



Chelsio Unified Wire for FreeBSD

Installation and User's Guide



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

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

All third party trademarks are copyright of their respective owners.

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

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



Chelsio Communications (Headquarters)

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

www.chelsio.com

Tel: 408.962.3600
Fax: 408.962.3661

Chelsio (India) Private Limited

Subramanya Arcade, Floor 3, Tower B
No. 12, Bannerghatta Road,
Bangalore-560029
Karnataka,
India

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

Chelsio KK (Japan)

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

Sales

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

Support

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

Copyright © 2017. 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	03/04/2013
1.0.1	08/29/2013
1.0.2	11/11/2016
1.0.3	12/05/2016
1.0.4	12/28/2016
1.0.5	01/27/2017
1.0.6	02/24/2017
1.0.7	05/10/2017
1.0.8	05/23/2017

TABLE OF CONTENTS

I.	CHELSIO UNIFIED WIRE	5
1.	Introduction	6
1.1.	Features	6
1.2.	Hardware Requirements	6
1.3.	Software Requirements	7
1.4.	Package Contents	7
2.	Hardware Installation	8
3.	Software/Driver Installation	11
3.1.	Compiling and Installing the kernel	11
3.2.	Userland Tools	12
3.3.	Firmware Update	13
4.	Software/Driver Update	14
II.	NETWORK (NIC-TOE)	15
1.	Introduction	16
1.1.	Hardware Requirements	16
1.2.	Software Requirements	17
2.	Software/Driver Loading	18
2.1.	Loading NIC driver	18
2.2.	Loading TOE driver	18
3.	Software/Driver Configuration and Fine-tuning	19
3.1.	Enabling network interface	19
3.2.	Enabling TCP Offload	19
3.3.	Enabling DDP	20
3.4.	Connection Offload Policies (COP)	20
3.5.	Network Device Configuration	21
4.	Software/Driver Unloading	22
4.1.	Unloading the NIC driver module	22
4.2.	Unloading the TOE driver module	22
III.	IWARP	23
1.	Introduction	24
1.1.	Hardware Requirements	24
1.2.	Software Requirements	24
2.	Software/Driver Loading	25
3.	Software/Driver Configuration and Fine-tuning	26
3.1.	Verifying IWARP Functionality	26
4.	Software/Driver Unloading	27
IV.	ISCSI PDU OFFLOAD TARGET	28
1.	Introduction	29
1.1.	Hardware Requirements	29
1.2.	Software Requirements	30
2.	Software/Driver Loading	31

2.1.	Module Installation	31
2.2.	Module Loading	31
3.	Software/Driver Configuration and Fine-tuning	33
3.1.	Starting iSCSI Target	33
4.	Software/Driver Unloading	34
V.	CLASSIFICATION AND FILTERING	35
1.	Introduction	36
1.1.	Hardware Requirements	36
1.2.	Software Requirements	37
2.	Usage	38
2.1.	Configuration	38
2.2.	Creating Filter Rules	38
2.3.	Listing Filter Rules	40
2.4.	Removing Filter Rules	40
VI.	OFFLOAD IPV6	41
1.	Introduction	42
1.1.	Hardware Requirements	42
1.2.	Software Requirements	42
2.	Software/Driver Configuration and Fine-tuning	43
VII.	APPENDIX	44
	Chelsio End-User License Agreement (EULA)	45

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 data centers with high performance solutions for various Network and Storage related requirements.

The **Terminator** series is Chelsio's next generation of highly integrated, hyper-virtualized 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, the 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 adapters bring a new level of performance metrics and functional capabilities to the computer networking industry.

1.1. Features

Chelsio Unified Wire driver package for FreeBSD contains Network, iSCSI and iWARP drivers for Chelsio's Unified Wire adapters.

The Unified Wire Installer installs the following Chelsio drivers, features and software:


- **Network driver (NIC-TOE)**
- **iWARP**
- **iSCSI PDU Offload Target**
- **Classification and Filtering**
- **Offload IPv6**

Please refer to their respective sections for more information.

1.2. Hardware Requirements

The Chelsio Unified Wire supports all x64 architectures supporting PCIE (x4, x8, x16) slots.

