Dolphin
INTERCONNECT SOLUTIONS

PCI Express High Speed Networking

A complete solution for demanding Network Applications

Introduction

Maximizing application performance is a combination of processing, communication, and software. PCI Express Networks combine all these elements. A type of Local Area Network (LAN), PCI Express Networks are used to connect Processors, I/O devices, FPGAs and GPUs into an intelligent network. They can be used with flexible cabling or fixed backplanes. The main goal of PCI Express® networks is to eliminate system bottlenecks related to communication, allowing applications to reach their potential. To accomplish this, they deliver the lowest latency possible and combining that low latency with high data rates.

Dolphin's PCI Express® networking solution consists of standardized computer hardware and software that obtains the most out of applications. Our standard form factor boards reduce time to market, enabling customers to rapidly develop PCI Express® networking solutions for data centers and embedded systems.

Dolphin's software ensures reuse of existing applications but with better response times and data accessibility. Developers can quickly deploy applications with improved overall performance with our socket and IPoPCIe supported development tools. Further tuning is available with our low level APIs that delivers maximum performance.

Dolphin uses standard PCI Express® components as a road map to high performance hardware, currently deploying products at 64 Gbps. PCI Express®, as a standard, will continue to deliver high performance, while providing access to a complete low cost infrastructure. Dolphin's software takes advantage of this infrastructure enabling customers to deliver next generation systems and maximize performance of existing applications. Dolphin's PCI Express solutions is easy to implement and deploy. Customer will not need to change their existing applications to take advantage of PCI Express performance.

Table of Contents

Why use PCI Express® ?

Performance

PCI Express® solutions deliver outstanding performance compared to other interconnects in latency and throughput. When compared to standard 10 Gigabit Ethernet, PCI Express® latency is 1/10 the measured latency. This lower latency is achieved without special tuning or complex optimization schemes.

In addition, Dolphin takes advantage of the high throughput of PCI Express®. Our current solution utilizes Gen 2 and Gen 3 PCI Express® with 40 /64 Gbps of raw bandwidth. Future products will support easy upgrade to the next PCI Express® generation, doubling bandwidth. The upgrade to future PCI Express® generations requires no software changes and maintains the low latency device characteristic of PCI Express®. The investment in low latency high performance Dolphin products will yield dividends today and into the future.

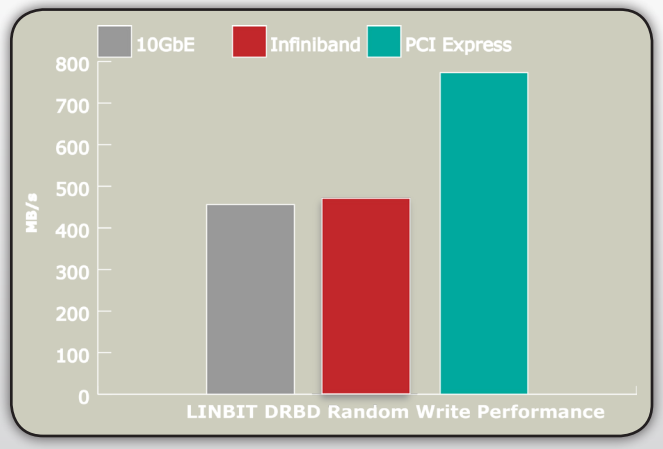

Eliminate Software Bottlenecks

Dolphin's PCI Express Software is aimed at performance critical applications. Advanced performance improving tools , such as the SuperSockets™ API, remove traditional network bottlenecks. Sockets, IP, and custom applications take advantage of the low latency PIO and DMA operations within PCI Express®, improving performance and reducing system overhead.

Applications utilizing SuperSockets™ expect latencies under 2 µs and throughput at 22 Gigabit/s. Other software components include an optimized TCP/IP driver for IP applications and the SISCI shared memory API. The SISCI API offers further optimization by using remote memory and replicated/ reflective memory. Customers benefit from even lower latencies in the range of 0.74µs latency with higher throughput of over 3500 MB/s.

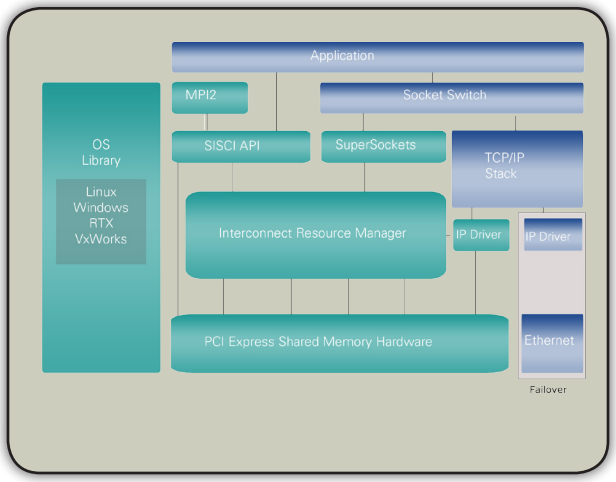

Key Applications

- Financial Trading Applications
- High Availability Systems
- Real Time Simulators
- Databases and Clustered Databases
- Network File Systems
- High Speed Storage
- Video Information Distribution
- Virtual reality systems
- Range and telemetry systems
- Medical Equipment
- Distributed sensor-to-processor systems
- High speed video systems
- Distributed shared memory systems

Robust Features

- Lowest host to host latency and low jitter with 0.74µs for fast connections and data movement
- DMA capabilities to move large amounts of data between nodes with low system overhead and low latency. Application to application transfers exceeding 3500 MB/s throughput.
- Management software to enable and disable connection and fail over to other connections
- Direct access to local and remote memory, hardware based uni- and multi-cast capabilities
- Set up and manage PCIe® peer to peer device transfers
- High speed sockets and TCP/IP application support
- Provides ease of installation and plug and play migration using standard network interfaces

High Performance Hardware

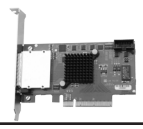	Low profile PCI Express® adapters card provide high data rates over a standard PCIe external cabling system. These cards feature transparent and non-transparent bridging (NTB) operation, along with clock isolation. Compatible with standard PCs and servers, they are used as a standard interface for servers and PCs in high performance low latency applications.
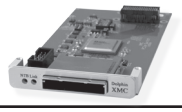	XMC Adapters bring PCIe data rates and advanced connection features to embedded computers supporting standard XMC slots, VPX, VME or cPCI carrier boards. PCI Express® adapters expand the capabilities of embedded systems by enabling very low latency, high throughput cabled expansion. Standard PCs can easily be connected to embedded systems.
	PCI Express switch boxes are used for scaling out PCI Express® Networks. By linking both transparent and non-transparent devices to a PCIe switch, distributed systems increase both I/O and processing capacity. These low latency switches enables systems to scale, while maintaining high throughput.

