

StarWind Virtual SAN: Configuration Guide for VMware vSphere [ESXi], VSAN Deployed as a Controller VM using Web UI

2026

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) and Proxmox version 8 and version 9

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



1. Make sure that a oVirt engine is installed on a separate host.
2. Deploy oVirt on each server and add them to oVirt engine.
3. Define at least 2x network interfaces on each node that will be used for the Synchronization and iSCSI/StarWind heartbeat traffic. Do not use iSCSI/Heartbeat and Synchronization channels over the same physical link. Synchronization and iSCSI/Heartbeat links can be connected either via redundant switches or directly between the nodes (see diagram above).
4. Separate Logical Networks should be created for iSCSI and Synchronization traffic based on the selected before iSCSI and Synchronization interfaces. Using oVirt engine Networking page create two Logical Networks: one for the iSCSI/StarWind Heartbeat channel (iSCSI) and another one for the Synchronization channel (Sync).
5. Add physical NIC to Logical network on each host and configure static IP addresses. In this document, the 172.16.10.x subnet is used for iSCSI/StarWind heartbeat traffic, while

172.16.20.x subnet is used for the Synchronization traffic.

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

Deploying Starwind Virtual San Cvm

1. Download StarWind VSAN CVM


KVM: <https://www.starwindsoftware.com/vsan#download>

2. Extract the VM CVM.qcow2 file from the downloaded archive.

3. Upload CVM.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-xr-x	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)

Storage: local
ISO image:

Guest OS:

Type: Linux
Version: 6.x - 2.6 Kernel

☐ 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 at least 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):

? 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:
Model:

VLAN Tag:
MAC address:

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	CVM01
net0	virtio,bridge=vbr0,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

NOTE: Make sure that the VM number is noted before proceeding with the next step.

