



Chelsio Unified Wire for VMware ESXi 6.7

Installation and User's Guide



This document and related products are distributed under licenses restricting their use, copying, distribution, and reverse-engineering.

No part of this document may be reproduced in any form or by any means without prior written permission by Chelsio Communications.

All third-party trademarks are copyright of their respective owners.

THIS DOCUMENTATION IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

THE USE OF THE SOFTWARE AND ANY ASSOCIATED MATERIALS (COLLECTIVELY THE “SOFTWARE”) IS SUBJECT TO THE SOFTWARE LICENSE TERMS OF CHELSIO COMMUNICATIONS, INC.



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

Copyright © 2018. Chelsio Communications. All Rights Reserved.

Chelsio® is a registered trademark of Chelsio Communications.

All other marks and names mentioned herein may be trademarks of their respective companies.

Version History

Version	Revision Date
1.0.0	02/13/2018
1.0.1	02/27/2018
1.0.2	03/27/2018
1.0.3	07/06/2018
1.0.4	10/05/2018
1.0.5	10/08/2018

TABLE OF CONTENTS

I.	CHELSIO UNIFIED WIRE	5
1.	Introduction	6
1.1.	Features	6
1.2.	Hardware Requirements	6
1.3.	Software Requirements	7
1.4.	Package Contents	7
2.	Hardware Installation	8
3.	Software/Driver Update	10
II.	NATIVE NETWORK DRIVER WITH SR-IOV SUPPORT	11
1.	Introduction	12
1.1.	Hardware Requirements	12
1.2.	Software Requirements	12
2.	Software/Driver Installation	14
3.	Software/Driver Loading	16
4.	Software/Driver Configuration and Fine-tuning	17
4.1.	cxgbtool	17
4.2.	Adapter Configuration	17
4.3.	Firmware Update	17
4.4.	Connecting a Virtual Machine	18
4.5.	Tuning vMotion for vSAN	18
4.6.	Virtual Functions (SR-IOV)	19
4.7.	VXLAN	22
5.	Software/Driver Unloading	43
6.	Software/Driver Uninstallation	44
III.	ISCSI/ISER OFFLOAD INITIATOR DRIVER	46
1.	Introduction	47
1.1.	Hardware Requirements	47
1.2.	Software Requirements	48
2.	Software/Driver Installation	49
3.	Software/Driver Loading	51
4.	Software/Driver Configuration and Fine-tuning	52
4.1.	Configuring Initiator	52
4.2.	Connecting to Target	55
4.3.	Configurable Options	62
5.	Software/Driver Unloading	63
6.	Software/Driver Uninstallation	64
IV.	APPENDIX	65
	Chelsio End-User License Agreement (EULA)	66

I. Chelsio Unified Wire

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 Unified Wire software for VMware ESXi and its various components.

1.1. Features

The Chelsio Unified Wire package installs various drivers and utilities and consists of the following software:

- **Native Network (NIC) driver with SR-IOV support**
- **iSCSI/iSER Offload Initiator Driver**

 **Note** *Drivers are not VMware certified.*

For detailed instructions on loading, unloading and configuring the drivers/tools please refer to their respective sections.

1.2. Hardware Requirements

The Chelsio Unified Wire software supports Chelsio Terminator series of Unified Wire adapters. To know more about the list of adapters supported by each driver, please refer to their respective sections.

1.3. Software Requirements

The Chelsio Unified Wire software has been developed to run on 64-bit ESXi based platforms and therefore it is a base requirement for running the driver. To know more about the complete list of operating systems supported by each driver, please refer to their respective sections.

1.4. Package Contents

The Chelsio Unified Wire driver package consists of the following files/directories:

- **cxl-*.vib**: Native Network driver VIB file.
- **cheiscsi-*.vib**: iSCSI/iSER Offload Initiator driver VIB file.
- **chelsio-esx-libcheiscsi_ima.so-*.vib**: IMA driver file.
- **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.



Note

All 4-ports of T6425-CR adapter will be functional only if PCIe x8 -> 2x PCIe x4 slot bifurcation is supported by the system and enabled in BIOS. Otherwise, only 2-ports will be functional.

- 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@ ~]# 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 NIC, SR-IOV functions of the adapter
- Physical function 4: Currently not used
- Physical function 5: for iSCSI
- Physical function 6: for FCoE
- Physical function 7: Currently not assigned

Once native network driver (`cxl`) is installed and loaded, examine the output of `dmesg` to see if the card is discovered. You should see a similar output:


```
2017-09-26T04:09:20.207Z cpu6:66032)cxl1.0: cxl_port_init:874: mbox 0 pf 0 chan 0 viid c0
2017-09-26T04:09:20.209Z cpu6:66032)DMA: 646: DMA Engine 'cxl-0000:04:00.0' created using mapper 'DMANull'.
2017-09-26T04:09:20.209Z cpu6:66032)cxl1.0: cxl_config_queues:1091: max_filters 120
2017-09-26T04:09:20.209Z cpu6:66032)VMK PCI: 765: device 0000:04:00.0 allocated 32 MSI-X interrupts
2017-09-26T04:09:20.209Z cpu6:66032)cxl1.0: cxl_intr_alloc_msix:2581: net q 14 rss q 16 non rss q 13 tx q 8
2017-09-26T04:09:20.211Z cpu6:66032)cxl1.0: cxl_rss_do_init:5221: pool 0 rss viid c1
2017-09-26T04:09:20.212Z cpu6:66032)cxl1.0: cxl_rss_init:2501: pool 0 rss mode 31
2017-09-26T04:09:20.212Z cpu6:66032)Chelsio T6225-CR rev 0 25G NIC PCIe 8 GT/s x8 MSI-X S/N: RE35160002, P/N: 11012096002
```

The above outputs indicate the hardware configuration of the adapter as well as serial number.



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 Update

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

Please visit [Chelsio Download Center](#) for regular updates on various software/drivers. You can also subscribe to our newsletter for the latest software updates.

II. Native Network Driver with SR-IOV Support

1. Introduction

Chelsio's Unified Wire adapters provide extensive support for NIC operation. A high performance fully offloaded and fully featured TCP/IP stack meets or exceeds software implementations in RFC compliance. Chelsio's Terminator engine provides unparalleled performance through a specialized data flow processor implementation and a host of features designed for high throughput and low latency in demanding conditions and networking environments.

1.1. Hardware Requirements

1.1.1. 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
- T61100-OCF
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR
- T580-SO-CR
- T580-LP-CR
- T580-CR
- T540-CR
- T540-LP-CR
- T540-SO-CR
- T520-SO-CR
- T520-CR
- T520-LL-CR
- T520-BT
- T540-BT

1.2. Software Requirements

1.2.1. ESXi Requirements

The Chelsio native network driver has been developed to run on 64-bit ESXi platforms. Currently the driver is available for the following versions:

- **Host:**
 - ESXi 6.7

- **Virtual Machine (with VFs):**
 - RHEL 7.5, 3.10.0-862.el7
 - RHEL 7.4, 3.10.0-693.el7
 - RHEL 6.9, 2.6.32-696.el6
 - SLES 15, 4.12.14-23-default
 - SLES 12 SP3, 4.4.73-5-default
 - SLES 12 SP2, 4.4.21-69-default
 - Ubuntu 18.04.1, 4.15.0-29-generic
 - Ubuntu 16.04.4, 4.4.0-116-generic
 - Kernel.org linux-4.9 (minimum 4.9 kernel version supported is 4.9.13)
 - Kernel.org linux-4.14.67



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

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.

2. Software/Driver Installation

- i. Download the driver package from [Chelsio Download Center](#).
- 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.

3. 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)
```


4. Software/Driver Configuration and Fine-tuning


4.1. 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.

4.2. Adapter Configuration

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

- Run the following `cxgbtool` command:

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

 **Note** Not supported on T6 adapters.

- Reboot the system for changes to take effect.

4.3. 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
```

4.4. 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.

4.5. 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
```

4.6. Virtual Functions (SR-IOV)

4.6.1. 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@    :~] 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]
```

ESXi 6.0.0, 2494585 | Evaluation (45 days remaining)

Resource Allocation Performance Configuration Users Events Permissions

DirectPath I/O Configuration

Warning: Configuring host hardware without special virtualization features for virtual machine passthrough will make it unavailable for use except a device needed for normal host boot or operation can make normal host boot impossible and may require significant effort to undo. See the online help for more details.

Each listed device is available for direct access by the virtual machines on this host.

0000:05:01.0 | Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller
 0000:05:01.1 | Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller
 0000:05:01.4 | Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller
 0000:05:01.5 | Chelsio Communications Inc T580-LP-CR Unified Wire Ethernet Controller

Device Details			
Device Name	--	Vendor Name	--
ID	--	Class ID	--
Device ID	--	Subdevice ID	--
Vendor ID	--	Subvendor ID	--
Function	--	Slot	--
Bus	--		

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:

$$pciPassthru<VF_ID>.maxMSIXvectors = \text{<Number of CPUs in Win VM>} + 2$$

For more information refer to [VMware documentation](#).

- Windows Guest is not supported with SR-IOV.

4.6.2. 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](#).

4.6.3. 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](#).
- iv. Untar the tarball using the following command:


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

- v. Change your 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*.

4.6.4. 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.

4.7. 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:

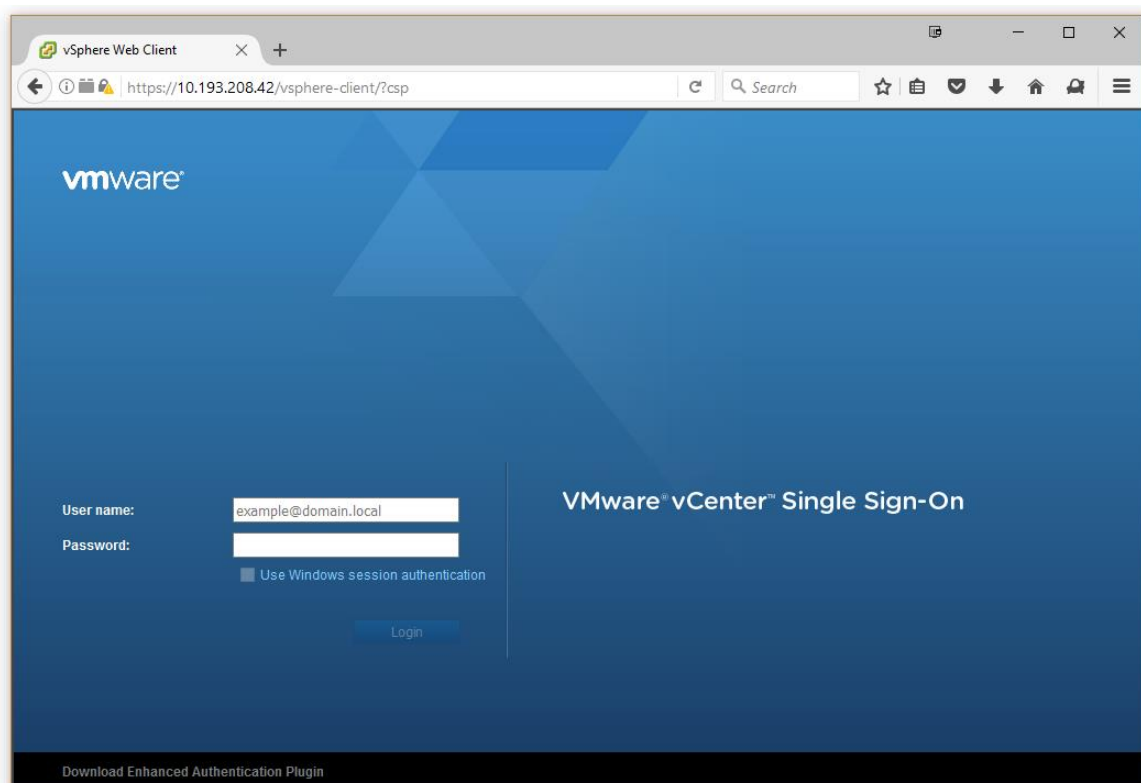
4.7.1. ESXi 6.7

- **Configuring Adapters**

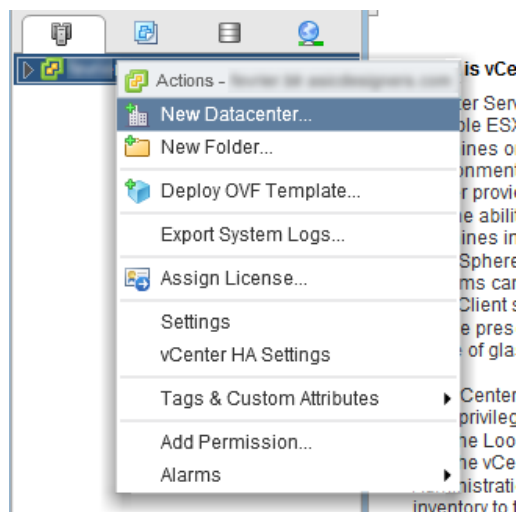
For T6 adapters, VXLAN offload is enabled by default on loading the driver.

