StarWind Virtual SAN® 2-node Compute and Storage Separated Scenario with Windows Server 2012 R2

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.

IMPORTANT NOTE: Windows Server 2012 R2 is no longer supported by Microsoft (Mainstream End Date is Oct 9, 2018). This guide should be used for informational, test, and educational purposes only.

Starting from the 14869 build, the Asynchronous replication and LSFS features were excluded and deprecated. Please remove all Asynchronous replicas and LSFS devices before the update. See more information in the release notes:
https://www.starwindsoftware.com/release-notes-build

Introduction To Starwind Virtual San For Hyper-V

StarWind Virtual SAN is a native Windows hypervisor-centric hardware-less VM storage solution. It creates a fully fault-tolerant and high performing storage pool built for the virtualization workloads by mirroring the existing server's storage and RAM between the participating storage cluster nodes. The mirrored storage resources are then connected to all cluster nodes and treated just as a local storage by all hypervisors and clustered applications. High Availability (HA) is achieved by providing multipath access to all storage nodes. StarWind Virtual SAN delivers supreme performance compared to any dedicated SAN solution since it runs locally on the hypervisor and all I/O is processed by local RAM, SSD cache, and disks. This way it never gets bottlenecked by storage fabric.

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 2012 R2 Hosts

The network interconnection diagram is demonstrated below:
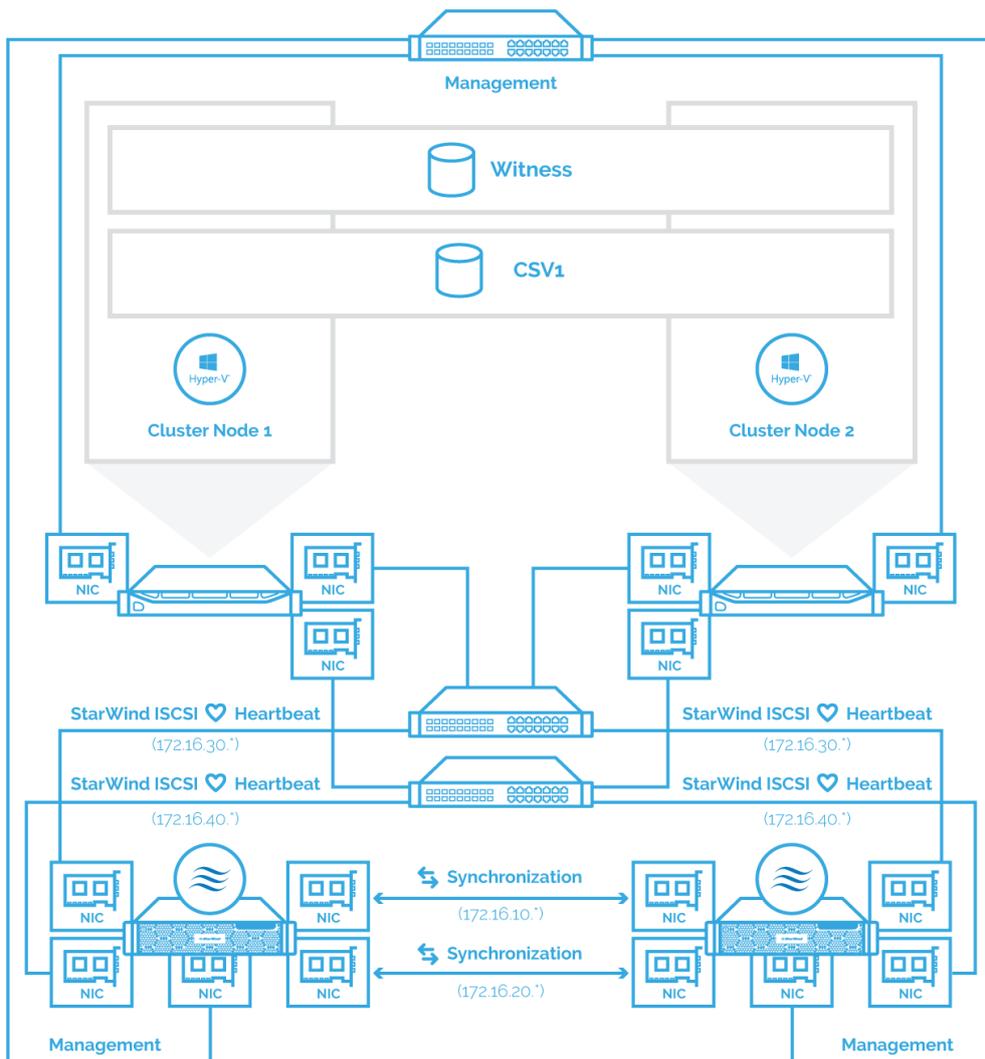
1. Make sure that a domain controller is configured and the servers are added to the domain.
NOTE: Please follow the recommendation in KB article on how to place a DC in case of StarWind Virtual SAN usage.
2. Install Failover Clustering and Multipath I/O features, as well as the Hyper-V role on both servers. This can be done through the Server Manager (Add Roles and Features) menu item.
3. Configure network interfaces on each node to make sure that Synchronization and iSCSI/StarWind heartbeat interfaces are in different subnets and connected according to the network diagram above. In this document, In this document, 172.16.10.x and 172.16.20.x subnets are used for the Synchronization traffic, while 172.16.30.x and 172.16.40.x subnets are used for iSCSI/StarWind heartbeat traffic.
4. In order to allow iSCSI Initiators to discover all StarWind Virtual SAN interfaces, the StarWind configuration file (StarWind.cfg) should be changed after stopping the StarWind service on the node where it will be edited. Locate the StarWind Virtual SAN configuration file (the default path is "C:\Program Files\StarWind

Software\StarWind\StarWind.cfg") and open it via WordPad as Administrator. Find the *<iScsiDiscoveryListInterfaces value="0"/>* string and change the value from 0 to 1 (should look as follows: *<iScsiDiscoveryListInterfaces value="1"/>*). Save the changes and exit Wordpad. Once StarWind.cfg is changed and saved, the StarWind service can be restarted.

Enabling Multipath Support

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

```
mpiocpl
```

6. In the Discover Multi-Paths tab, choose the Add support for iSCSI devices checkbox and click Add.
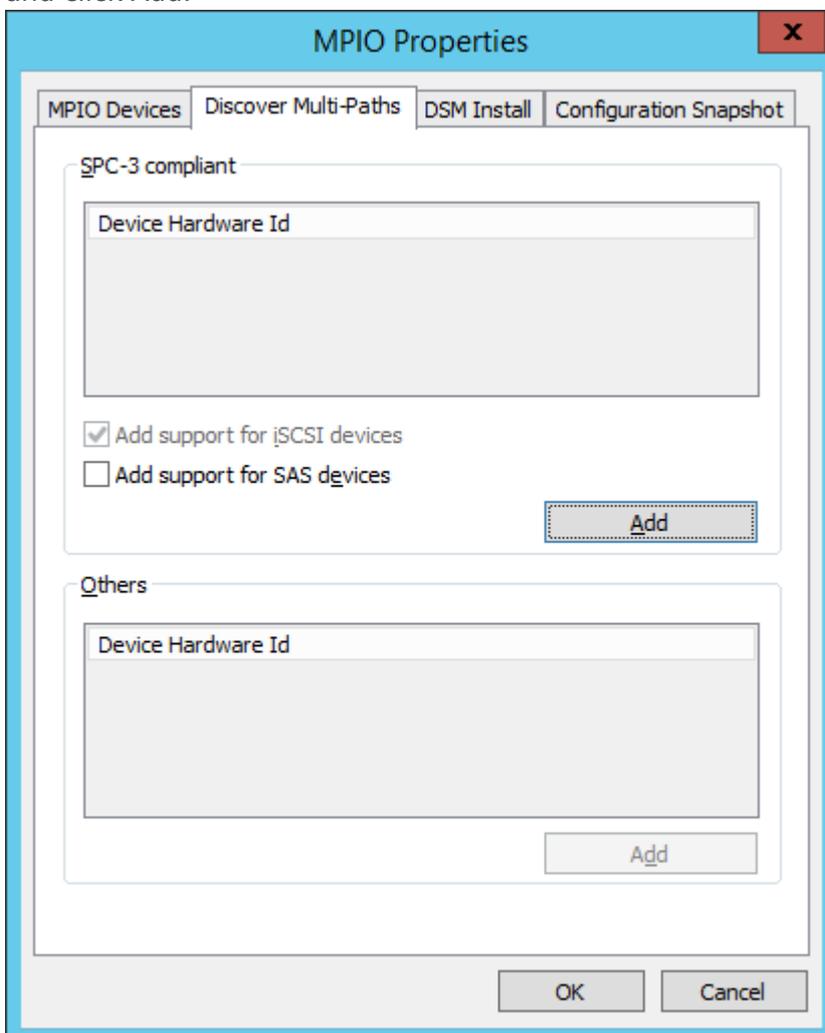
7. When prompted to restart the server, click Yes to proceed.
8. Repeat the same procedure on the other server.

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
NOTE: The setup file is the same for x86 and x64 systems, as well as for all Virtual SAN deployment scenarios.

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.
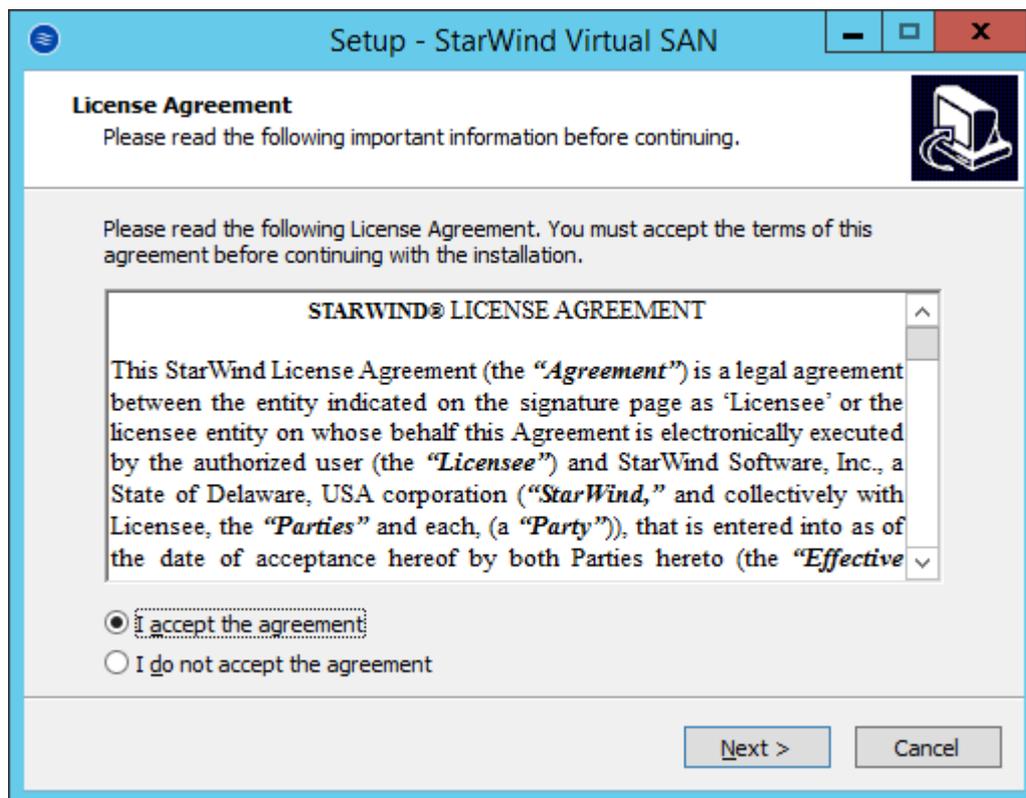
3. Carefully read the information about the new features and improvements. Red text indicates warnings for users who are updating the existing software installations.
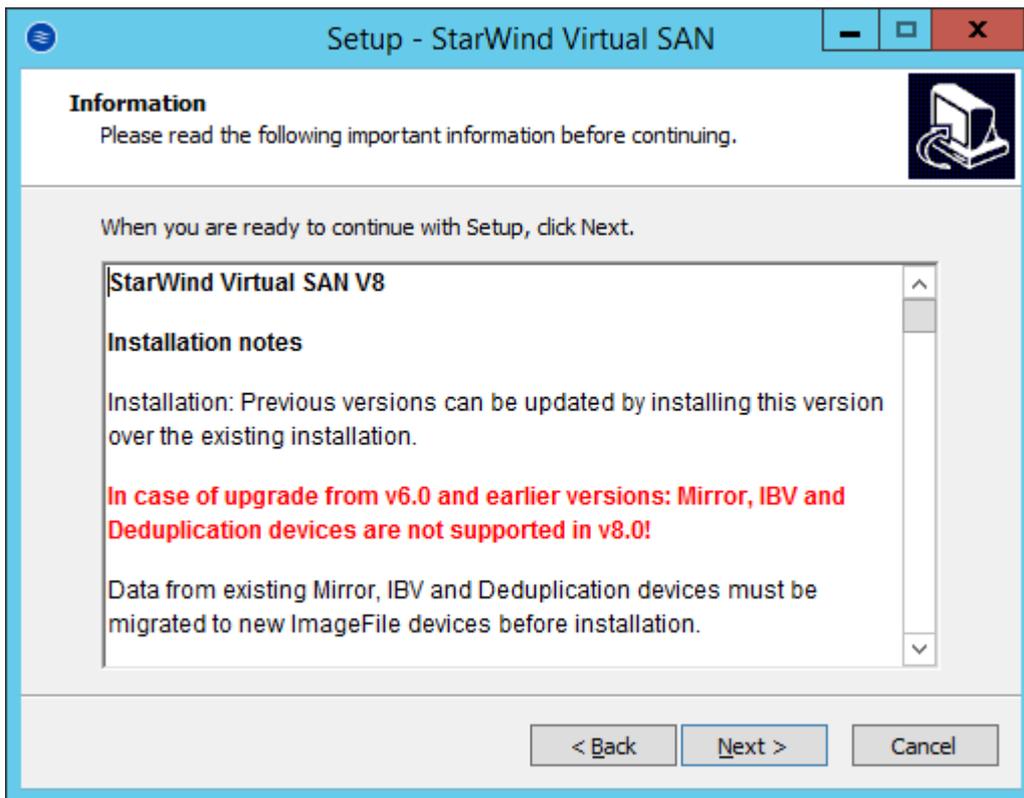
4. Select Browse to modify the installation path if necessary. Click Next to continue.
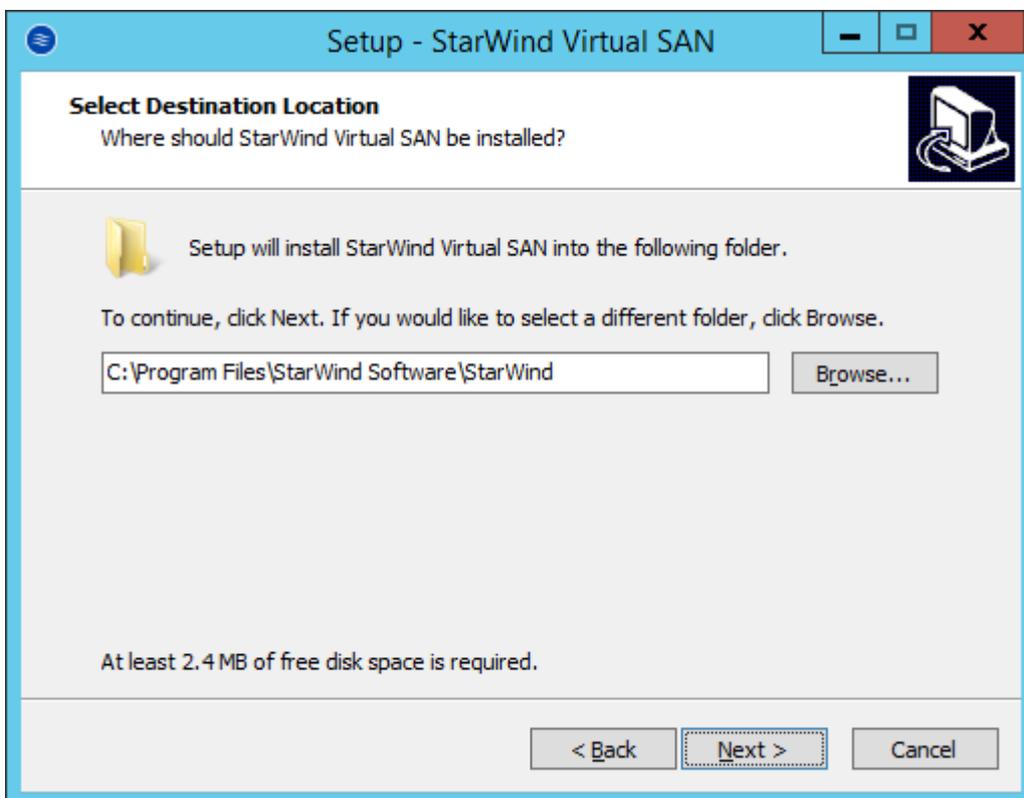
5. Select the following components for the minimum setup:

- StarWind Virtual SAN Service. StarWind 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 or VSA 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.
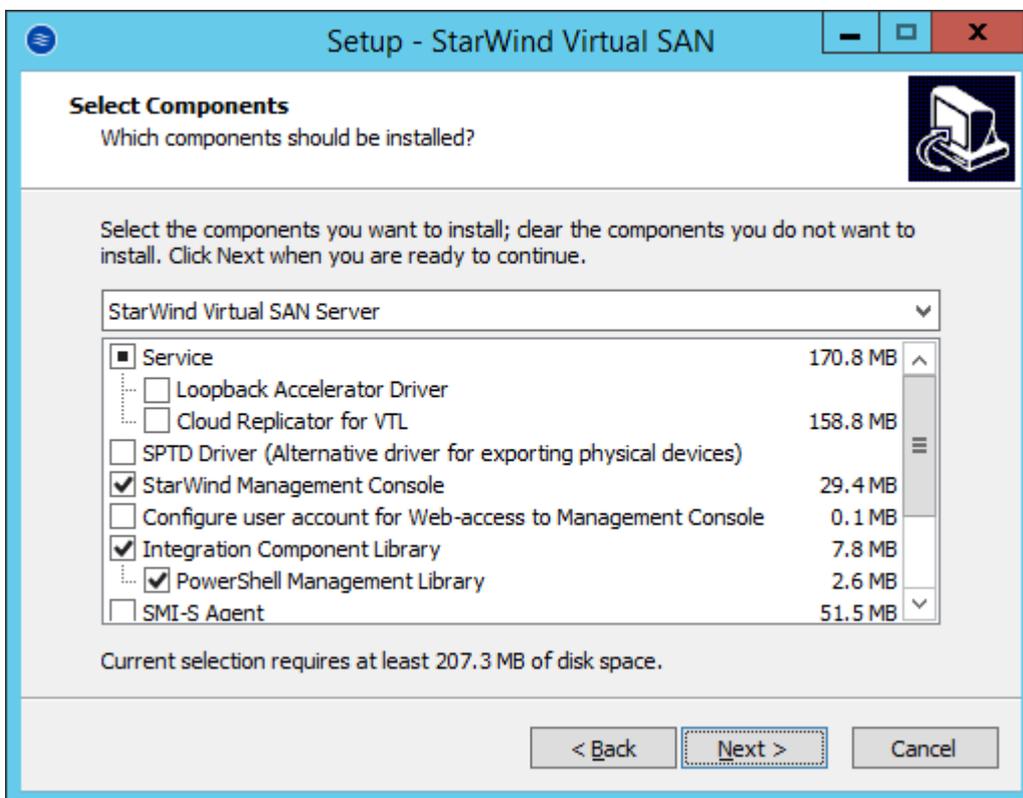
6. Specify Start Menu Folder.
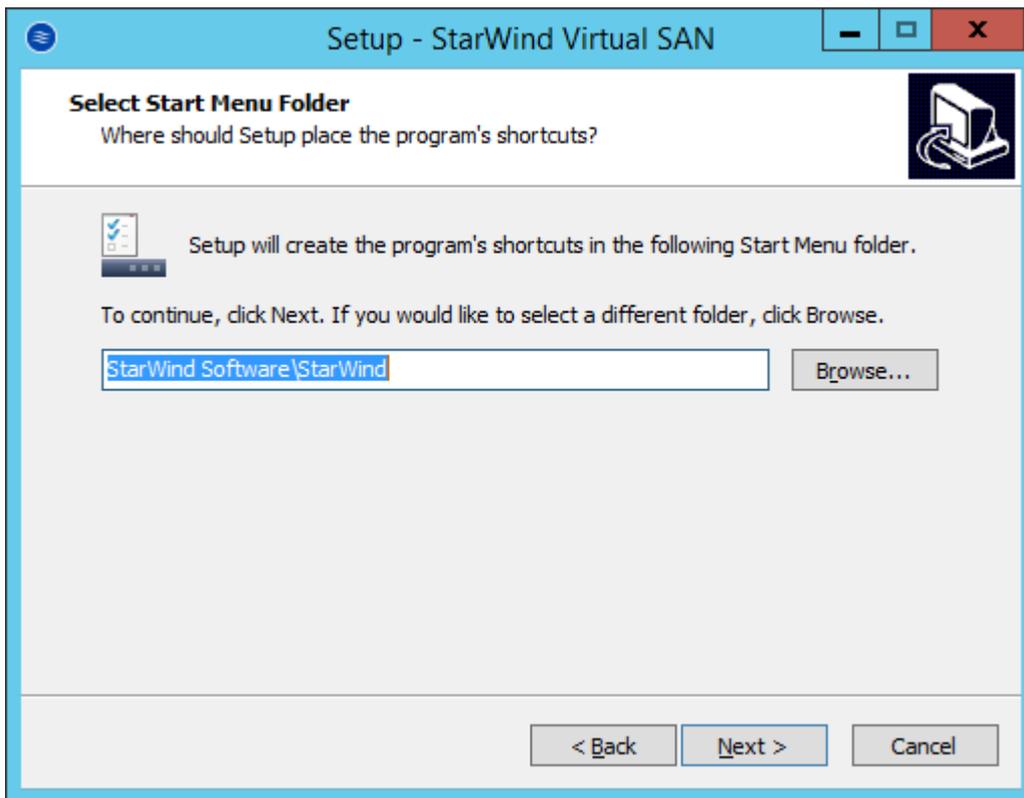
7. Enable the checkbox if a desktop icon needs to be created. Click Next to continue.
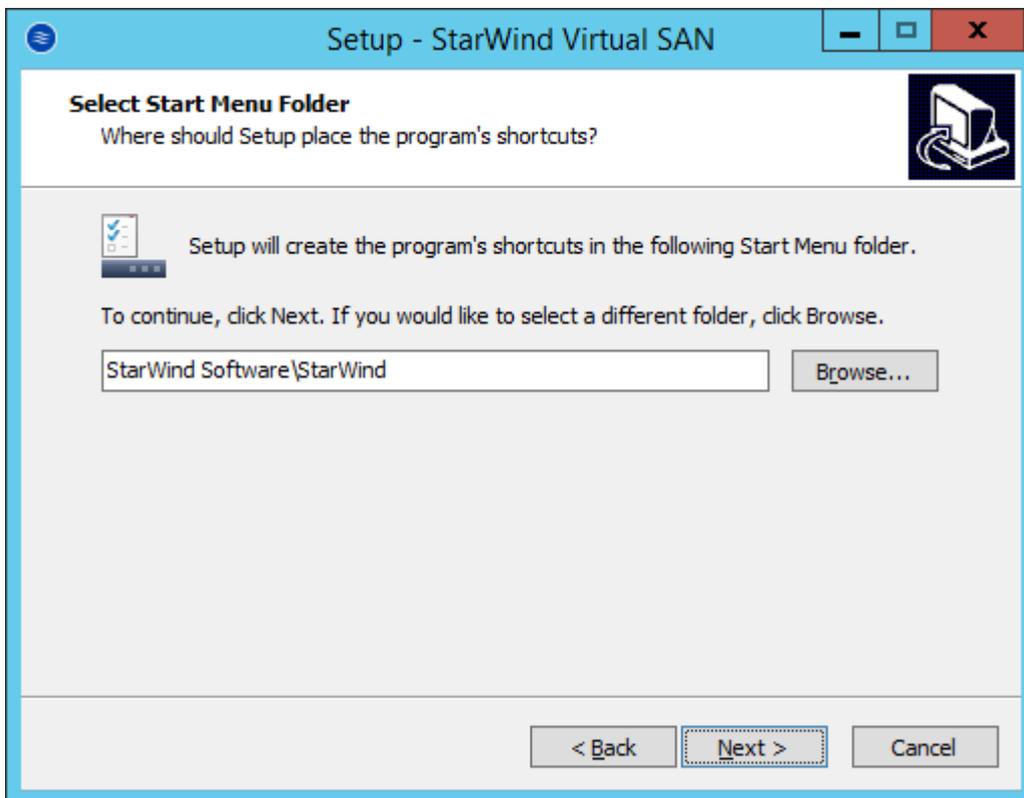
8. When the license key prompt appears, choose the appropriate option:

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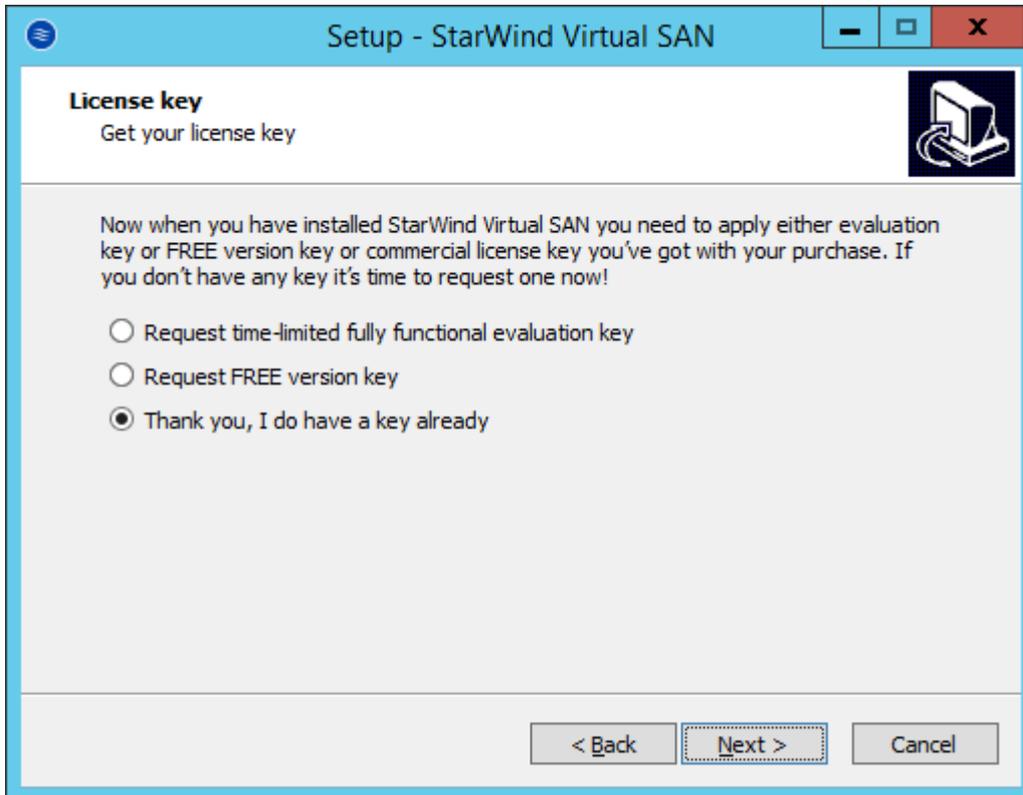
9. Click Browse to locate the license file. Click Next to continue.
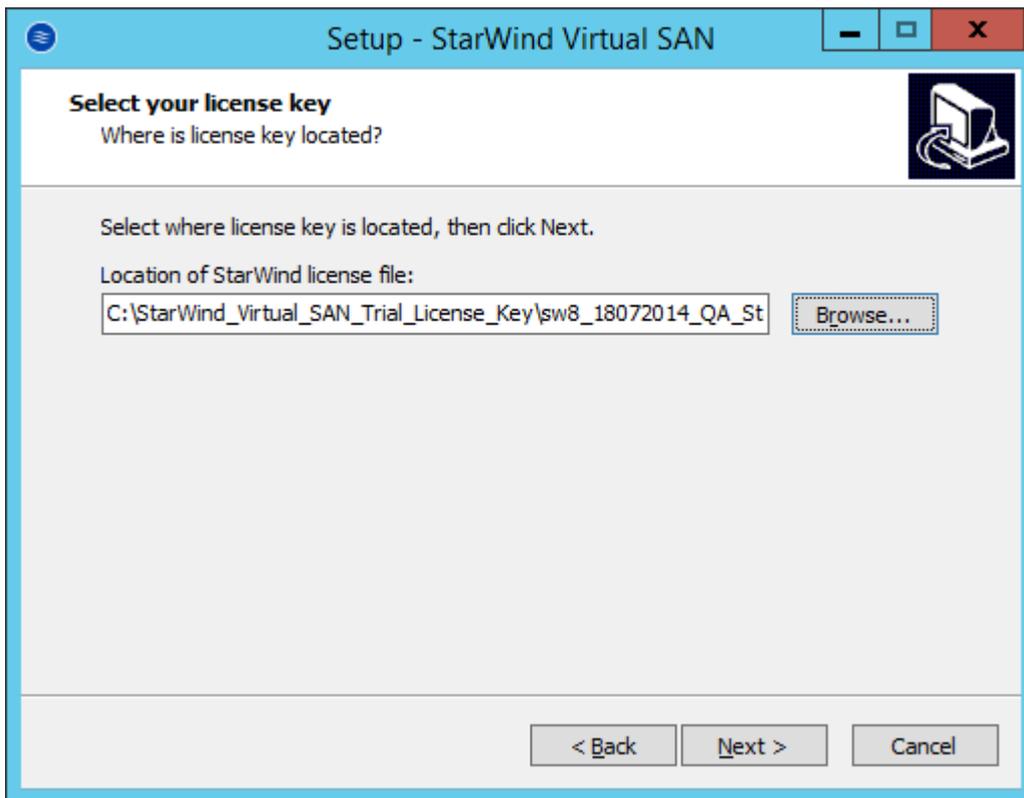
10. Review the licensing information. Click Next to continue.
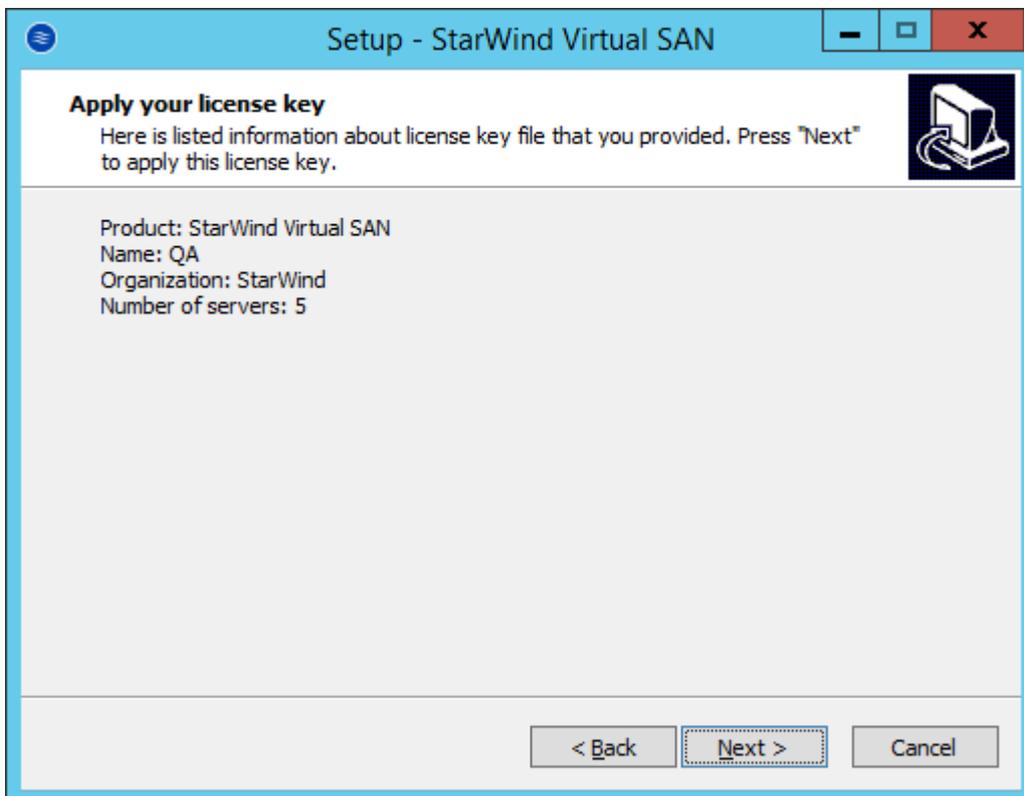
11. Verify the installation settings. Click Back to make any changes or Install to proceed with the installation.
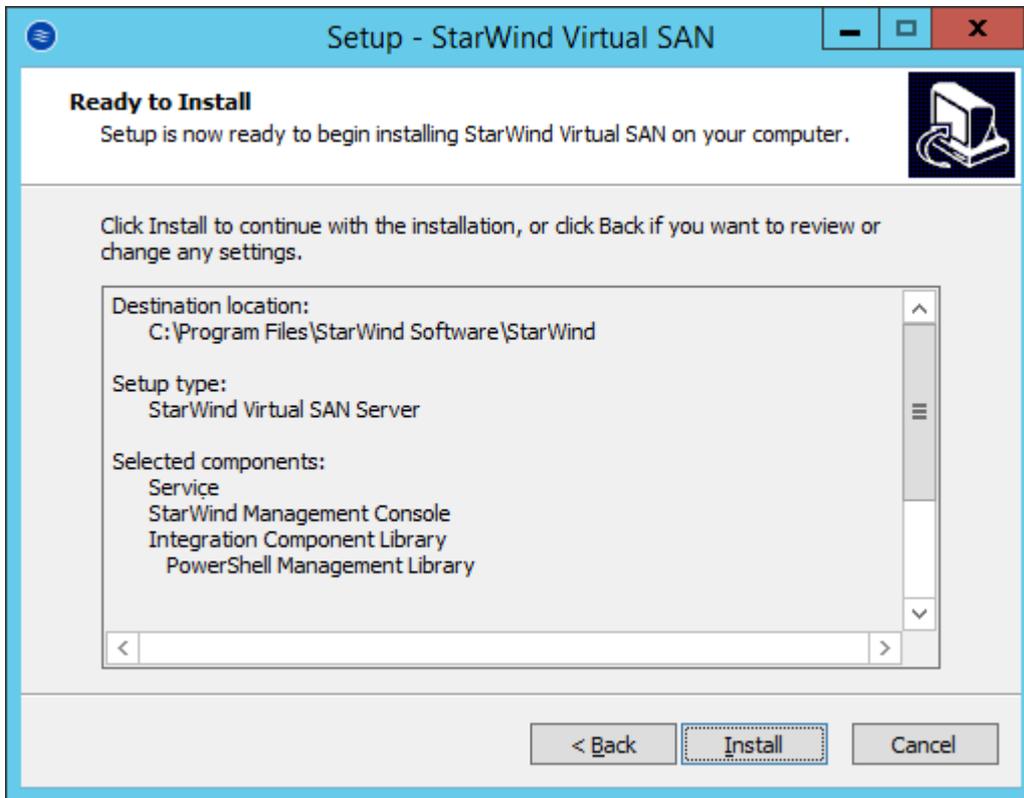
12. Enable the appropriate checkbox to launch StarWind Management Console right after the setup wizard is closed and click Finish.
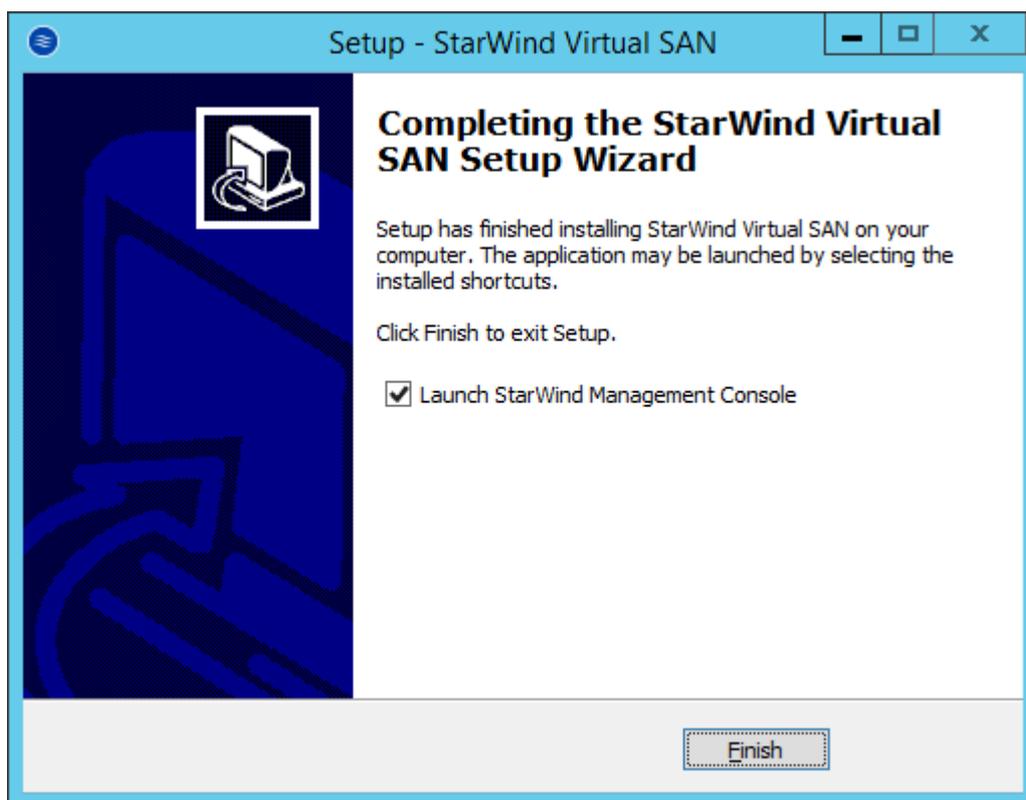
13. Repeat the installation steps on the partner node.

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.
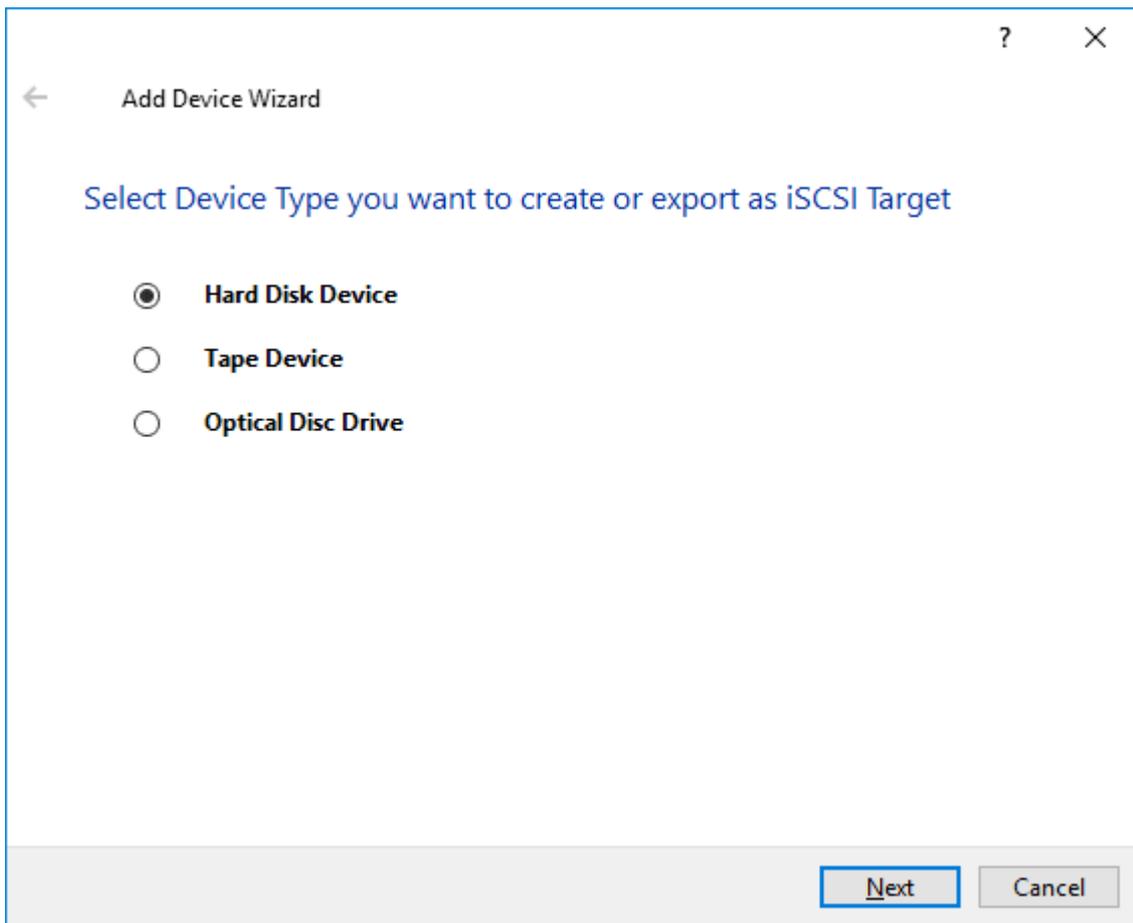
3. Select Virtual Disk.
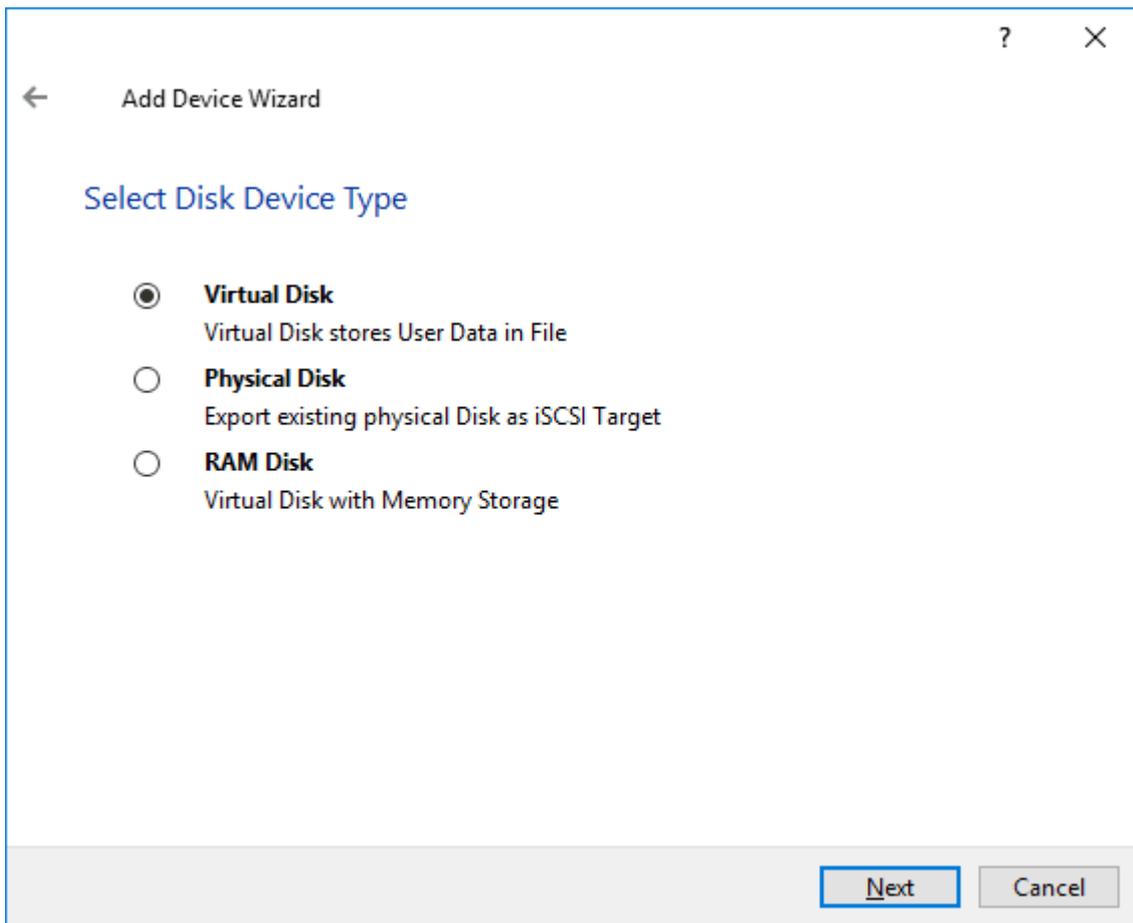
4. Specify a virtual disk Name, Location, and Size.
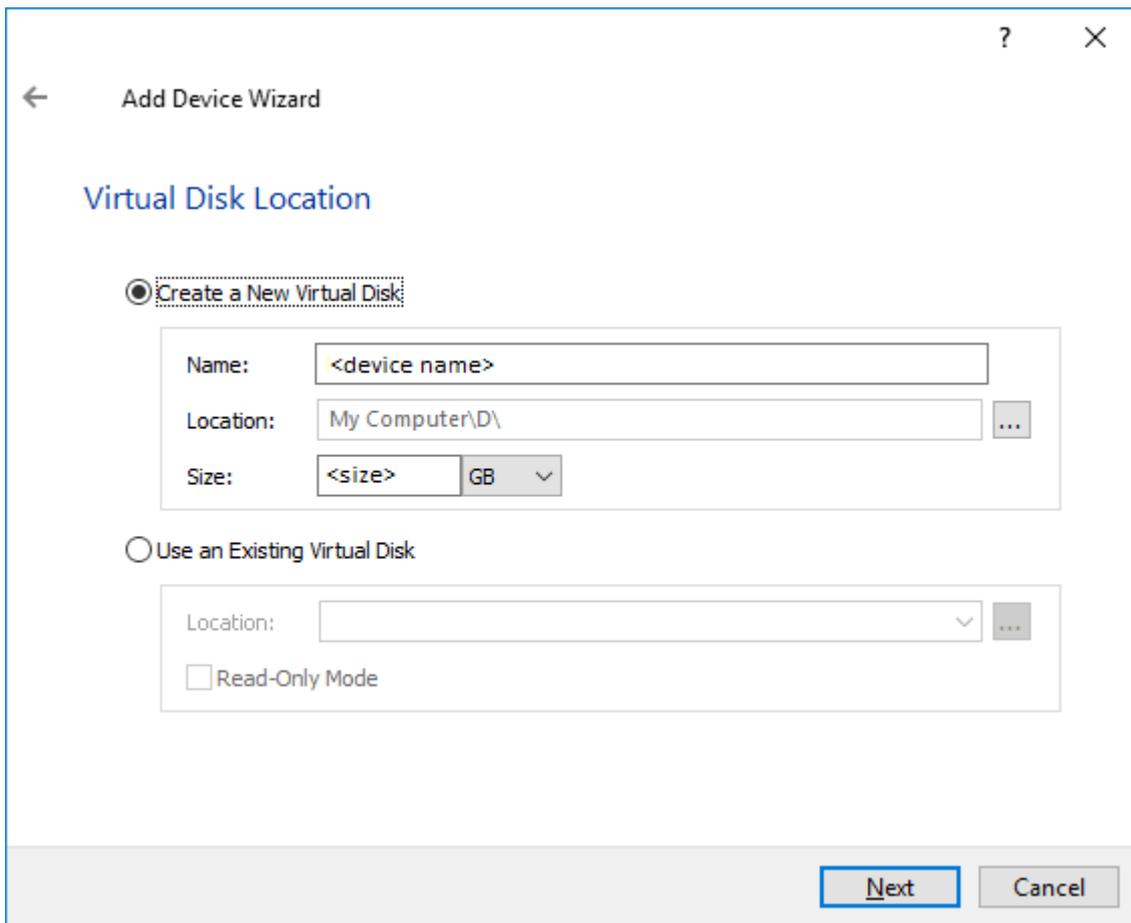
5. Select the Thick provisioned disk type and block size.
NOTE: Use 4096 sector size for targets, connected on Windows-based systems and 512 bytes sector size for targets, connected on Linux-based systems (ESXi/Xen/KVM).

