StarWind Virtual SAN® 3-node Compute and Storage Separated Scenario with VMware vSphere 6.5/6.7

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.

IMPORTANT NOTE: ESXi6.5 and ESXi 6.7 are no longer supported by VMware (see more details here: https://lifecycle.vmware.com/). This guide should be used for informational, test, and educational purposes only.

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

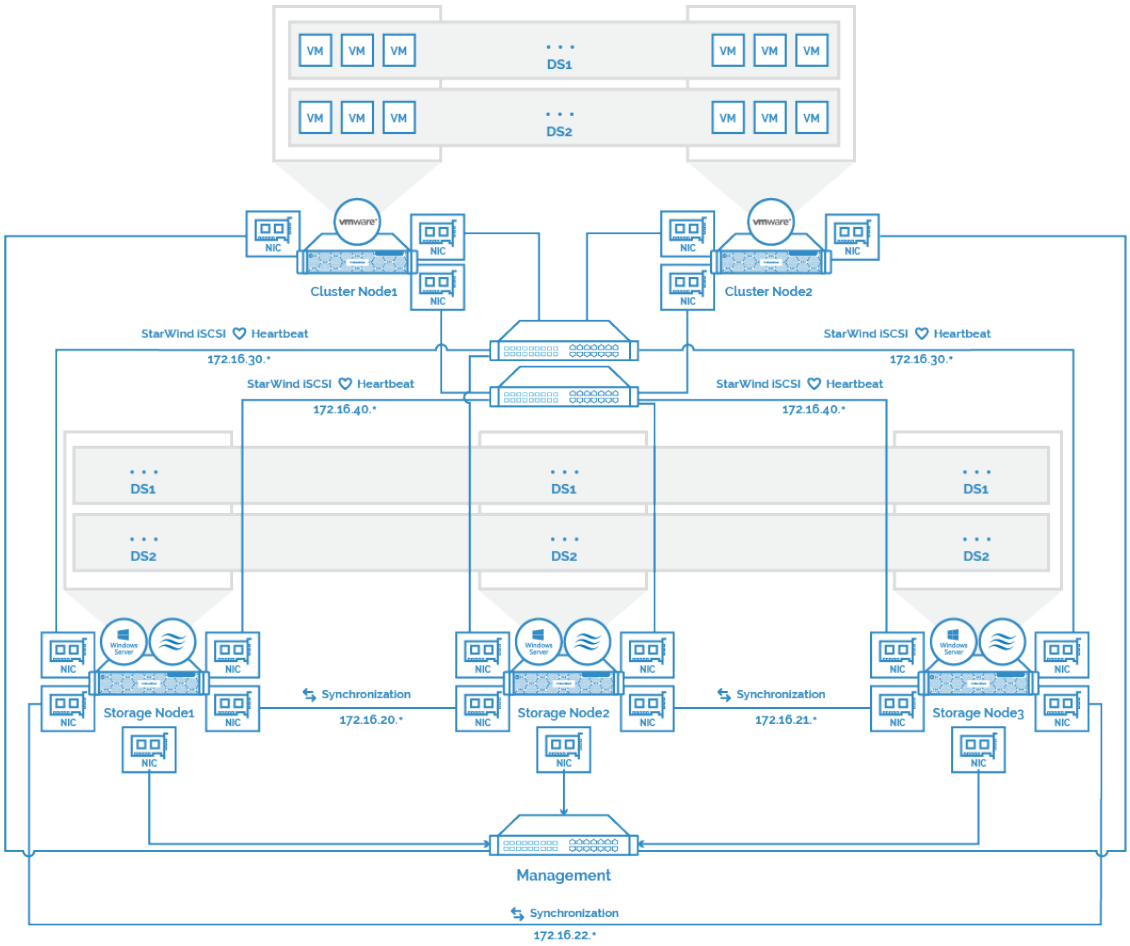
Prior to installing StarWind Virtual SAN, please make sure that the system meets the requirements, which are available via the following link:
https://www.starwindsoftware.com/system-requirements

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 Servers

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



1. ESXi hypervisor should be installed on compute hosts.

2. Windows Server Operating System should be installed on each storage node.

3. StarWind VSAN should be installed on the Windows Server operating system on each storage node.

4. On compute nodes and storage nodes, network interfaces to be used for Synchronization and iSCSI/StarWind heartbeat should be in different subnets and connected directly according to the network diagram above. Here, the 172.16.30.x, 172.16.40.x subnets are used for the iSCSI/StarWind heartbeat traffic on compute nodes, while the 172.16.20.x, 172.16.21.x, 172.16.22.x subnets are used for the Synchronization traffic on storage nodes.

NOTE: 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.

Installing Starwind Vsan For Hyper-V

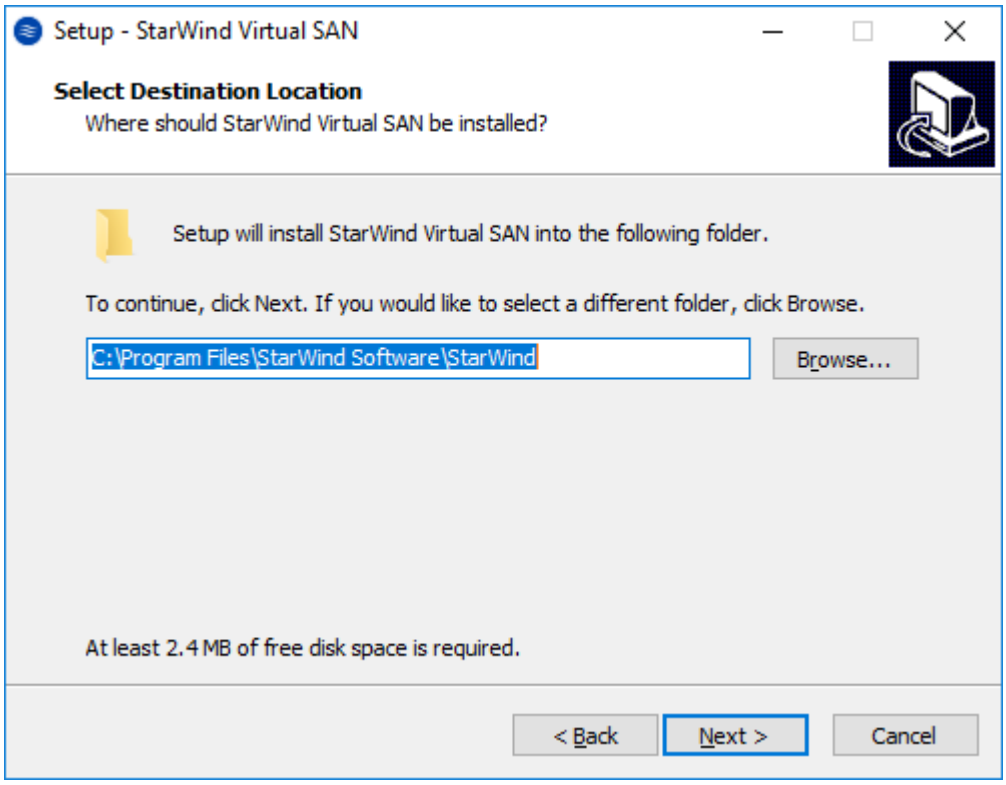
1. Download the StarWind setup executable file from the StarWind website: https://www.starwind.com/registration-starwind-virtual-san

2. Launch the downloaded setup file on the server to install StarWind Virtual SAN or one of its components. The Setup wizard will appear. Read and accept the License Agreement.



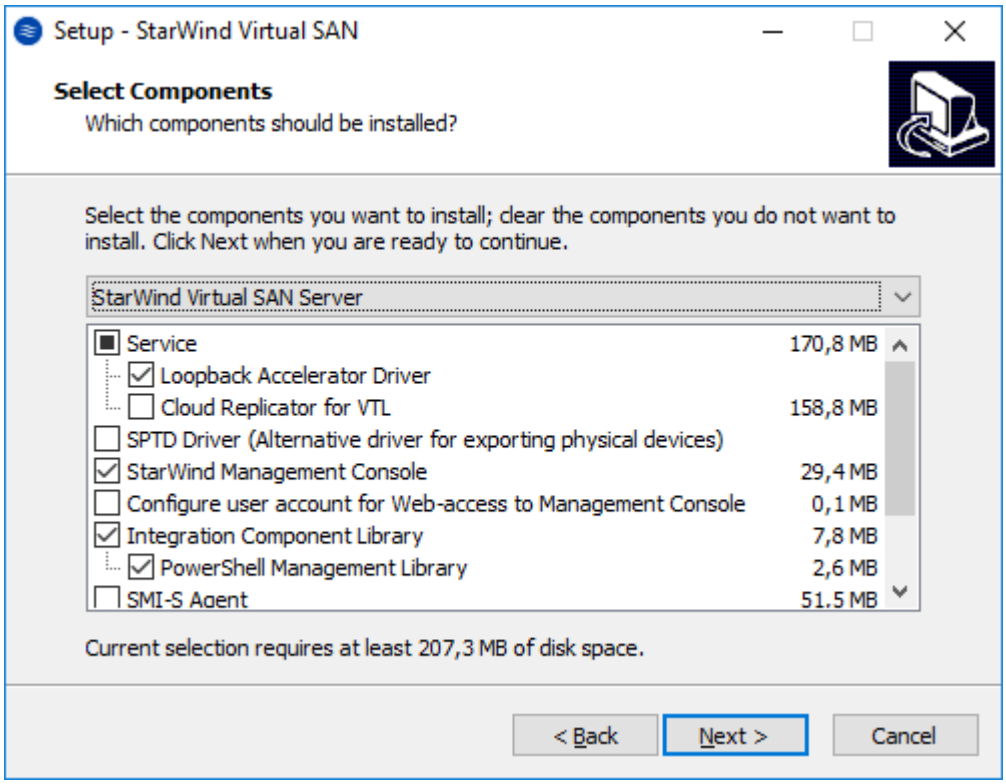
3. Carefully read the information about the new features and improvements. Red text indicates warnings for users that are updating the existing software installations.

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

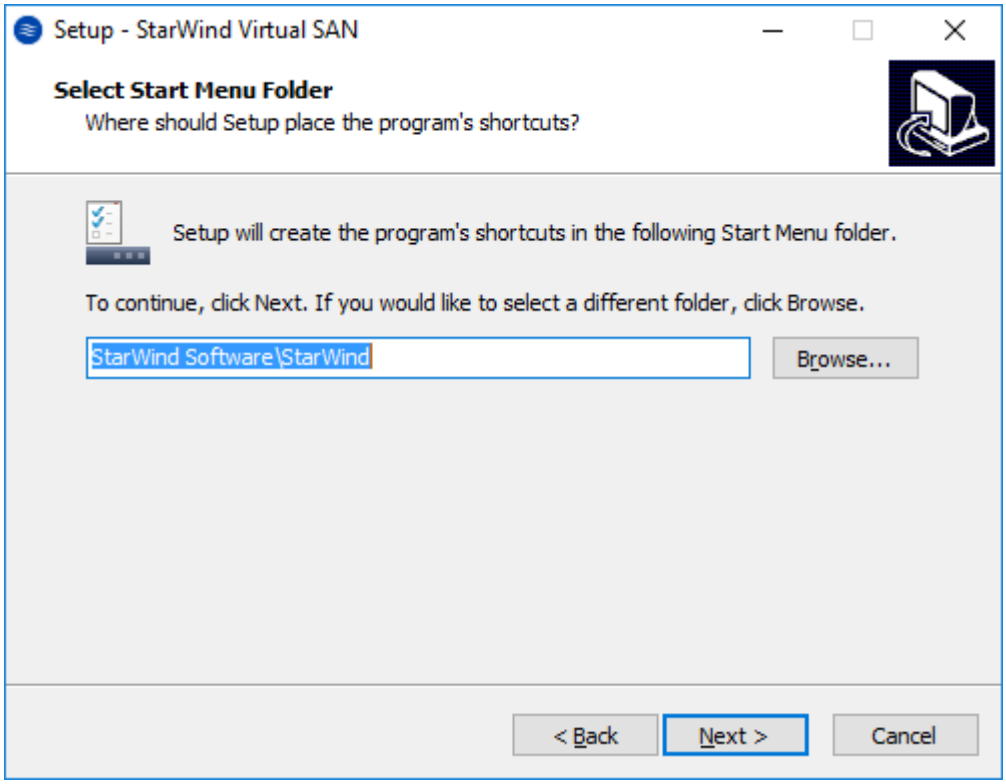


5. Select the following components for the minimum setup:

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



6. Specify Start Menu Folder.



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

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

- request time-limited fully functional evaluation key.
- request FREE version key.
- relect the previously purchased commercial license key.

9. Click on the Browse button to locate the license file.

10. Review the licensing information.

11. Verify the installation settings. Click on Back to make any changes or Install to proceed with installation.

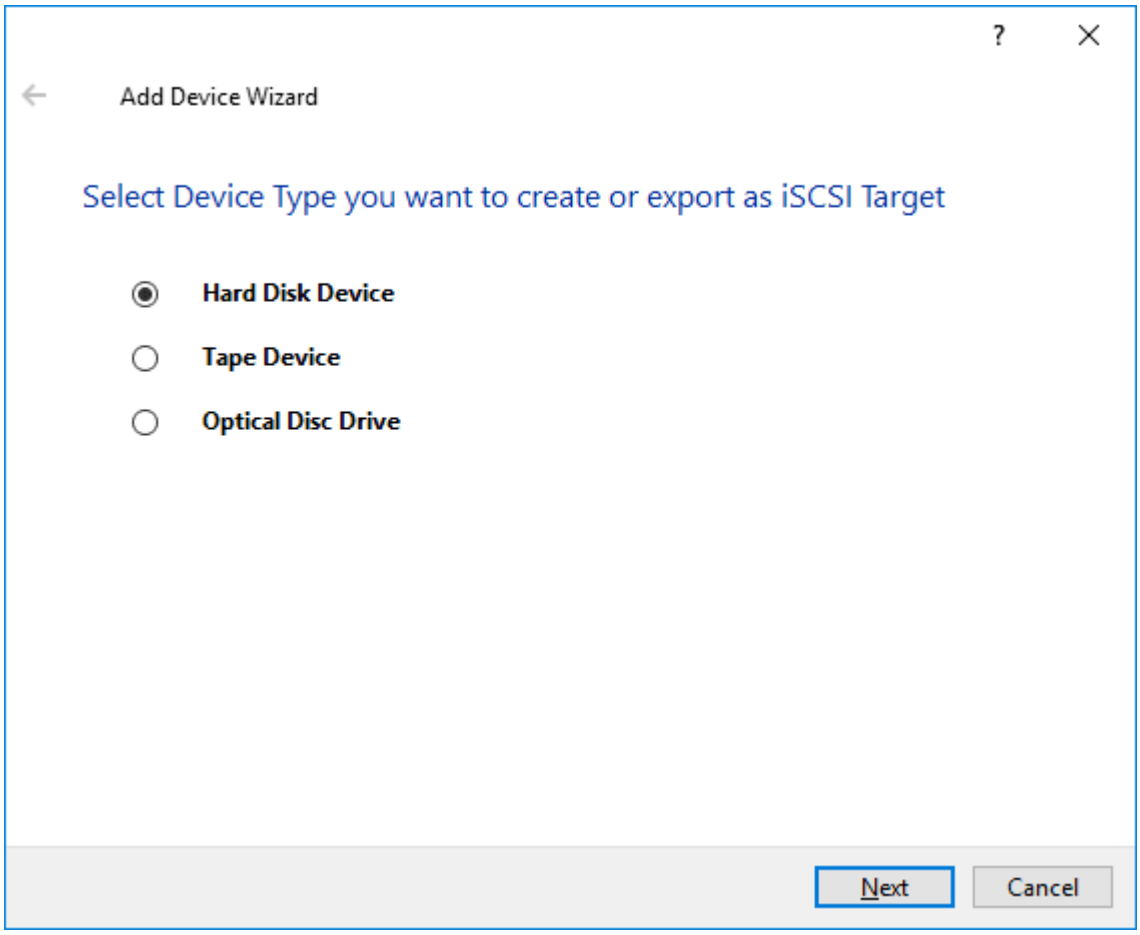
12. Enable the appropriate checkbox to launch StarWind Management Console right after the setup wizard is closed and click on Finish.