- AMD CPUs, 64-bit (x86_64/amd64)
- Intel CPUs, 64-bit (x86_64)

 **Note** *The Chelsio Unified Wire supports 3.3v PCI bus only. Running an adapter on a PCI x4 slot is not recommended as performance will be significantly reduced by the limitations of PCI.*

1.3. Software Requirements

The Chelsio Unified Wire software has been developed to run on FreeBSD based platforms. To know more about the complete list of versions supported by each driver, please refer to their respective sections.

1.4. Package Contents

The Chelsio Unified Wire Source Package consists of the following files/directories:

- **kernel-head:** Source code for different drivers/software.
- **docs:** The *docs* directory contains support documents - README, Release Notes and User's Guide (this document) for the driver.
- **firmware:** Firmware binaries and firmware configuration files.
- **EULA:** Chelsio's End User License Agreement.

2. Hardware Installation

1. Shutdown/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 Ethernet adapter.
6. For maximum performance, it is highly recommended to install the adapter into a PCIe x8/x16 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 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/x16 slot as sometimes the process of securing the card causes the card to become unseated.
10. 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.
11. Power on your system.
12. Verify if the adapter was installed successfully by using the following command:
 - a. For T6 adapters:

```

t6iov0@pci0:129:0:0:    class=0x020000 card=0x00001425 chip=0x60071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t6iov1@pci0:129:0:1:    class=0x020000 card=0x00001425 chip=0x60071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t6iov2@pci0:129:0:2:    class=0x020000 card=0x00001425 chip=0x60071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t6iov3@pci0:129:0:3:    class=0x020000 card=0x00001425 chip=0x60071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t6nex0@pci0:129:0:4:    class=0x020000 card=0x00001425 chip=0x64071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
none150@pci0:129:0:5:   class=0x010000 card=0x00001425 chip=0x65071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Storage Controller'
    class       = mass storage
    subclass    = SCSI
none151@pci0:129:0:6:   class=0x0c0400 card=0x00001425 chip=0x66071425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T62100-LP-CR Unified Wire Storage Controller'
    class       = serial bus
    subclass    = Fibre Channel

```

b. For T5 adapters:

```

t5iov0@pci0:3:0:0:      class=0x020000 card=0x00001425 chip=0x50011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t5iov1@pci0:3:0:1:      class=0x020000 card=0x00001425 chip=0x50011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t5iov2@pci0:3:0:2:      class=0x020000 card=0x00001425 chip=0x50011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t5iov3@pci0:3:0:3:      class=0x020000 card=0x00001425 chip=0x50011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
t5nex0@pci0:3:0:4:      class=0x020000 card=0x00001425 chip=0x54011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Ethernet Controller'
    class       = network
    subclass    = ethernet
none65@pci0:3:0:5:      class=0x010000 card=0x00001425 chip=0x55011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Storage Controller'
    class       = mass storage
    subclass    = SCSI
none66@pci0:3:0:6:      class=0x0c0400 card=0x00001425 chip=0x56011425 rev=0x00 hdr=0x00
    vendor      = 'Chelsio Communications Inc'
    device      = 'T520-CR Unified Wire Storage Controller'
    class       = serial bus
    subclass    = Fibre Channel

```

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

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

Once driver is installed and loaded, run the following commands and examine the output to see if the card is discovered.

```

[root@host~]# sysctl dev.t5nex.0.%desc
dev.t5nex.0.%desc: Chelsio T520-LL RNIC (rev 0), S/N:RE16130313,
E/C:0000000000000000

[root@host~]# dmesg
t5nex0: <Chelsio T520-LL-CR> mem 0xe5300000-0xe537ffff,0xe4000000-
0xe4ffffff,0xe5884000-0xe5885fff irq 16 at device 0.4 on pci1
t5nex0: PCIe x8, 2 ports, 14 MSI-X interrupts, 31 eq, 13 iq

```

The above outputs indicate 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.



Note

Network device names for Chelsio's physical ports are assigned using the following convention: the port farthest from the motherboard will appear as the first Ethernet network interface. However, for T5 40G and T420-BT adapters, the association of physical Ethernet ports and their corresponding network device names is opposite. For these adapters, the port nearest to the motherboard will appear as the first network interface.

