# StarWind Virtual SAN: Configuration Guide for Proxmox Virtual Environment, StarWind NVMe-oF/TCP Deployed as a Controller VM

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 NVMe-oF/TCP HA starting from V8 build 20104

# Purpose

This document outlines how to configure a Proxmox VE Cluster using StarWind Virtual SAN (VSAN), with VSAN running as a Controller Virtual Machine (CVM). The guide includes steps to prepare Proxmox VE 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 SAN: Architecture Overview

# 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

Following this guide will result in a fully configured, highly available Proxmox VE Cluster with virtual machine shared storage provided by StarWind VSAN.
NOTE: This guide applies to clusters with 2 nodes. Be sure to follow the notes within each configuration step to ensure you perform the actions appropriate for your specific cluster size.

# 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

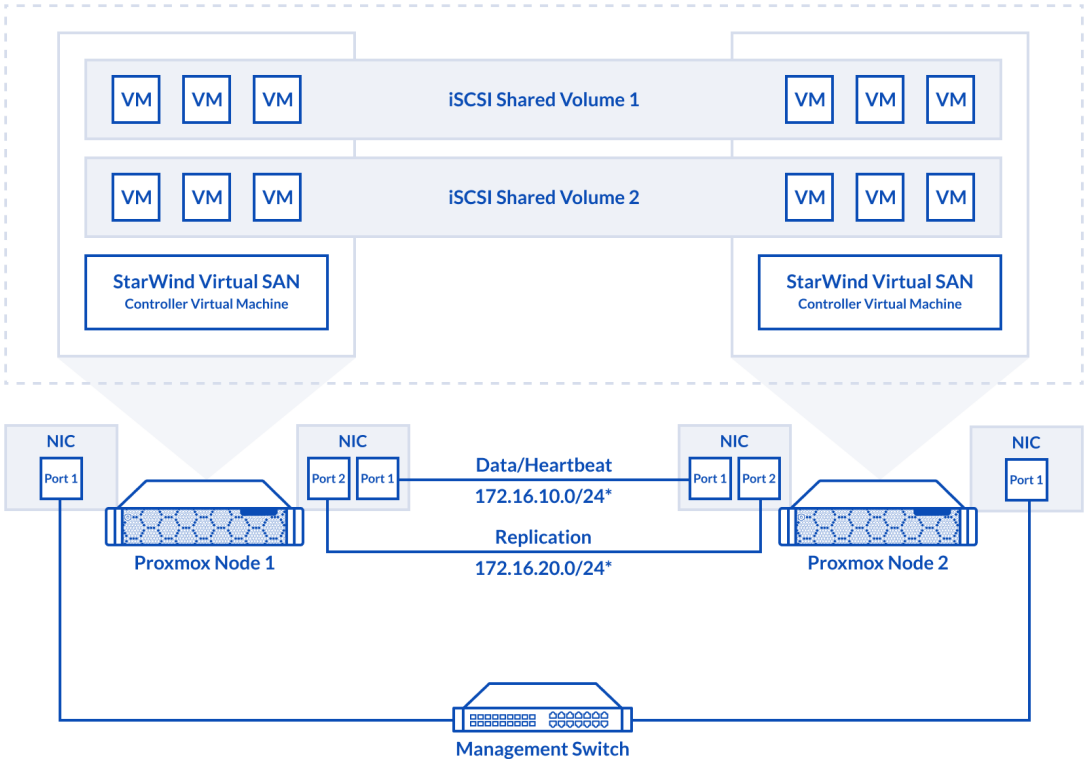
Recommended RAID settings for HDD and SSD disks:

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

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:

```
apt install corosync-qnetd
```

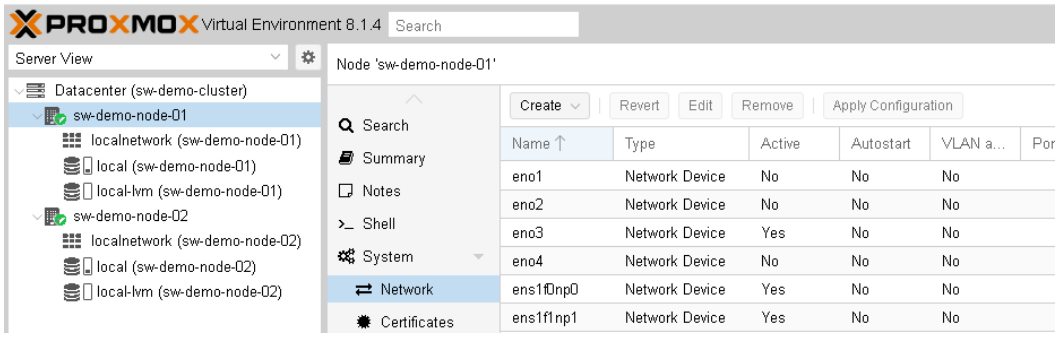
4. Install qdevice on both cluster nodes:

```
apt install corosync-qdevice
```

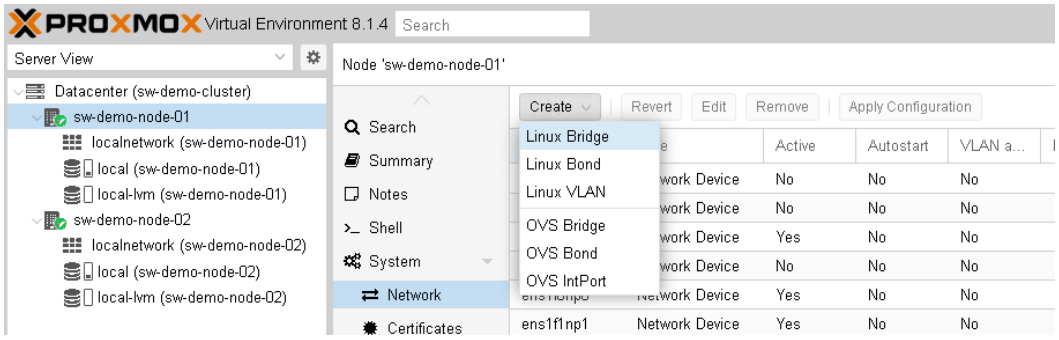
5. Configure quorum running the following command on one of the ProxMox node (change IP address)

```
pvecm qdevice setup %IP_Address_Of_Qdevice%
```

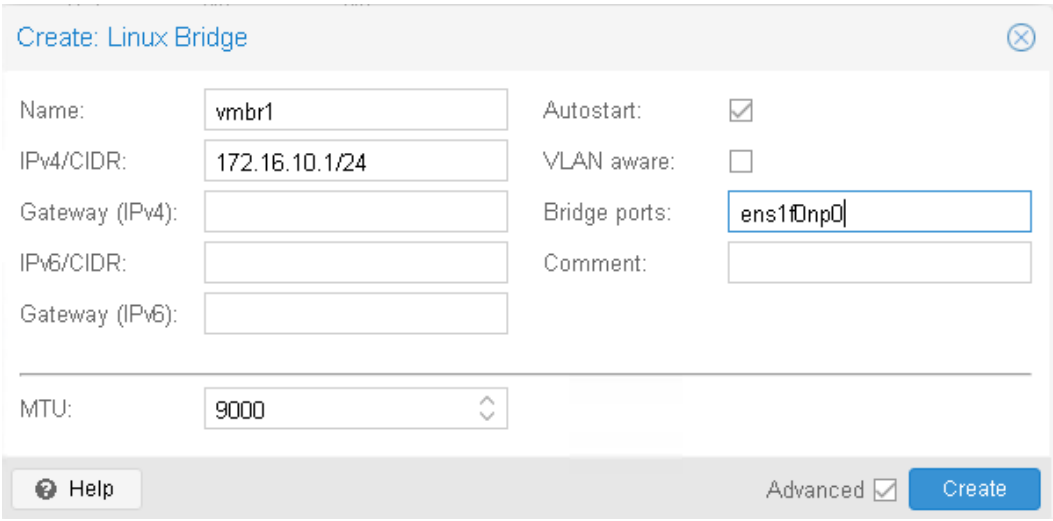
6. Configure network interfaces on each node to make sure that Synchronization and Data/StarWind heartbeat interfaces are in different subnets and connected according to the network diagram above. In this document, 172.16.10.x subnet is used for Data/StarWind heartbeat traffic, while 172.16.20.x subnet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




9. Repeat step 8 for all network adapters, which will be used for Synchronization and Data/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 passthourgh 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. Install nvmecli on both cluster nodes:

```
apt install nvme-cli -y
```

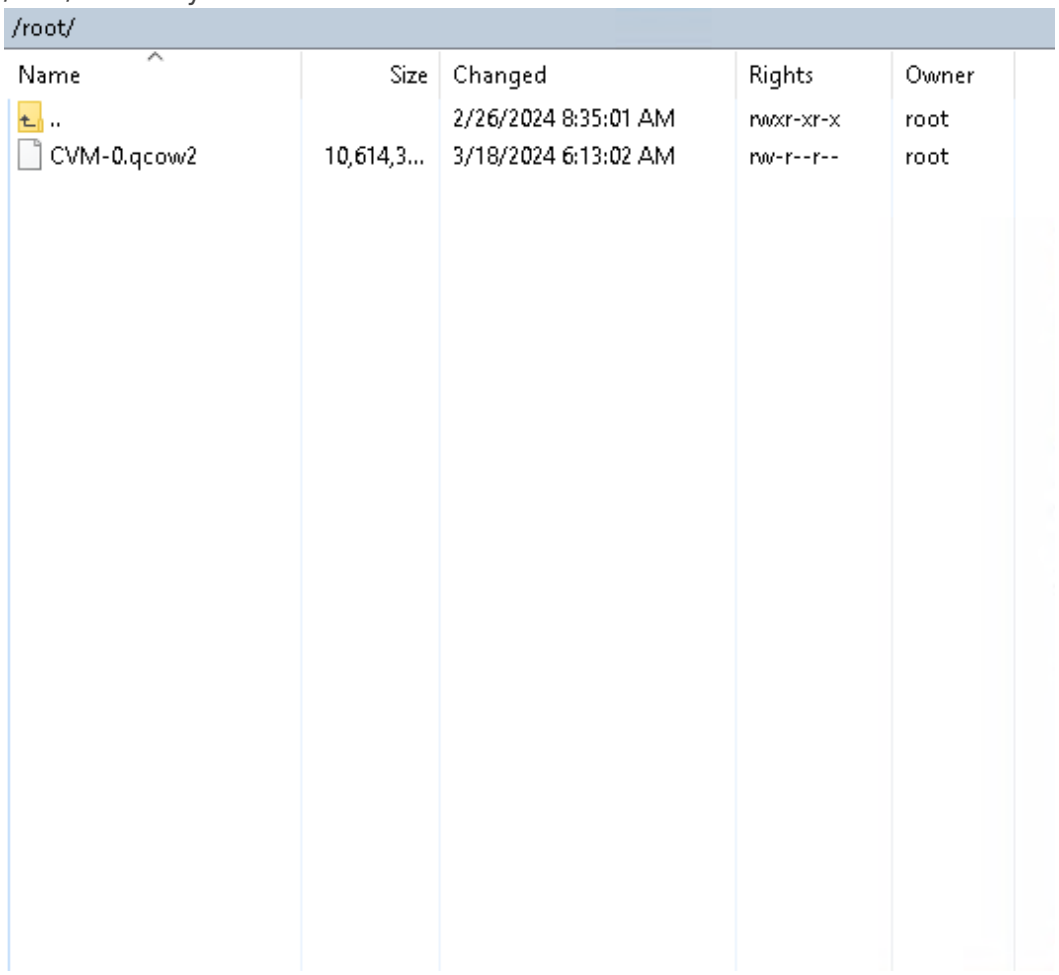
13. Reboot the host.
14. Repeat steps 6-12 an all nodes.

# Deploying Starwind Virtual San Cvm

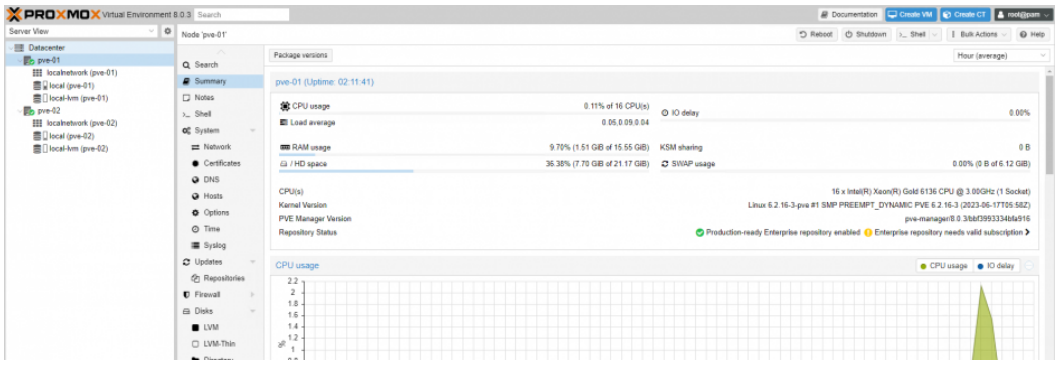
1. Download StarWind VSAN CVM KVM: VSAN by StarWind: Overview

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

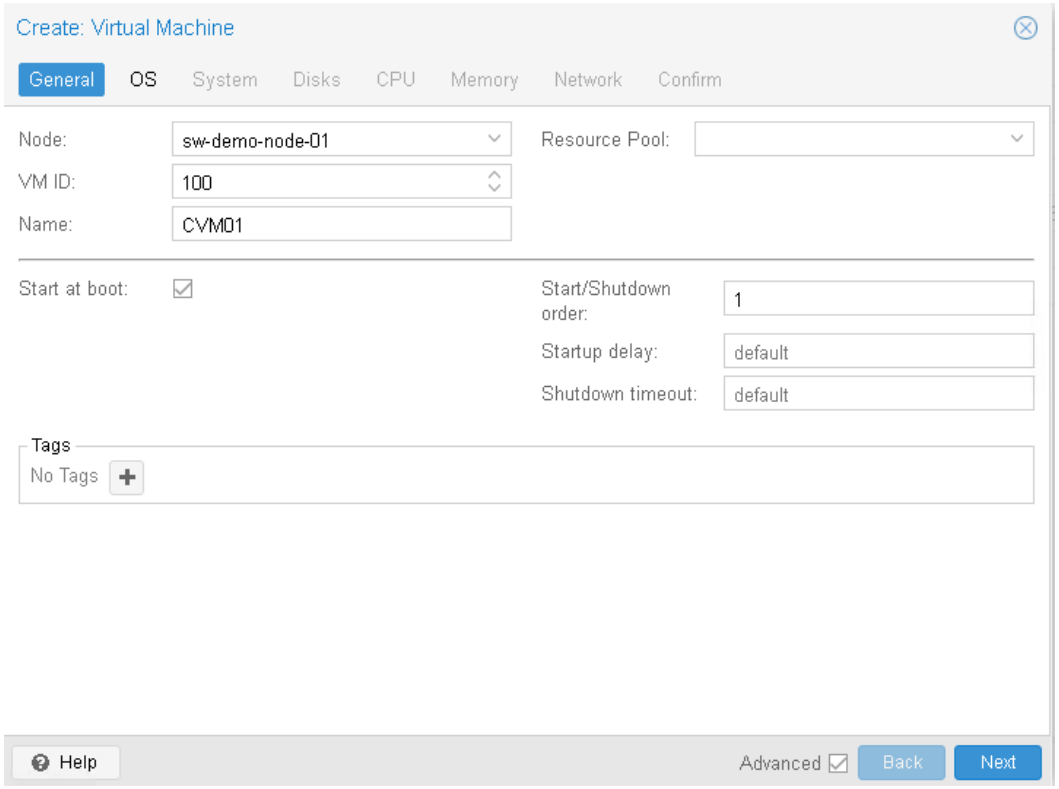
3. Upload the CVM.qcow2 file to the Proxmox Host via any SFTP client (e.g. WinSCP) to /root/ directory.



