

# StarWind Virtual SAN: Stretched Cluster Configuration Guide for Microsoft Windows Server [Hyper-V], StarWind Deployed as a Windows Application using GUI

2026

TECHNICAL PAPERS



## Trademarks

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

## Changes

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

## Technical Support and Services

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

## About StarWind

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

## Copyright ©2009-2018 StarWind Software Inc.

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

## Annotation

### Relevant Products

StarWind Virtual SAN (VSAN)

### Purpose

The guide provides a comprehensive step-by-step process of configuring a stretched cluster with StarWind Virtual SAN Windows-based application, ensuring a highly available, fault-tolerant, and efficient virtualized storage solution that integrates seamlessly with Microsoft Windows Server environments.

### Audience

This guide is intended for system administrators, IT professionals, and infrastructure architects planning to implement or manage stretched clusters in virtualized environments using StarWind VSAN and Microsoft Windows Server.

### Expected Result

The end result of following this guide was a successfully set up of a stretched cluster configuration with StarWind Virtual SAN, achieving a resilient and high-performance storage solution.

## 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>

## Infrastructure Design

### Prerequisites For Stretched Cluster

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:

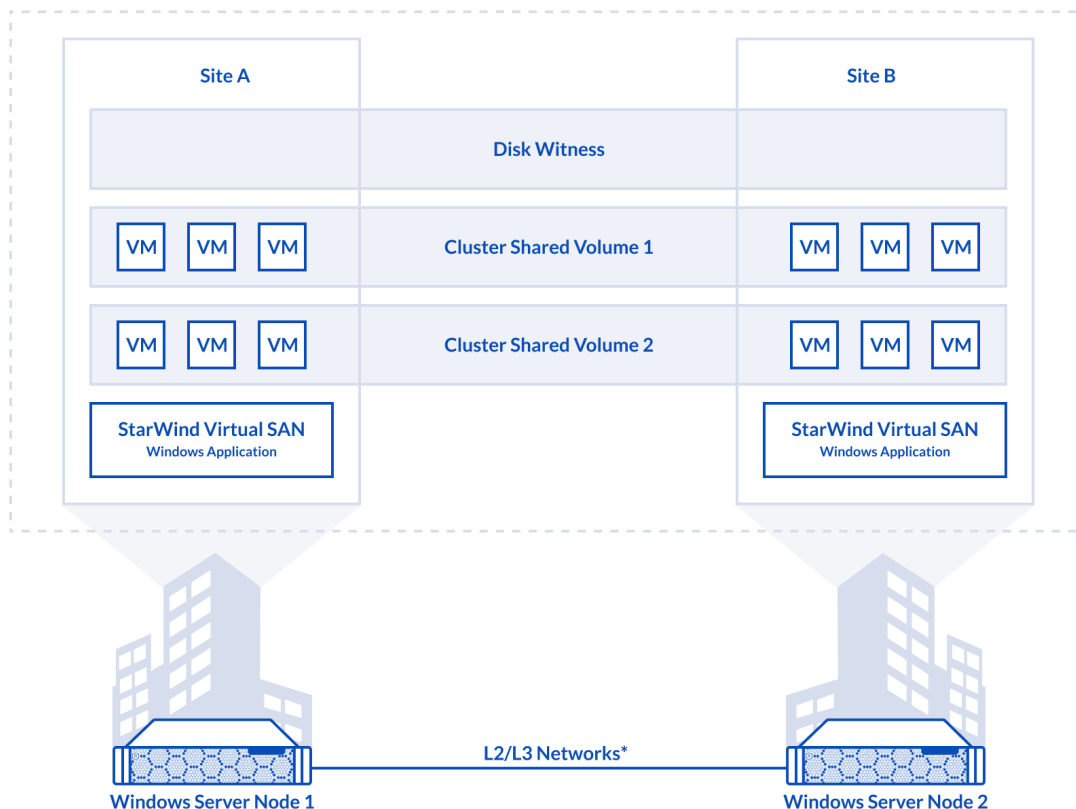
<https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/>

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 diagram below illustrates the connection scheme of the StarWind stretched cluster configuration described in this guide.



Make sure that the prerequisites for deploying StarWind stretched Microsoft Windows cluster are met:

- Design L2/L3 multisite network according to the appropriate StarWind failover strategy
- Each iSCSI and Synchronization network channel throughput should be at least 1Gbps
- Up to 10 ms Round-trip time (RTT) network latency between locations
- Deployed on-premises Active Directory structure and DNS
- Windows Server is installed on the server that is going to be clustered

## Heartbeat Failover Strategy For Stretched Cluster



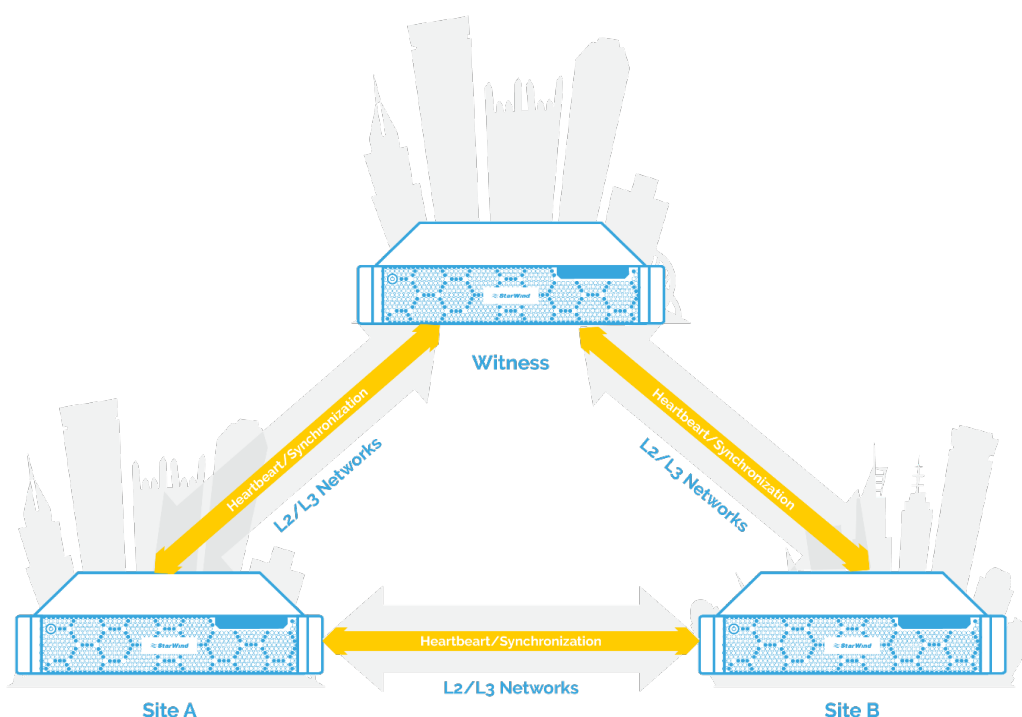
Heartbeat is a technology that allows avoiding the so-called “split-brain” scenario when the HA cluster nodes are unable to synchronize but continue to accept write commands from the initiators independently. It can occur when all synchronization and heartbeat channels disconnect simultaneously, and the other 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 in a single-node mode using the data written to it.

If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The services mark the device with the lowest priority as not-synchronized one. Subsequently it gets blocked for further read and write operations until the synchronization channel resumption. Then, 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 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.

#### Heartbeat Failover Strategy Network Design

- Management / Heartbeat - 100Mbps network or higher.
- iSCSI / Heartbeat - 1Gbps network or higher. The 10Gbps or higher bandwidth link is highly recommended.
- Synchronization - 1Gbps network or higher. The 10Gbps or higher bandwidth link is highly recommended.

## Node Majority Failover Strategy For Stretched Cluster



This strategy ensures synchronization connection without any additional heartbeat links. The failure-handling process occurs when the node has detected the absence of connection with the partner. The main requirement for keeping the node operational is an active connection with more than a half of the HA device's nodes. Calculation of the available partners bases on their "votes". In case of a two-node HA storage, all nodes disconnect if there is a problem with the node itself, or with communication within the cluster. Therefore, the Node Majority failover strategy does not work in case if only two synchronous nodes are available. To apply this strategy, the third entity is required. It can be a Witness node which participates in the nodes count for the majority, but neither contains data nor processes clients' requests. Node Majority failover strategy allows tolerating failure of only one node. If two nodes fail, the third one will also become unavailable to clients' requests. If replicated between 2 nodes, the Witness node requires additional configuration for an HA device that uses Node Majority failover strategy. Replication of an HA device among 3 nodes requires no Witness nodes.

**Node Majority Failover Strategy Network Design**

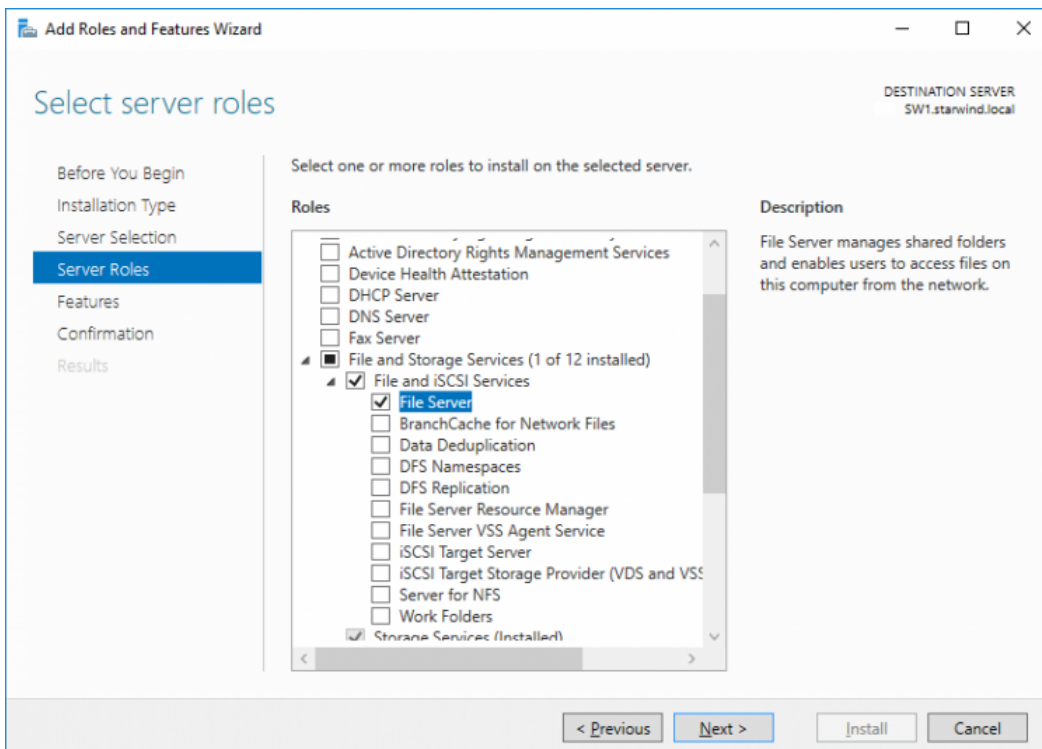
- Management / Heartbeat /Synchronization - 1Gbps network or higher. The 10Gbps or higher bandwidth link is highly recommended.

## 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

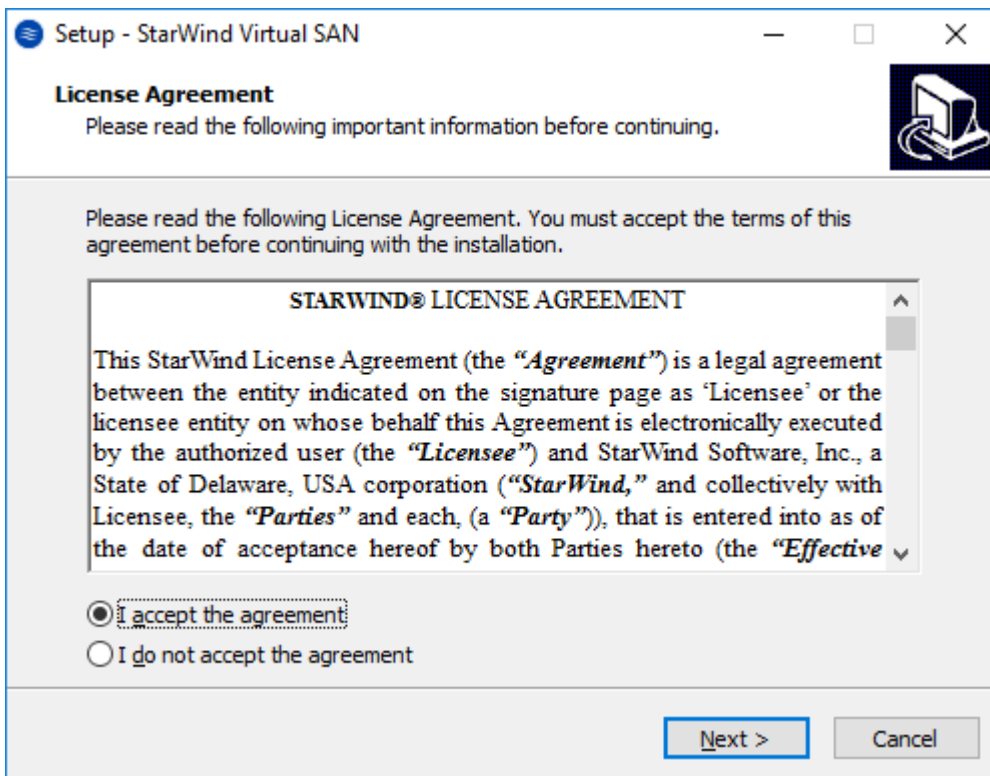
1. Open Server Manager: Start -> Server Manager.
2. Select: Manage -> Add Roles and Features.
3. Follow the installation wizard steps to install the roles selected in the screenshot below:



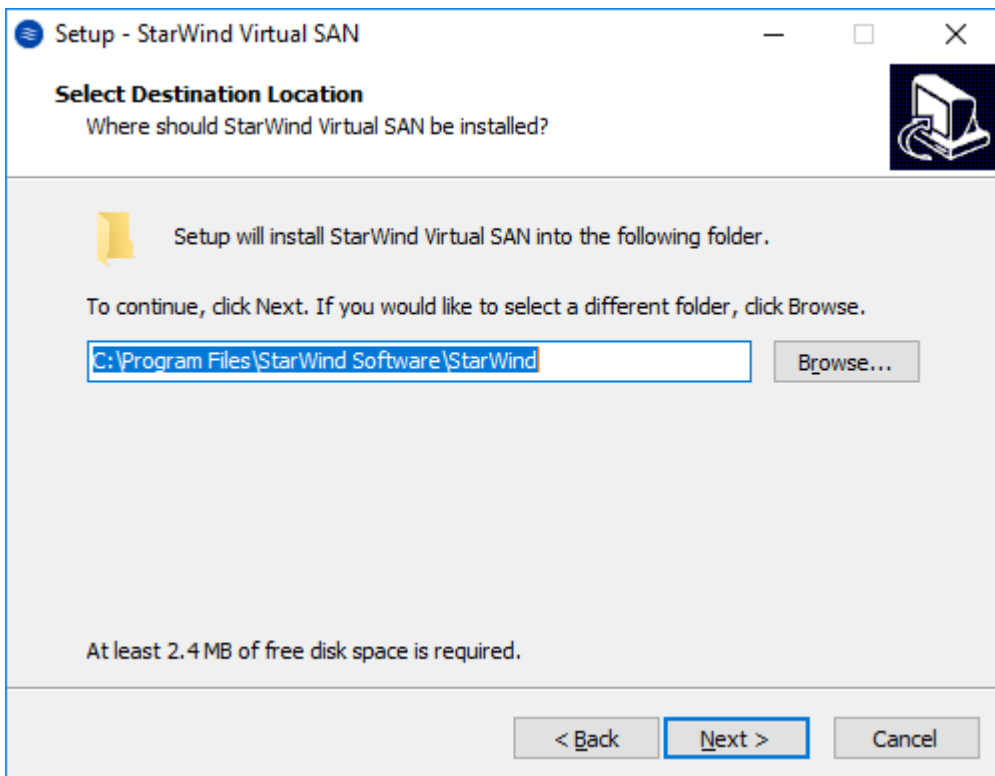
4. Restart the server after installation is completed and perform steps above on each server.

