



Chelsio T3 Open-iSCSI Data Path Accelerator for Linux

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)

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

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 ©2011.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	06/20/2011
1.0.1	10/19/2011

TABLE OF CONTENTS

1. Introduction	1
1.1. Features	1
1.2. Hardware Requirements	2
1.3. Software Requirements	2
1.4. Package Contents	2
2. Hardware Installation	3
3. Software/Driver Installation	4
3.1. Installing Network (cxgb3toe v1.5.1.6) Driver	4
3.2. Installing Offload (cxgb3i) Driver	5
4. Software/Driver Uninstallation	7
5. Software/Driver Configuration and Fine-tuning	8
5.1. Using the cxgb3i Driver	8
6. Software/Driver Update	11
7. Appendix A	12
7.1. Chelsio End-User License Agreement (EULA)	12

1. Introduction

The Chelsio T3 based Adapters (S310, S320, S302, S304, Mezz cards, etc. series of products) 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.

1.1. Features

Chelsio's open-iSCSI Data Path Accelerator interfaces with open-iSCSI initiator and provides the iSCSI acceleration through Chelsio hardware wherever applicable.

This package comes with the following driver versions:

- cxgb3i v2.0.0.1
- libcxgbi v0.9.2

1.2. Hardware Requirements

1.2.1. Supported Adapters

The following are the currently shipping Chelsio Adapters that are compatible with Chelsio Open-iSCSI Data Path Accelerator:

- S302E
- S320E-CR
- S320E-LP-CR
- S320EM-BCH
- S320EM-BS

1.3. Software Requirements

Open iSCSI initiator (v2.0-872) and cxgb3-toe (v1.5.1.6) should be installed before installing the open-iSCSI data path accelerator.

For RHEL 5.6, open iSCSI initiator is available in the Operating System installation DVD.

1.3.1. Linux Requirements

The Chelsio Open-iSCSI Data Path Accelerator has been developed to run on 64-bit Linux based platforms. Currently the driver is available for following version(s):

- Redhat Enterprise Linux 5 update 6 kernel (RHEL5.6), 2.6.18-238.el5

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

1.4. Package Contents

The driver package comes with Chelsio Offload driver for open-iSCSI Initiator v2.0.0.1

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 Ethernet 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. For maximum performance, it is highly recommended to install the adapter into a PCI-X slot running at 133Mhz or PCIE x8 slot.
6. Holding the Chelsio Ethernet 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 XFP/SFP (optics) modules prior to inserting the adapter.
7. Secure the Chelsio Ethernet adapter with a screw, or other securing mechanism, as described by the system manufacturer's instructions. Replace the case cover.
8. 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.

3. Software/Driver Installation

The Chelsio Open-iSCSI offload solution works with two drivers and a Chelsio T3 based adapter. The two drivers consist of the lower level NIC driver (cxgb3) and the higher level Open-iSCSI (cxgb3i) driver. This section will be a guide for installing the cxgb3i driver on a Linux system.

The newer versions of the major Linux distributions (RHEL and SLES) already have these drivers in-boxed. However, it is recommended by Chelsio that the latest updated drivers be downloaded from the Chelsio's support website <http://service.chelsio.com> and used. They have the latest bug fixes, features, and performance optimizations that will improve the overall experience with Chelsio's Open-iSCSI solution.



*The installation instruction for network (cxgb3) and iSCSI transport (cxgb3i) driver provided in this document is for default version of the kernel available in major distributions (RHEL and SLES).
In case the kernel is updated either by the user manually or through the updates, the network and offload driver have to be re-compiled and installed for the new kernel.*

3.1. Installing Network (cxgb3toe v1.5.1.6) Driver

The Chelsio driver package cxgb3toe-1.5.1.6.tar.gz that is downloaded from website needs to be untarred by executing the command:

```
~# tar mzxvf cxgb3toe-1.5.1.6.tar.gz
```

Change directory to source directory of cxgb3toe-1.5.1.6 to compile, install and load the Chelsio network driver:

```
~# cd cxgb3toe-1.5.1.6/src
~/cxgb3toe-1.5.1.6/src # make && make install
~/cxgb3toe-1.5.1.6/src # modprobe cxgb3
```

Although `make` and `make install` is enough, the following command is recommended which will compile and load the driver without support for bonding:

```
~/cxgb3toe-1.5.1.6/src # make disable_bonding=1 && make install
```



Even though the `cxgb3toe` package has “toe” in the package name, installation of the offload NIC driver is harmless since it is not required by the Chelsio iSCSI transport driver (`cxgb3i`).

3.2. Installing Offload (`cxgb3i`) Driver

Typically most Linux distributions have an in-boxed version of the Open-iSCSI libraries installed. This is not to be confused with the Chelsio iSCSI transport driver (`cxgb3i`), but are instead the Open-iSCSI initiator libraries themselves. So installing the latest Open-iSCSI package from source is not necessary except for special cases where a new version fixes critical issues. These cases will be outlined in detail further in this guide.



Driver must be installed by the root user. Any attempt to install the drivers as a regular user will fail.

3.2.1. Installing `cxgb3i` on RHEL 5U6

Inbox open-iSCSI utilities package (`iscsi-initiator-utils`) available on the installer DVD needs to be installed for `cxgb3i` to work on these Operating systems.

For example, on RHEL 5U6, the `iscsi-initiator-utils` available is `iscsi-initiator-utils-6.2.0872-6.el5` which is available under Base System during OS installation.

If the open-iSCSI package is being installed after OS installation, the `iscsi-initiator-utils-6.2.0872-6.el5.rpm` is available under the “Server” folder in the DVD.