3. Software/Driver Installation

3.1. Compiling and Installing the kernel

1. Install FreeBSD 10.X/11.X from <http://ftp.freebsd.org/pub/FreeBSD/releases/amd64/amd64/ISO-IMAGES/>
2. If you haven't done so already, download the tar-ball *ChelsioUwire-FBSD-x.x.x.x.tar.gz* from Chelsio Download Center, <http://service.chelsio.com>
3. Untar the tar-ball using the following command:

```
[root@host~]# tar xvf ChelsioUwire-FBSD-x.x.x.x.tar.gz
```

4. This package has the entire FreeBSD world (operating system). You need to build entire FreeBSD world since infiniband core driver compiles only within FreeBSD kernel and not separately as a module.

Build the FreeBSD world (operating system) from the package as follows:

- a. Add the following entry in */etc/src.conf* file:

```
WITH_OFED='yes'
```

- b. Create a backup of existing kernel source.

```
[root@host~]# mv /usr/src /usr/src_bkp
```

- c. Copy the contents of *ChelsioUwire-FBSD-x.x.x.x/kernel-head* to */usr/src/*
- d. Build and install the world as follows:

```
[root@host~]# cd usr/src/  
[root@host~]# make buildworld  
[root@host~]# make installworld
```

- e. Build and install the kernel as follows:

```
[root@host~]# make buildkernel
[root@host~]# make installkernel
```

- f. Add the following entries to */boot/loader.conf* to enable capabilities:

```
hw.cxgbe.config_file="flash"
hw.cxgbe.iscsicaps_allowed=0xf
```

- g. Reboot the machine for changes to take effect.
h. After reboot, check if you have booted to the right kernel by executing `uname -a`.
You should see a similar output:

```
[root@kick ~]# uname -a
FreeBSD kick 12.0-CURRENT FreeBSD 12.0-CURRENT #0 f061800bf4e8 (freebsd-head) tip: Thu Nov 10 20:02:53
IST 2016      root@kick:/usr/obj/usr/src/sys/GENERIC  amd64
```

5. Verify that the kernel supports TOE:

```
[root@host~]# config -x /boot/kernel/kernel | grep TCP_OFFLOAD
options TCP_OFFLOAD
```

3.2. Userland Tools

Build and install *cxgbetool*

```
[root@host~]# cd /usr/src/usr.sbin/cxgbetool
[root@host~]# make && make install
```

3.3. Firmware Update

Firmware will have to be loaded manually. Run the following commands for T6 adapters:

```
[root@host~]# kldload if_cxgbe
[root@host~]# cxgbetool t6nexX loadfw ChelsioUwire-FBSD-
x.x.x.x/firmware/t6fw-x.xx.xx.x.bin
[root@host~]# cxgbetool t6nexX loadcfg ChelsioUwire-FBSD-
x.x.x.x/firmware/t6-config.txt
```

Reboot system for changes to take effect.

Firmware version can be verified using:

```
[root@host~]# sysctl -n dev.t6nex.X.firmware_version
```



Note

For T5 adapters, use **t5nex** as the parent device, **t5fw-x.xx.xx.x.bin** as the firmware binary file and **t5-config.txt** as the firmware configuration file. Similarly, for T4 adapters, use **t4nex**, **t4fw-x.xx.xx.x.bin** and **t4-config.txt** respectively.

4. 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-TOE)

1. Introduction

Chelsio's 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 aka TSO, and assist mechanisms for accelerating LRO. 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 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 Network (*cxgbe*) driver provides support for PCI Express Ethernet adapters based on the Chelsio Terminator ASIC.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the Chelsio adapters that are compatible with Chelsio 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
- T520-CR
- T520-LL-CR
- T520-SO-CR*
- T420-CR
- T420-LL-CR
- T440-CR
- T440-LP-CR
- T422-CR
- T420-SO-CR*
- T420-CX
- T420-BT
- T404-BT

*Only NIC supported

1.2. Software Requirements

Currently the Network driver is available for the following version(s):

- FreeBSD 10.3-RELEASE
- FreeBSD 11.0-RELEASE

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

2. Software/Driver Loading

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

2.1. Loading NIC driver


Load the NIC driver using the command given below:

```
[root@host~]# kldload if_cxgbe
```

2.2. Loading TOE driver

The TOE driver module (*t4_tom.ko*) is needed only if you plan to use the TOE and iWARP capabilities, and must be loaded manually. Run the following command to do so:

```
[[root@host~]# kldload t4_tom
```

 **Note** *The NIC module (*if_cxgbe*) by itself can be loaded / unloaded at any time, but once the TOE driver module (*t4_tom*) is loaded, it will prevent both the modules from being unloaded.*

3. Software/Driver Configuration and Fine-tuning

3.1. Enabling network interface

Enable any Chelsio T6 network interface using the following command:

```
[root@host~]# ifconfig ccX <ip_address> up
```

For T5 and T4, use *cxlX* and *cxgbeX* respectively.

