

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

2024

TECHNICAL PAPERS



Trademarks

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

Changes

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

Technical Support and Services

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

About StarWind

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

Copyright ©2009-2018 StarWind Software Inc.

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

Annotation

Relevant products

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

Purpose

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

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

Audience

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

Expected result

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

Prerequisites

StarWind Virtual SAN system requirements

Prior to installing StarWind Virtual SAN, please make sure that the system meets the requirements, which are available via the following link:
<https://www.starwindsoftware.com/system-requirements>

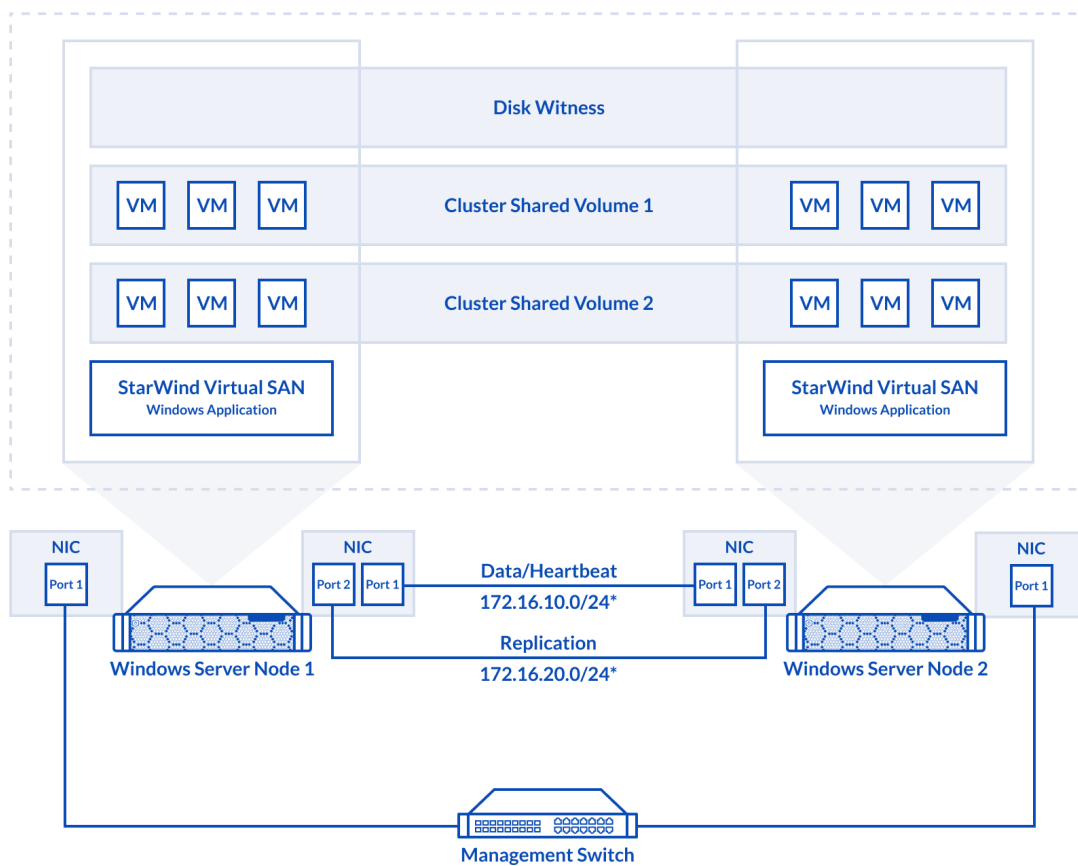
Recommended RAID settings for HDD and SSD disks:

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


Please read StarWind Virtual SAN Best Practices document for additional information:
<https://www.starwindsoftware.com/resource-library/starwind-virtual-san-best-practices>

Solution diagram

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



2-node cluster



3-node cluster

Preconfiguring cluster nodes

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

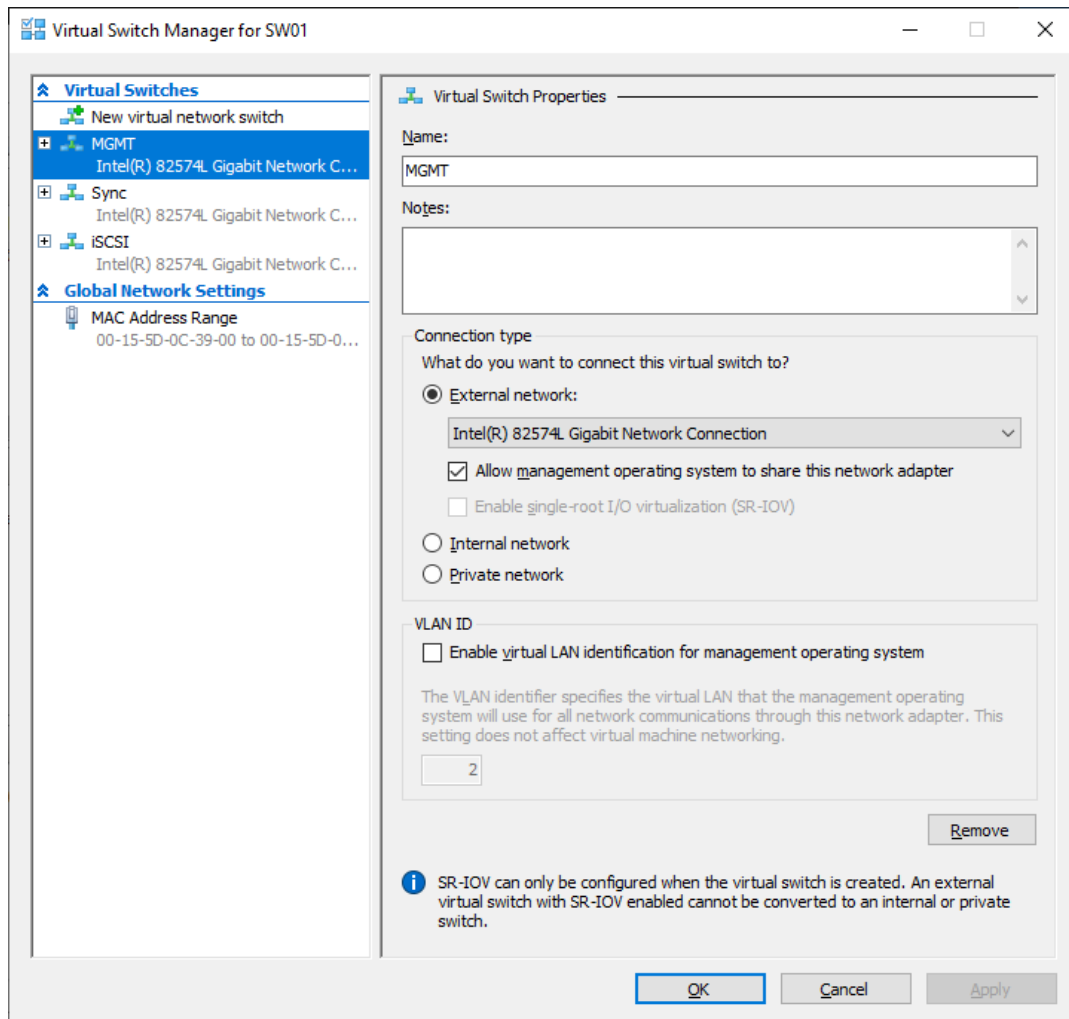
NOTE: Please follow the recommendation in [KB article](#) on how to place a DC in case of StarWind Virtual SAN usage.

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

3. Define at least 2x network interfaces (2 node scenario) or 4x network interfaces (3 node scenario) on each node that will be used for the Synchronization and iSCSI/StarWind heartbeat traffic. Do not use iSCSI/Heartbeat and Synchronization channels over the same physical link. Synchronization and iSCSI/Heartbeat links can be connected either via redundant switches or directly between the nodes (see diagram above).

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, 172.16.1x.x subnets are used for iSCSI/StarWind heartbeat traffic, while 172.16.2x.x subnets are 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.

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.


11. Enable Remote Desktop connection to the servers and join them to the domain by selecting the corresponding option in the Server Configuration window.



NOTE: Rename and configure a static IP address for network adapters.