13. Repeat the installation steps on the partner node.

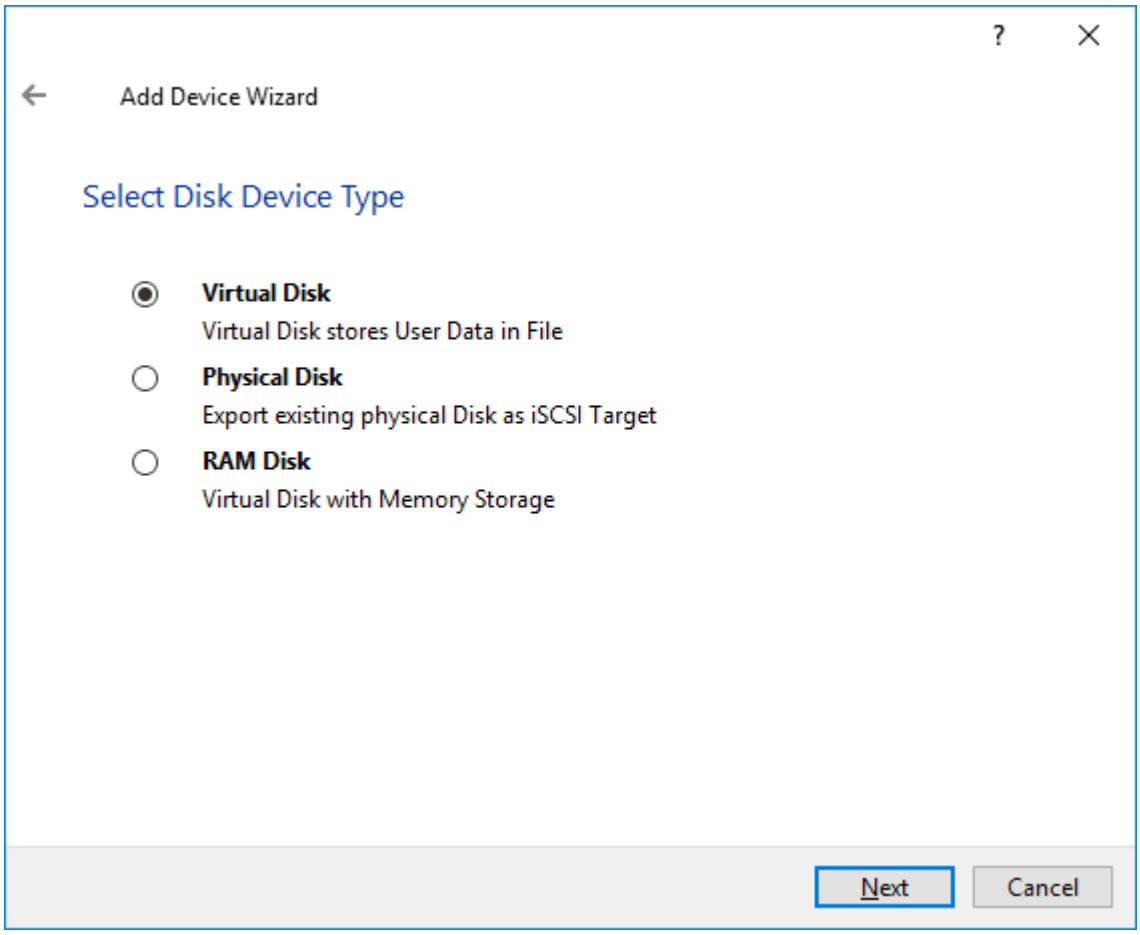
Creating Starwind Devices

1. In the StarWind Management Console click to Add Device (advanced) button and open Add Device (advanced) Wizard.

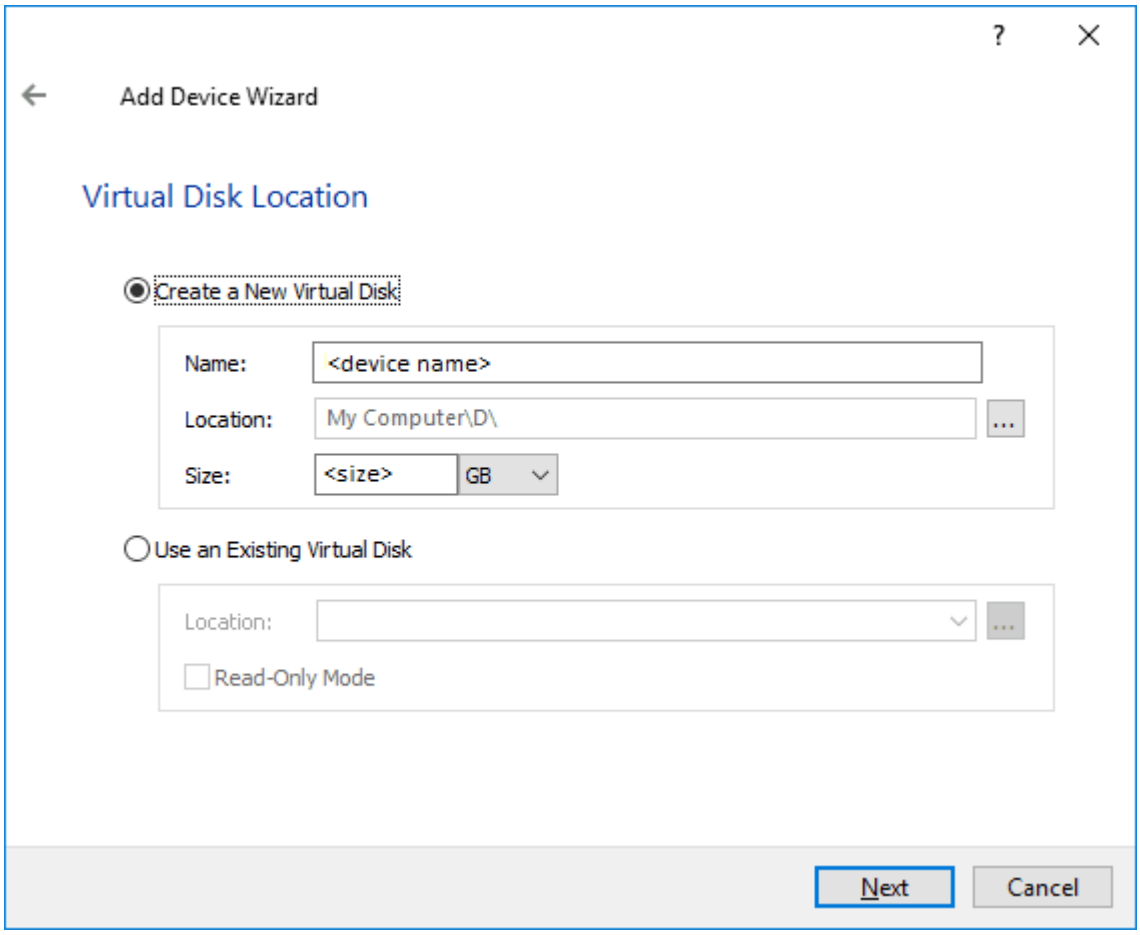
2. Select Hard Disk Device as the type of device to be created.



3. Select Virtual Disk.

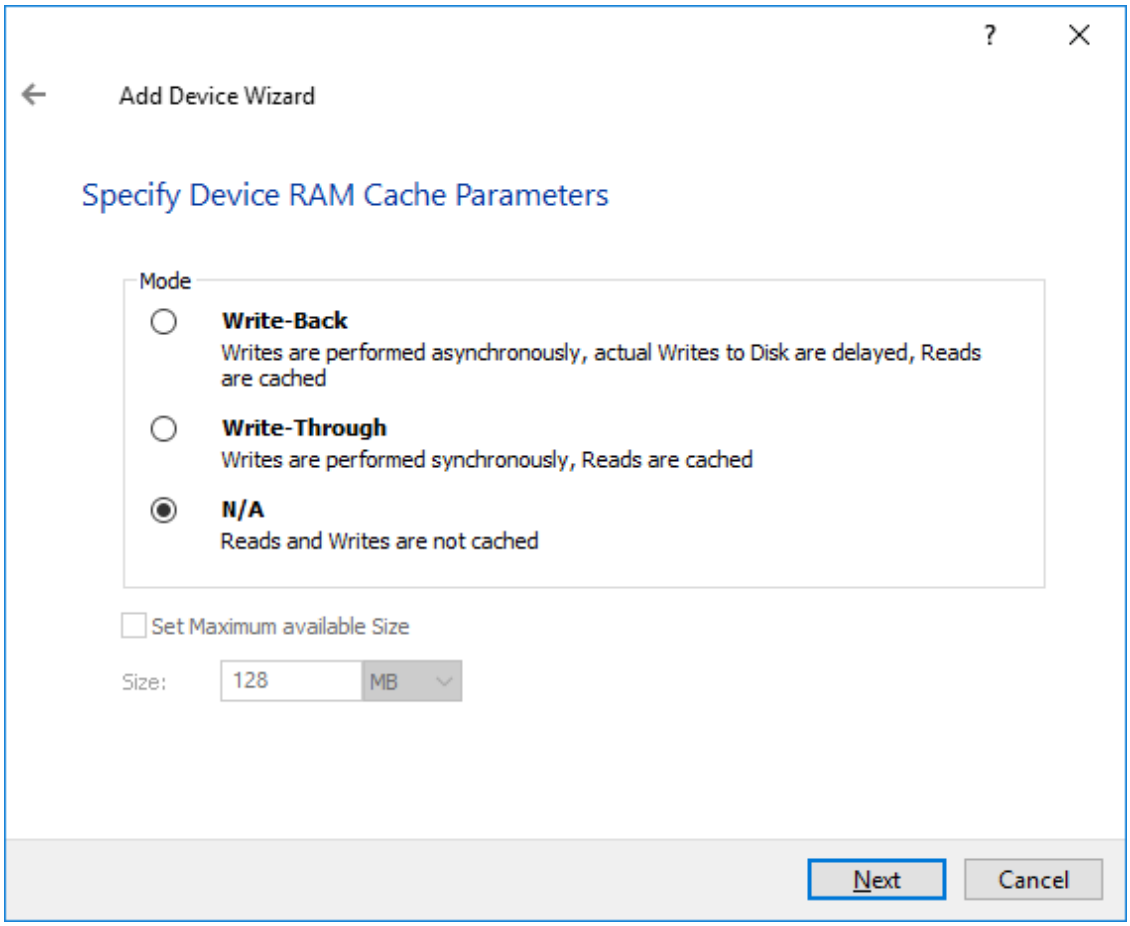


4. Specify a virtual disk Name, Location, and Size.

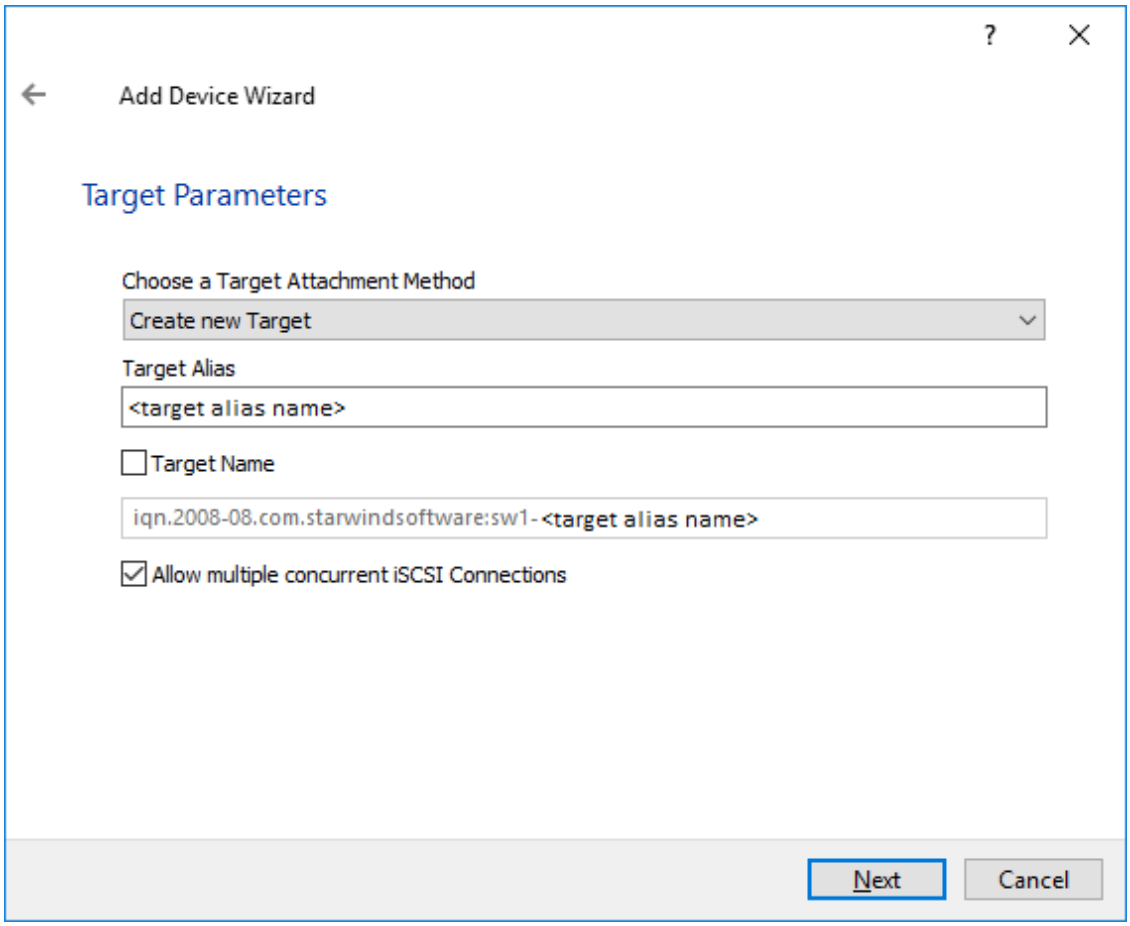


5. Select the Thick provisioned disk type and block size.
NOTE: Use 4096 sector size for targets, connected on Windows-based systems and 512 bytes sector size for targets, connected on Linux-based systems (ESXi/Xen/KVM).

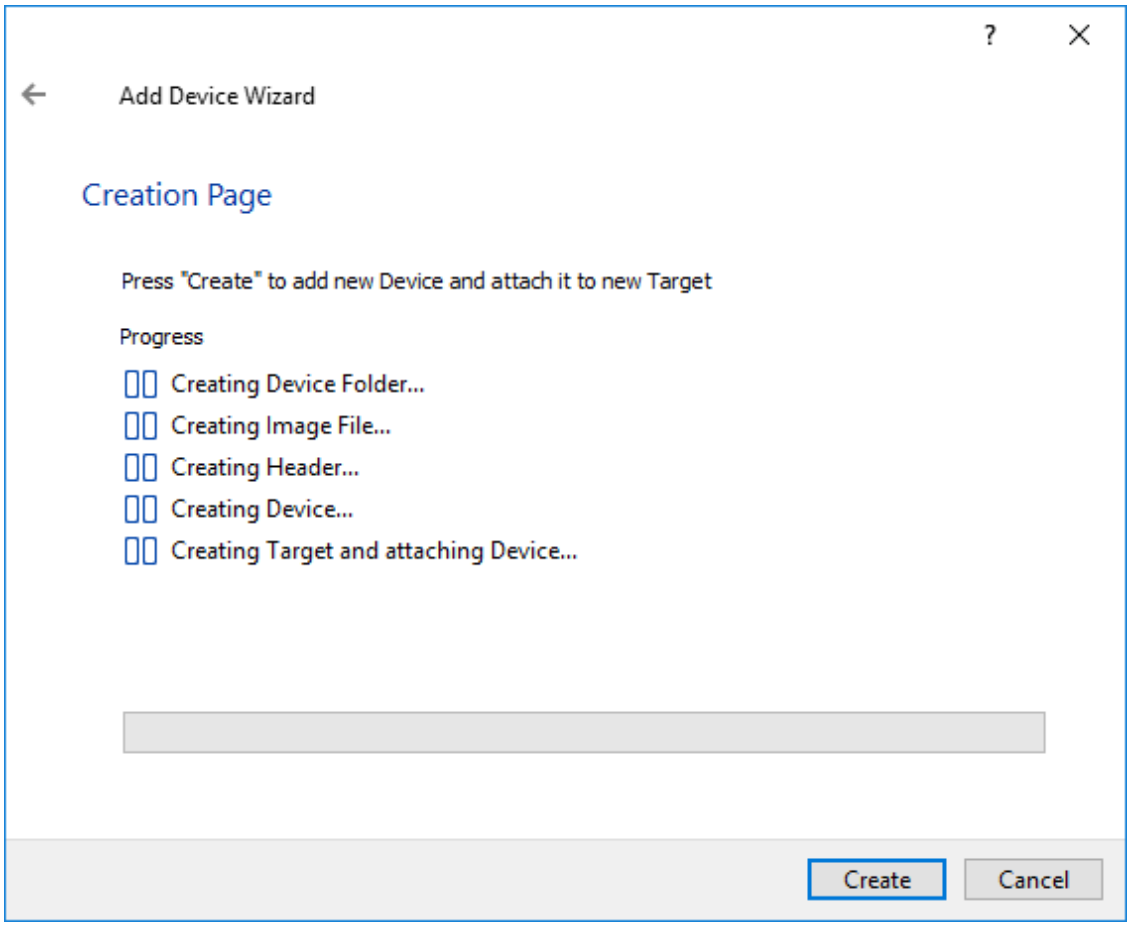
6. Define a caching policy and specify a cache size (in MB). Also, the maximum available cache size can be specified by selecting the appropriate checkbox. Optionally, define the L2 caching policy and cache size.



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



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



9. Click Close to finish the device creation.

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

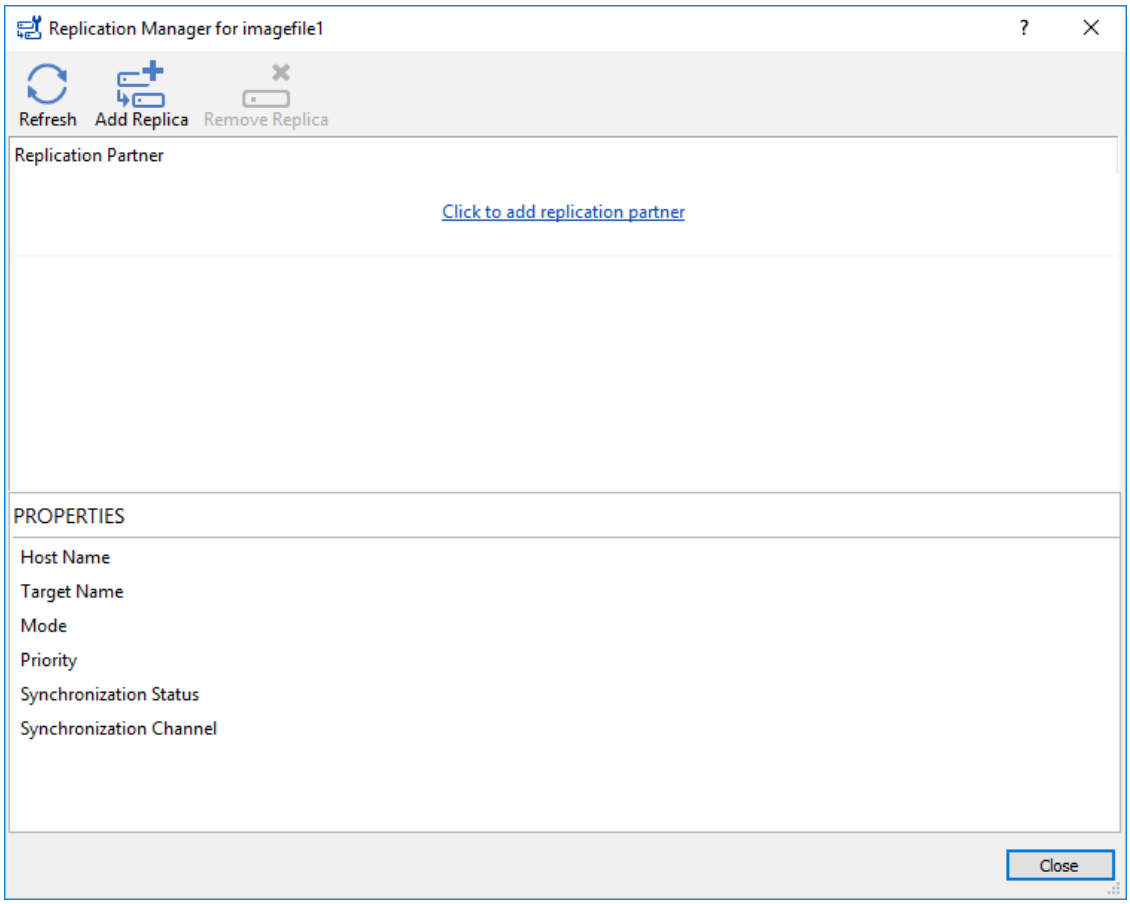
Select The Required Replication Mode

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

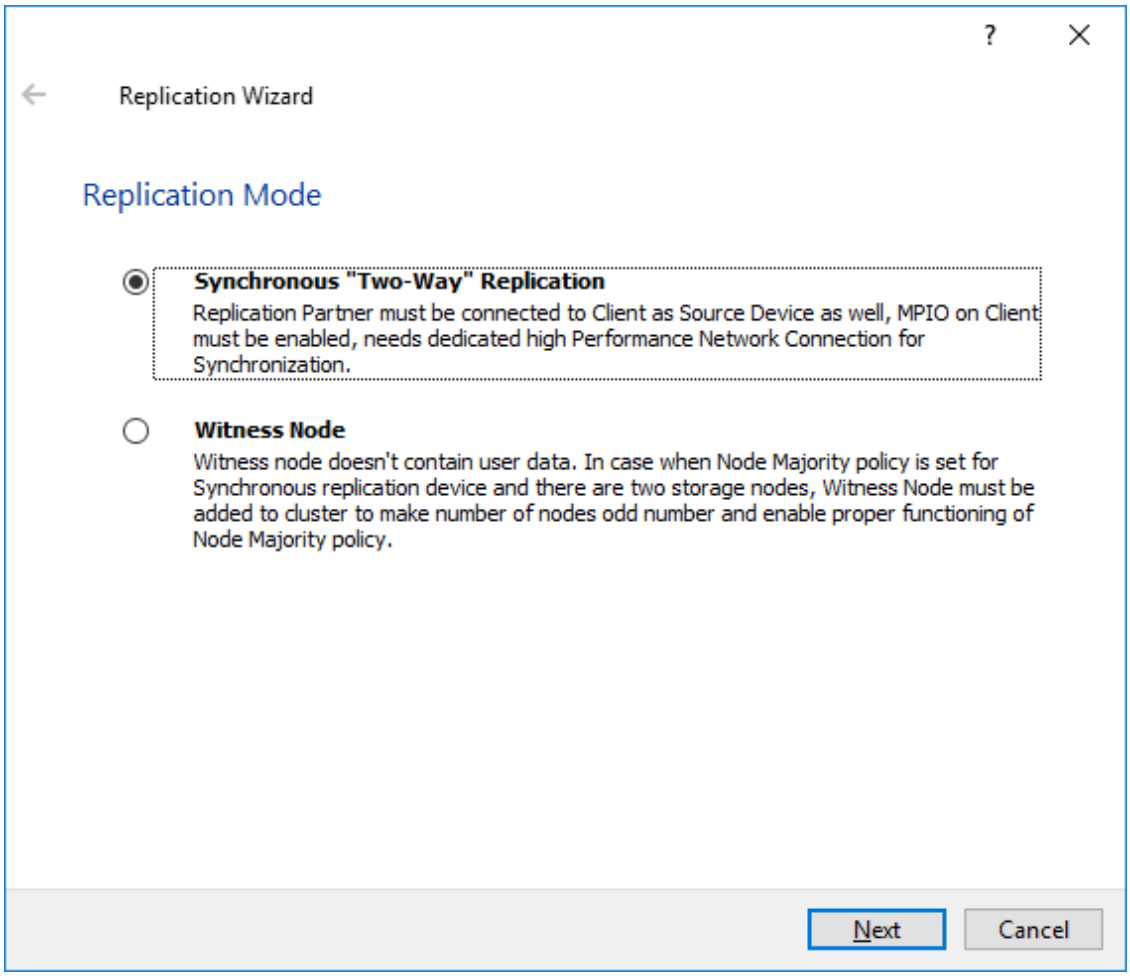
Synchronous “Two-Way” Replication

1. Right-click the recently created device and select Replication Manager from the shortcut menu.

2. Select the Add Replica button in the top menu.



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



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

Selecting The Failover Strategy

StarWind provides 2 options for configuring a failover strategy:

Heartbeat

The Heartbeat failover strategy allows avoiding the “split-brain” scenario when the HA cluster nodes are unable to synchronize but continue to accept write commands from the initiators independently. It can occur when all synchronization and heartbeat channels disconnect simultaneously, and the partner nodes do not respond to the node’s requests. As a result, StarWind service assumes the partner nodes to be offline and continues operations on a single-node mode using data written to it.
If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The device with the lowest priority will be marked as not synchronized and get subsequently blocked for the further read and write operations until the synchronization channel resumption. At the same time, the partner device on the

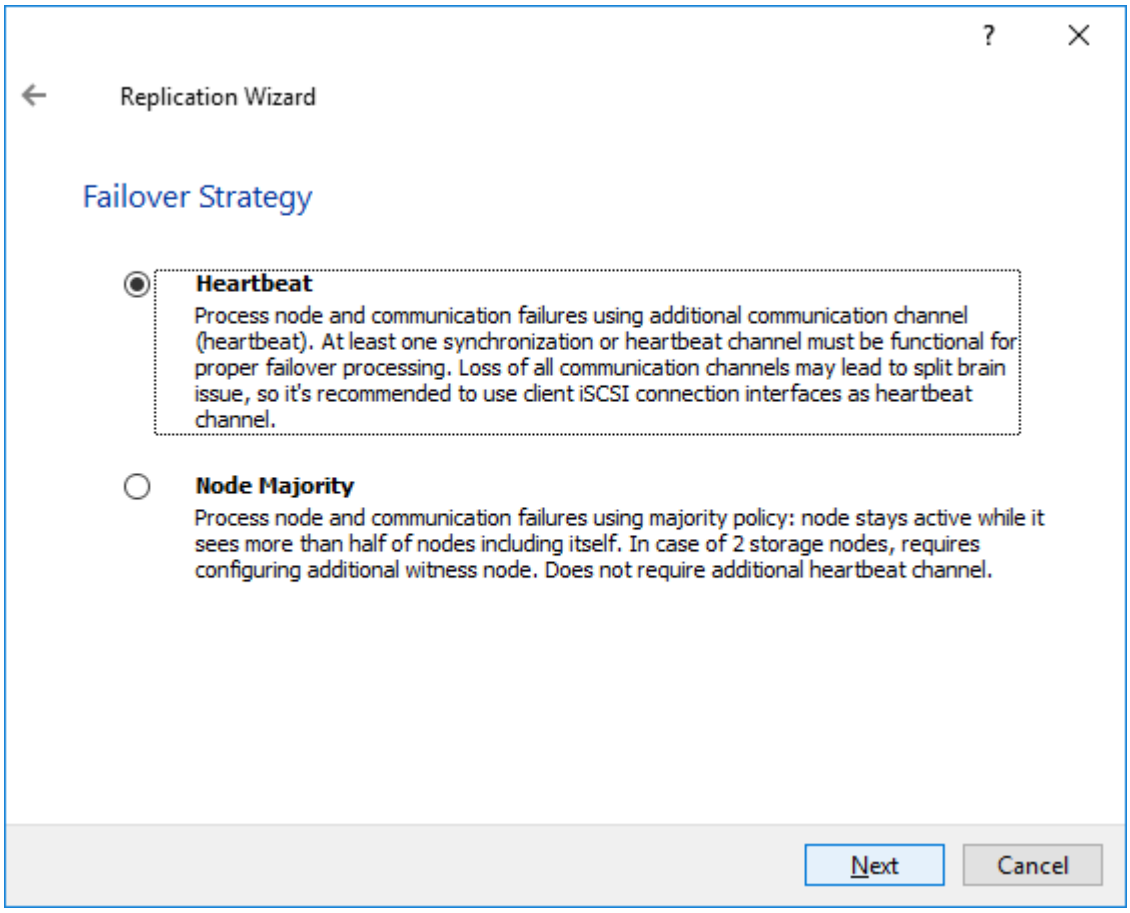
synchronized node flushes data from the cache to the disk to preserve data integrity in case the node goes down unexpectedly. It is recommended to assign more independent heartbeat channels during the replica creation to improve system stability and avoid the “split-brain” issue.
With the heartbeat failover strategy, the storage cluster will continue working with only one StarWind node available.

Node Majority

The Node Majority failover strategy ensures the synchronization connection without any additional heartbeat links. The failure-handling process occurs when the node has detected the absence of the connection with the partner.
The main requirement for keeping the node operational is an active connection with more than half of the HA device’s nodes. Calculation of the available partners is based on their “votes”.
In case of a two-node HA storage, all nodes will be disconnected if there is a problem on the node itself, or in communication between them. Therefore, the Node Majority failover strategy requires the addition of the third Witness node or file share (SMB) which participates in the nodes count for the majority, but neither contains data on it nor is involved in processing clients’ requests. In case an HA device is replicated between 3 nodes, no Witness node is required.
With Node Majority failover strategy, failure of only one node can be tolerated. If two nodes fail, the third node will also become unavailable to clients’ requests.
Please select the required option:

Heartbeat

1. Select Failover Strategy.



2. Select Create new Partner Device and click Next.

3. Select a partner device Location and click Next.

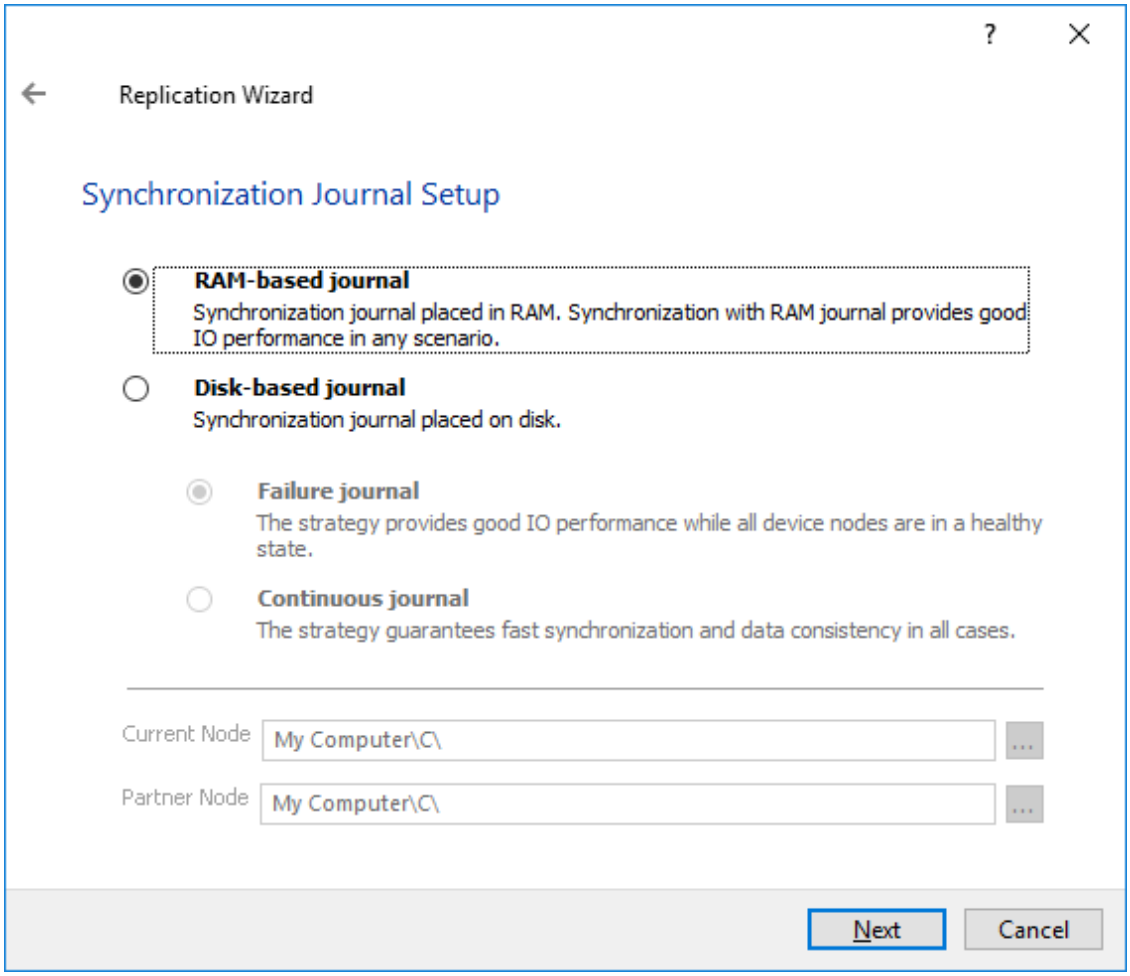
4. Select Synchronization Journal Strategy and click Next.
NOTE: There are several options – RAM-based journal (default) and Disk-based journal with failure and continuous strategy, that allow to avoid full synchronization cases.

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

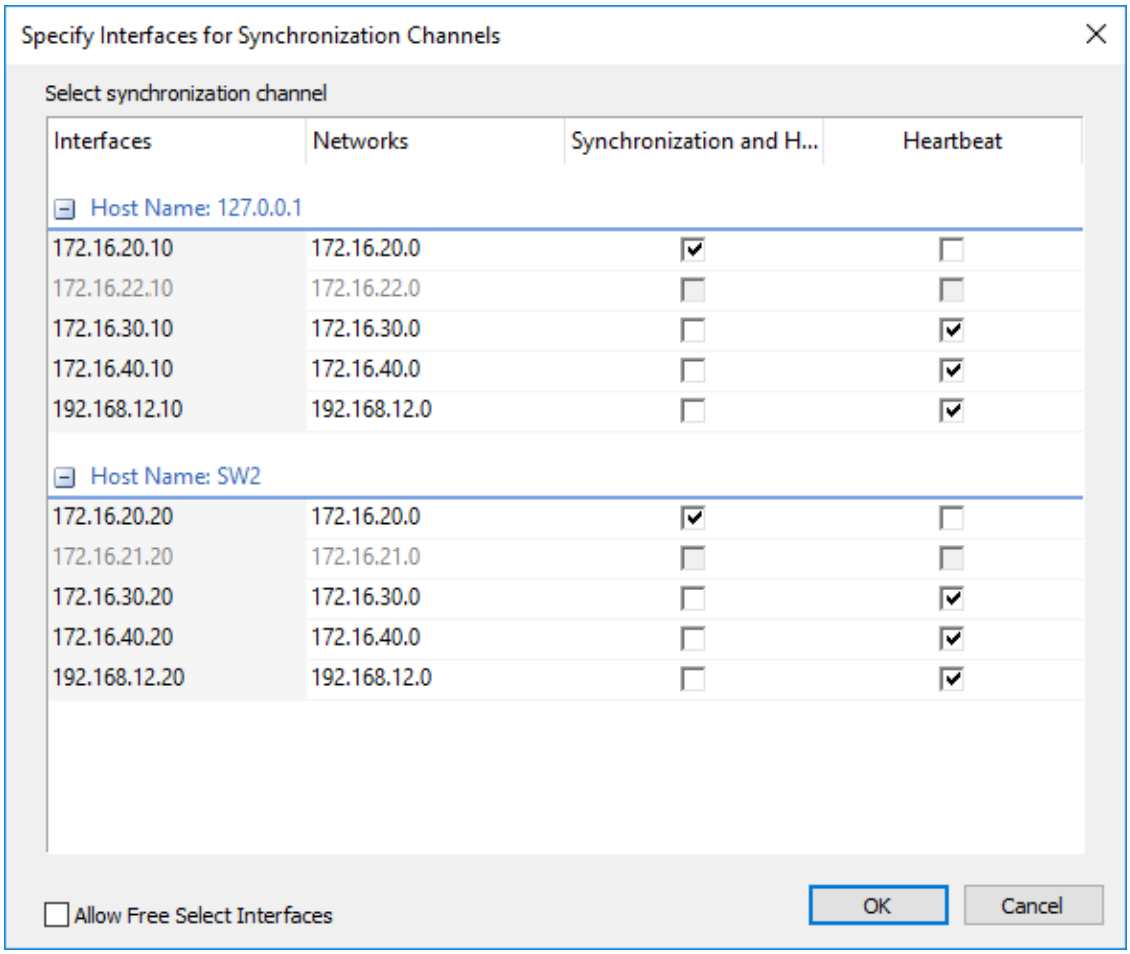
Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes.
Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind 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 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. Specify the interfaces for Synchronization and Heartbeat Channels. Click OK and then click Next.

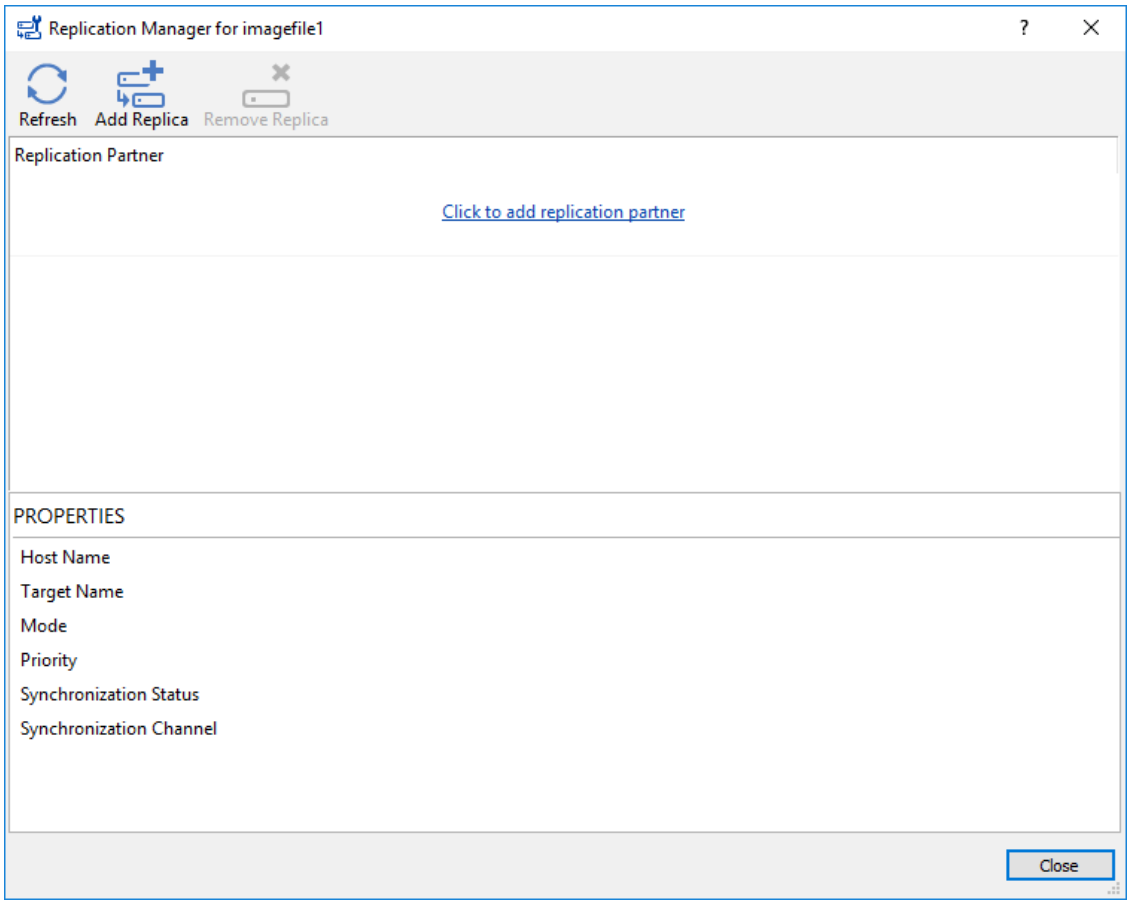


6. In Select Partner Device Initialization Mode, select Synchronize from existing Device and click Next.

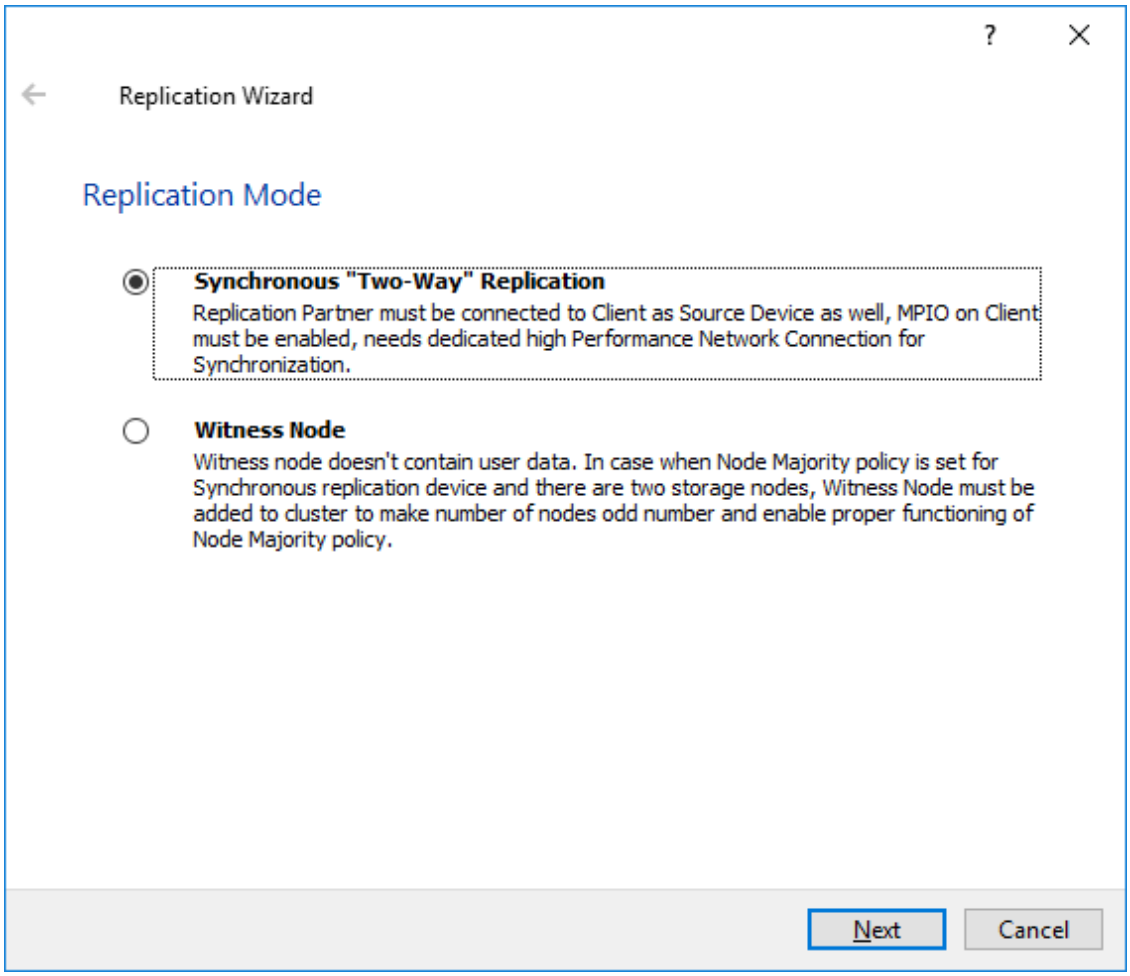
7. Click Create Replica. Click Finish to close the wizard.