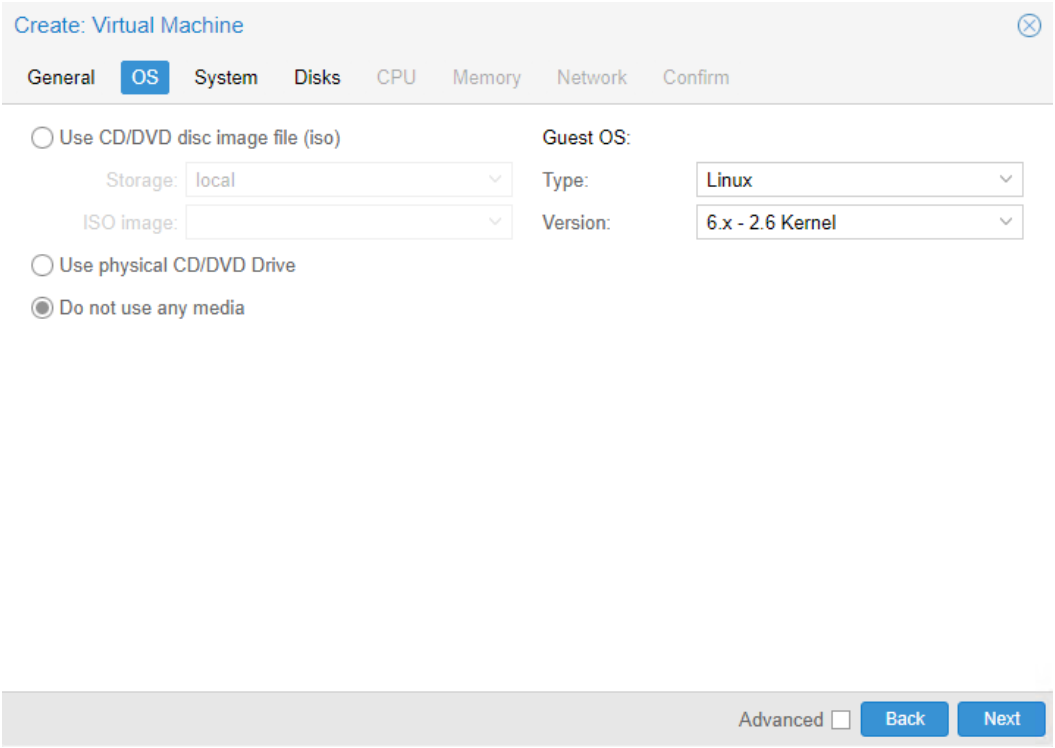

4. Create a VM without an OS. Log in to the Proxmox host via the Web GUI. Click Create VM.



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

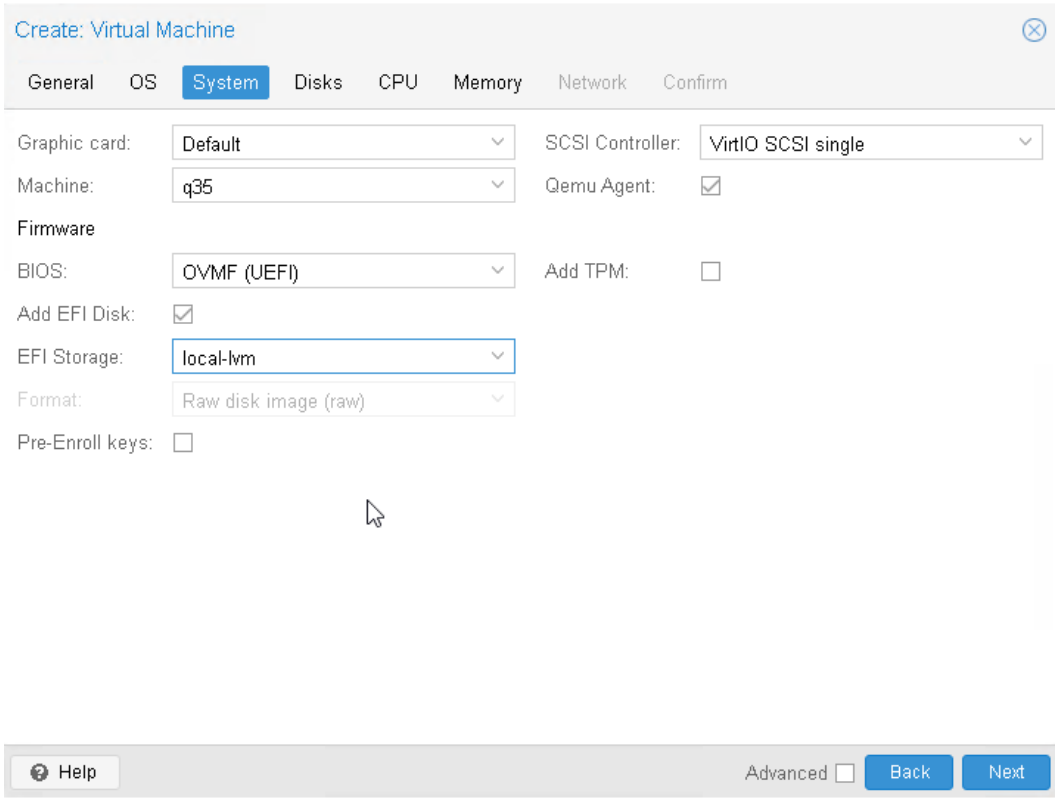



6. Choose Do not use any media and choose Guest OS Linux. Click Next.

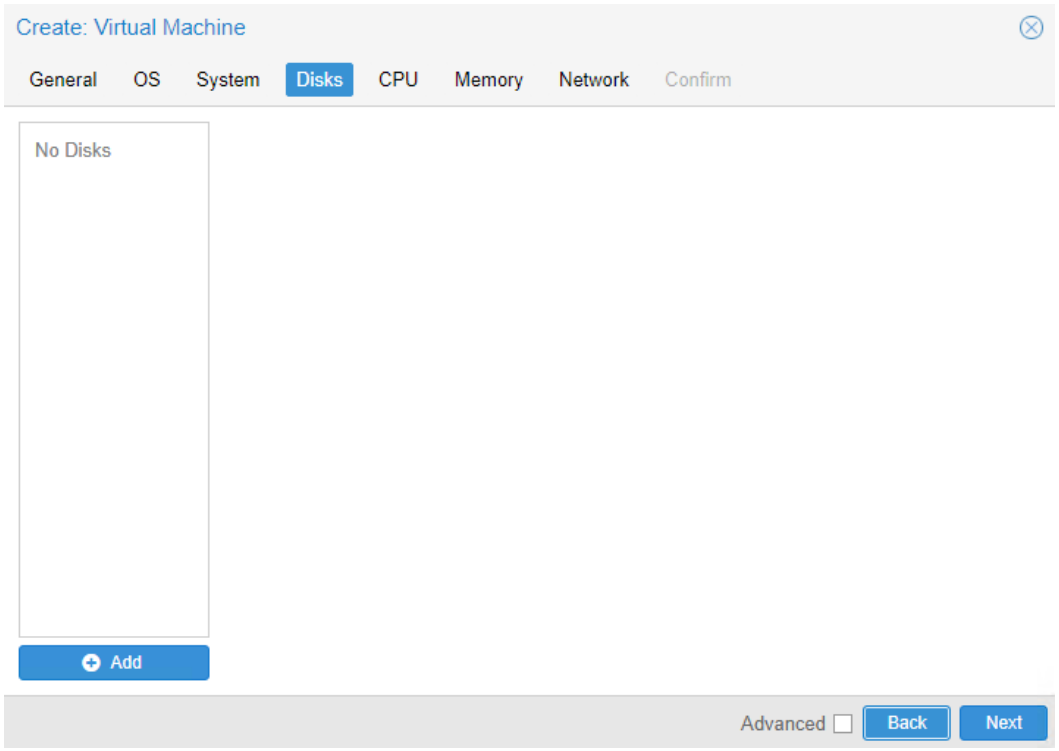



6. Specify system options. Choose Machine type q35, OVMF (UEFI), and check the Qemu Agent box. Click Next.

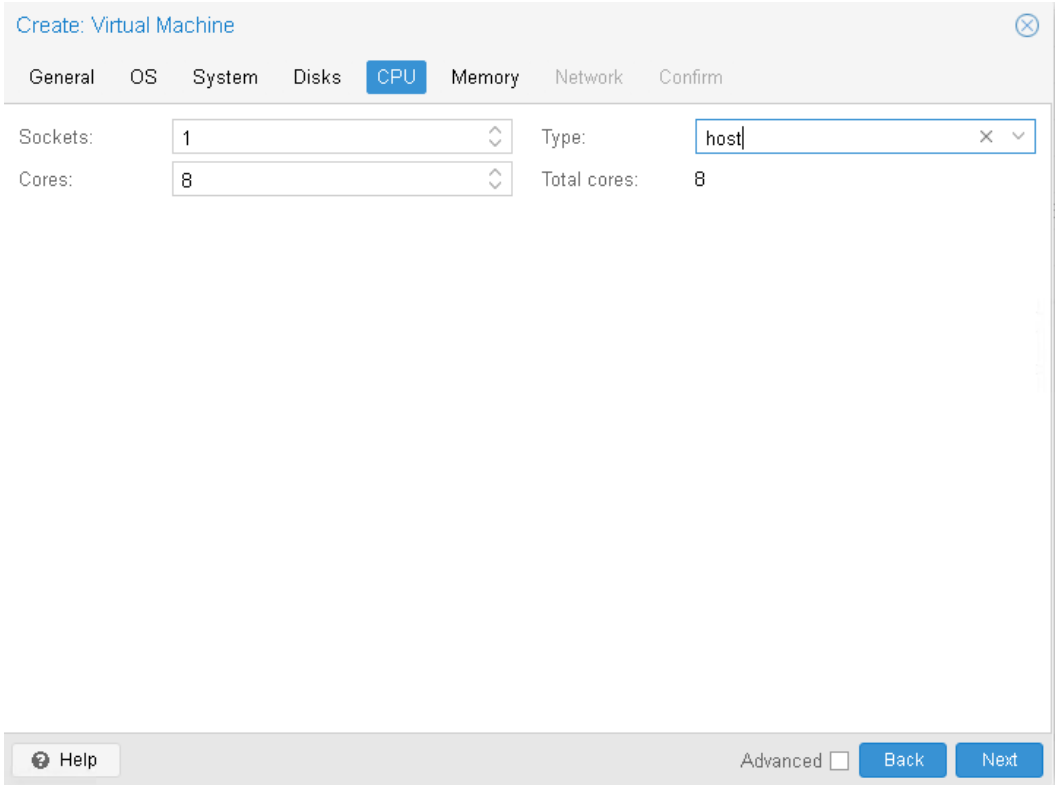
NOTE: Make sure that the “Pre-Enroll keys” checkbox is unchecked.



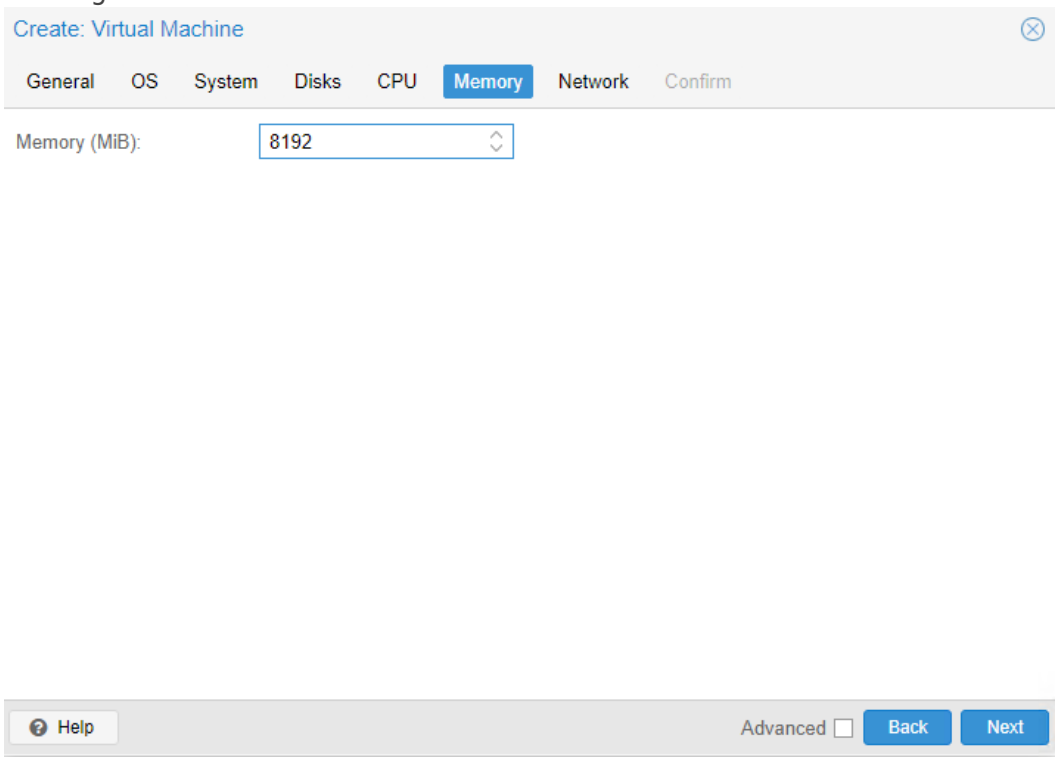

7. Remove all disks from the VM. Click Next.



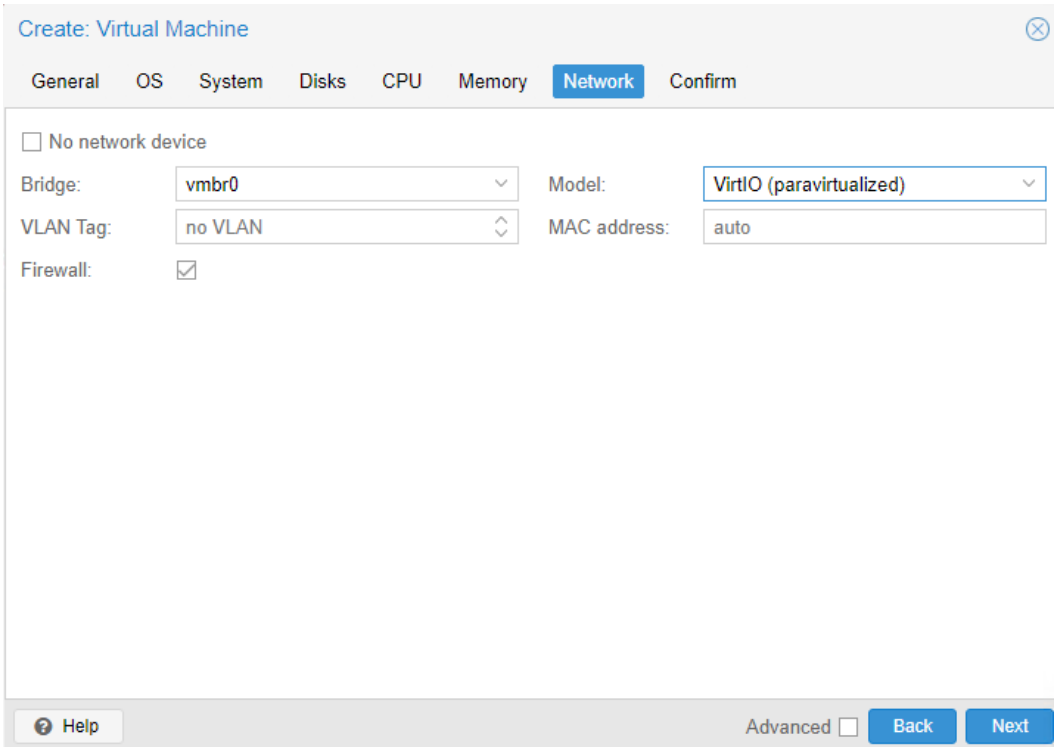

8. Assign 8 cores to the VM and choose Host CPU type. Click Next.



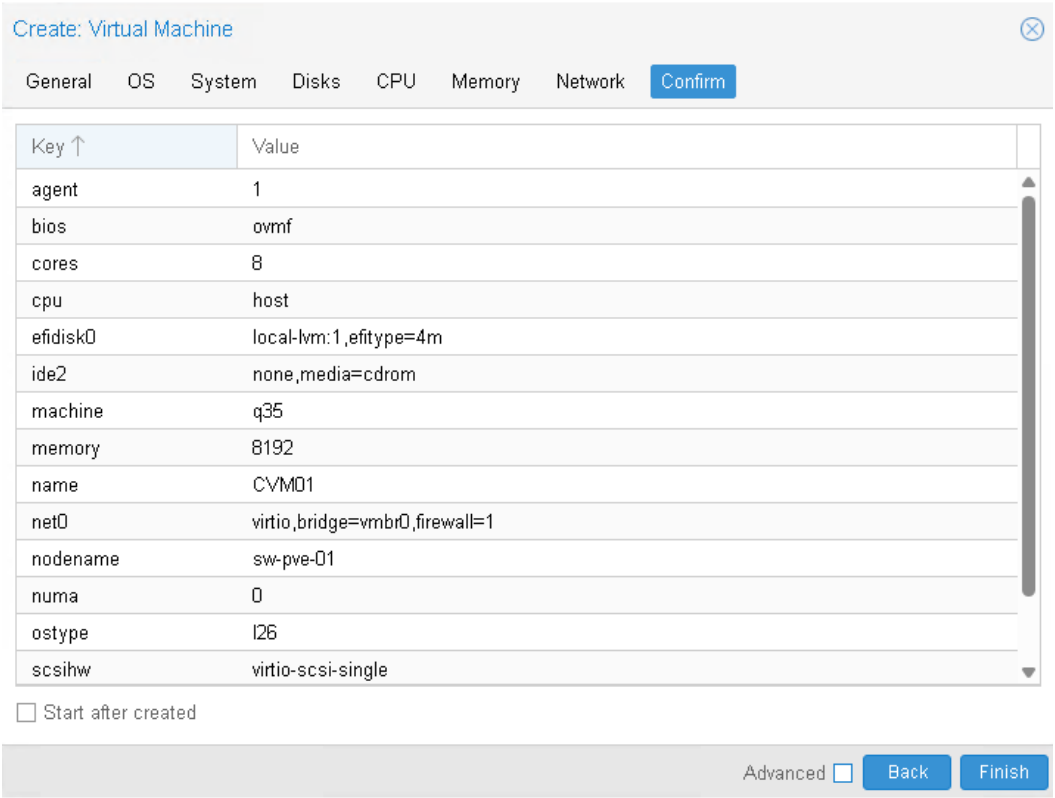

9. Assign at least 8GB of RAM to the VM. Click Next.




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




11. Confirm settings. Click Finish.

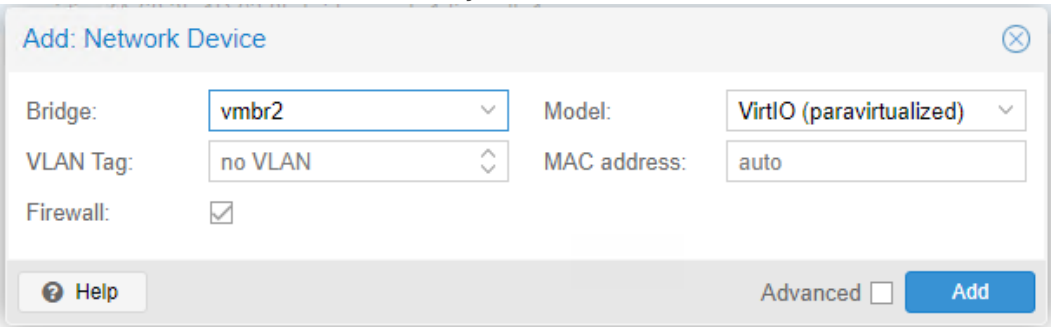



12. Connect to the Proxmox VE host via SSH. Attach the CVM.qcow2 file to the VM.

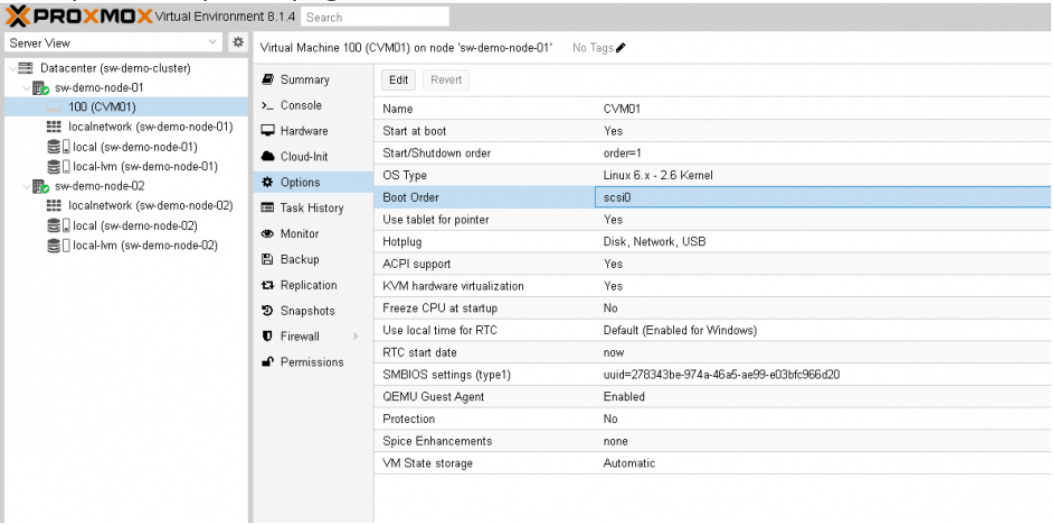
```
qm importdisk 100 /root/CVM.qcow2 local-lvm
```

