Chelsio Communications

Accelerate

Chelsio Native Network Driver with SR-IOV support for VMware ESXi 6.X

Installation and User's Guide

Chelsio Communications (Headquarters)
209 North Fair Oaks Avenue,
Sunnyvale, CA 94085
U.S.A

www.chelsio.com

Tel:408.962.3600
Fax: 408.962.3661

Chelsio (India) Private Limited
Subramanya Arcade, Floor 3, Tower B
No. 12, Bannerghatta Road,
Bangalore-560029
Karnataka,
India

Tel: +91-80-4039-6800

Chelsio KK (Japan)
Yamato Building 8F,
5-27-3 Sendagaya,
Shibuya-ku,
Tokyo 151-0051,
Japan

Sales
For all sales inquiries please send email to sales@chelsio.com

Support
For all support related questions please send email to support@chelsio.com

Version History

Version	Revision Date
1.0.0	03/24/2016
1.0.1	10/31/2016
1.0.2	11/11/2016
1.0.3	12/05/2016
1.0.4	12/30/2016
1.0.5	01/27/2017
1.0.6	02/24/2017

TABLE OF CONTENTS

1. Introduction

Thank you for choosing Chelsio Unified Wire adapters. These high speed, single chip, single firmware cards provide enterprises and datacenters with high performance solutions for various Network and Storage related requirements.

The **Terminator** series of adapters is Chelsio's next generation of highly integrated, hyper-virtualized 1/10/25/40/50/100GbE controllers. The adapters are built around a programmable protocol-processing engine, with full offload of a complete Unified Wire solution comprising NIC, TOE, iWARP RDMA, iSCSI, FCoE and NAT support. It scales to true 100Gb line-rate operation from a single TCP connection to thousands of connections, and allows simultaneous low latency and high bandwidth operation thanks to multiple physical channels through the ASIC.

Ideal for all data, storage and high performance clustering applications, Chelsio adapters enable a unified fabric over a single wire by simultaneously running all unmodified IP sockets, Fibre Channel and InfiniBand applications over Ethernet at line rate.

Designed for deployment in virtualized datacenters, cloud service installations and high performance computing environments, Chelsio adapters bring a new level of performance metrics and functional capabilities to the computer networking industry.

This document describes the installation, use and maintenance of Native Network driver for VMware ESXi and its various components.

1.1. Features

Chelsio's Native Network driver for VMware is developed for Chelsio's Unified Wire family of adapters.

1.2. Hardware Requirements

1.2.1. Supported Adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio Native Network driver:

- T62100-CR
- T62100-LP-CR
- T62100-SO-CR
- T6225-CR
- T6225-SO-CR
- T580-SO-CR
- T580-LP-CR
- T580-CR

- T540-CR
- T520-SO-CR
- T520-CR
- T520-LL-CR
- T520-BT

1.2.2. SR-IOV Requirements

- SR-IOV should be enabled in the BIOS.
- Intel Virtualization Technology for Directed I/O (VT-d) should be enabled in the BIOS.
- PCI Express Slot should be ARI capable.

1.3. Software Requirements

1.3.1. ESXi Requirements

The Chelsio Native Network driver has been developed to run on ESXi platforms. Currently the driver is available for the following versions:

- **Host:**
  - ESXi 6.5a
  - ESXi 6.5
  - ESXi 6.0
  - ESXi 6.0 U1
  - ESXi 6.0 U2

- **Virtual Machine (with VFs):**
  - RHEL 7.3, 3.10.0-514.el7
  - RHEL 7.2, 3.10.0-327.el7
  - RHEL 6.8, 2.6.32-642.el6
  - RHEL 6.7, 2.6.32-573.el6
  - SLES 12 SP2, 4.4.21-69-default
  - SLES 12 SP1, 3.12.49-11-default
  - SLES 11 SP4, 3.0.101-63-default
  - Ubuntu 16.04.1, 4.4.0-31-generic
  - Ubuntu 14.04.4, 4.2.0-27-generic
  - Kernel.org linux-4.9
  - Kernel.org linux-4.8

Note *Windows Guest is not supported with SR-IOV.*

1.4. Package Contents

The Chelsio Native Network driver package consists of the following files/directories:

- **cxl-*.vib**: Driver VIB file required for installation.
- **docs:** This directory contains support documents - README, Release Notes and User's Guide (this document) for the software.
- **EULA:** Chelsio's End User License Agreement.

2. Hardware Installation

Follow these steps to install Chelsio Adapter in your system:

i. Shutdown/power off your system.
ii. Power off all remaining peripherals attached to your system.
iii. Unpack the Chelsio adapter and place it on an anti-static surface.
iv. Remove the system case cover according to the system manufacturer's instructions.
v. Remove the PCI filler plate from the slot where you will install the Ethernet adapter.
vi. For maximum performance, it is highly recommended to install the adapter into a PCIe x8/x16 slot.
vii. Holding the Chelsio adapter by the edges, align the edge connector with the PCI connector on the motherboard. Apply even pressure on both edges until the card is firmly seated. It may be necessary to remove the transceiver modules prior to inserting the adapter.
viii. Secure the Chelsio adapter with a screw, or other securing mechanism, as described by the system manufacturer's instructions. Replace the case cover.
ix. After securing the card, ensure that the card is still fully seated in the PCIE x8/x16 slot as sometimes the process of securing the card causes the card to become unseated.
x. Connect a fiber/twinax cable, multi-mode for short range (SR) optics or single-mode for long range (LR) optics, to the Ethernet adapter or regular Ethernet cable for the 1Gb Ethernet adapter.
xi. Power on your system.
xii. Verify if the adapter was installed successfully by using the *lspci* command

```
[root@wayne ~]# lspci | grep -i Chelsio
06:00.0 Ethernet controller: Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller
06:00.1 Ethernet controller: Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller
06:00.2 Ethernet controller: Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller
06:00.3 Ethernet controller: Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller
06:00.4 Ethernet controller: Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller
06:00.5 SCSI storage controller: Chelsio Communications Inc T6225-CR Unified Wire Storage Controller
06:00.6 Fibre Channel: Chelsio Communications Inc T6225-CR Unified Wire Storage Controller
```

For Chelsio adapters, the physical functions are currently assigned as:

- Physical functions 0 - 3: for the SR-IOV functions of the adapter
- Physical function 4: for all NIC functions of the adapter
- Physical function 5: for iSCSI
- Physical function 6: for FCoE
- Physical function 7: Currently not assigned

Note

Network device names for Chelsio's physical ports are assigned using the following convention: the port farthest from the motherboard will appear as the first network interface. However, for T5 40G adapters, the association of physical Ethernet ports and their corresponding network device names is opposite. For these adapters, the port nearest to the motherboard will appear as the first network interface.

3. Software/Driver Installation

i. Download the driver package from Chelsio Download Center, http://service.chelsio.com.
ii. Unzip the driver package:

```
[root@host:~] unzip <driver_package>.zip
```

iii. Put the host in maintenance mode using the vSphere (desktop or web) Client.

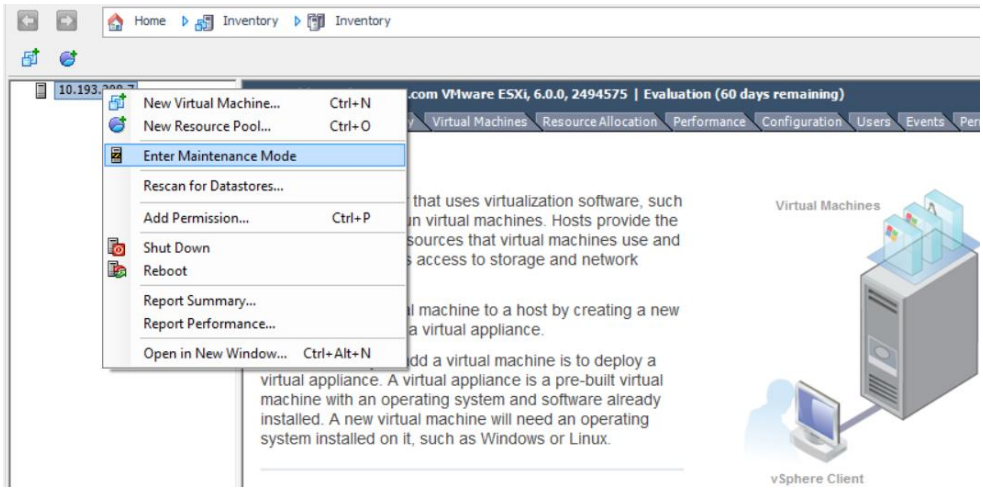

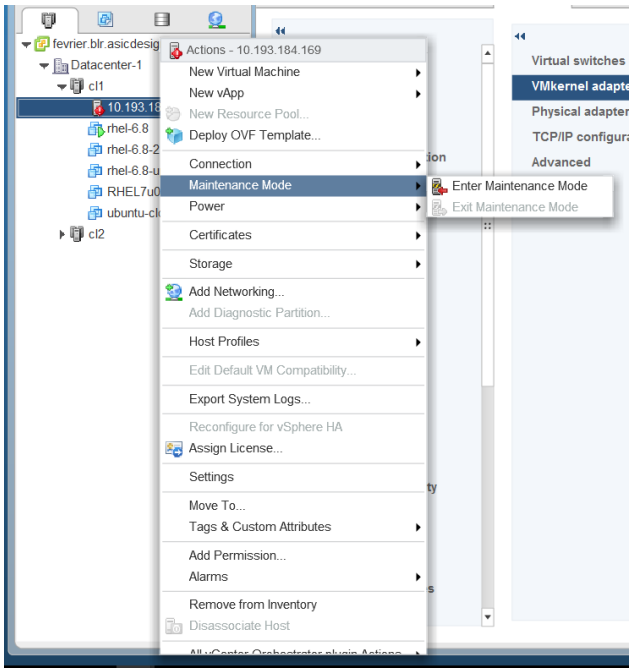

iv. Install the Native Network driver (*cxl*):

```
[root@host:~] cp cxl-*.vib /productLocker/
[root@host:~] cp cxl-*.vib /var/log/vmware/
[root@host:~] esxcli software vib install -v /productLocker/cxl-*.vib --no-
sig-check
```

v. After installation completes successfully, exit from maintenance mode and reboot the host.

4. Software/Driver Loading

After rebooting the ESXi Host, the driver will load automatically. However, it is possible to manually load the driver by using the command below:

```
[root@host:~] vmkload_mod cxl
```

Execute the below command so that device manager performs a rescan:

```
[root@host:~] kill -SIGHUP $(cat /var/run/vmware/vmkdevmgr.pid)
```

5. Software/Driver Configuration and Fine-tuning

5.1. cxgbtool

The *cxgbtool* command queries or sets various aspects of Chelsio network interface cards. It complements standard tools used to configure network settings and provides functionality not available through such tools. Some of the commands provided can be used to query running statistics to aid in debugging. The tool will be installed by default on installing the driver VIB.

Syntax & Usage

To use cxgbtool, use the syntax:

```
[root@host:~] /opt/chelsio/bin/cxgbtool <options>
```

Note *For information on available parameters and their usage, refer to cxgbtool help by running the* `/opt/chelsio/bin/cxgbtool -h` *command.*

5.2. Adapter Configuration

The adapter's configuration should be updated for optimal performance in ESXi environment. Follow the steps mentioned below:

i. Run the following *cxgbtool* command:

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c esxcfg -set
```

Note *Not supported on T6 adapters.*

ii. Reboot the system for changes to take effect.

5.3. Firmware update

The driver will auto-load the firmware if an update is required. The version can be verified using:

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c version
```

5.4. Connecting a Virtual Machine

Follow the steps mentioned below to connect Chelsio adapter to a virtual machine:

i. Create a new virtual switch:

```
[root@host:~] esxcfg-vswitch -a vSwitchN
```

ii. Link a Chelsio adapter to the newly created virtual switch:

```
[root@host:~] esxcfg-vswitch -L vmnicN vSwitchN
```

iii. Create a new port group on the vSwitch:

```
[root@host:~] esxcfg-vswitch -A <port group name> vSwitchN
```

iv. From the vSphere client, right-click on the virtual machine, select the virtual network adapter to be used, and attach the newly created port group.

5.5. Tuning vMotion for vSAN

While creating a vSAN you may encounter out of memory issues, resulting in vMotion timing out. To avoid this, follow the steps mentioned below on all ESXi hosts in the vSAN:

i. Set the low-memory packet heap value to the maximum supported value of 512:

```
[root@host:~] esxcli system settings kernel set -s netPktHeapMaxMBPerGB -v
512
```

ii. Set the networking packet buffer pool size to the maximum supported value of 200:

```
[root@host:~] esxcli system settings kernel set -s netPktPoolMaxMBPerGB -v
200
```

iii. Reboot the host.

iv. Enable maximum RSS queues:

```
[root@host:~] vmkload_mod cxl rss_q=1
```

5.6. Virtual Functions (SR-IOV)

5.6.1. Instantiate VFs

Follow the steps mentioned below to instantiate virtual functions:

i. *max_vfs* is a comma separated module parameter that specifies the maximum number of VFs per port. Load the Native Network driver (cxl) with *max_vfs* parameter and set it to a non-zero value. In case of multiple adapters, use ',,' to separate the number of VFs per adapter:

```
[root@host:~] esxcfg-module cxl -s max_vfs=W,X,,Y,Z
```

Where,

W: Number of VFs per port 0 of adapter 0.
X: Number of VFs per port 1 of adapter 0.
Y: Number of VFs per port 0 of adapter 1.
Z: Number of VFs per port 1 of adapter 1.

Note *A maximum of 16 VFs can be instantiated per port*.

E.g. - To instantiate 3 VFs for port 1 of adapter 0 & 4 VFs for port 0 of adapter 1:

```
[root@host:~] esxcfg-module cxl -s max_vfs=0,3,,4,0
```

ii. Verify *max_vfs* setting using the *-g* option:

```
[root@host:~] esxcfg-module -g cxl
```

E.g. :

```
[root@host:~] esxcfg-module -g cxl
cxl enabled = 1 options = 'max_vfs=2,2'
```

iii. Reboot the ESXi host for changes to take effect.
iv. Check if VFs were instantiated successfully on the PCI bus by either using the shell prompt (using *lspci*) or vSphere GUI (under *Host > Configuration>Advanced setting*)

```
[root@austen:~] lspci | grep Chelsio
0000:05:00.0 Network controller: Chelsio Communications Inc. T580-LP-CR Unified Wire Ethernet Controller [vmnic4]
0000:05:00.1 Network controller: Chelsio Communications Inc. T580-LP-CR Unified Wire Ethernet Controller [vmnic5]
0000:05:00.2 Network controller: Chelsio Communications Inc. T580-LP-CR Unified Wire Ethernet Controller [vmnic6]
0000:05:00.3 Network controller: Chelsio Communications Inc. T580-LP-CR Unified Wire Ethernet Controller [vmnic7]
0000:05:00.4 Network controller: Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller
0000:05:00.5 Mass storage controller: Chelsio Communications Inc T580-LP-CR Unified Wire Storage Controller
0000:05:00.6 Serial bus controller: Chelsio Communications Inc T580-LP-CR Unified Wire Storage Controller
0000:05:01.0 Network controller: Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller [PF_0.5.0_VF_0]
0000:05:01.1 Network controller: Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller [PF_0.5.1_VF_0]
0000:05:01.4 Network controller: Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller [PF_0.5.0_VF_1]
0000:05:01.5 Network controller: Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller [PF_0.5.1_VF_1]
```

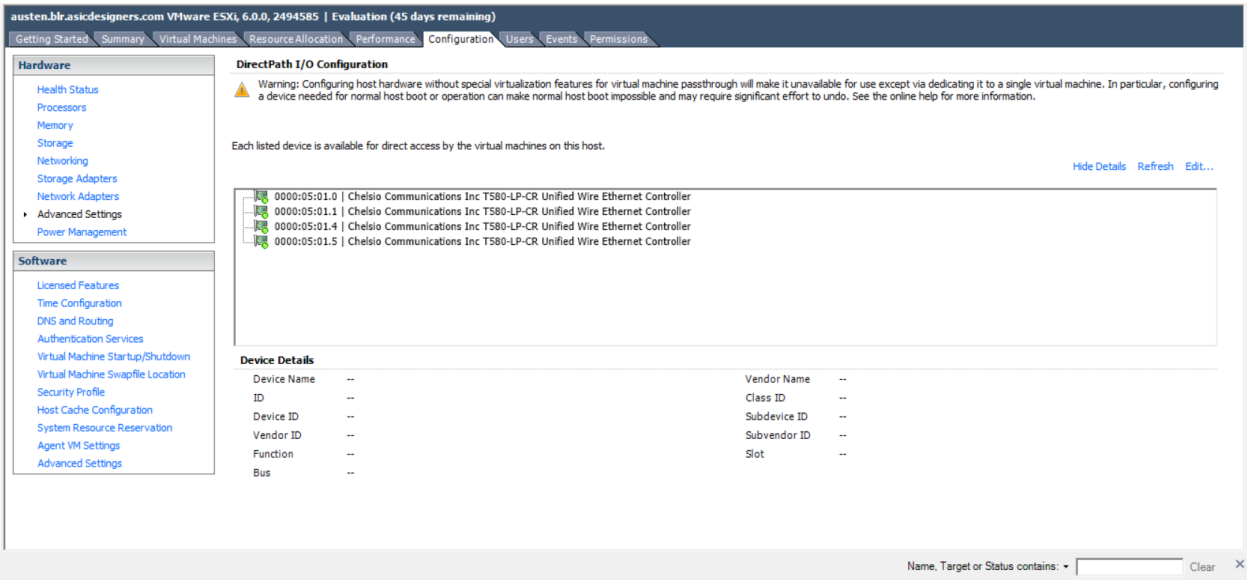

Note

- *Unloading driver when VFs are attached to VMs is not supported by VMware.*
- *VMs with SRIOV interface might not power on with "out of MSI-X vectors" message in vmkernel.log. To resolve this issue, you need to add "pciPassthru<VF_ID>.maxMSIXvectors" parameter to VMs configuration file. Maximum value allowed for this param is 31. It is recommended to set the value according to the following equation:*
  pciPassthru<VF_ID>.maxMSIXvectors = <Number of CPUs in Win VM> + 2

  For more information refer to VMware documentation.

- Windows Guest is not supported with SR-IOV.

5.6.2. Assigning VFs to VMs

Once the VFs are instantiated successfully, it's time to attach them to the virtual machine. For instructions on how to assign virtual functions to a virtual machine, please refer to VMware's official documentation.

5.6.3. Using VFs in Linux VM

To use the newly attached VFs in a virtual machine, follow the steps mentioned below:

i. Power-on the Virtual Machine with VF attached to it.

ii. Verify that the Chelsio VF shows up in the VM using the *lspci* command. You should see a similar output:

```
[root@host~]# lspci | grep Chelsio
00:08:0 Ethernet controller: Chelsio Communications Inc T580-LP-CR Unified
Wire Ethernet Controller [VF]
```

iii. Download the latest *Chelsio Unified Wire for Linux* driver package, from Chelsio Download Center, http://service.chelsio.com/

iv. Untar the tarball using the following command:

```
[root@host~]# tar zxvfm ChelsioUwire-x.xx.x.x.tar.gz
```

v. Change you current working directory to *ChelsioUwire-x.xx.x.x* directory and install the VF driver using the following command:

```
[root@host~]# make vnic_install
```

vi. Load the VF driver in the VM using the below command:

```
[root@host~]# modprobe cxgb4vf
```

Note *To know more about Chelsio Virtual Function driver, please refer Chelsio Unified Wire for Linux User's Guide.*

5.6.4. Example

i. In this example, 2 VFs are instantiated per port, hence a total of 4 VFs instantiated on the host. The host is then rebooted.

```
[root@host:~] esxcfg-module cxl -s max_vfs=2,2
[root@host:~] reboot
```

ii. 4 VMs are setup in the following combination:
- VF0 of PF0 (VF marked with the bus-id <PCIslot.01.0>) is assigned to VM1
- VF1 of PF0 (VF marked with the bus-id <PCIslot.01.4>) is assigned to VM2
- VF0 of PF1 (VF marked with the bus-id <PCIslot.01.1>) is assigned to VM3
- VF1 of PF1 (VF marked with the bus-id <PCIslot.01.5>) is assigned to VM4

iii. VMs are powered up one after another.

iv. VF driver (*cxgb4vf)* is installed and loaded in all the VMs.

The above configuration will result in the following connectivity:

- VFs of the same port can communicate with each other. i.e. VM1 can communicate with VM2, and VM3 can communicate with VM4.
- VFs of port 0 (VM1 and VM2) will be able to communicate with any peer connected to port 0 of the network adapter.
- VFs of port 1 (VM3 and VM4) will be able to communicate with any peer connected to port 1 of the network adapter.

5.7. VXLAN

Virtual Extensible LAN (VXLAN) is a network virtualization technique that uses overlay encapsulation protocol to provide Ethernet Layer 2 network services with extended scalability and flexibility. VXLAN extends the virtual LAN (VLAN) address space by adding a 24-bit segment ID and increasing the number of available logical networks from 4096 to 16 million, thereby addressing the scalability and network segmentation issues associated with large cloud computing deployments. Furthermore, VXLAN provides a cost effective software-defined networking (SDN) solution for migration of a large number of VMs over large distances using existing hardware and software resources.

The following sections describe the method to setup a VXLAN using Chelsio adapter:

5.7.1. ESXi 6.5a

- **Configuring Adapters**

For T6 adapters, VXLAN offload is enabled by default on loading the driver. In case of T5 adapters, this feature is enabled for Tx but disabled for Rx. To enable it, load the driver with *vxlan_rx_offload=1* option:

```
[root@host:~] vmkload_mod cxl vxlan_rx_offload=1
```

- **Creating Cluster**

i. Log in to vCenter Server through vSphere Web Client using a web browser.

