

StarWind Virtual SAN: Configuration Guide for Red Hat oVirt [KVM], 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.

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 SAN (VSAN)

Purpose

The guide details the configuration of StarWind Virtual SAN (VSAN) as a Controller Virtual Machine (CVM) on Red Hat oVirt environments. It includes steps for pre-configuring KVM hosts, enabling multipath support, creating NFS shares, deploying and configuring the StarWind Virtual SAN CVM, and setting up storage on oVirt hosts.

Audience

The intended audience includes system administrators and IT professionals who are responsible for managing and configuring virtualized storage solutions in KVM environments.

Expected Result

Upon following the guide, users should have a fully configured StarWind VSAN CVM in their Red Hat oVirt environment, resulting in a robust and efficient storage solution for their VM clusters.

Introduction To Starwind Virtual San Cvm

StarWind Virtual SAN Controller Virtual Machine (CVM) comes as a prepackaged Linux Virtual Machine (VM) to be deployed on any industry-standard hypervisor. It creates a VM-centric and high-performing storage pool for a VM cluster.

This guide describes the deployment and configuration process of the StarWind Virtual SAN CVM.

Starwind Vsan System Requirements

Prior to installing StarWind Virtual SAN, please make sure that the system meets the requirements, which are available via the following link:

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

Recommended RAID settings for HDD and SSD disks:

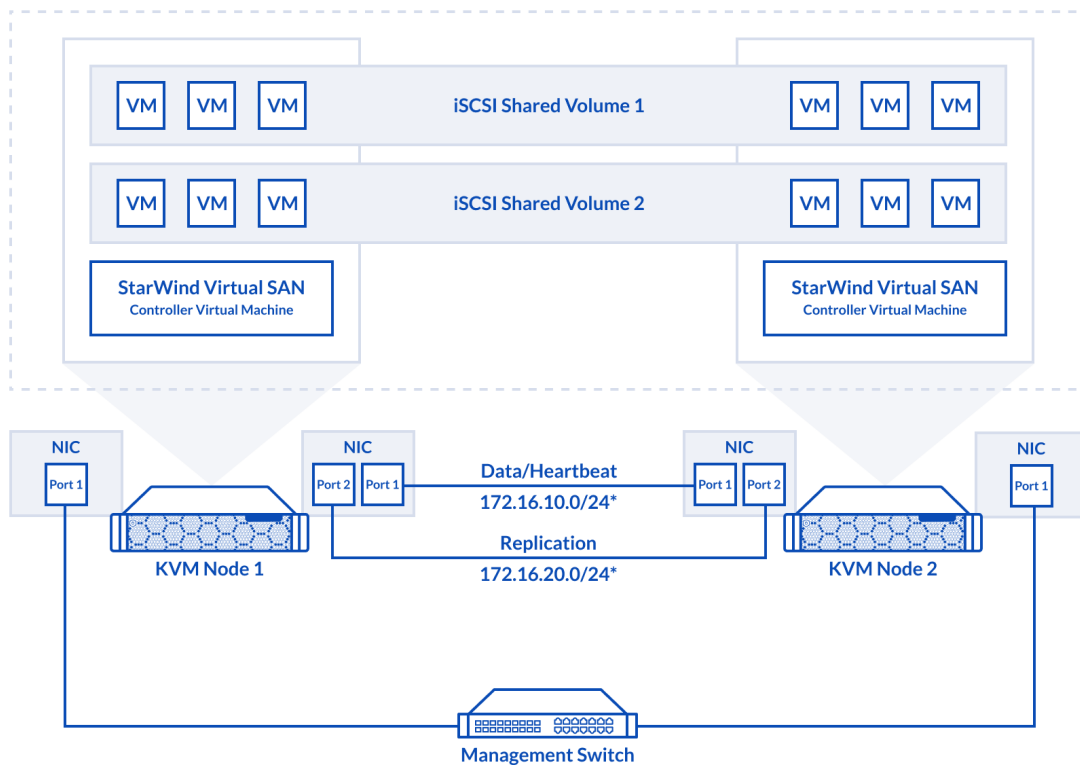
<https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/>

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

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

Pre-Configuring The Kvm Hosts

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



1. Make sure that a oVirt engine is installed on a separate host.
2. Deploy oVirt on each server and add them to oVirt engine.
3. Define at least 2x network interfaces on each node that will be used for the Synchronization and iSCSI/StarWind heartbeat traffic. Do not use iSCSI/Heartbeat and Synchronization channels

over the same physical link. Synchronization and iSCSI/Heartbeat links can be connected either via redundant switches or directly between the nodes (see diagram above).

4. Separate Logical Networks should be created for iSCSI and Synchronization traffic based on the selected before iSCSI and Synchronization interfaces. Using oVirt engine Networking page create two Logical Networks: one for the iSCSI/StarWind Heartbeat channel (iSCSI) and another one for the Synchronization channel (Sync).

5. Add physical NIC to Logical network on each host and configure static IP addresses. In this document, the 172.16.10.x subnet is used for iSCSI/StarWind heartbeat traffic, while 172.16.20.x subnet is used for the Synchronization traffic.

NOTE: In case NIC supports SR-IOV, enable it for the best performance. Contact support for additional details.

Enabling Multipath Support

8. Connect to server via ssh.

9. Create file /etc/multipath/conf.d/starwind.conf with the following content:

```
devices{
    device{
        vendor "STARWIND"
        product "STARWIND*"
        path_grouping_policy multibus
        path_checker "tur"
        failback immediate
        path_selector "round-robin 0"
        rr_min_io 3
        rr_weight uniform
        hardware_handler "1 alua"
    }
}
```

10. Restart multipathd service.

```
systemctl restart multipathd
```

11. Repeat the same procedure on the other server.

Creating NFS share

1. Make sure that each host has free storage to create NFS share.
2. Enable nfs server and rpcbind services.

```
systemctl enable --now nfs-server rpcbind
```

3. Create directory for NFS share.

```
mkdir -p /mnt/nfs
```

4. Change rights and owner of the share to KVM

```
chmod 0775 /mnt/nfs/  
chown -R nobody:users /mnt/nfs/
```

5. Add NFS share to /etc/exports file.

```
vi /etc/exports  
/mnt/nfs/ *(rw,anonuid=36,anongid=36)
```

6. Restart NFS server service.

```
systemctl restart nfs-server
```

7. Check that share has been exported.

```
exportfs -rvv
```

8. Add firewall rules for NFS.

```
firewall-cmd --add-service={nfs,nfs3,rpc-bind} --permanent  
firewall-cmd --reload
```

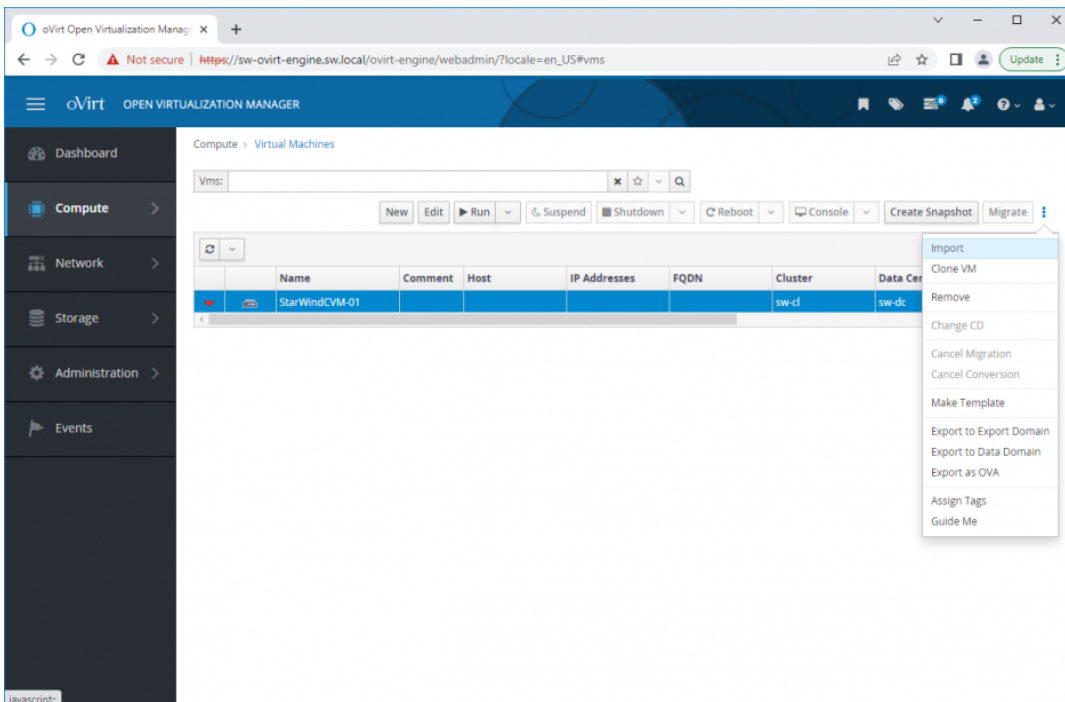
Deploying Starwind Virtual San Cvm

1. Download StarWind VSAN CVM KVM: [VSAN by StarWind: Overview](#)


2. Extract the VM StarWindCVM.ova file from the downloaded archive.
3. Upload StarWindCVM.ova file to the oVirt Host via any SFTP client.
4. Change owner of the StarWindCVM.ova.

```
chown -R nobody:users /mnt/nfs/
```

