

StarWind Virtual SAN: Enable Deduplication and Compression

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 provides detailed instructions for deploying and configuring StarWind Virtual SAN in a hyperconverged scenario. It focuses on setting up deduplication and compression, covering requirements, pre-configuration steps, and installation processes.

Audience

The primary audience includes system administrators, IT professionals, and network engineers who are responsible for implementing and managing virtualized storage solutions.

Expected Result

Upon successful implementation, users should have a fully functional StarWind Virtual SAN setup with optimized storage performance and capacity, thanks to deduplication and compression. This setup should enhance the efficiency and performance of VMware clusters.

Introduction To Starwind Virtual San For Vsphere

StarWind Virtual SAN for vSphere comes as a prepackaged Linux virtual machine to be installed as a VM on vSphere. It creates a VM-centric and high performing storage pool for a VMware cluster.

This guide describes the deployment and configuration process of the StarWind Virtual SAN with VMware vSphere.

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>

NOTE: Enabling deduplication may impact the resultant storage performance. Please make sure that storage performance meets the requirements for the intended use.

Additional StarWind Virtual SAN VM requirements for deduplication:

Each VDO volume has the following additional memory requirements:

- 370 MB plus an additional 268 MB per each 1 TB of physical storage
- 250 MB per each 1 TB of physical storage, if deduplication is enabled (UDS index)

Example: $370\text{MB} + 4 * 268\text{MB} + 4 * 250\text{MB} = 2442\text{ MB RAM}$ for 4 TB underlying usable physical storage. We recommend assigning StarWind Virtual SAN VM at least 8 GB of RAM (4GB for VM + 4 GB for VDO).

A single VDO volume can be configured to use up to 256 TB of physical storage. All-flash arrays or NVMe drives are strictly recommended as underlying storage when using a combination of StarWind VSAN for vSphere with VDO.


As a general rule, certain storage layers should be placed under VDO and others on top of VDO:

- under VDO: Hardware and software RAID (LVM or mdraid)
- on top of VDO: StarWind Virtual SAN Thick-Provisioned devices (stand-alone and HA)

VDO volume is a thinly provisioned block device, which can run out of space on underlying storage device. It is recommended to use physical storage, which can be extended in future (MD RAID, LVM).

Pre-Configuring The Servers

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

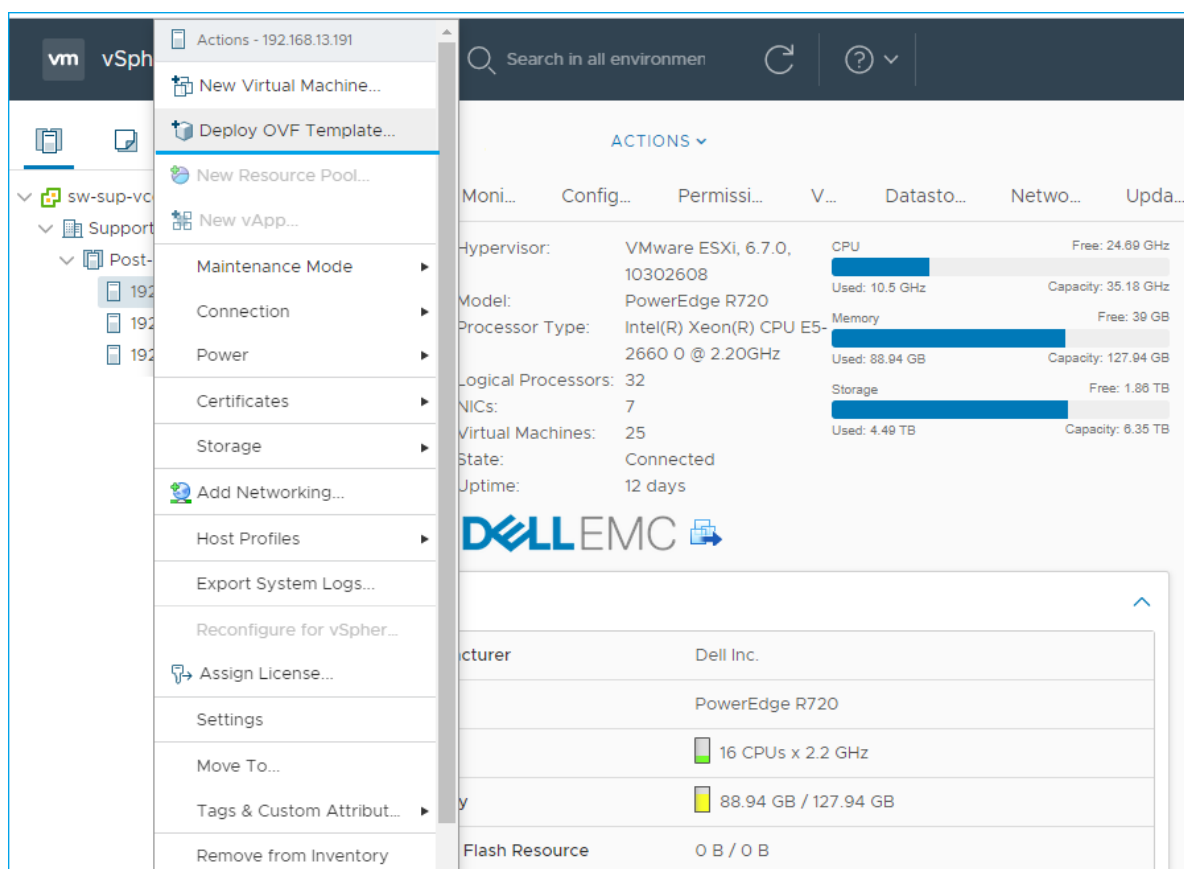


1. ESXi hypervisor should be installed on each host.
2. StarWind Virtual SAN for vSphere VM should be deployed on each ESXi host from an OVF template, downloaded on this page:
<https://www.starwindsoftware.com/release-notes-build-vsan-for-vmware-vsphere>
3. The network interfaces on each node for Synchronization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

Installing 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.
3. Deploy a virtual machine to the vSphere. Right-click on the host and select “Deploy OVF template” from a drop-down menu.



4. In the first step of the wizard, point to the location of the OVF template. Select VM files and click Next.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

☐ URL

[http](#) | [https://remoteserver-address/filetodeploy.ovf](#) | .ova

☒ Local file

Choose Files

No file chosen

⚠

Select a template to deploy. Use multiple selection to select all the files associated with an OVF template (.ovf, .vmdk, etc.)

×

CANCEL

BACK

NEXT

5. Specify the name and location for the StarWind Virtual SAN VM.

Deploy OVF Template

✓ 1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name:

Select a location for the virtual machine.

sw-sup-vcenter.starwind.local

Support

CANCEL

BACK

NEXT

6. Select a resource for the StarWind Virtual SAN VM.

Deploy OVF Template

✓ 1 Select an OVF template

✓ 2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a compute resource

Select the destination compute resource for this operation

Support

Post-Sale

192.168.12.10

192.168.12.20

192.168.12.30

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

7. Review the information about the VM.

8. Select the storage for the VM.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- 6 Select storage**
- 7 Select networks
- 8 Ready to complete

Select storage
Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed ▾

VM Storage Policy: Datastore Default ▾

Name	Capacity	Provisioned	Free	Type
HDD	4.55 TB	3.41 TB	1.25 TB	VM
ISO	851.87 GB	702.82 GB	149.06 GB	NF
SSD	893 GB	485.18 GB	432.23 GB	VM
vCenter-DS	99.75 GB	54.88 GB	44.87 GB	VM

Compatibility

✓ Compatibility checks succeeded.

CANCEL
BACK
NEXT

9. Select networks for the VM.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- 7 Select networks**
- 8 Ready to complete

Select networks

Select a destination network for each source network.

Source Network	Destination Network
ISCSI	ISCSI_VMs
Management	VM Network
Sync	Sync_VMs

3 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL

BACK

NEXT

10. Click Finish to start the deployment process.

11. Add additional network interfaces for iSCSI and SYNC (can be configured for redundancy or 3-way replica, if required)

Edit Settings | sw1

Virtual Hardware
VM Options

ADD NEW DEVICE

> CPU	4		
> Memory	8	GB	
> Hard disk 1	16	GB	
> SCSI controller 0	LSI Logic SAS		
> Network adapter 1	VM Network		<input checked="" type="checkbox"/> Connected
> Network adapter 2	iSCSI_VMs		<input checked="" type="checkbox"/> Connected
> Network adapter 3	Sync_VMs		<input checked="" type="checkbox"/> Connected
> New Network *	Sync_VMs		<input checked="" type="checkbox"/> Connected
Status	<input checked="" type="checkbox"/> Connect At Power On		
Adapter Type	VMXNET 3		
DirectPath I/O	<input checked="" type="checkbox"/> Enable		
MAC Address		Automatic	
> CD/DVD drive 1	Client Device		<input type="checkbox"/> Connected

CANCEL
OK

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

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.

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.

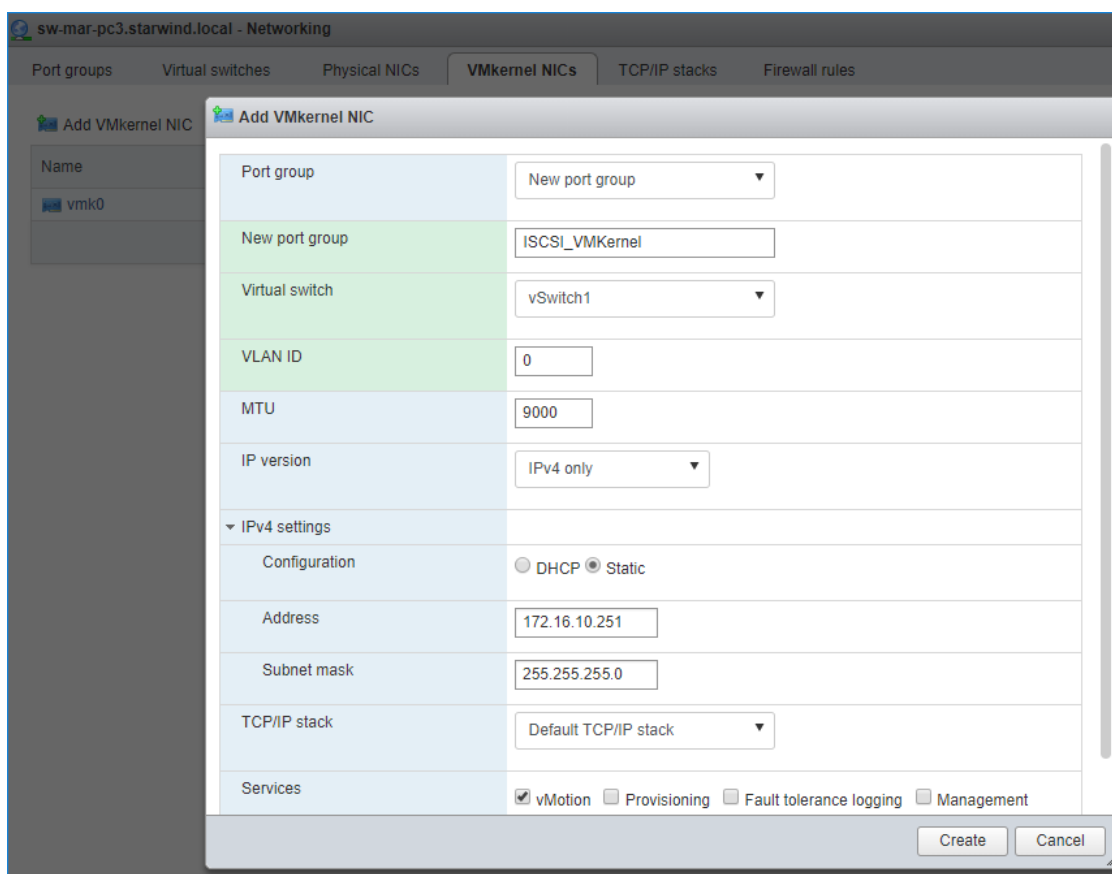
NOTE: Virtual Machine Port Group should be created for both iSCSI/ StarWind Heartbeat and the Synchronization vSwitches. VMKernel port should be created only for iSCSI traffic. Static IP addresses should be assigned to VMKernel ports.

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.

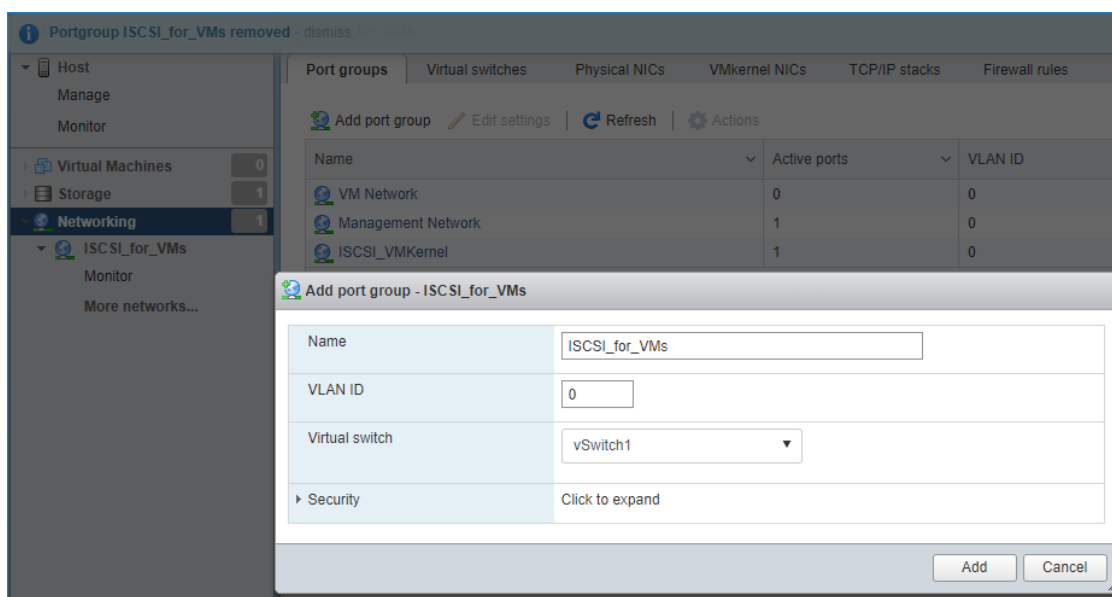
1. Using the VMware ESXi web console, create two standard vSwitches: one for the iSCSI/ StarWind Heartbeat channel (vSwitch1) and the other one for the Synchronization channel (vSwitch2).



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.

