

## XGE 1Gb Ethernet Software Support Overview

### Abstract

Critical I/O's XGE 1Gb Ethernet products are designed specifically for data intensive real-time applications. The XGE hardware implements full TCP/UDP/IP/iSCSI/RDMA offload in silicon, providing the highest performance, low host CPU loading, and highly deterministic operation.

This whitepaper discusses the support software offered by Critical I/O for the XGE Ethernet product family. This includes XGE Drivers for Linux and VxWorks, and the operating system as well as the operating system and platform independent XGE Software Libraries.

## XGE Silicon Stack Software Support

Critical I/O provides extensive software support for its XGE Silicon Stack products that is specifically designed to provide optimal performance in embedded and real-time applications. The focus of Critical I/O's hardware and support software is to maximize performance (as measured by throughput, latency, and determinism), while minimizing the loading on host CPUs. A second goal is to provide simple and standardized application programming interfaces that allow application developers to easily leverage these capabilities. This paper describes the Critical I/O software library and driver features, and how application programmers can use these features.

Critical I/O's XGE software support is provided in the form of Libraries and Drivers:

- Critical I/O's XGE Drivers for VxWorks and Linux are operating system specific, and connect directly into standard operating system I/O interfaces providing such standard interfaces as Sockets and SCSI. Critical I/O Drivers are provided in object form.
- Critical I/O's XGE Libraries provide an easy to use API that can be easily accessed directly from application programs, without the need for an intervening driver. The libraries are designed to be operating system, processor, and endian independent. Critical I/O Software Libraries are normally provided in object code form, although they can also be provided in operating system independent source code form in certain situations.

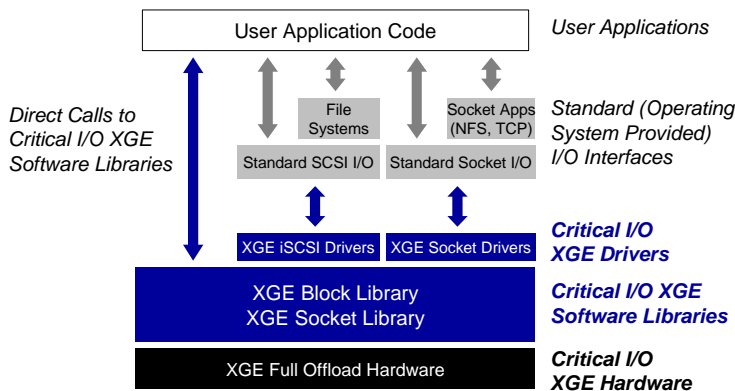


Figure 1: Diagram Illustrating Library vs. Driver Support

## XGE Gigabit Ethernet

Critical I/O's XGE product family provides a full TCP/UDP/Ethernet offload hardware engine. Full offload of the iSCSI protocol is also supported for storage, data collection, and instrumentation applications. Unlike partial-offload methods, the full-offload approach supports moving data at Gigabit wire speeds, while using just a few percent of the host CPUs processing power. Critical I/O's software support for the XGE product is designed specifically to allow application programmers to fully exploit the offload capability of the XGE interfaces, while providing the real-time performance, controllability, and characteristics that embedded real-time systems demand.

XGE software support consists of two multi-platform software libraries, and several platform specific (VxWorks/PPC, Linux/PPC, Linux/x86) sockets drivers. Figure 1 illustrates how the Critical I/O support software elements are layered between user application code and the XGE hardware, and also illustrates how the XGE drivers tie into standard operating system provided I/O interfaces.

## XGE Sockets Drivers for VxWorks & Linux

A socket is a communication end-point that established between two Ethernet devices, and is associated with a specific pair of TCP or UDP ports. Socket connections are established between two nodes over a network, with the IP address/port pair for each node defining the socket connection. Sockets may be created as stream sockets, using the TCP protocol, or as datagram sockets, using the UDP protocol.

The Critical I/O VxWorks and Linux Sockets Drivers allow socket based applications easy access to the high performance data transfer capabilities of the XGE interface. The sockets driver allows applications to make standard sockets I/O calls to create sockets, and send and receive data, while offloading all of the TCP/UDP/IP processing to the XGE hardware. This allows standard sockets based applications (such as NFS, FTP, etc.) to take advantage of XGE offload capabilities, as well as supplying the standardized, well understood socket network programming model for new and existing user developed applications.

The XGE sockets drivers provide an industry standard sockets programming interface but with slightly more host CPU overhead and less control and visibility, as compared to direct use of the Sockets Library.

