



# PCI Express High Speed Products

COMPLETE SOLUTION FOR HIGH SPEED CONNECTIVITY


# Introduction

Dolphin engineers are experts in PCI Express advanced uses. Whether it is connecting Processors, I/O devices, FPGAs or GPUs in complex systems or simple I/O expansion applications. Dolphin has advanced solutions that solve complex problems and deliver real world systems. Our products include both software and hardware. eXpressWare software is an advanced middleware that runs on Dolphin hardware or can be ported to third party hardware. Our standard hardware products include host adapters and switches suitable for a number of transparent and non-transparent applications. Overall Dolphin provides high performance solution based on standard PCI Express.

# Contents

<b>Introduction.....</b>	<b>2</b>
<b>Applications.....</b>	<b>3</b>
PCIe Cluster .....	4
Composable Architecture .....	5
PCIe Multi-cast .....	6
NVMe over PCIe .....	7
PCIe Expansion.....	8
Performance/ Topologies .....	9
<b>eXpressWare .....</b>	<b>12</b>
SISCI API .....	13
Device Lending.....	14
PCIe Hot Add .....	15
SuperSockets.....	16
IPoPCIe.....	16
<b>Hardware .....</b>	<b>18</b>
eXpressWare Host Adapters .....	19
PCIe Host/Target Adapters .....	23
PCIe Switches .....	27
PCIe Backplanes and Cables.....	29
<b>Gen4 Products.....</b>	<b>31</b>

# Applications




PCI Express is used in a number of industries and is the most ubiquitous interconnect on the market. Dolphin leverages PCI Express to provide lower latency and higher performing solutions. By extending the capabilities of PCI Express, we have solved system problems that have plagued developers. Our solutions leverage the low latency and performance of PCI Express with a complete software suite that makes these capabilities accessible.

# PCIe Cluster

## High Performance Networking

eXpressWare Software enables the deployment of a very low latency PCIe fabric. This fabric can connect nodes ranging from just two for high speed replication to over 100 for networking applications.

eXpressWare software supports multiple protocols for clustering including TCP/IP, Sockets, and the SISI API for low latency shared memory. These various methods of communication can be used for processor to processor communication. The advantage of PCI Express is based on using PIO and DMA transfers that can be as low as 540ns and can be combined with throughput levels of 13GT/s. Our Gen3 products include support for copper and fiber cables. The fiber connections can be as long as 100 meters.



## Benefits

- Extremely low latency communication at  $<1\mu\text{s}$  host to host
- Performance scaling with PCI Express without changing application
- Support for standard APIs such as sockets and TCP/IP
- Low latency shared memory API for custom applications

## Features

Shared memory API for clusters and SMP's

Distributed shared memory

Easy deployment of DMA transfers

Direct remote CPU load and stores, memcopy()

Create and trigger remote application level interrupts

Caching and error checking support

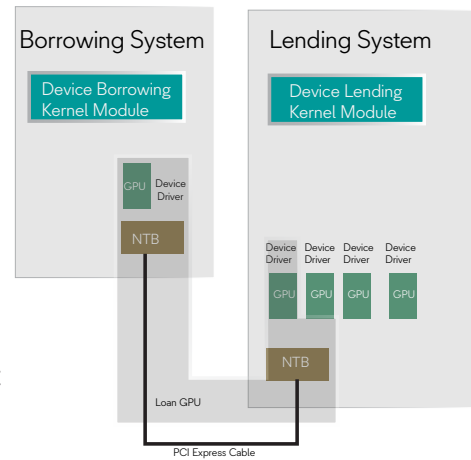
Events and callbacks

Windows, RTX, VxWorks and Linux OS support

# Composable Architecture

## Sharing Devices

eXpressWare software now includes SmartIO technology. This software enables the development of composable infrastructures. Devices within the PCIe Fabric can be moved and shared between systems. This creates a pool of resources that can be borrowed between systems in the fabric. For example, GPUs located in one system can be borrowed and used by another system as if they were locally installed. This enables building systems on the fly without the need to reboot the system. Most PCIe devices can be added and released as needed with the SmartIO software and some devices can be shared between systems.



## Benefits

- Multiple sharing methods including device lending and SISI smartIO
- Low overhead sharing without the need for translation layers
- MR-IOV capable with the use of SR-IOV devices
- Near native performance between remote devices and local device

## Features


Sharing of SR-IOV devices	Borrow remote functions from SR-IOV devices
Borrow remote devices located in remote systems	Share non-SR-IOV devices with SISI API
No translation layer overhead	No power-on sequencing when borrowing or lending devices
Linux only support	Support most modern Linux distributions

# PCIe Multi-cast

## Reflective Memory

### Solution

Our reflective memory or multi-cast solution reinterprets traditional reflective memory offerings by using a modern high speed switched architecture. A PCIe switched architecture employs multi-cast and switches vs the traditional ring approach. A single bus write transaction is sent to multiple remote targets through the switched network. PCI Express multi-cast results in a lower latency and higher bandwidth reflective memory solution.



### Benefits

- Minimal delay enables real-time applications to benefit from a significantly reduced total communication time
- Uses main memory or device memory. Reading data in main memory is significantly faster than solutions storing data in specialized PCIe device memory located in the computer IO system
- Multi-casting of data to all switch ports simultaneously
- Extremely low latency PCIe Fabric - Less than 1 ms of application-to-application latency

### Features

Data in main memory or device memory	Fair arbitration and sharing of bandwidth
Data is multicast by a centralized switch fabric	Eliminates significant delays between the first and the last node in the network receiving the data
Each PCI Express switch will send data out on all connected ports simultaneously	Hardware based CRC and re-transmission

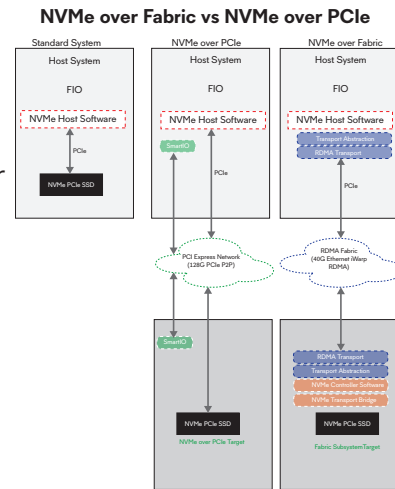
# NVMe over PCIe

## PCIe Device lending

NVM Express (NVMe) over Fabrics defines a common architecture that supports a range of networking hardware (e.g. Infiniband, RoCE, iWARP) for a NVMe block storage protocol over a networking fabric.