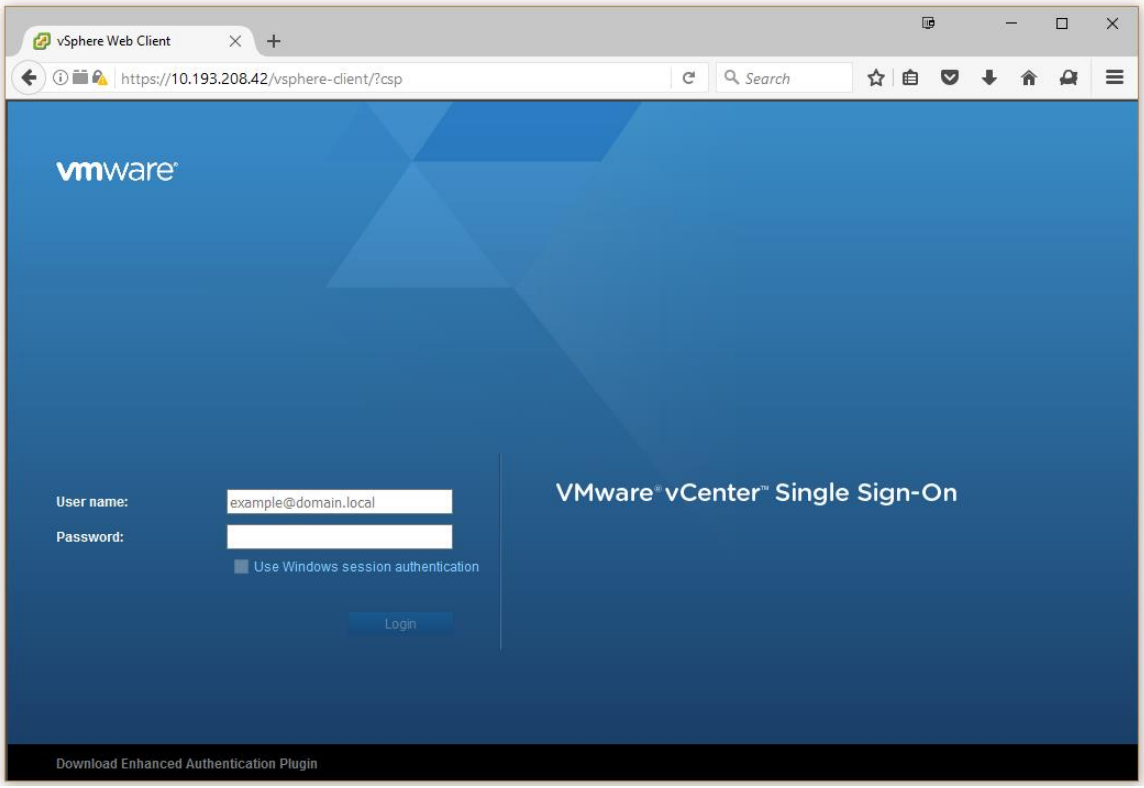

ii. Click on **Hosts and Clusters** tab in the **Object Navigator**. Right-click on your vCenter server and select **New Datacenter**.

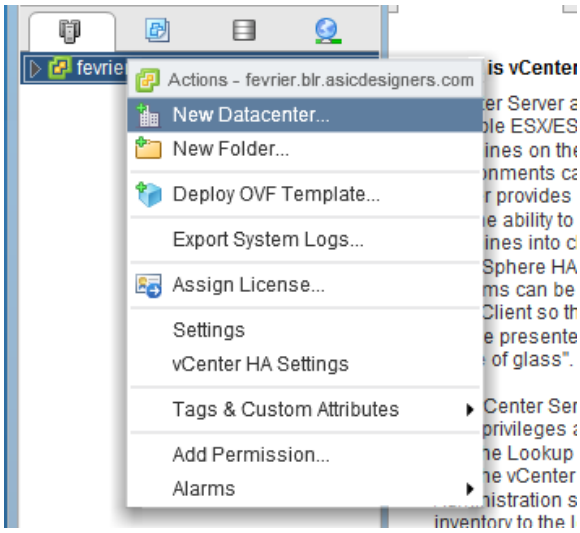

iii. Provide a name for the new datacenter and click **OK**.

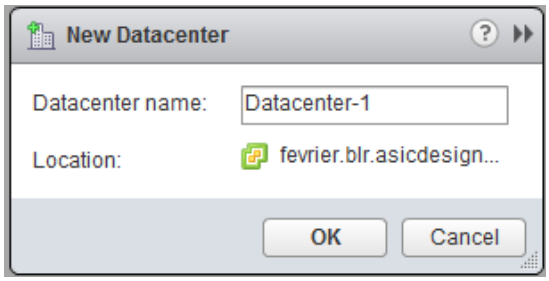

iv. Right-click on the newly created datacenter and select **New Cluster**.

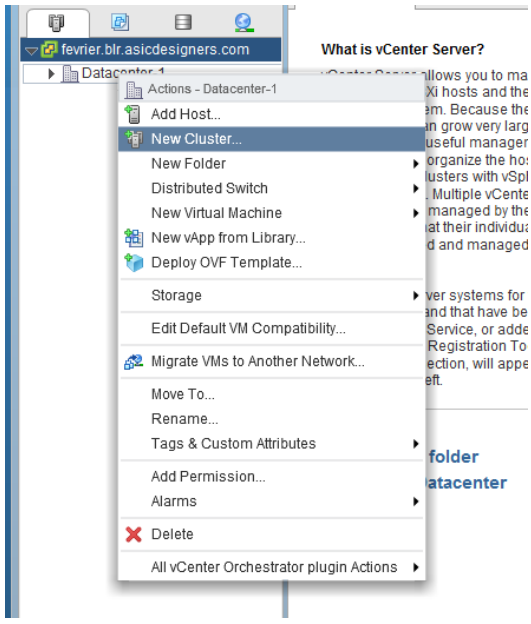

v. Enter a name for the new cluster and select options for cluster features. Click **OK**.

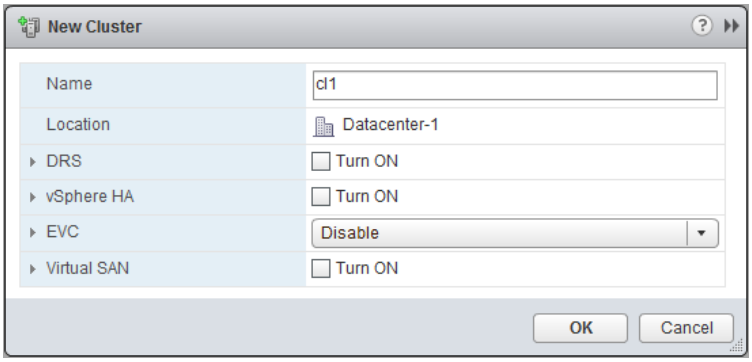

vi. The newly created cluster will appear in the inventory. Repeat the above step to add more clusters to the datacenter.

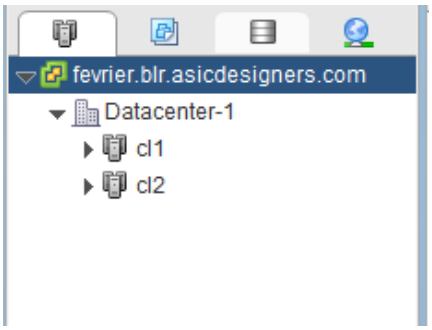

vii. Right-click on the newly created cluster and select **Add Host**. This will invoke the **Add Host Wizard** allowing you to add host(s) to the newly created cluster.

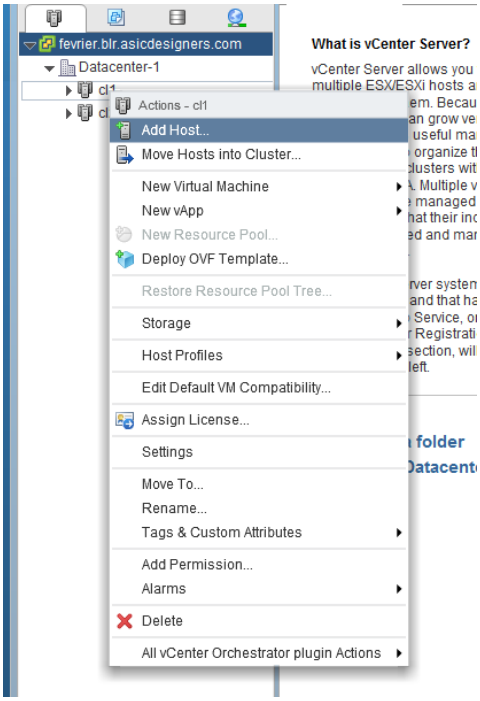

viii. Provide name or IP address of the host, username and password, and other configuration details. Click **Next**.

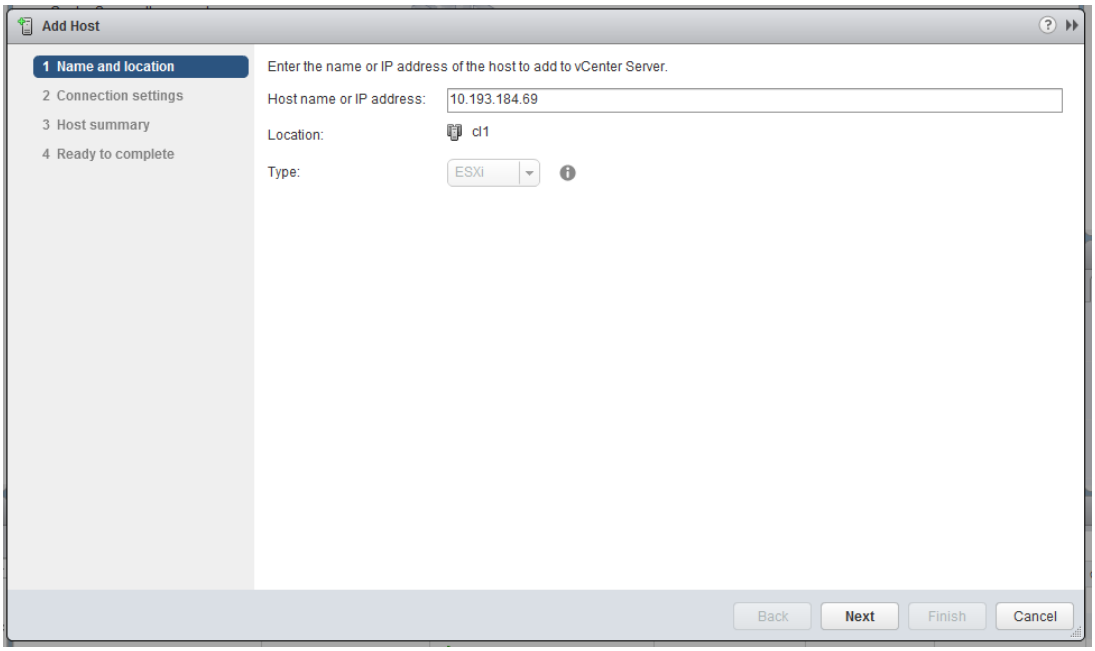

ix. Review the information provided and click **Finish**.

Creating Distributed Switch

i. Log in to vCenter Server through vSphere Web Client using a web browser.
ii. Click on **Networking** tab.
iii. Right-click on your datacenter and select **Distributed Switch > New Distributed Switch**.

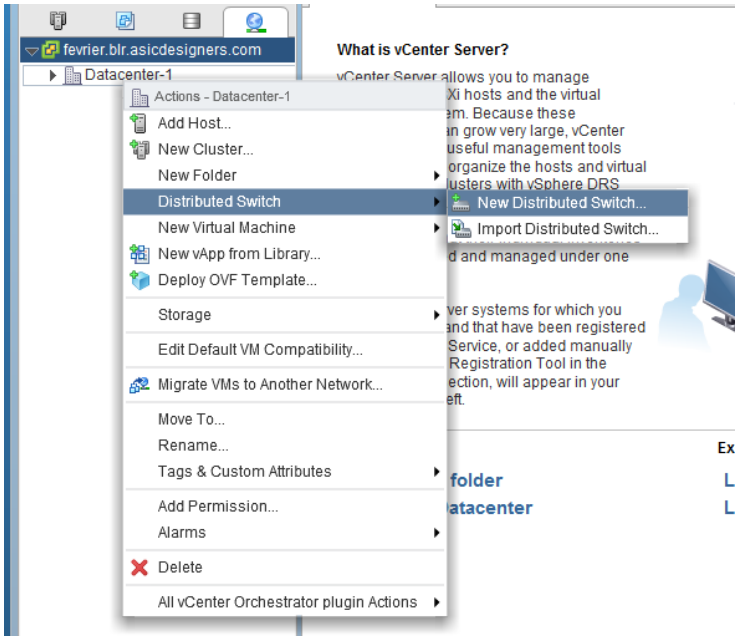

iv. Enter a name for the switch and click next.

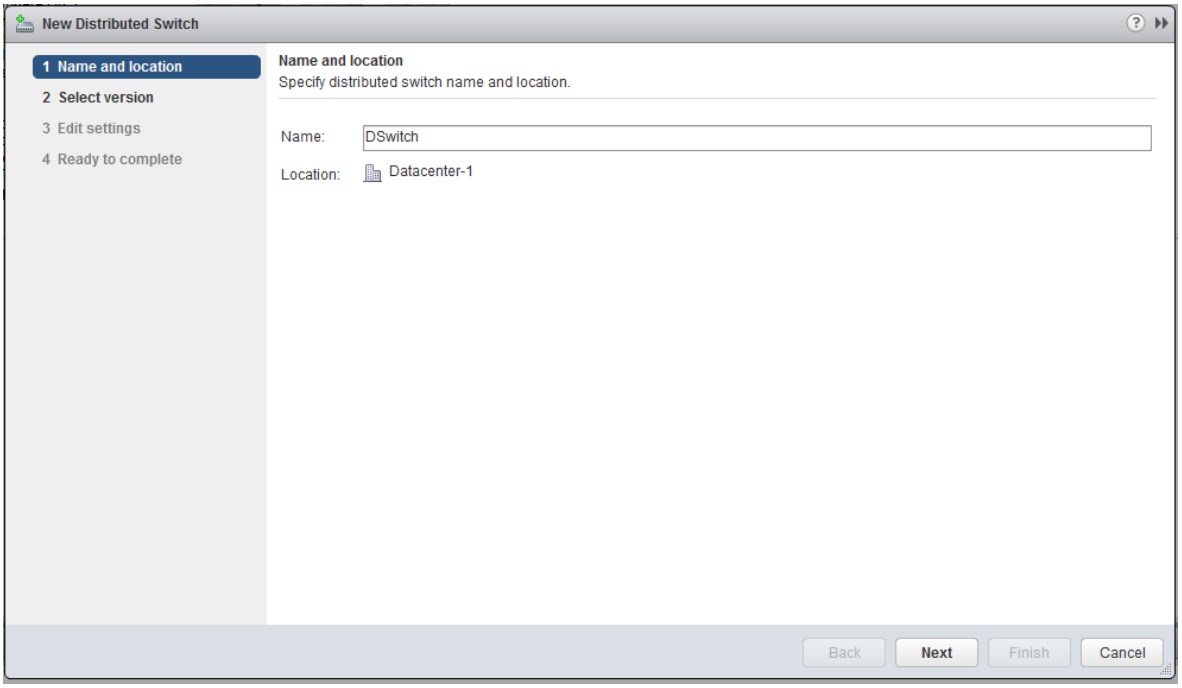

v. Select the required switch version and click **Next**.

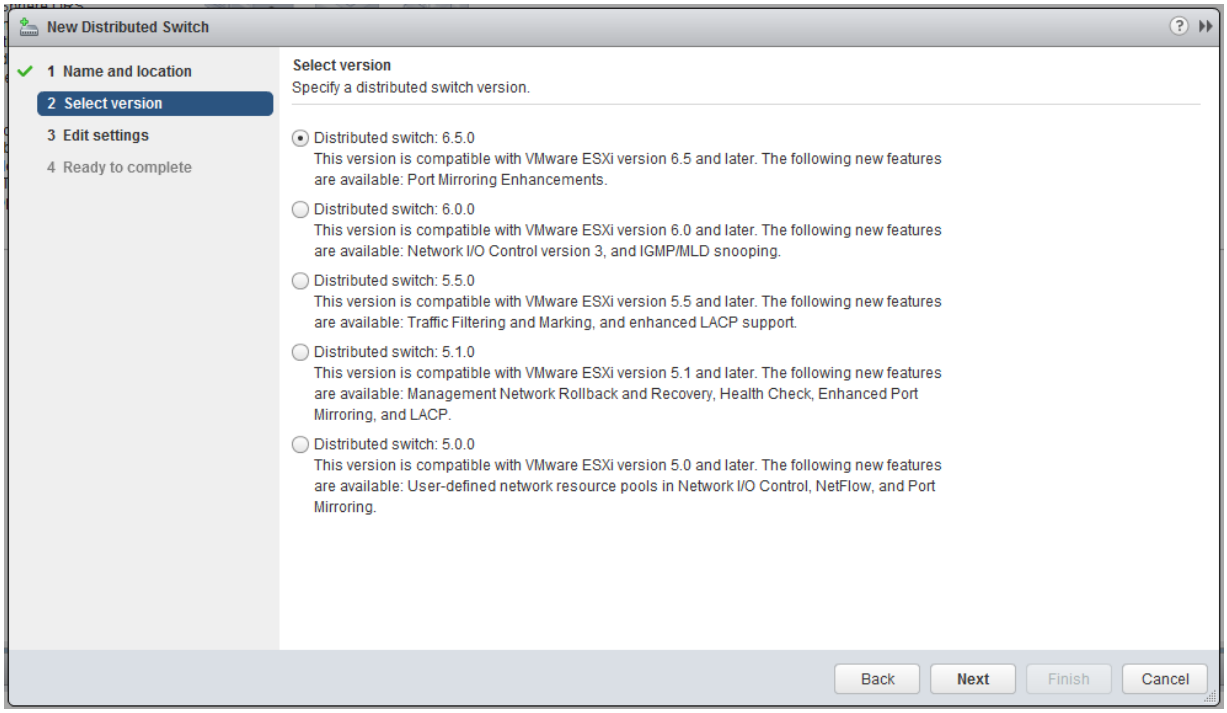

vi. Select the required number of uplinks for your setup. Please note that the number of uplinks must be equal or less than the number of adapters to be used for each host. Click **Next**.

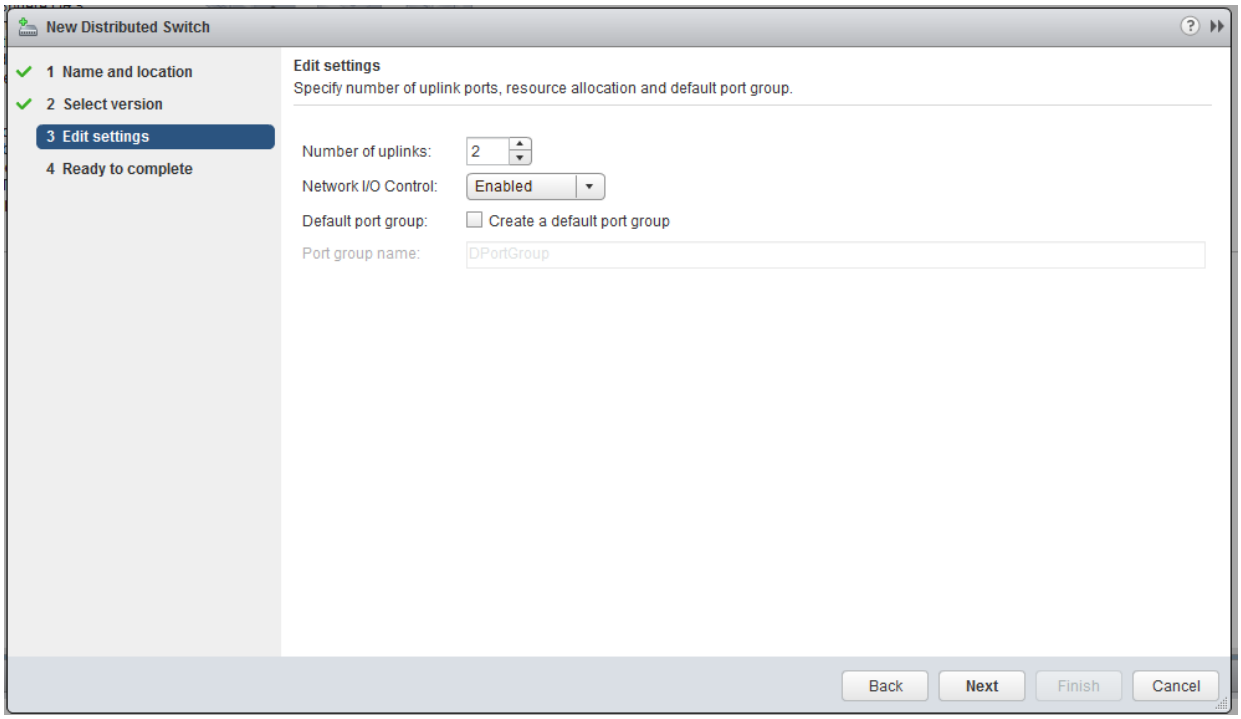

vii. Review the settings and click **Finish**.

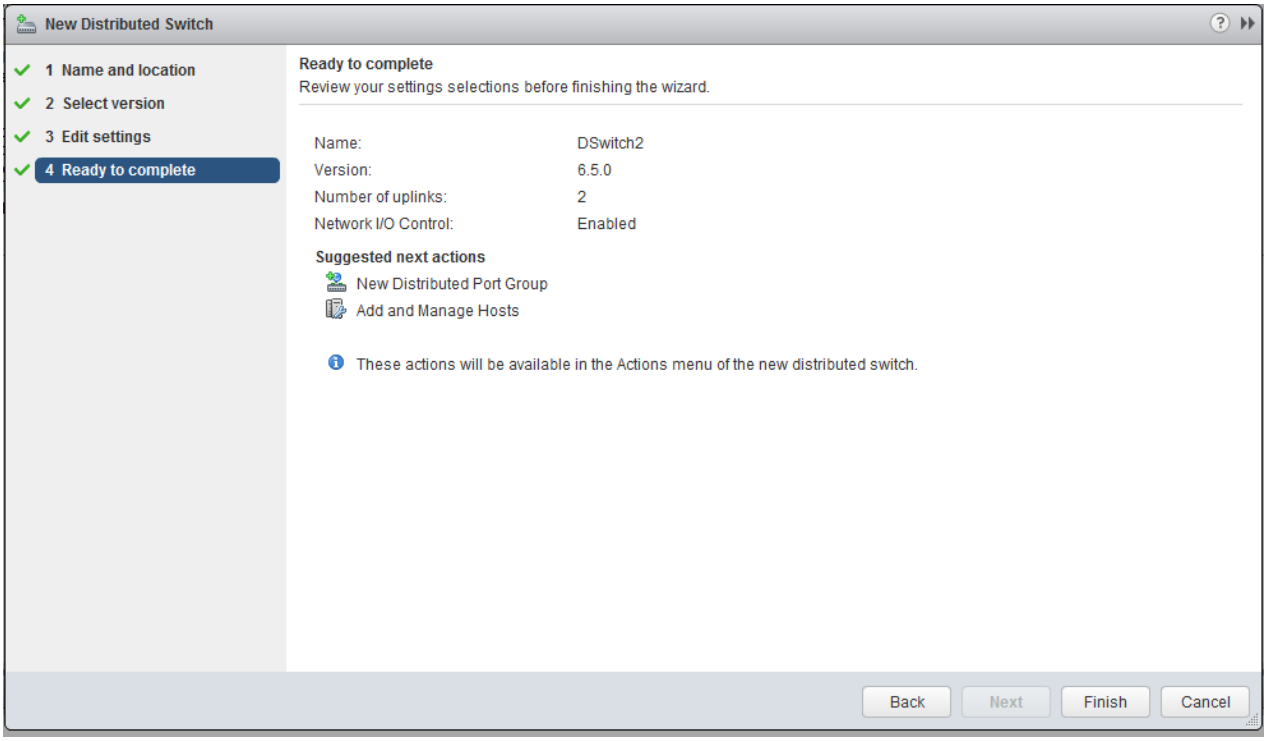

viii. The newly created switch will appear in the inventory.