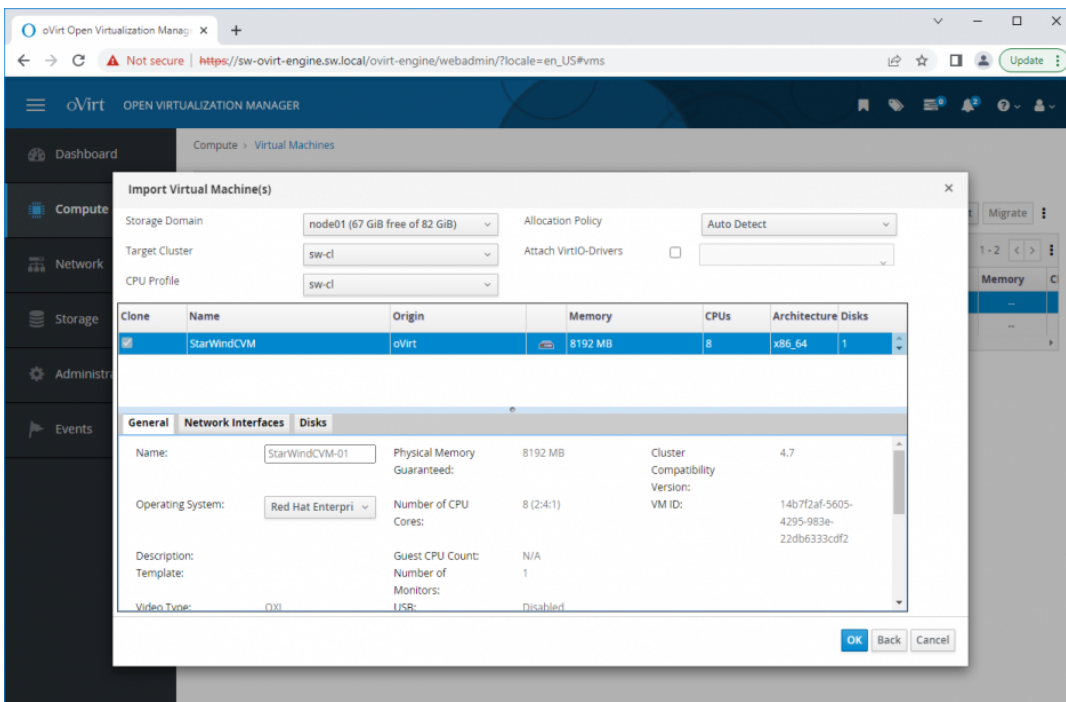
5. Login to oVirt and open Compute -> Virtual Machines page. Choose Import.



6. Specify path to .ova file and choose VM to import. Click Next.




7. Verify VM settings and configure networks. Click OK.



8. Repeat all the steps from this section on other oVirt hosts.

Configuring Starwind Virtual San Vm Settings

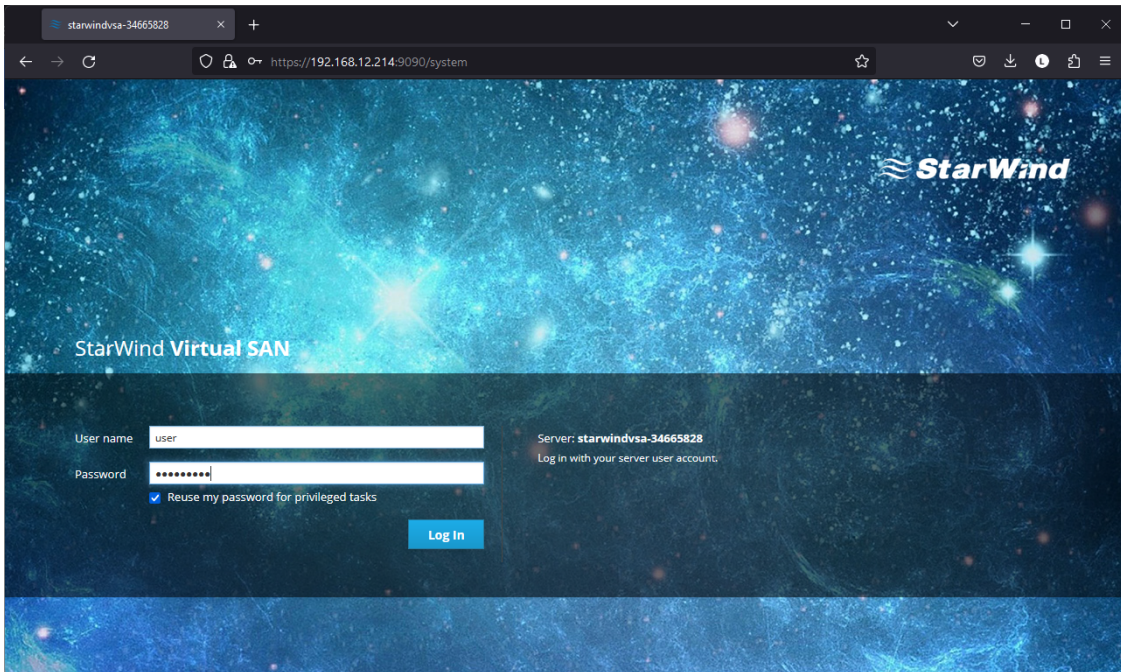
1. Open the VM console and check the IP address received via DHCP (or which was assigned manually).



Another alternative is to log into the VM via console and assign static IP using nmcli if there is no DHCP.

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

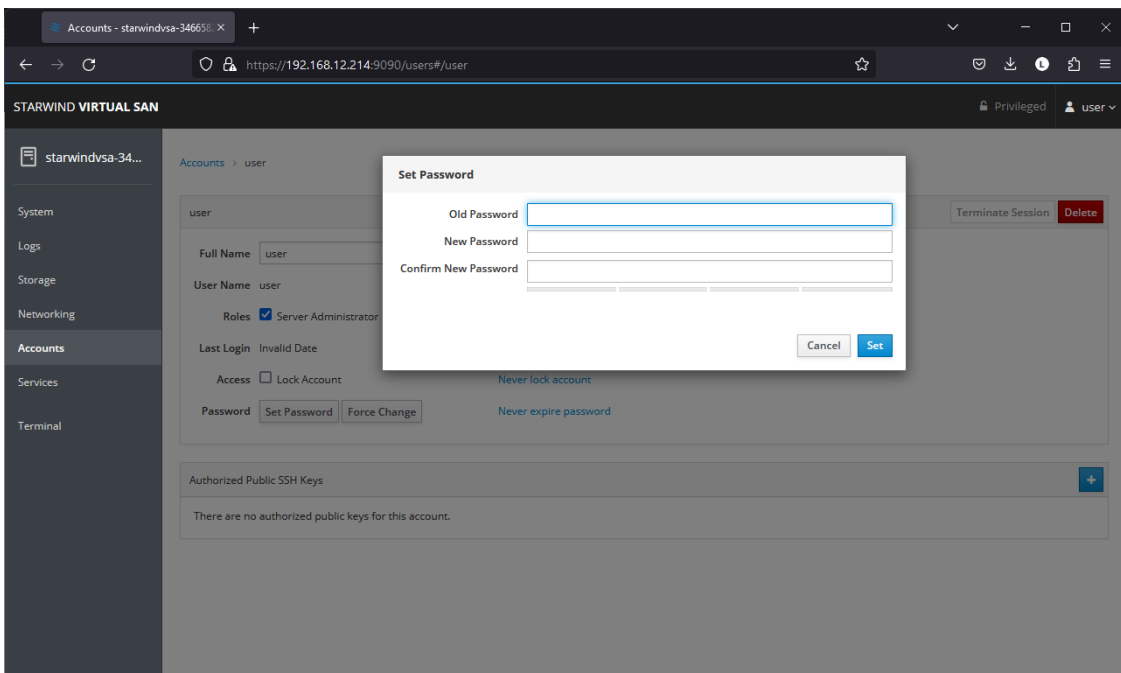
- Username: user
- Password: rds123RDS
- NOTE: Make sure to check the “Reuse my password for privileged tasks” box.



