

StarWind Virtual SAN[®]

Providing HA storage repositories for XenServer 6.5

APRIL, 2019

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.

Introduction

Traditionally XenServer Pool requires shared storage to guarantee data safety, allow virtual machines migration, enable continuous application availability and eliminate any single point of failure within an IT environment.

NOTE: Hyperconverged scenario is not officially supported with Citrix XenServer.

The network diagram below illustrates the system configuration.



Starwind Vsan For Hyper-V System Requirements

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


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

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


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

Installing Starwind Vsan For Hyper-V

1. Download the StarWind setup executable file from the StarWind website:
<https://www.starwind.com/registration-starwind-virtual-san>
2. Launch the downloaded setup file on the server to install StarWind Virtual SAN or one of its components. The Setup wizard will appear. Read and accept the License Agreement.



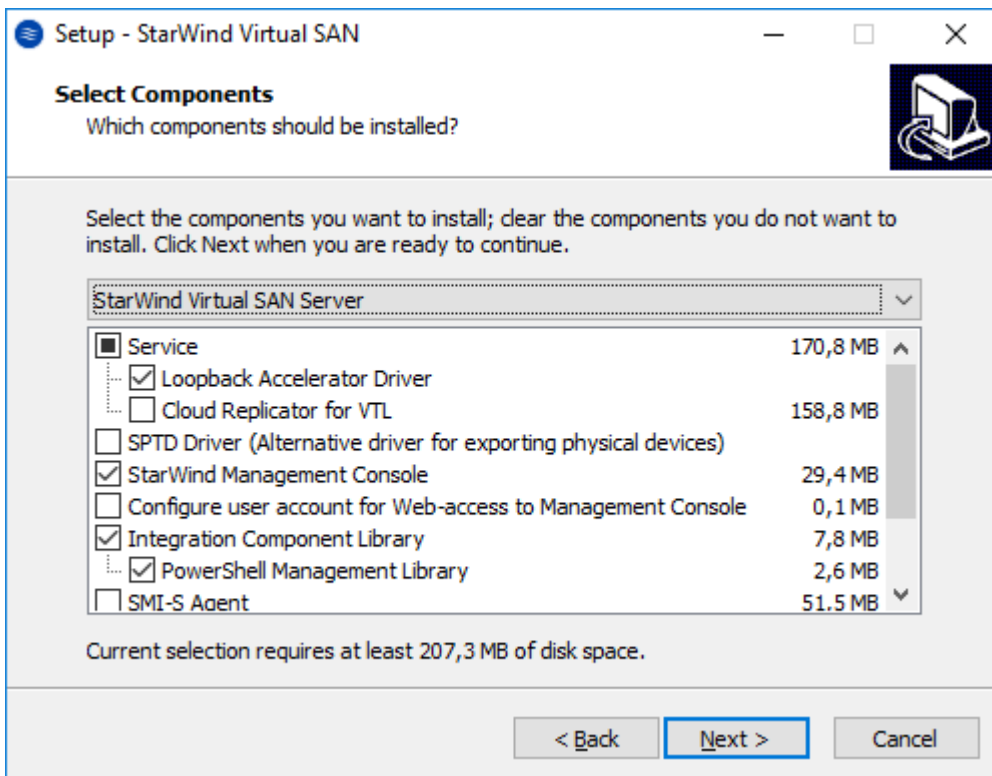
3. Carefully read the information about the new features and improvements. Red text indicates warnings for users that are updating the existing software installations.
4. Select **Browse** to modify the installation path if necessary. Click on **Next** to continue.



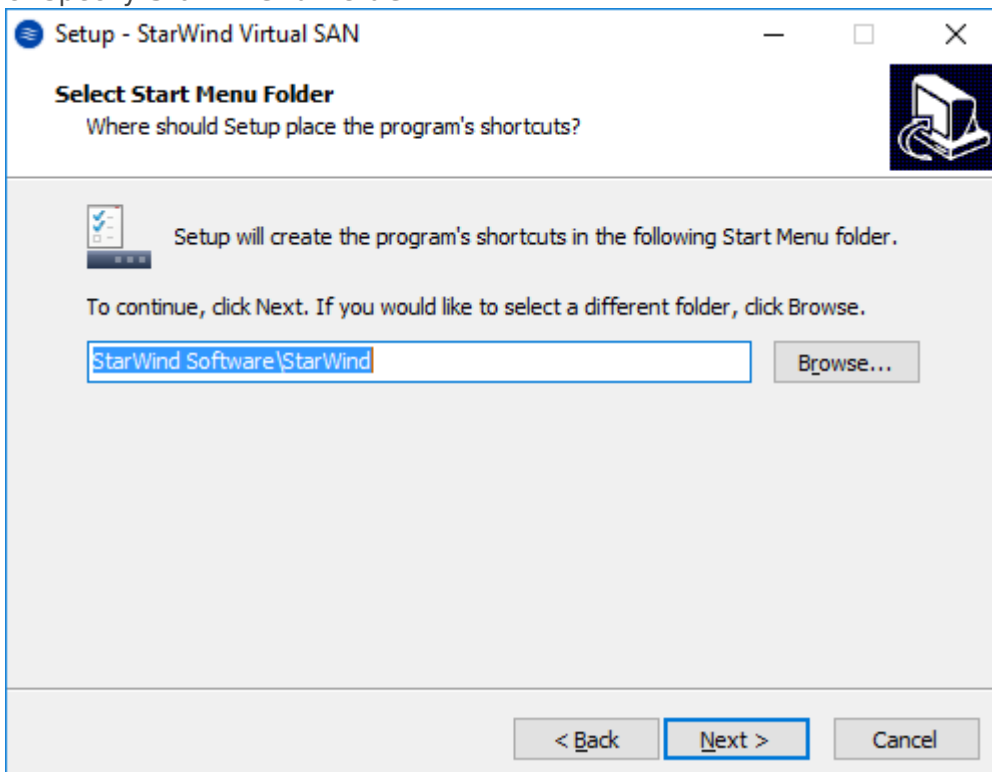
5. Select the following components for the minimum setup:

- **StarWind Virtual SAN Service.** The StarWind Virtual SAN service is the “core” of the software. It can create iSCSI targets as well as share virtual and physical devices. The service can be managed from StarWind Management Console on any Windows computer that is on the same network. Alternatively, the service can be managed from StarWind Web Console deployed separately.
- **StarWind Management Console.** Management Console is the Graphic User Interface (GUI) part of the software that controls and monitors all storage-related operations (e.g., allows users to create targets and devices on StarWind Virtual SAN servers connected to the network).

