

StarWind Virtual SAN: Configuration Guide for Proxmox Virtual Environment (VE) [KVM], VSAN Deployed as a Controller 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

This guide applies to StarWind Virtual SAN, StarWind Virtual SAN Free (starting from version 1.2xxx – Oct. 2023).

Purpose

This document outlines how to configure a Proxmox Cluster using StarWind Virtual SAN (VSAN), with VSAN running as a Controller Virtual Machine (CVM). The guide includes steps to prepare Proxmox hosts for clustering, configure physical and virtual networking, and set up the Virtual SAN Controller Virtual Machine.

For more information about StarWind VSAN architecture and available installation options, please refer to the [StarWind Virtual \(vSAN\) Getting Started Guide](#).

Audience

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

Expected result

The end result of following this guide will be a fully configured high-availability Proxmox Cluster that includes virtual machine shared storage provided by StarWind VSAN.

Prerequisites

StarWind Virtual SAN system requirements

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

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

Recommended RAID settings for HDD and SSD disks:

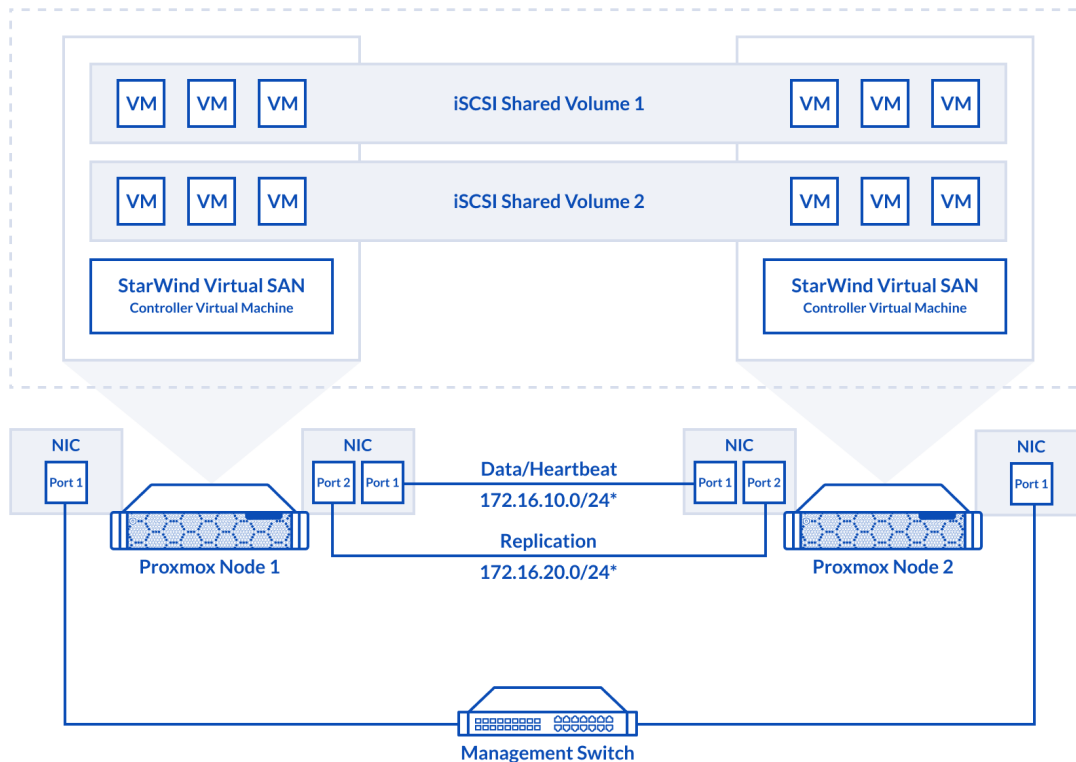
<https://knowledgebase.starwindsoftware.com/guidance/recommended-raid-settings-for-hdd-and-ssd-disks/>

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

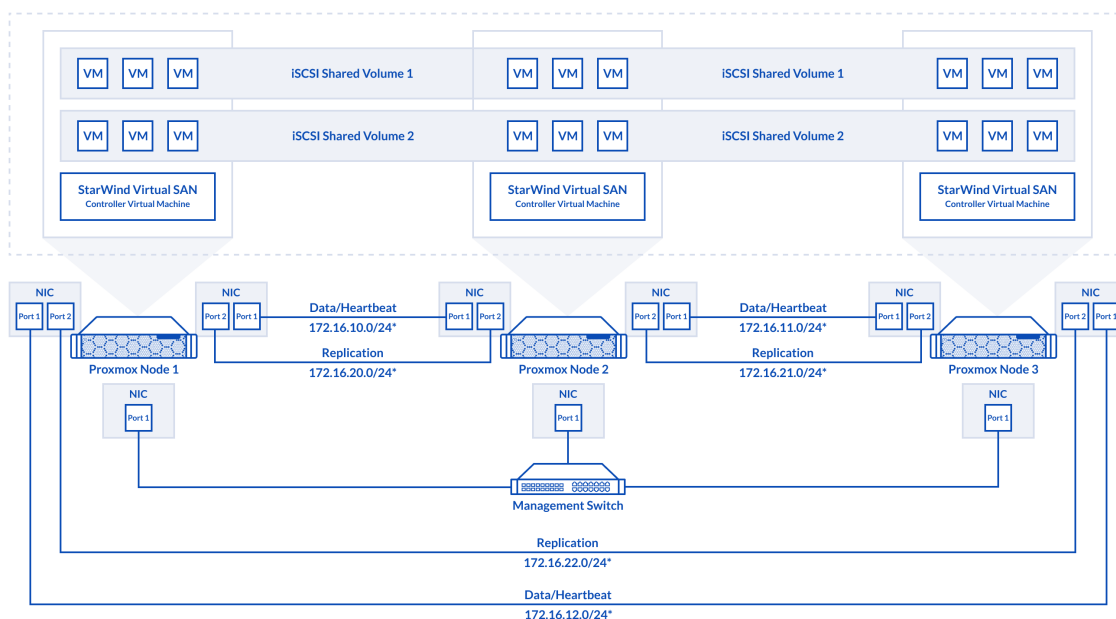
<https://www.starwindsoftware.com/resource-library/starwind-virtual-san-best-practices>

Solution diagram

The diagrams below illustrate the network and storage configuration of the solution:



2-node cluster



3-node cluster

Preconfiguring cluster nodes

1. ProxMox cluster should be created before deploying any virtual machines.
2. 2-nodes cluster requires quorum. iSCSI/SMB/NFS cannot be used for this purposes. QDevice-Net package must be installed on 3rd Linux server, which will act as a witness.
https://pve.proxmox.com/wiki/Cluster_Manager#_corosync_external_vote_support
3. Install qdevice on witness server:

```
ubuntu# apt install corosync-qnetd
```


4. Install qdevice on both cluster nodes:

```
pve# apt install corosync-qdevice
```


5. Configure quorum running the following command on one of the ProxMox node (change IP address)

```
pve# pvecm qdevice setup %IP_Address_Of_Qdevice%
```

6. Configure network interfaces on each node to make sure that Synchronization and iSCSI/StarWind heartbeat interfaces are in different subnets and connected according to the network diagram above. In this document, 172.16.10.x subnet is used for iSCSI/StarWind heartbeat traffic, while 172.16.20.x subnet is used for the Synchronization traffic. Choose node and open System -> Network page.



7. Click Create. Choose Linux Bridge.



8. Create

Linux Bridge and set IP address. Set MTU to 9000. Click Create.

