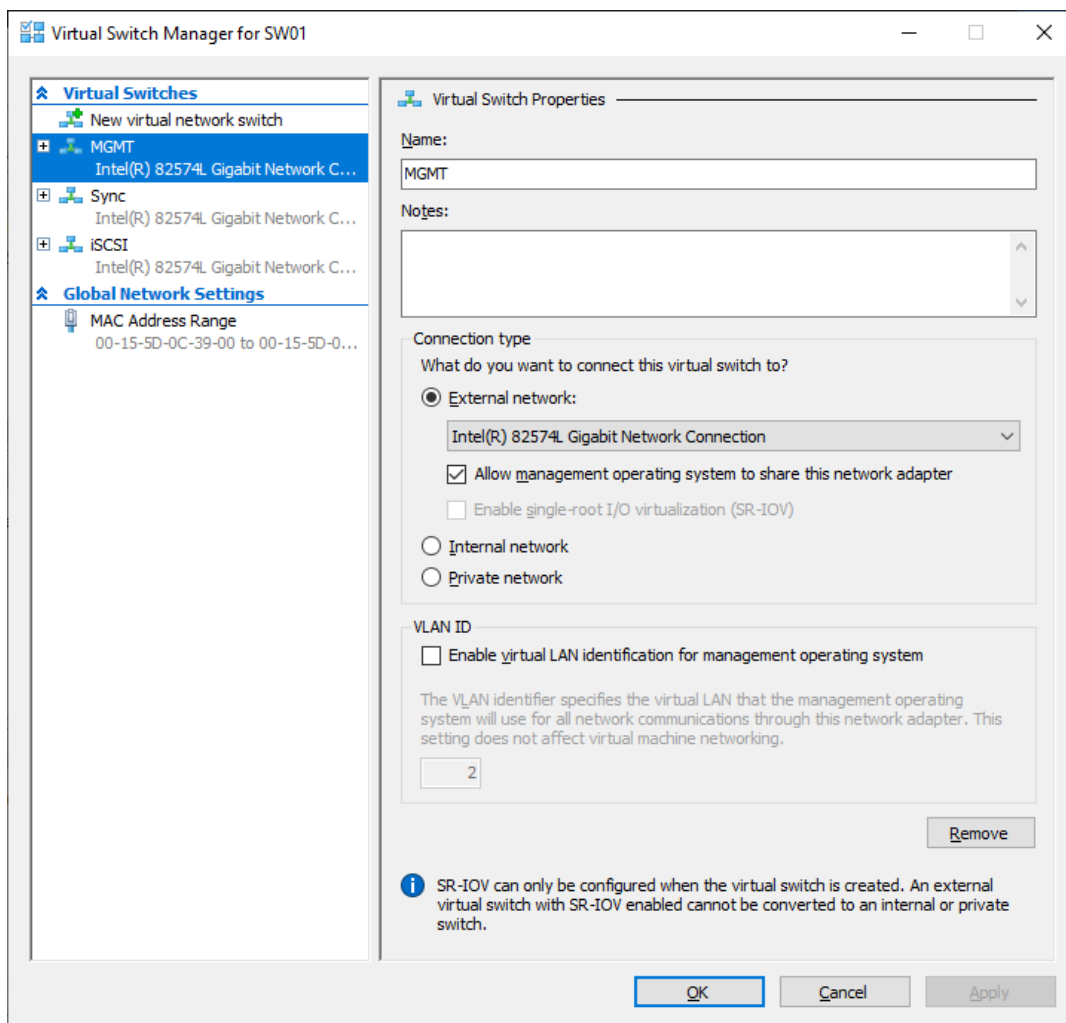
1. Make sure that a domain controller is configured and the servers are added to the domain.

NOTE: Please follow the recommendation in [KB article](#) on how to place a DC in case of StarWind Virtual SAN usage.

2. Deploy Windows Server on each server and install Failover Clustering and Multipath I/O features, as well as the Hyper-V role on both servers. This can be done through Server Manager (Add Roles and Features menu item).

3. Define at least 2x network interfaces on each node that will be used for the Synchronization and iSCSI/StarWind heartbeat traffic. Do not use iSCSI/Heartbeat and Synchronization channels over the same physical link. Synchronization and iSCSI/Heartbeat links can be connected either via redundant switches or directly between the nodes (see diagram above).

4. Separate external Virtual Switches should be created for iSCSI and Synchronization traffic based on the selected before iSCSI and Synchronization interfaces. Using Hyper-V Manager open Virtual Switch Manager and create two external Virtual Switches: one for the iSCSI/StarWind Heartbeat channel (iSCSI) and another one for the Synchronization channel (Sync).



5. Configure and set the IP address on each virtual switch interface. In this document, the 172.16.10.x subnet is used for iSCSI/StarWind heartbeat traffic, while 172.16.20.x subnet is used for the Synchronization traffic.

NOTE: In case NIC supports SR-IOV, enable it for the best performance. An additional internal switch is required for iSCSI Connection. Contact support for additional details.

6. Set MTU size to 9000 on iSCSI and Sync interfaces using the following Powershell script.

```
$iSCSIs = (Get-NetAdapter -Name "*iSCSI*").Name
$SynCs = (Get-NetAdapter -Name "*Sync*").Name
foreach ($iSCSI in $iSCSIs) {
Set-NetAdapterAdvancedProperty -Name "$iSCSI" -RegistryKeyword
"*JumboPacket" -Registryvalue 9014
Get-NetAdapterAdvancedProperty -Name "$iSCSI" -RegistryKeyword
"*JumboPacket"
}
}
```

```
foreach ($Sync in $Syncs) {  
Set-NetAdapterAdvancedProperty -Name "$Sync" -RegistryKeyword  
"*JumboPacket" -Registryvalue 9014  
Get-NetAdapterAdvancedProperty -Name "$Sync" -RegistryKeyword  
"*JumboPacket"  
}
```

It will apply MTU 9000 to all iSCSI and Sync interfaces if they have iSCSI or Sync as part of their name.

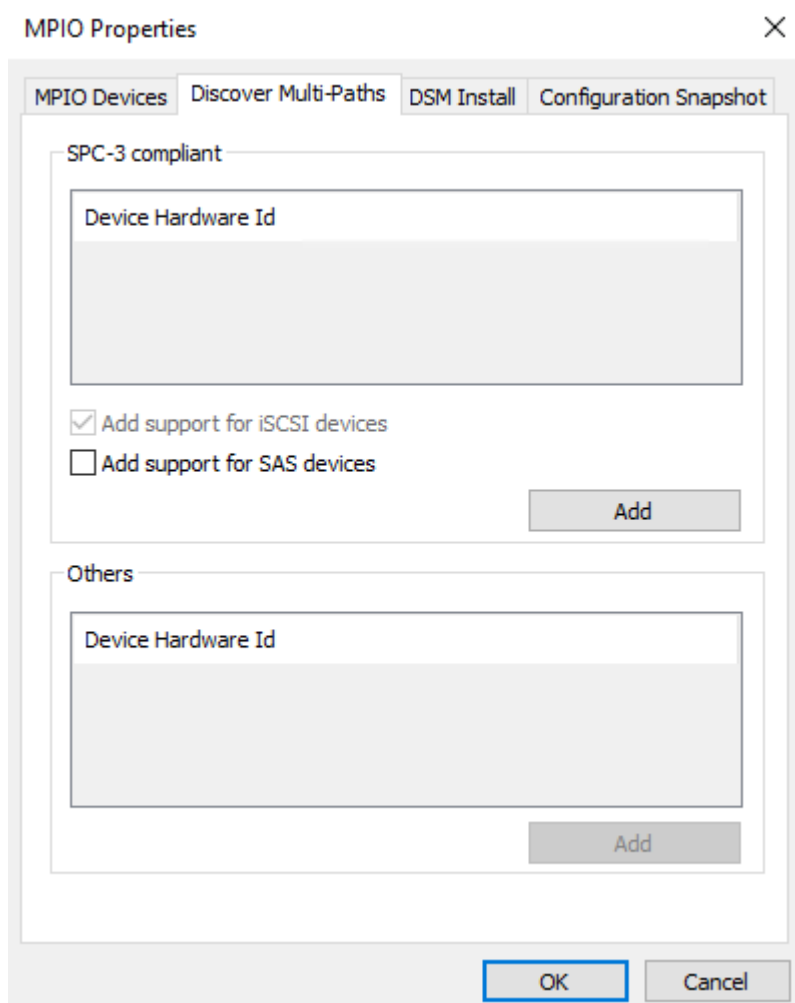
NOTE: MTU setting should be applied on the adapters only if there is no live production running through the NICs.

### Enabling Multipath Support

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

```
mpioctl
```

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



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

10. 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

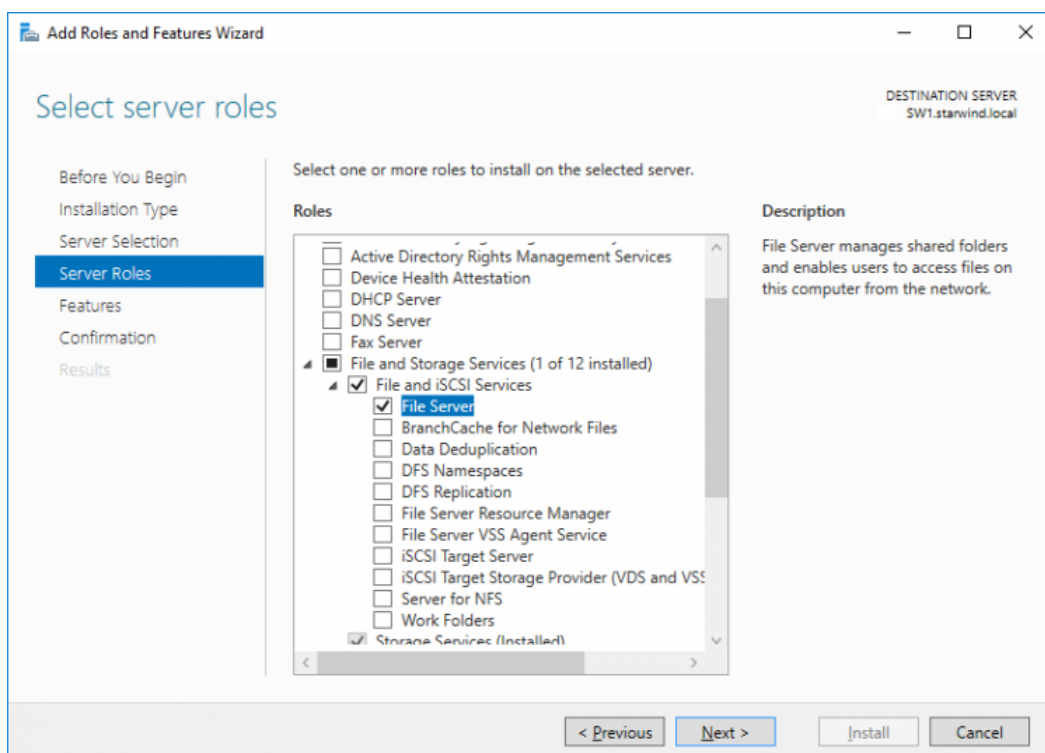
below:



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

## File Server For General Use With Smb Share

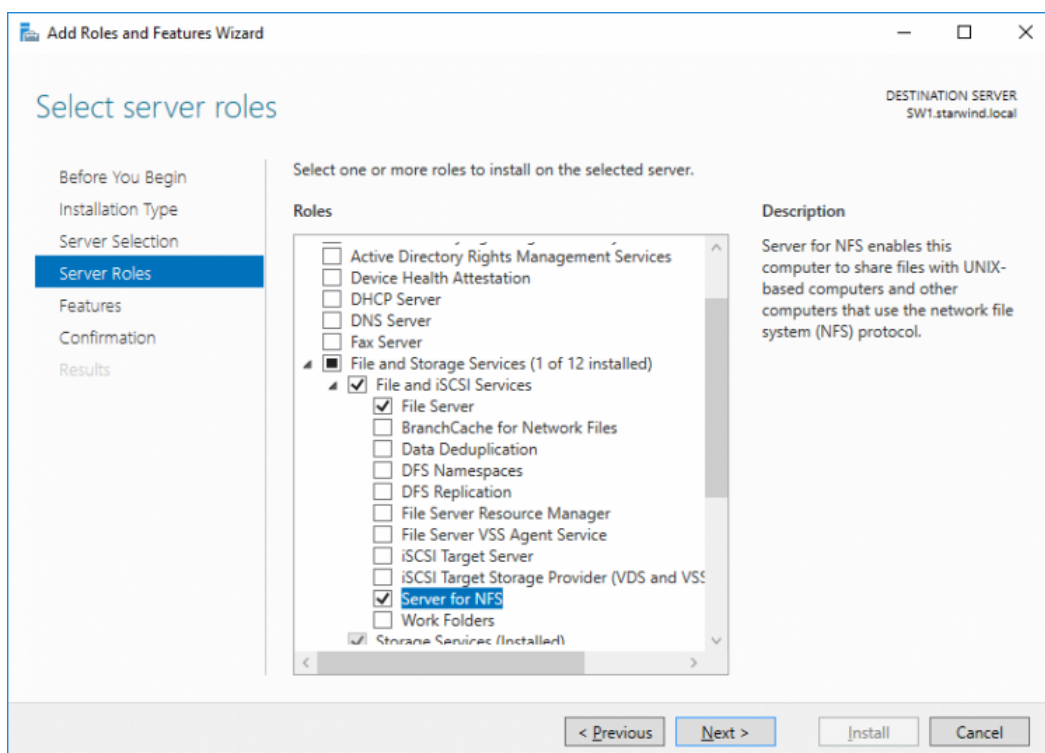
1. Open Server Manager: Start -> Server Manager.
2. Select: Manage -> Add Roles and Features.
3. Follow the installation wizard steps to install the roles selected in the screenshot below:



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.

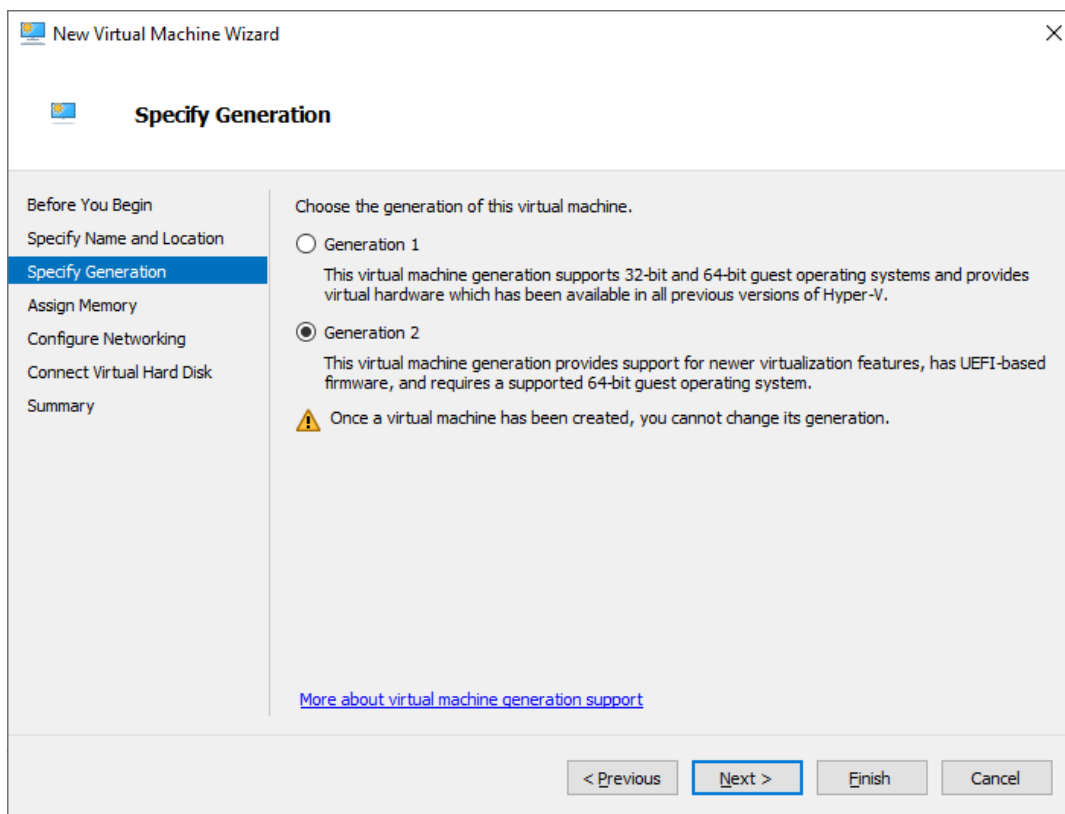
## Deploying Starwind Virtual San Cvm

1. Download StarWind VSAN CVM to the Hyper-V server:  
[VSAN by StarWind: Overview](#)
2. Extract the VM .vhdx file from the downloaded archive.
3. Open Hyper-V Manager and create a new Hyper-V VM.
4. Specify the name of the virtual machine with StarWind VSAN, and choose the location for the VM. For example: C:\SW1\

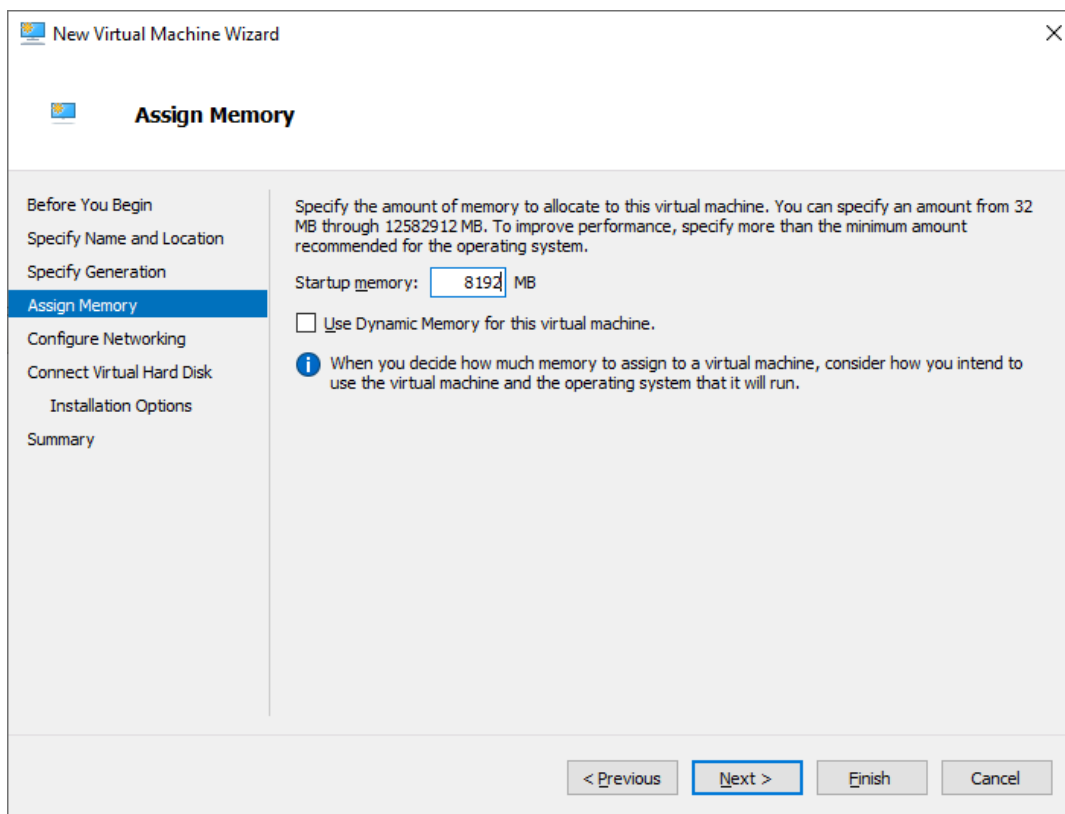


5. Copy .vhdx to the folder where the VM was created. In this case, its C:\ SW1\Virtual Hard Disks\


6. Choose Generation 2 VM.



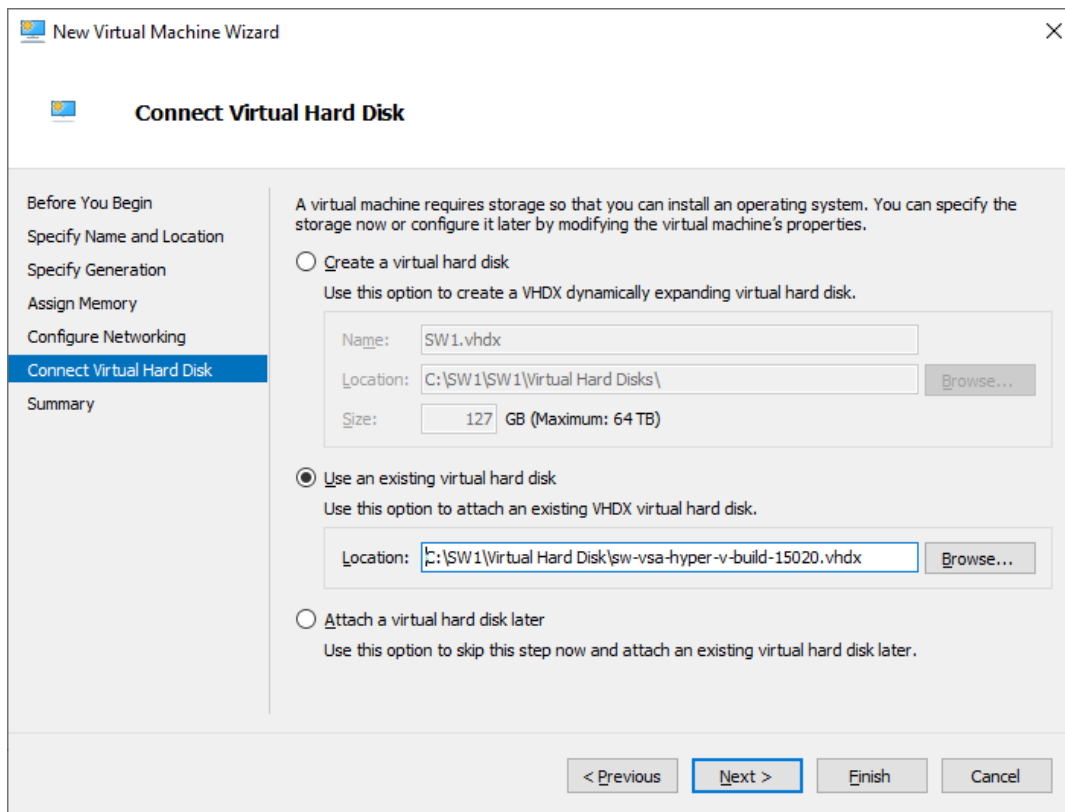
7. Assign memory to the VM. We recommend allocating at least 8GB of RAM for StarWind CVM. If StarWind L1 cache is used, an appropriate amount of RAM should be assigned.



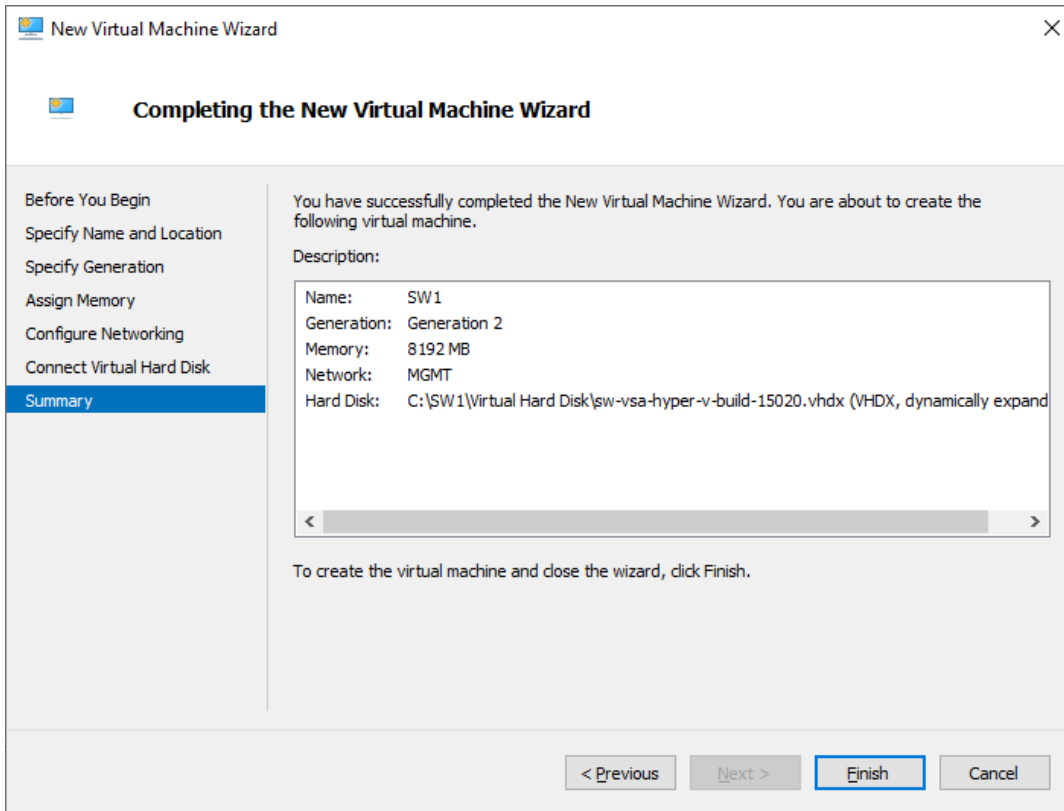
8. Choose the management network for the VM.



