StarWind Virtual SAN: Stretched Cluster Configuration Guide for VMware vSphere [ESXi], VSAN Deployed as a Controller VM 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.

Annotation

Relevant products

This guide is applicable to StarWind Virtual SAN and StarWind Virtual SAN Free (OVF Version 20230901 Version V8 (build 15260) or earlier).

Purpose

This document outlines how to configure a stretched cluster based on VMware vSphere ESXi, with StarWind Virtual SAN (VSAN) running as a Controller Virtual Machine (CVM). The guide includes steps to prepare ESXi hosts for clustering, configure physical and virtual networking, and set up the Virtual SAN Controller Virtual Machine.

Audience

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

Expected result

The end result of following this guide will be a fully configured high-availability StarWind Virtual HCI Appliance (VHCA) powered by VMware ESXi that includes virtual machine shared storage provided by StarWind VSAN.

Infrastructure Design

Prerequisites For Stretched Cluster

StarWind Virtual SAN system requirements

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

Recommended RAID settings for HDD and SSD disks:
https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-h

dd-and-ssd-disks/

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

Solution diagram

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

Make sure that the prerequisites for deploying StarWind stretched cluster on VMware are met:

- L2/L3 multisite network applied according to the appropriate StarWind failover strategy
- each iSCSI and Synchronization network channel throughput should be at least 1Gbps
  the 10Gbps or higher link bandwidth is highly recommended
- the maximum supported latency for StarWind synchronous storage replication should be 10ms round-trip time (RTT)

- the maximum supported latency between the ESXi ethernet networks should be 10ms round-trip time (RTT)
- vSphere 6.5 or newer installed on the servers to be clustered
- StarWind Virtual SAN installed on Windows Server VMs

Heartbeat Failover Strategy For Stretched Cluster

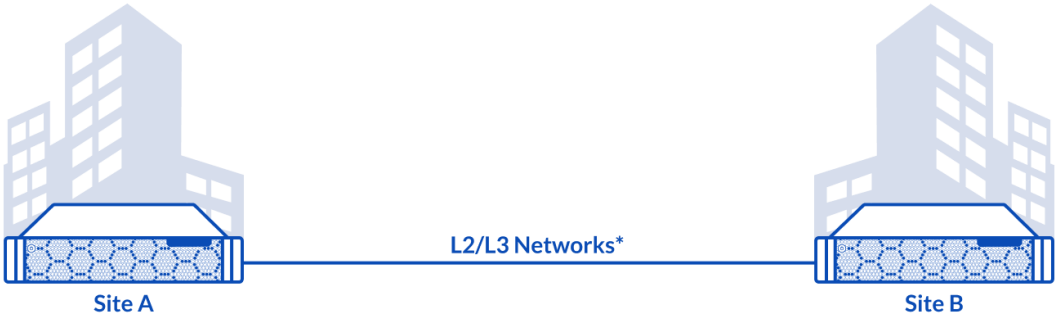

Heartbeat is a technology that allows avoiding the so-called “split-brain” scenario when the HA cluster nodes are unable to synchronize but continue to accept write commands from the initiators independently. It can occur when all synchronization and heartbeat channels disconnect simultaneously, and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

If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The services mark the device with the lowest priority as not-synchronized one. Subsequently it gets blocked for further read and write operations until the synchronization channel resumption. Then, the partner device on the synchronized node flushes data from the cache to the disk to preserve data integrity in case the node goes down unexpectedly. It is recommended to assign more independent heartbeat channels during replica creation to improve system stability and avoid the “split-brain” issue. With the Heartbeat Failover Strategy, the storage cluster will continue working with only one StarWind node available.

Heartbeat Failover Strategy Network Design

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

highly recommended.

Node Majority Failover Strategy For Stretched Cluster

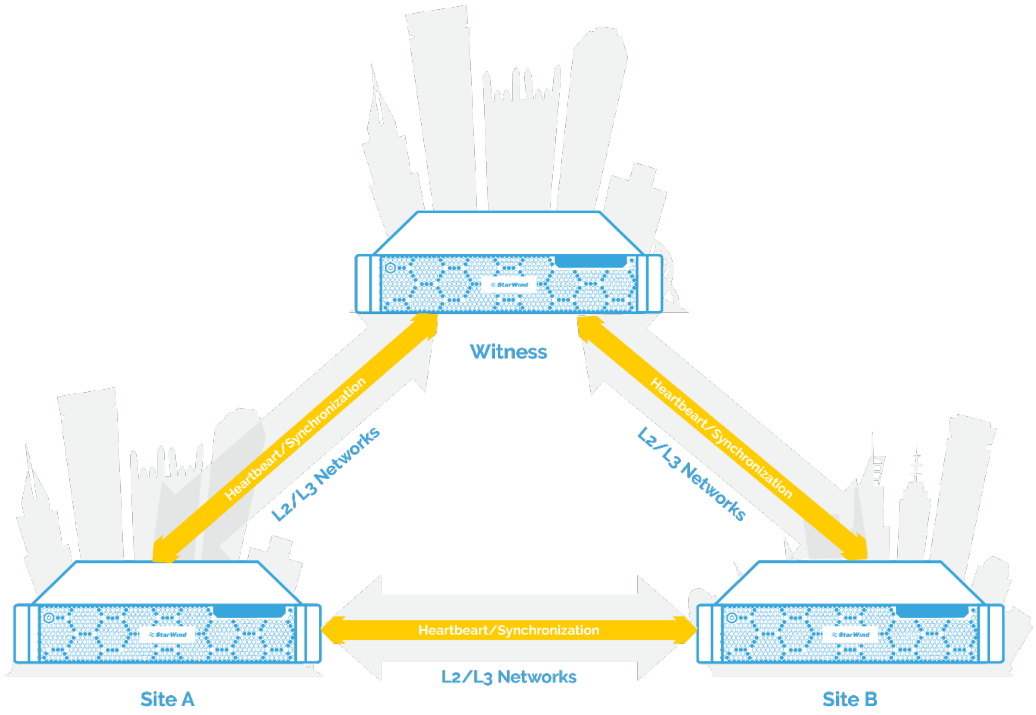

This strategy ensures synchronization connection without any additional heartbeat links. The failure-handling process occurs when the node has detected the absence of connection with the partner. The main requirement for keeping the node operational is an active connection with more than a half of the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

- Management / Heartbeat /Synchronization – 1Gbps network or higher. The 10Gbps

or higher bandwidth link is highly recommended.

Preparing Environment For Starwind Vsan Deployment

Configuring Networks

Configure network interfaces on each node to make sure that Synchronization and iSCSI/StarWind heartbeat interfaces are in different subnets and connected physically according to the network diagram above. All actions below should be applied to each ESXi server.
NOTE: Virtual Machine Port Group should be created for both iSCSI/ StarWind Heartbeat and the Synchronization vSwitches. VMKernel port should be created only for iSCSI traffic. Static IP addresses should be assigned to VMKernel ports.
NOTE: It is recommended to set MTU to 9000 on vSwitches and VMKernel ports for iSCSI and Synchronization traffic. Additionally, vMotion can be enabled on VMKernel ports.

1. Using the VMware ESXi web console, create two standard vSwitches: one for the iSCSI/ StarWind Heartbeat channel (vSwitch1) and the other one for the Synchronization channel (vSwitch2).

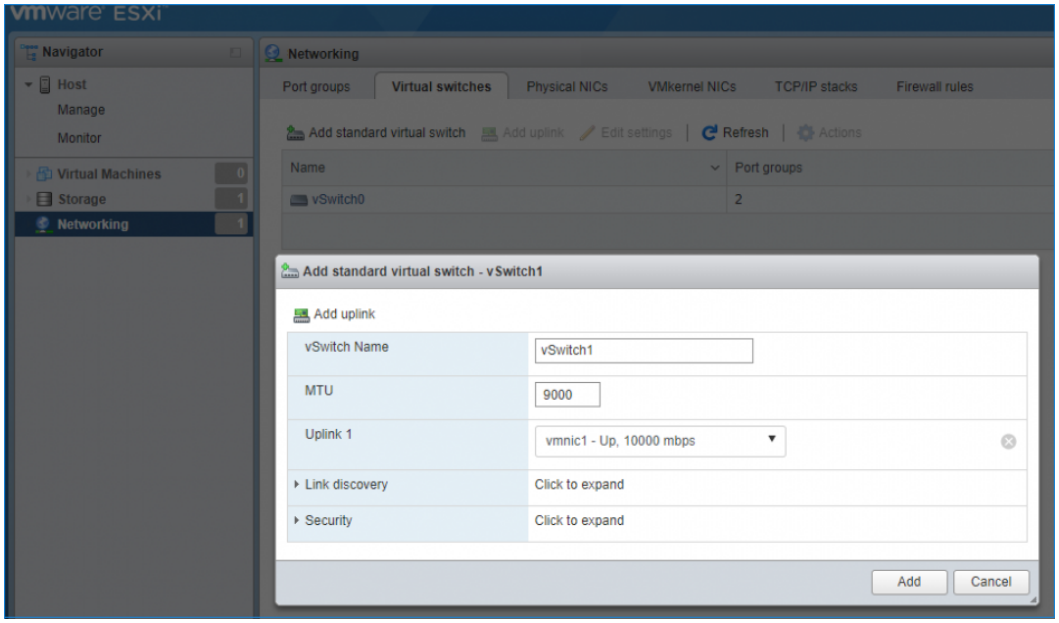

2. Create a VMKernel port for the iSCSI/ StarWind Heartbeat channel.

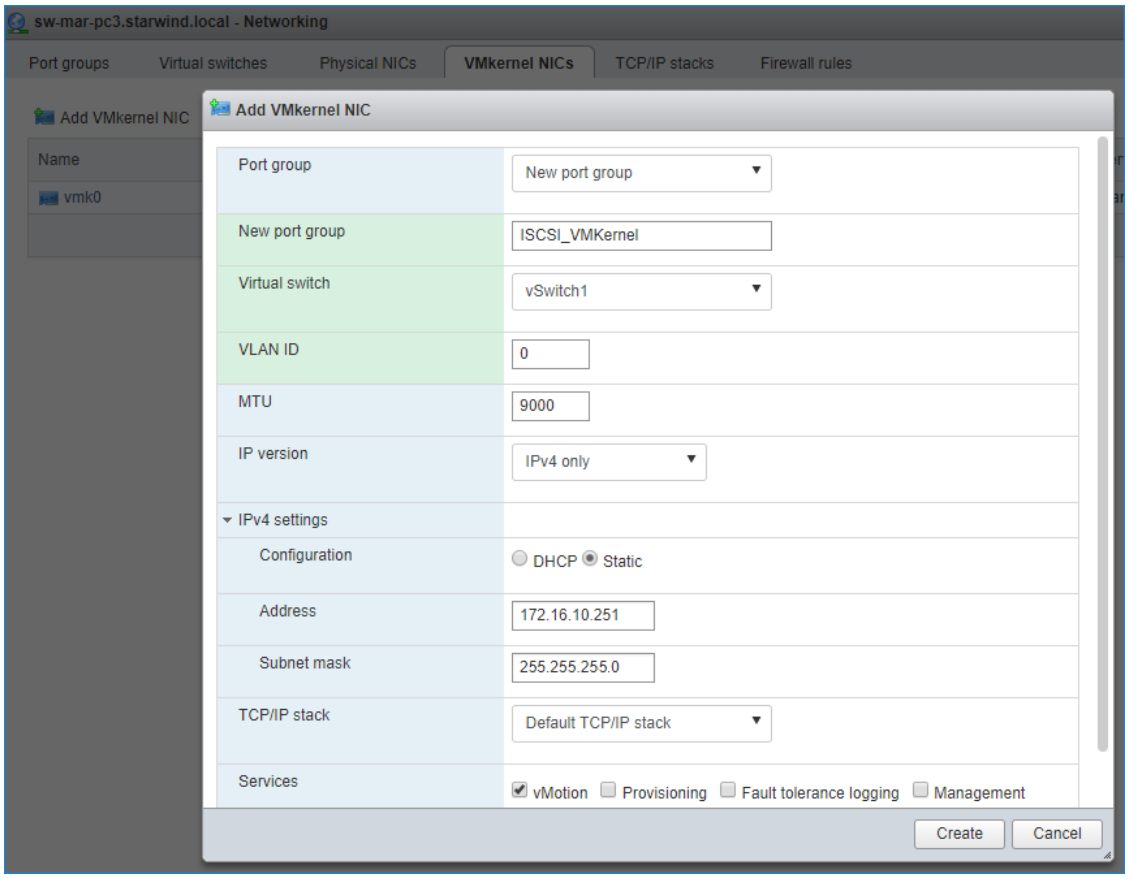

3. Add a Virtual Machine Port Groups on the vSwitch for iSCSI traffic (vSwtich1) and on the vSwitch for Synchronization traffic (vSwitch2).

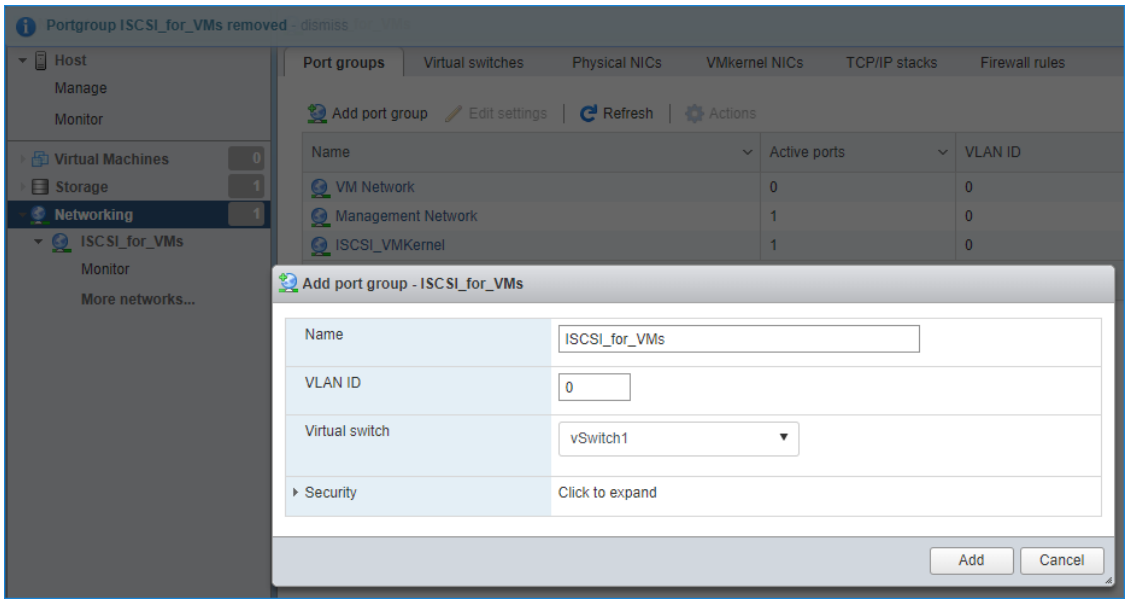

4. Repeat steps 1-3 for any other links intended for Synchronization and iSCSI/Heartbeat

traffic on ESXi hosts.

Preparing Starwind Virtual Machines

Create Virtual Machines (VMs) on each ESXi host with supported Windows Server OS for StarWind VSAN (see system requirements: https://www.starwindsoftware.com/system-requirements) and further StarWind VSAN installation.
StarWind VMs on ESXi hosts should be configured with the following settings:

RAM: at least 4 GB (plus the size of the RAM cache if it is planned to be used) reserved for the VM;

CPUs: at least 4 virtual processors with 2 GHz reserved;

Network adapter 1: Management

Network adapter 2: iSCSI

Network adapter 3: Sync
NOTE: Network adapters for iSCSI and Sync should be of the VMXNET3 type. Network adapter for Management should be of the E1000 type.

Hard disk 1: 100 GB for OS (recommended) – Thick Provisioned Eager Zeroed.

Hard disk 2: Depends on the storage volume to be used as shared storage – Thick Provisioned Eager Zeroed.
NOTE: Alternatively, the disk can be added to StarWind VSAN VM as RDM. The link to VMware documentation is below:
https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-4236E44E-E11F-4EDD-8CC0-12BA664BB811.html

Also, if a separate RAID controller is available, it can be used as dedicated storage for StarWind VM, and RAID controller can be added to StarWind VM as a PCI device. In this case RAID volume will be available as a virtual disk in the Drives section in the Web console. Follow the instructions in the section below on how to add RAID controller as PCI device to StarWind VM.

NOTE: The Active Directory Domain Services role can be added on StarWind Virtual Machine (VM) if necessary, thus it can serve as a domain controller.
NOTE: When using StarWind with the synchronous replication feature inside of a Virtual

Machine, it is recommended not to make backups and/or snapshots of the Virtual Machine with the StarWind VSAN service installed, as this could pause the StarWind Virtual Machine. Pausing the Virtual Machines while the StarWind VSAN service in under load may lead to split-brain issues in synchronous replication devices, thus to data corruption.

Configuring Starwind Vms Startup/shutdown

1. Setup the VMs startup policy on ESXi hosts from the Manage -> System tab in the ESXi web console. In the popup window, select Yes to enable the option and set the stop action to Shut down. Click on Save to proceed.

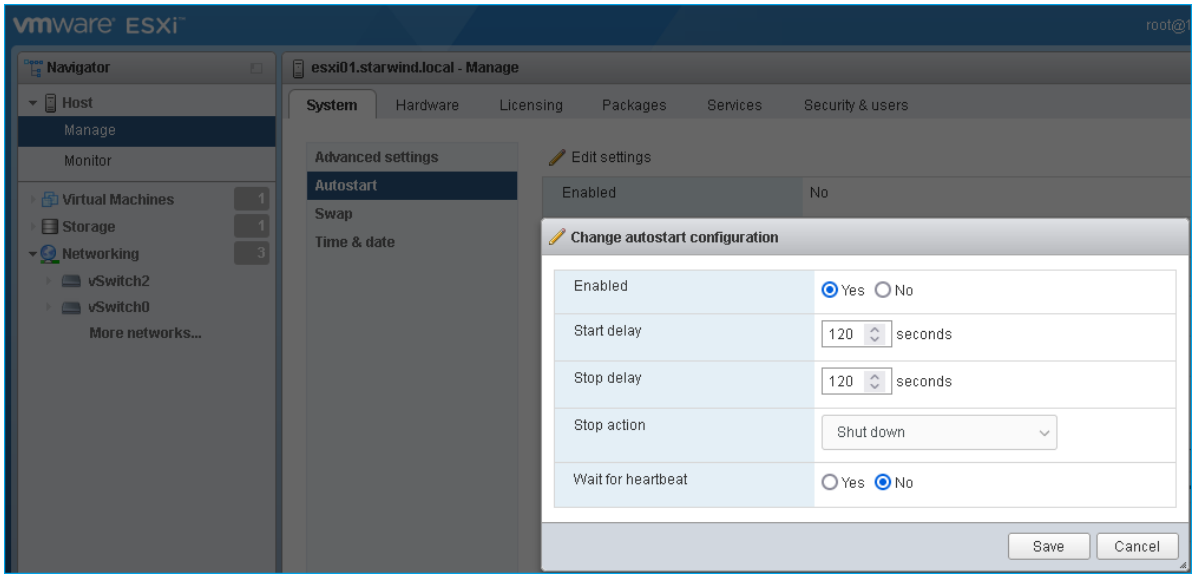

2. To configure VM autostart, right-click on it, navigate to Autostart and click on Enable.

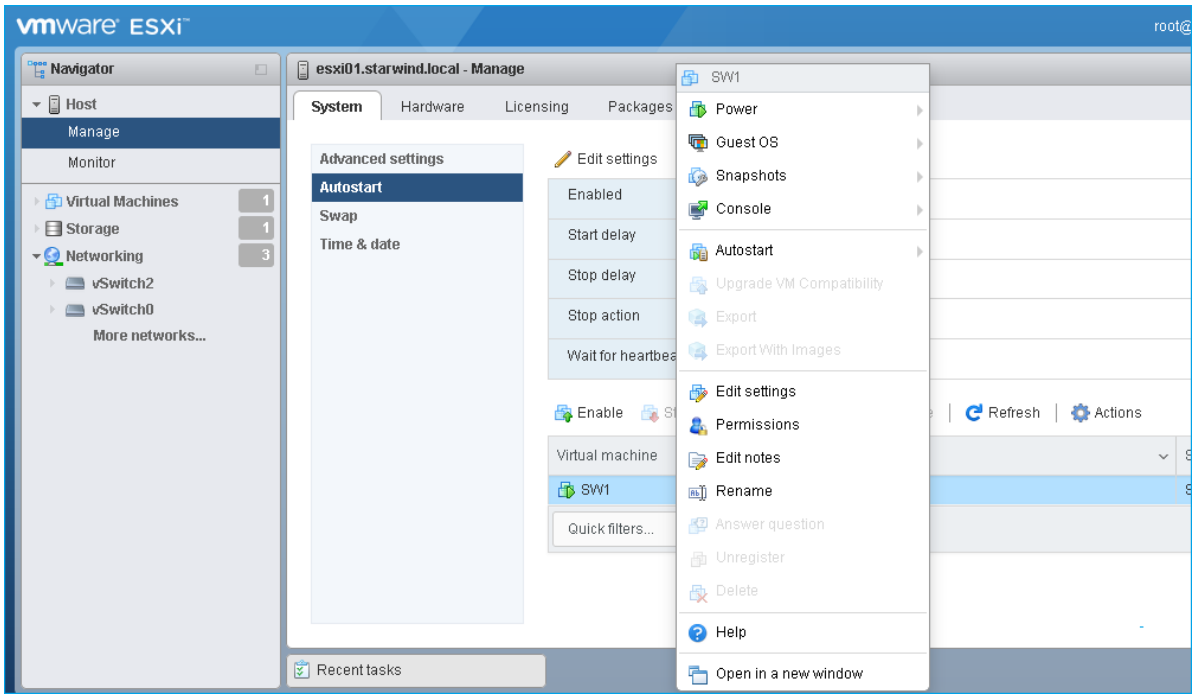

3. Complete the actions above on the StarWind VM located on another host.

4. Start virtual machines, install Windows Server OS and StarWind Virtual SAN.

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.

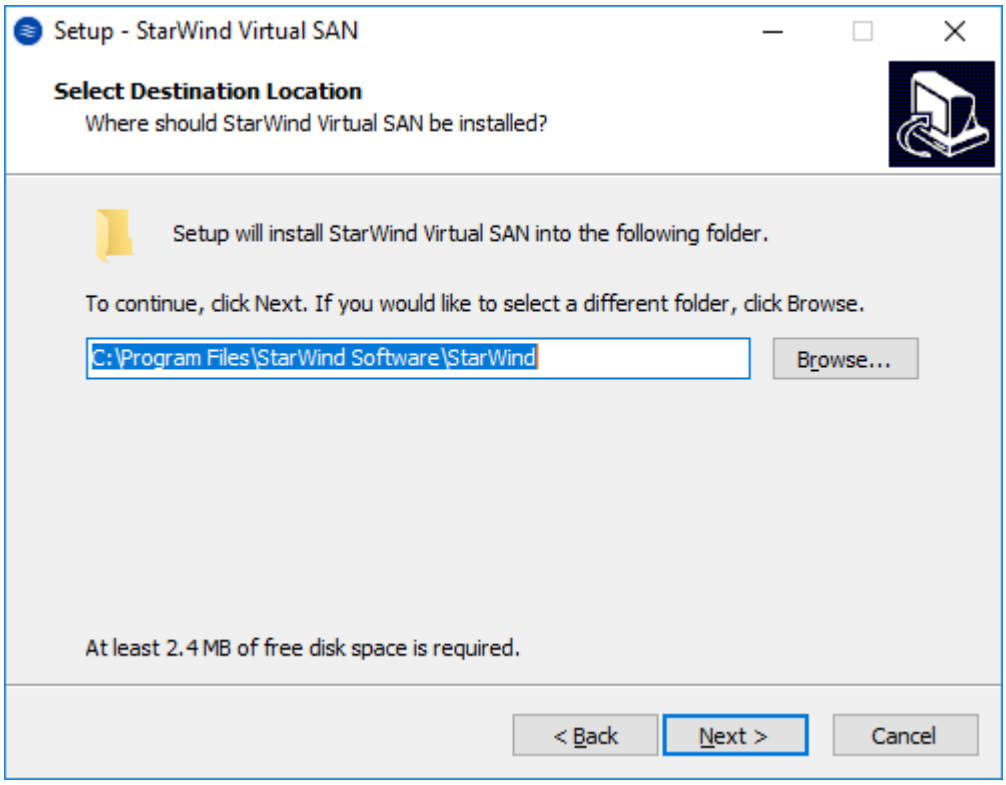

5. Select the following components for the minimum setup:

- 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.