PCI Express® Software

Reflective Memory / Multi-cast

One example of a PCI Express® application is Dolphin's Reflective Memory or Multi-cast solution. Reflective Memory solutions have been available in the market for years. With PCI Express, Dolphin implements a reflective memory architecture in a modern switched architecture.

PCI Express® multicast enables a single bus write transaction to be sent to multiple remote targets or in PCI Express technical terms - multicast capability enables a single TLP to be forwarded to multiple destinations. Unlike other reflective memory concepts, Dolphin implements a switch based solution. The implementation of a PCI Express switched architecture results in lower latency and higher bandwidth. Dolphin benchmarks show end-to-end latencies as low as 0.99µs and over 2,650 Megabytes /sec dataflow at the application level. By using PCI Express based reflective memory functionality, customers can easily solve their real time, distributed computing performance requirements.

Dolphin combines PCI Express multicast with our SISCI (Software Infrastructure for Shared-memory Cluster Interconnect) API. The combination allows customers to easily implement applications that directly access and utilize PCI Express' reflective memory functionality. Applications can be built without the need to write device drivers or spend time studying PCI Express chipset specifications.

One major difference in Dolphin solution is the use of cacheable main system memory to store data. The use of cacheable main memory provides a significant performance and cost benefit. Remote interrupts or polling is used to signal the arrival of data from a remote node. Since the memory segments are normal cacheable main memory, polling is very fast and consumes no memory bandwidth. The CPU polls for changes in its local cache. When new data arrives from the remote node, the I/O system automatically invalidates the cache and the new value is cached.

FPGAs and GPUs can also take advantage of this reflective memory mechanism. Customers can use the SISCI API to configure and enable GPUs, FPGAs etc (any PCIe master device) to send data directly to reflective memory. Thus, avoiding the need to first store the data in local memory. Data can be written directly from an FPGA to multiple end points for processing or data movement. FPGAs can also be receive data from multiple end points.

Reflective memory solutions are known for their simplicity, just read and write into a shared distributed memory. Our high-performance network includes easy installation and operation. The Dolphin's SISCI Developers Kit manages the system. The kit includes all the tools and flexibility to setup your reflected memory system. Once setup, your application simply reads and writes to remote memory.

Features

- High-performance, ultra low-latency switched 40-Gbps data rate interconnect
- Up to 2650 MB/s data throughput
- FPGA support
- Hardware based multicast
- Configurable shared memory regions
- Fiber-Optic and copper cabling support
- Scalable switched architecture
- SISCI developers kit
- Built in CRC, 8b/10b encoding
- PCI Express®host adapters
- Expandable switch solutions

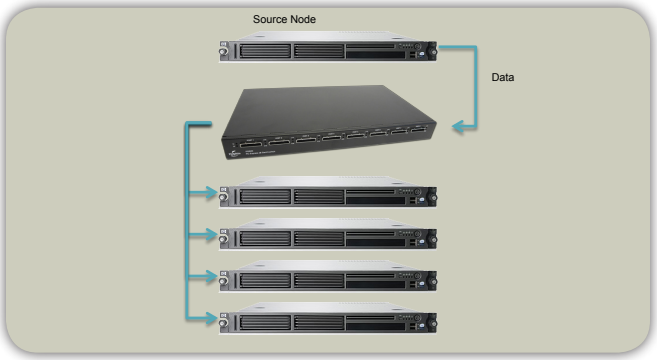

figure 2: Reflective memory diagram

Why use Dolphin Reflective Memory?

The Dolphin solution creates a reflective memory address space between nodes. The solution offers significantly higher performance at a much lower cost than other reflective memory solutions and is available with both low cost copper and long distance fiber cabling.

The Dolphin reflective memory functionality is implemented in the Dolphin Express switch. Topology sizes range from 2 node systems without a switch to multi-switch large system configurations. For larger systems, switches are cascaded. All hosts needs to install a Dolphin Express adapter card.

The solution offers significantly faster access to local data than other reflective memory solutions. One key point is that Dolphin uses system main memory as reflective memory . Other solutions use expensive device memory, limiting system configuration options and increasing cost. Main memory solutions benefit from CPU caching and very high performance internal memory buses. Traditional reflective memory device memory is non-cacheable and memory access is very expensive as the CPU must fetch the data from the card through the I/O system.

The Dolphin IXH adapters come with a x8 PCI Express® link enabling customer applications to take advantage of the exceptional 40 Gb/s link bandwidth.

The fully hardware based memory mapped data transmission does not rely on any operating system service or kernel driver functionality. The solution includes an extensive software library that makes configuration and setup easy. This software is inactive during application runtime.

Memory Addressing

Reflective memory can be addressed in 3 different ways.

The CPU can do direct PIO/Store operations to a reflective memory address using a pointer or memcopy() function. The access is write posted and will typically be completed in one CPU instruction.

The PCI Express DMA controller can be used to move data to a reflected memory address. A locally attached PCI Express device, FPGA, GPU etc can directly send data to a reflective memory address. Data does not need to go to local memory first.

Reflective Memory Performance

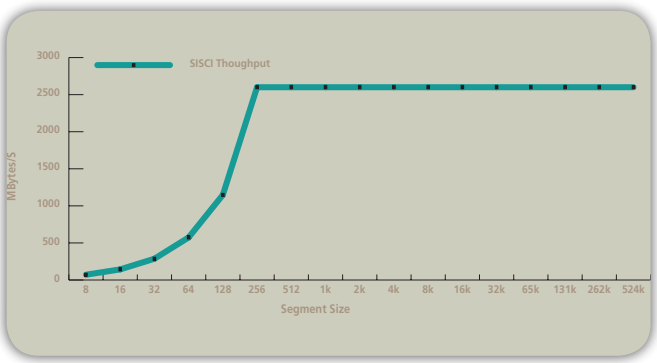

figure 1: Reflective memory throughput

The Dolphin replicated shared memory system delivers ultra fast performance. The SISCI low level API enables the lowest latency. Dolphin benchmarks show latencies as low as 0.99µs. This is accomplished while delivering high throughput. The reflective memory solution delivers up to 2650 MB/s of data throughput. The chart on Figure 2 show the throughput at various message sizes.

