

Chelsio
Communications
Accelerate



WELCOME TO

NVMe™ Developer Days



We put the iWARP
in
NVMe-oF

NVMe-oF with iWARP

- NVMe/iWARP RDMA over Ethernet
- Extend storage fabric beyond PCIe reach
- High Performance, scalable low latency API
- RDMA performance and efficiency across standard Ethernet without the DCB tax and complexity



Why is this important?

- iWARP is....
 - Easier to deploy
 - Uses legacy infrastructure
 - Performance
 - Price
 - Stable
 - Scalable



Supported across the Chelsio product line

- 10 Gb/sec to 100 Gb/sec
- Kernel mode part of the Linux kernel today!
- SPDK mode coming soon
- Supported on multiple platforms
 - x86_x64/EPYC
 - ARM



NVMe-oF Kernel Mode Bandwidth



Chelsio T62100-CR 100Gb/sec

1 Target

2 Initiators

1GB RAM DISK

8 Total Connections

2 Target devices per initiator

Complete set of Benchmark results
including competitive comparison
available upon request

Chelsio
Communications
Accelerate



NVMe-oF SPDK Bandwidth



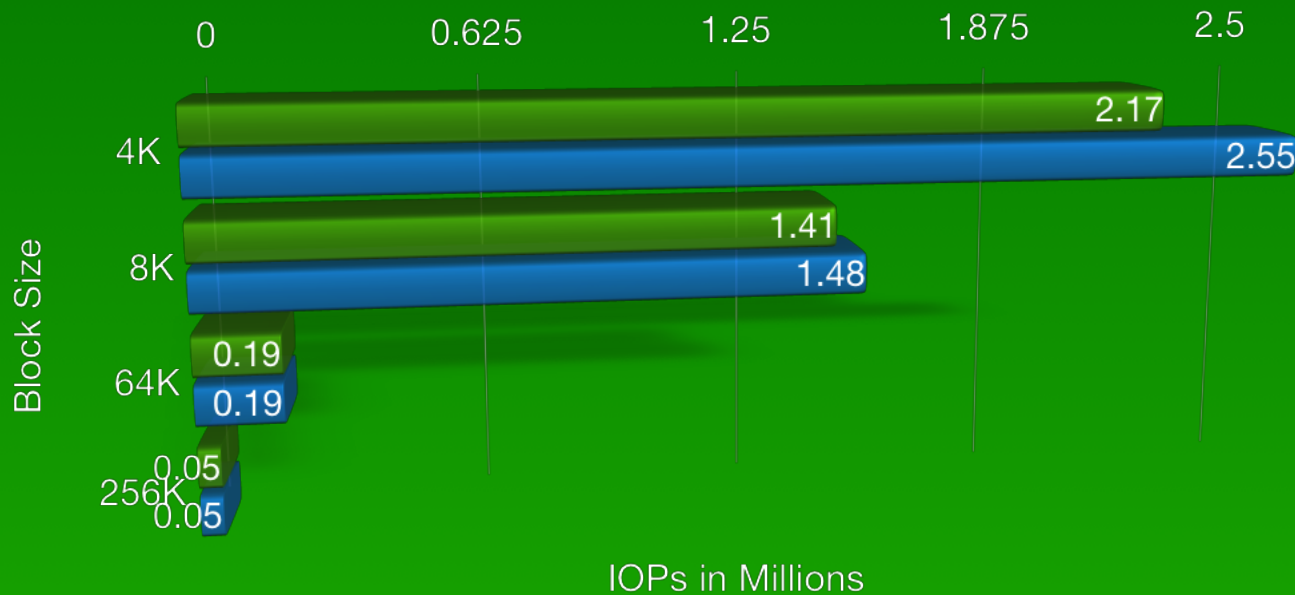
Chelsio T62100-CR 100Gb/sec
1 Target
2 Initiators
1GB RAM DISK
8 Total Connections
2 Target devices per initiator

Complete set of Benchmark results
including competitive comparison
available upon request