- **Creating Cluster**


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



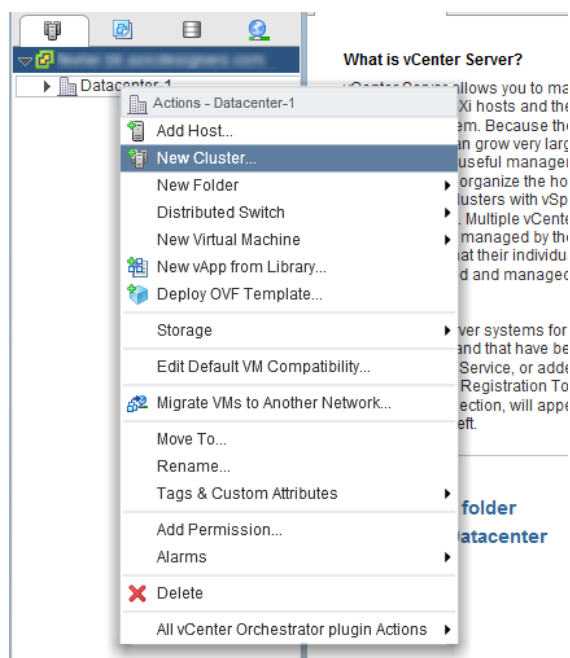
- ii. Click on **Hosts and Clusters** tab in the **Object Navigator**. Right-click on your vCenter server and select **New Datacenter**.



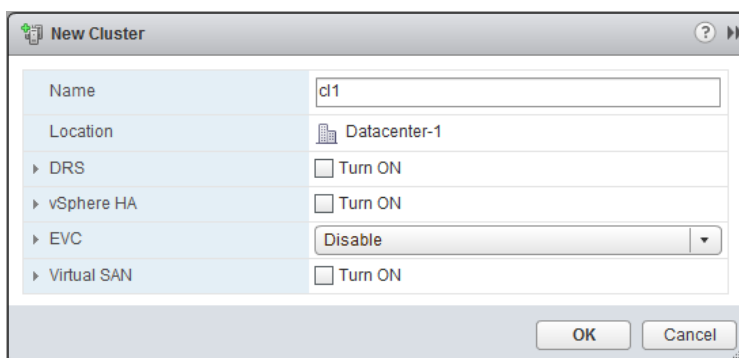
- iii. Provide a name for the new datacenter and click **OK**.




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



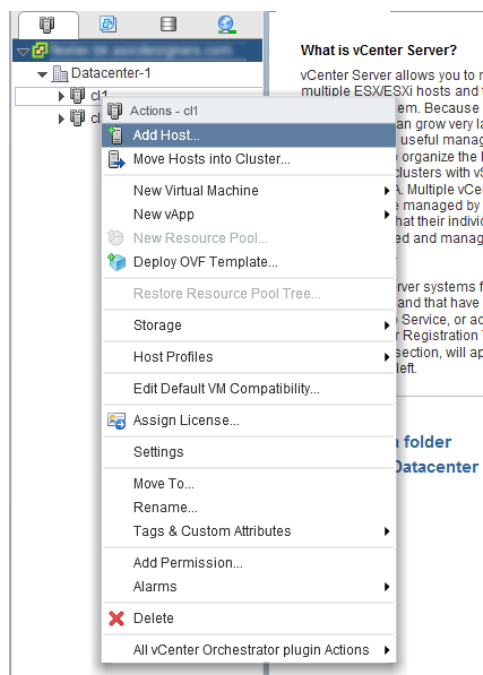
- v. Enter a name for the new cluster and select options for cluster features. Click **OK**.



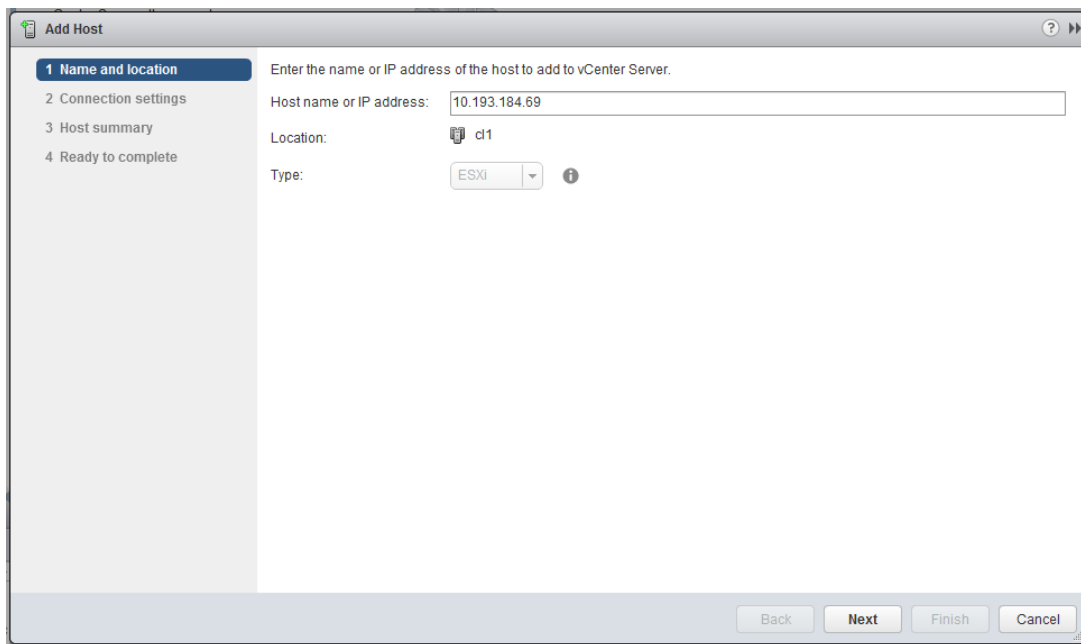
- vi. The newly created cluster will appear in the inventory. Repeat the above step to add more clusters to the datacenter.



- 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.



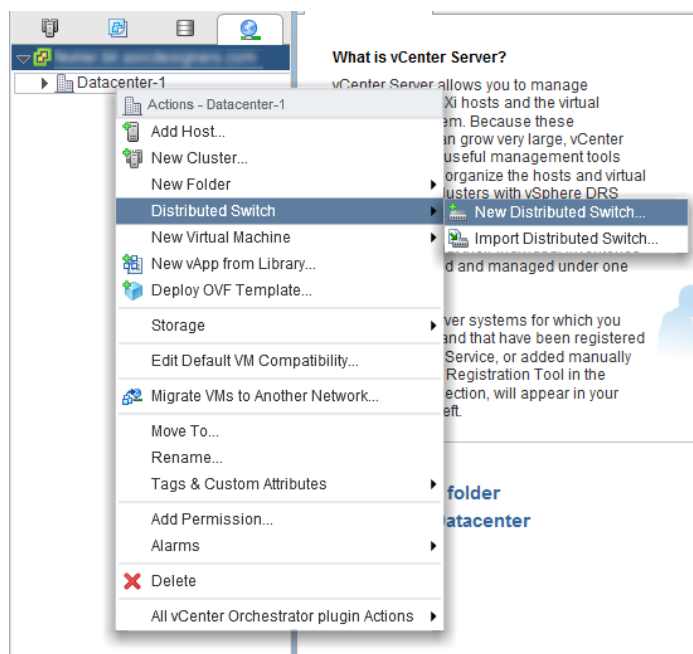
- viii. Provide name or IP address of the host, username and password, and other configuration details. Click **Next**.



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

• Creating Distributed Switch

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



iv. Enter a name for the switch and click **Next**.

The screenshot shows the 'New Distributed Switch' wizard window. On the left, a sidebar lists four steps: 1 Name and location (selected), 2 Select version, 3 Edit settings, and 4 Ready to complete. The main area is titled 'Name and location' and contains the instruction 'Specify distributed switch name and location.' Below this, there are two input fields: 'Name:' with the text 'DSwitch' entered, and 'Location:' with a dropdown menu showing 'Datacenter-1'. At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

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

The screenshot shows the 'New Distributed Switch' wizard window at Step 2: Select version. The sidebar on the left shows the first step '1 Name and location' with a green checkmark, and the second step '2 Select version' is selected. The main area is titled 'Select version' and contains the instruction 'Specify a distributed switch version.' Below this, there are five radio button options, each with a version number and a description of compatibility and features:

- ☒ Distributed switch: 6.5.0
This version is compatible with VMware ESXi version 6.5 and later. The following new features are available: Port Mirroring Enhancements.
- ☐ Distributed switch: 6.0.0
This version is compatible with VMware ESXi version 6.0 and later. The following new features are available: Network I/O Control version 3, and IGMP/MLD snooping.
- ☐ Distributed switch: 5.5.0
This version is compatible with VMware ESXi version 5.5 and later. The following new features are available: Traffic Filtering and Marking, and enhanced LACP support.
- ☐ Distributed switch: 5.1.0
This version is compatible with VMware ESXi version 5.1 and later. The following new features are available: Management Network Rollback and Recovery, Health Check, Enhanced Port Mirroring, and LACP.
- ☐ Distributed switch: 5.0.0
This version is compatible with VMware ESXi version 5.0 and later. The following new features are available: User-defined network resource pools in Network I/O Control, NetFlow, and Port Mirroring.

At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

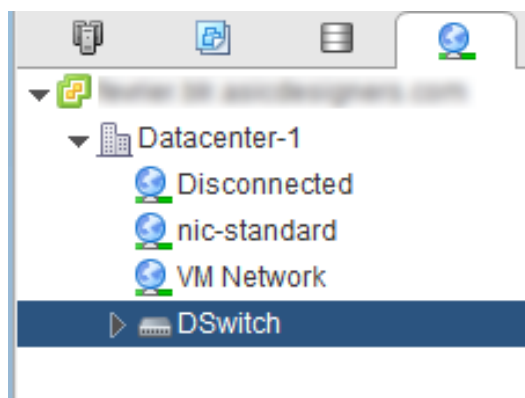
- vi. Select the required number of uplinks for your setup. Please note that the number of uplinks must be equal or less than the number of adapters to be used for each host. Click **Next**.

