

StarWind Virtual HCI Appliance: Configuration Guide for Red Hat oVirt [KVM], VHCA Deployed using Automated Installation

2024

TECHNICAL PAPERS



Trademarks

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

Changes

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

Technical Support and Services

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

About StarWind

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

Copyright ©2009-2018 StarWind Software Inc.

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

Annotation

Relevant products

StarWind Virtual HCI Appliance (VHCA)

Purpose

This document outlines how to configure a StarWind Virtual HCI Appliance (VHCA) based on Red Hat oVirt, with VSAN running as a Controller Virtual Machine (CVM). The guide includes steps to prepare oVirt 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 Red Hat oVirt that includes virtual machine shared storage provided by StarWind VSAN.

Prerequisites

StarWind Virtual HCI Appliance (VHCA) system requirements

Prior to configuring StarWind Virtual HCI Appliance (VHCA), please make sure that the system meets the requirements, which are available via the following link:

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

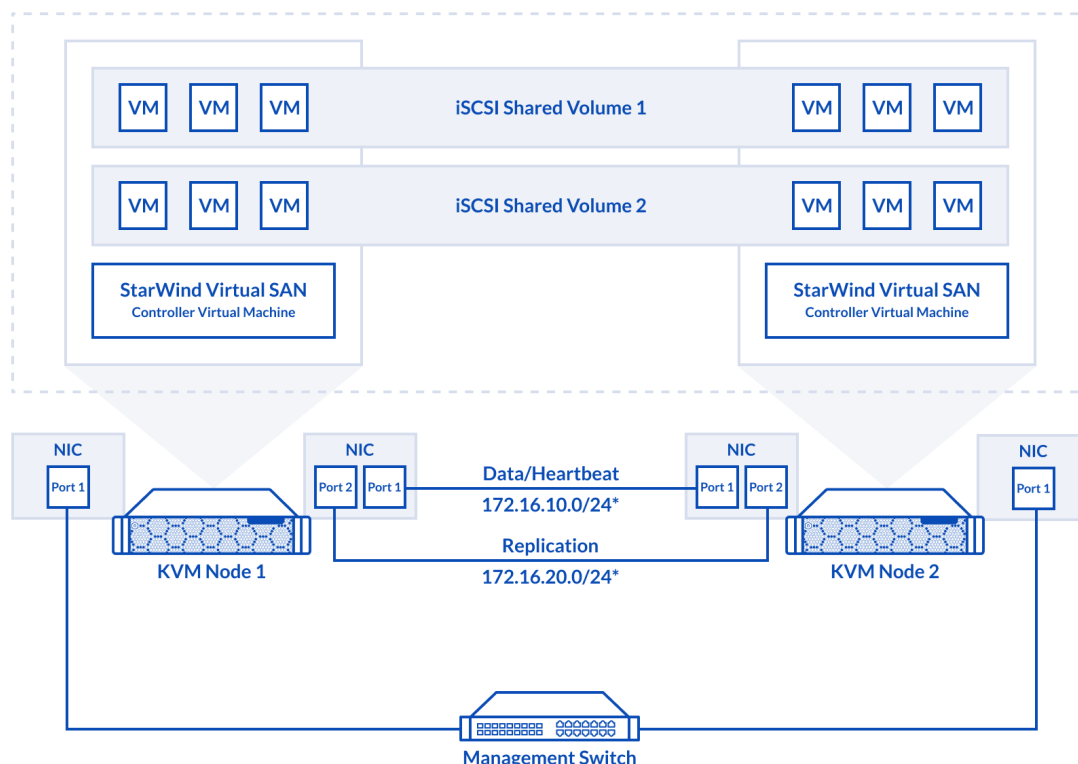
Recommended RAID settings for HDD and SSD disks:

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

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

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

Solution diagram



Prerequisites:

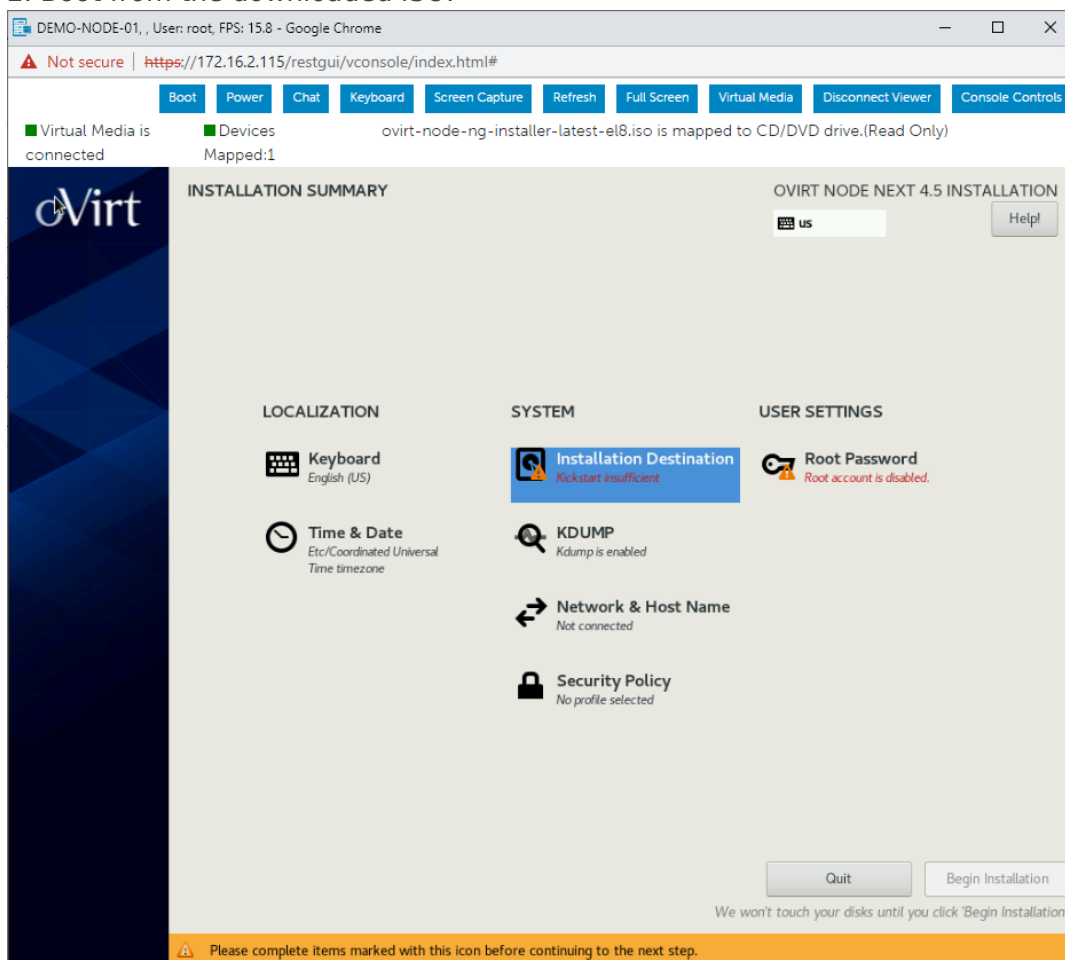
1. 2 servers with local storage, which have direct network connections for Synchronization and iSCSI/StarWind heartbeat traffic.
2. Servers should have local storage available for OS installation and StarWind VSAN Controller Virtual Machine. CVM utilizes local storage to create replicated shared storage connected to oVirt nodes via iSCSI.
3. DNS records of the oVirt Nodes and oVirt Engine should be added to DNS Server. It is required to use static IP addresses for both oVirt node and oVirt Engine.
4. StarWind HA devices require at least 2 separate network links between the nodes. The first one is used for iSCSI traffic, the second one is used for Synchronization traffic.
Note. The network interfaces on each node for Synchronization and iSCSI/StarWind heartbeat interfaces should be in different subnets and connected directly according to the network diagram above. Here, the 172.16.10.x subnet is used for the iSCSI/StarWind heartbeat traffic, while the 172.16.20.x subnet is used for the Synchronization traffic.

Deploying Ovirt Node

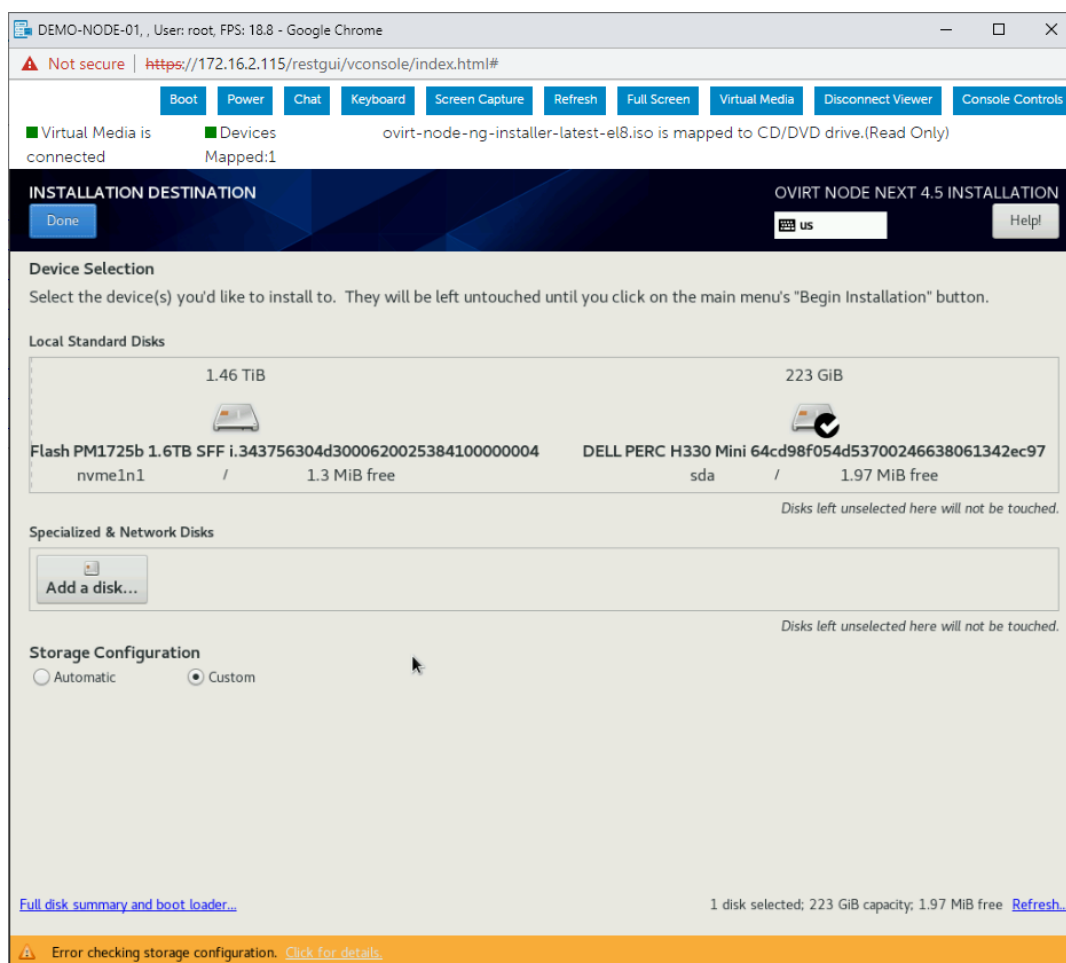
1. Download oVirt (CentOS-based) node iso.

