T420-BCH

Chelsio Communications
Accelerate

All-In-One daughter card for IBM BladeCenter

Chelsio 10GbE Unified Wire Adapter

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

Highlights

- PCI Express Gen2 x8
- Low Latency
- Supports Up to 1M connections
- Full TCP & UDP offload
- TCP Chimney
- Full iSCSI, FCoE offload
- Full iWARP RDMA offload
- EVB, VEPA, Flex10, VNTag
- PCI-SIG SR-IOV
- Integrated media streaming offload
- HW based firewall in the cloud
- Traffic filtering & management

Applications

Data-Center Networking

- Scale up servers and NAS systems
- Link servers in multiple facilities to synchronize data centers
- Consolidate LAN, SAN, and cluster networks

Cloud Computing

- Virtualization features to maximize cloud scaling and utilization
- Runs InfiniBand, Fibre Channel apps unmodified on Ethernet
- Cloud ready functional and management features
- QoS and Traffic management

Networked Storage

- Enable high performance NAS systems and Ethernet-based IP SANs
- Develop shared-storage systems providing both file- and block-level services

High Performance Computing

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

Overview

Chelsio's T420-BCH is a dual port 10 Gigabit Ethernet Unified Wire adapter with PCI Express 2.0 host bus interface for IBM BladeCenter, optimized for cloud computing, HPC, virtualization, storage, and other data center applications.

The fourth-generation (T4) technology from Chelsio provides the highest 10GbE performance available and dramatically lowers host-system CPU communications overhead with on-board hardware that off-loads TCP/IP, iSCSI, FCoE and iWARP RDMA processing from its host system. T420-BCH frees up host CPU cycles for useful applications. The system achieves increased bandwidth, lower latency, and lower power.

The Unified Wire Solution

With the T420-BCH, Chelsio is enabling a unified wire for LAN, SAN and cluster traffic. This unified wire was made possible by the high bandwidth and low latency of 10GbE combined with storage and cluster protocols operating over TCP/IP (iSCSI, FCoE and iWARP respectively). In parallel, operating systems and hypervisors have incorporated native support for iSCSI and database applications are now supporting file-based storage protocols such as NFS as an alternative to SANs.

T420-BCH includes a full-fledged integrated Traffic Manager for robust flow control, traffic management and QoS.

Fibre Channel over Ethernet (FCoE) provides a transition path from legacy SANs to converged networks. Expanding its unified wire approach, Chelsio has added FCoE hardware support to the new T420-BCH adapter.

The adapter's two ports and IEEE 802.3ad link aggregation/failover features are ideal for critical network applications that require redundancy and high-availability capabilities.

T420-BCH Ethernet-only networking reduces the data center's cost in network adapters, cables, switches, rack space, power, equipment spares, management tools, planning, networking skills, and installation.

T4 - Fourth-Generation Protocol Offload Engine

The T4 represents Chelsio's fourth-generation TCP offload (TOE) design, third-generation iSCSI design, and second-generation iWARP (RDMA) implementation. For the server connection, the T4 chip includes a PCI Express v2.0 ×8 host interface. With support for the 5Gbps Gen2 data rate, the PCIe interface provides up to 32Gbps of bandwidth to the server. T4 also adds support for PCIe I/O virtualization. Most of T4 features are enhanced versions of those found in T3. Features that carry over from T3 include stateless offloads, packet filtering (firewall offload), and traffic shaping (media streaming).

Complete and Flexible TCP Offload

The T4 has hundreds of programmable registers for protocol configuration, RFCs, and offload control. The T420-BCH can offload TCP processing per connection, per-server, per-interface, and globally and simultaneously tunnel traffic from non-offloaded connections to the host processor for the native TCP/IP stack to process. The T420-BCH provides a flexible zero copy capability for regular TCP connections, requiring no changes to the sender, to deliver line rate performance at minimal CPU utilization.

Robust, Proven Solution

Subjected to thousands of hours of compatibility testing, over multiple years of stress testing by several OEM test suites and production deployments in servers, storage systems and cluster computing, Chelsio's robust, stable protocol offload technology delivers proven performance in a wide range of environments. The T420-BCH is generations ahead of competing products.

Packet Switching and Routing

T420-BCH 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).

Software Drivers

Chelsio offers a full suite of protocol software and drivers with the T420-BCH adapters. See www.chelsio.com/support for latest info. The software supports operation in both protocol-offload and non-offload modes.

Ordering info

Model: T420-BCH
Daughter card (CFFh) for IBM BladeCenter

Physical interface: PCIe x8, dual XAUI ports

Specifications

Host Interface
- PCI Express Gen2 x8
- MSI-X, MSI and support for legacy pin interrupts

High Performance RDMA
- Low latency and line rate bandwidth
- Enhanced RDMA primitives
- iWARP support on Linux OFED
- Microsoft Network Direct support

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

Virtualization
- PCI-SIG SR-IOV
- 128 Virtual Functions (VF), 8 PF
- 144 port virtual switch
- EVB, VEPA, Flex10, VNTag
- 336 virtual MAC addresses

FCoE
- Full HBA FCoE (Initiator or Target)
- Open-FCoE Offload (Initiator)
- CRC32 offload generation & verification
- Ingress & Egress ACL (Access Control List)

iSCSI Acceleration
- Full iSCSI initiator and target mode stack
- CRC32 offload generation & verification
- iSCSI proxy switching based on SCSI CDB
- Full HBA offload

TCP/IP Full Offload
- Full TCP implementation including IPv4 & IPv6
- Microsoft Chimney support
- Extensive RFC compliance, fully featured stack
- Full TCP Proxy between any set of connections
- VLAN support up to 4096 different VLAN IDs
- Load balancing and Fail-over capabilities

Stateless Offloads
- TCP/UDP checksum offload for IPv4 & IPv6
- TSO, LRO, LSO, and GSO for IPv4 & IPv6
- VLAN filtering, insertion & extraction
- Line rate packet filtering and attack protection
- Fine granularity time stamping (down to 4ns)
- Ethernet Routing (packet header rewrite)
- Packet Tracing and Packet Sniffing

Ethernet
- IEEE 802.3ae (10 GbE)
- IEEE 802.3z (1 GbE) compliant
- IEEE 802.IP priority
- IEEE 802.IQ VLAN tagging
- IEEE 802.1Qbg EVB/VEPA
- IEEE 802.1Qbh VNTag
- IEEE 802.3x flow control
- IEEE 802.3ad load-balancing and failover
- Ether II and 802.3 encapsulated frames
- Multiple MAC addresses per interface
- Jumbo Frames up to 9.6Kbytes

Physical and Environmental
- Dimensions without bracket:
  4.92" x 6.26" or 12.49 cm x 15.90 cm
- Operating Temp: 0 to 55ºC or 32 to 131ºF
- Operating Humidity: 5 to 95%
- Typical power consumption: 12W