

StarWind Virtual SAN: Configuration Guide for XCP-ng [Xen], StarWind 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 XCP-ng version ≥ 8.3 .

Purpose

This document outlines how to configure an XCP-ng Cluster using StarWind Virtual SAN (VSAN), with VSAN running as a Controller Virtual Machine (CVM). The guide includes steps to prepare XCP-ng 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

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

Before you start

This guide was written on XCP-ng 8.3 and used XOA v5 UI elements. As XCP-ng is an evolving product, the visuals might be different on your end.

Multi-Node Cluster Configuration

StarWind Virtual SAN (VSAN) provides a flexible architecture that supports clusters of varying sizes, depending on the chosen hypervisor and deployment needs. The maximum size of a cluster is typically determined by the hypervisor vendor's limitations. For instance, VMware and Hyper-V impose specific constraints on the number of nodes that can be part of a single cluster. However, from StarWind's perspective, there are no

inherent restrictions on the number of nodes in a cluster.

While StarWind VSAN imposes no limits on the cluster's overall size, the synchronization for highly-available (HA) devices is limited to 2 or 3 nodes. To utilize StarWind VSAN effectively in larger clusters, devices can be replicated with an offset between nodes, in so-called GRID architecture. Replicating HA devices across different nodes with 2 or 3-way replica ensures that all nodes in the cluster contribute to storage redundancy and failover capabilities. This approach also maximizes storage resource utilization, data redundancy, and enhances the overall fault tolerance of the system.

To ensure optimal performance and fault tolerance, it is essential to establish resilient networking between cluster nodes. This involves:

- Redundant network links: Implementing multiple paths to avoid single points of failure.
- Sufficient bandwidth: Ensuring the network can handle synchronization traffic.
- Low latency: Minimizing delays to enhance HA replication efficiency.

Below is the conceptual diagram with StarWind VSAN deployment in a multi-node cluster.




Figure 1. Multi-node cluster configuration

The links between the cluster nodes can be connected via network switches or directly (for example for the replication links). As it was mentioned already, StarWind devices can be replicated in a 2-way or 3-way replica with an offset between nodes and connected on all cluster nodes. Data and Replication channels should not be used over the same physical link to avoid potential performance issues.

For example, in 4-nodes cluster the 1st HA device is replicated between the 1st and 2nd nodes, the 2nd HA device between the 2nd and 3rd nodes, the 3rd HA device – between the 3rd and 4th nodes and the 4th HA device – between the 4th and 1st nodes. The storage configuration could be different depending on the requirements.

Please do not hesitate to contact StarWind support to design the setup with multi-nodes.

2-Node Cluster Configuration

StarWind VSAN allows the creation of a minimalistic 2-node XCP-ng Cluster with shared storage while ensuring data redundancy and fault tolerance at the storage level.

To deploy a 2-node setup, both nodes must have identical hardware and meet the minimum StarWind VSAN system requirements, including supported OS, memory size, dedicated network interfaces, latency specifications, and minimum connection speed.

In a StarWind VSAN-based hyperconverged setup the compute, storage, and networking are integrated into a single solution. Shared storage, managed by StarWind VSAN running as a Controller Virtual Machine (CVM), is replicated between the same hosts where the StarWind CVM itself, the cluster, and the clustered VMs reside. StarWind VSAN CVM (can be installed on the local storage where the OS is installed) should be deployed on each XCP-ng node to configure the cluster. Local storage, such as physical disks, NVMe, or RAID arrays on each server intended for data, should be added to the StarWind Controller Virtual Machine (CVM). The StarWind CVM manages virtual disks replicated with the same StarWind CVM on another server, creating a shared storage pool. This highly available (HA) storage is connected to each server and used as cluster storage. Data replication in active-active mode between StarWind CVMs is performed on a block level over dedicated links between the servers. Production VMs managed by the cluster should be placed on Storage Repositories based on StarWind HA devices.

The diagram below illustrates the network and storage configuration of the hyperconverged setup.




Figure 2. 2-node cluster configuration

In this diagram, the 172.16.10.0/24 subnet is used to discover and connect the iSCSI targets (Data) and for StarWind heartbeat traffic, while the 172.16.20.0/24 subnet is used for the Replication traffic. The Data/Heartbeat and Replication links could be connected either through redundant switches or directly between nodes. The Data/Heartbeat and Replication networks should not share the same physical link to avoid potential performance issues. The management links should also be used for the redundant StarWind heartbeat connections.

Synchronous (active-active) replication ensures real-time synchronization and load balancing between cluster nodes. Each write operation requires confirmation from both storage nodes, guaranteeing reliable data transfers, which demands low-latency, high-bandwidth networks, as mirroring could fail on high-latency connections. In case when the node gets all the replication connections broken to the partner nodes (got isolated) and there is no way to get the partner state, the node may “assume” that the partner nodes are offline and could continue operations on a single-node mode, able to accept write commands from the initiators and use data written to it. A scenario when the data is written independently to each node during the replication connection failure is called a

“split-brain,” which can result in data corruption or inconsistency. To minimize or prevent split-brain scenarios, StarWind VSAN provides two failover strategies:

Heartbeat Failover Strategy

This strategy utilizes an alternate network connection, referred to as a heartbeat link, to monitor the availability of partner nodes. If replication networks fail but the heartbeat link remains active, StarWind services communicate through it to determine the partner node's state or exchange commands. The high-availability (HA) device with the lowest priority is then marked as unsynchronized and blocked from further read/write operations until the replication network is restored. Simultaneously, the partner device on the synchronized node flushes data from its cache to disk, ensuring data integrity in the event of an unexpected node failure.

Given the critical role of heartbeat links in maintaining data consistency, it is strongly recommended to configure multiple independent heartbeat channels during replica creation. This enhances stability and minimizes the risk of a “split-brain” scenario.

With the heartbeat failover strategy, the storage cluster can continue operating with only one StarWind node available.

Node Majority Failover Strategy

This strategy ensures that the system continues operating only when more than half of the nodes are active. Nodes that fall into the minority automatically stop accepting read and write requests, and their HA devices are marked as unsynchronized. Node Majority does not require heartbeat links; instead, it calculates the majority based on “votes.” The failure-handling process occurs when the node has detected the absence of the replication 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. In a two-node HA setup (which has only 2 “votes”), all nodes will be disconnected if there is a problem on the node itself, or in communication between them. Therefore, to keep at least one node active, the Node Majority failover strategy requires the addition of the third Witness node or file share (SMB). The witness participates in the nodes count for the majority, but neither contains data on it nor is involved in processing clients' requests.