12. Run powershell.exe in the Command Prompt to check the network adapters availability in the system:

Get-NetAdapter



13. To change the name and set the static IP for Heartbeat/iSCSI and Synchronization channel, run the next commands via PowerShell:

```
Get-NetAdapter "Ethernet1" | Rename-NetAdapter -NewName "Sync"
```

```
Get-NetAdapter "iSCSI" | New-NetIPAddress -IPAddress
```

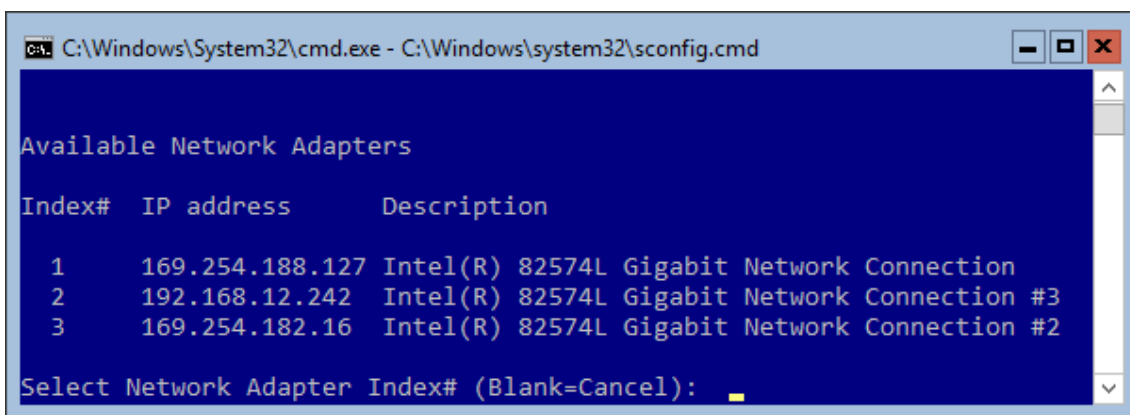
```
172.16.10.10 -PrefixLength 24

Get-NetAdapter "Ethernet2" | Rename-NetAdapter -NewName "Sync"

Get-NetAdapter "Sync" | New-NetIPAddress -IPAddress
172.16.20.10 -PrefixLength 24
```

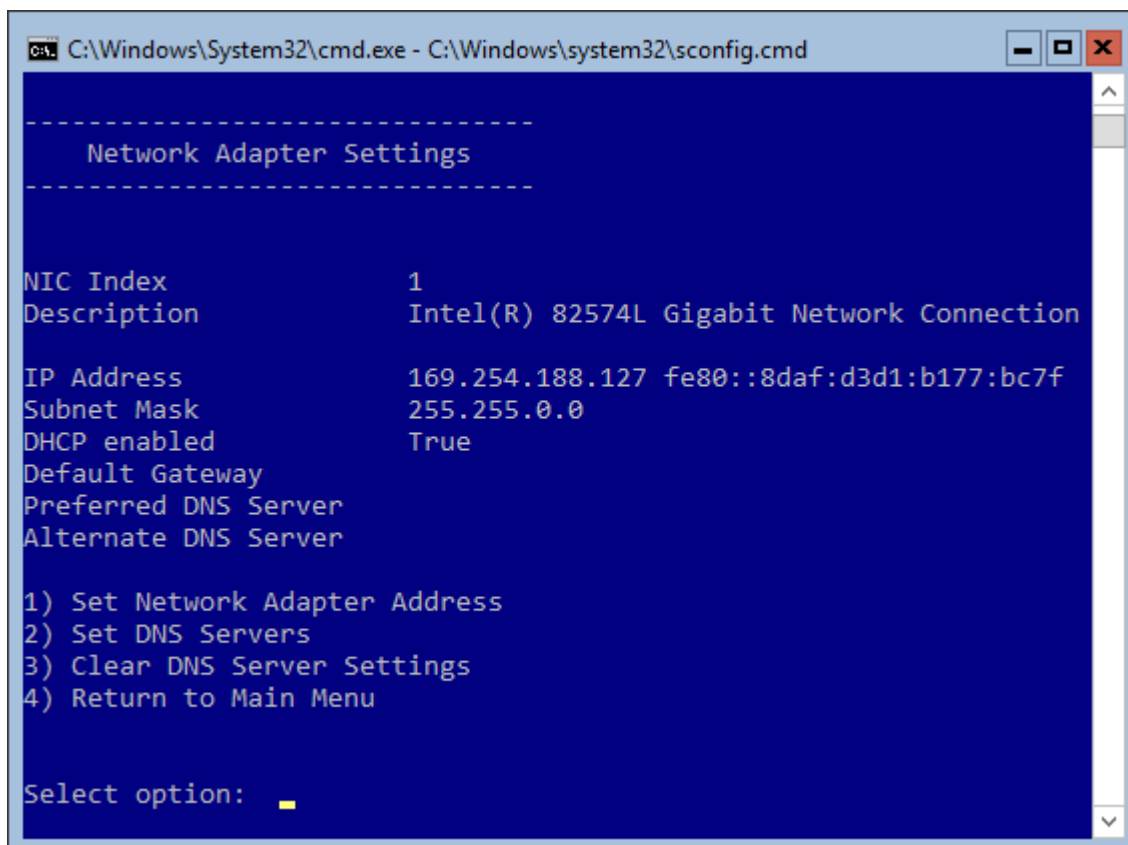
NOTE: The corresponding IP addresses should be configured on the partner node.

14. Alternatively, the network settings can be changed through the sconfig.cmd window. In Network Settings (option 8), select the Index of the NIC which should be edited:



The following actions are possible:

- set Network Adapter Address - selection between DHCP or Static IP (recommended);
- set DNS Servers - providing DNS settings;
- clear DNS Server Settings;
- return to Main Menu.



15. It is highly recommended to enable jumbo frames (9014) on the Synchronization and iSCSI networks. This can be done via PowerShell in two ways:


- Directly on the Synchronization / iSCSI adapter:

```
Set-NetAdapterAdvancedProperty -Name "Sync" -RegistryKeyword
"*JumboPacket" -Registryvalue 9014
```

```
Set-NetAdapterAdvancedProperty -Name "iSCSI" -RegistryKeyword
"*JumboPacket" -Registryvalue 9014
```

- For all adapters available in the system:

```
Set-NetAdapterAdvancedProperty -Name "*" -RegistryKeyword
"*JumboPacket" -Registryvalue 9014
```



16. Ping each node with jumbo frames (change the asterisk to the corresponding partner node's IP address):

```
ping 172.16.20.* -f -l 8900 - for Synchronization network;
ping 172.16.10.* -f -l 8900 - for iSCSI network
```

17. To create a virtual switch on the Management interface, run the next command via PowerShell:

```
New-VMSwitch -Name "vSwitch" -NetAdapterName "management"
```

NOTE: The Virtual Switch name must be the same on both nodes.

18. To disable firewall, please run the following command via PowerShell on each server:

```
Get-NetFirewallProfile | Set-NetFirewallProfile -Enabled False
```

19. Install Failover Clustering and Multipath I/O features on both servers using PowerShell:

```
Install-WindowsFeature Failover-Clustering
-IncludeAllSubFeature -Restart

Enable-WindowsOptionalFeature -Online -FeatureName MultiPathIO

Enable-MSDSMAutomaticClaim -BusType iSCSI
```

20. Preparing storage for StarWind devices

NOTE: Please refer to the KB article about recommended RAID settings before

proceeding:

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

Any storage array intended to be used by StarWind Virtual SAN for storing virtual disk images should meet the following requirements:

- initialized as GPT;
- have a single NTFS-formatted partition;
- have a drive letter assigned.

21. To create a Local Partition for the storage drive, run the commands below in the CMD window:

```
Diskpart

list disk

select disk X //where X is the number of the disk to be
processed

online disk

attributes disk clear readonly

convert GPT

create partition Primary

format fs=ntfs label=X quick //where X is the name for the
Volume

list volume


select volume X

assign letter X //where X is the letter for the Volume


list Volume
```

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.



