Chelsio Communications

Accelerate

Chelsio T4 Unified Wire for XenServer

Installation and User's Guide

Chelsio Communications (Headquarters)
370 San Aleso Ave.
Suite 100
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-560076
Karnataka,
India

Tel: +1-91-80-4039-6800

Chelsio KK (Japan)

SHIMA Akasaka Bldg.
Minato-ku, Tokyo
Japan 107-0052

Tel: 03-6234-4353

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

Document History

Version	Revision Date
1.0.0	03/15/2013
1.0.1	04/01/2013

TABLE OF CONTENTS

I. Chelsio Unified Wire

1. Introduction

Thank you for choosing Chelsio T4 Unified Wire adapters. These high speed, single chip, single firmware cards provide enterprises and data centers with high performance solutions for various Network and Storage related requirements.

The T4 adapters can fully offload TCP, UDP, iSCSI, iWARP and FCoE over a single Unified Wire. The adapters also fully support SR-IOV, EVB/VNTag, DCB, Traffic Management and Filtering.

Ideal for all data, storage and high performance clustering applications, the T4 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 data centers, cloud service installations and high performance computing environments, Chelsio T4 adapters bring a new level of performance metrics and functional capabilities to the computer networking industry.

1.1. Features

Chelsio's Unified Wire software is an ISO image which installs Network, iSCSI PDU Offload Initiator and FCoE Full Offload Initiator drivers for Chelsio's T4 10/1 G adapters on XenServer.

1.2. Hardware Requirements

The Chelsio T4 Unified Wire software supports Chelsio T4 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 T4 Unified Wire software has been developed to run on XenServer based platforms. 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 software supplied is an ISO image. It consists of the following:

- RPM packages containing drivers and associated tools, firmware, etc.

- Shell script to install the driver package

2. Hardware Installation

Follow the steps to install Chelsio Adapter in your system:

1. Shutdown and power off your system.
2. Power off all remaining peripherals attached to your system.
3. Unpack the Chelsio adapter and place it on an anti-static surface.
4. Remove the system case cover according to the system manufacturer's instructions.
5. Remove the PCI filler plate from the slot where you will install the 10Gb Ethernet adapter.
6. For maximum performance, it is highly recommended to install the adapter into a PCIE x8 slot.
7. 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 SFP (transceiver) modules prior to inserting the adapter.
8. Secure the Chelsio adapter with a screw, or other securing mechanism, as described by the system manufacturer's instructions. Replace the case cover.
9. After securing the card, ensure that the card is still fully seated in the PCIE x8 slot as sometimes the process of securing the card causes the card to become unseated.
10. Connect a fiber cable, multi-mode for short range (SR) optics or single-mode for long range (LR) optics, to the 10Gb Ethernet adapter or regular Ethernet cable for the 1Gb Ethernet adapter.
11. Power on your system.
12. Verify if the T4 adapter was installed successfully by using the following command:

```
[root@host]# lspci | grep –i Chelsio
03:00.0 Ethernet controller: Chelsio Communications Inc T420-CR Unified Wire
Ethernet Controller
03:00.1 Ethernet controller: Chelsio Communications Inc T420-CR Unified Wire
Ethernet Controller
03:00.2 Ethernet controller: Chelsio Communications Inc T420-CR Unified Wire
Ethernet Controller
03:00.3 Ethernet controller: Chelsio Communications Inc T420-CR Unified Wire
Ethernet Controller
03:00.4 Ethernet controller: Chelsio Communications Inc T420-CR Unified Wire
Ethernet Controller
03:00.5 SCSI storage controller: Chelsio Communications Inc T420-CR Unified
Wire Storage Controller
03:00.6 Fibre Channel: Chelsio Communications Inc T420-CR Unified Wire
Storage Controller
03:00.7 Ethernet controller: Chelsio Communications Inc Device 0000
```

For Chelsio T4 adapters, the physical functions are currently assigned as:

- Physical functions 0 - 3: for the SR-IOV functions of the T4
- Physical function 4: for all NIC functions of the card
- Physical function 5: for iSCSI
- Physical function 6: for FCoE
- Physical function 7: Currently not assigned

Once the driver is installed and loaded, examine the output of `dmesg` to see if the card is discovered:

```
eth0: Chelsio T420-CR rev 2 10GBASE-SFP RNIC PCIe x8 MSI-X
0000:04:00.4: S/N: PT18111226, P/N: 110112140D0
```

This output indicates the hardware configuration of the card as well as the Serial number of the card. As observed by the x8, the card is properly installed in an x8 slot on the machine and the card is using MSI-X interrupts.

3. Software/Driver Installation

Chelsio Unified Wire must be installed by the root user. Any attempt to install the software as a regular user will fail.

There are two methods of installation: During operating system installation and Runtime.

3.1. During OS installation

1. If you haven't done so already, download the Unified Wire software (ISO image) from Chelsio Download Center, http://service.chelsio.com
2. Burn the image to a CD/DVD.
3. Start an interactive install of XenServer.
4. When the installer prompts, opt for installing drivers through Supplemental Packs. Choose "Yes" and hit [Enter].