The screenshot shows the 'New Distributed Switch' wizard in the 'Edit settings' step. The left sidebar has four steps: '1 Name and location', '2 Select version', '3 Edit settings' (selected), and '4 Ready to complete'. The main area is titled 'Edit settings' with the instruction 'Specify number of uplink ports, resource allocation and default port group.' It contains the following fields: 'Number of uplinks' set to 2, 'Network I/O Control' set to 'Enabled', 'Default port group' with an unchecked checkbox for 'Create a default port group', and a 'Port group name' field containing 'DPortGroup'. At the bottom right are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

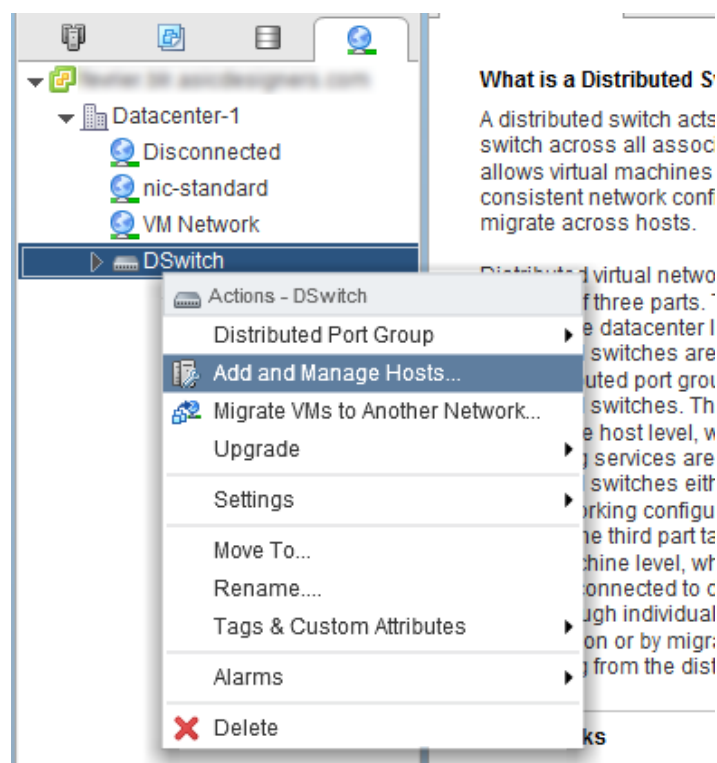
- vii. Review the settings and click **Finish**.

The screenshot shows the 'New Distributed Switch' wizard in the 'Ready to complete' step. The left sidebar has four steps: '1 Name and location', '2 Select version', '3 Edit settings', and '4 Ready to complete' (selected). The main area is titled 'Ready to complete' with the instruction 'Review your settings selections before finishing the wizard.' It displays a summary of the settings: Name: DSwitch2, Version: 6.5.0, Number of uplinks: 2, and Network I/O Control: Enabled. Below this, under 'Suggested next actions', are two items: 'New Distributed Port Group' and 'Add and Manage Hosts'. A note at the bottom states: 'These actions will be available in the Actions menu of the new distributed switch.' At the bottom right are buttons for 'Back', 'Next', 'Finish', and 'Cancel'.

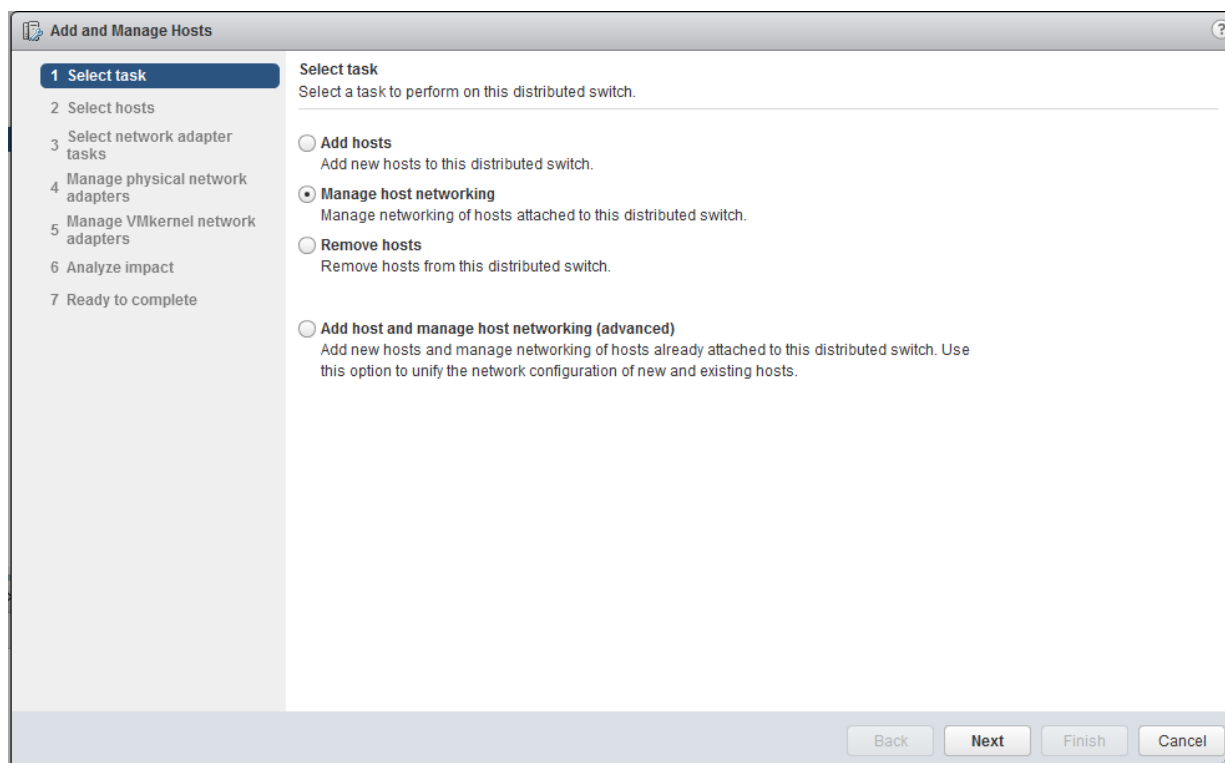
viii. The newly created switch will appear in the inventory.




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



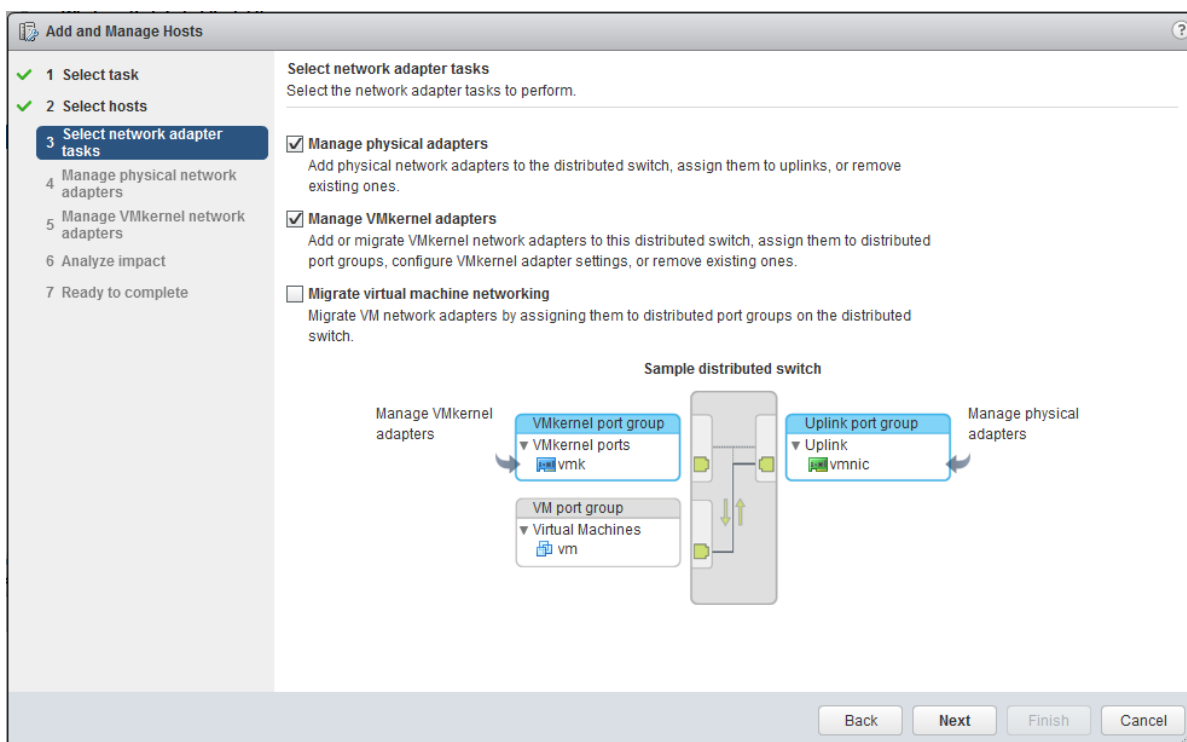
- x. Select **Manage host networking** and click **Next**.



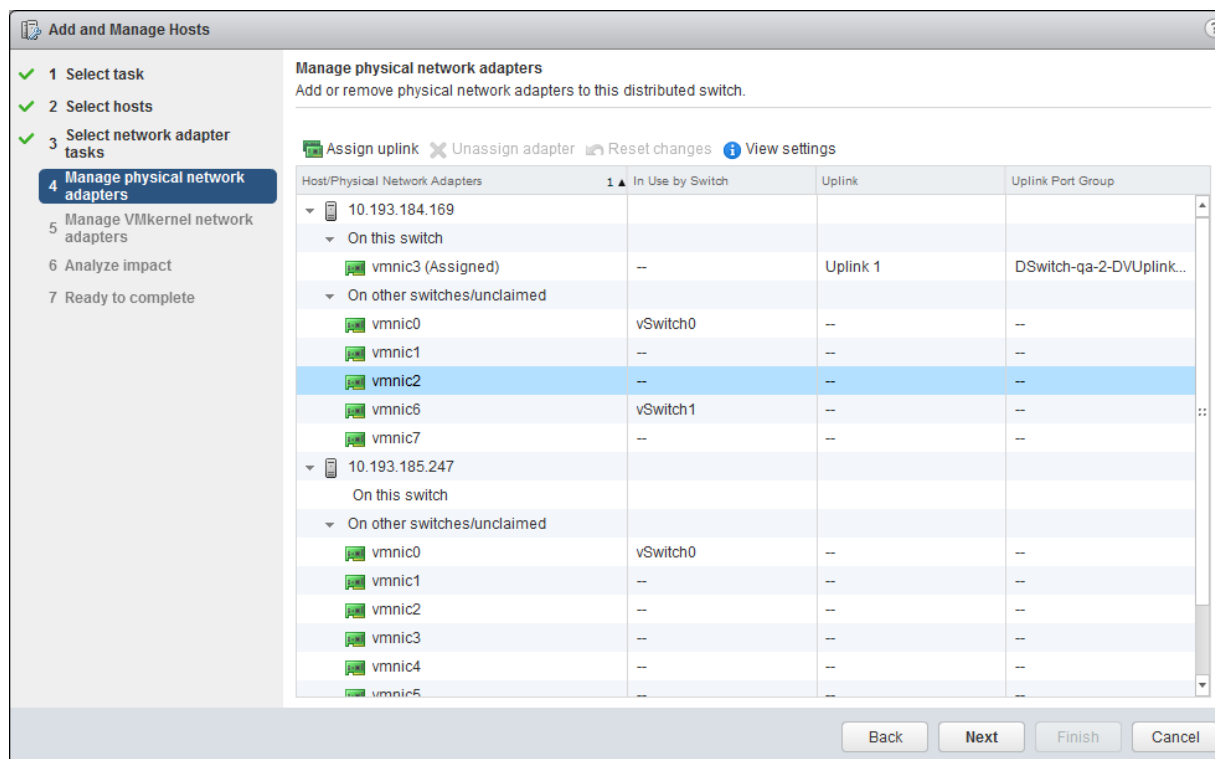
- xi. Click **+ Attached hosts**. Select member hosts to add and click **OK**.