Configuring Starwind Virtual San Vm Settings

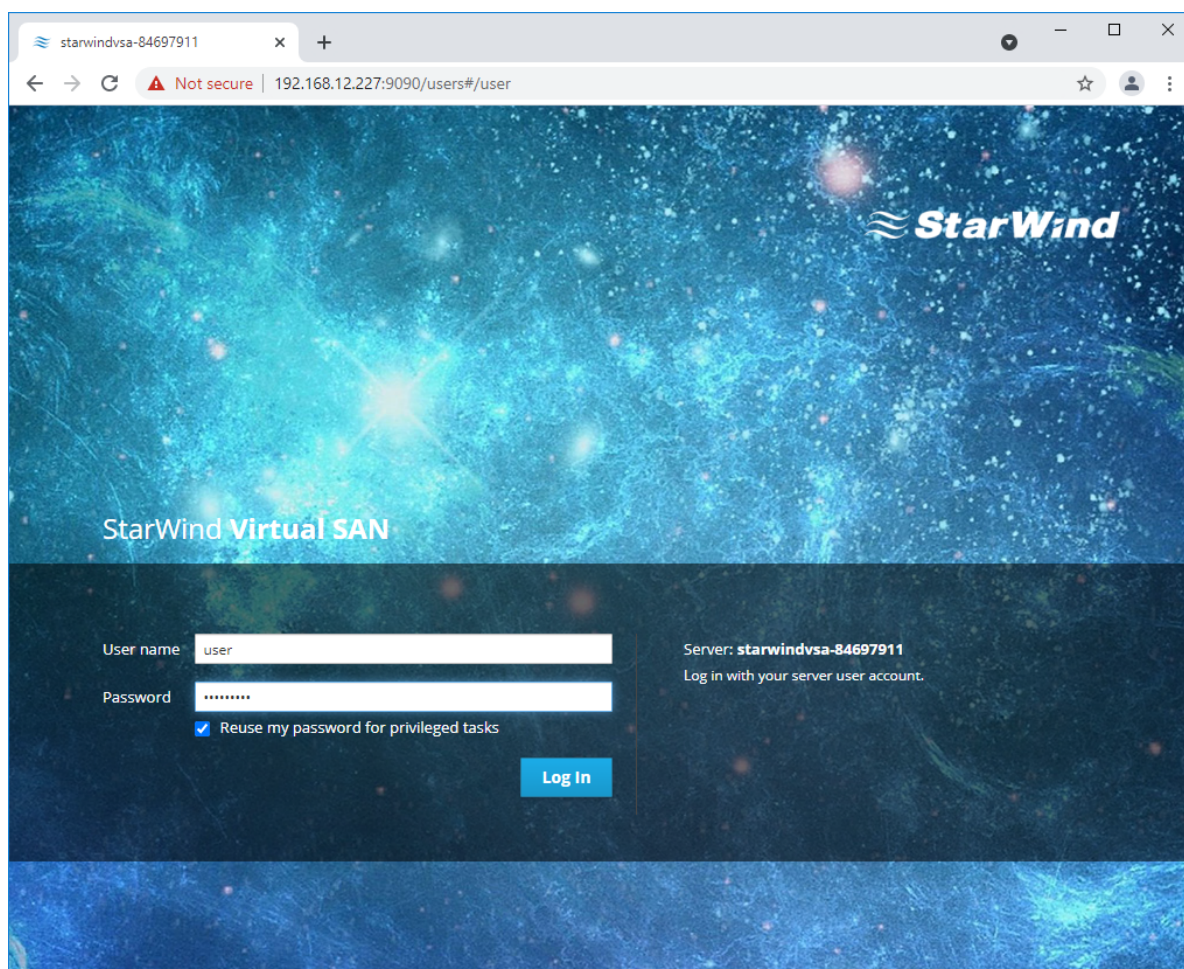
By default, the StarWind Virtual SAN virtual machine receives an IP address automatically via DHCP. It is recommended to create a DHCP reservation and set a static IP address for this VM. In order to access StarWind Virtual SAN VM from the local network, the virtual machine must have access to the network. 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

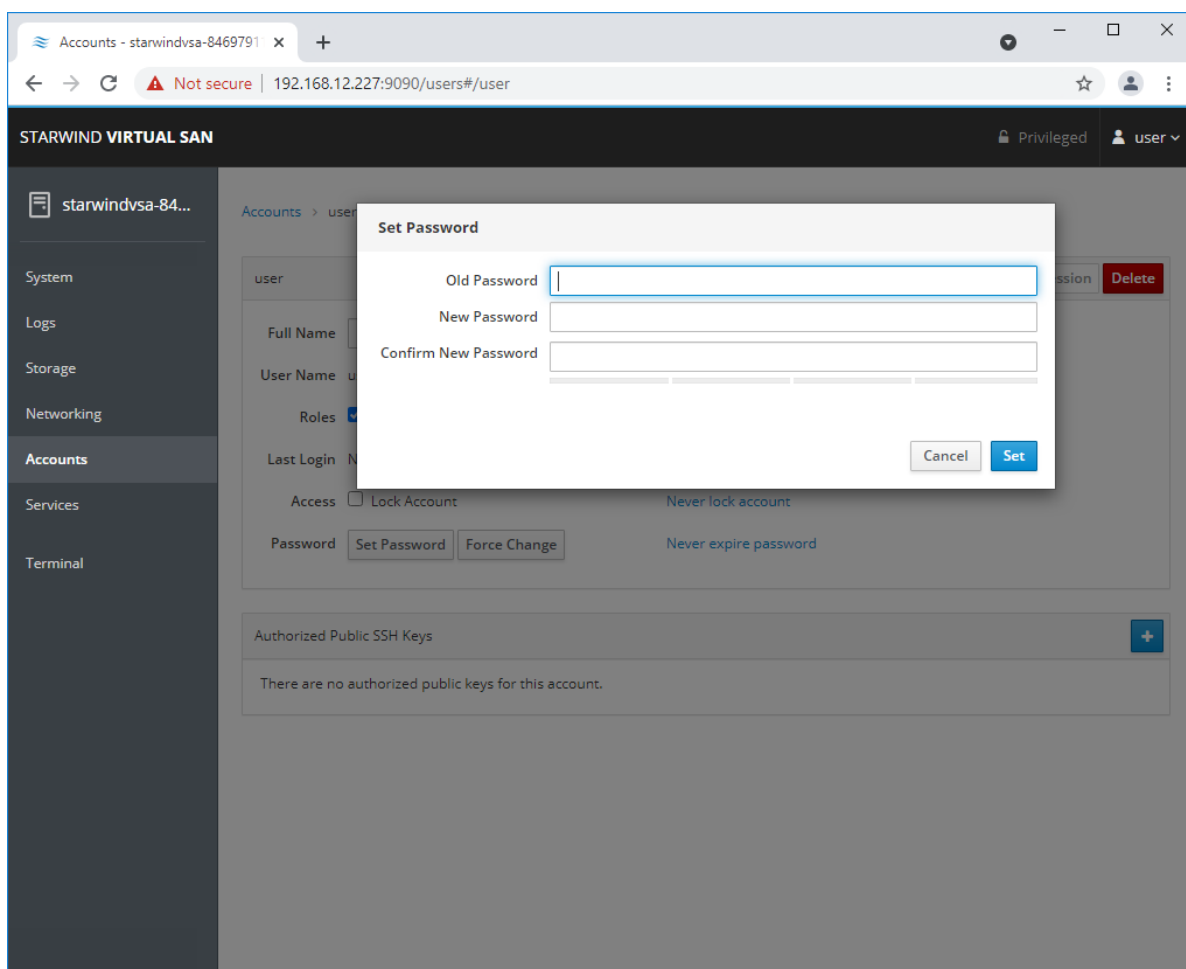
Password: rds123RDS

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.



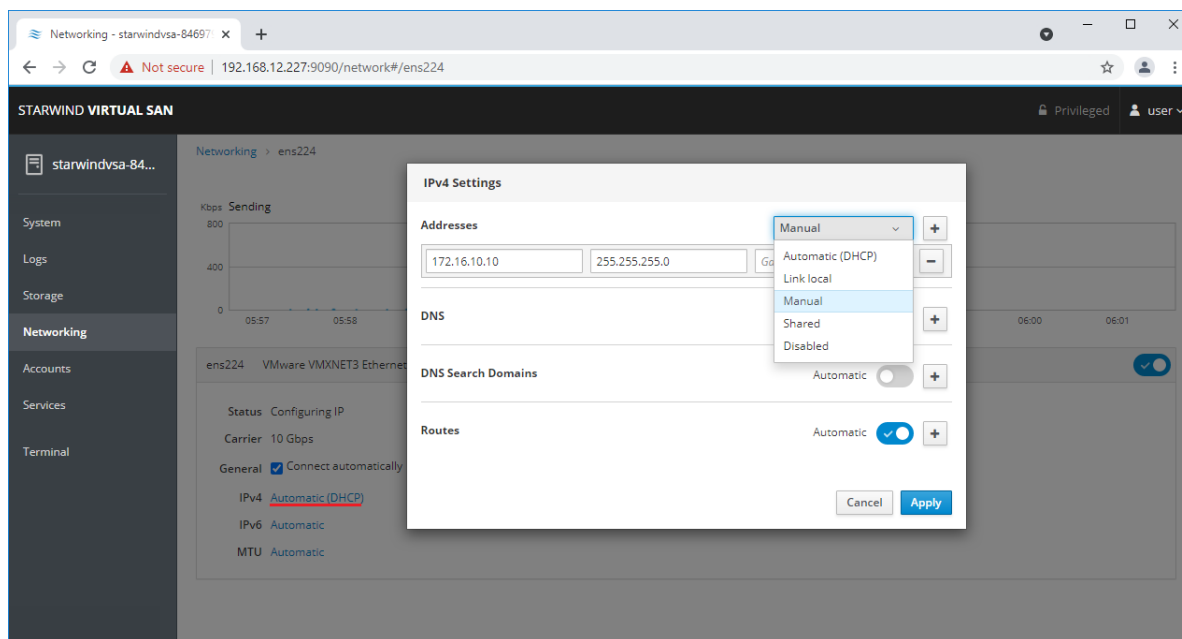
4. On the left sidebar, click Networking.



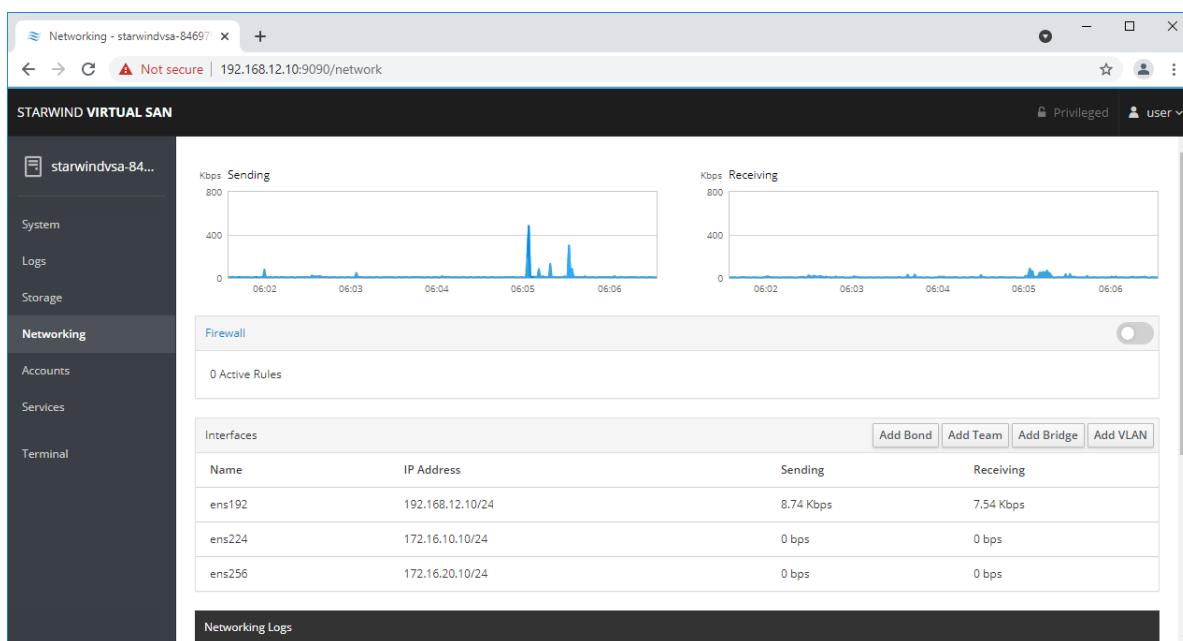
Here, the Management IP address of the StarWind Virtual SAN Virtual Machine, as well as IP addresses for iSCSI and Synchronization networks can be configured. In case the Network interface is inactive, click on the interface, turn it on, and set it to "Connect automatically".



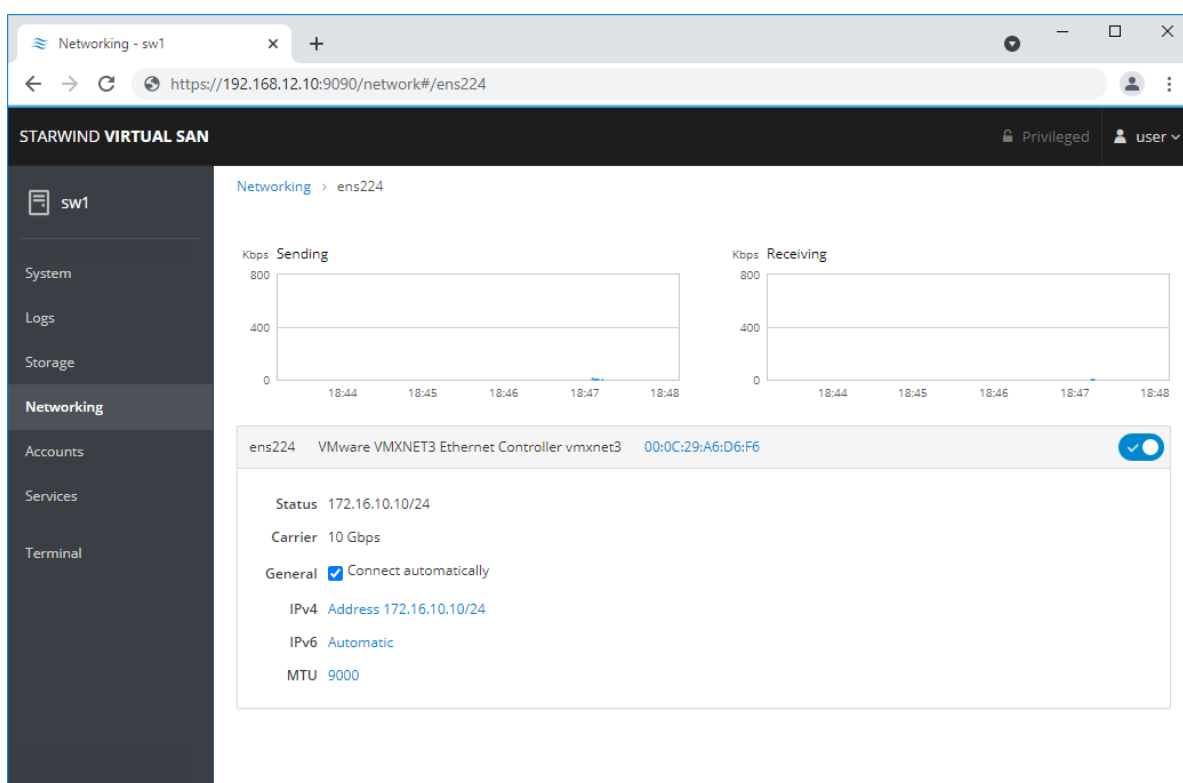
5. Click on Automatic (DHCP) to set the IP address (DNS and gateway – for Management).