PCI Express® Software
SISCI Developers Kit

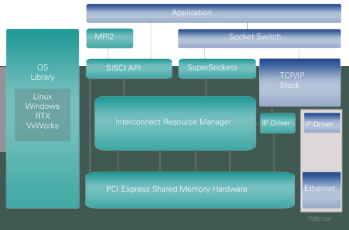

System architects seeking to maximize distributed application performance are exploring PCI Express® as a cable interconnect. Dolphin's Software Infrastructure Shared-Memory Cluster Interconnect (SISCI) API makes developing PCI Express® Network applications faster and easier. The SISCI API is a well established API for shared memory environments. In a multiprocessing architecture with PCI Express®, the SISCI API enables PCI Express® based applications to use distributed resources like CPUs, I/O, and memory. The resulting application features reduced system latency and increased data throughput.

For inter-process communication, PCI Express® supports both CPU driven programmed IO (PIO) and Direct Memory Access (DMA) as transports through Non-transparent bridging (NTB). Dolphin's SISCI API utilizes these components in creating a development and runtime environment for system architects seeking maximum performance. The environment is very deterministic, low latency and low jitter. Ideal for traditional high performance applications like real time simulators, reflective memory applications, high availability servers with fast fail-over, and high speed trading applications.

The SISCI API supports data transfers between applications and processes running in an SMP environment as well as between independent servers.

It's capabilities include managing and triggering of application specific local and remote interrupts and catching and managing events generated by the underlying PCI Express® system (such as a cable being unplugged). The SISCI API makes extensive use of the "resource" concept. Resources can be items such as virtual devices, memory segments, and DMA queues.

The API removes the need to understand and manage low level PCI Express® chip registers at the application level, easily enabling developers to utilize these resources in their applications without sacrificing performance. Programming features include allocating memory segments, mapping local and remote memory segments into the addressable space of their program, and manage and transfer data with DMA. The SISCI API includes advanced features to improve overall system performance and availability. Caching techniques can be exploited to improve performance and the API can be used to check for data transfer errors and correct them.

Features

- Shared memory API
- PCI Express Peer to Peer support
- Replicated/reflective memory support
- Distributed shared memory and DMA support
- Low latency messaging API
- Interrupt management
- Direct memory reads and writes
- Windows, RTX, VxWorks, and Linux OS support
- Caching and error checking support
- Events and callbacks
- Example code available

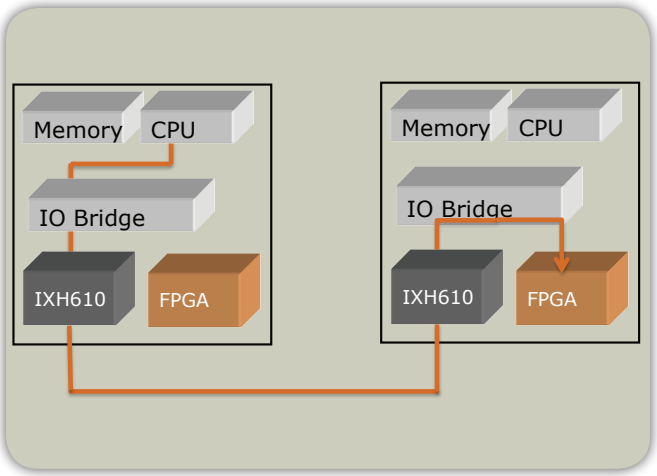

figure 3: Device to device transfers

Why use SISCI?

The SISCI software and underlying drivers simplifies the process of building shared memory based applications. The built in resource management enables multiple concurrent SISCI programs and other users of the PCI Express® network to coexist and operate independent of each other. The SISCI API is available in user space. A similar API is available in kernel space.

For PCI Express® based application development, the API utilizes PCI Express® Non-transparent bridging to maximum application performance. The shared memory API drivers allocate memory segments on the local node and make this memory available to other nodes. The local node then connects to memory segments on remote nodes.

Once available, a memory segment is accessed in two ways, either mapped into the address space of your process and accessed as a normal memory access, e.g. via pointer operations, or use the DMA engine in the PCI Express® chipset to transfer data. Figure 4 illustrates both data transfer options.

Mapping the remote address space and using PIO may be appropriate for control messages and data transfers up to e.g. 1k bytes, since the processor moves the data with very low latency. PIO optimizes small write transfers by requiring no memory lock down, data may already exist in the CPU cache, and the actual transfer is just a single CPU instruction – a write posted store instruction. A DMA implementations saves CPU cycles for larger transfers, enabling overlapped data transfers and computations. DMA has a higher setup cost so latencies usually increase slightly because of the time required to lock down memory and setup the DMA engine and interrupt completion time. However, more data transfers joined and sent together to the PCI Express® Switch in order amortizes the overhead.