With Node Majority failover strategy, the failure of only one node can be tolerated. The Witness node can be deployed on any Windows server running StarWind VSAN service. The Witness node also can be deployed as StarWind CVM on a supported hypervisor (Hyper-V, ESXi, Proxmox, XCP-ng, etc.). More details about the Node Majority failover strategy are available

here: <https://www.starwindsoftware.com/blog/whats-split-brain-and-how-to-avoid-it/>

3-Node Cluster Configuration

StarWind VSAN supports 3-way storage replica, so in case of 3-node XCP-ng Cluster the shared storage can be replicated on each of the nodes ensuring extra redundancy at the storage level.

To deploy a 3-node setup, all nodes must have identical hardware and meet the minimum StarWind VSAN system requirements, including supported OS, memory size, dedicated network interfaces, latency specifications, and minimum connection speed.

In a StarWind VSAN-based hyperconverged infrastructure, compute, storage, and networking are combined into a unified solution. Shared storage, managed by StarWind VSAN running as a Controller Virtual Machine (CVM), is replicated between the nodes hosting the StarWind CVM, the cluster, and the clustered VMs. StarWind VSAN CVMs, which can be installed on the local storage where the OS is installed, must be deployed on each XCP-ng node to configure the cluster. Local storage, such as physical disks, NVMe drives, or RAID arrays used for data, should be assigned to the StarWind Controller Virtual Machine (CVM). The StarWind CVM manages virtual disks that are replicated with CVMs on other nodes, creating a shared storage pool. This highly available (HA) storage is connected to each server and utilized as cluster storage. Data replication in active-active mode is performed at the block level over dedicated links between the servers. Production VMs managed by the cluster should reside on Storage Repositories based on StarWind HA devices.

The diagram below illustrates the network and storage configuration of the 3-node hyperconverged setup.




Figure 3. 3-node cluster configuration

In this diagram, 172.16.10.0/24, 172.16.11.0/24, 172.16.12.0/24 subnets are used to discover and connect the iSCSI targets (Data) and for StarWind heartbeat traffic, while 172.16.20.0/24, 172.16.21.0/24, 172.16.22.0/24 subnets are used for the Replication traffic. The Data/Heartbeat and Replication links could be connected either through redundant switches or directly between nodes. The Data/Heartbeat and Replication networks should not share the same physical link to avoid potential performance issues. The management links should also be used for the redundant StarWind heartbeat connections.

Synchronous (active-active) replication ensures real-time synchronization and load balancing between cluster nodes. Each write operation requires confirmation from both storage nodes, guaranteeing reliable data transfers, which demands low-latency, high-bandwidth networks, as mirroring could fail on high-latency connections. In case when the node gets all the replication connections broken to the partner nodes (got isolated) and there is no way to get the partner state, the node may “assume” that the partner nodes are offline and could continue operations on a single-node mode, able to accept write commands from the initiators and use data written to it. A scenario when the data is written independently to each node during the replication connection failure is called a “split-brain,” which can result in data corruption or inconsistency. To minimize or prevent split-brain scenarios, StarWind VSAN provides two failover strategies:

Heartbeat Failover Strategy

This strategy uses an alternate network connection, known as a heartbeat link, to monitor partner node availability. If replication links fail but the heartbeat remains online, StarWind services communicate via the heartbeat link to identify the partner node state or send the commands to each other. The HA device with the lowest priority is marked as unsynchronized and blocked for further read/write operations until the replication 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.

Since heartbeat links are critical for data consistency, it is recommended to assign multiple independent heartbeat channels during replica creation for improved stability and minimize the risk of the “split-brain” issue.

With the heartbeat failover strategy, the storage cluster will continue working with only one StarWind node available.

Node Majority Failover Strategy

This strategy ensures that the system continues operating only when more than half of the nodes are active. Nodes that fall into the minority automatically stop accepting read and write requests, and their HA devices are marked as unsynchronized. Node Majority does not require heartbeat links; instead, it calculates the majority based on “votes.”

The failure-handling process occurs when the node has detected the absence of the replication 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.

In a three-node HA setup (which has 3 “votes”), the system can tolerate the failure of one node while maintaining quorum. However, if two nodes fail, the remaining node will lose the majority and cease operations. Unlike in two-node configurations, a Witness node is not necessary in a three-node setup, as the majority is naturally achieved. More details about the Node Majority failover strategy are available

here: <https://www.starwindsoftware.com/blog/whats-split-brain-and-how-to-avoid-it/>

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>

IMPORTANT NOTE: StarWind VSAN on XCP-ng only supports PCI-E pass-through and currently doesn't support Virtual Disks, meanin that you need to have separate disk for OS and separate disk for StarWind storage (i.e separated by controllers). Before proceeding to the next part, make sure your storage can be passed through to the Virtual Machine.

Preconfiguring cluster nodes

1. Deploy XCP-ng image on each physical server you want to use for clustering.
2. The XenOrchestra cluster orchestrator should be deployed and configured before deploying any virtual machines and configuring HA storage.

XenOrchestra (or XOA for short) is very similar to vCenter in its nature and can be installed either as a standalone VM (manually) or automated using the online XOA deploy method.

Navigate to <https://xen-orchestra.com/#!/member/products> and select the preferred deployment method.

XOA quick deploy

Deploy now

OR

Alternative download

Download the XVA file and then **follow instructions**.



Download Xen Orchestra

3. For simplicity, this guide will use the online deployment method.



Deploy XOA in your infrastructure

You're about to start the deployment of Xen Orchestra in your infrastructure.

**This is a safe place**

Everything that happens from now on is between your browser and your host. Nothing is sent to our servers.


Issue with the deployment? Find alternate method in our [Deployment documentation](#)

Got it, let's go!

Your browser acts as a proxy for online deployment services, so you need to open this wizard on the PC that has access to the local network.

4. Open the Web console (XOA Lite) of the XCP-ng node on which you want to deploy XOA and accept the HTTPS certificate.

5. Specify the IP address of one of your XCP-ng nodes, login, and password. To learn primary-secondary meaning in XCP-ng clustering, read more here: <https://docs.xcp-ng.org/management/ha/>



The image shows a web-based configuration interface for StarWind. At the top, there is a progress bar with four steps: 'Connect' (active), 'Configure', 'Secure', and 'Deploy'. Below the progress bar, the main heading is 'Connect to your primary host'. The form contains three input fields: 'Primary host' with a location pin icon and the value '172.16.50.81', 'Login' with a user icon and the value 'root', and 'Password' with a lock icon and masked characters '.....'. A small information icon and text 'IP address or URL' are located below the Primary host field. At the bottom of the form are two buttons: 'Go back' and 'Continue'.

Connect Configure Secure Deploy

Connect to your primary host

Primary host

📍 172.16.50.81

ℹ IP address or URL

Login

👤 root

Password

🔒

Go back Continue

6. Select local storage (for example, BOSS card or OS RAID-1 disk) where XOA will be installed, and which network it should use.