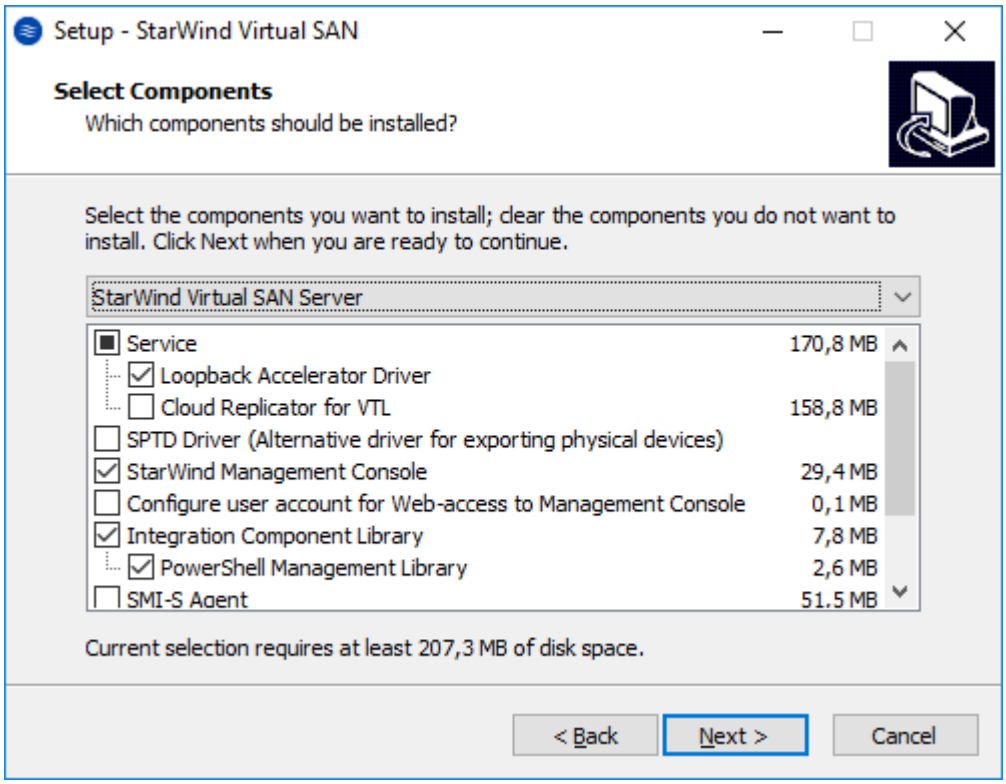

6. Specify Start Menu Folder.

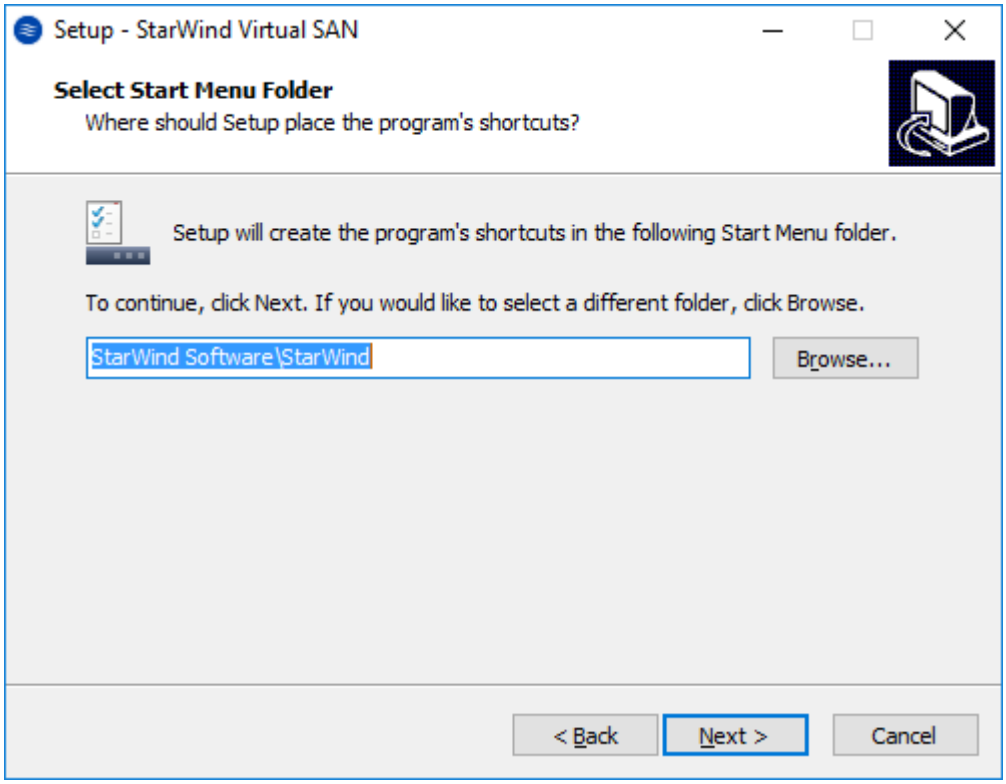

7. Enable the checkbox if a desktop icon needs to be created. Click on Next to continue.

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.

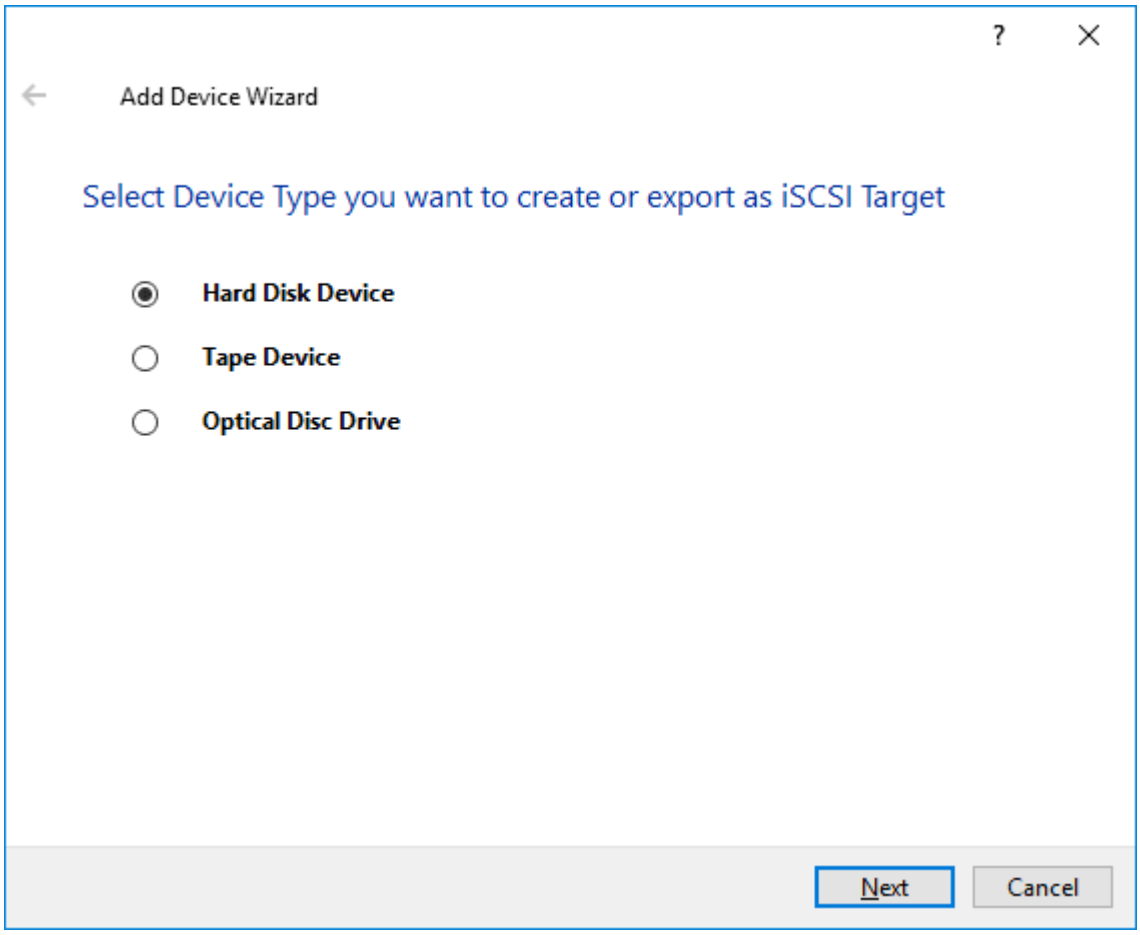

3. Select Virtual Disk.

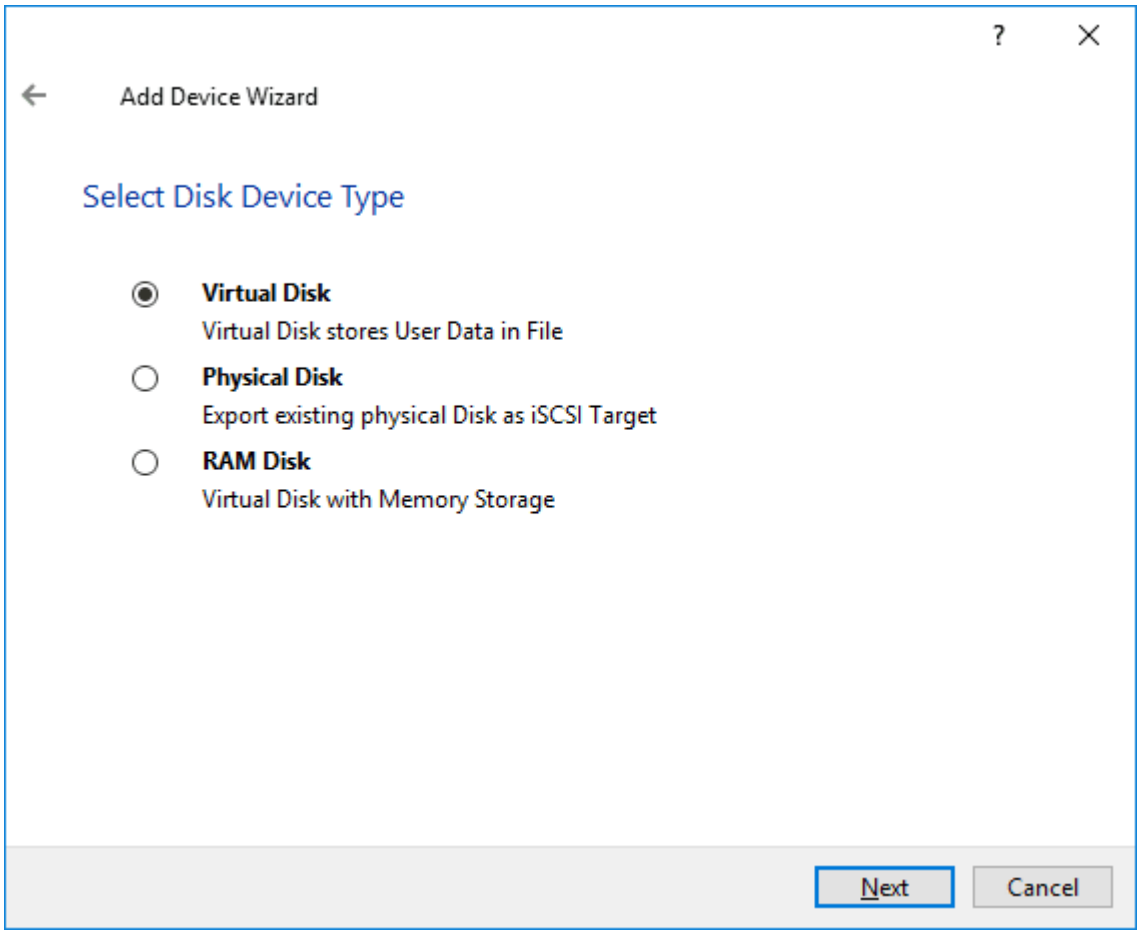

4. Specify a virtual disk Name, Location, and Size.

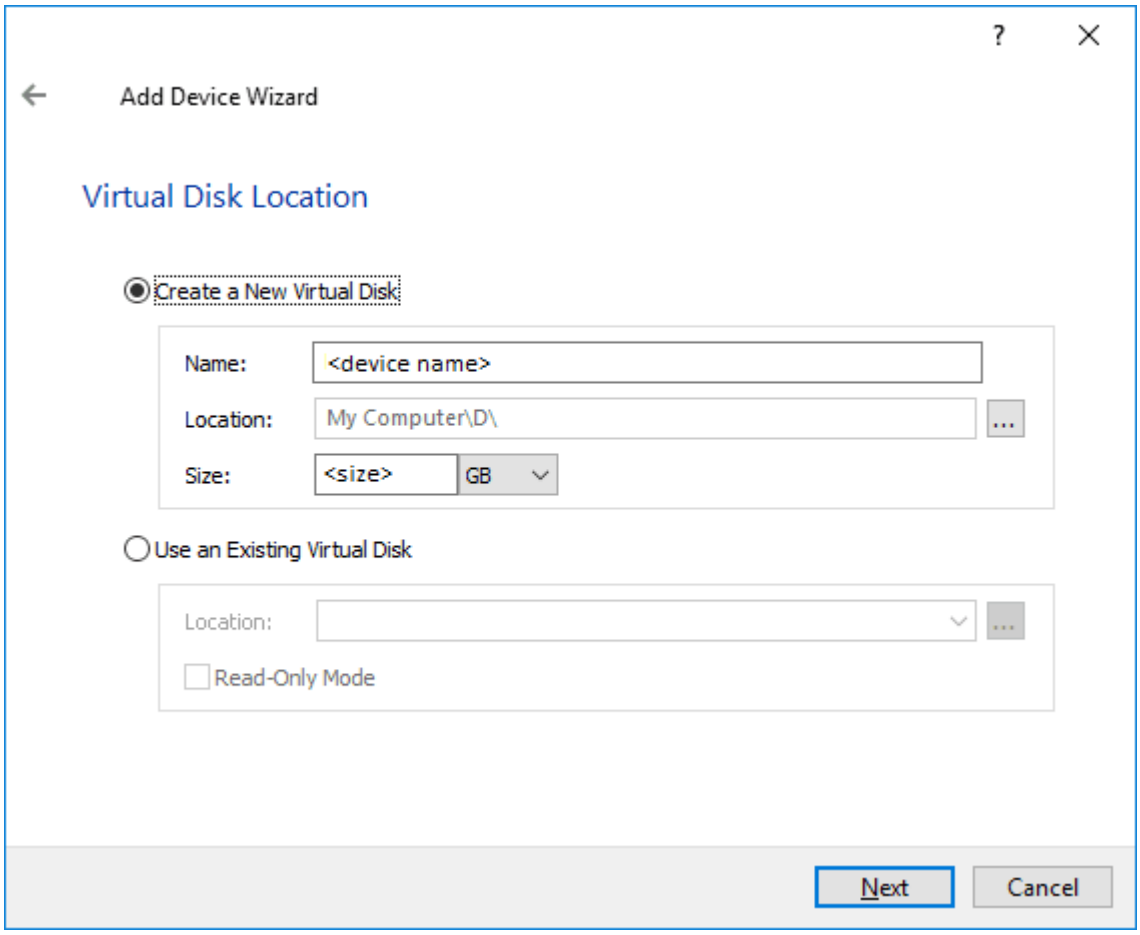

5. Select the Thick provisioned disk type and block size.
NOTE: Use 4096 sector size for targets, connected on Windows-based systems and 512 bytes sector size for targets, connected on Linux-based systems (ESXi/Xen/KVM).

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.

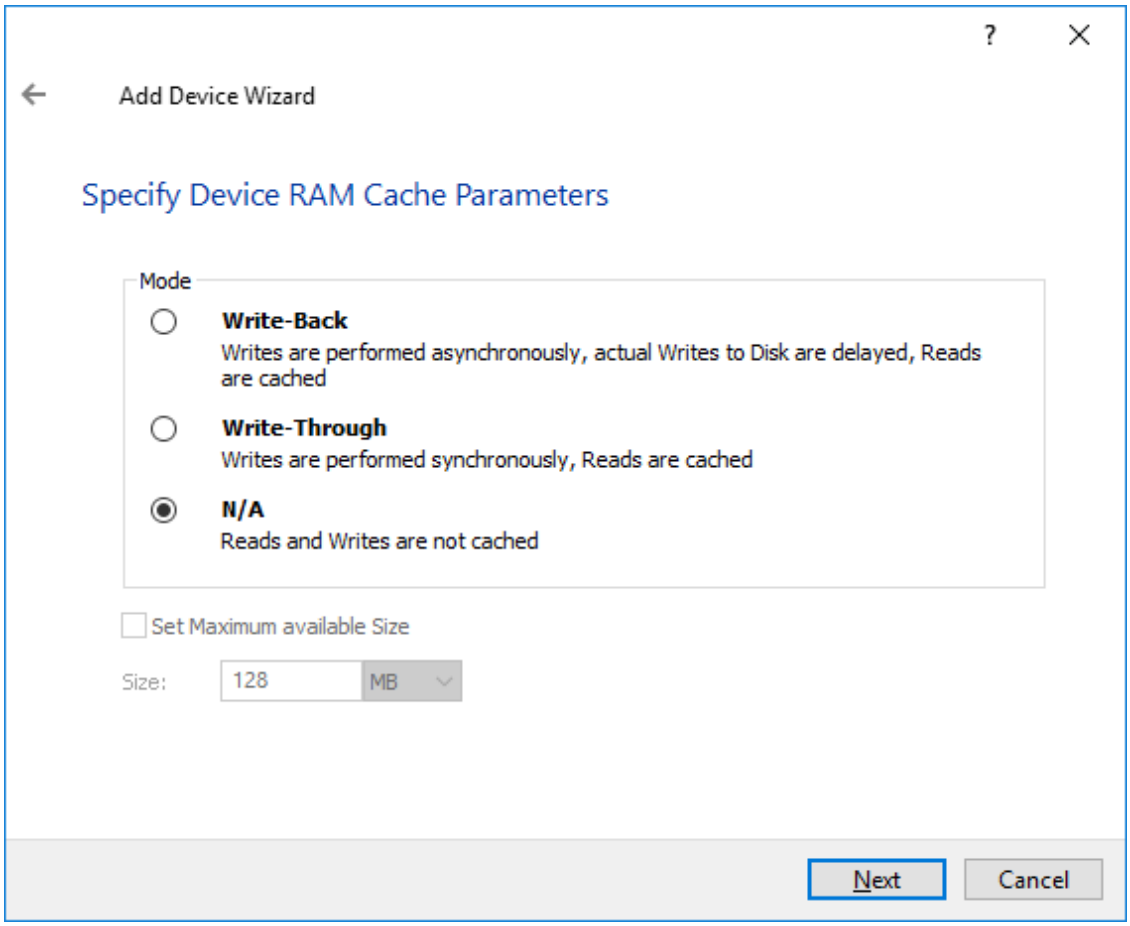

7. Specify Target Parameters. Select the Target Name checkbox to enter a custom target name. Otherwise, the name is generated automatically in accordance with the specified target alias.

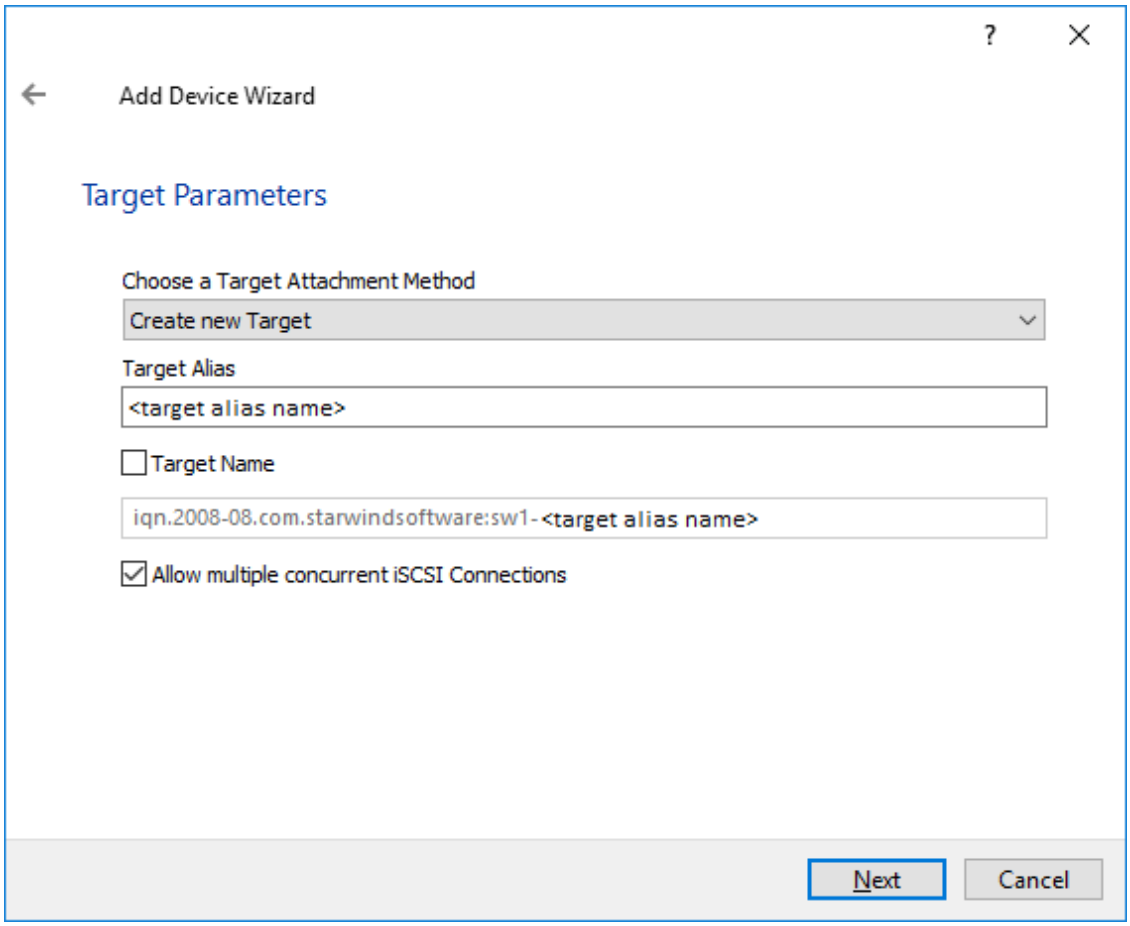

8. Click Create to add a new device and attach it to the target.

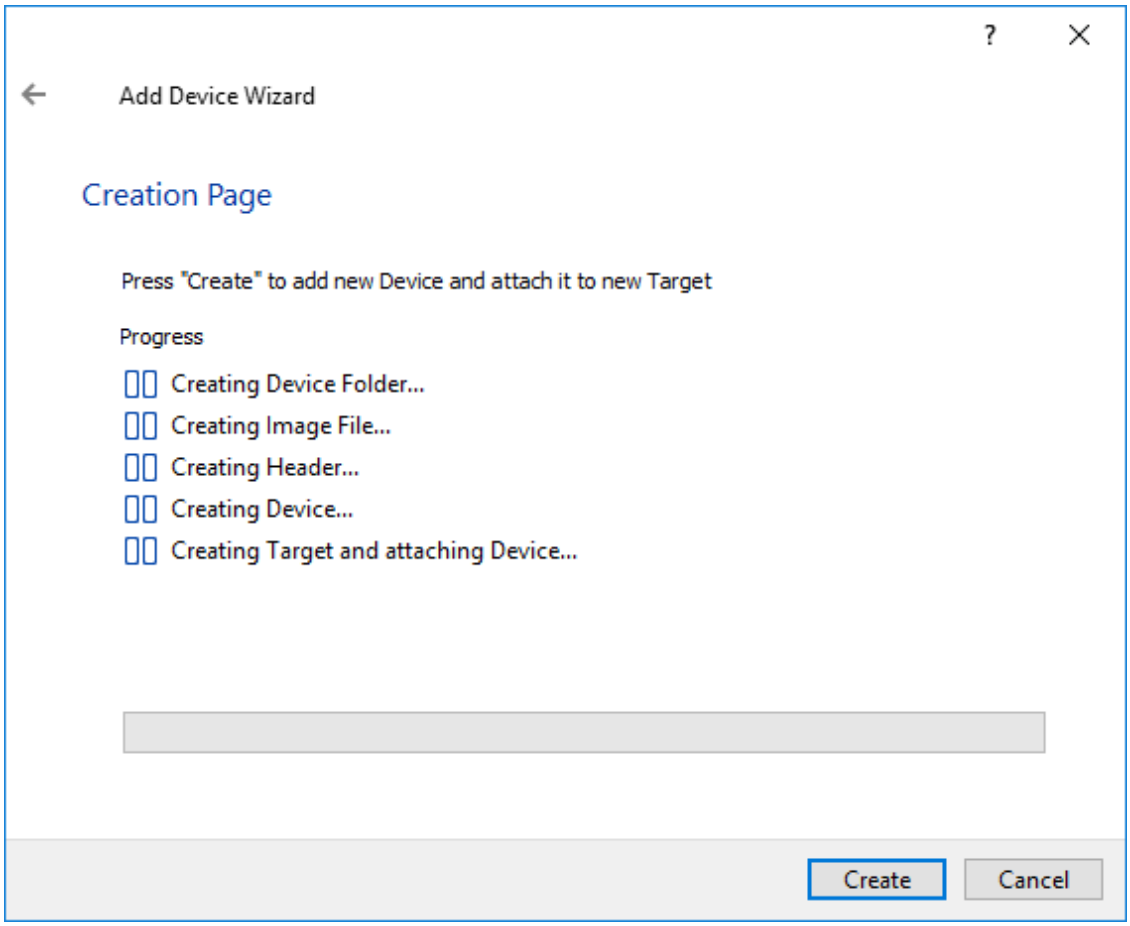

9. Click Close to finish the device creation.

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.