Create: Linux Bridge

Name:
vmbr1

IPv4/CIDR:
172.16.10.1/24

Gateway (IPv4):

IPv6/CIDR:

Gateway (IPv6):

Autostart:
☒

VLAN aware:
☐

Bridge ports:
ens1f0np0

Comment:

MTU:
9000

Help

Advanced ☒

Create

9. Repeat step 8 for all network adapters, which will be used for Synchronization and iSCSI/StarWind heartbeat traffic.

10. Verify network configuration in /etc/network/interfaces file. Login to the node via SSH and check the contents of the file.

```

auto lo
iface lo inet loopback

iface enp1s0 inet manual

iface enp7s0 inet manual
    mtu 9000
iface enp8s0 inet manual
    mtu 9000
auto vmbr0
iface vmbr0 inet static
    address 172.16.2.37/24
    gateway 172.16.2.1
    bridge-ports enp1s0
    bridge-stp off
    bridge-fd 0

auto vmbr1
iface vmbr1 inet static
    address 172.16.10.1/24
    bridge-ports enp7s0
    bridge-stp off
    bridge-fd 0
    mtu 9000

auto vmbr2
iface vmbr2 inet static
    address 172.16.20.1/24
    bridge-ports enp8s0
    bridge-stp off
    bridge-fd 0
    mtu 9000
-- INSERT --
  
```

11. Enable IOMMU support in kernel, if PCIe passthrough will be used to pass RAID Controller, HBA or NVMe drives to the VM. Update grub configuration file.

For Intel CPU:

Add "intel_iommu=on iommu=pt" to GRUB_CMDLINE_LINUX_DEFAULT line in /etc/default/grub file.

For AMD CPU:

Add "iommu=pt" to GRUB_CMDLINE_LINUX_DEFAULT line in /etc/default/grub file.

```

root@sw-demo-node-01:~# cat /etc/default/grub
# If you change this file, run 'update-grub' afterwards to update
# /boot/grub/grub.cfg.
# For full documentation of the options in this file, see:
#   info -f grub -n 'Simple configuration'

GRUB_DEFAULT=0
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR=`lsb_release -i -s 2> /dev/null || echo Debian`
GRUB_CMDLINE_LINUX_DEFAULT="quiet intel_iommu=on iommu=pt"
GRUB_CMDLINE_LINUX=""

# If your computer has multiple operating systems installed, then you
# probably want to run os-prober. However, if your computer is a host
# for guest OSes installed via LVM or raw disk devices, running
# os-prober can cause damage to those guest OSes as it mounts
# filesystems to look for things.
#GRUB_DISABLE_OS_PROBER=false

# Uncomment to enable BadRAM filtering, modify to suit your needs
# This works with Linux (no patch required) and with any kernel that obtains
# the memory map information from GRUB (GNU Mach, kernel of FreeBSD ...)
#GRUB_BADRAM="0x01234567,0xfefefefe,0x89abcdef,0xefefefef"

# Uncomment to disable graphical terminal
#GRUB_TERMINAL=console

# The resolution used on graphical terminal
# note that you can use only modes which your graphic card supports via VBE
# you can see them in real GRUB with the command `vbeinfo'
#GRUB_GFXMODE=640x480

# Uncomment if you don't want GRUB to pass "root=UUID=xxx" parameter to Linux
#GRUB_DISABLE_LINUX_UUID=true

# Uncomment to disable generation of recovery mode menu entries
#GRUB_DISABLE_RECOVERY="true"

# Uncomment to get a beep at grub start
#GRUB_INIT_TUNE="480 440 1"
root@sw-demo-node-01:~# █

```

12. Reboot the host.

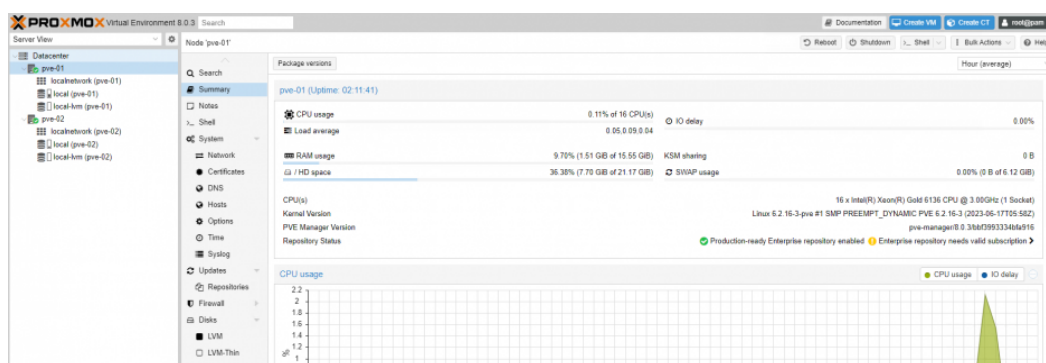
13. Repeat steps 6-12 an all nodes.

Deploying Starwind Virtual San Cvm

1. Download StarWind VSAN CVM KVM: [VSAN by StarWind: Overview](#)
2. Extract the VM StarWindAppliance.qcow2 file from the downloaded archive.
3. Upload StarWindAppliance.qcow2 file to the Proxmox Host via any SFTP client (e.g. WinSCP) to /root/ directory.

/root/					
Name	Size	Changed	Rights	Owner	
..		2/26/2024 8:35:01 AM	rw-r--r--	root	
CVM-0.qcow2	10,614,3...	3/18/2024 6:13:02 AM	rw-r--r--	root	

4. Create a VM without OS. Login to Proxmox host via Web GUI. Click Create VM.



5. Choose node to create VM. Enable Start at boot checkbox and set Start/Shutdown order to 1. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

Node:
sw-demo-node-01
Resource Pool:
VM ID:
100
Name:
CVM01

Start at boot:
☒
Start/Shutdown order:
1
Startup delay:
default
Shutdown timeout:
default

Tags
No Tags
+

Help
Advanced
☒
Back
Next

6. Choose Do not use any media and choose Guest OS Linux. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

☐ Use CD/DVD disc image file (iso)
Guest OS:
Type:
Linux
Version:
6.x - 2.6 Kernel
Storage:
local
ISO image:
☐ Use physical CD/DVD Drive
☒ Do not use any media

Advanced
☐
Back
Next

6. Specify system options. Choose Machine type q35 and check the Qemu Agent box. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

Graphic card:
Default
SCSI Controller:
VirtIO SCSI single
Machine:
q35
Qemu Agent:
☒
Firmware
BIOS:
Default (SeaBIOS)
Add TPM:
☐

Help
Advanced ☒
Back
Next

7. Remove all disks from the VM. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

No Disks

Add

Advanced ☐
Back
Next

8. Assign 8 cores to the VM and choose Host CPU type. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

Sockets: 1
Type: host
Cores: 8
Total cores: 8

Help
Advanced
Back
Next

9. Assign at least 8GB of RAM to the VM. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

Memory (MiB): 8192

Help
Advanced
Back
Next

10. Configure Management network for the VM. Click Next.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

☐ No network device

Bridge:
vmbr0
Model:
VirtIO (paravirtualized)
VLAN Tag:
no VLAN
MAC address:
auto
Firewall:
☒

? Help
Advanced ☐
Back
Next

11. Confirm settings. Click Finish.

Create: Virtual Machine

General
OS
System
Disks
CPU
Memory
Network
Confirm

Key ↑	Value
agent	1
cores	8
cpu	host
ide2	none,media=cdrom
machine	q35
memory	8192
name	CVMD1
net0	virtio,bridge=vmbr0,firewall=1
nodename	sw-demo-node-01
numa	0
onboot	1
ostype	l26
scsihw	virtio-scsi-single
sockets	1

☐ Start after created

Advanced ☐
Back
Finish

12. Connect to Proxmox host via SSH. Attach StarWindAppliance.qcow2 file to the VM.

```
qm importdisk 100 /root/StarWindAppliance.qcow2 local-lvm
```

13. Open VM and go to Hardware page. Add unused SCSI disk to the VM.

14. Attach Network interfaces for Synchronization and iSCSI/Heartbeat traffic.

Add: Network Device

Bridge:
vmbr2

Model:
VirtIO (paravirtualized)

VLAN Tag:
no VLAN

MAC address:
auto

Firewall:
☒

Help

Advanced
☐

Add

15. Open Options page of the VM. Select Boot Order and click Edit.

Server View

Datacenter (sw-demo-cluster)

sw-demo-node-01

100 (CVM01)

localnetwork (sw-demo-node-01)
local (sw-demo-node-01)
local-lvm (sw-demo-node-01)
sw-demo-node-02
localnetwork (sw-demo-node-02)
local (sw-demo-node-02)
local-lvm (sw-demo-node-02)

Virtual Machine 100 (CVM01) on node 'sw-demo-node-01'
No Tags

Summary
Edit
Revert

Console
Hardware
Cloud-Init
Options
Task History
Monitor
Backup
Replication
Snapshots
Firewall
Permissions

