

StarWind Virtual SAN: Configuration Guide for VMware vSphere [ESXi], VSAN Deployed as a VM (CVM) 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

This guide outlines the process for deploying and configuring StarWind Virtual SAN (VSAN) as a virtual machine (CVM) within VMware vSphere environments. It focuses on creating a VM-centric and high-performing storage pool for VMware clusters, detailing requirements, pre-configuration steps, and specific settings for successful implementation.

Audience

This guide is intended for system administrators, virtualization specialists, and IT professionals responsible for managing and optimizing VMware vSphere environments, especially those looking to implement advanced storage solutions like StarWind Virtual SAN.

Expected Result

By following the guide, users will successfully deploy StarWind Virtual SAN as a VM on vSphere, achieving a high-performance and fault-tolerant storage pool that integrates seamlessly with VMware clusters.

Starwind Virtual San For Vsphere Vm Requirements

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

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

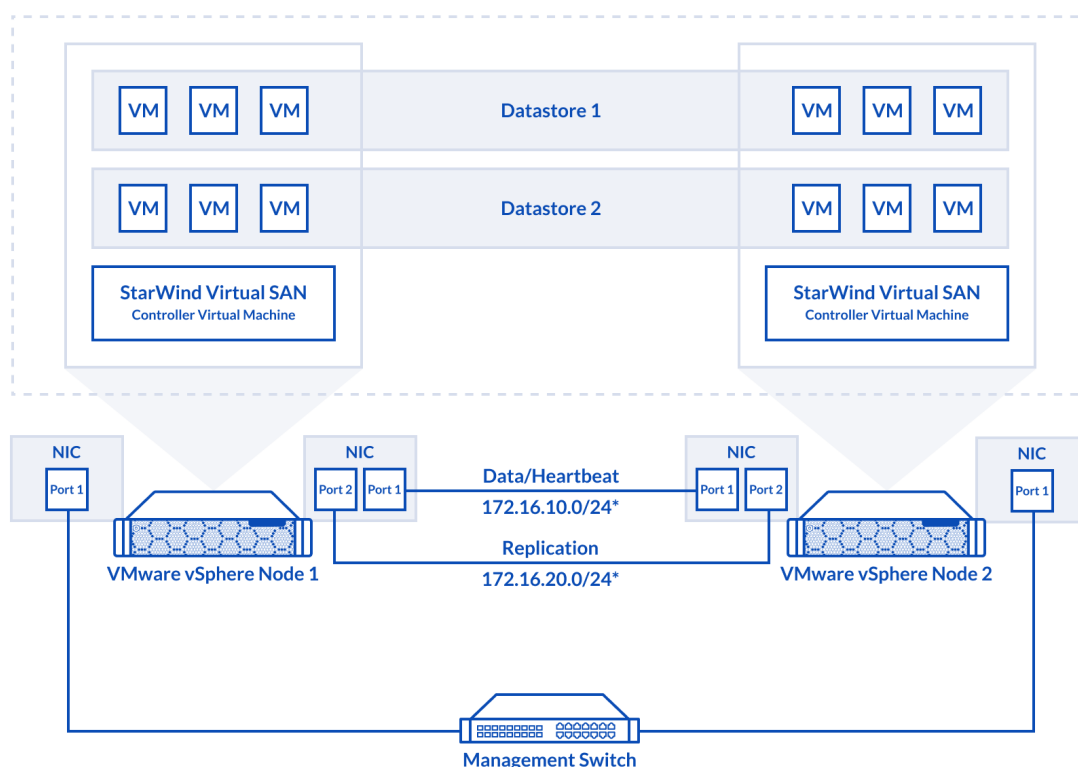
Recommended RAID settings for HDD and SSD disks:

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

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

Pre-Configuring The Servers

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



1. ESXi hypervisor should be installed on each host.
2. StarWind Virtual SAN for vSphere VM should be deployed on each ESXi host from an OVF template, downloaded on this page:
<https://www.starwindsoftware.com/release-notes-build-vsan-for-vmware-vmware>
3. 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.

NOTE: Do not use iSCSI/Heartbeat and Synchronization channels over the same physical link. Synchronization and iSCSI/Heartbeat links can be connected either via redundant switches or directly between the nodes.

vCenter Server can be deployed separately on another host or as VCSA on StarWind VSAN highly-available storage, created in this guide.

Preparing Environment For Starwind Vsan Deployment

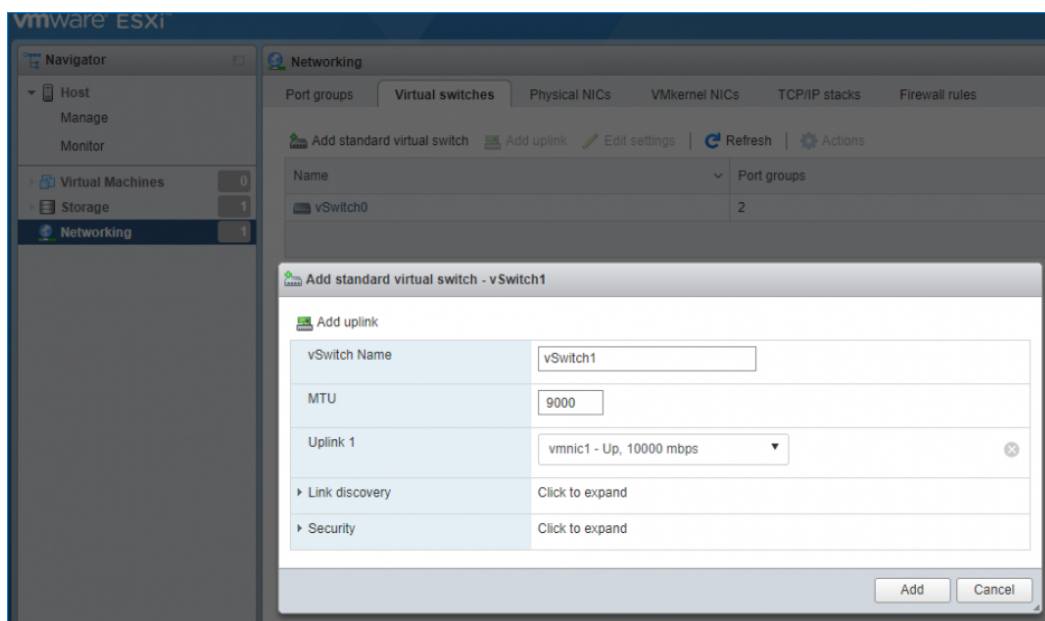
Configuring Networks

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.

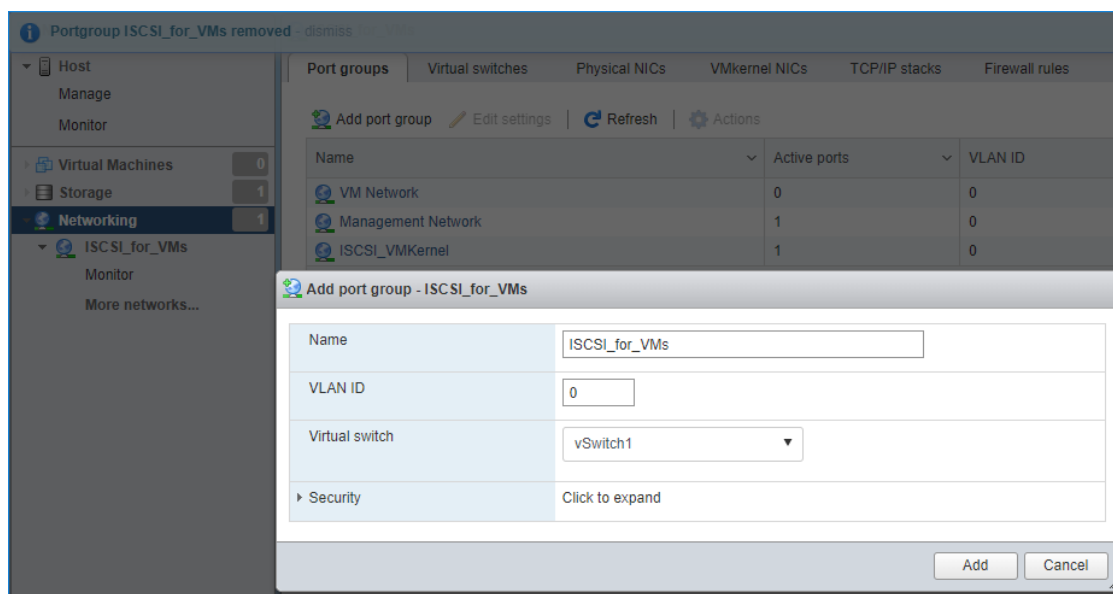
1. Using the VMware ESXi web console, create two standard vSwitches: one for the iSCSI/StarWind Heartbeat channel (vSwitch1) and the other one for the Synchronization channel (vSwitch2).



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



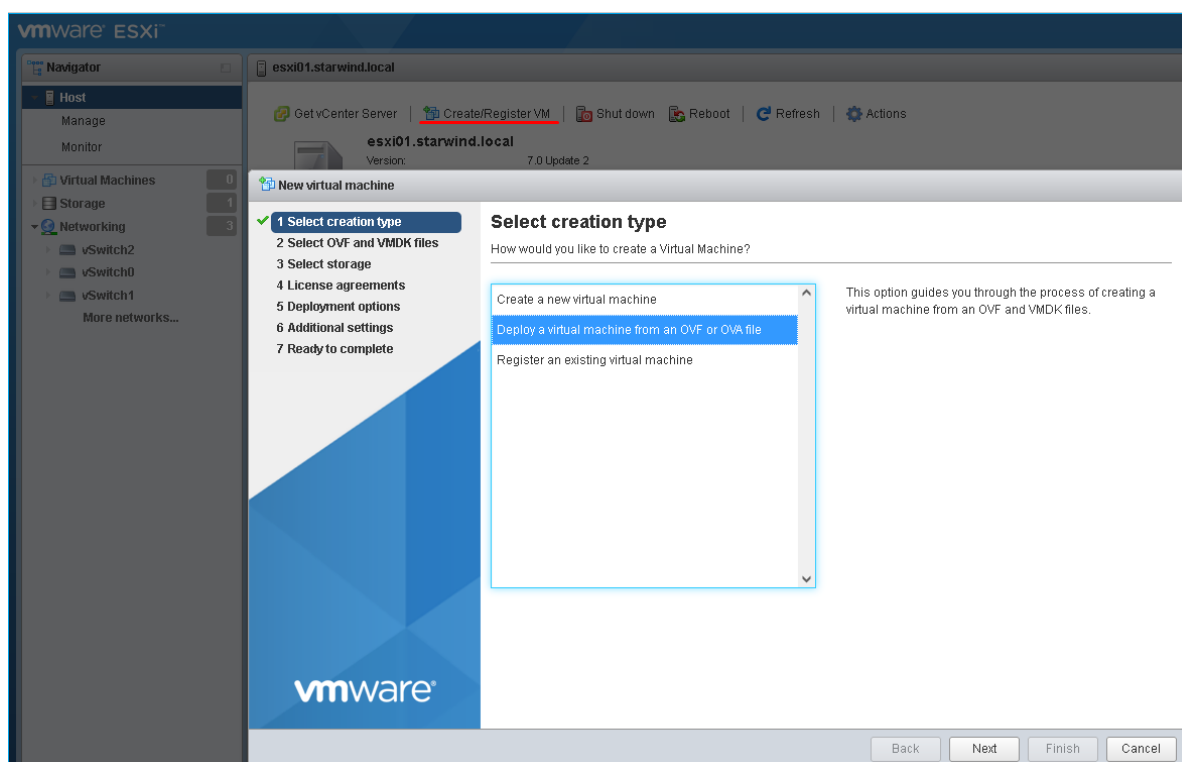
3. Add a Virtual Machine Port Groups on the vSwitch for iSCSI traffic (vSwitch1) and on the vSwitch for Synchronization traffic (vSwitch2).



