

StarWind Virtual SAN: Configuration Guide for Microsoft Windows Server [Hyper-V], VSAN Deployed as a Windows Application

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.

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

Annotation

Relevant products

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

Purpose

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

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

Audience

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

Expected result

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

Prerequisites

StarWind Virtual SAN system requirements

Prior to installing StarWind Virtual SAN, please make sure that the system meets the requirements, which are available via the following link:

<https://www.starwindsoftware.com/system-requirements>

Recommended RAID settings for HDD and SSD disks:

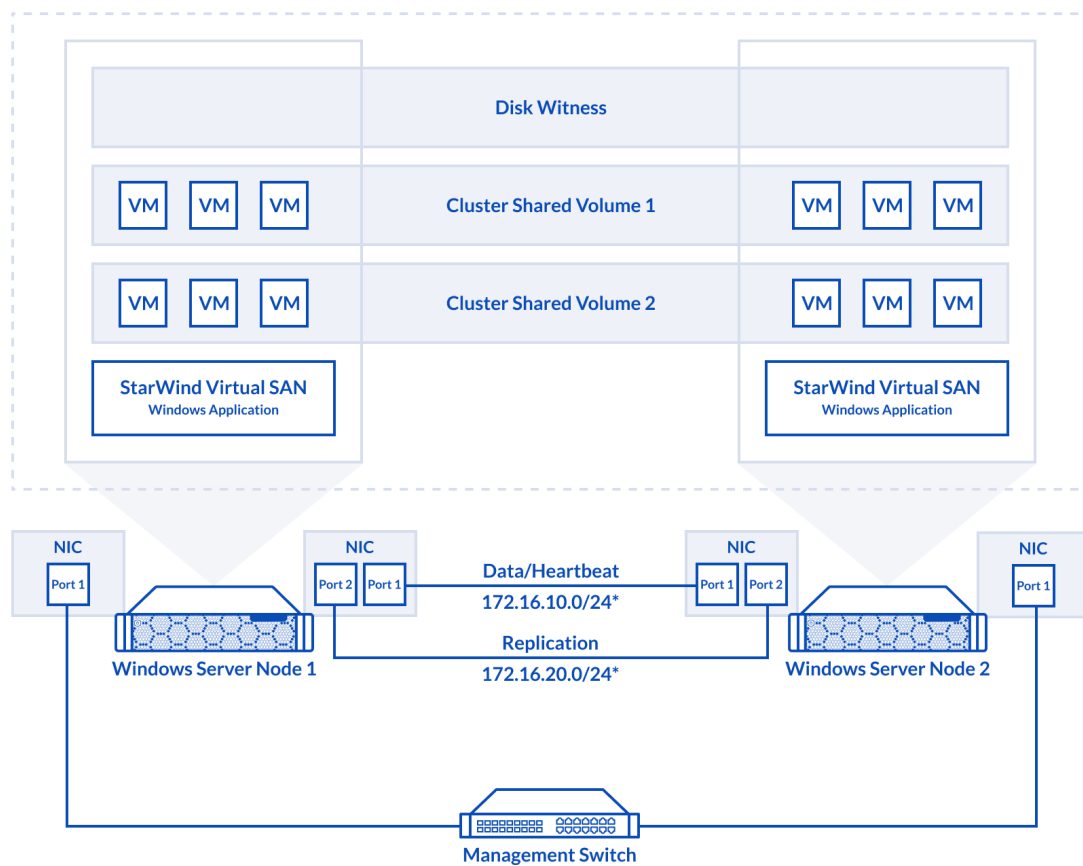
<https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/>

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


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

Solution diagram

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



2-node cluster



3-node cluster

Preconfiguring cluster nodes

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

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

2. 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 (2 node scenario) or 4x network interfaces (3 node scenario) 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).

For 2-node scenario, 172.16.10.x subnet is used for iSCSI/StarWind heartbeat traffic, while 172.16.20.x subnet is used for the Synchronization traffic.

For 3-node scenario, 172.16.10.x, 172.16.11.x, 172.16.12.x subnets are used for the

iSCSI/StarWind heartbeat traffic, while 172.16.20.x,172.16.21.x, 172.16.22.x subnets are used for the Synchronization traffic.

4. Set MTU size to 9014 or 9000 depending on network cards vendor recommendations 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 9014 (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.

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

```
mpiocpl
```

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



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

8. Repeat the same procedure on the other server.

Automated Storage Tiering Configuration

In case of using Automated Storage Tiering, the individual spinning disks and flash drives can be either connected into OS directly in Pass-Through mode or preconfigured into separate virtual “SSD” and “HDD” LUNs, and then attached into the host OS.

NOTE: Please make sure that the final configuration which is expected to be put into production is fully supported by either the software vendor (i.e., StarWind, Microsoft, VMware, etc.) or particular MSP/ISV who is in charge of the whole project. Please ensure that specific workload expectations can be met with both existing and future deployments. Depending on the configuration it can be either self-supported, partially supported, or fully supported.

NOTE: Simple Tier has no redundancy built-in, so in case of either individual spinning disk and/or flash drive failure there is a risk of losing the data. It is recommended to

configure resilient storage to be used as the underlying storage for the Tier. Please make sure that such a configuration is supported by either the software vendor (i.e., StarWind, Microsoft, VMware, etc.) or/and MSP/ISV.

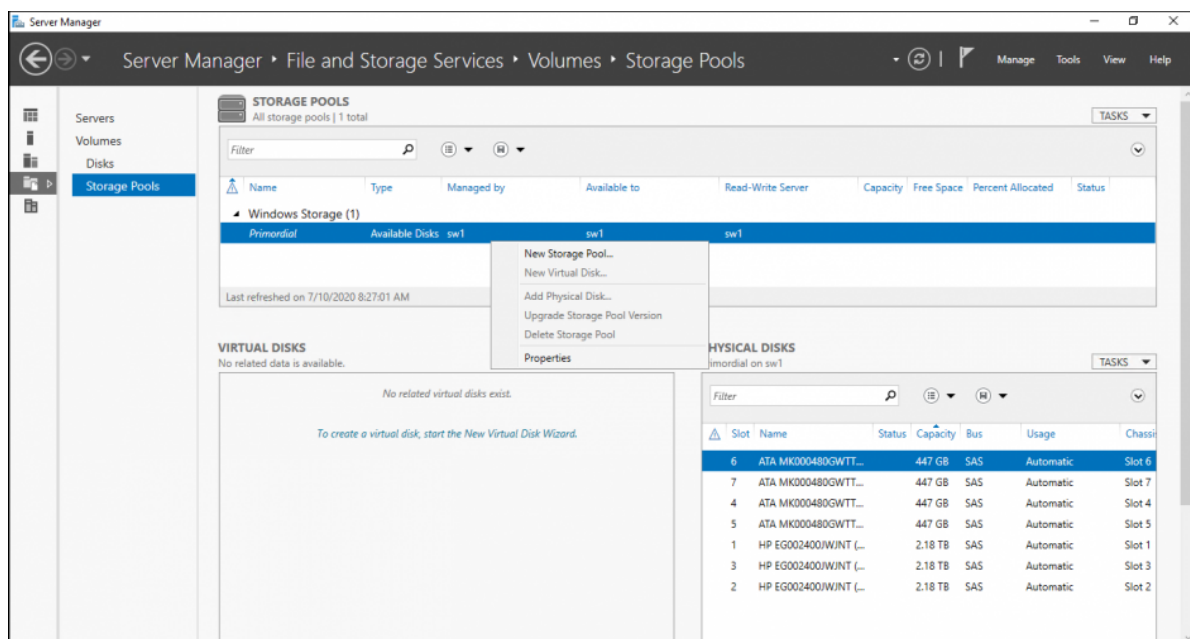
Automated Storage Tier creation

There are two ways to configure Automated Storage Tiering. It can be done via Server Manager and via the PowerShell console.

The first level of Storage Tier is Storage pools. At this level, separate physical disks are united into a single pool, providing the ability to flexibly expand the capacity and delegate administration.

The upper level is Storage Spaces. At this level, virtual disks are created using the available capacity of a storage pool. Storage Spaces feature the following characteristics: resiliency level, storage tiers, fixed provisioning, and precise administrative control.

1. Launch Server Manager->File and Storage Services->Volumes->Storage Pools. All disks available for Storage Pool are listed in Physical Disks. Click New Storage Pool.



NOTE:

`Get-PhysicalDisk`

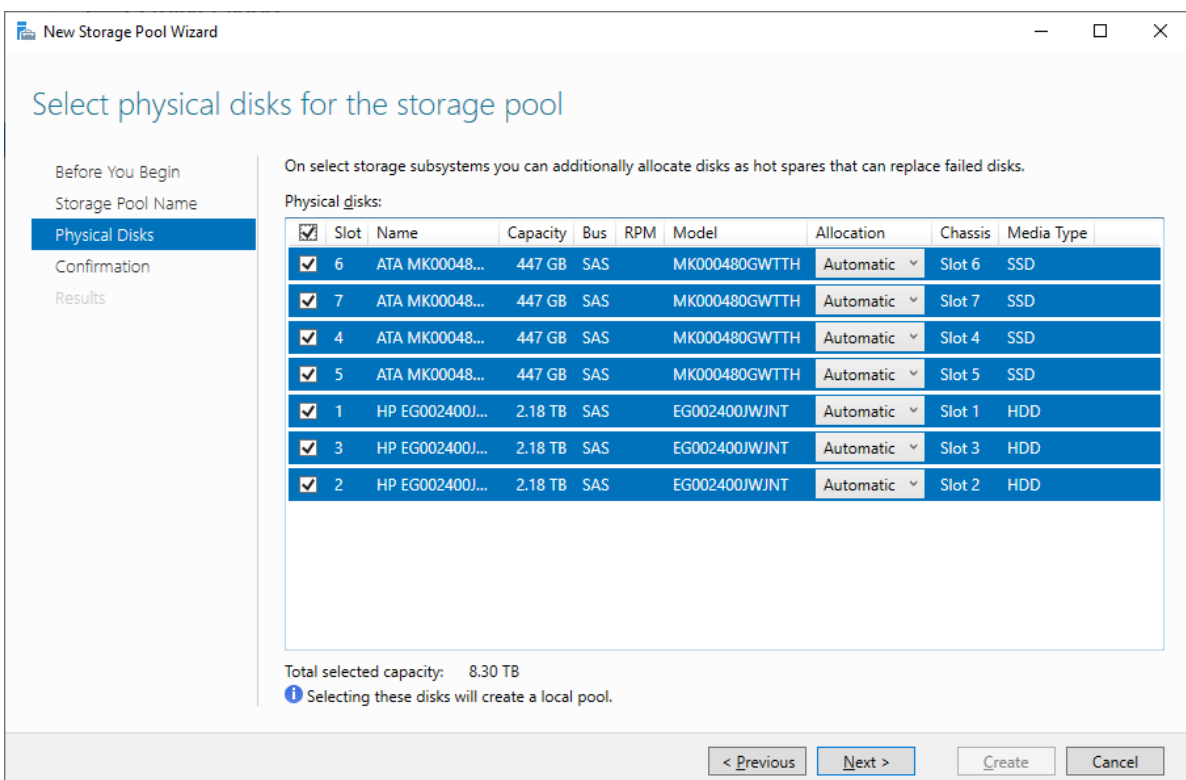
is a PowerShell command that allows checking the disks available for Storage Pool.


```
Get-PhysicalDisk | sort-object SlotNumber | select SlotNumber,
FriendlyName, Manufacturer, Model, PhysicalSectorSize,
LogicalSectorSize | ft
```


is a PowerShell command that allows checking the parameters of physical disks.

2. Specify a Storage Pool name:

3. Select the disks for Storage Pool and then press Next. In case of using Storage Tiers with both SSDs and HDDs, all these disks need to be added into the Storage Pool.



4. Confirm the correct settings and click Create to create Storage Pool.



NOTE: There might be cases when the MediaType parameter should be changed manually. It can be done with the following PowerShell commands:

Assign SSD MediaType for the disk with size less than []GB:

```
Get-PhysicalDisk | where Size -lt [ ]GB | Set-PhysicalDisk -
MediaType SSD
```

Assign HDD MediaType for the disk with size more than []GB:

```
Get-PhysicalDisk | where Size -gt [ ]GB | Set-PhysicalDisk -
MediaType HDD
```

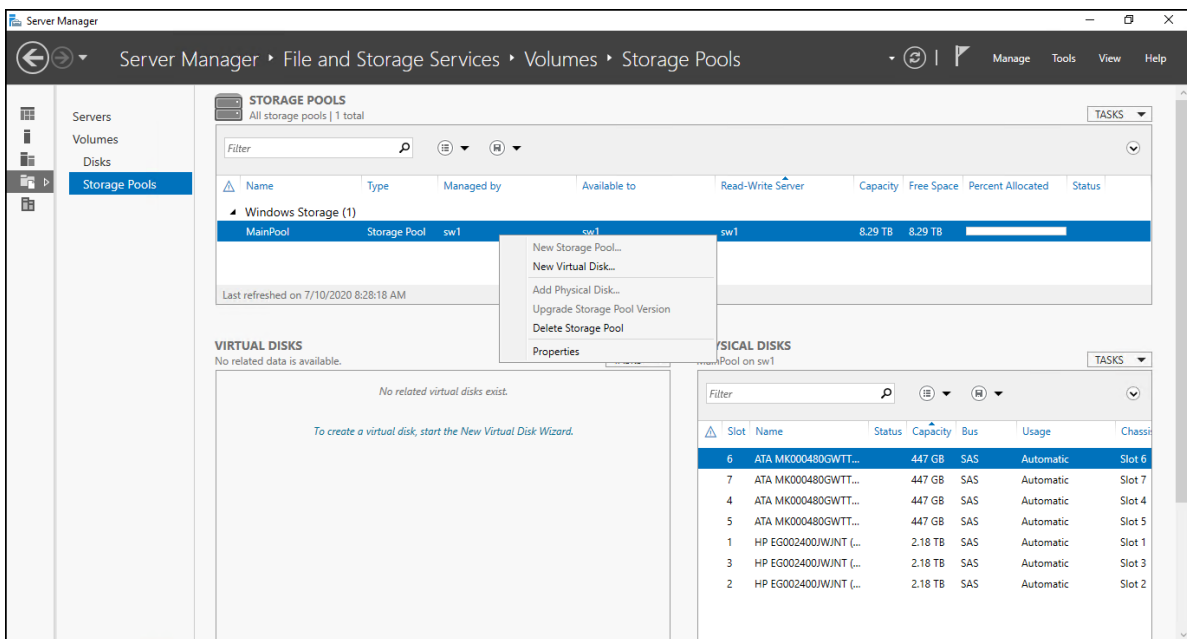
Additionally, the following commands can be used:

```
Get-PhysicalDisk | ft FriendlyName,CanPool,Size,MediaType
Set-PhysicalDisk -FriendlyName [disk name] -MediaType [SSD or
HDD]
```

