

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

2026

TECHNICAL PAPERS



Trademarks

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

Changes

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

Technical Support and Services

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

About StarWind

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

Copyright ©2009-2018 StarWind Software Inc.

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

Annotation

Relevant Products

This guide applies to StarWind Virtual SAN and StarWind Virtual SAN Free, specifically Controller Virtual Machine (CVM) version 20231016 and later.

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

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

Purpose

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

For more information about StarWind VSAN architecture and available installation options, please refer to the:

[StarWind Virtual \(VSAN\) Getting Started Guide.](#)

Audience

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

Expected Result

Following this guide will result in a fully configured, highly available Windows Failover Cluster with virtual machine shared storage provided by StarWind VSAN.

NOTE: This guide applies to clusters with 2, 3, or more nodes. Be sure to follow the notes within each configuration step to ensure you perform the actions appropriate for your specific cluster size.

Introduction To Starwind Virtual San Cvm

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

Starwind Vsan System Requirements

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

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

Recommended RAID settings for HDD and SSD disks:


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

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

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

Pre-Configuring The Windows Server Hosts

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




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

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

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

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

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



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

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

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

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

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

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


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

Enabling Multipath Support

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

```
mpiocpl
```

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



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

10. Repeat the same procedure on the other server.


Installing File Server Roles

Please follow the steps below if file shares configuration is required

Scale-Out File Server (Sofs) For Application Data

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


below:



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

File Server For General Use With Smb Share


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



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

File Server For General Use With Nfs Share


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



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

Deploying Starwind Virtual San Cvm


1. Download the zip archive containing StarWind Virtual SAN CVM from <https://www.starwindsoftware.com/vsan#download>
2. Extract the virtual machine files on the server where StarWind Virtual SAN CVM will be deployed.
3. Open Hyper-V Manager and launch the Import Virtual Machine wizard.




4. On the Locate Folder step of the wizard, navigate to the location of the StarWind Virtual SAN CVM template folder, select the folder, and then click Next.

The screenshot shows a web-based wizard titled "Import Virtual Machine". The current step is "Locate Folder", which is highlighted in the left-hand navigation pane. The main area of the wizard contains the instruction "Specify the folder containing the virtual machine to import." Below this, there is a text input field labeled "Folder:" containing the path "C:\Users\Administrator\Downloads\StarWindAppliance\". To the right of the input field is a "Browse..." button. At the bottom of the wizard, there are four buttons: "< Previous", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

5. On the Select Virtual Machine step, click Next.



6. On the Choose Import Type step, select the Copy the virtual machine import type and click Next.




7. On the Choose Destination and Choose Storage Folders steps, specify folders to store virtual machine files, such as configuration, snapshots, smart paging, and virtual disks. Then, click Next.

NOTE: It's recommended to deploy StarWind Virtual SAN CVM on a separate storage device accessible to the hypervisor host, e.g., a BOSS card, etc.

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

16



The image shows a screenshot of the 'Import Virtual Machine' wizard in the StarWind Virtual SAN web UI. The window title is 'Import Virtual Machine'. The current step is 'Choose Folders to Store Virtual Hard Disks', which is indicated by a folder icon and the step name. On the left, a vertical navigation pane lists the steps: 'Before You Begin', 'Locate Folder', 'Select Virtual Machine', 'Choose Import Type', 'Choose Destination', 'Choose Storage Folders' (highlighted in blue), and 'Summary'. The main area of the wizard asks 'Where do you want to store the imported virtual hard disks for this virtual machine?'. Below this question, there is a text field labeled 'Location:' containing the path 'C:\Hyper-V\StarWindAppliance\Virtual Hard Disks\'. To the right of the text field is a 'Browse...' button. At the bottom of the wizard, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Import Virtual Machine


Choose Folders to Store Virtual Hard Disks

Before You Begin
Locate Folder
Select Virtual Machine
Choose Import Type
Choose Destination
Choose Storage Folders
Summary

Where do you want to store the imported virtual hard disks for this virtual machine?

Location: C:\Hyper-V\StarWindAppliance\Virtual Hard Disks\ Browse...

< Previous **Next >** Finish Cancel





8. On the Connect Network step, the wizard will validate the network settings.

Below are the default names for virtual switches:


- The Management network virtual switch is named "Management vSwitch"
- The Data/Heartbeat network virtual switch is named "Data/iSCSI vSwitch"
- The Replication network virtual switch is named "Replication/Sync vSwitch"

If the existing virtual switches have different names, manually select the appropriate ones from the drop-down list. Then, click Next.





