



Chelsio Unified Wire for VMware ESXi 8.0

Installation and User 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)

735 N Pastoria 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

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 © 2025. 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	02/26/2023

TABLE OF CONTENTS

I. CHELSIO UNIFIED WIRE	6
1. Introduction	7
1.1. Features	7
1.2. Hardware Requirements	7
1.3. Software Requirements	8
1.4. Package Contents	8
2. Hardware Installation	9
3. Software/Driver Installation	11
4. Software/Driver Uninstallation	12
5. Software/Driver Update	13
II. NATIVE NETWORK DRIVER WITH SR-IOV SUPPORT	14
1. Introduction	15
1.1. Hardware Requirements	15
1.2. Software Requirements	15
2. Software/Driver Loading	17
3. Software/Driver Configuration and Fine-tuning	18
3.1. Multiple Adapters	18
3.2. cxgbtool	18
3.3. Adapter Configuration	18
3.4. Firmware Update	19
3.5. Connecting a Virtual Machine	19
3.6. Virtual Functions (SR-IOV)	19
4. Software/Driver Unloading	26
III. ISCSI OFFLOAD INITIATOR DRIVER	27
1. Introduction	28
1.1. Hardware Requirements	28
2. Software/Driver Loading	29
3. Software/Driver Configuration and Fine-tuning	30
3.1. Configuring Initiator	30
3.2. Connecting to Target	32
3.3. Configurable Options	35
4. Software/Driver Unloading	37
IV. ISER OFFLOAD INITIATOR DRIVER	38
1. Introduction	39
1.1. Hardware Requirements	39
2. Software/Driver Loading	40
3. Software/Driver Configuration and Fine-tuning	41
3.1. Configuring Initiator	41
3.2. Connecting to Target	45
3.3. Configurable Options	48

4. Software/Driver Unloading	50
V. NVME-OF OFFLOAD INITIATOR DRIVER	51
1. Introduction	52
1.1. Hardware Requirements	52
2. Software/Driver Loading	53
3. Software/Driver Configuration and Fine-tuning	54
3.1. Connecting to NVMe target	54
3.2. Disconnecting from NVMe target	56
4. Software/Driver Unloading	57
VI. APPENDIX	58
1. Troubleshooting	59
2. Chelsio End-User License Agreement (EULA)	60

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 Offload Initiator Driver**
- **iSER Offload Initiator Driver**
- **NVMe-oF Offload Initiator Driver**

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

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

1.2. Hardware Requirements

The Chelsio Unified Wire software supports the Chelsio Terminator series of Unified Wire adapters. To know more about the list of adapters supported by each driver, 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.

- ESXi 8.0

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, NVMe-oF Offload Initiator driver VIB file.
- **cheiwarp-*.vib**: iSER Offload Initiator driver VIB file.

2. Hardware Installation

Follow these steps to install the 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 MSIX 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.

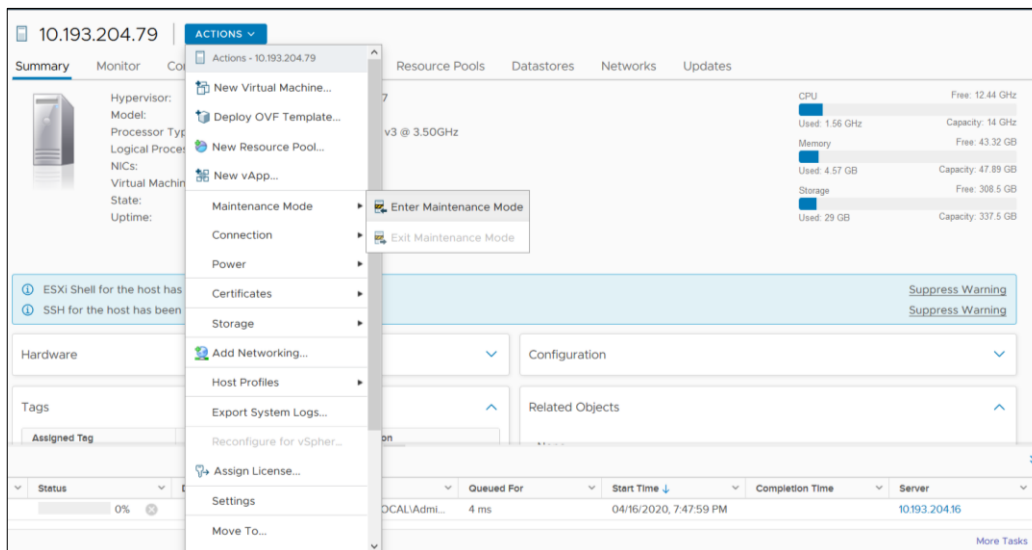


Note

Network device names for Chelsio's physical ports are assigned using the following convention: the port farthest from the motherboard appears 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 appears as the first network interface.

3. Software/Driver Installation

- i. Download the driver package from [Chelsio Download Center](#).
- ii. Put the host in maintenance mode using the vSphere (desktop or web) Client.



- iii. Install the drivers.

```
[root@host:~] cp *.zip /productLocker/
[root@host:~] cp *.zip /var/log/vmware/
[root@host:~] esxcli software component apply --depot=/productLocker/CHL-
esx-8.0.0-Chelsio-Drivers-5.3.0.33-10EM.800.1.0.20613240.zip --no-sig-check
```

```
[root@KGF2:~] esxcli software component apply --depot=/productLocker/CHL-esx-8.0.0-Chelsio-Drivers-5.3.0.33-10EM.800.1.0.20613240.zip --no-sig-check
Installation Result
  Components Installed: Chelsio-Drivers_5.3.0.33-10EM.800.1.0.20613240
  Components Removed:
  Components Skipped:
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
```

- iv. Once the installation or update is successfully completed, exit from maintenance mode and reboot the host.
- v. Verify that the drivers are installed successfully.

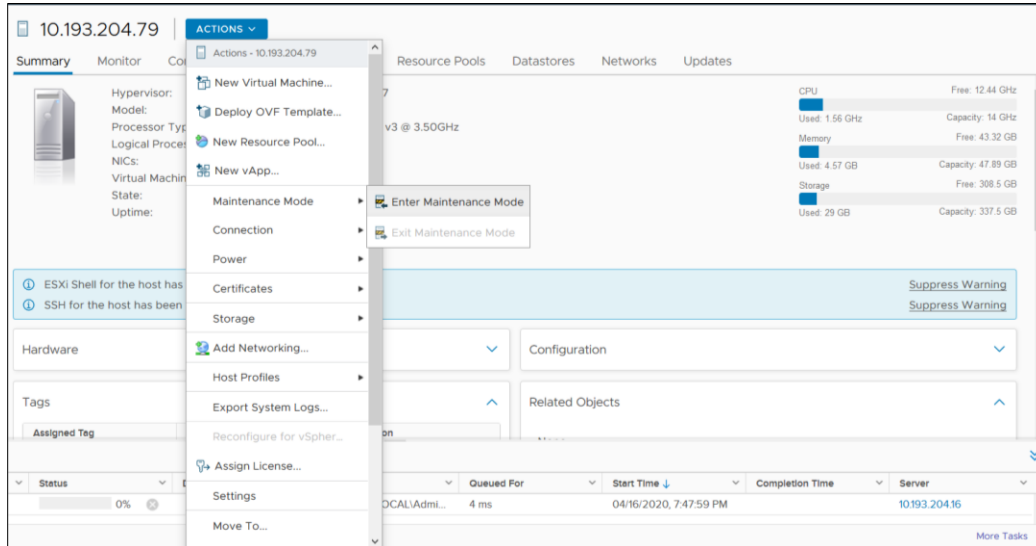
```
[root@host:~] esxcli software component vib list --component=Chelsio-Drivers
```

```
[root@KGF2:~] esxcli software component vib list --component=Chelsio-Drivers
Name          Version                               Vendor      Acceptance Level  Install Date  Platforms
-----
cheiscsi      5.3.0.33-10EM.800.1.0.20613240      CHL        VMwareCertified   2025-02-25   host
cheiwarp      5.3.0.33-10EM.800.1.0.20613240      CHL        VMwareCertified   2025-02-25   host
cxl           5.3.0.33-10EM.800.1.0.20613240      CHL        VMwareCertified   2025-02-25   host
```

4. Software/Driver Uninstallation

Note Before proceeding with the uninstallation, ensure that no iSCSI, iSER or NVMe-oF session or connection is active and running.

- i. Use the **vSphere Client (desktop or web)** to place the host in maintenance mode:



- ii. Uninstall the drivers.

```
[root@host:~] esxcli software component remove --component=Chelsio-Drivers
```

```
[root@KGF2:~] esxcli software component remove --component=Chelsio-Drivers
Removal Result
Components Installed:
Components Removed: Chelsio-Drivers_5.3.0.33-10EM.800.1.0.20613240
Components Skipped:
Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
Reboot Required: true
```

- iii. Reboot the host.