<https://resources.ovirt.org/pub/ovirt-4.5/iso/ovirt-node-ng-installer/ovirt-node-ng-installer-latest-el8.iso>

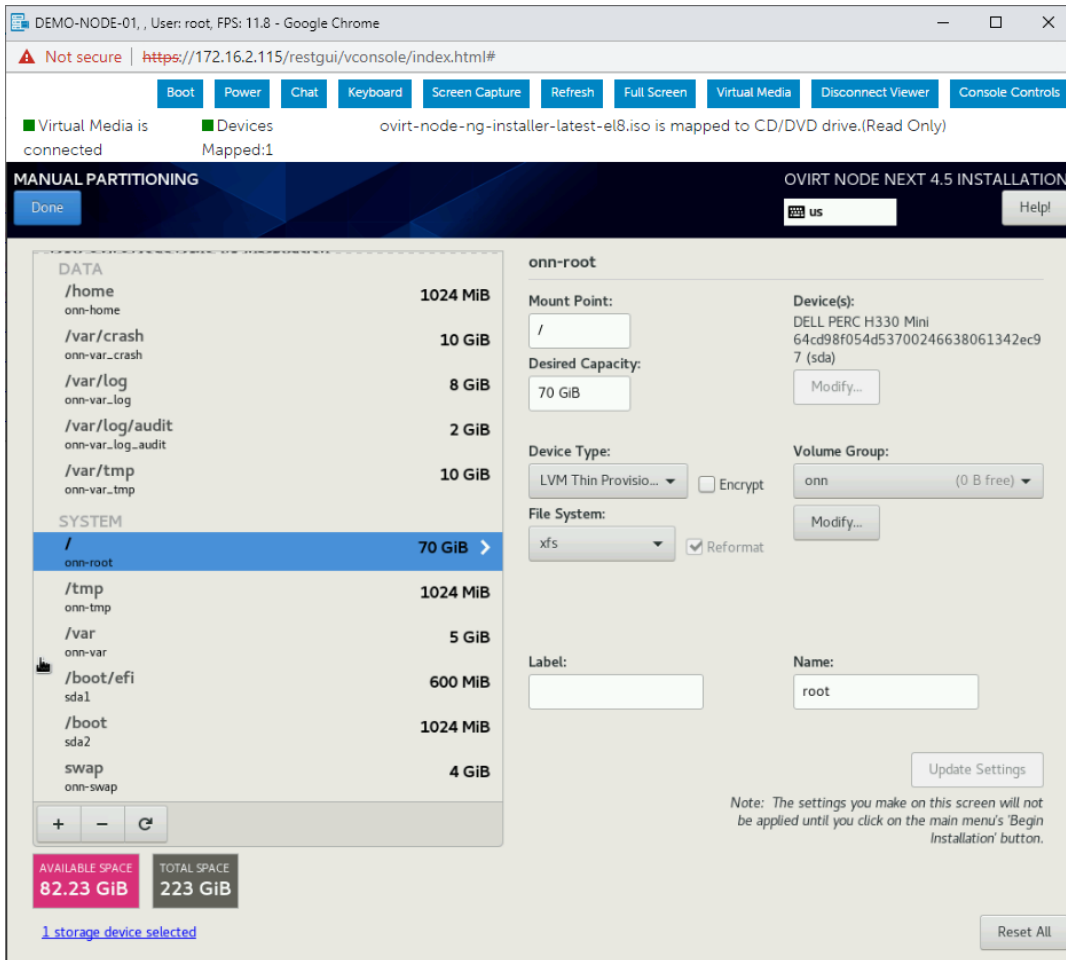
2. Boot from the downloaded ISO.



3. Choose “Installation Destination” and choose drive, which will be used for OS installation. Choose “Custom” storage configuration.



4. Click "Done" and partition your drive. At least 70 GiB should be left as available space. It will be used for StarWind VSAN Controller VM deployment.



5. Install oVirt Node on your host.

Preconfiguring Ovirt Hosts

1. Connect to the node via SSH or to a local console.
2. Download vHCI archive with StarWind CVM and configuration files to the oVirt node.

```
curl -k -O
https://tplink.starwind.com/StarWind_vHCI_KVM_scripts.zip
```

3. Unzip the archive.

```
mkdir StarWind_vHCI_KVM_scripts
unzip StarWind_vHCI_KVM_scripts.zip -d
StarWind_vHCI_KVM_scripts
```

3. Run node_predeploy.sh script.

```
cd ./StarWind_vHCI_KVM_scripts/
chmod +x node_predeploy.sh
./node_predeploy.sh
```

4. Type node number (01 or 02). Press Enter.

```
[root@sw-demo-node-01 StarWind_vHCI_KVM_scripts]# chmod +x node_predeploy.sh
[root@sw-demo-node-01 StarWind_vHCI_KVM_scripts]# ./node_predeploy.sh
Input node number (01 or 02)
01
```

5. Type the name of the Management interface. Press Enter.

```
NAME      UUID                                TYPE      DEVICE
eno3      5717bfd9-19e2-4c18-b2ec-5318a4107d76 ethernet  eno3
virbr0    faaa7902-2b26-4971-a3a7-85fd6935edaa bridge    virbr0
eno1      5aa0d20e-adf3-468a-9732-17ad27fe11aa ethernet  --
eno2      a14f45a8-b8d6-4c1b-ae8a-b56abbd054ba ethernet  --
eno4      796650d9-37ba-463a-a528-766301c6f8e5 ethernet  --
ens1f0    96773f81-946d-454b-bd6c-c036a7c26eb6 ethernet  --
ens1f1    203837de-9864-405b-8539-7c7ca5f49000 ethernet  --
#####
Input the name of Management Connection (e.g eno1, ens1f0)
#####
eno3
```

6.

Type the name of the Data interface. Press Enter.

```
NAME      UUID                                TYPE      DEVICE
ovirtmgmt 5c2b2e22-e3aa-43e4-90f0-4415ae792231 bridge    ovirtmgmt
virbr0    faaa7902-2b26-4971-a3a7-85fd6935edaa bridge    virbr0
bridge-slave-eno3 7535b8ea-7253-412f-8a9f-1d6b46c4e17e ethernet  eno3
eno1      5aa0d20e-adf3-468a-9732-17ad27fe11aa ethernet  --
eno2      a14f45a8-b8d6-4c1b-ae8a-b56abbd054ba ethernet  --
eno3      5717bfd9-19e2-4c18-b2ec-5318a4107d76 ethernet  --
eno4      796650d9-37ba-463a-a528-766301c6f8e5 ethernet  --
ens1f0    96773f81-946d-454b-bd6c-c036a7c26eb6 ethernet  --
ens1f1    203837de-9864-405b-8539-7c7ca5f49000 ethernet  --
#####
Input the name of Data Connection (e.g eno1, ens1f0)
#####
ens1f0
```

7. Type the name of the Replication interface. Press Enter.


```

NAME                UUID                                TYPE    DEVICE
iSCSI               14cbe1b0-7b69-4999-83c3-86bbfb39cd68 bridge  iSCSI
ovirtmgmt           5c2b2e22-e3aa-43e4-90f0-4415ae792231 bridge  ovirtmgmt
virbr0              faaa7902-2b26-4971-a3a7-85fd6935edaa bridge  virbr0
bridge-slave-eno3   7535b8ea-7253-412f-8a9f-1d6b46c4e17e ethernet eno3
bridge-slave-ensf10 f53ade06-0a28-4ae7-8766-7a286ea91e9c ethernet --
eno1                 5aa0d20e-adf3-468a-9732-17ad27fe11aa ethernet --
eno2                 a14f45a8-b8d6-4c1b-ae8a-b56abb054ba ethernet --
eno3                 5717bfd9-19e2-4c18-b2ec-5318a4107d76 ethernet --
eno4                 796650d9-37ba-463a-a528-766301c6f8e5 ethernet --
enslf0              96773f81-946d-454b-bd6c-c036a7c26eb6 ethernet --
enslf1              203837de-9864-405b-8539-7c7ca5f49000 ethernet --
#####
Input the name of Sync Connection (e.g eno1, enslf0)
#####
enslf1

```

8. If DHCP server is available on the network, IP of the StarWind CVM will be displayed in the script out put.

```

Domain 'CVM01' marked as autostarted
Domain 'CVM01' started

IP of the VM can be found here (if DHCP server is available)
{"return":{"name":"lo","ip-addresses":[{"ip-address-type":"ipv4","ip-address":"127.0.0.1","prefix":8}],"statistics":{"tx-packets":0,"tx-errs":0,"rx-bytes":0,"rx-dropped":0,"rx-packets":0,"rx-errs":0,"tx-bytes":0,"tx-dropped":0},"hardware-address":"00:00:00:00:00:00"},"name":"eth0","ip-addresses":[{"ip-address-type":"ipv4","ip-address":"172.16.2.38","prefix":24}],"statistics":{"tx-packets":39,"tx-errs":0,"rx-bytes":10512,"rx-dropped":13,"rx-packets":124,"rx-errs":0,"tx-bytes":5005,"tx-dropped":0},"hardware-address":"52:54:00:e6:34:57"},"name":"eth1","hardware-address":"52:54:00:8b:6e:f8"},"name":"eth2","hardware-address":"52:54:00:c5:a2:28"},"name":"eth3","hardware-address":"52:54:00:57:66:d4"},"name":"eth4","hardware-address":"52:54:00:1b:a0:aa"}}

```

Note. If DHCP is not available, static IP address should be configured.

9. Repeat steps 1-8 on every node.

Configuring Starwind Vsan Controller Vm

1. Login to node via SSH.
2. Check VM display using the following command.

```
virsh -c qemu:///system?authfile=/etc/ovirt-hosted-engine/virsh_auth.conf domdisplay CVM01
```

```

[root@sw-demo-node-01 ~]# virsh -c qemu:///system?authfile=/etc/ovirt-hosted-engine/virsh_auth.conf domdisplay cvm01
spice://127.0.0.1:5906?tls-port=5907
[root@sw-demo-node-01 ~]#

```

NOTE. Depending on the output you should connect to the VM display via Spice or VNC console.

3. Create tunnel to the host using the following command using SSH client.

```
ssh -L 5906:127.0.0.1:5906 root@%hostipaddress%
```

4. Connect to VM using Virt-Viewer and login using the following credentials.

- Username: user

- Password: rds123RDS

Note. Virt-Manager can be downloaded here:

<https://virt-manager.org/download.html>

5. Open the file corresponding to the Management interface using a text editor, for example: `sudo nano /etc/sysconfig/network-scripts/ifcfg-eth0`

6. Edit the file:

Change the line `BOOTPROTO=dhcp` to: `BOOTPROTO=static`

Add the IP settings needed to the file:

`IPADDR=192.168.12.10`

`NETMASK=255.255.255.0`

`GATEWAY=192.168.12.1`

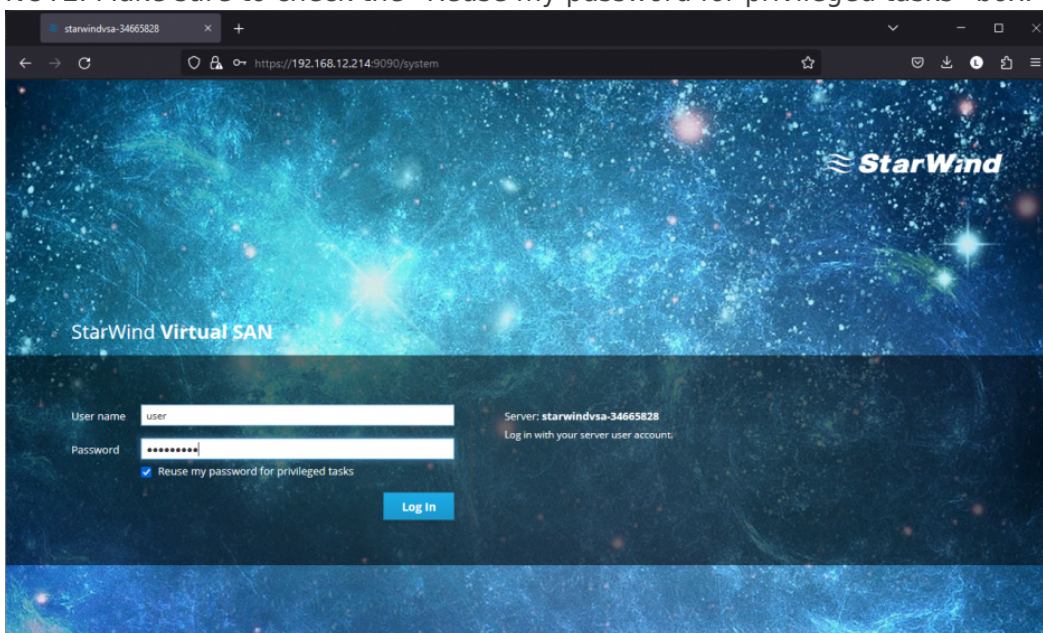
`DNS1=192.168.1.1`

7. Now, open the web browser and enter the IP address of the VM. Log into the VM using the following default credentials:

- Username: user

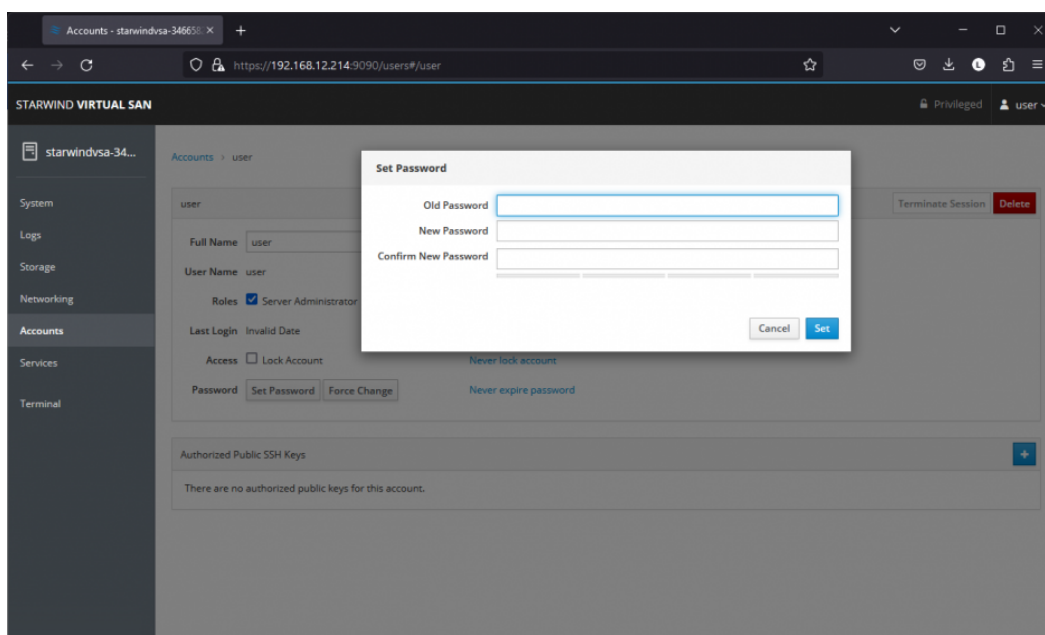
- Password: rds123RDS

NOTE: Make sure to check the “Reuse my password for privileged tasks” box.