## 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.
  - select 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

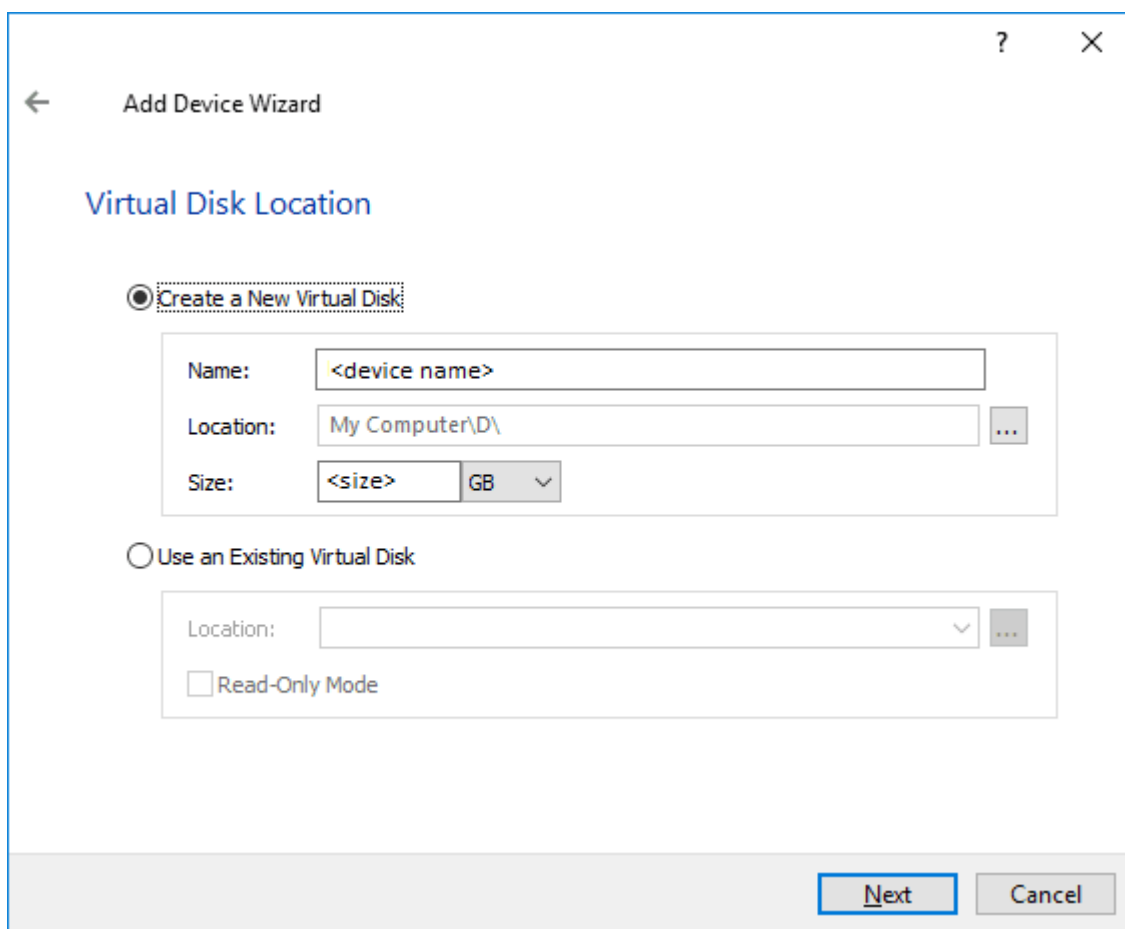
1. In the StarWind Management Console click to Add Device (advanced) button and open Add Device (advanced) Wizard.
2. Select Hard Disk Device as the type of device to be created.



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

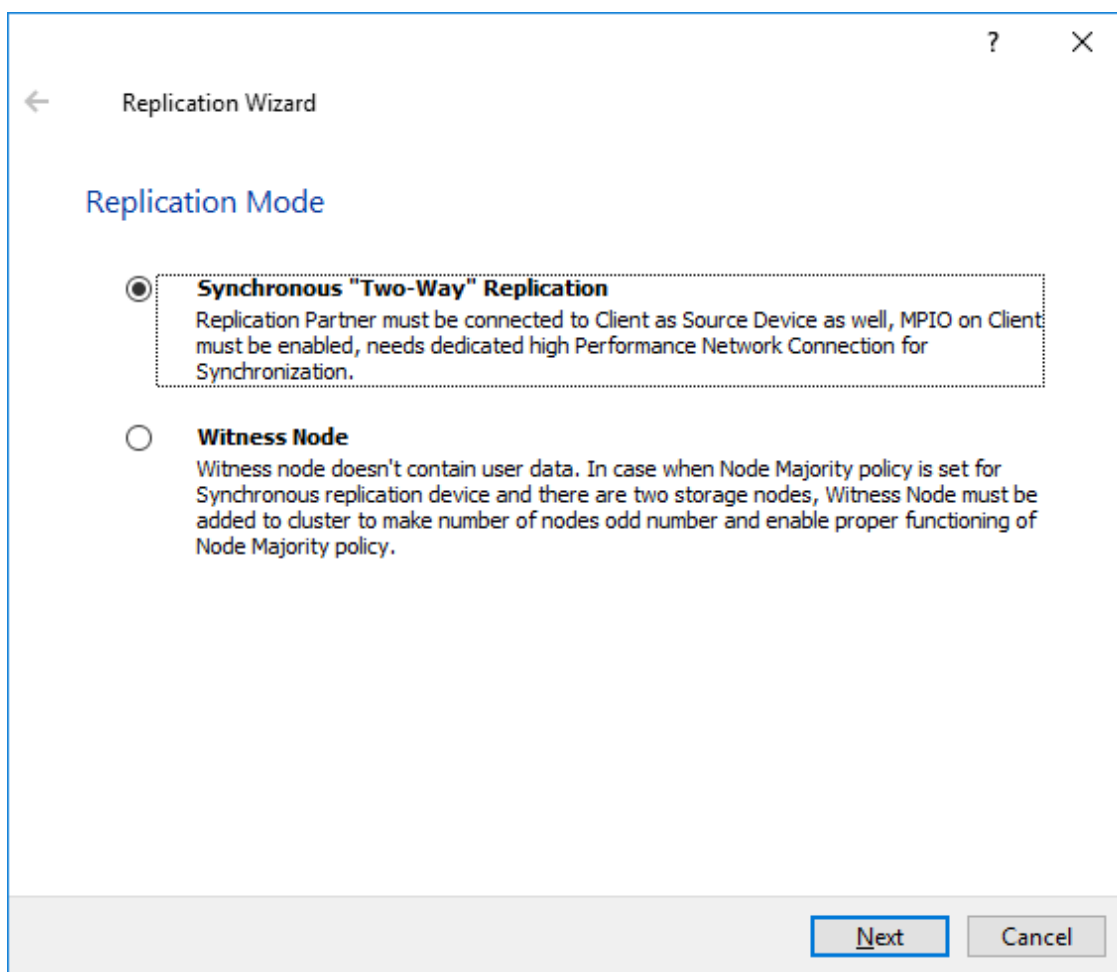
The replication can be configured using Synchronous “Two-Way” Replication mode: Synchronous or active-active replication ensures real-time synchronization and load balancing of data between two or three cluster nodes. Such a configuration tolerates the failure of two out of three storage nodes and enables 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.

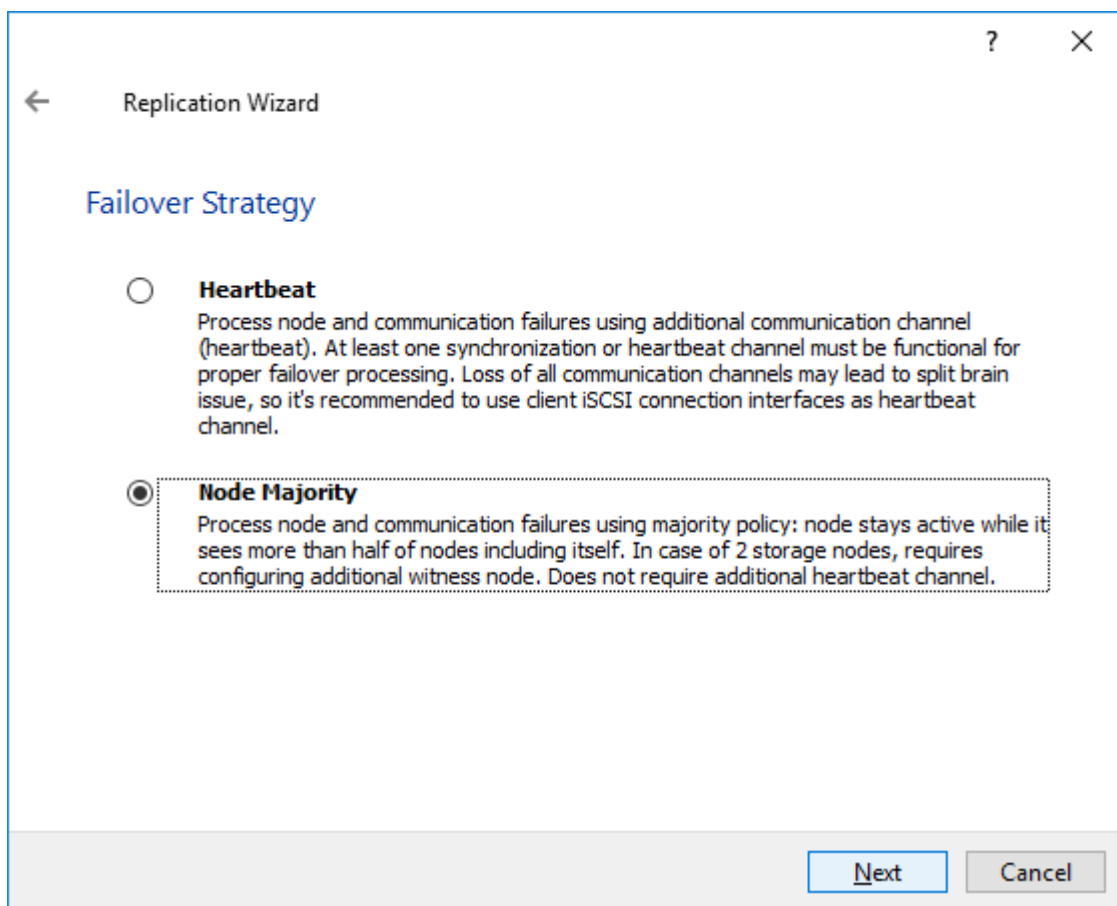
- Creating HA device with File SHare(SMB) as Witness:

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:

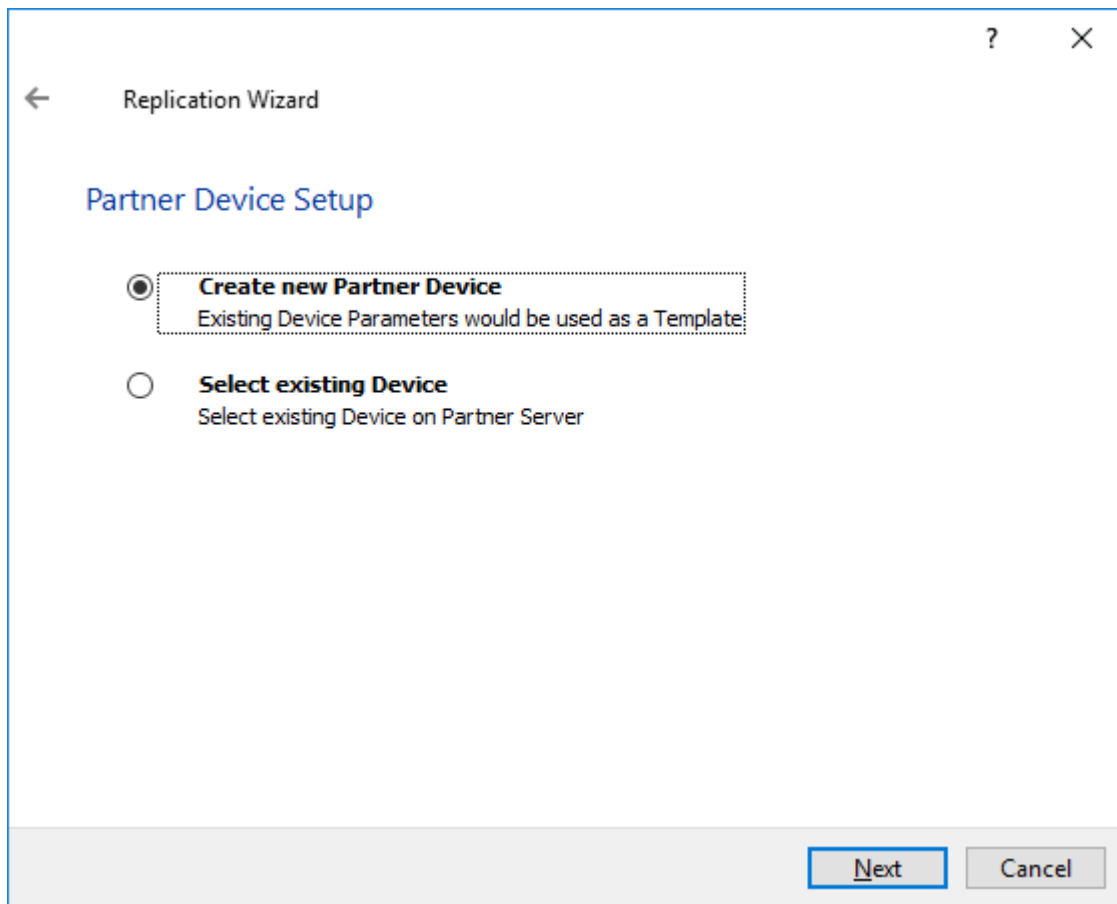
<https://www.starwindsoftware.com/help/ConfiguringFileShareSMBasWitness.html>

- Creating HA device with Witness Node:

1. Check Node Majority Failover Strategy according to the network design. Click Next to continue.



2. Select Create new Partner Device. Click Next.



3. Specify the partner device location if necessary and/or modify the target name of the device. Click Next.

4. Select Synchronization Journal Strategy and click Next.

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

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

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

Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes. Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are

recommended as the device performance is defined by the disk speed, where the journal is located. For example, it can be placed on the OS boot volume.

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

#### Failure journal

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

#### Continuous journal

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

Replication Wizard

### Synchronization Journal Setup

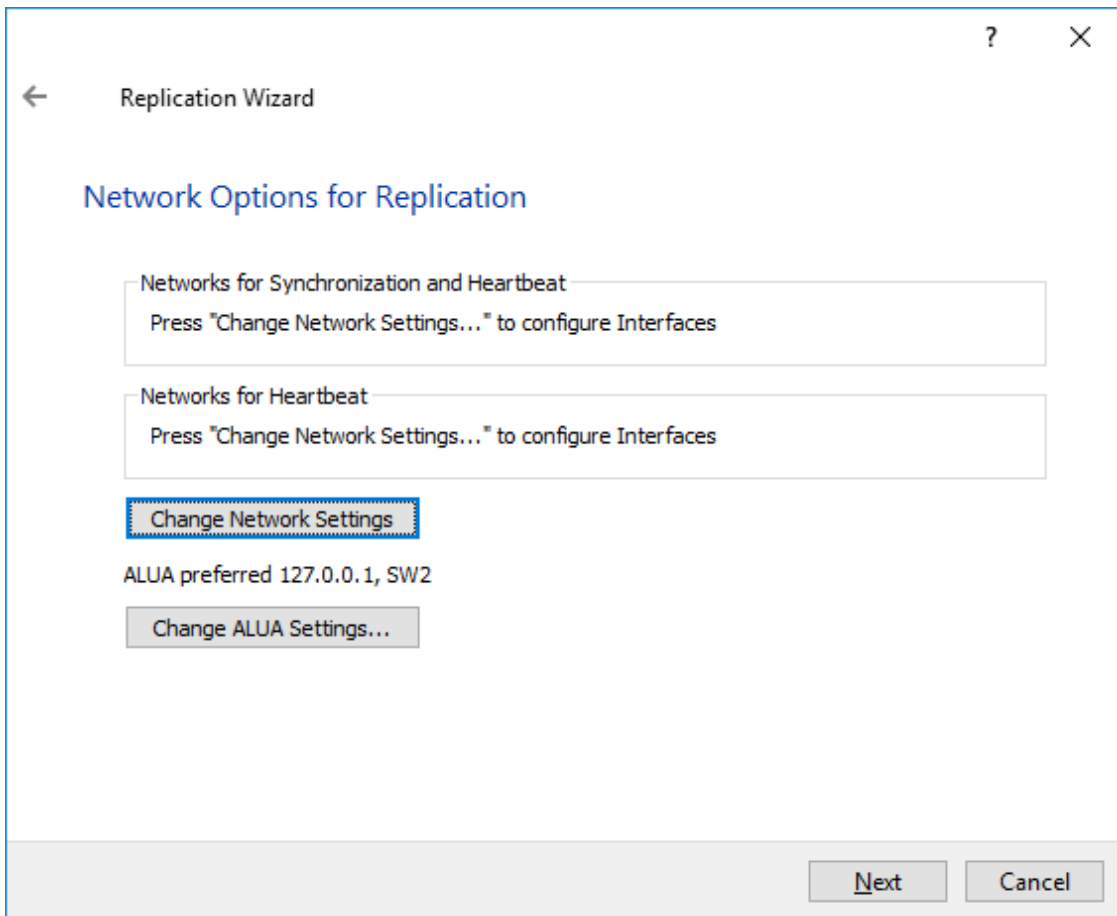
- RAM-based journal**  
Synchronization journal placed in RAM. Synchronization with RAM journal provides good IO performance in any scenario.
- Disk-based journal**  
Synchronization journal placed on disk.
- Failure journal**  
The strategy provides good IO performance while all device nodes are in a healthy state.
- Continuous journal**  
The strategy guarantees fast synchronization and data consistency in all cases.