3.2. Enabling TCP Offload

Enable TCP offload on a T6 interface using *ifconfig*. An error usually indicates that *t4_tom* is not loaded.

```
[root@host~]# ifconfig ccX toe
```

For T5 and T4, use *cxlX* and *cxgbeX* respectively.

If successful, TOE4 and TOE6 will be displayed in the list of the enabled options as shown in example below:

```
[root@ ~]# ifconfig cc0 toe
[root@ ~]# ifconfig cc0
cc0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1500
    options=ecc7bb<RXCSUM, TXCSUM, VLAN_MTU, VLAN_HWTAGGING, JUMBO_MTU, VLAN_HWCSUM, TSO4, TSO6, LRO, TOE4,
    TOE6, VLAN_HWTSO, LINKSTATE, RXCSUM_IPV6, TXCSUM_IPV6>
    ether 00:07:43:04:b2:30
    inet 102.1.1.112 netmask 0xffffffff broadcast 102.1.1.255
    nd6 options=29<PERFORMNUD, IFDISABLED, AUTO_LINKLOCAL>
    media: Ethernet 100GBase-CR4 <full-duplex>
    status: active
```

When a TOE capable interface is the outbound path to a destination to which a new connection is being established, the connection should be offloaded automatically. Verify with *netstat*:

```
[root@host~]# netstat -np tcp | grep toe
```

T6 SO (Server Offload) adapters support limited number of offload connections. To see the maximum number of offload connections, run the following command:

```
[root@host~]# sysctl dev.t6nex.X.misc.tids
```

```
root@ ~ # sysctl dev.t6nex.0.misc.tids
dev.t6nex.0.misc.tids: ATID range: 0-127, in use: 0
TID range: 0-255, in use: 0
STID range: 256-319, in use: 0
FTID range: 320-815
HW TID usage: 0 IP users, 0 IPv6 users
```

Where,

TID is the number of offload connections.

STID is the number of offload servers.

3.3. Enabling DDP

DDP is disabled by default. To enable it, run the *sysctl* command.

For T6 adapters:

```
[root@host~]# sysctl -w dev.t6nex.X.toe.ddp=1
```

For T5 and T4, use *t5nex* and *t4nex* respectively.

3.4. Connection Offload Policies (COP)

Connection Offload Policies (COP) are policies applicable to offloaded connections that are created through TOM driver. COP compiles offload policies into a simple program using BPF, and that generated byte code can be loaded into the kernel and interpreted. These offload policies are used to determine the settings to be used for various connections based on matching filter specifications with program rules.

The following section explains the method to create and apply COP.

1. If not done already, load NIC and TOE drivers:

```
[root@host~]# kldload if_cxgbe
[root@host~]# kldload t4_tom
```

2. For T6, bring up the interface and enable TOE:

```
[root@host~]# ifconfig ccX up
[root@host~]# ifconfig ccX toe
```

For T5 and T4, use *cx/X* and *cxgbeX* respectively.

3. Create the policy_file.

Example:


The below policy will offload only connections to port number 5001.

```
[L] port 5001 => offload
[P] port 5001 => offload
```

4. Apply the policy using the *cxgbetool*.

```
[root@host~]# cxgbetool t6nexY policy <policy_file>
```


For T5 and T4, use *t5nexY* and *t4nexY* respectively.



Note For more information on additional parameters and examples, refer *cxgbetool* manual by running the `man cxgbetool` command

3.5. Network Device Configuration

Please refer to the operating system documentation for administration and configuration of network devices.




Some operating systems may attempt to auto-configure the detected hardware and some may not detect all ports on a multi-port adapter. If this happens, please refer to the operating system documentation for manually configuring the network device.