NVMe devices have a direct PCIe interface. NVMe over Fabrics defines a software stack to implement transport abstraction layers at both sides of the fabric interfaces. These layers translate native PCIe transactions and disk operations over the fabric.

A PCIe fabric is different. Native PCIe transactions (TLPs) are forwarded automatically over the fabric with no protocol conversion. Standard PCIe NT technology is used to route the PCIe traffic from the host computer to the NVMe device. Device Interrupts are also automatically routed through the PCIe fabric. This image illustrates the difference between using a PCIe Fabric and other Fabrics. PCIe Fabrics eliminate the transport abstraction, thus providing a much lower latency. It also still supports features such as RDMA.



## Benefits

- Share SRIOV NVMe drives between systems
- Fast access to remote drivers
- Use native NVMe driver
- Near native performance with no translation

## Features

No modifications to Linux or device drivers. Works with all PCIe based storage solutions

Send and receive native NVMe commands directly to and from the fabric.

Optimized for low latency. Latency as low as 540 ns data transfers.

Supports direct memory region access for applications.

Scales to 100s of devices or more.

Multiple ports for simultaneous communication.

Self-throttling, guaranteed delivery at hardware level, no dropped frames or packets due to congestion.


Support for hot plugging of cables, full error containment and transparent recovery.

# PCIe Expansion

## Transparent Expansion

Dolphin offers a line of standard host and target adapters for use in traditional IO expansion. These cards require no software and rely on standard PCI Express bridging technology to connect downstream targets. Dolphin adapter cards can be used to connect expansion boxes or other standard PCIe IO boxes, such as JBOFs, cameras, or sensors. Dolphin also offers a transparent board management software, that manages and monitors our transparent cards. This software provides reporting of firmware revision, board serial number, active configuration setting, and PCIe link and host slot status.

Our transparent offerings include host and target adapters along with switches. This figure illustrates various configurations possible with our hardware.



## Benefits

- Gen3 Host and Target adapters
- Support for large topologies using PCIe Switch boxes
- Board and status monitoring
- iPass, MiniSAS-HD, and PCIe 3.0 cable support
- Fiber optic support

## Features

Fiber optic support up to 100 meters

Gen3 and Gen2 adapters, Gen4 under development

iPass, MiniSAS-HD and PCIe 3.0 cable support

Transparent board management software

PCIe link and host slot error monitoring


Automatic link training for both speed and width - (Gen3 -Gen1, x16 -x1)

Advanced EEPROM configurations

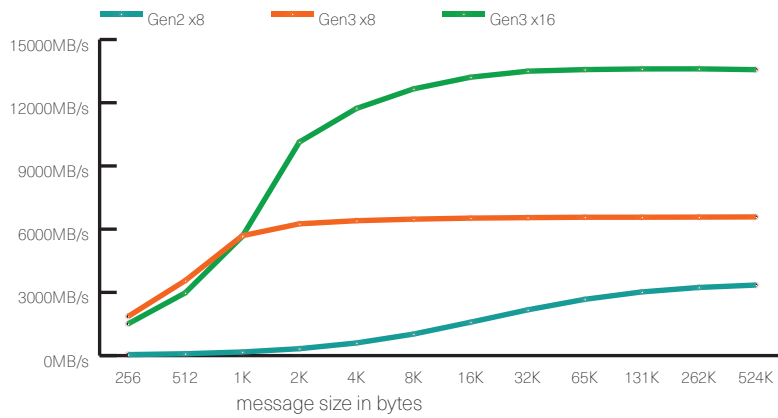
Long cable support



# *Performance/ Topologies*



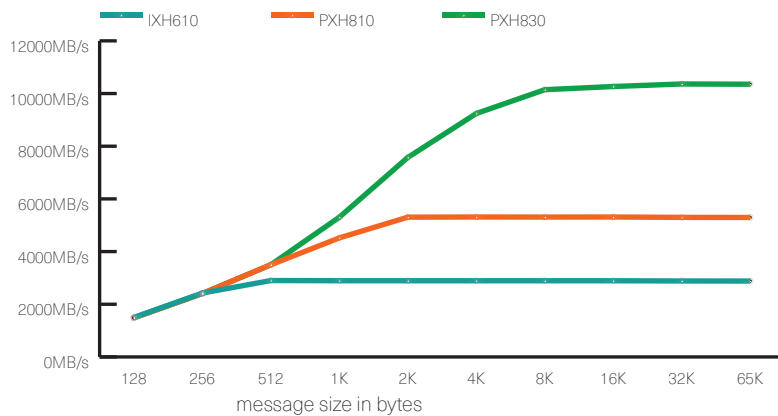
PCI Express provides various aspects of performance including latency, throughput, and scalability. Dolphin's products expose these aspects of performance with our eXpressWare software and high performance hardware. The combination enables building various system topologies that deliver high performance results.



## Throughput

### DMA Throughput

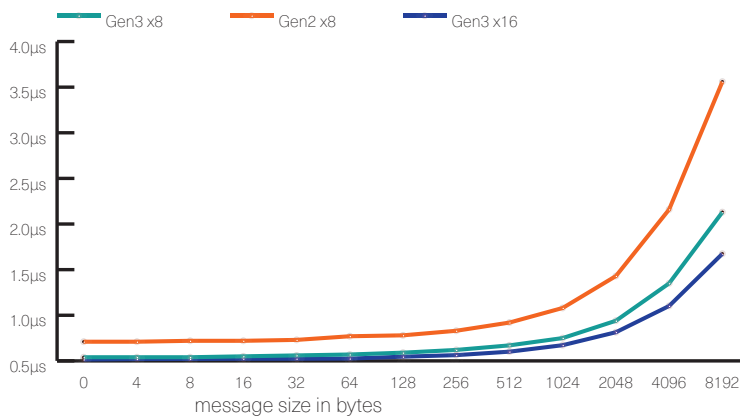
A DMA benchmark is included with the eXpressWare suite. The graph on the left shows the throughput of different PCIe generations. The current Gen3 PXH830 can achieve 13 GBps with DMA enabled.