---

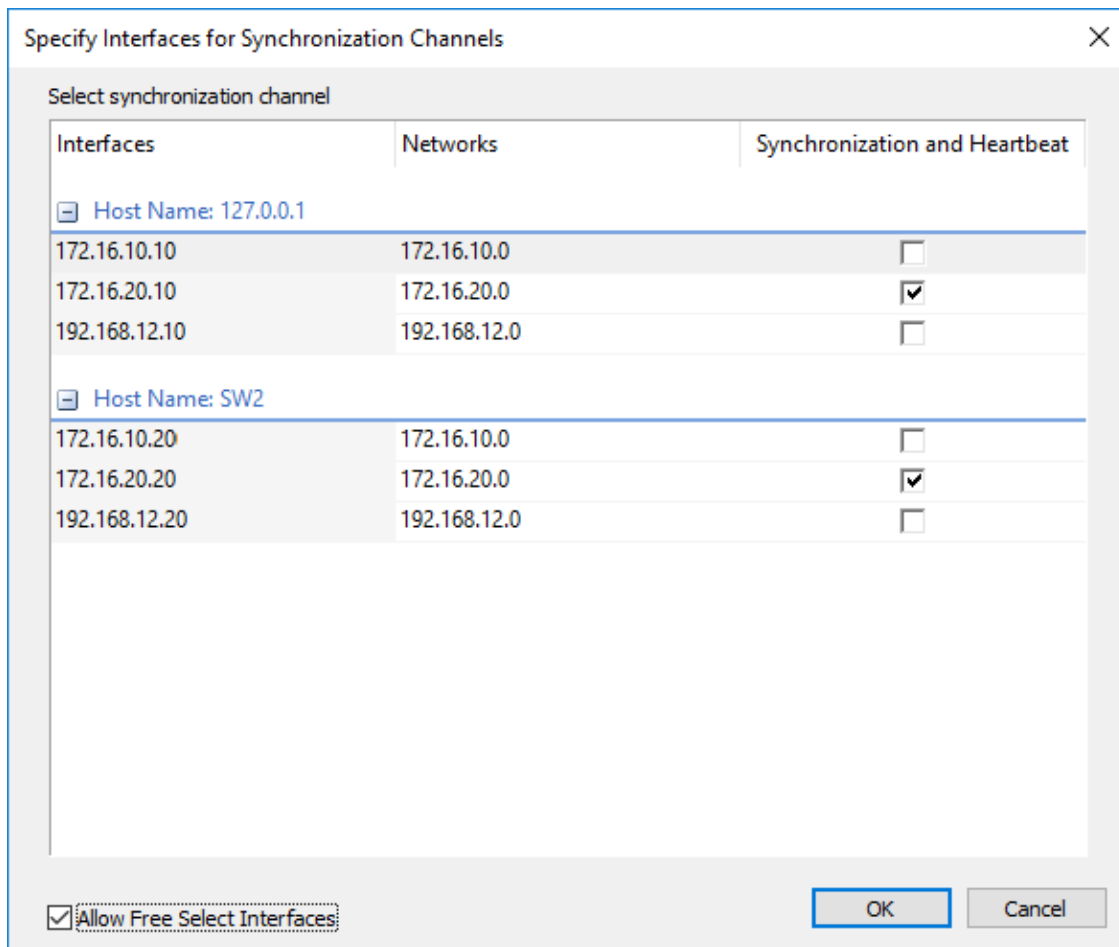
Current Node:

Partner Node:


5. Select the Synchronization and Heartbeat networks for the HA device by clicking Change Network Settings.



6. Specify the interfaces for Synchronization and Heartbeat. Press OK. Then, click Next.



7. Select Synchronize from existing Device for the partner device initialization mode. Click Next.

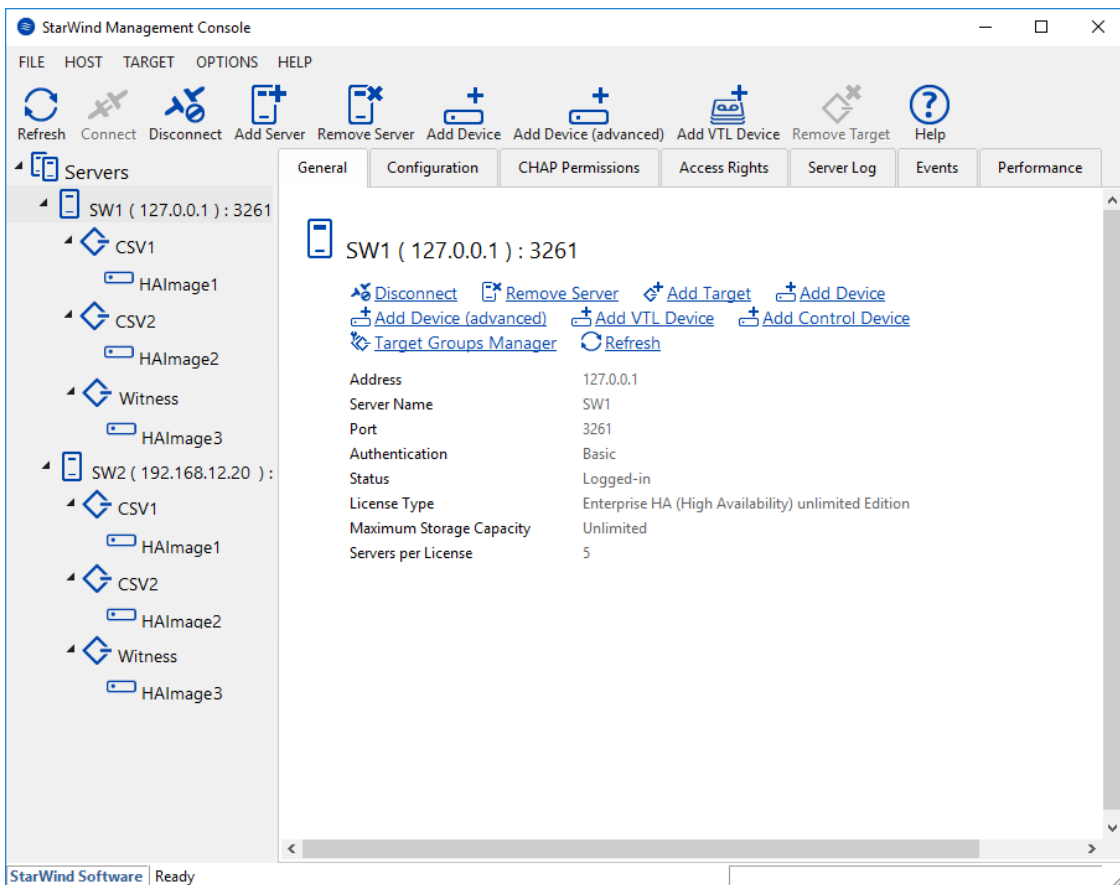


8. Press the Create Replica button. Then click Close.

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

Repeat HA device creation steps for any virtual disks that will be further used as a Cluster Shared Volumes.

Once created, the devices appear in the left pane of the Management Console as shown in the screenshot below.

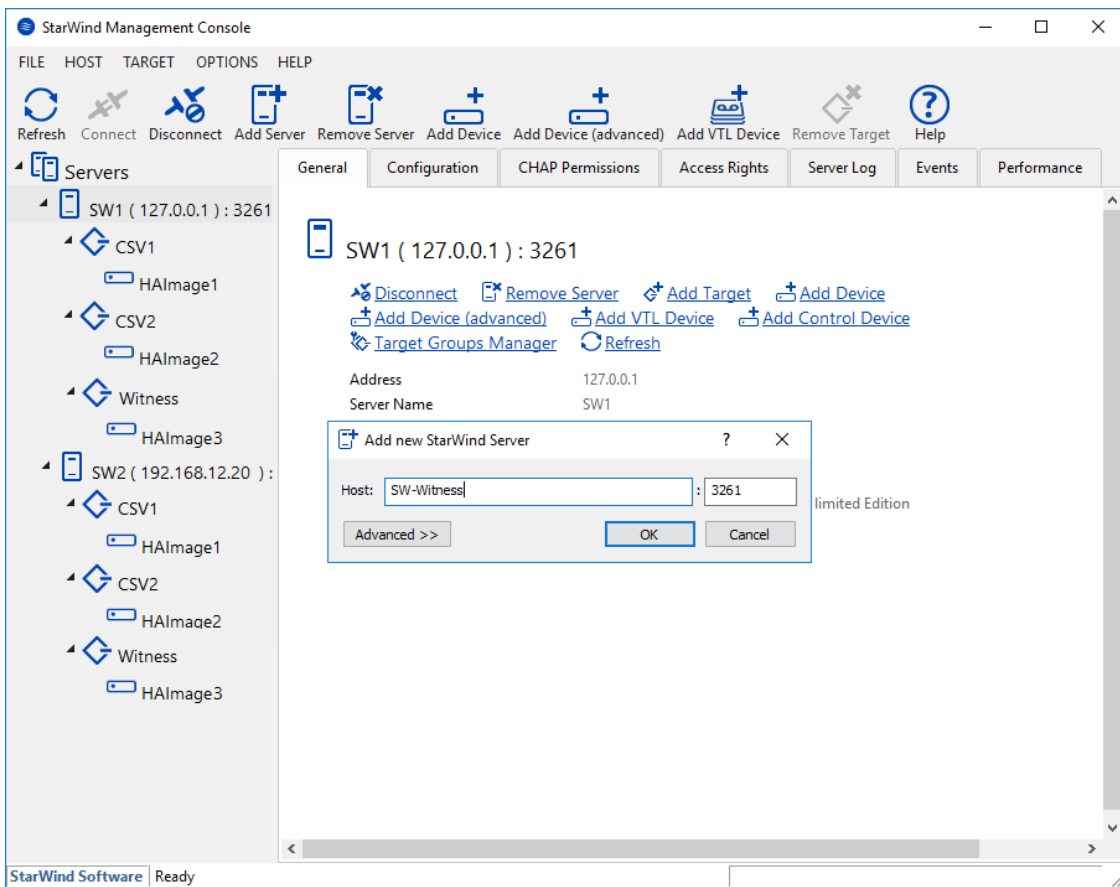


## Adding Witness Node

This section describes adding a Witness node. Witness nodes count for the majority, but neither contain data nor process any clients' requests.

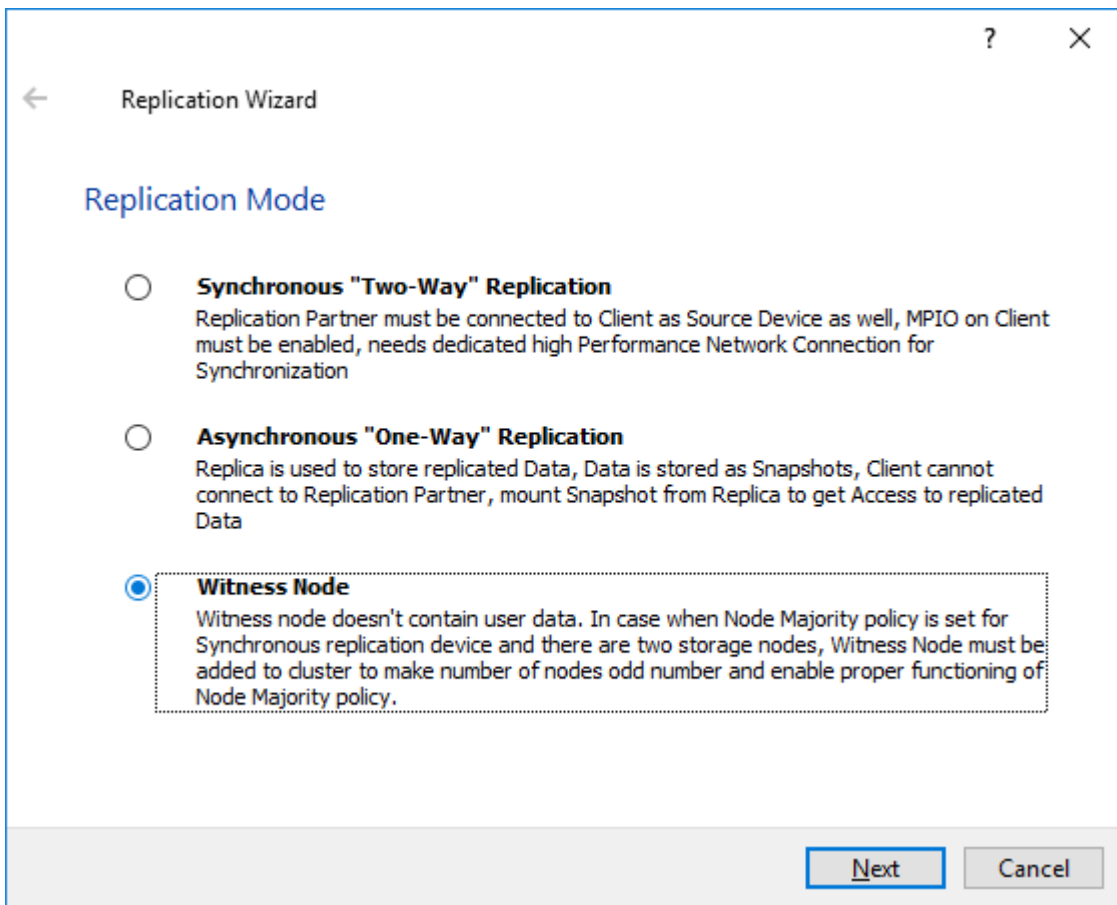
Configure the Witness node at a separate location. There are two options to do so: it can either be a virtual machine run in the cloud or a host at another site. Witness node should have StarWind Virtual SAN service installed on it.

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



2. Right-click the HA device with the configured Node Majority failover policy and select Replication Manager. The Replication Manager window will appear. Press the Add Replica button.

3. Select Witness Node and click Next




4. Specify the Witness node name or its IP address.


The screenshot shows a window titled "Replication Wizard" with a back arrow on the left and help and close buttons on the right. The main heading is "Add Partner Node". Below this, there is a prompt: "Specify Partner Host Name or IP Address where Replication Node would be created". There are two input fields: "Host Name or IP Address" with a dropdown menu containing "SW-Witness" and a small downward arrow, and "Port Number" with a text box containing "3261". At the bottom right, there are two buttons: "Next" (highlighted with a blue border) and "Cancel".

5. Specify the Witness device location and its target name if necessary. Click Next.


6. For the HA device, select the synchronization channel with the Witness node by clicking on the Change network settings button.