9. On the Summary step, review the configuration settings and click Finish to begin the import process.



10. Repeat steps 1 through 9 on each host where StarWind CVM should be deployed.

Initial Configuration Wizard


1. Start the StarWind Virtual SAN Controller Virtual Machine.
2. Launch the VM console to view the VM boot process and obtain the IPv4 address of the Management network interface.

NOTE: If the VM does not acquire an IPv4 address from a DHCP server, use the Text-based User Interface (TUI) to set up the Management network manually.


Default credentials for TUI: user/rds123RDS

3. Using a web browser, open a new tab and enter the VM's IPv4 address to access the


StarWind VSAN Web Interface. On the Your connection is not private screen, click Advanced and then select Continue to...




4. On the Welcome to StarWind Appliance screen, click Start to launch the Initial Configuration Wizard.



5. On the License step, upload the StarWind Virtual SAN license file.



6. On the EULA step, read and accept the End User License Agreement to continue.



7. On the Management network step, review or edit the network settings and click Next.

IMPORTANT: The use of Static IP mode is highly recommended.

StarWind Appliance Initial configuration

- ✓ License
- ✓ EULA
- **Management network**
- Static hostname
- Administrator account
- Summary
- Configuration

Configure management network

Specify the unique IP address (static is recommended) and configure other network settings.

i The Management network is used to communicate with services such as DNS and NTP and to access the appliance web UI from external clients.

IP mode: **Static**

NIC	Model	Bandwidth	MAC address	IP address	Netmask	Gateway
ens160	82574L Gigabit Ne...	1 Gbit	00:50:56:9C:E...	192.168.12.206	255.255.254.0	192.168.12.1

Name servers (optional):

DNS 1: 192.168.12.17 DNS 2:


Time settings (optional):

NTP server: Time zone: UTC


Separate servers with commas, maximum 3 servers

Back **Next**


8. On the Static hostname, specify the hostname for the virtual machine and click Next.




9. On the Administrator account step, specify the credentials for the new StarWind Virtual SAN administrator account and click Next.




10. Wait until the Initial Configuration Wizard configures StarWind Virtual SAN for you.



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



12. After the configuration process is completed, click Finish to install the StarWind vCenter Plugin immediately, or uncheck the checkbox to skip this step and proceed to the Login page.




13. Repeat steps 1 through 12 on each Windows Server host.

Add Appliance


To create replicated, highly available storage, add partner appliances that use the same StarWind Virtual SAN license key.

1. Navigate to the Appliances page and click Add to open the Add appliance wizard.
2. On the Credentials step, enter the IP address and credentials for the partner StarWind Virtual SAN appliance, then click Next.



3. Wait for the connection to be established and the settings to be validated


4. On the Summary step, review the properties of the partner appliance, then click Add Appliance.




Configure Ha Networking

1. Navigate to the Network page and open the Configure HA networking wizard.
2. On the Appliances step, select either 2 partner appliances to configure two-way replication or 3 appliances for three-way replication, then click Next.

NOTE: The number of appliances in the cluster is limited by your StarWind Virtual SAN license.



3. Set Network Channels step combines configuration of Data and Replication Networks. Select the network interfaces designated to carry iSCSI or NVMe-oF storage traffic. Assign and configure at least one interface on each appliance with a static IP address in a unique network (subnet), specify the subnet mask and Cluster MTU size.
4. Select the network interfaces designated to carry the traffic for synchronous replication. Assign and configure at least one interface on each appliance with a static IP address in a unique network (subnet), specify the subnet mask and Cluster MTU size.




IMPORTANT: For a redundant, high-availability configuration, configure at least 2 network interfaces on each appliance. Ensure that the Data Network interfaces are interconnected between appliances through multiple direct links or via redundant switches.


5. Assign an MTU value on all selected network adapters, e.g., 1500 or 9000 bytes. If you are using network switches with the selected Data Network adapters, ensure that they are configured with the same MTU size value. In case of an MTU settings mismatch, stability and performance issues might occur on the whole setup.

NOTE: Setting MTU to 9000 bytes on some physical adapters (like Intel Ethernet Network Adapter X710, Broadcom network adapters, etc.) might cause stability and performance issues depending on the installed network driver. To avoid them, use 1500 bytes MTU size or install the stable version of the driver.


6. Once configured, click Next to validate network settings.




7. The warning might appear if a single data interface is configured. Click Yes, continue to proceed with the configuration.