## Throughput

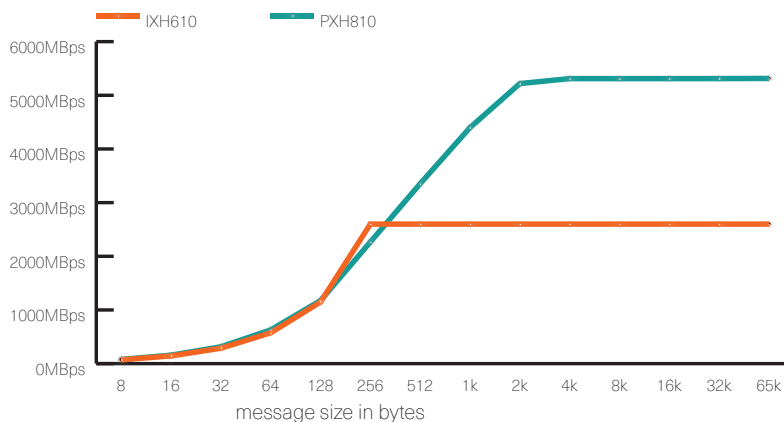
### Programmed IO (PIO) Throughput

The PIO benchmark provides a CPU only throughput analysis. This performance varies based on CPU. This chart shows that over 10 GBps can be achieved using only PIO.




## Latency

PCIe latencies are some of the lowest in the industry. This chart illustrates consistently low latency. Our latencies start at 540 ns between nodes. We test this with a loop back test, so this is actual application latency.




## Reflective Memory

Dolphin's reflective memory solution improves performance vs older generations of reflective memory solutions. It provides higher throughput along with the lower latency of PCI Express.



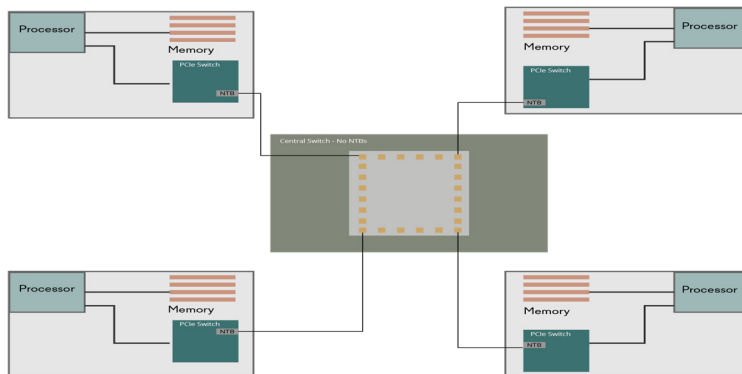
## Two Node

PCI Express is a point-to-point protocol. Using non-transparent bridging (NTB), two systems can be connected through a PCI Express connection. This connection can be a cable or through a backplane.



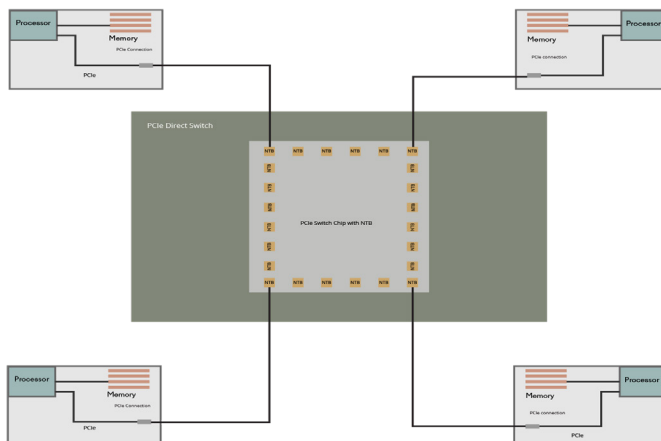
## Switchless

Depending on the number of NTBs, three or more nodes can be connected without using an external switch. These topologies are high performance as they don't introduce the latency of another switch chip.



## Central Switch

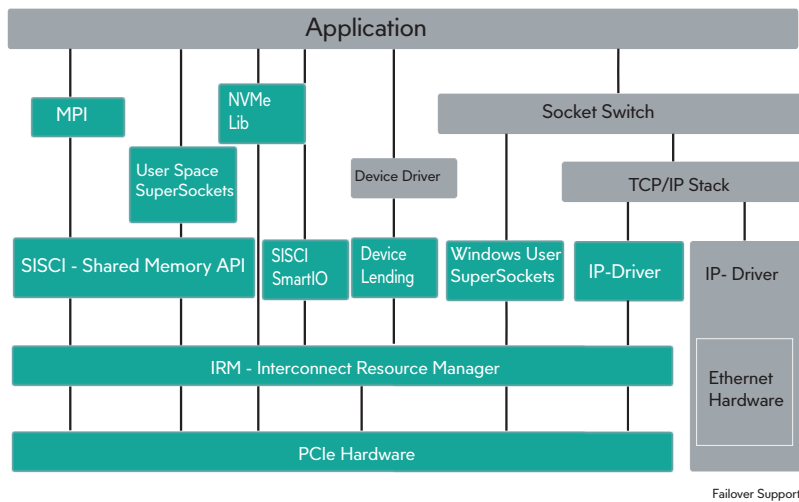
To scale out topologies, an external switch can be introduced. The NTBs are located on the host adapters to provide higher availability. Multiple switches can be cascaded to create large clusters.