Name
CVM01

Start at boot
Yes

Start/Shutdown order
order=1

OS Type
Linux 6.x - 2.6 Kernel

Boot Order
scsi0

Use tablet for pointer
Yes

Hotplug
Disk, Network, USB

ACPI support
Yes

KVM hardware virtualization
Yes

Freeze CPU at startup
No

Use local time for RTC
Default (Enabled for Windows)

RTC start date
now

SMBIOS settings (type1)
uuid=278343be-974a-46a5-ae99-e03bfc966d20

QEMU Guest Agent
Enabled

Protection
No

Spice Enhancements
none

VM State storage
Automatic

16. Move scsi0 device as #1 to boot from.

Edit: Boot Order

#	Enabled	Device	Description
1	<input checked="" type="checkbox"/>	scsi0	local-lvm:vm-100-disk-0,iouthread=1,size=30G
2	<input type="checkbox"/>	hostpci0	0000:88:00.0
3	<input type="checkbox"/>	hostpci1	0000:89:00.0
4	<input type="checkbox"/>	ide2	none,media=cdrom
5	<input type="checkbox"/>	net0	virtio=BC:24:11:A2:9F:D3,bridge=vmbr0,firewall=1

Drag and drop to reorder

Help

OK

Reset


17. Repeat all the steps from this section on other Proxmox hosts.

StarWind Virtual SAN: Configuration Guide for Proxmox Virtual Environment (VE) [KVM], VSAN
Deployed as a Controller VM (CVM) using GUI

14

Configuring Starwind Virtual San Vm Settings


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



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

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

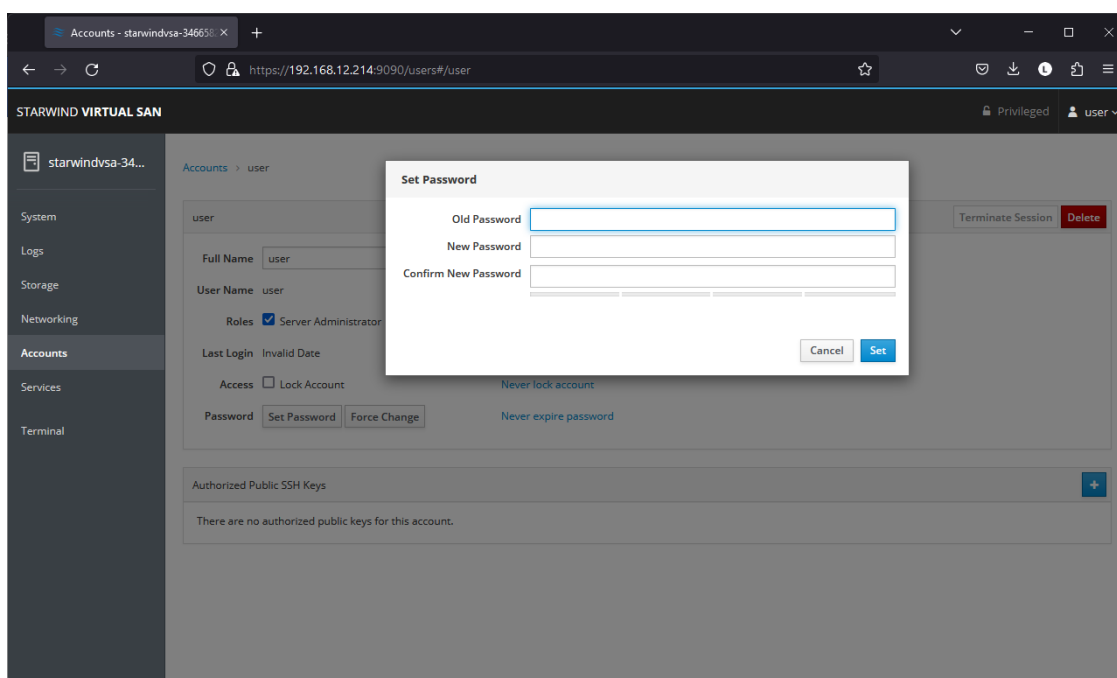
- Username: user
- Password: rds123RDS
- NOTE: Make sure to check the “Reuse my password for privileged tasks” box.



3. Af


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

4. Select a user and click Set Password.




5.

On the left sidebar, click Networking.

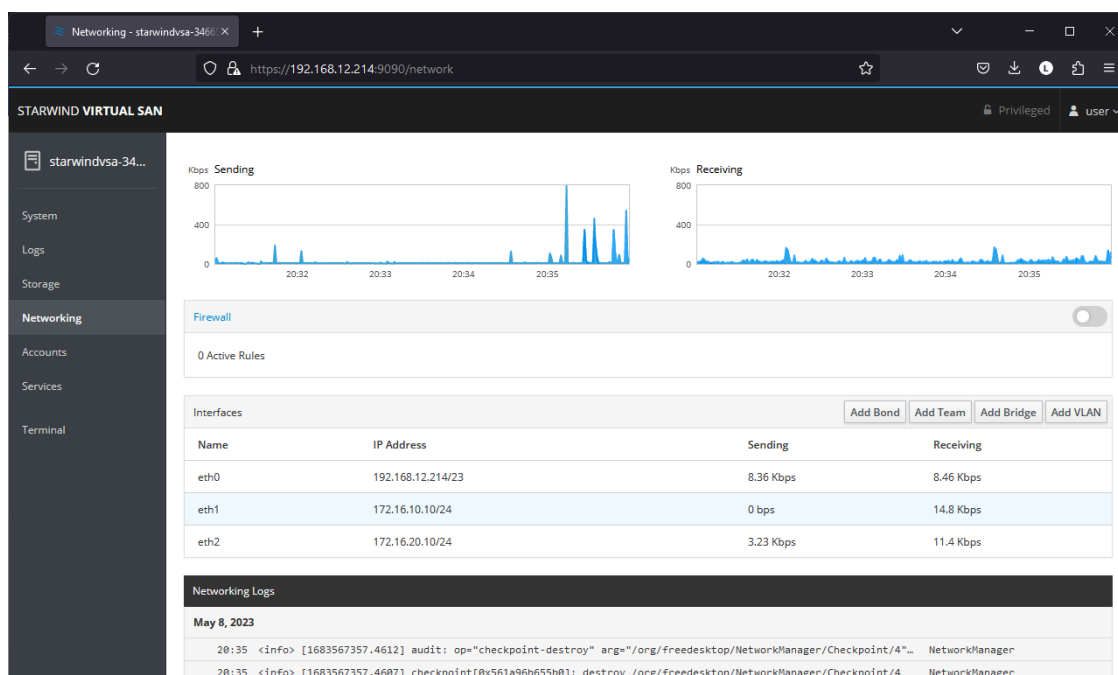


Here, the Management IP address of the StarWind Virtual SAN Virtual Machine can be configured, as well as IP addresses for iSCSI and Synchronization networks. In case the Network interface is inactive, click on the interface, turn it on, and set it to Connect automatically.