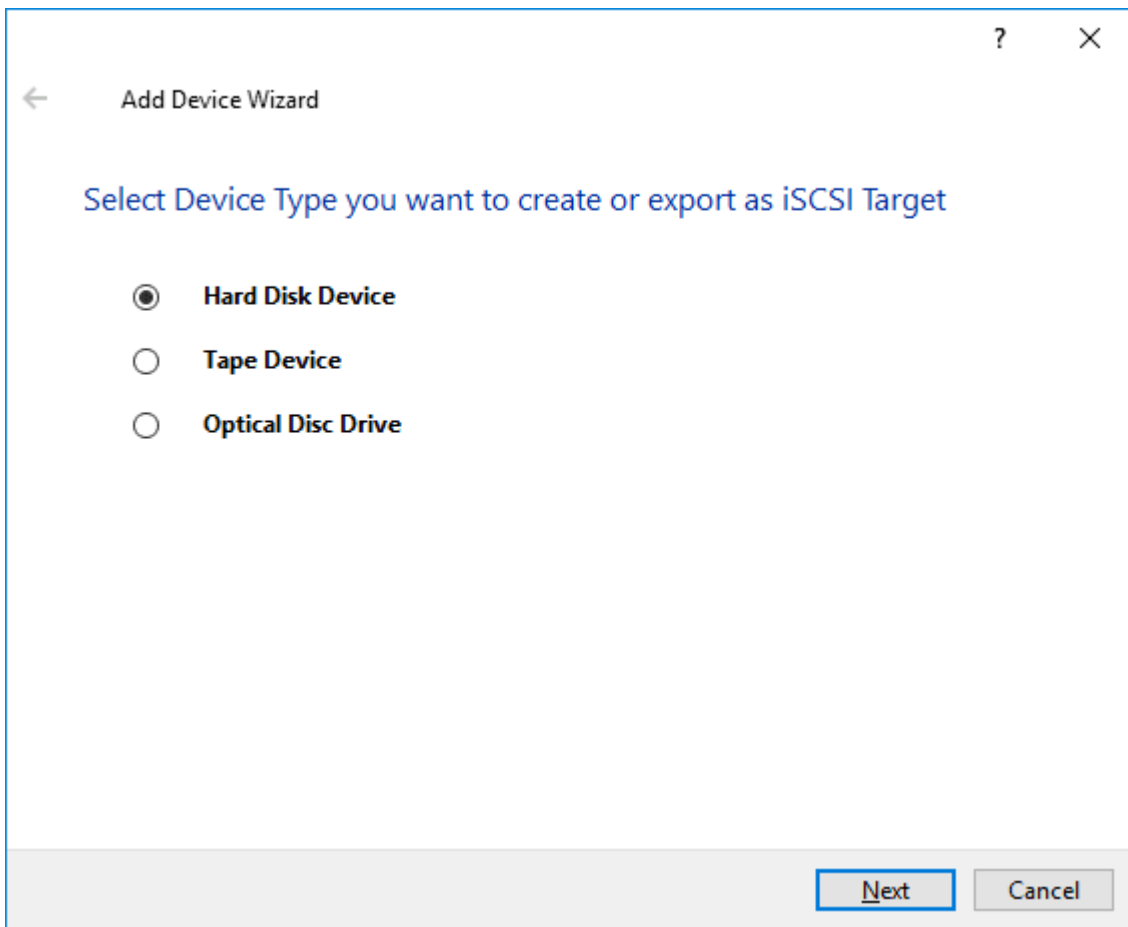
7. Enable the checkbox if a desktop icon needs to be created. Click on Next to continue.

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.



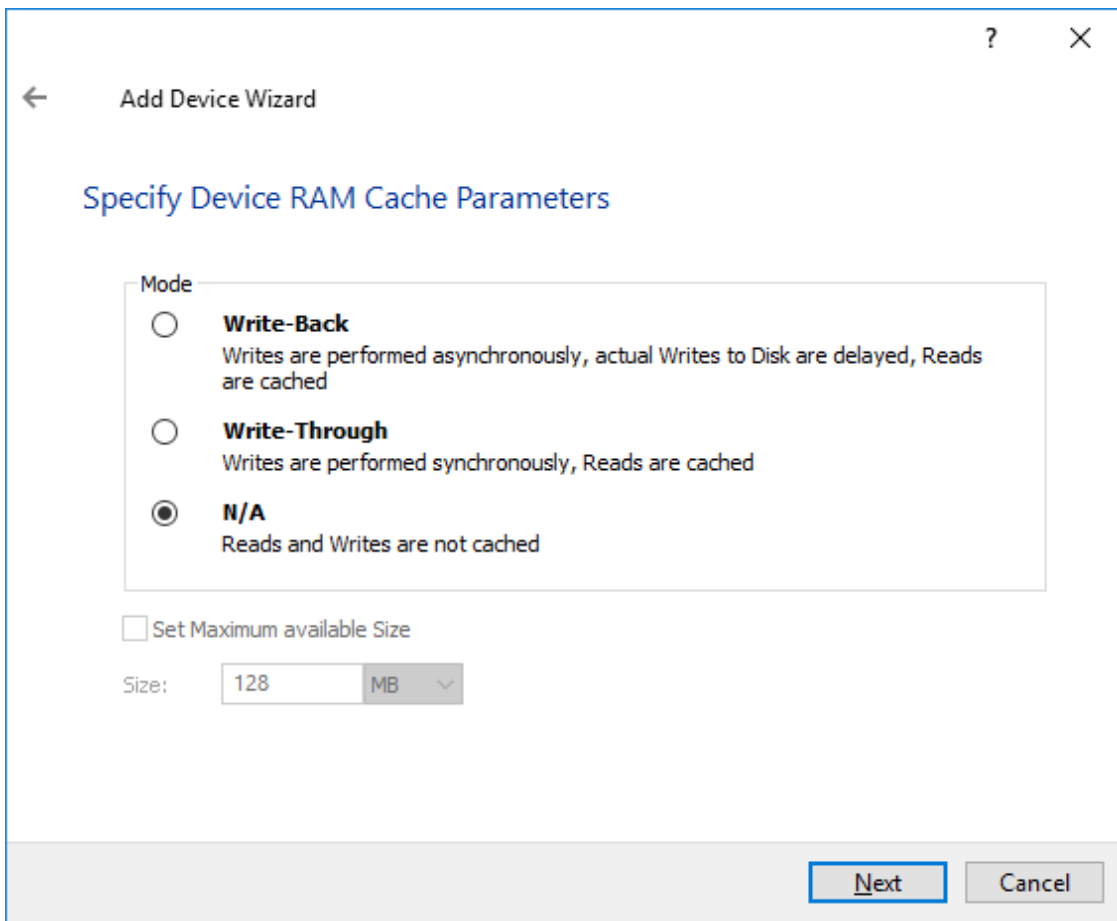
4. Specify a virtual disk Name, Location, and Size.




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).

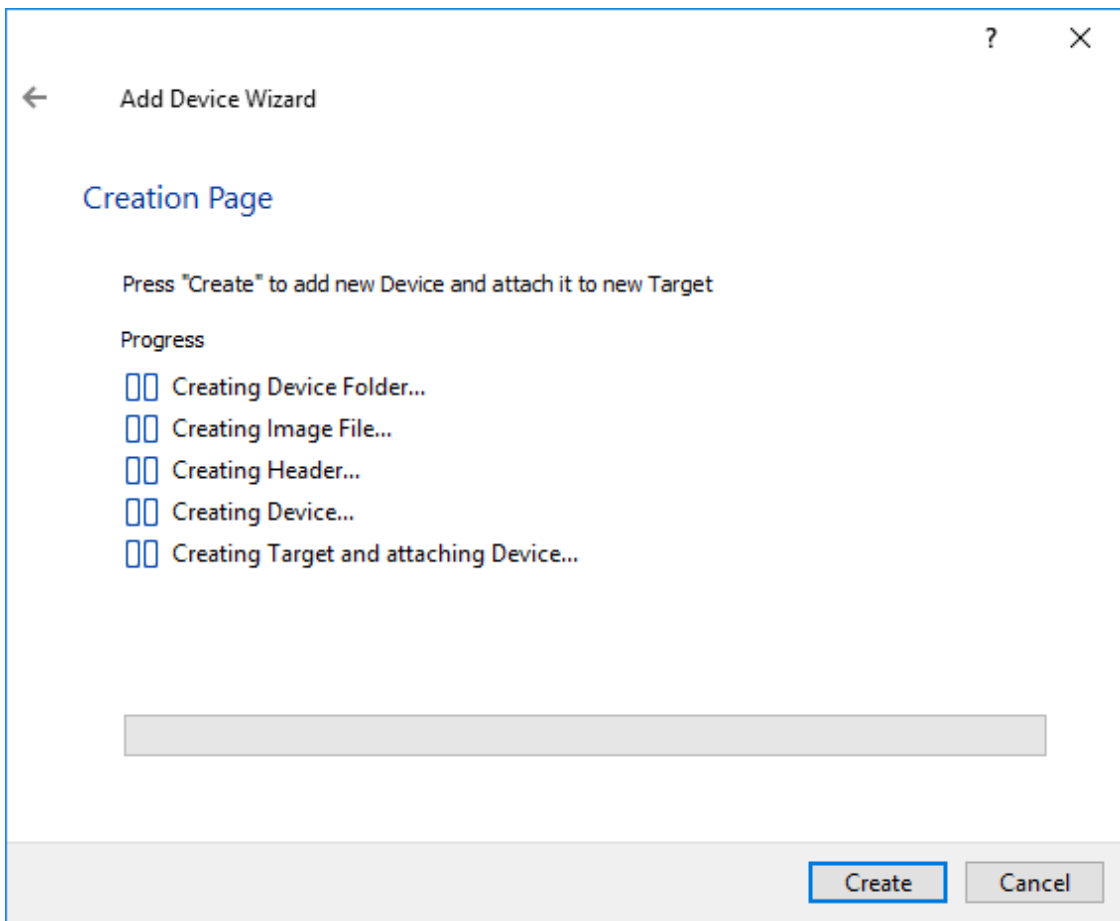
6. Define a caching policy and specify a cache size (in MB). Also, the maximum available cache size can be specified by selecting the appropriate checkbox. Optionally, define the L2 caching policy and cache size.



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.