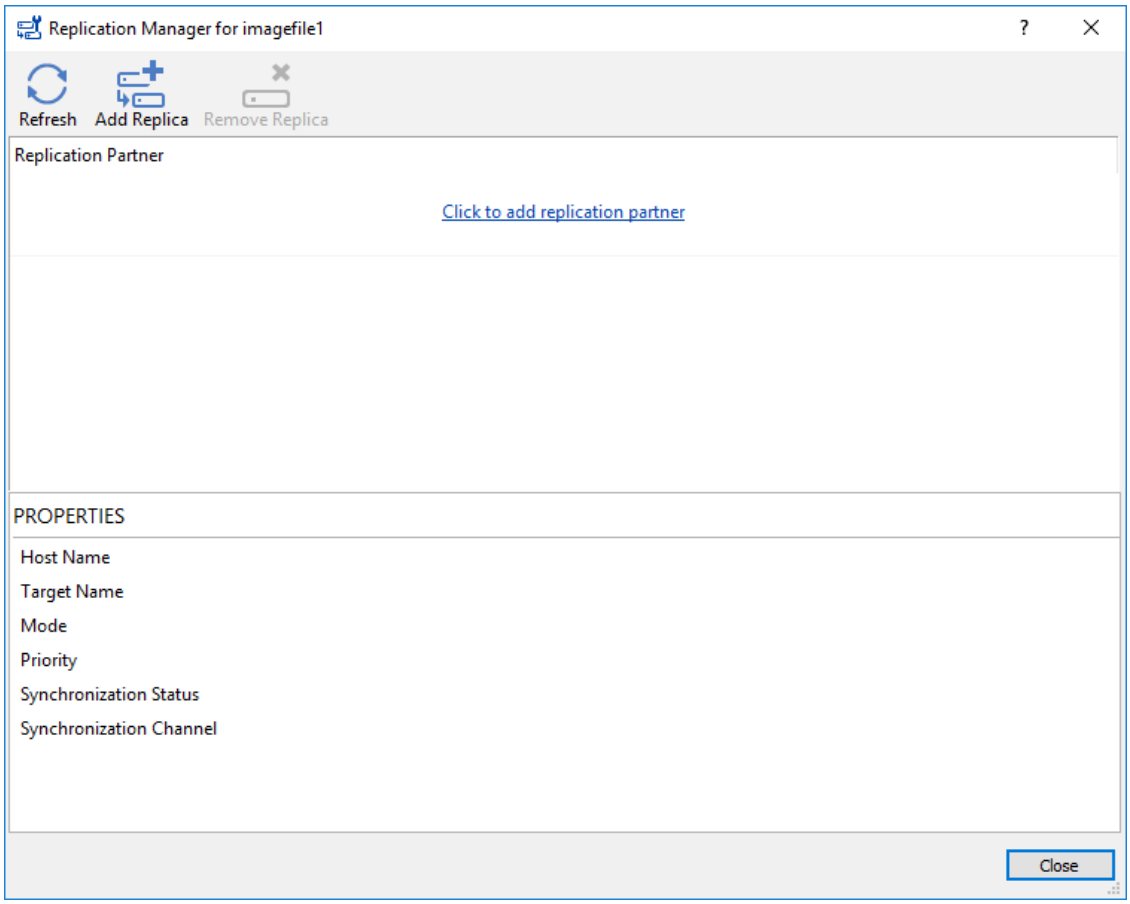

3. Select Synchronous “Two-Way” replication as a replication mode.

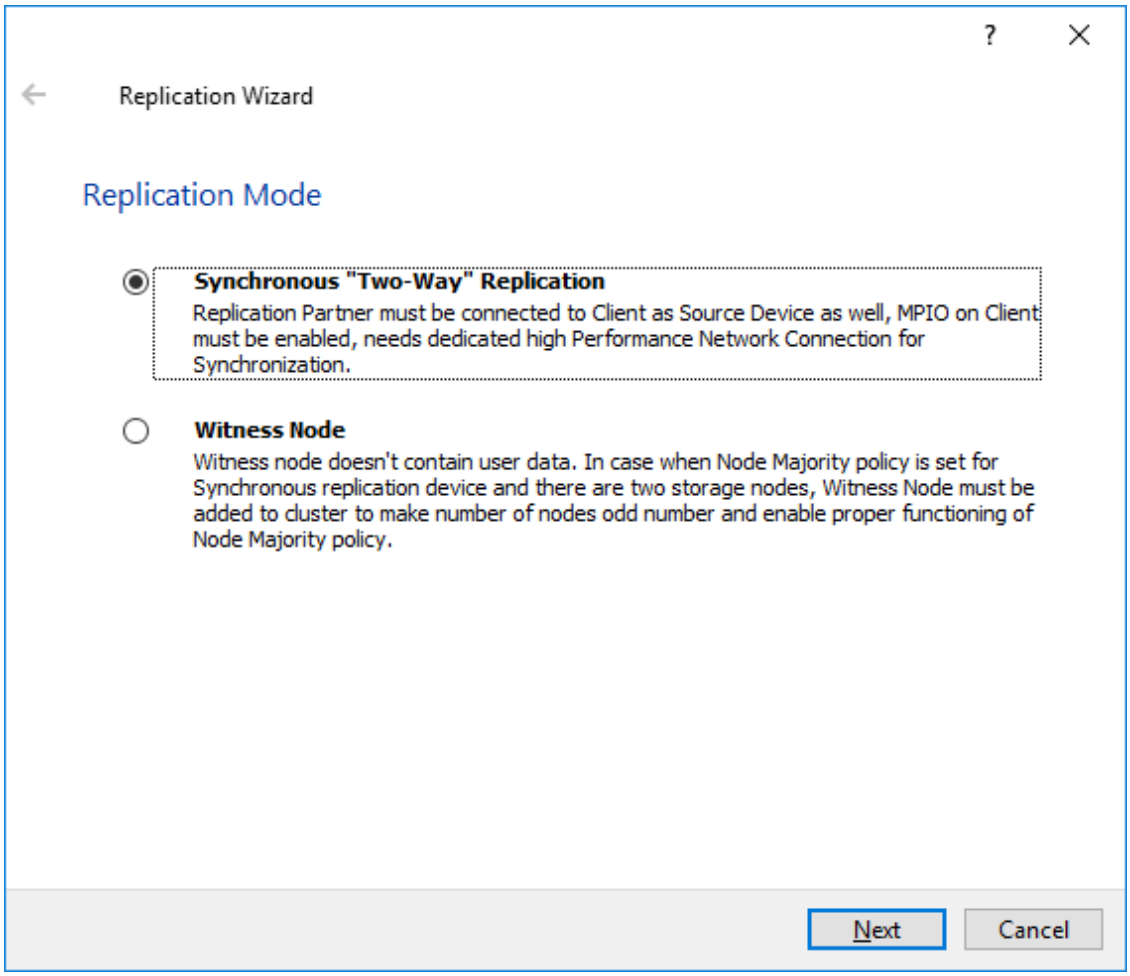

4. Specify a partner Host name or IP address and Port Number.

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.

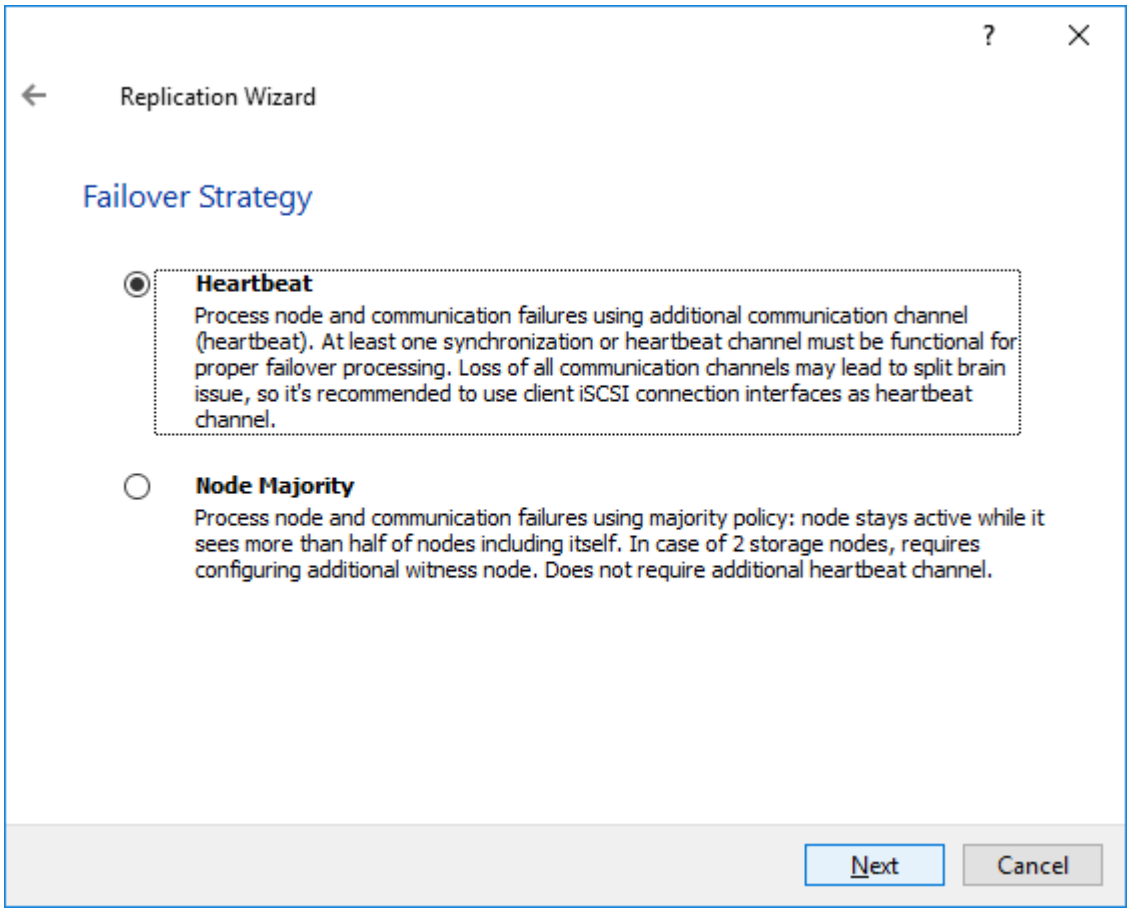

2. Select Create new Partner Device and click Next.

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.

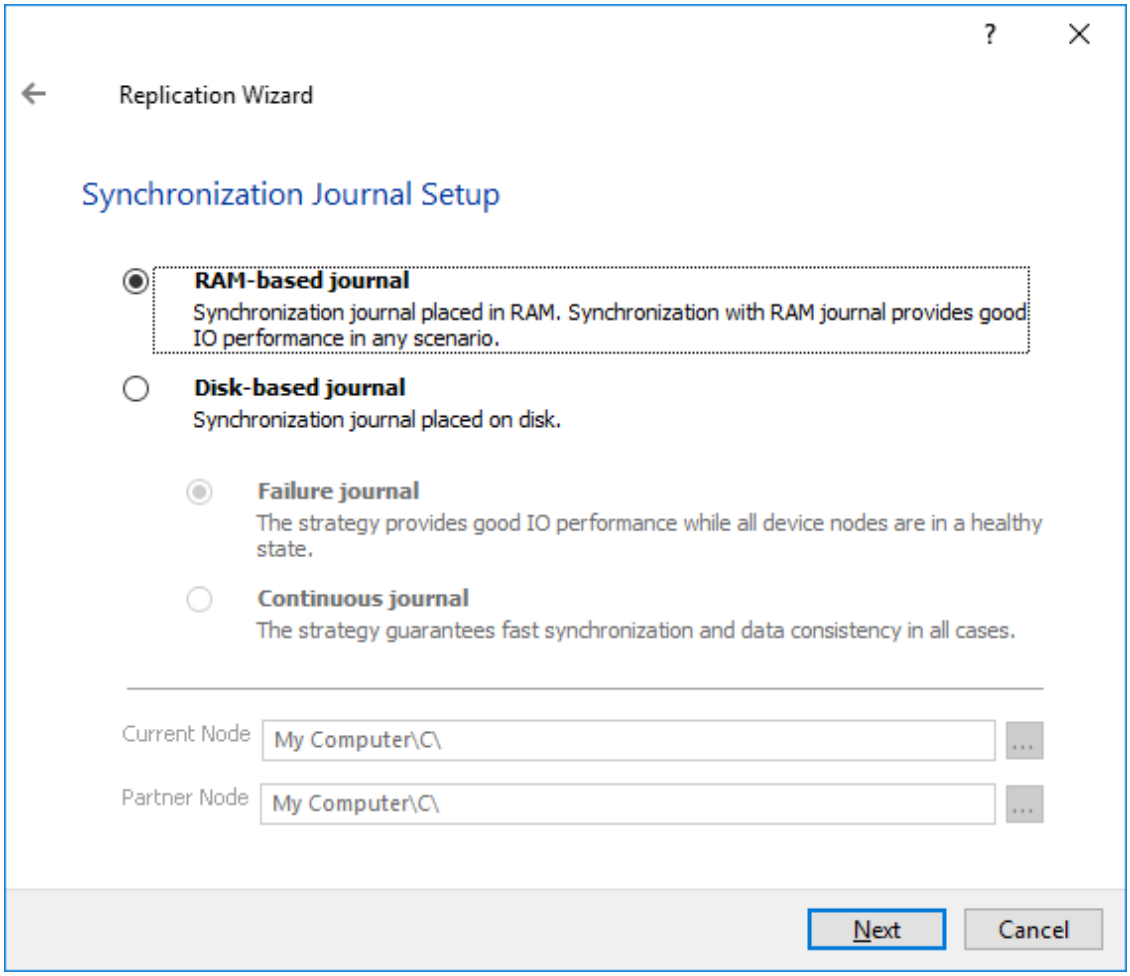

5. Click Change Network Settings.

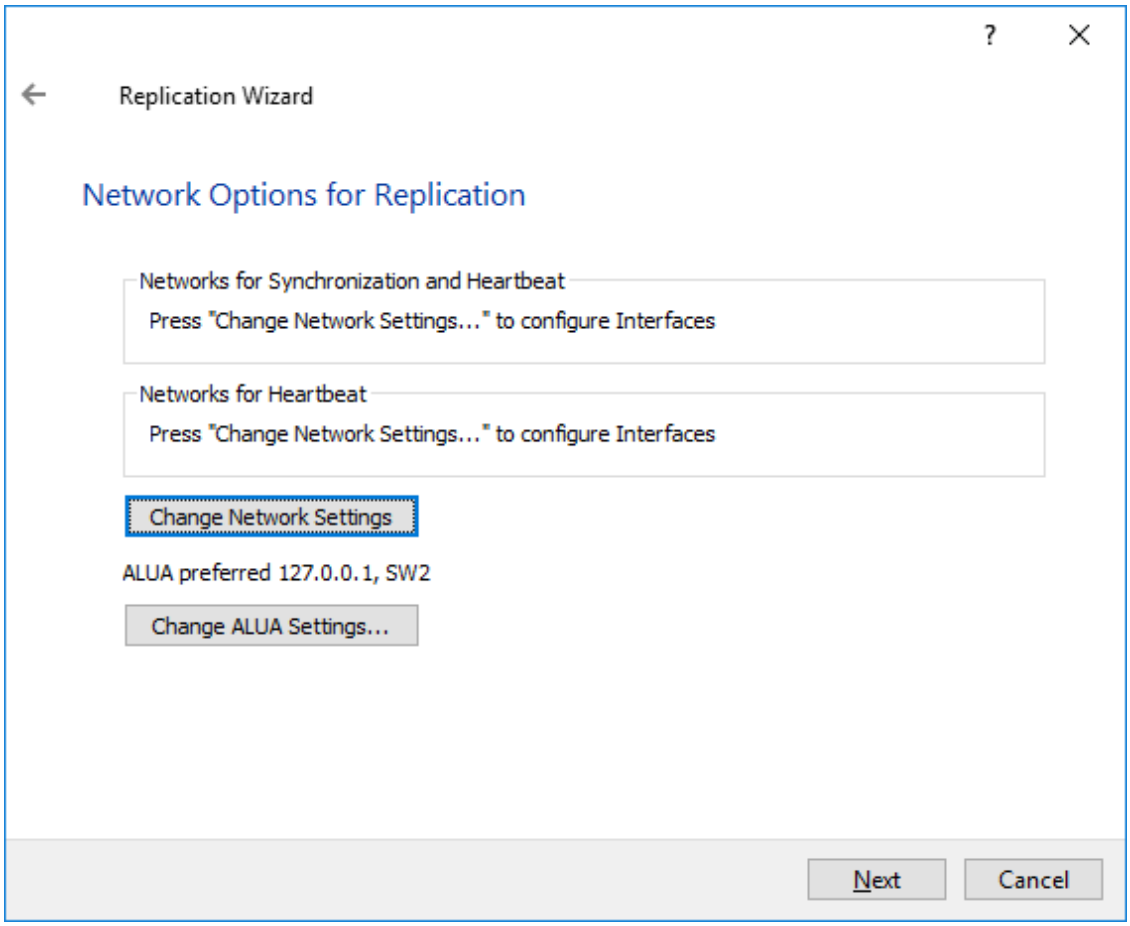

6. Specify the interfaces for Synchronization and Heartbeat Channels. Click OK and then click Next.

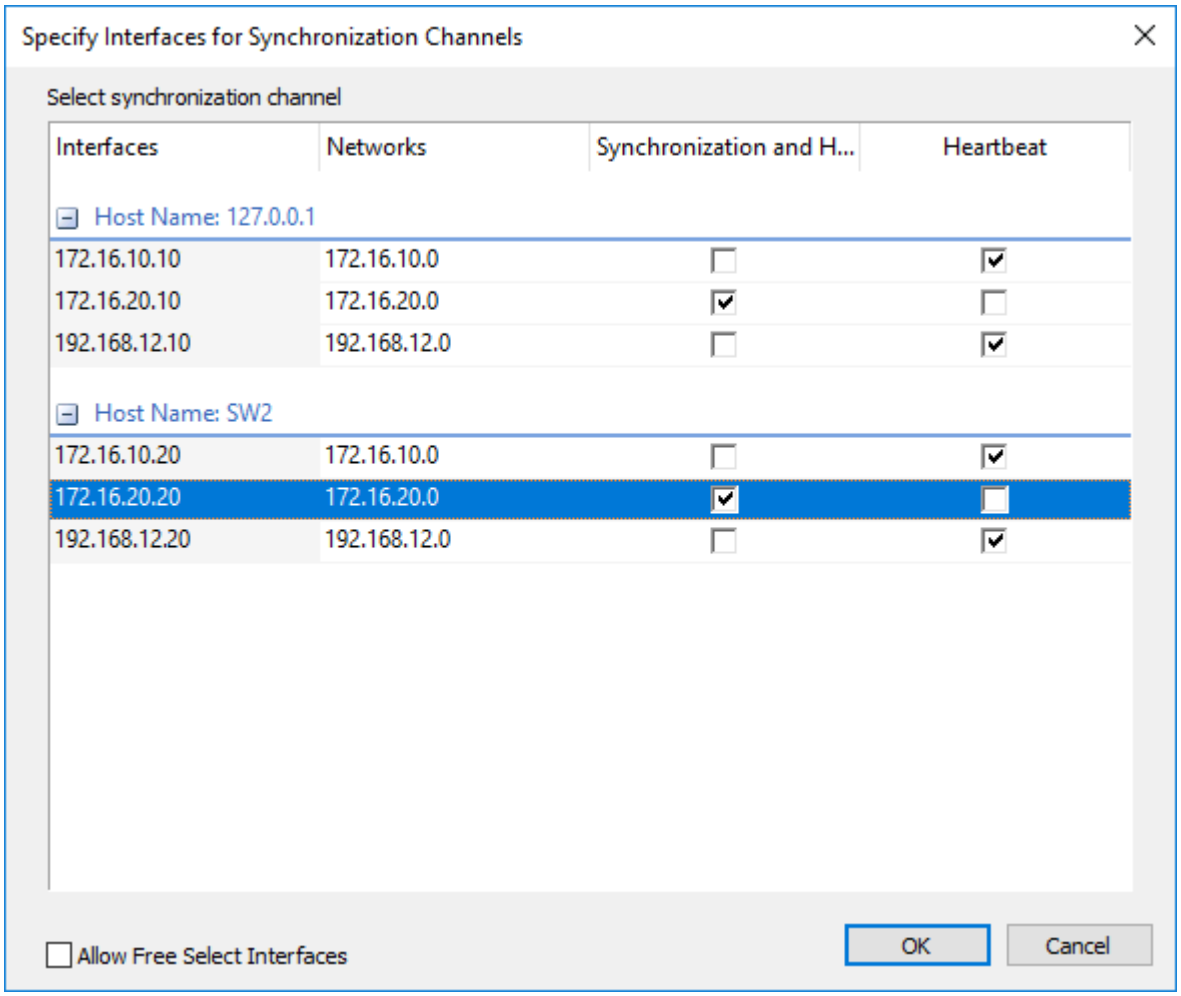

7. In Select Partner Device Initialization Mode, select Synchronize from existing Device and click Next.

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

9. Follow the similar procedure for the creation of other virtual disks that will be used as storage repositories.
NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:

How to extend Image File or High Availability device

Node Majority

There are two ways to configure Witness for 2-nodes StarWind HA device, created with Node Majority Failover Strategy: File Share (SMB) as Witness and additional server as Witness Node.

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

SMB Witness is a file, located on SMB share, which can be accessed by both nodes and help them to eliminate the split-brain issue in case of synchronization connection interruption between the nodes. To set up the SMB file share as a Witness for 2-nodes HA device with Node Majority Failover Strategy, perform the actions, described on this page:
https://www.starwindsoftware.com/help/ConfiguringFileShareSMBasWitness.html

– Creating HA device with Witness Node:

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

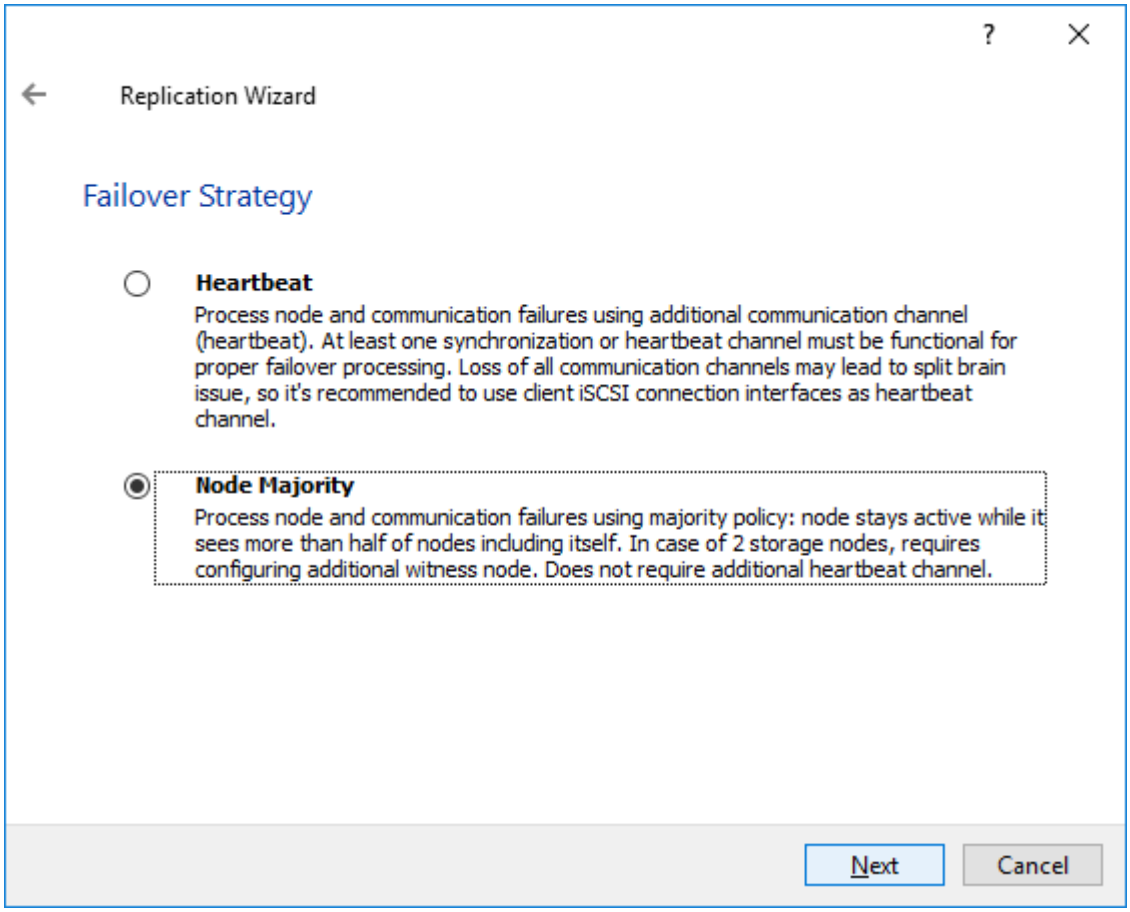

2. Select Create new Partner Device. Click Next.

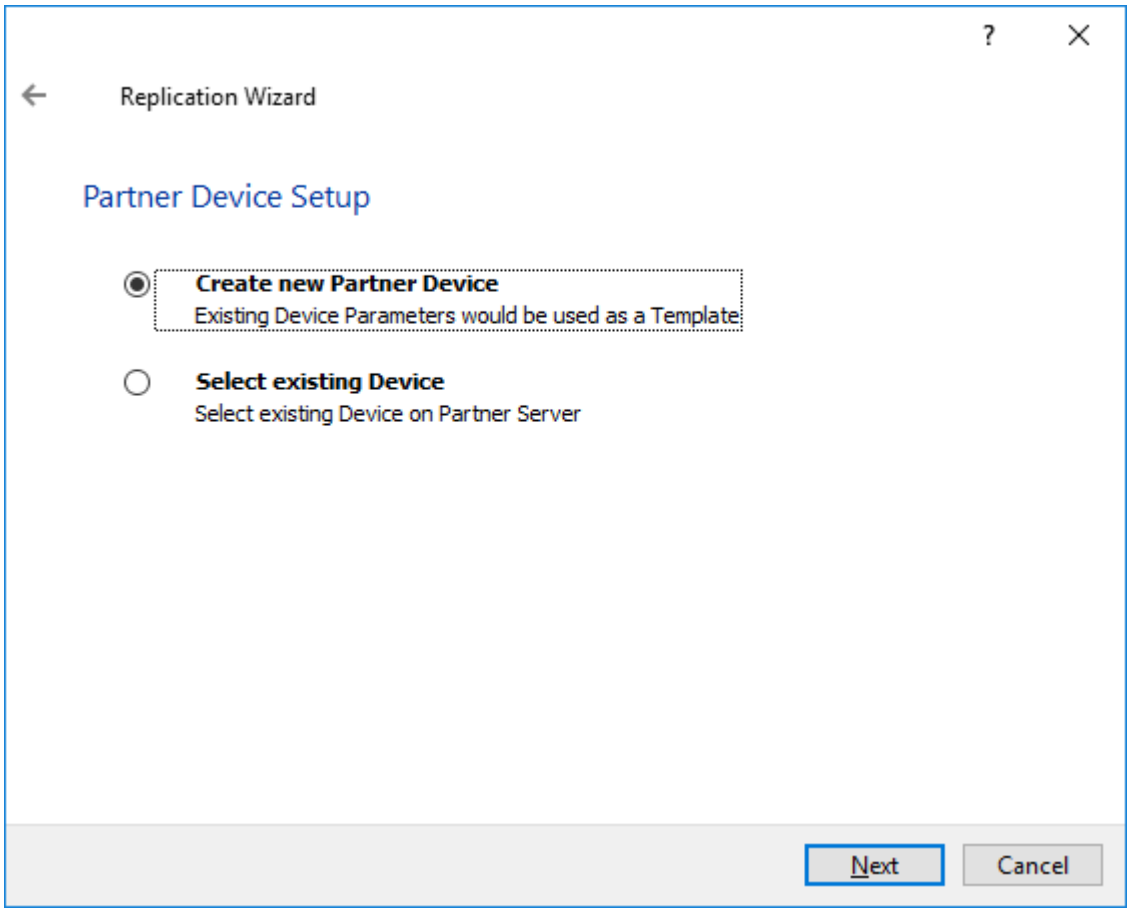

3. Specify the partner device location if necessary and/or modify the target name of the device. Click Next.

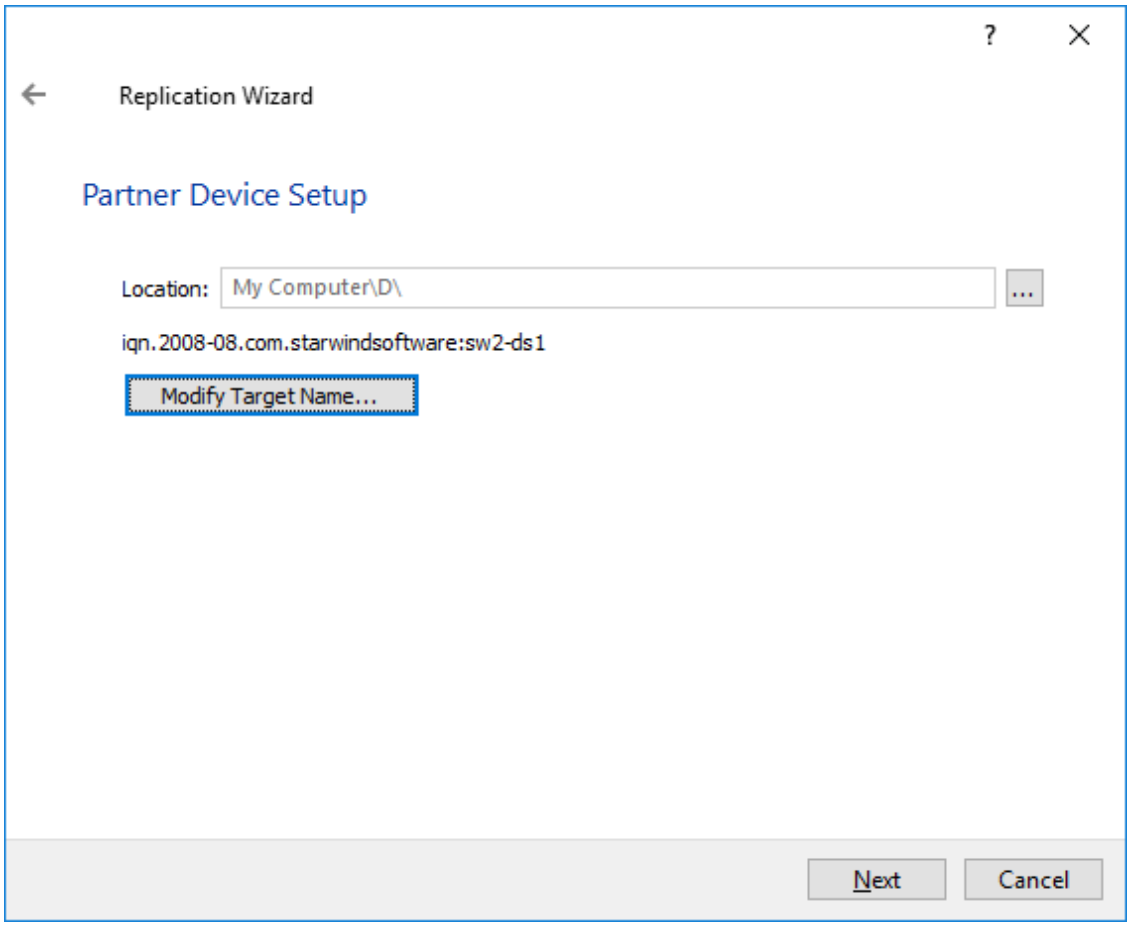

4. Select Synchronization Journal Strategy and click Next.
NOTE: There are several options – RAM-based journal (default) and Disk-based journal with failure and continuous strategy, that allow to avoid full synchronization cases.

