



Chelsio Unified Wire for Windows

Installation and User's Guide



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

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

All third-party trademarks are copyright of their respective owners.

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

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



Chelsio Communications (Headquarters)

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

www.chelsio.com

Tel: 408.962.3600
Fax: 408.962.3661

Chelsio (India) Private Limited

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

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

Chelsio KK (Japan)

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

Sales

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

Support

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

Copyright © 2020. 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
1.7.8	01/31/2019
1.7.9	03/05/2019
1.8.0	03/13/2019
1.8.1	03/29/2019
1.8.2	04/11/2019
1.8.3	04/24/2019
1.8.4	05/21/2019
1.8.5	05/23/2019
1.8.6	05/30/2019
1.8.7	07/03/2019
1.8.8	07/24/2019
1.8.9	12/09/2019
1.9.0	01/30/2020
1.9.1	02/10/2020
1.9.2	06/23/2020
1.9.3	07/07/2020
1.9.4	08/31/2020

TABLE OF CONTENTS

I.	CHELSIO UNIFIED WIRE	7
1.	Introduction	8
1.1.	Features	8
1.2.	Hardware Requirements	9
1.3.	Software Requirements	9
1.4.	Package Contents	9
2.	Hardware Installation	10
3.	Software/Driver Installation	11
3.1.	Prerequisites	11
3.2.	Unified Wire Installer	11
3.3.	Zip Package	16
4.	Firmware Update	20
5.	Software/Driver Update	21
5.1.	Windows Update	21
5.2.	Unified Wire Installer	28
5.3.	Zip Package	29
6.	Configuring Chelsio Network Interfaces	30
6.1.	Configuring 40G Adapters	30
6.2.	Assigning IP Address	31
6.3.	Deadman Watchdog Timer	32
6.4.	Chelsio Management	33
7.	Performance Tuning	39
8.	<i>cxgbtool</i> help	40
9.	Software/Driver Uninstallation	47
9.1.	Installer	47
9.2.	Zip Package	50
II.	NDIS FUNCTION	52
1.	Introduction	53
1.1.	Hardware Requirements	53
1.2.	Software Requirements	53
2.	Software/Driver Configuration and Fine-tuning	54
2.1.	Advanced Configuration	54
2.2.	NVGRE Offload	62
2.3.	VXLAN Task Offload	62
2.4.	Timestamps	62
2.5.	PacketDirect	63
2.6.	Virtual Environments	66
III.	IWARP (ND)	67

1. Introduction	68
1.1. Hardware Requirements	68
1.2. Software Requirements	68
2. Software/Driver Configuration and Fine-tuning	69
2.1. Registering iWARP(ND) driver	69
2.2. chrping application	69
2.3. HMA	70
IV. SMB DIRECT	71
1. Introduction	72
1.1. Hardware Requirements	72
1.2. Software Requirements	72
2. Software/Driver Configuration and Fine-tuning	73
2.1. Enabling SMB Direct	73
2.2. Verifying RDMA	73
2.3. RDMA with Virtual Switch (Mode 2)	74
2.4. Storage Replica	75
2.5. Storage Spaces Direct (S2D)	75
2.6. HMA	80
2.7. Troubleshooting	80
V. MESH TOPOLOGY	81
1. Introduction	82
1.1. Hardware Requirements	82
1.2. Software Requirements	82
1.3. Mesh topology	83
2. Software/Driver Installation	84
3. Software/Driver Configuration and Fine-tuning	85
VI. NDIS SR-IOV	86
1. Introduction	87
1.1. Hardware Requirements	87
1.2. Software Requirements	88
1.3. Prerequisites	88
2. Software/Driver Configuration and Fine-tuning	89
2.1. Enabling SR-IOV	89
2.2. Host Configuration	89
2.3. Guest (VM) Configuration	95
2.4. Guest RDMA (Mode 3)	96
VII. ISCSI STORPORT MINIPORT	101
1. Introduction	102

1.1.	Hardware Requirements	102
1.2.	Software Requirements	103
2.	Software/Driver Configuration and Fine-tuning	104
2.1.	Configuring iSCSI Initiator	104
2.2.	iSCSI Target Discovery and Login	111
2.3.	HMA	119
VIII.	ISER STORPORT MINIPORT	120
1.	Introduction	121
1.1.	Hardware Requirements	121
1.2.	Software Requirements	121
2.	Software/Driver Configuration and Fine-tuning	122
2.1.	Configuring iSER Target	122
2.2.	Configuring iSER Initiator	122
2.3.	iSER Target Discovery and Login	122
2.4.	HMA	128
IX.	NVME STORPORT MINIPORT	129
1.	Introduction	130
1.1.	Hardware Requirements	130
1.2.	Software Requirements	130
2.	Software/Driver Configuration and Fine-tuning	131
2.1.	Connecting to NVMe target	131
2.2.	Disconnecting from NVMe target	134
2.3.	HMA	134
X.	DATA CENTER BRIDGING (DCB)	135
1.	Introduction	136
1.1.	Hardware Requirements	136
1.2.	Software Requirements	137
2.	Software/Driver Configuration and Fine-tuning	138
2.1.	Network QoS (optional)	138
XI.	APPENDIX	141
1.	Troubleshooting	142
2.	Chelsio End-User License Agreement (EULA)	143

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 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 Terminator adapters bring a new level of performance metrics and functional capabilities to the computer networking industry.

1.1. Features

Chelsio Unified Wire is an easy to use utility developed to provide installation of 64-bit Windows based drivers and tools for Chelsio's adapters. It consists of the following components:

- NDIS Function driver (NIC)
- NVGRE Offload
- VXLAN Task Offload
- Timestamps (Hardware and Software)
- PacketDirect (Kernel Mode and ARM Moderation)
- VMMQ/vRSS
- iWARP (ND)
- SMB Direct
- RDMA with Virtual Switch (Mode 2)
- Storage Spaces Direct (S2D)
- Storage Replica (SR)
- Mesh Topology
- NDIS SR-IOV
- Guest RDMA (Mode 3)
- iSCSI Storport Miniport
- iSER Storport Miniport
- NVMe Storport Miniport
- Data Center Bridging (DCB)

1.2. Hardware Requirements

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

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

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

Note *T4 family of adapters are not supported.*

1.3. Software Requirements

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

Note *Drivers are WHQL certified for Windows Server 2019, 2016, 2012 R2 and 10 Client.*

1.4. Package Contents

- **Unified Wire Installer**


Chelsio Unified Wire comes with an interactive installer and support documentation. The documentation, consisting of README, Release Notes and User's Guide (this document), can be found in the <system_drive>\ChelsioUwire\docs\ directory after installing Chelsio Unified Wire.

- **Zip Package**

The zip package contains driver files, firmware binaries, adapter configuration binaries and support documentation (docs folder).

2. Hardware Installation


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

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

- vii. Holding the Chelsio adapter by the edges, align the edge connector with the PCI connector on the motherboard. Apply even pressure on both edges until the card is firmly seated. It may be necessary to remove the 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. Verify if the adapter was installed successfully. To do so, open **Device Manager** in **Control Panel**.
- xiii. Under **Other devices** section, Chelsio adapter should be listed as **Ethernet Controller**. If the adapter is not listed, right-click on the system name or click on the **Actions** menu and select **Scan for hardware changes**

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

- Physical functions 0 - 3: for the SR-IOV functions
 - Physical function 4: for all NIC functions of the card
 - Physical function 5: for iSCSI
 - Physical function 6: for FCoE
 - Physical function 7: Currently not assigned
- xiv. Once the Unified Wire package is installed, open **Device Manager** again. Expand **Network adapters** section and now Chelsio adapter should be listed.

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

3. Software/Driver Installation

Chelsio Unified Wire can be installed using the Installer or the zip package. Refer the relevant section below depending on the method of installation selected.

! Important *If you are planning to install Unified Wire in a cluster, it should be done on all the members. Using different versions of Unified Wire is not supported.*

3.1. Prerequisites

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

- Latest Microsoft Visual C++ Redistributable packages for Visual Studio (x86 and x64) are required if you are planning to use installer.
- Windows update KB2685811 should be installed on Windows 7 Client and 2008 R2 Server. It will be installed if *Automatic Updates* is configured or can be installed manually from <https://www.catalog.update.microsoft.com/Search.aspx?q=KB2685811>

3.2. Unified Wire Installer

Chelsio Unified Wire Installer provides two methods of installation: **GUI** or **CLI** mode. GUI mode provides an interactive GUI installer with customizable options. Whereas, CLI mode enables unattended installation of Chelsio drivers and software thereby relieving the user from monitoring the installation process and providing input via dialog boxes.

3.2.1. Enabling Test Signing

In case the drivers in the package are not WHQL certified, follow the steps mentioned below to enable test signing on your system:

- i. Goto Start->Run command option, enter "cmd" and press OK. This will open the command prompt utility.
- ii. Run the following command:

```
C:\Users\Administrator> bcdedit /set testsigning on
```

- iii. Reboot the machine for the changes to take effect.

3.2.2. GUI mode (Installer)

- i. Run the **ChelsioUwire_x.x.x_x_WIN_yyy.y.yy.exe** installer application.
- ii. Click **Next** for the Chelsio End User License Agreement Window.

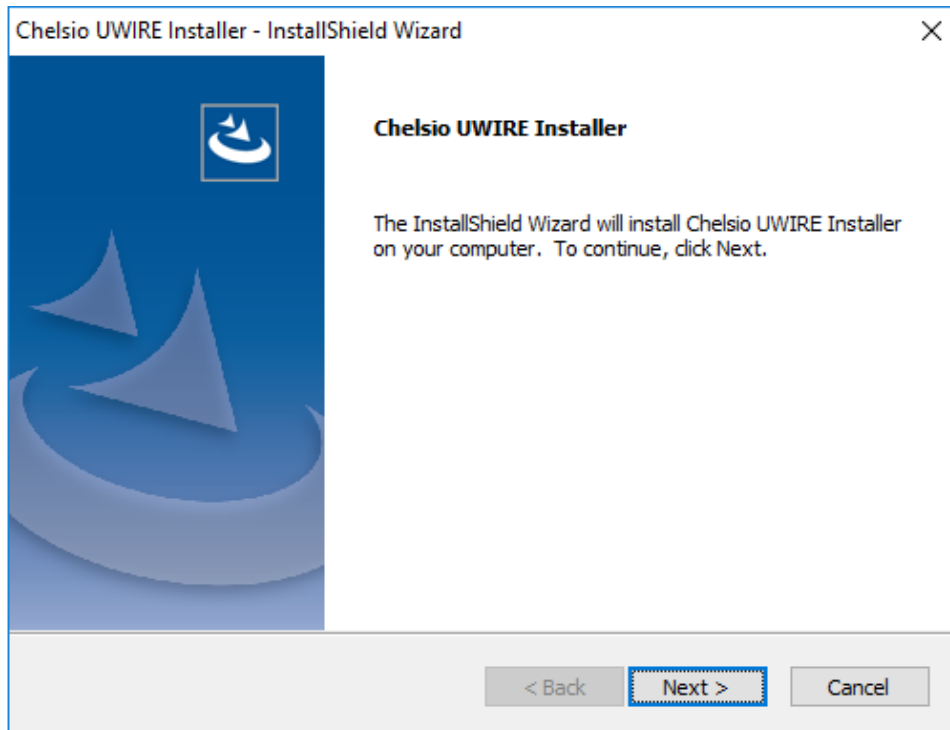


Figure 1 - Unified Wire installer welcome window

- iii. Select the radio button **I accept the terms of the license agreement** and click **Next**

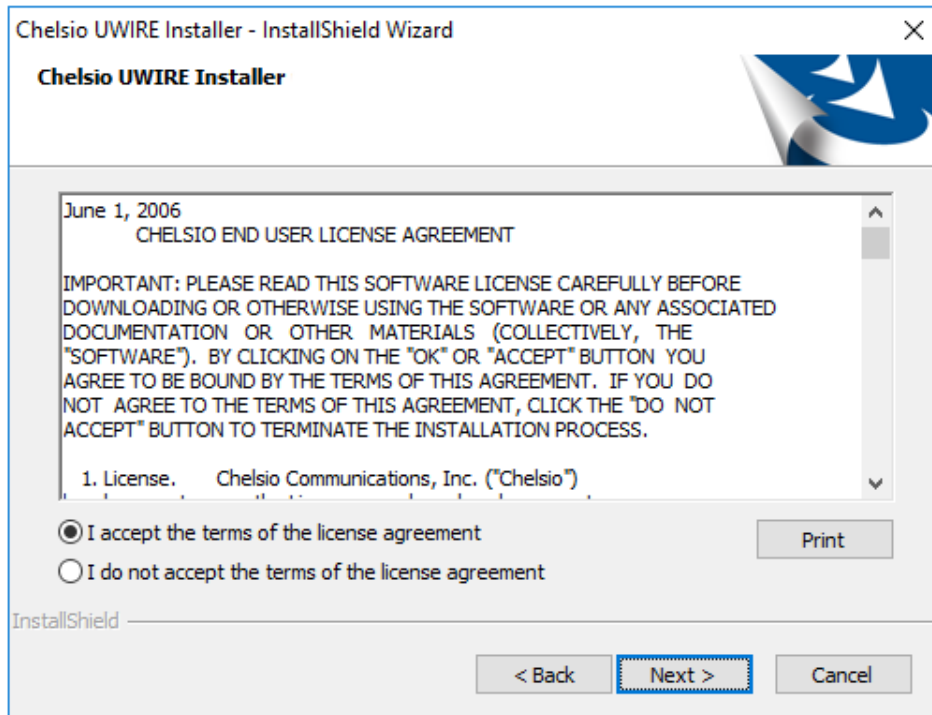


Figure 2 - Chelsio EULA window

- iv. Now, either select **Complete** for complete package installation or else select **Custom** radio button to customize the installation.




Figure 3 - Select setup (installation) type

Complete option will not install iSCSI, iSER and NVMe Storport Miniport drivers. Use the **Custom** option to install them. Virtual Bus Driver (VBD) and NDIS (NIC) drivers are required for iSER and NVMe Storport Miniport drivers to work, whereas iSCSI Storport Miniport needs VBD to function.

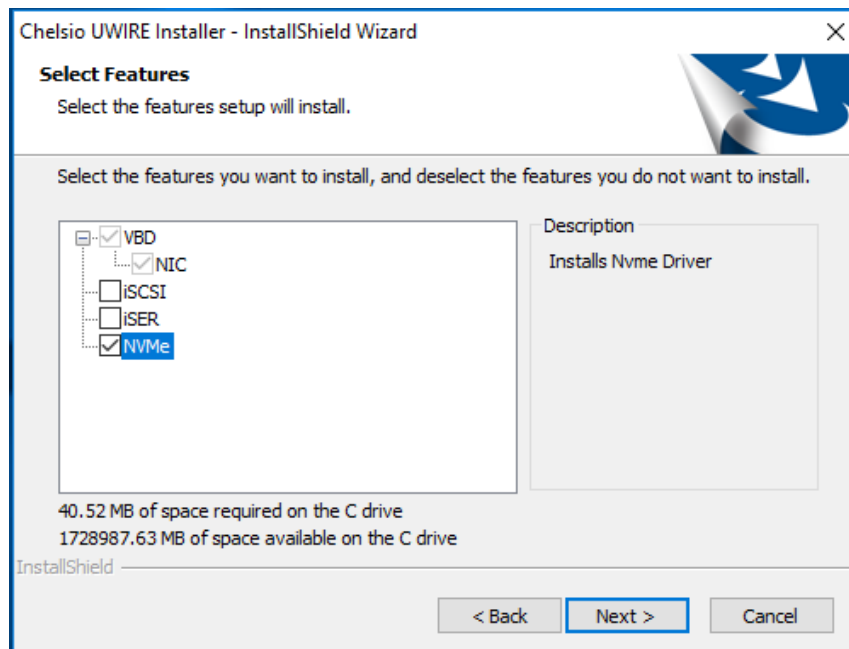


Figure 4 – Custom installation

- v. Click **Install** to start the installation.




Figure 5 - Start installation

- vi. Selected driver components will now be installed.



Figure 6 - Installation in progress

vii. Click **Finish** to exit from the Unified Wire Installer.




Figure 7 - Finishing Unified Wire installation

3.2.3. CLI mode (Silent Installation)

To install drivers/software using this feature, open **command prompt** and execute the following command:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -in <driver(s)>
```

E.g.:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -in all
```

The above command will install all drivers, except iSCSI, iSER and NVMe Storport Miniport drivers.

To install a particular driver (network driver in this case), execute the following command:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -in nic
```

To add a driver to an existing list of already installed drivers, use the following command:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -add  
<driver(s)>
```

E.g.: To add iSCSI Storport Miniport driver,

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -add iSCSI
```

To know more about other parameters and options, execute the following command:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -help
```

The above command will create a help file, *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy_help.txt*, in the present working directory which contains the complete list of command line syntax required for performing all the necessary CLI operations.

Note A log file, *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.log*, is created in the same directory which keeps a record of all the commands executed and their results.

3.3. Zip Package

Chelsio Unified Wire zip package provides two methods of installation: **Manual** and using **Offline Windows Image**. In Manual method, the driver components must be individually installed using the Device Manager. Using the second method, you can add Chelsio driver components to an offline OS image. The image can then be used to install Windows on your system.

3.3.1. Enabling Test Signing

In case the drivers in the package are not WHQL certified, please follow the steps mentioned below to enable test signing before installing Unified Wire:

- i. Open command prompt and run the following command to enable test signing:

```
C:\Users\Administrator>bcdedit /set testsigning on
```

- ii. Change your working directory to *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy/Selfsign/* and run the following commands:

```
C:\ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\Selfsign> .\certmgr.exe -add  
.\chelsiocert1.cer -s -r LocalMachine root -all  
C:\ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\Selfsign> .\certmgr.exe -add  
.\chelsiocert1.cer -s -r LocalMachine trustedpublisher -all
```

- iii. Reboot the machine for the changes to take effect.

3.3.2. Drivers location

The driver files should be manually picked from:

- *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\drivers\Win19* for Server 2019.
- *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\drivers\Win10* for Server 2016 and 10 Client.
- *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\drivers\Win7* for Server 2012 R2.

3.3.3. Manual Installation

- ! Important**
- *The driver components need to be installed strictly in the order described here. Using any other order is not supported.*
 - *Before proceeding, open Device Manager and disable NDIS Miniport driver (Network adapters->Chelsio Network Adapter) followed by Virtual Bus Driver (System devices->Chelsio Bus Enumerator).*

• Virtual Bus Driver

- Open **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **System Devices**, right click on **Chelsio Bus Enumerator** and select **Update Driver Software**.
- Select **Browse my computer for driver software**.
- Select **Let me pick from a list of device drivers on my computer** and click **Next**.
- Click on **Have Disk** Button and on the next screen browse for **chvbdx64.inf** file and click **Open** and then **OK**.
- Click **Next** and driver installation will progress. Click **Close** once the installation is complete.

• NDIS Miniport driver

- Open **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **Network Adapters**, right click on the **Chelsio network adapter** and select **Update Driver Software**.
- Select **Browse my computer for driver software**.
- Select **Let me pick from a list of device drivers on my computer**.
- Click on **Have Disk** Button and on the next screen browse for **chnetx64.inf** and click **Open** and then **OK**.
- Click **Next** and driver installation will progress. Click **Close** once the installation is complete.

Repeat the above steps for the other adapters.

• iSCSI Storport Miniport driver

- Open **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **System Devices**, right click on the **Chelsio Bus Enumerator** and select **Properties**.
- Under **Advanced** tab, select **iSCSI Instances** and set the *Value* to **default**. Click **OK**.
- Under **Other Devices**, select **Chelsio iScsi Function on Port#00**, right-click and select **Update Driver Software**.
- Select **Browse my computer for driver software**.

- v. Select **Let me pick from a list of device drivers on my computer** and click **Next**.
- vi. Click on **Have Disk** Button and on the next screen, browse for **cht4iscsi.inf** file and click **Open** and then **OK**.
- vii. Click **Next** and driver installation will progress. Click **Close** once the installation is complete.

Repeat the same procedure for other ports.

- **iSER Storport Miniport driver**

- i. Open **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **System Devices**, right click on the **Chelsio Bus Enumerator** and select **Properties**.
- ii. Under **Advanced** tab, select **iSER Instances** and set the *Value* to **default**. Click **OK**.
- iii. Under **Other Devices**, select **Chelsio iSer Function on Port#00**, right-click and select **Update Driver Software**.
- iv. Select **Browse my computer for driver software**.
- v. Select **Let me pick from a list of device drivers on my computer** and click **Next**.
- vi. Click on **Have Disk** Button and on the next screen, browse for **chiserx64.inf** file and click **Open** and then **OK**.
- vii. Click **Next** and driver installation will progress. Click **Close** once the installation is complete.

Repeat the same procedure for other ports.

- **NVMe Storport Miniport driver**

- i. Open **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **System Devices**, right click on the **Chelsio Bus Enumerator** and select **Properties**.
- ii. Under **Advanced** tab, select **NVMeOF Instances** and set the *Value* to **default**. Click **OK**.
- iii. Under **Other Devices**, select **Chelsio NVMeOF Function on Port#00**, right-click and select **Update Driver Software**.
- iv. Select **Browse my computer for driver software**.
- v. Select **Let me pick from a list of device drivers on my computer** and click **Next**.
- vi. Click on **Have Disk** Button and on the next screen, browse for **chnvmfx64.inf** file and click **Open** and then **OK**.
- vii. Click **Next** and driver installation will progress. Click **Close** once the installation is complete.

Repeat the same procedure for other ports.

- **Generic Function**

- i. Open **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **System Devices**, right click on **Chelsio Generic Function** and select **Update Driver Software**.
- ii. Select **Browse my computer for driver software**.
- iii. Select **Let me pick from a list of device drivers on my computer**.
- iv. Click on **Have Disk** Button and on the next screen browse for **chnulx64.inf** file and click **Open** and then **OK**.
- v. Click **Next** and driver installation will progress. Click **Close** once the installation is complete.

3.3.4. Offline Windows Image

You can use Windows Deployment Image Servicing and Management (DISM) tool to add Chelsio driver components (.inf files) to an offline Windows image. For more information, visit the following links:

- <http://technet.microsoft.com/en-us/library/hh825070.aspx>
- <http://blogs.technet.com/b/heyscriptingguy/archive/2012/09/27/use-the-powershell-dism-commandlets-to-manage-windows-8.aspx>

4. Firmware Update

The driver will auto-load the firmware if an update is required. The firmware version can be verified using *cxgbtool*:

```
C:\Users\Administrator>cxgbtool nic0 debug version
```

In case of installer, the firmware binaries will be copied to `<system_drive>\ChelsioUwire\Firmware\` folder during installation.

For zip package, the binaries will be present in `ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\firmware\` folder.

5. Software/Driver Update

Chelsio Unified Wire can be updated using Windows update, installer or zip package. The update method for each is described below:

Important *If you are planning to upgrade Unified Wire in a cluster, it should be done on all the members. Using different versions of Unified Wire is not supported.*

5.1. Windows Update

- i. Make sure that you have an active internet connection.
- ii. Open **Device Manager**, click on **System Devices**, right-click on the **Chelsio Bus Enumerator**, and then select **Update Driver Software..**

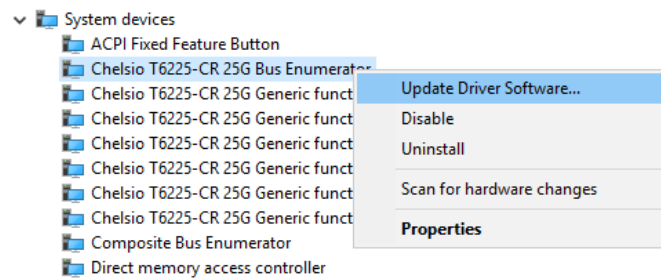


Figure 8 – Updating VBD driver

- iii. Select **Search automatically for updated driver software**.

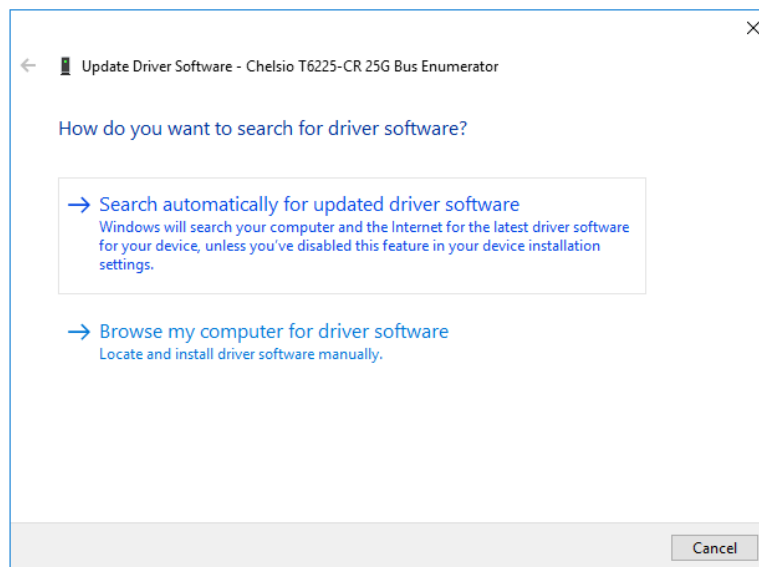


Figure 9 – Searching for driver using Windows update

iv. Windows will search and install the latest driver. Click **Close**.