Chelsio
Communications
Accelerate



NVMe-oF Kernel Mode IOPs



Chelsio T62100-CR 100Gb/sec

1 Target

2 Initiators

1GB RAM DISK

8 Total Connections

2 Target devices per initiator

Complete set of Benchmark results
including competitive comparison
available upon request

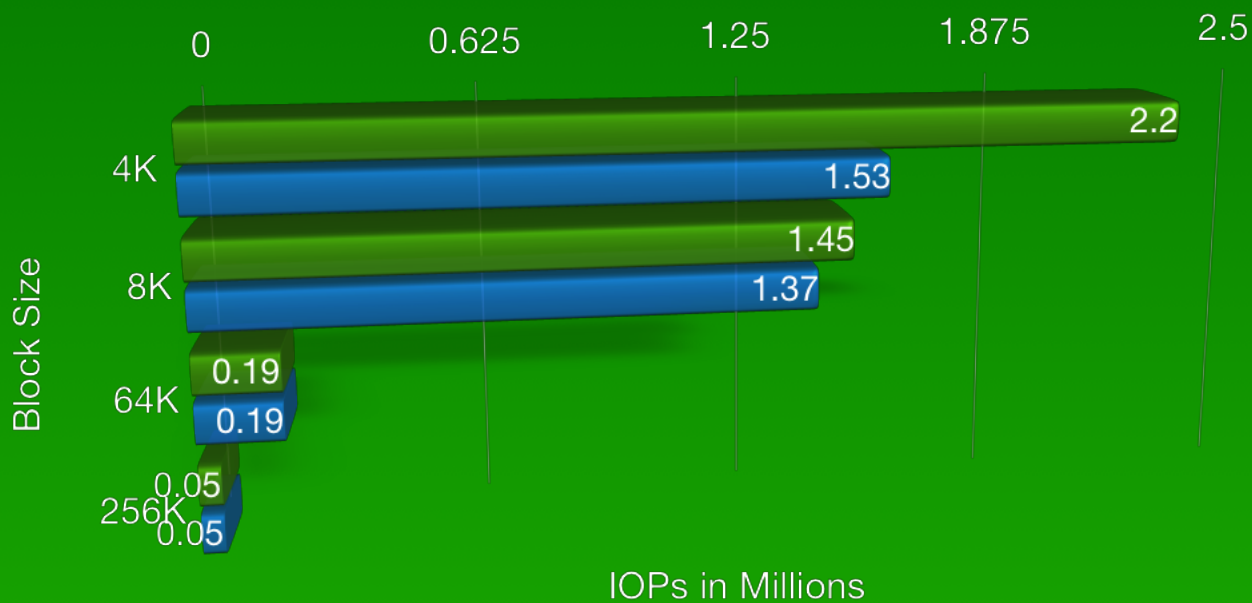
Read

Write

Chelsio
Communications
Accelerate



NVMe-oF SPDK IOPs



Chelsio T62100-CR 100Gb/sec
1 Target
2 Initiators
1GB RAM DISK
8 Total Connections
2 Target devices per initiator

Complete set of Benchmark results
including competitive comparison
available upon request