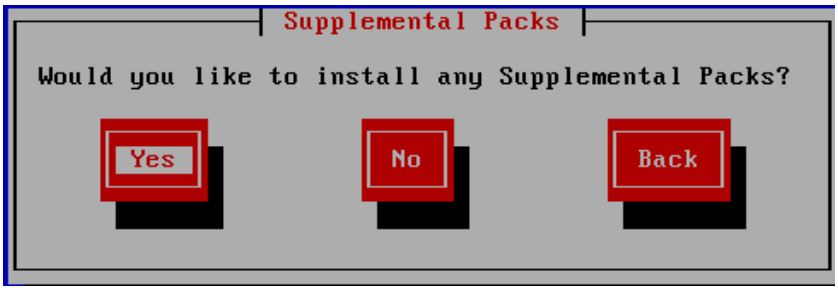

5. Insert the disc into disc drive. Choose "Ok" and hit [Enter].

6. The installer will now search and display the drivers. Choose “Use” and hit [Enter].

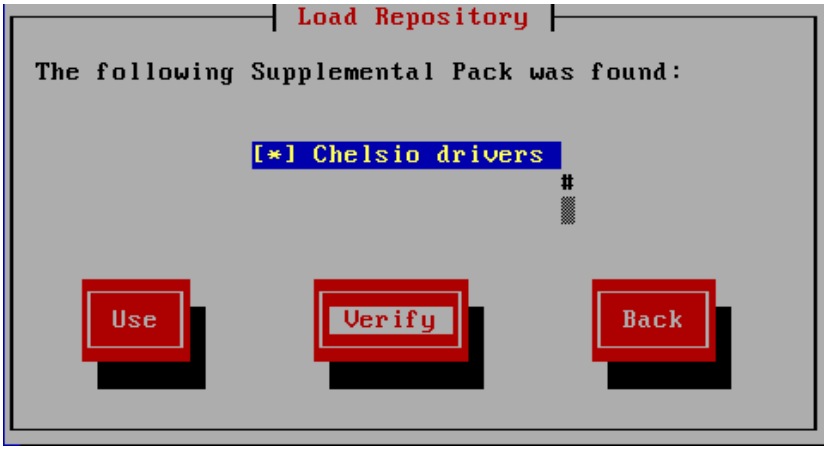

7. The selected drivers will now be installed.

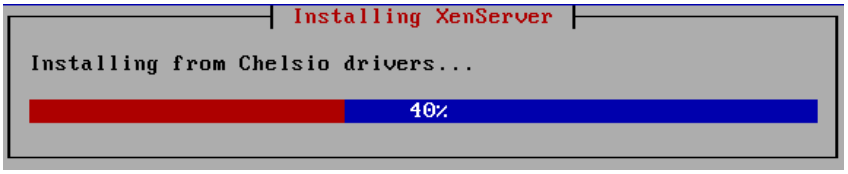

8. To install any additional drivers follow the same procedure mentioned above. Or choose “Skip” and hit [Enter] to complete the driver installation.

3.2. Runtime Installation

1. If you haven’t done so already, download the Unified Wire software (ISO image) from Chelsio Download Center, http://service.chelsio.com
2. Copy the ISO image on to a host running XenServer.
3. Mount the image onto a directory by running the following commands:

```
[root@host]# mkdir /tmp/iso
[root@host]# mount -o loop <driver.iso> /tmp/iso
[root@host]# cd /tmp/iso
```

4. Run the `install.sh` script to install the drivers and tools：

```
[root@host]# ./install.sh
```

5. Finally, unmount the ISO image:

```
[root@host]# cd
[root@host]# umount /tmp/iso
```

3.3. Firmware Update

The T4 firmware is installed on the system, typically in `/lib/firmware/cxgb4`, and the driver will auto-load the firmware if an update is required.

The firmware version can be verified using *ethtool*:

```
[root@host]# ethtool -i <iface>
```

4. Software/Driver Uninstallation

1. Use the following query command to determine the name of the driver RPM:

```
[root@host]# rpm -qa | grep -i <driver_name>
```

2. Now, with the result from the above query, execute the following command to uninstall all related packages:

```
[root@host]# rpm -e <package_name>
```

E.g.:

```
[root@host~]# rpm -qa | grep -i cxgb4
firmware-chelsio-cxgb4-01.04.10-1
cxgb4-t4fw-1.8.4-9
cxgb4-modules-xen-2.6.32.43-0.4.1.xs1.6.10.734.170748-1.2.4-9
cxgb4-tools-1.16-9
cxgb4-phyfw-1.2.4-9
cxgb4-modules-kdump-2.6.32.43-0.4.1.xs1.6.10.734.170748-1.2.4-9

[root@host ~]# rpm -e cxgb4-modules-kdump-2.6.32.43-
0.4.1.xs1.6.10.734.170748-1.2.4-9
[root@host ~]# rpm -e cxgb4-t4fw-1.8.4-9
[root@host ~]# rpm -e cxgb4-phyfw-1.2.4-9
[root@host ~]# rpm -e cxgb4-modules-xen-2.6.32.43-0.4.1.xs1.6.10.734.170748-
1.2.4-9
[root@host ~]# rpm -e cxgb4-tools-1.16-9
```

5. Software/Driver Update

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

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

II. Network (NIC)

1. Introduction

Chelsio's T4 series of Unified Wire Adapters provide extensive support for NIC operation, including all stateless offload mechanisms for both IPv4 and IPv6 (IP, TCP and UDP checksum offload, LSO - Large Send Offload aka TSO - TCP Segmentation Offload, and assist mechanisms for accelerating LRO - Large Receive Offload).