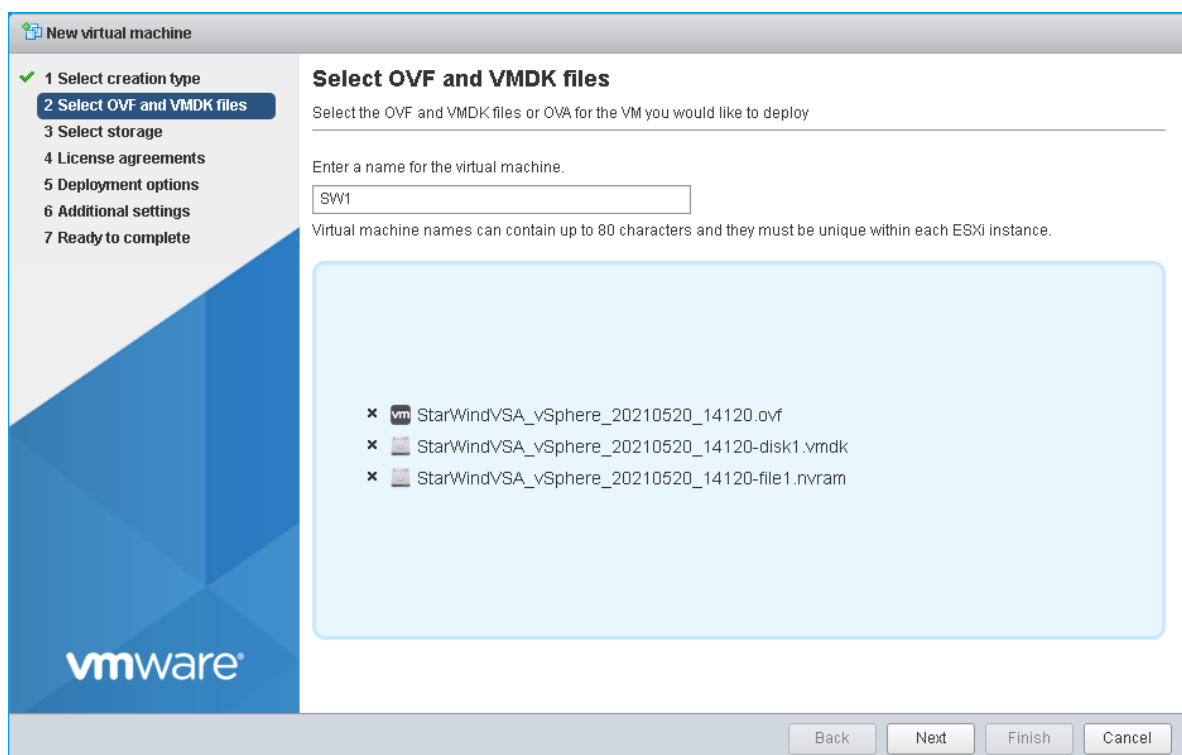
4. Repeat steps 1-3 for any other links intended for Synchronization and iSCSI/Heartbeat traffic on ESXi hosts.

Deploying Starwind Virtual San For Vsphere

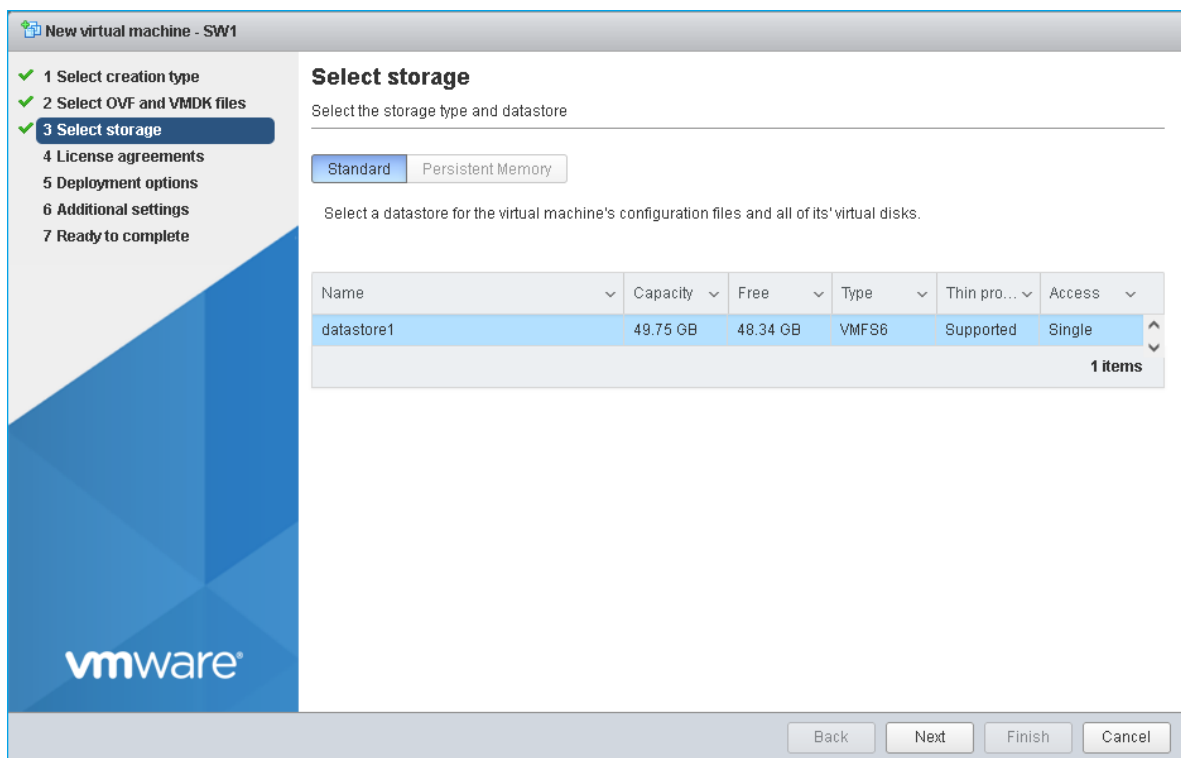
1. Download zip archive that contains StarWind Virtual SAN for vSphere:
<https://www.starwindsoftware.com/starwind-virtual-san#download>
2. Extract virtual machine files from the downloaded archive.
3. Deploy a virtual machine on each ESXi host using the "Create/Register VM" button. Select "Deploy a virtual machine from an OVF or OVA file" in the Select creation type section and press Next.



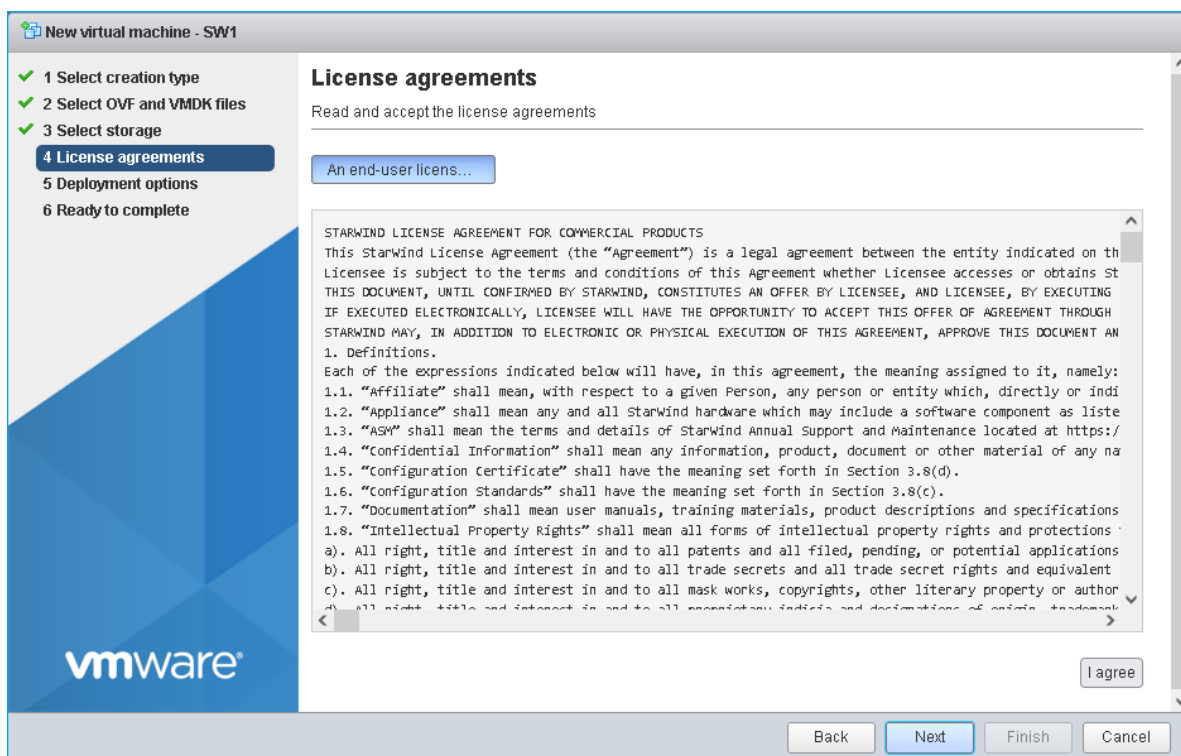
4. Specify the name for the virtual machine with StarWind VSAN, drag and drop the extracted files to the wizard, and press Next.



5. Specify the location for the StarWind Virtual SAN VM and press Next.



6. Read and accept the license agreements by pressing on “I agree” button. Press Next to continue.



7. Choose the proper networks for the VM by their purpose and other options.

New virtual machine - SW1

- ✓ 1 Select creation type
- ✓ 2 Select OVF and VMDK files
- ✓ 3 Select storage
- ✓ 4 License agreements
- 5 Deployment options**
- 6 Ready to complete

Deployment options

Select deployment options

Network mappings	Management: VM Network
	iSCSI: iSCSI_for_VMs
	Sync: SYNC_for_VMs
Disk provisioning	<input type="radio"/> Thin <input checked="" type="radio"/> Thick
Power on automatically	<input checked="" type="checkbox"/>

Back Next Finish Cancel

8. Review the settings and click Finish to start the deployment process.

New virtual machine - SW1

- ✓ 1 Select creation type
- ✓ 2 Select OVF and VMDK files
- ✓ 3 Select storage
- ✓ 4 License agreements
- ✓ 5 Deployment options
- 6 Ready to complete**

Ready to complete

Review your settings selection before finishing the wizard

Product	StarWindVSA_vSphere_20210520_14120
VM Name	SW1
Files	StarWindVSA_vSphere_20210520_14120-disk1.vmdk StarWindVSA_vSphere_20210520_14120-file1.nvram
Datastore	datastore1
Provisioning type	Thick
Network mappings	Management: VM Network, iSCSI: iSCSI_for_VMs, Sync: SYNC_for_VMs
Guest OS Name	Red Hat Enterprise Linux 7 (64-bit)

Do not refresh your browser while this VM is being deployed.