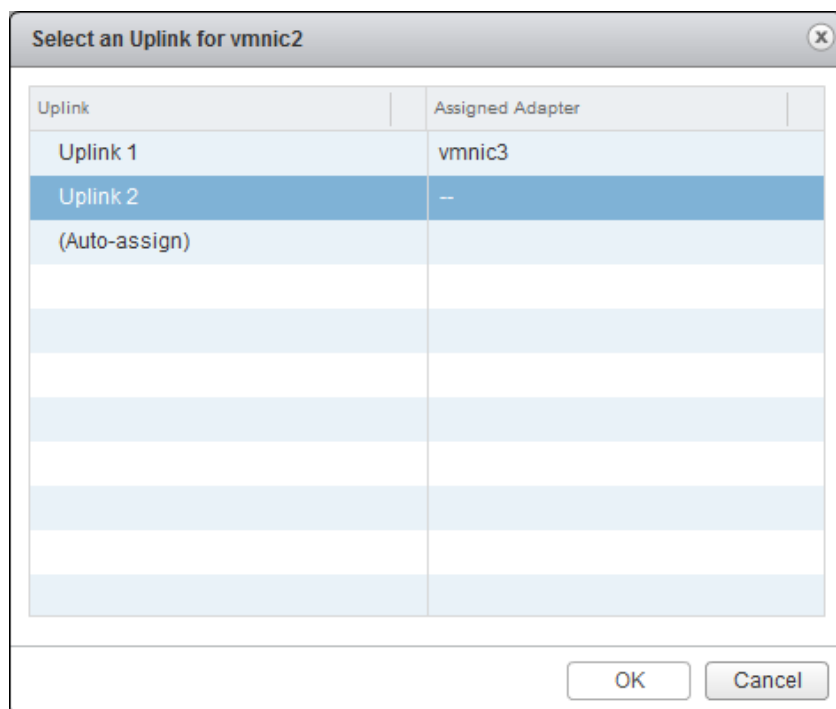
xii. Click **Next** and select the appropriate network adapter tasks to perform. Click **Next**.



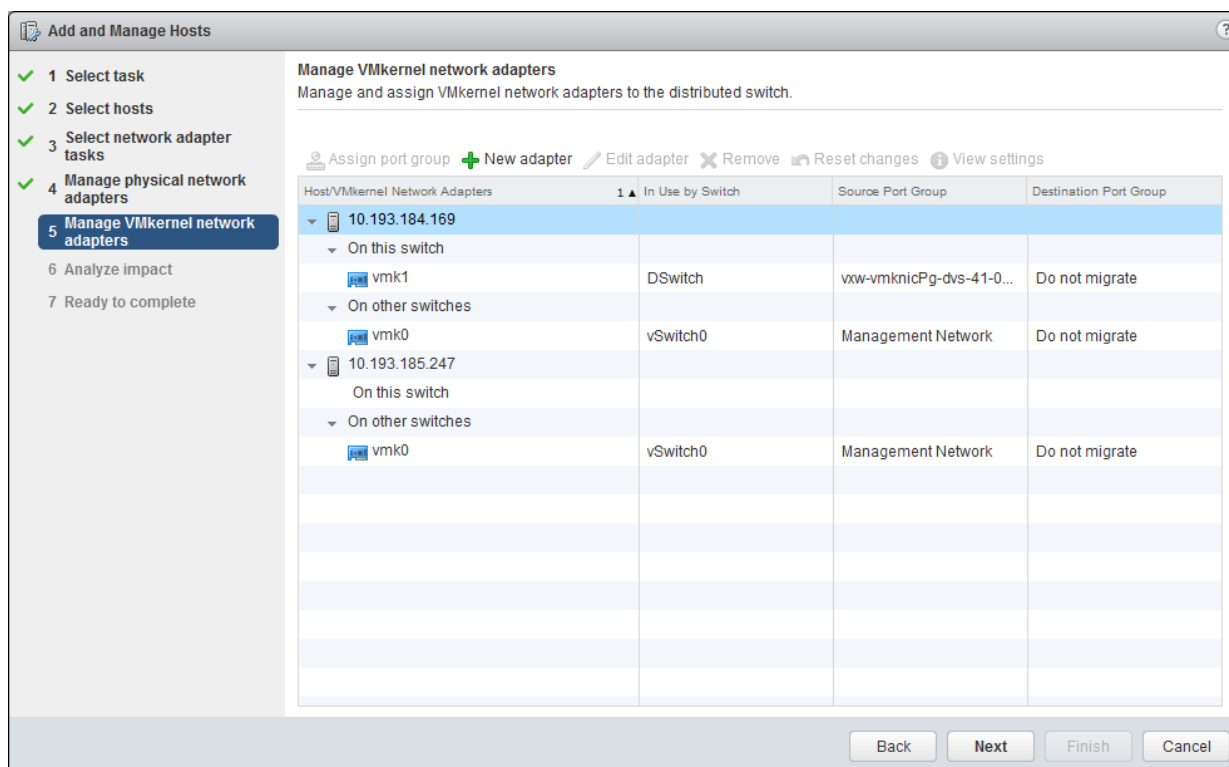
xiii. Here you can add physical network adapters to the switch. Select the adapter and click **Assign uplink**.



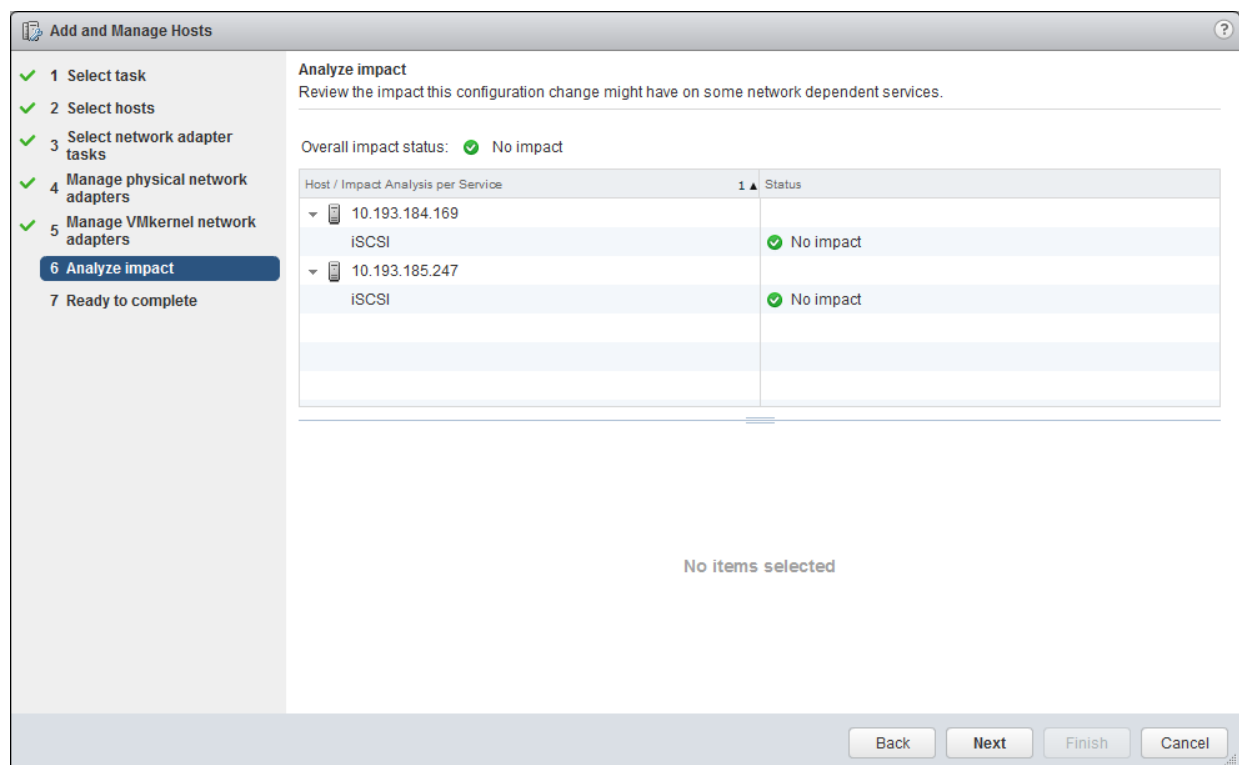
- xiv. Select an uplink for the adapter and click **OK** and then **Next**.



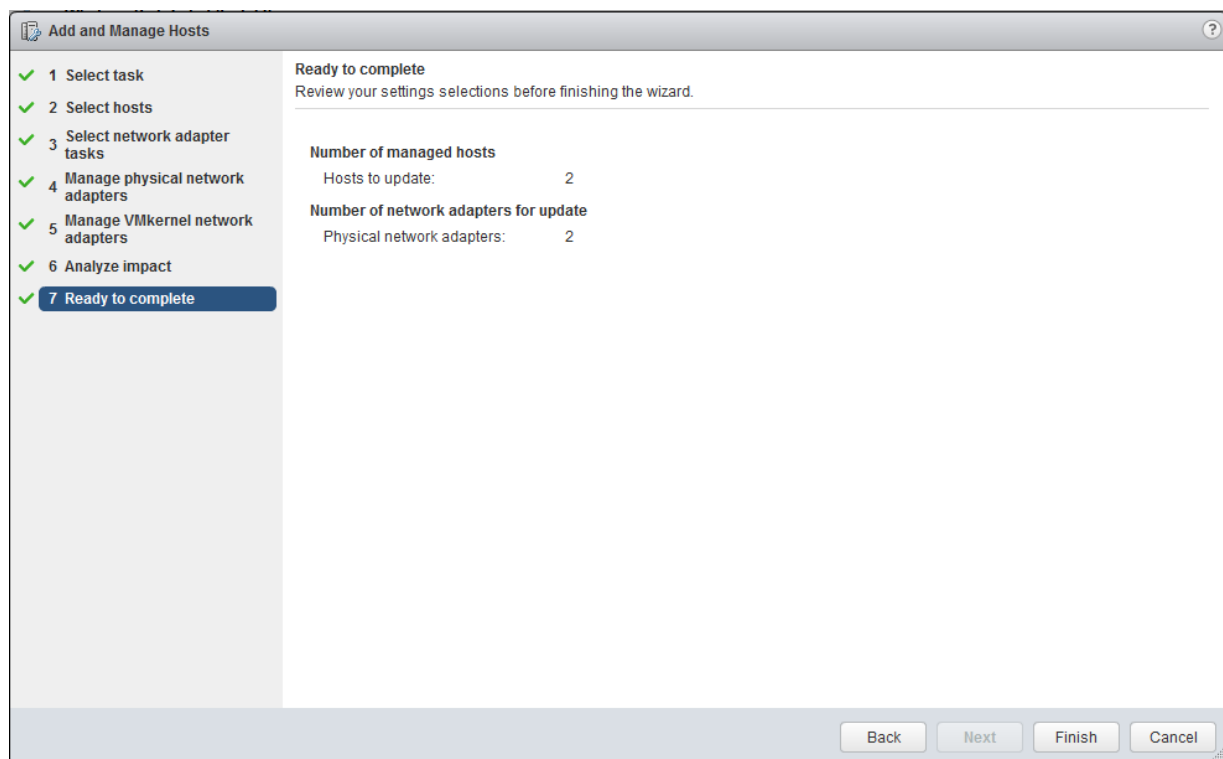
- xv. In this step, you can view the VMkernel adapters automatically added to the hosts. If VXLAN modules are already installed using NSX Manager, you can also view the VXLAN Port group associated to the respective host in distributed switch. Click **Next**.