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

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

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

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

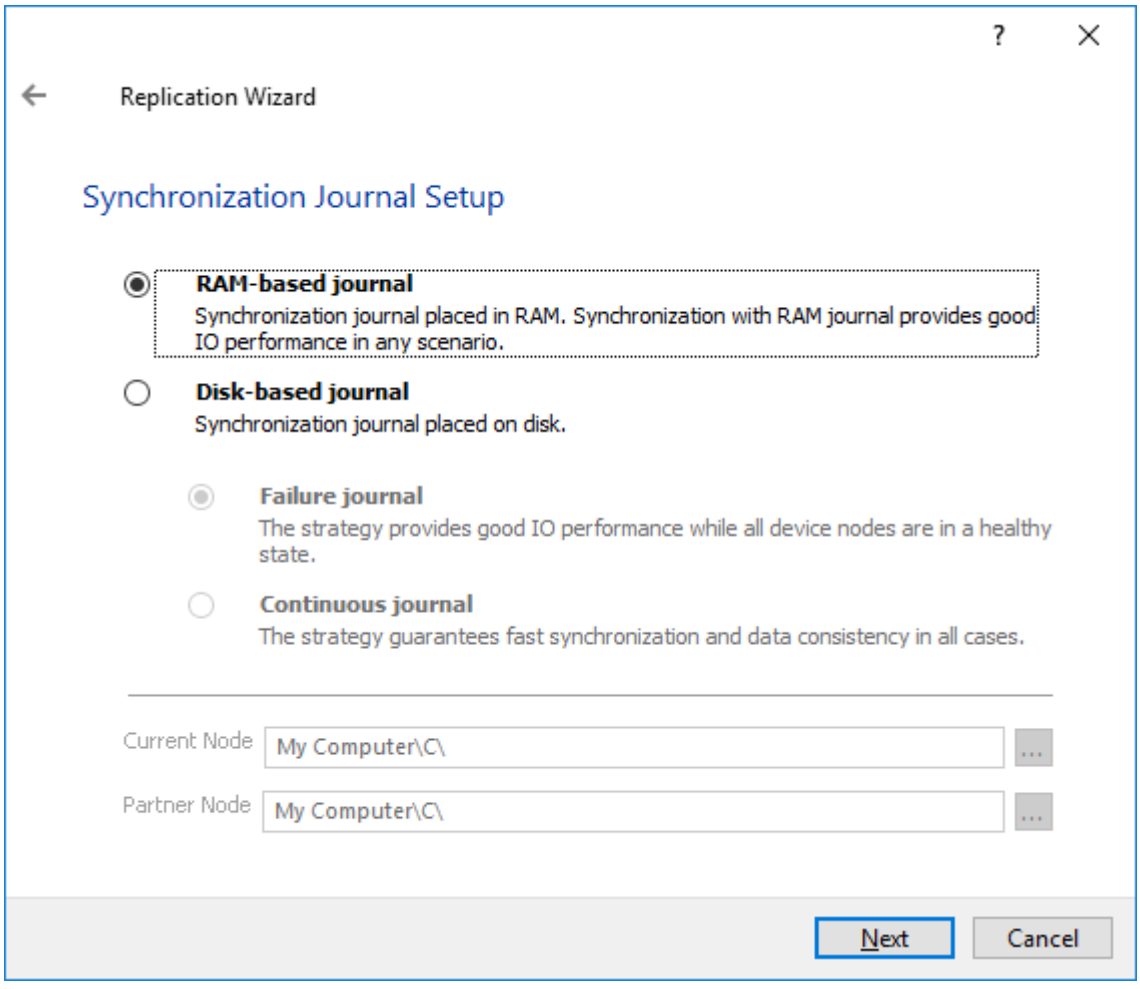

5. Select the Synchronization and Heartbeat networks for the HA device by clicking Change Network Settings.

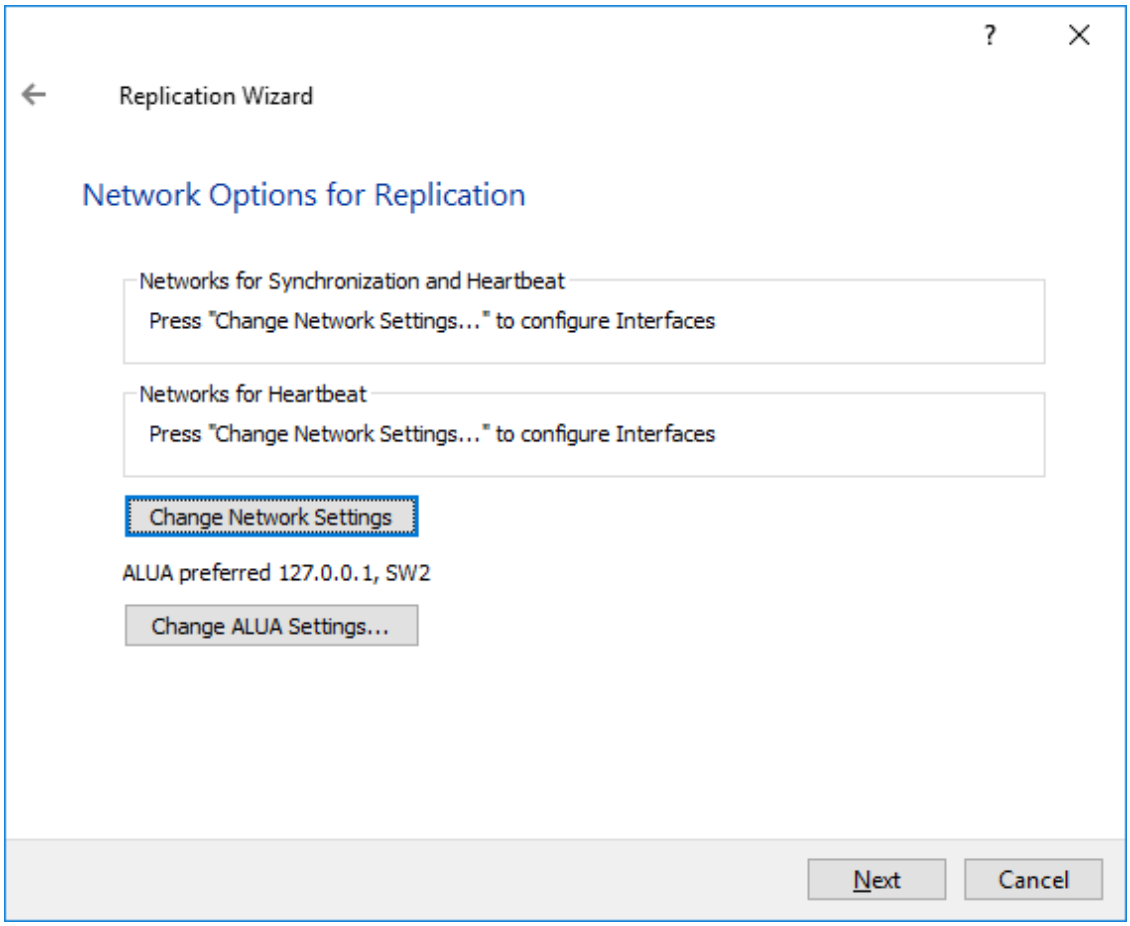

6. Specify the interfaces for Synchronization and Heartbeat. Press OK. Then, click Next.

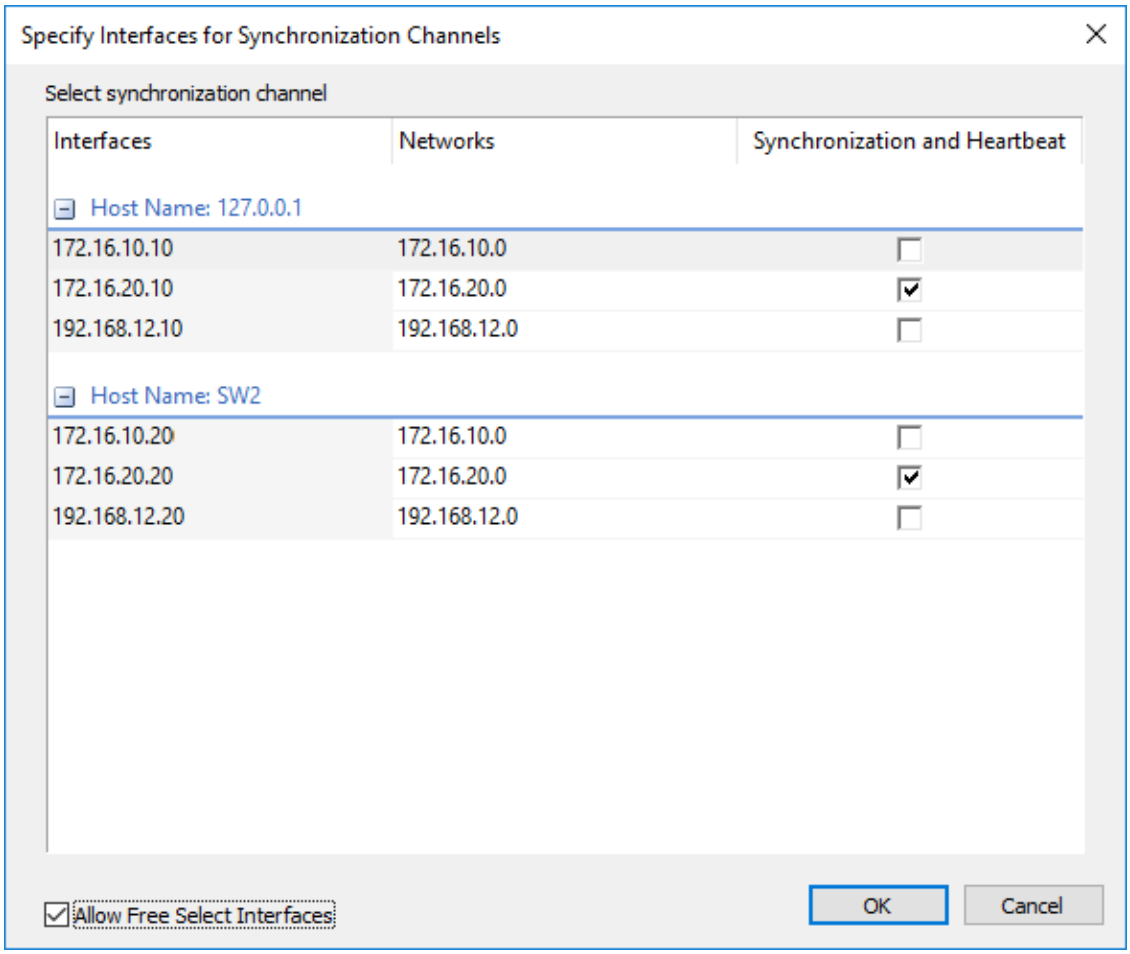

7. Select Synchronize from existing Device for the partner device initialization mode. Click Next.

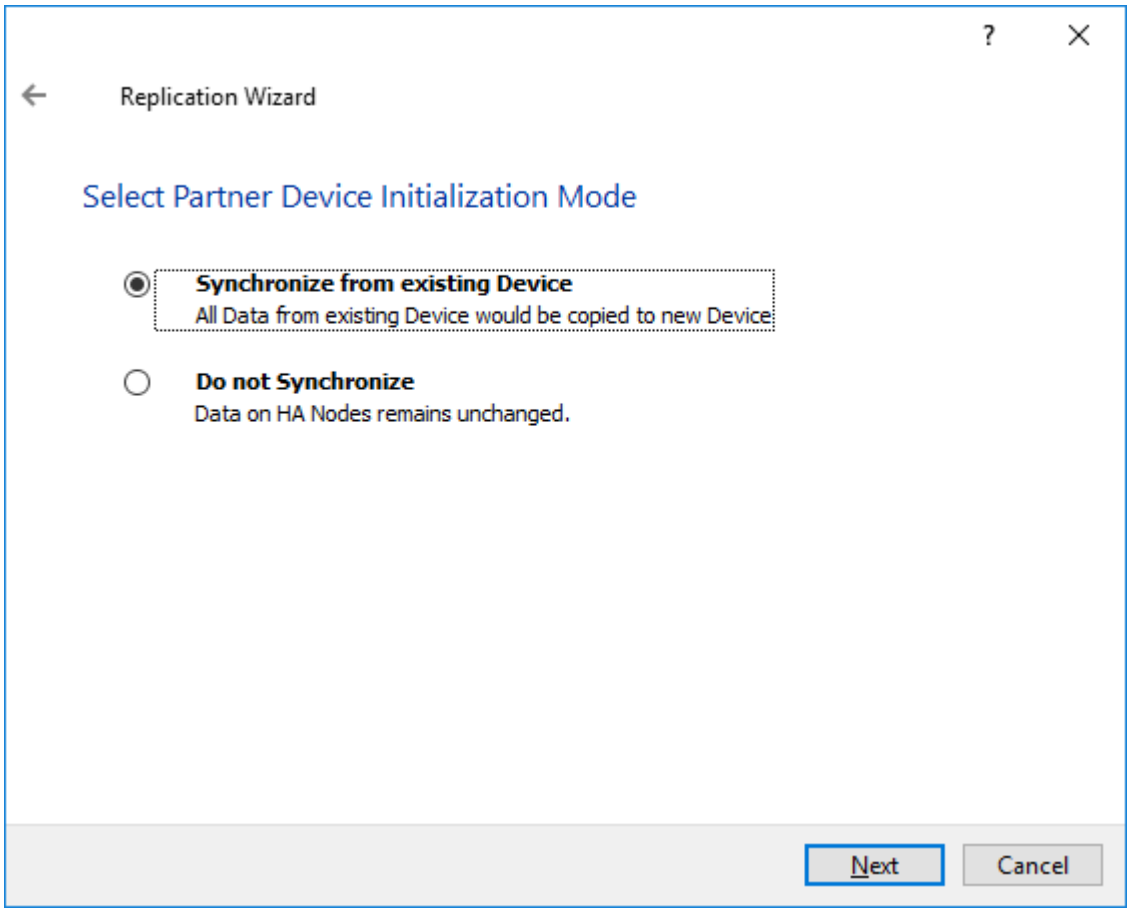

8. Press the Create Replica button. Then click Close.

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

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

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

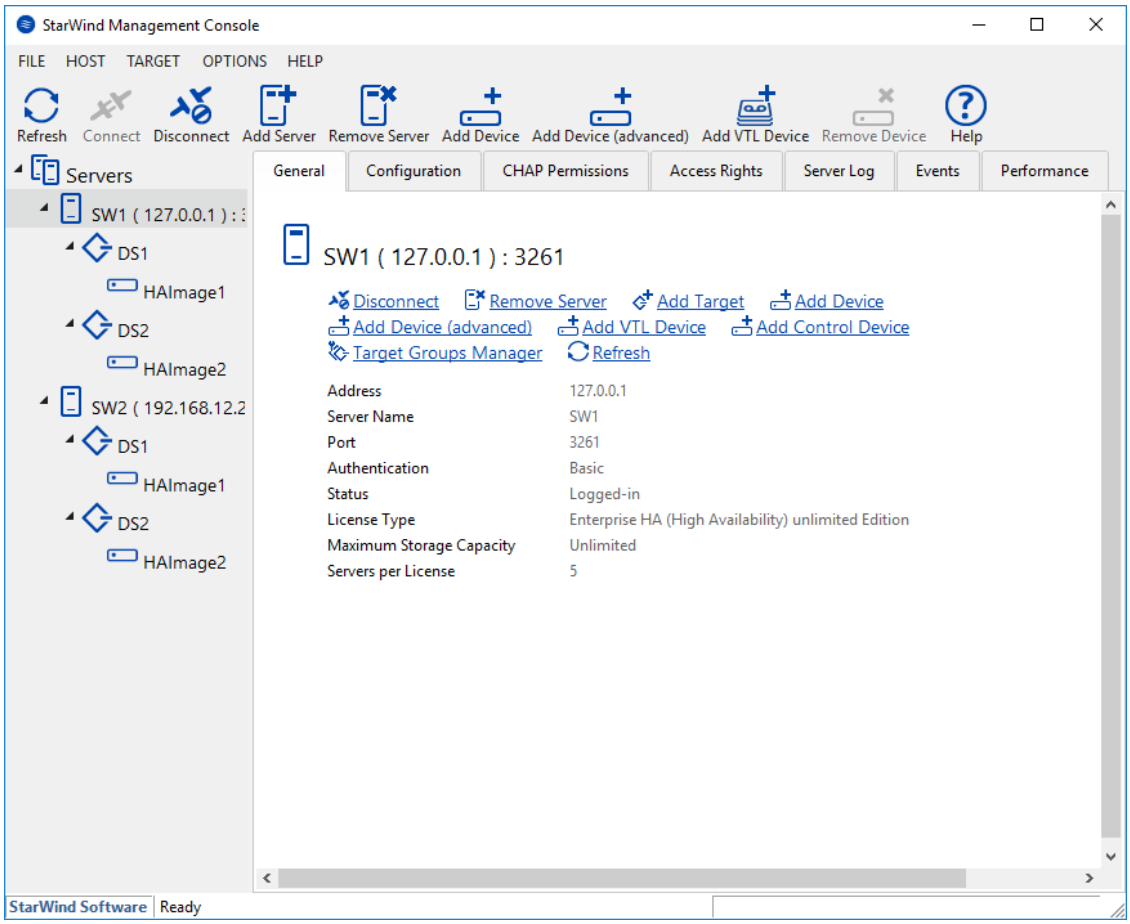

Adding Witness Node

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

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

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

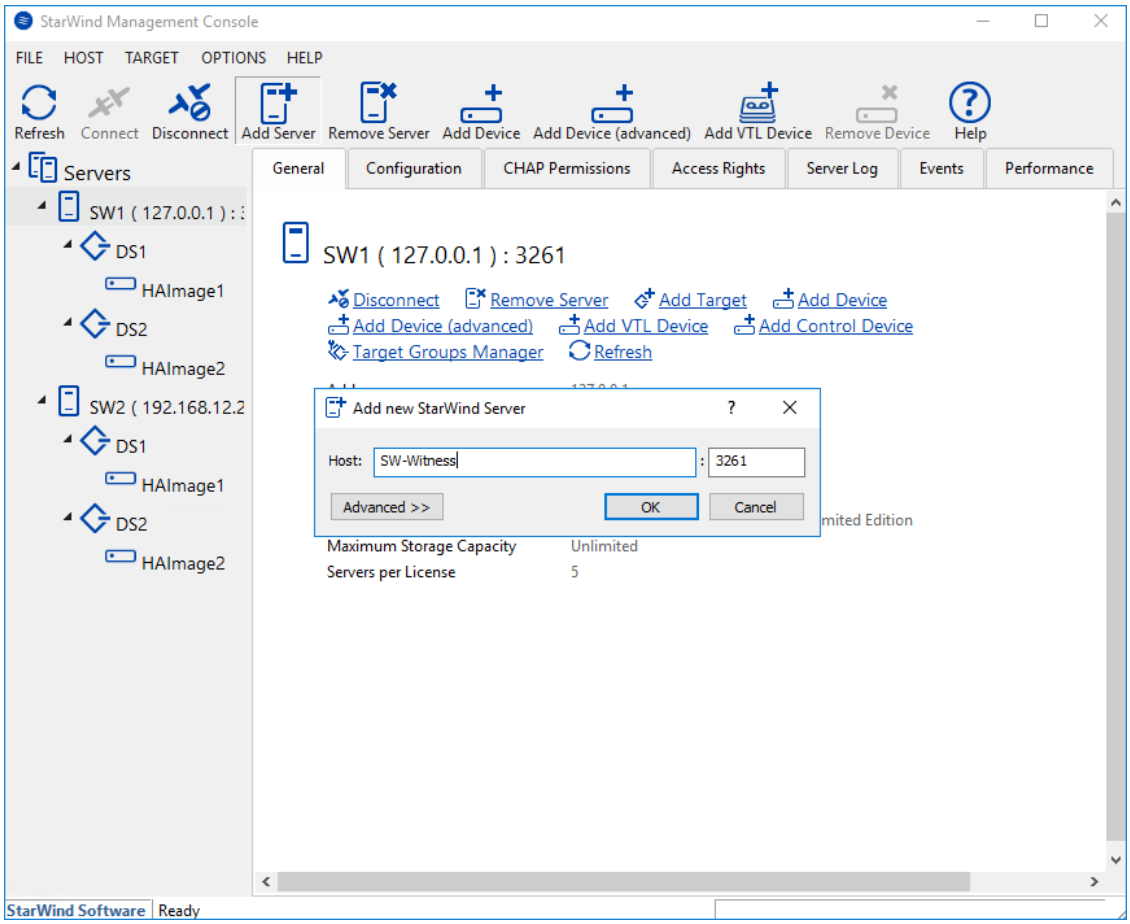

2. Right-click the HA device with the configured Node Majority failover policy and select Replication Manager. The Replication Manager window will appear. Press the Add Replica button.