NOTE: To manage **StarWind Virtual SAN** installed on a Windows Server Core edition with no GUI, StarWind Management Console should be installed on a different computer running the GUI-enabled Windows edition.



6. Specify **Start Menu Folder**.



7. Enable the checkbox if a desktop icon needs to be created. Click on **Next** to continue.

8. When the license key prompt appears, choose the appropriate option:

- Request time-limited fully functional evaluation key.

- Request FREE version key.
 - Select the previously purchased commercial license key.
9. Click on the **Browse** button to locate the license file.
 10. Review the licensing information.
 11. Verify the installation settings. Click on **Back** to make any changes or **Install** to proceed with installation.
 12. Enable the appropriate checkbox to launch **StarWind Management Console** right after the setup wizard is closed and click on **Finish**.
 13. Repeat the installation steps on the partner node.

Creating Starwind Ha Devices

1. Open **Add Device (advanced) Wizard**.
2. Select **Hard Disk Device** as the type of device to be created.
3. Select **Virtual Disk**.
4. Specify a virtual disk **Name**, **Location**, and **Size**.
5. Select the **Thick provisioned** disk type.
6. Define a caching policy and specify a cache size (in MB). Also, the maximum available cache size can be specified by selecting the appropriate checkbox. Optionally, define the L2 caching policy and cache size.
7. Specify **Target Parameters**. Select the **Target Name** checkbox to enter a custom target name. Otherwise, the name is generated automatically in accordance with the specified target alias.
8. Click **Create** to add a new device and attach it to the target.
9. Click **Close** to finish the device creation.
10. Right-click the recently created device and select **Replication Manager** from the shortcut menu.
11. Select the **Add Replica** button in the top menu.
21. The successfully added devices appear in the **StarWind Management Console**.

Select The Required Replication Mode

The replication can be configured in one of two modes:

Synchronous “Two-Way” Replication

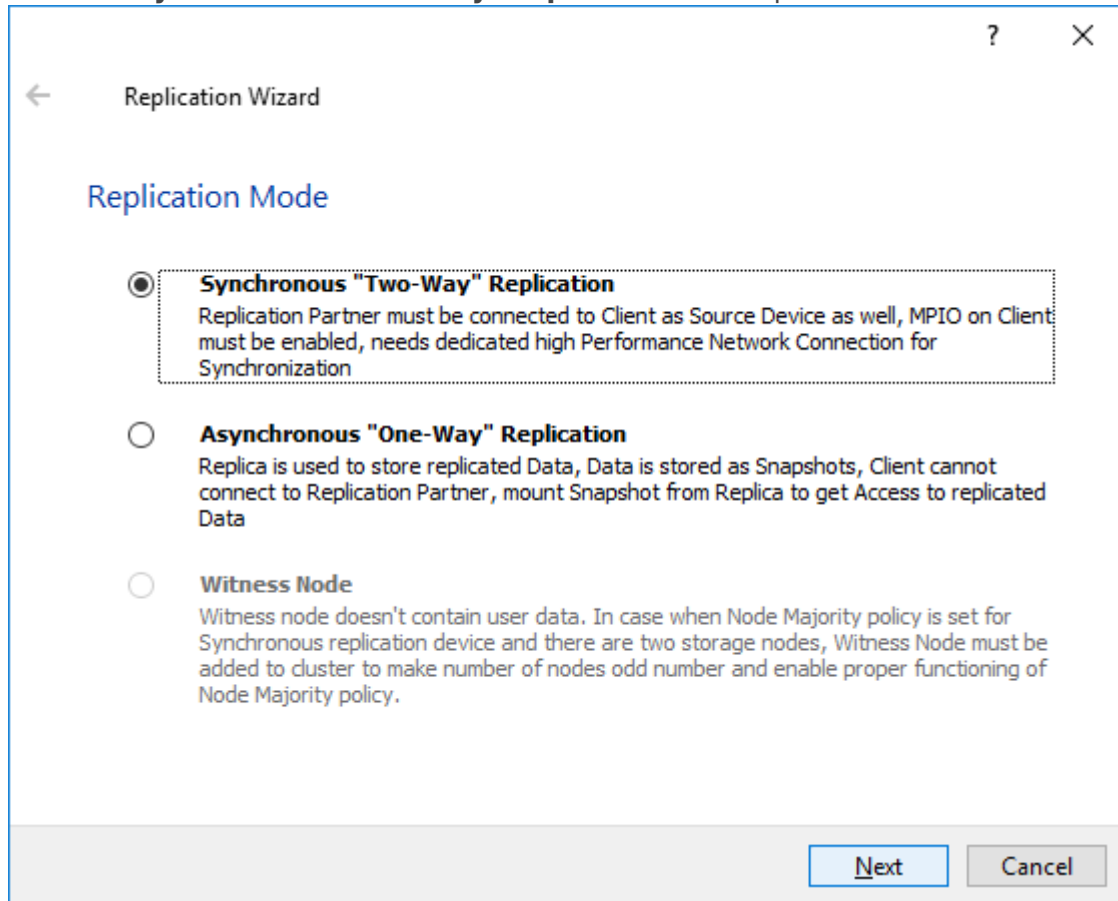
Synchronous or active-active replication ensures real-time synchronization and load balancing of data between two or three cluster nodes. Such a configuration tolerates the failure of two out of three storage nodes and enables the creation of an effective business continuity plan. With synchronous mirroring, each write operation requires control confirmation from both storage nodes. It guarantees the reliability of data transfers but is demanding in bandwidth since mirroring will not work on high-latency networks.

Asynchronous “One-Way” Replication

Asynchronous Replication is used to copy data over a WAN to a separate location from the main storage system. With asynchronous replication, confirmation from each storage node is not required during the data transfer. Asynchronous replication does not guarantee data integrity in case of storage or network failure; hence, some data loss may occur, which makes asynchronous replication a better fit for backup and disaster recovery purposes where some data loss is acceptable. The Replication process can be scheduled in order to prevent the main storage system and network channels overloads. Please select the required option:

Synchronous “Two-Way” Replication

1. Select **Synchronous “Two-Way” replication** as a replication mode.



The screenshot shows a window titled "Replication Wizard" with a back arrow and help/cancel icons. Under the heading "Replication Mode", there are three radio button options:

- Synchronous "Two-Way" Replication** (Selected): Replication Partner must be connected to Client as Source Device as well, MPIO on Client must be enabled, needs dedicated high Performance Network Connection for Synchronization
- Asynchronous "One-Way" Replication**: Replica is used to store replicated Data, Data is stored as Snapshots, Client cannot connect to Replication Partner, mount Snapshot from Replica to get Access to replicated Data
- Witness Node**: Witness node doesn't contain user data. In case when Node Majority policy is set for Synchronous replication device and there are two storage nodes, Witness Node must be added to cluster to make number of nodes odd number and enable proper functioning of Node Majority policy.

