VPX PCI-Express Interface

P1/P2/P5 Port Configurations

Document #:	**1000-6001-V1.00**
Prepared By:	Terry Laviolette
Approved By:	
Version #:	**V104**
Date:	**18-May-2017**

DOCUMENT HISTORY

Revision	Date	Description of Changes	By
V100	09-Jan-2017	Initial Draft	TL
V101	11-Jan-2017	Added scenarios, text adjustments	TL
V102	12-Jan-2017	Feedback corrections	TL
V103	25-Apr-2017	Added Slot “Type” considerations	TL
V104	18-May-2017	Prepare for public release	SB

TABLE OF CONTENTS

TABLE OF FIGURES

1 PREAMBLE

1.1 WOLF ADVANCED TECHNOLOGY

WOLF Advanced Technology is a North American engineering design center, specializing in the rapid development and manufacture of real time video capture and processing systems targeting rugged, embedded, low power, small form factor solutions for the harshest military and aerospace environments.

WOLF's industry-proven modular technology solutions integrate high resolution graphics capture, process, encode and display technologies for rapidly available COTS solutions, feature tailored MCOTS solutions or full custom design systems. All solutions are designed to operate in harsh environmental conditions using the latest generation high-speed digital signal processing technologies while conforming to demanding embedded systems size, weight and power (SWaP) constraints.

All of ***WOLF's*** core technology IP can be design-targeted to a variety of architectures such as VPX, XMC, PMC, COMExpress, VME and Vita 65 including support for digital and legacy analog video standards such as DisplayPort, 3G/HD-SDI, TMDS, LVDS, RGBHV, STANAG 3350 and NTSC/PAL/SECAM.

The "***WOLF Pack***" strategy creates a unique team environment geared toward collaborating with our clients to create a high-performance engineering attack team. This business model provides OEM clients with the resources for complete or selective engineering and manufacturing services including: Custom board design, PCB manufacture, system integration, software driver development, mechanical design, technical support, procurement, fabrication, testing, certifications, sales, graphic design and technical writing.

2 OVERVIEW

WOLF's VPX 3U/6U modular solutions have been designed to facilitate rapid module integration into today's highly configurable rugged system designs. WOLF's VPX solutions combine advanced technology graphics and compute modules, legacy analog and high speed digital capture, process and display, and VPX interface configurability supporting multiple host interface options

This document describes WOLF's VPX PCI Express port configuration options as it applies to both the VPX3U and VPX6U system architectures.

In aggregate, WOLF's PCI Express port configuration options can be grouped as in Figure 1 below.

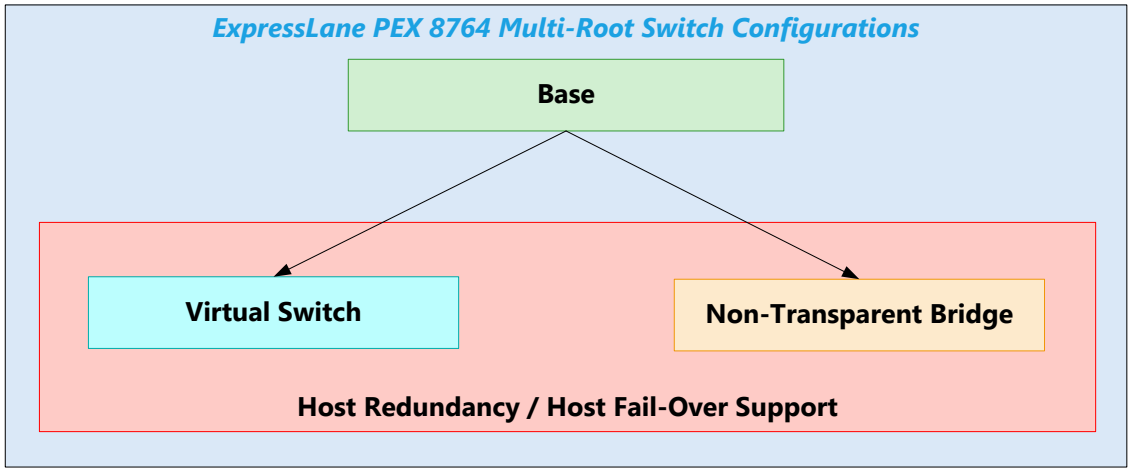

Figure 1: Multi-Root Switch Configurations

Several Host to Endpoint configuration options are possible however fundamentally all switch modes in the Base level act as a standard PCI Express switch, supporting a one Host hierarchy.