3. Af

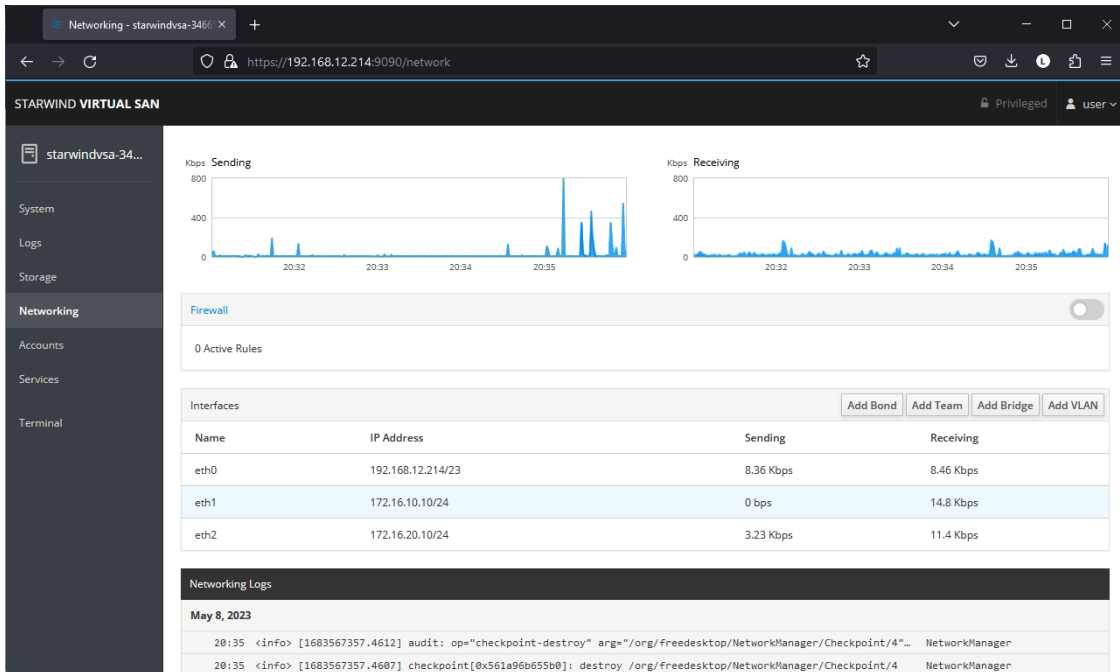
ter a successful login, click Accounts on the left sidebar.

4. Select a user and click Set Password.



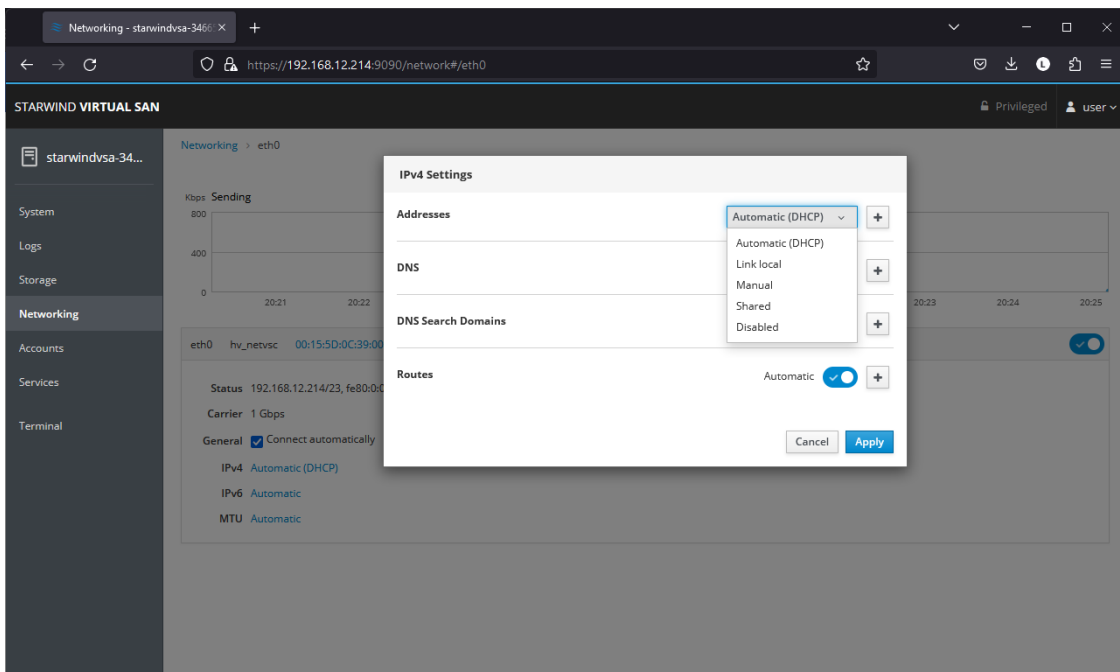
5.

On the left sidebar, click Networking.

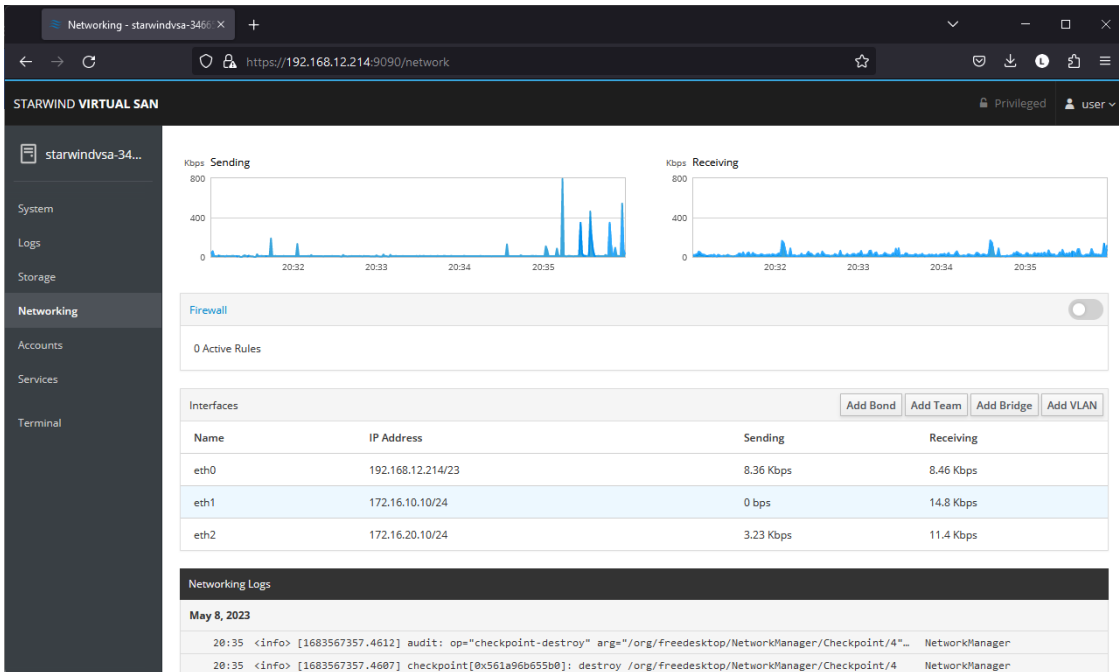


Here, the Management IP address of the StarWind Virtual SAN Virtual Machine can be configured, as well as IP addresses for iSCSI and Synchronization networks. In case the Network interface is inactive, click on the interface, turn it on, and set it to Connect automatically.