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.
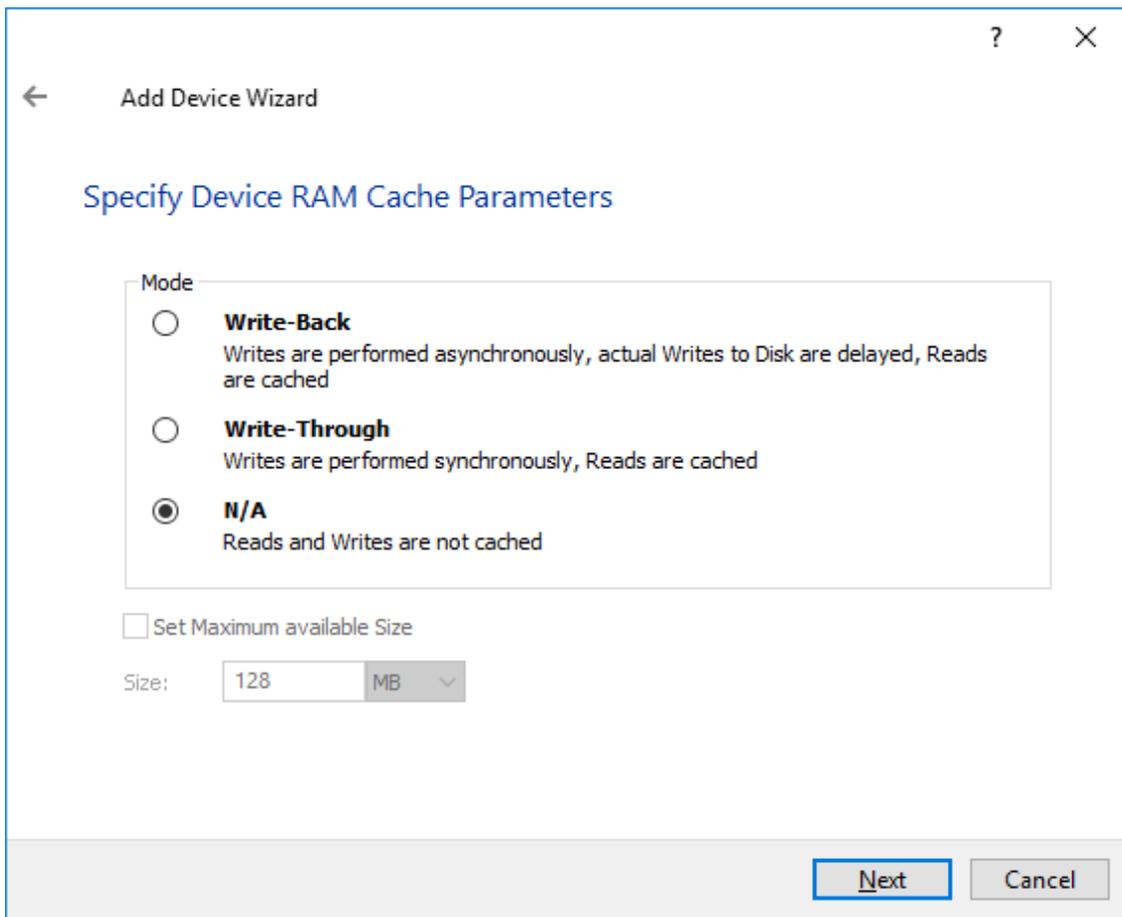
7. Specify Target Parameters. Select the Target Name checkbox to enter a custom target name. Otherwise, the name is generated automatically in accordance with the specified target alias.
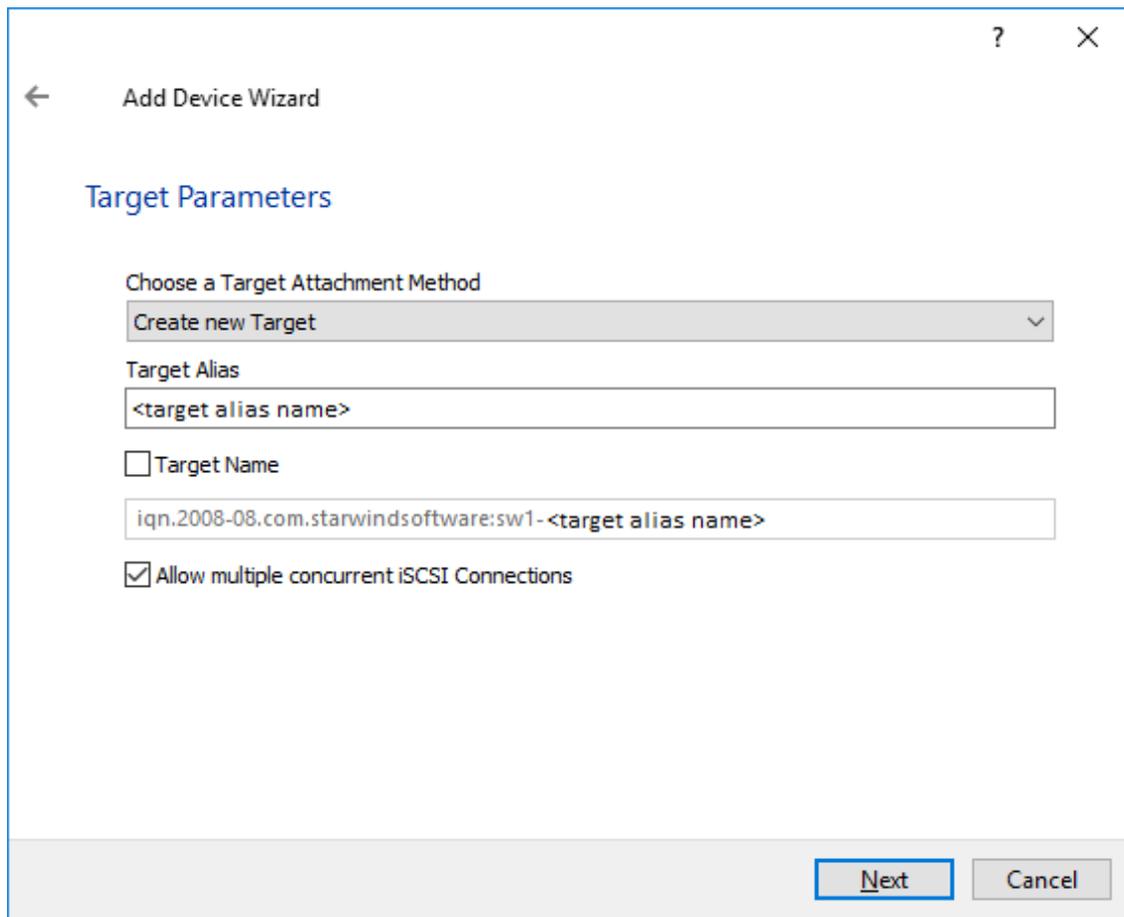
8. Click Create to add a new device and attach it to the target.
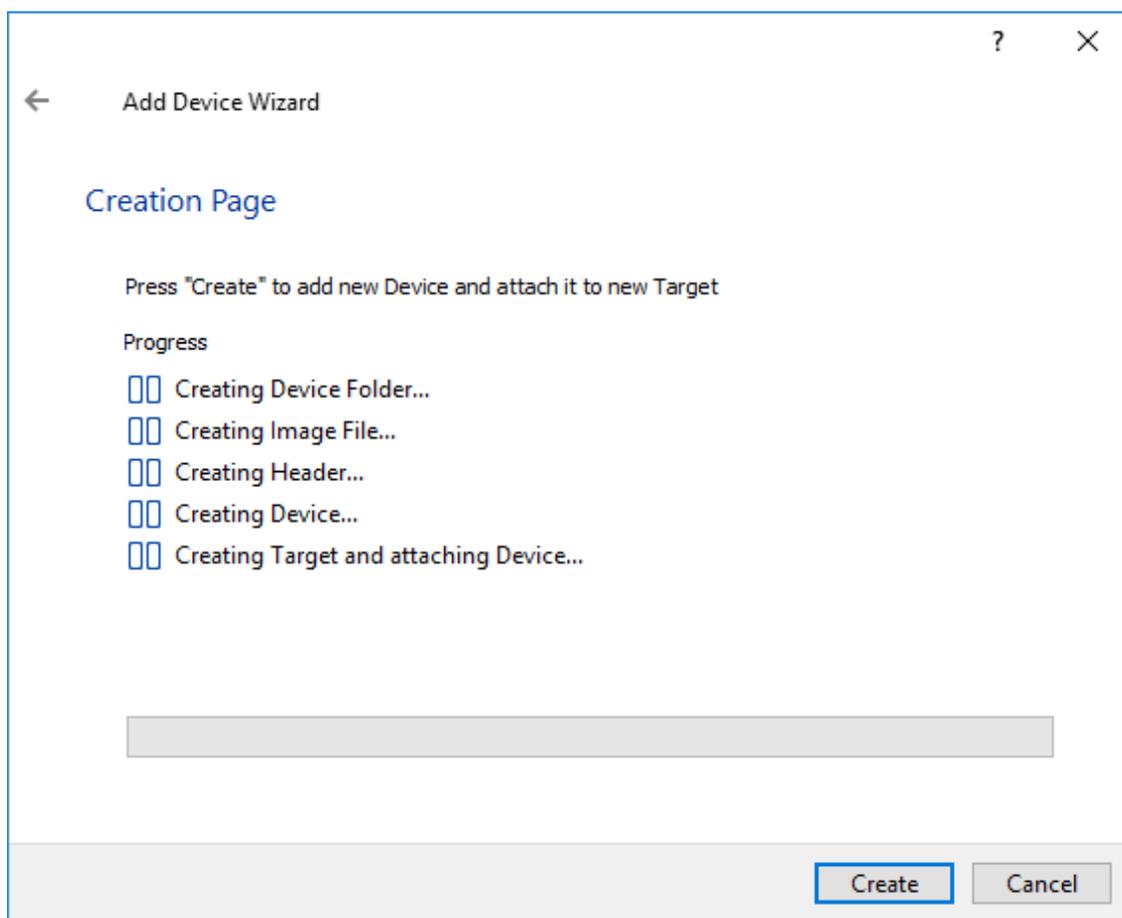
9. Click Close to finish the device creation.

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. Select Synchronous "Two-Way" replication as a replication mode.
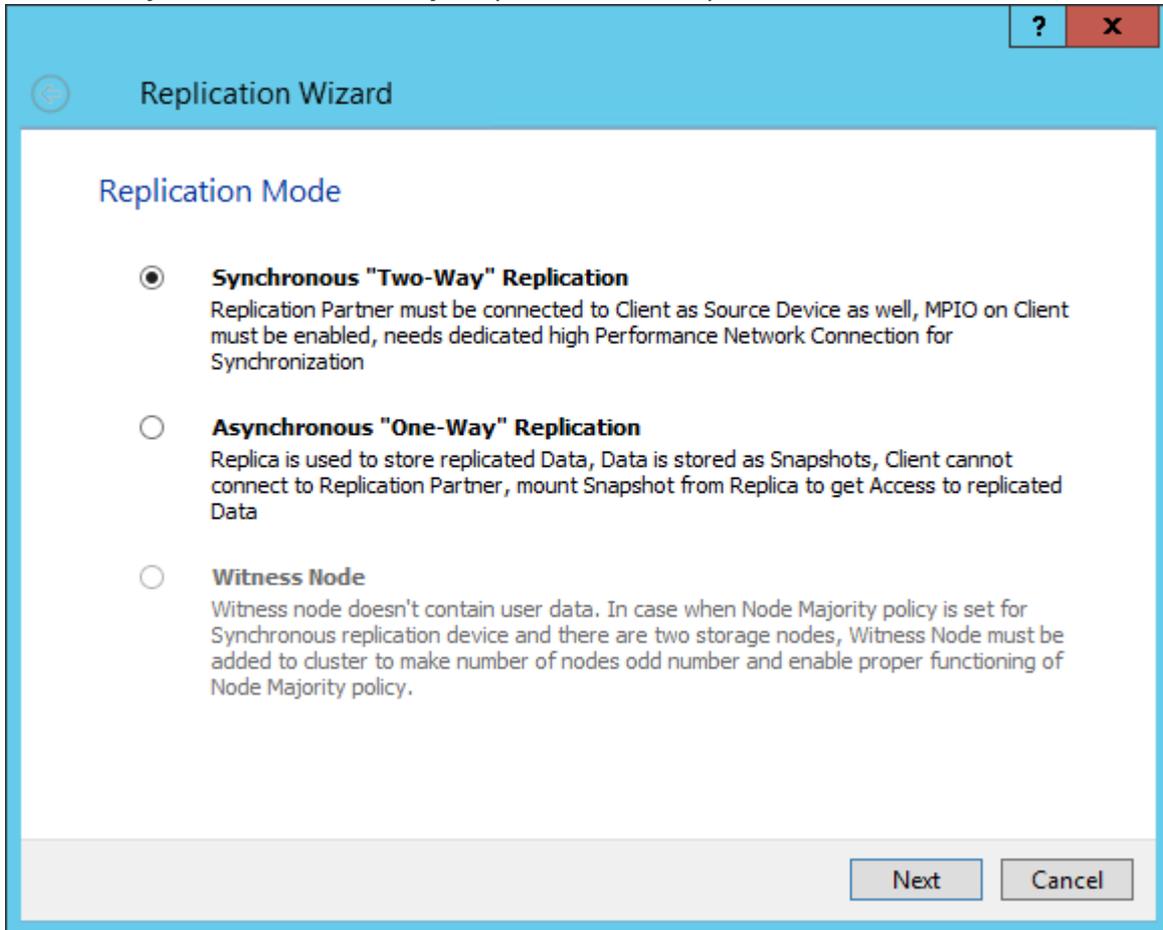
2. Specify a partner Host name or IP address and Port Number.

Asynchronous "One-Way" Replication

NOTE: Asynchronous replication requires minimum 100 MbE network bandwidth or higher. The Asynchronous node uses the LSFS device by design. Please, make sure that the Asynchronous node meets the LSFS device requirements:
https://knowledgebase.starwindsoftware.com/explanation/lsfs-container-technical-description/
1. Select Asynchronous "One-Way" Replication.
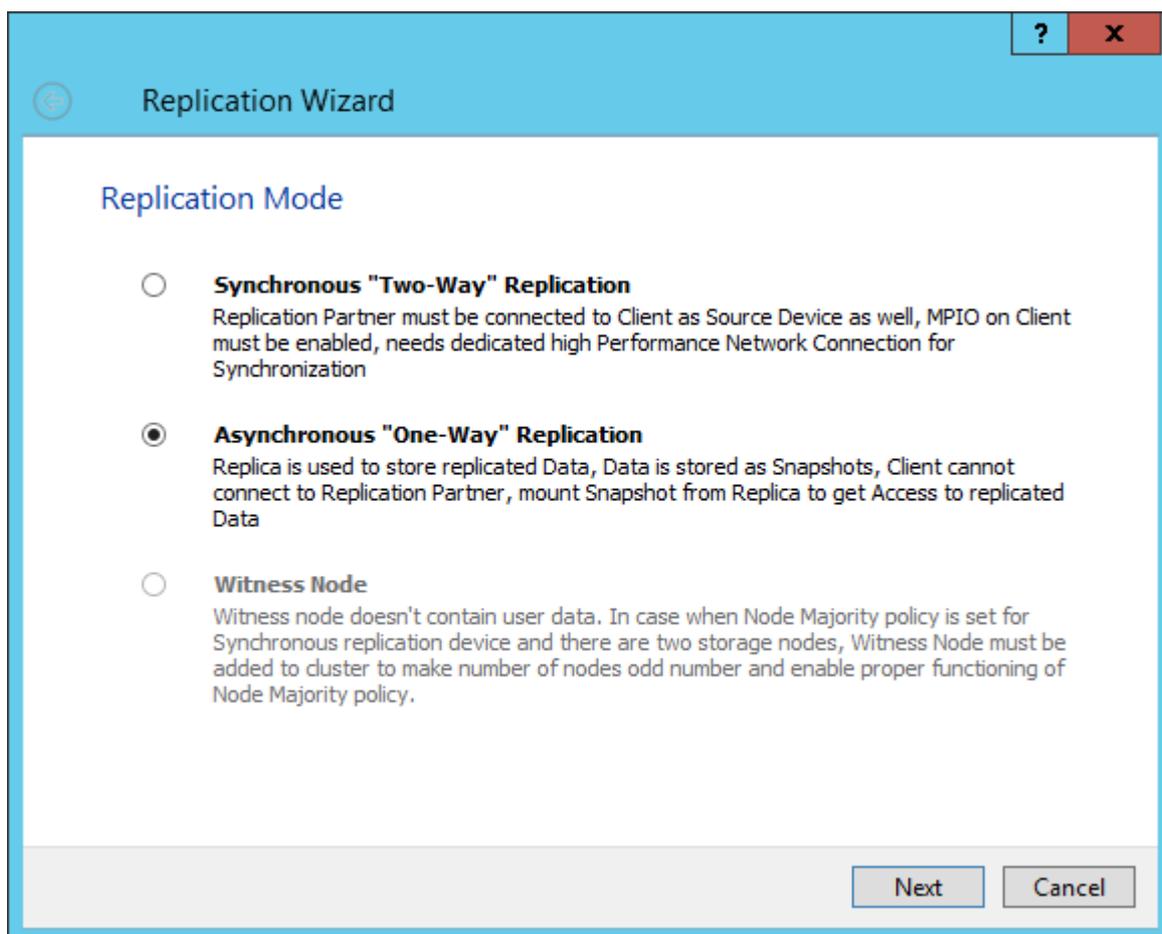
2. Enter Host name or IP address of the Asynchronous node.

3. Choose the Create New Partner Device option.

4. Specify the partner device Location . Optionally, modify the target name by clicking the appropriate button.

5. Click Change Network Settings.

6. Specify the network for asynchronous replication between the nodes. Click OK and then click Next.

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

8. Specify Scheduler Settings and click Next.

NOTE: The size of journal files and number of snapshots depends on the settings specified in this step.

9. Specify the path for journal files and click Next.

NOTE: By default, the journal files will be located on the node with the original device. However, it is highly recommended not to store journal files on the same drive where the original device is located. Additionally, the C:\ drive should not be used as the path for journal files to avoid any issues with Windows OS.

If the same drive where the StarWind device is located is selected, the warning message about possible performance issues will pop up. If there is no additional volume available for storing the journals, click I understand the potential problem. Use the selected path.
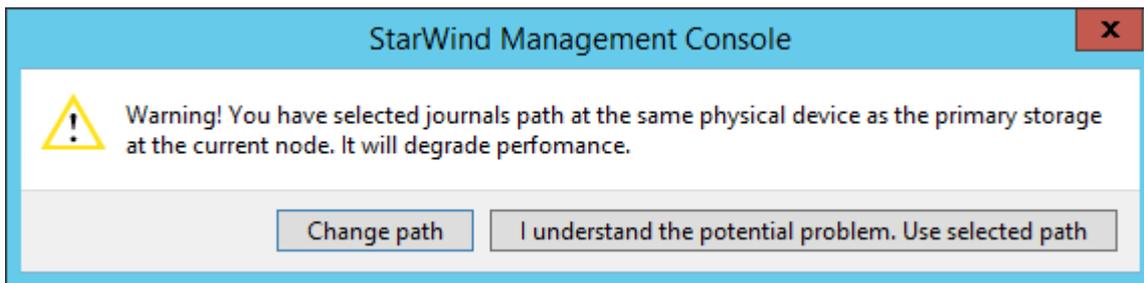
10. Press the Create Replica button.
11. Wait until StarWind service creates a device and click Close to complete the device creation.

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.
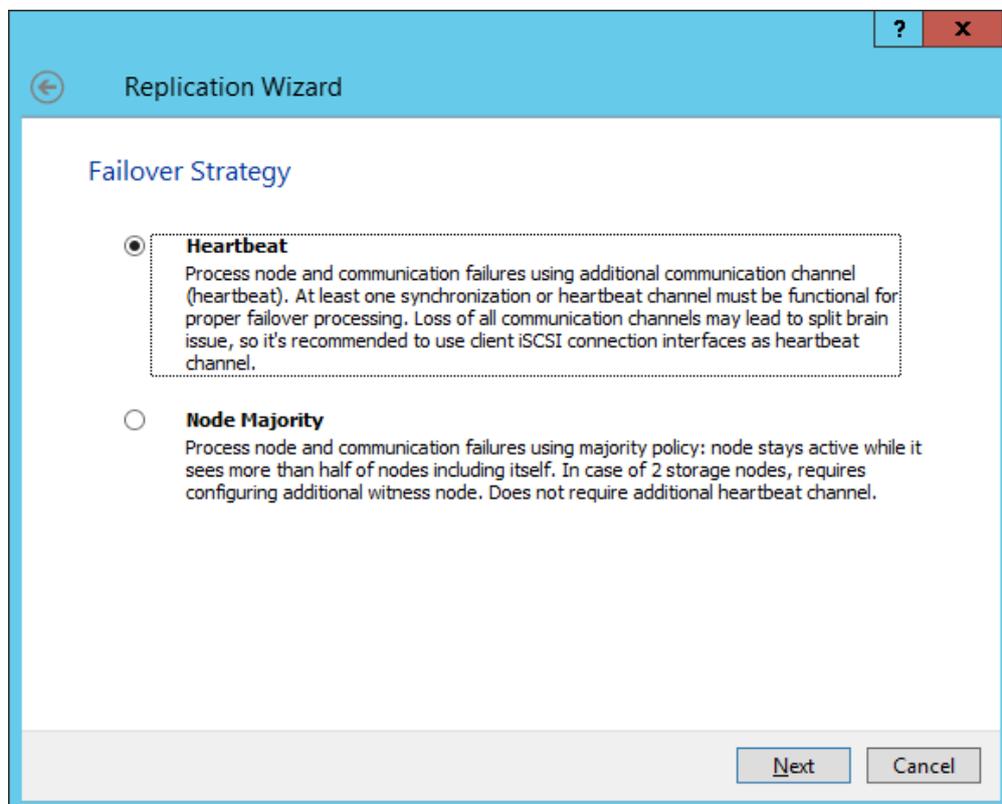
2. Select Create new Partner Device and click Next.

3. Select a partner device Location.

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

RAM-based (default) synchronization journal is placed in RAM. Synchronization with RAM journal provides good I/O performance in any scenario. Full synchronization could occur in the cases described in this KB:
https://knowledgebase.starwindsoftware.com/explanation/reasons-why-full-synchronization-may-start/

Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes. Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are recommended as the device performance is defined by the disk speed, where the journal is located. For example, it can be placed on the OS boot volume.
It is required to allocate 2 MB of disk space for the synchronization journal per 1 TB of HA device size with a disk-based journal configured with 2-way replication and 4MB per 1 TB of HA device size for 3-way replication.

Failure journal
The strategy provides good I/O performance, as a RAM-based journal, while all device nodes are in a healthy synchronized state. If a device on one node went into a not synchronized state, the disk-based journal activates and a performance drop could occur as the device performance is defined by the disk speed, where the journal is located. Fast synchronization is not guaranteed in all cases. For example, if a simultaneous hard reset of all nodes occurs, full synchronization will occur.
Continuous journal
The strategy guarantees fast synchronization and data consistency in all cases. Although, this strategy has the worst I/O performance, because of frequent write operations to the journal, located on the disk, where the journal is located.