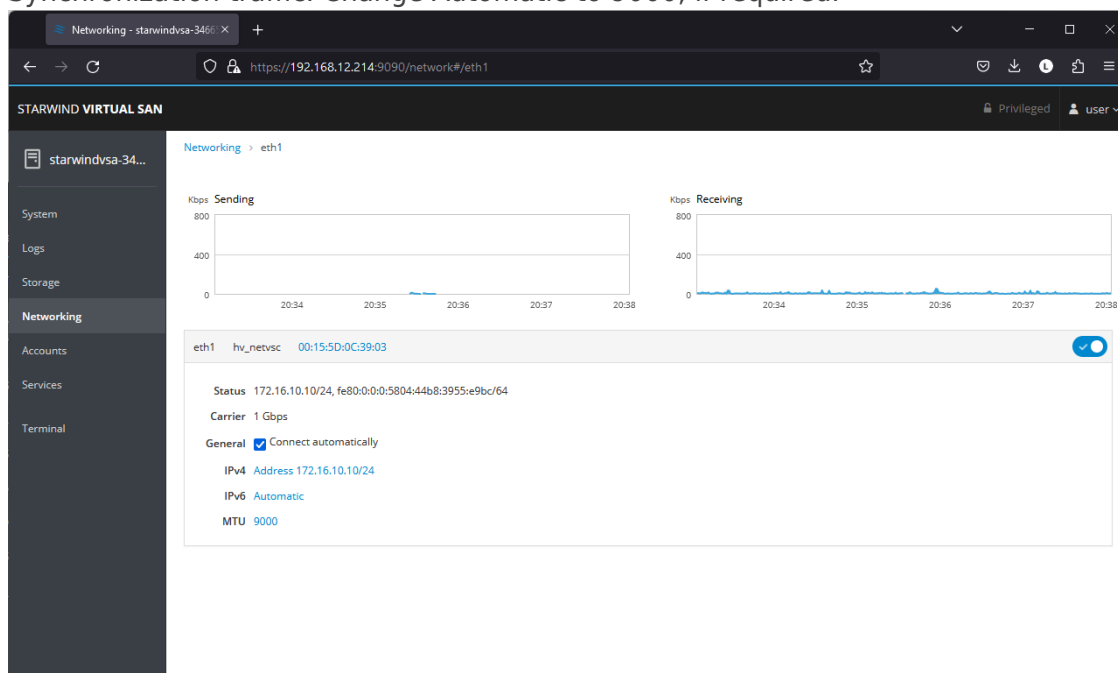
6. Click on Automatic (DHCP) to set the IP address (DNS and gateway – for Management).



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



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



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

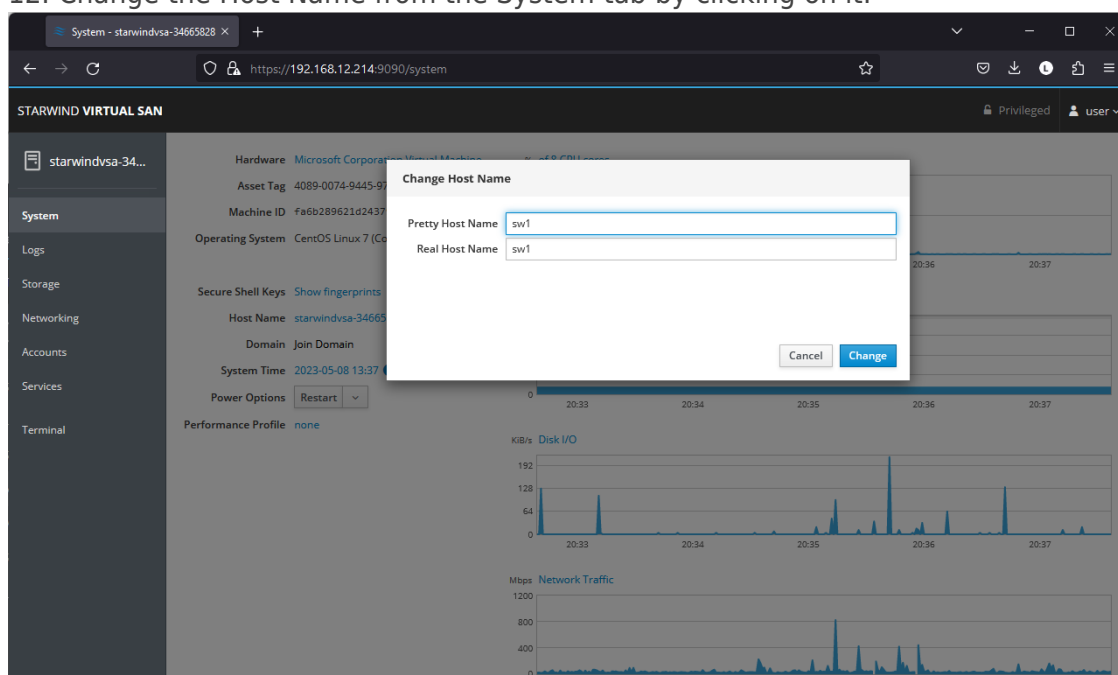
9. Open the file corresponding to the Management interface using a text editor, for example: `sudo nano /etc/sysconfig/network-scripts/ifcfg-eth0`

10. Edit the file:


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

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

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



13. Change System time and NTP settings if required:



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

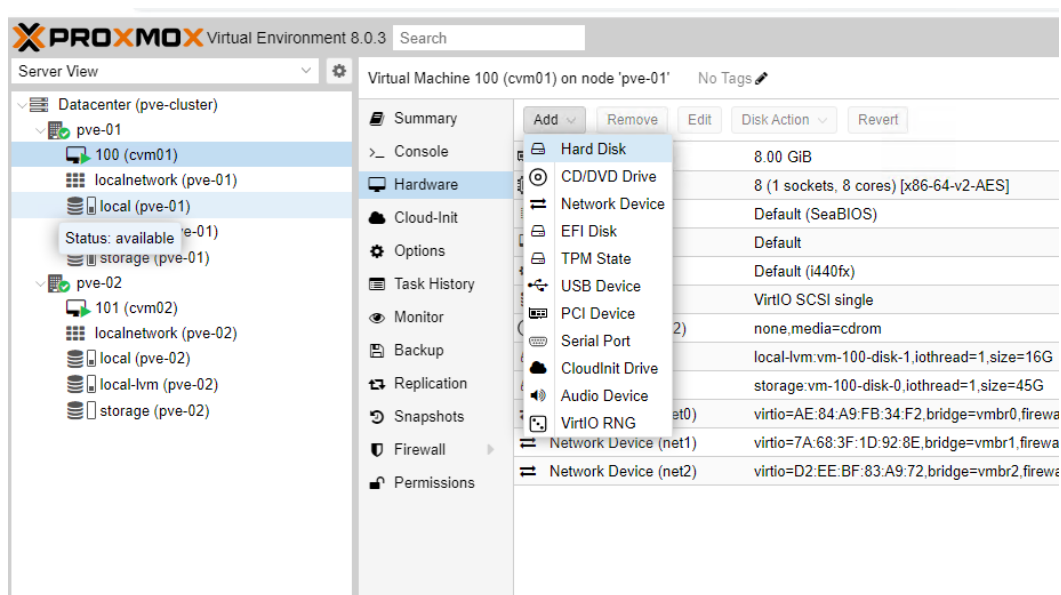
Configuring Storage

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

Please select the required option:

Configuring Storage With Hardware Raid

1. Open VM Hardware page. Click Add -> Hard Disk.



Note. RAID controller can be passed through to the VM.

[https://pve.proxmox.com/wiki/PCI\(e\)_Passthrough](https://pve.proxmox.com/wiki/PCI(e)_Passthrough)

2. Set disk size and storage location. Click Add.

Add: Hard Disk

Disk

Bandwidth

Bus/Device:
SCSI
2

Cache:
Default (No cache)

SCSI Controller:
VirtIO SCSI single

Discard:
☐

Storage:
storage

IO thread:
☒

Disk size (GiB):
100

Format:
Raw disk image (raw)

