

# StarWind Virtual SAN: Configuration Guide for Microsoft Azure Stack HCI [Hyper-V], VSAN Deployed as a Controller Virtual Machine (CVM) using Web UI

2024

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 is applicable to StarWind Virtual SAN and StarWind Virtual SAN Free Controller Virtual Machine (CVM) Version 20231016 and later).

For older versions of StarWind Virtual SAN (OVF Version 20230901 and earlier), refer to this configuration guide:

[StarWind Virtual SAN \(VSAN\): Configuration Guide for Microsoft Windows Server \[Hyper-V\], VSAN Deployed as a Windows Application using Legacy GUI.](#)

## Purpose

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

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.

NOTE: This guide universally applies to the clusters with 2-node, 3-node and more. Follow the quick notes within the configuration steps to carry out the necessary actions required for each cluster size.

## 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:

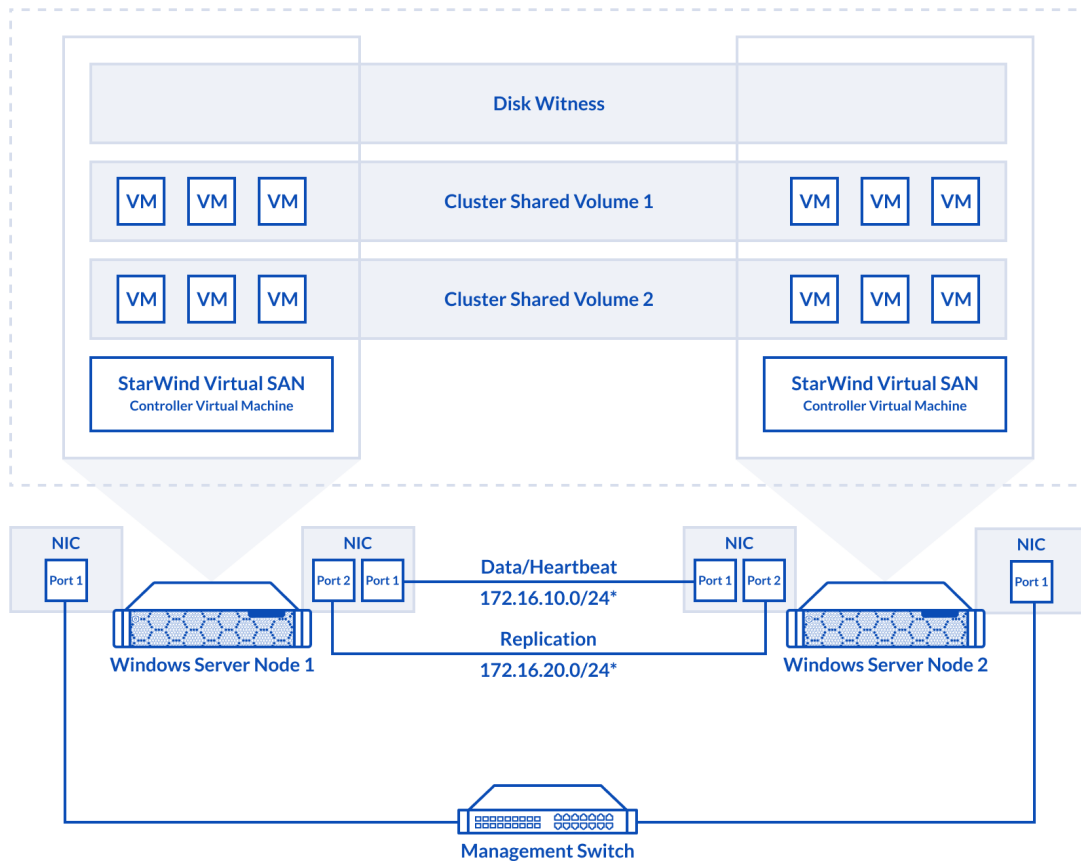
<https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/>

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 :

```
mpiocpl
```

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

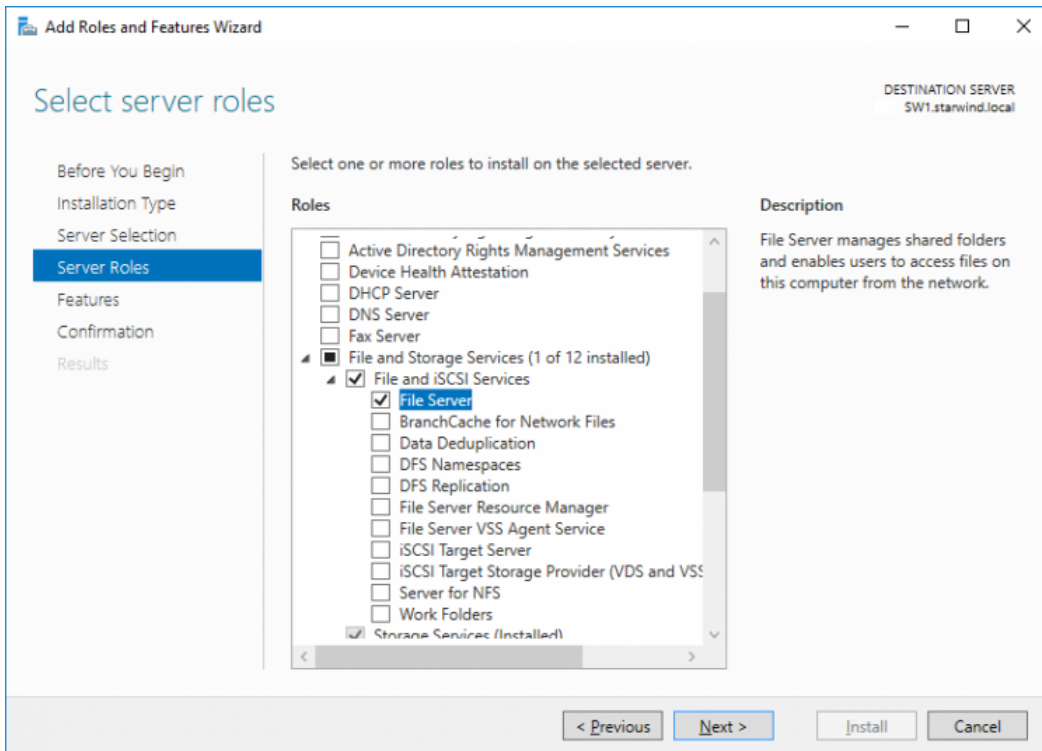
Please follow the steps below if file shares configuration is required

### 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

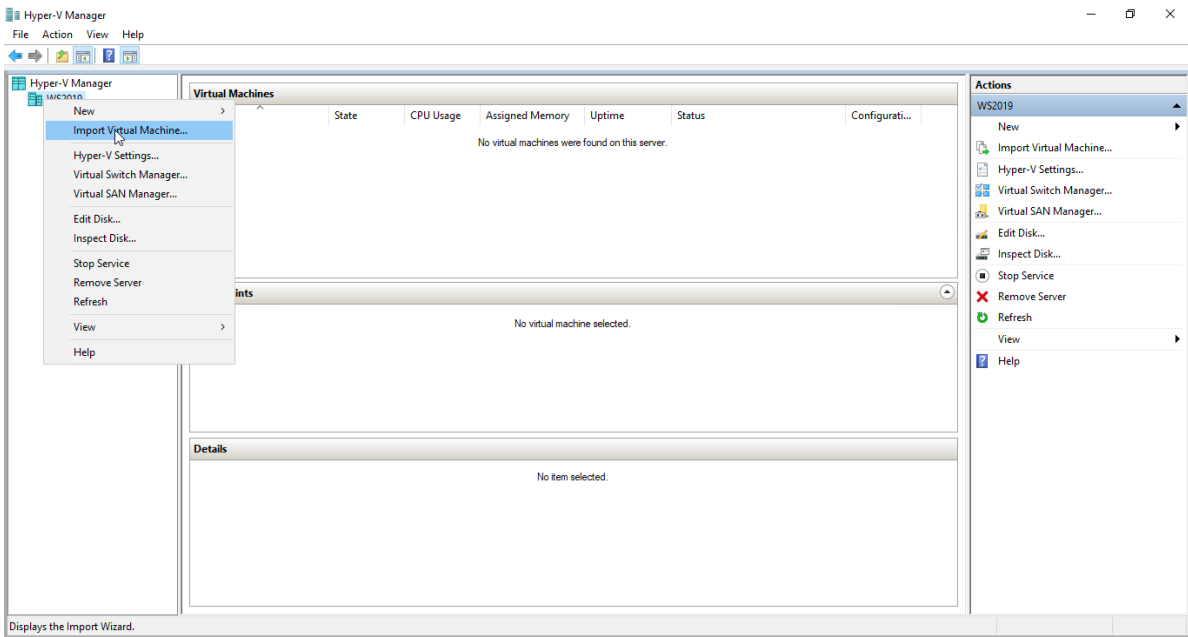
1. Open Server Manager: Start -> Server Manager.
2. Select: Manage -> Add Roles and Features.
3. Follow the installation wizard steps to install the roles selected in the screenshot below:



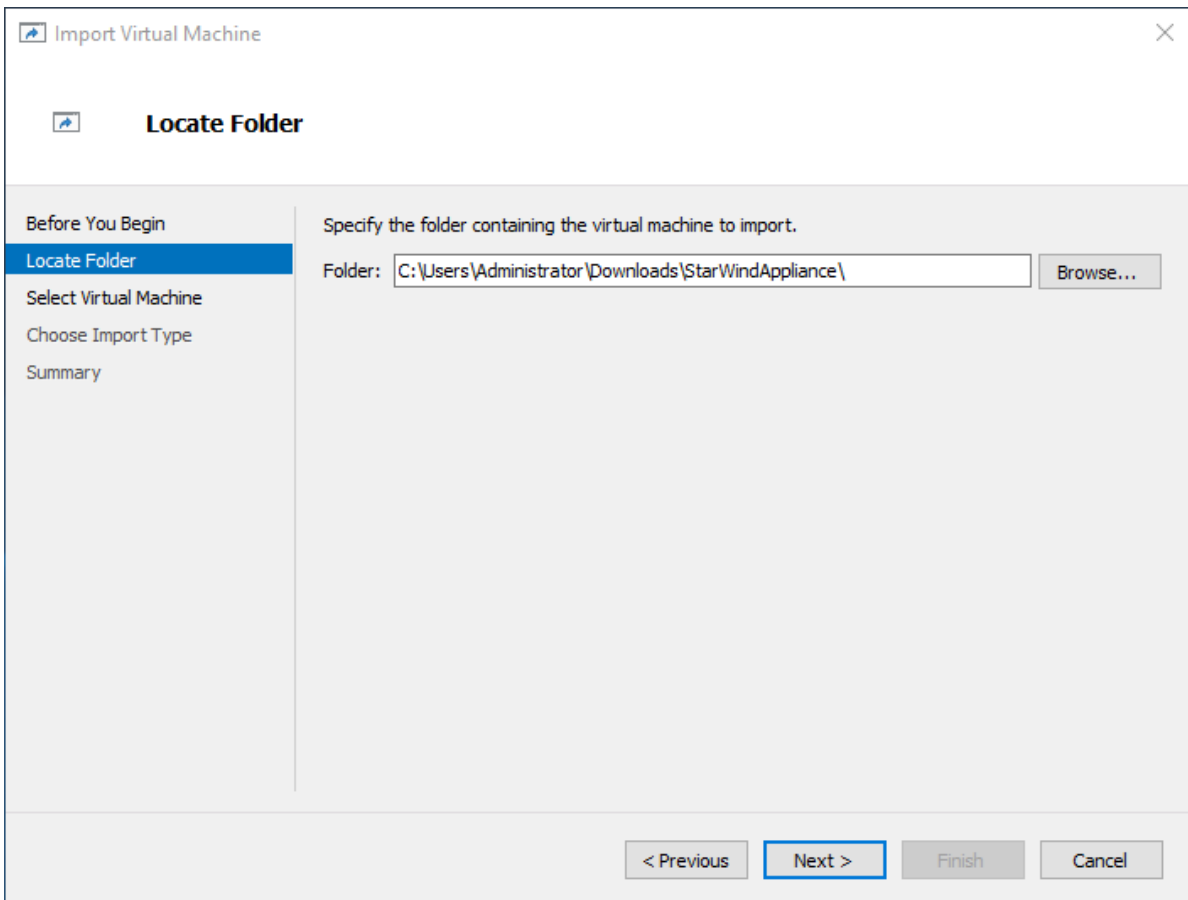
4. Restart the server after installation is completed and perform steps above on each server.

## Deploying Starwind Virtual San Cvm


1. Download the zip archive that contains StarWind Virtual SAN CVM  
<https://www.starwindsoftware.com/vsan#download>
2. Extract the virtual machine files.
3. Deploy the control virtual machine to the Microsoft Hyper-V Server using the “Import Virtual Machine” wizard in Hyper-V Manager.



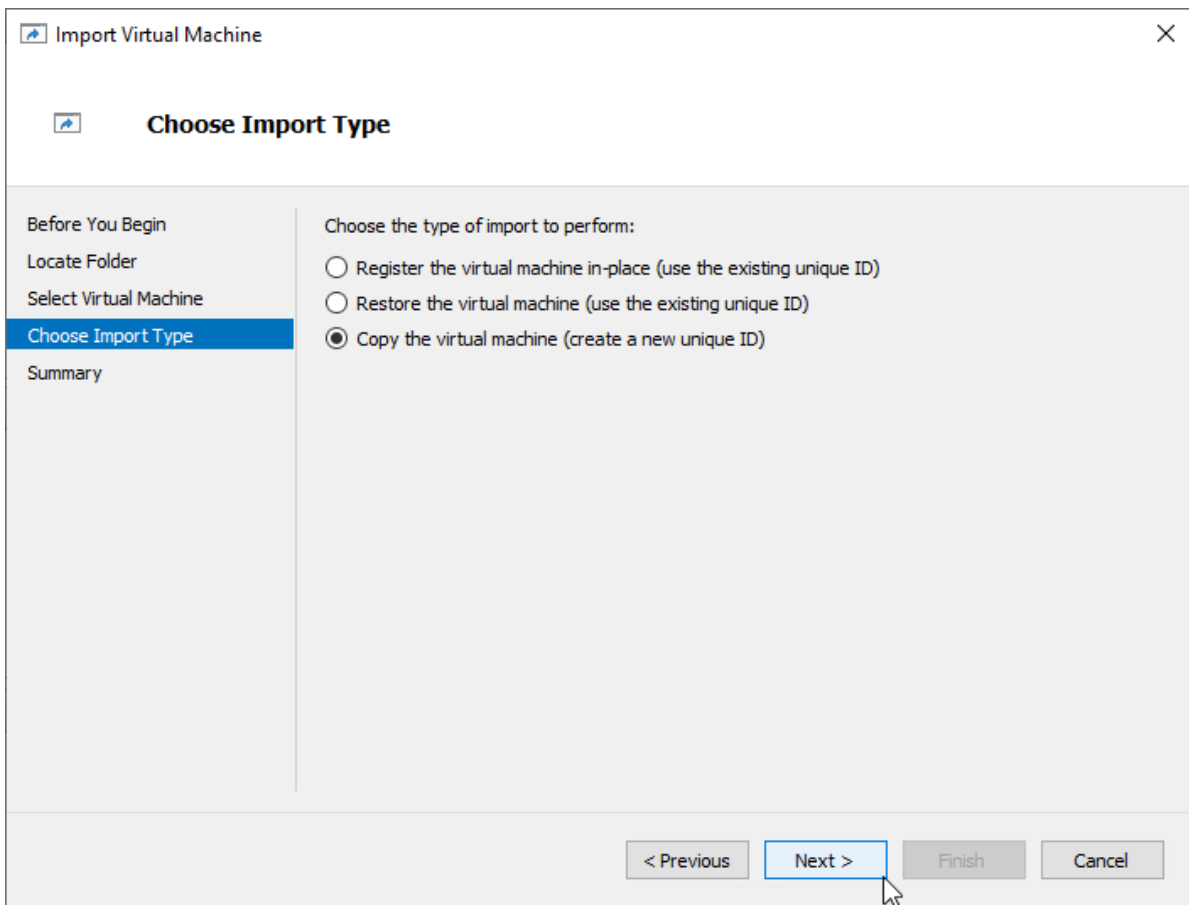
4. On the second page of the wizard, point to the location of the VM template. Select the VM folder and click Next.



