

MXH830 PCIe NTB Adapter

NTB x16 Gen3 SFF-8644 Adapter



The MXH830 Gen3 PCI Express NTB Host Adapter is our high performance clustering product. Based on Microsemi Switchtec® Gen3 PCI Express bridging architecture, the MXH830 host adapter includes advanced features for non-transparent bridging (NTB) and clock isolation.

For high performance application developers, the MXH830 host adapter combines 128 GT/s performance with an application to application latency starting just above 500 nanoseconds. Inter-processor communication benefits from the high throughput and low latency. Using the latest SmartIO technology software from Dolphin, GPU / Cuda applications can now stream data to remote GPUs at the same speed as to local GPUs.

The card implements a quad SFF-8644 connector which is compliant with the latest MiniSAS-HD and PCI SIG External cable specification 3.0 cables. Cable distance

extend to 9 meters with copper cables and 100 meters with fiber cables. These four ports can be aggregated into a single x16 port that delivers a 128 GT/s data link to remote systems or partitioned to two x8 ports or four x4 ports.

The MXH830 includes Dolphin's eXpressWare™ software suite license. eXpressWare provides a software infrastructure for developing PCIe applications. This software suite takes advantage of system DMA and PIO data transfer schemes to create a complete environment for customized and standardized applications. The software suite includes a Shared-Memory Cluster Interconnect (SISCI) API as well as a IPoPCle drive, SuperSockets software and SmartIO software. The SISCI API is a robust and powerful shared memory programming environment. This API enables the advanced features of PCIe such as PCIe multi-cast, remote

peer-to-peer transfers, and allocation of various size system or device memory for data transfers. The optimized IPoPCle driver and SuperSockets software remove traditional networking bottlenecks, allowing standard IP and sockets applications to take advantage of PCIe's high performance without modification. The overall framework is designed for rapid deployment of inter-processor communication systems. SmartIO software enables peer to peer transfers and device lending between systems.

The MXH830 support both switchless and switched configurations. For scaling out beyond five nodes, the MXH830 is combined with Dolphin MXS824 switch. The MXS824 is a 24 port 1U PCIe switch that can be configured a twenty four x4 ports, twelve x8 ports, or six x16 ports. This switch can also be cascaded to create large topology that can expand up to 64 or 128 nodes.

Features

- PCI Express Gen3 - 8.0 GT/s per lane
- Link compliant with Gen1, Gen2, and Gen3 PCIe
- Quad SFF-8644 connector
 - » PCI Express 3.0 cables
 - » MiniSAS -HD cables
- RDMA support through PIO and System DMA
- Copper and fiber-optic cable support
- Full host Clock isolation support. Automatic support for host running CFC or SSC.
- Low Profile PCIe form factor
- EEPROM for custom system configuration
- Link status LEDs through face plate
- <170ns - Cut Through latency
- MXH830 supports the following port configurations
 - » One - x16 PCIe port
 - » Two - x8 PCIe ports
 - » Four - x4 PCIe ports
- MXH830 can be configured into a 5 node cluster using 4 x4 ports to each host or a 3 node cluster using x8 connections, or a two node x16 connection
- Combined with the MXS824 can scale to up to 64 or 128 nodes



¹ System DMA is only available with select systems

eXpressWare PCIe Software

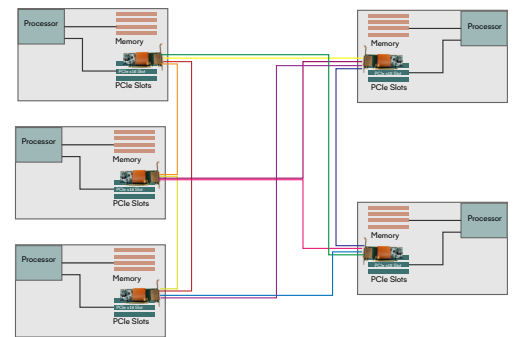

The MXH830 supports Dolphin's eXpressWare™ software suite. eXpressWare™ software enables customer to easily migrate applications to a PCIe Networks. A complete software infrastructure, eXpressWare enabling applications to network using standard PCI Express over cables. Several interfaces and APIs are supported including standard TCP/IP networking - IPoPCle driver, a low level direct remote memory access API - SISI API, and a sockets API -SuperSockets. Each API has its benefits and can be selected based on application requirements.

The SISI API enables customers to fully exploit the PCIe programming model without having to spend months developing device drivers. The API offers a C programming API for shared / remote memory access, including reflective memory/multi-cast functionality and RDMA capabilities.

SuperSockets enabled networked applications to benefit from a low latency, high throughput PCIe network without any modifications. With SuperSockets, a PCIe network can replace local Ethernet networks.

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

Dolphin SmartIO adds new features to standard PCIe IO. It enable device lending, which enables standard IO devices to be logically moved between systems. It also provides an infrastructure for peer to peer memory transfers. SmartIO supports direct FPGA to FPGA, GPU to GPU, or any combination of communication with FPGA,CPU's, GPU's or memory over PCIe.



Five node MXH830 cluster

Specifications

Link Speeds	32GT/s per port / 128 GT/s total
Application Performance	~500ns latency (application to application) Above 10 Gbytes/s throughput
Active Components	Microsemi Switchtec PFX Switch
PCI Express	Base Specification 3.1 Cable Specification 3.0 (preliminary) Card Electromechanical Specification 3.0
Topologies	Two nodes direct cable Three to five nodes mesh topology Switched Topology with MXS824
Cut-Through Latency	<170ns
Cable Connections	SFF-8644 connector for copper / fiber cables Supports 4 - x4/ 2 - x8 or 1 - x16 connections PCIe 3.0 copper and MiniSAS-HD copper / fiber up to 9m with copper, 100m with fiber
Power Consumption	12 Volt: Max 19.6 Watts 12 Volt: Max 24 Watts including AOC +3.3 Volt: Not connected +3.3 Volt AUX: Max 1 Watt
Mechanical Dimensions	Low profile - 69.2 mm (2.7 inches) x 168.2 mm (6.6 inches)

Operating Environment	Operating Temperature: 0°C - 55°C (32°F - 131°F) AirFlow: 150 LFM Relative Humidity: 5% -95% non-condensing
Dolphin Software	SuperSockets Berkley Sockets API Microsoft WinSock2/LSP support IPoPCle driver SISI API
Usage Modes	Non-transparent bridging
Regulatory	CE Mark FCC Class B UL94V-0 compliant RoHS Compliant
Configuration	DIP-switch x4,x8,x16 link / Safe Boot
Mounting Plates	Full height plate installed Half height plate included with board
Operating Systems supported	Windows, Linux, VxWorks
Product Codes	MXH830 Host NTB Adapter