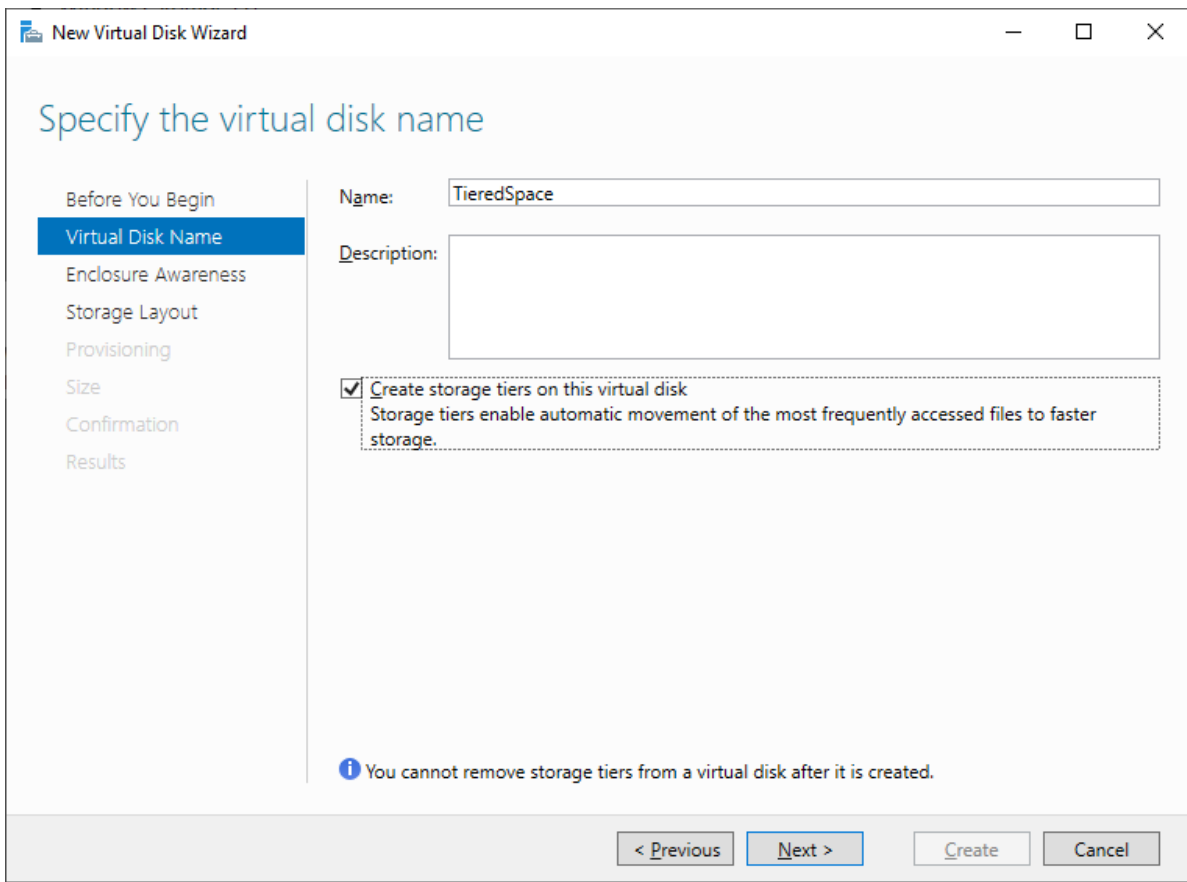
or

```
Get-PhysicalDisk | ft FriendlyName,CanPool,Size,MediaType
Get-PhysicalDisk | Where Size -EQ [disk size] | Set-PhysicalDisk -MediaType [SSD or HDD]
```

5. The next step is to create a virtual disk on the storage pool. It is possible to create multiple virtual disks that exist in the storage pool and then create multiple volumes that exist in each virtual disk. Create new virtual disk by right-clicking on the storage pool and selecting New Virtual Disk.

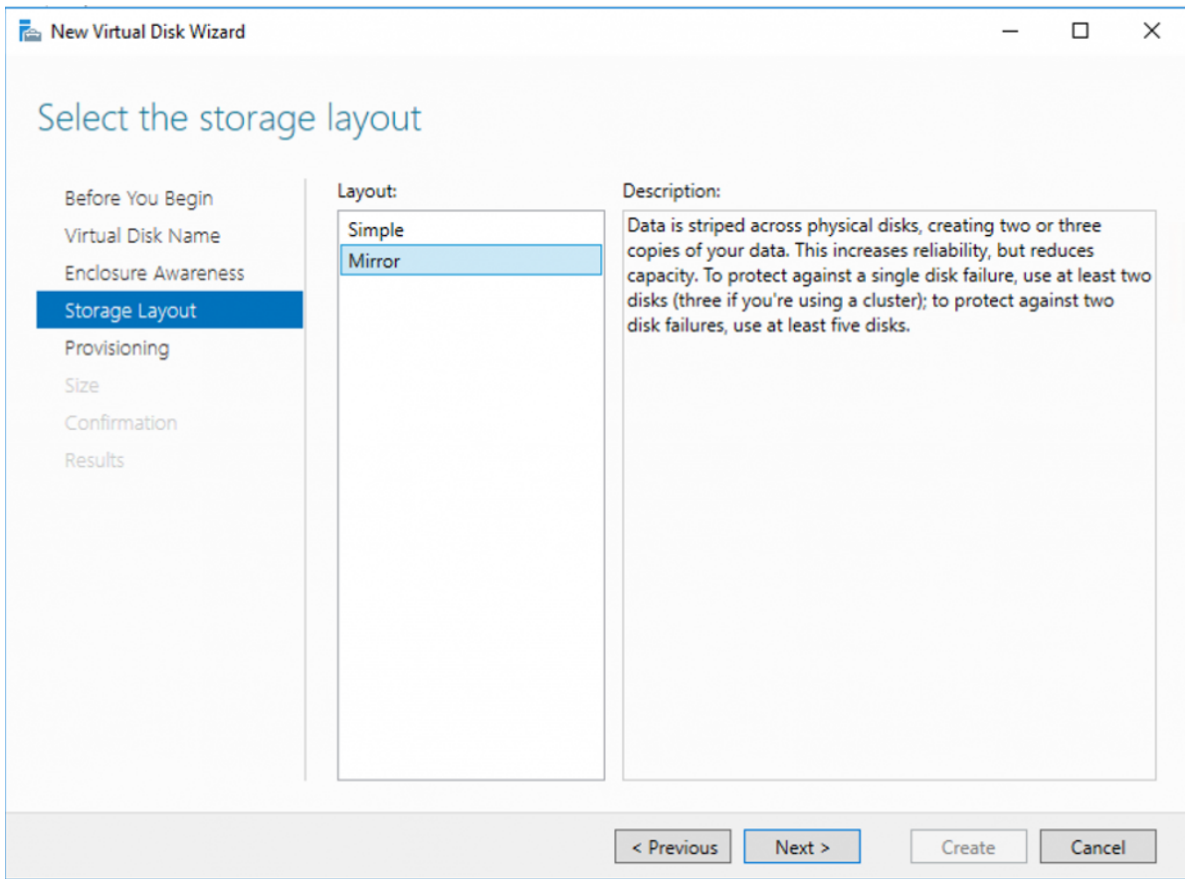


6. For Automated Storage Tiering, both HDD- and SSD-based disks or storage arrays should be in the storage pool to make use of Storage Tiers. In case of using Storage Tiers, Storage Layout can be only Simple and Mirror. Specify Virtual Disk Name and select Create storage tiers on this virtual disk.



NOTE: Simple Tier has no redundancy built-in, so in case of either individual spinning disk and/or flash drive failure there is a risk of losing the data. It is recommended to configure resilient storage to be used as the underlying storage for the Tier. Please make sure that such a configuration is supported by either the software vendor (i.e., StarWind, Microsoft, VMware, etc.) or/and MSP/ISV.

7. Select the storage layout type. Under the Simple layout, the data is striped across physical disks. This would be equivalent to a RAID-0 configuration. In case of using at least two disks, the Mirror configuration can be configured. The Mirror is equivalent to RAID-1. Once done, click next.




8. Specify the provisioning type.

Fixed. This provision type means that virtual disk cannot exceed the actual storage pool capacity.

Thin. This provision type means that there is a possibility to create a volume with a size exceeding the storage pool capacity and then add physical disks later.


Choose fixed disk provisioning since this type is required by Storage Tiers. Click Next.

9. Specify the size of the Virtual Disk.



NOTE: At least 8 GB of free space on each Tier should be provisioned to allow Automated Storage rebuilding in case of the disk loss.

10. Confirm the settings and click Create to create Virtual Disk.



NOTE: In case of using both SSD and HDD disks or virtual LUNs, automated Storage Tier consists of the so-called “hot” and “cold” Tiers. Automated Storage Tier elaborates a data map taking into account how often the certain data is used, thus defining how hot separate data blocks are. During the process of optimization that is launched automatically every day, the hot data, i.e. data that is used on the most frequent basis, is transferred to the fast SSD tier, with the data used less frequently, the so-called cold data, being transferred to the slower HDD tier.

As the SSD tier-based data gets updated only once a day, it is possible to manually optimize it with the help of the following CMD one-liner:

```
defrag.exe /C /H /K /G
```

This command should be run on all cluster nodes, as it optimizes only those virtual disks the owner node for which is the one where the command is running.

For certain files, it can be optimal to permanently stay on the SSD tier. An example is a VHDX file that is accessed frequently and requires minimum latency and high performance. Such a result can be achieved by pinning the file to the SSD tier.

The following recommendations should be taken into account before running the command:

- the command should be run from the node owning the storage (Cluster Shared Volume) with the file stored on it.
- local path to the storage (Cluster Shared Volume) on the node should be used.

After a file is pinned, it will stay in the tier until the next optimization process triggered either automatically or manually.

To pin files to the SSD tier, run the following PowerShell command:

```
Set-FileStorageTier -FilePath <localFilePath> -  
DesiredStorageTierFriendlyName<ssdTierName>
```

To unpin files from the SSD tier, run the following PowerShell command:

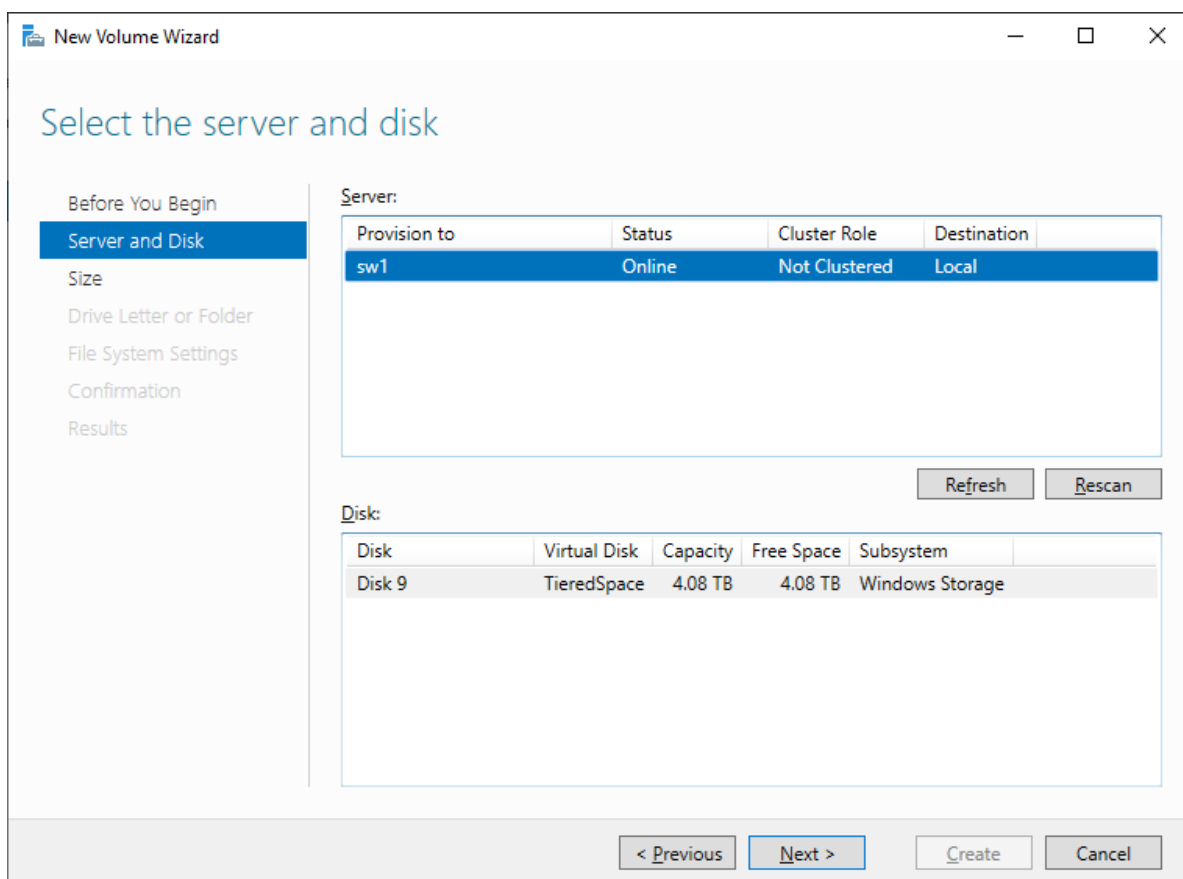
```
Set-FileStorageTier -FilePath <localFilePath>
```

The below PowerShell command lists all files that are currently pinned:

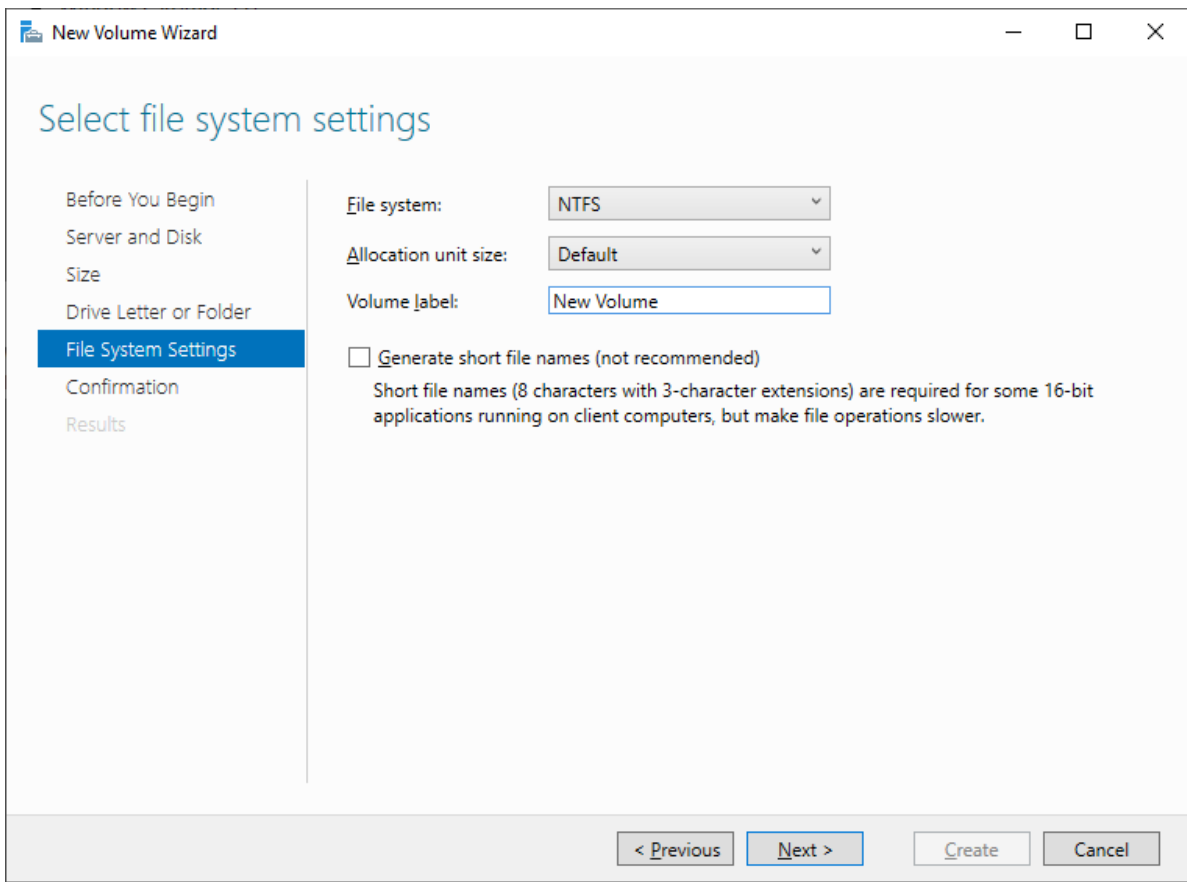
```
Get-FileStorageTier -VolumePath <csvVolumePath>
```

11. Create a New Volume using New Volume Wizard:

12. Select the server and disk and click Next.



13. Select the file system settings and click Next to proceed.



NOTE: The steps described above can be performed with help of PowerShell commands. Also, with help of PowerShell, additional parameters can be configured for better performance:

Set 64K size of interleave: *-Interleave 65536*.

Set *LogicalSectorSizeDefault 4096* instead of default 512.

The cache size can be changed with the help of *-WriteCacheSize [JGB* parameter. It is possible to set cache size only via PowerShell commands for creating Automated Storage Tier.

Set SSD tier in two-way mirror: *ResiliencySettingName Mirror -NumberOfDataCopies 2*

The number of threads can be set with *-NumberOfColumns* parameter. The recommended number is the number of SSDs divided by 2.

The example of the PowerShell commands for Storage Pool and Virtual Disk with Tiered Storage creation is provided below:

Get-StorageSubsystem – check the storage subsystem name before running the commands below.

```
Get-PhysicalDisk
```

```
$disks = Get-PhysicalDisk |? {$_.CanPool -eq $true}
```

```
New-StoragePool -StorageSubSystemFriendlyName "[ ]*" -
FriendlyName [ ] -PhysicalDisks $disks -
LogicalSectorSizeDefault 4096
```

```
Get-PhysicalDisk | where Size -lt [ ]GB | Set-PhysicalDisk -
MediaType SSD
```

```
Get-PhysicalDisk | where Size -gt [ ]GB | Set-PhysicalDisk -
MediaType HDD
```

```
Get-StoragePool -FriendlyName [ ]
```

```
New-StorageTier -MediaType SSD -StoragePoolFriendlyName [ ] -
FriendlyName SSDTier -ResiliencySettingName [Simple, Mirror or
Parity] -NumberOfDataCopies 2 -NumberOfColumns [ ] -Interleave
65536
```

```
New-StorageTier -MediaType HDD -StoragePoolFriendlyName [ ] -
FriendlyName HDDTier -ResiliencySettingName [Simple, Mirror or
Parity] -Interleave 65536
```

```
$SSD = Get-StorageTier -FriendlyName SSDTier
```

```
$HDD = Get-StorageTier -FriendlyName HDDTier
```

```
New-VirtualDisk -FriendlyName "[ ]" -StoragePoolFriendlyName [
] -StorageTiers $SSD, $HDD -StorageTierSizes [ ]GB, [ ]Gb -
ResiliencySettingName [Simple, Mirror or Parity] -
ProvisioningType fixed -WriteCacheSize [ ]GB
```

IMPORTANT NOTE: The virtual storage configuration created with PowerShell script must be reviewed by the software vendor (i.e., StarWind, Microsoft, VMware, etc.) for compliance and is fully supported and validated by either software vendor or MSP / ISV to satisfy and match the existing and future workload requirements and expectations of the specific production environment.


The operations specified in this section should be performed on each server.

Installing File Server Roles

Please follow the steps below if file shares configuration is required

Scale-Out File Server (Sofs) For Application Data

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

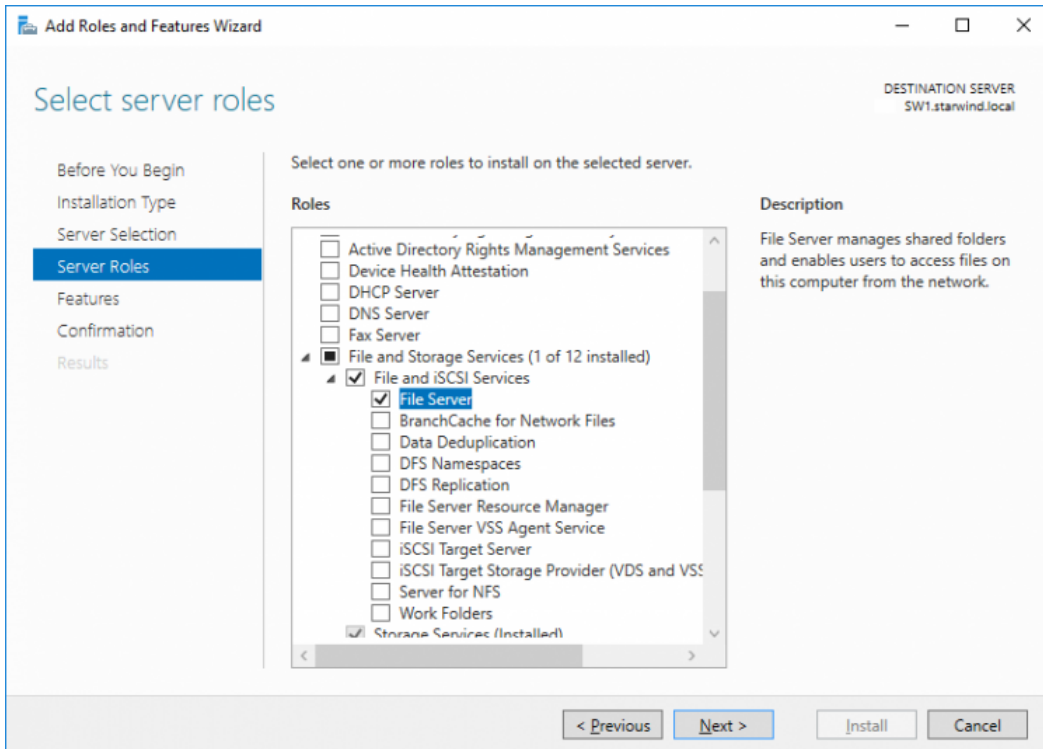


4. Restart the server after installation is completed and perform steps above on the each server.

File Server For General Use With Smb Share

1. Open Server Manager: Start -> Server Manager.
2. Select: Manage -> Add Roles and Features.


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



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

File Server For General Use With Nfs Share


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



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


Installing Starwind Vsan For Hyper-V

1. Download the StarWind setup executable file from the StarWind website: <https://www.starwind.com/registration-starwind-virtual-san>
2. Launch the downloaded setup file on the server to install StarWind Virtual SAN or one of its components. The Setup wizard will appear. Read and accept the License Agreement.



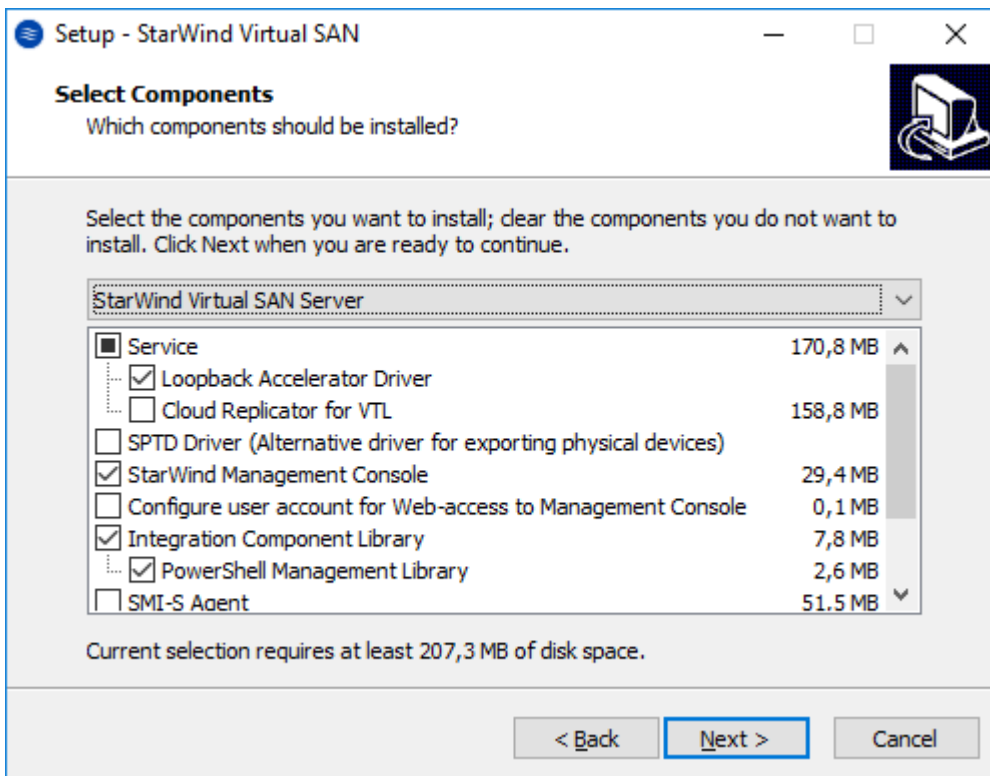
3. Carefully read the information about the new features and improvements. Red text indicates warnings for users that are updating the existing software installations.

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

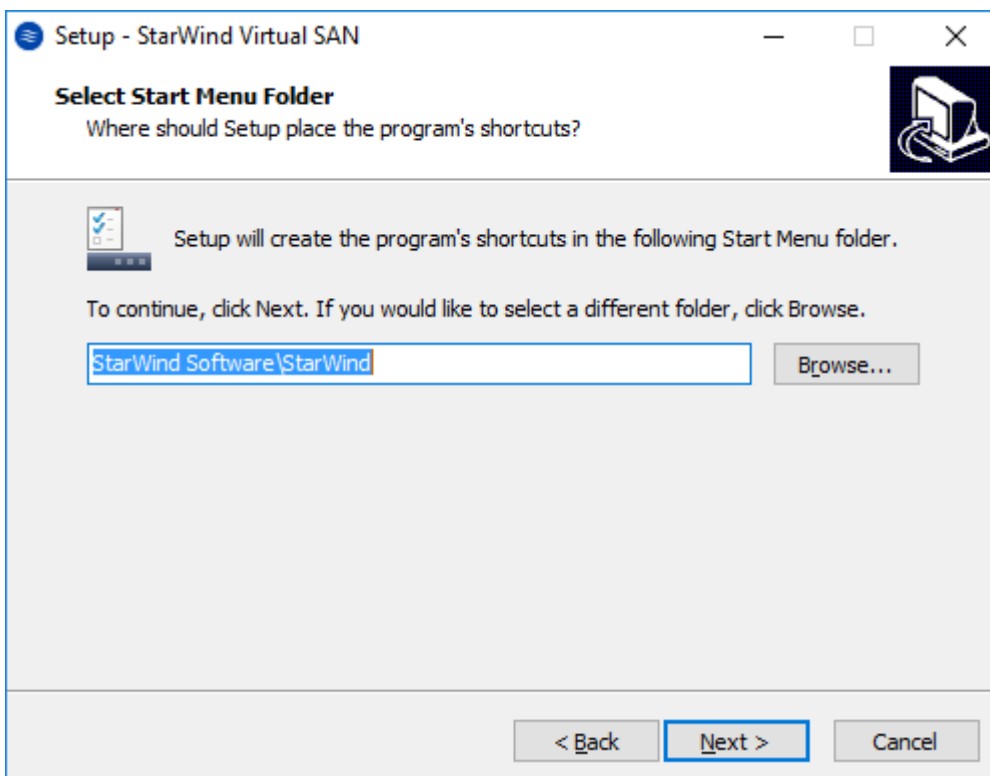


5. Select the following components for the minimum setup:

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



6. Specify Start Menu Folder.




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

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


- request time-limited fully functional evaluation key.
 - request FREE version key.
 - relect the previously purchased commercial license key.
9. Click on the Browse button to locate the license file.
 10. Review the licensing information.
 11. Verify the installation settings. Click on Back to make any changes or Install to proceed with installation.
 12. Enable the appropriate checkbox to launch StarWind Management Console right after the setup wizard is closed and click on Finish.
 13. Repeat the installation steps on the partner node.