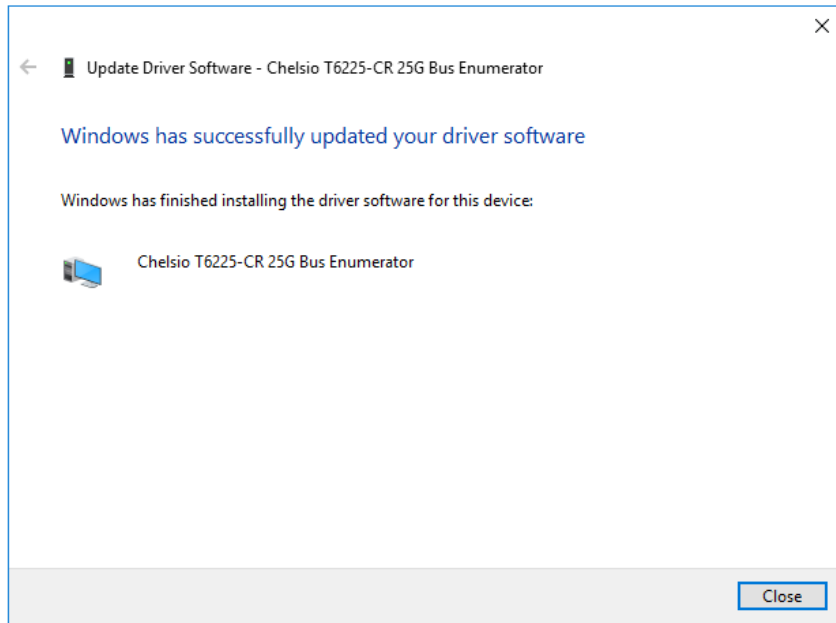


Figure 10 - Driver update successful

v. Click on **Network Adapters**, right-click on **Chelsio Network Adapter** and select **Update Driver Software..**

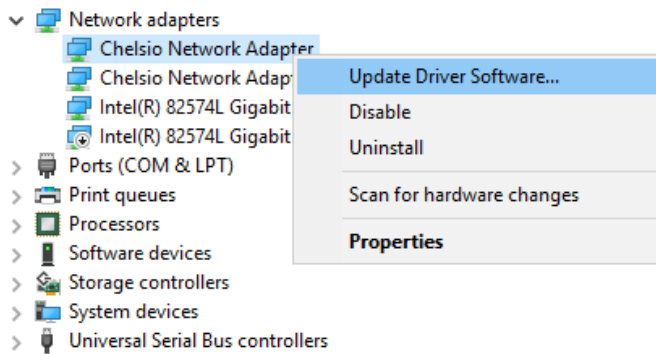


Figure 11 - Updating NDIS driver

vi. Select **Search automatically for updated driver software**

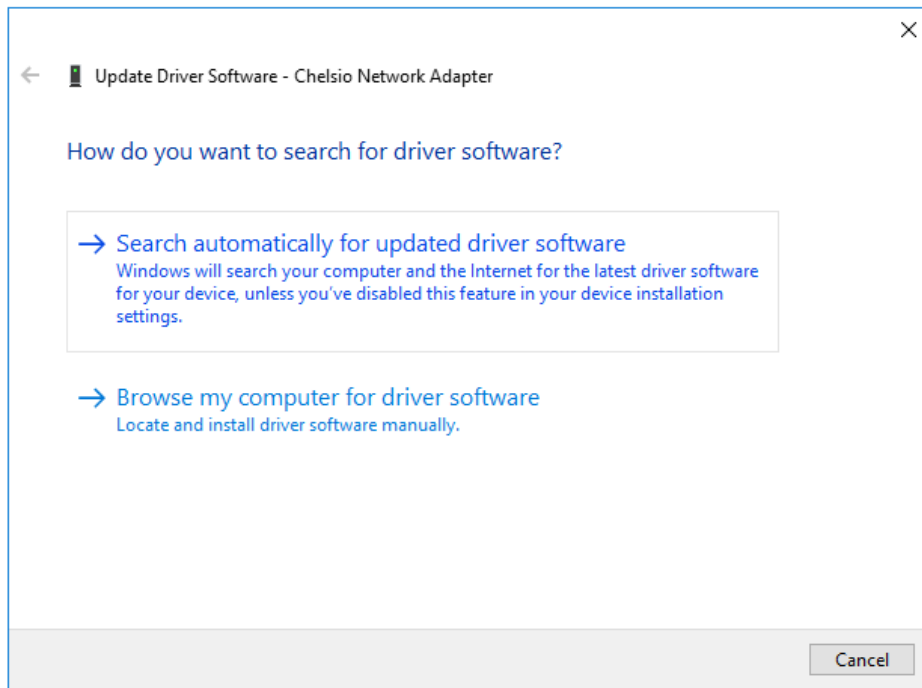


Figure 12- Searching for driver using Windows update

vii. Windows will search and install the latest driver. Click **Close**.




Figure 13 - Driver update successful

viii. Click on **Storage controllers**, right-click on **Chelsio T6 iSCSI Initiator** and select **Update Driver Software..**

Note *If Chelsio iSCSI interface is not listed, verify if you have specified the number of iSCSI instances in [VBD parameters](#).*




Figure 14 - Updating iSCSI Storport Miniport driver

ix. Windows will search and install the latest driver. Click **Close**.

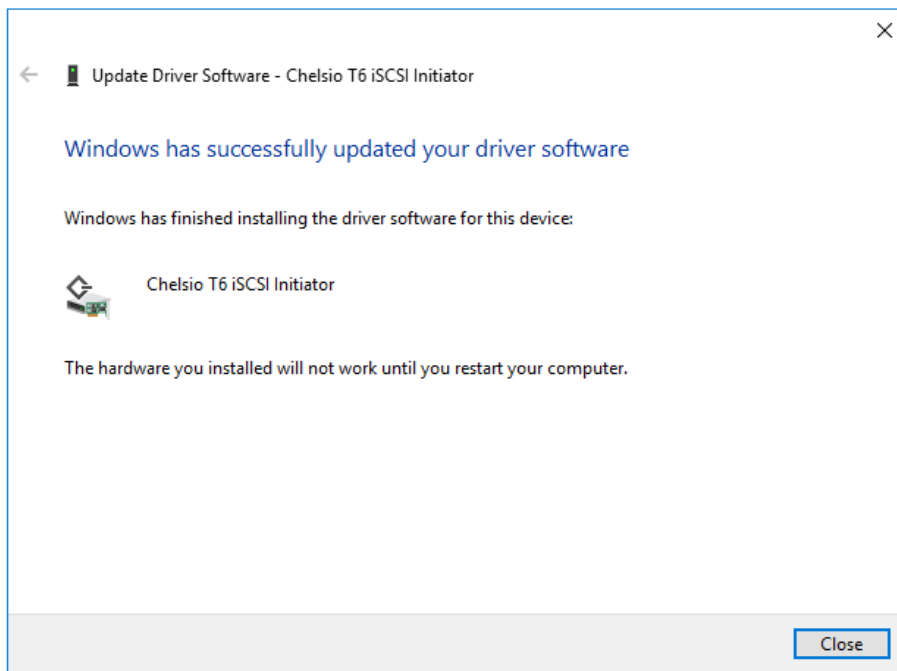


Figure 15 - Driver update successful

- x. Click on **Storage controllers**, right-click on **Chelsio T6 iSER Initiator** and select **Update Driver Software..**

Note *If Chelsio iSER interface is not listed, verify if you have specified the number of iSER instances in [VBD parameters](#).*




Figure 16 - Updating iSER Initiator driver

- xi. Select **Search automatically for updated driver software**

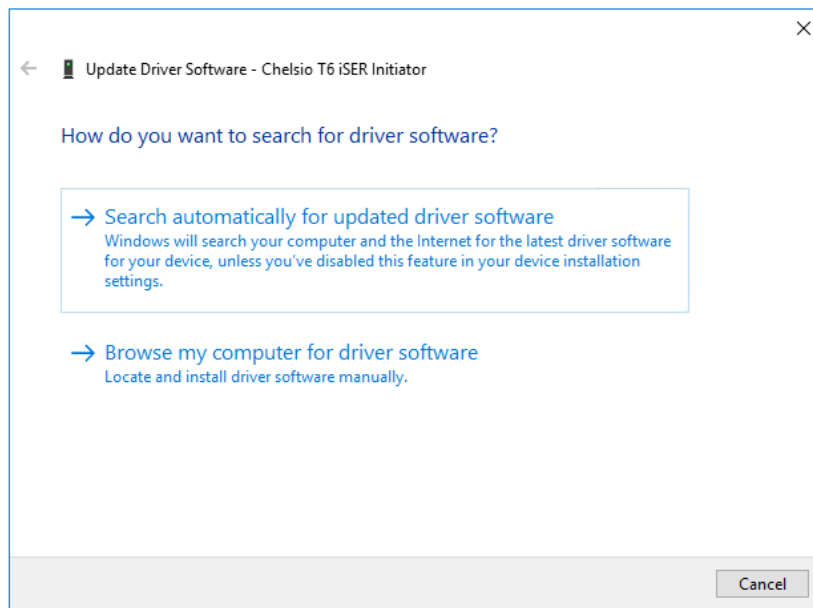


Figure 17- Searching for driver using Windows update

xii. Windows will search and install the latest driver. Click **Close**.




Figure 18 - Driver update successful

xiii. Open the **Device Manager**, click on **System devices**, right-click on **Chelsio Generic Function** and select **Update Driver Software..**

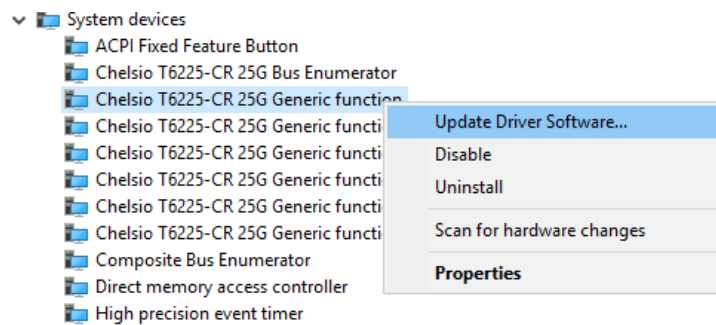


Figure 19 - Updating Generic function driver

xiv. Select **Search automatically for updated driver software**




Figure 20 - Searching for driver using Windows update

xv. Windows will search and install the latest driver. Click **Close**.




Figure 21 - Driver update successful

5.2. Unified Wire Installer

- i. Run the **ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe** installer application.
- ii. Click **Next** to start the update.




Figure 22 - Unified Wire installer welcome window

- iii. Click **Finish** to exit from the Unified Wire Installer.




Figure 23 - Finishing Unified Wire update

5.3. Zip Package

The method to update Unified Wire using zip package is same as installation. Refer the [installation section](#) for more information.

6. Configuring Chelsio Network Interfaces

In order to test Chelsio adapters' features, it is required to use at least two machines, each with Chelsio's network adapters. These machines can be connected directly (back-to-back) or with a switch.

6.1. Configuring 40G Adapters

You can use the **chelsio_adapter_config.ps1** script to configure Chelsio 40G adapters in any of the following three operational modes:

- **DEFAULT (2X40G)**: 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).
- **SPIDER (4X10G)**: In this mode, port 0 functions as 4 10Gbps links and port 1 is disabled.
- **QSA (2X10G)**: 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).

6.1.1. Windows GUI Machine (Installer & Zip Package)

- i. In case of Installer, the configuration script and `cxgbtool.exe` will be copied to `<system_drive>\Windows\System32` folder during installation.

If you are using the zip package, copy `cxgbtool.exe` from `ChelsioUwire_x.x.x.x_WIN_yyy.y.yy` folder to `<system_drive>\Windows\System32` and change your working directory to `ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\Adapter Configuration`.

- ii. Open PowerShell with administrative privileges.
- iii. Run the adapter configuration script and enter the index of the 40G adapter for which the configuration needs to be updated. Hit [Enter].

Installer:

```
PS C:\Users\Administrator>chelsio_adapter_config.ps1
```

Zip Package:

```
PS D:\ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\Adapter Configuration>
.\chelsio_adapter_config.ps1
```

```
Select the Adapter
  1. T580-S0-CR      S/N:PT39131112
Input : 1
```

Figure 24 - Selecting adapter

- iv. Select *Port settings* (option 3) as the configuration type. Hit [Enter].

```
Choose the configuration type:
  1. NON-SRIOV (Default)
  2. SRIOV
  3. Port settings
Input : 3
```

Figure 25 - Port settings

- v. Select the operation mode and then enter *y* to confirm.

```
Choose the type:
  1. DEFAULT (2 x 40G)
  2. SPIDER (4 x 10G)
  3. QSA (2 x 10G)
Input : 2
Do you want to continue (y/n): y
Successfully updated the selected configuration type.
Verification: Passed
```

Figure 26 - Operation mode

- vi. Open **Device Manager**, click on **System Devices** and right-click on the 40G Chelsio adapter selected in step (iii).
- vii. Select **Disable** and then **Enable** for changes to take effect.

6.2. Assigning IP Address

If your computer is connected to a network with a DHCP server, Chelsio network interfaces will acquire DHCP IPs automatically.

To assign static IPs, follow the steps mentioned below:

- i. Double click on the Network Connections icon and choose the Chelsio card entry and double click it.
- ii. Click on the Properties button from the Local Area Connection X Status.
- iii. Select "Internet Protocol (TCP/IP)" from the list and click on Properties button below it.
- iv. From the Internet Protocol (TCP/IP) Properties window, assign an IP Address (e.g. 192.169.1.10) and subnet mask (e.g. 255.255.255.0).
- v. Click on Ok and close on the other window.
- vi. Check to see if you can ping to some other address on this subnet.

6.3. Deadman Watchdog Timer

In some cases when a port is receiving traffic, if the host machine hits a BSOD or crashes, the port cleanup will not be complete, and the port will send indefinite pause frames. This might cause few switch ports to go into an error state. Using the Deadman Watchdog timer feature, the firmware will bring down the host machine ports after the specified time (user configurable). This feature is disabled by default. It can be configured using the **Set-DeadmanWatchdogTimer.ps1** script.

- i. In case of Installer, the configuration script will be copied to `<system_drive>\Windows\System32` folder during installation.

If you are using the zip package, change your working directory to `ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\DeadmanWatchdogTimer`.

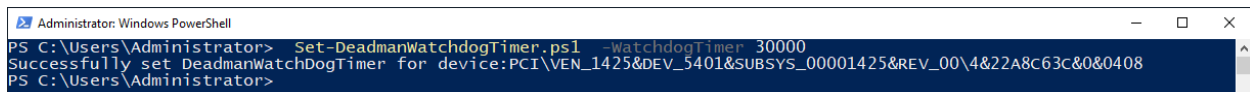
- ii. Open PowerShell with administrative privileges.
- iii. Run the configuration script specifying the time (in milliseconds).

Installer:

```
PS C:\Users\Administrator>Set-DeadmanWatchdogTimer.ps1 -WatchdogTimer
30000
```

Zip Package:

```
PS D:\ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\DeadmanWatchdogTimer> .\Set-
DeadmanWatchdogTimer.ps1 -WatchdogTimer 30000
```



```
Administrator: Windows PowerShell
PS C:\Users\Administrator> Set-DeadmanWatchdogTimer.ps1 -watchdogTimer 30000
Successfully set DeadmanWatchdogTimer for device:PCI\VEN_1425&DEV_5401&SUBSYS_00001425&REV_00\4&22A8C63C&0&0408
PS C:\Users\Administrator>
```

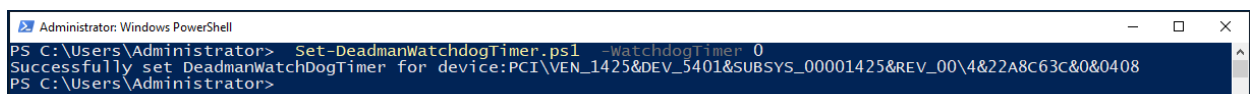
Figure 27- Setting Deadman Watchdog Timer

The command will set the Deadman Watchdog Timer value for all the adapters in the machine. To set only to a specific adapter, specify the Device Path using the **-DeviceId** option.

Note For more help and detailed examples, execute the below command:

```
PS C:\Users\Administrator>get-help Set-DeadmanWatchdogTimer.ps1 -full
```

- iv. To disable the feature, set 0 as the timer value.



```
Administrator: Windows PowerShell
PS C:\Users\Administrator> Set-DeadmanWatchdogTimer.ps1 -watchdogTimer 0
Successfully set DeadmanWatchdogTimer for device:PCI\VEN_1425&DEV_5401&SUBSYS_00001425&REV_00\4&22A8C63C&0&0408
PS C:\Users\Administrator>
```

Figure 28 - Disabling Deadman Watchdog Timer

6.4. Chelsio Management

Chelsio Management Feature allows to update the following properties of Chelsio adapters and to collect Debug logs.

- FW Config file
- Serial Init
- Performance tuning

Open the **Device Manager**, click on the **System Devices**, double click on the **Chelsio Bus Enumerator**, and then click the **Chelsio Management** tab.

- **FW Config file**

Appropriate Firmware Configuration file will already be used by the VBD driver. To use a custom Firmware Configuration file, click **Browse** in **Change Firmware Config** section and select the new FW Configuration file.

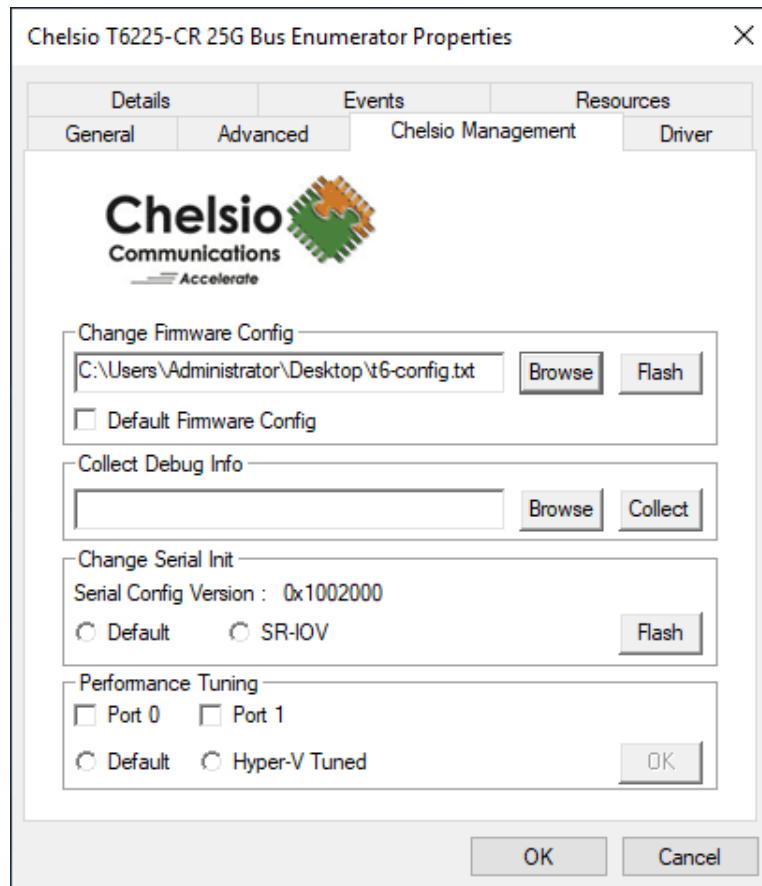


Figure 29 - Changing Firmware Config

Click **Flash** and then **OK** when asked for Confirmation.




Figure 30 - Confirmation

Click **OK** on the Chelsio Management Window for the changes to take effect. The VBD driver will reload and use the newly flashed Firmware Configuration.

To revert to default FW Configuration (the one used by Driver), select **Default Firmware Config** and click **Flash**.

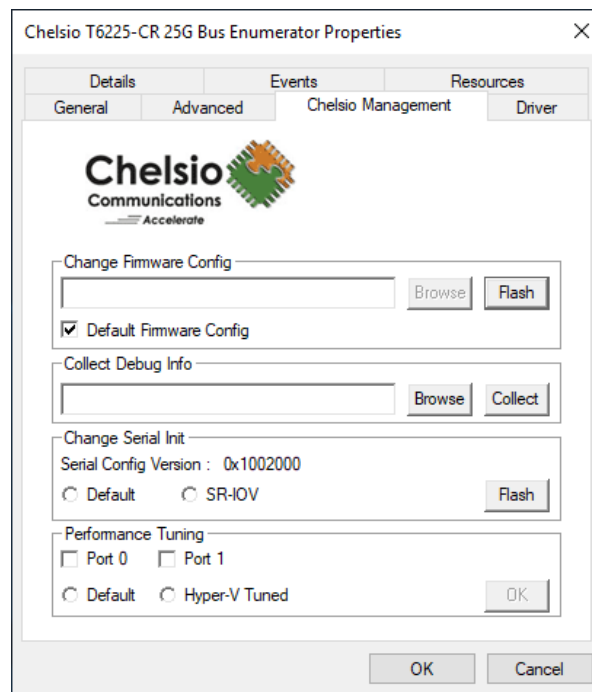


Figure 31 - Default Firmware Config

Click **OK** when asked for Confirmation.




Figure 32 - Confirmation

Click **OK** on the Chelsio Management Window for the changes to take effect. The VBD driver will reload and use the default Firmware Configuration.

- **Change Serial Init**

The adapter configuration needs to be updated to use NDIS SR-IOV features. To update it, select **SR-IOV** in Change Serial Init section and click **Flash**.

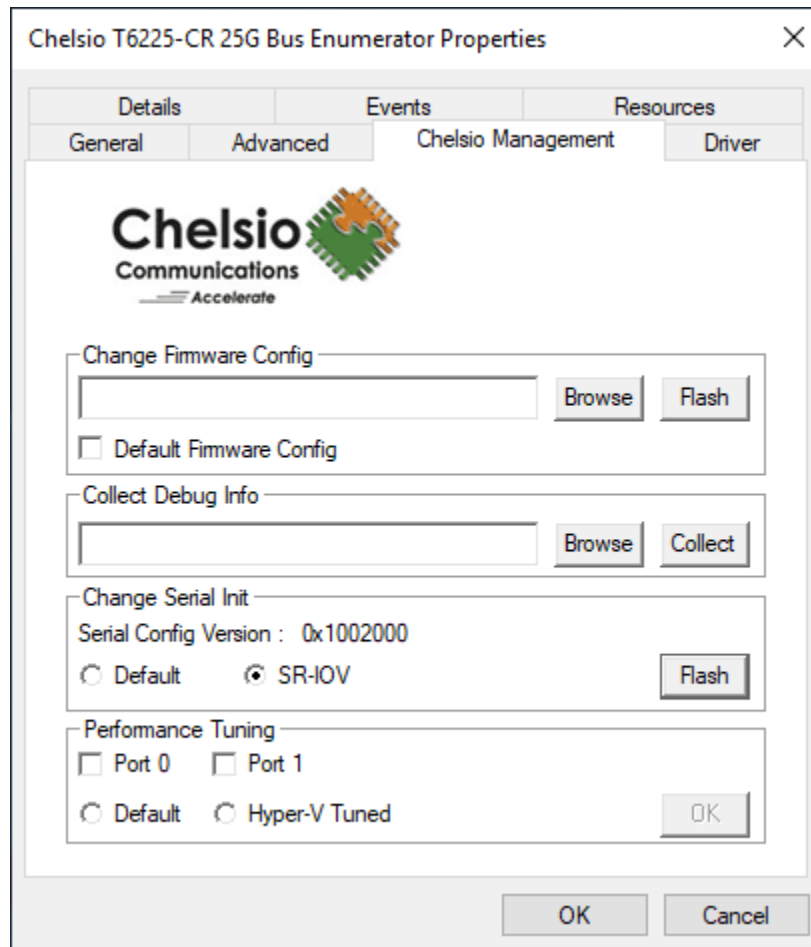


Figure 33 - Update Serial Init to SR-IOV

Click **OK** when asked for Confirmation.

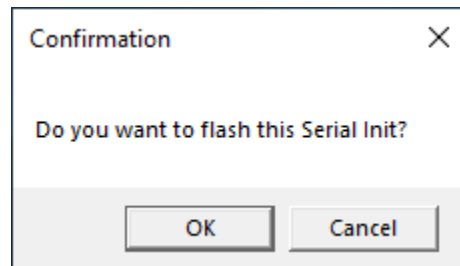


Figure 34 - Confirmation

Click **OK** on the Chelsio Management Window and reboot the machine for changes to take effect.

To revert to Default settings (non SR-IOV), select **Default** in Change Serial Init section and click **Flash**.