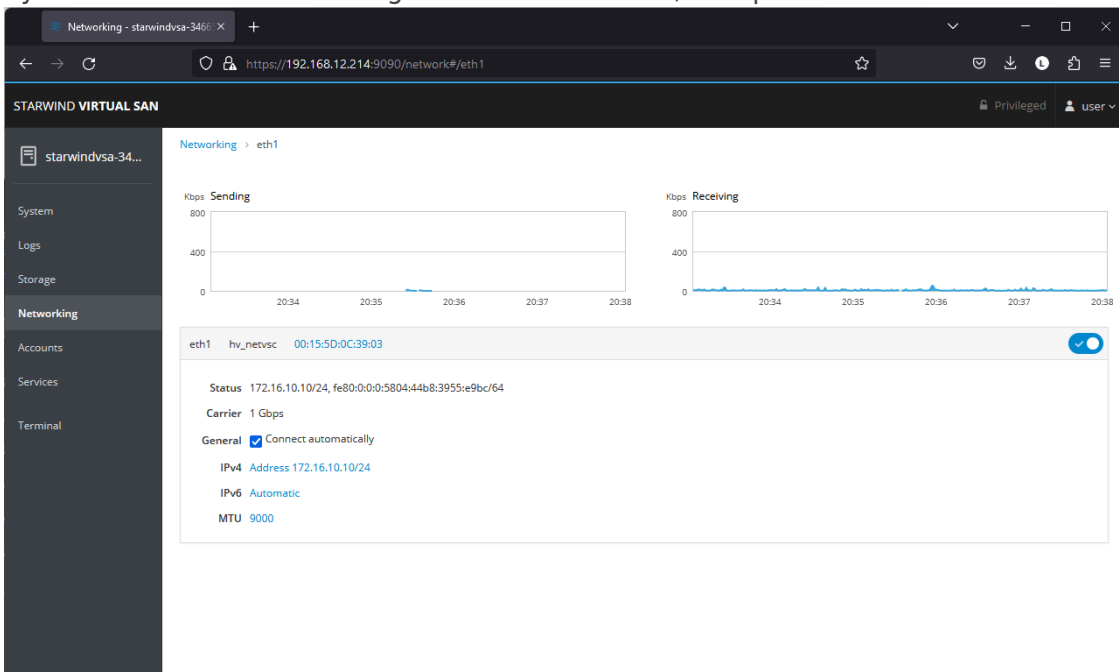
6. Click on Automatic (DHCP) to set the IP address (DNS and gateway - for Management).



7. 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.



8. Alternatively, log into the VM via the oVirt console and assign a static IP address by editing the configuration file of the interface located by the following path:
/etc/sysconfig/network-scripts

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

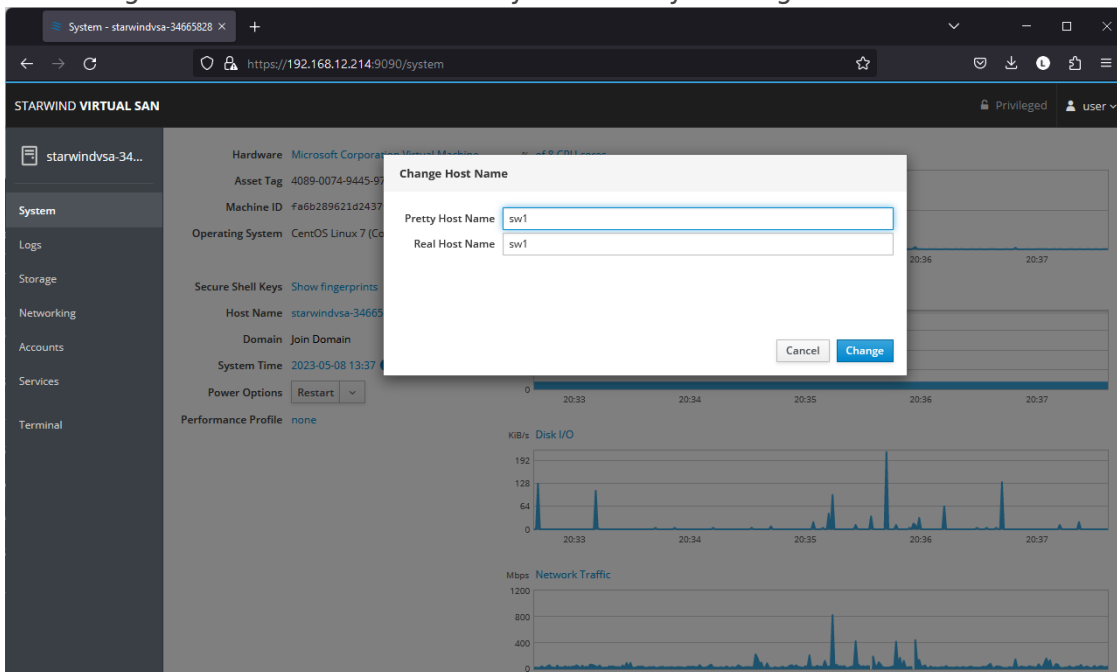
10.

1. Edit the file:

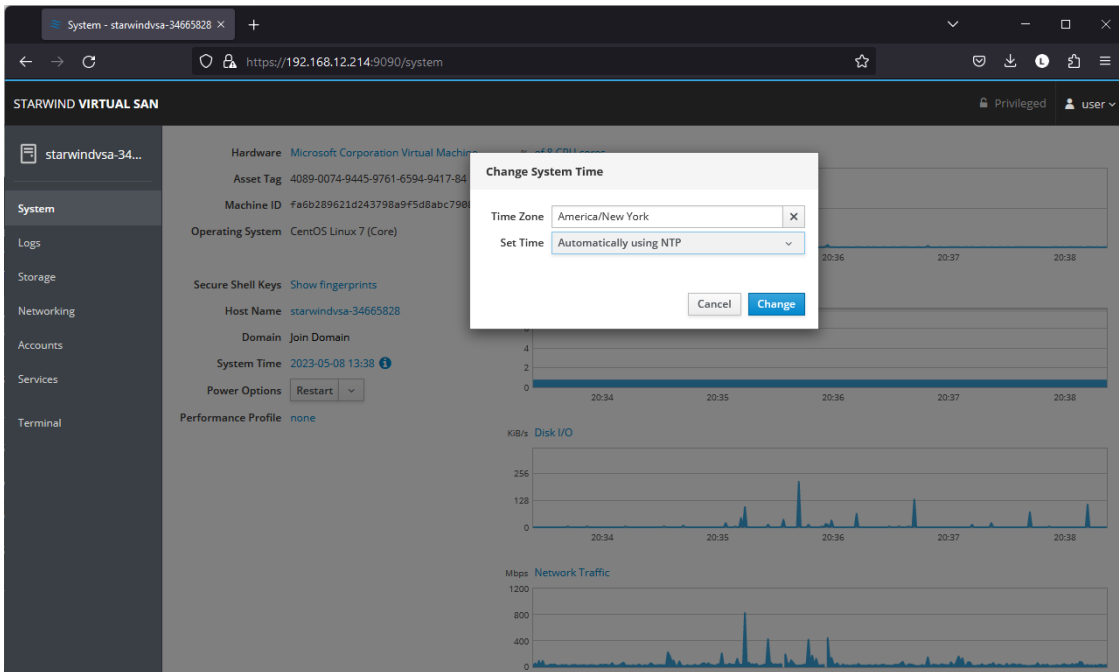
- Change the line BOOTPROTO=dhcp to: BOOTPROTO=static
- Add the IP settings needed to the file:
- IPADDR=192.168.12.10
- NETMASK=255.255.255.0
- GATEWAY=192.168.12.1
- DNS1=192.168.1.1

11. Restart the interface using the following cmdlet: `sudo ifdown eth0, sudo ifup eth0`, or restart the VM.

12. Change the Host Name from the System tab by clicking on it:



13. Change System time and NTP settings if required:



14. Repeat the steps above on each StarWind VSAN VM.

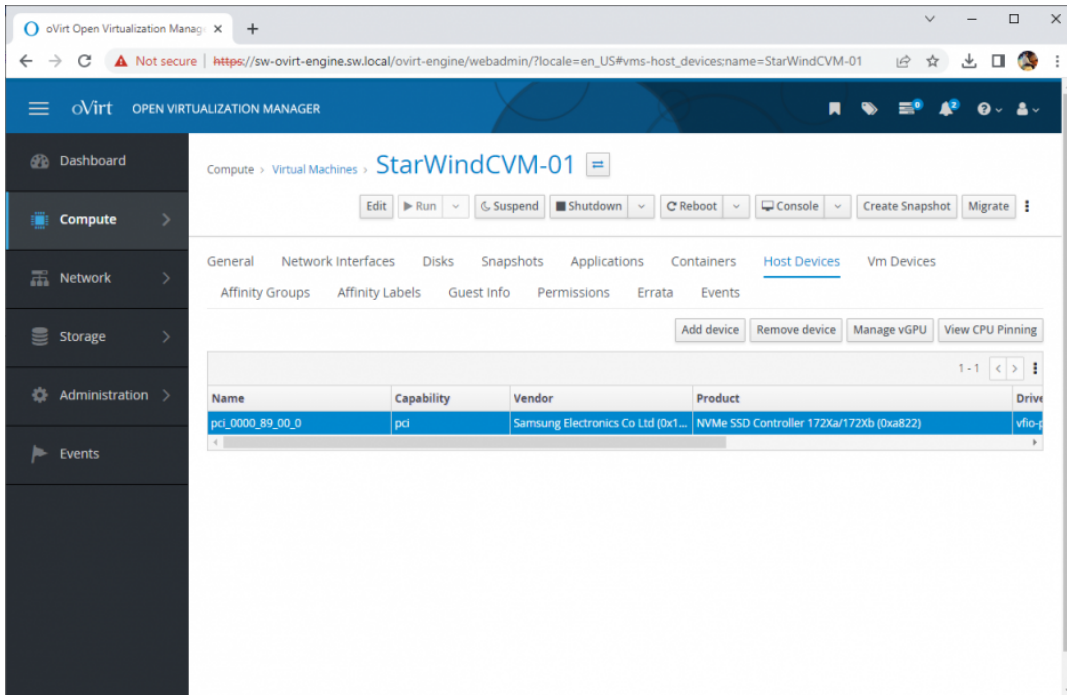
Configuring Storage

StarWind Virtual SAN for vSphere can work on top of Hardware RAID or Linux Software RAID (MDADM) inside of the Virtual Machine.