A high performance fully offloaded and fully featured TCP/IP stack meets or exceeds software implementations in RFC compliance. Chelsio's T4 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, using standard size Ethernet frames.

TCP offload is fully implemented in the hardware, thus freeing the CPU from TCP/IP overhead. The freed CPU can be used for any computing needs. The TCP offload in turn removes network bottlenecks and enables applications to take full advantage of the networking capabilities.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio T4 Adapters that are compatible with Chelsio Network Driver software:

- T420-CR
- T420-LL-CR
- T440-CR
- T440-LP-CR
- T420-BCH
- T422-CR
- T420-SO-CR
- T420-CX
- T420-BT
- T404-BT

1.2. Software Requirements

1.2.1. XenServer Requirements

The Chelsio Network driver has been developed to run on XenServer platforms. Currently the driver is available for the following version:

- Citrix XenServer Host 6.1.0
  Kernel version: 2.6.32.43-0.4.1.xs1.6.10.734.170748xen

Other versions have not been tested and are not guaranteed to work.

1.2.2. Resource Requirements

- **Control Domain (Dom0) memory**: Based on the number of 10G and 1G ports, driver allocates memory enough to run at the desired rate. It is suggested to increase the control domain memory to 2GB using the procedure explained in http://support.citrix.com/article/CTX134951

- **SW-IOMMU Size**: Increasing SW-IOMMU size by appending `swiotlb=128` to `/boot/extlinux.conf` avoids `PCI-DMA: Out of SW-IOMMU space for <size> bytes at device <BDF>` messages.

2. Software/Driver Loading

The driver must be loaded by the root user. Any attempt to load the driver as a regular user will fail.

To load the driver, execute the following:

```
[root@host]# modprobe cxgb4
```

III. iSCSI PDU Offload Initiator

1. Introduction

The Chelsio T4 series Adapters 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 transmitting, Chelsio h/w computes and inserts the Header and Data digest into the PDUs. On receiving, Chelsio h/w computes and verifies the Header and Data digest of the PDUs.

- **Direct Data Placement (DDP)**
  Chelsio h/w can directly place the iSCSI Data-In or Data-Out PDU's payload into pre-posted final 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 transmitting, Chelsio h/w 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 receiving, Chelsio h/w 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 payload data will be sent to the host too.

The cxgb4i driver interfaces with open-iSCSI initiator and provides the iSCSI acceleration through Chelsio hardware wherever applicable.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio Adapters that are compatible with iSCSI PDU Offload Initiator Software:

- T420-CR
- T420-LL-CR
- T440-CR
- T440-LP-CR
- T420-BCH
- T422-CR
- T420-CX

- T420-BT
- T404-BT

1.2. Software Requirements

1.2.1. XenSever Requirements

Currently the iSCSI PDU Offload Initiator software is available for the following version:

- Citrix XenServer Host 6.1.0
  Kernel version: 2.6.32.43-0.4.1.xs1.6.10.734.170748xen

Other versions have not been tested and are not guaranteed to work.

2. Software/Driver Loading

The driver must be loaded by the root user. Any attempt to load the driver as a regular user will fail.

Run the following command to load the driver:

```
[root@host]# modprobe cxgb4i
```

Note *If loading of driver fails, then kill iscsid, unload all open iSCSI modules and related/dependent modules (like multipath, bnx2i etc). Finally, load the cxgb4i driver.*

3. Software/Driver Unloading

To unload the driver, execute the following commands:

```
[root@host]# rmmod cxgb4i
[root@host]# rmmod libcxgbi
```

4. Software/Driver Configuration and Fine-tuning

4.1. Accelerating open-iSCSI Initiator

The following steps need to be taken to accelerate the open-iSCSI initiator:

4.1.1. Configuring iscsid.conf file

Edit the `iscsi/iscsid.conf` file and change the setting for `MaxRecvDataSegmentLength`:

```
node.conn[0].iscsi.MaxRecvDataSegmentLength = 8192
```

The login would fail for a normal session if `MaxRecvDataSegmentLength` is too big. A error message in the format of `ERR! MaxRecvSegmentLength <X> too big. Need to be <= <Y>.` would be logged to dmesg.

Important

Always take a backup of `iscsid.conf` file before installing Chelsio Unified Wire Package.

