

High Performance, Quad Port 10GbE Unified Wire Adapter

Enables TCP, UDP, iSCSI, iWARP and FCoE Offload Over Single Unified Wire with SR-IOV, EVB/VNTag and DCB

Overview

Chelsio's T540-BT is a quad port 10GBase-T 10 Gigabit Ethernet Unified Wire adapter with PCI Express 3.0 host bus interface, optimized for storage, cloud computing, HPC, virtualization and other datacenter applications.

The fifth-generation T5 ASIC technology from Chelsio provides the highest performance available and dramatically lowers host-system CPU communications overhead with on-board hardware that offloads TCP/IP, iSCSI, FCoE and iWARP RDMA processing from its host system. This will enable savings in host CPU acquisition costs, power and operational costs, and dramatically increases system performance.



A large portion of offloads enabled by T540-BT are based on standard TCP, IP, UDP protocols (such as iSCSI and iWARP), and thus can operate with a software peer, or be replaced with a software solution at lower performance, thus providing the requisite reliability for enterprise customers, and allowing incremental installs in the datacenter. It will work with any legacy switch infrastructure and does not rely on new features such as DCB, PFC, ETS, etc.

T540-BT supports IEEE standards-based link aggregation/failover features, as well as inter-adapter failover techniques that make it ideal for critical network applications requiring redundancy and high-availability capabilities. It also includes an integrated Traffic Manager and a built-in eSwitch for robust network traffic flow control, traffic classifier, load balancer, Firewall capabilities at the port level.

FCoE and iSCSI support in T540-BT benefit from high reliability features that include memory ECC, data path CRC and T10-DIX offload, in addition to the checksums and CRC available at different protocols layers. High performance iSCSI and FCoE provide a drop-in replacement upgrade from legacy SANs to converged networks.

The Unified Wire Solution

T540-BT enables a unified wire for LAN, SAN and cluster applications, built upon a high bandwidth and low latency architecture along with a complete set of storage and cluster protocols operating over Ethernet (iSCSI, NVMe-oF, FCoE and iWARP). A unified wire means having the ability to utilize all offload or non-offload protocols at the same time, over the same link, using the exact same firmware, host software, and adapter. Thus, a given system using T540-BT can be easily targeted at different vertical markets and a variety of useful functions enabled (i.e. dial the bandwidth to a given application and assign access control, while offloading only part of the traffic). The Ethernet-only networking thus reduces the infrastructure costs in network adapters, cables, switches, rack space, power, equipment spares, management tools, planning, networking staff and installation.

Highlights

- PCI Express Gen3 x8
- Low latency user-space I/O
- 32K concurrent offloaded conn. capacity
- Full TCP and UDP offload
- Full iSCSI, FCoE offload
- Full iWARP RDMA offload
- OVS Offload with OpenFlow support
- DPDK (Data Plane Development Kit)
- EVB, Flex10, VNTag
- PCI-SIG SR-IOV
- Integrated media streaming offload
- HW based firewall
- NAT offload
- Traffic filtering & management
- IEEE 1588v2 Precision Time Protocol

Applications

Datacenter Networking

- Scale out servers and NAS systems
- Consolidate LAN, SAN and cluster networks
- Enhanced network and server security

Cloud Computing

- Virtualization features to maximize cloud scaling and utilization
- Runs InfiniBand, FibreChannel applications unmodified over Ethernet
- Cloud-ready functional and management features
- QoS and Traffic Management

Networked Storage

- Enable high performance NAS systems and Ethernet-based SANs
- Develop shared-storage systems providing both file and block level services
- Build high performance storage backend fabrics using Ethernet

High Performance Computing

- Very low latency Ethernet
- High performance RDMA support
- Increase cluster fabric bandwidth
- Deploy Ethernet-only networking for cluster fabric, LAN and SAN

Specifications

Host Interface

- PCI-E Gen3 x8
- MSI-X, MSI and support for legacy pin interrupts

Virtualization

- PCI-SIG SR-IOV
- 128 Virtual and 8 Physical functions
- 140 port virtual switch
- EVB, VEPA, Flex10, VNTag
- 512 MAC addresses
- Offload 802.1 Qbg/h
- NVGRE, VXLAN and GENEVE support

Fifth-Generation Protocol Offload Engine

The T5 is Chelsio's fifth-generation TCP offload (TOE) design, fourth-generation iSCSI design, and third-generation iWARP (RDMA) implementation. With support for the 8 Gbps Gen 3 data rate, it provides 64 Gbps of raw bandwidth to the server. It also provides support for PCIe SR-IOV virtualization with embedded virtual switch.