8. The successfully added device appears in StarWind Management Console.

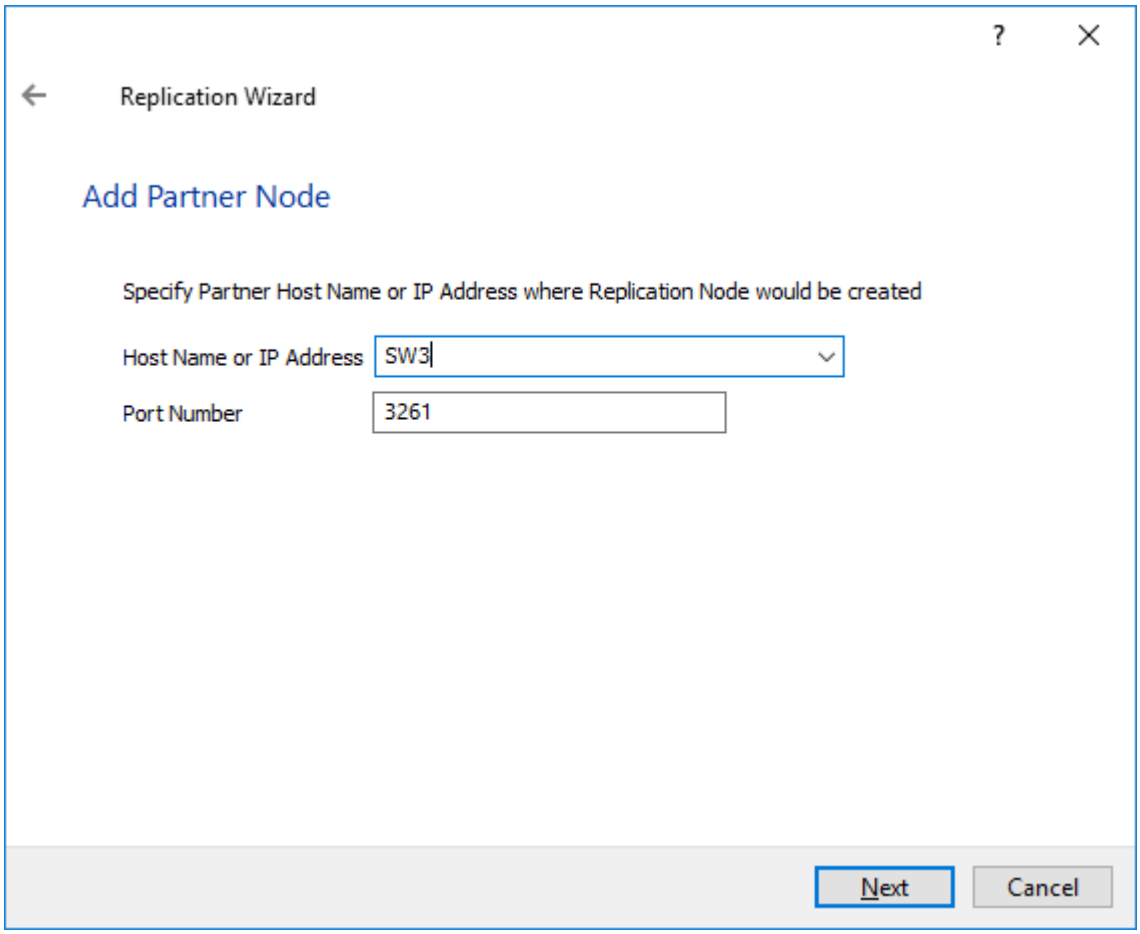
9. Choose device, open Replication Manager and click Add replica again.



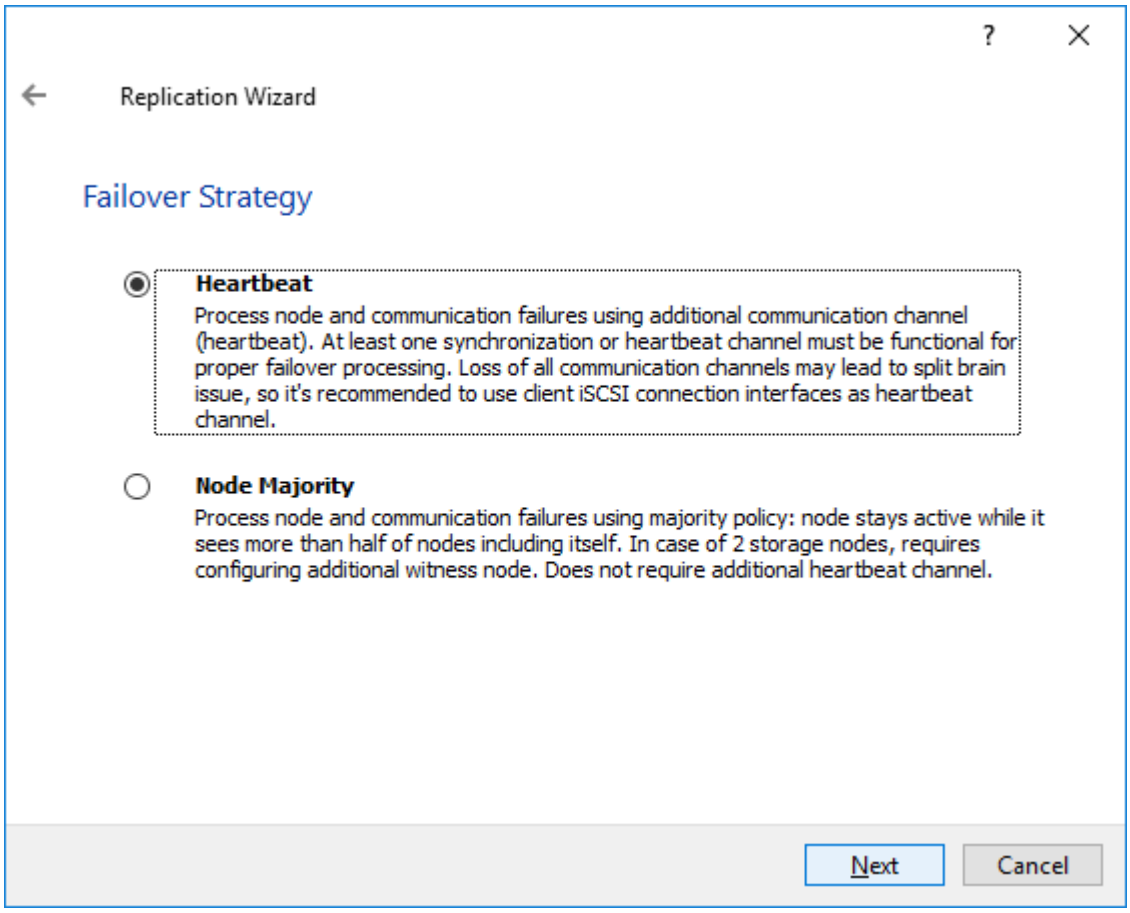
10. Select Synchronous “Two-Way” Replication as a replication mode. Click Next to proceed.



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



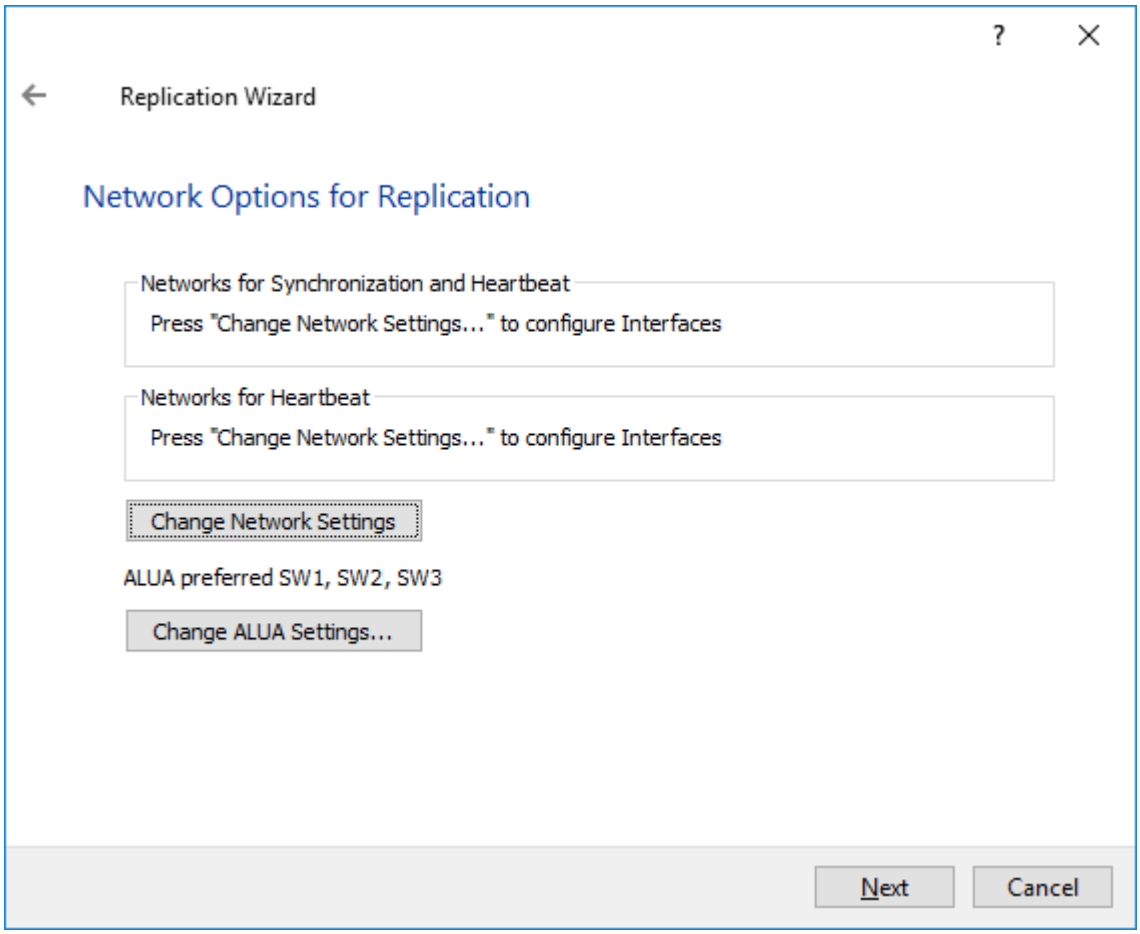
12. Select Failover Strategy.



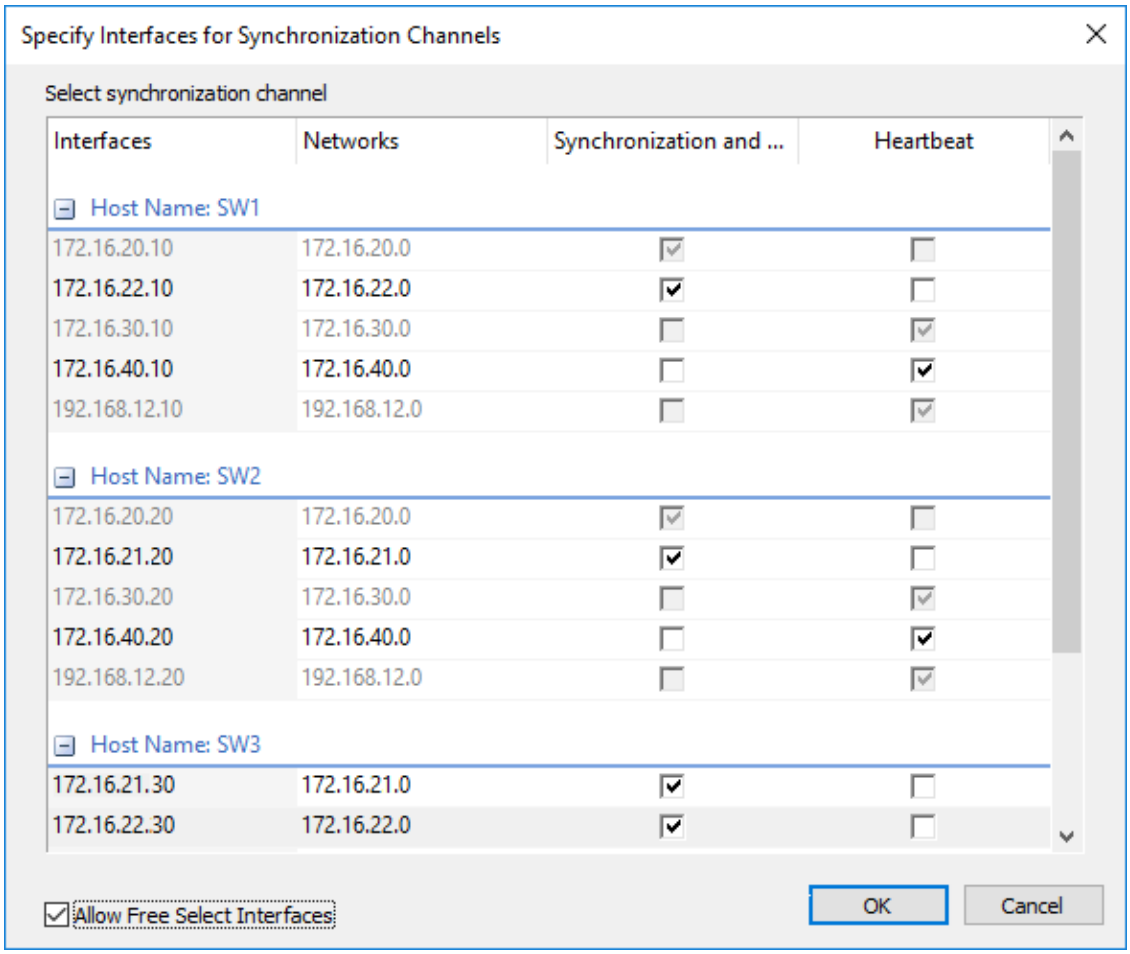
13. Select Create new Partner Device and click Next.

14. Select a partner device Location and Synchronization Journal Strategy. Set journal location if required.

15. Click Change Network Settings.



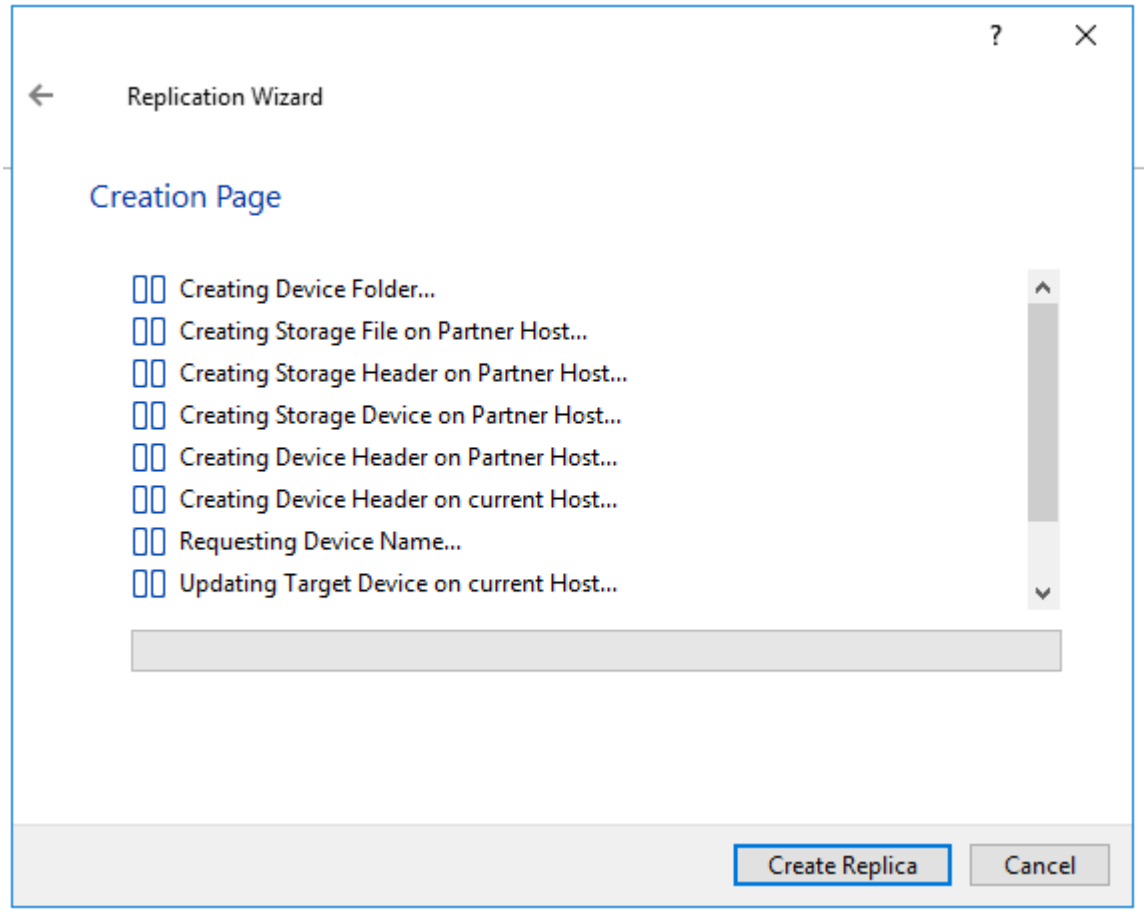
16. Specify the interfaces for Synchronization and Heartbeat Channels. Click OK and then click Next.