Back Next Finish Cancel

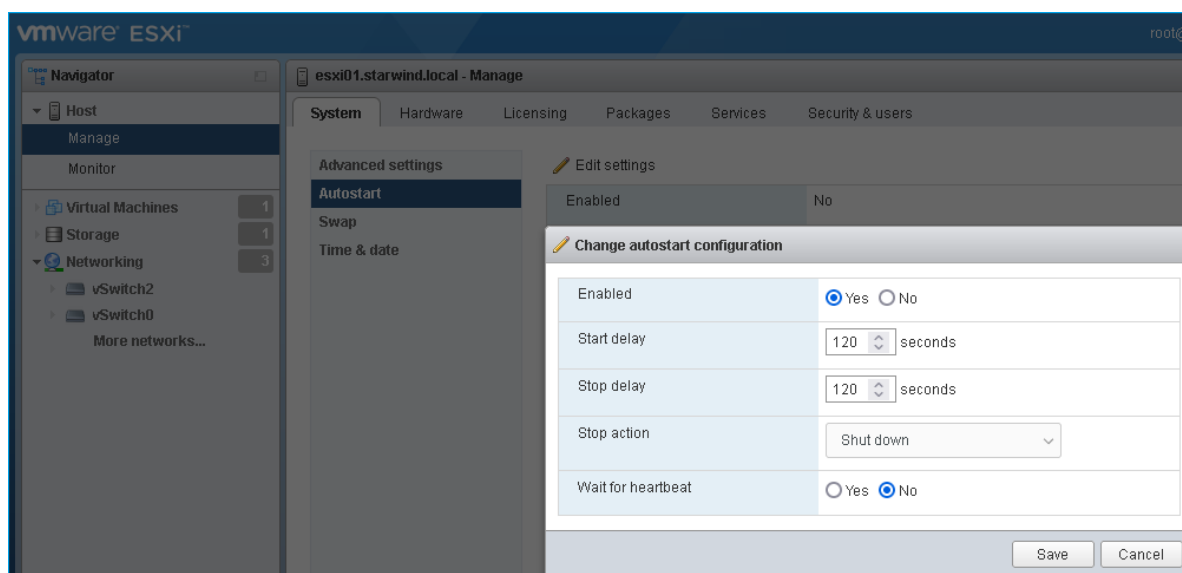
12. Repeat all the steps from this section on the other ESXi hosts.

NOTE: In some cases, it's recommended to reserve memory for StarWind VSAN VM.

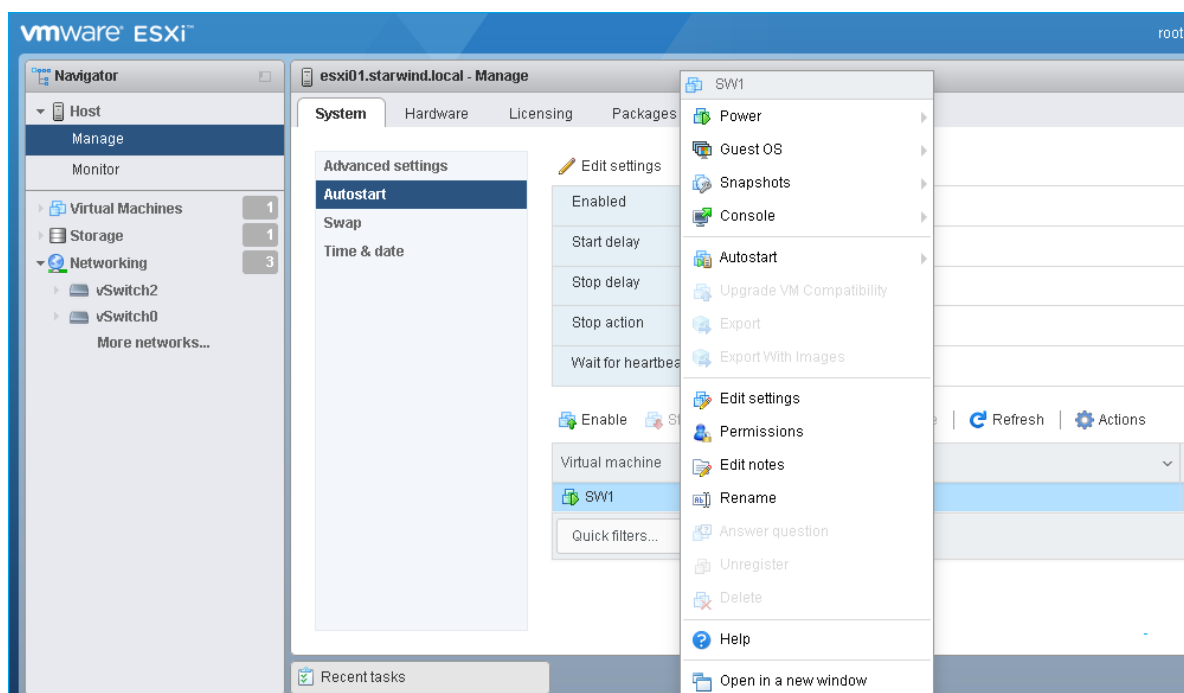
NOTE: When using StarWind with the synchronous replication feature inside of a Virtual Machine, it is recommended not to make backups and/or snapshots of the Virtual Machine with the StarWind VSAN service installed, as this could pause the StarWind Virtual Machine. Pausing the Virtual Machines while the StarWind VSAN service is 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 both ESXi hosts from Manage -> System tab in the ESXi web console. In the appeared window, check Yes to enable the option and choose the stop action as Shut down. Click Save to proceed.



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



3. Complete the actions above on StarWind VM located on all ESXi hosts.
4. Start the virtual machines on all ESXi hosts.

Configuring Starwind Virtual San Vm Settings

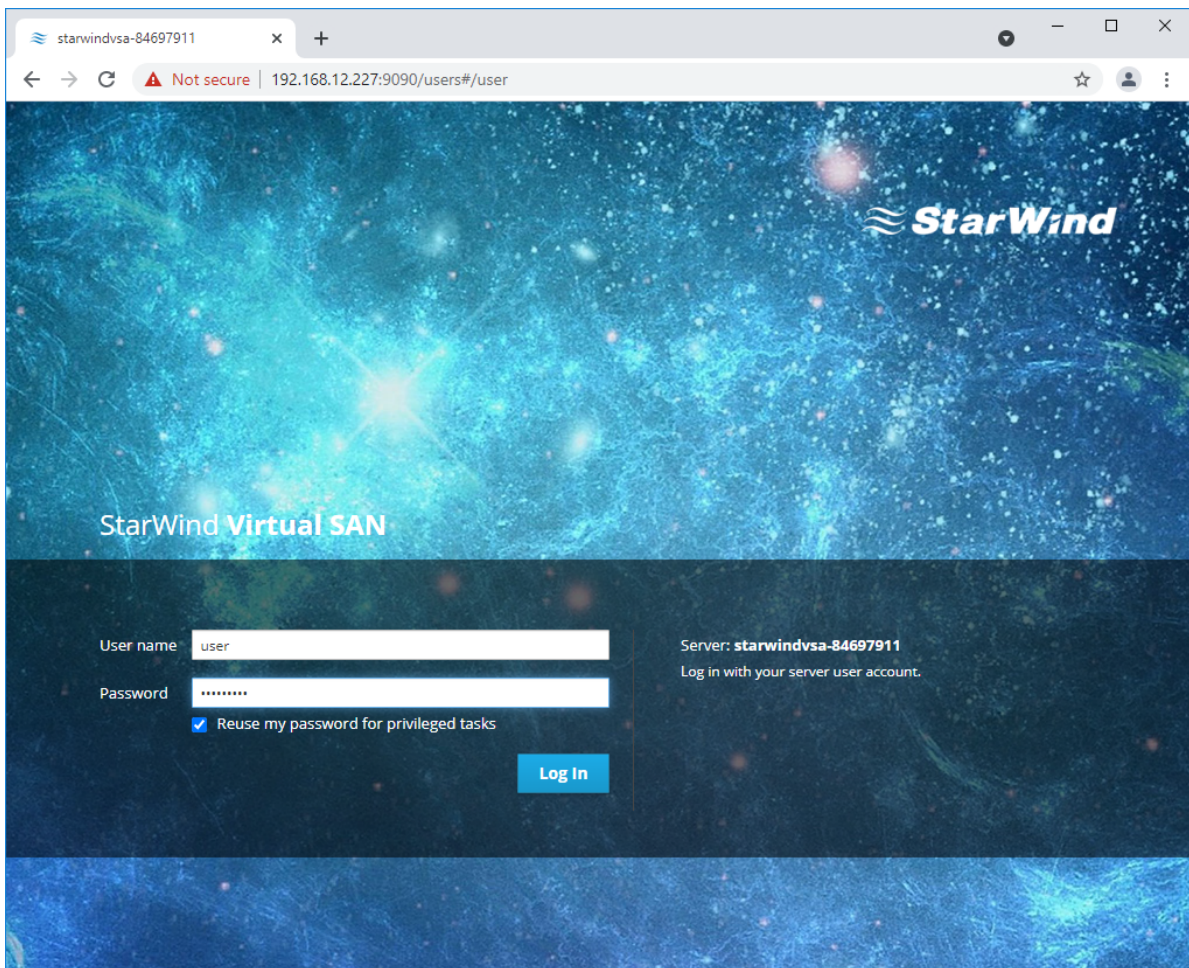
By default, the StarWind Virtual SAN virtual machine receives an IP address automatically via DHCP. It is recommended to create a DHCP reservation and set a static IP address for this VM. In order to access StarWind Virtual SAN VM from the local network, the virtual machine must have access to the network. In case there is no DHCP server, the connection to the VM can be established using the VMware console and static IP address can be configured manually.

1. Open a web browser and enter the IP address of the VM, which it had received via DHCP (or had it assigned manually), and log in to StarWind Virtual SAN for vSphere using the following default credentials:

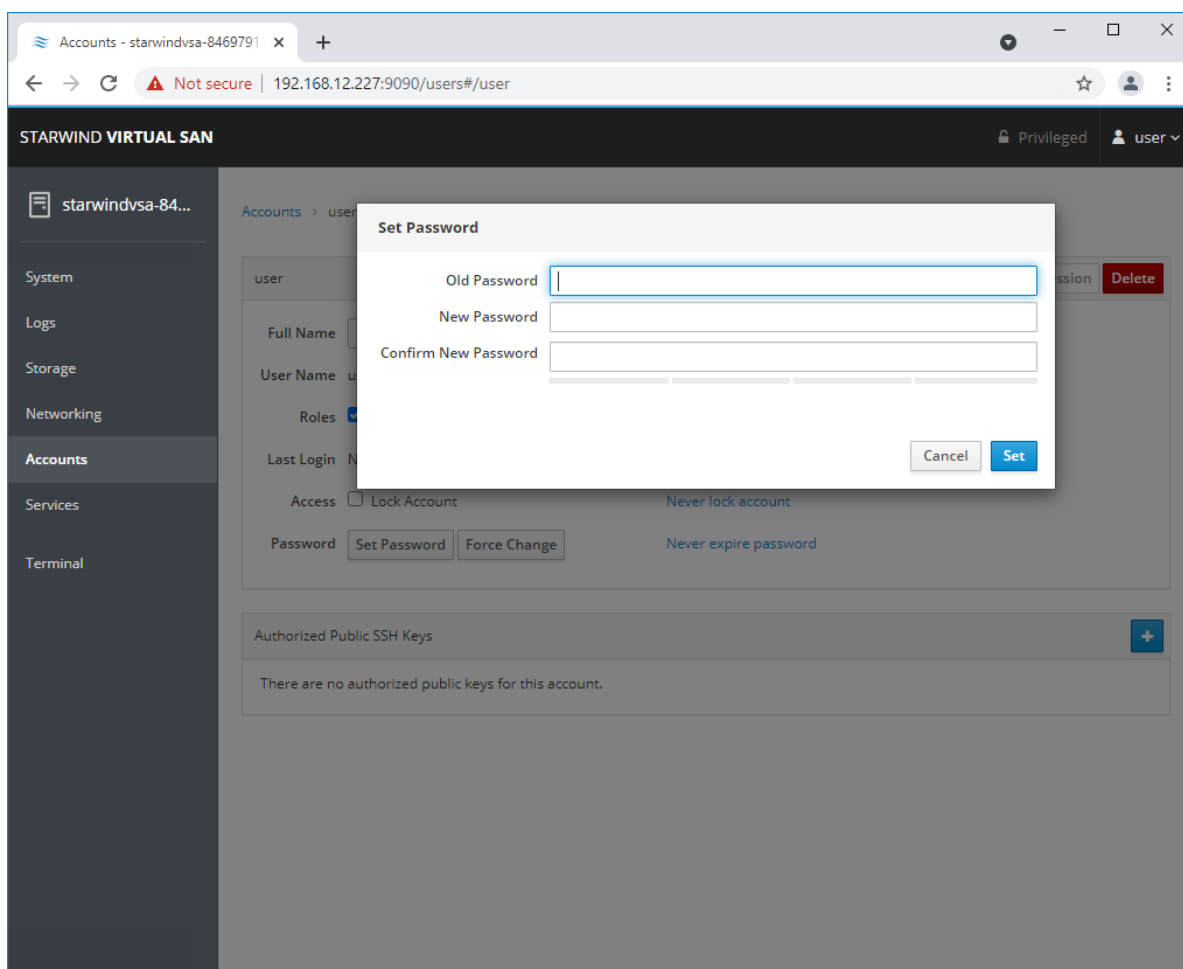
Username: user

Password: rds123RDS

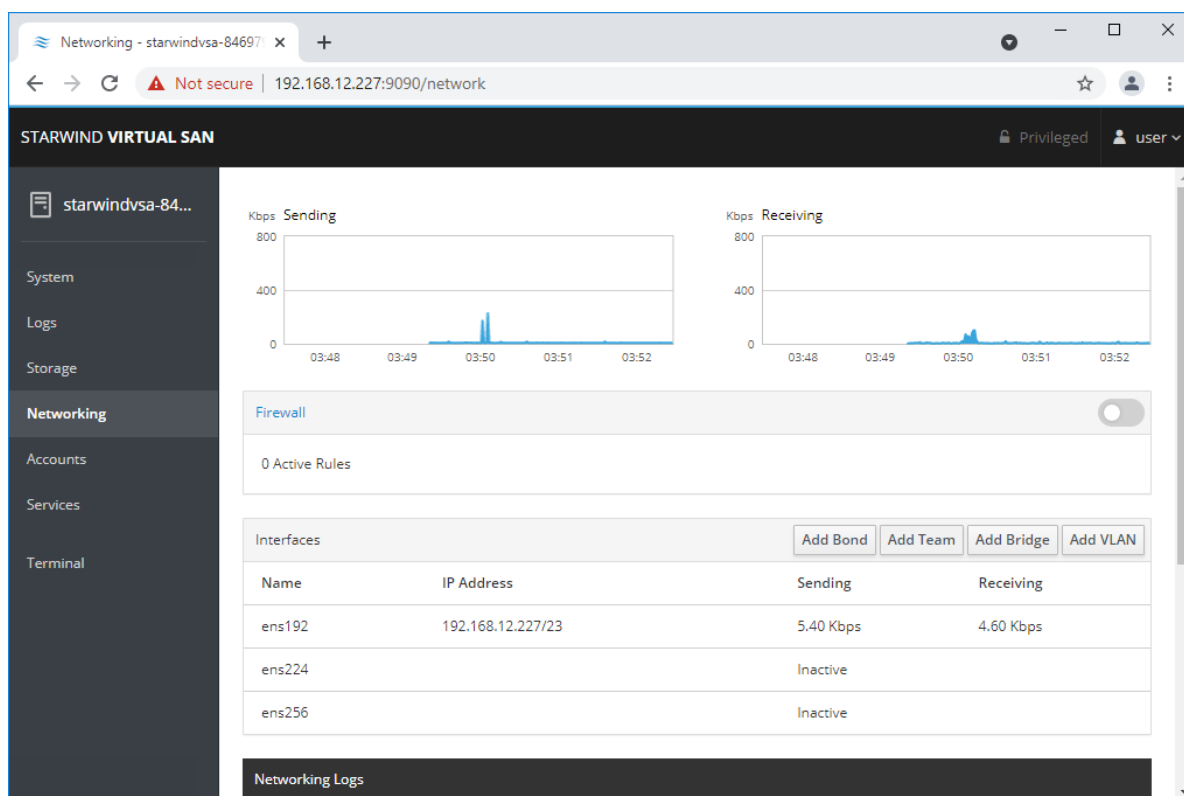
NOTE: Make sure to tick Reuse my password for privileged tasks check box.




2. After the successful login, on the left sidebar, click Accounts.
3. Select a user and click Set Password.




4. On the left sidebar, click Networking.




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



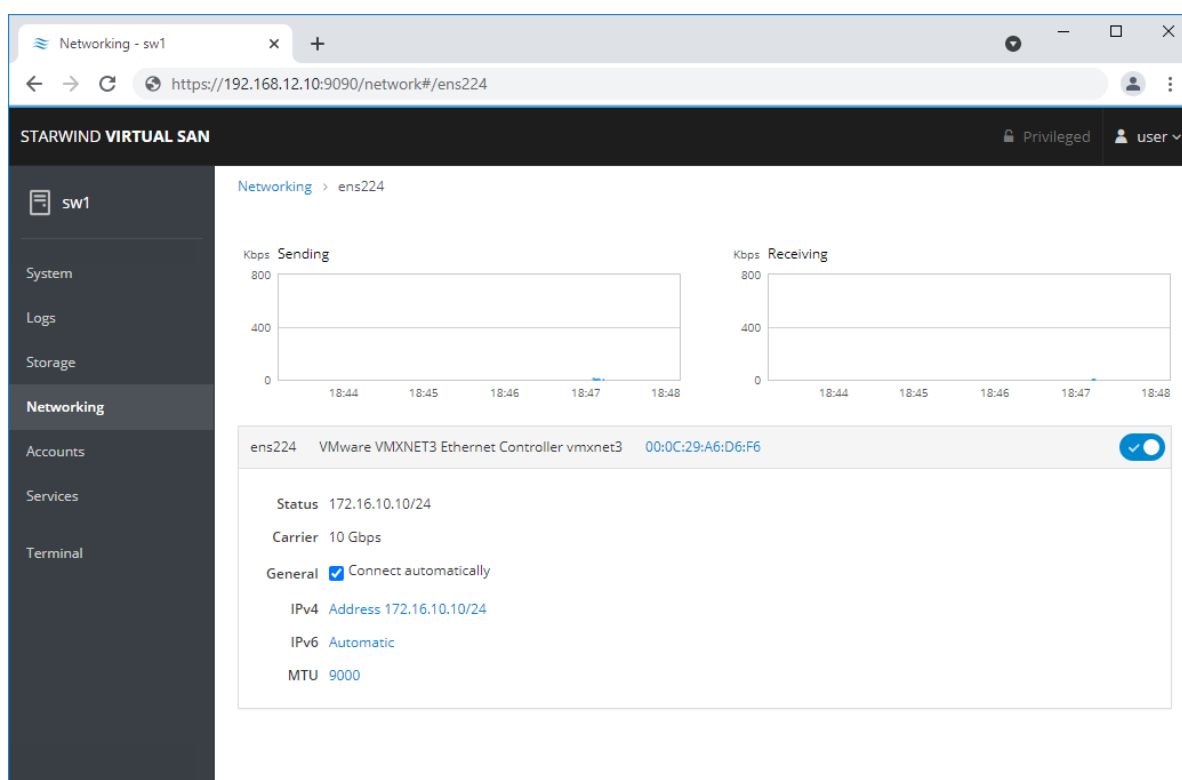
5. Click on Automatic (DHCP) to set the IP address (DNS and gateway – for Management).



6. The result should look like on 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.



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

```

Web console: https://starwindusa-84697911:9898/ or https://192.168.12.227:9898/

starwindusa-84697911 login: Password:
Last login: Tue Aug 17 04:43:59 on tty1
user@starwindusa-84697911 ~$ ls /etc/sysconfig/network-scripts/
ifcfg-ens192  ifdown      ifdown-ppp  ifdown-TeamPort  ifup-bnep  ifup-ipv6  ifup-post  ifup-Team  init.ipv6-global
ifcfg-ens224  ifdown-bnep  ifdown-ipv6  ifdown-routes  ifup-eth  ifup-isdn  ifup-ppp  ifup-TeamPort  network-functions
ifcfg-ens256  ifdown-eth  ifdown-isdn  ifdown-sit  ifup  ifup-ib  ifup-plip  ifup-routes  ifup-tunnel  network-functions-ipv6
ifcfg-lo      ifdown-ib  ifdown-post  ifdown-Team  ifup-aliases  ifup-ippv  ifup-plusb  ifup-sit  ifup-wireless

user@starwindusa-84697911 ~$ ifconfig
ens192: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.12.227 netmask 255.255.254.0 broadcast 192.168.13.255
    ether 00:0c:29:a6:d6:ec txqueuelen 10000 (Ethernet)
    RX packets 3057 bytes 277047 (271.3 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 1067 bytes 1554056 (1.4 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
    device interrupt 19 memory 0xf3a0000-fd3c0000

ens224: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    ether 00:0c:29:a6:d6:f6 txqueuelen 10000 (Ethernet)
    RX packets 2 bytes 120 (120.0 B)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 0 bytes 0 (0.0 B)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

ens256: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    ether 00:0c:29:a6:d6:00 txqueuelen 10000 (Ethernet)
    RX packets 2 bytes 120 (120.0 B)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 0 bytes 0 (0.0 B)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    loop txqueuelen 1000 (Local Loopback)
    RX packets 172 bytes 147586 (144.1 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 172 bytes 147586 (144.1 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

user@starwindusa-84697911 ~$ sudo nano /etc/sysconfig/network-scripts/ifcfg-ens192

```

7. Open the file, corresponding to the Management interface using text editor, for example:

```
sudo nano /etc/sysconfig/network-scripts/ifcfg-ens192
```

8. 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
```

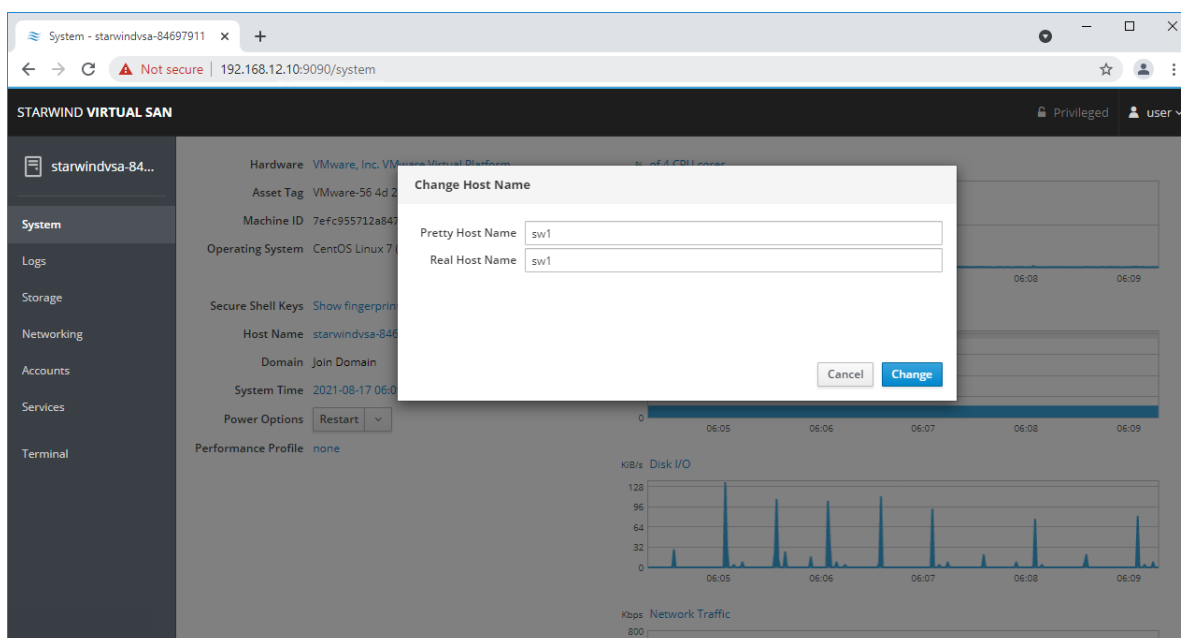
By default, the Management link should have an ens192 interface name. The configuration file should look as follows