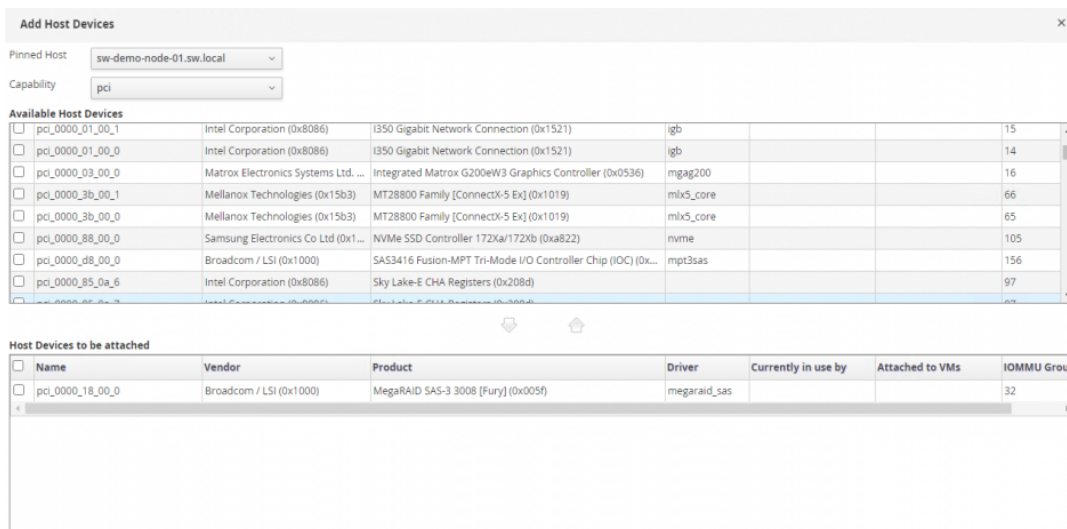
Please select the required option:

Configuring Storage With Hardware Raid

1. Open VM Settings on oVirt Virtual Machines page and choose Host Devices tab. It is recommended to pass the entire RAID array to the VM. Click Add device.



2. Choose RAID controller or HBA and pass it to the VM. Click OK.

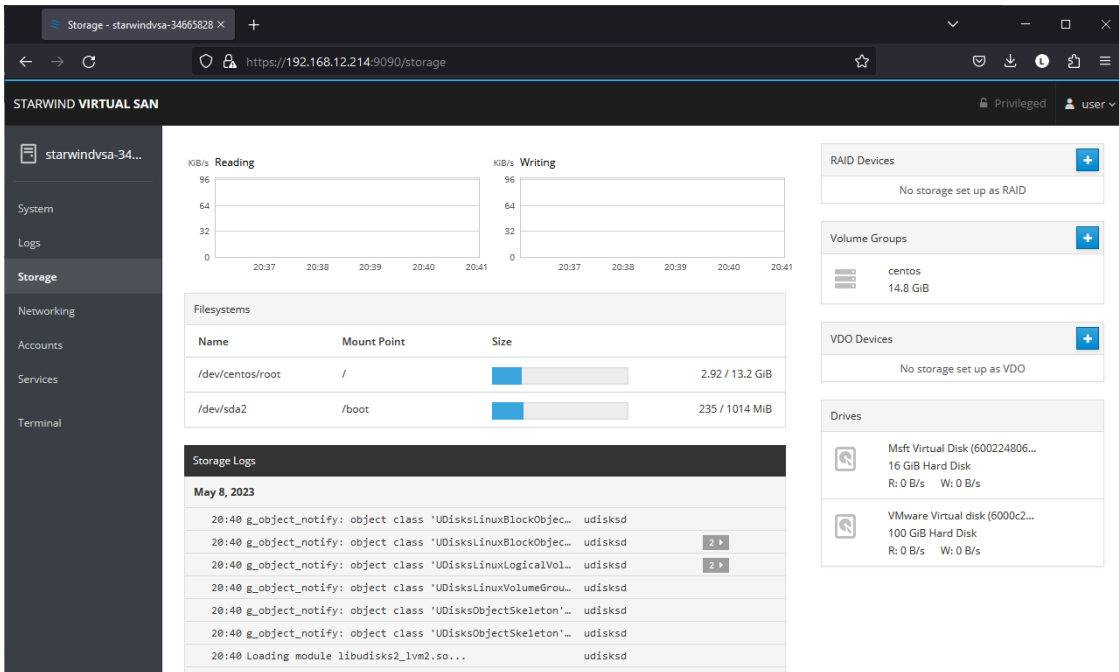


OK Cancel

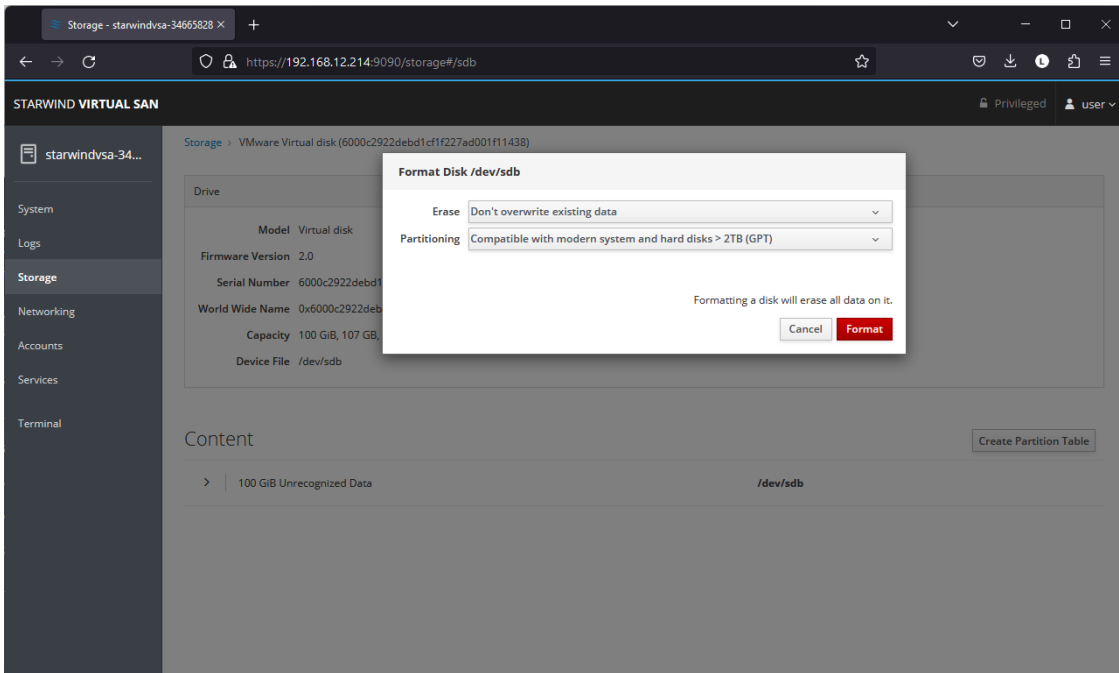
3. Start

StarWind VSAN CVM.


4. Login to StarWind VSAN VM web console and access the Storage section. Locate the recently added disk in the Drives section and choose it.



