Chelsio Communications

Accelerate

Chelsio Unified Wire 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)
209 North Fair Oaks Avenue,
Sunnyvale, CA 94085
U.S.A

www.chelsio.com

Tel: 408.962.3600
Fax: 408.962.3661

Chelsio (India) Private Limited
Subramanya Arcade, Floor 3, Tower B
No. 12, Bannerghatta Road,
Bangalore-560029
Karnataka,
India

Tel: +91-80-4039-6800

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

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

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

Copyright © 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.

Document History

Version	Revision Date
1.0.0	12/08/2011
1.0.1	01/09/2013
1.0.2	01/27/2013
1.0.3	03/26/2013
1.0.4	04/12/2013
1.0.5	06/20/2013
1.0.6	08/17/2013
1.0.7	10/22/2013
1.0.8	03/08/2013
1.0.9	05/15/2013
1.1.0	07/26/2013
1.1.1	08/14/2013
1.1.2	12/06/2013
1.1.3	12/19/2013
1.1.4	03/13/2014
1.1.5	05/02/2014
1.1.6	06/30/2014
1.1.7	10/22/2014
1.1.8	11/04/2014
1.1.9	02/05/2015
1.2.0	03/04/2015
1.2.1	03/25/2015
1.2.2	06/03/2015
1.2.3	08/05/2015
1.2.4	02/29/2016
1.2.5	04/27/2016
1.2.6	08/25/2016
1.2.7	10/07/2016
1.2.8	10/18/2016
1.2.9	11/11/2016
1.3.0	12/05/2016
1.3.1	12/30/2016
1.3.2	01/30/2017
1.3.3	02/27/2017
1.3.4	05/11/2017
1.3.5	09/05/2017

1.3.6	**09/29/2017**

TABLE OF CONTENTS

I. CHELSIO UNIFIED WIRE 15

1. Introduction 16

1.1. Features 16

1.2. Hardware Requirements 17

1.3. Software Requirements 17

1.4. Package Contents 17

2. Hardware Installation 21

3. Software/Driver Installation 23

3.1. Pre-requisites 23

3.2. Enabling RDMA on ARM Platforms 24

3.3. Mounting debugfs 24

3.4. Allowing unsupported modules on SLES 25

3.5. Installing Chelsio Unified Wire from source 25

3.6. Installing Chelsio Unified Wire from RPM 32

3.7. Firmware Update 35

4. Configuring Chelsio Network Interfaces 36

4.1. Configuring 40G adapters 36

4.2. Configuring network-scripts 37

4.3. Creating network-scripts 38

4.4. Checking Link 38

5. Performance Tuning 40

6. Software/Driver Update 42

7. Software/Driver Uninstallation 43

7.1. Uninstalling Chelsio Unified Wire from source 43

7.2. Uninstalling Chelsio Unified Wire from RPM 47

II. NETWORK (NIC/TOE) 48

1. Introduction 49

1.1. Hardware Requirements 49

1.2. Software Requirements 50

2. Software/Driver Installation 51

3. Software/Driver Loading 52

3.1. Loading in NIC mode (without full offload support) 52

3.2. Loading in TOE mode (with full offload support) 52

4. Software/Driver Configuration 53

4.1. Disabling Auto-negotiation 53

4.2. Enabling TCP Offload 53

4.3. Enabling Busy waiting 54

4.4. Precision Time Protocol (PTP) 54

4.5. Performance Tuning 57

5. Software/Driver Unloading **63**
5.1. Unloading the NIC driver 63
5.2. Unloading the TOE driver 63

III. VIRTUAL FUNCTION NETWORK (VNIC) **65**
1. Introduction **66**
1.1. Hardware Requirements 66
1.2. Software Requirements 67
2. Software/Driver Installation **68**
2.1. Pre-requisites 68
2.2. Installation 68
3. Software/Driver Loading **69**
3.1. Instantiate Virtual Functions (SR-IOV) 69
3.2. Loading the driver 69
4. Software/Driver Configuration and Fine-tuning **71**
5. Software/Driver Unloading **72**
5.1. Unloading the driver 72

IV. IWARP (RDMA) **73**
1. Introduction **74**
1.1. Hardware Requirements 74
1.2. Software Requirements 74
2. Software/Driver Installation **76**
2.1. Pre-requisites 76
2.2. Installation 76
3. Software/Driver Loading **77**
3.1. Loading iWARP driver 77
4. Software/Driver Configuration and Fine-tuning **78**
4.1. Testing connectivity with ping and rping 78
4.2. Enabling various MPIs 79
4.3. Setting up NFS-RDMA 87
4.4. Performance Tuning 88
5. Software/Driver Unloading **89**

V. ISER **90**
1. Introduction **91**
1.1. Hardware Requirements 91
1.2. Software Requirements 91
2. Kernel Installation **92**
3. Software/Driver Installation **93**
3.1. Pre-requisites 93
3.2. Installation 93

4. Software/Driver Loading 94
5. Software/Driver Configuration and Fine-tuning 96
5.1. HMA 97
5.2. Performance Tuning 98
6. Software/Driver Unloading 99

VI. WD-UDP 100
1. Introduction 101
1.1. Hardware Requirements 101
1.2. Software Requirements 101
2. Software/Driver Installation 103
3. Software/Driver Loading 104
4. Software/Driver Configuration and Fine-tuning 105
4.1. Accelerating UDP Socket communications 105
5. Software/Driver Unloading 111

VII. WD-TOE 112
1. Introduction 113
1.1. Hardware Requirements 113
1.2. Software Requirements 113
2. Software/Driver Installation 114
2.1. Pre-requisites 114
2.2. Installation 114
3. Software/Driver Loading 115
4. Software/Driver Configuration and Fine-tuning 116
4.1. Running the Application 116
5. Software/Driver Unloading 117

VIII. NVME-OF 118
1. Introduction 119
1.1. Hardware Requirements 119
1.2. Software Requirements 119
2. Kernel Installation 120
3. Software/Driver Installation 121
3.1. Pre-requisites 121
3.2. Installlation 121
4. Software/Driver Loading 122
5. Software/Driver Configuration and Fine-tuning 123
5.1. Target 123
5.2. Initiator 124
5.3. HMA 125
5.4. Performance Tuning 126

6. Software/Driver Unloading 127

IX. LIO ISCSI TARGET OFFLOAD 128

1. Introduction 129
1.1. Hardware Requirements 129
1.2. Software Requirements 129
2. Kernel Installation 131
3. Software/Driver Installation 132
3.1. Pre-requisites 132
3.2. Installlation 132
4. Software/Driver Loading 133
5. Software/Driver Configuration and Fine-tuning 134
5.1. Configuring LIO iSCSI Target 134
5.2. Offloading LIO iSCSI Connection 134
5.3. Running LIO iSCSI and Network Traffic Concurrently 135
5.4. Performance Tuning 136
6. Software/Driver Unloading 137
6.1. Unloading the LIO iSCSI Target Offload driver 137
6.2. Unloading the NIC driver 137

X. ISCSI PDU OFFLOAD TARGET 138

1. Introduction 139
1.1. Features 139
1.2. Hardware Requirements 140
1.3. Software Requirements 141
2. Software/Driver Installation 143
3. Software/Driver Loading 144
3.1. Latest iSCSI Software Stack Driver Software 144
4. Software/Driver Configuration and Fine-tuning 146
4.1. Command Line Tools 146
4.2. iSCSI Configuration File 146
4.3. A Quick Start Guide for Target 147
4.4. The iSCSI Configuration File 149
4.5. Challenge-Handshake Authenticate Protocol (CHAP) 160
4.6. Target Access Control List (ACL) Configuration 162
4.7. Target Storage Device Configuration 164
4.8. Target Redirection Support 166
4.9. The command line interface tools “iscsictl” & “chisns” 168
4.10. Rules of Target Reload (i.e. “on the fly” changes) 173
4.11. System Wide Parameters 174
4.12. Performance Tuning 175
5. Software/Driver Unloading 176

XI. ISCSI PDU OFFLOAD INITIATOR 177
1. Introduction 178
1.1. Hardware Requirements 178
1.2. Software Requirements 179
2. Software/Driver Installation 180
2.1. Pre-requisites 180
2.2. Installation 180
3. Software/Driver Loading 181
4. Software/Driver Configuration and Fine-tuning 182
4.1. Accelerating open-iSCSI Initiator 182
4.2. Auto login from cxgb4i initiator at OS bootup 184
4.3. Performance Tuning 185
5. Software/Driver Unloading 187
XII. CRYPTO OFFLOAD 188
1. Introduction 189
1.1. Hardware Requirements 189
1.2. Software Requirements 189
2. Kernel Installation 190
3. Software/Driver Installation 191
3.1. Pre-requisites 191
3.2. Installation 191
4. Software/Driver Loading 192
4.1. Co-processor 192
4.2. Inline 192
5. Software/Driver Configuration and Fine-tuning 193
5.1. Co-processor 193
5.2. Inline 194
5.3. Performance Tuning 197
6. Software/Driver Unloading 199
XIII. DATA CENTER BRIDGING (DCB) 200
1. Introduction 201
1.1. Hardware Requirements 201
1.2. Software Requirements 201
2. Software/Driver Installation 203
3. Software/Driver Loading 204
4. Software/Driver Configuration and Fine-tuning 206
4.1. Configuring Cisco Nexus 5010 switch 206
4.2. Configuring the Brocade 8000 switch 209
5. Running NIC & iSCSI Traffic together with DCBx 211

XIV. FCOE FULL OFFLOAD INITIATOR 212
1. Introduction 213
1.1. Hardware Requirements 213
1.2. Software Requirements 213
2. Software/Driver Installation 214
3. Software/Driver Loading 215
4. Software/Driver Configuration and Fine-tuning 216
4.1. Configuring Cisco Nexus 5010 and Brocade switch 216
4.2. FCoE fabric discovery verification 216
4.3. Formatting the LUNs and Mounting the Filesystem 220
4.4. Creating Filesystem 221
4.5. Mounting the formatted LUN 222
5. Software/Driver Unloading 223
XV. OFFLOAD BONDING 224
1. Introduction 225
1.1. Hardware Requirements 225
1.2. Software Requirements 225
2. Software/Driver Installation 227
3. Software/Driver Loading 228
4. Software/Driver Configuration and Fine-tuning 229
4.1. Offloading TCP traffic over a bonded interface 229
5. Software/Driver Unloading 230
XVI. OFFLOAD MULTI-ADAPTER FAILOVER (MAFO) 231
1. Introduction 232
1.1. Hardware Requirements 232
1.2. Software Requirements 233
2. Software/Driver Installation 234
3. Software/Driver Loading 235
4. Software/Driver Configuration and Fine-tuning 236
4.1. Offloading TCP traffic over a bonded interface 236
5. Software/Driver Unloading 238
XVII. UDP SEGMENTATION OFFLOAD AND PACING 239
1. Introduction 240
1.1. Hardware Requirements 241
1.2. Software Requirements 241
2. Software/Driver Installation 243
3. Software/Driver Loading 244
4. Software/Driver Configuration and Fine-tuning 245
4.1. Modifying the application 245

4.2. Configuring UDP Pacing 247
5. Software/Driver Unloading 249

XVIII.OFFLOAD IPV6 250
1. Introduction 251
1.1. Hardware Requirements 251
1.2. Software Requirements 251
2. Software/Driver Installation 253
2.1. Pre-requisites 253
2.2. Installation 253
3. Software/Driver Loading 254
4. Software/Driver Configuration and Fine-tuning 255
5. Software/Driver Unloading 256
5.1. Unloading the NIC driver 256
5.2. Unloading the TOE driver 256

XIX. WD SNIFFING AND TRACING 257
1. Theory of Operation 258
1.1. Hardware Requirements 259
1.2. Software Requirements 260
2. Software/Driver Installation 261
3. Usage 262
3.1. Installing basic support 262
3.2. Using Sniffer (wd_sniffer) 262
3.3. Using Tracer (wd_tcpdump_trace) 262

XX. CLASSIFICATION AND FILTERING 264
1. Introduction 265
1.1. Hardware Requirements 265
1.2. Software Requirements 266
2. LE-TCAM Filters 267
2.1. Configuration 267
2.2. Creating Filter Rules 270
2.3. Listing Filter Rules 272
2.4. Removing Filter Rules 272
2.5. Layer 3 example 273
2.6. Layer 2 example 276
3. Hash/DDR Filters 279
3.1. Configuration 279
3.2. Creating Filter Rules 281
3.3. Listing Filter Rules 284
3.4. Removing Filter Rules 285

3.5. Filter Priority 285
3.6. Swap MAC Feature 285
3.7. Traffic Mirroring 285
3.8. Packet Tracing and Hit Counters 287
4. NAT Filtering 290

XXI. OVS KERNEL DATAPATH OFFLOAD 291
1. Introduction 292
1.1. Hardware Requirements 292
1.2. Software Requirements 292
2. Software/Driver Installation 293
2.1. Pre-requisites 293
2.2. Installation 293
3. Software/Driver Configuration and Fine Tuning 294
3.1. Configuring OVS Machine 295
3.2. Creating OVS flows 297
3.3. Verifying OVS Flow Dump 301
3.4. Setting up ODL with OVS 301

XXII. TRAFFIC MANAGEMENT 303
1. Introduction 304
1.1. Hardware Requirements 304
1.2. Software Requirements 305
2. Software/Driver Loading 306
3. Software/Driver Configuration and Fine-tuning 307
3.1. Traffic Management Rules 307
3.2. Configuring Traffic Management 309
4. Usage 313
4.1. Non-Offloaded Connections 313
4.2. Offloaded Connections 313
4.3. Offloaded Connections with Modified Application 314
5. Software/Driver Unloading 315

XXIII.DPDK DRIVER 316
1. Introduction 317
1.1. Hardware Requirements 317
1.2. Software Requirements 317
2. Software/Driver Installation 319
2.1. Pre-requisites 319
2.2. Installation 319
3. Flashing Firmware Configuration File 320
4. Software/Driver Loading 321

5. Software/Driver Configuration and Fine Tuning 323
5.1. Huge Pages 323
5.2. Binding network ports 324
5.3. Unbinding network ports 325
5.4. Performance Tuning 325
6. Running DPDK Test Applications 326
6.1. Testpmd application 326
6.2. Pktgen application 329
7. Software/Driver Unloading 333
8. Software/Driver Uninstallation 334

XXIV. UNIFIED BOOT 335
1. Introduction 336
1.1. Hardware Requirements 336
1.2. Software Requirements 337
1.3. Pre-requisites 338
2. Secure Boot 339
3. Flashing firmware and option ROM 340
3.1. Preparing USB flash drive 340
3.2. Legacy 341
3.3. uEFI 344
3.4. HP Firmware Management Protocol (FMP) 350
3.5. Default Option ROM Settings 355
4. Configuring PXE Server 357
5. PXE boot process 358
5.1. Legacy PXE boot 358
5.2. uEFI PXE Boot 361
6. FCoE boot process 365
6.1. Legacy FCoE boot 365
6.2. uEFI FCoE Boot 370
7. iSCSI boot process 376
7.1. Legacy iSCSI boot 376
7.2. uEFI iSCSI Boot 384
8. Creating Driver Update Disk (DUD) 393
8.1. Creating DUD for RedHat Enterprise Linux 393
8.2. Creating DUD for Suse Enterprise Linux 393
9. OS Installation 395
9.1. Installation using Chelsio DUD 395
9.2. Installation on FCoE LUN 406
9.3. Installation on iSCSI LUN 409

XXV. APPENDIX A 421

1. Troubleshooting **422**
2. Chelsio End-User License Agreement (EULA) **424**

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 1/10/25/40/50/100GbE controllers. The adapters are built around a programmable protocol-processing engine, with full offload of a complete Unified Wire solution comprising NIC, TOE, iWARP RDMA, iSCSI, FCoE and NAT support. It scales to true 100Gb line rate operation from a single TCP connection to thousands of connections, and allows simultaneous low latency and high bandwidth operation thanks to multiple physical channels through the ASIC.

Ideal for all data, storage and high-performance clustering applications, the Unified Wire 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.

Chelsio Unified Wire software comes in two formats: Source code and RPM package forms. Installing from source requires compiling the package to generate the necessary binaries. You can choose this method when you are using a custom-built kernel. You can also install the package using the interactive GUI installer. In other cases, download the RPM package specific to your operating system and follow the steps mentioned to install the package. Please note that the OFED software required to install Chelsio iWARP driver comes bundled in both source as well as RPM packages.

This document describes the installation, use and maintenance of Unified Wire software and its various components.

1.1. Features

The Chelsio Unified Wire Package uses a single command to install various drivers and utilities. It consists of the following software:

- **Network (NIC/TOE)**
- **Virtual Function Network (vNIC)**
- **iWARP (RDMA)**
- **iSER**
- **WD-UDP**
- **WD-TOE**
- **NVMe-oF**

- **LIO iSCSI Target Offload**
- **iSCSI PDU Offload Target**
- **iSCSI PDU Offload Initiator**
- **Crypto Offload**
- **Data Center Bridging (DCB)**
- **FCoE full offload Initiator**
- **Offload Bonding**
- **Offload Multi-Adapter Failover(MAFO)**
- **UDP Segmentation Offload and Pacing**
- **Offload IPv6**
- **Classification and Filtering feature**
- **OVS Kernel Datapath Offload**
- **Traffic Management feature (TM)**
- **DPDK**
- **Unified Boot Software**
- **Utility Tools (cop, cxgbtool, t4_perftune, benchmark tools, sniffer & tracer)**
- **libs (iWARP, WD-UDP and WD-TOE libraries)**

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

1.2. Hardware Requirements

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

1.3. Software Requirements

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

1.4. Package Contents

1.4.1. Source Package

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

- **debrules**: This directory contains packaging specification files required for building Debian packages.

- **docs:** This directory contains support documents - README, Release Notes and User's Guide (this document) for the software.
- **libs:** This directory is for libraries required to install the WD-UDP, WD-TOE and iWARP drivers. The libibverbs library has implementation of RDMA verbs which will be used by iWARP applications for data transfers. The librdmacm library works as an RDMA connection manager. The libcxgb4 library works as an interface between the above mentioned generic libraries and Chelsio iWARP driver. The libcxgb4_sock library is a LD_PRELOAD-able library that accelerates UDP Socket communications transparently and without recompilation of the user application.
- **OFED**: This directory contains supported OFED packages.
- **RPM-Manager**: This directory contains support scripts used for cluster deployment.
- **scripts:** Support scripts used by the Unified Wire Installer.
- **specs:** The packaging specification files required for building RPM packages.
- **src:** Source code for different drivers.
- **support:** This directory contains source files for the dialog utility.
- **tools:**
    - **autoconf-x.xx**: This directory contains the source for Autoconf tool needed for WD-UDP and iWARP libraries.
    - **benchmarks:** This directory contains various benchmarking tools to measure throughput and latency of various networks.
    - **chelsio_adapter_config**: This directory contains scripts and binaries needed to configure Chelsio 40G adapters.
    - **cop:** The cop tool compiles offload policies into a simple program form that can be loaded into the kernel and interpreted. These offload policies are used to determine the settings to be used for various connections. The connections to which the settings are applied are based on matching filter specifications. Please find more details on this tool in its manual page (run `man cop` command).
    - **cudbg:** Chelsio Unified Debug tool which facilitates collection and viewing of various debug entities like register dump, Devlog, CIM LA, etc.
    - **cxgbtool:** The cxgbtool queries or sets various aspects of Chelsio network interface cards. It complements standard tools used to configure network settings and provides functionality not available through such tools. Please find more details on this tool in its manual page (run `man cxgbtool` command).

      Note *To use cxbtool for FCoE Initiator driver, use [root@host~]# cxgbtool stor -h*

    - **rdma_tools:** This directory contains iWARP benchmarking tools.
    - **t4_sniffer:** This directory contains sniffer tracing and filtering libraries. See WD Sniffing and Tracing chapter for more information.
    - **90-rdma.rules:** This file contains udev rules needed for running RDMA applications as a non-root user.
    - **chdebug:** This script collects operating system environment details and debug information which can be sent to the support team, to troubleshoot Chelsio hardware/software related issues.

- **chiscsi_set_affinity.sh**: This shell script is used for mapping iSCSI Worker threads to different CPUs.
- **chsetup:** The chsetup tool loads NIC, TOE and iWARP drivers, and creates WD-UDP configuration file.
- **chstatus:** This utility provides status information on any Chelsio NIC in the system.
- **t4_latencytune.sh**: Script used for latency tuning of Chelsio adapters.
- **t4_perftune.sh:** This shell script is to tune the system for higher performance. It achieves it through modifying the IRQ-CPU binding. This script can also be used to change Tx coalescing settings.
- **t4-forward.sh:** RFC2544 Forward test tuning script.
- **uname_r:** This file is used by *chstatus* script to verify if the Linux platform is supported or not.
- **wdload:** UDP acceleration tool.
- **wdunload:** Used to unload all the loaded Chelsio drivers.
- **target**: The directory contains *targetcli* installation files and dependent components.

- **install.py, dialog.py:** Python scripts needed for the GUI installer.
- **EULA:** Chelsio's End User License Agreement
- **install.log:** File containing installation summary.
- **Makefile:** The Makefile for building and installing from the source.
- **sample_machinefile:** Sample file used during iWARP installation on cluster nodes.
- **Uboot:** There are two sub-directories in the *Uboot directory: OptionROM* and *LinuxDUD.* The *OptionROM* directory contains Unified Boot Option ROM image (*cubt4.bin*), uEFI driver (*ChelsioUD.efi*), default boot configuration file (*bootcfg*) and a legacy flash utility (*cfut4.exe*), which can be used to flash the option ROM onto Chelsio adapters (CNAs).
  The *LinuxDUD* directory contains image (.img) files required to update drivers for Linux distributions.

1.4.2. RPM package

The Chelsio Unified Wire RPM package consists of the following:

- **config**: This directory contains firmware configuration files.
- **docs:** This directory contains support documents i.e. README, Release Notes and User's Guide (this document) for the software.
- **DRIVER-RPMS:** RPM packages of Chelsio drivers.
- **OFED-RPMS:** OFED RPM packages required to install iWARP driver.
- **scripts:** Support scripts used by the Unified Wire Installer.
- **EULA:** Chelsio's End User License Agreement.
- **install.py:** Python script that installs the RPM package. See Software/Driver Installation section for more information.
- **uninstall.py:** Python script that uninstalls the RPM package. See Software/Driver Uninstallation section for more information.
- **Uboot:** There are two sub-directories in the *Uboot directory: OptionROM* and *LinuxDUD.* The *OptionROM* directory contains Unified Boot Option ROM image (*cubt4.bin*), uEFI driver

(*ChelsioUD.efi*), default boot configuration file (*bootcfg*) and a legacy flash utility (*cfut4.exe*), which can be used to flash the option ROM onto Chelsio adapters (CNAs).

The *LinuxDUD* directory contains image (.img) files required to update drivers for Linux distributions.

2. Hardware Installation

Follow these steps to install Chelsio adapter in your system:

i. Shutdown/power off your system.
ii. Power off all remaining peripherals attached to your system.
iii. Unpack the Chelsio adapter and place it on an anti-static surface.
iv. Remove the system case cover as per the system manufacturer's instructions.
v. Remove the PCI filler plate from the slot where you will install the Ethernet adapter.
vi. For maximum performance, it is highly recommended to install the adapter into a PCIe x8/x16 slot.

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

vii. Holding the Chelsio adapter by the edges, align the edge connector with the PCI connector on the motherboard. Apply even pressure on both edges until the card is firmly seated. It may be necessary to remove the SFP (transceiver) modules prior to inserting the adapter.
viii. Secure the Chelsio adapter with a screw, or other securing mechanism, as described by the system manufacturer's instructions. Replace the case cover.
ix. After securing the card, ensure that the card is still fully seated in the PCIE x8/x16 slot as sometimes the process of securing the card causes the card to become unseated.
x. Connect a fiber/twinax cable, multi-mode for short range (SR) optics or single-mode for long range (LR) optics, to the Ethernet adapter or regular Ethernet cable for the 1Gb Ethernet adapter.
xi. Power on your system.
xii. Run *update-pciids* command to download the current version of PCI ID list

```
[root@ ~]# update-pciids
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  227k  100  227k    0     0  68592      0  0:00:03  0:00:03 --:--:-- 68610
Done.
```

xiii. Verify if the adapter was installed successfully by using the *lspci* command

```
[root@ ~]# lspci | grep -i Chelsio
81:00.0 Ethernet controller: Chelsio Communications Inc T62100-LP-CR Unified Wire Ethernet Controller
81:00.1 Ethernet controller: Chelsio Communications Inc T62100-LP-CR Unified Wire Ethernet Controller
81:00.2 Ethernet controller: Chelsio Communications Inc T62100-LP-CR Unified Wire Ethernet Controller
81:00.3 Ethernet controller: Chelsio Communications Inc T62100-LP-CR Unified Wire Ethernet Controller
81:00.4 Ethernet controller: Chelsio Communications Inc T62100-LP-CR Unified Wire Ethernet Controller
81:00.5 SCSI storage controller: Chelsio Communications Inc T62100-LP-CR Unified Wire Storage Controller
81:00.6 Fibre Channel: Chelsio Communications Inc T62100-LP-CR Unified Wire Storage Controller
```

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 adapter
- Physical function 5: for iSCSI

- Physical function 6: for FCoE
- Physical function 7: Currently not assigned

Once Unified Wire package is installed and loaded, examine the output of `dmesg` to see if the card is discovered. You should see a similar output:

```
[ 1119.854346] cxgb4 0000:81:00.4: Chelsio T62100-LP-CR rev 0
[ 1119.854347] cxgb4 0000:81:00.4: S/N: RE41160042, P/N: 11012106003
[ 1119.854348] cxgb4 0000:81:00.4: Firmware version: [redacted]
[ 1119.854349] cxgb4 0000:81:00.4: Bootstrap version: 255.255.255.255
[ 1119.854350] cxgb4 0000:81:00.4: TP Microcode version: 0.1.23.2
[ 1119.854351] cxgb4 0000:81:00.4: No Expansion ROM loaded
[ 1119.854351] cxgb4 0000:81:00.4: Serial Configuration version: 0x7002000
[ 1119.854352] cxgb4 0000:81:00.4: VPD version: 0x52
[ 1119.854354] cxgb4 0000:81:00.4: Configuration: NIC MSI-X, non-Offload capable
[ 1119.854355] eth0: Chelsio T62100-LP-CR (eth0) 100GBASE-CR4_QSFP
```

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

Note

Network device names for Chelsio's physical ports are assigned using the following convention: the port farthest from the motherboard will appear as the first network interface. However, for T5 40G 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

There are two main methods to install the Chelsio Unified Wire package: from source and RPM.
If you decide to use source, you can install the package using CLI or GUI mode.
If you decide to use RPM, you can install the package using Menu or CLI mode.

RPM packages support only distro base kernels. In case of updated/custom kernels, use source package. Irrespective of the method chosen for installation, the machine needs to be rebooted for changes to take effect.

The following table describes the various *configuration tuning options* available during installation and drivers/software installed with each option by default:

Configuration Tuning Option	Description	Driver/Software installed
Unified Wire(Default)	Default Configuration. Configures adapters to run all protocols simultaneously.	NIC/TOE, vNIC, iWARP, iSER, WD-UDP, NVMe-oF, LIO iSCSI Target, iSCSI Target, iSCSI Initiator, FCoE Initiator, Bonding, MAFO, IPv6, Sniffer & Tracer, Filtering, TM
Low latency Networking	Configures adapters to run TOE and iWARP traffic with low latency.	TOE, iWARP, WD-UDP, WD-TOE, IPv6, Bonding, MAFO
High capacity RDMA	Configures adapters to establish a large number of iWARP connections.	iWARP
RDMA Performance	Improves iWARP performance.	iWARP, iSER, NVMe-oF
High capacity TOE	Configures adapters to establish a large number of TOE connections.	TOE, Bonding, MAFO, IPv6
iSCSI Performance#	Improves iSCSI performance.	LIO iSCSI Target, iSCSI Target, iSCSI Initiator, Bonding, DCB
UDP Seg.Offload & Pacing*	Configures adapters to establish a large number of UDP Segmentation Offload connections.	UDP-SO, Bonding
Wire Direct Latency#	Configures adapters to provide low Wire Direct latency.	TOE, iWARP, WD-UDP, WD-TOE
High Capacity WD	Configures adapters to establish a large number of WD-UDP connections.	WD-UDP, WD-TOE
Hash Filter#	Configures adapters to create more filters.	Filtering

** Supported on T4/T5*

+ Supported only on T5

 # Supported on T5/T6

3.1. Pre-requisites

To install Unifided Wire using GUI mode (with Dialog utility), *ncurses-devel* package must be installed.

3.2. Enabling RDMA on ARM Platforms

RDMA is disabled by default in RHEL 7.X build of ARM architecture. To enable this feature, follow the steps mentioned below:

i. Download the kernel source package and extract it.
ii. Create a kernel configuration file.

```
[root@host~]# make oldconfig
```

iii. The above command will create a configuration file *.config* in the same location. Edit the file and enable the following parameters:

```
CONFIG_NET_VENDOR_CHELSIO=y
CONFIG_INFINIBAND=y
```

iv. Compile the kernel.
v. During kernel compilation, please ensure that the following parameters are set as follows:

```
CONFIG_CHELSIO_T1=m
CONFIG_CHELSIO_T1_1G=y
CONFIG_CHELSIO_T3=m
CONFIG_CHELSIO_T4=m
CONFIG_CHELSIO_T4VF=m
CONFIG_INFINIBAND_USER_MAD=m
CONFIG_INFINIBAND_USER_ACCESS=m
CONFIG_INFINIBAND_USER_MEM=y
CONFIG_INFINIBAND_CXGB3=m
CONFIG_INFINIBAND_CXGB3_DEBUG=y
CONFIG_INFINIBAND_CXGB4=m
CONFIG_SCSI_CXGB3_ISCSI=m
CONFIG_SCSI_CXGB4_ISCSI=m
```

vi. Install the kernel.
vii. Reboot into the newly installed kernel.

3.3. Mounting debugfs

All driver debug data is stored in debugfs, which will be mounted in most cases. If not, mount it manually using:

```
[root@host~]# mount -t debugfs none /sys/kernel/debug
```

3.4. Allowing unsupported modules on SLES

On SLES11 SPx platforms, edit the */etc/modprobe.d/unsupported-modules* file and change *allow_unsupported_modules* to 1.

On SLES12 SPx platforms, edit the */etc/modprobe.d/10-unsupported-modules.conf* file and change *allow_unsupported_modules* to 1.

3.5. Installing Chelsio Unified Wire from source

3.5.1. GUI mode (with Dialog utility)

i. Download the Unified Wire driver package (tarball) from Chelsio Download Center.
ii. Untar the tarball using the following command:

```
[root@host~]# tar zxvf <driver_package>.tar.gz
```

iii. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

iv. Run the following script to start the GUI installer:

```
[root@host~]# ./install.py
```

v. If **Dialog** utility is present, you can skip to step (vi). If not, press ‘y’ to install it when the installer prompts for input.

vi. Select “install” under “Choose an action”

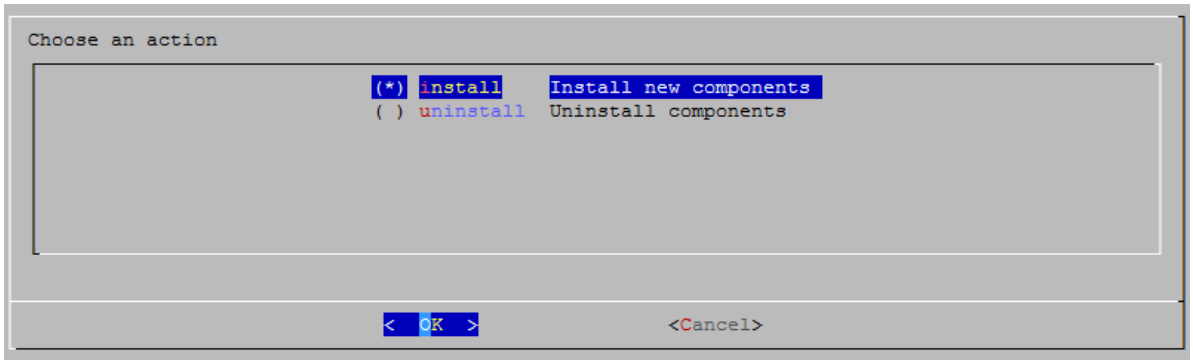

vii. Select *Enable IPv6-Offload* to install drivers with IPv6 Offload support or *Disable IPv6-offload* to continue installation without IPv6 offload support.

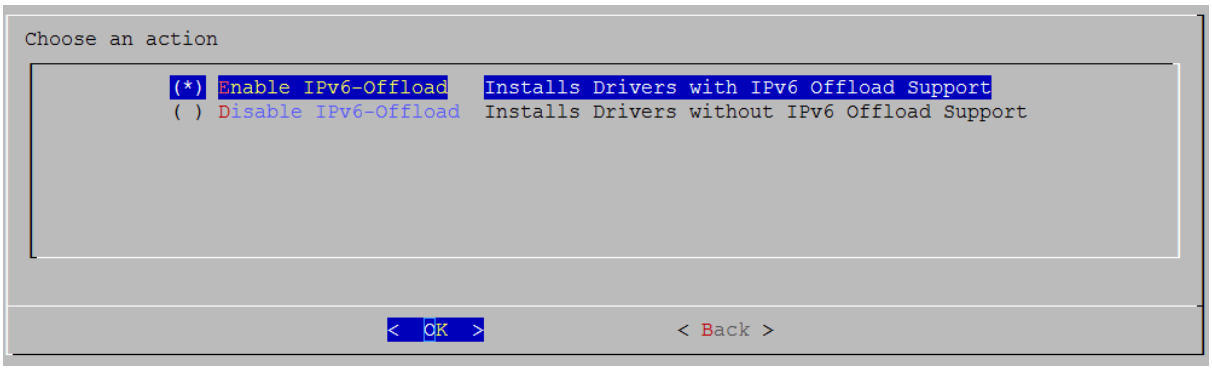

viii. Select the required configuration tuning option:

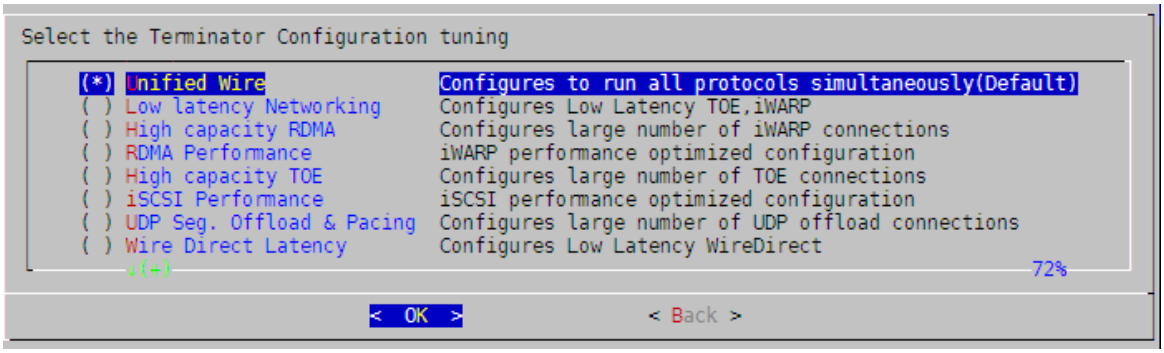

Note *The tuning options may vary depending on the Linux distribution.*

ix. Under “Choose install components”, select “all” to install all the related components for the option chosen in step (viii) or select “custom” to install specific components.

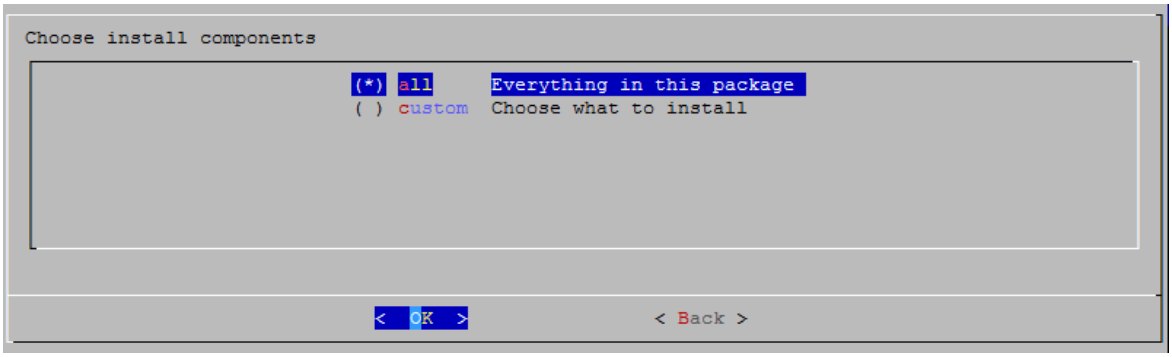

Important *To install Crypto Offload, WD-TOE, OVS, DPDK drivers and benchmark tools, please select “custom option”.*

x. Select the required performance tuning option.
   a. Enable Binding IRQs to CPUs: Bind MSI-X interrupts to different CPUs and disable IRQ balance daemon.
   b. *Retain IRQ balance daemon*: Do not disable IRQ balance daemon.
   c. *TX-Coalasce*: Write tx_coal=2 to modprobe.d/conf.

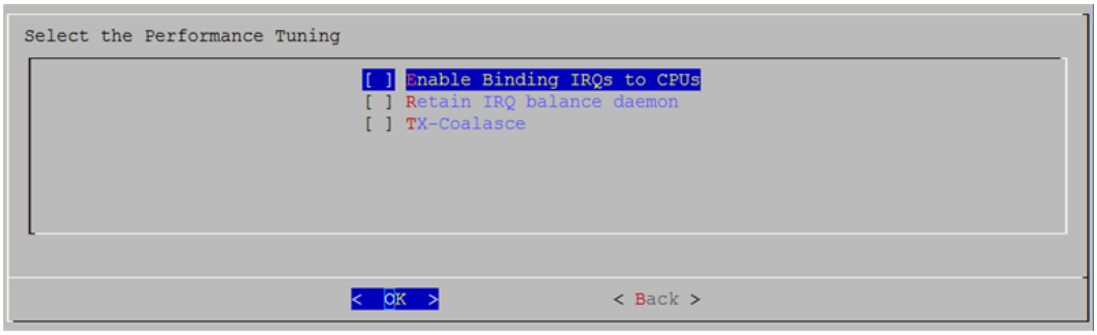

Note *For more information on the Performance tuning options, please refer to* Performance Tuning *section of the* ***Network (NIC/TOE)*** *chapter.*

xi. If you already have the required version of OFED software installed, select *Skip-OFED*. To install OFED-3.18-3 choose the *Install-OFED* option.

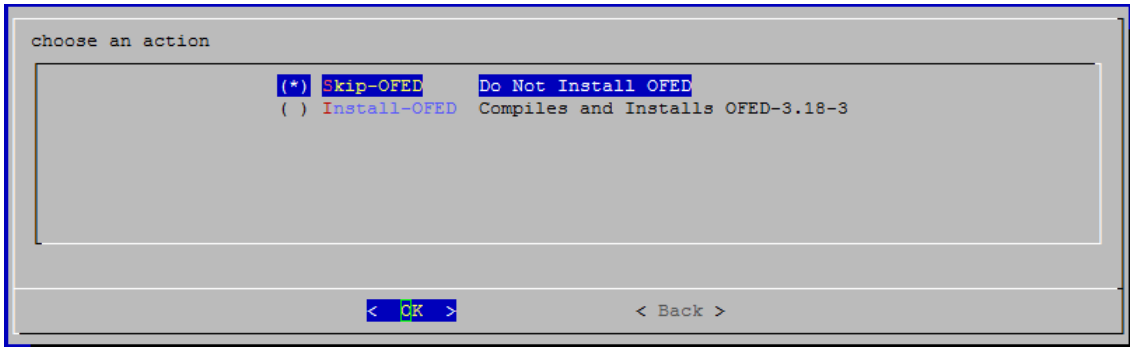

Note *This step will be prompted only for OFED supported platforms.*

xii. The selected components will now be installed:

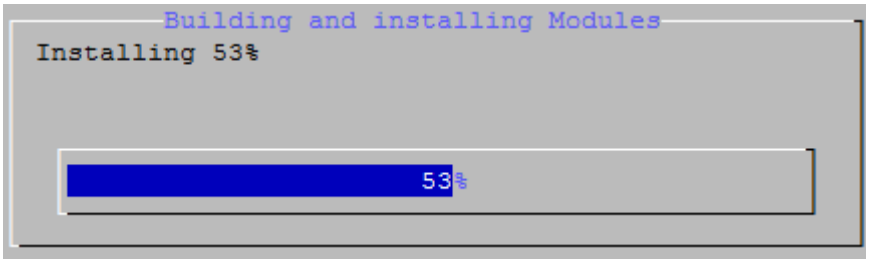

xiii. After successful installation, summary of installed components will be displayed.

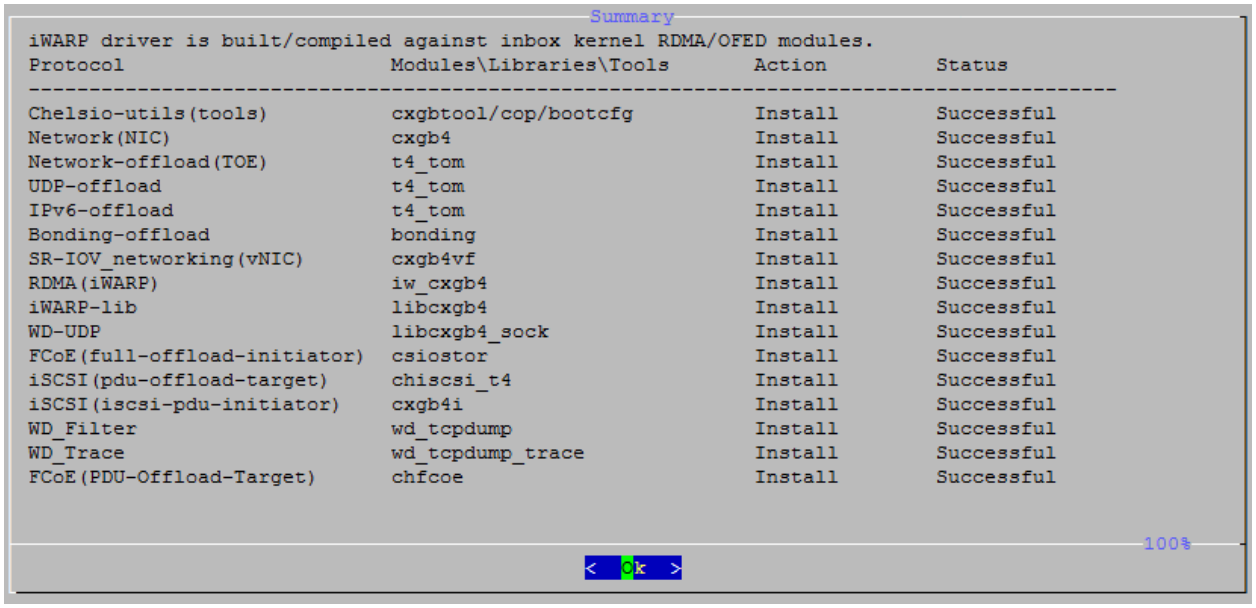

xiv. Select “View log” to view the installation log or “Exit” to continue.

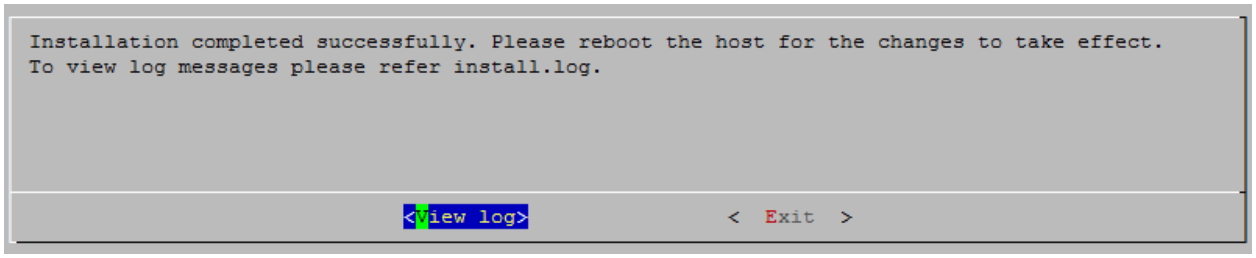

xv. Select “Yes” to exit the installer or “No” to go back.

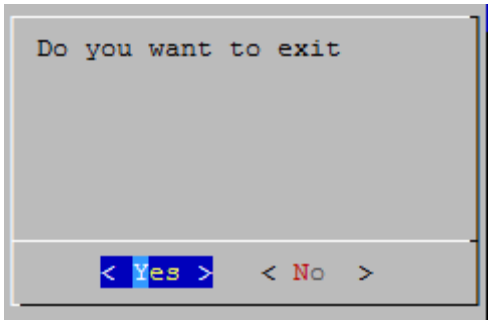

xvi. Reboot your machine for changes to take effect.

Note *Press Esc or Ctrl+C to exit the installer at any point of time.*

3.5.3. CLI mode (without Dialog utility)

If your system does not have **Dialog** or you choose not to install it, follow the steps mentioned below to install the Unified Wire package:

i. Download the Unified Wire driver package from Chelsio Download Center.
ii. Untar the tarball using the following command:

```
[root@host~]# tar zxvf <driver_package>.tar.gz
```

iii. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

iv. Run the following script to start the installer

```
[root@host~]# ./install.py -c <target>
```

v. Enter the number corresponding to the Configuration tuning option in the Input field and press Enter.
vi. If you already have the required version of OFED software installed, select Skip-OFED. To install OFED-3.18-3 choose the *Install-OFED* option.

Note *This step will be prompted only for OFED supported platforms.*

vii. The selected components will now be installed.
After successful installation you can press 1 to view the installation log. Press any other key to exit from the installer.

Important *To customize the installation, view the help by typing* `[root@host~]#./install.py -h`

viii. Reboot your machine for changes to take effect.

- **iWARP driver installation on Cluster nodes**

Important *Please make sure that you have enabled password less authentication with ssh on the peer nodes for this feature to work.*

Chelsio's Unified Wire package allows installing iWARP drivers on multiple Cluster nodes with a single command. Follow the procedure mentioned below:

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Create a file (*machinefilename*) containing the IP addresses or hostnames of the nodes in the cluster. You can view the sample file, *sample_machinefile,* provided in the package to view the format in which the nodes have to be listed.

iii. Now, execute the following command:

```
[root@host~]# ./install.py -C  -m <machinefilename>
```

iv. Select the required configuration tuning option. The tuning options may vary depending on the Linux distribution.

v. Select the required Cluster Configuration.

vi. If you already have the required version of OFED software installed, select *Skip-OFED*. To install OFED-3.18-3 choose the *Install-OFED* option.

vii. The selected components will now be installed.

The above commands will install iWARP (*iw_cxgb4*) and TOE (*t4_tom*) drivers on all the nodes listed in the *machinefilename* file.

3.5.4. CLI mode

i. Download the Unified Wire driver package from Chelsio Download Center.

ii. Untar the tarball using the following command:

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

iii. Change your current working directory to Chelsio Unified Wire package directory and build the source:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
[root@host~]# make
```

iv. Install the drivers, tools and libraries using the following command:

```
[root@host~]# make install
```

v. The default configuration tuning option is *Unified Wire.* The configuration tuning can be selected using the following commands:

```
[root@host~]# make CONF=<configuration_tuning>
[root@host~]# make CONF=<configuration_tuning> install
```

Important *Steps (iii) and (iv) mentioned above will NOT install Crypto, DCB, WD-TOE, OVS, DPDK drivers, and benchmark tools. They will have to be installed manually.*

Please refer to section CLI mode (individual drivers) *for instructions on installing them.*

Note *To view the different configuration tuning options, view help by typing* `[root@host~]# make help`

vi. Reboot your machine for changes to take effect.

3.5.5. CLI mode (individual drivers)

You can also choose to install drivers/features individually. Provided here are steps to build and install some of them. For the complete list, view help by running `make help`.

Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

- To build and install all drivers without IPv6 support:

```
[root@host~]# make ipv6_disable=1
[root@host~]# make ipv6_disable=1 install
```

- The default configuration tuning option is *Unified Wire.* The configuration tuning can be selected using the following commands:

```
[root@host~]# make CONF=<configuration_tuning> <Build Target>
[root@host~]# make CONF=<configuration_tuning> <Install Target>
```

- To build and install drivers along with benchmarks:

```
[root@host~]# make BENCHMARKS=1
[root@host~]# make BENCHMARKS=1 install
```

- The drivers will be installed as RPMs or Debian packages (for ubuntu). To skip this and install drivers:

```
[root@host~]# make SKIP_RPM=1 install
```

- The installer will remove the Chelsio specific drivers (inbox/outbox) from *initramfs*. To skip this and install drivers:

```
[root@host~]# make SKIP_INIT=1 install
```

- The installer will check for the required dependency packages and will install them if they are missing from the machine. To skip this and install drivers:

```
[root@host~]# make SKIP_DEPS=1 install
```

Note *To view the different configuration tuning options, view the help by typing* `[root@host~]#make help`

Note *If IPv6 is administratively disabled in the machine, the drivers will be built and installed without IPv6 Offload support by default.*

3.6. Installing Chelsio Unified Wire from RPM

Note
- *IPv6 should be enabled in the machine to use the RPM Packages.*
- *Drivers installed from RPM Packages do not have DCB support.*

3.6.1. Menu Mode

i. Download the tarball specific to your operating system and architecture from Chelsio Download Center.

ii. Untar the tarball:
E.g., for RHEL 6.9, untar using the following command:

```
[root@host~]# tar zxvf <driver_package>-RHEL6.9_x86_64.tar.gz
```

iii. Change your current working directory to Chelsio Unified Wire package directory

```
[root@host~]# cd ChelsioUwire-x.x.x.x-<OS>-<arch>
```

iv. Install Unified Wire:

```
[root@host~]# ./install.py
```

v. Select the Installation type as described below. Enter the corresponding number in the Input field and press Enter.
   a. *Unified Wire*: Install all the drivers in the Unified Wire software package. This option will not install OFED and drivers built against OFED.
   b. *Wire Direct Latency*: Install Wire Direct Latency drivers needed for Low latency applications.
   c. *Custom*: Customize the installation. Use this option to install drivers/software and related components (like OFED-3.18-3) as per the tuning option selected.
   d. *EXIT*: Exit the installer.

Note *The Installation options may vary depending on the Configuration tuning option selected.*

vi. The selected components will now be installed.
vii. Reboot your machine for changes to take effect.

Note *If the installation aborts with the message "Resolve the errors/dependencies manually and restart the installation", please go through the install.log to resolve errors/dependencies and then start the installation again.*

3.6.2. CLI mode

i. Download the tarball specific to your operating system and architecture from Chelsio Download Center.

ii. Untar the tarball:

E.g., for RHEL 6.9, untar using the following command:

```
[root@host~]# tar zxvf ChelsioUwire-x.x.x.x-RHEL6.9_x86_64.tar.gz
```

iii. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x-<OS>-<arch>
```

iv. Install Unified Wire:

```
 [root@host~]# ./install.py -i <nic_toe/all/udpso/wd/crypto/ovs>
```

Here,

nic_toe	: NIC and TOE drivers only
all	: All Chelsio drivers built against inbox OFED
udpso	: UDP segmentation offload capable NIC and TOE drivers only
wd	: Wire Direct drivers and libraries only
crypto	: Crypto drivers and Chelsio Openssl modules.
ovs	: OVS modules and NIC driver.

Note *The installation options may vary depending on the Linux distribution.*

v. The default configuration tuning option is *Unified Wire*. The configuration tuning can be selected using the following command:

```
[root@host~]# ./install.py –i <Installation mode> -c <configuration_tuning>
```

Note *To view the different configuration tuning options, view the help by typing* `[root@host~]# ./install.py –h`

vi. To install OFED and Chelsio drivers built against OFED, run the above command with *-o* option.

```
[root@host~]# ./install.py -i <Installation mode> -c <Configuration> -o
```

vii. Reboot your machine for changes to take effect.

- ### iWARP driver installation on cluster nodes

Important *Please make sure that you have enabled password less authentication with ssh on the peer nodes for this feature to work.*

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x-<OS>-<arch>
```

ii. Create a file (*machinefilename*) containing the IP addresses or hostnames of the nodes in the cluster. You can view the sample file, *sample_machinefile,* provided in the package to view the format in which the nodes have to be listed.

iii. Now, execute the following command:

```
[root@host~]# ./install.py -C  -m <machinefilename> -i
<nic_toe/all/udpso/wd> -c <configuration_tuning> -o
```

Here, -o parameter will install OFED and Chelsio drivers built against OFED.

The above command will install iWARP (*iw_cxgb4*) and TOE (*t4_tom*) drivers on all the nodes listed in the <machinefilename> file.

iv. Reboot your machine for changes to take effect.

3.7. Firmware Update

The 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 kernel must be configured to enable userspace firmware loading support:

Device Drivers -> Generic Driver Options -> Userspace firmware loading support

The firmware version can be verified using *ethtool*:

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

4. Configuring Chelsio Network Interfaces

To test Chelsio adapters' features it is required to use two machines both with Chelsio's network adapters installed. These two machines can be connected directly without a switch (back-to-back), or both connected to a switch. The interfaces have to be declared and configured. The configuration files for network interfaces on Red Hat Enterprise Linux (RHEL) distributions are kept under `/etc/sysconfig/network-scripts`.

Note *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.*

4.1. Configuring 40G adapters

Chelsio T5 40G adapters can be configured in the following three modes:

i. 2X40Gbps: This is the default mode of operation where each port functions as 40Gbps link. The port nearest to the motherboard will appear as the first network interface (Port 0).
ii. 4X10Gbps: In this mode, port 0 functions as 4 10Gbps links and port 1 is disabled.
iii. QSA: This mode adds support for QSA (QSFP to SFP+) modules, enabling smooth, cost-effective, connections between 40 Gigabit Ethernet adapters and 1 or 10 Gigabit Ethernet networks using existing SFP+ based cabling. The port farthest from the motherboard will appear as the first network interface (Port 0).

To configure/change the mode of operation, use the following procedure:

i. Unload all Chelsio drivers using the *rmmod* command:

```
[root@host~]# rmmod <chelsio_driver>
```

ii. Run the *chelsio_adapter_config* command to detect all T5 40G adapter(s) present in the system.

```
[root@host~]# chelsio_adapter_config
Chelsio T580 card detected
Chelsio T580 PCI devices :
|------------------------------|
| 1 T580-LP-CR    01:00.0      |
| 2 T580-CR       03:00.0      |
| 3 T580-LP-SO-CR 04:00.0      |
|------------------------------|
```

iii. Select the adapter to configure by specifying the adapter index.

iv. Select the required mode:

```
Possible T580 adapter modes:
|------------------------------|
| 1: 2x40G                     |
| 2: 4x10G                     |
| 3: QSA                       |
|------------------------------|
Select mode for adapter (1,2,3):
```

v. Reload the network driver for changes to take effect.

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

Note *In case of T580-SO-CR adapters, reboot the machine for changes to take effect.*

4.2. Configuring network-scripts

A typical interface network-script (e.g., eth0) on RHEL 6.X looks like the following:

```
# file: /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE="eth0"
HWADDR=00:30:48:32:6A:AA
ONBOOT="yes"
NM_CONTROLLED="no"
BOOTPROTO="static"
IPADDR=10.192.167.111
NETMASK=255.255.240.0
```

Note *On earlier versions of RHEL the NETMASK attribute is named IPMASK. Make sure you are using the right attribute name.*

In the case of DHCP addressing the last two lines should be removed and `BOOTPROTO="static"` should be changed to `BOOTPROTO="dhcp"`

The `ifcfg-ethX` files have to be created manually. They are required for bringing the interfaces up and down and attribute the desired IP addresses.

4.3. Creating network-scripts

To spot the new interfaces, make sure the driver is unloaded first. To that point `ifconfig -a | grep HWaddr` should display all non-chelsio interfaces whose drivers are loaded, whether the interfaces are up or not.

```
[root@host~]# ifconfig -a | grep HWaddr
eth0 Link encap:Ethernet HWaddr 00:30:48:32:6A:AA
```

Then load the driver using the `modprobe cxgb4` command (for the moment it does not make any difference whether we are using NIC-only or the TOE-enabling driver). The output of ifconfig should display the adapter interfaces as:

```
[root@host~]# ifconfig -a | grep HWaddr
eth0 Link encap:Ethernet HWaddr 00:30:48:32:6A:AA
eth1 Link encap:Ethernet HWaddr 00:07:43:04:6B:E9
eth2 Link encap:Ethernet HWaddr 00:07:43:04:6B:F1
```

For each interface you can write a configuration file in `/etc/sysconfig/network-scripts`. The `ifcfg-eth1` could look like:

```
# file: /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE="eth1"
HWADDR=00:07:43:04:6B:E9
ONBOOT="no"
NM_CONTROLLED="no"
BOOTPROTO="static"
IPADDR=10.192.167.112
NETMASK=255.255.240.0
```

From now on, the `eth1` interface of the adapter can be brought up and down through the `ifup eth1` and `ifdown eth1` commands respectively. Note that it is of course not compulsory to create a configuration file for every interface if you are not planning to use them all.

4.4. Checking Link

Once the network-scripts are created for the interfaces you should check the link i.e. make sure it is actually connected to the network. First, bring up the interface you want to test using `ifup eth1`.

You should now be able to ping any other machine from your network provided it has ping response enabled.

5. Performance Tuning

- In order to auto tune the system for best performance, Chelsio recommends: Disabling virtualization, c-state technology, VT-d, Intel I/O AT and SR-IOV in the BIOS settings
- Installing the adapter into a PCIe Gen3 x8/x16 slot. Ensure that T6 100G adapters are placed in x16 slots and not in x8_in_x16 slots.
- Add *intel_pstate=disable processor.max_cstate=1 intel_idle.max_cstate=0* to the kernel command line to prevent the system from entering power-saving/idle states and avoid CPU frequency changes.
- Set the below tuned-adm profile for RHEL platforms.

```
[root@host~]# tuned-adm profile throughput-performance
```

- Turn off *irqbalance*

```
[root@host~]# /etc/init.d/irqbalance stop
```

On RHEL7.X platforms, use the below command:

```
[root@host~]# systemctl stop irqbalance.service
```

- To lower latency:

i. Follow all the tuning instructions mentioned above.
ii. Disable SELinux
iii. Add *idle=poll* to the kernel command line.
iv. Set the below tuned-adm profile for RHEL7 platforms.

```
[root@host~]# tuned-adm profile network-latency
```

v. Disable few services.

```
[root@host~]# t4_latencytune.sh <interface>
```

vi. Set sysctl param *net.ipv4.tcp_low_latency* to 1

```
[root@host~]# sysctl -w net.ipv4.tcp_low_latency=1
```

To optimize your system for different protocols, please refer to their respective chapters.

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 Download Center for regular updates on various software/drivers. You can also subscribe to our newsletter for the latest software updates.

7. Software/Driver Uninstallation

Similar to installation, the Chelsio Unified Wire package can be uninstalled using two main methods: from the source and RPM, based on the method used for installation. If you decide to use source, you can uninstall the package using CLI or GUI mode.

7.1. Uninstalling Chelsio Unified Wire from source

7.1.1. GUI mode (with Dialog utility)

i. Change your current working directory to Chelsio Unified Wire package directory and run the following script to start the GUI installer:

```
[root@host~]# ./install.py
```

ii. Select “uninstall” , Under “Choose an action”

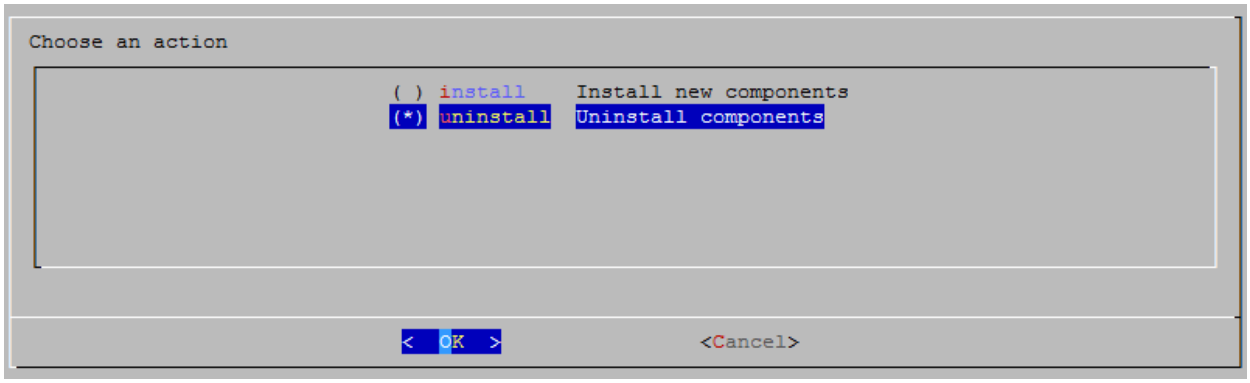

iii. Select “all” to uninstall all the installed drivers, libraries and tools or select “custom” to remove specific components.

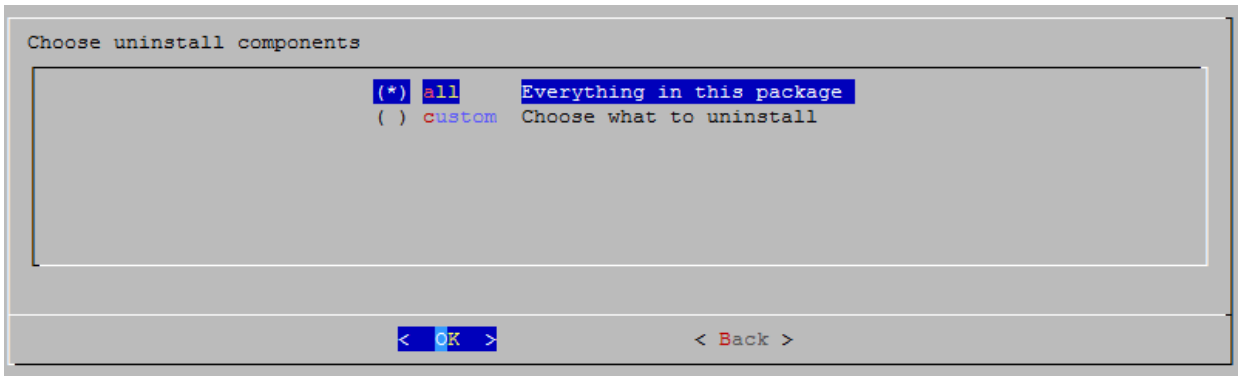

iv. The selected components will now be uninstalled.

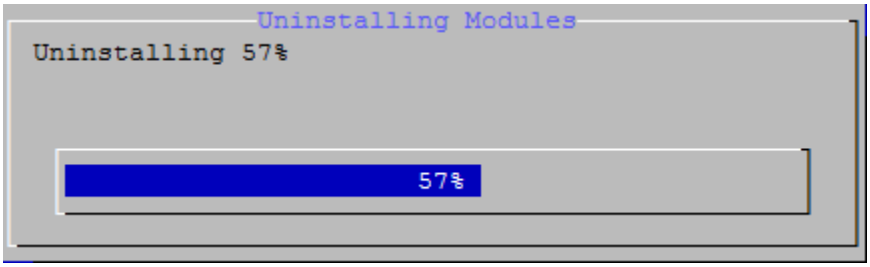

v. After successful uninstalltion, summary of the uninstalled components will be displayed.

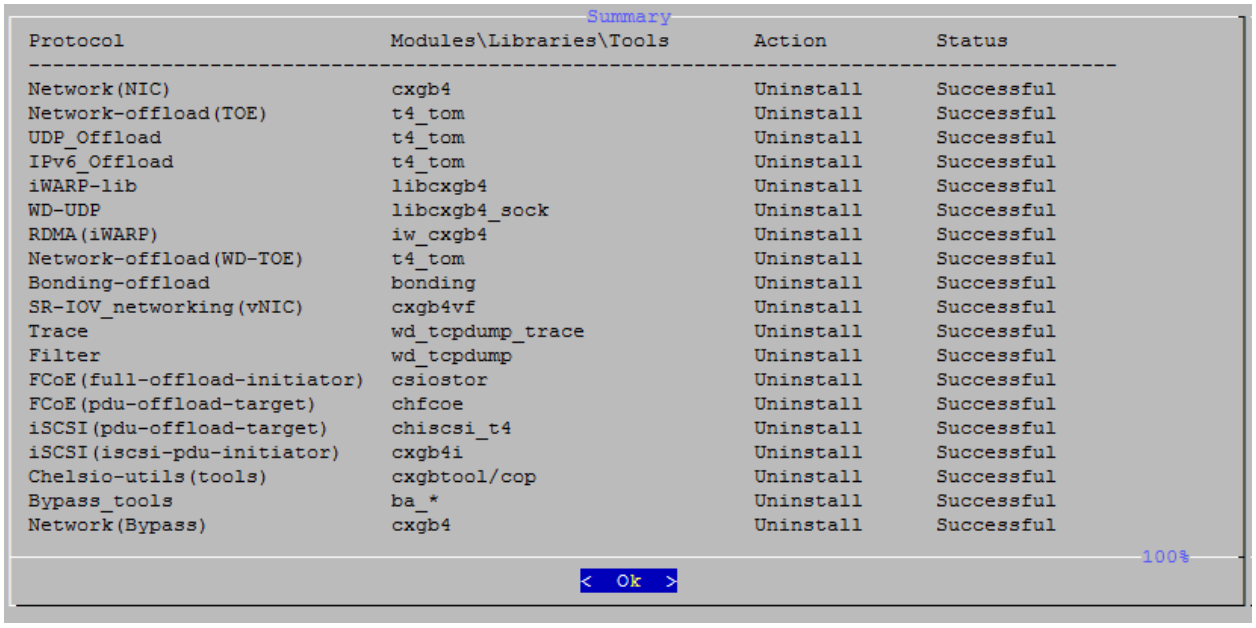

vi. Select “View log” to view uninstallation log or “Exit” to continue.

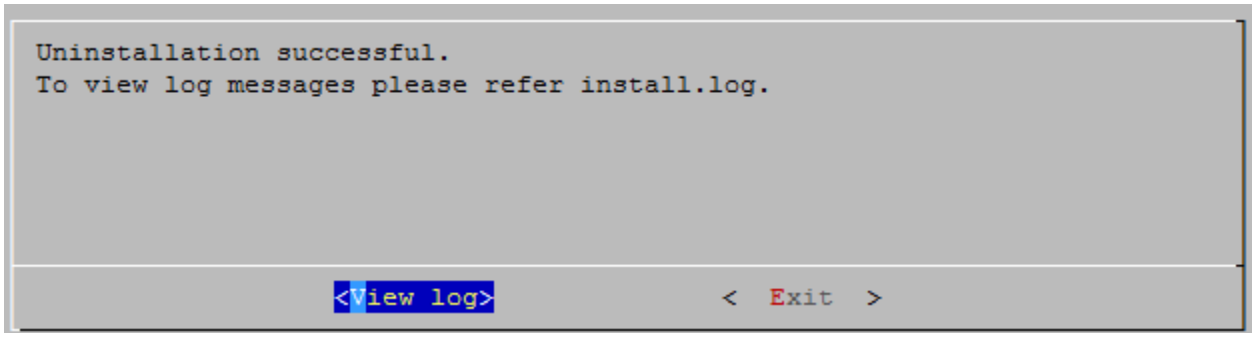

vii. Select “Yes” to exit the installer or “No” to go back.

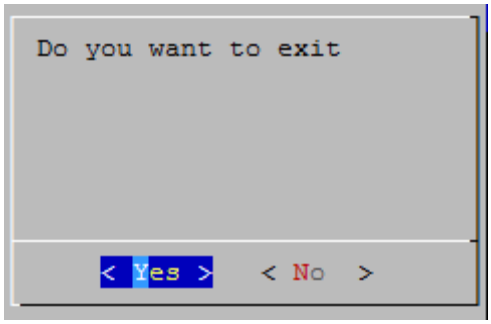

Note *Press Esc or Ctrl+C to exit the installer at any point of time.*

7.1.2. CLI mode (without Dialog utility)

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Run the following script with `-u` option to uninstall the Unified Wire Package:

```
[root@host~]# ./install.py -u <target>
```

Note *View help by typing* `[root@host~]# ./install.py -h` for more information

7.1.3. iWARP driver uninstallation on Cluster nodes

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Uninstall iWARP drivers on multiple Cluster nodes using:

```
[root@host~]# ./install.py -C -m <machinefilename> -u all
```

The above command will remove Chelsio iWARP (*iw_cxgb4*) and TOE (*t4_tom*) drivers from all the nodes listed in the *machinefilename* file.

7.1.4. CLI mode

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Uninstall using the following command:

```
[root@host~]# make uninstall
```

7.1.5. CLI mode (individual drivers/software)

You can also choose to uninstall drivers/software individually. Provided here are steps to uninstall few of them. For the complete list, view help by running `make help`

Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

- To uninstall NIC driver:

```
[root@host~]# make nic_uninstall
```

- To uninstall offload driver:

```
[root@host~]# make toe_uninstall
```

- To uninstall iWARP driver:

```
[root@host~]# make iwarp_uninstall
```

- To uninstall UDP Segmentation Offload driver:

```
[root@host~]# make udp_offload_uninstall
```

7.2. Uninstalling Chelsio Unified Wire from RPM

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x-<OS>-<arch>
```

ii. Uninstall Unified Wire:

```
[root@host~]# ./uninstall.py <inbox/ofed>
```

inbox : for removing all Chelsio drivers.
ofed : for removing OFED and Chelsio drivers.

Note *The uninstallation options may vary depending on Linux distribution. View help by typing* `[root@host~]# ./uninstall.py -h` *for more information.*

7.2.1. iWARP driver uninstallation on Cluster nodes

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x-<OS>-<arch>
```

ii. Uninstall iWARP drivers on multiple Cluster nodes using:

```
[root@host~]# ./install.py -C -m <machinefilename> -u
```

The above command will remove Chelsio iWARP (*iw_cxgb4*) and TOE (*t4_tom*) drivers from all the nodes listed in the *machinefilename* file.

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 - 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 Terminator engine provides unparalleled performance through a specialized data flow processor implementation and a host of features designed for high throughput and low latency in demanding conditions and networking environments.

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 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 (*Memory-free; limited number of offload connections supported*)
- T580-OCP-SO*
- T520-OCP-SO*
- T520-BT
- T580-CR
- T580-SO-CR*
- T580-LP-CR
- T520-LL-CR
- T520-SO-CR*
- T520-CR
- T540-CR
- T420-CR
- T440-CR
- T422-CR

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

*Only NIC driver supported

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the Network driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

- To build and install NIC only driver (without offload support):

```
[root@host~]# make nic_install
```

- To build and install drivers with offload support:

```
[root@host~]# make toe_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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

3.1. Loading in NIC mode (without full offload support)

To load the Network driver without full offload support, run the following command:

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

3.2. Loading in TOE mode (with full offload support)

To enable full offload support, run the following command:

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

Note *Offload support needs to be enabled upon each reboot of the system. This can be done manually as shown above.*

In VMDirect Path environment, it is recommended to load the offload driver using the following command:

```
[root@host~]# modprobe t4_tom vmdirectio=1
```

4. Software/Driver Configuration

4.1. Disabling Auto-negotiation

Auto-negotiation is enabled by default for T6 adapters. Few 100G switches still have open bugs around this functionality. If you face any link issue with such switches, disable this feature using the command below:

```
[root@host~]# ethtool -s ethX autoneg off
```

4.2. Enabling TCP Offload

Load the offload drivers and bring up the Chelsio interface.

```
[root@host~]# modprobe t4_tom
[root@host~]# ifconfig ethX <IP> up
```

All TCP traffic will be offloaded over the Chelsio interface now. To see the number of connections offloaded, run the following command:

```
[root@host~]# cat /sys/kernel/debug/cxgb4/<bus-id>/tids
```

```
[root@ ~]# cat /sys/kernel/debug/cxgb4/0000\:01\:00.4/tids
Connections in use: 8
TID range: 0..959/2048..18431, in use: 0/8
STID range: 960..1455, in use-IPv4/IPv6: 0/0
ATID range: 0..8191, in use: 0
FTID range: 1472..1967
UOTID range: 18432..19455, in use: 0
HW TID usage: 8 IP users, 0 IPv6 users
```

Where,

TID is the number of offload connections.
STID is the number of offload servers.

T6225-SO-CR adapter supports limited number of offload connections (256 IPv4/128 IPv6). Here is a sample output:

```
[root@ ~]# cat /sys/kernel/debug/cxgb4/0000\:03\:00.4/tids
Connections in use: 256
TID range: 0..255, in use: 256
STID range: 256..319, in use-IPv4/IPv6: 0/0
ATID range: 0..127, in use: 0
FTID range: 320..815
HW TID usage: 256 IP users, 0 IPv6 users
```

4.3. Enabling Busy waiting

Busy waiting/polling is a technique where a process repeatedly checks to see if an event has occurred, by spinning in a tight loop. By making use of similar technique, Linux kernel provides the ability for the socket layer code to poll directly on an Ethernet device's Rx queue. This eliminates the cost of interrupts and context switching, and with proper tuning allows to achieve latency performance similar to that of hardware.

Chelsio's NIC and TOE drivers support this feature and can be enabled on Chelsio supported devices to attain improved latency.

To make use of BUSY_POLL feature, follow the steps mentioned below:

i. Enable BUSY_POLL support in kernel config file by setting `CONFIG_NET_RX_BUSY_POLL=y`
ii. Enable BUSY_POLL globally in the system by setting the values of following sysctl parameters depending on the number of connections:

```
sysctl -w net.core.busy_read=<value>
sysctl -w net.core.busy_poll=<value>
```

Set the values of the above parameters to 50 for 100 or less connections; and 100 for more than 100 connections.

Note *BUSY_POLL can also be enabled on a per-connection basis by making use of SO_BUSY_POLL option in the socket application code. Refer socket man-page for more details.*

4.4. Precision Time Protocol (PTP)

Precision Time Protocol (PTP) standard defines a protocol for precise synchronization of clock between master and slave devices in a local area network. It can provide timing accuracies in nanosecond units. The protocol is based on time stamping and measuring the send and receive times. Most of the implementation relies on time stamping of the packets in the software which reduces the accuracy of the time measured. One possible solution to this problem is time stamping the packet in the NIC hardware itself.

Chelsio's Terminator hardware provides many features to support PTP implementations:

- High precision timers which can be read through PIO registers.
- Wall clock time based on the time of the day.
- Time stamping of selected PTP packets on both ingress and egress direction.