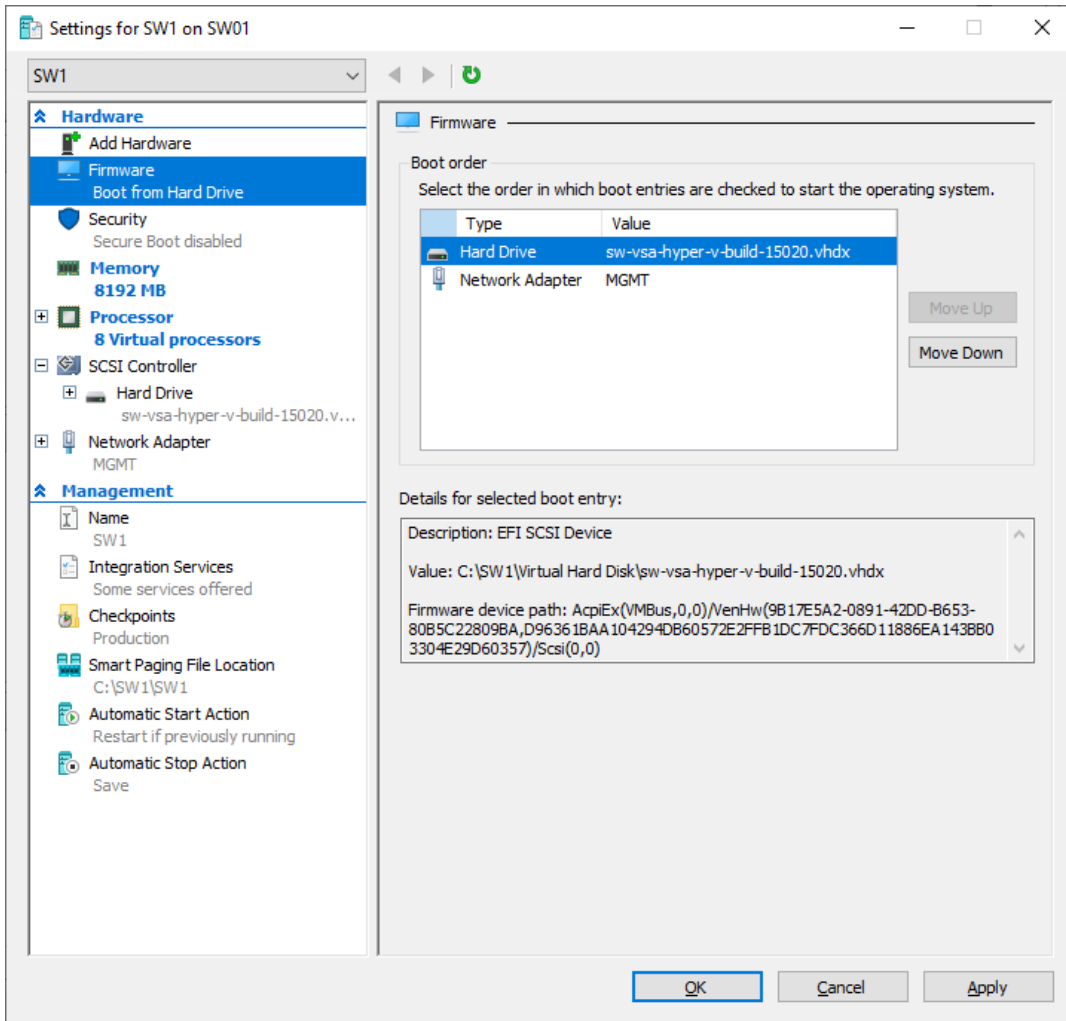
9. Connect .vhdx to the VM.



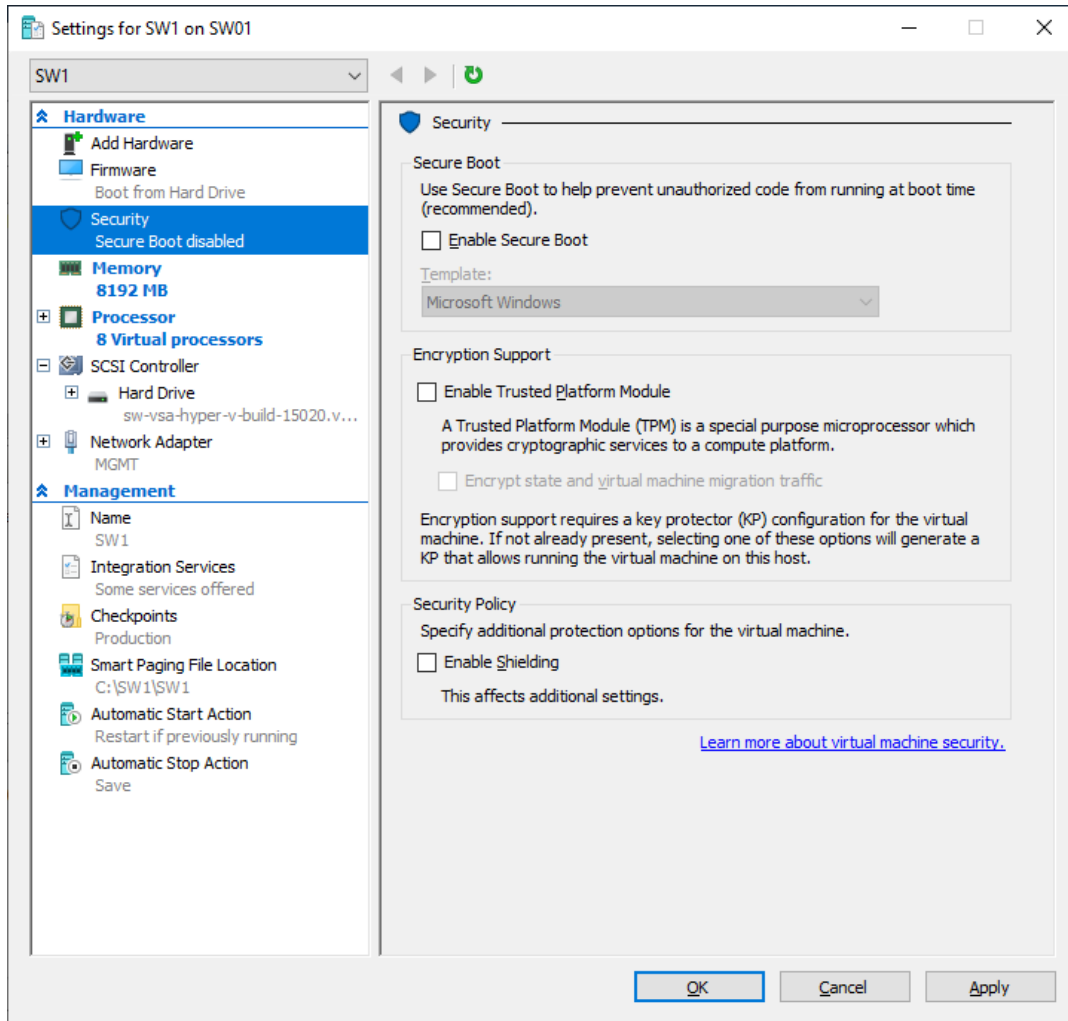
10. Review the Summary and click Finish to create the VM.



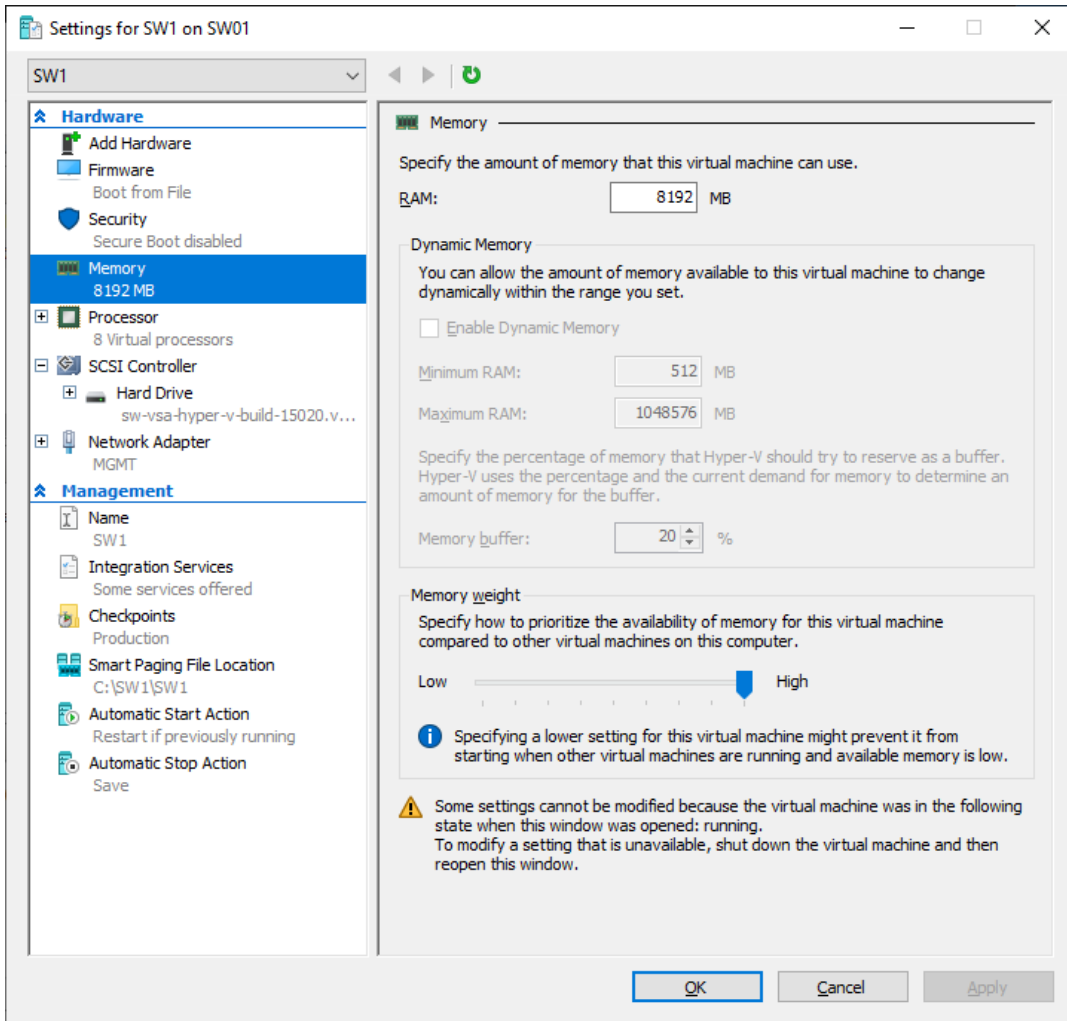
11. Right-click on the VM and choose Settings. Open Firmware and move the option Hard Drive to the first place in the list.



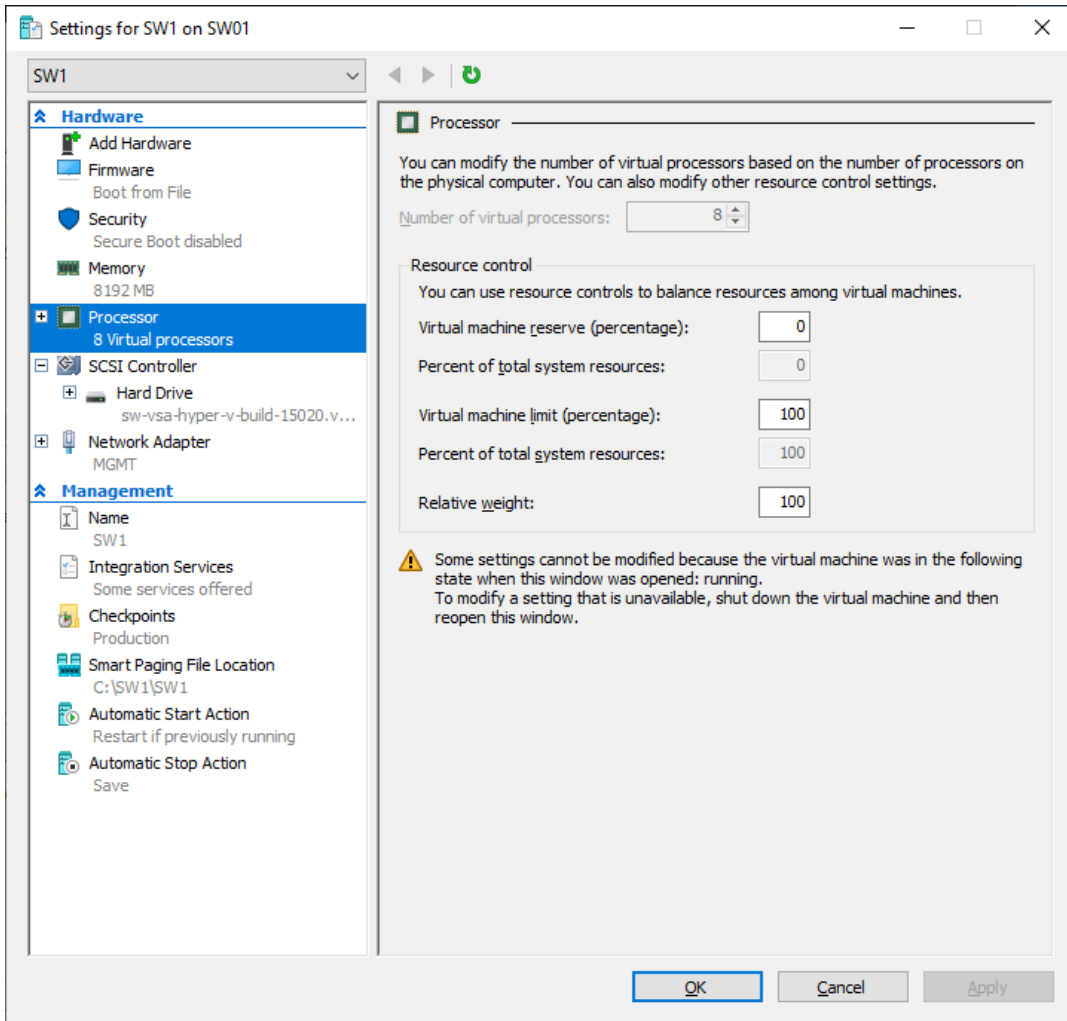
12. Go to the Security page and uncheck the Enable Secure Boot box.



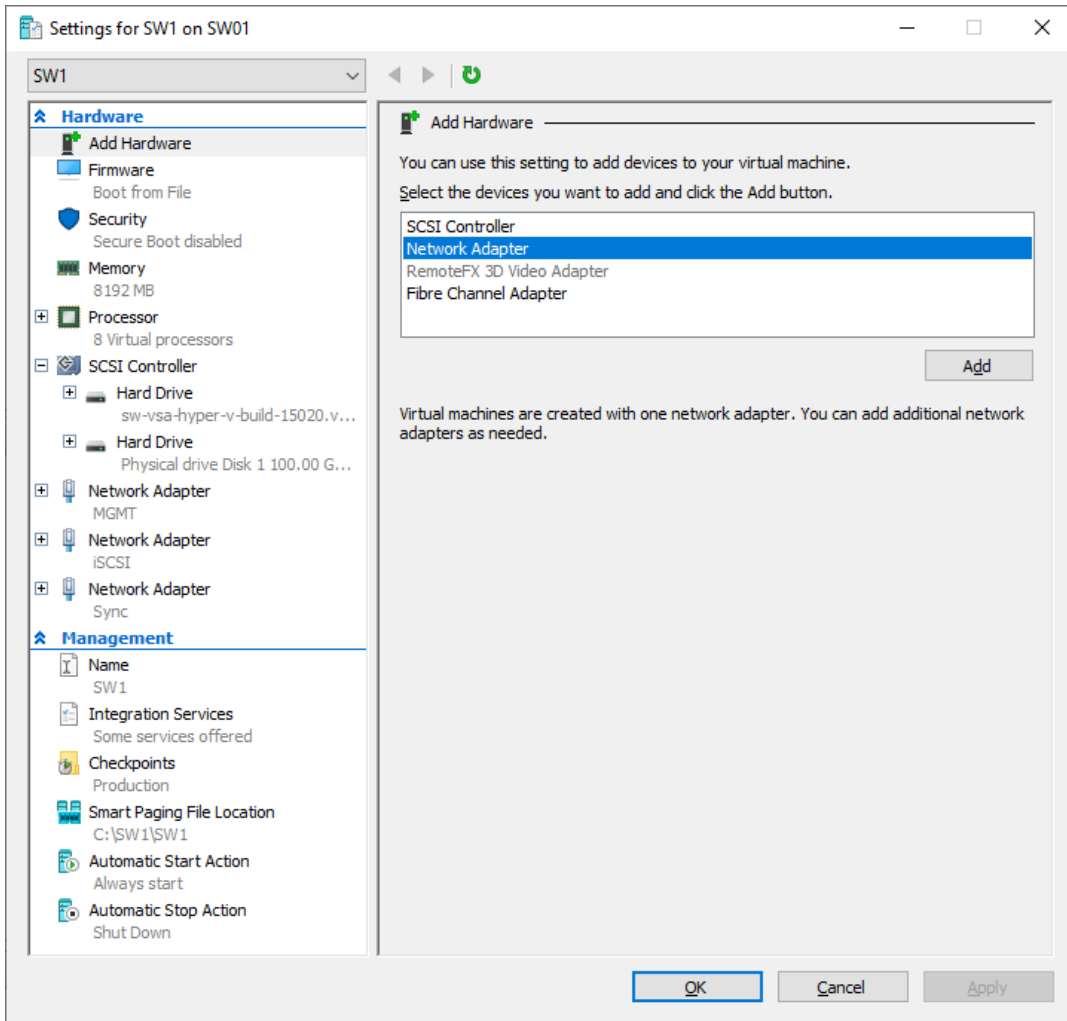
13. Go to Memory and move the slider for Memory weight to High.

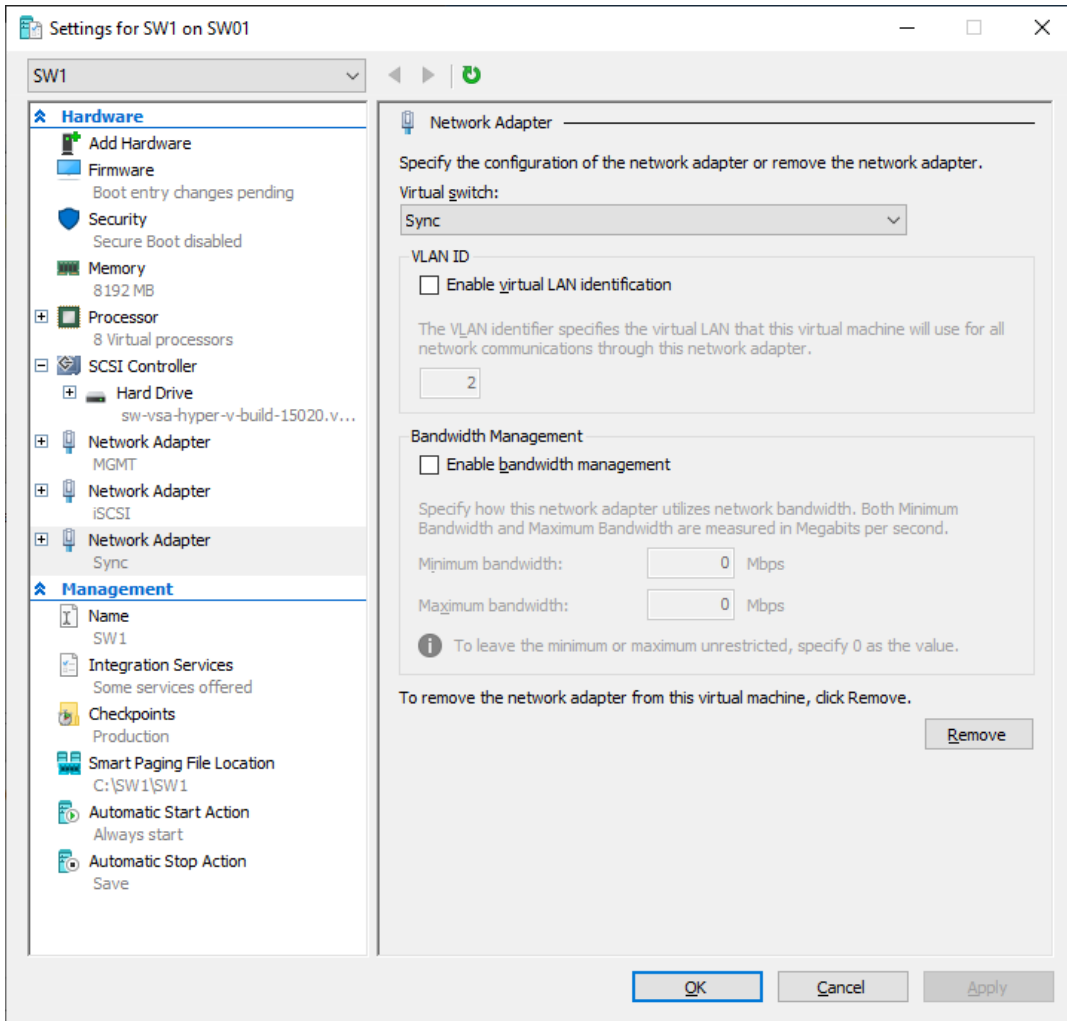


14. Go to Processor and assign 8 vCPUs to the VM.

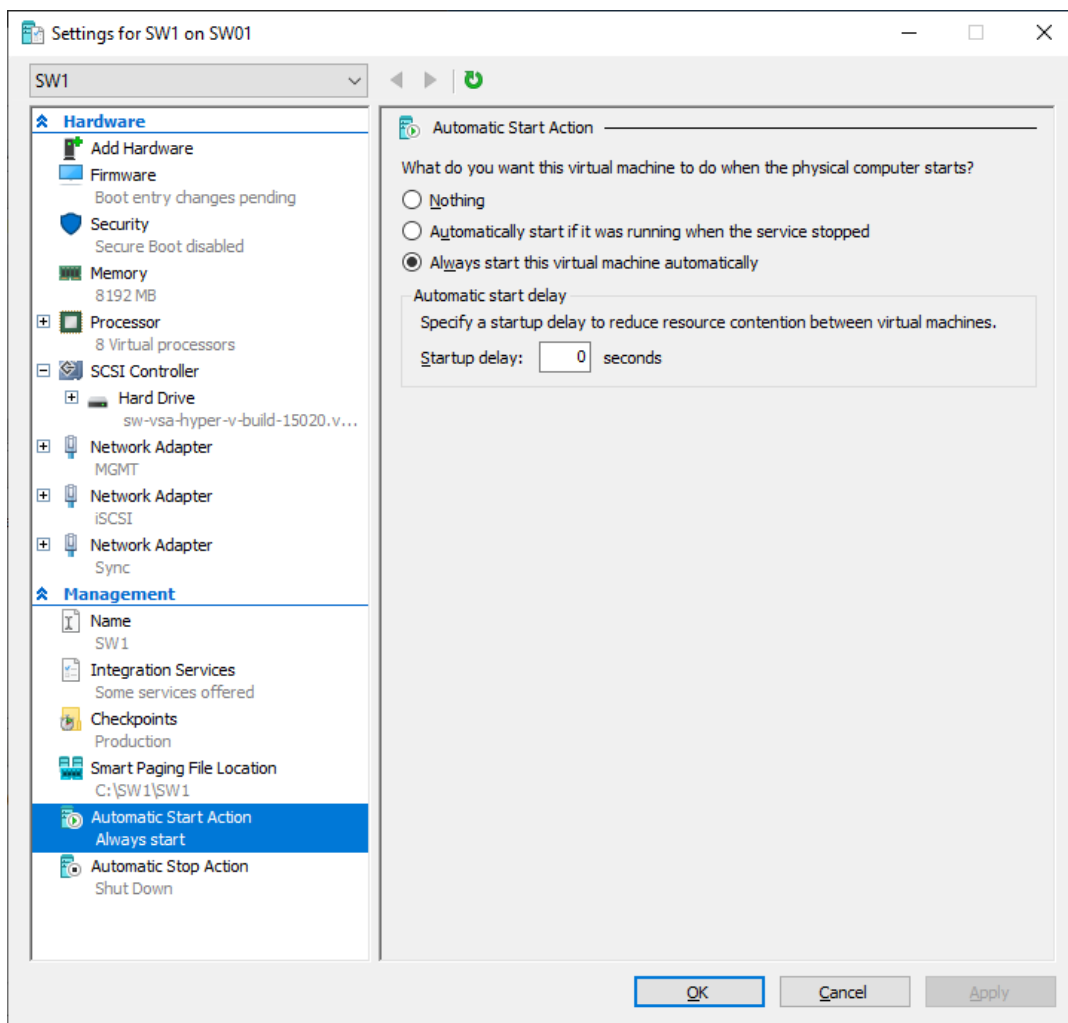


15. Click Add Hardware and add NICs for iSCSI and Synchronization traffic to the VM.

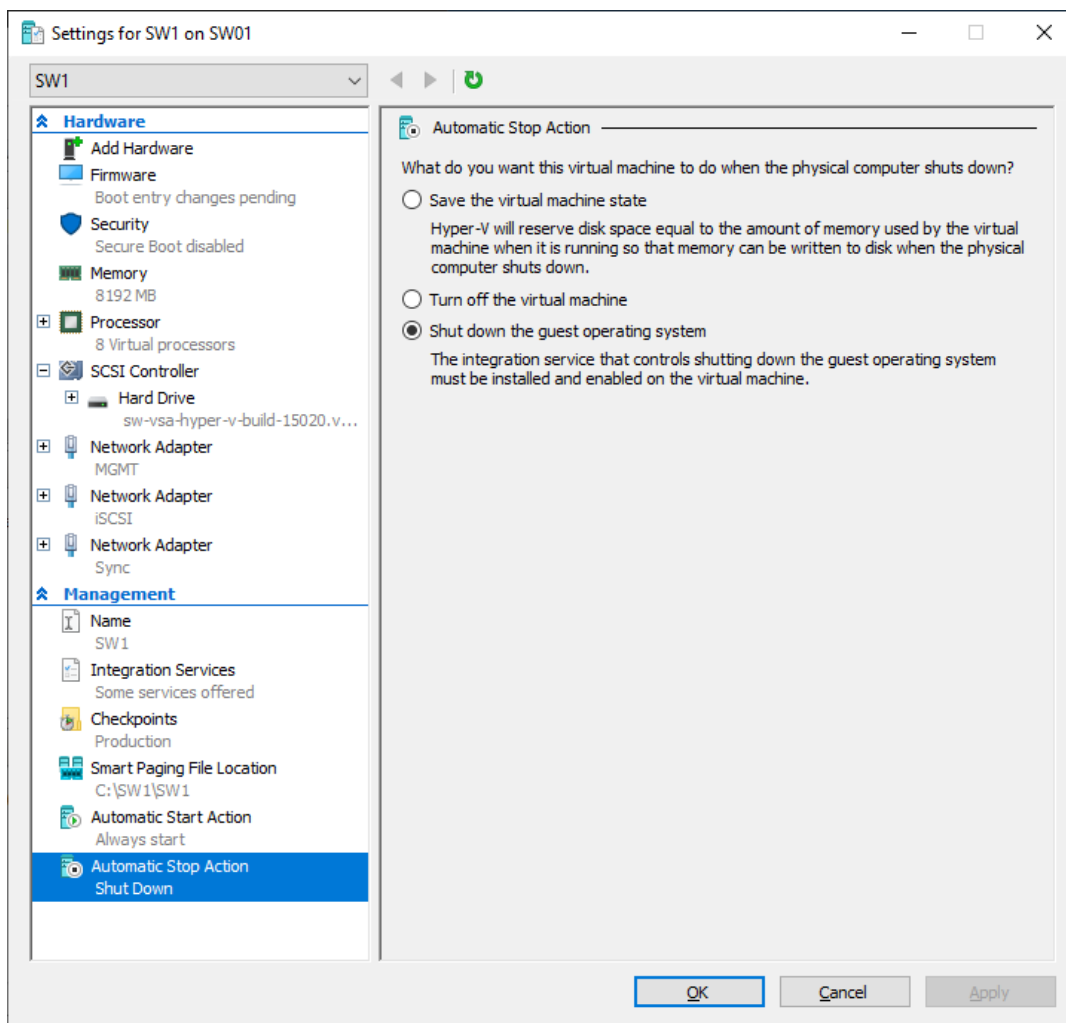




16. Configure Automatic Start and Stop actions, so the VM will always start automatically.




17. Choose Shut down the guest operating system as an Automatic Stop Action.



18. Repeat all the steps from this section on other Windows Server hosts.
19. Start virtual machines on all Windows Server hosts.