4. Software/Driver Unloading

4.1. Unloading the NIC driver module

To unload the NIC driver, stop all the running traffic and run the following command:

```
[root@host~]# kldunload if_cxgbe
```

 **Note** *NIC module cannot be unloaded if TOE module is loaded.*

4.2. Unloading the TOE driver module

To unload the TOE driver, reboot the system.

III. iWARP

1. Introduction

The iWARP (*iw_cxgbe*) driver provides support for iWARP including RDMA, DDP and MPA protocols.

1.1. Hardware Requirements

1.1.1. Supported Adapters

Following are the currently shipping Chelsio adapters that are compatible with the iWARP driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T580-CR
- T580-LP-CR
- T540-CR
- T520-CR
- T520-LL-CR
- T420-CR
- T440-CR
- T422-CR
- T404-BT

1.2. Software Requirements

Currently the iWARP driver is available for the following version(s):

- FreeBSD 10.3-RELEASE
- FreeBSD 11.0-RELEASE

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.

1. If not done already, load NIC and TOE drivers:

```
[root@host~]# kldload if_cxgbe
[root@host~]# kldload t4_tom
```

2. For T6, bring up the interface and enable TOE:

```
[root@host~]# ifconfig ccX up
[root@host~]# ifconfig ccX toe
```



Note For T5 and T4, use **cxIX** and **cxgbeX** respectively.

3. Load the iWARP driver using the following command:

```
[root@host~]# kldload iw_cxgbe
```

3. Software/Driver Configuration and Fine-tuning

3.1. Verifying iWARP Functionality

1. Load *krping* module using the following command:

```
[root@host~]# kldload krping
```

2. Start *krping* server on one machine as:

```
[root@host~]# echo  
"verbose,count=5000,server,addr=<chelsioInterface>,port=<port>" >  
/dev/krping
```

3. Start *krping* client on PEER machine as:

```
[root@host~]# echo  
"verbose,count=5000,client,addr=<Server_chelsioInterface>,port=<port>" >  
/dev/krping
```

4. After starting the client, read the */dev/krping* on server as:

```
[root@host~]# cat /dev/krping  
  
krping:num device snd bytes snd msgs rcv bytes rcv msgs wr bytes wr msgs rd bytes rd msgs  
krping: 1 listen 0 0 16 1 60085 70624 917184 0
```

4. Software/Driver Unloading

To unload the iWARP driver, stop all the running traffic and run the following command:

```
[root@host~]# kldunload iw_cxgbe
```

IV. iSCSI PDU Offload Target

1. Introduction

Chelsio iSCSI PDU Offload Target (*cxgbei.ko*) is the hardware accelerated iSCSI target module for Chelsio TOE capable adapters supported by NIC driver (*if_cxgbe*) and provides support for offloading iSCSI traffic.

On the host side the driver interacts with the storage stack via the ICL (iSCSI Common Layer) in the kernel. On the wire, standard iSCSI traffic (SCSI over TCP as per RFC 3720/7143 etc.) interoperates with all other standards compliant implementations. The driver is layered on top of the TOE driver (*t4_tom*) and promotes connections being handled by *t4_tom* in iSCSI ULP (Upper Layer Protocol) mode. The driver provides hardware assistance including:

- Full TCP processing.
- iSCSI PDU identification and recovery within the TCP stream.
- Header and/or data digest insertion (Tx) and verification (Rx).
- Zero copy (both Tx and Rx).



This product will be supported in the next release.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio adapters that are compatible with iSCSI PDU Offload Target software:

- T580-CR
- T580-LP-CR
- T540-CR
- T520-CR
- T520-LL-CR
- T420-CR
- T420-LL-CR
- T440-CR
- T440-LP-CR
- T422-CR
- T420-CX
- T420-BT
- T404-BT

1.2. Software Requirements

`cxgbei.ko` is the iSCSI PDU offload target mode driver.

`if_cxgbe`, `toecore` and `t4_tom` modules are required by `cxgbei.ko` module to work.