Important *This feature is currently supported on kernel 4.9.49 only.*

4.4.1. Synchronizing Clocks

ptp4l tool (installed during Unified Wire installation) is used to synchronise clocks:

i. Load the network driver on all master and slave nodes;

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

ii. Assign IP addresses and ensure that master and slave nodes are connected.
iii. Start the *ptp4l* tool on master using the Chelsio interface:

```
[root@host~]# ptp4l -i <interface> -H -m
```

```
[root@ ~]# ptp4l -i enp1s0f4 -H -m
ptp4l[16681.046]: selected /dev/ptp4 as PTP clock
ptp4l[16681.054]: port 1: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[16681.054]: port 0: INITIALIZING to LISTENING on INIT_COMPLETE
ptp4l[16681.055]: port 1: link up
ptp4l[16688.483]: port 1: LISTENING to MASTER on ANNOUNCE_RECEIPT_TIMEOUT_EXPIRES
ptp4l[16688.483]: selected best master clock 000743.fffe.293be0
ptp4l[16688.483]: assuming the grand master role
```

iv. Start the *ptp4l* tool on slave nodes:

```
[root@host~]# ptp4l -i <interface> -H -m -s
```

Note *To view the complete list of available options, refer ptp4l help manual.*

```
[root@ ~]# ptp4l -i enp6s0f4 -m -H -s
ptp4l[13393.931]: selected /dev/ptp3 as PTP clock
ptp4l[13393.939]: port 1: INITIALIZING to LISTENING on INITIALIZE
ptp4l[13393.940]: port 0: INITIALIZING to LISTENING on INITIALIZE
ptp4l[13394.360]: port 1: new foreign master 000743.fffe.293be0-1
ptp4l[13398.360]: selected best master clock 000743.fffe.293be0
ptp4l[13398.360]: port 1: LISTENING to UNCALIBRATED on RS_SLAVE
ptp4l[13399.362]: master offset       5597 s0 freq  -26920 path delay       159
ptp4l[13400.362]: master offset       5643 s2 freq  -26874 path delay       159
ptp4l[13400.363]: port 1: UNCALIBRATED to SLAVE on MASTER_CLOCK_SELECTED
ptp4l[13401.362]: master offset       5516 s2 freq  -21358 path delay       159
ptp4l[13402.362]: master offset      -7045 s2 freq  -32264 path delay      3370
ptp4l[13403.362]: master offset       1897 s2 freq  -25436 path delay      3370
ptp4l[13404.362]: master offset        992 s2 freq  -25772 path delay      1801
ptp4l[13405.362]: master offset        398 s2 freq  -26068 path delay       427
ptp4l[13406.362]: master offset      -1038 s2 freq  -27385 path delay       395
ptp4l[13407.362]: master offset       -248 s2 freq  -26906 path delay       318
ptp4l[13408.362]: master offset       -209 s2 freq  -26941 path delay       318
ptp4l[13409.362]: master offset        -92 s2 freq  -26887 path delay       237
ptp4l[13410.362]: master offset       -233 s2 freq  -27056 path delay       237
ptp4l[13411.362]: master offset        -40 s2 freq  -26933 path delay       237
ptp4l[13412.363]: master offset        -10 s2 freq  -26915 path delay       237
ptp4l[13413.363]: master offset         93 s2 freq  -26815 path delay       175
ptp4l[13414.363]: master offset        -46 s2 freq  -26926 path delay       175
ptp4l[13415.363]: master offset         -3 s2 freq  -26897 path delay       169
ptp4l[13416.363]: master offset       -144 s2 freq  -27038 path delay       169
ptp4l[13417.363]: master offset         61 s2 freq  -26877 path delay       169
ptp4l[13418.363]: master offset        -68 s2 freq  -26987 path delay       171
ptp4l[13419.363]: master offset        -32 s2 freq  -26972 path delay       171
ptp4l[13420.363]: master offset          3 s2 freq  -26946 path delay       176
ptp4l[13421.363]: master offset         50 s2 freq  -26898 path delay       174
```

v. Synchronize the system clock to a PTP hardware clock (PHC) on slave nodes.

```
[root@host~]# phc2sys -s <interface> -c CLOCK_REALTIME -w -m
```

```
[root@ ~]# phc2sys -s enp6s0f4 -c CLOCK_REALTIME -w -m
phc2sys[13406.672]: phc offset 36493387467 s0 freq  +29332 delay 1086689
phc2sys[13407.679]: phc offset 36493338184 s1 freq  -19723 delay 1084486
phc2sys[13408.684]: phc offset       1346 s2 freq  -18377 delay 1085776
phc2sys[13409.690]: phc offset      -2051 s2 freq  -21370 delay 1077105
phc2sys[13410.695]: phc offset       2070 s2 freq  -17864 delay 1078811
phc2sys[13411.701]: phc offset       5037 s2 freq  -14276 delay 1085488
phc2sys[13412.707]: phc offset      -2483 s2 freq  -20285 delay 1078173
phc2sys[13413.712]: phc offset        171 s2 freq  -18376 delay 1079112
phc2sys[13414.718]: phc offset        949 s2 freq  -17547 delay 1079990
phc2sys[13415.723]: phc offset      -1293 s2 freq  -19504 delay 1076567
phc2sys[13416.729]: phc offset     -15208 s2 freq  -33807 delay 1045916
phc2sys[13417.735]: phc offset      15711 s2 freq   -7450 delay 1076534
phc2sys[13418.740]: phc offset       5199 s2 freq  -13249 delay 1076439
phc2sys[13419.746]: phc offset       2017 s2 freq  -14871 delay 1080000
```

4.5. Performance Tuning

Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.

- **TOE**

i. Run the performance tuning script to map TOE queues to different CPUs:

```
[root@host~]# t4_perftune.sh -n -Q ofld
```

ii. Set the following *sysctl* parameter:

```
[root@host~]# sysctl -w toe.toe0_tom.delayed_ack=3
[root@host~]# sysctl -w net.ipv4.tcp_timestamps=0
[root@host~]# sysctl -w net.core.netdev_max_backlog=250000
[root@host~]# sysctl -w net.core.rmem_max=4194304
[root@host~]# sysctl -w net.core.wmem_max=4194304
[root@host~]# sysctl -w net.core.rmem_default=4194304
[root@host~]# sysctl -w net.core.wmem_default=4194304
[root@host~]# sysctl -w net.ipv4.tcp_rmem="4096 1048576 4194304"
[root@host~]# sysctl -w net.ipv4.tcp_wmem="4096 1048576 4194304"
```

For 100G performance, disable Nagle using the following steps:

a. Create a COP policy:

```
[root@host~]# cat <policy_file>
all=>offload !nagle
```

b. Compile the policy:

```
[root@host~]# cop -d -o <policy_out> <policy_file>
```

c. Apply the policy:

```
[root@host~]# cxgbtool ethX policy <policy_out>
```

- **NIC**

i. Run the performance tuning script to map NIC queues to different CPUs:

```
[root@host~]# t4_perftune.sh -n -Q nic
```

ii. Enable adaptive-rx

```
[root@host~]# ethtool -C enp2s0f4 adaptive-rx on
```

iii. Set the following sysctl parameters

```
[root@host~]# sysctl -w net.ipv4.tcp_timestamps=0
[root@host~]# sysctl -w net.core.netdev_max_backlog=250000
[root@host~]# sysctl -w net.core.rmem_max=4194304
[root@host~]# sysctl -w net.core.wmem_max=4194304
[root@host~]# sysctl -w net.core.rmem_default=4194304
[root@host~]# sysctl -w net.core.wmem_default=4194304
[root@host~]# sysctl -w net.ipv4.tcp_rmem="4096 1048576 4194304"
[root@host~]# sysctl -w net.ipv4.tcp_wmem="4096 1048576 4194304"
```

- **NIC/TOE Latency**

Enable BUSY_POLL feature:

```
[root@host~]# sysctl -w net.core.busy_poll = 50
[root@host~]# sysctl -w net.core.busy_read = 50
```

- **Receiver Side Scaling (RSS)**

Receiver Side Scaling enables the receiving network traffic to scale with the available number of processors on a modern networked computer. RSS enables parallel receive processing and dynamically balances the load among multiple processors. Chelsio's network controller fully supports Receiver Side Scaling for IPv4 and IPv6.

This script first determines the number of CPUs on the system and then each receiving queue is bound to an entry in the system interrupt table and assigned to a specific CPU. Thus, each receiving queue interrupts a specific CPU through a specific interrupt now. For example, on a 4-core system, `t4_perftune.sh` gives the following output:

```
[root@host~]# t4_perftune.sh
Discovering Chelsio T4/T5 devices ...
Configuring Chelsio T4/T5 devices ...
Tuning eth7
IRQ table length 4
Writing 1 in /proc/irq/62/smp_affinity
Writing 2 in /proc/irq/63/smp_affinity
Writing 4 in /proc/irq/64/smp_affinity
Writing 8 in /proc/irq/65/smp_affinity
eth7 now up and tuned
...
```

Because there are 4 CPUs on the system, 4 entries of interrupts are assigned. For other network interfaces, you should see similar output message.

Now the receiving traffic is dynamically assigned to one of the system's CPUs through a Terminator queue. This achieves a balanced usage among all the processors. This can be verified, for example, by using the **iperf** tool. First set up a server on the receiver host:

```
[root@receiver_host~]# iperf –s
```

Then on the sender host, send data to the server using the iperf client mode. To emulate a moderate traffic workload, use *-P* option to request 20 TCP streams from the server:

```
[root@sender_host~]# iperf -c receiver_host_name_or_IP -P 20
```

Then on the receiver host, look at interrupt rate at `/proc/interrupts`:

```
[root@receiver_host~]# cat /proc/interrupts | grep eth6

Id      CPU0      CPU1      CPU2      CPU3       type              interface
36:   115229         0         0         1   PCI-MSI-edge     eth6 (queue 0)
37:        0    121083         1         0   PCI-MSI-edge     eth6 (queue 1)
38:        0         0    105423         1   PCI-MSI-edge     eth6 (queue 2)
39:        0         0         0    115724   PCI-MSI-edge     eth6 (queue 3)
```

Now interrupts from eth6 are evenly distributed among the 4 CPUs.

Without Terminator's RSS support, the interrupts caused by network traffic may be distributed unevenly over CPUs. For your information, the traffic produced by the same iperf commands gives the following output in `/proc/interrupts`.

```
[root@receiver_host~]# cat /proc/interrupts | grep eth6

Id      CPU0      CPU1      CPU2      CPU3       type              interface
36:        0         9         0     17418  PCI-MSI-edge    eth6 (queue 0)
37:        0         0     21718      2063  PCI-MSI-edge    eth6 (queue 1)
38:        0         7    391519       222  PCI-MSI-edge    eth6 (queue 2)
39:        1         0        33     17798  PCI-MSI-edge    eth6 (queue 3)
```

Here there are 4 receiving queues from the eth6 interface, but they are not bound to a specific CPU or interrupt entry. Queue 2 has caused a very large number of interrupts on CPU2 while CPU0 and CPU1 are barely used by any of the four queues. Enabling RSS is thus essential for best performance.

Note

Linux's `irqbalance` *may take charge of distributing interrupts among CPUs on a multiprocessor platform. However,* `irqbalance` *distributes interrupt requests from all hardware devices across processors. For a server with Chelsio network card constantly receiving large volume of data at 40/10Gbps, the network interrupt demands are significantly high. Under such circumstances, it is necessary to enable RSS to balance the network load across multiple processors and achieve the best performance.*

- **Interrupt Coalescing**

The idea behind Interrupt Coalescing (IC) is to avoid flooding the host CPUs with too many interrupts. Instead of throwing one interrupt per incoming packet, IC waits for 'n' packets to be available in the Rx queues and placed into the host memory through DMA operations before an interrupt is thrown, reducing the CPU load and thus improving latency. It can be changed using the following command:

```
[root@host~]# ethtool –C ethX rx-frames n
```

Note

For more information, run the following command:

```
[root@host~]# ethtool -h
```

- **Large Receive Offload / Generic Receive Offload**

Large Receive Offload or Generic Receive Offload is a performance improvement feature at the receiving side. LRO/GRO aggregates the received packets that belong to same stream, and combines them to form a larger packet before pushing them to the receive host network stack. By doing this, rather than processing every small packet, the receiver CPU works on fewer packet headers but with same amount of data. This helps reduce the receive host CPU load and improve throughput in a 40/10Gb network environment where CPU can be the bottleneck.

LRO and GRO are different names to refer to the same receiver packets aggregating feature. LRO and GRO actually differ in their implementation of the feature in the Linux kernel. The feature was first added into the Linux kernel in version 2.6.24 and named Large Receive Offload (LRO). However, LRO only works for TCP and IPv4. As from kernel 2.6.29, a new protocol-independent implementation removing the limitation is added to Linux, and it is named Generic Receive Offload (GRO). The old LRO code is still available in the kernel sources but whenever both GRO and LRO are presented GRO is always the preferred one to use.

Please note that if your Linux system has IP forwarding enabled, i.e. acting as a bridge or router, the LRO needs to be disabled. This is due to a known kernel issue.

Chelsio's card supports both hardware assisted GRO/LRO and Linux-based GRO/LRO. `t4_tom` is the kernel module that enables the hardware assisted GRO/LRO. If it is not already in the kernel module list, use the following command to insert it:

```
[root@host~]# lsmod | grep t4_tom
[root@host~]# modprobe t4_tom
[root@host~]# lsmod | grep t4_tom
t4_tom 88378 0 [permanent]
toecore 21618 1 t4_tom
cxgb4 225342 1 t4_tom
```

Then Terminator's hardware GRO/LRO implementation is enabled.

If you would like to use the Linux GRO/LRO for any reason, first the `t4_tom` kernel module needs to be removed from kernel module list. Please note you might need to reboot your system.

After removing the `t4_tom` module, you can use `ethtool` to check the status of current GRO/LRO settings, for example:

```
[root@host~]# ethtool -k eth6
Offload parameters for eth6:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: on
large-receive-offload: off
```

Now the `generic-receive-offload` option is on. This means GRO is enabled. Please note that there are two offload options here: `generic-receive-offload` and `large-receive-offload`. This is because on this Linux system (RHEL6.0), the kernel supports both GRO and LRO. As mentioned earlier, GRO is always the preferred option when both of them are present. On other systems LRO might be the only available option. Then `ethtool` could be used to switch LRO on and off as well.

When Linux's GRO is enabled, Chelsio's driver provides two GRO-related statistics. They are displayed using the following command:

```
[root@host~]# ethtool -S eth6
...
GROPackets : 0
GROMerged : 897723
...
```

`GROPackets` is the number of held packets. Those are candidate packets held by the kernel to be processed individually or to be merged to larger packets. This number is usually zero. `GROMerged` is the number of packets that merged to larger packets. Usually this number increases if there is any continuous traffic stream present.
`ethtool` can also be used to switch off the GRO/LRO options when necessary:

```
[root@host~]# ethtool -K eth6 gro off
[root@host~]# ethtool -k eth6
Offload parameters for eth6:
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp-segmentation-offload: on
udp-fragmentation-offload: off
generic-segmentation-offload: on
generic-receive-offload: off
large-receive-offload: off
```

The output above shows a disabled GRO.

5. Software/Driver Unloading

5.1. Unloading the NIC driver

To unload the NIC driver, run the following command:

```
[root@host~]# rmmod cxgb4
```

5.2. Unloading the TOE driver

A reboot is required to unload the TOE driver. To avoid rebooting, follow the steps mentioned below:

i. Load *t4_tom* driver with *unsupported_allow_unload* parameter.

```
[root@host~]# modprobe t4_tom unsupported_allow_unload=1
```

ii. Stop all the offloaded traffic, servers and connections. Check for the reference count.

```
[root@host~]# cat /sys/module/t4_tom/refcnt
```

If the reference count is 0, the driver can be directly unloaded. Skip to step (iii)

If the count is non-zero, load a COP policy which disables offload using the following procedure:

a. Create a policy file which will disable offload

```
[root@host~]# cat policy_file
all => !offload
```

b. Compile and apply the output policy file

```
[root@host~]# cop -o no-offload.cop policy_file
[root@host~]# cxgbtool ethX policy no-offload.cop
```

iii. Unload the driver:

```
[root@host~]# rmmod t4_tom
[root@host~]# rmmod toecore
[root@host~]# rmmod cxgb4
```

III. Virtual Function Network (vNIC)

1. Introduction

The ever-increasing network infrastructure of IT enterprises has lead to a phenomenal increase in maintenance and operational costs. IT managers are forced to acquire more physical servers and other data center resources to satisfy storage and network demands. To solve the Network and I/O overhead, users are opting for server virtualization which consolidates I/O workloads onto lesser physical servers thus resulting in efficient, dynamic and economical data center environments. Other benefits of Virtualization include improved disaster recovery, server portability, cloud computing, Virtual Desktop Infrastructure (VDI), etc.

Chelsio's Unified Wire family of adapters deliver increased bandwidth, lower latency and lower power with virtualization features to maximize cloud scaling and utilization. The adapters also provide full support for PCI-SIG SR-IOV to improve I/O performance on a virtualized system. User can configure up to 64 Virtual and 8 Physical functions (with 4 PFs as SR-IOV capable) along with 336 virtual MAC addresses.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with the Chelsio vNIC driver:

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

- T420-CX

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the vNIC driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49(Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

The Virtual Function implementation for Chelsio adapters comprises of two modules:

- Standard NIC driver module, *cxgb4*, which runs on base Hypervisor and is responsible for instantiation and management of the PCIe Virtual Functions (VFs) on the adapter.
- VF NIC driver module, *cxgb4vf*, which runs on Virtual Machine (VM) guest OS using VFs "attached" via Hypervisor VM initiation commands.

2.1. Pre-requisites

Please make sure that the following requirements are met before installation:

- SR-IOV should be enabled in the machine.
- Intel Virtualization Technology for Directed I/O (VT-d) should be enabled in the BIOS.
- PCI Express Slot should be ARI capable.

2.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. On the host, install network driver:

```
[root@host~]# make nic_install
```

iii. On the guest (VM), install vNIC driver:

```
[root@host~]# make vnic_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

3.1. Instantiate Virtual Functions (SR-IOV)

To instantiate virtual functions on the host, run the following commands:

```
[root@host~]# modprobe cxgb4
[root@host~]# echo n >
/sys/class/net/ethX/device/driver/<bus_id>/sriov_numvfs
```

Here, *ethX* is the interface and *n* specifies the number of virtual functions to be instantiated per physical function (*bus_id*). A maximum of 64 virtual functions can be instantiated with 16 virtual functions per physical function.

Example:

```
[root@ ~]# modprobe cxgb4
[root@ ~]# echo 16 > /sys/class/net/eth0/device/driver/0000\:06\:00.3/sriov_numvfs
[root@ ~]#
```

Note *To get familiar with physical and virtual function terminologies, please refer the PCI Express specification.*

Unload the vNIC driver on the host (if loaded):

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

The virtual functions can now be assigned to virtual machines (guests).

3.2. Loading the driver

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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

To load the driver, run the following command:

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

4. Software/Driver Configuration and Fine-tuning

This section describes the method to rate-limit traffic passing through virtual functions (VFs).

i. The VF rate limit needs to be set on the Host (hypervisor). Apply rate-limiting using:

```
[root@host~]# ip link set dev mgmtpfXX vf <vf_number> rate <rate_in_mbps>
```

Here,

- *mgmtpfXX* is the management interface to be used. For each PF on which VFs are instantiated, 1 interface will be created (in "ifconfig -a").
- *vf_number* is VF on which rate-limiting is applied. Value 0-15.

ii. Run traffic over the VF and the throughput should be rate-limited as per the values set in the previous step.

Example:

i. 4 VFs are instantiated on PF0.

```
[root@host~]# modprobe cxgb4
[root@host~]# echo 4 >
/sys/class/net/ethX/device/driver/<bus_id>/sriov_numvfs
```

ii. 2 VMs are configured with 2 VFs each. 2 different networks are configured with the following IP configuration:

```
VM0: VF0 (102.1.1.2/24)
VF1 (102.2.2.2/24)
VM1: VF2 (102.1.1.3/24)
VF3 (102.2.2.3/24)
```

iii. VF Rate-limiting is configured on the host:

```
[root@host~]# ip link set dev mgmtpf10 vf 0 rate 2000
[root@host~]# ip link set dev mgmtpf10 vf 1 rate 3000
```

The traffic on 102.1.1.X network will be rate-limited to 2Gbps whereas traffic on 102.2.2.X network will be rate-limited to 3Gbps.

5. Software/Driver Unloading

5.1. Unloading the driver

The vNIC driver must be unloaded on the Guest OS by the root user. Any attempt to unload the driver as a regular user will fail.

To unload the driver, execute the following command:

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

IV. iWARP (RDMA)

1. Introduction

Chelsio's Terminator engine implements a feature rich RDMA implementation which adheres to the IETF standards with optional markers and MPA CRC-32C.

The iWARP RDMA operation benefits from the virtualization, traffic management and QoS mechanisms provided by Terminator engine. It is possible to ACL process iWARP RDMA packets. It is also possible to rate control the iWARP traffic on a per-connection or per-class basis, and to give higher priority to QPs that implement distributed locking mechanisms. The iWARP operation also benefits from the high performance and low latency TCP implementation in the offload engine.

1.1. Hardware Requirements

1.1.1. Supported adapters

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

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

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the iWARP driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

2.1. Pre-requisites

- *libnl-devel* package should be installed for *libiwpm* installation.
- If you are planning to upgrade OFED on one member of the cluster, the upgrade needs to be installed on all the members.
- If you want to install OFED with NFS-RDMA support, refer to Setting up NFS-RDMA section.

2.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install iWARP drivers and libraries:

```
[root@host~]# make iwarp_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

3.1. Loading iWARP driver

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

To load the iWARP driver we need to load the NIC driver and core RDMA drivers first. Run the following commands:

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe iw_cxgb4
[root@host~]# modprobe rdma_ucm
```

Optionally, you can start the iWARP Port Mapper daemon to enable port mapping:

```
[root@host~]# iwpmd
```

4. Software/Driver Configuration and Fine-tuning

4.1. Testing connectivity with *ping* and *rping*

Load the NIC, iWARP & core RDMA modules as mentioned in Software/Driver Loading section. After which, you will see two or four ethernet interfaces for the Terminator device. Configure them with an appropriate ip address, netmask, etc. You can use the Linux *ping* command to test basic connectivity via the Terminator interface. To test RDMA, use the *rping* command that is included in the librdmacm-utils RPM:

Run the following command on the server machine:

```
[root@host~]# rping -s -a server_ip_addr -p 9999
```

Run the following command on the client machine:

```
[root@host~]# rping -c -Vv -C10 -a server_ip_addr -p 9999
```

You should see ping data like this on the client:

```
ping data: rdma-ping-0: ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqr
ping data: rdma-ping-1: BCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrs
ping data: rdma-ping-2: CDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrst
ping data: rdma-ping-3: DEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstu
ping data: rdma-ping-4: EFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuv
ping data: rdma-ping-5: FGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvw
ping data: rdma-ping-6: GHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwx
ping data: rdma-ping-7: HIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxy
ping data: rdma-ping-8: IJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz
ping data: rdma-ping-9: JKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyzA
client DISCONNECT EVENT...
#
```

4.2. Enabling various MPIs

4.2.1. Setting shell for Remote Login

User needs to set up authentication on the user account on all systems in the cluster to allow user to remotely logon or executing commands without password.

Quick steps to set up user authentication:

i. Change to user home directory

```
[root@host~]# cd
```

ii. Generate authentication key

```
[root@host~]# ssh-keygen -t rsa
```

iii. Hit [Enter] upon prompting to accept default setup and empty password phrase
iv. Create authorization file

```
[root@host~]# cd .ssh
[root@host~]# cat *.pub > authorized_keys
[root@host~]# chmod 600 authorized_keys
```

v. Copy directory .ssh to all systems in the cluster

```
[root@host~]# cd
[root@host~]# scp -r /root/.ssh remotehostname-or-ipaddress:
```

4.2.2. Configuration of various MPIs (Installation and Setup)

- **Intel-MPI**

i. Download latest Intel MPI from the Intel website
ii. Copy the license file (.lic file) into `l_mpi_p_x.y.z` directory
iii. Create machines.LINUX (list of node names) in `l_mpi_p_x.y.z`
iv. Select advanced options during installation and register the MPI.
v. Install software on every node.

```
[root@host~]# ./install.py
```

vi. Set IntelMPI with mpi-selector (do this on all nodes).

```
[root@host~]# mpi-selector --register intelmpi --source-dir
/opt/intel/impi/3.1/bin/
[root@host~]# mpi-selector --set intelmpi
```

vii. Edit `.bashrc` and add these lines:

```
export RSH=ssh
export DAPL_MAX_INLINE=64
export I_MPI_DEVICE=rdssm:chelsio
export MPIEXEC_TIMEOUT=180
export MPI_BIT_MODE=64
```

viii. Logout & log back in.
ix. Populate `mpd.hosts` with node names.

Note

- *The hosts in this file should be Chelsio interface IP addresses.*
- `I_MPI_DEVICE=rdssm:chelsio` *assumes you have an entry in* `/etc/dat.conf` *named* `chelsio`*.*
- `MPIEXEC_TIMEOUT` *value might be required to increase if heavy traffic is going across the systems.*

x. Contact Intel for obtaining their MPI with DAPL support.
xi. To run Intel MPI over RDMA interface, DAPL 2.0 should be set up as follows:

Enable the Chelsio device by adding an entry at the beginning of the /`etc/dat.conf` file for the Chelsio interface. For instance, if your Chelsio interface name is eth2, then the following line adds a DAT version 2.0 device named “chelsio2" for that interface:

```
chelsio2 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "eth2 0" ""
```

Open MPI (Installation and Setup)

Open MPI iWARP support is only available in Open MPI version 1.3 or greater.

Open MPI will work without any specific configuration via the openib btl. Users wishing to performance tune the configurable options may wish to inspect the receive queue values. Those can be found in the "Chelsio T4" section of `mca-btl-openib-device-params.ini`. Follow the steps mentioned below to install and configure Open MPI.

i. If not alreay done, install *mpi-selector* tool.
ii. Download the latest stable/feature version of openMPI from OpenMPI website.

iii. Untar and change your current working directory to openMPI package directory.
iv. Configure and install as:

```
[root@host~]#./configure --with-openib=/usr  CC=gcc CXX=g++ F77=gfortran
FC=gfortran --enable-mpirun-prefix-by-default --prefix=/usr/mpi/gcc/openmpi-
x.y.z/ --with-openib-libdir=/usr/lib64/ --libdir=/usr/mpi/gcc/openmpi-
x.y.z/lib64/ --with-contrib-vt-flags=--disable-iotrace
[root@host~]# make
[root@host~]# make install
```

The above step will install openMPI in */usr/mpi/gcc/openmpi-x.y.z/*

Note *To enable multithreading, add "--enable-mpi-thread-multiple" and "--with-threads=posix" parameters to the above configure command.*

v. Next, create a shell script, *mpivars.csh,* with the following entry:

```
# path
if ("" == "`echo $path | grep /usr/mpi/gcc/openmpi-x.y.z/bin`") then
   set path=(/usr/mpi/gcc/openmpi-x.y.z/bin $path)
endif


# LD_LIBRARY_PATH
if ("1" == "$?LD_LIBRARY_PATH") then
    if ("$LD_LIBRARY_PATH" !~ */usr/mpi/gcc/openmpi-x.y.z/lib64*) then
    setenv LD_LIBRARY_PATH /usr/mpi/gcc/openmpi-
x.y.z/lib64:${LD_LIBRARY_PATH}
    endif
else
   setenv LD_LIBRARY_PATH /usr/mpi/gcc/openmpi-x.y.z/lib64
endif
# MPI_ROOT
setenv MPI_ROOT /usr/mpi/gcc/openmpi-x.y.z
```

vi. Simlarly, create another shell script, *mpivars.sh*, with the following entry:

```
# PATH
if test -z "`echo $PATH | grep /usr/mpi/gcc/openmpi-x.y.z/bin`"; then
    PATH=/usr/mpi/gcc/openmpi-x.y.z/bin:${PATH}
    export PATH
fi
# LD_LIBRARY_PATH
if test -z "`echo $LD_LIBRARY_PATH | grep
/usr/mpi/gcc/openmpi-          x.y.z/lib64`"; then
   LD_LIBRARY_PATH=/usr/mpi/gcc/openmpi-   x.y.z/lib64${LD_LIBRARY_PATH:+:}$
{LD_LIBRARY_PATH}
   export LD_LIBRARY_PATH
fi


# MPI_ROOT
MPI_ROOT=/usr/mpi/gcc/openmpi-x.y.z
export MPI_ROOT
```

vii. Next, copy the two files created in steps (v) and (vi) to */usr/mpi/gcc/openmpi-x.y.z/bin* and */usr/mpi/gcc/openmpi-x.y.z/etc*

viii. Register OpenMPI with MPI-selector:

```
[root@host~]# mpi-selector --register openmpi --source-dir
/usr/mpi/gcc/openmpi-x.y.z/bin
```

ix. Verify if it is listed in mpi-selector:

```
[root@host~]# mpi-selector --l
```

x. Set OpenMPI:

```
[root@host~]# mpi-selector --set openmpi –yes
```

xi. Logut and log back in.

MVAPICH2 (Installation and Setup)

i. Download the latest MVAPICH2 software package from http://mvapich.cse.ohio-state.edu/
ii. Untar and change your current working directory to MVAPICH2 package directory.
iii. Configure and install as:

```
[root@host~]# ./configure --prefix=/usr/mpi/gcc/mvapich2-x.y/ --with-
device=ch3:mrail --with-rdma=gen2 --enable-shared --with-ib-
libpath=/usr/lib64/ -enable-rdma-cm --libdir=/usr/mpi/gcc/mvapich2-x.y/lib64
[root@host~]# make
[root@host~]# make install
```

The above step will install MVAPICH2 in */usr/mpi/gcc/mvapich2-x.y/*

iv. Next, create a shell script , mpivars.csh, with the following entry:

```
# path
if ("" == "`echo $path | grep /usr/mpi/gcc/mvapich2-x.y/bin`") then
   set path=(/usr/mpi/gcc/mvapich2-x.y/bin $path)
endif


# LD_LIBRARY_PATH
if ("1" == "$?LD_LIBRARY_PATH") then
    if ("$LD_LIBRARY_PATH" !~ */usr/mpi/gcc/mvapich2-x.y/lib64*) then
    setenv LD_LIBRARY_PATH /usr/mpi/gcc/mvapich2-
x.y/lib64:${LD_LIBRARY_PATH}
    endif
else
   setenv LD_LIBRARY_PATH /usr/mpi/gcc/mvapich2-x.y/lib64
endif


# MPI_ROOT
setenv MPI_ROOT /usr/mpi/gcc/mvapich2-x.y
```

v. Simlarly, create another shell script, *mpivars.sh*, with the following entry:

```
# PATH
if test -z "`echo $PATH | grep /usr/mpi/gcc/ mvapich2-x.y/bin`"; then
    PATH=/usr/mpi/gcc/mvapich2-x.y/bin:${PATH}
    export PATH
fi
# LD_LIBRARY_PATH
if test -z "`echo $LD_LIBRARY_PATH | grep /usr/mpi/gcc/mvapich2-
x.y/lib64`"; then
   LD_LIBRARY_PATH=/usr/mpi/gcc/mvapich2-
x.y/lib64${LD_LIBRARY_PATH:+:}${LD_LIBRARY_PATH}
   export LD_LIBRARY_PATH
fi


# MPI_ROOT
MPI_ROOT=/usr/mpi/gcc/mvapich2-x.y
export MPI_ROOT
```

vi. Next, copy the two files created in steps (iv) and (v) *to /usr/mpi/gcc/mvapich2-x.y/bin* and */usr/mpi/gcc/mvapich2-x.y/etc*

vii. Add the following entries in *.bashrc* file:

```
export MVAPICH2_HOME=/usr/mpi/gcc/mvapich2-x.y/
export MV2_USE_IWARP_MODE=1
export MV2_USE_RDMA_CM=1
```

viii. Register MPI:

```
[root@host~]# mpi-selector --register mvapich2 --source-dir
/usr/mpi/gcc/mvapich2-x.y/bin/
```

ix. Verify if it is listed in mpi-selector:

```
[root@host~]# mpi-selector --l
```

x. Set MVAPICH2:

```
[root@host~]# mpi-selector --set mvapich2 –yes
```

xi. Logut and log back in.
xii. Populate `mpd.hosts` with node names.
xiii. On each node, create /etc/mv2.conf with a single line containing the IP address of the local adapter interface. This is how MVAPICH2 picks which interface to use for RDMA traffic.

4.2.3. Building MPI tests

i. Download *Intel's MPI Benchmarks* from http://software.intel.com/en-us/articles/intel-mpi-benchmarks
ii. Untar and change your current working directory to *src* directory.
iii. Edit *make_mpich* file and set *MPI_HOME* variable to the MPI which you want to build the benchmarks tool against. For example, in case of openMPI-1.6.4 set the variable as:

```
MPI_HOME=/usr/mpi/gcc/openmpi-1.6.4/
```

iv. Next, build and install the benchmarks using:

```
[root@host~]# gmake -f make_mpich
```

The above step will install IMB-MPI1, IMB-IO and IMB-EXT benchmarks in the current working directory (i.e. *src*).

v. Change your working directory to the MPI installation directory. In case of OpenMPI, it will be */usr/mpi/gcc/openmpi-x.y.z/*
vi. Create a directory called *tests* and then another directory called *imb* under *tests.*
vii. Copy the benchmarks built and installed in step (iv) to the *imb* directory.
viii. Follow steps (v), (vi) and (vii) for all the nodes.

4.2.4. Running MPI applications

- Run Intel MPI applications as:

```
mpdboot -n <no_of_nodes_in_cluster> -r ssh
mpdtrace
mpiexec -ppn -n 2 /opt/intel/impi/3.1/tests/IMB-3.1/IMB-MPI1
```

The performance is best with NIC MTU set to 9000 bytes.

- Run Open MPI application as:

```
mpirun --host  node1,node2 -mca btl openib,sm,self /usr/mpi/gcc/openmpi-
x.y.z/tests/imb/IMB-MPI1
```

Note

For OpenMPI/RDMA clusters with node counts greater than or equal to 8 nodes, and process counts greater than or equal to 64, you may experience the following RDMA address resolution error when running MPI jobs with the default OpenMPI settings:

```
The RDMA CM returned an event error while attempting to make a connection.
This type of error usually indicates a network configuration error.
Local host:   core96n3.asicdesigners.com
Local device: Unknown
Error name:   RDMA_CM_EVENT_ADDR_ERROR
Peer:         core96n8
```

***Workaround**: Increase the OpenMPI rdma route resolution timeout. The default is 1000, or 1000ms. Increase it to 30000 with this parameter:*

```
--mca btl_openib_connect_rdmacm_resolve_timeout 30000
```

Important

openmpi-1.4.3 can cause IMB benchmark stalls due to a shared memory BTL issue. This issue is fixed in openmpi-1.4.5 and later releases. Hence, it is recommended that you download and install the latest stable release from Open MPI's official website, http://www.open-mpi.org

- Run MVAPICH2 application as :

```
mpirun_rsh -ssh -np 8 -hostfile mpd.hosts $MVAPICH2_HOME/tests/imb/IMB-MPI1
```

4.3. Setting up NFS-RDMA

4.3.1. Starting NFS-RDMA

- **Server-side settings**

Follow the steps mentioned below to set up an NFS-RDMA server.

i. Make entry in `/etc/exports` file for the directories you need to export using NFS-RDMA on server as:

```
/share/rdma      *(fsid=0,async,insecure,no_root_squash)
/share/rdma1     *(fsid=1,async,insecure,no_root_squash)
```

Note that for each directory you export, you should have DIFFERENT fsid's.

ii. Load the iwarp modules and make sure peer2peer is set to 1.
iii. Load `xprtrdma` and `svcrdma` modules as:

```
[root@host~]# modprobe xprtrdma
[root@host~]# modprobe svcrdma
```

iv. Start the nfs service as:

```
[root@host~]# service nfs start
```

All services in NFS should start without errors.

v. Now we need to edit the file portlist in the path `/proc/fs/nfsd/`
Include the rdma port 2050 into this file as:

```
[root@host~]# echo rdma 2050 > /proc/fs/nfsd/portlist
```

vi. Run `exportfs` to make local directories available for Network File System (NFS) clients to mount.

```
[root@host~]# exportfs
```

Now the NFS-RDMA server is ready.

- **Client-side settings**

Follow the steps mentioned below at the client side.

i. Load the iwarp modules and make sure peer2peer is set to 1. Make sure you are able to ping and ssh to the server Chelsio interface through which directories will be exported.

ii. Load the `xprtrdma` module.

```
[root@host~]# modprobe xprtrdma
```

iii. Run the `showmount` command to show all directories from server as:

```
[root@host~]# showmount –e <server-chelsio-ip>
```

iv. Once the exported directories are listed, mount them as:

```
[root@host~]# mount.nfs <serverip>:<directory> <mountpoint-on-client> -o
vers=3,rdma,port=2050,wsize=65536,rsize=65536
```

4.4. Performance Tuning

i. Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.
ii. Run the performance tuning script to map iWARP queues to different CPUs.

```
[root@host~]# t4_perftune.sh -Q rdma -n
```

5. Software/Driver Unloading

To unload the iWARP driver, run the following command:

```
[root@host~]# rmmod iw_cxgb4
```

V. iSER

1. Introduction

The iSCSI Extensions for RDMA (iSER) protocol is a translation layer for operating iSCSI over RDMA transports, such as iWARP/Ethernet or InfiniBand.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio iSER driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory-free; limited number of offload connections supported*)
- T520-BT
- T580-CR
- T580-LP-CR
- T520-LL-CR
- T520-CR
- T540-CR

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the iSER driver is available for the following kernels compiled on RHEL 7.3:

- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Kernel Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. To install 4.9.49 kernel with iSER components enabled, use the following command:

```
[root@host~]# make kernel_install
```

Note

If you wish to use custom 4.8.X/4.9.X kernel, enable the following iSER parameters in the kernel configuration file and then proceed with kernel installation:

```
CONFIG_ISCSI_TARGET=m
CONFIG_INFINIBAND_ISER=m
CONFIG_INFINIBAND_ISERT=m
```

iii. Boot into the new kernel and install Chelsio Unified Wire.

3. Software/Driver Installation

3.1. Pre-requisites

- *libnl-devel* package should be installed for libiwpm installation.
- Python v2.7 or above is required for targetcli installation. If Python v2.7 is not already present in the system, or if an older version exists, v2.7.10 provided in the package will be installed.

3.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install Chelsio iSER driver, libraries and targetcli utilities:

```
[root@host~]# make iser_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

4. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

Follow the steps mentioned below on both target and initiator machines:

i. Unload Chelsio iWARP driver if previously loaded:

```
[root@host~]# rmmod iw_cxgb4
```

ii. Load the following modules:

```
[root@host~]# modprobe iw_cxgb4 mpa_rev=2
[root@host~]# modprobe rdma_ucm
```

iii. Start the iWARP Port Mapper Daemon:

```
[root@host~]# iwpmd
```

iv. Bring up the Chelsio interface(s):

```
[root@host~]# ifconfig ethX x.x.x.x up
```

v. Run the following script for IRQ mapping:

```
[root@host~]# t4_perftune.sh
```

vi. On target, run the following command:

```
[root@host~]# modprobe ib_isert
```

On initiator, run the following command:

```
[root@host~]# modprobe ib_iser
```

5. Software/Driver Configuration and Fine-tuning

i. Configure LIO target with iSER support, using ramdisk as LUN:

```
[root@host~]# targetcli /backstores/ramdisk create name=ram0 size=1GB
[root@host~]# targetcli /iscsi create wwn=iqn.2003-01.org.lun0.target
[root@host~]# targetcli /iscsi/iqn.2003-01.org.lun0.target/tpg1/luns create
/backstores/ramdisk/ram0
[root@host~]# targetcli /iscsi/iqn.2003-01.org.lun0.target/tpg1 set
attribute authentication=0 demo_mode_write_protect=0 generate_node_acls=1
cache_dynamic_acls=1
[root@host~]# targetcli saveconfig
```

ii. Discover LIO target using OpeniSCSI initiator:

```
[root@host~]# iscsiadm -m discovery -t st -p  102.10.10.4
```

iii. Enable iSER support in LIO target:

```
[root@host~]# targetcli /iscsi/iqn.2003-
01.org.lun0.target/tpg1/portals/0.0.0.0:3260 enable_iser boolean=True
```

iv. Login from the initiator with iSER as transport:

```
[root@host~]# iscsiadm -m node -p 102.10.10.4 -T iqn.2003-01.org.lun0.target
--op update -n node.transport_name -v iser
[root@host~]# iscsiadm -m node -p 102.10.10.4 -T iqn.2003-01.org.lun0.target
--login
```

5.1. HMA

To use HMA, please ensure that Unified Wire is installed using the *Unified Wire(Default)* configuration tuning option.

Currently 182 IPv4/91 IPv6 iSER connections are supported on T6225-SO-CR adapter. The following image shows the HMA reserved memory.

```
[root@[redacted] ~]# cat /sys/kernel/debug/cxgb4/0000\:81\:00.4/meminfo
EDC0:             0x0-0x3fffff [4.00 MiB]
EDC1:             0x400000-0x7fffff [4.00 MiB]
HMA:              0x800000-0x47fffff [64.0 MiB]

RQUDP region:     0xffffffff-0xfffffffe [0 B]
DBQ contexts:     0x4d2b80-0x4e87ff [87.1 KiB]
IMSG contexts:    0x4e8800-0x4fa7ff [72.0 KiB]
FLM cache:        0x4fa800-0x53d3bf [267 KiB]
ULPTX state:      0x53d3c0-0x53da3f [1.63 KiB]
ULPRX state:      0x53da40-0x53db3f [256 B]
Timers:           0x53db40-0x54bb7f [56.1 KiB]
TCBs:             0x54bb80-0x5abfff [385 KiB]
Tx payload:       0x5ac000-0x66bfff [768 KiB]
Rx payload:       0x66c000-0x72bfff [768 KiB]
Pstructs:         0x72c000-0x72e9ff [10.5 KiB]
Rx FL:            0x72ea00-0x72ea7f [128 B]
Tx FL:            0x72ea80-0x72eaff [128 B]
Pstruct FL:       0x72eb00-0x72ecff [512 B]
TDDP region:      0x72ed00-0x75047f [134 KiB]
iSCSI region:     0x750480-0x77047f [128 KiB]
TLSKey region:    0x770480-0x78047f [64.0 KiB]
TPT region:       0x800000-0x11248ff [9.14 MiB]
STAG region:      0x800000-0x11248ff [9.14 MiB]
TXPBL region:     0x1124900-0x27fff7f [22.9 MiB]
PBL region:       0x1124900-0x27fff7f [22.9 MiB]
RQ region:        0x27fff80-0x47fff7f [32.0 MiB]

uP RAM:           0x0-0xffffffff [4.00 GiB]
uP Extmem2:       0x0-0xffffffff [4.00 GiB]

48 Rx pages of size 16KiB for 1 channels
48 Tx pages of size 16KiB for 2 channels
168 p-structs

Port 0 using 0 pages out of 768 allocated
Port 1 using 0 pages out of 0 allocated
Port 2 using 0 pages out of 768 allocated
Port 3 using 0 pages out of 0 allocated
Loopback 0 using 0 pages out of 128 allocated
Loopback 1 using 0 pages out of 128 allocated
```

The following image is of *dmesg* log showing HMA memory being reserved during adapter initialization:

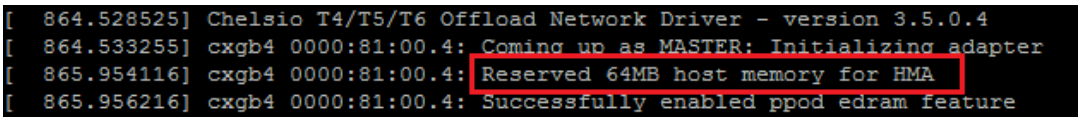

The following image shows the number of offloaded iSER connections.

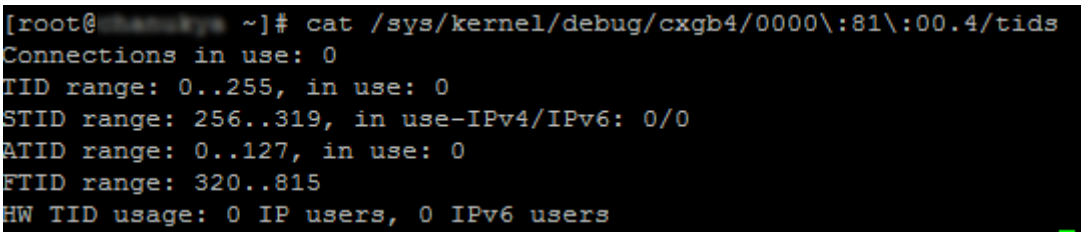

5.2. Performance Tuning

i. Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.
ii. Run the performance tuning script to map iWARP queues to different CPUs.

```
[root@host~]# t4_perftune.sh -Q rdma -n
```

6. Software/Driver Unloading

To unload iSER driver:

On target, run the following commands:

```
[root@host~]# rmmod ib_isert
[root@host~]# rmmod iw_cxgb4
```

On initiator, run the following commands:

```
[root@host~]# rmmod ib_iser
[root@host~]# rmmod iw_cxgb4
```

VI. WD-UDP

1. Introduction

Chelsio WD-UDP (Wire Direct-User Datagram Protocol) with Multicast is a user-space UDP stack with Multicast address reception and socket acceleration that enables users to run their existing UDP socket applications unmodified.

It features software modules that enable direct wire access from user space to the Chelsio network adapter with complete bypass of the kernel, which results in an ultra-low latency Ethernet solution for high frequency trading and other delay-sensitive applications.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio WD-UDP driver:

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

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the WD-UDP driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.2, 3.10.0-327.el7

- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install WD-UDP driver and libraries:

```
[root@host~]# make wdudp_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers:*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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

RDMA core modules from the OFED package should be loaded before proceeding. To load the WD-UDP driver, use the following commands which will automatically load RDMA core modules:

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe iw_cxgb4
[root@host~]# modprobe rdma_ucm
```

4. Software/Driver Configuration and Fine-tuning

4.1. Accelerating UDP Socket communications

The *libcxgb4_sock* library is a LD_PRELOAD-able library that accelerates UDP Socket communications transparently and without recompilation of the user application. This section describes how to use libcxgb4_sock.

By preloading *libcxgb4_sock*, all sockets created by the application are intercepted and possibly accelerated based on the user's configuration. Once accelerated, data for the UDP endpoint are transmitted or received via HW queues allocated specifically for the accelerated endpoint, bypassing the kernel, the host networking stack and sockets framework, and enabling ultra-low latency and high bandwidth utilization.

Due to HW resource limitations, only a small number of queues can be allocated for UDP acceleration. Therefore, only performance critical UDP applications should use *libcxgb4_sock*.

Only 64 IPv4 UDP / 28 IPv6 UDP sockets can be accelerated per Chelsio device, with *Unified Wire Configuration* tuning option. If you want more sockets to be accelerated, please use *Low Latency* or *High Capacity WD* tuning option.

4.1.1. Application Requirements

Certain application behavior is not supported by *libcxb4_sock* in this release. If your application does any of the following, it will not work with *libcxgb4_sock*:

- Calling *fork()* after creating UDP sockets and using the UDP socket in the child process.
- Using multiple threads on a single UDP socket without serialization. For instance, having one thread sending concurrently with another thread receiving. If your application does this, you need to serialize these paths with a spin or mutex lock.
- Only 1 UDP endpoint is allowed to bind to a given port per host. So, if you have multiple processes on the same host binding to the same UDP port number, you cannot use *libcxgb4_sock*.
- Applications must have root privileges to use *libcxgb4_sock*.
- Applications requiring bonded adapter interfaces are not currently supported.

The performance benefit observed with *libcxgb4_sock* will vary based on your application's behavior. While all UDP I/O is handled properly, only certain datagrams are accelerated. Non-accelerated I/O is handled by *libcxgb4_sock* via the host networking stack seamlessly. Both Unicast and Multicast datagrams can be accelerated, but the datagrams must meet the following criteria:

- Non-fragmented. In other words, they fit in a single IP datagram that is <= the adapter device MTU.

- Routed through the Terminator acceleration device. If the ingress datagram arrives via a device other than the Terminator acceleration device, then it will not utilize the acceleration path. On egress, if the destination IP address will not route out via the Terminator device, then it too will not be accelerated.

4.1.2. Using *libcxgb4_sock*

The *libcxgb4_sock* library utilizes the Linux RDMA Verbs subsystem, and thus requires the RDMA modules be loaded. Ensure that your systems load the *iw_cxgb4* and *rdma_ucm* modules:

```
[root@host~]# modprobe iw_cxgb4
[root@host~]# modprobe rdma_ucm
```

Now, preload *libcxgb4_sock*, using one of the methods mentioned below when starting your application:

- **Preloading using *wdload* script**:

```
[root@host~]# PROT=UDP wdload <pathto>/your_application
```

The above command will generate an end point file, *libcxgb4_sock.conf* at /etc/. Parameters like interface name and port number can be changed in this file.

The following example shows how to run Netperf with WD-UDP:

server:

```
[root@host~]# PROT=UDP wdload netserver -f -p <port_num>
```

client:

```
[root@host~]# PROT=UDP wdload netperf -H <hostIp> -p <port_num> -t UDP_RR
```

- **Preloading manually**

Create a configuration file that defines which UDP endpoints should be accelerated, their vlan and priority if any, as well as which Terminator interface/port should be used. The file `/etc/libcxgb4_sock.conf` contains these endpoint entries. Create this file on all systems using *libcxgb4_sock*. Here is the syntax:

```
#
# Syntax:
#
# endpoint {attributes} ...
# where attributes include:
#                interface = interface-name
#                port = udp-port-number
#                vlan = vlan-id
#                priority = vlan-priority
#
# e.g.
# endpoint {
#                interface=eth2.5
#                port = 8000 vlan = 5 prio rity=1
#   }
# endpoint { interface=eth2 port=9999}
#
# endpoints that bind to port 0 (requesting the host allocate a port)

# can be accelerated with port=0:
#
# endpoint {interface=eth1 port=0}
#
```

Assume your Terminator interface is eth2. To accelerate all applications that preload *libcxgb4_sock* using eth2, you only need one entry in `/etc/libcxgb4_sock.conf`:

```
endpoint {interface=eth2 port=0}
```

If you have eth2 and eth3 configured for example, you can define certain endpoints to eth2 and others to eth3:

```
endpoint {interface=eth2 port=9999}
endpoint {interface=eth3 port=8888}
```

For VLAN support, create your VLANs using the normal OS service (like vconfig, for example), then add entries to define the VLAN and priority for each endpoint to be accelerated:

```
endpoint {interface = eth2.5 port=10000}
endpoint {interface = eth2.7 priority=3 port=9000}
```

Now, preload *libcxgb4_sock*:

```
[root@host~]# CXGB4_SOCK_CFG=<path to config file>
LD_PRELOAD=libcxgb4_sock.so <pathto>/your_application
```

Note *In order to offload IPv6 UDP sockets, please select "low latency networking" as configuration tuning option during installation.*

4.1.3. Running WD-UDP in debug mode

To use *libcxgb4_sock*'s debug capabilities, use the *libcxgb4_sock_debug* library provided in the package. Follow the steps mentioned below:

i. Make the following entry in the */etc/syslog.conf* file:

```
*.debug                                /var/log/cxgb4.log
```

ii. Restart the service:

```
[root@host~]# /etc/init.d/syslog restart
```

iii. Finally, preload *libcxgb4_sock_debug* using the command mentioned below when starting your application:

```
[root@host~]# LD_PRELOAD=libcxgb4_sock_debug.so CXGB4_SOCK_DEBUG=-1
<pathto>/your_application
```

4.1.4. Running WD-UDP with larger I/O size

If the I/O size is > 3988, execute the commands mentioned below:

```
[root@host~]# echo 1024 > /proc/sys/vm/nr_hugepages
[root@host~]# CXGB4_SOCK_HUGE_PAGES=1 PROT=UDP wdload
<pathto>/your_application
```

4.1.5. Example with hpcbench/udp

The udp benchmark from the hpcbench suite can be used to show the benefits of libcxgb4_sock. The hpcbench suite can be found at:

Source: http://hpcbench.sourceforge.net/index.html

Sample: http://hpcbench.sourceforge.net/udp.html

The nodes in this example, r9 and r10, have Terminator eth1 configured and the ports are connected point-to-point.

```
[root@r9 ~]# ifconfig eth1|grep inet
          inet addr:192.168.2.111  Bcast:192.168.2.255  Mask:255.255.255.0
          inet6 addr: fe80::7:4300:104:465a/64 Scope:Link

[root@r9 ~]#
[root@r10 ~]# ifconfig eth1|grep inet
          inet addr:192.168.2.112  Bcast:192.168.2.255  Mask:255.255.255.0
          inet6 addr: fe80::7:4300:104:456a/64 Scope:Link
[root@r10 ~]#
```

For this benchmark, we need a simple “accelerate all” configuration on both nodes:

```
[root@r9 ~]# cat /etc/libcxgb4_sock.conf
endpoint {interface=eth1 port=0}
[root@r9 ~]#

[root@r10 ~]# cat /etc/libcxgb4_sock.conf
endpoint {interface=eth1 port=0}
[root@r10 ~]#
```

On R10, we run udpserver on port 9000 without *libcxgb4_sock* preloaded, and on port 90001 with preload:

```
[root@r10 ~]# /usr/local/src/hpcbench/udp/udpserver -p 9000 &
[1] 11453
[root@r10 ~]# TCP socket listening on port [9000]

[root@r10 ~]# LD_PRELOAD=libcxgb4_sock.so
/usr/local/src/hpcbench/udp/udpserver -p 9001 &
[2] 11454
[root@r10 ~]# TCP socket listening on port [9001]
```

Then on r9, we run udptest to port 9000 to see the host stack UDP latency:

```
[root@r9 ~]# /usr/local/src/hpcbench/udp/udptest -r 5 -a -h 192.168.1.112 -p
9000
```

Running the same test with *libcxgb4_sock*:

```
[root@r9 ~]# LD_PRELOAD=libcxgb4_sock.so /usr/local/src/hpcbench/udp/udptest
-r 5 -a -h 192.168.1.112 -p 9001
```

4.1.6. Determining if the application is being offloaded

To see if the application is being offloaded, open a window on one of the machines, and run `tcpdump` against the Chelsio interface. If you see minimal UDP output on the interface, then the UDP traffic is being properly offloaded.

5. Software/Driver Unloading

To unload the WD-UDP driver, run the following command:

```
[root@host~]# rmmod iw_cxgb4
```

VII. WD-TOE

1. Introduction

Chelsio WD-TOE (Wire Direct-Transmission Control Protocol) with a user-space TCP stack enables users to run their existing TCP socket applications unmodified.

It features software modules that enable direct wire access from user space to the Chelsio network adapter with complete bypass of the kernel, which results in a low latency Ethernet solution for high frequency trading and other delay-sensitive applications.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the currently shipping Chelsio adapters that are compatible with the WD-TOE driver.

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the WD-TOE driver is available for the following version(s):

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

2.1. Pre-requisites

Before proceeding, please ensure that the offload drivers are installed with WD-TOE support. They are installed by default with *Low Latency Networking*, *Wire Direct Latency* or *High Capacity WD configuration* tuning Options. With any other configuration tuning option, the installation needs to be customized.

2.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Run the following command to install WD-TOE drivers:

```
[root@host~]# make wdtoe_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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

Run the following commands to load the driver:

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe t4_tom
```

4. Software/Driver Configuration and Fine-tuning

4.1. Running the Application

To run an application with WD-TOE, use the following command:

```
[root@host~]# PROT=TCP wdload <path to application>
```

Note

- *Netperf, sockperf and IPerf only supported.*
- *Maximum of 10 connections supported.*
- *Abrupt killing of connections is not supported.*
- *Message size should be less than 1500.*
- *Short-lived connections not supported.*
- *Multi-threaded applications not supported.*
- *Jumbo frames not supported.*
- *Multiple adapters not supported.*

Example: To run Netperf application with WD-TOE.

i. Start netserver at the PEER, using the following command:

```
[root@host~]# PROT=TCP wdload netserver -D -4
```

ii. On the Test Machine, run the following command to run netperf application.

```
[root@host~]# PROT=TCP wdload netperf -H <PEER_IP> -t TCP_RR -l 10
```

5. Software/Driver Unloading

Reboot the system to unload the driver.

VIII. NVMe-oF

1. Introduction

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

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio NVMe-oF driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory-free; limited number of offload connections supported*)
- T580-CR
- T580-LP-CR
- T540-CR
- T520-CR
- T520-LL-CR
- T520-BT

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the NVMe-oF driver is available for the following kernels compiled on RHEL 7.3:

- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Kernel Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install 4.9.49 kernel with NVMe-oF components enabled, use the following command:

```
[root@host~]# make kernel_install
```

Note *If you wish to use custom 4.8.X/4.9.X kernel, enable the following parameters in the kernel configuration file and then proceed with kernel installation:*

```
CONFIG_BLK_DEV_NVME=m
CONFIG_NVME_RDMA=m
CONFIG_NVME_TARGET=m
CONFIG_NVME_TARGET_RDMA=m
CONFIG_NVME_RDMA=m
CONFIG_BLK_DEV_NULL_BLK=m
CONFIG_CONFIGFS_FS=y
```

iii. Boot into the new kernel and install Chelsio Unified Wire.

3. Software/Driver Installation

3.1. Pre-requisites

- *bzip2-devel, zlib-devel, ncurses-devel, sqlite-devel, libudev-devel* packages for *nvmetcli* utility should be installed.
- Python v2.7 or above should be installed. If this version is not already present in the system, or if an older version exists, v2.7.10 provided in the package will be installed.

3.2. Installlation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install iWARP driver and NVMe utilities:

```
[root@host~]# make nvme_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

4. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

Follow the steps mentioned below on both target and initiator machines:

i. Load the following modules:

```
[root@host~]# modprobe iw_cxgb4
[root@host~]# modprobe rdma_ucm
```

ii. Bring up the Chelsio interface(s):

```
[root@host~]# ifconfig ethX x.x.x.x up
```

iii. Mount configfs by running the below command:

```
[root@host~]# mount -t configfs none /sys/kernel/config
```

iv. On target, run the following commands:

```
[root@host~]# modprobe null_blk
[root@host~]# modprobe nvmet
[root@host~]# modprobe nvmet-rdma
```

On initiator, run the following commands:

```
[root@host~]# modprobe nvme
[root@host~]# modprobe nvme-rdma
```

5. Software/Driver Configuration and Fine-tuning

The following sections describe the method to configure target and initiator:

5.1. Target

i. The following commands will configure target using *nvmetcli* with a LUN:

```
[root@host~]# nvmetcli
/> cd subsystems
/subsystems> create nvme-ram0
/subsystems> cd nvme-ram0/namespaces
/subsystems/n...m0/namespaces> create nsid=1
/subsystems/n...m0/namespaces> cd 1
/subsystems/n.../namespaces/1> set device path=/dev/ram1
/subsystems/n.../namespaces/1> cd ../..
/subsystems/nvme-ram0> set attr allow_any_host=1
/subsystems/nvme-ram0> cd namespaces/1
/subsystems/n.../namespaces/1> enable
/subsystems/n.../namespaces/1> cd ../../../..
/> cd ports
/ports> create 1
/ports> cd 1/
/ports/1> set addr adrfam=ipv4.
/ports/1> set addr trtype=rdma
/ports/1> set addr trsvcid=4420
/ports/1> set addr traddr=102.1.1.102
/ports/1> cd subsystems
/ports/1/subsystems> create nvme-ram0
```

ii. Save the target configuration to a file:

```
/ports/1/subsystems> saveconfig /root/nvme-target_setup
/ports/1/subsystems> exit
```

iii. Clear the targets:

```
[root@host~]# nvmetcli clear
```

5.2. Initiator

i. Discover the target

```
[root@host~]# nvme discover -t rdma -a <target_ip> -s 4420
```

ii. Connect to target

- Connecting to a specific target:

```
[root@host~]# nvme connect -t rdma -a <target_ip> -s 4420 -n <target_name>
```

- Connecting to all targets configured on a portal

```
[root@host~]# nvme connect-all -t rdma -a <target_ip> -s 4420
```

iii. List the connected targets

```
[root@host~]# nvme list
```

iv. Format and mount the NVMe disks shown with the above command.

v. Disconnect from the target and unmount the disk:

```
[root@host~]# nvme disconnect -d <nvme_disk_name>
```

Note *nvme_disk_name is the name of the device (e.g., nvme0n1) and not the device path.*

5.3. HMA

To use HMA, please ensure that Unified Wire is installed using the *Unified Wire(Default)* configuration tuning option.

Currently 128 IPv4 connections are supported on T6225-SO-CR adapter, provided I/O queues is set to 1 while logging into the target.

```
[root@host~]# nvme connect -i 1 -t rdma -a <target_ip> -s 4420 -n
<target_name>
```

The following image shows the HMA reserved memory.

```
[root@[redacted] ~]# cat /sys/kernel/debug/cxgb4/0000\:81\:00.4/meminfo
EDC0:              0x0-0x3fffff [4.00 MiB]
EDC1:              0x400000-0x7fffff [4.00 MiB]
HMA:               0x800000-0x47fffff [64.0 MiB]

RQUDP region:      0xffffffff-0xfffffffe [0 B]
DBQ contexts:      0x4d2b80-0x4e87ff [87.1 KiB]
IMSG contexts:     0x4e8800-0x4fa7ff [72.0 KiB]
FLM cache:         0x4fa800-0x53d3bf [267 KiB]
ULPTX state:       0x53d3c0-0x53da3f [1.63 KiB]
ULPRX state:       0x53da40-0x53db3f [256 B]
Timers:            0x53db40-0x54bb7f [56.1 KiB]
TCBs:              0x54bb80-0x5abfff [385 KiB]
Tx payload:        0x5ac000-0x66bfff [768 KiB]
Rx payload:        0x66c000-0x72bfff [768 KiB]
Pstructs:          0x72c000-0x72e9ff [10.5 KiB]
Rx FL:             0x72ea00-0x72ea7f [128 B]
Tx FL:             0x72ea80-0x72eaff [128 B]
Pstruct FL:        0x72eb00-0x72ecff [512 B]
TDDP region:       0x72ed00-0x75047f [134 KiB]
iSCSI region:      0x750480-0x77047f [128 KiB]
TLSKey region:     0x770480-0x78047f [64.0 KiB]
TPT region:        0x800000-0x11248ff [9.14 MiB]
STAG region:       0x800000-0x11248ff [9.14 MiB]
TXPBL region:      0x1124900-0x27fff7f [22.9 MiB]
PBL region:        0x1124900-0x27fff7f [22.9 MiB]
RQ region:         0x27fff80-0x47fff7f [32.0 MiB]

uP RAM:            0x0-0xffffffff [4.00 GiB]
uP Extmem2:        0x0-0xffffffff [4.00 GiB]

48 Rx pages of size 16KiB for 1 channels
48 Tx pages of size 16KiB for 2 channels
168 p-structs
```

The following image is of *dmesg* log showing HMA memory being reserved during adapter initialization:

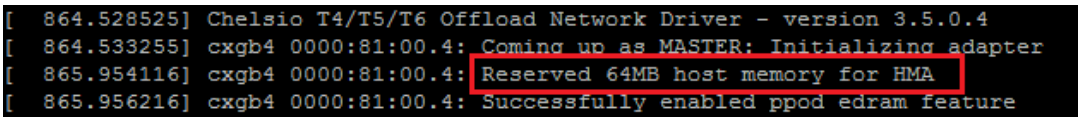

The following image shows the number of offloaded NVMe-oF connections:

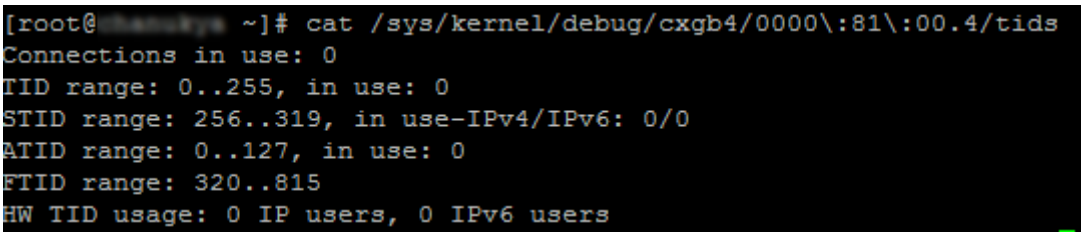

5.4. Performance Tuning

i. Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.
ii. Run the performance tuning script to map iWARP queues to different CPUs.

```
[root@host~]# t4_perftune.sh -Q rdma -n
```

6. Software/Driver Unloading

Follow the steps mentioned below to unload the drivers:

On target, run the following commands:

```
[root@host~]# rmmod nvmet-rdma
[root@host~]# rmmod nvmet
[root@host~]# rmmod iw_cxgb4
```

On initiator, run the following commands:

```
[root@host~]# rmmod nvme-rdma
[root@host~]# rmmod nvme
[root@host~]# rmmod iw_cxgb4
```

IX. LIO iSCSI Target Offload

1. Introduction

Linux-IO Target (LIO) is the in-kernel SCSI target implementation in Linux. This open-source standard supports common storage fabrics, including Fibre Channel, FCoE, iEEE 1394, iSCSI, NVMe-oF, iSER, SRP, USB, vHost, etc. The LIO iSCSI fabric module implements many advanced iSCSI features that increase performance and resiliency. The LIO iSCSI Target Offload driver provides the following high-level features:

- Offloads TCP/IP.
- Offloads iSCSI Header and Data Digest Calculations.
- Offload Speeds at 10/25/40/100Gb.
- Supports Direct Data Placement (DDP).
- Supports iSCSI Segmentation Offload and iSCSI PDU recovery.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio LIO iSCSI Target Offload driver:

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

1.2. Software Requirements

cxgb4, iscsi_target_mod, target_core_mod, ipv6 modules are required by LIO iSCSI Target Offload (`cxgbit.ko`) module to work.

1.2.1. Linux Requirements

Currently the LIO iSCSI Target Offload driver is available for the following version(s):

- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Kernel Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install 4.9.49 kernel with LIO iSCSI Target Offload enabled, use the following command:

```
[root@host~]# make kernel_install
```

iii. Boot into the new kernel and install Chelsio Unified Wire.

In case of 3.14.57 kernel, follow the below steps:

i. Download the kernel from kernel.org.
ii. Untar the tar-ball.
iii. Change your working directory to kernel package directory and run the following command to invoke the installation menu.

```
[root@host~]# make menuconfig
```

iv. Select **Device Drivers** > **Generic Target Core Mod (TCM) and ConfigFS Infrastructure**.
v. Enable **Linux-iSCSI.org iSCSI Target Mode Stack**.
vi. Select **Save**.
vii. Exit from the installation menu.
viii. Untar the patch file:

```
[root@host~]# cp /root/<driver_package>/src/cxgbit/patch/linux_3-14.a .
[root@host~]# ar xvf linux_3-14.a
```

ix. Apply all the patches to kernel source one by one:

```
[root@host~]# patch -p1 < <file_name>.patch
```

x. Continue with installation as usual.
xi. Reboot to the newly installed kernel. Verify by running *uname -a* command.
xii. Install the Unified Wire package as mentioned in the Chelsio Unified Wire chapter.

3. Software/Driver Installation

3.1. Pre-requisites

The LIO iSCSI Target Offload driver requires the following components to be installed to function:

- Python* (v2.7.10 provided in the package)
- TargetCLI* (v2.1 provided in the package)
- OpenSSL (Download from https://www.openssl.org/source/)

If not already present in the system, the component provided in the package will be installed along with the kernel.

3.2. Installlation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install LIO driver and targetcli utilities:

```
[root@host~]# make lio_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

4. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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

i. Load network driver (*cxgb4*):

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

ii. Bring up the interface.

```
[root@host~]# ifconfig ethX <IP address> up
```

iii. Load the LIO iSCSI Target Offload driver (*cxgbit)*:

```
 [root@host~]# modprobe cxgbit
```

5. Software/Driver Configuration and Fine-tuning

5.1. Configuring LIO iSCSI Target

The LIO iSCSI Target needs to be configured before it can become useful. Please refer the user manual at http://www.linux-iscsi.org/Doc/LIO Admin Manual.pdf to do so.

5.1.1. Sample Configuration

Here is a sample iSCSI configuration listing a target configured with 1 RAM disk LUN and ACL not configured:

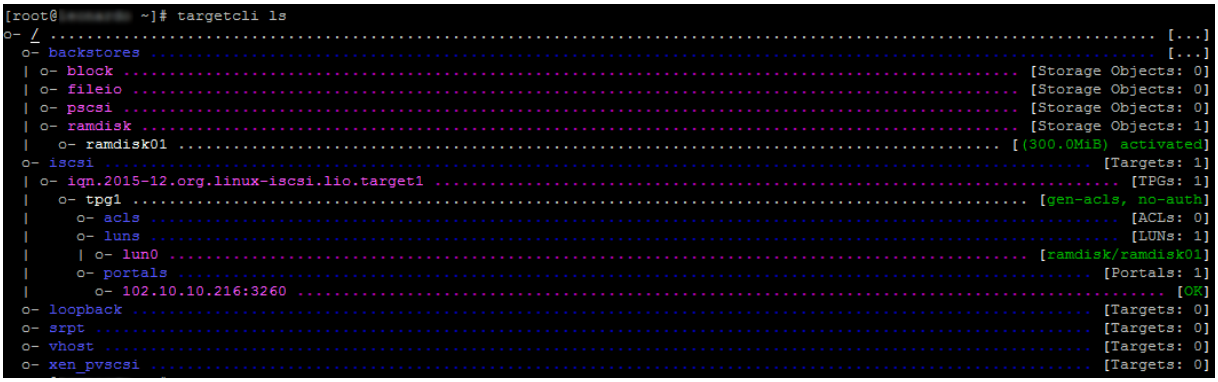

5.2. Offloading LIO iSCSI Connection

To offload the LIO iSCSI Target use the following command:

```
[root@host~]# echo 1 >
/sys/kernel/config/target/iscsi/<target_iqn>/tpgt_1/np/<target_ip>\:3260/cxg
bit
```

Example:

```
[root@ ~]#
[root@ ~]# echo 1 > /sys/kernel/config/target/iscsi/iqn.2015-12.org.linux-iscsi.lio.target1/tpgt_1/np/102.10.10.216\:3260/cxgbit
[root@ ~]#
```

Execute the above command for every portal address listening on Chelsio interface.

5.3. Running LIO iSCSI and Network Traffic Concurrently

If you wish to run network traffic with offload support (TOE) and LIO iSCSI traffic together, follow the steps mentioned below:

i. If not done already, load network driver with offload support (TOE):

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

ii. Create a new policy file:

```
[root@host~]# cat <new_policy_file>
```

iii. Add the following lines to offload all traffic except LIO iSCSI:

```
listen && src port <target_listening_port> && src host <target_listening_ip>
=> !offload
all => offload
```

iv. Compile the policy:

```
[root@host~]# cop -d -o <output_policy_file> <new_policy_file>
```

v. Apply the policy:

```
[root@host~]# cxgbtool ethX policy <output_policy_file>
```

Example:

Using multiple listening servers:

```
[root@ ~]# modprobe t4_tom
[root@ ~]# cat policy
listen && src port 3260 && src host 102.11.11.216 => !offload
listen && src port 3260 && src host 102.22.22.216 => !offload
listen && src port 3260 && src host 0.0.0.0 => !offload
all => offload
[root@ ~]# cop -o /root/policy.o /root/policy
[root@ ~]# cxgbtool eth2 policy /root/policy.o
[root@ ~]# cat /proc/net/offload/devices
```

5.4. Performance Tuning

i. Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.
ii. Run the performance tuning script to map LIO Target queues to different CPUs.

```
[root@host~]# t4_perftune.sh -Q iSCSIT -n
```

iii. Configure MTU 9000 on all interfaces.

For maximum performance, it is recommended to use iSCSI PDU offload initiator.

6. Software/Driver Unloading

6.1. Unloading the LIO iSCSI Target Offload driver

To unload the LIO iSCSI Target Offload kernel module, follow the steps mentioned below:

i. Log out from the initiator.
ii. Run the following command:

```
[root@host~]# echo 0 >
/sys/kernel/config/target/iscsi/<target_iqn>/tpgt_1/np/<target_ip>\:3260/cxg
bit
```

Execute the above command for every portal address listening on Chelsio interface.

iii. Unload the driver:

```
[root@host~]# rmmod cxgbit
```

6.2. Unloading the NIC driver

To unload the NIC driver, run the following command:

```
[root@host~]# rmmod cxgb4
```

X. iSCSI PDU Offload Target

1. Introduction

This section describes how to install and configure iSCSI PDU Offload Target software for use as a key element in your iSCSI SAN. The software runs on Linux-based systems that use Chelsio or non-Chelsio based Ethernet adapters. However, to guarantee highest performance, Chelsio recommends using Chelsio adapters. Chelsio's adapters include offerings that range from stateless offload adapters (regular NIC) to the full line of TCP/IP Offload Engine (TOE) adapters.

The software implements RFC 3720, the iSCSI standard of the IETF. The software has been fully tested for compliance to that RFC and others and it has been exhaustively tested for interoperability with the major iSCSI vendors.

The software implements most of the iSCSI protocol in software running in kernel mode on the host with the remaining portion, which consists of the entire fast data path, in hardware when used with Chelsio's TOE adapters. When standard NIC adapters are used the entire iSCSI protocol is executed in software.

The performance of this iSCSI stack is outstanding and when used with Chelsio's hardware it is enhanced further. Because of the tight integration with Chelsio's TOE adapters, this software has a distinct performance advantage over the regular NIC. The entire solution, which includes this software, Chelsio TOE hardware, an appropriate base computer system – including a high end disk subsystem, has industry leading performance. This can be seen when the entire solution is compared to others based on other technologies currently available on the market in terms of throughput and IOPS.

1.1. Features

Chelsio's iSCSI driver stack supports the iSCSI protocol in the Target mode. From henceforth "iSCSI Software Entity" term refers to the iSCSI target.

The Chelsio iSCSI PDU Offload Target software provides the following high level features:

- Expanded NIC Support
  - Chelsio TCP Offload Engine (TOE) Support
    - T6/T5/T4 Based HBAs (T6/T5/T4xx Series cards)
  - Non-Chelsio
    - Runs on regular NICs
- Chelsio Terminator ASIC Support
  - Offloads iSCSI Fast Data Path with Direct Data Placement (DDP)
  - Offloads iSCSI Header and Data Digest Calculations
  - Offload Speeds at 1Gb, 10Gb, 25Gb, 40Gb and 100Gb
  - Offloads TCP/IP for NAS simultaneously with iSCSI
- Target Specific features
  - Full compliance with RFC 3720

- Error Recovery Level 0 (ERL 0)
- CHAP support for both discovery and login including mutual authentication
- Internet Storage Name Service (iSNS) Client
- Target Access Control List (ACL)
- Multiple Connections per Session
- Multiple Targets
- Multiple LUNs per Target
- Multi Path I/O (MPIO)
- Greater than 2 TB Disk Support
- Reserve / Release for Microsoft Cluster© Support
- Persistent Reservation
- Dynamic LUN Resizing
- iSCSI Target Redirection
- Multiple Target device types
  - Block
  - Virtual Block (LVM, Software RAID, EVMS, etc.)
  - Built in RAM Disk
  - Built in zero copy RAM Disk
- Supports iSCSI Boot Initiators
- An Intuitive and Feature Rich Management CLI

This chapter will cover these features in detail.

