

PRODUCT BRIEF



Chelsio WireDirect™

Ultra Low Latency WireDirect™ for Queue Pair, UDP With Multicast and Packet Tracing Applications

WireDirect™ technology consists of three main software modules that enable direct wire access from user space to the Chelsio T4 network adapter with complete bypass of the kernel. This results in a 10Gb Ethernet solution with an ultra-low 3µs deterministic latency perfectly suited for high frequency trading and other delay-sensitive applications. The three **Chelsio WireDirect** (WD) modules are: **WD-QP** (Queue Pair), **WD-UDP with Multicast**, and **WD-TRACE**.

WireDirect™ Technology

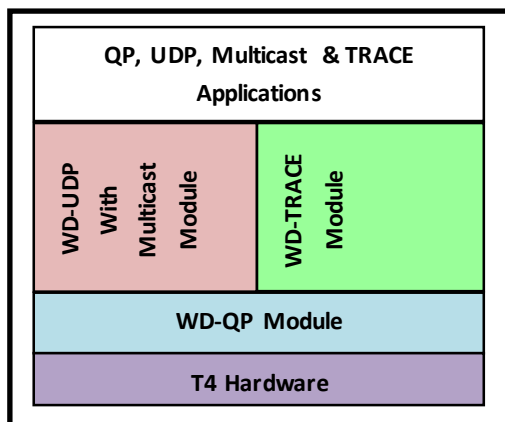


Figure 1: Chelsio WireDirect™ Architecture

At the core of Chelsio's WireDirect technology is the WireDirect Queue Pair (**WD-QP**) module (Figure 1). **WD-QP** provides user space applications direct access to Ethernet Packets on the wire via the Chelsio T4 ASIC without having to go through any of the I/O layers of the host computing stack. It provides 100 percent kernel bypass and zero copy with packet polling for data to and from the network wire. All the other WD modules are built on top of WD-QP to provide seamless I/O acceleration to applications.

WD-QP

The **WD-QP** API provides an asynchronous RDMA Verbs-like interface allowing zero-copy low latency Ethernet communication. It is intended to be used by middleware to implement a full user mode Ethernet interface, as well as by custom applications needing low latency/high packet rates. The API provides a QP interface for sending and receiving Ethernet Packets.

WD-UDP with Multicast

The **WD-UDP with Multicast** module provides an accelerated Sockets interface for sending and receiving both unicast and multicast UDP datagrams. No application changes are required to take advantage of **WD-UDP with Multicast**.

The WireDirect `libcxgb4_sock` library accelerates the following Socket calls: `write()`, `send()`, `sendto()`, `sendmsg()`, `read()`, `recv()`, `recvfrom()`, `recvmsg()`, `select()`, `poll()`, and `epoll()`. Applications that use these calls and are bound to the **WD-UDP** library via the `LD_PRELOAD` facility will exercise the **WD-UDP** logic to provide seamless I/O acceleration. WD-UDP also supports end-to-end traffic prioritization and application I/O QoS through the NIC and switch.

WD-TRACE

WD-TRACE allows bi-directional (egress + ingress) direct on-the-wire packet tracing, packet capture and packet filtering at line rates up to 20Gbps by bypassing the host OS kernel and I/O stack and going directly to user space that packet capture applications like `tcpdump`, `wireshark`, and `ethereal` can use. Additionally, WD-TRACE includes hardware packet timestamp capabilities with an ultra-precise 3.6ns (nanosecond) resolution. Additionally, WD-TRACE allows concurrent low-latency WD-UDP operation, i.e. when required the trace traffic travels on a separate path through the NIC, and goes to a different DMA queue that's handled by a dedicated core (thereby not disrupting the regular ingress traffic).

WireDirect™ Performance