## Direct Switch

In a direct switch topology, the NTBs are located in the central switch and not on the host adapters. This is ideal for backplane applications since it eliminates the need for power-on sequencing and provides better performance.

# eXpressWare



## Overview

eXpressWare is an advanced software suite for PCI Express hardware. It supports the non-transparent features in many PCIe switches from leading PCIe vendors.

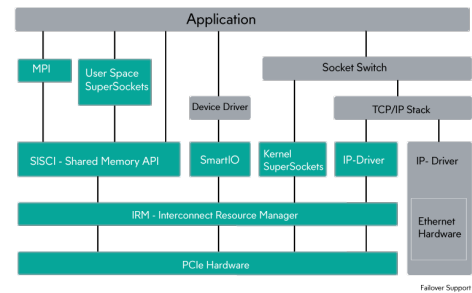
eXpressWare includes a low level API called SISCi, a sockets API - SuperSockets, an IPoPCiE driver, and advanced IO software - SmartIO for sharing PCIe devices.

eXpressWare comes with all Dolphin eXpressWare NTB adapters. It also can be licensed for third party PCIe hardware.

# SISCI API

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 and can be used for both intra-and inter-system communication.

In a multiprocessing architecture with PCI Express®, the SISCI API enables PCI Express® based applications to use distributed resources such as CPUs, I/O, and memory. The resulting application features reduced system latency and increased data throughput. The Dolphin SISCI Developers kit consists of driver and API software, tools, documentation and source files needed to develop your own embedded application. The development kit provides a C system call interface to ease customer integration to the PCI Express Network. SISCI enables customer application to easily bypass the limitations of traditional network solutions, avoiding time consuming operating system calls, and network protocol software overhead. The SISCI software supports clusters of hundreds of nodes.



## Benefits

- Low Latency API designed for direct memory access between multiple systems and devices
- Supports Windows, Linux, VxWorks, and RTX
- Cross O/S support between Windows and Linux
- Example code and performance test suite for latency and throughput
- Peer-to-Peer support and memory writes into device memory of GPUs and FPGAs

## Features

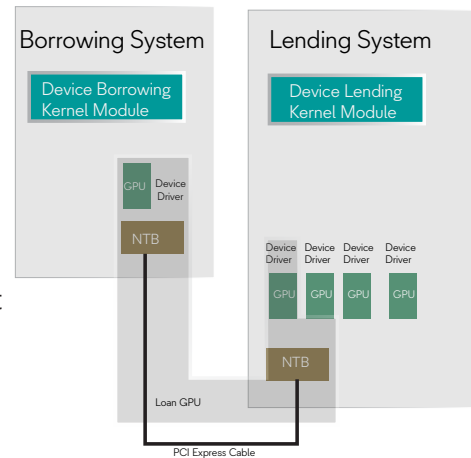
Shared memory API for clusters and SMP's	Create and trigger remote application level interrupts
Reflective Memory / PCIe Multicast	Direct remote CPU load and stores, memcpy()
PCI Express Peer-to-Peer support	Windows, RTX, VxWorks and Linux OS support
Distributed shared memory	Caching and error checking support
Easy deployment of DMA transfers	Events and callbacks

# Device Lending

PCIe Device Lending offers a flexible way for NVMeS, FPGAs, GPUs, and other devices to be accessed within a PCIe Fabric. PCIe devices can be borrowed between nodes in a PCIe Fabric. This is done without any software overhead and at PCIe data rates.

Device Lending is a simple way to reconfigure systems and reallocate resources. GPUs, NVMe drives or FPGAs can be logically added from and to any system in the fabric. The result is a flexible simple method to maximize usage by creating a pool of devices. Since this solution uses standard PCIe, it doesn't add any software overhead to the communication path. Standard PCIe transactions are used between the systems.

Device lending also enables a SR-IOV device to be shared as a MR-IOV device. SR-IOV functions can be borrowed by any system in the PCIe Fabric. Thereby enabling the device to be shared by multiple systems. This maximizes the use of SR-IOV devices such as 100 Gbit Ethernet cards.



## Benefits

- Dynamic reallocation of NVMe Drives and other resources
- Flexible use of GPUs
- Sharing of SR-IOV devices

## Features

Linux OS support	Lending and Borrowing Software
Command line interface	API to enable scripting
Any system can be lender and borrower	Supports GPUs, FPGAs, Ethernet cards, NVMe drivers and other devices
Supports Hot add supported devices only	Enables sharing of devices with SR-IOV between multiple system.


# PCIe Hot Add

PCIe NTB technology enables software based resource management and enumeration of standard PCIe devices. This process of adding or removing devices can be automatic or managed through some basic command line utilities.


The main difference between PCIe hot add and traditional expansion setups is that the PCIe host uplink card is replaced with a Dolphin PCIe NTB card. eXpressWare software is used to enumerate downstream devices vs the system BIOS. The solution works with directly attached PCIe devices as well as PCIe devices installed into PCI Express expansion systems.

The solution does not require any special BIOS version. Hot adding many devices or devices that require a large PCIe address space may require the BIOS to support 64bit encoding and large PCIe BARs.


Single device:



Multiple devices in an expansion box:



Using a switch to connect multiple expansion boxes:



## Benefits

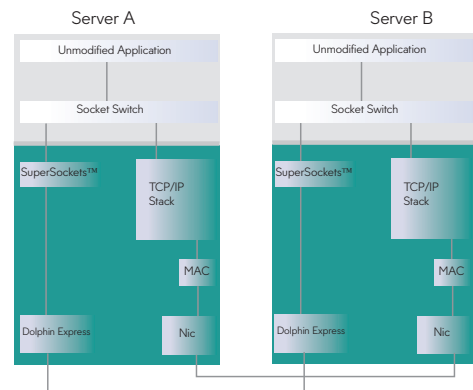
- Add multiple devices or new devices to a running system
- No power-on sequencing
- Dynamic Hot Add and replace
- Works with standard PCIe devices (NVMe drives)

## Features

NTB host adapter	Dolphin eXpressWare software
Software based resource management and enumeration	Automatic management of device adding and removing
Scalability	Add or remove a number of devices
No prior installation required	No changes to device drivers

# SuperSockets

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. It transparently achieves performance gains for existing socket-based network applications. Dolphin 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.