## Configuring Starwind Virtual San Vm Settings

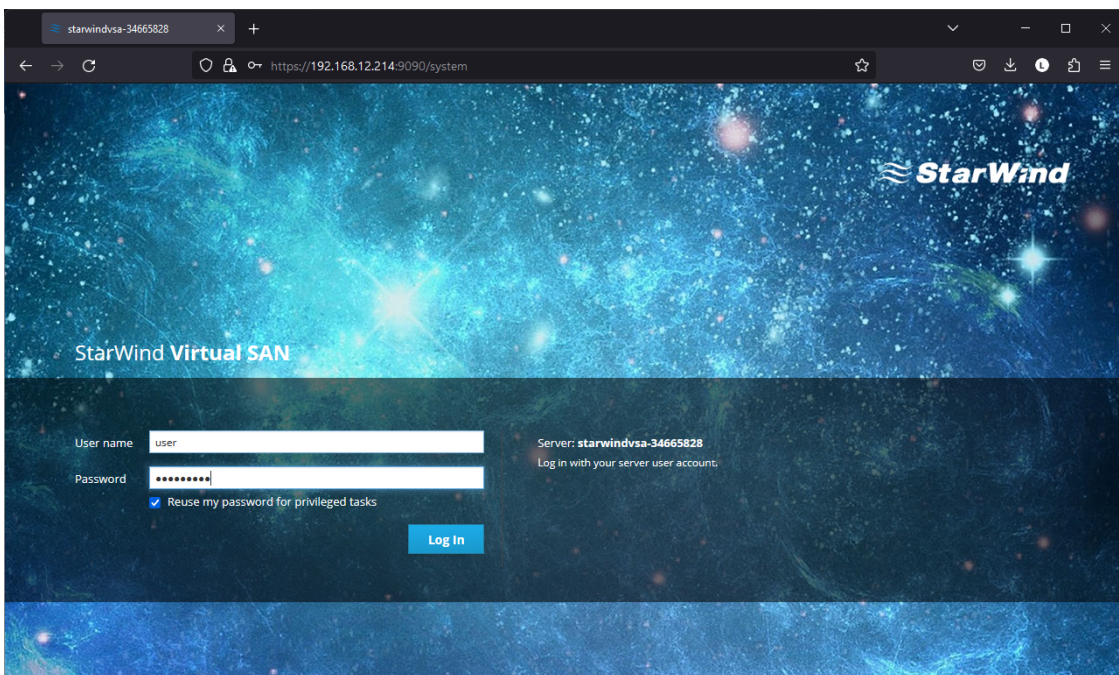
1. Open the VM console and check the IP address received via DHCP (or which was assigned manually).



Another alternative is to log into the VM via console and assign static IP using nmcli if there is no DHCP,

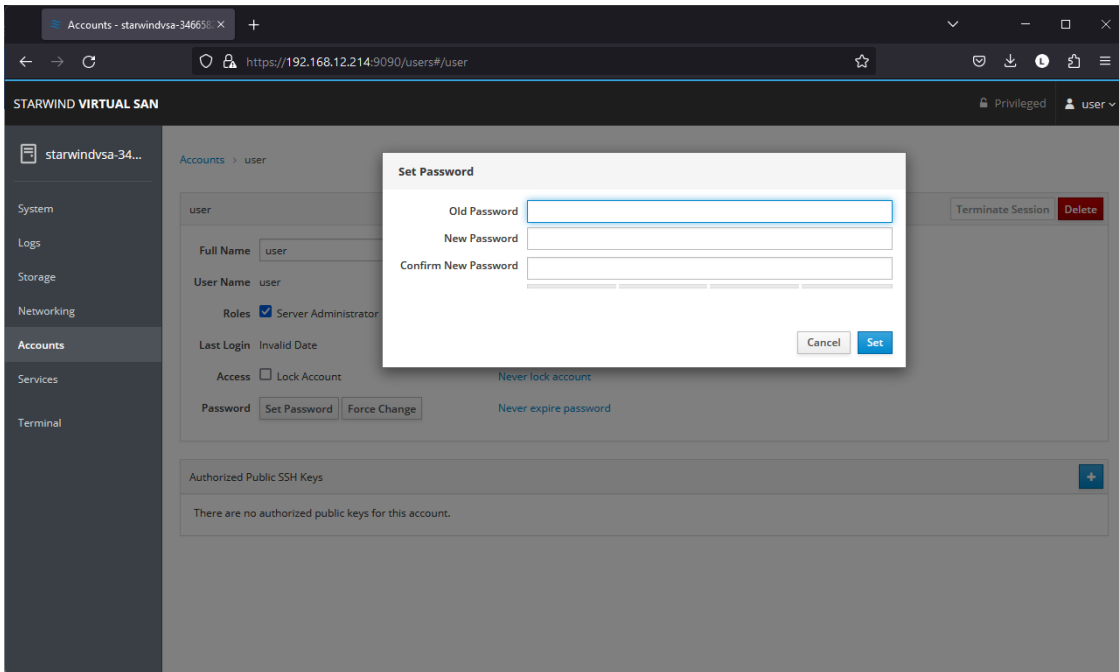
2. Now, open the web browser and enter the IP address of the VM. Log into the VM using the following default credentials:

- username: user
- password: rds123RDS
- NOTE: Make sure to check the “Reuse my password for privileged tasks” box




3. After a successful login, click Accounts on the left sidebar.

4. Select a user and click Set Password.

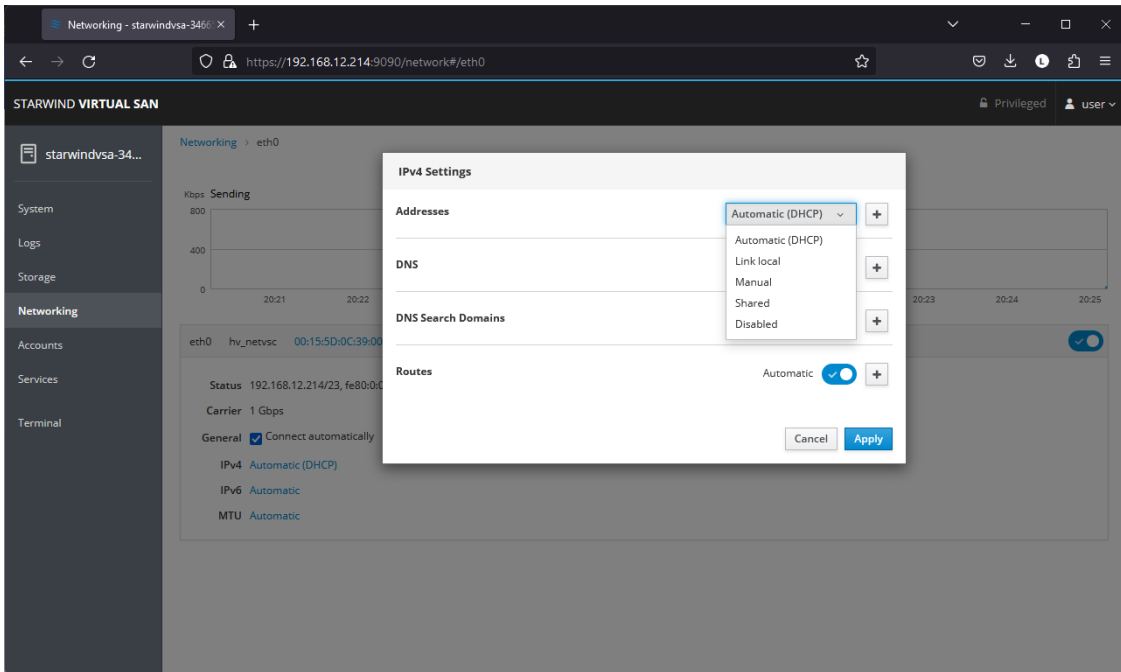


5. On the left sidebar, click Networking.

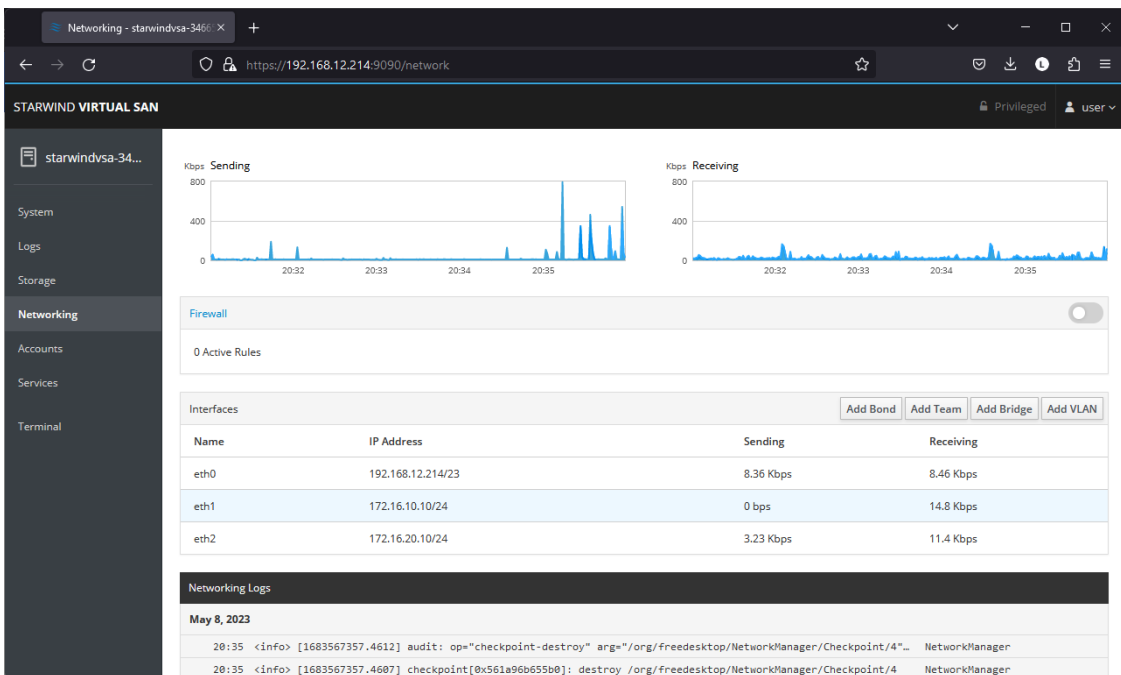


Here, the Management IP address of the StarWind Virtual SAN Virtual Machine can be configured, as well as IP addresses for iSCSI and Synchronization networks. In case the Network interface is inactive, click on the interface, turn it on, and set it to Connect automatically.


6. Click on Automatic (DHCP) to set the IP address (DNS and gateway – for Management).



7. The result should look like in the picture below:



**NOTE:** It is recommended to set MTU to 9000 on interfaces dedicated for iSCSI and Synchronization traffic. Change Automatic to 9000, if required.



8. Alternatively, log into the VM via the VMware console and assign a static IP address by editing the configuration file of the interface located by the following path:  
/etc/sysconfig/network-scripts


9. Open the file corresponding to the Management interface using a text editor, for example: `sudo nano /etc/sysconfig/network-scripts/ifcfg-eth0`

10. Edit the file:


- change the line `BOOTPROTO=dhcp` to: `BOOTPROTO=static`
- add the IP settings needed to the file:
- `IPADDR=192.168.12.10`
- `NETMASK=255.255.255.0`
- `GATEWAY=192.168.12.1`
- `DNS1=192.168.1.1`

11. Restart the interface using the following cmdlet: `sudo ifdown eth0, sudo ifup eth0`, or restart the VM.

12. Change the Host Name from the System tab by clicking on it:



13. Change System time and NTP settings if required:



14. Repeat the steps above on each StarWind VSAN VM.

## Configuring Storage


StarWind Virtual SAN for vSphere can work on top of Hardware RAID or Linux Software RAID (MDADM) inside of the Virtual Machine.

Please select the required option:


### Configuring Storage With Hardware Raid

1. Open VM Settings in Hyper-V and add drive to the VM, which going to be used by StarWind service. It is recommended to pass the entire RAID array to the VM by selecting a physical hard disk option.


NOTE: Using virtual hard disks is not recommended due to potential performance issues. Thick provisioned virtual disks should be used to improve storage performance.



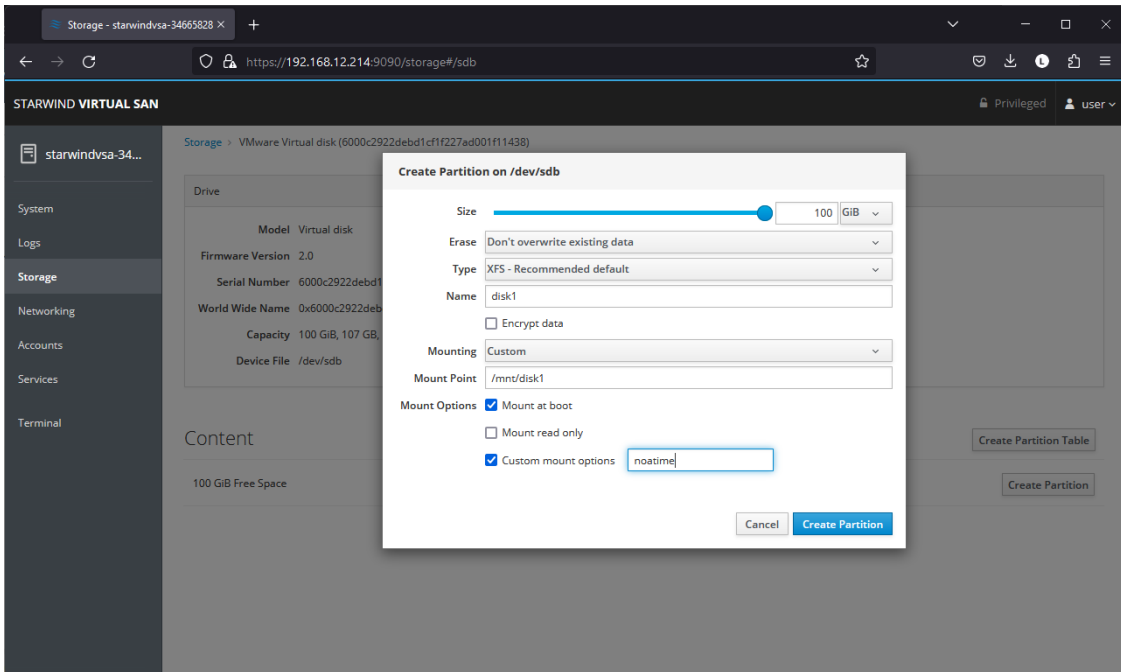
2. Login to StarWind VSAN VM web console and access the Storage section. Locate the recently added disk in the Drives section and choose it.



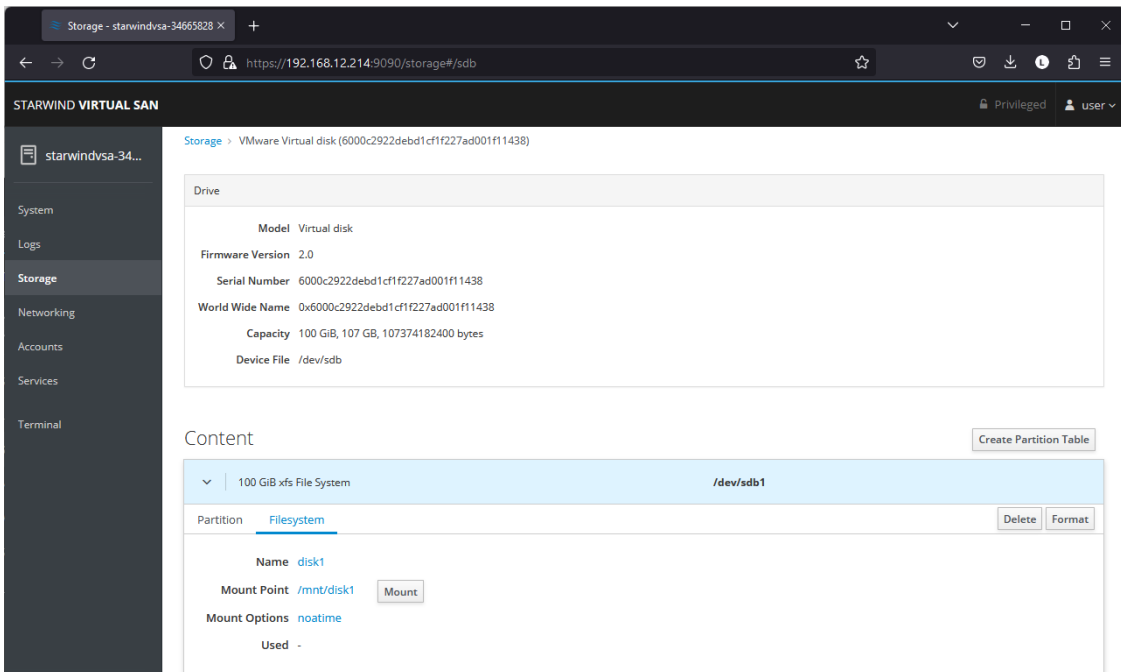
3. The added disk does not have any partitions and filesystem. Press the Create Partition Table button to create the partition.



4. Press Create Partition to format the disk and set the mount point. The mount point should be as follows: /mnt/%yourdiskname%



5. On the Storage section, under Content, navigate to the Filesystem tab. Click Mount.



## Configuring Starwind Management Console


1. Install StarWind Management Console on each server or on a separate workstation with Windows OS (Windows 7 or higher, Windows Server 2008 R2 and higher) using the installer available here.

NOTE: StarWind Management Console and PowerShell Management Library components are required.


2. Select the appropriate option to apply the StarWind License key.

Once the appropriate license key has been received, it should be applied to StarWind Virtual SAN service via Management Console or PowerShell.

3. Open StarWind Management Console and click Add Server.




4. Type the IP address of the StarWind Virtual SAN in the pop-up window and click OK.



5. Select the server and click Connect.

6. Click Apply Key... on the pop-up window.



7. Select Load license from file and click the Load button.

8. Select the appropriate license key.

As an alternative, PowerShell can be used. Open StarWind InstallLicense.ps1 script with PowerShell ISE as administrator. It can be found here:

C:\Program Files\StarWind

Software\StarWind\StarWindX\Samples\powershell\InstallLicense.ps1

Type the IP address of StarWind Virtual SAN VM and credentials of StarWind Virtual SAN service (defaults login: root, password: starwind).

Add the path to the license key.

```

1 #
2 # The following example shows how to apply license on a server
3 #
4 Import-Module StarWindX
5
6 Enable-SWXLog
7
8 $server = New-SWServer -host 127.0.0.1 -port 3261 -user root -password starwind
9
10 try
11 {
12     $server.Connect()
13
14     Get-SWLicense $server
15
16     Remove-SWLicense $server
17
18     #apply license key
19     Set-SWLicense $server "C:\License\licensekey.swk"
20 }
21 catch
22 {
23     Write-Host $_ -foreground red
24 }
25 finally
26 {
27     $server.Disconnect()
28 }
29

```

PS C:\Program Files\StarWind Software\StarWind\StarWindX\Samples\powershell>

9. After the license key is applied, StarWind devices can be created.  
NOTE: In order to manage StarWind Virtual SAN service (e.g. create ImageFile devices, VTL devices, etc.), StarWind Management Console can be used.

## Creating Starwind Ha Luns Using Powershell

1. Open PowerShell ISE as Administrator.
2. Open StarWindX sample CreateHA\_2.ps1 using PowerShell ISE. It can be found here: C:\Program Files\StarWind Software\StarWind\StarWindX\Samples\

NOTE: The script below creates a 1TB size 2-node HA device with a heartbeat failover strategy on StarWind nodes with management IP addresses 192.168.0.1 and 192.168.0.2 correspondingly.

The IP addresses 172.16.10.1 and 172.16.10.2 are used as heartbeat interfaces along with 192.168.0.1 and 192.168.0.2 for redundancy.

The IP addresses 172.16.20.1 and 172.16.20.2 on each node correspondingly as well as 172.16.21.1 and 172.16.21.2 are used for the devices synchronization between the nodes.

The script does not create a directory. Make sure you create the directory listed as \$imagePath value before running the script.  
Make sure to open 3261 and 3260 ports.

The same approach applies to CreateHA\_3.ps1 which allow to create a 3-way replica HA device.

3. Configure script parameters according to the following example:

```
param($addr="192.168.0.1", $port=3261, $user="root",
      $password="starwind",
      $addr2="192.168.0.2", $port2=$port, $user2=$user,
      $password2=$password,
      #common
      $initMethod="NotSynchronize",
      $size=1048576,
      $sectorSize=512,
      $failover=0,
      $bmpType=1,
      $bmpStrategy=0,
      #primary node
      $imagePath="/mnt/sdb1/volume1",
      $imageName="device1",
      $createImage=$true,
      $storageName="",
      $targetAlias="target1",
      $poolName="pool1",
      $syncSessionCount=1,
      $aluaOptimized=$true,
      $cacheMode="none",
      $cacheSize=0,
      $syncInterface="#p2=172.16.20.2:3260,172.16.21.2:3260",
      $hbInterface="#p2=172.16.10.2:3260,192.168.0.2:3260",
      $createTarget=$true,
      $bmpFolderPath="",
      #secondary node
      $imagePath2="/mnt/sdb1/volume1",
      $imageName2="device1",
      $createImage2=$true,
      $storageName2="",
      $targetAlias2="target1",
      $poolName2="pool1",
```

```

    $syncSessionCount2=1,
    $aluaOptimized2=$false,
    $cacheMode2=$cacheMode,
    $cacheSize2=$cacheSize,
    $syncInterface2="#p1=172.16.20.1:3260,172.16.21.1:3260",
    $hbInterface2="#p1=172.16.10.1:3260,192.168.0.1:3260",
    $createTarget2=$true,
    $bmpFolderPath2=""
)
Import-Module StarWindX

try
{
    Enable-SWXLog -level SW_LOG_LEVEL_DEBUG

    $server = New-SWServer -host $addr -port $port -user
$user -password $password

    $server.Connect()

    $firstNode = new-Object Node

    $firstNode.HostName = $addr
    $firstNode.HostPort = $port
    $firstNode.Login = $user
    $firstNode.Password = $password
    $firstNode.ImagePath = $imagePath
    $firstNode.ImageName = $imageName
    $firstNode.Size = $size
    $firstNode.CreateImage = $createImage
    $firstNode.StorageName = $storageName
    $firstNode.TargetAlias = $targetAlias
    $firstNode.AutoSynch = $autoSynch
    $firstNode.SyncInterface = $syncInterface
    $firstNode.HBInterface = $hbInterface
    $firstNode.PoolName = $poolName
    $firstNode.SyncSessionCount = $syncSessionCount
    $firstNode.ALUAOptimized = $aluaOptimized
    $firstNode.CacheMode = $cacheMode
    $firstNode.CacheSize = $cacheSize
    $firstNode.FailoverStrategy = $failover
    $firstNode.CreateTarget = $createTarget
    $firstNode.BitmapStoreType = $bmpType
    $firstNode.BitmapStrategy = $bmpStrategy

```

```

    $firstNode.BitmapFolderPath = $bmpFolderPath
    #
    # device sector size. Possible values: 512 or 4096(May
be incompatible with some clients!) bytes.
    #
    $firstNode.SectorSize = $sectorSize
    $secondNode = new-Object Node

    $secondNode.HostName = $addr2
    $secondNode.HostPort = $port2
    $secondNode.Login = $user2
    $secondNode.Password = $password2
    $secondNode.ImagePath = $imagePath2
    $secondNode.ImageName = $imageName2
    $secondNode.CreateImage = $createImage2
    $secondNode.StorageName = $storageName2
    $secondNode.TargetAlias = $targetAlias2
    $secondNode.AutoSynch = $autoSynch2
    $secondNode.SyncInterface = $syncInterface2
    $secondNode.HBInterface = $hbInterface2
    $secondNode.SyncSessionCount = $syncSessionCount2
    $secondNode.ALUAOptimized = $aluaOptimized2
    $secondNode.CacheMode = $cacheMode2
    $secondNode.CacheSize = $cacheSize2
    $secondNode.FailoverStrategy = $failover
    $secondNode.CreateTarget = $createTarget2
    $secondNode.BitmapFolderPath = $bmpFolderPath2
    $device = Add-HADevice -server $server -firstNode
$firstNode -secondNode $secondNode -initMethod $initMethod
    while ($device.SyncStatus -ne
[SwHaSyncStatus]::SW_HA_SYNC_STATUS_SYNC)
    {
        $syncPercent =
$device.GetPropertyValue("ha_synch_percent")
        Write-Host "Synchronizing: $($syncPercent)%" -
foreground yellow

        Start-Sleep -m 2000

        $device.Refresh()
    }
}
catch
{

```

```

        Write-Host $_ -foreground red
    }
    finally
    {
        $server.Disconnect()
    }

```

Detailed explanation of script parameters:

-addr, -addr2 — partner nodes IP address.

Format: string. Default value: 192.168.0.1, 192.168.0.2

allowed values: localhost, IP-address

-port, -port2 — local and partner node port.

Format: string. Default value: 3261

-user, -user2 — local and partner node user name.

Format: string. Default value: root

-password, -password2 — local and partner node user password.

Format: string. Default value: starwind

#common

-initMethod - set the initial synchronization option.

Format: string.

Values:

Clear - default

NotSynchronize - skips synchronization (works ONLY IF THERE IS NO DATA TO SKIP THE ORIGINAL SYNCHRONIZATION).

SyncFromFirst or SyncFromSecond or SyncFromThird - runs full synchronization from the specific node. Use it for recreating replicas.

-size - set size for HA-device (in MB)

Format: integer. Default value: 12

-sectorSize - set sector size for HA-device

Format: integer. Default value: 512

allowed values: 512, 4096

-failover - set type failover strategy

Format: integer. Default value: 0 (Heartbeat)

allowed values: 0, 1 (Node Majority)

-bmpType - set bitmap type, is set for both partners at once

Format: integer. Default value: 1 (RAM)

allowed values: 1, 2 (DISK)

-bmpStrategy – set journal strategy, is set for both partners at once

Format: integer. Default value: 0

allowed values: 0, 1 – Best Performance (Failure), 2 – Fast Recovery (Continuous)

-storageName is used only if you plan to add the partner to the existing device. For CreateHA\_2.ps1 use, leave it as is.

#primary node

-imagePath – set path to store the device file

Format: string. Default value: My computer\C\starwind". For Linux the following format should be used: "VSA Storage\mnt\mount\_point"

-imageName – set name device

Format: string. Default value: masterImg21

-createImage – set create image file

Format: boolean. Default value: true

-targetAlias – set alias for target

Format: string. Default value: targetha21

-poolName – set storage pool. Do not change it and keep default value.

Format: string. Default value: pool1

-aluaOptimized – set Alua Optimized

Format: boolean. Default value: true

-cacheMode – set type L1 cache (optional parameter)

Format: string. Default value: wb

allowed values: none, wb, wt

-cacheSize – set size for L1 cache in MB (optional parameter)

Format: integer. Default value: 128

allowed values: 1 and more

-syncInterface – set sync channel IP-address from partner node

Format: string. Default value: "#p2={0}:3260"

-hbInterface – set heartbeat channel IP-address from partner node

Format: string. Default value: ""

-createTarget – set creating target

Format: string. Default value: true

Even if you do not specify the parameter -createTarget, the target will be created automatically.