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

5. Software/Driver Update

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

For regular updates on various software/drivers, visit [Chelsio Download Center](#). 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 adapters that are compatible with Chelsio native network driver:

- T62100-CR
- T62100-LP-CR
- T62100-SO-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR
- T580-CR
- T580-LP-CR
- T580-SO-CR
- T540-CR
- T540-LP-CR
- T540-SO-CR
- T540-BT
- T520-CR
- T520-LL-CR
- T520-SO-CR
- T520-BT

1.2. Software Requirements

1.2.1. ESXi Requirements

The native network driver has been developed to run on the following 64-bit ESXi platforms:

- **Host:**
 - ESXi 8.0

- **Virtual Machine (with VFs):**
 - RHEL/Rocky Linux 9.5, 5.14.0-503.11.1.el9_5.x86_64
 - RHEL/Rocky Linux 9.4, 5.14.0-427.14.1.el9_4.x86_64
 - RHEL/Rocky Linux 8.10, 4.18.0-553.el8_10.x86_64
 - RHEL/Rocky Linux 8.9, 4.18.0-513.5.1.el8_9.x86_64
 - Kernel.org 6.12.16
 - Kernel.org 6.6.79
 - Windows Server 2025
 - Windows Server 2022

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 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 following command, so that the device manager performs a rescan:

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

3. Software/Driver Configuration and Fine-tuning

3.1. Multiple Adapters

By default, the cxl driver will initialize eight Chelsio ports. In case of using multiple adapters, set the `max_ports` module parameter and reboot the machine.

```
[root@host:~] esxcfg-module -s max_ports=N cxl
[root@host:~] reboot
```

Note *This setting is persistent across reboots and need not be applied every time.*

Example: To use three numbers of T540-CR (4-port) adapters, with a total of 12 Chelsio ports:

```
[root@host:~] esxcfg-module -s max_ports=12 cxl
[root@host:~] reboot
```

3.2. 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.

Syntax and 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.*

3.3. Adapter Configuration

The adapter's configuration should be updated for optimal performance in ESXi environment.

Note *This configuration is not required for T6 adapters.*

Run the following `cxgbtool` command and reboot the machine.

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

3.4. 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
```

3.5. Connecting a Virtual Machine

Follow the steps mentioned below to connect the 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.

3.6. Virtual Functions (SR-IOV)

3.6.1. 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.*

Example: 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
```

iii. Reboot the ESXi host to apply the changes.

iv. Check if VFs were instantiated successfully on the PCI bus by either using the shell prompt (using *lspci*) or GUI.

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

Address	Description	SR-IOV	Passthrough	Hardware Label
0000:00:02.0	Intel Corporation Xeon E7 v2/Keen E5 v2/Core i7 PCI Express Root Port 2a	Not capable	Not capable	
0000:04:08.5	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:08.4	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:08.1	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:08.0	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:07.5	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:07.4	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:07.1	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:07.0	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:06.5	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	
0000:04:06.4	Chelsio Communications Inc T6225-CR Unified Wire Ethernet Controller (vF)	Not capable	Active	

3.6.2. Assigning VFs to VMs

Once the SR-IOV VFs are enabled/instantiated successfully on the host, they can be associated with Virtual Machines (VMs). The VFs can be assigned to VMs as:

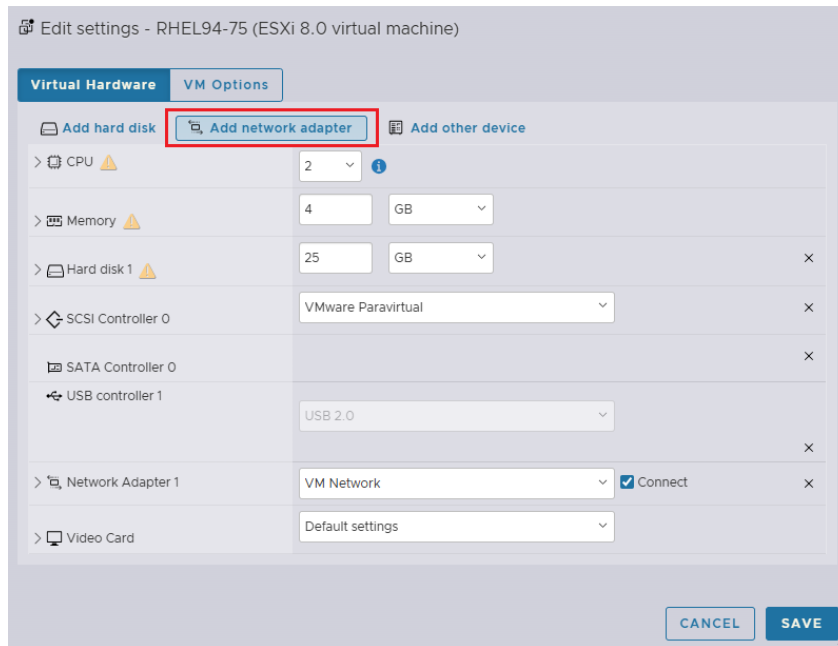
- SR-IOV passthrough network adapters
- PCI Devices



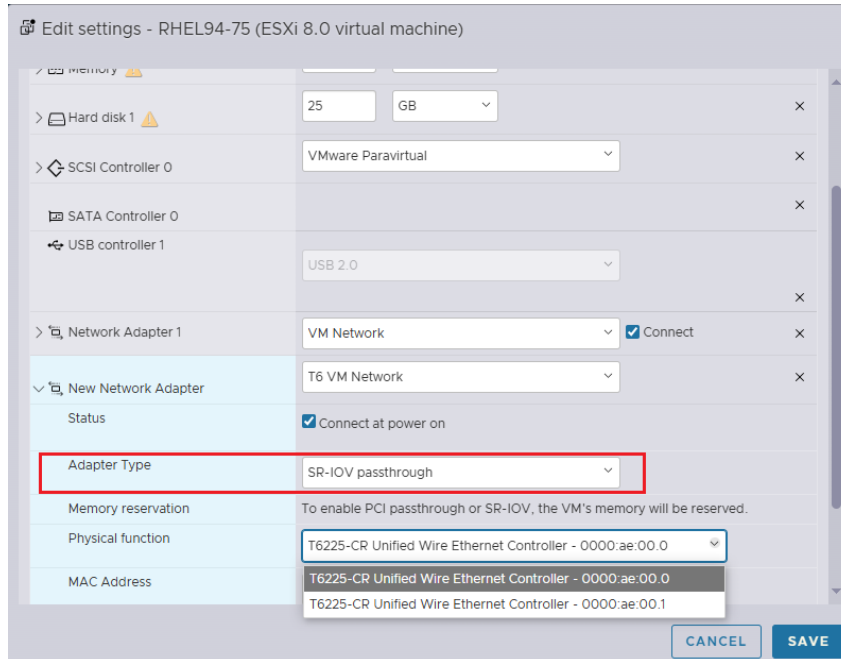
- *Unloading host network driver (cxl) when VFs are attached to VMs is not supported by VMware.*
- *T5 Adapters are not supported in Windows VMs with SR-IOV.*

- **SR-IOV passthrough network adapters**

i. Right-click on the VM and select **Edit Settings**. Click **Add network adapter**.



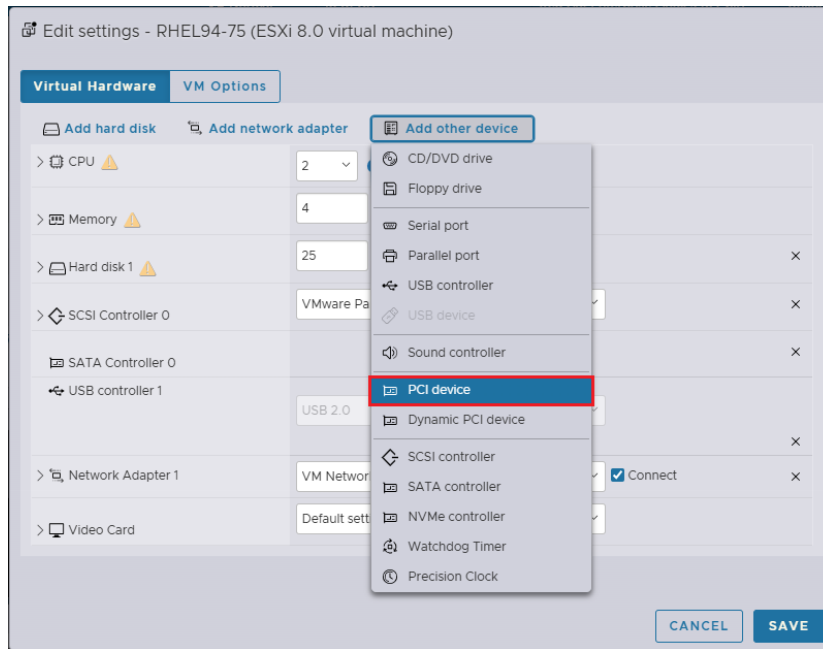
ii. Select the required Port Group and select **Adapter type** as **SR-IOV passthrough**.



