




Dolphin PCI Express MXH930 Adapter



MXH930 NTB Adapter Users Guide
Version 1.7

Date: 30th November 2023

Table of Contents

DISCLAIMER	4
TERMS AND ACRONYMS	4
MXH930 HIGH LEVEL SPECIFICATION	5
MTBF BY TEMPERATURE AND ENVIRONMENT	6
PACKAGING	7
PRE-INSTALLATION QUESTIONS	7
PCIE SLOT DETERMINATION.....	7
MXH930 NTB CONFIGURATION.....	7
AIRFLOW - OPERATING ENVIRONMENT	7
<i>Thermal shutdown</i>	7
CABLE CONNECTIONS.....	8
<i>Active Optical Cables (AOC)</i>	8
INSTALLATION	9
STEP 1 - UNPACK BOARD	9
STEP 2 - CHANGE PCIE BRACKET IF NECESSARY	9
STEP 3 - CONFIGURE THE BOARD FOR PROPER OPERATION	9
STEP 4 - ENSURE PROPER AIRFLOW	9
STEP 5 - INSTALL THE ADAPTER CARD.....	9
STEP 6 - INSTALLING AND REMOVING THE CABLE.....	9
<i>Connecting the Cable</i>	9
<i>Cable stain relief</i>	10
<i>Supported Cables</i>	10
<i>Disconnecting the Cable</i>	10
STEP 7 - INSTALLING THE DOLPHIN EXPRESSWARE SOFTWARE	10
STEP 7 – VERIFY INSTALLATION & LEDs	10
OPERATION	11
CONFIGURATION AND DIP SWITCHES.....	11
<i>DIP Switch Bank – Configuration</i>	11
<i>Thermal shutdown</i>	11
<i>DIP Switch NTB settings for MXH930</i>	12
USE CASE A - 2 NODE CONFIGURATION – NTB MODE.....	13
<i>Connecting the cables for 2 Node Configurations, x16 link</i>	13
USE CASE B - 3 NODE CONFIGURATION – NTB MODE.....	14
<i>Connecting the cables for 3 Node Configurations, x8 link</i>	14
USE CASE C - 5 NODE CONFIGURATION – NTB MODE.....	14
<i>Connecting the cables for 5 Node Configurations, x4 link</i>	14
USE CASE D – MXS924 SWITCH CONFIGURATIONS – NTB MODE	15
<i>6 Node, x16 Link Width using MXS924</i>	15
<i>12 Node, x8 Link Width using MXS924</i>	15
<i>24 Node, x4 Link Width using MXS924</i>	16
<i>Larger PCIe networks using the MXS924</i>	16
EEPROM AND FIRMWARE UPGRADE	17
IDENTIFYING THE CARD.....	17
SUPPORT	ERROR! BOOKMARK NOT DEFINED.
TECHNICAL INFORMATION	19

BOARD REVISION HISTORY	19
SIMPLIFIED SCHEMATICS	19
PCIe CABLE PORT SIGNALS	20
EXTERNAL PCIe x4 CABLE CONNECTOR PIN-OUT	20
PCIe CABLE PORT MAPPING	20
COMPLIANCE AND REGULATORY TESTING.....	21
LIMITED WARRANTY.....	22
WARRANTY PERIOD.....	ERROR! BOOKMARK NOT DEFINED.
COVERAGE.....	ERROR! BOOKMARK NOT DEFINED.
SERVICE PROCEDURE	ERROR! BOOKMARK NOT DEFINED.

DISCLAIMER

DOLPHIN INTERCONNECT SOLUTIONS RESERVES THE RIGHT TO MAKE CHANGES WITHOUT FURTHER NOTICE TO ANY OF ITS PRODUCTS TO IMPROVE RELIABILITY, FUNCTION, OR DESIGN.

DO NOT USE INFORMATION IN THIS GUIDE TO DESIGN YOUR OWN PCI-SIG COMPLIANT PRODUCT, ALWAYS REFERENCE THE ORIGINAL PCI-SIG SPECIFICATION FOR DETAILS.

TO THE FULLEST EXTENT PERMITTED BY LAW, DOLPHIN WILL NOT BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS, LOST DATA, OR LOSS OF USE) ARISING OUT OF ANY USE OF DOLPHIN'S PRODUCTS, SOFTWARE OR SERVICE PROVIDED. DOLPHIN'S MAXIMUM LIABILITY WILL NOT EXCEED THE TOTAL AMOUNT PAID FOR THE PRODUCT BY THE PURCHASER.

LIFE SUPPORT POLICY

DOLPHIN INTERCONNECT SOLUTIONS' PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES.

ENVIRONMENTAL POLICY

Dolphin is minimizing the amount of printed documentation and software CDs in its shipments; please download additional documentation and software from www.dolphinics.com.