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

```
qm importdisk 100 /root/CVM.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.



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



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


Initial Configuration Wizard

1. Start the StarWind Virtual SAN Controller Virtual Machine.
2. Launch the VM console to view the VM boot process and obtain the IPv4 address of the Management network interface.


NOTE: If the VM does not acquire an IPv4 address from a DHCP server, use the Text-based User Interface (TUI) to set up the Management network manually.

Default credentials for TUI: user/rds123RDS


3. Using a web browser, open a new tab and enter the VM's IPv4 address to access the StarWind VSAN Web Interface. On the Your connection is not private screen, click Advanced and then select Continue to...




4. On the Welcome to StarWind Appliance screen, click Start to launch the Initial Configuration Wizard.



5. On the License step, upload the StarWind Virtual SAN license file.




6. On the EULA step, read and accept the End User License Agreement to continue.




7. On the Management network step, review or edit the network settings and click Next.

IMPORTANT: The use of Static IP mode is highly recommended.




8. On the Static hostname, specify the hostname for the virtual machine and click Next.




The screenshot shows the 'StarWind Appliance Initial configuration' window. On the left sidebar, the steps are: License (checked), EULA (checked), Management network (checked), Static hostname (selected), Administrator account, Summary, and Configuration. The main panel is titled 'Verify hostname' and contains the instruction 'Check the current appliance hostname and modify it if required'. Below this is a note: 'Use Latin letters, numbers, and dash'. A text input field labeled 'Hostname' contains the value 'SW1'. At the bottom right, there are 'Back' and 'Next' buttons.

9. On the Administrator account step, specify the credentials for the new StarWind Virtual SAN administrator account and click Next.




The screenshot shows the 'StarWind Appliance Initial configuration' window. On the left sidebar, the steps are: License (checked), EULA (checked), Management network (checked), Static hostname (checked), Administrator account (selected), Summary, and Configuration. The main panel is titled 'Create administrator account' and contains the instruction 'Specify new credentials for the appliance administrator account'. It features three input fields: 'Username' with the value 'admin', 'Password' (masked with dots), and 'Confirm password' (masked with dots). Below these are optional fields for 'Full name' and 'E-mail'. At the bottom right, there are 'Back' and 'Next' buttons.


10. Wait until the Initial Configuration Wizard configures StarWind Virtual SAN for you.



11. Please standby until the Initial Configuration Wizard configures StarWind VSAN for you.



12. After the configuration process is completed, click Finish to install the StarWind vCenter Plugin immediately, or uncheck the checkbox to skip this step and proceed to the Login page.




13. Repeat steps 1 through 12 on each Windows Server host.


Add Appliance

To create replicated, highly available storage, add partner appliances that use the same StarWind Virtual SAN license key.


1. Navigate to the Appliances page and click Add to open the Add appliance wizard.
2. On the Credentials step, enter the IP address and credentials for the partner StarWind Virtual SAN appliance, then click Next.



3. Provide credentials of partner appliance.




4. Wait for the connection to be established and the settings to be validated




5. On the Summary step, review the properties of the partner appliance, then click Add

Appliance.




Configure Ha Networking

1. Navigate to the Network page and open Configure HA networking wizard.



2. On the Appliances step, select either 2 partner appliances to configure two-way replication, or 3 appliances for three-way replication, then click Next.

NOTE: The number of appliances in the cluster is limited by your StarWind Virtual SAN license.




3. On the Data Network step, select the network interfaces designated to carry iSCSI or NVMe-oF storage traffic. Assign and configure at least one interface on each appliance (in our example: 172.16.10.10 and 172.16.10.20) with a static IP address in a unique network (subnet), specify the subnet mask and Cluster MTU size.

IMPORTANT: For a redundant, high-availability configuration, configure at least 2 network interfaces on each appliance. Ensure that the Data Network interfaces are interconnected between appliances through multiple direct links or via redundant switches.


4. Assign MTU value on all selected network adapters, e.g. 1500 or 9000 bytes. If you are using network switches with the selected Data Network adapters, ensure that they are configured with the same MTU size value. In case of MTU settings mismatch, stability and performance issues might occur on the whole setup.

NOTE: Setting MTU to 9000 bytes on some physical adapters (like Intel Ethernet Network Adapter X710, Broadcom network adapters, etc.) might cause stability and performance issues depending on the installed network driver. To avoid them, use 1500 bytes MTU size or install the stable version of the driver.

5. Once configured, click Next to validate network settings.



6. The warning might appear if a single data interface is configured. Click Yes, continue to proceed with the configuration.




7. On the Replication Network step, select the network interfaces designated to carry the traffic for synchronous replication. Assign and configure at least one interface on each appliance with a static IP address in a unique network (subnet), specify the subnet mask and Cluster MTU size.

IMPORTANT: For a redundant, high-availability configuration, configure at least 2 network interfaces on each appliance. Ensure that the Replication Network interfaces are interconnected between appliances through multiple direct links or via redundant switches.


8. Assign MTU value on all selected network adapters, e.g. 1500 or 9000 bytes. If you are using network switches with the selected Replication Network adapters, ensure that they are configured with the same MTU size value. In case of MTU settings mismatch, stability and performance issues might occur on the whole setup.

NOTE: Setting MTU to 9000 bytes on some physical adapters (like Intel Ethernet Network Adapter X710, Broadcom network adapters, etc.) might cause stability and performance issues depending on the installed network driver. To avoid them, use 1500 bytes MTU size or install the stable version of the driver.


9. Once configured, click Next to validate network settings.




10. If only one Replication Network interface is configured on each partner appliance, a warning message will pop up. Click Yes, continue to acknowledge the warning and proceed.