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-ipp  ifdown-ppp  ifdown-TeamPort  ifup-bnep  ifup-ipv6  ifup-post  ifup-Team  init.ipv6-global
ifcfg-ens224  ifdown-bnep  ifdown-ipv6  ifdown-routes  ifdown-tunnel  ifup-eth  ifup-isd  ifup-ppp  ifup-TeamPort  network-functions
ifcfg-ens256  ifdown-eth  ifdown-isd  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-ipp  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 08:0c: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 0xd3a0000-fd3c0000

ens224: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    ether 08:0c: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 08:0c: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




```

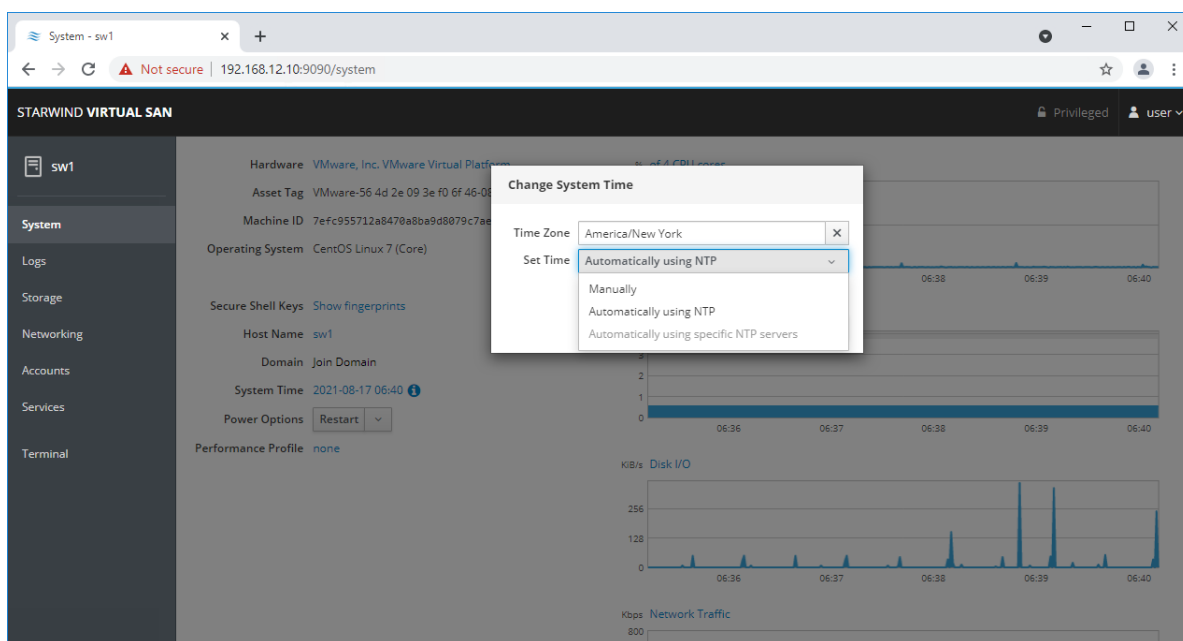
GNU nano 2.3.1 File: /etc/sysconfig/network-scripts/ifcfg-ens192 Modified
TYPE=Ethernet
PROXY_METHOD=none
BROWSER_ONLY=no
BOOTPROTO=static
DEFROUTE=yes
IPV4_FAILURE_FATAL=no
IPV6_INIT=yes
IPV6_AUTOCONF=yes
IPV6_DEFROUTE=yes
IPV6_FAILURE_FATAL=no
IPV6_ADDR_GEN_MODE=stable-privacy
NAME=ens192
UUID=bcf81314-58a7-4bc3-a3ba-76e317a7bd1e
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 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.

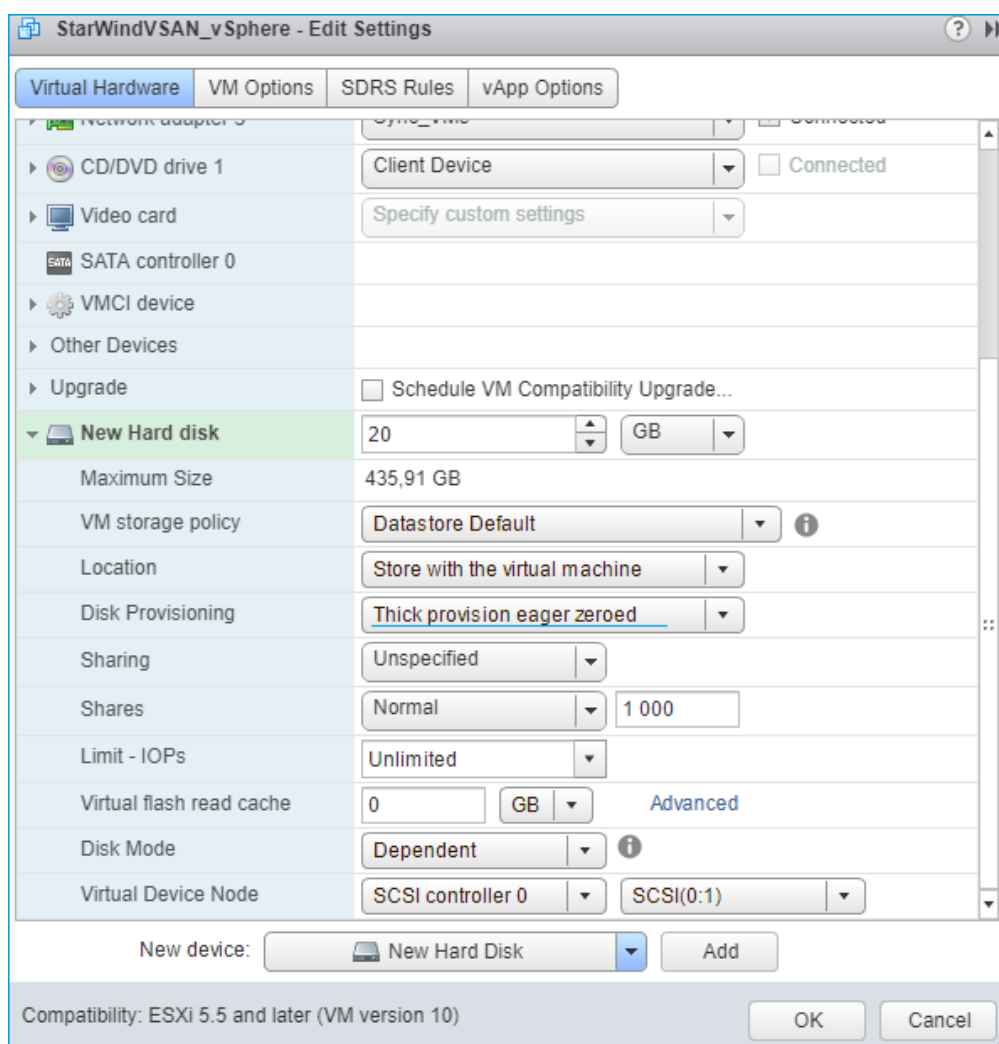
Configuring Storage

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. Make sure it is Thick Provisioned Eager Zeroed. Virtual Disk should be located on the datastore provided by hardware RAID.



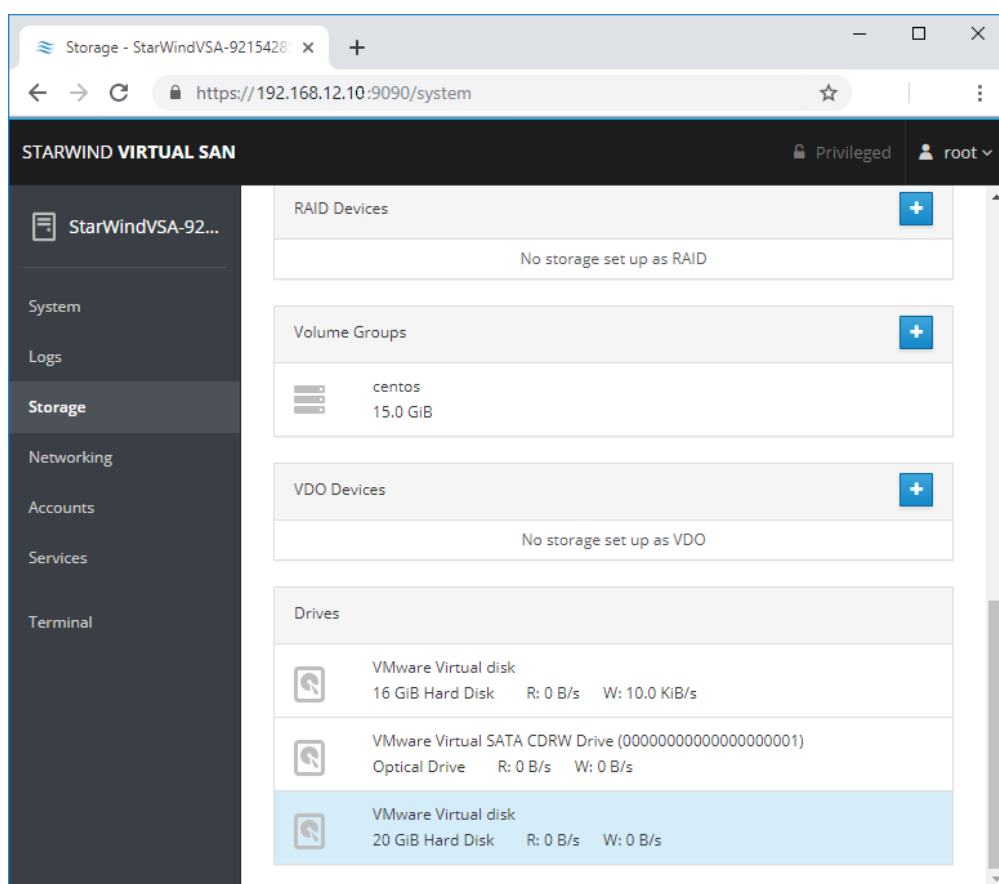
NOTE: Alternatively, the disk can be added to StarWind VSAN VM as RDM. The link to

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.

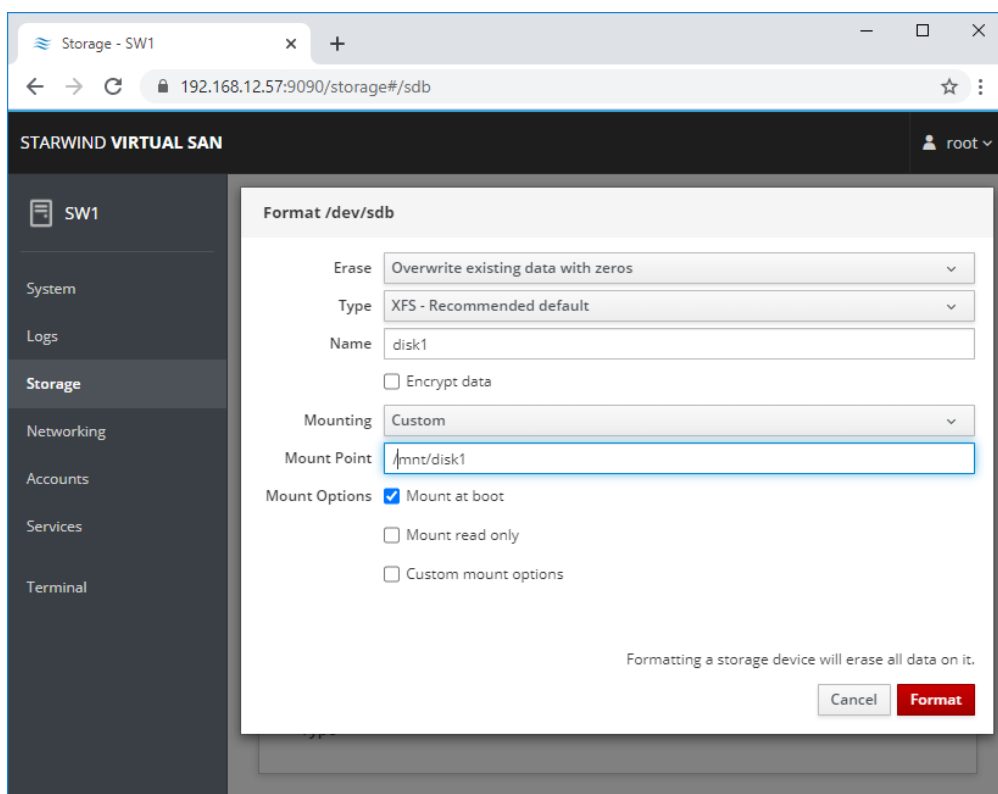


3. The added disk does not have any partitions and filesystem. Press Create partition table and press Format afterward to create the partition and format it.