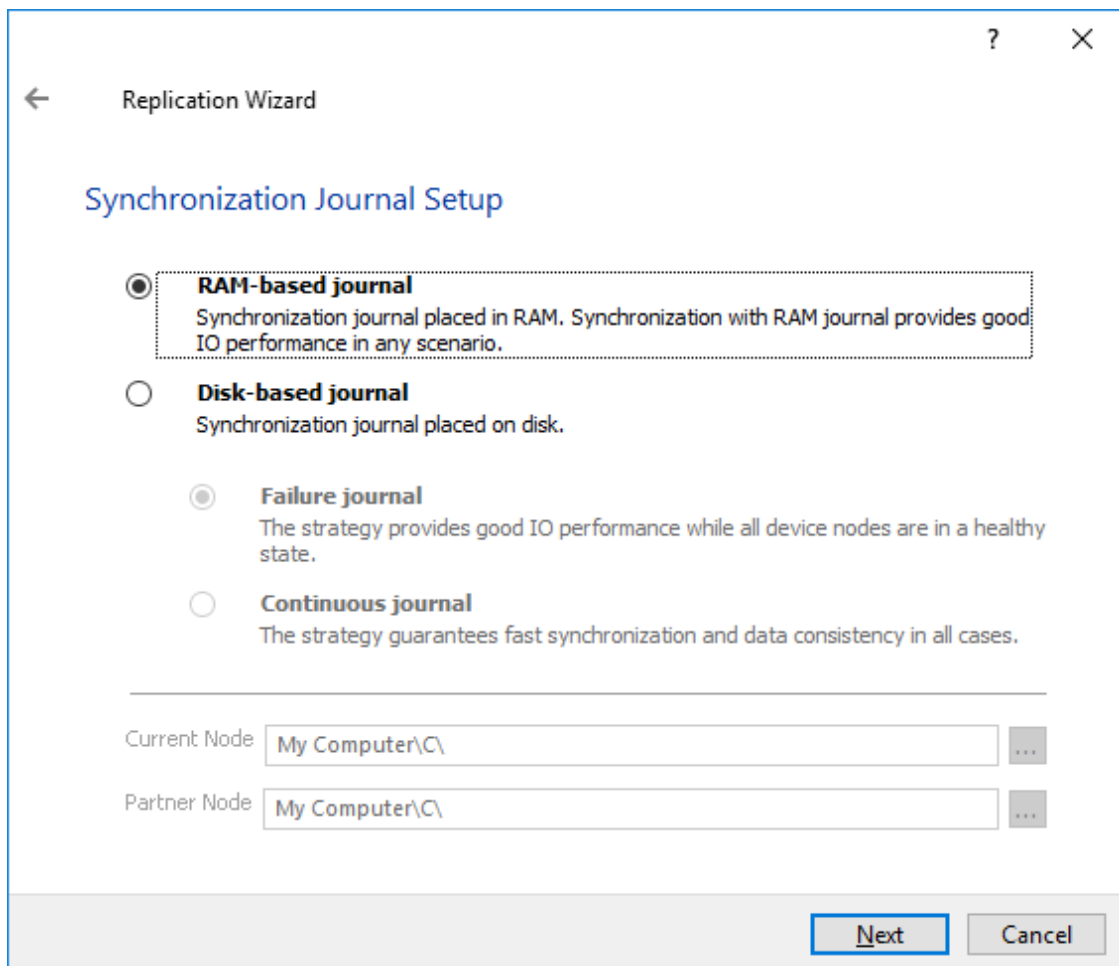
Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are recommended as the device performance is defined by the disk speed, where the journal is located. For example, it

can be placed on the OS boot volume.

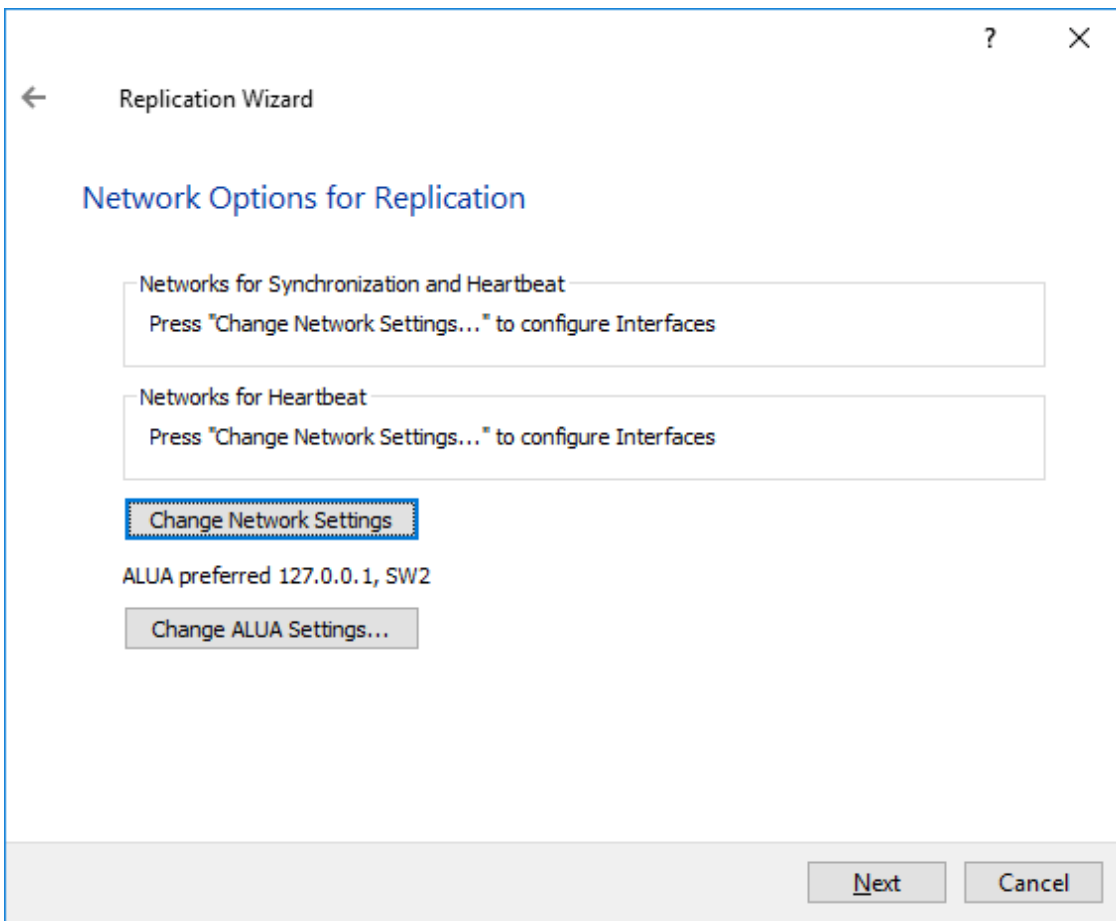
It is required to allocate 2 MB of disk space for the synchronization journal per 1 TB of HA device size with a disk-based journal configured and 2-way replication and 4MB per 1 TB of HA device size for 3-way replication.

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.