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

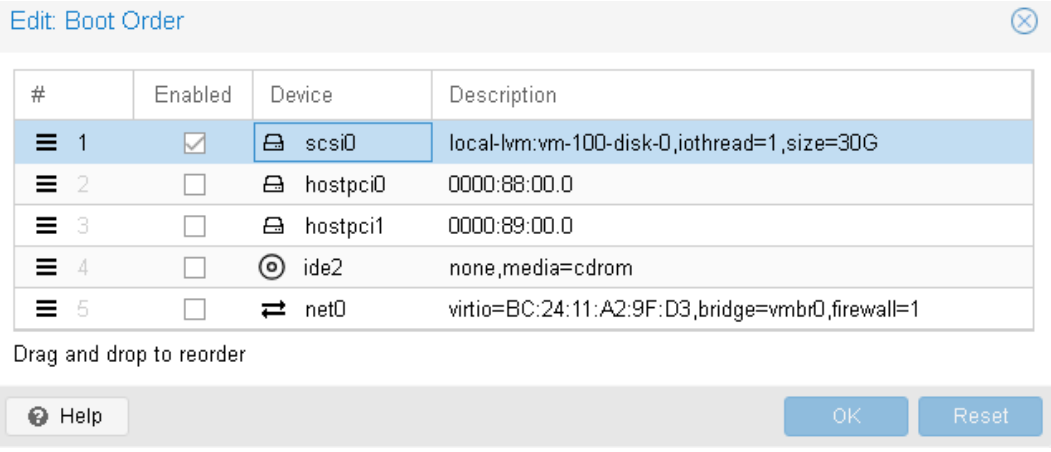
14. Attach Network interfaces for Synchronization and Data/Heartbeat traffic.




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




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



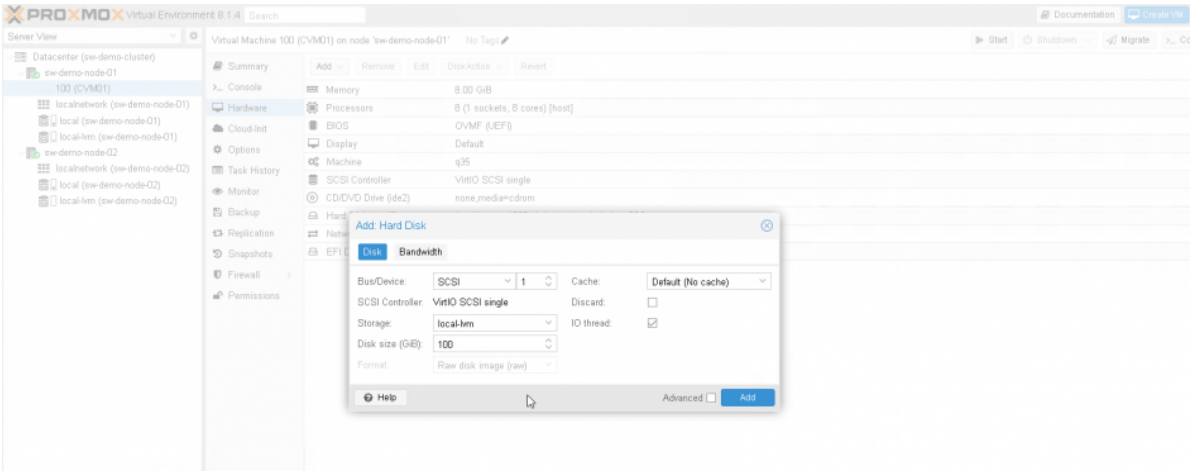

17. Repeat all the steps from this section on other Proxmox hosts.

# Attaching Storage To Starwind Virtual San Cvm

Please follow the steps below to attach desired storage type to the CVM

# Attaching Virtual Disk To Starwind Virtual San Cvm

1. Open the VM Hardware page in Proxmox and add a drive to the VM, which will be used by the StarWind service. Specify size of the Virtual disk and click OK.




Note. It is recommended to use the VirtIO SCSI single controller for better performance. If multiple virtual disks are needed to be used in a software RAID inside the CVM, the VirtIO SCSI controller should be used.

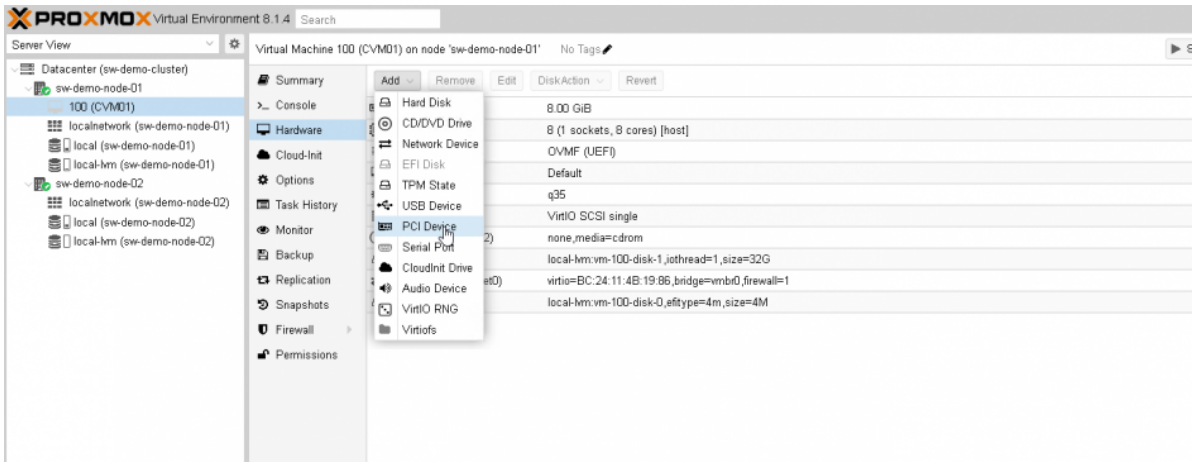
2. Repeat step 1 to attach additional Virtual Disks.

3. Start the VM.

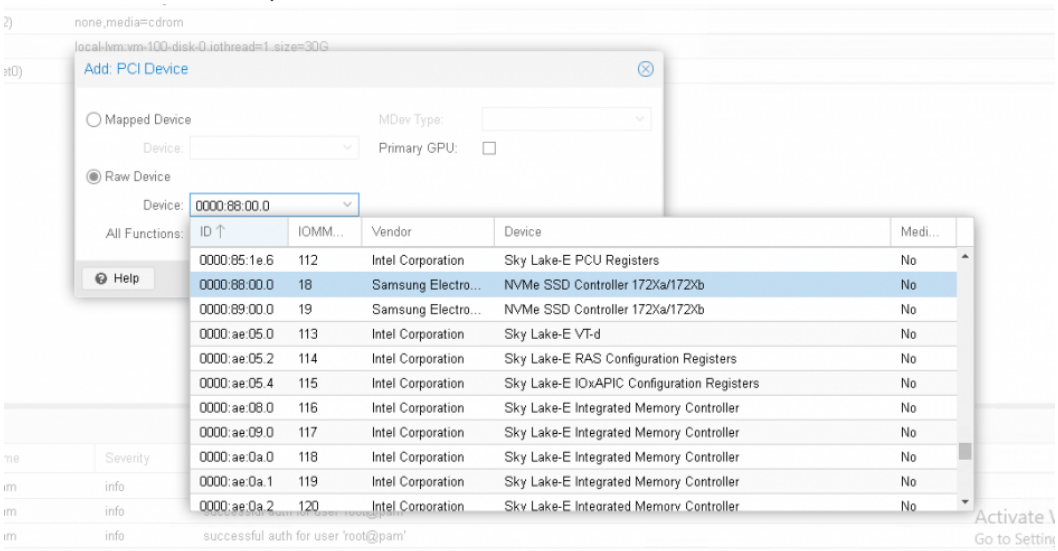
4. Repeat steps 1-2 on all nodes.

# Attaching Pcie Device To Starwind Virtual San Cvm

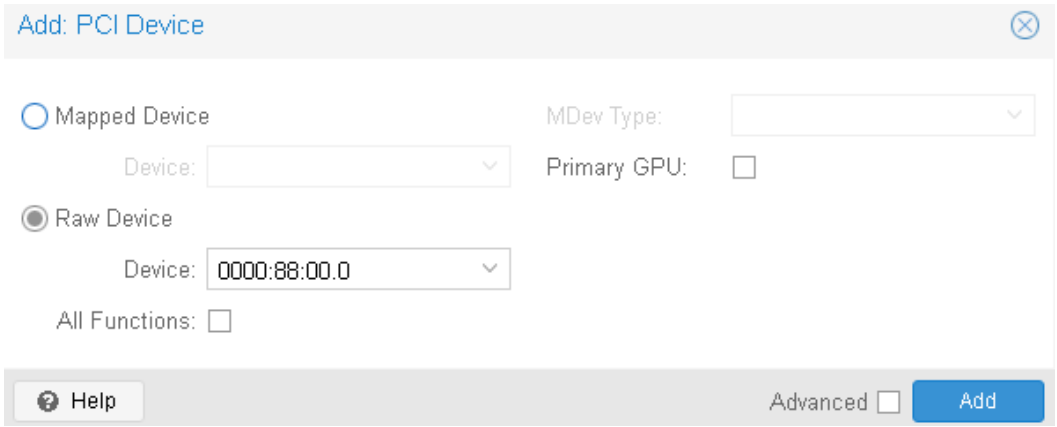
1. Shut down the StarWind VSAN CVM.



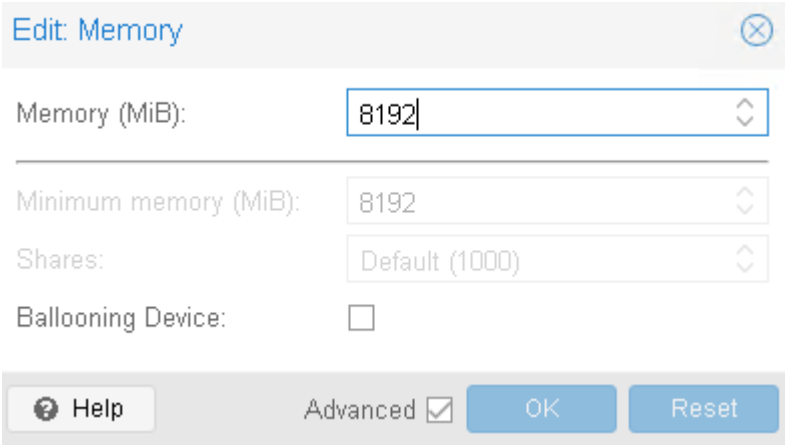

2. Open the VM Hardware page in Proxmox and click Add -> PCI Device.3. Choose PCIe Device from the drop-down list.




4. Click Add.




5. Edit Memory. Uncheck Ballooning Device. Click OK.




6. Start VM.

7. Repeat steps 1-6 on all nodes.

## Initial Configuration Wizard

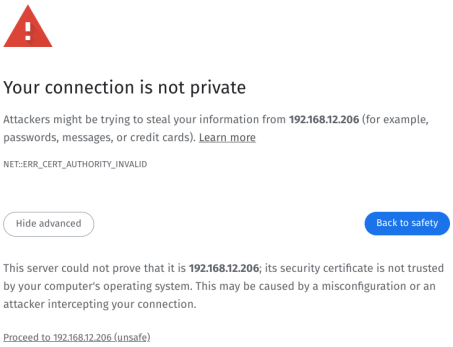
1. Start the StarWind Virtual SAN Controller Virtual Machine.

2. Launch the VM console to view the VM boot process and obtain the IPv4 address of the Management network interface.

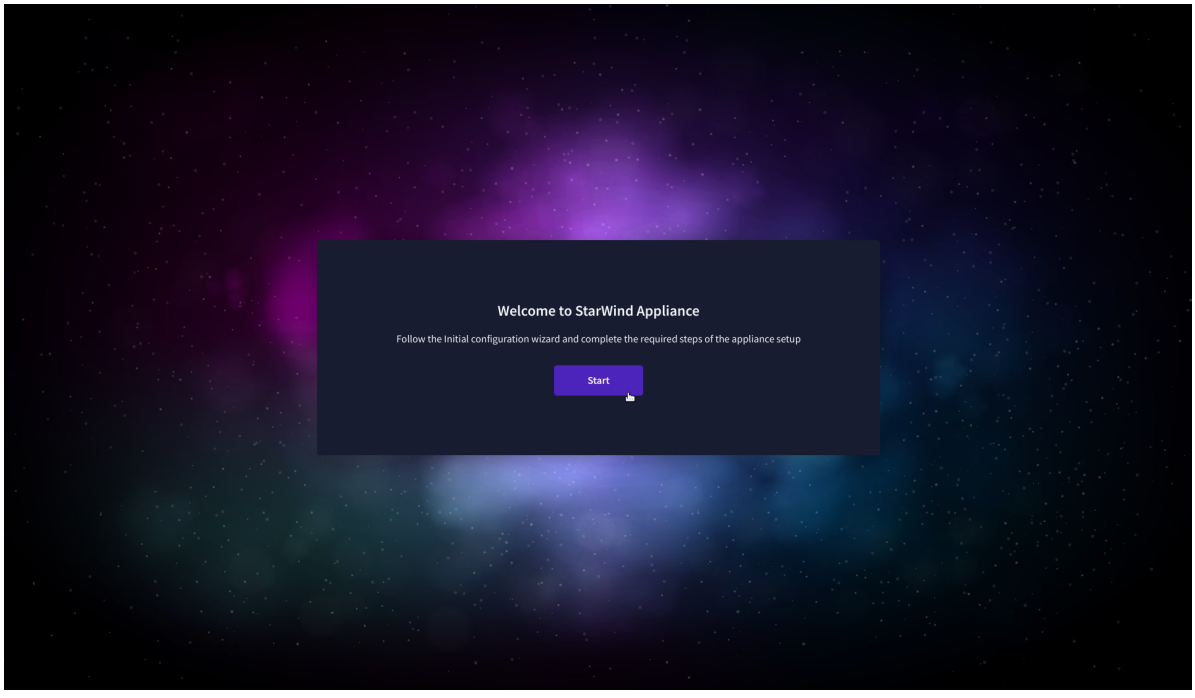
NOTE: If the VM does not acquire an IPv4 address from a DHCP server, use the Text-based User Interface (TUI) to set up the Management network manually.

Default credentials for TUI: user/rds123RDS

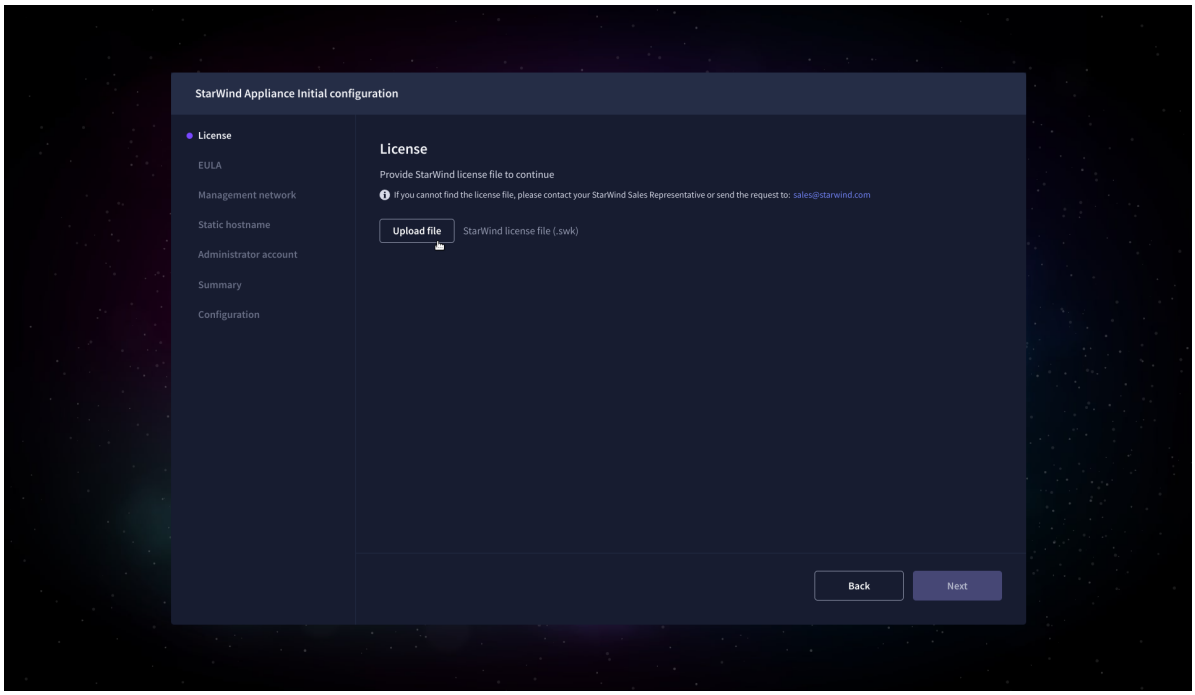
3. Using a web browser, open a new tab and enter the VM's IPv4 address to access the StarWind VSAN Web Interface. On the Your connection is not private screen, click Advanced and then select Continue to...