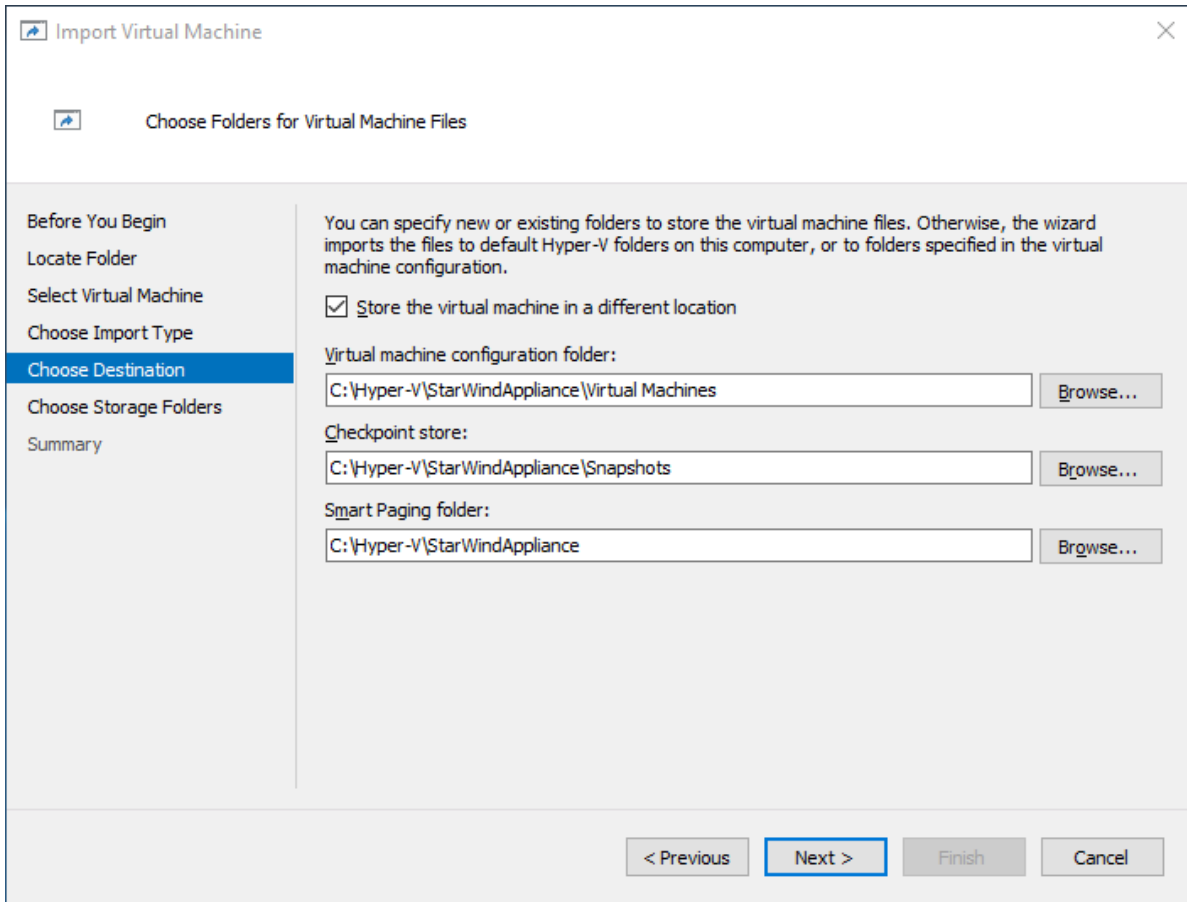
5. Click Next on the “Select Virtual Machine” step.

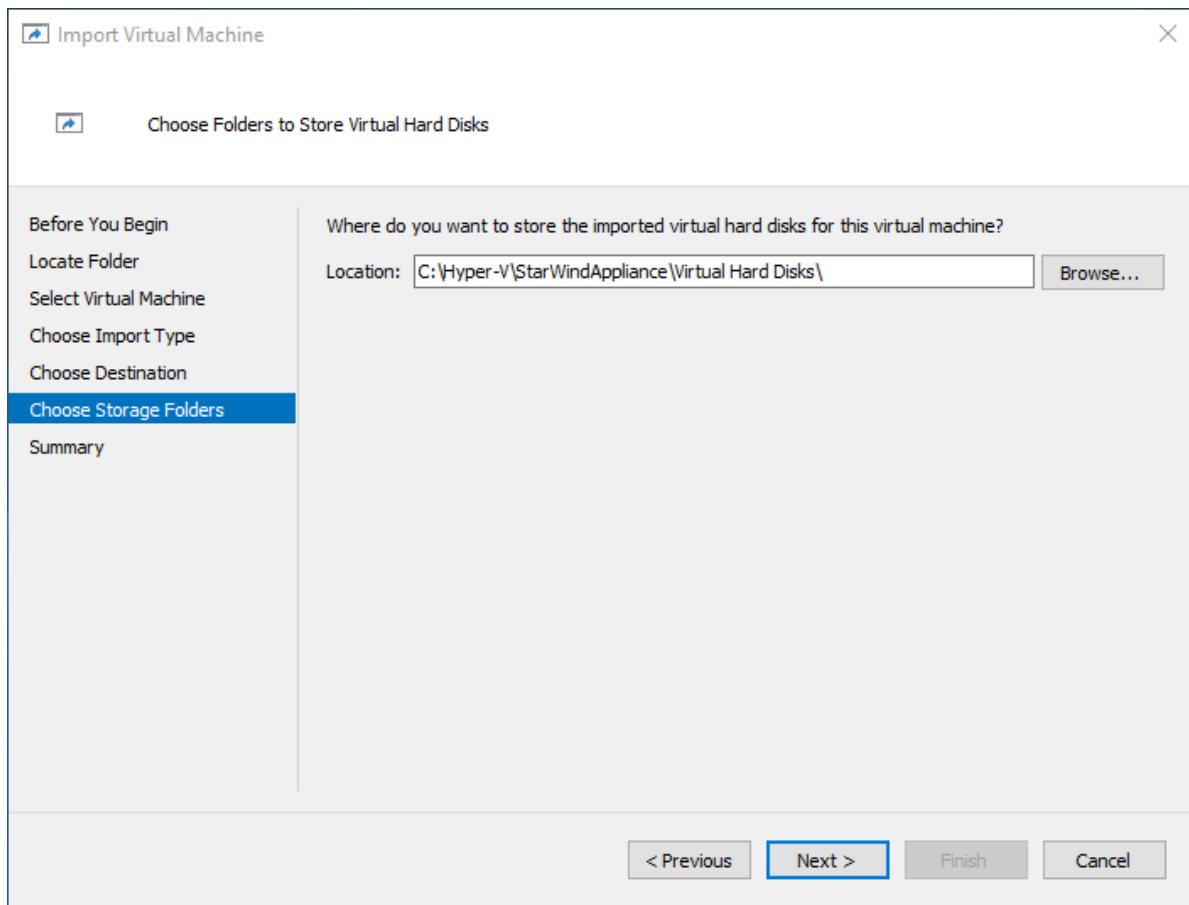


6. Select the “Copy the virtual machine” import type and click Next.




7. Specify new or existing folders to store virtual machine files, such as configuration, snapshots, smart paging, and virtual disk. Click Next.










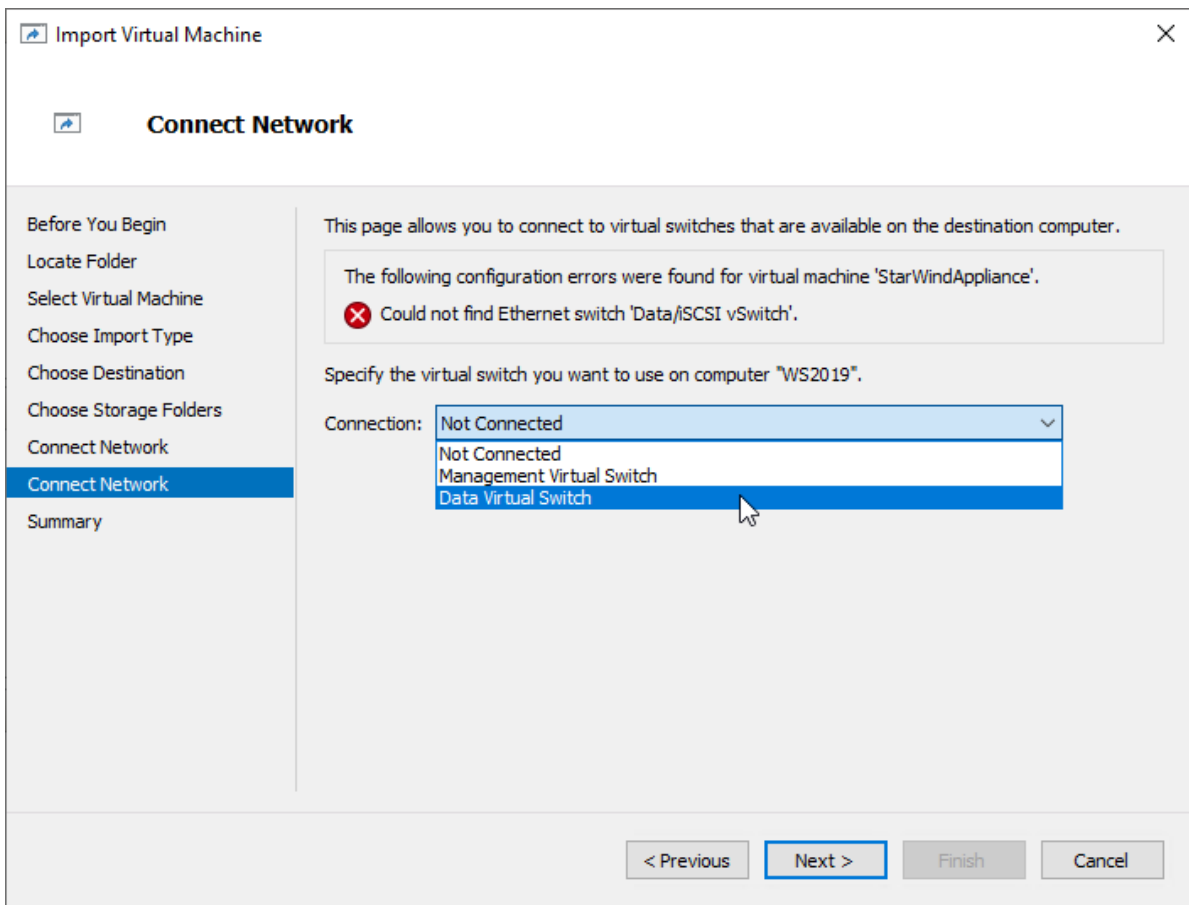
8. In the second step of the wizard, the “VM import” wizard will validate the network.

The default naming for virtual switches:

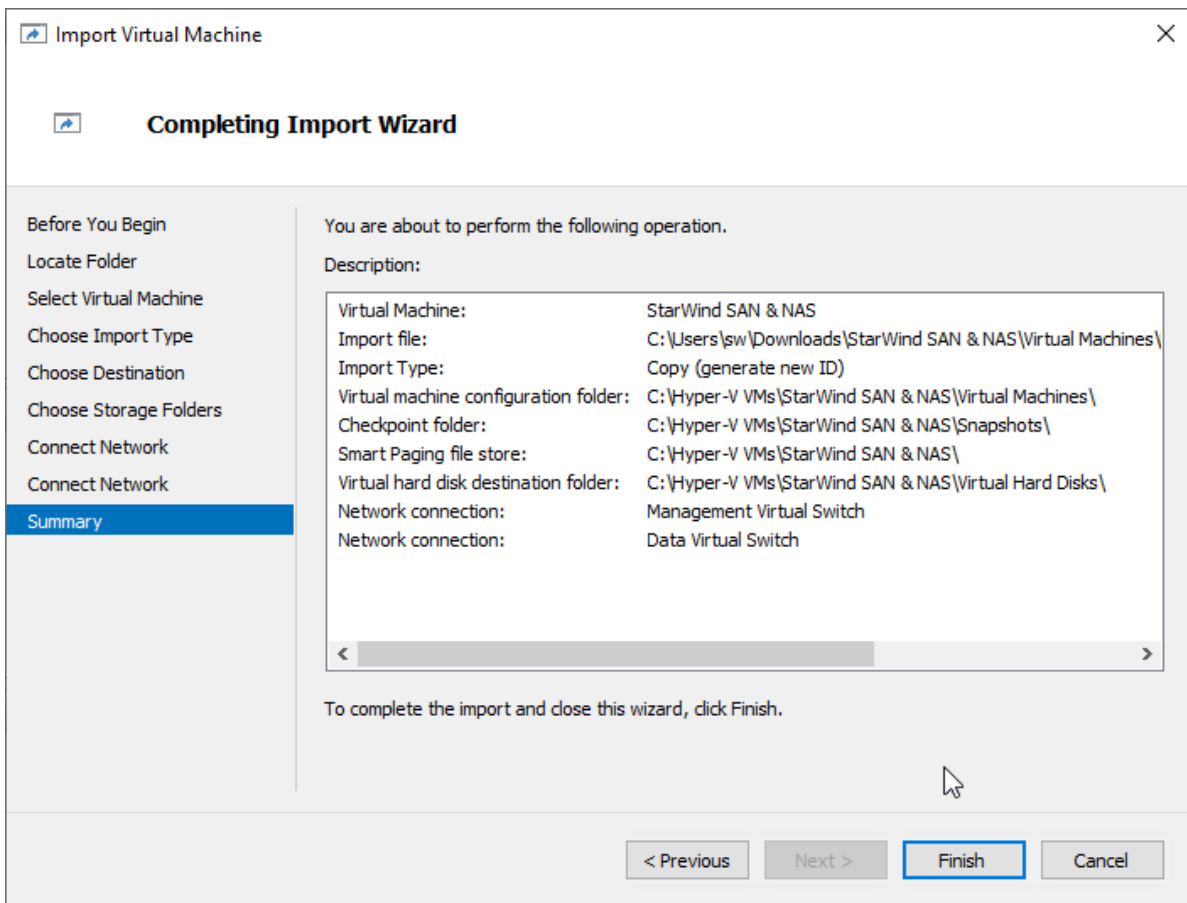
- the Management virtual switch is “Management vSwitch”
- the iSCSI virtual switch is “Data/iSCSI vSwitch”
- the Synchronization virtual switch is “Replication/Sync vSwitch “

If existing virtual switches have different names, specify corresponding network connections. Click Next.