1.2. Hardware Requirements

1.2.1. Supported adapters

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

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

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

1.2.2. Adapter Requirements

The Chelsio iSCSI PDU Offload Target software can be used with or without hardware protocol offload technology. There are two modes of operation using the iSCSI PDU Offload Target software on Ethernet-based adapters:

- Regular NIC – The software can be used in non-offloaded (regular NIC) mode. Please note however that this is the least optimal mode of operating the software in terms of performance.
- iSCSI HW Acceleration – In addition to offloading the TCP/IP protocols in hardware (TOE), this mode also takes advantage of Chelsio's ASIC capability of hardware assisted iSCSI data and header digest calculations as well as using the direct data placement (DDP) feature.

1.2.3. Storage Requirements

When using the Chelsio iSCSI target, a minimum of one hardware storage device is required. This device can be any of the device types that are supported (block, virtual block, RAM disk). Multiple storage devices are allowed by configuring the devices to one target or the devices to multiple targets. The software allows multiple targets to share the same device but use caution when doing this.

Chelsio's implementation of the target iSCSI stack has flexibility to accommodate a large range of configurations. For quick testing, using a RAM Disk as the block storage device works nicely. For deployment in a production environment a more sophisticated system would be needed. That typically consists of a system with one or more storage controllers with multiple disk drives attached running software or hardware based RAID.

1.3. Software Requirements

`chiscsi_base.ko` is iSCSI non-offload target mode driver and `chiscsi_t4.ko` is iSCSI PDU offload target mode driver.

cxgb4, toecore, t4_tom and chiscsi_base modules are required by `chiscsi_t4.ko` module to work in offloaded mode. Whereas in iscsi non-offloaded target (NIC) mode, only `cxgb4` is needed by `chiscsi_base.ko` module.

1.3.1. Linux Requirements

Currently the iSCSI PDU Offload Target software is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7

- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

1.3.2. Requirements for Installing the iSCSI Software

When installing the iSCSI software, it is required that the system have Linux kernel source or its headers installed in order to compile the iSCSI software as a kernel module. The source tree may be only header files, as for RHEL6 as an example, or a complete tree. The source tree needs to be configured and the header files need to be compiled. Additionally, the Linux kernel must be configured to use modules.

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install iSCSI-target driver,firmware and utilities:

```
[root@host~]# make iscsi_pdu_target_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

There are two main steps to installing the Chelsio iSCSI PDU Offload Target software. They are:

1. **Installing the iSCSI software** – The majority of this section deals with how to install the iSCSI software.
2. **Configuring the iSCSI software** – Information on configuring the software can be found in a section further into this user's guide.

3.1. Latest iSCSI Software Stack Driver Software

The iSCSI software stack comes bundled in the Chelsio Unified Wire package which can be downloaded from the Chelsio Download Center.

The iSCSI software is available for use with most installations of the Linux kernel. The software is dependent on the underlying NIC adapter driver and thus the limitation on what version of the Linux kernel it can run on is mostly dependent on the NIC driver's limitations.

The iSCSI module will be installed in the `/lib/modules/<linux_kernel_version>/updates/kernel/drivers/scsi/chiscsi` directory. The modules database will be updated by the installer. This allows the iSCSI module to be located when using the `modprobe` utility. The actual module `chiscsi_t4.ko` can be found inside the package under /build/src/chiscsi/t4.

The `iscsictl` tool and the `chisns` tool will be installed in `/sbin`. The `chisns` tool starts the iSNS client. The `iscsictl` tool is provided for configuring and managing the iSCSI targets and iSNS client. It also provides control for iSCSI global settings.

1. Loading the Kernel module

- Run modprobe as follows:

```
[root@host~]# modprobe chiscsi_t4
```

Note

i. *While using rpm-tar-ball for installation*
   a. *Uninstallation will result into chiscsi.conf file renamed into chiscsi.conf.rpmsave, but if again uninstallation is done then it will lead to overwriting of the old chiscsi.rpmsave file.*
   b. *Its advised to take a backup of chiscsi.conf file before you do an uninstallation and installation of new/same unified wire package. As re-installing/upgrading unified-wire package may lead to loss of chiscsi.conf file.*

ii. *Installation/uninstallation using source-tar-ball will neither remove the conf file nor rename it. It will always be intact. However it's recommended to always take a backup of your configuration file for both methods of installation.*

A sample iSCSI configuration file will be installed in `/etc/chelsio-iscsi/chiscsi.conf`. This file should be edited using a standard text editor and customized to fit your environment.

2. Set iSCSI service to automatically start at bootup

The `chelsio-target` service scripts are installed to `/etc/init.d` and the parameters for the script are installed at `/etc/sysconfig/chiscsi`. The script is installed as a system service.

To auto-start the iSCSI target service at a certain runlevel, e.g., runlevel 3, `chkconfig` can be used on Red Hat and Novell / SuSE based systems as follows:

```
[root@host~]# chkconfig --level 3 chelsio-target on
```

The `chelsio-target` service scripts do basic checks before starting the iSCSI target service, loads the kernel module, and starts all the targets configured by default. It can also be used to stop the targets, and restart/reload configuration.

Note

For the script to execute properly, make sure the following flag is set on all kernel.org kernels.

 # CONFIG_MODULE_FORCE_LOAD=y

4. Software/Driver Configuration and Fine-tuning

The Chelsio iSCSI software needs configuration before it can become useful. The following sections describe how this is done.

There are two main components used in configuring the Chelsio iSCSI software: the **configuration file** and the **iSCSI control tool**. This section describes in some detail what they are and their relationship they have with one another.

4.1. Command Line Tools

There are two command line tools, one for control of the iSNS client and one for control of the iSCSI target nodes.

4.1.1. iscsictl

The Chelsio iSCSI control tool, `iscsictl`, is a Command Line Interface (CLI) user space program that allows administrators to:

- Start/Stop the iSCSI Target
- Start the iSNS client
- Get/Set the iSCSI driver global settings
- Get/Set/Remove the iSCSI Target configuration settings
- Retrieve active sessions' information of an iSCSI Target
- Manually flush data to the iSCSI Target disks
- Reload the iSCSI configuration file
- Update the iSCSI configuration file
- Save the current iSCSI configuration to a file

4.1.2. chisns

The Chelsio iSNS client, `chisns`, can be started independently of `iscsictl`.

4.2. iSCSI Configuration File

The iSCSI configuration file is the place where information about the Chelsio iSCSI software is stored. The information includes global data that pertains to all targets as well as information on each specific iSCSI target node. Most of the information that can be placed in the configuration file has default values that only get overwritten by the values set in the configuration file.There are only a few global configuration items that can be changed.

There are many specific parameters that can be configured, some of which are iSCSI specific and the rest being Chelsio specific. An example of an iSCSI specific item is "HeaderDigest" which is defaulted to "None" but can be overridden to "CRC32C". An example of a Chelsio specific

configurable item is “ACL” (for Access Control List). “ACL” is one of the few items that have no default.

Before starting any iSCSI target, an iSCSI configuration file must be created. An easy way to create this file is to use the provided sample configuration file and modify it. This file can be named anything and placed in any directory but it must be explicitly specified when using iscsictl by using the -f option. To avoid this, put configuration file in the default directory (`/etc/chelsio-iscsi`) and name it the default file name (`chiscsi.conf`).

4.2.1. “On the fly” Configuration Changes

Parameters for the most part can be changed while an iSCSI node is running. However, there are exceptions and restrictions to this rule that are explained in a later section that describes the details of the iSCSI control tool `iscsictl`.

4.3. A Quick Start Guide for Target

This section describes how to get started quickly with a Chelsio iSCSI target. It includes:

- Basic editing of the iSCSI configuration file.
- Basic commands of the iSCSI control tool including how to start and stop a target.

4.3.1. A Sample iSCSI Configuration File

The default Chelsio iSCSI configuration file is located at `/etc/chelsio-iscsi/chiscsi.conf`. If this file doesn’t already exist, then one needs to be created.

To configure an iSCSI target, there are three required parameters (in the form of key=value pairs) needed as follows:

- TargetName – A worldwide unique iSCSI target name.
- PortalGroup – The portal group tag associating with a list of target IP address (es) and port number(s) that service the login request. The format of this field is a Chelsio specific iSCSI driver parameter which is described in detail in the configuration file section.
- TargetDevice – A device served up by the associated target. A device can be:
    - A block device (e.g., /dev/sda)
    - A virtual block device (e.g., /dev/md0)
    - A RAM disk
    - A regular file

A target can serve multiple devices, each device will be assigned a Logical Unit Number (LUN) as per the order it is specified (i.e., the first device specified is assigned LUN 0, the second one LUN 1, …, and so on and so forth). Multiple TargetDevice key=value pairs are needed to indicate multiple devices.

Here is a sample of a minimum iSCSI target configuration located at `/etc/chelsio-iscsi/chiscsi.conf`:

```
target:
        TargetName=iqn.2006-02.com.chelsio.diskarray.san1
        TargetDevice=/dev/sda
        PortalGroup=1@192.0.2.178:3260
```

The `TargetDevice` value must match with the storage device in the system. The `PortalGroup` value must have a matching IP address of the Ethernet adapter card in the system.

For more information about TargetDevice configuration see Target Storage Device Configuration.

4.3.2. Basic iSCSI Control

Control of the Chelsio iSCSI software is done through `iscsictl`, the command line interface control tool. The following are the basic commands needed for effective control of the target.

Start Target: To start all of the iSCSI targets specified in the iSCSI configuration file, execute `iscsictl` with the "`-S`" option followed by "`target=ALL`".

```
[root@host~]# iscsictl -S target=ALL
```

To start a specific target execute `iscsictl` with "`-S`" followed by the target.

```
[root@host~]# iscsictl -S target=iqn.2006-02.com.chelsio.diskarray.san1
```

Stop Target: To stop the all the iSCSI target(s), execute `iscsictl` with "`-s`" option followed by "`target=ALL`".

```
[root@host~]# iscsictl -s target=ALL
```

To stop a specific target execute `iscsictl` with "`-s`" followed by the target name.

```
[root@host~]# iscsictl -s target=iqn.2006-02.com.chelsio.diskarray.san1
```

View Configuration: To see the configuration of all the active iSCSI targets, execute `iscsictl` with "`-c`" option.

```
[root@host~]# iscsictl -c
```

To see the more detailed configuration settings of a specific target, execute `iscsictl` with "`-c`" option followed by the target name.

```
[root@host~]# iscsictl -c target=iqn.2006-02.com.chelsio.diskarray.san1
```

View Global Settings: To see Chelsio global settings, execute `iscsictl` with "`-g`" option.

```
[root@host~]# iscsictl -g
```

Change Global Settings: To change Chelsio global settings, execute `iscsictl` with "`-G`" option.

```
[root@host~]# iscsictl -G iscsi_login_complete_time=300
```

View Help: To print help to stdout, execute `iscsictl` with "`-h`" option.

```
[root@host~]# iscsictl -h
```

4.4. The iSCSI Configuration File

The iSCSI configuration file consists of a series of blocks consisting of the following types of iSCSI entity blocks:

1. global
2. target

There can be only one global entity block whereas multiple target entity blocks are allowed. The global entity block is optional but there must be at least one target entity block.

An entity block begins with a block type (global or target). The content of each entity block is a list of parameters specified in a "key=value" format. An entity block ends at the beginning of the next entity block or at the end-of-file.

The parameter list in an entity block contains both:

- iSCSI parameters that override the default values
- Parameters that facilitate passing of control information to the iSCSI module

All lines in the configuration file that begin with "#" character are treated as comments and will be ignored. White space is not significant except in key=value pairs.

For the "key=value" parameters the <value> portion can be a single value or a list of multiple values. When <value> is a list of multiple values, they must be listed on one line with a comma "," to separate their values. Another way to list the values instead of commas is to list their values as key=value pairs repeatedly, each on a new line, until they are all listed.

There are three categories of "key=value" parameter, the first category belongs to the global entity block whereas the second and third categories belong to target entity block:

1. The Chelsio Global Entity Settings of key=value pairs
2. The iSCSI Entity Settings of key=value pairs
3. The Chelsio Entity Settings of key=value pairs

The following sub-sections describe these three categories and list in tables the details of their key=value parameters.

4.4.1. Chelsio System Wide Global Entity Settings

Description

Chelsio System Wide Global Entity Parameters pass system control information to the iSCSI software which affects all targets in the same way. More detail of these parameters below can be found in a later section entitled "System Wide Parameters".

Table of Chelsio Global Entity Settings

Key	Valid Values	Default Value	Multiple Values	Description
iscsi_auth_order	"ACL" "CHAP"	"CHAP"	No	Authorization order for login verification on the target. Valid only when a target's ACL_Enable=Yes **ACL**: ACL first then CHAP **CHAP**: CHAP first then ACL *Applies to Target(s) Only*
DISC_AuthMethod	"CHAP" "NONE"	None	No	To choose an authentication method for discovery phase.
DISC_Auth_CHAP_Policy	"Oneway" "Mutual"	"Oneway"	No	Oneway or Mutual (two-way) CHAP
DISC_Auth_CHAP_Target	"<user id>" :"<secret>"		Yes	CHAP user id and secret for the target. **<user id>** must be less than 256 characters. Commas "," are not allowed. **<secret>** must be between 6 and 255 characters. Commas "," are not allowed. The target user id and secret are used by the initiator to authenticate the target while doing mutual chap. *NOTE: The double quotes are required as part of the format.*
DISC_Auth_CHAP_Initiator	"<user id>" :"<secret>"		Yes	CHAP user id and secret for the initiator. **<user id>** must be less than 256 characters. Commas "," are not allowed. **<secret>** must be between 6 and 255 characters. Commas "," are not allowed. The initiator user id and secret are used by the target to authenticate the initiator. *NOTE: The double quotes are required as part of the format.*
iscsi_chelsio_ini_idstr	a string of maximum of 255 characters	"cxgb4i"	No	To enable additional optimization when Chelsio adapters and drivers are used at both ends (initiator and target) systems. Make sure the initiator name contain the substring set in iscsi_chelsio_ini_idstr when

				using Chelsio iscsi initiator driver.
iscsi_target_vendor_id	a string of maximum of 8 characters	"CHISCSI"	No	The target vendor ID part of the device identification sent by an iSCSI target in response of SCSI Inquiry command.
iscsi_login_complete_time	0 to 3600	300	No	Time allowed (in seconds) for the initiator to complete the login phase. Otherwise, the connection will be closed *NOTE: value zero means this check is NOT performed.*

4.4.2. iSCSI Entity Settings

Description

iSCSI Entity Parameters pass iSCSI protocol control information to the Chelsio iSCSI module. This information is unique for each entity block. The parameters follow the IETF iSCSI standard RFC 3720 in both definition and syntax. The descriptions below are mostly from this RFC.

Table of iSCSI Entity Settings

Key	Valid Values	Default Value	Multiple Values	Description
MaxConnections	1 to 65535	1	No	Initiator and target negotiate the maximum number of connections requested/acceptable.
InitialR2T	"Yes" "No"	"Yes"	No	To turn on or off the default use of R2T for unidirectional and the output part of bidirectional commands.
ImmediateData	"Yes" "No"	"Yes"	No	To turn on or off the immediate data.
FirstBurstLength	512 to 16777215 (2^{24} - 1)	65536	No	The maximum negotiated SCSI data in bytes of unsolicited data that an iSCSI initiator may send to a target during the execution of a single SCSI command.
MaxBurstLength	512 to 16777215 (2^{24} - 1)	262144	No	The maximum negotiated SCSI data in bytes, of a Data-In or a solicited Data-Out iSCSI sequence between the initiator and target.
DefaultTime2Wait	0 to 3600	2	No	The minimum time, in seconds, to wait before attempting an explicit / implicit logout or connection reset between initiator and target.
DefaultTime2Retain	0 to 3600	20	No	The maximum time, in seconds, after an initial wait.
MaxOutstandingR2T	1 to 65535	1	No	The maximum number of outstanding R2Ts per task.

DataPDUInOrder	"Yes" "No"	"Yes"	No	To indicate the data PDUs with sequence must be at continuously increasing order or can be in any order. *Chelsio only supports "Yes".*
DataSequenceInOrder	"Yes" "No"	"Yes"	No	To indicate the Data PDU sequences must be transferred in continuously non-decreasing sequence offsets or can be transferred in any order. *Chelsio only supports "Yes".*
ErrorRecoveryLevel	0 to 2	0	No	To negotiate the recovery level supported by the node. *Chelsio only supports 0.*
HeaderDigest	"None" "CRC32C"	"None"	Yes	To enable or disable iSCSI header Cyclic integrity checksums.
DataDigest	"None" "CRC32C"	"None"	Yes	To enable or disable iSCSI data Cyclic integrity checksums.
AuthMethod	"CHAP" and "None"	"None, CHAP"	Yes	To choose an authentication method during login phase.
TargetName	"<target name>"		No	A worldwide unique iSCSI target name. *Target only.*
TargetAlias	"<target alias>"		No	A human-readable name or description of a target. It is not used as an identifier, nor is it for authentication. *Target only.*
MaxRecvDataSegmentLength	512 to 16777215 (2^{24} - 1)	8192	No	To declare the maximum data segment length in bytes it can receive in an iSCSI PDU.
OFMarker	"Yes" "No"	"No"	No	To turn on or off the initiator to target markers on the connection. *Chelsio only supports "No".*
IFMarker	"Yes" "No"	"No"	No	To turn on or off the target to initiator markers on the connection. *Chelsio only supports "No".*
OFMarkInt	1 to 65535	2048	No	To set the interval for the initiator to target markers on a connection.
IFMarkInt	1 to 65535	2048	No	To set the interval for the target to initiator markers on a connection.

4.4.3. Chelsio Entity Settings

Description

Chelsio entity parameters pass control information to the Chelsio iSCSI module. The parameters are specific to Chelsio's implementation of the iSCSI node (target or initiator) and are unique for each entity block. The parameters consist of information that can be put into three categories:

1. Challenge Handshake Authentication Protocol (CHAP).
2. Target specific settings. All of the following parameters can have multiple instances in one target entity block (i.e., they can be declared multiple times for one particular target):

- Portal Group
- Storage Device

3. Access Control List (ACL)

Table of Chelsio Entity Settings

Key	Valid Values	Default Value	Multiple Values	Description
Chelsio CHAP Parameter (Target)				
Auth_CHAP_Target	"<user id>" :"<secret>"		No	CHAP user id and secret for the target. **<user id>** must be less than 256 characters. Commas "," are not allowed. **<secret>** must be between 6 and 255 characters. Commas "," are not allowed. The target user id and secret are used by the initiator to authenticate the target while doing mutual chap. *NOTE: The double quotes are required as part of the format.*
Auth_CHAP_Initiator	"<user id>" :"<secret>"		Yes	CHAP user id and secret for the initiator. **<user id>** must be less than 256 characters. Commas "," are not allowed. **<secret>** must be between 6 and 255 characters. Commas "," are not allowed. The initiator user id and secret are used by the target to authenticate the initiator. *NOTE: The double quotes are required as part of the format.*
Auth_CHAP_ChallengeLength	16 to 1024	16	No	CHAP challenge length
Auth_CHAP_Policy	"Oneway" or "Mutual"	"Oneway"	No	Oneway or Mutual (two-way) CHAP
Chelsio Target Specific Parameter				
PortalGroup	<portal group tag> @<target IP address> [:<port number>] . . . [,<target IP address> [:<port number>]] [,timeout=		Yes	The portal group name associates the given target with the given list of IP addresses (and optionally, port numbers) for servicing login requests. It"s required to have at least one per target. **<portal group tag>** is a unique tag identifying the portal group. It must be a positive integer. **<target IP address>** is the IP address associated with the portal

	<timeout value in milliseconds>] [,[portalgrouptag1, portalgrouptag2, … portalgrouptagn]			group tag. **<port number>** is the port number associated with the portal group tag. It is optional and if not specified the well-known iSCSI port number of 3260 is used. **<timeout>** is optional, it applies to all the portals in the group. The timeout value is in milliseconds and needs to be multiple of 100ms. It is used to detect loss of communications at the iSCSI level. *NOTE: There can be multiple target IP address/port numbers per portal group tag. This enables a target to operate on multiple interfaces for instance.* **<portalgrouptagX>**The portalgroup to which login requests should be redirected to. *NOTE: There can be multiple redirection target portalgroups specified for a particular target portal group and the redirection will happen to these in a round robin manner.*
ShadowMode	"Yes" "No"	"No"	No	To turn ShadowMode on or off for iSCSI Target Redirection
TargetSessionMaxCmd	1 to 2048	64	No	The maximum number of outstanding iSCSI commands per session.
TargetDevice*	<path/name> [,FILE|MEM|BLK] [,NULLRW] [,SYNC] [,RO] [,size=xMB] [,ID=xxxxxx] [,WWN=xxxxxxxxxx] [,SN= xxxxxx]		No	A device served up by the associated target. The device mode can be a: • Block Device (e.g., /dev/sda) • Virtual Block Device (e.g., /dev/md0) • RamDisk • Regular File **<path/name>** is the path to the device - with the exception of when a RAM Disk is specified, where it is a unique name given to the device. If multiple RAM Disks are used for a target then each name must be unique within the target. **NULLRW** specifies that random data is returned for reads, and for writes data is dropped. Useful for testing network performance.

				SYNC specifies that the device will function in the write-through mode (i.e., the data will be flushed to the device before the response is returned to the initiator). *NOTE: SYNC is only applicable with FILE mode.* **RO** specifies the device as a read-only device. **FILE** specifies this device should be accessed via the kernel‟s VFS layer. This mode is the most versatile, and it is the default mode in the cases where there is no mode specified. **BLK** specifies this device should be accessed via the kernel‟s block layer. This mode is suitable for high-speed storage device such as RAID Controllers. **MEM** specifies this device should be created as a RAM Disk. **size=**xMB is used with “MEM”, to specify the RamDisk size. If not specified, the default RamDisk size is 16MB (16 Megabytes). The minimum value of x is 1 (1MB) and the maximum value is limited by system memory. **SN** is a 16 character unique value. **ID** is a 24 character unique value. **WWN** is a 16 character unique value. It is recommended when using a multipath aware initiator , the optional ID (short form for SCSI ID), SN and WWN values should be set manually for the TargetDevice. These values will be returned in Inquiry response (VPD 0x83). Multiple TargetDevice key=value pairs are needed to indicate multiple devices. There can be multiple devices for any particular target. Each device will be assigned a Logical Unit Number (LUN) as per the order it is specified (i.e., the first device specified is assigned LUN 0, the second one LUN 1, …, and so on and so forth).

				*NOTE: **FILE** mode is the most versatile mode, if in doubt use **FILE** mode.*
ACL_Enable	"Yes" "No"	"No"	No	Defines if Chelsio'sAccess Control List (ACL) method will be enforced on the target: **Yes:** ACL is enforced on the target **No:** ACL is not enforced on the target *NOTE*: ACL flag is not allowed to be updated on the fly. Target must be restarted for new ACL flag to take effect.
ACL	[iname=<name1>][;<sip=<sip1>][;dip=<dip1>][;lun=<lun_list:permissions>]		Yes	The ACL specifies which initiators and how they are allowed to access the LUNs on the target. **iname=<Initiator Name>** specifies one or more initiator names, the name must be a fully qualified iSCSI initiator name. **sip=<Source IP address>** specifies one or more IP addresses the initiators are connecting from. **Dip=<Destination IP address>** specifies one or more IP addresses that the iSCSI target is listening on (i.e., the target portal IP addresses). *NOTE: when configuring an ACL at least one of the above three must be provided:* • *iname, and/or* • *sip, and/or* • *dip.* **lun=<lun list>:<permission>** controls how the initiators access the luns. The supported value for **<lun list>** is **ALL**. **<permissions>** can be: **R**: Read Only **RW** or **WR**: Read and Write If permissions are specified then the associated LUN list is required. If no **lun=<lun list>:[R|RW]** is specified then it defaults to **ALL:RW**. *NOTE: For the Chelsio Target Software release with lun-masking included,* ***<lun list>** is in the format of **<0..N | 0~N | ALL>*** *Where:*

				0..N**: only one value from **0** through **N ***0~N**: a range of values between **0** through **N*** ***ALL**: all currently supported LUNs.* *Multiple lists of LUN numbers are allowed. When specifying the list separate the LUN ranges by a comma.*
RegisteriSNS	"Yes" "No"	"Yes"	No	To turn on or off exporting of target information via iSNS

4.4.4. Sample iSCSI Configuration File

Following is a sample configuration file. While using iSCSI node (target), irrelevant entity block can be removed or commented.

```
#
# Chelsio iSCSI Global Settings
#
global:
       iscsi_login_complete_time=300
       iscsi_auth_order=CHAP
       DISC_AuthMethod=None
       DISC_Auth_CHAP_Policy=Oneway
       DISC_Auth_CHAP_Target="target_id1":"target_secret1"
       DISC_Auth_CHAP_Initiator="initiator_id1":"initiator_sec1"
#
# an iSCSI Target "iqn.2006-02.com.chelsio.diskarray.san1"
# being served by the portal group "5". Setup as a RAM Disk.
#
target:
       TargetName=iqn.2006-02.com.chelsio.diskarray.san1
       # lun 0: a ramdisk with default size of 16MB
       TargetDevice=ramdisk,MEM
       PortalGroup=5@192.0.2.178:3260
#
# an iSCSI Target "iqn.2005-8.com.chelsio:diskarrays.san.328"
# being served by the portal group "1" and "2"
#
target:
       #
```

```
# iSCSI configuration
#
TargetName=iqn.2005-8.com.chelsio:diskarrays.san.328
TargetAlias=iTarget1
MaxOutstandingR2T=1
MaxRecvDataSegmentLength=8192
HeaderDigest=None,CRC32C
DataDigest=None,CRC32C
ImmediateData=Yes
InitialR2T=No
FirstBurstLength=65535
MaxBurstLength=262144

#
# Local block devices being served up
# lun 0 is pointed to /dev/sda
# lun 1 is pointed to /dev/sdb

TargetDevice=/dev/sda,ID=aabbccddeeffgghh,WWN=aaabbbcccdddeeef
TargetDevice=/dev/sdb

#
# Portal groups served this target
#

PortalGroup=1@102.50.50.25:3260
PortalGroup=2@102.60.60.25:3260

#
# CHAP configuration
#
Auth_CHAP_Policy=Mutual
Auth_CHAP_target="iTarget1ID":"iTarget1Secret"

Auth_CHAP_Initiator="iInitiator1":"InitSecret1"
Auth_CHAP_Initiator="iInitiator2":"InitSecret2"
Auth_CHAP_ChallengeLength=16

#
```

```
        # ACL configuration
        #
        # initiator "iqn.2006-02.com.chelsio.san1" is allowed full access
        # to this target
        ACL=iname=iqn.2006-02.com.chelsio.san1

        # any initiator from IP address 102.50.50.101 is allowed full
access                                  # of this target
        ACL=sip=102.50.50.101

        # any initiator connected via the target portal 102.60.60.25
is        # allowed full access to this target
        ACL=dip=102.60.60.25

        # initiator "iqn.2005-09.com.chelsio.san2" from 102.50.50.22 and
        # connected via the target portal 102.50.50.25 is allowed read
only        # access of this target
        ACL=iname=iqn.2006-
02.com.chelsio.san2;sip=102.50.50.22;dip=102.50.50.25;lun=ALL:R
```

4.5. Challenge-Handshake Authenticate Protocol (CHAP)

CHAP is a protocol that is used to authenticate the peer of a connection and uses the notion of a challenge and response, (i.e., the peer is challenged to prove its identity).

The Chelsio iSCSI software supports two CHAP methods: **one-way** and **mutual**. CHAP is supported for both login and discovery sessions.

4.5.1. Normal Session CHAP Authentication

For a normal Session, the CHAP authentication is configured on a per-target basis.

4.5.2. Oneway CHAP authentication

With **one-way** CHAP (also called unidirectional CHAP) the target uses CHAP to authenticate the initiator. The initiator does not authenticate the target. This method is the default method.

For **one-way** CHAP, the initiator CHAP id and secret are configured and stored on a per-initiator with Chelsio Entity parameter "Auth_CHAP_Initiator".

4.5.3. Mutual CHAP authentication

With **mutual** CHAP (also called bidirectional CHAP), the target and initiator use CHAP to authenticate each other.

For **mutual** CHAP, in addition to the initiator CHAP id and secret, the target CHAP id and secret are required. They are configured and stored on a per target basis with Chelsio Entity parameter “Auth_CHAP_Target”.

4.5.4. Adding CHAP User ID and Secret

A single `Auth_CHAP_Target` key and multiple `Auth_CHAP_Initiator` keys could be configured per target:

```
target:
      TargetName=iqn.2006-
02.com.chelsio.diskarray.san1       TargetDevice=/dev/sda
PortalGroup=1@192.0.2.178:8000       Auth_CHAP_Policy=Oneway
      Auth_CHAP_Initiator=”remoteuser1”:”remoteuser1_secret”       Auth_CHAP
_Initiator=”remoteuser2”:”remoteuser2_secret”       Auth_CHAP_Target=“target
id1”:”target1_secret”
```

In the above example, target `iqn.2005-com.chelsio.diskarray.san1` has been configured to authenticate two initiators, and its own id and secret are configured for use in the case of mutual CHAP.

4.5.5. AuthMethod and Auth_CHAP_Policy Keys

By setting the iSCSI keys `AuthMethod` and `Auth_CHAP_Policy`, a user can choose whether to enforce CHAP and if mutual CHAP needs to be performed.

The `AuthMethod` key controls if an initiator needs to be authenticated or not. The default setting of `AuthMethod` is `None,CHAP`

The `Auth_CHAP_Policy` key controls which CHAP authentication (one-way or mutual) needs to be performed if CHAP is used. The default setting of `Auth_CHAP_Policy` is `Oneway`

On an iSCSI node, with `AuthMethod=None,CHAP`

- `Auth_CHAP_Policy=Oneway`, Chelsio iSCSI target will accept a relevant initiator if it does
    a) no CHAP
    b) CHAP Oneway or
    c) CHAP Mutual

- `Auth_CHAP_Policy=Mutual`, the Chelsio iSCSI target will accept a relevant initiator if it does
    a) no CHAP or
    b) CHAP Mutual

With `AuthMethod=None`, regardless the setting of the key Auth_CHAP_Policy, the Chelsio iSCSI target will only accept a relevant initiator if it does no CHAP.

With `AuthMethod=CHAP`, CHAP is enforced on the target:

i. `Auth_CHAP_Policy=Oneway`, the iSCSI target will accept a relevant initiator only if it does
    a) CHAP Oneway or
    b) CHAP Mutual
ii. `Auth_CHAP_Policy=Mutual`, the iSCSI node will accept a relevant initiator only if it does
    a) CHAP Mutual

4.5.6. Discovery Session CHAP

CHAP authentication is also supported for the discovery sessions where an initiator queries of all available targets.

Discovery session CHAP is configured through the global section in the configuration file. List of keys to provision discovery CHAP are:

- DISC_AuthMethod: None or CHAP.
- DISC_Auth_CHAP_Policy: Oneway or Mutual (i.e., two-way) authentication.
- DISC_Auth_CHAP_Target: target CHAP user id and secret
- DISC_Auth_CHAP_Initiator: initiator CHAP user id and secret.

Sample:

```
#
# Chelsio iSCSI Global Settings
#
global:
        DISC_AuthMethod=CHAP
        DISC_Auth_CHAP_Policy=Mutual
        DISC_Auth_CHAP_Target="target_id1":"target_secret1"
        DISC_Auth_CHAP_Initiator="initiator_id1":"initiator_sec1"
```

4.6. Target Access Control List (ACL) Configuration

The Chelsio iSCSI target supports iSCSI initiator authorization via an Access Control List (ACL).

ACL configuration is supported on a per-target basis. The creation of an ACL for a target establishes:

- Which iSCSI initiators are allowed to access it
- The type of the access: read-write or read-only
- Possible SCSI layer associations of LUNs with the initiator

More than one initiator can be allowed to access a target and each initiator‟s access rights can be independently configured.

The format for ACL rule is as follows:

```
ACL=[iname=<initiator name>][;<sip=<source ip addresses>]
    [;dip=<destination ip addresses>][;lun=<lun_list>:<permissions>]

 target:
      TargetName=iqn.2006-02.com.chelsio.diskarray.san1
      TargetDevice=/dev/sda

      PortalGroup=1@102.50.50.25:3260
      PortalGroup=2@102.60.60.25:3260

      # initiator "iqn.2006-02.com.chelsio.san1" is allowed
      # full read-write access to this target
      ACL=iname=iqn.2006-02.com.chelsio.san1

      # any initiator from IP address 102.50.50.101 is allowed full
      # read-write access of this target
      ACL=sip=102.50.50.101

      # any initiator connected via the target portal 102.60.60.25
      # is allowed full read-write access to this target
      ACL=dip=102.60.60.25

      # initiator "iqn.2005-09.com.chelsio.san2" from 102.50.50.22
      # and connected via the target portal 102.50.50.25 is allowed
      # read only access of this target
      ACL=iname=iqn.2006-      02.com.chelsio.san2;sip=102.50.50.22;dip=102.
50.50.25;lun=ALL:R
```

4.6.1. ACL Enforcement

To toggle ACL enforcement on a per-target base, a Chelsio keyword `ACL_Enable` is provided:

- Setting `ACL_Enable=Yes` enables the target to perform initiator authorization checking for all the initiators during login phase. And in addition, once the initiator has been authorized to access the target, the access rights will be checked for each individual LU the initiator trying to access.
- Setting `ACL_Enable=No` disable the target to perform initiator authorization checking.

When a target device is marked as read-only (RO), it takes precedence over ACL's write permission (i.e., all of ACL write permission of an initiator is ignored).

4.7. Target Storage Device Configuration

An iSCSI Target can support one or more storage devices. The storage device can either be the built-in RAM disk or actual backend storage.

Configuration of the storage is done through the Chelsio configuration file via the key-value pair `TargetDevice`.

When option `NULLRW` is specified, on writes the data is dropped without being copied to the storage device, and on reads the data is not actually read from the storage device but instead random data is used. This option is useful for measuring network performance.

The details of the parameters for the key `TargetDevice` are found in the table of Chelsio Entity Settings section earlier in this document.

4.7.1. RAM Disk Details

For the built-in RAM disk:

- The minimum size of the RAM disk is 1 Megabyte (MB) and the maximum is limited by system memory.
- To use a RAM disk with a Windows Initiator, it is recommended to set the size >= 16MB.

To configure an ramdisk specify `MEM` as the device mode:

```
TargetDevice=<name>,MEM,size=xMB
```

Where:	`<name>`	Is a unique name given to the RAM disk. This name identifies this particular ramdisk. If multiple RAM disks are configured for the same target, the name must be unique for each RAM Disk.
	`x`	Is the size of the RAM disk in MB. It's an integer between (1-max), where max is limited by system memory. If not specified, the default value is 16 MB.

```
target:
#<snip>
  # 16 Megabytes RAM Disk named ramdisk1
  TargetDevice=ramdisk1,MEM,size=16MB
#<snip>
```

4.7.2. FILE Mode Storage Device Details

The FILE mode storage device is the most common and versatile mode to access the actual storage attached to the target system:

- The FILE mode can accommodate both block devices and virtual block devices.
- The device is accessed in the exclusive mode. The device should not be accessed (or active) in any way on the target system.
- Each device should be used for one and only one iSCSI target.
- "SYNC" can be used with FILE mode to make sure the data is flushed to the storage device before the Target responds back to the Initiator.

To configure a FILE storage device specify `FILE` as the device mode:

```
TargetDevice=<path to the storage device>[,FILE][,SYNC]
```

Where:	`<path>`	Is the path to the actual storage device, such as /dev/sdb for a block device or /dev/md0 for a software RAID. The path must exist in the system.
	`SYNC`	When specified, the Target will flush all the data in the system cache to the storage driver before sending response back to the Initiator.

4.7.3. Example Configuration of FILE Mode Storage

Below is an example:

```
target:
#<snip>
 # software raid /dev/md0 is accessed in FILE mode
TargetDevice=/dev/md0,FILE
#<snip>
```

4.7.4. BLK Mode Storage Device Details

The BLK mode storage device is suitable for high-speed storage attached to the target system:

- The BLK mode can accommodate only block devices.
- The device is accessed in the exclusive mode. The device should not be accessed (or active) in any way on the target system.
- Each device should be used for one and only one iSCSI target.

To configure a block storage device specify BLK as the device mode:

```
TargetDevice=<path to the storage device>,BLK
```

Where: `<path>` Is the path to the actual storage device, such as /dev/sdb. The path must exist in the system.

```
target:
#<snip>
 # /dev/sdb is accessed in BLK mode
 TargetDevice=/dev/sdb,BLK
#<snip>
```

4.7.5. Multi-path Support

To enable multi-path support on the initiator, it is highly recommended that the following options are specified:

- *[,ID=xxxxxx]*: SCSI ID, a twenty-four (24) bytes alpha-numeric string
- *[,WWN=xxxxxxxxx]*: SCSI World Wide Name (WWN), a sixteen (16) bytes alpha-numeric string
- *[,SN= xxxxxx]*: SCSI SN, a sixteen (15) bytes alpha-numeric string.

The user should make sure the three values listed above are the same for the target LUNs involved in the multipath.

4.8. Target Redirection Support

An iSCSI Target can redirect an initiator to use a different IP address and port (often called a portal) instead of the current one to connect to the target. The redirected target portal can either be on the same machine, or a different one.

4.8.1. ShadowMode for Local vs. Remote Redirection

The *ShadowMode* setting specifies whether the Redirected portal groups should be present on the same machine or not. If *ShadowMode* is enabled, the redirected portal groups are on a different system. If it is disabled, then the redirected portal groups must be present on the same system otherwise the target would fail to start.

Below is an example with *ShadowMode* enabled:

```
target:
 #<snip>
     # any login requests received on 10.193.184.81:3260 will be
     # redirected to 10.193.184.85:3261.

     PortalGroup=1@10.193.184.81:3260,[2]
     PortalGroup=2@10.193.184.85:3261

     # the PortalGroup "2" is NOT presented on the same
system.          ShadowMode=Yes
 #<snip>
```

Below is an example with *ShadowMode* disabled:

```
target:
 #<snip>
     # any login requests received on 10.193.184.81:3260 will be
     # redirected to 10.193.184.85:3261

    PortalGroup=1@10.193.184.81:3260,[2]
    PortalGroup=2@10.193.184.85:3261
    # the PortalGroup "2" IS present on the same system
    ShadowMode=No
#<snip>
```

4.8.2. Redirecting to Multiple Portal Groups

The Chelsio iSCSI Target Redirection allows redirecting all login requests received on a particular portal group to multiple portal groups in a round robin manner.

Below is an example Redirection to Multiple Portal Groups:

```
target:
 #<snip>
     # any login requests received on 10.193.184.81:3260 will be
     # redirected to 10.193.184.85:3261 and 10.193.184.85:3262 in a
     # Round Robin Manner.

     PortalGroup=1@10.193.184.81:3260,[2,3]
     PortalGroup=2@10.193.184.85:3261
     PortalGroup=3@10.193.184.85:3262
     ShadowMode=No
#<snip>
```

4.9. The command line interface tools “iscsictl” & “chisns”

4.9.1. iscsictl

`iscsictl` is the tool Chelsio provides for controlling the iSCSI target. It is a Command Line Interface (CLI) that is invoked from the console. Its usage is as follows:

```
iscsictl <options> <mandatory parameters> [optional parameters]
```

The mandatory and optional parameters are the **key=value** pair(s) defined in RFC3720, or the **var=const** pair(s) defined for Chelsio iSCSI driver implementation. In this document, the key=value is referred to as “pair”, and var=const is referred to as “parameter” to clarify between iSCSI protocol‟s pair value(s), and Chelsio iSCSI driver‟s parameter value(s). Note that all **value** and **const** are case sensitive.

4.9.2. chisns

`chisns` is the command line tool for controlling the iSNS client. This is a simple tool that starts the iSNS client with a client and server parameter.

4.9.3. iscsictl options

Options	Mandatory Parameters	Optional Parameters	Description
`-h`			Display the help messages.
`-v`			Display the version.
`-f`	`<[path/] filename>`		Specifies a pre-written iSCSI configuration text file, used to start, update, save, or reload the iSCSI node(s). This option must be specified with one of the following other options: “**-S**” or “**-W**”. For the “**-S**” option “**-f**” must be specified first. All other options will ignore this “**-f**” option. If the “**-f**” option is not specified with the commands above the default configuration file will be used. It‟s name and location is: `/etc/chelsio-iscsi/chiscsi.conf` The configuration file path and filename must conform to Linux standards. For the format of the iSCSI configuration file, please see “Format of The iSCSI Configuration File” section earlier in this document.
`-k`	`<key>[=<val>]`		Specifies an iSCSI Entity or Chelsio Entity parameter. This option can be specified after “**-c**” option to retrieve a parameter setting

-c	target=<name> [,name2 . . . ,<nameN>]		Display the Chelsio iSCSI target configuration. **target=<name>** parameter: Where **name** is the name of the node whose information will be returned. **name** can be one or more string of names, separated by a comma, **<name1[,name2,…,nameN] | ALL>** A **name** of **ALL** returns information on all targets. **ALL** is a reserved string that must be uppercase. Example: **iscsictl -c target=iqn.com.cc.it1** **Iscsictl -c target=iqn.com.cc.target1 -k TargetAlias** The **<name>** parameter can also be specified as one or more parameter on the same command line, separated by a comma, **target=<name1>, <name2>, … ,<nameN>** The **target=<name>** parameter(s) are optional and if not specified all active Chelsio iSCSI targets(s) configuration(s) will be displayed. If **target=ALL** is specified or no parameters are specified the output will be abbreviated. Specify specific targets to get detailed configuration data. If the **target=<name>** option is specified, the -k <key> option can optionally be specified along with this option to display only the selected entity parameter setting. Example: **iscsictl -c target=iqn.com.cc.target1 -k HeaderDigest**
-F		target=<name> -k lun=<value>	Flush the cached data to the target disk(s). **target=<name>** parameter: Where **name** is the name of the target to be flushed. **name** can be one or more string of names, separated by a comma, **<name1[,name2,…,nameN] | ALL>** A **name** of **ALL** will cause all the target data to be flushed. **ALL** is a reserved string that must be uppercase. The **target=name** parameter is optional. If no **target=name** parameter is specified, it is the same as specifying **target=ALL**. The -k lun=<value> option is optional. It can be used to further specify a particular lun to be flushed. Example:

			To flush all the targets in the system: **iscsictl -F** To flush a particular target: **iscsictl -F target=iqn.com.cc.it1** To flush only the lun 0 of a particular target: **iscsictl -F target=iqn.com.cc.it1 -k lun=0**
-g			Display the Chelsio iSCSI Global Entity Settings.
-G	<var=const>		Set the Chelsio iSCSI Global Entity Settings. **var=const** parameter: Where **var=const** can be any key listed under Chelsio Global Entity Settings. Example: **iscsictl -G iscsi_auth_order=ACL** The **var=const** parameter(s) are mandatory. If the **var=const** parameter is not specified, the command will be denied. If any of the specified **var=const** parameter is invalid, the command will reject only the invalid parameters, but will continue on and complete all other valid parameters if any others are specified.
-s	target=<name>		Stop the specified active iSCSI targets. **target=<name>** parameter: *See the description of option -c for the target=<name> parameter definition.* The **target=<name>** parameter is mandatory. If no **target=<name>** parameter is specified, the command will be denied. If the **target=<name>** parameter is specified, only the specified targets from the **target=<name>** parameters will be stopped. If **target=ALL** is specified, all active targets will stop.
-S	target=<name>		Start or reload the iSCSI targets. **target=<name>** parameter: Where **name** is the name of the target(s) that will be started or reloaded. The **target=<name>** parameter can be specified as one or more parameter on the same command line, separated by a space, **target=<name1> target=<name2> … target=<nameN>** The **target=<name>** parameter can also be, **target=ALL**

			A **name** of **ALL** starts or reloads all targets specified in the configuration file. **ALL** is a reserved string that must be uppercase. The **target=<name>** parameter is optional. If this command line option is specified without the -f option, the default configuration file `/etc/chelsio-iscsi/chiscsi.conf` will be used. Rules, 1. If the **target=<name>** parameter is specified, only the targets from the list will be started or reloaded. 2. If **target=ALL** is specified, all targets specified from the iSCSI configuration file will be started or reloaded. 3. If the **target=<name>** parameter is not specified, all active targets configurations will be reloaded from the configuration file while those targets are running. All non-active targets specified will not be loaded / started. For Rules 1-3, if the specified targets are currently active (running), they will get reloaded. For Rules 1 & 2, if the specified targets are not currently active, they will be started. For Rules 2 & 3, please note the differences – they are not the same! The global settings are also reloaded from the configuration file with this option.
`-r`	`target=<name>`	`-k initiator=<name>`	Retrieve active iSCSI sessions under a target. **target=<name>** parameter: Where **name** must be a single target name. If **target=<name>** parameter is specified as target=<name>, the sessions can be further filtered based on the remote node name with optional –k initiator=<name> option. Examples: **iscsictl -r target=iqn.com.cc.it1** **iscsictl -r target=iqn.com.cc.it1 -k initiator=iqn.com.cc.ii1** The first **target=<name>** parameter is mandatory. If it is not specified, the command will be denied.
`-D`	`<Session handle in hex>`		Drop initiator session. This option should be specified with the handle of the session (in hex) that needs to be dropped. The session handle can be retrieved using the previous mentioned iscsictl option (-r used to retrieve active iSCSI sessions under a target).
`-W`			Overwrite the specified iSCSI configuration file with ONLY the current iSCSI global settings and the active iSCSI targets" configuration to the

			specified iSCSI configuration file. *Will delete any non-active targets' configuration from the specified file.* The -f option MUST be specified along with this option.
-h			Display the help messages.
	server=<IP address> [:<port>]	id=<isns entity id> query=<query interval>	Start the Chelsio iSNS client. **server=<IP address>[:<port>]** where **server** is the iSNS server address. The port is optional and if it"s not specified it defaults to 3205. The server with the ip address is mandatory and if it"s not specified the, the command will be denied. **id=<isns entity id>** where **id** is the iSNS entity ID used to register with the server. It defaults to <hostname>. **query=<query interval>** where **query** is the initiator query interval (in seconds). It defaults to 60 seconds. Examples: chisns server=192.0.2.10 chisns server=192.0.2.10:3205 id=isnscln2 query=30 In the first example the minimum command set is given where the IP address of the iSNS server is specified. In the second example a fully qualified command is specified by also setting three optional parameters. Here, the mandatory IP address and the corresponding optional port number are specified. Also set is the iSNS entity ID to "isnscln2" as well as the query interval to 30 seconds.

4.9.4. chisns options

Options	Mandatory Parameters	Optional Parameters	Description
-h			Display the help messages.
	server=<IP address> [:<port>]	id=<isns entity id> query=<query interval>	Start the Chelsio iSNS client. **server=<IP address>[:<port>]** where **server** is the iSNS server address. The port is optional and if it's not specified it defaults to 3205. The server with the ip address is mandatory and if it's not specified the, the command will be denied. **id=<isns entity id>** where **id** is the iSNS entity ID used to register with the server. It defaults to <hostname>. **query=<query interval>** where **query** is the initiator query interval (in seconds). It defaults to 60 seconds. Examples:

			chisns server=192.0.2.10 chisns server=192.0.2.10:3205 id=isnscln2 query=30 In the first example the minimum command set is given where the IP address of the iSNS server is specified. In the second example a fully qualified command is specified by also setting three optional parameters. Here, the mandatory IP address and the corresponding optional port number are specified. Also set is the iSNS entity ID to 'isnscln2' as well as the query interval to 30 seconds.

4.10. Rules of Target Reload (i.e. "on the fly" changes)

After a target has been started its settings can be modified via reloading of the configuration file (i.e., `iscsictl -S`).

The following parameters cannot be changed once the target is up and running otherwise the target reload would fail:

- TargetName
- TargetSessionMaxCmd
- ACL_Enable
- ACL

The following parameters **CAN** be changed by reloading of the configuration file. The new value will become effective **IMMEDIATELY** for all connections and sessions:

- TargetDevice

PortalGroupThe following parameter **CAN** be changed by reloading of the configuration file. The new value will **NOT** affect any connections and sessions that already completed login phase:

- TargetAlias
- MaxConnections
- InitialR2T
- ImmediateData
- FirstBurstLength
- MaxBurstLength
- MaxOutstandingR2T
- HeaderDigest
- DataDigest
- MaxRecvDataSegmentLength
- AuthMethod
- Auth_CHAP_Initiator
- Auth_CHAP_Target

- Auth_CHAP_ChallengeLength
- Auth_CHAP_Policy

The following parameters **SHOULD NOT** be changed because only one valid value is supported:

- DataPDUInOrder (support only "Yes")
- DataSequenceInOrder (support only "Yes")
- ErrorRecoveryLevel (support only "0")
- OFMarker (support only "No")
- IFMarker (support only "No")

The following parameters can be changed but would not have any effect because they are either not supported or they are irrelevant:

- DefaultTime2Wait (not supported)
- DefaultTime2Retain (not supported)
- OFMarkInt (irrelevant because OFMarker=No)
- IFMarkInt (irrelevant because IFMarker=No)

4.11. System Wide Parameters

The Chelsio Global Entity Settings are system wide parameters that can be controlled through the configuration file or the use of the command line "`iscsictl -G`". The finer points of some of these parameters are described in detail here:

4.11.1. iscsi_login_complete_time

Options: An integer value between 0 and 3600 (seconds). Default value is 300 (seconds).

This is the login timeout check. The parameter controls the maximum time (in seconds) allowed to the initiator to complete the login phase. If a connection has been in the login phase longer than the set value, the target will drop the connection.

Value zero turns off this login timeout check.

4.11.2. iscsi_auth_order

Options: "ACL" or "CHAP", defaults to "CHAP"

On an iSCSI target when `ACL_Enable` is set to `Yes`, `iscsi_auth_order` decides whether to perform CHAP first then ACL or perform ACL then CHAP.

- **ACL:** When setting `iscsi_auth_order=ACL`, initiator authorization will be performed at the start of the login phase for an iSCSI normal session: upon receiving the first iscsi_login_request, the target will check its ACL. If this iSCSI connection does not match any ACL provisioned, the login attempt will be terminated.

- **CHAP**: When setting `iscsi_auth_order=CHAP`, initiator authorization will be performed at the end of the login phase for an iSCSI normal session: before going to the full feature phase, the target will check its ACL. If this iSCSI connection does not match any ACL provisioned, the login attempt will be terminated.

Note *iscsi_auth_order has no meaning when* `ACL_Enable` *is set to* `No` *on a target.*

4.11.3. iscsi_target_vendor_id

Options: A string of maximum of 8 characters, defaults to `CHISCSI`

The `iscsi_target_vendor_id` is part of the device identification sent by an iSCSI target in response of a SCSI Inquiry request.

4.11.4. iscsi_chelsio_ini_idstr

Options: A string of maximum of 255 characters, defaults to “cxgb4i”.

For an iscsi connection, more optimization can be done when both initiator and target are running Chelsio adapters and drivers.

This string is used to verify the initiator name received, and identify if the initiator is running Chelsio drivers: if the initiator name contains the same substring as `iscsi_chelsio_ini_idstr` it is assumed the initiator is running with the Chelsio iscsi initiator driver and additional offload optimization is performed.

4.12. Performance Tuning

i. Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.
ii. Ensure that Unified Wire is installed with *iSCSI Performance* configuration tuning.
iii. For T6 adapters, set *ImmediateData=No* in iSCSI target configuration file (*/etc/chelsio-iscsi/chiscsi.conf*).
iv. Next, load the iSCSI PDU offload target driver (*chiscsi_t4*) and run the *chiscsi_set_affinity.sh* script to map iSCSI worker threads to different CPUs.

```
[root@host~]# chiscsi_set_affinity.sh
```

v. Configure MTU 9000 on all interfaces.

For maximum performance, it is recommended to use iSCSI PDU offload initiator.

5. Software/Driver Unloading

Use the following command to unload the module:

```
[root@host~]# rmmod chiscsi_t4
```

XI. iSCSI PDU Offload Initiator

1. Introduction

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

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

- **PDU Transmit and Recovery**
  On transmit-side, Chelsio hardware accepts the complete PDU (header + data) from the host driver, computes and inserts the digests, decomposes the PDU into multiple TCP segments if necessary, and transmit all the TCP segments onto the wire. It handles TCP retransmission if needed.
  On receive-side, Chelsio hardware recovers the iSCSI PDU by reassembling TCP segments, separating the header and data, calculating and verifying the digests, then forwarding the header to the host. The payload data, if possible, will be directly placed into the pre-posted host DDP buffer. Otherwise, the data will be sent to the host too.

The *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:

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

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

1.2. Software Requirements

1.2.1. Linux Requirements

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

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

2.1. Pre-requisites

Please make sure that the following requirements are met before installation:

- The iSCSI PDU Offload Initiator driver (cxgb4i) runs on top of NIC module (cxgb4) and open-iscsi-2.0-872/873/874 only, on a Chelsio card.
- *openssl-devel* package should be installed.

2.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install open-iSCSI,iSCSI-initiator,firmware and utilities:

```
[root@host~]# make iscsi_pdu_initiator_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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
```

If loading of cxgb4i displays "unkown symbols found" error in dmesg, follow the steps mentioned below:

i. Kill iSCSI daemon *iscsid*
ii. View all the loaded iSCSI modules

```
[root@host~]# lsmod | grep iscsi
```

iii. Now, unload them using the following command:

```
[root@host~]# rmmod <modulename>
```

iv. Finally reload the *cxgb4i* driver.

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 interface (*iface)* file

Create the file automatically by loading *cxgb4i* driver and then executing the following command:

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

Alternatively, you can 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>
```

Here,

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

Example:

```
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
```

Note

i. The interface file needs to be created in `/etc/iscsi/ifaces/` *directory.*
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.*

4.1.2. Discovery and Login

i. Starting iSCSI Daemon

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

```
[root@host~]# iscsid
```

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 discover an iSCSI target, execute the command in the following format:

```
[root@host~]# iscsiadm -m discovery -t st -p <target ip address>:<target
port no> -I <cxgb4i iface file name>
```

Example:

```
[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:

```
[root@host~]# iscsiadm -m node -T <iqn name of target> -p <target ip
address>:<target port no> -I <cxgb4i iface file name> -l
```

Example:

```
[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
```

If the login fails with an error message in the format of `ERR! MaxRecvSegmentLength <X> too big. Need to be <= <Y>`. in dmesg, edit the `iscsi/iscsid.conf` file and change the setting for `MaxRecvDataSegmentLength`:

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

Important *Always take a backup of `iscsid.conf` file before installing Chelsio Unified Wire Package. Although the file is saved to `iscsid.rpmsave` after uninstalling the package using RPM, you are still advised to take a backup.*

iv. Logging out from an iSCSI Target

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

```
[root@host~]# iscsiadm -m node -T <iqn name of target> -p <target ip
address>:<target port no> -I <cxgb4i iface file name> -u
```

Example:

```
[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
```

Example:

```
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
}
```

4.3. Performance Tuning

Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.

In case iSCSI Initiator IRQs pose a bottleneck for multiple connections, you can improve IOPS performance using the steps mentioned below.

i. Enable iSCSI multi-queue. In 3.18+ kernels, add the below entry to the grub configuration file and reboot the machine:

```
scsi_mod.use_blk_mq=1
```

ii. Run the performance tuning script to map iSCSI Initiator queues to different CPUs:

```
[root@host~]# t4_perftune.sh -Q iSCSI -n
```

iii. Load initiator driver:

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

iv. Configure MTU 9000 on all interfaces.
v. Login to multiple targets.
vi. Run IOPS test.

5. Software/Driver Unloading

To unload the driver, execute the following commands:

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

XII. Crypto Offload

1. Introduction

Chelsio's Terminator 6 (T6) Unified Wire ASIC enables concurrent secure communication and secure storage with support for integrated TLS/SSL/DTLS and inline cryptographic functions, leveraging the proprietary TCP/IP offload engine. Chelsio's full offload TLS/SSL/DTLS is uniquely capable of 100Gb line-rate performance. In addition, the accelerator can be used in a traditional co-processor Lookaside mode to accelerate TLS/SSL, IPsec, SMB 3.X crypto, data at rest encryption/decryption, and data-deduplication fingerprint computation.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio's Crypto Offload driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR

1.2. Software Requirements

1.2.1. Linux Requirements

Following is the list of Crypto Offload components and supported Linux versions:

Linux Version	Crypto Components
RHEL 7.3, 3.10.0-514.el7	Co-processor
Ubuntu 16.04.1, 4.4.0-31-generic	Inline-TLS
RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)	Inline-TLS, Co-processor
Kernel.org linux-4.9.49 compiled on RHEL 7.3	

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

2. Kernel Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. In case of 4.9.49 kernel, use the following command to build and install the kernel with Crypto components enabled by default:

```
[root@host~]# make kernel_install
```

Note

- *To support digests on RHEL 7.3 x86_64 (base 3.10.0-514.el7 kernel), the kernel needs to be patched with https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=4afa5f9617927453ac04b24b584f6c718dfb4f45*
- *To use Co-processor/Inline-TLS on RHEL7.3 ARM (base 4.5.0-15.el7.aarch64 kernel), no extra kernel configuration is needed.*
- *To use Inline-TLS on Ubuntu 16.04.1, 4.4.0-31-generic kernel, no extra kernel configuration is needed.*

iii. Reboot the machine and boot into the new kernel.

3. Software/Driver Installation

3.1. Pre-requisites

Please make sure that SELinux and firewall are disabled.

3.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install Crypto driver and Chelsio openSSL modules:

```
[root@host~]# make crypto_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

4. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

4.1. Co-processor

i. To load the Crypto Offload driver in Co-processor mode (*chcr*), run the following command:

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

ii. Bring up the Chelsio network interface:

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

Where ethX is the Chelsio interface.

4.2. Inline

To load the Crypto Offload driver in Inline mode, load the Network driver in TOE mode:

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

5. Software/Driver Configuration and Fine-tuning

5.1. Co-processor

To view the complete list of supported cryptographic algorithms, use the following command:

```
[root@host~]# cat /proc/crypto|grep -i chcr
```

The following applications can be offloaded by Chelsio Co-processor:

- Apache
- Dmcrypt
- SFTP
- OpenVPN
- Strongswan

To enable Digests in RHEL 7.3 x86_64 (base 3.10.0-514.el7 kernel), uncomment the "DIGESTS" line in *af_alg_engine* section of openSSL configuration file, */etc/pki/tls/openssl.cnf*:

```
[af_alg_engine]
CIPHERS=aes-128-cbc aes-192-cbc aes-256-cbc aes-128-ctr aes-192-ctr aes-256-
ctr
#DIGESTS=sha1 sha224 sha256 sha384 sha512
```

To verify if Chelsio Co-processor is used by the above applications, run the following command:

```
[root@host~]# cat /sys/kernel/debug/cxgb4/<PF4_id>/crypto
Chelsio Crypto Co-processor Stats
aes_ops: 1016
digest_ops: 323
aead_ops: 2739611
comp: 2740950
error: 0
Fallback: 9
```

5.2. Inline

- **Configuring TLS Offload ports**

To configure TLS offload ports, COP rule for each TCP port must be defined. Follow the steps mentioned below:

i. Create a new policy file and add the following line for each TCP port:

```
src or dst port <tcp_port> => offload tls mss 32 bind random
```

Example:

To offload TCP ports 443, 989, 1000, 1001 and 1002, a policy file *new_policy_file* is created with the following lines :

```
src or dst port 444 => offload tls mss 32 bind random
src or dst port 989 => offload tls mss 32 bind random
src or dst port 1000 => offload tls mss 32 bind random
src or dst port 1001 => offload tls mss 32 bind random
src or dst port 1002 => offload tls mss 32 bind random
```

ii. Compile the policy:

```
[root@host~]# cop -d -o <policy_out> <new_policy_file>
```

Example:

```
[root@ ~]# cop -d -o policy_out new_policy_file
policy rules read:
  rule 0: src or dst port 444 => offload tls mss 32 bind random
  rule 1: src or dst port 989 => offload tls mss 32 bind random
  rule 2: src or dst port 1000 => offload tls mss 32 bind random
  rule 3: src or dst port 1001 => offload tls mss 32 bind random
  rule 4: src or dst port 1002 => offload tls mss 32 bind random

classifier program:
   0    8/01bc0000%ffff0000  yes->[0]     no->step 1
   1    8/000001bc%0000ffff  yes->[0]     no->step 2
   2    8/03dd0000%ffff0000  yes->[1]     no->step 3
   3    8/000003dd%0000ffff  yes->[1]     no->step 4
   4    8/03e80000%ffff0000  yes->[2]     no->step 5
   5    8/000003e8%0000ffff  yes->[2]     no->step 6
   6    8/03e90000%ffff0000  yes->[3]     no->step 7
   7    8/000003e9%0000ffff  yes->[3]     no->step 8
   8    8/03ea0000%ffff0000  yes->[4]     no->step 9
   9    8/000003ea%0000ffff  yes->[4]     no->[5]
optimized classifier program:
   0  8 #1  ffff0000  yes->[0]     no->step 5
   4    01bc0000
   5  8 #1  0000ffff  yes->[0]     no->step 10
   9    000001bc
  10  8 #1  ffff0000  yes->[1]     no->step 15
  14    03dd0000
  15  8 #1  0000ffff  yes->[1]     no->step 20
  19    000003dd
  20  8 #1  ffff0000  yes->[2]     no->step 25
  24    03e80000
  25  8 #1  0000ffff  yes->[2]     no->step 30
  29    000003e8
  30  8 #1  ffff0000  yes->[3]     no->step 35
  34    03e90000
  35  8 #1  0000ffff  yes->[3]     no->step 40
  39    000003e9
  40  8 #1  ffff0000  yes->[4]     no->step 45
  44    03ea0000
  45  8 #1  0000ffff  yes->[4]     no->[5]
  49    000003ea
offload settings:
   0: offload 1, ddp -1, coalesce -1, cong_algo -1, queue -2, class -1, tstamp -1, sack -1, tls 1, nagle -1, mss 32
   1: offload 1, ddp -1, coalesce -1, cong_algo -1, queue -2, class -1, tstamp -1, sack -1, tls 1, nagle -1, mss 32
   2: offload 1, ddp -1, coalesce -1, cong_algo -1, queue -2, class -1, tstamp -1, sack -1, tls 1, nagle -1, mss 32
   3: offload 1, ddp -1, coalesce -1, cong_algo -1, queue -2, class -1, tstamp -1, sack -1, tls 1, nagle -1, mss 32
   4: offload 1, ddp -1, coalesce -1, cong_algo -1, queue -2, class -1, tstamp -1, sack -1, tls 1, nagle -1, mss 32
   5: offload 0, ddp -1, coalesce -1, cong_algo -1, queue -2, class -1, tstamp -1, sack -1, tls 0, nagle -1, mss 0
```

iii. Apply the policy:

```
[root@host~]# cxgbtool <iface> policy <policy_out>
```

Example:

```
[root@ ~]# cxgbtool enp129s0f4 policy policy_out
[root@cool1 ~]#
```

- **Configuring Chelsio OpenSSL**

OpenSSL which supports Chelsio inline offload is installed as part of Unified Wire package. It is installed in */usr/chssl/bin*

i. Start TLS offload Server

```
[root@host~]# cd /usr/chssl/bin
[root@host~]# ./openssl s_server -key <key_file> -cert <cert_file> -accept
443 -cipher AES128-GCM-SHA256  -WWW
```

Example:

```
[root@ ~]# ./openssl s_server -key /root/server.key -cert /root/server.crt
-accept 443 -cipher AES128-GCM-SHA256  -WWW
```

ii. Start TLS offload Client:

```
[root@host~]# cd /usr/chssl/bin
[root@host~]# ./openssl s_time -connect <tls_server_ip>:<tls_server_port>  -
www /<file>
```

Example:

```
[root@ bin]# ./openssl s_time -connect 102.1.1.154:443  -www /1G
No CIPHER specified
Collecting connection statistics for 30 seconds
**********************************************

47 connections in 12.42s; 3.78 connections/user sec, bytes read 50465867843
47 connections in 31 real seconds, 1073741869 bytes read per connection


Now timing with session id reuse.
starting
rrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrr

47 connections in 12.32s; 3.81 connections/user sec, bytes read 50465867843
47 connections in 31 real seconds, 1073741869 bytes read per connection
```

To compile applications using Chelsio OpenSSL library:

```
[root@host~]# export LD_LIBRARY_PATH=/usr/chssl/lib/
[root@host~]# gcc -g -o <server/client output file> <server/client file> -
lcrypto -lssl -L/usr/chssl/lib/
```

Example:

Client:

```
[root@localhost ~]# export LD_LIBRARY_PATH=/usr/chssl/lib/
[root@localhost ~]# gcc -g -o client client.c -lcrypto -lssl -L/usr/chssl/lib/
```

Server:

```
[root@localhost ~]# export LD_LIBRARY_PATH=/usr/chssl/lib/
[root@localhost ~]# gcc -g -o server server.c -lcrypto -lssl -L/usr/chssl/lib/
```

- **Inline TLS Counters**

To verify if Chelsio Inline is used, run the following command:

```
[root@host~]# cat /sys/kernel/debug/cxgb4/<PF4_id>/tls
Chelsio Inline TLS Stats
TLS PDU Tx: 32661534
TLS PDU Rx: 231039210
TLS Keys (DDR) Count: 48
```

5.3. Performance Tuning

Apply the generic performance settings mentioned in the Performance Tuning section in the **Unified Wire** chapter before proceeding.

Inline-TLS

i. Disable Nagle for offload connections:

a. Create a COP policy:

```
[root@host~]# cat <policy_file>
all=>offload !nagle
```

b. Compile the policy:

```
[root@host~]# cop -d -o <policy_out> <policy_file>
```

c. Apply the policy:

```
[root@host~]# cxgbtool ethX policy <policy_out>
```

ii. Configure TLS ports as described in Configuring TLS Offload ports section.

iii. Run the performance tuning script to map TLS queues to different CPUs:

```
[root@host~]# t4_perftune.sh -n -Q tls
```

iv. Ensure that the application sends 8k PDU for best performance.

- **Co-processor**

i. Use only AES128-GCM algorithm
ii. Run the performance tuning script to map crypto queues to different CPUs:

```
[root@host~]# t4_perftune.sh -n -Q crypto
```

6. Software/Driver Unloading

To unload Crypto Offload driver in Co-processor mode, run the following command:

```
[root@host~]# rmmod chcr
```

To unload Crypto Offload driver in Inline mode, unload the network driver in TOE mode. See Software/Driver Unloading section in **Network (NIC/TOE)** chapter for more information.

XIII. Data Center Bridging (DCB)

1. Introduction

Data Center Bridging (DCB) refers to a set of bridge specification standards, aimed to create a converged Ethernet network infrastructure shared by all storage, data networking and traffic management services. An improvement to the existing specification, DCB uses priority-based flow control to provide hardware-based bandwidth allocation and enhances transport reliability.

One of DCB's many benefits includes low operational cost, due to consolidated storage, server and networking resources, reduced heat and noise, and less power consumption.
Administration is simplified since the specifications enable transport of storage and networking traffic over a single unified Ethernet network.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with Chelsio's DCB feature:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T520-LL-CR
- T520-CR
- T580-CR
- T580-LP-CR
- T420-CR
- T420-LL-CR

1.2. Software Requirements

1.2.1. Linux Requirements

Currently Chelsio's DCB feature is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default

- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Build and install all drivers with DCB support:

```
[root@host~]# make dcbx=1 install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers:*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

Before proceeding, please ensure that Unified Wire is installed with DCB support as mentioned in previous section. The switch ports need to be enabled with DCBX configuration (Class mapping, ETS and PFC).

Upon loading the network/storage driver and interface bringup, firmware completes DCBX negotiation with the switch.

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe t4_tom
[root@host~]# ifconfig ethX up
[root@host~]# modprobe csiostor
```

The negotiated DCBX parameters can be reviewed at */sys/kernel/debug/cxgb4/<PF4_id>/dcb_info*

Example:

```
[root@[redacted] ~]# cat /sys/kernel/debug/cxgb4/0000\:04\:00.4/dcb_info
Data Center Bridging Information

Port: 0 (DCB negotiated: yes)
[ DCBx Version DCBx-CEE 1.01 ]

  Index                 :        0   1   2   3   4   5   6   7

  Priority Group IDs    :        6   6   0   6   6   6   1   6
  Priority Group BW(%)  :       60  40   0   0   0   0   0   0
  Max PG Traffic Classes  [  8  ]

  Priority Flow Control   :      0   0   1   0   0   0   1   0
  Max PFC Traffic Classes [  8  ]

  Application Information:
  App     Priority    Selection         Protocol
  Index   Map         Field             ID
    0     0x40        Socket TCP (1)    0x0cbc (3260)
    1     0x04        Socket TCP (1)    0xc350 (50000)

Port: 1 (DCB negotiated: no)
```

The storage driver (FCoE Full Offload Initiator) uses the DCBX negotiated parameters (ETS, PFC etc.) without any further configuration. The network drivers (*cxgb4*, *t4_tom*) and iSCSI drivers (*cxgb4i*, *chiscsi*) needs further VLAN configuration to be setup, which is explained in the Running NIC & iSCSI Traffic together with DCBx section.