8. Wait for the configuration to complete.



9. On the Summary step, review the specified network settings and click Configure to apply the changes.



Add Physical Disks


Attach physical storage to StarWind Virtual SAN Controller VM:

- Ensure that all physical drives are connected through an HBA or RAID controller.
- To get optimal storage performance, add HBA, RAID controllers, or NVMe SSD drives to StarWind CVM via a passthrough device.

For detailed instructions, refer to [Microsoft's documentation on DDA](#). Also, find the storage provisioning guidelines in the [KB article](#).

Create Storage Pool


1. Navigate to the Storage pools page and click the + button to open the Create storage pool wizard .



2. On the Appliance step, select partner appliances on which to create new storage pools, then click Next.


NOTE: Select 2 appliances for configuring storage pools if you are deploying a two-node cluster with two-way replication, or select 3 appliances for configuring a three-node

cluster with a three-way mirror.




3. On the Physical disks step, select physical disks to be pooled on each node, then click Next.

IMPORTANT: Select an identical type and number of disks on each appliance to create storage pools with a uniform configuration.



4. On the Profile step, select one of the preconfigured storage profiles, or choose Manual to configure the storage pool manually based on your redundancy, capacity, and performance requirements, then click Next.




NOTE: Hardware RAID, Linux Software RAID, and ZFS storage pools are supported. To

simplify the configuration of storage pools, preconfigured storage profiles are provided. These profiles recommend a pool type and layout based on the attached storage:

- High capacity – creates Linux Software RAID-5 to maximize storage capacity while maintaining redundancy.
- High performance – creates Linux Software RAID-10 to maximize storage performance while maintaining redundancy.
- Hardware RAID – configures a hardware RAID virtual disk as a storage pool. This option is available only if a hardware RAID controller is passed through to the StarWind Virtual SAN.
- Better redundancy – creates ZFS Striped RAID-Z2 (RAID 60) to maximize redundancy while maintaining high storage capacity.
- Manual – allows users to configure any storage pool type and layout with the attached storage.

5. On the Summary step, review the storage pool settings and click Create to configure new storage pools on the selected appliances.




NOTE: The storage pool configuration may take some time, depending on the type of pooled storage and the total storage capacity. Once the pools are created, a notification will appear in the upper right corner of the Web UI.

IMPORTANT for CVM only: In some cases, additional tweaks are required to optimize the storage performance of the disks added to the Controller Virtual Machine. Please follow the steps in this KB to change the scheduler type depending on the disks type:
<https://knowledgebase.starwindsoftware.com/guidance/starwind-vsan-for-vsphere-changi>

[ng-linux-i-o-scheduler-to-optimize-storage-performance/](#)


Create Volume

1. Navigate to the Volumes page and click the + button to open the Create volume wizard.




2. On the Storage pool step, select partner appliances on which to create new volumes, then click Next.


NOTE: Select 2 appliances for configuring volumes if you are deploying a two-node cluster with two-way replication, or select 3 appliances for configuring a three-node cluster with a three-way mirror.




3. On the Settings step, specify the volume name and size, then click Next.



4. On the Filesystem type step, select Standard, then click Next.



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




Create Ha Lun Using Webui


This section describes how to create LUN in Web UI. This option is available for the setups with Commercial, Trial, and NFR licenses applied.

For setups with a Free license applied, the PowerShell script should be used to create the LUN – please follow the steps described in the section: Create StarWind HA LUNs using PowerShell

1. Navigate to the LUNs page and click the + button to open the Create LUN wizard.




2. On the Protocols step, select the preferred storage protocol and click Next.



3. On the LUN availability step, select the High availability and click Next.


NOTE: The availability options for a LUN can be Standalone (without replication) or High Availability (with 2-way or 3-way replication), and are determined by the StarWind Virtual SAN license.

Below are the steps for creating a high-availability iSCSI LUN.




4. On the Appliances step, select partner appliances that will host new LUNs and click Next.

IMPORTANT: Selected partner appliances must have identical hardware configurations, including CPU, RAM, storage, and networking.




5. On the Volumes step, select the volumes for storing data on the partner appliances and click Next.

IMPORTANT: For optimal performance, the selected volumes must have identical underlying storage configurations.




6. On the Failover strategy step, select the preferred failover strategy and click Next.

NOTE: The failover strategies for a LUN can be Heartbeat or Node Majority. In case of 2-nodes setup and None Majority failover strategy, Node witness (requires an additional third witness node), or File share witness (requires an external file share) should be configured. These options are determined by StarWind Virtual SAN license and setup configuration. Below are the steps for configuring the Heartbeat failover strategy in a two-node cluster.