9. Review the import configuration and click Finish to complete the import.



10. Repeat the VM deployment on each partner server which is used for configuring 2-node or 3-node highly available storage according to your licensing.

## Initial Configuration Wizard

1. Start StarWind Virtual SAN CVM.

2. Launch VM console to see the VM boot process and get the IPv4 address of the Management network interface.

NOTE: in case VM has no IPv4 address obtained from a DHCP server, use the Text-based User Interface (TUI) to set up a Management network.

Default credentials for TUI: user/rds123RDS

3. Using the web browser, open a new tab and enter the VM IPv4 address to open StarWind VSAN Web Interface. Click “Advanced” and then “Continue to...”



### Your connection is not private

Attackers might be trying to steal your information from **192.168.12.206** (for example, passwords, messages, or credit cards). [Learn more](#)

NET:ERR\_CERT\_AUTHORITY\_INVALID


Hide advanced

Back to safety


This server could not prove that it is **192.168.12.206**; its security certificate is not trusted by your computer's operating system. This may be caused by a misconfiguration or an attacker intercepting your connection.

[Proceed to 192.168.12.206 \(unsafe\)](#)

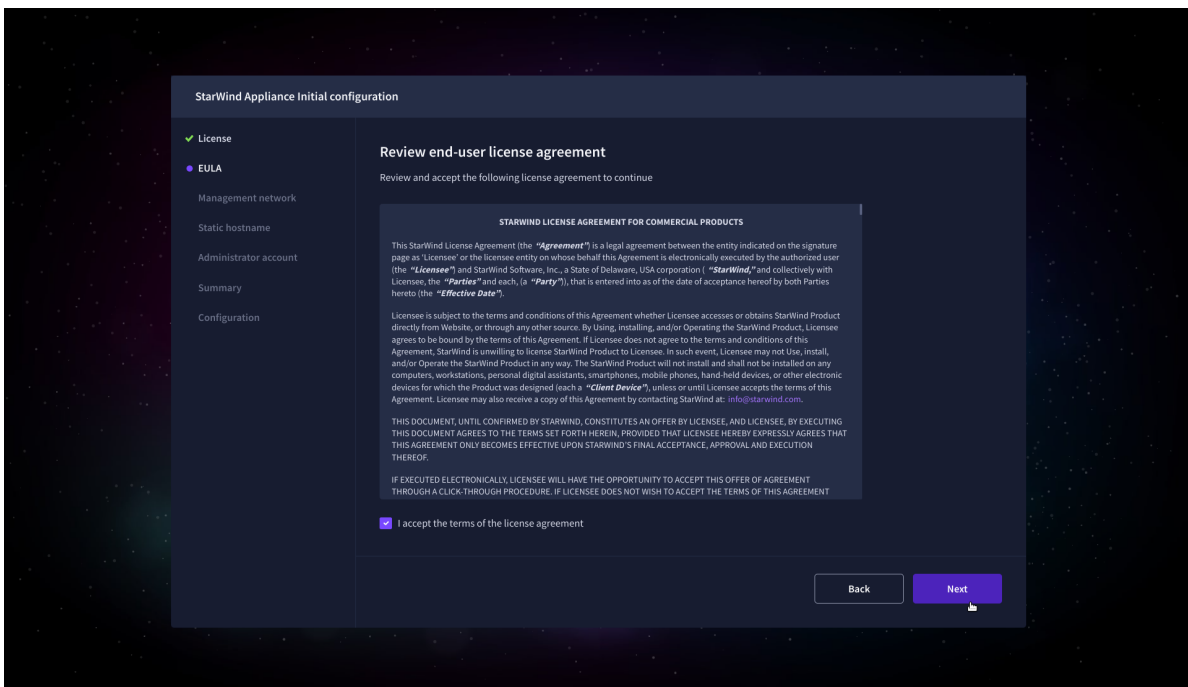
4. StarWind VSAN web UI welcomes you, and the “Initial Configuration” wizard will guide you through the deployment process.



5. In the following step, upload the license file.




6. Read and accept the End User License Agreement to proceed.




7. Review or edit the Network settings and click Next.


NOTE: Static network settings are recommended for the configuration.



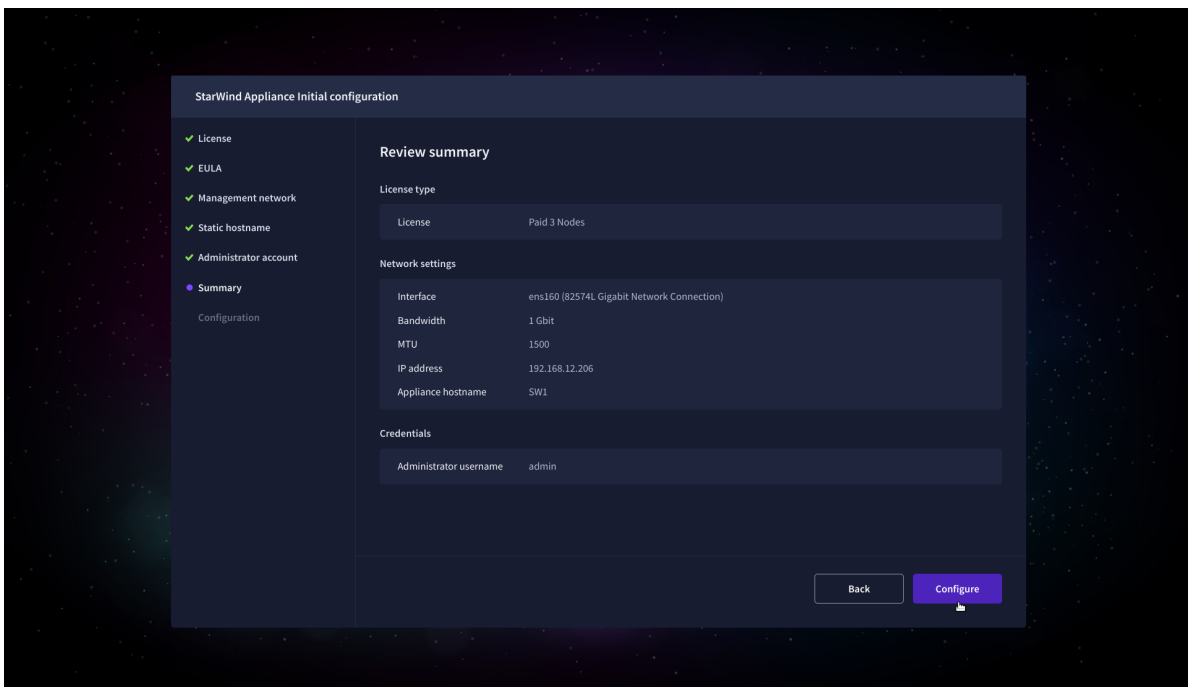
8. Specify the hostname for the virtual machine and click Next.



9. Create an administrator account. Click Next.




10. Review your settings selection before setting up StarWind VSAN.

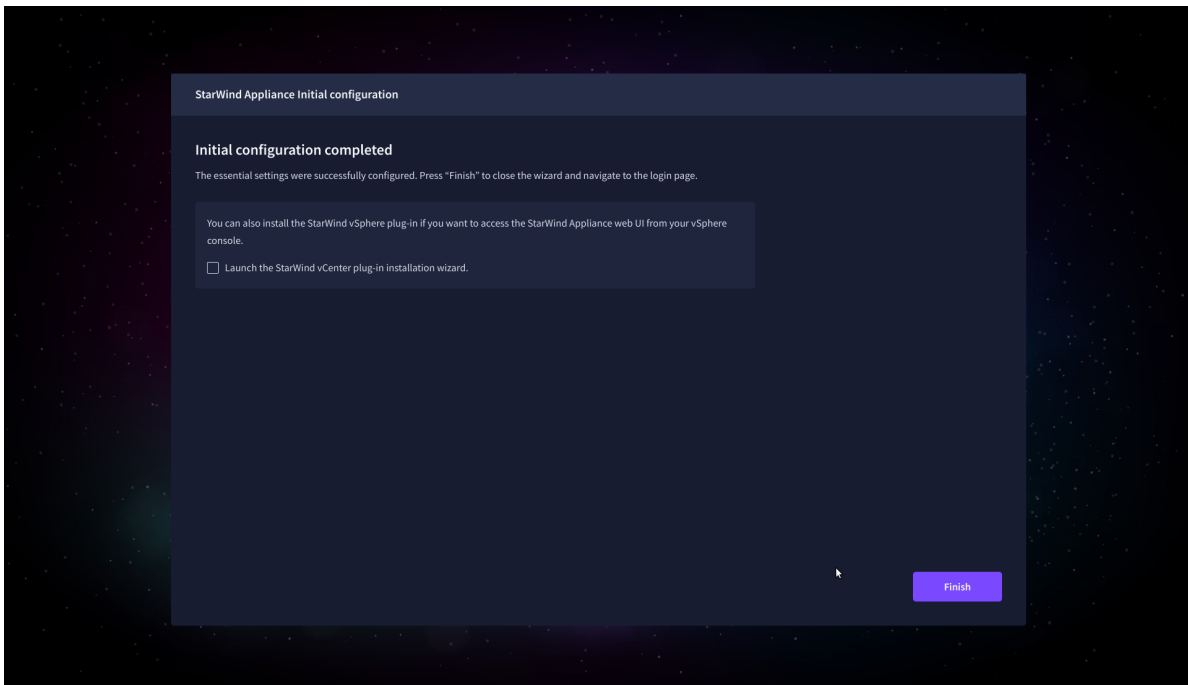


11. Please standby until the Initial Configuration Wizard configures StarWind VSAN for you.





12. The appliance is set and ready. Click on the Done button to install the StarWind vCenter Plugin right now or uncheck the checkbox to skip this step and proceed to the [Login page](#).




13. Repeat the initial configuration on other StarWind CVMs that will be used to create 2-node or 3-node HA shared storage.

## Add Appliance


To create 2-way or 3-way synchronously replicated highly available storage, add partner appliances that use the same license key.

1. Add StarWind appliance(s) in the web console, on the Appliances page.


NOTE: The newly added appliance will be linked to already connected partners.




2. Provide credentials of partner appliance.



3. Wait for connection and validation of settings.




4. Review the summary and click "Add appliance".




## Configure Ha Networking

1. Launch the “Configure HA Networking” wizard.



2. Select appliances for network configuration.


NOTE: the number of appliances to select is limited by your license, so can be either two or three appliances at a time.




3. Configure the “Data” network. Select interfaces to carry storage traffic, configure them with static IP addresses in unique networks, and specify subnet masks:

- assign and configure at least one interface on each node
- for redundant configuration, select two interfaces on each node
- ensure interfaces are connected to client hosts directly or through redundant switches

4. Assign MTU value to all selected network adapters, e.g. 1500 or 9000. Ensure the switches have the same MTU value set.



5. Click Next to validate Data network settings.




6. Configure the “Replication” network. Select interfaces to carry storage traffic, configure them with static IP addresses in unique networks, and specify subnet masks:


- assign and configure at least one interface on each node
- for redundant configuration, select two interfaces on each node


- ensure interfaces are connected to client hosts directly or through redundant switches

7. Assign MTU value to all selected network adapters, e.g. 1500 or 9000. Ensure the switches have the same MTU value set.




8. Click Next to validate the Replication network settings completion.





9. Review the summary and click Configure.





## Add Physical Disks

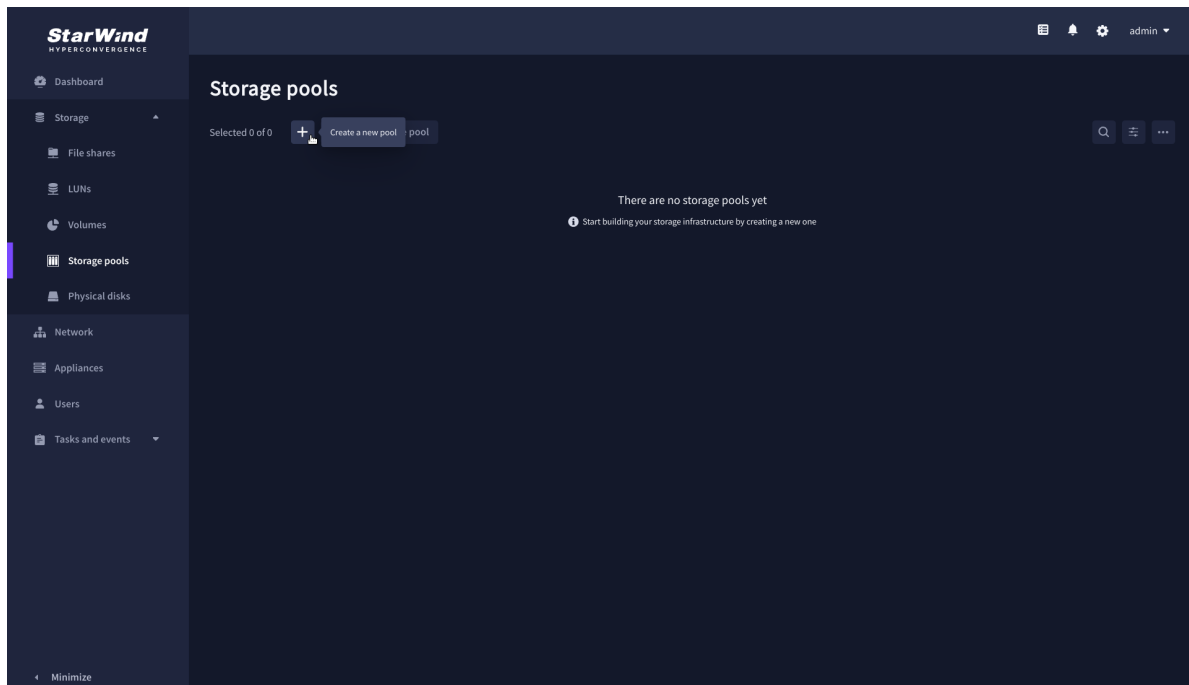
Attach physical storage to StarWind Virtual SAN Controller VM:


- Ensure that all physical drives are connected through an HBA or RAID controller.
- Deploy StarWind VSAN CVM on each server that will be used to configure fault-tolerant standalone or highly available storage.
- Store StarWind VSAN CVM on a separate storage device accessible to the hypervisor host (e.g., SSD, HDD).
- Add HBA, RAID controllers, or NVMe SSD drives to StarWind CVM via a passthrough device.