4.1.2. Configuring interface (*iface*) file

Create an interface file located under *iface* directory for the new transport class *cxgb4i* in the following format:

iface.iscsi_ifacename = <iface file name>
iface.hwaddress = <MAC address>
iface.transport_name = cxgb4i
iface.net_ifacename = <ethX>
iface.ipaddress = <iscsi ip address>

E.g.:-

```
iface.iscsi_ifacename = cxgb4i.00:07:43:04:5b:da
iface.hwaddress = 00:07:43:04:5b:da
iface.transport_name = cxgb4i
iface.net_ifacename = eth3
iface.ipaddress = 102.2.2.137
```

Alternatively, you can create the file automatically by executing the following command:

```
[root@host]# iscsiadm -m iface
```

Here,

- `iface.iscsi_ifacename` denotes the name of interface file in `/etc/iscsi/ifaces/`.
- `iface.hwaddress` denotes the MAC address of the Chelsio interface via which iSCSI traffic will be running.
- `iface.transport_name` denotes the transport name, which is cxgb4i.
- `iface.net_ifacename` denotes the Chelsio interface via which iSCSI traffic will be running.
- `iface.ipaddress` denotes the IP address which is assigned to the interface.

Note

i. The interface file needs to be created in `/etc/iscsi/iscsid.conf`.

ii. If `iface.ipaddress` is specified, it needs to be either the same as the ethX's IP address or an address on the same subnet. Make sure the IP address is unique in the network.

4.1.3. Discovery and Login

i. Starting iSCSI Daemon

Start Daemon from */sbin* by using the following command:

```
[root@host]# iscsid -f -d 3
```

Note *if `iscsid` is already running, then kill the service and start it as shown above after installing the Chelsio Unified Wire package.*

ii. Discovering iSCSI Targets

To discovery an iSCSI target execute a command in the following format:

iscsiadm -m discovery -t st -p <target ip address>:<target port no> -I <cxgb4i iface file name>

E.g.:-

```
[root@host]# iscsiadm -m discovery -t st -p 102.2.2.155:3260 -I
cxgb4i.00:07:43:04:5b:da
```

iii. Logging into an iSCSI Target

Log into an iSCSI target using the following format:

iscsiadm -m node -T <iqn name of target> -p <target ip address>:<target port no> -I <cxgb4i iface file name> -l

E.g.:-

```
[root@host]# iscsiadm -m node -T iqn.2004-05.com.chelsio.target1 -p
102.2.2.155:3260,1 -I cxgb4i.00:07:43:04:5b:da -l
```

iv. Logging out from an iSCSI Target

Log out from an iSCSI Target by executing a command in the following format:

iscsiadm -m node -T <iqn name of target> -p <target ip address>:<target port no> -I <cxgb4i iface file name> -u

E.g.:-

```
[root@host]# iscsiadm -m node -T iqn.2004-05.com.chelsio.target1 -p
102.2.2.155:3260,1 -I cxgb4i.00:07:43:04:5b:da -u
```

Note *Other options can be found by typing* `iscsiadm --help`

4.2. Auto login from cxgb4i initiator at OS bootup

For iSCSI auto login (via cxgb4i) to work on OS startup, please add the following line to `start()` in `/etc/rc.d/init.d/iscsid` file on RHEL:

```
modprobe -q cxgb4i
```

E.g.:-

```
force_start() {
    echo -n $"Starting $prog: "
    modprobe -q iscsi_tcpmodprobe -q ib_iser
    modprobe -q cxgb4i
    modprobe -q cxgb3i
    modprobe -q bnx2i
    modprobe -q be2iscsi
    daemon brcm_iscsiuio
    daemon $prog
    retval=$?
    echo
    [ $retval -eq 0 ] && touch $lockfile
    return $retval
}
```

IV. FCoE Full Offload Initiator

1. Introduction

Fibre Channel over Ethernet (FCoE) is a mapping of Fibre Channel over selected full duplex IEEE 802.3 networks. The goal is to provide I/O consolidation over Ethernet, reducing network complexity in the Datacenter.

Chelsio FCoE initiator maps Fibre Channel directly over Ethernet while being independent of the Ethernet forwarding scheme. The FCoE protocol specification replaces the FC0 and FC1 layers of the Fibre Channel stack with Ethernet. By retaining the native Fibre Channel constructs, FCoE will integrate with existing Fibre Channel networks and management software.

Chelsio's FCoE full offload Initiator provides enhanced SAN Boot features. The capability of booting from external storage resources compared to the traditional internal devices provides better system management, reduced downtime during system failure, enhanced performance, high availability, etc.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio Adapters that are compatible with FCoE full offload Initiator driver:

- T420-CR
- T420-LL-CR
- T440-CR
- T422-CR
- T440-LP-CR
- T404-BT

1.2. Software Requirements

1.2.1. XenServer Requirements

The Chelsio FCoE Full Offload Initiator driver has been developed to run on XenServer platforms. Currently the driver is available for the following version:

- Citrix XenServer Host 6.1.0
  Kernel version: 2.6.32.43-0.4.1.xs1.6.10.734.170748xen

Other versions have not been tested and are not guaranteed to work.

2. Software/Driver Loading

The driver must be loaded by the root user. Any attempt to load the driver as a regular user will fail.

To load the driver, execute the following:

```
[root@host]# modprobe csiostor
```

3. Software/Driver Configuration and Fine-tuning

3.1. Configuring Cisco Nexus 5010 switch

3.1.1. Configuring the DCBX parameters

Note *By default the Cisco Nexus switch enables DCBX functionality and configures PFC for FCoE traffic making it no drop with bandwidth of 50% assigned to FCoE class of traffic and another 50% for the rest(like NIC). If you wish to configure custom bandwidth, then follow the procedure below.*

In this procedure, you may need to adjust some of the parameters to suit your environment, such as VLAN IDs, Ethernet interfaces, and virtual Fibre Channel interfaces.