7. On the LUN settings step, specify the LUN name, size, block size, then click Next.

NOTE: For high-availability configurations, ensure that MPIO checkbox is selected.



8. On the Summary step, review the LUN settings and click Create to configure new LUNs on the selected volumes.



Creating Starwind Ha Luns Using Powershell

1. Open PowerShell ISE as Administrator.
2. Open StarWindX sample CreateHA_2.ps1 using PowerShell ISE. It can be found here: C:\Program Files\StarWind Software\StarWind\StarWindX\Samples\

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

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

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

The script does not create a directory. Make sure you create the directory listed as \$imagePath value before running the script.

Make sure to open 3261 and 3260 ports.

The same approach applies to CreateHA_3.ps1 which allow to create a 3-way replica HA

device.

3. Configure script parameters according to the following example:

```
param($addr="192.168.0.1", $port=3261, $user="root",
$password="starwind",
    $addr2="192.168.0.2", $port2=$port, $user2=$user,
$password2=$password,
#common
    $initMethod="NotSynchronize",
    $size=1048576,
    $sectorSize=512,
    $failover=0,
    $bmpType=1,
    $bmpStrategy=0,
#primary node
    $imagePath="/mnt/sdb1/volume1",
    $imageName="device1",
    $createImage=$true,
    $storageName="",
    $targetAlias="target1",
    $poolName="pool1",
    $syncSessionCount=1,
    $luaOptimized=$true,
    $cacheMode="none",
    $cacheSize=0,
    $syncInterface="#p2=172.16.20.2:3260,172.16.21.2:3260",
    $hbInterface="#p2=172.16.10.2:3260,192.168.0.2:3260",
    $createTarget=$true,
    $bmpFolderPath="",
#secondary node
    $imagePath2="/mnt/sdb1/volume1",
    $imageName2="device1",
    $createImage2=$true,
    $storageName2="",
    $targetAlias2="target1",
    $poolName2="pool1",
    $syncSessionCount2=1,
    $luaOptimized2=$false,
    $cacheMode2=$cacheMode,
    $cacheSize2=$cacheSize,
    $syncInterface2="#p1=172.16.20.1:3260,172.16.21.1:3260",
```

```

    $hbInterface2="#p1=172.16.10.1:3260,192.168.0.1:3260",
    $createTarget2=$true,
    $bmpFolderPath=""
  )
Import-Module StarWindX

try
{
    Enable-SWXLog -level SW_LOG_LEVEL_DEBUG

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

    $server.Connect()

    $firstNode = new-Object Node

    $firstNode.HostName = $addr
    $firstNode.HostPort = $port
    $firstNode.Login = $user
    $firstNode.Password = $password
    $firstNode.ImagePath = $imagePath
    $firstNode.ImageName = $imageName
    $firstNode.Size = $size
    $firstNode.CreateImage = $createImage
    $firstNode.StorageName = $storageName
    $firstNode.TargetAlias = $targetAlias
    $firstNode.AutoSynch = $autoSynch
    $firstNode.SyncInterface = $syncInterface
    $firstNode.HBInterface = $hbInterface
    $firstNode.PoolName = $poolName
    $firstNode.SyncSessionCount = $syncSessionCount
    $firstNode.ALUAOptimized = $aluaOptimized
    $firstNode.CacheMode = $cacheMode
    $firstNode.CacheSize = $cacheSize
    $firstNode.FailoverStrategy = $failover
    $firstNode.CreateTarget = $createTarget
    $firstNode.BitmapStoreType = $bmpType
    $firstNode.BitmapStrategy = $bmpStrategy
    $firstNode.BitmapFolderPath = $bmpFolderPath
    #
    # device sector size. Possible values: 512 or 4096(May
be incompatible with some clients!) bytes.
    #

```

```

    $firstNode.SectorSize = $sectorSize
    $secondNode = new-Object Node

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

        Start-Sleep -m 2000

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

```

```
}
```

Detailed explanation of script parameters:

-addr, -addr2 — partner nodes IP address.

Format: string. Default value: 192.168.0.1, 192.168.0.2

allowed values: localhost, IP-address

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

Format: string. Default value: 3261

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

Format: string. Default value: root

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

Format: string. Default value: starwind

#common

-initMethod – set the initial synchronization option.

Format: string.

Values:

Clear – default

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

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

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

Format: integer. Default value: 12

-sectorSize – set sector size for HA-device

Format: integer. Default value: 512

allowed values: 512, 4096

-failover – set type failover strategy

Format: integer. Default value: 0 (Heartbeat)

allowed values: 0, 1 (Node Majority)

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

Format: integer. Default value: 1 (RAM)

allowed values: 1, 2 (DISK)

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

Format: integer. Default value: 0

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

-storageName is used only if you plan to add the partner to the existing device. For

CreateHA_2.ps1 use, leave it as is.