7. Specify Interfaces for Synchronization Channels, confirm, and click Next.



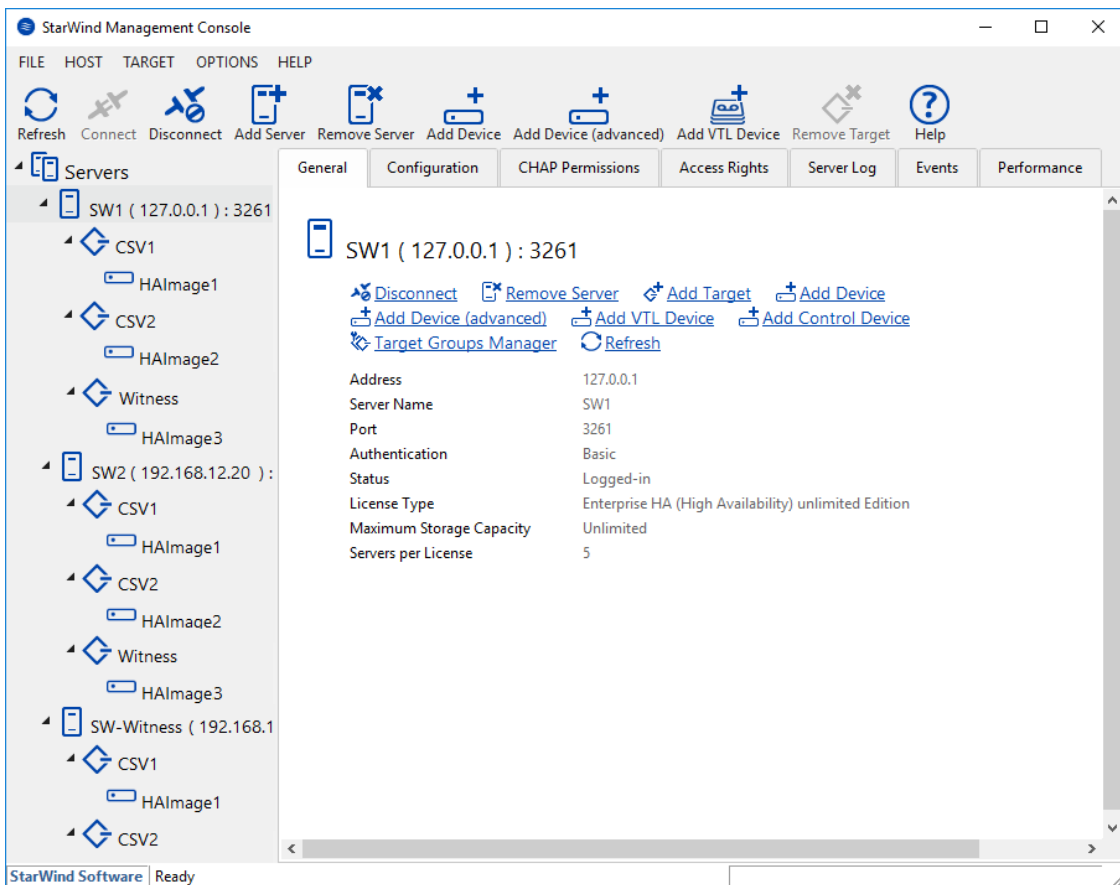
8. Click Create Replica.



9. After the device creation is completed, close the Wizard by pressing the Close button.

10. Repeat the steps above to create other virtual disks.

11. The added device will appear in StarWind Management Console. The list of HA devices should look as follows:



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](#)

## 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.

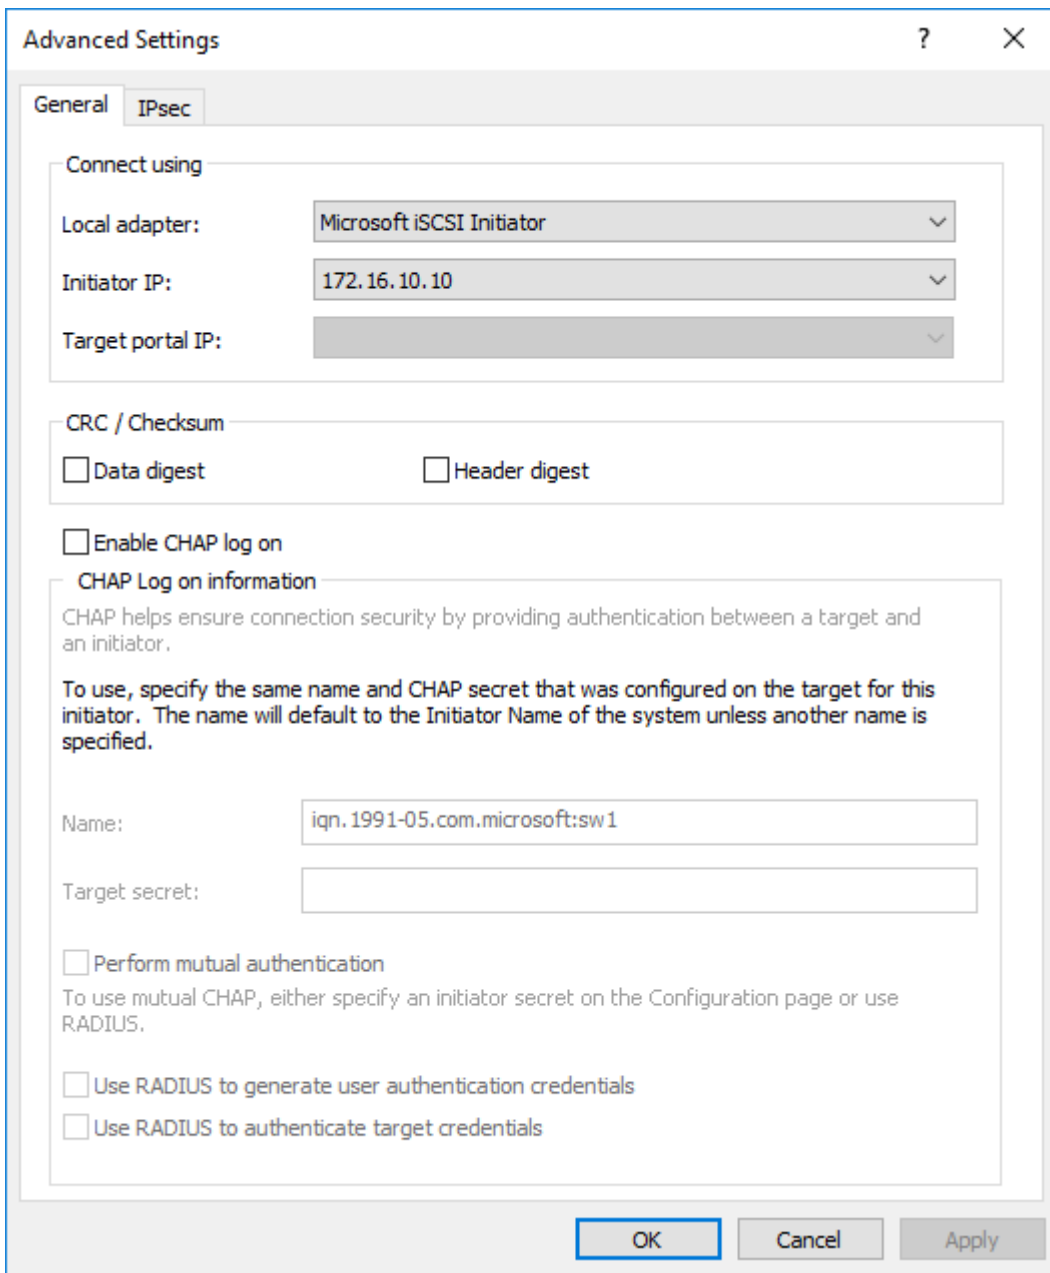


5. Click the Discover Portal... button once again.


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.

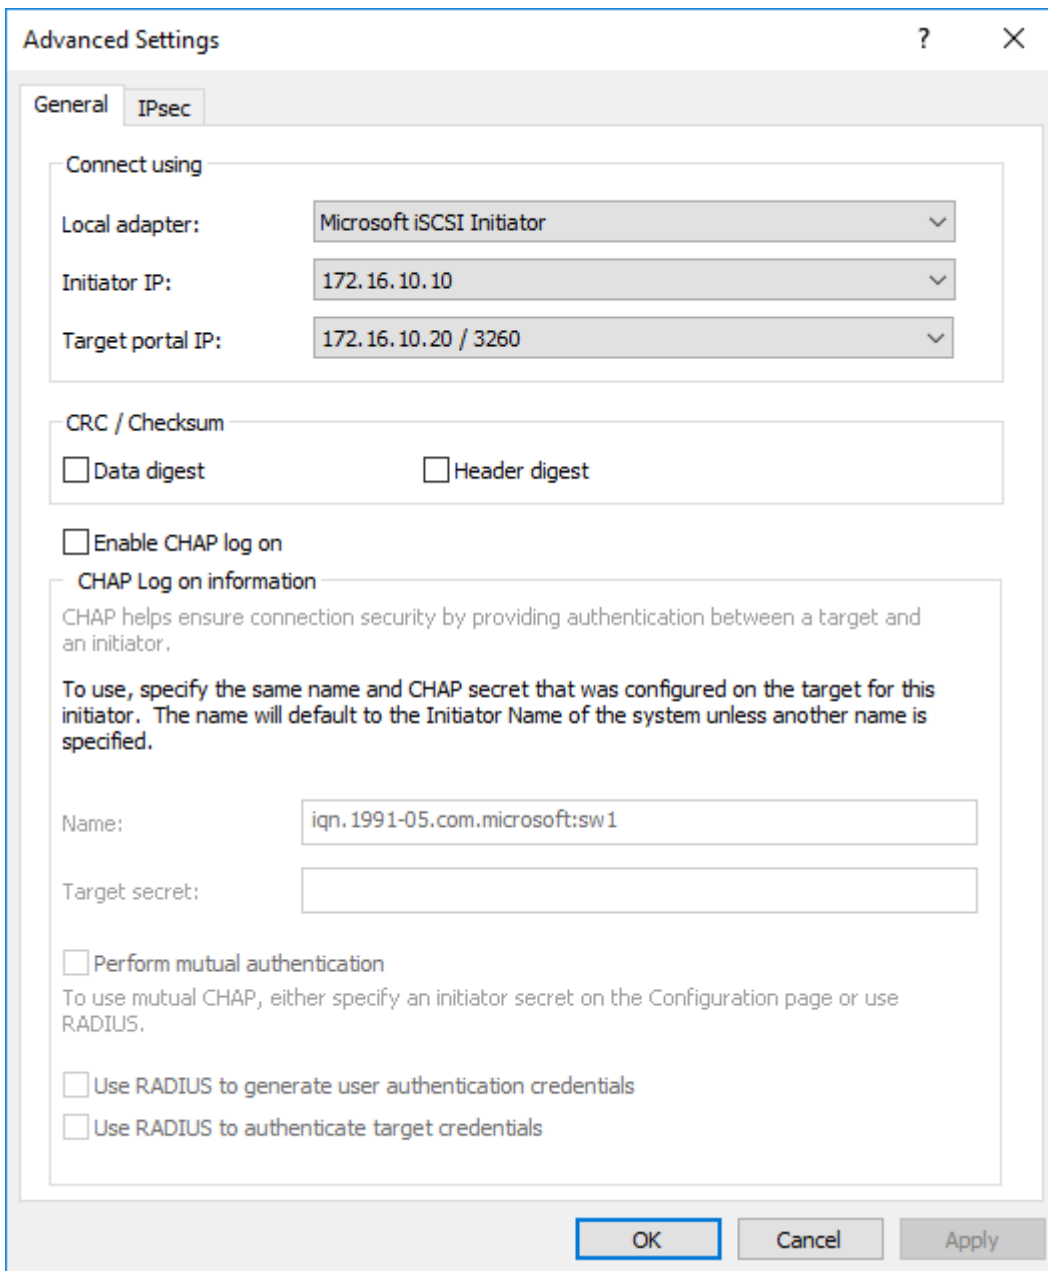


7. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In Target portal IP, select 127.0.0.1. Confirm the actions.

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.



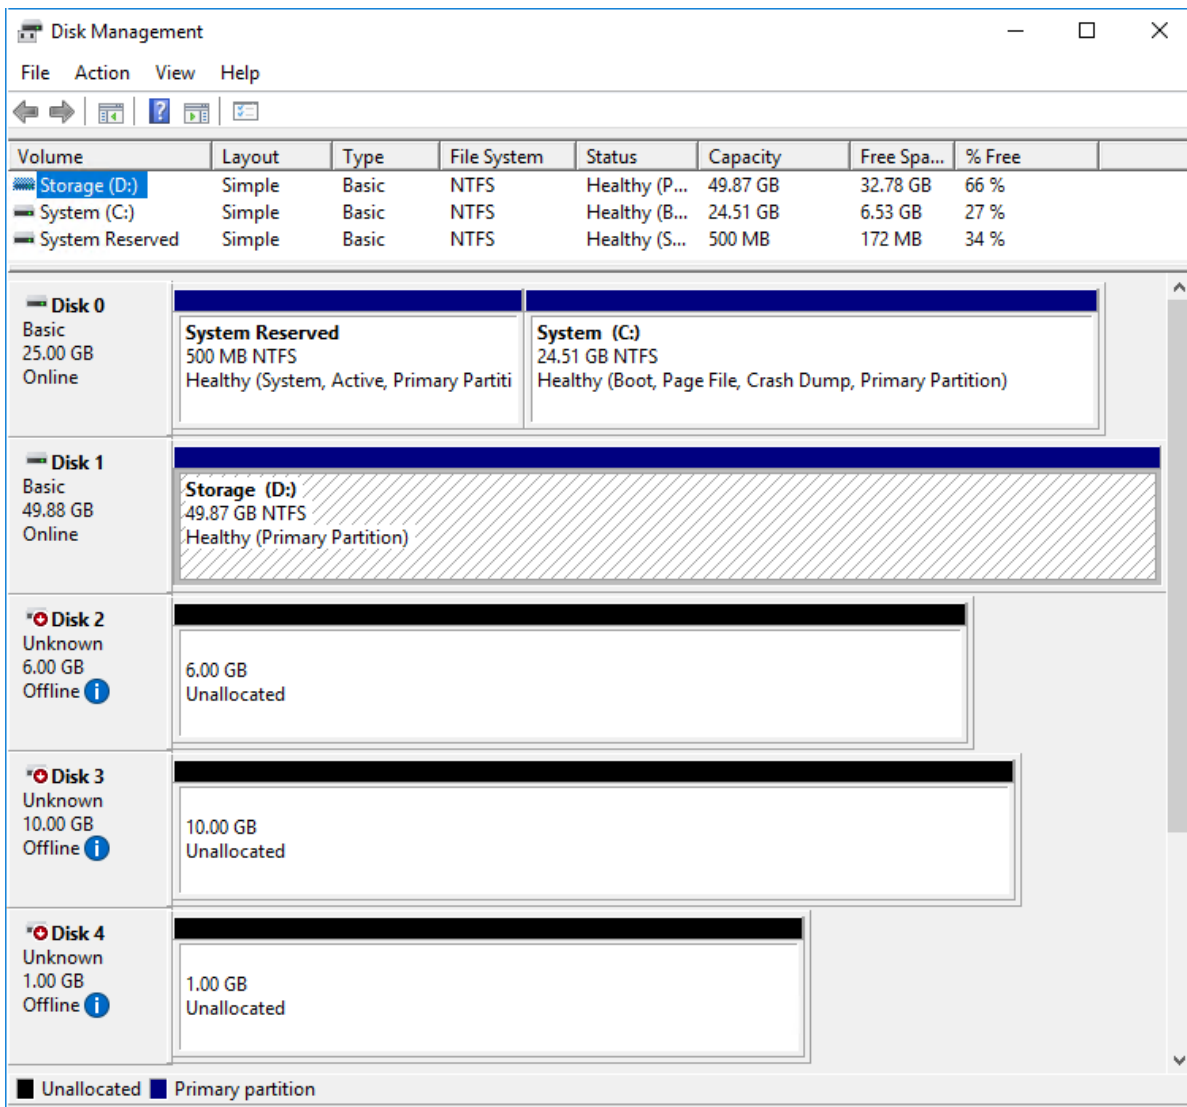
4. For the Witness target, set the load balance policy to Failover Only.

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.