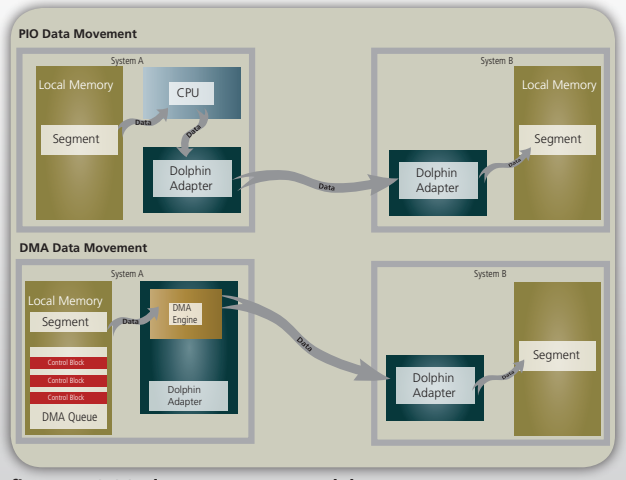

figure 4: SISCI data movement model

SISCI Performance

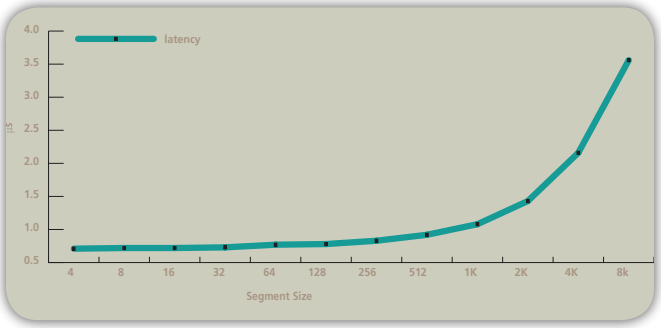

figure 5: SISCI latency

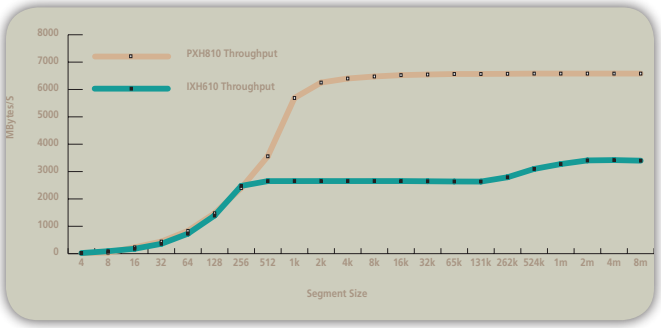

figure 6: SISCI PIO/DMA Throughput

The SISCI API provides applications direct access to the low latency messaging enabled by PCI Express®. Dolphin SISCI benchmarks show latencies as low as 0.74µs. The chart on Figure 5 show the latency at various message sizes.

The SISCI API enables high throughput applications. This high performance API takes advantage of the PCI Express® hardware performance to deliver over 6500 MB/s for Gen 3 and 3500 MB/s for Gen 2 of real application data throughput. Figure 6 shows the throughput at various message sizes using Dolphin IXH and PXH host adapters.

PCI Express® Software

IP over PCI Express

Dolphins performance optimized TCP IP driver for PCIe (IPoPCIe) provides a fast and transparent way for any networked applications to dramatically improve network throughput.

The software is highly optimized to reduce system load (e.g. system interrupts) and uses both PIO and RDMA operations to implement most efficient transfers for all message sizes.

The major benefits are plug and play and much high bandwidth and lower latency than network technologies like 10G Ethernet.

At the hardware level, the TCP/IP driver provides a very low latency connection, but the operating system networking protocols is typically introducing a significant delay for safe networking (required for non reliable networks like Ethernet). User space applications seeking lowest possible network latency should utilize the Dolphin SuperSockets technology for lowest possible latency, the IPoPCIe driver will typically provide 2-3 times better latency than 10G Ethernet.

The optimized TCP/IP driver is recommended for applications like

Windows:

- Microsoft Hyper-V live migration
- Network file sharing (map network drive)
- Applications that requires UDP which is not support by SuperSockets™ yet.

Linux:

- General networking
- NFS
- Cluster file systems not supported by SuperSockets
- iSCSI

Features

- All networked, users space and kernel space applications are supported
- 100% compliant with Linux Socket library, Berkeley Socket API and Windows WinSock2
- No OS patches or application modifications required. Just install and run
- Routing between networks
- ARP support
- Both TCP and UDP supported. (UDP multicast/broadcast is not supported yet using Linux, but SuperSockets™ for Linux supports UDP multicast)
- Supports hot-pluggable links for high availability operation
- Easy to install

IPoPCIe Performance

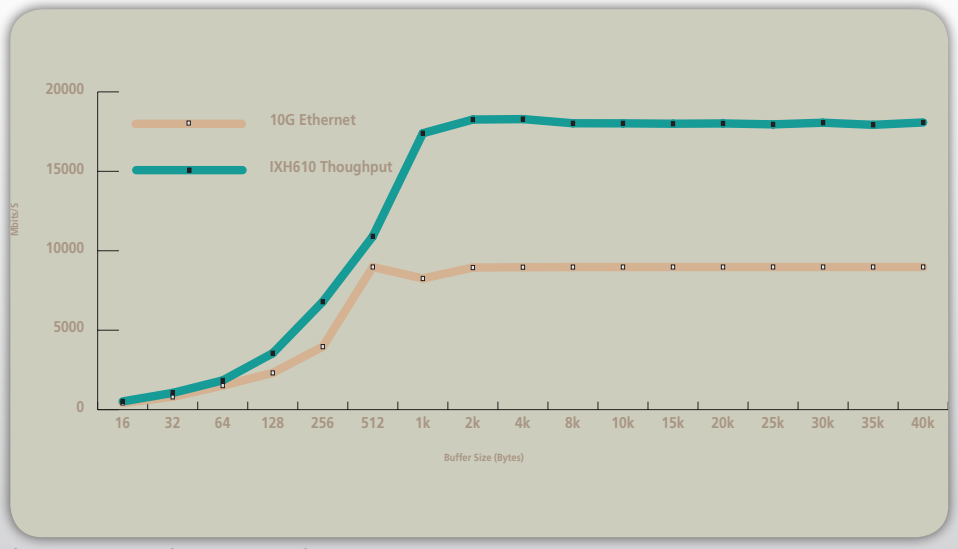

figure 7: Sockperf Benchmark of throughput with IXH610 and Intel x540

PCI Express® Software

SuperSockets™

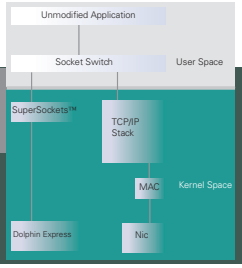

PCI Express® can replace local Ethernet networks with a high speed low latency network. The combination of Dolphin's PCI Express® hardware and Dolphin's SuperSockets™ delivers maximum application performance without necessitating application changes. SuperSockets™is a unique implementation of the Berkeley Sockets API that capitalizes on the PCI Express® transport to transparently achieve performance gains for existing socket-based network applications. Dolphin PCI Express® hardware and the SuperSockets™software layer create an ultra-low latency, high-bandwidth, low overhead, and high availability platform to support the most demanding sockets based applications. With support for both Linux and Windows Operating systems, new and existing applications can easily be deployed on the high performance Dolphin platform without any modifications.

Traditional implementations of TCP socket require two major CPU consuming tasks: data copy between application buffers and NIC buffers along with TCP transport handling (segmentation, reassembly, check summing, timers, acknowledgments, etc). These operations turn into performance bottlenecks as I/O interconnect speeds increase. SuperSockets™eliminates the protocol stack bottlenecks, delivering superior latency performance. Our ultra low latency is achieved through the use of an efficient remote memory access mechanism. This mechanism is based on a combination of PIO (Programmed IO) for short transfers and RDMA (Remote Direct Memory Access) for longer transfers, allowing both control and data messages to experience performance improvements.