## Benefits

- Increased application performance without changing your application
- Extremely low latency with under 1µs of latency
- Highly available network with failover to Ethernet
- High bandwidth and throughput

## Features

Compliant with Linux Socket library and Berkeley Sockets	No OS patches or application modifications required
All popular Linux distributions supported	Supports both user space and kernel space clients
Both TCP and UDP support	Transparent failover to Ethernet if high speed connection is down.

# IPoPCIE

Dolphin’s 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 higher bandwidth than network technologies like 10G Ethernet.



Feature	SISCI	SuperSockets	IPoPCle Driver	SISCI SmartIODevice Lending	
Linux Platform Support	x86, x86_64	x86, x86_64	x86, x86_64	x86, x86_64	
Kernel 2.6 - 4.x	ARM 32/64. PPC	ARM 64 PPC	ARM 64 PPC	ARM 32/64. PPC	x86, x86_64
Windows Vista, 7, 10, Server 2008-2019 Platform Support	x86, x86_64	x86, x86_64	x86, x86_64	coming	2020
VxWorks 6.9 , 7.0	x86, x86_64 PPC	No	No	No	No
Latency	0.54 $\mu$ s Linux	0.98 $\mu$ s Linux	5.6 $\mu$ s Linux	Native PCIe	Native PCIe
Max Bandwidth	13 GB/s Linux	65 Gb/s	65 Gb/s	Native PCIe	Native PCIe
Reflective Memory Support	Yes	N/A	N/A	N/A	N/A
Address based Multi-cast support	Yes	N/A	N/A	N/A	N/A
TCP and UDP Support	N/A	TCP -Yes UDP Linux - Yes UDP Windows - No	Yes	N/A	N/A
UDP Multi-cast Support	N/A	Linux -Yes Windows -No	Linux - No Windows - Yes	N/A	N/A
Application modifications required or unmodified device driver	Yes	No	No	Yes	No
Cross O/S and Platform data transfer support	All	User space Linux / Windows	No	N/A	2020
Peer to Peer Transfers support	Yes	N/A	N/A	Yes	Yes
Sharing of NON-SR-IOV devices				Yes	No
Sharing SR-IOV devices to multiple hosts				Yes	Yes
Direct Transfers NVME to GPU / FPGA buffers				Yes	No
Flexible device memory management				Yes,	No
Support for GPU-Direct				Yes	Yes
Support for PCIe Multicast - device				Yes	No
PCIe DMA support				Yes	No
Device DMA support				Yes	Yes

# Hardware



Dolphin develops PCIe hardware for use with standard servers, workstations, and industrial form factor systems. This hardware supports various PCIe specifications along with advanced features offered by a variety of PCIe switch vendors.

Dolphin hardware enables its customers to quickly deploy and evaluate the latest PCIe standards for performance and features. These products are available for volume applications to allow customers to focus their effort on other components of the system.

Dolphin hardware supports advanced features based on non-transparent bridging. These features include clustering, sharing of PCIe devices, and PCIe multi-cast. The system illustrated in the picture is a 60 node PCIe cluster using Dolphin hardware.


# eXpressWare Host Adapters

## Gen3 Host Adapters

PCI Express NTB Host Adapters are high performance cabled interface to external processor sub-systems. Based on standard PCI Express bridging architectures, the host adapters include advanced features for non-transparent bridging (NTB) and clock isolation. These host adapters combine high performance with an application-to-application latency starting at sub-1 microsecond. They support both Remote Direct Memory Access (RDMA) and Programmed IO (PIO) transfers, effectively supporting both large and small data packets. RDMA transfers result in efficient larger packet transfers and processor off-load. PIO transfers optimize small packet transfers at the lowest latency. The combination of RDMA and PIO creates a highly potent data transfer system.

Our host adapters are carefully designed for maximum cable length. They support copper cables at various lengths at full PCI Express speeds. Fiber optics are available on some models to extend cable distances up to 100 meters.


Our eXpressWare™ software suite is included with all eXpressWare NTB host adaptercards.



### PXH830

*PCIe Gen3 x16 NTB Host Adapter*


- One x16 PCIe edge port
- Four x4 PCIe cable ports
- PCIe 3.0 or MiniSAS-HD cables
- Low Profile Design
- Quad SFF-8644 Connectors
- Full Dolphin Software Stack
- 128 Gigabit/s performance



### PXH840

*PCIe Gen3 x16 NTB Host Adapter*


- One x16 PCIe edge port
- Configurable up to 4 fiber optic ports
- Low Profile Design
- 128 Gigabit/s performance
- PIO and DMA RDMA
- Full Dolphin Software Stack



### PXH820

*PCIe Gen3 XMC NTB Adapter*


- x4, x8 or x16 PCIe ports
- PIO and DMA RDMA
- Quad SFF-8644 cable port
- VITA 42.0 XMC 1.0
- Up to 9M copper, 100M fiber cables
- Transparent & Non Transparent bridging



## **PXH810**

*Gen3 PCI Express NTB Host Adapter*

- Gen3 64 Gbps
- PCI Express iPass Connectors
- Supports Gen1,2,and 3 PCIe
- RDMA through PIO and DMA
- Clock isolation for CFC or SSC
- Low Profile Design
- Copper cables up to 5m
- Fiber-optic cables up to 100m
- Non-transparent bridging
- EEPROM configurable
- Link status LEDs
- Full Dolphin Software stack




## **MXH830**

*PCI Express Gen 3 x16 NTB Host Adapter*

- One x16 PCIe edge port
- Four x4 PCIe Cable ports
- PCIe 3.0 or MiniSAS-HD cables
- 0.54 microsecond latency
- Low Profile Design
- Quad SFF-8644 Connectors
- Non Transparent Bridging
- 128 Gbit/s Performance
- Full Dolphin Software Stack