#primary node

-imagePath – set path to store the device file

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

-imageName – set name device

Format: string. Default value: masterImg21

-createImage – set create image file

Format: boolean. Default value: true

-targetAlias – set alias for target

Format: string. Default value: targetha21

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

Format: string. Default value: pool1

-aluaOptimized – set Alua Optimized

Format: boolean. Default value: true

-cacheMode – set type L1 cache (optional parameter)

Format: string. Default value: wb

allowed values: none, wb, wt

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

Format: integer. Default value: 128

allowed values: 1 and more

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

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

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

Format: string. Default value: ""

-createTarget – set creating target

Format: string. Default value: true

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

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

-bmpFolderPath – set path to save bitmap file

Format: string.

#secondary node

-imagePath2 – set path to store the device file

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

-imageName2 – set name device

Format: string. Default value: masterImg21

-createImage2 – set create image file

Format: boolean. Default value: true

-targetAlias2 – set alias for targetFormat: string.
Default value: targetha22

-poolName2 – set storage pool. Do not change it and keep default value.
Format: string. Default value: pool1

-aluaOptimized2 – set Alua Optimized
Format: boolean. Default value: true

-cacheMode2 – set type L1 cache (optional parameter)
Format: string. Default value: wb
allowed values: wb, wt

-cacheSize2 – set size for L1 cache in MB (optional parameter)
Format: integer. Default value: 128
allowed values: 1 and more

-syncInterface2 – set sync channel IP-address from partner node
Format: string. Default value: “#p1={0}:3260”

-hbInterface2 – set heartbeat channel IP-address from partner node
Format: string. Default value: “”

-createTarget2 – set creating target
Format: string. Default value: true
Even if you do not specify the parameter -createTarget, the target will be created automatically. If the parameter is set as -createTarget \$false, then an attempt will be made to create the device with existing targets, the names of which are specified in the -targetAlias (targets must already be created)

-bmpFolderPath2 – set path to save bitmap file
Format: string.

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

1. Stop StarWind Service:

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

2. Go to /opt/starwind/starwind-virtual-san/drive_c/starwind/headers and delete the headers you do not need.
3. Go to the underlying storage specified as \$imagePath and delete the header and imagefile there.
4. Go to the folder with StarWind.cfg (/opt/starwind/starwind-virtual-san/drive_c/starwind/StarWind.cfg) and copy it.
5. Edit StarWind.cfg:

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

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

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

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

Selecting The Failover Strategy

StarWind provides 2 options for configuring a failover strategy:

Heartbeat

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

If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The device with the lowest priority will be marked as not synchronized and get subsequently blocked for the further read and write operations until the synchronization channel resumption. At the same time, the partner device on the synchronized node flushes data from the cache to the disk to preserve data integrity in case the node goes down unexpectedly. It is recommended to assign more independent heartbeat channels during the replica creation to improve system stability and avoid the “split-brain” issue.

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

Node Majority

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

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

In case of a two-node HA storage, all nodes will be disconnected if there is a problem on

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

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

Please select the required option:


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

Discovery of StarWind iSCSI targets on Windows Server hosts


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

```
iscsicpl
```

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



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




The image shows a 'Discover Target Portal' dialog box. It has a title bar with a close button (X). The main text area contains instructions: 'Enter the IP address or DNS name and port number of the portal you want to add.' and 'To change the default settings of the discovery of the target portal, click the Advanced button.' Below this, there are two input fields: 'IP address or DNS name:' with the value '172.16.10.10' and 'Port: (Default is 3260.)' with the value '3260'. At the bottom, there are three buttons: 'Advanced...', 'OK', and 'Cancel'.

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


The screenshot shows the 'Advanced Settings' dialog box with the 'IPsec' tab selected. The 'Connect using' section has three dropdown menus: 'Local adapter' set to 'Microsoft iSCSI Initiator', 'Initiator IP' set to '172.16.10.1', and 'Target portal IP' which is empty. The 'CRC / Checksum' section has two unchecked checkboxes: 'Data digest' and 'Header digest'. The 'Enable CHAP log on' checkbox is also unchecked. The 'CHAP Log on information' section contains a text box for 'Name' with the value 'iqn.1991-05.com.microsoft:sw01' and an empty 'Target secret' text box. Below this, there are three more unchecked checkboxes: 'Perform mutual authentication', 'Use RADIUS to generate user authentication credentials', and 'Use RADIUS to authenticate target credentials'. The dialog box has 'OK', 'Cancel', and 'Apply' buttons at the bottom right.