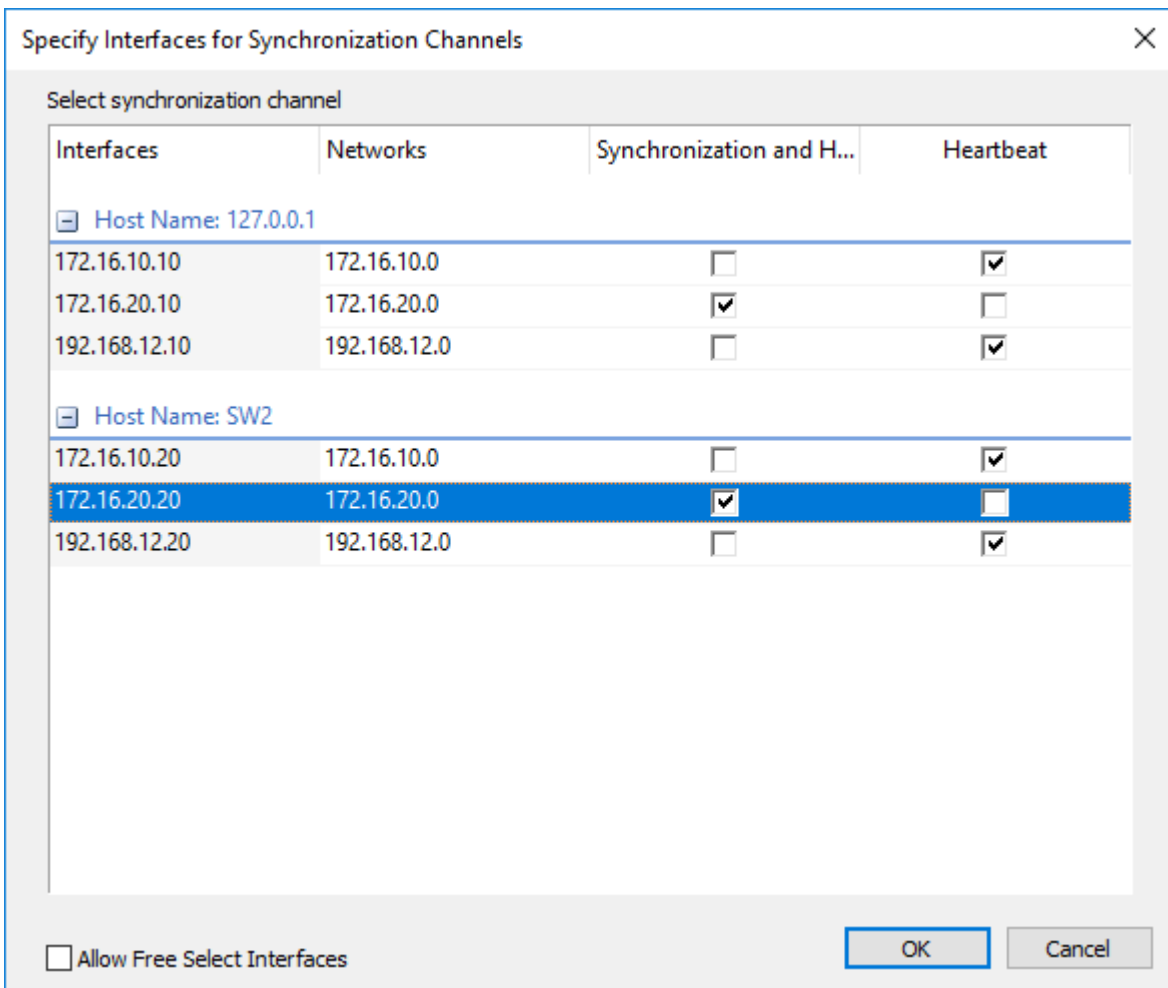
Continuous journal – guarantees fast synchronization and data consistency in all cases. Although, this strategy has the worst I/O performance, because of frequent write operations to the journal, located on the disk, where the journal is located.



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. Select the Node Majority failover strategy and click Next.



2. Choose Create new Partner Device and click Next.

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

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

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

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

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

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


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

Adding Witness Node

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


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

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




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

3. Select Witness Node.



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



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

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

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

8. Click Create Replica and then close the wizard.

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

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


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

Provisioning Starwind Ha Storage To Windows Server Host


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

```
iscsicpl
```

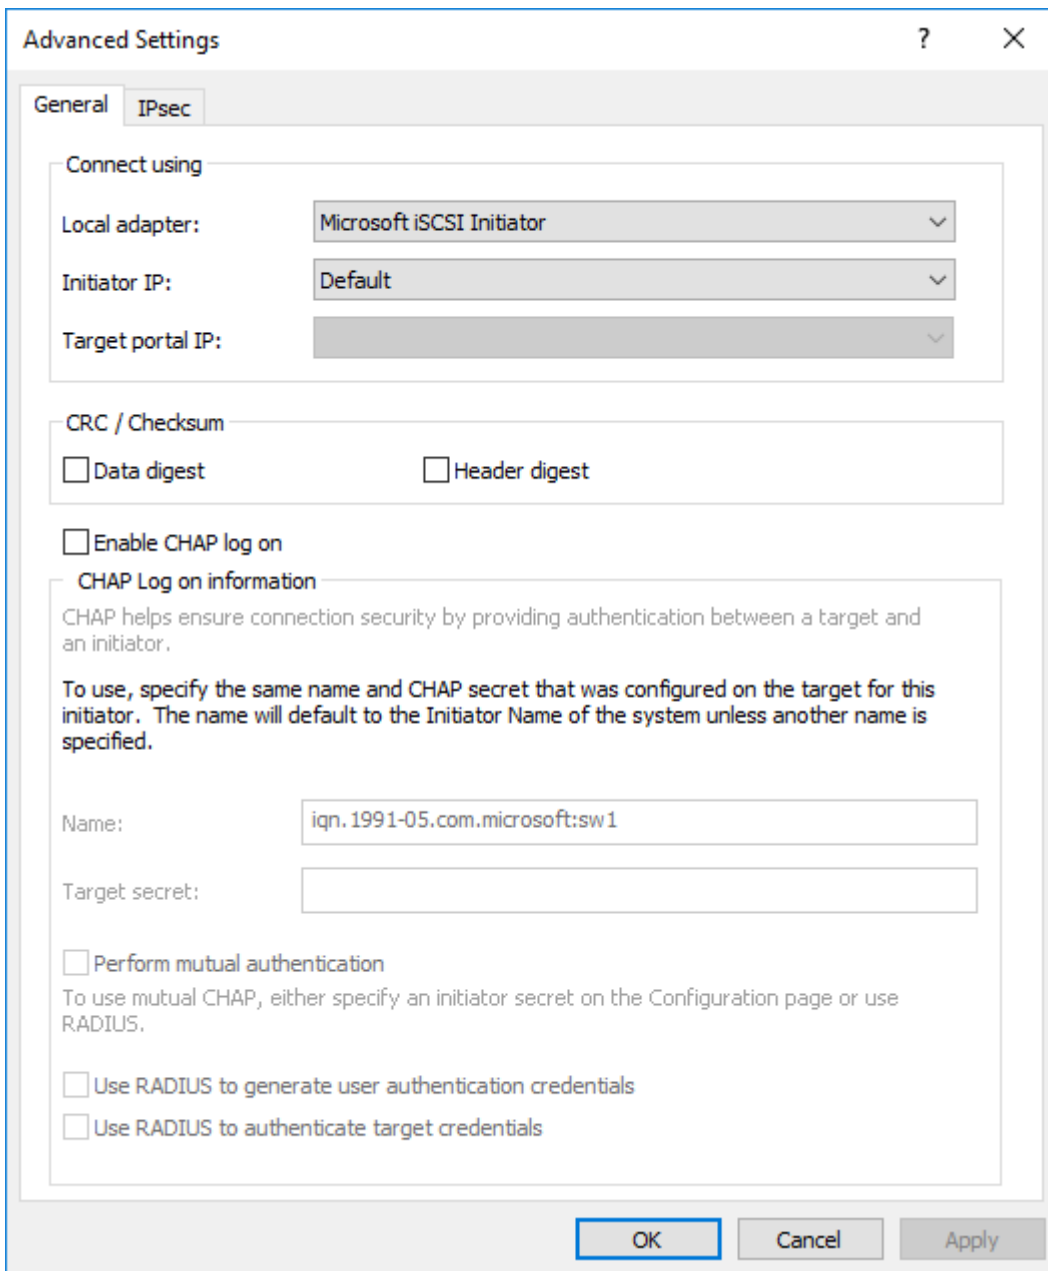
2. Navigate to the Discovery tab.



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




4. Click the Advanced button. Select Microsoft iSCSI Initiator as a Local adapter and select Initiator IP (leave default for 127.0.0.1). Confirm the actions to complete the Target Portal discovery.