Terms and acronyms

Important terms and acronyms used in this manual

AOC	Active Optical Cable. PCIe fiber cable assembly available from Dolphin.
CMI	Cable Management Interface. The 2-wire management interface for communication between subsystems connected by a PCIe 3.0 cable. Details can be found in the PCI-SIG External Cabling Specification 3.0.
CMI Controller	Microcontroller on the MXH930 used to manage and implement CMI communications.
eXpressWare	Dolphin's software stack for PCIe clustering and IO. Please visit www.dolphinics.com for more information.
Host mode	The card is installed in a root complex / Host PC.
Lane	One PCI Express Lane contains a differential pair for transmission and a differential pair for reception.
Link	A collection of one or more PCI Express Lanes providing the communication path between an Upstream and Downstream Port.
PCIe 4.0 cable	Cable compliant to the preliminary PCI-SIG External Cabling Specification 4.0. Support for CMI.
Port	PCIe Cable port. The MXH930 has four x4 ports, named P1, P2, P3, P4. These ports can be merged to two x8 ports, or one x16 port. The physical ports are identified by text on the PCIe brackets.

MXH930 High Level Specification

The MXH930 is a low profile, half-length PCIe adapter that can be used for NTB and Transparent configurations. Once installed in any server's or PC's PCI Express compliant slot, a MXH930 can connect to another MXH930, MXH932 or the MXS924 switch. The MXH930 supports PCIe Gen1, Gen2, Gen3 and Gen4 speeds and x1, x2, x4, x8 and x16 link-widths. The card will operate at the highest common speed shared between the slot and the card (Gen4) and the widest common link-width (x16).

- PCI Express Base Specification, Rev. 4.0.
- PCI Express CEM Specification, Rev. 4.0.
- PCI Express External Cabling specification 4.0 (draft)
- PCI Express Gen4 16.0 GT/s per lane signaling – 256 GT/s total signaling.
- PCI Express Gen4 x16 edge connector. The card installs in any PCI Express slot that has a physical x16 connector.
- Compliant with PCI Express Gen1 through Gen4 computers and IO systems, auto detection.
- Quad SFF-8644 cable connector:
 - Durability max total 250 mating cycles
- Cable port configurations:
 - One x16 NTB port
 - Two x8 NTB ports
 - Four x4 NTB ports
- Microsemi Switchtec PM40036 PFX PCI Express Gen4 chipset with DMA.
- Pre-loaded NTB configurations, DIP switch selectable.
- Synthetic End Point mode for Hot Add operation (currently disabled, will be enabled by firmware update).
- 100 nanosecond cut-through latency port to port.
- Support for PCIe 3.0 and PCIe 4.0 SFF-8644 copper cables with CMI and PCIe Active Optical cables.
- Low profile, Half-length - PCI Express Electromechanical Specification, Rev 4.0.
- Dimensions 167.65mm (6.600 inches) x 68.90 mm (2.731 inches).
- Comes with both low profile and standard profile PCI Express bracket.
- Host clock isolation. Automatic support for host running CFC or SSC mode.
- VAUX powered board management controllers for flexible configuration and cable management.
- Flash recovery option. PFX Multiconfiguration support.
- Maximum power ratings:
 - 12 Volt: 1.66 A (no port power) 2.66 A (max port power)
 - +3.3 Volt: Not connected
 - +3.3 Vaux: 100 mA
- Typical power ratings:
 - 12 Volt: 1.5 A (no port power)
 - +3.3 Volt: Not connected
 - +3.3 Vaux: 100 mA
- Port power supply (per cable port):
 - VPWR: 3.2V - 3.4V, 0.86 A
 - VMGTPWR: 3.2V - 3.4V, 89 mA
- Operating Temperature: 0°C - 55°C (32°F - 131°F), Airflow: 350 LFM (Preliminary)
- Operating Temperature with AOC: TBD
- Relative Humidity: 5% - 95% (non- condensing)
- CE and FCC marking

MTBF by Temperature and Environment

The MTBF (in hours) for the MXH930 can be found in the table below. The numbers are calculated using the Telcordia SR-332 issue 2 (2006) standard.

Ambient Temp [°C]	Environment		
	Ground fixed, controlled	Ground fixed, uncontrolled	Ground mobile
0	4.101.016	2.050.508	683.503
5	3.543.448	1.771.724	590.575
10	3.028.513	1.514.257	504.752
15	2.560.400	1.280.200	426.733
20	2.141.924	1.070.962	356.987
25	1.774.116	887.058	295.686
30	1.456.142	728.071	242.690
35	1.185.515	592.758	197.586
40	958.455	479.227	159.742
45	770.360	385.180	128.393
50	616.260	308.130	102.710
55	491.191	245.595	81.865

Table 1: MTBF vs. Temperature and Environment

Packaging

The MXH930 includes the following components.

- MXH930 Adapter Board with full height profile PCIe bracket mounted.
- Low profile PCIe bracket.
- Anti-static bag.
- Getting started guide with serial number for quick download of Dolphins eXpressWare PCIe software.

Pre-Installation Questions

Certain steps should be taken prior to installing the MXH930. You should determine the following configuration requirements.