Read

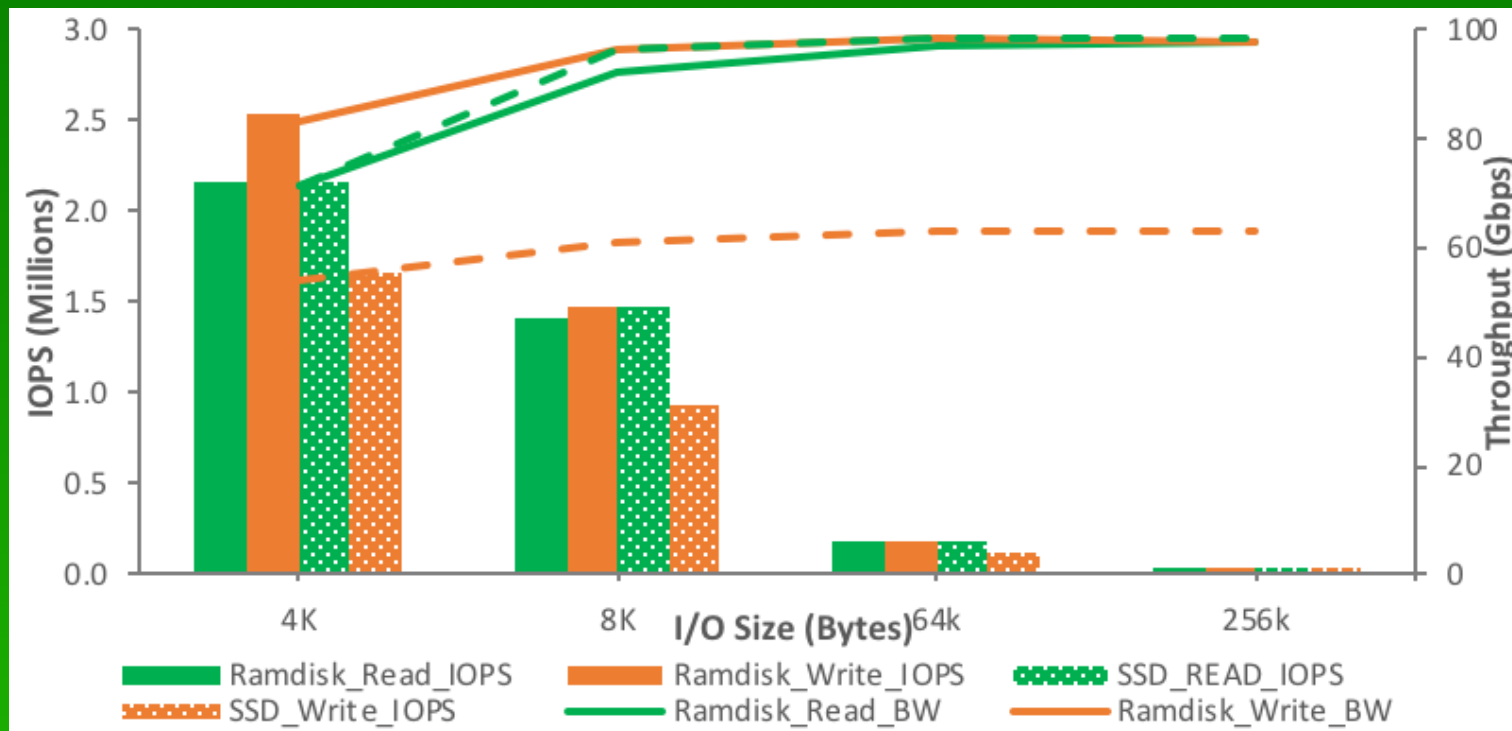
Write

Chelsio
Communications
Accelerate




NVMe kernel mode over 100GbE iWARP Fabric

RamDisk and SSD IOPs & Throughput vs I/O size




NVMe kernel mode over 100GbE iWARP Fabric

Latency vs # of connections



Benchmark: 100G NVMe/TCP Host Stack vs TOE



Test Setup

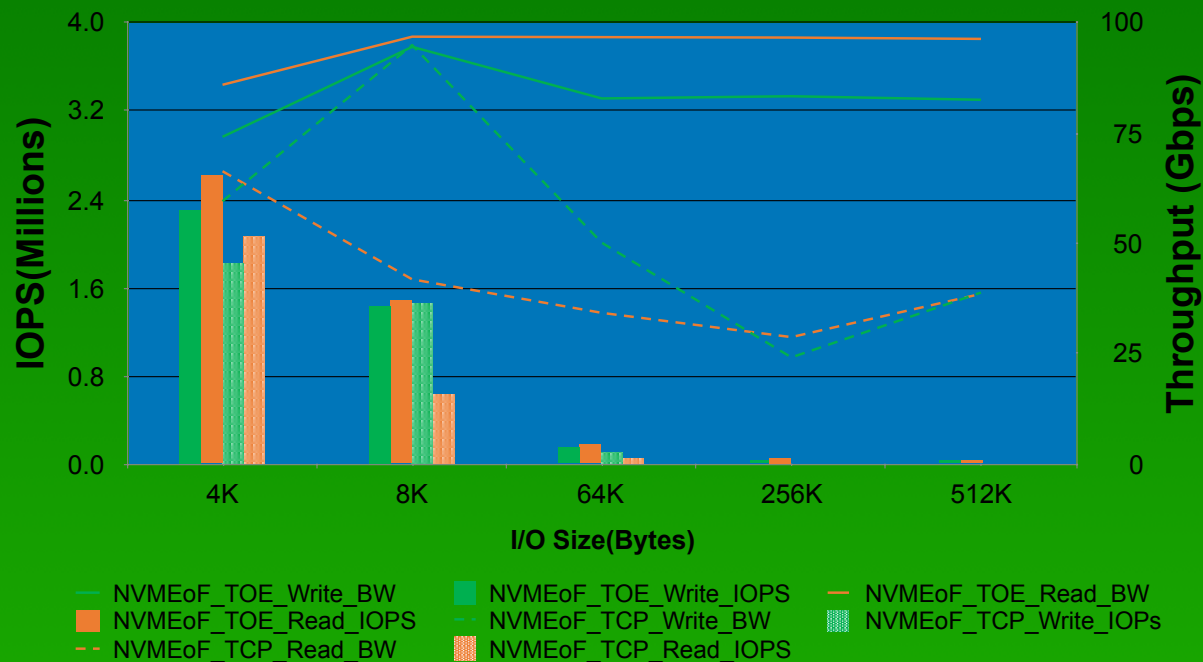
• The target machine

- 2 Intel Xeon E5-2687W v4
 - 12-core @ 3.00GHz (HT disabled)
- 128GB of RAM
- Chelsio T62100-CR (2 x 100Gbps)
- RHEL 7.3 (4.18.0-rc6 kernel)

• The initiator machines

- 1 Intel Xeon E5-1620 v4
 - 4-core processor @ 3.50GHz (HT enabled)
- 32GB of RAM
- Chelsio T62100-CR (2x100 Gbps)
- RHEL 7.3 (4.18.0-rc6 kernel)

Benchmark: 100G NVMe/TCP Host Stack vs TOE



Summary

- Read throughput is line rate for TOE
- Write throughput near line rate for TOE
- 2.6M IOPS at 4K I/O size for TOE
- ~12 μSec delta latency between local and remote storage

What's New? SoftiWARP!

- NVMe/iWARP RDMA over Ethernet in Software
- Can run on any L2 NIC (No RNIC Required)
- High Performance, scalable API
- RDMA performance and efficiency across standard Ethernet without the DCB tax and complexity
- Get the Benefits of RDMA across your cluster without having to equip all nodes with RNICs

Chelsio
Communications
Accelerate




More on SoftiWARP

- Builds in an automatic second source to any offloaded solution
- Unlike RoCE, this solution interoperates with ALL iWARP Cards
- Enables NVMe-oF and iSER to benefit from ubiquity of a software solution
- SoftiWARP is available under GPL and BSD Licenses
- Surprisingly Good Performance