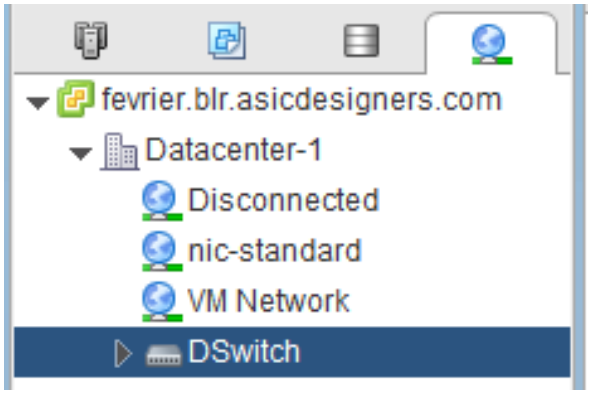

ix. Right-click on the switch and select **Add and Manage Hosts**.

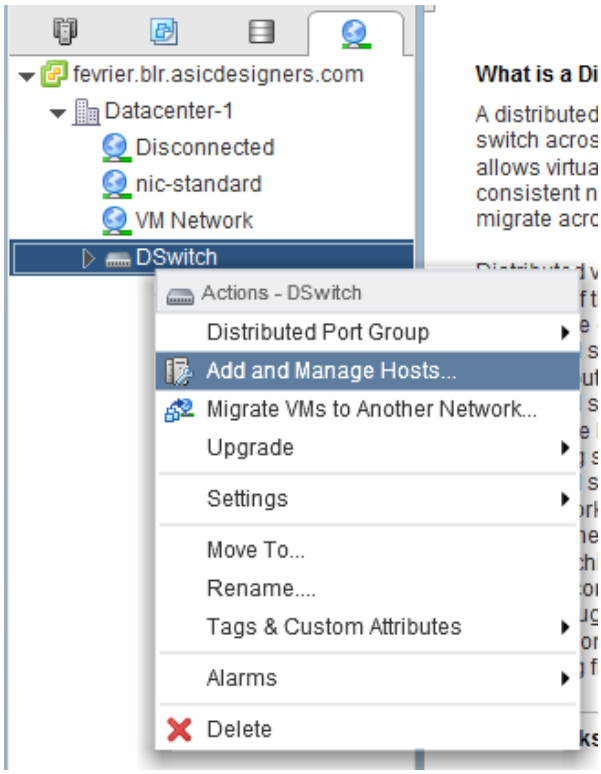

x. Select **Manage host networking** and click **Next**.

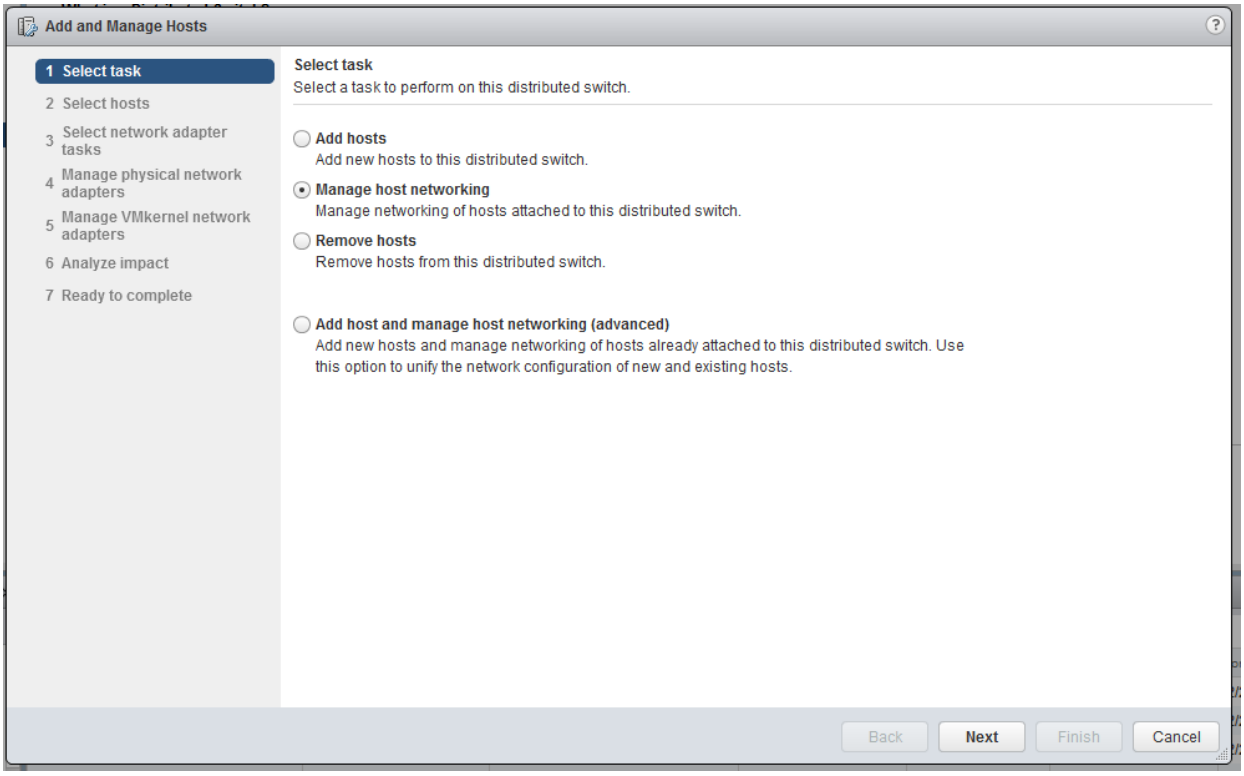

xi. Click **+ Attached hosts**. Select member hosts to add and click **OK**.

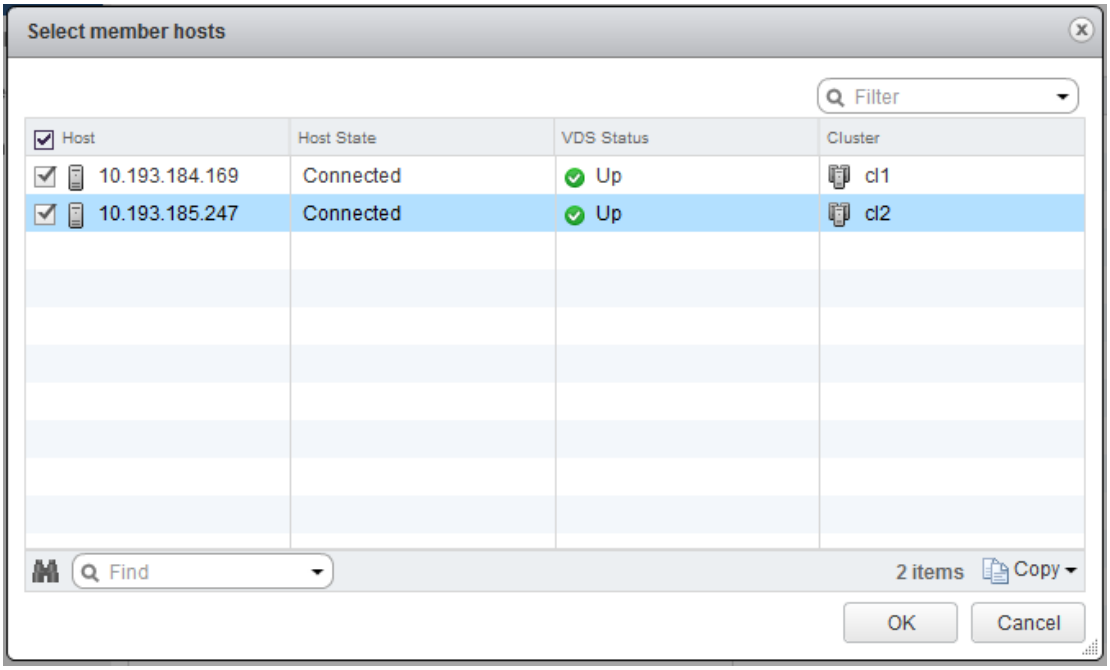

xii. Click **Next** and select the appropriate network adapters tasks to perform. Click **Next**.

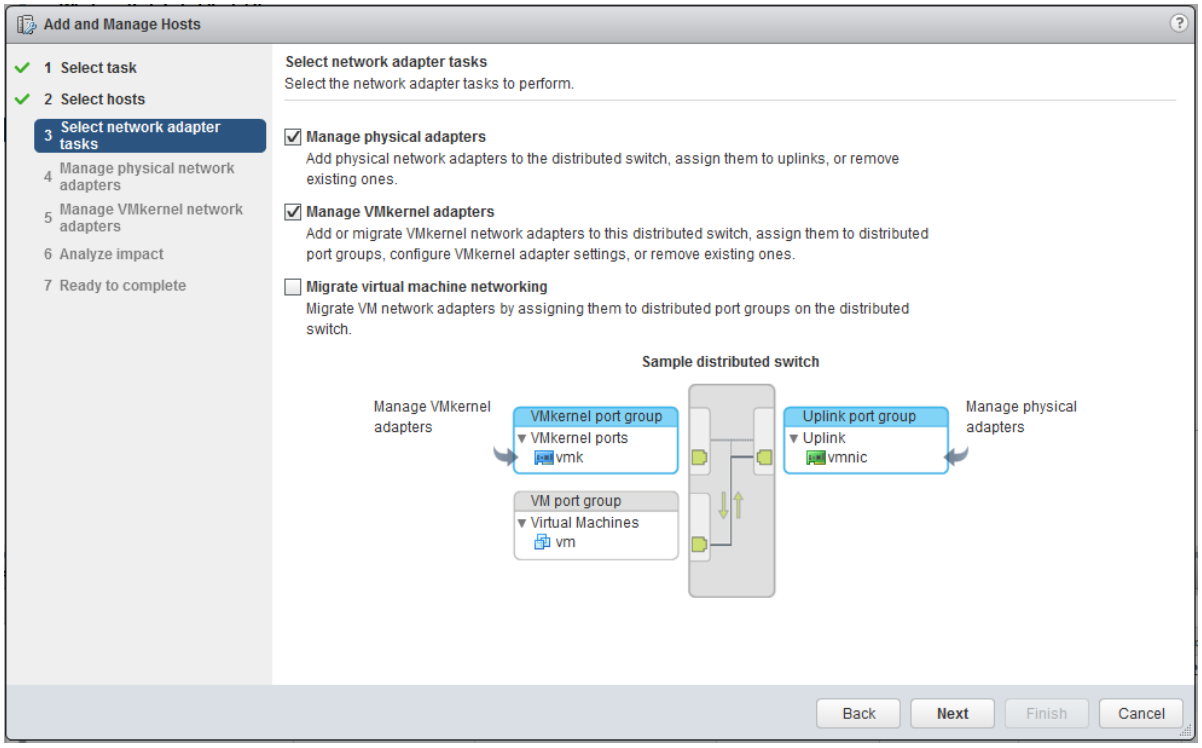

xiii. Here you can add physical network adapters to the switch. Select the adapter and click **Assign uplink**.

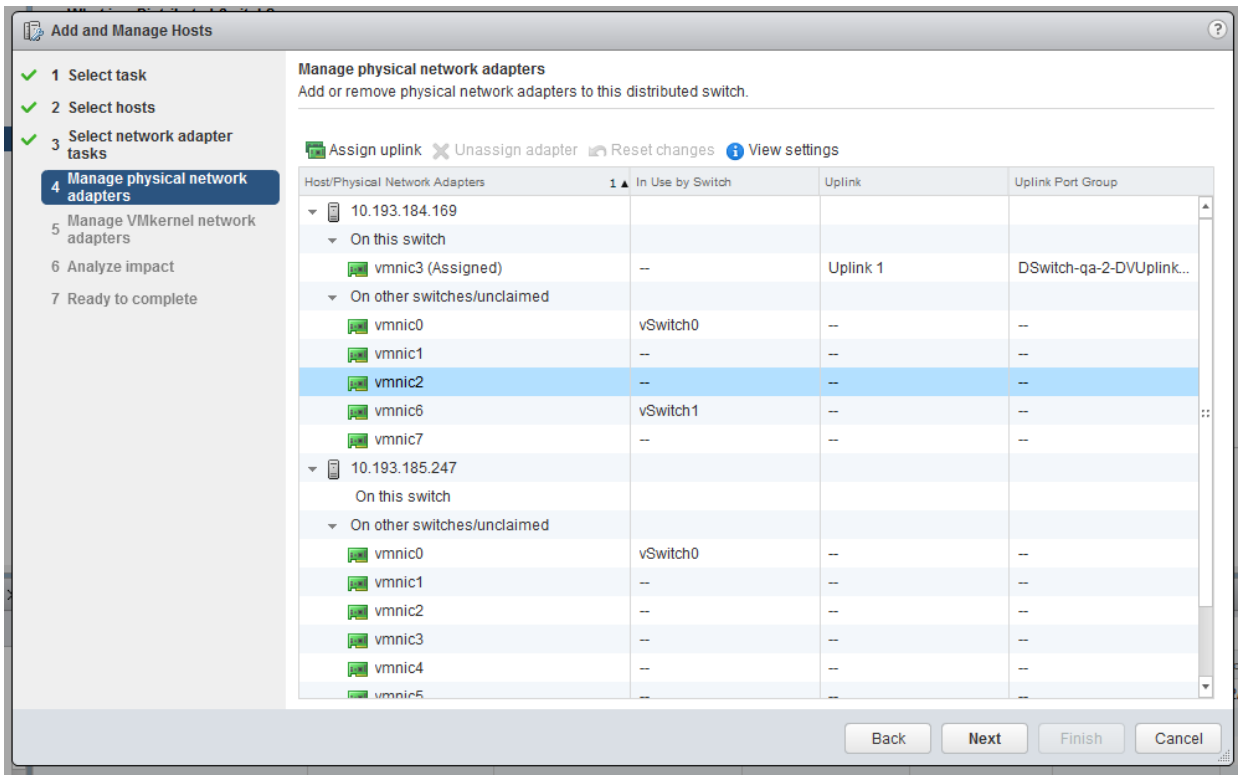

xiv. Select an uplink for the adapter and click **OK** and then **Next**.

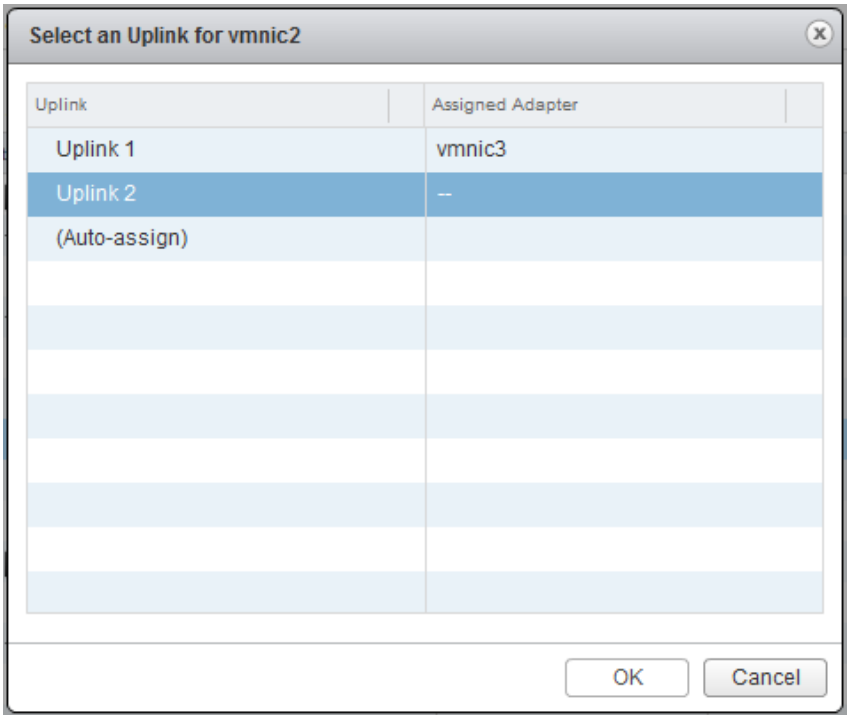

xv. In this step, you can view the VMkernel adapters automatically added to the hosts. If VXLAN modules are already installed using NSX Manager, you can also view the VXLAN Port group associated to the respective host in distributed switch. Click **Next**.

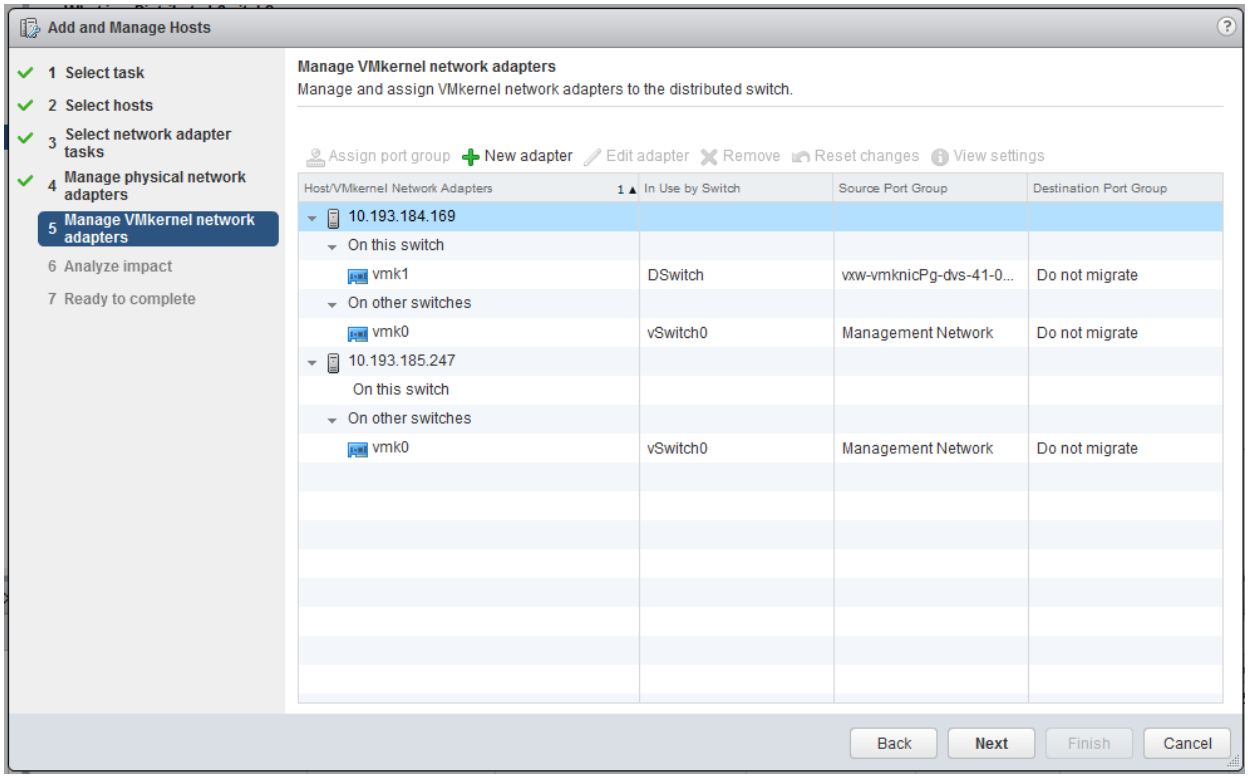

xvi. Review the impact status of the configuration. Click **Next**.

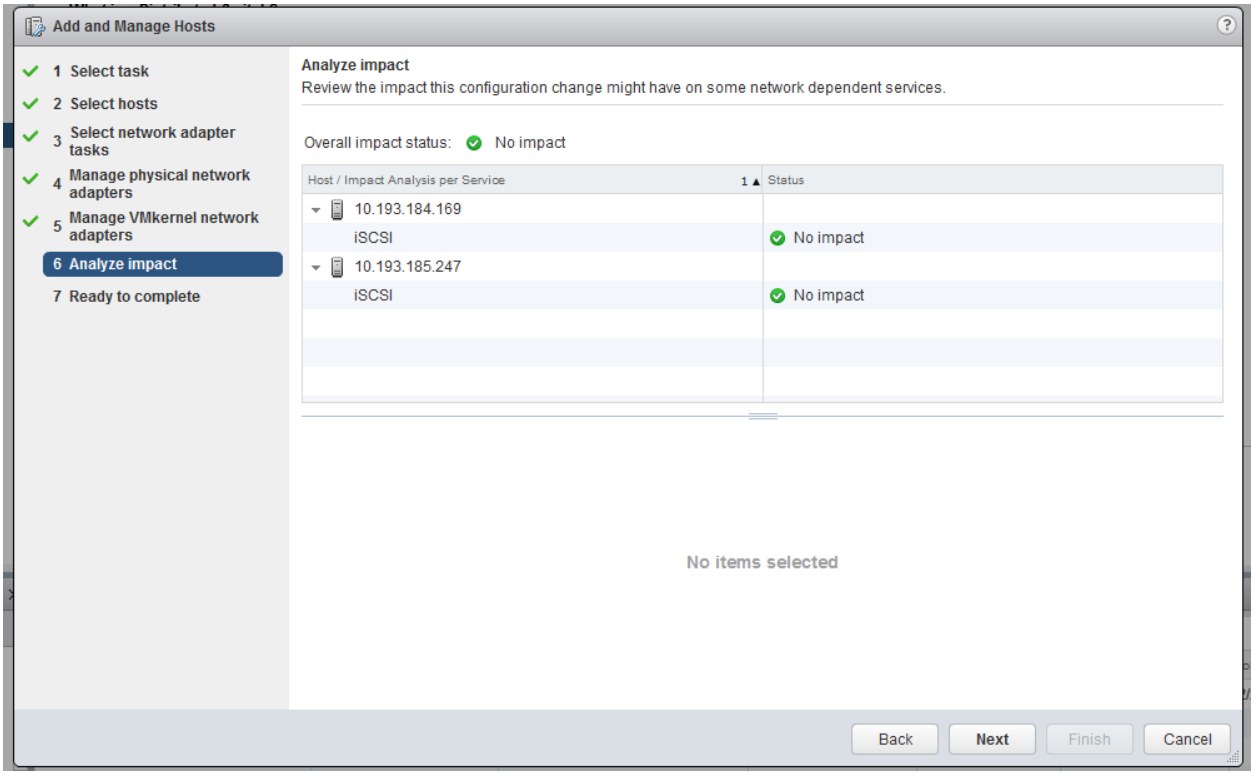

xvii. Review the settings and click **Finish**.