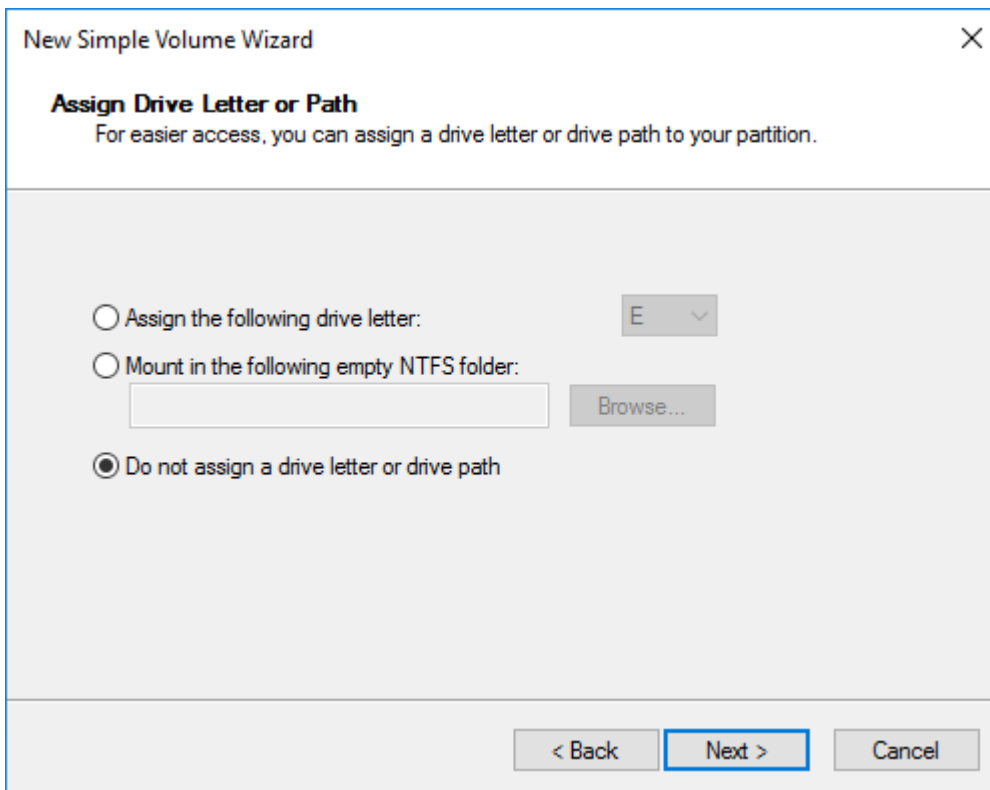


8. Select NTFS in the File System dropdown menu. Keep Allocation unit size as Default. Set the Volume Label of choice. Click Next.



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

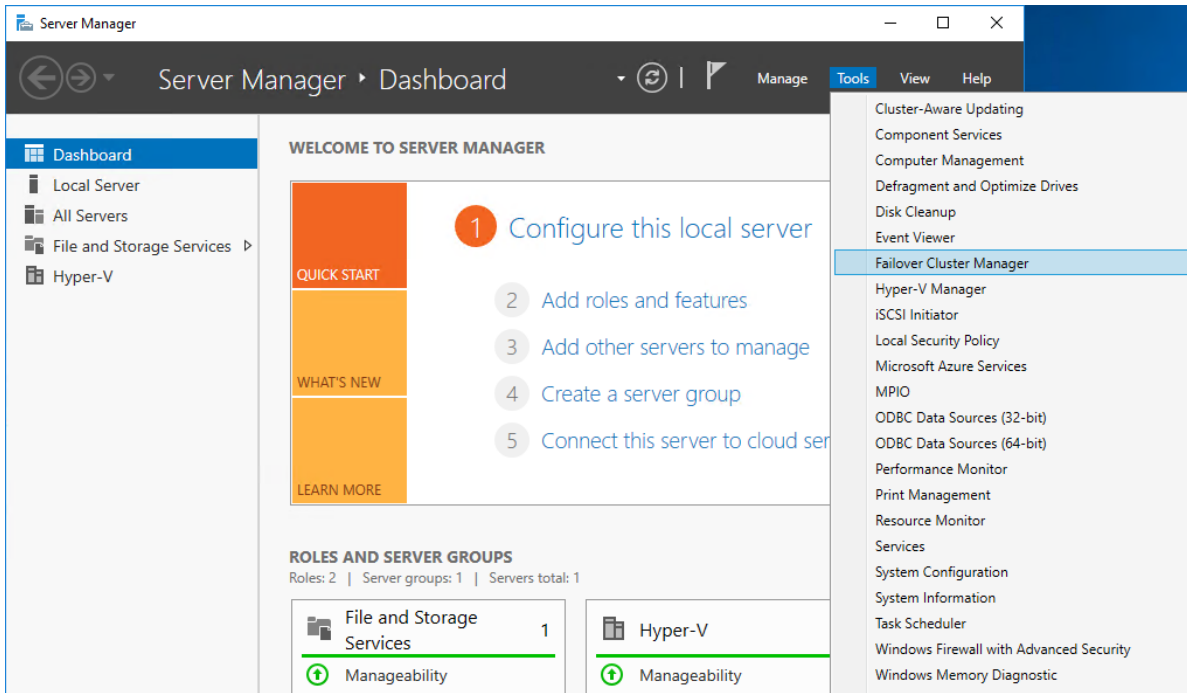
NOTE: To avoid issues during the cluster validation configuration, it is recommended to install the latest Microsoft updates on each node.

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:

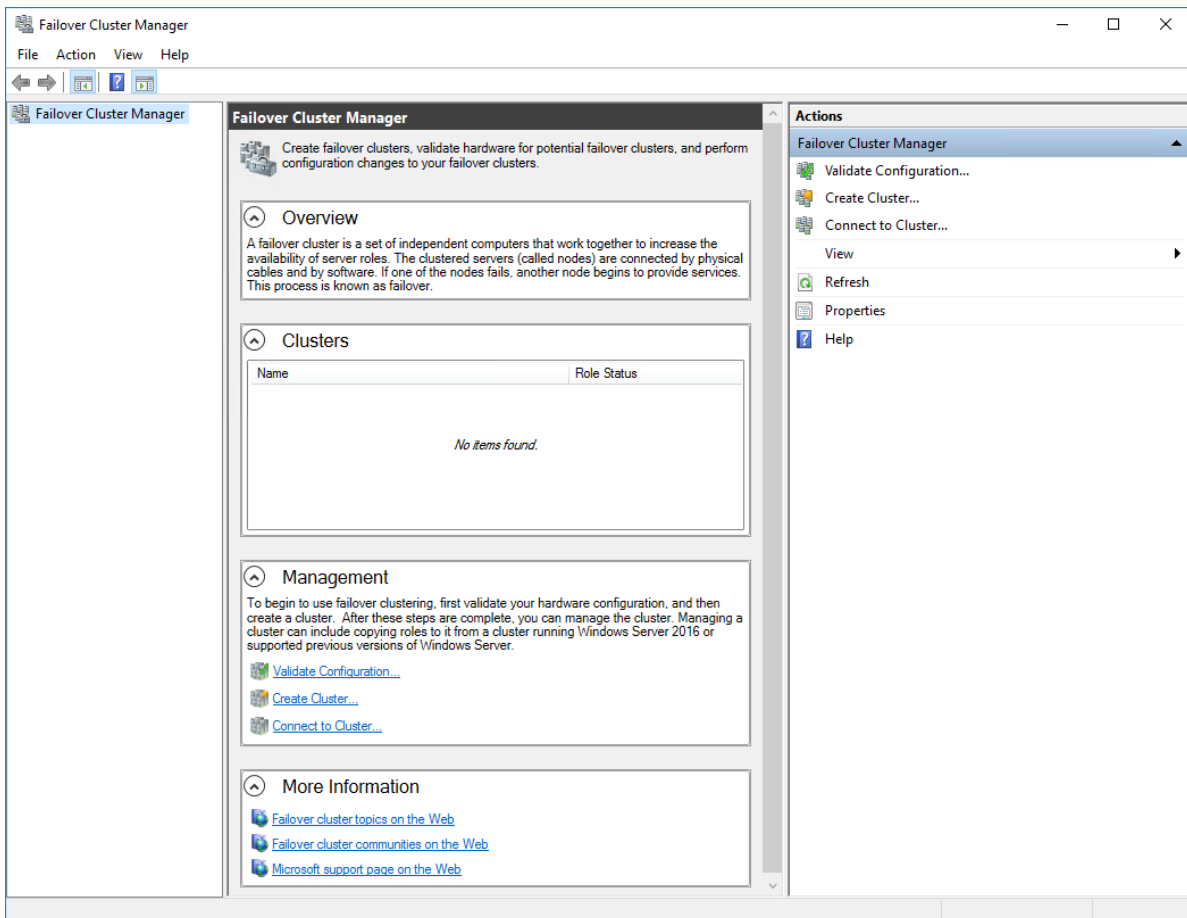
<https://docs.microsoft.com/en-us/windows-server/remote/remote-server-administration-tools>

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.

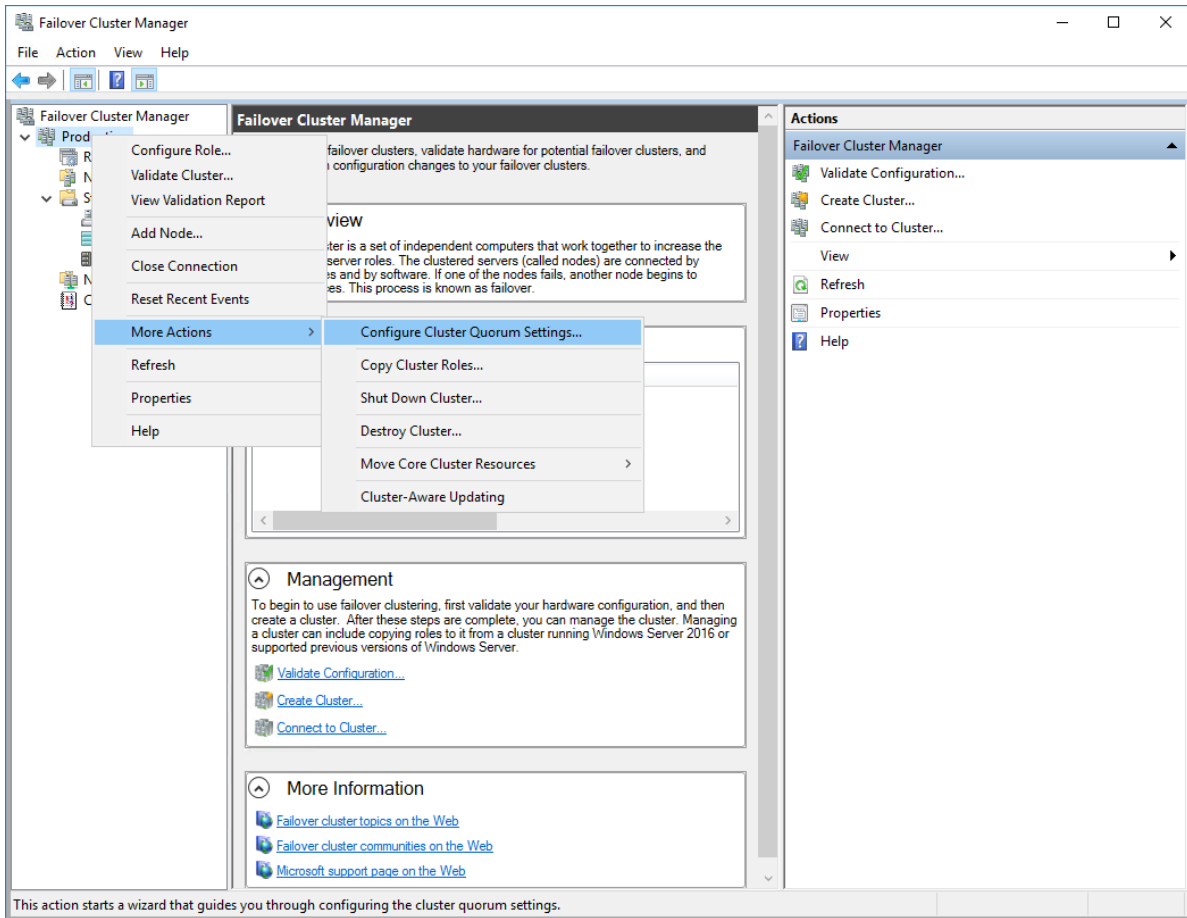


## Adding Storage to the Cluster


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.



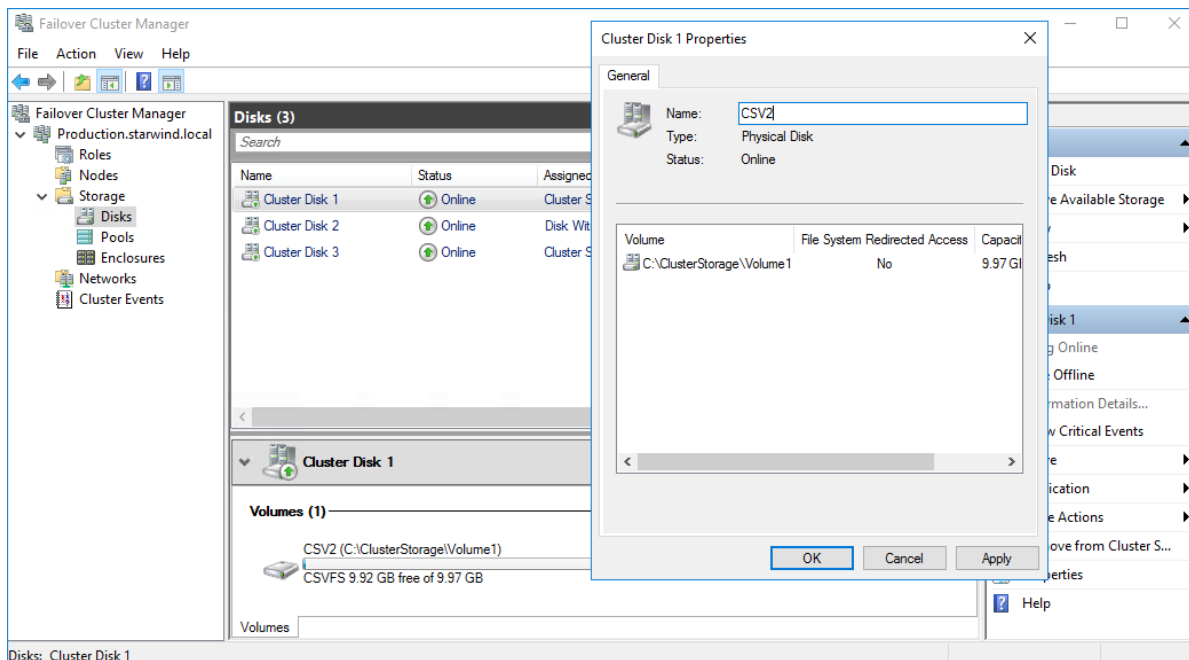
5. Select the Witness disk to be assigned as the cluster witness disk. Click Next and press Finish to complete the operation.




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: 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.



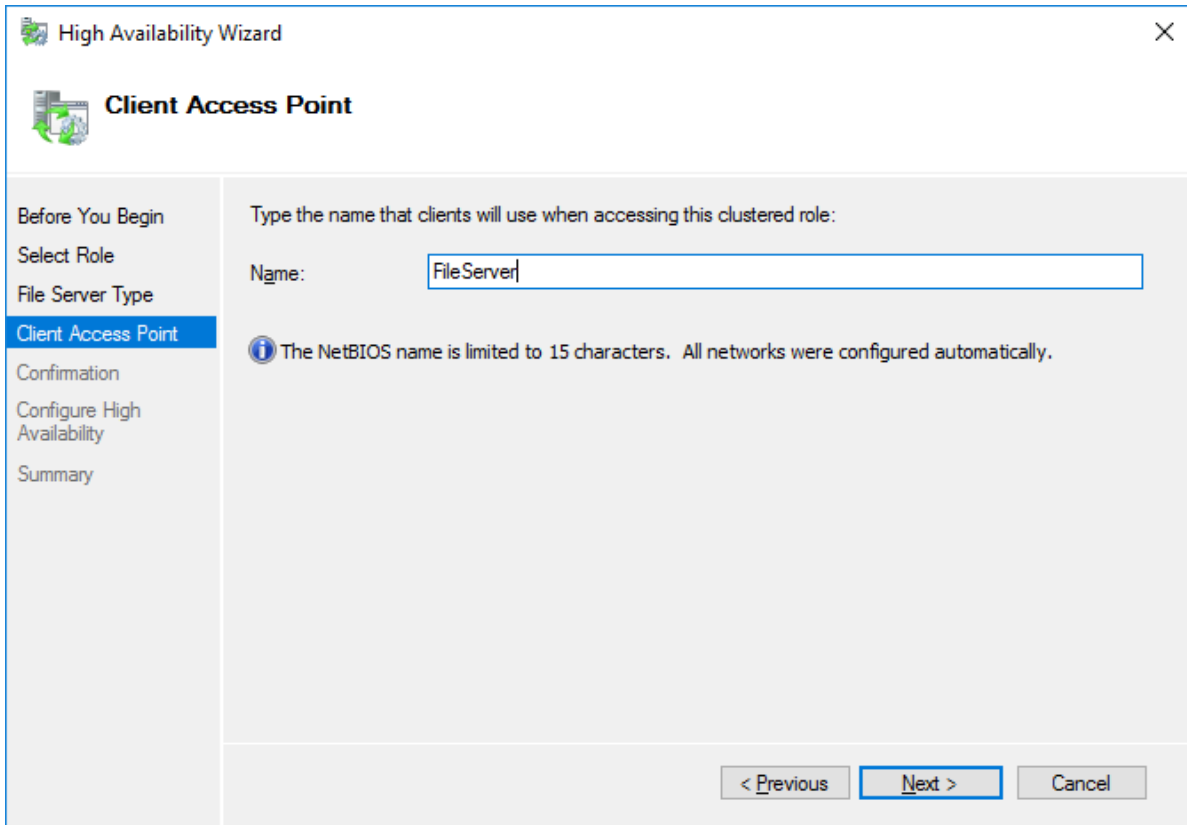
3. Select the File Server item from the list in High Availability Wizard and click Next to continue.