For more advanced Host to Endpoint designs, several Virtual Switch configurations are possible each supporting enhanced feature support such as Non-Transparent Bridging, Host Redundancy and Host Fail-Over.

All PCI Express port options as described in this document are factory configured based on the specific client configuration established through WOLF's product part number information in discussion with the customer.

Specific signal definitions and pin mapping designations are described in Appendix 'A'.

3 PORT CONFIGURATIONS

3.1 CONSIDERATIONS

3.1.1 VPX6U Specific

The VPX6U interface is limited to 2 PCI Express upstream interfaces to the host processor configured as 1 interface on P2 and 1 interface on P5

WOLF VPX6U PCI Express interface uses the Expansion plane of the P2 and P5 connectors.

All WOLF VPX6U products which include an integrated PEX Switch [such as the PLX 8764 referenced in this document] can be installed into OpenVPX Payload slots of either slot Type ‘A’, Type ‘B’ or Type ‘C’.

Figure 2 depicts various Slot Type configurations of a typical Host SBC and WOLF Carrier module chassis installation. For these combinations, the WOLF Carrier is assumed to be factory configured to either a x16 or x8 uplink on P2. Given this it can be observed that certain Slot Type combinations will having varying compatibility results.

WOLF VPX6U carriers support multiple factory-programmable configurations including…

- x16 P2 Up, x16 P5 Non-Transparent Bridge
- x16 P2 Up, x16 P5 Down
- x8 P2 Up, x8 P2 Down, NC on P5
- x16 P2 Up, NC on P5
- x8 P2, NC on P5

Given that multiple factory configurations are possible; consultation with WOLF Technical Support is recommended to ensure optimal Carrier PEX Switch and Host SBC / Chassis compatibility.