Chelsio
Communications
Accelerate



Ask for your Demo NOW



Test Setup

- **The target machine**

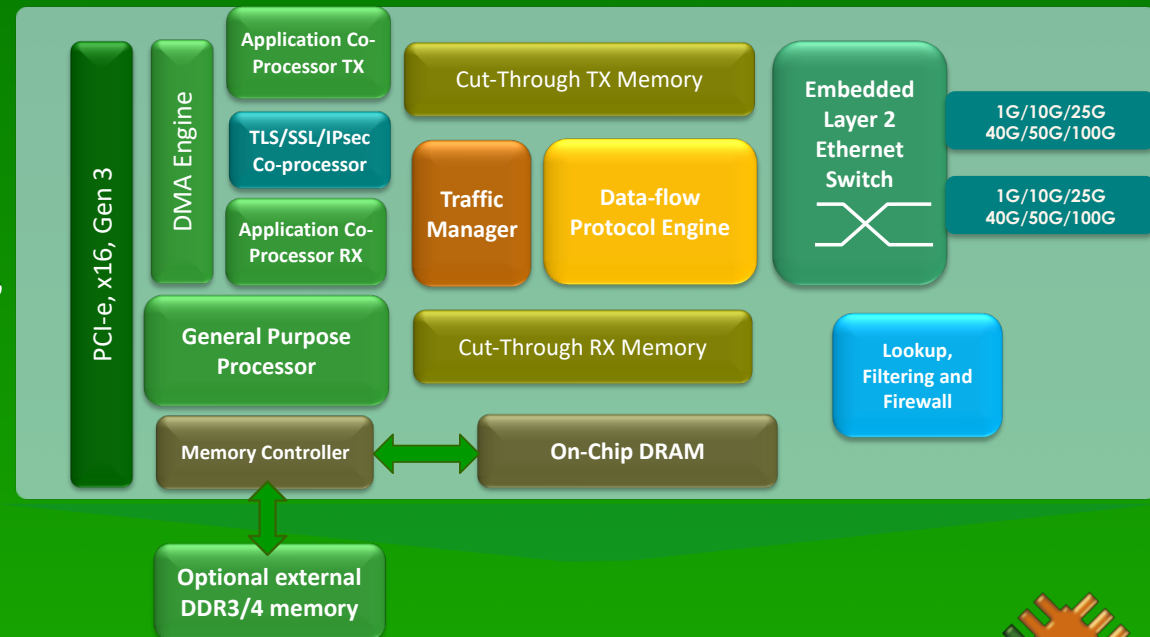
- 2 Intel Xeon E5-2687W v3
 - 8-core @ 3.00GHz (HT enabled)
 - 64 GB of RAM
- Chelsio T62100-CR (2 x 100Gbps)
- CentOS 7.6 (4.14.85 kernel)

- **The initiator machines**

- 2 Intel Xeon E5-2687W v3
 - 8-core processor @ 3.00GHz (HT enabled)
- 64GB of RAM
- Chelsio T62100-CR (2x100 Gbps)
- CentOS 7.6 (5.0 rc5 kernel)

Chelsio T6 ASIC

- 2 ports - 1/10/25/40/50/100 Gbit
- PCI Gen 3 x16
- On-Board TLS/SSL/IPsec Co-processor
- 264 Port Embedded Ethernet Switch
- On-NIC Filtering, Traffic Management, Traffic Pacing
- Single processor data-flow pipelined architecture
- Up to 1M connections
- Concurrent Multi-protocol Operation