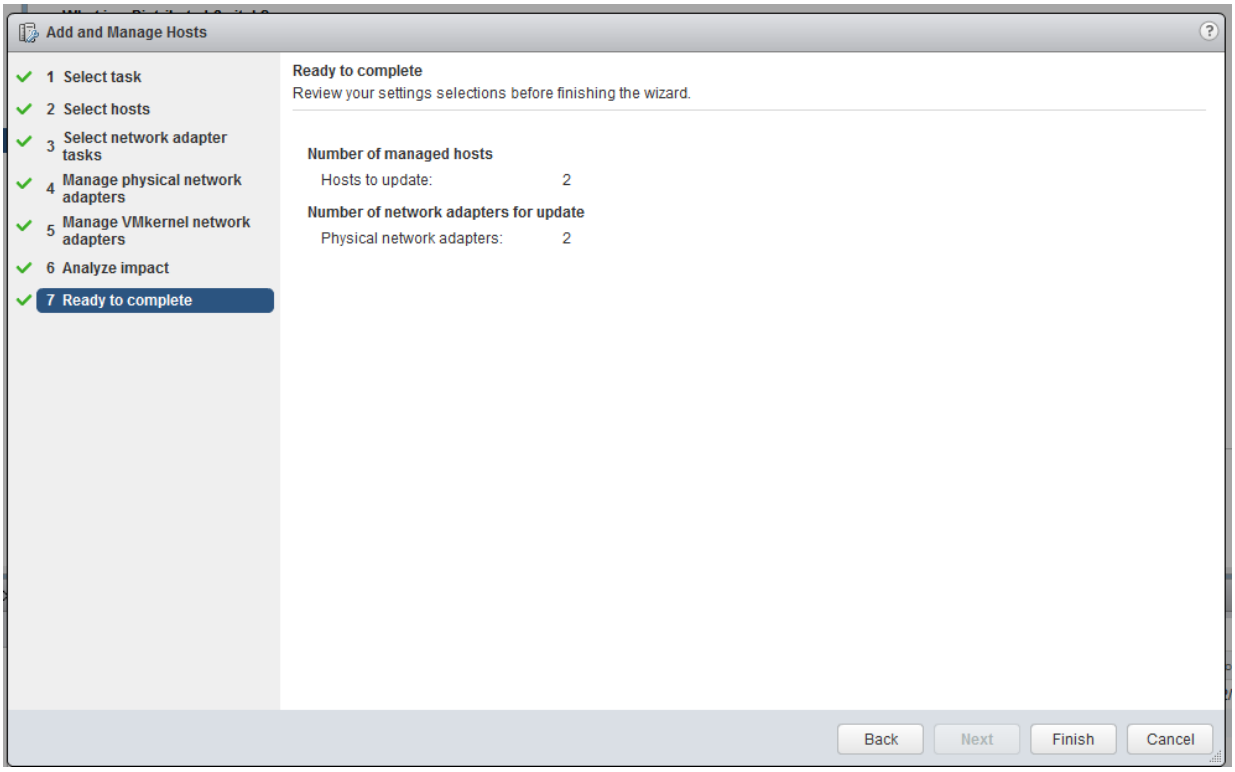

Setting up VXLAN

i. Follow the steps mentioned in VMware's NSX installation guide to install and configure NSX Manager.

ii. Make certain that vCenter Server is registered with NSX manager successfully. **Networking & Security** link should appear in both **Home** and **Navigator** panes. Click on the link in any of the two panes.

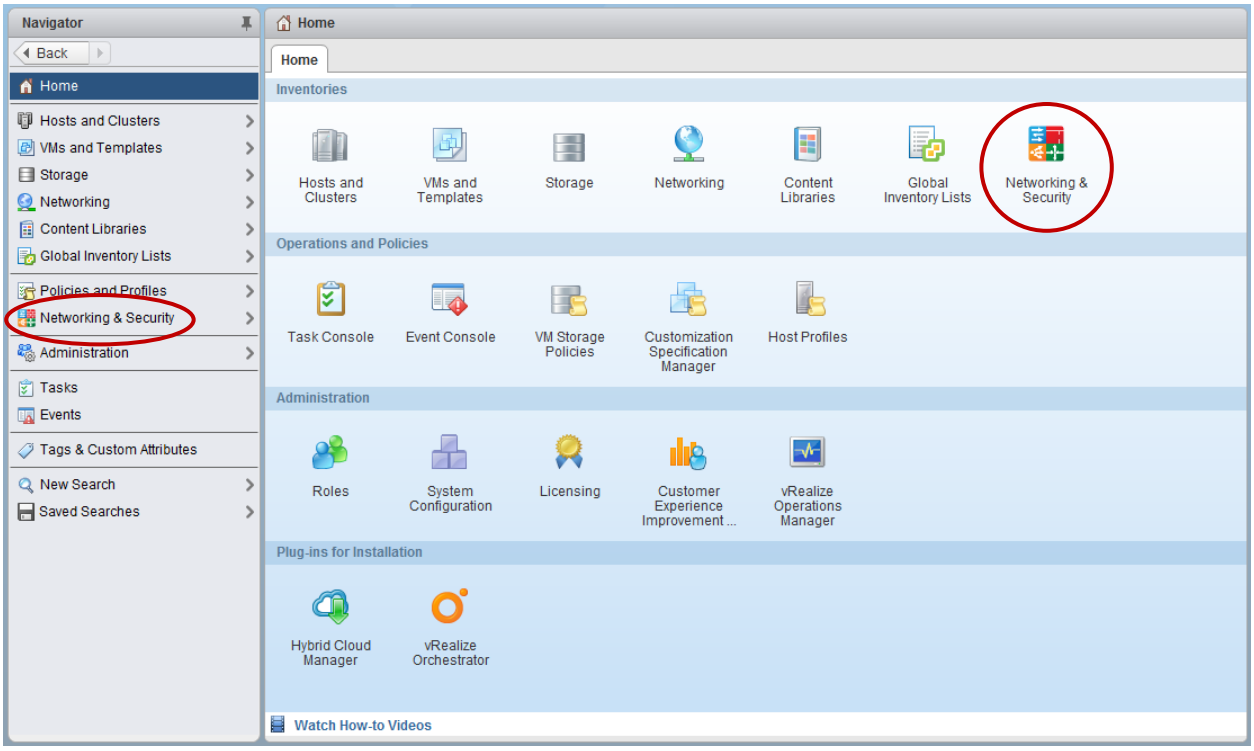

iii. In the **Navigator** pane, click on **Installation** and then **Host Preparation** tab. Select the IP of the required NSX manager from the drop-down menu.

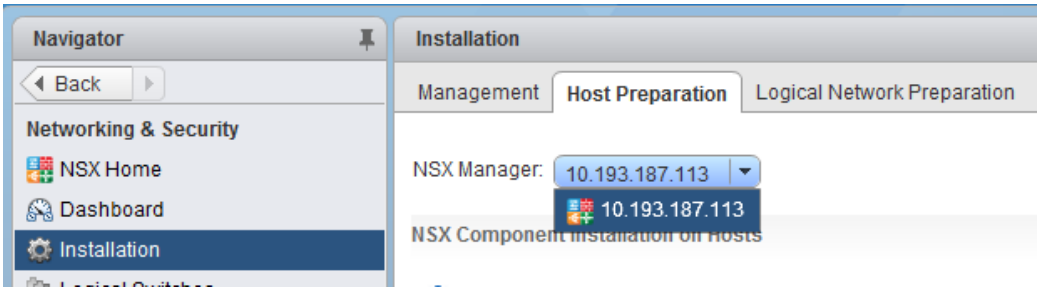

iv. In the *Installation Status* column for the required cluster, click on options (⚙) and select **Install**.

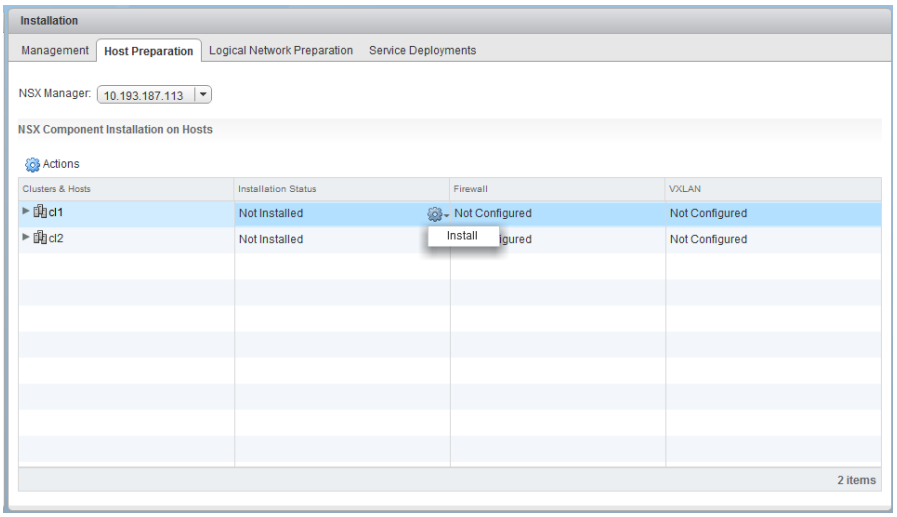

v. Click **Yes** to confirm. Installation progress will be shown in the **Recent Tasks** pane.

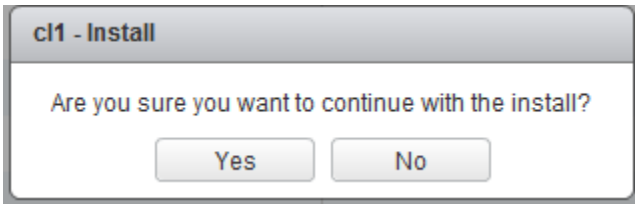

vi. In the *VXLAN* column for the required cluster, click on options (⚙) and select **Configure VXLAN**.

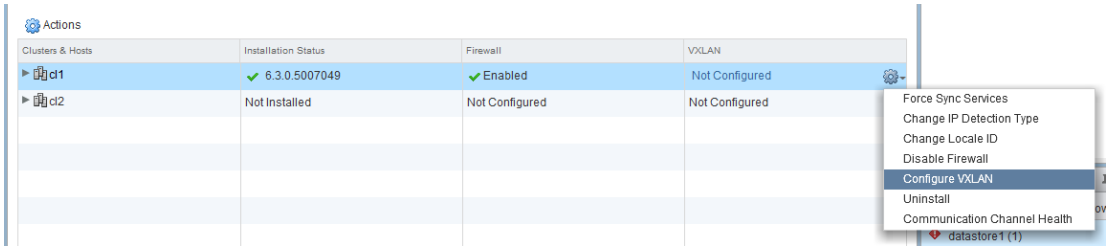

vii. Select the distributed switch created previously from the drop-down menu. Provide details like VLAN ID, IP addressing and Teaming policy and click **OK**.

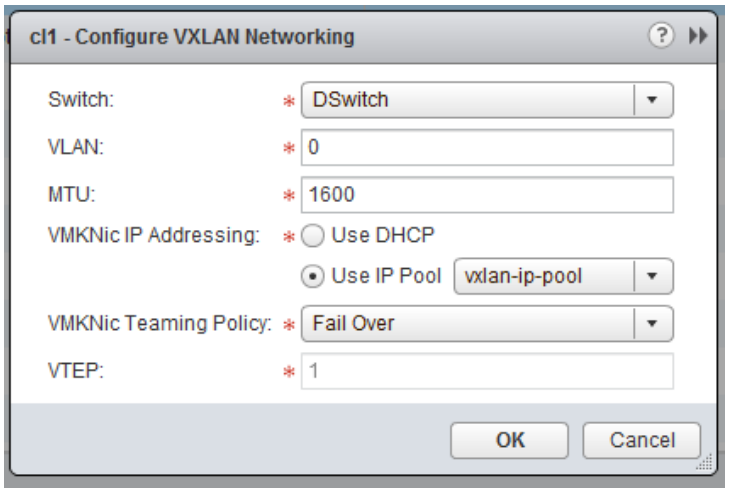

viii. Click on the **Logical Network Preparation** tab and then **Segment ID**. Set the range for *Segment ID pool* and *Multicast addresses*. Click **OK**.

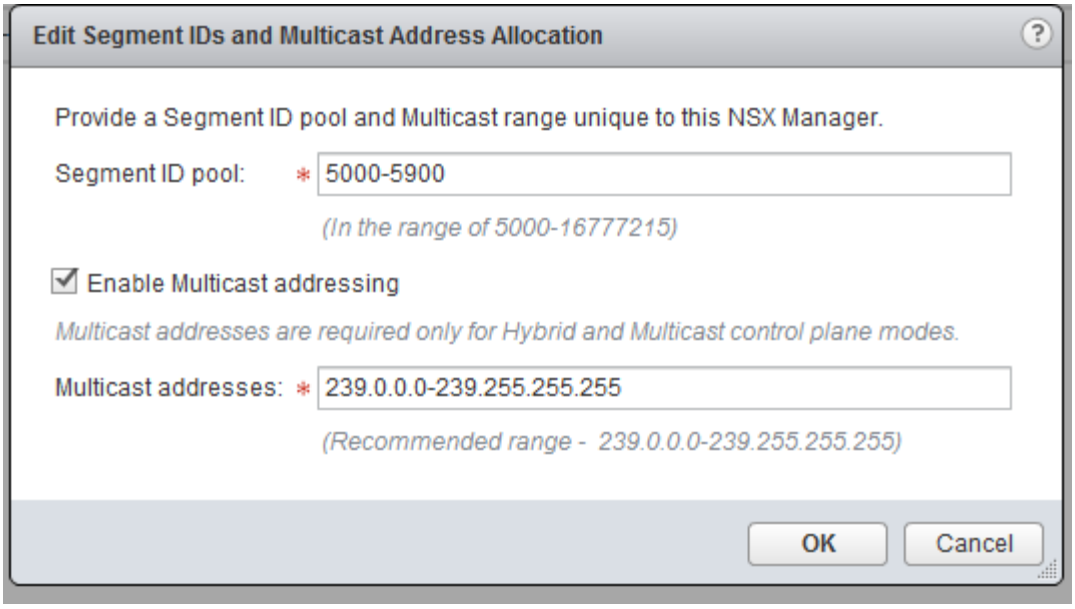

ix. Click on **Transport Zones** and then ✚ link.

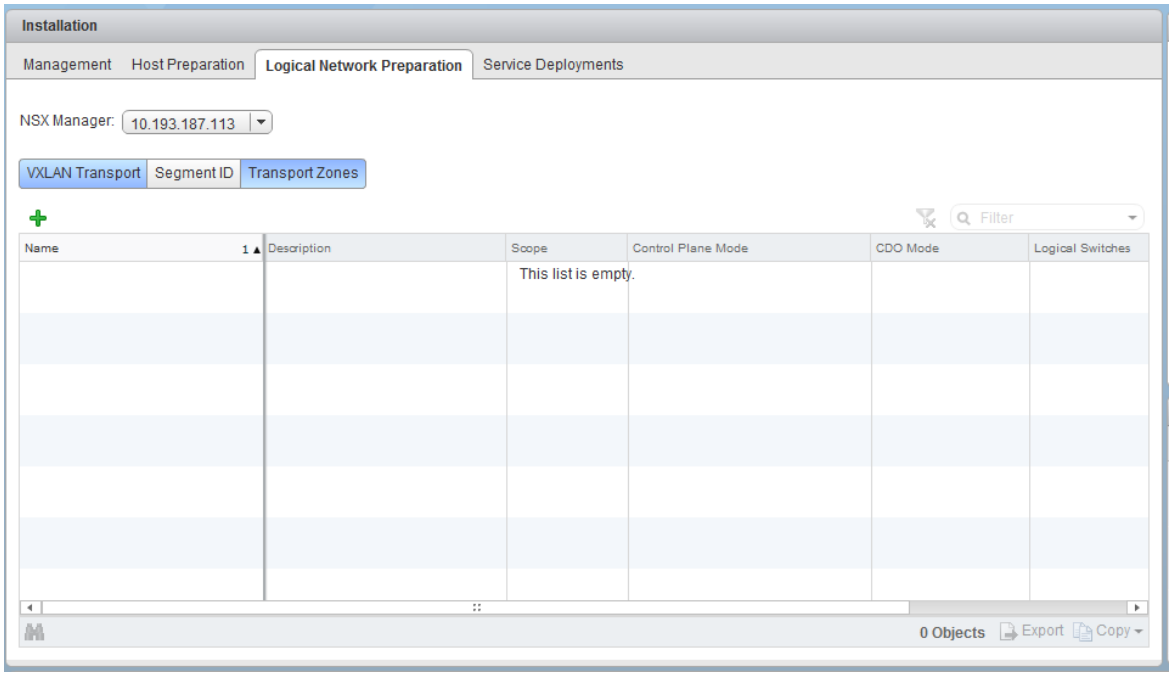

x. Provide a name for the transport zone. Select the replication mode and clusters that will be part of the transport zone. Click **OK**.

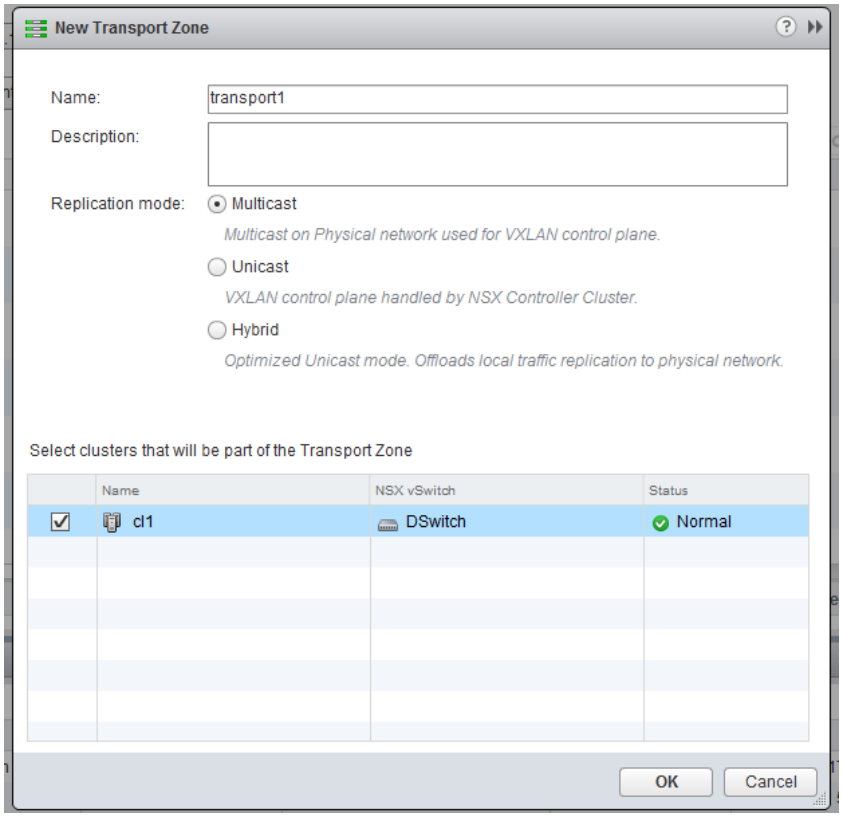

xi. In the **Navigator** pane, select **Logical Switches** under **Networking & Security** and then ✚ link.

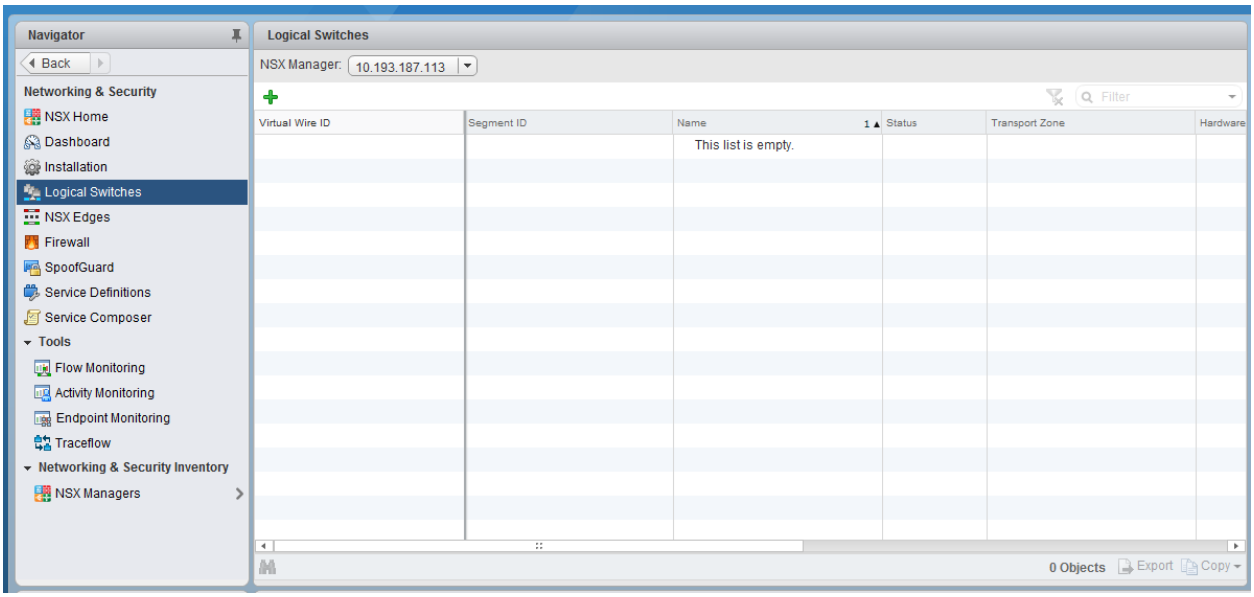

xii. Provide a name for the logical switch and select the transport zone that was previously created. Change the replication mode if required and Click **OK**.