5. Click Change Network Settings.
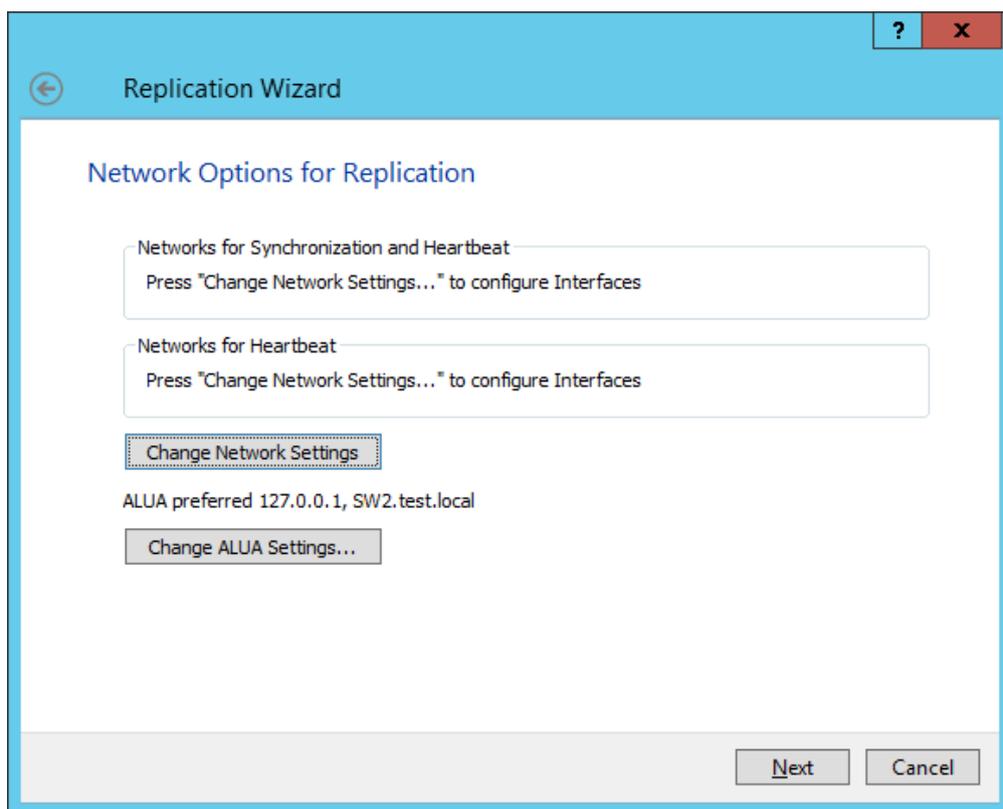
6. Specify the interfaces for Synchronization and Heartbeat Channels. Click OK and then click Next.
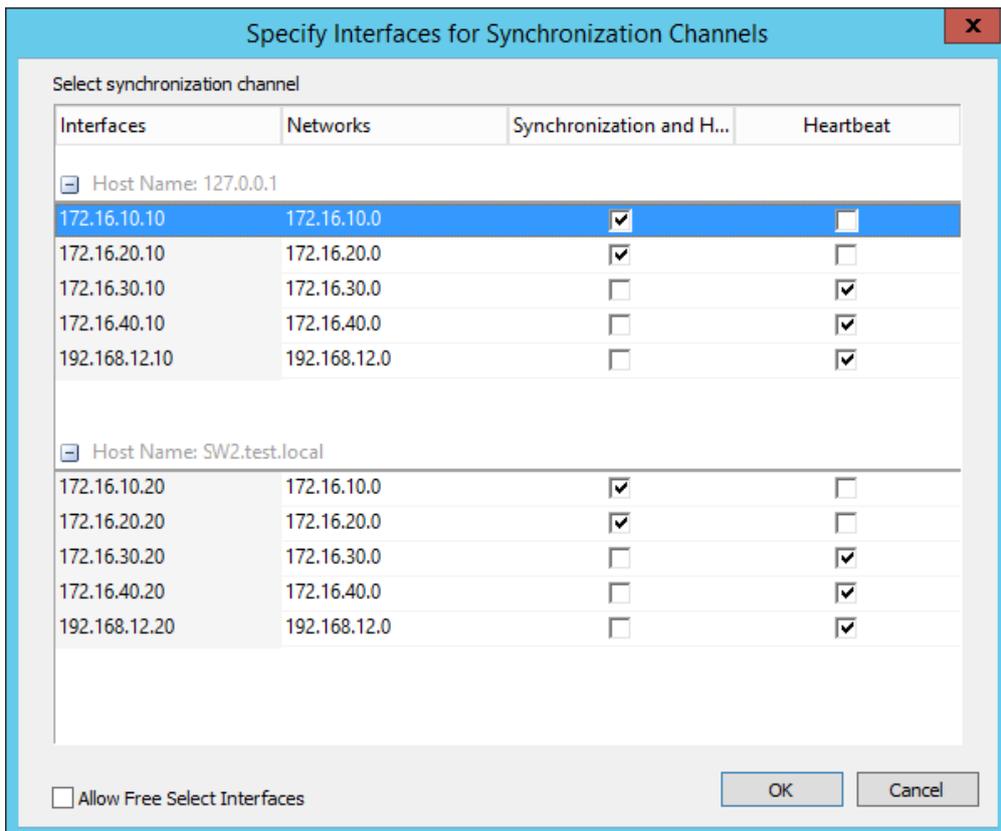
7. In Select Partner Device Initialization Mode, select Synchronize from existing Device and click Next.

8. Click Create Replica. Click Finish to close the wizard.
The successfully added device appears in StarWind Management Console.

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

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.
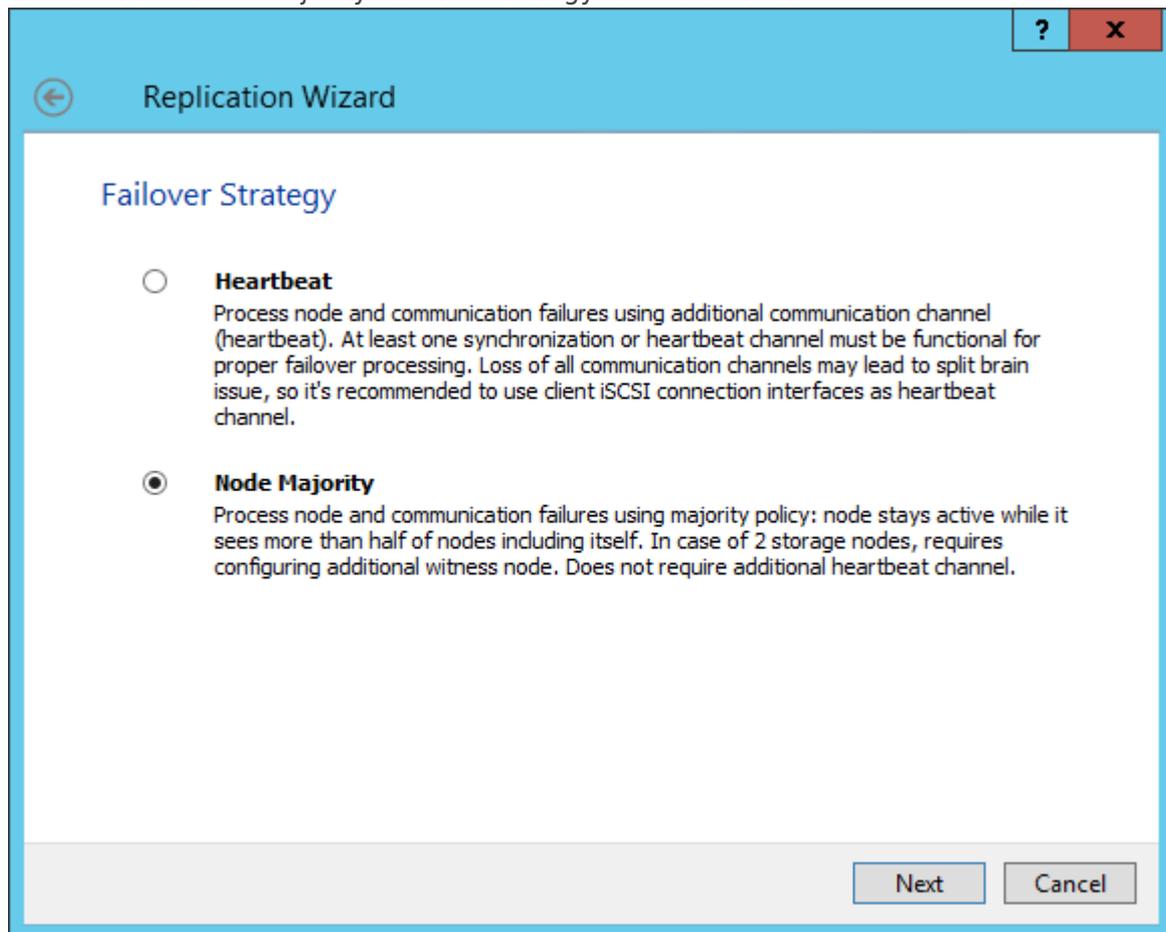
2. Choose Create new Partner Device and click Next.
3. Specify the partner device Location and modify the target name if necessary.
Click Next. Select Synchronization Journal strategy and location and click Next.
4. In Network Options for Replication, press the Change network settings button and select 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.
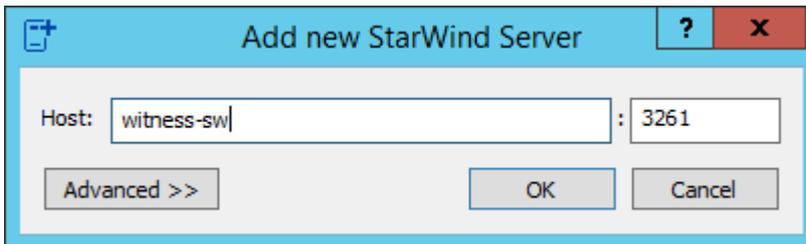
2. Right-click on the HA device with the configured Node Majority failover policy and select Replication Manager and press the Add Replica button.

3. Select Witness Node.
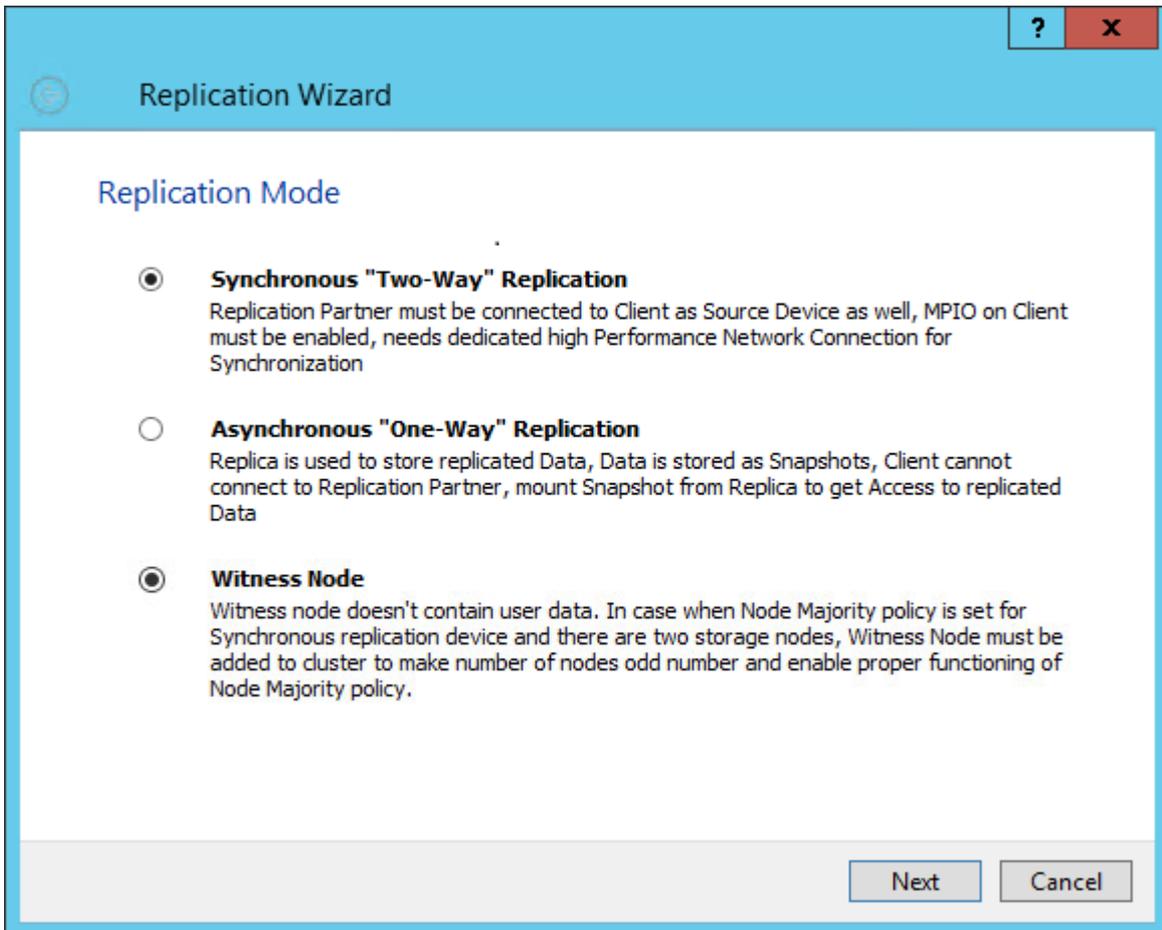
4. Specify the Witness node Host Name or IP address. The default Port Number is 3261.
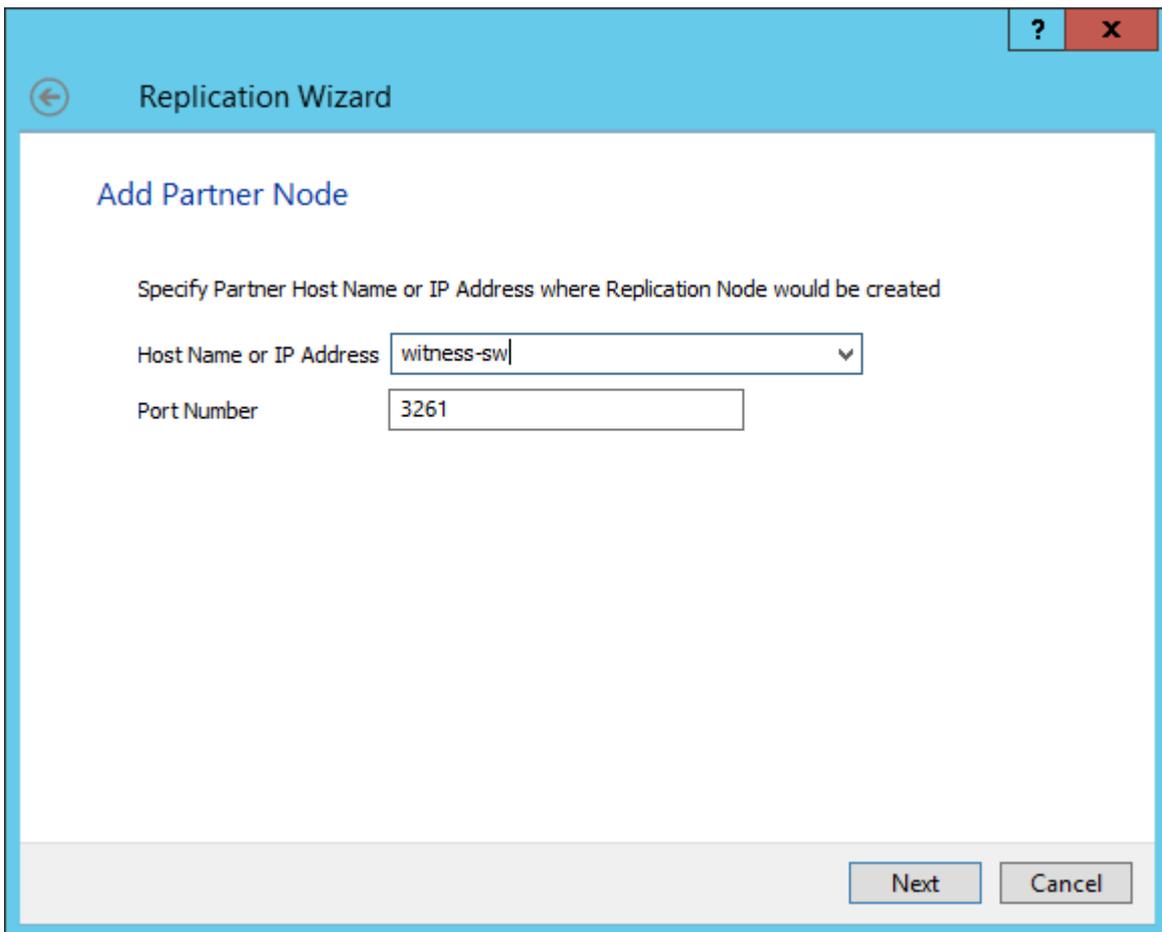
5. In Partner Device Setup, specify the Witness device Location. Optionally, modify the target name by clicking the appropriate button.

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 Hosts

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

interface:

```
iscsicpl
```

2. Navigate to the Discovery tab.

3. Click the Discover Portal button. The Discover Target Portal dialog appears. Type 172.16.30.10
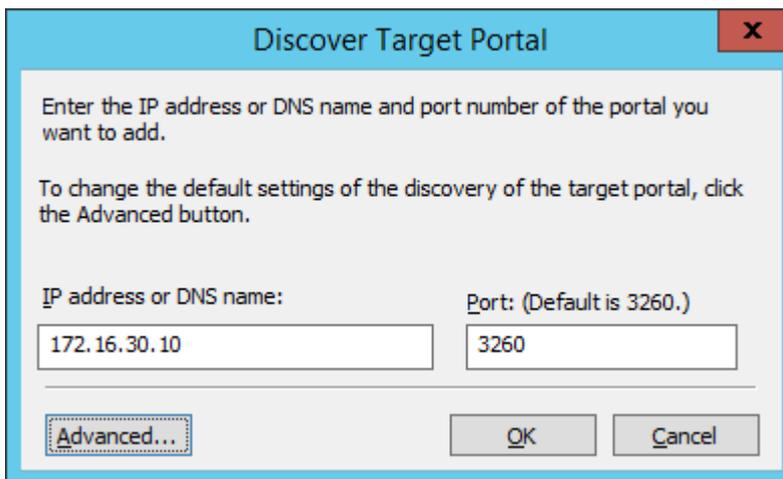
Click the Advanced button.

4. Select Microsoft iSCSI Initiator as Local adapter and select Cluster Node 1 initiator IP address from the same subnet. Click OK twice to add the Portal.
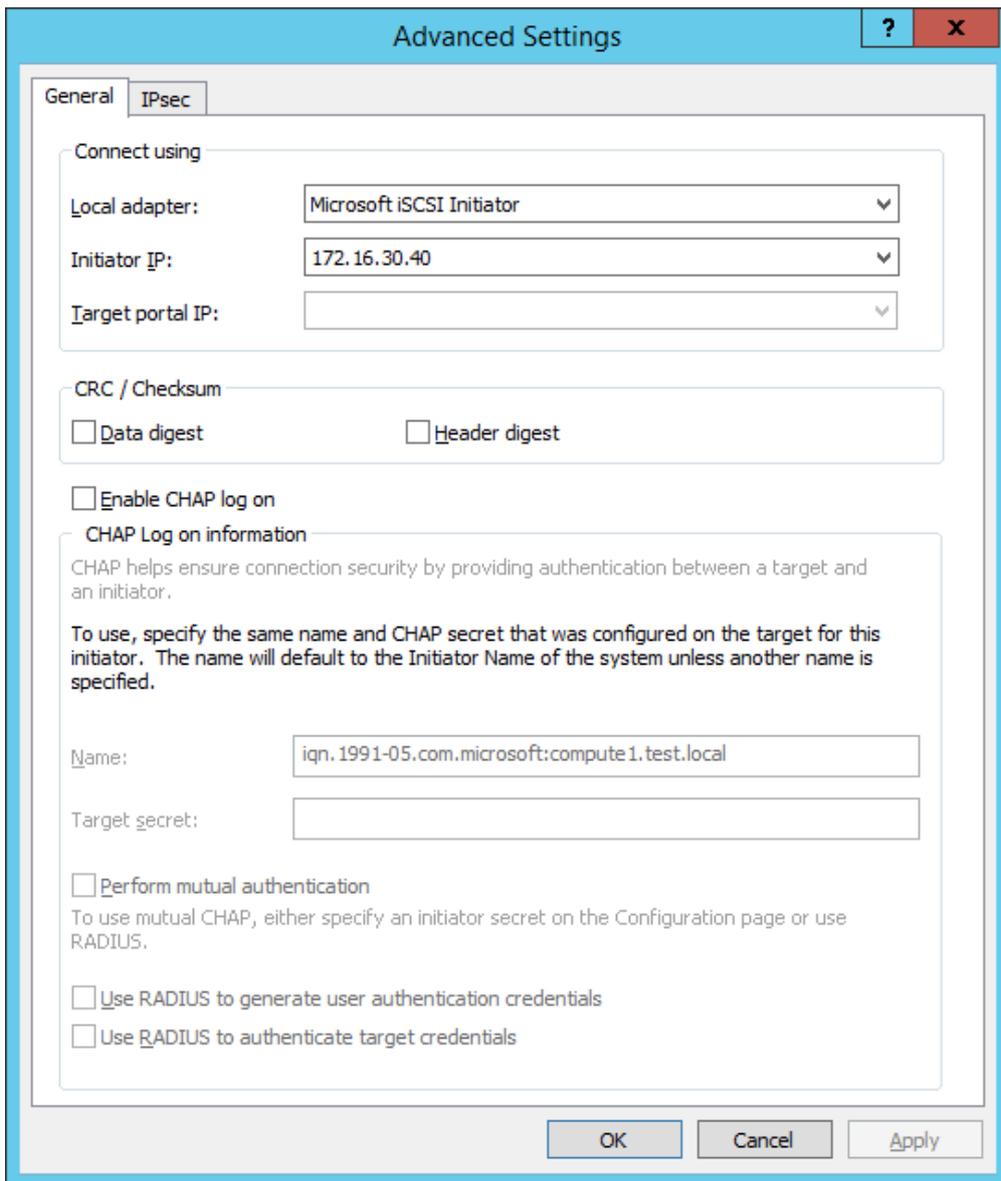
5. Click the Discover Portal button once again.