At the bottom right, there are "Next" and "Cancel" buttons.

2. Specify a partner **Host name or IP address** and **Port Number**.

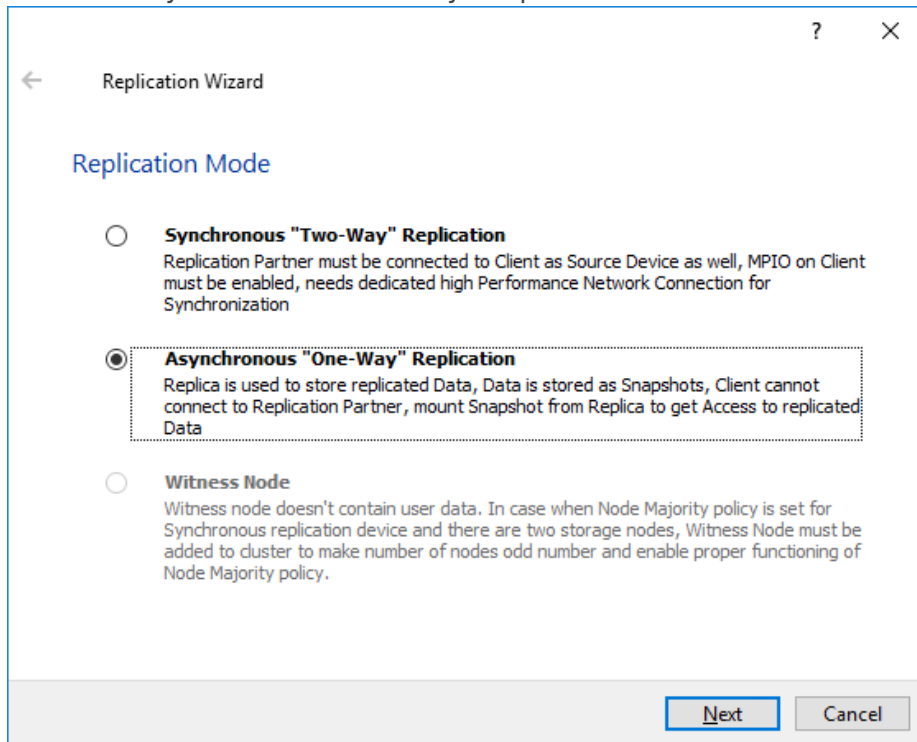
Asynchronous "one-Way" Replication

NOTE: Asynchronous replication requires minimum 100 MbE network bandwidth or higher. The Asynchronous node uses the LSFS device by design. Please, make sure that

the Asynchronous node meets the LSFS device requirements:

<https://knowledgebase.starwindsoftware.com/explanation/lsfs-container-technical-description/>

1. Select Asynchronous “One-Way” Replication.



2. Enter **Host name or IP address** of the Asynchronous node.

3. Choose the **Create New Partner Device** option.

4. Specify the partner device **Location** . Optionally, modify the target name by clicking the appropriate button.

5. Click **Change Network Settings**.

6. Specify the network for asynchronous replication between the nodes. Click **OK** and then click **Next**.

7. In **Select Partner Device Initialization Mode**, select **Synchronize from existing Device** and click **Next**.


8. Specify **Scheduler Settings** and click **Next**.

NOTE: The size of journal files and number of snapshots depends on the settings specified in this step.

9. Specify the path for journal files and click **Next**.

NOTE: By default, the journal files will be located on the node with the original device. However, it is highly recommended not to store journal files on the same drive where the original device is located. Additionally, the C:\ drive should not be used as the path for journal files to avoid any issues with Windows OS.

If the same drive where the StarWind device is located is selected, the warning message about possible performance issues will pop up. If there is no additional volume available for storing the journals, click **I understand the potential problem. Use the selected path**.



10. Press the **Create Replica** button.

11. Wait until StarWind service creates a device and click **Close** to complete the device creation.

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 which participates in the nodes

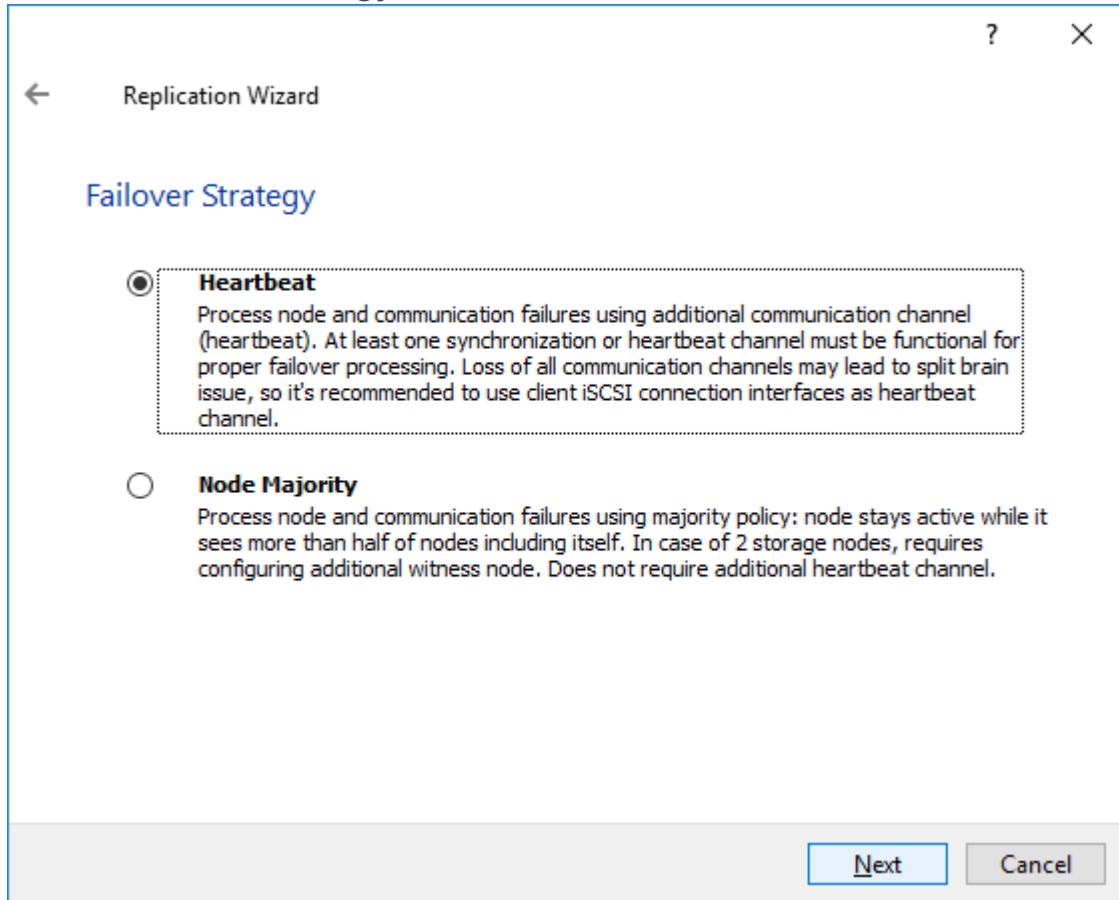
count for the majority, but neither contains data on it nor is involved in processing clients' requests. In case an HA device is replicated between 3 nodes, no Witness node is required.