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.

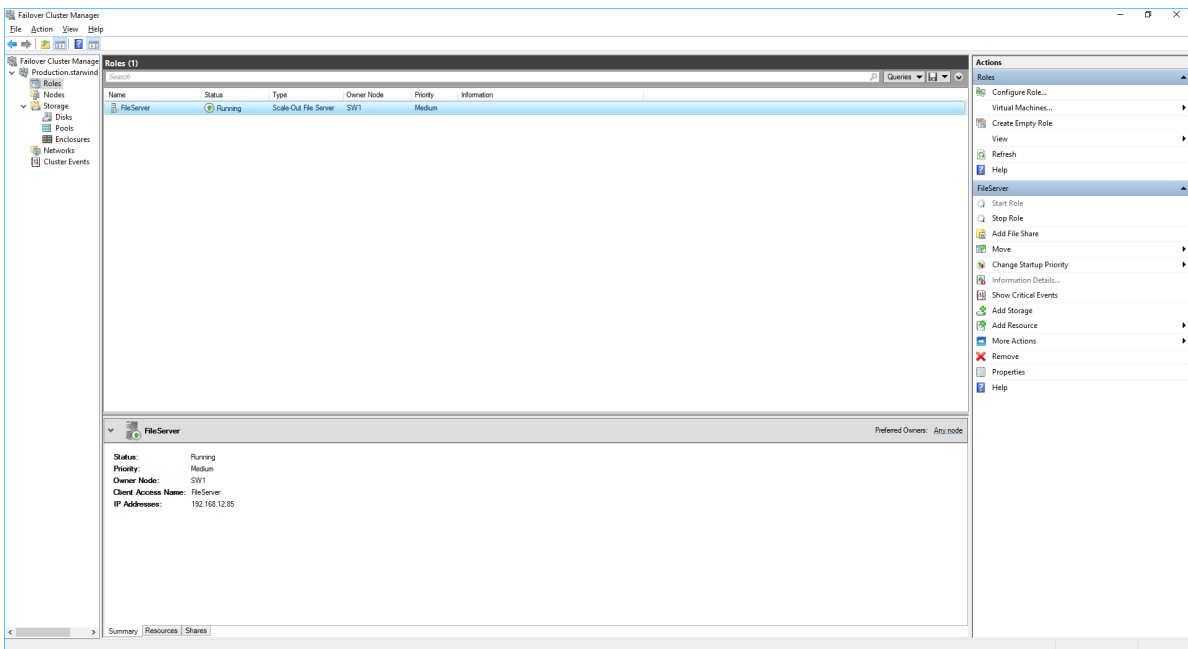


7. Once the installation is finished successfully, the Wizard should now look like the screenshot below.

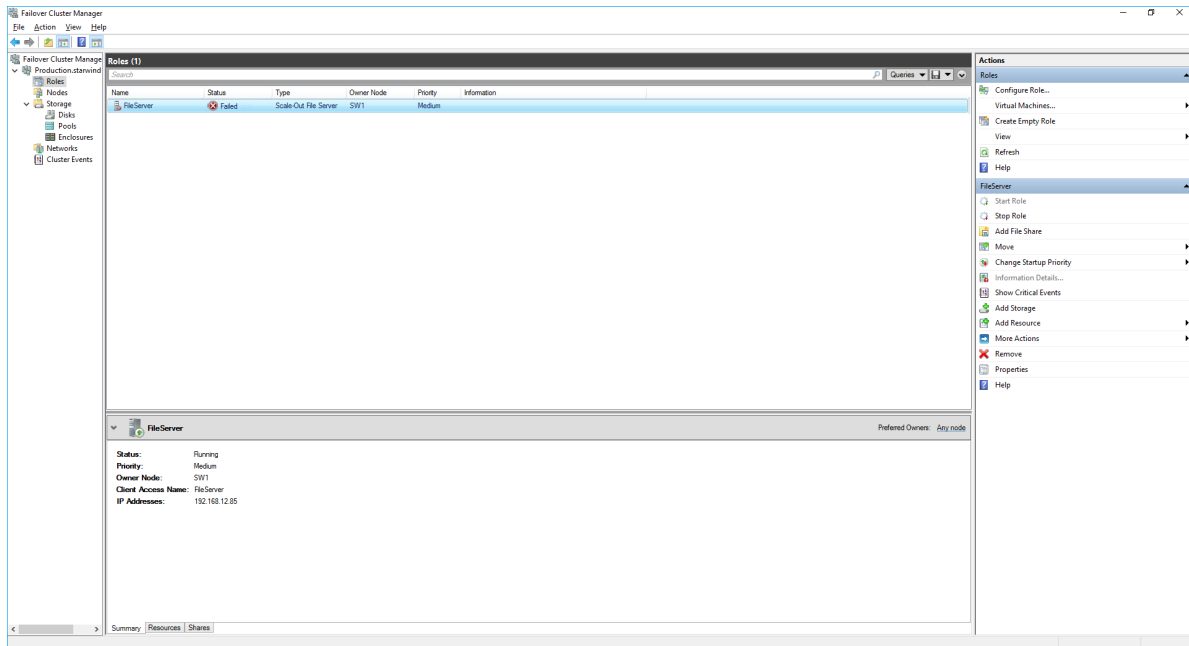
Click Finish to close the Wizard.



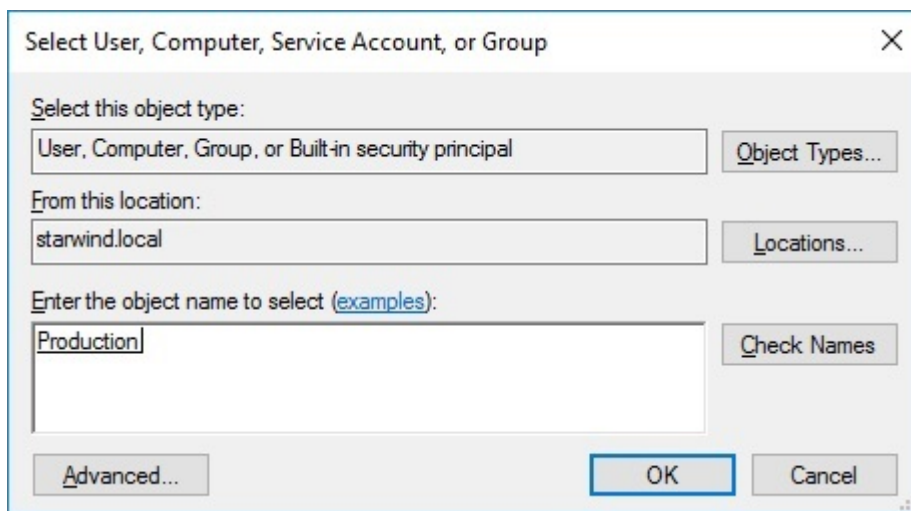
8. The newly created role should now look like the screenshot below.



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

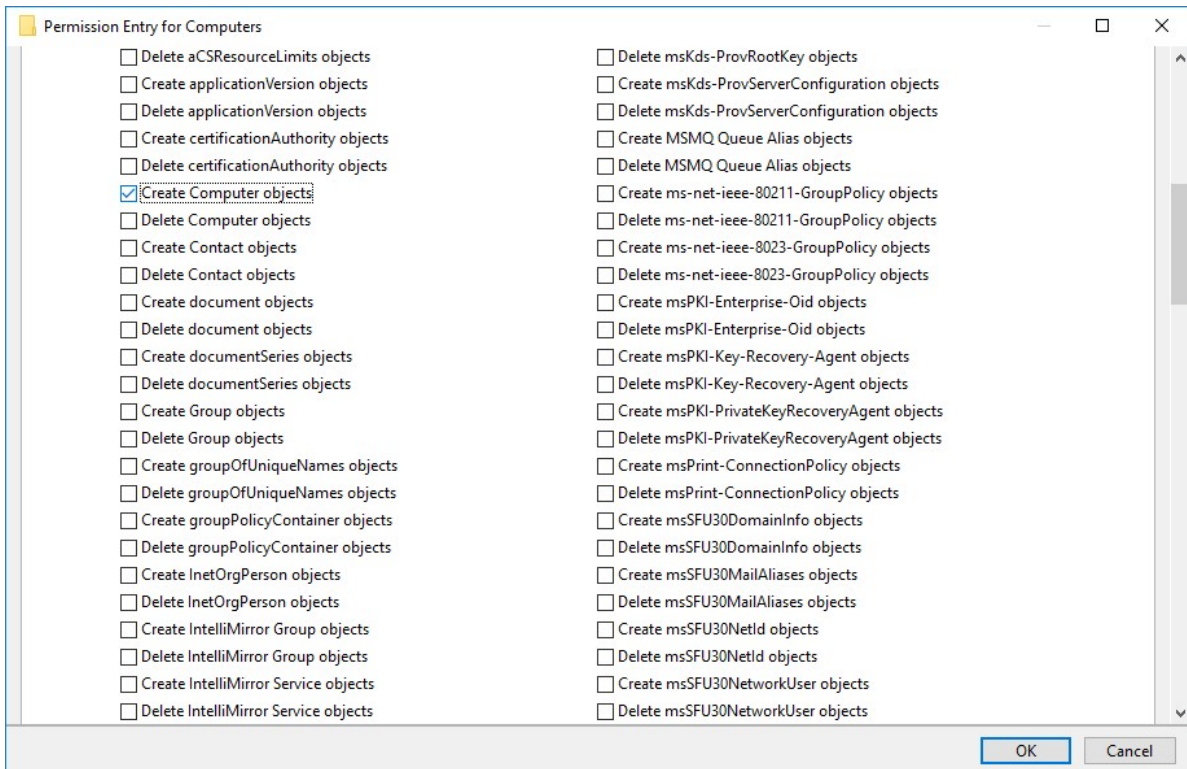


- open Active Directory Users and Computers
- enable the Advanced view if it is not enabled
- edit the properties of the OU containing the cluster computer object (in this case - Production)
- open the Security tab and click Advanced
- in the appeared window, press Add (the Permission Entry dialog box 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

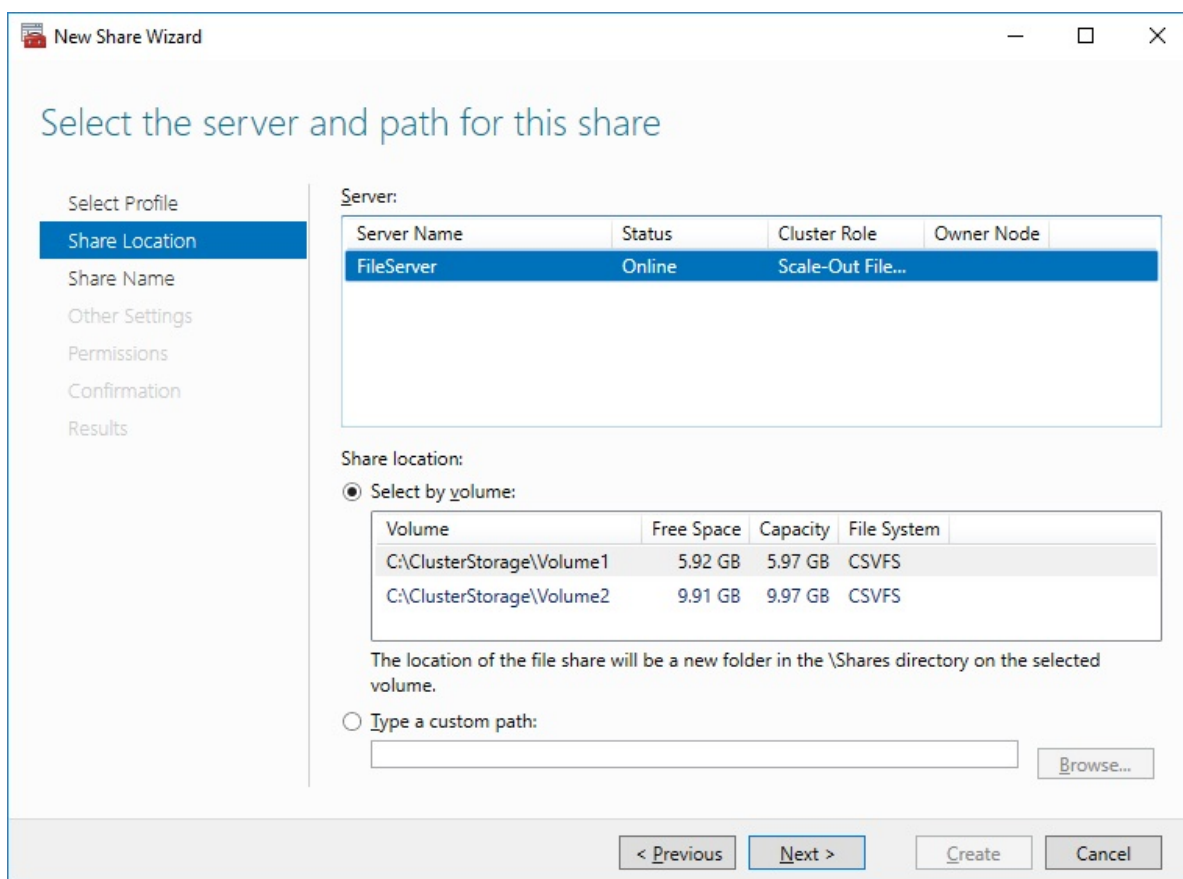
Configuring File Share

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



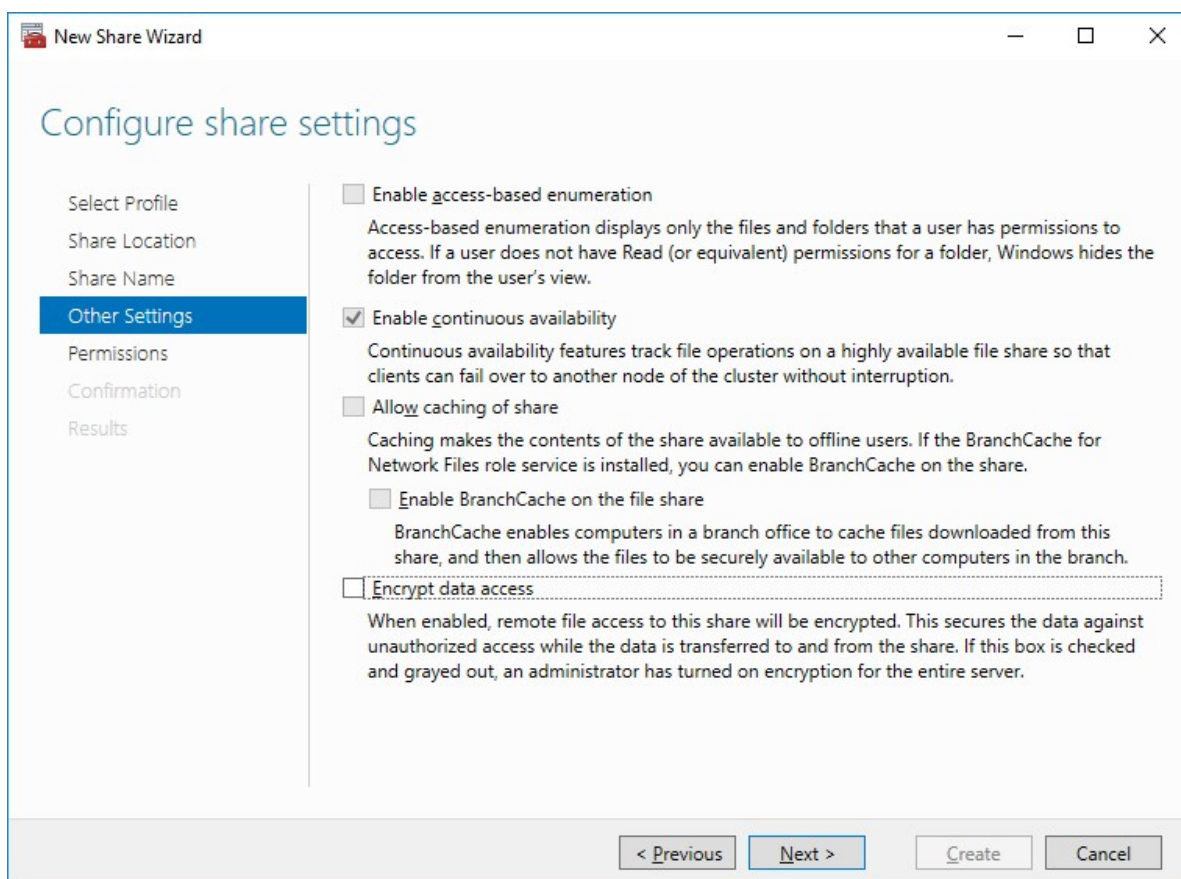
5. Select a CSV to host the share. Click Next to proceed.