```

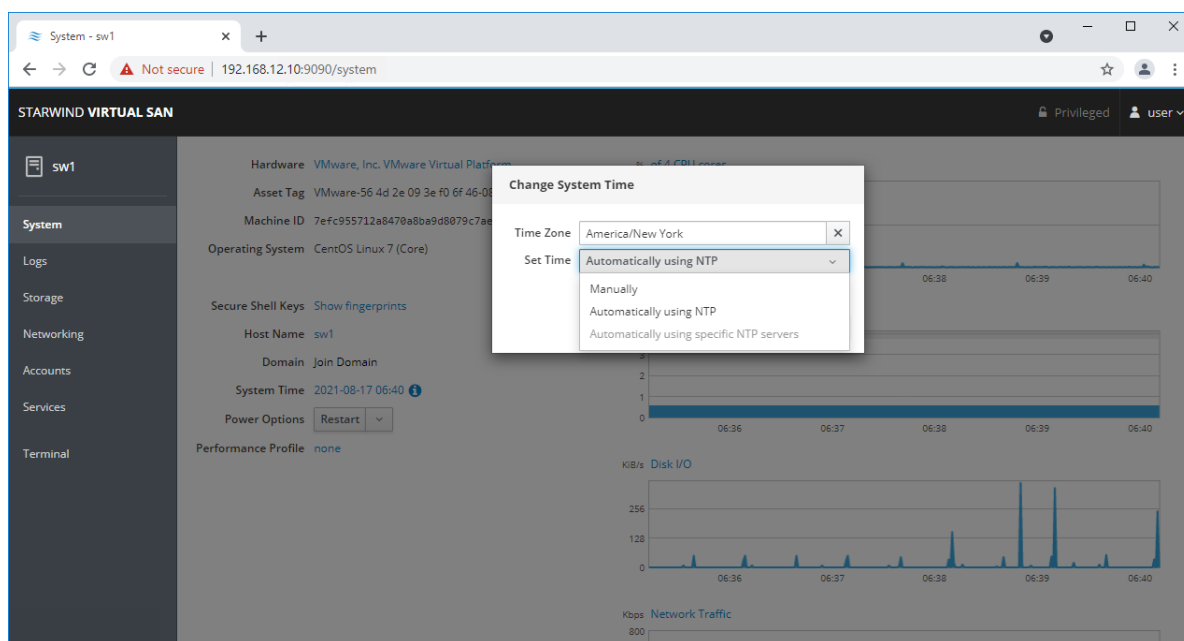
GNU nano 2.3.1 File: /etc/sysconfig/network-scripts/ifcfg-ens192 Modified
TYPE=Ethernet
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=static
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6_INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
IPV6_ADDR_GEN_MODE=stable-privacy
NAME=ens192
UUID=bcf81314-58a7-4bc3-a3ba-76e317a7bd1e
DEVICE=ens192
ONBOOT=yes
IPV6_PRIVACY=no
IPADDR=192.168.12.10
NETMASK=255.255.255.0
GATEWAY=192.168.12.1
DNS1=192.168.1.1
  
```

9. Restart interface using the following cmdlet: `sudo ifdown ens192` , `sudo ifup ens192` or restart the VM.

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



11. Change System time and NTP settings if required




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

Configuring Starwind Management Console


1. Install StarWind Management Console on a 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.

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 Starwind Storage On Top Of Hardware Raid

1. Add a new virtual disk to the StarWind Virtual SAN VM by editing its settings. Make sure it is Thick Provisioned Eager Zeroed. Virtual Disk should be located on the datastore provided by hardware RAID.

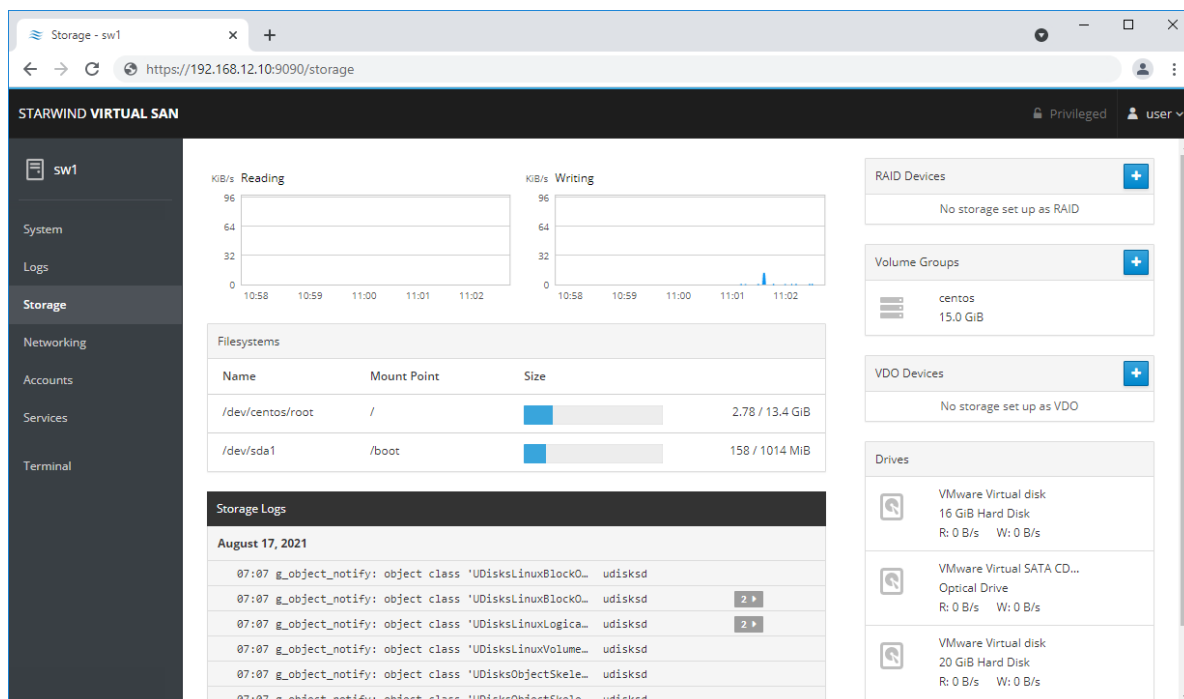
NOTE: Alternatively, the disk can be added to StarWind VSAN VM as RDM. The link to VMware documentation is below:

https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-4236E44E-E11F-4EDD-8CC0-12BA664BB811.html

NOTE: 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.

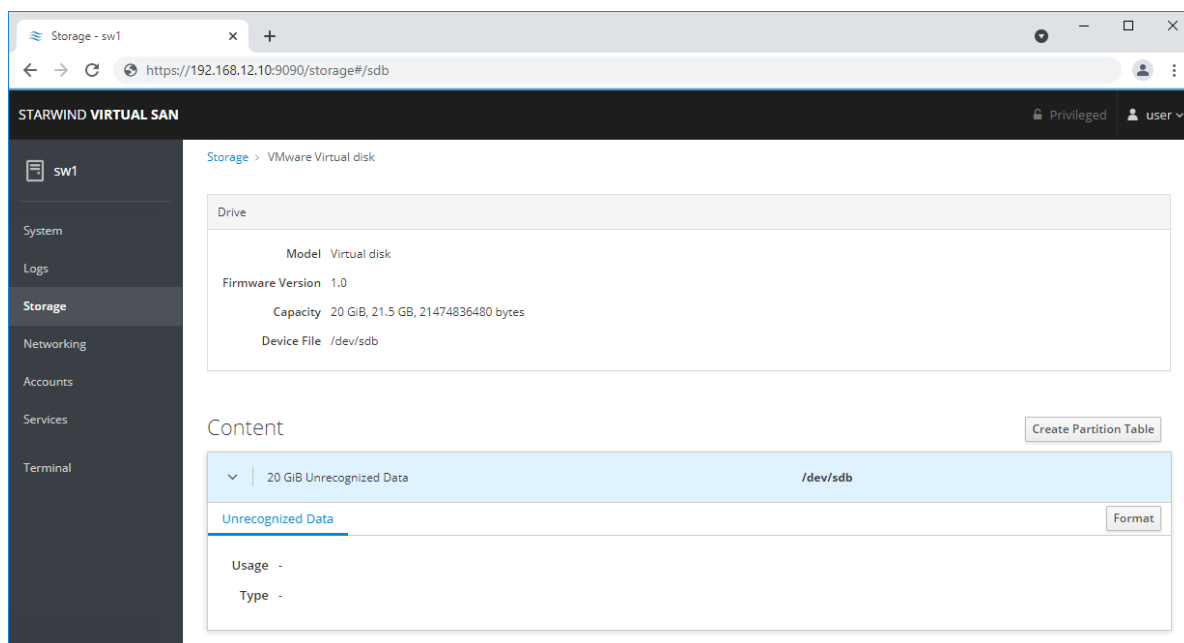
2. Login to StarWind VSAN VM web console and find in the Storage section under Drives

the Virtual Disk that was recently added and choose it.



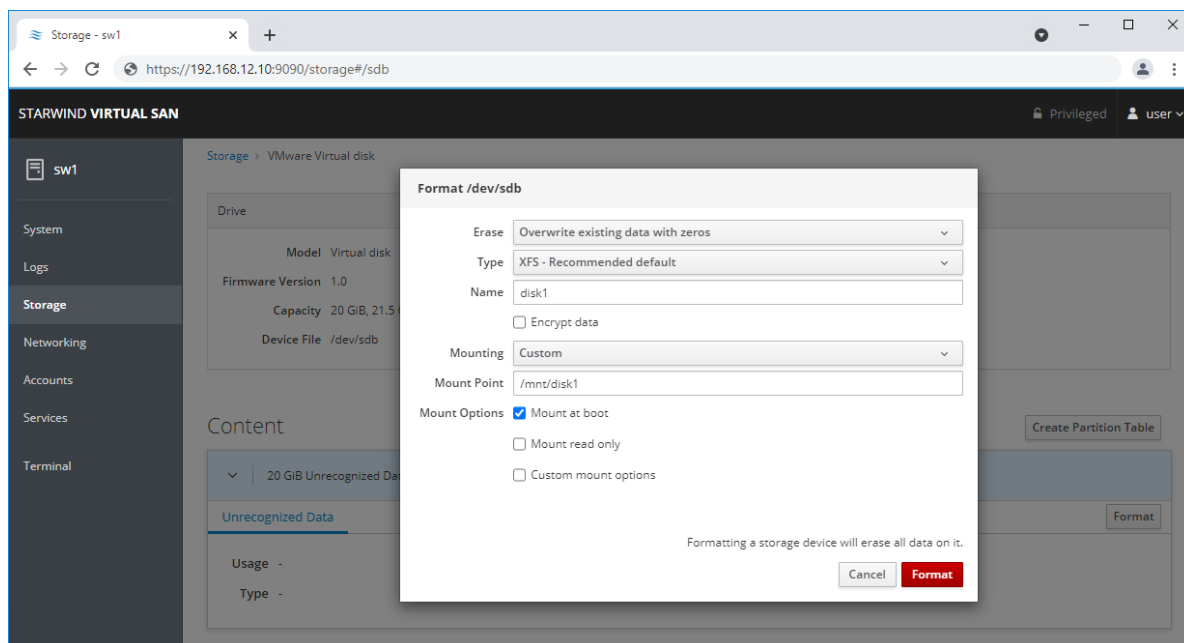
3. The added disk does not have any partitions and filesystem. Press Create partition table and press Format afterward to create the partition and format it.

NOTE: It is not necessary to overwrite data while creating partition.

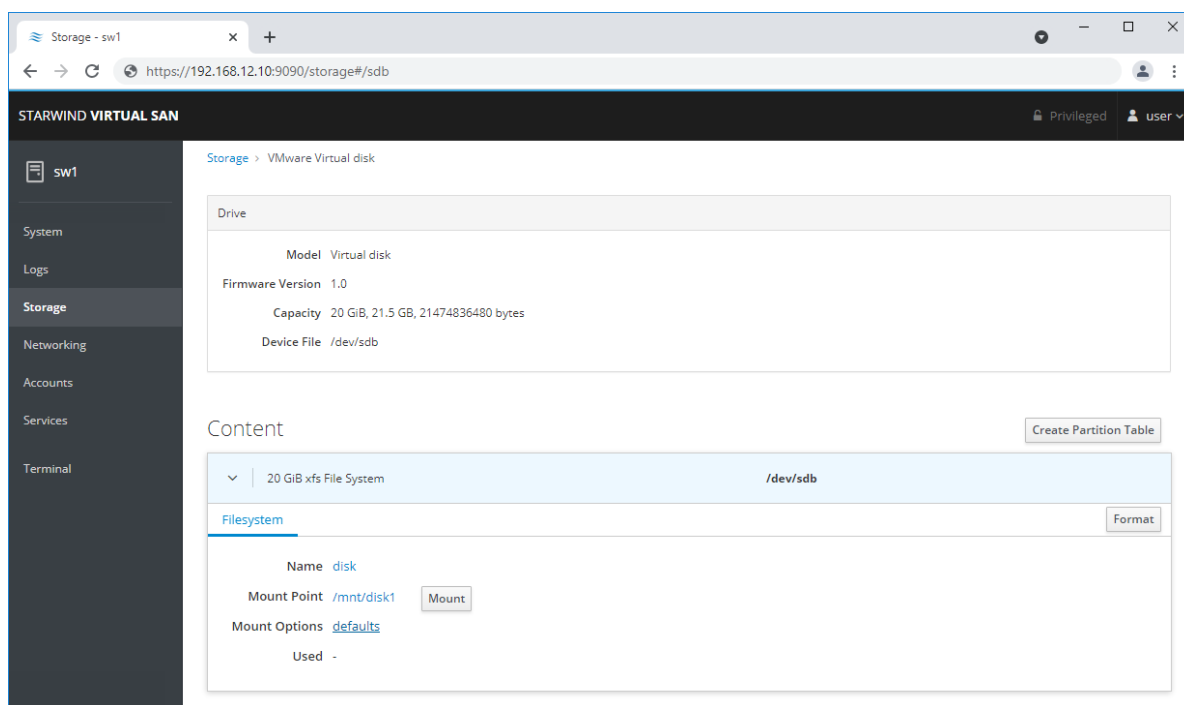


4. Create the XFS partition. Specify the name and erase option. The mount point should be as following: /mnt/%yourdiskname% . Click Format. To enable OS boot when mount


point is missing (e.g., hardware failure), add *nofail* as a boot option.




5. On the storage page of the disk, navigate to the Filesystem tab. Click Mount.



6. Using StarWind Management Console, connect to StarWind Virtual SAN VM and configure storage pool (default storage for StarWind devices) by clicking Yes.



7. Select the disk which was recently mounted.



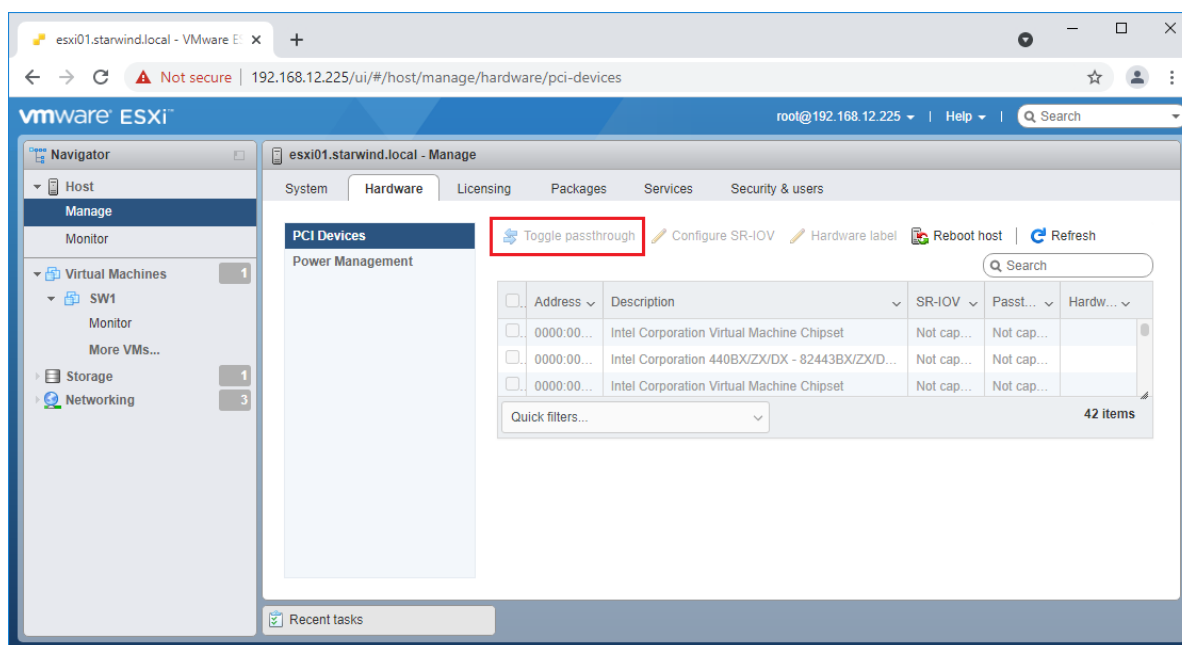
Configuring Starwind Storage On Top Of Software Raid

Make sure that the prerequisites for deploying Software RAID with StarWind Virtual SAN are met:


- The ESXi hosts have all the drives connected through HBA or RAID controller in HBA mode
- StarWind Virtual SAN for vSphere VM is deployed on the ESXi server and turned off
- StarWind Virtual SAN VM is installed on a separate storage device available to the ESXi host (e.g. SSD, HDD etc.)
- HBA or RAID controller can be added via a DirectPath I/O passthrough device to a StarWind Virtual SAN VM without affecting ESXi host work

PCI Device Configuration

1. Login to the ESXi host where StarWind Virtual SAN VM is installed.