With Node Majority failover strategy, failure of only one node can be tolerated. If two nodes fail, the third node will also become unavailable to clients' requests.

Please select the required option:

Heartbeat

1. Select **Failover Strategy**.



Replication Wizard

Failover Strategy

Heartbeat
 Process node and communication failures using additional communication channel (heartbeat). At least one synchronization or heartbeat channel must be functional for proper failover processing. Loss of all communication channels may lead to split brain issue, so it's recommended to use client iSCSI connection interfaces as heartbeat channel.

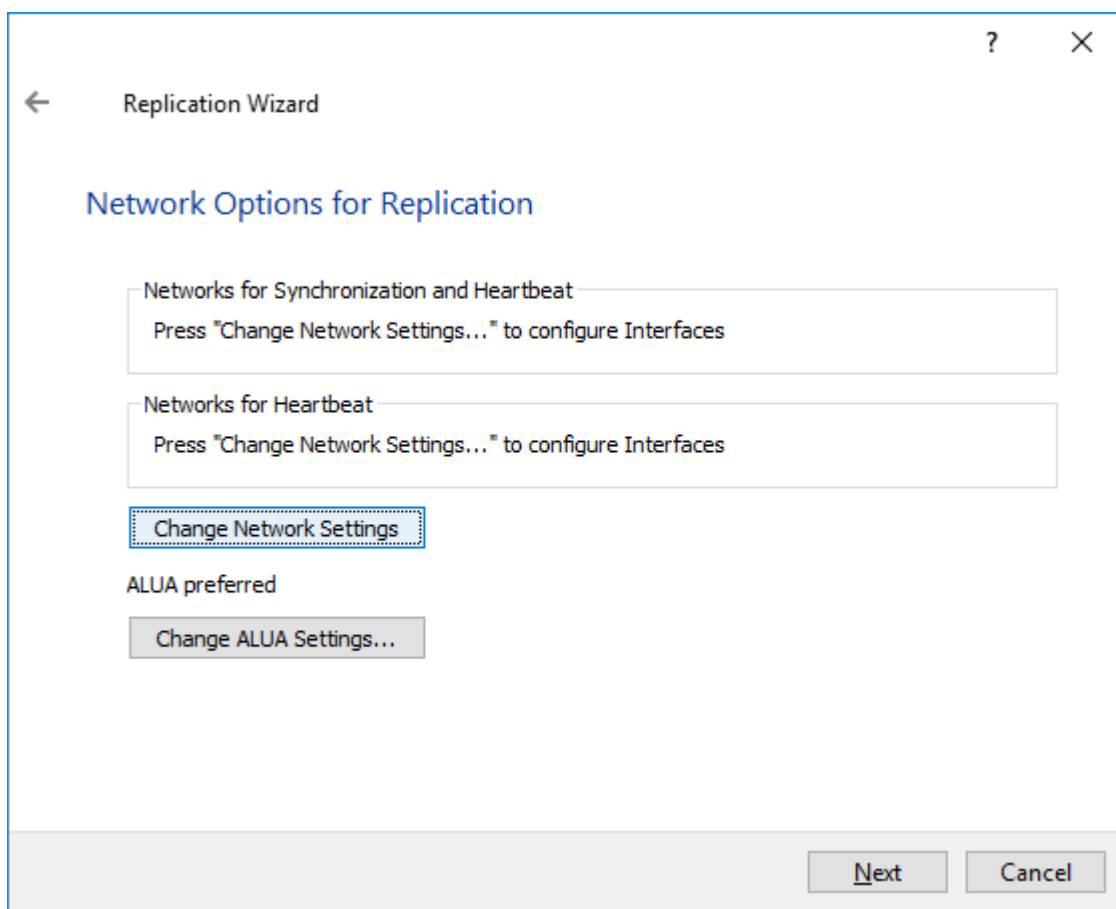
Node Majority
 Process node and communication failures using majority policy: node stays active while it sees more than half of nodes including itself. In case of 2 storage nodes, requires configuring additional witness node. Does not require additional heartbeat channel.

Next Cancel

2. Select **Create new Partner Device** and click **Next**.

3. Select a partner device **Location**.

4. Click **Change Network Settings**.



5. Specify the interfaces for Synchronization and Heartbeat Channels. Click **OK** and then click **Next**.

Specify Interfaces for Synchronization Channels ✕

Select synchronization channel

Interfaces	Networks	Synchronization and H...	Heartbeat
<input type="checkbox"/> Host Name: 127.0.0.1			
172.16.10.10	172.16.10.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
172.16.20.10	172.16.20.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
172.16.30.10	172.16.30.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.40.10	172.16.40.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
192.168.12.10	192.168.12.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="checkbox"/> Host Name: SW2			
172.16.10.20	172.16.10.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
172.16.20.20	172.16.20.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
172.16.30.20	172.16.30.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
172.16.40.20	172.16.40.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>
192.168.12.20	192.168.12.0	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Allow Free Select Interfaces