Learn more about storage provisioning guidelines in the [KB article](#).


## Create Storage Pool

1. Click the “Add” button to create a storage pool.
2. Select two storage nodes to create a storage pool on them simultaneously.

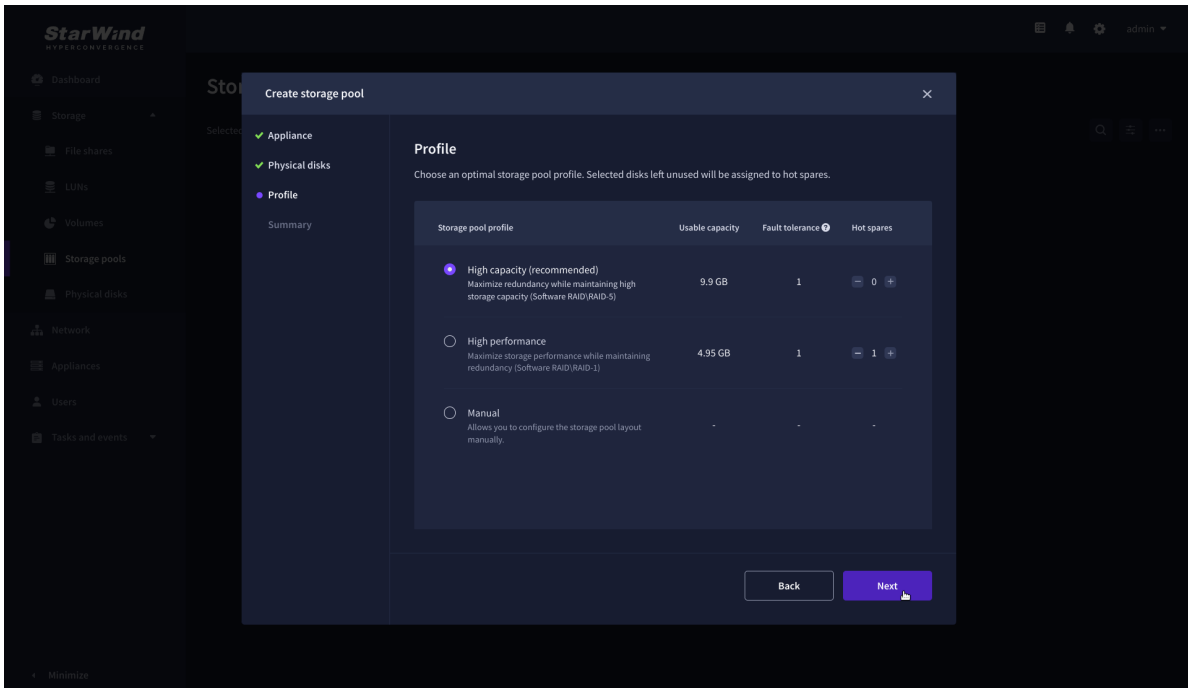




3. Select physical disks to include in the storage pool name and click the “Next” button.  
NOTE: Select identical type and number of disks on each storage node to create identical storage pools.




4. Select one of the preconfigured storage profiles or create a redundancy layout for the new storage pool manually according to your redundancy, capacity, and performance requirements.



Hardware RAID, Linux Software RAID, and ZFS storage pools are supported and integrated into the StarWind CVM web interface. To make easier the storage pool configuration, the preconfigured storage profiles are provided to configure the recommended pool type and layout according to the direct-attached storage:



- hardware RAID – configures Hardware RAID’s virtual disk as a storage pool. It is available only if a hardware RAID controller is passed through to the CVM
- high performance – creates Linux Software RAID-10 to maximize storage performance while maintaining redundancy
- high capacity – creates Linux Software RAID-5 to maximize storage capacity while maintaining redundancy
- better redundancy – creates ZFS Stripped RAID-Z2 (RAID 60)) to maximize redundancy while maintaining high storage capacity
- manual – allows users to configure any storage pool type and layout with attached storage

5. Review “Summary” and click the “Create” button to create the pools on storage servers simultaneously.




## Create Volume

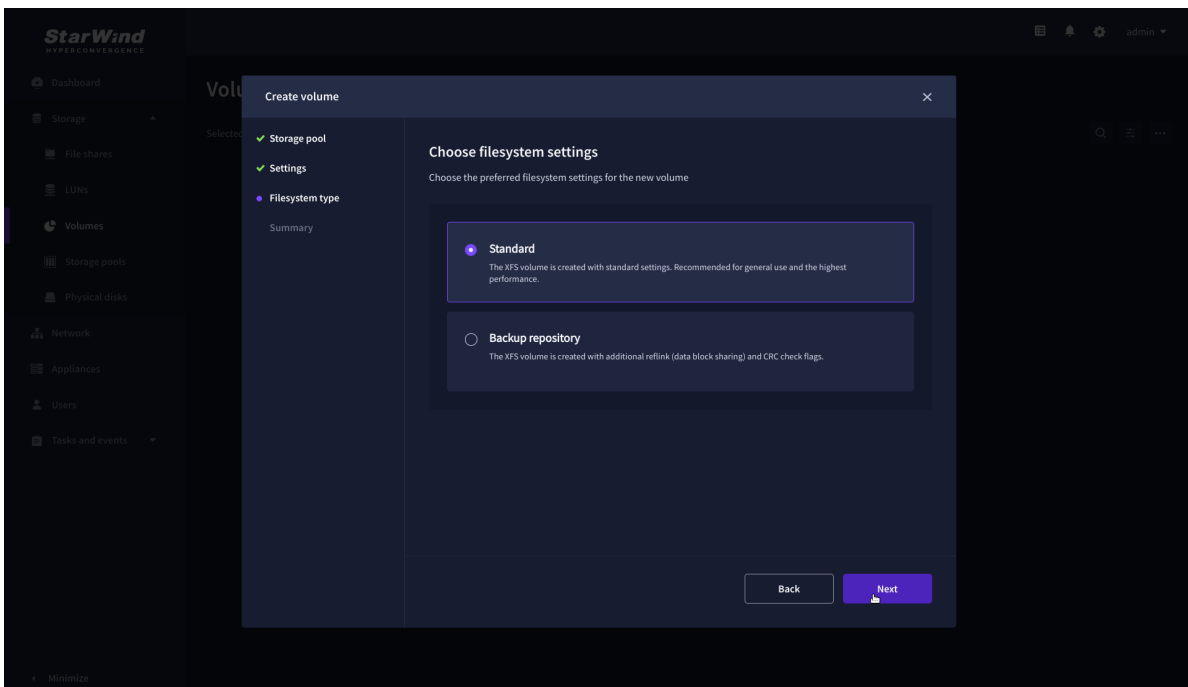
1. To create volumes, click the “Add” button.
2. Select two identical storage pools to create a volume simultaneously.




3. Specify volume name and capacity.



4. Select the Standard volume type.




5. Review "Summary" and click the "Create" button to create the pool.




## Create Ha Lun

The LUN availability for StarWind LUN can be Standalone and High availability (2-way or 3-way replication) and is narrowed by your license.

1. To create a virtual disk, click the Add button.




2. Select the protocol.

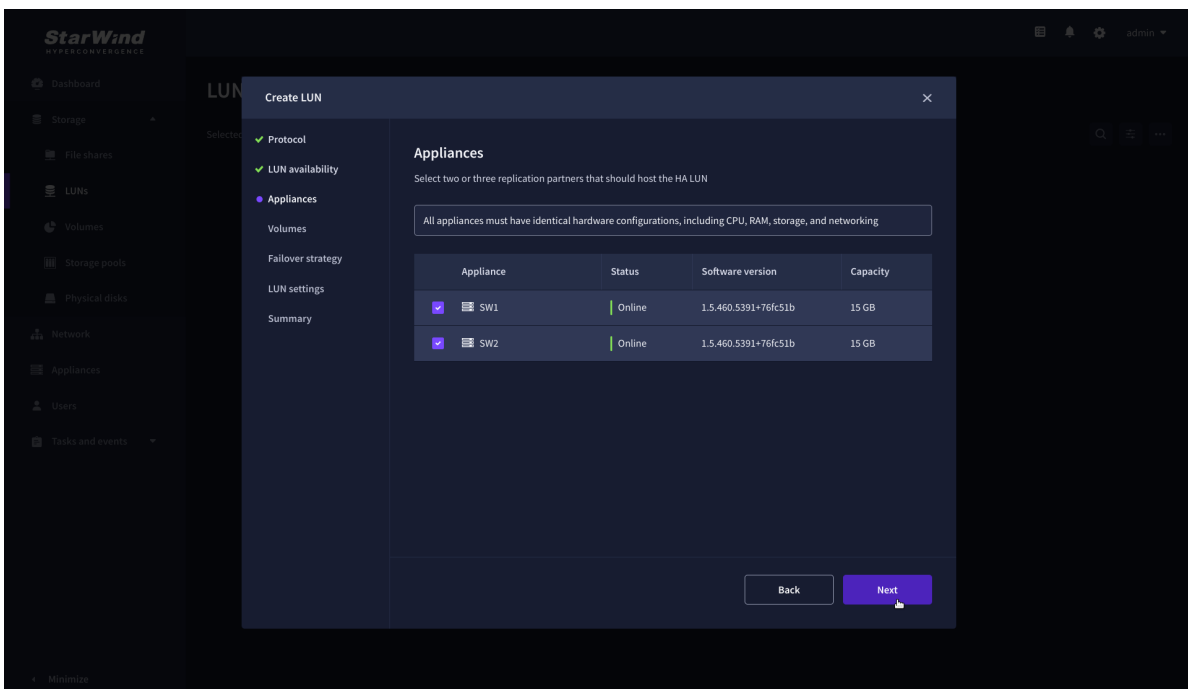


3. Choose the “High availability” LUN availability type.

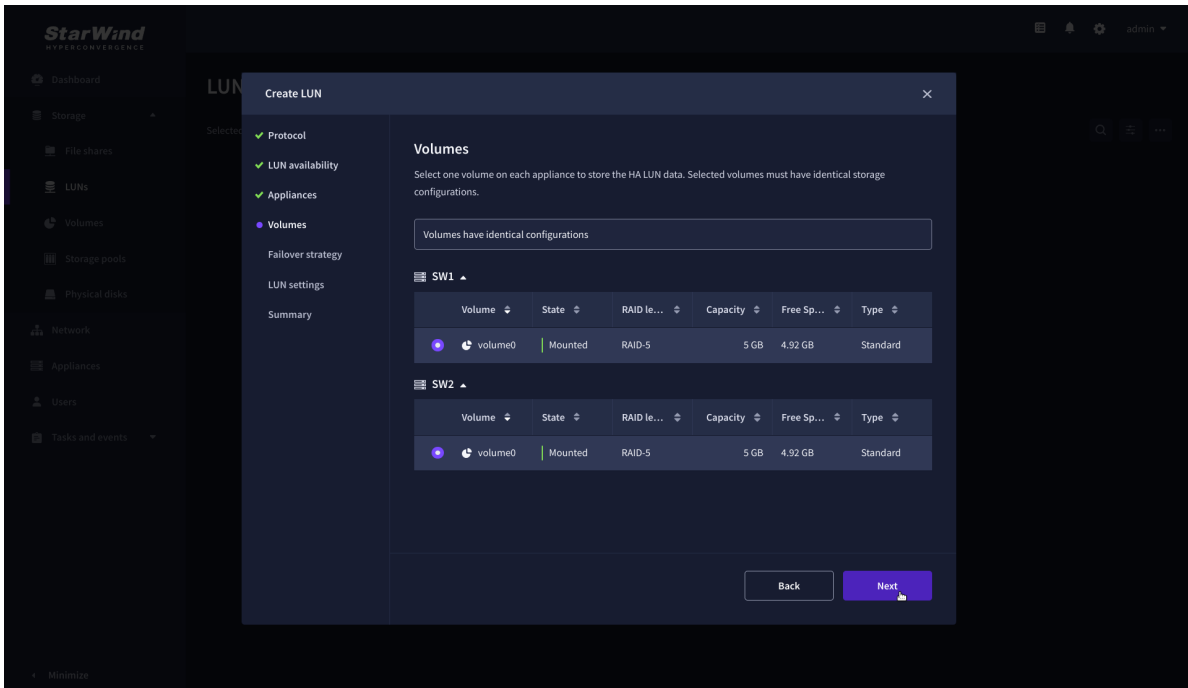




4. Select the appliances that will host the LUN. Partner appliances must have identical hardware configurations, including CPU, RAM, storage, and networking.




5. Select a volume to store the LUN data. Selected volumes must have identical storage configurations.




6. Select the “Heartbeat” failover strategy.


NOTE: To use the Node witness or the File share witness failover strategies, the appliances should have these features licensed.