NOTE: It is not recommended to configure the Heartbeat and iSCSI channels on the same interfaces to avoid the split-brain issue. If the Synchronization and Heartbeat interfaces are located on the same network adapter, it is recommended to assign one more Heartbeat interface to a separate adapter.

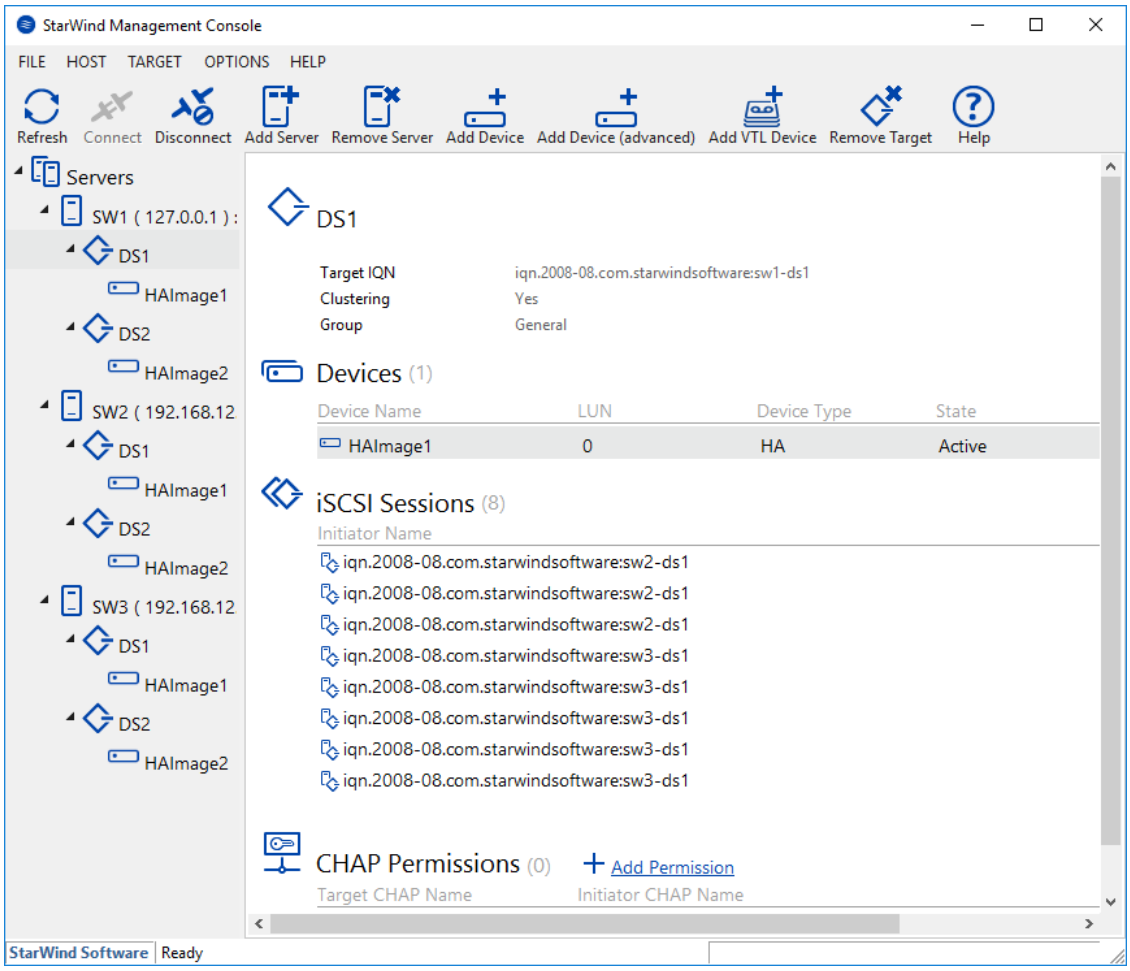
17. In Select Partner Device Initialization Mode, select Synchronize from existing Device and click Next.

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



The successfully added device appears in StarWind Management Console.

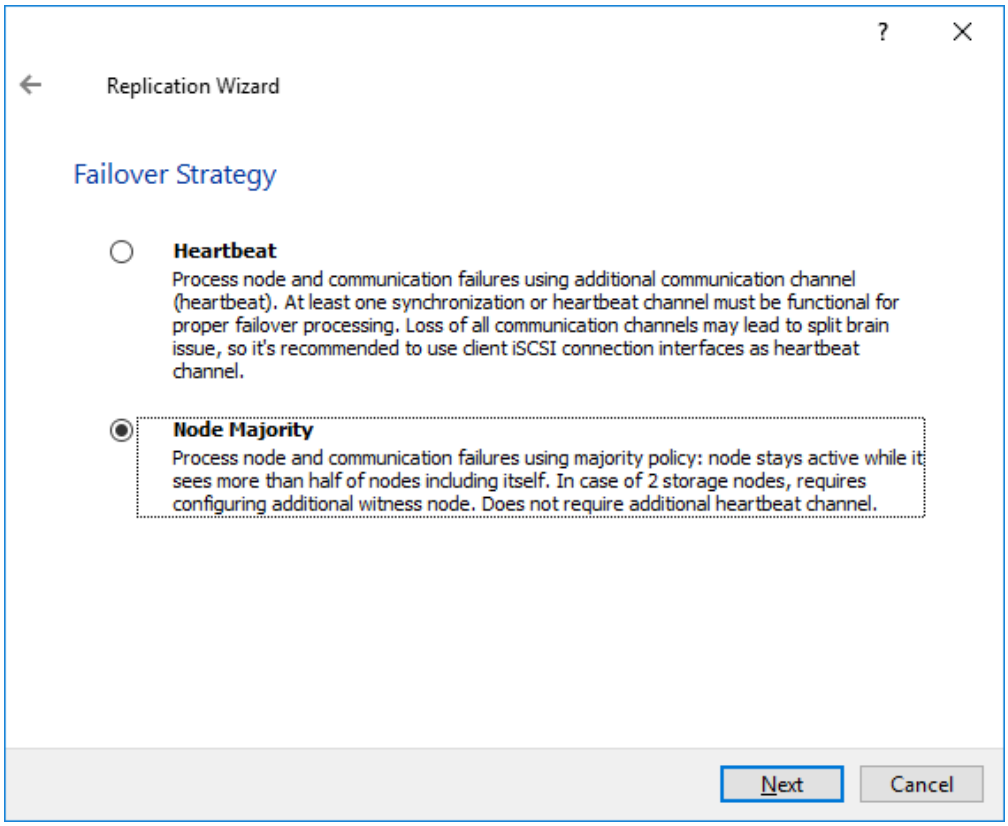
19. Follow the similar procedure for the creation of other virtual disks that will be used as storage repositories.



NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below: https://knowledgebase.starwindsoftware.com/maintenance/how-to-extend-image-file-or-high-availability-device/

Node Majority

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.

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

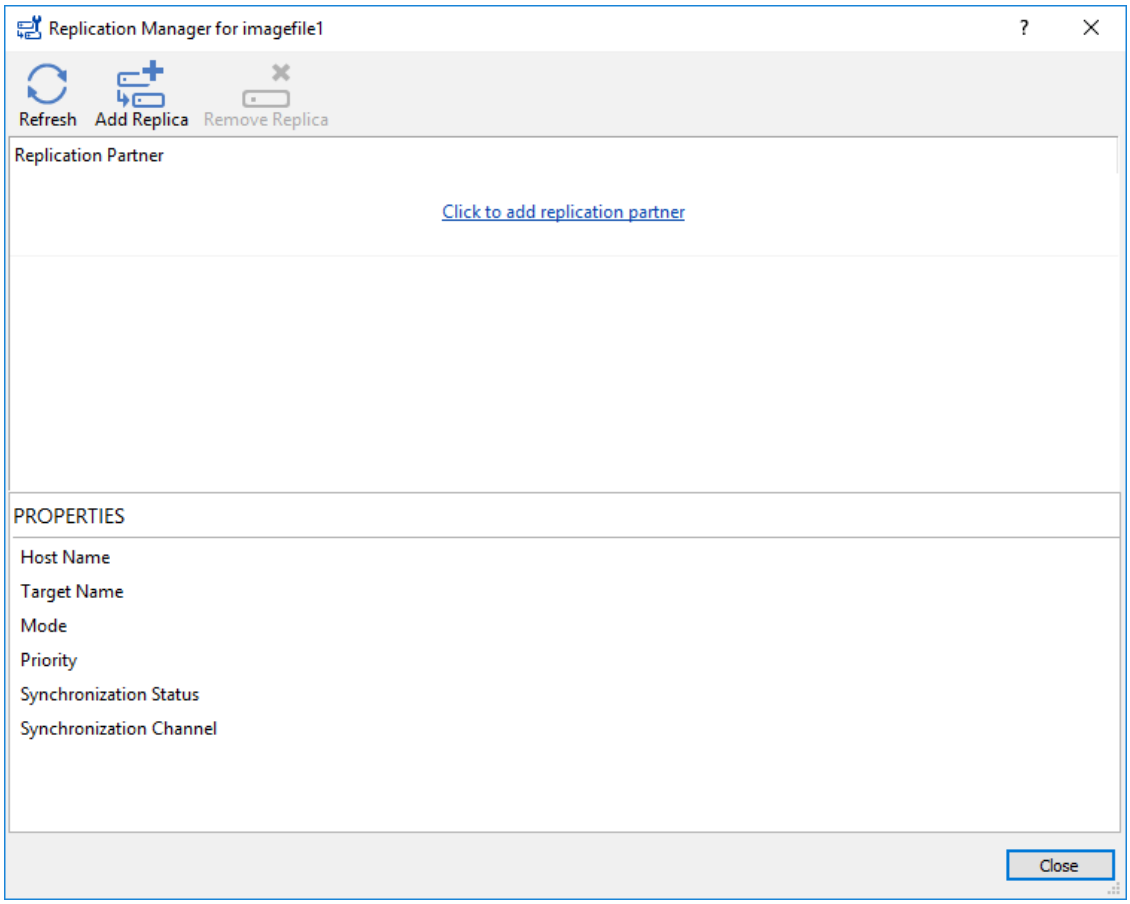
6. In Specify Interfaces for Synchronization Channels, select the checkboxes with the appropriate networks and click OK. Then click Next.

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

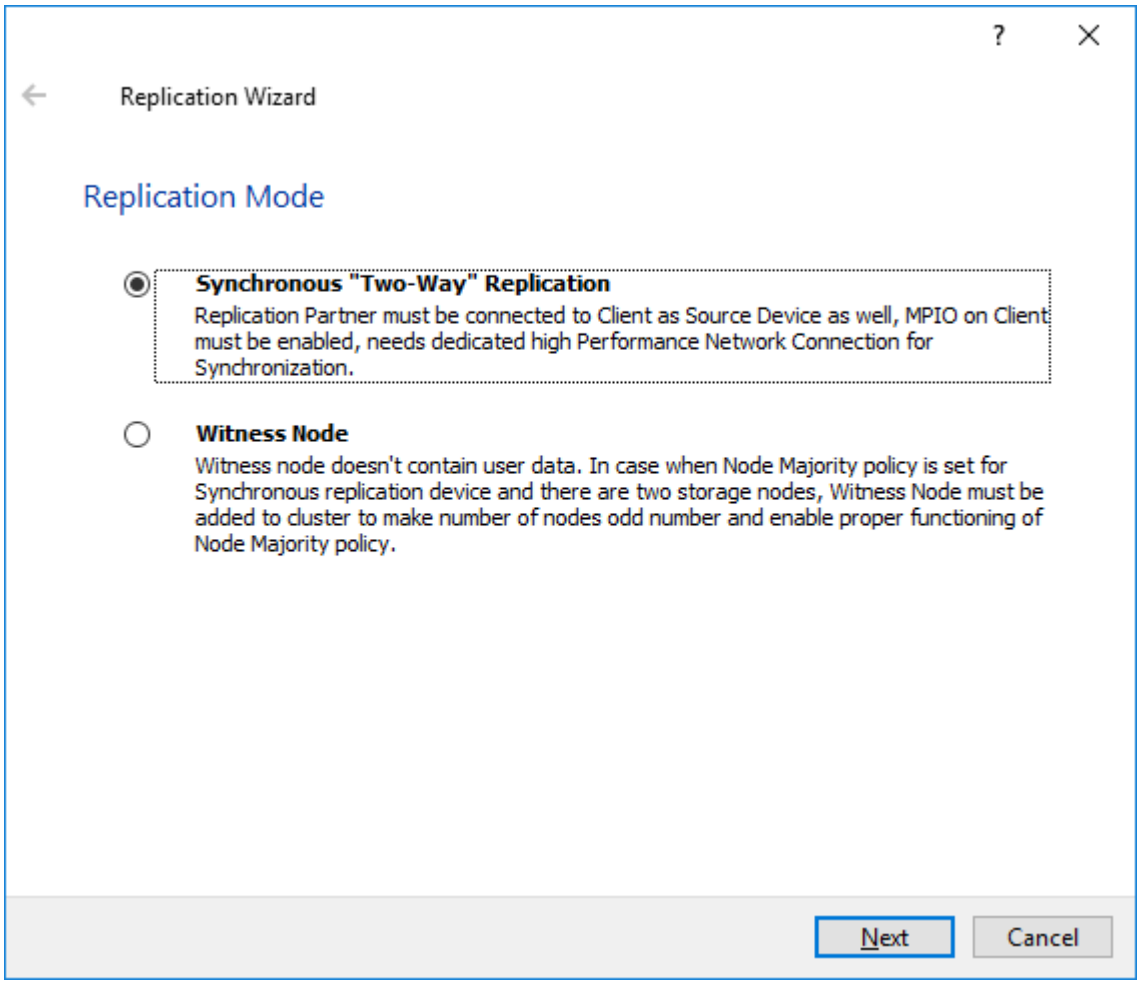
8. Press the Create Replica button and close the wizard.

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

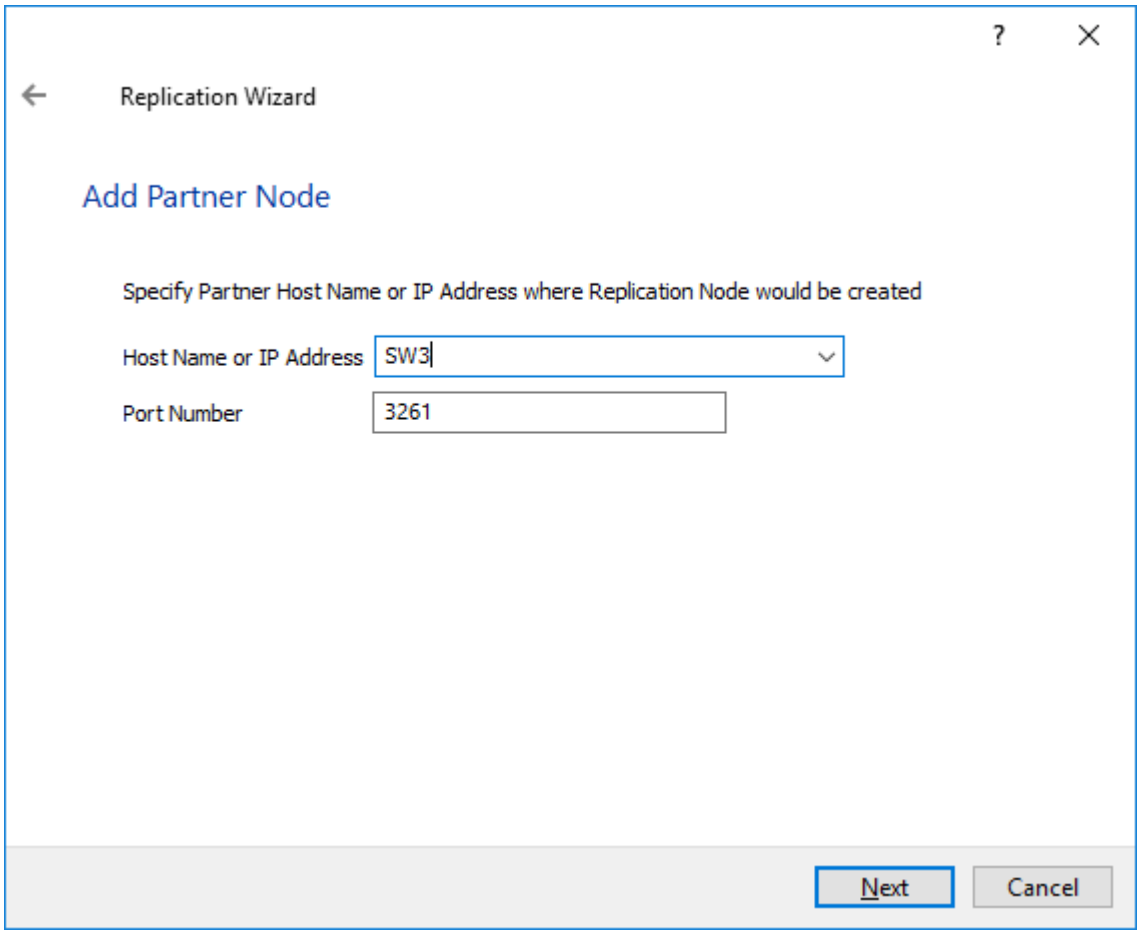
10. Choose device, open Replication Manager and click Add replica again.



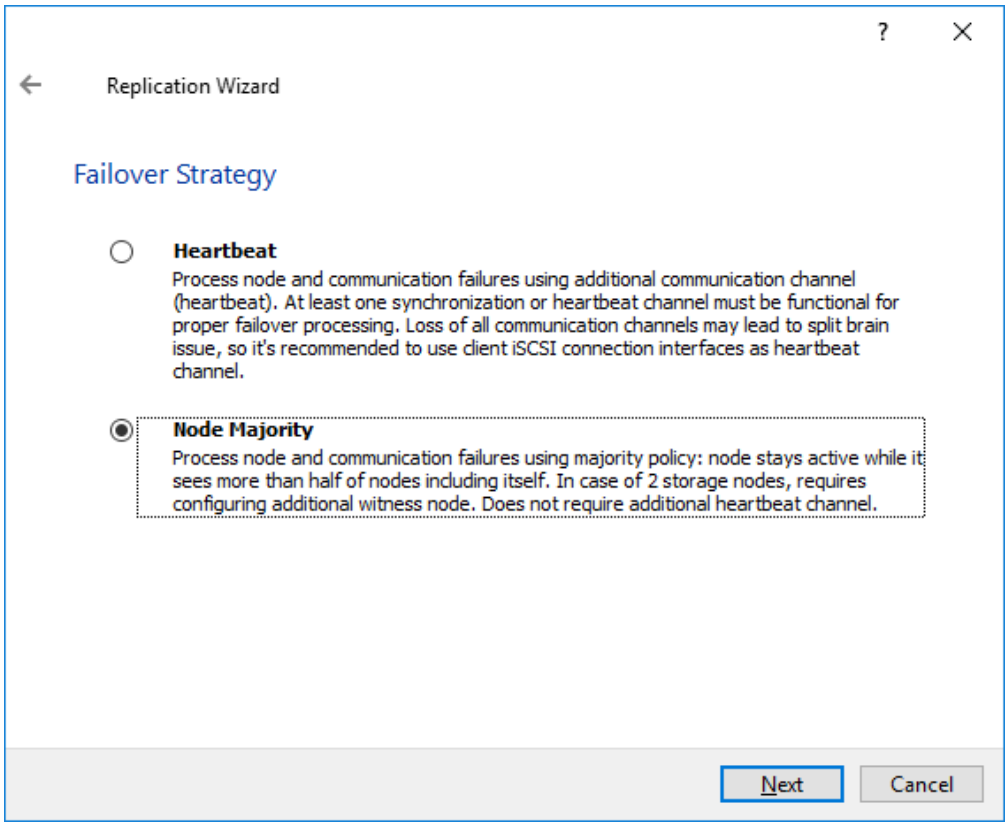
11. Select Synchronous “Two-Way” Replication as a replication mode. Click Next to proceed.