## Gen2 Host Adapters



### IXH610

#### *PCI Express Gen2 Host Adapter*

- x8 PCI Express port
- PIO and DMA RDMA
- iPass Connector
- Low Profile Design
- Up to 7m Copper Cables
- Transparent HOST bridging
- Non Transparent Bridging
- 40 Gbit/s Performance



### IXH620

#### *XMC PCI Express Gen2 Adapter*

- x8 PCI Express port
- PIO and DMA RDMA
- iPass Connector
- VITA 42.0-2005,ANSI/VITA 42.3-2006
- Up to 7m Copper Cables
- Transparent bridging
- Non Transparent Bridging
- 40 Gbit/s Performance

Feature	PXH830	PXH840	PXH820	MXH830	PXH810
Port Speeds	32 GT/s per port / 128 GT/s max	32 GT/s per port / 128 GT/s max	32 GT/s per port / 128 GT/s max	32 GT/s per port / 128 GT/s max	64 GT/s per port/ 64 GT/s max
Max Application Performance	540 nanoseconds latency 13.2 GB/s throughput	540 nanoseconds latency 13.2 GB/s throughput	540 nanoseconds latency 13.2 GB/s throughput	600 nanoseconds latency 13.2 GB/s throughput	540 nanoseconds latency 6.8 GB/s throughput
Active Components	Broadcom/PLX Gen 3 PCIe Switch with DMA	Broadcom/PLX Gen 3 PCIe Switch with DMA	Broadcom/PLX Gen 3 PCIe Switch with DMA	Microsemi Switchtec PFX Gen 3 PCIe Switch	Broadcom/PLX Gen 3 PCIe Switch with DMA
Specifications	Base Spec 3.0 Cable Spec 3.0 CEM Spec 3.0	Base Spec 3.0 Cable Spec 3.0 CEM Spec 3.0	Base Spec 3.0 Cable Spec 3.0 VITA 41	Base Spec 3.1 Cable Spec 3.0 CEM Spec 3.0	Base Spec 3.0 Cable Spec 2.0 CEM Spec 3.0
Topologies	Point to point 3 node Mesh Topology Switch topology with MXS824	Point to point 3 node Mesh Topology	Point to point 3 node Mesh Topology Switch topology with MXS824	Point to point 3 node Mesh Topology 5 node Mesh Topology Switch topology with MXS824	Point to point 3 node Mesh Topology IXH600 Switch topology
Cable Connections	SFF-8644 connector Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper and MiniSAS-HD copper / PCIe fiber	Fiber optic cable (AOC) support MPO Fiber cabling	SFF-8644 connector Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper and MiniSAS-HD copper / PCIe fiber	SFF-8644 connector Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper and MiniSAS-HD copper / PCIe fiber	x8 iPass copper cable, fiber optic cable support Supports x8 to x4 transition cables
Max Cable Distance	9m with copper 100m with fiber	100 meters with fiber	9m with copper 100m with fiber	9m with copper 100m with fiber	5m with copper 100m with fiber
Mechanical Dimensions	Low profile, Half Length - 68.90 mm (2.731 in) x 167.65 mm (6.600 in)	Low profile, Half Length - 68.90 mm (2.731 inches) x 167.65 mm (6.600 inches)	Single width XMC	Low profile, Half Length - 69.2 mm (2.7 inches) x 168.2 mm (6.6 inches)	Low profile - 68.90 mm (2.731 inches) x 120 mm (6.600 inches)
Bracket	Full height bracket mounted Half height bracket included	Full height bracket mounted Half height bracket not included, only available for x8 and x4 configurations	XMC front panel face plate	Full height bracket mounted Half height bracket included	Full height bracket mounted Half height bracket included
Product Codes	PXH830 NTB Adapter	PXH840 - 2x8 - NTB x16/ NTB 2x8 PXH841 - 1x8 - NTB x8 PXH843 - 2x4 - NTB x8/ NTB 2x4 PXH844 - 1x4 - NTB x4	PXH820 Vita 41 NTB Adapter	MXH830 NTB Adapter	PXH810 NTB Adapter


# PCIe Host/Target Adapters

## Standard Gen3 Adapters

Dolphin's PCIe Host/Target Adapters are designed for use in IO expansion applications. All transparent host and target adapters support standard features for transparent bridging (TB), along with host and clock isolation.

These adapters support various cabling specifications including iPass or the new PCI Sig 3.0 specification. Some adapters can also be used with a standard MiniSAS-HD cable when the Dolphin card is both the host and target adapter. SFF-8644 adapters support the new CMI functionality when used with a PCIe 3.0 cable.


The adapters are carefully designed for maximum cable length. Some boards support copper cables up to 9 meters at full PCI Express Gen3 speed. Fiber optics extends this distance to 100 meters.



### PXH832

*Gen3 PCI Express Transparent Host and Target Adapter*


- x16 PCI Express edge port
- 4 x4 cable ports
- Total 128 Gbit/s Throughput
- 138 ns Cut-Through Latency
- Low Profile Design
- Quad SFF-8644 connector
- MiniSAS-HD / PCIe 3.0 cables



### PXH842

*Gen3 PCI Express Transparent Host and Target Adapter*



- x16 PCI Express edge port
- Two x8/x4 PCIe optical cable ports
- Total 128 Gbit/s Throughput
- 138 ns Cut-Through Latency
- Low Profile Design
- Dual MTP connectors



### PXH822

*PCIe Gen3 XMC Host/Target Adapter*

- x4, x8 or x16 PCI Express host port
- Quad SFF-8644 cable port
- VITA 42.0 XMC 1.0
- Prepared for Rear IO
- x4, x8 or x16 PCI Express cable port



## PXH812

*PCI Express Gen3 Host and Target Adapter*

- 8.0 Gbps per lane
- PCI Express iPass Connectors
- Supports Gen1,2,and 3 PCIe
- Clock isolation for CFC or SSC
- Low Profile Design
- Copper connection up to 5m
- Fiber-optic - up to 100m
- Transparent host and target
- EEPROM configurable
- Link status LEDs