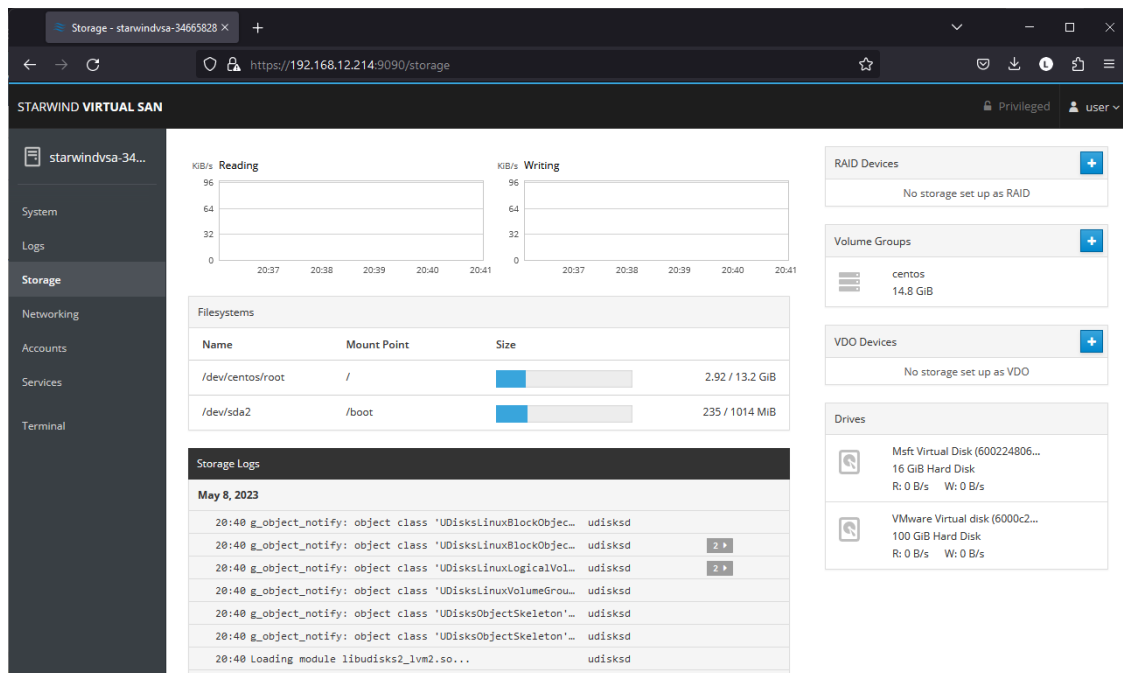
Help

Advanced ☐

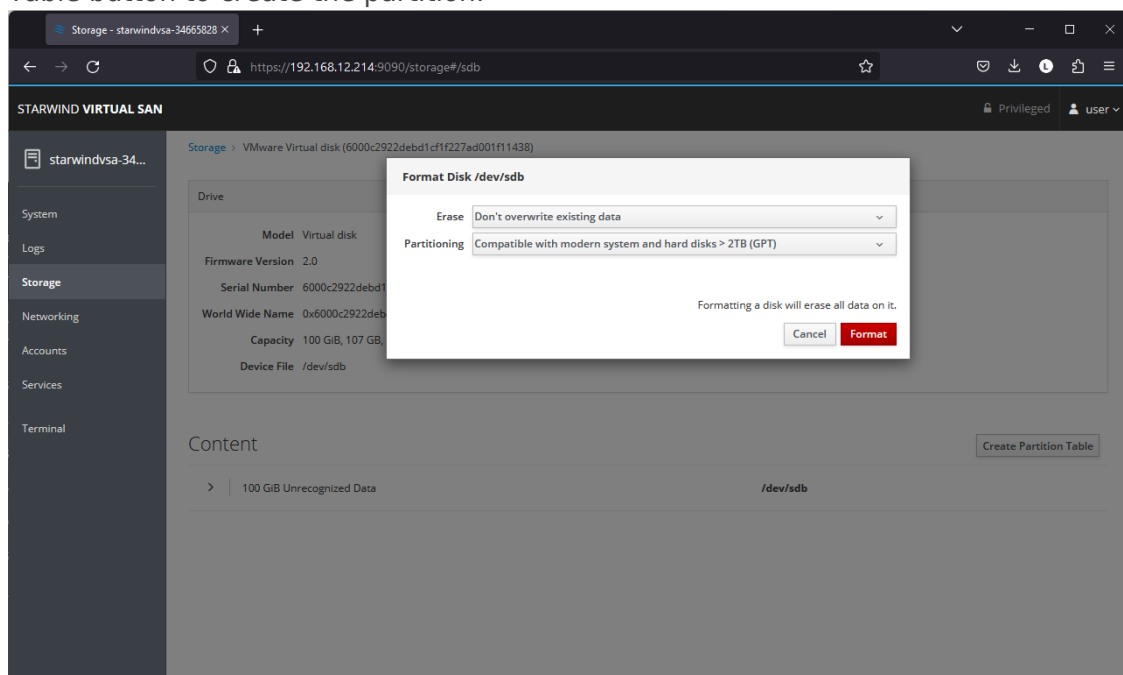
Add

3. Start StarWind VSAN CVM.

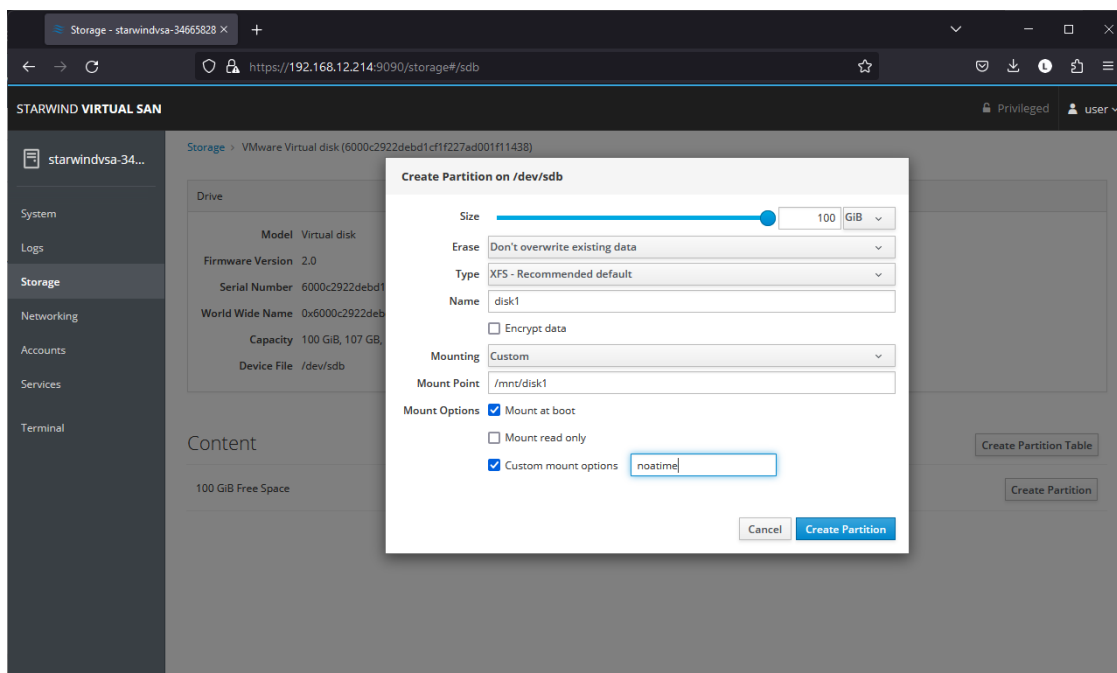
4. Login to StarWind VSAN VM web console and access the Storage section. Locate the recently added disk in the Drives section and choose it.



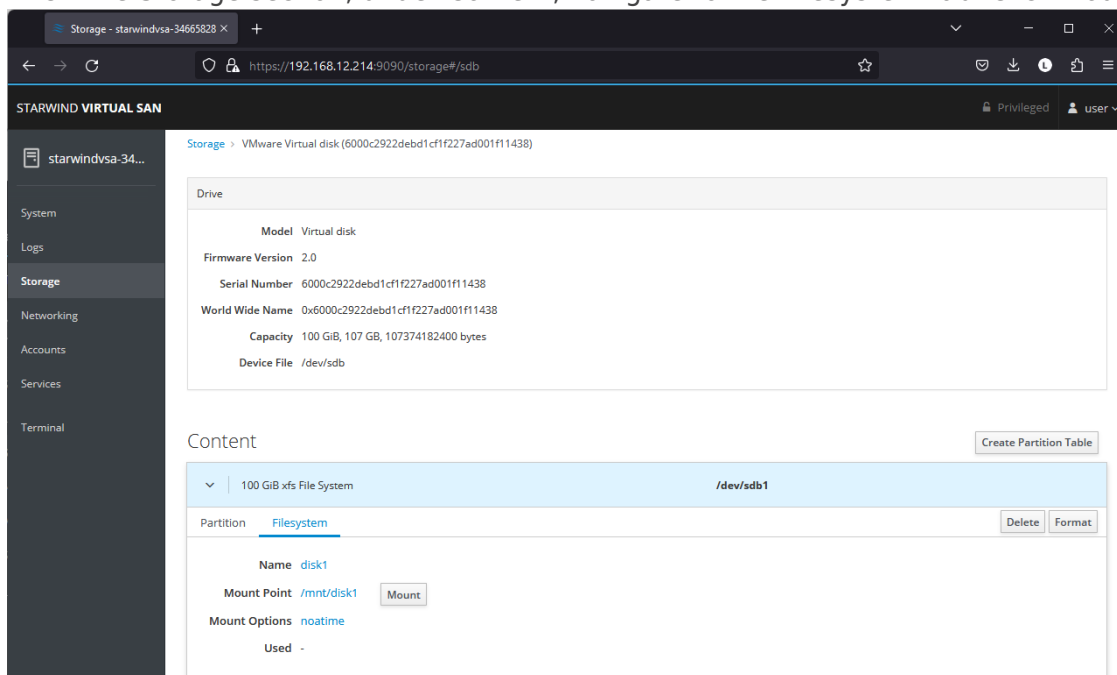
5. The added disk does not have any partitions and filesystem. Press the Create Partition Table button to create the partition.