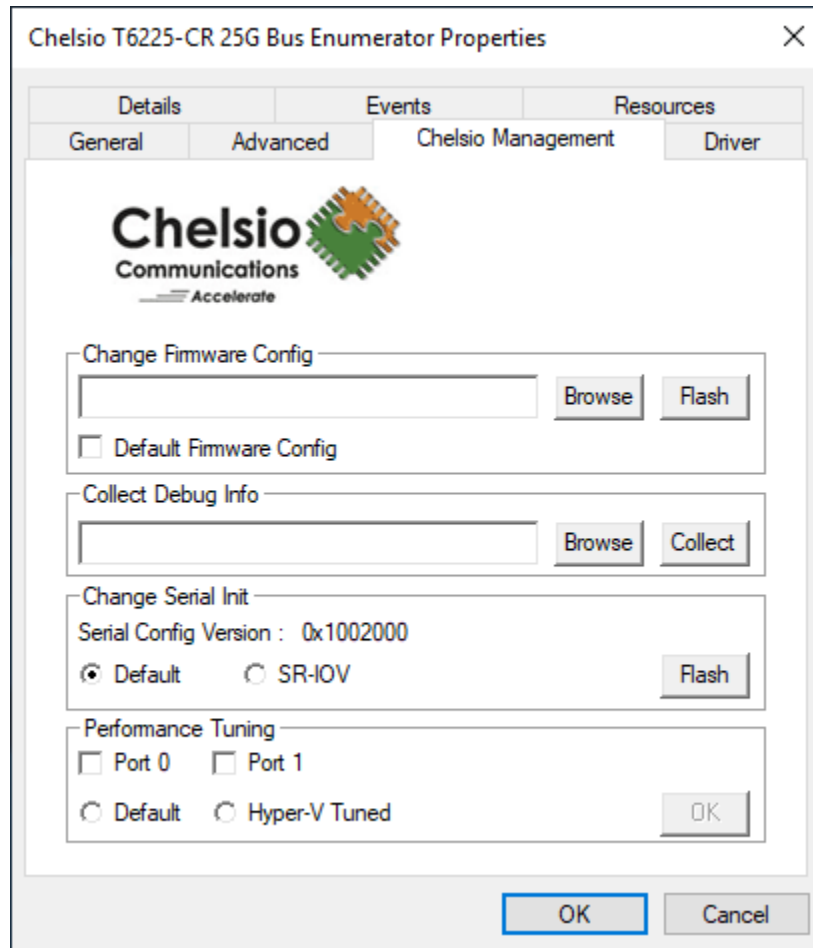


Figure 35 - Update to Default Serial Init

Click **OK** when asked for Confirmation.




Figure 36 - Confirmation

Click **OK** on the Chelsio Management Window and reboot the machine for changes to take effect.

- **Performance Tuning**

To tune the adapter for best Performance, select the required Ports and select **Hyper-V Tuned** in Performance Tuning section. Click **OK**

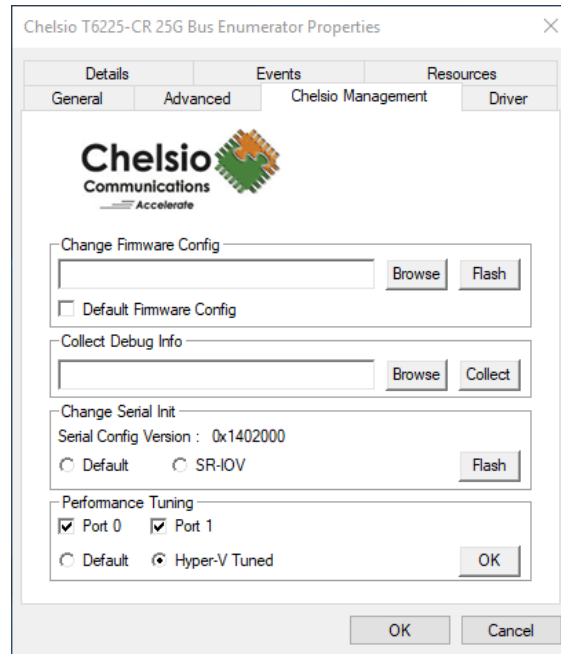


Figure 37 - Update Performance tuning

To revert the Performance tuning settings, select the required Ports and click **Default** in Performance Tuning section. Click **OK**

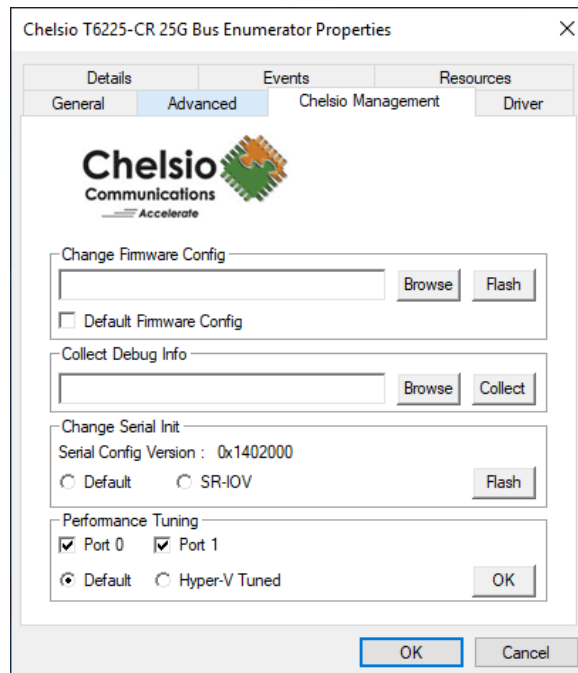


Figure 38 - Default Performance tuning

- **Collect Debug Info**

In case of any Driver/Firmware issue, Chelsio hardware debug information needs to be collected and the same should be provided to Chelsio for further analysis. To collect the required debug logs from the adapter, click **Browse** in Collect Debug Info section and select the destination folder.

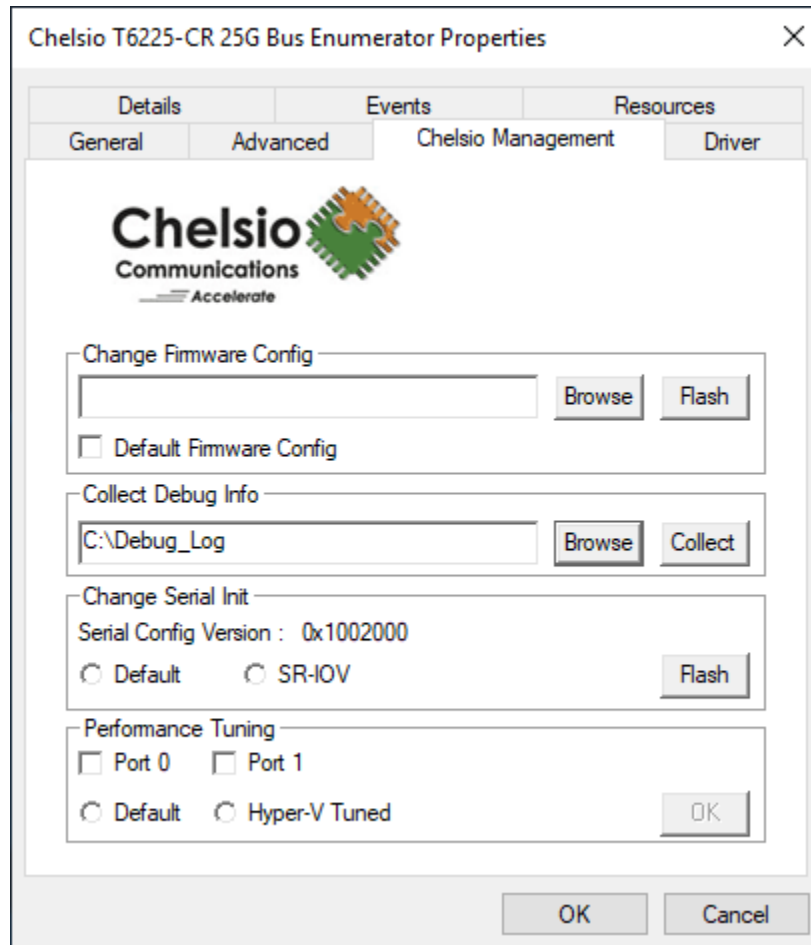



Figure 39 - Collecting Debug Information

Click **Collect**. The Logs will be collected in the selected folder with name *ChelsioDebug_<ddMMyyyHHmmss>.DMP*

7. Performance Tuning

The following section lists the various steps to tune the system for optimal performance:

- Install 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.
- Disable virtualization, c-state technology, VT-d, Intel I/O AT and SR-IOV in the BIOS settings.

 **Note** *NDIS SR-IOV will require virtualization, VT-d and SR-IOV to be enabled in the BIOS.*

- Set **Power Options** in Control Panel to *High Performance*.
- Tune the adapter for best performance using Chelsio Management feature as explained [here](#).

8. *cxgbtool* help

The *cxgbtool* command queries or sets various aspects of Chelsio network interface cards. It complements standard tools used to configure network settings and provides functionality not available through such tools.

Some of the commands provided can be used to query running statistics to aid in debugging.

Syntax

```
cxgbtool <vbdIface|nicIface|iSCSIIface|rdma> <category> <command> <options>
```

Definitions

nicIface is the name of the network device to work on, given in the format *nic0*, *nic1*.

vbdIface is the name of the Chelsio Bus Enumerator instance to work on, given in the format *vbd0*, *vbd1*

iscsIface is the name of the Chelsio iSCSI interface to work on, given in the format *iSCSI0*, *iSCSI1*

rdma is the name of the Chelsio RDMA interface to work on.

category is one of the *cxgbtool* categories, i.e., *debug*, *hardware*, *cudbg*, *firmware*, *boot*

command is the action you wish to perform on the adapter or property you wish to view/change.

option is the optional parameter to be provided with *command*.

- **Help**

- Displaying *cxgbtool* help: `cxgbtool`

- Displaying list of available *cxgbtool* commands: `cxgbtool -h all`

- Displaying category specific help:

```
cxgbtool <vbdIface|nicIface|iSCSIIface> <category> -h
```

- Displaying list of *cudbg* entities: `cxgbtool <vbdIface> cudbg dbg -h`

- Displaying RDMA help: `cxgbtool rdma -h`

- **Debug**

- **version**

Description: Display adapter information like part number, serial number, device ID, firmware version, TP, etc.

Syntax: cxgbtool <vbdIface|nicIface|iscsiIface> debug version

Example:

```
C:\Users\Administrator>cxgbtool nic0 debug version
P/N:          11012096004
S/N:          PT43160304
DeviceId:     6401 (T6225-CR)
FW Ver:       1.16.38.0
TP Ver:       0.1.23.2
Driver Ver:   6.5.4.10
Nic File Ver: 6.5.4.10
VBD File Ver: 6.5.4.10
Cudbg lib Ver: 1.14.0
SCFG Ver:     0x1402000
```

Figure 40 - Display adapter information

- **inst**

Description: Install driver package specified in the setup information file (.inf).

Syntax: cxgbtool debug inst <file>

Example:

```
C:\Users\Administrator>cxgbtool debug inst Desktop\chvbdx64.inf
ENTER: DriverPackageInstallA
ENTER: DriverPackageInstallW
Installing INF file 'C:\Users\Administrator\Desktop\chvbdx64.inf' (Plug and Play).
Looking for Model Section [Chelsio.NTamd64]...
Installing devices with Id "PCI\VEN_1425&DEV_5410&SUBSYS_00001425&REV_00" using INF "C:\Windows\System32\DriverStore\FileRepository\chvbdx64.inf_amd64_db5d99a4eb08b3f\chvbdx64.inf".
Installation did not occur because the hardware isn't currently present.
Installing devices with Id "PCI\VEN_1425&DEV_6401&SUBSYS_00001425&REV_00" using INF "C:\Windows\System32\DriverStore\FileRepository\chvbdx64.inf_amd64_db5d99a4eb08b3f\chvbdx64.inf".
ENTER: UpdateDriverForPlugAndPlayDevices...
RETURN: UpdateDriverForPlugAndPlayDevices...
Installation was successful.
Marked Phantom Device with Hardware/Compatible Id 'PCI\VEN_1425&DEV_5410&SUBSYS_00001425&REV_00' for reinstall on next plug-in.
Install completed
RETURN: DriverPackageInstallW (0x0)
RETURN: DriverPackageInstallA (0x0)
```

Figure 41 - Install driver package

- **update**

Description: Update driver package.

Syntax: cxgbtool debug update <file>

Example:

```
C:\Users\Administrator>cxgbtool debug update C:\Users\Administrator\Desktop\v6.5.8.0\bin\fre\x64\chvbdx64.inf
Driver for HwId:PCI\VEN_1425&DEV_6407&SUBSYS_00001425&REV_00 updated successfully!!!
```

Figure 42 - Update driver package

- **uninst**

Description: Uninstall driver package specified in the setup information file (.inf).

Syntax: cxgbtool debug uninst <file>

Example:

```
C:\Users\Administrator>cxgbtool debug uninst Desktop\chvbdx64.inf
Removing device instance PCI\VEN_1425&DEV_6401&SUBSYS_00001425&REV_00\4&2A43D483&080410
Removing device instance PCI\VEN_1425&DEV_5410&SUBSYS_00001425&REV_00\4&31024D2F&080418
ENTER: DriverPackageGetPathA
RETURN: DriverPackageGetPathA (0x7A)
ENTER: DriverPackageGetPathA
RETURN: DriverPackageGetPathA (0x0)
ENTER: DriverPackageUninstallA
ENTER: DriverPackageUninstallW
Uninstalling driver package C:\Windows\System32\DriverStore\FileRepository\chvbdx64.inf_amd64_db5d99a4eb080b3f\chvbdx64.inf...
Successfully uninstalled 'C:\Windows\INF\oem6.inf'.
No devices found for C:\Windows\System32\DriverStore\FileRepository\chvbdx64.inf_amd64_db5d99a4eb080b3f\chvbdx64.inf uninstall.
Successfully deleted properties for driver store entry 'C:\Windows\System32\DriverStore\FileRepository\chvbdx64.inf_amd64_db5d99a4eb080b3f\chvbdx64.inf'.
Uninstall completed.
RETURN: DriverPackageUninstallW (0x0)
RETURN: DriverPackageUninstallA (0x0)
```

Figure 43 - Uninstall driver package

- **rescan**

Description: Scan for hardware changes in the device manager.

Syntax: cxgbtool debug rescan

Example:

```
C:\Users\Administrator>cxgbtool debug rescan
```

Figure 44 - Rescan for hardware changes

- **qstats**

Description: Display statistics for each Tx & Rx queue.

Syntax: cxgbtool <nicIface|iscsiIface> debug qstats <queueType|clr>

clr: Clear queue statistics.

Queue Type	Description
<i>txeth</i>	Tx tunnel queue statistics.
<i>rxeth</i>	Rx tunnel queue statistics.
<i>txvirt</i>	Tx VM queue statistics.
<i>rxvirt</i>	Rx VM queue statistics.
<i>txtoe</i>	Tx Chimney queue statistics.
<i>rxtoe</i>	Rx Chimney queue statistics.
<i>txrdma</i>	Tx RDMA queue statistics.
<i>rxrdma</i>	Rx RDMA queue statistics.
<i>txctrl</i>	Chimney control queue statistics.
<i>txfwd</i>	Chimney forwarding queue statistics.
<i>txnvgre *</i>	Tx NVGRE statistics.
<i>rxnvgre *</i>	Rx NVGRE statistics.

<i>txiscsi</i>	Tx iSCSI queue statistics.
<i>rxiscsi</i>	Rx iSCSI queue statistics.
<i>txpd</i>	Tx PacketDirect queue statistics.
<i>rxpd</i>	Rx PacketDirect queue statistics.

* not supported on T6 adapters

Qstats Parameter	Description
<i>BcBytes</i>	Broadcast packet size in bytes.
<i>BcPkts</i>	Number of broadcast packets.
<i>CoalBytes</i>	Coalesced packet size in bytes.
<i>CoalPkts</i>	Number of coalesced packets.
<i>DropPkts</i>	Number of dropped packets.
<i>EncapBytes</i> #	NVGRE/VXLAN packet size in bytes.
<i>EncapPkts</i> #	Number of NVGRE/VXLAN packets.
<i>VlanMismatch</i>	Number of packets with VLAN mismatch.
<i>Fw4Ack</i>	Number of SGE update requests.
<i>LargeRxDropPkts</i>	Dropped packets due to large size counter.
<i>LSO</i>	Number of large send offloaded packets.
<i>McBytes</i>	Multicast packet size in bytes.
<i>McPkts</i>	Number of multicast packets.
<i>NblsRecv</i>	Total Nbls received.
<i>NblsSent</i>	Total Nbls sent.
<i>NoRxBufs</i>	Packets couldn't process because of lack of memory.
<i>NvPkts</i>	Send NVGRE packets.
<i>PeerAbort</i>	Peer abort request.
<i>CqComp</i>	Number of CQ completions.
<i>PktsDefrag</i>	Number of defragged packets.
<i>PktsDrop</i>	Total packets dropped.
<i>PktsQued</i>	Total packets queued.
<i>PktsSent</i>	Total packets sent.
<i>Quefull</i>	Number of time queue full happened.
<i>RecvNbls</i>	Total Nbls received.
<i>RssGroup</i>	RSS group number allocated to queue.
<i>VPortId</i>	Vport ID of queue.
<i>RssProc</i>	RSS processor number allocated to queue.
<i>RxCsumErr</i>	Number of received packets with checksum error.
<i>RxCsumErrIp</i>	Number of received IP packets with checksum error.
<i>RxCsumErrTcp</i>	Number of received TCP packets with checksum error.
<i>SysThreadEvts</i>	Number of system thread events.
<i>TxCsumOfld</i>	Number of checksum offloaded Tx packets.
<i>UcBytes</i>	Unicast packets size in bytes.
<i>UcPkts</i>	Number of unicast packets.
<i>VlanEx</i>	Number of VLAN extracted packets.
<i>VLANin</i>	Number of VLAN inserted packets.

supported only on T6 adapters

Example of rxeth qstats

```
C:\Users\Administrator>cxgbtool nic0 debug qstats rxeth
```

StatsType	RxEth#36	RxEth#37	RxEth#38	RxEth#39
UcPkts	21	3	3	1
UcBytes	1324	260	239	93
McPkts	0	2	3	0
McBytes	0	180	258	0
BcPkts	6	0	45	0
BcBytes	360	0	4572	0
CoalPkts	0	0	0	0
CoalBytes	0	0	0	0
DropPkts	0	0	0	0
RecvNbls	27	5	51	1
RxCsumErr	0	0	0	0
RxCsumErrTcp	0	0	0	0
RxCsumErrIp	0	0	0	0
VLANex	0	0	0	0
SysThreadEvs	0	0	0	0
LargeRxDropPkts	0	0	0	0
RssProc	0	16	24	32
RssGroup	0	0	0	0

Figure 45 - Queue Statistics

- **hardware**

- **sensor**

Description: Display sensor data.

Syntax: cxgbtool <nicIface> hardware sensor

Example:

```
C:\Users\Administrator>cxgbtool nic0 hardware sensor
Temperature: 80C
Core VDD: 855mV
```

Figure 46 - Sensor data

- **fec**

Description: Set/Display Forward Error Correction (FEC) parameters.

Syntax:

```
cxgbtool <nicIface> hardware fec [<config params>]
```

Example:

- Displaying currently set FEC mode.

```
C:\Users\Administrator>cxgbtool nic0 hardware fec
supported: : Base-R/Reed-Solomon Reed-Solomon
advertising:: no FEC
```

Figure 47 - Read FEC

- Setting FEC mode:

```
C:\Users\Administrator>cxgbtool nic0 hardware fec rs
C:\Users\Administrator>cxgbtool nic0 hardware fec
supported: : Base-R/Reed-Solomon Reed-Solomon
advertising:: Reed-Solomon
```

Figure 48 - Set FEC

▪ link_cfg

Description: Set/Display link configuration parameters.

Syntax:

```
cxgbtool <nicIface> hardware link_cfg <options>
```

Example:

- Displaying link configuration information:

```
C:\Users\Administrator>cxgbtool nic0 hardware link_cfg
Link Configuration
Supported ports           : [ FIBRE ]
Supported link modes     : 1000baseT/Full
                        : 10000baseT/Full
                        : 25000baseCR/Full
Supported auto-negotiation : Yes
Advertised link modes   : 1000baseT/Full
                        : 10000baseT/Full
                        : 25000baseCR/Full
Advertised auto-negotiation : Yes
Speed                   : 25Gbps
port                   : Other
Auto-negotiation       : On
```

Figure 49 - Read Link Configuration

- Enabling/Disabling auto-negotiation:

Auto-negotiation is enabled by default. To disable it, use the *autoneg off* parameter on all connected machines of the link.

```
C:\Users\Administrator>cxgbtool nic0 hardware link_cfg autoneg off
C:\Users\Administrator>cxgbtool nic0 hardware link_cfg
Link Configuration
Supported auto-negotiation : Yes
Advertised auto-negotiation : No
Speed                   : 25Gbps
Auto-negotiation       : Off
Link State             : Disconnected
```

Figure 50 - Disable auto-negotiation

- Changing link speed:
25G port can be set to 10G and 100G port can be set to 25G/40G/50G. Auto-negotiation should be disabled on all connected machines while changing link speed.

```
C:\Users\Administrator>cxgbtool nic0 hardware link_cfg autoneg off speed 10000
C:\Users\Administrator>cxgbtool nic0 hardware link_cfg
Link Configuration
Supported auto-negotiation : Yes
Advertised auto-negotiation : No
Speed : 10Gbps
Auto-negotiation : Off
Link State : connected
```

Figure 51 - Change Link Speed



Important

Before setting 40G or 10G speeds, FEC should be disabled on the port using:

```
C:\Users\Administrator>cxgbtool nic0 hardware fec off
```

- **cudbg**

- **collect**

Description: Collect Chelsio adapter debug logs to a compressed file.

Syntax: cxgbtool <vbdIface> cudbg collect <entities> <outputFile> <options>

Examples:

- Collecting debug log for all entities

```
C:\Users\Administrator>cxgbtool vbd0 cudbg collect all dump_file
cxgbtool: This may take a while. Please be patient
cxgbtool: Writing cudbg block to file dump_file... size: 12634452 bytes
cxgbtool: Done writing cudbg data to file dump_file
```

Figure 52 - Collect all logs

- Collecting debug log skipping specific debug entities

```
C:\Users\Administrator>cxgbtool vbd0 cudbg collect all dump_file skip edc1,edc0
cxgbtool: This may take a while. Please be patient
cxgbtool: Writing cudbg block to file dump_file... size: 12082984 bytes
cxgbtool: Done writing cudbg data to file dump_file
```

Figure 53 - Skip few entities in log collection

- Collecting debug log avoiding entities that can affect running traffic

```
C:\Users\Administrator>cxgbtool vbd0 cudbg collect all dump_file safe
cxgbtool: This may take a while. Please be patient
cxgbtool: Writing cudbg block to file dump_file... size: 48484 bytes
cxgbtool: Done writing cudbg data to file dump_file
```

Figure 54 - Safe mode log collection

9. Software/Driver Uninstallation

Similar to installation, Chelsio Unified Wire can be uninstalled using the Installer or zip package. Refer the relevant section depending on the method of installation used to install drivers.

9.1. Installer

Chelsio Unified Wire Installer can be uninstalled using two methods: **GUI** or **CLI** mode. GUI mode requires user interaction and uninstallation occurs with options specified by the user. Whereas, CLI mode does not require any user input.

9.1.1. GUI mode (Installer)

- i. Run the **ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe** application.
- ii. Select **Modify** to add or remove features. Select **Repair** to repair the previous installation. Select **Remove** to uninstall the application. After you have selected the appropriate option, click **Next**.




Figure 55 - Adding or removing features using 'Modify' option




Figure 56 - Repairing previous installation using 'Repair' option

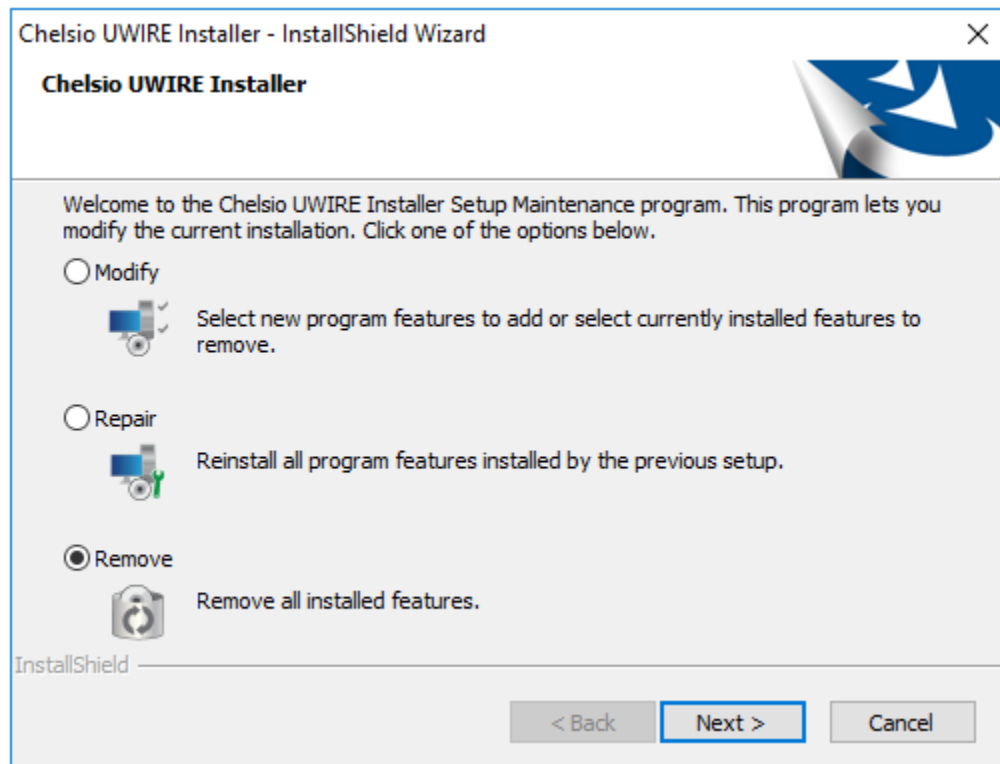


Figure 57 – Uninstalling Unified Wire using 'Remove' option

- iii. Click on the **Finish** button to exit from the installer.