3. Select Witness Node and click Next

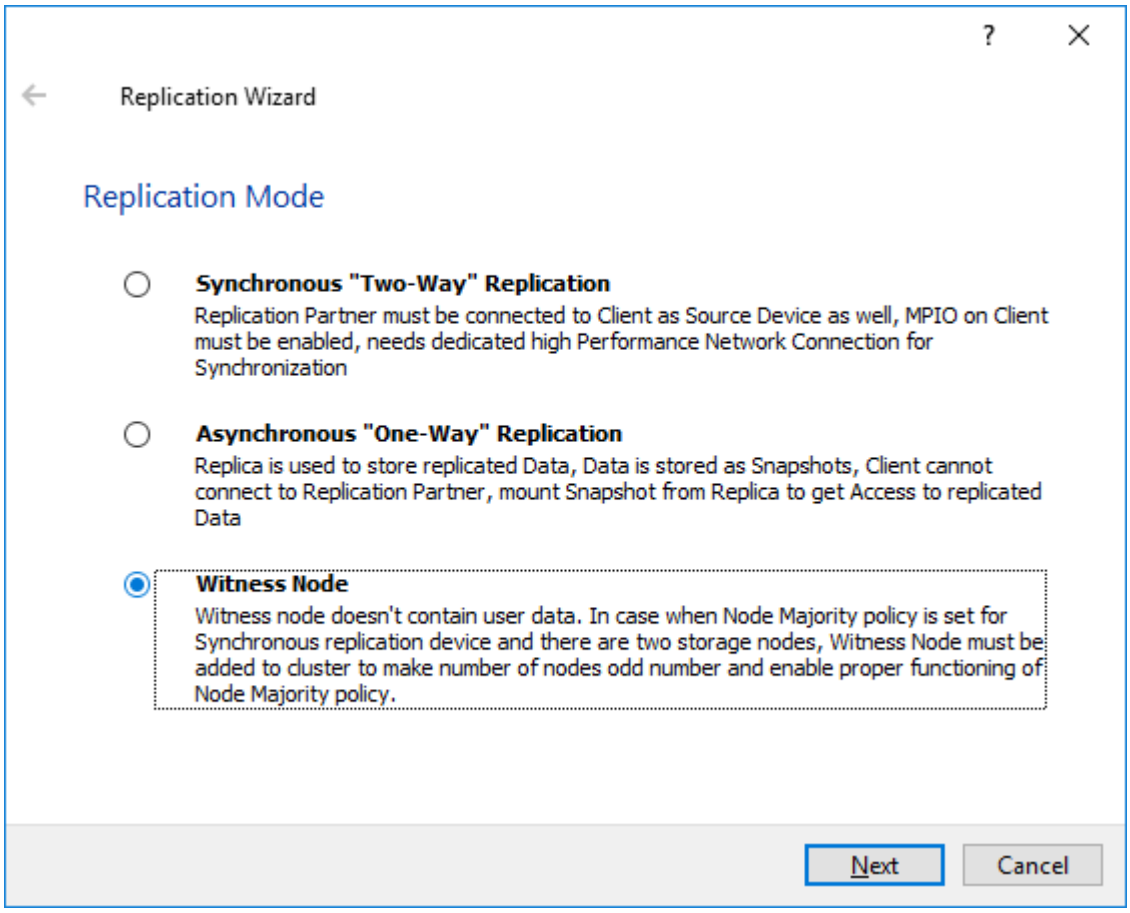

4. Specify the Witness node name or its IP address.

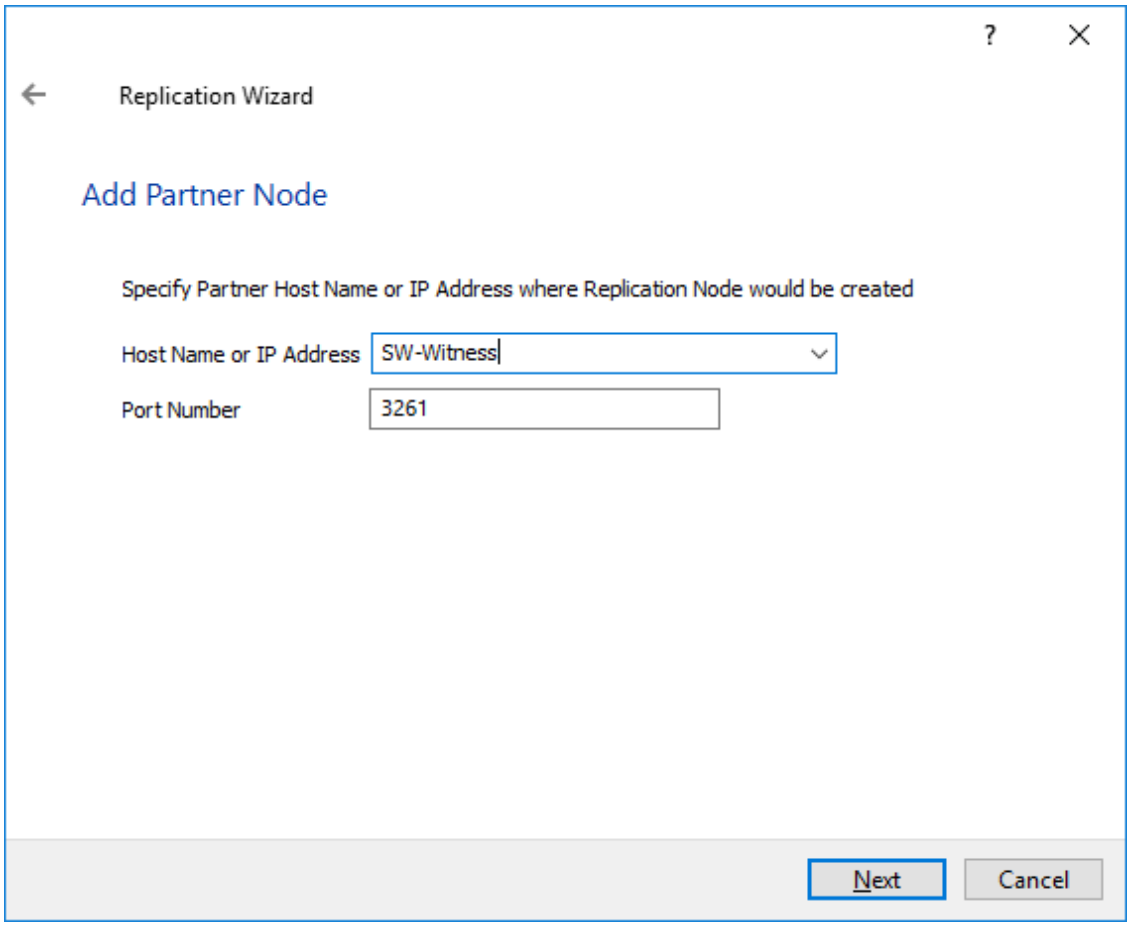

5. Specify the Witness device location and its target name if necessary. Click Next.

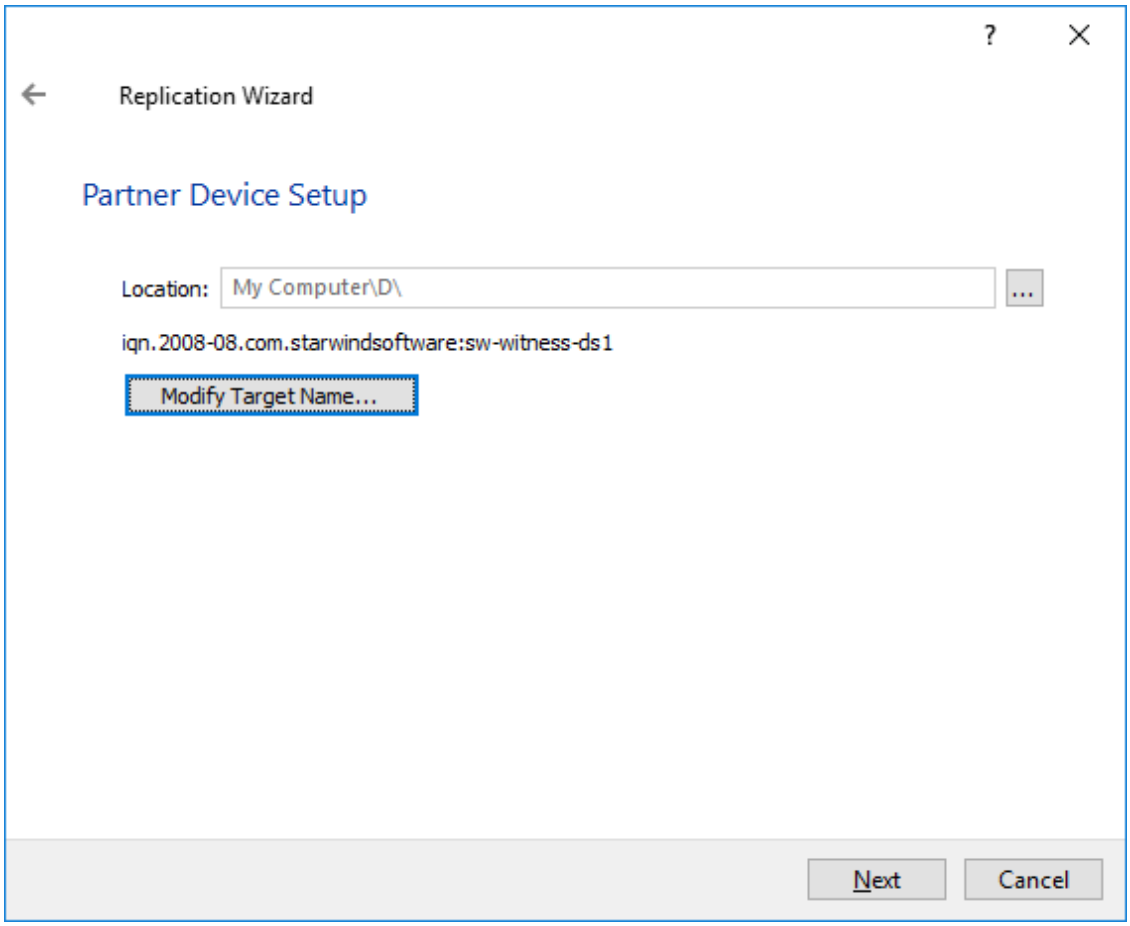

6. For the HA device, select the synchronization channel with the Witness node by clicking on the Change network settings button.

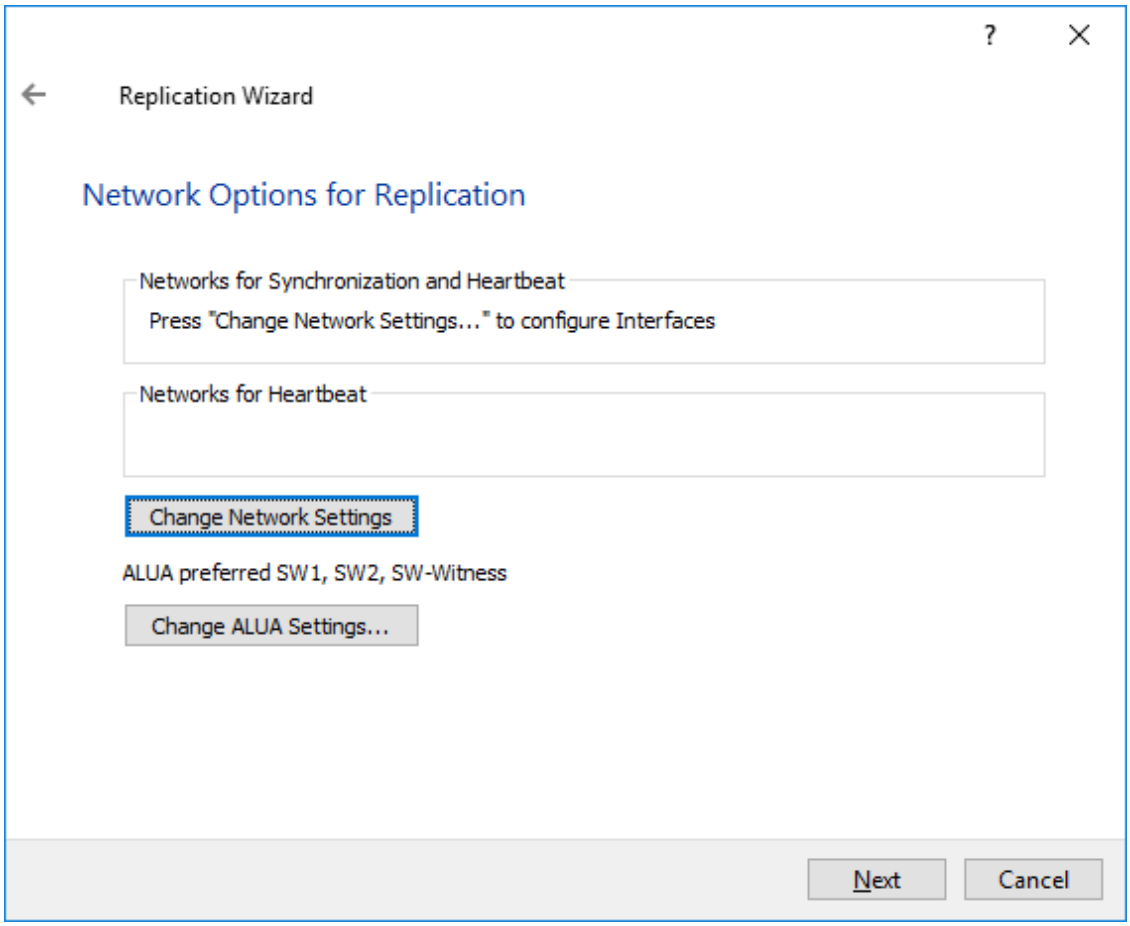

7. Specify Interfaces for Synchronization Channels, confirm, and click Next.

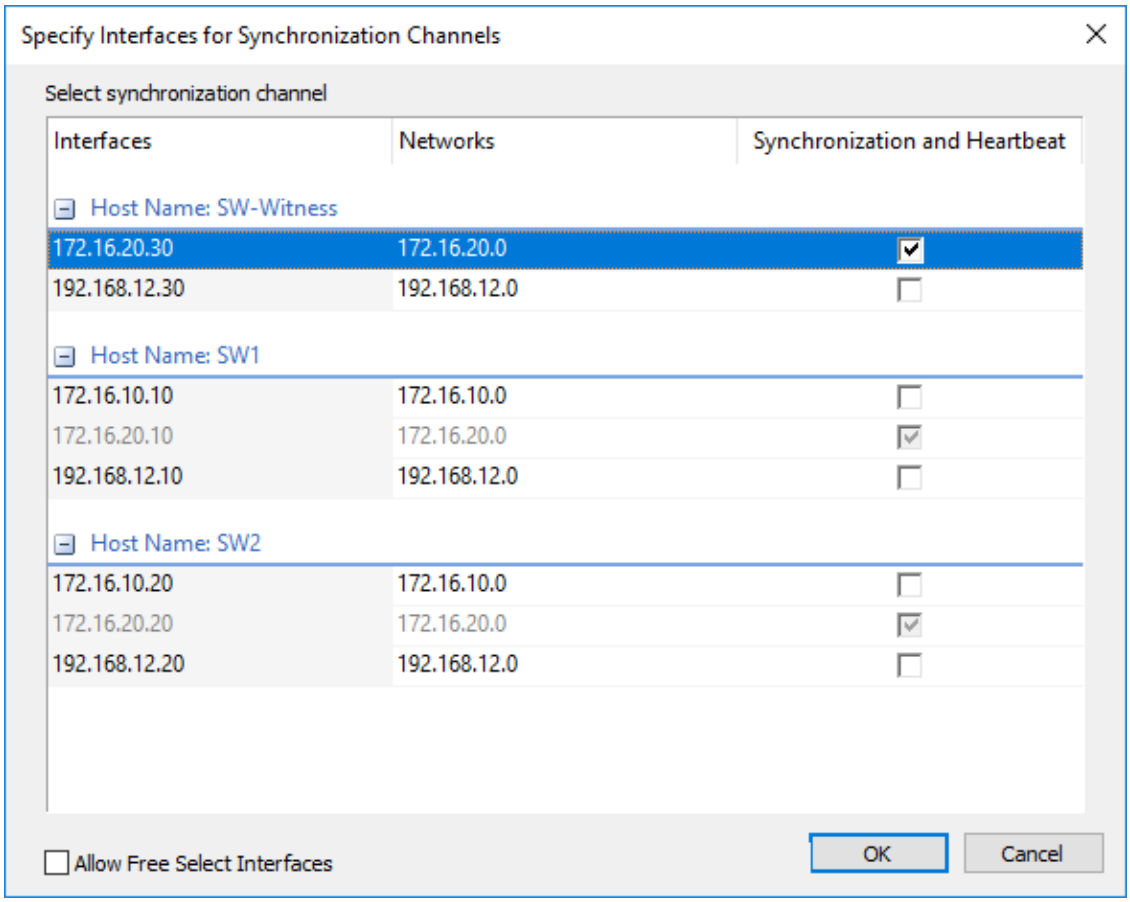

8. Click Create Replica.

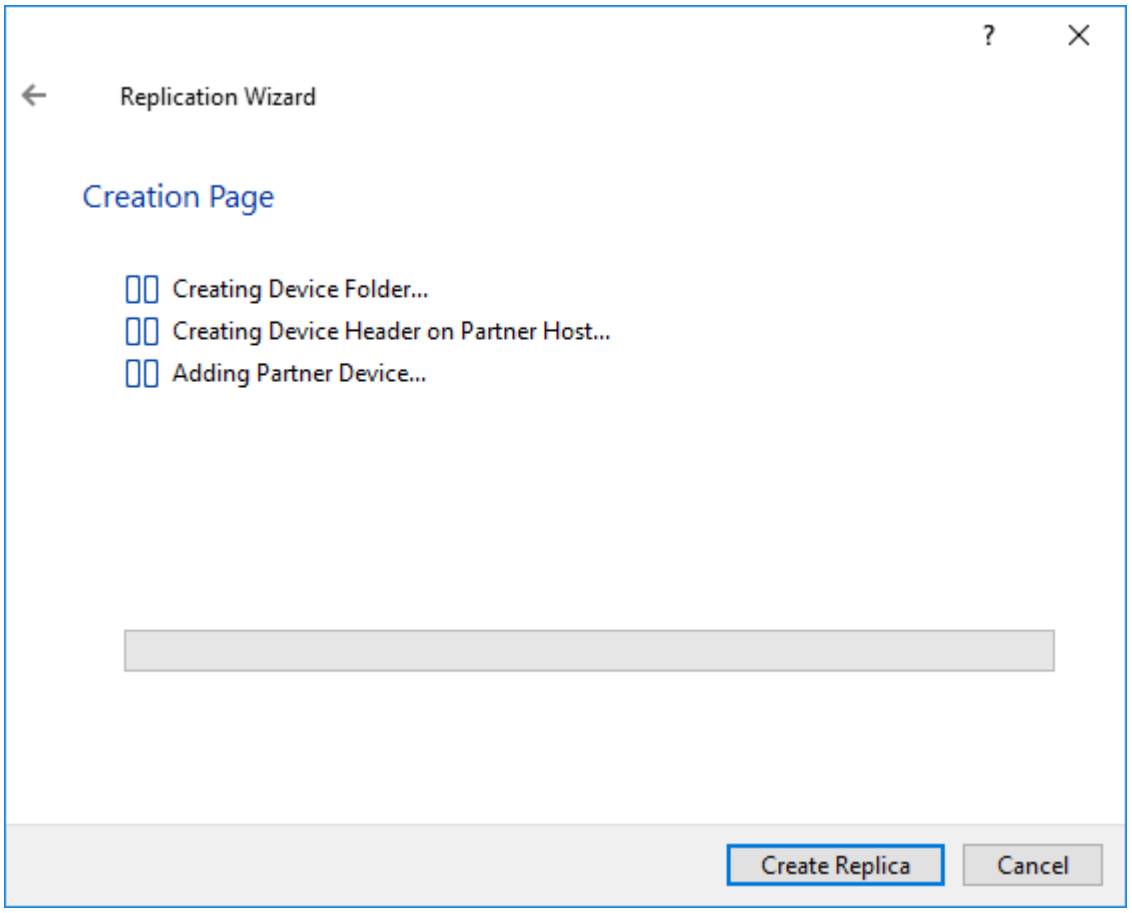

9. After the device creation is completed, close the Wizard by pressing the Close button.

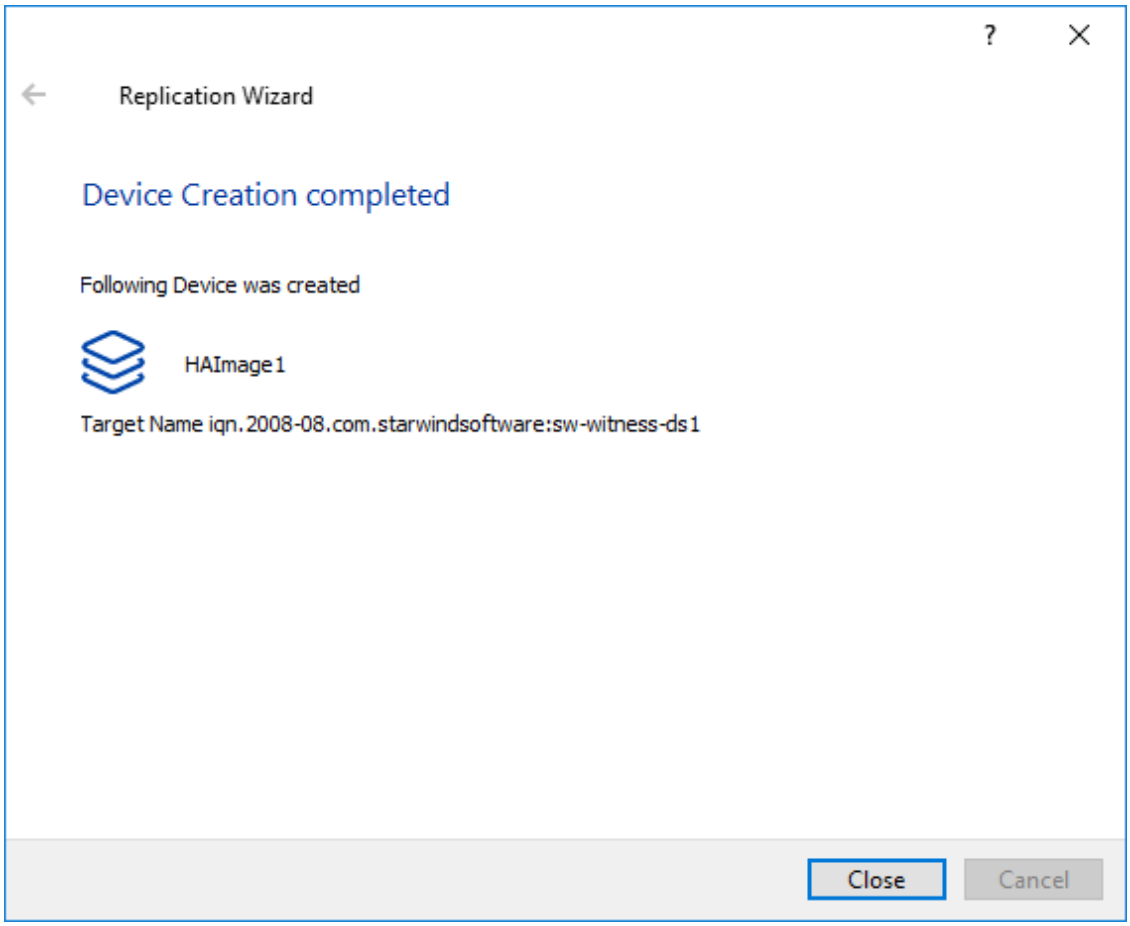

10. Repeat the steps above to create other virtual disks.

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

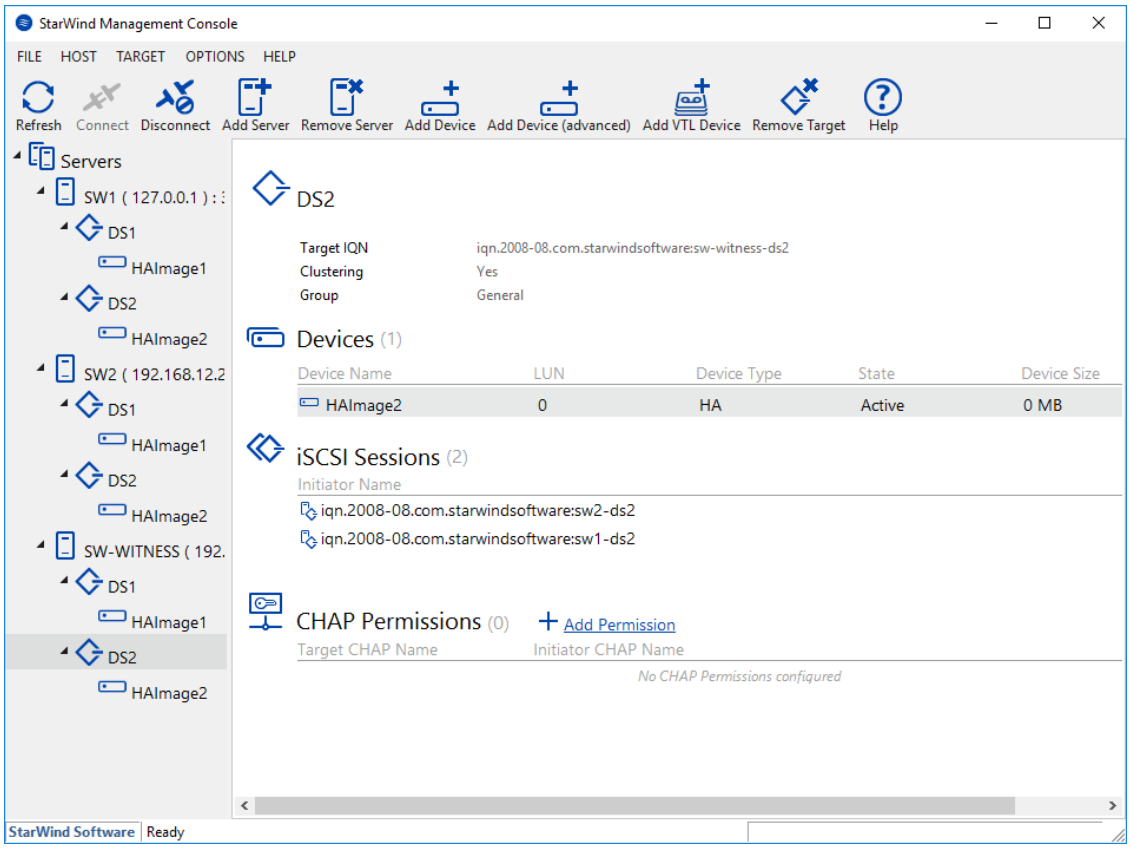

Configuring Automatic Storage Rescan

For faster paths recovery, for example, after StarWind HA device synchronization, configuring automatic storage rescan is required for each ESXi host.