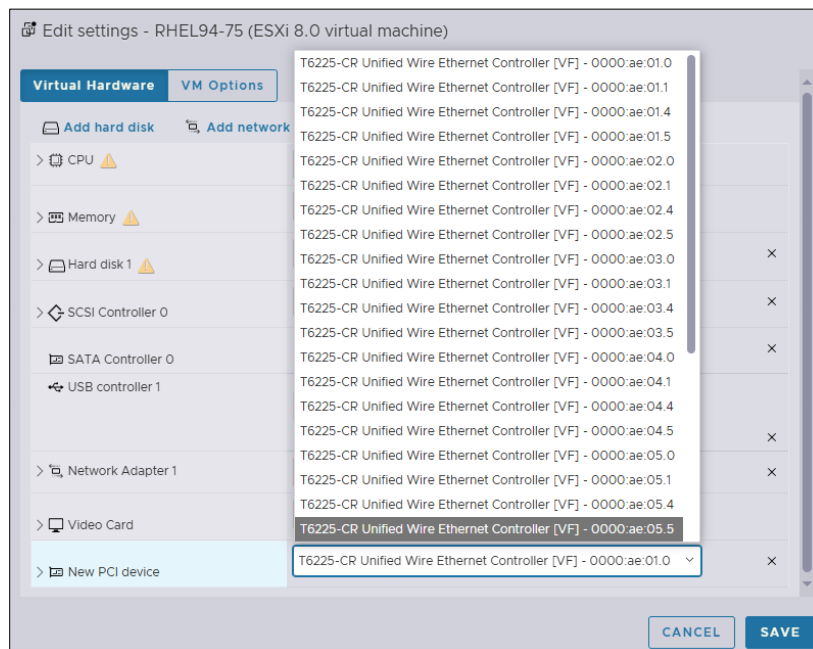
iii. Select the required Physical Function (Port) to use and click **Save**.

- **PCI Devices**

- Right-click on the VM and select **Edit Settings**. Click on **Virtual Hardware** tab. Click **ADD NEW DEVICE** and select **PCI Device**.



- Select the required Chelsio VF and click **Save**.



For more information on configuring SR-IOV, refer to VMware's official [documentation](#).

3.6.3. 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.

```
[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* 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, refer to the Chelsio Unified Wire for Linux User Guide.*

- vii. Bring up the VF interface with the IP address and it will communicate with other VFs or hosts.

```
[root@host~]# ifconfig ethX <IPv4/IPv6 address> up
```

3.6.4. Using VFs in Windows 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. Download the latest *Chelsio Unified Wire* driver package, from [Chelsio Download Center](#).
- iii. Install the Chelsio Unified Wire using the installer or zip package.

 **Note** *For detailed instructions, refer to the Chelsio Unified Wire for Windows User Guide.*

- iv. Assign the required IP addresses for the VF in Network Connections.

 **Important** *T5 Adapters are not supported in Windows VMs with SR-IOV.*

3.6.5. VF Link State

VF link state depends on the physical port link status by default. To override this and always enable the VF link, follow the below procedure. This will enable VF to VF communication irrespective of the physical port link status.

- i. Update the *vfstate* using the following command on the host.

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c vfopt -idx <VF ID> -vfstate 2 -a <adap> -p <port>
```

```
[root@zajila:~] /opt/chelsio/bin/cxgbtool -c vfopt -idx 2 -vfstate 2 -a 0 -p 1
Setting vf link status successful
```

- ii. View the current *vfstate*.

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c vfopt -idx <VF ID> -a <adap> -p <port>
```

```
[root@zajila:~] /opt/chelsio/bin/cxgbtool -c vfopt -idx 2 -a 0 -p 1
vf link state: 2
```

- iii. Toggle the VF interface on the VM to apply the changes.

```
[root@host~]# ifconfig ethX down
[root@host~]# ifconfig ethX up
```

vfstate can be set to 1 for default behavior.

3.6.6. 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 set up 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 results in the following connectivity:

- VFs of the same port can communicate with each other. That is 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.

3.6.7. Configuring VLANs

If the VFs are attached as SR-IOV passthrough network adapters to the VMs, VLAN can be configured at the port group. After configuring the VLAN, the VF network driver should be reloaded on the VM to apply the changes.


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

For Windows VMs, disable and enable the **Chelsio Bus Enumerator [Virtual Function]** in System Devices of the Device Manager to apply the changes.

4. Software/Driver Unloading

Execute the following command to unload the Native Network driver:

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

 **Note** *If iSCSI, iSER or NVMe-oF Offload Initiator Driver is loaded, unload it before unloading the native network driver.*

III. iSCSI 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.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the adapters that are compatible with Chelsio iSCSI Offload Initiator driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory Free; 256 IPv4/128 IPv6 offload connections supported*)
- T580-CR
- T580-LP-CR
- T540-CR
- T540-LP-CR
- T540-BT
- T520-CR
- T520-LL-CR
- T520-BT


2. 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
```

Execute the following command, so that device manager performs a rescan:

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

 **Note** *Execute the following command to restore the Advanced Options of storage adapter after cheiscsi reload:*

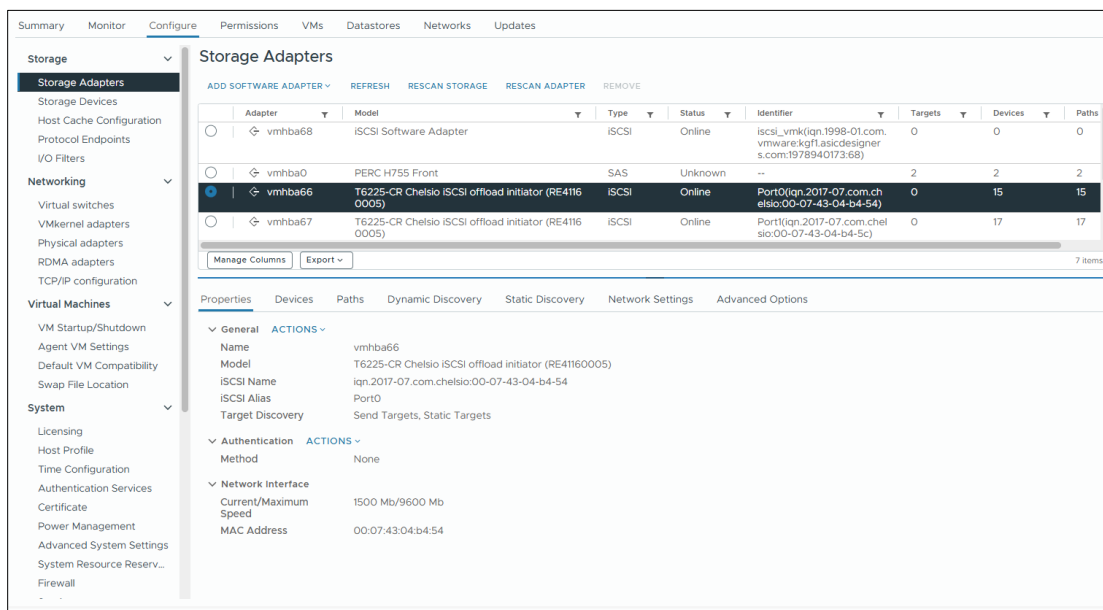
```
[root@host:~] esxcfg-rescan -A
```

3. Software/Driver Configuration and Fine-tuning

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

3.1. 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 the onscreen instructions and provide information to add the host. Click **Finish**.
- iii. Select the host and under the **Configure** tab, select **Storage Adapters**. This displays the list of available Chelsio iSCSI adapters.



You can also view the list in CLI using:

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

```
[root@KGF1:~] esxcli iscsi adapter list
Adapter  Driver  State  UID  Description
-----
vmhba66  cheiscsi  online  iscsi.vmhba66  T6225-CR Chelsio iSCSI offload initiator (RE41160005)
vmhba67  cheiscsi  online  iscsi.vmhba67  T6225-CR Chelsio iSCSI offload initiator (RE41160005)
```

- iv. In the **Adapter Details** section, click **Network Settings** tab and then **Edit**.

- v. Configure IPv4 address for the adapter and click **OK**.