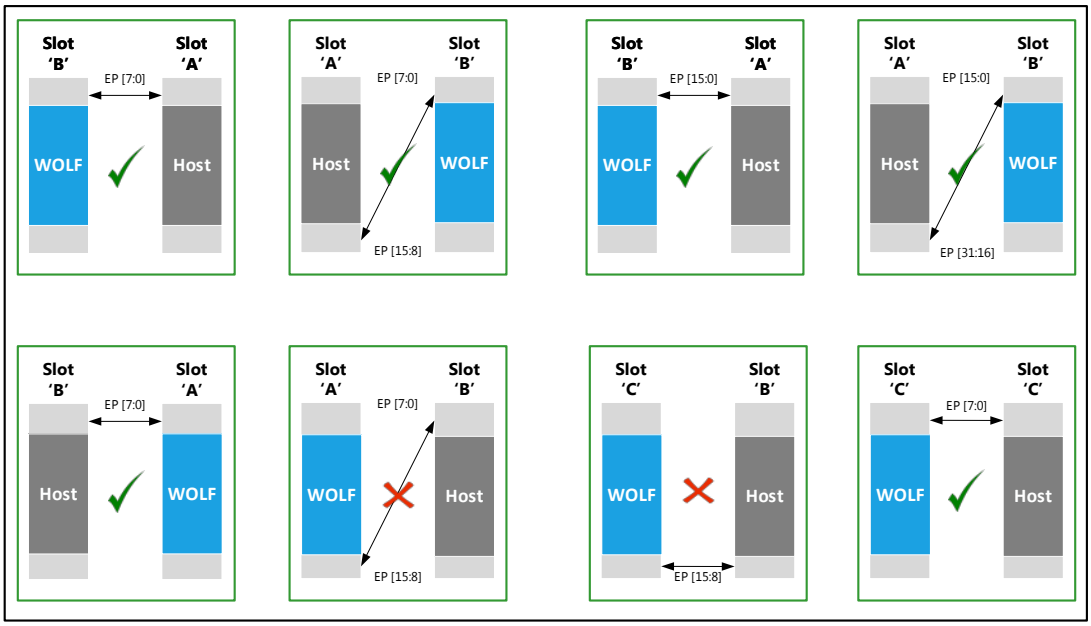

Figure 2: P2 Quad-Pipe Slot Type Configuration Compatibility

3.1.2 VPX3U Specific

VPX3U Carrier modules can be configured to x4, x8 or x16 lanes. The PCI Express bus can only operate as an end-point and must be connected to a root complex.

Whether the PCI Express bus operates in Gen2 or Gen3 mode depends on the rating of the MXM module that is integrated.

P1 Configuration	Diagram	Description
x4 PCI Express	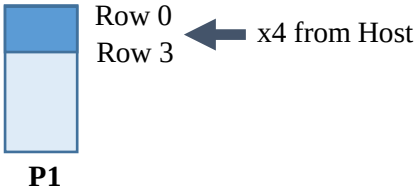	Only the first four rows of P1 are used when configured for x4 PCI Express mode.
x8 PCI Express	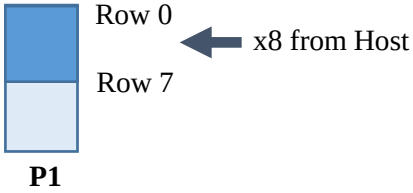	Only the first eight rows of P1 are used when configured for x8 PCI Express mode. Warning: Using this mode in a system with x4 backplane mapping can have undesirable effect.
x16 PCI Express	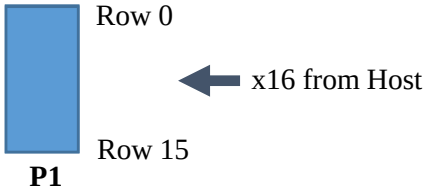	All rows of the P1 connector are used when configured for x16 PCI Express mode. Note: This mode is extremely rare and must be supported by your system and SBC.

Table 1: VPX3U P1 PCI Plane Usage Options

The VPX3U P1 interface is limited to 1 PCI Express upstream interface to the host processor.

WOLF VPX3U PCI Express interface uses the Data plane of the P1 connector.

All WOLF VPX3U products which include an integrated PEX Switch can be installed into OpenVPX Payload slots of either slot Type ‘A’, Type ‘B’ or Type ‘C’. It is expected that typically a host SBC will be installed into a Type ‘A’ slot of the chassis with the WOLF carrier installed into a directly neighboring Type ‘B’ slot. Consultation with WOLF Technical Support is recommended to ensure optimal compatibility.

All WOLF VPX3U products which do not include an integrated PEX Switch must only be installed into a slot Type ‘A’, Type ‘B or Type ‘C’ which supports Host DP[n:0] to WOLF DP[n:0]. Configurations that map Host DP[n:0] to WOLF DP[7:4] as an example are not compatible. WOLF modules must be mapped such that WOLF lane 0 of the PEX interface connects with the lowest lane of the Host PEX interface.

3.1.3 Common

PCI Express lane reversing and/or polarity swapping is supported in accordance with the PCI specification.

All downstream PCI Express ports on the WOLF module are designed to support X16 lanes of GEN3 PCI Express unless otherwise stated.

For cases where ANSI/VITA 65 OpenVPX compliance is a consideration, it should be noted that the specification is defined to X8 (or less) upstream lanes (VPX3U) at GEN1/GEN2 data rates although WOLF's design can support X16 upstream lanes at GEN3 data rates with the appropriate system design considerations.

Field-based reconfiguration of the PCI Express bridge port definition is not supported however WOLF can accommodate lab-based reconfiguration if needed by the customer.