6. In **Select Partner Device Initialization Mode**, select **Synchronize from existing Device** and click **Next**.

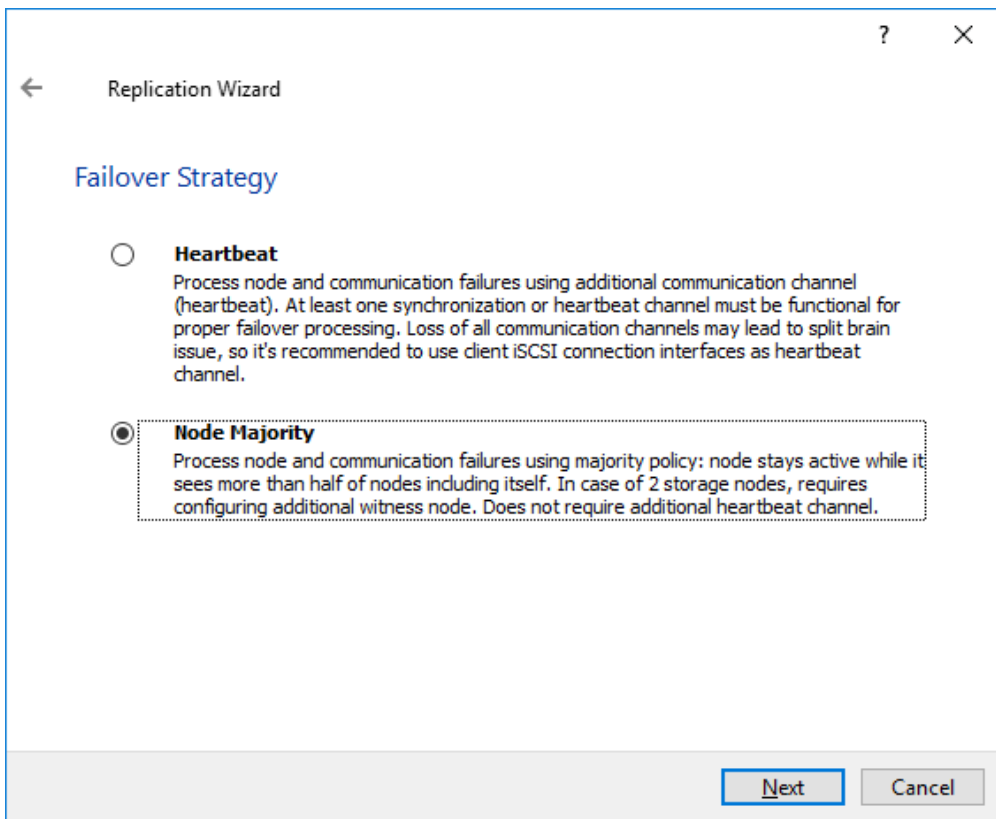
7. Click **Create Replica**. Click **Finish** to close the wizard.

The successfully added device appears in **StarWind Management Console**.

8. Follow the similar procedure for the creation of other virtual disks that will be used as storage repositories.

Node Majority

1. Select the **Node Majority** failover strategy and click **Next**.



2. Choose **Cr**

reate new Partner Device and click **Next**.

3. Specify the partner device **Location** and modify the target name if necessary. Click **Next**.

4. In **Network Options for Replication**, press the **Change network settings** button and select the synchronization channel for the HA device.

5. In **Specify Interfaces for Synchronization Channels**, select the checkboxes with the appropriate networks and click **OK**. Then click **Next**.

6. Select **Synchronize from existing Device** as the partner device initialization mode.

7. Press the **Create Replica** button and close the wizard.

8. The added devices will appear in **StarWind Management Console**.


Repeat the steps above to create other virtual disks if necessary.

Adding Witness Node

Witness node can be configured on a separate host or as a virtual machine in a cloud. It requires StarWind Virtual SAN service installed on it.


NOTE: Since the device created in this guide is replicated between 2 active nodes with the Node Majority failover strategy, a Witness node must be added to it.

1. Open **StarWind Management Console**, right-click on the Servers field and press the **Add Server** button. Add a new StarWind Server which will be used as the Witness node and click **OK**.

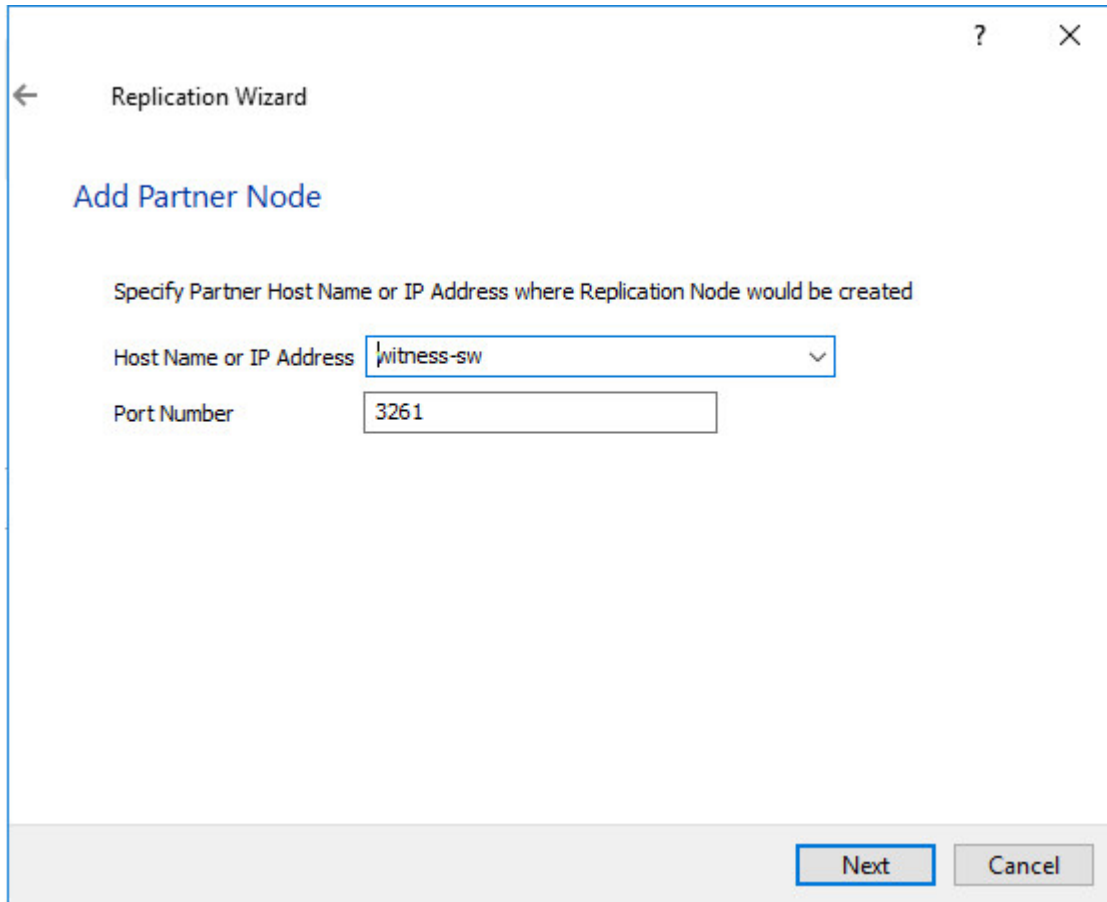


2. Right-click on the HA device with the configured Node Majority failover policy and select **Replication Manager** and press the **Add Replica** button.