To enable PFC, ETS, and DCB functions on a Cisco Nexus 5000 series switch:

i. Open a terminal configuration setting.

```
switch# config terminal
switch(config)#
```

ii. Configure qos class-maps and set the traffic priorities: NIC uses priority 0 and FcoE uses priority 3.

```
switch(config)#class-map type qos class-nic
switch(config-cmap-qos)# match cos 0
switch(config-cmap-qos)# class-map type qos class-fcoe
switch(config-cmap-qos)# match cos 3
```

iii. Configure queuing class-maps.

```
switch(config)#class-map type queuing class-nic
switch(config-cmap-que)#match qos-group 2
```

iv. Configure network-qos class-maps.

```
switch(config)#class-map type network-qos class-nic
switch(config-cmap-nq)#match qos-group
```

v. Configure qos policy-maps.

```
switch(config)#policy-map type qos policy-test
switch(config-pmap-qos)#class type qos class-nic
switch(config-pmap-c-qos)#set qos-group 2
```

vi. Configure queuing policy-maps and assign network bandwidth. Divide the network bandwidth between FcoE and NIC traffic.

```
switch(config)#policy-map type queuing policy-test
switch(config-pmap-que)#class type queuing class-nic
switch(config-pmap-c-que)#bandwidth percent 50
switch(config-pmap-c-que)#class type queuing class-fcoe
switch(config-pmap-c-que)#bandwidth percent 50
switch(config-pmap-c-que)#class type queuing class-default
switch(config-pmap-c-que)#bandwidth percent 0
```

vii. Configure network-qos policy maps and set up the PFC for no-drop traffic class.

```
switch(config)#policy-map type network-qos policy-test
switch (config-pmap-nq)#class type network-qos class-nic
switch(config-pmap-nq-c)#pause no-drop
```

Note *By default FCoE is set to pause no drop. In such a trade off, one may want to set NIC to drop instead.*

viii. Apply the new policy (PFC on NIC and FcoE traffic) to the entire system.

```
switch(config)#system qos
switch(config-sys-qos)#service-policy type qos input policy-test
switch(config-sys-qos)#service-policy type queuing output policy-test
switch(config-sys-qos)#service-policy type queuing input policy-test
switch(config-sys-qos)#service-policy type network-qos policy-test
```

3.1.2. Configuring the FCoE/FC ports

In this procedure, you may need to adjust some of the parameters to suit your environment, such as VLAN IDs, Ethernet interfaces, and virtual Fibre Channel interfaces

i. Following steps will enable FCoE services on a particular VLAN and does a VSAN-VLAN mapping. Need not do these steps every time, unless a new mapping has to be created.

```
switch(config)# vlan 2
switch(config-vlan)# fcoe vsan 2
switch(config-vlan)#exit
```

ii. Following steps help in creating a virtual fibre channel (VFC) and binds that VFC to a Ethernet interface so that the Ethernet port begins functioning as a FCoE port.

```
switch(config)# interface vfc 13
switch(config-if)#bind interface ethernet 1/13
switch(config-if)# no shutdown
switch(config-if)#exit
switch(config)#vsan database
switch(config-vsan-db)# vsan 2
switch(config-vsan-db)#vsan 2 interface vfc 13
switch(config-vsan-db)#exit
```

Note *If you are binding the VFC to a MAC address instead of an ethernet port then make sure the ethernet port is part of both default VLAN and FCoE VLAN.*

iii. Assign VLAN ID to the Ethernet port on which FCoE service was enabled in step1.

```
switch(config)# interface ethernet 1/13
switch(config-if)#switchport mode trunk
switch(config-if)#switchport trunk allowed vlan 2
switch(config-if)#no shutdown
switch(config)#exit
```

iv. Enabling DCBX:

```
switch(config)# interface ethernet 1/13
switch(config-if)# priority-flow-control mode auto
switch(config-if)# flowcontrol send off
switch(config-if)# flowcontrol receive off
switch(config-if)# lldp transmit
switch(config-if)# lldp receive
switch(config-if)# no shutdown
```

v. On the FC Ports, if a FC target is connected then perform the following steps -

```
switch(config)#vsan database
switch(config-vsan-db)#vsan 2
switch(config-vsan-db)# vsan 2 interface fc 2/2
switch(config-vsan-db)#exit
switch(config)interface fc 2/2

switch(config-if)# switchport mode auto
switch(config-if)# switchport speed auto
switch(config-if)# no shutdown.
```

vi. If you have not created a zone then make sure the default-zone permits the VSAN created, otherwise the initiator and the target on that particular VSAN although FLOGI'd into the switch will not talk to each other. To enable it, execute the below command.

```
switch(config)# zone default-zone permit vsan 2
```