3.2 BASE CONFIGURATIONS

Each of the VPX port configuration options as described in the Base configuration grouping assumes a single host interface in either of VPX3U or VPX6U.

See the tables in the Appendix for pin mapping details.

3.2.1 Base 1A

The Base 1A configuration supports a single upstream X8 lane PCI Express link to a single host CPU.

Figure 1 references an example VPX6U implementation. It should be noted that while this interface can support from X1 to X8 lanes; it will not support P2 interfaces that are configured for 2 upstream links. It is important that the link interface to P2 be a single, dedicated link to the Host. This same constraint applies to VPX3U configurations using the P1 interface.

OpenVPX ANSI/VITA 65 3U Slot Profile Support: SLT3-PAY-1D-14.2.6, SLT3-PAY-2F-14.2.7, SLT3-PAY-1F4U-14.2.8, SLT3-PAY-8U-14.2.9
OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

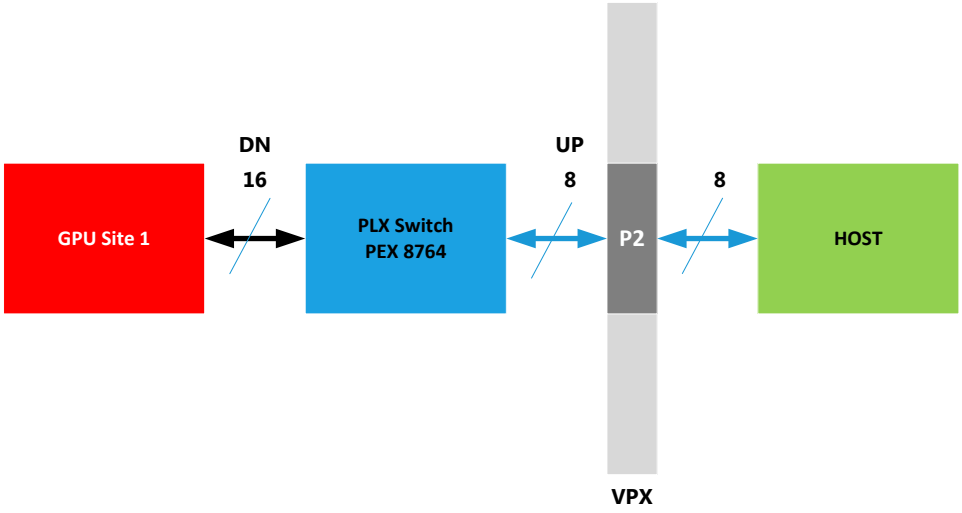

Figure 3: Base 1A

3.2.2 Base 1B

The Base 1B configuration supports a single upstream X16 lane PCI Express link to a single host CPU.

Figure 2 references the VPX6U implementation. It should be noted that while this interface can support from X1 to X16 lanes; it will not support P2 interfaces that are configured for 2 upstream links. It is important that the link interface to P2 be a single, dedicated link to the Host. This same constraint applies to VPX3U configurations using the P1 interface.

OpenVPX ANSI/VITA 65 3U Slot Profile Support: SLT3-PAY-1D-14.2.6, SLT3-PAY-2F-14.2.7, SLT3-PAY-1F4U-14.2.8, SLT3-PAY-8U-14.2.9
OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