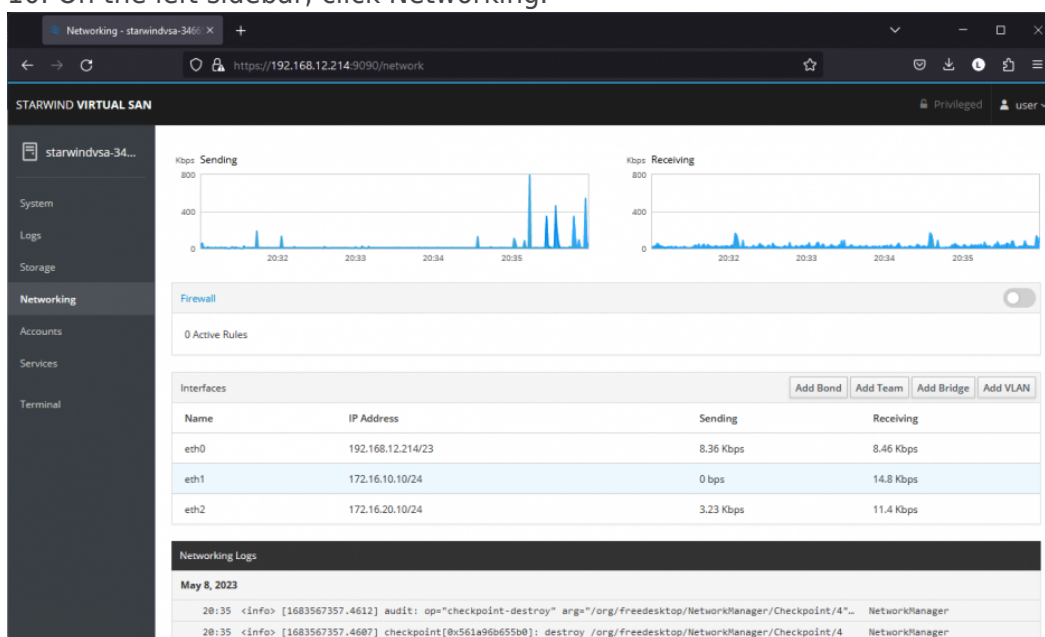


8. After a successful login, click Accounts on the left sidebar.

9. Select a user and click Set Password.



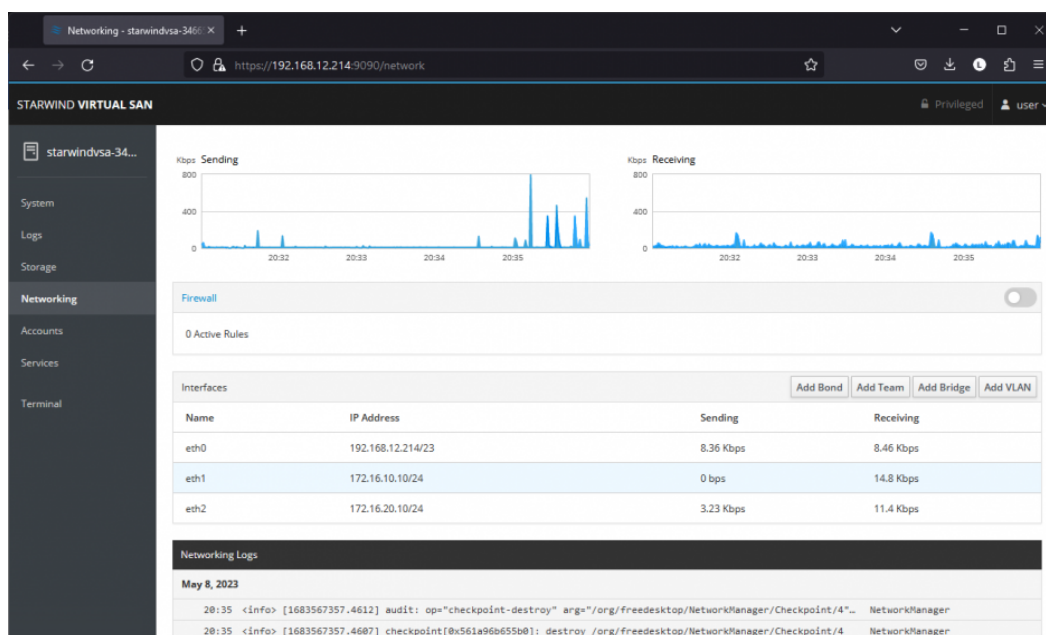
10. On the left sidebar, click Networking.



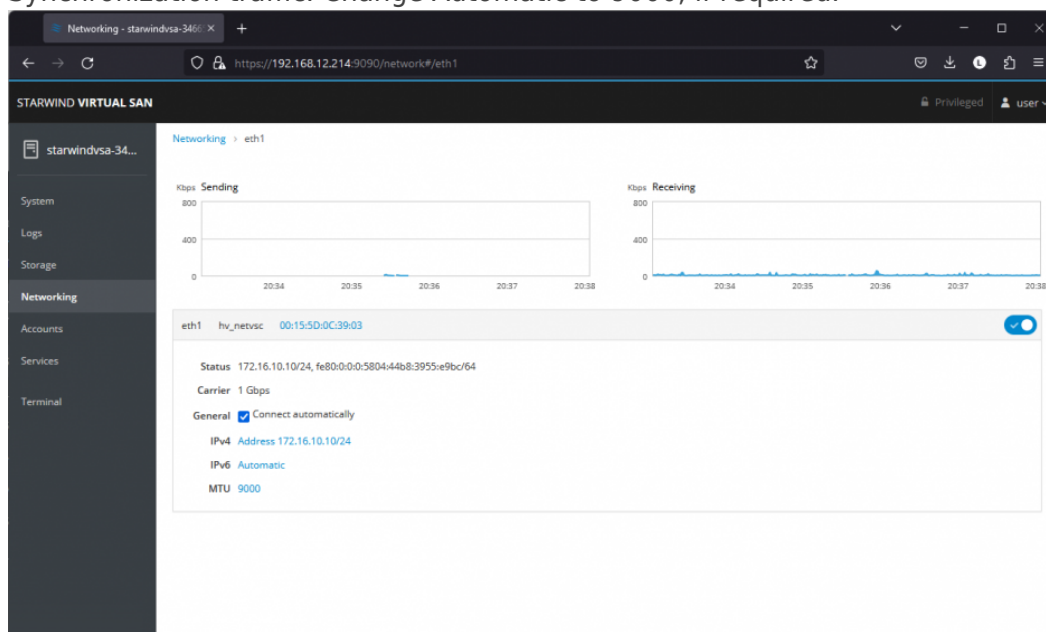
Here, the Management IP address of the StarWind Virtual SAN Virtual Machine can be configured, as well as IP addresses for iSCSI and Synchronization networks. In case the Network interface is inactive, click on the interface, turn it on, and set it to Connect automatically.

11. Configure IP addresses for each interface.

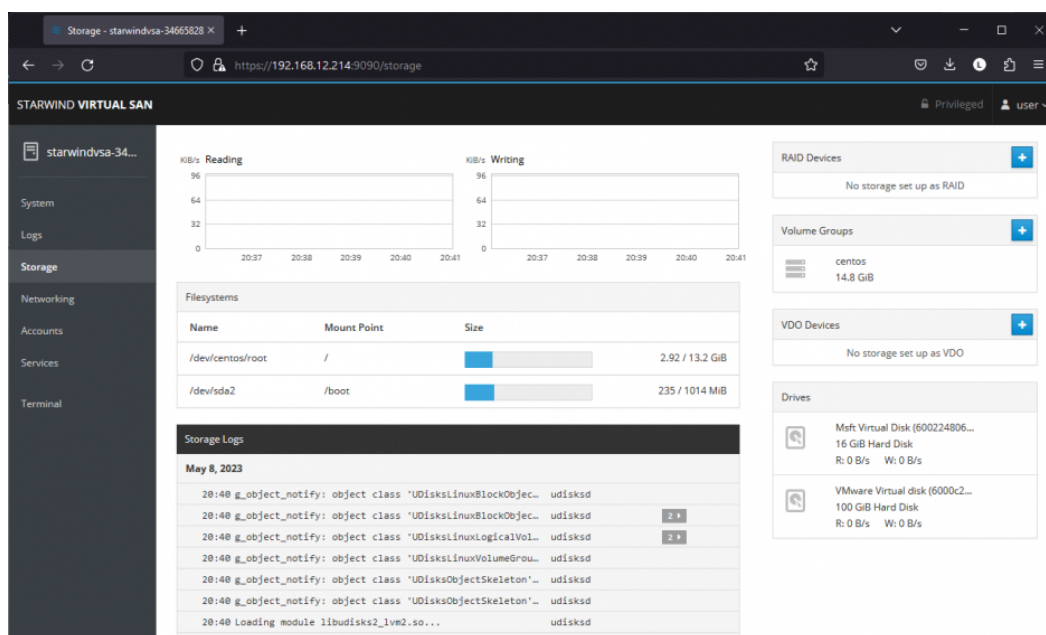
12. The result should look like in the picture below:



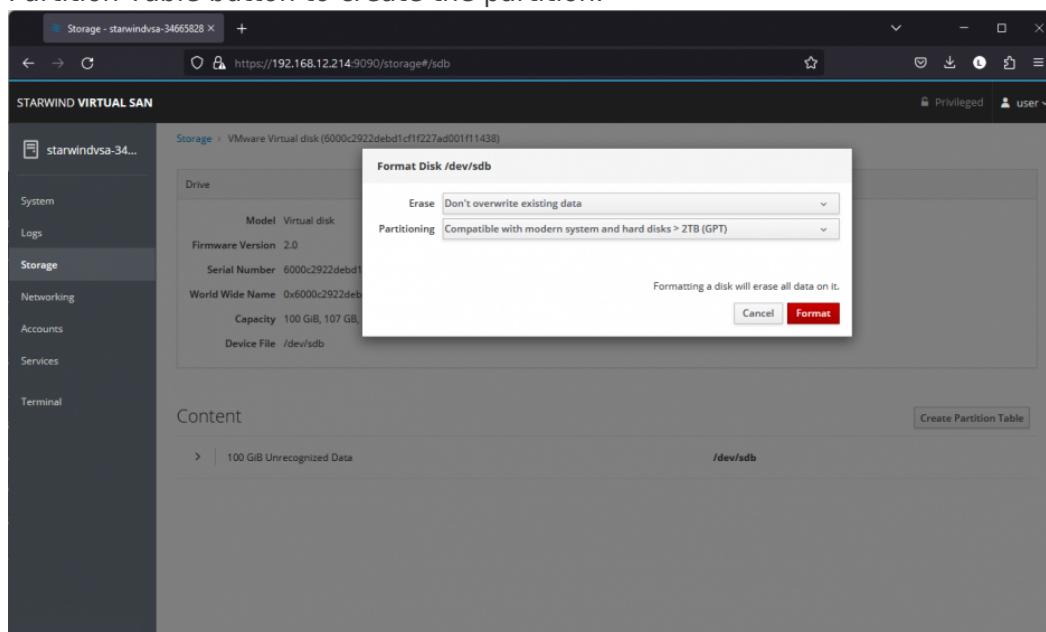
NOTE: It is recommended to set MTU to 9000 on interfaces dedicated for iSCSI and Synchronization traffic. Change Automatic to 9000, if required.