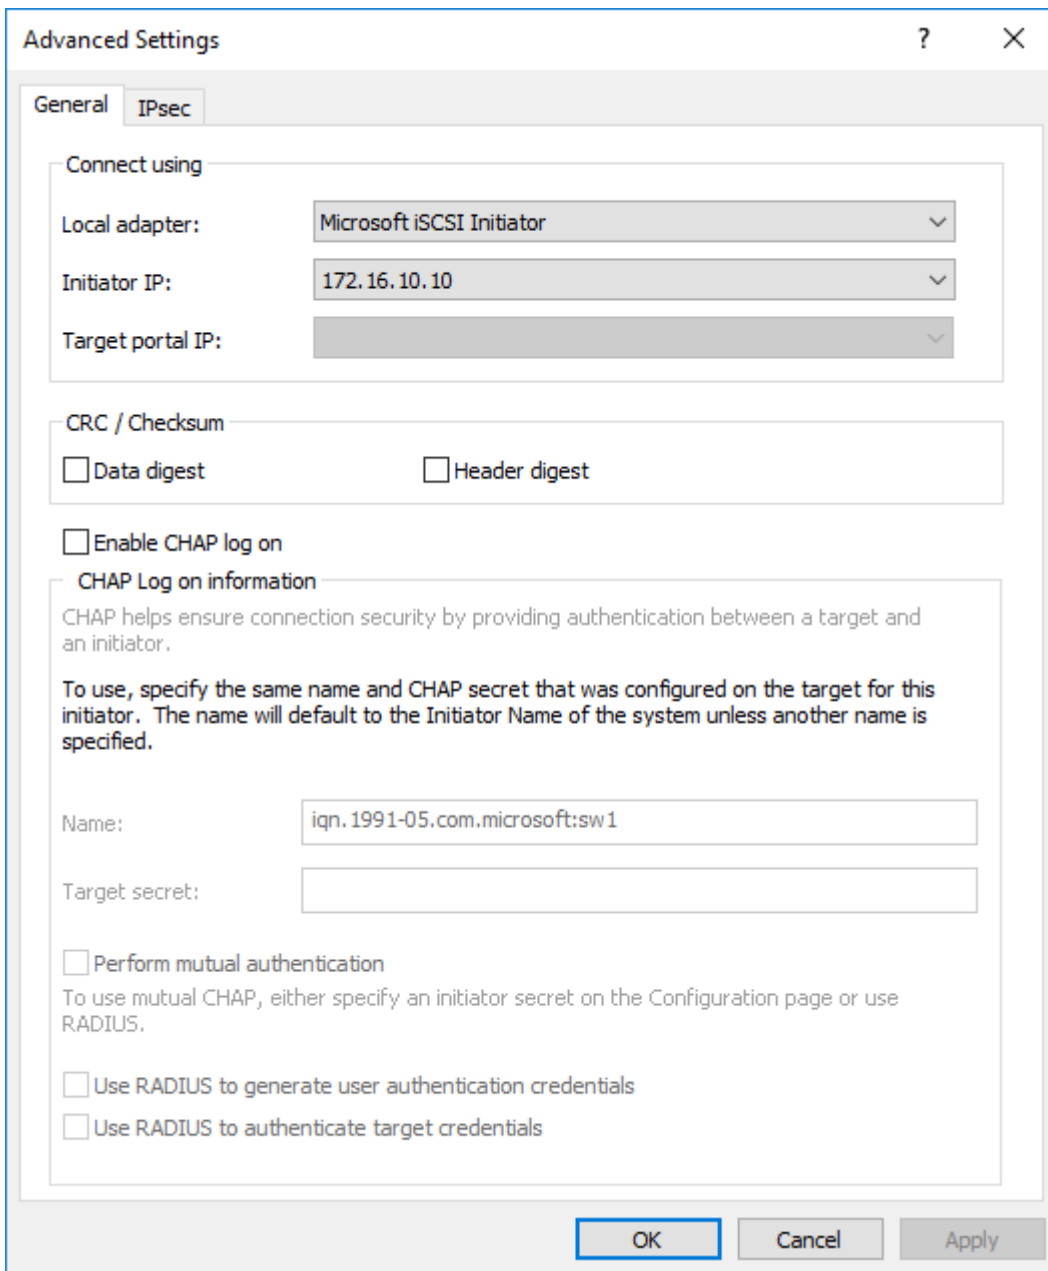


5. Click the Discover Portal... button once again.

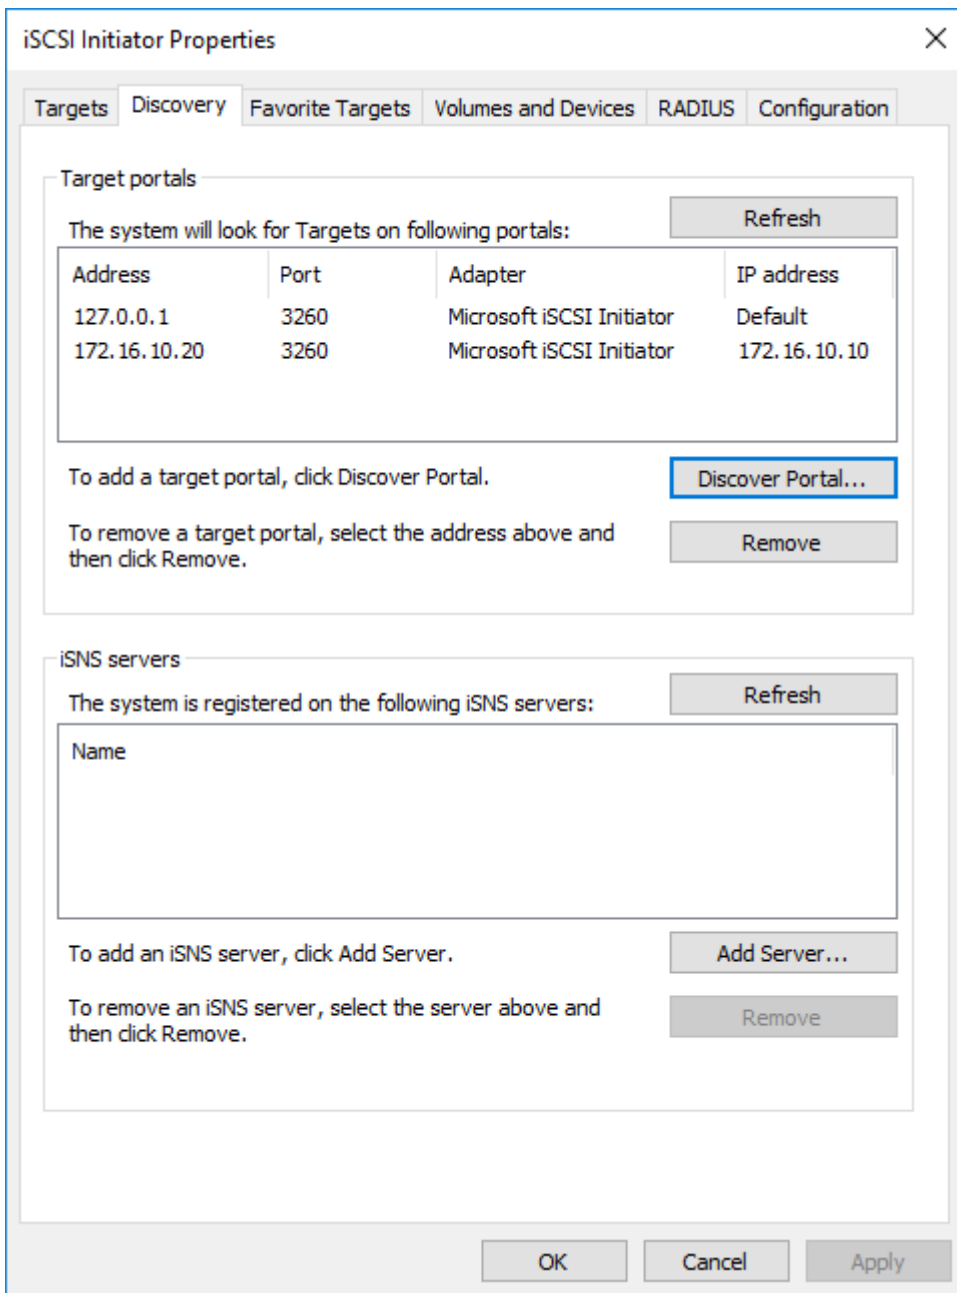
6. In Discover Target Portal dialog, type in the iSCSI interface IP address of the partner node that will be used to connect the StarWind provisioned targets. Click Advanced.



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



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




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

Connecting Targets


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

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


for any restrictions.



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




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



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

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




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

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

9. Repeat the step 6.

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



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


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

Configuring Multipath

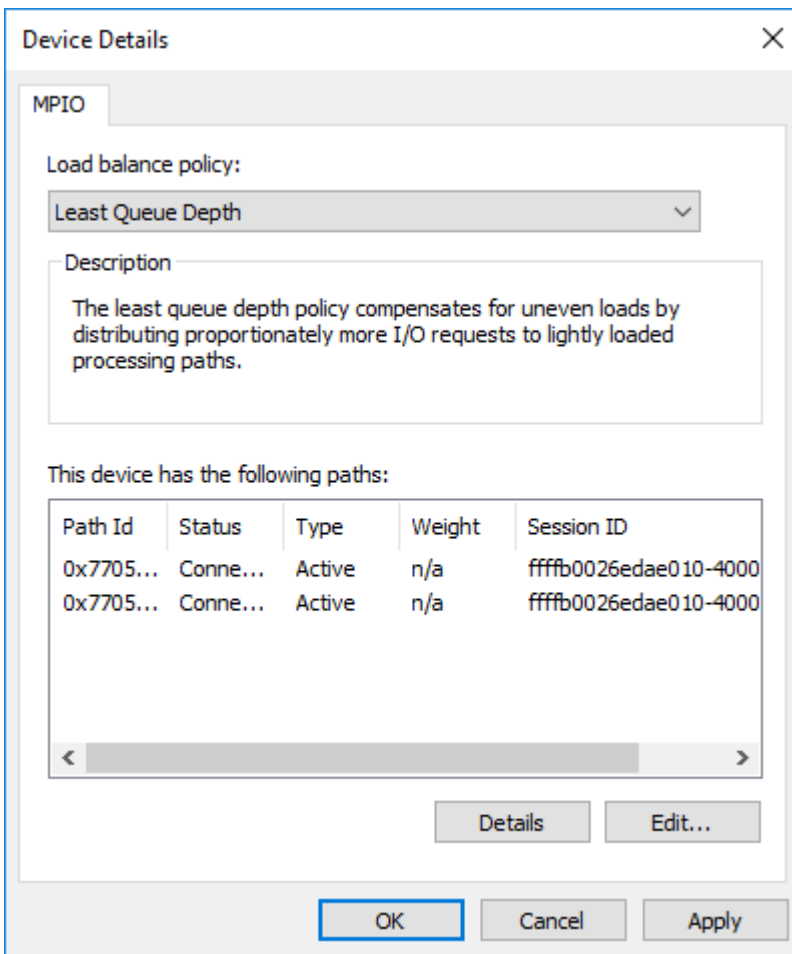
NOTE: It is recommended to configure the different MPIO policies depending on iSCSI channel throughput. For 1 Gbps iSCSI channel throughput, it is recommended to set Failover Only or Least Queue Depth MPIO load balancing policy. For 10 Gbps iSCSI channel throughput, it is recommended to set Round Robin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

load balancing policy.