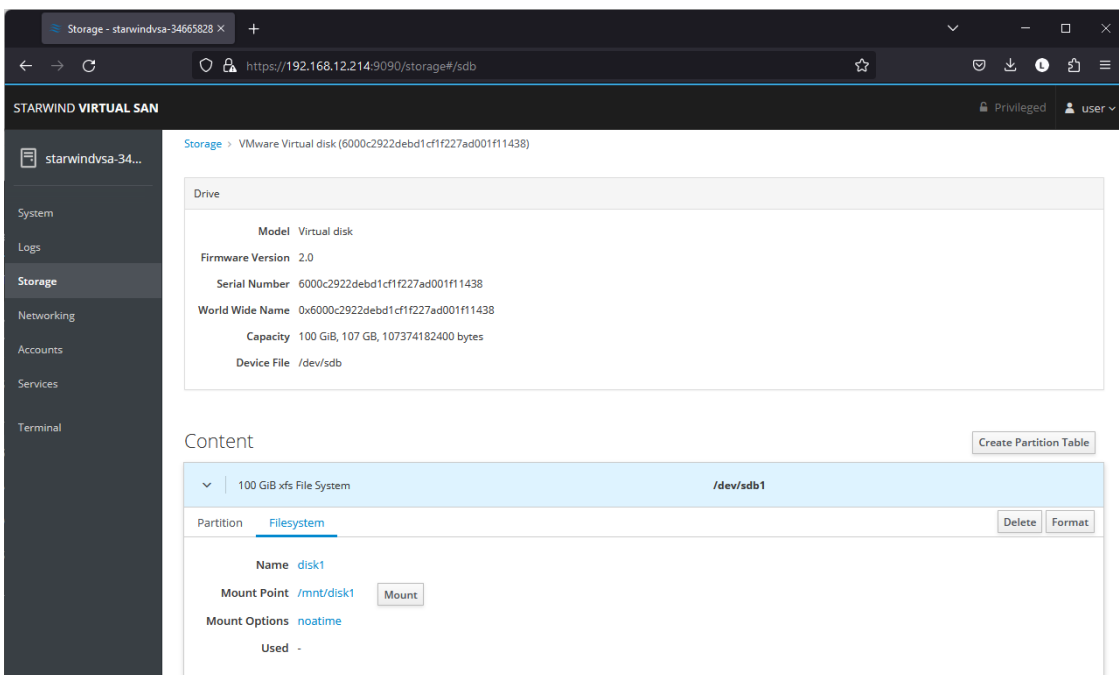
5. The added disk does not have any partitions and filesystem. Press the Create Partition Table button to create the partition.



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



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



Configuring Starwind Management Console

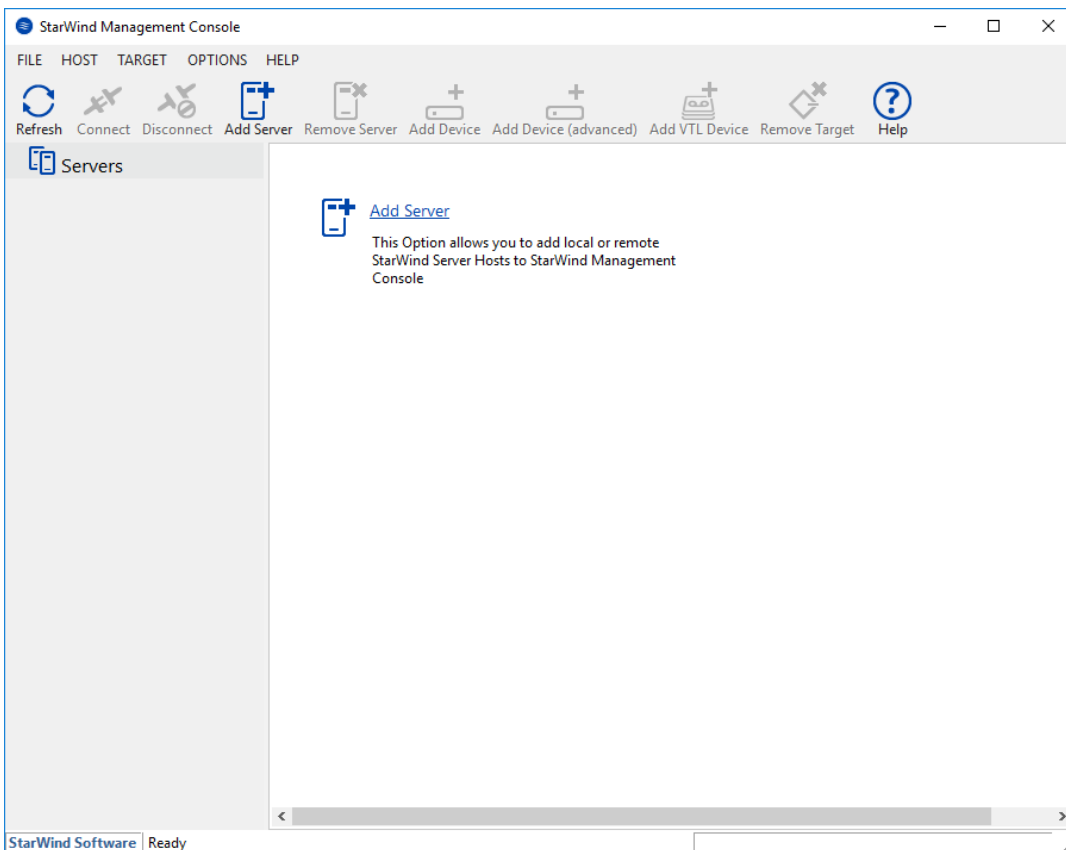
1. Install StarWind Management Console on each server or on a separate workstation with Windows OS (Windows 7 or higher, Windows Server 2008 R2 and higher) using the