4. Software/Driver Configuration and Fine-tuning

4.1. Configuring Cisco Nexus 5010 switch

4.1.1. Configuring the DCB parameters

Note *By default, the Cisco Nexus switch enables DCB 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 5010 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 2
```

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
```

4.1.2. 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 DCB:

```
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
```

4.2. 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
```

5. Running NIC & iSCSI Traffic together with DCBx

Note *Please refer iSCSI PDU Offload Initiator chapter to configure iSCSI Initiator.*

Use the following procedure to run NIC and iSCSI traffic together with DCBx enabled.

i. Identify the VLAN priority configured for NIC and iSCSI class of traffic on the switch.
ii. Create VLAN interfaces for running NIC and iSCSI traffic, and configure corresponding VLAN priority.

Example:

Switch is configured with a VLAN priority of 2 and 5 for NIC and iSCSI class of traffic respectively. NIC traffic is run on VLAN10 and iSCSI traffic is run on VLAN20.

Assign proper VLAN priorities on the interface (here eth5), using the following commands on the host machine:

```
[root@host~]# vconfig set_egress_map eth5.10 0 2
[root@host~]# vconfig set_egress_map eth5.20 5 5
```

XIV. 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.

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:

- T6225-CR
- T6225-LL-CR
- T520-BT
- T520-LL-CR
- T520-CR

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the FCoE full offload Initiator driver is available for the following version(s):

- RHEL 7.3, 3.10.0-514.el7

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

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install FCoE PDU offload target driver:

```
[root@host~]# make fcoe_full_offload_initiator_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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
```

4. Software/Driver Configuration and Fine-tuning

4.1. Configuring Cisco Nexus 5010 and Brocade switch

To configure various Cisco and Brocade switch settings, please refer Software/Driver Configuration and Fine-tuning section of **Data Center Bridging (DCB)** chapter.

4.2. FCoE fabric discovery verification

4.2.1. 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@      ~]# cat /sys/class/fc_host/host0/port_name
0x500074304639f080
[root@      ~]# cat /sys/class/fc_host/host0/port_state
Online
```

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~]# cat /sys/kernel/debug/csiostor/<pci_id>/lnodes
```

```
[root@ ~]# cat /sys/kernel/debug/csiostor/0000\:81\:00.6/lnodes
*****************[Index:  0]***************
device id: 1613956
vnpi: 41092
fcfi: 41120
mac: 0efcfa340020
nport id: 340020
wwnn: 50007433cadb6000
wwpn: 50007433cadb6080
num rnodes:   4
npiv: SUPPORTED
common service params:
        hi ver:00
        low ver:00
        bb credit:10
        word1(31:16) flags:8000
        rcv size:2068
        maxsq_reloff:16711711
        ratov:16711711
        edtov: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 exch:0

class 4:NOT SUPPORTED
```

4.2.2. Verifying the target discovery

To view the list of targets discovered on a particular FCoE port, follow the below mentioned steps:

i. Determine 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
```

ii. 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@      ~]# cat /sys/class/fc_remote_ports/rport-0\:0-0/roles
Fabric Port
[root@      ~]# cat /sys/class/fc_remote_ports/rport-0\:0-1/roles
Directory Server
[root@      ~]# cat /sys/class/fc_remote_ports/rport-0\:0-2/roles
Management Server
[root@      ~]# cat /sys/class/fc_remote_ports/rport-0\:0-3/roles
FCP Initiator
[root@      ~]# cat /sys/class/fc_remote_ports/rport-0\:0-4/roles
FCP Initiator
[root@      ~]# cat /sys/class/fc_remote_ports/rport-0\:0-9/roles
FCP Target
```

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

Alternatively, the local ports can also be found using*:*

```
[root@host~]# cat /sys/kernel/debug/csiostor/<pci_id>/lnodes
```

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

```
[root@host~]# cat /sys/kernel/debug/csiostor/<pci_id>/rnodes
```

```
[root@ ~]# cat /sys/kernel/debug/csiostor/0000\:81\:00.6/rnodes
******************Index :  0**************
ssni: 40064
vnpi: 41092
fcfi: 41120
wwnn: 2058000decb1bd41
wwpn: 2003000decb1bd7f
nport id: fffffe
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: 40065
vnpi: 41092
fcfi: 41120
wwnn: 2058000decb1bd41
wwpn: 250d000decb1bd40
nport id: fffffc
fcp flags: 0
role: nameserver
class service params:
class 1:NOT SUPPORTED
class 2:NOT SUPPORTED
class 3:SUPPORTED
class 4:NOT SUPPORTED
******************Index :  2**************
ssni: 40067
vnpi: 41092
fcfi: 41120
wwnn: 2058000decb1bd41
wwpn: 250b000decb1bd40
nport id: fffffa
fcp flags: 0
role: nport
class service params:
class 1:NOT SUPPORTED
class 2:NOT SUPPORTED
class 3:SUPPORTED
class 4:NOT SUPPORTED
******************Index :  3**************
ssni: 40068
vnpi: 41092
fcfi: 41120
wwnn: 500a0980892bb831
wwpn: 500a0981992bb831
nport id: 340000
fcp flags: 0
role: target
class service params:
class 1:NOT SUPPORTED
class 2:NOT SUPPORTED
class 3:SUPPORTED
class 4:NOT SUPPORTED
```

4.3. Formatting the LUNs and Mounting the Filesystem

Use *lsscsi -g* to list the LUNs discovered by the initiator

```
[root@host~]# lsscsi –g
```

```
[root@      ~]# lsscsi -g
[0:0:0:0]    disk    NETAPP   LUN             8010  /dev/sda   /dev/sg0
[0:0:0:1]    disk    NETAPP   LUN             8010  /dev/sdb   /dev/sg1
[0:0:0:2]    disk    NETAPP   LUN             8010  /dev/sdc   /dev/sg2
[0:0:0:3]    disk    NETAPP   LUN             8010  /dev/sdd   /dev/sg3
[0:0:0:4]    disk    NETAPP   LUN             8010  /dev/sde   /dev/sg4
[0:0:0:5]    disk    NETAPP   LUN             8010  /dev/sdf   /dev/sg5
[0:0:0:6]    disk    NETAPP   LUN             8010  /dev/sdg   /dev/sg6
[0:0:0:7]    disk    NETAPP   LUN             8010  /dev/sdh   /dev/sg7
[0:0:0:8]    disk    NETAPP   LUN             8010  /dev/sdi   /dev/sg8
[0:0:0:9]    disk    NETAPP   LUN             8010  /dev/sdj   /dev/sg9
[0:0:0:10]   disk    NETAPP   LUN             8010  /dev/sdk   /dev/sg10
[0:0:0:11]   disk    NETAPP   LUN             8010  /dev/sdl   /dev/sg11
[0:0:0:12]   disk    NETAPP   LUN             8010  /dev/sdm   /dev/sg12
[0:0:0:13]   disk    NETAPP   LUN             8010  /dev/sdn   /dev/sg13
[0:0:0:14]   disk    NETAPP   LUN             8010  /dev/sdo   /dev/sg14
[0:0:0:15]   disk    NETAPP   LUN             8010  /dev/sdp   /dev/sg15
[0:0:0:16]   disk    NETAPP   LUN             8010  /dev/sdq   /dev/sg16
[0:0:0:17]   disk    NETAPP   LUN             8010  /dev/sdr   /dev/sg17
[0:0:0:18]   disk    NETAPP   LUN             8010  /dev/sds   /dev/sg18
[0:0:0:19]   disk    NETAPP   LUN             8010  /dev/sdt   /dev/sg19
[1:0:0:0]    disk    NETAPP   LUN             8010  /dev/sdu   /dev/sg20
[1:0:0:1]    disk    NETAPP   LUN             8010  /dev/sdv   /dev/sg21
[1:0:0:2]    disk    NETAPP   LUN             8010  /dev/sdw   /dev/sg22
[1:0:0:3]    disk    NETAPP   LUN             8010  /dev/sdx   /dev/sg23
[1:0:0:4]    disk    NETAPP   LUN             8010  /dev/sdy   /dev/sg24
[1:0:0:5]    disk    NETAPP   LUN             8010  /dev/sdz   /dev/sg25
[1:0:0:6]    disk    NETAPP   LUN             8010  /dev/sdaa  /dev/sg26
[1:0:0:7]    disk    NETAPP   LUN             8010  /dev/sdab  /dev/sg27
[1:0:0:9]    disk    NETAPP   LUN             8010  /dev/sdac  /dev/sg28
[1:0:0:10]   disk    NETAPP   LUN             8010  /dev/sdad  /dev/sg29
[1:0:0:11]   disk    NETAPP   LUN             8010  /dev/sdae  /dev/sg30
[1:0:0:12]   disk    NETAPP   LUN             8010  /dev/sdaf  /dev/sg31
[1:0:0:15]   disk    NETAPP   LUN             8010  /dev/sdag  /dev/sg32
[3:0:0:0]    disk    NETAPP   LUN             8010  /dev/sdah  /dev/sg33
[3:0:0:1]    disk    NETAPP   LUN             8010  /dev/sdai  /dev/sg34
[3:0:0:2]    disk    NETAPP   LUN             8010  /dev/sdaj  /dev/sg35
[3:0:0:3]    disk    NETAPP   LUN             8010  /dev/sdak  /dev/sg36
[3:0:0:4]    disk    NETAPP   LUN             8010  /dev/sdal  /dev/sg37
[3:0:0:5]    disk    NETAPP   LUN             8010  /dev/sdam  /dev/sg38
[3:0:0:6]    disk    NETAPP   LUN             8010  /dev/sdan  /dev/sg39
[3:0:0:7]    disk    NETAPP   LUN             8010  /dev/sdao  /dev/sg40
[3:0:0:8]    disk    NETAPP   LUN             8010  /dev/sdap  /dev/sg41
[3:0:0:9]    disk    NETAPP   LUN             8010  /dev/sdaq  /dev/sg42
[3:0:0:10]   disk    NETAPP   LUN             8010  /dev/sdar  /dev/sg43
[3:0:0:11]   disk    NETAPP   LUN             8010  /dev/sdas  /dev/sg44
[3:0:0:12]   disk    NETAPP   LUN             8010  /dev/sdat  /dev/sg45
[3:0:0:13]   disk    NETAPP   LUN             8010  /dev/sdau  /dev/sg46
[3:0:0:14]   disk    NETAPP   LUN             8010  /dev/sdav  /dev/sg47
[3:0:0:15]   disk    NETAPP   LUN             8010  /dev/sdaw  /dev/sg48
[3:0:0:16]   disk    NETAPP   LUN             8010  /dev/sdax  /dev/sg49
```

Alternatively, the LUNs discovered by the Chelsio FCoE initiators can be accessed via easily-identifiable ‘udev’ path device files like:

```
[root@host~]# ls /dev/disk/by-path/pci-0000:04:00.0-csio-fcoe
<local_wwpn>:<remote_wwpn>:<lun_wwn>
```

```
[root@      ~]# mount /dev/disk/by-path/pci-0000:03:00.6-csio-fcoe-0x50007430463b7380:0x500a0981892bb831:0x0000000000000000 /mnt/
[root@      ~]# mount
/dev/sdbe5 on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sdbe1 on /boot type ext3 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
/tmp on /tmp type none (rw,bind)
/var/tmp on /var/tmp type none (rw,bind)
/home on /home type none (rw,bind)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
none on /sys/kernel/debug type debugfs (rw)
gvfs-fuse-daemon on /root/.gvfs type fuse.gvfs-fuse-daemon (rw,nosuid,nodev)
/dev/sdah on /mnt type ext3 (rw)
```

4.4. Creating Filesystem

Create an ext3 filesystem using the following command:

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

```
[root@      ~]# mkfs.ext3 /dev/sdah
mke2fs 1.41.12 (17-May-2010)
/dev/sdah 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)
Stride=1 blocks, Stripe width=16 blocks
327680 inodes, 1310720 blocks
65536 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=1342177280
40 block groups
32768 blocks per group, 32768 fragments per group
8192 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 25 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

4.5. 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@      ~]# mount /dev/sdah /mnt/
[root@      ~]# mount
/dev/sdbo5 on / type ext4 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sdbo1 on /boot type ext3 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
/tmp on /tmp type none (rw,bind)
/var/tmp on /var/tmp type none (rw,bind)
/home on /home type none (rw,bind)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
none on /sys/kernel/debug type debugfs (rw)
gvfs-fuse-daemon on /root/.gvfs type fuse.gvfs-fuse-daemon (rw,nosuid,nodev)
/dev/sdah on /mnt type ext3 (rw)
```

5. Software/Driver Unloading

To unload the driver:

```
[root@host~]# modprobe -r  csiostor
```

Note *If multipath services are running, unload of FCoE driver is not possible. Stop the multipath service and then unload the driver.*

XV. Offload Bonding

1. Introduction

The Chelsio Offload bonding driver provides a method to aggregate multiple network interfaces into a single logical bonded interface effectively combining the bandwidth into a single connection. It also provides redundancy in case one of link fails.

The traffic running over the bonded interface can be fully offloaded to the adapter, thus freeing the CPU from TCP/IP overhead.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with the Chelsio Offload Bonding driver:

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

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the Offload Bonding driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install Chelsio Offload bonding driver:

```
[root@host~]# make bonding_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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 (with offload support), run the following command:

```
[root@host~]# modprobe bonding
```

4. Software/Driver Configuration and Fine-tuning

4.1. Offloading TCP traffic over a bonded interface

The Chelsio Offload Bonding driver supports all the bonding modes in NIC Mode. In offload mode (t4_tom loaded) however, only the **balance-rr (mode=0)**, **active-backup (mode=1)**, **balance-xor (mode=2)** and **802.3ad (mode=4)** modes are supported.

To offload TCP traffic over a bonded interface, use the following method:

i. Load the network driver with TOE support:

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

ii. Create a bonded interface:

```
[root@host~]# modprobe bonding mode=1 miimon=100
```

iii. Bring up the bonded interface and enslave the interfaces to the bond:

```
[root@host~]# ifconfig bond0 up
[root@host~]# ifenslave bond0 ethX ethY
```

Note ***ethX*** *and* ***ethY*** *are interfaces of the same adapter.*

iv. Assign IPv4/IPv6 address to the bonded interface:

```
[root@host~]# ifconfig bond0 X.X.X.X/Y
[root@host~]# ifconfig bond0 inet6 add <128-bit IPv6 Address> up
```

v. Disable FRTO on the PEER:

```
[root@host~]# sysctl -w net.ipv4.tcp_frto=0
```

All TCP traffic will be offloaded over the bonded interface now.

5. Software/Driver Unloading

To unload the driver, run the following command:

```
[root@host~]# rmmod bonding
```

XVI. Offload Multi-Adapter Failover (MAFO)

1. Introduction

Chelsio's adapters offer a complete suite of high reliability features, including adapter-to-adapter failover. The patented offload Multi-Adapter Failover (MAFO) feature ensures all offloaded traffic continue operating seamless in the face of port failure.

MAFO allows aggregating network interfaces across multiple adapters into a single logical bonded interface, providing effective fault tolerance.

The traffic running over the bonded interface can be fully offloaded to the adapter, thus freeing the CPU from TCP/IP overhead.

Important

- *Portions of this software are covered under US Patent, Failover and migration for full-offload network interface devices: US 8346919 B1*
- *Use of the covered technology is strictly limited to Chelsio ASIC-based soutions.*

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with the Offload Multi-Adapter Failover driver.

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

- T420-CX

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the Offload Multi-Adapter Failover driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

1.2.2. Driver Requirements

Multi-adapter failover feature will work for *Link Down* events caused by:

- Cable unplug on bonded interface.
- Bringing corresponding switch port down.

Note *The feature will not work if the bonded interfaces are administratively taken down.*

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install MAFO feature:

```
[root@host~]# make bonding_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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 (with offload support), run the following command:

```
[root@host~]# modprobe bonding
```

4. Software/Driver Configuration and Fine-tuning

4.1. Offloading TCP traffic over a bonded interface

The Chelsio MAFO driver supports only the **active-backup (mode=1)** mode.

To offload TCP traffic over a bonded interface, use the following method:

i. Load the network driver with TOE support:

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

ii. Create a bonded interface:

```
[root@host~]# modprobe bonding mode=1 miimon=100
```

iii. Bring up the bonded interface and enslave the interfaces to the bond:

```
[root@host~]# ifconfig bond0 up
[root@host~]# ifenslave bond0 ethX ethY
```

Note ***ethX*** *and* ***ethY*** *are interfaces of different adapters.*

iv. Assign IPv4/IPv6 address to the bonded interface

```
[root@host~]# ifconfig bond0 X.X.X.X/Y
[root@host~]# ifconfig bond0 inet6 add <128-bit IPv6 Address> up
```

vi. Disable FRTO on the PEER:

```
[root@host~]# sysctl -w net.ipv4.tcp_frto=0
```

vii. Disable TCP timestamps:

```
[root@host~]# sysctl -w net.ipv4.tcp_timestamps=0
```

All TCP traffic will be offloaded over the bonded interface now and fail-over will happen in case of link-down event.

5. Software/Driver Unloading

To unload the driver, run the following command:

```
[root@host~]# rmmod bonding
```

XVII. UDP Segmentation Offload and Pacing

1. Introduction

Chelsio's Terminator series of adapters provide UDP segmentation offload and per-stream rate shaping to drastically lower server CPU utilization, increase content delivery capacity, and improve service quality.

Tailored for UDP content, UDP Segmentation Offload (USO) technology moves the processing required to packetize UDP data and rate control its transmission from software running on the host to the network adapter. USO increases performance and dramatically reduces CPU overhead, allowing significantly higher capacity using the same server hardware. Without USO support, UDP server software running on the host needs to packetize payload into frames, process each frame individually through the network stack and schedule individual frame transmission, resulting in millions of system calls, and packet traversals through all protocol layers in the operating system to the network device. In contrast, USO implements the network protocol stack in the adapter, and the host server software simply hands off unprocessed UDP payload in large I/O buffers to the adapter.

The following figure compares the traditional datapath on the left to the USO datapath on the right, showing how per-frame processing is eliminated. In this example, the video server pushes 5 frames at a time. In an actual implementation, a video server pushes 50 frames or more in each I/O, drastically lowering the CPU cycles required to deliver the content.

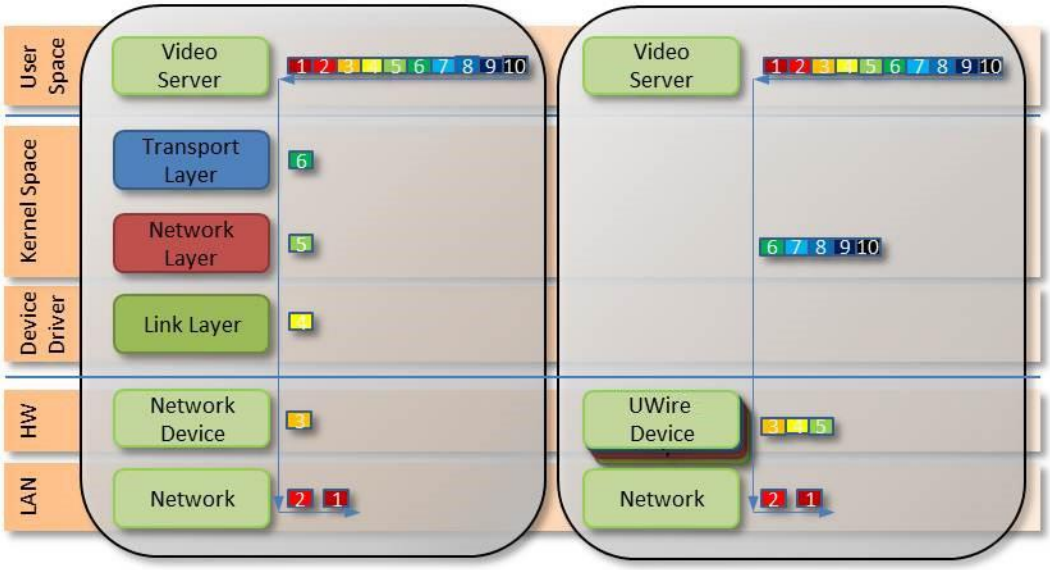

Pacing is beneficial for several reasons, one example is for Content Delivery Networks (CDNs)/Video On Demand (VOD) providers to avoid receive buffer overflows, smooth out network traffic, or to enforce Service Level Agreements (SLAs).

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with the UDP Segmentation Offload and Pacing driver:

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

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the UDP Segmentation Offload and Pacing driver is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)

- Kernel.org linux-4.8

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

2. Software/Driver Installation

The offload drivers support UDP Segmentation Offload with limited number of connections (1024 connections).To build and install UDP Offload drivers which support large number of offload connections (approx 10K):

Note *10K UDP Segmentation offload connections currently not supported on T6.*

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Run the following command:

```
[root@host~]# make udp_offload_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

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

Run the following commands to load the driver:

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe t4_tom
```

Though normally associated with the Chelsio TCP Offload engine, the *t4_tom* module is required in order to allow for the proper redirection of UDP socket calls.

4. Software/Driver Configuration and Fine-tuning

4.1. Modifying the application

To use the UDP offload functionality, the application needs to be modified. Follow the steps mentioned below:

i. Determine the UDP socket file descriptor in the application through which data is sent

ii. Declare and initialize two variables in the application:

```
int fs=1316;
int cl=1;
```

Here,

- *fs* is the UDP packet payload size in bytes that is transmitted on the wire. The minimum value of fs is 256 bytes.
- *cl* is the UDP traffic class(scheduler-class-index) that the user wishes to assign the data stream to. This value needs to be in the range of 0 to 14 for T4 series of adapters and a range of 0 to 15 for T5/T6 series of adapters.

The application will function as per the parameters set for that traffic class.

iii. Add socket option definitions:

In order to use *setsockopt()* to set the options to the UDP socket, the following three definitions need to be made:

- SO_FRAMESIZE used for setting frame size, which has the value 291.
- SOL_SCHEDCLASS used for setting UDP traffic class, which has the value 290.
- IPPROTO_UDP used for setting the type of IP Protocol.

```
# define SO_FRAMESIZE 291
# define SOL_SCHEDCLASS 290
# define IPPROTO_UDP 17
```

iv. Use the *setsockopt()* function to set socket options:

```
//Get the UDP socket descriptor variable
setsockopt (sockfd , IPPROTO_UDP, SO_FRAMESIZE, &fs, sizeof(fs));
setsockopt (sockfd , IPPROTO_UDP, SOL_SCHEDCLASS, &cl, sizeof(cl));
```

Here:

- *sockfd* : The file descriptor of the UDP socket
- *&fs* / *&cl* : Pointer to the framesize and class variables
- *sizeof(fs)* / *sizeof(cl)* : The size of the variables

v. Now, compile the application.

4.1.1. UDP offload functionality for RTP data

In case of RTP data, the video server application sends the initial sequence number and the RTP payload. The USO engine segments the payload data, increments the sequence number and sends out the data.

In order to use the UDP offload functionality for RTP data, make the following additions to the steps mentioned above:

1. In step (ii), declare and initialize a new variable in the application:

```
int rtp_header_size=16;
```

Here, *rtp_header_size* is the RTP header size in bytes that the application sends.

2. In step (iii), define a new macro, *UDP_RTPHEADERLEN* used for setting RTP header length with the value 292.

```
# define UDP_RTPHEADERLEN 292
```

3. In step (iv), define a new socket option:

```
setsockopt (sockfd,17,UDP_RTPHEADERLEN,&rtp_header_size,
sizeof(rtp_header_size));
```

Here,

- *&rtp_header_size* : pointer to the RTP header length variable
- *sizeof(rtp_header_size)* : the size of the RTP header length variable

4.2. Configuring UDP Pacing

Now that the application has been modified to associate the application's UDP socket to a particular UDP traffic class, the pacing of that socket's traffic can be set using the *cxgbtool* utility.

i. Bring up the network interface:

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

ii. Run the following command:

```
[root@host~]# cxgbtool <ethX> sched-class params  type packet  level cl-rl
mode flow  rate-unit bits  rate-mode absolute channel <Channel No.> class
<scheduler-class-index> max-rate <maximum-rate>  pkt-size <Packet size>
```

Here,

- *ethX* is the Chelsio interface
- *Channel No.* is the port on which data is flowing (0-3)
- *scheduler-class-index* is the UDP traffic class (0-14 for T4 series of adapters and 0-15 for T5/T6 series of adapters) set in the SOL_SCHEDCLASS socket option in the application in section 4.1.
- *maximum-rate* is the bit rate (Kbps) for this UDP stream. This value should be in the range of 50 Kbps to 50 Mbps for T4 adapters. For T5/T6 adapters, the value should be in the range of 100 kbps to 1 Gbps.
- *Packet size* is the UDP packet payload size in bytes; it should be equal to the value set in the *SO_FRAMESIZE* socket option in the application in section 4.1.

Example:
The user wants to transfer UDP data on port 0 of the adapter using the USO engine. The application has been modified as shown in section 4.1. In order to set a bit rate of 10Mbps for traffic class 1 with payload size of 1316 on port 0, the following invocation of *cxgbtool* is used:

```
[root@host~]# cxgbtool ethX sched-class params  type packet  level cl-rl
mode flow  rate-unit bits  rate-mode absolute channel 0  class 1 max-rate
10000  pkt-size 1316
```

Note

To get an accurate bit rate per class, data sent by the application to the sockets should be a multiple of the value set for the "pkt-size" parameter. In above example, IO size sent by application should be a multiple of 1316.

Note

Linux Unified Wire currently supports 10240 offload UDP connections. If the application needs to establish more than 10240 UDP connections, it can check the return code of ENOSPC from a send() or sendto() call and close this socket and open a new one that uses the kernel UDP stack.

5. Software/Driver Unloading

Reboot the system to unload the driver. To unload without rebooting, refer Unloading the TOE driver section of **Network (NIC/TOE)** chapter.

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

1.2. Software Requirements

1.2.1. Linux Requirements

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

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

2.1. Pre-requisites

Please make sure that the following requirements are met before installation:

- IPv6 must be enabled in your system (enabled by default).
- Unified Wire must be installed with IPv6 support as explained in the Unified Wire chapter.

2.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install Unified Wire with IPv6 support:

```
[root@host~]# make install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

After installing Unified Wire package and rebooting the host, load the NIC (*cxgb4*) and TOE (*t4_tom*) drivers. The drivers must be loaded by the root user. Any attempt to load the drivers as a regular user will fail.

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe t4_tom
```

4. Software/Driver Configuration and Fine-tuning

i. Load the Offload capable drivers.

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

ii. Assign IPv6 address and bring up the interface.

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

iii. All the IPv6 traffic over the Chelsio interface will be offloaded now. To see the number of connections offloaded, run the following command:

```
[root@host~]# cat /sys/kernel/debug/cxgb4/<bus-id>/tids
```

5. Software/Driver Unloading

5.1. Unloading the NIC driver

To unload the NIC driver, run the following command:

```
[root@host~]# rmmod cxgb4
```

5.2. Unloading the TOE driver

Please reboot the system to unload the TOE driver. To unload without rebooting, refer Unloading the TOE driver section of **Network (NIC/TOE)** chapter.

XIX. WD Sniffing and Tracing

1. Theory of Operation

The objective of these utilities (*wd_sniffer* and *wd_tcpdump_trace*) is to provide sniffing and tracing capabilities by making use of T5/T4's hardware features.

- Sniffer is a tool to measure bandwidth and involves targeting specific multicast traffic and sending it directly to user space.
    a) Get a Queue (raw QP) idx.
    b) Program a filter to redirect specific traffic to the raw QP queue.

- Tracer - All tapped traffic is forwarded to user space and also pushed back on the wire via the internal loop back mechanism
    a) Get a Queue (raw QP) idx
    b) Set the T4 adapter in loop back
    c) Connect Client A and B to ports 0 and 1 or ports 2 and 3.
    d) Enable tracing.

In either mode, the targeted traffic bypasses the kernel TCP/IP stack and is delivered directly to user space by means of an RX queue.

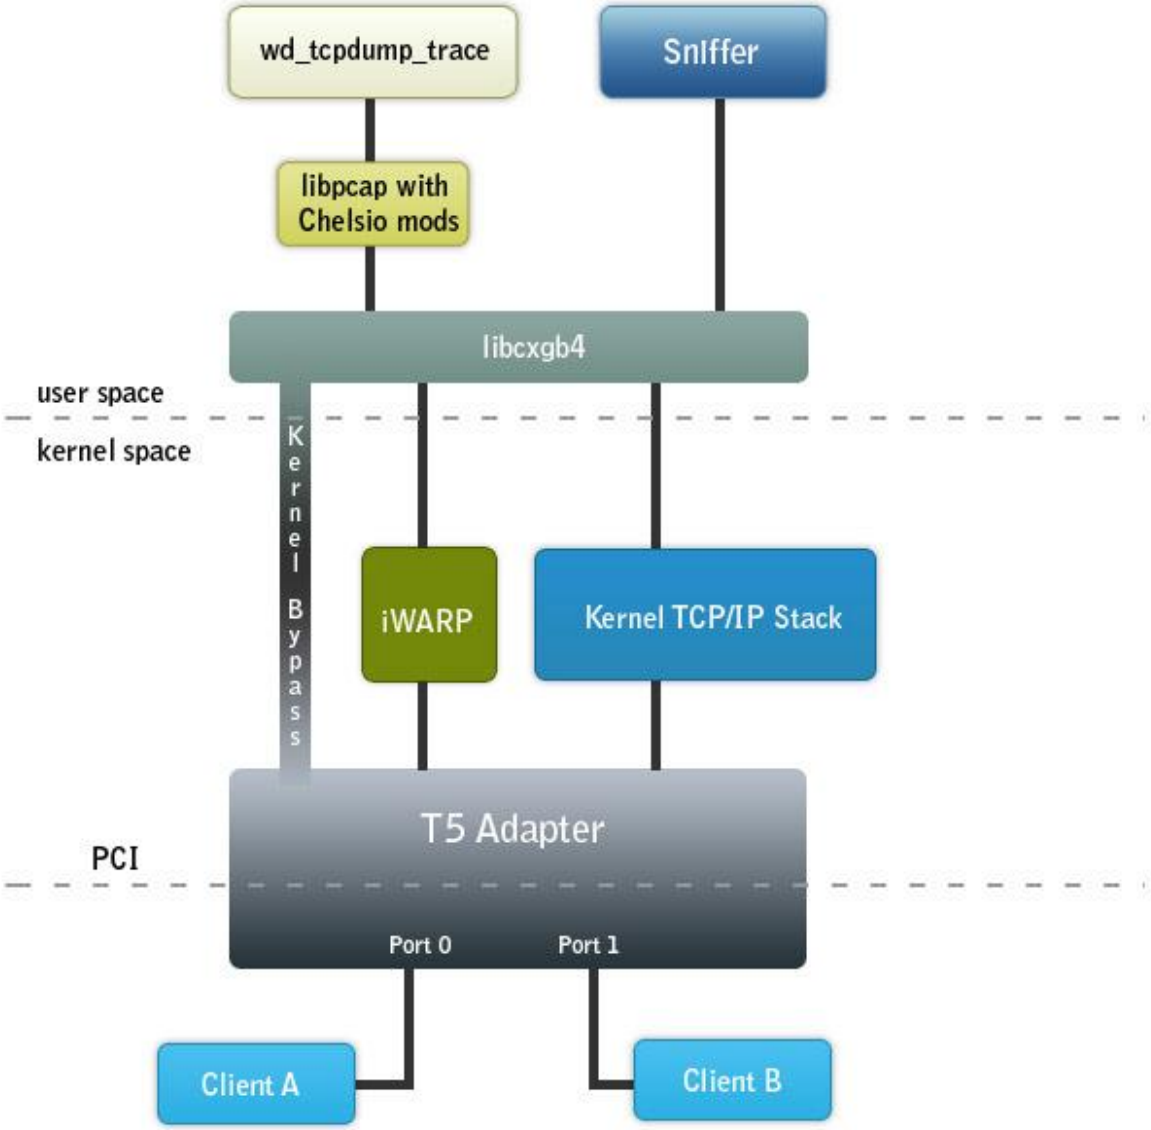

Schematic diagram of sniffer and tracer

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with the tools:

- T520-BT
- T580-CR
- T580-LP-CR
- T520-LL-CR
- T520-CR

- T540-CR
- T420-CR
- T440-CR
- T422-CR
- T404-BT
- T420-BCH
- T440-LP-CR
- T420-BT
- T420-LL-CR
- T420-CX

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the WD Sniffing and Tracing utility is available for the following version:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install Sniffer tracing & filtering tcpdump and iWARP driver:

```
[root@host~]#make sniffer_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Usage

3.1. Installing basic support

iw_cxgb4 (Chelsio iWARP driver) and *cxgb4* (Chelsio NIC driver) drivers should be compiled and loaded before running the utilities. Refer to the **Software/Driver Loading** section for each driver and follow the instructions mentioned before proceeding.

3.2. Using Sniffer (wd_*sniffer)*

1. **Setup**:

Wire filter sniffing requires 2 systems with one machine having a T5/T4 card.

The machines should be setup in the following manner:

Machine A <---------> Machine B

192.168.1.100 192.168.1.200

2. **Procedure**:

On the Device Under Test (DUT), start sniffer.

```
[root@host~]# wd_sniffer -T 20 -s 1000 -I <MAC address of interface to
sniff>
```

Start traffic on the PEER and watch the sniffer.

The sniffer will receive all packets as fast as possible, update the packet count, and then discard the data. Performance is a full 10Gbps for packet size 1000.

3.3. Using Tracer (wd_tcpdump_trace)

1. **Setup**:

Wire tapping requires 3 systems with one machine having a T5/T4 card with two or more ports. The machines should be setup in the following manner:

DUT: Machine B

PEER: Machine A <-----> (port 0) (port 1) <-----> PEER: Machine C

192.168.1.100 IP-dont-care IP-dont-care 192.168.1.200

2. Procedure:

Run `wd_tcpdump_trace -i iface` on the command prompt where *iface* is one of the interfaces whose traffic you want to trace. In the above diagram its port 0 or port 1.

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

Use any tool (like ping or ssh) to run traffic between machines A and B. The traffic should successfully make it from end to end and wd_tcpdump_trace on the DUT should show the tapped traffic.

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

1.2. Software Requirements

1.2.1. Linux Requirements

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

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. LE-TCAM Filters

The default (*Unified Wire)* configuration tuning option allows you to create LE-TCAM filters, which has a limit of 496 for T5, and 560 for T6 adapters. For T5 adapters, all available filter indices can be optionally configured as high priority. In case of T6 adapters, the following filter indices are available:

- High priority indices (PRIO): 0 to X
- Normal indices: X+1 to X+1+495

Where *X* is the upper limit in the *HPFTID range*, mentioned in the tids file (*/sys/kernel/debug/cxgb4/<bus-id>/tids*).

```
[root@    ~]# cat /sys/kernel/debug/cxgb4/0000\:02\:00.4/tids
Connections in use: 0
TID range: 64..2047/3072..19455, in use: 0/0
STID range: 2048..2543, in use-IPv4/IPv6: 0/0
ATID range: 0..8191, in use: 0
FTID range: 2560..3055
HPFTID range: 0..63
UOTID range: 19456..20479, in use: 0
HW TID usage: 0 IP users, 0 IPv6 users
```

For example, if the upper hpftid limit is 63, then high priority indices will be from 0 to 63 and normal indices will be from 64 to 559.

Note *It is mandatory to add **prio 1** option* when creating high priority filter rule.

2.1. Configuration

2.1.1. Filter Modes

The Classification and Filtering feature is configured by specifying the filter modes in the firmware configuration file (*t6-config.txt* for T6 adapters; *t5-config.txt* for T5 adapters.) located in */lib/firmware/cxgb4/* .

Adapter initialization will fail if *filterMask* contains a tuple which is not present in *filterMode*.

The following filter tuples are present in filter modes:

fcoe	: Fibre Channel over Ethernet frames
port	: Packet ingress physical port number
vnic_id	:VF ID in MPS TCAM (*Currently not supported*) and outer VLAN ID, Encapsulation
vlan	: Inner VLAN ID
Tos	: Type of Service
protocol	: IP protocol number (ICMP=1, TCP=6, UDP=17, etc)
Ethertype	: Layer 2 EtherType

Macmatch :MAC index in MPS TCAM
mpshittype :MAC address "match type" (none, unicast, multicast, promiscuous, broadcast)
fragmentation :Fragmented IP packets

2.1.2. Supported Filter Combinations

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

- For T6:

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

- For T5:

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

- For T4:

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

Serial #	Filter Combination
1	*fragmentation, mpshittype, macmatch, ethertype, protocol, port*
2	*fragmentation, mpshittype, macmatch, ethertype, protocol, fcoe*
3	*fragmentation, mpshittype, macmatch, ethertype, tos, port*
4	*fragmentation, mpshittype, macmatch, ethertype, tos, fcoe*
5	*fragmentation, mpshittype, macmatch, ethertype, port, fcoe*
6	*fragmentation, mpshittype, macmatch, protocol, tos, port, fcoe*
7	*fragmentation, mpshittype, macmatch, protocol, vlan, fcoe*
8	*fragmentation, mpshittype, macmatch, protocol, vnic_id, fcoe*
9	*fragmentation, mpshittype, macmatch, tos, vlan, fcoe*
10	*fragmentation, mpshittype, macmatch, tos, vnic_id, fcoe*
11	*fragmentation, mpshittype, macmatch, vlan, port, fcoe*
12	*fragmentation, mpshittype, macmatch, vnic_id, port, fcoe*
13	*fragmentation, mpshittype, ethertype, protocol, tos, port, fcoe*
14	*fragmentation, mpshittype, ethertype, vlan, port*
15	*fragmentation, mpshittype, ethertype, vlan, fcoe*
16	*fragmentation, mpshittype, ethertype, vnic_id, port*
17	*fragmentation, mpshittype, ethertype, vnic_id, fcoe*
18	*fragmentation, mpshittype, protocol, tos, vlan, port*
19	*fragmentation, mpshittype, protocol, tos, vlan, fcoe*

20	*fragmentation, mpshittype, protocol, tos, vnic_id, port*
21	*fragmentation, mpshittype, protocol, tos, vnic_id, fcoe*
22	*fragmentation, mpshittype, protocol, vlan, port, fcoe*
23	*fragmentation, mpshittype, protocol, vnic_id, port, fcoe*
24	*fragmentation, mpshittype, tos, vlan, port, fcoe*
25	*fragmentation, mpshittype, tos, vnic_id, port, fcoe*
26	*fragmentation, mpshittype, vlan, vnic_id, fcoe*
27	*fragmentation, macmatch, ethertype, protocol, port, fcoe*
28	*fragmentation, macmatch, ethertype, tos, port, fcoe*
29	*fragmentation, macmatch, protocol, vlan, port, fcoe*
30	*fragmentation, macmatch, protocol, vnic_id, port, fcoe*
31	*fragmentation, macmatch, tos, vlan, port, fcoe*
32	*fragmentation, macmatch, tos, vnic_id, port, fcoe*
33	*fragmentation, ethertype, vlan, port, fcoe*
34	*fragmentation, ethertype, vnic_id, port, fcoe*
35	*fragmentation, protocol, tos, vlan, port, fcoe*
36	*fragmentation, protocol, tos, vnic_id, port, fcoe*
37	*fragmentation, vlan, vnic_id, port, fcoe*
38	*mpshittype, macmatch, ethertype, protocol, port, fcoe*
39	*mpshittype, macmatch, ethertype, tos, port, fcoe*
40	*mpshittype, macmatch, protocol, vlan, port*
41	*mpshittype, macmatch, protocol, vnic_id, port*
42	*mpshittype, macmatch, tos, vlan, port*
43	*mpshittype, macmatch, tos, vnic_id, port*
44	*mpshittype, ethertype, vlan, port, fcoe*
45	*mpshittype, ethertype, vnic_id, port, fcoe*
46	*mpshittype, protocol, tos, vlan, port, fcoe*
47	*mpshittype, protocol, tos, vnic_id, port, fcoe*
48	*mpshittype, vlan, vnic_id, port*

Important *Using any other filter mode combination is strictly not supported.*

2.1.3. Changing default filter mode

Based on your requirement, you can change the default filter mode to any one of the combinations mentioned in the table above. To do so, replace the default mode with the chosen mode in firmware configuration file (*t6-config.txt* for T6 adapters; *t5-config.txt* for T5 adapters) located in */lib/firmware/cxgb4/*.

For example, if you want to filter traffic based on *ethtype* value in the packets for T6 adapters, replace the default filterMode,

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

with

```
filterMode = fragmentation, mpshittype, macmatch, ethertype, protocol, port
```

The network driver needs to be reloaded next using the following command:

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

2.2. Creating Filter Rules

Network driver (*cxgb4*) must be installed and loaded before setting the filter rule.

i. If you haven't done already, run the Unified Wire Installer with the appropriate configuration tuning option to install the network driver.
ii. Load the network driver and bring up the Chelsio interface:

```
[root@host~]# modprobe cxgb4
[root@host~]# ifconfig ethX up
```

iii. Now, create filter rules using *cxgbtool*:

```
[root@host~]#cxgbtool ethx filter <index> action [pass/drop/switch] <prio 1>
<hitcnts 1>
```

Where,

ethX	: Chelsio interface
index	: positive integer set as filter id. 0-495 for T5 adapters; 0-559 for T6 adapters.
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.
prio 1	: Optional for T5. Mandatory for T6 indices 0-63; Should not be added for T6 indices 64-559
hitcnts 1	: To enable hit counts in cxgbtool filter show output.

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~]# cxgbtool ethX filter 100 action drop fip 192.168.1.5
```

The above filter rule will drop all ingress packets from source IP 192.168.1.5. Remaining packets will be sent to the host.

- ***pass* action**

```
[root@host~]# cxgbtool ethX filter 100 action pass lport 10001 fport 355
queue 2
```

The above filter rule will pass all ingress packets that match destination port 10001 and source port 355 to ingress queue 2 for load balancing. Remaining packets will be sent to the host.

- ***switch* action**

```
[root@host~]# cxgbtool ethX filter 100 action switch iport 0 eport 1 ivlan 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.

- ***prio* option**

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

```
[root@host~]# cxgbtool ethx filter <index> action <pass/drop/switch> prio 1
```

Where index is a positive integer set as filter id. 0-495 for T5 adapters and 0-63 for T6 adapters.

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

2.3. Listing Filter Rules

To list the filters set, run the following command:

```
[root@host~]# cxgbtool ethX filter show
```

OR

```
[root@host~]# cat /sys/kernel/debug/cxgb4/<bus-id>/filters
```

2.4. Removing Filter Rules

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

```
[root@host~]# cxgbtool ethX filter <index> <delete|clear>
```

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

2.5. Layer 3 example

Here's an example on how to achieve L3 routing functionality:

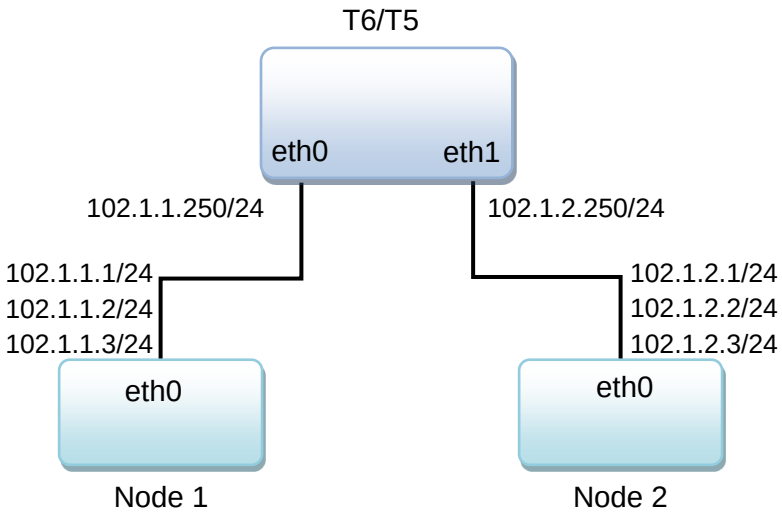

- **Follow these steps on Node 1**

i. Configure IP address and enable the 3 interfaces:

```
[root@host~]# ifconfig eth0 102.1.1.1/24 up
[root@host~]# ifconfig eth0:2 102.1.1.2/24 up
[root@host~]# ifconfig eth0:3 102.1.1.3/24 up
[root@host~]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:07:43:04:7D:50
          inet addr:102.1.1.1  Bcast:102.1.1.255  Mask:255.255.255.0
          inet6 addr: fe80::207:43ff:fe04:7d50/64 Scope:Link
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:14372 errors:0 dropped:0 overruns:0 frame:0
          TX packets:62203 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1585952 (1.5 MiB)  TX bytes:4798122 (4.5 MiB)
          Interrupt:16
```

ii. Setup a static OR default route towards T6/T5 router to reach 102.1.2.0/24 network

```
[root@host~]# route add -net 102.1.2.0/24 gw 102.1.1.250
```

- **Follow these steps on Node 2**

i. Configure IP address and enable the 3 interfaces:

```
[root@host~]# ifconfig eth0 102.1.2.1/24 up
[root@host~]# ifconfig eth0:2 102.1.2.2/24 up
[root@host~]# ifconfig eth0:3 102.1.2.3/24 up
[root@host~]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:07:43:12:D4:88
          inet addr:102.1.2.1  Bcast:102.1.2.255  Mask:255.255.255.0
          inet6 addr: fe80::7:4300:112:d488/64 Scope:Link
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:1961 errors:0 dropped:2 overruns:0 frame:0
          TX packets:141 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:218606 (213.4 KiB)  TX bytes:17483 (17.0 KiB)
          Interrupt:17
```

ii. Setup a static OR default route towards T6/T5 router to reach 102.1.1.0/24 network

```
[root@host~]# route add -net 102.1.1.0/24 gw 102.1.2.250
```

- **Follow these steps on machine with T6/T5 adapter**

i. Configure IP address and enable the 2 interfaces:

```
[root@host~]# ifconfig eth0 102.1.1.250/24 up
[root@host~]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:07:43:04:96:40
          inet addr:102.1.1.250  Bcast:102.1.1.255  Mask:255.255.255.0
          inet6 addr: fe80::207:43ff:fe04:9640/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:114 errors:0 dropped:0 overruns:0 frame:0
          TX packets:535 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:11880 (11.6 KiB)  TX bytes:61729 (60.2 KiB)
          Interrupt:16
[root@host~]# ifconfig eth1 102.1.2.250/24 up
[root@host~]# ifconfig eth1
eth1      Link encap:Ethernet  HWaddr 00:07:43:04:96:48
          inet addr:102.1.2.250  Bcast:102.1.2.255  Mask:255.255.255.0
          inet6 addr: fe80::7:4300:104:9648/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:31 errors:0 dropped:0 overruns:0 frame:0
          TX packets:433 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:3181 (3.1 KiB)  TX bytes:49134 (47.9 KiB)
          Interrupt:16
```

ii. Create filter rule to send packets for 102.1.2.0/24 network out via eth1 interface:

```
[root@host~]# cxgbtool eth0 filter 100 lip 102.1.2.0/24 hitcnts 1 action
switch  eport 1 smac 00:07:43:04:96:48 dmac 00:07:43:12:D4:88
```

Where, `smac` is the MAC address of eth1 interface on T6/T5 adapter machine and `dmac` is the MAC address of eth0 interface on Node 2.

iii. Create filter rule to send packets for 102.1.1.0/24 network out via eth0 interface

```
[root@host~]# cxgbtool eth0 filter 101 lip 102.1.1.0/24 hitcnts 1 action
switch eport 0 smac 00:07:43:04:96:40 dmac 00:07:43:04:7D:50
```

Where, `smac` is the MAC address of eth0 interface on T6/T5 adapter machine and `dmac` is the MAC address of eth0 interface on Node 1.

2.6. Layer 2 example

Here's an example on how to achieve L2 switching functionality. The following will only work on kernel 3.10 and above.

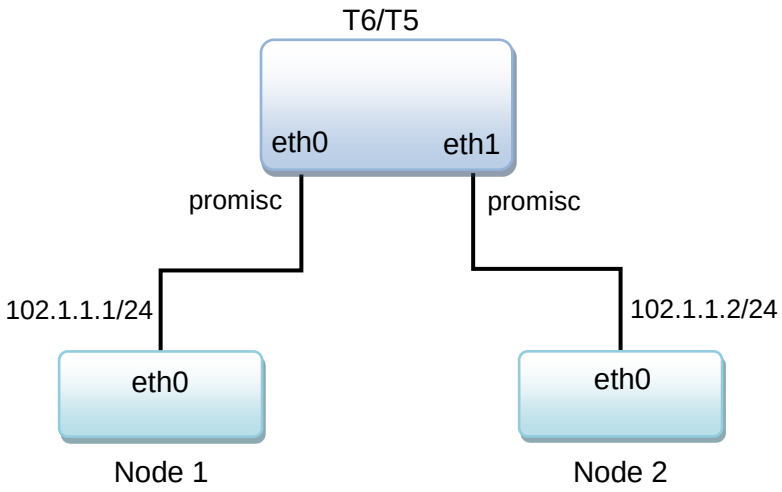

- **Follow these steps on Node 1**

i. Configure IP address and enable the interface:

```
[root@host~]# ifconfig eth0 102.1.1.1/24 up
[root@host~]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:07:43:04:7D:50
          inet addr:102.1.1.1  Bcast:102.1.1.255  Mask:255.255.255.0
          inet6 addr: fe80::207:43ff:fe04:7d50/64 Scope:Link
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:14372 errors:0 dropped:0 overruns:0 frame:0
          TX packets:62203 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1585952 (1.5 MiB)  TX bytes:4798122 (4.5 MiB)
          Interrupt:16
```

ii. Setup ARP entry to reach 102.1.1.2

```
[root@host~]# arp -s 102.1.1.2 00:07:43:12:D4:88
```

- **Follow these steps on Node 2**

i. Configure IP address and enable the interface:

```
[root@host~]# ifconfig eth0 102.1.1.2/24 up
[root@host~]# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:07:43:12:D4:88
          inet addr:102.1.1.2  Bcast:102.1.1.255  Mask:255.255.255.0
          inet6 addr: fe80::7:4300:112:d488/64 Scope:Link
          UP BROADCAST RUNNING PROMISC MULTICAST  MTU:1500  Metric:1
          RX packets:1961 errors:0 dropped:2 overruns:0 frame:0
          TX packets:141 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:218606 (213.4 KiB)  TX bytes:17483 (17.0 KiB)
          Interrupt:17
```

ii. Setup ARP entry to reach 102.1.1.1

```
[root@host~]# arp -s 102.1.1.1 00:07:43:04:7D:50
```

- **Follow these steps on machine with T6/T5 adapter**

i. Update filterMode value with below combination in */lib/firmware/cxgb4/t6-config.txt* to enable matching based on macidx (use *t5-config.txt* for T5 adapters).

```
filterMode = fragmentation, macmatch, mpshittype, protocol, tos, port, fcoe
```

ii. Unload and re-load the *cxgb4* driver.

iii. Enable promiscuous mode on both the interfaces on T6/T5 adapter machine:

```
[root@host~]# ifconfig eth0 up promisc
[root@host~]# ifconfig eth1 up promisc
```

iv. Build and install latest iproute2 package

v. Add fdb entry corresponding to Node-2 on T6/T5's eth0 interface:

```
[root@host~]# bridge fdb add 00:07:43:12:D4:88 dev eth0 self
```

vi. Add fdb entry corresponding to Node-1 on T6/T5's eth1 interface:

```
[root@host~]# bridge fdb add 00:07:43:04:7D:50 dev eth1 self
```

vii. Both MAC entries should show up in MPS table. Run the following command to view the table and note the index (*idx* field) of the entries:

```
[root@host~]# cat /sys/kernel/debug/cxgb4/0000\:01\:00.4/mps_tcam | more
```

viii. Create a filter to match incoming packet's *dst-mac 00:07:43:12:d4:88* with particular mac-idx and switch it out via eport 1:

```
[root@host~]# cxgbtool eth0 filter 100 macidx 5 action switch eport 1
hitcnts 1
```

ix. Create a filter to match incoming packet's *dst-mac 00:07:43:04:7d:50* with particular mac-idx and switch it out via eport 0:

```
[root@host~]# cxgbtool eth0 filter 101 macidx 7 action switch eport 0
hitcnts 1
```

3. Hash/DDR Filters

If you wish to create more filters, select *T5/T6 Hash Filter* configuration tuning option during installation which allows you to create ~0.5 million filter rules. You can create both LE-TCAM and Hash/DDR filters in this configuration.

Hash filters are created based on *filterMask* tuples in firmware configuration file (*t6-config.txt* for T6 adapters; *t5-config.txt* for t5 adapters) located in */lib/firmware/cxgb4/*. *filterMask* tuples should be either subset of or equal to *filterMode* tuples.

Hash filters are exact match filters. Hence, when you enable more fields(tuples) in *filterMask*, you must create a filter rule with exactly same tuples as mentioned in *filterMask*.

3.1. Configuration

3.1.1. Filter Modes

The Classification and Filtering feature is configured by specifying the filter modes in the firmware configuration file located in */lib/firmware/cxgb4/*

Adapter initialization will fail if *filterMask* contains a tuple which is not present in *filterMode*.

The following filter tuples are present in filter modes:

fcoe	: Fibre Channel over Ethernet frames
port	: Packet ingress physical port number
vnic_id	: VF ID in MPS TCAM (*Currently not supported*) and outer VLAN ID
vlan	: Inner VLAN ID
tos	: Type of Service
protocol	: IP protocol number (ICMP=1, TCP=6, UDP=17, etc)
ethertype	: Layer 2 EtherType
macmatch	: MAC index in MPS TCAM
mpshittype	: MAC address "match type" (none,unicast,multicast,promiscuous,broadcast)
fragmentation	: Fragmented IP packets

3.1.2. Supported Filter Combinations

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

- For T6:

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

- For T5:

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

Serial #	Filter Combination
1	*fragmentation, mpshittype, macmatch, ethertype, protocol, port*
2	*fragmentation, mpshittype, macmatch, ethertype, protocol, fcoe*
3	*fragmentation, mpshittype, macmatch, protocol, tos, port, fcoe*
4	*fragmentation, mpshittype, macmatch, protocol, vlan, fcoe*
5	*fragmentation, mpshittype, macmatch, protocol, vnic_id, fcoe*
6	*fragmentation, mpshittype, ethertype, protocol, tos, port, fcoe*
7	*fragmentation, mpshittype, protocol, tos, vlan, port*
8	*fragmentation, mpshittype, protocol, tos, vlan, fcoe*
9	*fragmentation, mpshittype, protocol, tos, vnic_id, port*
10	*fragmentation, mpshittype, protocol, tos, vnic_id, fcoe*
11	*fragmentation, mpshittype, protocol, vlan, port, fcoe*
12	*fragmentation, mpshittype, protocol, vnic_id, port, fcoe*
13	*fragmentation, macmatch, ethertype, protocol, port, fcoe*
14	*fragmentation, macmatch, protocol, vlan, port, fcoe*
15	*fragmentation, macmatch, protocol, vnic_id, port, fcoe*
16	*fragmentation, protocol, tos, vlan, port, fcoe*
17	*fragmentation, protocol, tos, vnic_id, port, fcoe*
18	*mpshittype, macmatch, ethertype, protocol, port, fcoe*
19	*mpshittype, macmatch, protocol, vlan, port*
20	*mpshittype, macmatch, protocol, vnic_id, port*
21	*mpshittype, protocol, tos, vlan, port, fcoe*
22	*mpshittype, protocol, tos, vnic_id, port, fcoe*

Important *Using any other filter mode combination is strictly not supported.*

3.1.3. Changing default filter mode

Based on your requirement, you can change the default filter mode to any one of the combinations mentioned in the table above. To do so, replace the default mode with the chosen mode in firmware configuration file (*t6-config.txt* for T6 adapters; *t5-config.txt* for T5 adapters) located in */lib/firmware/cxgb4/*.

For example, if you want to filter traffic based on *ethtype* value in the packets for T6 adapters, replace the default filterMode,

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

with

```
filterMode = fragmentation, mpshittype, macmatch, ethertype, protocol, port
```

You can change the default filter mode to any one of the following combinations, based on your requirement. The network driver needs to be reloaded next using the following command:

```
[root@host~]# rmmod cxgb4
[root@host~]# modprobe cxgb4 use_ddr_filters=1
```

3.2. Creating Filter Rules

Network driver (*cxgb4*) must be installed and loaded before setting the filter rule.

i. If you haven't done already, run the Unified Wire Installer with the *T5/T6 Hash Filter* configuration tuning option to install the Network Driver.
ii. Load the network driver with DDR filters support and bring up the Chelsio interface :

```
[root@host~]# modprobe cxgb4 use_ddr_filters=1
[root@host~]# ifconfig ethX up
```

iii. Now, create filter rules using cxgbtool:

```
[root@host~]# cxgbtool ethX filter <index> action [pass/drop/switch] fip
<source_ip> lip <destination_ip> fport <source_port> lport
<destination_port> proto <protocol> <hitcnts 1> <cap maskless>
```

Where,

ethX	:	Chelsio interface.
index	:	Filter index. For LE-TCAM filters, filter index should be 0-495 for T5 adapters and 0-559 for T6 adapters. In case of Hash/DDR filter, the index will be ignored and replaced by an automatically computed value, based on the hash (4-tuple). The index will be displayed after the filter rule is created successfully.
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.
source_ip/port	:	Source IP/port of incoming packet.
destination_ip/port	:	Destination IP/port of incoming packet.
protocol	:	TCP by default. To change, specify the corresponding internet protocol number, e.g., use 17 for UDP.
hitcnts 1	:	To enable hit counts in cxgbtool filter show output.
cap maskless	:	This is mandatory for hash filter. If not provided, LE-TCAM filter will be created at the specified index.

Note

In case of Hash/DDR filters, ***source_ip****,* ***destination_ip****,* ***source_port*** *and* ***destination_port*** *are mandatory, since the filters don't support masks and hence, 4-tuple must always be supplied.* ***Proto*** *is also a mandatory parameter.*

3.2.1. Choosing filterMode and filterMask

As mentioned earlier filterMask tuples can be subset of or equal to filterMode tuples. Following are examples of how you can select filterMode and filterMask and create a Hash filter based on those values:

- **when all tuples from filterMode are enabled in filterMask**

i. Select a filterMode from supported filterMode table based on your requirement. E.g.,

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

ii. Select a filterMask so that it is a subset of or equal to filterMode based on application. E.g.;

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

Here, we have selected *fragmentation, mpshittype, protocol, vlan, port, fcoe* in filterMask so it is mandatory to create a filter based on all of those tuples mentioned in filterMask. Otherwise, filter rule will not honour.

iii. Now, to create a hash filter based on the filterMode and filterMask values selected above:

```
[root@host~]# cxgbtool eth18 filter 100 action drop lip 120.10.10.100 fip
120.10.10.200 lport 5001 fport 51549 proto 6 frag 0 matchtype 0 ivlan 10
iport 0 fcoe 0 hitcnts 1 cap maskless
Hash-Filter Index = 303760
```

- **When all tuples from filterMode are not enabled in filterMask:**

In case if you don't want to create filter rule based on any particular tuple from filterMode, don't select it in filterMask. For example, if you don't want to create a filter based on VLAN value, then don't select it from filterMode to filterMask.

i. Select a filterMode from supported filterMode table based on your requirement. E.g.,

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

ii. Select a filterMask so that it is a subset of or equal to filterMode based on application without VLAN tuple. E.g.;

```
filterMask = fragmentation, mpshittype, protocol, port, fcoe
```

Here, we have selected *fragmentation, mpshittype, protocol, port, fcoe* in filterMask so it is mandatory to create a filter based on only those tuples mentioned in filterMask. Otherwise, filter rule will not honour.

iii. Now, to create a hash filter based on the filterMode and filterMask values selected above:

```
[root@indus sw]# cxgbtool eth18 filter 100 action drop lip 120.10.10.100 fip
120.10.10.200 lport 5001 fport 51549 proto 6 frag 0 matchtype 0  iport 0
fcoe 0 hitcnts 1 cap maskless
Hash-Filter Index = 196568
```

3.2.2. Examples

- ***drop* action**

```
[root@host~]# cxgbtool ethX filter 496 action drop lip 102.1.1.1 fip
102.1.1.2 lport 12865 fport 20000 hitcnts 1 cap maskless iport 1 proto 17
Hash-Filter Index = 61722
```

The above filter rule will drop all UDP packets matching above 4 tuple coming on Chelsio port 1. Remaining packets will be sent to the host.

- ***pass* action**

```
[root@host~]# cxgbtool ethX filter 496 action pass lip 102.2.2.1 fip
102.2.2.2 lport 12865 fport 12000 hitcnts 1 cap maskless proto 6
Hash-Filter Index = 308184
```

The above filter rule will pass all TCP packets matching above 4 tuple. Remaining packets will be sent to the host.

- ***switch* action**

```
[root@host~]# cxgbtool ethX filter 496 action switch lip 102.3.3.1 fip
102.3.3.2 lport 5001 fport 16000 iport 0 eport 1 hitcnts 1 cap maskless
Hash-Filter Index = 489090
```

The above filter rule will switch all the packets matching above 4 tuple from Chelsio port 0 to Chelsio port 1. Remaining packets will be sent to the host.

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

3.3. Listing Filter Rules

- To list the Hash/DDR filters set, run the following command:

```
[root@host~]# cat /sys/kernel/debug/cxgb4/<bus-id>/hash_filters
```

- To list the both LE-TCAM and Hash/DDR filters set, run the following command:

```
[root@host~]# cxgbtool ethX filter show
```

3.4. Removing Filter Rules

To remove a filter, run the following command with *cap maskless* parameter and corresponding filter rule index:

```
[root@host~]# cxgbtool ethX filter <index> <delete|clear> cap maskless
```

Note

- *Filter rule index can be determined by referring the "hash_filters" file located in /sys/kernel/debug/cxgb4/<bus-id>/.*
- *For more information on additional parameters, refer cxgbtool manual by running the* `man cxgbtool` *command.*

3.5. Filter Priority

By default, Hash/DDR filter has priority over LE-TCAM filter. To override this, the LE-TCAM filter should be created with *prio* option. For example:

```
[root@host~]# cxgbtool ethx filter <index> action <pass/drop/switch> prio 1
```

Where *index* is a positive integer set as filter id. 0-495 for T5 adapters and 0-63 for T6 adapters.

3.6. Swap MAC Feature

Chelsio's T6/T5 Swap MAC feature swaps packet source MAC and destination MAC addresses. This is applicable only for switch filter rules. Here's an example:

```
[root@host~]# cxgbtool eth2 filter 100 action switch lip 102.2.2.1 fip
102.2.2.2 lport 5001 fport 14000 hitcnts 1 iport 1 eport 0 swapmac 1 proto
17 cap maskless
Hash-Filter Index = 21936
```

The above example will swap source and destination MAC addresses of UDP packets (matching above 4 tuple) received on adapter port 1 and then switch them to port 0.

3.7. Traffic Mirroring

On T5 adapters, when using *Hash Filter* configuration tuning option, Network driver (cxgb4) parameter *enable_mirror* can be used to enable mirroring of traffic running on physical ports. The

mirrored traffic will be received via Mirror PF/VF on Mirror Receive queues, which will then inject this traffic into network stack of Linux kernel.

3.7.1. Enabling Mirroring

To enable traffic mirroring, follow the steps mentioned below:

i. If not done already, install Unified Wire with *Hash Filter* configuration tuning option as mentioned in the Unified Wire chapter.
ii. Enable *vnic_id* match for filterMode in Hash filter config file, *t5-config.txt*, located in */lib/firmware/cxgb4/*

```
filterMode = fragmentation, mpshittype, protocol, vnic_id, port, fcoe
filterMask = port, protocol, vnic_id
```

iii. Unload network driver (*cxgb4*) and reload it with mirroring enabled.

```
[root@host~]# rmmod cxgb4
[root@host~]# modprobe cxgb4 enable_mirror=1 use_ddr_filters=1
```

iv. The traffic will now be mirrored and received via mirror PF/VF corresponding to each port.

3.7.2. Switch Filter with Mirroring

The following example explains the method to switch and mirror traffic simultaneously:

i. Obtain the PF and VF values of the incoming port from */sys/kernel/debug/cxgb4/<bus-id>/mps_tcam*
ii. Create the desired switch filter rule:

```
[root@host~]# cxgbtool ethX filter 100 fip 102.8.8.2 lip 102.8.8.1 fport
20000 lport 12865 pf 4 vf 64 action switch iport 0 eport 1 cap maskless
```

The above hash filter rule is created to switch traffic from *102.8.8.2,20000* to *102.8.8.1,12865* received on port 0 to be switched out from port 1. The traffic will be switched and simultaneously received on mirror queues and network stack of host as mirroring is enabled.

3.7.3. Filtered Traffic Mirroring

Once mirroring is enabled, all the traffic received on a physical port will be duplicated. The following example explains the method to filter out the redundant traffic and receive only specific traffic on mirror queues:

i. Obtain the mirror PF and VF values from *dmesg*. You should see a similar output:

```
[165299.356887] cxgb4 0000:02:00.4: Port 0 Traffic Mirror PF = 4; VF = 66
[165299.358004] cxgb4 0000:02:00.4: Port 1 Traffic Mirror PF = 4; VF = 67
```

ii. Create a DROP-ALL rule as below:

```
[root@host~]# cxgbtool ethX filter 255 pf 4 vf 66 action drop
```

Where, 255 is the last index of available TCAM filters. This will create a catch-all DROP filter for Mirror PF/VF of port 0. Similarly, create DROP filters for rest of Mirror PF/VF.

iii. Create specific filter rules to allow specific traffic to be received on mirror queues as below:

```
[root@host~]# cxgbtool ethX filter 101 lip 102.8.8.1 fip 102.8.8.2 lport
12865 fport 20000 pf 4 vf 66 action pass
```

Now, the above specific traffic (from *102.8.8.2,20000* to *102.8.8.1,12865*) will be received in mirror receive queues and network stack of host.

3.8. Packet Tracing and Hit Counters

For T5/T6 LE-TCAM and T6 Hash/DDR filters, *hit counters* will work simply by adding *hitcnts 1* parameter to the filter rule. However, for T5 Hash/DDR filters, you will have to make use of tracing feature and RSS queues. Here's a step-by-step guide to enable packet tracing and *hit counters* for T5 Hash/DDR filter rules:

i. Load network driver with the following parameters:

```
[root@host~]# modprobe cxgb4 use_ddr_filters=1 enable_traceq=1
```

ii. Configure the required filter rules.
iii. Enable tracing on T5 adapter.

```
[root@host~]# cxgbtool ethX reg 0x09800=0x13
```

iv. Setup a trace filter

```
[root@host~]# echo tx1 snaplen=40 > /sys/kernel/debug/cxgb4/<bus_id>/trace0
```

Here, *snaplen* is the length in bytes to be captured.

Note *Use "snaplen=60" in case of IPV6.*

The above step will trace all the packets transmitting from port1(tx1) to trace filter 0.

v. Configure RSS Queue to send trace packets. Determine the RspQ ID of the queues by looking at *Trace* QType in */sys/kernel/debug/cxgb4/<bus-id>/sge_qinfo* file

```
[root@host~]# cxgbtool ethX reg 0x0a00c=<Trace Queue0-RspQ ID>
```

Now the traced packets can be seen in tcpdump and the hit counters will also increment.

- **Multi-tracing**

To enable packet capture or *hit counters* for multiple chelsio ports in Tx/Rx direction enable Multi-tracing. Using this we can configure 4 different RSS Queues separately corresponding to 4 trace-filters.

i. Enable Tracing as well as MultiRSSFilter

```
[root@host~]# cxgbtool ethX reg 0x09800=0x33
```

ii. Setup a trace filter

```
[root@host~]# echo tx0 snaplen=40 > /sys/kernel/debug/cxgb4/<bus_id>/trace0
```

iii. Configure the RSS Queue corresponding to trace0 filter configured above. Determine the *RspQ ID* of the queues by looking at *Trace* QType in */sys/kernel/debug/cxgb4/<bus-id>/sge_qinfo* file.

```
[root@host~]# cxgbtool ethX reg 0x09808=<Trace-Queue0-RspQ ID>
```

iv. Similarly for other direction and for multiple ports run the follow commands:

```
[root@host~]# echo rx0 snaplen=40 > /sys/kernel/debug/cxgb4/<bus_id>/trace1
[root@host~]# echo tx1 snaplen=40 > /sys/kernel/debug/cxgb4/<bus_id>/trace2
[root@host~]# echo rx1 snaplen=40 > /sys/kernel/debug/cxgb4/<bus_id>/trace3
[root@host~]# cxgbtool ethX reg 0x09ff4=<Trace-Queue1-RspQ ID>
[root@host~]# cxgbtool ethX reg 0x09ffc=<Trace-Queue2-RspQ ID>
[root@host~]# cxgbtool ethX reg 0x0a004=<Trace-Queue3-RspQ ID>
```

Note *Use “snaplen=60” in case of IPV6.*

4. NAT Filtering

T6/T5 adapters can change source & destination IP addresses and source & destination port numbers achieving NAT functionality. This feature is supported with both LE-TCAM and Hash filters.

Note *This feature is only supported with filter action switch.*

Syntax:

```
[root@host~]# cxgbtool ethX filter <index> action switch fip <source_ip> lip
<destination_ip> fport <source_port> lport <destination_port> nat <mode>
nat_fip <new_source_ip> nat_lip <new_destination_ip> nat_fport
<new_source_port> nat_lport <new_destination_port>
```

Where,

ethX	: Chelsio interface.
source_ip/port	: Source IP/port of incoming packet.
destination_ip/port	: Destination IP/port of incoming packet.
new_source_ip/port	: Source IP/port to be translated to.
new_destination_ip/port	: Destination IP/port to be translated to.
mode	: Combination of IP/port to be translated. *all* will translate all 4-tuple fields. To see other modes, refer cxgbtool manual page.

Examples:

- Hash filter to translate all four tuples, viz. source IP, destination IP, source port and destination port to new values:

```
[root@ ~]# cxgbtool eth0 filter 101 action switch iport 0 eport 1 fip 102.10.10.100 lip 102.10.10.200 fpo
rt 50000 lport 60000 proto 17 nat all nat_fip 192.168.10.100 nat_lip 192.168.10.200 nat_fport 60000 nat_lport
61000 hitcnts 1 cap maskless
Hash-Filter Index = 232776
```

- Hash filter to translate source IP and source port to new values:

```
[root@ ~]# cxgbtool eth0 filter 101 action switch iport 0 eport 1 fip 102.10.10.100 lip 102.10.10.200 fpo
rt 50000 lport 60000 proto 17 nat sip-sp nat_fip 192.168.10.100  nat_fport 61000 hitcnts 1 cap maskless
Hash-Filter Index = 232776
```

- LE-TCAM filter to translate destination IP and destination port to new values:

```
[root@ ~]# cxgbtool eth0 filter 101 action switch iport 0 eport 1 lip 102.10.10.200 lport 60000 nat dip-d
p nat_lip 192.168.10.200 nat_lport 61000 hitcnts 1
```

XXI. OVS Kernel Datapath Offload

1. Introduction

Open vSwitch is a production quality, multilayer virtual switch licensed under the open source Apache 2.0 license. It is designed to enable massive network automation through programmatic extension, while still supporting standard management interfaces and protocols.

Chelsio's T6/T5 Unified Wire solution can offload OVS datapath flow match entries and action processing onto Chelsio adapter for hardware acceleration of OVS datapath flow processing.

Chelsio 1/10/25/40/50/100Gb Ethernet controllers and adapters are capable of offloading OpenFlow and non-OpenFlow network traffic simultaneously, including tunnel handling (e.g. VXLAN / IPsec), NAT, IP stack (ARP, route lookup, frag tracking, fragment / defragment) and other kernel functionalities. A high performance, scalable network I/O is delivered by leveraging built in eSwitch and traffic manager capabilities. In addition, features like traffic classifier, load balancer and firewall are supported at port level by all Chelsio adapters.

1.1. Hardware Requirements

The following are the currently shipping Chelsio adapters that are compatible with OVS driver:

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

1.2. Software Requirements

The Chelsio OVS driver has been developed to run on 64-bit Linux platforms. Following is the list of supported distributions:

- RHEL 7.3, 3.10.0-514.el7 kernel
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)

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

2. Software/Driver Installation

2.1. Pre-requisites

GCC 4.6+, Python 2.7+, Python-six, Autoconf 2.63+, Automake 1.10+, libtool 2.4+ packages should be installed. For the complete list of software required visit *http://docs.openvswitch.org/en/latest/intro/install/general/*

2.2. Installation

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install OVS driver:

```
[root@host~]# make ovs_install
```

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Software/Driver Configuration and Fine Tuning

Supported Fields

The following match fields are supported for offload:

- Input Port
- L2 Ethernet Type
- L3 IP Protocol Type
- L3 IPv4 address
- L3 IPv6 address
- L3 IPv4 TOS
- L3 IP Fragmentation
- L4 Ports (tcp/udp src-port, dst-port)
- Tunnel/Encapsulation VNI (only on T6)

Supported Actions

The following actions are supported for offload:

- Drop
- Switch (output to a port)
- L2 Rewrite: src-mac, dst-mac
- VLAN Rewrite: push, pop, modify
- L3 Rewrite: ip-src, ip-dst (IPv4 and IPv6)
- L4 Rewrite: src-port, dst-port (TCP/UDP)

3.1. Configuring OVS Machine

The following example explains the method to configure an OVS machine:

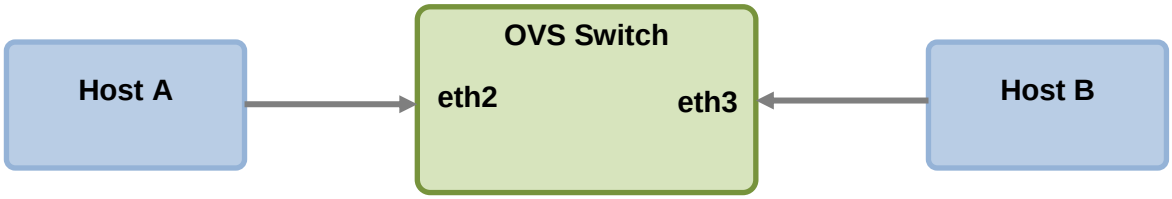

**eth2* and *eth3* are Chelsio interfaces.

i. Please install Unified Wire with ***Hash Filter (T5/T6)*** configuration tuning option.
ii. Update the *filterMode* and *filterMask* in the the hash config file in */lib/firmware/cxgb4/*. Use *t6-config.txt* for T6 adapters and *t5-config.txt* for T5 adapters:

```
filterMode = fragmentation,ethertype,protocol,tos,port
filterMask = fragmentation,ethertype,protocol,tos,port
```

Note *filterMask tuples can be subset of or equal to filterMode tuples.*

iii. Load NIC (cxgb4) driver with hash-filter support:

```
[root@host~]# modprobe cxgb4 use_ddr_filters=1
```

iv. Bring up the Chelsio interfaces in promiscous mode:

```
[root@host~]# ifconfig eth2 promisc up
[root@host~]# ifconfig eth3 promisc up
```

v. Load Open vSwitch module:

```
[root@host~]# modprobe openvswitch
```

vi. Configure OVS

```
[root@host~]# ovs-appctl exit
[root@host~]# pkill -9 ovs
[root@host~]# rm -rf /usr/local/etc/ovs-vswitchd.conf
[root@host~]# rm -rf /usr/local/var/run/openvswitch/db.sock
[root@host~]# rm -rf /usr/local/etc/openvswitch/conf.db
[root@host~]# touch /usr/local/etc/ovs-vswitchd.conf
[root@host~]# ovsdb-tool create /usr/local/etc/openvswitch/conf.db
<uwire_package>/src/openvswitch-x.x.x/vswitchd/vswitch.ovsschema
[root@host~]# ovsdb-server /usr/local/etc/openvswitch/conf.db --
remote=punix:/usr/local/var/run/openvswitch/db.sock --
remote=db:Open_vSwitch,Open_vSwitch,manager_options --bootstrap-ca-
cert=db:Open_vSwitch,SSL,ca_cert --pidfile --detach --log-file
[root@host~]# ovs-vsctl --no-wait init
[root@host~]# export DB_SOCK=/usr/local/var/run/openvswitch/db.sock
[root@host~]# ovs-vswitchd --pidfile --detach
```

vii. Create an OVS bridge and add Chelsio interfaces to it:

```
[root@host~]# ovs-vsctl add-br br0
[root@host~]# sleep 2
[root@host~]# ifconfig br0 up
[root@host~]# ovs-vsctl add-port br0 eth2
[root@host~]# sleep 5
[root@host~]# ovs-vsctl add-port br0 eth3
[root@host~]# sleep 5
[root@host~]# ovs-vsctl show
```

Note *Ports on OVS bridge must be added in the same order as the adapter, since there's no mapping between OVS and physical ports.*

viii. Now ping from Host A to Host B to verify that OVS is configured successfully.

ix. Stop the ping traffic and delete all the flows on switch:

```
[root@host~]# ovs-ofctl del-flows br0
```

3.2. Creating OVS flows

It is mandatory to specify L2 Ethernet Type (dl_type) to offload OVS flows. There are two types of flows:

- **exact-match**: Protocol and 4-tuple are mandatory to create an exact-match flow. ~0.5 million exact-match flows can be offloaded.
- **wild-card**: If any of 4-tuple and protocol are absent, wild-card flow is created. 496 wild-card flows can be offloaded.

3.2.1. Examples

3.2.1.1. Generic flows

Below are few example OVS Flows with the following *filterMode* and *filterMask* combination:

```
filterMode = fragmentation,ethertype,protocol,tos,port
filterMask = fragmentation,ethertype,protocol,tos,port
```

- Wild-card flow to drop incoming packets on first port:

```
[root@host~]# ovs-ofctl add-flow br0 in_port=1,dl_type=0x800,action=drop
```

- Wild-card flow to switch ARP packets (L2 EtherType=0x0806) on 1st port to 2nd port:

```
[root@host~]# ovs-ofctl add-flow br0
in_port=1,dl_type=0x0806,action=output:2
```

- Wild-card flow to switch TCP packets (L3 proto=6) on 1st port to 2nd port by rewriting source and destination MAC addresses:

```
[root@host~]# ovs-ofctl add-flow br0 in_port=1,dl_type=0x800,
nw_proto=6,action=mod_dl_src:00:07:43:28:E4:50,
mod_dl_dst:00:07:43:44:64:50,output:2
```

- Exact-match flow to drop matching 4-tuple traffic:

```
[root@host~]# ovs-ofctl add-flow br0
in_port=1,dl_type=0x800,nw_proto=6,nw_src=10.1.1.66,tp_src=11000,nw_dst=10.1
.1.58,tp_dst=11000,action=drop
```

- Exact-match flow to match 4-tuple IPv4 traffic and do NAT rewrite:

```
[root@host~]# ovs-ofctl add-flow br0
dl_type=0x800,nw_proto=6,nw_src=10.1.1.66,tp_src=11000,nw_dst=10.1.1.58,tp_d
st=21000,action=mod_nw_src=10.2.2.66,mod_tp_src=11005,mod_nw_dst:10.2.2.62,m
od_tp_dst:12345,output:2
```

- Exact-match flow to match 4-tuple IPv6 traffic and do NAT rewrite:

```
[root@host~]# ovs-ofctl add-flow br0
in_port=1,dl_type=0x86dd,nw_proto=6,ipv6_src=2000::66,tp_src=11000,ipv6_dst=
2000::58,tp_dst=11000,action=set_field:2001::66-
\>ipv6_src,mod_tp_src=15000,output:2
```

- Wild-card flow to drop fragmented packets:

```
[root@host~]# ovs-ofctl add-flow br0 dl_type=0x800,ip_frag=yes,action=drop
```

- Exact-match flow to switch 4-tuple IPv4 traffic with TOS 0xE0:

```
[root@host~]# ovs-ofctl add-flow br0 in_port=1,dl_type=0x800,nw_proto=6,
nw_src=10.1.1.66,tp_src=11000,nw_dst=10.1.1.58,tp_dst=21000,
nw_tos=0xE0,action=output:2
```

- If a wild-card and exact match flow both exist for the same traffic pattern, the flow that is created first will take priority. In the below example, the wild-card flow will take priority as it was created first.

```
[root@host~]# ovs-ofctl add-flow br0
dl_type=0x800,nw_src=10.1.1.58,nw_dst=10.1.1.66,tp_src=15000,tp_dst=15000,ac
tion=output:1
[root@host~]# ovs-ofctl add-flow br0
dl_type=0x800,nw_proto=6,nw_src=10.1.1.58,nw_dst=10.1.1.66,tp_src=15000,tp_d
st=15000,action=output:1
```

3.2.1.2. VLAN flows

Below are few example VLAN flows with the following *FilterMode* and *FiletrMask* combination.

```
filterMode = fragmentation,mpshittype,ethertype,vlan,port
filterMask = ethertype,vlan,port
```

Reload cxgb4 driver after updating *filterMode* and *filterMask*.