3.1.3. Configuring the Brocade 8000 switch

i. Configure LLDP for FCoE.Example of configuring LLDP for 10-Gigabit Ethernet interface.

```
switch(config)#protocol lldp
switch(conf-lldp)#advertise dcbx-fcoe-app-tlv
switch(conf-lldp)#advertise dcbx-fcoe-logical-link-tlv
```

ii. Create a CEE Map to carry LAN and SAN traffic if it does not exist. Example of creating a CEE map.

```
switch(config)# cee-map default
switch(conf-cee-map)#priority-group-table 1 weight 40 pfc
switch(conf-cee-map)#priority-group-table 2 weight 60
switch(conf-cee-map)#priority-table 2 2 2 1 2 2 2 2
```

iii. Configure the CEE interface as a Layer 2 switch port. Example of configuring the switch port as a 10-Gigabit Ethernet interface.

```
switch(config)#interface tengigabitethernet 0/16
switch(config-if-te-0/16)#switchport
switch(config-if-te-0/16)#no shutdown
switch(config-if)#exit
```

iv. Create an FCoE VLAN and add an interface to it. Example of creating a FCoE VLAN and adding a single interface.

```
switch(config)#vlan classifier rule 1 proto fcoe encap ethv2
switch(config)#vlan classifier rule 2 proto fip encap ethv2
switch(config)#vlan classifier group 1 add rule 1
switch(config)#vlan classifier group 1 add rule 2
switch(config)#interface vlan 1002
switch(conf-if-vl-1002 )#fcf forward
switch(conf-if-vl-1002 )#interface tengigabitethernet 0/16
switch(config-if-te-0/16)#switchport
switch(config-if-te-0/16)#switchport mode converged
switch(config-if-te-0/16)#switchport converged allowed vlan add 1002
switch(config-if-te-0/16)#vlan classifier activate group 1 vlan 1002
switch(config-if-te-0/16)#cee default
switch(config-if-te-0/16)#no shutdown
switch(config-if-te-0/16)#exit
```

Note *Unlike cisco, only one VLAN ID can carry FCoE traffic for now on Brocade 8000. It is their limitation.*

v. Save the Configuration

```
switch#copy running-config startup-config
```

3.2. FCoE fabric discovery verification

The verification is done using *cxgbtool*.

3.2.1. Verifying the DCBX parameters

To verify the current DCBX information being exchanged, execute the below commands using cxgbtool.

i. First step is to find the adapter number. Find it using the following command

```
[root@host]# cxgbtool stor -s
```

```
[root@sparks ~]# cxgbtool stor -s
Chelsio Adapter device file: /dev/csiostor0
[root@sparks ~]#
```

ii. Now execute the following command to check the DCBX information.

```
[root@host]# cxgbtool stor -a <adapter_no> --dcb-params
```

```
[root@sparks ~]# cxgbtool stor -a /dev/csiostor0 --dcb-params
******************** DCBX Paramters[Port:  0] ********************
Priority Group ID of Priority 0          : 0
Priority Group ID of Priority 1          : 0
Priority Group ID of Priority 2          : 0
Priority Group ID of Priority 3          : 0
Priority Group ID of Priority 4          : 0
Priority Group ID of Priority 5          : 0
Priority Group ID of Priority 6          : 0
Priority Group ID of Priority 7          : 0

Bandwidth Percentage :
---------------------
Bandwidth Percentage of Priority Group 0: 0
Bandwidth Percentage of Priority Group 1: 0
Bandwidth Percentage of Priority Group 2: 0
Bandwidth Percentage of Priority Group 3: 0
Bandwidth Percentage of Priority Group 4: 0
Bandwidth Percentage of Priority Group 5: 0
Bandwidth Percentage of Priority Group 6: 0
Bandwidth Percentage of Priority Group 7: 0


Number of Traffic Classes Supported     : 0

Strict Priorate :
-----------------
Strict Priorate for Priority 0           : 0
Strict Priorate for Priority 1           : 0
Strict Priorate for Priority 2           : 0
Strict Priorate for Priority 3           : 0
Strict Priorate for Priority 4           : 0
Strict Priorate for Priority 5           : 0
Strict Priorate for Priority 6           : 0
Strict Priorate for Priority 7           : 0

PFC Enabled/Disabled :
----------------------
PFC for Priotity 0                       : Disabled
PFC for Priotity 1                       : Disabled
PFC for Priotity 2                       : Disabled
PFC for Priotity 3                       : Disabled
PFC for Priotity 4                       : Disabled
PFC for Priotity 5                       : Disabled
PFC for Priotity 6                       : Disabled
PFC for Priotity 7                       : Disabled

User priority map                        : 0x0

Selector field                           : 0

Application Protocol ID                  : 0x0
```

3.2.2. Verifying Local Ports

Once connected to the switch, use the following command to see if the FIP has gone through and a VN_Port MAC address has been assigned.