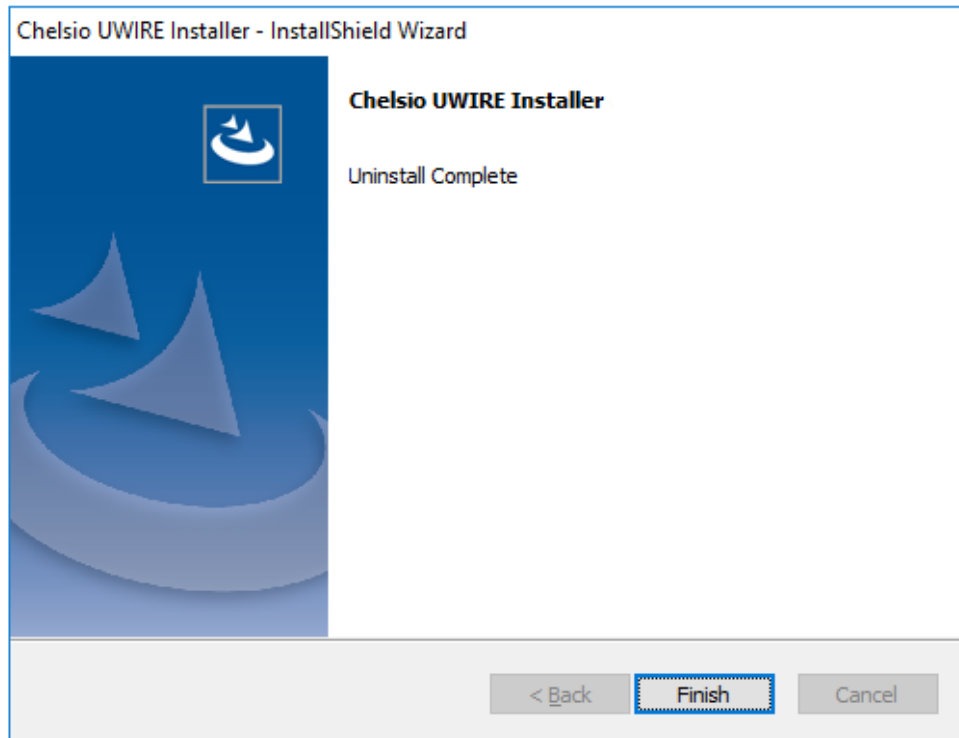


Figure 58 - Finishing uninstallation

Note This method of uninstallation is possible only if the drivers were installed using Unified Wire Installer.

9.1.2. CLI Mode

To uninstall all the drivers, execute the following command:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -un all
```

To uninstall a particular driver, execute the following command:

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -rm <driver(s)>
```

Example: To uninstall iSCSI Storport Miniport driver

```
C:\Users\Administrator>ChelsioUwire_x.x.x.x_WIN_yyy.y.yy.exe -rm iSCSI
```

9.2. Zip Package

If the driver components were installed using the zip package, they will have to be uninstalled manually. The following section lists the various components and their respective methods of uninstallation.

Important *The driver components need to be uninstalled strictly in the order described here. Using any other order is not supported.*

- **NDIS Miniport driver**

- i. Open the **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on the **Network Adapters**, right click on the **Chelsio Network Adapter** and select **Uninstall**.

Note *In case of Virtual Functions in VM, use Chelsio VF Network Adapter.*

- ii. Select the **Delete the driver software for this device** checkbox when **Confirm Device Uninstall** is prompted and click **OK**.

Repeat the same procedure for the other ports.

- **iSCSI Storport Miniport driver**

- i. Open the **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **Storage controllers**, right-click on **Chelsio iSCSI Initiator** and select **Uninstall**.
- ii. Select **Delete the driver software for this device** checkbox when **Confirm Device Uninstall** is prompted and click **OK**.

Repeat the same procedure for other ports.

- **iSER Storport Miniport driver**

- i. Open the **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **Storage controllers**, right-click on **Chelsio iSER Initiator** and select **Uninstall**.
- ii. Select **Delete the driver software for this device** checkbox when **Confirm Device Uninstall** is prompted and click **OK**.

Repeat the same procedure for other ports.


- **NVMe Storport Miniport driver**

- i. Open the **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on **Storage controllers**, right-click on **Chelsio NVMeOF Initiator** and select **Uninstall**.
- ii. Select the **Delete the driver software for this device** checkbox when **Confirm Device Uninstall** is prompted and click **OK**.

Repeat the same procedure for the other ports.

- **Virtual Bus driver**

- i. Open the **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on the **System Devices**, right click on the **Chelsio Bus Enumerator** and select **Uninstall**.

 **Note** *In case of Virtual Functions in VM, use **Chelsio Bus Enumerator [Virtual Function]**.*

- ii. Select the **Delete the driver software for this device** checkbox when **Confirm Device Uninstall** is prompted and click **OK**.

- **Generic Function**

- i. Open the **Device Manager** (Control Panel -> System & Security-> System -> Device Manager), click on the **Network Adapters**, right click on the **Chelsio Generic Function** and select **Uninstall**.
- ii. Select the **Delete the driver software for this device** checkbox when **Confirm Device Uninstall** is prompted and click **OK**.

Repeat the same procedure for the other **Generic Functions**.

II. NDIS Function

1. Introduction

Chelsio’s Terminator series of Unified Wire adapters provide extensive support for NIC operation, including all stateless offload mechanisms for both IPv4 and IPv6 (IP, TCP and UDP checksum offload, LSO - Large Send Offload aka TSO - TCP Segmentation Offload, Network Direct and assist mechanisms for accelerating RSC - Receive Segment Coalescing).

1.1. Hardware Requirements

1.1.1. Supported Adapters

Following is the list of NDIS components and supported Chelsio adapters:

- 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
- T580-OCP-SO
- T540-CR
- T540-LP-CR
- T540-SO-CR
- T540-BT
- T520-CR
- T520-LL-CR
- T520-SO-CR
- T520-OCP-SO
- T520-BT

1.2. Software Requirements

Following is the list of NDIS components and supported Windows versions:

Table 1 – Windows version and supported NDIS components

Windows Version	NDIS Components
Server 2019	NIC, NVGRE, VXLAN, Timestamps, PacketDirect, VMMQ/vRSS
Server 2016	NIC, NVGRE, VXLAN, PacketDirect, VMMQ/vRSS
10 Client	NIC
Server 2012 R2	NIC, NVGRE

2. Software/Driver Configuration and Fine-tuning


2.1. Advanced Configuration

The below sections provide the advanced configuration options for Chelsio devices.

2.1.1. VBD Driver Parameters

To see all tunable VBD driver parameters, open the **Device Manager**, click on the **System Devices**, double click on the **Chelsio Bus Enumerator**, and then click the **Advanced** tab.

- **iSCSI Instances**
 - *Description:* Specify the number of iSCSI instances.
 - *Value:* 0-2 | default
 - *Default:* 0
- **iSER Instances**
 - *Description:* Specify the number of iSER instances.
 - *Value:* 0-2 | default
 - *Default:* 0
- **NIC Instances**
 - *Description:* Specify the number of NIC instances.
 - *Value:* 0-8 | default
 - *Default:* default
- **NVMeoF Instances**
 - *Description:* Specify the number of NVMeoF instances.
 - *Value:* 0-2 | default
 - *Default:* 0

 **Note** Value will be 0-4 | default for T5 4-port adapters.

2.1.2. NDIS Miniport Driver Parameters

To see or change the tunable NDIS Driver parameters, open the **Device Manager**, click on the **Network adapters**, double click on any **Chelsio Network Adapter**, and then click the **Advanced** tab. They can also be changed from CLI using the below commands, especially in case of Server Core installations where there is no GUI.

```
PS C:\Users\Administrator> Get-NetAdapterAdvancedProperty -Name <NetAdapter>
PS C:\Users\Administrator> Set-NetAdapterAdvancedProperty -Name <NetAdapter>
-DisplayName <ParameterName> -DisplayValue <Value>
```

- **Encapsulated Task Offload**

- *Description:* Configure stateless offloads (checksums, LSO, VMQ and filtering) for encapsulated traffic.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **Encapsulated Overhead**

- *Description:* Specify the amount of overhead required in Ethernet frames due to virtual network overlay encapsulation such as VXLAN and NVGRE.
- *Value:*
 - 32
 - 64
 - 96
 - 128
 - 160
 - 192
 - 224
 - 256
- *Default:* 0

- **Flow Control**

- *Description:* Enable or disable the receipt or transmission of PAUSE frames. If Rx is enabled, all incoming PAUSE frames will be honored. If Tx is enabled, the HW will send PAUSE frames when its MAC Rx FIFO usage is beyond a high-watermark. If Disabled option is selected, no PAUSE will be sent. When Rx MAC FIFO gets full, further ingress packets will be dropped. If Rx & Tx Enabled is selected, both are enabled.
- *Value:*
 - Auto Negotiation
 - Disabled
 - Rx & Tx Enabled
 - Rx Enabled
 - Tx Enabled
- *Default:* Auto Negotiation

- **Interrupt Moderation**

- *Description:* Control the interrupts generated by NIC hardware.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **Interrupt Moderation Rate**

- *Description:* Set the interrupt rate.
- *Value:*
 - Extreme
 - High

- Low
 - Minimal
 - Moderate
 - *Default:* Low
- **IPv4 Checksum Offload**
 - *Description:* Allow the checksum to be computed by the adapter for IPV4 packets.
 - *Value:*
 - Disabled
 - Rx & Tx Enabled
 - Rx Enabled
 - Tx Enabled
 - *Default:* Rx & Tx Enabled
- **Jumbo Packet (Maximum Transmission Unit)**
 - *Description:* Specify the Maximum Transmission Unit (MTU) value.
 - *Value:*
 - 4088 Bytes
 - 9014 Bytes
 - Disabled
 - *Default:* Disabled
- **Large Send Offload V2 (IPv4)**
 - *Description:* Configuring Large Send Offload (LSO) using version 2 for the IPv4 traffic.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **Large Send Offload V2 (IPv6)**
 - *Description:* Configuring Large Send Offload (LSO) using version 2 for the IPv6 traffic.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **Locally Administered Address**
 - *Description:* Specify a new MAC address for the port. If Not Present, the default MAC from EEPROM is used. If specified, the new MAC overwrites the default MAC. This specified MAC is persistent across machine reboot.
 - *Value:* Not Present | Any legal 6-byte MAC address
 - *Default:* Not Present
- **Maximum Number of RSS Processors**
 - *Description:* Change the number of RSS Processors.
 - *Value:* 1-64
 - *Default:* 8

- **Maximum Number of RSS Queues**
 - *Description:* Change the number of RSS Queues.
 - *Value:* 1-16
 - *Default:* 8
- **Maximum RSS Processor Number**
 - *Description:* Set the number of RSS processors to help the overall performance.
 - *Value:* 0-63
 - *Default:* 63
- **NDIS QoS**
 - *Description:* Enable or disable NDIS Quality of Service for DCB interface.
 - *Value:* Disabled | Enabled
 - *Default:* Disabled
- **NetworkDirect Functionality**
 - *Description:* Enable or disable Network Direct Functionality.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **NetworkDirect Interrupt Moderation**
 - *Description:* Control the interrupts generated during NDK traffic.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled

For more information, refer <http://msdn.microsoft.com/en-us/library/windows/hardware/ff556017%28v=vs.85%29.aspx>

- **Nvgre Task Offload**
 - *Description:* Enable or disable offloading processing of NVGRE encapsulated frames.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **Packet Direct**
 - *Description:* Enable or disable PacketDirect feature.
 - *Value:* Disabled | Enabled
 - *Default:* Disabled

 **Note** *This feature is available for experimental usage only.*

- **Preferred NUMA node**
 - *Description:* The NUMA node that the adapter can allocate memory from.
 - *Value:*
 - Default Settings

- Node 0
- Node 1
- Node 2
- Node 3
- Node 4
- Node 5
- Node 6
- Node 7
- *Default:* Default Settings
- **PTP Hardware Timestamp**
 - *Description:* Enable or Disable Hardware timestamping for transmit and receive packets.
 - *Value:* Disabled | Enabled
 - *Default:* Disabled
- **Receive Side Scaling (RSS)**
 - *Description:* Control the RSS functions. If on, the Microsoft RSS function is enabled. If off, the Chelsio RSS function is enabled.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **Recv Segment Coalescing (IPv4)**
 - *Description:* Enable or disable the Receive Segment Coalescing feature for IPv4 traffic. This is a receiver packets aggregating feature which helps reduce the receive host CPU load and improve throughput in a network environment where CPU can be the bottleneck.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **Recv Segment Coalescing (IPv6)**
 - *Description:* Enable or disable the Receive Segment Coalescing feature for IPv6 traffic.
 - *Value:* Disabled | Enabled
 - *Default:* Enabled
- **RSS Base processor**
 - *Description:* Specify the number of the base RSS processors.
 - *Value:* 0-63
 - *Default:* 0
- **RSS load balancing profile**
 - *Description:* Control the RSS load balancing profile.
 - *Value:*
 - Closest Processor
 - Closest Processor Static

- Conservative Scaling
- NUMA Scaling
- NUMA Scaling Static
- *Default:* NUMA Scaling Static


For more information, refer <http://technet.microsoft.com/en-us/library/hh997036.aspx>

- **Rx Ethernet Queue Size**

- *Description:* Set different Rx Queue sizes. Driver ships with reasonable defaults.
- *Value:* 512-4096
- *Default:* 2048


- **Rx Offload Queue Size**

- *Description:* Set different Offload Rx Queue sizes. Driver ships with reasonable defaults.
- *Value:* 256-2048
- *Default:* 512

 **Note** *This feature is not fully tested and available for experimental usage only.*

- **Software Timestamp**

- *Description:* Configure the driver to generate software timestamps.
- *Value:*
 - Disabled
 - RxAll
 - RxAll & TaggedTx
 - RxAll & TxAll
 - TaggedTx
 - TxAll
- *Default:* Disabled

 **Note** *Currently, only RxAll and TaggedTx are supported.*

- **SR-IOV**

- *Description:* Enable or disable SR-IOV feature.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **TCP Checksum Offload (IPv4)**

- *Description:* Enable or disable computation of TCP checksum by the adapter for IPv4 packets.
- *Value:*
 - Disabled
 - Rx & Tx Enabled
 - Rx Enabled
 - Tx Enabled

- *Default:* Rx & Tx Enabled

- **TCP Checksum Offload (IPv6)**


- *Description:* Configure computation of TCP checksum by the adapter for IPv6 packets.
- *Value:*
 - Disabled
 - Rx & Tx Enabled
 - Rx Enabled
 - Tx Enabled
- *Default:* Rx & Tx Enabled

- **Tx Ethernet Queue Size**

- *Description:* Set different Tx Queue sizes. Driver ships with reasonable defaults.
- *Value:* 512-4096
- *Default:* 1024

- **Tx Offload Queue Size**

- *Description:* Set different Offload Tx Queue sizes. Driver ships with reasonable defaults.
- *Value:* 256-4096
- *Default:* 512



Note

This feature is not fully tested and available for experimental usage only.

- **UDP Checksum Offload (IPv4)**

- *Description:* Configure computation of UDP checksum by the adapter for IPv4 packets.
- *Value:*
 - Disabled
 - Rx & Tx Enabled
 - Rx Enabled
 - Tx Enabled
- *Default:* Rx & Tx Enabled

- **UDP Checksum Offload (IPv6)**

- *Description:* Configure computation of UDP checksum by the adapter for IPv6 packets.
- *Value:*
 - Disabled
 - Rx & Tx Enabled
 - Rx Enabled
 - Tx Enabled
- *Default:* Rx & Tx Enabled

- **User Mode NetworkDirect**

- *Description:* Enable or disable Microsoft's Remote Direct memory Access (RDMA) interface for high speed, low latency networks.


- *Value:* Disabled | Enabled
- *Default:* Enabled

- **Virtual Machine Queues**

- *Description:* Enable or disable the virtual machine queue feature on the interface.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **Virtual Switch RSS**

- *Description:* Enable or disable Virtual Switch Receive Side Scaling.
- *Value:* Disabled | Enabled
- *Default:* Enabled

 **Note** *If NDIS function driver is installed using Unified Wire installer, the vRSS feature will be enabled by default. However, if the driver is installed manually (using zip package), it will be disabled. This feature is not available on Windows Client OS.*

- **VLAN Identifier**

- *Description:* Specify the VLAN ID to be inserted.
- *Value:* 0-4095
- *Default:* 0

- **VMQ LookAhead Split**

- *Description:* Enable or disable the ability to split receive buffers into lookahead and post-lookahead buffers.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **VMQ VLAN ID Filtering**

- *Description:* Enable or disable the ability to filter network packets by using the VLAN identifier in the media access control (MAC) header.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **Vxlan Task Offload**

- *Description:* Enable or disable offloading processing of VXLAN encapsulated frames.
- *Value:* Disabled | Enabled
- *Default:* Enabled

- **Vxlan UDP Port Number**

- *Description:* Specify the current VXLAN UDP destination port number that is currently operational in the NIC.
- *Value:* 1-65535
Default: 4789

2.2. NVGRE Offload

Chelsio's Terminator based adapters are uniquely capable of offloading the processing of NVGRE encapsulated frames such that all stateless offloads (checksums, LSO, VMQ, RSS and filtering) are preserved, resulting in significant performance benefits. This feature is enabled by default. Configure the relevant customer and provider network settings on the host.

2.3. VXLAN Task Offload

Virtual Extensible LAN (VXLAN), an extension of the existing VLAN protocol, is an encapsulation protocol which creates an overlay network on the existing Layer 3 infrastructure. VXLAN not only extends VLAN's Ethernet Layer 2 network services but also provides more extensibility, flexibility and scalability for large scale cloud deployments than VLAN.

Chelsio's Terminator based adapters are capable of offloading the processing of VXLAN encapsulated frames such that all stateless offloads (checksums, LSO, VMQ, RSS and filtering) are preserved, resulting in significant performance benefits. This feature is enabled by default.

2.4. Timestamps

Timestamping feature is used to maintain a precise and accurate time for the workloads. This is essential in many markets, especially in the financial where transactions need to be timestamped with minimal accuracy.

 **Note** *This feature is currently supported on Windows Server 2019 only.*

2.4.1. Hardware Timestamps

Chelsio adapters support hardware timestamping capabilities which can be used by PTP applications. With high precision times, timestamping is done in hardware for both egress and ingress packets without affecting the performance. This feature is disabled by default. Use **NDIS Miniport Driver Parameters** (Device manager → Chelsio Network Adapter Properties → Advanced Tab → PTP Hardware Timestamp) to enable it. The following capabilities are currently supported:

- In receive, hardware detects the packets with PTP event message and generates a timestamp corresponding to when such a packet was received. Both IPv4 and IPv6 are supported.
- In transmit, the hardware generates timestamps for specific packets (tagged by the OS).

2.4.2. Software Timestamps

Chelsio NDIS function driver supports generating software timestamps for the packets. This feature is disabled by default. Use **NDIS Miniport Driver Parameters** (Device manager →

Chelsio Network Adapter Properties → Advanced Tab → Software Timestamp) to enable it. Currently the following capabilities are supported:

- In receive, driver generates timestamps for all packets.
- In transmit, driver generates timestamps for specific packets (tagged by the OS).

2.5. PacketDirect

PacketDirect Provider Interface (PDPI), an extension to the existing NDIS miniport driver model, gives applications complete control over buffers, poll processors, and sending and receiving packets over a miniport adapter. Chelsio’s Unified Wire adapters offer extensive support for this feature leading to a higher throughput performance.

Note *This feature is available for experimental usage only.*

This feature is disabled by default. Follow the steps mentioned below to enable and configure it:

- Open **Device Manager**, click on **Network adapters**, double click on **Chelsio Network Adapter**, and then click the **Advanced** tab.
- Select the *Virtual Switch RSS* property and enable it (if not already enabled).
- Select the *Packet Direct* property and enable it. Click OK

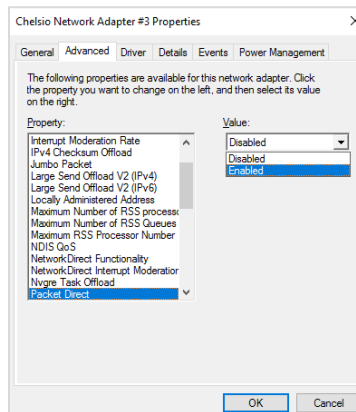


Figure 59 - Enabling Packet Direct

- Create a new virtual switch and enable PacketDirect on the host:

```
C:\Users\Administrator> New-VMSwitch -SwitchName <virtual_switch>
-NetAdapterName <chelsio_interface> -EnablePacketDirect $true
```

```
PS C:\Users\Administrator> New-VMSwitch -SwitchName sw0 -NetAdapterName "Ethernet 3" -EnablePacketDirect $true
Name SwitchType NetAdapterInterfaceDescription
-----
sw0 External Chelsio Network Adapter
```

Figure 60 - Enable PacketDirect

- v. Enable Azure Virtual Filtering Platform (VFP) Switch Extension on the switch.

```
C:\Users\Administrator> Enable-VMSwitchExtension -VMSwitchName
<virtual_switch> -Name "Microsoft Azure VFP Switch Extension"
```

```
PS C:\Users\Administrator> Enable-VMSwitchExtension -VMSwitchName sw0 -Name "Microsoft Azure VFP Switch Extension"
Id                : E9859CFA-2BE1-4B21-828F-B6FBDBDDC017
Name              : Microsoft Azure VFP Switch Extension
Vendor            :
Version           :
ExtensionType     : Forwarding
ParentExtensionId :
ParentExtensionName :
SwitchId         : 48f433e2-10ad-447a-8566-bb4864dbcb12
SwitchName       : sw0
Enabled           : True
Running          : True
CimSession       : CimSession: .
ComputerName     : DUKE2
IsDeleted        : False
```

Figure 61 - Enable VFP Switch Extension

- vi. Add a virtual network adapter to the virtual machine and connect it to the virtual switch:

```
C:\Users\Administrator> Add-VMNetworkAdapter -VMName <virtual_machine>
-SwitchName <virtual_switch> -Name <virtual_adapter>
```

```
C:\Users\Administrator> Add-VMNetworkAdapter -VMName VM1 -SwitchName sw0 -Name pd_nw1
C:\Users\Administrator>
```

Figure 62 - Add virtual network adapter

- vii. Set the number of PacketDirect processors:

```
C:\Users\Administrator> Set-VMNetworkAdapter -VMName <virtual_machine>
-PacketDirectNumProcs <1-8>
```

```
PS C:\Users\Administrator> Set-VMNetworkAdapter -VMName VM1 -PacketDirectNumProcs 4
PS C:\Users\Administrator>
```

Figure 63 - Set PacketDirect processors

- viii. Start the VM:

```
C:\Users\Administrator> Start-VM <virtual_machine>
```


ix. On the host, create a PowerShell script (with `.ps1` extension) with the following code:

```
param(
    [string]$switchName = $(throw "please specify a switch name")
)

$switches = Get-WmiObject -Namespace root\virtualization\v2 -Class
Msvm_VirtualEthernetSwitch
foreach ($switch in $switches) {
    if ( $switch.ElementName -eq $switchName) {
        $ExternalSwitch = $switch
        break
    }
}

$vfpcCtrlExe = "vfpcctrl.exe"
$ports = $ExternalSwitch.GetRelated("Msvm_EthernetSwitchPort",
"Msvm_SystemDevice", $null, $null, $null, $null, $false, $null)
foreach ($port in $ports) {
    #if ($port.ElementName -eq "Dynamic Ethernet Switch Port")
    #{
        $portGuid = $port.Name
        echo "Disabling VFP on port: " $portGuid
        & $vfpcCtrlExe /disable-port /port $portGuid
    #}
}
```

x. Run the script:

```
PS C:\Users\Administrator> <script>.ps1 <virtual_switch>
```

```
PS C:\Users\Administrator> .\disableVFP.ps1 sw0
Disabling VFP on port:
0D46B798-B405-40FE-86BC-6C67A870F368
ERROR: failed to execute disable-port
Error (1): Incorrect function.
Disabling VFP on port:
670970D3-23AF-45B0-8741-422C79F2B3DC
ERROR: failed to execute disable-port
Error (1): Incorrect function.
Disabling VFP on port:
C0C0B3ED-F097-4815-AD12-66B0492D360D
Command disable-port succeeded!
Disabling VFP on port:
D9049B8A-D9AA-4D08-B746-D0960EC25848
Command disable-port succeeded!
PS C:\Users\Administrator> _
```

Figure 64 - Running the script

Note *The errors seen in the output are expected.*

xi. Use system performance monitoring utilities, such as *PerfMon*, to add PacketDirect counters and verify if they are updated.