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



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



14. Choose Create new Partner Device and click Next.

15. Specify the partner device Location and modify the target name if necessary. Click Next.

16. Select Synchronization Journal Strategy and click Next.

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

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

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

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

21. The added devices will appear in StarWind Management Console.

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:

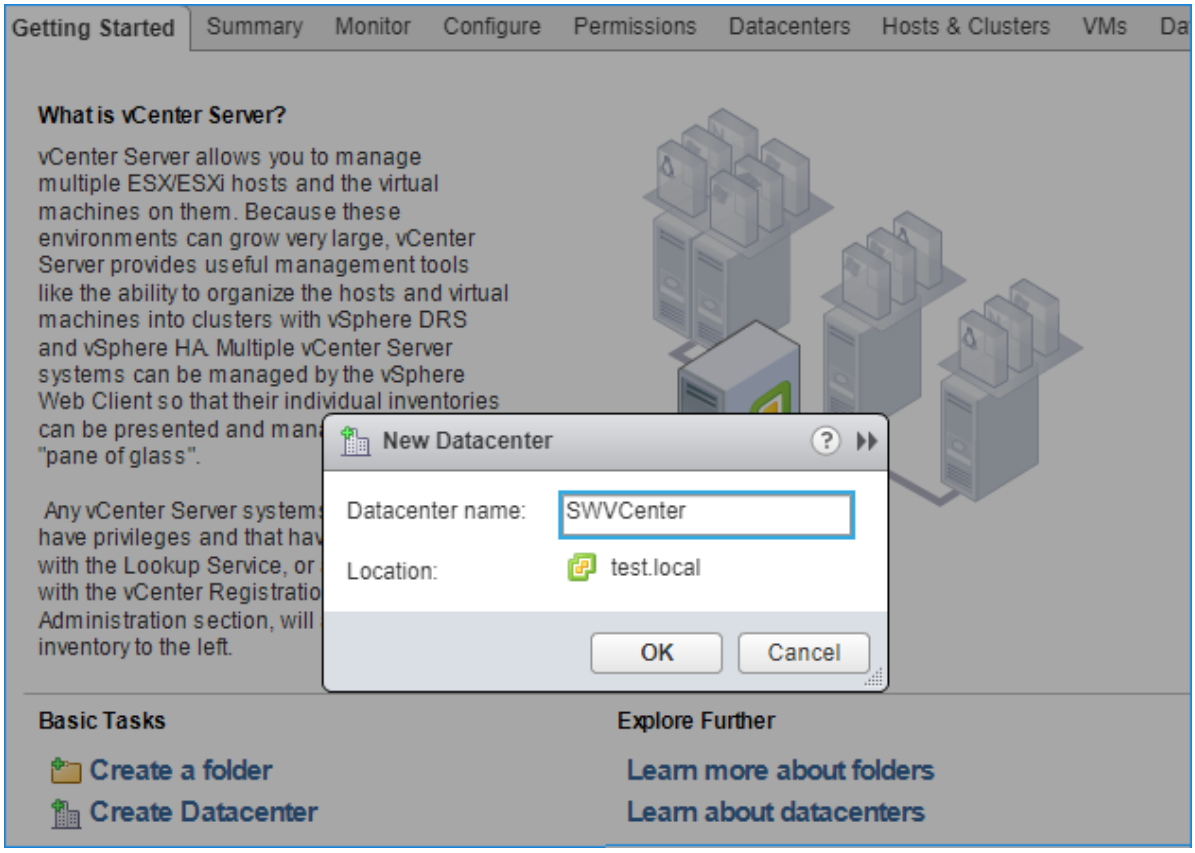
https://knowledgebase.starwindsoftware.com/maintenance/how-to-extend-image-file-or-high-availability-device/

Datacenter And Cluster Creation

Creating Datacenter

NOTE: Creating Datacenter requires deployment of vCenter prior to this step. In case of deploying vCenter as a virtual machine (vCenter Server Appliance) located on the local storage it is recommended to move vCenter to a StarWind HA datastore after its creation.

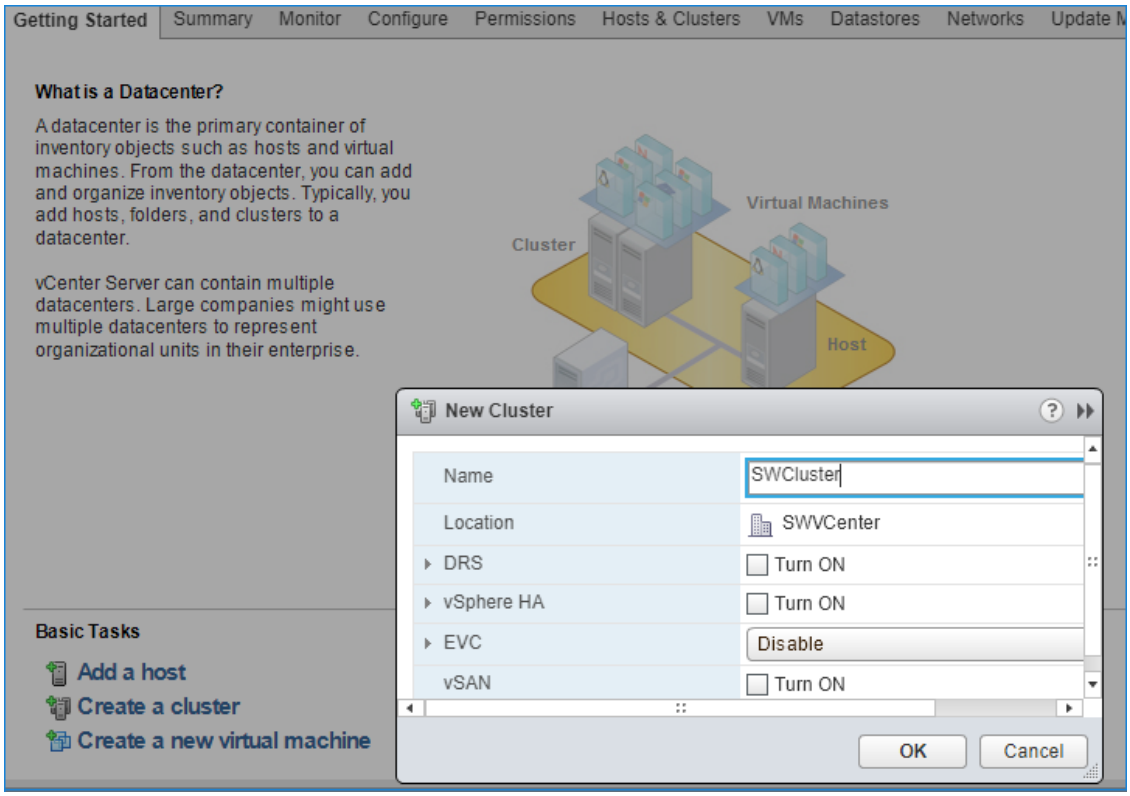
1. Connect to vCenter, select the Getting Started tab, click on Create Datacenter, and enter the Datacenter name.



Creating Cluster

2. Click the Datacenter's Getting Started tab and click Create a cluster. Enter the name

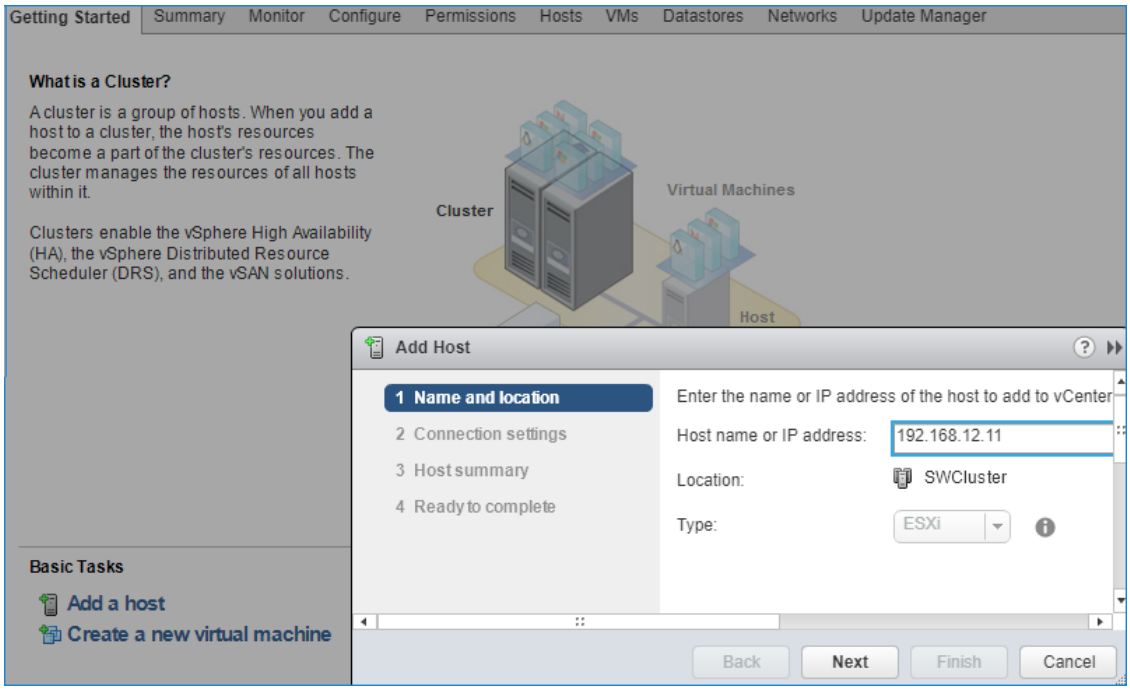
of the cluster and click Next.



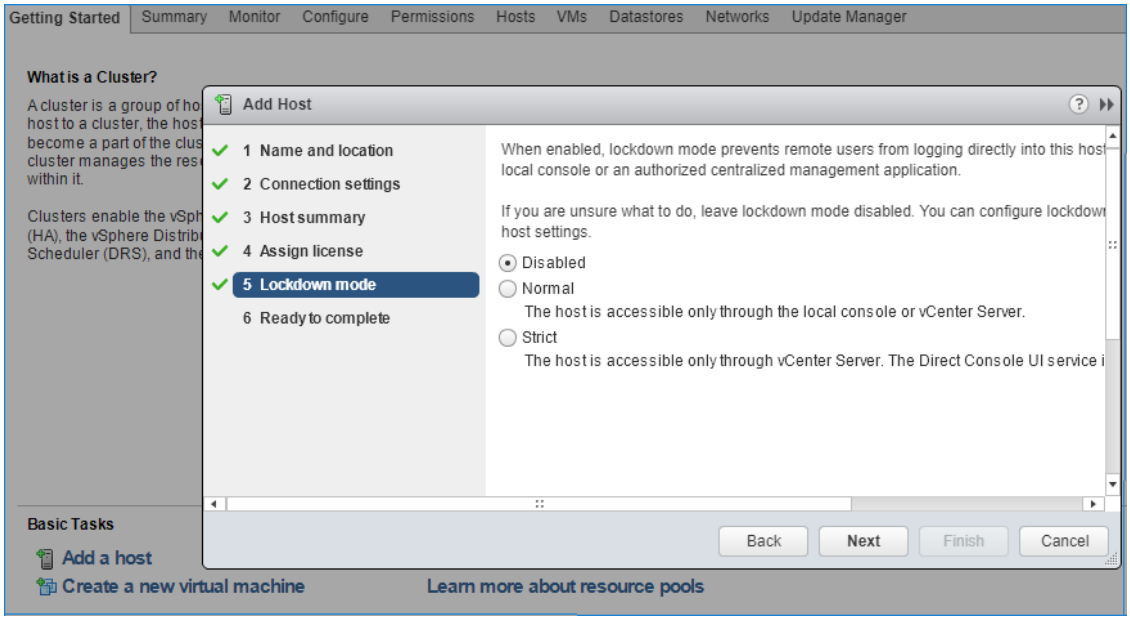
Adding Hosts to Cluster

3. Open the Cluster tab and click on Add a host.

4. Enter the name or IP address of the ESXi host and its administrative account.

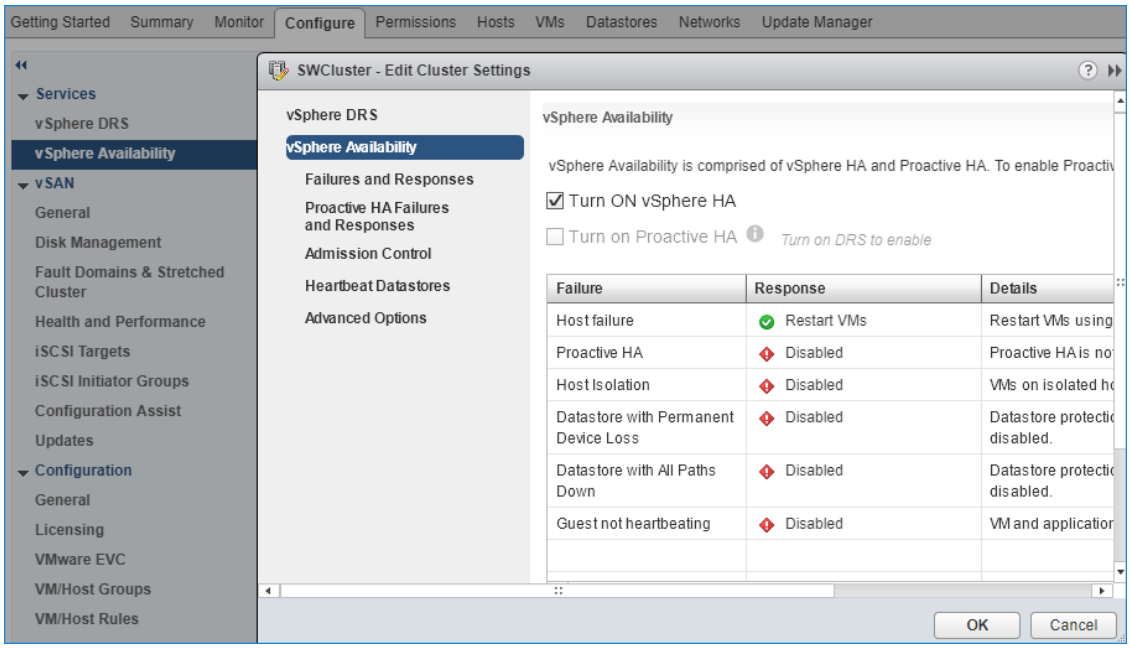


5. Lockdown mode is not enabled by default.



6. Assign the License from the Licensing tab.

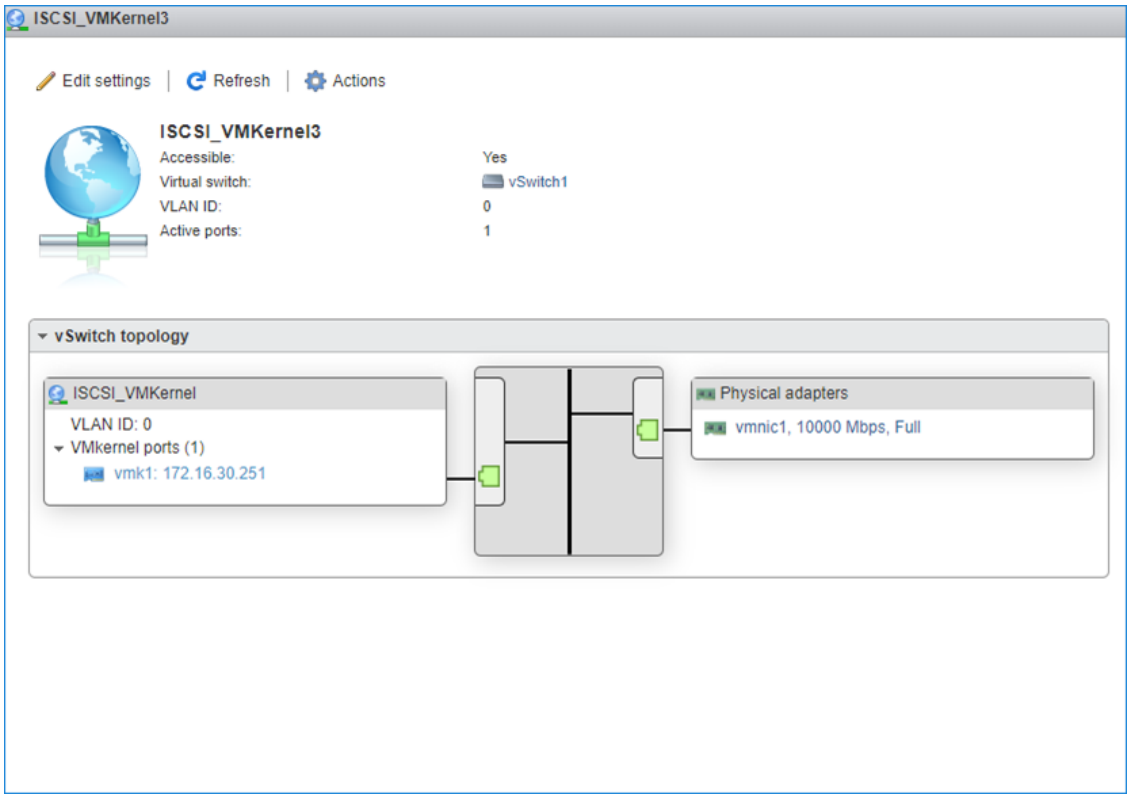
7. Turn on vSphere HA by clicking on Cluster -> Configure -> Edit



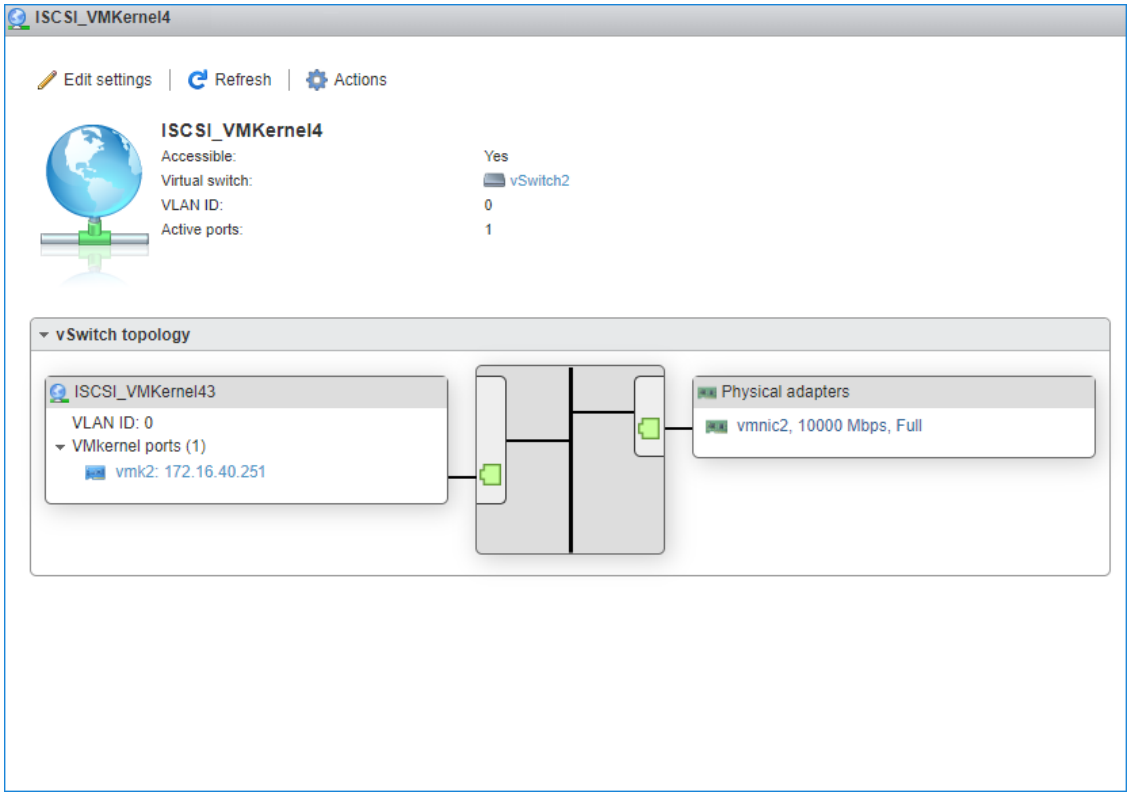
Configuring Networks

Perform the network configuration on each ESXi host.