7. Specify the HA LUN settings, e.g. name, size, and block size. Click Next.



8. Review “Summary” and click the “Create” button to create the LUN.



## 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\

```

1 param($addr="192.168.12.10", $port=3261, $user="root", $password="starwind",
2       $addr2="192.168.12.11", $port2=$port, $user2=$user, $password2=$password,
3
4       #common
5       $initMethod="Clear",
6       $size=2048,
7       $sectorSize=512,
8       $failover=0,
9       $bmpType=1,
10      $bmpStrategy=0,
11
12      #primary node
13      $ImagePath="VSA_Storage\mnt\crypted1",
14      $imageName="testha02",
15      $createImage=$true,
16      $storageName="",
17      $targetAlias="target02",
18      $autoSynch=$true,
19      $poolName="pool1",
20      $syncSessionCount=1,
21      $saluaOptimized=$true,
22      $cacheMode="none",
23      $cacheSize=0,
24      $syncInterface="#p2={0}:3260" -f "172.16.10.20",
25      $hbInterface="#p2=172.16.20.20:3260",
26      $createTarget=$true,
27      $bmpFolderPath="",
28
29      #secondary node
30      $ImagePath2="VSA_Storage\mnt\crypted1",
31      $imageName2="testha02",
32      $createImage2=$true,
33      $storageName2="",
34      $targetAlias2="target02",
35      $autoSynch2=$true,
36      $poolName2="pool1",
37      $syncSessionCount2=1,
38      $saluaOptimized2=$false,
39      $cacheMode2=$cacheMode,
40      $cacheSize2=$cacheSize,
41      $syncInterface2="#p1={0}:3260" -f "172.16.10.10",
42      $hbInterface2="#p1=172.16.10.10:3260",
43      $createTarget2=$true,
44      $bmpFolderPath2="",
45      )
46
47 Import-Module StarWindX
48
49 try
50 {
51     Enable-SWLog -level SW_LOG_LEVEL_DEBUG
52     $server = New-SWServer -host $addr -port $port -user $user -password $password
53     $server.Connect()
54 }
55 catch {
56     #Error handling
57 }
58

```

2. Configure script parameters according to the following example:

```

param($addr="192.168.12.10", $port=3261, $user="root",
$password="starwind",
      $addr2="192.168.12.11", $port2=$port, $user2=$user,
$password2=$password,
#common
      $initMethod="Clear",
      $size=2048,
      $sectorSize=512,
      $failover=0,
      $bmpType=1,

```

```

        $bmpStrategy=0,
#primary node
    $imagePath="VSA Storage\mnt\crypted1",
    $imageName="testha02",
    $createImage=$true,
    $storageName="",
    $targetAlias="target02",
    $autoSynch=$true,
    $poolName="pool1",
    $syncSessionCount=1,
    $luaOptimized=$true,
    $cacheMode="none",
    $cacheSize=0,
    $syncInterface="#p2={0}:3260" -f "172.16.20.20",
    $hbInterface="#p2={0}:3260" -f "172.16.10.20",
    $createTarget=$true,
    $bmpFolderPath="",
#secondary node
    $imagePath2="VSA Storage\mnt\crypted1",
    $imageName2="testha02",
    $createImage2=$true,
    $storageName2="",
    $targetAlias2="target02",
    $autoSynch2=$true,
    $poolName2="pool1",
    $syncSessionCount2=1,
    $luaOptimized2=$false,
    $cacheMode2=$cacheMode,
    $cacheSize2=$cacheSize,
    $syncInterface2="#p1={0}:3260" -f "172.16.20.10",
    $hbInterface2="#p1={0}:3260" -f "172.16.10.10",
    $createTarget2=$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, 192.168.0.1

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 -  
Format: string. Default value: Clear

-size - set size for HA-devcie (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)

#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  
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

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.

## 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:

## Provisioning Starwind Ha Storage To Windows Server Hosts


1. Launch Microsoft iSCSI Initiator: Start -> Windows Administrative Tools -> iSCSI Initiator. Alternatively, launch it using the command below in the command line interface:

```
iscsicpl
```

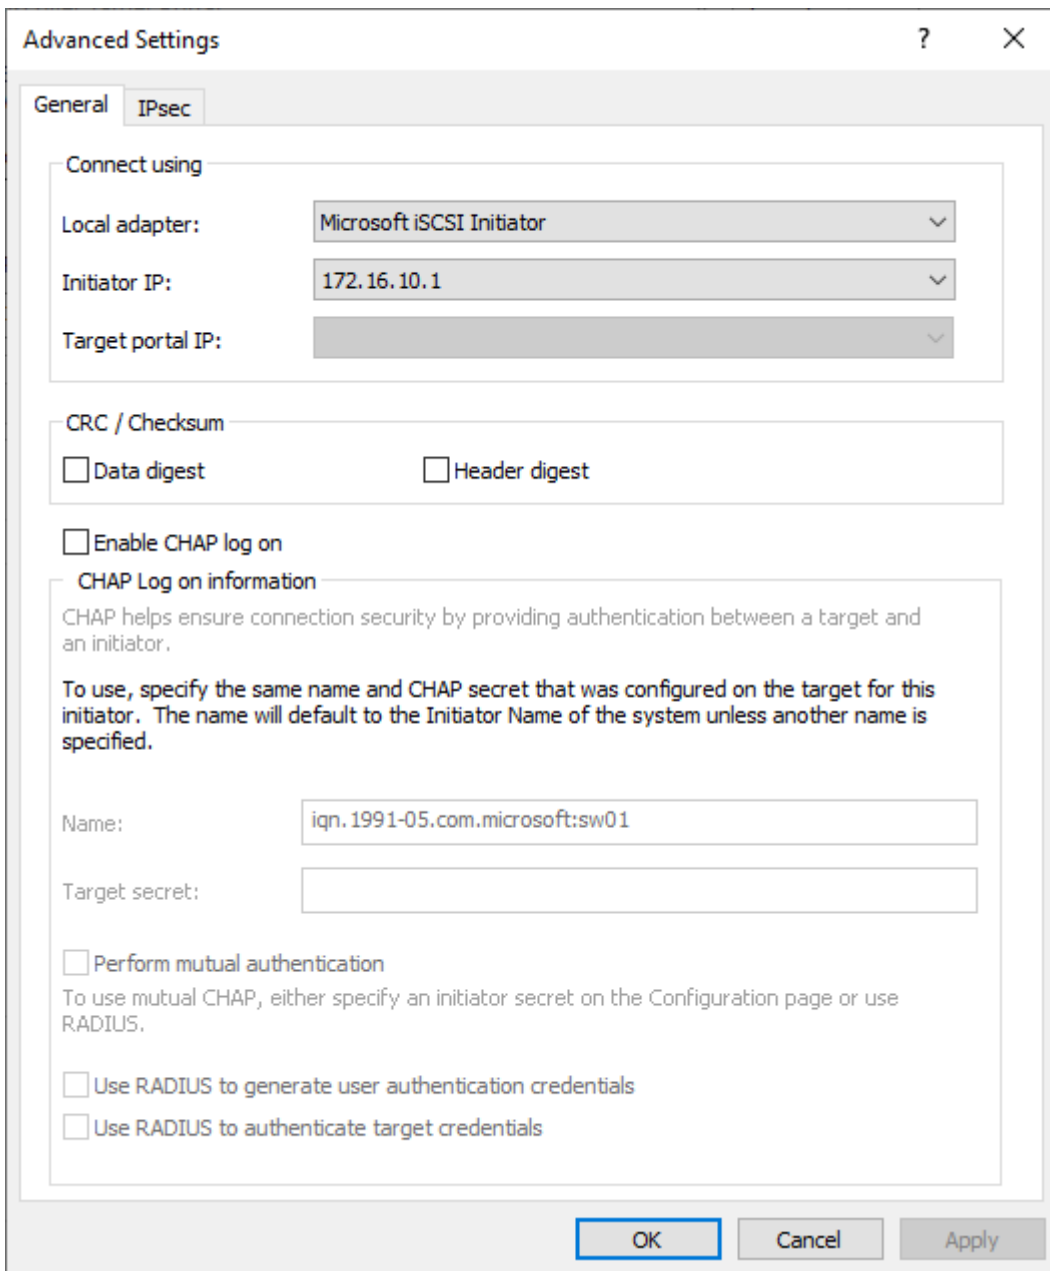
2. Navigate to the Discovery tab.



3. Click the Discover Portal button. The Discover Target Portal dialog appears. Type 172.16.10.10.



4. Click the Advanced button. Select Microsoft iSCSI Initiator as a Local adapter and select Initiator IP. Confirm the actions to complete the Target Portal discovery.



5. Click the Discover Portal... button once again.

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

Discover Target Portal


Enter the IP address or DNS name and port number of the portal you want to add.

To change the default settings of the discovery of the target portal, click the Advanced button.

IP address or DNS name:  Port: (Default is 3260.)

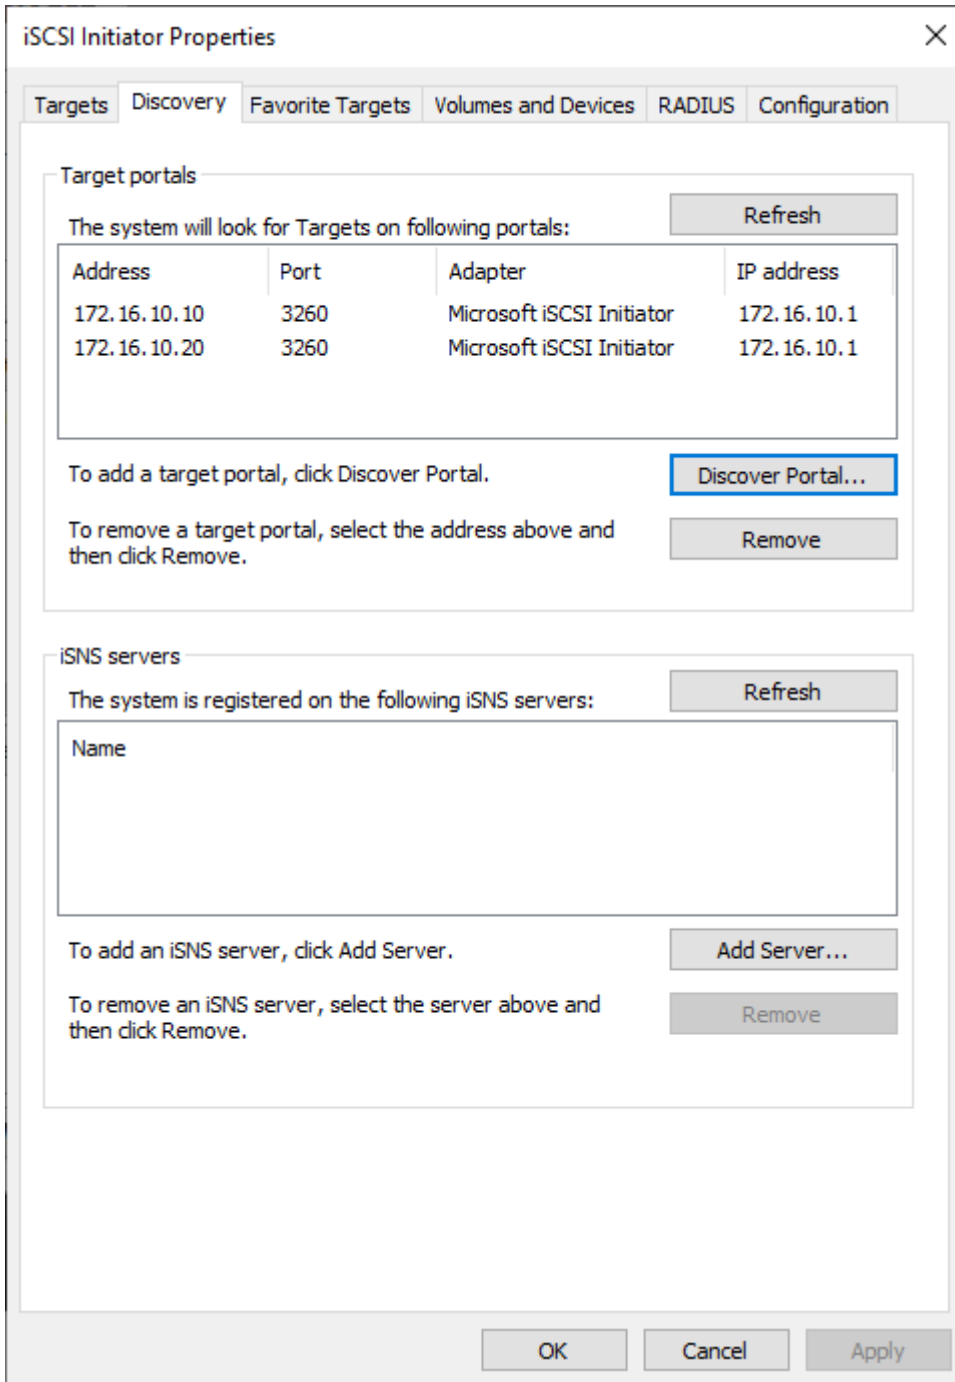
Advanced... OK Cancel

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



8. Now, all the target portals are added on the first node.






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

### 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 Server as well as the list of networks served by the StarWind Server (go to StarWind

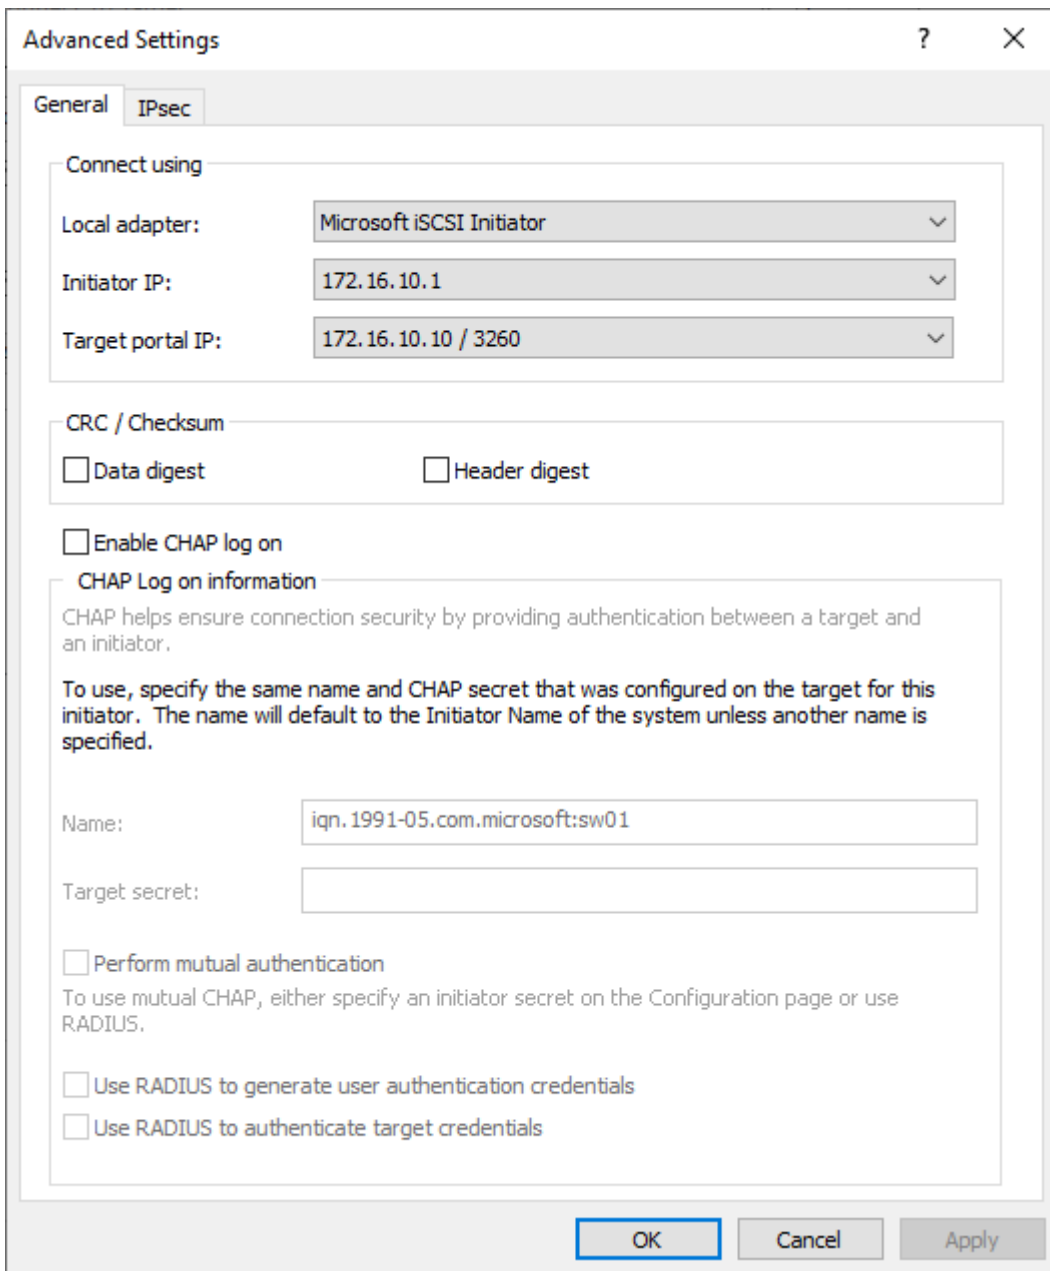
Management Console -> Configuration -> Network). Alternatively, check the Access Rights tab on the corresponding StarWind VSAN server in StarWind Management Console for any restrictions.




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 iSCSI channel. In the Target portal IP, select the corresponding portal IP from the same subnet. Confirm the actions.



5. Repeat the steps 2-4 to connect to partner node.
6. Select the CSV1 target discovered from the local server and click Connect.
7. Enable checkboxes as shown in the image below. Click Advanced.

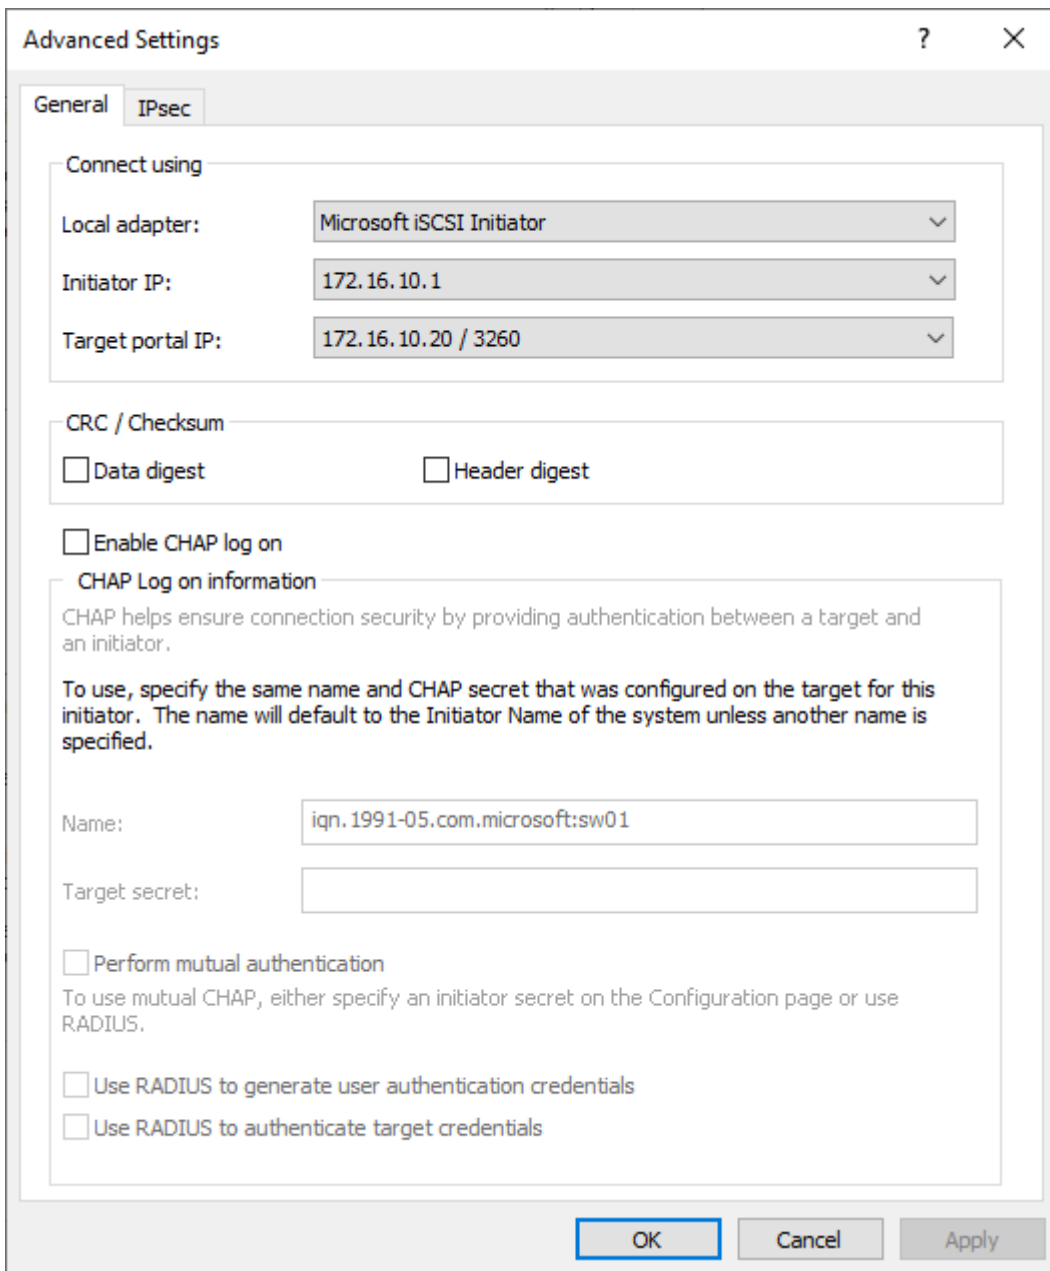


8. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Target portal IP, select 172.16.10.10. Confirm the actions.