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.
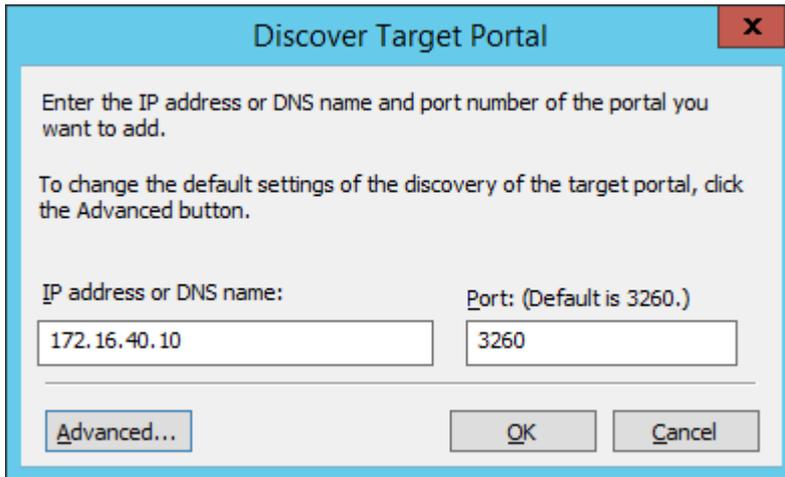
7. Select Microsoft iSCSI Initiator as the Local adapter, select the Initiator IP in the same subnet as the IP address of the partner server from the previous step. Confirm the actions to complete the Target Portal discovery.
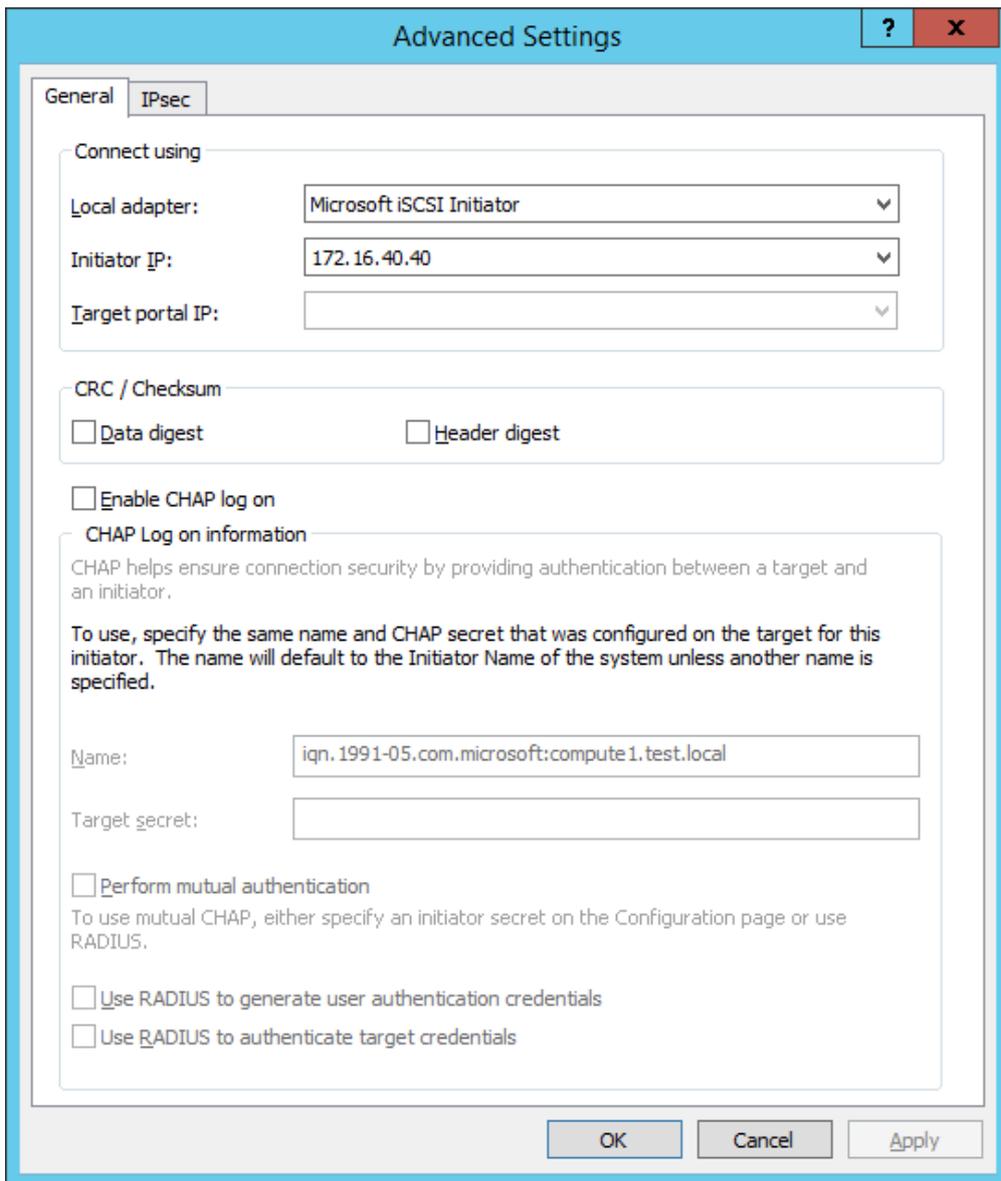
8. Now, all the target portals are added on the first node.

9. To Discover Targets Portals from the second StarWind node, click the Discover Portal button once again, enter iSCSI IP address for the second StarWind Node (172.16.30.20). Click the Advanced button.
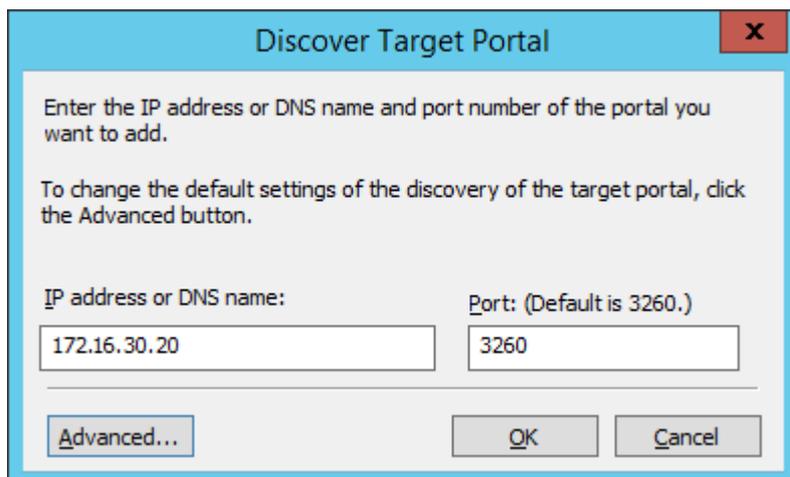
10. Select Microsoft iSCSI Initiator as Local adapter and select Cluster Node 1 initiator IP address from the same subnet. Click OK twice to add the Portal.
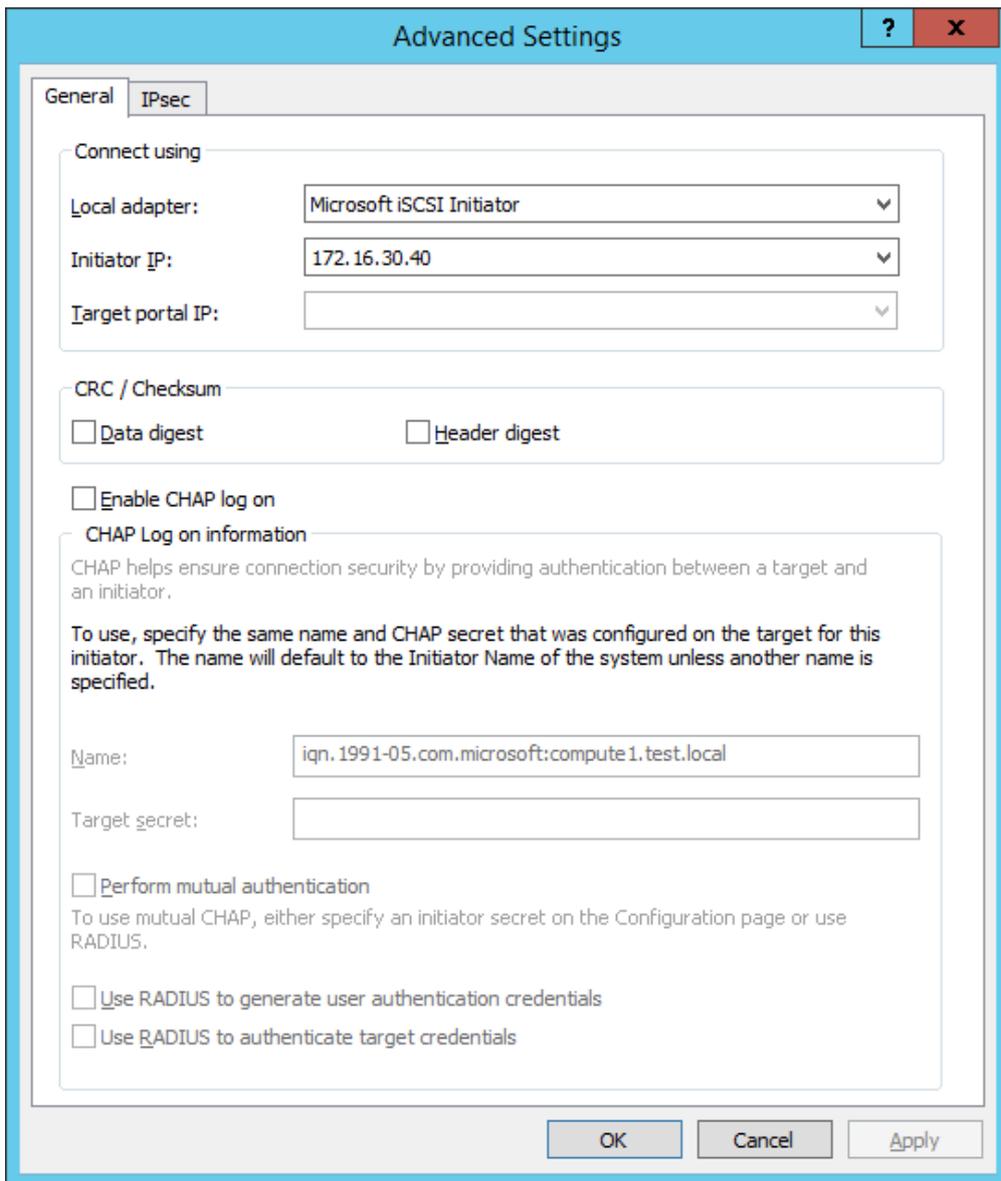
11. Click the Discover Portal button once again. In Discover Target Portal dialog, enter another iSCSI IP address of the second StarWind Node (172.16.40.20). Click the Advanced button.
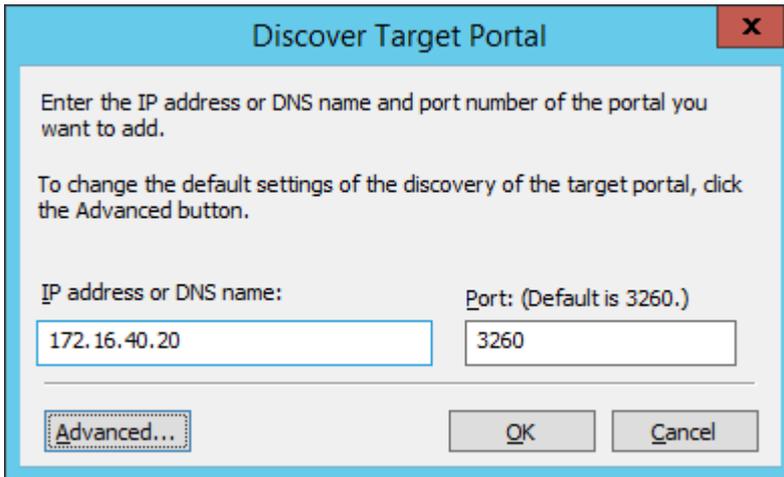
12. Select Microsoft iSCSI Initiator as Local adapter and select the Cluster Node 1 initiator IP address from the same subnet. Click OK twice to add the Portal.
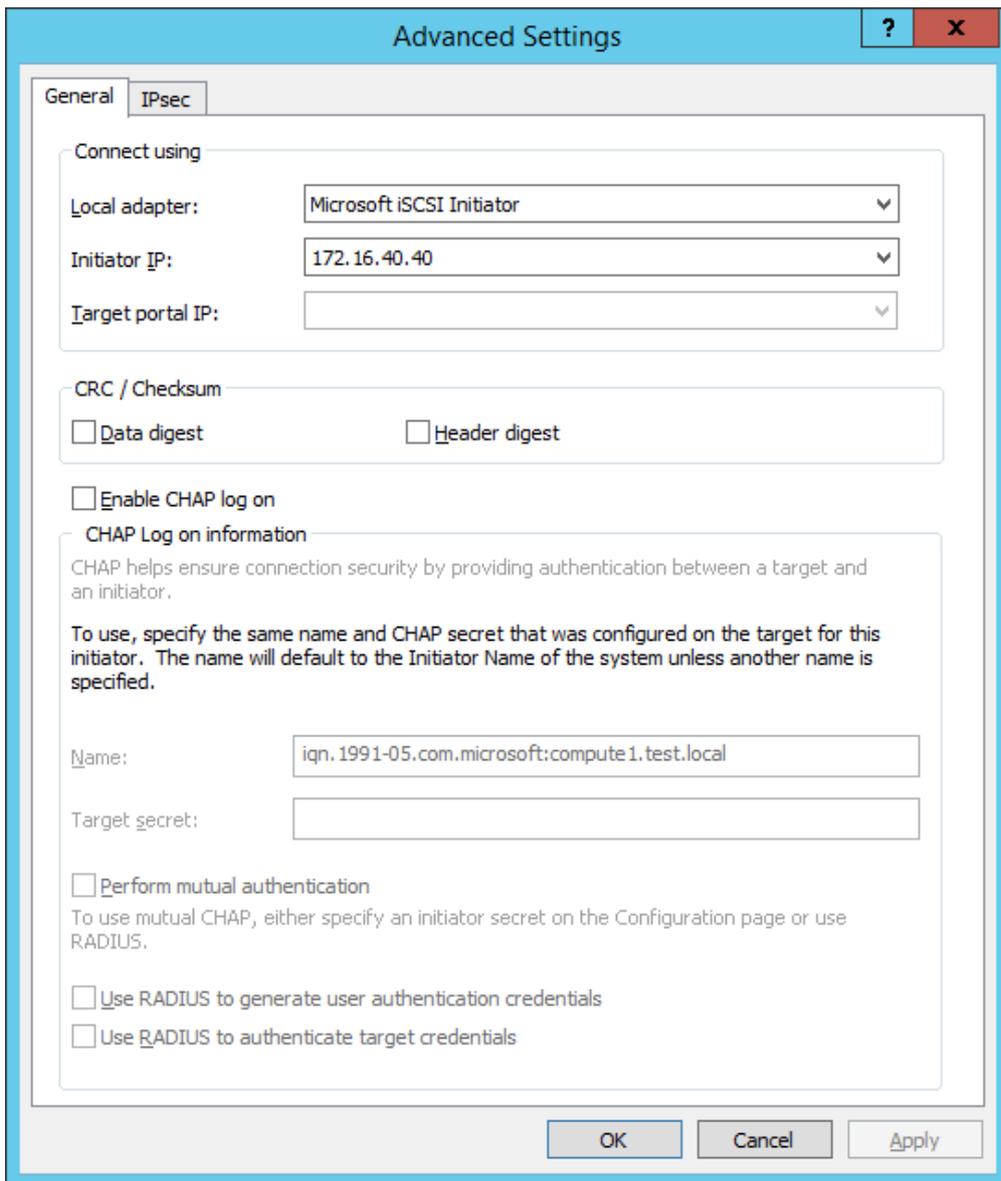
13. All target portals are successfully added to the Cluster Node 1.
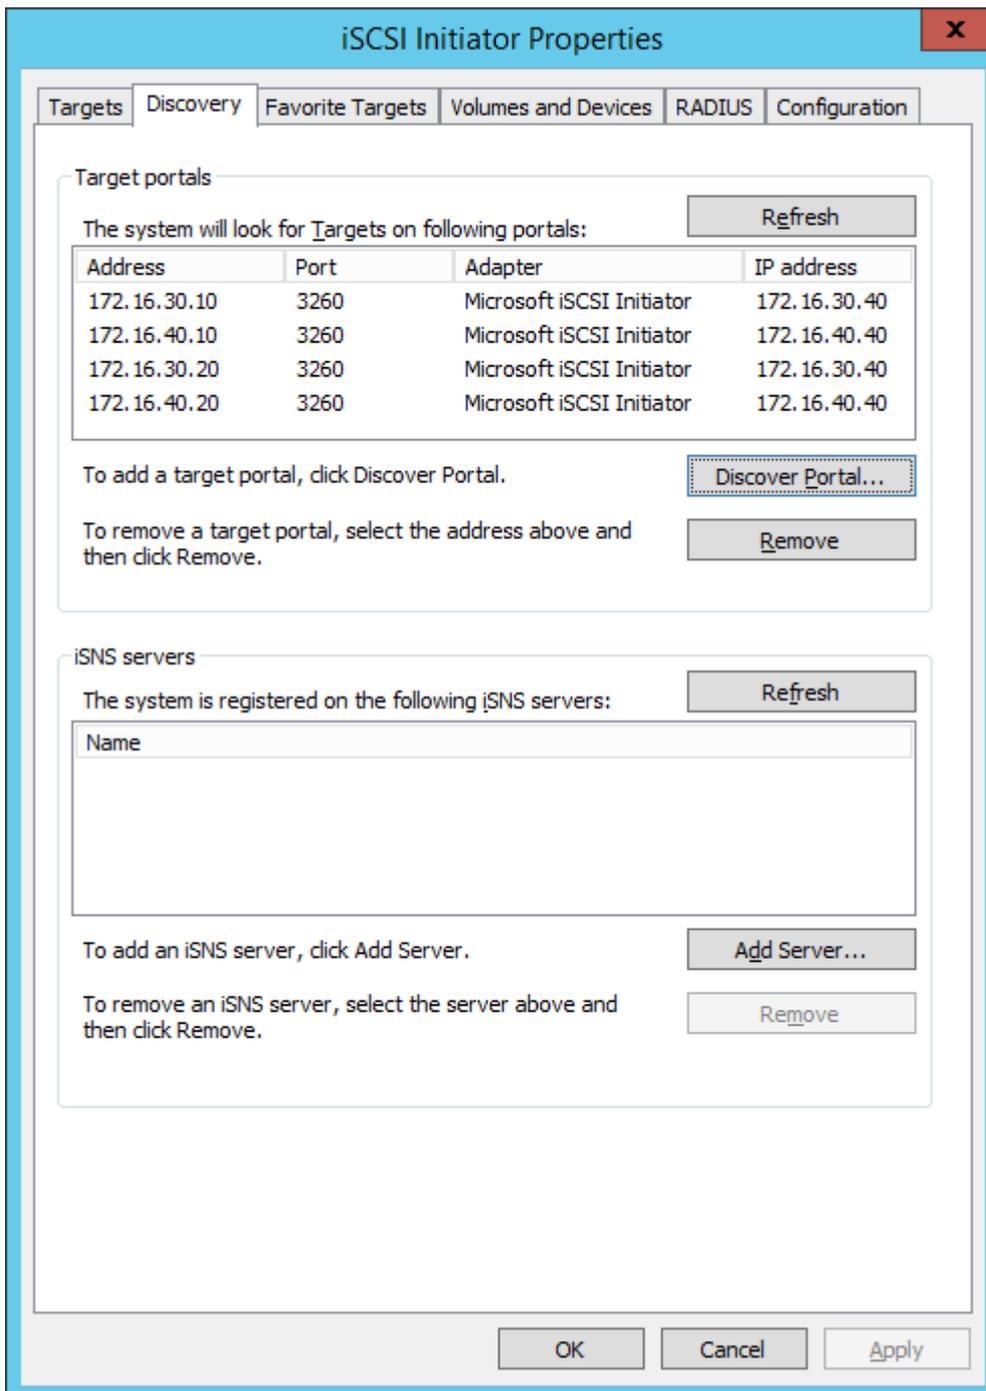
14. Repeat the steps 1-13 on the partner node.

Connecting Targets

1. Launch Microsoft iSCSI Initiator on Cluster Node 1 and click on the Targets tab. The previously created targets should be 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).

2. Select a target discovered from the first StarWind Node and click Connect.
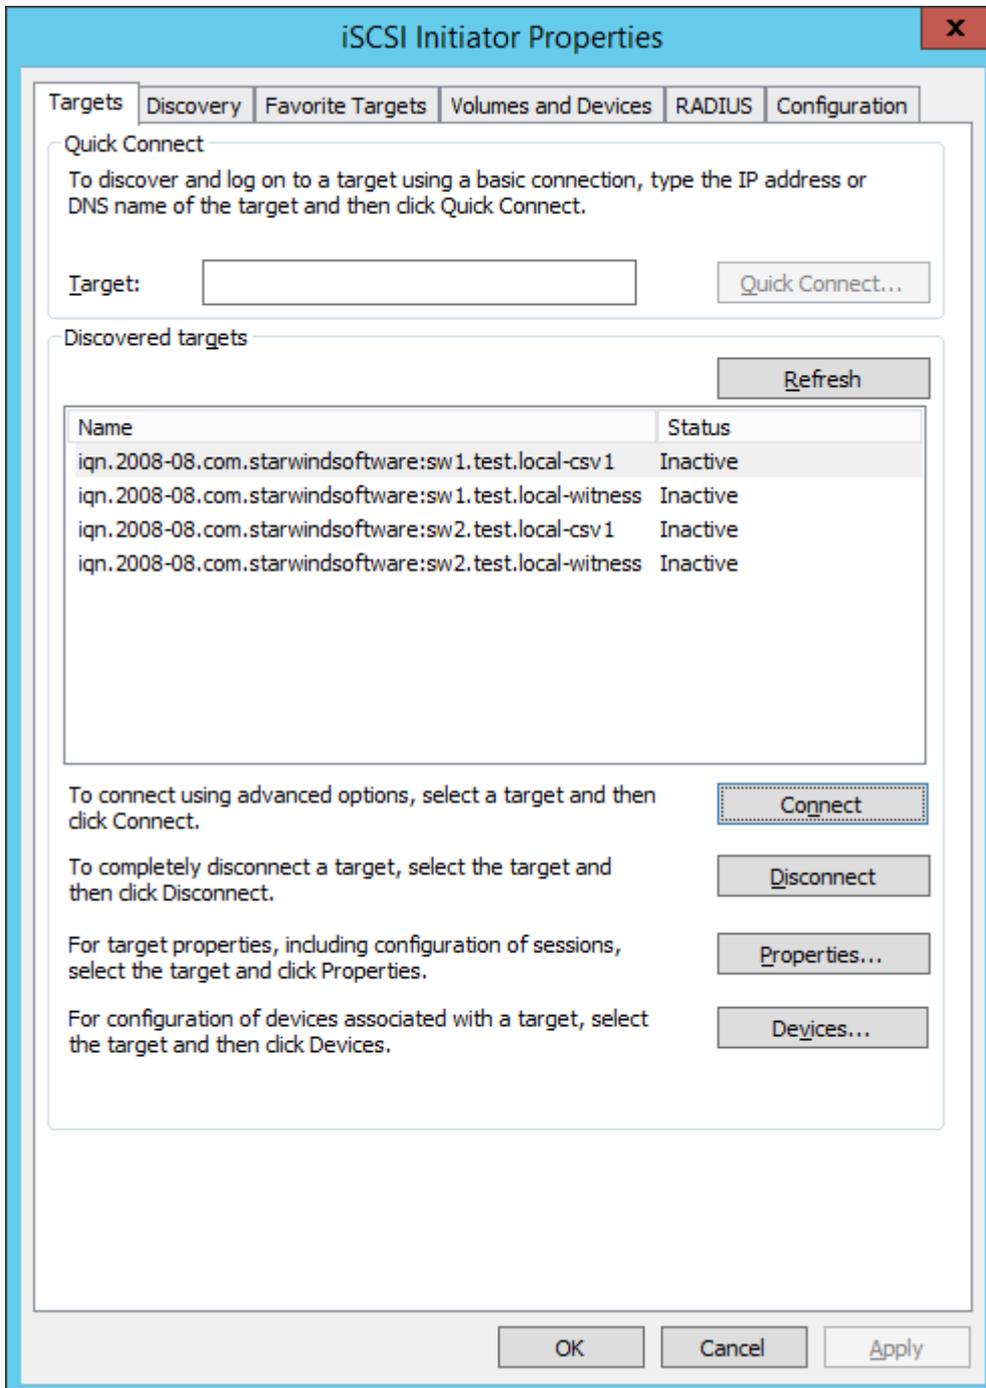
3. Enable checkboxes as shown in the image below. Click Advanced.
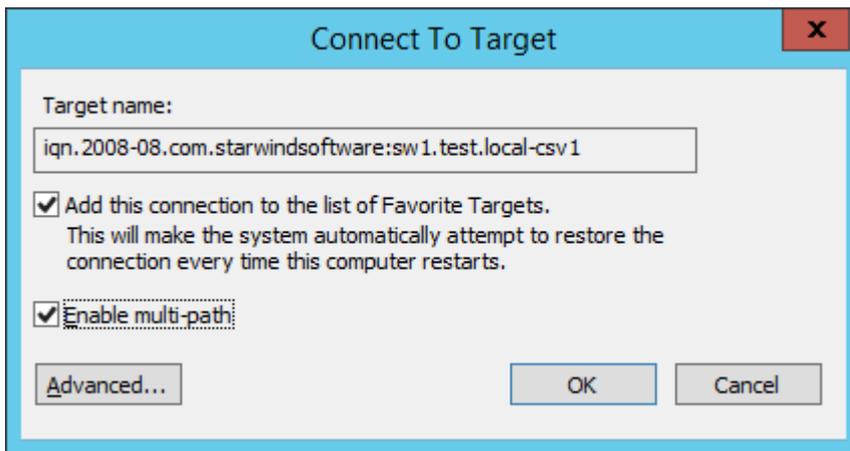
4. Select Microsoft iSCSI Initiator in the Local adapter dropdown menu. In the Target portal IP, select the IP address of the first StarWind Node and Initiator IP address from the same subnet. Click OK twice to connect the target.
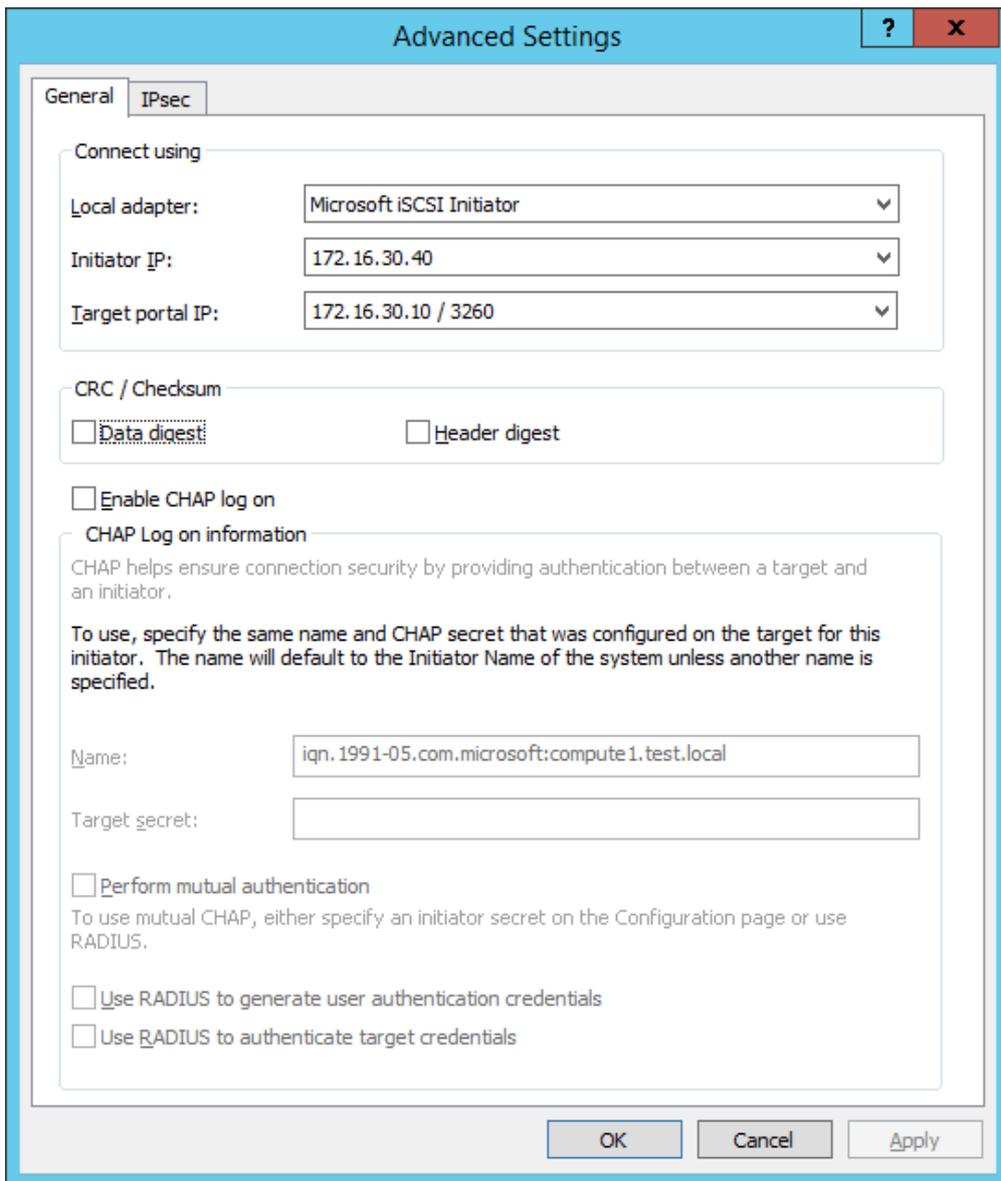
5. To connect the same target via another subnet, select it once again and click Connect.