1. Create a vSwitch for the iSCSI traffic on the first physical adapter.



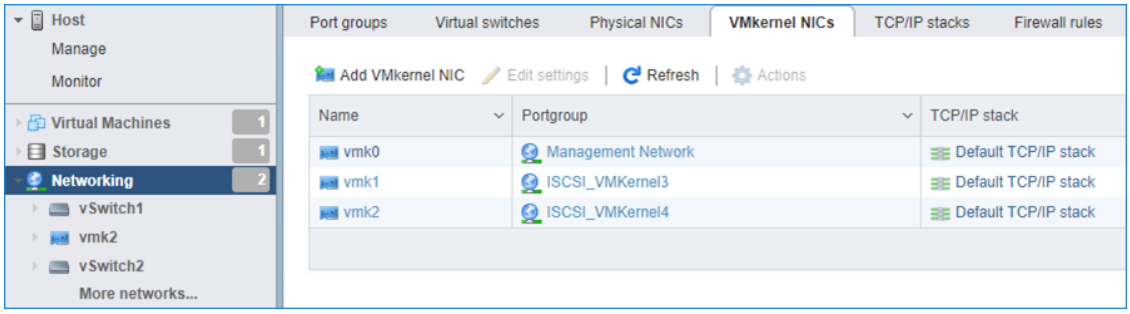
2. Create a vSwitch for the iSCSI traffic on the second physical adapter.



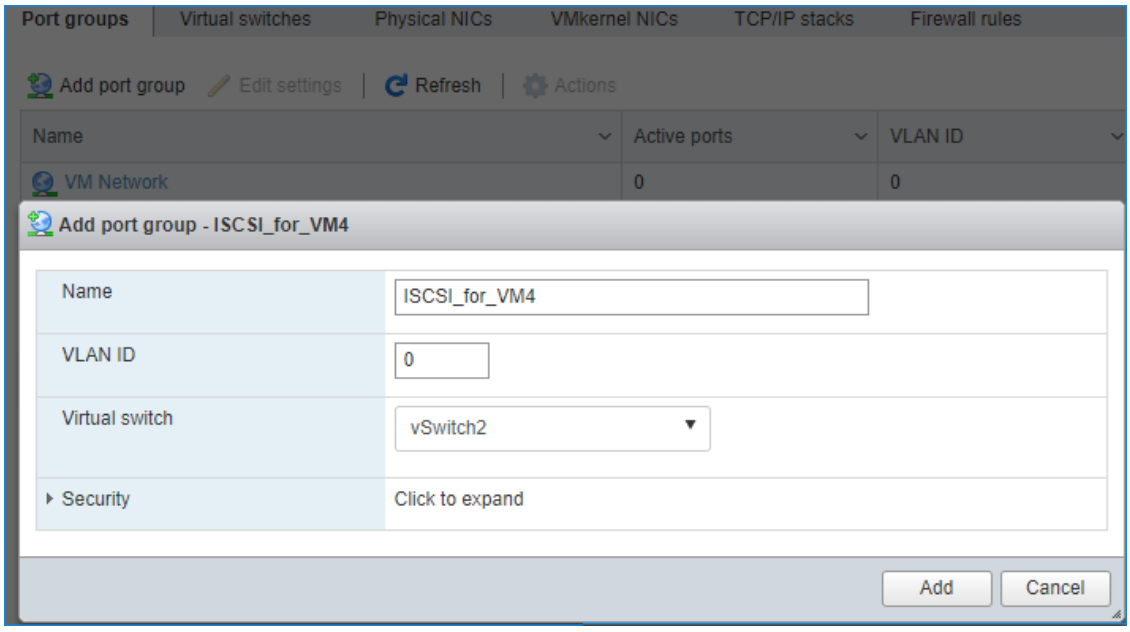
NOTE: Virtual Machine Port Group and VMKernel should be created for every iSCSI vSwitch. Static IP address should be assigned to VMKernel port.

NOTE: It is recommended to set MTU to 9000 on vSwitches and VMKernel ports for the iSCSI traffic. Additionally, vMotion can be enabled on VMKernel port.

3. Create VMKernel ports one per each iSCSI channel.



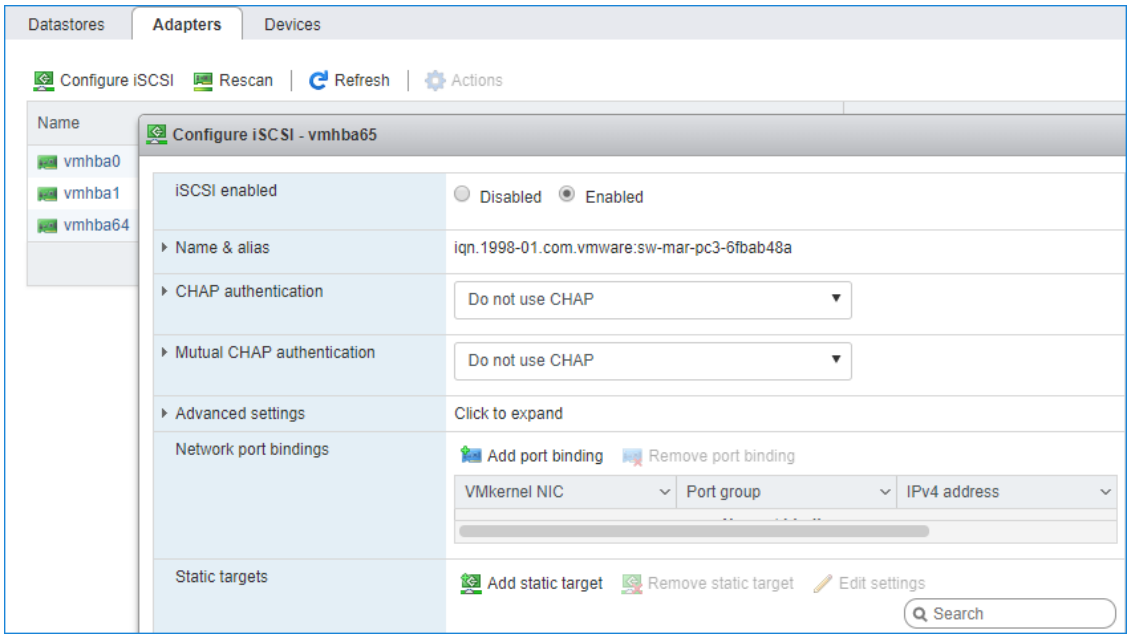
4. In VMkernel adapters pane, click Add Networking and Add Virtual Machine Port Group on the vSwitch for the iSCSI traffic.



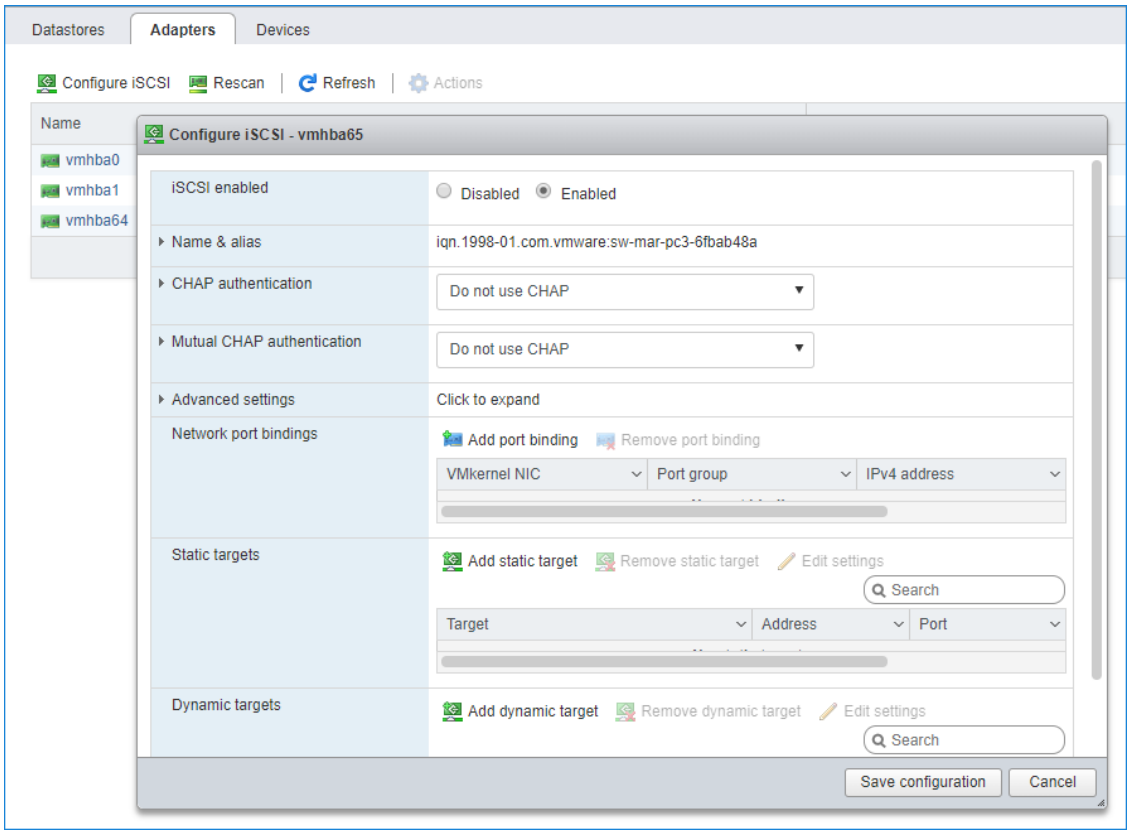
Preparing Datastores

Adding Discover Portals

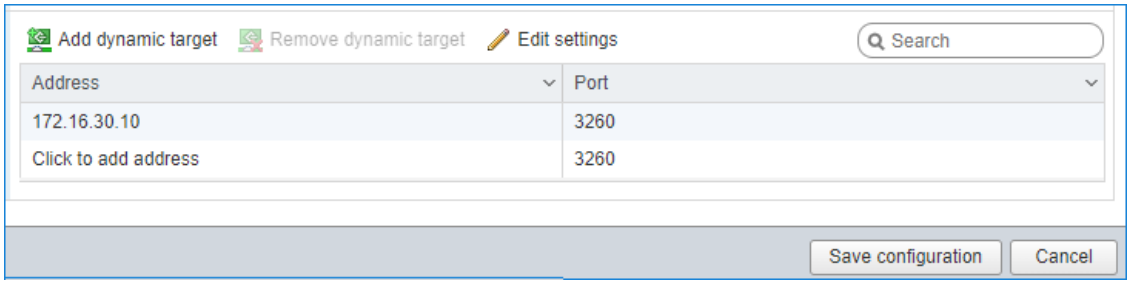
1. To connect the previously created devices to the ESXi host, click on the Storage -> Adapters -> Configure iSCSI and choose the Enabled option to enable Software iSCSI storage adapter.



2. In the Configure iSCSI window, under Dynamic Targets, click on the Add dynamic target button to specify iSCSI interfaces.

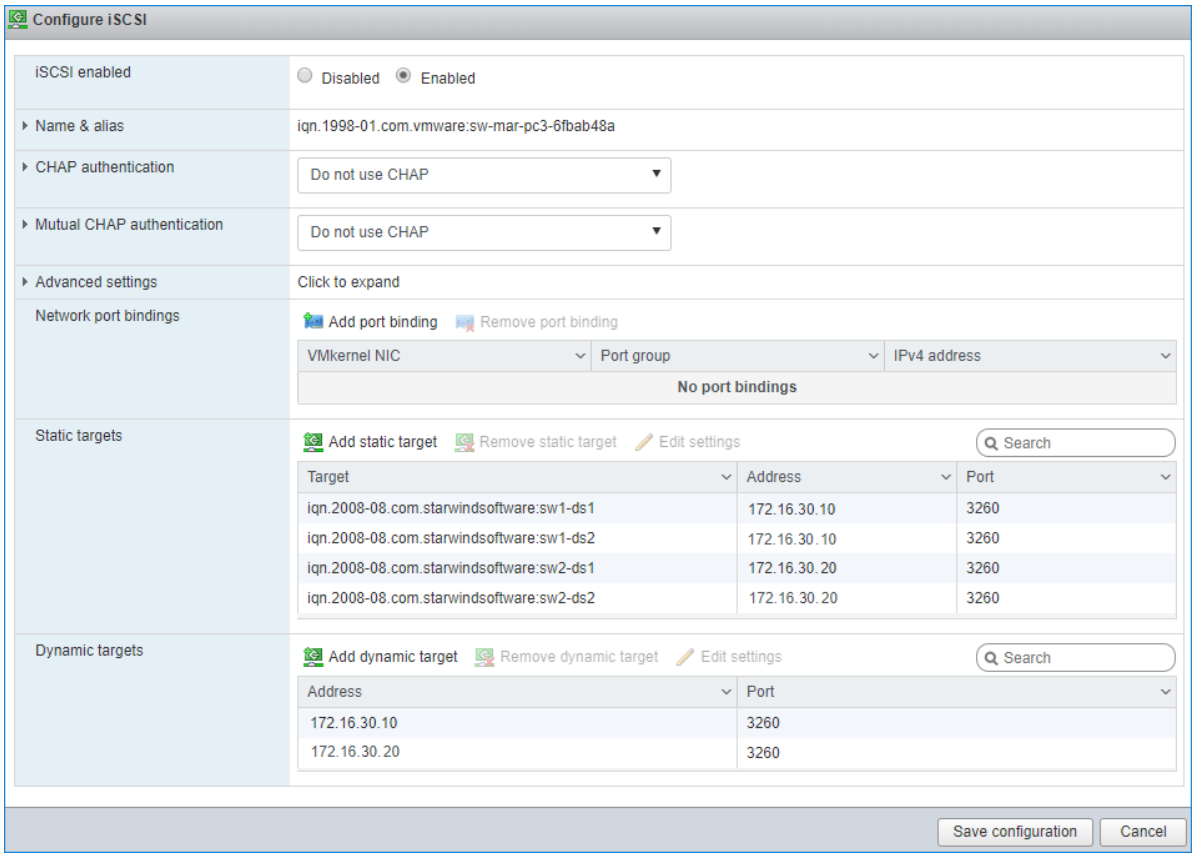


3. Enter the iSCSI IP addresses of all StarWind nodes for the iSCSI traffic.

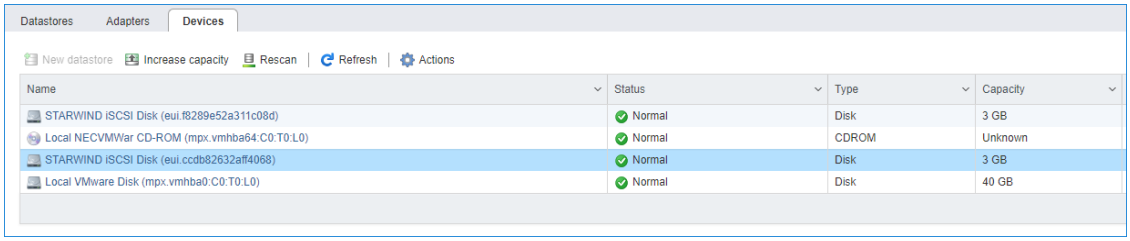


Confirm the actions by pressing Save configuration.