7. Uncheck the DHCP box if you want a static IP address for the XOA appliance.

Connect **Configure** Secure Deploy

Configure your server

STORAGE

Storage

Local storage - 168 GiB left

20 GiB required.

NETWORK

Network

Pool-wide network associated with eth0

☐ DHCP

XOA IP address

172.16.50.83

Netmask

255.255.255.0

DNS

172.16.50.8

Gateway

172.16.50.1

[+ Add another DNS](#)

[Go back](#) [Continue](#)

8. Specify login credentials, password, and SSH access to the XOA appliance.

Connect Configure **Secure** Deploy

Create accounts

XOA ADMIN ACCOUNT

Login

admin@swx.local

Password

.....

XOA SSH ACCOUNT

Login

xoa

Password

.....

[Go back](#) [Deploy](#)

9. XOA Appliance will begin its deployment. Make sure to keep the tab open until deployment is complete.

10. Once deployed, click Access XOA and proceed to log into your freshly deployed XOA.


11. Navigate to the XOA tab, click Updates, register your XOA appliance and install all the latest updates.

12. To add more servers to XOA, navigate to Settings -> Servers. Specify label, IP address, and login information. Click Connect once filled.

1 Items							
Label	Host	Username	Password	Status	Read Only	Unauthorized Certificates	Pool
swx-xcpng-pool	172.16.50.81	root	password	Enabled	☐	☒	swx-xcpng-
swx-xcp-02	172.16.50.82	root	*****	☐	☐	HTTP proxy URL	Connect


13. Click on the error message and accept the self-signed certificate.

14. Navigate to Home -> Pools, and there you will see individual pools, one per added host to XOA. Select one with XOA installed and click the Add Host button on the right-hand side.



15. Select recently added host(s) from the drop-down menu.

16. Click on the Pool name at the top and name it as you wish to call your cluster.



17. Now, multipath needs to be configured on all hosts

before proceeding to StarWind configuration. To do so, navigate to Home -> Hosts and select the recently added node.

18. Navigate to the Advanced tab within host information and click on the Disabled button to enable Multipath.

The screenshot shows the StarWind Web UI interface for a host named 'xcp-ng-02'. The 'Advanced' tab is selected, displaying various Xen settings. The 'Multipath' setting is currently 'Disabled', indicated by a red button with a red dot and the word 'Disabled'.

Xen settings	
Address	172.16.50.82
Status	Enabled
IOMMU	Enabled
Power on mode	☐
Control domain memory	4.19 GiB
Host uptime	Started 5 days ago
Toolstack uptime	Started 15 minutes ago
Version	8.3.0 ()
Build number	cloud
ISCSI IQN	iqn.2026-02.com.example:ffd6a334
Multipathing	● Disabled
Scheduler granularity	CPU Reboot to apply updates
Remote syslog	Click to edit
Custom fields	

19. You will be prompted with a warning message concerning the existing storage repositories. Since the recently added node(s) don't host the XOA appliance, you can click OK to continue.

NOTE: if you have VMs deployed on these nodes, you need to either migrate them to other node(s) or power them off before doing this operation.

⚠ Enable multipathing



The host will lose their connection to the SRs. Do you want to continue?

i This action will not be fulfilled if a VM is in a running state. Please ensure that all VMs are evacuated or stopped before performing this action!

Click to see concerned SRs ^

- 📀 XCP-ng Tools - swx-xcpng-pool
- 📀 Removable storage - xcp-ng-02
- 📀 DVD drives - xcp-ng-02
- 📀 Local storage (178.1 GiB free - thin) - xcp-ng-02


OK

Cancel

20. Repeat steps on other recently added nodes before proceeding to the Master node.

21. Since Master nodes are running the XOA appliance, to enable multipath on Master nodes, you first need to migrate the XOA appliance to a different node within a Pool. Navigate to Home -> VMs and click on XOA VM.

22. Click Migrate in the top right corner.



23. Specify the destination host and the local SR and click OK. You can check the progress in Tasks tab.

⚠ Migrate VM

×

Select a destination host:

xcp-ng-02 (58.57 GiB RAM f.⌵

Select a destination SR

Local storage (178.1 GiB free.⌵

For each VDI, select an SR (optional) ▼

Select a migration network:

Select network(s)...

▼

* optional

OK

Cancel

24. Once migration is complete, navigate to Home -> Hosts, then select the Master node to complete enabling Multipath.


25. Additionally, in the Advanced settings of each host in a pool, you can change the iSCSI Initiator IQNs from automatically generated to what you prefer. Alternatively, you can also do it by changing the iSCSI configuration file. /etc/iscsi/initiatorname.iscsi to set the initiator name.

```

GNU nano 2.3.1      File: /etc/iscsi/initiatorname.iscsi
InitiatorName=iqn.2026-02.com.example:xcp-ng-01
InitiatorAlias=xcp-ng-01
  
```

26. To prepare networking for StarWind HCI installation, navigate to Home -> Hosts and select the first node in the list.

27. 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. Navigate to Home -> Hosts, select the first node, and go to the Network tab.



#	Device	Network	VLAN	Address	Mode	MAC	MTU	Speed	Default locking mode	NBD	Status
0	eth0	Pool-wide network associated with eth0	None	172.16.50.01	Static	96:80:c8:1c:39:ee	1500	1 Gb/s	Off		Connected
1	eth1	Pool-wide network associated with eth1	None		None	96:80:c8:1c:39:ef	1500	0 b/s	Off		Connected
2	eth2	Pool-wide network associated with eth2	None		None	96:80:c8:1c:39:ec	1500	0 b/s	Off		Connected
3	eth3	Pool-wide network associated with eth3	None		None	96:80:c8:1c:39:ed	1500	0 b/s	Off		Connected
4	eth4	Pool-wide network associated with eth4	None		None	80:30:e9:42:d7:08	1500	10 Gb/s	Off		Connected
5	eth5	Pool-wide network associated with eth5	None		None	80:30:e9:42:d7:0c	1500	10 Gb/s	Off		Connected

28. Identify networks that you want to use for StarWind traffic and iSCSI. Click None in the Mode column on the network that you want to use as iSCSI/StarWind heartbeat and select Static. You will be prompted with a wizard for specifying the IP address. This IP will be used by the Initiator to connect iSCSI storage presented by the VM.

📍 Configure IP address



Static IP address

Static IPv6 address

Netmask

DNS

Gateway

OK

Cancel

29. Navigate to Home -> Pools -> Select Pool -> Network. Here you can rename network interfaces and set the MTU size to 9000.

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.