Creating Starwind Devices


1. In the StarWind Management Console click to Add Device (advanced) button and open Add Device (advanced) Wizard.
2. Select Hard Disk Device as the type of device to be created.



3. Select Virtual Disk.




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




5. Select the Thick provisioned disk type and block size.

NOTE: Use 4096 sector size for targets, connected on Windows-based systems and 512 bytes sector size for targets, connected on Linux-based systems (ESXi/Xen/KVM).


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



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



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



9. Click Close to finish the device creation.


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

Select The Required Replication Mode


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

Synchronous “Two-Way” Replication

1. Right-click the recently created device and select Replication Manager from the shortcut menu.
2. Select the Add Replica button in the top menu.



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



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

Selecting The Failover Strategy

StarWind provides 2 options for configuring a failover strategy:

Heartbeat

The Heartbeat failover strategy allows avoiding the “split-brain” scenario when the HA cluster nodes are unable to synchronize but continue to accept write commands from the initiators independently. It can occur when all synchronization and heartbeat channels disconnect simultaneously, and the partner nodes do not respond to the node’s requests. As a result, StarWind service assumes the partner nodes to be offline and continues operations on a single-node mode using data written to it.

If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The device with the lowest priority will be marked as not synchronized and get subsequently blocked for the further read and write operations until the synchronization channel resumption. At the same time, the partner device on the

synchronized node flushes data from the cache to the disk to preserve data integrity in case the node goes down unexpectedly. It is recommended to assign more independent heartbeat channels during the replica creation to improve system stability and avoid the “split-brain” issue.

With the heartbeat failover strategy, the storage cluster will continue working with only one StarWind node available.

Node Majority

The Node Majority failover strategy ensures the synchronization connection without any additional heartbeat links. The failure-handling process occurs when the node has detected the absence of the connection with the partner.

The main requirement for keeping the node operational is an active connection with more than half of the HA device’s nodes. Calculation of the available partners is based on their “votes”.


In case of a two-node HA storage, all nodes will be disconnected if there is a problem on the node itself, or in communication between them. Therefore, the Node Majority failover strategy requires the addition of the third Witness node or file share (SMB) which participates in the nodes count for the majority, but neither contains data on it nor is involved in processing clients’ requests. In case an HA device is replicated between 3 nodes, no Witness node is required.

With Node Majority failover strategy, failure of only one node can be tolerated. If two nodes fail, the third node will also become unavailable to clients’ requests.

Please select the required option:

Heartbeat

1. Select Failover Strategy.



2. Select Create new Partner Device and click Next.

3. Select a partner device Location and click Next.

4. Select Synchronization Journal Strategy and click Next.

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

RAM-based (default) synchronization journal is placed in RAM. Synchronization with RAM journal provides good I/O performance in any scenario. Full synchronization could occur in the cases described in this KB:

<https://knowledgebase.starwindsoftware.com/explanation/reasons-why-full-synchronizati-on-may-start/>

Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes.


Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are recommended as the device performance is defined by the disk speed, where the journal is located. For example, it

can be placed on the OS boot volume.


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

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


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



5. Click Change Network Settings.



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



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

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

The successfully added device appears in StarWind Management Console.

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

NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:

[How to extend Image File or High Availability device](#)

Node Majority

There are two ways to configure Witness for 2-nodes StarWind HA device, created with Node Majority Failover Strategy: File Share (SMB) as Witness and additional server as Witness Node.

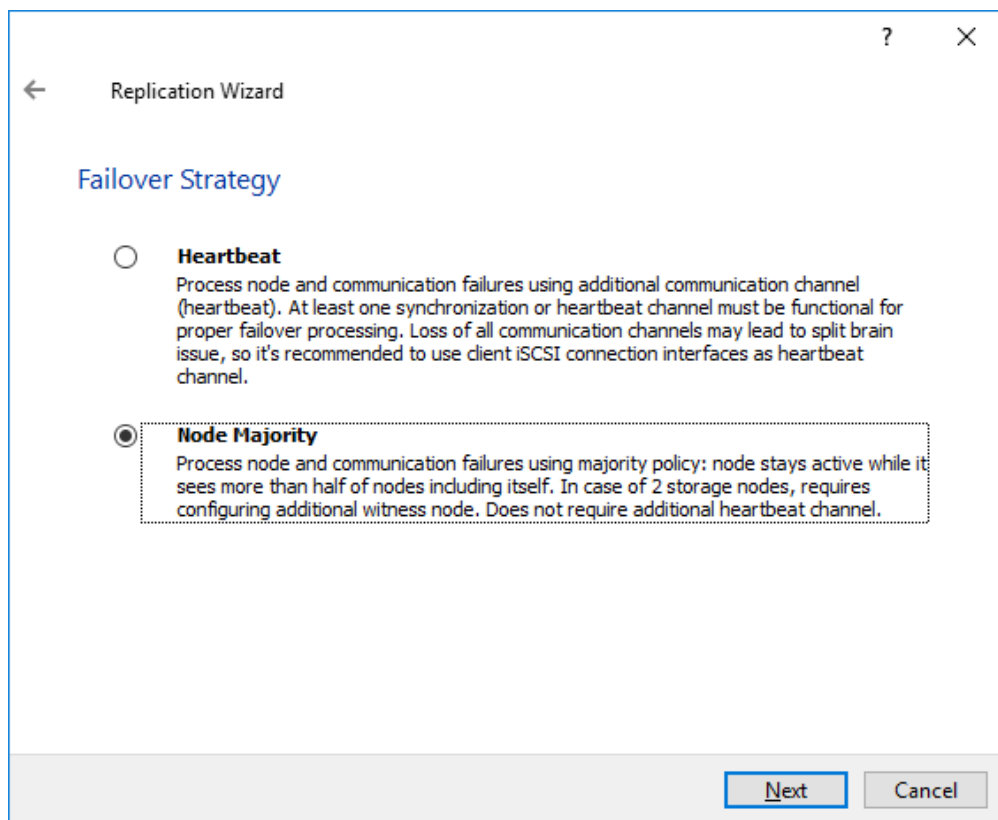
- Creating HA device with File SHare(SMB) as Witness:

SMB Witness is a file, located on SMB share, which can be accessed by both nodes and help them to eliminate the split-brain issue in case of synchronization connection interruption between the nodes. To set up the SMB file share as a Witness for 2-nodes HA device with Node Majority Failover Strategy, perform the actions, described on this page:

<https://www.starwindsoftware.com/help/ConfiguringFileShareSMBasWitness.html>

- Creating HA device with Witness Node:

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



2. Choose Create new Partner Device and click Next.

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

Click Next. Select Synchronization Journal strategy and location and click Next.

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

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

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

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


8. The added devices will appear in StarWind Management Console. Repeat the steps above to create other virtual disks if necessary.

Adding Witness Node

Witness node can be configured on a separate host or as a virtual machine in a cloud. It requires StarWind Virtual SAN service installed on it.


NOTE: Since the device created in this guide is replicated between 2 active nodes with the Node Majority failover strategy, a Witness node must be added to it.

1. Open StarWind Management Console, right-click on the Servers field and press the Add Server button. Add a new StarWind Server which will be used as the Witness node and click OK.




2. Right-click on the HA device with the configured Node Majority failover policy and select Replication Manager and press the Add Replica button.

3. Select Witness Node.



4. Specify the Witness node Host Name or IP address. The default Port Number is 3261.



5. In Partner Device Setup, specify the Witness device Location. Optionally, modify the target name by clicking the appropriate button.

6. In Network Options for Replication, select the synchronization channel with the Witness node by clicking the Change Network Settings button.

7. Specify the interface for Synchronization and Heartbeat and click OK.

8. Click Create Replica and then close the wizard.

9. Repeat the steps above to create other virtual disks if necessary.

NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:

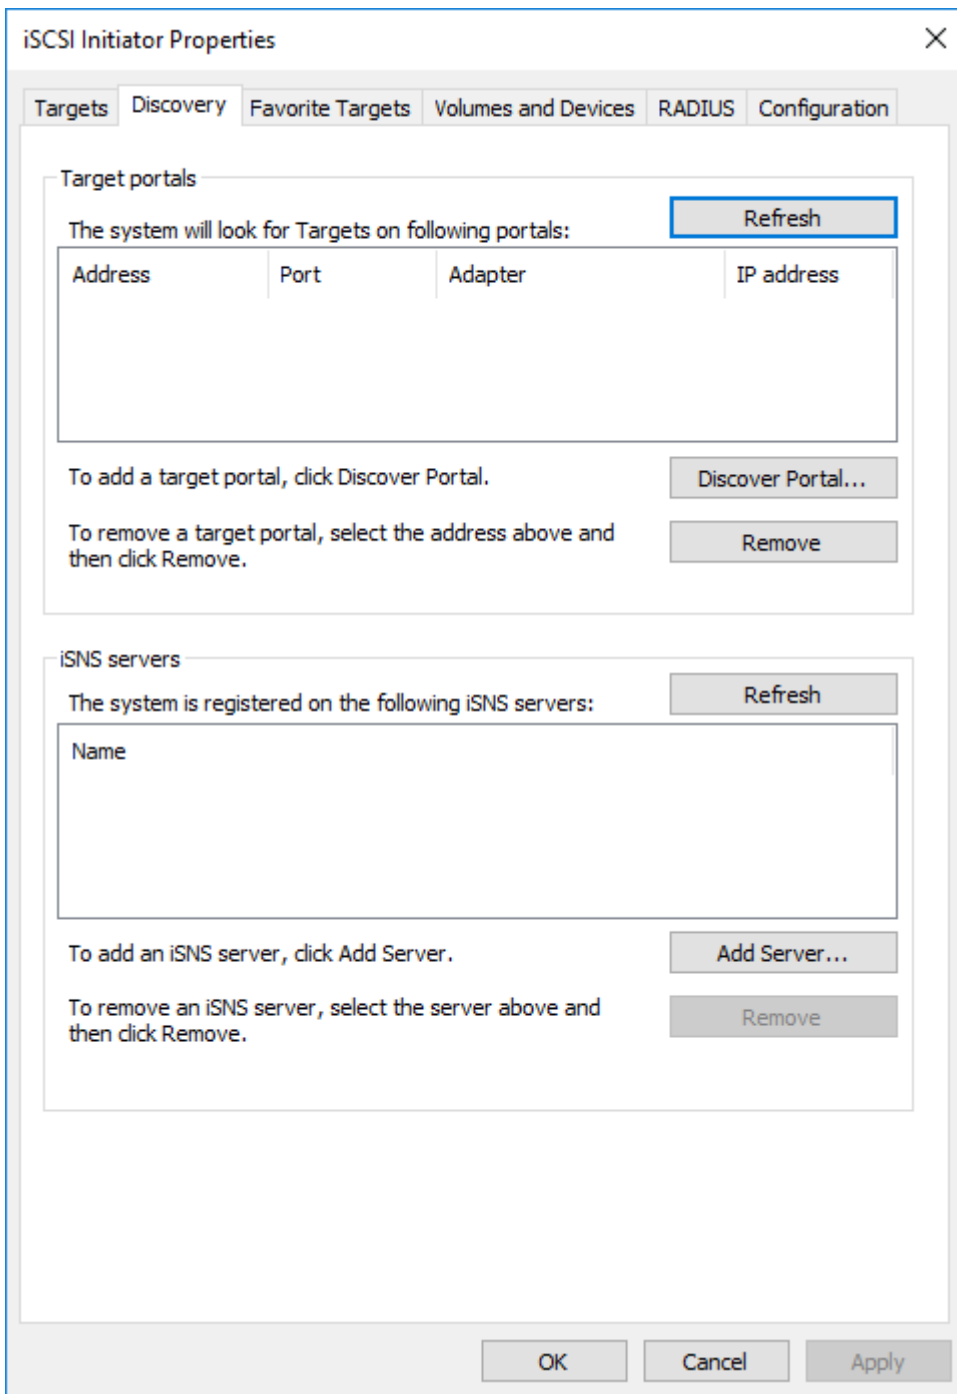
<https://knowledgebase.starwindsoftware.com/maintenance/how-to-extend-image-file-or-high-availability-device/>

Provisioning Starwind Ha Storage To Windows Server Host

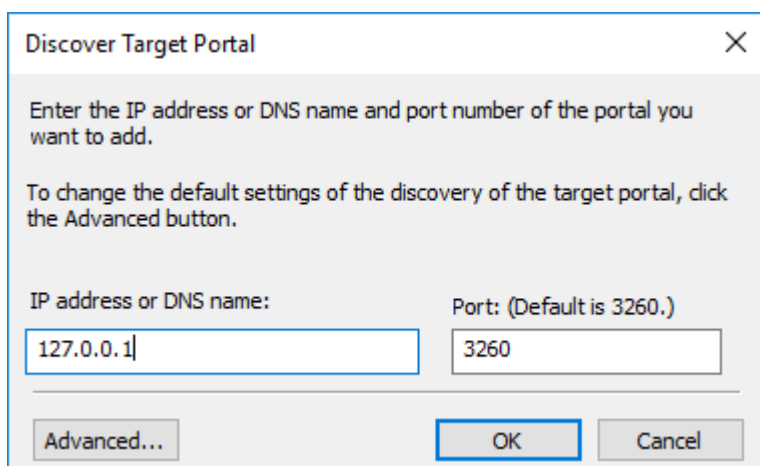
1. Launch Microsoft iSCSI Initiator: Start -> Windows Administrative Tools -> iSCSI Initiator. Alternatively, launch it using the command below in the command line interface:

```
iscsicpl
```