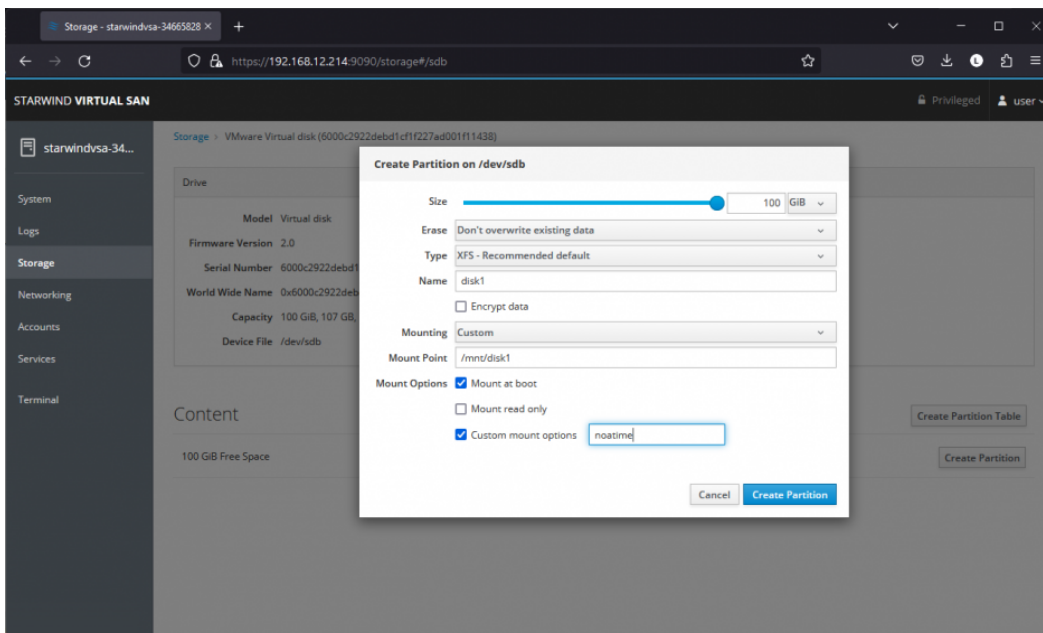
13. Access the Storage section. Locate the recently added disk in the Drives section and choose it.



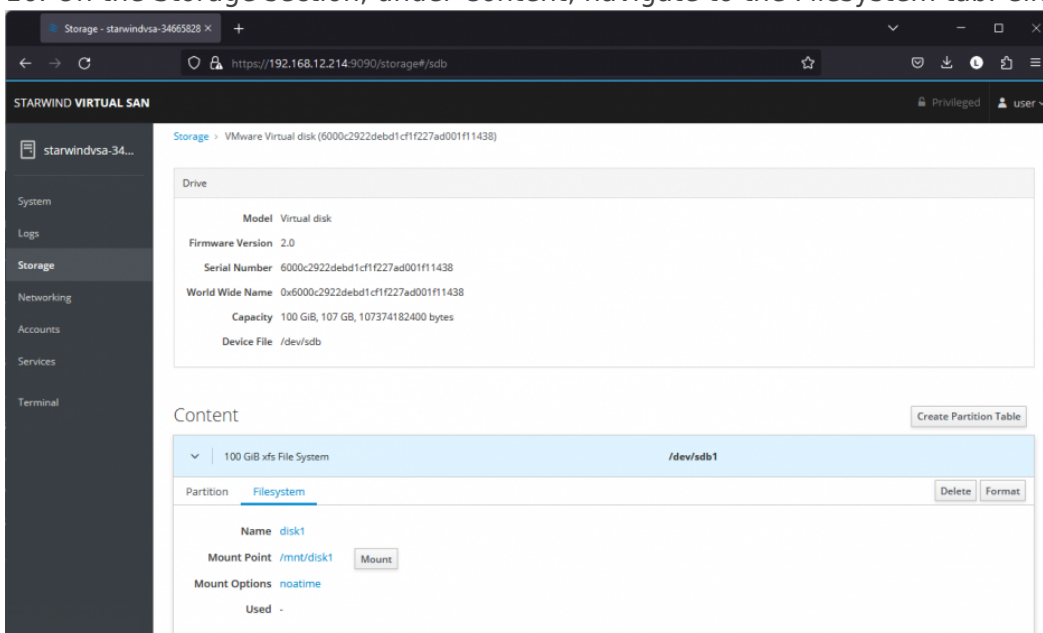
14. The added disk does not have any partitions and filesystem. Press the Create Partition Table button to create the partition.



15. Press Create Partition to format the disk and set the mount point. The mount point should be as follows: /mnt/%yourdiskname%



16. On the Storage section, under Content, navigate to the Filesystem tab. Click Mount.



17. Repeat steps 1-16 on every node.

Configuring Starwind Management Console

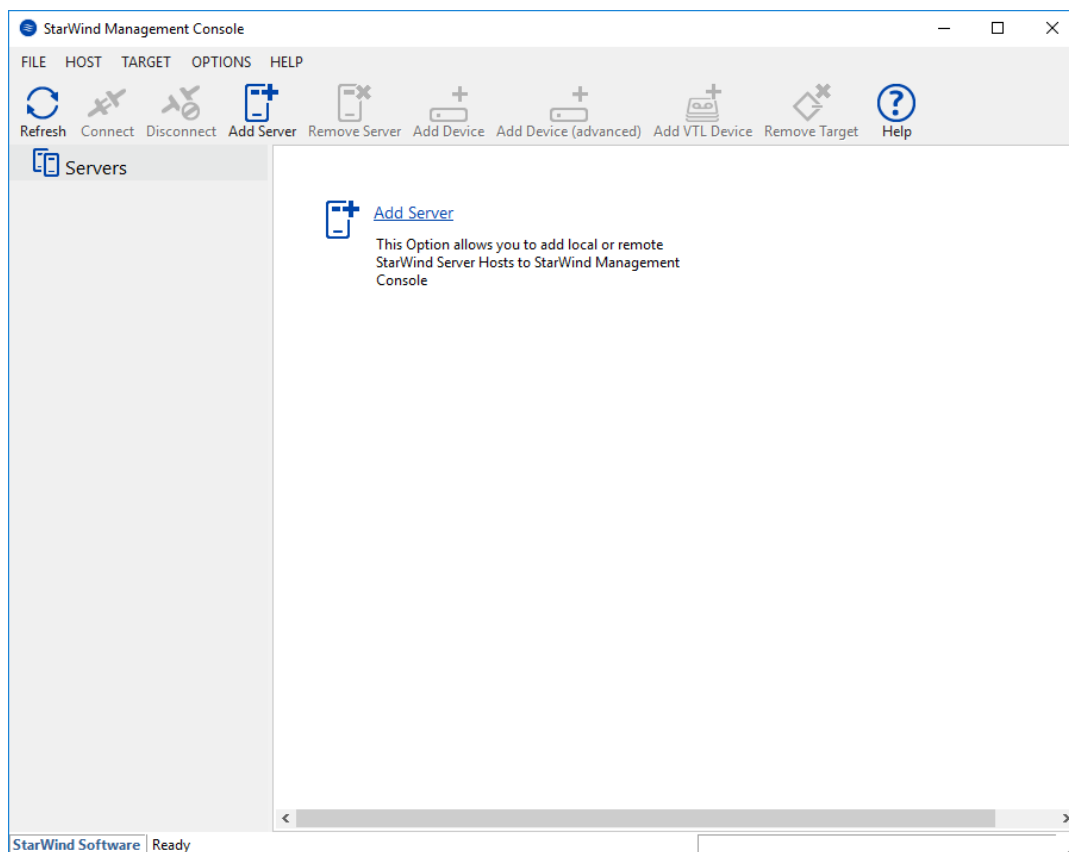
1. Install StarWind Management Console on each server or on a separate workstation with Windows OS (Windows 7 or higher, Windows Server 2008 R2 and higher) using the installer available [here](#).

NOTE: StarWind Management Console and PowerShell Management Library components

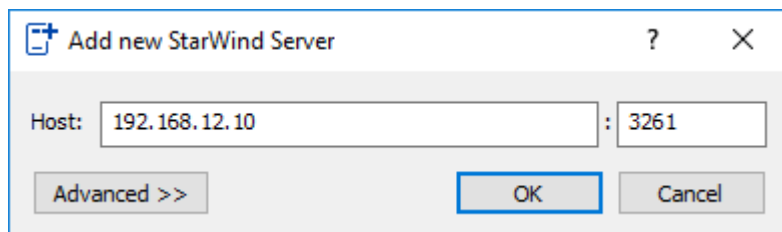
are required.

2. Select the appropriate option to apply the StarWind License key.
Once the appropriate license key has been received, it should be applied to StarWind Virtual SAN service via Management Console or PowerShell.

3. Open StarWind Management Console and click Add Server.

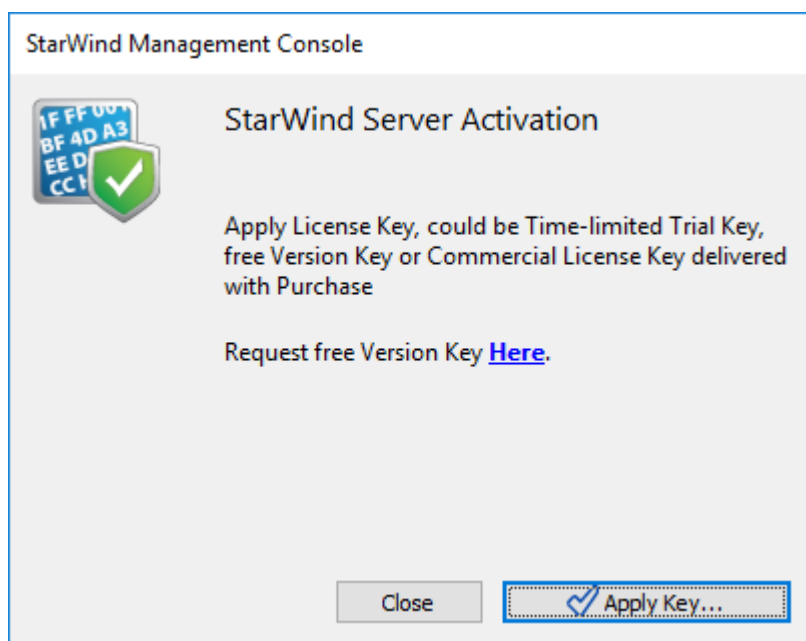


4. Type the IP address of the StarWind Virtual SAN in the pop-up window and click OK.



5. Select the server and click Connect.

6. Click Apply Key... on the pop-up window.



7. Select Load license from file and click the Load button.

8. Select the appropriate license key.

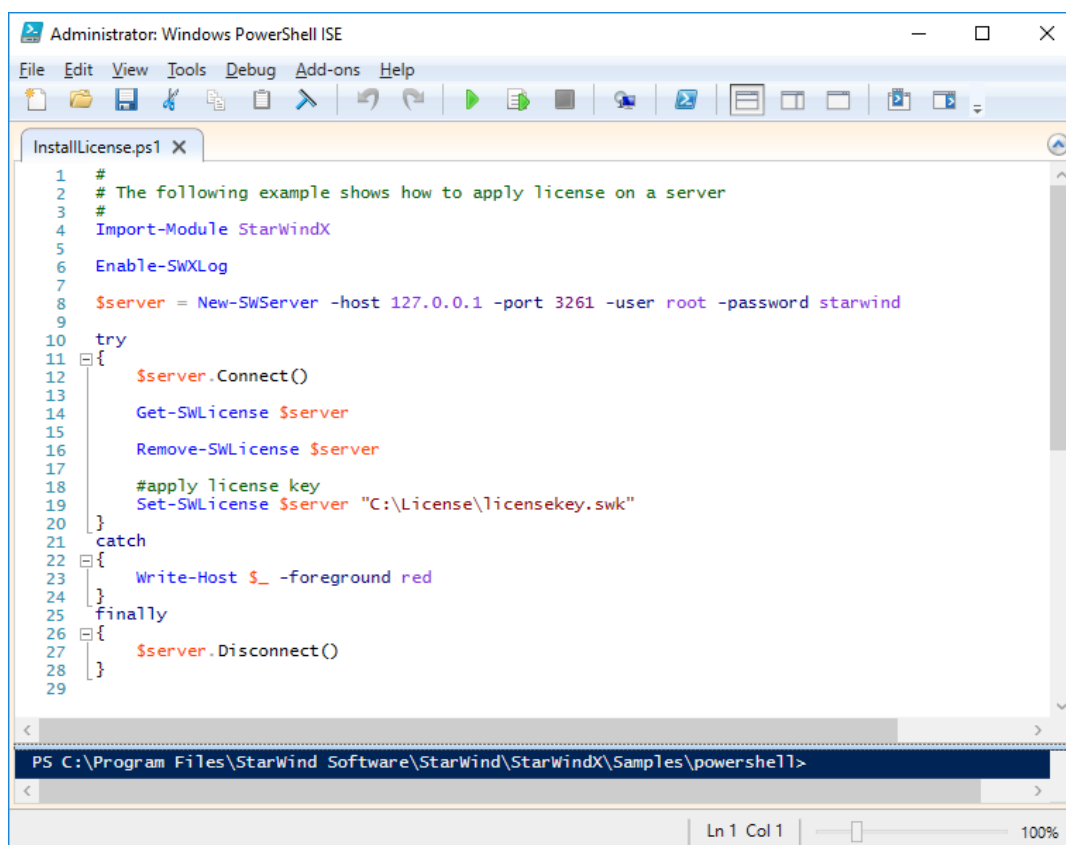
As an alternative, PowerShell can be used. Open StarWind InstallLicense.ps1 script with PowerShell ISE as administrator. It can be found here:

C:\Program Files\StarWind

Software\StarWind\StarWindX\Samples\powershell\InstallLicense.ps1

Type the IP address of StarWind Virtual SAN VM and credentials of StarWind Virtual SAN service (defaults login: root, password: starwind).