NOTE: If MTU is changed on the PIFs to 9000, please make sure that it's changed on the corresponding adapters. Ping the IP address of the partner node using the command: `ping -M do -s 8000 172.16.10.254` via SSH or Console within XenOrchestra.

30. Repeat step 27 for all network adapters. which will be used for Synchronization and iSCSI/StarWind heartbeat traffic.

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

```
node.startup = automatic
```

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

32. Edit /etc/multipath/conf.d/custom.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"
    }
}
```

NOTE: It is recommended to reboot hosts after applying the custom multipath config.

33. Repeat steps on all XCP-ng hosts.


Deploying Starwind Virtual San Cvm

1. Download StarWind VSAN CVM OVA: [VSAN by StarWind: Overview](#)
 2. Extract the VM CVM-0.vmdk file from the downloaded archive.
 3. Navigate to Import -> Disk. Select Pool and local storage SR. Upload CVM-0.vmdk file to the XOA. Repeat for all other host(s) local storage. Alternatively, you can use an ISO repository and an ISO image of StarWind VSAN CVM.
- NOTE: VMDK files will be converted to VHD format.


4. Navigate to Home and click +VM in the top right corner.

5. Select Pool, Select Generic Linux BIOS Template, Name VM, Assign 8 vCPU (recommended) and at least 8 GB of RAM (16 GB recommended), Select PXE install settings, add 3 networks (management network, iSCSI, and Sync), and remove disk created by the template. In advanced settings uncheck Boot VM after creation and click Create.


6. You will be redirected to the created VM. Navigate to disk, click Attach disk, and proceed to attach the recently imported disk.




7. Navigate to the Advanced tab. Reorder Boot order by dragging the hard drive to the top and clicking Save.




8. In the Advanced tab, change the NIC type to Intel e1000 and click Save.



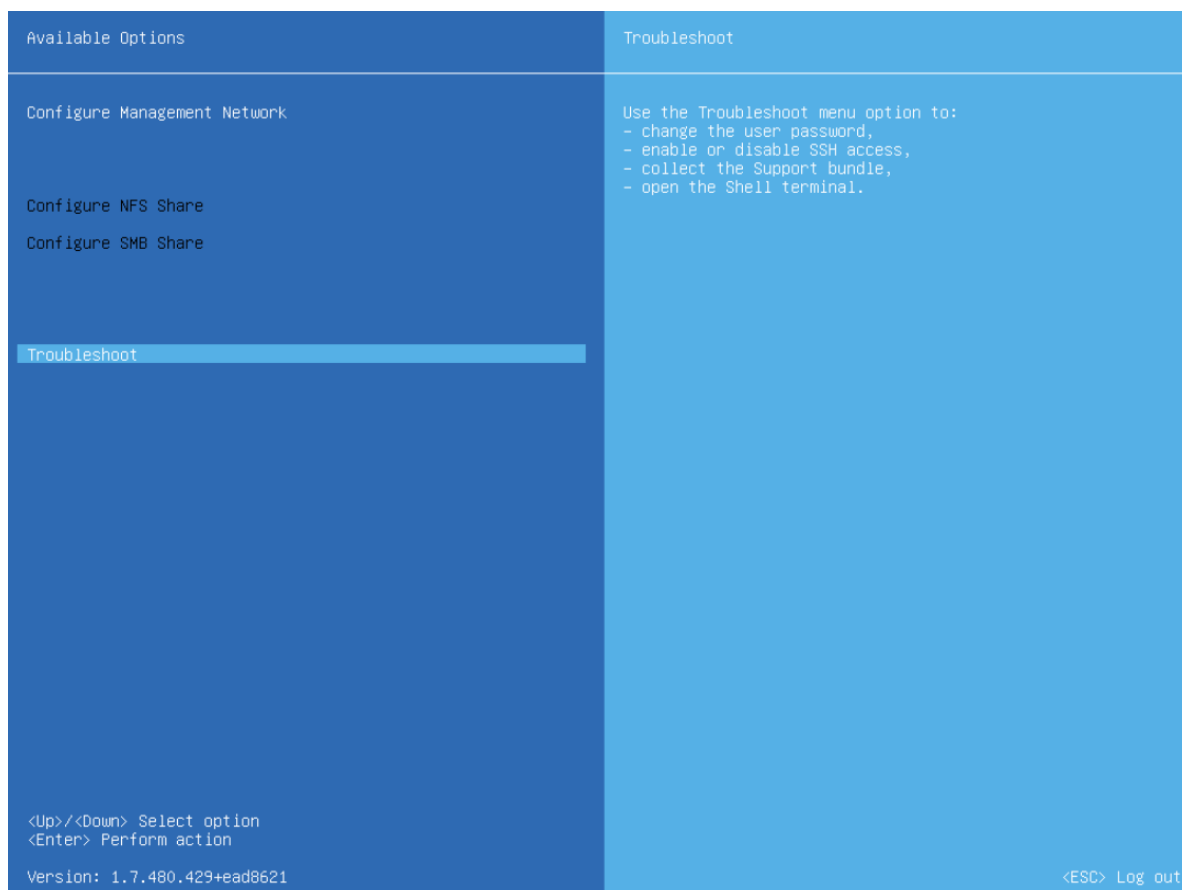
9. Navigate to Disk and select guest-tools.iso under primary virtual disk.




10. Power on the VM and navigate to Console.

[illegible]

12. Navigate to Troubleshoot and select Open Shell Terminal.





13. Inside the terminal, perform the following actions:

```
sudo mount /dev/cdrom /mnt
sudo bash /mnt/Linux/install.sh
sudo umount /dev/cdrom
```

This will install and enable XCP-ng guest tools

14. Shutdown Virtual Machine.

15. Navigate to the Disks tab and remove guest-tools.iso.

16. Repeat all the steps from this section on other XCP-ng hosts.

Attaching Storage To Starwind Virtual San Cvm

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

IMPORTANT NOTE: StarWind VSAN on XCP-ng only supports PCI-E pass-through.

Attaching Pcie Device To Starwind Virtual San Cvm

1. Shutdown StarWind VSAN CVM.
2. Navigate to Home -> Hosts, select the host, and open the Advanced tab.

The screenshot shows the Xen Orchestra web interface. On the left is a sidebar with navigation links: Home, Dashboard, Self service, Backup, XOA, Settings, Jobs, Hub, Proxies, About, Tasks, XOSTOR, Import, New, REST API doc, Try XO 6, and Check XOA. The main content area is for host 'xcp-ng-01' (UUID: 47b7b63d-8abc-4dba-9f49-2e6025d676f1). The 'Advanced' tab is selected, showing 'Xen settings'.

Xen settings	
Address	172.16.50.81
Status	Enabled
IOMMU	Enabled
Power on mode	☐
Control domain memory	4.19 GiB
Host uptime	Started 2 months ago
Toolstack uptime	Started 2 months ago
Version	8.3.0 ()
Build number	cloud
iSCSI IQN	iqn.2026-02.com.example:xcp-ng-01
Multipathing	● Enabled

3. Scroll down to PCI devices and toggle Passthrough enabled for the RAID controller or NVMe devices you want to use for StarWind HA.

PCI Devices

1 2

12 items

ID	PCI ID	Class name	Device name	Passthrough enabled
0838	0000:01:00.0	System peripheral	Integrated Lights-Out Standard Slave Instrumentation & System Support	☐
d41a	0000:01:00.1	VGA compatible controller	MGA 0200eH3	☐
2364	0000:01:00.2	System peripheral	Integrated Lights-Out Standard Management Processor Support and Messaging	☐
cadf	0000:01:00.4	USB controller	iLO5 Virtual USB Controller	☐
fa77	0000:02:00.0	Ethernet controller	NetXtreme BCM5719 Gigabit Ethernet PCIe	☐
7a53	0000:02:00.1	Ethernet controller	NetXtreme BCM5719 Gigabit Ethernet PCIe	☐
035d	0000:02:00.2	Ethernet controller	NetXtreme BCM5719 Gigabit Ethernet PCIe	☐
2526	0000:02:00.3	Ethernet controller	NetXtreme BCM5719 Gigabit Ethernet PCIe	☐