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.



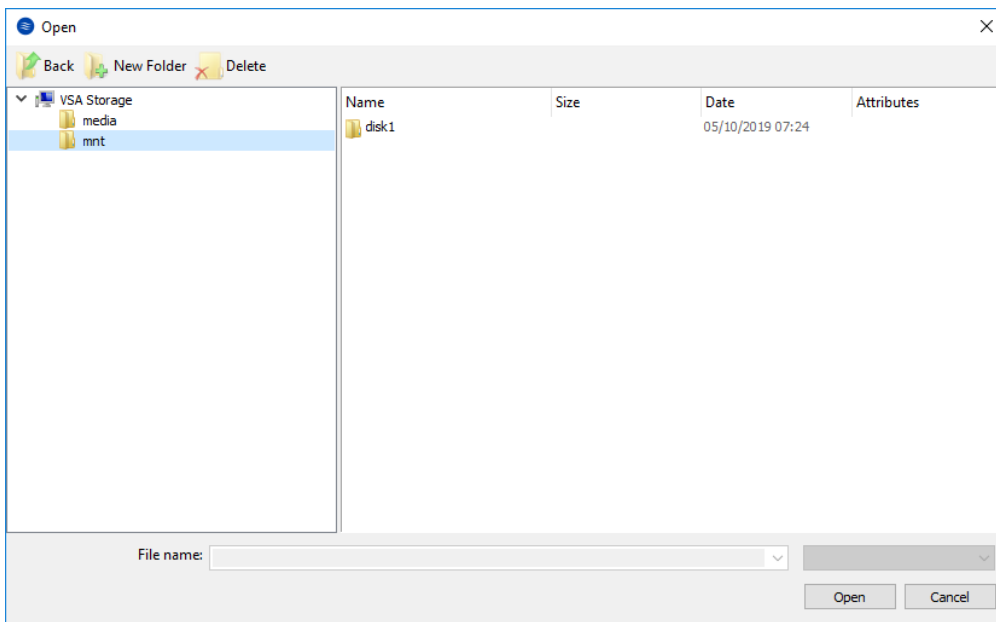
5. On the storage page of the disk, navigate to the Filesystem tab. Click Mount.



6. Connect to StarWind Virtual SAN from the StarWind Management Console. Click Yes.



7. Select the disk which was recently mounted.



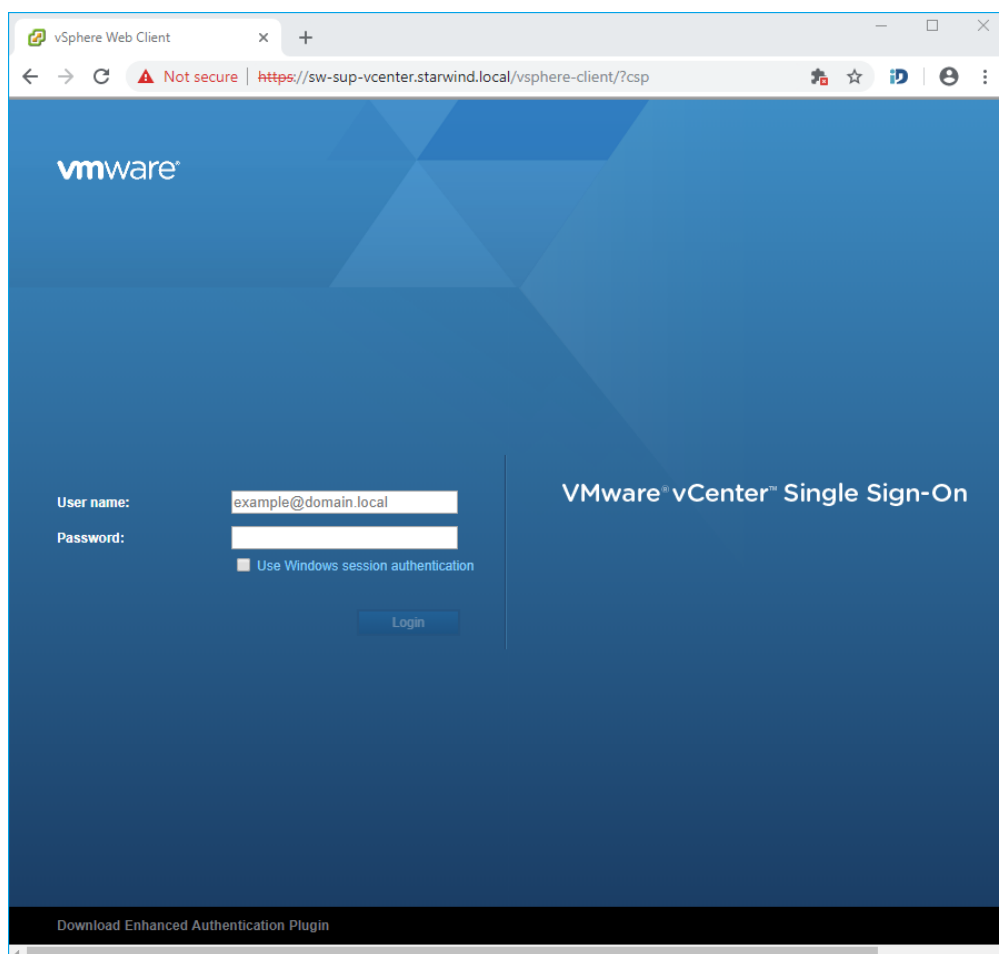
Configuring Starwind Storage On Top Of Software Raid

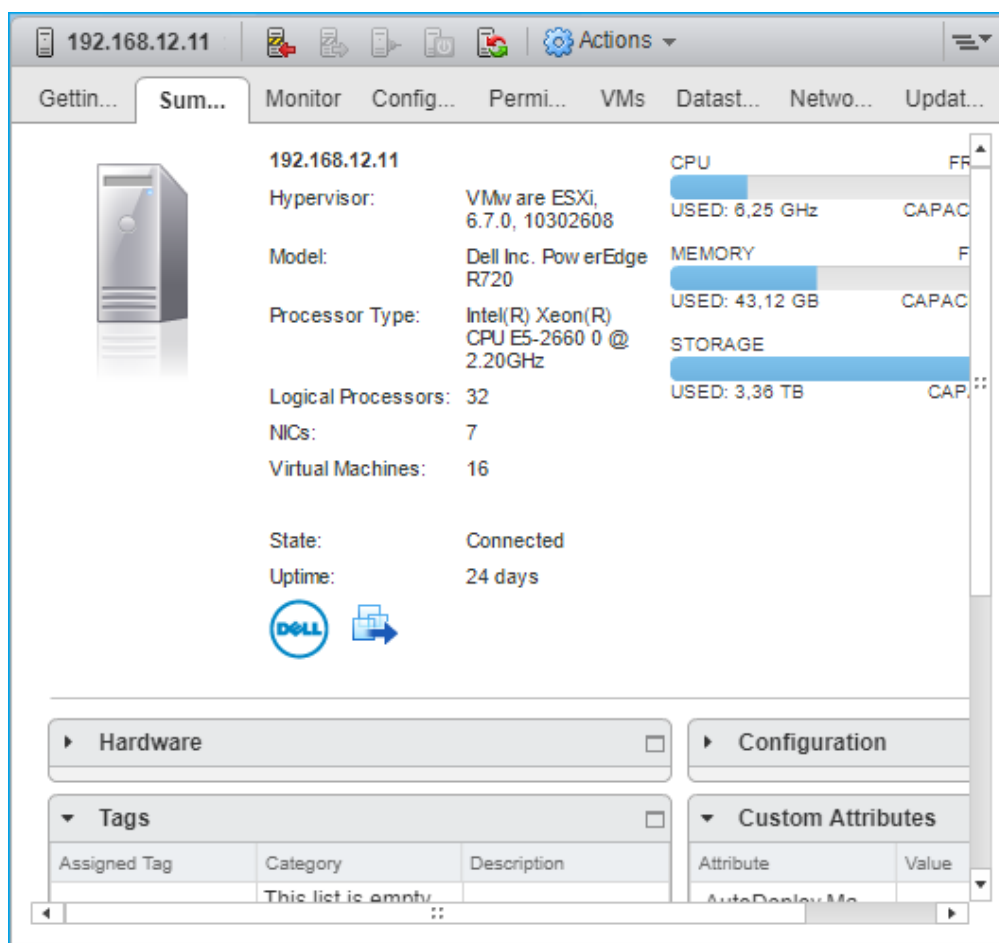
Make sure that the prerequisites for deploying Software RAID with StarWind Virtual SAN are met:

- the ESXi hosts have all the drives connected through HBA or RAID controller in HBA mode
- StarWind Virtual SAN for vSphere is installed on the ESXi server
- StarWind Virtual SAN must be installed on a separate storage device available to the ESXi host (e.g. SSD, HDD etc.)
- HBA or RAID controller will be added via a DirectPath I/O passthrough device to a StarWind VM
- vCenter is installed in the environment to manage ESXi hosts

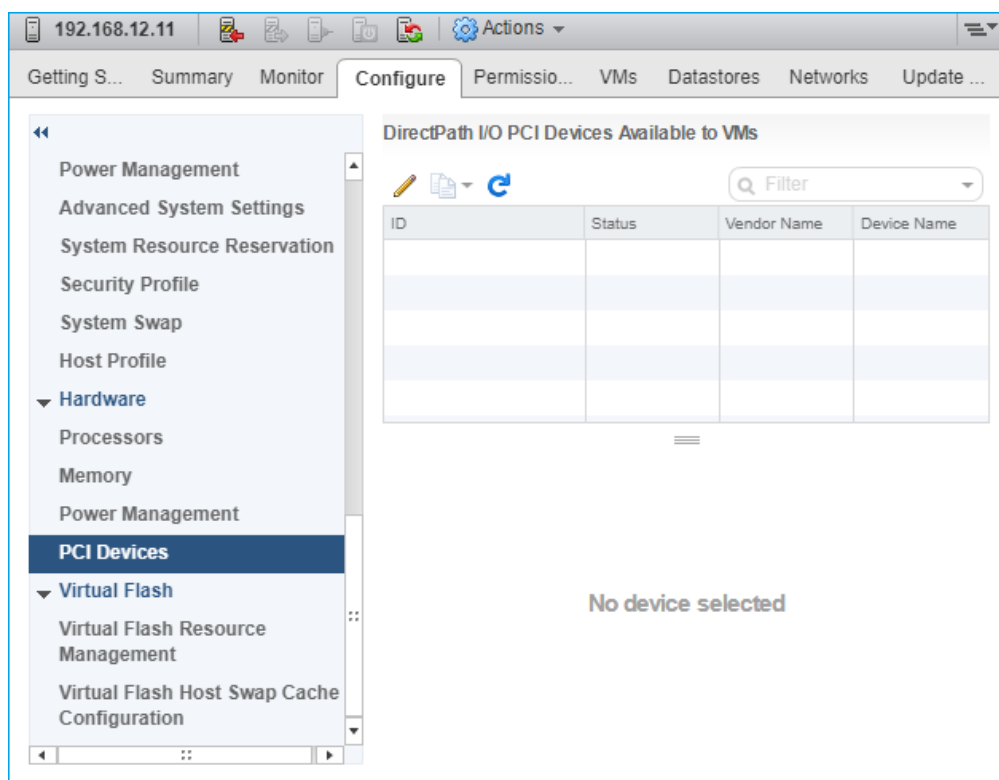
PCI Device Configuration

1. Login to the vCenter via vSphere Client. Select the host where StarWind Virtual SAN VM is installed.







2. Go to the Configure tab. In the Hardware section, select PCI Devices. Click Edit.



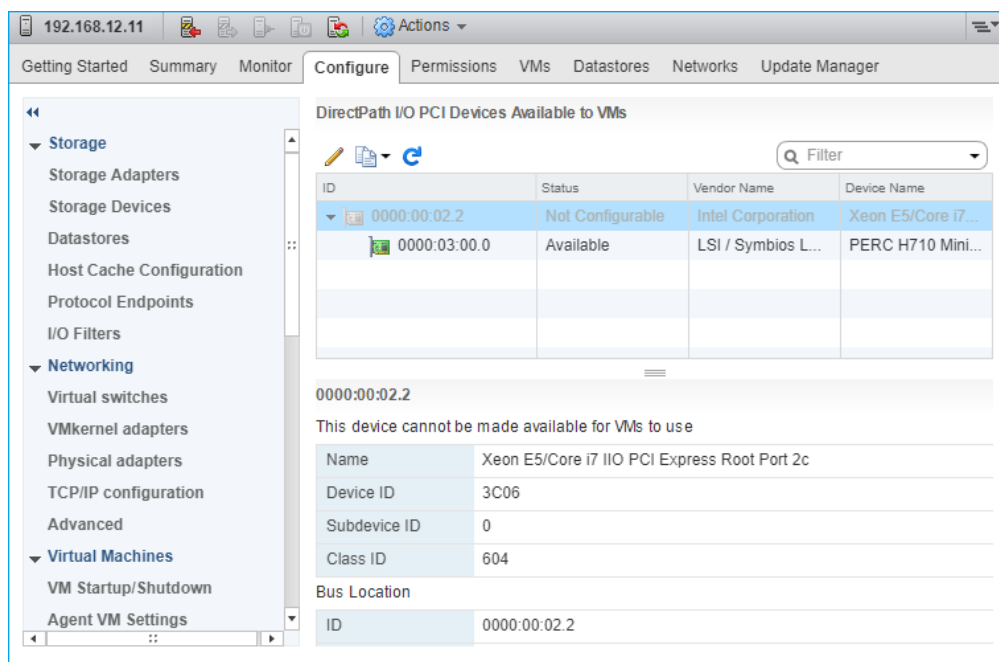
3. Locate the HBA/RAID Controller of the ESXi host. Check the box on the appropriate PCI device. Click OK.



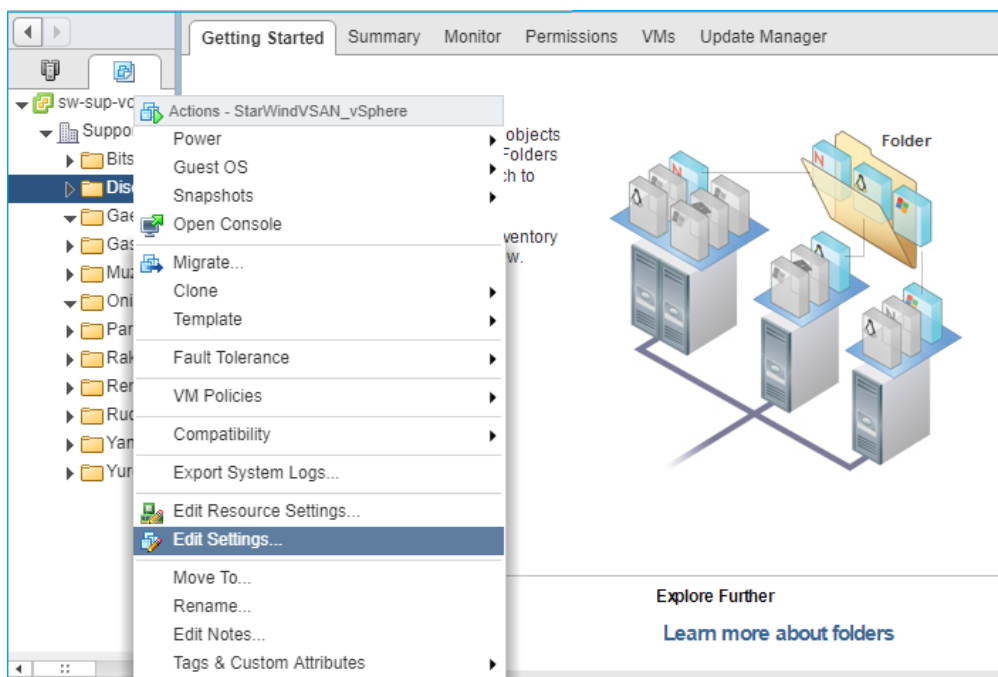
4. The device will appear in the Direct I/O PCI Devices Available to VMs table in the Available (pending) status.




5. Reboot ESXi host. After the reboot, the status changes to Available.



6. Right-click on the StarWind Virtual SAN VM. Select 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.