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



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

172.16.10.20), then click Advanced...




The image shows a 'Discover Target Portal' dialog box with a close button (X) in the top right corner. Inside the dialog, there is instructional text: 'Enter the IP address or DNS name and port number of the portal you want to add.' and 'To change the default settings of the discovery of the target portal, click the Advanced button.' Below this text are two input fields. The first field is labeled 'IP address or DNS name:' and contains the text '172.16.10.20'. The second field is labeled 'Port: (Default is 3260.)' and contains the text '3260'. At the bottom of the dialog, there are three buttons: 'Advanced...', 'OK', and 'Cancel'. The 'OK' button is highlighted with a blue border.

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

The screenshot shows the 'Advanced Settings' dialog box with the 'IPsec' tab selected. The 'Connect using' section has three dropdown menus: 'Local adapter' set to 'Microsoft iSCSI Initiator', 'Initiator IP' set to '172.16.10.1', and 'Target portal IP' is empty. The 'CRC / Checksum' section has two unchecked checkboxes: 'Data digest' and 'Header digest'. The 'Enable CHAP log on' checkbox is also unchecked. The 'CHAP Log on information' section contains a text box for 'Name' with the value 'iqn.1991-05.com.microsoft:sw01' and an empty 'Target secret' text box. Below this, there are three more unchecked checkboxes: 'Perform mutual authentication', 'Use RADIUS to generate user authentication credentials', and 'Use RADIUS to authenticate target credentials'. At the bottom right are 'OK', 'Cancel', and 'Apply' buttons.

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




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


Connecting StarWind iSCSI targets to Windows Server hosts

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

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



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



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

Advanced Settings ? X

General **IPsec**

Connect using

Local adapter: Microsoft iSCSI Initiator

Initiator IP: 172.16.10.1

Target portal IP: 172.16.10.10 / 3260

CRC / Checksum

☐ Data digest ☐ Header digest

☐ Enable CHAP log on

CHAP Log on information

CHAP helps ensure connection security by providing authentication between a target and an initiator.

To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.

Name: iqn.1991-05.com.microsoft:sw01

Target secret:

☐ Perform mutual authentication

To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.

☐ Use RADIUS to generate user authentication credentials

☐ Use RADIUS to authenticate target credentials

OK Cancel Apply


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

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

The screenshot shows the 'Advanced Settings' dialog box with the 'IPsec' tab selected. The 'Connect using' section has three dropdown menus: 'Local adapter' set to 'Microsoft iSCSI Initiator', 'Initiator IP' set to '172.16.10.1', and 'Target portal IP' set to '172.16.10.20 / 3260'. The 'CRC / Checksum' section has two unchecked checkboxes: 'Data digest' and 'Header digest'. The 'Enable CHAP log on' checkbox is also unchecked. The 'CHAP Log on information' section contains a text box for 'Name' with the value 'iqn.1991-05.com.microsoft:sw01' and an empty 'Target secret' text box. Below this, there are three more unchecked checkboxes: 'Perform mutual authentication', 'Use RADIUS to generate user authentication credentials', and 'Use RADIUS to authenticate target credentials'. At the bottom right are 'OK', 'Cancel', and 'Apply' buttons.

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

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



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

The screenshot shows the 'Advanced Settings' dialog box with the 'IPsec' tab selected. The 'Connect using' section has three dropdown menus: 'Local adapter' set to 'Microsoft iSCSI Initiator', 'Initiator IP' set to '172.16.10.1', and 'Target portal IP' set to '172.16.10.10 / 3260'. The 'CRC / Checksum' section has two unchecked checkboxes: 'Data digest' and 'Header digest'. The 'Enable CHAP log on' checkbox is also unchecked. The 'CHAP Log on information' section contains a text box for 'Name' with the value 'iqn.1991-05.com.microsoft:sw01' and an empty text box for 'Target secret'. Below this, there are three more unchecked checkboxes: 'Perform mutual authentication', 'Use RADIUS to generate user authentication credentials', and 'Use RADIUS to authenticate target credentials'. At the bottom right are 'OK', 'Cancel', and 'Apply' buttons.