1. Log in to StarWind VM and install VMware PowerCLI on each StarWind virtual machine by adding the PowerShell module (Internet connectivity is required). To do so, run the following command in PowerShell:

```
Install-Module -Name VMware.PowerCLI -AllowClobber
```

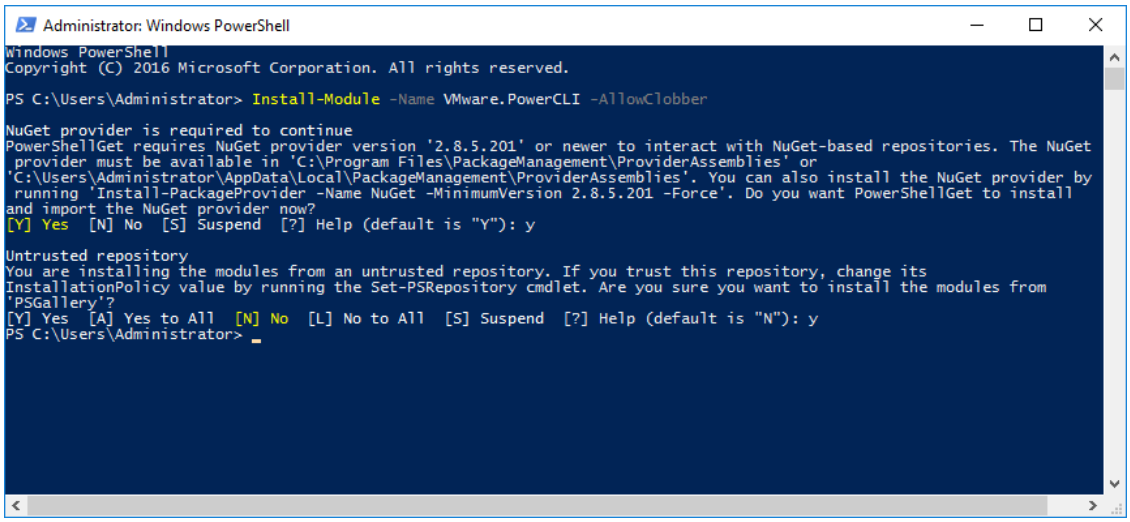

NOTE: In case of using Windows Server 2012 R2, online installation of PowerCLI requires Windows Management Framework 5.1 or upper version available on VMs. Windows Management Framework 5.1 can be downloaded from https://go.microsoft.com/fwlink/?linkid=839516

2. Open PowerShell and change the Execution Policy to Unrestricted by running the following command:

```
Set-ExecutionPolicy Unrestricted
```

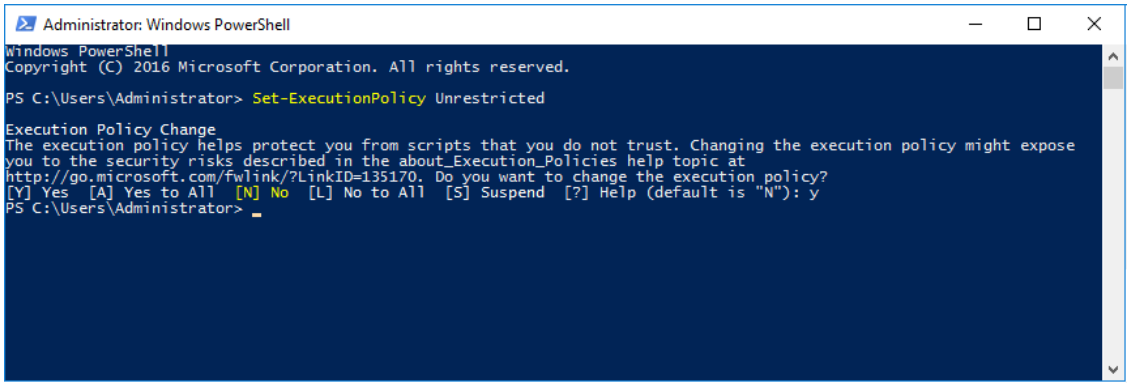

3. Create a PowerShell script which will perform an HBA rescan on the hypervisor host.

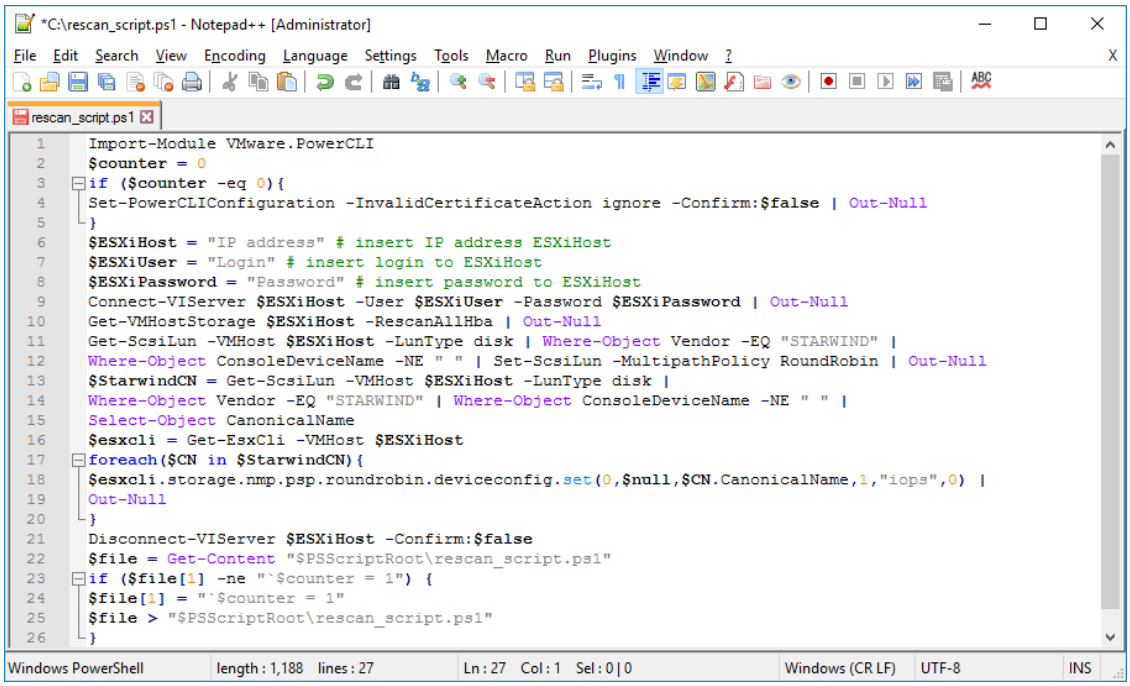


```
Import-Module VMware.PowerCLI
$counter = 0
if ($counter -eq 0){
        Set-PowerCLIConfiguration -InvalidCertificateAction
ignore -Confirm:$false | Out-Null
}
$ESXiHost = "IP address" # insert IP address ESXiHost
$ESXiUser = "Login" # insert login to ESXiHost
$ESXiPassword = "Password" # insert password ESXiHost
Connect-VIServer $ESXiHost -User $ESXiUser -Password
$ESXiPassword | Out-Null
Get-VMHostStorage $ESXiHost -RescanAllHba | Out-Null
Get-ScsiLun -VMHost $ESXiHost -LunType disk | Where-Object
Vendor -EQ "STARWIND"|
Where-Object ConsoleDeviceName -NE " " | Set-ScsiLun -
MultipathPolicy RoundRobin -CommandsToSwitchPath 1 |
Out-Null
Disconnect-VIServer $ESXiHost -Confirm:$false
$file = Get-Content "$PSScriptRoot\rescan_script.ps1"
if ($file[1] -ne "`$counter = 1") {
       $file[1] = "`$counter = 1"
       $file > "$PSScriptRoot\rescan_script.ps1"
}
```

In the appropriate lines, specify the IP address and login credentials of the ESXi host on

which the current StarWind VM is stored and running:

$ESXiHost = “IP address”
$ESXiUser = “Login”
$ESXiPassword = “Password”

Save the script as rescan_script.ps1 to the root of the C:\ drive of the Virtual Machine.

NOTE: In some cases the rescan script can be changed and storage rescan added for another ESXi host. Appropriate lines should be duplicated and changed with properly edited variables if required.
NOTE: In some cases, it makes sense to create a separate ESXi user for storage rescans. To create the user, please follow the steps below:

Log in to ESXi with the VMware Host Client. Click Manage, and under Security & users tab, in the Users section click Add user button. In the appeared window, enter a user name, and a password.

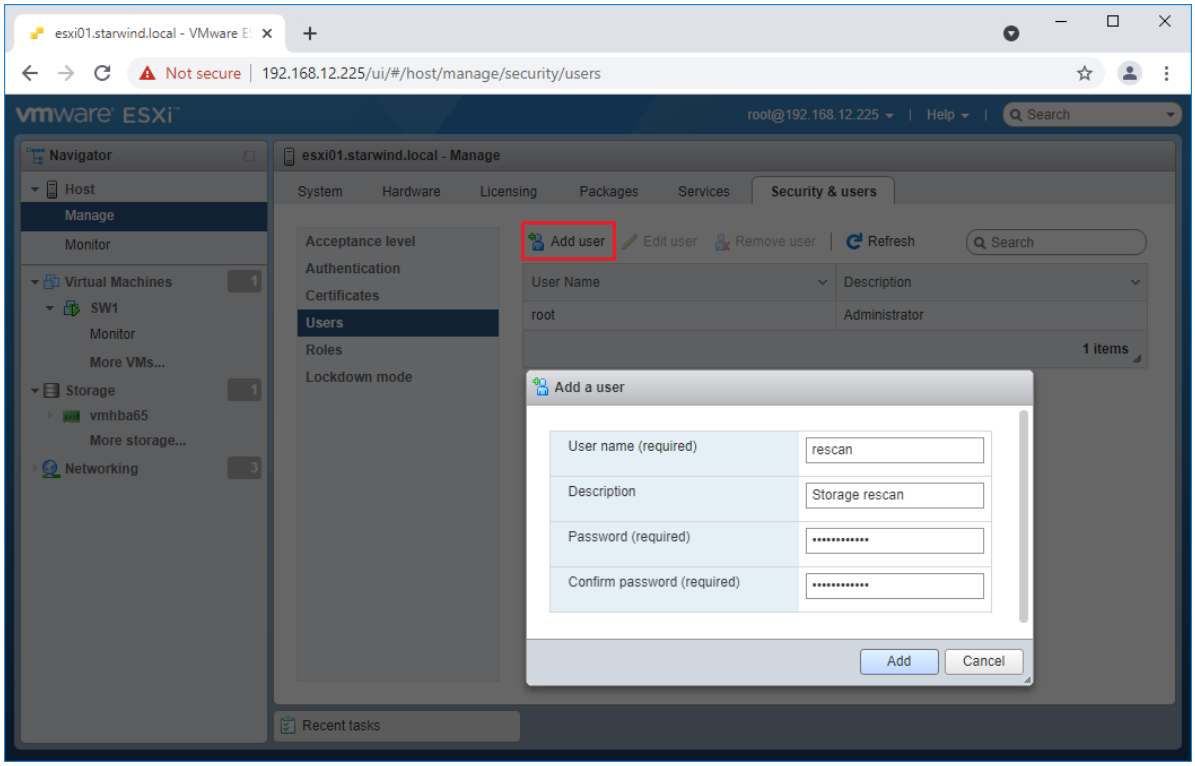

Create a new Role, under Roles section, and click New Role button. Type a name for the new role. Select privileges for the role and click OK.
The following privileges might be assigned: Host – Inventory, Config, Local Cim, and Global – Settings.

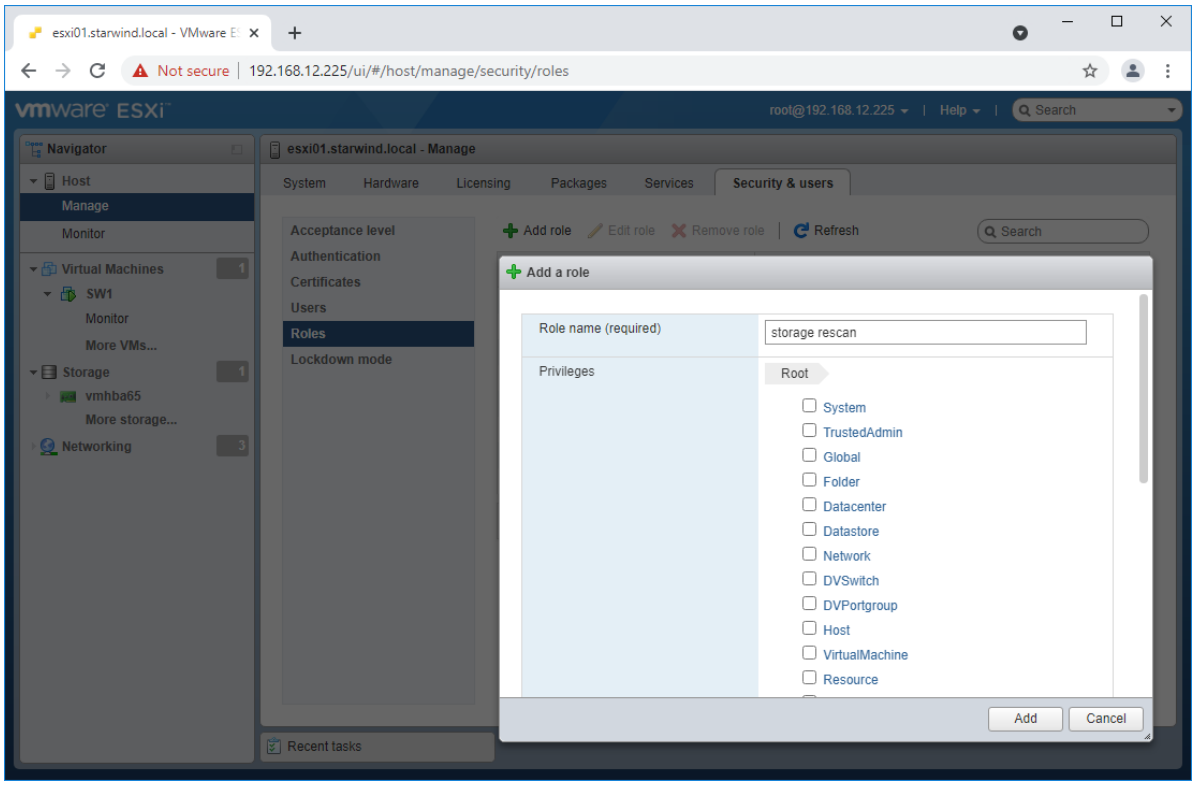

Assign permission to the storage rescan user for an ESXi host – right-click Host in the VMware Host Client inventory and click Permissions. In the appeared window click Add user.

Click the arrow next to the Select a user text box and select the user that you want to assign a role to. Click the arrow next to the Select a role text box and select a role from the list.

(Optional) Select Propagate to all children or Add as group. Click Add user and click Close.

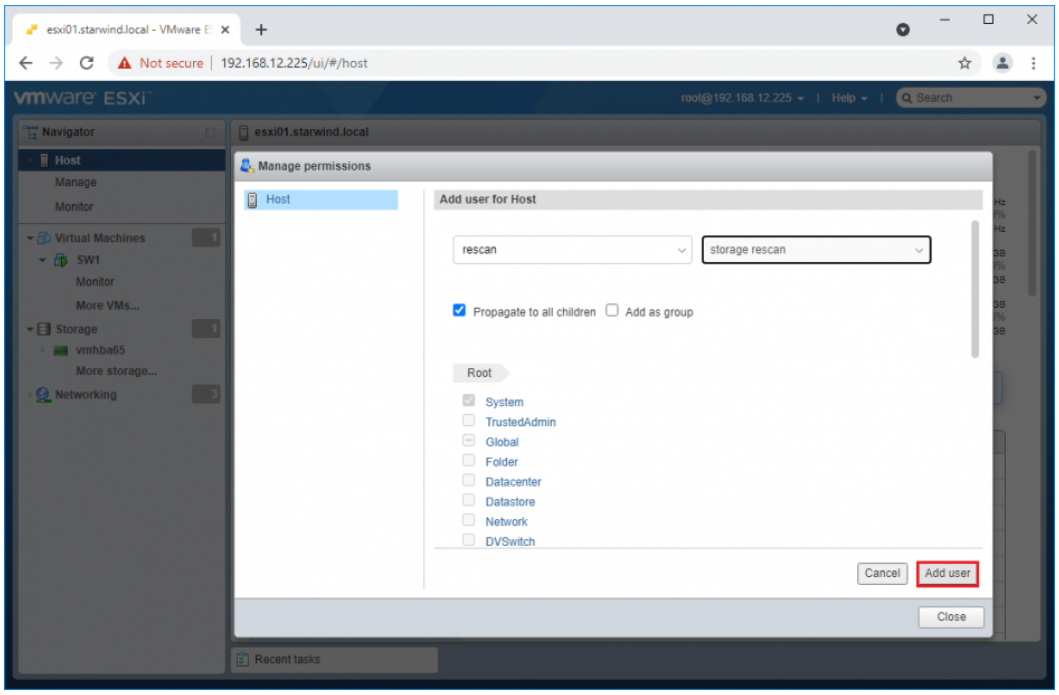

3. Repeat all steps from this section on the other ESXi hosts.

4. Perform the configuration steps above on the partner node.

5. Go to Control Panel -> Administrative Tools -> Task Scheduler -> Create Basic Task and follow the wizard steps:

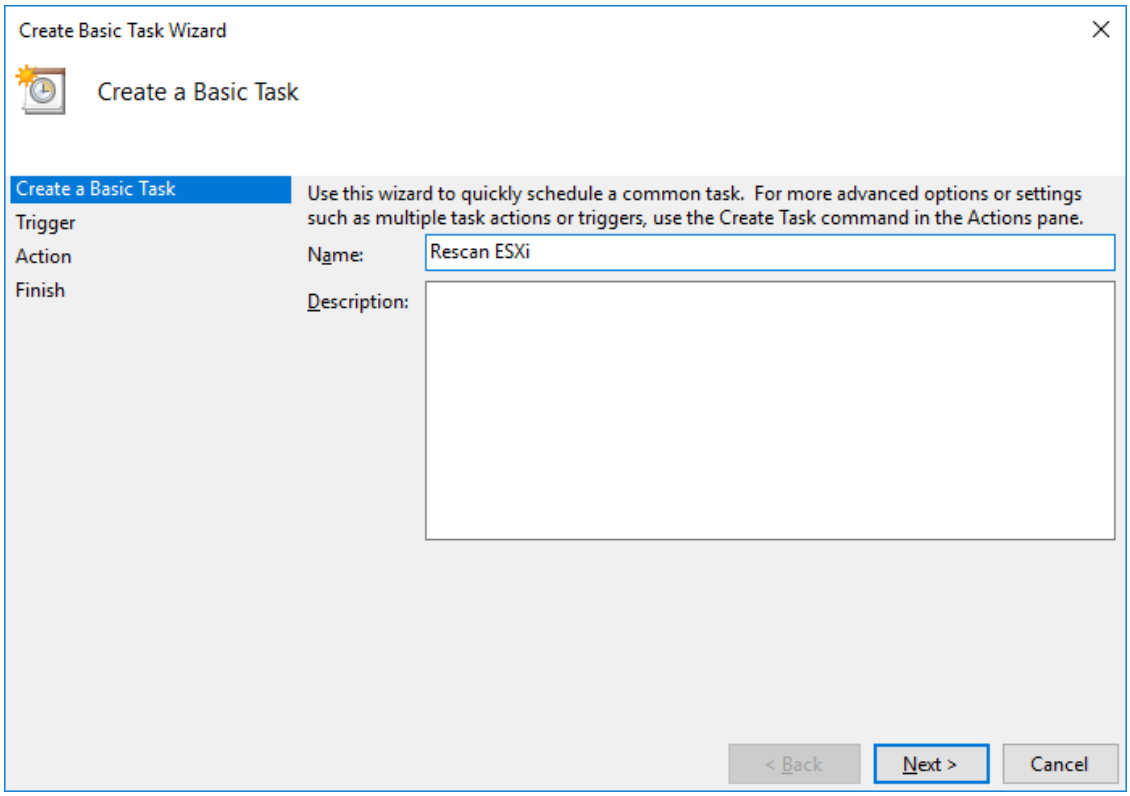

6. Specify the task name, select When a specific event is logged, and click on Next.

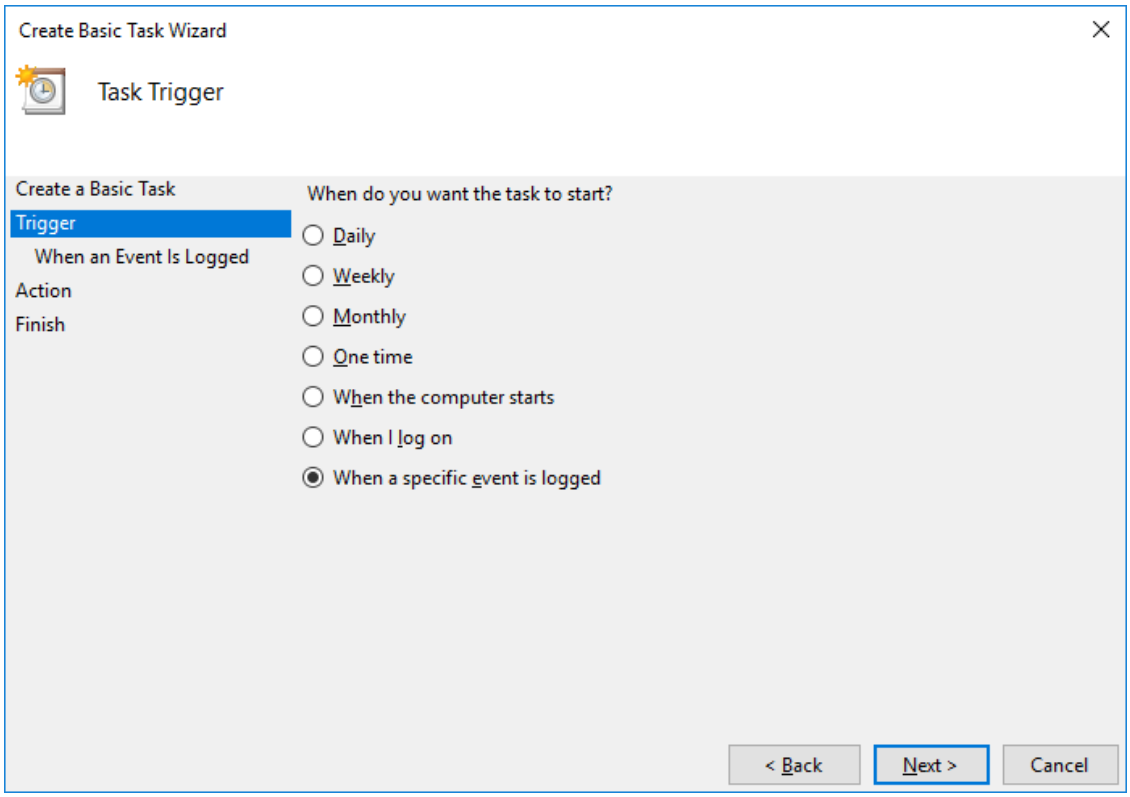

7. Select Application in the Log dropdown, type StarWindService for the event source and 788 as the event ID. Click the Next button.

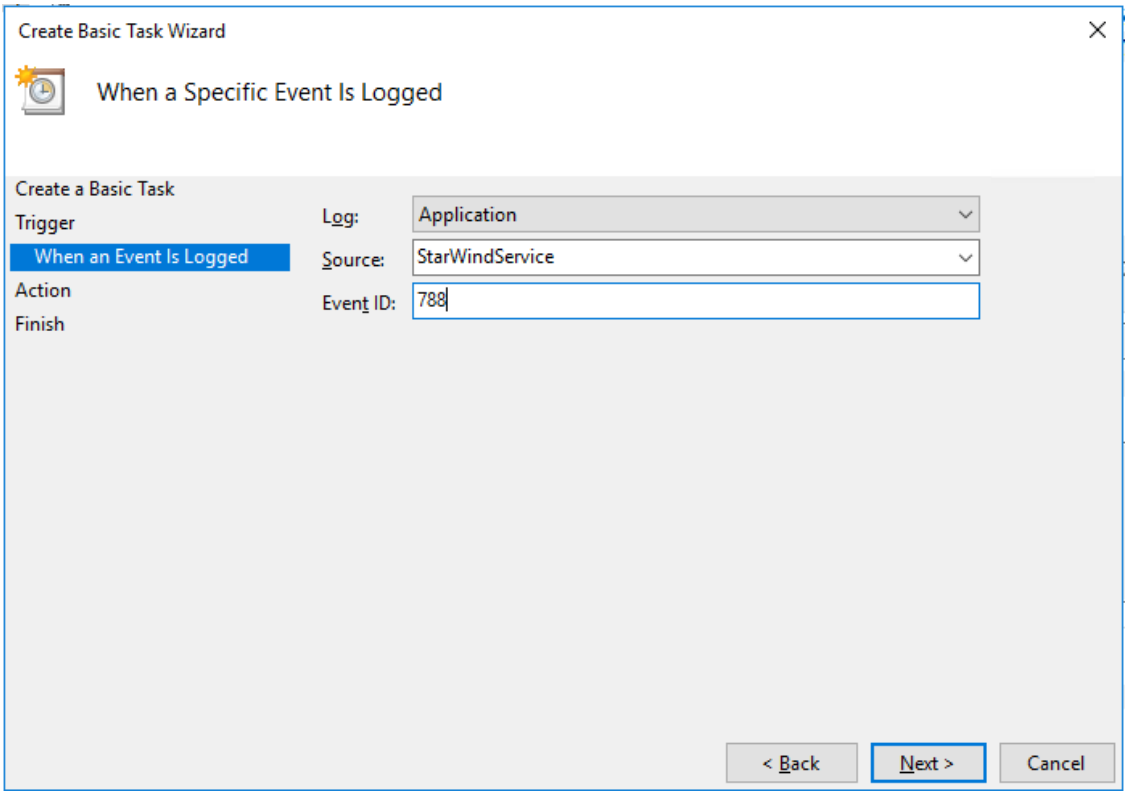

8. Choose Start a Program as the action that the task will perform and click on Next.

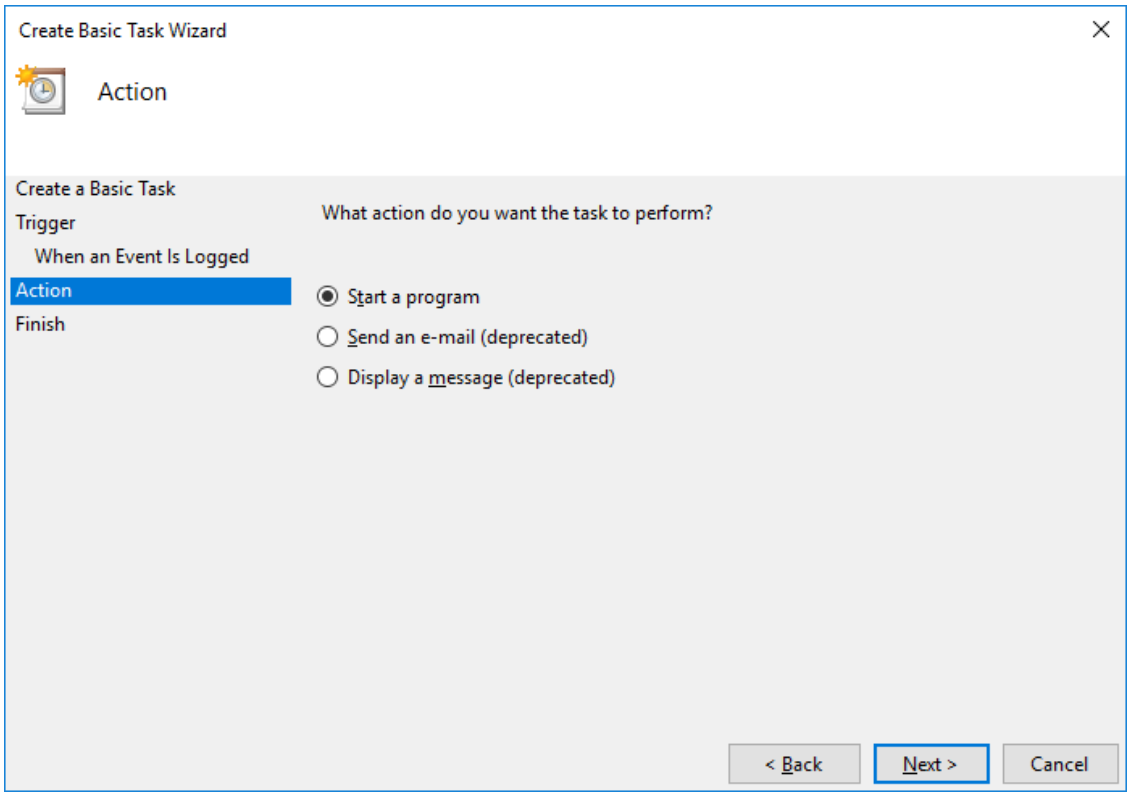

9. Type powershell.exe in the Program/script field. In the Add arguments field, type:

“ -ExecutionPolicy Bypass -NoLogo -NonInteractive -NoProfile -WindowStyle Hidden -File C:\rescan_script.ps1 ”

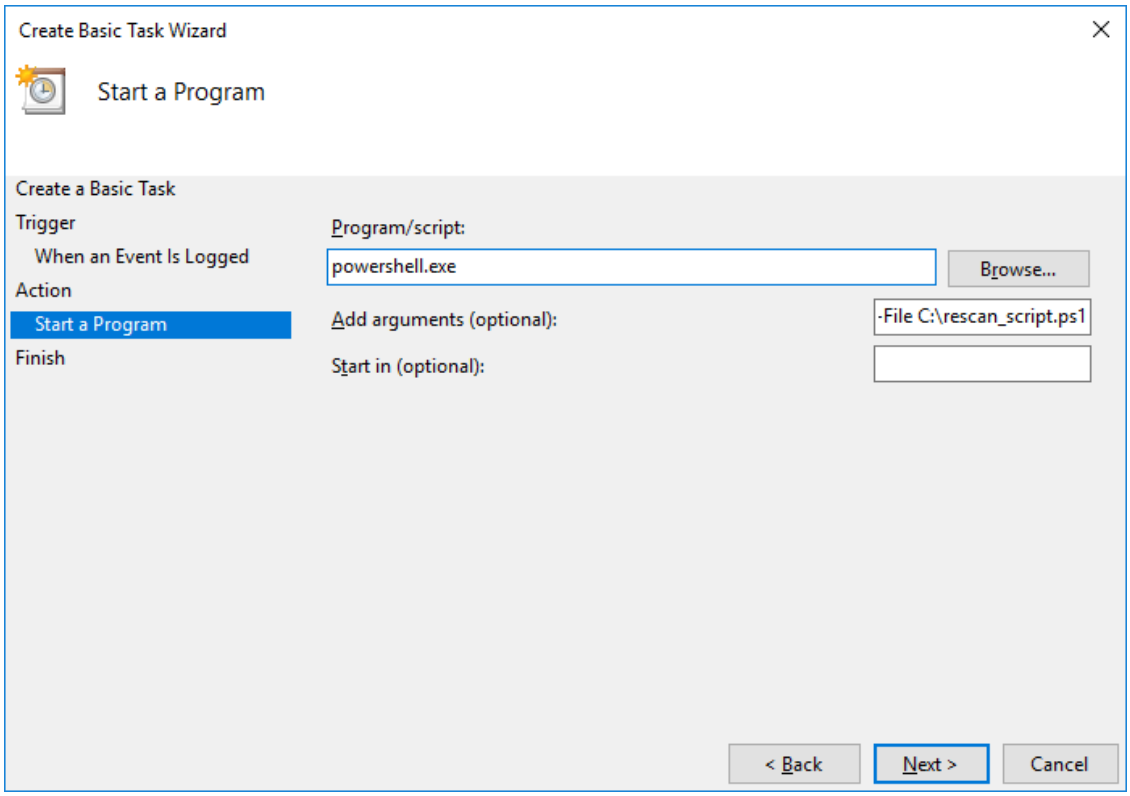

Click the Next button to continue.