NOTE: this action will reboot the host in order to enable PCI-E passthrough for your VMs. Make sure you don't have VMs running. You can also migrate XOA appliance between local storages for this

4. Navigate to StarWind VM and go to Advanced tab.

Xen Orchestra

view-xcp-ng-pool > xcp-ng-01

CVM-01
0b/1a/779-5ab6-d18a-de36-d387f13c71a9
Click to edit

General Stats Console Network Disks Snapshots Backup Logs **Advanced**

Pause Suspend Coalesce leaf Force reboot Force shutdown Warm migration Convert to template

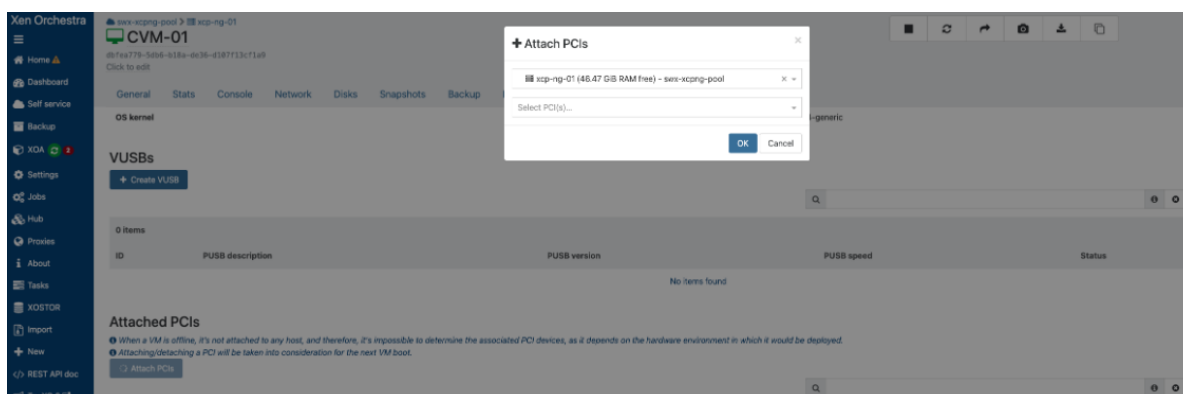
Boot order

- II Hard-Drive
- II Network
- II DVD-Drive

Xen settings

Virtualization mode	Hardware virtualization with paravirtualization drivers enabled (PVHVM) Convert to PV
Start delay (seconds)	0
CPU mask	None
CPU weight	Default (256)
CPU cap	Default (0)
Auto power on	☒
Protect from accidental deletion	☐
Protect from accidental shutdown	☐

5. Scroll down to the Attach PCIs section. Click the Attach PCIs button. In the wizard, select which PCI-E device to attach and click OK.



6. Repeat steps 1-5 on all nodes for all PCI-E devices you want to attach to StarWind CVM.

7. Navigate to the Advanced tab and enable Auto power on for StarWind VSAN CVM on each host.

Initial Configuration Wizard

1. Start the StarWind Virtual SAN Controller Virtual Machine.

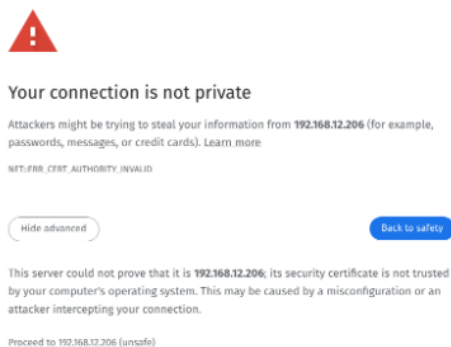
2. Launch the VM console to view the VM boot process and obtain the IPv4 address of the Management network interface.

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. Use of static IP addresses is highly recommended.


Default credentials for TUI: user/rds123RDS

NOTE: Please make sure to change the default password


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...



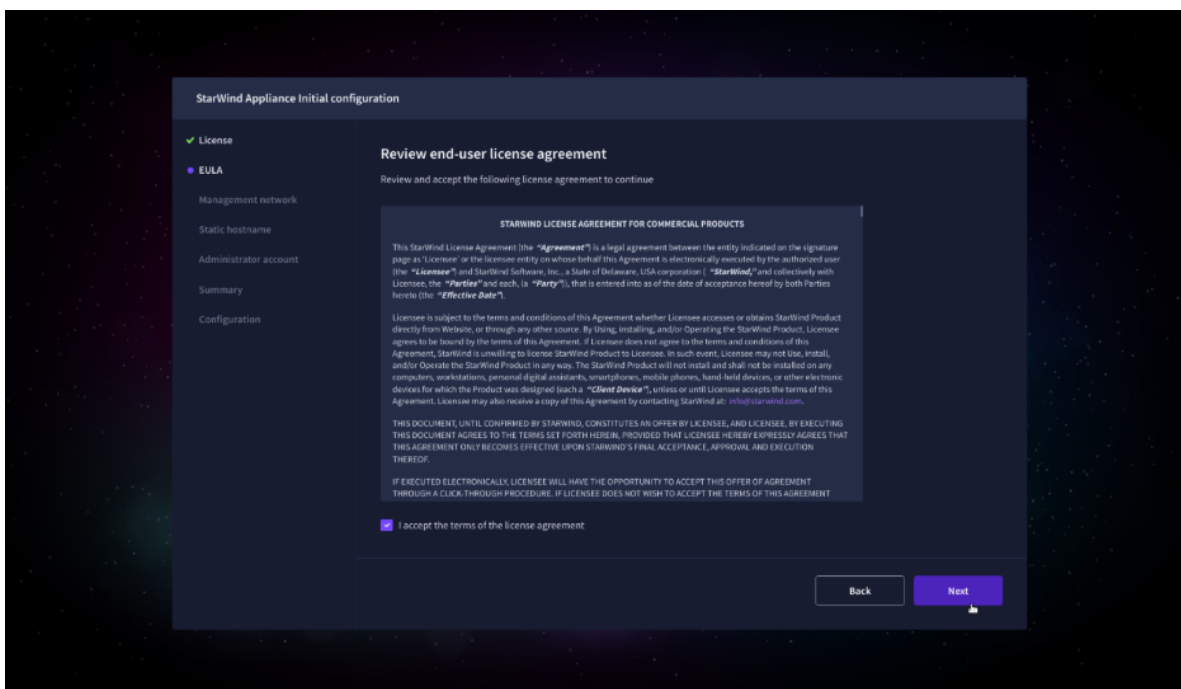
4. On the Welcome to StarWind Appliance screen, click Start to launch the Initial Configuration Wizard.



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.




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

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

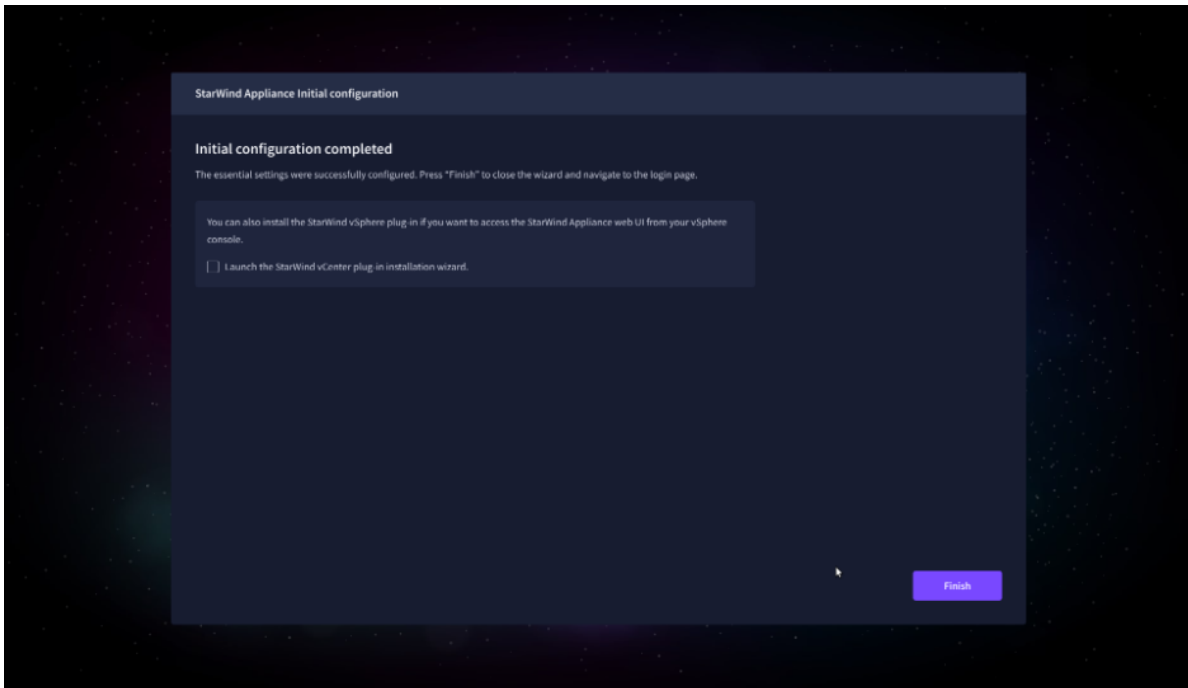