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

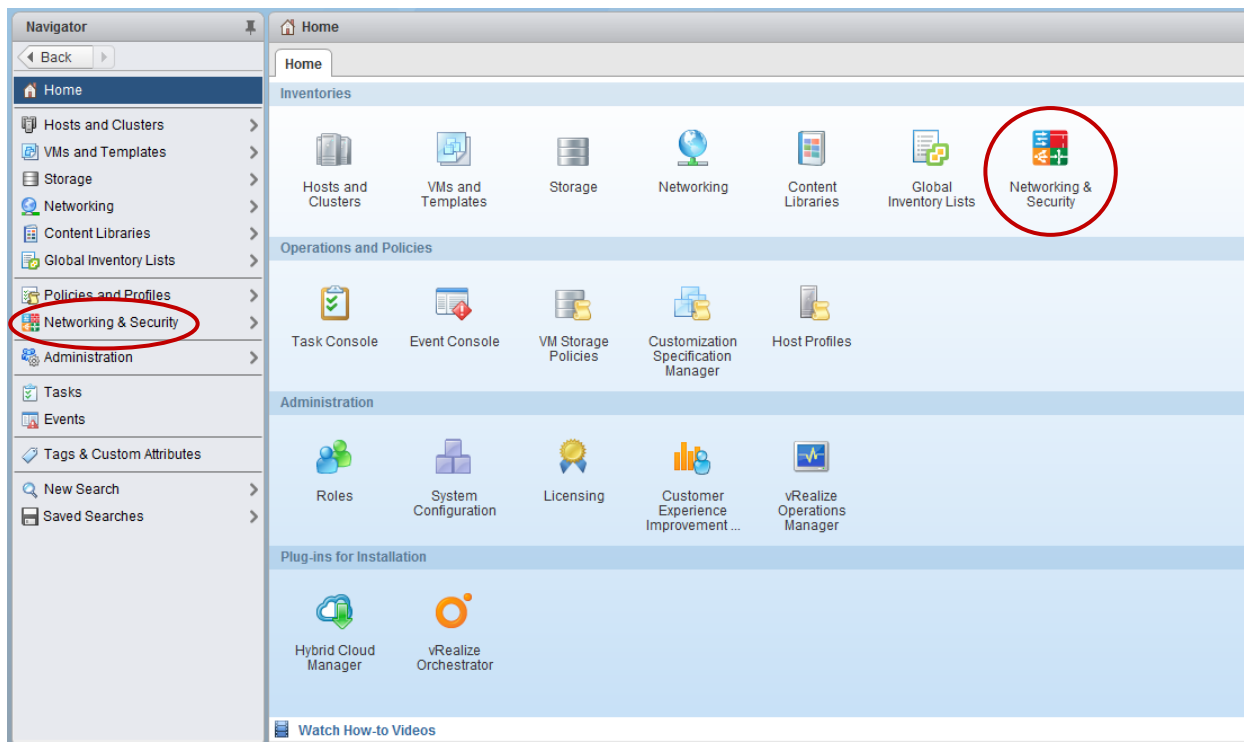


xvii. Review the settings and click **Finish**.

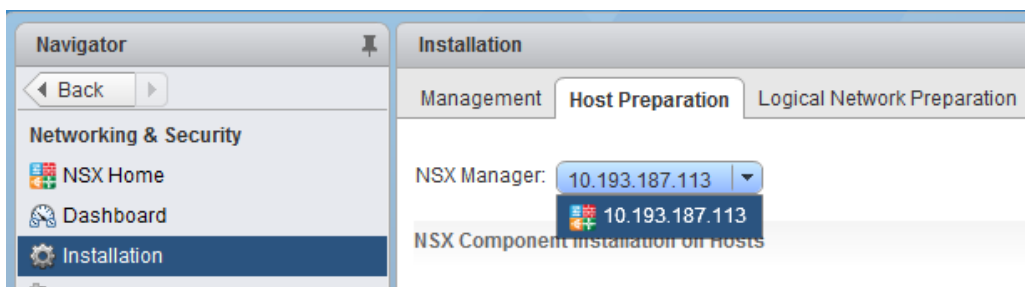


• Setting up VXLAN

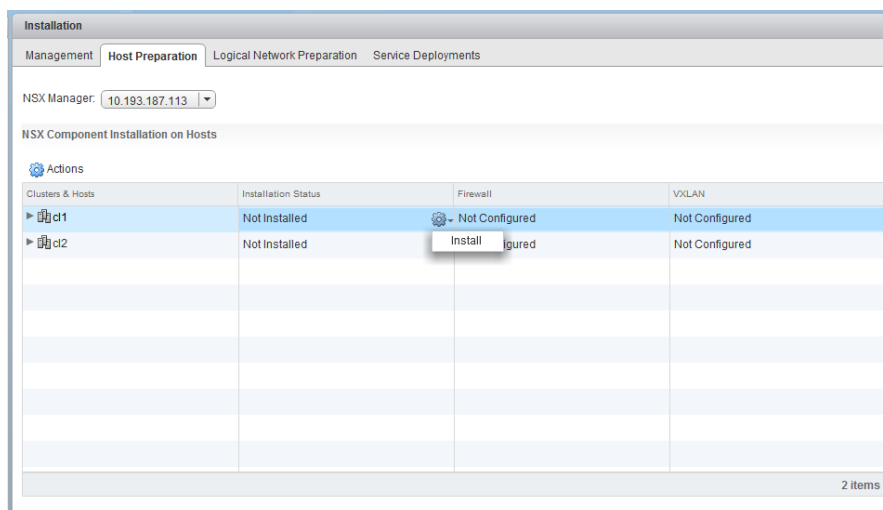
- i. Follow the steps mentioned in VMware's [NSX installation guide](#) to install and configure NSX Manager.
- ii. Make certain that vCenter Server is registered with NSX manager successfully. **Networking & Security** link should appear in both **Home** and **Navigator** panes. Click on the link in 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**.

The screenshot shows a dialog box titled "cl1 - Configure VXLAN Networking". It contains the following fields and options:

- Switch: * DSwitch (dropdown menu)
- VLAN: * 0 (text field)
- MTU: * 1600 (text field)
- VMKNic IP Addressing: * ☐ Use DHCP, ☒ Use IP Pool (radio buttons), vxlan-ip-pool (dropdown menu)
- VMKNic Teaming Policy: * Fail Over (dropdown menu)
- VTEP: * 1 (text field)

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

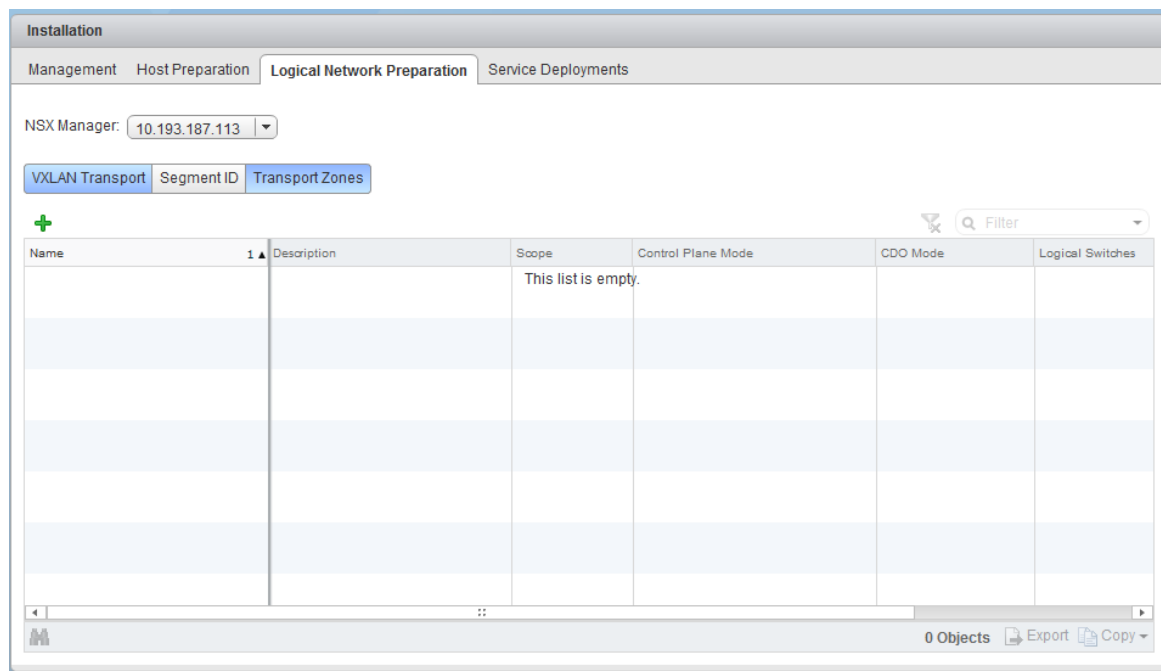
- viii. Click on the **Logical Network Preparation** tab and then **Segment ID**. Set the range for *Segment ID pool* and *Multicast addresses*. Click **OK**.

The screenshot shows a dialog box titled "Edit Segment IDs and Multicast Address Allocation". It contains the following fields and options:

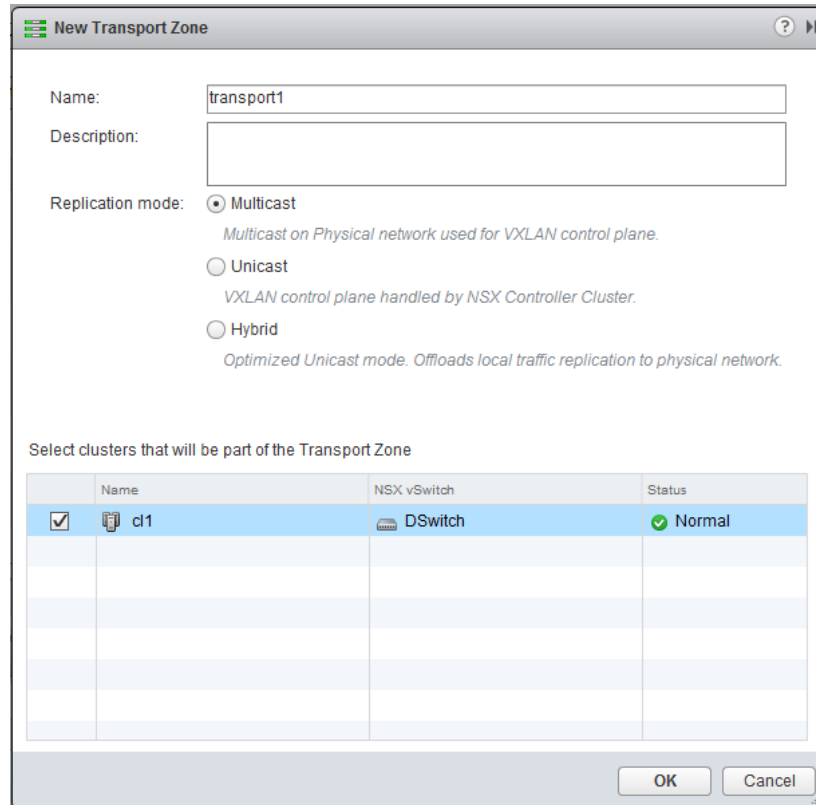
- Provide a Segment ID pool and Multicast range unique to this NSX Manager.
- Segment ID pool: * 5000-5900 (text field)
(In the range of 5000-16777215)
- ☒ Enable Multicast addressing
Multicast addresses are required only for Hybrid and Multicast control plane modes.
- Multicast addresses: * 239.0.0.0-239.255.255.255 (text field)
(Recommended range - 239.0.0.0-239.255.255.255)

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


- 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**.

New Logical Switch

Name: * wire1

Description:

Transport Zone: * transport1 [Change](#) [Remove](#)

Replication mode:

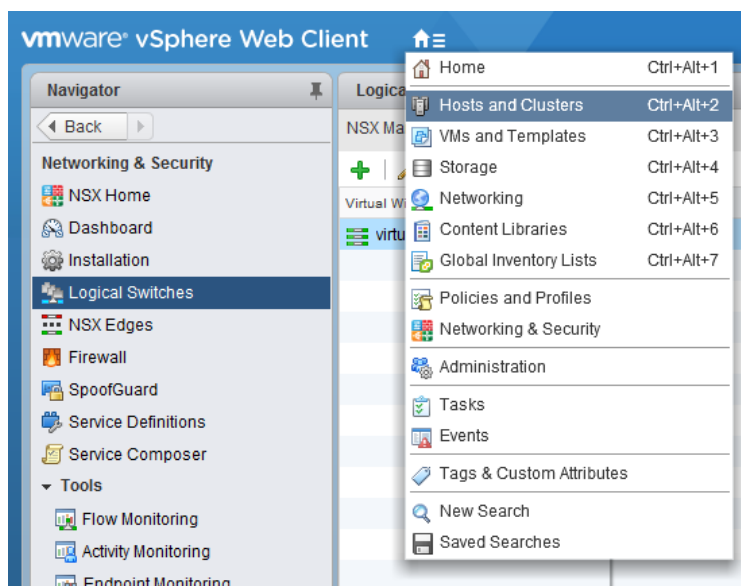
- ☒ Multicast
Multicast on Physical network used for VXLAN control plane.
- ☐ Unicast
VXLAN control plane handled by NSX Controller Cluster.
- ☐ Hybrid
Optimized Unicast mode. Offloads local traffic replication to physical network.