	<i>VxWorks/PPC</i>	<i>Linux/PPC</i>	<i>Linux/x86</i>	<i>Other RTOS</i>
<i>Socket Driver</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>-</i>
<i>Socket Library</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
<i>iSCSI Extensions</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>

*Table 1: XGE Driver and Library - Platform Support and Availability*

## **XGE Socket Drivers – Standard Socket API**

The XGE Socket Drivers leverage the standard operating system provided socket API. All standard socket functions are supported. The drivers also provide XGE unique socket options which allow users to specify zero copy send and receive operations. This allows the XGE hardware to DMA data directly to and from user application buffers for highest possible performance. The VxWorks zbuf socket API is also supported.

Some examples of the standard socket API functions include:

- socket – create as socket
- connect – form a TCP connection
- bind – associate a socket with IP addresses and ports
- accept – accept an incoming connection
- send, sendto, sendmsg, write – send data
- recv, recvfrom, recvmsg, read – receive data

## **Critical I/O XGE Software Libraries**

The Sockets and Block Libraries share a simple to use API that provides a common set of TCP, UDP, iSCSI, and Ethernet functionality that can be easily and directly accessed by user applications. The libraries are designed to be operating system, processor, and endian independent.

*Direct calls to the XGE Software Libraries provide the highest performance, lowest host CPU loading, and the highest degree of interface control and visibility.*

## **Critical I/O Sockets Software Library**

The CIO Sockets Library allows the user application to directly call CIO Sockets Library functions to (optionally, for TCP) set up connection(s) with the remote socket(s) then call additional Sockets Library routines to send and/or receive data to/from the remote node socket(s) using TCP, UDP, IP, or Ethernet protocols. The libraries also supports the ARP and ICMP protocols. The Sockets Library supports a novel feature called Named Buffer Receives, which allows the XGE hardware to place incoming TCP or UDP data directly into application buffers which can be specified on a per-connection (per port for UDP) basis.

Direct use of the Sockets Library provides the highest performance sockets compatible interface, sending and receiving data at 250 MByte/s per port using less than 1% of the host CPUs throughput, with *application to application* latencies of as low as 25 usec. The library is fully compatible with any remote nodes that use the standard TCP or UDP sockets protocol. Both polled and interrupt driven modes of operation are supported.

*The sockets library is fully compatible with remote nodes that use sockets communication,, and provides the highest possible sockets compatible interface performance.*

## XGE Software Libraries – A Common Application Programming Interface

The XGE Software Libraries utilize a simple, easy to use, common Application Programming Interface (API). The key API functions that are common to all Critical I/O Libraries include:

- XE\_SelfTest – Performs a full self test of the specified XGE interface. Detailed test results are provided.
- XE\_Initialization – Initializes the specified interface, including the download of the interface firmware.
- XE\_Transfer\_Buffer – Initiates a transfer of data. This is a very versatile function, which may be called to perform a variety of different types data transfers, including TCP, UDP, IP, ARP, Raw Ethernet, etc.
- XE\_Status – Collects and reports the status on an interface. This can be called as part of an on-line health monitor, to obtain current operating parameters, or just to determine link status (up or down).
- XE\_Service – This provides interface service functionality, and must be called regularly from the user's application. When called, this routine processes any outstanding completions or errors associated with initiator or target data transfers. In the polled application model, XE\_Service is called periodically from the application context, or from a periodic polling task context. In the interrupt application model, it is typically called from a user interrupt service routine (ISR) which is connected to the XGE interrupt.
- XE\_Close – Closes out and shuts down the specified interface.

The XGE Software Libraries allow the user to define a set of notification or callback functions, which are called from the XE\_Service context when certain XGE events occur. The key common user-defined callback functions are:

- userStatusNotificationHandler – A user defined function that is called when the interface status changes, for example, an link-up or link-down transition.
- userSendCompletionHandler – A user defined function that is called when a send to a remote node completes. The user may define a unique send notification handler for each send, if desired

- userReceiveHandler – A user defined function that is called when a transfer from a remote node is completed.

## Critical I/O iSCSI Extensions Library

The XGE iSCSI extensions to the XGE Library provide a simple, efficient, and versatile interface for exchanging data between iSCSI devices. The iSCSI extensions provide access to the silicon based iSCSI offload capabilities of the XGE hardware. This hardware implementation of the iSCSI protocol provides full iSCSI initiator and target support. Multiple concurrent Peer-Peer communications are supported as well as direct communication with iSCSI Devices such as disk arrays and RAID's.