1.2.1. FreeBSD Requirements

Currently the iSCSI PDU driver is available for the following version(s):

- FreeBSD 10.3-RELEASE
- FreeBSD 11.0-RELEASE

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

2. Software/Driver Loading

2.1. Module Installation

1. Compile and install iSCSI module

```
[root@host~]# cd /usr/src/sys/modules/iscsi
[root@host~]# make
[root@host~]# make install
```

2. Compile and install *ctl* drivers

```
[root@host~]# cd /usr/src/sys/modules/ctl
[root@host~]# make
[root@host~]# make install
```

3. Compile and install *ctld*

```
[root@host~]# cd /usr/src/usr.sbin/ctld
[root@host~]# make all install
```

4. Compile and install iSCSI PDU offload target module

```
[root@host~]# cd /usr/src/sys/modules/cxgbe/cxgbei
[root@host~]# make
[root@host~]# make install
```

2.2. Module Loading

1. If not done already, load NIC and TOE drivers:

```
[root@host~]# kldload if_cxgbe
[root@host~]# kldload t4_tom
```


2. For T6, bring up the interface and enable TOE:

```
[root@host~]# ifconfig ccX up  
[root@host~]# ifconfig ccX toe
```



Note For T5 and T4, use **cxIX** and **cxgbeX** respectively.

3. Load iSCSI PDU offload target module:

```
[root@host~]# kldload cxgbei
```

3. Software/Driver Configuration and Fine-tuning

3.1. Starting iSCSI Target

1. Make necessary changes in the iSCSI configuration file located in `/etc/ctl.conf`. In order to use `cxgbei` as offload module, add `offload cxgbei` to regular configuration.

Sample:

```
root@host:~ # vi /etc/ctl.conf

portal-group pg0 {
    discovery-auth-group no-authentication
    listen 102.11.11.102
    offload cxgbei
}

target ign.2017-05.com.chelsio.target:1 {
    auth-group no-authentication
    portal-group pg0

    lun 0 {
        path /dev/md0
    }
}
```

 **Note** Refer to the official FreeBSD documentation for more information.

2. Start iSCSI service:

```
[root@host~]# service ctld onestart
```

4. Software/Driver Unloading

Use the following command to unload the iSCSI PDU offload target module:

```
[root@host~]# kldunload cxgbei
```

V. Classification and Filtering

1. Introduction

Classification and Filtering feature enhances network security by controlling incoming traffic as they pass through network interface based on source and destination addresses, protocol, source and receiving ports, or the value of some status bits in the packet. This feature can be used in the ingress path to:

- Steer ingress packets that meet ACL (Access Control List) accept criteria to a particular receive queue.
- Switch (proxy) ingress packets that meet ACL accept criteria to an output port, with optional header rewrite.
- Drop ingress packets that meet ACL accept criteria.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio Adapters that are compatible with the Classification and Filtering feature:

- 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
- T520-CR
- T520-LL-CR
- T520-SO-CR
- T420-CR
- T420-LL-CR
- T440-CR
- T440-LP-CR
- T422-CR
- T420-SO-CR
- T420-CX
- T420-BT
- T404-BT

1.2. Software Requirements

1.2.1. FreeBSD Requirements

Currently the Classification and Filtering feature is available for the following versions:

- FreeBSD 10.3-RELEASE
- FreeBSD 11.0-RELEASE

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

2. Usage

2.1. Configuration

The Classification and Filtering feature is configured by specifying the filter selection combination set in the firmware configuration file *t6-config.txt* (*t5-config.txt* and *t4-config.txt* for T5 and T4 adapters respectively).

The following combination is set by default and packets will be matched accordingly:

```
filterMode = fragmentation, mpshittype, protocol, vlan, port, fcoe
```

Where,

fragmentation: Fragmented IP packets

mpshittype : MAC address "match type" (0=unicast, 1=unicast hash, 2=multicast, 3=multicast hash, 4=PROM, 5=hyper PROM, 6=broadcast, 7=none)

protocol : IP protocol number (ICMP=1, TCP=6, UDP=17, etc)

vlan : Inner VLAN Tag

port : Packet ingress port number

fcoe : Fibre Channel over Ethernet frames

2.2. Creating Filter Rules

Network driver (*if_cxgbe*) and *cxgbetool* must be installed before setting the filter rule.

Run the following command to load the network driver:

```
[root@host~]# kldload if_cxgbe
```

Now, create filter rules using *cxgbetool*:


```
[root@host~]# cxgbetool t6nexY filter <index> action [pass/drop/switch]
```

Where,

t6nexY : Chelsio T6 interface. Use *t5nexY* and *t4nexY* for T5 and T4 adapters respectively.

index : positive integer ranging from 0-495 set as filter id.

action : Ingress packet disposition.
pass : Ingress packets will be passed through set ingress queues.
switch : Ingress packets will be routed to an output port with optional header rewrite.
drop : Ingress packets will be dropped.