SuperSockets™is unique in its support for PIO. PIO has clear advantages for short messages, such as control messages for simulations systems, as the transfer is completed through a single CPU store operation that moves data from CPU registers into the memory of the remote node. In most cases, data transfers through SuperSockets™ are completed before alternative technologies have even managed to start their RDMA.

In addition to support for PIO, SuperSockets™ implements a high speed loopback device for accelerating local system sockets communication. This reduces local sockets latency to a minimum. For SMP systems, loopback performance is increased 10 times.

SuperSockets™comes with built in high availability, providing instantaneous switching for hardware node or network errors. If the Dolphin Express network is unavailable, socket communication will be served by the regular network stack. This automatic process is the most cost-effective way for building high availability systems. The Linux version comes with an instant fail-over and fail-forward mechanism that transparently will switch between Dolphin Express and regular networking. In conjunction, Dolphin's PCI Express® network and SuperSockets™provide an ultra fast and reliable transport mechanism for embedded and business applications.

Features

- Compliant with Linux Socket library, WinSock2, and Berkeley Sockets
- Windows and Linux OS support
- Both TCP and UDP support
- UDP multicast support
- Supports both user space and kernel space clients
- Full support for socket inheritance/duplication
- Transparent fail-over to Ethernet if high speed connection is down. Fail forward when problem is corrected
- Includes local loopback socket acceleration up to 10 times faster than standard Linux and Windows
- Supports multiple adapters per host for increased fault tolerance and speed
- No OS patches or application modifications required
- Easy to install with no application modifications

How Does SuperSockets™ Work?

In order to divert socket communication, without touching the application, the sockets API functions must be intercepted. This is done differently in the Windows and Linux environment.

Dolphin SuperSockets™ on Linux differs from regular sockets only in the address family. SuperSockets™ implement an AF_INET compliant socket transport called AF_SSOCK. The Linux LD_PRELOAD functionality is used to preload the standard socket library with a special SuperSockets™ library that intercepts the socket() call and replaces the AF_INET address family with AF_SSOCK. All other sockets calls follow the usual code path. Target address within the Dolphin Express Network are accelerated by the SuperSockets™ module.

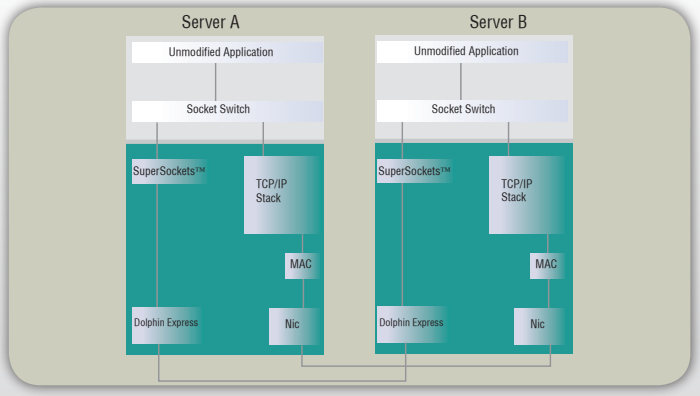

figure 8: SuperSockets™ vs. Ethernet Data Movement Model

For Windows applications or services, a Layered Service Provider module is installed and automatically configured. The LSP accelerates socket transfers initiated by AF_INET or AF_INET6, SOCK_STREAM endpoints. The SuperSockets stack provides a proxy application called dis_ssocks_run.exe that enables specific programs to use the PCI Express path. By default, the LSP is a pass-through module for all applications: the network traffic passes through the NDIS stack.

The network acceleration over PCI Express occurs when the interconnect topology is fully functional, the client and server programs are launched under the proxy application's control and both sides use the standard Winsock2 API calls. At runtime, a native socket is created and used for initial connection establishment. Therefore all connections are subject to typical network administrative policies.

The supported transfer modes are blocking, non-blocking, overlapped, asynchronous window and network events. The Service Provider balances the CPU consumption based on the traffic pattern. Dedicated operating system performance counters are additionally provided.

SuperSockets™ Performance

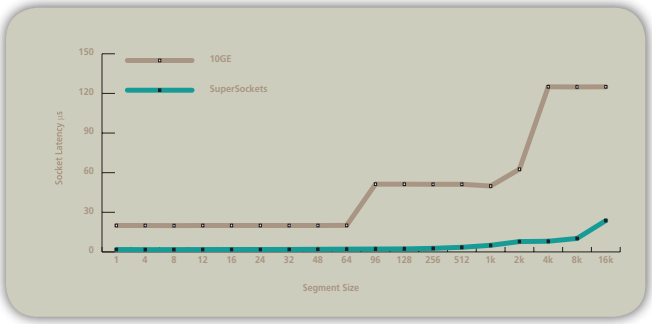

figure 9: SuperSockets latency

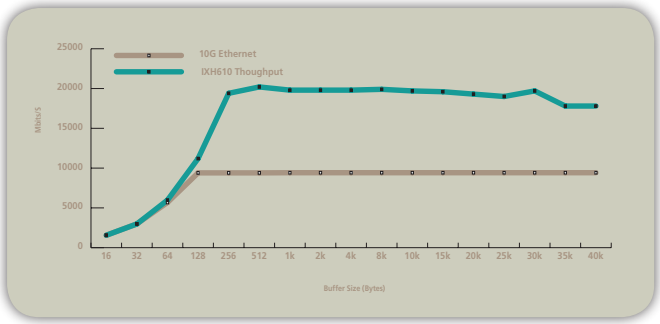

figure 10: SuperSockets™ data throughput

SuperSockets™ is optimized for high throughput, low latency communication by reducing system resource and interrupt usage in data transfers. The latency chart above shows performance results for two Intel Core2 2.00 GHz systems interconnected using PCI Express® vs 10 Gigabit Ethernet. The socket ping-pong test shows the half RTT (Round Trip Time). The minimum latency for Dolphin SuperSockets™ is under 1.5 microseconds. SuperSockets™ also delivers high throughput with over 2700 MB/s of data throughput.