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



3. Select the appropriate load balancing policy.



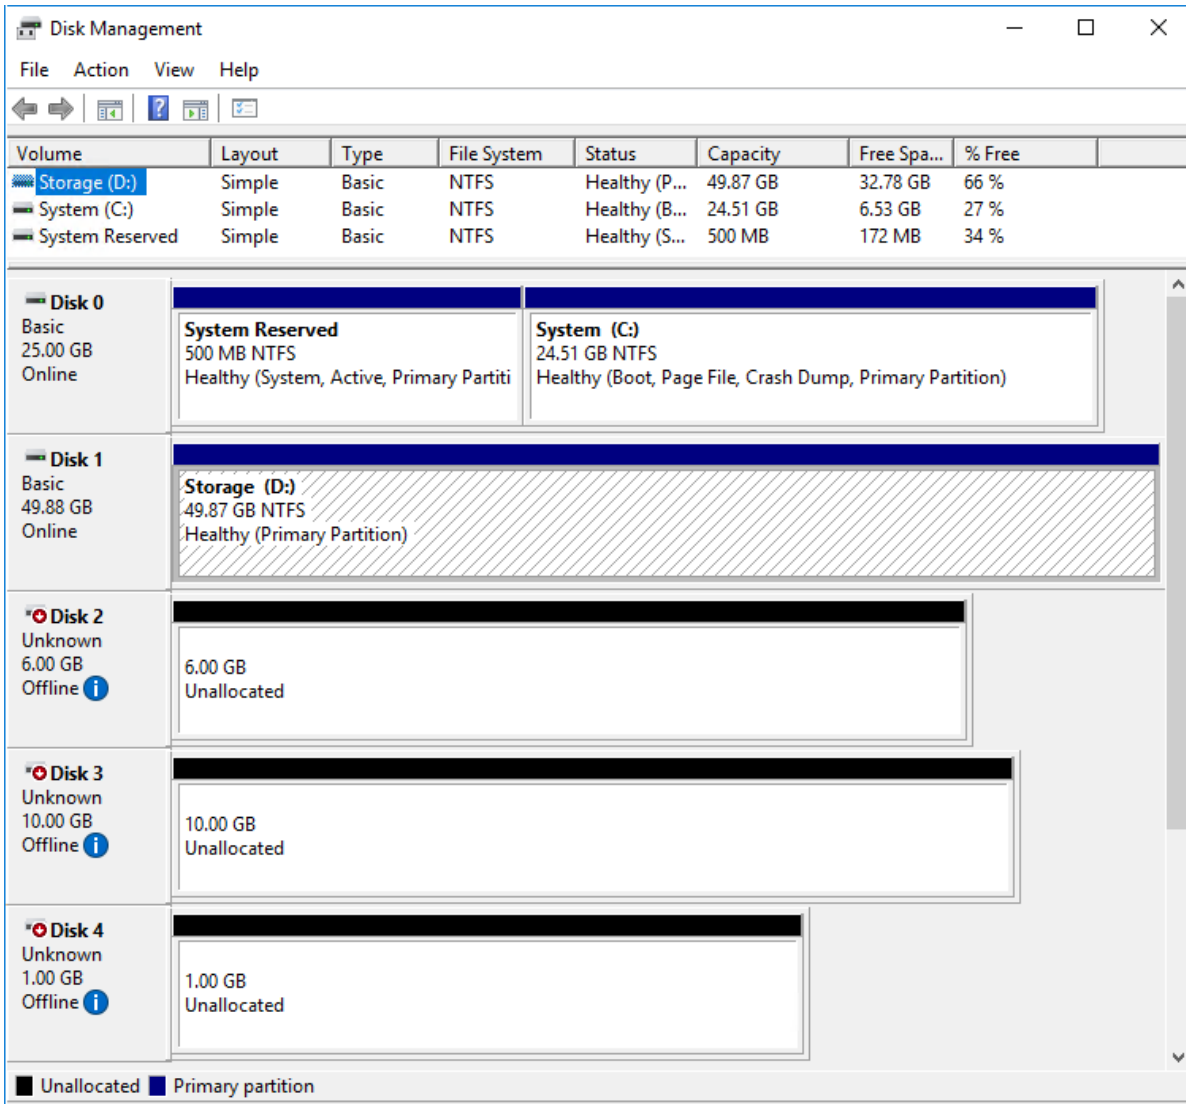
4. For the Witness target, set the load balance policy to Failover Only.

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


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

Connecting Disks to Servers


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




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



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




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



9. Press Finish to complete.

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



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

12. Repeat step 2 to bring all the remaining StarWind disks online.

Creating A Failover Cluster In Windows Server


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

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


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

NOTE: For converged deployment (SAN & NAS running as a dedicated storage cluster) the Microsoft Failover Cluster is deployed on separate computing nodes. Additionally, for the converged deployment scenario, the storage nodes that host StarWind SAN & NAS as CVM or bare metal do not require a domain controller and Failover Cluster to operate.


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




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



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




4. Validate the configuration by running the cluster validation tests: select Yes... and click Next to continue.

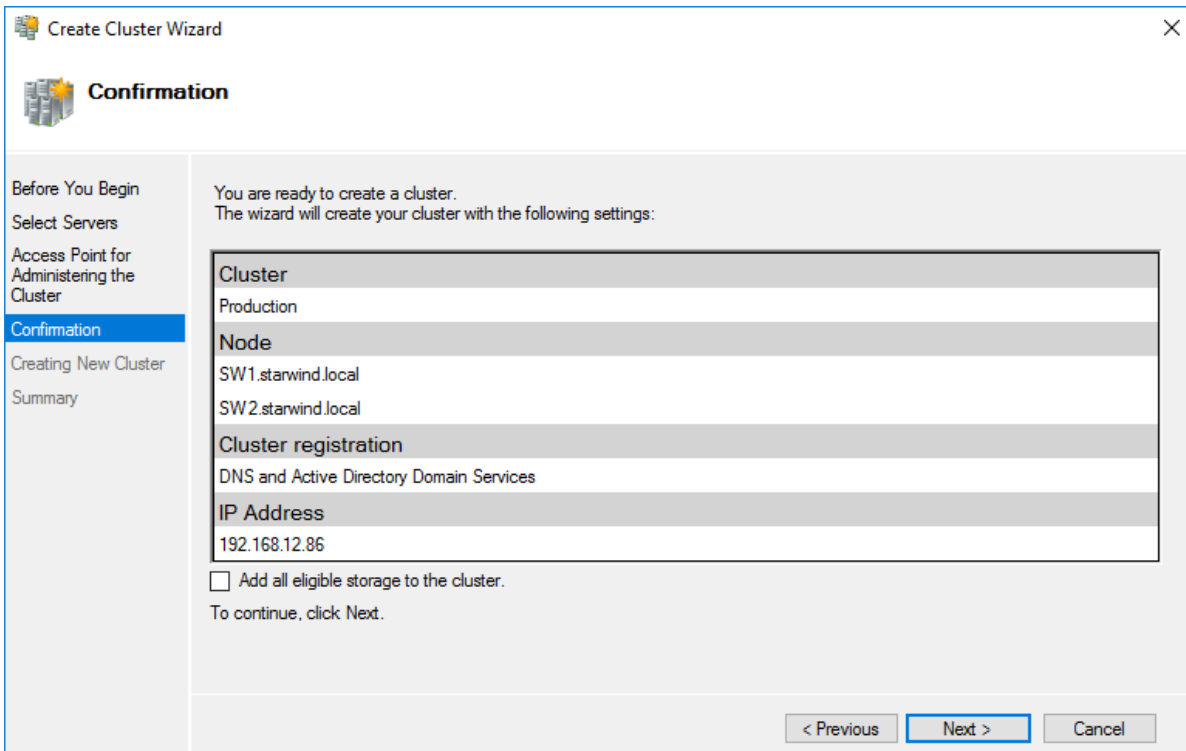


5. Specify the cluster name.

NOTE: If the cluster servers get IP addresses over DHCP, the cluster also gets its IP address over DHCP. If the IP addresses are set statically, set the cluster IP address manually.