## MXH832

*Gen3 x16 PCIe Transparent Host / Target Adapter*

- One x16 PCIe edge port
- Four x4 PCIe Cable ports
- PCIe 3.0 or MiniSAS-HD cables
- 170 nanoseconds cut through latency
- Quad SFF-8644 Connectors
- Host / Target Transparent Bridging
- Low Profile Design
- Optional board management software



## Gen2 Host Adapters



### IXH610

#### *PCI Express Gen2 Host Adapter*

- x8 PCI Express port
- PIO and DMA RDMA
- iPass Connector
- Low Profile Design
- Up to 7m Copper Cables
- Transparent HOST bridging
- Non Transparent Bridging
- 40 GT/s Performance



### IXH620

#### *XMC PCI Express Gen2 Adapter*

- x8 PCI Express port
- PIO and DMA RDMA
- iPass Connector
- VITA 42.0-2005,ANSI/VITA 42.3-2006
- Up to 7m Copper Cables
- Transparent bridging
- Non Transparent Bridging
- 40 GT/s Performance

Feature	PXH832	PXH842	PXH822	MXH832	PXH812
Port Speeds	32 GT/s per port / 128 GT/s max	32 GT/s per port / 128 GT/s max	32 GT/s per port / 128 GT/s max	32 GT/s per port / 128 GT/s max	64 GT/s per port/ 64 GT/s max
Port to port latency	138 nanoseconds	138 nanoseconds	138 nanoseconds	170 nanoseconds	138 nanoseconds
Active Components	Broadcom/PLX Gen 3 PCIe Switch with DMA	Broadcom/PLX Gen 3 PCIe Switch with DMA	Broadcom/PLX Gen 3 PCIe Switch with DMA	Microsemi Switchtec PFX Gen 3 PCIe Switch	Broadcom/PLX Gen 3 PCIe Switch with DMA
Specifications	Base Spec 3.0 Cable Spec 3.0 CEM Spec 3.0.	Base Spec 3.0 Cable Spec 3.0 CEM Spec 3.0.	Base Spec 3.0 Cable Spec 3.0 VITA 42	Base Spec 3.1 Cable Spec 3.0 CEM Spec 3.0	Base Spec 3.0 Cable Spec 2.0 CEM Spec 3.0
Topologies	Transparent Host/Target up to 4 exp systems MXS824 Switch topology	Transparent Host/Target up to 4 exp systems	Transparent Host/Target up to 4 exp systems MXS824 Switch topology	Transparent Host/Target up to 4 exp systems MXS824 Switch topology	Transparent Host/Target Switch topology with IXS600
Cable Connections	SFF-8644 connector Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper and MiniSAS-HD copper / fiber	Fiber optic cable (AOC) support MPO Fiber cabling	SFF-8644 connector Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper, MiniSAS-HD copper and PCIe fiber	SFF-8644 connector Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper, MiniSAS-HD copper and PCIe fiber	x8 iPass copper cable, fiber optic cable support Supports x8 to x4 transition cables
Max Cable Distance	9m with copper 100m with fiber	100 meters with fiber	9m with copper 100m with fiber	9m with copper 100m with fiber	5m with copper 100m with fiber
Power Consumption	12 Volt : 14 Watt Max 12 Volt: Max 18 Watts incl AOC +3.3 Volt : Max 3.3 Watt +3.3 VAUX : Max 1 Wattz	12 Volt : 14 Watt Max +3.3 Volt : Max 3.3 Watt +3.3 VAUX : Max 1 Watt	12 Volt : 14 Watt Max 12 Volt: Max 18 Watts incl AOC +3.3 Volt : Max 3.3 Watt +3.3 VAUX : Max 1 Watt	12 Volt: Max 19.6 Watts 12 Volt: Max 24 Watts incl AOC +3.3 Volt: Not connected +3.3 VAUX: Max 1 Watt	12 Volt: 12 Watt Max +3.3 VAUX : Max 40 milliwatt
Mechanical Dimensions	Low profile, Half Length - 68.90 mm (2.731 inches) x 167.65 mm (6.600 inches)	Low profile, Half Length - 68.90 mm (2.731 inches) x 167.65 mm (6.600 inches)	Single width XMC	Low profile, Half Length - 69.2 mm (2.7 inches) x 168.2 mm (6.6 inches)	Low profile - 68.90 mm ( 2.731 inches) x120 mm (6.600 inches)
PCIe Bracket	Full height bracket mounted Half height bracket included	Full height bracket mounted Half height bracket not included, only available x8 and x4 configurations	XMC front panel face plate	Full height bracket mounted Half height bracket included	Full height bracket mounted Half height bracket included
Product Codes	PXH832 Host/Target Adapter	PXH842- 2x8 x16/2x8 PXH845 - 4x4 x16/ 2x8/4x4 PXH846 - 1x8 x8 PXH847 - 2x4 x8/ 2x4 PXH848 - 1x4 x4	PXH822 Host/Target Adapter	MXH832 Host/Target Adapter	PXH812 Host/Target Adapter

# PCIe Switches

Dolphin PCI Express switch boxes provide low latency, highly efficient switching for high performance applications. These powerful switch products enable I/O scaling and inter-processor communication by combining transparent and non-transparent bridging capabilities with Dolphin's eXpressWare software and clustering technology. Users can connect multiple PCI Express end points or create a highly efficient compute cluster with PCs, servers, or SBCs with XMC sites.

Dolphin's switch boxes work as the switching element within Dolphin's product line. Each switch port delivers maximum bandwidth to each device while maintaining backwards compatibility with Gen1 and Gen2 components. These products utilize standard cable connectors such as iPass and MiniSAS-HD as well as PCIe 3.0 cables. They support both copper or fiber-optic cabling. Customers can link multiple standardized PCI Express products such as PXI chassis, storage, and I/O expansion units.