Chelsio
Communications
Accelerate



Chelsio Adapters



Chelsio Unified Wire Adapter Product Selection Guide

Single Adapter with concurrent Multi-Protocol Operation



Chelsio P/N	T62100-CR	T62100-LP-CR	T62100-SO-CR	T61100-OC	T6225-OC	T6225-SO-CR	T6225-LL-CR	
Ports	2x40/50/100G		1x40/50/100G		2x1/10/25G			
Connector Type	QSFP28				SFP28			
Host Interface	PCIe Gen3 x16				PCIe Gen3 x8			
Form Factor	Half Size	Low Profile		OCP		Low Profile		
Concurr. Conn.	32k		256 → 1k		32k			
Power (Typical)	22W	19W	16W	18W	15W	10W	16W	
Airflow	200 LFM (Linear Feet Per minute)							
Chelsio P/N	T580-OC	T580-SO-CR	T580-LP-CR	T580-CR	T6425-CR	T6225-CR		
Ports	2x40G				4x1/10/25G	2x1/10/25G		
Connector Type	QSFP+ Cage				SFP28			
Host Interface	PCIe Gen3 x8							
Form Factor	OCP	Low Profile		Half Size		Low Profile		
Concurr. Conn.	256 → 1k		32k					
Power (Typical)	12W		17W	21W	26W	13W		
Airflow	200 LFM (Linear Feet Per minute)							
Chelsio P/N	T520-OC	T520-SO-CR	T520-CR	T520-LL-CR	T540-CR	T540-BT	T520-BT	
Ports	2x1/10G				4x1/10G	4x100M/1G/10G	2x100M/1G/10G	
Connector Type	SFP+ Cage				RJ-45			
Host Interface	PCIe Gen3 x8							
Form Factor	OCP	Low Profile		Half Size		Low Profile		
Concurr. Conn.	256 → 1k		32k					
Power (Typical)	9W		12W	16W	16W	26W	20W	
Airflow	200 LFM (Linear Feet Per minute)							

Features

Windows Server 2012-R2 / 2016
 Windows Storage Spaces Direct (S2D)
 Windows Azure Stack Private Cloud
 Windows Network Direct
 Windows SMB Direct
 Windows Nano Server
 Windows Client RDMA
 Low Latency L2 NIC
 TCP Offload Engine (TOE)
 TLS/SSL, DTLS, IPsec, SMB 3.X crypto¹
 iSCSI Offload (Initiator and Target)
 FCoE Offload
 iWARP RDMA Offload
 NVMe Over Fabrics
 NVMe JBOD (Just a Bunch of Flash)
 GPUDirect RDMA
 Wire Direct Low Latency Suite
 SRIOV / Virtualization
 DPDK / PacketDirect / Netmap
 OVS / OpenFlow Offload
 NVGRE / VxLAN/ Geneve
 Media StreamEngine²
 Traffic Management / DCB
 Packet Classification
 Sniffer / Tracer
 Multi / Single adapter Failover

¹ Supported only on T6 family of adapters
² Not supported by SO (Server Offload) adapters