Add the path to the license key.



```

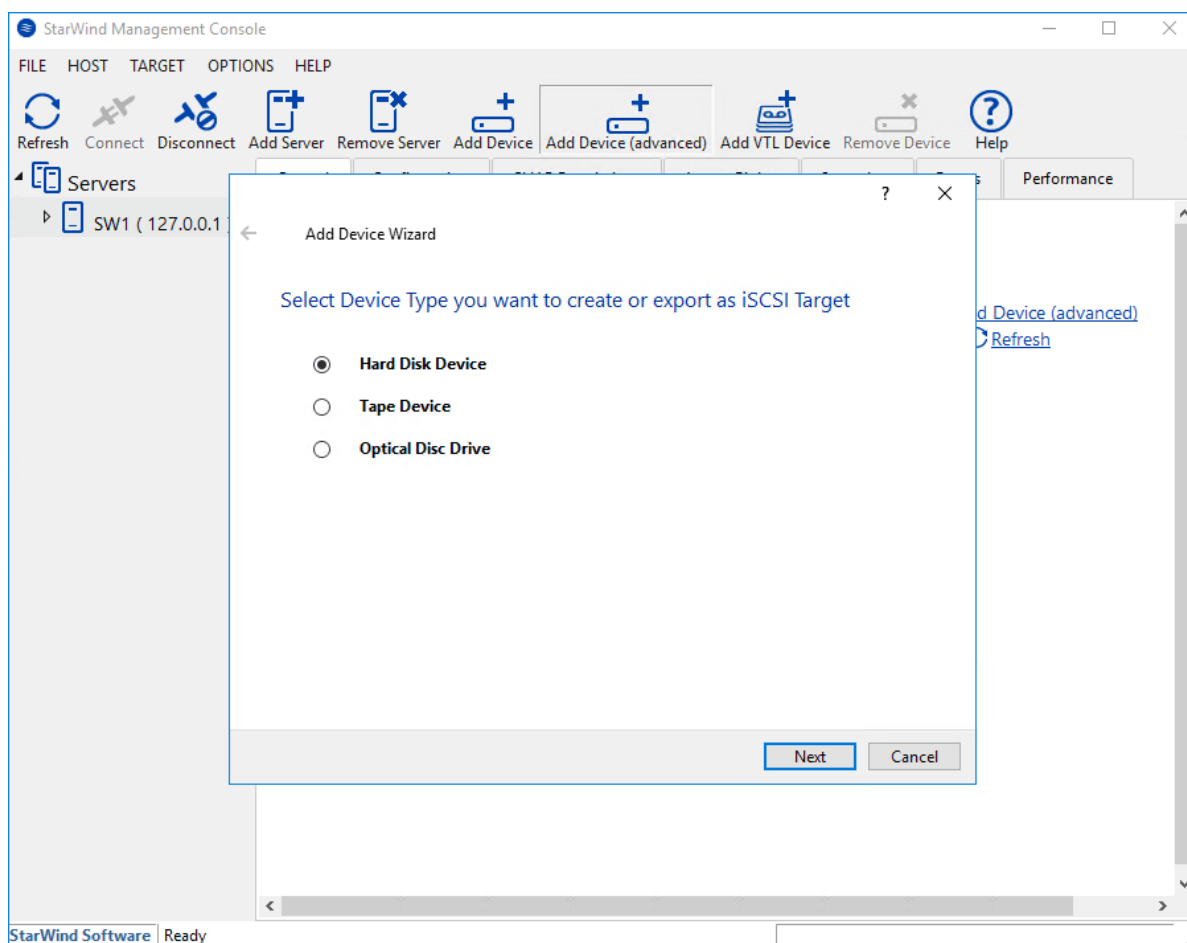
Administrator: Windows PowerShell ISE
File Edit View Tools Debug Add-ons Help
InstallLicense.ps1 X
1 #
2 # The following example shows how to apply license on a server
3 #
4 Import-Module StarWindX
5
6 Enable-SWXLog
7
8 $server = New-SWServer -host 127.0.0.1 -port 3261 -user root -password starwind
9
10 try
11 {
12     $server.Connect()
13
14     Get-SWLicense $server
15
16     Remove-SWLicense $server
17
18     #apply license key
19     Set-SWLicense $server "C:\License\licensekey.swk"
20 }
21 catch
22 {
23     Write-Host $_ -foreground red
24 }
25 finally
26 {
27     $server.Disconnect()
28 }
29
PS C:\Program Files\StarWind Software\StarWind\StarWindX\Samples\powershell>
Ln 1 Col 1 100%
  
```

9. After the license key is applied, StarWind devices can be created.
NOTE: In order to manage StarWind Virtual SAN service (e.g. create ImageFile devices, VTL devices, etc.), StarWind Management Console can be used.

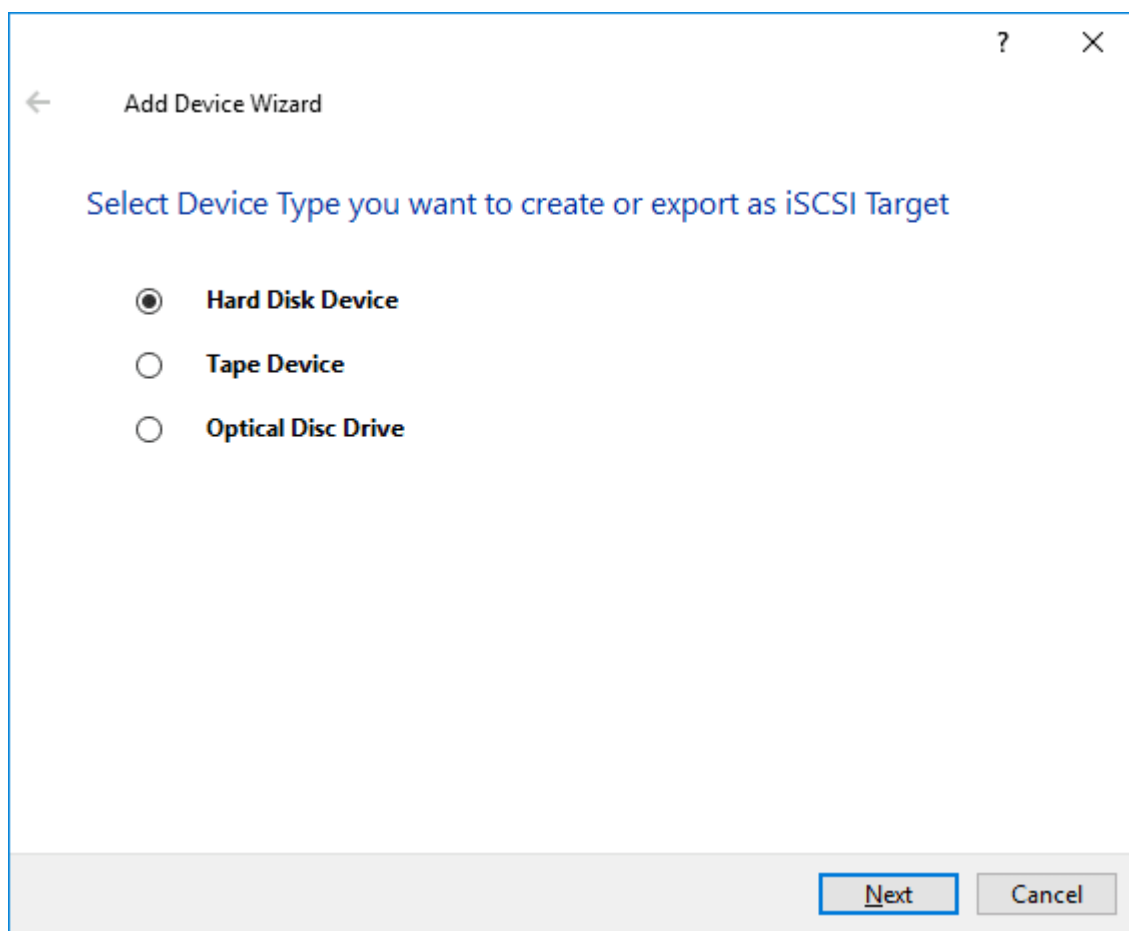
Creating Starwind Devices

Note. oVirt hosted engine requires shared storage to be deployed. Create StarWind HA device for hosted engine at this stage. HA devices for VMs should be created after hosted engine deployment.

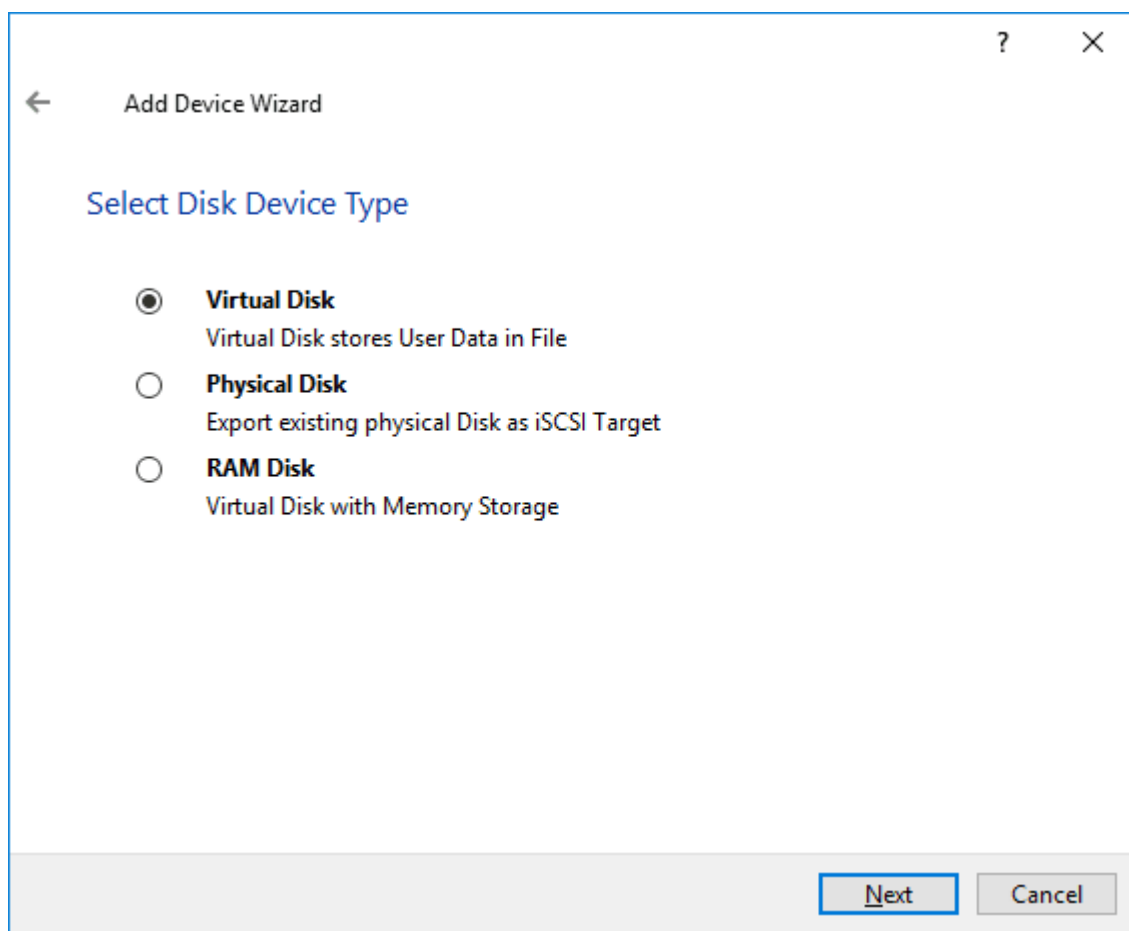
1. In the StarWind Management Console click to Add Device (advanced) button and open Add Device (advanced) Wizard.



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.

← Add Device Wizard

Virtual Disk Location

☒ Create a New Virtual Disk

Name:

Location:

Size: GB

☐ Use an Existing Virtual Disk

Location:

☐ Read-Only Mode

Next Cancel

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).