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.

```

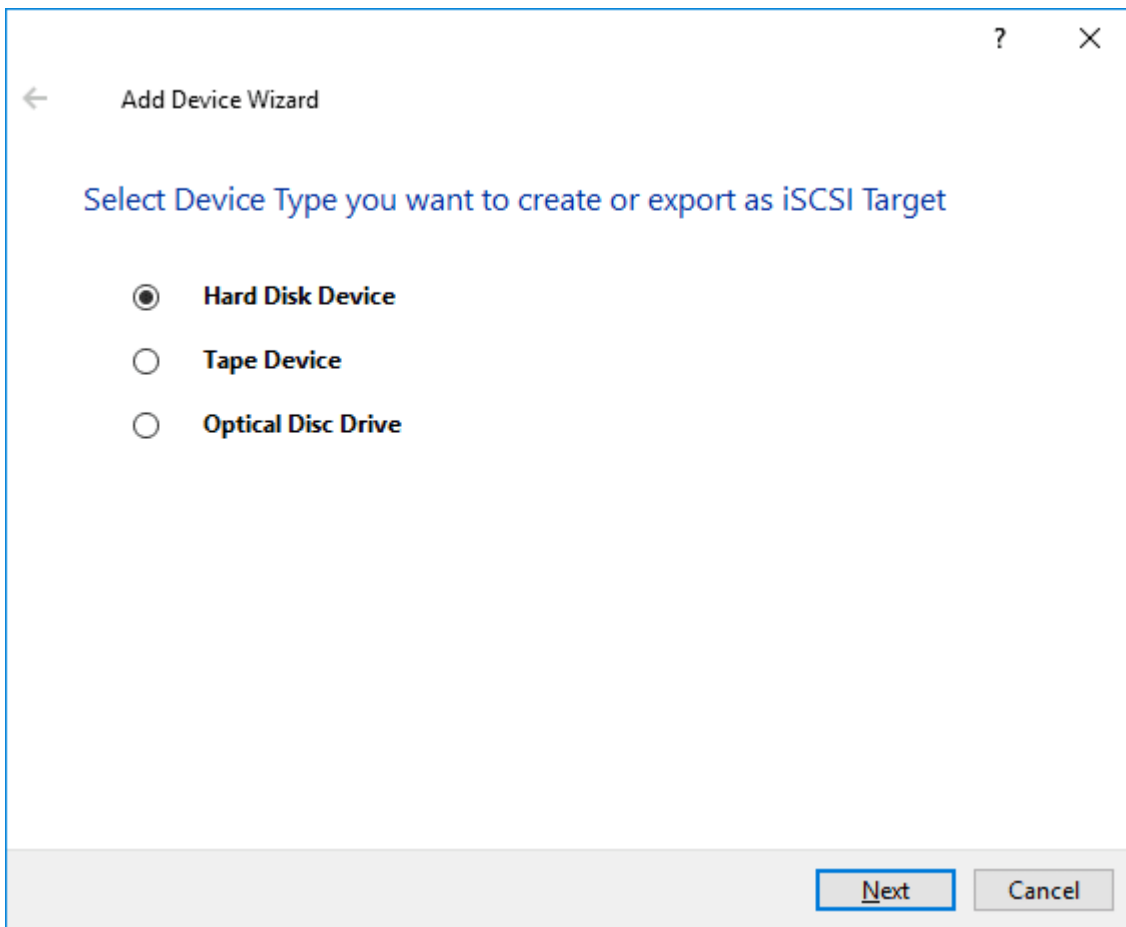
Administrator: Windows PowerShell ISE
File Edit View Tools Debug Add-ons Help
InstallLicense.ps1 X
1 #
2 # The following example shows how to apply license on a server
3 #
4 Import-Module StarWindX
5
6 Enable-SWXLog
7
8 $server = New-SWServer -host 127.0.0.1 -port 3261 -user root -password starwind
9
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

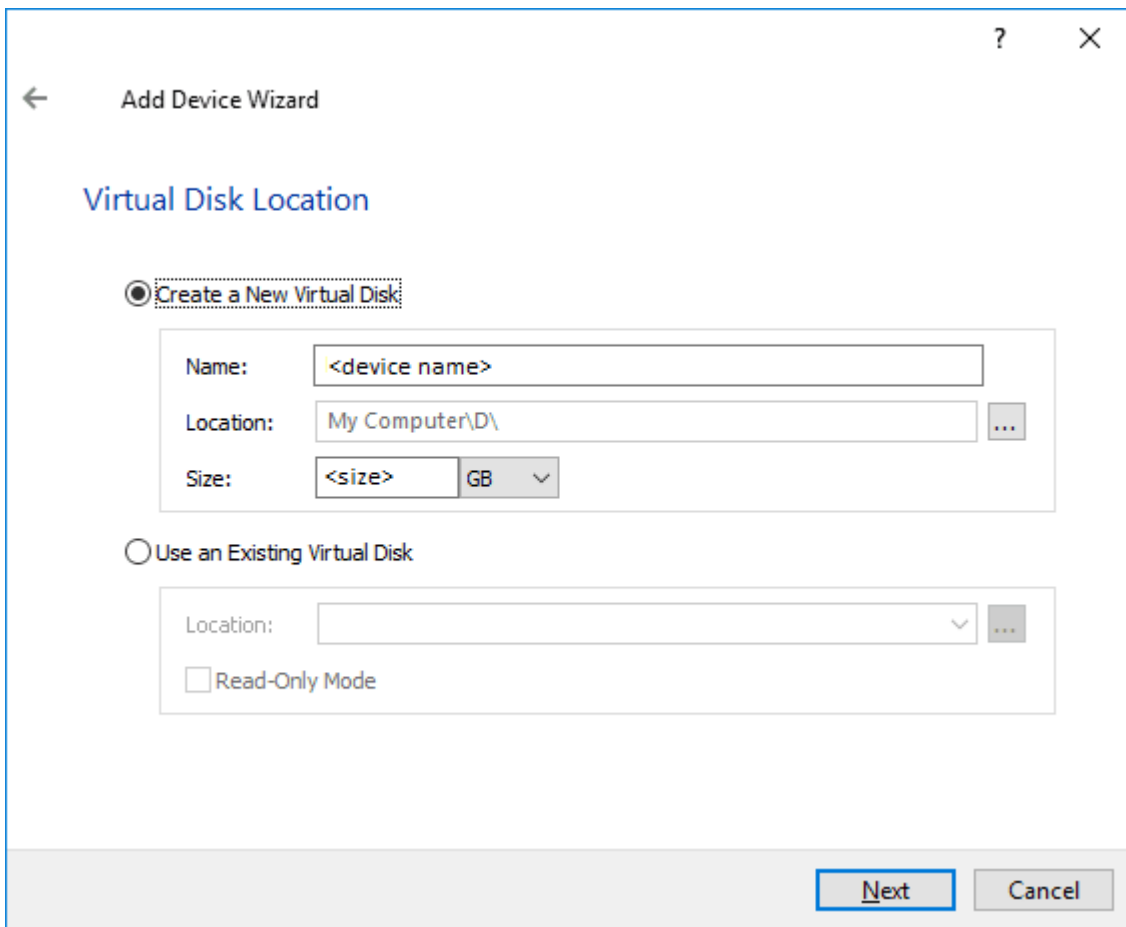
1. In the StarWind Management Console click to Add Device (advanced) button and open Add Device (advanced) Wizard.
2. Select Hard Disk Device as the type of device to be created.



3. Select Virtual Disk.




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




5. Select the Thick provisioned disk type and block size.

NOTE: Use 4096 sector size for targets, connected on Windows-based systems and 512 bytes sector size for targets, connected on Linux-based systems (ESXi/Xen/KVM).


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



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



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



9. Click Close to finish the device creation.


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

Select The Required Replication Mode


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

Synchronous “Two-Way” Replication

1. Right-click the recently created device and select Replication Manager from the shortcut menu.
2. Select the Add Replica button in the top menu.



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



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

Selecting The Failover Strategy

StarWind provides 2 options for configuring a failover strategy:

Heartbeat

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

If at least one heartbeat link is online, StarWind services can communicate with each other via this link. The device with the lowest priority will be marked as not synchronized and get subsequently blocked for the further read and write operations until the synchronization channel resumption. At the same time, the partner device on the

synchronized node flushes data from the cache to the disk to preserve data integrity in case the node goes down unexpectedly. It is recommended to assign more independent heartbeat channels during the replica creation to improve system stability and avoid the “split-brain” issue.

With the heartbeat failover strategy, the storage cluster will continue working with only one StarWind node available.

Node Majority

The Node Majority failover strategy ensures the synchronization connection without any additional heartbeat links. The failure-handling process occurs when the node has detected the absence of the connection with the partner.

The main requirement for keeping the node operational is an active connection with more than half of the HA device’s nodes. Calculation of the available partners is based on their “votes”.

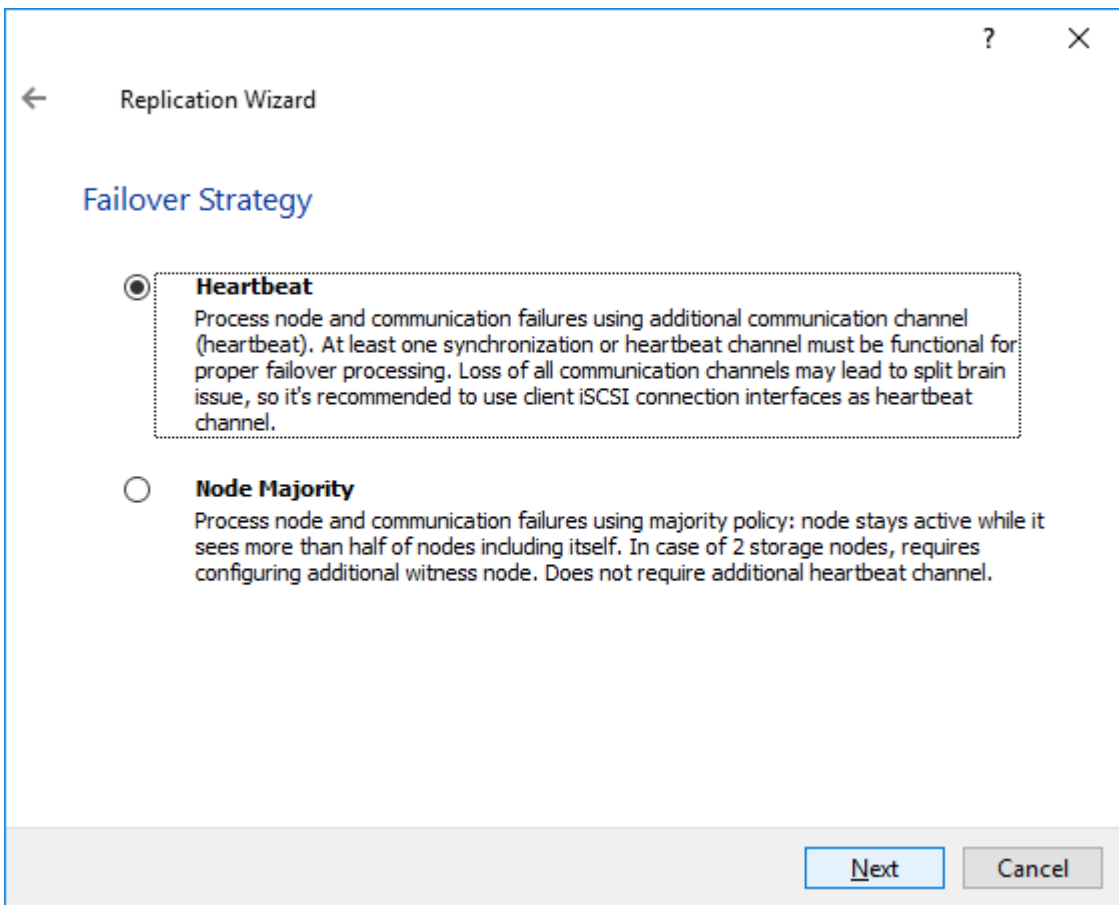