Complete and Flexible TCP Offload

The T540-BT transport engine executes programmable firmware and is configurable with hundreds of registers for protocol parameters, RFC compliance and offload control. It can offload protocol processing per connection, per-server, per-interface, while simultaneously providing complete stateless offload for non-offloaded connections (processed by operating systems stack running on host CPU). It also provides a flexible direct data placement capability for regular TCP sockets, with all the benefits of zero-copy and kernel bypass without rewriting the applications.

Packet Switching and Routing

T540-BT integrates a high performance packet switch, which allows switching traffic from any of the input ports to any of the output ports (wire-to-wire) and from any of the output ports to any of the input ports (host-to-host).

Robust, Proven Solution

Subjected to thousands of hours of compatibility testing, over a decade of stress testing by several OEM test suites and production deployment in servers, storage systems, and cluster computing, Chelsio's robust, stable protocol offload technology delivers proven performance in a wide range of environments.

Software Drivers

Chelsio offers a full suite of protocol software drivers with the T540-BT adapters. See <http://www.chelsio.com/support> for the latest information. The software suite supports operation in both stateful and stateless offload modes for all major operating systems.

Ordering Information

Model:	T540-BT
Physical Interface:	Quad port RJ-45 / 10GBase-T
Connector:	RJ-45 / Cat-6a
Media:	Twisted Pair

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH CHELSIO PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN CHELSIO'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, CHELSIO ASSUMES NO LIABILITY WHATSOEVER, AND CHELSIO DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND OR USE OF CHELSIO PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. CHELSIO PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS. CHELSIO MAY MAKE CHANGES TO SPECIFICATIONS AND PRODUCT DESCRIPTIONS AT ANY TIME WITHOUT NOTICE.

Copyright © 2017 - Chelsio Communications - All rights reserved.

High Performance RDMA

- Low latency and line rate bandwidth
- Enhanced RDMA primitives including Atomics & Immediate data
- iWARP support in standard OFED
- Microsoft Network Direct support
- Native support for Windows Server 2012-R2, 2016, Azure Stack, Storage Replica, Storage Spaces Direct, Client RDMA, SMB-Direct
- Support for NFS-RDMA, Lustre-RDMA, NVIDIA's GPU-Direct, Hadoop-RDMA

TCP/IP Full Offload

- Full TCP stack including IPv4 & IPv6
- Extensive RFC compliance, fully featured
- Full TCP Proxy between any set of connections
- VLAN support up to 4096 VLAN IDs
- Load balancing and failover capabilities

UDP & Multicast Offload

- UDP Sockets API
- Low user-to-user latency
- Multicast replication on ingress or egress

iSCSI

- iSCSI initiator and target mode stack
- CRC32 offload generation verification
- iSCSI proxy switching based on SCSI CDB
- Full HBA offload
- T10 DIF/DIX support

FCoE

- Full FCoE offload (Initiator or Target)
- Open FCoE offload (Initiator)
- CRC32 offload generation & verification
- Ingress & Egress ACL (Access Control List)
- T10 DIF/DIX support

Stateless Offloads

- TCP/UDP IPv4/6 checksum offload
- TSO, LSO and GSO for IPv4 & IPv6
- VLAN filtering, insertion & extraction
- Line rate packet filtering and attack protection
- Fine granularity time stamping (down to 2ns)
- Nanosecond granularity
- 64b timestamping
- Ethernet Routing (packet header rewrite)
- Packet Tracing and Packet Sniffing

Ethernet

- IEEE 802.3ae (10 GbE)
- IEEE 802.3az Energy Efficient Ethernet
- IEEE 802.3z (1GbE)
- IEEE 802.1p Priority
- IEEE 802.1Q VLAN Tagging
- IEEE 802.1Qbg EVB/VEPA
- IEEE 802.1BR Bridge Port Extension
- IEEE 802.1Qau Congestion Notification
- IEEE 802.3x Flow Control
- IEEE 802.3ad Load-balancing & Failover
- Ethernet II and 802.3 encapsulated frames
- Multiple MAC addresses per interface
- Jumbo Frames up to 9.6 Kbytes

Physical and Environmental

- Fully RoHS Compliant
- Operating Temp: 0° to 45° C or 32° to 113° F
- Operating Humidity: 5 to 95%
- Airflow: 200 lf/m
- Typical power: 26W
- Dimensions without bracket:
6.6in. x 3.95in. or 16.76cm x 10.03cm