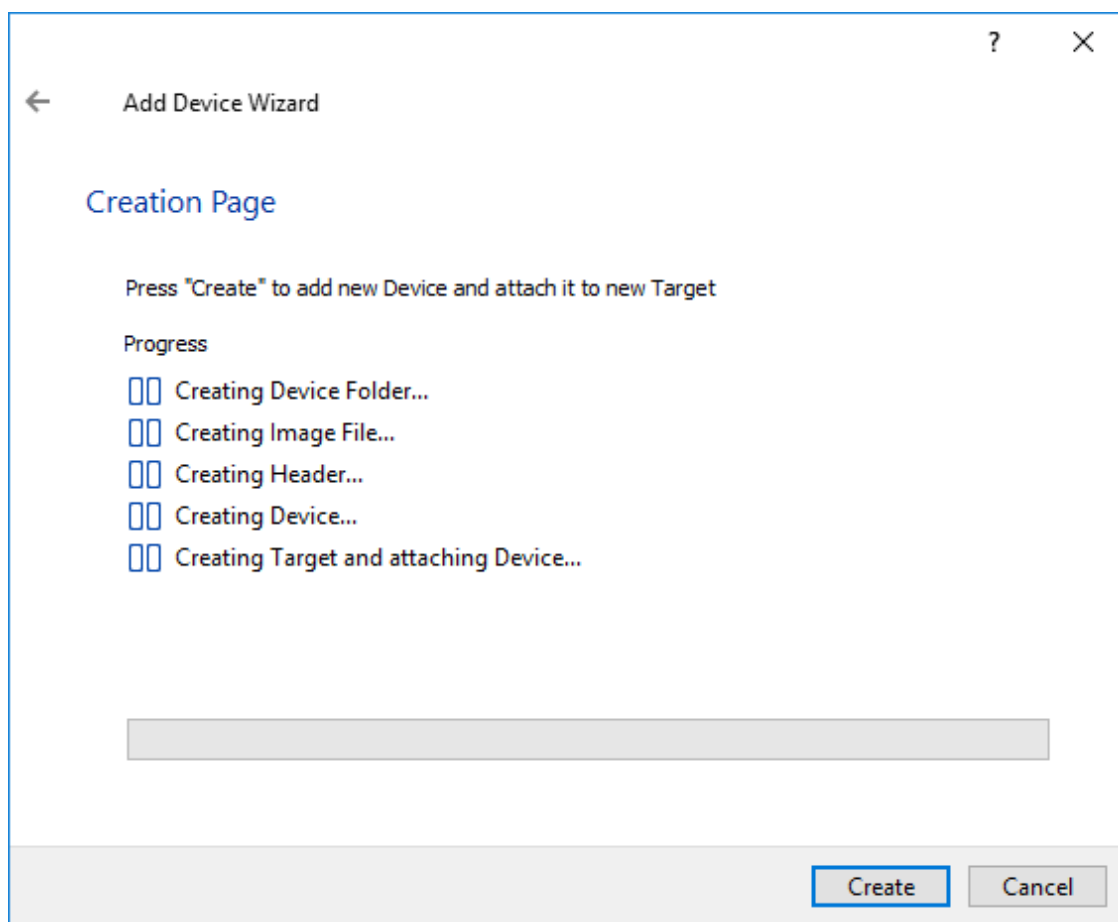
6. Define a caching policy and specify a cache size (in MB). Also, the maximum available cache size can be specified by selecting the appropriate checkbox. Optionally, define the L2 caching policy and cache size.

The screenshot shows a window titled "Add Device Wizard" with a back arrow and help/close icons. The main heading is "Specify Device RAM Cache Parameters". Below this is a "Mode" section with three radio button options: "Write-Back" (described as asynchronous writes, delayed disk writes, cached reads), "Write-Through" (described as synchronous writes, cached reads), and "N/A" (selected, described as no caching for reads or writes). Below the modes is a checkbox for "Set Maximum available Size" which is unchecked. At the bottom left is a "Size:" label followed by a text input containing "128" and a unit dropdown menu currently showing "MB". At the bottom right are "Next" and "Cancel" buttons.

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.

The screenshot shows a window titled "Add Device Wizard" with a back arrow and help/close icons. The main section is "Target Parameters". It contains a dropdown menu "Choose a Target Attachment Method" with "Create new Target" selected. Below it is a text field "Target Alias" with the placeholder "<target alias name>". There is an unchecked checkbox "Target Name" and a corresponding text field containing "iqn.2008-08.com.starwindsoftware:sw1- <target alias name>". At the bottom, there is a checked checkbox "Allow multiple concurrent iSCSI Connections". The "Next" button is highlighted with a blue border, and a "Cancel" button is also present.

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



9. Click Close to finish the device creation.

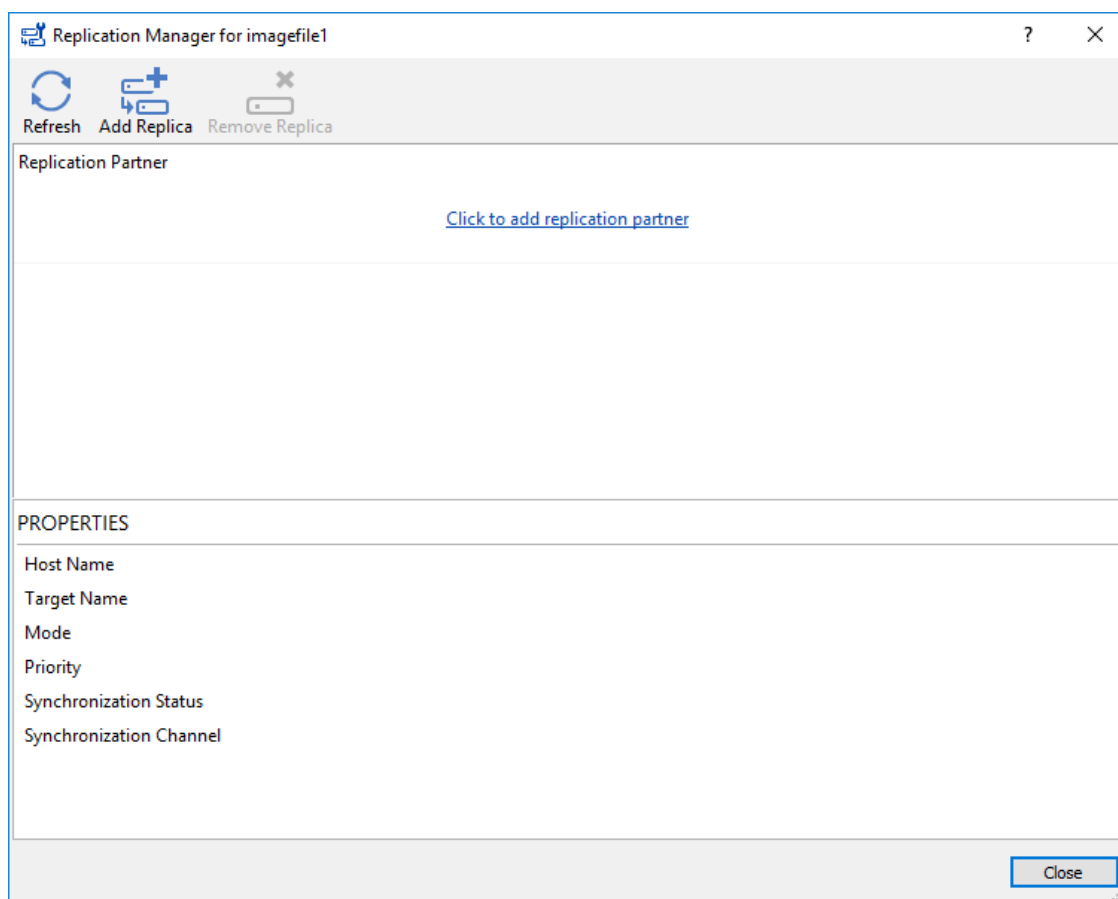
10. The successfully added devices appear in the StarWind Management Console.

Select The Required Replication Mode

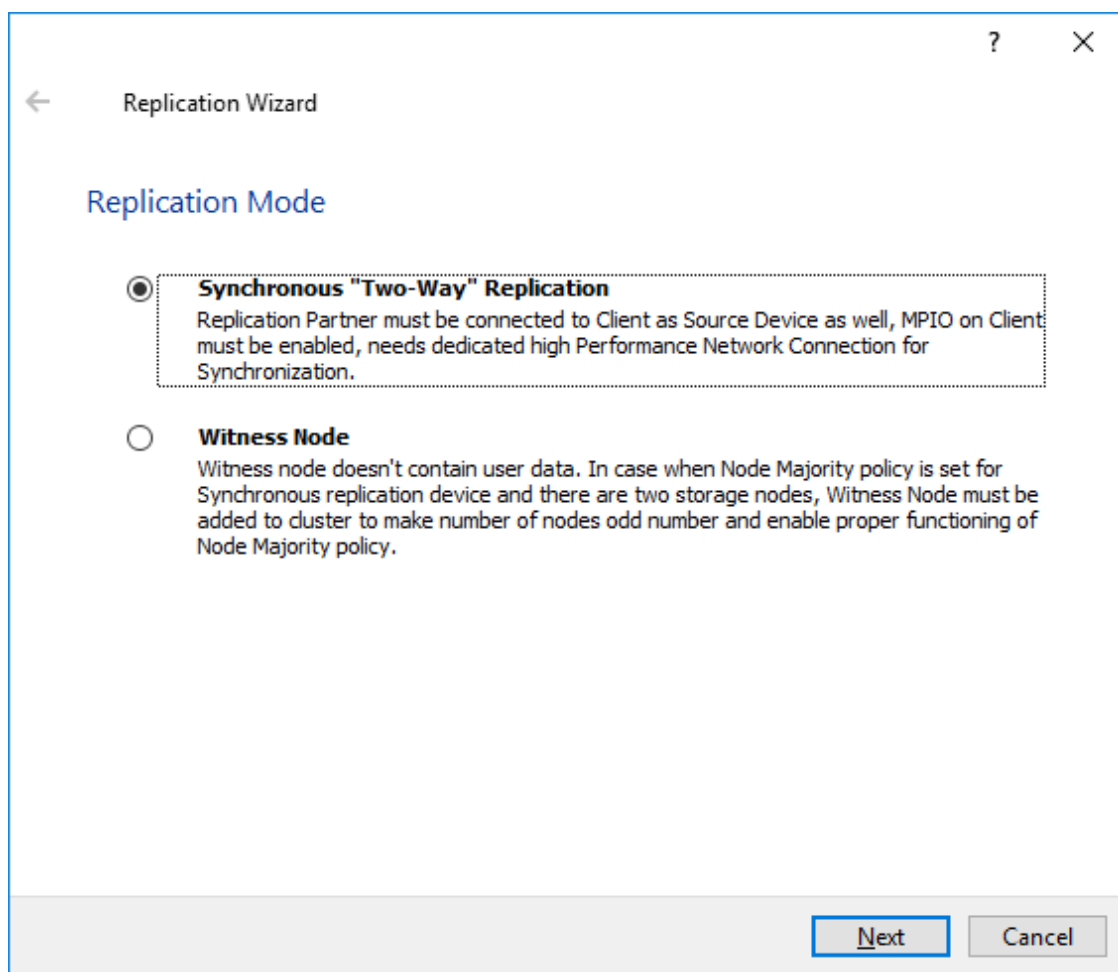
The replication can be configured using Synchronous “Two-Way” Replication mode: Synchronous or active-active replication ensures real-time synchronization and load balancing of data between two or three cluster nodes. Such a configuration tolerates the failure of two out of three storage nodes and enables 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”.

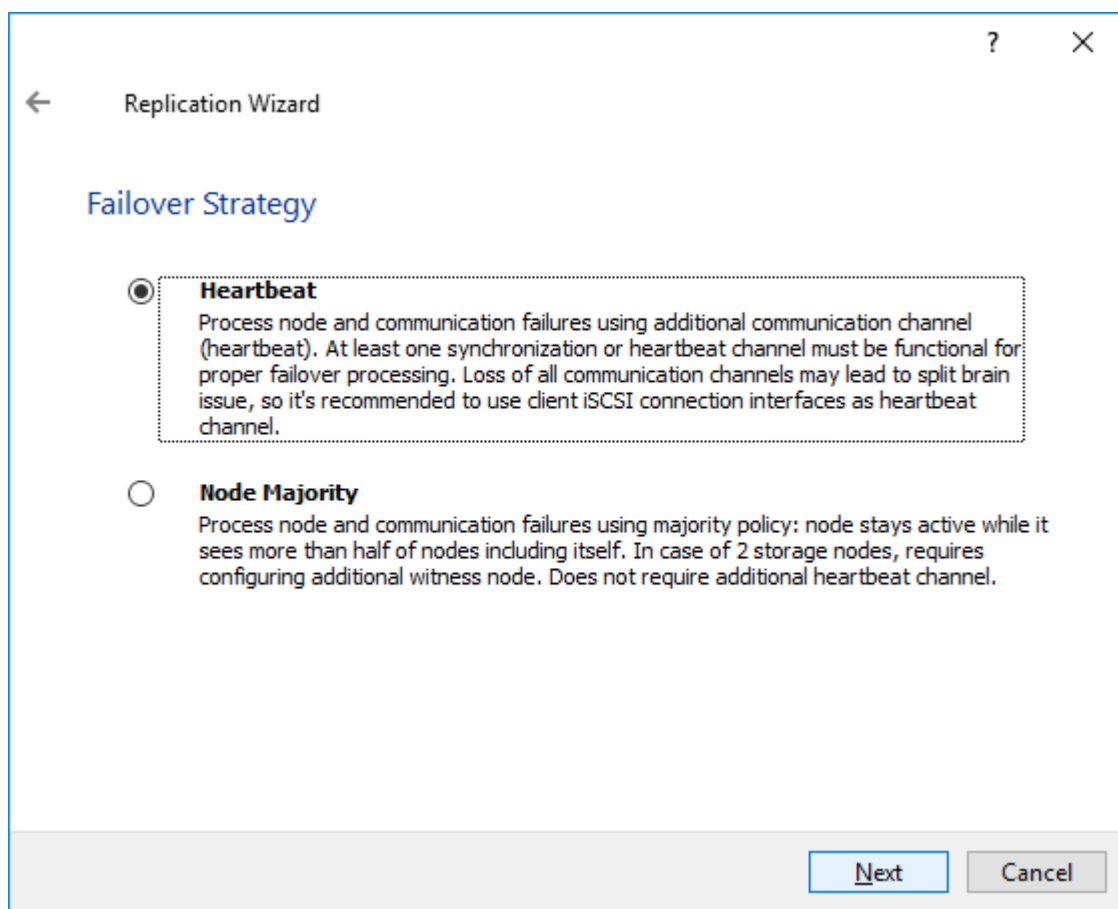
In case of a two-node HA storage, all nodes will be disconnected if there is a problem on the node itself, or in communication between them. Therefore, the Node Majority failover strategy requires the addition of the third Witness node or file share (SMB) which participates in the nodes count for the 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.