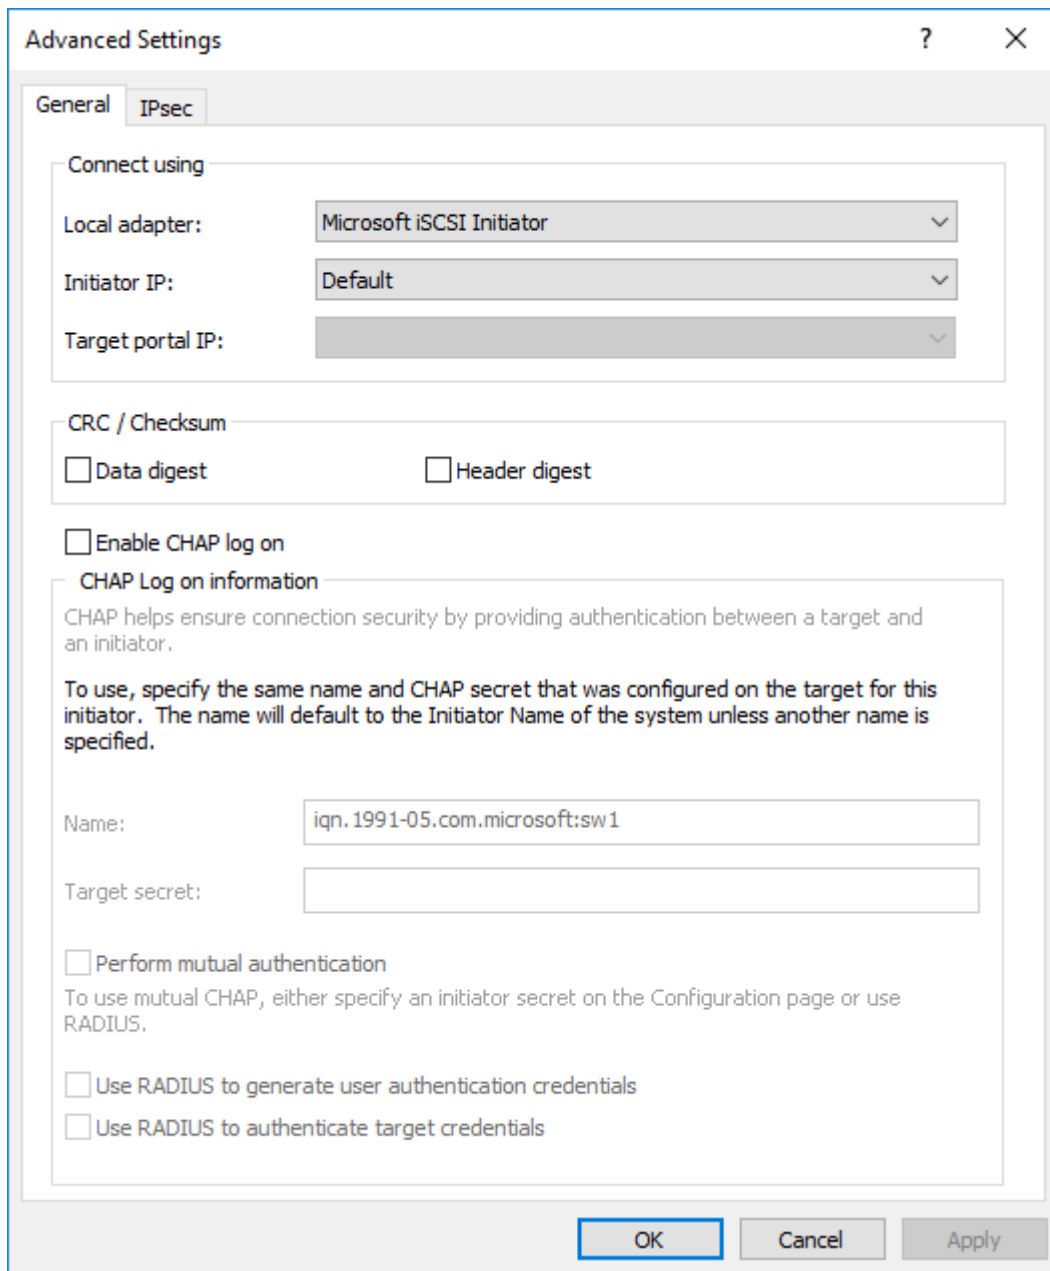
2. Navigate to the Discovery tab.



3. Click the Discover Portal button. The Discover Target Portal dialog appears. Type 127.0.0.1.




4. Click the Advanced button. Select Microsoft iSCSI Initiator as a Local adapter and select Initiator IP (leave default for 127.0.0.1). 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.



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

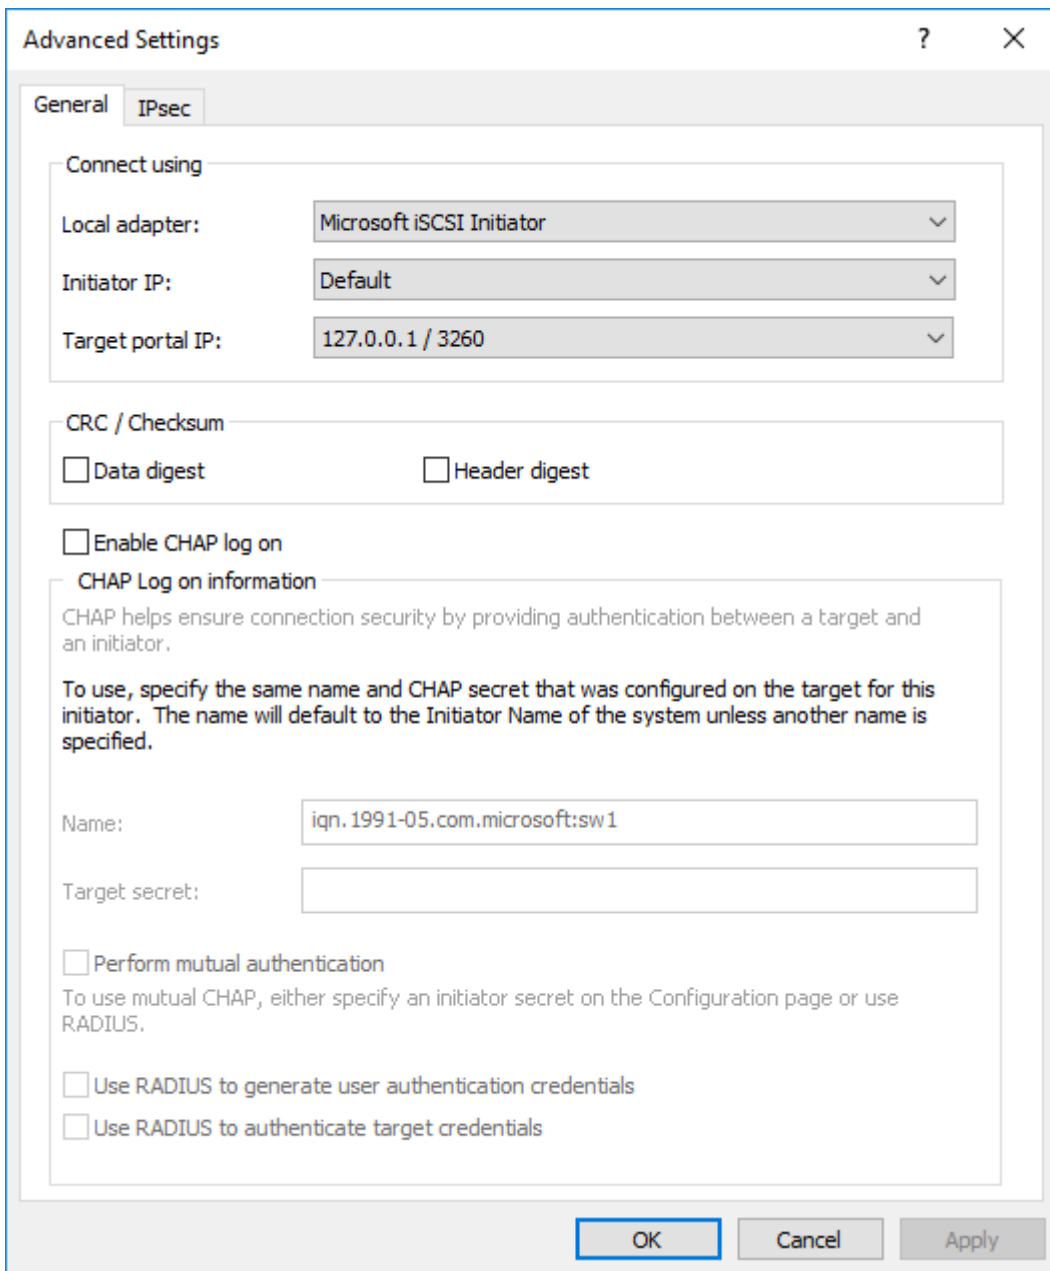
for any restrictions.



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 Target portal IP, select 127.0.0.1. Confirm the actions.



NOTE: It is recommended to connect the Witness device only by loopback (127.0.0.1) address. Do not connect the target to the Witness device from the partner StarWind node.

5. Select the CSV1 target discovered from the local server and click Connect.
6. Enable checkboxes as shown in the image below. Click Advanced.




7. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Target portal IP, select 127.0.0.1. Confirm the actions.

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

9. Repeat the step 6.

10. 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 local and 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 except for Witness 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. For the Witness target, set the load balance policy to Failover Only.

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


NOTE: In case the Failover Only MPIO policy is used, make sure to check that the local path (127.0.0.1) is set to Active, while the partner connection is set to Standby.

Connecting Disks to Servers

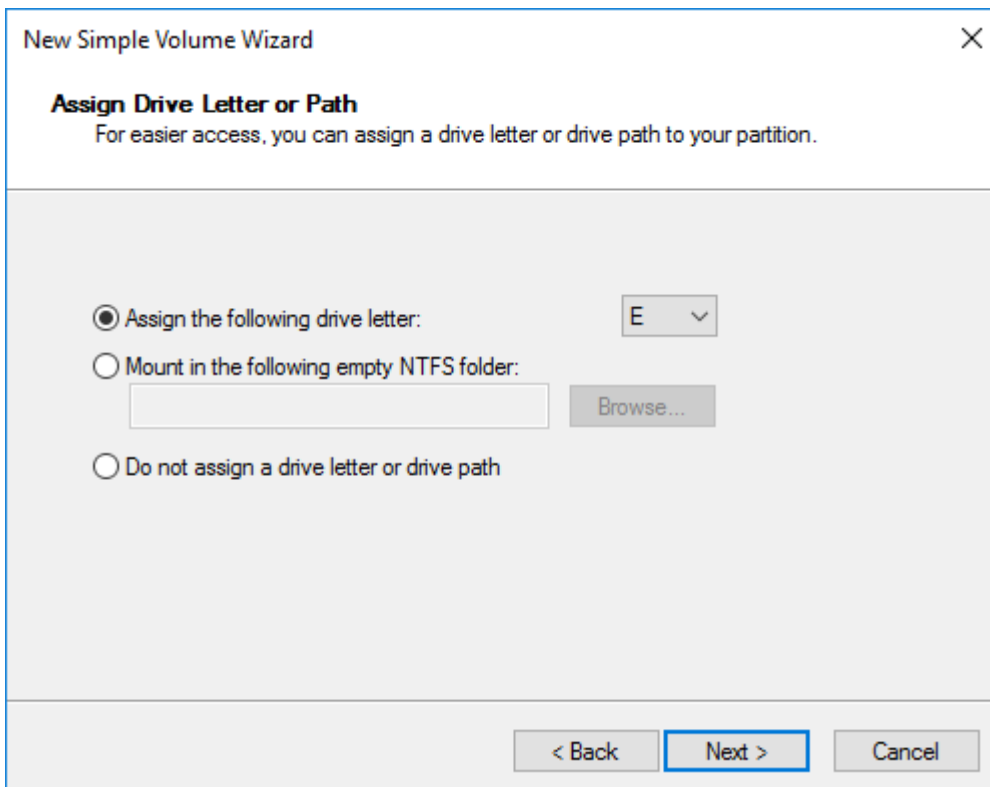
1. Open the Disk Management snap-in. The StarWind disks will appear as unallocated and offline.



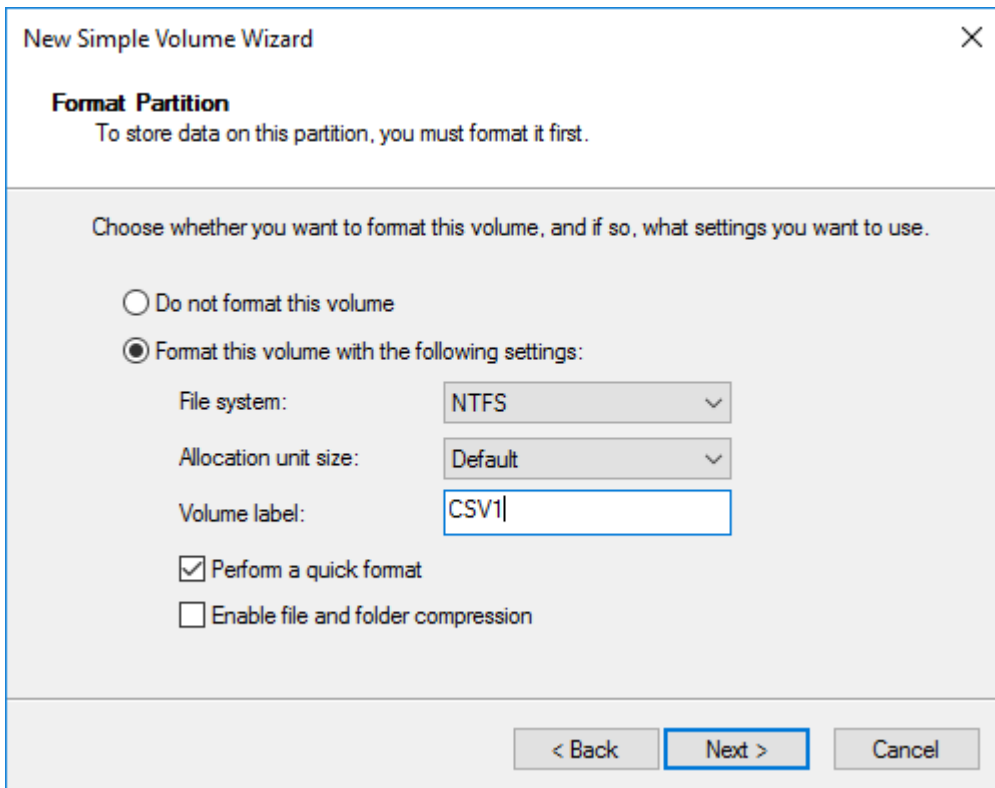
2. Bring the disks online by right-clicking on them and selecting the Online menu option.
3. Select the CSV disk (check the disk size to be sure) and right-click on it to initialize.
4. By default, the system will offer to initialize all non-initialized disks. Use the Select Disks area to choose the 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