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.



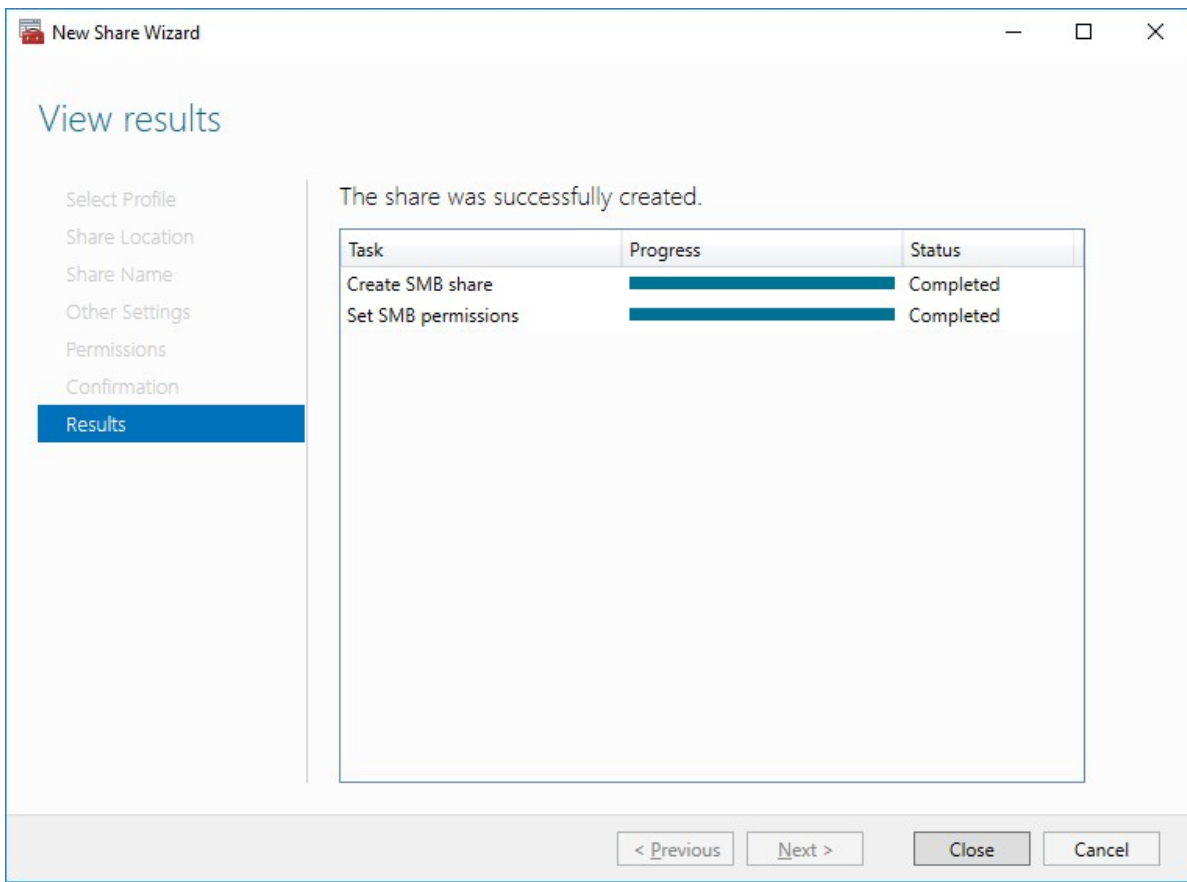
**NOTE:**

- 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.



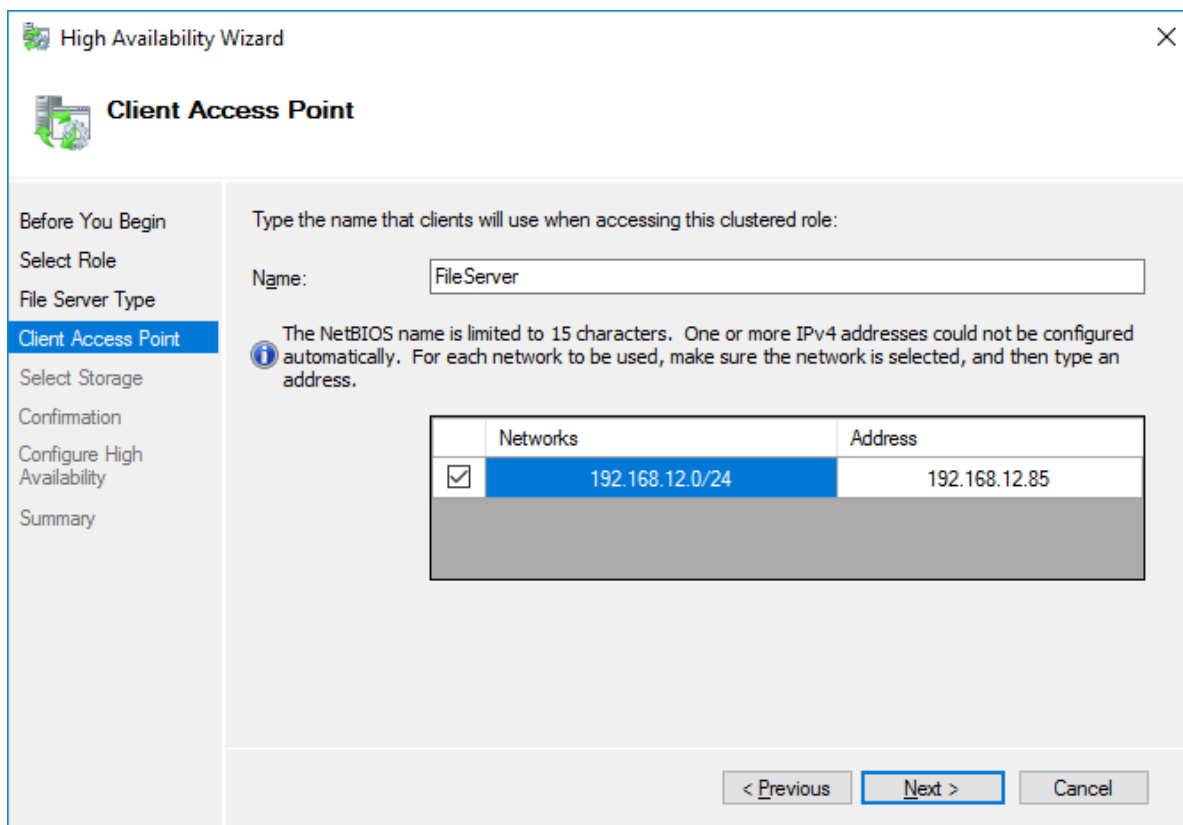
3. Select the File Server item from the list in High Availability Wizard and click Next to continue.



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

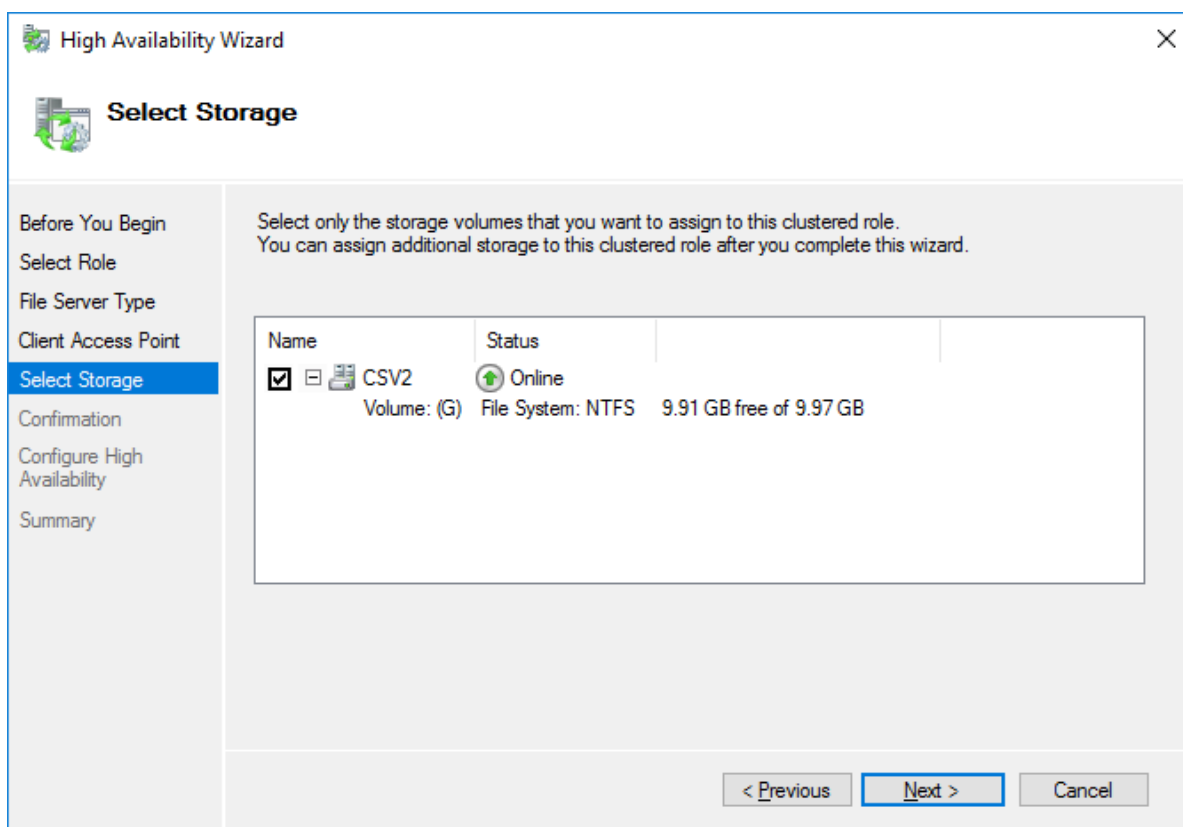


5. On the Client Access Point page, in the Name text field, type the NETBIOS name that will be used to access the File Server and IP for it.




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.

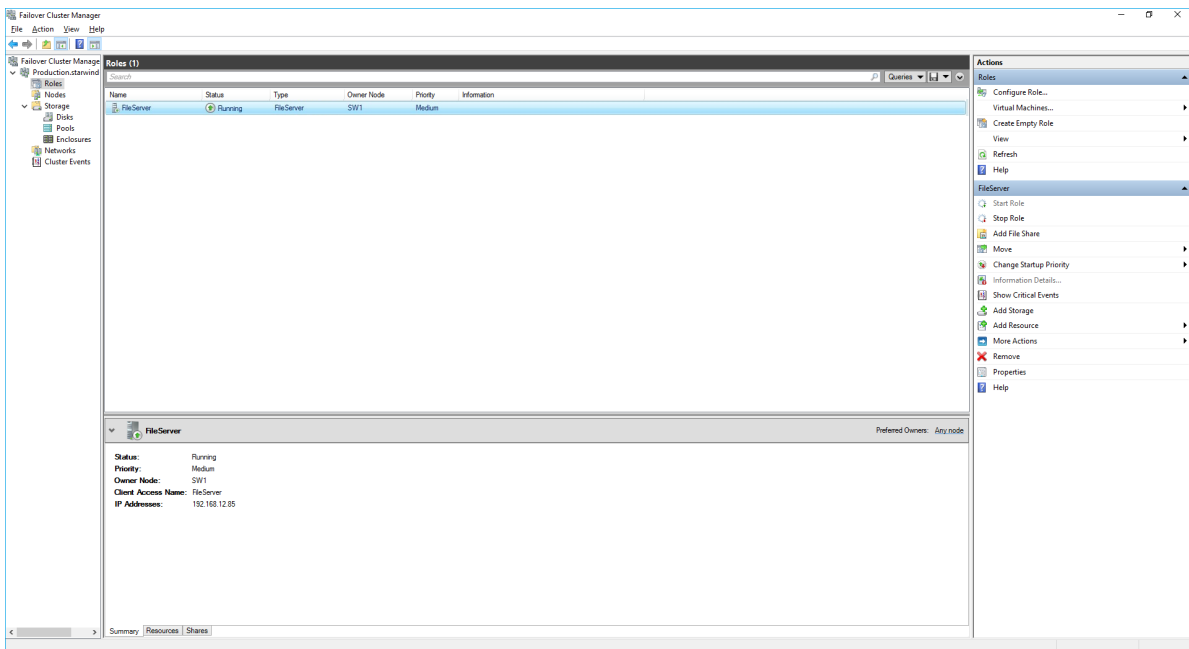


8. Once the installation has been finished successfully, the Wizard should now look like the screenshot below.

Click Finish to close the Wizard.

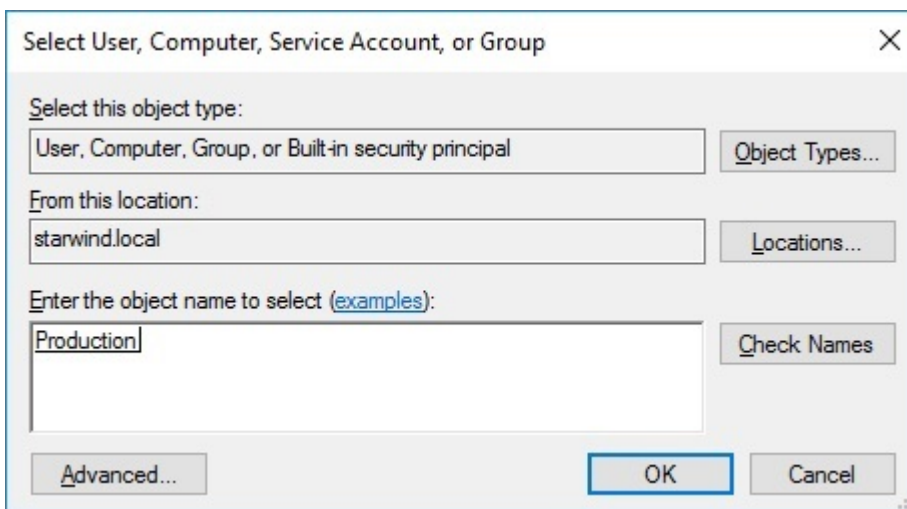


9. The newly created role should now look like the screenshot below.

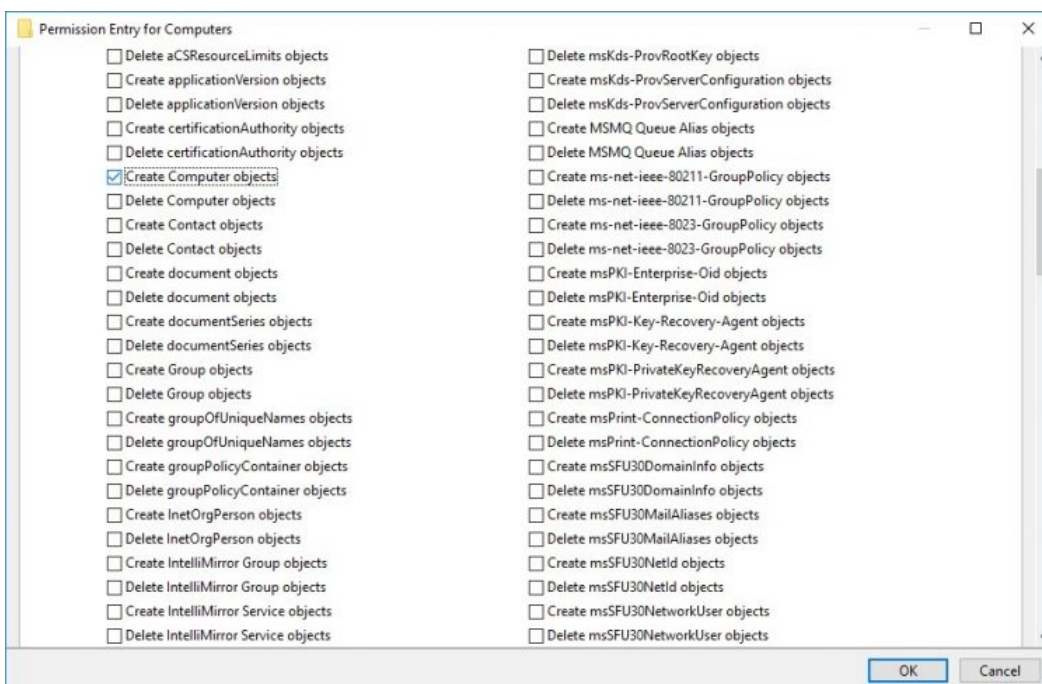


NOTE: If the role status is Failed and it is unable to Start, please, follow the next steps:

- 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

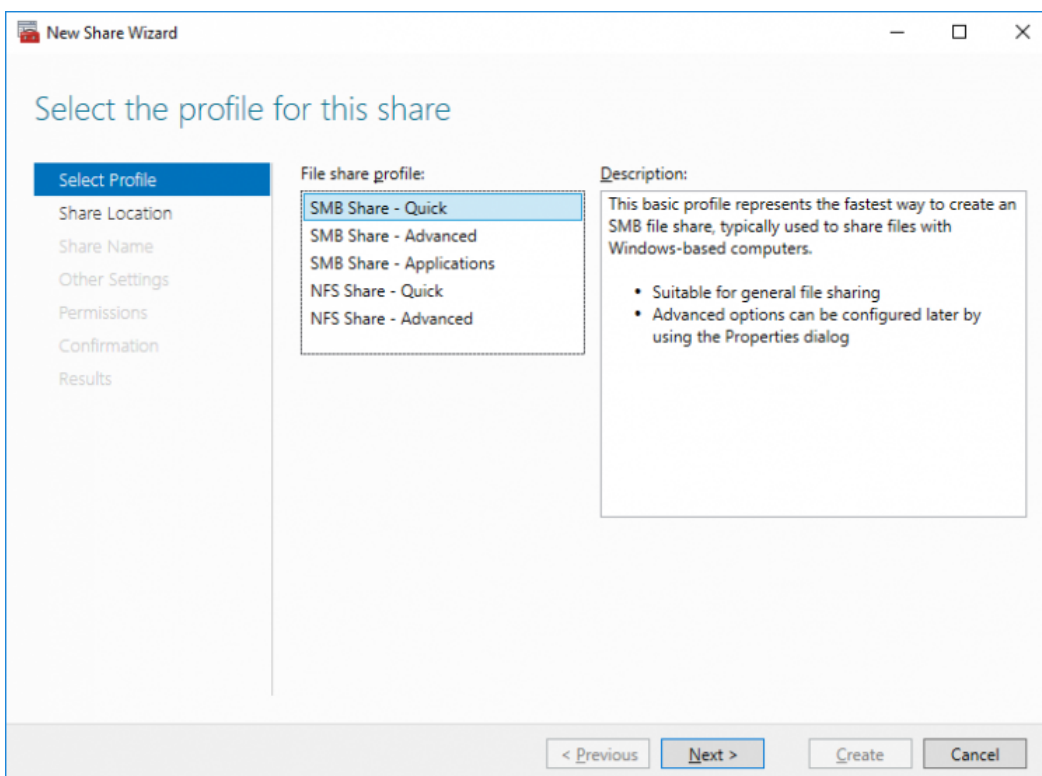


- click OK on all opened windows to confirm the changes
- open Failover Cluster Manager, right-click File Share role and click Start Role

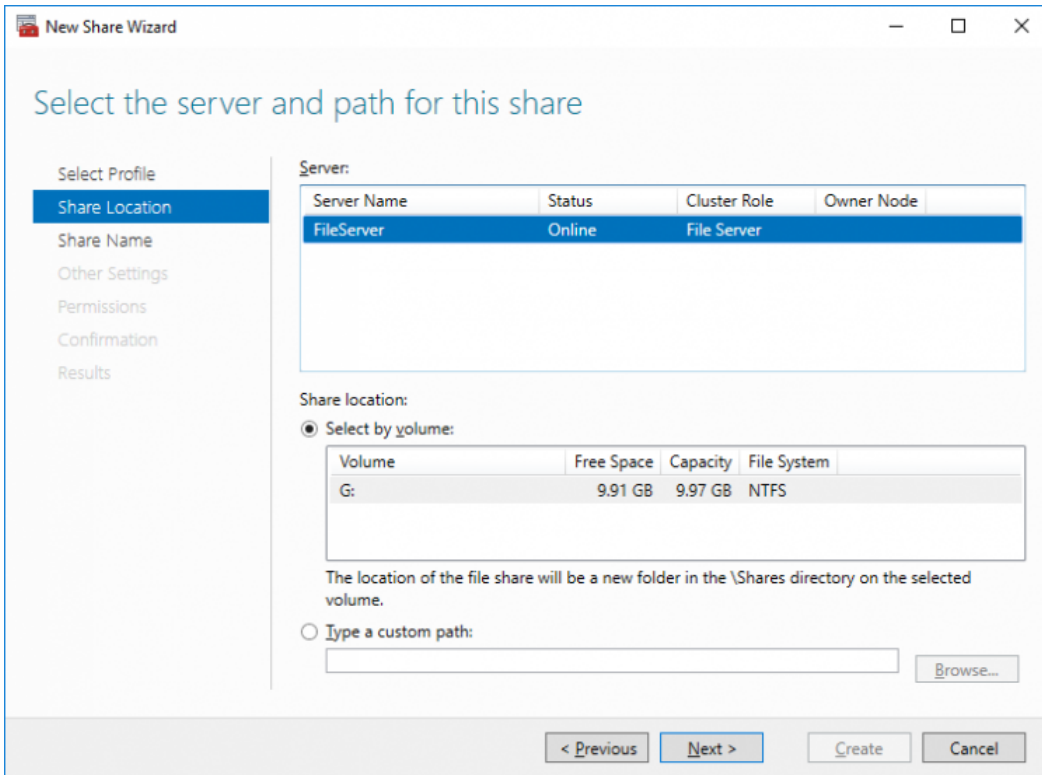
## Configuring Smb File Share

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.

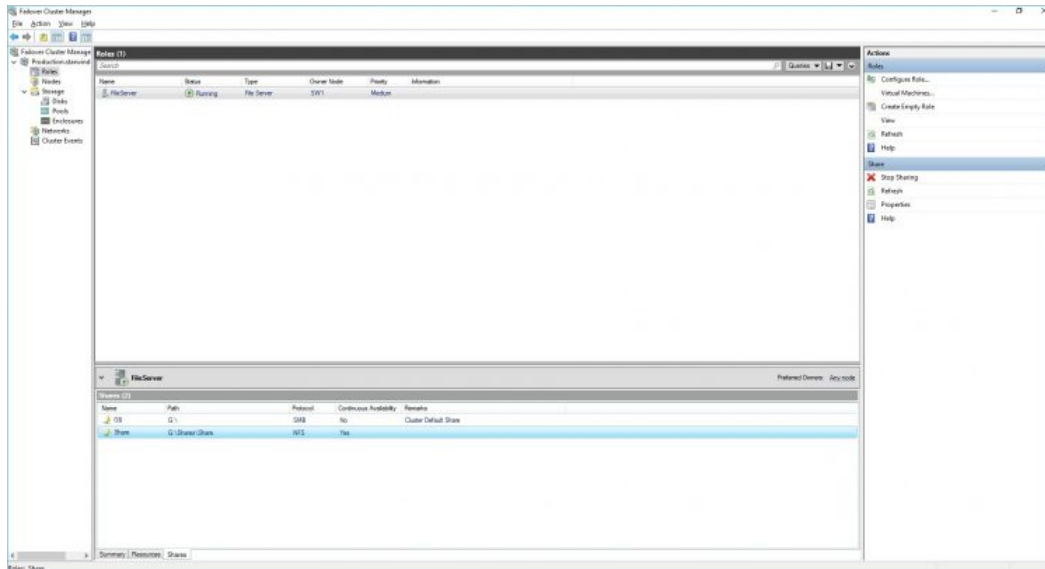


To manage created SMB File Shares

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.



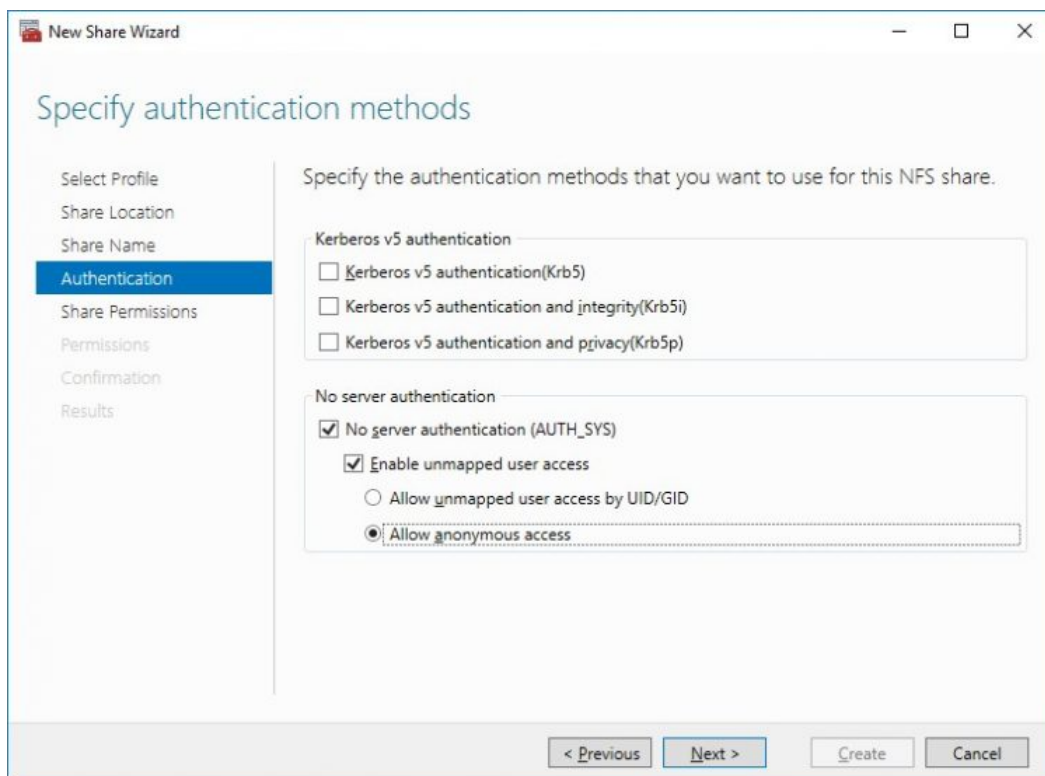
5. Select available storage to host the share. Click Next to continue.



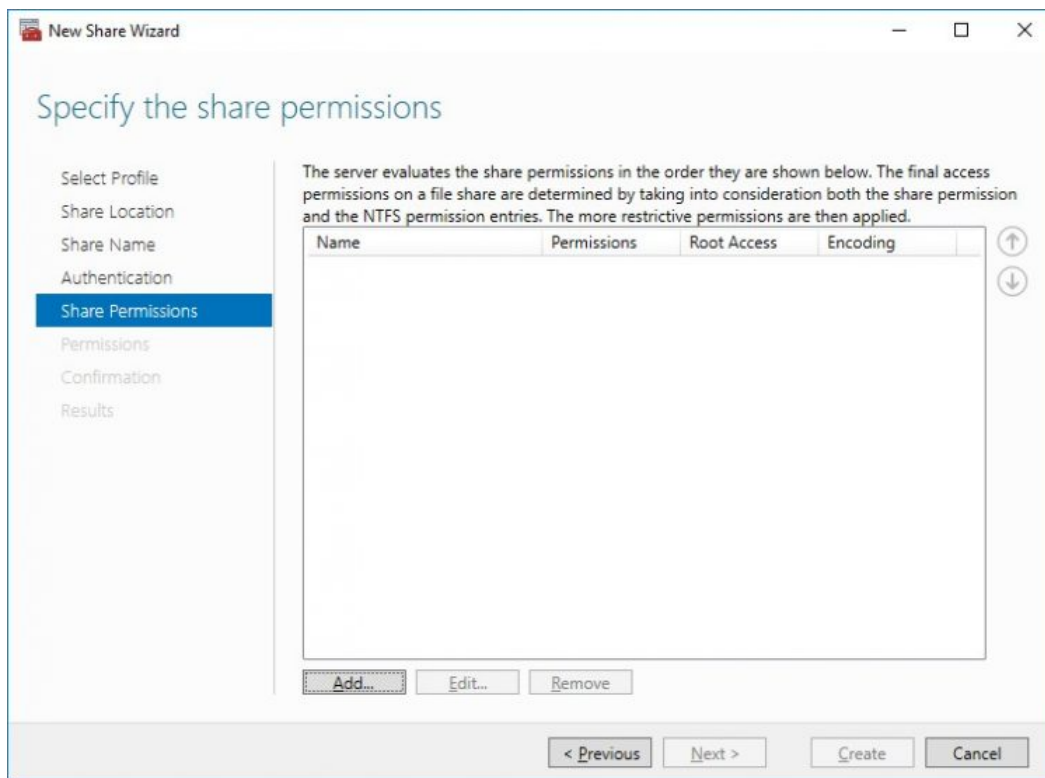
6. Type in the file share name and click Next.




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




8. Click Add and specify Share Permissions.






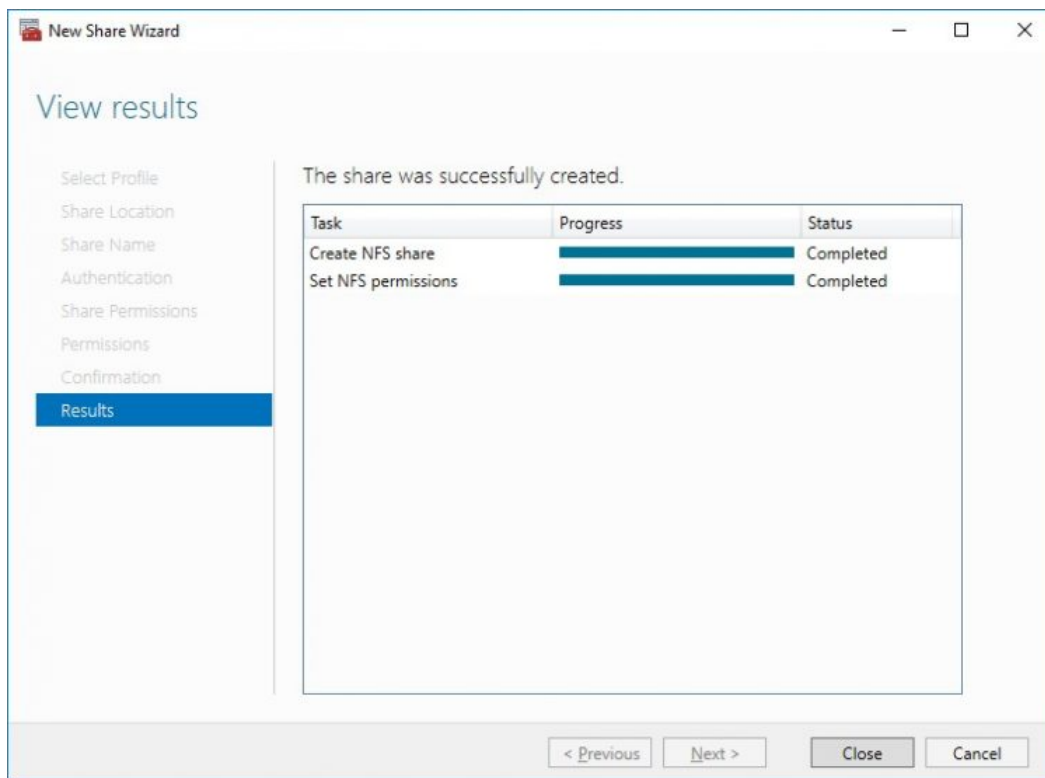
9. Specify the access permissions for the file share.



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, a robust virtualized storage solution was deployed, providing a scalable, high-availability storage platform suitable for a wide range of virtualization workloads in the stretched Microsoft Failover Cluster.

## Contacts

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

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

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

Sales: [sales@starwind.com](mailto:sales@starwind.com)

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



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