The screenshot shows the 'StarWind Appliance Initial configuration' window at the 'Review summary' step. The sidebar on the left now highlights 'Summary'. The main area displays a summary of the configuration: 'License type' is 'Paid 3 Nodes'; 'Network settings' include Interface 'ens160 (82574L Gigabit Network Connection)', Bandwidth '1 Gbit', MTU '1500', IP address '192.168.12.206', and Appliance hostname 'SW1'; 'Credentials' show 'Administrator username' as 'admin'. 'Back' and 'Configure' buttons are at the bottom right.

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



12. After the configuration process is completed, uncheck the checkbox to skip vCenter plug-in installation step, click Finish, and proceed to the Login page.




13. Repeat steps 1 through 12 on each XCP-ng 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.

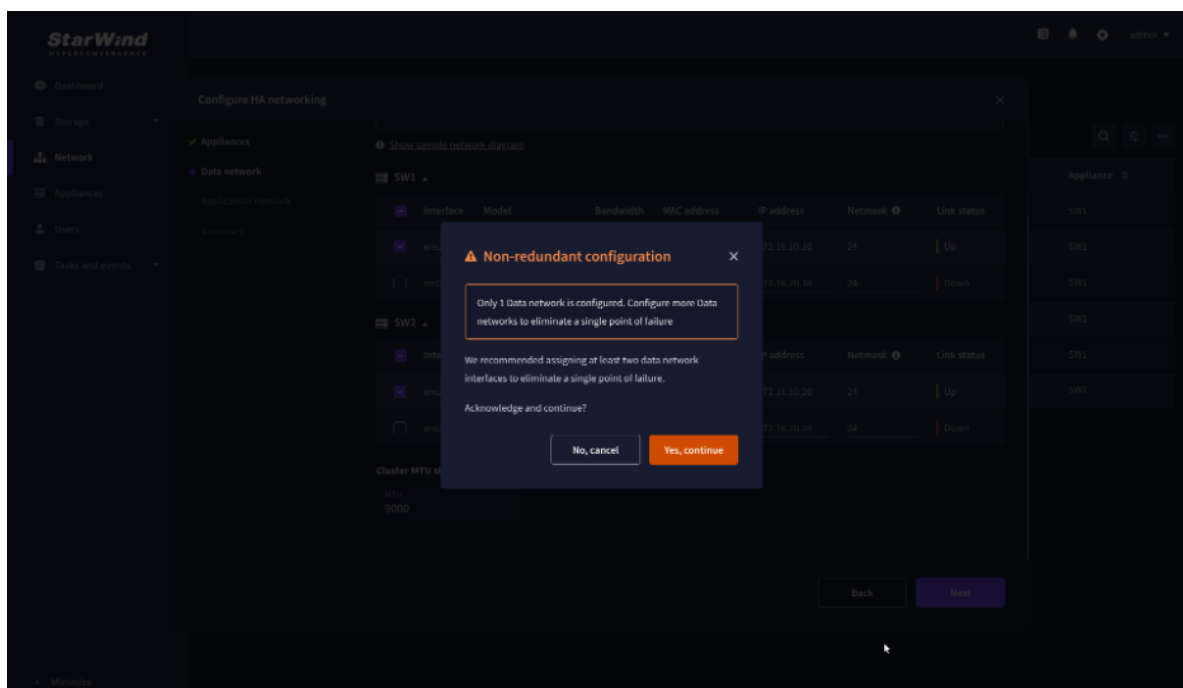
4. Assign MTU value on all selected network adapters, e.g. 1500 or 9000 bytes. If you are using network switches with the selected Data Network adapters, ensure that they are configured with the same MTU size value. In case of MTU settings mismatch, stability and performance issues might occur on the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



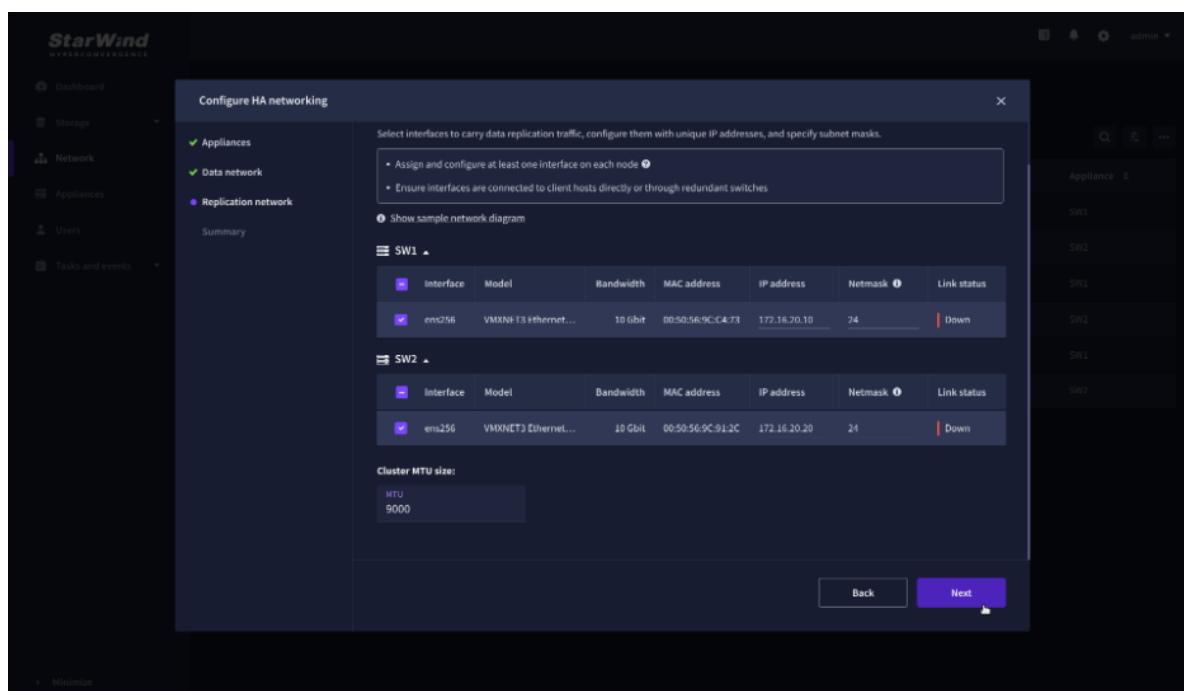
7. On the Replication Network step, select the network interfaces designated to carry the traffic for synchronous replication. Assign and configure at least one interface on each appliance with a static IP address in a unique network (subnet), specify the subnet mask and Cluster MTU size.

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.

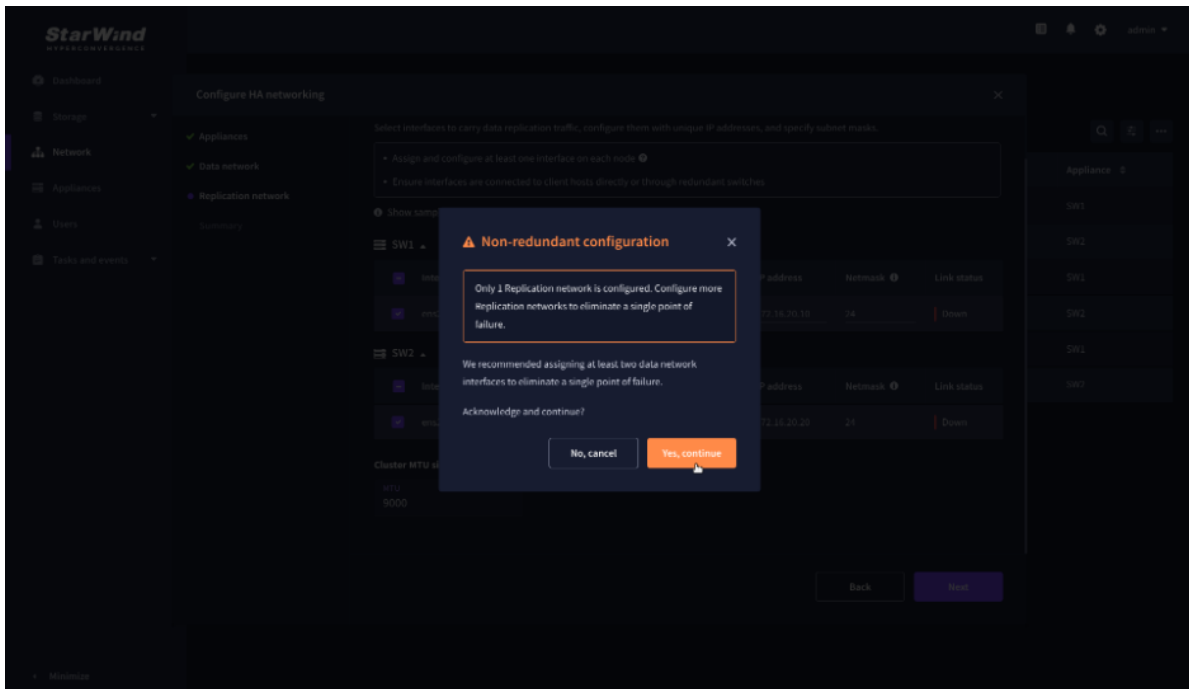
8. Assign MTU value on all selected network adapters, e.g. 1500 or 9000 bytes. If you are using network switches with the selected Replication Network adapters, ensure that they are configured with the same MTU size value. In case of MTU settings mismatch, stability and performance issues might occur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




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-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 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

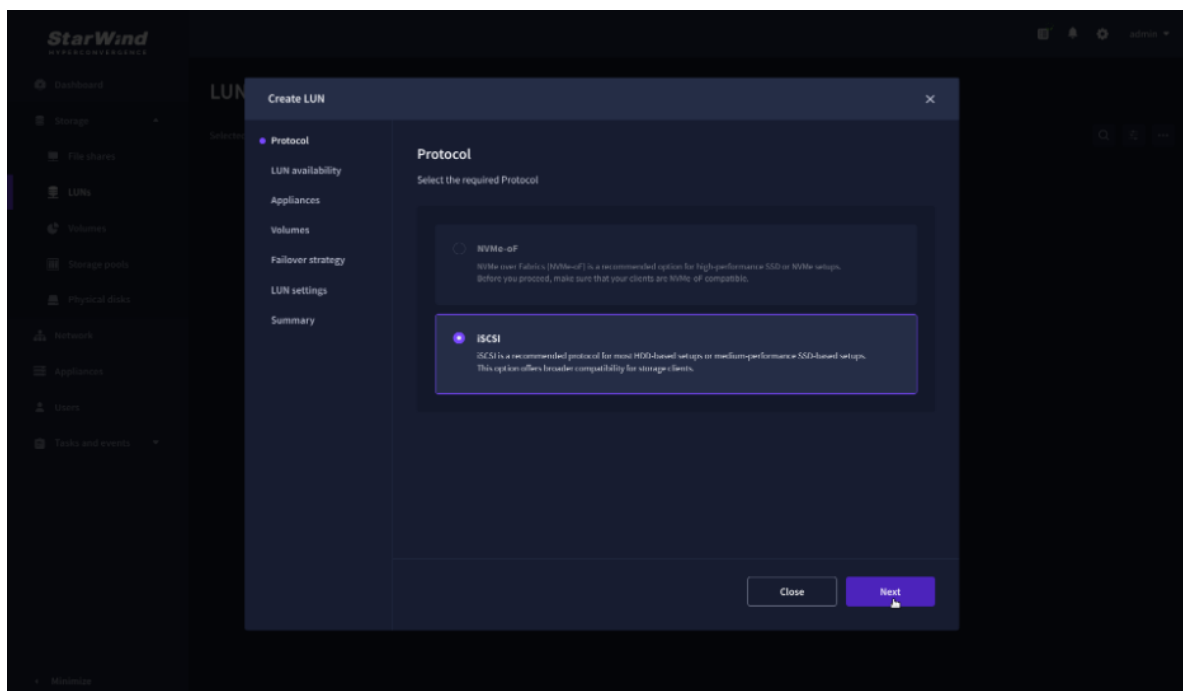