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

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

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

The screenshot shows the 'Advanced Settings' dialog box with the 'IPsec' tab selected. The 'Connect using' section has three dropdown menus: 'Local adapter' set to 'Microsoft iSCSI Initiator', 'Initiator IP' set to '172.16.10.1', and 'Target portal IP' set to '172.16.10.20 / 3260'. The 'CRC / Checksum' section has two unchecked checkboxes: 'Data digest' and 'Header digest'. The 'Enable CHAP log on' checkbox is also unchecked. The 'CHAP Log on information' section contains a text box for 'Name' with the value 'iqn.1991-05.com.microsoft:sw01' and an empty text box for 'Target secret'. Below this, there are three more unchecked checkboxes: 'Perform mutual authentication', 'Use RADIUS to generate user authentication credentials', and 'Use RADIUS to authenticate target credentials'. At the bottom right are 'OK', 'Cancel', and 'Apply' buttons.

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

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


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

Configuring Multipath

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

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

1. Configure the MPIO policy for each target with the load balance policy of choice. Select the Target located on the local server and click Devices.
2. In the Devices dialog, click MPIO.




3. Select the appropriate load balancing policy.

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

Connecting Disks to Servers

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




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


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

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

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



5. Right-click on the selected disk and choose New Simple Volume.
6. In New Simple Volume Wizard, indicate the volume size. Click Next.
7. Assign a drive letter to the disk. Click Next.



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

New Simple Volume Wizard [X]

Format Partition
To store data on this partition, you must format it first.

Choose whether you want to format this volume, and if so, what settings you want to use.

☐ Do not format this volume

☒ Format this volume with the following settings:

File system: NTFS

Allocation unit size: Default

Volume label: CSV1


☒ Perform a quick format

☐ Enable file and folder compression

< Back Next > Cancel

9. Press Finish to complete.

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



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

12. Repeat steps from the 2nd to bring all the remaining StarWind disks online.

Creating A Failover Cluster In Windows Server

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


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

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


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

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

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



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



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

Create Cluster Wizard

Select Servers

Before You Begin
Select Servers
 Validation Warning
 Access Point for Administering the Cluster
 Confirmation
 Creating New Cluster
 Summary


Add the names of all the servers that you want to have in the cluster. You must add at least one server.

Enter server name:

Selected servers:


- SW1.starwind.local
- SW2.starwind.local

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.



Create Cluster Wizard

Access Point for Administering the Cluster

Before You Begin
Select Servers
Access Point for Administering the Cluster
Confirmation
Creating New Cluster
Summary

Type the name you want to use when administering the cluster.


Cluster Name:

i The NetBIOS name is limited to 15 characters. One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.

	Networks	Address
☒	192.168.12.0/23	192.168.12.86

< Previous **Next >** Cancel

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



Create Cluster Wizard

Confirmation

Before You Begin
Select Servers
Access Point for Administering the Cluster
Confirmation
Creating New Cluster
Summary

You are ready to create a cluster.
The wizard will create your cluster with the following settings:

Cluster	Production
Node	SW1.starwind.local SW2.starwind.local
Cluster registration	DNS and Active Directory Domain Services
IP Address	192.168.12.86


☐ Add all eligible storage to the cluster.

To continue, click Next.

< Previous **Next >** Cancel


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.

Configure Cluster Quorum Wizard

Select Quorum Configuration Option

Before You Begin
Select Quorum Configuration Option
 Select Quorum Witness
 Confirmation
 Configure Cluster Quorum Settings
 Summary

Select a quorum configuration for your cluster.

☐ Use default quorum configuration
 The cluster determines quorum management options, including the quorum witness.


☒ Select the quorum witness
 You can add or change the quorum witness. The cluster determines the other quorum management options.

☐ Advanced quorum configuration
 You determine the quorum management options, including the quorum witness.


[Failover Cluster Quorum and Witness Configuration Options](#)

< Previous **Next >** Cancel


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.



The screenshot shows the 'High Availability Wizard' window, specifically the 'Client Access Point' step. On the left is a navigation pane with the following items: 'Before You Begin', 'Select Role', 'File Server Type', 'Client Access Point' (which is highlighted in blue), 'Confirmation', 'Configure High Availability', and 'Summary'. The main area of the wizard has a title bar with a close button (X) and a subtitle 'Client Access Point'. Below the subtitle, it says 'Type the name that clients will use when accessing this clustered role:'. There is a text input field labeled 'Name:' containing the text 'FileServer'. Below the input field, there is an information icon (i) followed by the text: 'The NetBIOS name is limited to 15 characters. All networks were configured automatically.' At the bottom right of the main area, there are three buttons: '< Previous', 'Next >' (which is highlighted with a blue border), and 'Cancel'.


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.