PCI Express Hardware
IXH610 Host Adapter

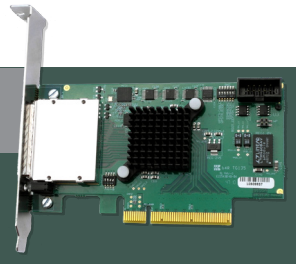

Dolphin Express provides an optimized PCI Express intercommunication platform for financial, industrial, medical, and military systems, utilizing standard, low cost PCI Express components. The IXH610 Gen2 PCI Express Host Adapter is our high performance cabled interface to external processor subsystems or I/O subsystems. Based on Gen2 PCI Express® bridging architecture, the IXH610 host adapter includes advanced features such as non-transparent bridging (NTB) and clock isolation.

For high performance application developers, the IXH610 host adapter combines 40 Gbit/s performance with less than one microsecond latency, significantly improving overall inter-system communication. Connecting remote I/O subsystems in transparent mode requires no special drivers, so deployment is fast and easy. Inter-processor communication also benefits from the high throughput and low latency.

The IXH610 performs both Direct Memory Access (DMA) and Programmed IO (PIO) transfers, effectively supporting both large and small data packets. DMA transfers result in efficient larger packet transfers and processor off-load. PIO transfers optimize small packet transfers at the lowest latency. The combination of DMA and PIO creates a highly potent data transfer system.

Dolphin's software suite takes advantage of this data transfer scheme. Delivering a complete deployment environment for customized and standardized applications. The Shared-Memory Cluster Interconnect (SISCI) API is a robust and powerful shared memory programming environment. The optimized TCP/IP driver and SuperSockets™ software remove traditional networking bottlenecks. IP and sockets applications take advantage of the high performance PCI Express interconnect without modification. The overall framework is designed to meet all the demands for rapid development of inter-processor communication systems.

With the implementation of clock isolation, the IXH610's signal quality is excellent. By isolating the system clock and transmitting an extremely low jitter high quality clock to downstream devices, the IXH610 offers users high signal quality and increased cable distances. Signal quality is essential for applications such as test and measurement equipment, medical equipment, and storage subsystem seeking high performance and data quality.

Features

- PCI Express® 2.1 compliant - 5.0 Gbps per lane
- x8 PCI Express port - 40 Gbit/s
- Link compliant with Gen1 and Gen2 PCI Express
- Support Gen1, Gen2, and Gen3 PCIe Slots
- RDMA support through PIO and DMA
- PCI Express® External Cabling Specification
- PCI Express x8 iPass® Connectors
- Copper and Fiber-optic cable connection up to 5 meters copper connections, up to 300 meters fiber optic
- Clock isolation support
- Transparent bridging to cabled I/O devices
- Non-transparent bridging to cabled PCI Express systems
- Low Profile PCI Express form factor
- EEPROM for custom system configuration
- Link and status LEDs through face plate

Inter-processor connections

When used for inter-host connections, the IXH610 adapter is capable of node to node connections or connections through a IXS600 Switch as shown in figure 11. Adding industrial systems is done by connecting to the IXH620 XMC adapter. Each connection supports 40 Gbps with latencies as low as 0.74 microseconds. Designed for x8 PCI Express Systems, the IXH610 supports any system with a standard x8 or x16 PCI Express slot.

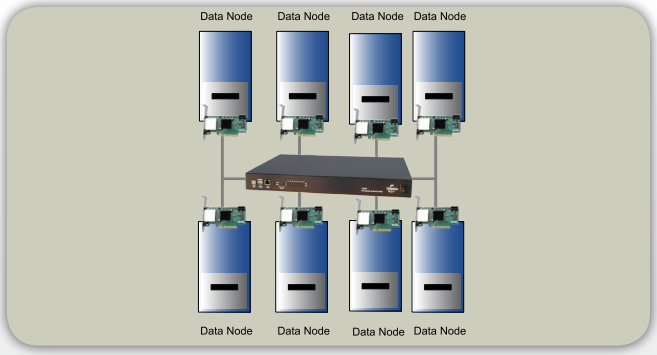

figure 11: Eight node PCI Express Cluster

Remote I/O Connections

The IXH610 functions as a high quality transparent connection to remote PCI Express I/O subsystems. These subsystems include test equipment, I/O expansion systems, specialized equipment, and storage systems. The IXH610 is specially designed for higher signal quality and support for spread spectrum clocking. The IXH611 is used as a target adapter in I/O expansion applications.

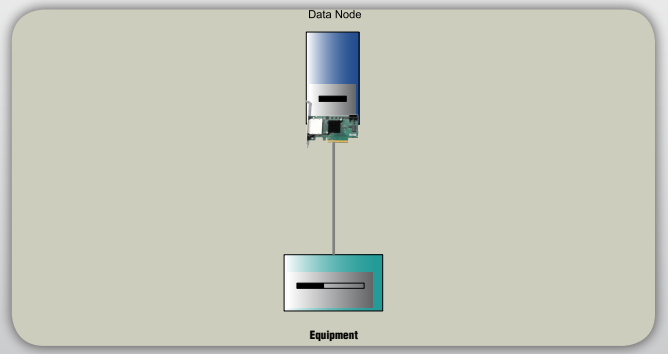

figure 12: I/O expansion with PCI Express

Specifications

Link Speeds	40 Gbit/s
Application Performance	0.74 microsecond latency (application to application)
PCI Express®	Base Specification 2.1
Topologies	Point to point, Switched
Cable Connections	One x8 standard PCI Express copper cable, fiber optic cable support Supports x8 to x4 transition cables
Power Consumption	7 watts
Mechanical Dimensions	PCI Express® Card Electromechanical Specification 2.0
Operating Environment	Operating Temperature: -10˚C -60˚C Relative Humidity: 5% -95% non-condensing
Dolphin Software	SuperSockets™ Berkeley Sockets API Microsoft WinSock2/LSP support SISCI API IPoPCIe

User Configuration Modes	Transparent/non-transparent(NTB)
Regulatory	CE Mark EN 55022,EN 55024-A1&A2, EN 61000-6-2 FCC Class A UL94V-0 compliant RoHS Compliant
Operating Systems	Windows Linux VxWorks RTX
Product Codes	IXH610 Host Adapter IXH611 Target Adapter

PCI Express Hardware

IXH620 XMC Host Adapter

As performance needs increase in embedded systems, high performance embedded system designers are implementing PCI Express for inter-system communication. The IXH620 XMC cable adapter enables high speed PCI Express inter-system connections to external systems including servers, single board computers, and I/O subsystems.