☒ Enable IP Discovery

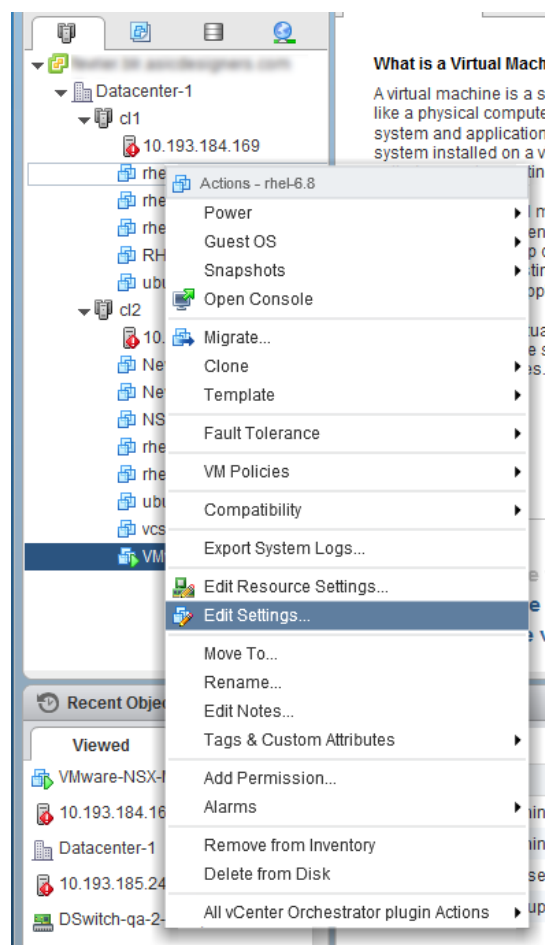
☐ Enable MAC Learning

OK **Cancel**


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



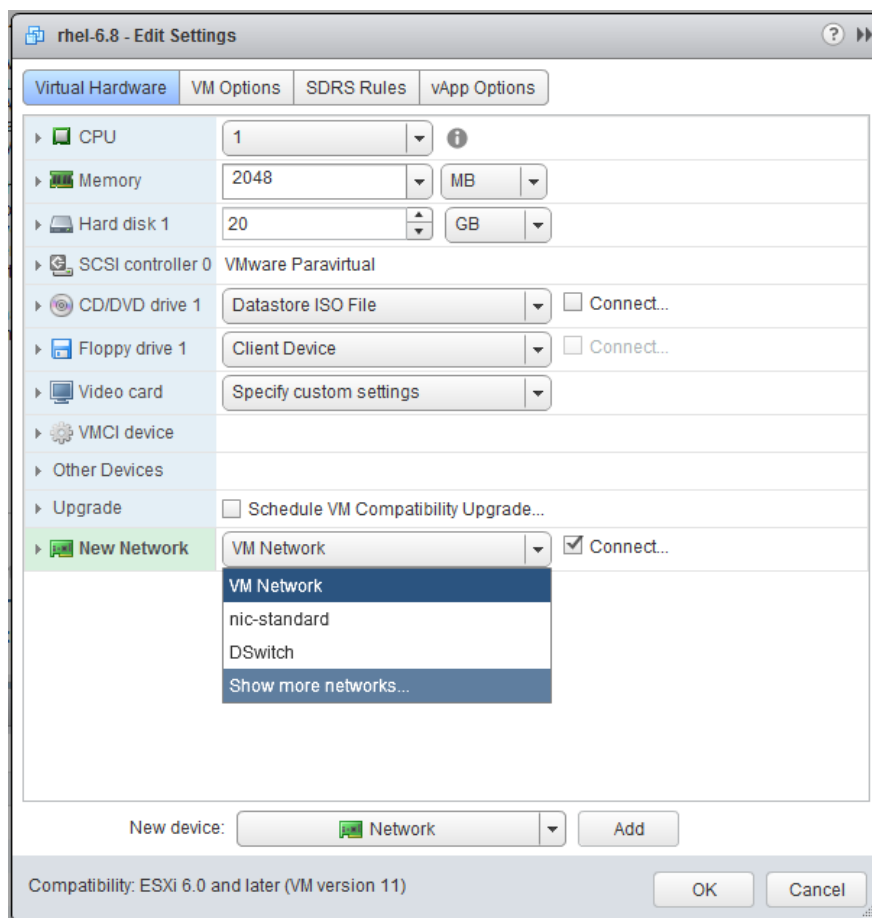
- xiv. Select VM under **Datacenter > Cluster**, right-click and select **Edit Settings**.




- 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. Software/Driver Unloading

Execute the command below to unload the Native Network driver:

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



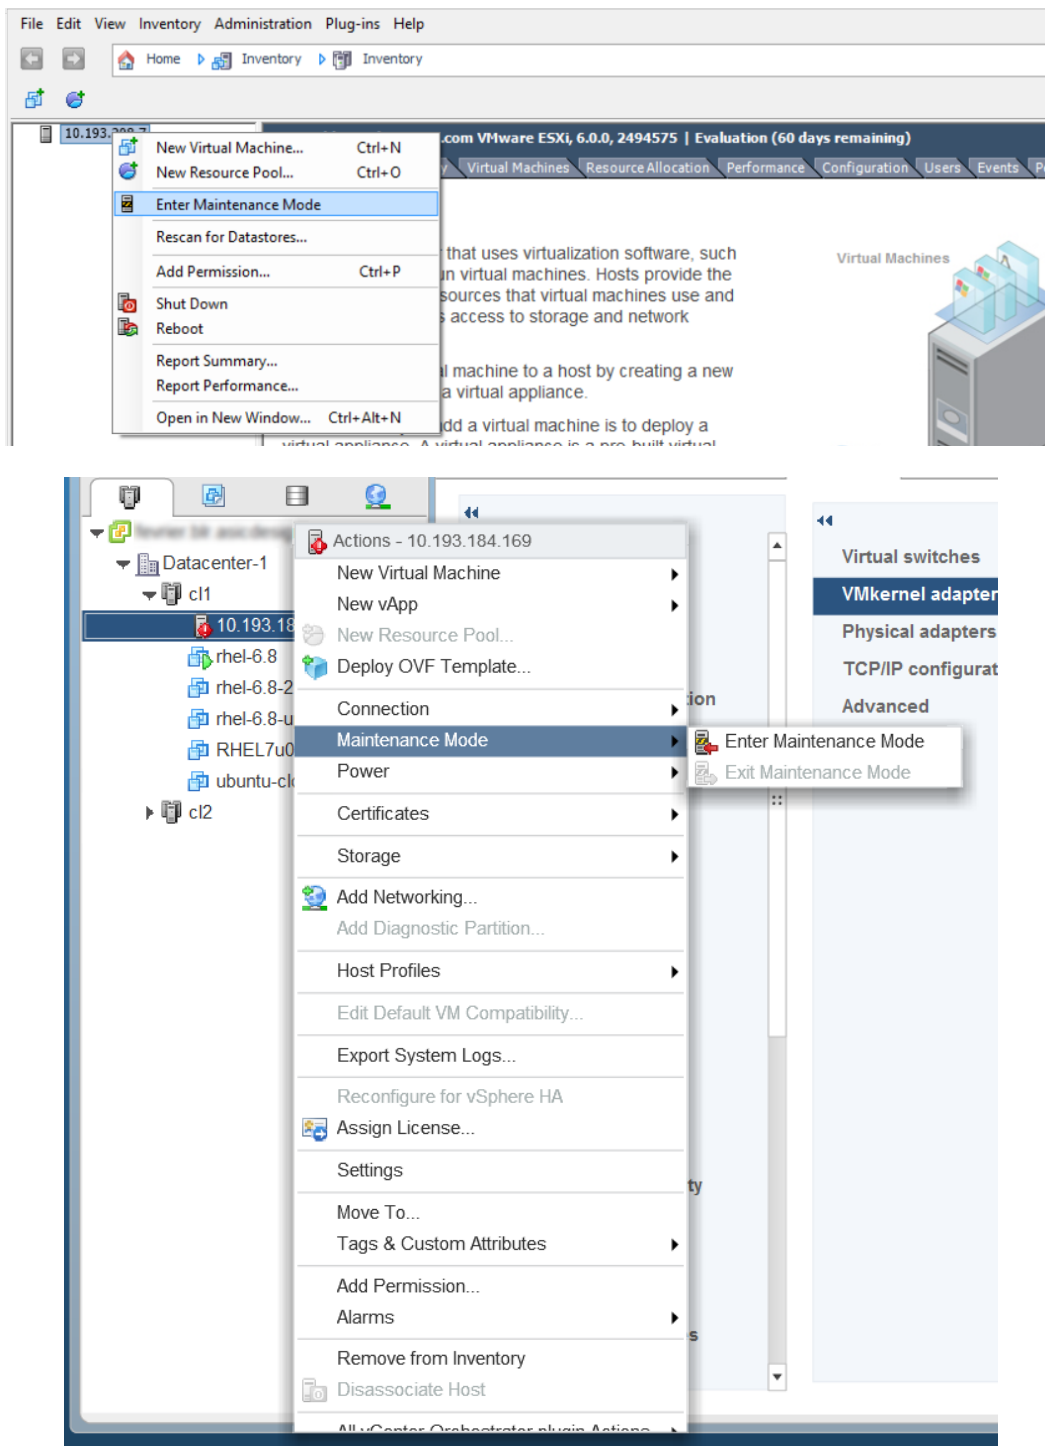
Note

If iSCSI/iSER Offload Initiator Driver (cheiscsi) is loaded, it needs to be unloaded before unloading the native network driver.