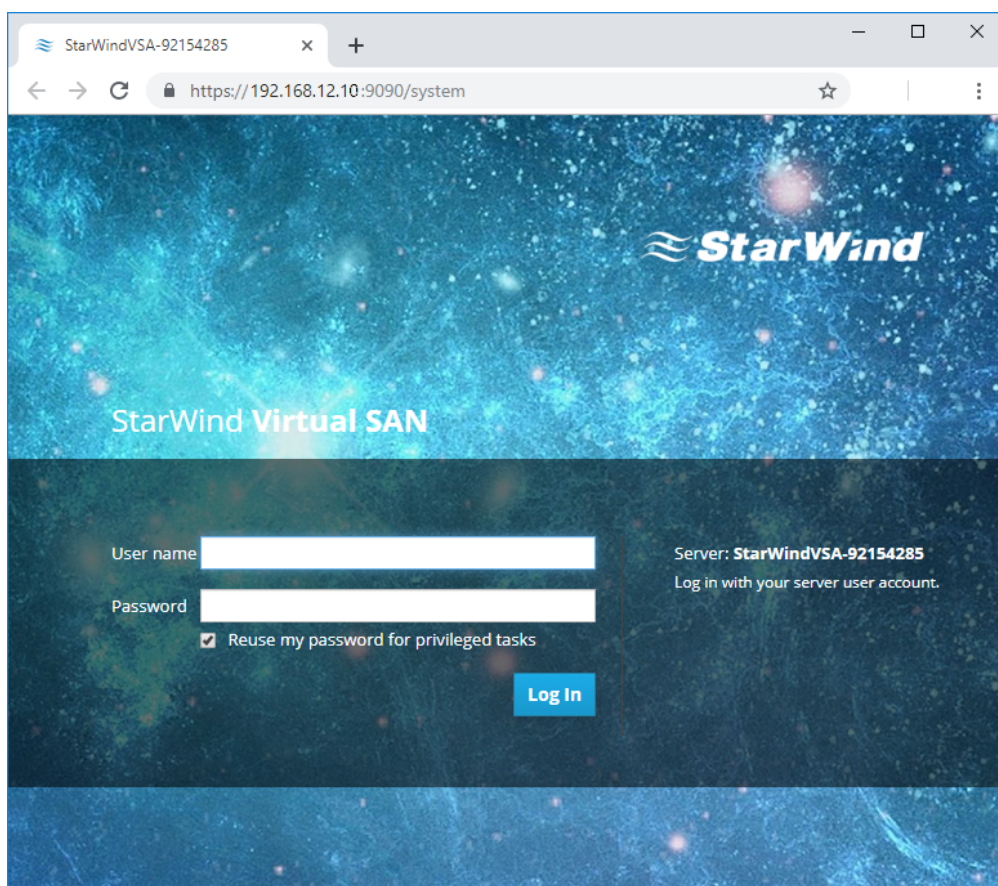
10. Repeat steps 1-8 for all hosts where StarWind Virtual SAN for vSphere is deployed.

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

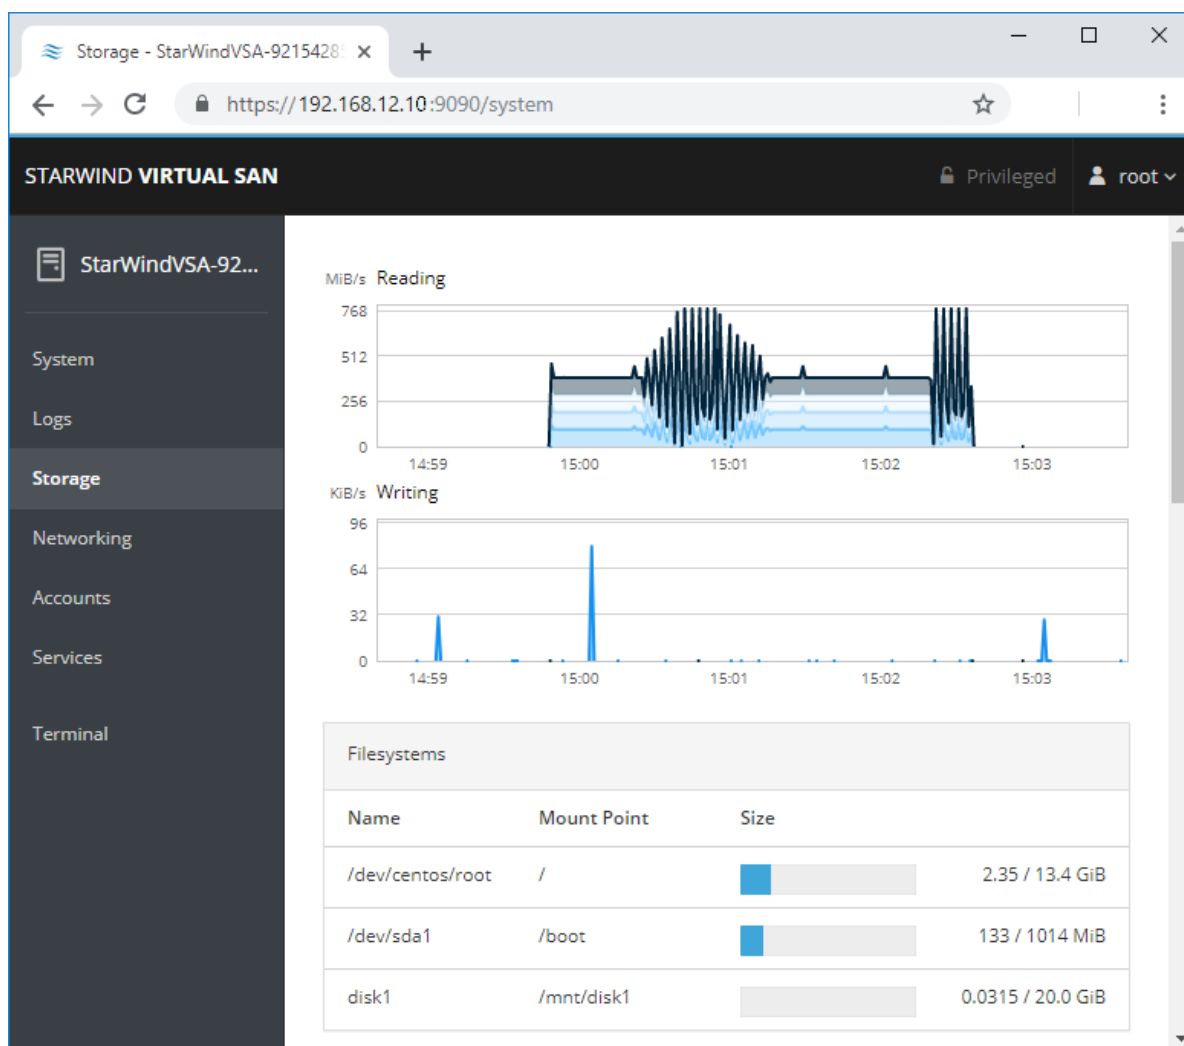
Login: user

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). StarWind recommendations of RAID configurations depending 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.



Configuring Deduplication

To configure VDO please choose a disk, where to place vdo device, logical size (recommended up to x3 of physical storage) and index memory size.

NOTE: It is recommended to use a sparse UDS index for all production use cases. This is an extremely efficient indexing data structure, requiring approximately one-tenth of a byte of DRAM per block in its deduplication window. On disk, it requires approximately 72 bytes of disk space per block. The minimum configuration of this index uses 256 MB of DRAM and approximately 25 GB of space on disk. To use this configuration, specify the `---sparseIndex=enabled` `---indexMem=0.25` options to the vdo create command. This configuration results in a deduplication window of 2.5 TB (meaning it will remember a history of 2.5 TB). For most use cases, a deduplication window of 2.5 TB is appropriate for deduplicating storage pools that are up to 10 TB in size.

1. Open Terminal and paste the following command with your parameters:

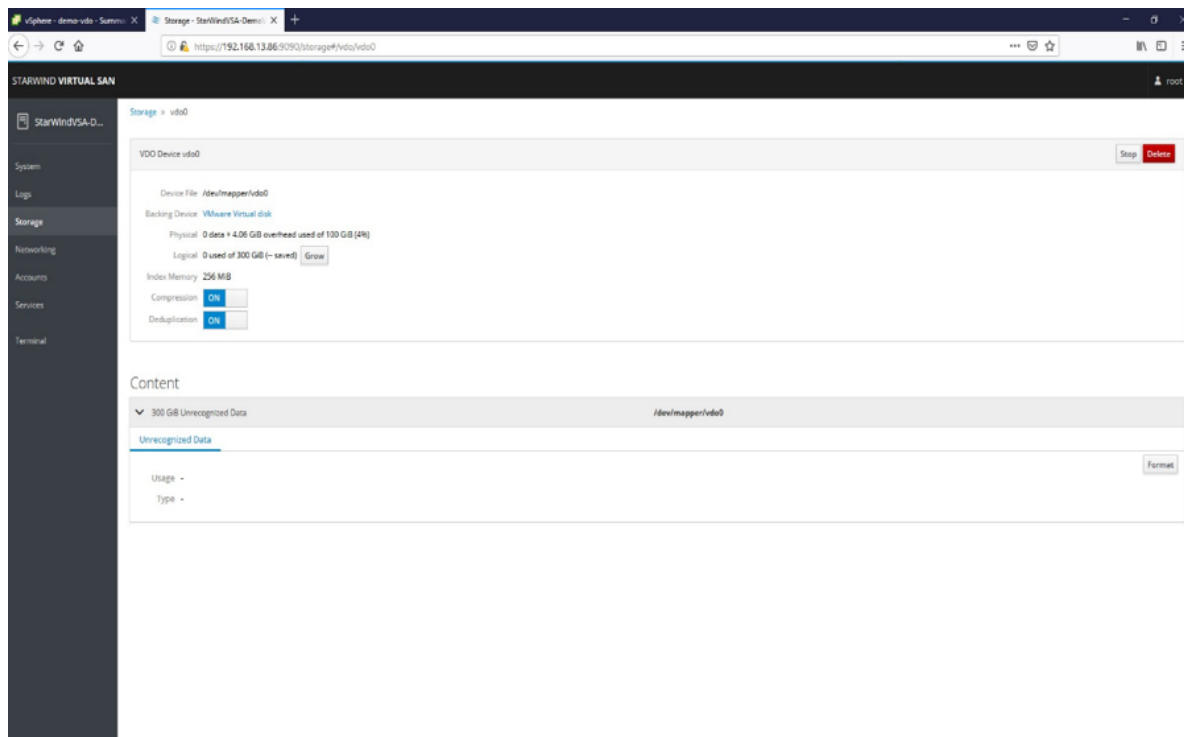
```
vdo create --name=%name% --device=%yourdevice% --
vdoLogicalSize=%size% --verbose --deduplication=enabled --
emulate512=enabled --compression=disabled
vdo modify -n %name% --vdoAckThreads=4 --vdoCpuThreads=4 --
vdoLogicalThreads=4 --vdoPhysicalThreads=4 --
vdoHashZoneThreads=2 --vdoBioThreads=4 --
blockMapCacheSize=2048M
```