The IXH620's 40Gbit/s throughput and sub one microsecond latency deliver superior performance and low latency to systems supporting standard XMC slots or XMC VPX, VME or cPCI carrier boards. The IXH620's x8 cable connection uses an iPass™connector as a standardized connection method. The XMC adapter supports either upstream or downstream target configurations. To connect remote I/O or processing, the IXH620 implemented IDT®'s transparent or non-transparent bridging (NTB) functions. Used in transparent mode, standard PCI Express devices and drivers require no modifications. In NTB mode, the adapter facilitates inter-processor communication through Programmed IO (PIO) or Remote Direct Memory Access (RDMA).

The IXH620 implements a system clock isolation system for excellent signal quality. By isolating the system clock and transmitting an extremely low jitter high quality clock to downstream devices, IXH620 users benefit from improved signal quality, reliability, and cable distances. .

The IXH620 comes with Dolphin's comprehensive software suite that reduces time to market for customer applications. The Shared-Memory Cluster Interconnect (SISCI) API is a robust and powerful programming environment for easy development of shared memory applications. Shared memory applications benefit from the 0.74 microsecond inter-system latency and more than 3500 Megabytes/s throughput. The optimized TCP/IP driver and SuperSockets™software remove traditional networking bottlenecks. IP and Sockets applications can take advantage of the high performance PCI Express interconnect. Sockets applications experience 1.25 microsecond latency and 23 Gigabit/second user payload throughput.

These powerful features make XMC adapter an ideal interconnect for applications such as military and industrial systems that seek high performance and flexibility.

Features

- PCI Express® 2.1 compliant - 5.0 Gbps per lane
- x8 PCI Express port - 40 Gbit/s
- Link compliant with Gen1 PCI Express
- VITA 42.0-2005, ANSI/VITA 42.3-2006 compliant
- RDMA support through PIO and DMA
- PCI Express® External Cabling Specification
- PCI Express x8 iPass® Connectors
- Copper and fiber-optic cable connections up to 5 meters copper connections, up to 300 meters fiber optic
- Clock isolation support
- Transparent bridging to cabled I/O devices
- Non-transparent bridging to cabled PCI Express systems
- Short XMC form factor
- EEPROM for custom system configuration
- XMC P15 connector
- Link and status LEDs through front panel

Inter-processor communication

The IXH620 connects single board computers or systems running a distributed processing environment. Industrial or military customers requiring redundancy or increased compute resources use the IXH620 by itself or in conjunction with the IXS600 switch. Figure 13 illustrates connecting two single board computers with the IXH620. Fast data transfers are enabled through Dolphin's shared memory mechanism.

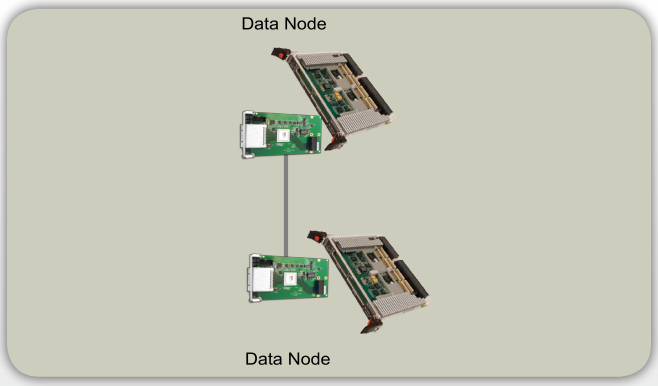

figure 13: two node connection with XMC adapter

Remote I/O subsystems

To extend the I/O capabilities of an XMC enabled system, the IXH620 supports host adapter or target adapter configurations. Industrial and Military customers requiring increased I/O bandwidth for graphics, processing, or data collection can add IXH620 enabled XMC carrier cards and chassis. Figure 14 illustrates connecting a standard server to and XMC enabled chassis to attach additional XMC cards for test system or increased functionality.

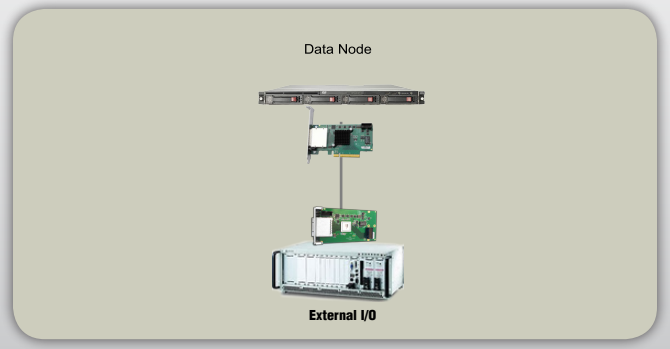

figure 14: connect embedded equipment to hosts

Specifications

Link Speeds	40 Gbit/s
Application Performance	0.74 microsecond latency (application to application)
Specifications	PCI Express Base Specification 2.1 VITA 42.0-2005, ANSI/VITA 42.3-2006
Topologies	Point to point, Switched
Cable Connections	One x8 Standard PCI Express copper cable, fiber optic cable support Supports x8 to x4 transition cables
Power Consumption	7 watts
Mechanical Dimensions	XMC Short form factor
Operating Environment	Operating Temperature: -10˚C -60˚C Relative Humidity: 5% -95% non-condensing

Dolphin Software	SuperSockets™ Berkeley Sockets API Microsoft WinSock2/LSP support SISCI API IPoPCIe
User Configuration Modes	Transparent/non-transparent(NTB)
Regulatory	CE Mark EN 55022,EN 55024-A1&A2, EN 61000-6-2 FCC Class A UL94V-0 compliant RoHS Compliant
Operating Systems	Windows Linux VxWorks RTX
Product Codes	IXH620 host /target adapter

PCI Express Hardware

IXS600 PCI Express® Switch

PCI Express provides low latency, highly efficient switching for high performance applications. The IXS600 Gen3 PCI Express switch delivers a powerful, flexible, Gen3 switching solution. This powerful switch enables I/O scaling and inter-processor communication by combining transparent and non-transparent bridging capabilities with Dolphin's software and clustering technology. IXS600 users can connect multiple PCI Express devices or create a highly efficient compute cluster with PCs, servers, or SBCs with XMC sites.