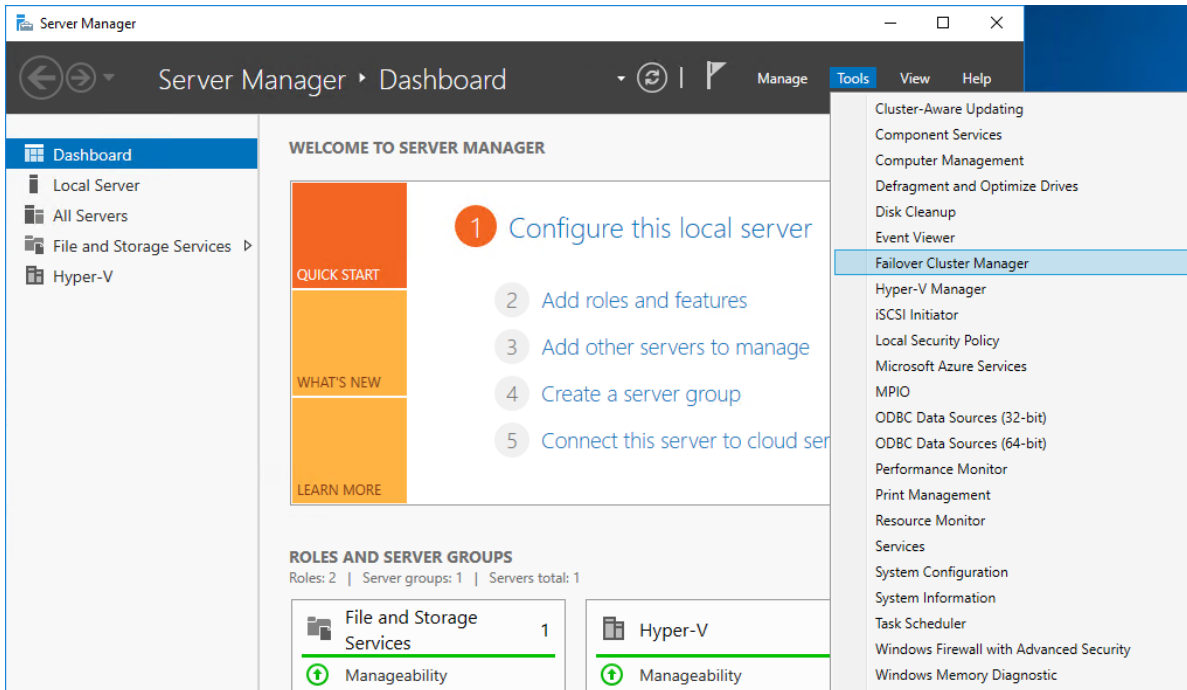
NOTE: To avoid issues during the cluster validation configuration, it is recommended to install the latest Microsoft updates on each node.

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.



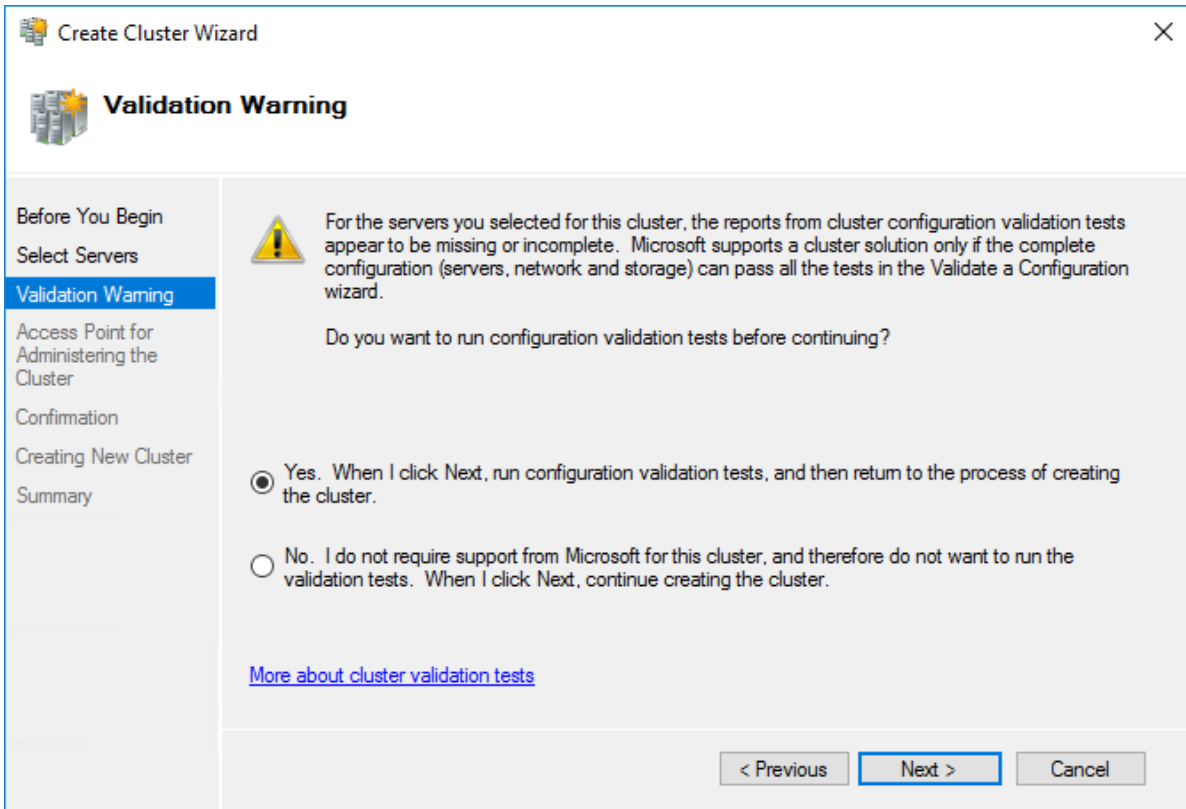
2. Click the Create Cluster link in the Actions section of Failover Cluster Manager.



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

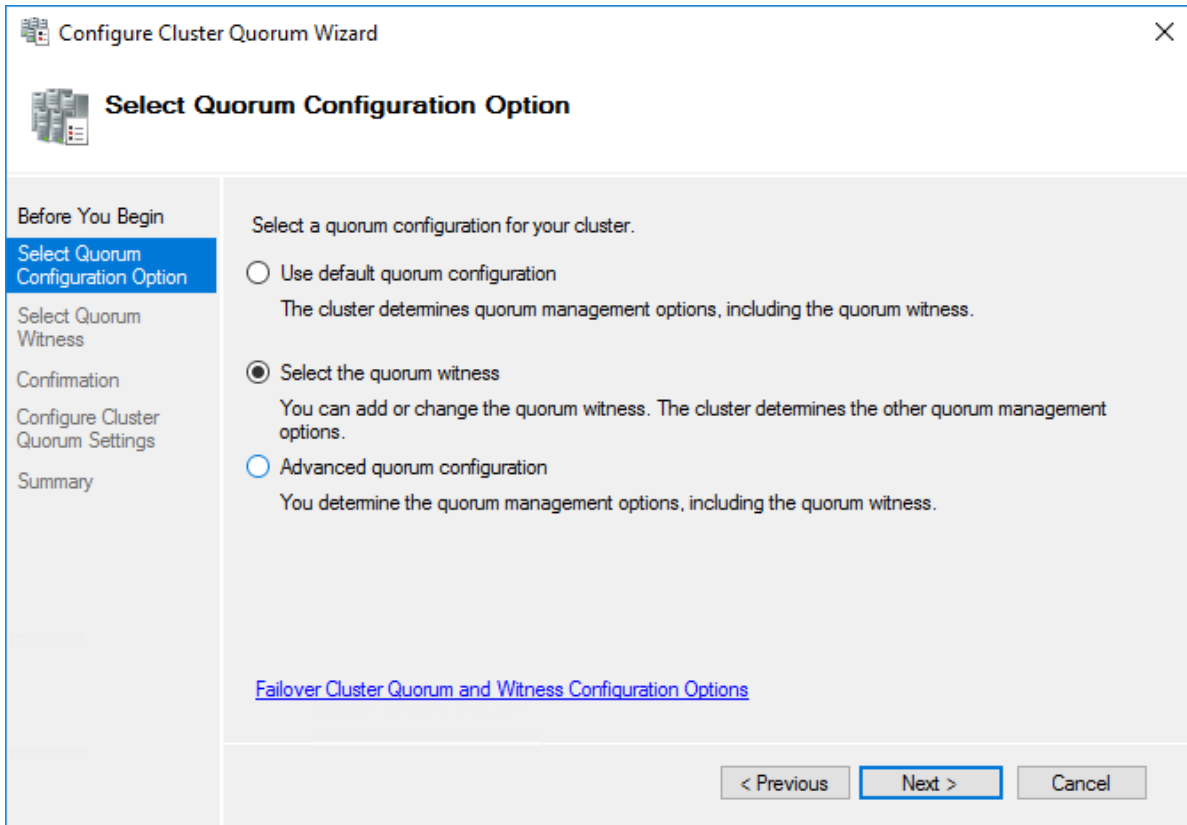
1. In Failover Cluster Manager, navigate to Cluster -> Storage -> Disks. Click Add Disk in the Actions panel, choose StarWind disks from the list and confirm the selection.



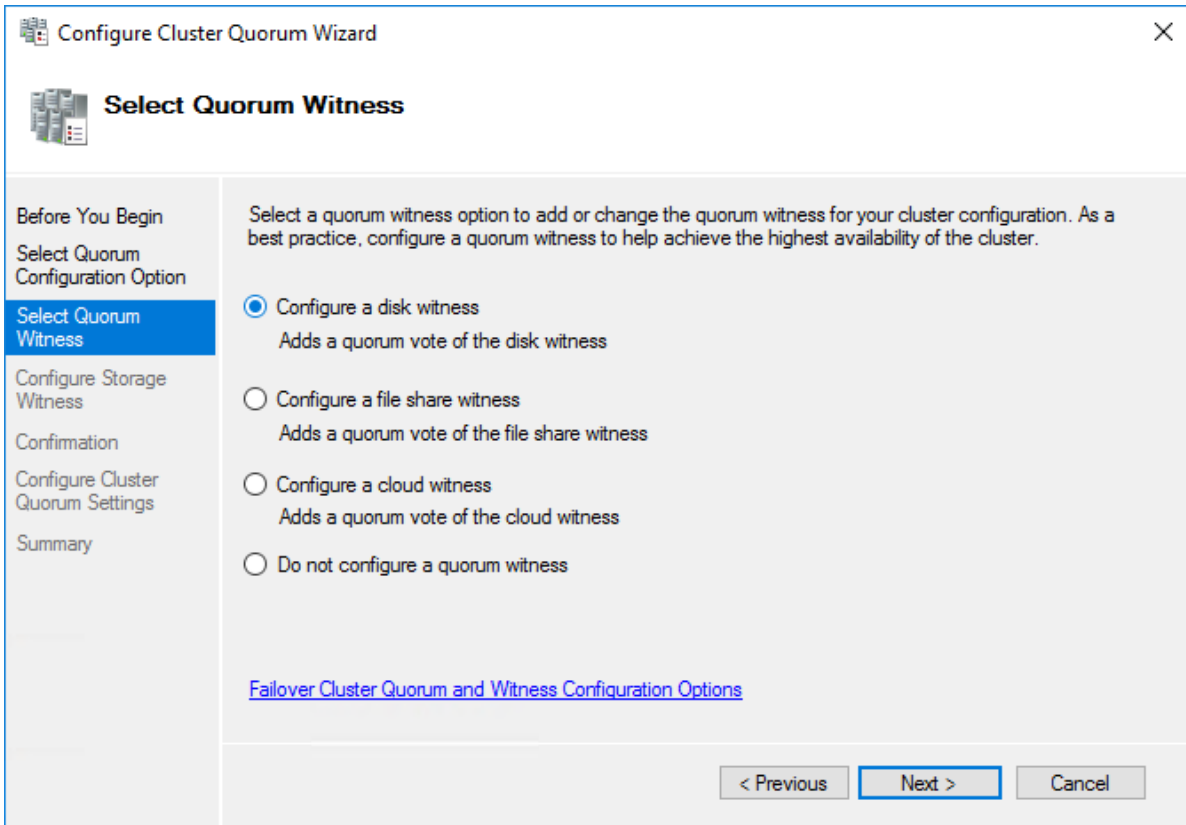
2. To configure the cluster witness disk, right-click on Cluster and proceed to More Actions -> Configure Cluster Quorum Settings.