If the parameter is set as -createTarget \$false, then an attempt will be made to create the device with existing targets, the names of which are specified in the -targetAlias (targets must already be created)

-bmpFolderPath – set path to save bitmap file

Format: string.

#secondary node

-imagePath2 – set path to store the device file

Format: string. Default value: "My computer\C\starwind". For Linux the following format

should be used: "VSA Storage\mnt\mount\_point"

-imageName2 - set name device

Format: string. Default value: masterImg21

-createImage2 - set create image file

Format: boolean. Default value: true

-targetAlias2 - set alias for targetFormat: string.

Default value: targetha22

-poolName2 - set storage pool. Do not change it and keep default value.

Format: string. Default value: pool1

-aluaOptimized2 - set Alua Optimized

Format: boolean. Default value: true

-cacheMode2 - set type L1 cache (optional parameter)

Format: string. Default value: wb

allowed values: wb, wt

-cacheSize2 - set size for L1 cache in MB (optional parameter)

Format: integer. Default value: 128

allowed values: 1 and more

-syncInterface2 - set sync channel IP-address from partner node

Format: string. Default value: "#p1={0}:3260"

-hbInterface2 - set heartbeat channel IP-address from partner node

Format: string. Default value: ""

-createTarget2 - set creating target

Format: string. Default value: true

Even if you do not specify the parameter -createTarget, the target will be created automatically. If the parameter is set as -createTarget \$false, then an attempt will be made to create the device with existing targets, the names of which are specified in the -targetAlias (targets must already be created)

-bmpFolderPath2 - set path to save bitmap file

Format: string.

**IMPORTANT:** If the script should be executed again with the same parameters, (for example, the first time execution has failed) make sure to do the following for one node at a time before the next attempt to execute the script:

1. Stop StarWind Service:

```
sudo systemctl stop starwind-virtual-san
```

2. Go to /opt/starwind/starwind-virtual-san/drive\_c/starwind/headers and delete the headers you do not need.

3. Go to the underlying storage specified as \$imagePath and delete the header and imagefile there.

4. Go to the folder with StarWind.cfg (/opt/starwind/starwind-virtual-san/drive\_c/starwind/StarWind.cfg ) and copy it.

5. Edit StarWind.cfg:

```
sudo nano /opt/starwind/starwind-virtual-  
san/drive_c/starwind/StarWind.cfg
```

6. Navigate under <targets>, remove target you do not need.
7. Navigate under <devices>, and remove the device entry you do not need.
8. Start the service:

```
sudo systemctl start starwind-virtual-san
```

9. Wait for the devices synchronization.
10. Repeat for the remaining StarWind VSAN instance.

## 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:


## 2-Node: Provisioning Starwind Ha Storage To Windows Server Hosts

Discovery of StarWind iSCSI targets on Windows Server hosts


1. Open Start/Search and type "iSCSI Initiator" to Launch Microsoft iSCSI Initiator. Alternatively, start Command Prompt (CMD) and execute the following command:

```
iscsicpl
```

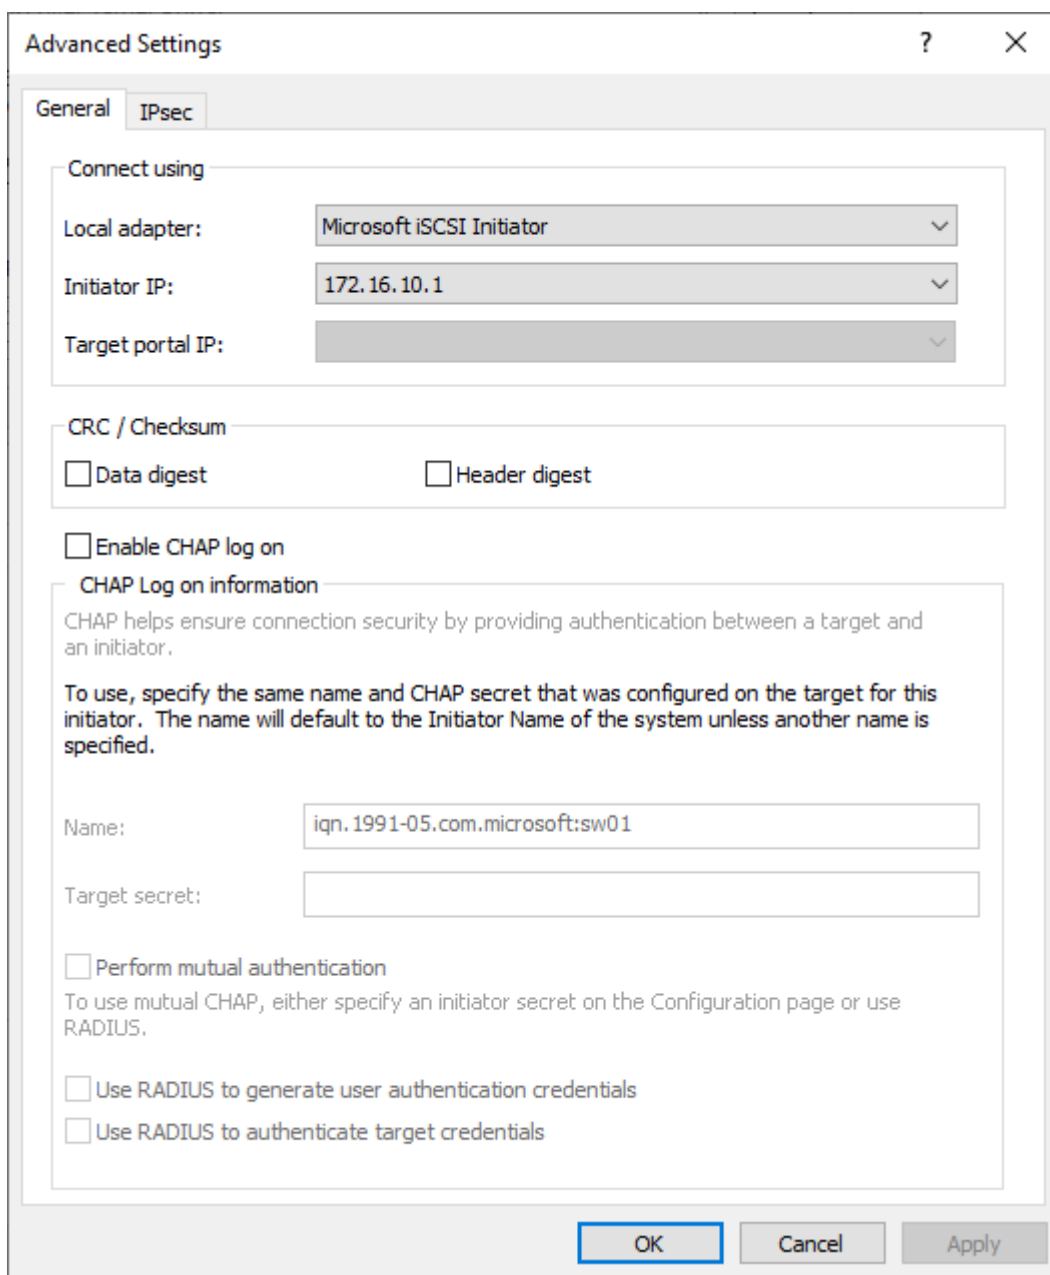
2. In the iSCSI Initiator Properties window, navigate to the Discovery tab and click Discover Portal...



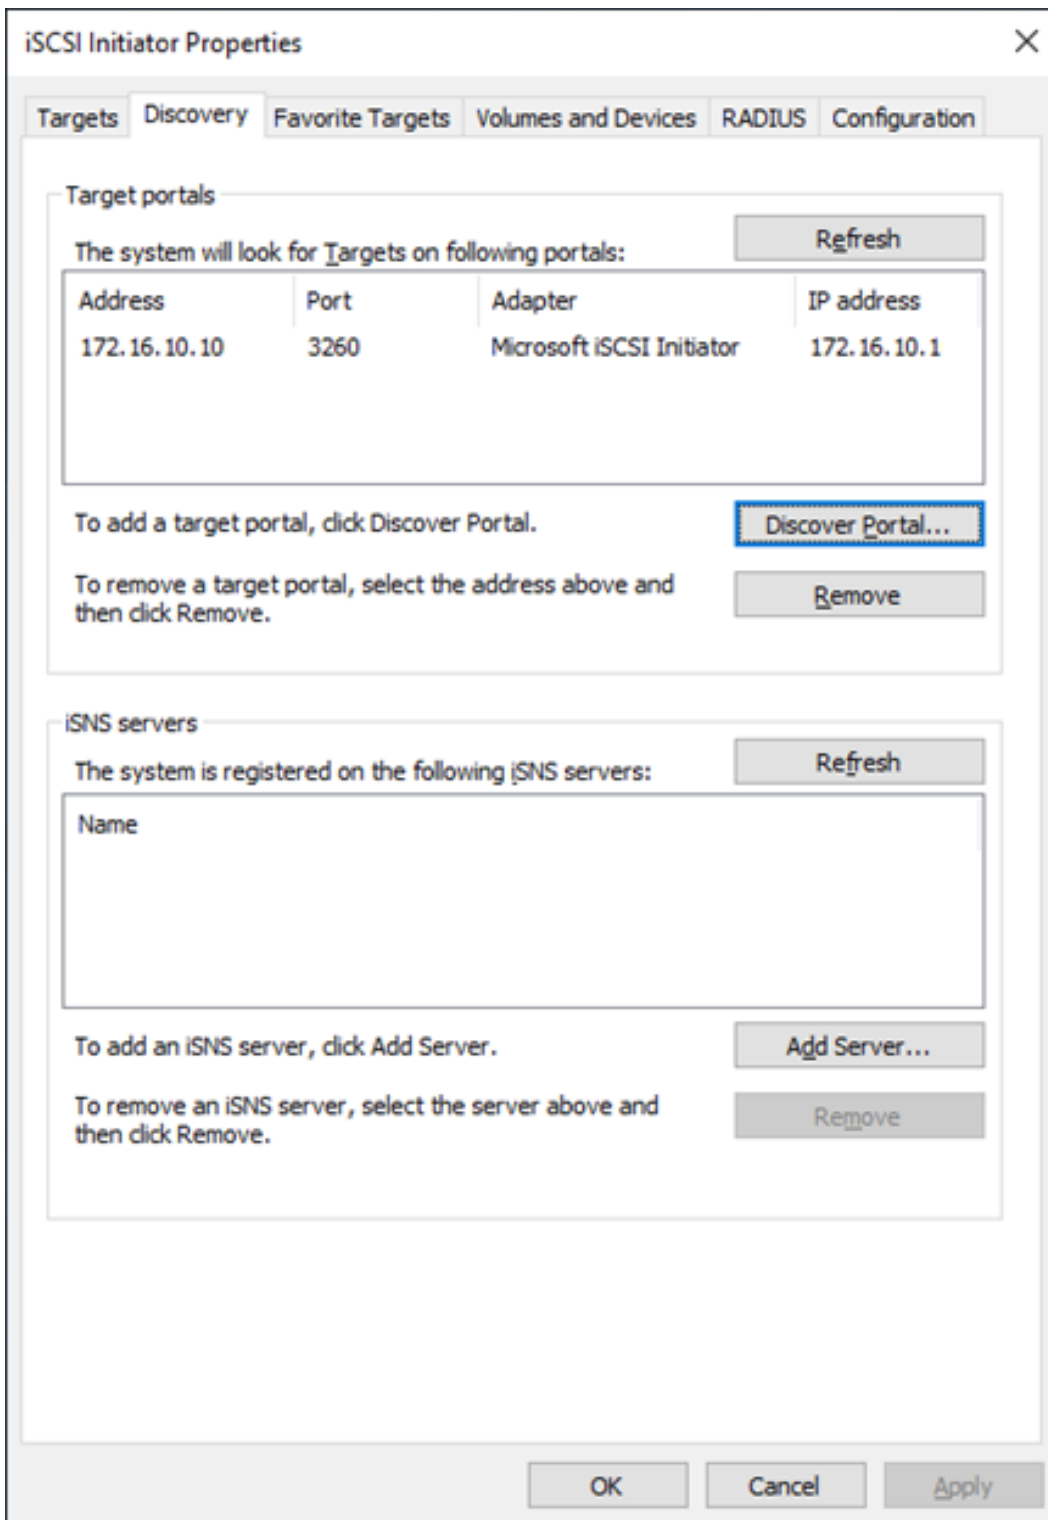
3. In the Discover Target Portal window, enter the IP address of the first StarWind Appliance’s Data network interface (for example, appliance “SW1” with the IP address 172.16.10.10), then click Advanced...



4. In the Advanced Settings window, under the Connect using section, select Microsoft iSCSI Initiator as the Local adapter. Then, select the IP address of the Data/Heartbeat network of the current Windows Server host as the Initiator IP, and click Apply.

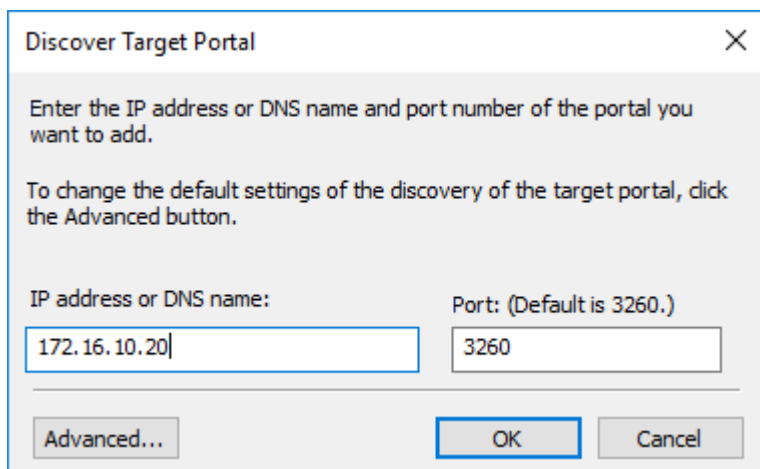


5. In the Discovery tab, click Discover Portal...

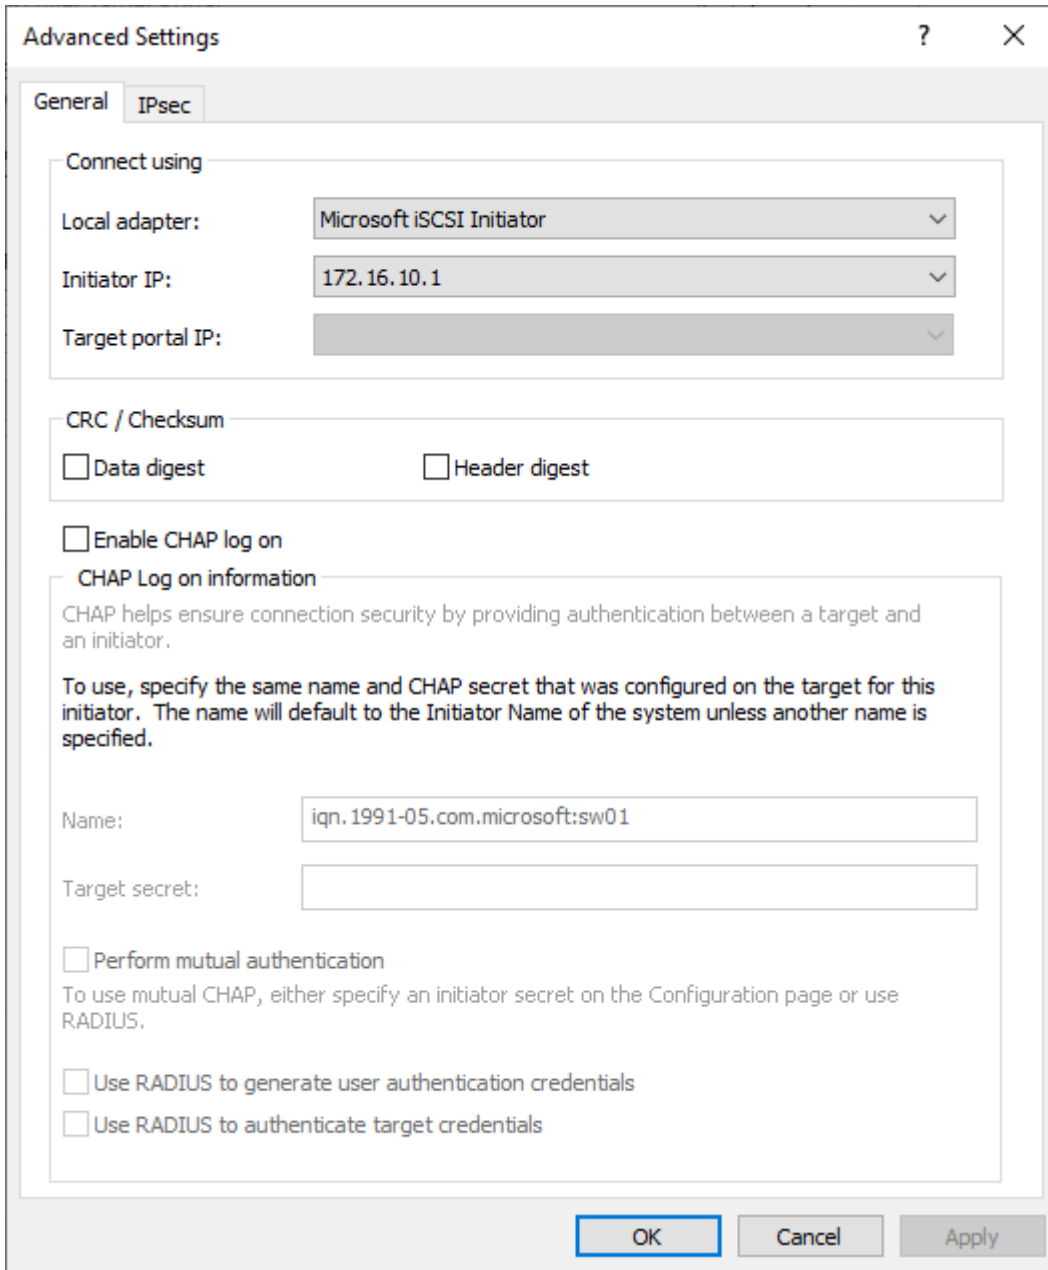


6. In the Discover Target Portal window, enter the IP address of the second StarWind Appliance’s Data network interface (for example, appliance “SW2” with the IP address


172.16.10.20), then click Advanced...



7. In the Advanced Settings window, under the Connect using section, select Microsoft iSCSI Initiator as the Local adapter. Then, select the IP address of the Data/Heartbeat network of the current Windows Server host as the Initiator IP, and click Apply.



8. Now, all iSCSI target portals have been added to the first Windows Server host.

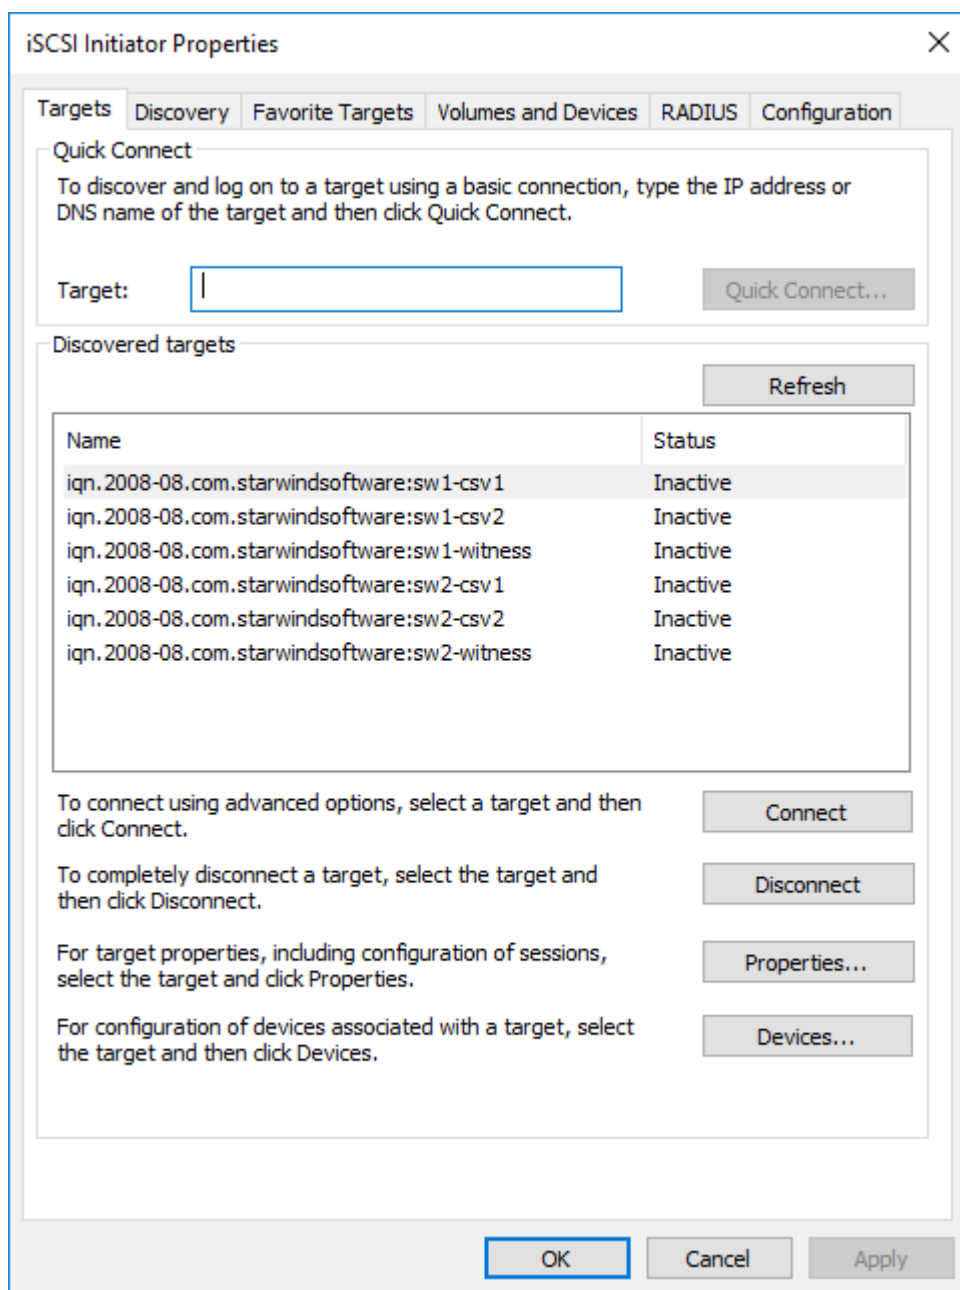


9. Repeat the steps 1-8 on the partner node.


### Connecting StarWind iSCSI targets to Windows Server hosts

1. Click the Targets tab. The previously created targets are listed in the Discovered Targets section.

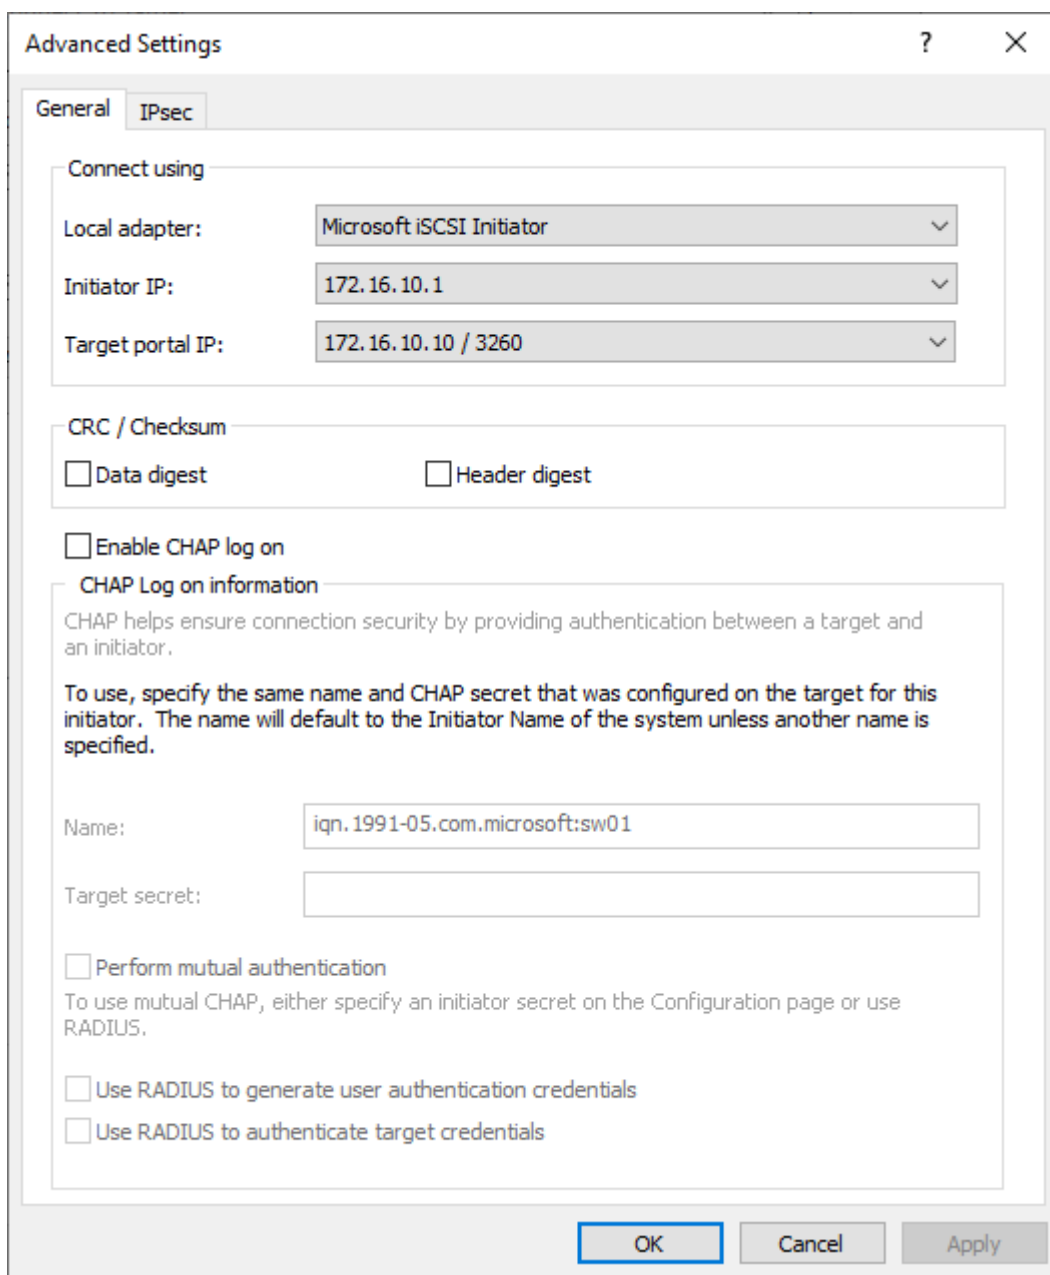
NOTE: If the created targets are not listed, check the firewall settings of the StarWind Server as well as the list of networks served by the StarWind Server (go to StarWind Management Console -> Configuration -> Network). Alternatively, check the Access Rights tab in StarWind Management Console for any restrictions.