11. Wait for the configuration completion.



12. On the Summary step, review the specified network settings and click Configure to apply the changes.



Add Physical Disks


Attach physical storage to StarWind Virtual SAN Controller VM:

- Ensure that all physical drives are connected through an HBA or RAID controller.
- To get the optimal storage performance, add HBA, RAID controllers, or NVMe SSD drives to StarWind CVM via a passthrough device.

For detailed instructions, refer to [Microsoft's documentation on DDA](#). Also, find the storage provisioning guidelines in the [KB article](#).

Create Storage Pool


1. Navigate to the Storage pools page and click the + button to open the Create storage pool wizard .



2. On the Appliance step, select partner appliances on which to create new storage pools, then click Next.


NOTE: Select 2 appliances for configuring storage pools if you are deploying a two-node cluster with two-way replication, or select 3 appliances for configuring a three-node

cluster with a three-way mirror.




3. On the Physical disks step, select physical disks to be pooled on each node, then click Next.

IMPORTANT: Select an identical type and number of disks on each appliance to create storage pools with a uniform configuration.



4. On the Profile step, select one of the preconfigured storage profiles, or choose Manual to configure the storage pool manually based on your redundancy, capacity, and performance requirements, then click Next.




NOTE: Hardware RAID, Linux Software RAID, and ZFS storage pools are supported. To

simplify the configuration of storage pools, preconfigured storage profiles are provided. These profiles recommend a pool type and layout based on the attached storage:

- High capacity – creates Linux Software RAID-5 to maximize storage capacity while maintaining redundancy.
- High performance – creates Linux Software RAID-10 to maximize storage performance while maintaining redundancy.
- Hardware RAID – configures a hardware RAID virtual disk as a storage pool. This option is available only if a hardware RAID controller is passed through to the StarWind Virtual SAN.
- Better redundancy – creates ZFS Striped RAID-Z2 (RAID 60) to maximize redundancy while maintaining high storage capacity.
- Manual – allows users to configure any storage pool type and layout with the attached storage.

5. On the Summary step, review the storage pool settings and click Create to configure new storage pools on the selected appliances.



NOTE: The storage pool configuration may take some time, depending on the type of pooled storage and the total storage capacity. Once the pools are created, a notification will appear in the upper right corner of the Web UI.


IMPORTANT: In some cases, additional tweaks are required to optimize the storage performance of the disks added to the Controller Virtual Machine. Please follow the steps in this KB to change the scheduler type depending on the disks type:

<https://knowledgebase.starwindsoftware.com/guidance/starwind-vsan-for-vsphere-changi>

[ng-linux-i-o-scheduler-to-optimize-storage-performance/](#)


Create Volume

1. Navigate to the Volumes page and click the + button to open the Create volume wizard.




2. On the Storage pool step, select partner appliances on which to create new volumes, then click Next.


NOTE: Select 2 appliances for configuring volumes if you are deploying a two-node cluster with two-way replication, or select 3 appliances for configuring a three-node cluster with a three-way mirror.




3. On the Settings step, specify the volume name and size, then click Next.



4. On the Filesystem type step, select Standard, then click Next.



5. Review Summary and click the Create button to create the pool.




Create Ha Lun Using Webui


This section describes how to create LUN in Web UI. This option is available for the setups with Commercial, Trial, and NFR licenses applied.

For setups with a Free license applied, the PowerShell script should be used to create the LUN – please follow the steps described in the section: Create StarWind HA LUNs using PowerShell

1. Navigate to the LUNs page and click the + button to open the Create LUN wizard.




2. On the Protocols step, select the preferred storage protocol and click Next.



3. On the LUN availability step, select the High availability and click Next.


NOTE: The availability options for a LUN can be Standalone (without replication) or High Availability (with 2-way or 3-way replication), and are determined by the StarWind Virtual SAN license.

Below are the steps for creating a high-availability iSCSI LUN.




4. On the Appliances step, select partner appliances that will host new LUNs and click Next.

IMPORTANT: Selected partner appliances must have identical hardware configurations, including CPU, RAM, storage, and networking.




5. On the Volumes step, select the volumes for storing data on the partner appliances and click Next.

IMPORTANT: For optimal performance, the selected volumes must have identical underlying storage configurations.




6. On the Failover strategy step, select the preferred failover strategy and click Next.

NOTE: The failover strategies for a LUN can be Heartbeat or Node Majority. In case of 2-nodes setup and None Majority failover strategy, Node witness (requires an additional third witness node), or File share witness (requires an external file share) should be configured. These options are determined by StarWind Virtual SAN license and setup configuration. Below are the steps for configuring the Heartbeat failover strategy in a two-node cluster.




7. On the LUN settings step, specify the LUN name, size, block size, then click Next.

NOTE: For high-availability configurations, ensure that MPIO checkbox is selected.



8. On the Summary step, review the LUN settings and click Create to configure new LUNs on the selected volumes.



Creating Starwind Ha Luns Using Powershell

1. Open PowerShell ISE as Administrator.
2. Open StarWindX sample CreateHA_2.ps1 using PowerShell ISE. It can be found here: C:\Program Files\StarWind Software\StarWind\StarWindX\Samples\

NOTE: The script below creates a 1TB size 2-node HA device with a heartbeat failover strategy on StarWind nodes with management IP addresses 192.168.0.1 and 192.168.0.2 correspondingly.

The IP addresses 172.16.10.1 and 172.16.10.2 are used as heartbeat interfaces along with 192.168.0.1 and 192.168.0.2 for redundancy.

The IP addresses 172.16.20.1 and 172.16.20.2 on each node correspondingly as well as 172.16.21.1 and 172.16.21.2 are used for the devices synchronization between the nodes.

The script does not create a directory. Make sure you create the directory listed as \$imagePath value before running the script.

Make sure to open 3261 and 3260 ports.

The same approach applies to CreateHA_3.ps1 which allow to create a 3-way replica HA

device.

3. Configure script parameters according to the following example:

```
param($addr="192.168.0.1", $port=3261, $user="root",
$password="starwind",
    $addr2="192.168.0.2", $port2=$port, $user2=$user,
$password2=$password,
#common
    $initMethod="NotSynchronize",
    $size=1048576,
    $sectorSize=512,
    $failover=0,
    $bmpType=1,
    $bmpStrategy=0,
#primary node
    $imagePath="/mnt/sdb1/volume1",
    $imageName="device1",
    $createImage=$true,
    $storageName="",
    $targetAlias="target1",
    $poolName="pool1",
    $syncSessionCount=1,
    $aluaOptimized=$true,
    $cacheMode="none",
    $cacheSize=0,
    $syncInterface="#p2=172.16.20.2:3260,172.16.21.2:3260",
    $hbInterface="#p2=172.16.10.2:3260,192.168.0.2:3260",
    $createTarget=$true,
    $bmpFolderPath="",
#secondary node
    $imagePath2="/mnt/sdb1/volume1",
    $imageName2="device1",
    $createImage2=$true,
    $storageName2="",
    $targetAlias2="target1",
    $poolName2="pool1",
    $syncSessionCount2=1,
    $aluaOptimized2=$false,
    $cacheMode2=$cacheMode,
    $cacheSize2=$cacheSize,
    $syncInterface2="#p1=172.16.20.1:3260,172.16.21.1:3260",
```