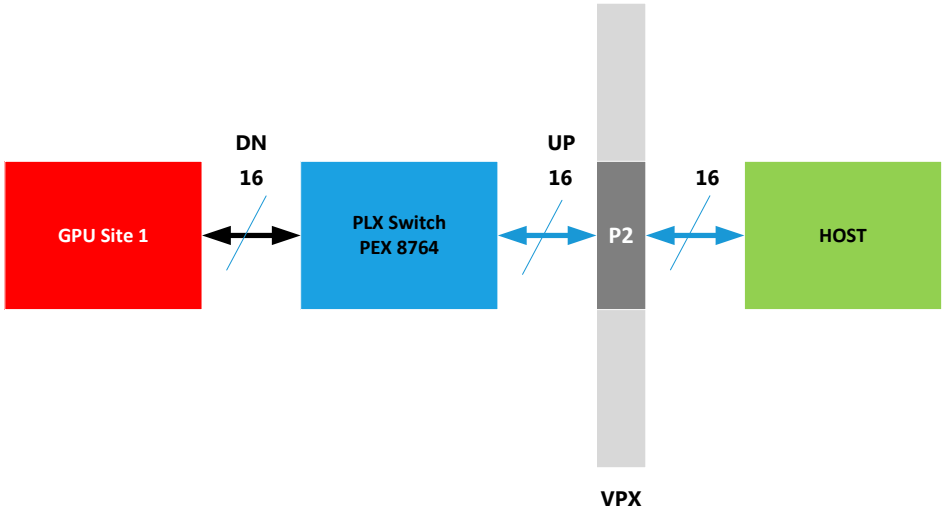

Figure 4: Base 1B

3.2.3 Base 1C

The Base 1C configuration supports a single upstream X16 lane PCI Express link to a single host CPU supporting 2 GPU sites through dedicated X16 GEN3 (system dependent) PCI Down Stream links from the PLX switch.

Figure 4 references the VPX6U implementation. It should be noted that while this interface can support from X1 to X8 lanes; it will not support P2 interfaces that are configured for 2 upstream links. It is important that the link interface to P2 be a single, dedicated link to the Host. This same constraint applies to VPX3U configurations using the P1 interface.

OpenVPX ANSI/VITA 65 3U Slot Profile Support: SLT3-PAY-1D-14.2.6, SLT3-PAY-2F-14.2.7, SLT3-PAY-1F4U-14.2.8, SLT3-PAY-8U-14.2.9

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

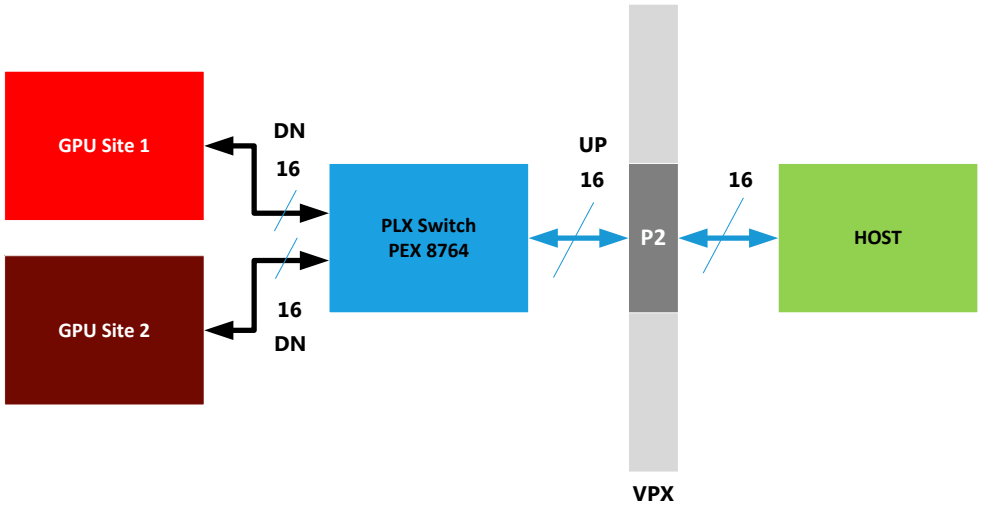

Figure 5: Base 1C

3.2.4 DAISY CHAIN

3.2.4.1 Chain 1A

The Daisy Chain 1A configuration supports a single upstream X8 lane PCI Express link to a single host CPU supporting 1 GPU sites through a dedicated X16 GEN3 lanes PCI downlink from the PLX switch. This configuration also supports a X8 downlink to P2 facilitating host daisy chaining to a neighboring module.

Figure 5 references the VPX6U implementation.