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

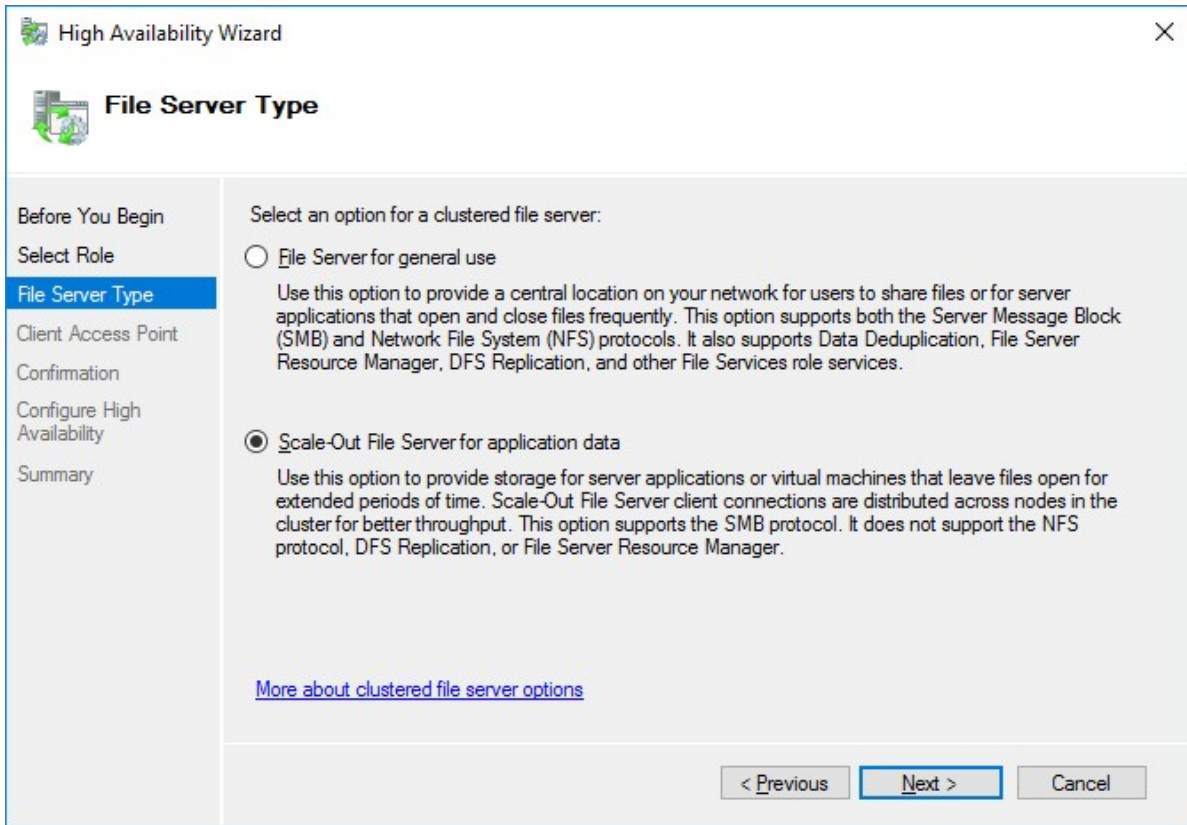
1. To configure the Scale-Out File Server Role, open Failover Cluster Manager.
2. Right-click the cluster name, then click Configure Role and click Next to continue.




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.

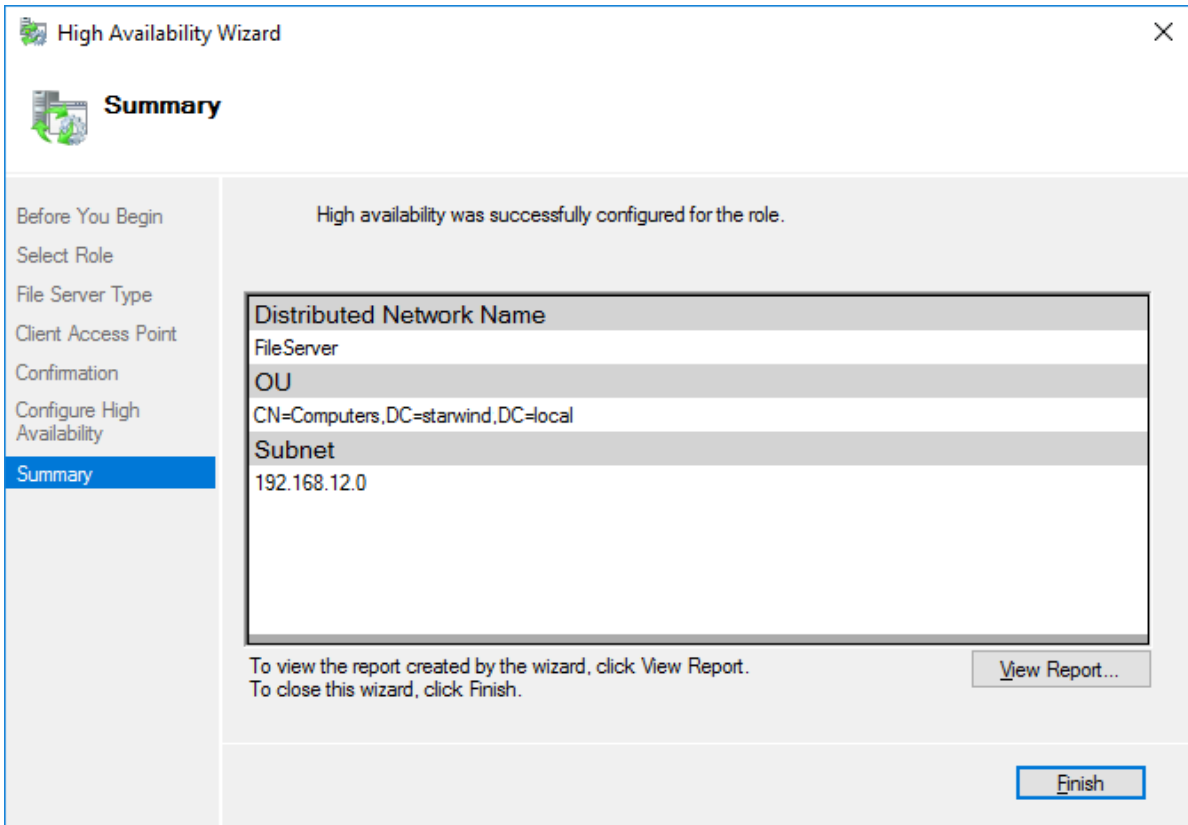


Click Next to continue.

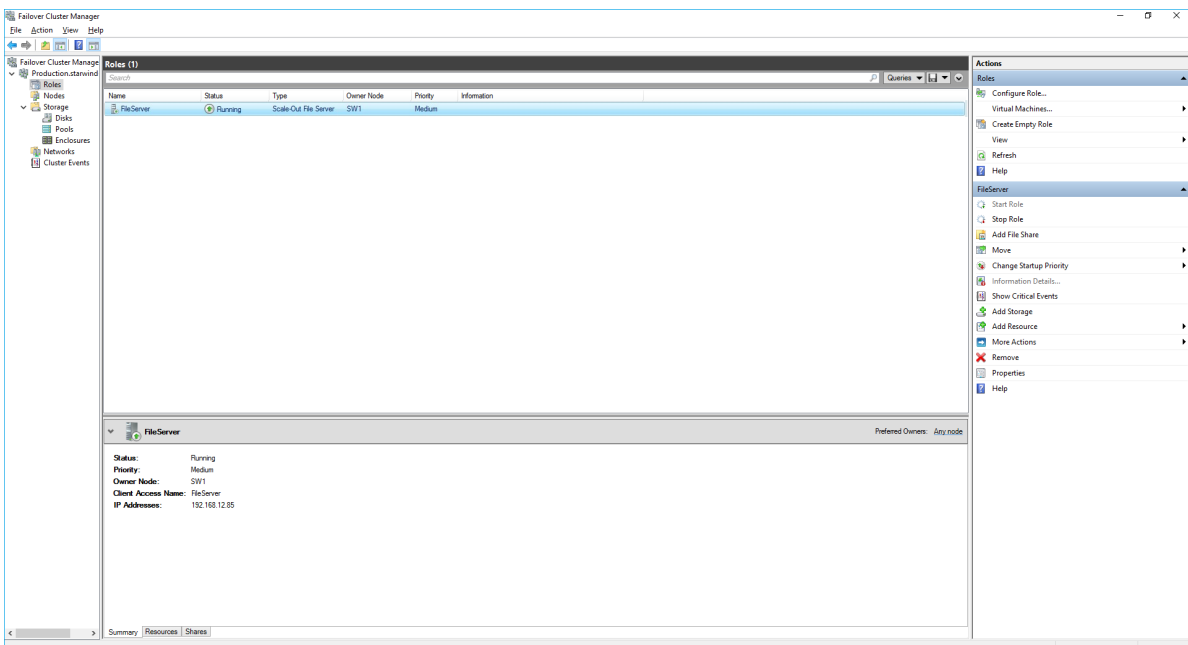
6. Check whether the specified information is correct. Click Next to continue or Previous to change the settings.




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



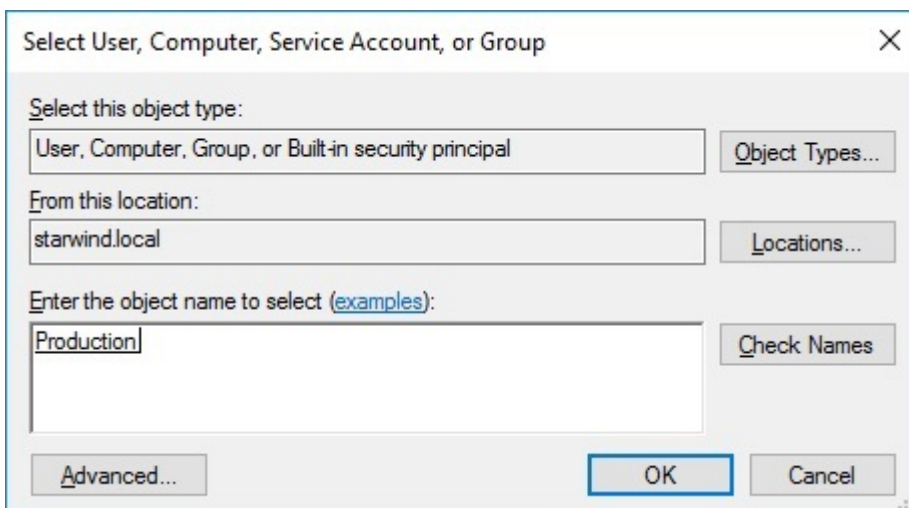
8. The newly created role should now look like the screenshot below.




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:



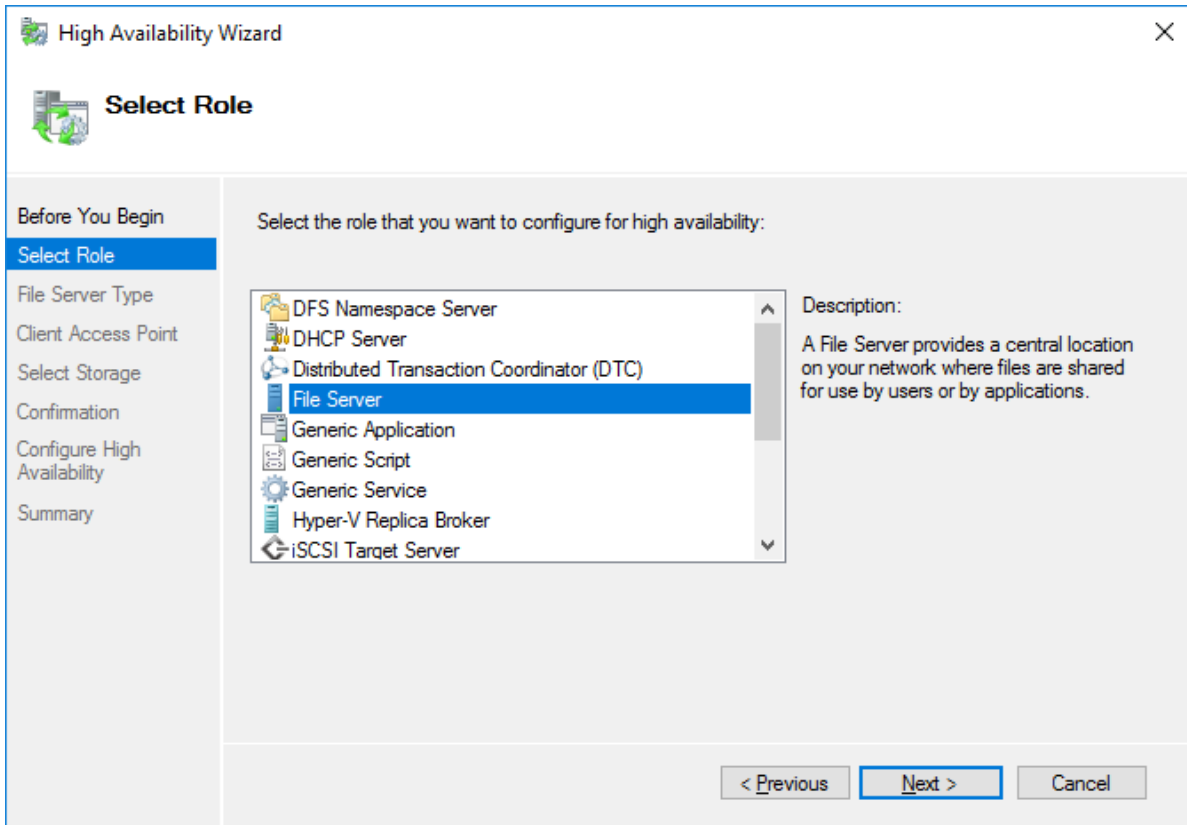
Configuring The File Server For General Use Role

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

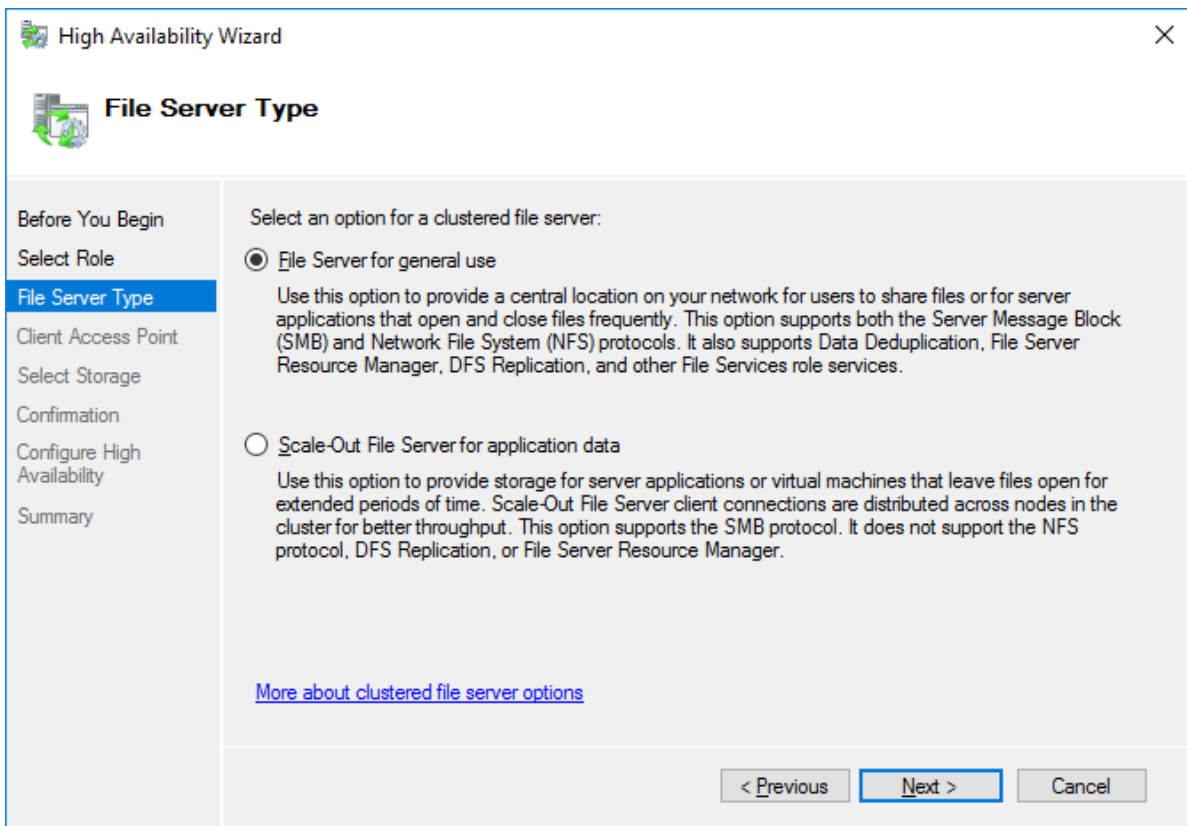
1. To configure the File Server for General Use role, open Failover Cluster Manager.
2. Right-click on the cluster name, then click Configure Role and click Next to continue.