- Strip VLAN tag and switch traffic:

```
[root@host~]# ovs-ofctl add-flow br0
in_port=1,dl_type=0x800,action=strip_vlan,output:2
```

- Insert VLAN tag 100 and switch traffic:

```
[root@host~]# ovs-ofctl -O OpenFlow11 add-flow br0
in_port=1,dl_type=0x800,action=push_vlan:0x8100,set_field:100-
\>vlan_vid,output:2
```

- Modify VLAN tag 30 to 50 and switch traffic:

```
[root@host~]# ovs-ofctl -O OpenFlow11 add-flow br0
in_port=1,dl_type=0x800,vlan_vid=30,action=mod_vlan_vid=50,output:2
```

If *vlan_vid* is not specieifed, the *mod_vlan_vid* tag will be added with a priority of 0.

- Modify VLAN priority and switch traffic:

```
[root@host~]# ovs-ofctl -O OpenFlow11 add-flow br0
in_port=1,dl_type=0x800,vlan_pcp=4,action=mod_vlan_pcp=3,output:2
```

3.2.1.3. VXLAN Flows

The following steps describe the method to configure VXLAN using OVS with single port on Server and Client machines.

Note

- *Only wild-card flows are currently supported with VXLAN matches.*
- *VxLAN VNI rewrite is not supported.*

- **Server**

i. Update the firmware configuration file, *t6-config.txt*, located at */lib/firmware/cxgb4/* :

```
filterMode = fragmentation, mpshittype, ethertype, vnic_id, port
filterMask = ethertype, vnic_id, port
vnicMode = encapsulation
```

ii. Load NIC (cxgb4) driver with hash-filter support:

```
[root@host~]# modprobe cxgb4 use_ddr_filters=1
```

iii. Bring up Chelsio interfaces in promiscous mode:

```
[root@host~]# ifconfig ethX <ip_address> promisc up
[root@host~]# ifconfig ethY <ip_address> promisc up
```

iv. Load Open vSwitch module:

```
[root@host~]# modprobe openvswitch
```

v. Configure OVS:

```
[root@host~]# ovs-appctl exit
[root@host~]# pkill -9 ovs
[root@host~]# rm -rf /usr/local/etc/ovs-vswitchd.conf
[root@host~]# rm -rf /usr/local/var/run/openvswitch/db.sock
[root@host~]# rm -rf /usr/local/etc/openvswitch/conf.db
[root@host~]# touch /usr/local/etc/ovs-vswitchd.conf
[root@host~]# ovsdb-tool create /usr/local/etc/openvswitch/conf.db
<uwire_package>/src/openvswitch-x.x.x/vswitchd/vswitch.ovsschema
[root@host~]# ovsdb-server /usr/local/etc/openvswitch/conf.db --
remote=punix:/usr/local/var/run/openvswitch/db.sock --
remote=db:Open_vSwitch,Open_vSwitch,manager_options --bootstrap-ca-
cert=db:Open_vSwitch,SSL,ca_cert --pidfile --detach --log-file
[root@host~]# ovs-vsctl --no-wait init
[root@host~]# export DB_SOCK=/usr/local/var/run/openvswitch/db.sock
[root@host~]# ovs-vswitchd --pidfile -detach
[root@host~]# ovs-vsctl add-br br0
[root@host~]# ifconfig br0 <server_ip>/24 up
[root@host~]# ovs-vsctl add-port br0 <chelsio_port0_name> -- set interface
<chelsio_port0_name> type=vxlan options:remote_ip=<peer_chelsio_ip>
options:local_ip=<local_chelsio_ip> options:key=flow
[root@host~]# ovs-vsctl add-port br0 <chelsio_port1_name>
```

vi. Disable GRO on the chelsio adapter, bridge and VXLAN interfaces:

```
[root@host~]# ethtool -K <chelsio_port0_name> gro off
[root@host~]# ethtool -K <chelsio_port1_name> gro off
[root@host~]# ethtool -K br0 gro off
[root@host~]# ethtool -K vxlan_sys_* gro off
```

vii. Set a rule to match packets with VNI=42 and drop:

```
[root@host~]# ovs-ofctl add-flow br0 in_port=2,tun_id=0x2a,action=drop
```

- **Client**

i. Follow steps (i)-(vi) described in the previous section **Server**.
ii. Set a rule to set VNI to 42 and send traffic:

```
[root@host~]# ovs-ofctl add-flow br0
in_port=LOCAL,action=set_tunnel:0x2a,output:2
```

3.3. Verifying OVS Flow Dump

OVS flow dump can be verified using:

```
[root@host~]# ovs-ofctl dump-flows br0
```

Run traffic between hosts which matches the flow and verify if the *n_packet* counter is incrementing.

3.4. Setting up ODL with OVS

The following example explains the method to set up OpenDaylight(ODL) using OVS:

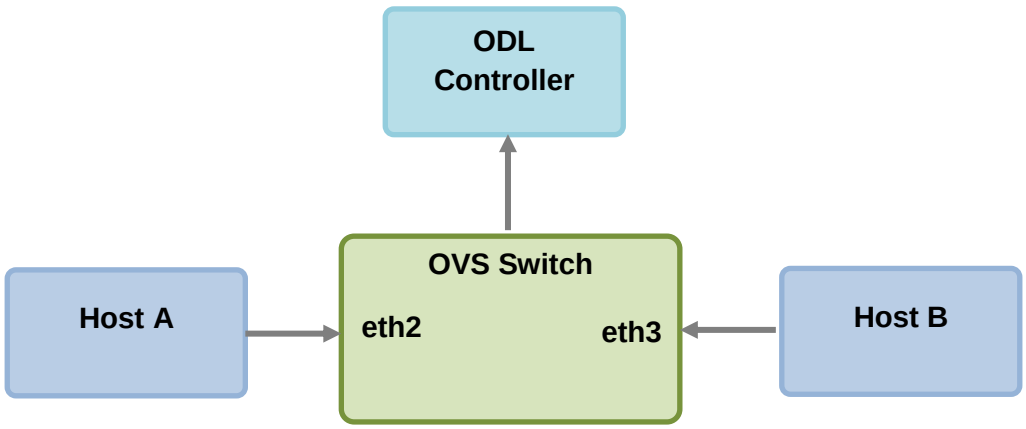

**eth2* and *eth3* are Chelsio interfaces.

On the ODL controller setup,

i. Download latest Java Development Kit.
ii. Untar the tar file.

iii. Create an entry in *.bashrc* which points to the extracted folder:

```
export JAVA_HOME=<path>/jdk1.8.0_92
export PATH=$PATH:$JAVA_HOME
```

iv. Logout & log back in.
v. Download ODL controller pre-built zip package.
vi. Unzip the package and change your working directory to *opendaylight.*
vii. Run the script *run.sh* and wait for ~3 minutes for the controller to be setup.
viii. Open a web browser and enter the address *http://localhost:8080*
ix. Login with *admin* keyword for both username and password.
x. On the OVS machine, add the bridge to the controller and disable in-band:

```
[root@host ~]# ovs-vsctl set-controller br0 tcp:<ODL Controller IP>:6633
[root@host ~]# ovs-vsctl set bridge br0 other-config:disable-in-band=true
```

xi. Refresh the webpage on the ODL controller and you should see the OVS details.
xii. Goto **Flows** tab, add and install a flow.
xiii. Verify the flow dump on the OVS machine:

```
[root@host ~]# ovs-ofctl dump-flows br0
```

Run traffic between hosts which matches the flow and verify if the *n_packet* counter is incrementing.

XXII. Traffic Management

1. Introduction

Traffic Management capabilities built-in to Chelsio adapters can shape transmit data traffic through the use of sophisticated queuing and scheduling algorithms built-in to the ASIC hardware which provides fine-grained software control over latency and bandwidth parameters such as packet rate and byte rate. These features can be used in a variety of data center application environments to solve traffic management problems.

Traffic Management features in Chelsio‘s adapters allow the user to control three main things:

- Guarantee low latency in the presence of other traffic
- Control max bandwidth that a connection or a flow (a group of connections) can use
- Allocate available bandwidth to several connection or flows based on desired levels of performance

Once the offload transmit traffic shaping classes have been configured, individual offloaded connections (flows) may be assigned to a traffic shaping class in order to manage the flows as per the class configuration. The mechanism to accomplish this "flow to class" mapping assignment is the Connection Offload Policy (COP) configuration system.

1.1. Hardware Requirements

1.1.1. Supported adapters

The following are the currently shipping Chelsio adapters that are compatible with the Traffic Management feature.

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

- T420-CX

1.2. Software Requirements

1.2.1. Linux Requirements

Currently the Traffic Management feature is available for the following versions:

- RHEL 7.3, 3.10.0-514.el7
- RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)
- RHEL 7.3, 4.5.0-15.el7.aarch64 (ARM64)
- RHEL 7.2, 3.10.0-327.el7
- RHEL 6.9, 2.6.32-696.el6
- RHEL 6.8, 2.6.32-642.el6
- SLES 12 SP2, 4.4.21-69-default
- SLES 12 SP1, 3.12.49-11-default
- SLES 11 SP4, 3.0.101-63-default
- Ubuntu 16.04.1, 4.4.0-31-generic
- Ubuntu 14.04.4, 4.2.0-27-generic
- Kernel.org linux-4.9.49 (Minimum 4.9 kernel version supported is 4.9.13)
- Kernel.org linux-4.8

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

2. Software/Driver Loading

Important *Please ensure that all inbox drivers are unloaded before proceeding with unified wire drivers.*

```
[root@host~]# rmmod csiostor cxgb4i cxgbit iw_cxgb4 chcr cxgb4vf cxgb4
libcxgbi libcxgb
```

Traffic Management can be performed on non-offloaded connections as well as on offloaded connections.

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

Run the following commands to load the TOE driver:

```
[root@host~]# modprobe cxgb4
[root@host~]# modprobe t4_tom
```

3. Software/Driver Configuration and Fine-tuning

3.1. Traffic Management Rules

Traffic Management supports the following types of scheduler hierarchy levels which can be configured using the *cxgbtool* utility:

i. Class Rate Limiting
ii. Class Weighted Round Robin
iii. Channel Rate Limiting

3.1.1. Class Rate Limiting

This scheduler hierarchy level can be used to rate limit individual traffic classes or individual connections (flow) in a traffic class.

Class rate limiting can be configured using the following command:

```
[root@host~]# cxgbtool <ethX> sched-class params type packet level cl-rl
mode <scheduler-mode>  rate-unit <scheduler-rate-unit>  rate-mode
<scheduler-rate-mode> channel <Channel No.> class <scheduler-class-index>
max-rate <maximum-rate>  pkt-size <Packet size>
```

Here,

ethX	: Chelsio interface
scheduler-mode	: specifies whether the rule is configured for individual traffic classes or individual connections (flow) in a traffic class. Possible values include *flow* or *class*.
scheduler-rate-unit	: Specifies whether the rule is configured for bit-rate or packet rate. Possible values include *bits* or *pkts.*
scheduler-rate-mode	: Specifies whether the rule is configured to support a percent of the channel rate or an effective rate. Possible values include *relative* or *absolute.*
Channel No.	: Port on which data is flowing (0-3).
scheduler-class-index	: TCP traffic class (0-14 for T4 series of adapters and 0-15 for T5/T6 series of adapters).
maximum-rate	: Bit rate (Kbps) for this TCP stream.
Packet size	: TCP mss size in bytes; for example - for an MTU of 1500, use a packet size of 1460.

3.1.2. Class Weighted Round Robin

Incoming traffic flows from various applications can be prioritized and provisioned using a weighted round-robin scheduling algorithm.

Class weighted round robin can be configured using the following command:

```
[root@host~]# cxgbtool <ethX> sched-class params type packet level cl-wrr
channel <Channel No.>  class <scheduler-class-index> weight <Y>
```

Here,

ethX	: Chelsio interface.
Channel No.	: Port on which data is flowing (0-3).
scheduler-class-index	: TCP traffic class (0-14 for T4 adapters; 0-15 for T5/T6 adapters).
weight	: Weight to be used for a weighted-round-robin scheduling hierarchy. Possible values include 1 to 99.

3.1.3. Channel Rate Limiting

This scheduler hierarchy level can be used to rate limit individual channels.

Channel rate limiting can be configured using the following command:

```
[root@host~]# cxgbtool eth6 sched-class params type packet level ch-rl rate-
unit <scheduler-rate-unit>  rate-mode <scheduler-rate-mode>  channel 1 max-
rate <maximum-rate>
```

Here,

ethX	: Chelsio interface.
scheduler-rate-unit	: Specifies whether the traffic management rule is configured for bit-rate or packet-rate. Possible values include *bits* or *pkts.*
scheduler-rate-mode	: Specifies whether the traffic management rule is configured to support a percent of the channel rate or an effective rate. Possible values include *relative* or *absolute.*
Channel No.	: Port on which data is flowing (0-3).
maximum-rate	: Bit rate (Kbps) for this TCP stream. The lower limit is 1 Gbps.

3.2. Configuring Traffic Management

3.2.1. For Non-offloaded connections

Traffic Management of non-offloaded connections is a 2-step process. In the first step bind connections to indicated NIC TX queue using *tc* utility from *iproute2-3.9.0* package. In the second step bind the indicated NIC TX queue to the specified TC Scheduler class using the *cxgbtool* utility.

i. Load the network driver and bring up the interface:

```
[root@host~]# modprobe cxgb4
[root@host~]# ifconfig ethX up
```

ii. Bind connections to queues:

```
[root@host~]# tc qdisc add dev ethX root handle 1: multiq
[root@host~]# tc filter add dev ethX parent 1: protocol ip prio 1 u32  match
ip dst <IP address of destination> action skbedit queue_mapping <queue>
```

Note *For additional binding options, run [root@host~]# man tc*

iii. Now, bind the NIC TX queue with traffic class:

```
[root@host~]# cxgbtool ethX sched-queue <queue> <class>
```

Here,

ethX : Chelsio interface
queue : NIC TX queue
class : TX scheduler class

Note *If the TX queue is all, * or any negative value, the binding will apply to all of the TX queues associated with the interface. If the class is unbind, clear or any negative value, the TX queue(s) will be unbound from any current TX Scheduler Class binding.*

3.2.2. For Offloaded connections

Traffic Management of offloaded connections can be configured either by applying *COP* policies that associate offloaded connections to classes or by modifying the application.

Both the methods have been described below:

- **Applying *COP* policy**

i. Load the TOE driver and bring up the interface:

```
[root@host~]# modprobe t4_tom
[root@host~]# ifconfig ethX up
```

ii. Create a new policy file (say *new_policy_file*) and add the following line to associate connections with the given scheduling class.

Example:

src host 102.1.1.1 => offload class 0

The above example will associate all connections originating from IP address 102.1.1.1 with scheduling class 0

Note *If no specified rule matches a connection, a default setting will be used which disables offload for that connection. That is, there will always be a final implicit rule following all the rules in the input rule set of:*

all => !offload

iii. Compile the policy file using *COP*

```
[root@host~]# cop -d –o <output_policy_file> <new_policy_file>
```

iv. Apply the *COP* policy:

```
[root@host~]# cxgbtool ethX policy <output_policy_file>
```

Where,

ethX: Chelsio interface

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

- **Modifying the application**

The application can also be modified to associate connections to scheduling classes. Follow the steps mentioned below:

i. Determine the TCP socket file descriptor in the application through which data is sent.
ii. Declare and initialize a variable in the application:

```
int cl=1;
```

Here,

cl is the TCP traffic class(scheduler-class-index) that the user wishes to assign the data stream to. This value needs to be in the range of 0 to 7.

The application will function as per the parameters set for that traffic class.

iii. Add socket option definitions:

In order to use *setsockopt()* to set the options to the TCP socket, the following two definitions need to be made:

- SOL_SCHEDCLASS used for setting TCP traffic class, which has the value 290.
- IPPROTO_TCP used for setting the type of IP Protocol.

```
# define SOL_SCHEDCLASS 290
# define IPPROTO_TCP 6
```

iv. Use the *setsockopt()* function to set socket options:

The *setsockopt()* call must be mentioned after the connect() call.

```
//Get the TCP socket descriptor variable
setsockopt (sockfd , IPPROTO_TCP, SOL_SCHEDCLASS, &cl, sizeof(cl));
```

Here,

sockfd : The file descriptor of the TCP socket.
&cl : Pointer to the class variables.
sizeof(cl) : The size of the variable.

v. Now, compile the application.

4. Usage

4.1. Non-Offloaded Connections

The following example demonstrates the method to rate limit all TCP connections on class 0 to a rate of 300 Mbps for Non-offload connections:

i. Load the network driver and bring up the interface:

```
[root@host~]# modprobe cxgb4
[root@host~]# ifconfig eth0 up
```

ii. Bind connections with destination IP address 192.168.5.3 to NIC TX queue 3

```
[root@host~]# tc qdisc add dev eth0 root handle 1: multiq
[root@host~]# tc filter add dev eth0 parent 1: protocol ip prio 1 u32  match
ip dst 192.168.5.3 action skbedit queue_mapping 3
```

iii. Bind the NIC TX queue to class 0

```
[root@host~]# cxgbtool eth0 sched-queue 3 0
```

iv. Set the appropriate rule for class 0

```
[root@host~]# cxgbtool eth0 sched-class params  type packet  level cl-rl
mode class  rate-unit bits  rate-mode absolute channel 0  class 0 max-rate
300000  pkt-size 1460
```

4.2. Offloaded Connections

The following example demonstrates the method to rate limit all TCP connections on class 0 to a rate of 300 Mbps for offloaded connections:

i. Load the TOE driver and bring up the interface

```
[root@host~]# modprobe t4_tom
[root@host~]# ifconfig eth0 up
```

ii. Create a new policy file (say *new_policy_file*) and add the following line to associate connections with the given scheduling class.:

```
src host 102.1.1.1 => offload class 0
```

iii. Compile the policy file using *COP*

```
[root@host~]# cop -d –o <output_policy_file> <new_policy_file>
```

iv. Apply the *COP* policy:

```
[root@host~]# cxgbtool eth0 policy <output_policy_file>
```

v. Set the appropriate rule for class 0

```
[root@host~]# cxgbtool ethX sched-class params  type packet  level cl-rl
mode class  rate-unit bits  rate-mode absolute channel 0  class 0 max-rate
300000  pkt-size 1460
```

4.3. Offloaded Connections with Modified Application

The following example demonstrates the method to rate limit all TCP connections on class 0 to a rate of 300 Mbps for for offloaded connections with modified application.

i. Load the TOE driver and bring up the interface.

```
[root@host~]# modprobe t4_tom
[root@host~]# ifconfig eth0 up
```

ii. Modify the application as mentioned in the Configuring Traffic Management section.

iii. Set the appropriate rule for class 0

```
[root@host~]# cxgbtool ethX sched-class params  type packet  level cl-rl
mode class  rate-unit bits  rate-mode absolute channel 0  class 0 max-rate
300000  pkt-size 1460
```

5. Software/Driver Unloading

Reboot the system to unload the driver. To unload without rebooting, refer Unloading the TOE driver section of **Network (NIC/TOE)** chapter.

XXIII. DPDK Driver

1. Introduction

Chelsio Data Plane Development Kit (DPDK) driver package is a collection of data plane libraries and NIC drivers optimized for running in the Linux user space to boost packet processing.

The driver has support for:

- Multiple queues for Tx and Rx
- Receive Side Scaling (RSS)
- VLAN filtering
- Checksum offload
- Promiscuous mode
- All multicast mode
- Port hardware statistics
- Jumbo frames (only for UIO)
- Classification and Filtering (only for UIO)

Note *Chelsio DPDK driver is built using open-source DPDK driver v17.02 with added support for Chelsio adapters.*

1.1. Hardware Requirements

The following are the currently shipping Chelsio adapters that are compatible with DPDK 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
- T520-BT

1.2. Software Requirements

The Chelsio DPDK driver has been developed to run on 64-bit Linux platforms. Following is the list of supported distributions:

Linux Distribution	Driver/Software
RHEL 7.3, 3.10.0-514.el7	UIO, VFIO
RHEL 7.2, 3.10.0-327.el7	
RHEL 7.3, 3.10.0-514.el7.ppc64le (POWER8 LE)	VFIO

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

Note *flex, bison, byacc, patch, patchutils, autoconf, automake and rpm-build packages must be present in the machine (required for libpcap installation).*

2. Software/Driver Installation

2.1. Pre-requisites

flex, *bison*, *byacc*, *patch*, *patchutils*, *autoconf*, *automake* and *rpm-build* packages must be present in the machine (required for *libpcap* installation).

2.2. Installation

Important *Please ensure that any existing version of DPDK driver is uninstalled before proceeding.*

i. Change your current working directory to Chelsio Unified Wire package directory:

```
[root@host~]# cd ChelsioUwire-x.x.x.x
```

ii. Install DPDK source, testpmd, tools and Pktgen:

```
[root@host~]# make dpdk_install
```

Note *The above command will*

- *install the driver with the following build configuration:*
  arch: x86_64/ppc_64
  machine: native/power8
  execenv: linuxapp
  toolchain: gcc
- *create a target environment directory "x86_64-native-linuxapp-gcc/ppc_64-power8-linuxapp-gcc" in the driver package directory.*

Note *The pktgen application has been modified to disable packet classification, to reflect correct Rx rate as packet classification is an expensive operation.*

Note *For more installation options, please run* `make help` *or* `install.py -h`

3. Flashing Firmware Configuration File

i. If not done already, load the kernel mode NIC driver (cxgb4):

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

ii. Flash the firmware configuration file (for T6 adapters):

```
[root@host~]# cxgbtool <iface> loadcfg /lib/firmware/cxgb4/t6-config.txt
```

Note *For T5 adapters, use t5-config.txt.*

4. Software/Driver Loading

Important *Before proceeding, unload all Chelsio inbox and outbox kernel mode drivers using the following commands.*

```
[root@host~]# rmmod cxgb4
[root@host~]# rmmod csiostor
```

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

UIO Support

Follow the steps mentioned below to load DPDK driver with UIO support:

i. Disable Intel VT-d in system BIOS.
ii. Turn off Intel iommu by adding below entry in Kernel grub/grub2 menu:

```
intel_iommu=off
```

iii. Reboot the system for changes to take effect.
iv. Load the UIO module:

```
[root@host~]# modprobe uio
[root@host~]# insmod Chelsio-DPDK-x.x.x.x/src/DPDK/x86_64-native-linuxapp-
gcc/kmod/igb_uio.ko
```

VFIO Support

Follow the steps mentioned below to load DPDK driver with VFIO support:

i. Enable Intel VT-d in system BIOS.
ii. Add the following entry to grub/grub2 menu:

```
intel_iommu=on vfio_iommu_type1.allow_unsafe_interrupts=1
```

E.g.:

```
kernel /vmlinuz-3.13.0-32-generic ro root=UUID=5149fae1-c52b-42a9-a48c-
b0a70937e8fb intel_iommu=on vfio_iommu_type1.allow_unsafe_interrupts=1
rd_NO_LUKS rd_NO_LVM LANG=en_US.UTF-8 rd_NO_MD SYSFONT=latarcyrheb-sun16
crashkernel=128M KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
```

iii. Reboot the system for changes to take effect.
iv. Load the VFIO module:

```
[root@host~]# modprobe vfio-pci
```

5. Software/Driver Configuration and Fine Tuning

5.1. Huge Pages

5.1.1. Using script

Run the *res_hugepages.sh* shell script (copied to */sbin* during installation):

```
[root@host~]# res_hugepages.sh
```

5.1.2. Manual

i. Mount *hugetlbfs*:

```
[root@host~]# mkdir -p /mnt/huge
[root@host~]# mount -t hugetlbfs nodev /mnt/huge
```

ii. Reserve Huge Page memory manually:

- x86_64

```
[root@host~]# echo 1024 >
/sys/devices/system/node/node0/hugepages/hugepages-2048kB/nr_hugepages
```

- POWERPC64

```
[root@host~]# echo 512 > /sys/devices/system/node/node0/hugepages/hugepages-
16384kB/nr_hugepages
```

In case of dual socket machines, run the above command for other CPU nodes.

5.2. Binding network ports

Note *Please make sure that adapter's physical function 4 is used, since it is assigned for all NIC functions of the adapter.*

- Run the following command to bind network ports to DPDK environment, with UIO support:

```
[root@host~]# dpdk_nic_bind.py --bind=igb_uio <PCI-ID-PF4>
```

Now, verify using:

```
[root@host~]# dpdk_nic_bind.py -status
```

Example:

```
[root@localhost ~]# dpdk_nic_bind.py --bind=igb_uio 0000:03:00.4
[root@localhost ~]# dpdk_nic_bind.py --status

Network devices using DPDK-compatible driver
============================================
0000:03:00.4 'T6225-CR Unified Wire Ethernet Controller' drv=igb_uio unused=
```

- Run the following command to bind network ports to DPDK environment, with VFIO support:

```
[root@host~]# dpdk_nic_bind.py --bind=vfio-pci <PCI-ID-PF4>
```

Now, verify using:

```
[root@host~]# dpdk_nic_bind.py -status
```

Example:

```
[root@localhost ~]# dpdk_nic_bind.py --bind=vfio-pci 0000:03:00.4
[root@localhost ~]# dpdk_nic_bind.py --status

Network devices using DPDK-compatible driver
============================================
0000:03:00.4 'T6225-CR Unified Wire Ethernet Controller' drv=vfio-pci unused=
```

5.3. Unbinding network ports

i. Run the following command to unbind network ports from DPDK environment:

```
[root@host~]# dpdk_nic_bind.py -u <PCI-ID-PF4>
```

ii. Run the following command to verify:

```
[root@host~]# dpdk_nic_bind.py –status
```

5.4. Performance Tuning

In order to auto tune the system for best performance, Chelsio recommends:

- disabling virtualization, c-state technology, Intel I/O AT and SR-IOV in the BIOS settings.
- setting the Power Profile to Maximum Performance in BIOS settings.
- installing the adapter into a PCIe Gen3 x8/x16 slot.
- using traffic with multiple tuples. The 'range' option of DPDK-Pktgen app can be used to create traffic with different 4-tuple values, so as to take advantage of RSS and hence, allow the traffic to be distributed across different queues/CPUs in Rx direction.

6. Running DPDK Test Applications

6.1. Testpmd application

The **Testpmd** application is provided as a part of the Chelsio DPDK driver package, and can be used to test the DPDK in a packet forwarding mode and also to access NIC hardware features such as Flow Director.

6.1.1. Syntax

Execute the following command to test DPDK using the **Testpmd** application:

```
[root@host~]# testpmd -c <coremask> -n <channels> -- -i --nb-cores=<cores> --txq=<TxQueue> --rxq=<RxQueue> --max-pkt-len=<PacketSize>
```

Where,

-c <coremask>	: A hexadecimal bit mask of the cores to run on. Note that core numbering can change between platforms and should be determined beforehand.
-n <channels>	: Number of memory channels per processor socket.
-i	: Enable interactive mode.
--nb-cores=<cores>	: Number of forwarding cores.
--txq=<TxQueue>	: Number of Tx queues.
--rxq=<RxQueue>	: Number of Rx queues.
--max-pkt-len=<PacketSize>	:Enable Jumbo mode and specify the packet size allowed for forwarding.

Note *In case of POWERPC64, if testpmd is not initializing, reallocate the huge pages using the setup script "setup.sh" present in <driver-package>/src/DPDK/tools/ directory.*

6.1.2. Examples

i. To run traffic using single Tx and Rx queue:

```
[root@host~]# testpmd -c f -n 4 -- -i --nb-cores=2 --txq=1 --rxq=1
```

ii. To run traffic using 4 Tx and Rx queues:

```
[root@host~]# testpmd -c 3ff -n 4 -- -i --nb-cores=8 --txq=4 --rxq=4
```

6.1.3. Flow Control

Flow control pause TX/RX is disabled by default and can be enabled via *Testpmd* as follows:

```
testpmd> set flow_ctrl rx on tx on 0 0 0 0 mac_ctrl_frame_fwd off autoneg
off 0
testpmd> set flow_ctrl rx on tx on 0 0 0 0 mac_ctrl_frame_fwd off autoneg
off 1
```

To disable again, run:

```
testpmd> set flow_ctrl rx off tx off 0 0 0 0 mac_ctrl_frame_fwd off autoneg
on 0
testpmd> set flow_ctrl rx off tx off 0 0 0 0 mac_ctrl_frame_fwd off autoneg
on 1
```

6.1.4. RSS

RSS is enabled by default and can be disabled via *Testpmd* as follows:

```
testpmd> port config all rss none
```

To enable again, run:

```
testpmd> port config all rss [all|ip|tcp|udp|sctp|ether|none]
```

6.1.5. Jumbo Mode

There are multiple ways to enable sending and receiving of jumbo frames: the first method involves using the command prompt and the other two via testpmd application.

- **Command Prompt**

Run the following command at the prompt:

```
[root@host~]# testpmd -c ffff -n 3 -- -i --nb-cores=8 --txq=4 --rxq=4 --max-
pkt-len=9018
```

- **Testpmd**

i. This method involves using the mtu command, which changes the MTU of an individual port without having to stop the selected port. To configure each port individually, run the mtu command as follows:

```
testpmd> port config mtu 0 9000
testpmd> port config mtu 1 9000
```

ii. This method involves stopping all the ports first and then running *max-pkt-len* command to configure the MTU of all the ports with a single command.To configure all the ports at once, stop all the ports first and run the *max-pkt-len* command as follows:

```
testpmd> port stop all
testpmd> port config all max-pkt-len 9000
```

6.2. Pktgen application

Pktgen is a traffic generator application capable of displaying real time metrics for ports and can handle packets with UDP, TCP, ARP, ICMP, GRE, MPLS and Queue-in-Queue. The application can run command scripts to set up repeatable test cases.

Note *The pktgen application has been modified to disable packet classification, to reflect correct Rx rate as packet classification is an expensive operation.*

6.2.1. Syntax

Execute the following command to run **Pktgen** application. Observe unidirectional pkts/sec and throughput performance under tool output:

```
[root@host~]# pktgen -c <coremask> -J -n <channels> -- -T -P -m
"<[cores]>.<port_id>" -N
```

Where,

-c <coremask>	: A hexadecimal bit mask of the cores to run on. Note that core numbering can change between platforms and should be determined beforehand.
-n <channels>	: Number of memory channels per processor socket.
-T	: Enable pktgen themes.
-P	: Enable PROMISCUOUS mode on all ports.
-m <[cores]>.<port_id>	: Map CPU logical cores to Chelsio ports.
-N	: Enable NUMA support.
-J	: Enable Jumbo mode.

6.2.2. Examples

- **Running single Tx and Rx queue traffic**

i. Configure Pktgen tool to use single Tx and Rx queue, mapping core 1 to Rx and core 2 to Tx queue on port 0:

```
[root@host~]# pktgen -c f -n 4 -- -T -P -m "[1:2].0" -N
.
.
.
Pktgen created by Keith Wiles -- >>> Powered by IntelÂ® DPDK <<<
- Ports 0-3 of 4   ** Main Page **  Copyright (c) <2010-2015>, Wind River
Systems, Inc. All rights reserved. Powered by IntelÂ® DPDK
  Flags:Port    :   P-------------:0
```

```
Link State       :      <UP-10000-FD>                          ---TotalRate---
Pkts/s  Rx       :                  0                                      0
        Tx       :           14880593                               14880593
MBits/s Rx/Tx    :             0/9999                                 0/9999
Broadcast        :                  0
Multicast        :                  0
  64 Bytes       :                  0
  65-127         :                  0
  128-255        :                  0
  256-511        :                  0
  512-1023       :                  0
  1024-1518      :                  0
Runts/Jumbos     :                0/0
Errors Rx/Tx     :                0/0
Total Rx Pkts    :                  0
      Tx Pkts    :          175605364
      Rx MBs     :                  0
      Tx MBs     :             118006
ARP/ICMP Pkts    :                0/0
                 :
Tx Count/% Rate  :       Forever/100%
PktSize/Tx Burst:              64/32
Src/Dest Port    :          1234/5678
Pkt Type:VLAN ID:      IPv4/TCP:0001
Dst  IP Address  :        192.168.1.1
Src  IP Address  :     192.168.0.1/24
Dst MAC Address  :  00:00:00:00:00:00
Src MAC Address  :  00:07:43:29:3c:40
-- Pktgen Ver:2.8.5(DPDK-2.1.0) ------------------------------------------
```

ii. Set packet size to 64B on all ports

```
Pktgen> set all size 64
```

iii. Start Tx traffic on port 0

```
Pktgen> start 0
```

iv. Stop Tx traffic on port 0

```
Pktgen> stop 0
```

- **Running multiple Tx and Rx queue traffic**

i. Configure Pktgen tool to use 4 Tx and Rx queues, mapping cores 1-4 to Rx and 5-8 to Tx queues, on port 0:

```
[root@host~]# pktgen -c 3ff -n 4 -- -T -P -m "[1-4:5-8].0" -N
```

ii. Set packet size to 64B on all ports

```
Pktgen> set all size 64
```

iii. Start Tx traffic on port 0

```
Pktgen> start 0
```

iv. Stop Tx traffic on port 0

```
Pktgen> stop 0
```

- **Running Tx and Rx queue traffic in Jumbo mode**

i. Configure Pktgen tool to use 4 Tx and Rx queues, mapping cores 1-4 to Rx and 5-8 to Tx queues, on port 0 in Jumbo mode:

```
[root@host~]# pktgen -c 3ff -n 4 -- -T -J -P -m "[1-4:5-8].0" -N
```

ii. Set packet size to 9018B on all ports

```
Pktgen> set all size 9018
```

iii. Start Tx traffic on port 0

```
Pktgen> start 0
```

iv. Stop Tx traffic on port 0

```
Pktgen> stop 0
```

7. Software/Driver Unloading

- ## UIO Support

Run the following commands to unload DPDK driver with UIO support:

```
[root@host~]# rmmod igb_uio
[root@host~]# rmmod uio
```

If unloading *uio* module reports an error, unload the following dependent modules and try again:

```
[root@host~]# rmmod bnx2fc
[root@host~]# rmmod bnx2i
[root@host~]# rmmod cnic
```

- ## VFIO Support

Run the following command to unload DPDK driver with VFIO support:

```
[root@host~]# rmmod vfio-pci
```

8. Software/Driver Uninstallation

i. Change your working directory to Chelsio-DPDK-x.x.x.x directory:

```
[root@host~]# cd Chelsio-DPDK-x.x.x.x
```

ii. Uninstall the DPDK driver using the following command:

```
[root@host~]# make uninstall
```

XXIV. Unified Boot

1. Introduction

PXE is short for Preboot eXecution Environment and is used for booting computers over an Ethernet network using a Network Interface Card (NIC). FCoE SAN boot process involves installation of an operating system to an FC/FCoE disk and then booting from it. iSCSI SAN boot process involves installation of an operating system to an iSCSI disk and then booting from it.

This section of the guide explains how to configure and use Chelsio Unified Boot Option ROM which flashes PXE, iSCSI and FCoE Option ROM onto Chelsio's converged network adapters (CNAs). It adds functionalities like PXE, FCoE and iSCSI SAN boot.

This section of the guide also describes the use and configuration of Chelsio's DUD for OS installations via PXE server on FC/FCoE LUN and iSCSI LUN. This solution can be used for installing operating systems over an Ethernet network/SAN using Chelsio's Terminator based Converged Network adapters (CNAs).

1.1. Hardware Requirements

1.1.1. Supported platforms

Following is the list of hardware platforms supported by Chelsio Unified Boot software:

- Dell T5600
- DELL PowerEdge 2950
- DELL PowerEdge T110
- DELL PowerEdge T710
- DELL PowerEdge R220
- DELL PowerEdge R720
- IBM X3650 M2
- IBM X3650 M4*
- HP Proliant DL180 gen9
- HP ProLiant DL385G2
- Supermicro X7DWE
- Supermicro X8DTE-F
- Supermicro X8STE
- Supermicro X8DT6
- Supermicro X9SRL-F
- Supermicro X9SRE-3F
- Supermicro-X10DRi
- ASUS P5KPL
- ASUS P8Z68
- Lenovo X3650 M5
- Intel DQ57TM

* If system BIOS version is lower than 1.5 and both Legacy and uEFI are enabled, system will hang during POST. Please upgrade the BIOS version to 1.5 or higher to avoid this issue.

1.1.2. Supported Switches

Following is the list of network switches supported by Chelsio Unified Boot software:

- Cisco Nexus 5010 with 5.1(3) N1 (1a) firmware.
- Arista DCS-7124S-F
- Mellanox SX_PPC_M460EX

Other platforms/switches have not been tested and are not guaranteed to work.

1.1.3. Supported Adapters

Following are the currently shipping Chelsio adapters that are compatible with Chelsio Unified Boot software:

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

** Only PXE supported*

1.2. Software Requirements

Chelsio Unified Boot Option ROM software requires Disk Operating System to flash PXE ROM onto Chelsio adapters.

The Chelsio Driver Update Disk driver has been developed to run on 64-bit Linux platforms. Following is the list of Drivers/Software and supported Linux distributions:

Linux Distribution	Driver/Software (DUDs)
RHEL 7.3, 3.10.0-514.el7	PXE, FCoE, iSCSI
RHEL 7.2, 3.10.0-327.el7	PXE, iSCSI
RHEL 6.9, 2.6.32-696.el6	
SLES 12 SP2, 4.4.21-69-default	
SLES 12 SP1, 3.12.49-11-default	
SLES 11 SP4, 3.0.101-63-default	

1.3. Pre-requisites

A DOS bootable USB flash drive or Floppy Disk is required for updating firmware, option ROM, creating DUD, etc.

2. Secure Boot

Secure Boot, a high-performance computing software solution is a method to restrict which binaries can be executed to boot the system. With Secure Boot, the system BIOS will only allow the execution of boot loaders that carry the cryptographic signature of trusted entities. In other words, anything run in the BIOS must be “signed” with a key that the system knows is trustworthy. With each reboot of the server, every executed component is verified.

The following example describes the method to enable Secure Boot on HP ProLiant servers. Steps may differ slightly on other platforms:

i. During system boot, press F9 to run the **System Utilities**.

ii. Select **System Configuration**.

iii. Select **BIOS/Platform Configuration (RBSU)**.

iv. Select **Server Security**.

v. Select **Secure Boot Settings**.

vi. Select **Advanced Secure Boot Options**.

vii. Provide the Platform Key (PK), Key Exchange Key (KEK) and Allowed Signature Database (DB) to the respective uEFI NVRAM variables.

- **Windows**:
  - PK: Will be generated at the discretion of the platform owner (OEM). Click here for more information.
  - KEK: http://www.microsoft.com/pkiops/certs/MicCorKEKCA2011_2011-06-24.crt
  - Windows DB: http://www.microsoft.com/pkiops/certs/MicWinProPCA2011_2011-10-19.crt
  - uEFI DB: http://www.microsoft.com/pkiops/certs/MicCorUEFCA2011_2011-06-27.crt
  - Signature GUID for all the above keys: 77fa9abd-0359-4d32-bd60-28f4e78f784b
- **Linux**:
  - Use the same values for PK, KEK, Windows DB, uEFI DB and Signature ID as mentioned above.
  - In addition, provide the following values:
    - chcert.cer: *Provided in ChelsioUwire-x.x.x.x/Uboot/chelsio_key/*
    - Signature GUID for chcert.cer: 0b74ace7-6136-a493-19a9-6104d6d1e432

viii. Reboot the system, run **System Utilities** and go to **Secure Boot Settings**.

ix. Select and enable **Secure Boot Enforcement**.

x. Reboot the system.

3. Flashing firmware and option ROM

Depending on the boot mode selected, Chelsio Unified Boot provides two methods to flash firmware and option ROM onto Chelsio adapters: Flash utility *cfut4* for Legacy mode and *HII* for uEFI mode. On HP machines, you can also use Firmware Manager Protocol (FMP), in addition to HII. These methods also provide the functionality to update/erase Hardware configuration and Phy Firmware files.

3.1. Preparing USB flash drive

This document assumes that you are using a USB flash drive as a storage media for the necessary files. Follow the steps below to prepare the drive:

xi. Create a DOS bootable USB flash drive. (Click here for instructions)

xii. Create a directory *CHELSIO* on the USB flash drive.

xiii. If you haven't done already, download the driver package from Chelsio Download Center.

xiv. Untar the downloaded package and change your working directory to *OptionROM* directory.

```
[root@host~]# cd <driver_package>/Uboot/OptionROM
```

xv. Copy all the files and place them in the *CHELSIO* directory created on the USB flash drive.

xvi. Plug-in the USB flash drive in the system on which the Chelsio CNA is installed.

xvii. Reboot the system and go into the BIOS setup.

xviii. Make the USB flash drive as the primary boot device.

xix. Save the changes.

3.2. Legacy

i. In BIOS, configure the system having Chelsio CNA to boot in Legacy mode.

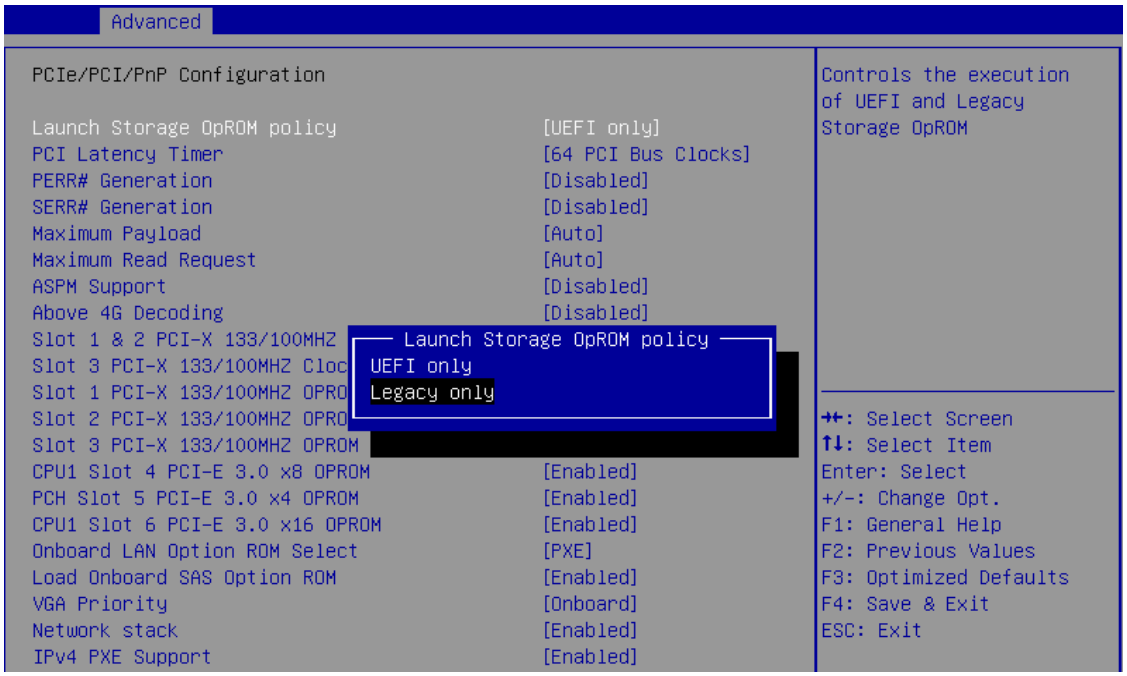

ii. Once the system boots from the USB flash drive, change your working directory to *CHELSIO* directory:

```
C:\>cd CHELSIO
```

iii. Run the following command to list all Chelsio CNAs present on the system. The list displays a unique index for each CNA found.

```
C:\CHELSIO>cfut4 -l
```

iv. Delete any previous version of Option ROM flashed onto the CNA:

```
C:\CHELSIO>cfut4 -d <idx> -xb
```

Here, `idx` is the CNA index found in step iii (0 in this case)

```
C:\CHELSIO>cfut4 -d 0 -xb

Chelsio T5/T6 Flash Utility v1.5

Erasing serial flash sector(s) ... Done
Reboot machine for changes to take effect
```

v. Delete any previous firmware using the following command:

```
C:\CHELSIO>cfut4 -d <idx> -xh -xf
```

```
C:\CHELSIO>cfut4 -d 0 -xh -xf

Chelsio T5/T6 Flash Utility v1.5

Erasing serial flash sector(s) ... Done
Erasing serial flash sector(s) ... Done
Reboot machine for changes to take effect
```

vi. Delete any previous Option ROM settings:

```
C:\CHELSIO>cfut4 -d <idx> -xc
```

```
C:\CHELSIO>cfut4 -d 0 -xc

Chelsio T5/T6 Flash Utility v1.5

Erasing serial flash sector(s) ... Done
Reboot machine for changes to take effect
```

vii. Run the following command to flash the appropriate firmware.

```
C:\CHELSIO>cfut4 -d <idx> -uf <firmware_file>.bin
```

Here, `firmware_file` is the firmware image file present in the *CHELSIO* directory.

```
C:\CHELSIO>cfut4 -d 0 -uf T6FW-1~1.BIN

Chelsio T5/T6 Flash Utility v1.5

Erasing serial flash sector(s) ... Done
Writing Image at Base 00080000 ... Done
Writing Image at Base 00088000 ... Done
Writing Image at Base 00090000 ... Done
Writing Image at Base 00098000 ... Done
Writing Image at Base 000a0000 ... Done
Writing Image at Base 000a8000 ... Done
Writing Image at Base 000b0000 ... Done
Writing Image at Base 000b8000 ... Done
Writing Image at Base 000c0000 ... Done
Writing Image at Base 000c8000 ... Done
Writing Image at Base 000d0000 ... Done
Writing Image at Base 000d8000 ... Done
Writing Image at Base 000e0000 ... Done
Writing Image at Base 000e8000 ... Done
Writing Image at Base 000f0000 ... Done
Writing Image at Base 000f8000 ... Done
Reboot machine for changes to take effect
```

viii. Flash the unified option ROM onto the Chelsio CNA using the following command:

```
C:\CHELSIO>cfut4 -d <idx> -ub cubt4.bin
```

Here, `cubt4.bin` is the unified option ROM image file present in the *CHELSIO* directory.

```
C:\CHELSIO>cfut4 -d 0 -ub cubt4.bin

Chelsio T5/T6 Flash Utility v1.5

Erasing serial flash sector(s) ... Done
Writing Image at Base 00000000 ... Done
Writing Image at Base 00008000 ... Done
Writing Image at Base 00010000 ... Done
Writing Image at Base 00018000 ... Done
Writing Image at Base 00020000 ... Done
Writing Image at Base 00028000 ... Done
Writing Image at Base 00030000 ... Done
Writing Image at Base 00038000 ... Done
Writing Image at Base 00040000 ... Done
Writing Image at Base 00048000 ... Done
Writing Image at Base 00050000 ... Done
Writing Image at Base 00058000 ... Done
Writing Image at Base 00060000 ... Done
Writing Image at Base 00068000 ... Done
Erasing serial flash sector(s) ... Done
Writing Image at Base 00070000 ... Done
Reboot machine for changes to take effect
```

ix. Reboot the system for changes to take effect.

x. To configure the base MAC address (optional), use the below command:

```
C:\CHELSIO>cfut4 -d <idx> -um <Hex MAC Address>
```

Here, *idx* is the CNA index found in step (c)

Example:

```
C:\CHELSIO>cfut4 -d 0 -um 000743000123
```

3.3. uEFI

i. Reboot the system and go into BIOS setup.

ii. Configure the system having Chelsio CNA to boot in uEFI mode.

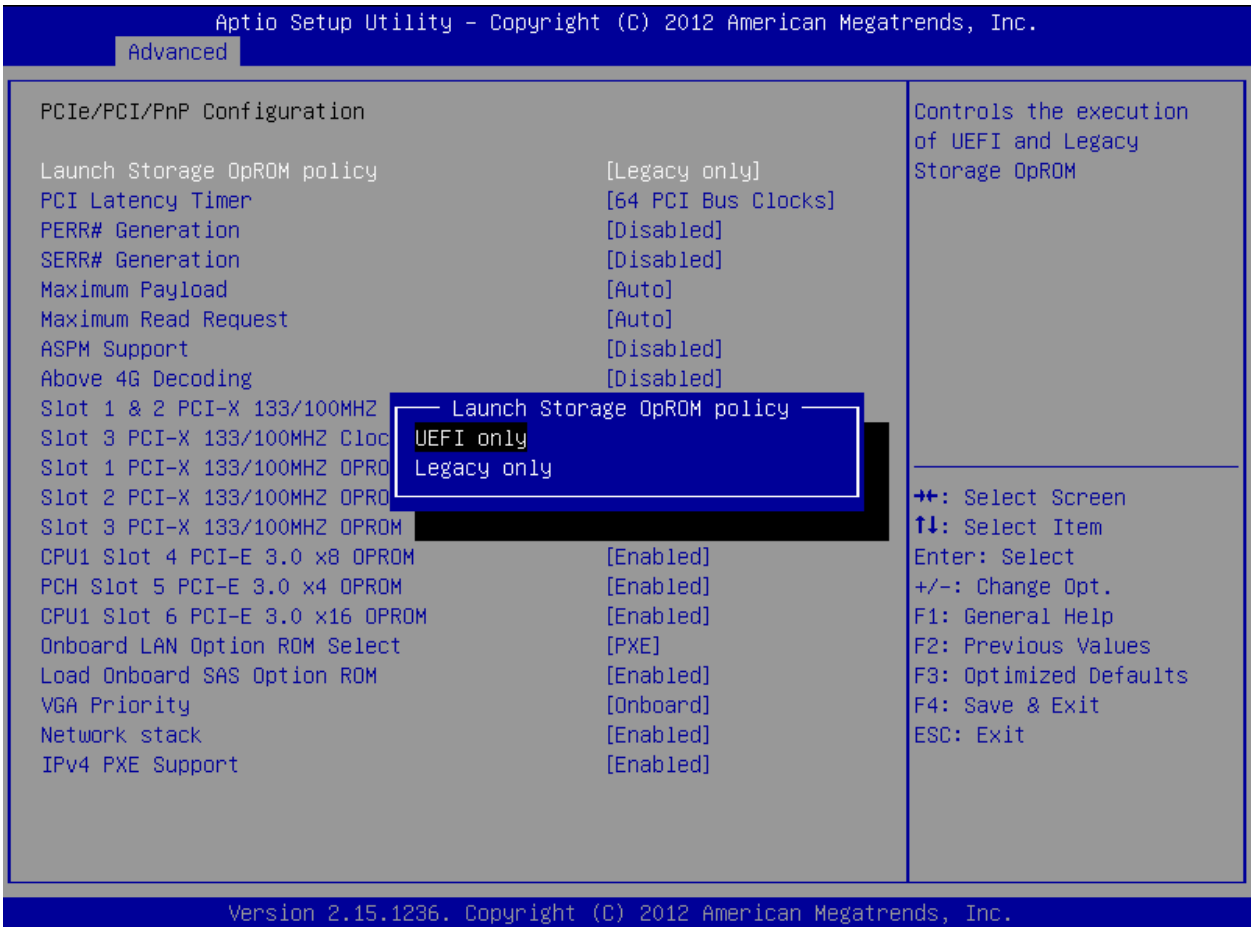

Note *For Supermicro systems, enable **Network Stack*** as well before proceeding.

iii. Boot to EFI Shell.

```
EFI Shell version 2.31 [4.654]
Current running mode 1.1.2
Device mapping table
  fs0  :Removable HardDisk - Alias hd83b0f0b blk0
        PciRoot(0x0)/Pci(0x1d,0x0)/USB(0x1,0x0)/USB(0x5,0x0)/HD(1,MBR,0x0fdb738d,0x800,0x78b800)
  blk0 :Removable HardDisk - Alias hd83b0f0b fs0
        PciRoot(0x0)/Pci(0x1d,0x0)/USB(0x1,0x0)/USB(0x5,0x0)/HD(1,MBR,0x0fdb738d,0x800,0x78b800)
  blk1 :HardDisk - Alias (null)
        PciRoot(0x0)/Pci(0x1f,0x2)/Sata(0x0,0x0)/HD(1,MBR,0x00092b0c,0x3f,0x9c25fe)
  blk2 :HardDisk - Alias (null)
        PciRoot(0x0)/Pci(0x1f,0x2)/Sata(0x0,0x0)/HD(2,MBR,0x00092b0c,0x9c263d,0x88b8fdc)
  blk3 :HardDisk - Alias (null)
        PciRoot(0x0)/Pci(0x1f,0x2)/Sata(0x0,0x0)/HD(3,MBR,0x00000000,0x927be19,0x14019e7)
  blk4 :HardDisk - Alias (null)
        PciRoot(0x0)/Pci(0x1f,0x2)/Sata(0x0,0x0)/HD(4,MBR,0x00000000,0xa67d83f,0x13fe849)
  blk5 :BlockDevice - Alias (null)
        PciRoot(0x0)/Pci(0x1f,0x2)/Sata(0x0,0x0)
  blk6 :Removable BlockDevice - Alias (null)
        PciRoot(0x0)/Pci(0x1d,0x0)/USB(0x1,0x0)/USB(0x5,0x0)

Press ESC in 1 seconds to skip startup.nsh, any other key to continue.
Shell> _
```

iv. Issue command `drivers` to determine if Chelsio uEFI driver is loaded. If the driver is loaded (as shown in the image below), continue to step (v)

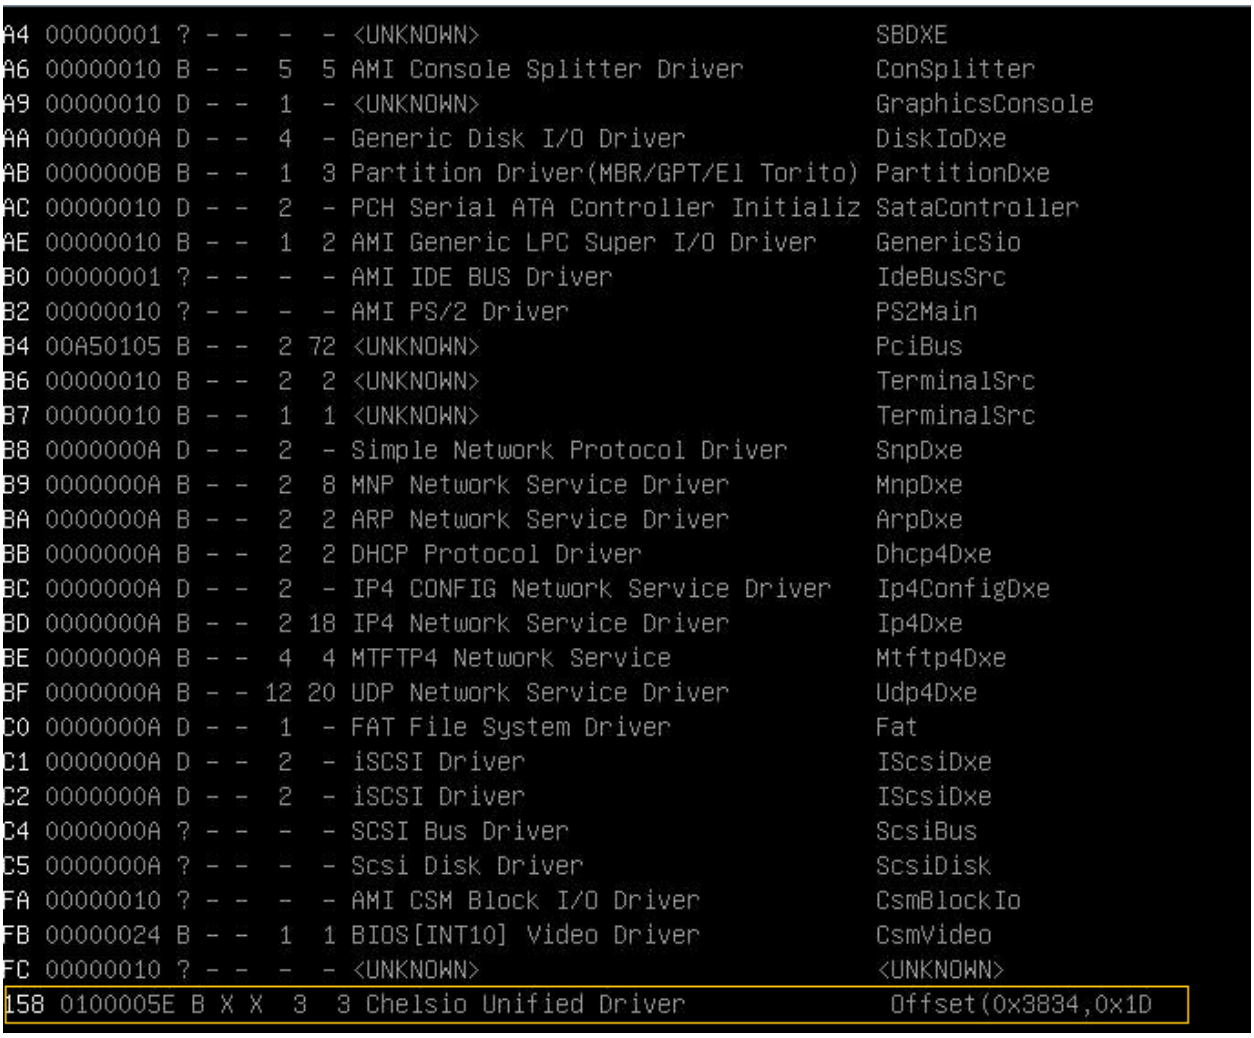

If the driver is not loaded, load the uEFI driver (*ChelsioUD.efi*) found in the CHELSIO directory, and try again.

```
fs0:\CHELSIO> load ChelsioUD.efi
load: Image fs0:\CHELSIO\ChelsioUD.efi loaded at 7F2BA000 - Success
```

v. Reboot the system and go into BIOS setup.

vi. Chelsio HII should be listed as ***Chelsio T5/T6***. Highlight it and press [Enter].

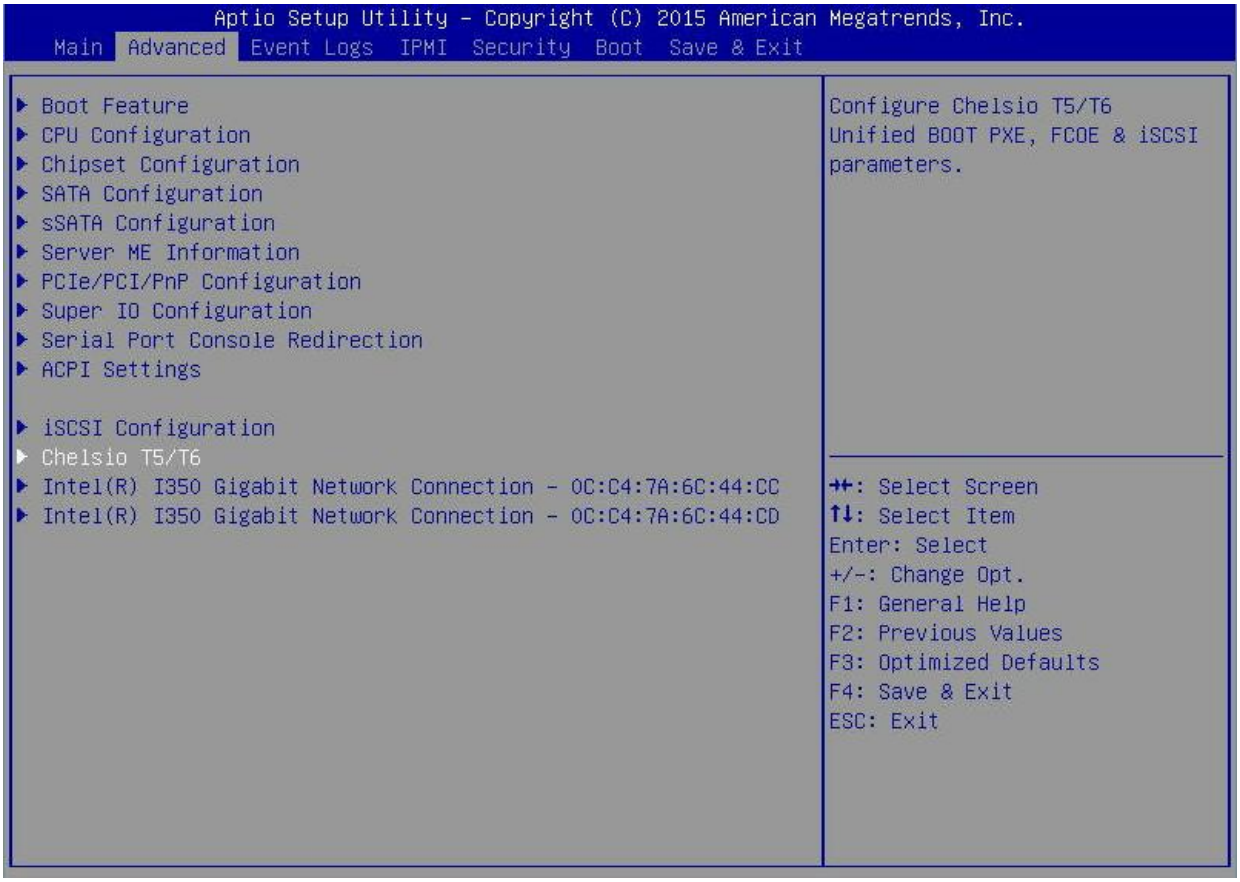

Note *If Chelsio T5/T6 is not listed, please ensure that Chelsio uEFI driver is loaded correctly as mentioned here in the* ***Flashing Firmware and Option ROM*** *section.*

vii. Highlight the Chelsio adapter to be configured and press [Enter].

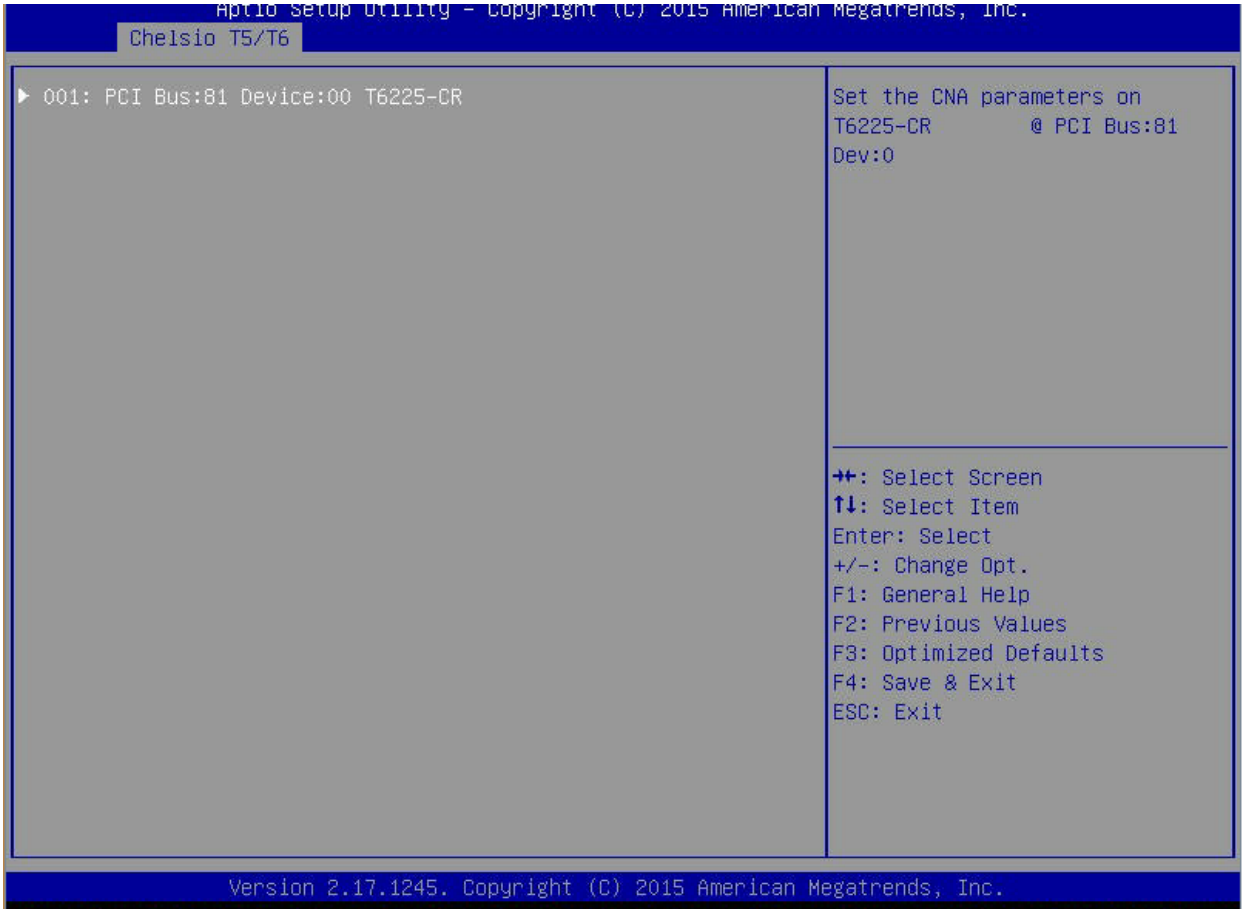

viii. Highlight **Flash Utility** and press [Enter].

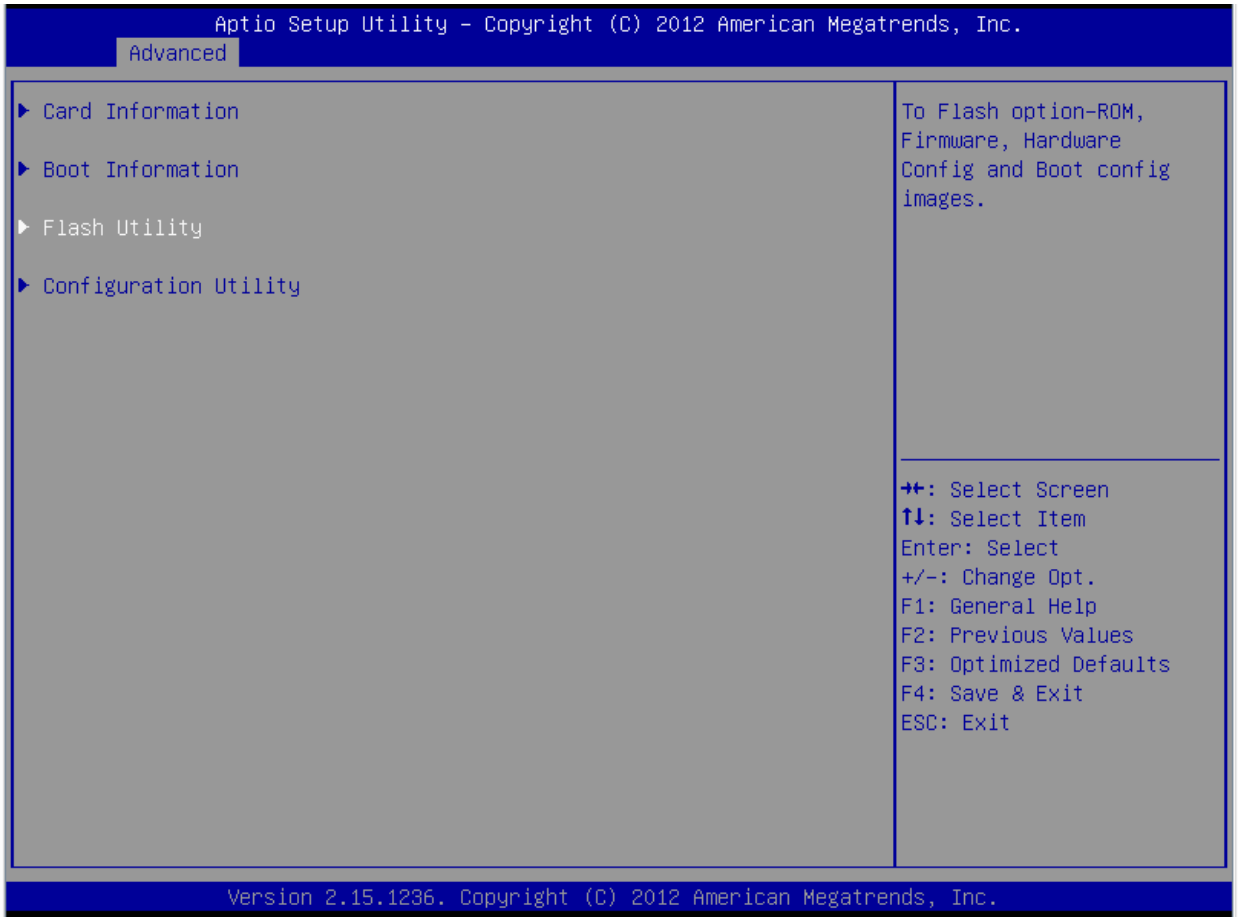

ix. Erase or update firmware using the methods explained below:

- **Erase existing firmware**
  a. Select [Erase] as Flash Operation
  b. Select [FW File] as Flash File Type
  c. Select Update/Erase
  d. Press [Y] to confirm.
  e. Reboot system.
- **Update firmware**
  a. Select [Update] as Flash Operation
  b. Select [FW File] as Flash File Type
  c. Enter full path to the firmware file for Enter File Name, e.g., CHELSIO\t6fw-1.16.29.0.bin.
  d. Press [Enter]
  e. Select Update/Erase
  f. Press [Y] to confirm.
  g. Reboot system

Similarly, you can use the above method to update/erase Option ROM, Hardware Configuration and Phy Firmware file.

3.4. HP Firmware Management Protocol (FMP)

Enabling FMP

For HP machines with Unified Boot v1.0.0.99 and below installed, FMP must be manually enabled. For v2.0.0.1 and above, you can skip to the next sub-section for instructions on how to upgrade firmware using FMP.

i. Prepare USB flash drive as described here.
ii. Boot system to EFI shell and change your working directory to *CHELSIO*.
iii. Issue command `drivers` to determine if Chelsio uEFI driver is loaded. If loaded, you should see a similar output:

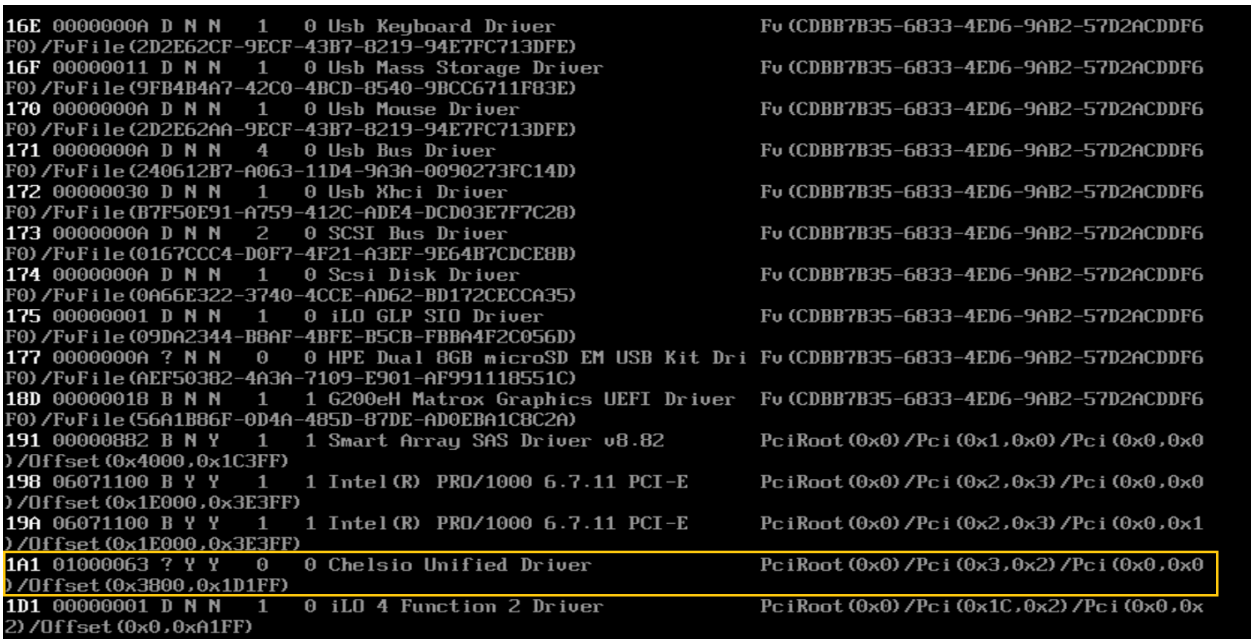

iv. Note the handle and unload the driver using the following syntax:

```
FS1:\CHELSIO\> unload -n <driver_handle>
```

Example:

v. Load the uEFI driver (*ChelsioUD.efi*) present in the CHELSIO directory.

```
FS1:\CHELSIO\> load ChelsioUD.efi
Image 'FS1:\CHELSIO\ChelsioUD.efi' loaded at 77588000 - Success
```

vi. If Chelsio uEFI driver was loaded successfully in the previous step, run the command `fwupdate -l` and Chelsio T6 CNA should be listed as shown below:

```
FS1:\CHELSIO\> fwupdate -l
* [BIOS] System ROM - U20 v2.20 (05/05/2016)
* [RAID.Slot.2.1]Slot 2 : Smart HBA H240 Controller - V2.52_B0
* [NIC.LOM.1.3]Embedded LOM 1 : HP Ethernet 1Gb 2-port 361i Adapter - NIC - 1.1067.0
* [NIC.Slot.3.1]Slot 3 : Chelsio T6 Controller - NIC -
```

Upgrading Firmware

Using CLI

i. Use the adapter's device name to update the firmware:

```
FS1:\CHELSIO\> fwupdate -d <device_name> -f cubt4.bin
```

Example:

```
FS1:\CHELSIO\> fwupdate -d NIC.Slot.3.1 -f cubt4.bin
Loading firmware file 'cubt4.bin'. It might take several minutes.
Current Firmware Version is .
Continue with firmware update? (y/n):y
Firmware update completed successfully.
```

ii. Reboot machine for changes to take effect.