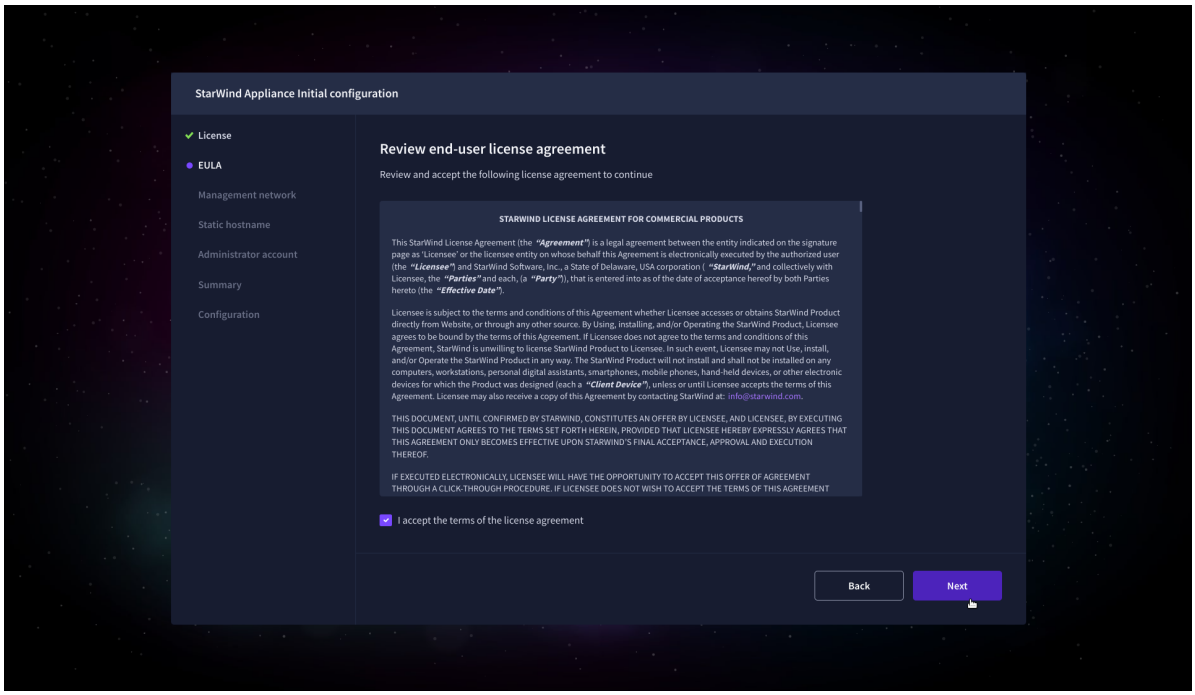

4. On the Welcome to StarWind Appliance screen, click Start to launch the Initial Configuration Wizard.




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

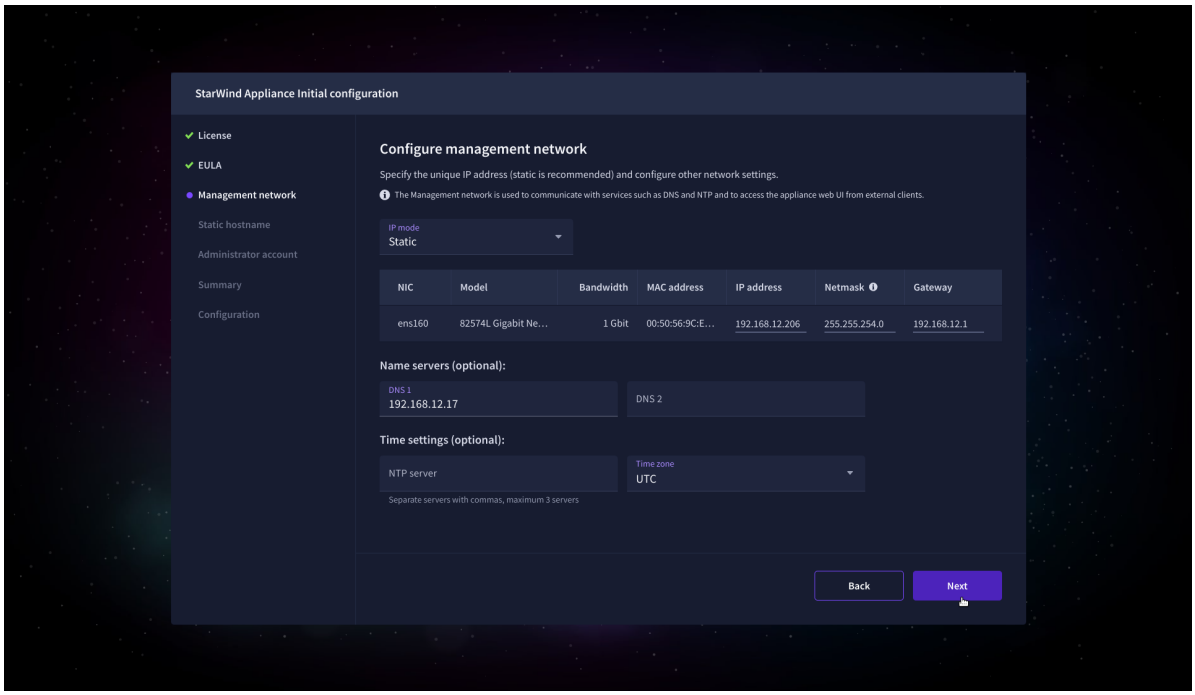



6. On the EULA step, read and accept the End User License Agreement to continue.

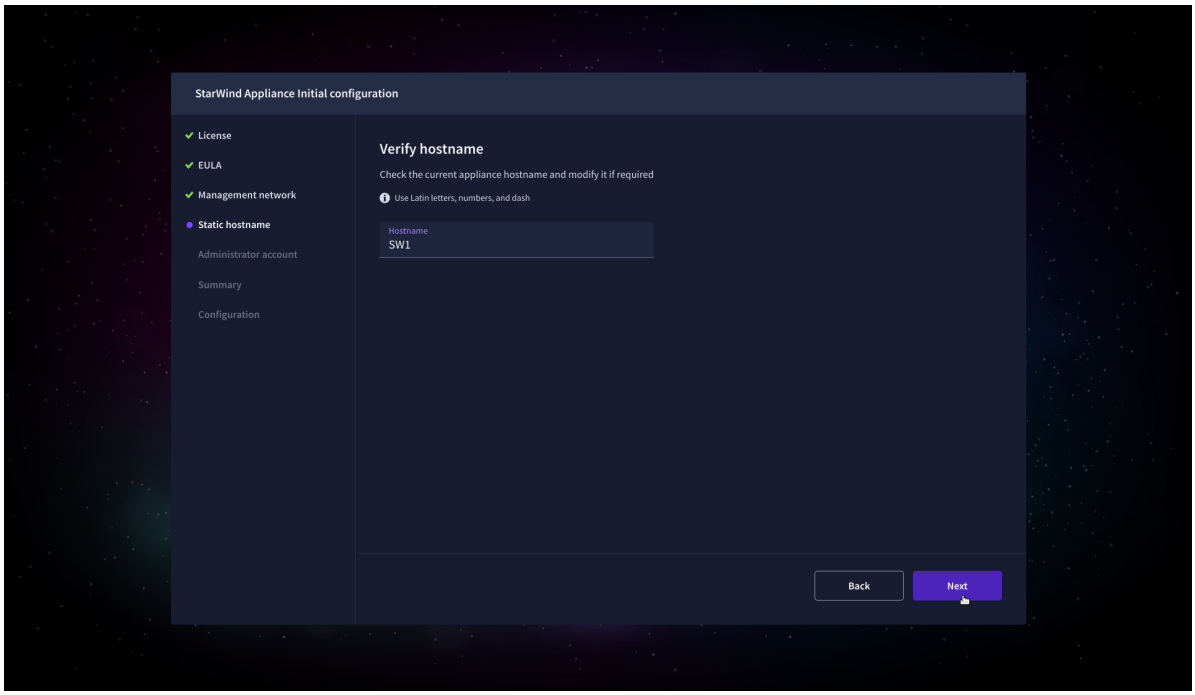



7. On the Management network step, review or edit the network settings and click Next.

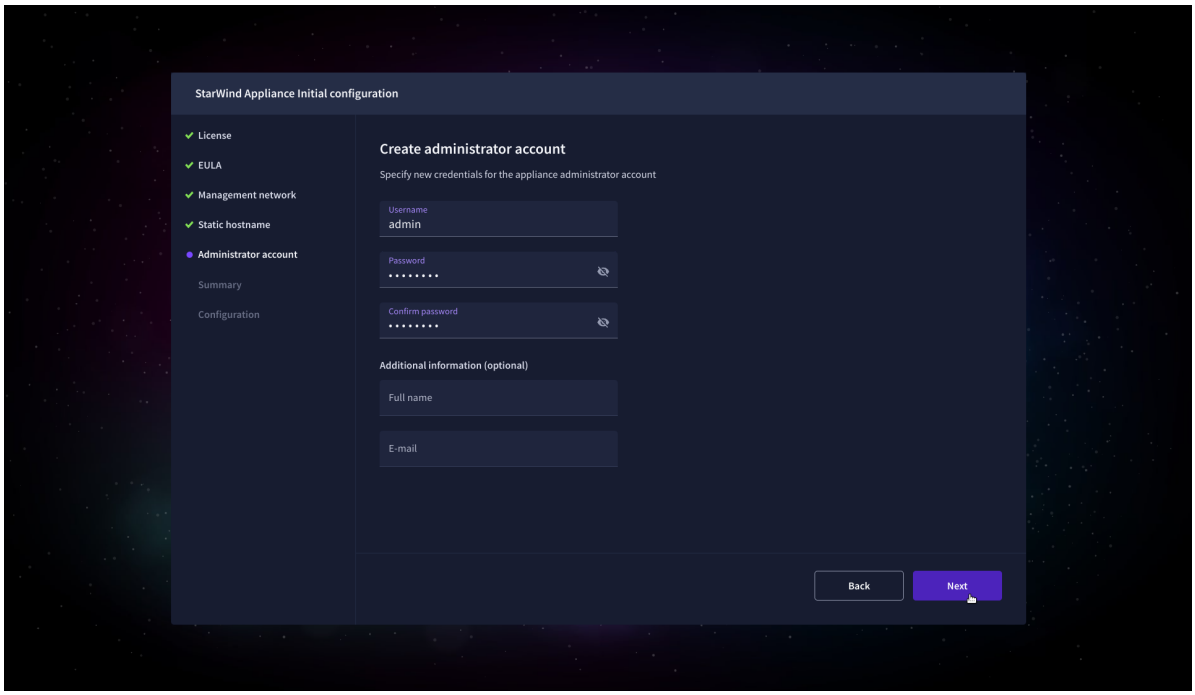
IMPORTANT: The use of Static IP mode is highly recommended.



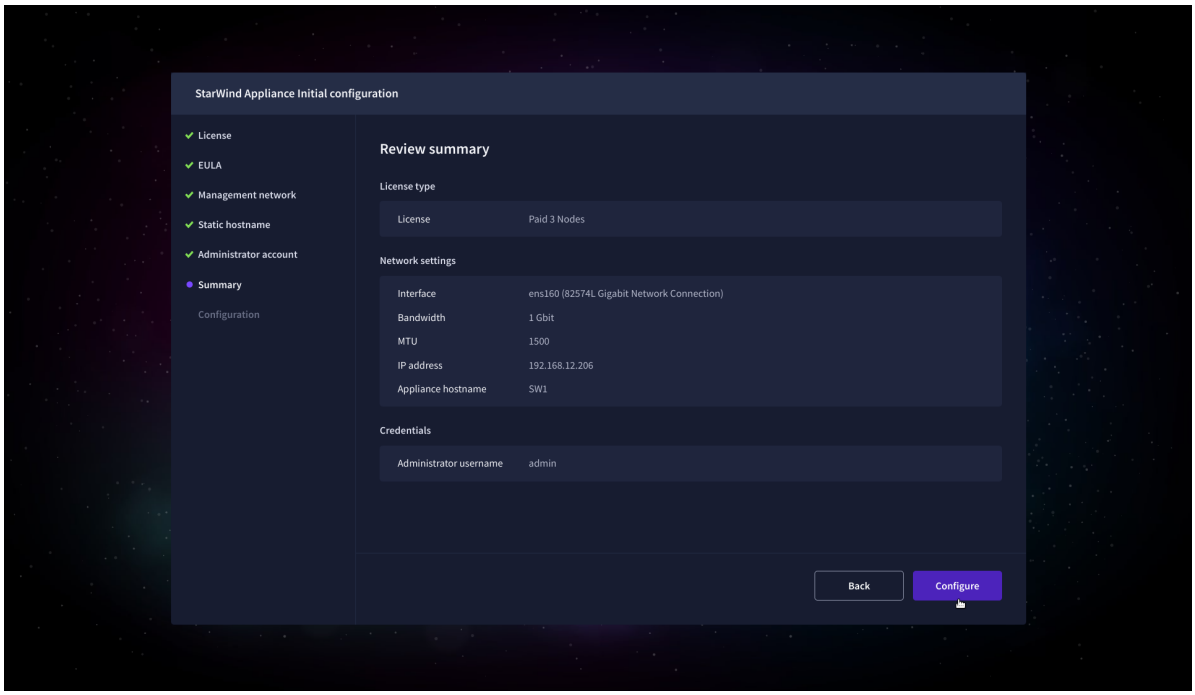

8. On the Static hostname, specify the hostname for the virtual machine and click Next.



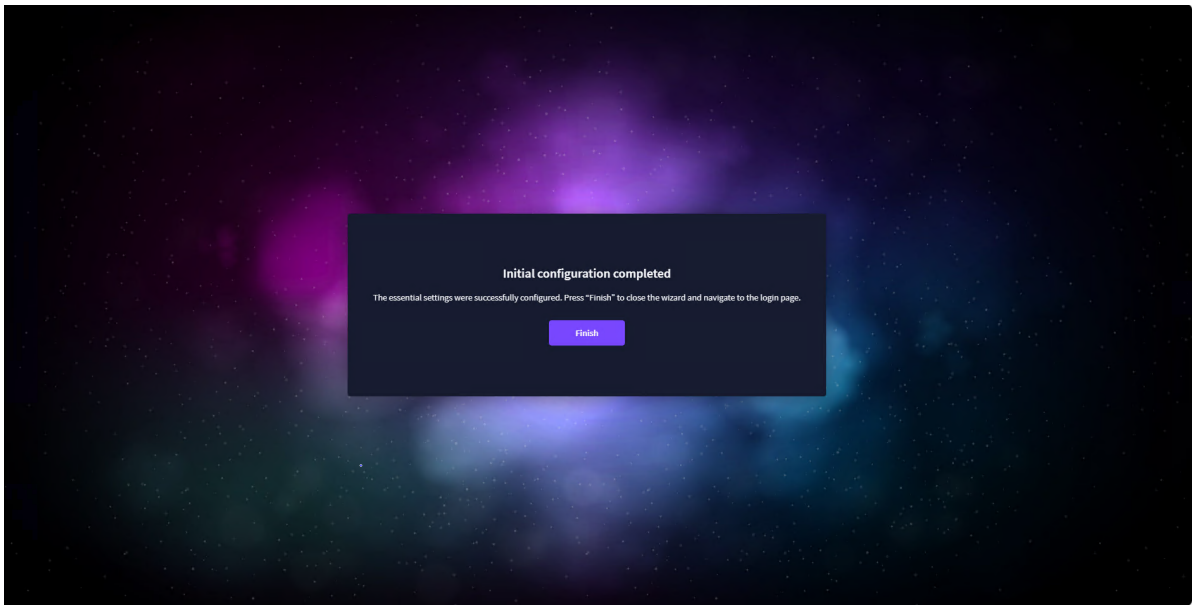

9. On the Administrator account step, specify the credentials for the new StarWind Virtual SAN administrator account and click Next.