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



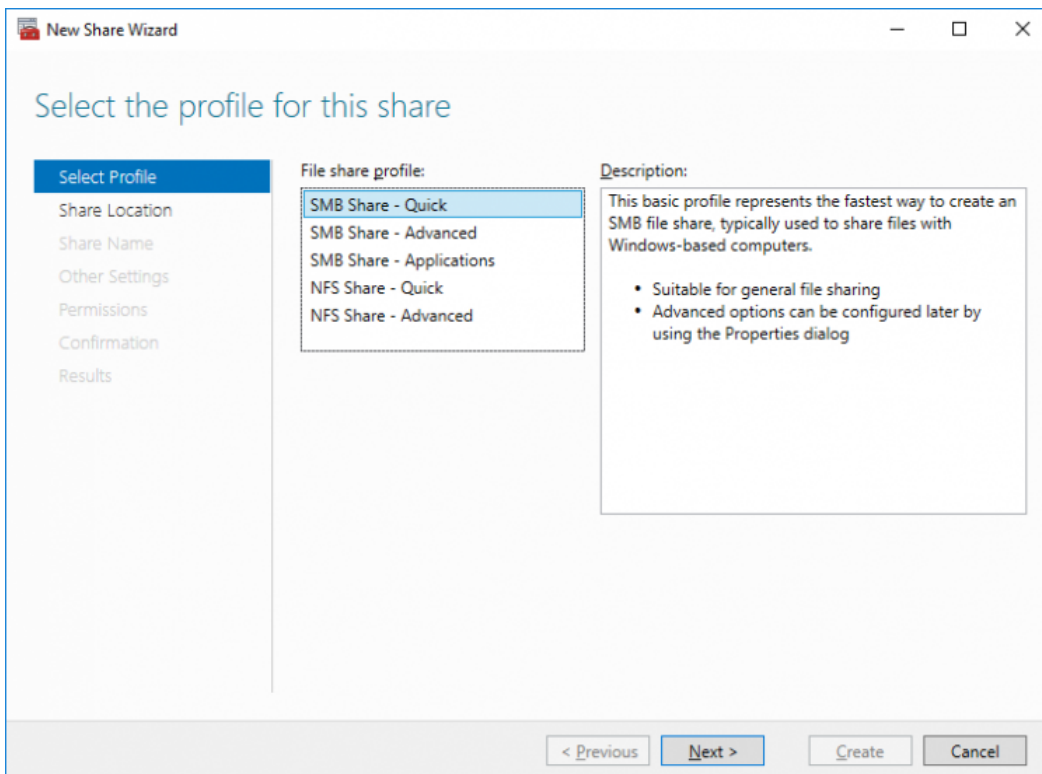
- click OK on all opened windows to confirm the changes

- 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

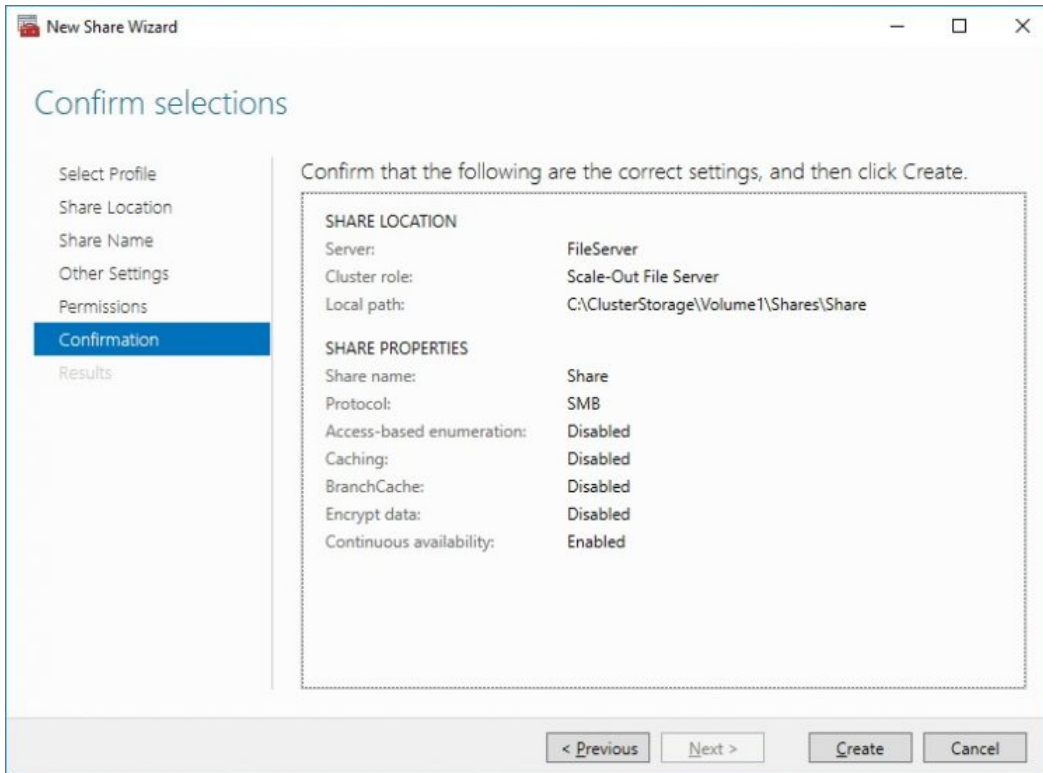
continue.




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.

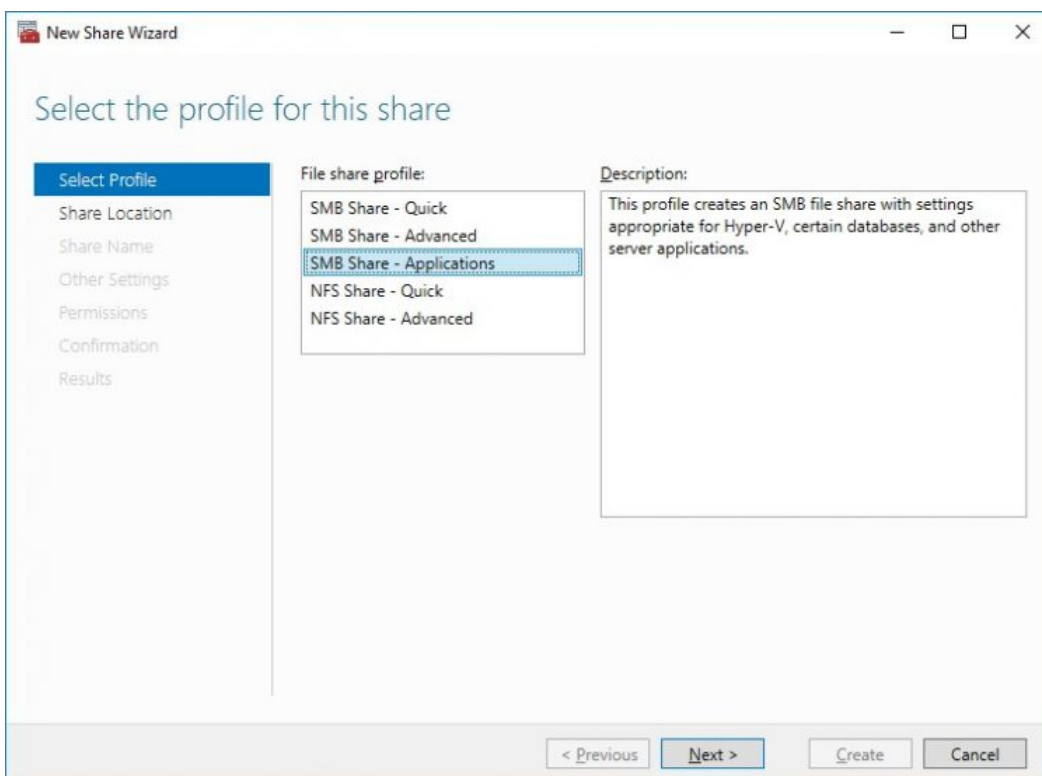
13. Choose the File Share role, select the Shares tab, right-click the created file share, and select Properties.




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.




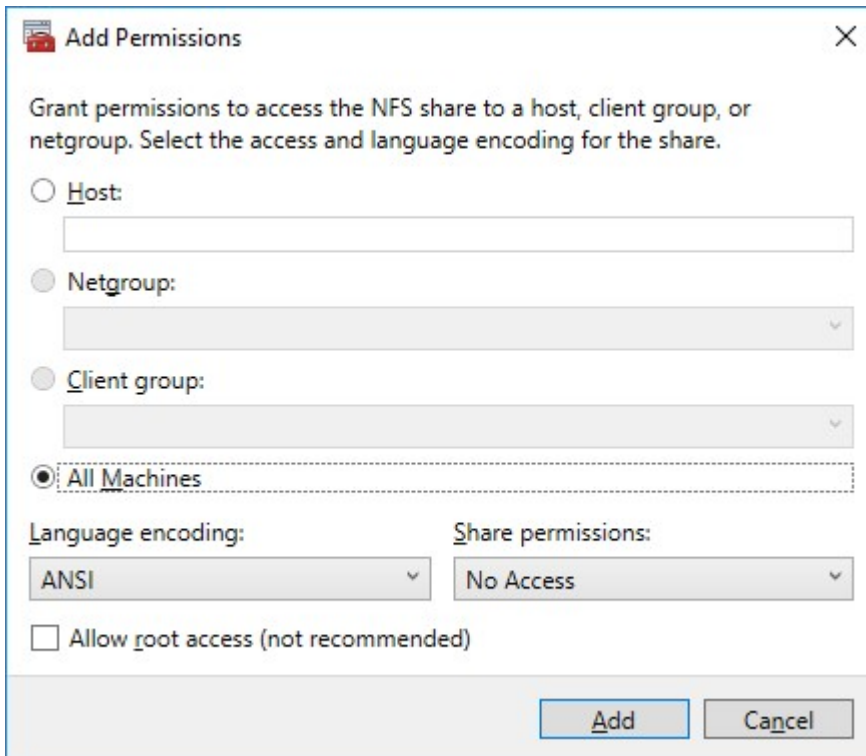
7. Specify the Authentication. Click Next and confirm the message in pop-up window to

continue.

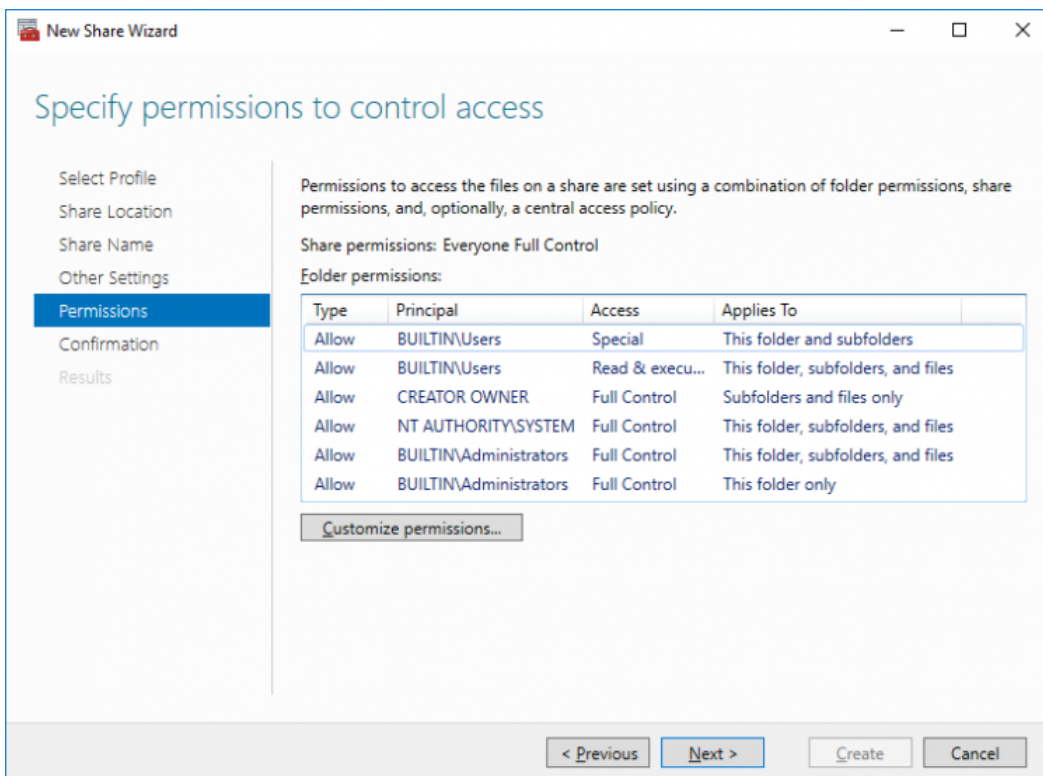


8. Click Add and specify Share Permissions.







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



Conclusion

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

Contacts

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

Customer Support Portal: <https://www.starwind.com/support>

Support Forum: <https://www.starwind.com/forums>

Sales: sales@starwind.com

General Information: info@starwind.com



StarWind Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA
www.starwind.com ©2024, StarWind Software Inc. All rights reserved.