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




NOTE: If checkbox Add all eligible storage to the cluster is selected, the wizard will add all disks to the cluster automatically. The device with the smallest storage volume will be assigned as a Witness. It is recommended to uncheck this option before clicking Next and add cluster disks and the Witness drive manually.

7. The process of the cluster creation starts. Upon the completion, the system displays the summary with the detailed information. Click Finish to close the wizard.




Adding Storage to the Cluster


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



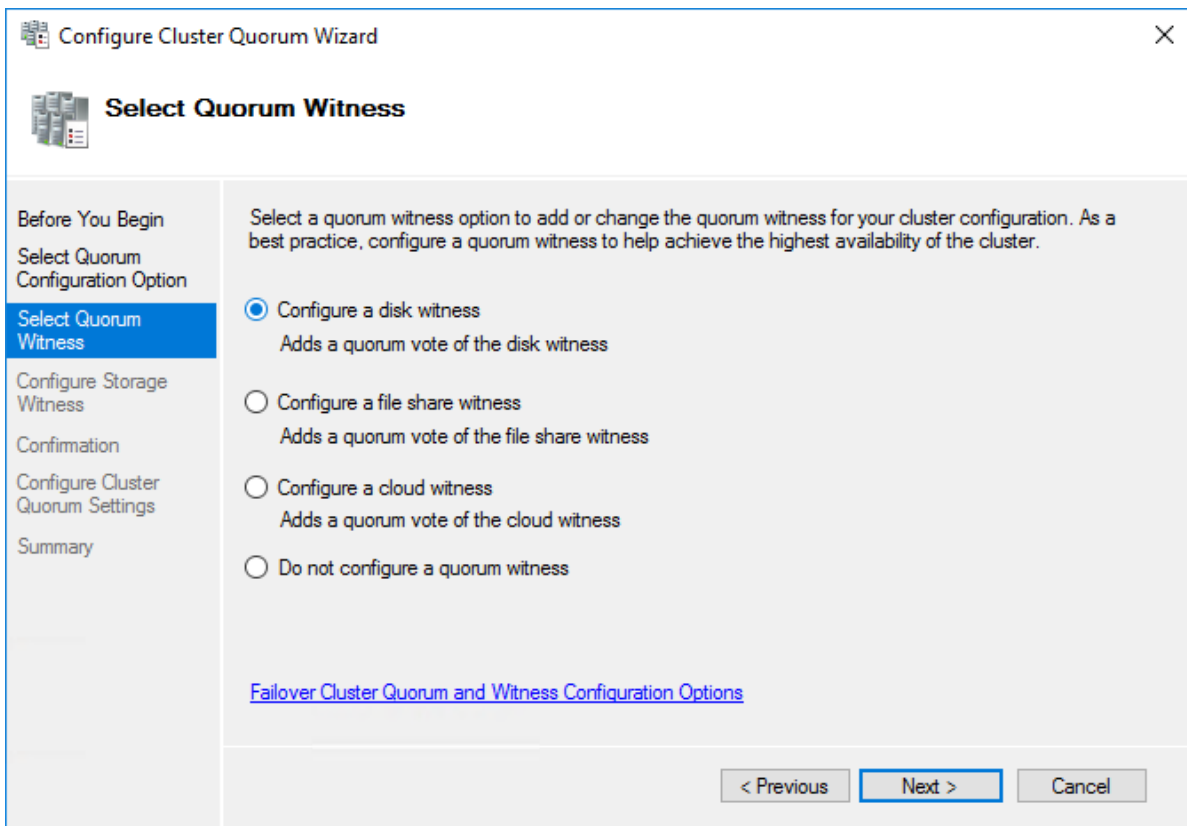
2. To configure the cluster witness disk, right-click on Cluster and proceed to More Actions -> Configure Cluster Quorum Settings.




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



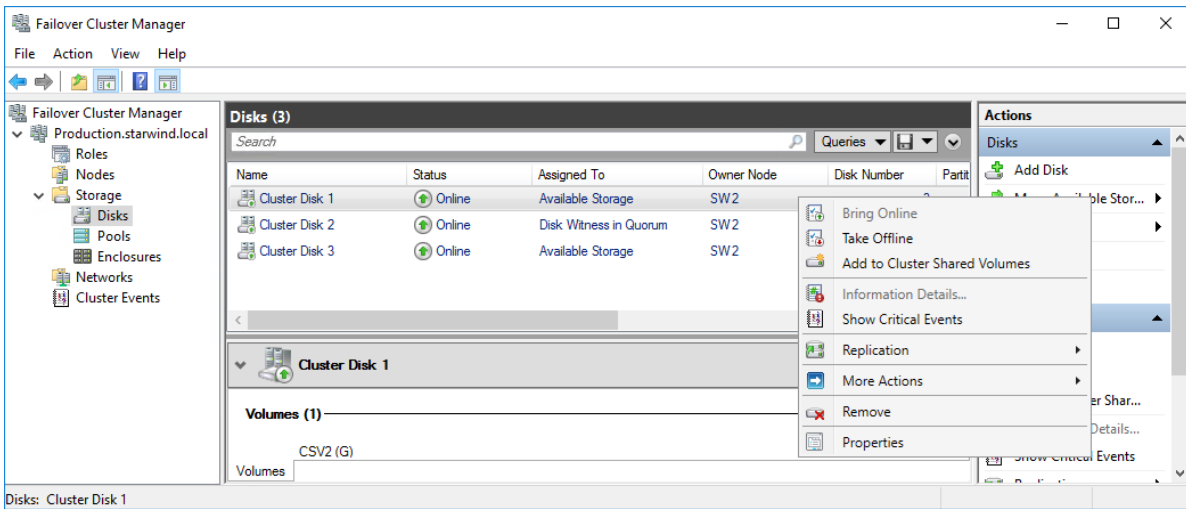
4. Select Configure a disk witness. Click Next.




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



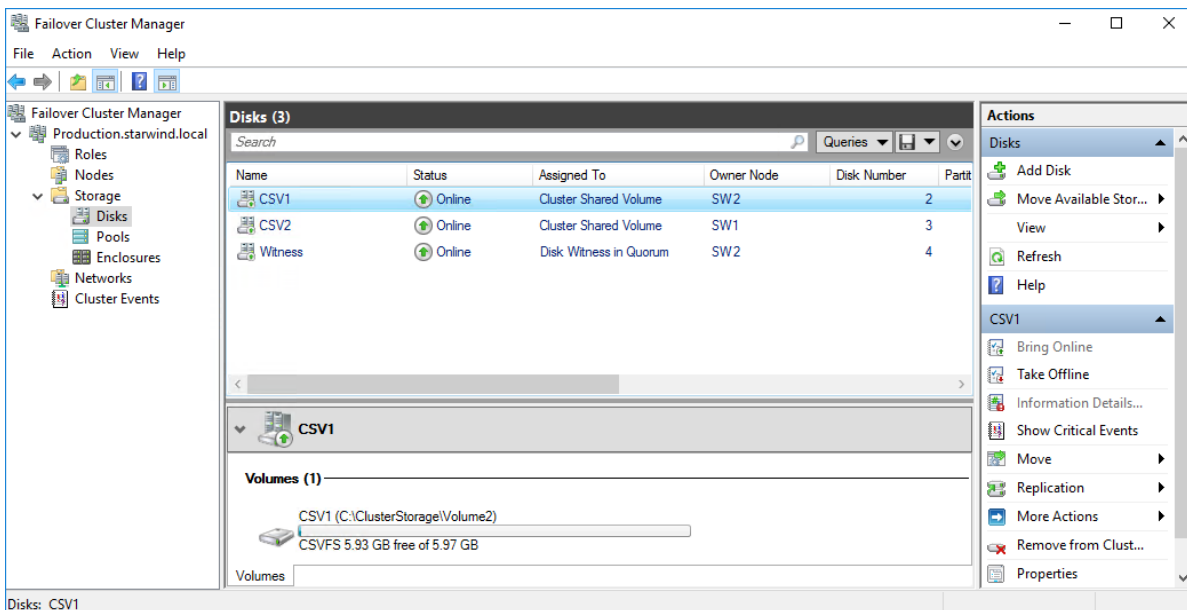
6. In Failover Cluster Manager, Right-click the disk and select Add to Cluster Shared Volumes.



7. If renaming of the cluster shared volume is required, right-click on the disk and select Properties. Type the new name for the disk and click Apply followed by OK.



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




Configuring Cluster Network Preferences


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

NOTE: Please double-check that cluster communication is configured with redundant networks:


<https://docs.microsoft.com/en-us/windows-server/failover-clustering/smb-multichannel>



2. Rename other networks as described above, if required.



3. In the Actions tab, click Live Migration Settings. Uncheck the synchronization network, while the iSCSI network can be used if it is 10+ Gbps. Apply the changes and click OK.










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

Conclusion

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

Contacts

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

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

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

Sales: sales@starwind.com

General Information: info@starwind.com



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