```
vdo stop -n %name%; vdo start -n %name%
```

2. Click Create.

Wait until creation is finished. And find new VDO devices in the corresponding section.

Hint: Created VDO device will not have any partitions and filesystem.



3. Click Format

Format /dev/md/RAID10

Erase

Overwrite existing data with zeros ▼

Type

XFS - Red Hat Enterprise Linux 7 default ▼

Name

Disk1

Mounting

Custom ▼

Mount Point

/mnt/disk1

Mount options

☒ Mount at boot

☐ Mount read only

☐ Custom mount options

Formatting a storage device will erase all data on it.

Cancel

Format

4. Create XFS partition. Mount point should be as following: /mnt/%yourdiskname%

Add following custom mounting options to keep mount persistent across reboots: defaults,discard,x-systemd.requires=vdo.service and click Format

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

Content

300 GiB xfs File System /dev/mapper/vdo0

Filesystem

Format

Name: VDOdisk1

Mount Point: /mnt/vdisk1 Mount

Mount Options: defaults

Used: -

Format

Provisioning Storage With Starwind Vsan

Creating StarWind HA devices

1. Open StarWind Management Console and click on the Add Device (advanced) button.
2. Select Hard disk device as the type of device to be created. Click Next to continue.
3. Select Virtual disk. Click Next to continue.
4. Specify 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

NOTE: Image file device and HA device can be extended later according to the steps in KB article:

[How to extend Image File or High Availability device](#)

5. Specify virtual disk options.

The screenshot shows a window titled "Add Device Wizard" with a back arrow and help/close icons. The main section is "Virtual Disk Options". It contains three radio button options: "Thick-provisioned" (selected), "Thick-provisioned with Write Log (experimental)", and "LSFS". Below these is a checkbox for "Deduplication". A "StarPack Cache Size" field is set to "16" with a unit dropdown set to "MB". A "Block Size" section has two radio button options: "Use 512 bytes sector size" (selected) and "Use 4096 bytes sector size. May be incompatible with some clients". At the bottom right are "Next" and "Cancel" buttons.

NOTE: Sector size should be 512 bytes in case of using ESXi.

6. Define the RAM caching policy and specify the cache size (in the corresponding units) and click Next to continue.

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" (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 an unchecked checkbox for "Set Maximum available Size". Underneath that is a "Size:" label followed by a text input field containing "128" and a unit dropdown menu currently set to "MB". At the bottom right, there are "Next" and "Cancel" buttons.

7. Define the Flash caching policy and the cache size. Click Next to continue.

The screenshot shows a window titled "Add Device Wizard" with a back arrow and help/close icons. The main heading is "Specify Flash Cache Parameters". There are two radio button options: "No Flash Cache" (which is selected) and "Use Flash Cache". Below these options is a form with three fields: "Name:" with the value "Flash-DS1", "Location:" with the value "VSA Storage\mnt\ssd\" and a browse button "...", and "Size:" with the value "1" and a unit dropdown menu set to "GB". At the bottom right, there are "Next" and "Cancel" buttons.

8. Specify Target Parameters and select the Allow multiple concurrent iSCSI connections checkbox to enable several clients to connect simultaneously to the target.

The screenshot shows a window titled "Add Device Wizard" with a back arrow and help/close icons. The main heading is "Target Parameters". Below this, there are four fields: a dropdown menu for "Choose a Target Attachment Method" with "Create new Target" selected; a text box for "Target Alias" containing "DS1"; a checkbox for "Target Name" which is unchecked; and a text box for "iqn.2008-08.com.starwindsoftware:sw1-ds1" which is checked. At the bottom right, there are "Next" and "Cancel" buttons.

← Add Device Wizard

Target Parameters

Choose a Target Attachment Method
Create new Target

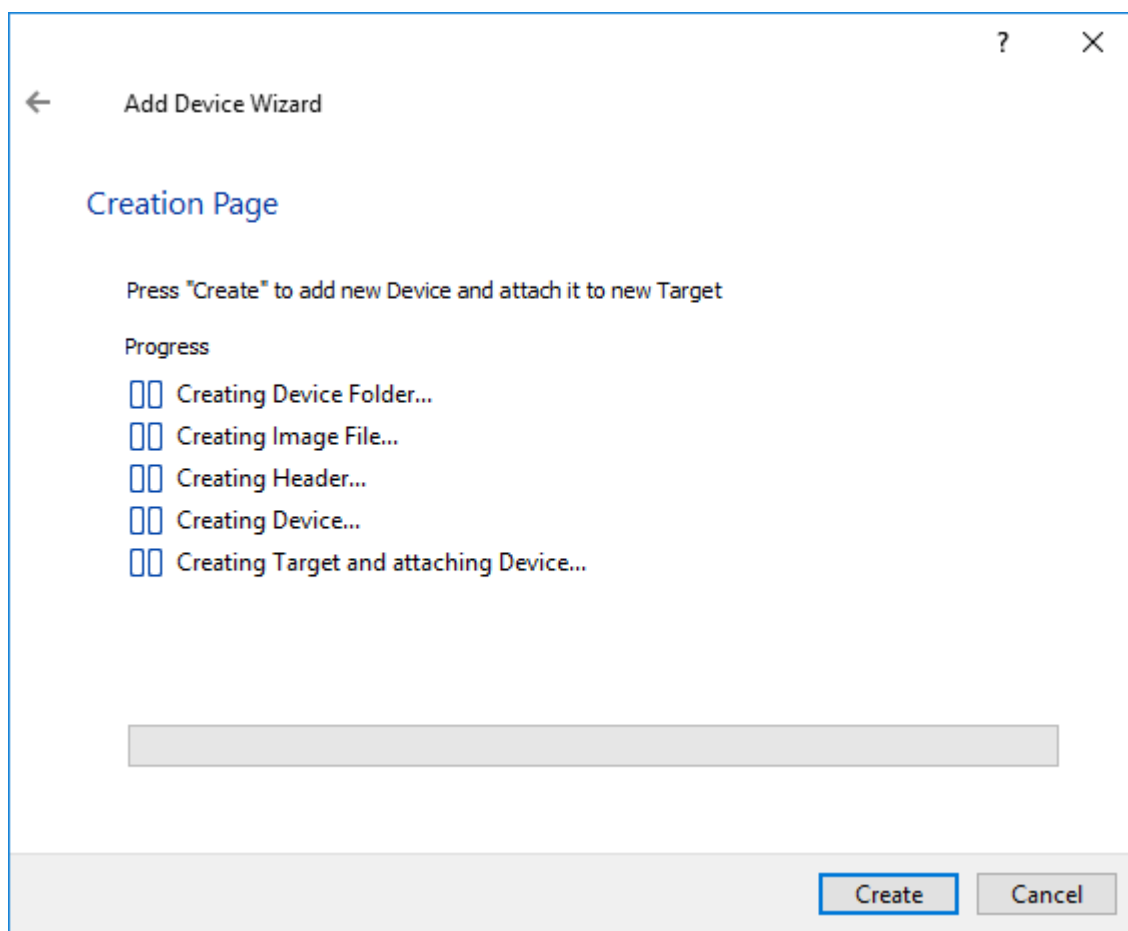
Target Alias
DS1

☐ Target Name
iqn.2008-08.com.starwindsoftware:sw1-ds1

☒ Allow multiple concurrent iSCSI Connections

Next Cancel

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

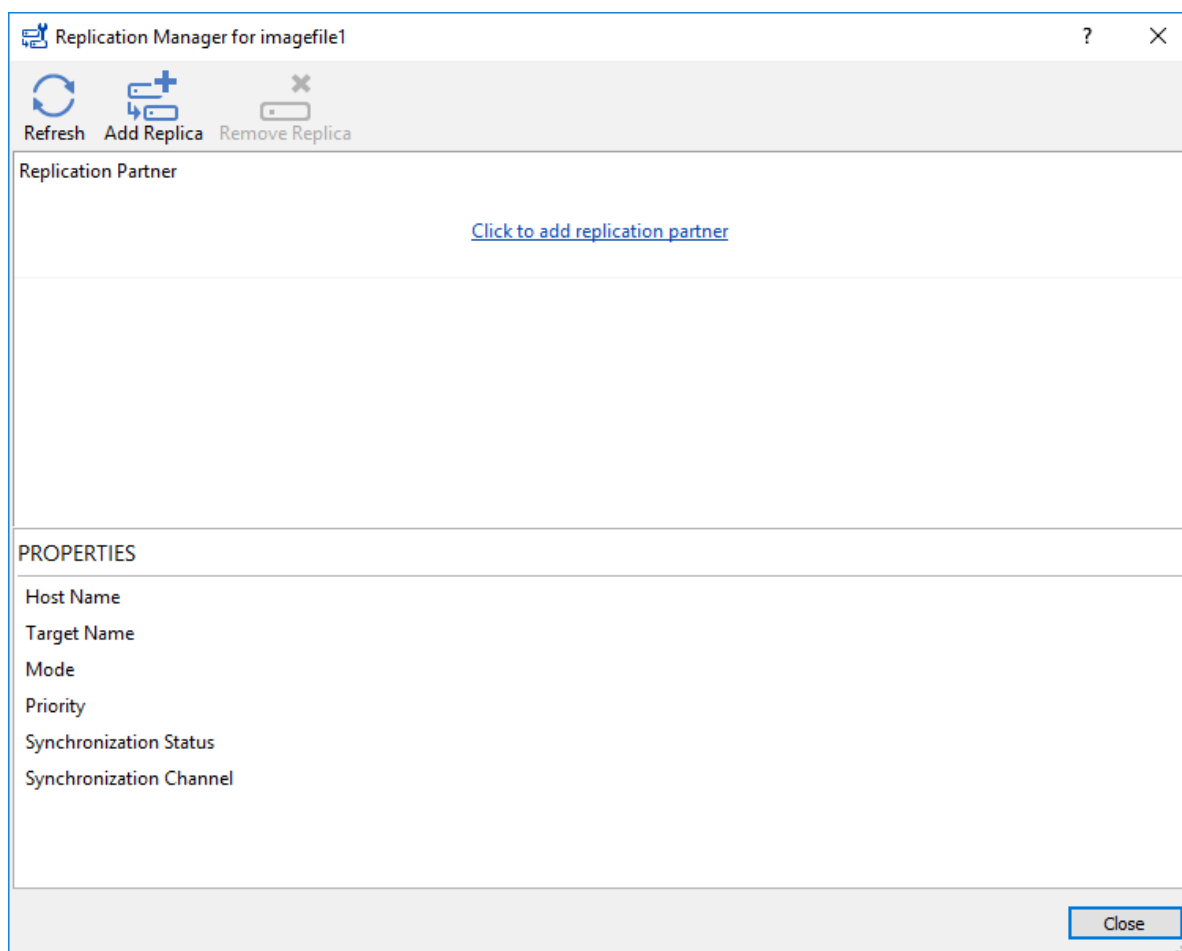


10. Click Create to add a new device and attach it to the target.