LL: Low Latency LP: Low Profile
 SO: Server Offload

• Chelsio Communications

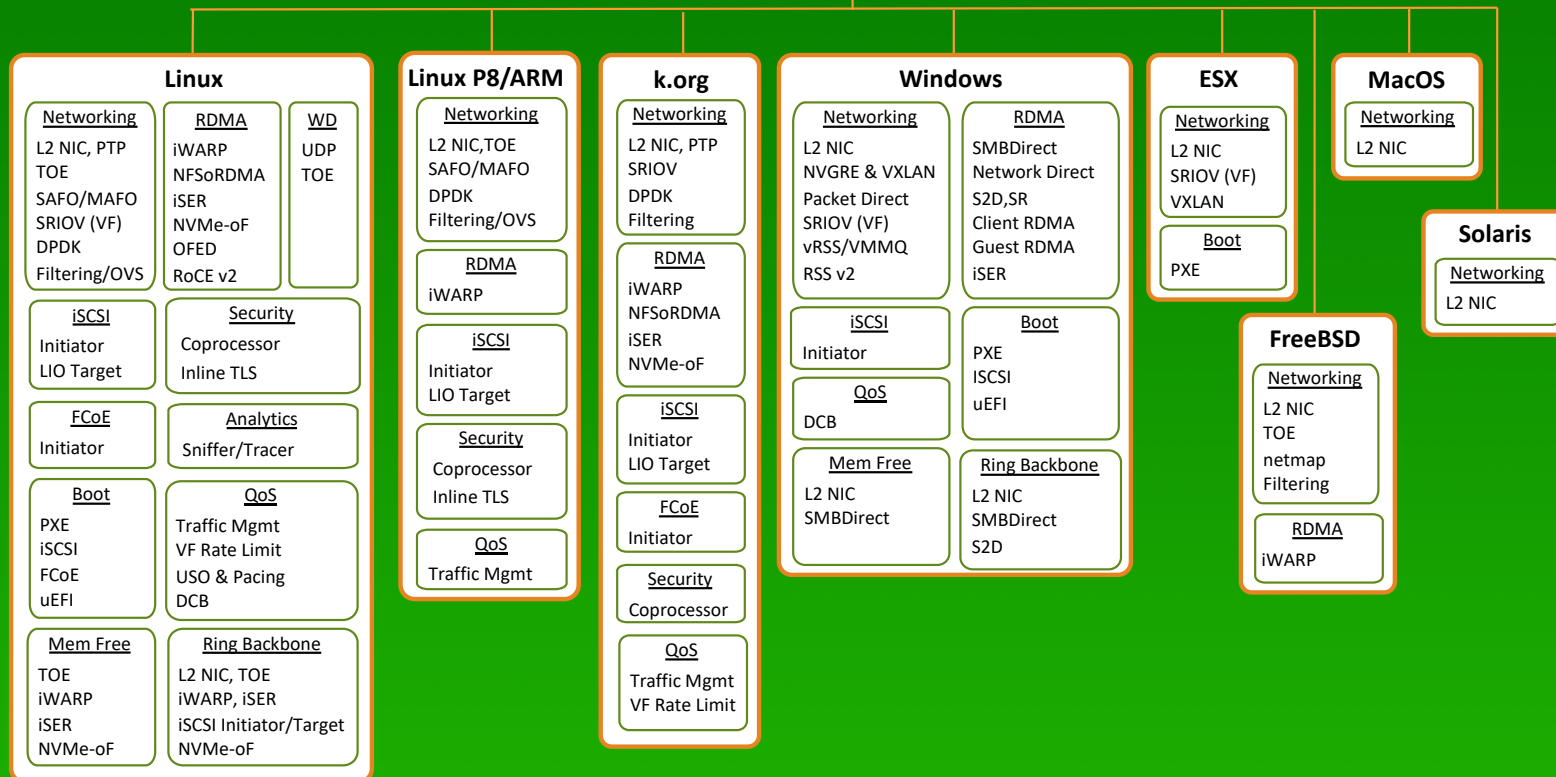
• www.chelsio.com

• sales@chelsio.com

• +1-408-962-3600

Software Ecosystem

Chelsio Software Suite



store.chelsio.com

Ocean OEM My Optimum Voice MSCU Ocean TV Monmouth Cou...stry | Login Mac Chelsio Travel Hockey HP Mac News My sites Apple eBay Amazon Next Inbox - Outloo...light version Apple Yahoo!

Chelsio Online Store for 10/25/40/100Gb Ethernet Adapters/Accessories

Chelsio
Communications
Accelerate

Adapters ▾ Accessories ▾ Product Selector FAQs

ADAPTERS

- T6 Adapters +
- T5 Adapters +
- Legacy Adapters

ACCESSORIES

- Cables
- Transceivers

FEATURED COLLECTIONS [View all Collections](#)

- T6 Unified Wire Adapters
- T6 Server Offload Adapters
- T5 Unified Wire Adapters
- T5 Server Offload Adapters
- Accessories

T6 UNIFIED WIRE ADAPTERS [View all Items](#)

- SALE** T62100-CR: 2-port Half Size 40/50/100GbE Unified Wire, Enhanced TOE & iSCSI Adapter with PCIe 3.0 x16
- SALE** T62100-LP-CR: 2-port Low Profile 40/50/100GbE Unified Wire Adapter with PCIe 3.0 x16 Interface, 32K
- SALE** T6225-CR: 2-port Low Profile 10/25GbE Unified Wire Adapter with PCIe 3.0 x8 Interface, 32K
- SALE** T6225-LL-CR: 2-port Low Latency Low Profile 10/25GbE Unified Wire Adapter with PCIe 3.0 x8 Interface, 32K

Chelsio
Communications
Accelerate

*Check out
the
Chelsio Webstore
store.chelsio.com*



sales@chelsio.com

www.chelsio.com

For support and other information

support@chelsio.com