Verify if all the FCoE ports are online/ready and a successful FIP has taken place using the following command. The **wwpn** and **state** of the initiator local port can be found under sysfs.

```
[root@host]# cat /sys/class/fc_host/hostX/port_name
```

```
[root@sparks ~]# cat /sys/class/fc_host/host3/port_name
0x50007431027e6080
[root@sparks ~]# cat /sys/class/fc_host/host3/port_state
Online
[root@sparks ~]#
```

Note

- *The hosts under fc_host depends on the number of ports on the adapter used.*
- *Inorder to identify chelsio fc_host from other vendor fc_host, the WWPN always begins with **0x5000743***

Alternatively, the local port information can also be found using:

```
[root@host]# cxgbtool stor -a <adapter_no> --show-lnode
```

```
[root@sparks ~]# cxgbtool stor -a /dev/csiostor0 --show-lnode
*****************[Index:  0]***************
LNode Device ID: 235524
VNPI    : 0x9804
FCFI    : 0x9820
MAC     : 0E-FC-00-48-01-0F

Port Id : 0
Nport id: 48010f
State   : READY

WWPN    : 50:00:74:31:02:7e:60:80
WWNN    : 50:00:74:31:02:7e:60:00
NPIV    : SUPPORTED
Total VPorts :  0

No.of RNodes :  8

Common Service Params:
        Rcv size: 2068
        ED-TOV  : 2000

Class Service Params:
Class 1: NOT SUPPORTED
Class 2: NOT SUPPORTED
Class 3: SUPPORTED
        Initiator ctl           : 0
        Recipient ctl           : 0
        Rcv size                : 2068
        Total concurrent seq    : 0
        EE Credit               : 0
        Open Sequence per Exchange: 0

Class 4: NOT SUPPORTED
*****************[Index:  1]***************
LNode Device ID: 301061
VNPI    : 0x9805
FCFI    : 0x9821
MAC     : 0E-FC-00-48-01-0E

Port Id : 1
Nport id: 48010e
State   : READY

WWPN    : 50:00:74:31:02:7e:e1:80
WWNN    : 50:00:74:31:02:7e:e1:00
NPIV    : SUPPORTED
Total VPorts :  0

No.of RNodes :  8

Common Service Params:
```

3.2.3. Verifying the target discovery

To check the targets being discovered use *cxgbtool*. To check the list of targets that are being discovered from a particular FCoE port, use the following commands from *cxgbtool*.

i. Check for the adapter number using the following command.

```
[root@host]# cxgbtool stor -s
```

ii. To check the list of targets discovered from a particular FCoE port, first find out the wwpn of the initiator local port under sysfs. The hosts under fc_host depends on the number of ports on the adapter used.

```
[root@host]# cat /sys/class/fc_host/hostX/port_name
```

iii. After finding the localport, go to the corresponding Remote port under sysfs # **cat /sys/class/fc_remote_ports/rport-X:B:R** where X is the Host ID, B is the bus ID and R is the remote port.

```
[root@sparks ~]# cat /sys/class/fc_remote_ports/rport-3:0-0/roles
Fabric Port
[root@sparks ~]# cat /sys/class/fc_remote_ports/rport-3:0-1/roles
Directory Server
[root@sparks ~]# cat /sys/class/fc_remote_ports/rport-3:0-2/roles
Management Server
[root@sparks ~]# cat /sys/class/fc_remote_ports/rport-3:0-3/roles
FCP Initiator
[root@sparks ~]# cat /sys/class/fc_remote_ports/rport-3:0-6/roles
FCP Target
[root@sparks ~]#
```

Note *R can correspond to NameServer, Management Server and other initiator ports logged in to the switch and targets.*

Alternatively the locaports can also be found using cxgbtool

```
[root@host]# cxgbtool stor -a  <adapter no>   --show-lnode
```

After finding out the wwpn of the local node, to verify the list of targets being discovered, use the following command.

```
[root@host]# cxgbtool stor -a <adapter_no> --show-rnode --wwn=<wwpn of
lnode>
```

```
[root@sparks ~]# cxgbtool stor -a /dev/csiostor0 --show-rnode --wwn=50:00:74:31:02:7e:60:80

****************[Index:   0]***************
SSNI            : 0x9400
VNPI            : 0x9804
FCFI            : 0x9820
WWPN            : 20:11:00:0d:ec:b1:bd:7f
WWNN            : 20:0a:00:0d:ec:b1:bd:41
Nport id        : FFFFFE
State           : READY
FCP Flags       : 0
Role            : Fabric

Class Service Params:
Class 1: NOT SUPPORTED
Class 2: NOT SUPPORTED
Class 3: SUPPORTED
Class 4: NOT SUPPORTED

****************[Index:   1]***************
SSNI            : 0x9401
VNPI            : 0x9804
FCFI            : 0x9820
WWPN            : 25:0d:00:0d:ec:b1:bd:40
WWNN            : 20:0a:00:0d:ec:b1:bd:41
Nport id        : FFFFFC
State           : READY
FCP Flags       : 0
Role            : Name-Server

Class Service Params:
Class 1: NOT SUPPORTED
Class 2: NOT SUPPORTED
Class 3: SUPPORTED
Class 4: NOT SUPPORTED

****************[Index:   2]***************
SSNI            : 0x9403
VNPI            : 0x9804
FCFI            : 0x9820
WWPN            : 00:00:00:00:00:00:00:00
WWNN            : 20:0a:00:0d:ec:b1:bd:41
Nport id        : FFFFFA
State           : READY
FCP Flags       : 0
Role            : N-Port

Class Service Params:
Class 1: NOT SUPPORTED
Class 2: NOT SUPPORTED
Class 3: SUPPORTED
Class 4: NOT SUPPORTED
```