- Which PCIe slot and system will the card be installed in?
- What is the speed and link width of the slot that the card will be installed in?
- What is the operating environment in which the card will be installed?
- What type and length of cables will be used?
- How to establish proper operational conditions, temperature, and airflow.

PCIe Slot Determination

The MXH930 supports PCIe Gen1, Gen2, Gen3 and Gen4 speeds and x1, x2, x4, x8 and x16 link-widths. The slot width and speed will affect the performance of the card. The card can be physically installed in a x4, x8 or x16 connector. The card will auto configure the slot speed and width.

- Install the board in a slot connecting directly to the CPU for optimal shared memory performance.
- Install the board behind a PCIe switch in the IO system for optimal PCIe peer to peer performance (e.g. between GPUs, FPGAs)

TIP: After installing the eXpressWare software, run the “sisci_benchmarks” benchmark suite to verify the expected performance. Please contact Dolphin support if you need assistance to reach the expected performance.

MXH930 NTB Configuration

The MXH930 has a DIP switch bank to control the main configuration of the card. The DIP switch labeled SW1 can be found close to the upper edge of the board. The default DIP switch setting is single link x16 operations. Alternative configuration settings are two x8 links or four x4 links.

Airflow - Operating Environment

To maximize the lifetime of the product and maintain the warranty, please honor the specified operating temperature, and make sure the specified airflow is present. Special care should be considered when the MXH930 is used in office type cabinets with poor cooling in combination with other high energy consuming PCIe devices, e.g. not active cooled GPUs:

Operating Temperature: 0°C - 55°C (32°F - 131°F), Airflow: 350 LFM (Preliminary)
Operating Temperature with AOC: TBD

TIP: Immediately after software installation, please use the dis_diag tool to verify the temperature. If the chip temperature reported by dis_diag exceeds 95°C, it is strongly recommended to improve the airflow.

Thermal shutdown

The MXH930 card includes overheat protection. The BMC will automatically shut down the MXH930 card if a PCIe switch temperature beyond 105°C (221°F) is detected. Similarly, if an AOC is plugged in, the MXH930 card will automatically shut down if the max operating temperature announced by the connected AOC is violated by 15°C (59°F). Please consult the AOC data sheet for details.

NOTE: The MXH930 link LEDs will start flashing yellow when a PCIe switch temperature of 95°C (203°F) is reached or when the AOC has reached the max operating temperature. Please consult the AOC data sheet for details.

NOTE: The automatic shutdown can be disabled by setting DIP-Switch OPT6. Prior to disabling the automatic shutdown, please ensure you have established proper airflow.

NOTE: Operating the card or AOC above the maximum specified operating temperature will void the warranty.

NOTE: If the card shuts down due to an overtemperature situation – you must improve the cooling before completely power cycling the system.

Cable Connections

The MXH930 is designed to support both long and short copper cables as well as active optical cables (AOC).

The MXH930 cable ports are compliant to the SFF-8644 industry specification and supports PCIe cables compliant to the PCIe External Cabling Specification 4.0. Four x4 cables are needed for full PCIe x16 connectivity.

Active Optical Cables (AOC)

Dolphin offers active optical PCIe fiber cables up to 100 meters for the MXH930 cards. Both pigtails with MPO connectors and complete cable assemblies are available.

Installation

Step 1 - Unpack board

The MXH930 card is shipped in an anti-static bag to prevent static electricity damage. The card should only be removed from the bag after ensuring that anti-static precautions are taken. Static electricity from your clothes or work environment can damage your PCI Express adapter card or your PC. Always wear a grounded anti-static wrist strap while opening the PC and when the MXH930 is removed from the anti-static bag. Unpack the MXH930 from the anti-static bag using proper anti-static procedures.



Step 2 - Change PCIe Bracket if necessary

The MXH930 package includes a standard and low-profile PCI Express bracket. By default, the standard height bracket is installed on the board. If you need to replace the mounted bracket with a low-profile bracket, carefully unscrew the two mounting screws to remove the full height bracket. Save the two mounting screws and replace the bracket with the low-profile bracket. Use the two mounting screws to install the low-profile bracket. The screws should be carefully tightened but be careful not to overtighten. Make sure you are properly grounded to avoid static discharges that may destroy the adapter card before performing this procedure.

Step 3 - Configure the Board for Proper Operation

Set the DIP switch settings for proper operation depending on the firmware. Please refer to the section Configuration and DIP Switches on page 11 for details.

Step 4 - Ensure proper Airflow

Please pay proper attention to ensure the selected server / slot position provides the minimum required airflow.

TIP: Immediately after software installation, please use the dis_diag tool to verify the temperature. If the chip temperature reported by dis_diag exceeds 95°C, it is strongly recommended to improve the airflow.

Step 5 - Install the Adapter Card

Before installing the adapter card, make sure you are properly grounded to avoid static discharges that may destroy your computer or the adapter card. Ensure you are properly grounded before opening your computer or the anti-static bag containing the MXH930. Please follow your computer's or expansion chassis' manual on how to install a PCI Express card.