## MXS824

*PCIe Gen3 24 port switch*

- 24 PCI Express Gen3 x4 ports
- NTB or Transparent use
- SFF-8644 Connectors
- 19 Inch 1U rack mountable
- Copper and Fiber-optic cables
- PCIe 3.0 or MiniSAS-HD cables



## IXS600


*PCIe Gen3 8 port switch*

- 8 PCI Express Gen3 x8 ports
- NTB and Transparent use
- iPass x8 Connectors
- 19 Inch 1U rackmountable
- iPass x8 Copper cables

Feature	MXS824 PCIe Switch	IXS600 PCIe Switch
PCIe	Gen3	Gen3
Link Speed	32 GT/s per x4	64 GT/s per x8
Port Latency	170ns	200ns
Active Components	Microsemi Switchtec® PCI Express Gen3 Switch	IDT PCI Express Gen3 Switch
PCI Express	Base Specification 3.1	Base Specification 3.0
Ports	24 - x4 non blocking PCIe Gen3 ports	8 - x8 Non blocking PCIe Gen3 ports
Port configurations	Flexible port configurations - E.g 24 -x4, 12 -x8, 6 - x16 ports or 2x16+2x+12x4	8- x8 ports
Operating Modes	NTB and Transparent	NTB and Transparent
Connector Type	SFF-8644	iPass
Cable Connections	x4 MiniSAS HD or PCIe 3.0 cable connections	x8 - iPass Cable connections
Cable Types	Copper cables - up to 5 meters at Gen3 speeds Fiber-Optic cables - up to 100 meters at Gen3 speeds	Copper Cables - up to 3 meter at Gen3 speeds
Dimensions	1U 19 inch rack mount	1U 19 inch rack mount
Power consumption	Auto-sensing power supply 110-240v AC 50- 60 Hz	Auto-sensing power supply 110-240v AC 50- 60 Hz
Additional ports	1 Ethernet Management Port, 1 USB port for firmware, 1 SD-card slot	1 Ethernet Management Port, 1 USB port for firmware
Power management	Advanced power management - hibernation, wake via Ethernet	N/A

# PCIe Backplanes and Cables

Dolphin offers a variety of cable and backplane offerings. These products are qualified and tested to work with Dolphin host adapters and switch products. The cables come in various lengths. Dolphin supports MiniSAS-HD, PCIe 3.0 and iPASS cables in copper and PCIe fiber cables. We also sell MPO/ MTP cables for our PXH84X based products.



## IBP-G3X16-2

*Two slot Gen3 PCIe backplane*

- Two x16 Gen3 slots
- Gen1, Gen2 and Gen3 compliant
- Compliant to all Dolphin PXH transparent host cards
- Compliant to ATX & MicroATX footprint



## PCIe3C-xM

*PCIe 3.0 cables*

- Cable Management Interface functionality (CMI)
- PCIe External cabling 3.0 spec
- Lengths 1 - 9 Meters
- Supports PXH and MXH SFF-8644 transparent boards



## MSCxM

*PCIe MiniSAS-HD cables*


- No CMI support
- Lengths 1-9 Meters
- Supports PXH, MXH and MXS SFF-8644 transparent and non-transparent boards.



### **MSFCxM**

*PCIe Gen3 SFF-8644 Fiber cables*


- No CMI Support
- Lengths 10, 50, 100 meters
- Supports PXH, MXH and MXS SFF-8644 transparent and non-transparent boards.



### **MTFCXFF-xM**

*MTP/MPO Fiber Cables*

- MPO/MTP compatible cables
- Lengths 0.5 - 100 meters
- Supports all PXH84X products
- 12 strand and 24 strand fiber



### **IXCxM**

*PCIe iPass Cables*

- PCIe External cabling 2.0 spec
- Lengths 0.5- 7 meters
- Supports all IXH, IXS and PXH iPass transparent and non-transparent boards

# Gen4 Products

Dolphin is working on a number of Gen4 products. These products are expected to be in production in 2020. Gen4 products will provide a new level of performance, doubling the performance of Gen3 in regards to throughput. We will provide several host adapters for both transparent and non-transparent bridging. We will also upgrade our switching capacity with a new Gen4 1U switch.



## MXH930

*PCI Express Gen 4 x16 NTB Host Adapter*

- One x16 PCIe edge port
- Four x4 PCIe Cable ports
- PCIe 4.0 or MiniSAS-HD cables
- <0.5 microsecond latency
- Low Profile Design
- Quad SFF-8644 Connectors
- Non Transparent Bridging
- 256 Gbit/s Performance
- Full Dolphin Software Stack



## MXH932

*Gen4 x16 PCIe Transparent Host / Target Adapter*

- One x16 PCIe edge port
- Four x4 PCIe Cable ports
- PCIe 4.0 or MiniSAS-HD cables
- Copper and Fiber cables
- Quad SFF-8644 Connectors
- H/T Transparent Bridging
- Low Profile Design
- Optional board management software



## MXS924

*Gen4 PCI Express Switch*


- 24 PCI Express Gen4 x4 ports
- Gen4 8 GT/S per lane
- 32 GT/s per port
- NTB or Transparent use
- SFF-8644 Connectors
- PCIe 3.0 or MiniSAS-HD cables
- Hot Plug cabling support
- 19 Inch 1U rack mountable
- Redundant Fans

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 25 years. For more information, please visit [www.dolphinics.com](http://www.dolphinics.com).

Visit us on LinkedIn and Facebook:

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

<https://www.facebook.com/dolphinInterconnect>



©2019 Dolphin Interconnect Solutions AS All rights reserved. Dolphin ICS and the Dolphin Logo are registered trademarks. PCI, PCI Express, PCIe, Express Module, and PCI-SIG are trademarks or registered trademarks of PCI-SIG. All other trademarks are the property of their respective owners. All statements herein are based on normal operating conditions, are provided for information purposes only and are not intended to create any implied warranty of merchantability or fitness for a particular purpose. Dolphin Interconnect AS reserves the right to modify at any time without prior notice these statements, our products, their performance specifications, availability, price, and warranty and post-warranty programs.