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

10. Repeat the step 6.

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



11. Repeat the steps 1-10 for all remaining HA device targets.

12. Repeat the steps 1-11 on the other StarWind node, specifying corresponding data channel IP addresses.


### Configuring Multipath

NOTE: It is recommended to configure the different MPIO policies depending on iSCSI channel throughput. For 1 Gbps iSCSI channel throughput, it is recommended to set Failover Only or Least Queue Depth MPIO load balancing policy. For 10 Gbps iSCSI channel throughput, it is recommended to set Round Robin or Least Queue Depth MPIO

load balancing policy.

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 step 2 to bring all the remaining StarWind disks online.

## Creating A Failover Cluster In Windows Server

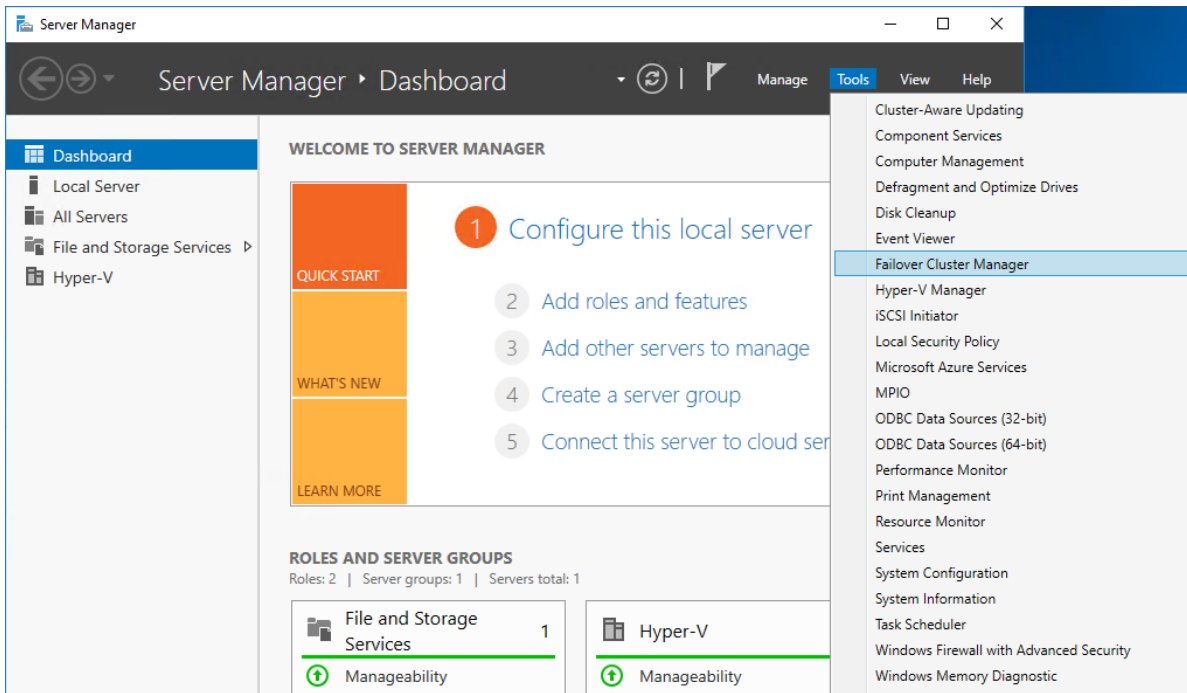
NOTE: To avoid issues during the cluster validation configuration, it is recommended to install the latest Microsoft updates on each node.

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: 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.





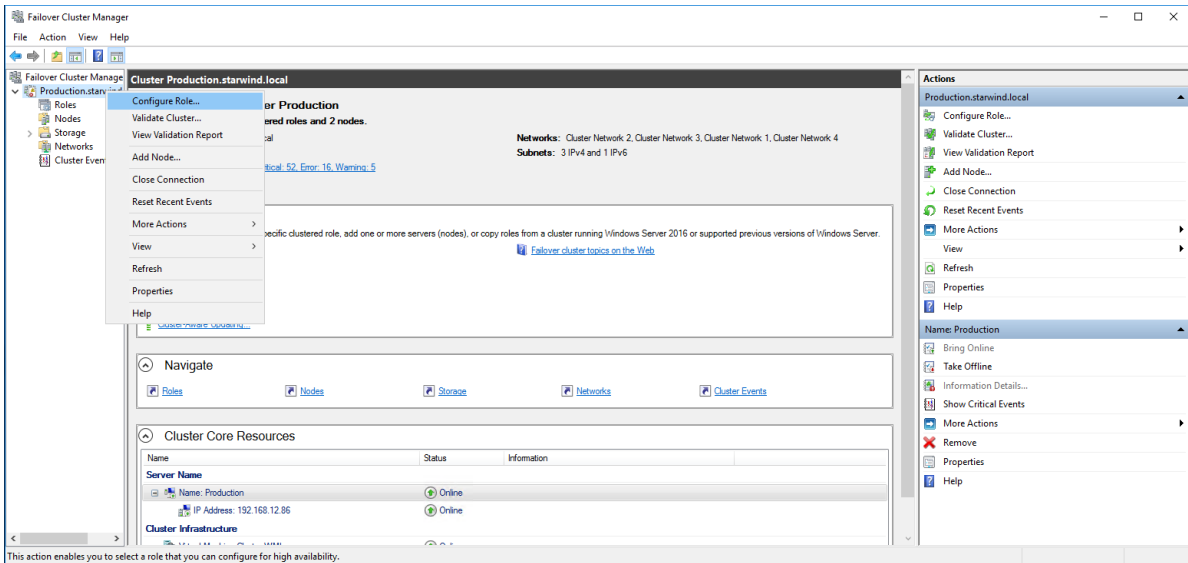
The cluster configuration is completed and it is ready for virtual machines deployment. Select Roles and in the Action tab, click Virtual Machines -> New Virtual Machine. Complete the wizard.

## 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.

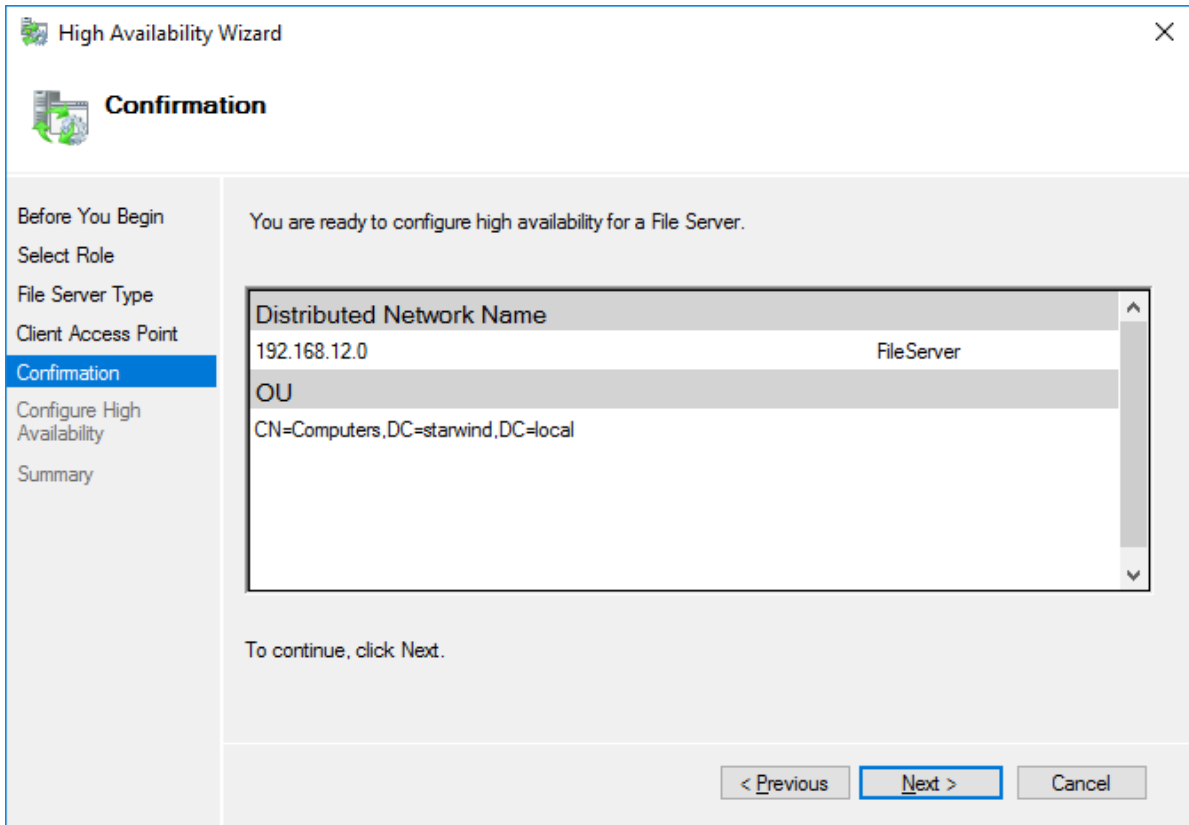


5. On the Client Access Point page, in the Name text field, type the NetBIOS name that will be used to access a Scale-Out File Server.




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






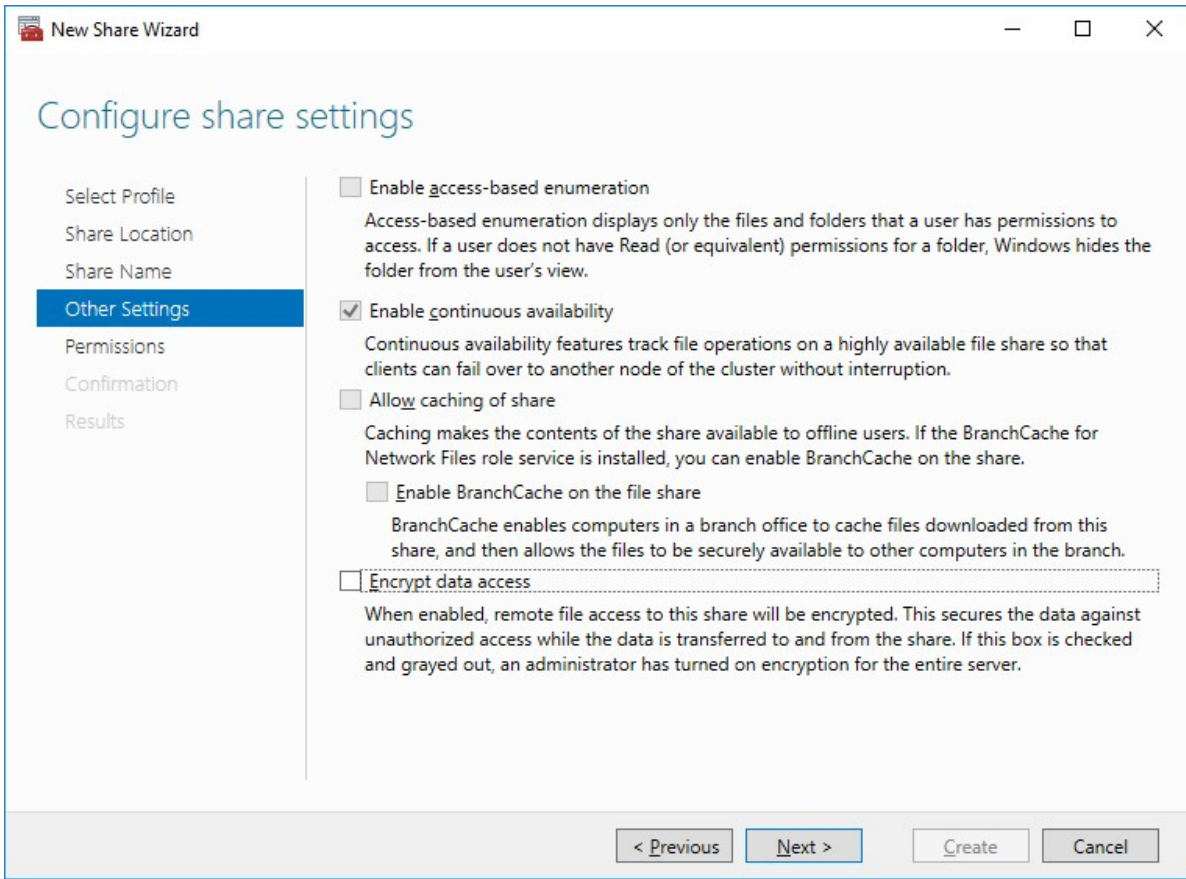
5. Select a CSV to host the share. Click Next to proceed.