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

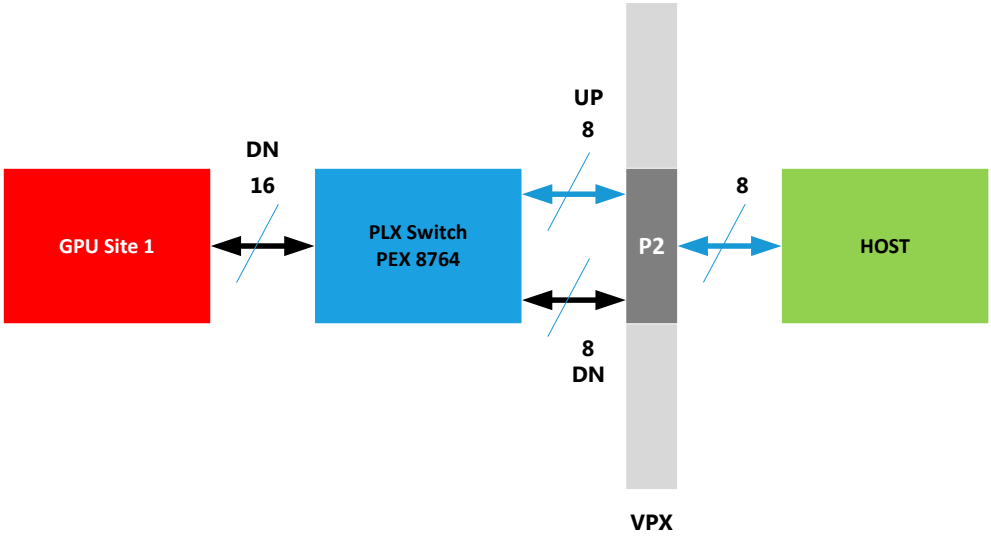

Figure 6: Chain 1A

3.2.4.2 Chain 1B

The Daisy Chain 1B configuration supports a single upstream X16 lane PCI Express link to a single host CPU supporting 1 GPU sites through a dedicated X16 GEN3 lanes PCI downlink from the PLX switch. This configuration also supports a X16 downlink to P5 facilitating host daisy chaining to a neighboring module.

Figure 6 references the VPX6U implementation.

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

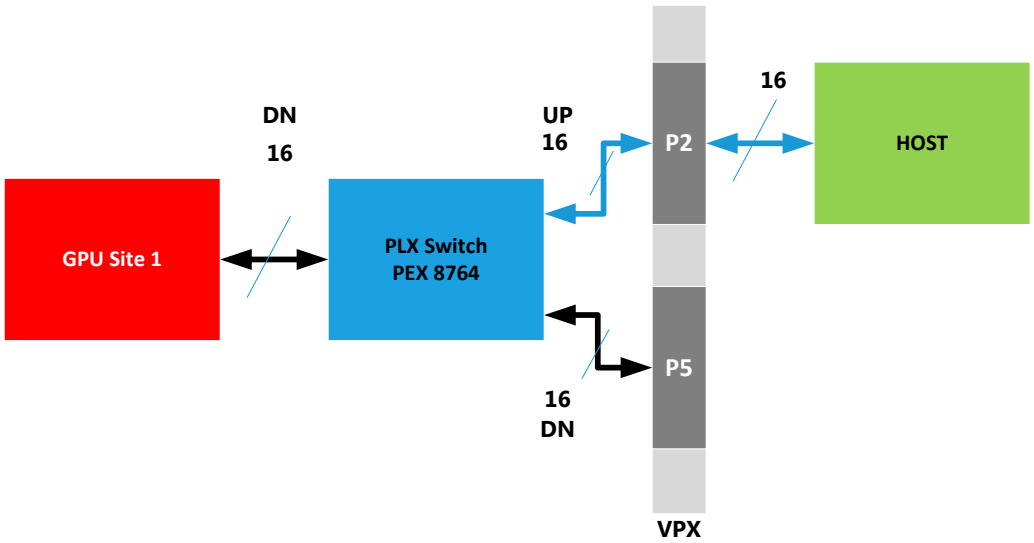

Figure 7: Chain 1B

3.2.4.3 Chain 1C

The Base 1F configuration supports a single upstream X16 lane PCI Express link to a single host CPU supporting 1 GPU sites through a dedicated X16 GEN3 lanes PCI downlink from the PLX switch. This configuration also supports a X16 downlink to P2 facilitating host daisy chaining to a neighboring module.