Replication Wizard

Synchronization Journal Setup

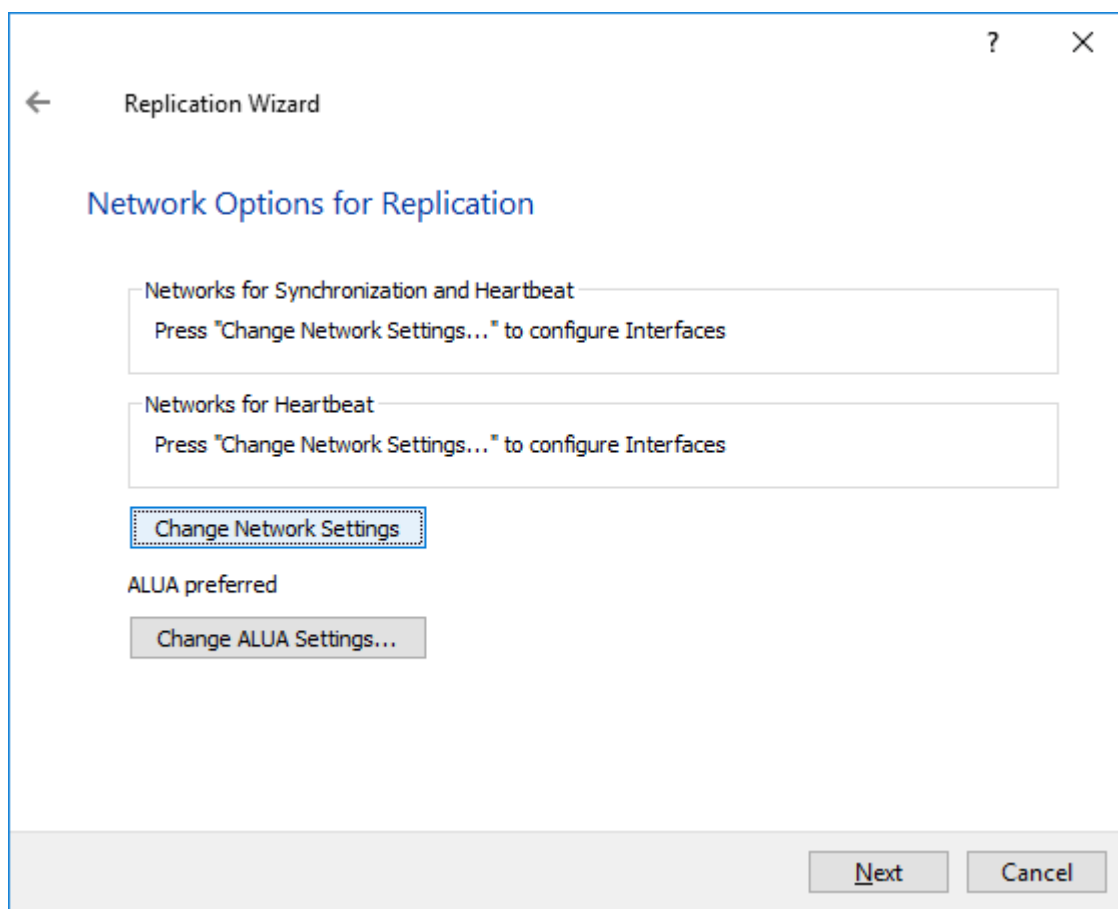
- ☒ **RAM-based journal**
Synchronization journal placed in RAM. Synchronization with RAM journal provides good IO performance in any scenario.
- ☐ **Disk-based journal**
Synchronization journal placed on disk.
- ☐ **Failure journal**
The strategy provides good IO performance while all device nodes are in a healthy state.
- ☐ **Continuous journal**
The strategy guarantees fast synchronization and data consistency in all cases.

Current Node: ...

Partner Node: ...

[Next](#) [Cancel](#)

5. Click Change Network Settings.



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

Specify Interfaces for Synchronization Channels

Select synchronization channel

Interfaces	Networks	Synchronization and H...	Heartbeat
Host Name: 127.0.0.1			
172.16.10.10	172.16.10.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.20.10	172.16.20.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
192.168.12.10	192.168.12.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Host Name: SW2			
172.16.10.20	172.16.10.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.20.20	172.16.20.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
192.168.12.20	192.168.12.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>

☐ Allow Free Select Interfaces

OK Cancel

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 same 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:

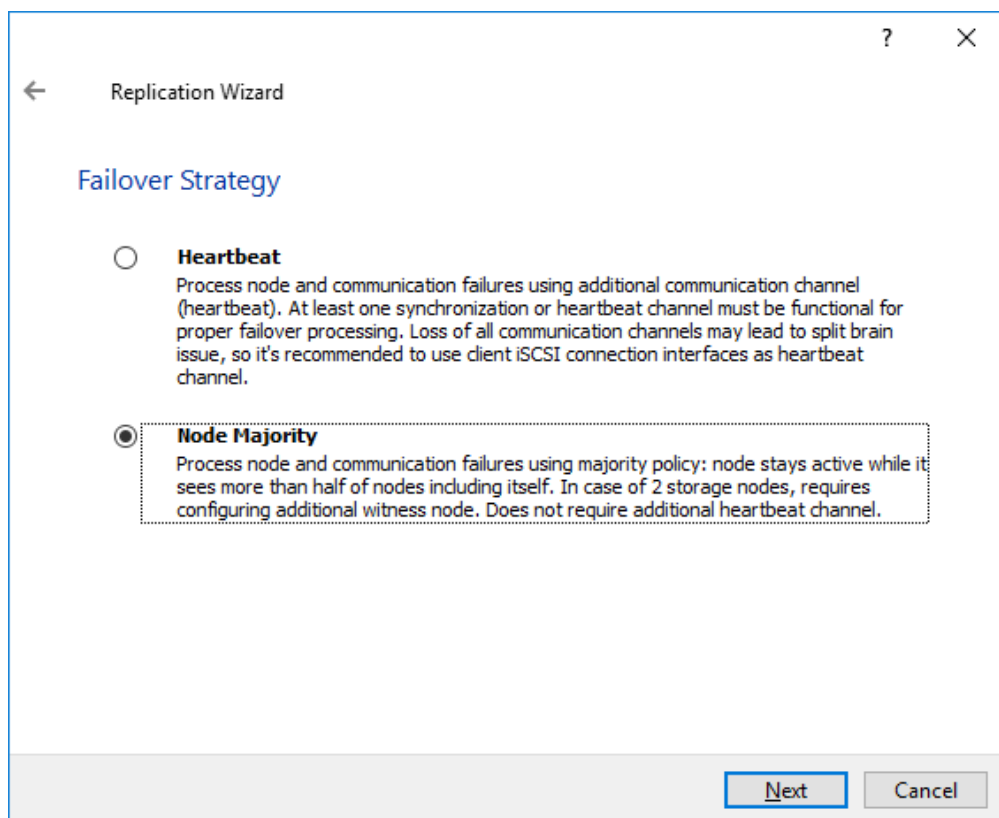
SMB Witness is a file, located on SMB share, which can be accessed by both nodes and help them to eliminate the split-brain issue in case of synchronization connection

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.

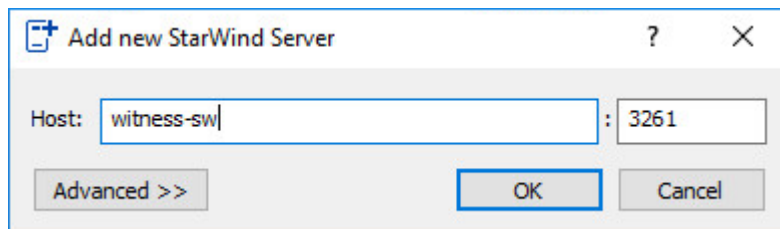
7. Press the Create Replica button and close the wizard.
8. The added devices will appear in StarWind Management Console. Repeat the steps above to create other virtual disks if necessary.

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.

Replication Wizard

Replication Mode

- ☐ **Synchronous "Two-Way" Replication**
Replication Partner must be connected to Client as Source Device as well, MPIO on Client must be enabled, needs dedicated high Performance Network Connection for Synchronization
- ☐ **Asynchronous "One-Way" Replication**
Replica is used to store replicated Data, Data is stored as Snapshots, Client cannot connect to Replication Partner, mount Snapshot from Replica to get Access to replicated Data
- ☒ **Witness Node**
Witness node doesn't contain user data. In case when Node Majority policy is set for Synchronous replication device and there are two storage nodes, Witness Node must be added to cluster to make number of nodes odd number and enable proper functioning of Node Majority policy.

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.

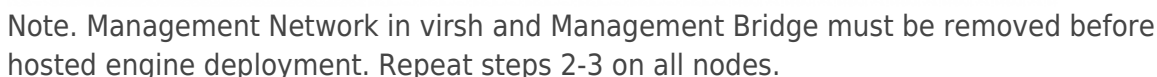
9. Repeat the steps above to create other virtual disks if necessary.

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/>

1. Login to oVirt Node, which will be used for self-hosted engine deployment, via SSH. Note. Hosted engine will be deployed as a VM on a shared shared storage of the node chosen for the deployment. Second node should be added separately to oVirt Engine.
2. Remove management network from virsh.

3. Login to oVirt node via web. Go to Networking page and choose management bridge. Click Delete.

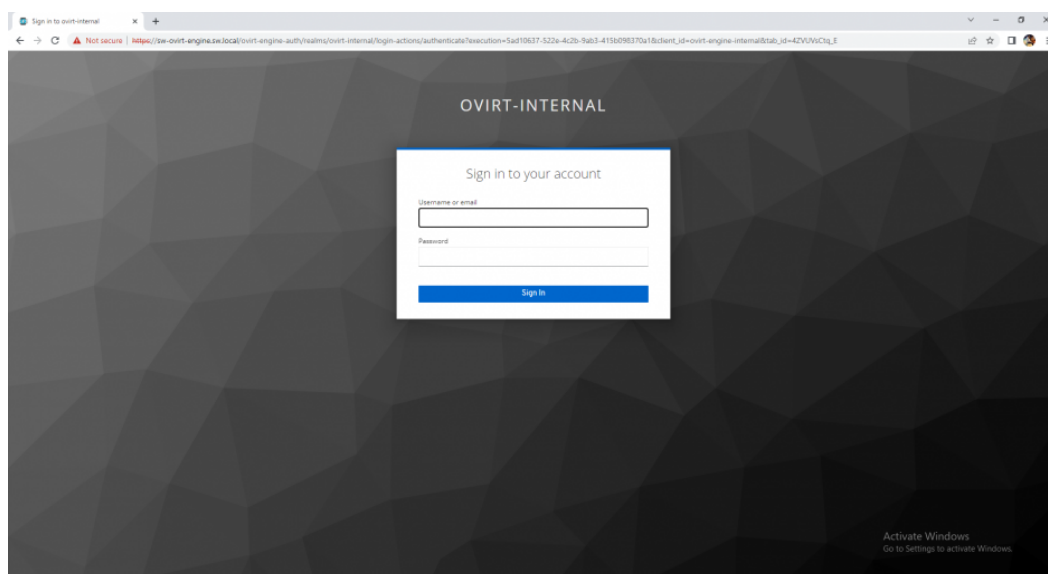


```
dnf -y install ovirt-engine-appliance --nogpgcheck
hosted-engine --deploy --4
```

Note. DNS records of the oVirt Node and oVirt Engine should be added before oVirt engine deployment. It is required to use static IP addresses for both oVirt node and oVirt Engine.