This section describes how to create LUN in Web UI. This option is available for the setups with Commercial, Trial, and NFR licenses applied.

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 a 2-node setup and the Node Majority failover strategy, the Node witness (requires an additional third witness node), or the 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.
 NOTE: For XCP-ng configuration, ensure that the 512 bytes block size is selected.



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



Adding Starwind Storage To Xcp-ng

There are several ways to add StarWind storage to an XCP-ng cluster. One of which is via UI. Navigate to New and select Storage. Select one of the nodes within your pool and select iSCSI in Storage type.

Xen Orchestra

Home ⚠

Dashboard

Self service

Backup

XOA 🔄 2

Settings

Jobs

Hub

Proxies

About

Tasks

XOSTOR

Import

+ New

</> REST API doc

📢 Try XO 6 🔗

📁 Check XOA

VM ▾

5x 🖥 (of 6)

☐ ● CVM-01

☐ ● CVM-02

☐ ● Debian Bookworm 12

☐ ● Debian test VM

☐ ● XOA

🖥 VM

📀 Storage

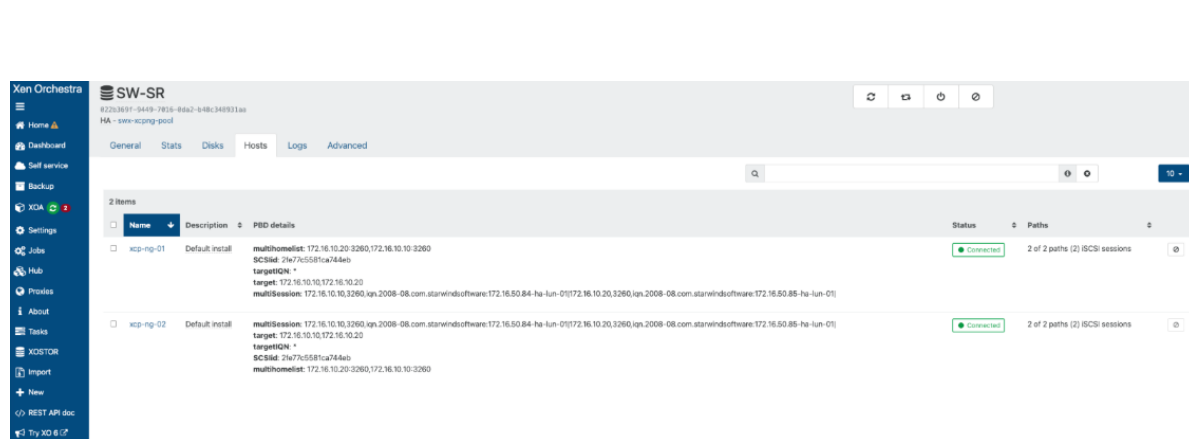
🌐 Network

☁ Server


Enter CVMs iSCSI IP addresses with a comma between them, let XCP-ng scan the storage and click Create.

NOTE: Please make sure that storage is discovered and connected via all iSCSI interfaces from each CVM with an HA device created.

Once StarWind HA is presented as a Storage Repository to your pool, double-check iSCSI multipath configuration by navigating to Home -> Storage -> Select created SR -> Hosts. Find below the example of paths for the Storage Repository connected to a 2-node XCP-ng cluster.



Navigate to Home -> Pools, select your pool, go to the Advanced tab, and enable High Availability. Select StarWind HA as heartbeat SR and click OK.



Now, you can deploy your VMs on top of StarWind HA. To achieve HA for XOA appliance, move it to one of your StarWind HA devices.

Conclusion

Following this guide, an XCP-ng cluster was deployed and configured with StarWind Virtual SAN (VSAN) running in a CVM on each host. As a result, a StarWind highly available storage was created, connected to all XCP-ng clustered nodes to store virtual machines on a highly available Storage Repository.

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.