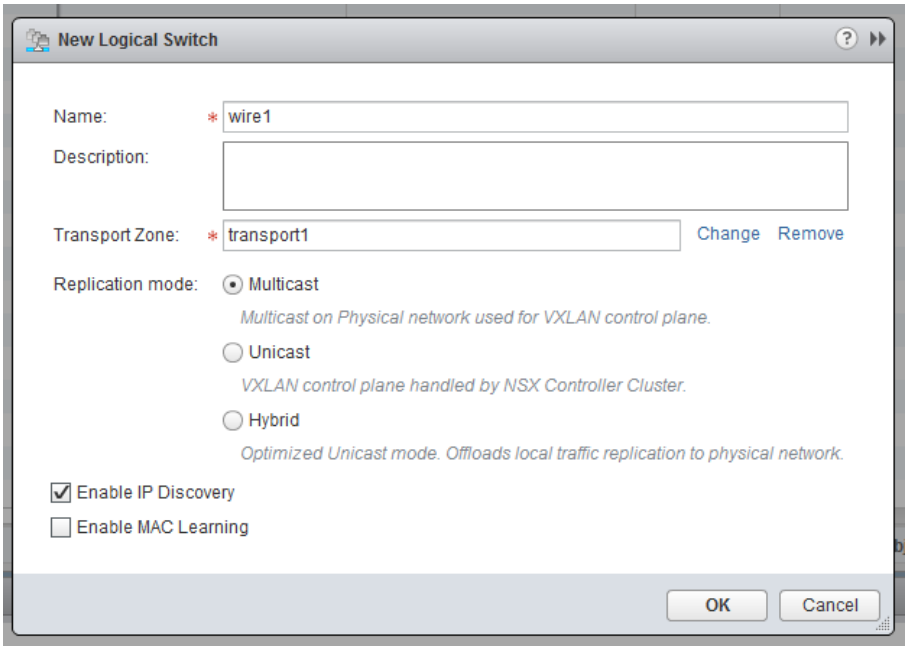

xiii. Using the main menu (or shortcut Ctrl+Alt+2), access the **Hosts and Clusters** section.

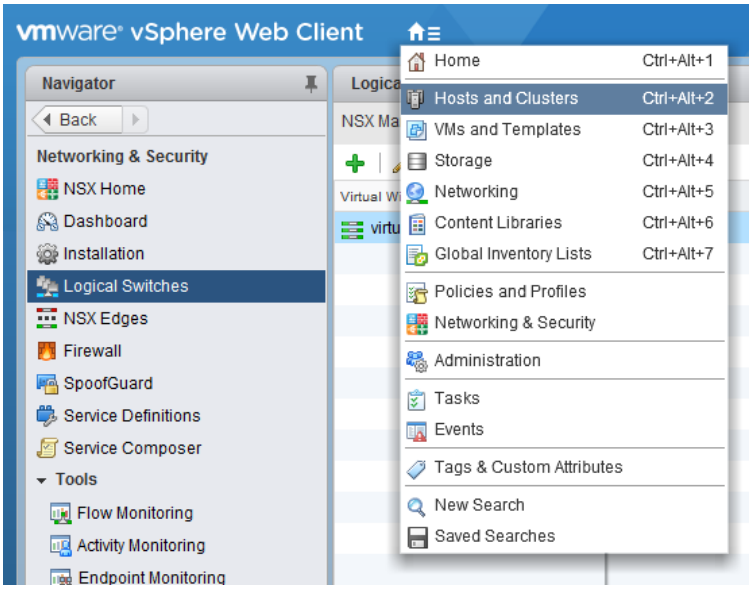

xiv. Select VM under **Datacenter > Cluster,** right-click and select **Edit Settings**.

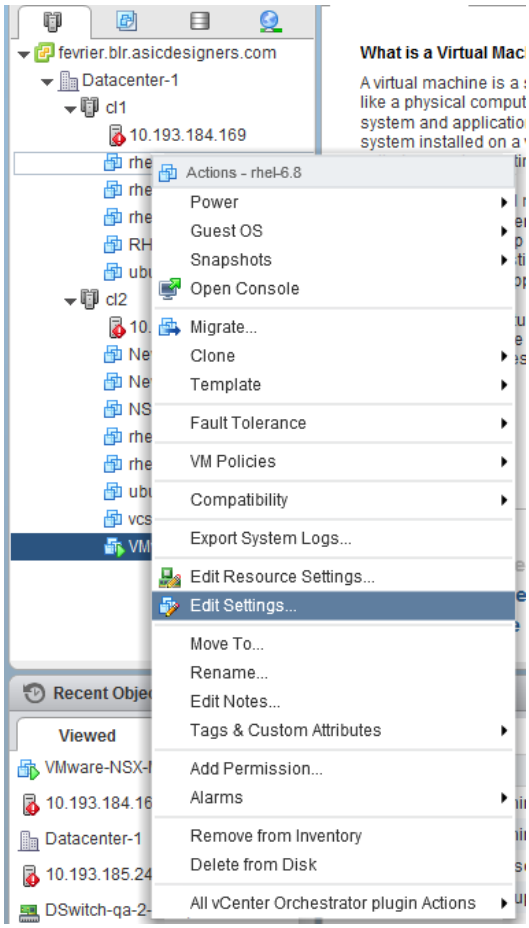

xv. In the **New device** drop-down menu, select **Network** and then click **Add**. This will add a **New Network** entry under **Virtual Hardware** tab.

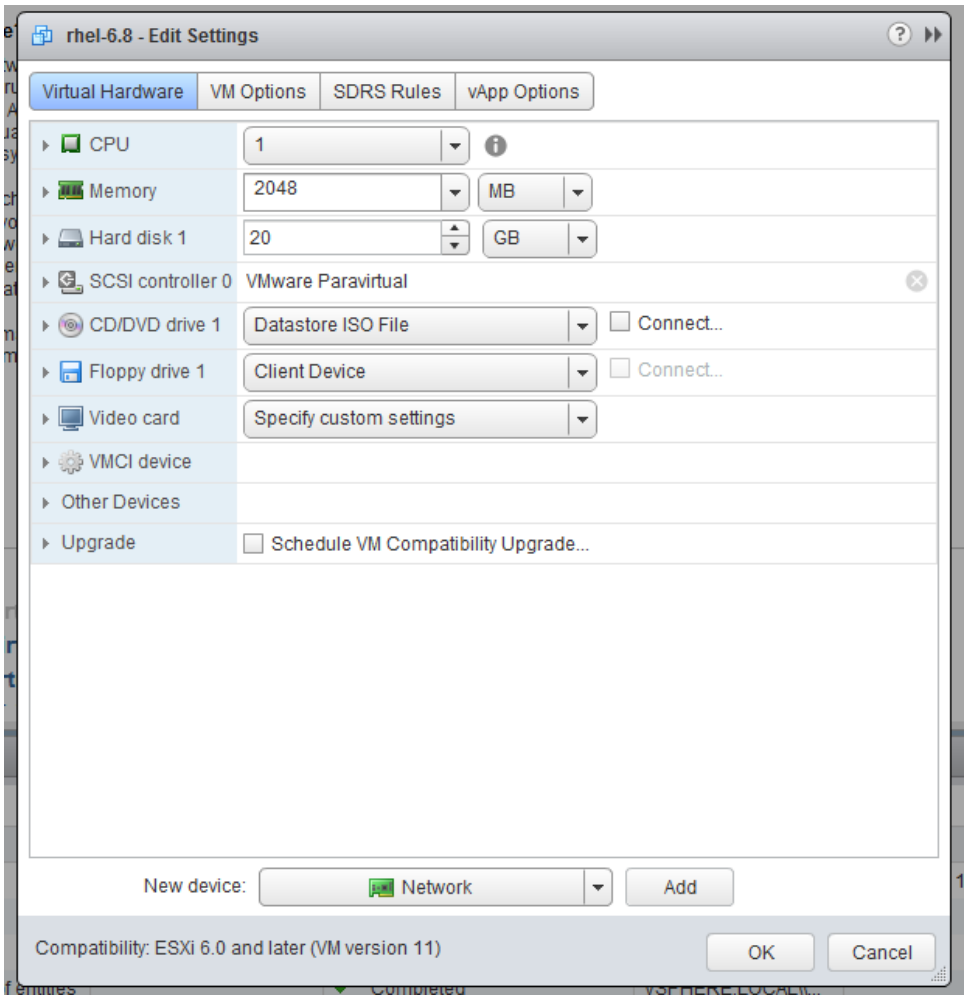

xvi. In the **New Network** drop-down menu, select *Show more networks*. Add the logical switch created previously. Click **OK** and then **OK** again.

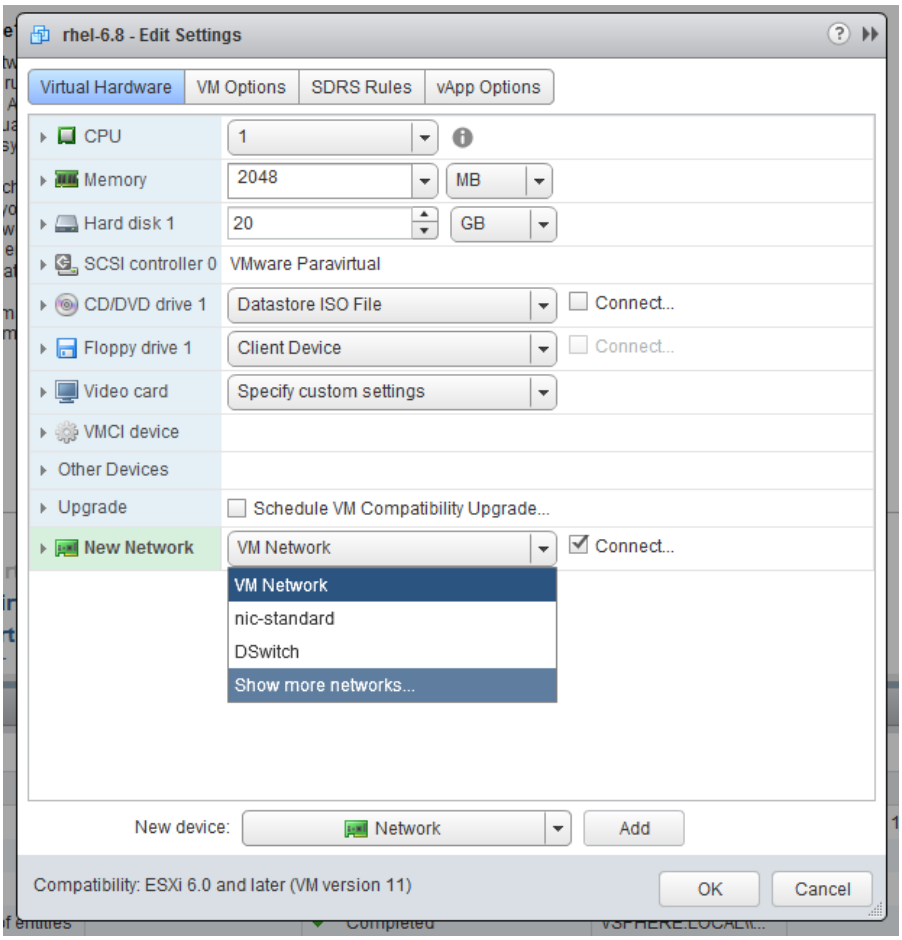

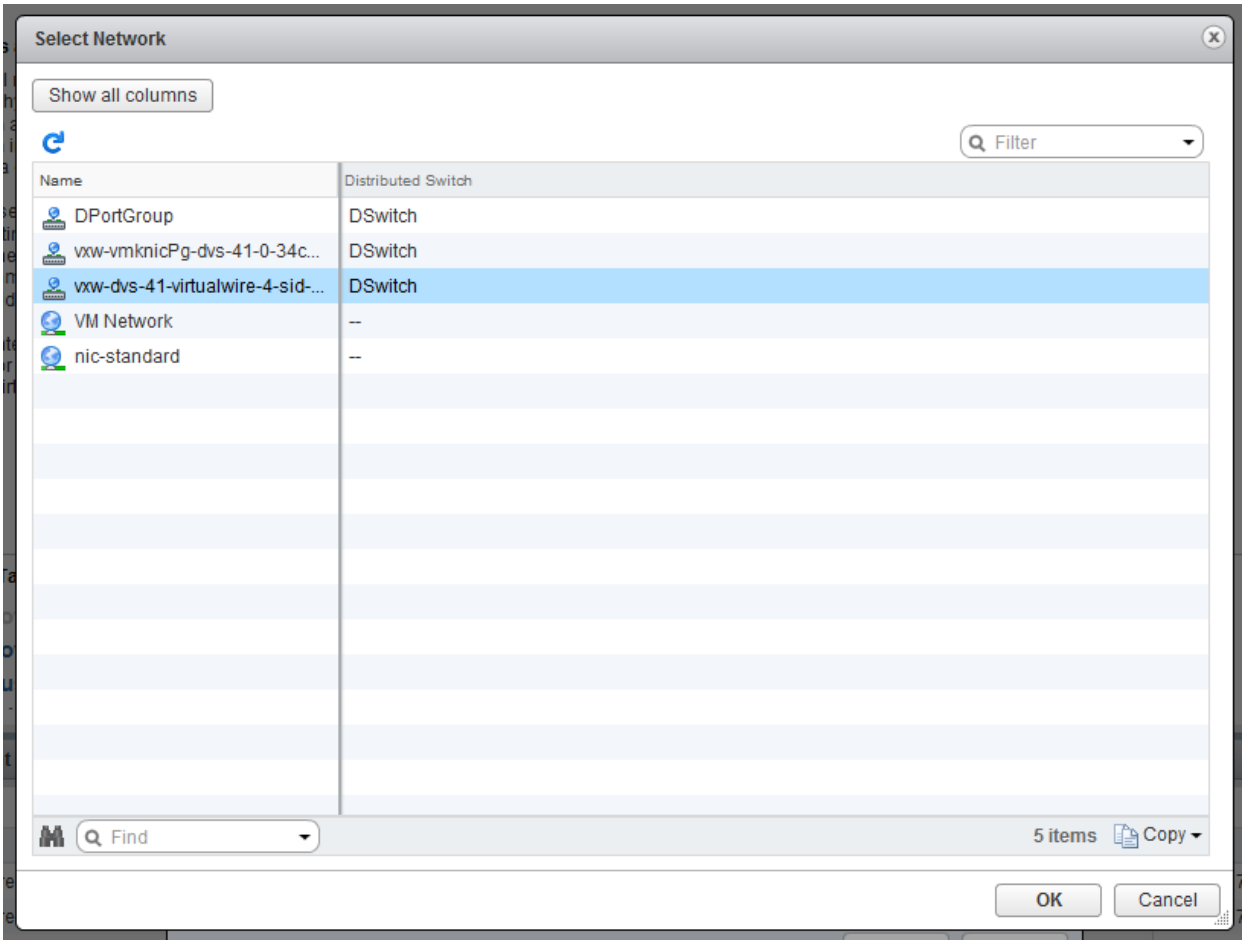

xvii. Repeat the above step for all the VMs in the cluster. All the VMs added to the VXLAN will now be able to communicate successfully.

5.7.2. ESXi 6.0

- **Configuring Adapters**

Please see **Configuring Adapters** sub-section for ESXi 6.5a.

- **Creating Cluster**

Note *Instructions provided here are for desktop client. Steps for web client are similar with minor variations.*

i. Open a vSphere Client session to a vCenter Server.
ii. Select **View > Inventory > Hosts and Clusters.**

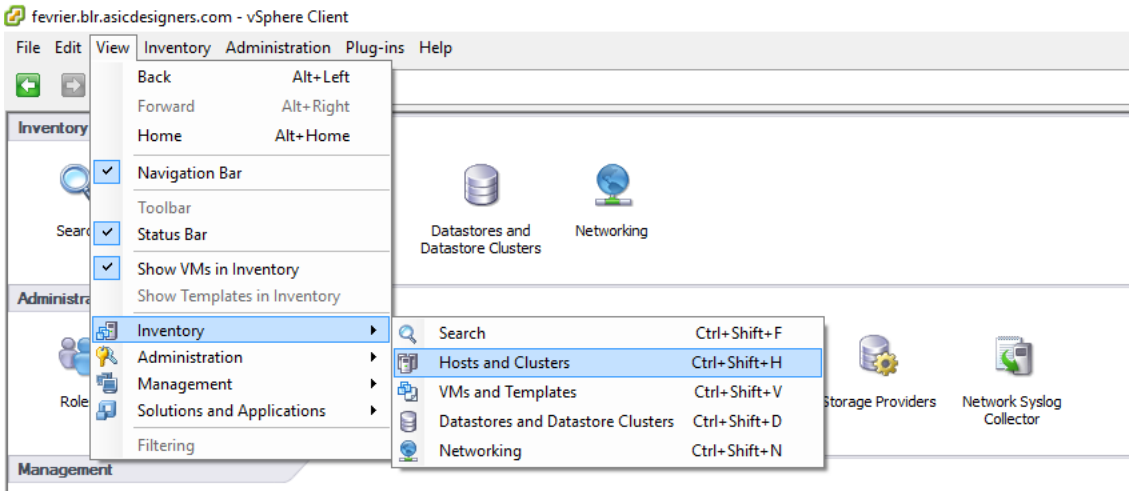

iii. Right-click on your vCenter Server and select **New Datacenter**. Provide a name for the new datacenter and hit [Enter].

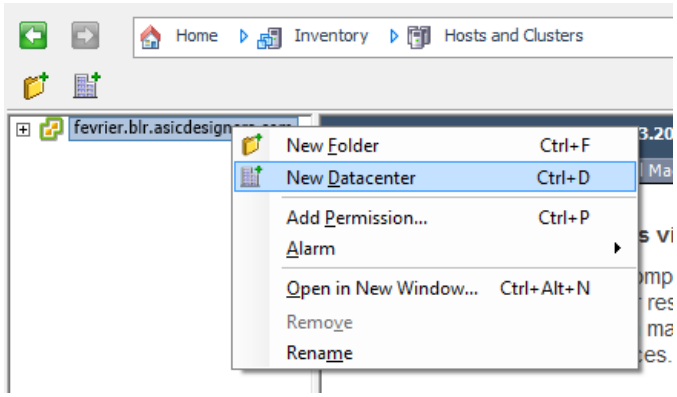

iv. Right-click on your datacenter and select **New Cluster**.

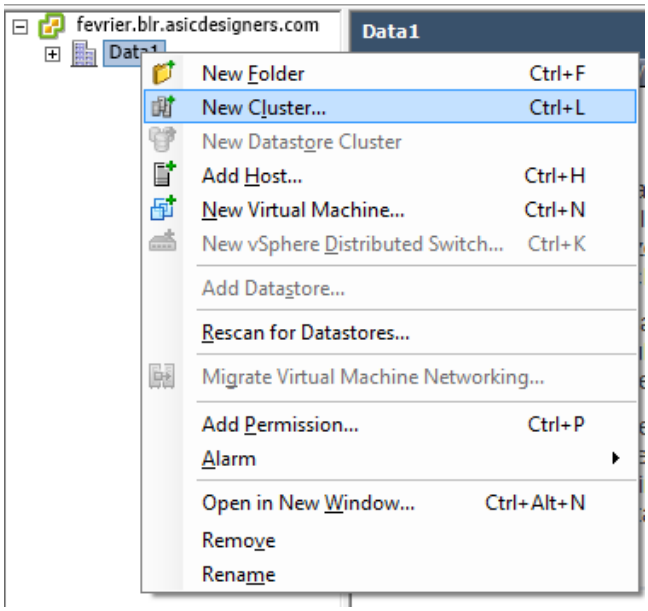

v. Enter a name for the new cluster and select options for cluster features, Enhanced vMotion Compatibility (EVC) and swap file policy in the subsequent screens.

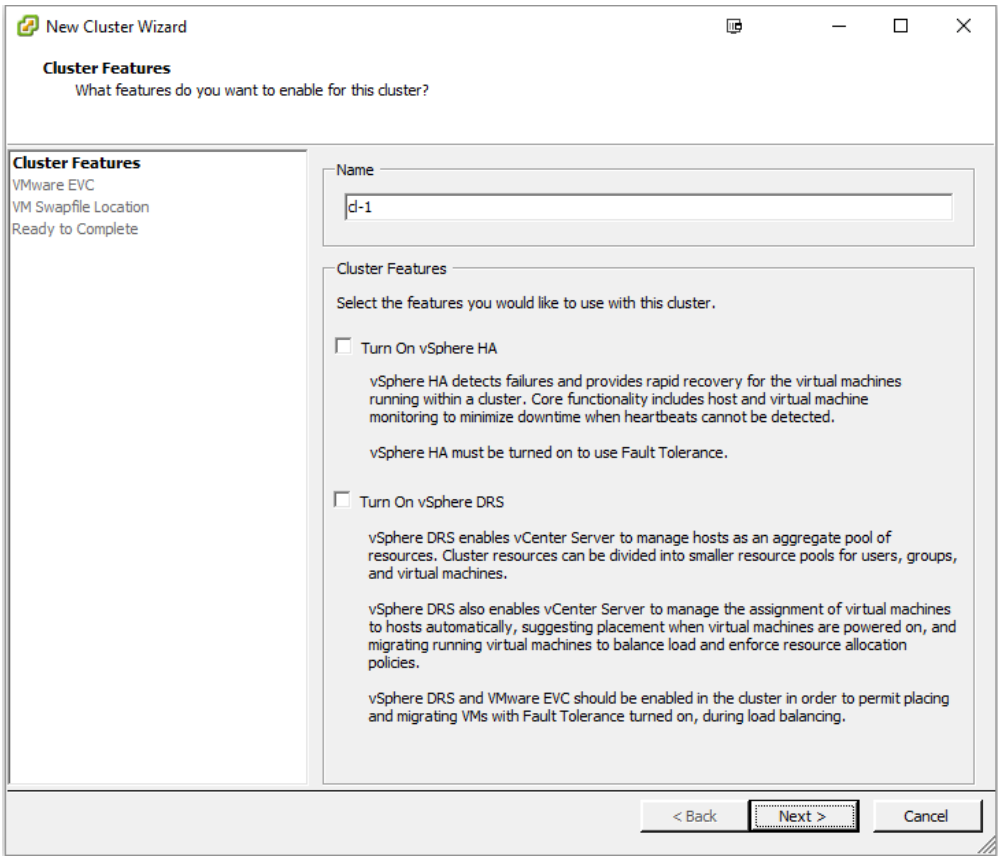

vi. Review the options selected and click **Finish**.