The IXS600 is the switching element of Dolphin's product line. This eight port, 1U cluster switch delivers 64 Gbps of non-blocking bandwidth per port at ultra low latencies. Each x8 PCI Express port delivers maximum bandwidth to each device while maintaining backwards compatibility with Gen1 and Gen2 components. As with other Dolphin products, the IXS600 utilizes standard iPass™ connectors to string components via copper or fiber-optic cabling. IXS600 customer can link multiple standardized PCI Express products such as PXI chassis, storage, and I/O expansion units.

For Non Transparent Bridging (NTB) or clustering applications, the IXS600 integrates with Dolphin's PCI Express Host Adapters or XMC Adapter. The total NTB solution solves many of the problems related to PCI Express operations, such as power sequencing and standard software. Dolphin customers avoid the severe power-on requirements normally associated with PCI Express. Hosts, cables, and the switch can be hot-swapped and power cycled in any sequence for real plug and play.

The IXS600 switch can also be partitioned into several virtually independent partitions, e.g. mixing NTB and Transparent functionality on separate ports.

The IXS600 switch supports Dolphin's comprehensive software suite. The Shared-Memory Cluster Interconnect (SISCI) API is a robust and powerful programming environment for easy development of shared memory applications. The optimized TCP/IP driver and SuperSockets™ software remove traditional networking bottlenecks. IP and Sockets applications can take advantage of the high performance PCI Express interconnect. This comprehensive solution is ideal for real-time, technical and high performance computing, cloud computing, and enterprise business applications.

Features

- PCI Express® 3.0 compliant -8.0 Gbps per lane
- Eight PCI Express Gen3 x8 ports
- Auto-training to lower lane widths
- Supports x4 lanes with a transition cable
- Link compliant with Gen1 and Gen2 PCI Express®
- Transparent and Non Transparent support
- PCI Express® External Cabling Specification
- PCI Express x8 iPass® Connectors
- Fiber-optic and copper cable support
- Hot Plug PCI Express cabling support in NTB mode
- Built in management processor
- Boot configuration and user data support for cluster and system configuration
- 19 Inch 1U rack mountable chassis
- Front and rear operational status and alert LEDs
- Redundant Fans

Inter-processor Switching

The IXS600 is a key element in Dolphin's inter-processor connectivity strategy. Industrial, military, or enterprise customer can create diverse multiprocessing configurations. The hybrid configuration illustrated shows single board computers and data nodes connected through the switch. The IXH610 host adapter and IXH620 XMC adapter are used to connect to different compute nodes.

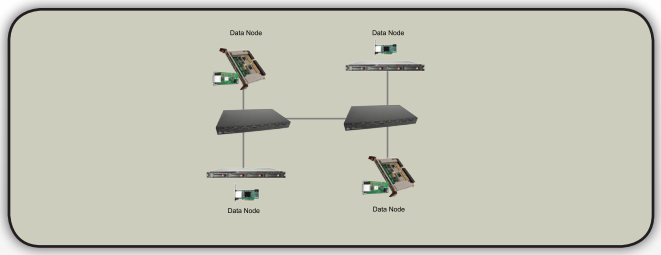

figure 15: hybrid configuration

I/O Expansion Interconnect

Increasing the number of I/O components in a system is accomplished by using the IXS600 with PCI Express I/O expansion boxes. Figure 16 illustrates the IXS600 connecting 7 additional I/O expansion boxes. These boxes can accommodate components such as sensors, graphics, coprocessors, video capture cards, and other I/O devices.

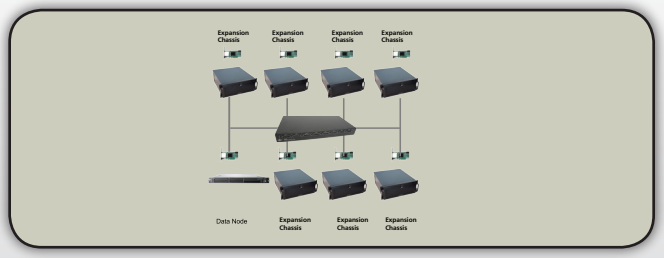

figure 16: I/O expansion

Scalability

Scalability is achieved by connecting multiple switches. Multiple IXS600 switches are used for larger reflective memory applications or creating larger processing clusters.

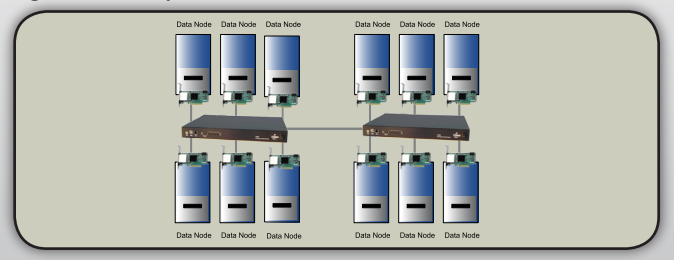

figure17: Scalability of switches

Specifications

Link Speeds	64 Gbits/s duplex
Application Performance	3Gbytes/s application data rate 150ns port latency
Ports	8 - x8 non-blocking 64 Gbps ports
Cable connections	PCI Express copper Cables -up to 5 meters Fiber-Optic cables- up to 300 meters
Management	RJ45 and USB management port Operational status and alert LEDs
Power	Auto-sensing power supply 110 - 240 V AC 50-60Hz Power consumption (including power to the optional fiber transceiver cables) : Max 80 Watts

Mechanical	1U, 19 inch rackmountable chassis 440mm (W) x 300mm (D) x 45mm(H) Redundant Fans
User Configuration Modes	Transparent/non-transparent(NTB)
Operating Environment	Operating Temperature: 0°C -55°C Relative Humidity: 5% -95% non-condensing
Regulatory	CE Mark EN 55022,EN 55024-A1&A2, EN 61000-6-2 FCC Class A UL94V-0 compliant RoHS
Product Codes	IXS600

Dolphin Interconnect Solutions AS is a wholly-owned subsidiary of Dolphin Group ASA which is listed at Oslo Stock Exchange (OSE ticker:DOLP). Dolphin Interconnect Solutions has been a global provider of ultra-low latency, high-bandwidth computer interconnect solutions for high speed real-time systems, clustered databases, general networking, web services and industrial applications for more than 20 years. For more information, please visit www.dolphinics.com.

Visit us on linked in and Facebook:

https://www.linkedin.com/company/dolphin-interconnect-solutions

https://www.facebook.com/dolphinInterconnect