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

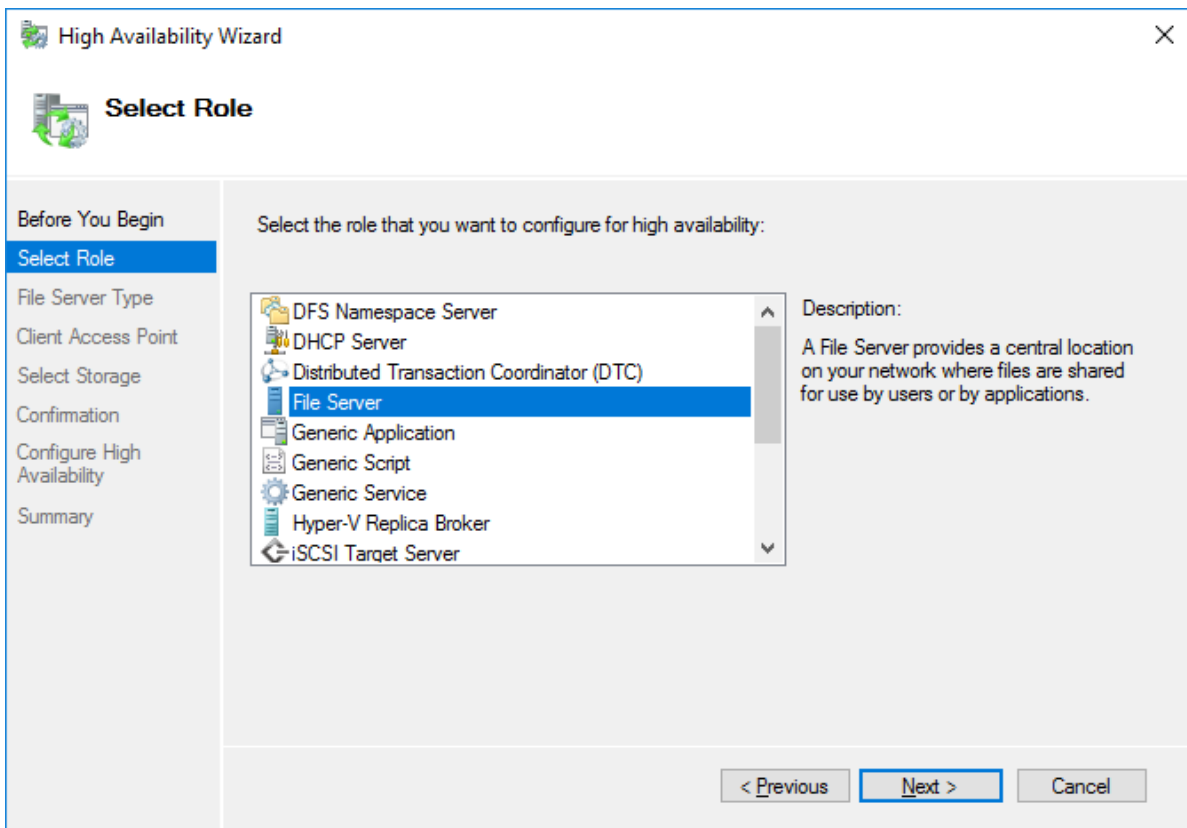
NOTE: To configure File Server for General Use, the cluster should have available storage

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.

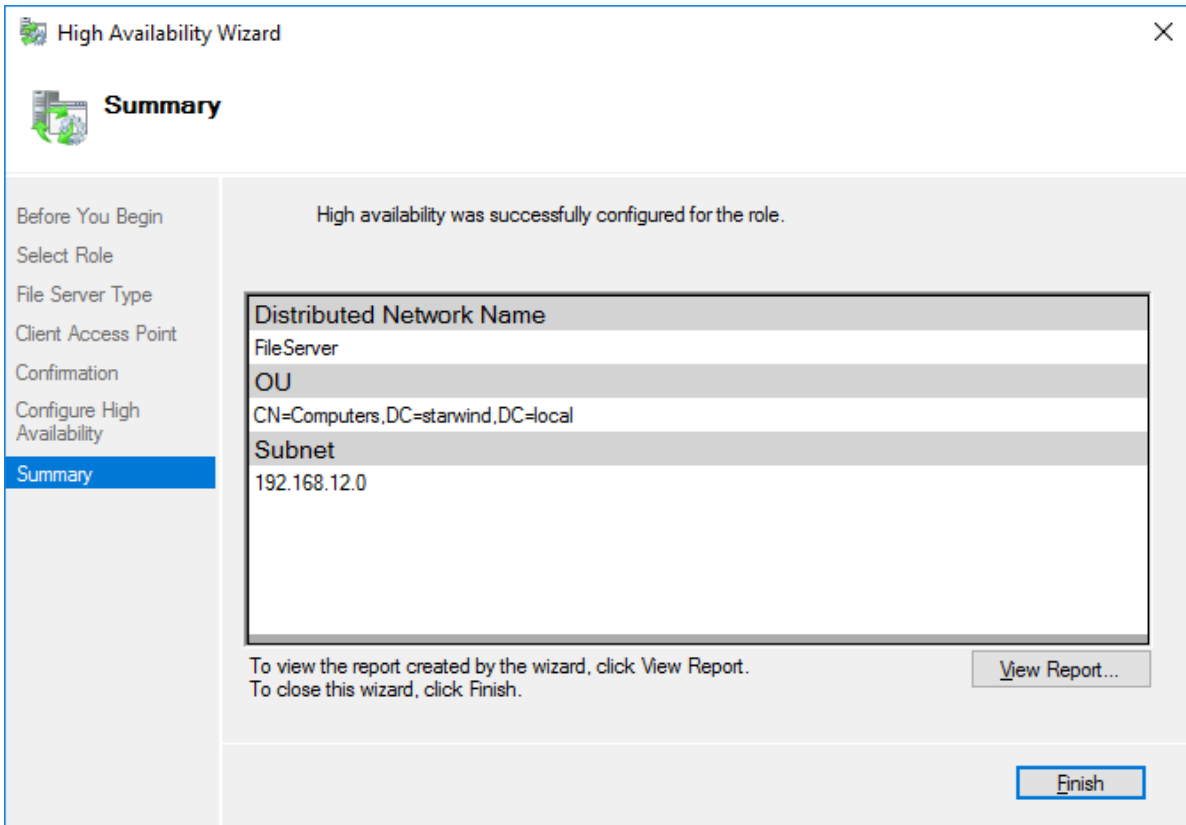


7. Check whether the specified information is correct. Click Next to proceed or Previous to change the settings.

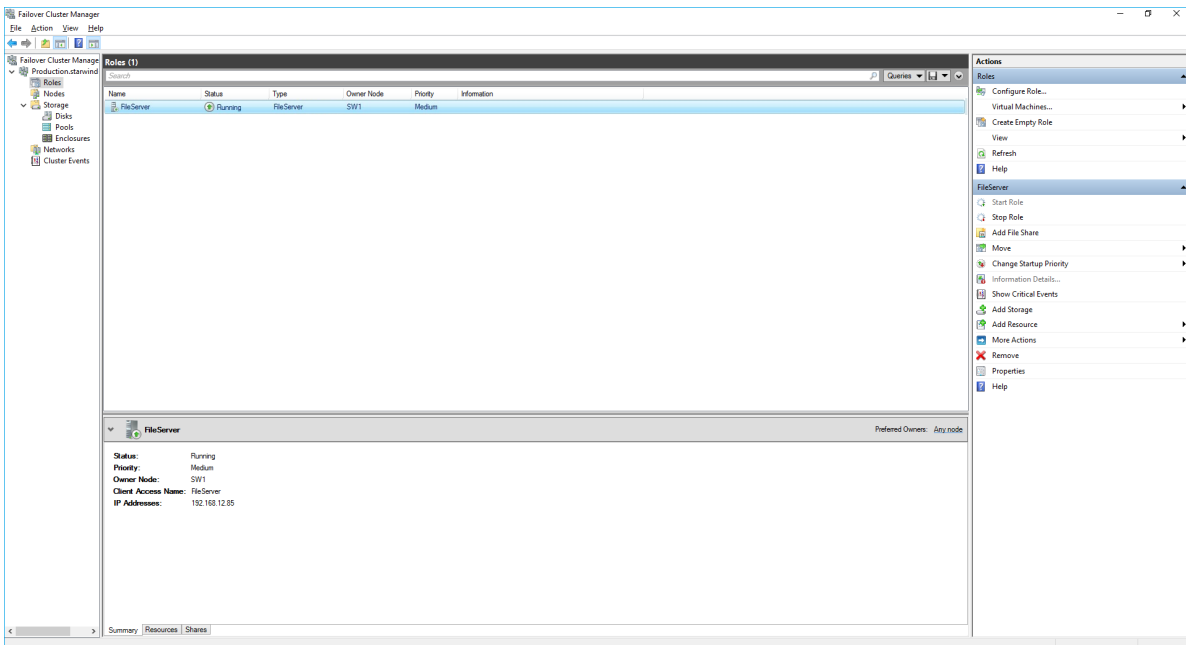


8. Once the installation has been finished successfully, the Wizard should now look like the screenshot below.

Click Finish to close the Wizard.



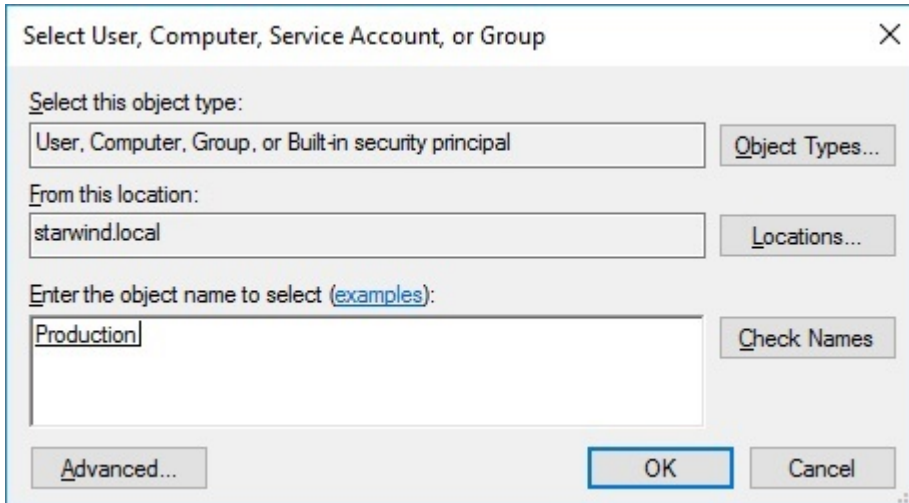
9. The newly created role should now look like the screenshot below.



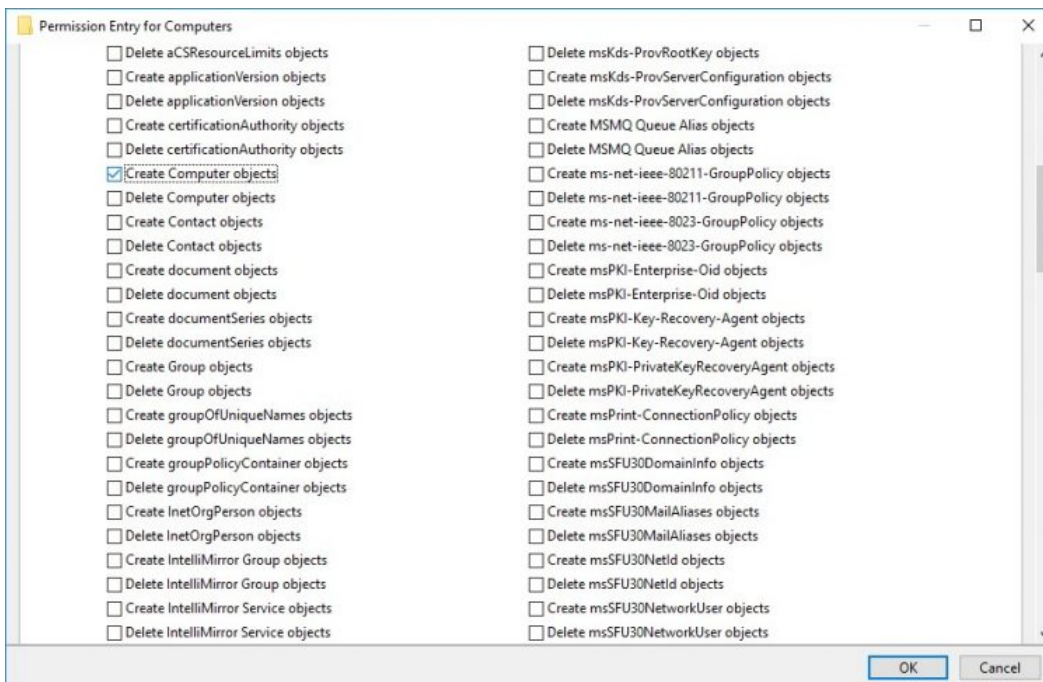
NOTE: If the role status is Failed and it is unable to Start, please, follow the next steps:

- 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



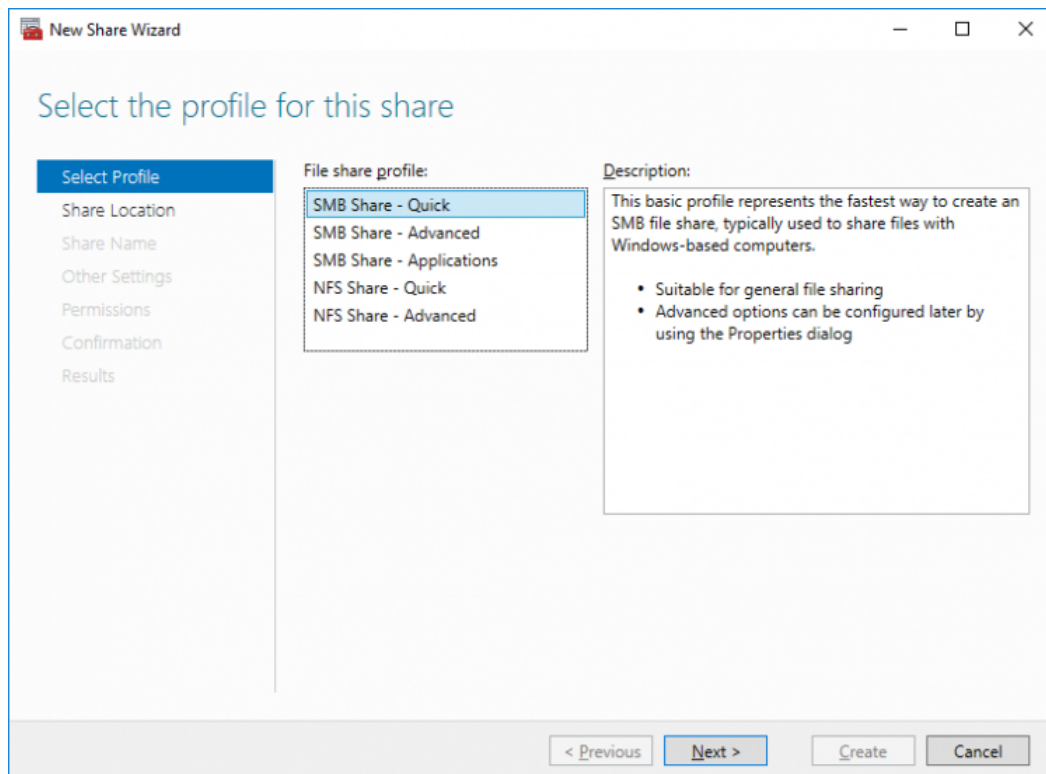
- click OK on all opened windows to confirm the changes

- 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.



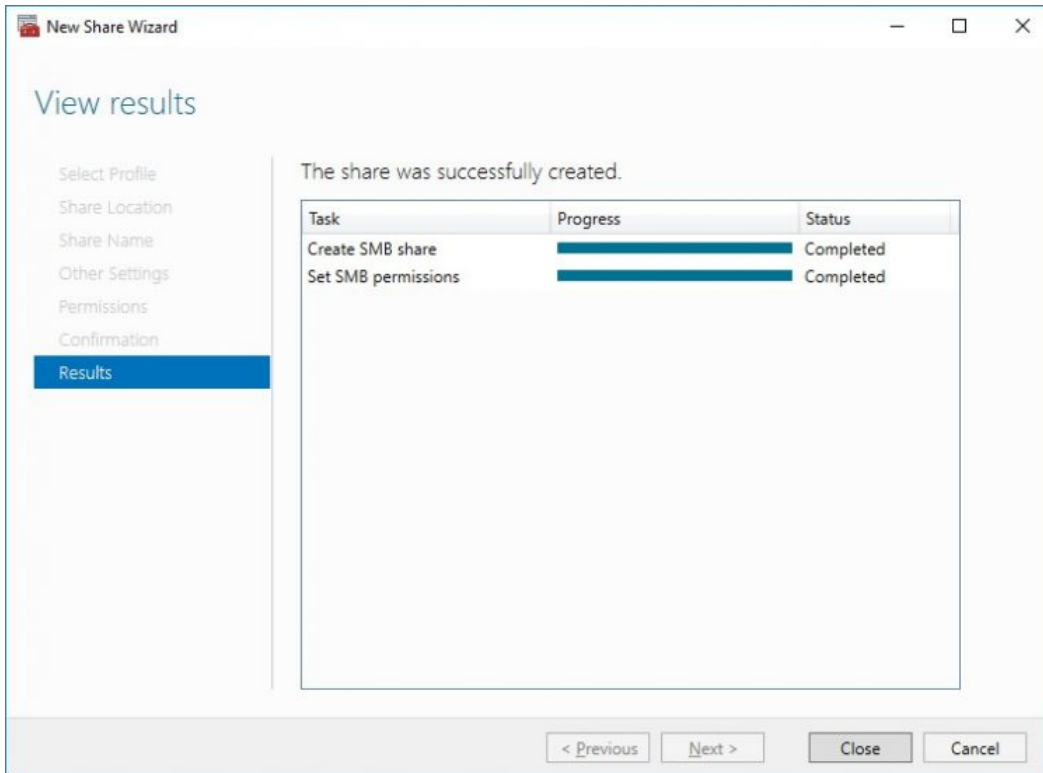
8. Specify the access permissions for the file share.



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.

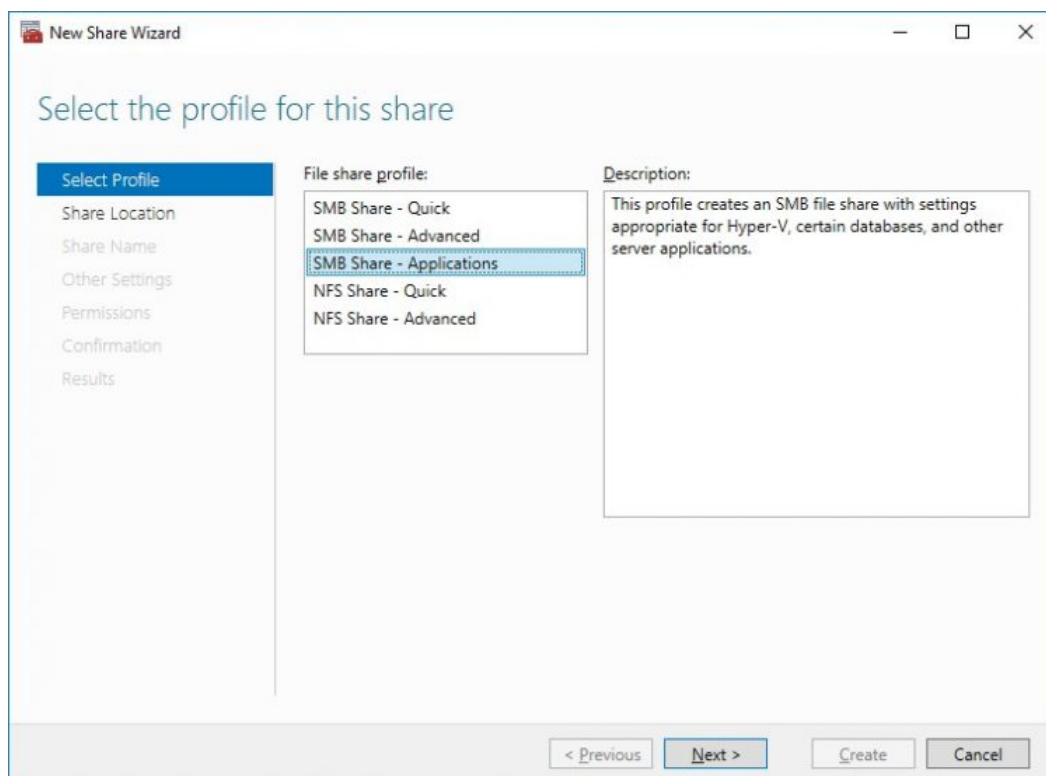
13. Choose the File Share role, select the Shares tab, right-click the created file share, and select Properties.




## 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.

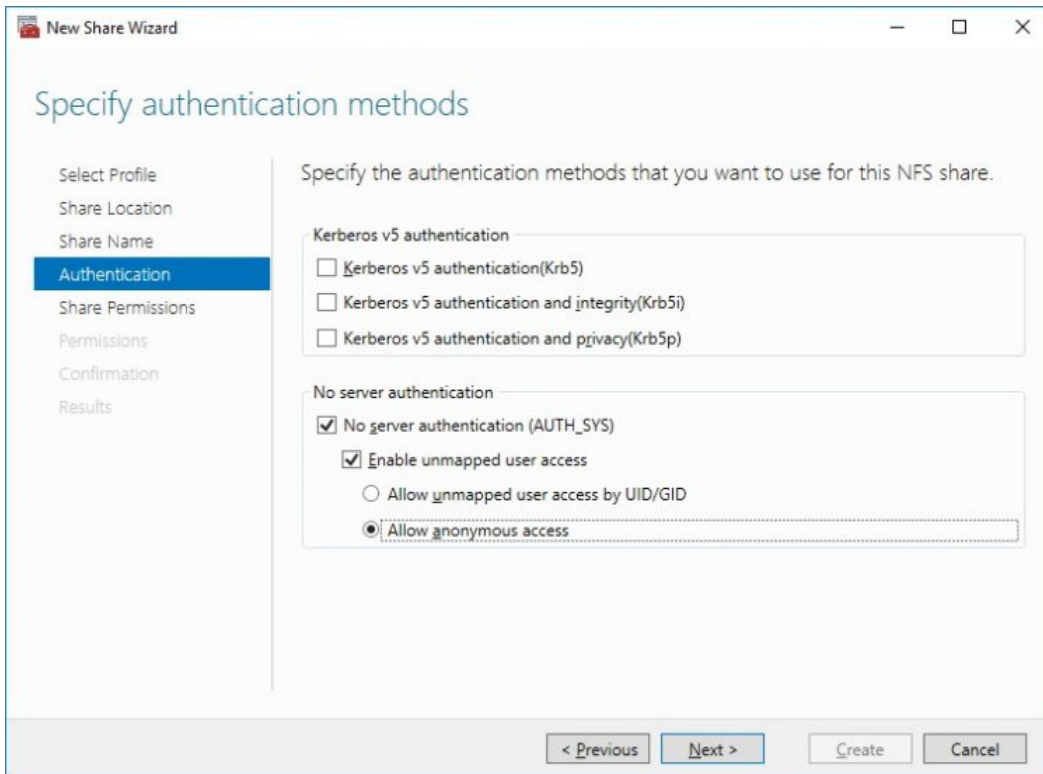


6. Type in the file share name and click Next.




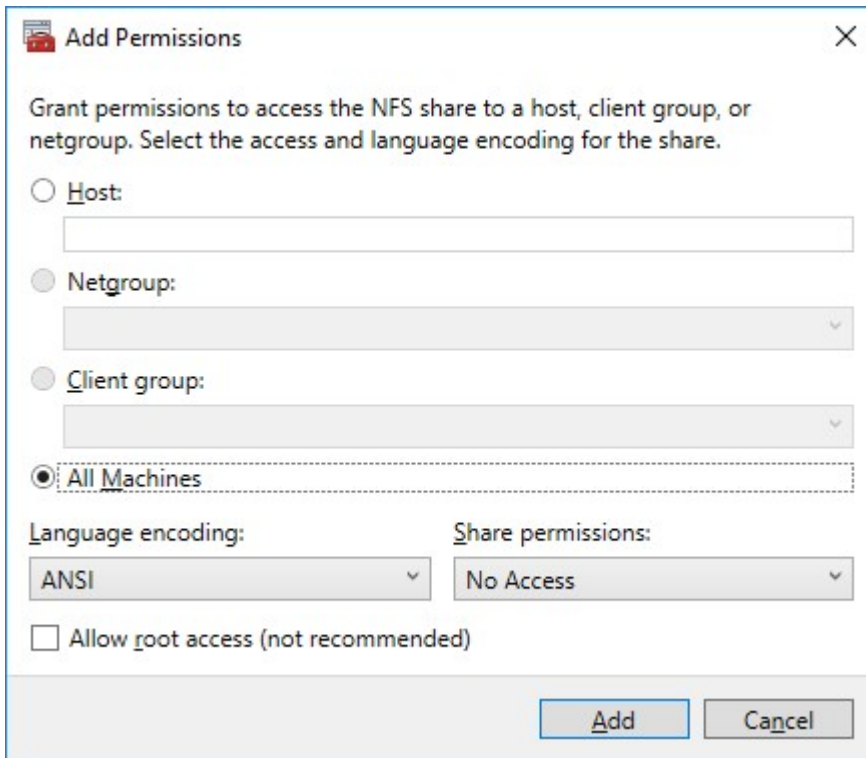
7. Specify the Authentication. Click Next and confirm the message in pop-up window to

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



## 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) ©2024, StarWind Software Inc. All rights reserved.