2. Select the Witness target from the local server and click Connect.
3. Enable checkboxes as shown in the image below. Click Advanced.

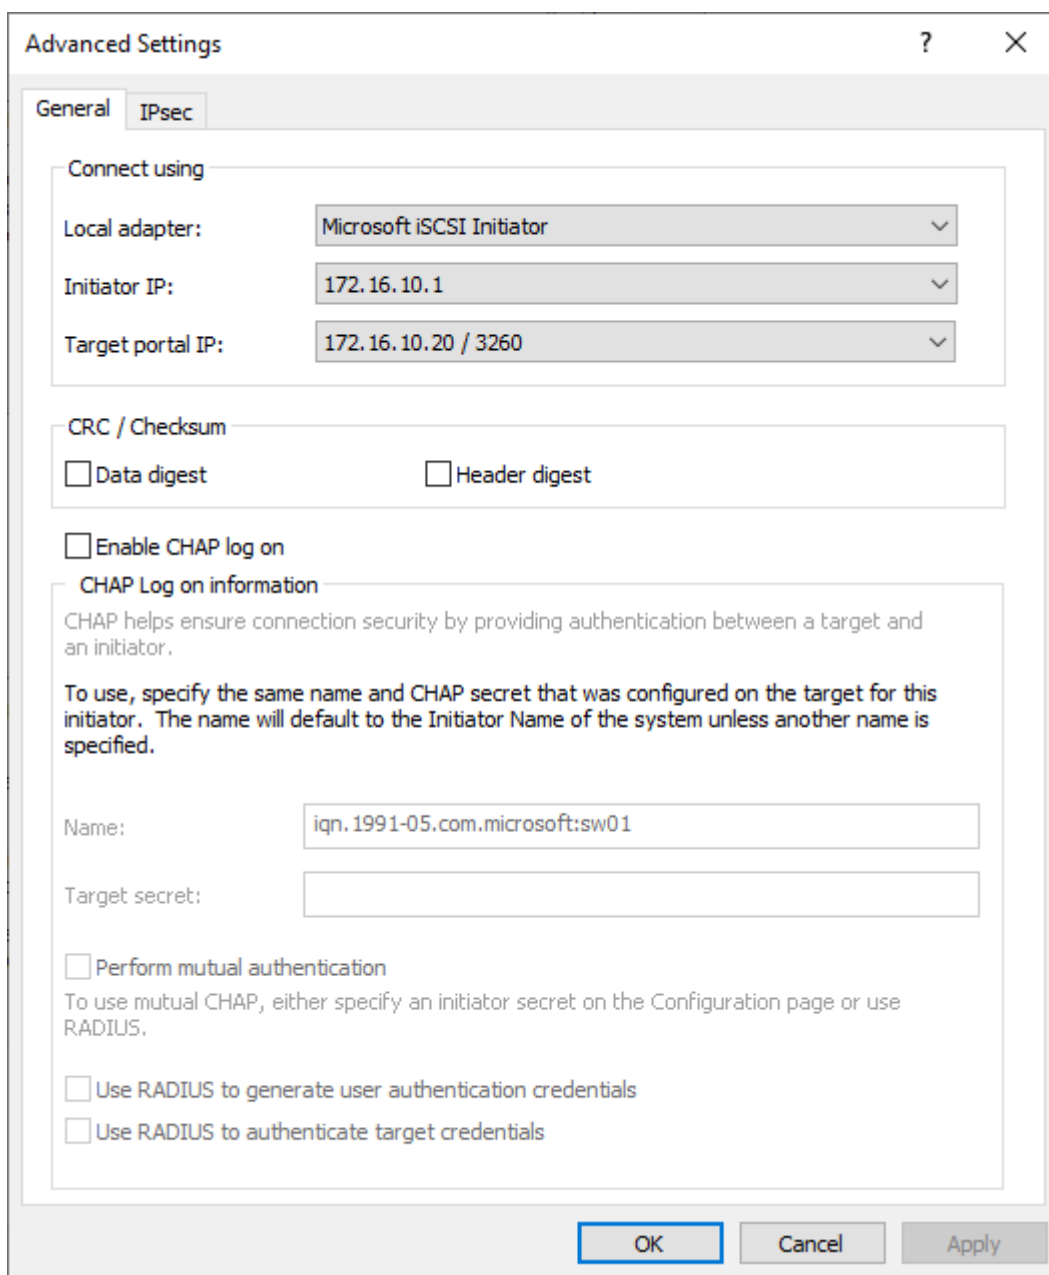


4. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In the Initiator IP field, select the IP address for the Data network. In the Target portal IP, select the portal IP from the same subnet. Confirm the actions.




5. Repeat the steps 2-4 to connect the Witness target from the partner node.

Advanced settings should look like in the picture below with the target portal IP 172.16.10.20:




6. Select the CSV1 target discovered from the local server and click Connect.

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




8. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Target portal IP, select 172.16.10.10. Confirm the actions.



9. Select the CSV1 target from the partner StarWind node and click Connect.

10. Enable checkboxes Add this connection to the list of Favorite Targets... and Enable multi-path. Click Advanced.

11. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In the Initiator IP field, select the IP address for the Data network. In the Target portal IP, select the portal IP from the same subnet. Confirm the actions.



12. Repeat the steps 1-11 for all remaining HA device targets.

13. Repeat steps 1-12 on the other StarWind node, specifying corresponding Data network IP addresses.


NOTE: Multiple (2,3 or 4) identical iSCSI connections to the same target might be configured depending on the network throughput used for Data and Replication, underlying storage performance, or CPU performance.

## Configuring Multipath

It is recommended to configure the different MPIO policies depending on iSCSI channel throughput. It is recommended to set Least Queue Depth MPIO load balancing policy.

NOTE: Different MPIO policies can be set depending on the network throughput, storage performance or setup peculiarities.

1. Configure the MPIO policy for each target 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.

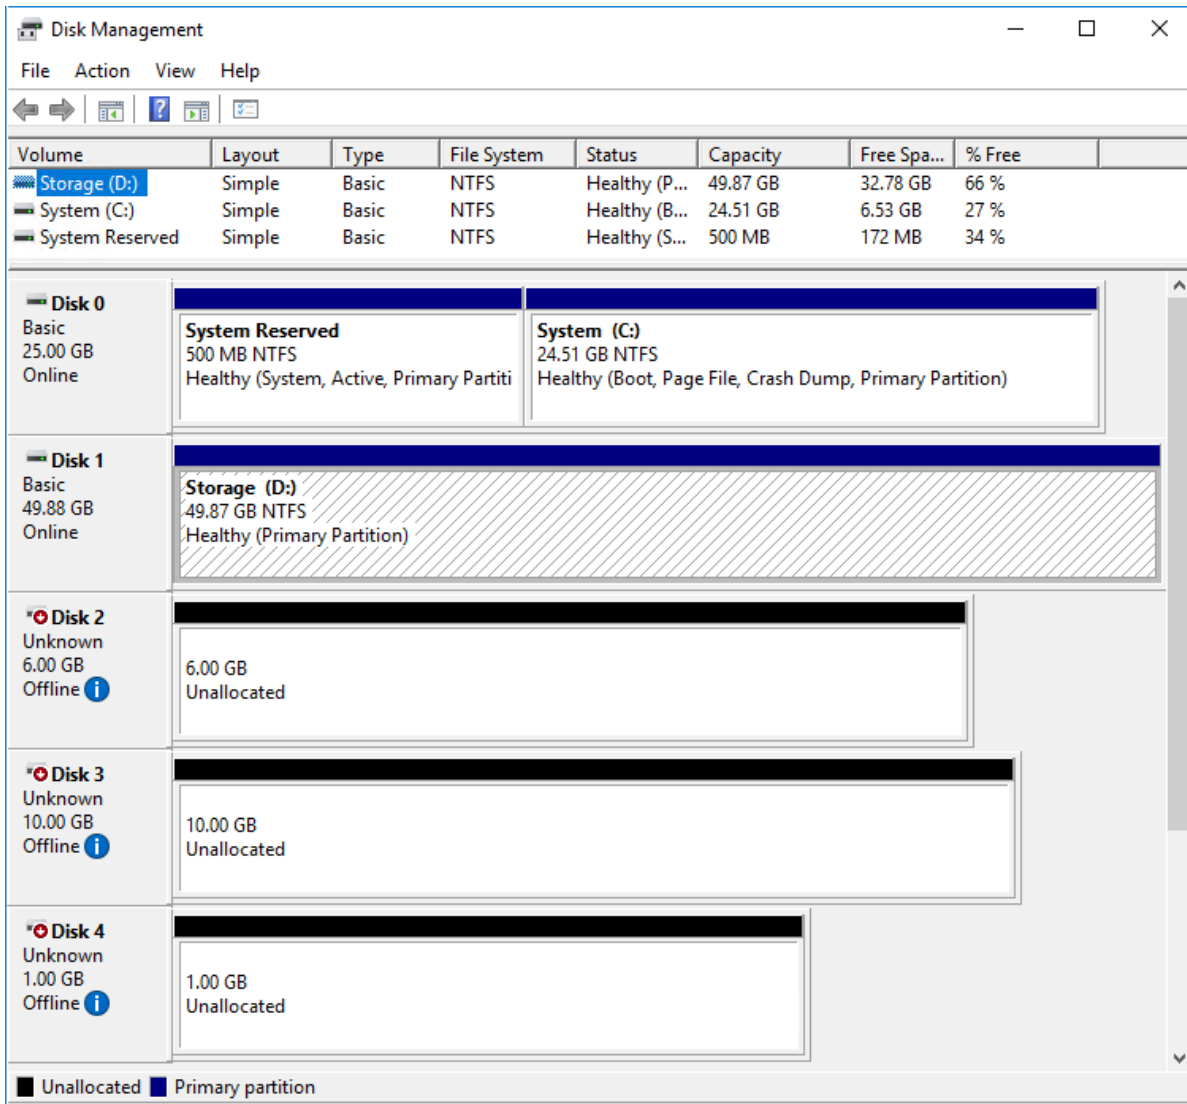


3. Select the appropriate load balancing policy.

4. Repeat the steps 1-3 for configuring the MPIO policy for each remaining device on the current node and on the partner node.

### Connecting Disks to Servers

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

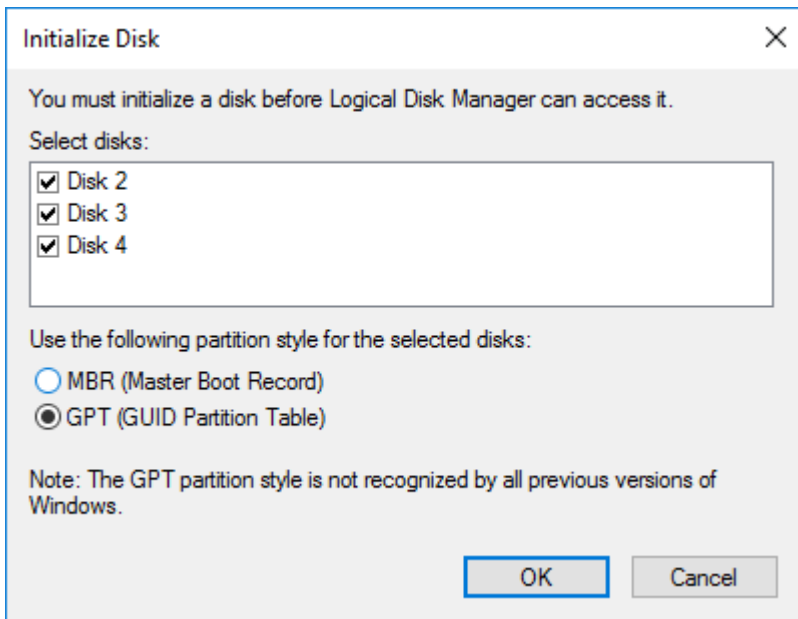


2. Bring the disks online by right-clicking on them and selecting the Online menu option.

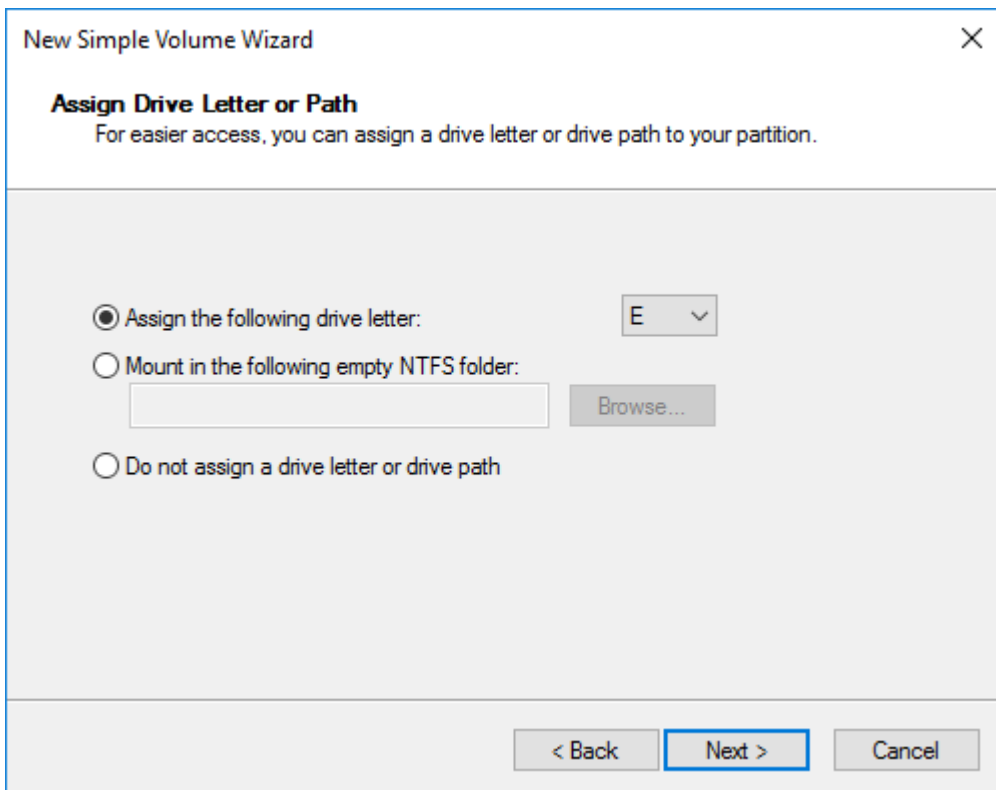
3. Select the CSV disk (check the disk size to be sure) and right-click on it to initialize.

4. By default, the system will offer to initialize all non-initialized disks. Use the Select

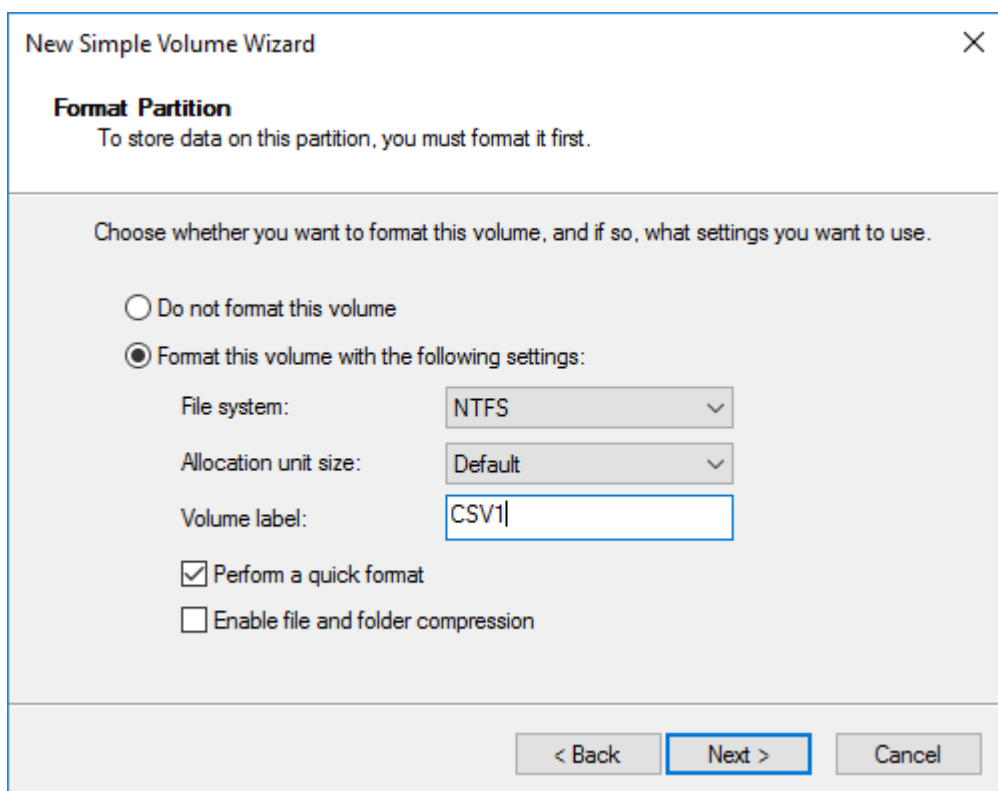
Disks area to choose the disks. Select GPT (GUID Partition Style) for the partition style to be applied to the disks. Press OK to confirm.



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.

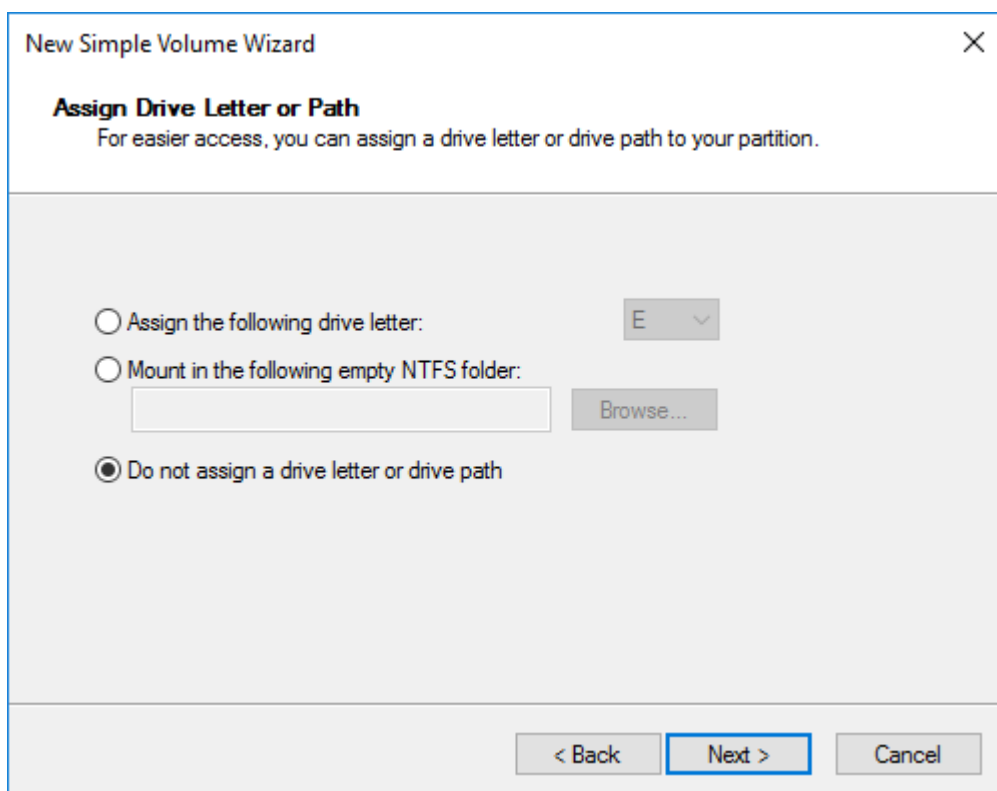


8. Select NTFS in the File System dropdown menu. Keep Allocation unit size as Default. Set the Volume Label of choice. Click Next.



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 steps from the 2nd 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.

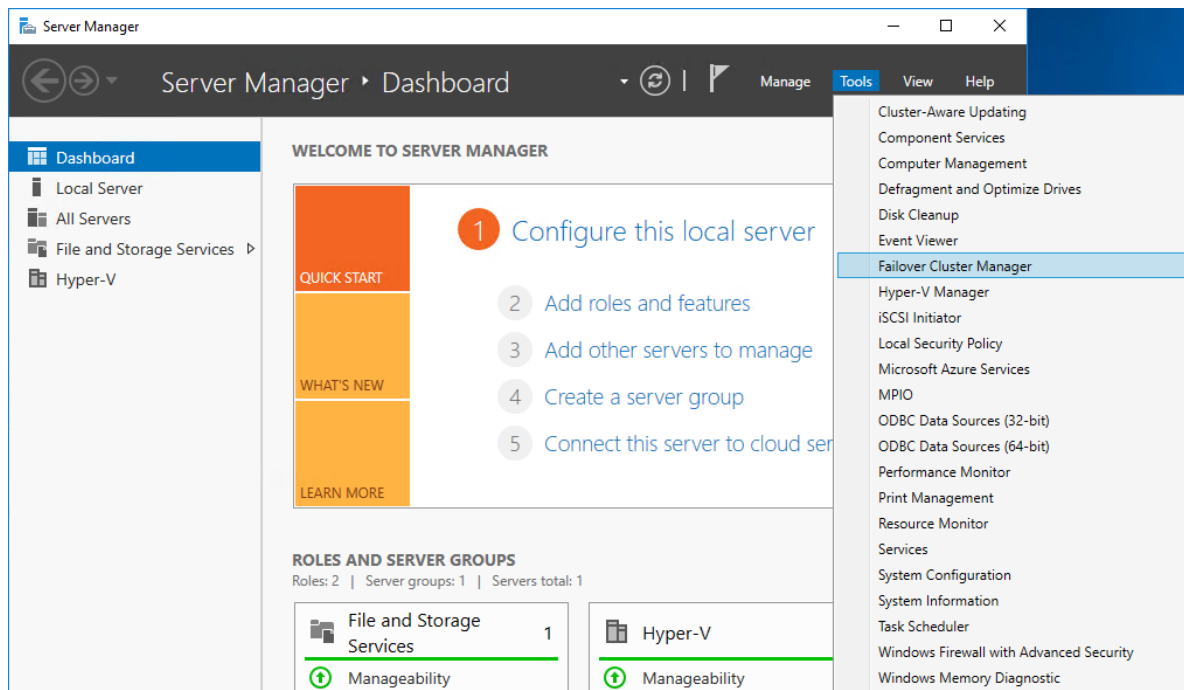
NOTE: Server Manager can be opened on the server with desktop experience enabled (necessary features should be installed). Alternatively, the Failover cluster can be managed with Remote Server Administration Tools:

<https://docs.microsoft.com/en-us/windows-server/remote/remote-server-administration-tools>

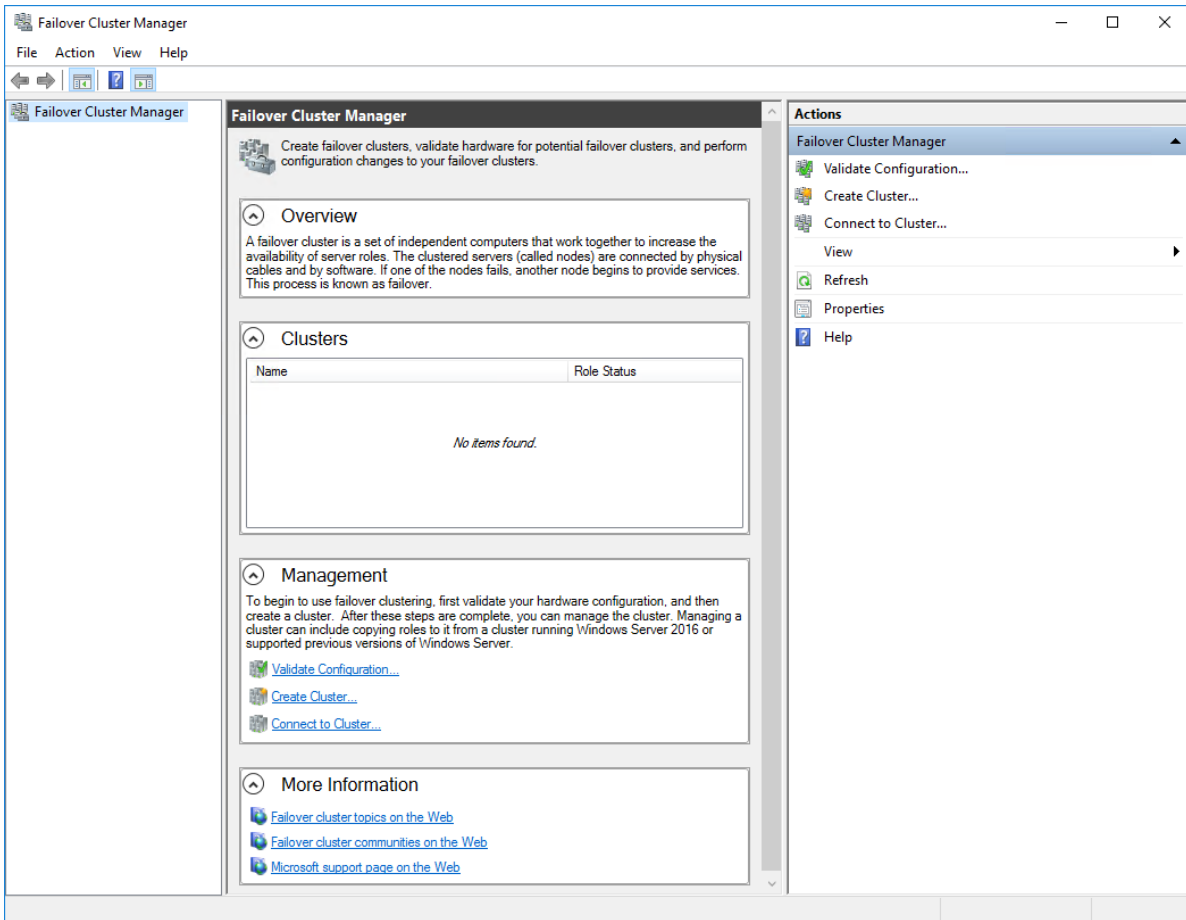
NOTE: For converged deployment (SAN & NAS running as a dedicated storage cluster) the Microsoft Failover Cluster is deployed on separate computing nodes. Additionally, for the converged deployment scenario, the storage nodes that host StarWind SAN & NAS as

CVM or bare metal do not require a domain controller and Failover Cluster to operate.