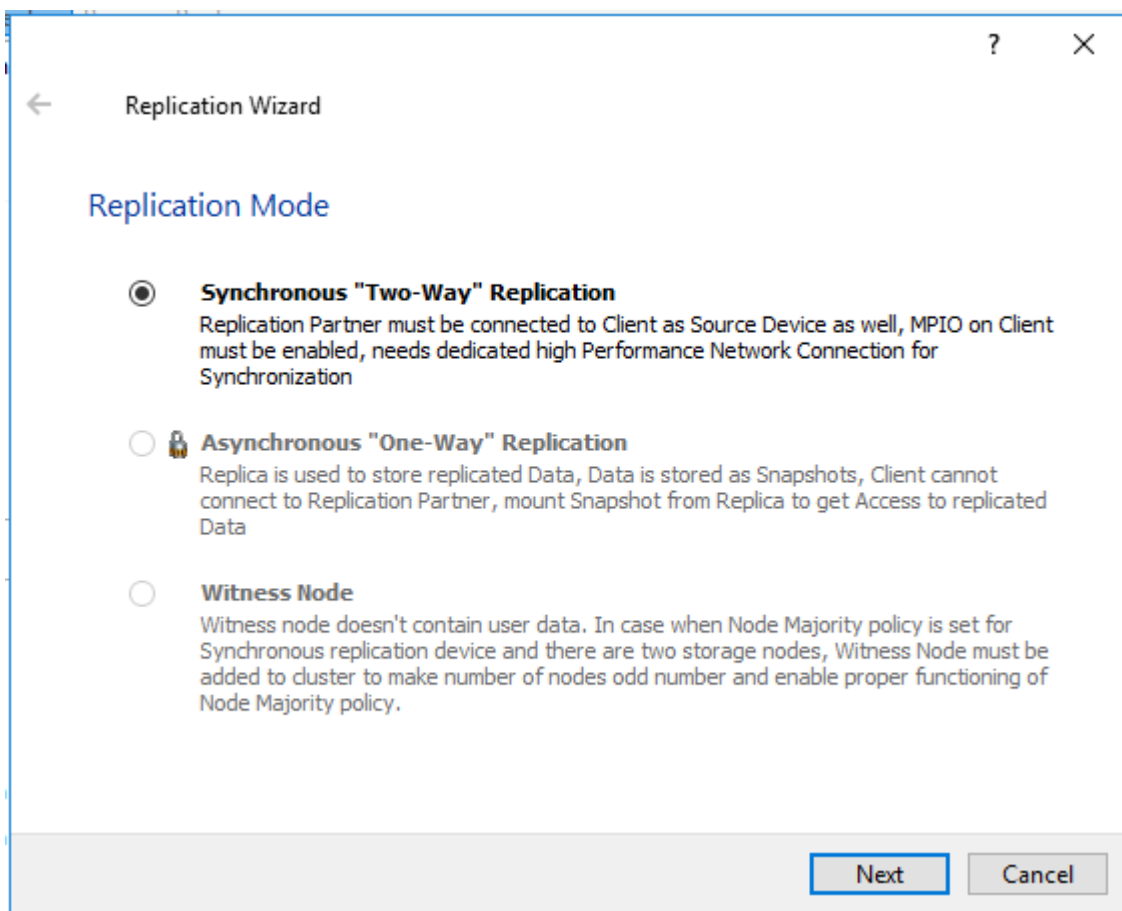
11. Click Finish to close the wizard.

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

13. Click Add replica.



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



15. Specify Partner Hostname or IP Address and Port Number. Click Next.

Replication Wizard

Add Partner Node

Specify Partner Host Name or IP Address where Replication Node would be created

Host Name or IP Address

Port Number

Next Cancel

16. Choose Create new Partner Device and click Next.

Replication Wizard

Partner Device Setup

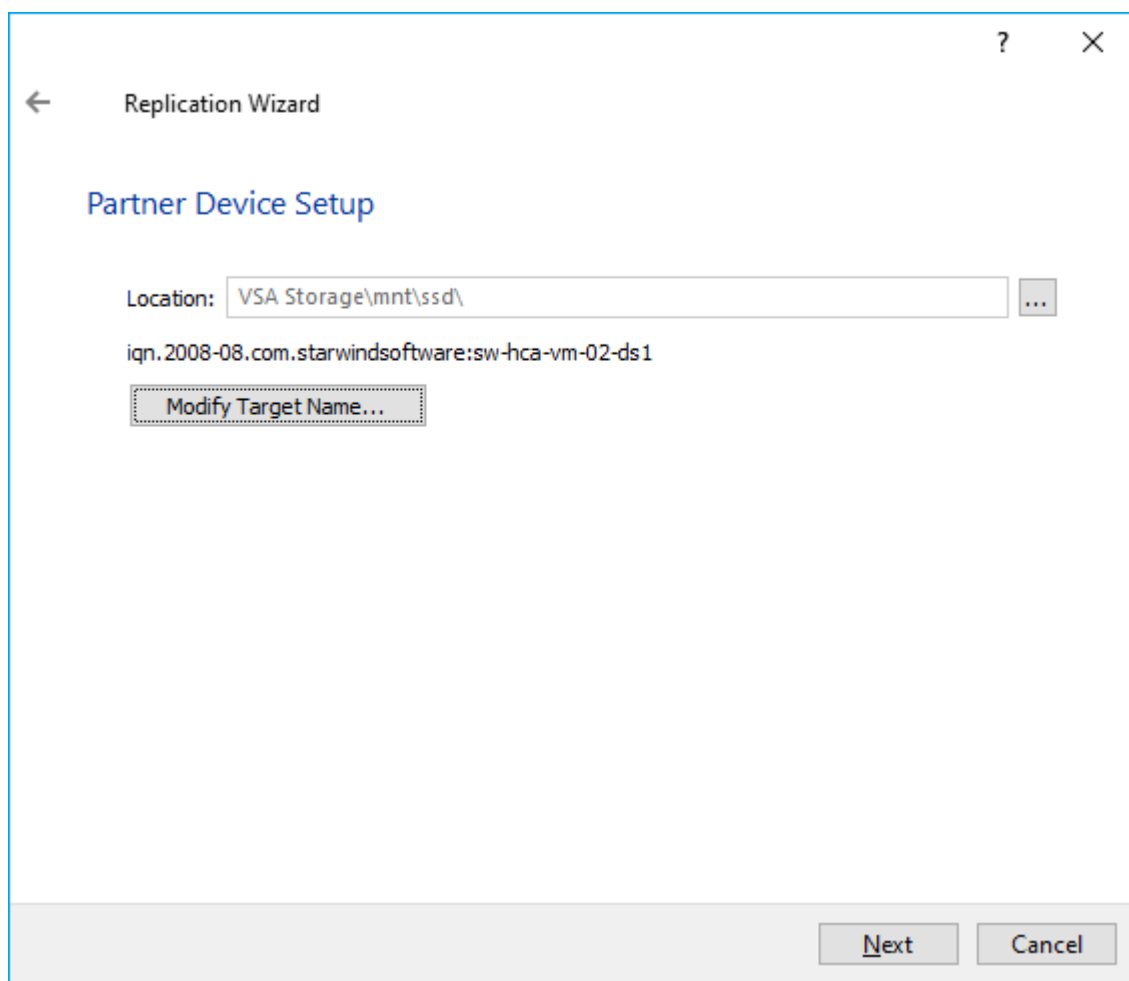
☒ **Create new Partner Device**
Existing Device Parameters would be used as a Template

☐ **Select existing Device**
Select existing Device on Partner Server

Numa Node: Auto

Next Cancel

17. Choose the device Location and specify the target name if required. Otherwise, the name is generated automatically in accordance with the specified target alias.



18.

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.

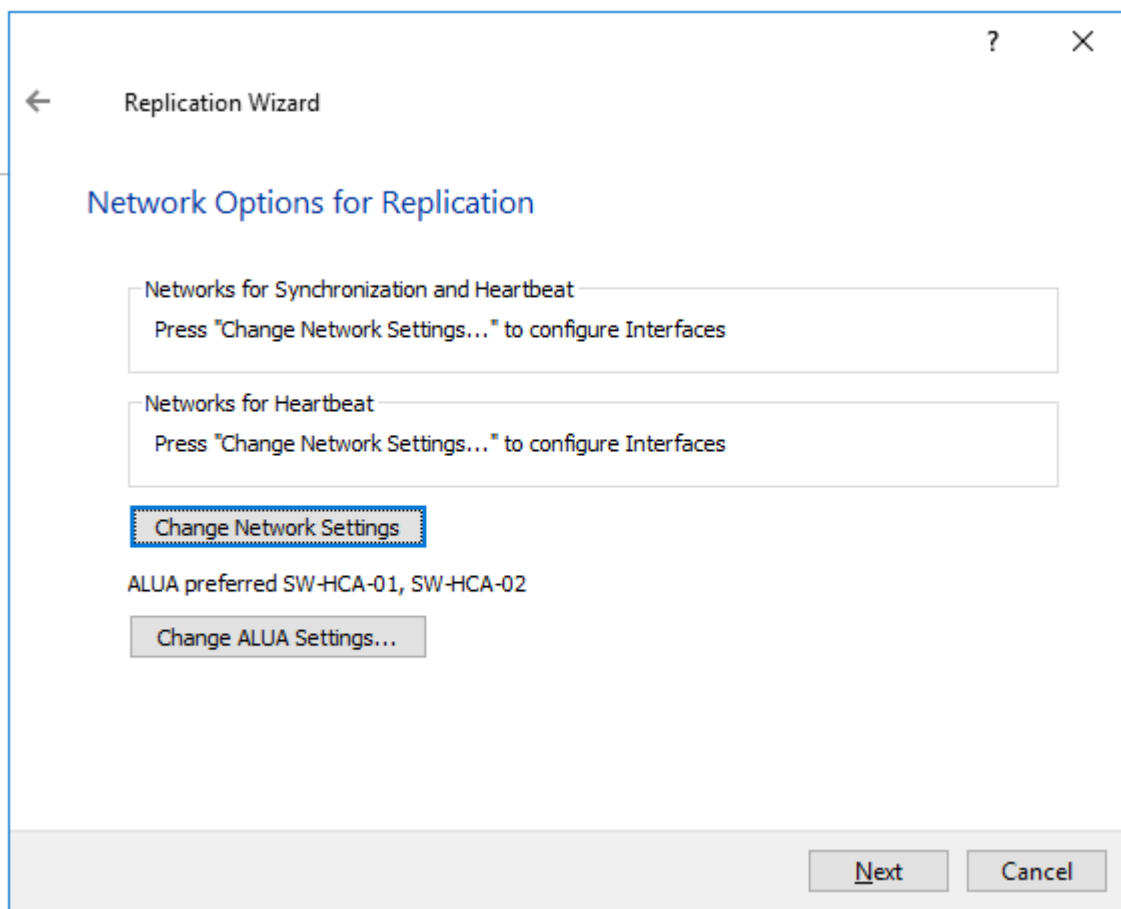
☐ **Continuous journal**
The strategy guarantees fast synchronization and data consistency in all cases.

Current Node:

Partner Node:

Next **Cancel**

19. Click Change Network Settings.



20. Specify the interfaces for Synchronization and Heartbeat channels. Click OK. Then click Next.

Specify Interfaces for Synchronization Channels


Select synchronization channel

Interfaces	Networks	Synchronization and H...	Heartbeat
<div> Host Name: SW-HCA-01 </div>			
172.16.10.1	172.16.10.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.2.111	172.16.2.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.20.1	172.16.20.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
<div> Host Name: SW-HCA-02 </div>			
172.16.10.2	172.16.10.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.2.112	172.16.2.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.20.2	172.16.20.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>

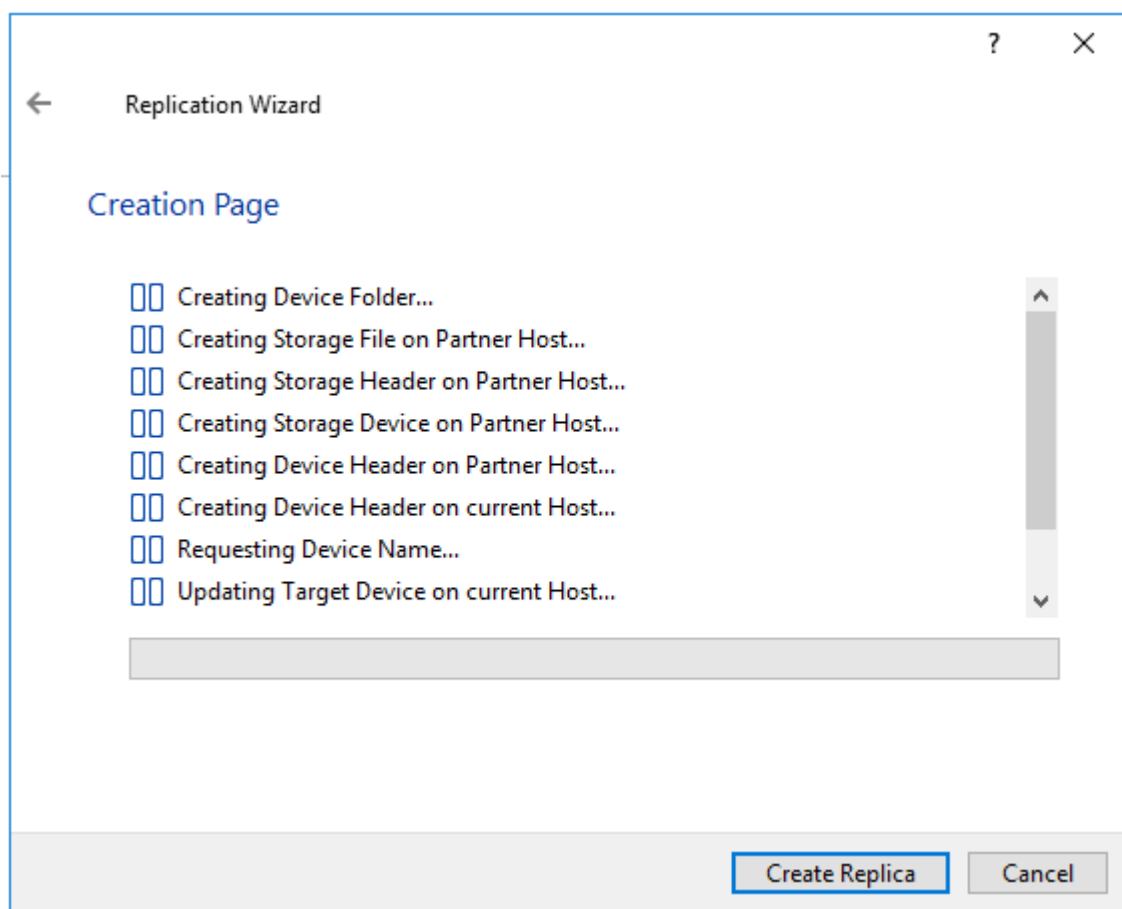
☐ Allow Free Select Interfaces

OK
Cancel

21. Choose Synchronize from existing Device

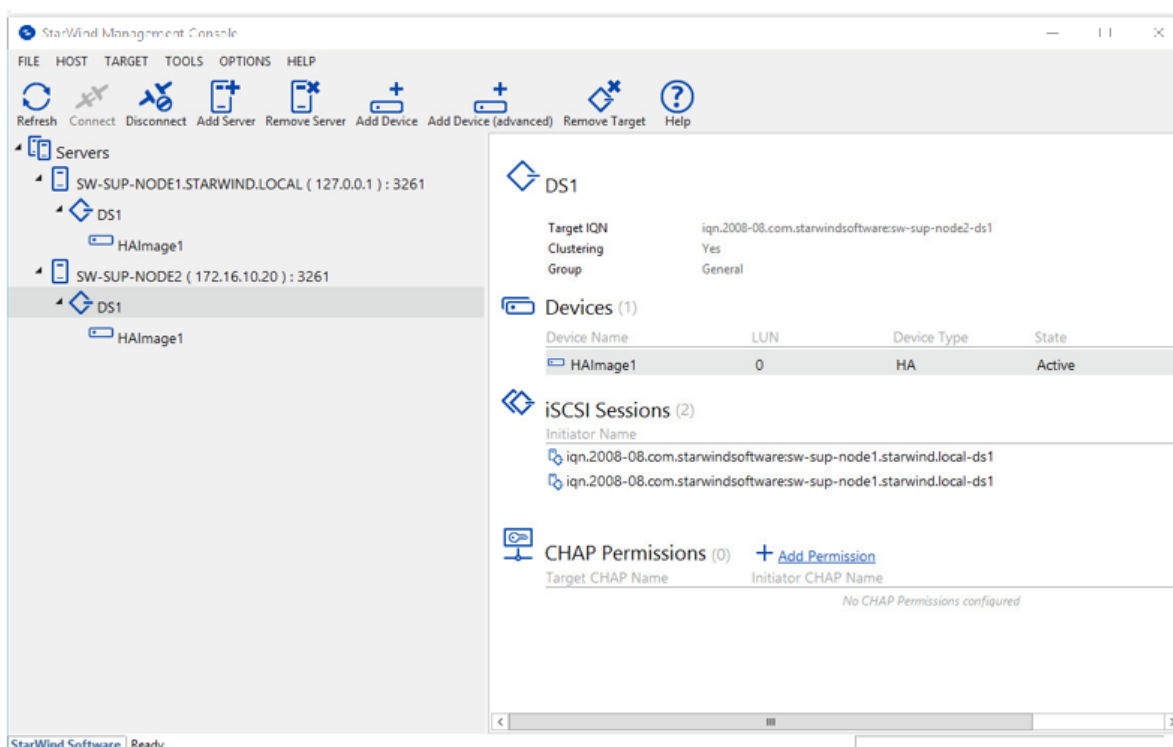


22. Click Create Replica.

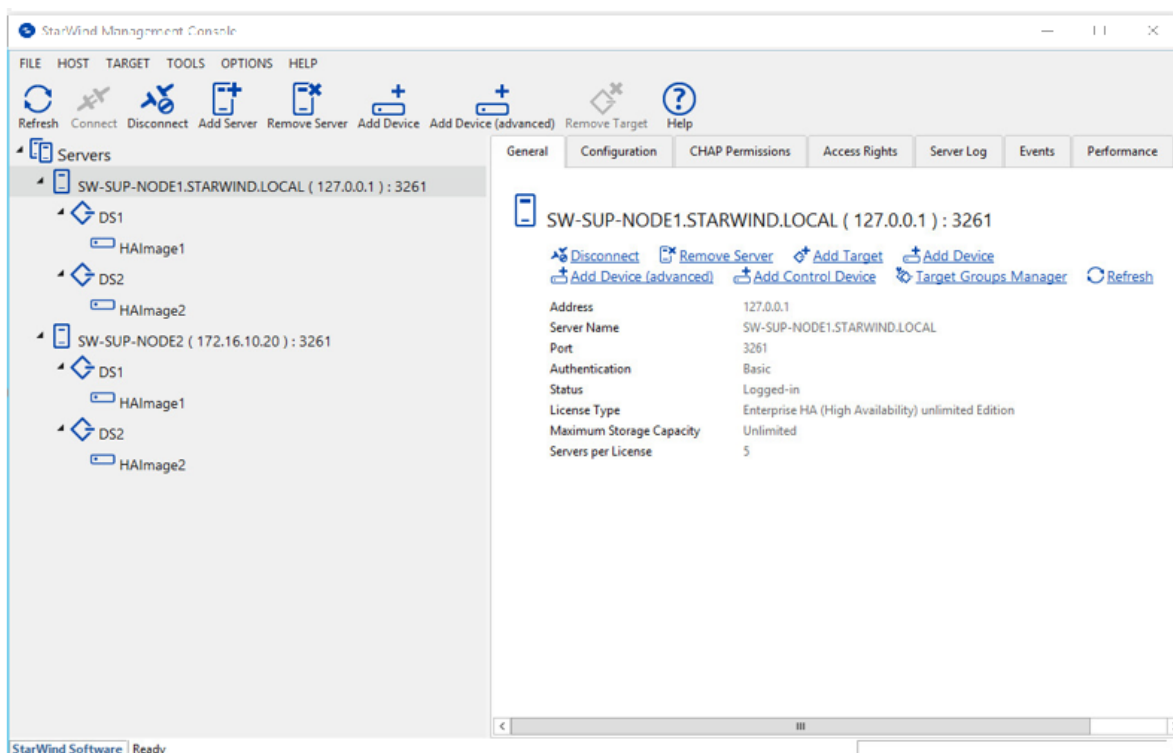


Click Finish to close the wizard.

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




24. Follow the similar procedure to create other virtual disks that will be used as datastores.




Preparing Datastores

Adding Discover Portals



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



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



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



Confirm the actions by pressing Save configuration.

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

Configure iSCSI

iSCSI enabled: ☐ Disabled ☒ Enabled

Name & alias: iqn.1998-01.com.vmware:sw-mar-pc3-6fbab48a

CHAP authentication: Do not use CHAP

Mutual CHAP authentication: Do not use CHAP

Advanced settings: Click to expand

Network port bindings: Add port binding Remove port binding

VMkernel NIC	Port group	IPv4 address
No port bindings		

Static targets: Add static target Remove static target Edit settings Search

Target	Address	Port
iqn.2008-08.com.starwindsoftware:sw1-ds1	172.16.10.10	3260
iqn.2008-08.com.starwindsoftware:sw1-ds2	172.16.10.10	3260
iqn.2008-08.com.starwindsoftware:sw2-ds1	172.16.10.20	3260
iqn.2008-08.com.starwindsoftware:sw2-ds2	172.16.10.20	3260

Dynamic targets: Add dynamic target Remove dynamic target Edit settings Search

Address	Port
172.16.10.10	3260
172.16.10.20	3260

Save configuration Cancel

5. Click on the Rescan button to rescan storage.

Devices

New datastore Increase capacity Rescan Refresh Actions


Name	Status	Type	Capacity
STARWIND iSCSI Disk (eui.f8289e52a311c08d)	Normal	Disk	3 GB
Local NECVMWare CD-ROM (mpx.vmhba64.C0:T0:L0)	Normal	CDROM	Unknown
STARWIND iSCSI Disk (eui.ccd882632aff4068)	Normal	Disk	3 GB
Local VMWare Disk (mpx.vmhba0.C0:T0:L0)	Normal	Disk	40 GB

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

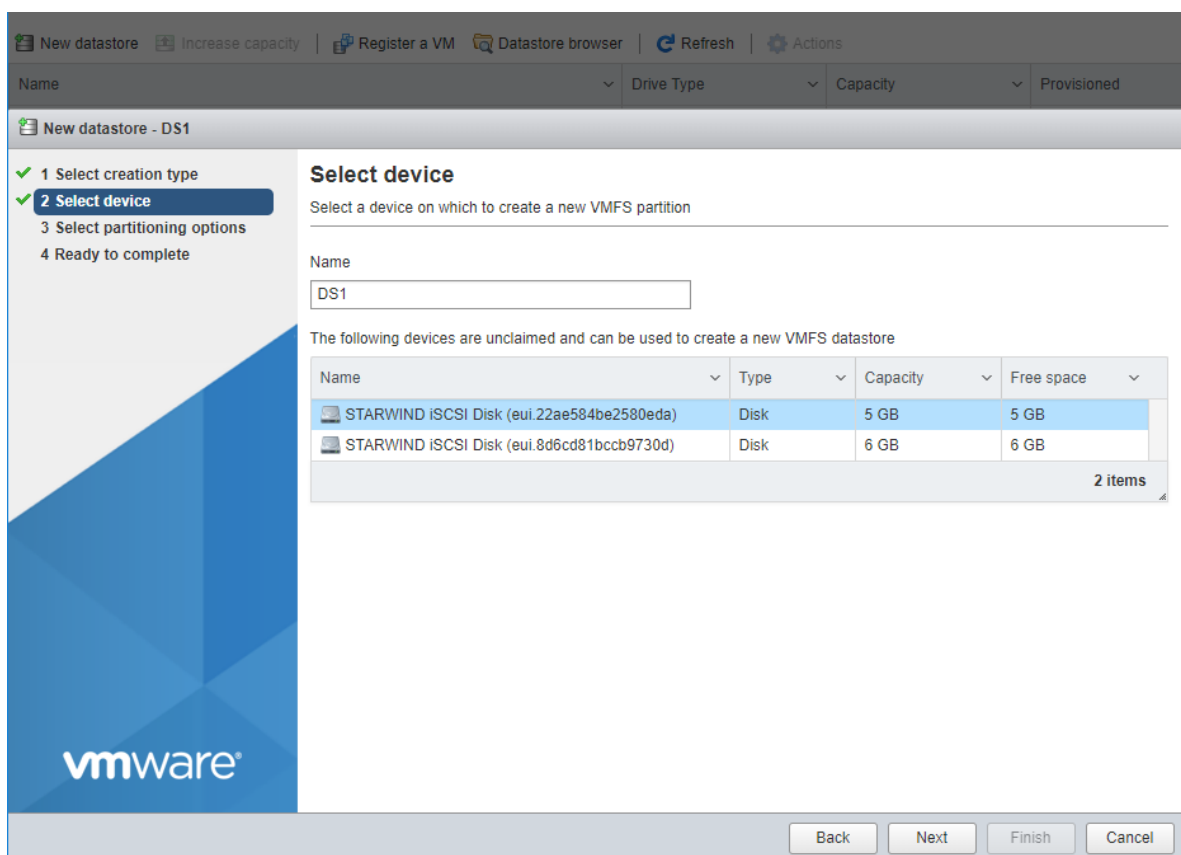
7. Repeat all the steps from this section on the other ESXi host, specifying corresponding IP addresses for the iSCSI subnet.

Creating Datastores

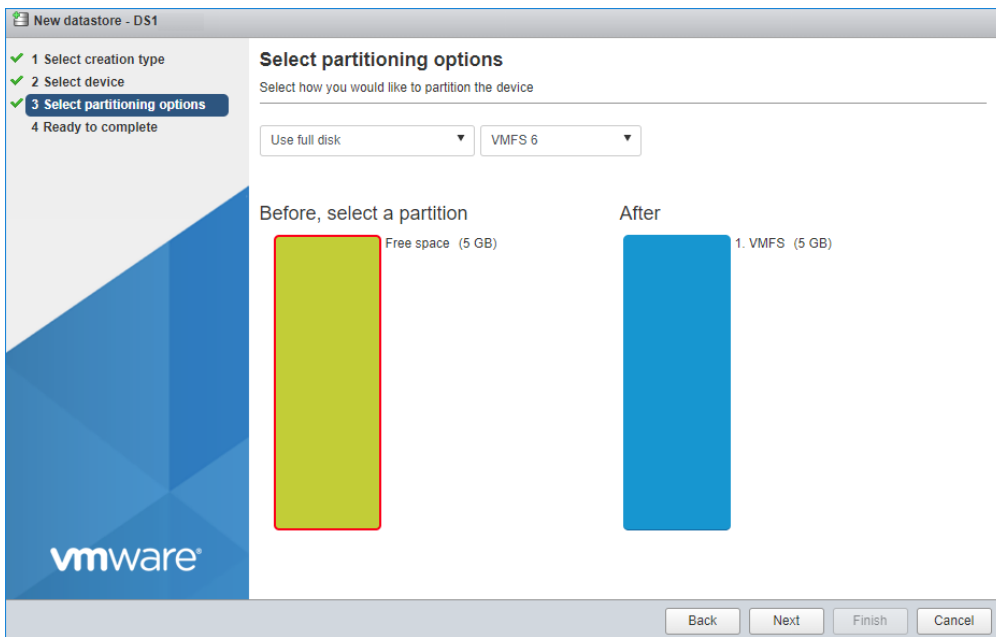
1. Open the Storage tab on one of your hosts and click on New Datastore.



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				
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`