The MXH930 Adapter card can be installed into any PCI Express x16 slot. The MXH930 supports PCI Express Gen1, Gen2, Gen3 and Gen4 signaling. NOTE: A Gen4 slot is recommended as it typically doubles the performance compared to a Gen3 slot. The MXH930 is an x16 card, so maximum performance will only be attained if the slot provides full electrical x16 signaling.

The MXH930 supports hosts using either spread spectrum or constant frequency clocking. The card implements clock isolation.

Step 6 - Installing and Removing the Cable

Installing and removing cables should be done with both host and expansion system powered off. Please contact your Dolphin representative if you intend to continuously connect and disconnect the PCI Express cables.

Connecting the Cable

Please carefully install the cable connector into the connector housing on the MXH930 adapter card. Cable port 1 is located at the top of the PCIe bracket, ref Figure 1: Full length PCIe bracket to the right. To install the cable, match the cable house with the connector on the MXH930 adapter card. Use even pressure to insert the connector until it is secure. Adhere to ESD guidelines when installing the cables to ensure you don't damage the board. Note that for wider than x4 connections, the same cable-ports (i.e. port 1 to port 1, port 2 to port 2 through port 4) should be used on both host and target for each individual cable, to ensure that the cards properly link up as x8 or x16.




Figure 1: Full length PCIe bracket

Cable strain relief

Strain-relief on the PCIe cables are important to protect the PCIe card from fatigue loading and damage. The cable should be mechanically connected to the rack or cabinet. This is especially important for long heavy cables and cables going between cabinets and racks.

Supported Cables

The MXH930 supports both copper and PCIe active optical cables (AOC) available from Dolphin. Maximum lengths between two MXH930 cards can be found in Table 2 below. The max distance may change when connecting to other PCIe products.

Cable	Speed	Distance
Copper PCIe 4.0 cable	Gen4	4 meters. (5 meters being qualified)
PCIe Active Optical cables (AOC)	Gen4	Up to 100 meters

Table 2: Cable Specifications

Disconnecting the Cable

Please carefully pull the release tab to release the cable from the locking latches and gently pull the cable out of the connector guides.

Step 7 - Installing the Dolphin eXpressWare Software

The MXH930 adapter card is compliant with Dolphin's eXpressWare software package for the MX product line. Please visit <http://www.dolphinics.com/mx> to register and download the latest documentation and software. It is recommended to follow the quick installation guide found on the web page above to install the software for your operating system.

TIP: Dolphin provides software and documentation for several product families; **please remember to select the MX product family before downloading.** The MXH930 requires Dolphin software version DIS 5.20.0 or higher to operate. The software download requires a password to log in. **The password will automatically be emailed to you if you follow the instructions found on the getting started document bundled with the adapter card.** If you fail to provide the correct software serial number found on the getting started document, your request will be managed manually.

Note: Please consult the eXpressWare release note for detailed information on available software functionality. Please contact Dolphin for more information.

Step 7 – Verify Installation & LEDs

The MXH930 comes with 4 bi-color LEDs which show the corresponding cable port status according to Table 3: LED below.

The LEDs are visible through cut-outs in the PCIe bracket on each side of the cable connector block.

LED color	Function
Off	No cable installed
Yellow	Cable installed; no link
Yellow blink	Link reset
Green	Cable installed; link operational
Blinking green	Link operating at lower speed
Blinking Yellow	Temperature overheat warning, please immediately improve cooling to avoid thermal damages or system shutdown.

Table 3: LED behavior

Operation

Configuration and DIP Switches

The MXH930 has one bank of 8 DIP switches. The default factory setting for the MXH930 is x16 NTB link connection suitable for a two-node configuration or connection to the MXS824 switch.

The MXH930 has DIP switches for setting special modes or operations, the meaning of each DIP switch depends on the loaded firmware. Please carefully read the documentation shipping with the card before modifying any DIP switch settings. Please pay close attention to ON and OFF positions written on the DIP switch.

DIP Switch Bank – Configuration




Figure 2: DIP Switch shows the DIP switch for the MXH930. It is used to configure the adapter card. Please leave all undocumented DIP switches in the default position. Table 4: DIP Switch settings shows all the various DIP switch NTB settings for the MXH930.

Figure 2: DIP Switch

DIP no.	Name	Description	ON	OFF	Default
1-5	OPT1-5	Configuration selector, details below			Opt4: ON
6	OPT6	Thermal Shutdown Disable	Will disable thermal shutdown	Card will shut down on thermal overheat	OFF
7	URES	Holds the management processors in reset. Do NOT ACTIVATE without instructions from Dolphin.	Board management is held in reset	Normal operation	OFF
8	Reserved	Do NOT ACTIVATE without instructions from Dolphin.	Do not activate without instructions from Dolphin	Normal operation	OFF

Table 4: DIP Switch settings

Note: DIP switch configuration options may be changed in the future versions. Please always consult the latest user guide for details. The table above is valid for BMC firmware version 2.8 and newer.