2.6. Virtual Environments

Network throughput of a multi-core virtual machine is limited by the processing power of a single virtual processor, since only one processor handles all the interrupts from a virtual network adapter. Virtual Machine Multi-Queue (VMMQ)/Virtual Switch Receive-side Scaling (vRSS) and Dynamic VMMQ/RSSv2 solve this bottleneck by distributing the network traffic across multiple virtual processors. Dynamic VMMQ/RSSv2, makes RSS algorithm dynamic by taking CPU measurement and making balancing decision at DISPATCH level. The feature also avoids taking global locks by changing only the subset of indirection table which affects current CPUs only.

If NDIS function driver is installed using Unified Wire installer, these features will be enabled by default. However, if the driver is installed manually (using zip package), they will be disabled. Use **NDIS Miniport Driver Parameters** (Device manager → Chelsio Network Adapter Properties → Advanced Tab → Virtual Switch RSS) to enable them.

! Important *This is currently supported for 28 VM Network Adapters per Chelsio adapter.*

- **VMMQ/vRSS**

Additionally, VMMQ/vRSS must be enabled for each VM on the host using:

```
PS C:\Users\Administrator> Set-VMNetworkAdapter -VMName <VM Name>
-VrssEnabled:$true -VmmqEnabled:$true -VmmqQueuePairs <No. of Queues>
```

```
PS C:\Users\Administrator> Set-VMNetworkAdapter -VMName vm1 -VrssEnabled 1 -VmmqEnabled 1 -VmmqQueuePairs 8
PS C:\Users\Administrator> Get-VMNetworkAdapter -VMName vm1 | fl | findstr "Enabled"
DynamicMacAddressEnabled           : True
VrssEnabledRequested              : True
VrssEnabled                       : True
VmmqEnabledRequested              : True
VmmqEnabled                       : True
```

Figure 65 - Enabling VMMQ/vRSS for VM

- **Dynamic VMMQ/RSSv2**

i Note *Available on Windows Server 2016 version 1709 and later.*

Additionally, Dynamic VMMQ/RSSv2 must be enabled for each VM on the host using:

```
PS C:\Users\Administrator> Set-VMNetworkAdapter -VMName <VM Name>
-VrssEnabled 1 -VmmqEnabled 1 -VmmqQueuePairs <No. of Queues>
-VrssQueueSchedulingMode Dynamic
```

```
PS C:\Users\Administrator> Set-VMNetworkAdapter -VMName vm1 -VrssEnabled 1 -VmmqEnabled 1 -VmmqQueuePairs 8
-VrssQueueSchedulingMode Dynamic
PS C:\Users\Administrator> Get-VMNetworkAdapter -VMName vm1 | fl | findstr "VrssQueueSchedulingMode"
VrssQueueSchedulingModeRequested   : Dynamic
VrssQueueSchedulingMode            : Dynamic
```

Figure 66 - Enabling Dynamic VMMQ/RSSv2 for VM

III. iWARP (ND)

1. Introduction

ND or Network Direct is Microsoft's new Remote Direct memory Access (RDMA) interface for high speed, low-latency networks such as those running on 10/25/40/50/100 Gigabit Ethernet or InfiniBand, with an architecture that directly bypasses OS and TCP/IP overhead. ND achieves better performance for massively parallel programs that can utilize very low-latency, high-bandwidth, and enables efficient CPU utilization.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the adapters that are compatible with Chelsio iWARP (ND) driver:

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

1.2. Software Requirements

1.2.1. Windows Requirement

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

- Server 2019
- Server 2016
- Server 2012 R2
- 10 Client (Enterprise and "Pro for Workstations" editions)

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

2. Software/Driver Configuration and Fine-tuning

2.1. Registering iWARP(ND) driver

Chelsio Unified Wire Installer registers Chelsio iWARP provider automatically during installation.

In case of zip package, run the following command to register:

```
C:\Users\Administrator>chinstallsp.exe -i
```

The iWARP provider should show up using:

```
C:\Users\Administrator>chinstallsp.exe -l
```

You should see a similar output:

```
0000001019 - Chelsio Network Direct provider
```

2.2. chrping application

chrping application can be used to test RDMA connectivity between the hosts. Configure the Chelsio interfaces on the hosts with the appropriate IP Address, netmask etc. Use the basic ping command to verify connectivity via the Terminator interface. In case of Installer, *chrping.exe* will be copied to `<system_drive>\Windows\System32` folder during installation.

If you are using the zip package, copy *chrping.exe* from `ChelsioUwire_x.x.x.x_WIN_yyy.y.yy\chrping` folder to `<system_drive>\Windows\System32`

The following steps verify RDMA connectivity between a Windows and a Linux host.

! Important *Disable iWARP Port Mapper (iwpm) service on Linux hosts. On RHEL7.X machines, use the below command.*

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

i. Start rping server on Linux host.

```
[root@host~]# rping -s -a <server_ip> -p <port #> -S <data size>
```

ii. Start chrping client on Windows host.

```
C:\Users\Administrator>chrping.exe c <server_ip> t <port #> S <data size>
```

Example:

Connecting to a Linux IPv4 Server.

```
PS C:\Users\Administrator\Desktop> chrping.exe c 10.1.1.5 t 8000 S 1024
Chelsio rdma interop utility v0.1
Bind Succeeded!!StartBuf Base:0x189574a6fd0 RKey:0x12fd85 Len:1024
Bind Succeeded!!RdmaBuf Base:0x189574a73e0 RKey:0x130986 Len:1024
Test complete
```

Figure 67 - IPv4 chrping

Connecting to a Linux IPv6 Server.

```
PS C:\Users\Administrator\Desktop> chrping.exe c 1000::5 t 8000 S 1024
Chelsio rdma interop utility v0.1
Bind Succeeded!!StartBuf Base:0x19437dec4e0 RKey:0x132d89 Len:1024
Bind Succeeded!!RdmaBuf Base:0x19437dec8f0 RKey:0x13368a Len:1024
Test complete
```

Figure 68 - IPv6 chrping

2.3. HMA

The Host Memory Access(HMA) implements a bridge between the Memory Arbiter (MA) and the PCIE so that the Host Memory is available to all clients. The translation from the MA memory map to the Host Memory map is performed based on the client number and the address of the transaction, both of which are used as index to look up a 64-bit offset within the Host Memory.

Currently 256 IPv4/128 IPv6 Network Direct offload connections are supported on T6225-SO-CR adapter. You can view the number of connections offloaded by using

```
C:\Users\Administrator> cxgbtool <nicIface> hardware tid_info
```

```
PS C:\> cxgbtool nic0 hardware tid_info
TID range: 0..255, in use: 128
STID range: 256..319, in use: 10
ATID range: 0..127, in use: 0
FTID range: 320..687
HW TID usage: 0 IP users, 128 IPv6 users
```

Figure 69 - HMA Offload connections

IV. SMB Direct

1. Introduction

SMB Direct is an extension of the Server Message Block (SMB) technology by Microsoft used for file operations. The *Direct* part implies the use of various high speed Remote Data Memory Access (RDMA) methods to transfer large amounts of data with little CPU intervention. By using RDMA as a transport medium for SMB, unprecedented levels of performance and efficiency can be achieved. With fully offloaded RDMA support, Chelsio Unified Wire adapters deliver large performance and efficiency gains to Windows users in a seamless, plug and play fashion.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the adapters that are compatible with Chelsio SMB Direct and RDMA with Virtual Switch (Mode 2) driver:

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

1.2. Software Requirements

1.2.1. Windows Requirement

Following is the list of SMB Direct components and supported Windows versions:

Windows Version	SMB Direct Components
Server 2019	SMB Direct*, SR, S2D
Server 2016	
10 Client [^]	SMB Direct
Server 2012 R2	

[^] Enterprise and "Pro for Workstations" editions

* SMB Direct and RDMA with Virtual Switch (Mode 2)

2. Software/Driver Configuration and Fine-tuning

2.1. Enabling SMB Direct

- **Windows Server 2019/2016/2012 R2**

SMB Direct is installed and enabled by default on Windows Server 2019/2016/2012 R2. RDMA functionality on Chelsio adapters will be enabled after installing the driver from the package.

- **Windows 10 Client**

On Windows 10 Client systems, SMB Direct feature is installed but disabled by default. To enable it, follow the steps mentioned below:

- Open **Control Panel** and click **Programs**.
- Under **Programs and Features**, click **Turn Windows features on or off**.
- In the **Windows Features** window, scroll down and select the *SMB Direct* checkbox.
- Click **OK**.

RDMA functionality on Chelsio adapters will be enabled after installing the driver from the package.

- **Configuration**

You can use the PowerShell command `Get-NetAdapterRdma` to get a list of RDMA capable interfaces in your system. At least one adapter should indicate “RDMA Capable” true to take advantage of SMB RDMA capability.

The next step is to configure your network so that SMB server and client machines can communicate with each other. Please configure the Microsoft firewall to allow traffic between the SMB clients and server.



Important *When using with Linux machines, disable iWARP Port Mapper (iwmpmd) service on them. On RHEL7.X machines, use the below command.*

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

2.2. Verifying RDMA

There are number of ways to verify if RDMA is working:

2.2.1. Verify active connections using PowerShell

Open PowerShell window and execute the command `Get-SmbConnection`

2.2.2. Use Event Viewer

You can use Event Viewer to view SMB connection events. To do this:

- i. Open Event Viewer
- ii. In the console tree, expand **Event Viewer**.
- iii. Navigate to **Application and Service Logs-> Microsoft->Windows->SMBClient ->Operational**

2.3. RDMA with Virtual Switch (Mode 2)

In this mode, you can run RDMA traffic using a virtual switch. This feature is enabled by default in the driver, but must be enabled on the Hyper-V Virtual Ethernet adapter. Follow the steps mentioned below to enable the feature:

- i. If you haven't done already, run the Chelsio Unified Wire Installer which will install NDIS function driver.
- ii. Verify if the driver is installed and loaded using the *Get-NetAdapter* command. The *InterfaceDescription* field should list Chelsio adapter as shown in the image below:

```
PS C:\Users\Administrator> Get-NetAdapter
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
vEthernet (p0)	Hyper-V Virtual Ethernet Adapter	18	Up	00-07-43-04-B4-20	25 Gbps
Ethernet 4	Chelsio Network Adapter #2	9	Up	00-07-43-04-B4-28	25 Gbps
Ethernet 3	Chelsio Network Adapter	2	Up	00-07-43-04-B4-20	25 Gbps
Ethernet 2	Intel(R) 82574L Gigabit Network Co...#2	8	Disconnected	00-25-90-35-95-FB	0 bps
Ethernet	Intel(R) 82574L Gigabit Network Conn...	11	Not Present	00-25-90-35-95-FA	0 bps
Local Area Connection* 1	Microsoft Kernel Debug Network Adapter	10	Up	00-25-90-35-95-FA	1 Gbps

Figure 70 - Verifying Chelsio driver

- iii. Next, create a virtual switch using the following syntax:

```
PS C:\Users\Administrator> New-VMSwitch -Name <virtual_switch>
-NetAdapterName <chelsio_interface>
```

```
PS C:\Users\Administrator> New-VMSwitch -Name switch0 -NetAdapterName "Ethernet 4"
```

Name	SwitchType	NetAdapter	InterfaceDescription
switch0	External	Chelsio Network Adapter #2	

Figure 71 - Creating virtual switch

- iv. Using *Get-NetAdapter* command, verify if the virtual switch is created successfully. The Mac address of the switch and the port on which the switch was created, should be same.

```
PS C:\Users\Administrator> Get-NetAdapter
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
vEthernet (switch0)	Hyper-V Virtual Ethernet Adapter #4	3	Up	00-07-43-04-B4-28	25 Gbps
vEthernet (p0)	Hyper-V Virtual Ethernet Adapter	18	Up	00-07-43-04-B4-20	25 Gbps
Ethernet 4	Chelsio Network Adapter #2	9	Up	00-07-43-04-B4-28	25 Gbps
Ethernet 3	Chelsio Network Adapter	2	Up	00-07-43-04-B4-20	25 Gbps
Ethernet 2	Intel(R) 82574L Gigabit Network Co...#2	8	Disconnected	00-25-90-35-95-FB	0 bps
Ethernet	Intel(R) 82574L Gigabit Network Conn...	11	Not Present	00-25-90-35-95-FA	0 bps
Local Area Connection* 1	Microsoft Kernel Debug Network Adapter	10	Up	00-25-90-35-95-FA	1 Gbps

Figure 72 - Verifying virtual switch

- v. Enable RDMA on the virtual switch using `Enable-NetAdapterRdma` command.

```
PS C:\Users\Administrator> Enable-NetAdapterRdma
PS C:\Users\Administrator> _
```

Figure 73 - Enabling RDMA on virtual switch

- vi. Run `Get-NetAdapterRdma` command to verify if RDMA is enabled on the switch.

```
PS C:\Users\Administrator> Get-NetAdapterRdma
```

Name	InterfaceDescription	Enabled
vEthernet (switch0)	Hyper-V Virtual Ethernet Adapter #4	True
vEthernet (p0)	Hyper-V Virtual Ethernet Adapter	True
Ethernet 4	Chelsio Network Adapter #2	True
Ethernet 3	Chelsio Network Adapter	True

Figure 74 - Verifying RDMA on virtual switch

2.4. Storage Replica

Storage Replica (SR) is a Windows Server 2016 feature which enables block-level replication between clusters or individual servers for disaster recovery, and stretching of failover clusters to metropolitan (MAN) and wide area (WAN, US coast-to-coast) distances for high availability. SR provides two modes of operation: *synchronous* and *asynchronous* replication. Synchronous replication enables mirroring of data with zero data loss at the volume level, whereas asynchronous replication trades off full data replication guarantees for reduced latency by locally completing I/O operations.

Currently, Unified Wire supports Server to Server replication mode. In this mode, data is synced between two servers and each server keeps an identical copy of the same volume. For more information on how to setup and configure in this mode, please refer to [Microsoft's official documentation](#).

2.5. Storage Spaces Direct (S2D)

Chelsio's high performance iWARP RDMA adapters enable incremental, non-disruptive server installs, and support the ability to work with any standard Ethernet switch, delivering a brownfield strategy to enable high performance, low cost, scalable Azure Stack Hyper Converged Infrastructure (HCI) deployments. Major benefits include cost savings on switches at higher speeds with each deployment. Windows SMB Direct over iWARP RDMA provides higher performance by giving direct access to the data residing on hyper-converged or disaggregated storage, while the CPU reduction enables a larger number of VMs per Hyper-V server, enabling savings in power dissipation, system configuration and deployment scale throughout the life of the installation. They prove to be a best fit for both networking and virtualization requirements, as well as hyper-converged scalable storage solutions like Storage Spaces Direct (S2D), a core storage feature for Azure Stack solution.

2.5.1. Configuring S2D

Connect all the ports of the Chelsio adapters to a Switch (Please refer [Switch Configuration](#) section for sample configurations). Install Unified Wire on all the machines in the Mesh Topology. Refer to [Chelsio Unified Wire](#) chapter. To configure S2D, follow the steps mentioned below:

- i. Install Hyper-V and Failover cluster roles on all the nodes.
- ii. Set the following Registry entries for Server 2019 and reboot the nodes to enable S2D.

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ClusSvc\Parameters]
"S2D"=dword:1
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\NetworkController]
"Enabled"=dword:1
```

- iii. Add the Nodes to a domain.
- iv. RDMA will be enabled by default for the Chelsio Ports. Verify using the following command.

```
PS C:\> Get-NetAdapterRdma
```

```
PS C:\> Get-NetAdapterRdma

Name                InterfaceDescription      Enabled  PFC    ETS
----                -
ch--p1              Chelsio Network Adapter #10 True     False  False
ch--p0              Chelsio Network Adapter #9  True     False  False
```

Figure 75 - RDMA for Chelsio Ports

- v. The disks intended to be used for S2D need to be empty and without partitions or other data. If a disk has partitions or other data, it will not be included in the S2D system. Check the status of all the disks using the below command.

```
PS C:\> Get-Disk
```

```
PS C:\> Get-Disk

Number Friendly Name                Serial Number                HealthStatus  OperationalStatus  Total Size Partition Style
-----
0      DELL PERC H710                    00fb24aea747143c2200ea91fb60f681 Healthy       Online             3.82 TB MBR
1      INTEL SSDPECKME016T4              CVF85475002B1P68GN-1_00000001 Healthy       Offline            745.21 GB RAW
2      INTEL SSDPECKME016T4              CVF85475002B1P68GN-2_00000001 Healthy       Offline            745.21 GB RAW
3      INTEL SSDPECKME016T4              CVF8547500181P68GN-1_00000001 Healthy       Offline            745.21 GB RAW
4      INTEL SSDPECKME016T4              CVF8547500181P68GN-2_00000001 Healthy       Offline            745.21 GB RAW
5      INTEL SSDPECKE020T7               0100_0000_0100_0000_5CD2_E4F2... Healthy       Offline            1.82 TB RAW
6      MTFDAX2T4MCF-1AN1ZABYY            CFE0_0003_0F24_B300...       Healthy       Offline            2.18 TB RAW
```

Figure 76 - Check Disks status

Note *Storage Spaces Direct does not support disks connected via multiple paths, and the Microsoft Multipath MPIO software stack.*

- vi. Before creating the cluster, validate the nodes using the cluster validation tool.

```
PS C:\> Test-Cluster -Node <Node1,Node2,...> -Include "Storage Spaces Direct",Inventory,Network,"System Configuration"
```

```
PS C:\> Test-Cluster -Node azure1.chddc.com,azure2.chddc.com,azure3.chddc.com -Include "Storage Spaces Direct",Inventory,Network,"System Configuration"
WARNING: System Configuration - Validate All Drivers Signed: The test reported some warnings..
WARNING: System Configuration - Validate Software Update Levels: The test reported some warnings..
WARNING: Storage Spaces Direct - Verify Node and Disk Configuration: The test reported some warnings..
WARNING:
Test Result:
HadUnselectedTests, ClusterConditionallyApproved
Testing has completed for the tests you selected. You should review the warnings in the Report. A cluster solution is supported by Microsoft only if you run all cluster validation tests, and all tests succeed (with or without warnings).
Test report file path: C:\Users\administrator.CHDDC\AppData\Local\Temp\Validation Report 2018.11.01 At 23.48.34.htm

Mode                LastWriteTime         Length Name
----                -
-a----             11/1/2018  11:49 PM           1209359 Validation Report 2018.11.01 At 23.48.34.htm
```

Figure 77 - Cluster Validation

- vii. Create a cluster using the nodes validated in the previous step.

```
PS C:\> New-Cluster -Name <ClusterName> -Node <Node1,Node2,...> -NoStorage -StaticAddress <static_ip> -Verbose
```

```
PS C:\> New-Cluster -Name cluster-s2d -Node "azure1.chddc.com", "azure2.chddc.com", "azure3.chddc.com" -NoStorage -StaticAddress 10.192.195.50 -Verbose
VERBOSE: Adding static network 10.192.192.0/20.

Name
----
cluster-s2d


PS C:\> Get-Cluster

Name
----
cluster-s2d

PS C:\> Get-ClusterNode

Name State Type
----
azure1 Up Node
azure2 Up Node
azure3 Up Node
```

Figure 78 - Cluster Creation



Note Without the `-NoStorage` parameter, the disks may be automatically added to the cluster and you will need to remove them before enabling S2D. Otherwise they will not be included in the S2D pool.

- viii. Enable S2D and create a storage pool.

```
PS C:\> Enable-ClusterS2D -CacheState <State> -Verbose
```

```

PS C:\> Enable-ClusterS2D -CacheState Disabled -SkipEligibilityChecks -Verbose
VERBOSE: 2018/11/01-23:58:56.414 Ensuring that all nodes support S2D
VERBOSE: 2018/11/01-23:58:56.439 Querying storage information
VERBOSE: 2018/11/01-23:58:56.806 Sorted disk types present (fast to slow): NVMe. Number of types present: 1
VERBOSE: 2018/11/01-23:58:56.807 Checking that nodes support the desired cache state

Confirm
Are you sure you want to perform this action?
Performing operation 'Enable Cluster Storage Spaces Direct' on Target 'cluster-s2d'.
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): A
VERBOSE: 2018/11/01-23:59:10.841 Creating health resource
VERBOSE: 2018/11/01-23:59:11.179 Setting cluster property
VERBOSE: 2018/11/01-23:59:11.180 Setting default fault domain awareness on clustered storage subsystem
VERBOSE: 2018/11/01-23:59:11.244 Waiting until physical disks are claimed
VERBOSE: 2018/11/01-23:59:14.251 Number of claimed disks on node 'azure1': 0/6
VERBOSE: 2018/11/01-23:59:17.261 Number of claimed disks on node 'azure2': 0/6
VERBOSE: 2018/11/01-23:59:20.271 Number of claimed disks on node 'azure3': 6/6
VERBOSE: 2018/11/01-23:59:23.280 Number of claimed disks on node 'azure1': 6/6
VERBOSE: 2018/11/01-23:59:26.289 Number of claimed disks on node 'azure2': 6/6
VERBOSE: 2018/11/01-23:59:26.298 Node 'azure1': Waiting until cache reaches desired state (HDD:'Disabled' SSD:'Disabled')
VERBOSE: 2018/11/01-23:59:26.302 SBL disks initialized in cache on node 'azure1': 6 (6 on all nodes)
VERBOSE: 2018/11/01-23:59:27.307 Node 'azure2': Waiting until cache reaches desired state (HDD:'Disabled' SSD:'Disabled')
VERBOSE: 2018/11/01-23:59:27.311 SBL disks initialized in cache on node 'azure2': 6 (12 on all nodes)
VERBOSE: 2018/11/01-23:59:28.316 Node 'azure3': Waiting until cache reaches desired state (HDD:'Disabled' SSD:'Disabled')
VERBOSE: 2018/11/01-23:59:28.320 SBL disks initialized in cache on node 'azure3': 6 (18 on all nodes)
VERBOSE: 2018/11/01-23:59:29.324 Waiting until SBL disks are surfaced
VERBOSE: 2018/11/01-23:59:32.346 Disks surfaced on node 'azure1': 18/18
VERBOSE: 2018/11/01-23:59:32.365 Disks surfaced on node 'azure2': 18/18
VERBOSE: 2018/11/01-23:59:32.404 Disks surfaced on node 'azure3': 18/18
VERBOSE: 2018/11/01-23:59:35.655 Waiting until all physical disks are reported by clustered storage subsystem
VERBOSE: 2018/11/01-23:59:38.954 Physical disks in clustered storage subsystem: 18
VERBOSE: 2018/11/01-23:59:38.955 Querying pool information
VERBOSE: 2018/11/01-23:59:39.223 Starting health providers
VERBOSE: 2018/11/01-23:59:50.357 Checking that all disks support the desired cache state
VERBOSE: 2018/11/01-23:59:50.401 Required steps for this action completed successfully

Node   EnableReportName
-----
azure1 C:\Windows\Cluster\Reports\EnableClusterS2D on 2018.11.01-23.59.50.htm

```

Figure 79 - Enable S2D

- ix. Create virtual disks on the storage pool created.