```

GNU nano 2.3.1 File: /opt/StarWind/StarWindVSA/drive_c/StarWind/hba_rescan.ps1

#!/usr/bin/pwsh-preview
$env:PSModulePath="/root/.local/share/powershell/Modules:/usr/local/share/powershell/Modules:/opt/microsoft/powershell/6-preview/Modules"
$env:HOME="/root"
Set-PowerCLIConfiguration -InvalidCertificateAction ignore -Confirm:$false | Out-Null
$ESXiHost = "SET_ESXi_HOST_IP"
$ESXiUser = "SET_ESXi_USER"
$ESXiPassword = "SET_PASSWORD"
Connect-VIServer $ESXiHost -User $ESXiUser -Password $ESXiPassword | Out-Null
Get-VMHostStorage $ESXiHost -RescanAllHba | Out-Null
Get-ScsiLun -VMHost $ESXiHost -LunType disk | Where-Object Vendor -EQ "STARWIND" |
Where-Object ConsoleDeviceName -NE " " | Set-ScsiLun -MultipathPolicy RoundRobin -CommandsToSwitchPath 1 |
Out-Null
Disconnect-VIServer $ESXiHost -Confirm:$false
  
```

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:

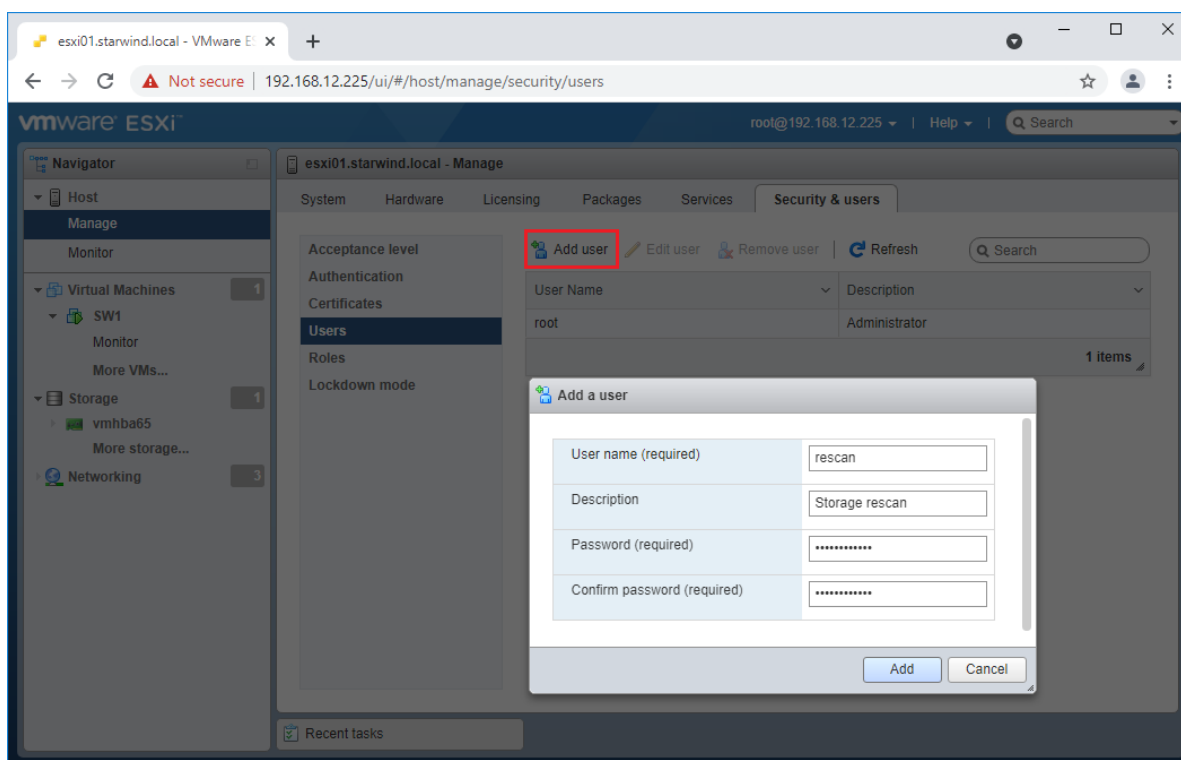
```

$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.

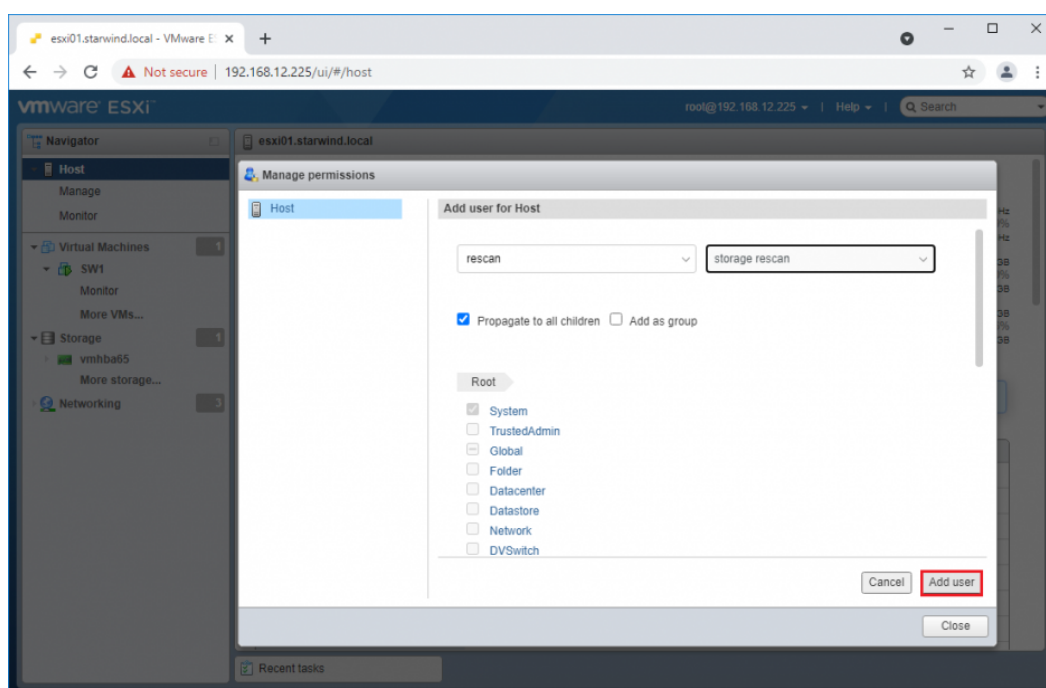
The following privileges might be assigned: Host – Inventory, Config, Local Cim, and Global – Settings.



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.

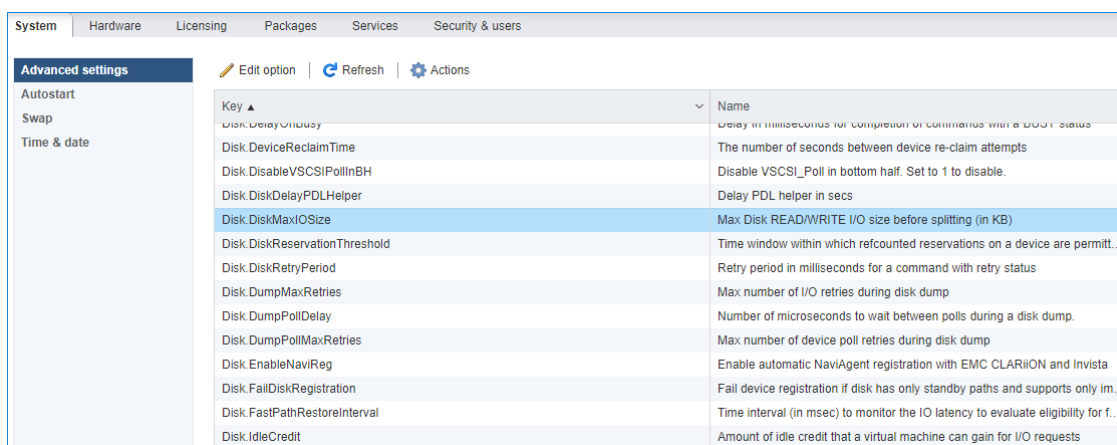


Make sure that rescan script is working and execute it from the VM: `sudo /opt/StarWind/StarWindVSA/drive_c/StarWind/hba_rescan.ps1`

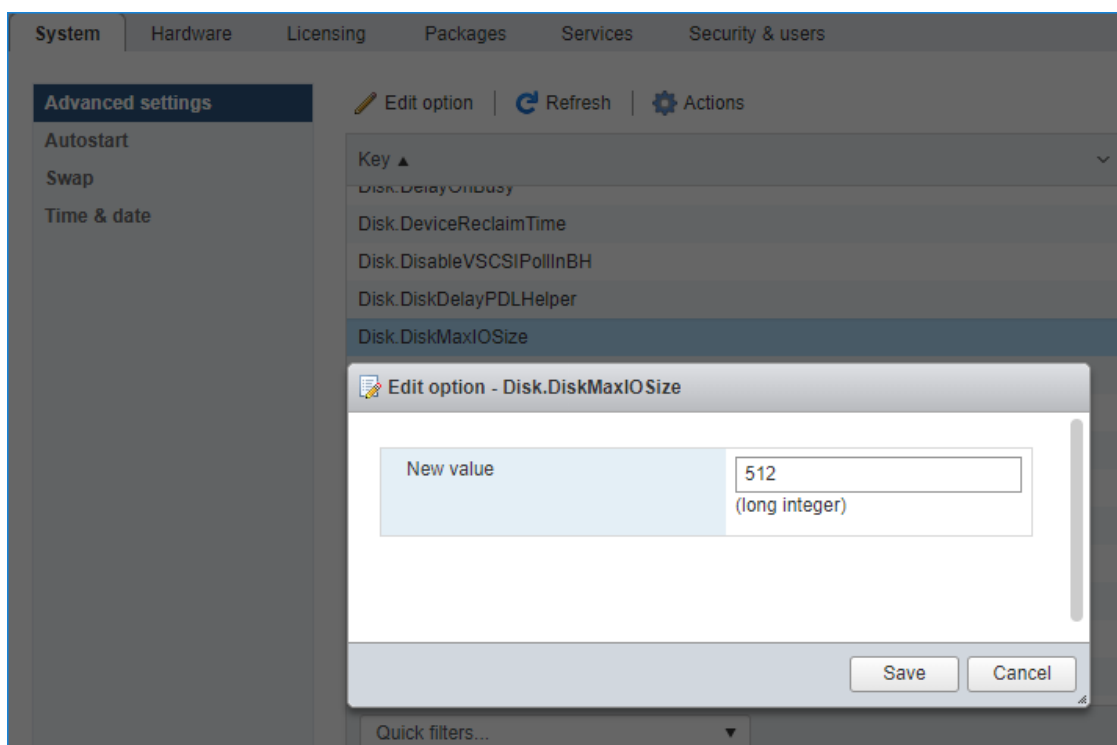
4. Repeat all steps from this section on the other ESXi hosts.

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.



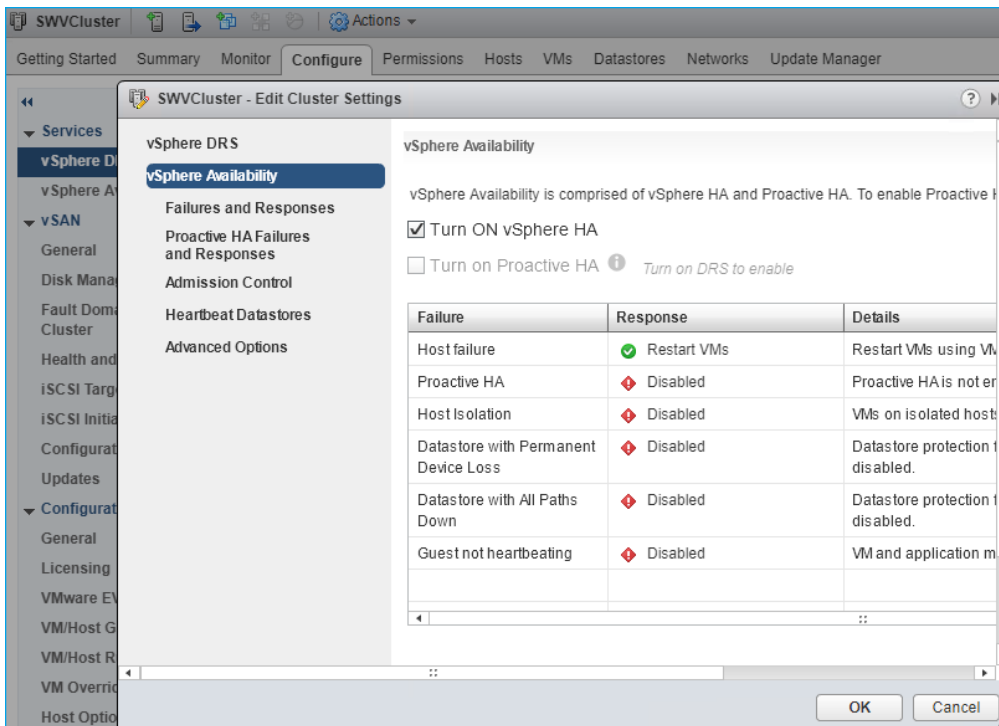
3. To optimize performance change I/O scheduler options according to the article below:
<https://knowledgebase.starwindsoftware.com/guidance/starwind-vsan-for-vsphere-changing-linux-i-o-scheduler-to-optimize-storage-performance/>

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

This guide provides comprehensive instructions for setting up StarWind Virtual SAN with deduplication and compression in hyperconverged scenario. It offers a detailed path for IT professionals to enhance their VMware storage infrastructure, leading to improved storage optimization and performance.

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.