10. Click Finish to exit the Wizard.

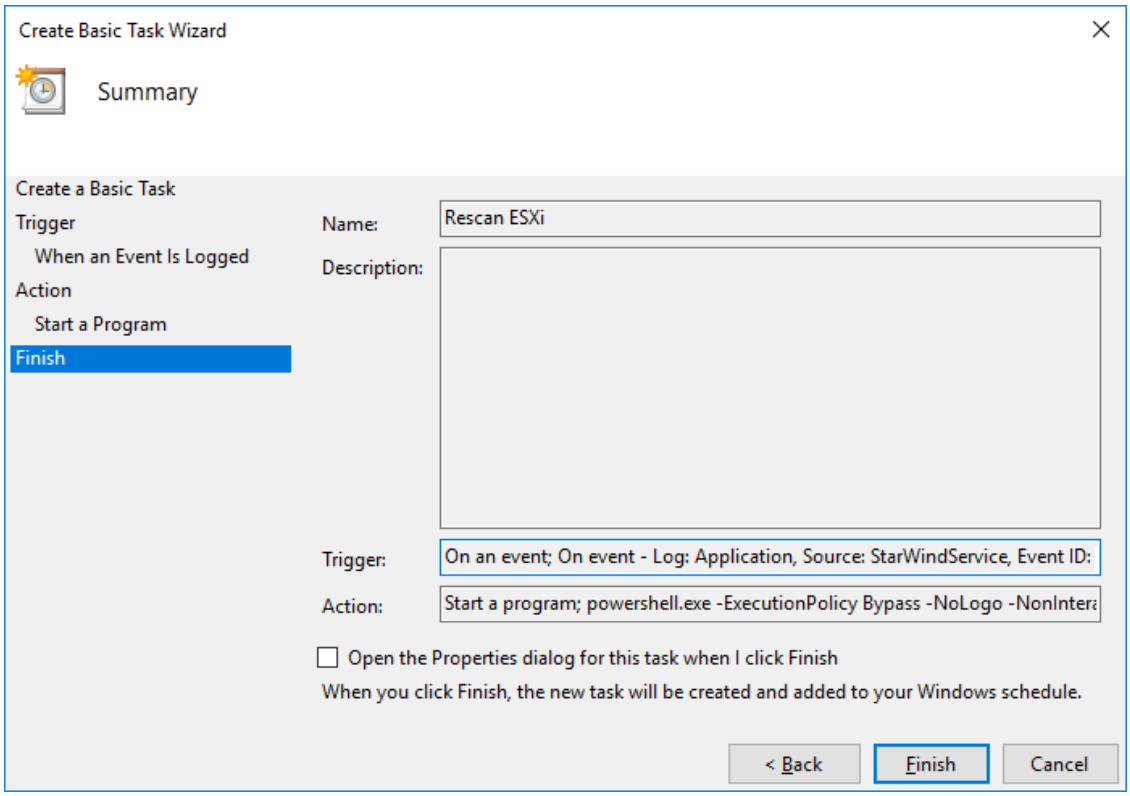

11. Configure the task to run with highest privileges by enabling the checkbox at the bottom of the window. Also, make sure that the “Run whether user is logged on or not” option is selected.

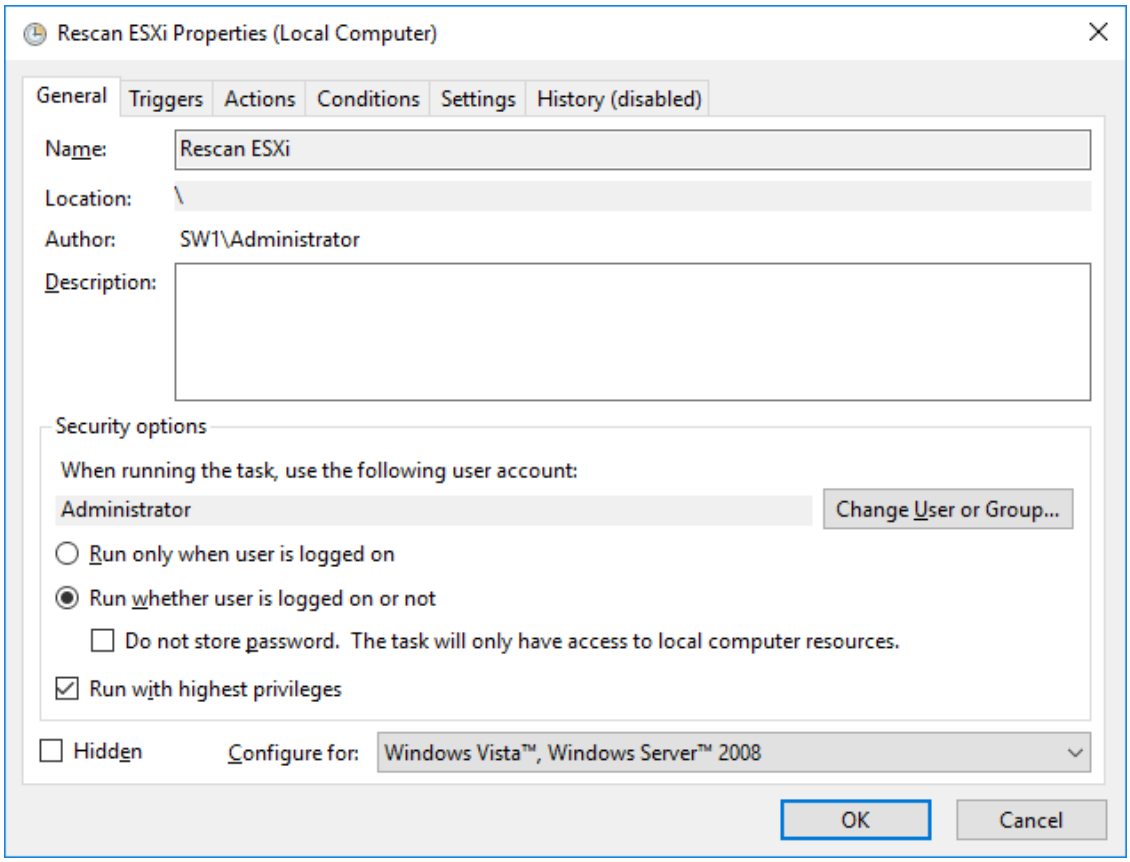

12. Switch to the Triggers tab. Verify that the trigger on event 788 is set up correctly.

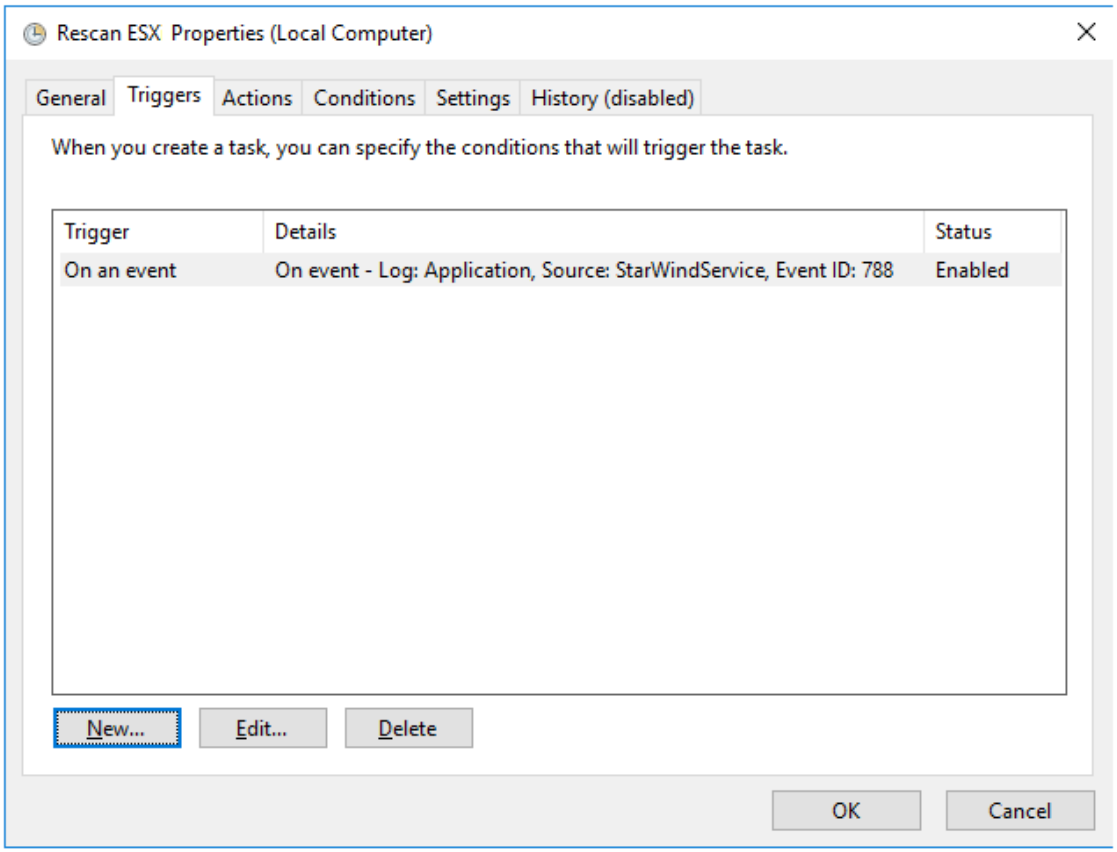

13. Click New and add other triggers by Event ID 782, 769, 773, and 817.

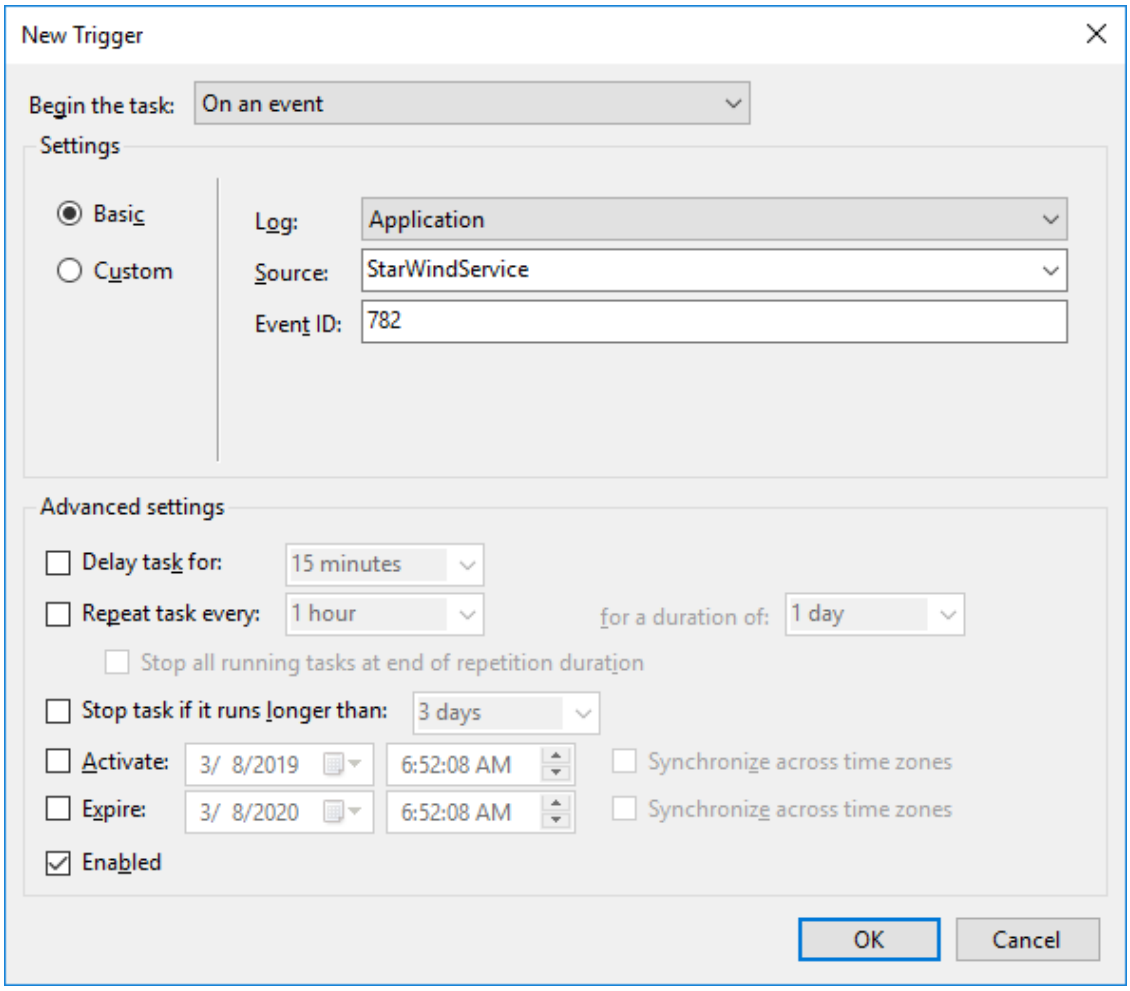

14. All added triggers should look like in the picture below.

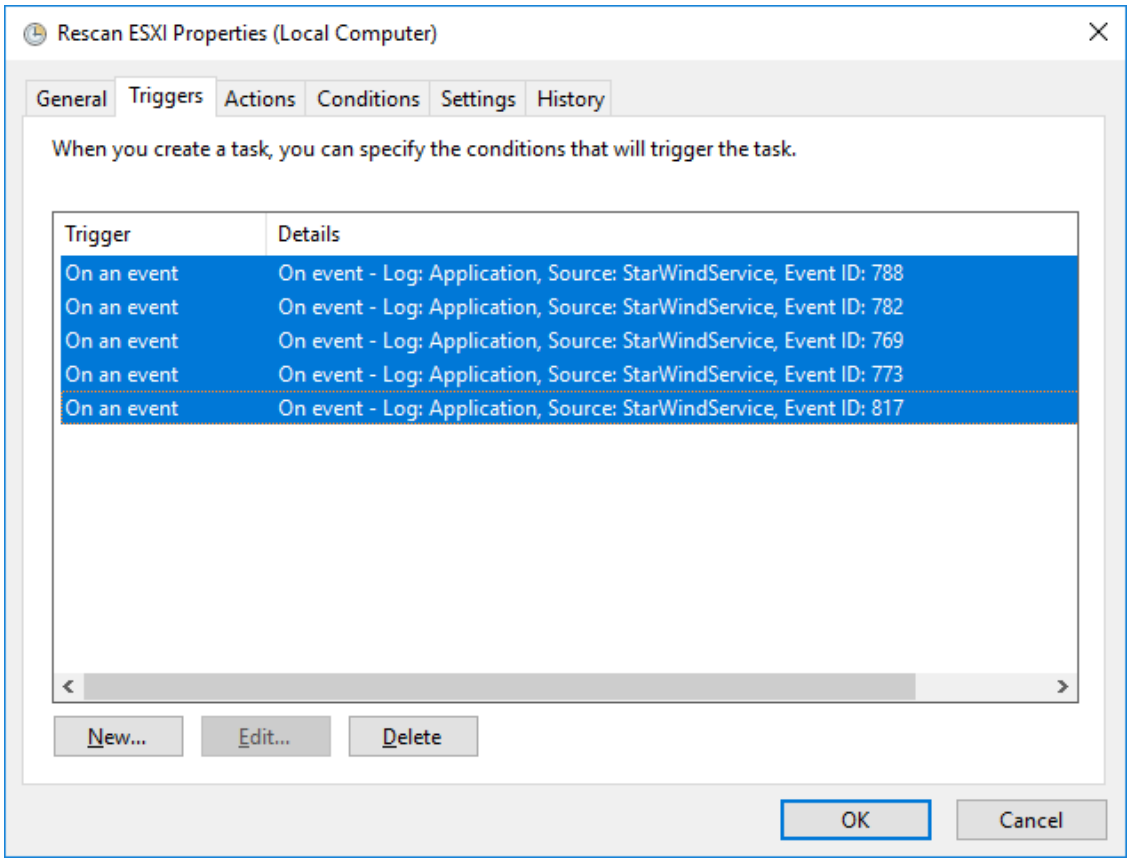

15. Switch to the Actions tab and verify the parameters for the task.

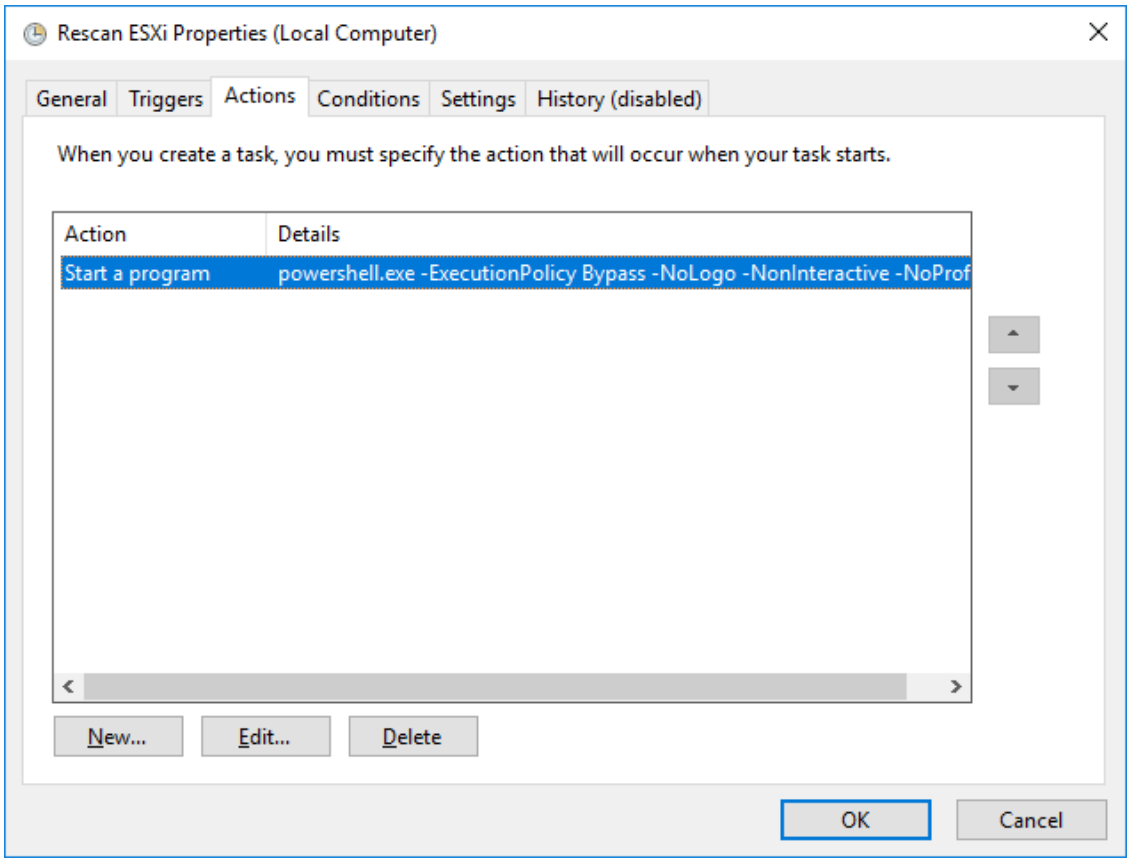

Press OK and type in the credentials for the user whose rights will be used to execute the command.

16. Perform the same steps on another StarWind VM, specifying the corresponding settings.

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.

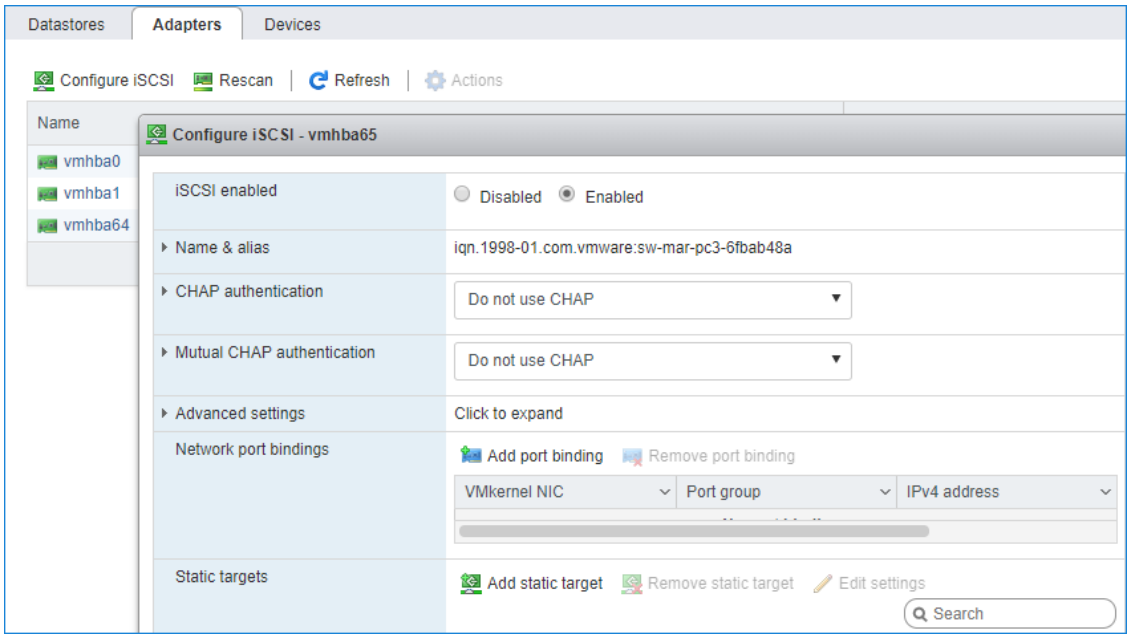

2. In the Configure iSCSI window, under Dynamic Targets, click on the Add dynamic target button to specify iSCSI interfaces.

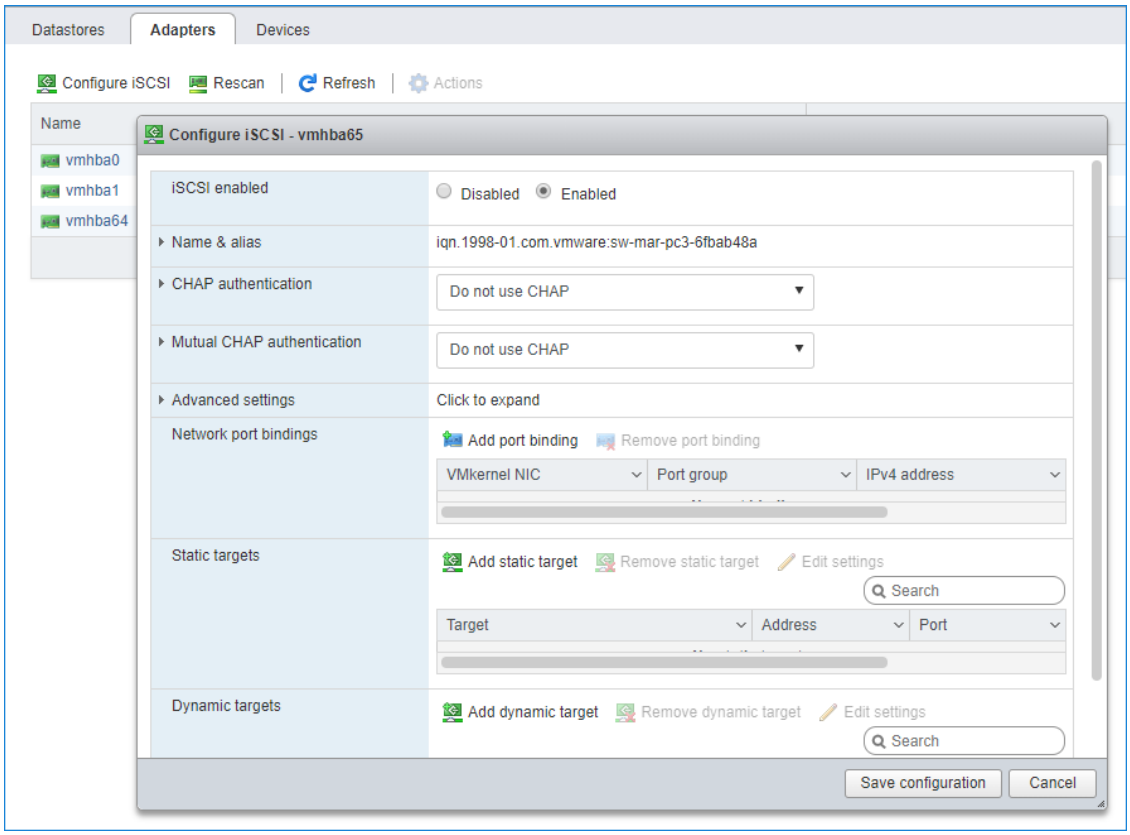

3. Enter the iSCSI IP addresses of all StarWind nodes for the iSCSI traffic.

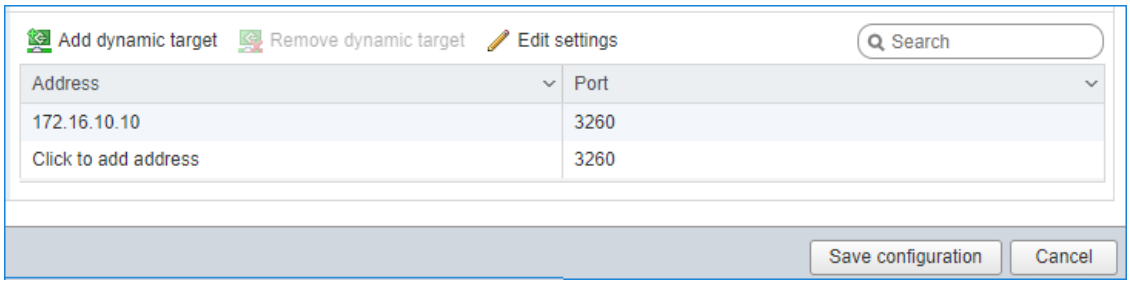

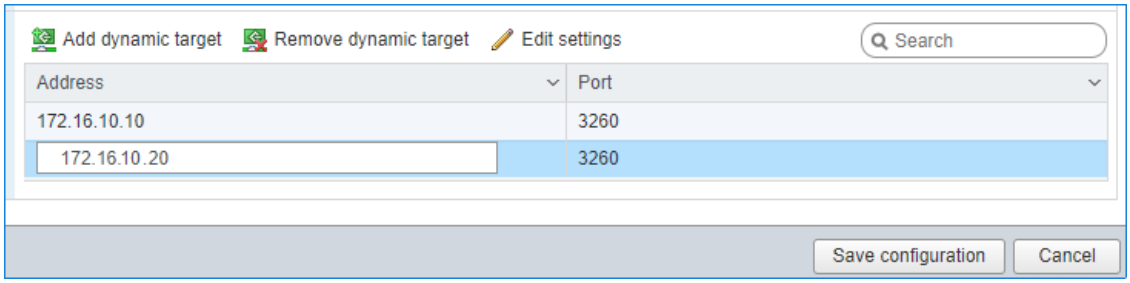

Confirm the actions by pressing Save configuration.