6. Enable checkboxes like in the image below and click Advanced.
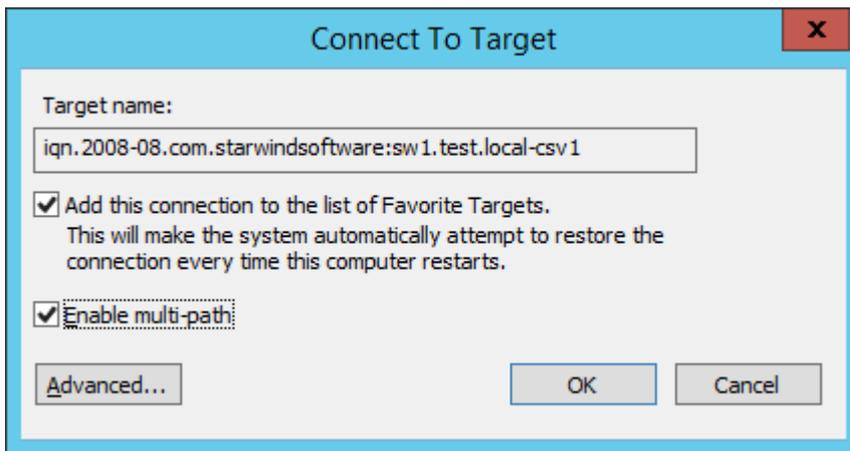
7. Select Microsoft iSCSI Initiator in the Local adapter text field.

8. In the Target portal IP select another IP address of the first StarWind
Node and Initiator IP address from the same subnet. Click OK twice to connect the target.
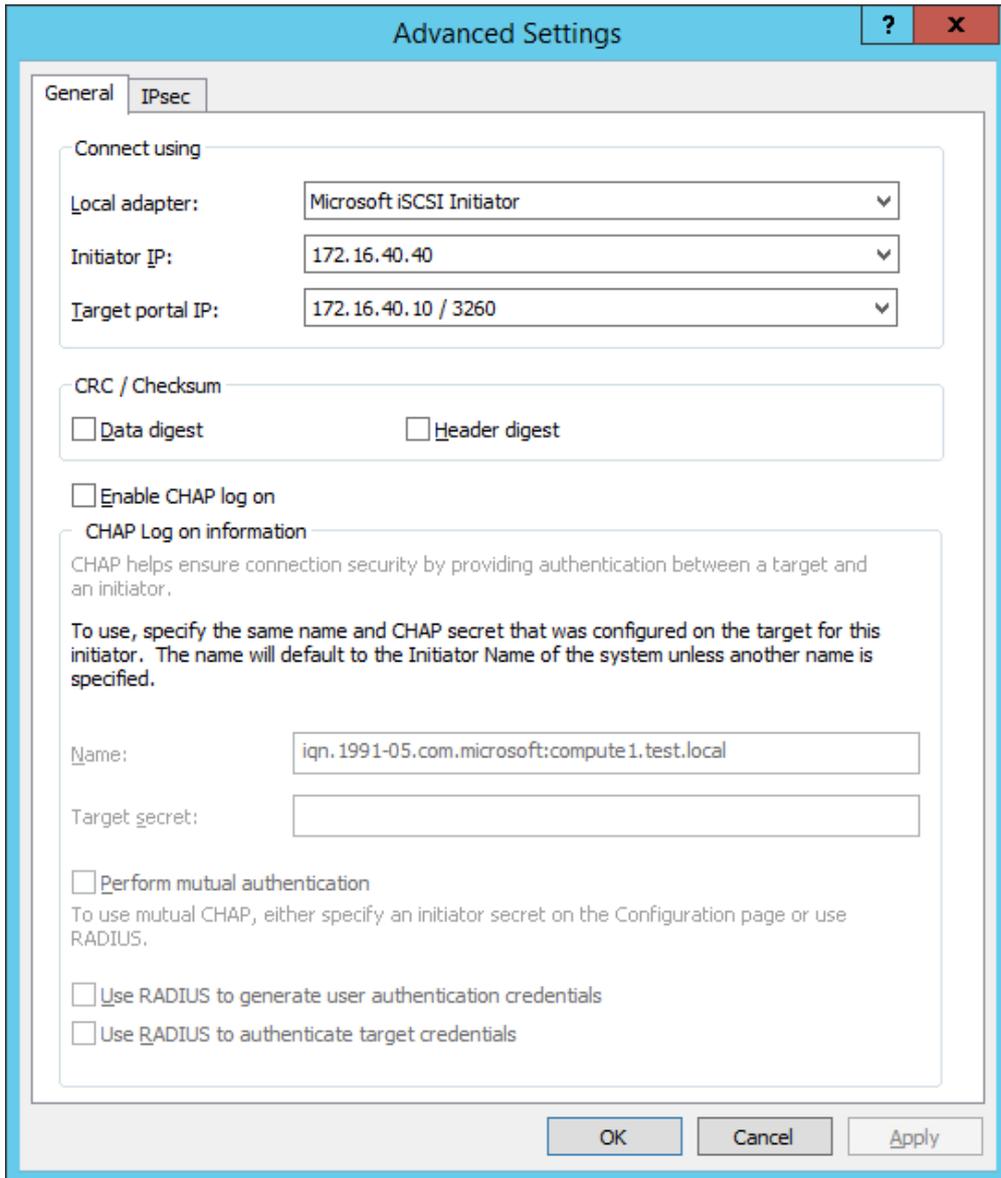
9. Select the partner target discovered from the second StarWind node and click Connect.
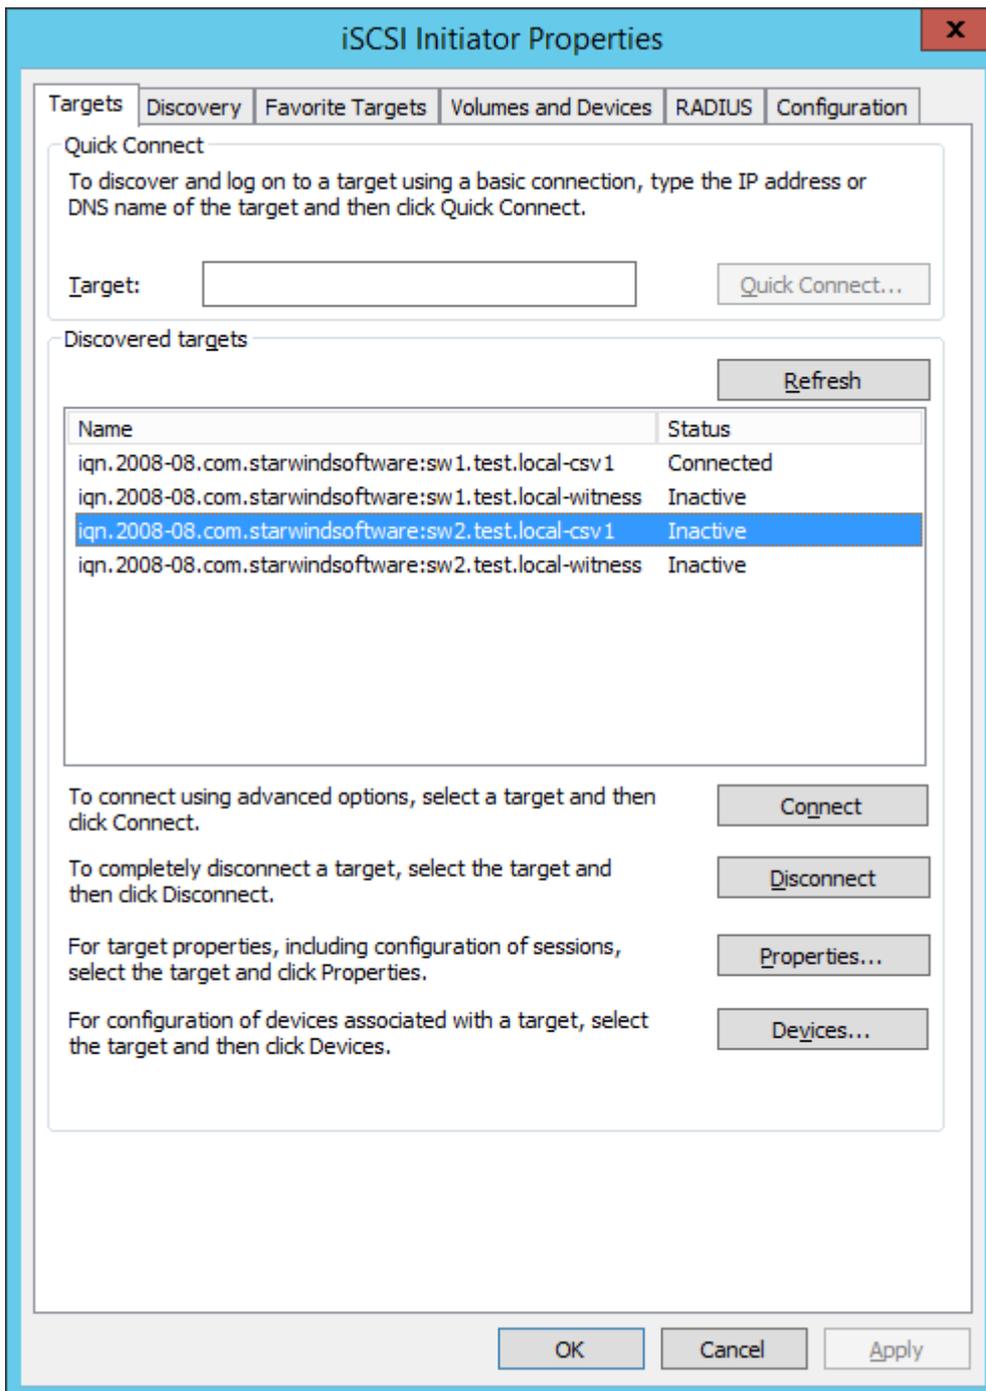
10. Enable checkboxes like in the image below and click Advanced.
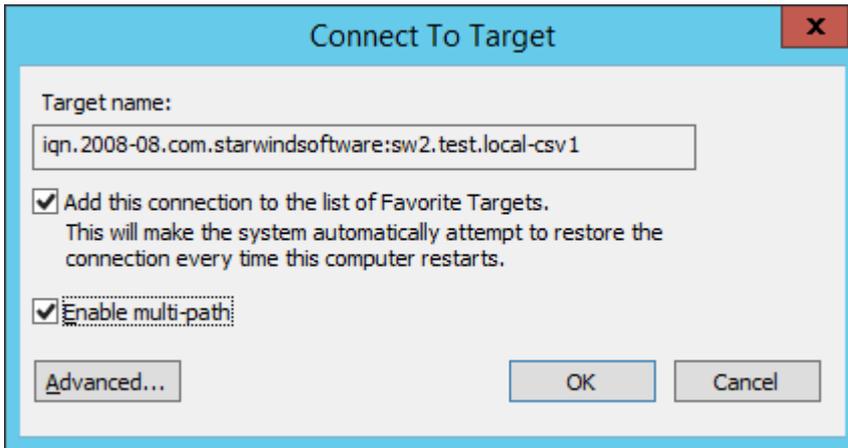
11. Select Microsoft iSCSI Initiator in the Local adapter text field.

12. In Target portal IP, select the IP address of the second StarWind Node and Initiator IP address from the same subnet. Click OK twice to connect the target.
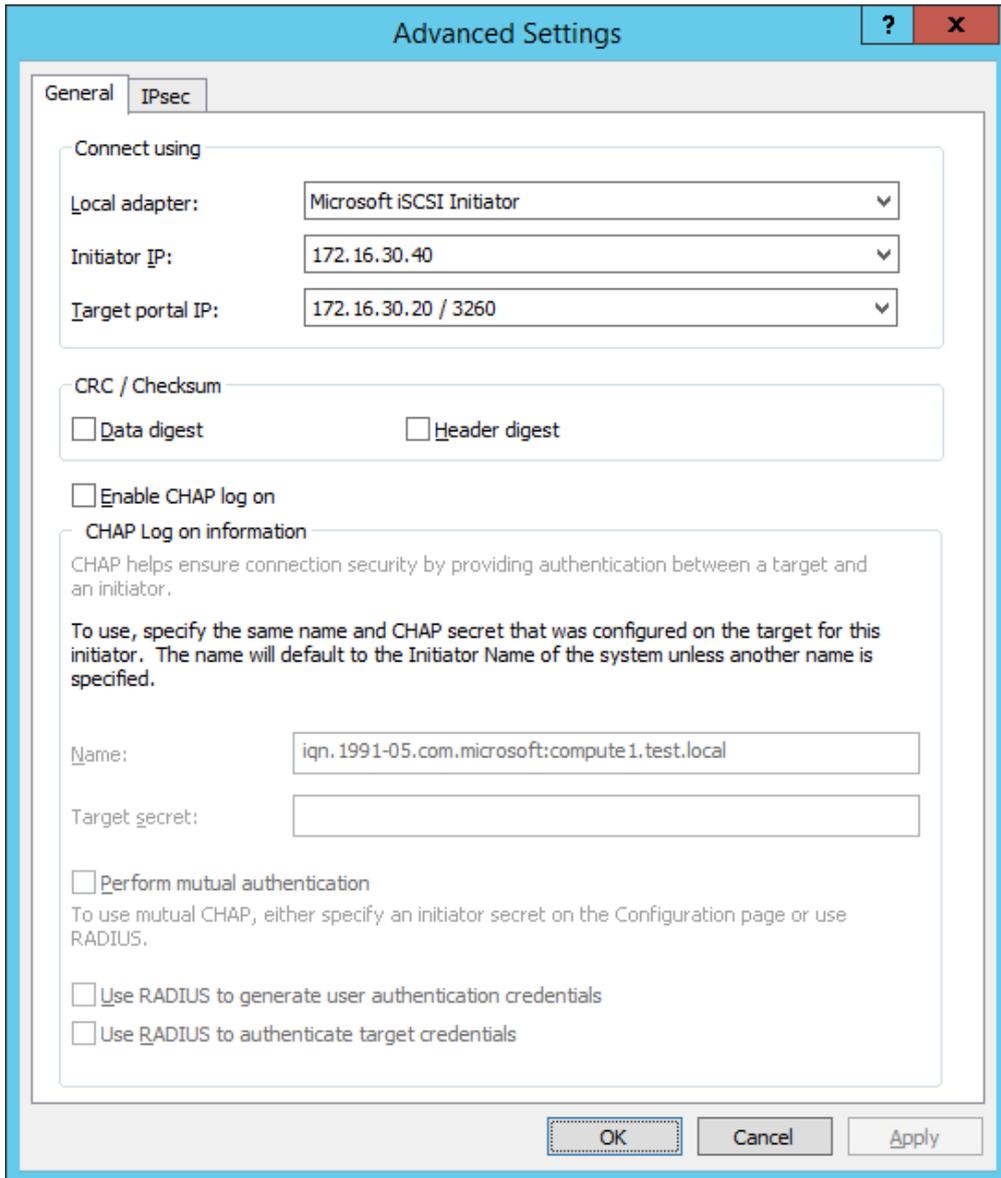
13. To connect the same target via another subnet, select it once again and click Connect.