2. In the Navigator, go to Manage, and in the Hardware tab, select PCI Devices.
3. Locate the HBA/RAID Controller of the ESXi host. Check the box on the appropriate PCI device. Click Toggle passthrough.



4. Restart ESXi host to make PCI device available to VMs.
5. Right-click on the StarWind Virtual SAN VM to Edit Settings.
7. Click ADD NEW DEVICE. Select PCI Device.



8. Add HBA/RAID Controller to the VM. Reserve memory for the StarWind Virtual Machine. Click OK.

9. Boot StarWind Virtual SAN VM.

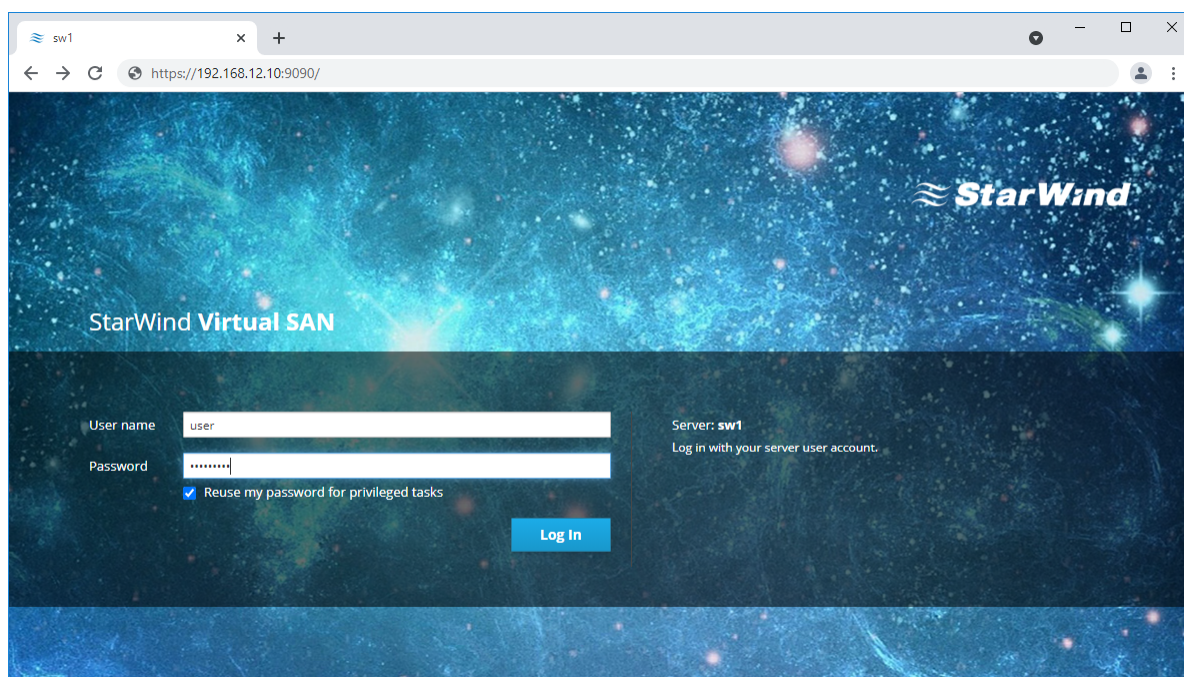
10. Repeat steps 1-8 for all hosts where StarWind Virtual SAN for vSphere is deployed.

11. Login to StarWind Virtual SAN VM via IP. The default credentials:

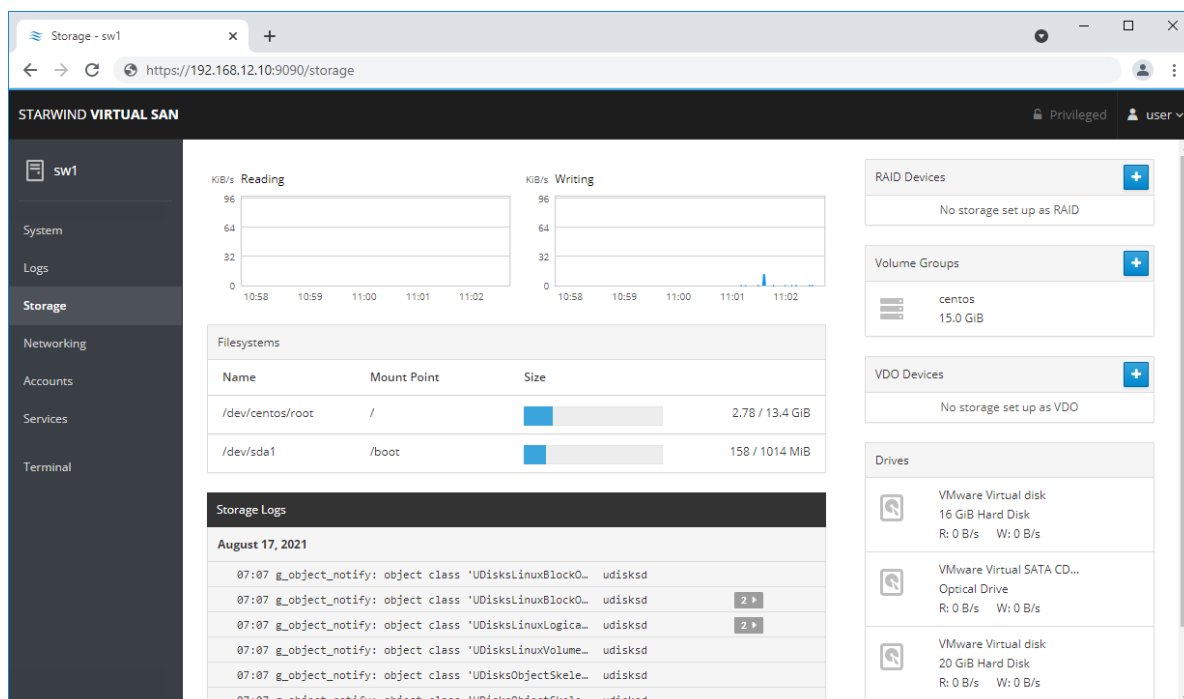
Login: user

Password: rds123RDS

NOTE: Please make sure that the default password is changed.



12. Go to the Storage page. The Drives section shows the drives connected to HBA/RAID Controller (if available). For each disk, create partition table.




13. Click “+” in the RAID Devices section to create Software RAID. (In the current example, RAID 10 will be created with 4 HDD drives). The RAID configuration depends on the number of disks, chunk size, and array level are shown in the table below:

RAID Level	Chunk size for HDD Arrays	Chunk size for SSD Arrays
0	Disk quantity * 4Kb	Disk quantity * 8Kb
5	(Disk quantity - 1) * 4Kb	(Disk quantity - 1) * 8Kb
6	(Disk quantity - 2) * 4Kb	(Disk quantity - 2) * 8Kb
10	(Disk quantity * 4Kb)/2	(Disk quantity * 8Kb)/2

StarWind Software RAID recommended settings can be found here:


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

14. Select the drives to add to the array.




15. After the synchronization is finished, find the RAID array created. Press Create partition table and press Format afterward to create the partition and format it.


NOTE: It is not necessary to overwrite data while creating a partition.




16. Create the XFS partition. Mount point should be as follows: `/mnt/%yourdiskname%` . Select the Custom mounting option and type `noatime`. To enable OS boot when mount point is missing (e.g., hardware failure), add `nofail` as a boot option. Click Format.




17. On the storage page of the disk, navigate to the Filesystem tab. Click Mount.



18. Connect to StarWind Virtual SAN from StarWind Management Console or from Web Console. Click Yes.




19. Select the disk recently mounted.




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.

← Add Device Wizard

Virtual Disk Location

☒ Create a New Virtual Disk

Name: <device name>

Location: My Computer\D\

Size: <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, there is a "Mode" section with three radio button options:

- Write-Back**: Writes are performed asynchronously, actual Writes to Disk are delayed, Reads are cached.
- Write-Through**: Writes are performed synchronously, Reads are cached.
- N/A** (selected): Reads and Writes are not cached.

 Below the mode section is a checkbox labeled "Set Maximum available Size" which is currently unchecked. Underneath that is a "Size:" label followed by a text input field containing "128" and a dropdown menu currently set to "MB". At the bottom right of the dialog 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 "Target Parameters" section contains the following fields and options:

- Choose a Target Attachment Method:** A dropdown menu with "Create new Target" selected.
- Target Alias:** A text input field containing the placeholder "<target alias name>".
- Target Name:** An unchecked checkbox followed by a text input field containing the placeholder "iqn.2008-08.com.starwindsoftware:sw1- <target alias name>".
- Allow multiple concurrent iSCSI Connections:** A checked checkbox.

At the bottom right, there are "Next" and "Cancel" buttons.

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.



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.

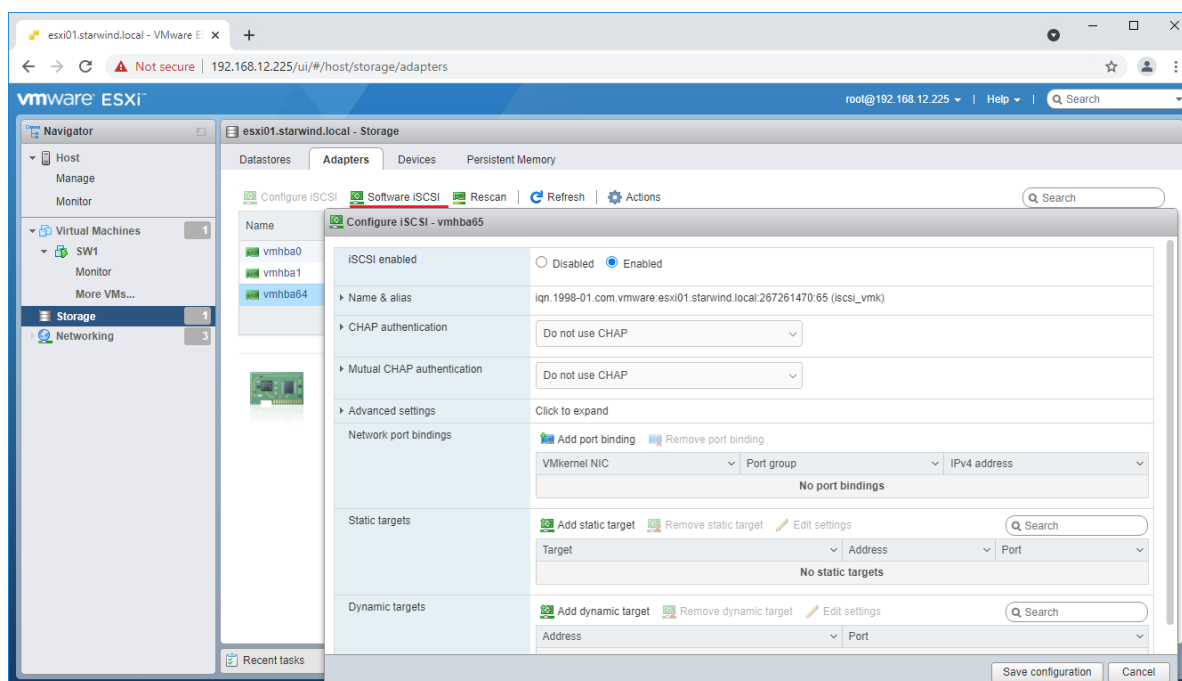
NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:

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

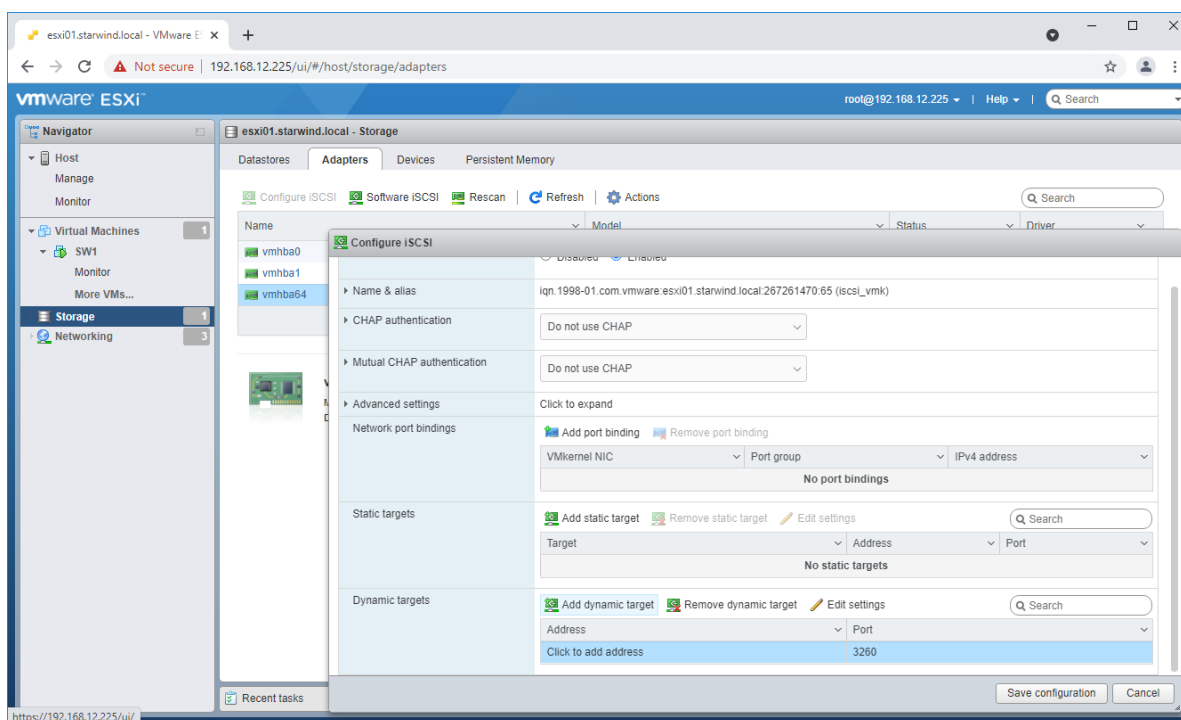
Preparing Datastores

Adding Discover Portals

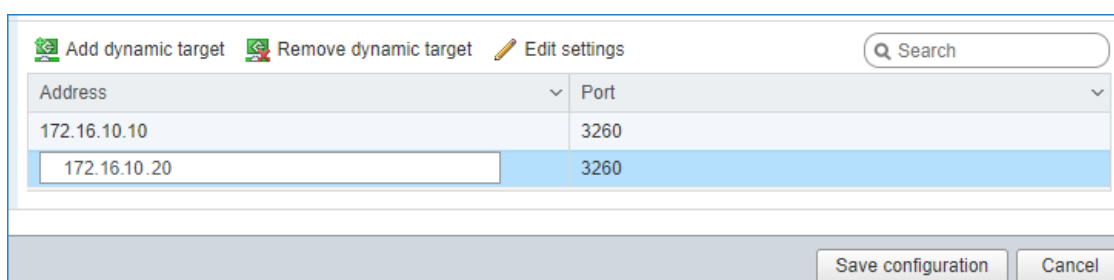
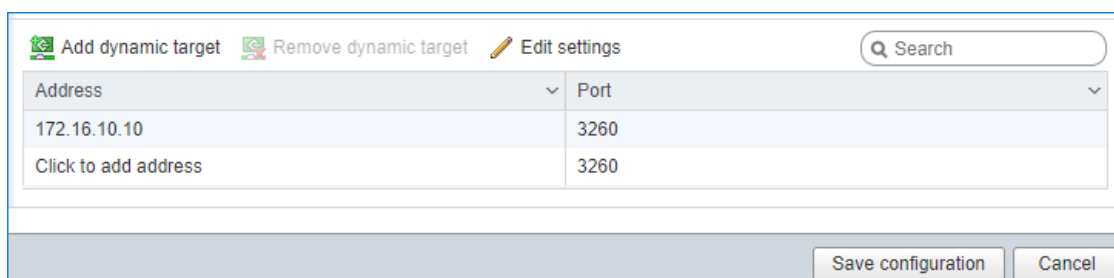
1. To connect the previously created devices to the ESXi host, click on the Storage -> Adapters -> Software iSCSI and in the appeared window choose the Enabled option to enable Software iSCSI storage adapter. Push the Save configuration button.



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.


Dastores Adapters Devices Persistent Memory				
New datastore Increase capacity Rescan Refresh Actions				
Name	Status	Type	Capacity	
Local NECVMWar CD-ROM (mpx.vmhba1:C0:T0:L0)	✓ Normal	CDROM	Unknown	
Local VMware Disk (mpx.vmhba0:C0:T1:L0)	✓ Normal	Disk	50 GB	
Local VMware Disk (mpx.vmhba0:C0:T0:L0)	✓ Normal	Disk	40 GB	
STARWIND iSCSI Disk (eui.60e7edcceeae5b4e)	✓ Normal	Disk	3 GB	
STARWIND iSCSI Disk (eui.4649d8befbe79223)	✓ Normal	Disk	3 GB	

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

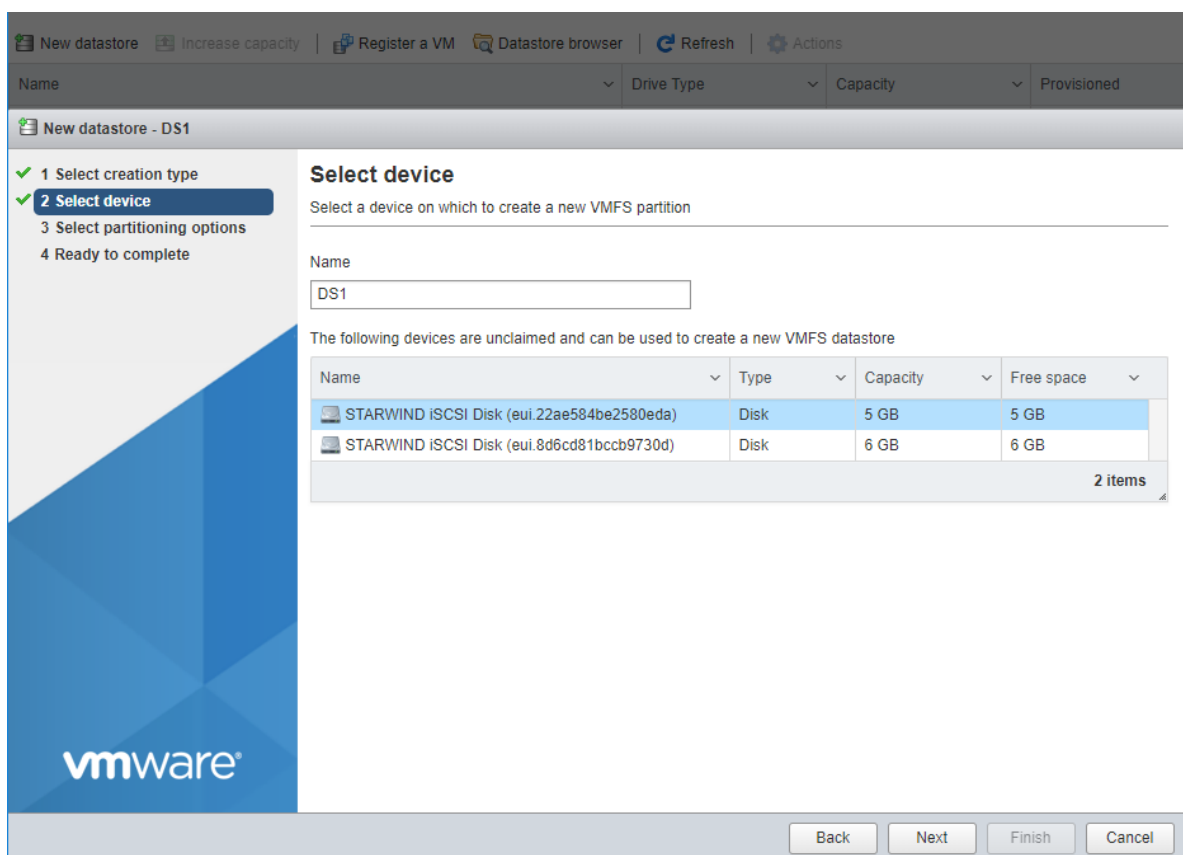
7. Repeat all the steps from this section on the other ESXi host, specifying corresponding IP addresses for the iSCSI subnet.

Creating Datastores


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



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



3. Enter datastore size and click Next.



4. Verify the settings and click Finish.


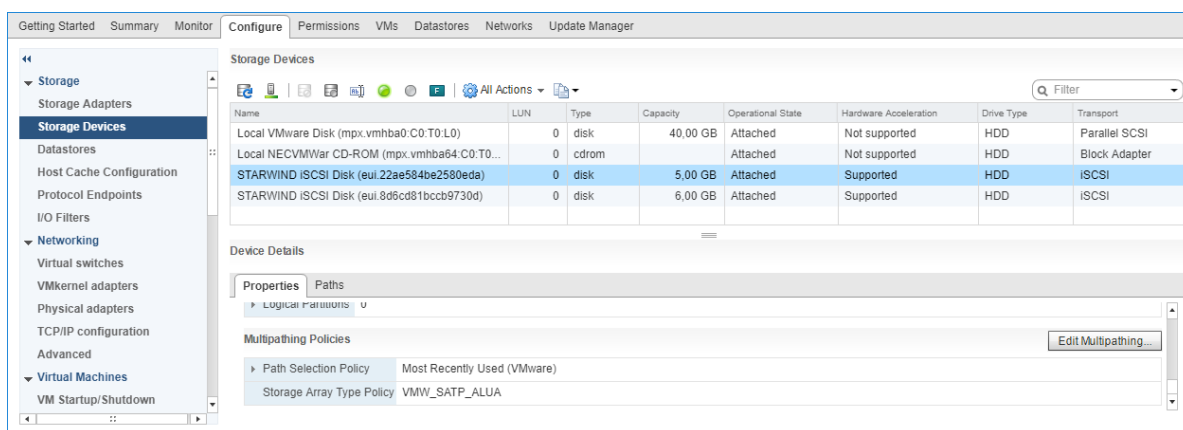
5. Add another Datastore (DS2) in the same way but select the second device for the second datastore.

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

Datastores				
Adapters				
Devices				
New datastore Increase capacity Register a VM Datastore browser Refresh Actions				
Name	Drive Type	Capacity	Provisioned	Free
datastore1 (1)	Non-SSD	32.5 GB	972 MB	31.55 GB
DS1	Non-SSD	4.75 GB	1.41 GB	3.34 GB
DS2	Non-SSD	5.75 GB	1.41 GB	4.34 GB

NOTE: Path Selection Policy changing for Datastores from Most Recently Used (VMware) to Round Robin (VMware) is 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.

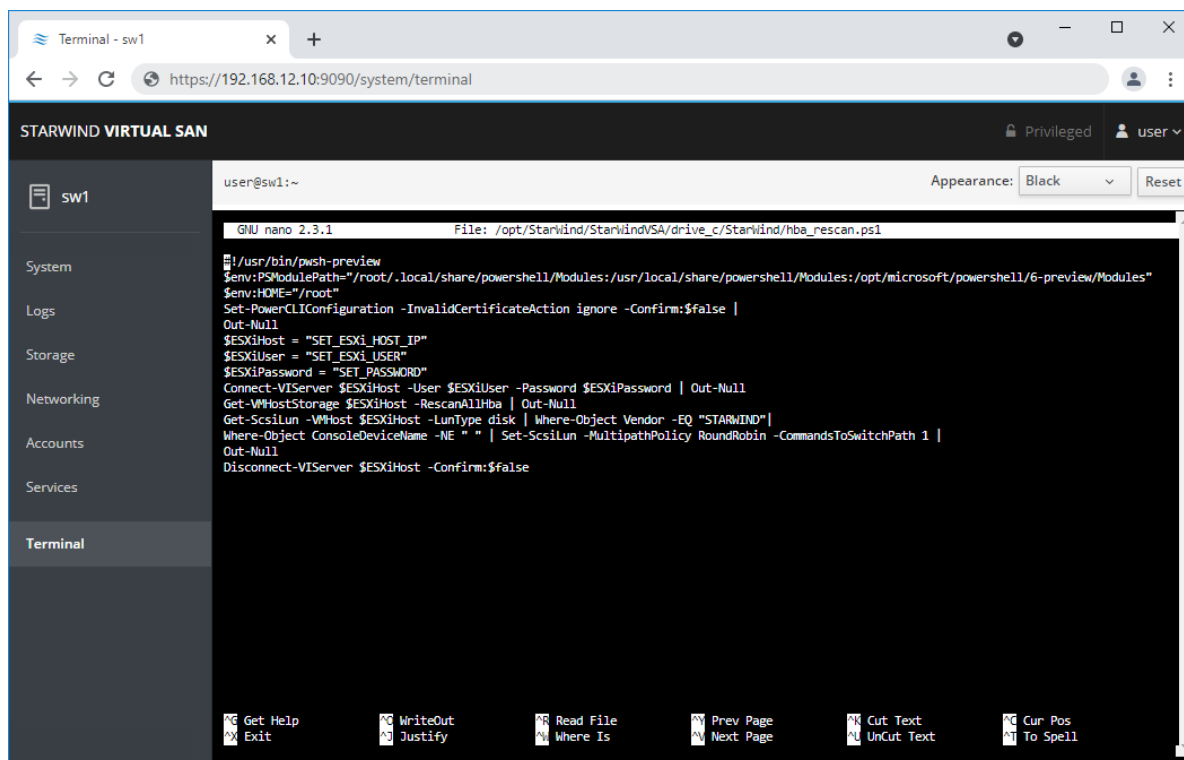
Multipathing configuration can be checked only from vCenter. To check it, click the Configure button, choose the Storage Devices tab, select the device, and click the Edit Multipathing button.



Configuring An Automatic Storage Rescan

1. Open the Terminal page.
2. Edit file `/opt/StarWind/StarWindVSA/drive_c/StarWind/hba_rescan.ps1` with the following command:

`sudo nano /opt/StarWind/StarWindVSA/drive_c/StarWind/hba_rescan.ps1`




3. In the appropriate lines, specify the IP address and login credentials of the ESXi host (see NOTE below) on which the current StarWind VM is stored and running:

```
$ESXiHost = "IP address"
$ESXiUser = "Login"
$ESXiPassword = "Password"
```


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.




Make sure that rescan script is working and execute it from the VM: `sudo /opt/StarWind/StarWindVSA/drive_c/StarWind/hba_rescan.ps1`


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

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.



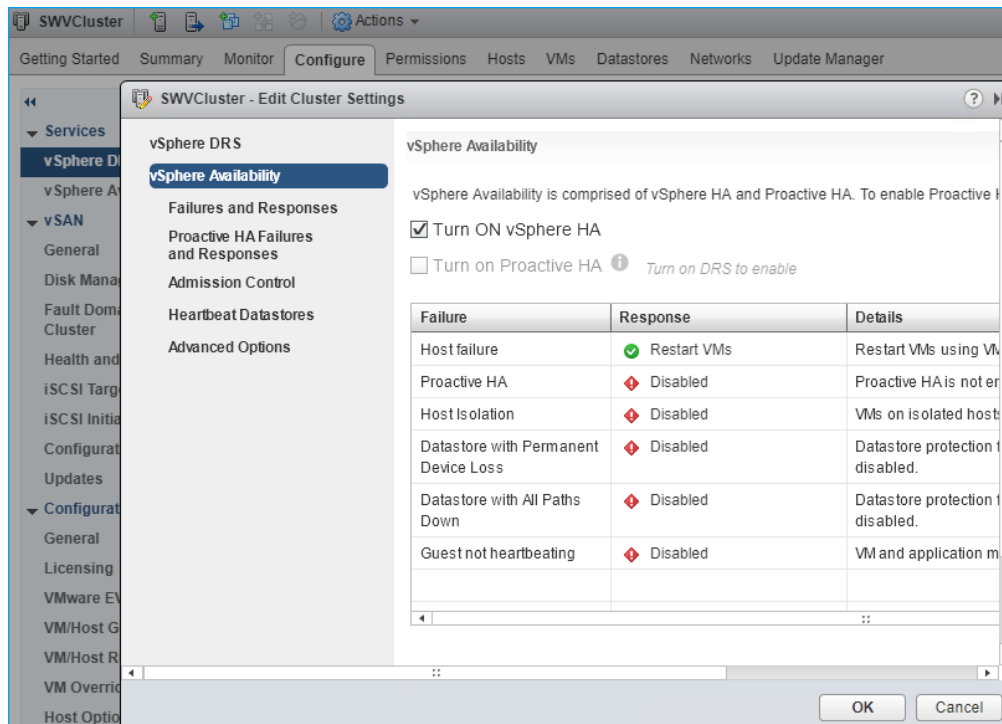
3. To optimize performance change I/O scheduler options according to the article below:
<https://knowledgebase.starwindsoftware.com/guidance/starwind-vsan-for-vsphere-changing-linux-i-o-scheduler-to-optimize-storage-performance/>

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: To provide high availability for clustered VMs, deploy vCenter and add ESXi hosts to the cluster.

Click on Cluster -> Configure -> Edit and check the turn on vSphere HA option if it's






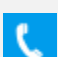
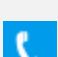
licensed.



Conclusion

The Configuration Guide for StarWind Virtual SAN on VMware vSphere provides a detailed roadmap for IT professionals to enhance their virtualized infrastructure with advanced storage solutions. By leveraging StarWind VSA, organizations can achieve significant improvements in storage performance, efficiency, and resilience, ensuring a robust foundation for their virtualization workloads.

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.