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

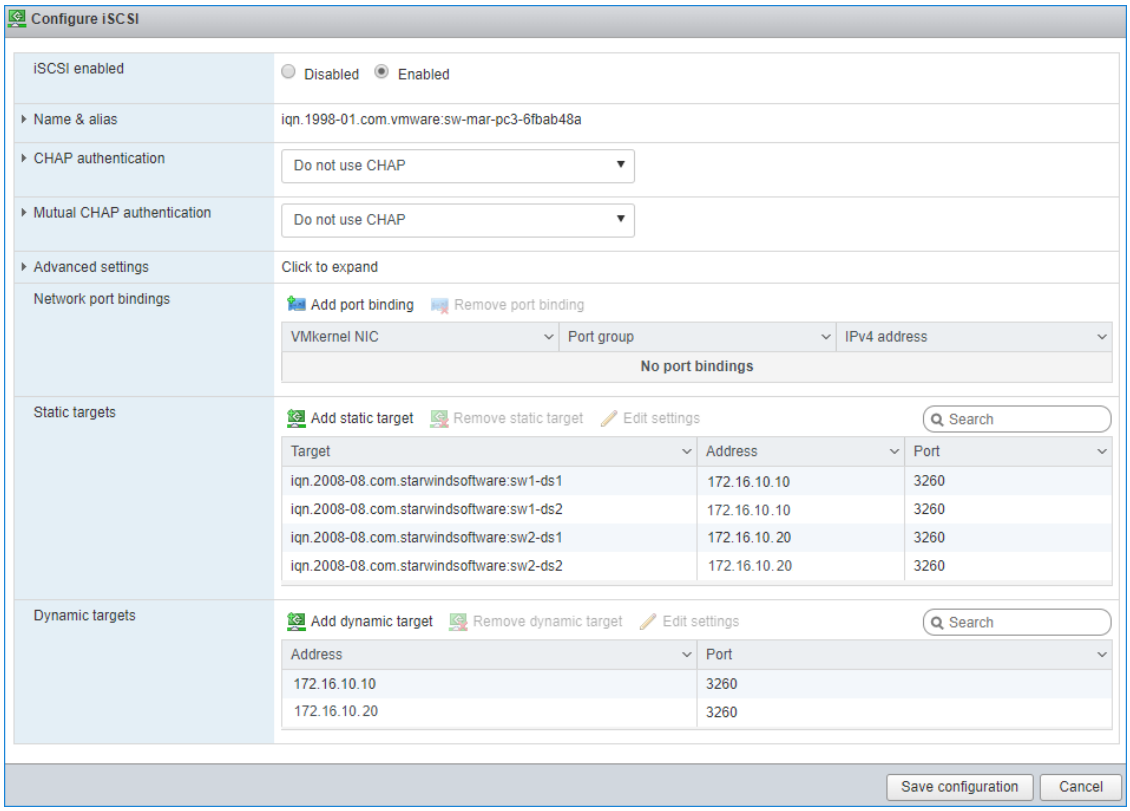

5. Click on the Rescan button to rescan storage.

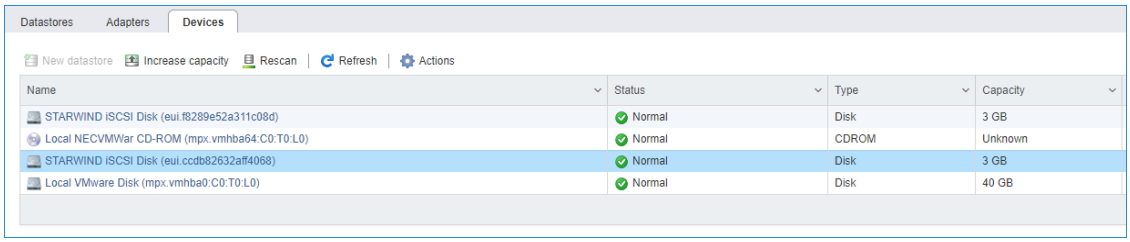

6. Now, the previously created StarWind devices are visible to the system.

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.

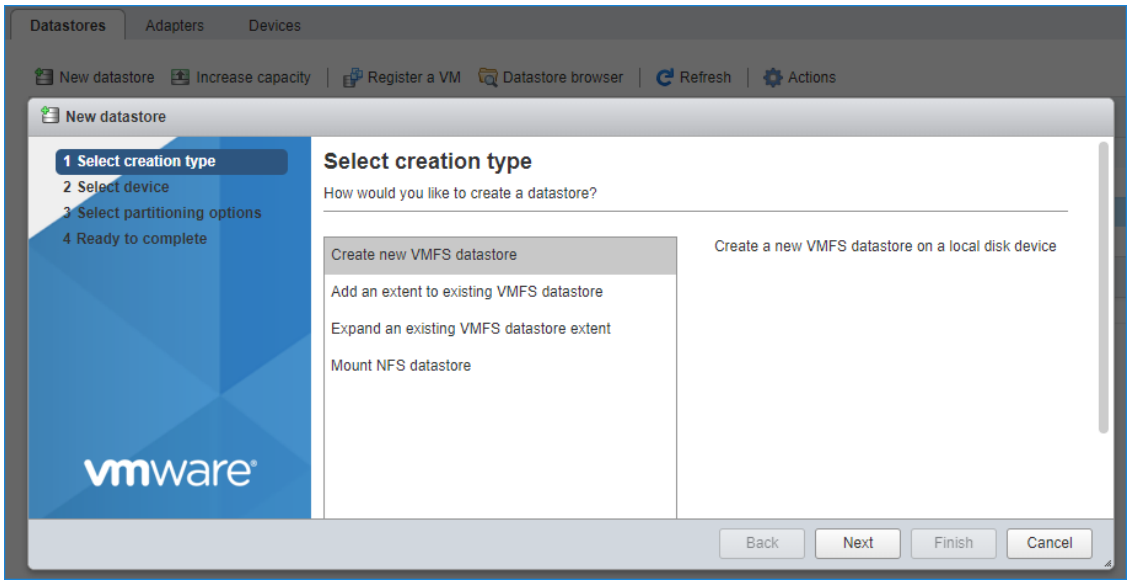

2. Specify the datastore name, select the previously discovered StarWind device, and click on Next.

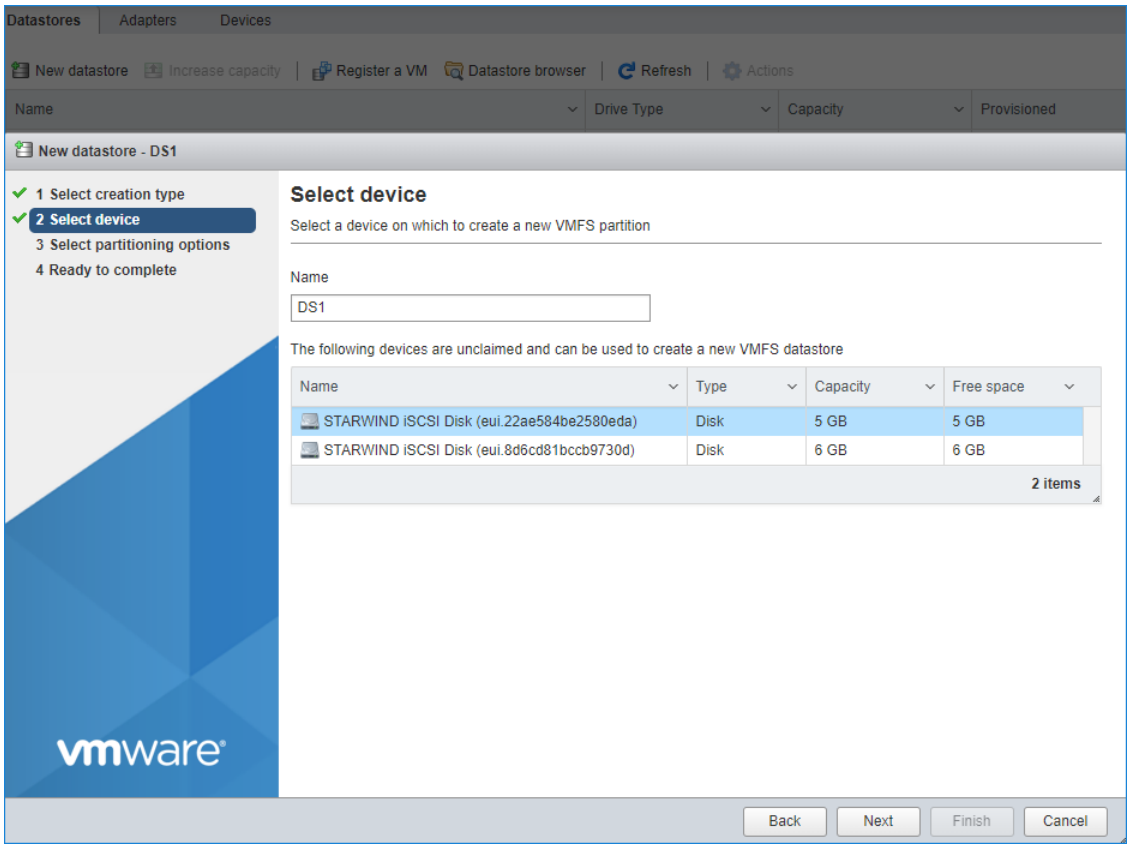

3. Enter datastore size. Click on Next.

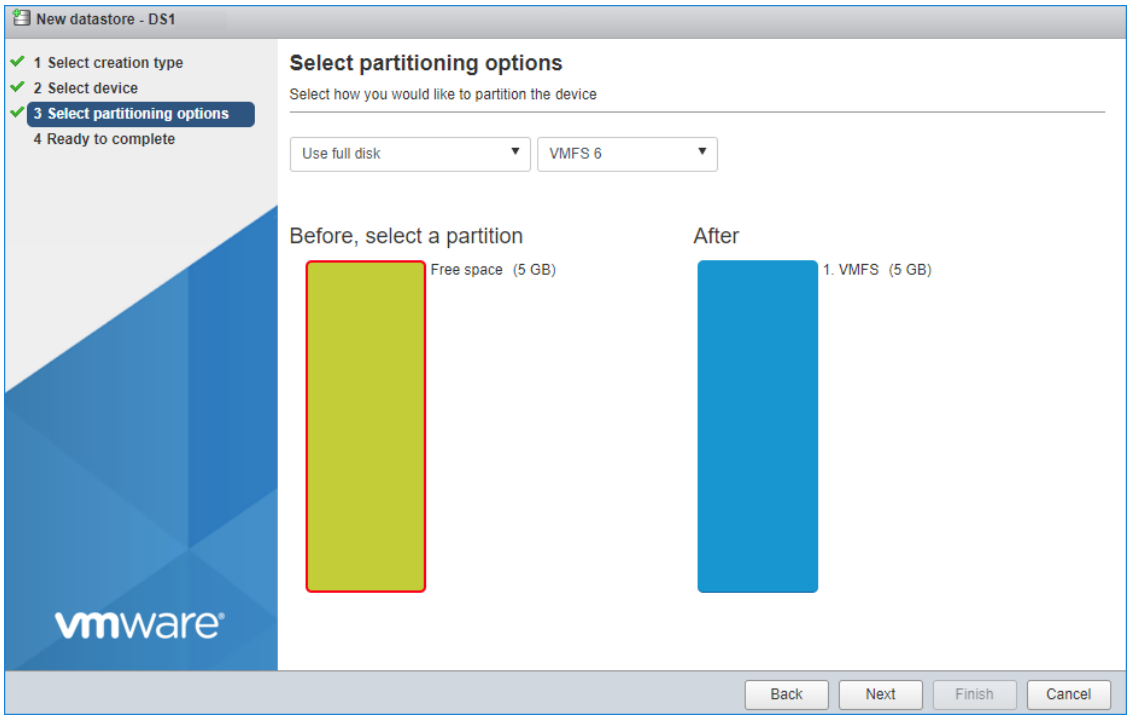

4. Verify the settings. Click on Finish.

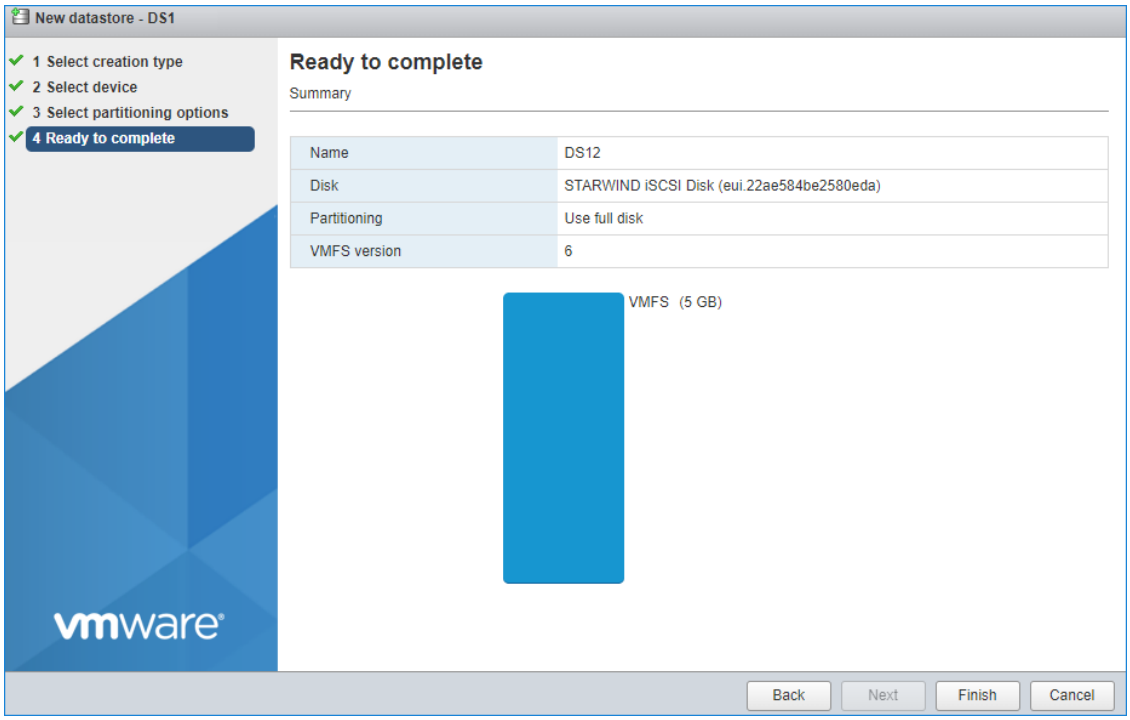

5. Add another datastore (DS2) in the same way but select the second device for it.

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

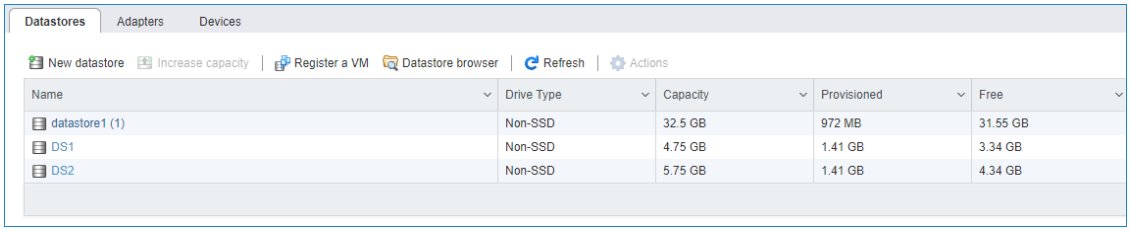

7. Path Selection Policy changing for Datastores from Most Recently Used (VMware) to Round Robin (VMware) has been already added into the Rescan Script, and this action is performed automatically. For checking and changing this parameter manually, the hosts should be connected to vCenter.

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.

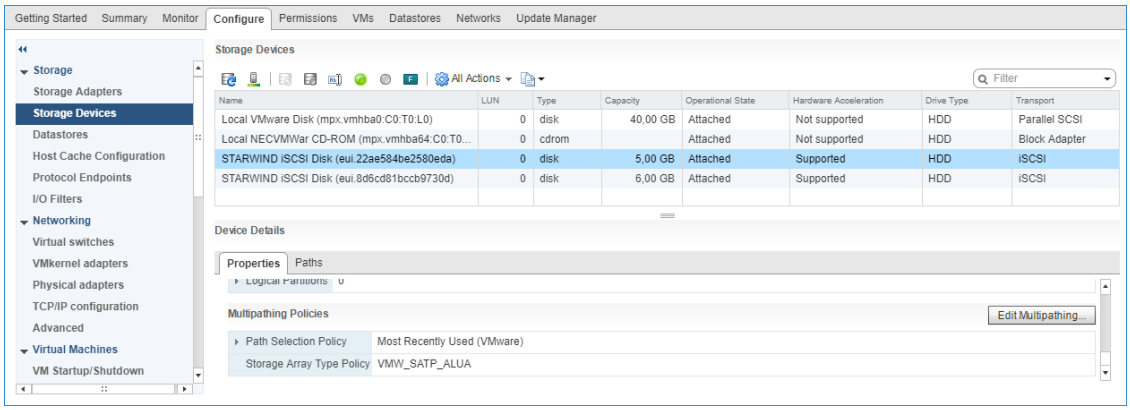

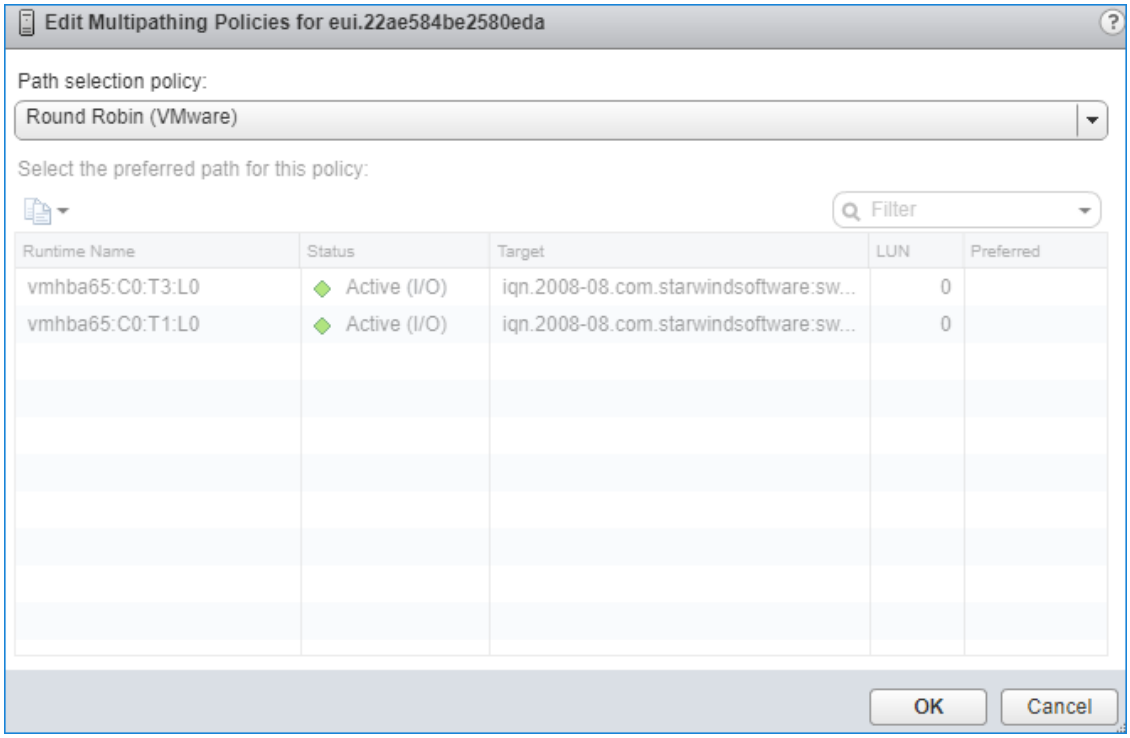

Performance Tweaks

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

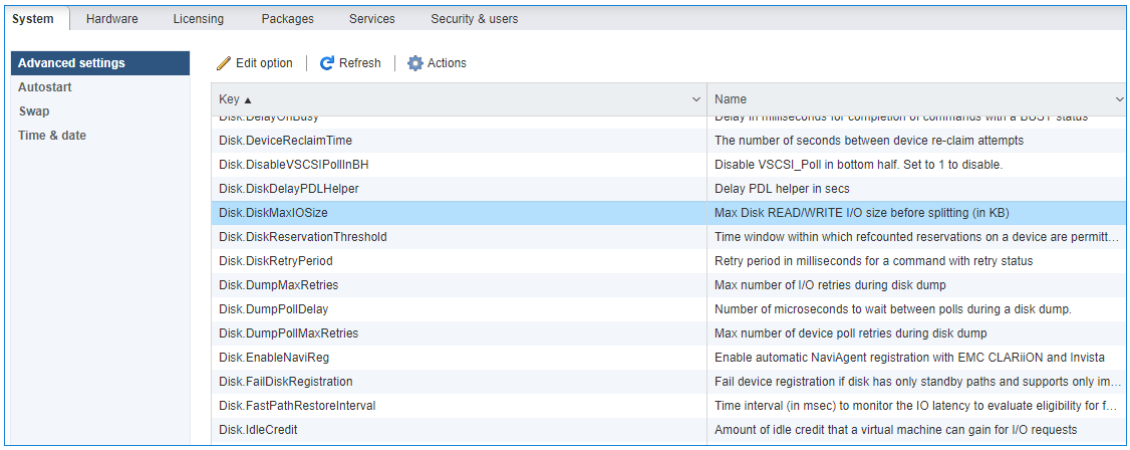

2. Select Disk and change the Disk.DiskMaxIOSize parameter to 512.

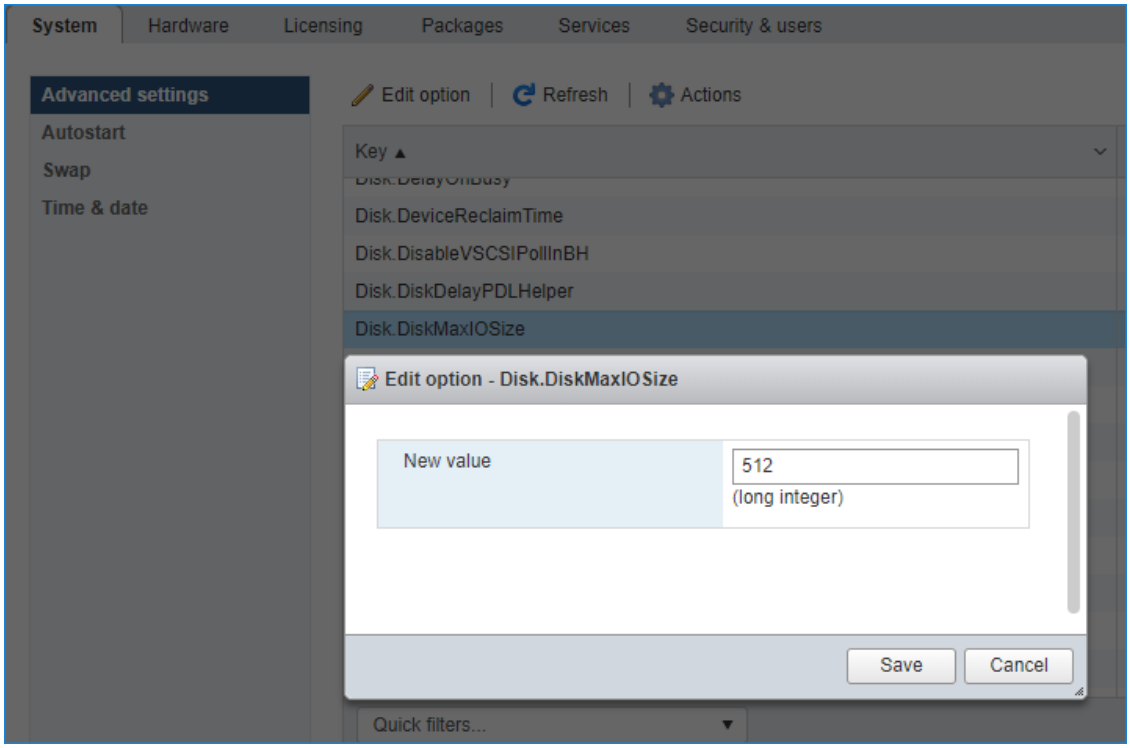

NOTE: Changing Disk.DiskMaxIOSize to 512 might cause startup issues with Windows-based VMs, located on the datastore where specific ESX builds are installed. If the issue with VMs start appears, leave this parameter as default or update the ESXi host to the next available build.
NOTE: In certain cases, 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.

Conclusion

Following this guide, a VMware ESXI Cluster was configured with StarWind Virtual SAN (VSAN) running in a CVM 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