```

PS C:\Users\Administrator>Get-ClusterNode |% { New-Volume -
StoragePoolFriendlyName s2d -FriendlyName $_ -FileSystem CSVFS_ReFS -Size
500GB -Verbose }
PS C:\Users\Administrator>New-Volume -StoragePoolFriendlyName s2d -
FriendlyName Collect -FileSystem CSVFS_ReFS -Size 100GB -Verbose

```

- x. Create or deploy VMs. The VM's files should be stored on the virtual disks.

2.5.2. Switch Configuration

Chelsio iWARP RDMA does not require any configuration of DCB, PFC, ETC, ECN etc. on the Switch. It is recommended to disable them and enable regular flow control on switch ports. The following section shows the sample configuration on few switches.

- **Dell/Force 10 S4810 Switch**

```

Forcel0#configure
Forcel0(conf)#no dcb enable
Forcel0(conf)#interface fortyGigE 0/48
Forcel0(conf-if-fo-0/48)#flowcontrol rx on tx on
Forcel0(conf-if-fo-0/48)#shutdown
Forcel0(conf-if-fo-0/48)#no shutdown

```

- **Dell EMC S5148F-ON**

```
OS10# configure terminal
OS10(config)# interface ethernet 1/1/6
OS10(conf-if-eth1/1/6)# no lldp transmit
OS10(conf-if-eth1/1/6)# no lldp receive
OS10(conf-if-eth1/1/6)# no priority-flow-control
OS10(conf-if-eth1/1/6)# flowcontrol receive on
OS10(conf-if-eth1/1/6)# flowcontrol transmit on
OS10(conf-if-eth1/1/6)# shutdown
OS10(conf-if-eth1/1/6)# no shutdown
```

- **Cisco Nexus 5010**

```
ciscoswitchcert2# configure
ciscoswitchcert2(config)# interface ethernet 1/1
ciscoswitchcert2(config-if)# no lldp transmit
ciscoswitchcert2(config-if)# no lldp receive
ciscoswitchcert2(config-if)# priority-flow-control mode off
ciscoswitchcert2(config-if)# flowcontrol receive on
ciscoswitchcert2(config-if)# flowcontrol send on
ciscoswitchcert2(config-if)# shutdown
ciscoswitchcert2(config-if)# no shutdown
```

- **Mellanox 2410 Switch**

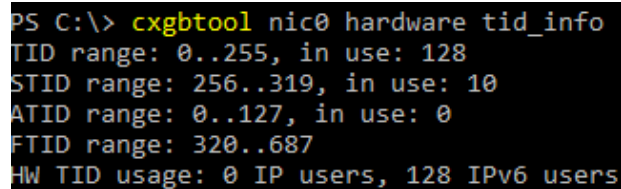
```
isn2410 [standalone: master] > enable
isn2410 [standalone: master] # configure terminal
isn2410 [standalone: master] (config) # interface ethernet 1/1
isn2410 [standalone: master] (config interface ethernet 1/1) # shutdown
isn2410 [standalone: master] (config interface ethernet 1/1) # no lldp
transmit
isn2410 [standalone: master] (config interface ethernet 1/1) # no lldp
receive
isn2410 [standalone: master] (config interface ethernet 1/1) # no dcb
priority-flow-control mode
isn2410 [standalone: master] (config interface ethernet 1/1) # flowcontrol
send on
isn2410 [standalone: master] (config interface ethernet 1/1) # flowcontrol
receive on
isn2410 [standalone: master] (config interface ethernet 1/1) # no shutdown
```

2.6. HMA

The Host Memory Access(HMA) implements a bridge between the Memory Arbiter (MA) and the PCIE so that the Host Memory is available to all clients. The translation from the MA memory map to the Host Memory map is performed based on the client number and the address of the transaction, both of which are used as index to look up a 64-bit offset within the Host Memory.

Currently 256 IPv4/128 IPv6 SMB Direct offload connections are supported on T6225-SO-CR adapter. You can view the number of connections offloaded by using

```
C:\Users\Administrator> cxgbtool <nicIface> hardware tid_info
```



```
PS C:\> cxgbtool nic0 hardware tid_info
TID range: 0..255, in use: 128
STID range: 256..319, in use: 10
ATID range: 0..127, in use: 0
FTID range: 320..687
HW TID usage: 0 IP users, 128 IPv6 users
```

Figure 80 - HMA Offload connections

Once the offload connections limit is reached, connections will continue to establish in regular NIC mode.

2.7. Troubleshooting

If RDMA is not working in your system, please check for the following:

- Operating system version should be Windows Server 2019/2016/2012 R2 or 10 Client (Enterprise and "Pro for Workstations" editions).
- Verify network interface configurations on both sides.
- Please ensure that the interfaces are reachable using *ping* command.
- In the Event Viewer (on the client side), please ensure that the traffic running is not over TCP.
- Verify that SMB driver is loaded and working.

```
C:\Users\Administrator>driverquery.exe | findstr /I smb
```

- Verify that Chelsio driver is loaded and working.

```
C:\Users\Administrator>driverquery.exe | findstr /I chel
```

- Confirm that traffic is running over Chelsio adapter.
- Look for potential errors in Event Log.

V. Mesh Topology

1. Introduction

Chelsio's fifth/sixth generation (T5/T6), high performance 10/25/40/50/100GbE adapters enable incremental, non-disruptive server installs, and support the ability to work without requiring any discrete external network switch, delivering a brownfield strategy to enable high performance, low cost, scalable deployments. Major benefits include cost savings on switches at higher speeds with each deployment. Mesh topology involves connecting each node to every other node.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the adapters that are compatible with the Chelsio Mesh topology:

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

* NIC only supported

1.2. Software Requirements

Currently the Chelsio Mesh Topology is available for the following Windows version(s):

- Server 2019
- Server 2016
- 10 Client
- Server 2012 R2

1.3. Mesh topology

Each node should be connected to other node. Supported configs using this approach: N ports per node, N+1 node cluster. The below is a 5-node Mesh using 4-port Chelsio adapters. NIC ports on each server connected to each other (1<->2, 1<->3, 1<->4, 1<->5, 2<->3, 2<->4, 2<->5, 3<->4, 3<->5, 4<->5).




Figure 81 - Mesh Topology

2. Software/Driver Installation

Install Unified Wire on all the machines in the Mesh Topology. Refer to [Chelsio Unified Wire](#) chapter.

3. Software/Driver Configuration and Fine-tuning

Configure the required IP addresses on all the interfaces. You should be able to run NIC and SMB Direct traffic between the nodes of mesh topology.

Example:

3 nodes are connected to each other in mesh topology with the following IP addresses.




Figure 82 - Sample connectivity

Node3 can communicate to node1,

```
PS C:\Users\Administrator> ping 102.1.1.3
```

Node3 can communicate to node2,

```
PS C:\Users\Administrator> ping 102.1.2.3
```

VI. NDIS SR-IOV

1. Introduction

The ever-increasing network infrastructure of IT enterprises has led 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 economic data center environments. Other benefits of Virtualization include improved disaster recovery, server portability, cloud computing, Virtual Desktop Infrastructure (VDI), etc. Chelsio's Unified Wire adapters with full support for PCI-SIG SR-IOV and virtualization features deliver increased bandwidth, lower latency and lower power to maximize cloud scaling and utilization.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the adapters that are compatible with the Chelsio NDIS SR-IOV components:

Chelsio Adapter	NDIS SR-IOV Components
T62100-SO-CR	NDIS SR-IOV
T580-SO-CR	
T580-OCP-SO	
T540-SO-CR	
T520-SO-CR	
T520-OCP-SO	
T62100-CR	NDIS SR-IOV, Guest RDMA
T62100-LP-CR	
T6425-CR	
T6225-CR	
T6225-LL-CR	
T6225-SO-CR	
T580-CR	
T580-LP-CR	
T540-BT	
T540-CR	
T540-LP-CR	
T520-CR	
T520-LL-CR	
T520-BT	

1.2. Software Requirements

Following is the list of NDIS SR-IOV components and supported Windows versions:

Windows Version	NDIS Components
Server 2016	NDIS SR-IOV, Guest RDMA
Server 2019	
Server 2012 R2	NDIS SR-IOV

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

1.3. Prerequisites

SR-IOV should be enabled in the machine.

2. Software/Driver Configuration and Fine-tuning

Please ensure that Unified Wire is installed on the host before proceeding. Refer [Software/Driver Installation](#) section of the **Chelsio Unified Wire** chapter for step-by-step instructions.

2.1. Enabling SR-IOV

The adapter's configuration should be updated to make use of the SR-IOV feature. Change the Serial Init to SR-IOV using Chelsio Management feature as explained [here](#).

Important To use the adapter in non-SRIOV scenarios in Windows or with other operating systems, the configuration type must be changed to **NON-SRIOV (Default)**.

2.2. Host Configuration

- i. Create a vSwitch with SR-IOV enabled using Chelsio adapter:
 - a) Open *Hyper-V Manager* application.




Figure 83 - Hyper-V Manager

b) Click the *Virtual Switch Manager* link under *Actions*.

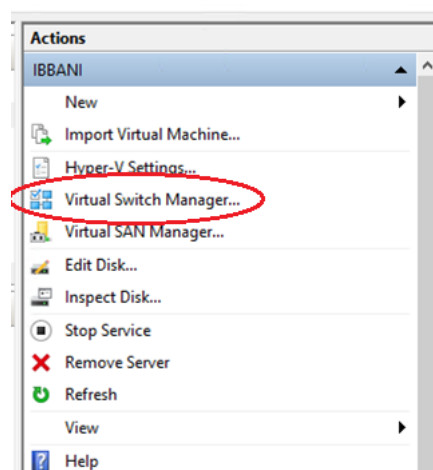


Figure 84 - Opening Virtual Switch Manager

c) Select *External* as the type of virtual switch to create. This will provide the VMs access to the external network. Click *Create Virtual Switch*.




Figure 85 - Creating Virtual Switch

- d) Provide a name for the vSwitch. Select the Chelsio adapter to connect the vSwitch to and select *Enable single-root I/O virtualization (SR-IOV)*. Click **OK**.

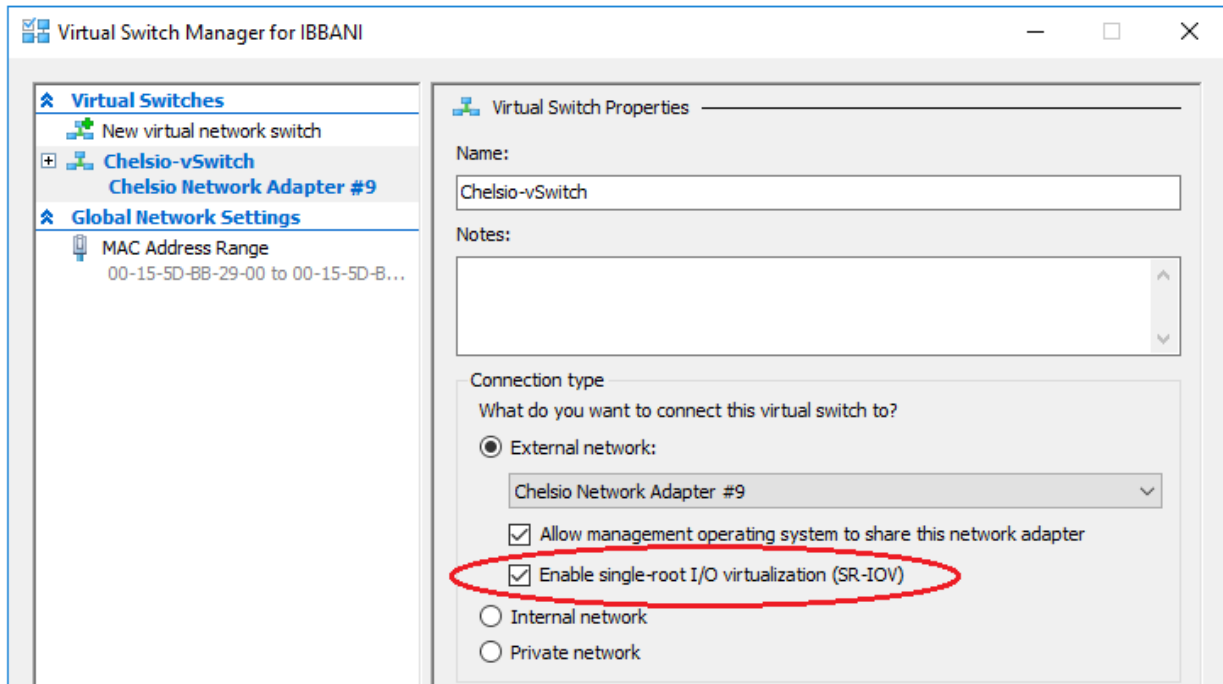


Figure 86 - Enabling SR-IOV on vSwitch

- e) Click **Yes** on the prompt that appears next.

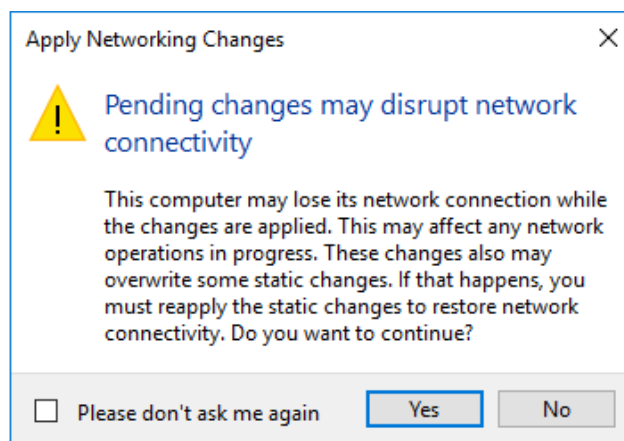


Figure 87 - Confirming Changes

- ii. Assign a virtual network adapter with SR-IOV enabled to the VM.
 - a) Open Hyper-V Manager. Right-click on the VM and select *Settings*.

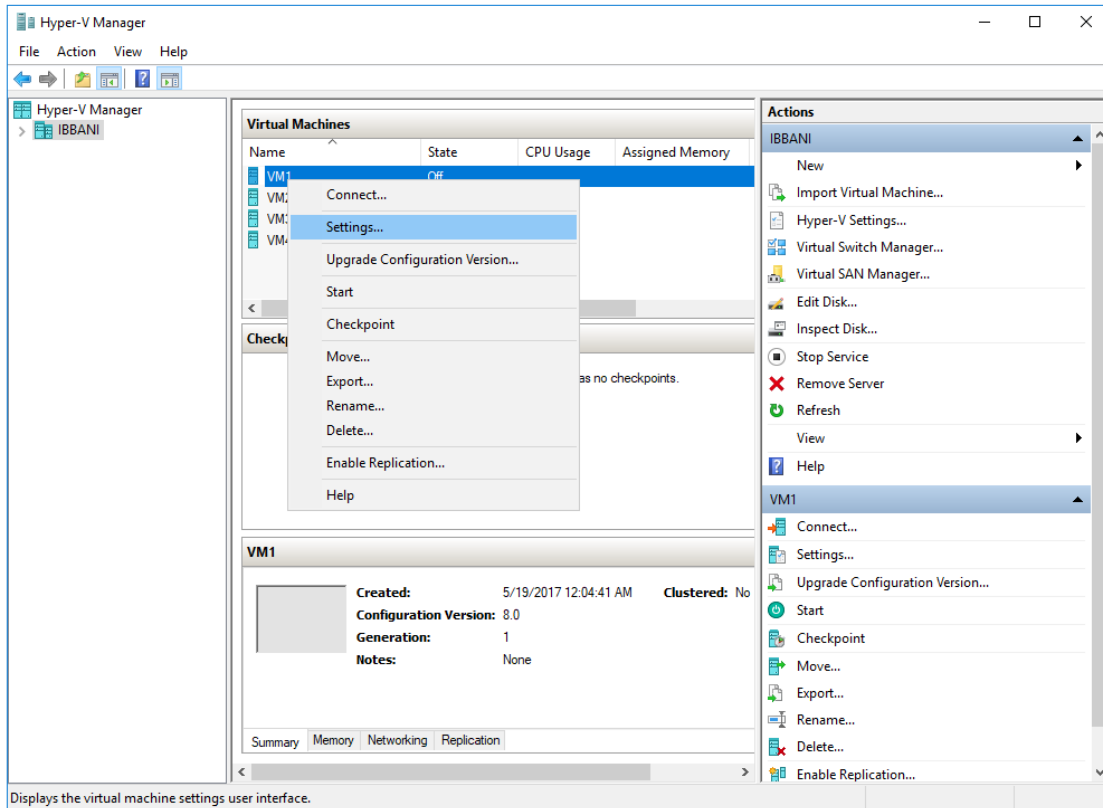


Figure 88 - Hyper-V Manager

- b) Expand the *Network Adapter* link, click *Hardware Acceleration* and select *Enable SR-IOV*. Click **OK**.

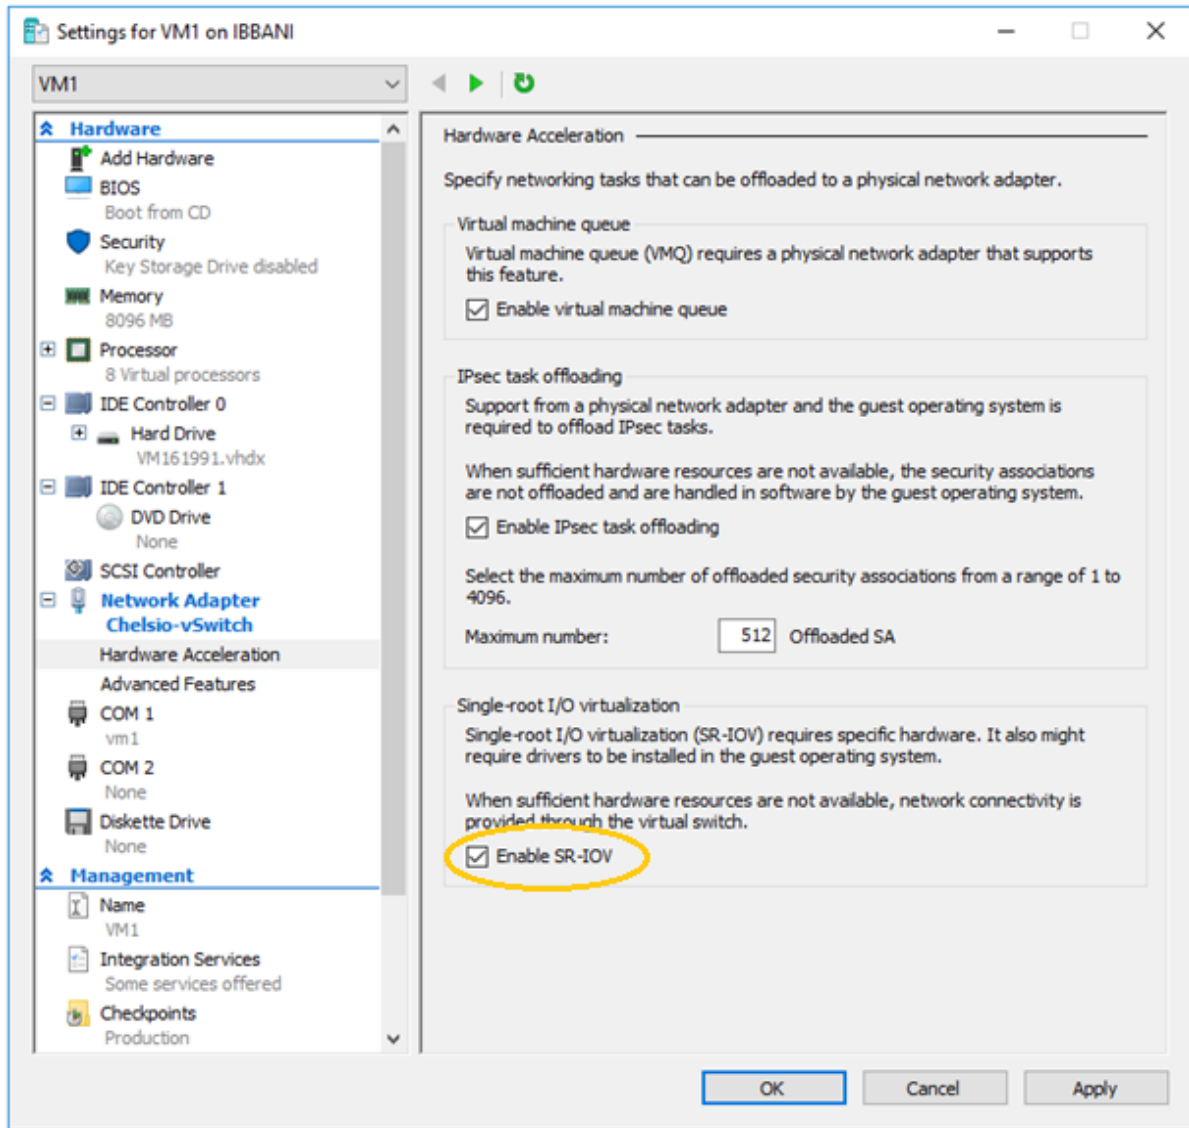


Figure 89 - Enabling SR-IOV on Virtual Network Adapter

c) Right-click on the VM and select *Start* to start the VM.




Figure 90 - Starting VM

2.3. Guest (VM) Configuration

- i. Use VM Connect as described [here](#) and access the Unified Wire package on the Guest (VM). Proceed with the Unified Wire installation as described in the [Software/Driver Installation](#) section of the **Unified Wire** chapter.
- ii. Verify if SR-IOV was enabled successfully on the host:

```
PS C:\Users\Administrator> Get-VMNetworkAdapter <vm_name> | fl
```

The `VFDataPathActive` parameter must display `True` as shown in the image below. This, however, is possible only if the physical port has an active link.


```
PS C:\Users\Administrator> Get-VMNetworkAdapter vml | fl
Name                : Network Adapter
Id                  : Microsoft:4C342213-5CC2-4FBD-93D7-CE036FD31CB3\C1018CF0-A3FF-4BEE-97B7-618C46A7461C
IsLegacy            : False
IsManagementOs     : False
ComputerName       : IBBANI
VMName             : VM1
VMId               : 4c342213-5cc2-4fbd-93d7-ce036fd31cb3
SwitchName         : Chelsio-vSwitch
SwitchId           : 09ab45b9-cd25-414f-8c18-01264f0b0bd2
Connected          : True
PoolName           :
MacAddress         : 00155D8B2918
DynamicMacAddressEnabled : True
AllowPacketDirect  : False
MacAddressSpoofing : Off
AllowTeaming       : Off
RouterGuard        : Off
DhcpGuard          : Off
StormLimit         : 0
PortMirroringMode  : None
IeeePriorityTag     : Off
VirtualSubnetId    : 0
DynamicIPAddressLimit : 0
DeviceNaming       : Off
VMQWeight          : 100
VMQUsage           : 0
IOVWeight          : 100
IOVUsage           : 1
IovQueuePairsRequested : 1
IovQueuePairsAssigned : 1
IOVInterruptModeration : Default
PacketDirectNumProcs : 0
PacketDirectModerationCount : 64
PacketDirectModerationInterval : 1000000
VrssEnabledRequested : True
VrssEnabled        : False
VmmqEnabledRequested : False
VmmqEnabled        : False
VmmqQueuePairsRequested : 16
VmmqQueuePairs     : 0
IPsecOffloadMaxSA  : 512
IPsecOffloadSAUsage : 0
VFDataPathActive   : True
MaximumBandwidth   :
MinimumBandwidthAbsolute :
MinimumBandwidthWeight :
BandwidthPercentage : 0%
```

Figure 91 - Verifying if SR-IOV was enabled successfully

To uninstall Unified Wire, please refer [Software/Driver Uninstallation](#) section of the Chelsio Unified Wire chapter for step-by-step instructions.

2.4. Guest RDMA (Mode 3)

In this mode, you can run RDMA traffic on Guest (VM) using a virtual switch. This feature is enabled by default in the driver, but must be enabled on the *Microsoft Hyper-V Virtual Ethernet* adapter.




Both Host and Guest should be installed with Windows Server 2016 version 1709 and later.

Follow the steps mentioned below to enable the feature:

- i. If you haven't done already, run the Chelsio Unified Wire Installer which will install NDIS function driver.
- ii. Configure Chelsio adapter as described in the [Enabling SR-IOV](#) section.
- iii. Enable SR-IOV support on the host as described in the [Host Configuration](#) section.
- iv. Enable Guest RDMA for VM using the following command on the host:

```
PS C:\Users\Administrator> Set-VMNetworkAdapterRdma -VMNetworkAdapterName
<adapter_name> -VMName <vm_name> -RdmaWeight 100
```



Make sure that VM is turned off before running the above command.

Example:

```
PS C:\Users\Administrator> Set-VMNetworkAdapterRdma -VMNetworkAdapterName "Network Adapter" -VMName VM1 -RdmaWeight 100
PS C:\Users\Administrator> _
```

Figure 92 - Enabling Guest RDMA for VM

To enable Guest RDMA on all VM network adapters, run the following command on the host:

```
PS C:\Users\Administrator> Set-VMNetworkAdapterRdma -VMNetworkAdapterName *
-VMName <VMName> -RdmaWeight 100
```

- v. Verify if Guest RDMA was enabled successfully using the *Get-VMNetworkAdapterRdma* command. The *Rdmaweight* parameter should show 100.