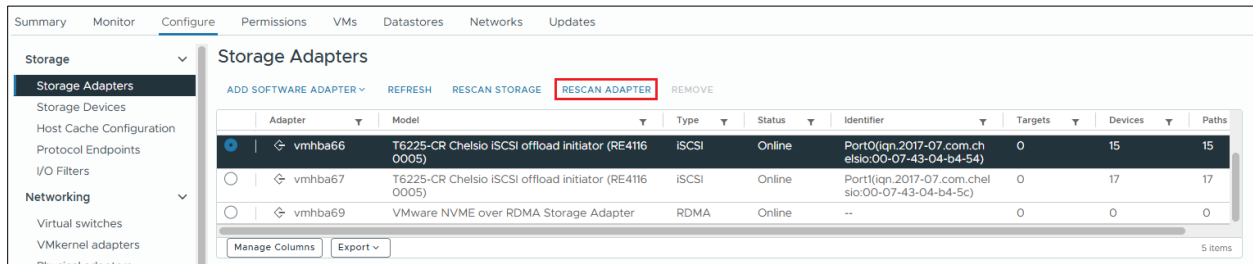
Properties	Devices	Paths	Dynamic Discovery	Static Discovery	Network Settings	Advanced Options
✓ IP Address and DNS Configuration ACTIONS						
IPv4 address					120.1.1.100 (static)	
Subnet mask for IPv4					255.255.255.0	
Default gateway for IPv4					120.1.1.1	
IPv6 address					Not enabled	
Preferred DNS server					::	
Alternate DNS server					::	

- vi. To use IPv6 address, use the following command:

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c chnet -set -6 -ipaddr <IPv6 address> -gw <IPv6 gateway> -plen <subnet mask> -p <port>
```

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c chnet -set -6 -ipaddr 2000::79 -gw 2000::1 -plen 64 -p 0
Setting chnet Configuration:
  Adapter name : vmhba64
  Node id      : 0
  Operation    : Setting Ipv6
  IP           : 2000::79
  prefix len   : 64
  Gateway      : 2000::1
  Status       : Success
```

vii. To apply the changes, click **RESCAN ADAPTER**.



3.2. Connecting to Target

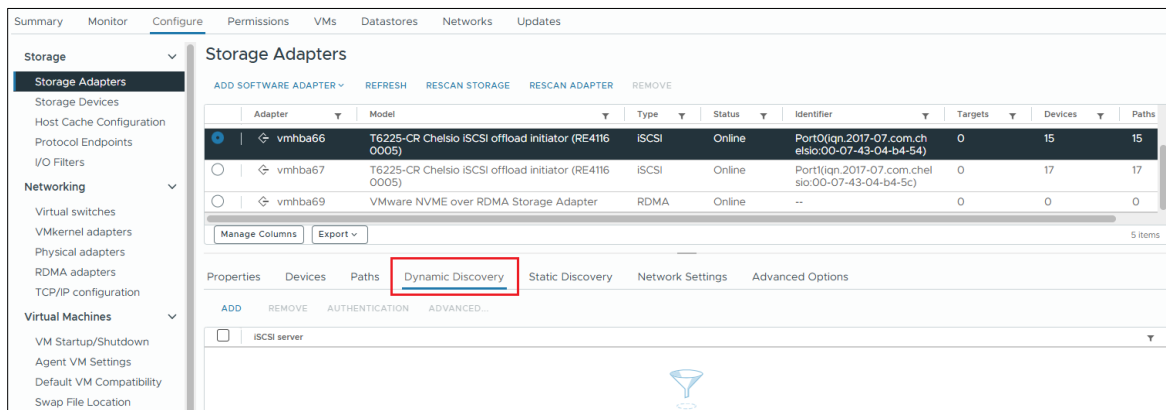
There are two methods to discover and connect to targets:

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

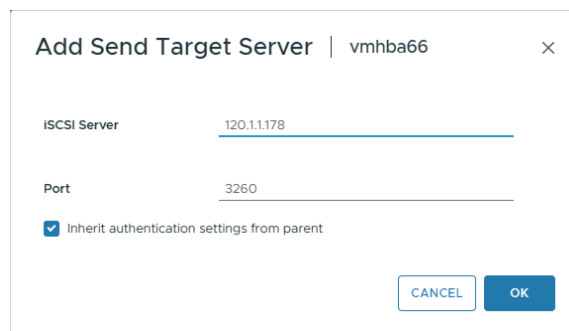
3.2.1. Dynamic Discovery

- **Adding Target Server**

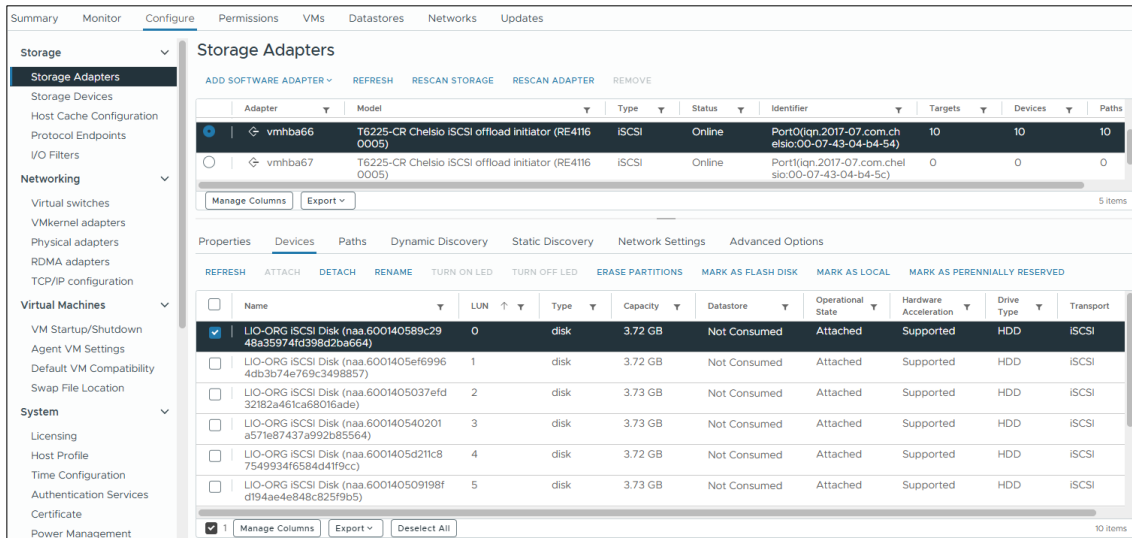
i. Select the iSCSI adapter to connect to the target and select **Dynamic Discovery**.




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



- iii. To apply the changes, rescan the iSCSI adapter.
- 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.




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



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

- **Removing Target Server**

- i. Under **Dynamic Discovery** tab, select target server, click **Remove** and then **OK** to confirm.

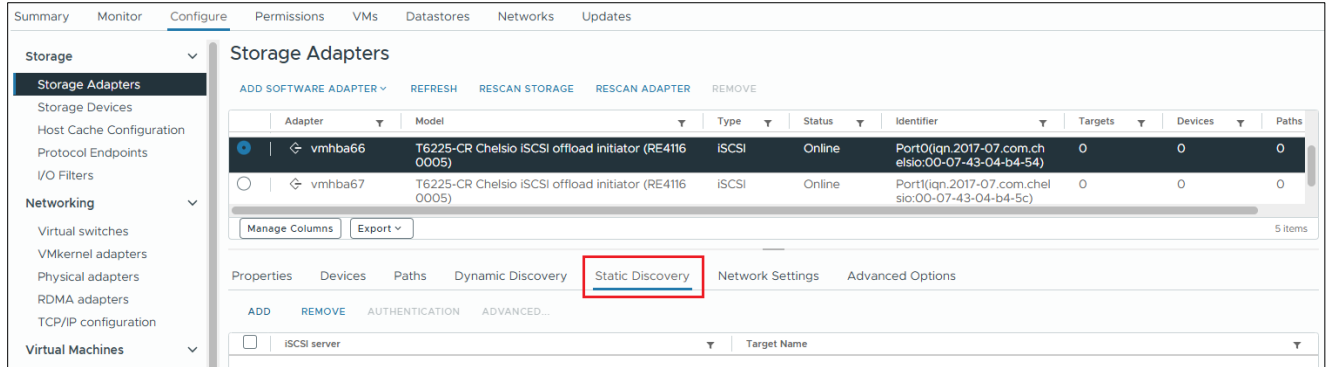


- ii. To apply the changes, rescan the iSCSI adapter.