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




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



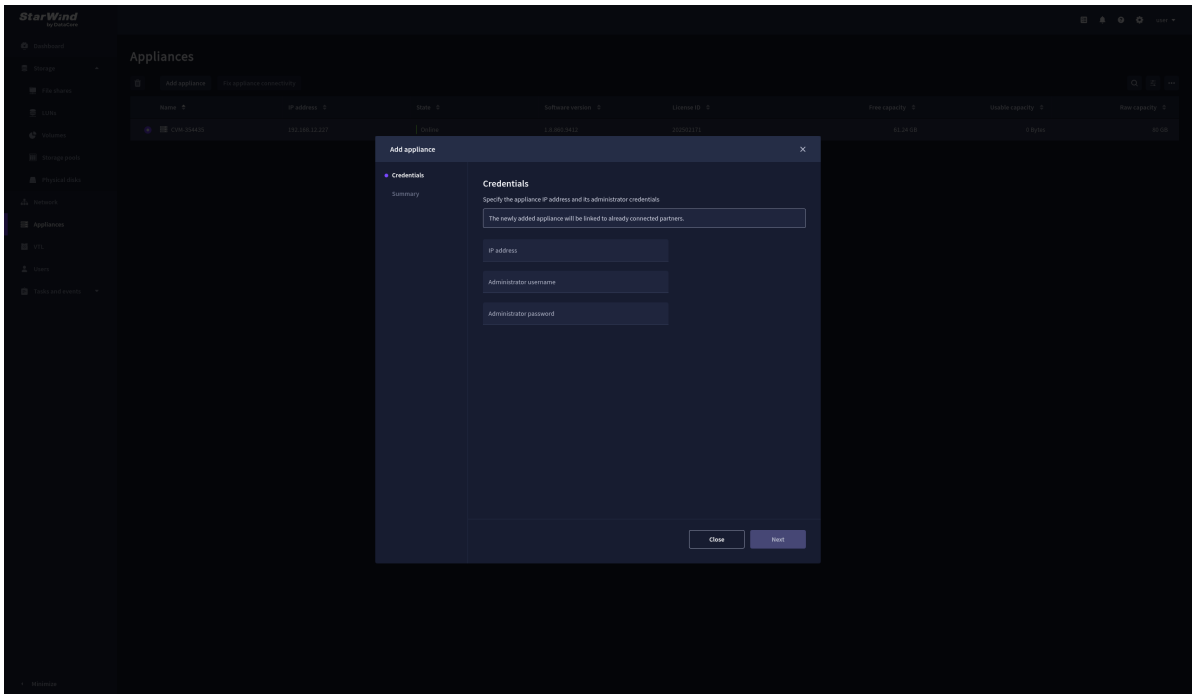

12. Repeat steps 1 through 12 on each Proxmox host.

## Add Appliance

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

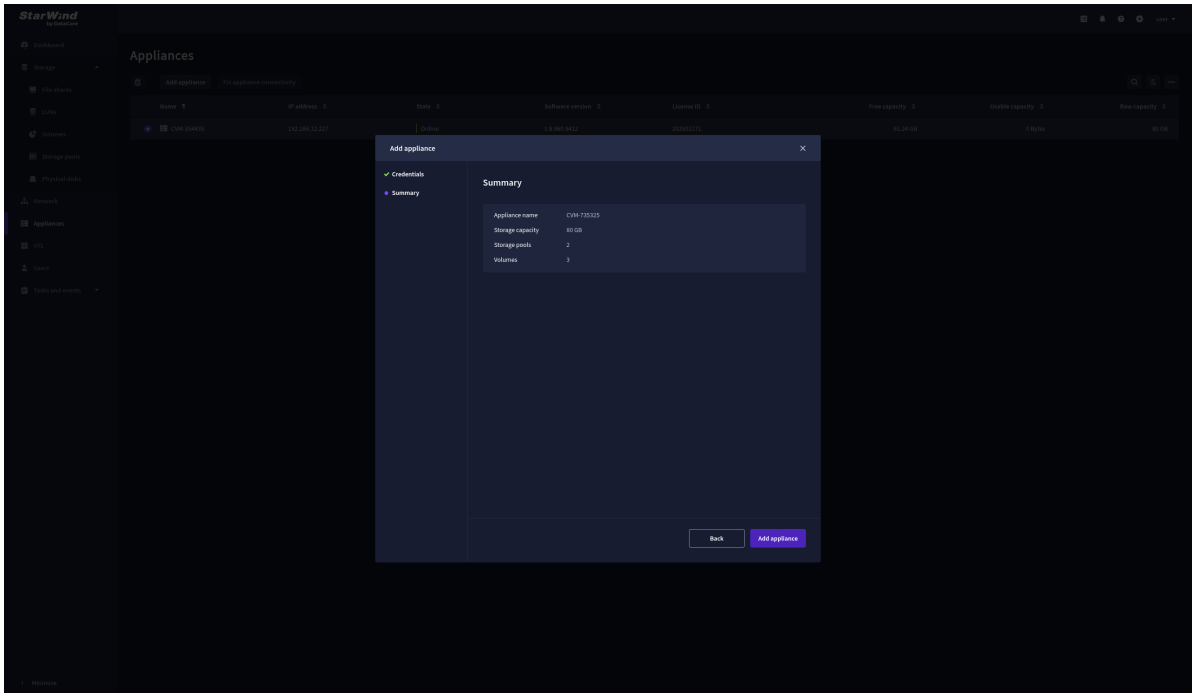
1. Navigate to the Appliances page and click Add to open the Add appliance wizard.

2. On the Credentials step, enter the IP address and credentials for the partner StarWind Virtual SAN appliance, then click Next.




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

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

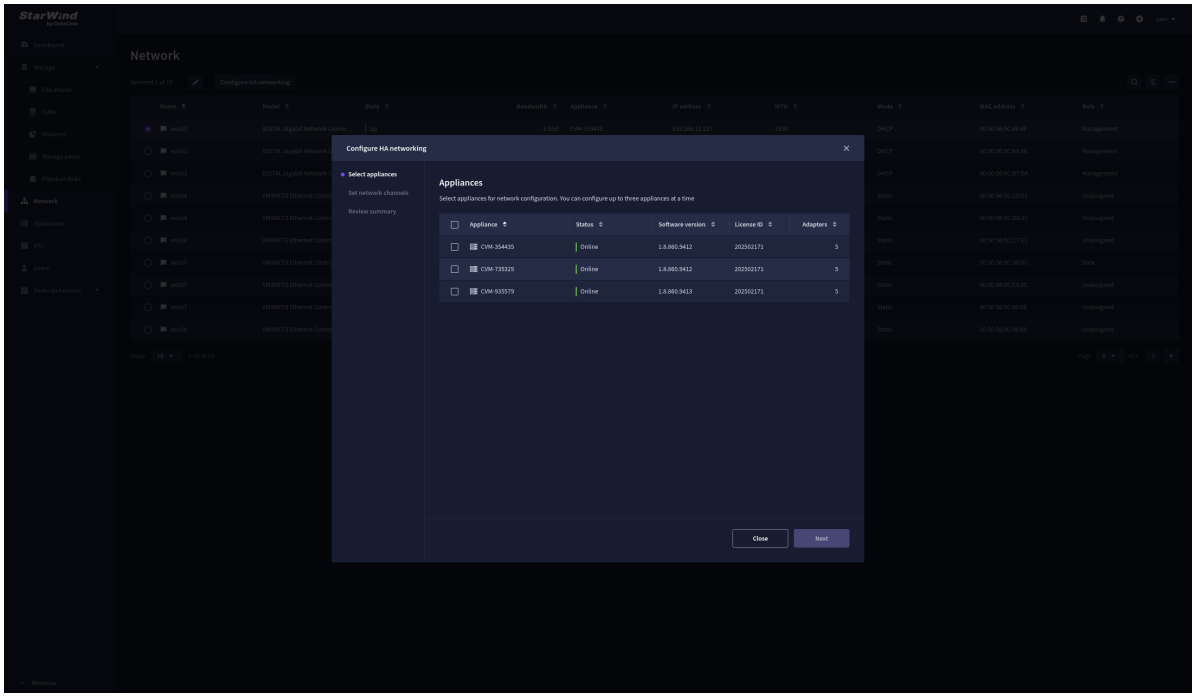


# Configure Ha Networking

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

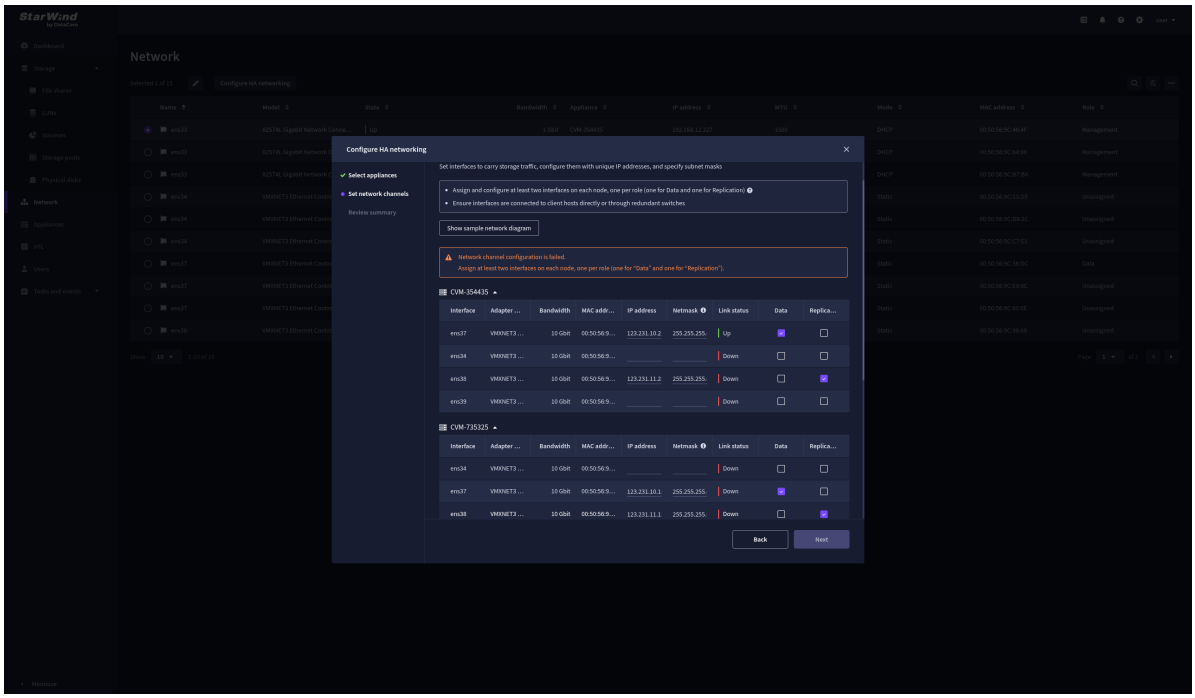
2. On the Appliances step, select either 2 partner appliances to configure two-way replication or 3 appliances for three-way replication, then click Next.

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




3. Set Network Channels step combines configuration of Data and Replication Networks. Select the network interfaces designated to carry iSCSI or NVMe-oF storage traffic. 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.

4. 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.

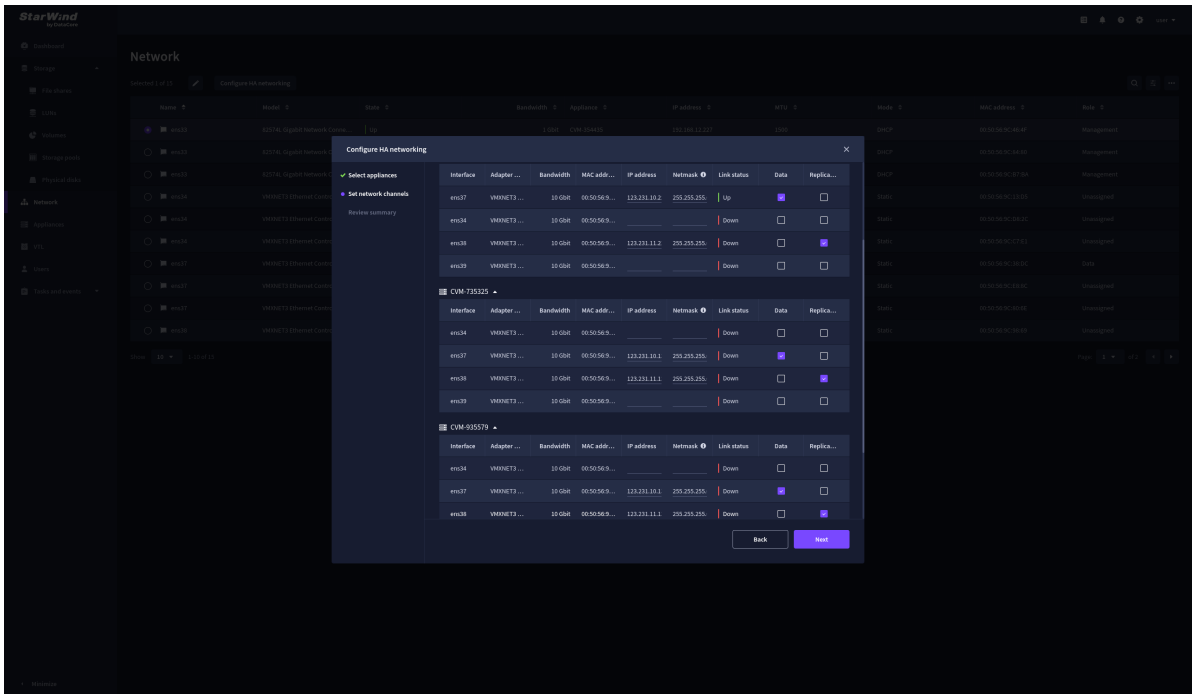



IMPORTANT: For a redundant, high-availability configuration, configure at least 2 network interfaces on each appliance. Ensure that the Data Network interfaces are interconnected between appliances through multiple direct links or via redundant switches.

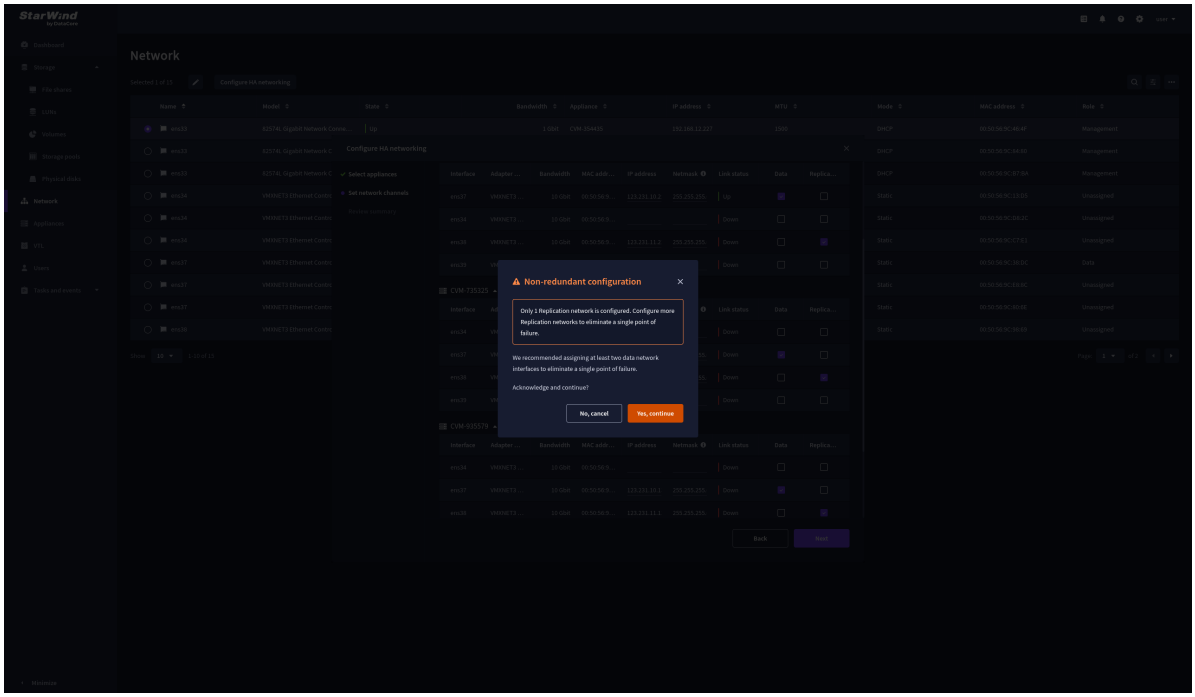
5. Assign an MTU value on all selected network adapters, e.g., 1500 or 9000 bytes. If you are using network switches with the selected Data Network adapters, ensure that they are configured with the same MTU size value. In case of an MTU settings mismatch, stability and performance issues might occur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

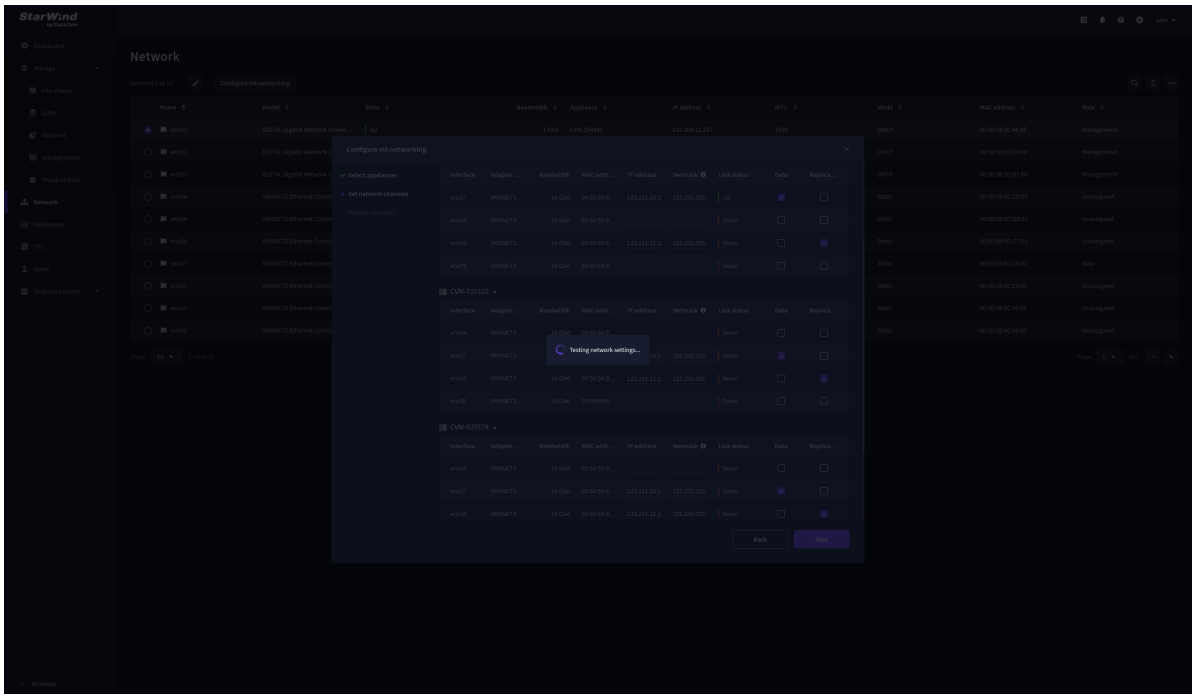
6. Once configured, click Next to validate network settings.