In case of a two-node HA storage, all nodes will be disconnected if there is a problem on the node itself, or in communication between them. Therefore, the Node Majority failover strategy requires the addition of the third Witness node or file share (SMB) which participates in the nodes count for the majority, but neither contains data on it nor is involved in processing clients’ requests. In case an HA device is replicated between 3 nodes, no Witness node is required.

With Node Majority failover strategy, failure of only one node can be tolerated. If two nodes fail, the third node will also become unavailable to clients’ requests.

Please select the required option:

Heartbeat

1. Select Failover Strategy.



2. Select Create new Partner Device and click Next.

3. Select a partner device Location and click Next.

4. Select Synchronization Journal Strategy and click Next.

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

RAM-based (default) synchronization journal is placed in RAM. Synchronization with RAM journal provides good I/O performance in any scenario. Full synchronization could occur in the cases described in this KB:

<https://knowledgebase.starwindsoftware.com/explanation/reasons-why-full-synchronizati-on-may-start/>

Disk-based journal placed on a separate disk from StarWind devices. It allows to avoid full synchronization for the devices where it's configured even when StarWind service is being stopped on all nodes.

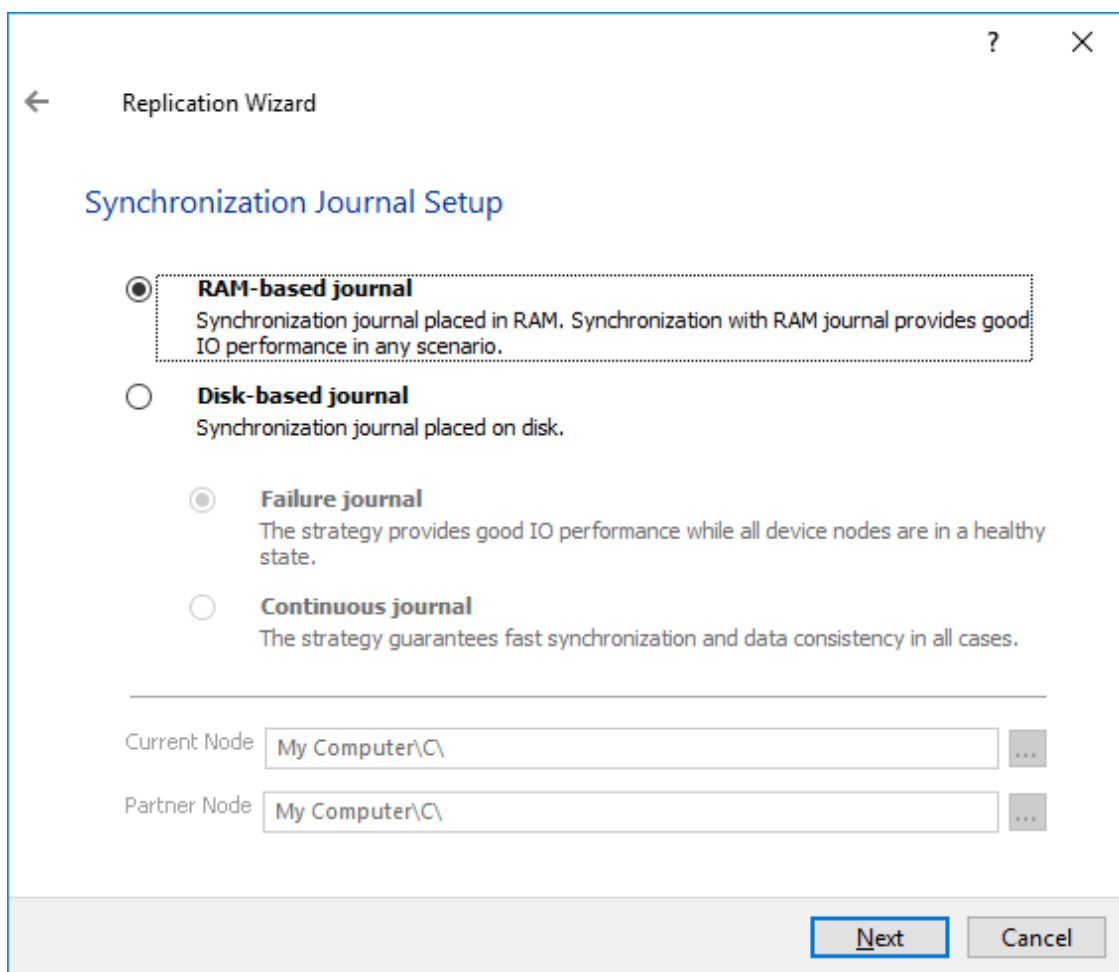
Disk-based synchronization journal should be placed on a separate, preferably faster disk from StarWind devices. SSDs and NVMe disks are recommended as the device

performance is defined by the disk speed, where the journal is located. For example, it can be placed on the OS boot volume.


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

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


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



5. Click Change Network Settings.



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



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

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

9. Follow the 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.



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



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

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

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

8. Click Create Replica and then close the wizard.


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

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


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

Provisioning Starwind Ha Storage To Hosts


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



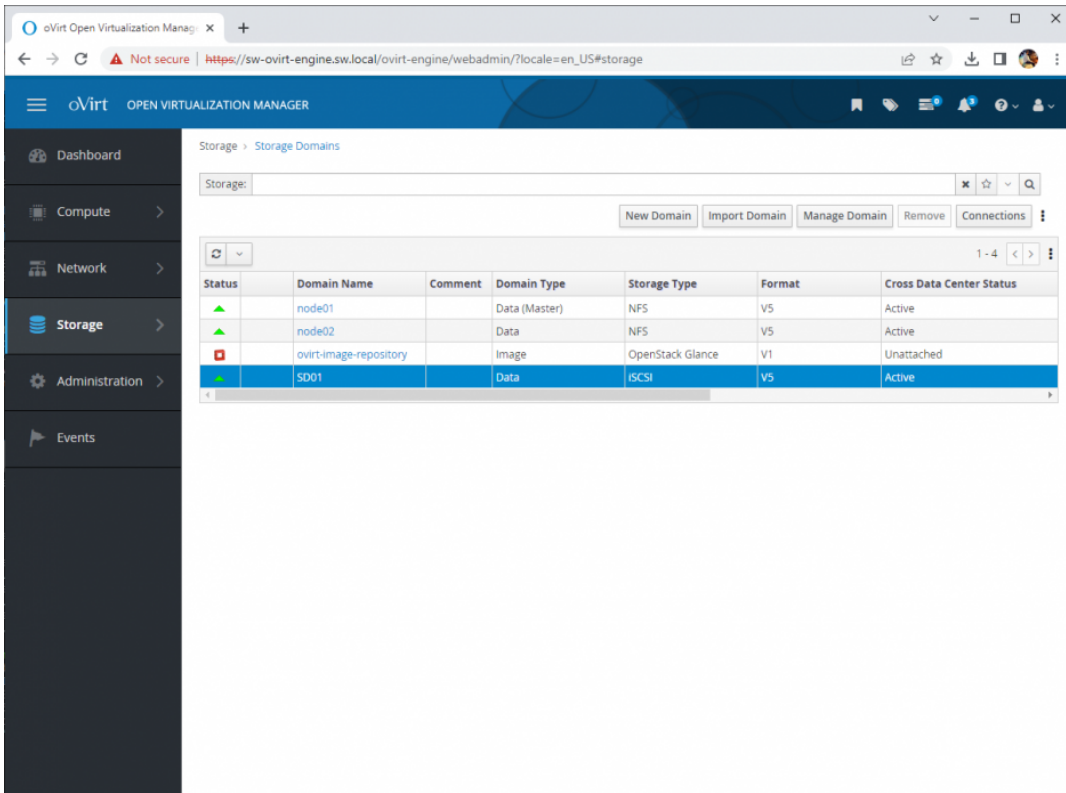
2. Choose Storage Type - iSCSI, Host and Name of Storage Domain. Discover targets via iSCSI links, which were previously configured. Click Login All.



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

This comprehensive guide provides a step-by-step approach to configuring StarWind Virtual SAN in a KVM environment, ensuring that users can efficiently set up a high-performance storage pool. The guide covers all necessary aspects, from pre-configuring hosts to finalizing storage settings, making it a valuable resource for professionals in virtualized storage management.

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.