6. Login to oVirt Engine after successful installation deployment using the following credentials:

Password: Configured during oVirt engine deployment.



7. After deployment connect to all paths of iSCSI LUN on every node.

```
iscsiadm -m discovery -t st -p 172.16.10.10
iscsiadm -m node -l -T iqn.2008-08.com.starwindsoftware:cvm01-lun01
iscsiadm -m discovery -t st -p 172.16.10.20
iscsiadm -m node -l -T iqn.2008-08.com.starwindsoftware:cvm02-lun01
```

Adding Hosts To Ovirt Engine

1. Login to oVirt Engine and go to Compute -> Hosts.
2. Type Name, Hostname or IP, Password of the host, which will be added. Click OK.

3. Repeat steps 1-2 for all hosts, which were not added to oVirt engine.

Defining Management Network For Starwind Cvm

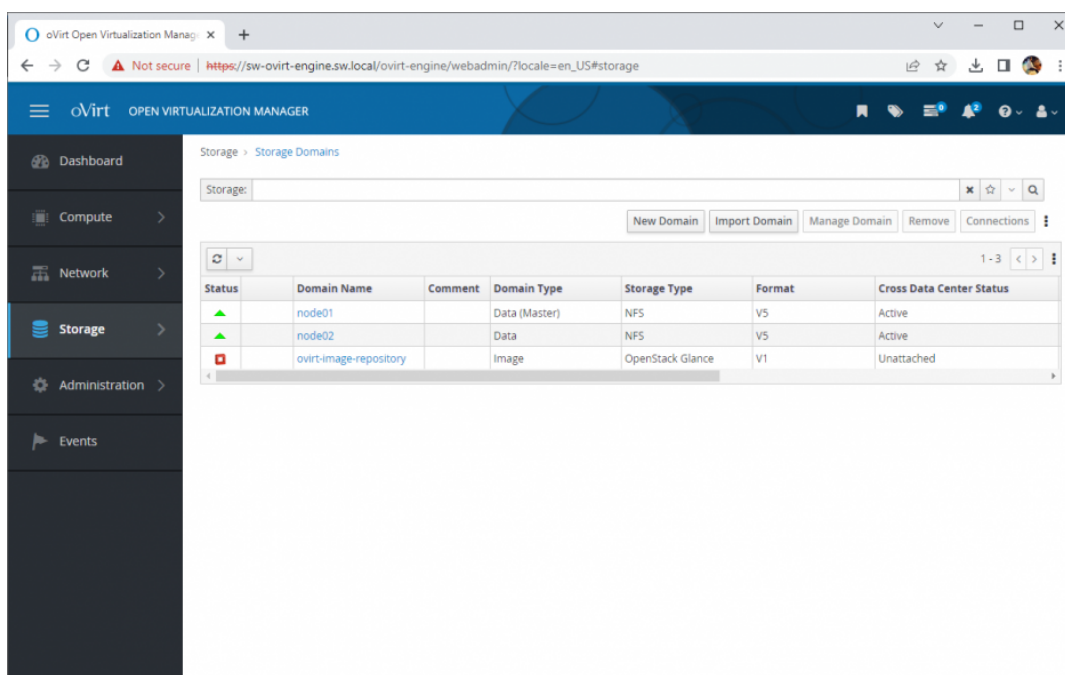
1. Login to the node via SSH
2. Attach Network interface to StarWind VM.

```
virsh -c qemu:///system?authfile=/etc/ovirt-hosted-engine/virsh_auth.conf attach-interface --domain cvm01 --type bridge --source vdsm-ovirtmgmt --target ovirtmgmt --model virtio --config --live
```

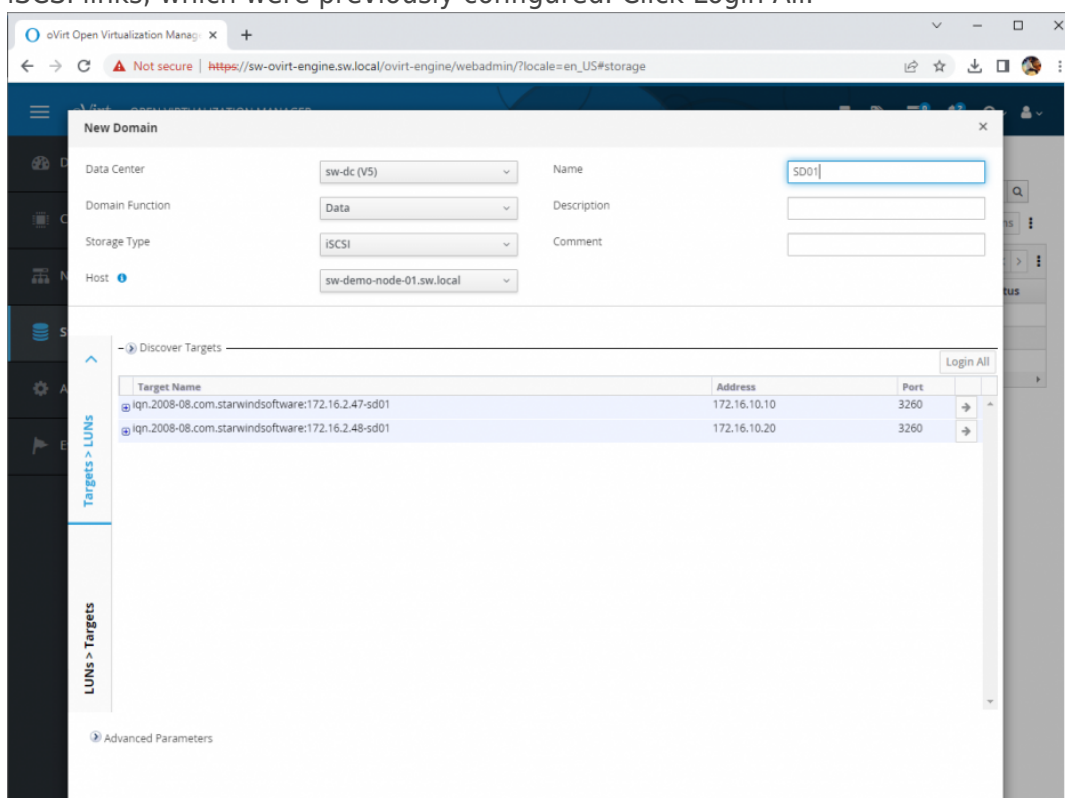
3. Repeat steps 1-2 on all nodes.

Provisioning Starwind Ha Storage To Hosts

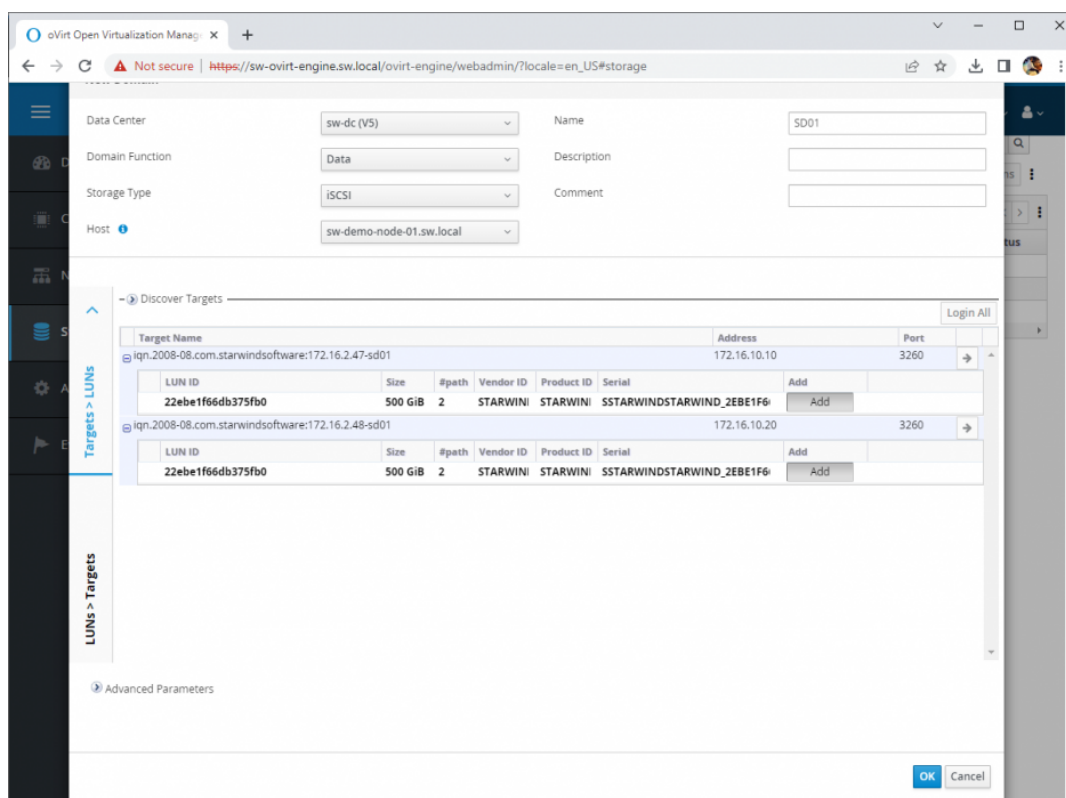
1. Login to Engine and open Storage -> Domain. Click New Domain.



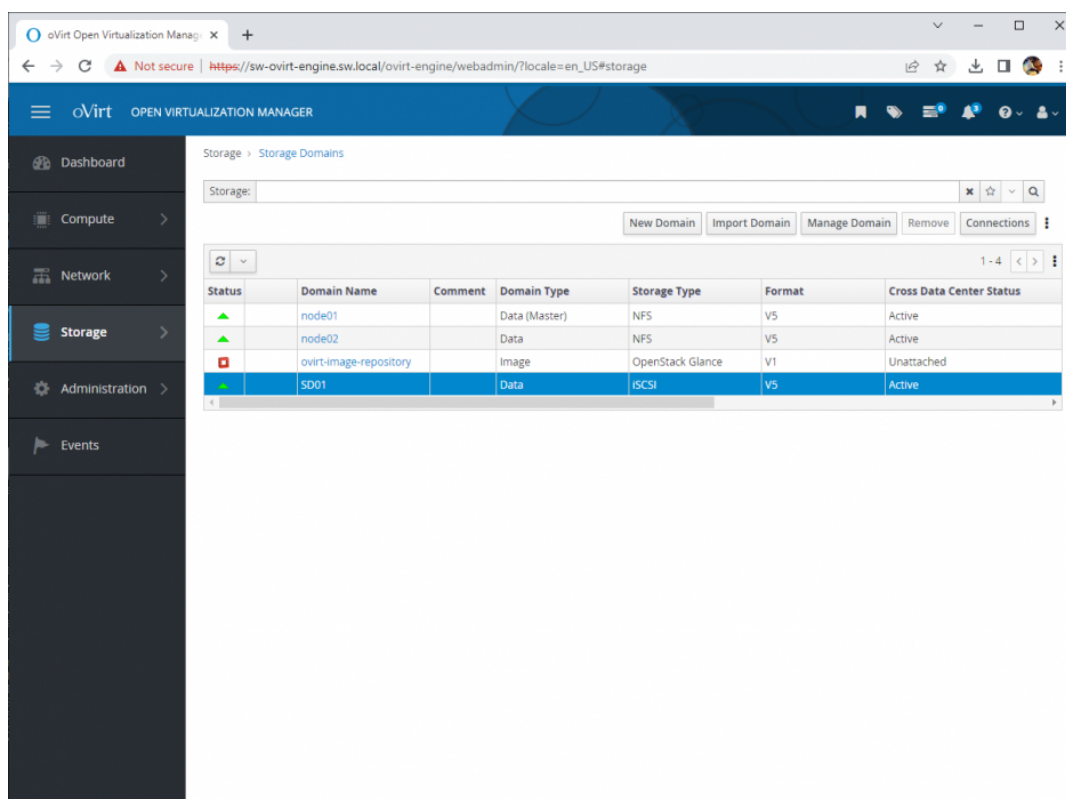
2. Choose Storage Type – iSCSI, Host and Name of Storage Domain. Discover targets via iSCSI links, which were previously configured. Click Login All.



3. Add LUN from each iSCSI target. Click OK.



4. Storage Domain will be added to the list of Domain and can be used as a storage for VMs.



5. Login to each host and verify that multipathing policy has been applied using the following command.






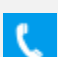
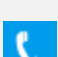
```
multipath -ll
```

```
[root@sw-demo-node-01 ~]# multipath -ll
22ebelf66db375fb0 dm-13 STARWIND,STARWIND
size=500G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
`-+- policy='round-robin 0' prio=50 status=active
   |- 16:0:0:0 sdb      8:16 active ready running
   `-- 17:0:0:0 sdc      8:32 active ready running
```

Conclusion

Following this guide, a StarWind Virtual HCI Appliance (VHCA) powered by Red Hat oVirt was deployed and 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 Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA
www.starwind.com ©2024, StarWind Software Inc. All rights reserved.