```
PS C:\Users\Administrator> Get-VMNetworkAdapterRdma vm1
RdmaWeight      : 100
ParentAdapter   : VMNetworkAdapter (Name = 'Network Adapter', VMName = 'VM1') [VMId = '4c342213-5cc2-4fbd-93d7-ce036fd31cb3']
IsTemplate      : False
CimSession      : CimSession: .
ComputerName    : IBBANI
IsDeleted       : False
```

Figure 93 - Verifying if Guest RDMA was enabled successfully

vi. In *Hyper-V Manager*, right-click on the VM and select *Start* to start the VM.




Figure 94 - Starting VM

vii. If you haven't done already, run the Chelsio Unified Wire Installer on VM which will install NDIS function driver.

viii. Open **Device Manager** on VM, click on **Network Adapters**, right click on **Microsoft Hyper-V Network Adapter** and select **Properties**.




Figure 95 - Device Manager on VM

- ix. In the **Advanced** tab, select the *Network Direct (RDMA)* property from the list and enable it. Click **OK**.




Figure 96 - Enabling Guest RDMA on Microsoft Hyper-V Network Adapter


- x. To verify if RDMA was enabled successfully, open PowerShell with administrative privileges and run `Get-NetAdapterRdma`, `Get-SmbClientNetworkInterface` and `Get-SmbServerNetworkInterface` commands. You should see a similar output:

```
PS C:\Users\Administrator> Get-NetAdapterRdma
Name                InterfaceDescription          Enabled
-----                -
Ethernet 21         Microsoft Hyper-V Network Adapter #18  True
Ethernet 22         Chelsio VF Network Adapter #4       True

PS C:\Users\Administrator> Get-SmbClientNetworkInterface
Interface Index RSS Capable RDMA Capable Speed      IpAddresses          Friendly Name
-----
12             True       True        25 Gbps {fe80::f0fa:848f:51f2:d5ec, 101.1.1.39} Ethernet 21

PS C:\Users\Administrator> Get-SmbServerNetworkInterface
Scope Name Interface Index RSS Capable RDMA Capable Speed      IpAddress
-----
*          12             True       True        25 Gbps fe80::f0fa:848f:51f2:d5ec
*          12             True       True        25 Gbps 101.1.1.39
```

Figure 97 - Verifying if Guest RDMA was enabled successfully



Note If the output of any one of the three commands is not *true*, then disable and enable the physical Chelsio network adapter on the host and try again.

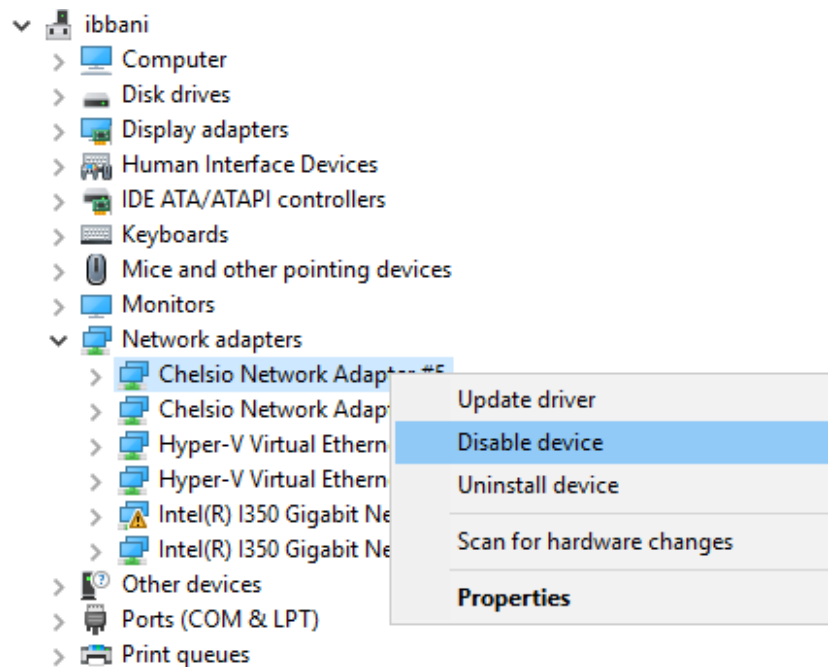


Figure 98 - Troubleshooting

You should be able to run RDMA traffic on the VM successfully now. Use a network traffic monitoring tool (like Windows Performance Monitor) to verify.



Important When using with Linux machines, disable *iWARP Port Mapper (iwpmid)* service on them. On RHEL7.X machines, use the below command.

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

2.4.1. Guest RDMA with VLAN

If VLAN is configured/updated for the VM network adapter (on the Host), Disable and Enable the **Chelsio Bus Enumerator [Virtual Function]** (on the VM: Device Manager -> System Devices) for the changes to take effect. Issue the below command on VM to check if the VF interface is listed as a Listener.

```
PS C:\Users\Administrator\Desktop> netstat -xan
```

If not listed, Disable and Enable the **Chelsio VF Network Adapter** (on the VM: Device Manager -> Network Adapters) and check if it is listed as a Listener.

VII. iSCSI Storport Miniport

1. Introduction

The Chelsio Terminator 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 transmitting, Chelsio h/w computes and inserts the Header and Data digest into the PDUs. On receiving, Chelsio h/w computes and verifies the Header and Data digest of the PDUs.
- **Direct Data Placement (DDP)**
Chelsio h/w can directly place the iSCSI Data-In or Data-Out PDU's payload into pre-posted final destination host-memory buffers based on the Initiator Task Tag (ITT) in Data-In or Target Task Tag (TTT) in Data-Out PDUs.
- **PDU Transmit and Recovery**
On transmitting, Chelsio h/w accepts the complete PDU (header + data) from the host driver, computes and inserts the digests, decomposes the PDU into multiple TCP segments if necessary, and transmit all the TCP segments onto the wire. It handles TCP retransmission if needed. On receiving, Chelsio h/w recovers the iSCSI PDU by reassembling TCP segments, separating the header and data, calculating and verifying the digests, then forwarding the header to the host. The payload data, if possible, will be directly placed into the pre-posted host DDP buffer. Otherwise, the payload data will be sent to the host too.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the Adapters that are compatible with Chelsio iSCSI Storport Miniport:

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

1.2. Software Requirements

1.2.1. Windows Requirement

Currently iSCSI Storport Miniport driver is available for Windows versions.

- Server 2019
- Server 2016
- 10 Client
- Server 2012 R2

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

2. Software/Driver Configuration and Fine-tuning

2.1. Configuring iSCSI Initiator

2.1.1. Configuring IP

There are two ways to assign IP to Chelsio iSCSI interface:

- **DHCP:** NIC driver (NDIS) must be installed (see [Software/Driver Installation](#)). Upon installing the driver, DHCP will be enabled by default and the NIC interface will be assigned an IP automatically. The same IP can be used by the iSCSI interface as well.
- **Static:** To assign a static IP to the iSCSI interface use the Device Manager, as explained below:

1. Open **Device Manager**, right click on **Chelsio iSCSI Initiator** under **Storage controllers** and click on **Properties**.




Figure 99 - Device Manager




Figure 100 - Chelsio adapter physical port properties

2. Click on **Advanced** tab and select IP option and add IP. Similarly add subnet mask and default gateway.




Figure 101 - Assigning IPv4 address

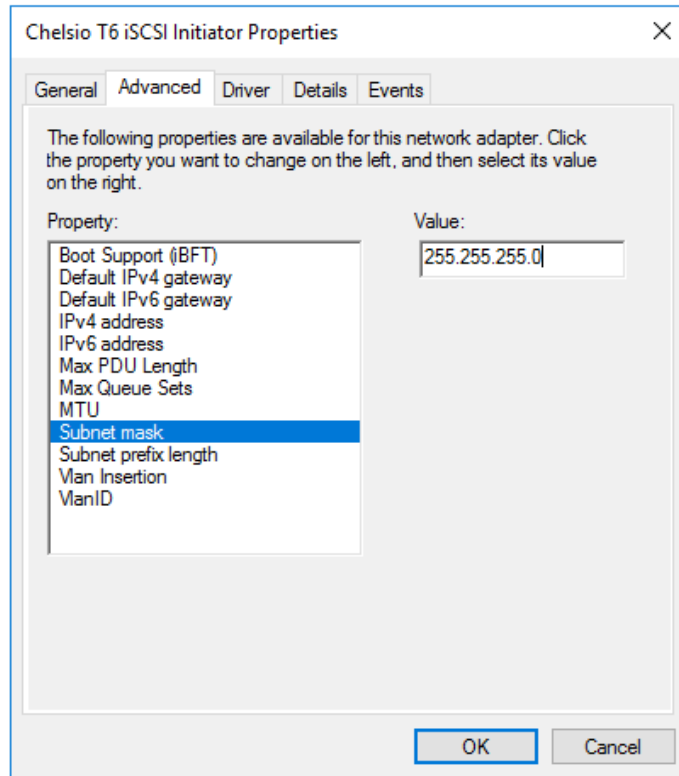


Figure 102 - Adding Subnet mask




Figure 103 - Adding Default IPv4 gateway




Figure 104 - Assigning IPv6 address




Figure 105 - Adding Subnet prefix length




Figure 106 - Adding Default IPv6 gateway

2.1.2. Configuring Optional Parameters

Configure the below optional parameters based on your requirement.

- The *Max PDU Length* parameter can be used to improve iSCSI performance. The default value of the parameter (indicated by 0) is set based on the MTU: 1024 for MTU 1500, 8192 for MTU 9000. However, these default values can be overridden by entering the value in increments of 512.

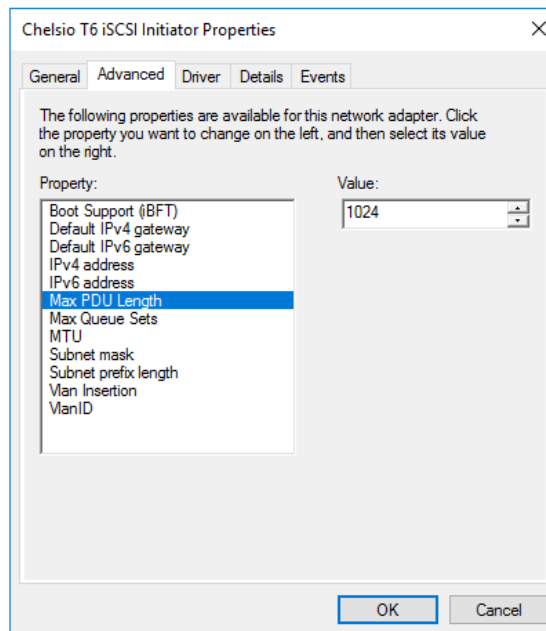


Figure 107 – Setting Max PDU Length

- Set the maximum iSCSI queue sets. This parameter can be used to improve performance.

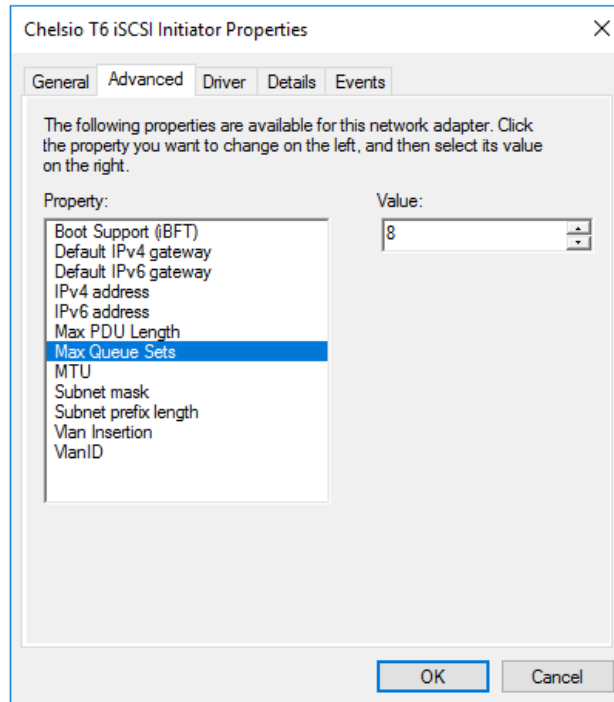


Figure 108 – Setting Max Queue Sets

- Select MTU and assign value as per requirement.




Figure 109 - Setting MTU

- VLAN is disabled by default. To enable it, click on “Vlan Insertion” and select value as “Enabled”.




Figure 110 - Enabling VLAN

- Assign the required VLAN ID and click **OK**.




Figure 111 - Assigning VLAN Id

Repeat the above-mentioned steps to assign these driver properties on other ports.

2.2. iSCSI Target Discovery and Login

Before target discovery, make sure an IP address, either static or using DHCP, is assigned to the respective Chelsio iSCSI node (see previous section).

- Note**
- *Connecting to more than 64 targets on single port using Chelsio iSCSI Initiator is currently not supported.*
 - *In case of Microsoft target, a maximum of 6 targets can be discovered by default. To increase this, change the maximum iSCSI PDU length to 8192 in **Device Manager** as explained [here](#).*

Following are the set of instructions to discover the target and login to it using the Chelsio iSCSI interface:

1. To start the Initiator configuration, go to Control Panel and click on the **iSCSI Initiator** option.

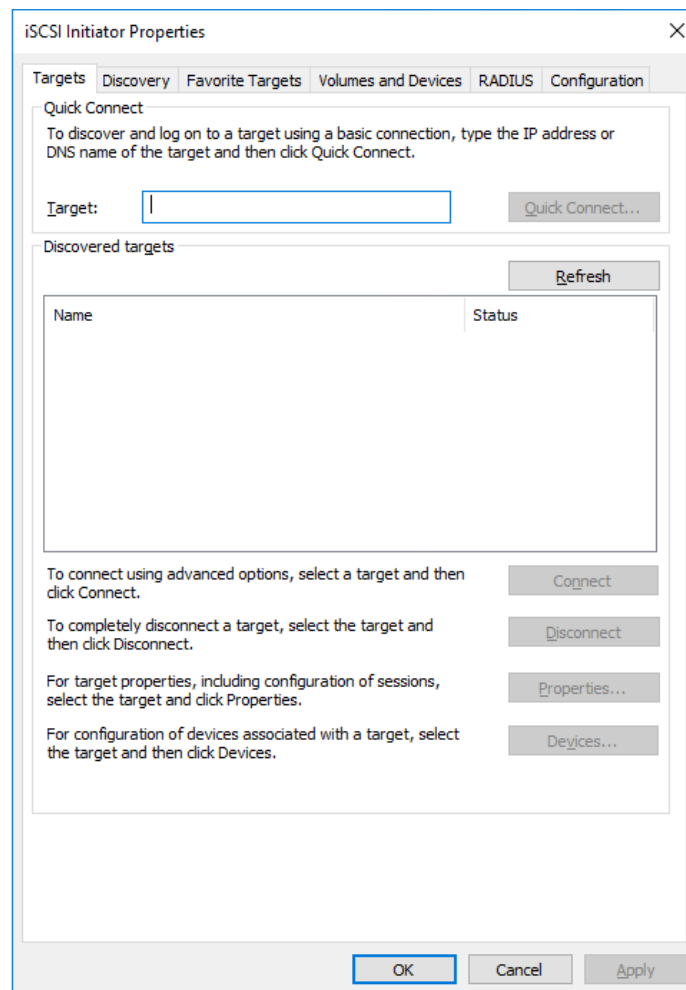


Figure 112 - iSCSI Initiator Properties: Targets tab

2. Choose the **Discovery** tab and click on **Discover Portal**.

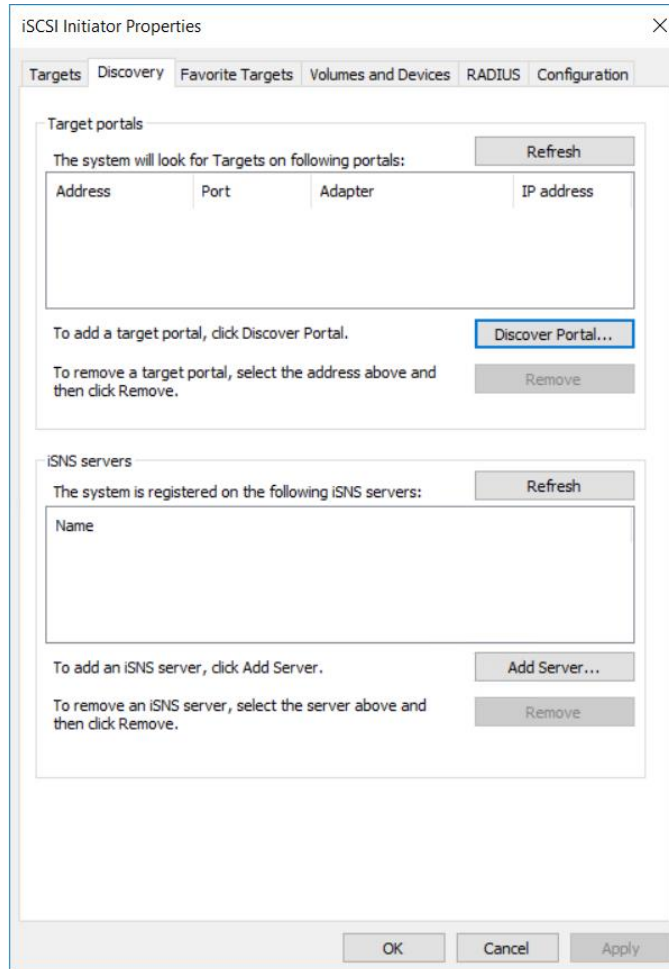


Figure 113 - iSCSI Initiator Properties: Discovery tab

3. Click on **Advanced** button.

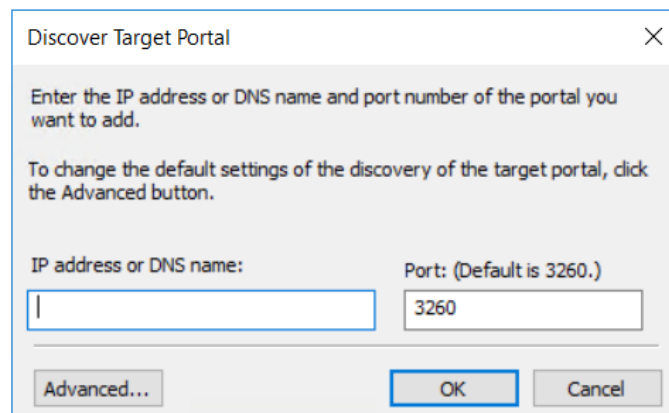


Figure 114 - Discovery Target Portal window

4. In the **Advanced Settings** window, select **Chelsio iSCSI interface** as the local adapter.




Figure 115 - Discovery Target Portal Advanced Settings: Selecting Local adapter

5. Select the corresponding Chelsio iSCSI Interface IP as the Initiator IP and click **OK**.

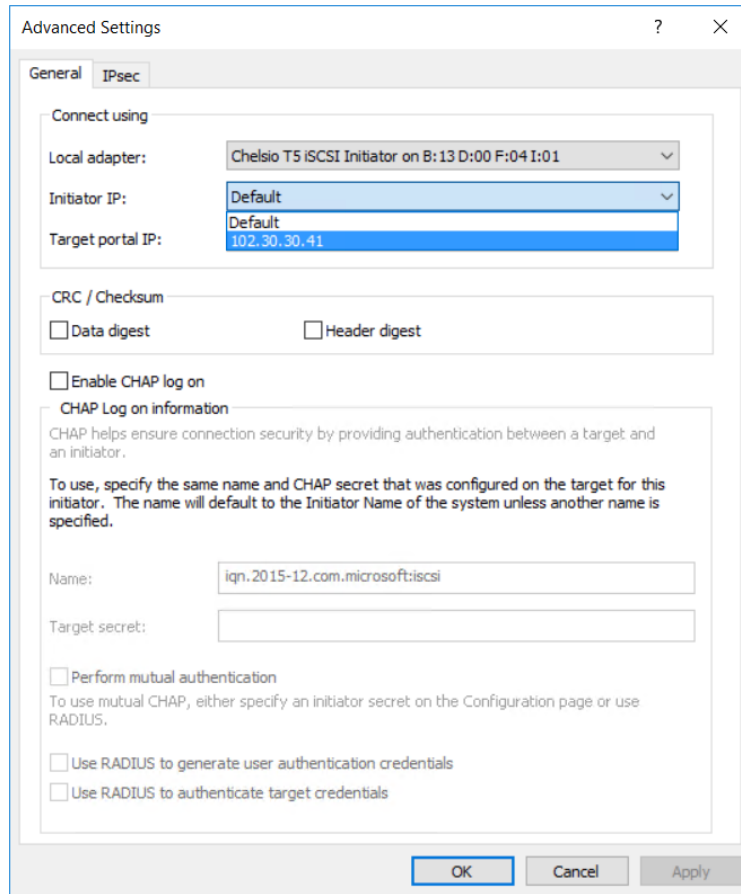


Figure 116 - Discovery Target Portal Advanced Settings: Specifying Initiator IP

6. Enter the IP address (DNS name not supported) of the target machine and the corresponding port number and click **OK** on the **Discover Target Portal** Window. Note that the default port number for iSCSI traffic is 3260.

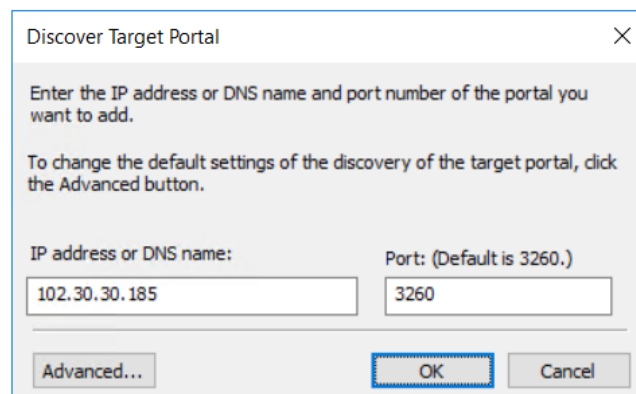


Figure 117 - Adding Target portal

7. Once target portal is added, details like target and initiator machine IP, Port number and Chelsio iSCSI interface IP will be displayed.




Figure 118 - Target portal added

- Click on the **Targets** tab to see the list of targets available, choose a particular target and click on **Connect**.




Figure 119 - Targets tab displaying list of available targets

- A window pops up showing the Target Name, Click **OK**.

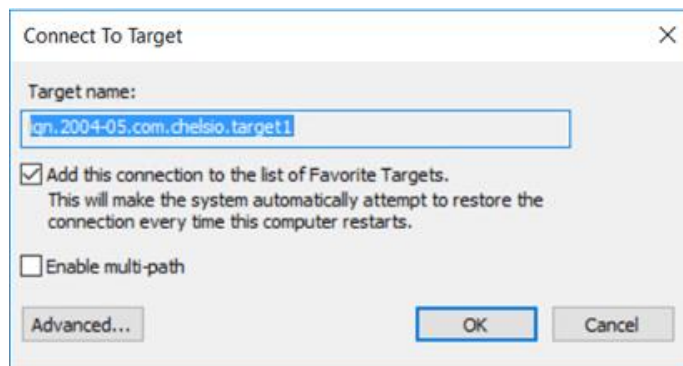


Figure 120 - Connecting to Target

10. After logging in to the target, the state of the target will change from *Inactive* to *Connected*.




Figure 121 - iSCSI target connected

11. After successful login, go to **Disk Management** and make the respective volume online.

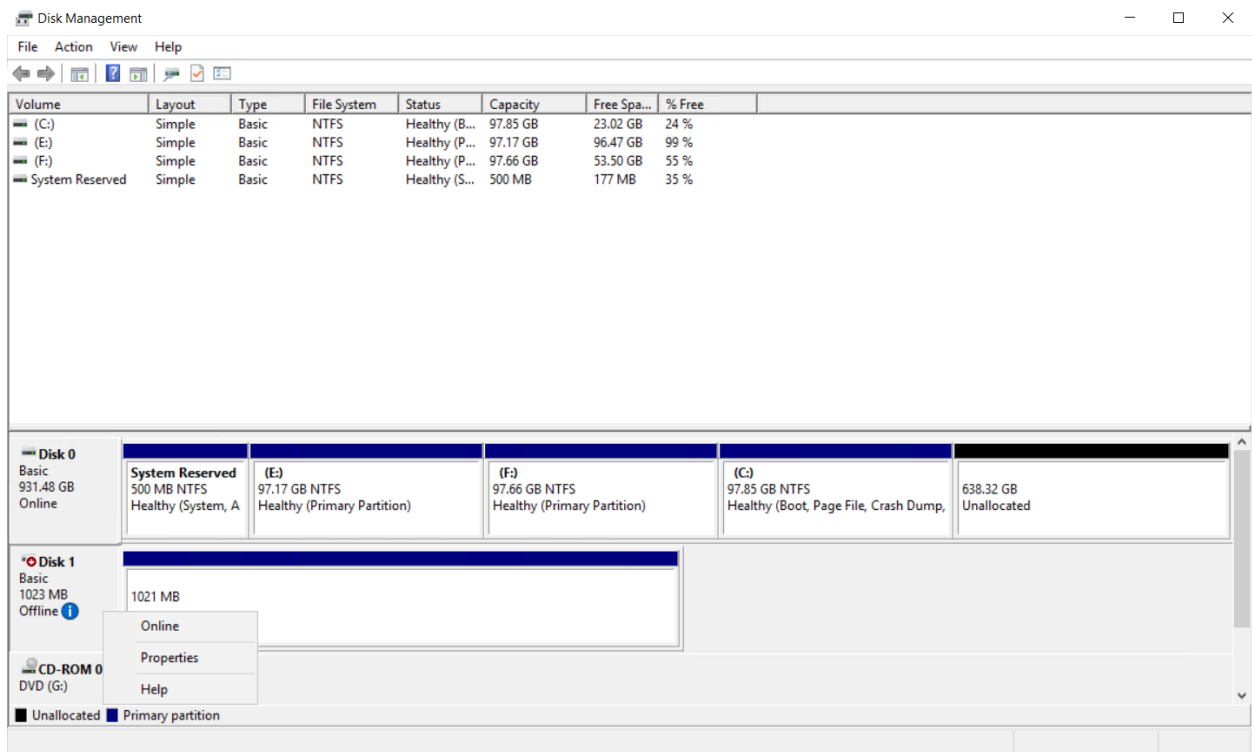


Figure 122 - Disk Management: making volume online

2.3. HMA

The Host Memory Access(HMA) implements a bridge between the Memory Arbiter (MA) and the PCIE so that the Host Memory is available to all clients. The translation from the MA memory map to the Host Memory map is performed based on the client number and the address of the transaction, both of which are used as index to look up a 64-bit offset within the Host Memory.

Currently 128 IPv4/IPv6 iSCSI Initiator Offload connections are supported on T6225-SO-CR adapter. You can view the number of connections offloaded by using

```
C:\Users\Administrator> cxgbtool <nicIface> hardware tid_info
```

```
PS C:\> cxgbtool nic0 hardware tid_info
TID range: 0..255, in use: 128
STID range: 256..319, in use: 10
ATID range: 0..127, in use: 0
FTID range: 320..687
HW TID usage: 0 IP users, 128 IPv6 users
```

Figure 123 - HMA Offload connections

VIII. iSER Storport Miniport

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 adapters that are compatible with Chelsio iSER Storport Miniport driver:

- T62100-CR
- T62100-LP-CR
- T6425-CR
- T6225-CR
- T6225-LL-CR
- T6225-SO-CR (*Memory-free; 32 IPv4/IPv6 offload connections supported*)
- T580-CR
- T580-LP-CR
- T540-CR
- T540-LP-CR
- T540-BT
- T520-CR
- T520-LL-CR
- T520-BT

1.2. Software Requirements

1.2.1. Windows Requirement

Currently iSER Storport Miniport driver is available for the following Windows versions.

- Server 2019
- Server 2016
- 10 Client

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

2. Software/Driver Configuration and Fine-tuning

2.1. Configuring iSER Target

The iSER target must be configured before establishing connection. For more information, please refer Unified Wire for Linux documentation available at the [Chelsio Download Center](#).

Important *Disable iWARP Port Mapper (iwpmid) service on the target machine. On RHEL7.X machines, use the below command.*

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

2.2. Configuring iSER Initiator

The iSER interface uses the same IP as the NIC interface. If not already done, assign an IP to Chelsio interface as described in the [Assigning IP Address](#) section, before proceeding.

If a vSwitch is created on the physical port, disable and enable the physical port using the below command for the iSER driver to pick up the IP address.