Thermal shutdown

The MXH930 card includes overheat protection. The BMC will automatically shut down the MXH930 card if a PCIe switch temperature beyond 105°C (221°F) is detected. Similarly, if an AOC is plugged in, the MXH930 card will automatically shut down if the max operating temperature announced by the connected AOC is violated by 15°C (59°F). Please consult the AOC data sheet for details.

NOTE: The MXH930 link LEDs will start flashing yellow when a PCIe switch temperature of 95°C (203°F) is reached or when the AOC has reached the max operating temperature. Please consult the AOC data sheet for details.







NOTE: The automatic shutdown can be disabled by setting DIP-Switch OPT6. Prior to disabling the automatic shutdown, please ensure you have established proper airflow. Operating the card or AOC above the maximum specified operating temperature will void the warranty.

NOTE: If the card shuts down due to an overtemperature situation – you must improve the cooling before completely power cycling the system.

DIP switch NTB Settings

NOTE: The following configurations are available with BMC firmware 2.8 and newer. Please consult the Firmware Release Note for details.

DIP Switch NTB settings for MXH930

MXH930 DIP settings	DIP switch view – 256MB BAR	32GB BAR Option <i>All systems may not support this.</i>
Single x16 port	 <p>Shipping Default, Opt 4 set</p>	 <p>Opt 1 and 4 set</p>
Dual x8 ports	 <p>Opt 2 and Opt 4 set</p>	 <p>Opt 1, 2 and 4 set</p>
Quad x4 ports	 <p>Opt 3 and opt 4 set</p>	 <p>Opt 1, 3 and 4 set</p>

NTB Use Cases Summary and DIP Settings

The MXH930 card may be used in the following NTB use cases. Please contact Dolphin if you have other requirements for BAR size.

Use case	Description	BAR SIZE	DIP switch setting
A	Two node NTB configuration. Both systems have MXH930 installed.	256 Megabytes	Opt4: ON
		32 Gigabytes	Opt4: ON Opt1: ON
B	Three node NTB configuration. All systems have a MXH930 installed	256 Megabytes	Opt4: ON Opt2: ON
		32 Gigabytes	Opt4: ON Opt2: ON Opt1: ON
C	Five node NTB configuration. All systems have a MXH930 installed	256 Megabytes	Opt4: ON Opt3: ON
		32 Gigabytes	Opt4: ON Opt3: ON Opt1: ON
D	Multi node with MXS924 switch (same DIP switch setting as use casa A)	256 Megabytes	Opt4: ON
		32 Gigabytes	Opt4: ON Opt1: ON

Table 5: Use Case DIP switch settings

Use Case A - 2 Node Configuration – NTB Mode

Each node has an MXH930 adapter and a direct x4, x8 or x16 cable connection to the other system. This configuration is fully supported by all Dolphin software, from version DIS 5.17.0 and newer. Both MXH930 adapter cards operate in NTB mode. Dip switches should be set as specified in Table 4: DIP Switch settings on page 11.


Connecting the cables for 2 Node Configurations, x16 link

To establish an x16 link, a given port number should be connected to the same port number on the other card.

Table 6: Required x16 cabling. Always connect a cable from Port #x to Port #x


Port, Host A	Port, Host B
P1	P1
P2	P2
P3	P3
P4	P4

A failure connecting any of the cables will cause the link to re-train to x8 or x4.



Use Case B - 3 Node Configuration – NTB Mode

Each node has a MXH930 adapter and a direct x4 or x8 cable connection is used between all systems. This configuration is fully supported by all Dolphin software, from version DIS 6.0.0 and newer. All MXH930 adapter cards operate in NTB mode. Dip switches should be set as specified in Table 4: DIP Switch settings on page 11.




Connecting the cables for 3 Node Configurations, x8 link

Connect the cables between two nodes as described above, connecting two nodes x8. Connect the third node to the two other cards using similar rules. Always connect P1 + P2 to P1 + P2 or P3 + P4

Use Case C - 5 Node Configuration – NTB Mode

Each node has a MXH930 adapter, and a direct x4 cable connection is used between all systems. This configuration is fully supported by all Dolphin software, from version DIS 6.0.0 and newer. All MXH930 adapter cards operate in NTB mode. Dip switches should be set as specified in Table 4: DIP Switch settings on page 11.



Connecting the cables for 5 Node Configurations, x4 link

Connect any cable/port to any port on the remote system ensuring there is a single x4 cable directly between all nodes.

Use Case D – MXS924 Switch Configurations – NTB Mode

The MXH930 card can be combined with the MXS924 24 port PCIe switch to create larger PCIe networks. Each node typically has a MXH930 adapter and a x4, x8 or x16 cable connection to the MXS924 switch. All MXH930 adapter cards operate in NTB mode. Please consult the MXS924 Users Guide for more details to configure the MXS924. Dip switches should be set as specified in Table 4: DIP Switch settings on page 11. Some configuration examples can be found below, please contact Dolphin for more information on how to build larger systems up to 60 nodes.