3. Select **Witness Node**.



4. Specify the Witness node **Host Name or IP address**. The default Port Number is 3261.



5. In **Partner Device Setup**, specify the Witness device **Location**. Optionally, modify the target name by clicking the appropriate button.

6. In **Network Options for Replication**, select the synchronization channel with the Witness node by clicking the **Change Network Settings** button.

7. Specify the interface for **Synchronization and Heartbeat** and click **OK**.

8. Click **Create Replica** and then close the wizard.

9. Repeat the steps above to create other virtual disks if necessary.

NOTE: To extend an Image File or a StarWind HA device to the required size, please check the article below:

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

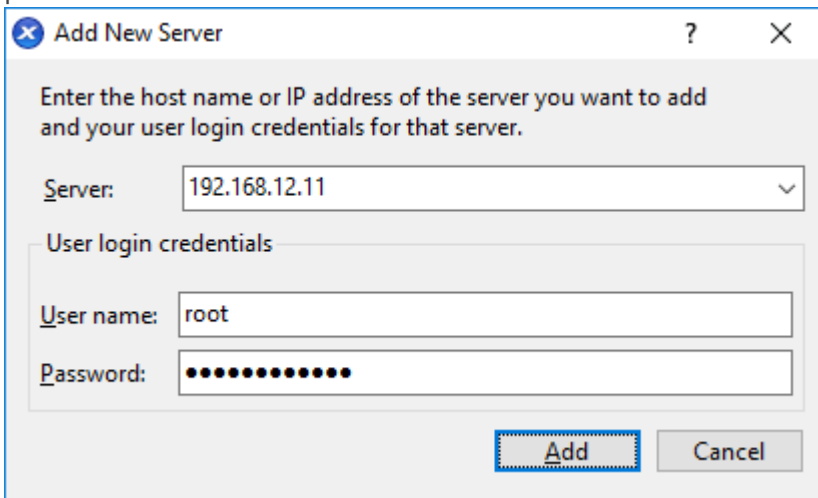
Configuring Hosts With Xenserver 6.5

1. Launch **XenCenter**.

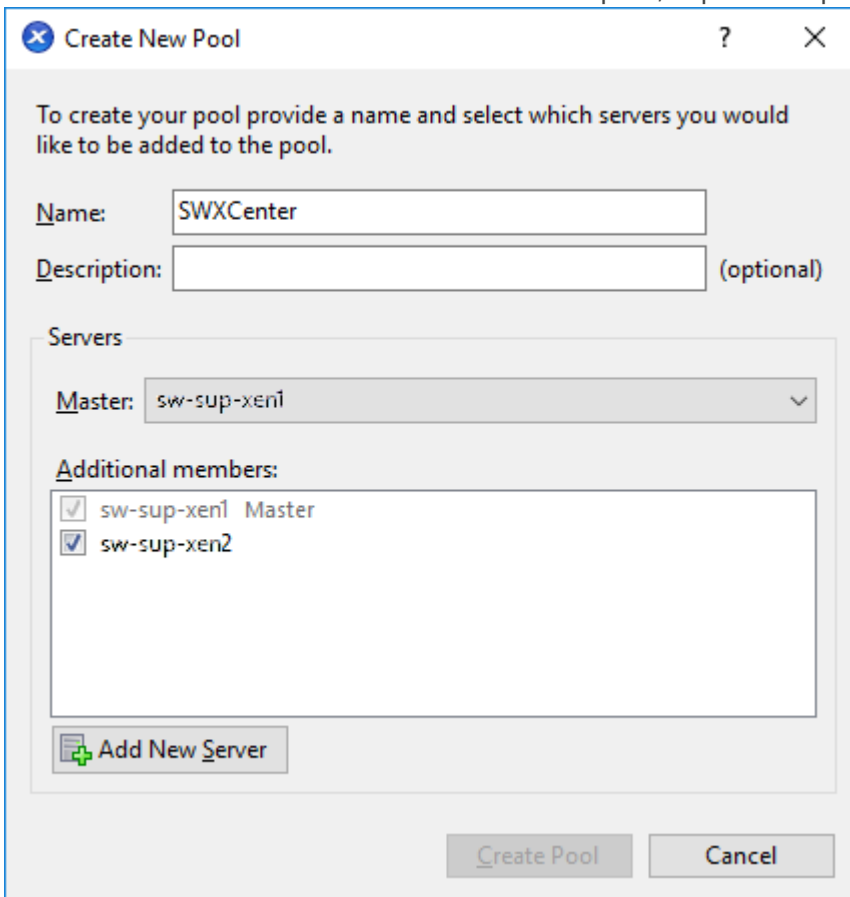
2. Click **New Pool**.

3. Specify the name and description of the pool and click **Add New Server**.

4. In the **Add New Server** window, specify the first XenServer host to be added to the pool as well as **User name** and **Password**.



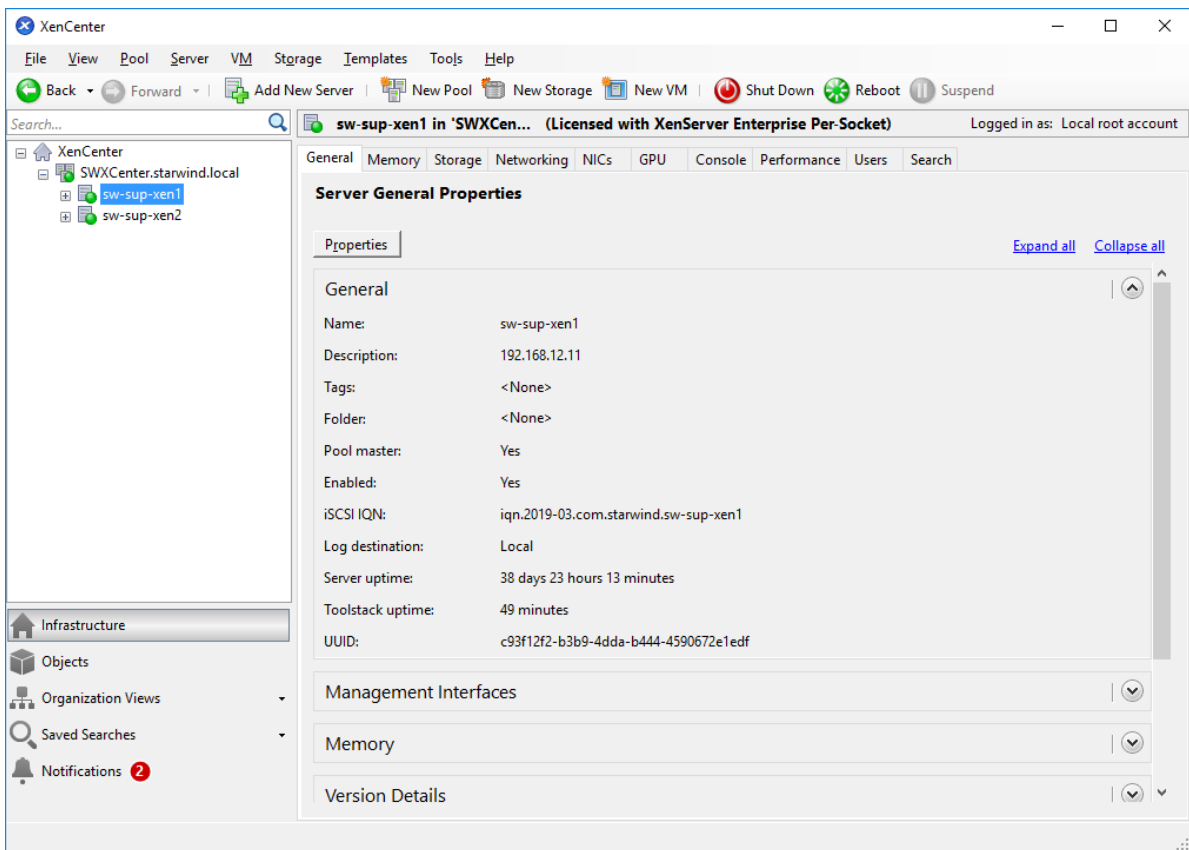
5. To add the second XenServer host to the pool, repeat the previous steps.