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



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



Configuring Starwind Management Console


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

NOTE: StarWind Management Console and PowerShell Management Library components are required.


2. Select the appropriate option to apply the StarWind License key.

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

3. Open StarWind Management Console and click Add Server.




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



5. Select the server and click Connect.

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



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

8. Select the appropriate license key.

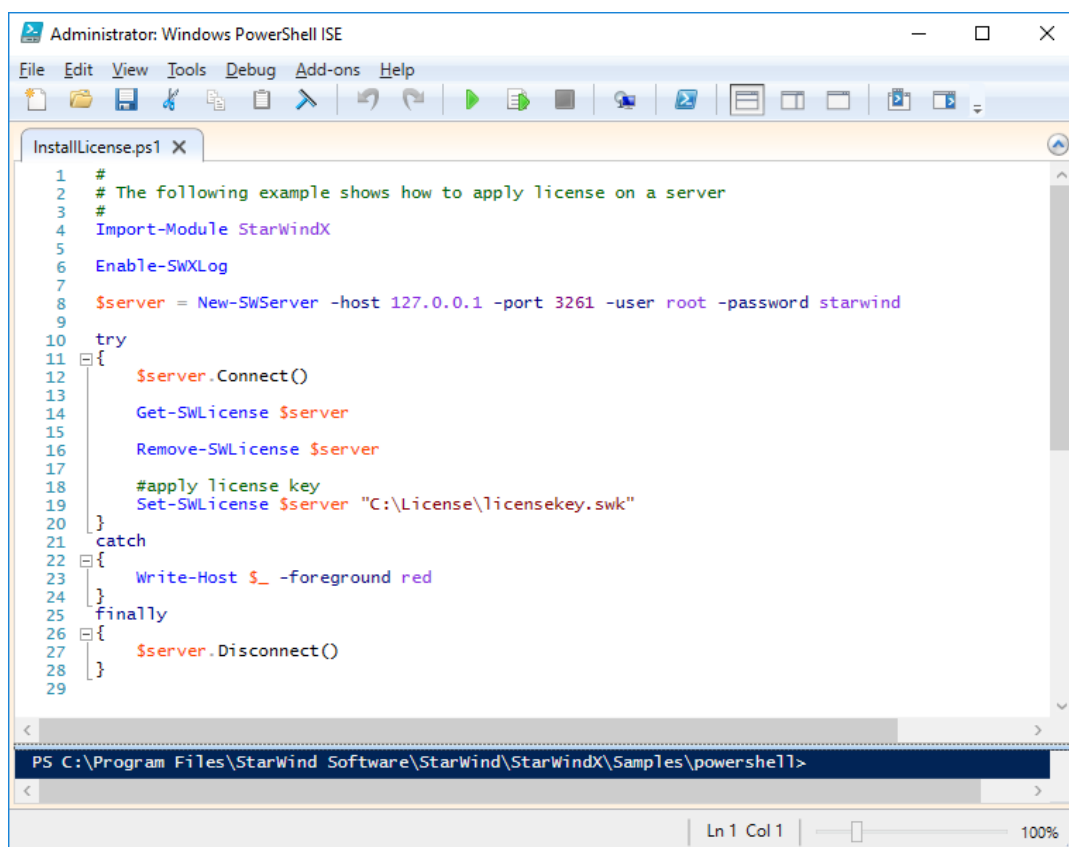
As an alternative, PowerShell can be used. Open StarWind InstallLicense.ps1 script with PowerShell ISE as administrator. It can be found here:

C:\Program Files\StarWind

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

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

Add the path to the license key.



```

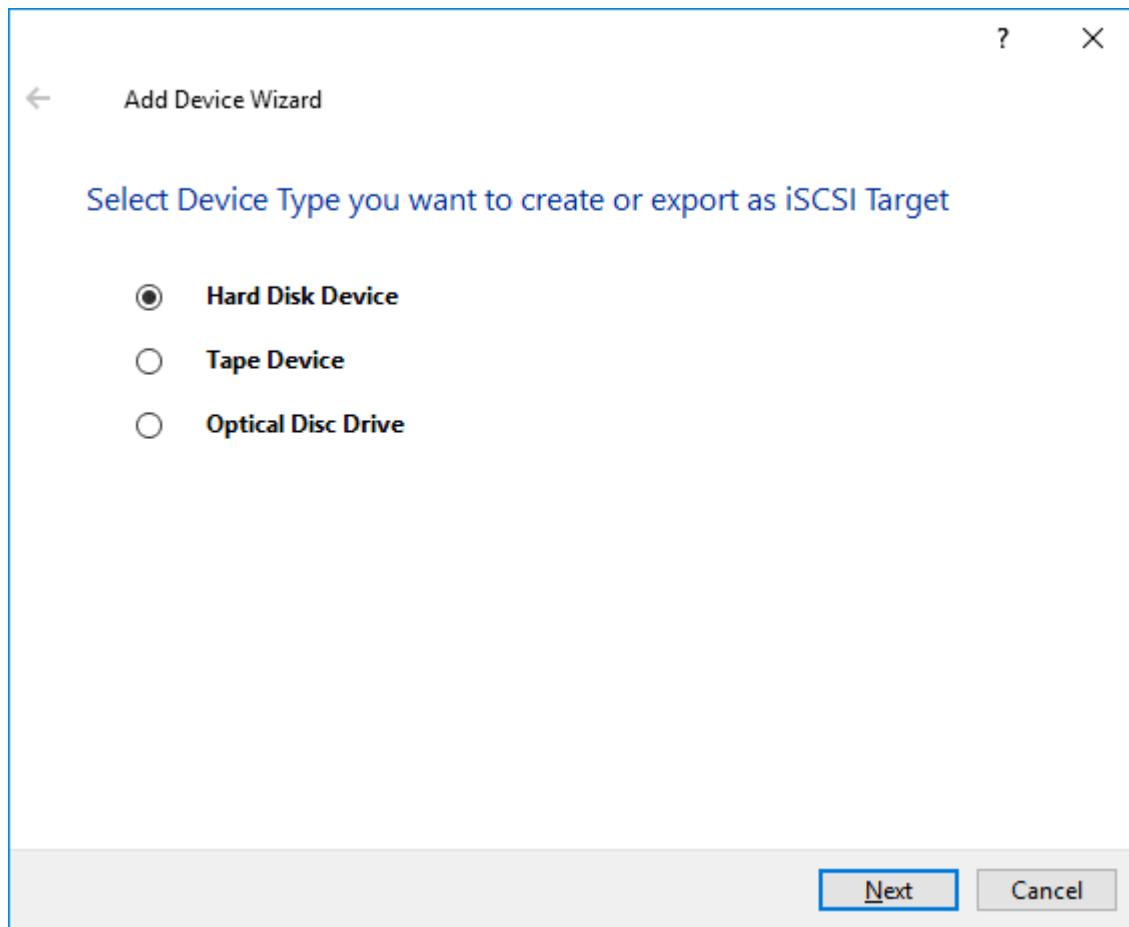
Administrator: Windows PowerShell ISE
File Edit View Tools Debug Add-ons Help
InstallLicense.ps1 X
1 #
2 # The following example shows how to apply license on a server
3 #
4 Import-Module StarWindX
5
6 Enable-SWXLog
7
8 $server = New-SWServer -host 127.0.0.1 -port 3261 -user root -password starwind
9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

```


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

Creating Starwind Devices

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



3. Select Virtual Disk.



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

← 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.