6 Node, x16 Link Width using MXS924

In this use case, each MXH930 card is connected to the switch using four x4 cables to create a x16 Gen4 link. The aggregated link speed is 256 GT/s. The graphics in Figure 3 below shows the detailed cabling for connecting the card to the MXS924 switch. The MXS924 is pre-configured to support this setup.




Figure 3: A 6 x16 network

12 Node, x8 Link Width using MXS924

In this use case, each MXH930 card is connected to the switch using two cables to create a x8 Gen4 link. The aggregated link speed is 128 GT/s. The graphics in Figure 4 below shows the detailed cabling for connecting the card to the MXS924 switch. The MXS924 must be re-configured through the Ethernet interface to support this setup. Please consult the MXS924 Users Guide for details.

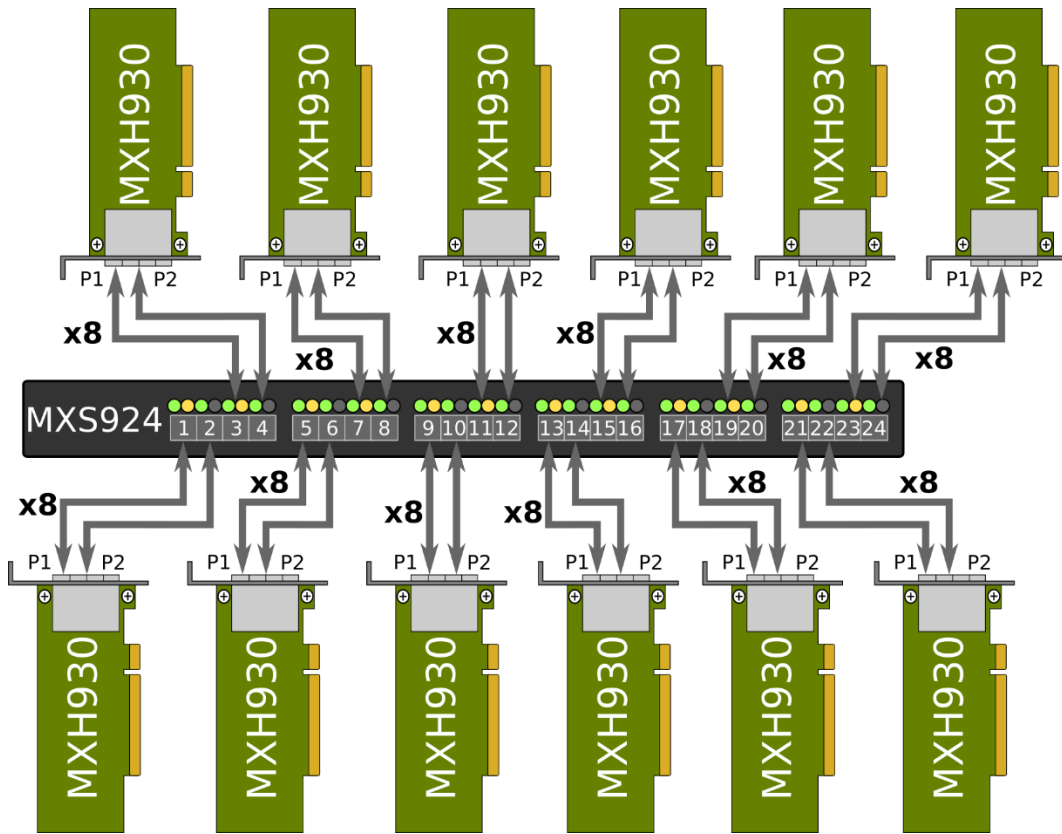


Figure 4: A 12 node PCIe x8 network using MXS924

24 Node, x4 Link Width using MXS924

In this use case, each MXH930 card is connected to the switch using one x4 cable to create a x4 Gen4 link. The aggregated link speed is 64 GT/s. The MXS924 must be re-configured through the Ethernet interface to support this setup. Please consult the MXS924 Users Guide for details.

Larger PCIe networks using the MXS924

Please consult the MXS924 Users Guide or contact Dolphin for information on larger topologies.




Figure 5: Example of larger MXS924 topology

Firmware Upgrade

The MXH930 design uses a microcontroller to implement the PCIe CMI protocol and other maintenance functions. Dolphin may from time to time publish updated firmware for the microcontroller, PCIe switch or EEPROM data for the card.

WARNING: Please note that standard Microsemi tools (ChipLink) cannot be used to upgrade the firmware as this will violate the warranty. Please contact Dolphin for instructions on how to upgrade the MXH930.

Note: Please consult the MXH930 firmware release note for details on supported configurations. This document is compliant to the BMC version 2.8.

Identifying the Card

The card has a label-sticker with the serial number in the format 'MXH930-YY-ZZZZZ', where YY denotes the card revision (e.g. BB) and ZZZZZ denotes the serialized production number (e.g. 012345) – this whole string makes up the serial number of the card (i.e. MXH930-CF-012345).