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
```

III. iSCSI/iSER Offload Initiator Driver

1. Introduction

The Chelsio Unified Wire series of adapters are Independent Hardware iSCSI adapters. They support iSCSI acceleration and iSCSI Direct Data Placement (DDP) where the hardware handles the expensive byte touching operations, such as CRC computation and verification, and direct DMA to the final host memory destination:

- **iSCSI PDU digest generation and verification**

On transmit -side, Chelsio hardware computes and inserts the Header and Data digest into the PDUs. On receive-side, Chelsio hardware computes and verifies the Header and Data digest of the PDUs.

- **Direct Data Placement (DDP)**

Chelsio hardware can directly place the iSCSI Data-In or Data-Out PDU's payload into pre-posted destination host-memory buffers based on the Initiator Task Tag (ITT) in Data-In or Target Task Tag (TTT) in Data-Out PDUs.

- **PDU Transmit and Recovery**

On transmit-side, Chelsio hardware accepts the complete PDU (header + data) from the host driver, computes and inserts the digests, decomposes the PDU into multiple TCP segments if necessary, and transmit all the TCP segments onto the wire. It handles TCP retransmission if needed.

On receive-side, Chelsio hardware recovers the iSCSI PDU by reassembling TCP segments, separating the header and data, calculating and verifying the digests, then forwarding the header to the host. The payload data, if possible, will be directly placed into the pre-posted host DDP buffer. Otherwise, the data will be sent to the host too.

The iSCSI Extensions for RDMA (iSER) protocol is a translation layer for operating iSCSI over RDMA transports, such as iWARP RDMA. Chelsio Unified Wire adapters supporting iWARP provide the higher bandwidth and lower latency required for block storage transfers. iSER is stable and provides benefits of the iSCSI protocol like security and high availability.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio iSCSI/iSER Offload Initiator driver:

- T62100-CR
- T62100-LP-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory Free; 256 IPv4/128 IPv6 offload connections supported*)

- T6425-CR
- T580-LP-CR
- T580-CR
- T540-CR
- T540-LP-CR
- T520-LL-CR
- T520-CR
- T520-BT
- T540-BT

1.2. Software Requirements

1.2.1. ESXi Requirements

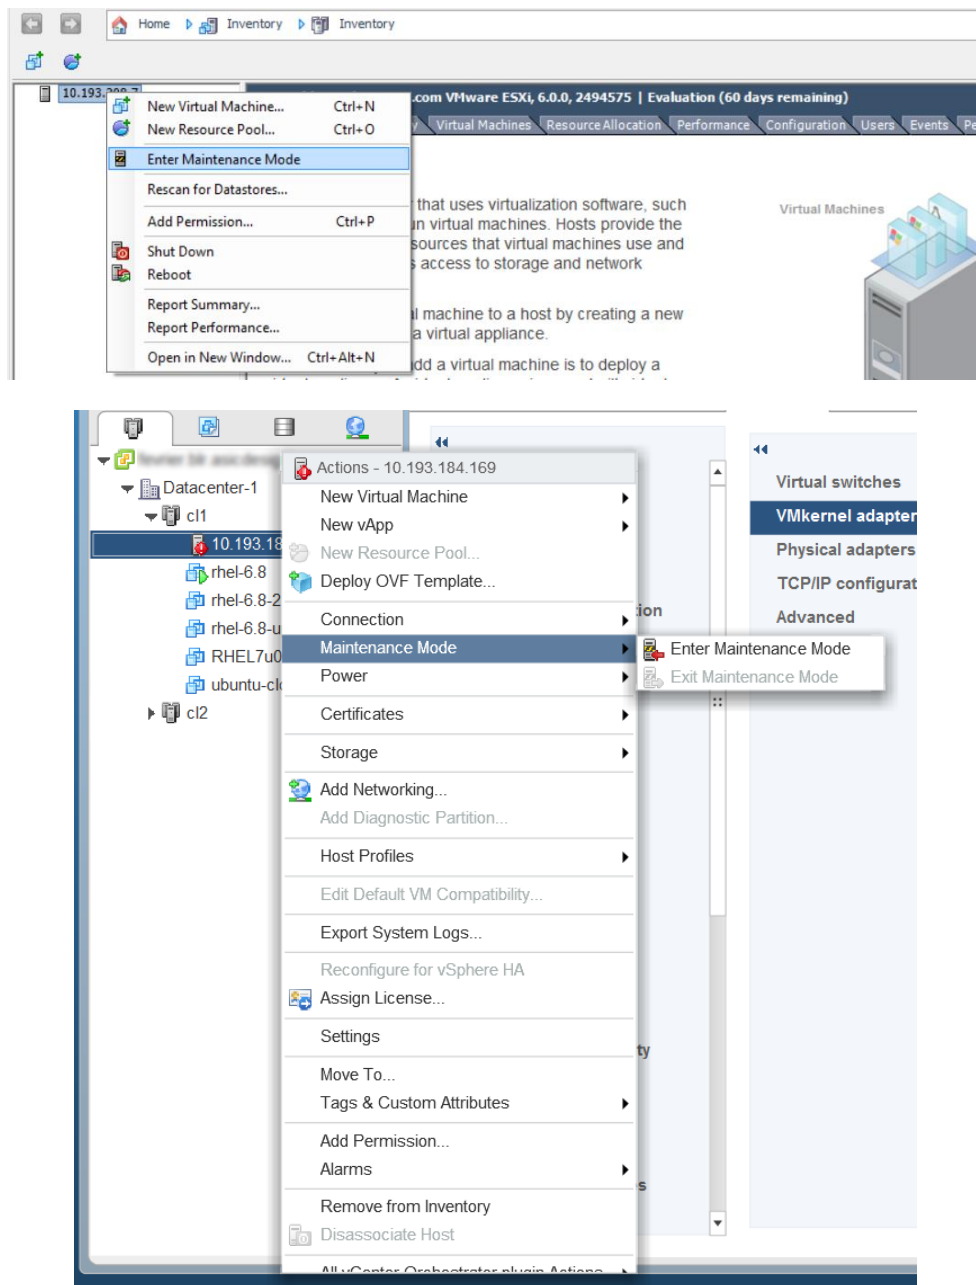
The Chelsio iSCSI/iSER Offload Initiator driver has been developed to run on 64-bit ESXi platforms. Currently the driver is available for the following version(s):

- ESXi 6.7

2. Software/Driver Installation

Native Network driver (cxl) is required for iSCSI/iSER Offload Initiator driver to work. Hence, ensure that the driver is installed and loaded as described in the [Native Network Driver](#) chapter before proceeding.

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




- ii. Install the iSCSI/iSER driver components *cheiscsi* and *chelsio-esx-libcheiscsi_ima.so*:

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

- iii. Remove the *chelsio* folder from the productLocker (if exists).

```
[root@host:~] cd /productLocker/  
[root@host:~] rm -rf chelsio/
```

 **Note** *All the persistent values like IP address, target lists etc. will need to be reconfigured with the new driver.*

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

3. Software/Driver Loading

After rebooting the ESXi Host, the driver will load automatically. However, it is possible to manually load the driver.

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

By default, the driver will load with iSER enabled. To disable iSER and use iSCSI,

```
[root@host:~] vmkload_mod cheiscsi iser=0
```

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

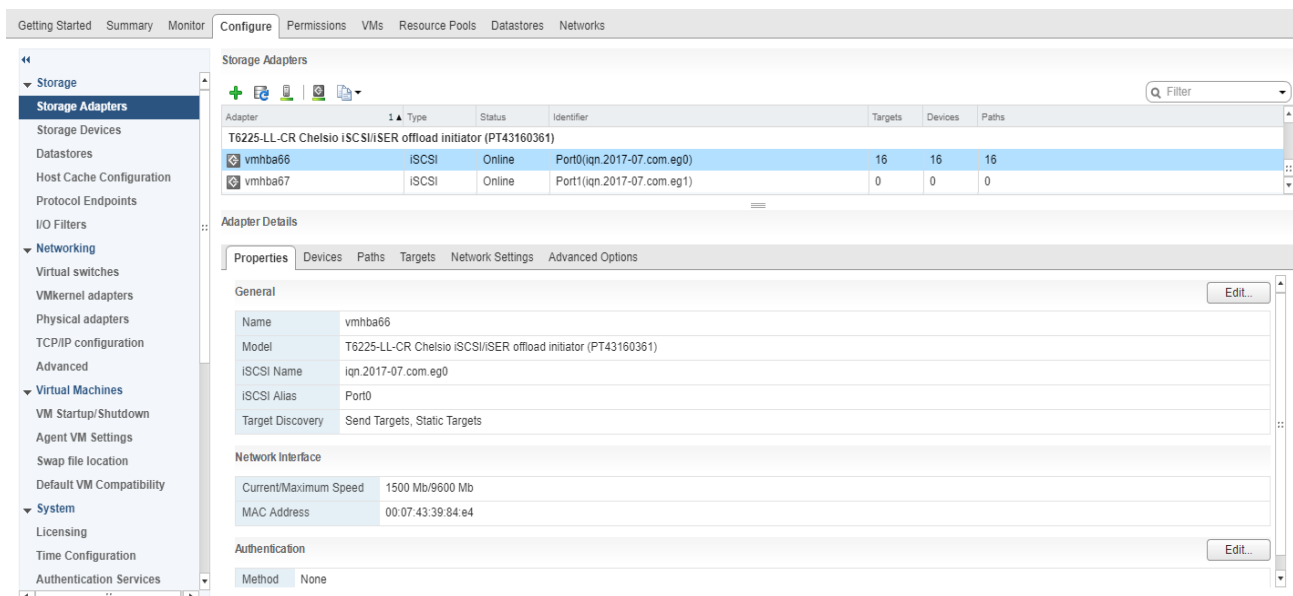
```
[root@host:~] kill -SIGHUP $(cat /var/run/vmware/vmkdevmgr.pid)
```

4. Software/Driver Configuration and Fine-tuning

The following sections describe the method to configure Chelsio iSCSI/iSER Offload Initiator and connect to target.

4.1. Configuring Initiator

- i. Log in to vCenter Server through vSphere Web Client using a web browser.
- ii. If you have already created and configured the host intended to be used as initiator, skip to step (iii)
 - a. Under **Hosts and Clusters**, right-click and click **New Datacenter...** Provide a name and Click **OK**.
 - b. Right-click on the newly created datacenter and click **Add Host...** Follow onscreen instructions and provide information to add the host. Click **Finish**.
- iii. Select the host and under the **Configure** tab, select **Storage Adapters**. This will display the list of available Chelsio iSCSI/iSER adapters.

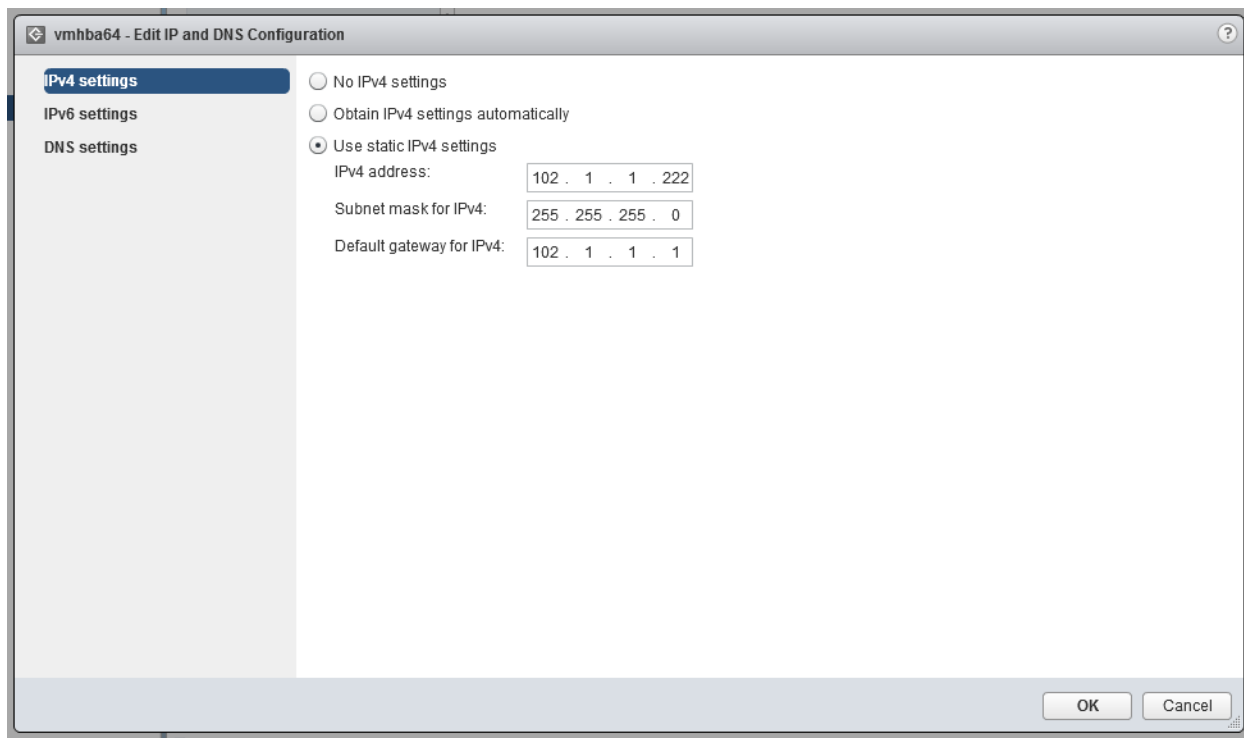


You can also view the list in CLI using:

```
[root@host:~] esxcli iscsi adapter list
```

```
[root@rattletrap:~] esxcli iscsi adapter list
Adapter  Driver  State  UID              Description
-----
vmhba66  cheiscsi online iscsi.vmhba66    T6225-LL-CR Chelsio iSCSI/iSER offload initiator (PT43160361)
vmhba67  cheiscsi online iscsi.vmhba67    T6225-LL-CR Chelsio iSCSI/iSER offload initiator (PT43160361)
```

- iv. In the **Adapter Details** section, click **Network Settings** tab and then **Edit**.
- v. Configure IP for the adapter and click **OK**. You can configure either IPv4 or IPv6 or both.

**Adapter Details**

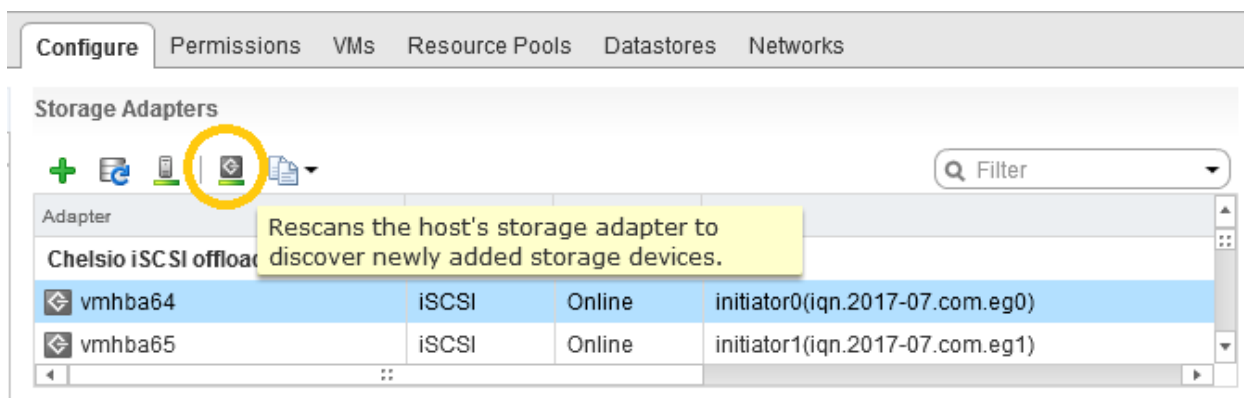
Properties Devices Paths Targets **Network Settings** Advanced Options

IP Address and DNS Configuration

Edit...

IPv4 address	102.1.1.222 (static)
Subnet mask for IPv4	255.255.255.0
Default gateway for IPv4	102.1.1.1
IPv6 settings	Not enabled
Preferred DNS server	Not supported
Alternate DNS server	Not supported

- vi. For changes to take effect, rescan the iSCSI/iSER adapter using **Rescan**.



4.2. Connecting to Target

There are two methods to discover and connect to targets:

- **Dynamic Discovery:** Discovers all the available targets for a given target server.
- **Static Discovery:** Discovers a specific target by manually entering target information.

4.2.1. Dynamic Discovery

- **Adding Target Server**

- Select the iSCSI/iSER adapter to connect to the target and under the **Targets** tab, select **Dynamic Discovery**.

The screenshot shows the vSphere Client interface for configuring a Chelsio iSCSI offload initiator. The top navigation bar includes tabs for 'Configure', 'Permissions', 'VMs', 'Resource Pools', 'Datastores', and 'Networks'. The 'Configure' tab is active, and the 'Storage Adapters' section is expanded. A table lists the available adapters:

Adapter	Type	Status	Identifier
Chelsio iSCSI offload initiator			
vmhba64	iSCSI	Online	initiator0(iqn.2017-07.com.eg0)
vmhba65	iSCSI	Online	initiator1(iqn.2017-07.com.eg1)

Below the table, the 'Adapter Details' section is visible. The 'Targets' sub-tab is selected, and the 'Dynamic Discovery' button is highlighted. The 'iSCSI server' list is empty, displaying the message 'This list is empty.'

- ii. Click **Add** and provide the target server IP. Click **OK**.

vmhba64 - Add Send Target Server

iSCSI Server: 102.1.1.18

Port: 3260

Authentication Settings

☒ Inherit settings from parent

OK Cancel

- iii. For changes to take effect, rescan the iSCSI/iSER adapter using **Rescan**.
- iv. All the available LUNs discovered on the given target server will be displayed in the **Devices** tab. You can perform various tasks like detach, rename, erase partitions, etc.

Adapter Details

Properties	Devices	Paths	Targets	Network Settings	Advanced Options
<div> All Actions ▾ ▾ </div> <div> <input type="text" value="Filter"/> </div>					
Name	LUN	Type	Capacity	Operational...	Hardware Acceleration
IET iSCSI Disk (t10.945445000...	0	disk	104.91 GB	Attached	Unknown
IET iSCSI Disk (t10.945445000...	1	disk	40.06 GB	Attached	Unknown

The **Paths** tab displays LUN status and allows you to enable or disable them.

Adapter Details

Properties

Devices

Paths

Targets

Network Settings

Advanced Options

Enable

Disable

Runtime Name	Target	LUN	Status
vmhba64:C0:T0:L0		0	◆ Active (I/O)
vmhba64:C0:T0:L1		1	◆ Active (I/O)

These LUNs can now be attached to VMs or can be used to store VMs.

- **Removing Target Server**

- i. Under **Targets** tab, select **Dynamic Discovery**.
- ii. Select the target server, click **Remove** and then **Yes** to confirm.
- iii. The iSCSI/iSER server will also be listed as static target under **Static Discovery**, hence needs to be removed from here as well. Select the server, click **Remove** and then **Yes** to confirm.

Adapter Details

Properties Devices Paths **Targets** Network Settings Advanced Options

Dynamic Discovery **Static Discovery**

Add...

Remove

Authentication...

Advanced...

iSCSI server	Target Name
102.1.1.18:3260	iqn.2001-04.com.example:storage.disk2.sys1.xyz

- iv. For changes to take effect, rescan the iSCSI/iSER adapter using **Rescan**.

- **Static Discovery**

- **Adding Target Server**

- i. Select the iSCSI/iSER interface to connect to the target and under the **Targets** tab, select **Static Discovery**.

ConfigurePermissionsVMsResource PoolsDatastoresNetworks

Storage Adapters

Filter

Adapter	Type	Status	Identifier
Chelsio iSCSI offload initiator			
vmhba64	iSCSI	Online	initiator0(iqn.2017-07.com.eg0)
vmhba65	iSCSI	Online	initiator1(iqn.2017-07.com.eg1)

Adapter Details

PropertiesDevicesPathsTargetsNetwork SettingsAdvanced Options

Dynamic DiscoveryStatic Discovery

Add...RemoveAuthentication...Advanced...

iSCSI server	Target Name

- ii. Click **Add** and provide the target server IP and target IQN. Click **OK**.

vmhba64 - Add Static Target Server

iSCSI Server: 102.1.1.18

Port: 3260

iSCSI Target Name: iqn.2001-04.com.example:storage.disk2.sys1.xyz

Authentication Settings

☒ Inherit settings from parent

OK Cancel

- iii. For changes to take effect, rescan the iSCSI/iSER adapter using **Rescan**.
- iv. All the available LUNs discovered on the given target server will be displayed in the **Devices** tab. You can perform various tasks like detach, rename, erase partitions, etc.

Adapter Details

PropertiesDevicesPathsTargetsNetwork SettingsAdvanced Options

All Actions

Filter

Name	LUN	Type	Capacity	Operational...	Hardware Accelera
IET iSCSI Disk (t10.945445000...	0	disk	104.91 GB	Attached	Unknown
IET iSCSI Disk (t10.945445000...	1	disk	40.06 GB	Attached	Unknown

The **Paths** tab displays LUN status and allows you to enable or disable them.

Adapter Details

PropertiesDevicesPathsTargetsNetwork SettingsAdvanced Options

EnableDisable

Runtime Name	Target	LUN	Status
vmhba64:C0:T0:L0		0	◆ Active (I/O)
vmhba64:C0:T0:L1		1	◆ Active (I/O)

These LUNs can now be attached to VMs or can be used to store VMs.

- **Removing Target Server**
 - i. Under **Targets** tab, select **Static Discovery**.
 - ii. Select the target server, click **Remove** and then **Yes** to confirm.

Adapter Details

Properties Devices Paths **Targets** Network Settings Advanced Options

Dynamic Discovery **Static Discovery**

Add...

Remove

Authentication...

Advanced...

iSCSI server	Target Name
102.1.1.18:3260	iqn.2001-04.com.example:storage.disk2.sys1.xyz

- iii. For changes to take effect, rescan the iSCSI/iSER adapter using **Rescan**.

4.3. Configurable Options

The option to edit general initiator properties like alias and name is available under the **Properties** tab, whereas advanced parameters like Digest, MTU, etc., can be changed in the **Advanced Options** tab.

To change target server's advanced parameters, use the **Advanced** option under the **Targets** tab.

5. Software/Driver Unloading

Logout all the existing iSCSI/iSER sessions. Execute the command below to unload the iSCSI/iSER Offload Initiator driver:

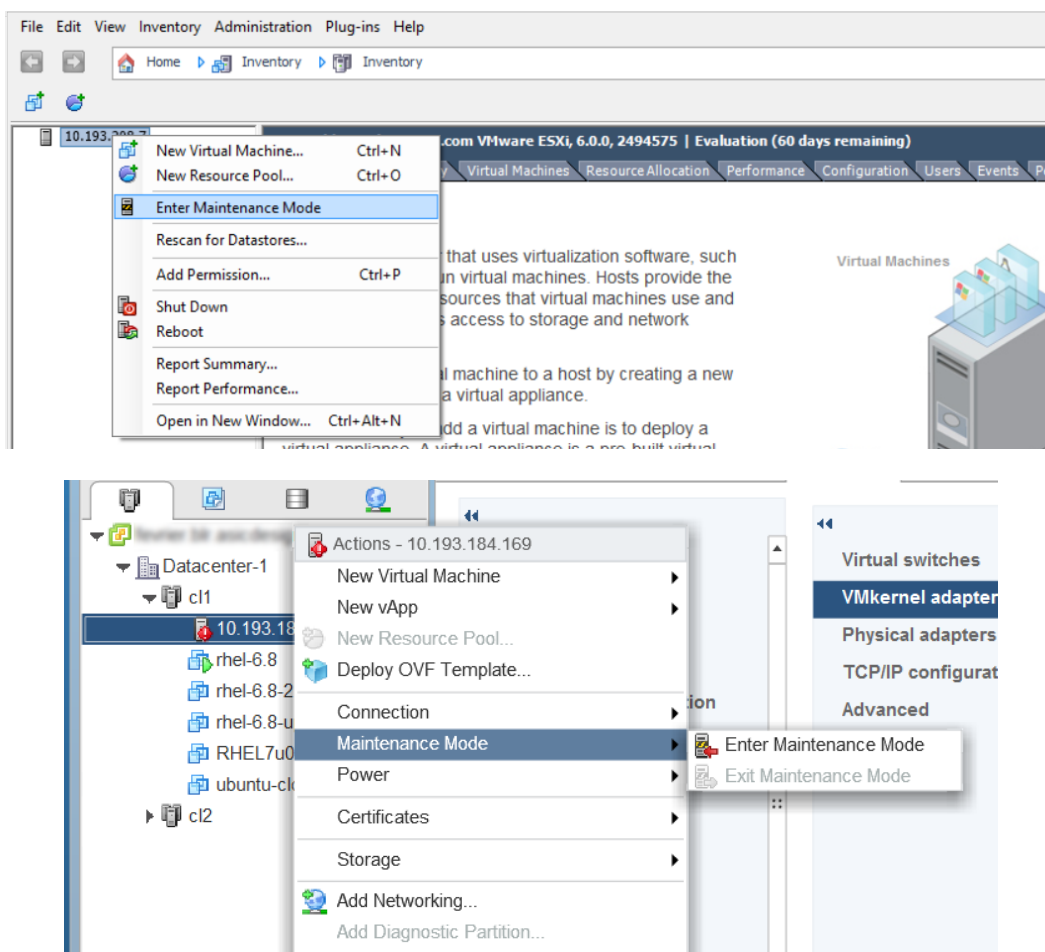
```
[root@host:~] vmkload_mod -u cheiscsi
```

6. Software/Driver Uninstallation

Note Before proceeding, please ensure that no iSCSI/iSER session or connection is active and running.

Follow the steps mentioned below to uninstall the iSCSI/iSER Offload Initiator 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 cheiscsi
[root@host:~] esxcli software vib remove -n chelsio-esx-libcheiscsi_ima.so
```

- iii. Reboot the host:

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


IV. Appendix

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