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 the 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

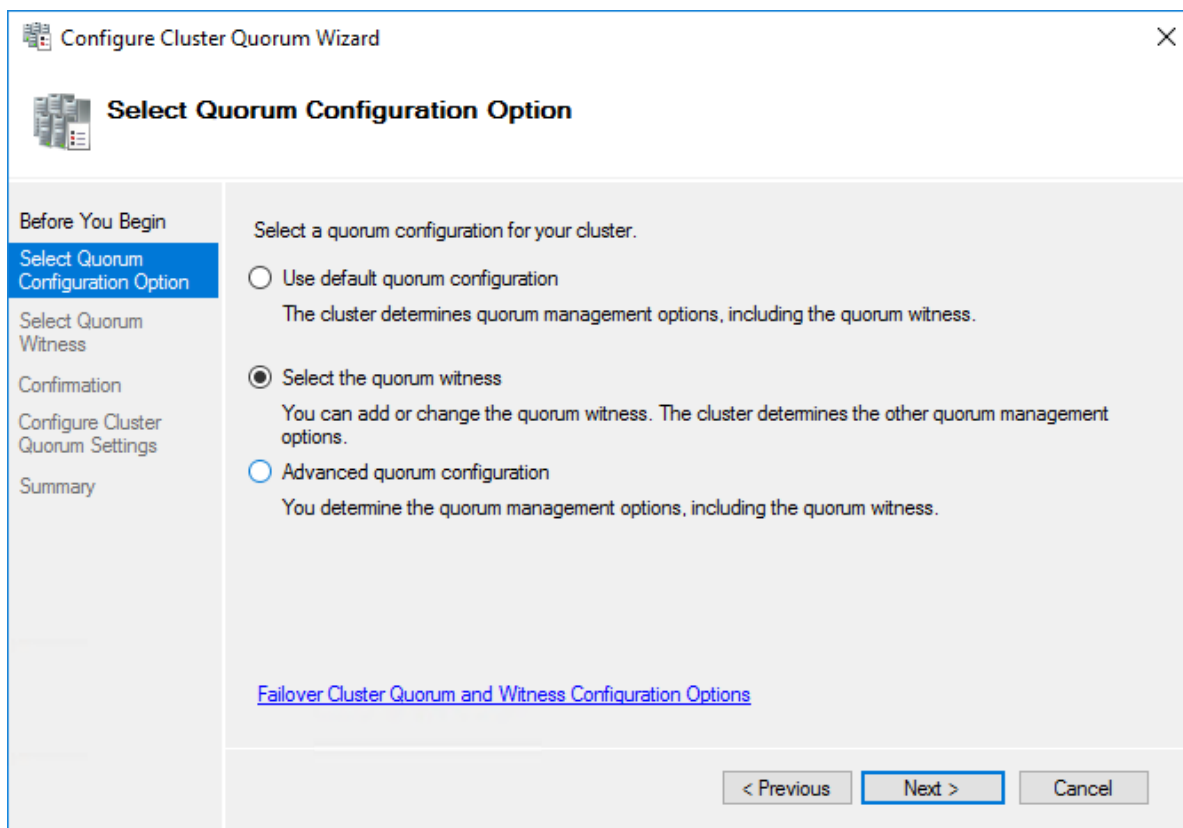
1. In Failover Cluster Manager, navigate to Cluster -> Storage -> Disks. Click Add Disk in the Actions panel, choose StarWind disks from the list and confirm the selection.



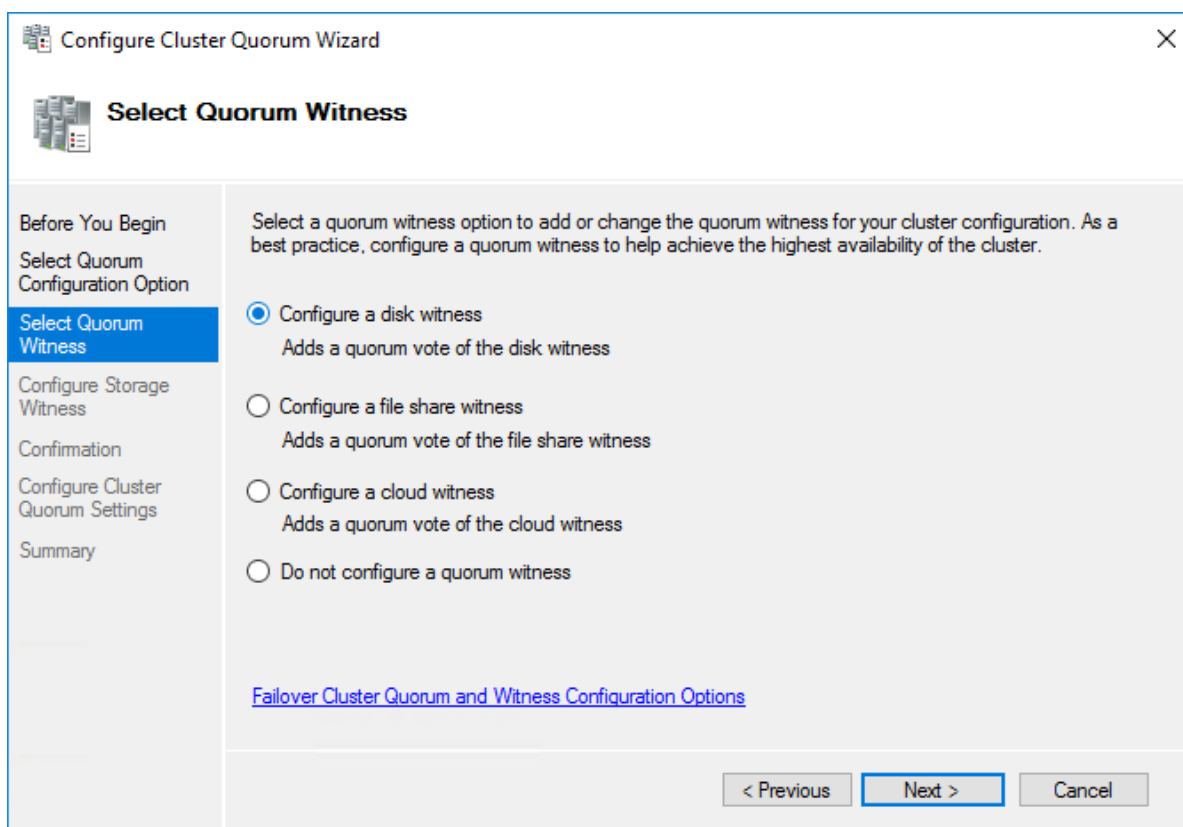
2. To configure the cluster witness disk, right-click on Cluster and proceed to More Actions -> Configure Cluster Quorum Settings.



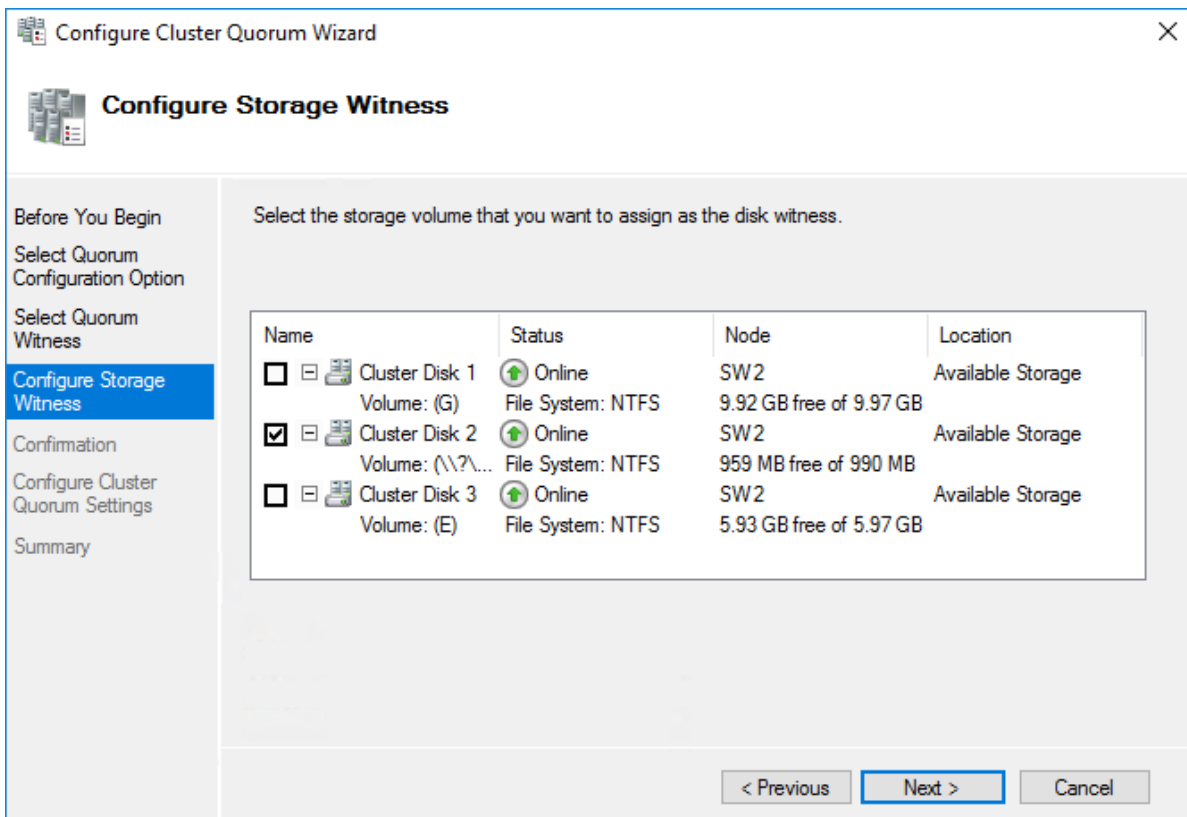
3. Follow the wizard and use the Select the quorum witness option. Click Next.




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.



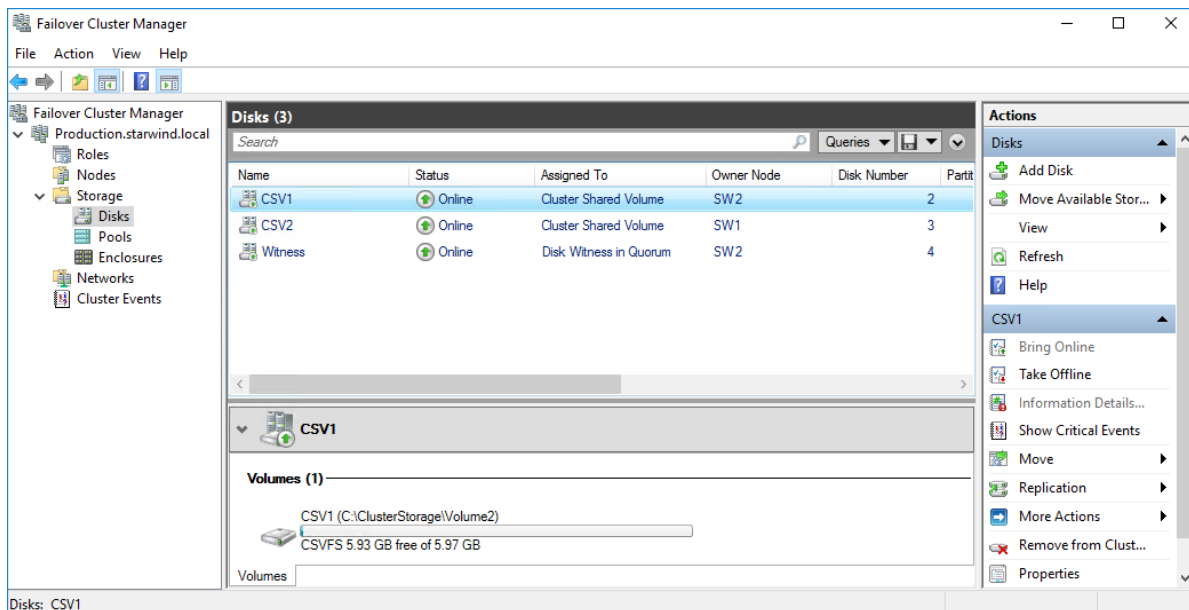
6. In Failover Cluster Manager, Right-click the disk and select Add to Cluster Shared Volumes.



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.



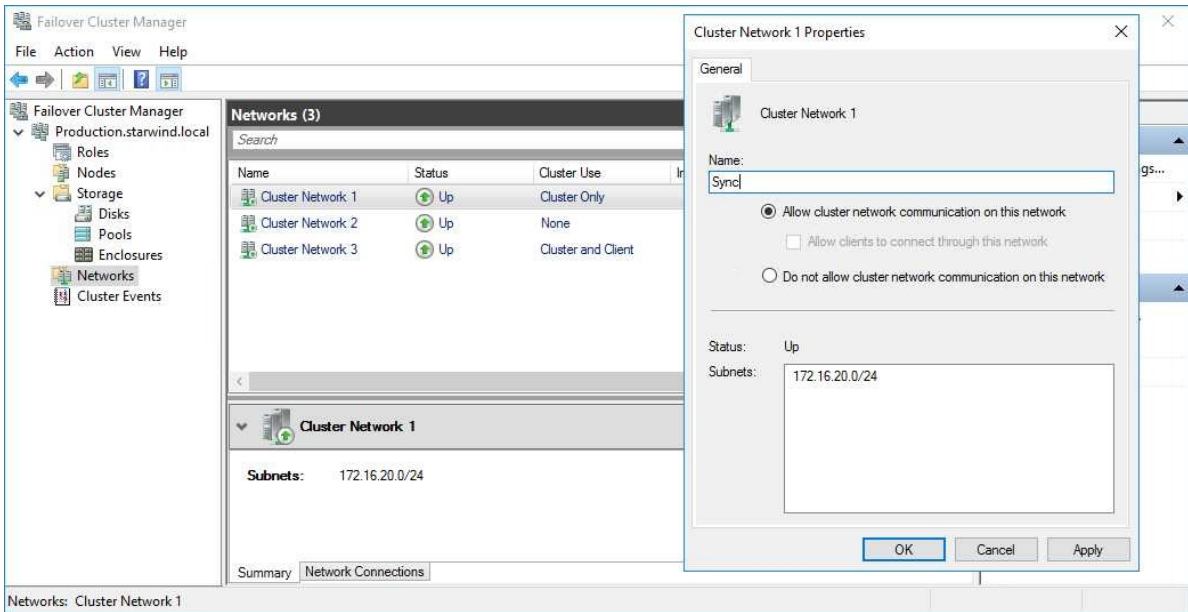
8. Perform the steps 6-7 for any other disk in Failover Cluster Manager. The resulting list of disks will look similar to the screenshot below.



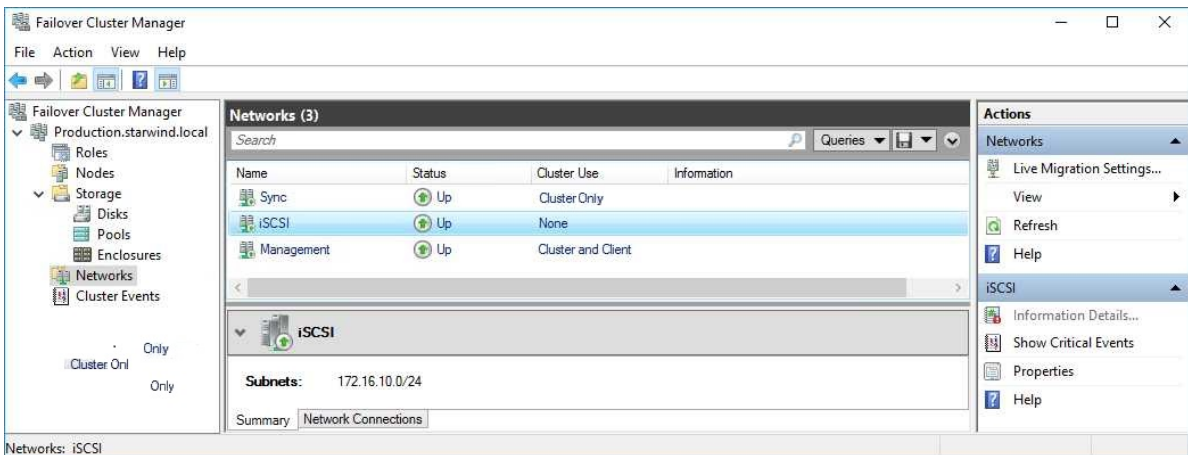
## Configuring Cluster Network Preferences

1. In the Networks section of the Failover Cluster Manager, right-click on the network from the list. Set its new name if required to identify the network by its subnet. Apply the change and press OK.

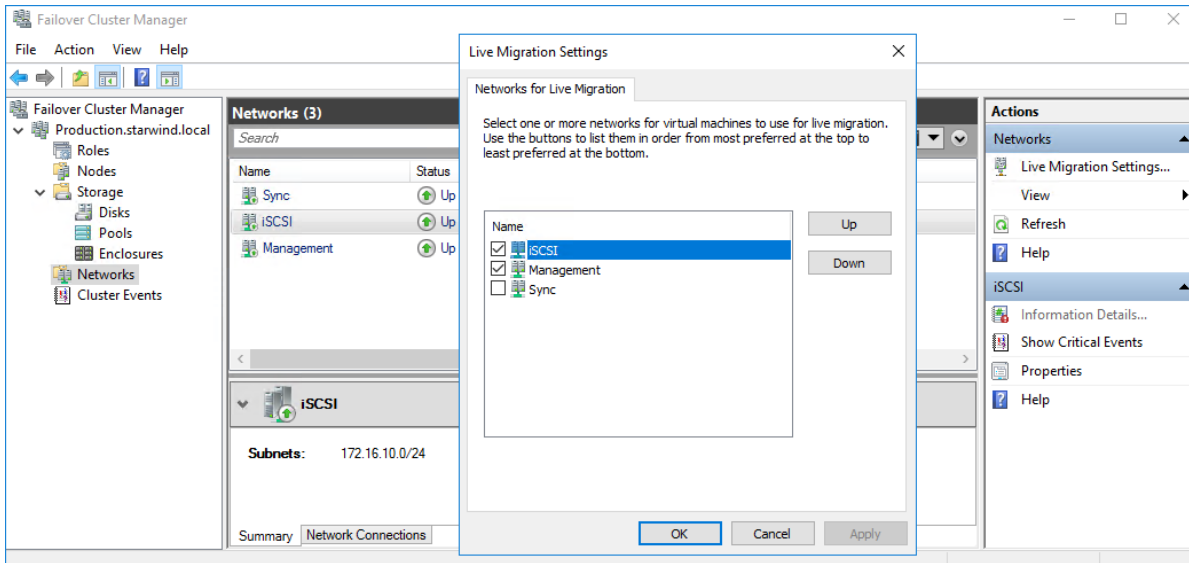
NOTE: Double-check that cluster communication is configured with redundant networks: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/smb-multichannel>



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

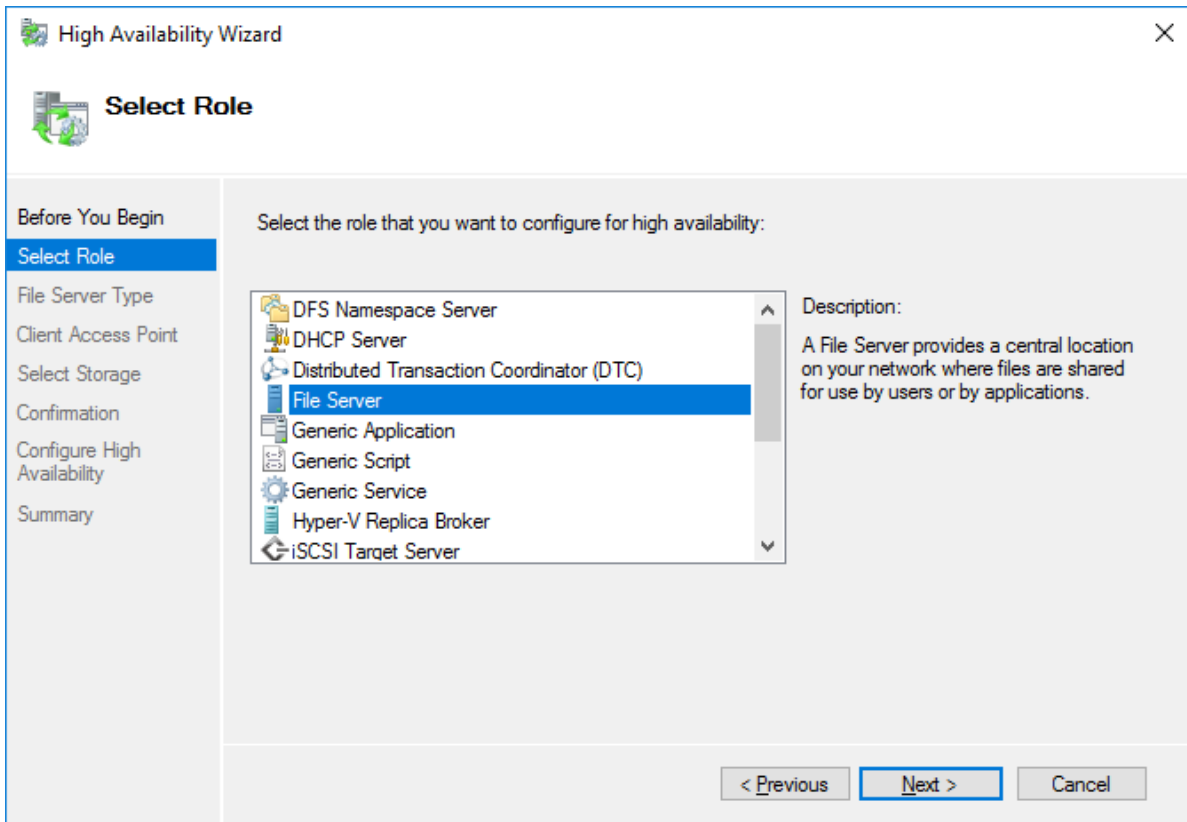
Please follow the steps below if file shares should be configured on cluster nodes.

## 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.

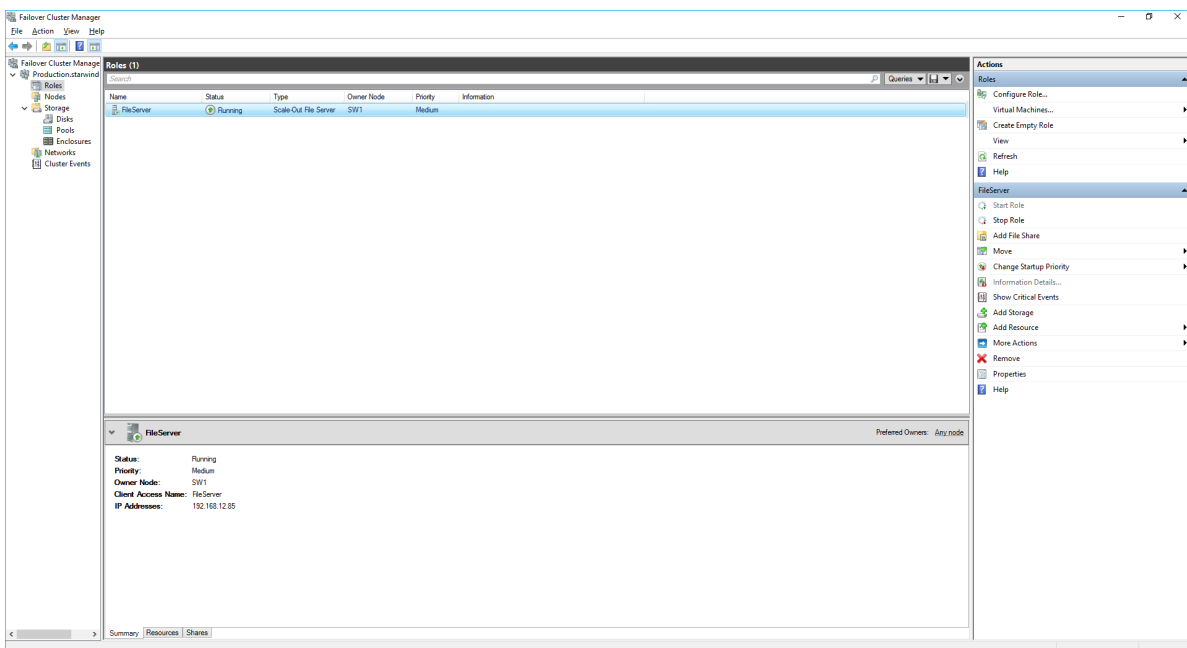


7. Once the installation is finished successfully, the Wizard should now look like the screenshot below.

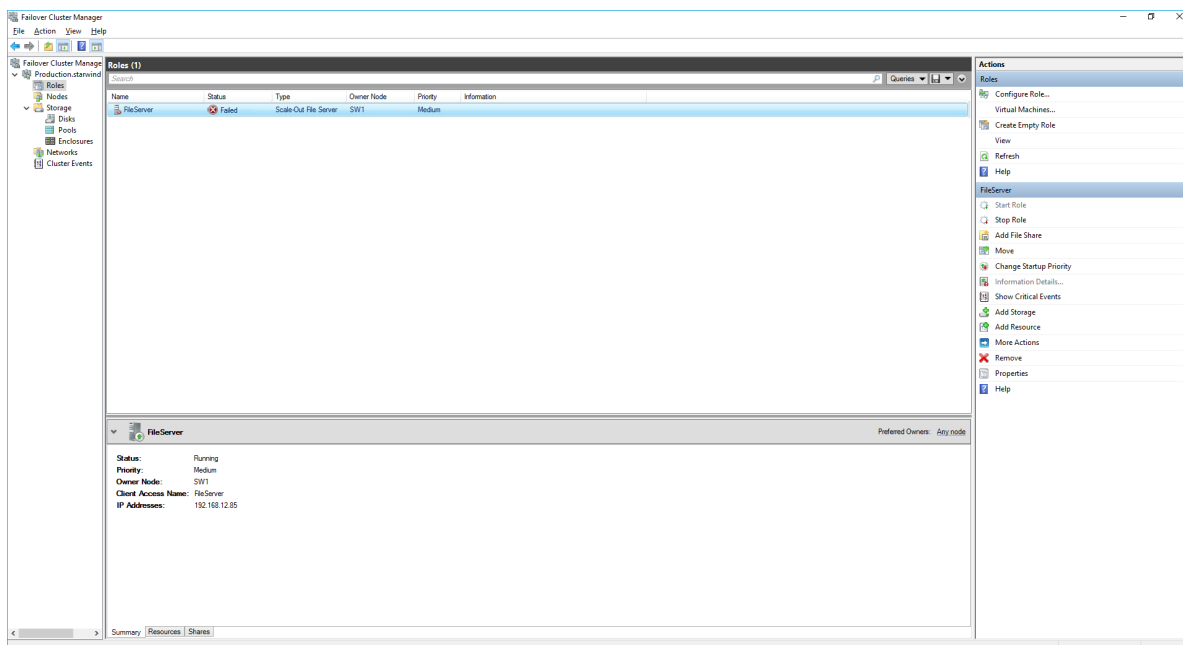
Click Finish to close the Wizard.