14. Enable checkboxes like in the screenshot below and click Advanced.
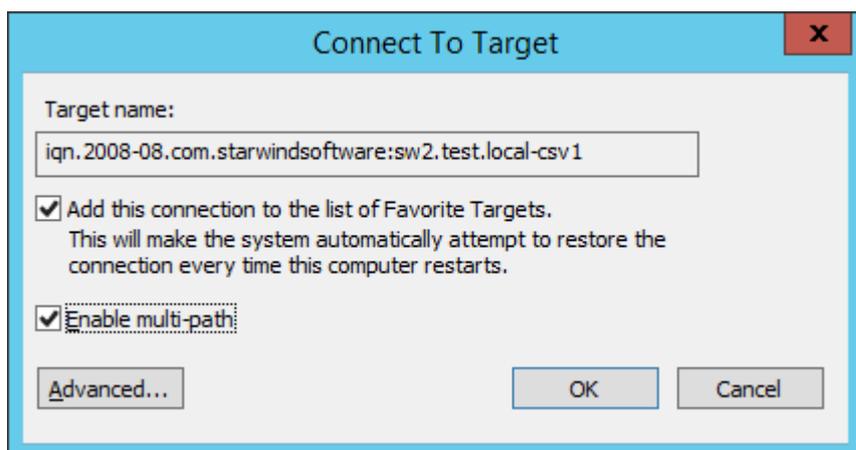
15. In the Target portal IP select another IP address of the second StarWind
Node and Initiator IP address from the same subnet. Click OK twice to connect the target.
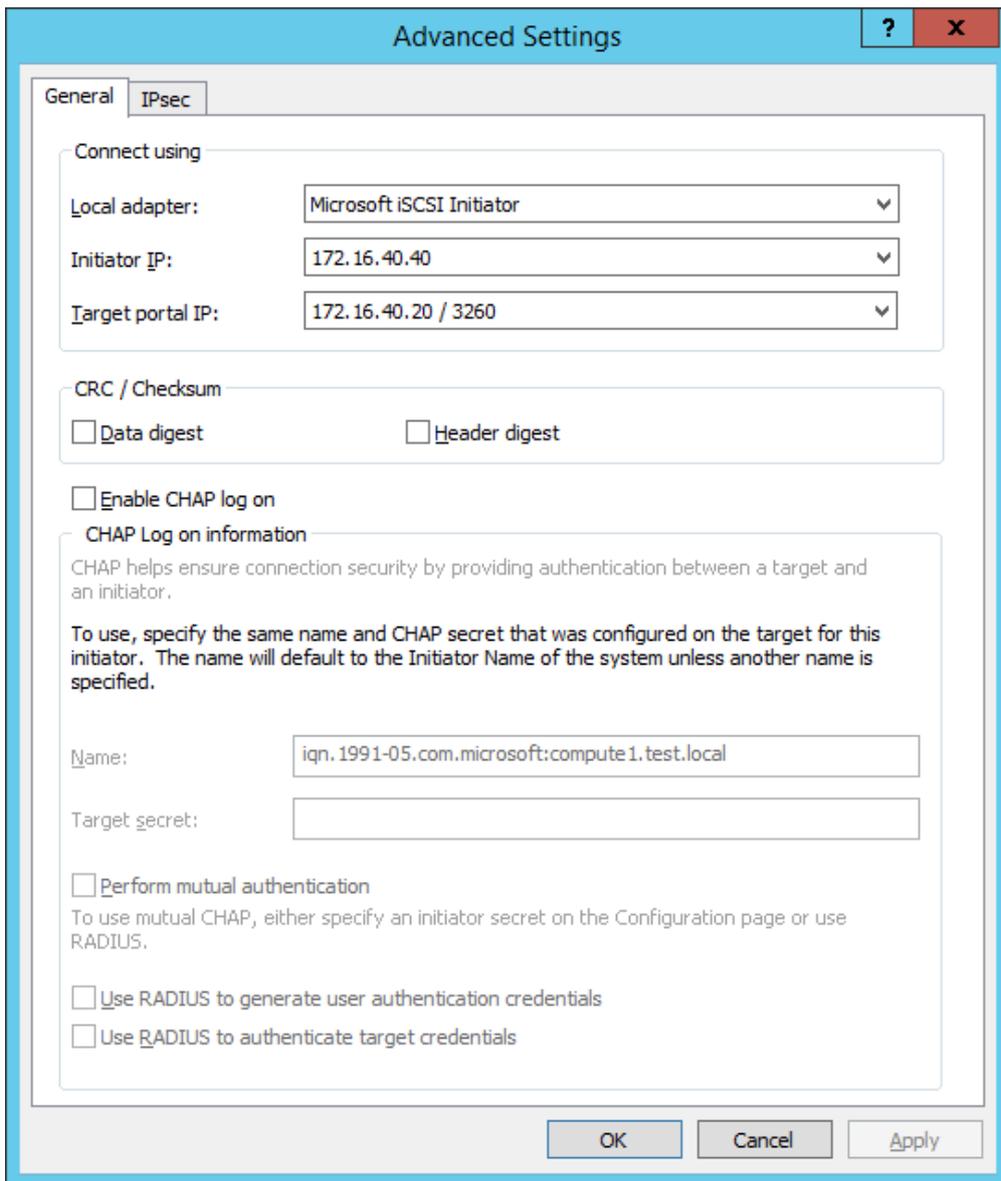
16. Repeat steps 1-15 for all HA device targets.
After that, repeat the steps 1-17 on the Cluster node 2, specifying corresponding IP addresses. The result should look like in the picture below.
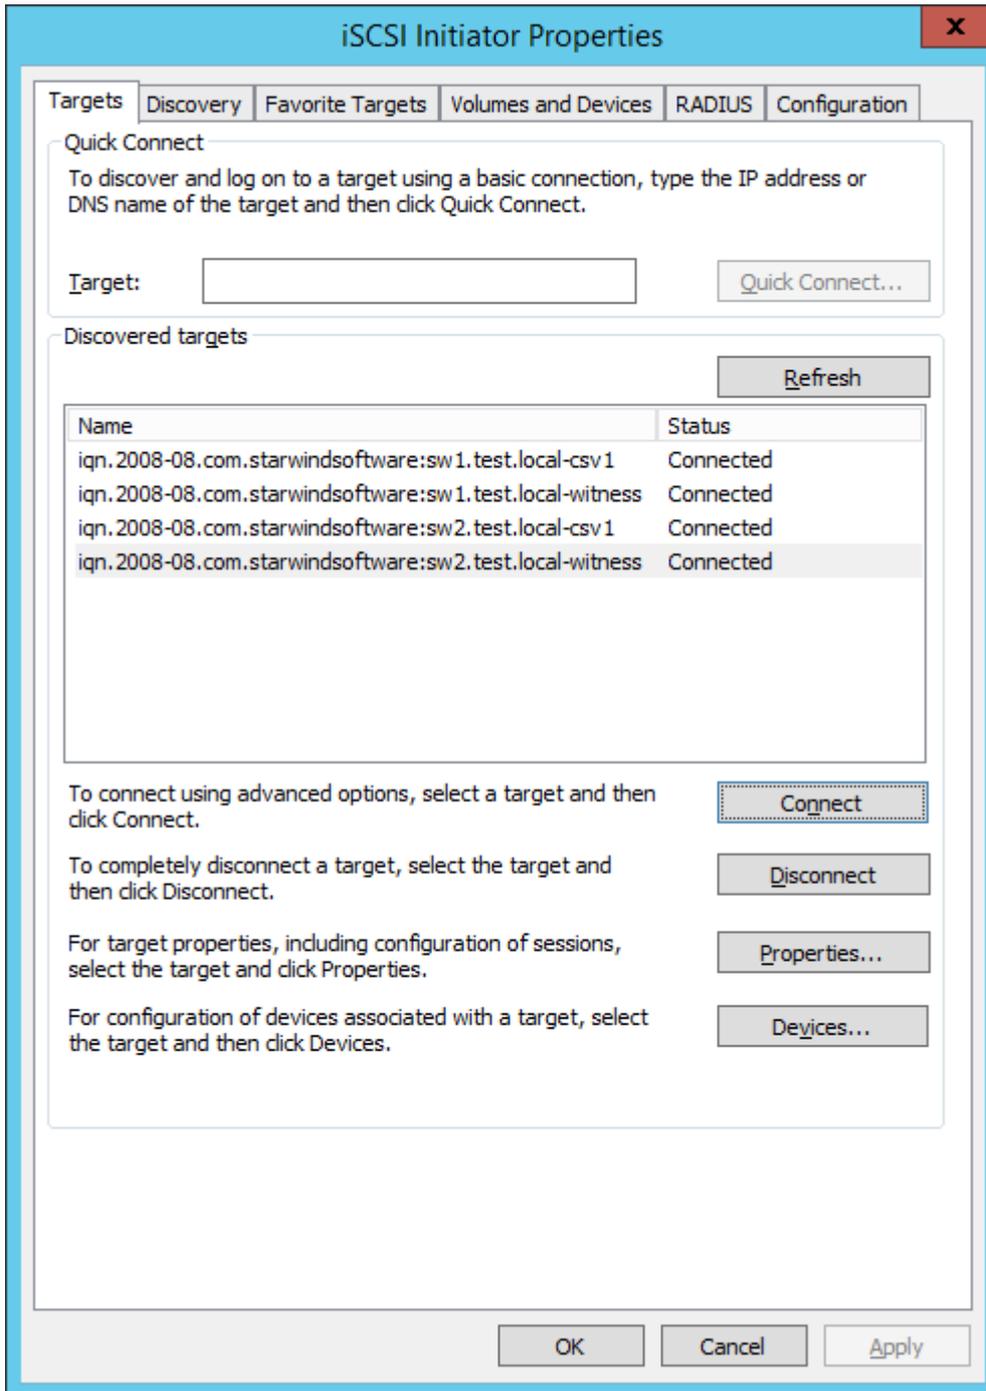
17. Initialize the disks and create partitions on them using the Disk Management snap-in. To create the cluster, the disk devices must be initialized and formatted on both nodes. NOTE: it is recommended to initialize the disks as GPT.

Creating A Failover Cluster In Windows Server 2012 R2

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

1. Open Server Manager. Select the Failover Cluster Manager item from the Tools menu.
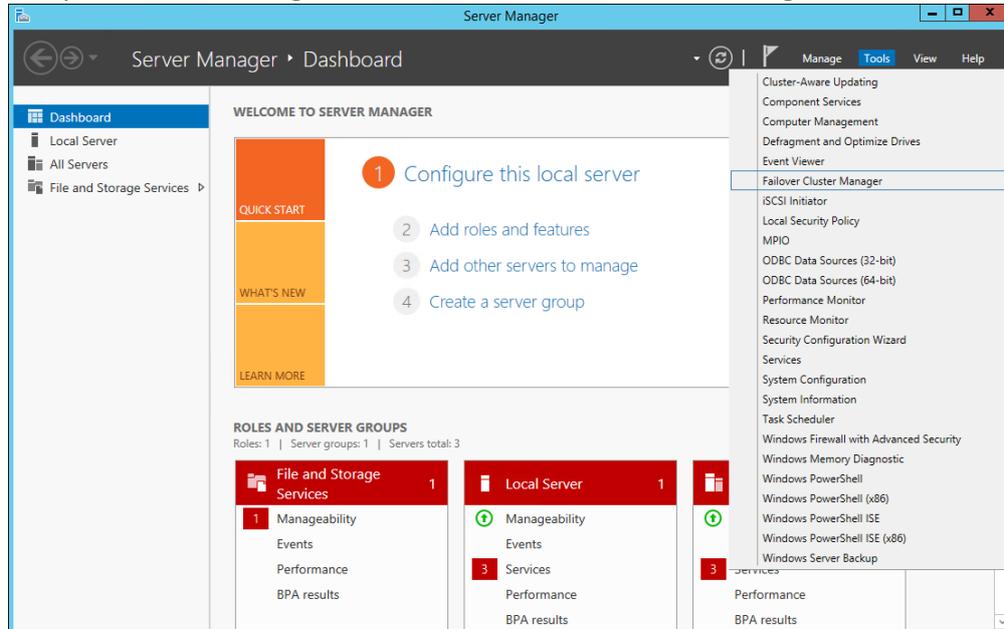
2. Click the Create Cluster link in the Actions section of Failover Cluster Manager.
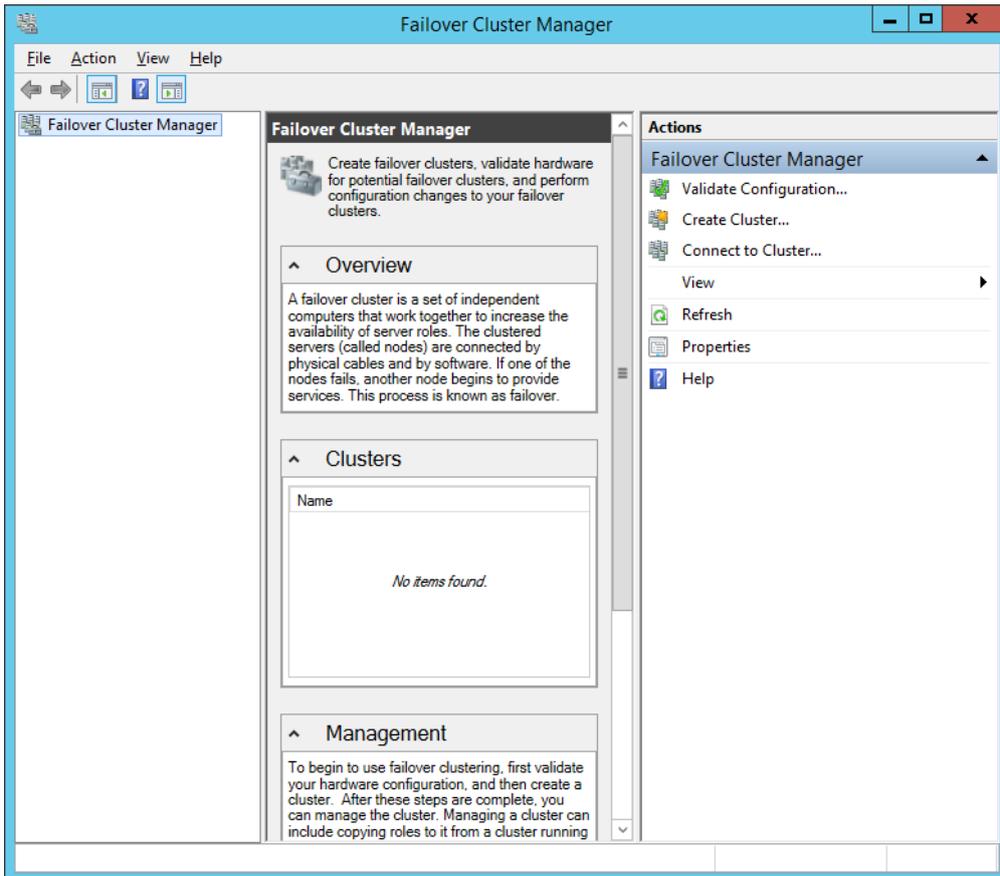
3. Specify the servers to be added to the cluster. Click Next to continue.
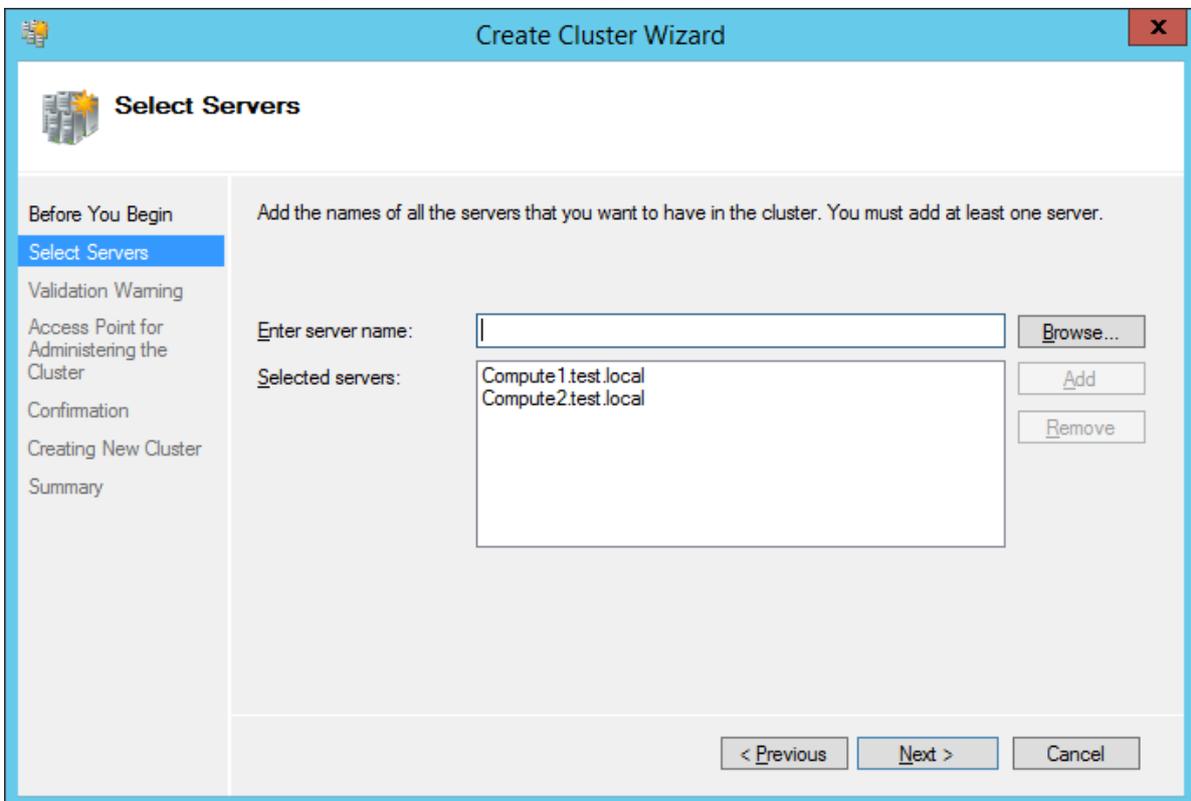
4. Validate the configuration by running the cluster validation tests: select Yes... and click Next to continue.
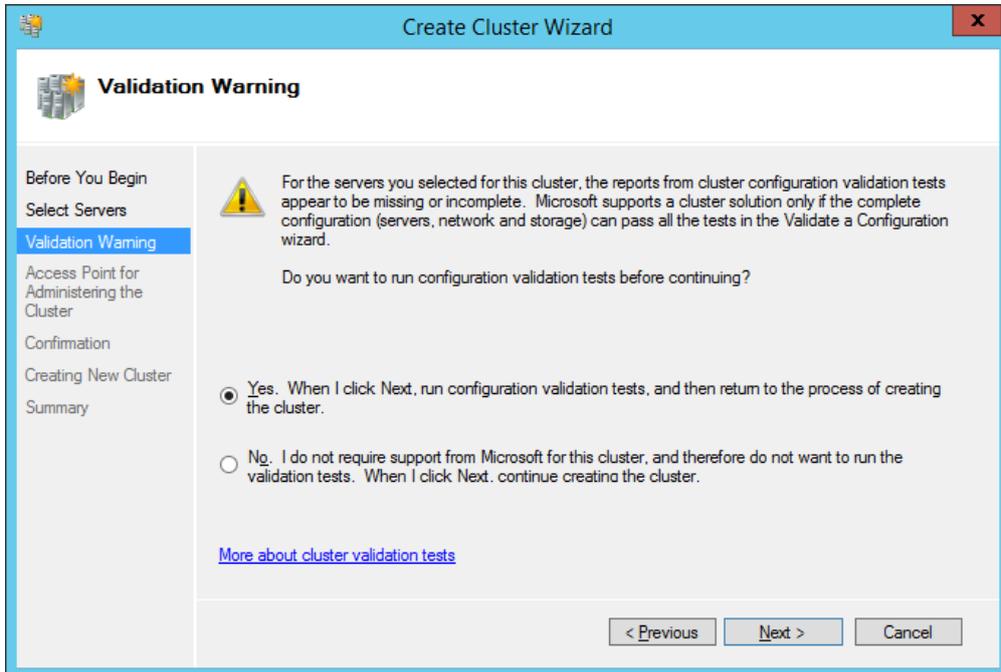
5. Specify Cluster Name.

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
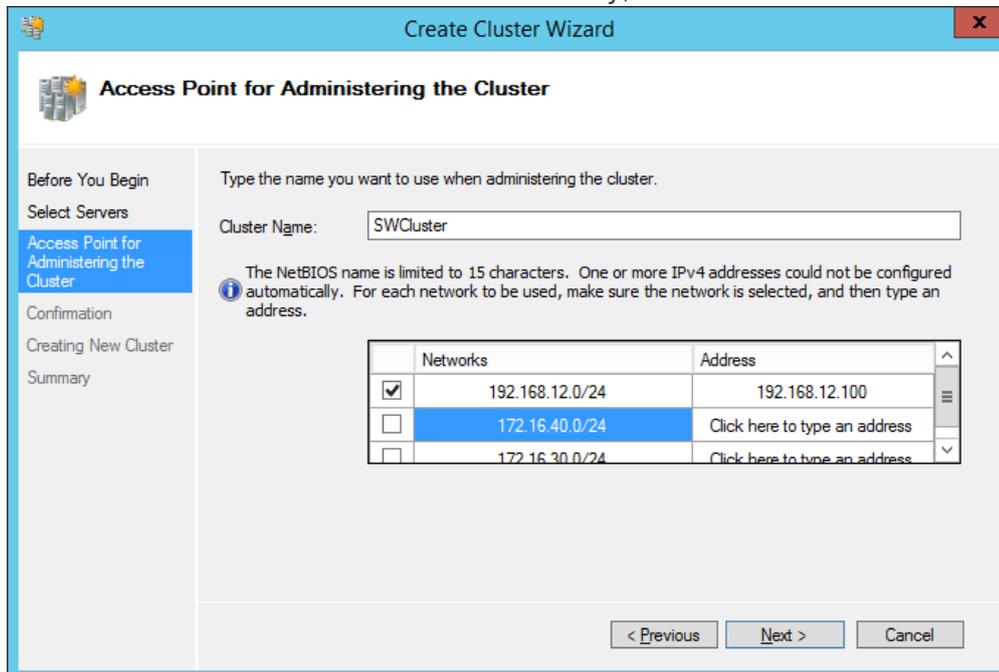
manually. 6. In Confirmation, make sure that all settings are correct. Click Previous to make any changes or Next to proceed.

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.
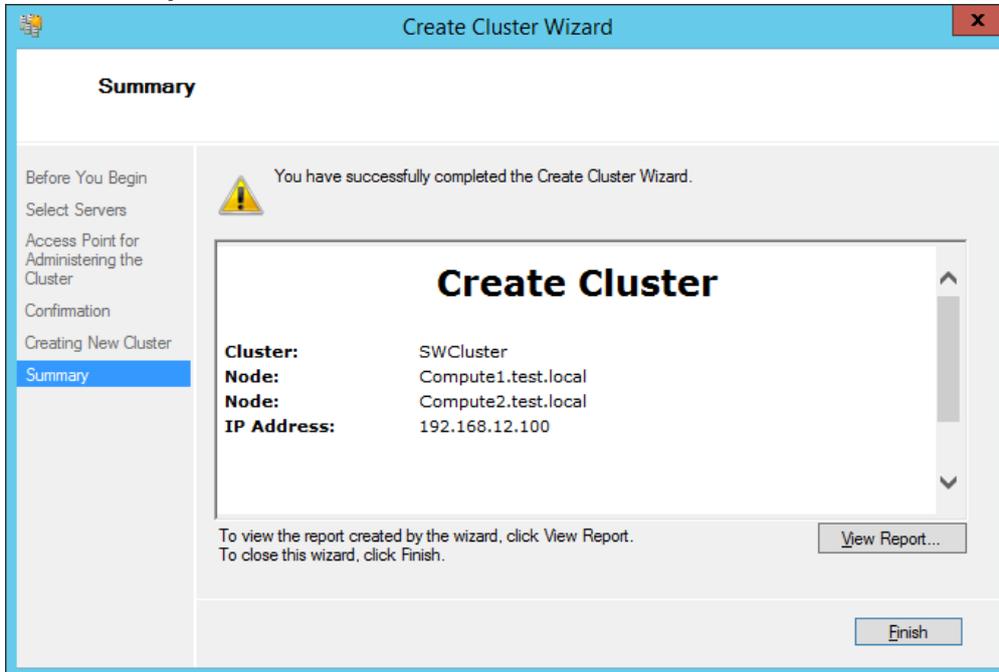
Adding Storage to the Cluster