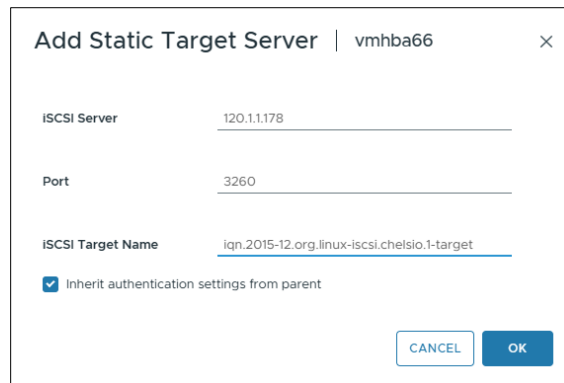
3.2.2. Static Discoverys

- **Adding Target Server**

i. Select the iSCSI interface to connect to the target and select **Static Discovery** tab.

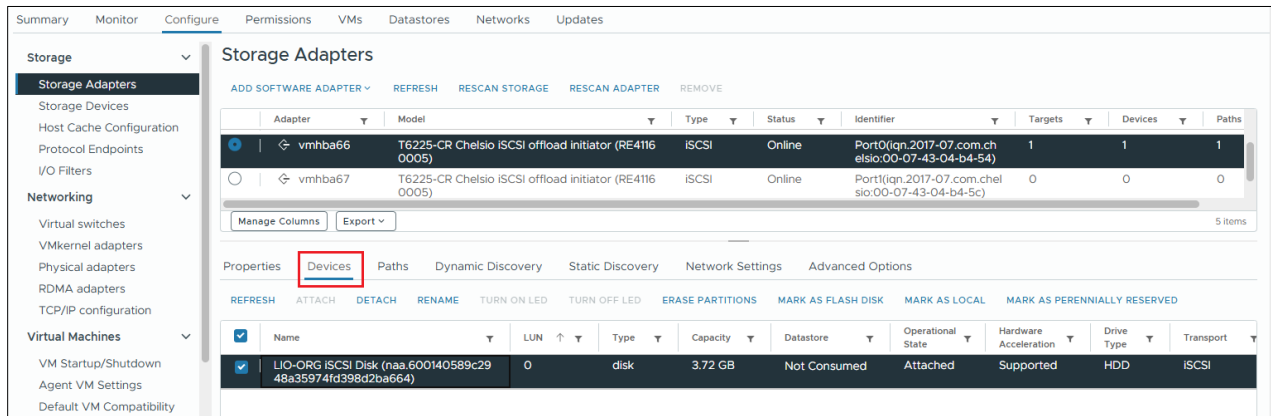


ii. Click **Add** and provide the target server IP and target IQN. Click **OK**.




iii. To apply the changes, rescan the iSCSI adapter.

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.




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



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

- **Removing Target Server**

i. Under **Static Discovery** tab, select the target server, click **Remove** and then **OK** to confirm.



ii. To apply the changes, rescan the iSCSI adapter.

3.3. Configurable Options

The option to edit general initiator properties like alias and name is available under the **Properties** tab.



Advanced parameters like Digest, MTU, etc., can be changed in the **Advanced Options** tab.

Advanced Options ✕

Option	Description	Value
Header Digest	iSCSI adapter option : Header Digest	Prohibited ▾
Data Digest	iSCSI adapter option : Data Digest	Prohibited ▾
ErrorRecoveryLevel	iSCSI option : iSCSI Error Recovery Level (ERL) value that the ESX initiator would negotiate during login.	0
LoginRetryMax	iSCSI option : Maximum number of times ESX initiator would retry login to a target, before giving up.	4
MaxOutstandingR2T	iSCSI option : Maximum number of R2T (Ready To Transfer) PDUs, that can be outstanding for a task.	1
FirstBurstLength	iSCSI option : Maximum unsolicited data in bytes initiator can send during the execution of a single SCSI command. It must be <= MaxBurstLength.	262144
MaxBurstLength	iSCSI option : Maximum SCSI data payload in bytes in a Data-In or a solicited Data-Out iSCSI sequence.	262144
MaxRecvDataSegLen	iSCSI option : Maximum data segment length in bytes that can be received in an iSCSI PDU. It is recommended to keep it <= MaxBurstLength.	131072
MaxCommands	iSCSI option : Maximum SCSI commands that can be queued on the iSCSI adapter.	128
DefaultTimeToWait	iSCSI option : Minimum seconds to wait before attempting a logout or an active task reassignment after an unexpected connection termination or reset.	2
DefaultTimeToRetain	iSCSI option : Maximum seconds that a connection and task allegiance reinstatement is still possible following a connection termination or reset.	0
LoginTimeout	iSCSI option : Time in seconds initiator will wait for the Login response	5

19 items

CANCEL
OK

4. Software/Driver Unloading

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

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

```
[root@KGF2:~] vmkload_mod -u cheiscsi  
Module cheiscsi successfully unloaded
```

IV. iSER Offload Initiator Driver

1. Introduction

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 adapters that are compatible with Chelsio iSER Offload Initiator driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory Free; 256 IPv4/128 IPv6 offload connections supported*)
- T580-CR
- T580-LP-CR
- T540-CR
- T540-LP-CR
- T540-BT
- T520-CR
- T520-LL-CR
- T520-BT

2. Software/Driver Loading

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

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

Execute the following command, so that the device manager performs a rescan:

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


3. Software/Driver Configuration and Fine-tuning

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

3.1. Configuring Initiator

- i. Enable the iSER adapter.

```
[root@host:~] esxcli rdma iser add
```

Note *This is not persistent across reboots. To make it persistent, add the above command to `/etc/rc.local.d/local.sh` file.*

- ii. Log in to vCenter Server through vSphere Web Client using a web browser.
- iii. If you have already created and configured the host intended to be used as initiator, skip to step (iv).
 - 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 the onscreen instructions and provide information to add the host. Click **Finish**.
- iv. Select the host and under the **Configure** tab, select **Storage Adapters**. This displays the list of available iSER adapters.

Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba64	VMware iSER over RDMA (iSER) Adapter	iSER	Unbound	iser-vmnic9(iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1161313288:64)	0	0	0
vmhba65	VMware iSER over RDMA (iSER) Adapter	iSER	Unbound	iser-vmnic10(iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1940944980:65)	0	0	0

Properties for selected adapter (vmhba64):

- General**
 - Name: vmhba64
 - Model: VMware iSER over RDMA (iSER) Adapter
 - iSER Name: iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1161313288:64
 - iSER Alias: iser-vmnic9
 - Target Discovery: Send Targets, Static Targets
- Authentication**
 - Method: None

v. Create a VMkernel adapter connected to Chelsio uplink by following the screenshots below.

VMkernel adapters							
Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services		
vmk0	Management Network	vSwitch0	10.193.204.114	Default	Management		
vmk1	VMkernel	vSwitch1	120.1.1.114	Default	--		

Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

Select connection type

Select a connection type to create.

- VMkernel Network Adapter**
The VMkernel TCP/IP stack handles traffic for ESXi services such as vSphere vMotion, iSCSI, NFS, FCoE, Fault Tolerance, vSAN, host management and etc.
- Virtual Machine Port Group for a Standard Switch**
A port group handles the virtual machine traffic on standard switch.
- Physical Network Adapter**
A physical network adapter handles the network traffic to other hosts on the network.

CANCEL **NEXT**

Add Networking

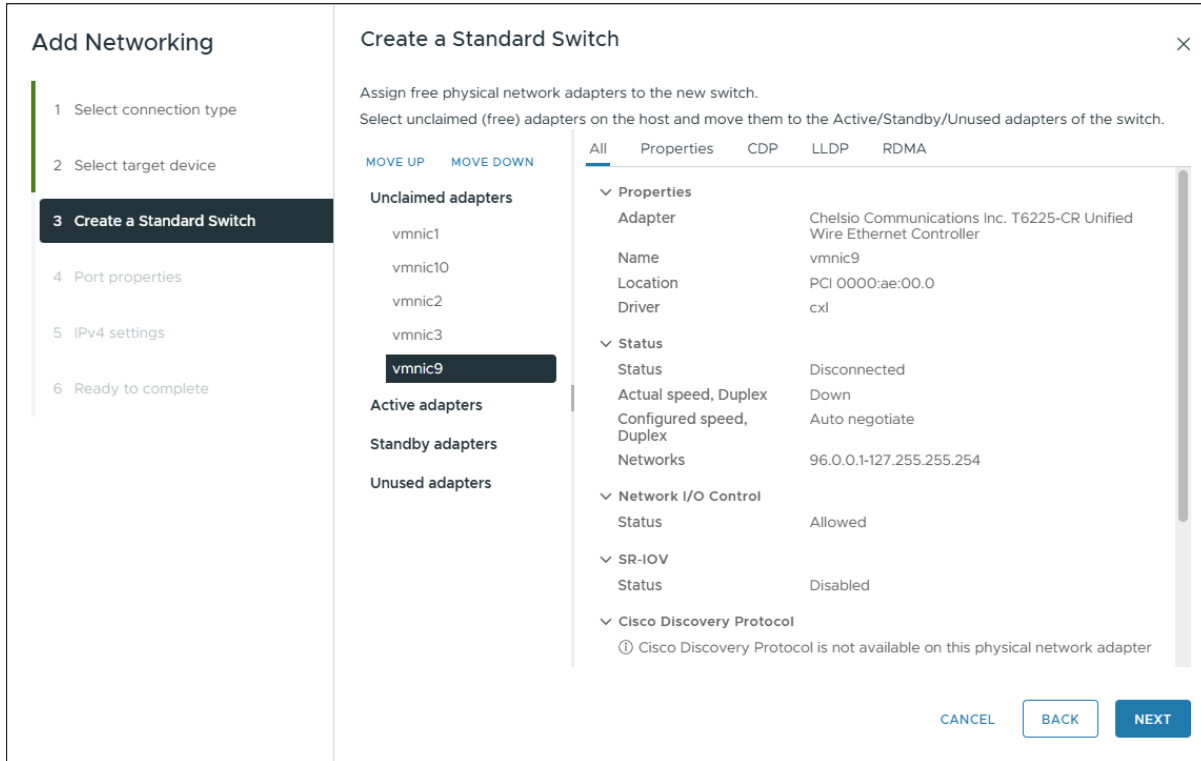
- 1 Select connection type
- 2 Select target device
- 3 Create a Standard Switch
- 4 Port properties
- 5 IPv4 settings
- 6 Ready to complete