```

    $hbInterface2="#p1=172.16.10.1:3260,192.168.0.1:3260",
    $createTarget2=$true,
    $bmpFolderPath=""
  )
Import-Module StarWindX

try
{
    Enable-SWXLog -level SW_LOG_LEVEL_DEBUG

    $server = New-SWServer -host $addr -port $port -user
$user -password $password

    $server.Connect()

    $firstNode = new-Object Node

    $firstNode.HostName = $addr
    $firstNode.HostPort = $port
    $firstNode.Login = $user
    $firstNode.Password = $password
    $firstNode.ImagePath = $imagePath
    $firstNode.ImageName = $imageName
    $firstNode.Size = $size
    $firstNode.CreateImage = $createImage
    $firstNode.StorageName = $storageName
    $firstNode.TargetAlias = $targetAlias
    $firstNode.AutoSynch = $autoSynch
    $firstNode.SyncInterface = $syncInterface
    $firstNode.HBInterface = $hbInterface
    $firstNode.PoolName = $poolName
    $firstNode.SyncSessionCount = $syncSessionCount
    $firstNode.ALUAOptimized = $aluaOptimized
    $firstNode.CacheMode = $cacheMode
    $firstNode.CacheSize = $cacheSize
    $firstNode.FailoverStrategy = $failover
    $firstNode.CreateTarget = $createTarget
    $firstNode.BitmapStoreType = $bmpType
    $firstNode.BitmapStrategy = $bmpStrategy
    $firstNode.BitmapFolderPath = $bmpFolderPath
    #
    # device sector size. Possible values: 512 or 4096(May
be incompatible with some clients!) bytes.
    #

```



```

    $firstNode.SectorSize = $sectorSize
    $secondNode = new-Object Node

    $secondNode.HostName = $addr2
    $secondNode.HostPort = $port2
    $secondNode.Login = $user2
    $secondNode.Password = $password2
    $secondNode.ImagePath = $imagePath2
    $secondNode.ImageName = $imageName2
    $secondNode.CreateImage = $createImage2
    $secondNode.StorageName = $storageName2
    $secondNode.TargetAlias = $targetAlias2
    $secondNode.AutoSynch = $autoSynch2
    $secondNode.SyncInterface = $syncInterface2
    $secondNode.HBInterface = $hbInterface2
    $secondNode.SyncSessionCount = $syncSessionCount2
    $secondNode.ALUAOptimized = $aluaOptimized2
    $secondNode.CacheMode = $cacheMode2
    $secondNode.CacheSize = $cacheSize2
    $secondNode.FailoverStrategy = $failover
    $secondNode.CreateTarget = $createTarget2
    $secondNode.BitmapFolderPath = $bmpFolderPath2
    $device = Add-HADevice -server $server -firstNode
    $firstNode -secondNode $secondNode -initMethod $initMethod
    while ($device.SyncStatus -ne
    [SwHaSyncStatus]::SW_HA_SYNC_STATUS_SYNC)
    {
        $syncPercent =
    $device.GetPropertyValue("ha_synch_percent")
        Write-Host "Synchronizing: $($syncPercent)%" -
    foreground yellow

        Start-Sleep -m 2000

        $device.Refresh()
    }
  }
catch
{
    Write-Host $_ -foreground red
}
finally
{
    $server.Disconnect()
  }

```