6. Specify the **Master** host.


7. Click **Create Pool**.

8. After creating the pool, the **XenCenter** window should look as shown in the screenshot below.

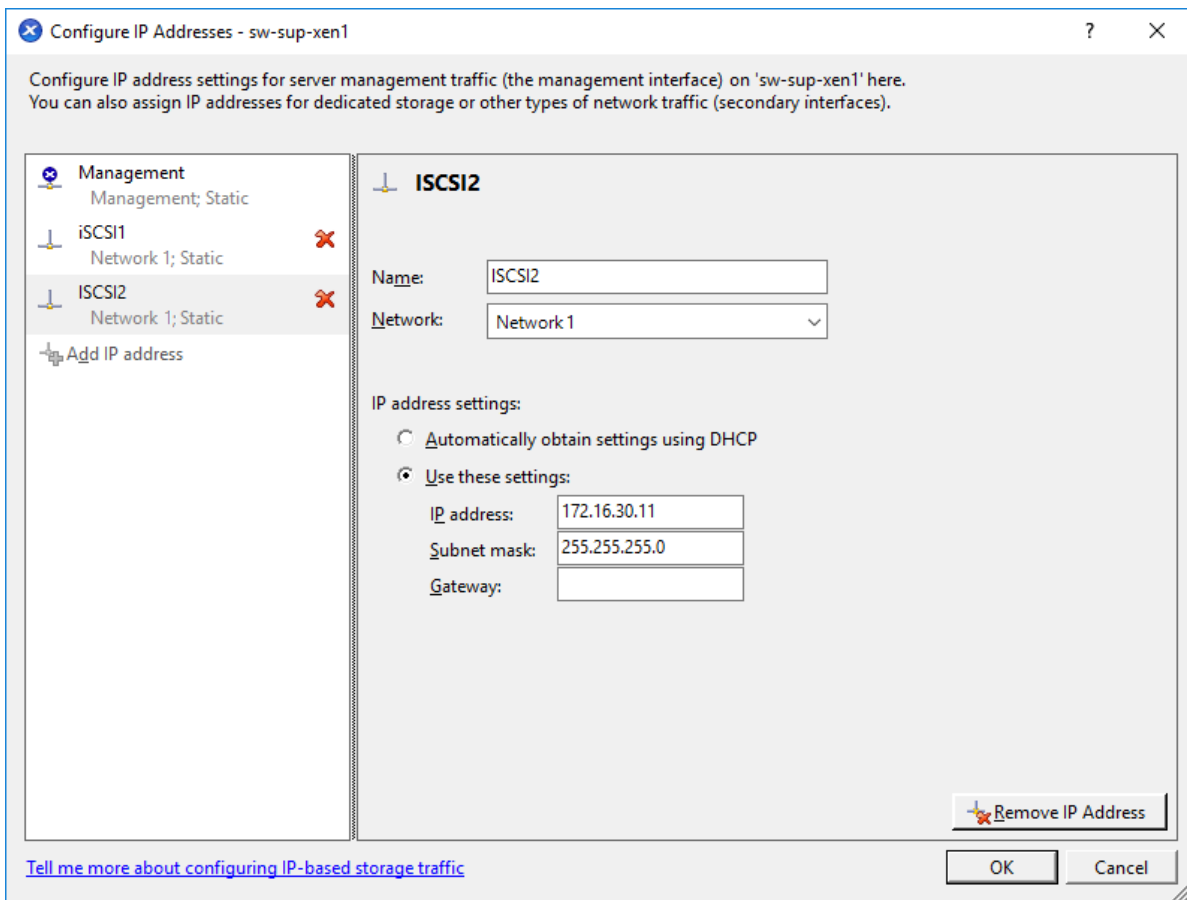


Configuring Network Settings

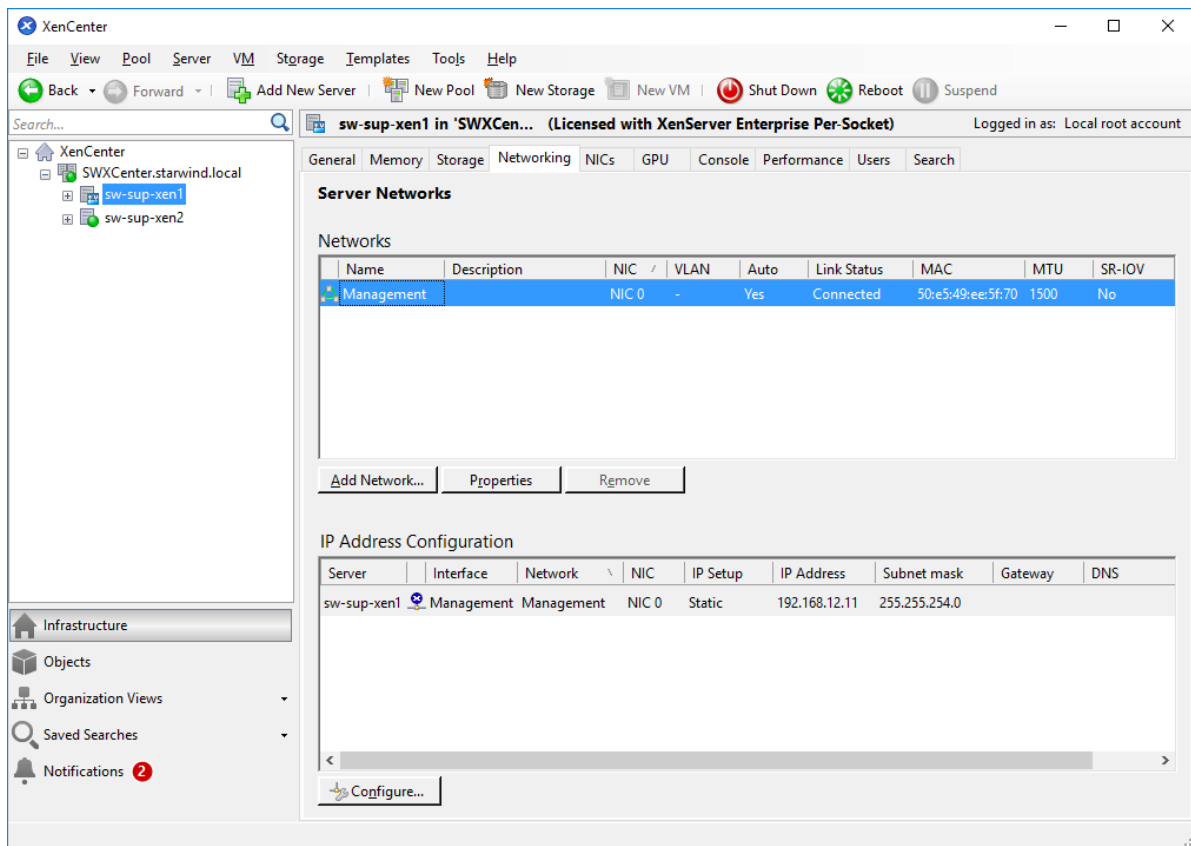
NOTE: To switch the XenServer host to Maintenance Mode for network configuration, right-click the XenServer and select **Enter Maintenance Mode...** from the shortcut menu.



1. Select the first XenServer host, switch to the **Networking** tab, and click **Configure**.
2. In the appeared window, click **Add IP address** to add networks.
3. Configure the required IP address and click **OK**.



4. The window should look as shown below.



5. Exit the Maintenance mode.
6. Repeat the same steps on all XenServer hosts.

Configuring iSCSI Initiator

NOTE: The following actions should be performed on each XenServer host in the pool.

1. Switch XenServer host to the Maintenance mode for iSCSI configuration.
2. Edit `/etc/iscsi/initiatorname.iscsi` and set the appropriate iqn for XenServer iSCSI initiator.
3. Edit the `/etc/iscsi/iscsid.conf` file.

Set the following values:

```
node.startup = automatic
node.session.iscsi.FirstBurstLength = 262144
node.session.iscsi.MaxBurstLength = 262144
```

4. Restart the software iSCSI service:

```
service iscsi restart
service iscsid restart
```

5. Set the iSCSI services to start automatically when launching XenServer:

```
chkconfig iscsi on
```

```
chkconfig iscsid on
```

Applying Multipathing

NOTE: The following actions should be performed on each XenServer host in the pool.

1. Switch XenServer host to maintenance mode for multipath configuration
2. Switch to console tab, edit the multipath configuration file by running “**vi /etc/multipath.conf**” command.
3. Insert the following block into the “**devices**” section:

```
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"
}
```

4. Restart the multipath service:


```
service multipathd restart
```

```
service multipathd reload
```


5. Set the multipath service to start automatically when launching XenServer:

```
chkconfig multipathd on
```

6. Right-click the **XenServer** and select **Properties** from the shortcut menu.