Select the server and path for this share

Select Profile
Share Location
 Share Name
 Other Settings
 Permissions
 Confirmation
 Results

Server:

Server Name	Status	Cluster Role	Owner Node
FileServer	Online	Scale-Out File...	

Share location:

☒ Select by volume:

Volume	Free Space	Capacity	File System
C:\ClusterStorage\Volume1	5.92 GB	5.97 GB	CSVFS
C:\ClusterStorage\Volume2	9.91 GB	9.97 GB	CSVFS

The location of the file share will be a new folder in the \Shares directory on the selected volume.

☐ Type a custom path:

< Previous Next > Create Cancel

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.

New Share Wizard

Configure share settings

- Select Profile
- Share Location
- Share Name
- Other Settings**
- Permissions
- Confirmation
- Results

☐ **Enable access-based enumeration**
 Access-based enumeration displays only the files and folders that a user has permissions to access. If a user does not have Read (or equivalent) permissions for a folder, Windows hides the folder from the user's view.

☒ **Enable continuous availability**
 Continuous availability features track file operations on a highly available file share so that clients can fail over to another node of the cluster without interruption.


☐ **Allow caching of share**
 Caching makes the contents of the share available to offline users. If the BranchCache for Network Files role service is installed, you can enable BranchCache on the share.

☐ **Enable BranchCache on the file share**
 BranchCache enables computers in a branch office to cache files downloaded from this share, and then allows the files to be securely available to other computers in the branch.

Encrypt data access
 When enabled, remote file access to this share will be encrypted. This secures the data against unauthorized access while the data is transferred to and from the share. If this box is checked and grayed out, an administrator has turned on encryption for the entire server.

< Previous Next > Create Cancel


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.



The screenshot shows the 'High Availability Wizard' window, specifically the 'Client Access Point' step. On the left is a navigation pane with steps: 'Before You Begin', 'Select Role', 'File Server Type', 'Client Access Point' (highlighted), 'Select Storage', 'Confirmation', 'Configure High Availability', and 'Summary'. The main area is titled 'Client Access Point' and contains the instruction: 'Type the name that clients will use when accessing this clustered role:'. Below this is a text box labeled 'Name:' containing 'FileServer'. An information icon (i) is followed by a note: 'The NetBIOS name is limited to 15 characters. One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.' Below the note is a table with two columns: 'Networks' and 'Address'. The first row shows a checked checkbox, the network '192.168.12.0/24', and the address '192.168.12.85'. At the bottom right are three buttons: '< Previous', 'Next >' (highlighted), and 'Cancel'.

High Availability Wizard

Client Access Point

Before You Begin
Select Role
File Server Type
Client Access Point
Select Storage
Confirmation
Configure High Availability
Summary

Type the name that clients will use when accessing this clustered role:

Name:


i The NetBIOS name is limited to 15 characters. One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.

	Networks	Address
☒	192.168.12.0/24	192.168.12.85


< Previous **Next >** Cancel

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.



New Share Wizard

Select the server and path for this share

Select Profile

- Share Location
- Share Name
- Other Settings
- Permissions
- Confirmation
- Results

Server:

Server Name	Status	Cluster Role	Owner Node
FileServer	Online	File Server	

Share location:


☒ Select by volume:

Volume	Free Space	Capacity	File System
G:	9.91 GB	9.97 GB	NTFS

The location of the file share will be a new folder in the \Shares directory on the selected volume.

☐ Type a custom path:

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



New Share Wizard

Specify share name

Select Profile

- Share Location
- Share Name
- Other Settings
- Permissions
- Confirmation
- Results

Share name:

Share description:


Local path to share:

Remote path to share:


Local path to share:
G:\Shares\Share
i If the folder does not exist, the folder is created.

Remote path to share:
\\FileServer\Share


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.

New Share Wizard

Specify share name

Select Profile
Share Location
Share Name
Authentication
Share Permissions
Permissions
Confirmation
Results

Share name:


Local path to share:

If the folder does not exist, the folder is created.

Remote path to share:

< Previous Next > Create Cancel

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



Add Permissions

Grant permissions to access the NFS share to a host, client group, or netgroup. Select the access and language encoding for the share.

☐ Host:

☐ Netgroup:

☐ Client group:

☒ All Machines


Language encoding: ANSI

Share permissions: No Access


☐ Allow root access (not recommended)

Add Cancel


9. Specify the access permissions for the file share.



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




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



To manage created NFS File Shares:

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



Contacts

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

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

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

Sales: sales@starwind.com

General Information: info@starwind.com



StarWind Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA
www.starwind.com ©2026, StarWind Software Inc. All rights reserved.