Select target device

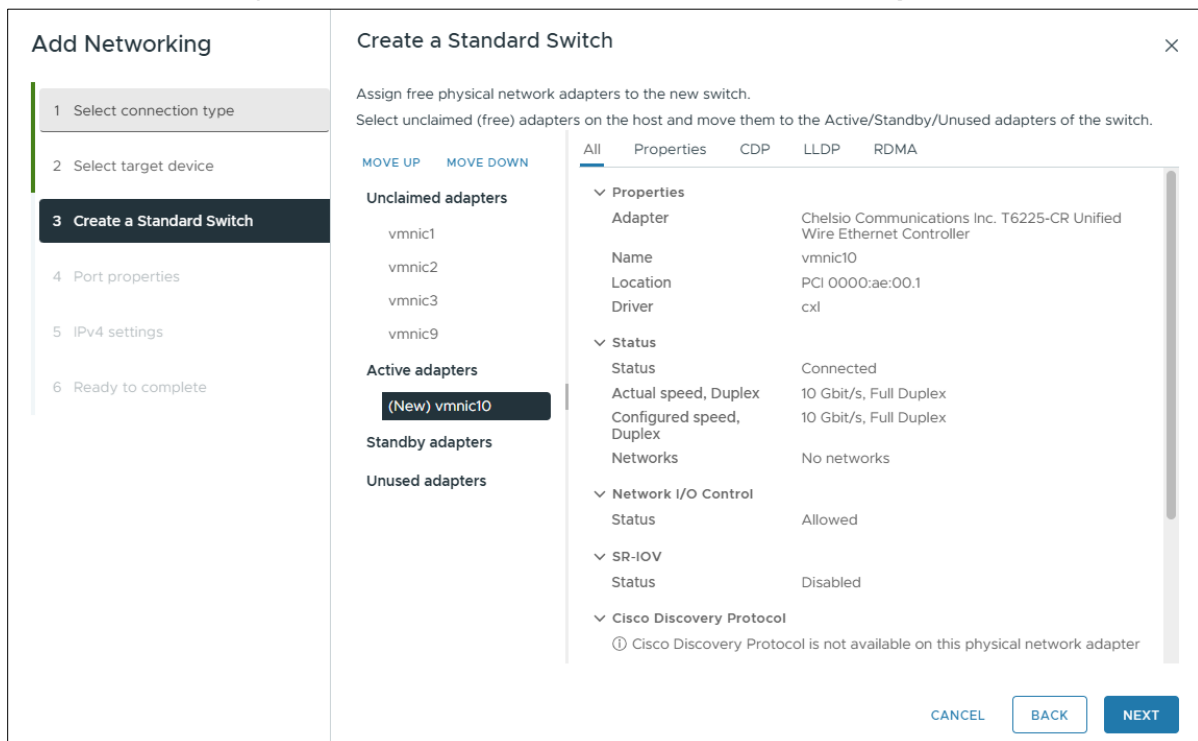
Select a target device for the new connection.

- Select an existing network
- Select an existing standard switch
- New standard switch**

MTU (Bytes) 1500



vi. Move the Chelsio ports to be used in the vSwitch to the **Active Adapters** section.



Networking		VMkernel adapters					
Virtual switches		ADD NETWORKING... REFRESH					
VMkernel adapters		Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services
Physical adapters		vmk0	Management Network	vSwitch0	10.193.204.114	Default	Management
RDMA adapters		vmk1	VMkernel	vSwitch1	120.1.1.114	Default	--
TCP/IP configuration							

vii. Bind the VMkernel adapter to the VMware iSER Adapter.

Storage		Storage Adapters							
Storage Adapters		ADD SOFTWARE ADAPTER... REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE							
Storage Devices		Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
Host Cache Configuration		vmhba64	VMware iSCSI over RDMA (iSER) Adapter	iSCSI	Unbound	iser-vmnic9(iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1161313288:64)	0	0	0
Protocol Endpoints		vmhba65	VMware iSCSI over RDMA (iSER) Adapter	iSCSI	Unbound	iser-vmnic10(iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1940944980:65)	0	0	0
I/O Filters		Manage Columns Export							
Networking		Properties Devices Paths Dynamic Discovery Static Discovery Network Port Binding Advanced Options							
Virtual switches		ADD REMOVE VIEW DETAILS							
VMkernel adapters		Port Group	VMkernel Adapter	Port Group Policy	Path Status	Physical Network Adapter			

Bind vmhba65 with VMkernel Adapter		
<input checked="" type="checkbox"/>	Port Group	VMkernel Adapter
<input checked="" type="checkbox"/>	VMkernel (vSwitch1)	vmk1
	Physical Network Adapter	vmnic10 (10 Gbit/s, Full)

viii. To apply the changes, click **RESCAN ADAPTER**. The status should show **Online**.

Storage		Storage Adapters							
Storage Adapters		ADD SOFTWARE ADAPTER... REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE							
Storage Devices		Due to recent configuration changes, a rescan of "vmhba65" is recommended.							
Host Cache Configuration		Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
Protocol Endpoints		vmhba64	VMware iSCSI over RDMA (iSER) Adapter	iSCSI	Unbound	iser-vmnic9(iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1161313288:64)	0	0	0
I/O Filters		vmhba65	VMware iSCSI over RDMA (iSER) Adapter	iSCSI	Online	iser-vmnic10(iqn.1998-01.com.vmware.kgfl.asicdesigners.com:1940944980:65)	0	0	0
Networking		Manage Columns Export							
Virtual switches		Properties Devices Paths Dynamic Discovery Static Discovery Network Port Binding Advanced Options							
VMkernel adapters		ADD REMOVE VIEW DETAILS							
Physical adapters		Port Group	VMkernel Adapter	Port Group Policy	Path Status	Physical Network Adapter			
RDMA adapters		VMkernel (vSwitch1)	vmk1	Compliant	Not used	vmnic10 (10 Gbit/s, Full)			
TCP/IP configuration									
Virtual Machines									
VM Startup/Shutdown									
Agent VM Settings									
Default VM Compatibility									
Swap File Location									

Note Ensure that the VMkernel NIC is associated with a vSwitch containing only a single physical vmnic, as each physical vmnic maps to a distinct iSER vmhba adapter in ESXi.

3.2. Connecting to Target

Configure the iSER target machine with the IP Address, Target name, disks etc. For information on how to configure the iSER Target, refer to the *Chelsio Unified Wire for Linux User Guide*.

Important Enable `iwpmid` service on the target machine. On RHEL7.X machines, use the following command:

```
[root@host~]# systemctl start iwpmid
```

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.

3.2.1. Dynamic Discovery

- **Adding Target Server**

i. Select the iSER adapter to connect to the target and select **Dynamic Discovery**.

The screenshot shows the vSphere Storage Adapters configuration page. The 'Dynamic Discovery' tab is selected and highlighted with a red box. Below the tab, there is an 'ADD' button and a list of iSCSI servers.