 **Note** *In case of multiple filter rules, the rule with the lowest filter index takes higher priority.*

2.2.1. Examples

- **Drop action**


```
[root@host~]# cxgbetool t6nexY filter 0 action drop sip 192.168.1.5
```

The above filter rule will drop all ingress packets from IP 192.168.1.5

- **Pass action**

```
[root@host~]# cxgbetool t6nexY filter 0 action pass sport 10001 dport 355  
queue 2
```

The above filter rule will pass all ingress packets that match source port 10001 and destination port 355 to the ingress queue with context id 2 for load balancing.

 **Note** *The context id for a particular ingress queue can be obtained using*
[root@host~]# sysctl dev.cxgbe.X.rxq.Y.cntxt_id

- **Switch action**

```
[root@host~]# cxgbetool t6nexY filter 0 action switch iport 0 eport 1 vlan 3
```

The above filter rule will route all ingress packets that match VLAN id 3 from port 0 of Chelsio adapter to port 1. Remaining packets will be sent to the host.

For offloaded ingress packets, use the `prio` argument with the above command:

```
[root@host~]# cxgbe tool t5nexY filter <index> action <pass|drop|switch>
prio 1
```



Note `prio` argument currently not supported with T6 adapters.

2.3. Listing Filter Rules

To list the filters set, run the following command:

```
[root@host~]# cxgbe tool t6nexY filter list
```

2.4. Removing Filter Rules

To remove a filter, run the following command with the corresponding filter rule index

```
[root@host~]# cxgbe tool t6nexY filter <index> delete
```



Note *For more information on additional parameters, refer `cxgbe tool` manual by running the `man cxgbe tool` command*

VI. Offload IPv6

1. Introduction

The growth of the Internet has created a need for more addresses than are possible with IPv4. Internet Protocol version 6 (IPv6) is a version of the Internet Protocol (IP) designed to succeed the Internet Protocol version 4 (IPv4). Chelsio's Offload IPv6 feature provides support to fully offload IPv6 traffic on to the Unified wire adapter.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio Offload IPv6 feature:

- T62100-CR
- T62100-LP-CR
- T62100-SO-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR
- T580-CR
- T580-LP-CR
- T540-CR
- T520-CR
- T520-LL-CR
- T420-CR
- T420-LL-CR
- T440-CR
- T440-LP-CR
- T422-CR
- T420-CX
- T420-BT
- T404-BT

1.2. Software Requirements

1.2.1. FreeBSD Requirements

Currently the Offload IPv6 feature is available for the following versions:

- FreeBSD 10.3-RELEASE
- FreeBSD 11.0-RELEASE

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

2. Software/Driver Configuration and Fine-tuning

- i. Add the following entry in `/etc/rc.conf` file:


```
ipv6_prefer="YES"
```

- ii. Reboot the system.
- iii. TOE driver (*t4_tom*) must be installed and loaded in order to use the Offload IPv6 feature. Also, the driver must be loaded by the root user. Any attempt to load the driver as a regular user will fail.
Load the TOE driver as:

```
[root@host~]# kldload t4_tom
```

- iv. For T6, bring up the interface and enable TOE:

```
[root@host~]# ifconfig ccX up  
[root@host~]# ifconfig ccX toe
```

 **Note** For T5 and T4, use **cxIX** and **cxgbeX** respectively.

All IPv6 traffic over Chelsio interfaces will be offloaded now.

VII. Appendix

Chelsio End-User License Agreement (EULA)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

11. Government Restricted Rights. The Software is subject to restricted rights as follows. If the Software is acquired under the terms of a GSA contract: use, reproduction or disclosure is subject to the restrictions set forth in the applicable ADP Schedule contract. If the Software is acquired under the terms of a DoD or civilian agency contract, use, duplication or disclosure by the Government is subject

to the restrictions of this Agreement in accordance with 48 C.F.R. 12.212 of the Federal Acquisition Regulations and its successors and 49 C.F.R. 227.7202-1 of the DoD FAR Supplement and its successors.

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

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

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