```
C:\Users\Administrator> restart-netadapter <physical port name>
```

2.3. iSER Target Discovery and Login

The following example describes the method to discover and login to iSER target using the Chelsio iSER interface:

- i. To start the initiator configuration, go to **Control Panel** and click **iSCSI Initiator** in **Administrative Tools**.




Figure 124 – Administrative Tools: iSCSI Initiator




Figure 125 - iSCSI Initiator Properties: Targets tab

ii. Click the **Discovery** tab and then **Discover Portal...**

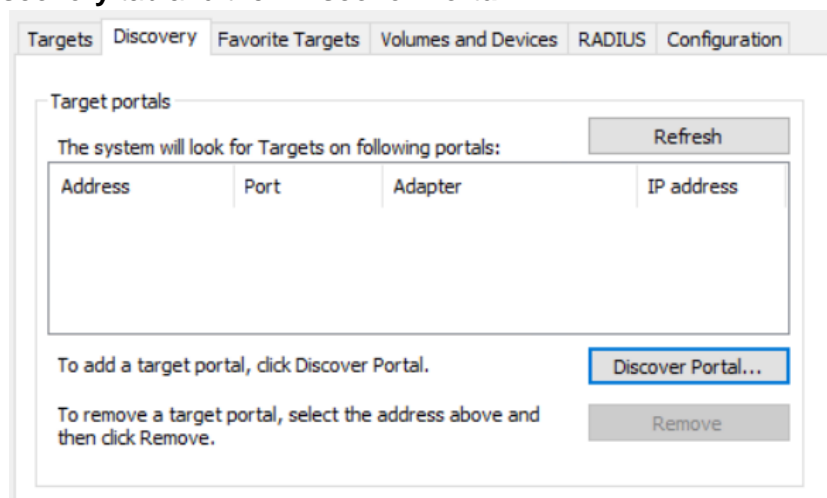


Figure 126 - iSCSI Initiator Properties: Discovery tab


iii. Click on **Advanced...**

Figure 127 - Discovery Target Portal window

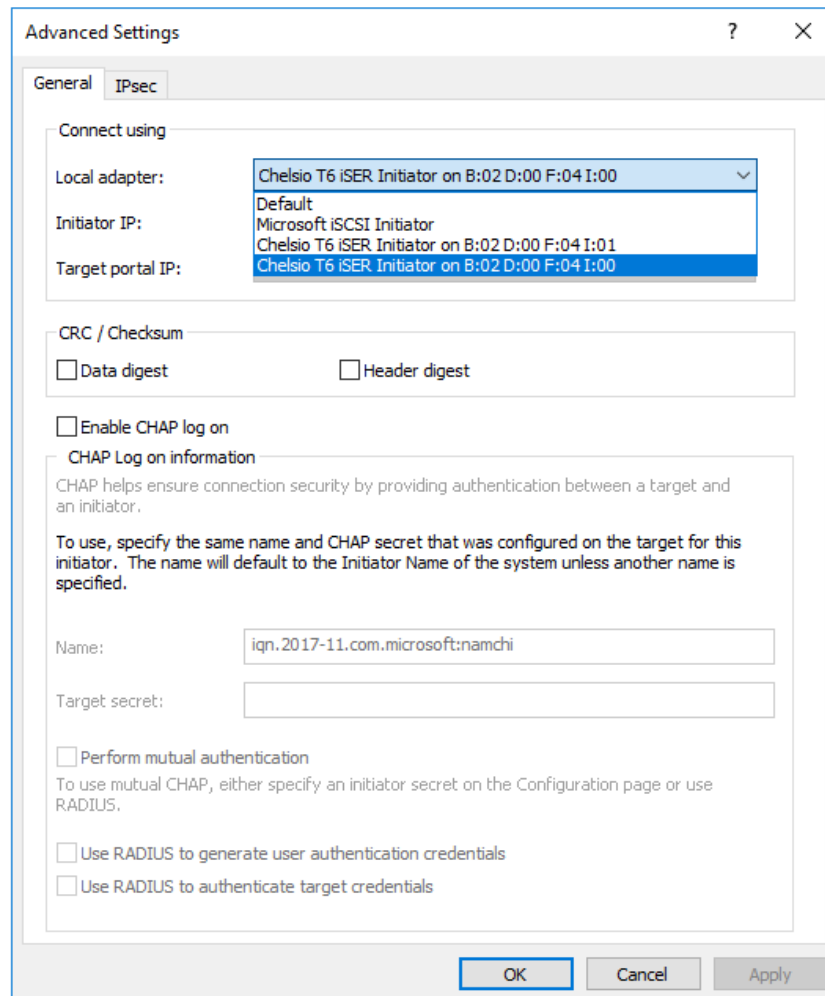
iv. In the **Advanced Settings** window, select **Chelsio iSER Initiator** as the local adapter, the corresponding Chelsio interface IP as the initiator IP and click **OK**.

Figure 128 - Discovery Target Portal Advanced Settings: Selecting Local adapter




Figure 129 - Discovery Target Portal Advanced Settings: Specifying Initiator IP

- v. On the **Discover Target Portal** Window, enter the IP address (DNS name not supported) of the target machine and the corresponding port number, and click **OK**. Note that the default port number for iSER traffic is 3260.




Figure 130 - Adding Target portal

- vi. Once target portal is added, details like target and initiator machine IP, port number and Chelsio interface IP will be displayed.

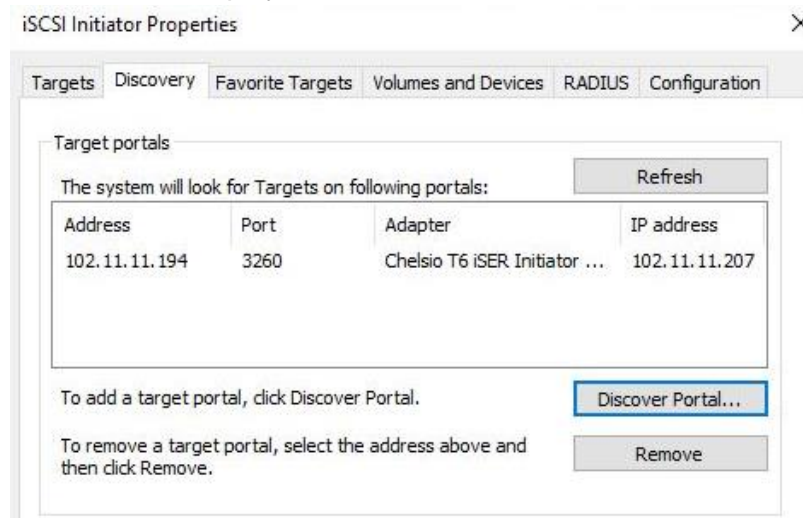


Figure 131 - Target portal added

vii. Click the **Targets** tab to see the list of targets available. Choose a target and click **Connect**.




Figure 132 - Targets tab displaying list of available targets

viii. **Connect To Target** window pops up displaying the Target name.

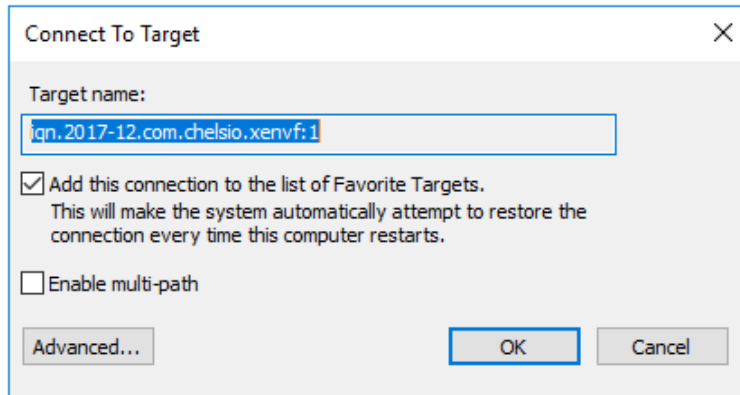


Figure 133 - Connecting to Target

ix. Click **Advanced...**, select Chelsio iSER interface and corresponding Initiator and Target portal IPs. Click **OK**.



Figure 134 - Advanced Settings

x. Click **OK** to close the **Connect To Target** window.

xi. After successful login, target's status will change from *Inactive* to *Connected*.

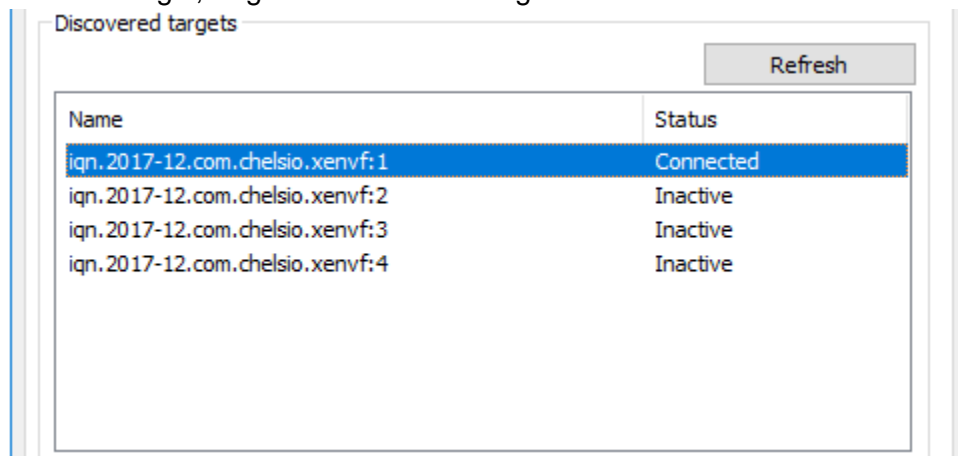


Figure 135 - iSER target connected

- xii. The discovered LUN will appear as an offline disk in the **Disk Management** application. Right-click the disk and select *Online*.




Figure 136 - Disk Management: making volume online

2.4. HMA

The Host Memory Access(HMA) implements a bridge between the Memory Arbiter (MA) and the PCIE so that the Host Memory is available to all clients. The translation from the MA memory map to the Host Memory map is performed based on the client number and the address of the transaction, both of which are used as index to look up a 64-bit offset within the Host Memory.

Currently 32 IPv4/IPv6 iSER Initiator Offload connections are supported on T6225-SO-CR adapter. You can view the number of connections offloaded by using

```
C:\Users\Administrator> cxgbtool <nicIface> hardware tid_info
```


IX. NVMe Storport Miniport

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 adapters that are compatible with Chelsio NVMe Storport Miniport driver:

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

1.2. Software Requirements

1.2.1. Windows Requirement

Currently Chelsio NVMe Storport Miniport driver is available for following Windows versions:

- Server 2019
- Server 2016
- 10 Client (Enterprise and "Pro for Workstations" editions)

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

2. Software/Driver Configuration and Fine-tuning

2.1. Connecting to NVMe target

Configure the NVMe target machine with the IP Address, Target name, disks etc.

! Important *Disable iWARP Port Mapper (iwpmid) service on the target machine. On RHEL7.X machines, use the below command.*

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

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

- i. Configure the Initiator with DHCP or static IP Address. If it is connected to a network with a DHCP server, Chelsio network interfaces will acquire DHCP IPs automatically. To assign static IPs, follow the steps mentioned below:
 - a. Double click on the Network Connections icon and choose the Chelsio card entry and double click it.
 - b. Click on the Properties button from the Local Area Connection X Status.
 - c. Select "Internet Protocol (TCP/IP)" from the list and click on Properties button below it.
 - d. From the Internet Protocol (TCP/IP) Properties window, assign an IP Address (e.g. 192.169.1.10) and subnet mask (e.g. 255.255.255.0).
 - e. Click on Ok and close on the other window.

Check to see if you can ping to the target machine.

- ii. *nvmetool* utility should be used to connect to the target. In case of Installer, it will be copied to <system_drive>\Windows\System32 folder during installation.

If you are using the zip package, copy *nvmetool.exe* from *ChelsioUwire_x.x.x.x_WIN_yyy.y.yy* folder to <system_drive>\Windows\System32 folder.

- iii. Scan the system for available NVMe initiator devices.

```
PS C:\Users\Administrator> nvmetool.exe findnvmehosts
```

```

PS C:\Users\Administrator> nvmetool.exe findnvmehosts
NVMeOF Initiator Utility. Version 1.0
Copyright (c) 2019 Chelsio Communications

NVMeOF Host found at PortId: 0
  Interface Name   : SupportDevice0
  IPv6 Address     : fe80::6153:b05:cd4d:e11c
  IPv4 Address     : 10.1.1.7
NVMeOF Host found at PortId: 1
  Interface Name   : vEthernet (p1)
  IPv6 Address     : fe80::489:8a2c:fa26:c7c6
  IPv4 Address     : 10.2.2.7
    
```

Figure 137 - List NVMe Adapters

- iv. Discover the target by specifying the PortId and Target IP Address.

```

PS C:\Users\Administrator> nvmetool.exe DiscoverTargets <PortId>
<TargetIPAddress>
    
```

```

PS C:\Users\Administrator> nvmetool.exe DiscoverTargets 0 10.1.1.5
NVMeOF Initiator Utility. Version 1.0
Copyright (c) 2019 Chelsio Communications

Connected to Target: nqn.2014-08.org.nvmeexpress.discovery, TargetId: 0
Number of Records      : 32
===== Discovery Log Entry 0 =====
NVMe Transport Type    : rdma
Address Family         : ipv4
Subsystem type         : nvme subsystem
Transport Req          : not specified
Port Id                : 1
Controller Id         : 65535
Admin Max SQ Size      : 32
Subsystem Qualified Name : nvme-ram0
Transport Address       : 10.1.1.5
Transport Service Id   : 4420
===== Discovery Log Entry 1 =====
NVMe Transport Type    : rdma
Address Family         : ipv4
Subsystem type         : nvme subsystem
Transport Req          : not specified
Port Id                : 1
Controller Id         : 65535
Admin Max SQ Size      : 32
Subsystem Qualified Name : nvme-ram1
Transport Address       : 10.1.1.5
Transport Service Id   : 4420
    
```

Figure 138 - Discover NVMe target

- v. Connect to the target by specifying the PortId, Target Name and Target IP Address.

```

PS C:\Users\Administrator> nvmetool.exe ConnectTarget <PortId> <Targetnqn>
<TargetIPAddress>
    
```

```
PS C:\Users\Administrator> nvmetool.exe connecttarget 0 nvme-ram0 10.1.1.5
NVMeOF Initiator Utility. Version 1.0
Copyright (c) 2019 Chelsio Communications

Connected to Target: nvme-ram0, TargetId: 0
```

Figure 139 - Connect to NVMe Target

On successful connection with the target, the NVMe disks will be available for use.

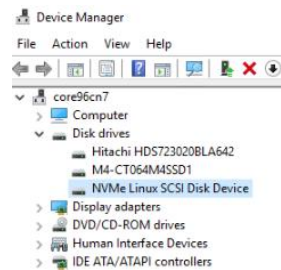


Figure 140 - NVMe Disks available

vi. List the successful connections.

```
PS C:\Users\Administrator> nvmetool.exe ListTargets <PortId>
```

```
PS C:\Users\Administrator> nvmetool.exe listtargets 0
NVMeOF Initiator Utility. Version 1.0
Copyright (c) 2019 Chelsio Communications

Target = 0
VID/SSVID           : 0x0/0x0
Serial Number       : d08b518b1c43ec1d
Model Number        : Linux
Firmware Revision   : 4.14.61
Max Data Transfer   : 0
Version             : 1.3
Number Of Namespace : 1
SubNQN              : nvme-ram0
```

Figure 141 - Listing NVMe connections

vii. For targets to be persistent across reboots, execute the below command:

```
PS C:\Users\Administrator> schtasks /create /sc ONLOGON /tn "addnvmerestore"
/tr "nvmetool restorepreviousconfig"
```

2.2. Disconnecting from NVMe target

To disconnect from the NVMe target, use the below command.

```
PS C:\Users\Administrator> nvmetool.exe RemoveTarget <PortId> <TargetId>
```

```
PS C:\Users\Administrator> nvmetool.exe removetarget 0 0
NVMeOF Initiator Utility. Version 1.0
Copyright (c) 2019 Chelsio Communications

Successfully disconnected from Target: nvme-ram0
```

Figure 142 - Disconnecting from NVMe Target

2.3. HMA

The Host Memory Access (HMA) implements a bridge between the Memory Arbiter (MA) and the PCIE so that the Host Memory is available to all clients. The translation from the MA memory map to the Host Memory map is performed based on the client number and the address of the transaction, both of which are used as index to look up a 64-bit offset within the Host Memory.

Currently 256 IPv4/128 IPv6 NVMe Initiator Offload connections are supported on T6225-SO-CR adapter. You can view the number of connections offloaded by using:

```
C:\Users\Administrator> cxgbtool <nicIface> hardware tid_info
```

Configure the number of IoQueues while connecting to more number of targets. By default, number of IoQueues will be equal to the number of CPU cores.

```
PS C:\Users\Administrator> nvmetool.exe ConnectTarget <PortId> <Targetnqn>
<TargetIPAddress> [<NumIoQueues>]
```

Example: While connecting to 64 NVMe targets from an NVMe Initiator machine with 8 CPU cores, by default 8 IoQueues will be used per target, a total of 512 (8x64) connections. While using T6225-SO-CR (with maximum 256 offload connections) NVMe Initiator, IoQueues should be adjusted to a value of 4 or lesser to successfully connect to all the 64 targets.

X. 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.

 **Note** Only IEEE configuration supported.

1.1. Hardware Requirements

1.1.1. Supported Adapters

The following are the adapters that are compatible with Chelsio's DCB feature:

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

1.2. Software Requirements

1.2.1. Windows Requirements

Currently Chelsio's DCB feature is available for the following Windows versions:

- Server 2019
- Server 2016
- 10 Client
- Server 2012 R2

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

2. Software/Driver Configuration and Fine-tuning

Chelsio network adapters work seamlessly with any legacy switches and do not require the complexity of configuring Data Center Bridging (DCB) protocols either on adapter or network switch. In case DCB is already configured on the switch, the adapter will automatically negotiate the PFC settings configured on the switch.

2.1. Network QoS (optional)

Network Quality of Service (QoS) is an advanced Windows feature that can be used to distribute bandwidth between different kinds of outgoing traffic. This feature ensures efficient usage of resources and minimizes the impact of bandwidth congestion. The bandwidth percentage can be configured either on switch or the host.

2.1.1. Enabling QoS on Host

Network QoS is disabled by default. Use **NDIS Miniport Driver Parameters** (Device manager → Chelsio Network Adapter Properties → Advanced Tab → NDIS QoS) to enable it.

Note Chelsio QoS requires a minimum of 8 cores to work. In case of multiple sockets, where the number of cores per socket is less than 8, Chelsio adapters must be configured to utilize all the available cores before enabling QoS.

Run the following command on all Chelsio ports:

```
PS C:\Users\Administrator> Set-NetAdapterRss -Name <interface name> -  
MaxProcessorGroup 3
```

2.1.2. Configuring Bandwidth Allocation

- **Switch**

Here is an example of setting bandwidth allocation on a Dell Force10 switch:

- i. Log in to the switch and enter the configuration mode:

```
Login: admin  
Password:  
Force10#enable  
Force10#configure
```

ii. Enable DCB and LLDP:

```
Force10(conf)#dcb enable
Force10(conf)#protocol lldp
Force10(conf-lldp)#no disable
```

iii. Set DCBx version to *auto*:

```
Force10(conf-lldp)#dcbx version auto
```

iv. Specify the intervals between hello packets:

```
Force10(conf-lldp)#hello 5
```

v. Create a DCB map. For example, here we are configuring 2 groups/classes:

- Group 0 has BW of 80% with *pfc* enabled.
- Group 1 had BW of 20% with *pfc* disabled.
- Assigning priority groups: Priority 5-7 are under group 0 and 0-4 are under group 1.

```
Force10(conf)#dcb-map win_dcb
Force10(conf-dcbmap-win_dcb)# priority-group 0 bandwidth 80 pfc on
Force10(conf-dcbmap-win_dcb)# priority-group 1 bandwidth 20 pfc off
Force10(conf-dcbmap-win_dcb)# priority-pgid 1 1 1 1 1 0 0 0
```

vi. Enter interface configuration mode and configure protocol LLDP:

```
Force10(conf)#interface tengigabitethernet 0/17
Force10(conf-if-te-0/17)#protocol lldp
Force10(conf-if-te-0/17-lldp)#no disable
```

vii. Set DCBx version to *auto*:

```
Force10(conf-if-te-0/17-lldp)#dcbx version auto
```

viii. Apply the DCB-MAP created:

```
Force10(conf-if-te-0/17)#dcb-map win_dcb
```

Now the host adapter will honor the bandwidth allocation settings configured on the switch.

- **Host**

Here is an example of setting bandwidth allocation on the host:

i. Open PowerShell with administrative privileges and run the following command:

```
PS C:\Users\Administrator> Install-WindowsFeature -Name "Data-Center-Bridging" -IncludeAllSubFeature -IncludeManagementTools
```

ii. Create new policy:

```
PS C:\Users\Administrator> New-NetQosPolicy -Name rdma -PriorityValue 5 -NetDirectPortMatchCondition 445
```

iii. Create new rule with bandwidth percentage:

```
PS C:\Users\Administrator> New-NetQosTrafficClass -Name rdma -Priority 5 -Algorithm ETS -BandwidthPercentage 80 -Verbose
```

The above command allocates 80% of the bandwidth to the policy named 'rdma'.

To change the percentage of an existing ETS rule:

```
PS C:\Users\Administrator> Set-NetQosTrafficClass -Name rdma -Priority 5 -Algorithm ETS -BandwidthPercentage 10 -Verbose
```

XI. Appendix

1. Troubleshooting

- **Installer issues**

In case of any failures while running the Chelsio Unified Wire Installer, please collect the below:

- Screenshot of the failure.
- Installer logs from `<system_drive>\Windows\Temp\ChelsioUwire_x.x.x.x_WIN_y.y.y_logs` directory. Compress the directory to a zip package.

- **Logs collection**

In case of any other issues, follow the below procedure to collect the standard logs:

- Download the *Get-NetView.psm1* file from <https://github.com/microsoft/Get-NetView> and run the below commands.

```
PS C:\Users\Administrator\Desktop> Set-ExecutionPolicy Unrestricted
PS C:\Users\Administrator\Desktop> Import-Module Get-NetView.psm1
```

- Run the below command and the logs will be collected in a zip package.

```
PS C:\Users\Administrator\Desktop> get-netview
```

```
Diagnostics Data:
-----
Get-NetView
Version: 2018.07.13.0
SHA256: C350E2F057669276AC1F77D8123A5805E82AD32E318E22395F48FA128A16F742
C:\Users\Administrator\Desktop\msdbg.CORE96CN2-2018.09.20_03.53.50.zip
Size: 12.01 MB
C:\Users\Administrator\Desktop\msdbg.CORE96CN2
Size: 157.06 MB
Dirs: 128
Files: 416
Execution Time:
-----
5 Min 24 Sec
```

Figure 143 - Get-NetView Logs

- Driver Event tracing logs can be collected using the below procedure:
 - Start the tracing.

```
PS C:\Users\Administrator\Desktop> cxgbtool debug trace nic 0xffff 5 vbd
0xffff 5 iscsi 0xffff 5 iser 0xffff 5 nvme 0xffff 5
```

- Rerun the test to hit the observed issue.
- Stop the tracing.

```
PS C:\Users\Administrator\Desktop> cxgbtool debug trace stop
```

cxgbtrace.etl file will be collected in the same path.

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

2. Chelsio End-User License Agreement (EULA)

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

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

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

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

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

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

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

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

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

6. Limited Warranty. If Chelsio furnishes the Software to you on media, Chelsio warrants only that the media upon which the Software is furnished will be free from defects in

material or workmanship under normal use and service for a period of thirty (30) days from the date of delivery to you.

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

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

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

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

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

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

11. Government Restricted Rights. The Software is subject to restricted rights as follows. If the Software is acquired under the terms of a GSA contract: use, reproduction or disclosure is subject to the restrictions set forth in the applicable ADP Schedule contract. If the Software is acquired under the terms of a DoD or civilian agency contract, use, duplication or disclosure by the Government is subject to the restrictions of this Agreement in accordance with 48 C.F.R. 12.212 of the Federal

Acquisition Regulations and its successors and 49 C.F.R. 227.7202-1 of the DoD FAR Supplement and its successors.

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

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

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