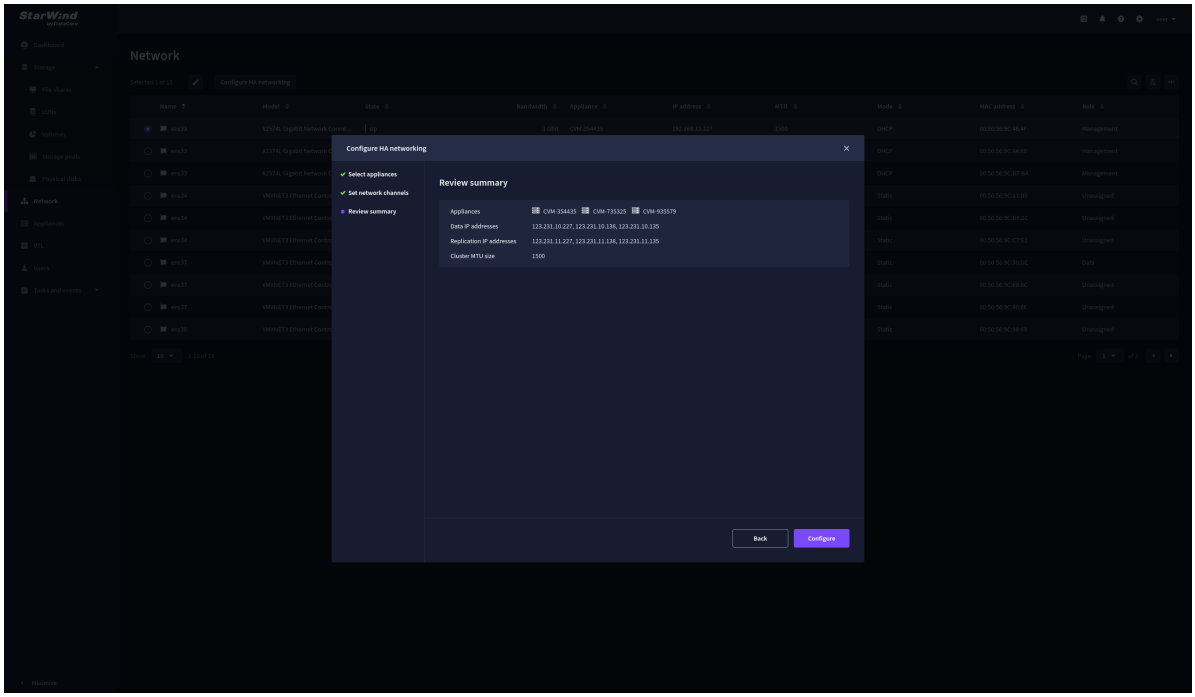

7. The warning might appear if a single data interface is configured. Click Yes, continue to proceed with the configuration.




8. Wait for the configuration to complete.




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



# Add Physical Disks

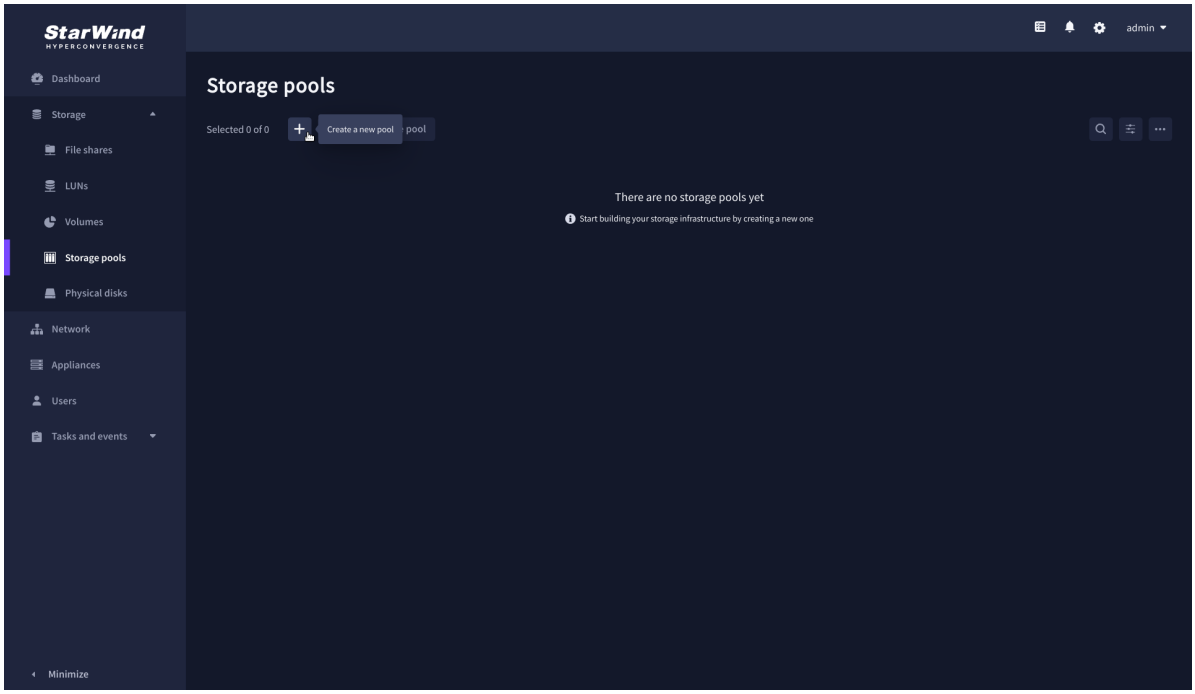
Attach physical storage to StarWind Virtual SAN Controller VM:

- 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.

Also, find the storage provisioning guidelines in the KB article.

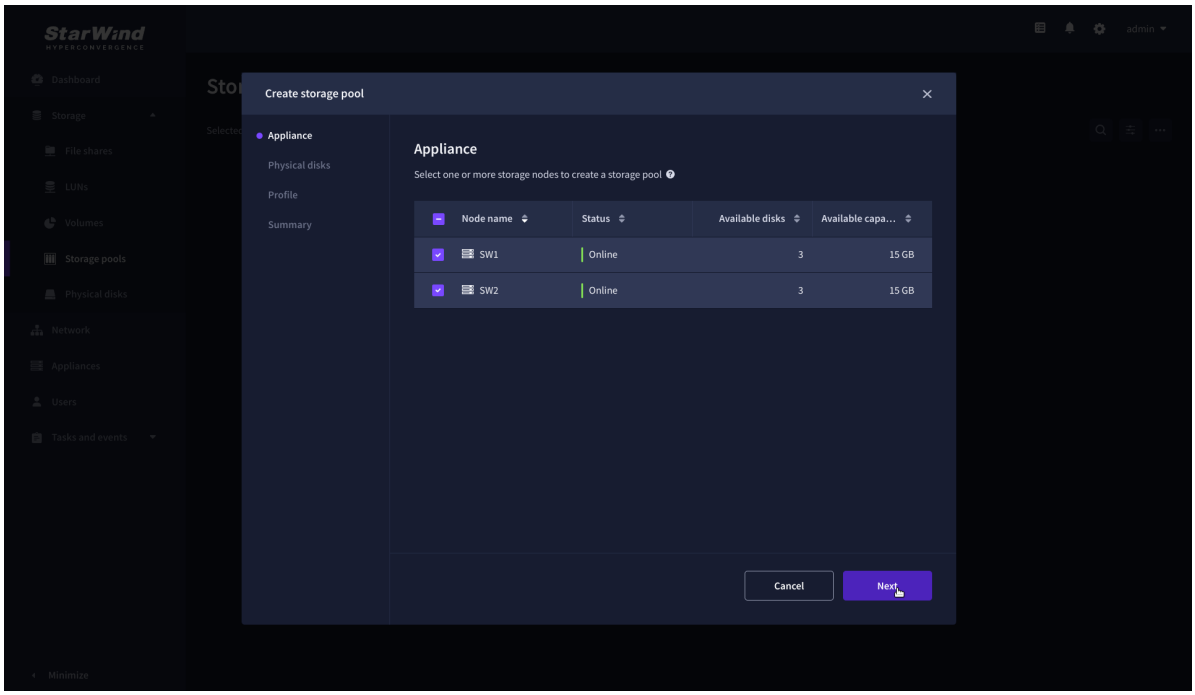
# Create Storage Pool

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



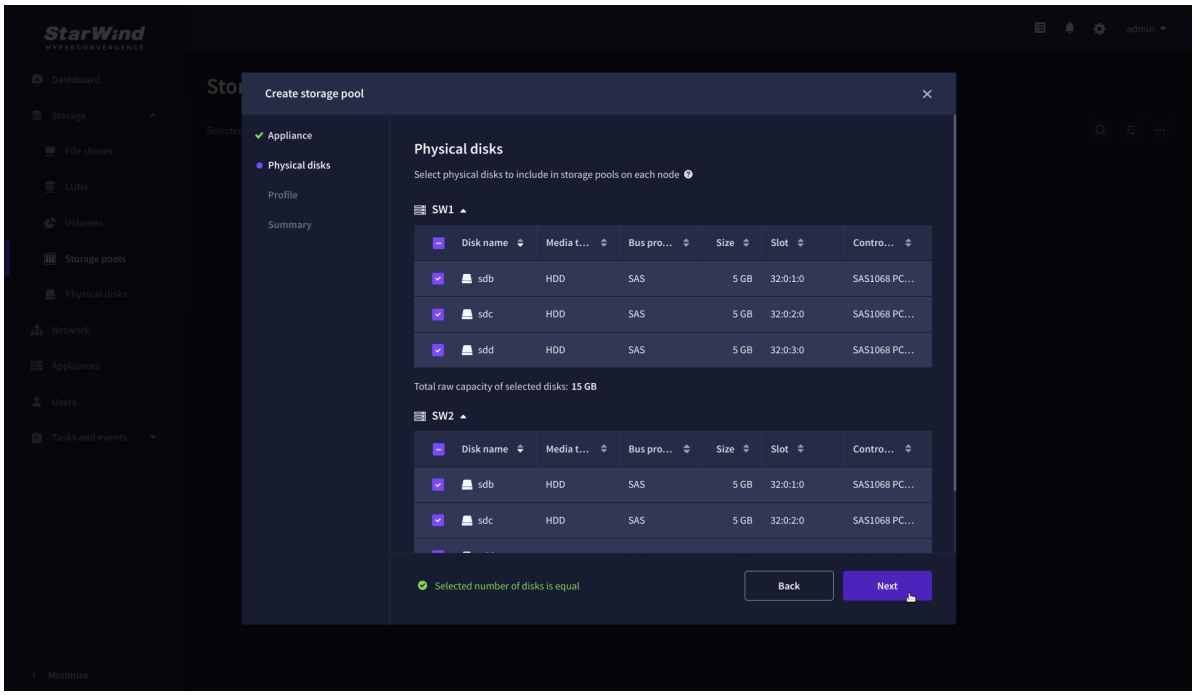

2. On the Appliance step, select partner appliances on which to create new storage pools, then click Next.

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.

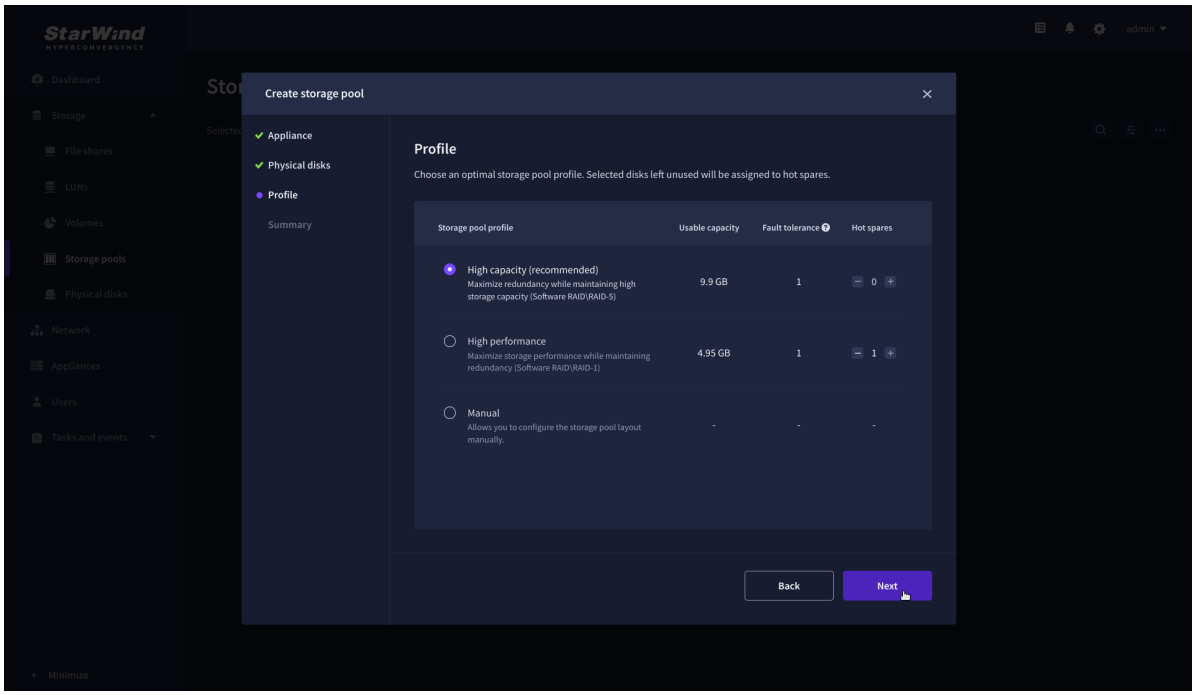



3. On the Physical disks step, select physical disks to be pooled on each node, then click Next.

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



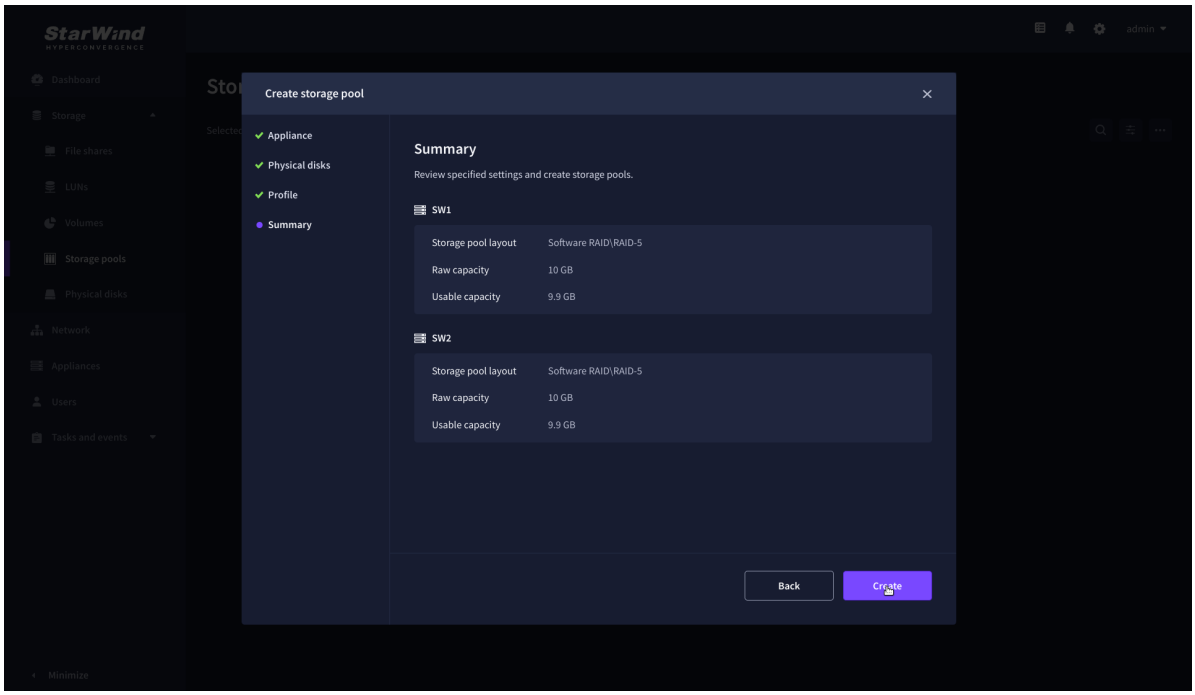
4. On the Profile step, select one of the preconfigured storage profiles, or choose Manual to configure the storage pool manually based on your redundancy, capacity, and performance requirements, then click Next.




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.

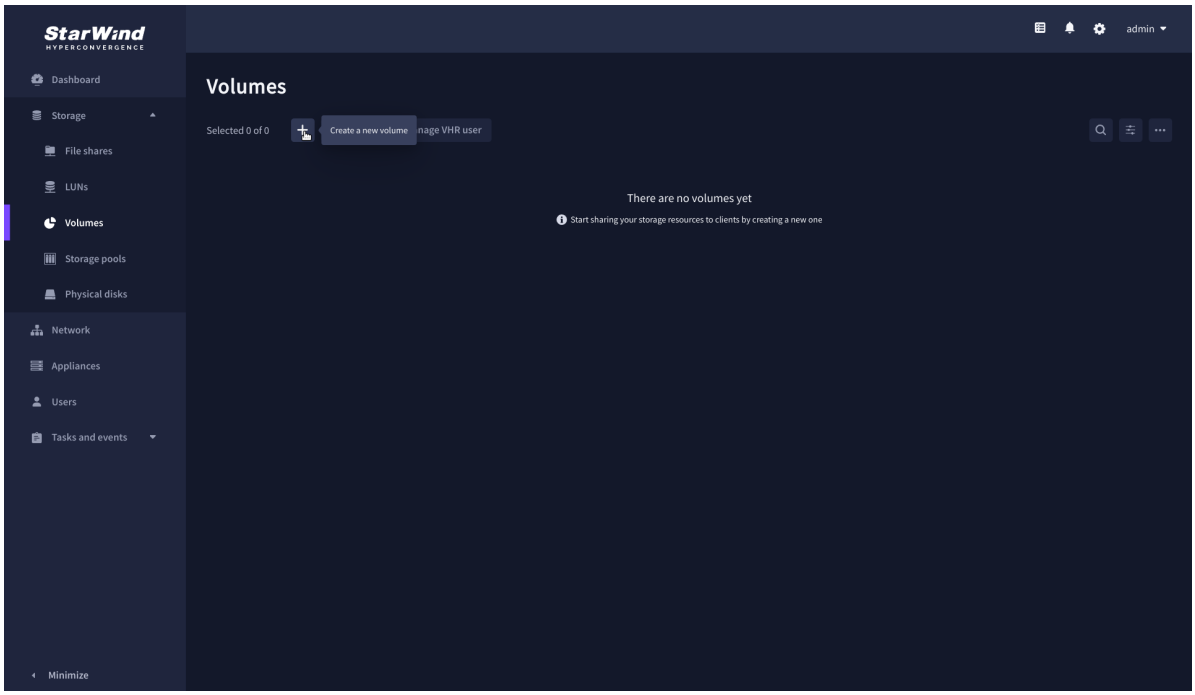



NOTE: The storage pool configuration may take some time, depending on the type of pooled storage and the total storage capacity. Once the pools are created, a notification will appear in the upper right corner of the Web UI.