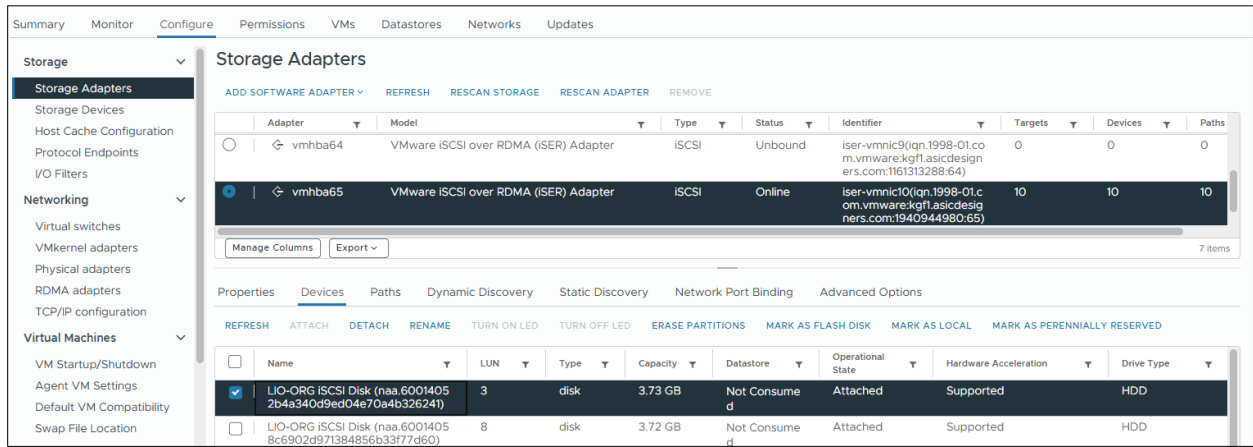
Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba64	VMware iSCSI over RDMA (iSER) Adapter	iSCSI	Unbound	iser-vmnic9(ign.1998-01.com.vmware.kgfl.asicdesigners.com:1161313298:64)	0	0	0
vmhba65	VMware iSCSI over RDMA (iSER) Adapter	iSCSI	Online	iser-vmnic10(ign.1998-01.com.vmware.kgfl.asicdesigners.com:1940944980:65)	0	0	0

Below the table, there are tabs for Properties, Devices, Paths, Dynamic Discovery (selected), Static Discovery, Network Port Binding, and Advanced Options. Under the Dynamic Discovery tab, there is an 'ADD' button and a list of iSCSI servers.


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

The screenshot shows the 'Add Send Target Server' dialog box for adapter vmhba65. The 'iSCSI Server' field contains '120.1.1.178' and the 'Port' field contains '3260'. The 'Inherit authentication settings from parent' checkbox is checked. There are 'CANCEL' and 'OK' buttons at the bottom right.

- iii. To apply the changes, click **RESCAN ADAPTER**.
- 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.



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




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


- **Removing Target Server**

During **Dynamic Discovery**, iSER targets also appear as static targets and need to be removed from the **Static** list.

- i. Go to the **Static Discovery** tab, and select all the target iSCSI servers.
- ii. Click **Remove** and then **OK** to confirm.



- iii. Go to the **Dynamic Discovery** tab, and select all the target iSCSI servers.
- iv. Click **Remove** and then **OK** to confirm.

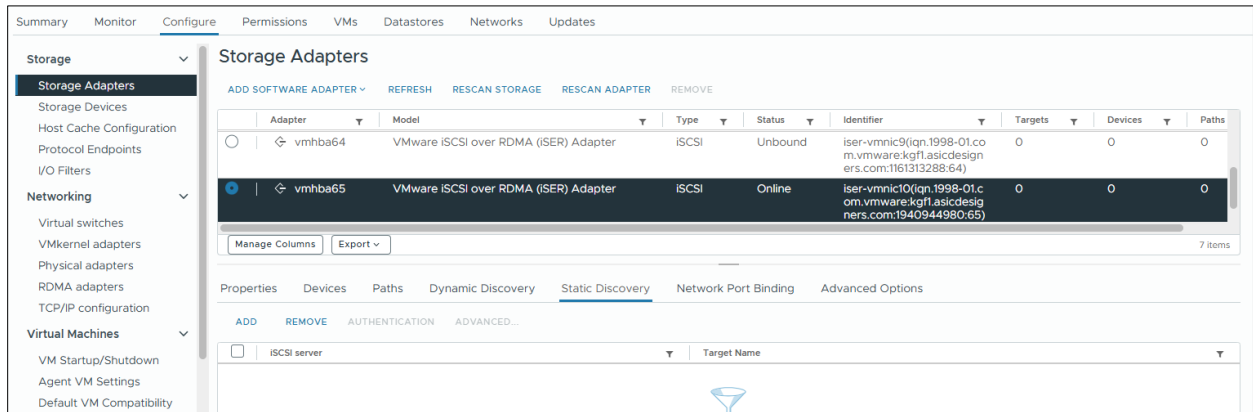


v. To apply the changes, rescan the iSER adapter.

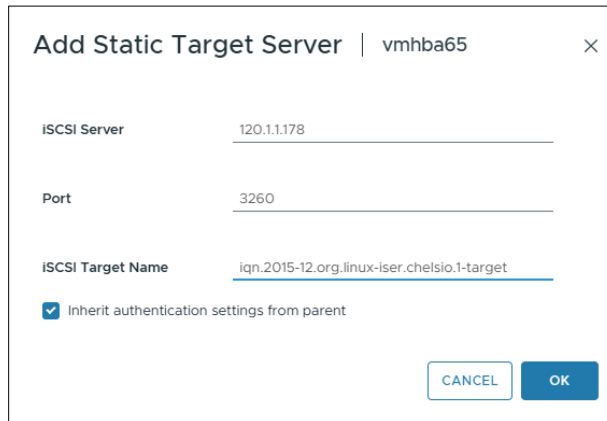
3.2.2. Static Discovery

- **Adding Target Server**

i. Select the iSER interface to connect to the target and select **Static Discovery** tab.



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



iii. To apply the changes, rescan the iSER adapter.

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.

Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type
LIO-ORG iSCSI Disk (naa.600140585856e8cdac746cd84585b310)	0	disk	3.72 GB	Not Consumed	Attached	Supported	HDD

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

Runtime Name	Target	LUN	Status
vmhba65:C0:T0:L0	iqn.2015-12.org.linux-iser.chelsio.1-target:120.1.1.178:3260	0	Active (I/O)

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

- **Removing Target Server**
 - i. Select the target server in **Static Discovery** tab.
 - ii. Click **Remove** and then **Yes** to confirm.

ISCSI server	Target Name
120.1.1.178:3260	iqn.2015-12.org.linux-iser.chelsio.1-target

- iii. To apply the changes, rescan the iSER adapter.

3.3. Configurable Options

The option to edit general initiator properties like alias and name is available under the **Properties** tab.

Edit General | vmhba65 ✕

ISCSI Name:

ISCSI Alias:

Advanced parameters like Digest, MTU, etc., can be changed in the **Advanced Options** tab.

Advanced Options ✕

Option	Description	Value
Header Digest	iSCSI adapter option : Header Digest	Prohibited ▾
Data Digest	iSCSI adapter option : Data Digest	Prohibited ▾
ErrorRecoveryLevel	iSCSI option : iSCSI Error Recovery Level (ERL) value that the ESX initiator would negotiate during login.	0
LoginRetryMax	iSCSI option : Maximum number of times ESX initiator would retry login to a target, before giving up.	4
MaxOutstandingR2T	iSCSI option : Maximum number of R2T (Ready To Transfer) PDUs, that can be outstanding for a task.	1
FirstBurstLength	iSCSI option : Maximum unsolicited data in bytes initiator can send during the execution of a single SCSI command. It must be <= MaxBurstLength.	262144
MaxBurstLength	iSCSI option : Maximum SCSI data payload in bytes in a Data-In or a solicited Data-Out iSCSI sequence.	262144
MaxRecvDataSegLen	iSCSI option : Maximum data segment length in bytes that can be received in an iSCSI PDU. It is recommended to keep it <= MaxBurstLength.	131072
MaxCommands	iSCSI option : Maximum SCSI commands that can be queued on the iSCSI adapter.	128
DefaultTimeToWait	iSCSI option : Minimum seconds to wait before attempting a logout or an active task reassignment after an unexpected connection termination or reset.	2
DefaultTimeToRetain	iSCSI option : Maximum seconds that a connection and task allegiance reinstatement is still possible following a connection termination or reset.	0
LoginTimeout	iSCSI option : Time in seconds initiator will wait for the Login response	5

19 items

CANCEL
OK

4. Software/Driver Unloading

Logout all the existing iSER sessions. Execute the following command to unload the iSER Offload Initiator driver:

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

V. NVMe-oF Offload Initiator Driver

1. Introduction

NVMe over Fabrics specification extends the benefits of NVMe to large fabrics, beyond the reach and scalability of PCIe. NVMe enables deployments with hundreds or thousands of SSDs using a network interconnect, such as RDMA over Ethernet. Thanks to an optimized protocol stack, an end-to-end NVMe solution is expected to reduce access latency and improve performance, particularly when paired with a low latency, high efficiency transport such as RDMA. This allows applications to achieve fast storage response times, irrespective of whether the NVMe SSDs are attached locally or accessed remotely across enterprise or datacenter networks. Chelsio Unified Wire adapters with iWARP RDMA offload provide the high bandwidth and low latency fabric needed for NVMe-oF.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio adapters that are supported:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory Free; 256 IPv4/128 IPv6 offload connections supported*)
- T580-CR
- T580-LP-CR
- T540-CR
- T540-LP-CR
- T540-BT
- T520-CR
- T520-LL-CR
- T520-BT

2. Software/Driver Loading

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

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

Execute the following command, so that the device manager performs a rescan:

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

3. Software/Driver Configuration and Fine-tuning

3.1. Connecting to NVMe target

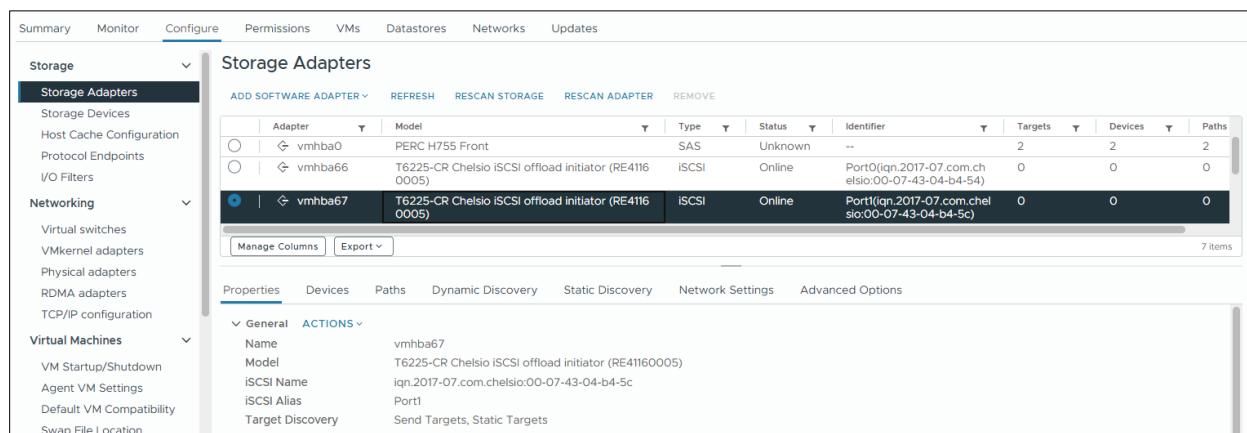
Configure the NVMe target machine with the IP Address, Target name, disks etc. For information on how to configure the NVMe Target, refer to the *Chelsio Unified Wire for Linux User Guide*.

Important Disable *iwpmid* service on the target machine. On RHEL7.X machines, use the following command:

```
[root@host~]# systemctl stop iwpmid
```

Follow the below procedure on NVMe Initiator machine to connect to the target:

- 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 displays the list of available Chelsio storage adapters.


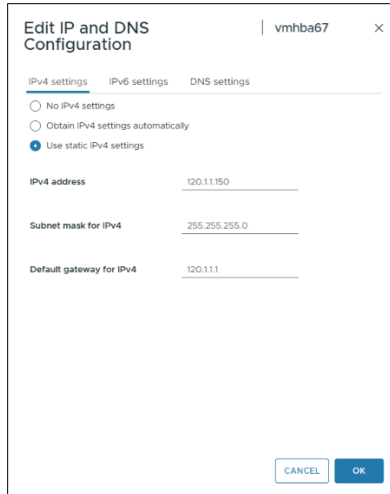


Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba0	PERC H755 Front	SAS	Unknown	--	2	2	2
vmhba66	T6225-CR Chelsio iSCSI offload initiator (RE4116 0005)	iSCSI	Online	Port0(iqn.2017-07.com.chelsio:00-07-43-04-b4-54)	0	0	0
vmhba67	T6225-CR Chelsio iSCSI offload initiator (RE4116 0005)	iSCSI	Online	Port1(iqn.2017-07.com.chelsio:00-07-43-04-b4-5c)	0	0	0

Properties: **General**

- Name: vmhba67
- Model: T6225-CR Chelsio iSCSI offload initiator (RE41160005)
- iSCSI Name: iqn.2017-07.com.chelsio:00-07-43-04-b4-5c
- iSCSI Alias: Port1
- Target Discovery: Send Targets, Static Targets

- iv. In the **Adapter Details** section, click **Network Settings** tab and then **Edit**.
- v. Configure IPv4 address for the adapter and click **OK**.



vi. Use the following command to configure the IPv6 address:

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c chnet -set -6 -ipaddr <IPv6 address> -gw <IPv6 gateway> -plen <subnet mask> -p <port>
```

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c chnet -set -6 -ipaddr 2000::79 -gw 2000::1 -plen 64 -p 0
Setting chnet Configuration:
Adapter name : vmhba64
Node id      : 0
Operation    : Setting Ipv6
IP           : 2000::79
prefix len   : 64
Gateway      : 2000::1
Status       : Success
```

vii. Discover the target.

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c nvme -tport <target_port> -ipaddr <target_ip_address> -p <Chelsio_Port #> -D
```

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c nvme -tport 4420 -ipaddr 10.200.200.5 -p 0 -D
process_nvme_resp status 0x0 nrec 0x10
0) nvme subsystem nvme-ram0 ctrl id 0xffff addr 10.200.200.5
1) nvme subsystem nvme-ram1 ctrl id 0xffff addr 10.200.200.5
```

If *-p* is not specified, by default Port 0 will be used.

While using IPv6, specify the target IPv6 address within [].



```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c nvme -tport 4420 -ipaddr [1000::146] -p 0 -D
```

Login to the target by specifying the target name.

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c nvme -tport <target_port> -
ipaddr <target_ip_address> -p <Chelsio_Port #> -L -subnqn <target_name>
```

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c nvme -tport 4420 -ipaddr 10.200.200.5 -p 0 -L -subnqn nvme-ram0
login status 0
target id 0
```

viii. Rescan the storage adapter and the target LUNs will be visible.

```
[root@host:~] esxcfg-rescan -A
```

ix. List the logged in targets.

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c nvme -tlist -p <Chelsio_Port #>
```

If `-p` is not specified, by default Port 0 will be used.

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c nvme -tlist
0) tgt_id: 0, state: 5, IP: 10.200.200.5, Port: 4420, sub_nqn: nvme-ram0
```

S

x. All the available LUNs will be displayed in the **Devices** tab. These LUNs can now be attached to VMs or can be used to store VMs.

The screenshot shows the 'Devices' tab in the vSphere Web Client. It displays a table of storage devices. The table has columns for Name, LUN, Type, Capacity, Datastore, Operational State, Hardware Acceleration, and Drive Type. A single device is listed: 'NVMe iSCSI Disk (10.NVMe Linux 73f0df9717c685c56c4c0000001)' with LUN 0, Type 'disk', Capacity '3.72 GB', Datastore 'Not Consumed', Operational State 'Attached', Hardware Acceleration 'Unknown', and Drive Type 'Flash'.

Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type
NVMe iSCSI Disk (10.NVMe Linux 73f0df9717c685c56c4c0000001)	0	disk	3.72 GB	Not Consumed	Attached	Unknown	Flash

3.2. Disconnecting from NVMe target

To logout or disconnect from the NVMe target:

```
[root@host:~] /opt/chelsio/bin/cxgbtool -c nvme -tport <target_port> -
ipaddr <target_ip_address> -p <Chelsio_Port #> -LT -all
```

If `-p` is not specified, by default Port 0 will be used.

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c nvme -tport 4420 -ipaddr 10.200.200.5 -p 0 -LT -all
Log out: tgt id 0
Logout status 0
```


4. Software/Driver Unloading

Logout of all the existing NVMe-oF sessions. Execute the following command to unload the driver:

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

VI. Appendix

1. Troubleshooting

- **Logs collection**

In case of any issues, collect the below logs:

- `/var/log/vmkernel.log`
- Adapter logs (*dump_file*) using the below command:

```
[root@host:~]/opt/chelsio/bin/cxgbtool -c cudbg -d all -f <dump_file> -a <adap>
```

```
[root@localhost:~] /opt/chelsio/bin/cxgbtool -c cudbg -d all -f /productLocker/cudbg.dmp -a 0  
Writing 51347516 bytes to /productLocker/cudbg.dmp
```

In case of a PSOD, additionally provide the vmkernel zdump from `/var/core/` directory.

Please contact Chelsio support at support@chelsio.com with all relevant logs for any issues.

2. 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.
735 N Pastoria Avenue,
Sunnyvale, CA 94085
U.S.A