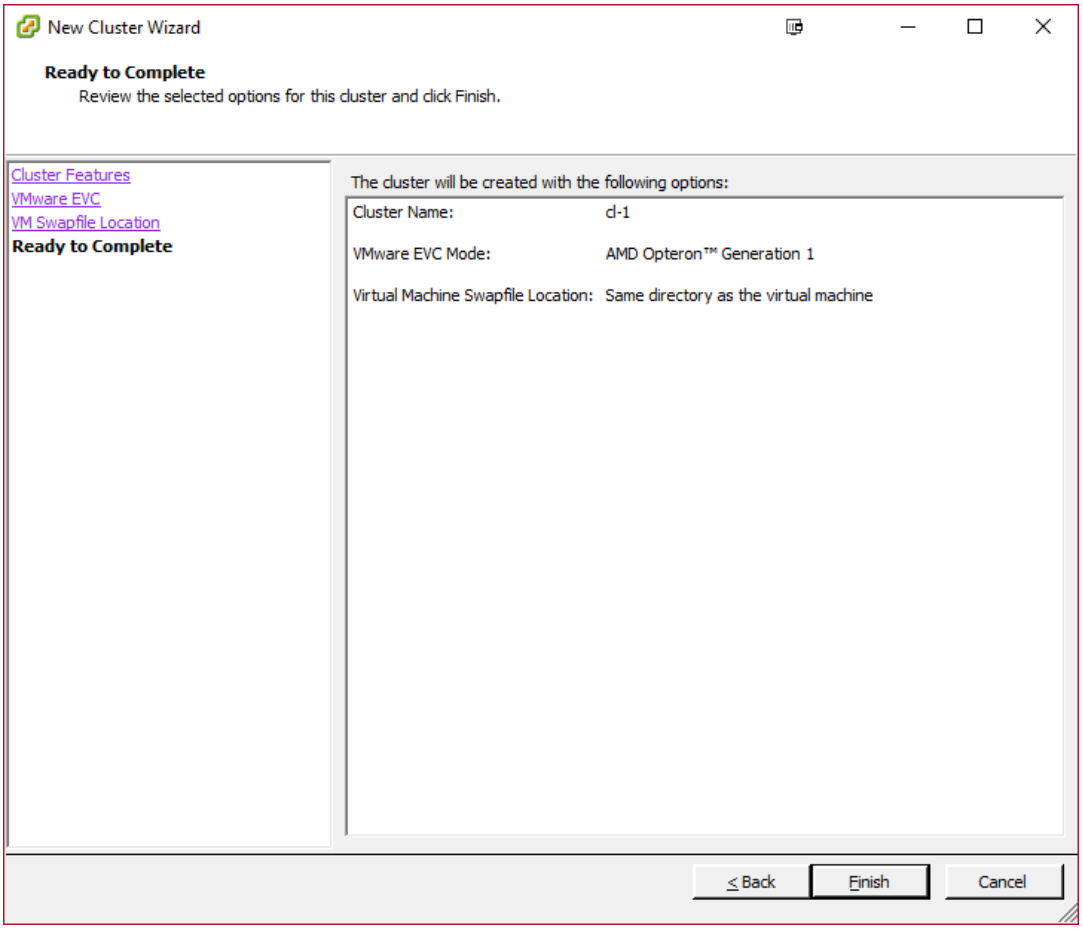

vii. The newly created cluster will appear in the inventory. Repeat the above step to add more clusters in the datacenter.

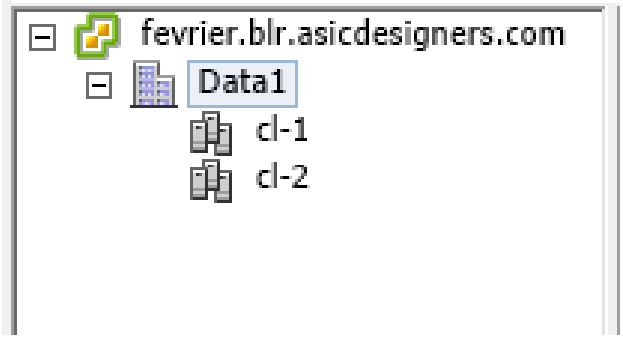

viii. Right-click on the newly created cluster and select **Add Host**. This will invoke the **Add Host Wizard** allowing you to add host(s) to the newly created cluster.

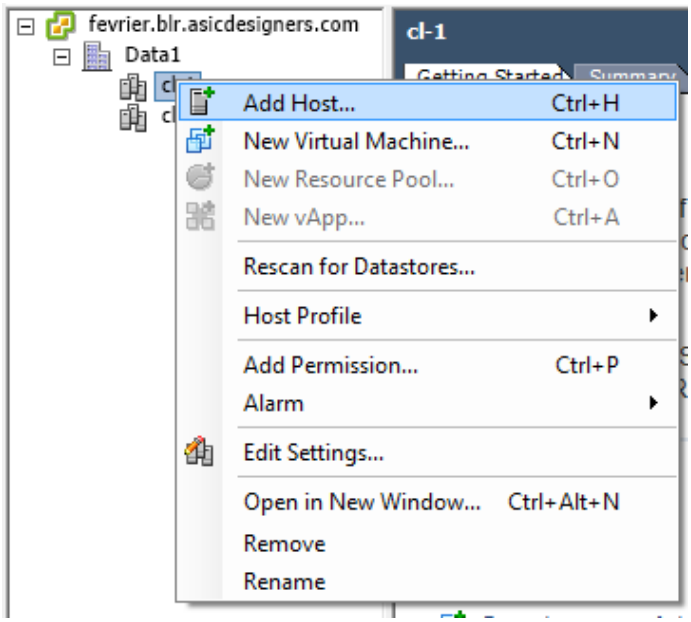

Creating Distributed Switch

Note *Instructions provided here are for desktop client. Steps for web client are similar with minor variations.*

i. Open a vSphere Client session to a vCenter Server.
ii. Select **View** > **Inventory** > **Networking**

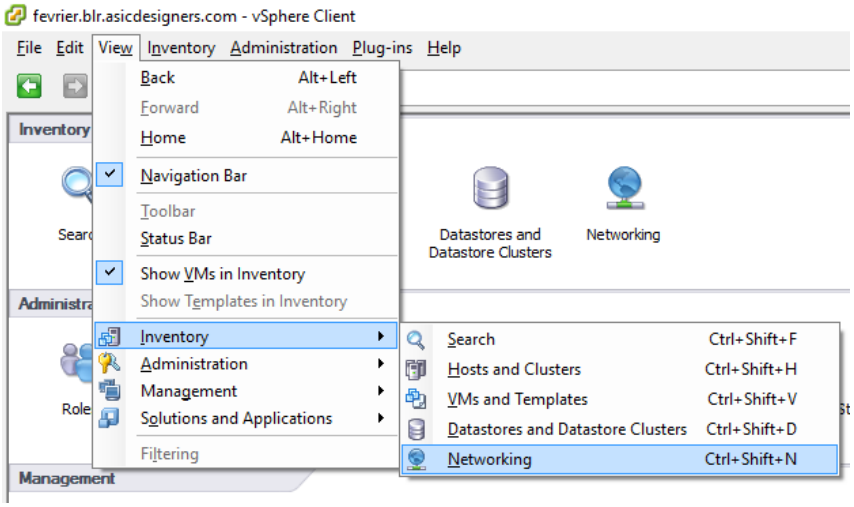

iii. Right-click on your datacenter and select **New vSphere Distributed Switch**.

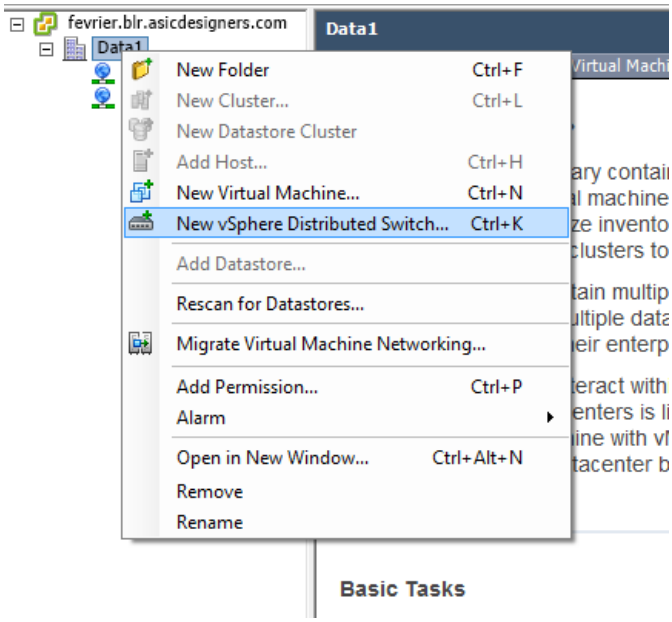

iv. In the setup wizard, select the required switch version and click **Next**.

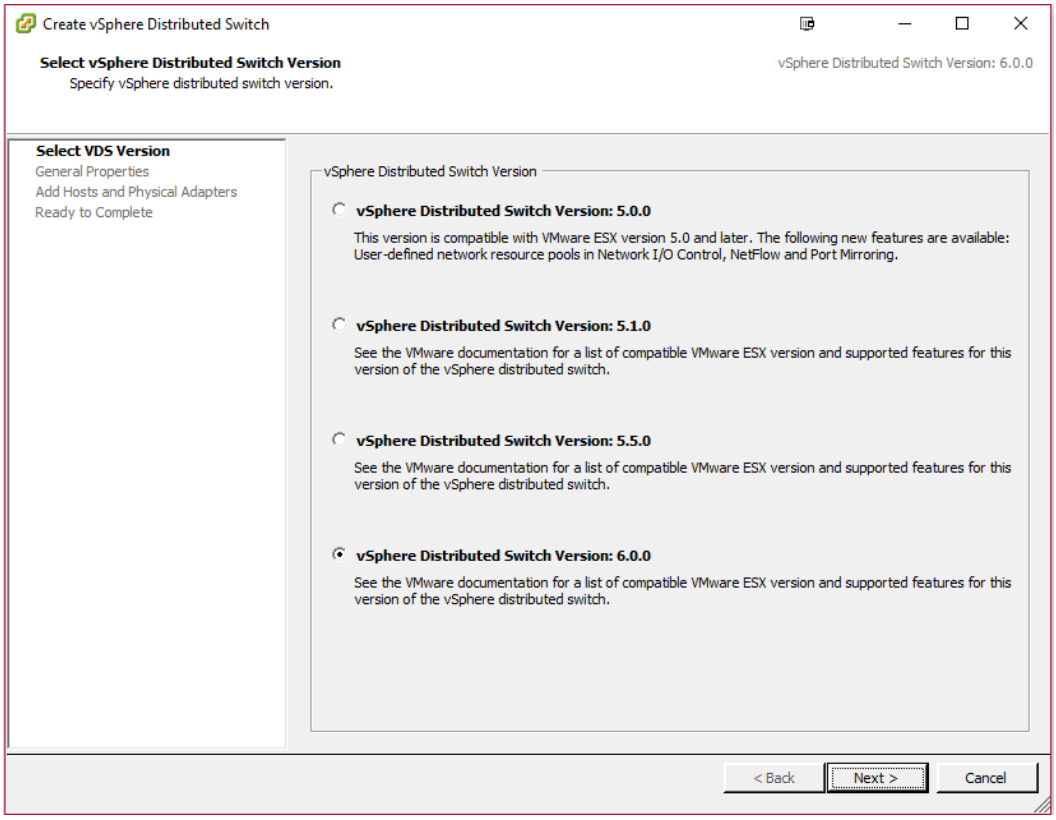

v. Select the required number of uplinks for your setup. Please note that the number of uplinks must be equal or less than the number of adapters to be used for each host. Click **Next**.

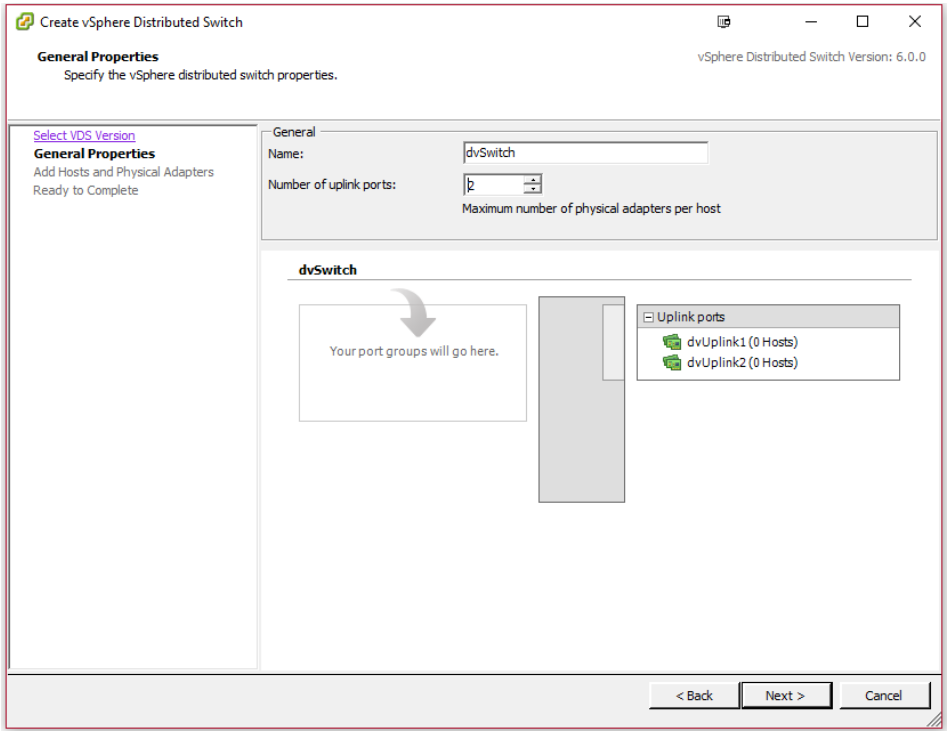

vi. Select the hosts and VMNICs for each host, to be added to the VXLAN network. Click **Next**.

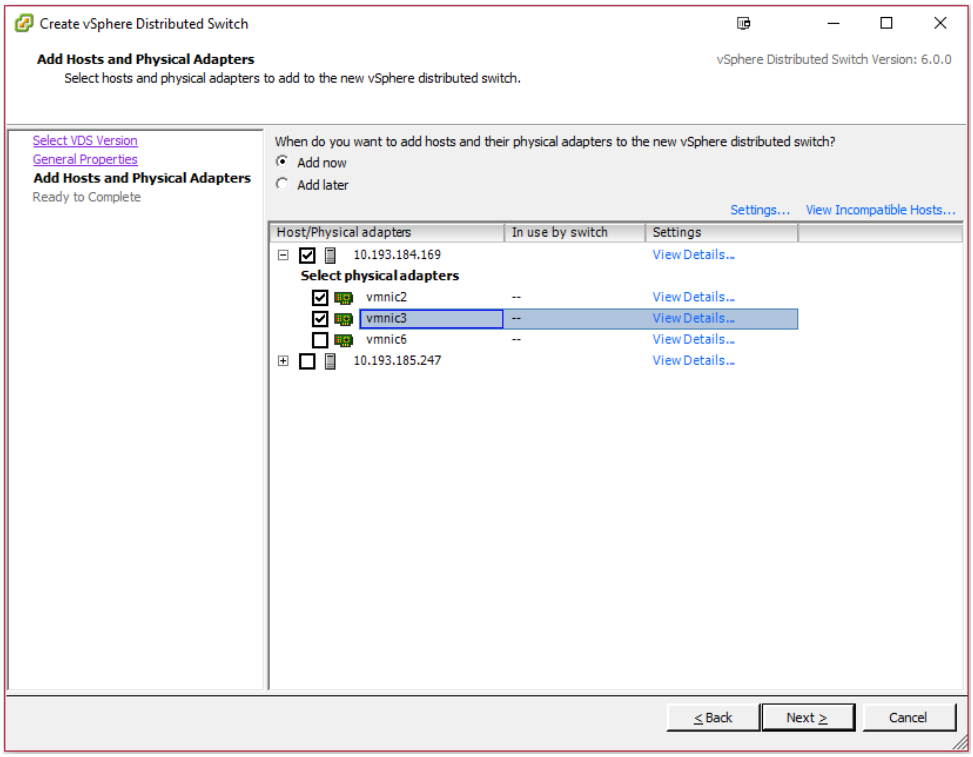

vii. Review the options selected and click **Finish**.

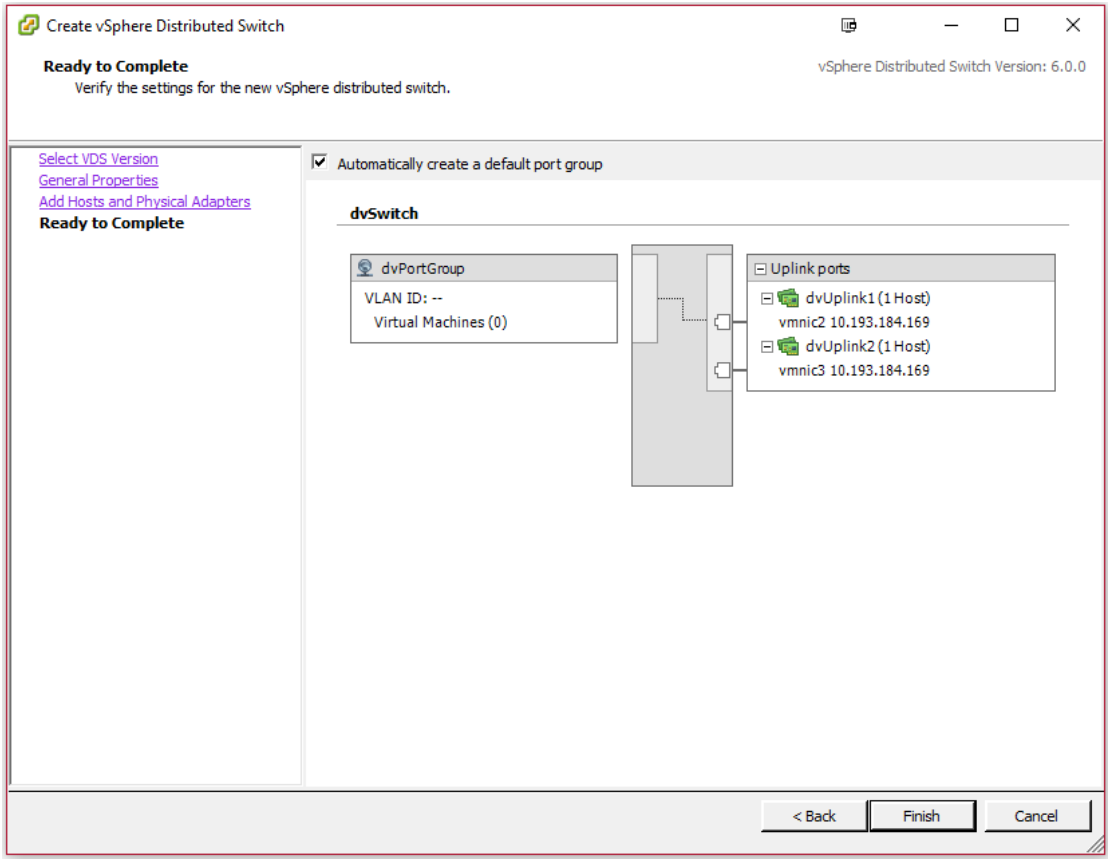

viii. The newly created switch will appear in the inventory.

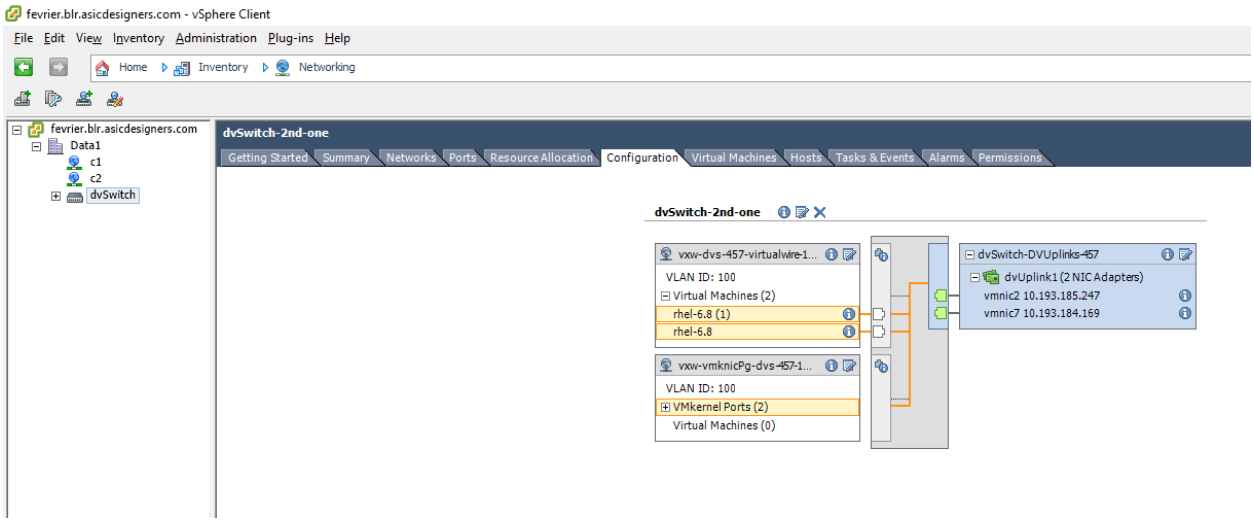

Setting up VXLAN

i. Before proceeding, please follow the steps mentioned in VMware's vShield installation guide to install and configure vShield Manager.

ii. Using a web browser, login to vShield Manager and select your datacenter in the left pane.