Figure 7 references the VPX6U implementation.

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

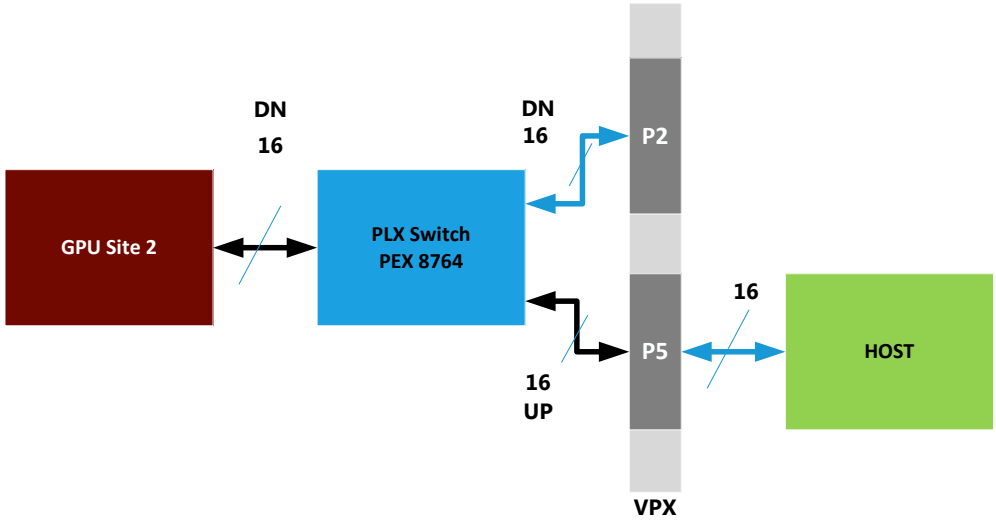

Figure 8: Chain 1C

3.3 VIRTUAL SWITCH

3.3.1 Dual Host 1A

The Dual Host 1A configuration supports 2 upstream X16 lane PCI Express link to 2 host CPUs supporting 1 GPU site through a dedicated X16 GEN3 lanes PCI downlink from the PLX switch. In this example Host 2 has no Downstream device. See *Virtual Switch Host Fail-Over Operation* and *Non-Transparent Bridge* for use case scenarios.

Figure 8 references the VPX6U implementation.