← Add Device Wizard

Specify Device RAM Cache Parameters

Mode

☐ **Write-Back**
Writes are performed asynchronously, actual Writes to Disk are delayed, Reads are cached

☐ **Write-Through**
Writes are performed synchronously, Reads are cached

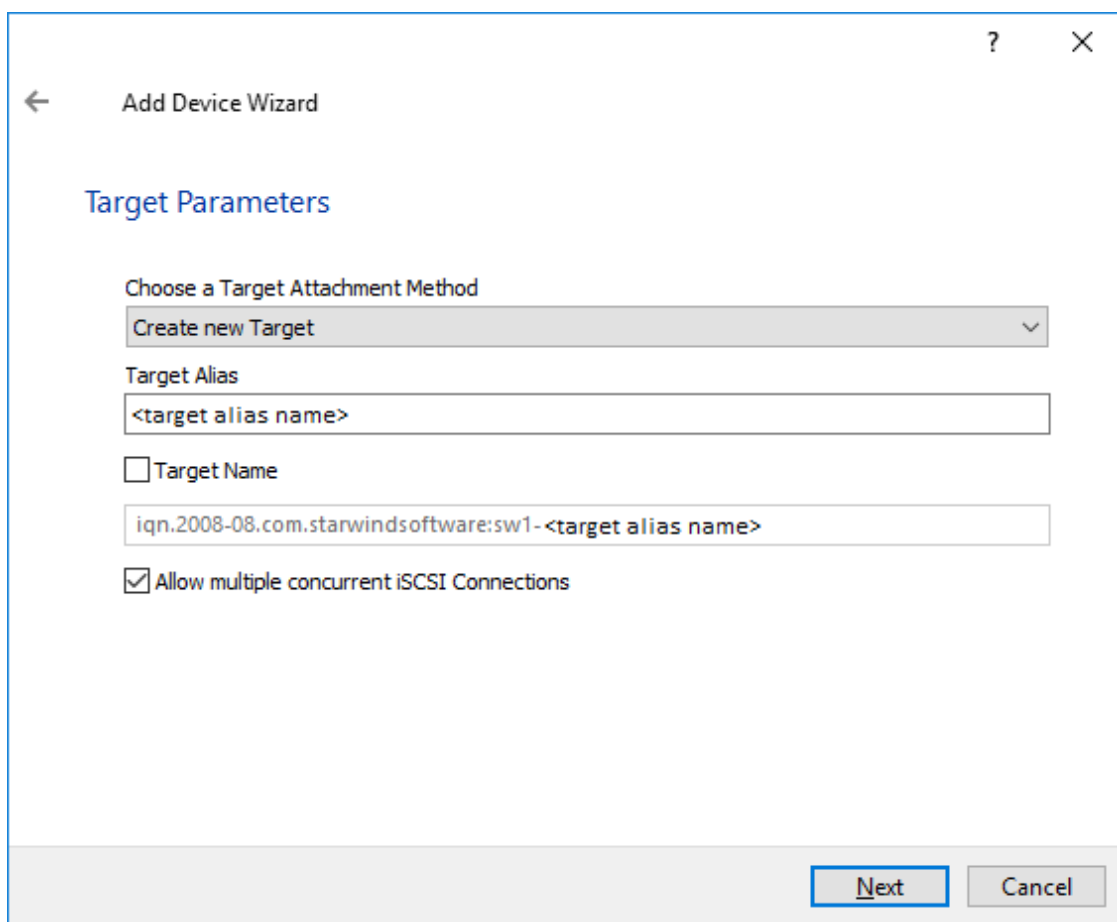
☒ **N/A**
Reads and Writes are not cached

☐ Set Maximum available Size

Size: MB

Next Cancel

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.



← Add Device Wizard

Target Parameters

Choose a Target Attachment Method

Create new Target

Target Alias

<target alias name>


☐ Target Name

iqn.2008-08.com.starwindsoftware:sw1-<target alias name>

☒ Allow multiple concurrent iSCSI Connections

Next Cancel

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



9. Click Close to finish the device creation.

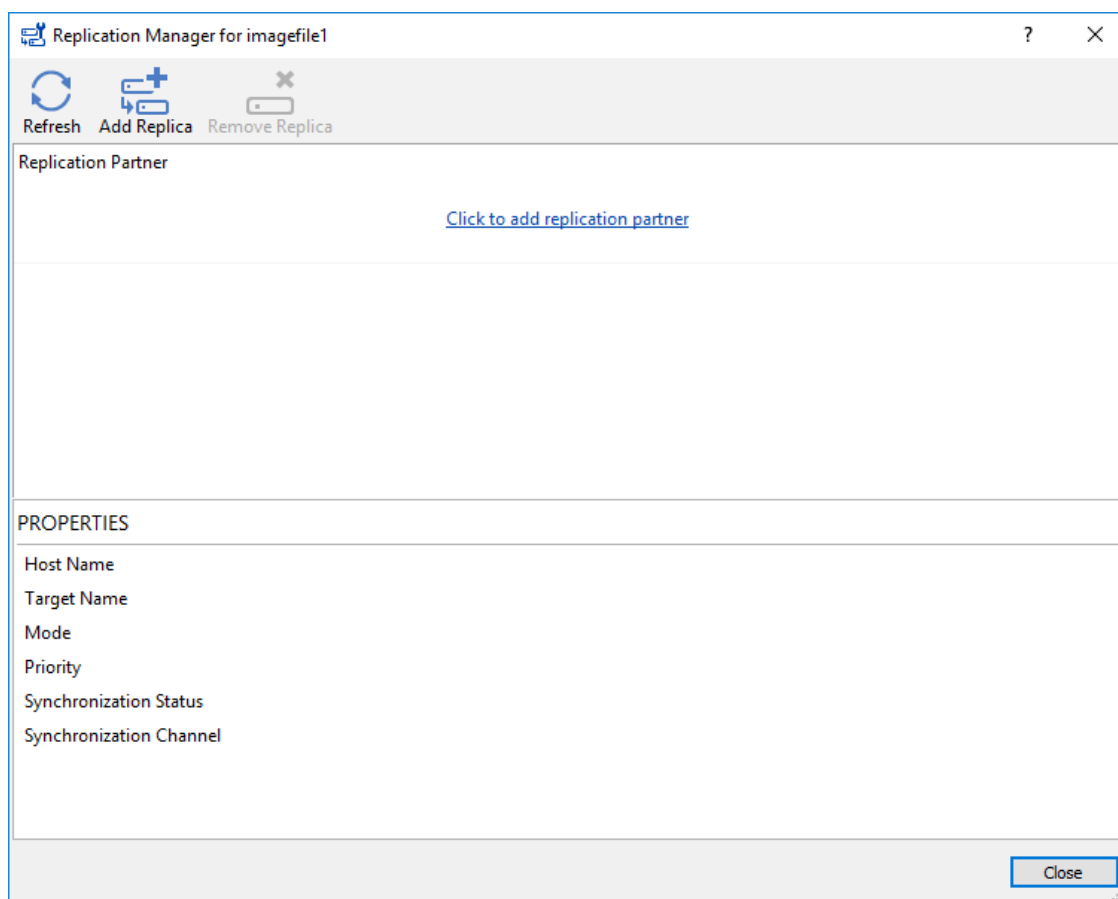
10. The successfully added devices appear in the StarWind Management Console.