3.3. Formatting the LUNs and Mounting the Filesystem

Use *fdisk -l* to list the LUNs discovered by the initiator

```
[root@host]# fdisk -l
```

```
[root@sparks ~]# fdisk -l

Disk /dev/sdb: 499.5 GB, 499558383616 bytes
255 heads, 63 sectors/track, 60734 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

Disk /dev/sdb doesn't contain a valid partition table

WARNING: GPT (GUID Partition Table) detected on '/dev/sda'! The util fdisk doesn't support GPT. Use GNU Parted.


Disk /dev/sda: 499.5 GB, 499558383616 bytes
256 heads, 63 sectors/track, 60497 cylinders
Units = cylinders of 16128 * 512 = 8257536 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1       60498   487849983+  ee  EFI GPT

Disk /dev/sdd: 4294 MB, 4294967296 bytes
133 heads, 62 sectors/track, 1017 cylinders
Units = cylinders of 8246 * 512 = 4221952 bytes

Disk /dev/sdd doesn't contain a valid partition table

Disk /dev/sde: 4294 MB, 4294967296 bytes
133 heads, 62 sectors/track, 1017 cylinders
Units = cylinders of 8246 * 512 = 4221952 bytes

Disk /dev/sde doesn't contain a valid partition table

Disk /dev/sdf: 4294 MB, 4294967296 bytes
133 heads, 62 sectors/track, 1017 cylinders
Units = cylinders of 8246 * 512 = 4221952 bytes

Disk /dev/sdf doesn't contain a valid partition table
```

In the above example, *sdd*, *sde* and *sdf* are the discovered LUNs

3.4. Creating Filesystem

Create an ext3 filesystem using the following command:

```
[root@host]# mkfs.ext3 /dev/sdx
```

```
[root@sparks ~]# mkfs.ext3 /dev/sdd
mke2fs 1.39 (29-May-2006)
/dev/sdd is entire device, not just one partition!
Proceed anyway? (y,n) y
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
524288 inodes, 1048576 blocks
52428 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=1073741824
32 block groups
32768 blocks per group, 32768 fragments per group
16384 inodes per group
Superblock backups stored on blocks:
        32768, 98304, 163840, 229376, 294912, 819200, 884736

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 37 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
[root@sparks ~]#
```

3.5. Mounting the formatted LUN

The formatted LUN can be mounted on the specified mountpoint using the following command:

```
[root@host]# mount /dev/sdx /mnt
```

```
[root@sparks ~]# mount /dev/sdd /mnt/
[root@sparks ~]# mount
/dev/sda1 on / type ext3 (rw)
none on /proc type proc (rw)
none on /sys type sysfs (rw)
none on /dev/pts type devpts (rw)
none on /dev/shm type tmpfs (rw)
/opt/xensource/packages/iso/XenCenter.iso on /var/xen/xc-install type iso9660 (ro,loop=/dev/loop0)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
/root/Drivers/cxgb4-6.1.0-59235p-1.2.4-8.iso on /mnt type iso9660 (ro,loop=/dev/loop1)
/dev/sdd on /mnt type ext3 (rw)
[root@sparks ~]#
```

Note

To create storage repository (SR) using XenCenter, open */opt/xensource/sm/devscan.py* file and add the following line in the *MODULE_INFO* section:

```
'csiostor': 'chelsio HBA Driver',
```

E.g.:

```
MODULE_INFO = {
  'csiostor': 'chelsio HBA Driver',
```

V. Appendix

1. Troubleshooting

- ***Configuring firewall for your application***

In many cases the firewall software on the systems may prevent the applications from working properly. Please refer to the appropriate documentation for the Linux distribution on how to configure or disable the firewall.

- ***FCoE link not up***

Always enable LLDP on the interfaces as FCoE link won't come up until and unless a successful LLDP negotiation happens.

- ***priority-flow-control mode on the switch***

On the switch, make sure priority-flow-control mode is always set to auto and flow control is disabled.

- ***Configuring Ethernet interfaces on Cisco switch***

Always configure Ethernet interfaces on Cisco switch in trunk mode.

- ***Binding VFC to MAC***

If you are binding the VFC to MAC address in case of Cisco Nexus switch, then make sure you make the Ethernet interface part of both Ethernet VLAN and FCoE VLAN.

- ***Cisco nexus switch reporting "pauseRateLimitErrDisable"***

If in any case the switch-port on the Cisco nexus switch is reporting "pauseRateLimitErrDisable", then perform an Ethernet port shut/no shut.

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, nontransferable license to:

(i) install and use the Software on a single computer system or on multiple workstations, systems and servers that incorporate a Chelsio network adapter and may be accessed by multiple users from multiple locations; and (ii) 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, as 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 for profit, 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 recorded on the original media 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. 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 not meeting the above limited warranty 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.

Chelsio reserves the right to modify this license agreement at any time without notice, and any modified version of this agreement shall supercede any earlier versions.

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

Chelsio Communications, Inc.
370 San Aleso Ave.
Sunnyvale, CA 94085