- **Using FMP**

- Reboot system and press F9 to access **System Utilities**

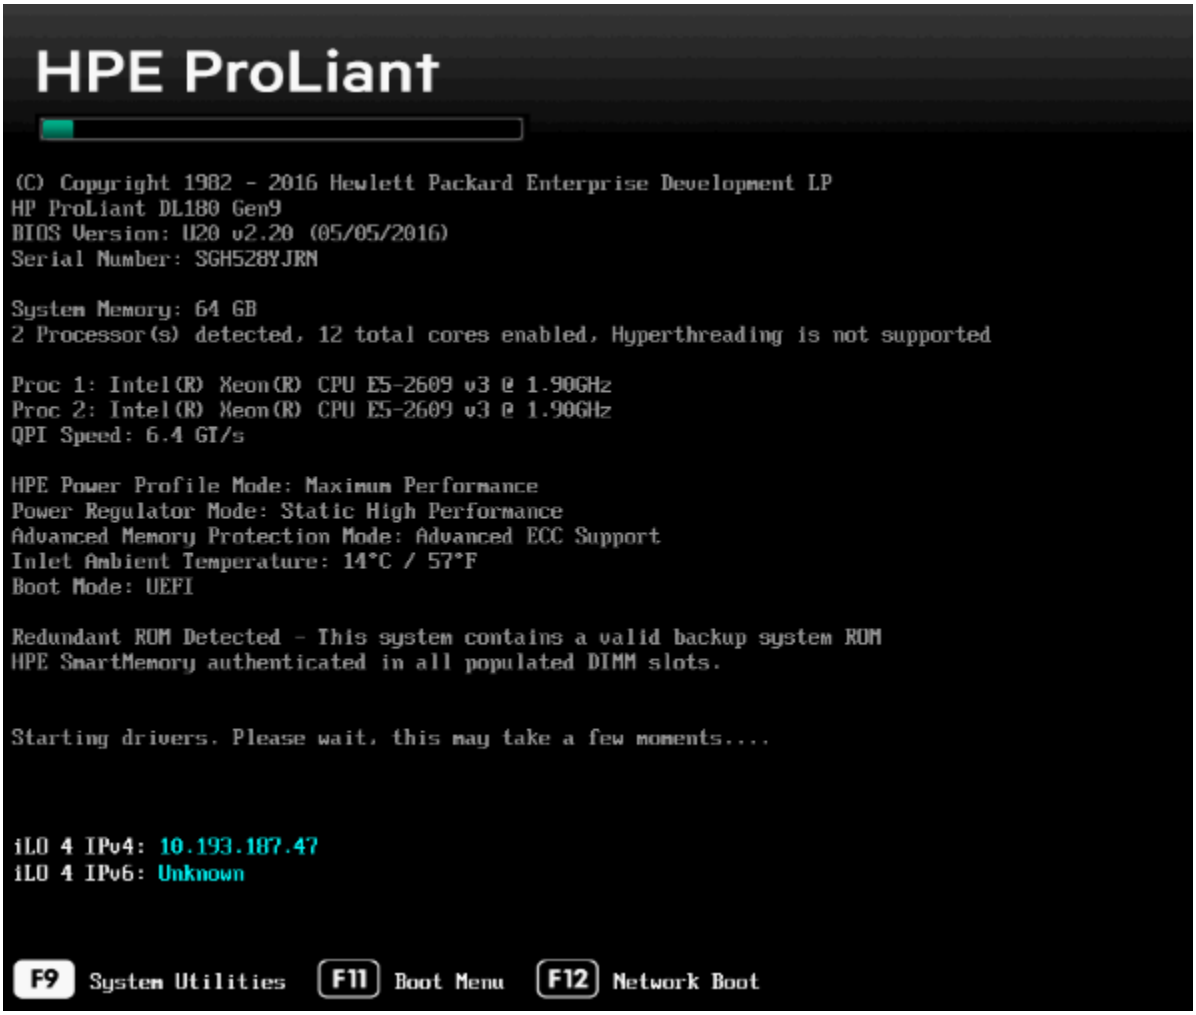

- Go to **Embedded Applications → Firmware Update → Chelsio T6 Controller**

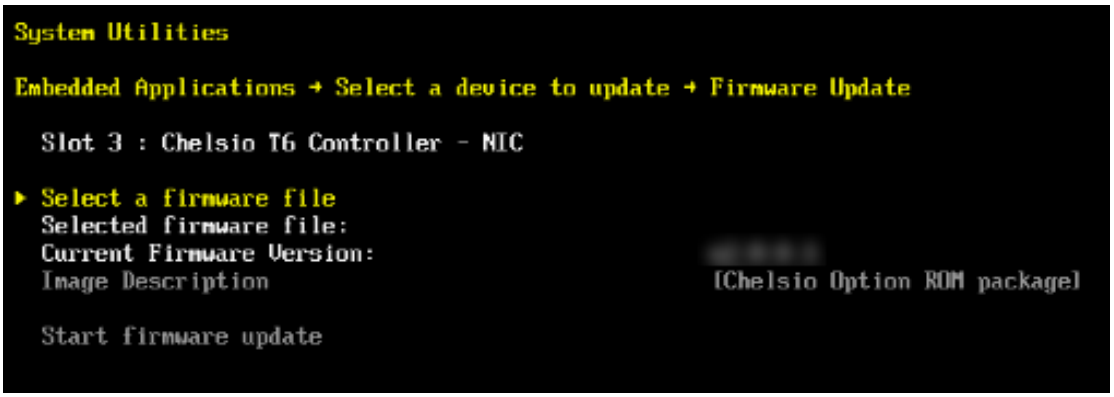

- Highlight ***Select a firmware file*** option and hit [Enter].

- Select the USB flash drive which contains the latest optionROM and hit [Enter].

- Select optionROM file *cubt4.bin* and hit [Enter]. The file should show up in the ***Selected firmware file*** field.

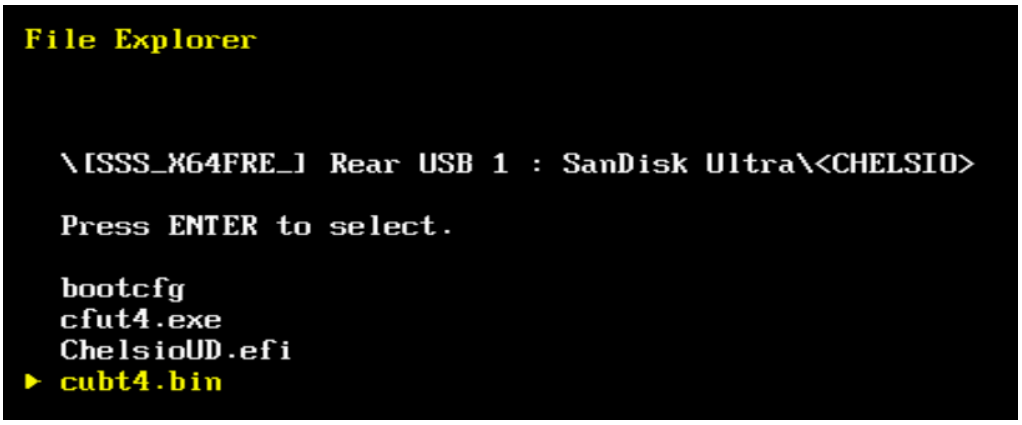

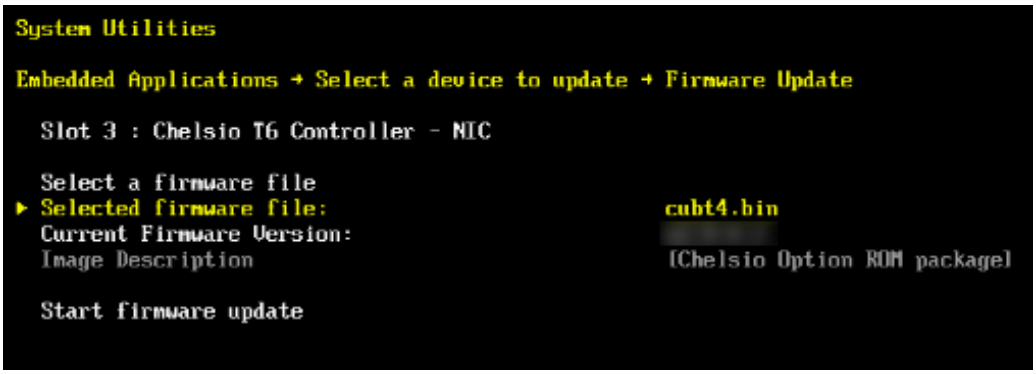

- Select ***Start firmware update*** and hit [Enter].

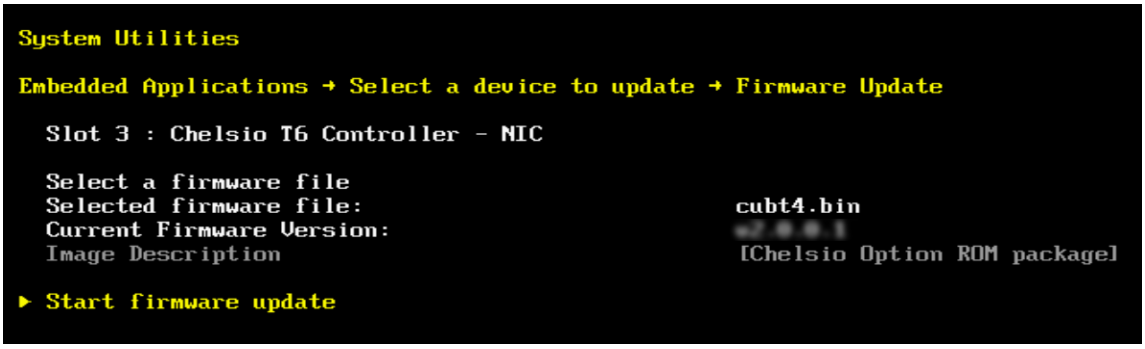

- After **Firmware update completed successfully** prompt appears, reboot the machine for changes to take effect.

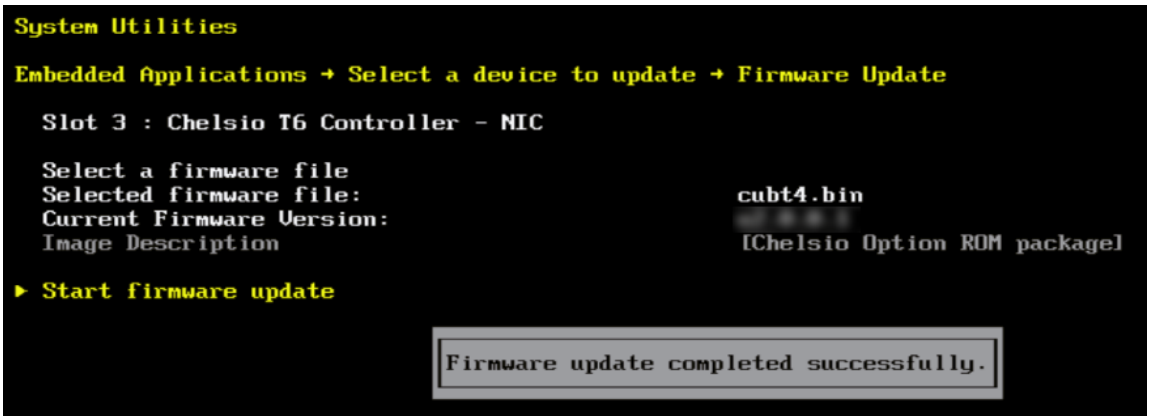

3.5. Default Option ROM Settings

If you wish to restore option ROM settings to their default values, i.e., PXE enabled, iSCSI and FCoE disabled, use any of the methods mentioned below:

3.5.1. Using Option ROM (boot level)

- **Legacy PXE**

Boot system into Chelsio's Unified Boot Setup utility and press F8.

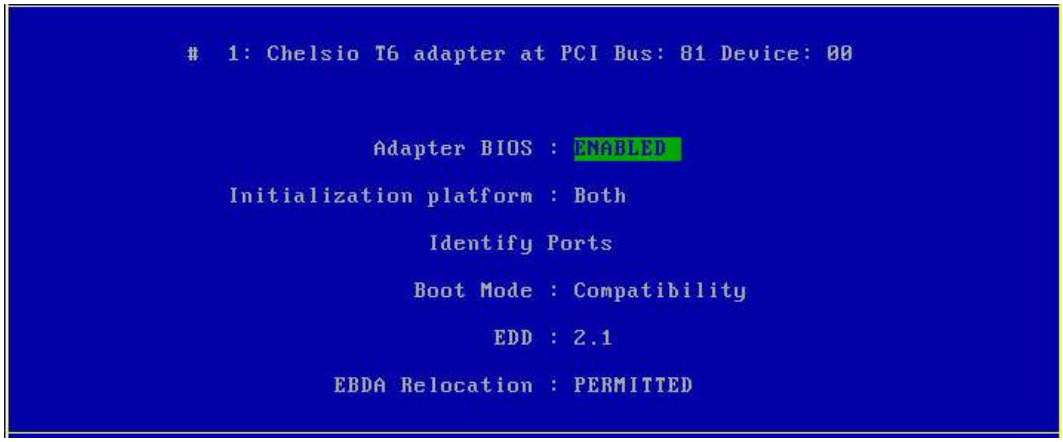

- **uEFI PXE**

Boot system into uEFI mode and press F3.

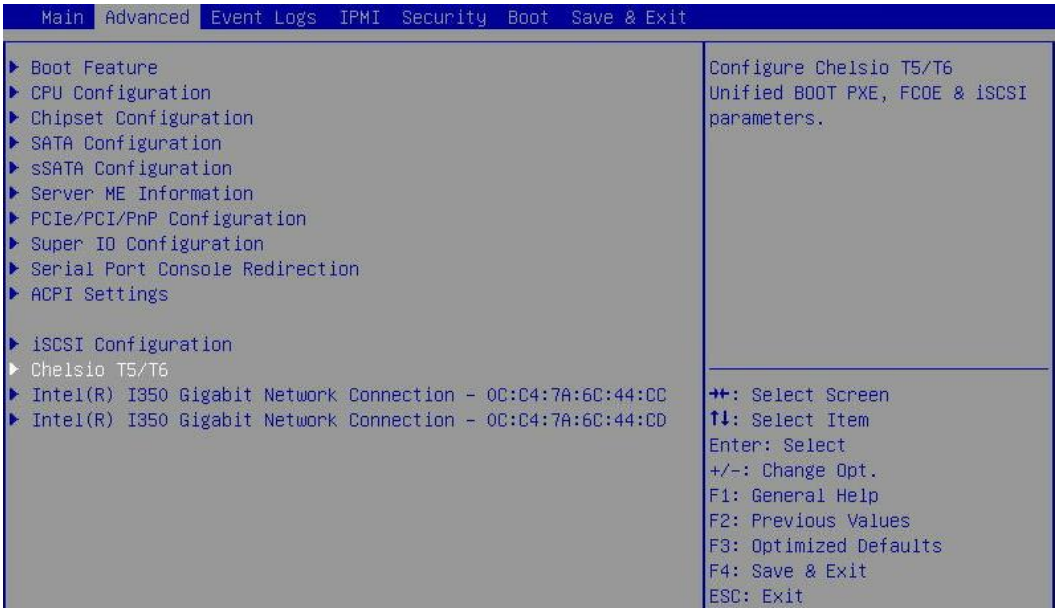

3.5.2. Using *cxgbtool* (OS level)

Change your working directory to *OptionROM* directory and use *cxgbtool* to flash the default boot configuration onto the adapter:

```
[root@host~]# cd <driver_pakage>/Uboot/OptionROM/
[root@host~]# cxgbtool <ethX> loadboot-cfg bootcfg
```

4. Configuring PXE Server

The following components are required to configure a PXE Server:

- DHCP Server
- TFTP Server

PXE server configuration steps for different operating systems can be found on following links:

Note *Chelsio Communications does not take any responsibility regarding contents given in below mentioned links. They are provided for example purposes only.*

Linux

- http://linux-sxs.org/internet_serving/pxeboot.html
- http://www.howtoforge.com/ubuntu_pxe_install_server

Windows

- http://technet.microsoft.com/en-us/library/cc771670%28WS.10%29.aspx
- http://tftpd32.jounin.net/ (Use port # 67, set PXE option and provide bootable file name in settings)
- http://unattended.sourceforge.net/pxe-win2k.html

5. PXE boot process

Before proceeding, please ensure that the Chelsio CNA has been flashed with the provided firmware and option ROM (See Flashing Firmware and option ROM).

5.1. Legacy PXE boot

i. After configuring the PXE server, make sure the PXE server works. Then reboot the client machine.

ii. Press [Alt+C] when the message to configure Chelsio adapters appears on the screen.

iii. The configuration utility will appear as below:

iv. Choose the CNA on which you flashed the option ROM image. Hit [Enter].

v. Enable the adapter BIOS using arrow keys if not already enabled. Hit [Enter].

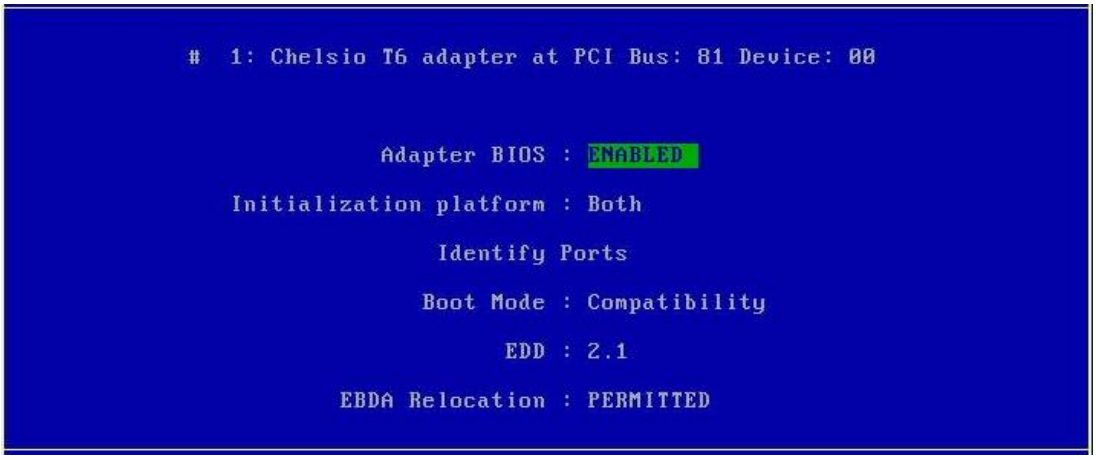

Note *Use the default values for Boot Mode, EDD and EBDA Relocation parameters, unless instructed otherwise.*

vi. Choose *PXE* from the list to configure. Hit [Enter].

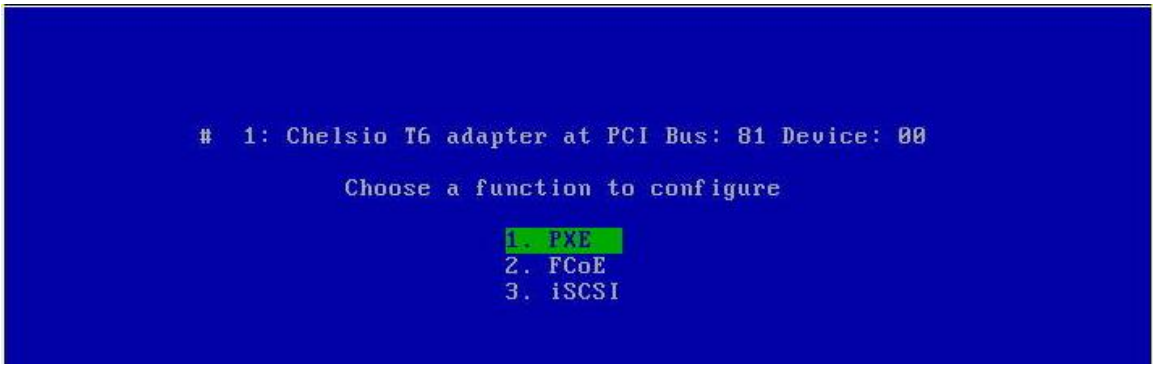

vii. Use the arrow keys to highlight the appropriate function among the supported NIC functions and hit [Enter] to select.

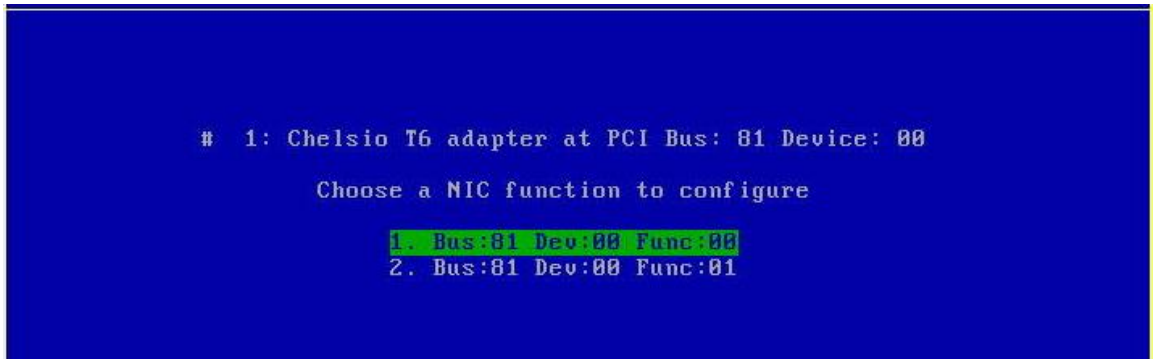

viii. Enable NIC function bios if not already enabled.

ix. Choose the boot port to try the PXE boot. It is recommended to only enable functions and ports which are going to be used. Please note that enabling NIC Func 00 will enable port 0 for PXE, enabling NIC Func 01 will enable port 1 and so on for NIC function.

x. Hit [F10] or [Esc] and then [Y] to save configuration changes.

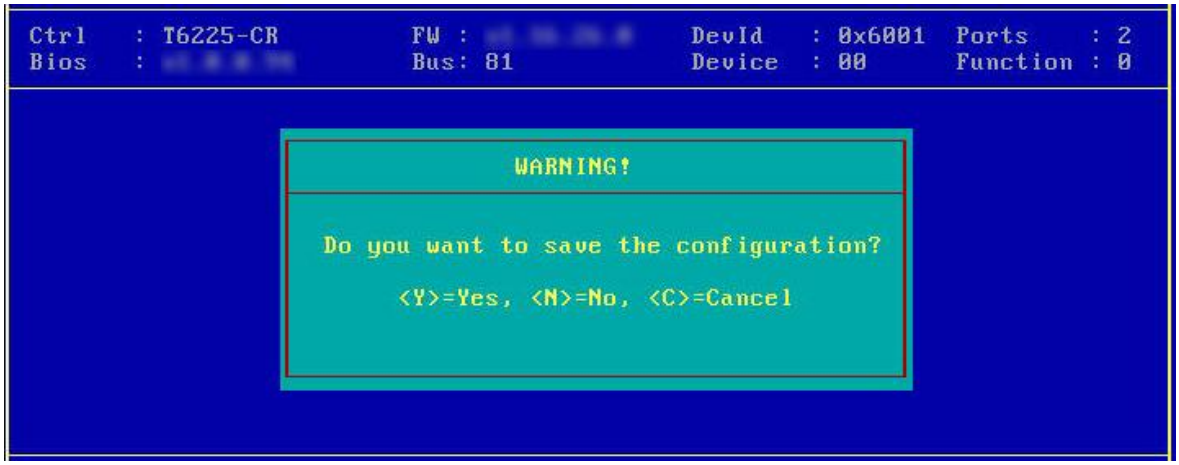

xi. Reboot the system.

xii. Allow the Chelsio option ROM to initialize and setup PXE devices. DO NOT PRESS ALT-S to skip Chelsio option ROM.

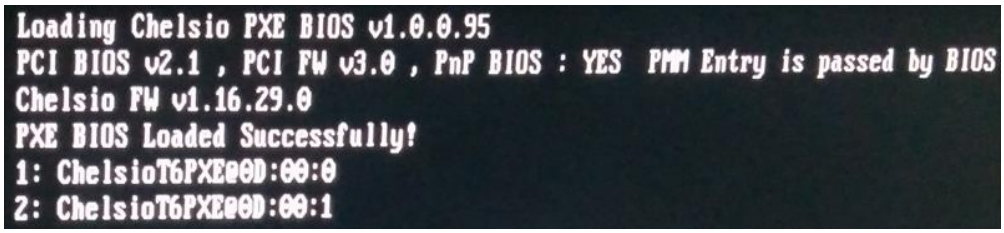

xiii. In the system setup, choose any of the Chelsio PXE devices as the first boot device.

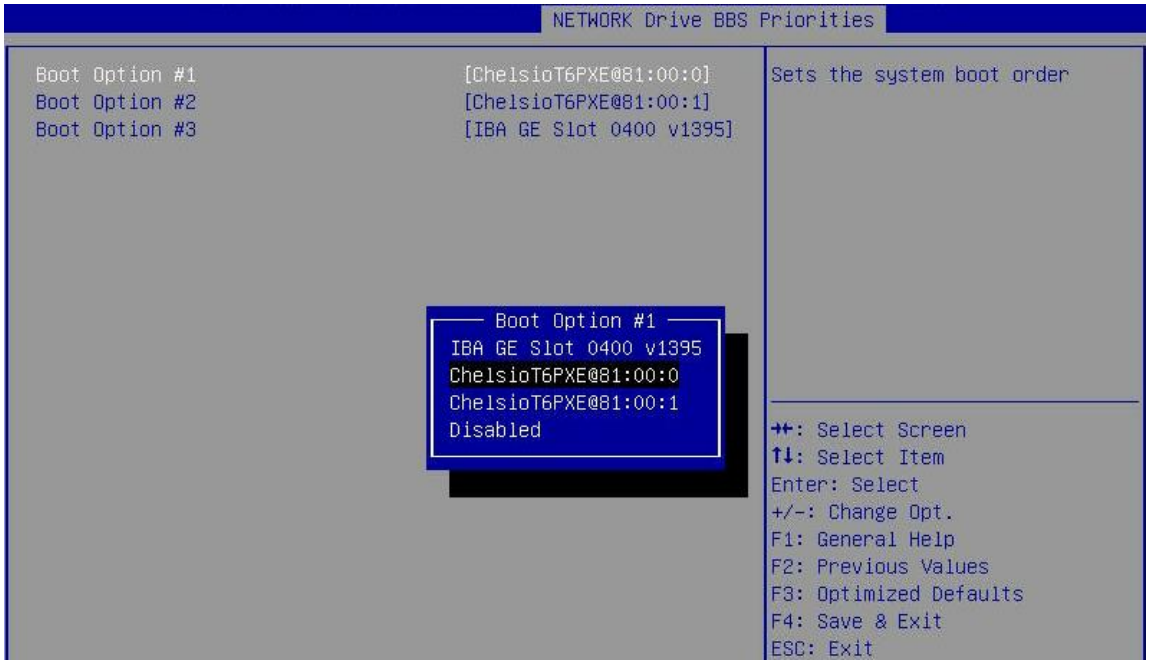

xiv. Reboot. DO NOT PRESS ALT-S to skip Chelsio option ROM, during POST.

xv. Hit [F12] key when prompted to start PXE boot.

5.2. uEFI PXE Boot

Important

- *Only uEFI v2.3.1, v2.4 and v2.5 supported.*
- *Any other uEFI version is NOT SUPPORTED and may render your system unusable.*

This section describes the method to configure and use Chelsio uEFI PXE interfaces.

i. Reboot the system and go into the BIOS setup.
ii. Chelsio HII should be listed as **Chelsio T5/T6**. Highlight it and press [Enter].

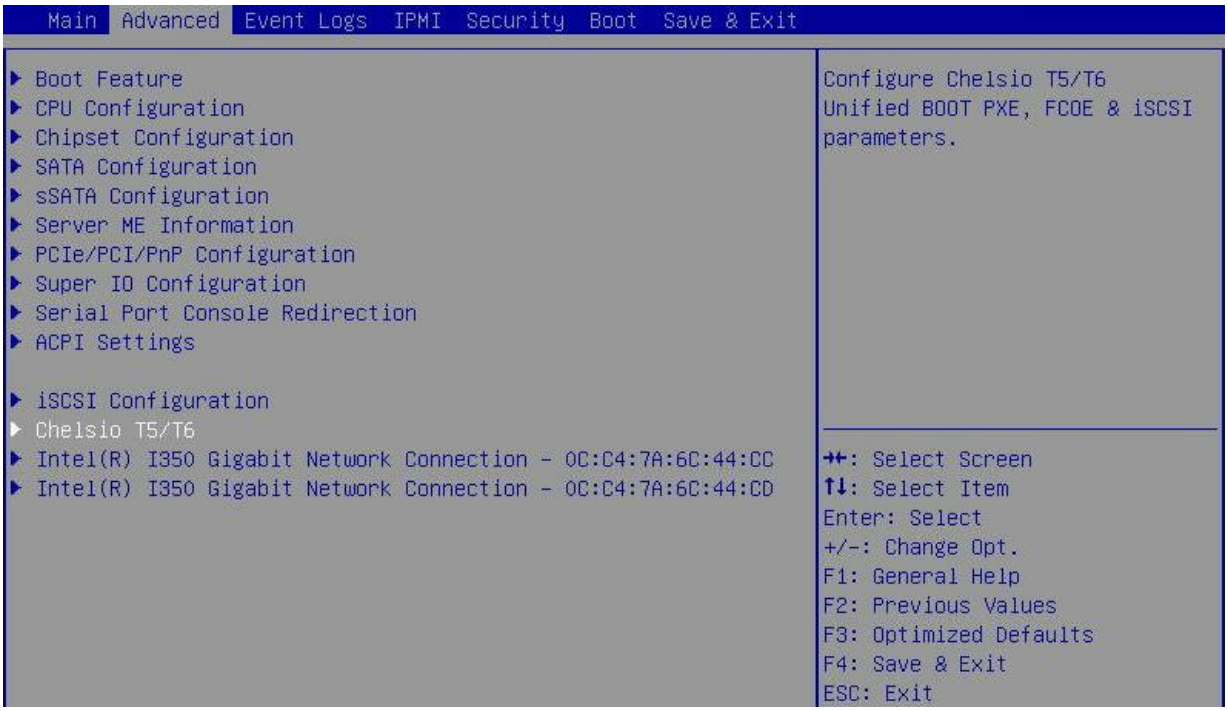

Note *If Chelsio T5/T6 is not listed, please ensure that Chelsio uEFI driver is loaded correctly as mentioned here in the **Flashing Firmware and Option ROM** section.*

iii. Select the Chelsio adapter to be configured and press [Enter].

iv. Select **Configuration Utility** and press [Enter].

v. Enable adapter BIOS if not already enabled.

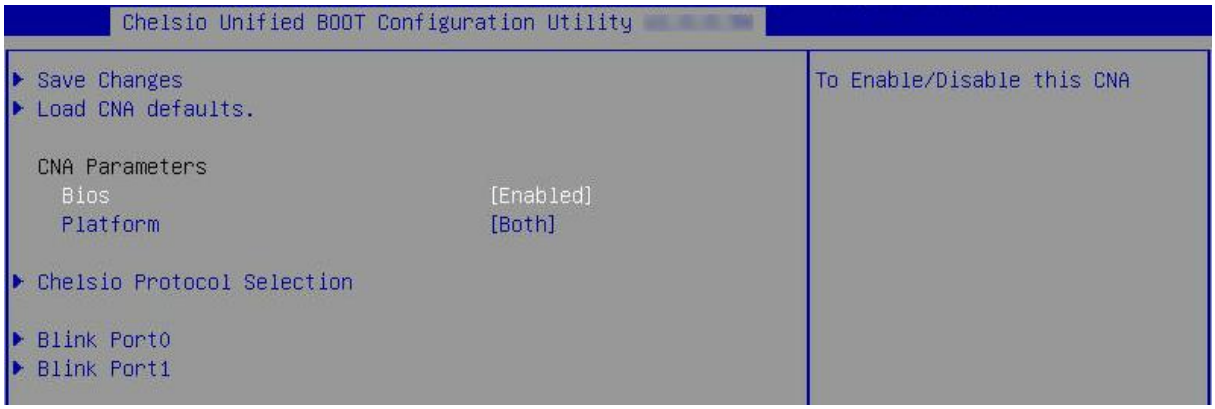

Note *It is highly recommended that you use the **Save Changes** option every time a parameter/option is changed.*

vi. Select **Chelsio Protocol Selection** and press [Enter].

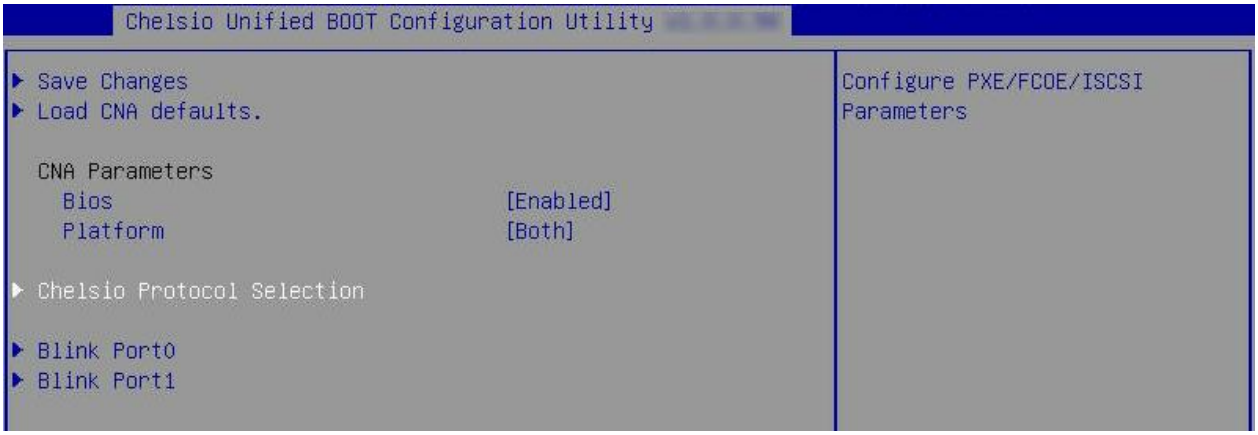

vii. Select **PXE** and press [Enter].

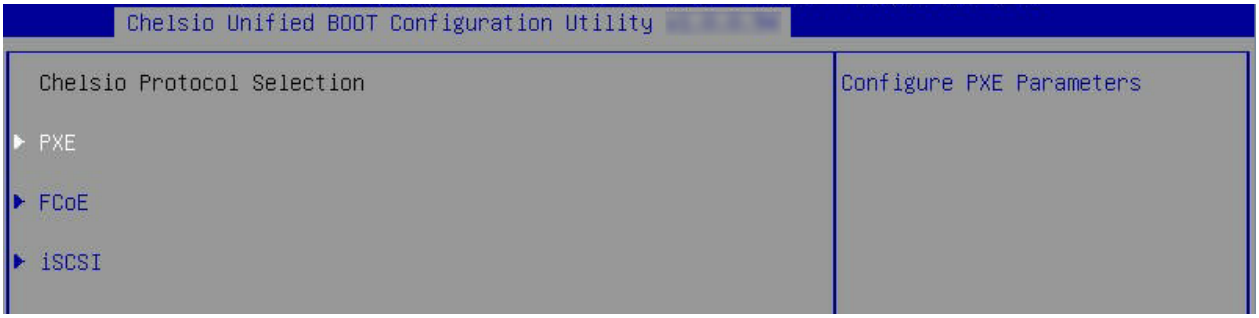

viii. Choose the boot port to try PXE boot. It is recommended to enable only those functions and ports which are going to be used. Please note that enabling PXE Function 0 will enable port 0 for PXE, enabling PXE Function 1 will enable port 1 and so on, for NIC function.

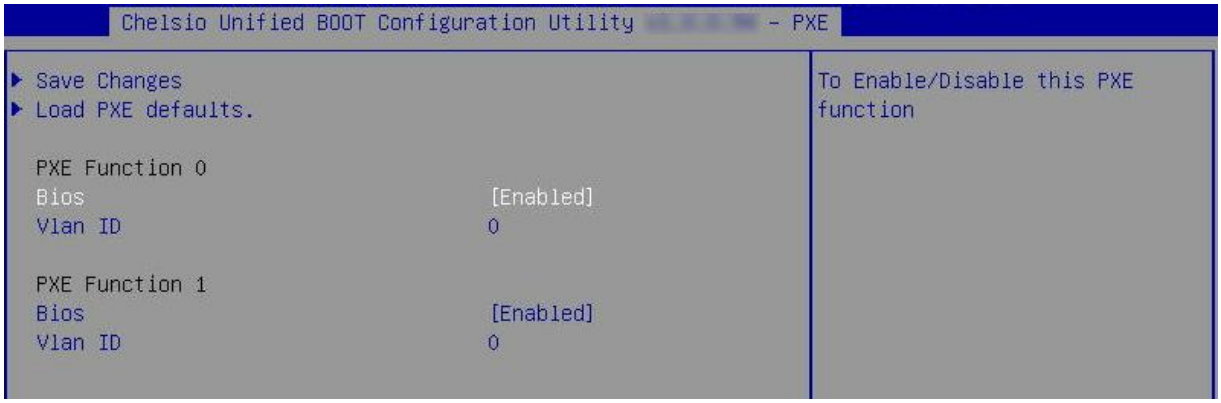

ix. Select **Save Changes** and press [Enter].

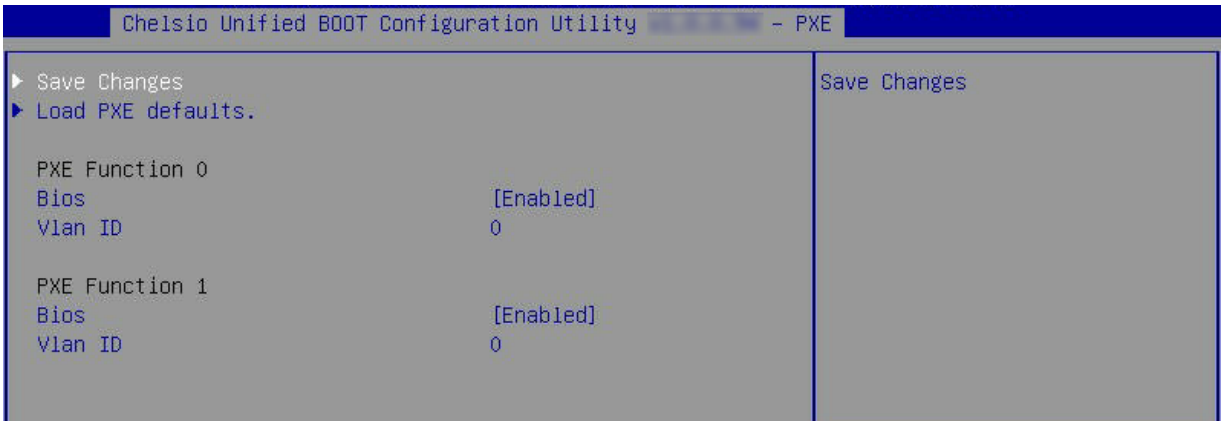

x. Reboot the system and in BIOS, choose any of the available Chelsio PXE devices for PXE boot.

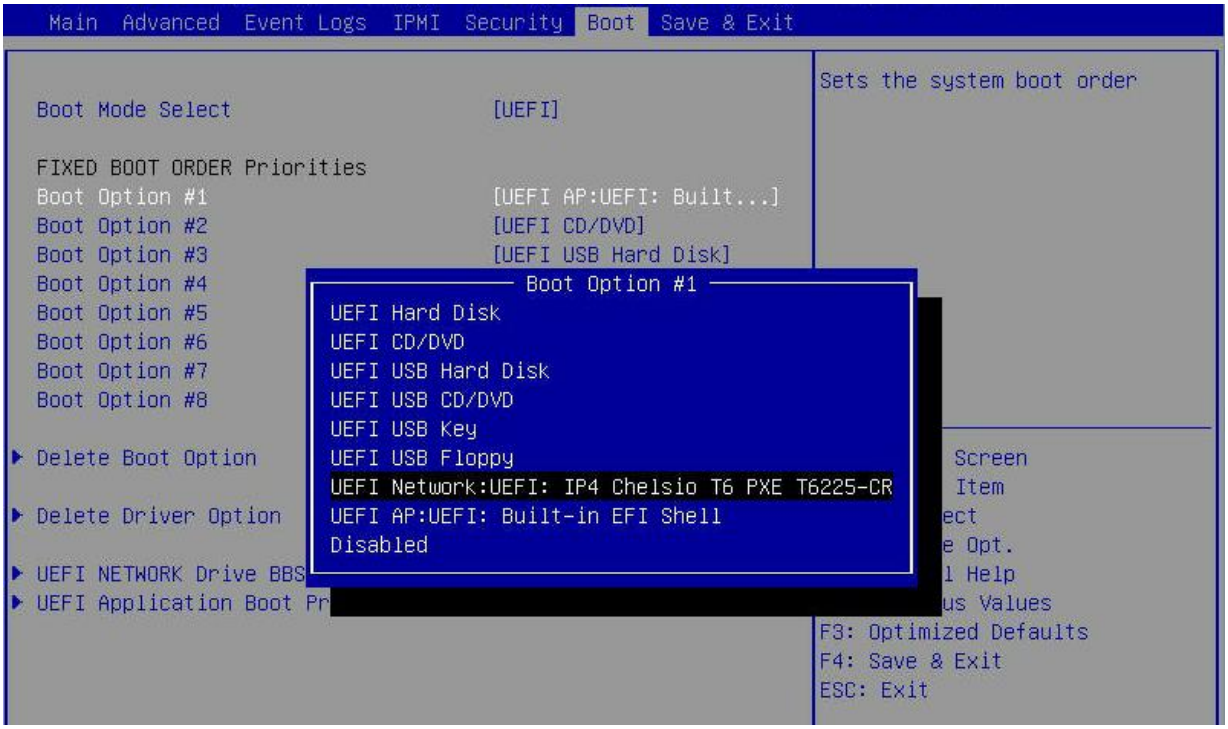

xi. Reboot and hit [F12] key when prompted to start PXE boot.

6. FCoE boot process

Before proceeding, please ensure that the Chelsio CNA has been flashed with the provided firmware and option ROM (See Flashing firmware and option ROM).

6.1. Legacy FCoE boot

i. Reboot the system.

ii. Press [Alt+C] when the message to configure Chelsio adapters appears on the screen.

iii. The configuration utility will appear as below:

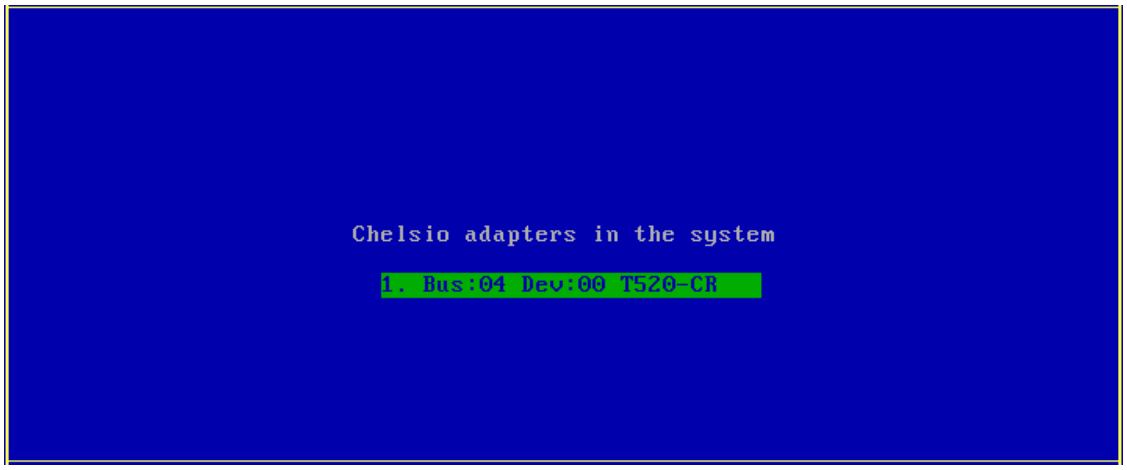

iv. Choose the CNA on which you flashed the option ROM image. Hit [Enter].

v. Enable the adapter BIOS if not already enabled. Hit [ENTER].

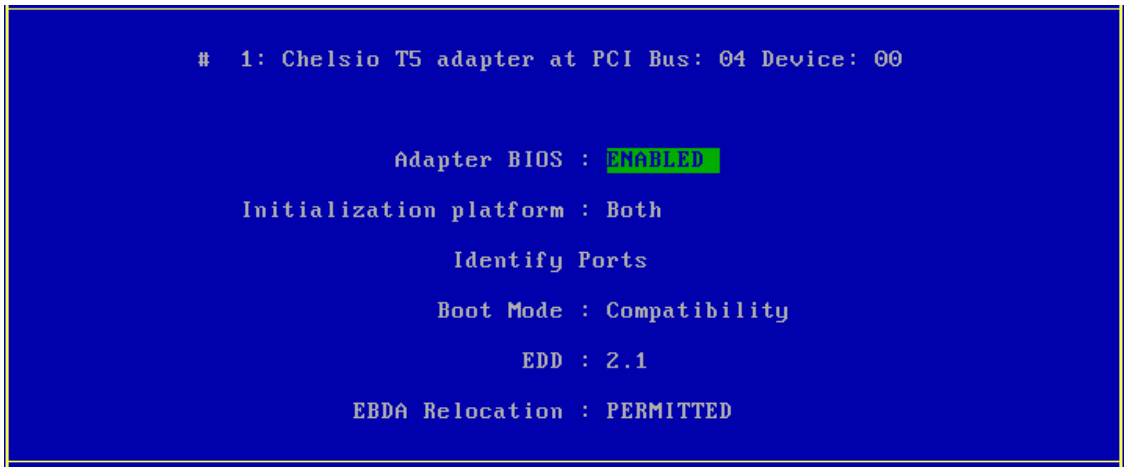

Note *Use the default values for Boot Mode, EDD and EBDA Relocation parameters, unless instructed otherwise.*

vi. Choose FCoE from the list to configure and hit [Enter].

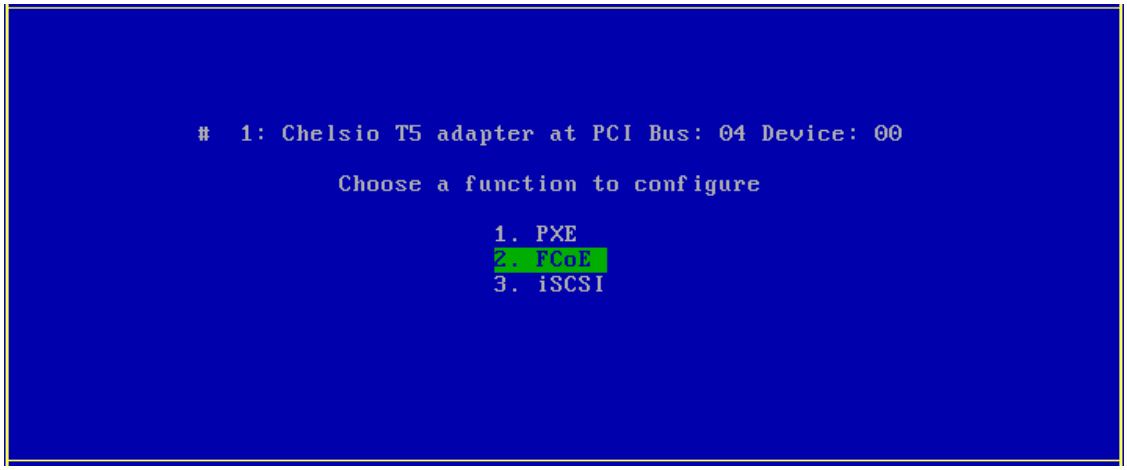

vii. Choose the first option, **Configure function parameters**, from the list of parameter type and hit [Enter].

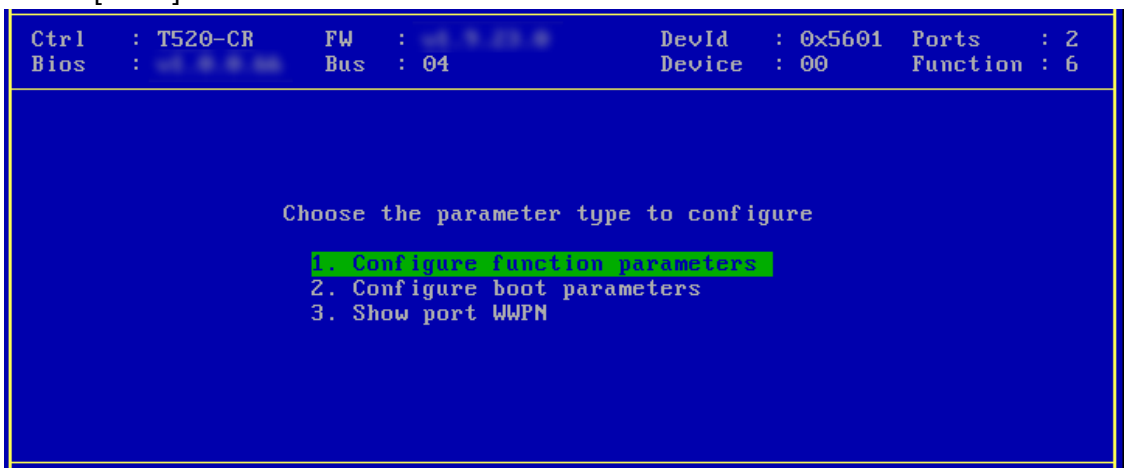

viii. Enable FCoE BIOS if not already enabled.

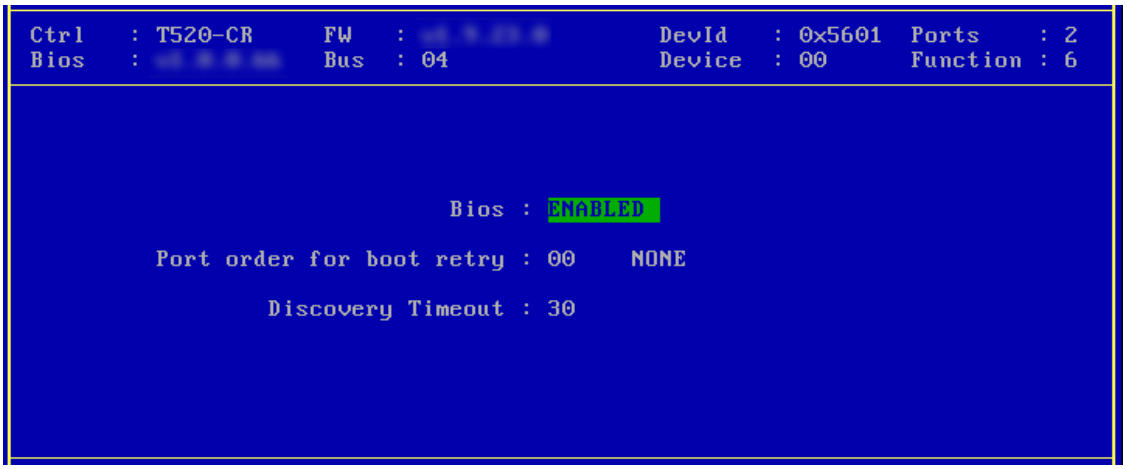

ix. Choose the order of the ports to discover FCoE targets.

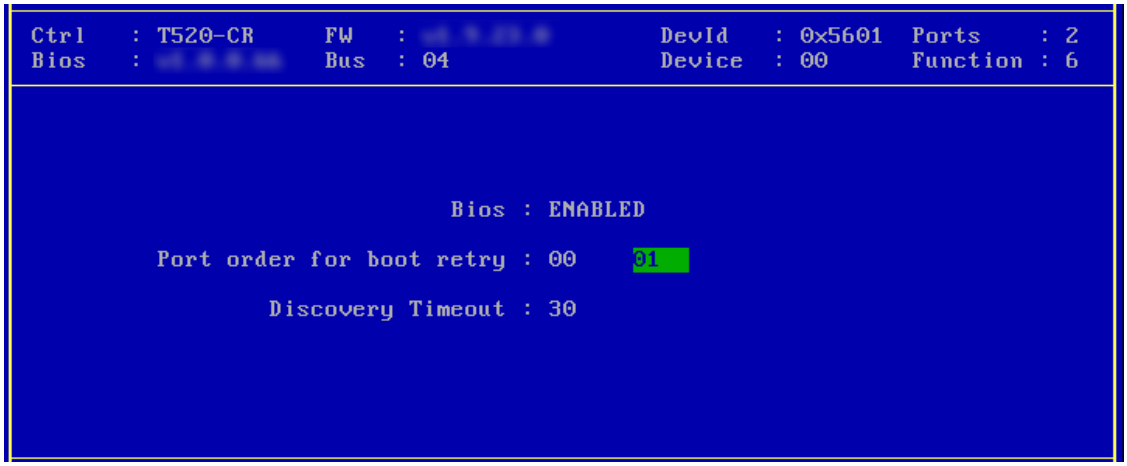

x. Set discovery timeout to a suitable value. Recommended value is >= 30.

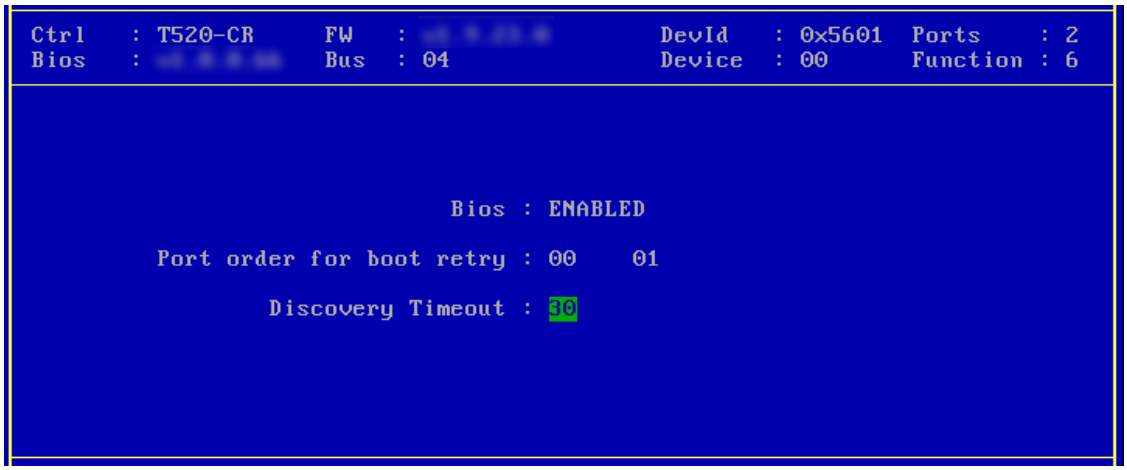

xi. Hit [F10] or [Esc] and then [Y] to save the configuration.

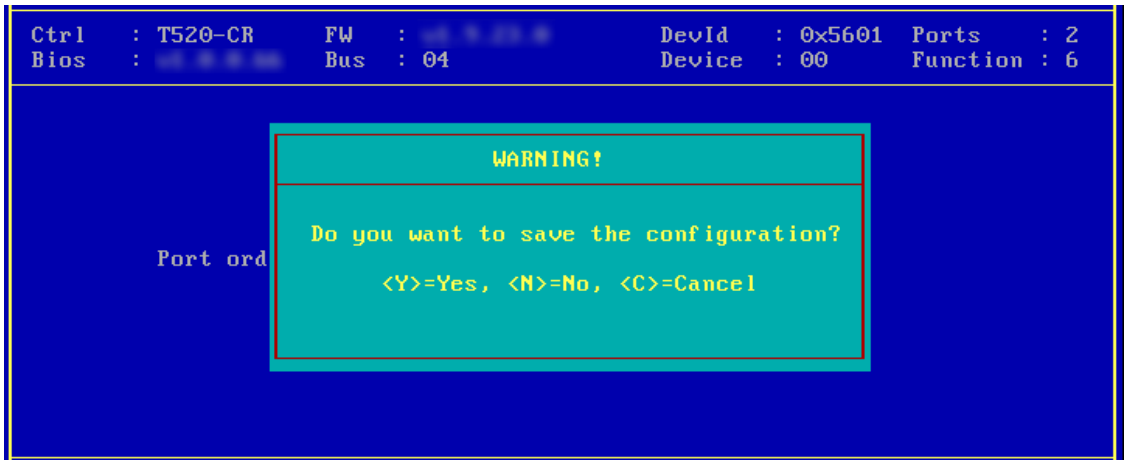

xii. Choose **Configure boot parameters**.

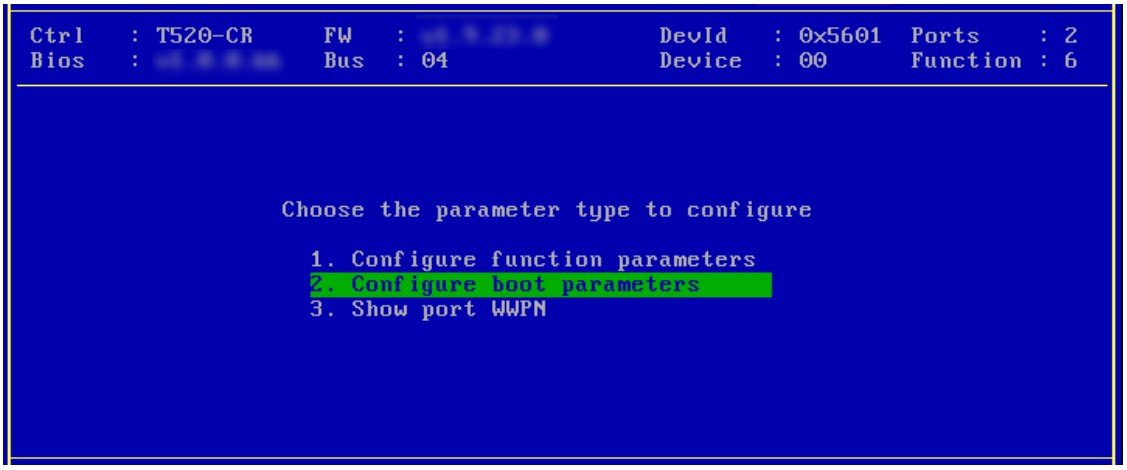

xiii. Select the first boot device and hit [Enter] to discover FC/FCoE targets connected to the switch. Wait till all reachable targets are discovered.

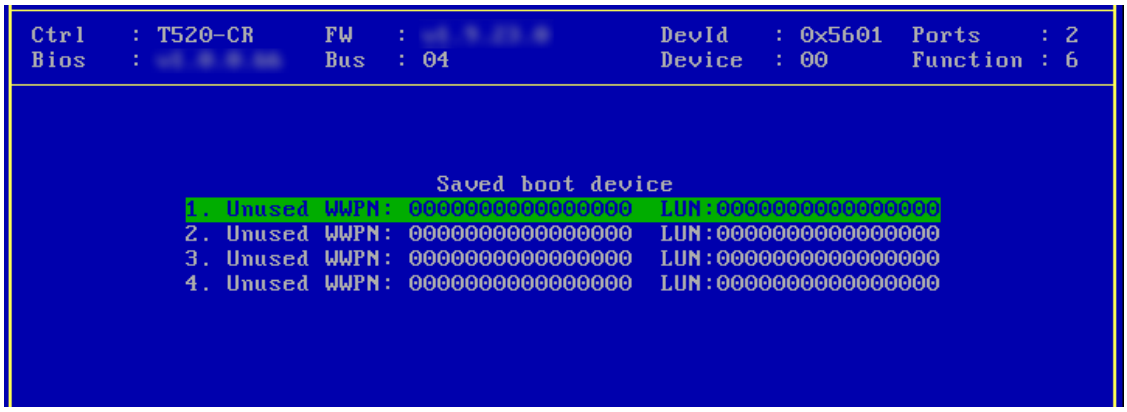

xiv. List of discovered targets will be displayed. Highlight a target using the arrow keys and hit [Enter] to select.

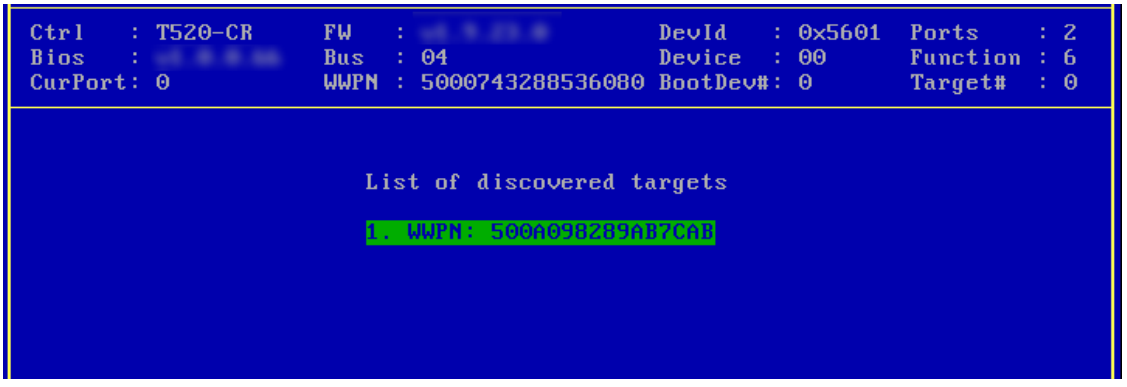

xv. From the list of LUNs displayed for the selected target, choose one on which operating system has to be installed. Hit [Enter].

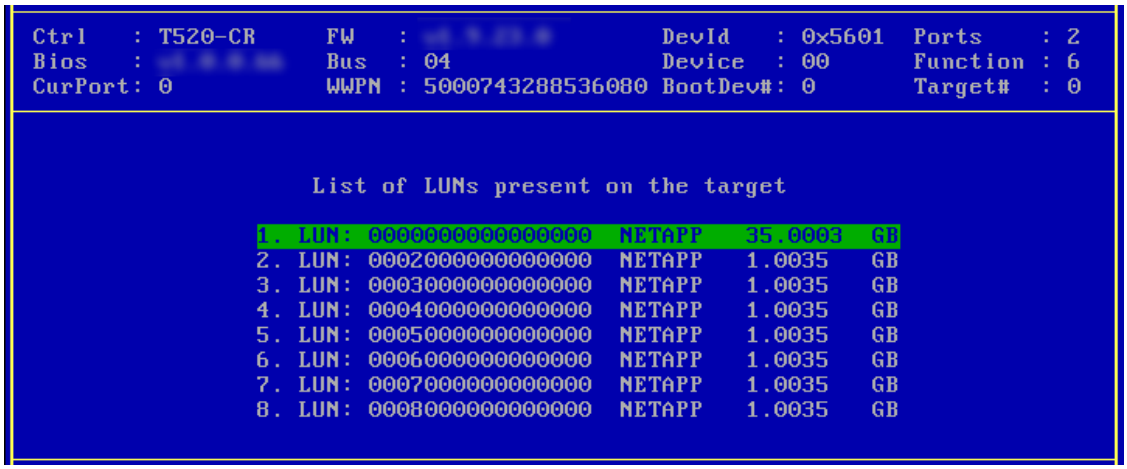

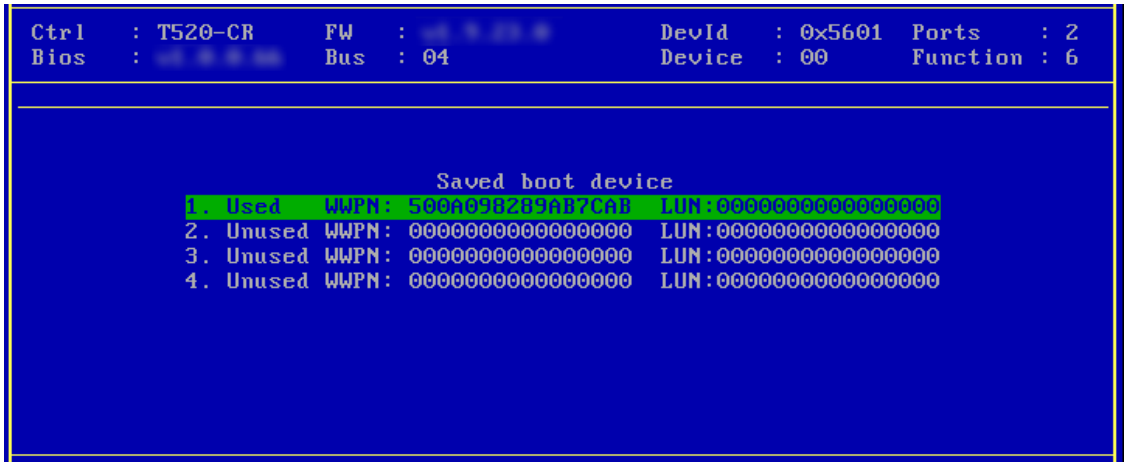

xvi. Hit [F10] or [Esc] and then [Y] to save the configuration.

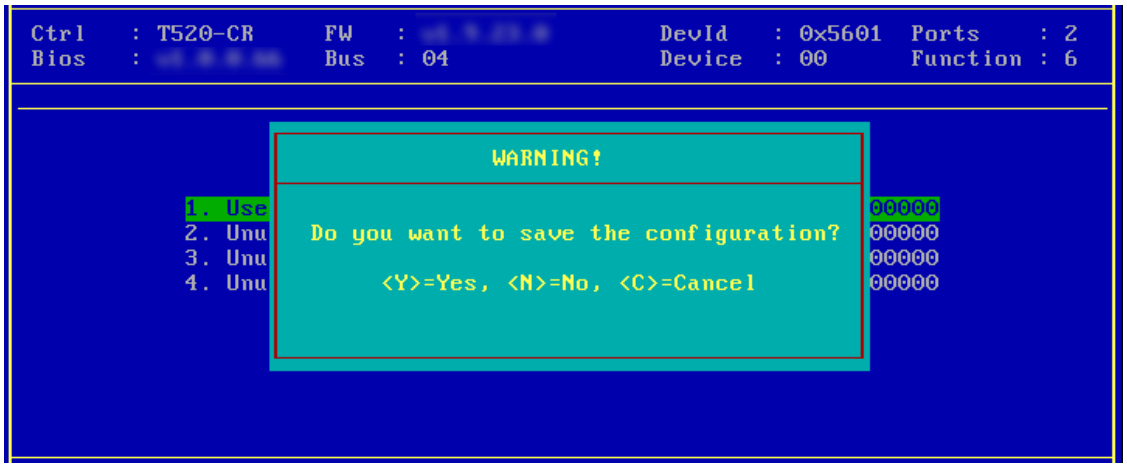

xvii. Reboot the machine.

xviii. During POST, allow the Chelsio option ROM to discover FCoE targets.

```
Installing Chelsio T5 Storage FCoE BIOS
PCI BIOSv3.0 PCI FWv2.1 PnP BIOS: YES PMM Entry is passed by BIOS
Bringing up link on PCI:04:00:6 Port 0 ... Done
Discovering FCoE Target(s) on PCI:04:00:6 Port 0 ... Done
sd(1): T520-CR         PCI:04:00:6 P(0) WWPN:500A098289AB7CAB Lun(00)
       NETAPP  LUN              8030 35.0003 GB
Storage FCoE BIOS Installed Successfully!
```

xix. Enter BIOS setup and choose FCoE disk discovered via Chelsio adapter as the first boot device.

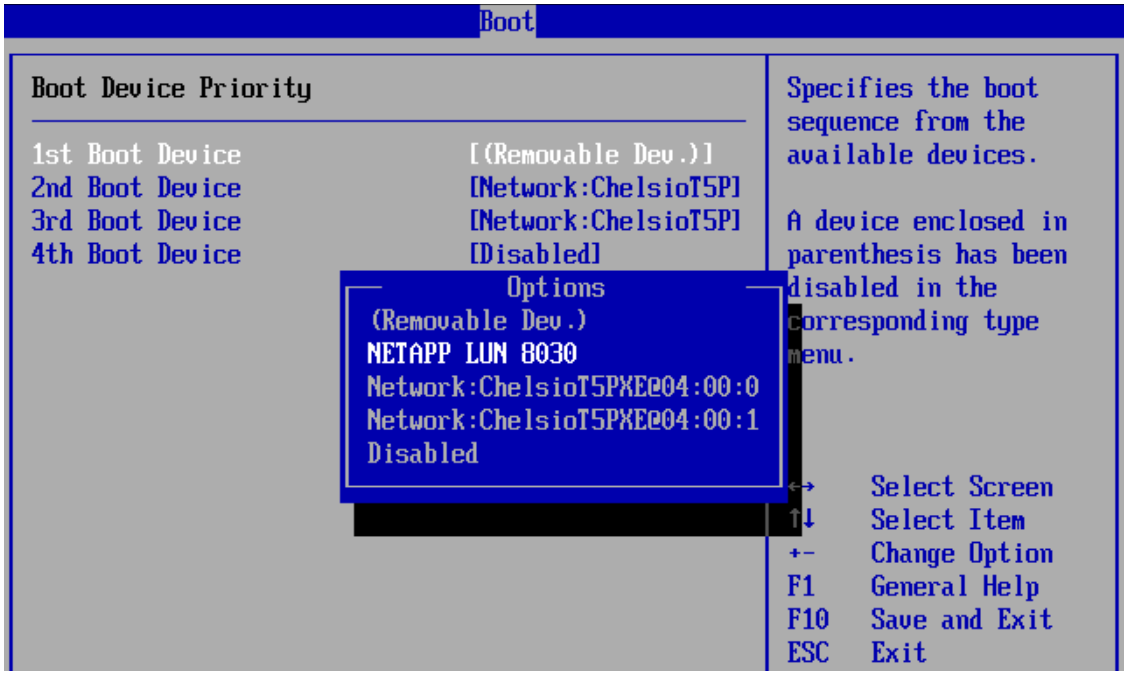

xx. Reboot and boot from the FCoE disk or install the required OS using PXE.

6.2. uEFI FCoE Boot

Important

- *Only uEFI v2.3.1, v2.4 and v2.5 supported.*
- *Any other uEFI version is NOT SUPPORTED and may render your system unusable.*

i. Reboot the system and go into BIOS setup.

ii. Select **Chelsio T5/T6** and press [Enter]

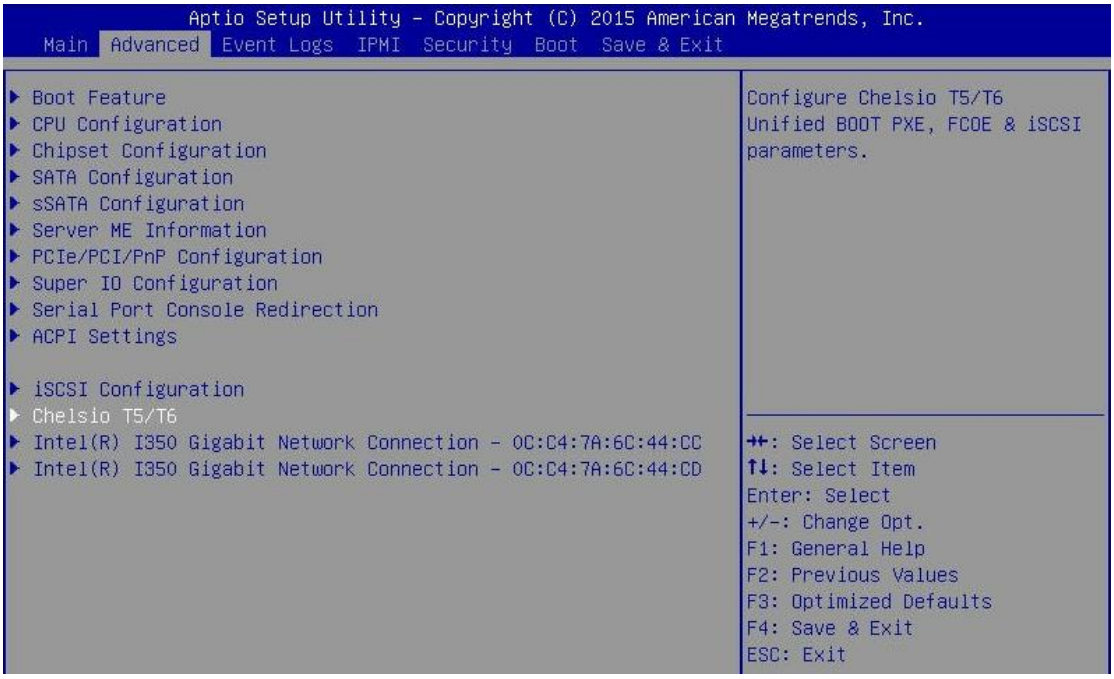

Note *If Chelsio T5/T6 is not listed, please ensure that Chelsio uEFI driver is loaded correctly as mentioned here in the **Flashing Firmware and Option ROM** section.*

iii. Select the Chelsio adapter to be configured and press [Enter].

iv. Select **Configuration Utility** and press [Enter].

v. Enable adapter BIOS if not already enabled.

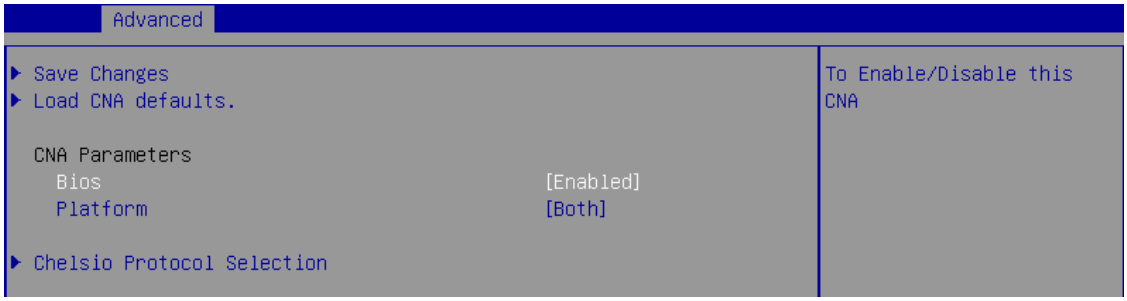

Note *It is highly recommended that you use the **Save Changes** option every time a parameter/option is changed.*

vi. Select **Chelsio Protocol Selection** and press [Enter].

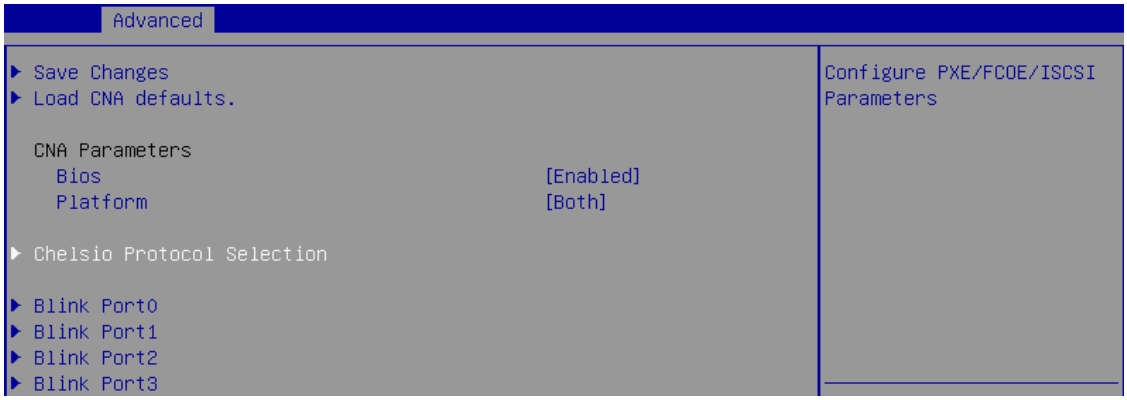

vii. Select **FCoE** and press [Enter].

viii. Under **Function Parameters**, enable FCoE BIOS, if not already enabled.