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:

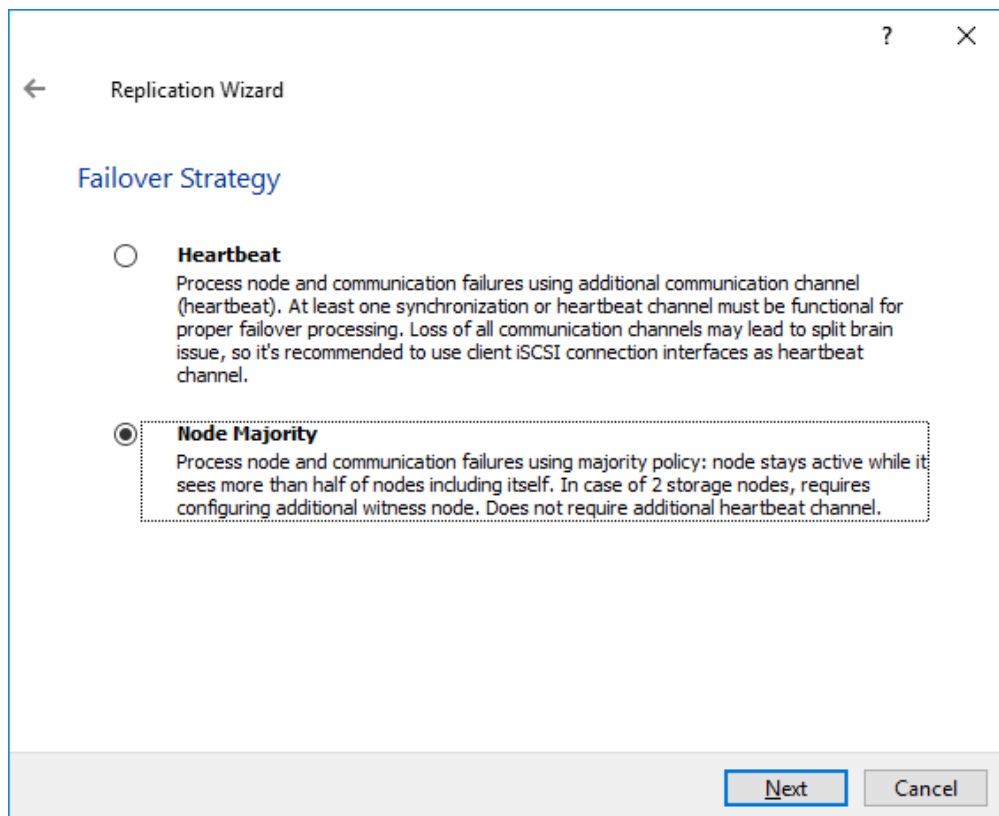
SMB Witness is a file, located on SMB share, which can be accessed by both nodes and help them to eliminate the split-brain issue in case of synchronization connection

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/>

Connecting Starwind Ha Storage To Proxmox Hosts

1. Connect to Proxmox host via SSH and install multipathing tools.

```
pve# apt-get install multipath-tools
```

2. Edit nano /etc/iscsi/initiatorname.iscsi setting the initiator name.

```
GNU nano 7.2 /etc/iscsi/initiatorname.iscsi
## DO NOT EDIT OR REMOVE THIS FILE!
## If you remove this file, the iSCSI daemon will not start.
## If you change the InitiatorName, existing access control lists
## may reject this initiator. The InitiatorName must be unique
## for each iSCSI initiator. Do NOT duplicate iSCSI InitiatorNames.
InitiatorName=iqn.1993-08.org.debian:CVMO1
```

3. Edit /etc/iscsi/iscsid.conf setting the following parameters:

```
node.startup = automatic
node.session.timeo.replacement_timeout = 15
node.session.scan = auto
```

Note. node.startup = manual is the default parameter, it should be changed to node.startup = automatic.

4. Create file /etc/multipath.conf using the following command:

```
touch /etc/multipath.conf
```

5. Edit /etc/multipath.conf adding the following content:

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

```

}

defaults {
    polling_interval          2
    path_selector             "round-robin 0"
    path_grouping_policy     multibus
    uid_attribute             ID_SERIAL
    rr_min_io                 100
    failback                  immediate
    no_path_retry             queue
    user_friendly_names       yes
}

```

6. Run iSCSI discovery on both nodes:

```

pve# iscsiadm -m discovery -t st -p 10.20.1.10
pve# iscsiadm -m discovery -t st -p 10.20.1.20

```

7. Connect iSCSI LUNs:

```

pve# iscsiadm -m node -T iqn.2008-08.com.starwindsoftware:sw1-
ds1 -p 10.20.1.10 -l
pve# iscsiadm -m node -T iqn.2008-08.com.starwindsoftware:sw2-
ds1 -p 10.20.1.20 -l

```

8. Get WWID of StarWind HA device:

```

/lib/udev/scsi_id -g -u -d /dev/sda

```

9. The wwid must be added to the file '/etc/multipath/wwids'. To do this, run the following command with the appropriate wwid:

```

multipath -a %WWID%

```

10. Restart multipath service.

```

systemctl restart multipath-tools.service

```

11. Check if multipathing is running correctly:

```
pve# multipath -ll
```

12. Repeat steps 1-11 on every Proxmox host.

13. Create LVM PV on multipathing device:

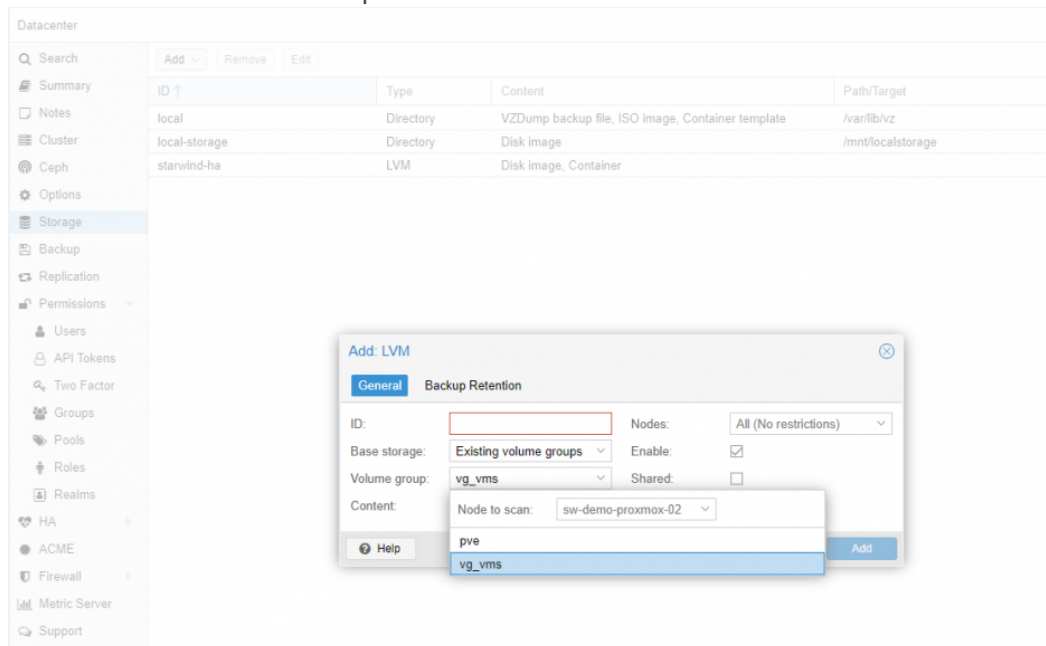
```
pve# pvcreate /dev/mapper/mpatha
```

where mpatha - alias for StarWind LUN

14. Create VG on LVM PV:

```
pve# vgcreate vg-vms /dev/mapper/mpath0
```

15. Login to Proxmox via Web and go to Datacenter -> Storage. Add new LVM storage based on VG created on top of StarWind HA Device. Enable Shared checkbox. Click Add.








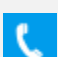
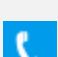
16. Login via SSH to all hosts and run the following command:

```
pvscan --cache
```

Conclusion

Following this guide, a Proxmox Cluster was deployed and configured with StarWind Virtual SAN (VSAN) running in a CVM on each host. As a result, a virtual shared storage “pool” accessible by all cluster nodes was created for storing highly available virtual machines.

Contacts

US Headquarters	EMEA and APAC
 +1 617 829 44 95	 +44 2037 691 857 (United Kingdom)
 +1 617 507 58 45	 +49 800 100 68 26 (Germany)
 +1 866 790 26 46	 +34 629 03 07 17 (Spain and Portugal)
	 +33 788 60 30 06 (France)

Customer Support Portal: <https://www.starwind.com/support>

Support Forum: <https://www.starwind.com/forums>

Sales: sales@starwind.com

General Information: info@starwind.com



StarWind Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA
www.starwind.com ©2024, StarWind Software Inc. All rights reserved.