```
}
```

Detailed explanation of script parameters:

-addr, -addr2 — partner nodes IP address.

Format: string. Default value: 192.168.0.1, 192.168.0.2

allowed values: localhost, IP-address

-port, -port2 — local and partner node port.

Format: string. Default value: 3261

-user, -user2 — local and partner node user name.

Format: string. Default value: root

-password, -password2 — local and partner node user password.

Format: string. Default value: starwind

#common

-initMethod – set the initial synchronization option.

Format: string.

Values:

Clear – default

NotSynchronize – skips synchronization (works ONLY IF THERE IS NO DATA TO SKIP THE ORIGINAL SYNCHRONIZATION).

SyncFromFirst or SyncFromSecond or SyncFromThird – runs full synchronization from the specific node. Use it for recreating replicas.

-size – set size for HA-device (in MB)

Format: integer. Default value: 12

-sectorSize – set sector size for HA-device

Format: integer. Default value: 512

allowed values: 512, 4096

-failover – set type failover strategy

Format: integer. Default value: 0 (Heartbeat)

allowed values: 0, 1 (Node Majority)

-bmpType – set bitmap type, is set for both partners at once

Format: integer. Default value: 1 (RAM)

allowed values: 1, 2 (DISK)

-bmpStrategy – set journal strategy, is set for both partners at once

Format: integer. Default value: 0

allowed values: 0, 1 – Best Performance (Failure), 2 – Fast Recovery (Continuous)

-storageName is used only if you plan to add the partner to the existing device. For

CreateHA_2.ps1 use, leave it as is.

#primary node

-imagePath – set path to store the device file

Format: string. Default value: My computer\C\starwind". For Linux the following format should be used: "VSA Storage\mnt\mount_point"

-imageName – set name device

Format: string. Default value: masterImg21

-createImage – set create image file

Format: boolean. Default value: true

-targetAlias – set alias for target

Format: string. Default value: targetha21

-poolName – set storage pool. Do not change it and keep default value.

Format: string. Default value: pool1

-aluaOptimized – set Alua Optimized

Format: boolean. Default value: true

-cacheMode – set type L1 cache (optional parameter)

Format: string. Default value: wb

allowed values: none, wb, wt

-cacheSize – set size for L1 cache in MB (optional parameter)

Format: integer. Default value: 128

allowed values: 1 and more

-syncInterface – set sync channel IP-address from partner node

Format: string. Default value: "#p2={0}:3260"

-hbInterface – set heartbeat channel IP-address from partner node

Format: string. Default value: ""

-createTarget – set creating target

Format: string. Default value: true

Even if you do not specify the parameter -createTarget, the target will be created automatically.

If the parameter is set as -createTarget \$false, then an attempt will be made to create the device with existing targets, the names of which are specified in the -targetAlias (targets must already be created)

-bmpFolderPath – set path to save bitmap file

Format: string.

#secondary node

-imagePath2 – set path to store the device file

Format: string. Default value: "My computer\C\starwind". For Linux the following format should be used: "VSA Storage\mnt\mount_point"

-imageName2 – set name device

Format: string. Default value: masterImg21

-createImage2 – set create image file

Format: boolean. Default value: true

-targetAlias2 – set alias for targetFormat: string.
Default value: targetha22

-poolName2 – set storage pool. Do not change it and keep default value.
Format: string. Default value: pool1

-aluaOptimized2 – set Alua Optimized
Format: boolean. Default value: true

-cacheMode2 – set type L1 cache (optional parameter)
Format: string. Default value: wb
allowed values: wb, wt

-cacheSize2 – set size for L1 cache in MB (optional parameter)
Format: integer. Default value: 128
allowed values: 1 and more

-syncInterface2 – set sync channel IP-address from partner node
Format: string. Default value: “#p1={0}:3260”

-hbInterface2 – set heartbeat channel IP-address from partner node
Format: string. Default value: “”

-createTarget2 – set creating target
Format: string. Default value: true
Even if you do not specify the parameter -createTarget, the target will be created automatically. If the parameter is set as -createTarget \$false, then an attempt will be made to create the device with existing targets, the names of which are specified in the -targetAlias (targets must already be created)

-bmpFolderPath2 – set path to save bitmap file
Format: string.

IMPORTANT: If the script should be executed again with the same parameters, (for example, the first time execution has failed) make sure to do the following for one node at a time before the next attempt to execute the script:

1. Stop StarWind Service:

```
sudo systemctl stop starwind-virtual-san
```

2. Go to /opt/starwind/starwind-virtual-san/drive_c/starwind/headers and delete the headers you do not need.
3. Go to the underlying storage specified as \$imagePath and delete the header and imagefile there.
4. Go to the folder with StarWind.cfg (/opt/starwind/starwind-virtual-san/drive_c/starwind/StarWind.cfg) and copy it.
5. Edit StarWind.cfg:

```
sudo nano /opt/starwind/starwind-virtual-  
san/drive_c/starwind/StarWind.cfg
```

6. Navigate under <targets>, remove target you do not need.
7. Navigate under <devices>, and remove the device entry you do not need.
8. Start the service:

```
sudo systemctl start starwind-virtual-san
```

9. Wait for the devices synchronization.
10. Repeat for the remaining StarWind VSAN instance.

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:

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
    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.

On Proxmox 8:

```
systemctl restart multipath-tools.service
```

On Proxmox 9:

```
systemctl restart multipathd
```

11. Check if multipathing is running correctly:

```
pve# multipath -ll
```

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

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 ©2026, StarWind Software Inc. All rights reserved.