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.
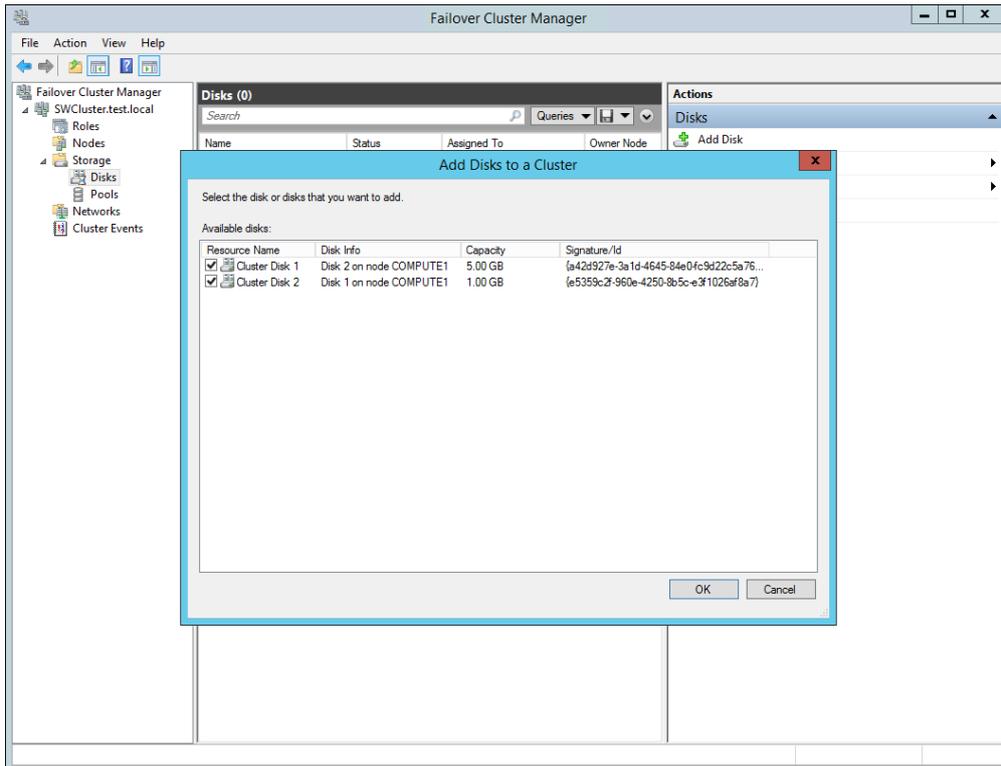
2. To configure the cluster witness disk, right-click on Cluster and proceed to More Actions -> Configure Cluster Quorum Settings.
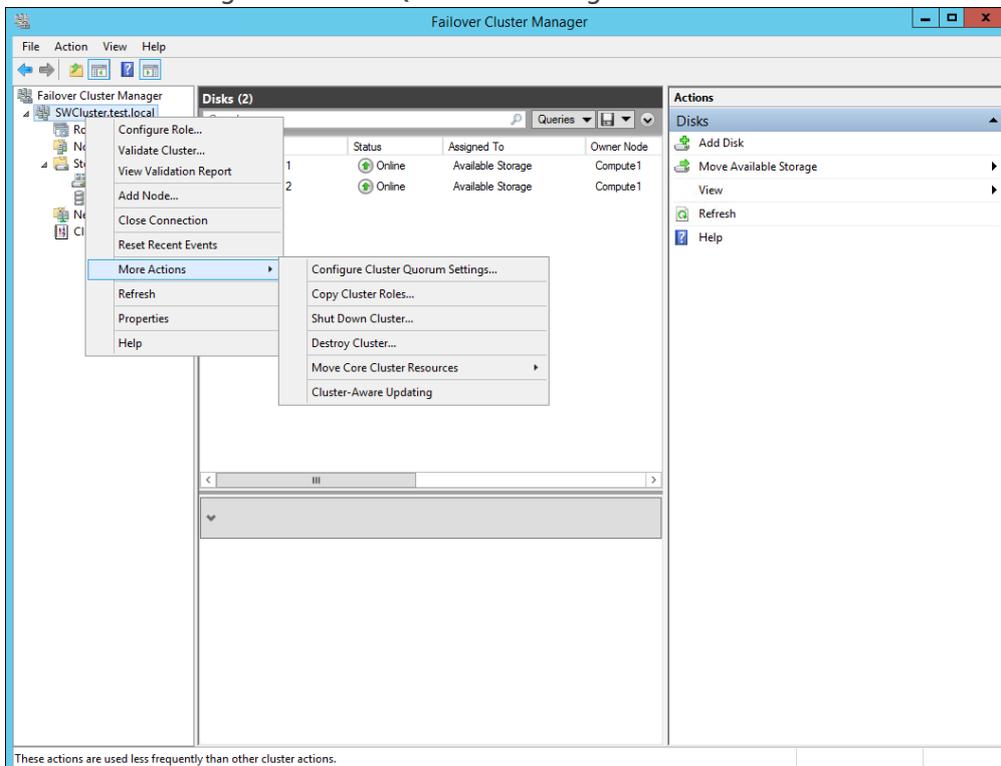
3. Follow the wizard and select the Select the quorum witness option. Click Next.
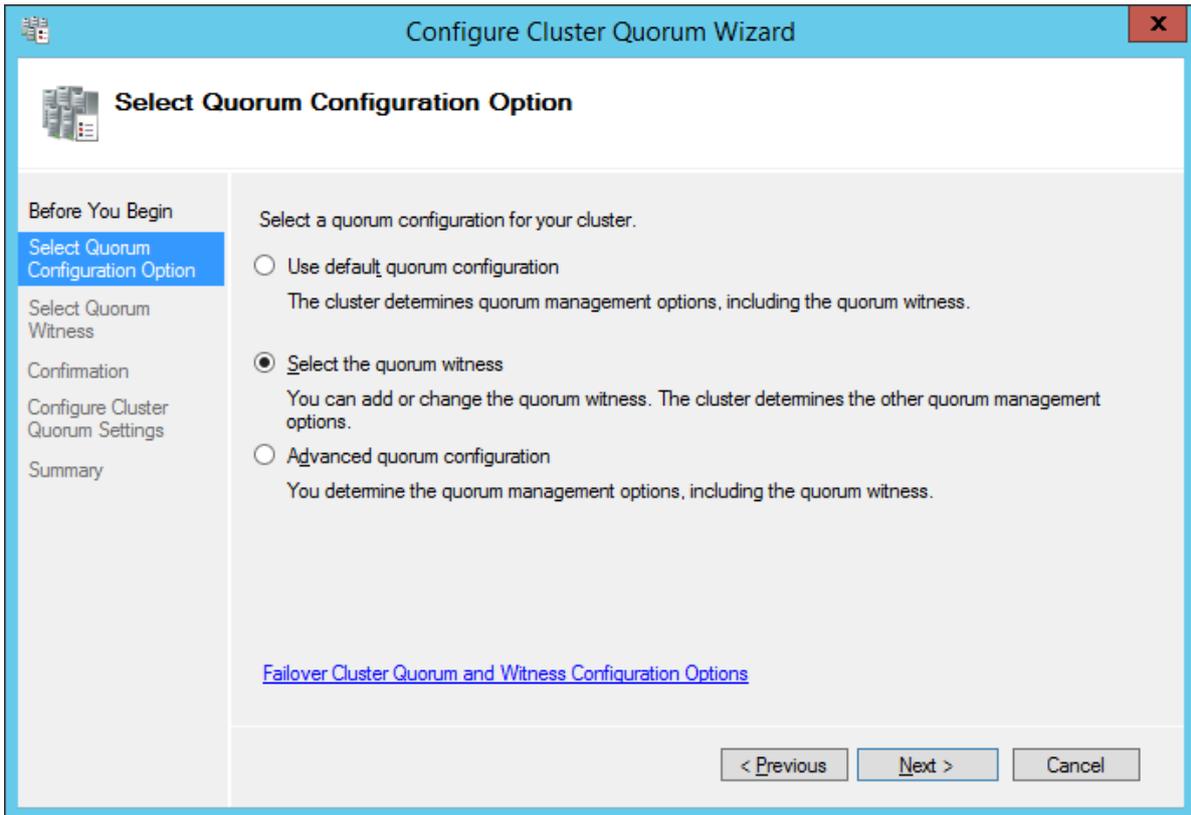
4. Select Configure a disk witness. Click Next.
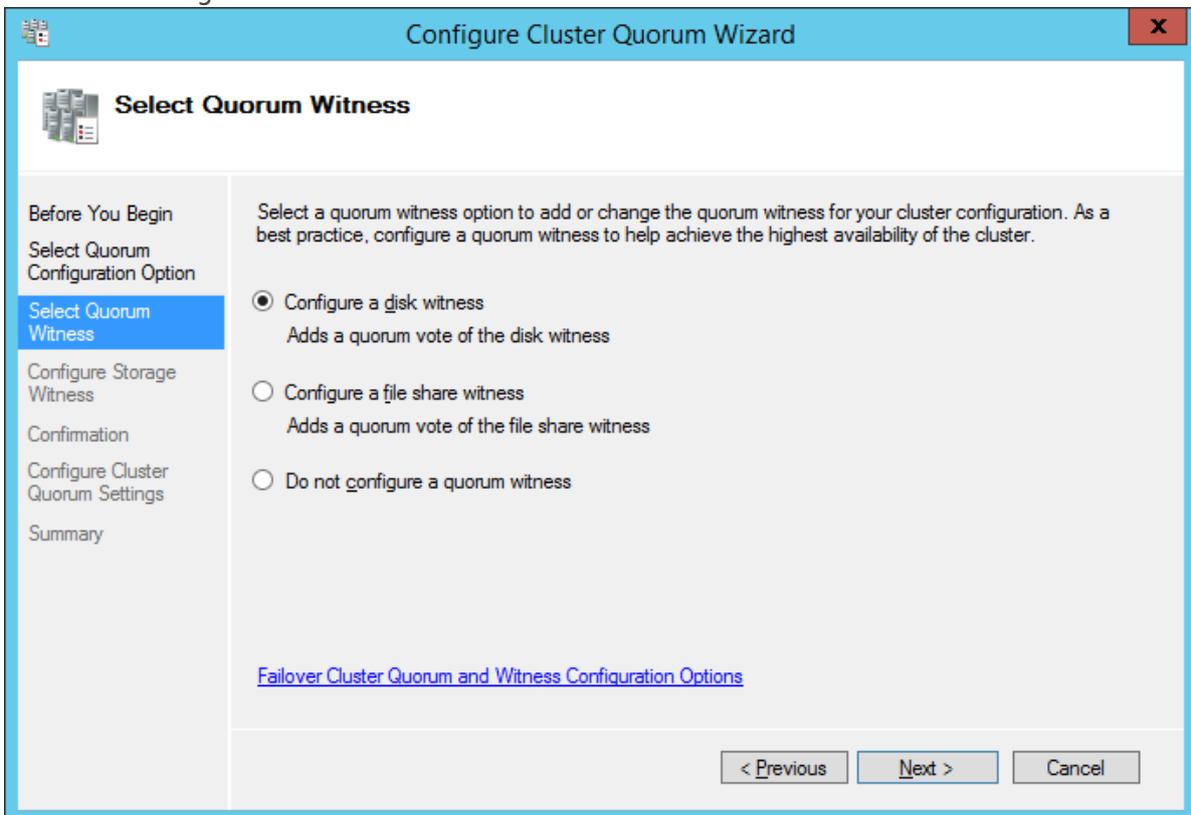
5. Select the Witness disk to be assigned as the cluster witness disk. Click Next and

press Finish to complete the operation.
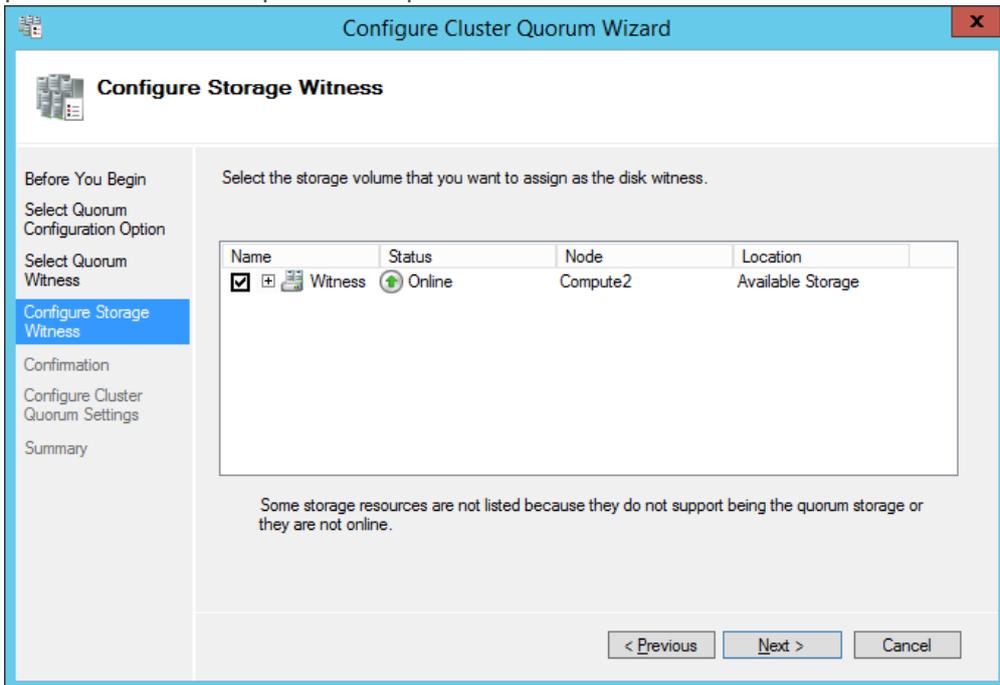
6. In Failover
Cluster Manager, right-click the disk and select Add to Cluster Shared Volumes.
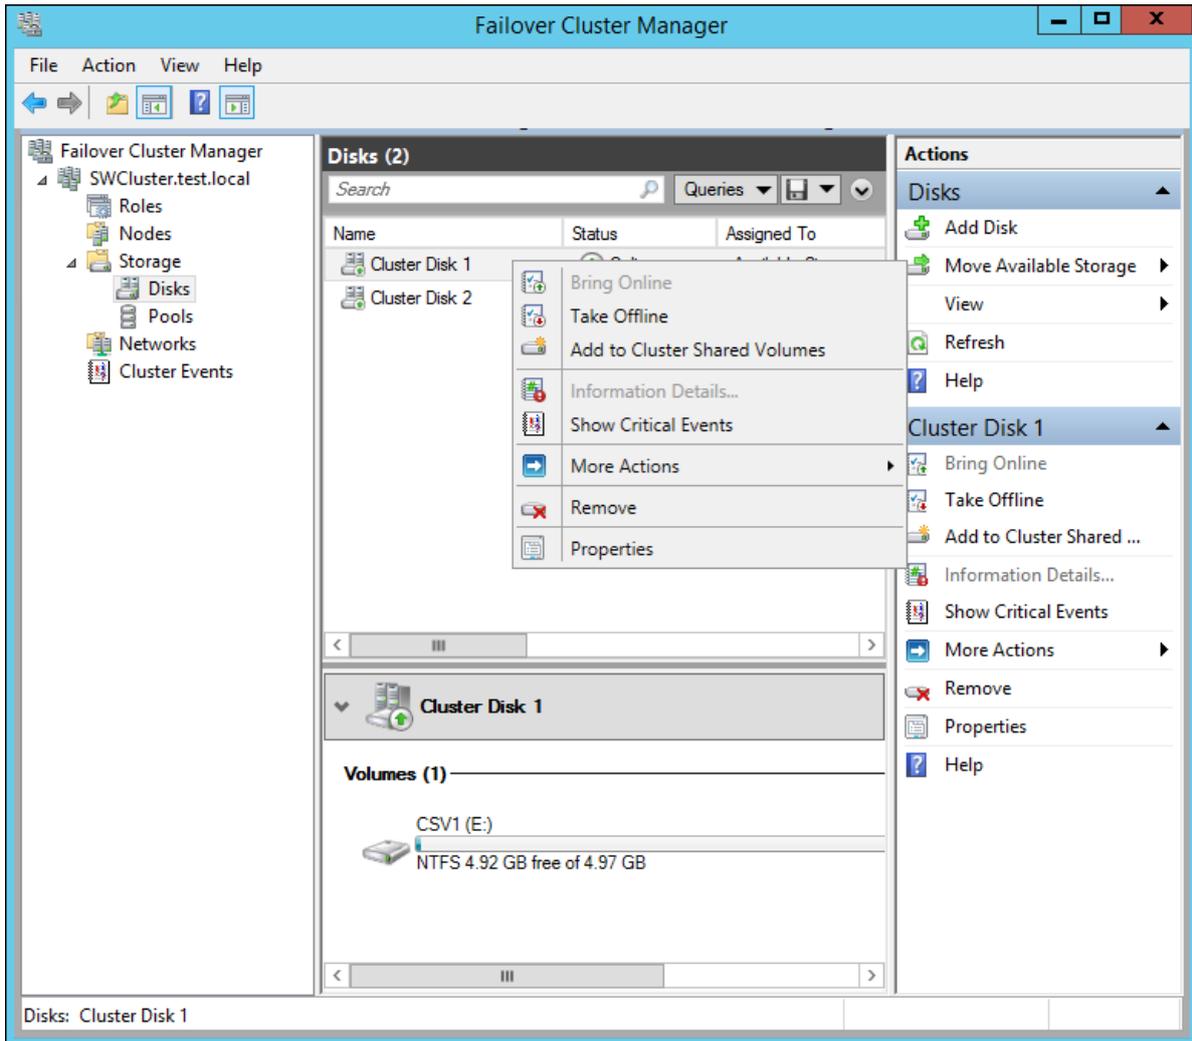
7. If renaming of the cluster shared volume is required, right-click on the disk and select Properties. Type the new name for the disk and click Apply followed by OK.
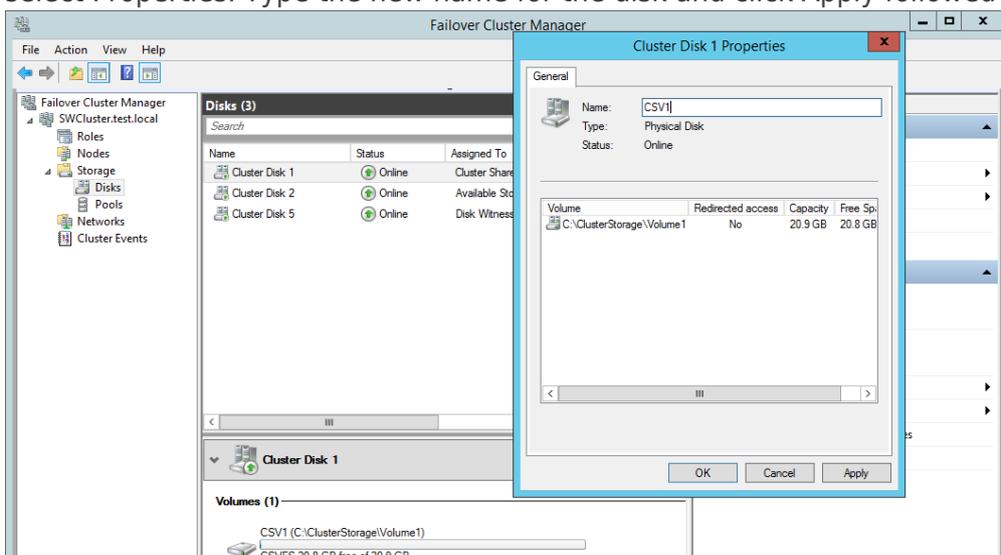
8. Perform the steps 6-7 for any other disk in Failover Cluster Manager. The resulting list of disks will

look similar to the screenshot below.
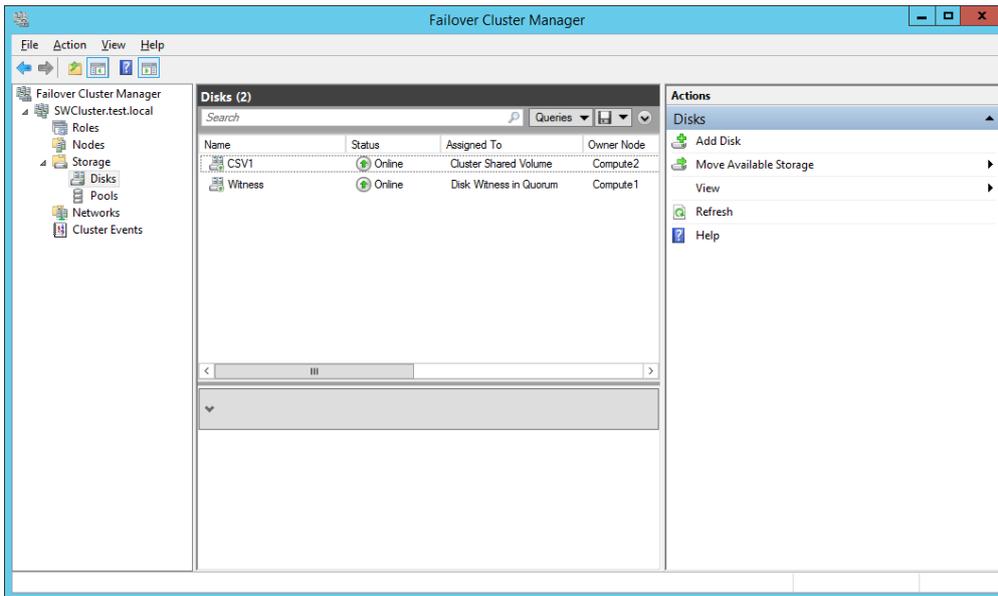
Configuring Cluster Network Preferences

1. In the Networks section of the Failover Cluster Manager, right-click on the network from the list. If required, set its new name 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
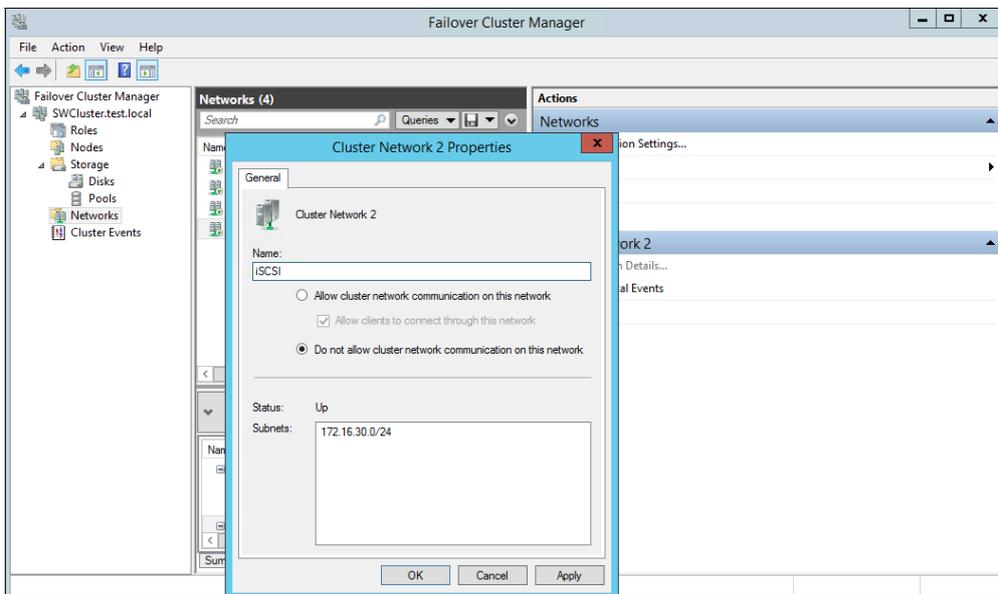
2. Rename other networks as described above, if required.
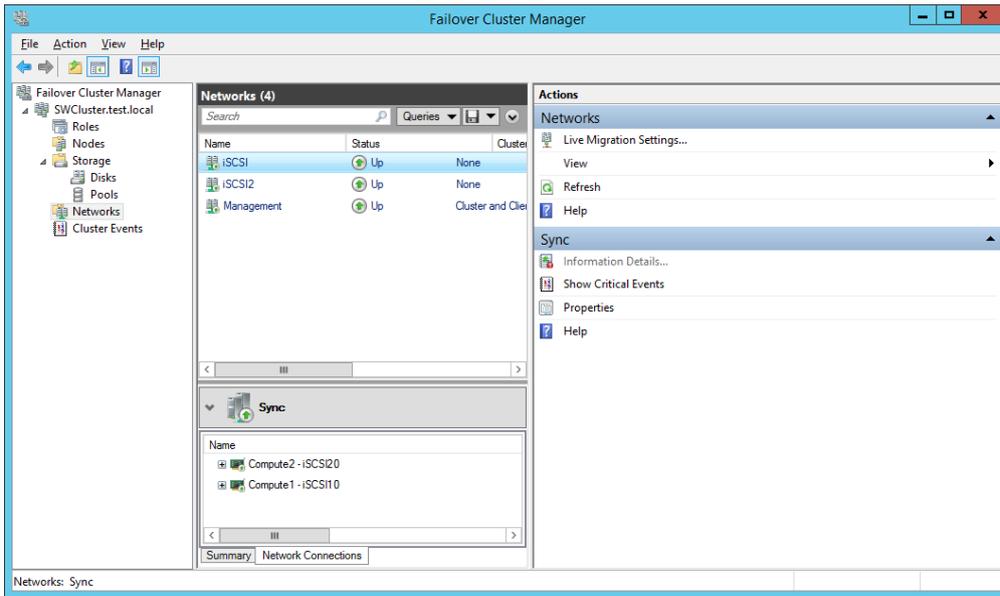
3. In the Actions tab, click Live Migration Settings. Uncheck the synchronization network, while the iSCSI network can be used if it is 10+ Gbps. Apply the changes and click OK.
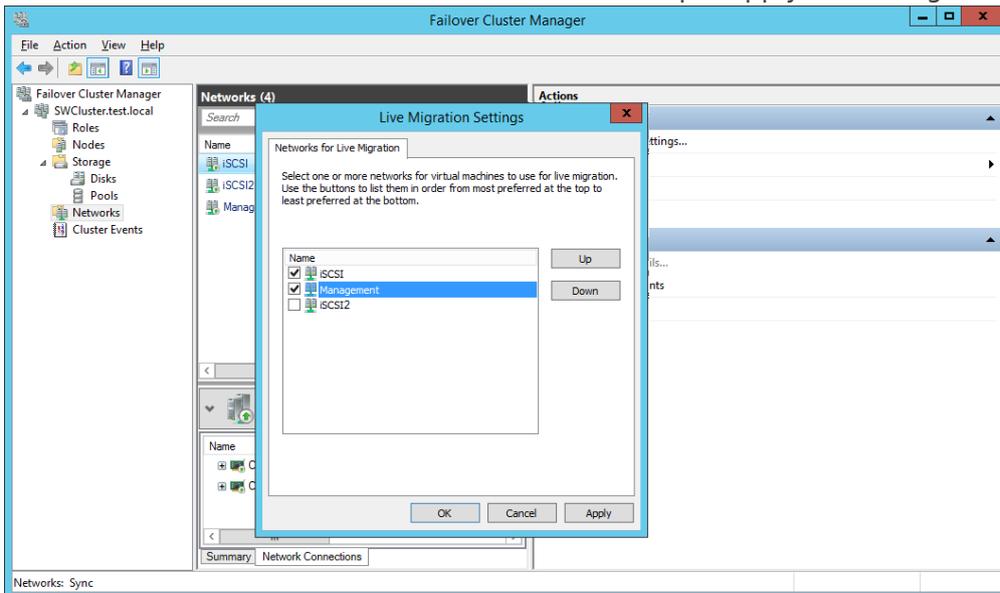
Contacts

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