4. The result should look like in the image below.



5. Click on the Rescan button to rescan storage.

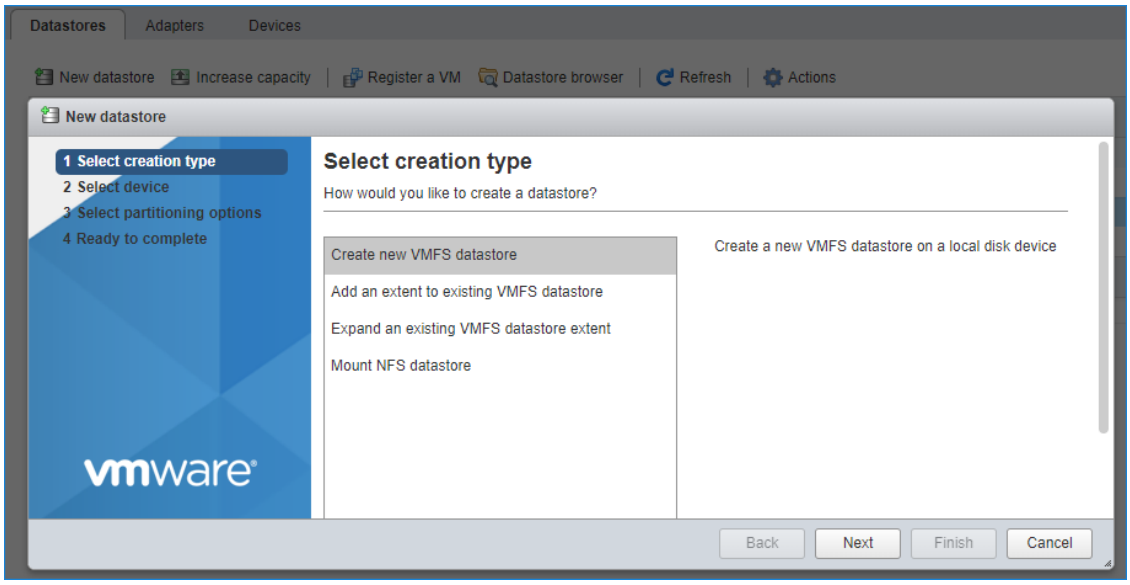


6. Now, the previously created StarWind devices are visible to the system.

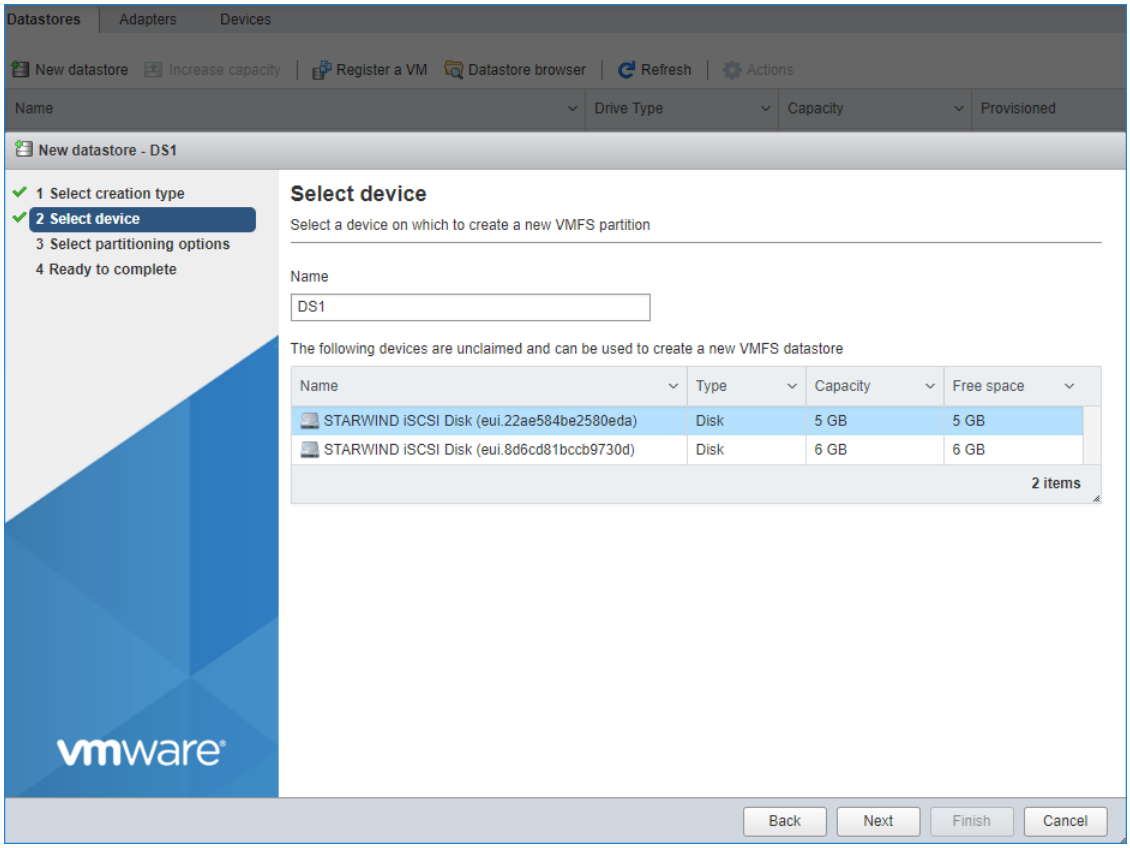
7. Repeat all the steps from this section on the other ESXi host, specifying corresponding IP addresses for the iSCSI subnet.

Creating Datastores

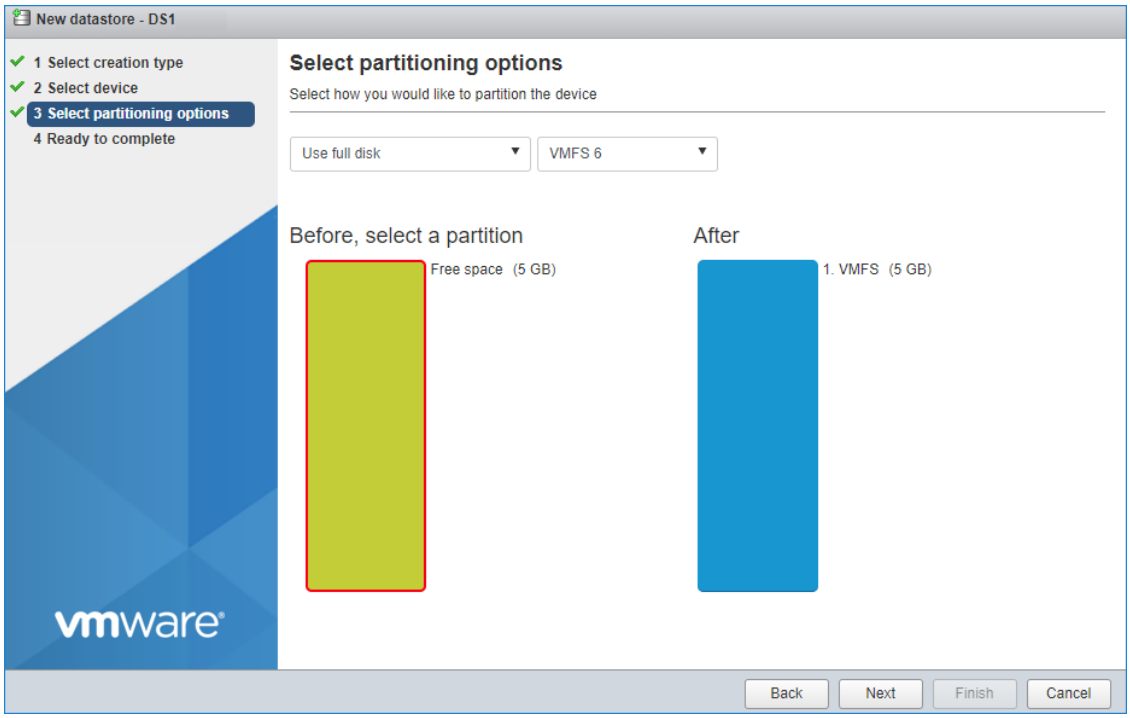
1. Open the Storage tab on one of the hosts and click on New Datastore.



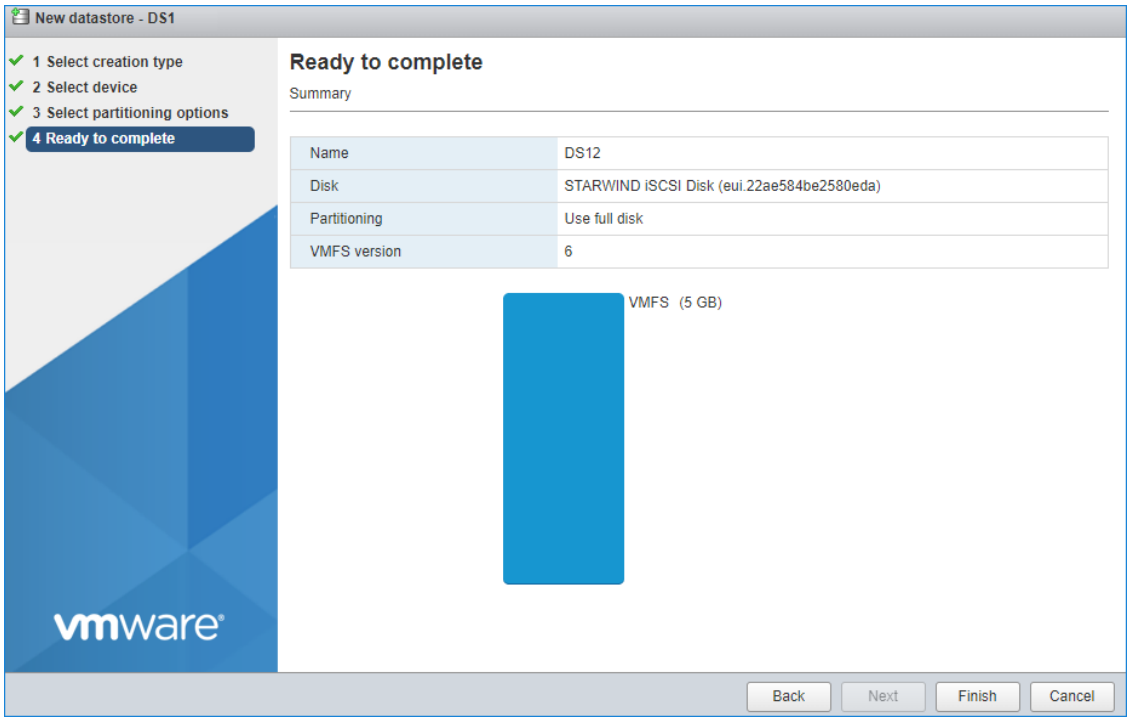
2. Specify the datastore name, select the previously discovered StarWind device, and click on Next.



3. Enter datastore size. Click on Next.

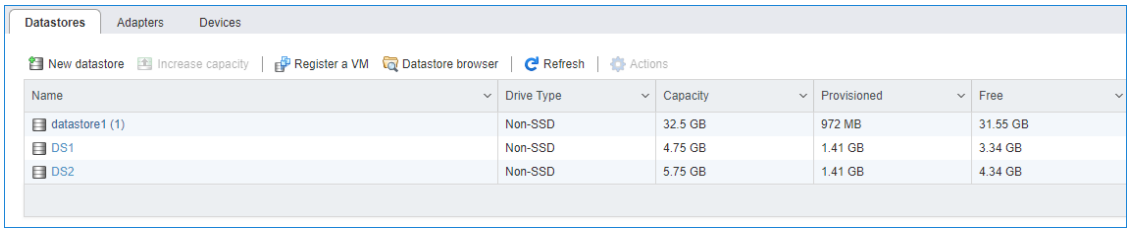


4. Verify the settings. Click on Finish.



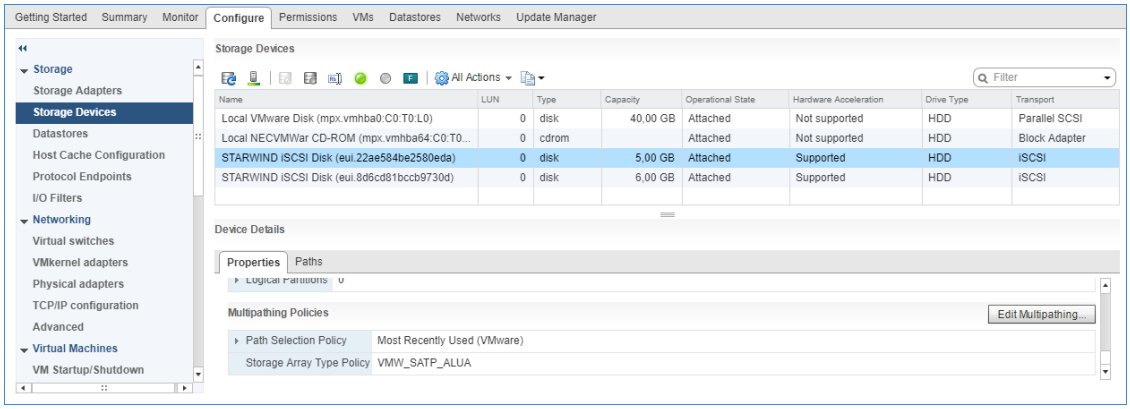
5. Add another datastore (DS2) in the same way but select the second device for it.

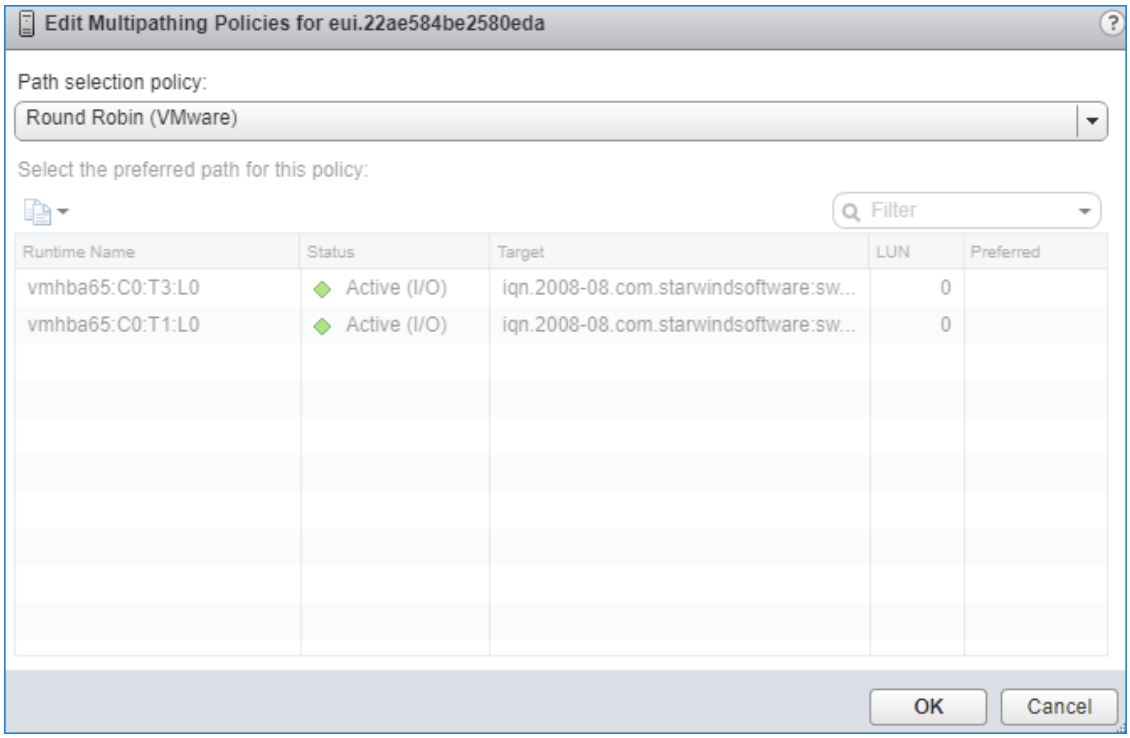
6. Verify that storage (DS1, DS2) is connected to both hosts. Otherwise, rescan the storage adapter.



7. Path Selection Policy changing for Datastores from Most Recently Used (VMware) to Round Robin (VMware) has been already added into the Rescan Script, and this action is performed automatically. For checking and changing this parameter manually, the hosts should be connected to vCenter.

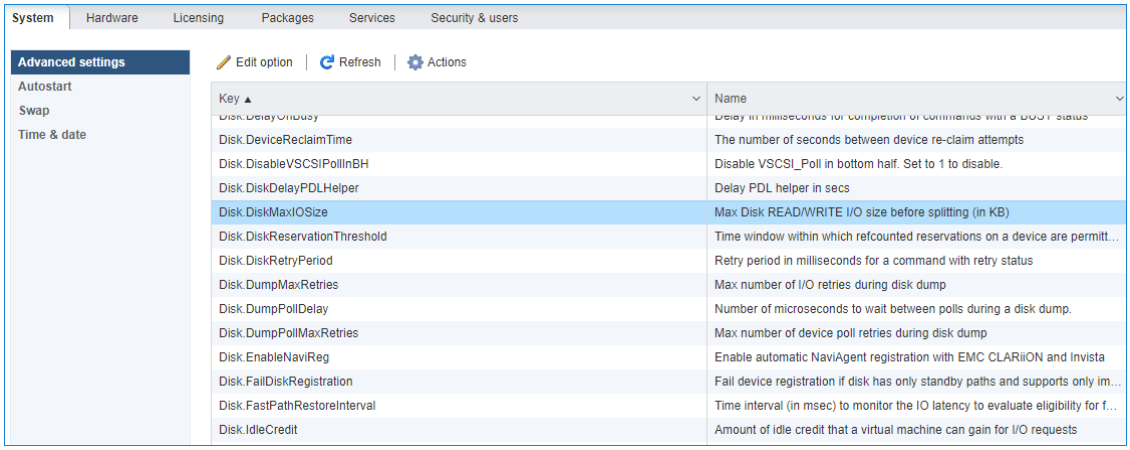
8. Multipathing configuration can be checked only from vCenter. To check it, click the Configure button, choose the Storage Devices tab, select the device, and click on the Edit Multipathing button.



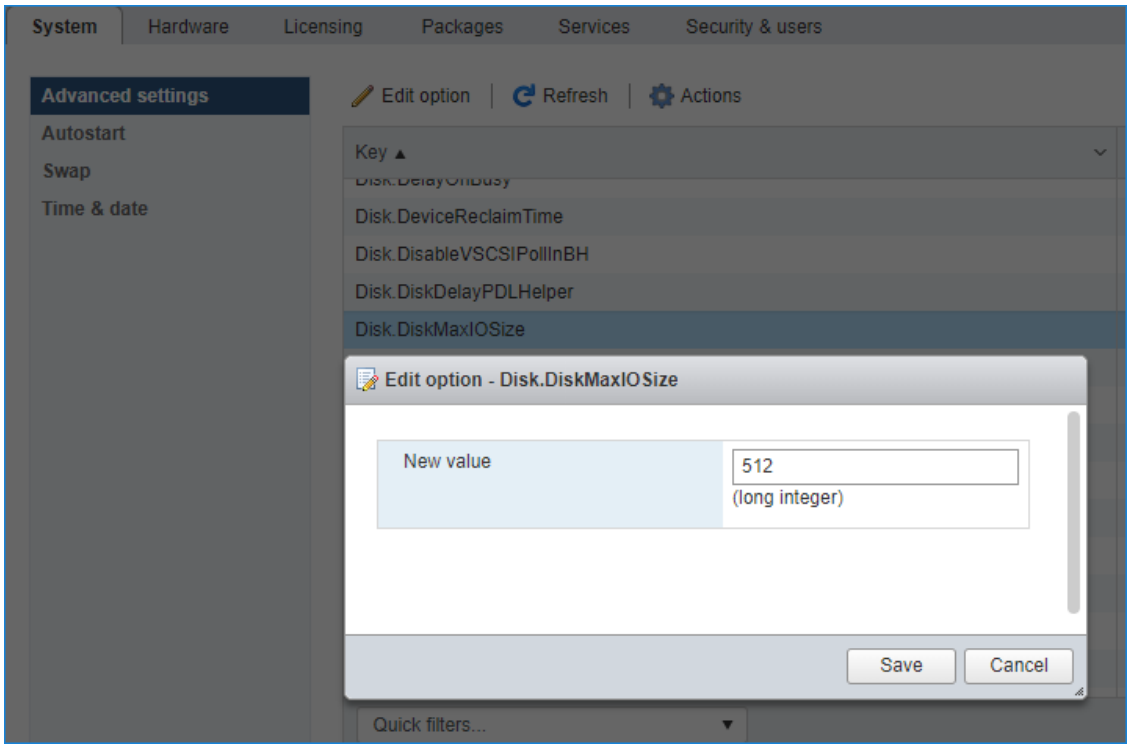


Performance Tweaks

1. Click on the Configuration tab on all of the ESXi hosts and choose Advanced Settings.



2. Select Disk and change the Disk.DiskMaxIOSize parameter to 512.



NOTE: Changing Disk.DiskMaxIOSize to 512 might cause startup issues with Windows-based VMs, located on the datastore where specific ESX builds are installed. If the issue with VMs start appears, leave this parameter as default or update the ESXi host to the next available build.
NOTE: In certain cases, in Virtual Machine, Windows event log may report an error similar to “Reset to device, \Device\RaidPort0, was issued”. Check this KB acticle for a possible solution.

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

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