IMPORTANT for CVM only: 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-changing-linux-i-o-scheduler-to-optimize-storage-performance/

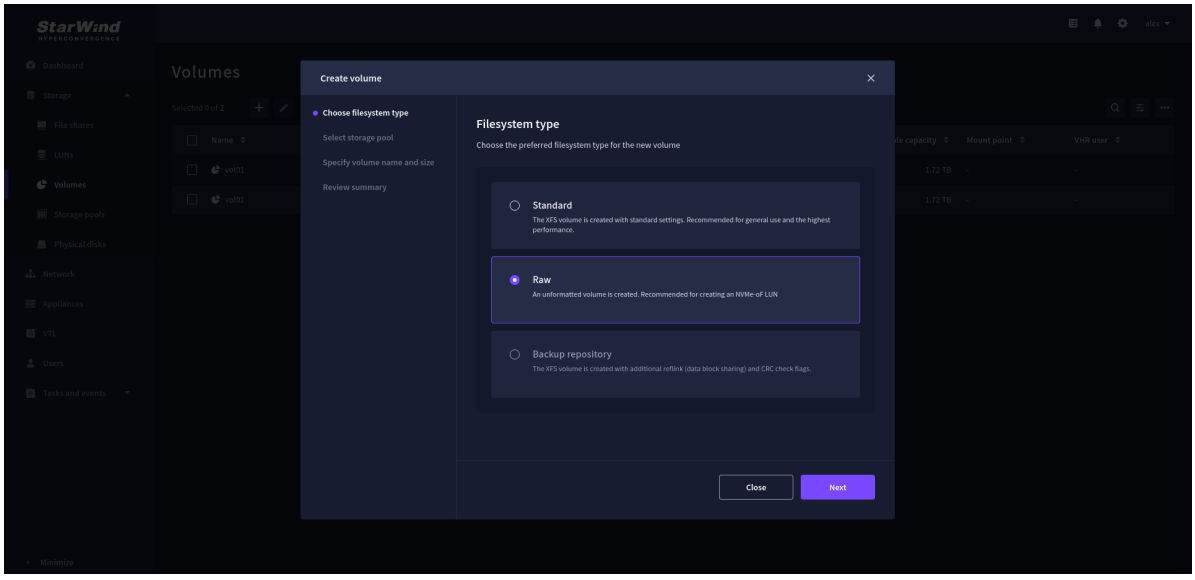
## Create Volume

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

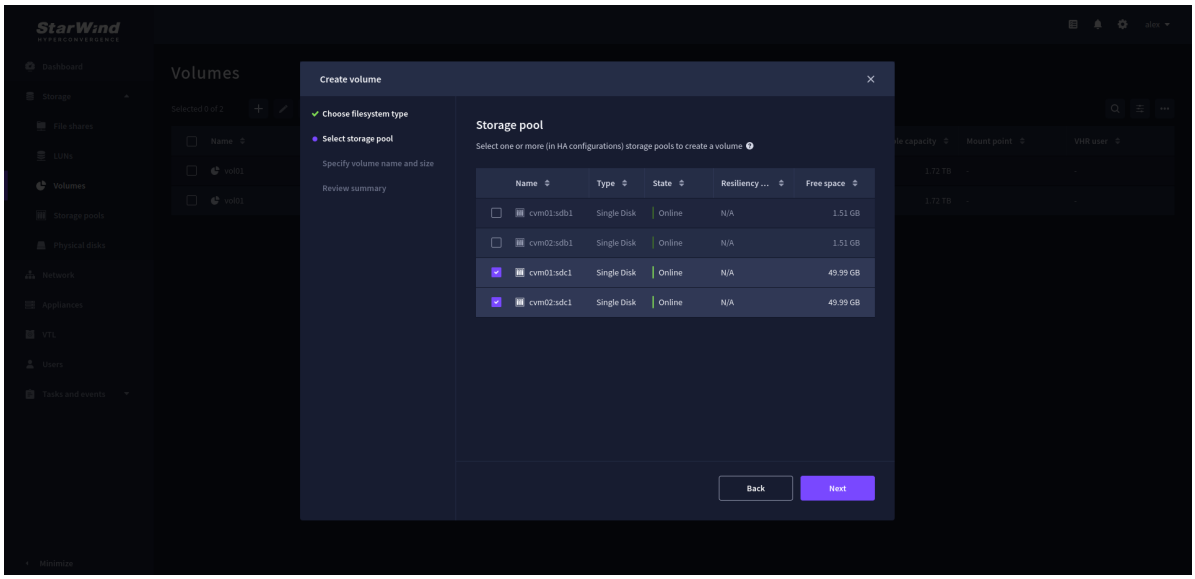



2. On the Filesystem type step, select Standard, then click Next.

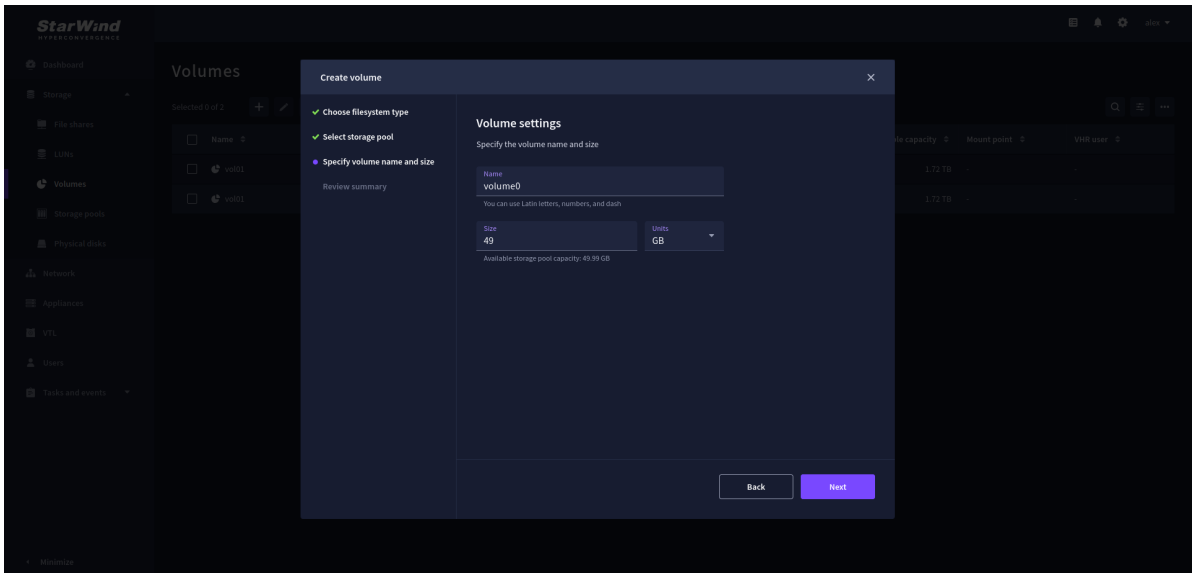
NOTE: Select 2 appliances for configuring volumes if you are deploying a two-node cluster with two-way replication, or select 3 appliances for configuring a three-node cluster with a three-way mirror.



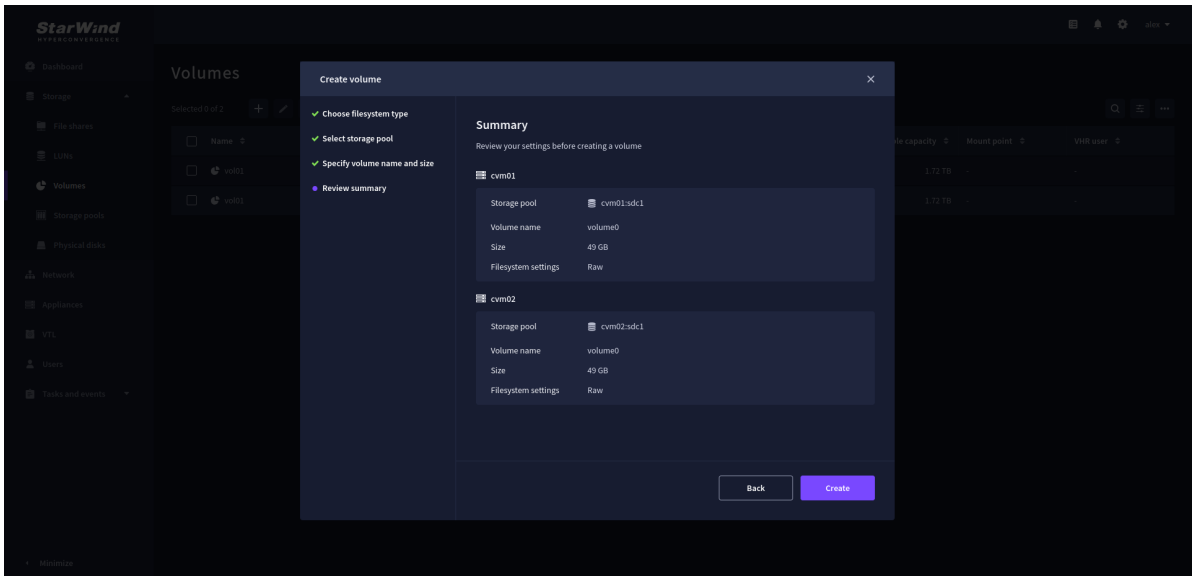

3. On the Storage pool step, select partner appliances on which to create new volumes, then click Next.




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

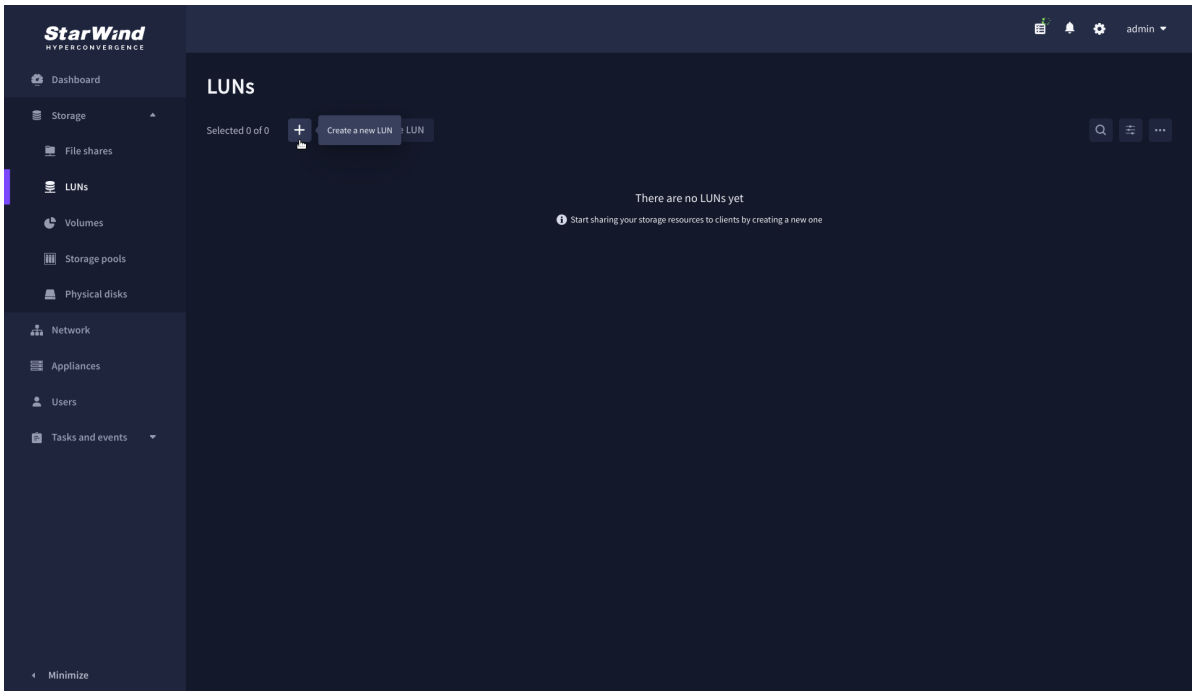



5. Review Summary and click the Create button to create the pool.

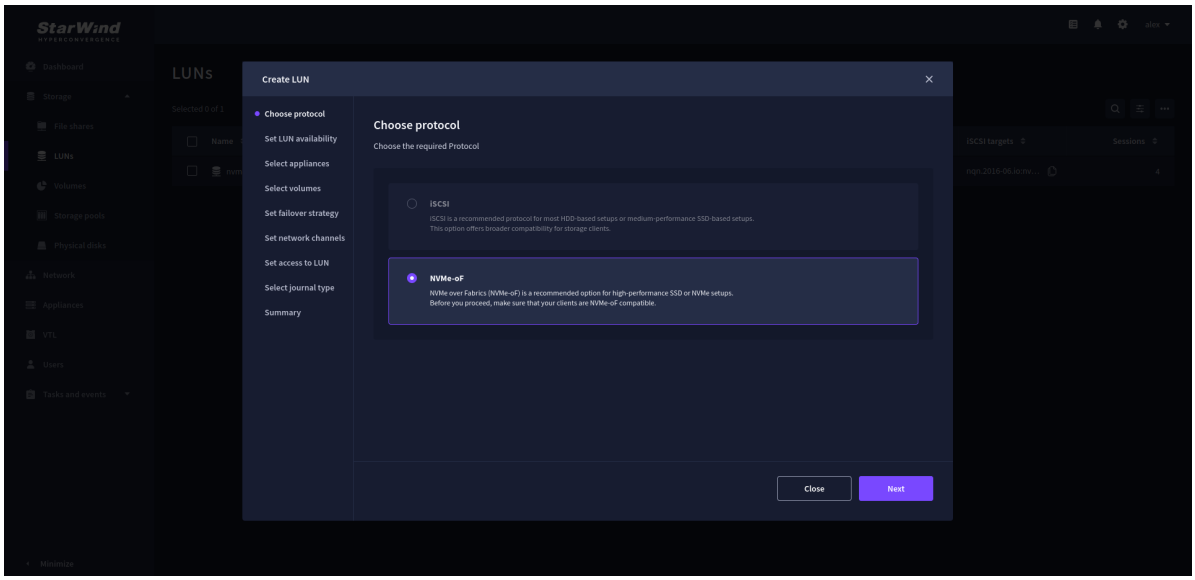



## Create Ha Lun Using Webui

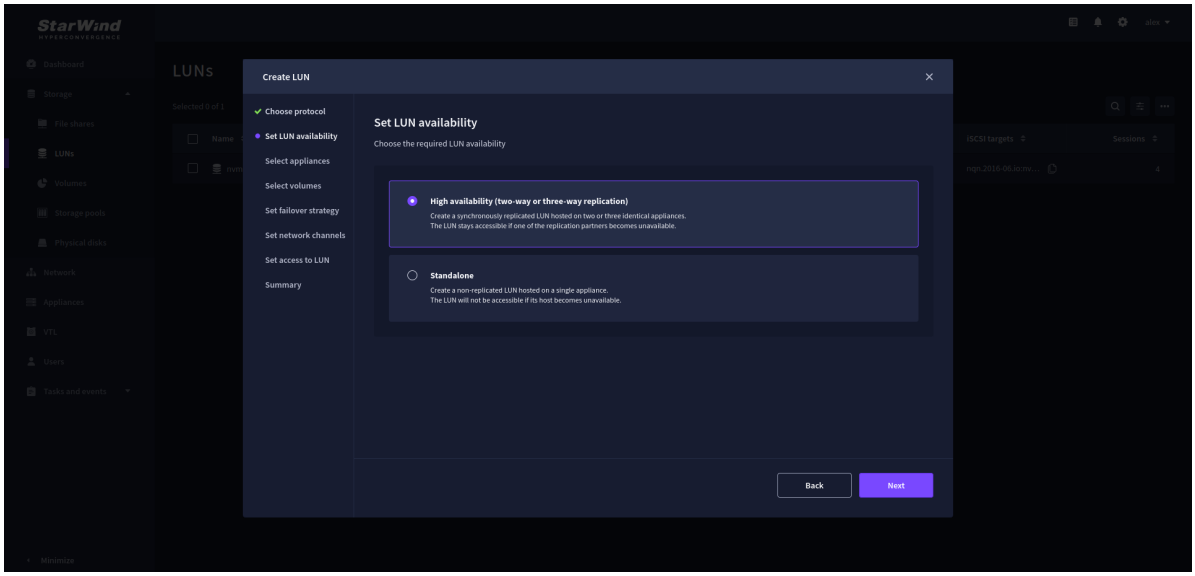
1. Navigate to the LUNs page and click the + button to open the Create LUN wizard.