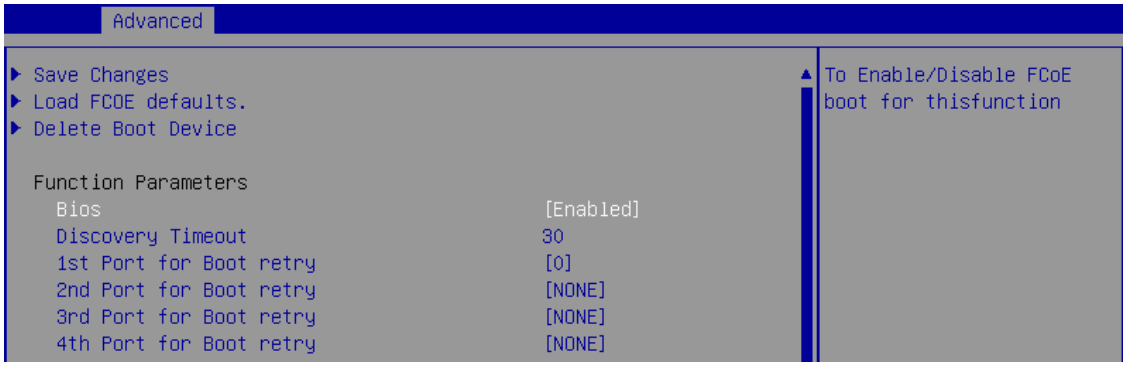

ix. Set discovery timeout to a suitable value. Recommended value is >= 30

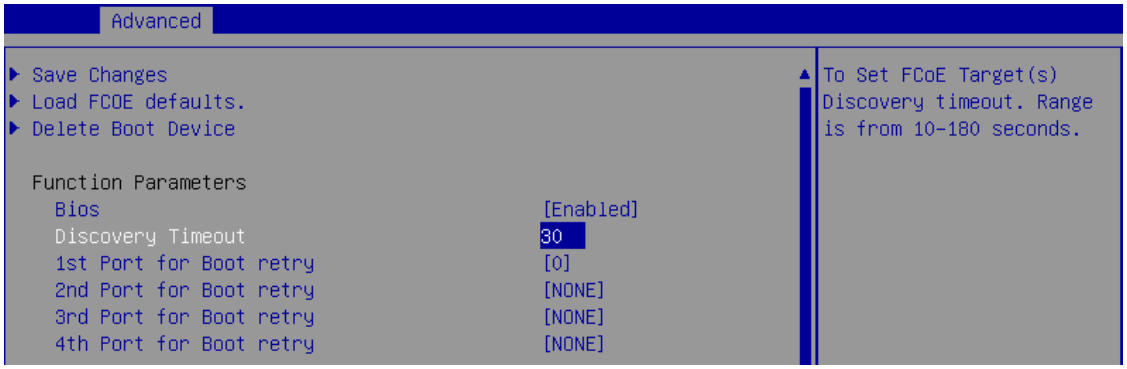

x. Choose the order of the ports to discover FCoE targets.

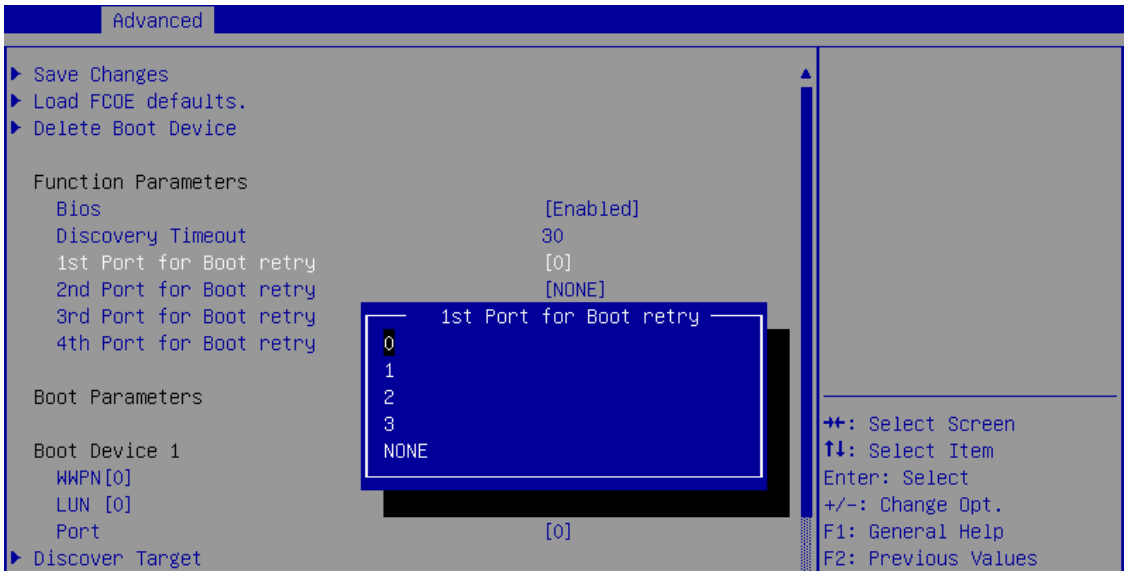

xi. Under the first boot device, select **Discover Target** and press [Enter] to discover FC/FCoE targets connected to the switch. Wait till all reachable targets are discovered.

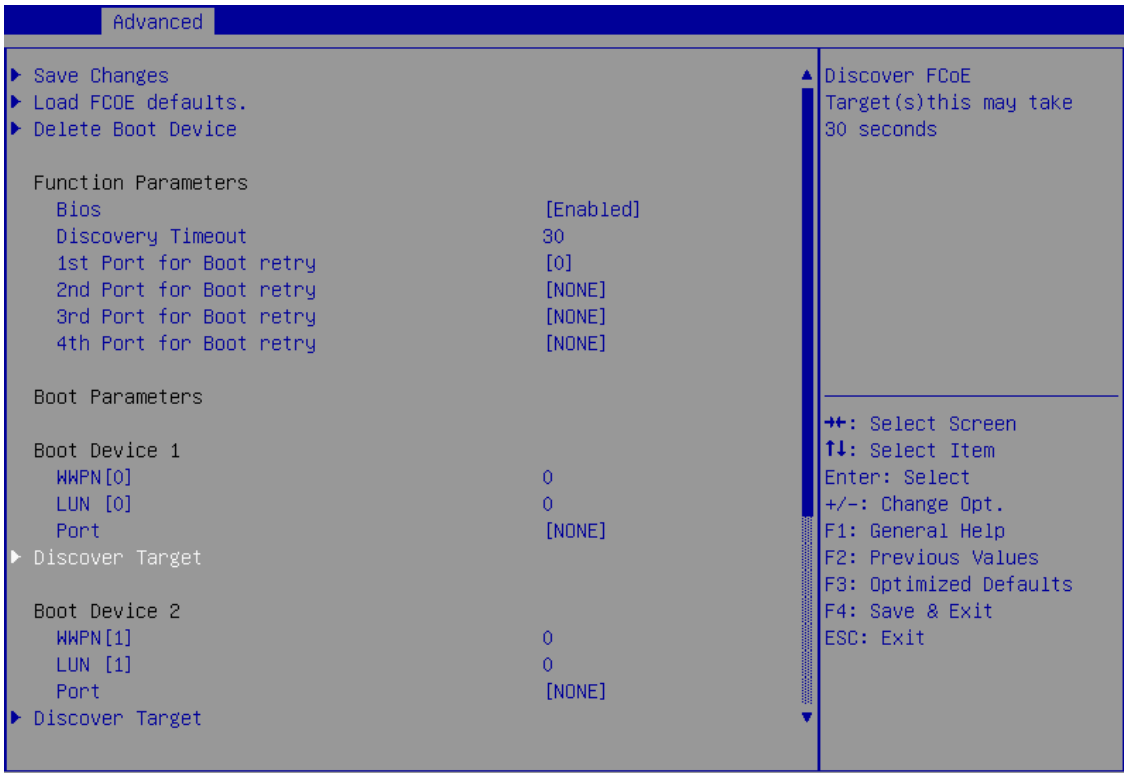

xii. List of discovered targets will be displayed. Highlight a target to select it and hit [Enter].

xiii. List of LUNs for the selected target will be displayed. Highlight a LUN to select it and hit [Enter].

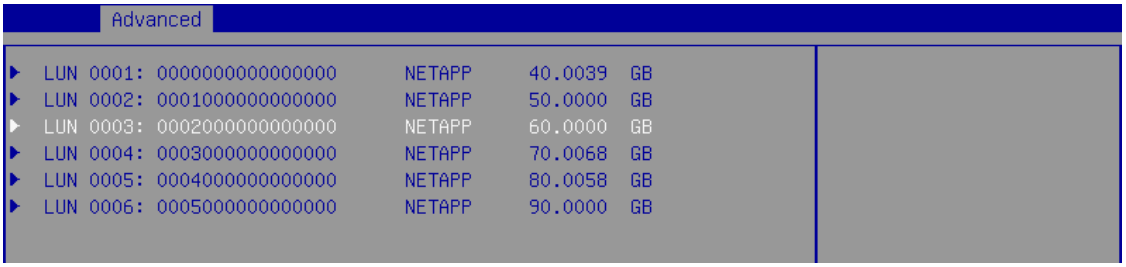

xiv. Select **Save Changes** and press [Enter].

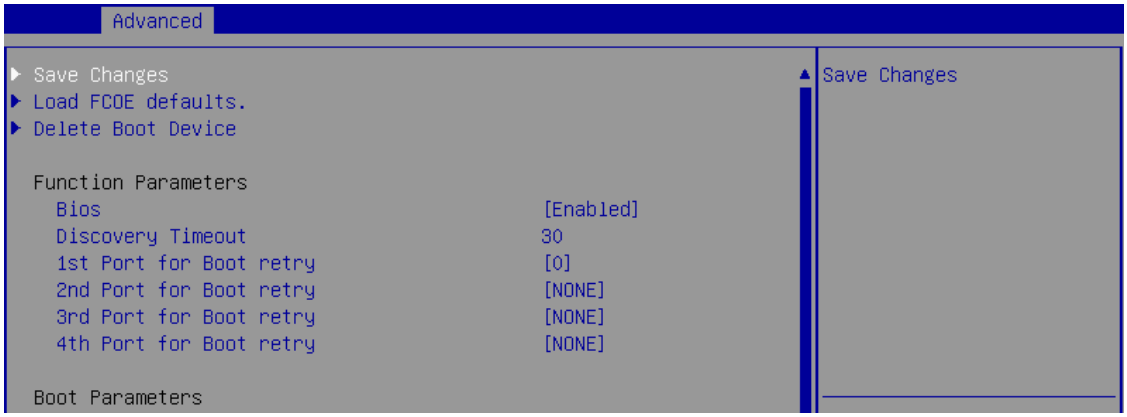

xv. Reboot the system for changes to take effect.

xvi. The discovered LUN should appear in the **Boot Configuration** section and system BIOS section.

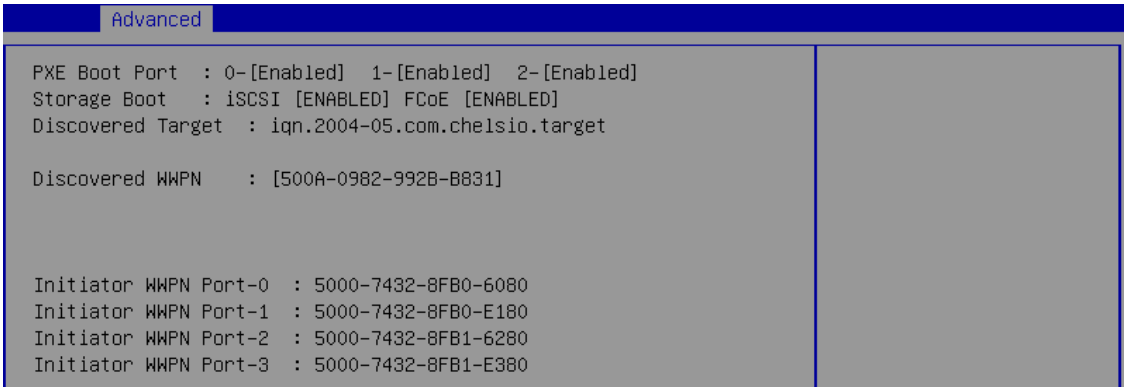

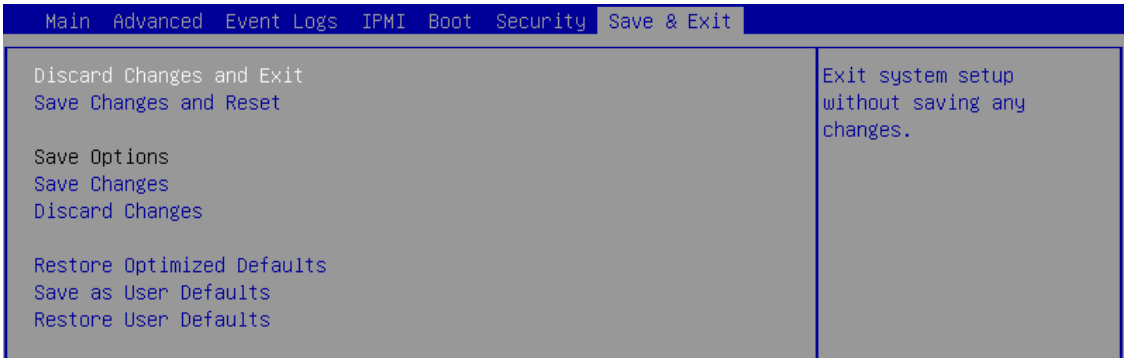

xvii. Select the LUN as the first boot device and exit from BIOS.

xviii. Either boot from the LUN or install the required OS.

7. iSCSI boot process

Before proceeding, please ensure that the Chelsio CNA has been flashed with the provided firmware and option ROM (See Flashing firmware and option ROM).

7.1. Legacy iSCSI boot

i. Reboot the system.

ii. Press [Alt+C] when the message to configure Chelsio adapters appears on the screen.

iii. The configuration utility will appear as below:

iv. Choose the CNA on which you flashed the option ROM image. Hit [Enter].

v. Enable the adapter BIOS if not already enabled. Hit [Enter].

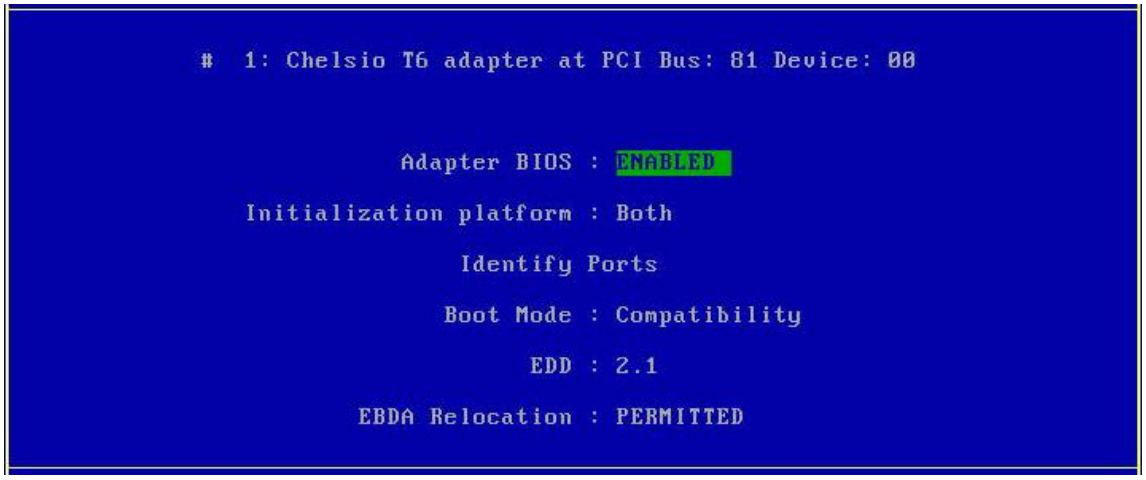

Note *Use the default values for Boot Mode, EDD and EBDA Relocation parameters, unless instructed otherwise.*

vi. Choose *iSCSI* from the list to configure and hit [Enter].

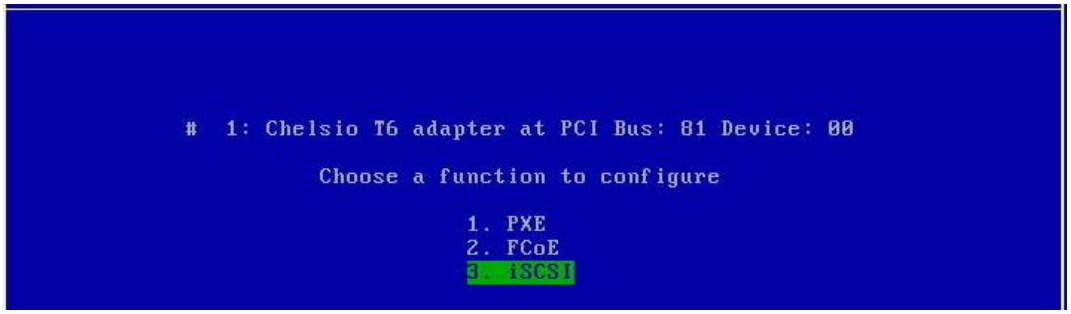

vii. Choose the first option, **Configure Function Parameters**, from the list of parameter type and hit [Enter].

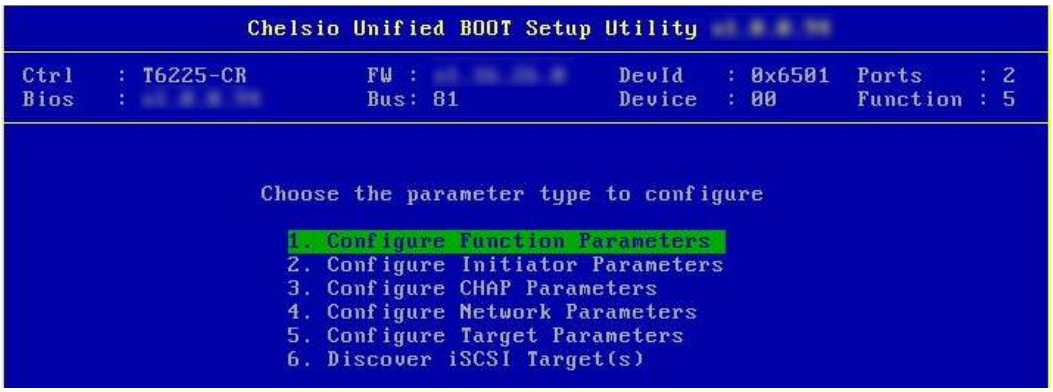

viii. Enable iSCSI BIOS if not already enabled. iBFT (iSCSI Boot Firmware Table) will be selected by default. You can also configure the number of iSCSI login attempts (retries) in case the network is unreachable or slow.

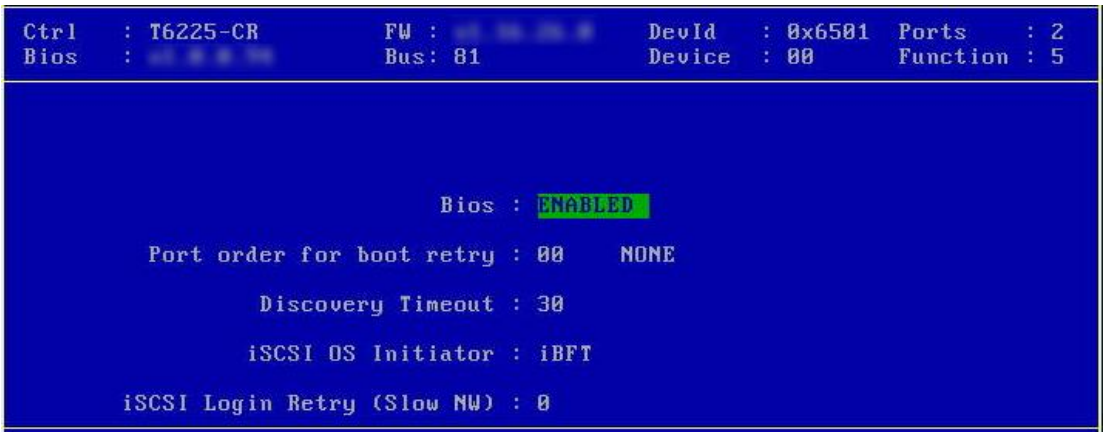

ix. Choose the order of the ports to discover iSCSI targets.

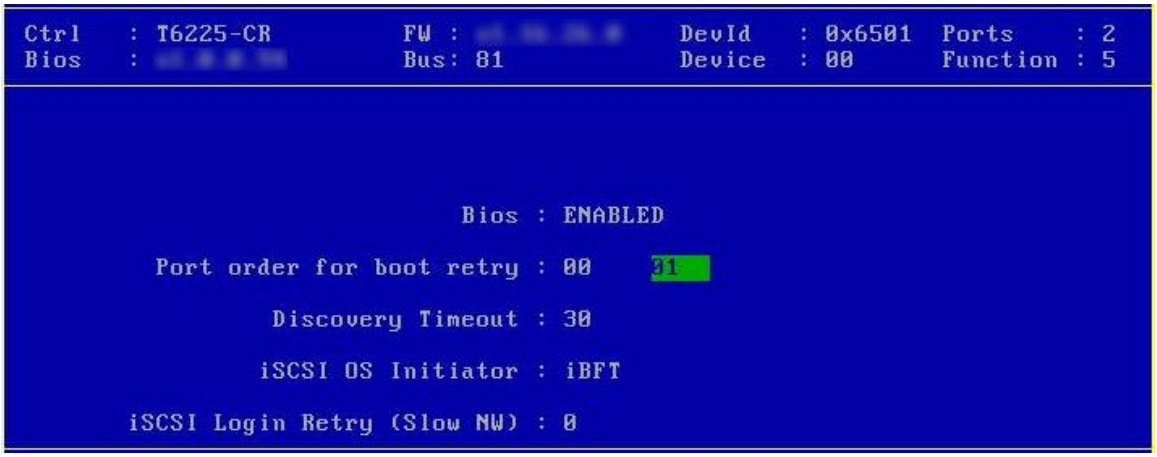

x. Set discovery timeout to a suitable value. Recommended value is >= 30.

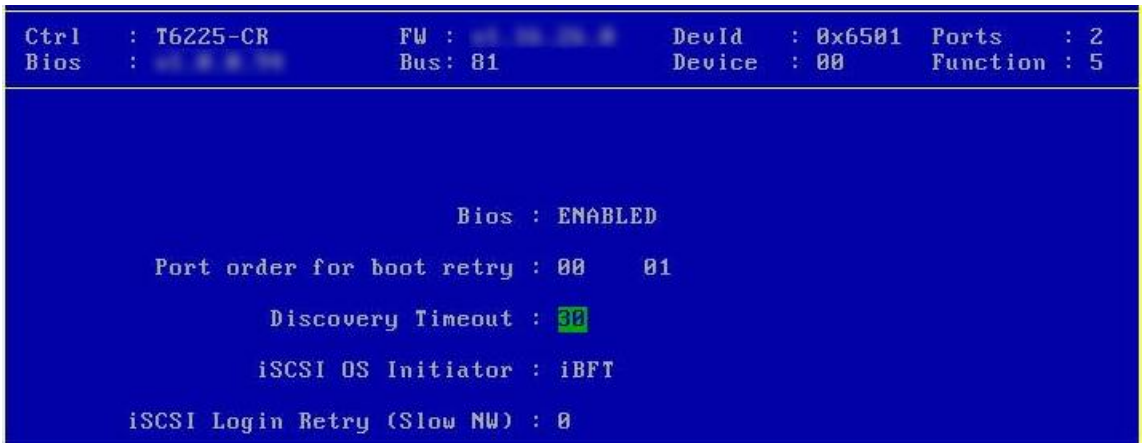

xi. Hit [Esc] and then [Y] to save the configuration.

xii. Go back and choose **Configure Initiator Parameters** to configure initiator related properties.

xiii. Initiator properties like IQN, Header Digest, Data Digest, etc. will be displayed. Change the values appropriately or continue with the default values. Hit [F10] to save.

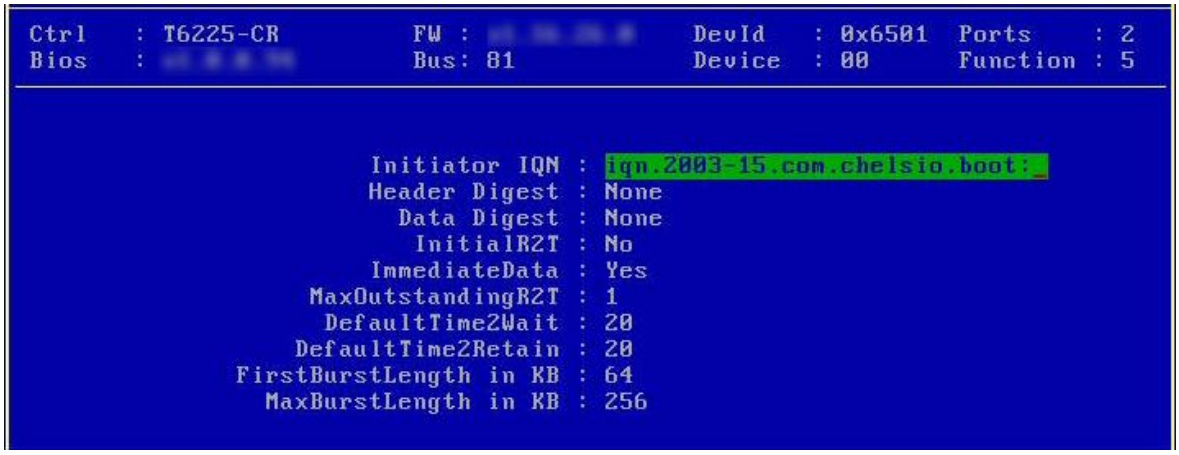

xiv. CHAP authentication is disabled by default. To enable and configure, go back and choose **Configure CHAP Parameters**

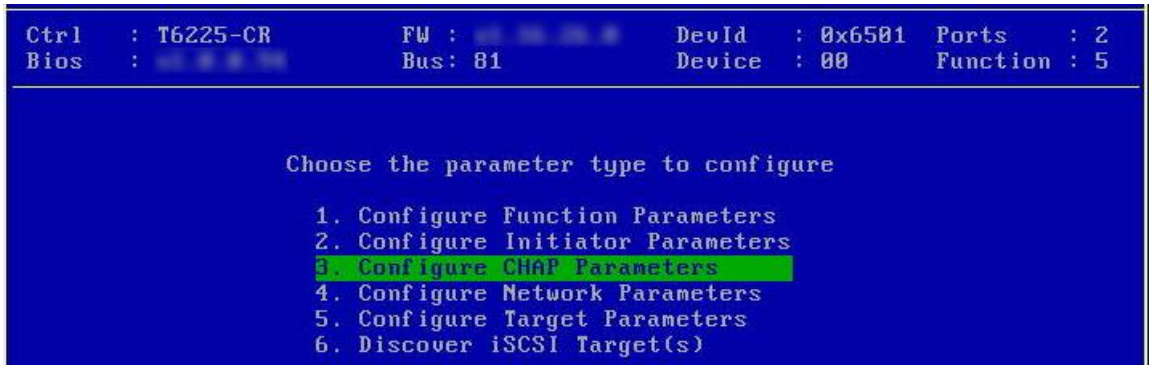

xv. Enable CHAP authentication by selecting ONE-WAY or MUTUAL in the **CHAP Policy** field. Next, choose the CHAP method. Finally, provide Initiator and Target CHAP credentials as per the authentication method selected. Hit [F10] to save.

xvi. Go back and choose **Configure Network Parameters** to configure iSCSI Network related properties.

xvii. Select the port using which you want to connect to the target. Hit [Enter].

xviii. Select *Yes* in the **Enable DHCP** field to configure port using DHCP or *No* to manually configure the port. Hit [F10] to save.

xix. Go back and choose **Configure Target Parameters** to configure iSCSI target related properties.

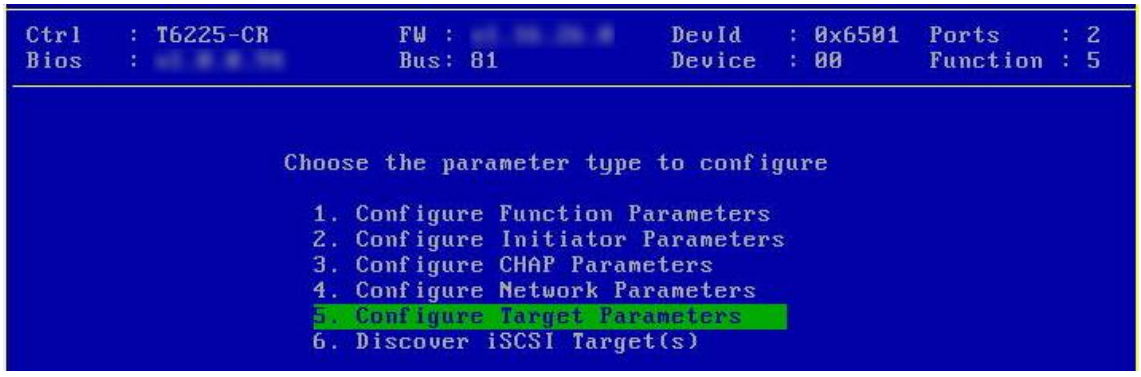

xx. If you want to discover target using DHCP, select *Yes* in the **Discover Boot Target via DHCP** field. To discover target via static IP, select *No* and provide the target IP and Hit [F10] to save. The default TCP port selected is 3260.

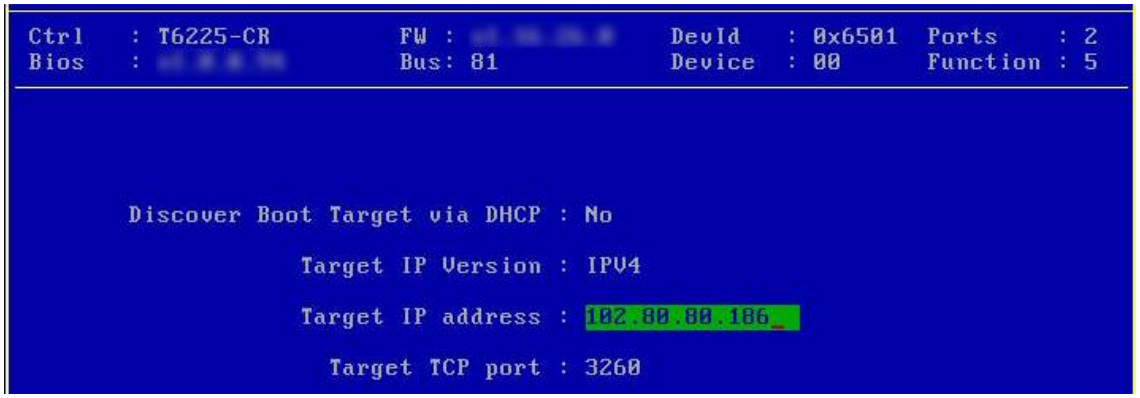

xxi. Go back and choose **Discover iSCSI Target (s)** to connect to a target.

xxii. Select the portal group on which iSCSI service is provided by the target.

xxiii. A list of available targets will be displayed. Select the target you wish to connect to and hit [Enter].

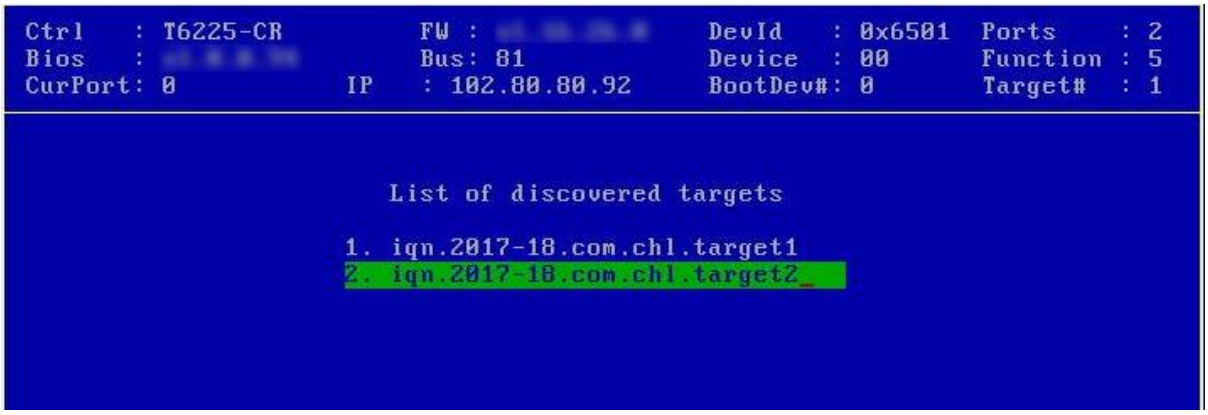

xxiv. A list of LUNs configured on the selected target will be displayed. Select the LUN you wish to connect to and hit [Enter].

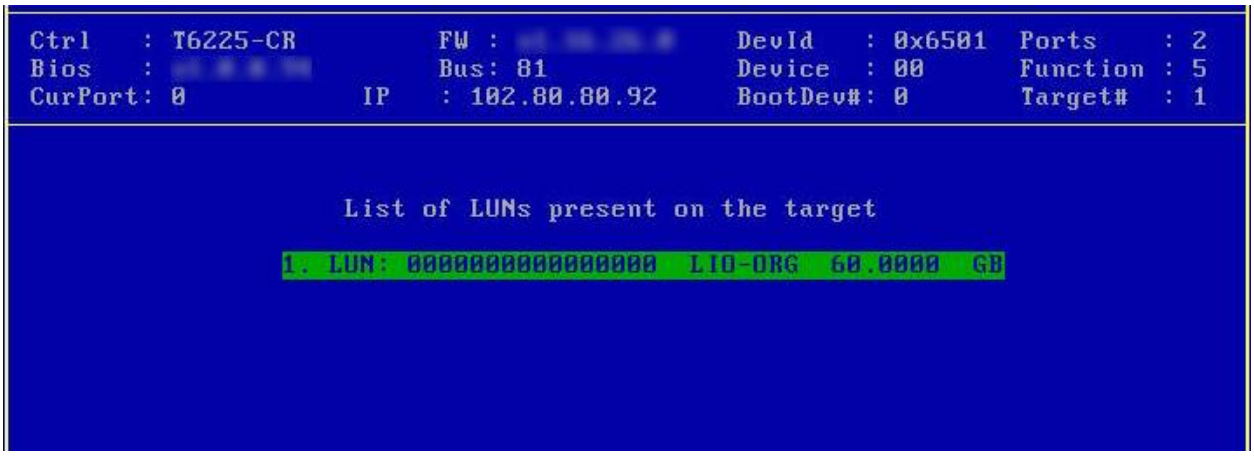

xxv. Hit [Esc] and then [Y] to save the configuration.

xxvi. Reboot the machine.

xxvii. During POST, allow the Chelsio option ROM to discover iSCSI targets.

xxviii. Enter BIOS setup and choose iSCSI target LUN discovered via Chelsio adapter as the first boot device.

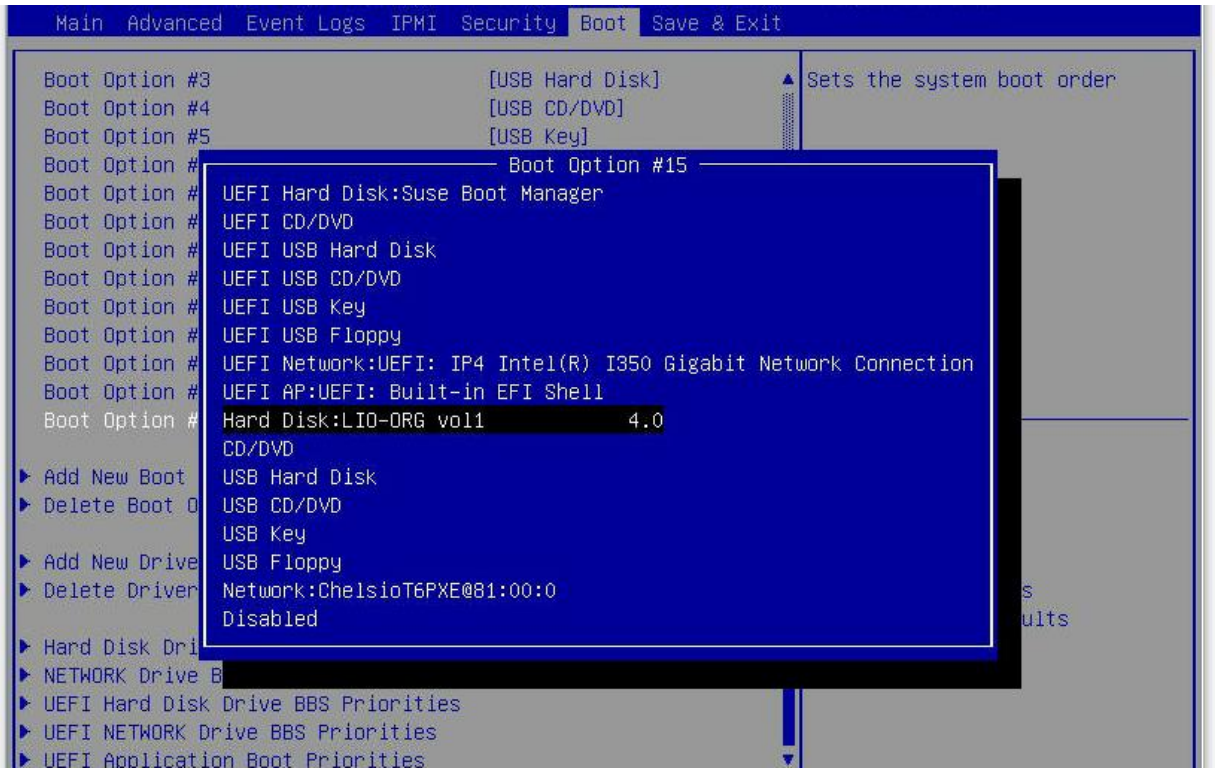

xxix. Reboot and boot from the iSCSI Target LUN or install the required OS using PXE.

7.2. uEFI iSCSI Boot

Important

- *Only uEFI v2.3.1, v2.4 and v2.5 supported.*
- *Any other uEFI version is NOT SUPPORTED and may render your system unusable.*

This section describes the method to perform iSCSI boot on uEFI platforms.

i. Reboot the system and go into BIOS setup.

ii. Select **Chelsio T5/T6** and press [Enter]

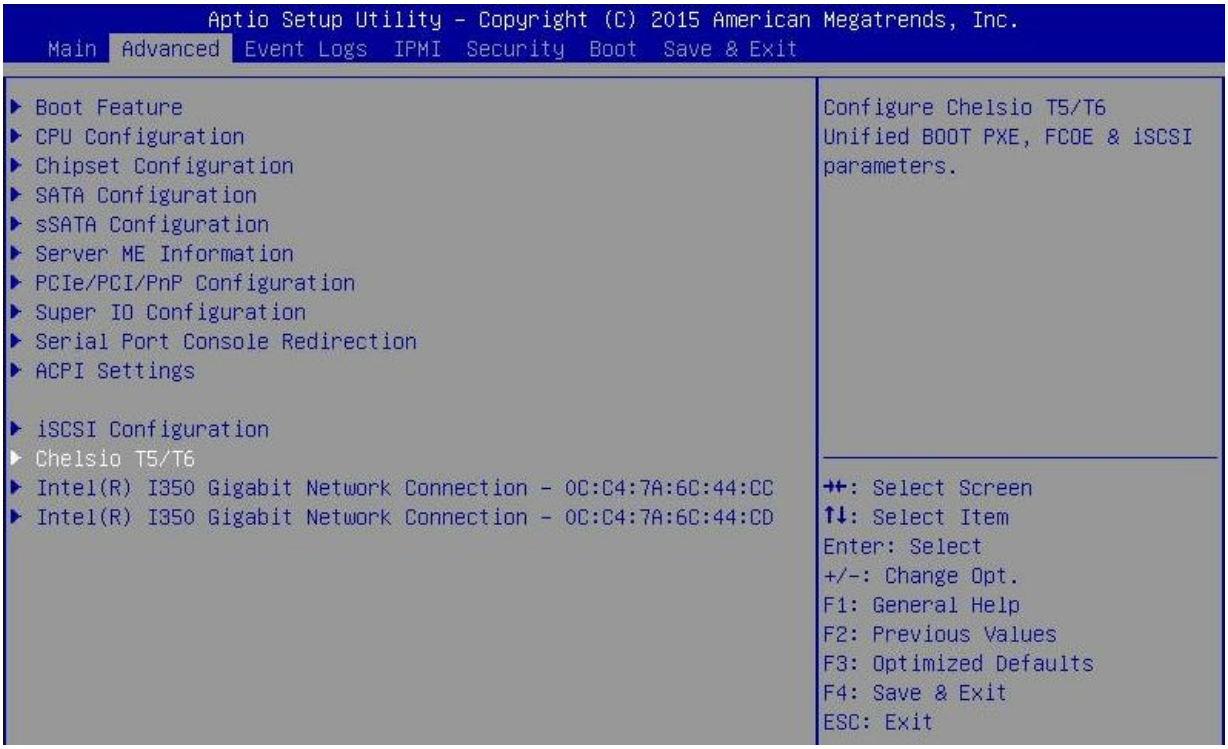

Note *If Chelsio T5/T6 is not listed, please ensure that Chelsio uEFI driver is loaded correctly as mentioned here in the **Flashing Firmware and Option ROM** section.*

iii. Select the Chelsio adapter to be configured and press [Enter].

iv. Select **Configuration Utility** and press [Enter].

v. Enable adapter BIOS if not already enabled.

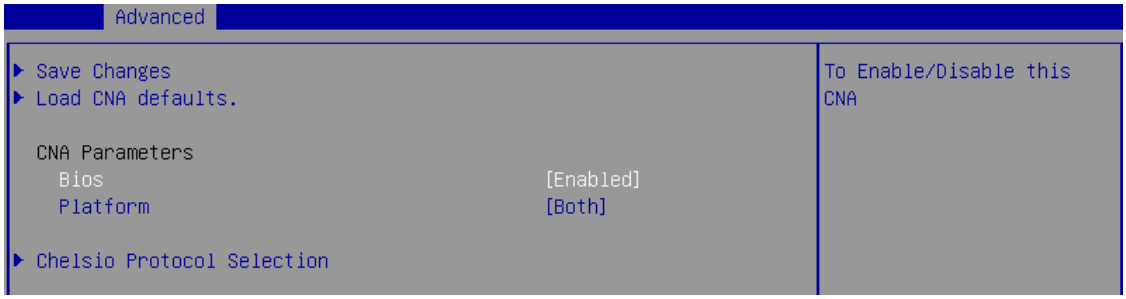

Note *It is highly recommended that you use the **Save Changes** option every time a parameter/option is changed.*

vi. Select **Chelsio Protocol Selection** and press [Enter].

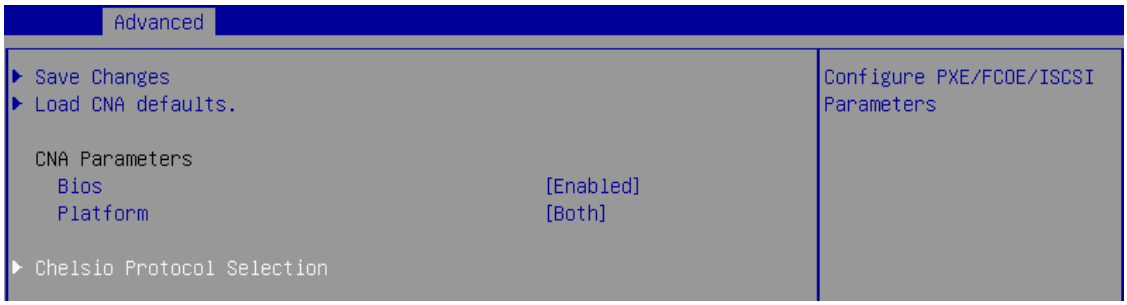

vii. Select **iSCSI** and press [Enter].

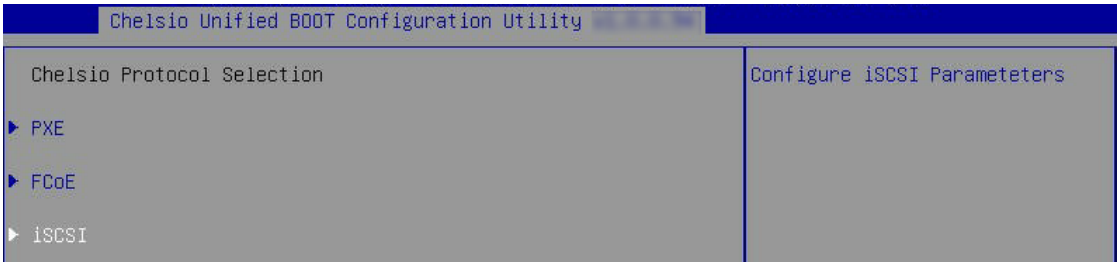

viii. Under **Function Parameters**, enable iSCSI BIOS, if not already enabled.

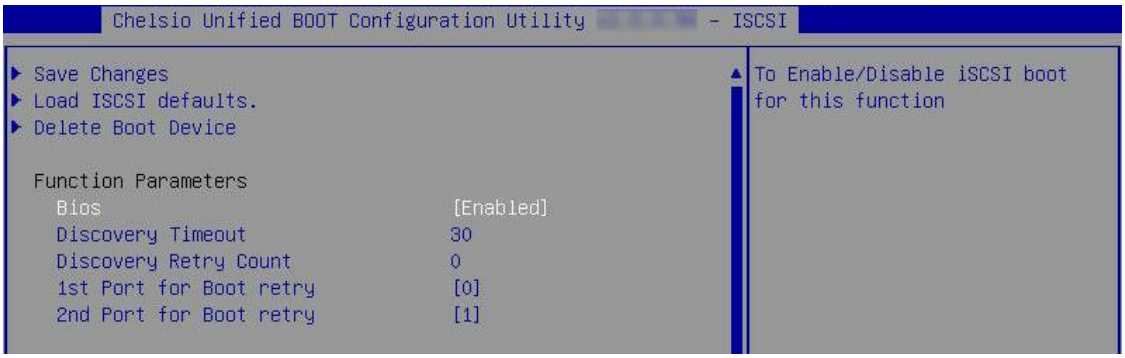

ix. Set discovery timeout to a suitable value. Recommended value is >= 30

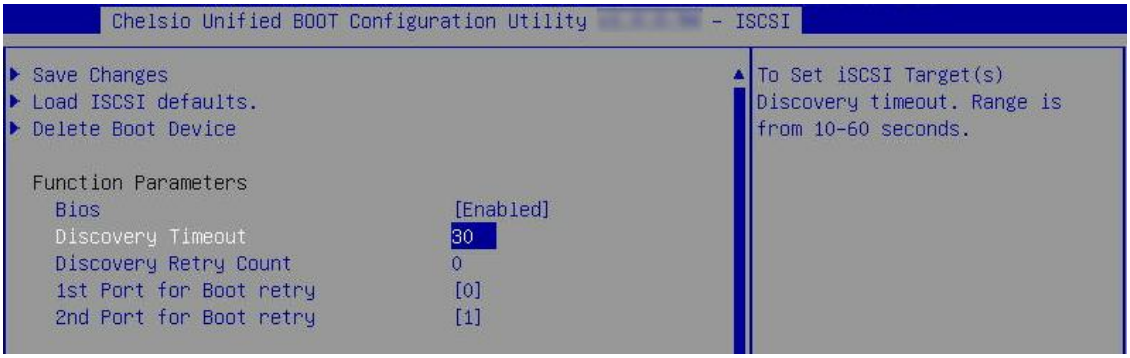

x. Choose the order of the ports to discover iSCSI targets.

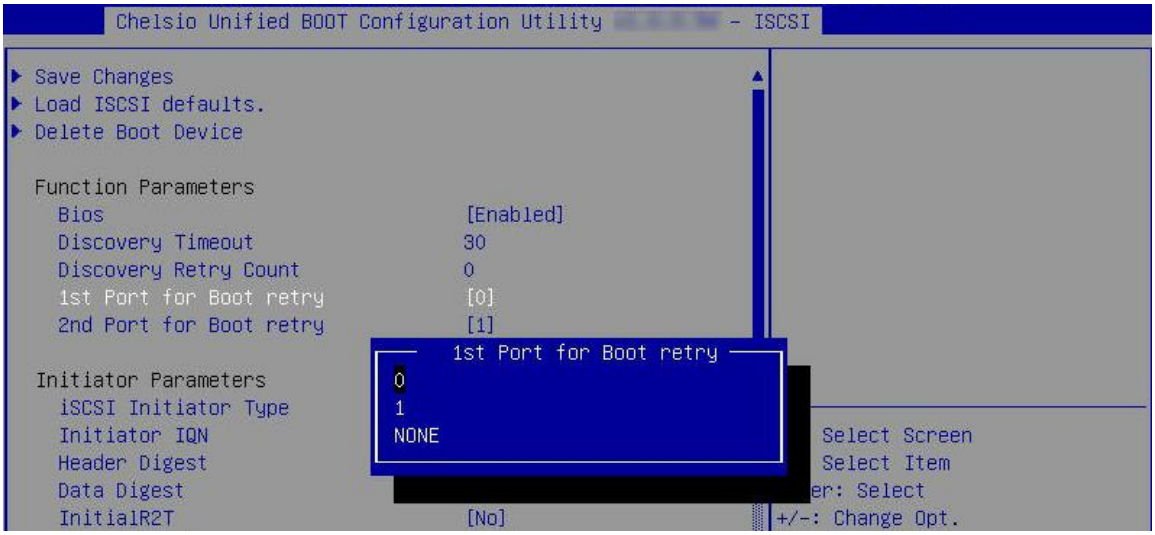

xi. Under **Initiator Parameters**, iSCSI Initiaitor properties like IQN, Header Digest, Data Digest, etc will be displayed. Change the values appropriately or continue with the default values.

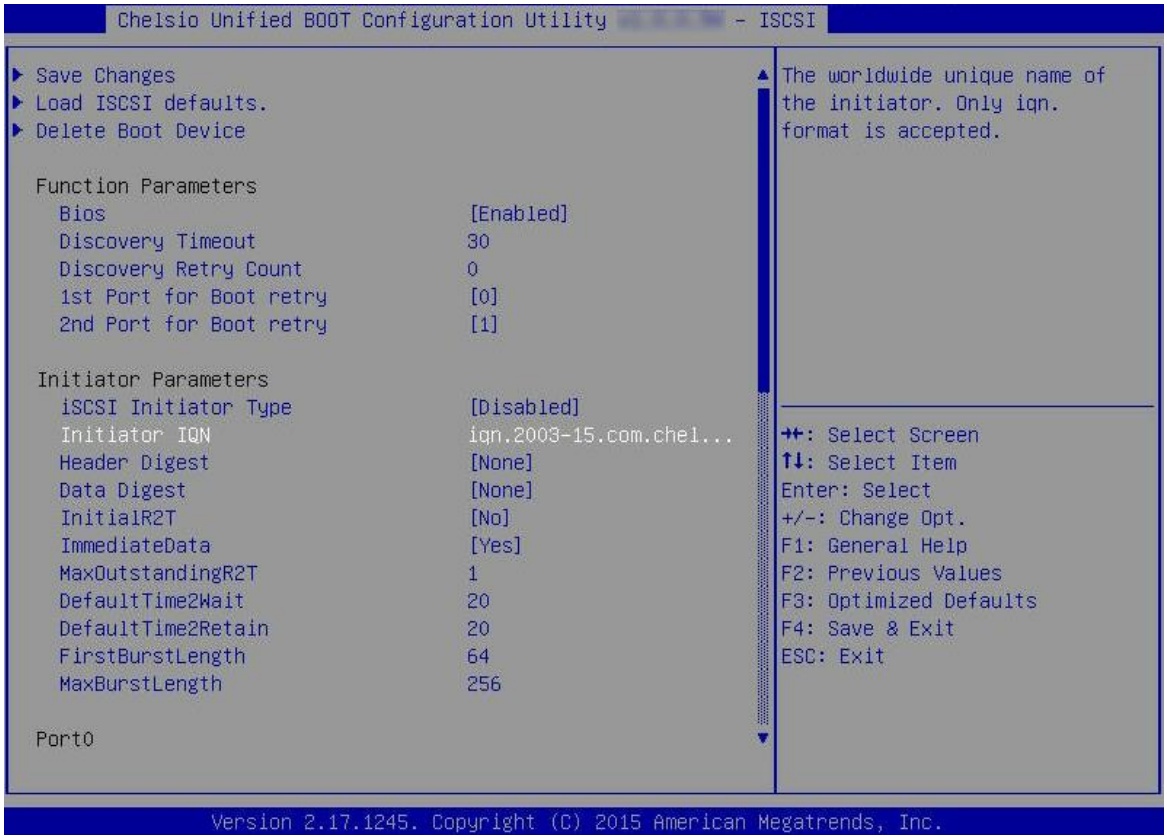

xii. Under the first port, select **Enable DHCP** field, hit [Enter] and select **Enabled**. This will configure port using DHCP. Select **Disabled** to manually configure the port.

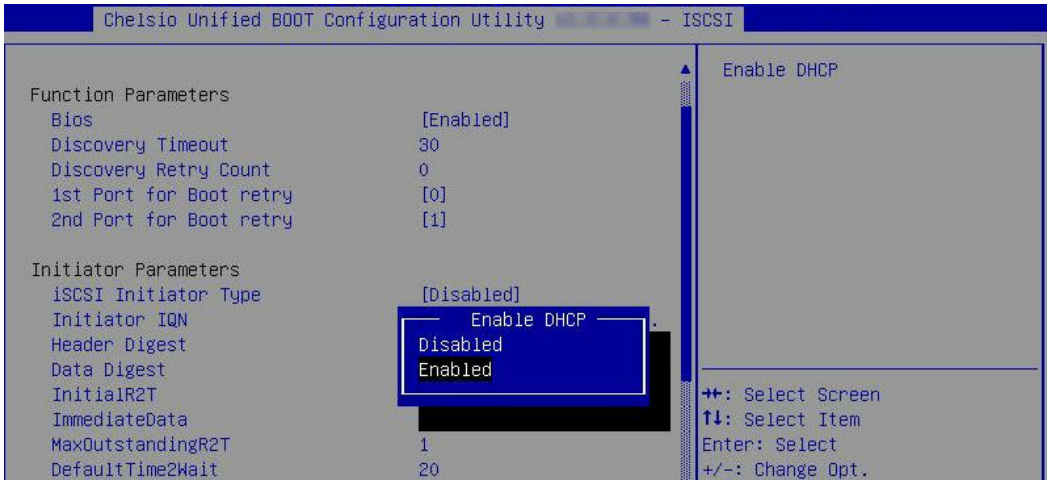

xiii. Under **Target Parameters**, select **Enabled** for the **Boot Target via DHCP** parameter to discover target using DHCP.

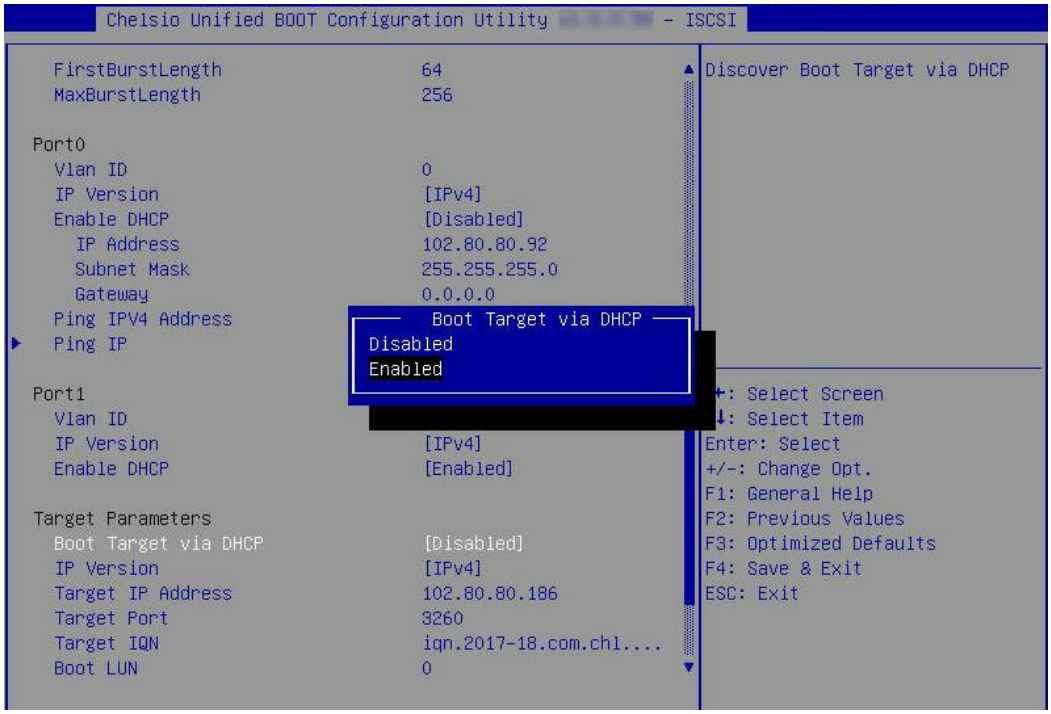

To discover target via static IP, select **Disabled** and provide the target IP.

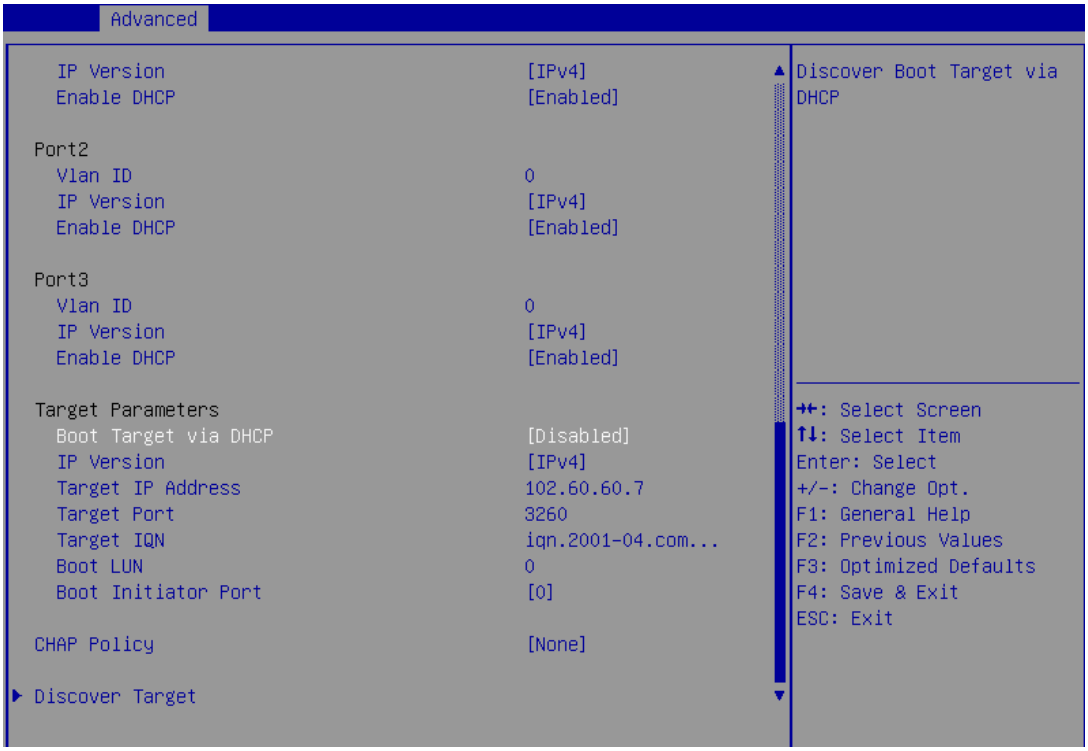

xiv. CHAP authentication is disabled by default. To enable and configure, highlight **CHAP Policy** and hit [Enter]. Select the policy type from the corresponding pop-up and hit [Enter] again.

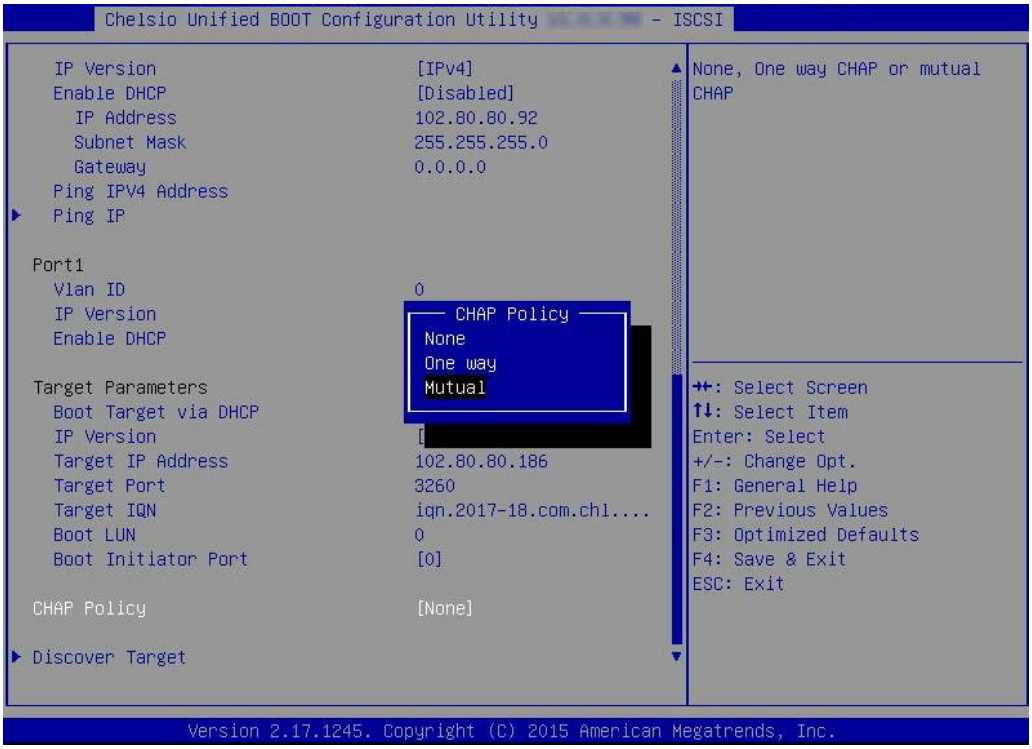

xv. Provide Initiator and Target CHAP credentials as per the CHAP policy selected.

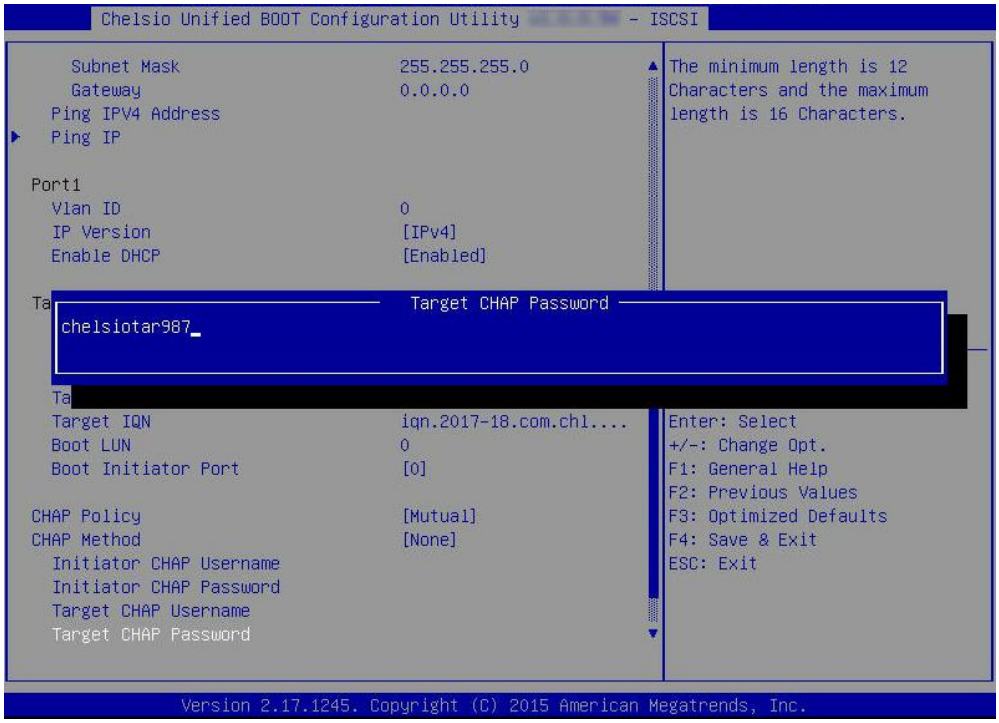

xvi. Select **Discover Target** and press [Enter] to discover iSCSI targets connected to the switch. Wait till all reachable targets are discovered.

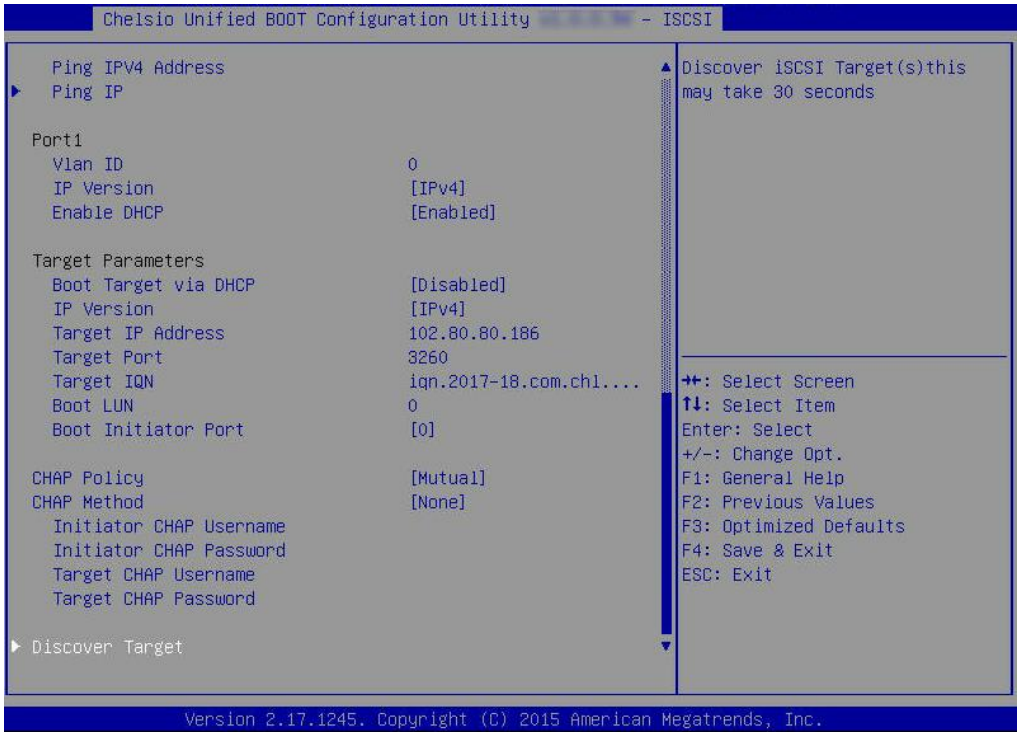

xvii. A list of available targets will be displayed. Select the target you wish to connect to and hit [Enter].

xviii. A list of LUNs configured on the selected target will be displayed. Select the LUN you wish to connect to and hit [Enter].

xix. Select **Save Changes** and press [Enter]

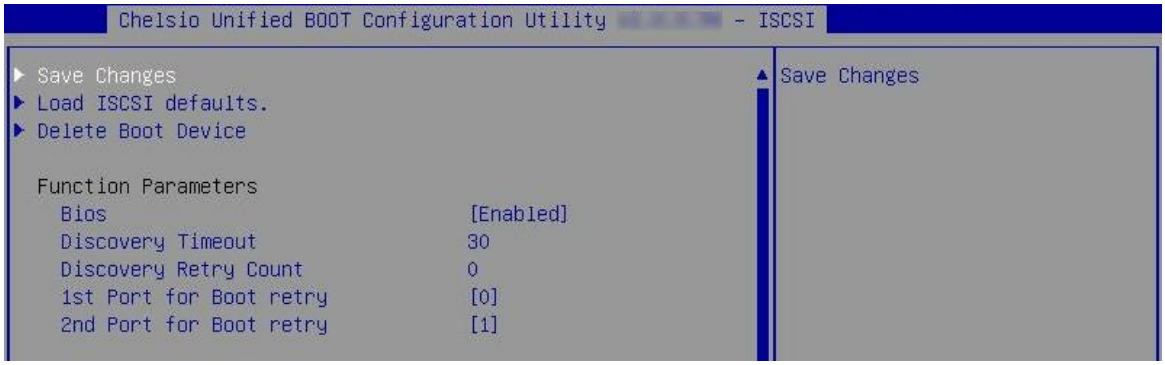

xx. Reboot the system for changes to take effect.

xxi. The discovered LUN should appear in the **Boot Configuration/ Boot Information** section and system BIOS.

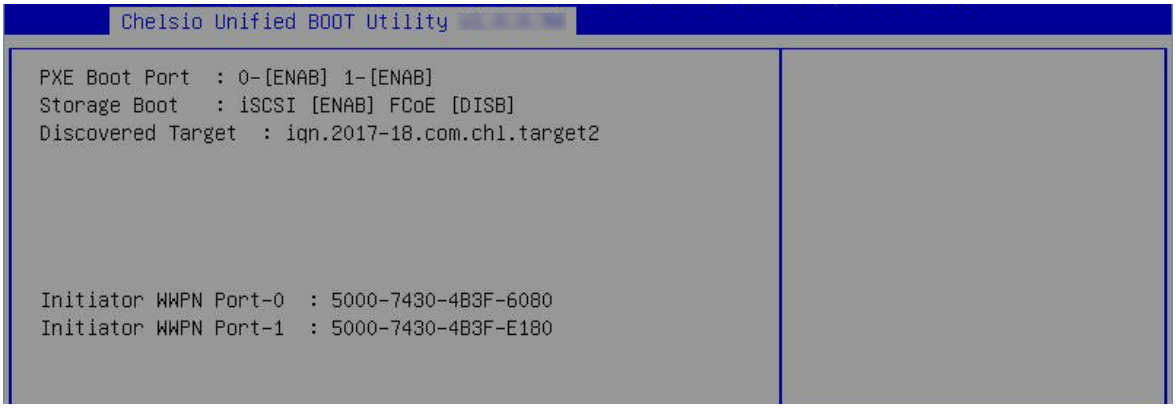

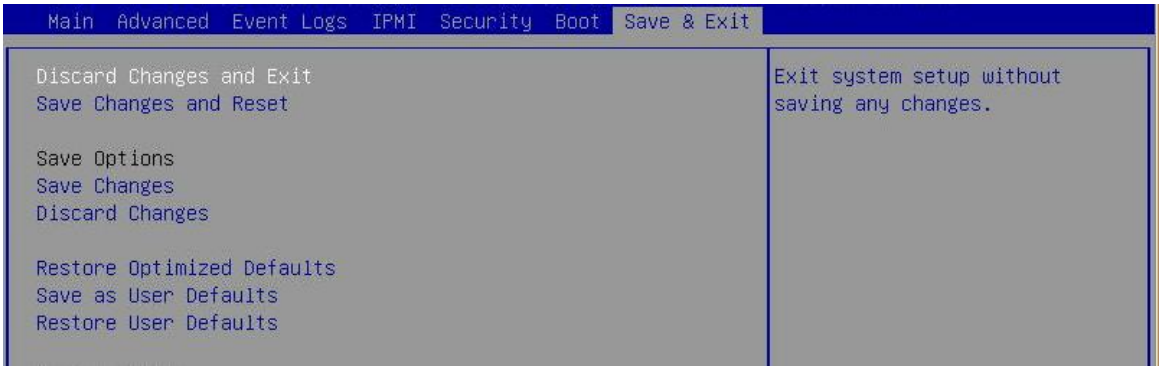

xxii. Select the LUN as the first boot device and exit from BIOS.

xxiii. Either boot from the LUN or install the required OS.

8. Creating Driver Update Disk (DUD)

The following section describes the procedure to create Driver Update Disks for RHEL and SLES distributions for Chelsio adapters. In case of T4 adapters, you can skip this step and use inbox drivers to install the operating system.

8.1. Creating DUD for RedHat Enterprise Linux

i. If you haven't done already, download the Chelsio Unified Wire driver package from Chelsio Download Center.

ii. Untar the package:

```
[root@host~]# tar zxvf <driver_package>.tar.gz
```

iii. Change your current working directory to *LinuxDUD* directory:

```
[root@host~]# cd <driver_package>/Uboot/LinuxDUD
```

iv. Insert a blank, formatted USB flash drive.

v. Depending on the distribution to be installed, copy the corresponding image file to the USB drive. For example, execute the following command for RHEL 6.6:

```
[root@host~]# cp Chelsio-DriverUpdateDisk-RHEL6.6-x86_64-x.xx.x.x.img <path
to USB drive>
```

Note *For RHEL 7.X, use Chelsio-DriverUpdateDisk-RHEL7.X-x86_64-x.xx.x.x.iso*

8.2. Creating DUD for Suse Enterprise Linux

i. If you haven't done already, download Chelsio Unified Wire driver package from Chelsio Download Center.

ii. Untar the package,

```
[root@host~]# tar zxvf <driver_package>.tar.gz
```

iii. Insert a blank USB flash drive.

iv. Format the USB drive

```
[root@host~]# mkfs.vfat /dev/sda1
```

v. Depending on the distribution to be installed, copy the corresponding image file to the USB stick. For example, execute the following command for SLES 11 sp4.

```
[root@host~]# dd if=/root/<driver_package>/Uboot/LinuxDUD/Chelsio-
DriverUpdateDisk-SLES11sp4-x86_64-x.xx.x.x.img of=/dev/sda1
```

9. OS Installation

9.1. Installation using Chelsio DUD

This is the recommended method for installing Linux OS using Chelsio PXE boot. The Chelsio Driver Update Disk (DUD) has support for all the new adapters. Use Network Boot (PXE Boot) media to install the OS, and provide the Driver Update Disk as per the detailed instructions for each OS.