1. To install the `cxgb3i` driver, download the driver package (`cxgb3i-2.0.0.1-2.tar.gz`) from the Chelsio website and untar the package:

```
~ # tar mzxvf cxgb3i-2.0.0.1-2.tar.gz
```

2. Change directory to `cxgb3i-2.0.0.1-2` for compiling and installing `cxgb3i` driver:

```
~ # cd cxgb3i-2.0.0.1-2
~/cxgb3i-2.0.0.1-2 # make CXGB3TOE_SRC=<path to cxgb3toe>
```

3. To install the compiled driver, the following command needs to be executed:

```
~/cxgb3i-2.0.0.1-2 # make CXGB3TOE_SRC=<path to cxgb3toe> install
```

Assuming that the driver package is located under root, the command to be executed for compiling and installing the offload driver is:

```
~/cxgb3i-2.0.0.1-2 # make CXGB3TOE_SRC=/root/cxgb3toe-1.5.1.6  
~/cxgb3i-2.0.0.1-2 # make CXGB3TOE_SRC=/root/cxgb3toe-1.5.1.6/install
```

4. Reboot the system so that firmware dependencies are addressed and then load the driver using the `modprobe` command.

```
~/cxgb3i-2.0.0.1-2 # modprobe cxgb3i
```

4. Software/Driver Uninstallation

Chelsio offload driver (cxgb3i) should be uninstalled before uninstalling the Chelsio Network driver.

1. To un-install the cxgb3i driver, change directory to the cxgb3i driver package folder and run `make uninstall` command.

For example,

```
~ # cd /root/cxgb3i-2.0.0.1-2
~/cxgb3i-2.0.0.1-2 # make uninstall
```

2. To un-install the cxgb3 network driver, change directory to the cxgb3 driver package folder and run `make uninstall` command.

For example,

```
~ # cd cxgb3toe-1.5.1.6
~/cxgb3toe-1.5.1.6 # make uninstall
```

5. Software/Driver Configuration and Fine-tuning

5.1. Using the cxgb3i Driver

Create an interface file located under `ifaces` directory for the new open-iSCSI transport class “cxgb3i” provided by the cxgb3i iSCSI transport. This is `/etc/iscsi/ifaces/` if you compiled open-iSCSI utilities yourself, or `/var/lib/iscsi/ifaces` for inbox open-iSCSI utilities.

The content of the file should be in the following format:

```
iface.iscsi_ifacename = <cxgb3i_file_name>
iface.transport_name = cxgb3i
iface.net_ifacename = <ethX>
iface.ipaddress = <iscsi_ip_address>
```

Note *If **iface.ipaddress** is specified, **<iscsi_ip_address>** 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.*

The iface files can be automatically created using the following command.

```
~ # iscsiadm -m iface
```

The cxgb3i interfaces are listed as “cxgb3i.<mac address>” under `/etc/iscsi/ifaces`. Edit the file to enter the IP address.

For example, the content of the iface file would like,

```
iface.iscsi_ifacename = cxgb3i.00:07:43:04:5b:da
iface.hwaddress = 00:07:43:04:5b:da
iface.transport_name = cxgb3i
iface.net_ifacename = eth3
iface.ipaddress = 102.2.2.137
```

Here,

- `iface.iscsi_ifacename` denotes the name of iface file in your applicable ifaces directory

- `iface.hwaddress` denotes the MAC address of the chelsio interface via which iSCSI traffic will be running.
- `iface.net_ifacename` is the chelsio interface name `ethX` (that appears on running `ifconfig`) through which you want to login. This is recommended when using vlan interfaces In case of login through VLAN interface use `iface.net_ifacename = ethX.Y` where X is chelsio interface id and Y is VLAN Id.
- `iface.transport_name` denotes the transport name, which is `cxgb3i`.
- `iface.ipaddress` denotes the IP address which is assigned to the interface.

For login to work properly, Please edit `/etc/iscsi/iscsid.conf` and change the setting for `MaxRecvDataSegmentLength` :

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

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

5.1.1. Starting the iSCSI daemon

The iSCSI daemon has to be started for the initiator to work. The daemon can be started by executing the following command:

```
# chkconfig <service name> on
# service <service name> start
```

Where `<service name>` is `iscsid` for inbox iSCSI utilities and `open-iscsi` if you compiled the utilities from source. The `chkconfig` command needs to be executed only on the first run.

5.1.2. Discovery, Login and Logout using the initiator

Using the initiator, iSCSI target should be discovered and then logged in to access the target disk.

1. To discover the target the command to be executed is:

```
# iscsiadm -m discovery -t st -p <target ip address>:<target port no> -I
<cxgb3i iface file name>
```

For example,

```
# iscsiadm -m discovery -t st -p 102.2.2.155:3260 -I  
cxgb3i.00:07:43:04:5b:da
```

2. The following command needs to be executed to login to the target :

```
iscsiadm -m node -T <iqn name of target> -p <target ip address>:<target port  
no> -I <cxgb3i iface file name> -l
```

For example,

```
# iscsiadm -m node -T iqn.2004-05.com.chelsio.target1 -p  
192.51.100.155:3260,1 -I cxgb3i.00:07:43:04:5b:da -l
```

3. To Logout, the command to be used is :

```
# iscsiadm -m node -T <iqn name of target> -p <target ip address>:<target  
port no> -I <cxgb3i iface file name> -u
```

For example,

```
# iscsiadm -m node -T iqn.2004-05.com.chelsio.target1 -p  
192.51.100.155:3260,1 -I cxgb3i.00:07:43:04:5b:da -u
```

 **Note** The `,1` in `192.51.100.155:3260,1` refers to the target portal group tag

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

6. 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.

7. Appendix A

7.1. Chelsio End-User License Agreement (EULA)

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

May 16, 2011

CHELSIO END USER LICENSE AGREEMENT

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