You can also get this information using `lspci` in Linux:

First, identify the devices for the Dolphin Host card:

```
# lspci | grep "Device 4036"

09:00.0 PCI bridge: PMC-Sierra Inc. Device 4036
09:00.1 Bridge: PMC-Sierra Inc. Device 4036
0a:00.0 PCI bridge: PMC-Sierra Inc. Device 4036
```

Then run `lspci` and identify the card. It will show up as something like

```
# lspci -s 9:0.0 -v
09:00.0 PCI bridge: PMC-Sierra Inc. Device 4036 (prog-if 00 [Normal decode])
  Flags: bus master, fast devsel, latency 0
  Bus: primary=01, secondary=02, subordinate=03, sec-latency=0
  I/O behind bridge: 00002000-00002fff
  Capabilities: [40] Express Upstream Port, MSI 00
  Capabilities: [7c] MSI: Enable- Count=1/8 Maskable- 64bit+
  Capabilities: [8c] Power Management version 3
  Capabilities: [94] Subsystem: Dolphin Interconnect Solutions AS Device 0930
  Capabilities: [100] Advanced Error Reporting
  Capabilities: [148] Power Budgeting <?>
  Capabilities: [158] #12
  Capabilities: [188] #19
  Capabilities: [1b4] Device Serial Number 00-00-43-43-00-00-00-21
  Capabilities: [1c0] Latency Tolerance Reporting
  Capabilities: [1c8] Access Control Services
  Capabilities: [1f0] #25
  Capabilities: [1fc] #26
  Capabilities: [23c] #27
  Capabilities: [7f8] Vendor Specific Information: ID=ffff Rev=1 Len=808 <
  Kernel driver in use: pcieport
  Kernel modules: shpchp
```

Second, do

```
# lspci -s 9:0.0 -v | grep -E "Subsystem|Serial"  
Capabilities: [a4] Subsystem: Dolphin Interconnect Solutions AS Device 0930  
Capabilities: [100] Device Serial Number 00-00-42-42-00-00-00-ff
```

This shows the card as revision 0x4242 (hexadecimal values of the 'BB' letters in the ASCII table), with the production number 0x000000ff (00000255 in decimal).

Support

More information about the product, support and software download can be found at <http://www.dolphinics.com/mx>.
For general support questions, please contact Dolphin via the Jira Service Management portal:
<https://www.dolphinics.com/csp>.

Technical Information

Board revision history

The following table gives a general overview of the hardware revision history. Please consult the MXH930 PCN available from Dolphin for details.

Adapter card revision	Capabilities
MXH930-CC	<ul style="list-style-type: none"> Initial product version
MXH930-CE	<ul style="list-style-type: none"> Adds support for AOC
MXH930-DD	<ul style="list-style-type: none"> New SFF-8644 connector due to EOL Replace PCIe oscillator with lower jitter variant (Si52254A)
MXH930-DE	<ul style="list-style-type: none"> Replaced AOC current limiter due to long lead time

Simplified schematics

The figure below shows the MXH930 simplified schematics. Only the main functions are shown.

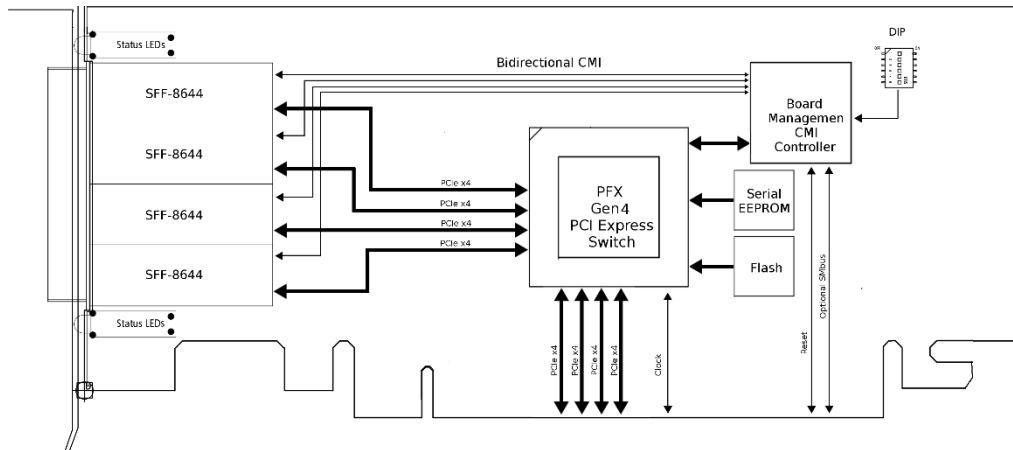


Figure 6: MXH930 Simplified Schematics

PCIe Cable Port Signals

The external PCI Express SFF-8644 cable connector supports the following signals:

- PETpN/PETnN: PCI Express Transmitter pairs, labeled where N is the Lane number (starting with 0); “p” is the true signal while “n” is the complement signal.
- PERpN/PERnN: PCI Express Receiver pairs, labeled where N is the Lane number (starting with 0); “p” is the true signal while “n” is the complement signal.
- PWR: Power to support AOC and signal conditioning components within the cable assembly.
- MGTPWR: Power supplied to the connector for cable management components that are needed while the link is not active. This needs to be active if the subsystem has power.
- CBLPRSNT#: Cable present detect, an active-low signal pulled-down by the cable when it is inserted into the MXH930 connector.
- CADDR: Signal used to configure the upstream cable management device address.
- CINT#: Signal asserted by the cable assembly to indicate a need for service via the CMI controller.
- CMISDA: Management interface data line. Used for both initial link setup and sideband messages when used with CMI compliant cables.
- CMISCL: Management interface clock line. Used for both initial link setup and sideband messages when used with CMI compliant cables.

External PCIe x4 Cable Connector Pin-Out

	Column								
Row	9	8	7	6	5	4	3	2	1
D	GND	PETn2	PETp2	GND	PETn1	PETp1	GND	MGTPWR	PWR
C	GND	PETn3	PETp3	GND	PETn0	PETp0	GND	CMISDA	CMISCL
B	GND	PERn2	PERp2	GND	PERn1	PERp1	GND	CBLPRSNT#	PWR
A	GND	PERn3	PERp3	GND	PERn0	PERp0	GND	CINT#	CADDR

Table 7; External PCIe x4 cable Pin-Out

PCIe Cable Port Mapping

The MXH930 card has a quad SFF-8644 connector. The table below shows the signal / port map.

Cable Port	x16	Dual x8	Quad x4	PCIe 3.0 Cable Pin
1	L15	L7	L3	TX3/RX3
	L14	L6	L2	TX2/RX2
	L13	L5	L1	TX1/RX1
	L12	L4	L0	TX0/RX0
2	L11	L3	L3	TX3/RX3
	L10	L2	L2	TX2/RX2
	L9	L1	L1	TX1/RX1
	L8	L0	L0	TX0/RX0
3	L7	L7	L3	TX3/RX3
	L6	L6	L2	TX2/RX2
	L5	L5	L1	TX1/RX1
	L4	L4	L0	TX0/RX0
4	L3	L3	L3	TX3/RX3
	L2	L2	L2	TX2/RX2
	L1	L1	L1	TX1/RX1
	L0	L0	L0	TX0/RX0
Lx – PCIe lane X, Cable port is ref PCIe bracket marking				

Table 8 : PCIe Cable Port Mapping

Compliance and Regulatory Testing

EMC Compliance

The Dolphin PCI Express MXH930 adapter is tested to PASS the following relevant test standards for PCI Express cards, telecommunication and industry equipment installed in a standard PC:

- EN 55032:2012,
- EN 55035:2017,
- EN 61000-3-2:2014,
- EN 61000-3-3:2013
- 47 CFR Part 15, Subpart B (Clause 15.107 and 15.109) in conjunction with ANSI C63.4:2014
- CISPR 35:2016 Edition 1.0 (CISPR/I/412/CDV) Korean Harmonized standard, KN 35



This does not ensure that it will comply with these standards in any random PC. It is the responsibility of the integrator to ensure that their products are compliant with all regulations where their product will be used.

FCC Class A

This equipment is tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules.



RoHS Compliance

The Dolphin MXH930 is RoHS compliant. A Compliance certificate issued by the manufacturer is available upon request.



WEEE Notice

The adapter card is labelled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to products to indicate that the product is not to be thrown away but returned to your local approved WEEE waste collector.



Limited Warranty

Dolphin Interconnect Solutions warrants this product to be free from manufacturing defects under the following terms:

Warranty Period

The warranty applies for one (1) year from the date of purchase. Extended warranty is available.

Coverage

To the extent permitted by applicable law, this warranty does not apply to:

- Damage caused by operator error or non-compliance with instructions available for the product.
- Use or attempt to use or program firmware not approved by Dolphin.
- Damage due to accidents, abuse, misuse, improper handling or installation, moisture, corrosive environments, high voltage surges, shipping, or abnormal working conditions.
- Damage resulting from violation of the specified operating or storage temperatures and airflow.
- Damage caused by acts of nature, e.g. floods, storms, fire, or earthquakes.
- Damage caused by any power source out of range or not provided with the product.
- Normal wear and tear.
- Attempts to repair, modify, open, or upgrade the product by personnel or agents not authorized by Dolphin.
- Products for which the serial number label has been tampered with or removed.
- Damage to the product caused by products not supplied by Dolphin.

Service Procedure

In the event that the product proves defective during the Warranty Period, you should contact the seller that supplied you with the product, or if you purchased it directly from Dolphin, visit <https://www.dolphinics.com/csp> to obtain a valid RMA number and instructions. Products returned to Dolphin without a proper RMA number will not be serviced under this warranty.