The DUD supports installation of Linux distributions using Chelsio adapters over Network. There may be built-in Chelsio driver in these distributions. The driver may or may not work with Chelsio adapters, depending on the adapter in use, and the version of the driver that shipped in that particular distribution. Please flash the firmware provided in the package.

9.1.1. RHEL 7.X installation

i. Please make sure that the USB drive with DUD image is inserted. Type *e* and then *dd* at the boot prompt for the installation media. The *dd* option specifies that you will be providing a Driver Update Disk during the installation.

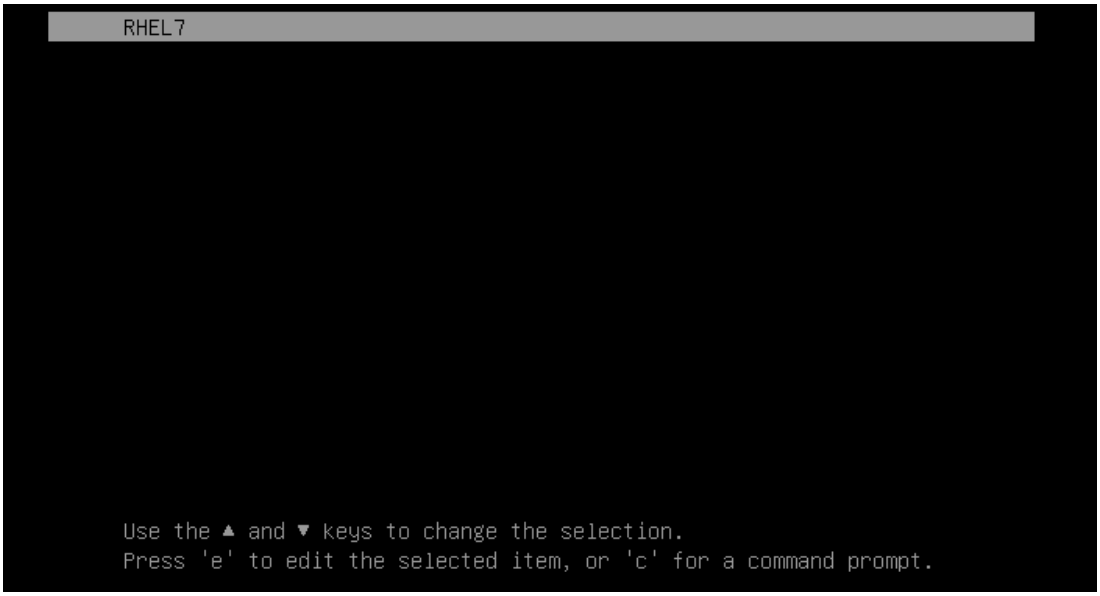

Note

- *In case of iSCSI boot, type **dd ip=ibft***
- *In case of T5 adapters with RHEL 7.2, use inbox drivers for installation. No DUDs required.*

ii. You will be asked to select the Driver Update Disk device from a list. USB drives usually show up as SCSI disks in Linux. Enter the index number of the device to be used and hit [Enter].

iii. The installer will search and display DUD image files found in the selected device. Enter the index number of the file to be used and hit [Enter].

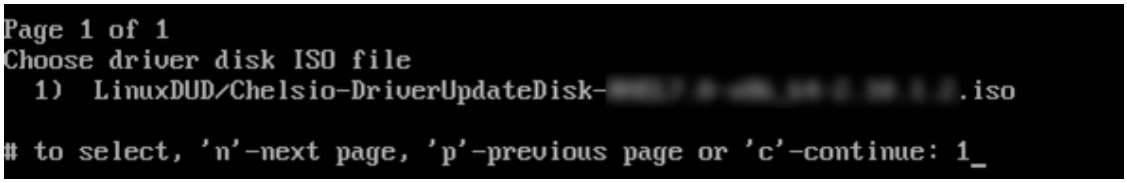

iv. Drivers provided in the DUD will be listed. Enter 1 to select FCoE driver (*csiostor*), or 2 to select Network driver (*cxgb4*). Hit [Enter]

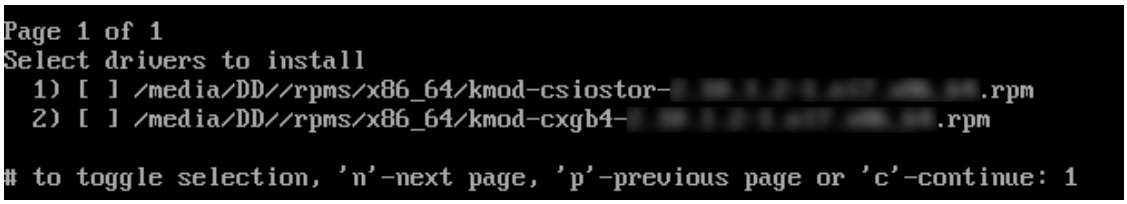

v. To select the next driver, enter the driver index or enter “c” to start the loading process. Hit [Enter]. The selected driver(s) will now be loaded.

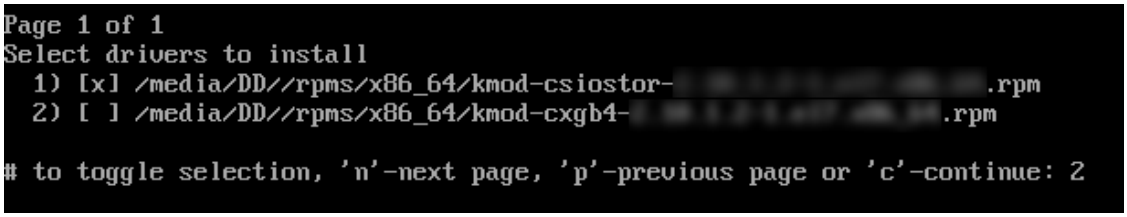

Note *To deselect a driver, enter the index of the selected driver and hit [Enter]*

vi. The **Driver disk prompt** will be displayed again. Follow the same procedure mentioned above to select any other drivers you wish to load or press “C” to skip and start the loading process.

vii. After the drivers are successfully loaded, OS installation will commence. Proceed as usual.

9.1.2. RHEL 6.X installation

i. Please make sure that the USB drive with DUD image is inserted. Press *Tab* and then type *dd* at the boot prompt for the installation media. The *dd* option specifies that you will be providing a Driver Update Disk during the installation.

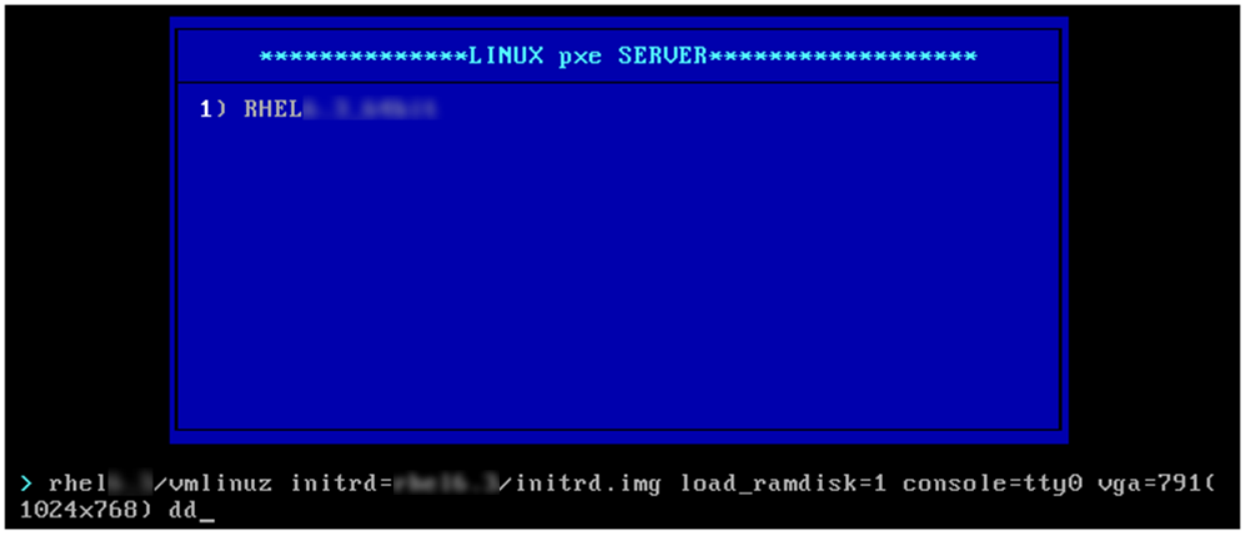

ii. The installer will load and prompt you for the driver update disk. Select “Yes” and hit [Enter] to proceed.

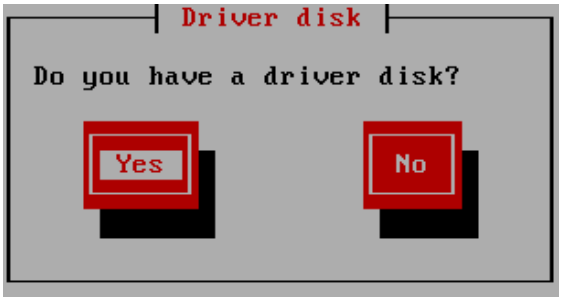

iii. You will be asked to select the Driver Update Disk device from a list. USB drives usually show up as SCSI disks in Linux. So, if there are no other SCSI disks connected to the system, the USB drive would assume the first drive letter “a”. Hence the drive name would be “sda”.

You can view the messages from the Linux kernel and drivers to determine the name of the USB drive, by pressing [Alt] + [F3] or [Alt] + [F4]. Press [Alt] + [F1] to get back to the list.

iv. Select the Appropriate image file and Choose “OK”. Now the installer will search for the appropriate drivers from the driver disk and load them. This step may take some time. Check on the [Alt] + [F3] or [Alt] + [F4] screens for log messages.

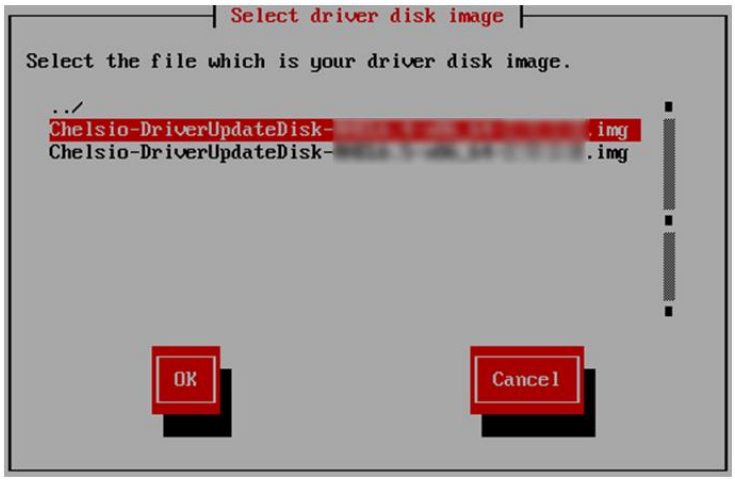

v. The installer will ask if you wish to load more drivers. Choose “Yes” to load if you have any other drivers to load. Otherwise choose “No”.

vi. Select the required language from the list.

vii. Select the type of keyboard you have from the list.

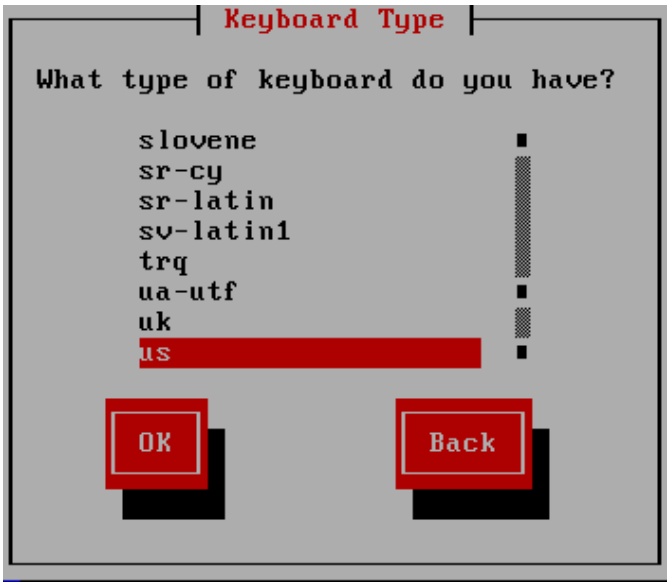

viii. In this step, you can choose the source which contains the OS installation ISO image. In this case, select “NFS directory”.

ix. The Chelsio detwork devices will be displayed. Select the appropriate Chelsio NIC interface to proceed with installation.

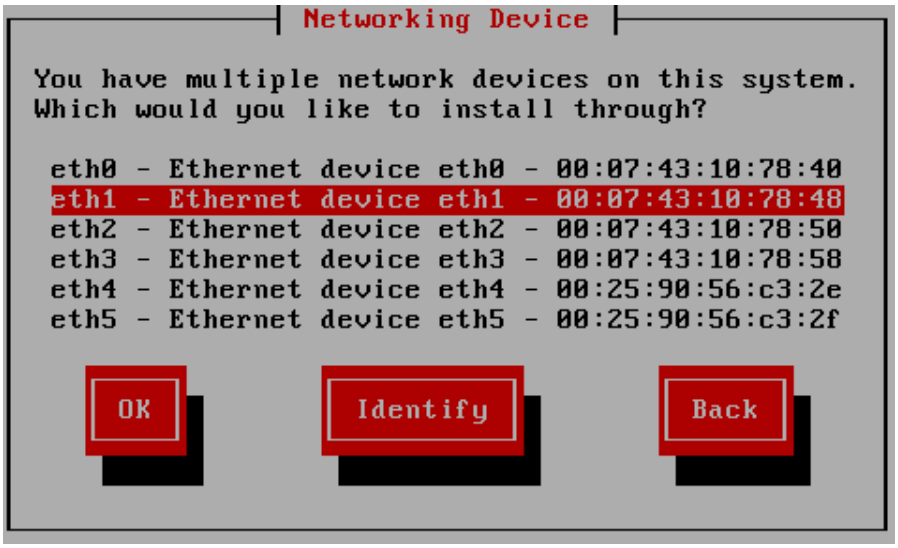

x. Here you can specify if you want to configure your network interfaces using DHCP or manually using IPv4. IPv6 is currently not supported. Hence disable IPv6 before proceeding.

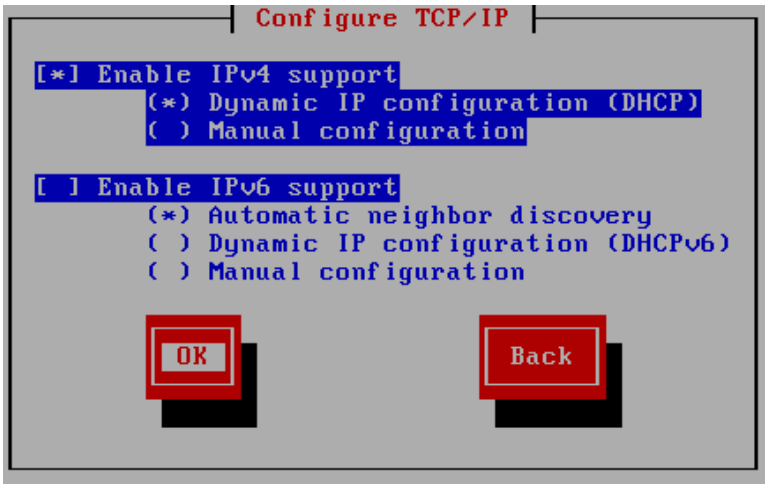

xi. Proceeding with the installation will get NFS/FTP/HTTP setup page. Here, provide NFS server details to proceed with the installation. Then the graphical Installation screens for RHEL will appear. Proceed with the installation as usual.

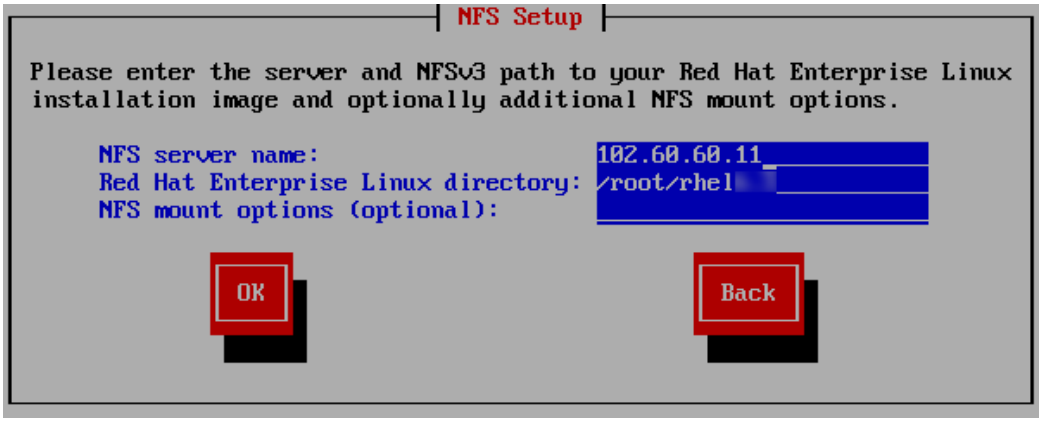

9.1.3. SLES 11 SPx/SLES 12/SLES 12 SPx installation

i. Please make sure that the USB drive with DUD image is inserted.
ii. Select the appropriate entry from the PXE menu and press [Enter].

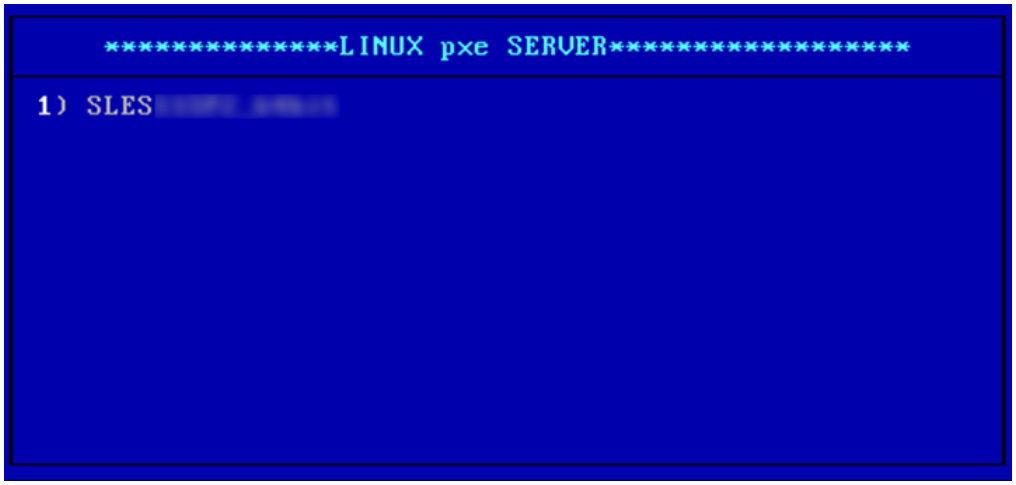


```
[    2.227429] hp_sw: device handler registered
[    2.252145] rdac: device handler registered

>>> SUSE Linux Enterprise Server 11 installation program v3.3.81 (c) 1996-2010 SUSE Linux Products GmbH <<<
Starting udev... ok
Loading basic drivers... ok
Starting hardware detection... ok
(If a driver is not working for you, try booting with brokenmodules=driver_name.)

Activating usb devices... ok
AMI Virtual CDROM
  drivers: usb_storage*
JetFlash Transcend 2GB
  drivers: usb_storage*
Logitech USB Multimedia Keyboard
  drivers: usbhid*
Chelsio Ethernet controller
  drivers: cxgb4*
Chelsio Ethernet controller
  drivers: cxgb4*
Chelsio Ethernet controller
  drivers: cxgb4*
Chelsio Ethernet controller
  drivers: cxgb4*
Chelsio Ethernet controller
  drivers: cxgb4*
Intel 82574L Gigabit Network Connection
  drivers: e1000e*
Intel 82574L Gigabit Network Connection
  drivers: e1000e*
Driver Update: Chelsio Network driver update Disk
Driver Update: Chelsio FCoE Initiator Driver Update Disk
Driver Updates added:
  Chelsio Network driver update Disk
  Chelsio FCoE Initiator Driver Update Disk
```

iii. Select “Start Installation” and then “Start Installation or Update”.

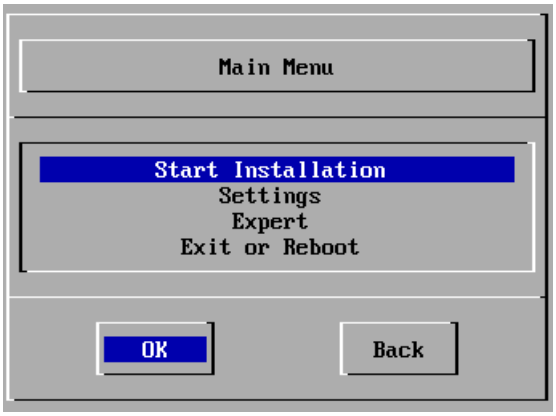

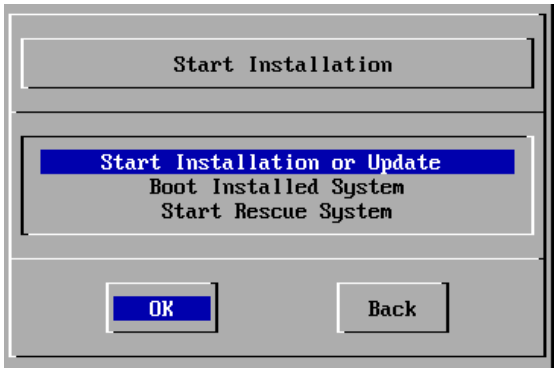

iv. Select “Network” as the source of medium to install the SLES Operating System.

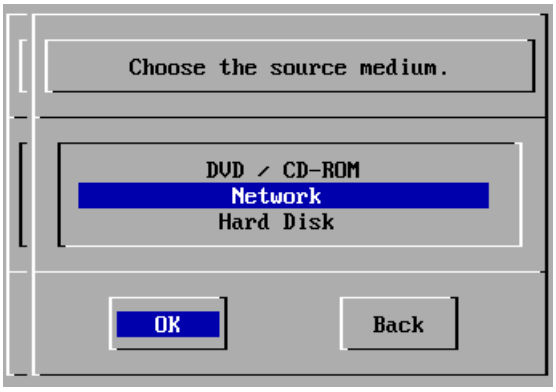

v. Select the desired Network protocol from the list presented.

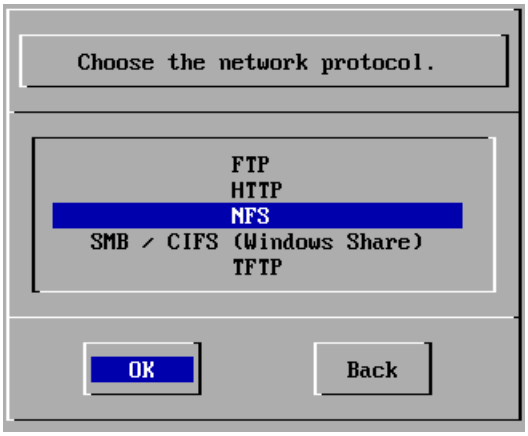

vi. Select the appropriate Chelsio interface from the list to proceed with installation. You can view the messages from the Linux kernel and drivers to determine the name of NIC interface by pressing [Alt] + [F3] or [Alt] + [F4]. Press [Alt] + [F1] to get back to the list.

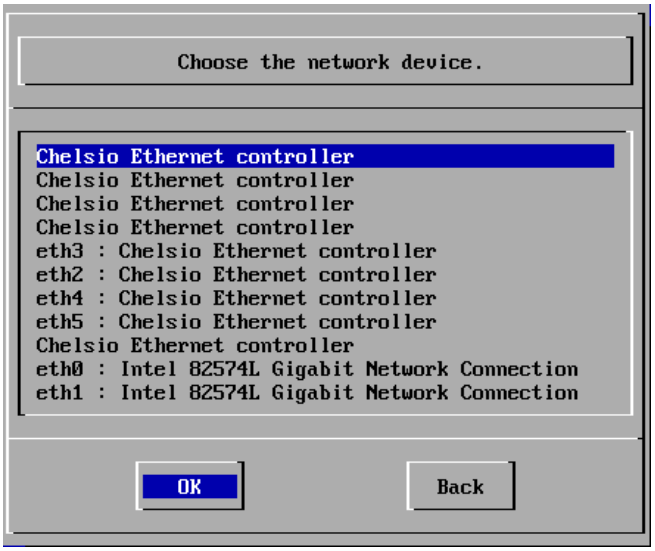

vii. Select “Yes” to configure the network interface selected in the previous step using DHCP.

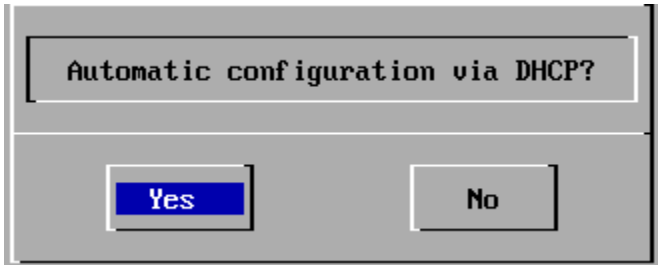

viii. Provide a valid NFS/FTP/HTTP/TFTP Server IP address to proceed.

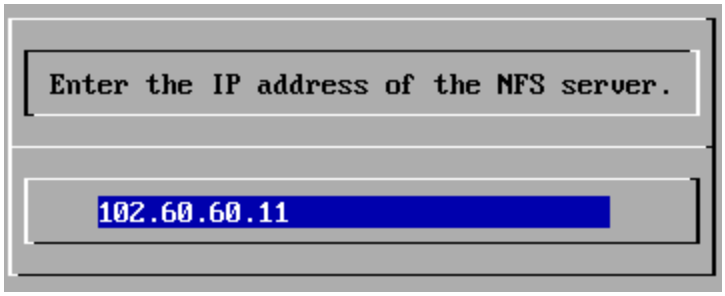

ix. Provide a valid directory path to the operating system to be installed.

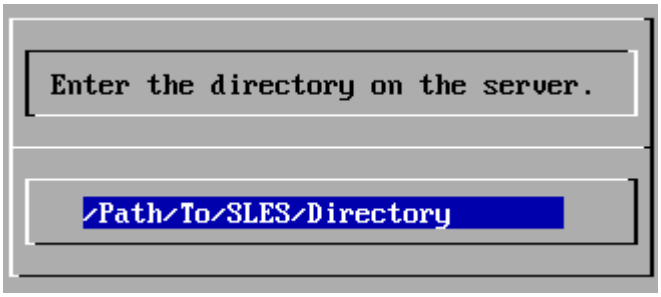

x. Proceed with the installation as usual.

9.2. Installation on FCoE LUN

- If you are installing using CD/DVD, please make sure that the USB drive with DUD image is inserted. Also, change the boot priority to boot from CD/DVD in the BIOS setup.
    i. Insert the OS installation disc into your CD/DVD ROM.
    ii. On the Grub menu, choose *Install or upgrade an existing system* option if not already selected.
    iii. Type *e* and then *dd* at the boot prompt for RHEL 7.
    iv. Load Chelsio Driver Update Disk depending on the Linux distribution (Click here for RHEL 7.3)
- If you are installing from a PXE server, please refer **8.1. Installation using Chelsio DUD** (Click here for RHEL 7.3) section to load Chelsio Driver Update Disk.

After successfully loading Chelsio DUD, follow the procedure mentioned below to continue installation, based on the distribution.

9.2.1. RHEL 7.3

i. Choose your installation language and click **Continue**

ii. Click **INSTALLATION DESTINATION** under **SYSTEM.**

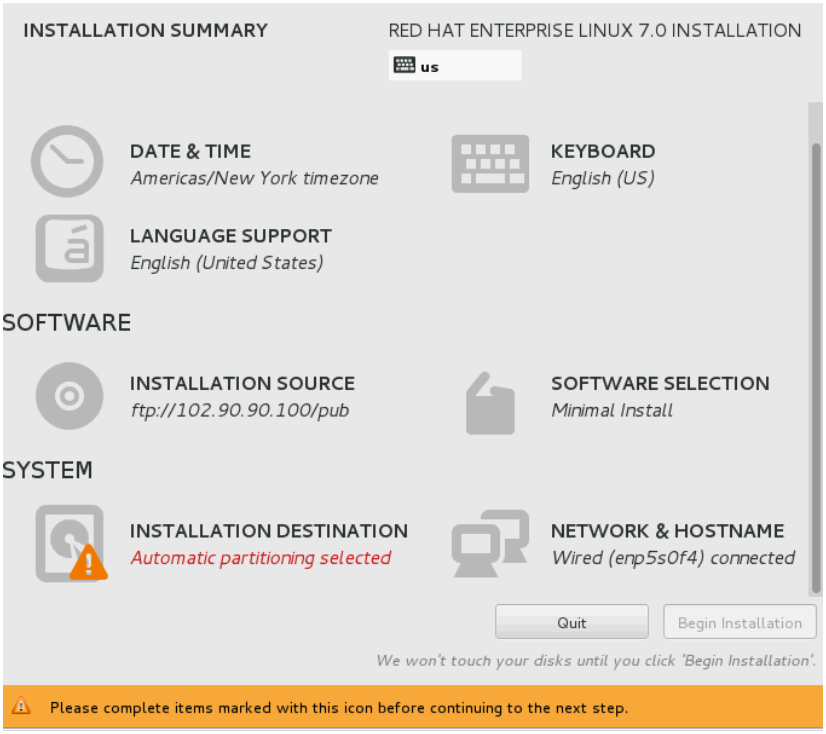

iii. The discovered FC/FCoE LUNs will appear as local storage in the **Local Standard Disks** section**.** Select the LUN which was saved as boot device in system BIOS.

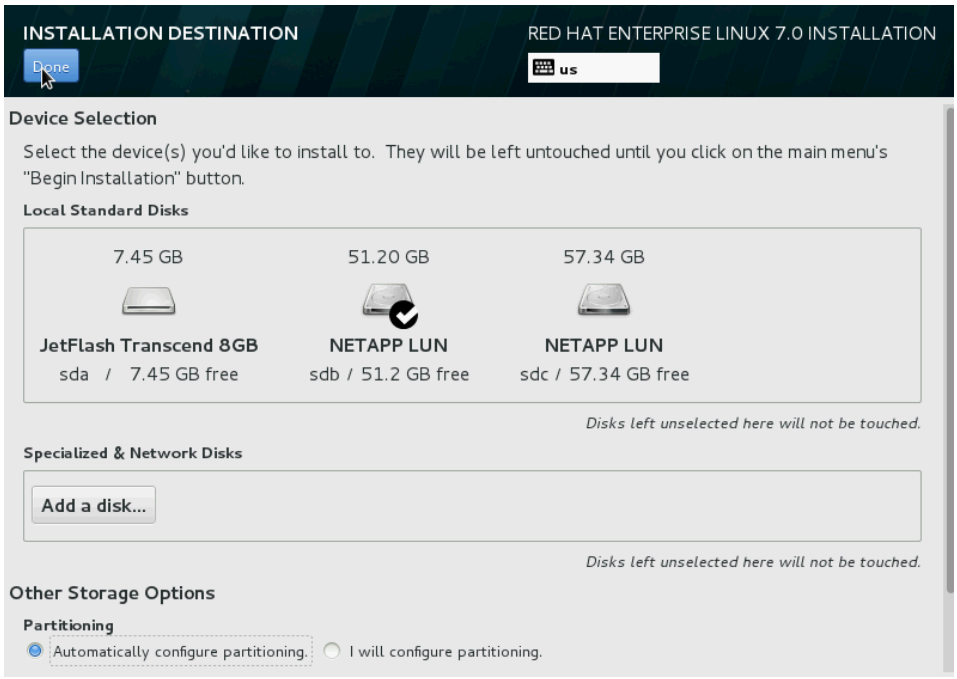

Note *Make sure the same LUN discovered at the Option ROM stage is selected for OS installation.*

iv. Under **Other Storage Options**, you can either chose to configure partition automatically or manually. Select the appropriate option and click **Done**. Then proceed with the installation as usual.

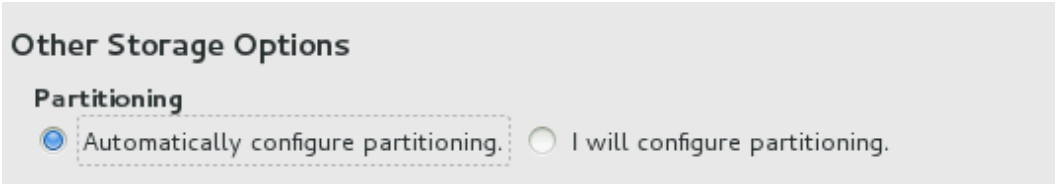

9.3. Installation on iSCSI LUN

- If you are installing using CD/DVD, please make sure that the USB drive with DUD image is inserted. Also, change the boot priority to boot from CD/DVD in the BIOS setup.
    i. Insert the OS installation disc into your CD/DVD ROM.
    ii. On the Grub menu, choose *Install or upgrade an existing system* option if not already selected.
    iii. For RHEL 6 press *Tab* and then type *dd.* For RHEL 7, type *e* and then *dd ip=ibft*. This will ensure that Chelsio iSCSI Initiator driver is used as SCSI transport medium. For SLES distributions, press *Tab* and then type *dd.*
    iv. Load Chelsio Driver Update Disk depending on the Linux distribution (Click here for RHEL 7.x; Click here for RHEL 6.x; Click here for SLES 11 SPx/SLES 12/SLES 12 SPx.

- If you are installing from a PXE server, please refer **8.1. Installation using Chelsio DUD**(Click here for RHEL 7.x; Click here for RHEL 6.x; Click here for SLES 11 SPx/SLES 12/SLES 12 SPx) section to load Chelsio Driver Update Disk.

After successfully loading Chelsio DUD, follow the procedure mentioned below to continue installation, based on the distribution.

9.3.1. RHEL 7.x

i. On the installer, welcome screen, choose your installation language and click **Continue**

ii. Click **Installation Destination** under **SYSTEM.**

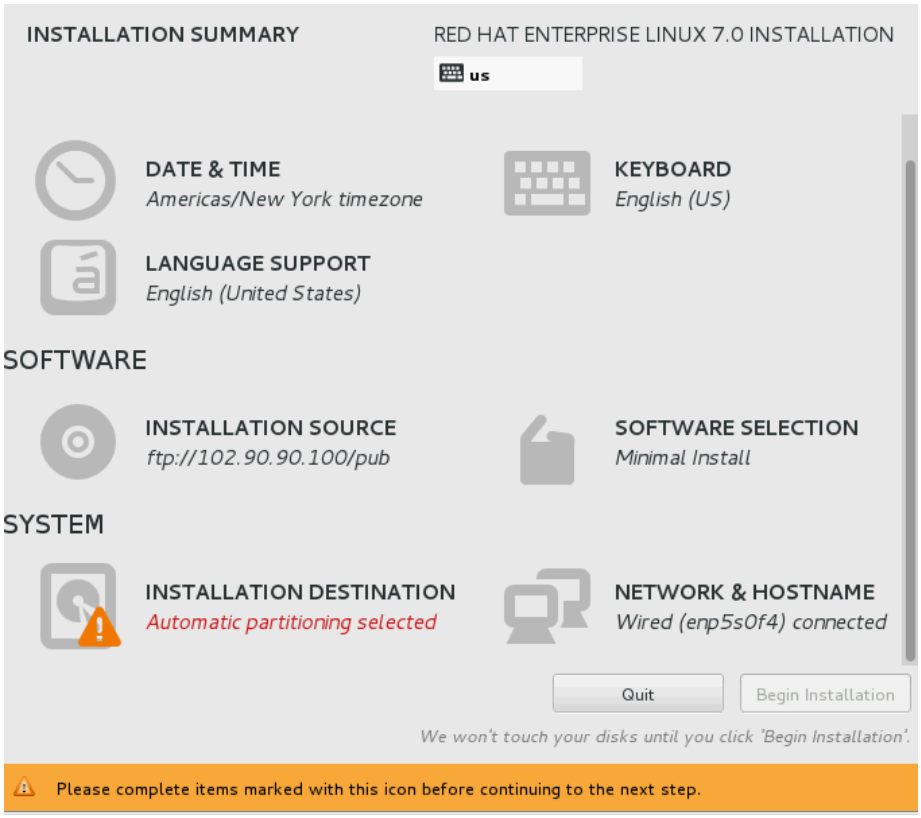

iii. Click **Add a disk**

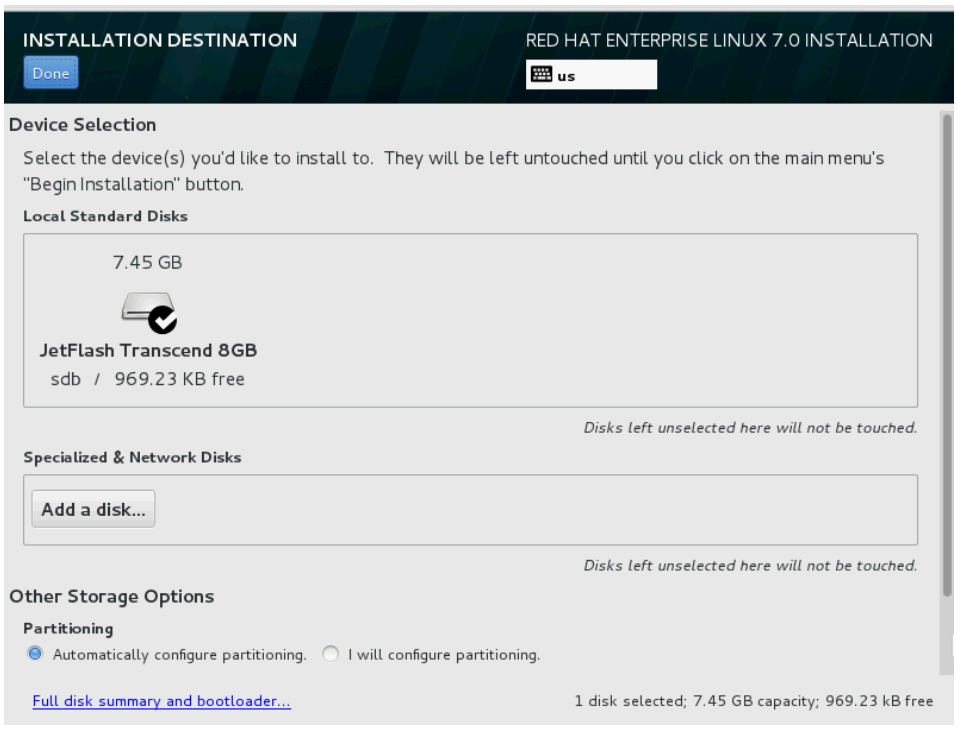

iv. The discovered iSCSI LUNs will appear in the **Search** tab. Select it and click **Done**.

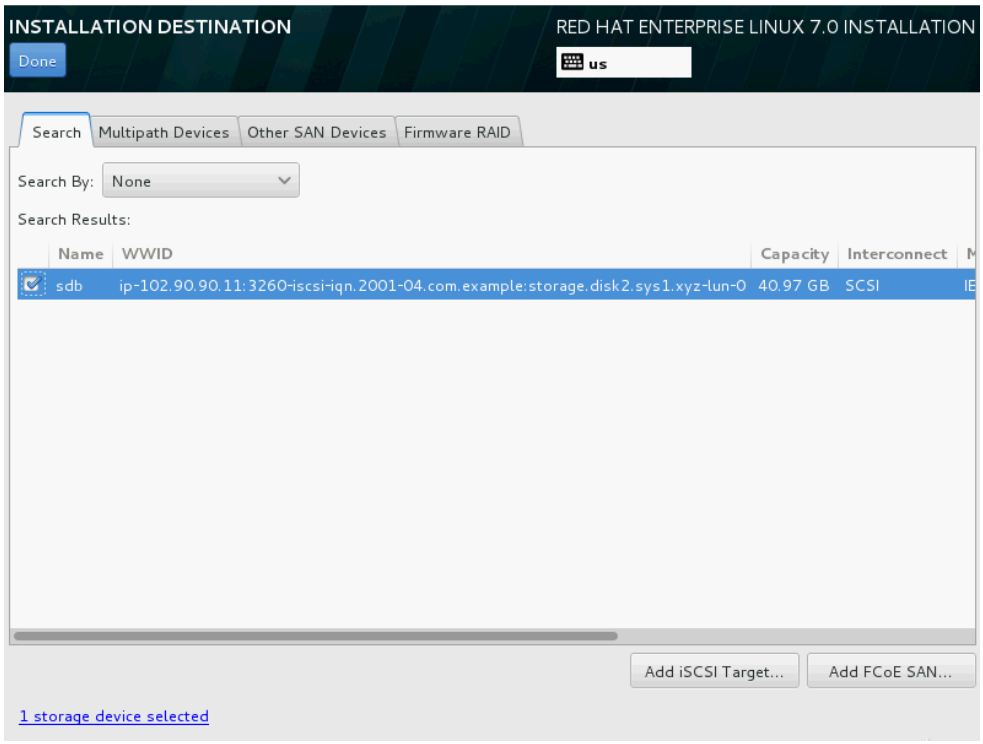

Note *Make sure the same LUN discovered at the Option ROM stage is selected for OS installation.*

v. Under **Other Storage Options**, you can either chose to configure partition automatically or manually. Select the appropriate option and click **Done**. Then proceed with the installation as usual.

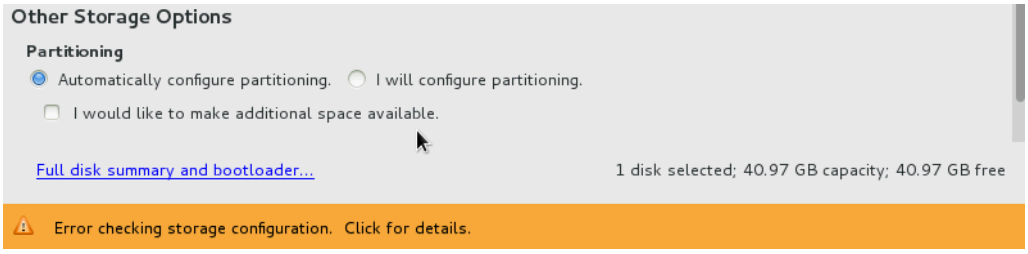

9.3.2. RHEL 6.x

i. Click **Next** when the graphical installer screen appears.

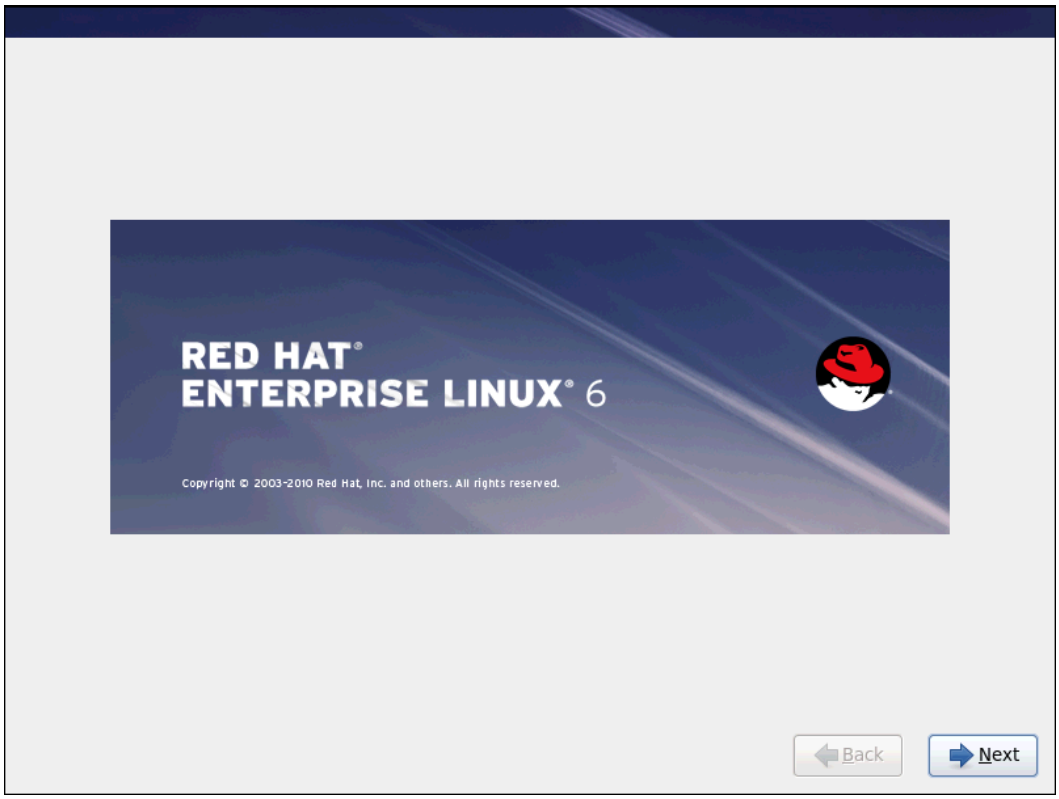

ii. Select **Specialized Storage Devices** radio button and click **Next**.

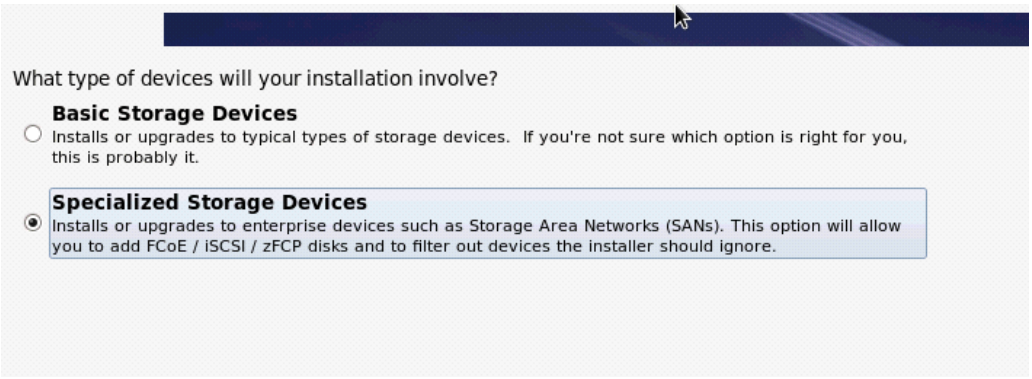

iii. The discovered LUNs will appear in the **Basic Devices** tab. Select the LUN which was saved as boot device in system BIOS and click **Next**.

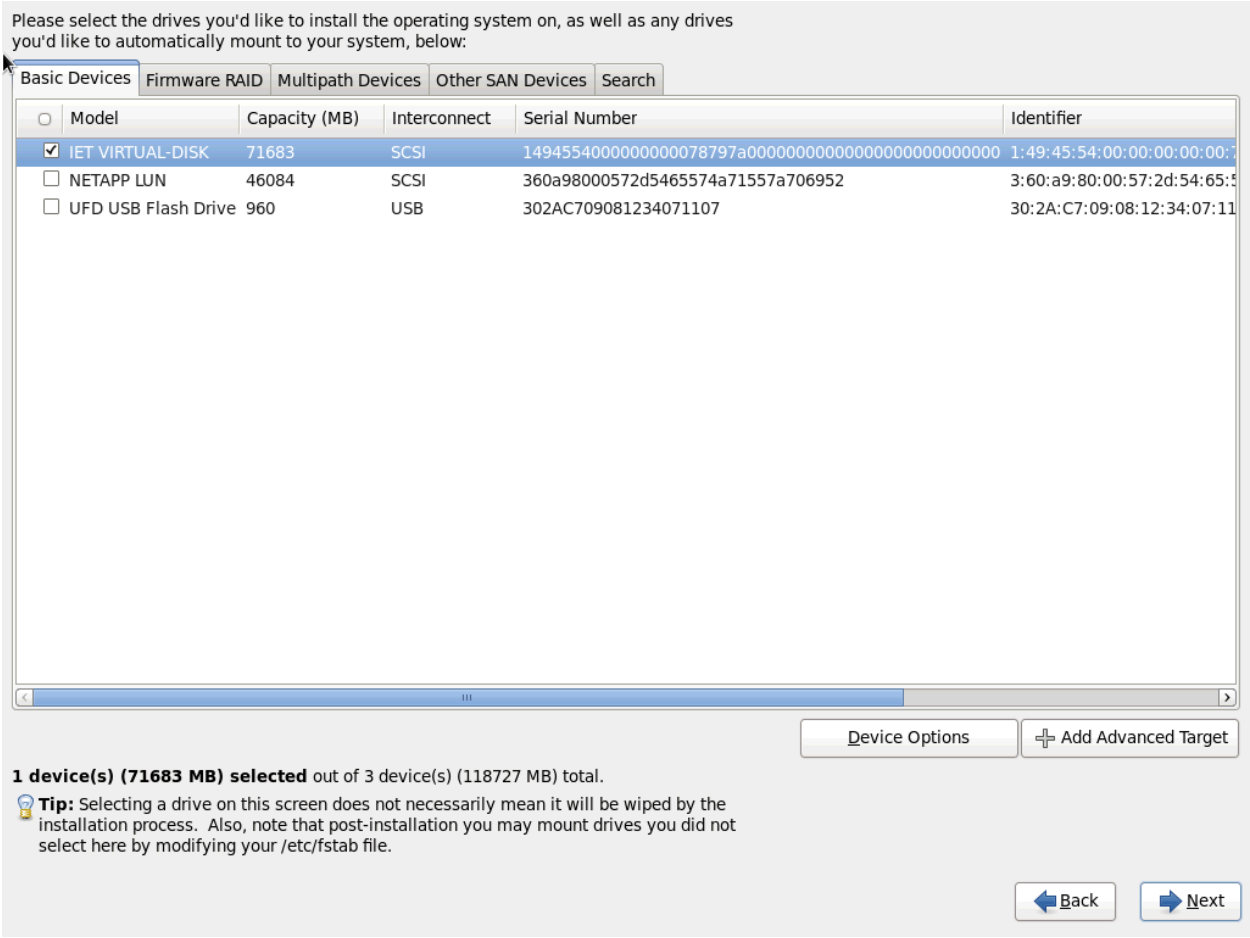

Note *Make sure the same LUN discovered at the Option ROM stage is selected for OS installation.*

iv. Proceed with the installation as usual.

9.3.3. SLES 11 SPx installation

i. Choose installation language and Keyboard layout type. Select the checkbox **I Agree to the License terms** and click **Next**.

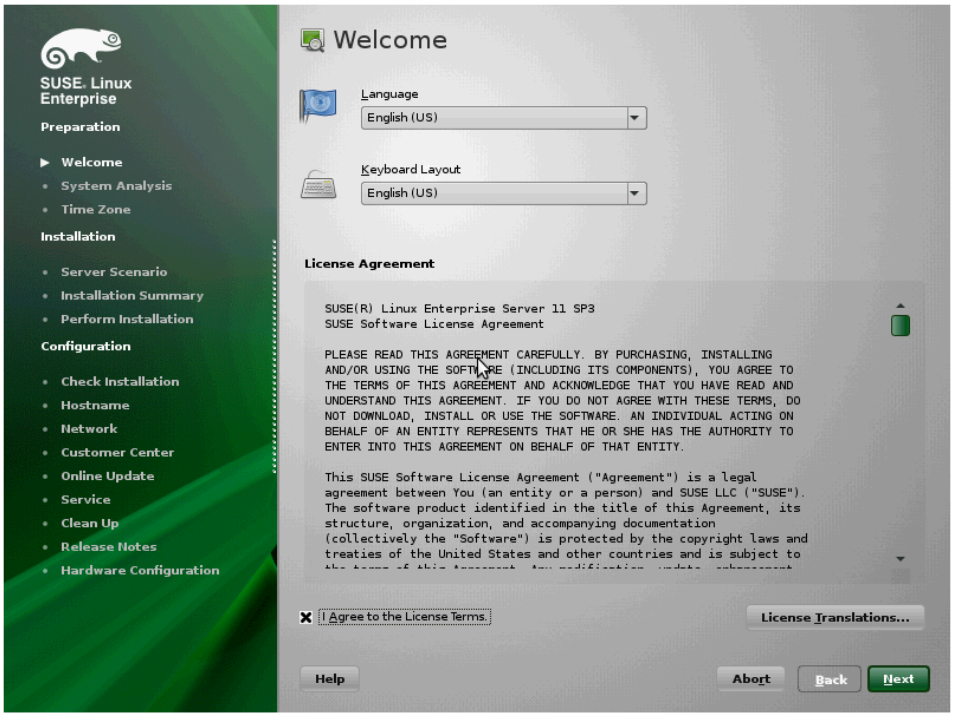

ii. Click **Configure iSCSI Disks** in the **Disk Activation** screen.

iii. The discovered LUNs will appear in the **Connected Targets** tab. Select the LUN which was saved as boot device in system BIOS and click **OK**.

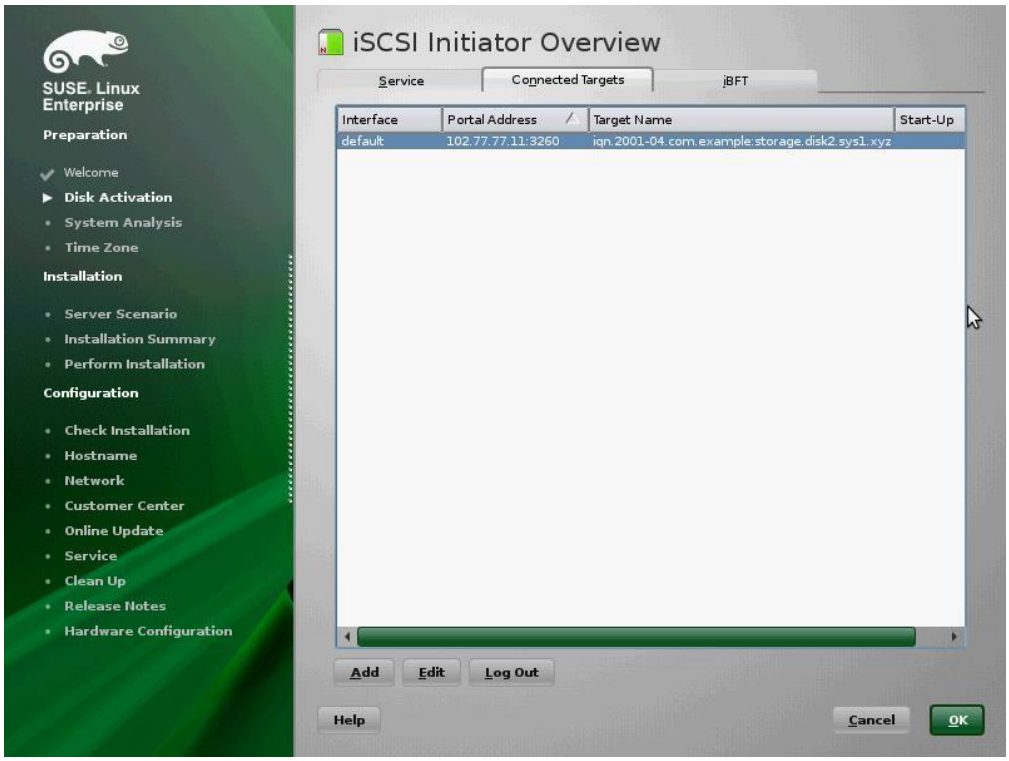

iv. Select **New Installation** to perform a fresh installation and click **Next**.

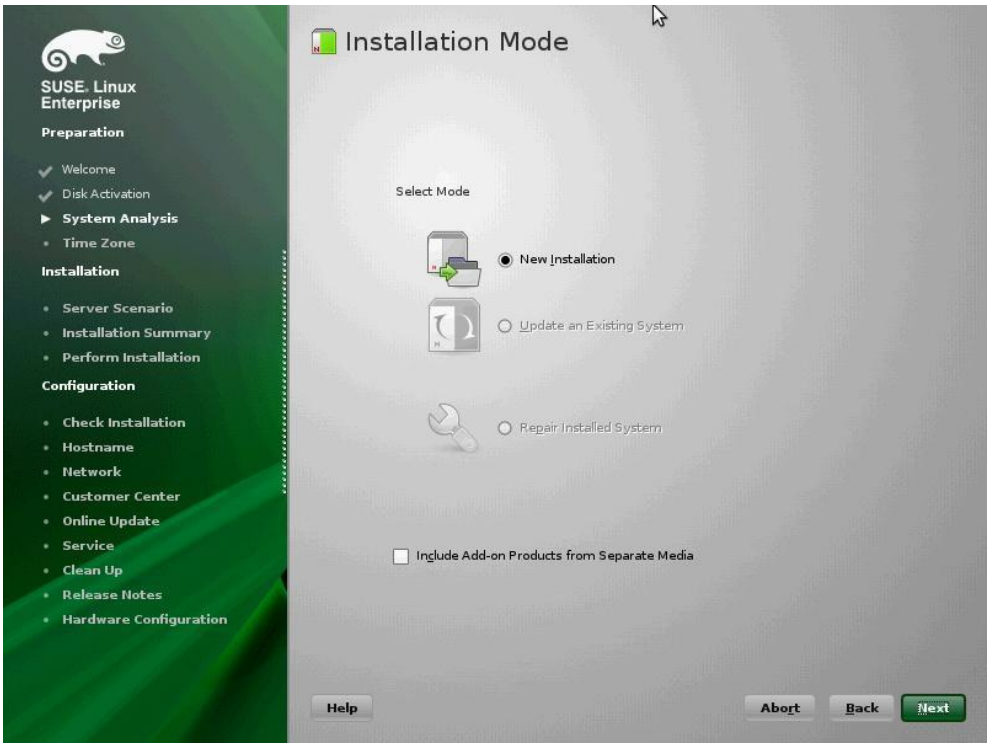

v. Configure Clock and Time Zone settings. Click **Next**.

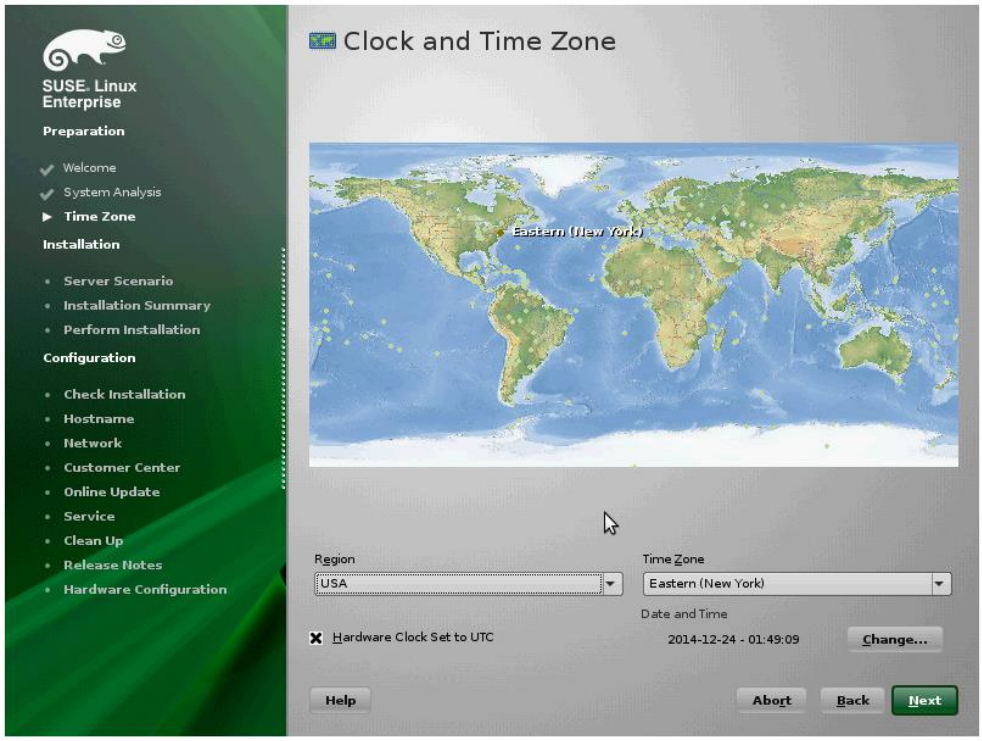

vi. Choose from the available server base scenarios and click **Next**.

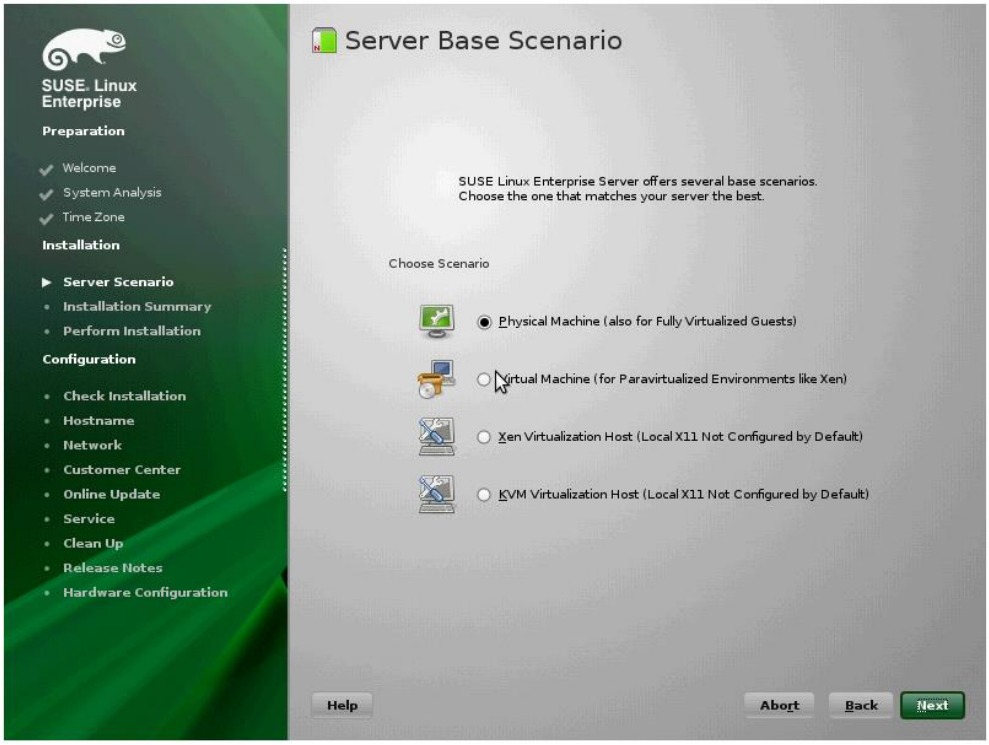

vii. The **Installation Settings** screen displays the summary of user-selected and YaST-suggested options for the installation. You can review and modify them if required. Basic settings can be changed in the **Overview** tab and advanced settings can be changed in the **Expert** tab. To change, click on one of the headlines or click **Change** and select the category. Finally, click **Next**.

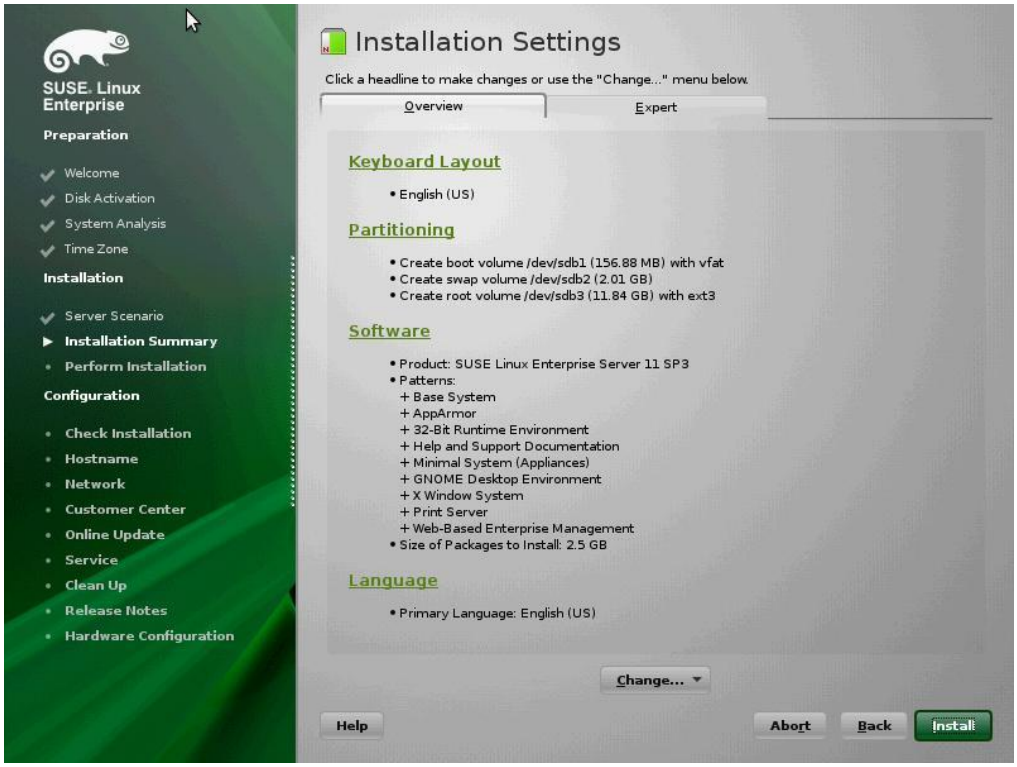

viii. Proceed with installation as usual.

9.3.4. SLES 12/SLES 12 SPx Installation

i. Choose installation language and Keyboard layout type. Select the checkbox **I Agree to the License terms** and click **Next**.

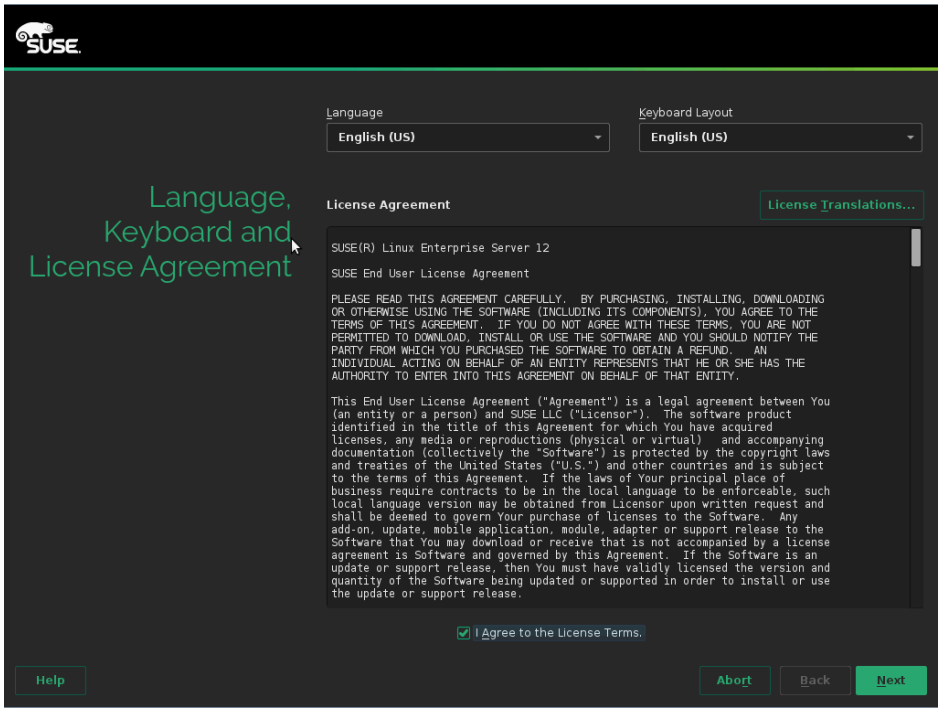

ii. Click **Configure iSCSI Disks** in the **Disk Activation** screen.

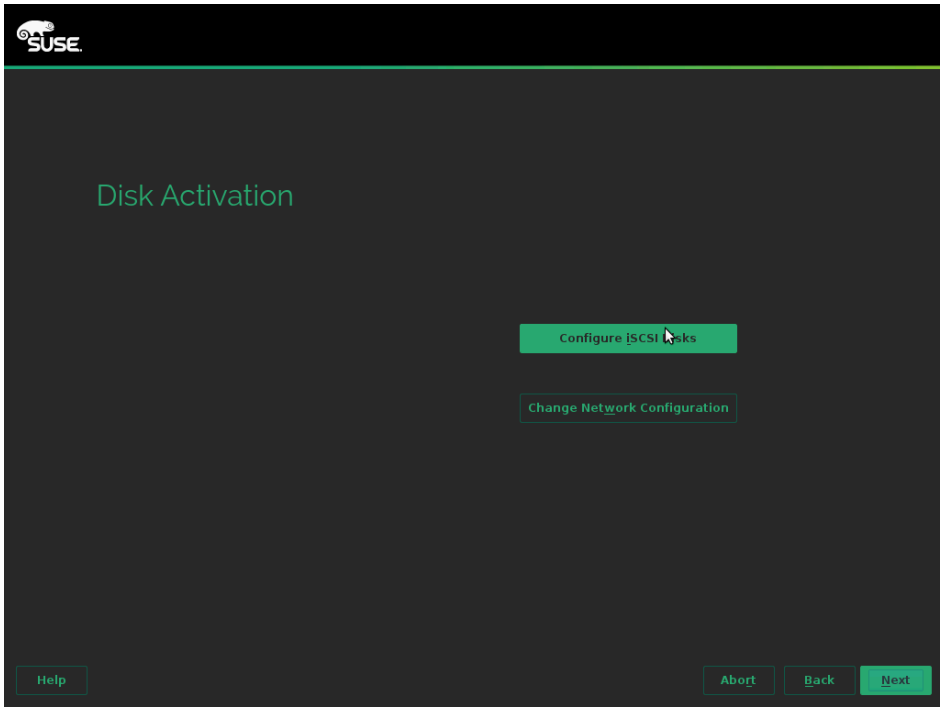

iii. The discovered LUNs will appear in the **Connected Targets** tab. Select the LUN which was saved as boot device in system BIOS and click **OK**.

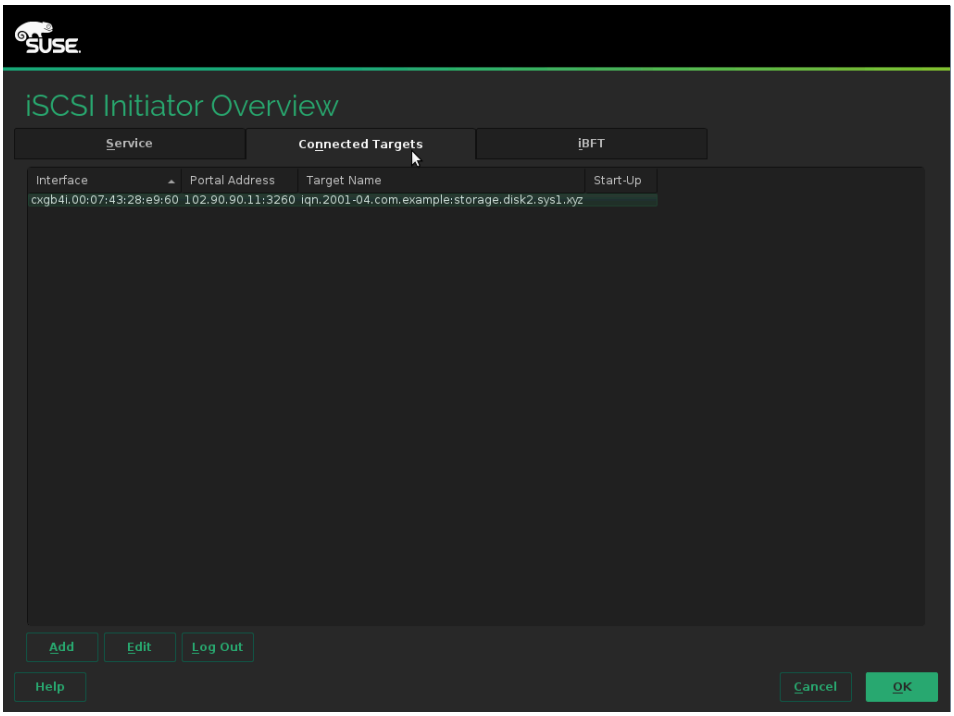

Note *Make sure the same LUN discovered at the Option ROM stage is selected for OS installation.*

iv. Proceed with the installation as usual.

XXV. Appendix A

1. Troubleshooting

- ***Cannot bring up Chelsio interface***

Make sure you have created the corresponding network-script configuration file as stated in **Cheslsio Unified Wire** chapter (See Creating network-scripts). If the file does exist, make sure the structure and contents are correct. A sample is given in the **Chelsio Unified Wire** chapter (See Configuring network-scripts). Another reason may be that the IP address mentioned in the configuration file is already in use on the network.

- ***Cannot ping through Chelsio interface***

First, make sure the interface was successfully brought up using ifup ethX (where ethX is your interface) and that it is linked to an IP address, either static or obtained through DHCP.

You then may want to check whether the destination host (i.e. the machine you are trying to ping) is up and running and accepts ICMP requests such as ping. If you get a return value of 0 when doing a `cat /proc/sys/net/ipv4/icmp_echo_ignore_all` on the remote host that means it is configured to reply to incoming pings. Change `ipv4` to `ipv6` in the path if you are using IPv6. Note that this is a Linux-only tip.

If you have more than one interface wired to the network, make sure you are using the right one for your outgoing ping requests. This can be done by using the `-I` option of the ping command, as shown in the following example:

```
[root@host~]# ping -I eth1 10.192.167.1
```

Where 10.192.167.1 is the machine you want to ping.

- ***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.

- ***"unexpected CM event" messages seen with iWARP traffic***

One of the reason for this could be port number collisions. To fix this, use iWARP port mapper (iwpmd).

```
[root@host~]# iwpmd
```

Note iWARP port mapper (iwpmd) has its own issues.

2. Chelsio End-User License Agreement (EULA)

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

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

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

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

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

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

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

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

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

6. Limited Warranty. If Chelsio furnishes the Software to you on media, Chelsio warrants only that the media upon which the Software is furnished will be free from defects in

material or workmanship under normal use and service for a period of thirty (30) days from the date of delivery to you.

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

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

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

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

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

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

11. Government Restricted Rights. The Software is subject to restricted rights as follows. If the Software is acquired under the terms of a GSA contract: use, reproduction or disclosure is subject to the restrictions set forth in the applicable ADP Schedule contract. If the Software is acquired under the terms of a DoD or civilian agency contract, use, duplication or disclosure by the Government is subject to the restrictions of this Agreement in accordance with 48 C.F.R. 12.212 of the Federal

Acquisition Regulations and its successors and 49 C.F.R. 227.7202-1 of the DoD FAR Supplement and its successors.

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

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

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