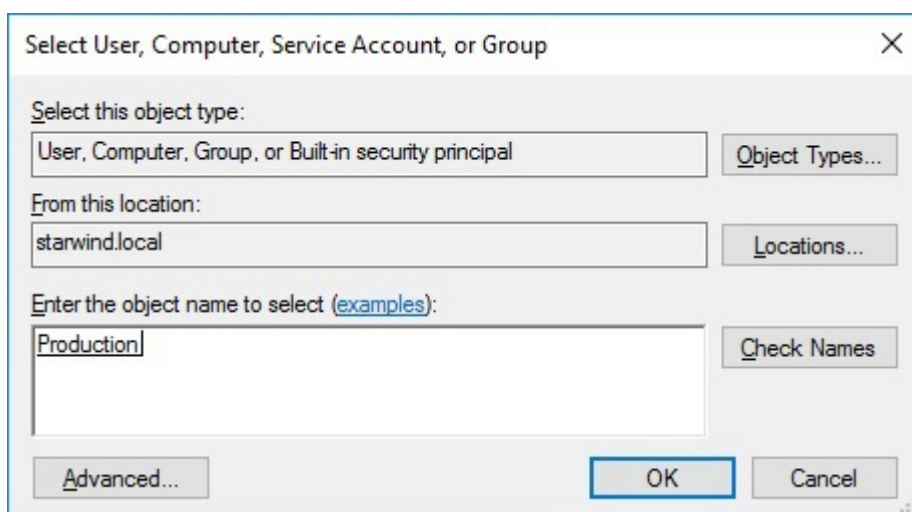
8. The newly created role should now look like the screenshot below.



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,

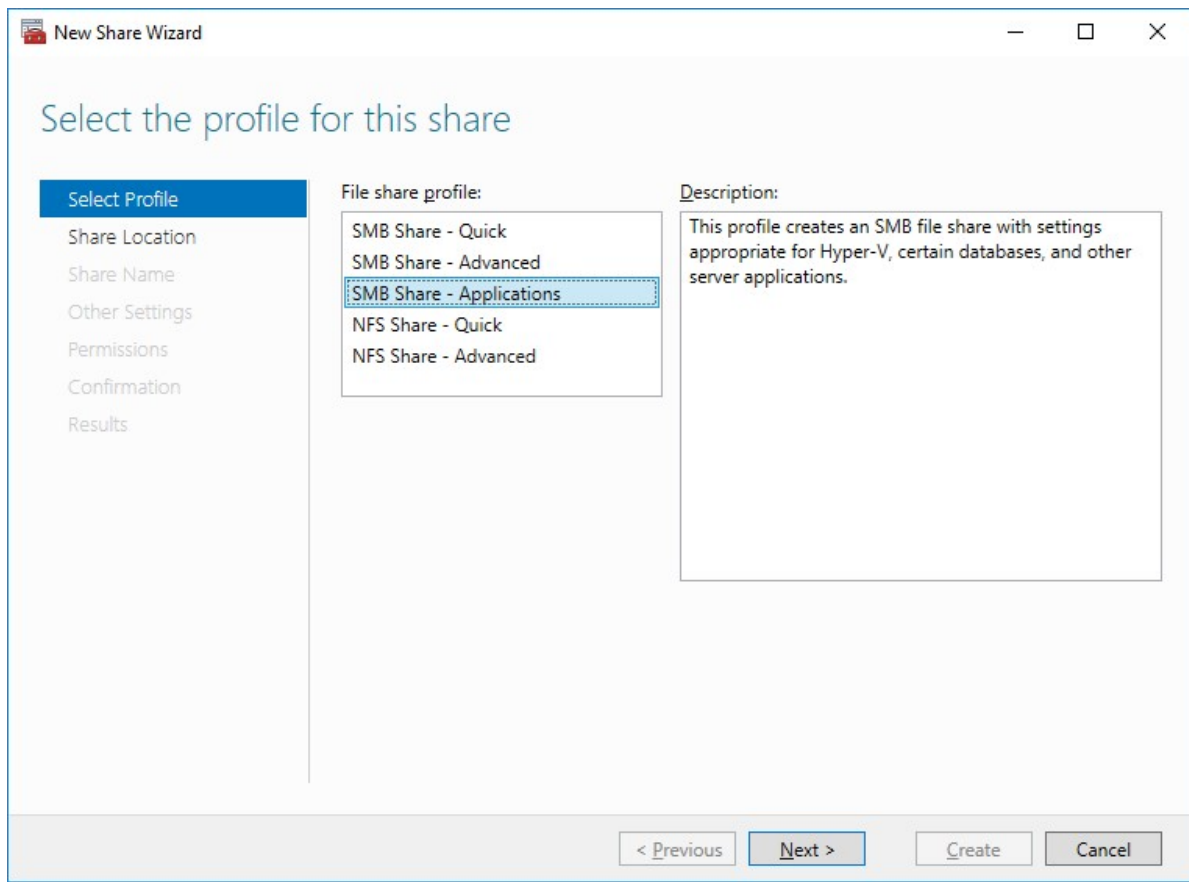


- click OK on all opened windows to confirm the changes
- open Failover Cluster Manager, right-click SOFS role and click Start Role


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



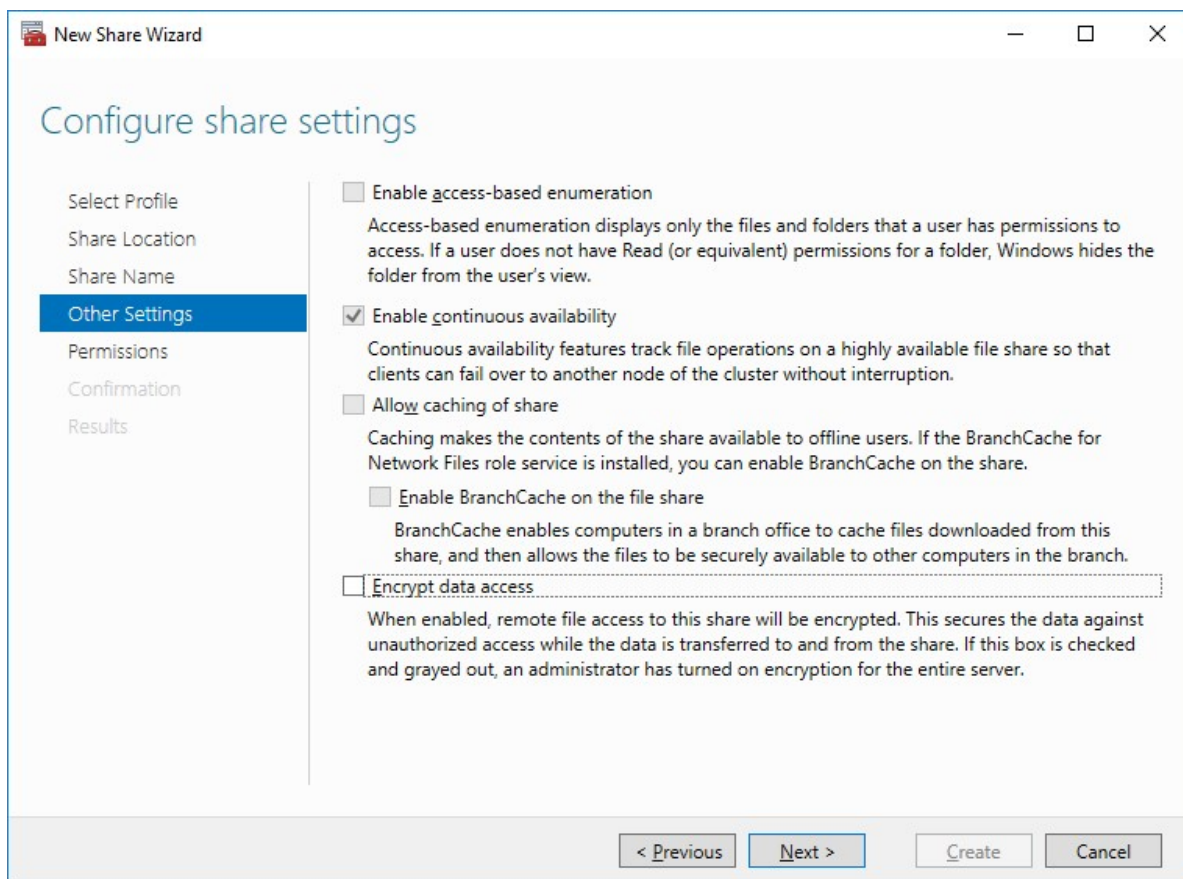
5. Select a CSV to host the share. Click Next to proceed.




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:

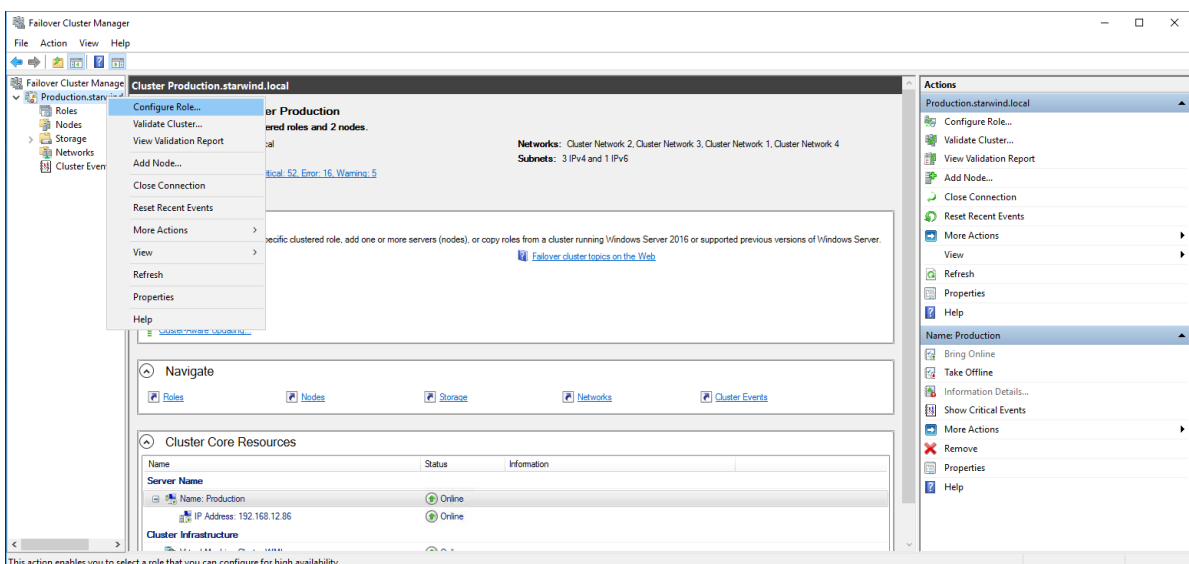
- open Failover Cluster Manager
- expand the cluster and click Roles
- choose the file share role, select the Shares tab, right-click the created file share, and select Properties:



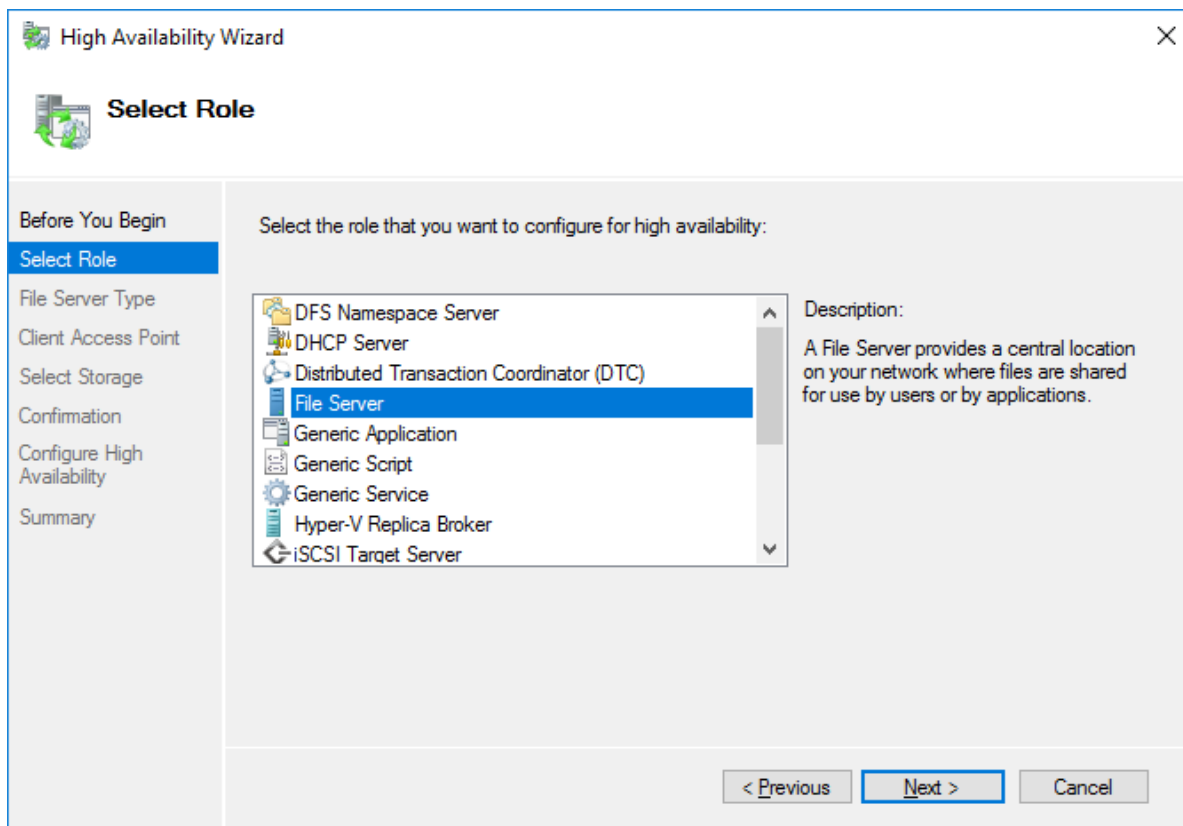
## Configuring The File Server For General Use Role

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

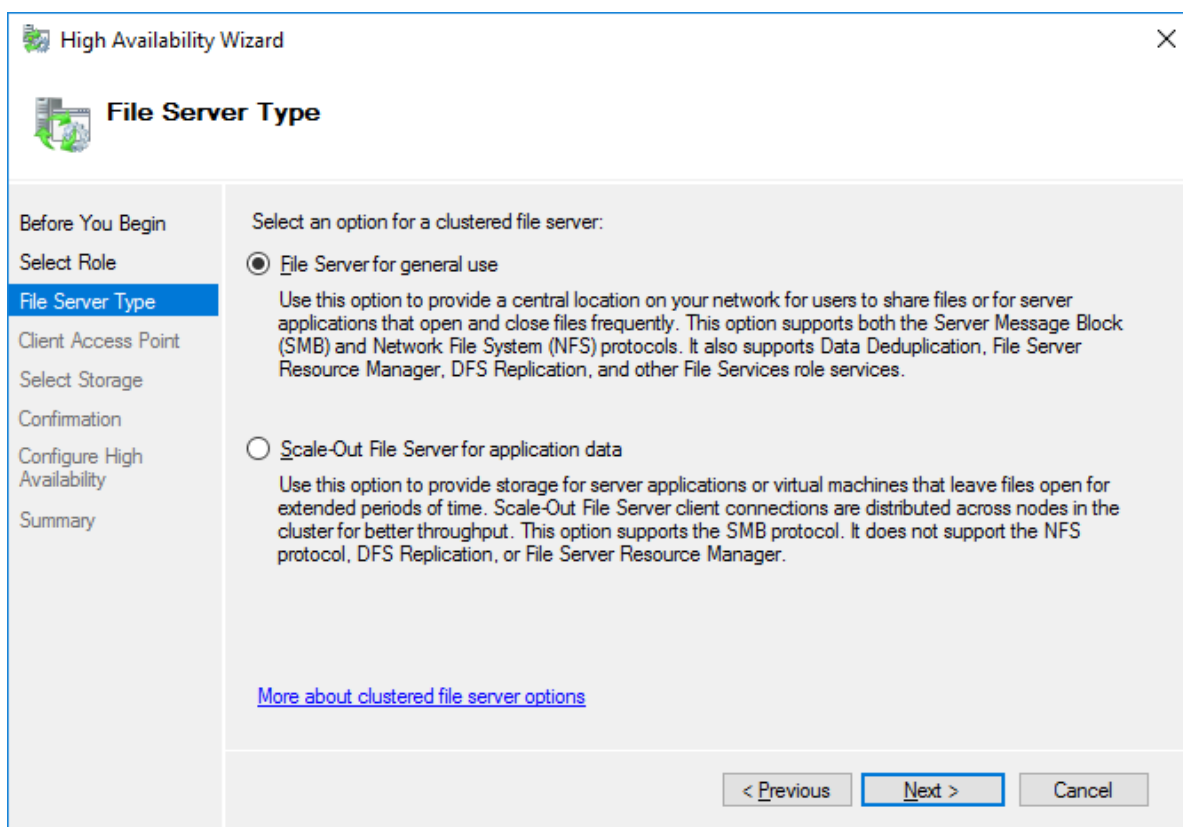
1. To configure the File Server for General Use role, open Failover Cluster Manager.
2. Right-click on the cluster name, then click Configure Role and click Next to continue.



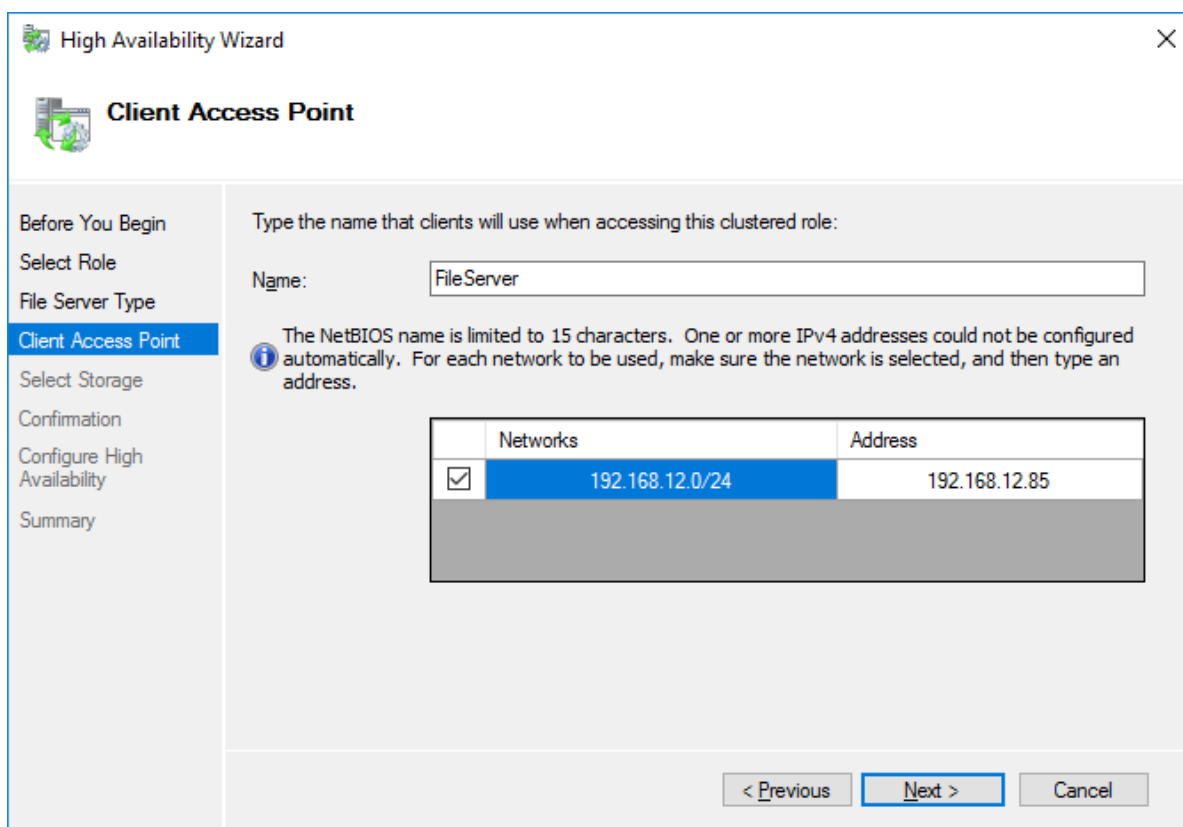
3. Select the File Server item from the list in High Availability Wizard and click Next to continue.



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.

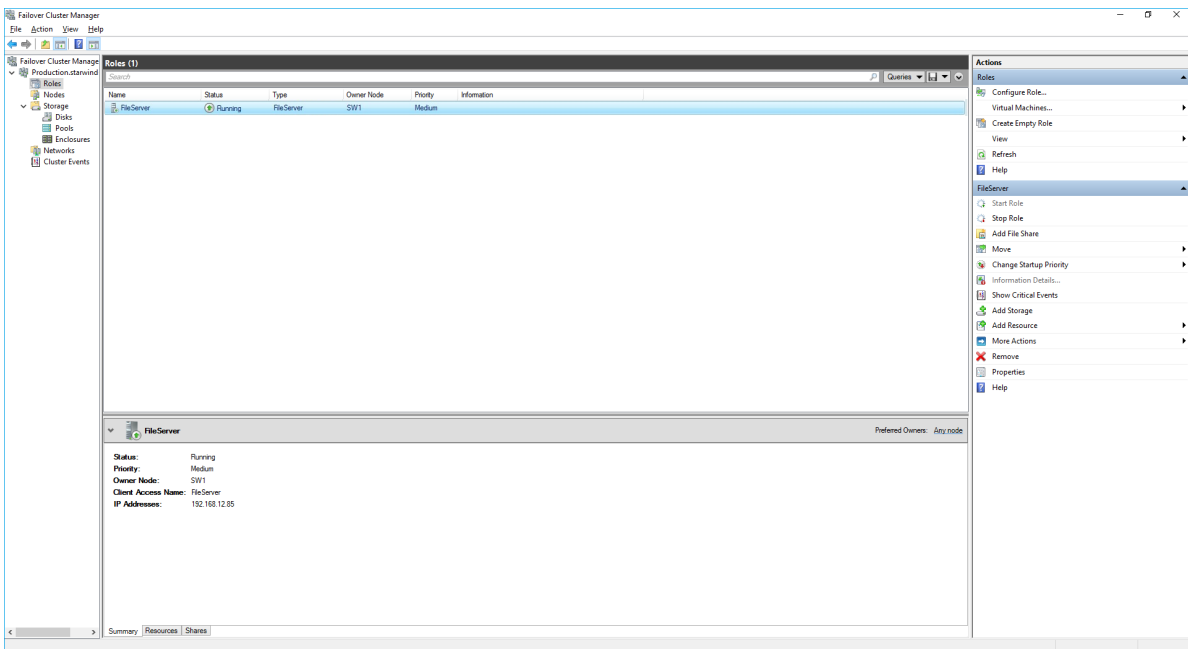


8. Once the installation has been finished successfully, the Wizard should now look like the screenshot below.

Click Finish to close the Wizard.

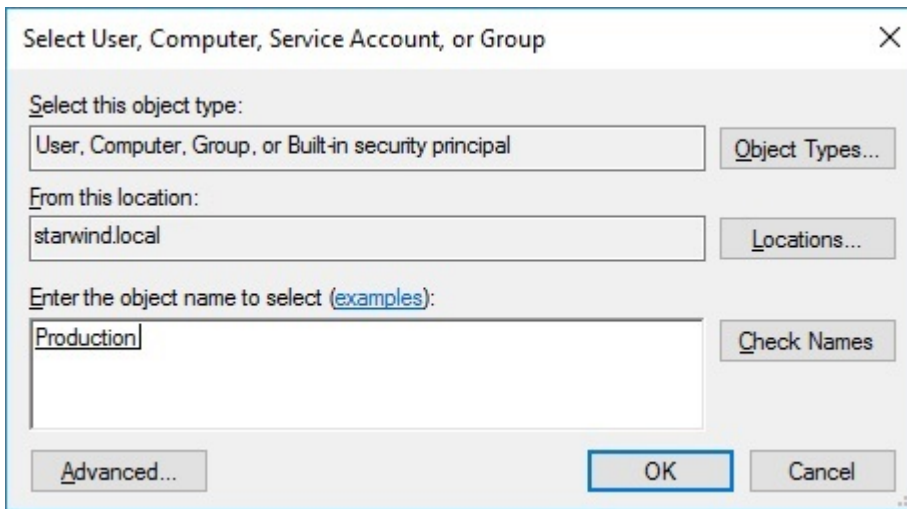


9. The newly created role should now look like the screenshot below.

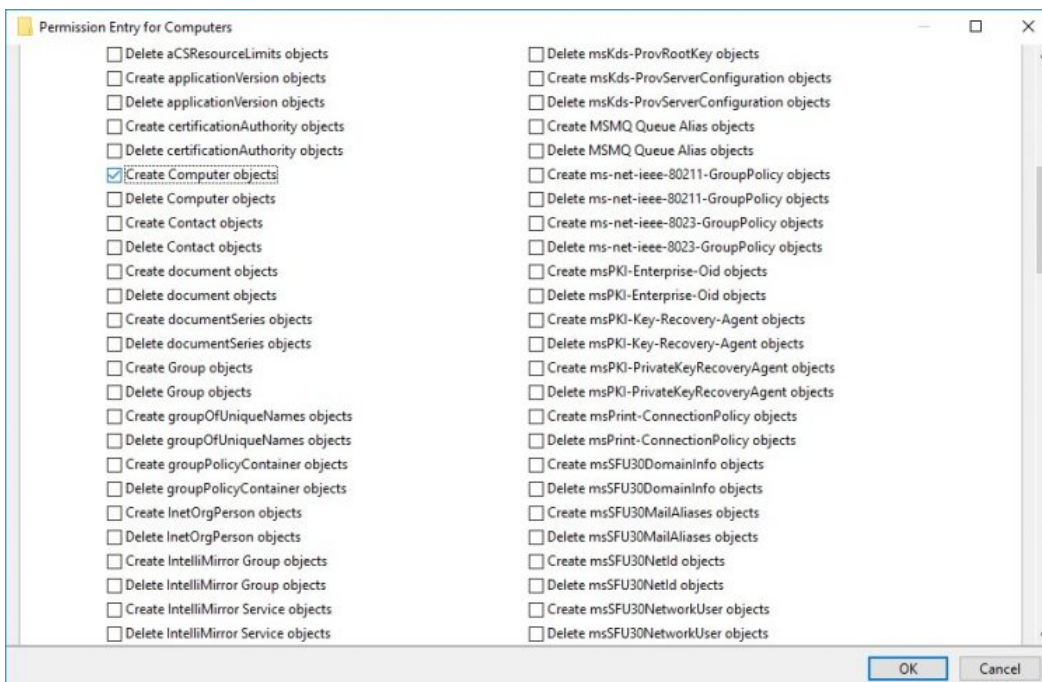


NOTE: If the role status is Failed and it is unable to Start, please, follow the next steps:

- 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

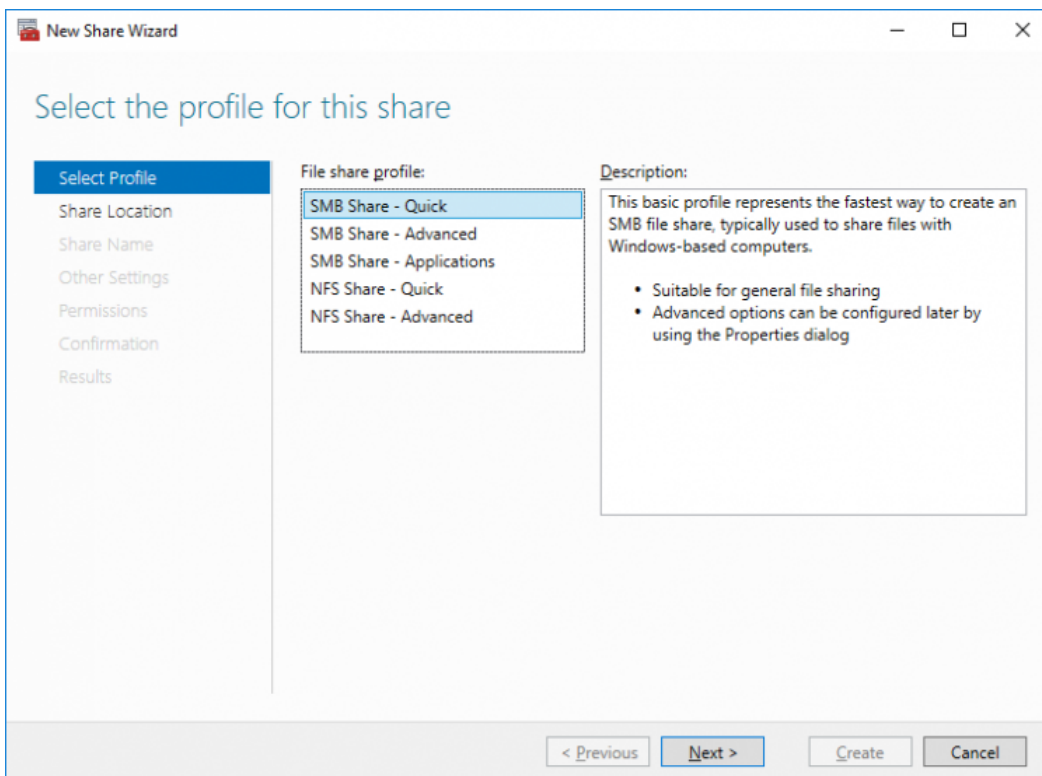


- click OK on all opened windows to confirm the changes
- open Failover Cluster Manager, right-click File Share role and click Start Role

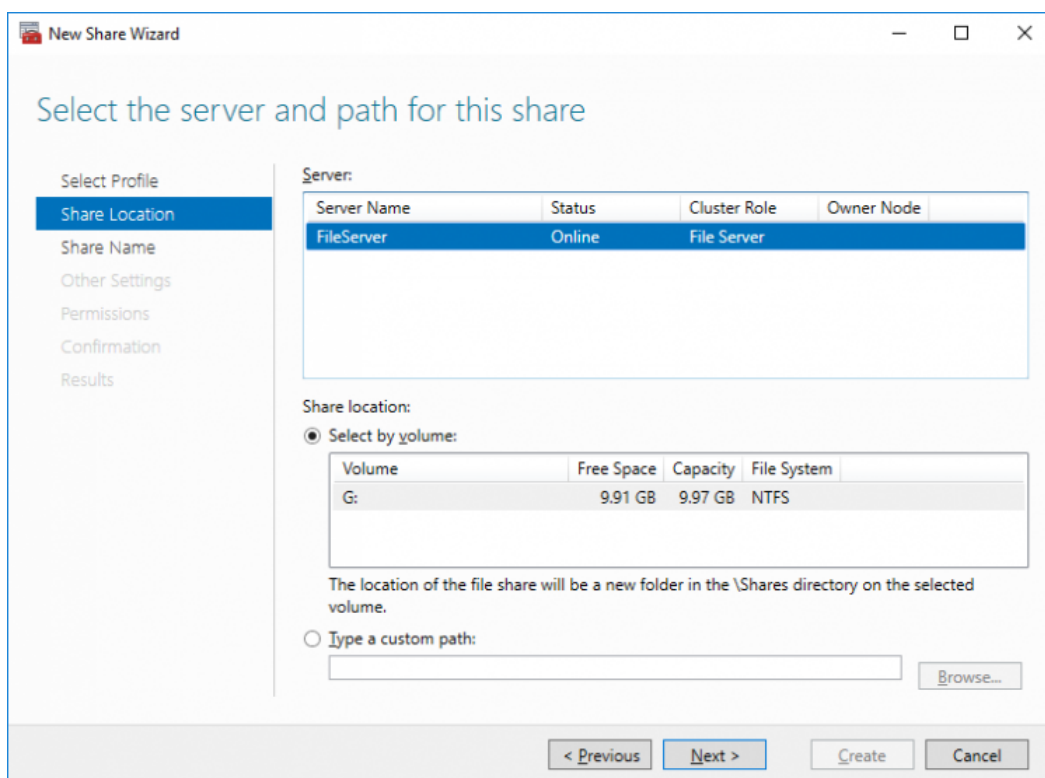
## Configuring Smb File Share

To Add SMB File Share

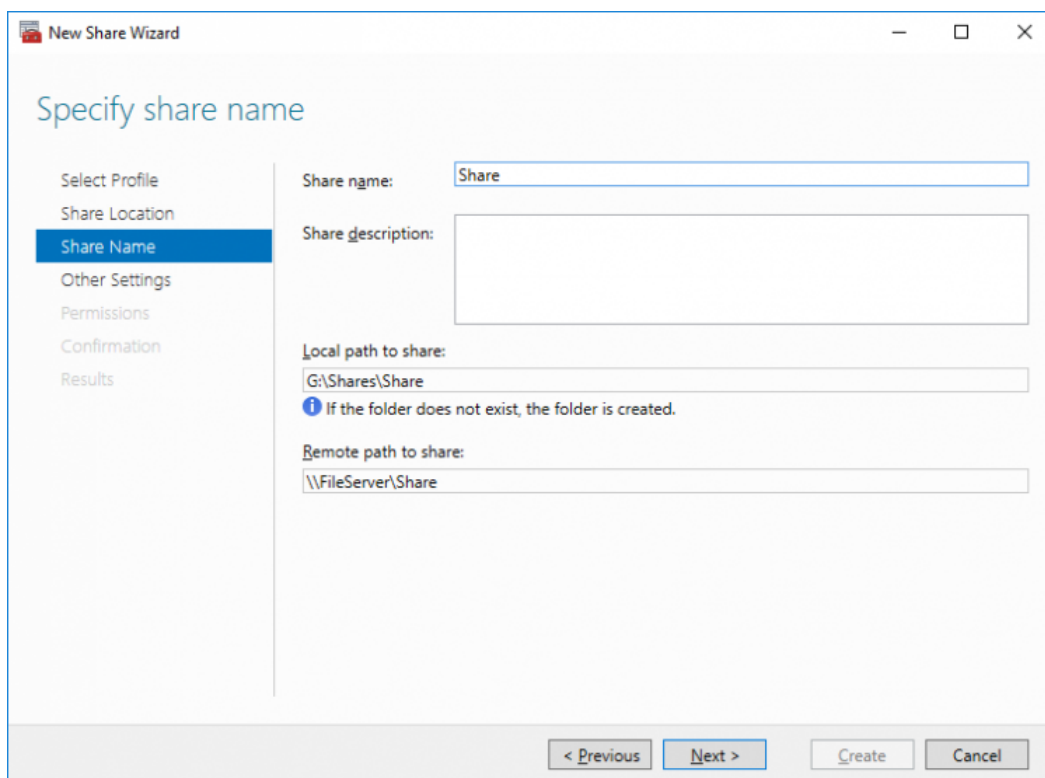
1. Open Failover Cluster Manager.
2. Expand the cluster and then click Roles.
3. Right-click the File Server role and then press Add File Share.
4. On the Select the profile for this share page, click SMB Share – Quick and then click Next.




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




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



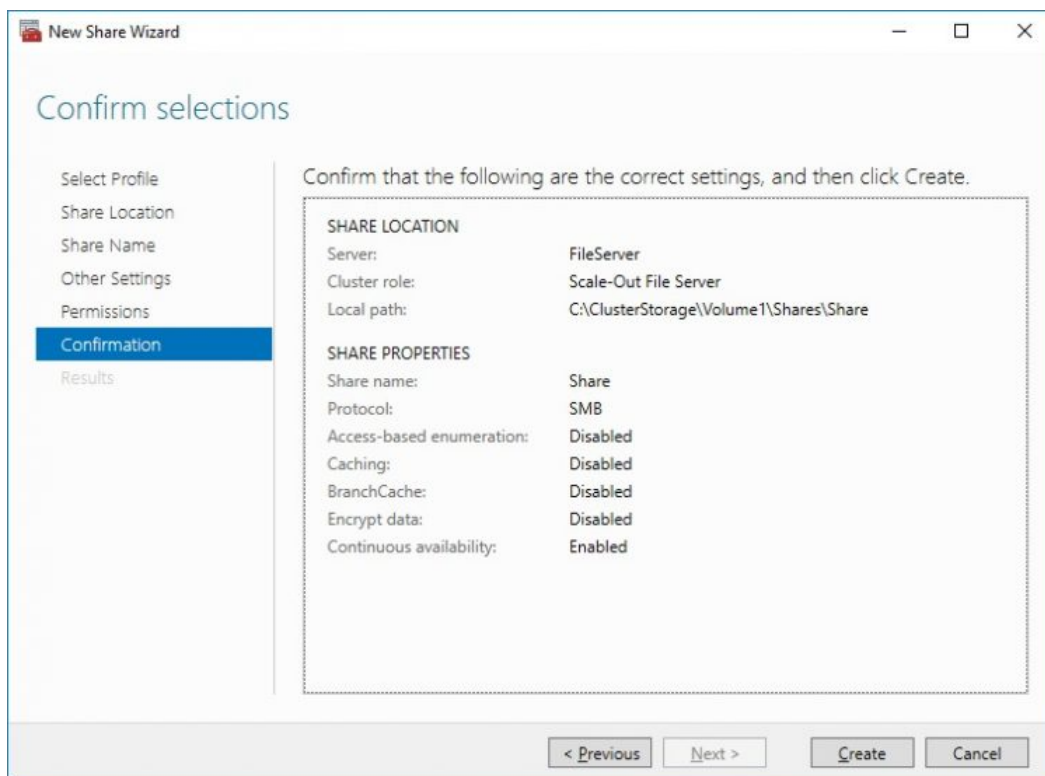
7. Make sure that the Enable Continuous Availability box is checked. Click Next to continue.




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

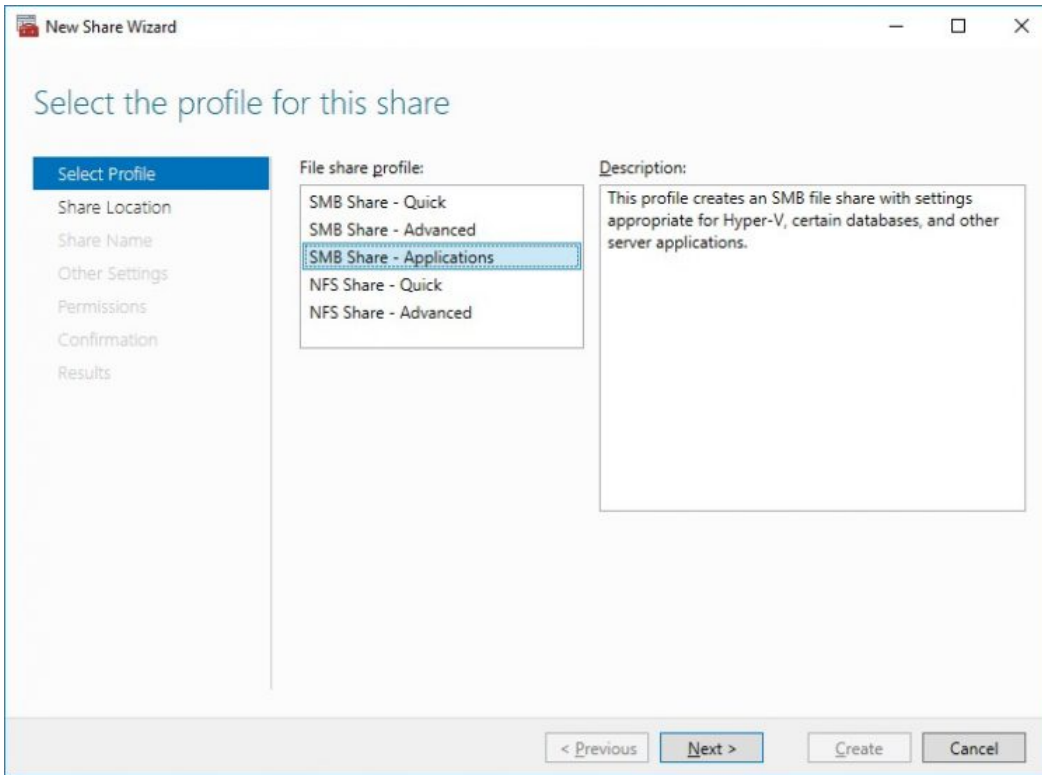
11. Open Failover Cluster Manager.
12. Expand the cluster and click Roles.
13. Choose the File Share role, select the Shares tab, right-click the created file share, and select Properties.



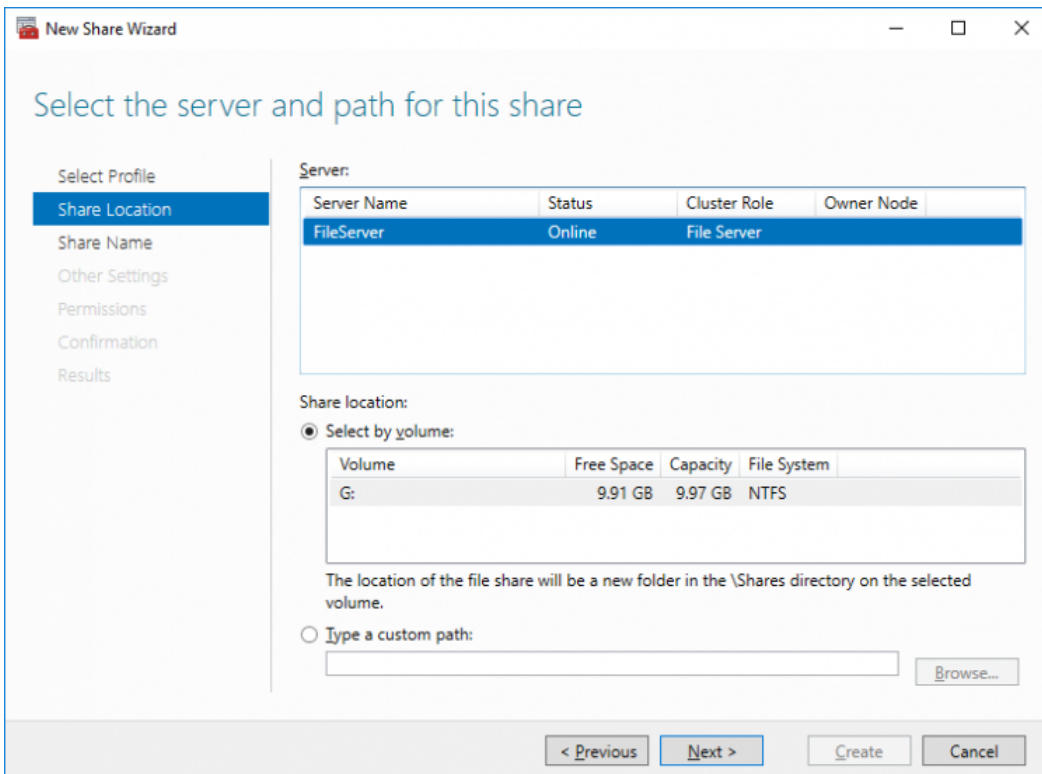
## Configuring Nfs File Share

To Add NFS File Share

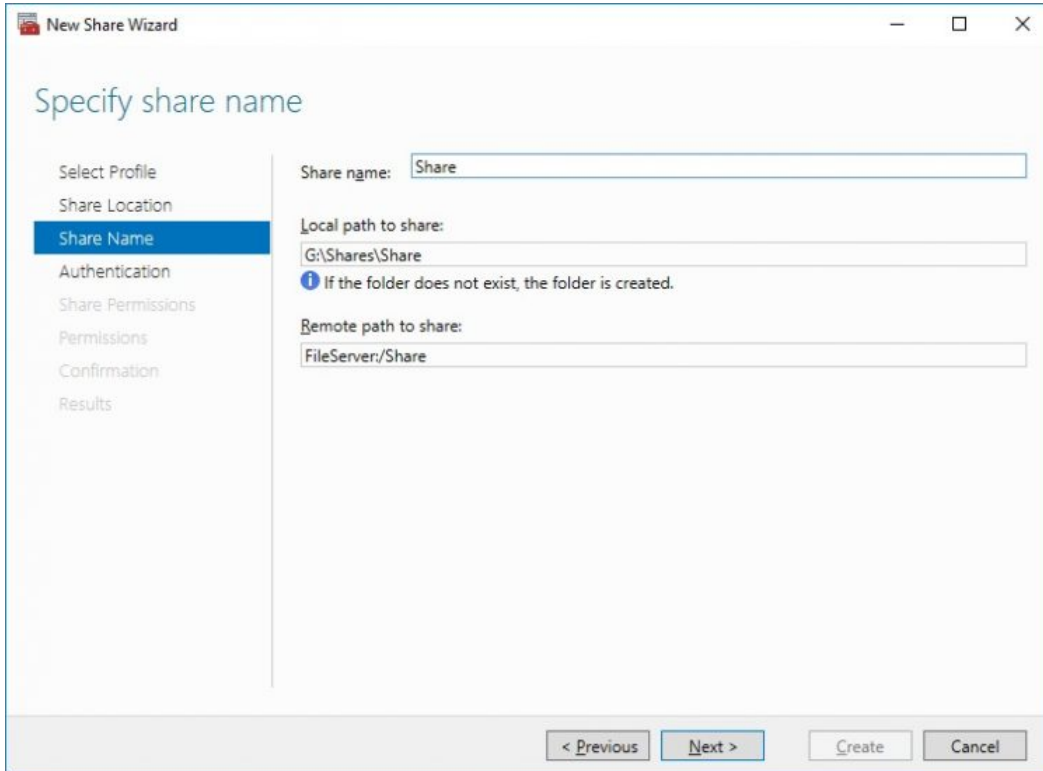
1. Open Failover Cluster Manager.
2. Expand the cluster and then click Roles.
3. Right-click the File Server role and then press Add File Share.
4. On the Select the profile for this share page, click NFS Share – Quick and then click Next.



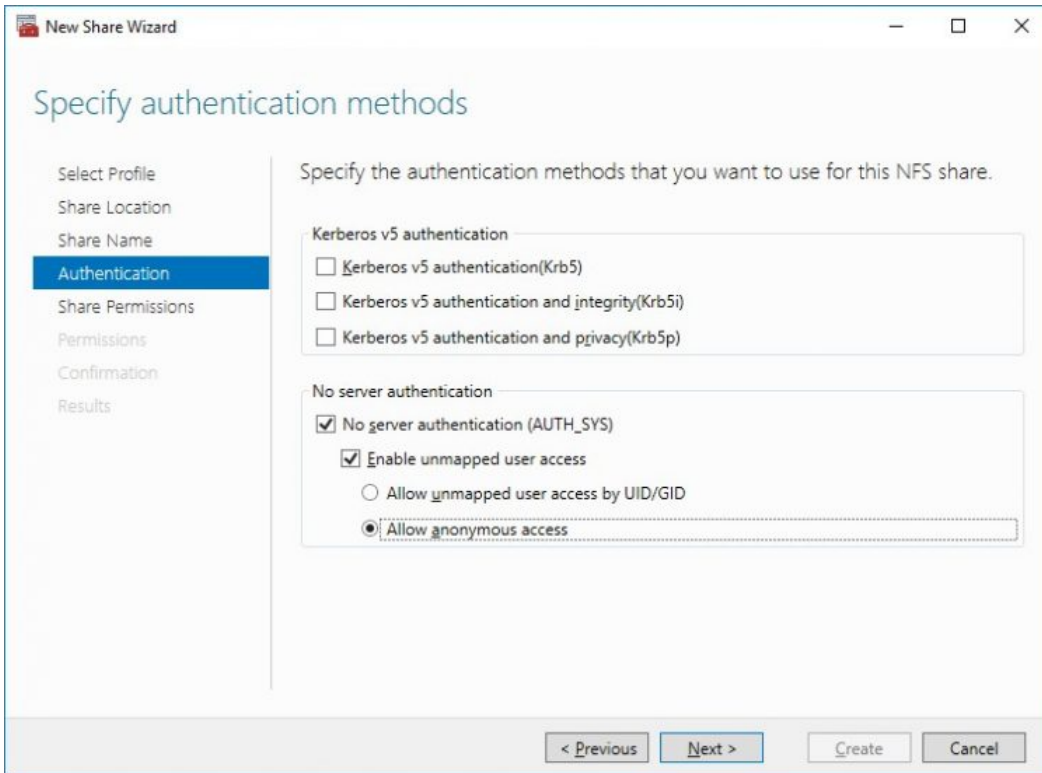
5. Select available storage to host the share. Click Next to continue.



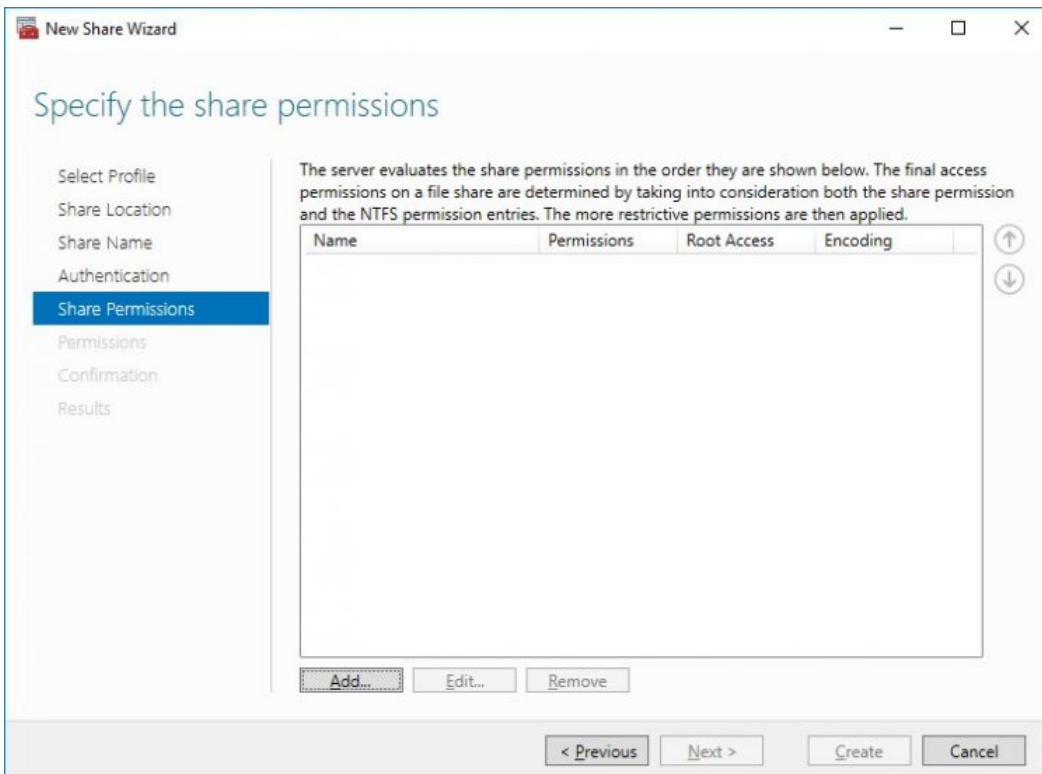
6. Type in the file share name and click Next.




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



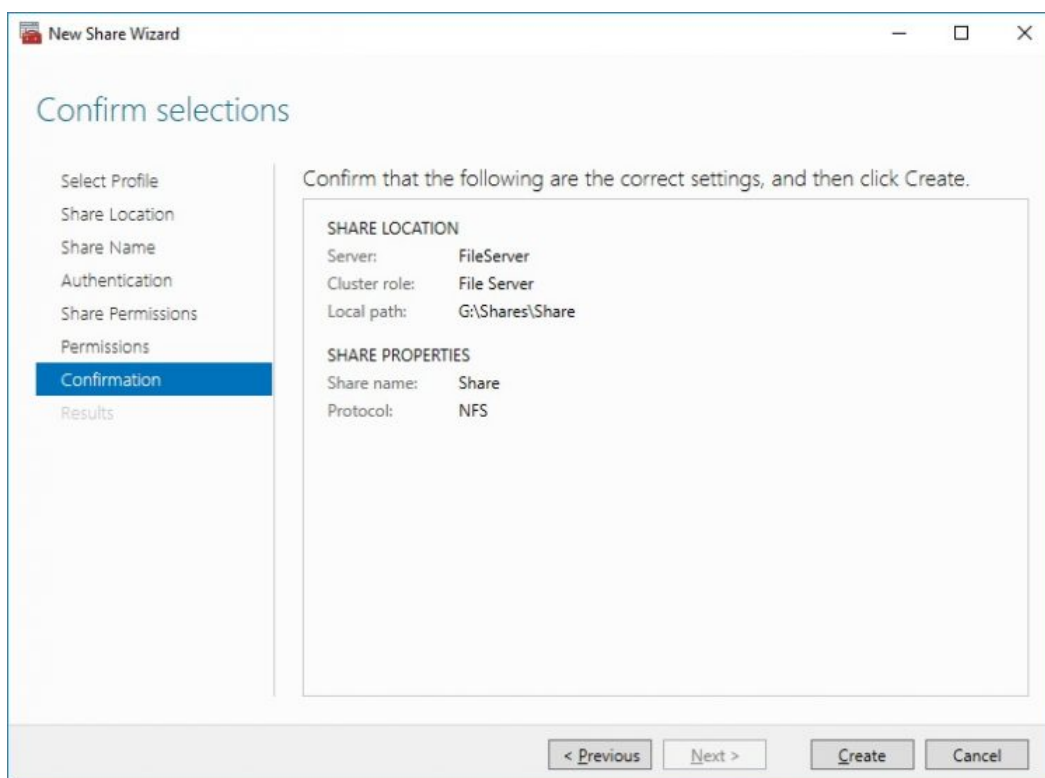
8. Click Add and specify Share Permissions.








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

Deploying and configuring Failover Cluster on Microsoft Azure Stack HCI is a crucial process for organizations looking to achieve a high-performing, VM-centric environment. With this guide, IT professionals will have all the necessary resources for a successful cluster deployment with a shared highly-available storage, provided by StarWind VSAN.

## 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](mailto:sales@starwind.com)

General Information: [info@starwind.com](mailto:info@starwind.com)



StarWind Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA  
[www.starwind.com](http://www.starwind.com) ©2026, StarWind Software Inc. All rights reserved.