7. Click **Multipathing Active** in the left pane. Select the **Enable multipathing on this server** checkbox and click **OK**.



8. Turn off the Maintenance mode: right-click the XenServer and select **Exit Maintenance Mode** from the shortcut menu.

9. Reboot the XenServer host.

Adding Shared Storage manually

1. Select the XenServer host. Switch to the **Console** tab in XenCenter.

2. Run discovery:

```
iscsiadm -m discovery -t st -p 172.16.30.10
iscsiadm -m discovery -t st -p 172.16.40.10
iscsiadm -m discovery -t st -p 172.16.30.20
iscsiadm -m discovery -t st -p 172.16.40.20
```

3. Connect the targets using their IQNs found by discovery:

```
iscsiadm -m node -T iqn.2008-08.com.starwindsoftware:sw1-ds1 -p
172.16.30.10 -l
```

```
iscsiadm -m node -T iqn.2008-08.com.starwindsoftware:sw1-ds1 -p
172.16.40.10 -l
```

```
iscsiadm -m node -T iqn.2008-08.com.starwindsoftware:sw2-ds1 -p
172.16.30.20 -l
```

```
iscsiadm -m node -T iqn.2008-08.com.starwindsoftware:sw2-ds1 -p
172.16.40.20 -l
```

4. Perform the steps 1-3 for all XenServer hosts in the Pool.

5. On **Pool Master**, run the following commands to add SR:

```
xe sr-probe type=lvmoiscsi device-
config:target=172.16.30.10,172.16.40.10,172.16.30.20,172.16.40.20
device-config:targetIQN=*

xe sr-create type=lvmoiscsi \
shared=true \
name-label="sw-sr" \
device-
config:target=172.16.30.10,172.16.40.10,172.16.30.20,172.16.40.20 \
device-config:targetIQN=* \
device-
config:multiSession="172.16.30.10,3260,iqn.2008-08.com.starwindsoft
ware:sw1-
ds1|172.16.40.10,3260,iqn.2008-08.com.starwindsoftware:sw1-
ds1|172.16.30.20,3260,iqn.2008-08.com.starwindsoftware:sw2-
ds1|172.16.40.20,3260,iqn.2008-08.com.starwindsoftware:sw2-ds1|" \
```



```
device-  
config:multihomelist=172.16.30.10:3260,172.16.40.10:3260,172.16.30.  
20:3260,172.16.40.20:3260 \  
device-config:port=3260 \  
device-config:SCSIid=20cfc3b5a7c6bcef2
```

6. In the **Console** tab, check if the multipath works properly on both hosts.

Contacts

US Headquarters	EMEA and APAC
 1-617-449-77 17	 +44 203 769 18 57 (UK)
 1-617-507-58 45	 +34 629 03 07 17
 1-866-790-26 46	(Spain and Portugal)

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. 35 Village Rd., Suite 100, Middleton, MA 01949 USA

www.starwind.com

©2019, StarWind Software Inc. All rights reserved.