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




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

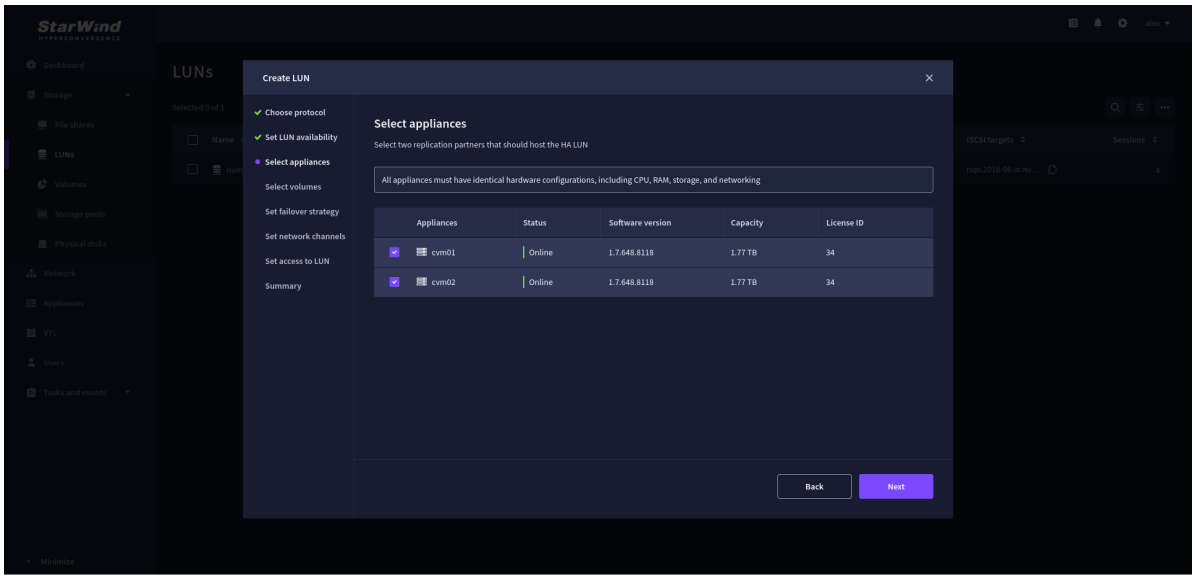



NOTE: The availability options for a LUN can be Standalone (without replication) or High Availability (with 2-way or 3-way replication), and are determined by the StarWind Virtual SAN license.

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

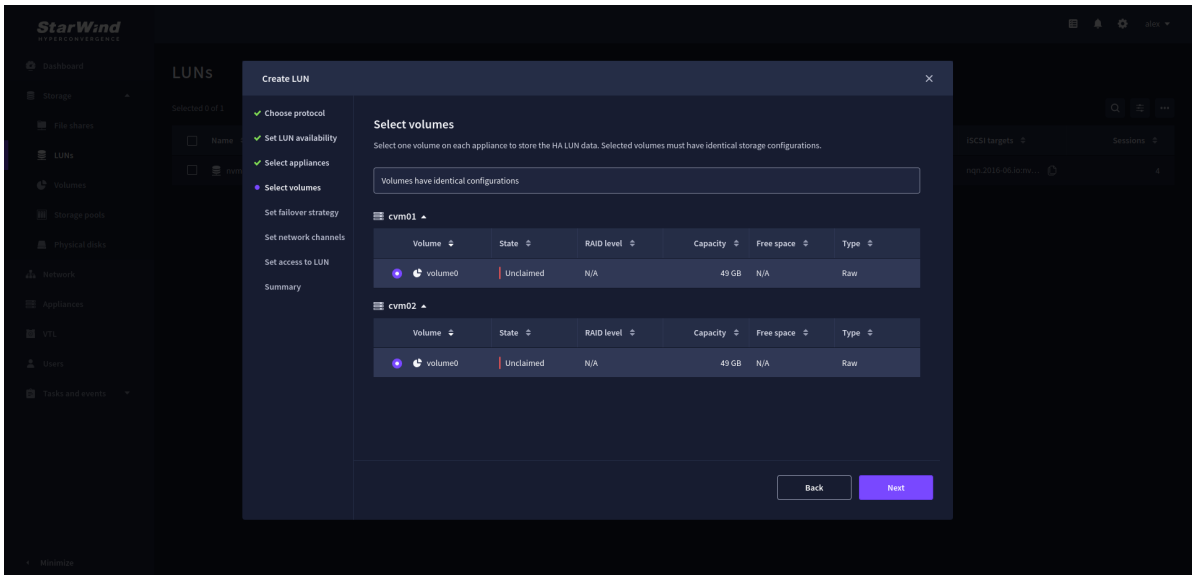
4. On the Appliances step, select partner appliances that will host new LUNs and click Next.

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.

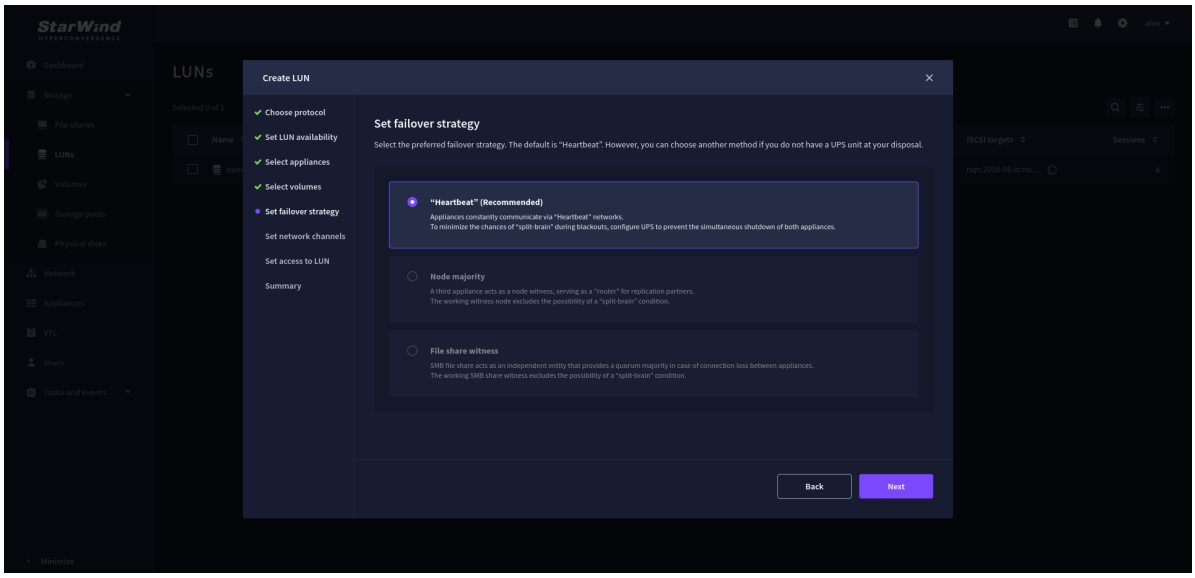
IMPORTANT: For optimal performance, the selected volumes must have identical underlying storage configurations.




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

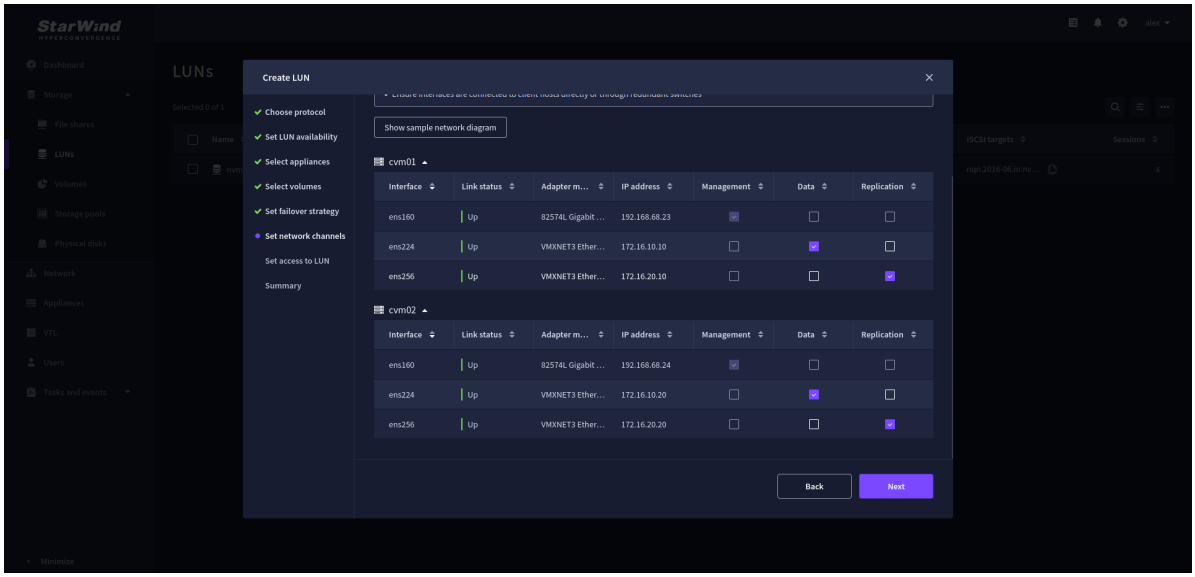
NOTE: The failover strategies for a LUN can be Heartbeat or Node Majority. In case of 2-nodes setup and None Majority failover strategy, Node witness (requires an additional third witness node), or File share witness (requires an external file share) should be configured. These options are determined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

two-node cluster.

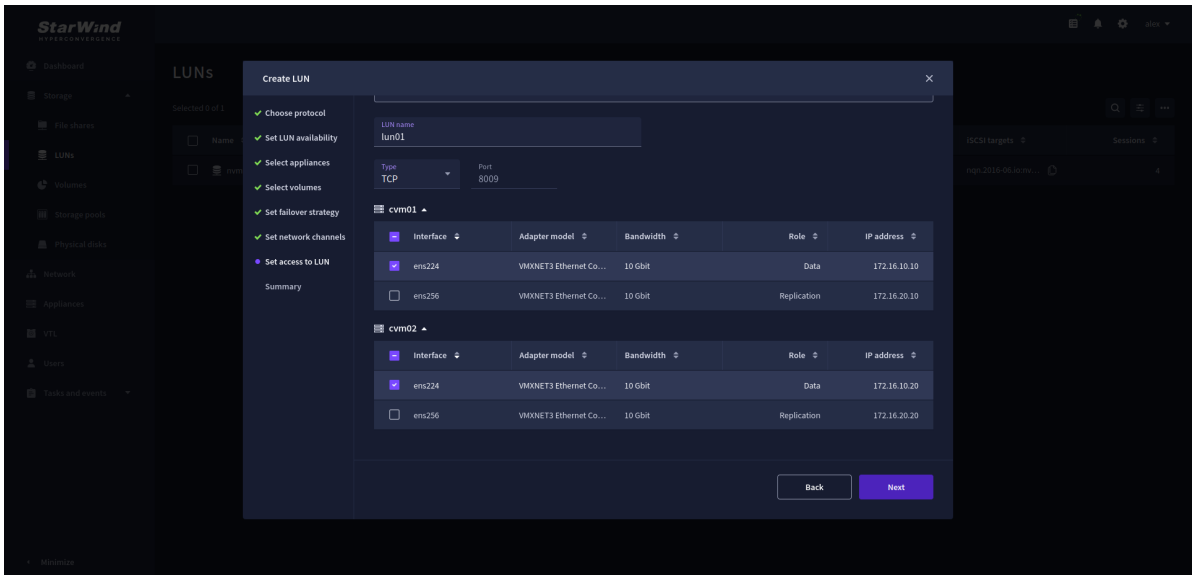



7. On the Network channels step, specify the Networks for Data and Replication, then click Next.

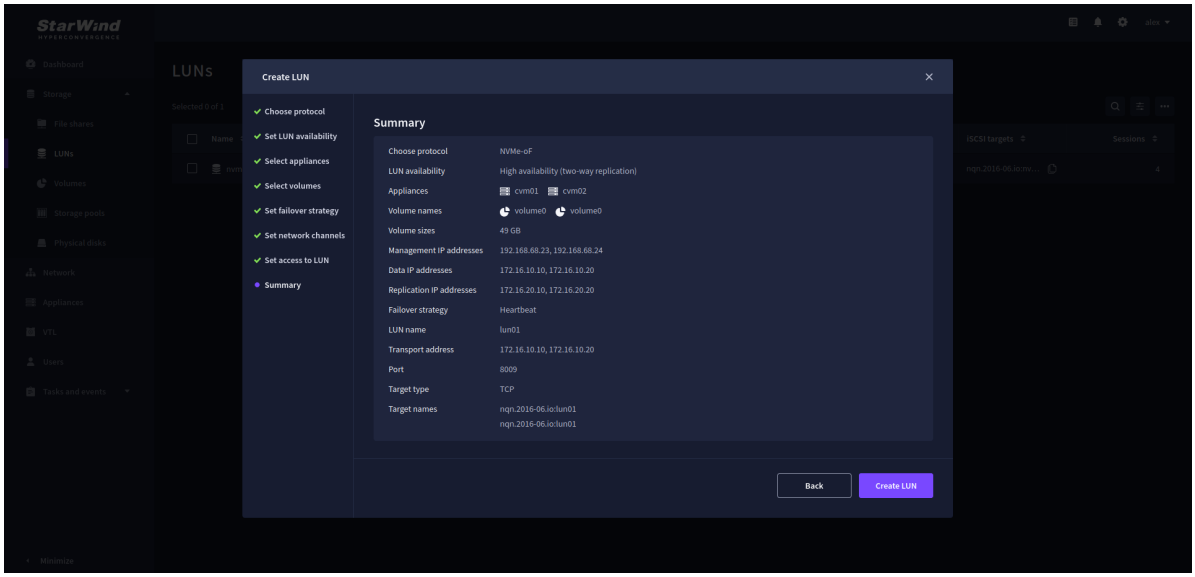
NOTE: Assign and configure at least two interface on each node, one per role (Data and Replication).




8. On the Access to LUN step, set LUN name and choose network interfaces to create NVMe-oF listeners, then click Next.




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




## Connecting Starwind Ha Storage To Proxmox Hosts

1. Connect to Proxmox host via SSH and create folder for nvmf-autoconnect.service.

```
# create a directory
mkdir -p /opt/starwind/starwind-nvmf-target
```

2. Download nvme-autoconnect.service here.

3. Upload starwind-nvmf-autoconnect.service to /etc/systemd/system/ directory.

4. Upload starwind-nvmf-autoconnect.py to /opt/starwind/starwind-nvmf-target/ directory and allow execution of the script.

```
chmod +x /opt/starwind/starwind-nvmf-target/starwind-nvmf-
autoconnect.py
```

5. Create the configuration file:

```
touch /etc/nvme/starwind-nvmf-autoconnect.conf
```

Add the following configuration to the configuration file (change nqn to the one created on the previous steps):

```
-t tcp -n nqn.2016-06.io:nvme -a 172.16.10.10 -s 8009 -i 8 -c
60 --ctrl-loss-tmo=-1
-t tcp -n nqn.2016-06.io:nvme -a 172.16.10.20 -s 8009 -i 8 -c
60 --ctrl-loss-tmo=-1
-t tcp -n nqn.2016-06.io:nvme02 -a 172.16.10.10 -s 8009 -i 8 -c
60 --ctrl-loss-tmo=-1
-t tcp -n nqn.2016-06.io:nvme02 -a 172.16.10.20 -s 8009 -i 8 -c
60 --ctrl-loss-tmo=-1
```

6. Enable starwind-nvmf-autoconnect.service service.

```
systemctl enable --now starwind-nvmf-autoconnect.service
```

Prerequisites Integration (StarLVM) for Proxmox VE

- Proxmox VE 8 or 9 cluster
- Shared SAN storage (iSCSI / NVMe-oF / FC) accessible from all nodes
- SSH access to cluster nodes
- Root privileges
- Identical device mapping for SAN LUN on each node
- Network access to the StarWind APT repository (repo.starwind.com) to download installation packages and scripts

Important

The shared SAN LUN must be connected to each node in the Proxmox cluster. It must appear on all nodes using the same persistent device identifier (for example,

/dev/disk/by-id/...).

Using inconsistent device paths across nodes (such as /dev/sdX) may result in storage activation failures, migration issues, or data inconsistencies.

To create a starlvm volume, shared LUNs should not already have existing LVM.

Installing StarWind Integration on Proxmox VE 8

1. Add StarWind Proxmox repo to the Proxmox VE 8 host.

```
echo "deb [arch=amd64 trusted=yes]
http://repo.starwind.com/proxmox/ master main" >
/etc/apt/sources.list.d/starwind-proxmox.list

curl --silent https://repo.starwind.com/keys/repo_public.key |
gpg --dearmor --yes --output /etc/apt/trusted.gpg.d/starwind-
proxmox.gpg
```

2. Run update.

```
apt update
```

3. Install StarWind Integration deb package.

```
apt install starwind-proxmox-plugin
```

4. Repeat steps 1-3 on all nodes in the cluster.

Installing StarWind Integration on Proxmox VE 9

1. Download the StarWind Integration installation script by executing command on each Proxmox VE 9 host:

```
curl -fsSL -o install-starwind-proxmox-plugin-pve9.sh
http://repo.starwind.com/san-and-nas/scripts/install-starwind-p
roxmox-plugin-pve9.sh
```

2. Execute the script:

```
bash install-starwind-proxmox-plugin-pve9.sh
```

The script will add the StarWind Proxmox repository to Proxmox VE 9 sources, validate the currently installed version, and install or update the latest StarWind Integration for Proxmox VE 9.

3. Repeat steps 1-2 on all nodes in the cluster.

Configure the attached storage

1. Create an LVM Physical Volume on the shared SAN LUN.

```
pvcreate /dev/mapper/mpatha
vgcreate swvg /dev/mapper/mpatha
vgs
```

Where:

/dev/mapper/mpatha – shared multipath device  swvg – LVM Volume Group name

2. Add a new storage to the PVE cluster using the starlvm plugin.

```
pvesm add starlvm swlun -vgname swvg -content images,rootdir -
shared 1
```

Parameters explanation:

| Parameter | Description |
|---|---|
| starlvm | Storage plugin type |
| swlun | Storage name in Proxmox |
| -vgname | LVM Volume Group used by the plugin |
| -content | Allowed content types (VM disks and containers) |
| -shared | Enables shared storage across the cluster |

3. Repeat steps 1-2 on all nodes in the cluster.

# 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

StarWind Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA
www.starwind.com ©2026, StarWind Software Inc. All rights reserved.