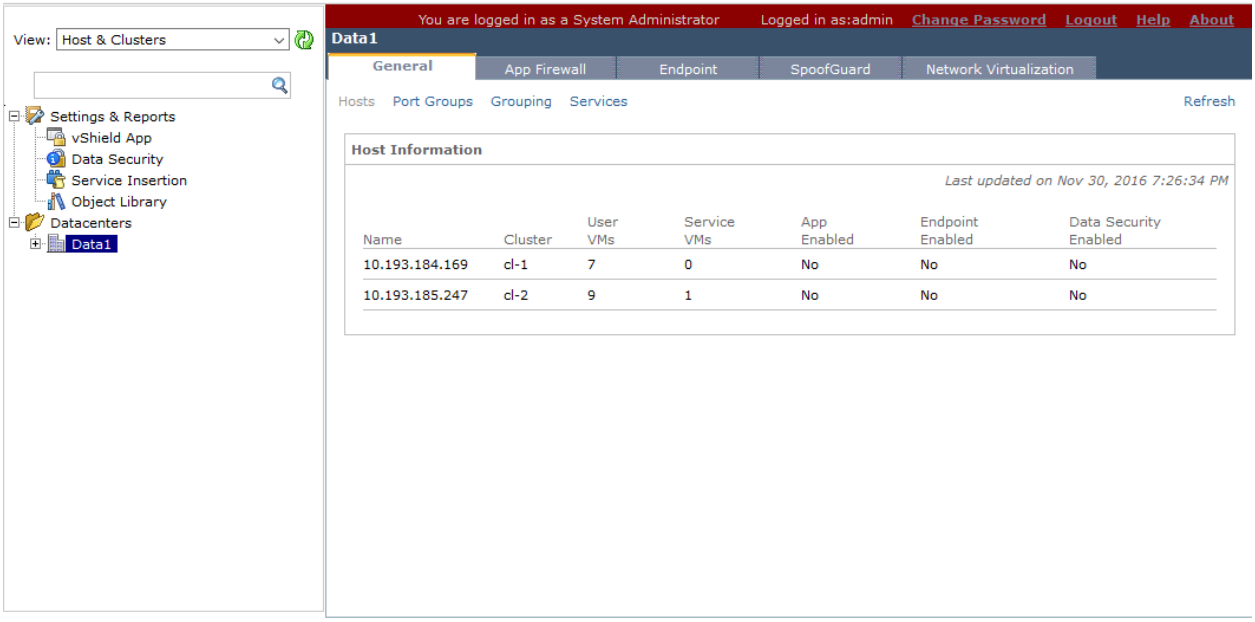

iii. Select the **Network Virtualization** tab on the right pane, then click on the **Preparation** link and then **Edit**. Select the clusters you want to add to the VXLAN. For each of the cluster added, select the designated distributed switch to transport the VXLAN traffic, and specify the VLANID. Click **Next**.

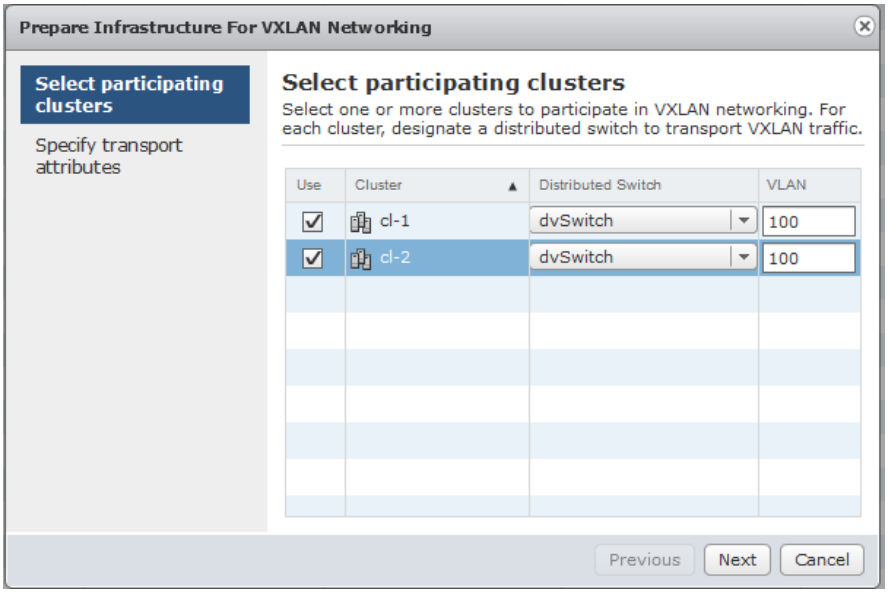

iv. For the distributed switch, select the desired teaming policy and set the MTU to 1600 (to accommodate VXLAN header). Click **Finish**.

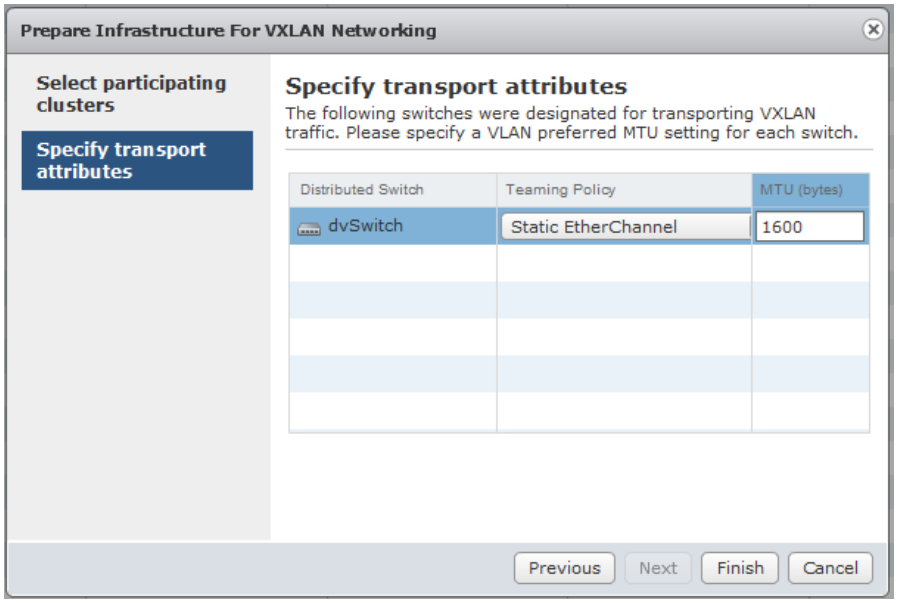

v. The newly added clusters will be initialized now.

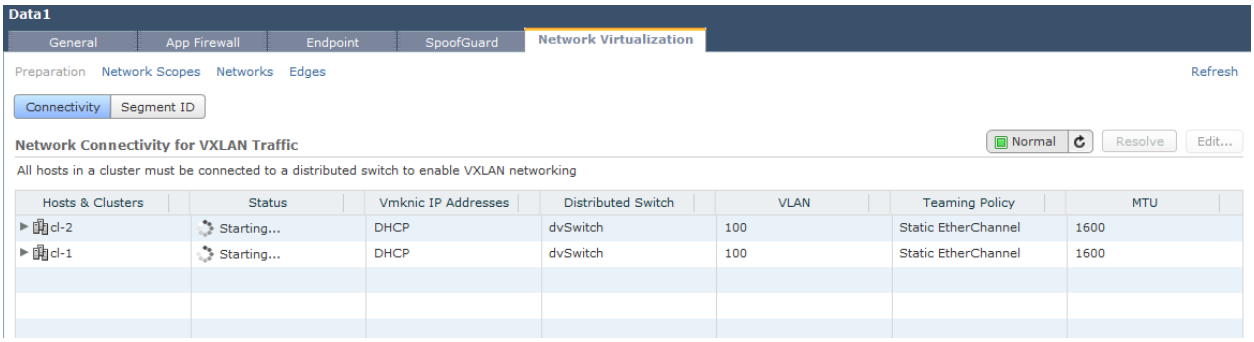

vi. In case DHCP server is not available for clusters, use vSphere client to configure static IPs for the hosts.

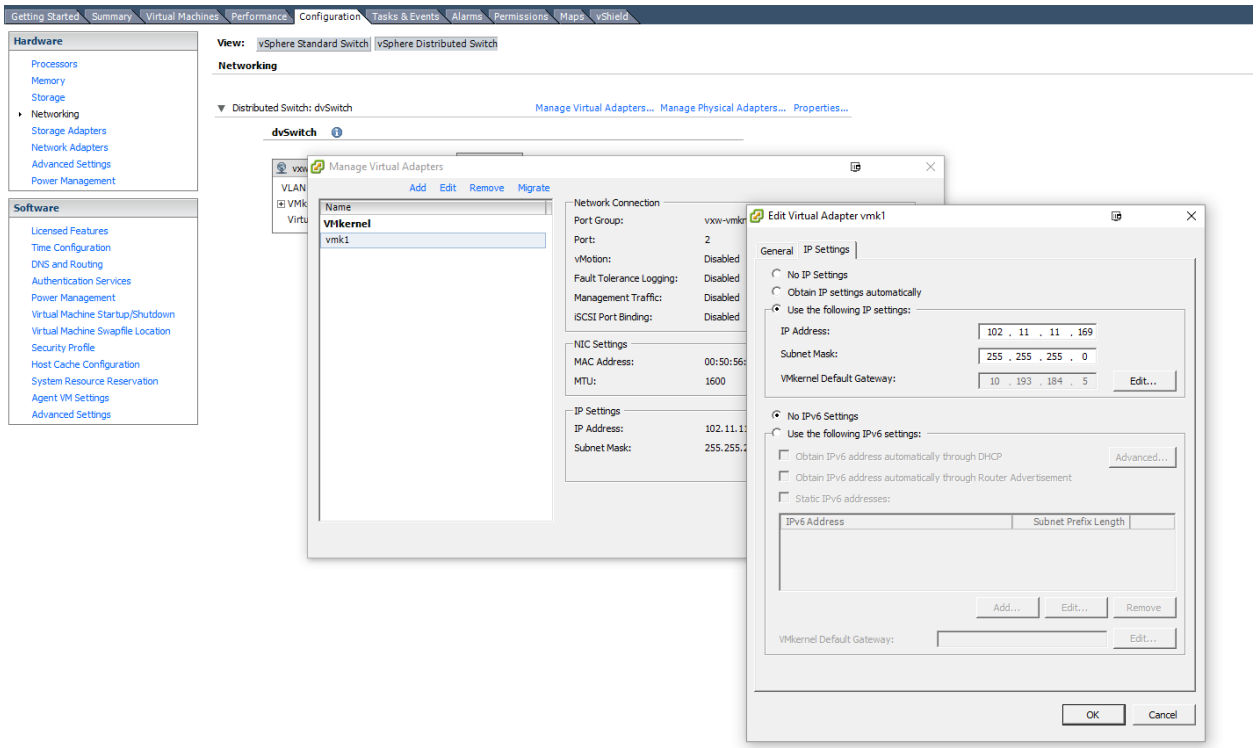

vii. While still in the **Preparation** section, click the **Segment ID** button and then **Edit**. Set the range for *Segment ID pool* and *Multicast addresses*. Click **OK**.

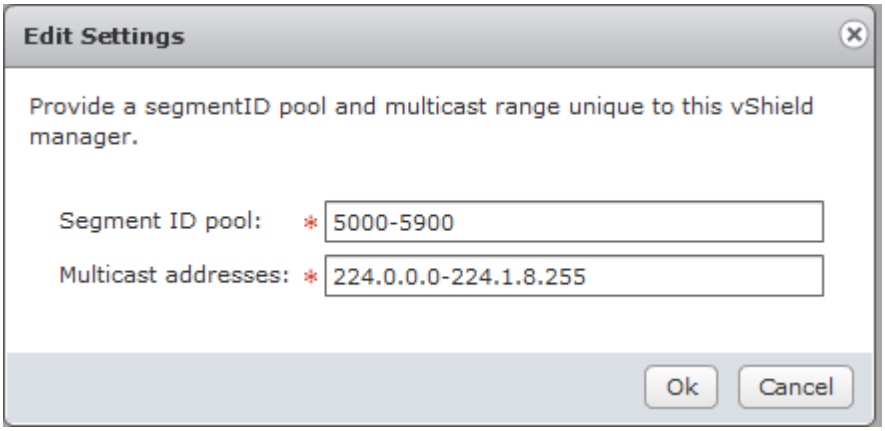

viii. Click the **Network Scopes** link and then ✚. Provide a name and description (optional) for the new network scope. Select the clusters you want to add in this network scope. Click **Ok**.

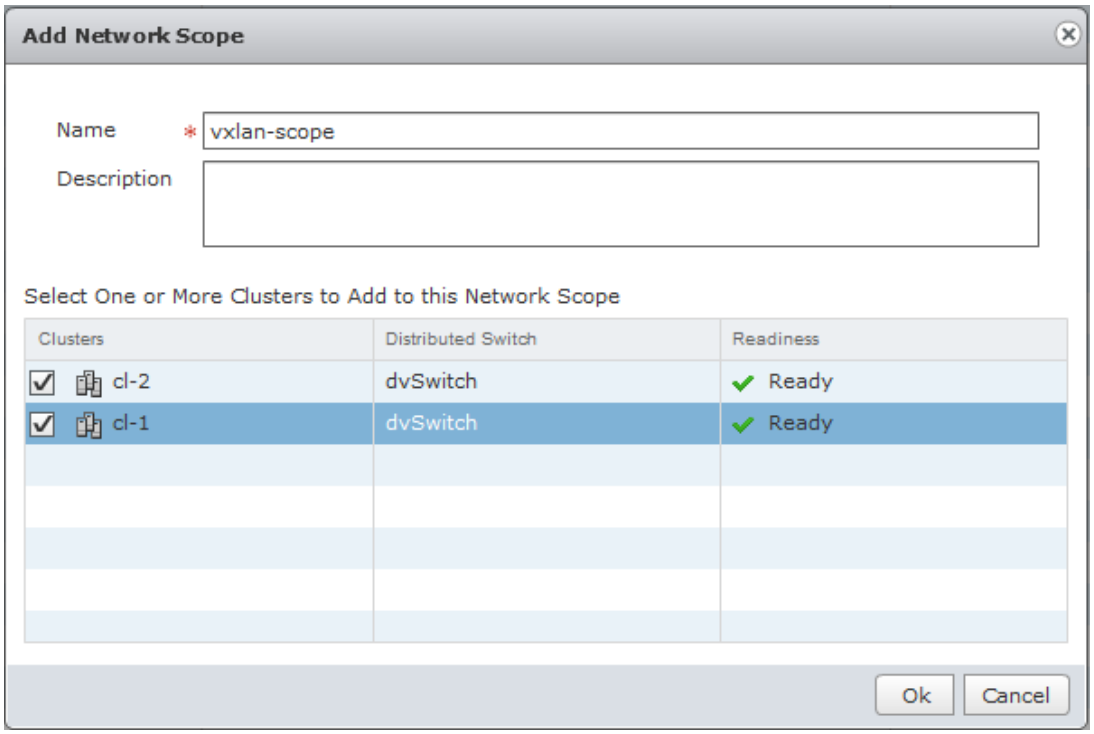

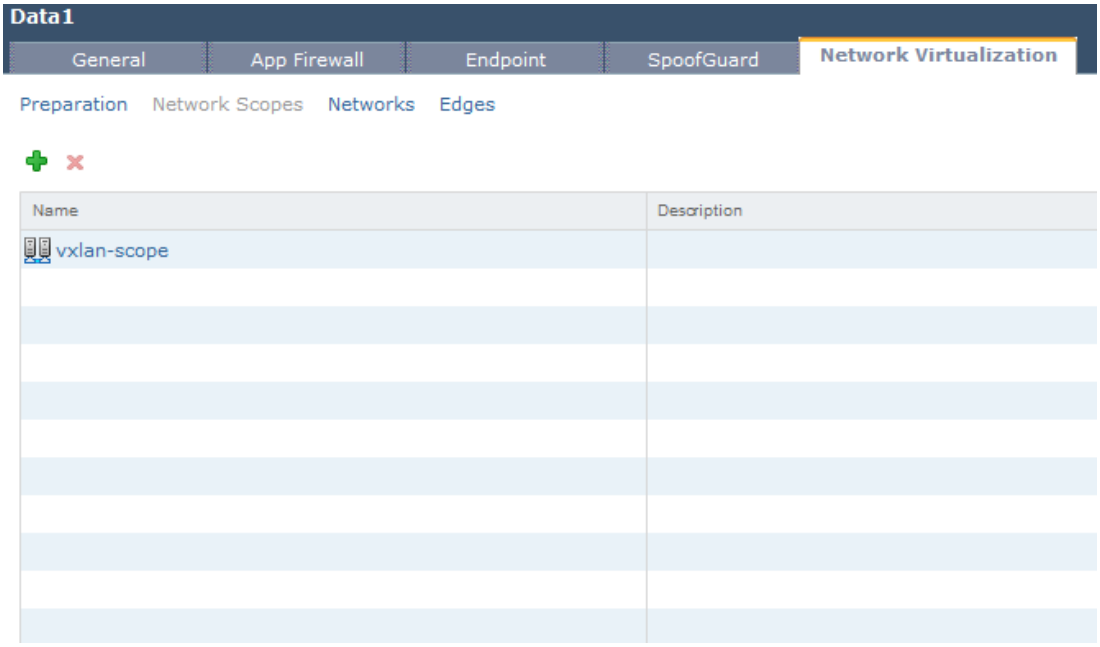

ix. Click the **Networks** link and press the ✚. Provide a name and description (optional) for the new VXLAN. Click on *Network Scope and s*elect the network scope created in the previous step from the drop-down. Click **Ok**.

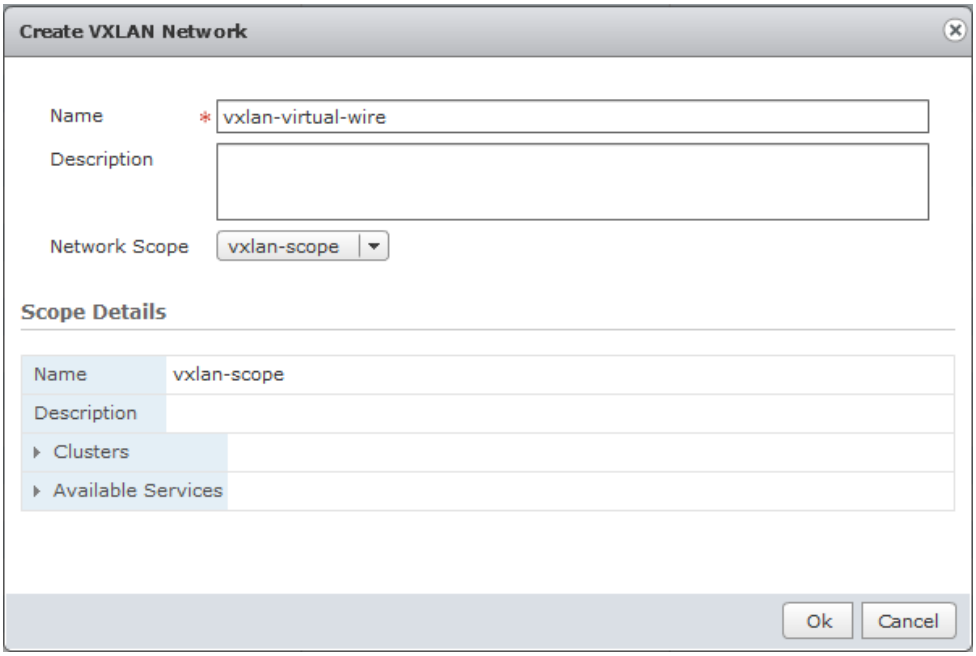

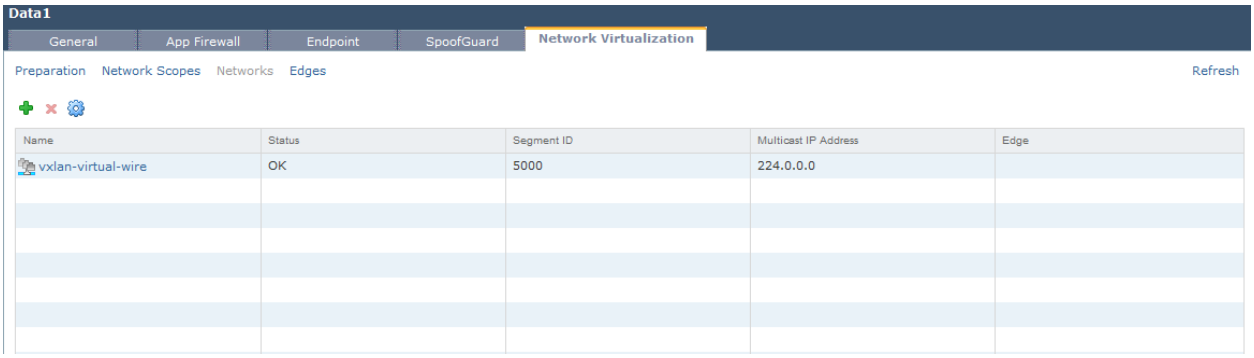

x. In vSphere client, right-click on the VM you want to add to the VXLAN and select *Edit Settings*.

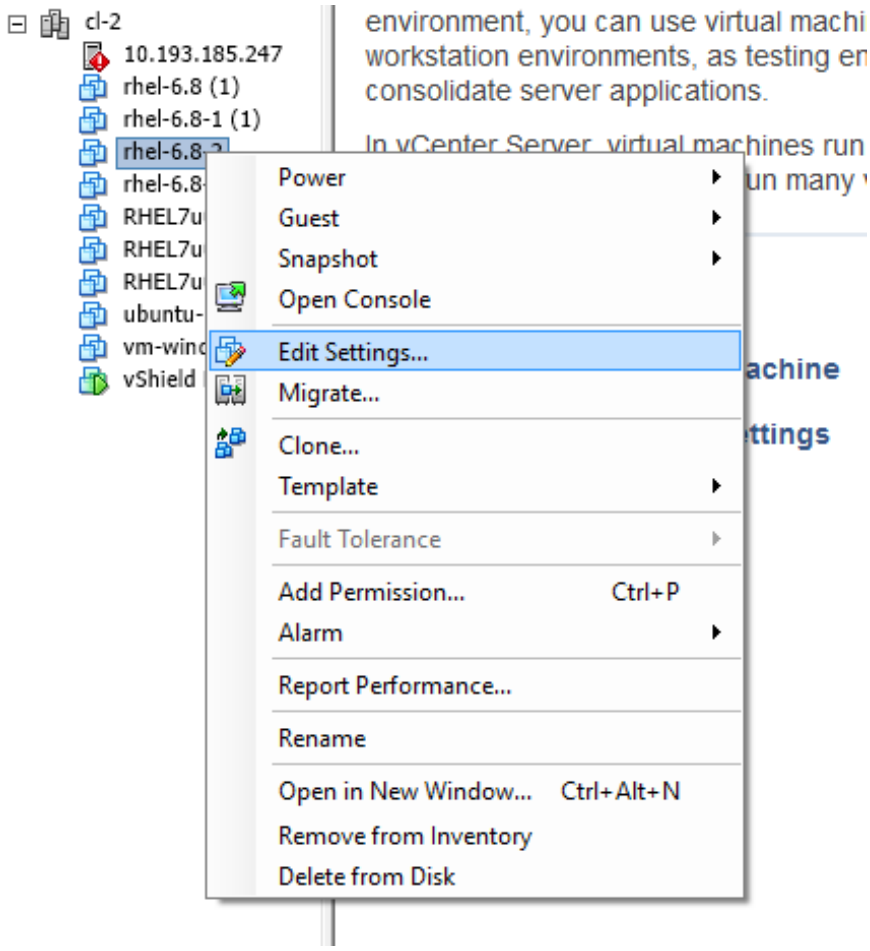

xi. Click **Add**, select *Ethernet Adapter* and then click **Next**.

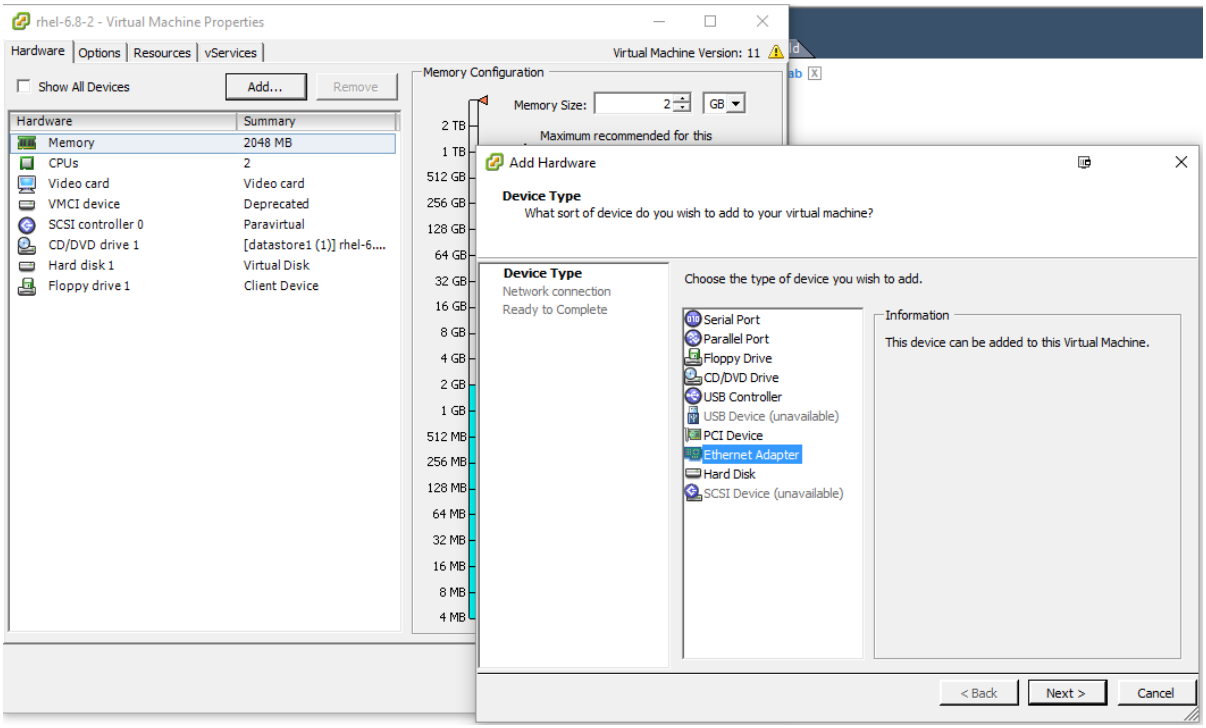

xii. Select the VXLAN created previously from the *Network label* drop-down. Click **Next**.

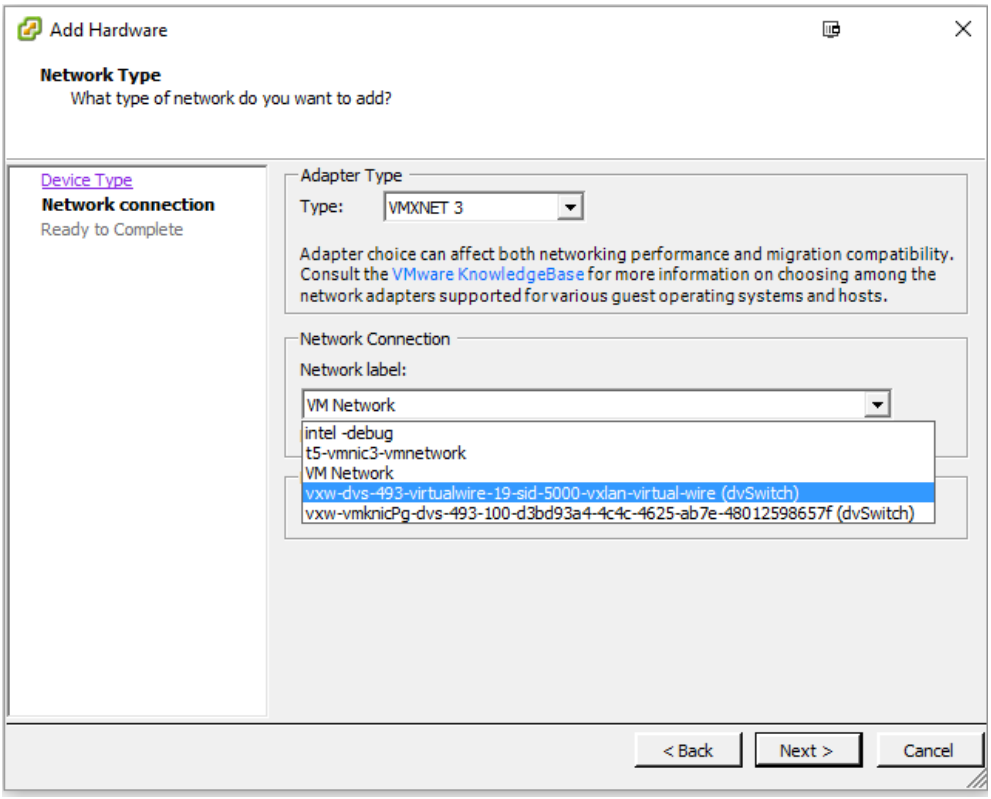

xiii. Review the options selected and click **Finish**.

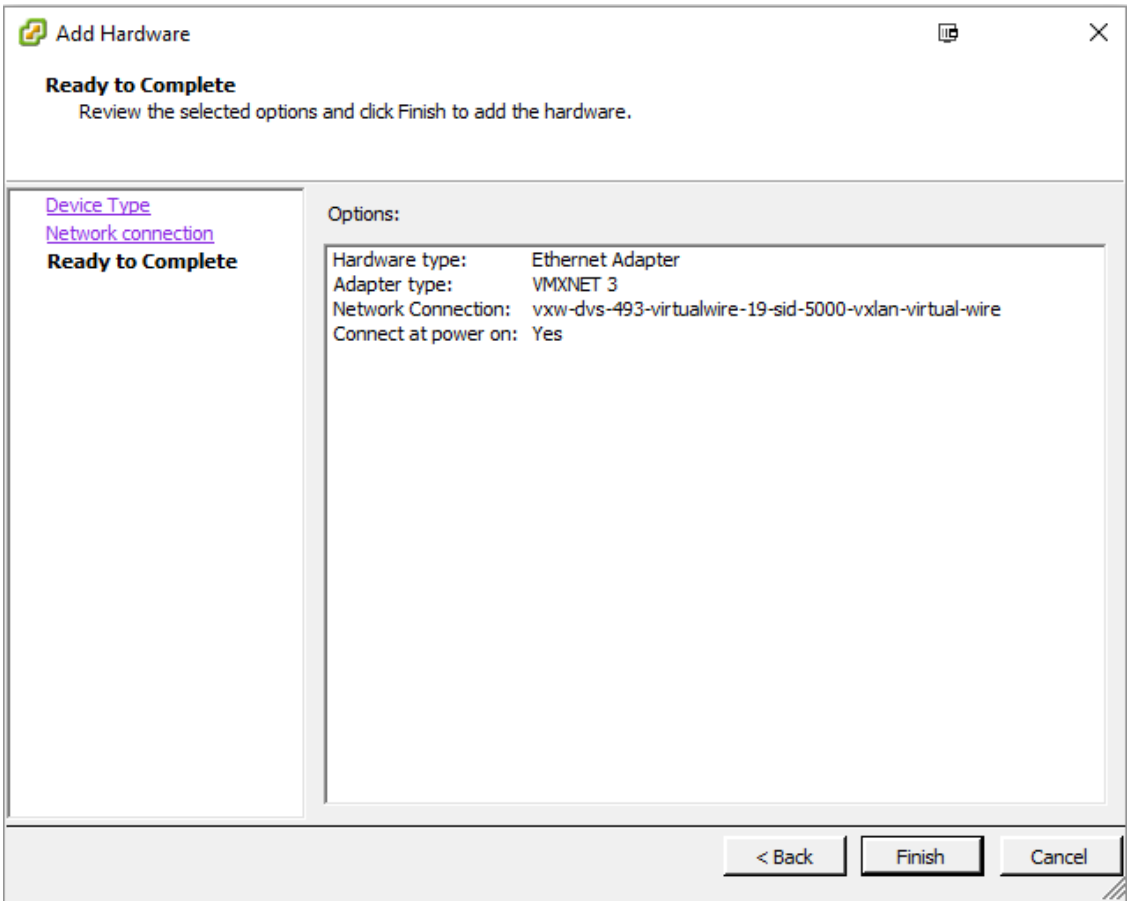

xiv. All the VMs added to the VXLAN will now be able to communicate successfully.

6. Software/Driver Uninstallation

Follow the steps mentioned below to uninstall the driver:

i. Put the host in maintenance mode using the vSphere (desktop or web) Client:

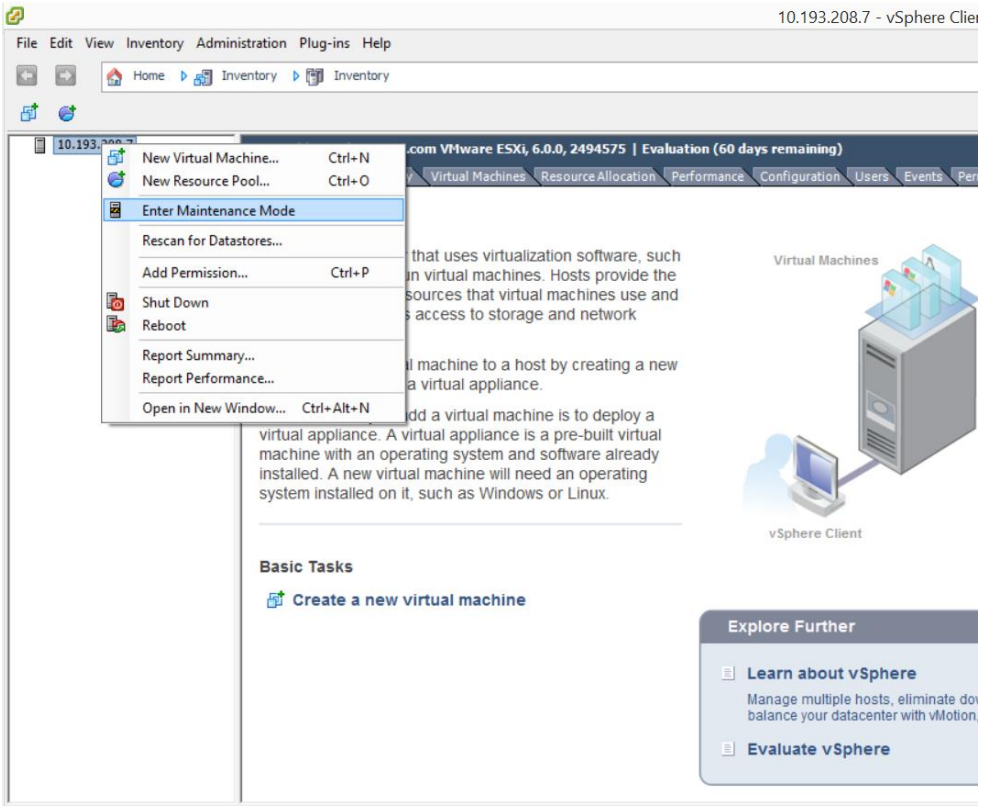

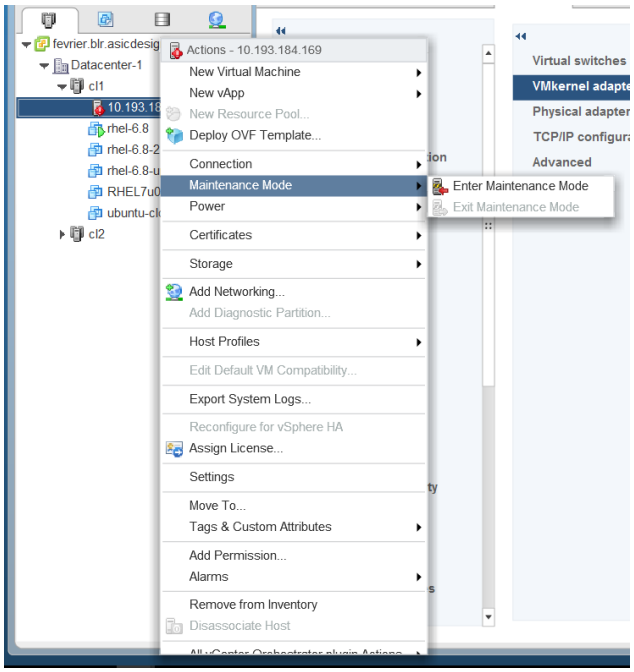

ii. Uninstall the driver:

```
[root@host:~] esxcli software vib remove -n cxl
```

iii. Reboot the host:

```
[root@host:~] reboot
```

7. Software/Driver Update

For any distribution specific problems, please check README and Release Notes included in the release for possible workaround.

Please visit Chelsio support web site http://service.chelsio.com/ for regular updates on various software/drivers. You can also subscribe to our newsletter for the latest software updates.

8. Appendix

8.1. Chelsio End-User License Agreement (EULA)

Installation and use of the driver/software implies acceptance of the terms in the Chelsio End-User License Agreement (EULA).

IMPORTANT: PLEASE READ THIS SOFTWARE LICENSE CAREFULLY BEFORE DOWNLOADING OR OTHERWISE USING THE SOFTWARE OR ANY ASSOCIATED DOCUMENTATION OR OTHER MATERIALS (COLLECTIVELY, THE "SOFTWARE"). BY CLICKING ON THE "OK" OR "ACCEPT" BUTTON YOU AGREE TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, CLICK THE "DO NOT ACCEPT" BUTTON TO TERMINATE THE INSTALLATION PROCESS.

1. License. Chelsio Communications, Inc. ("Chelsio") hereby grants you, the Licensee, and you hereby accept, a limited, non-exclusive, non-transferable license to install and use the Software with one or more Chelsio network adapters on a single server computer for use in communicating with one or more other computers over a network. You may also make one copy of the Software in machine readable form solely for back-up purposes, provided you reproduce Chelsio's copyright notice and any proprietary legends included with the Software or as otherwise required by Chelsio.

2. Restrictions. This license granted hereunder does not constitute a sale of the Software or any copy thereof. Except as expressly permitted under this Agreement, you may not:

(i) reproduce, modify, adapt, translate, rent, lease, loan, resell, distribute, or create derivative works of or based upon, the Software or any part thereof; or

(ii) make available the Software, or any portion thereof, in any form, on the Internet. The Software contains trade secrets and, in order to protect them, you may not decompile, reverse engineer, disassemble, or otherwise reduce the Software to a human-perceivable form. You assume full responsibility for the use of the Software and agree to use the Software legally and responsibly.

3. Ownership of Software. As Licensee, you own only the media upon which the Software is recorded or fixed, but Chelsio retains all right, title and interest in and to the Software and all subsequent copies of the Software, regardless of the form or media in or on which the Software may be embedded.

4. Confidentiality. You agree to maintain the Software in confidence and not to disclose the Software, or any information or materials related thereto, to any third party without the express written consent of Chelsio. You further agree to take all reasonable precautions to limit access of the Software only to those of your employees who reasonably require such access to perform their employment obligations and who are bound by confidentiality agreements with you.

5. Term. This license is effective in perpetuity, unless terminated earlier. You may terminate the license at any time by destroying the Software (including the related documentation), together with all copies or modifications in any form. Chelsio may terminate this license, and this license shall be deemed to have automatically terminated, if you fail to comply with any term or condition of this Agreement. Upon any termination, including termination by you, you must destroy the Software

(including the related documentation), together with all copies or modifications in any form.

6. Limited Warranty. If Chelsio furnishes the Software to you on media, Chelsio warrants only that the media upon which the Software is furnished will be free from defects in material or workmanship under normal use and service for a period of thirty (30) days from the date of delivery to you.

CHELSIO DOES NOT AND CANNOT WARRANT THE PERFORMANCE OR RESULTS YOU MAY OBTAIN BY USING THE SOFTWARE OR ANY PART THEREOF. EXCEPT FOR THE FOREGOING LIMITED WARRANTY, CHELSIO MAKES NO OTHER WARRANTIES, EXPRESS OR IMPLIED, AND HEREBY DISCLAIMS ALL OTHER WARRANTIES, INCLUDING, BUT NOT LIMITED TO, NON-INFRINGEMENT OF THIRD PARTY RIGHTS, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow the exclusion of implied warranties or limitations on how long an implied warranty may last, so the above limitations may not apply to you. This warranty gives you specific legal rights and you may also have other rights which vary from state to state.

7. Remedy for Breach of Warranty. The sole and exclusive liability of Chelsio and its distributors, and your sole and exclusive remedy, for a breach of the above warranty, shall be the replacement of any media furnished by Chelsio not meeting the above limited warranty and which is returned to Chelsio. If Chelsio or its distributor is unable to deliver replacement media which is free from defects in materials or workmanship, you may terminate this Agreement by returning the Software.

8. Limitation of Liability. IN NO EVENT SHALL CHELSIO HAVE ANY LIABILITY TO YOU OR ANY THIRD PARTY FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, HOWEVER CAUSED, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR RELATED TO THE LICENSE OR USE OF THE SOFTWARE, INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR LOSS OF ANTICIPATED PROFITS, EVEN IF CHELSIO HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL CHELSIO'S LIABILITY ARISING OUT OF OR RELATED TO THE LICENSE OR USE OF THE SOFTWARE EXCEED THE AMOUNTS PAID BY YOU FOR THE LICENSE GRANTED HEREUNDER. THESE LIMITATIONS SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

9. High Risk Activities. The Software is not fault-tolerant and is not designed, manufactured or intended for use or resale as online equipment control equipment in hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life support machines, or weapons systems, in which the failure of the Software could lead directly to death, personal injury, or severe physical or environmental damage. Chelsio specifically disclaims any express or implied warranty of fitness for any high risk uses listed above.

10. Export. You acknowledge that the Software is of U.S. origin and subject to U.S. export jurisdiction. You acknowledge that the laws and regulations of the United States and other countries may restrict the export and re-export of the Software. You agree that you will not export or re-export the Software or documentation in any form in violation of applicable United States and foreign law. You agree to comply with all applicable international and national laws that apply to the Software, including the U.S.

Export Administration Regulations, as well as end-user, end-use, and destination restrictions issued by U.S. and other governments.

11. Government Restricted Rights. The Software is subject to restricted rights as follows. If the Software is acquired under the terms of a GSA contract: use,

reproduction or disclosure is subject to the restrictions set forth in the applicable ADP Schedule contract. If the Software is acquired under the terms of a DoD or civilian agency contract, use, duplication or disclosure by the Government is subject to the restrictions of this Agreement in accordance with 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors and 49 C.F.R. 227.7202-1 of the DoD FAR Supplement and its successors.

12. General. You acknowledge that you have read this Agreement, understand it, and that by using the Software you agree to be bound by its terms and conditions. You further agree that it is the complete and exclusive statement of the agreement between Chelsio and you, and supersedes any proposal or prior agreement, oral or written, and any other communication between Chelsio and you relating to the subject matter of this Agreement. No additional or any different terms will be enforceable against Chelsio unless Chelsio gives its express consent, including an express waiver of the terms of this Agreement, in writing signed by an officer of Chelsio. This Agreement shall be governed by California law, except as to copyright matters, which are covered by Federal law. You hereby irrevocably submit to the personal jurisdiction of, and irrevocably waive objection to the laying of venue (including a waiver of any argument of forum non conveniens or other principles of like effect) in, the state and federal courts located in Santa Clara County, California, for the purposes of any litigation undertaken in connection with this Agreement. Should any provision of this Agreement be declared unenforceable in any jurisdiction, then such provision shall be deemed severable from this Agreement and shall not affect the remainder hereof. All rights in the Software not specifically granted in this Agreement are reserved by Chelsio. You may not assign or transfer this Agreement (by merger, operation of law or in any other manner) without the prior written consent of Chelsio and any attempt to do so without such consent shall be void and shall constitute a material breach of this Agreement.

Should you have any questions concerning this Agreement, you may contact Chelsio by writing to:

Chelsio Communications, Inc.
209 North Fair Oaks Avenue,
Sunnyvale, CA 94085
U.S.A