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

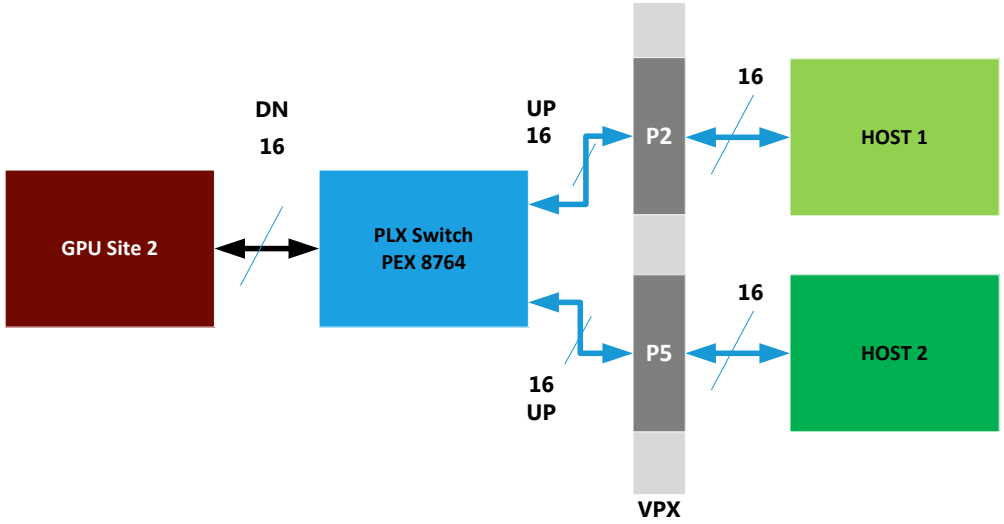

Figure 9: Dual Host 1A

3.3.2 Dual Host 1B

The Dual Host 1B configuration supports 2 upstream X16 lane PCI Express link to 2 host CPUs supporting 2 GPU sites through dedicated X16 GEN3 lanes PCI Down Stream links from the PLX switch. In this configuration, Host 1 and GPU 1 are completely isolated into 2 separate PCI Express address domains. No communication is shared across Host 1 domain and Host 2 domain.

Figure 9 references the VPX6U implementation.

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

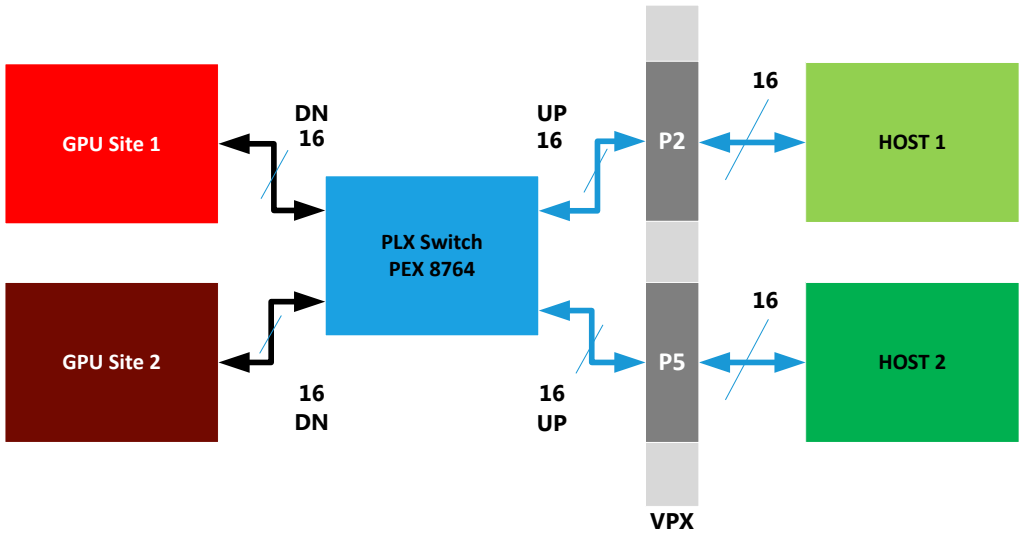

Figure 10: Dual Host 1B

3.3.3 Virtual Switch Host Fail-Over Operation

In this example, Host 1 assumes the role of the primary (Active) Host managing PCI Express communication via P2 of the VPX connector to the GPU Site 2. Host 2 is designated as the Redundant Host and has no downstream device via the PLX Switch.

A typical Virtual Switch Fail-Over software-driven program flow would consist of…

1. Redundant Host monitors the Active Host for "**heartbeat**" status through the Redundant Management Port. See Figure 10 *Virtual Switch Host Fail-Over "Before"*
2. If the "heartbeat" status fails, the Redundant Management Port is reassigned to be the Active Management Port. Conversely the previous Active Management Port is reassigned to be the Redundant Management Port. Both Ports are held in the Enable mode.
3. Active Management Port scans the Virtual Switch Upstream to Management Upstream **Doorbell Request** and Management Upstream to Virtual Switch Upstream Doorbell Request registers for any pending interrupts to service. See Figure 11 *Virtual Switch Host Fail-Over "After"*
4. Active Management Port informs other Hosts that a new Active Management Host has been promoted.

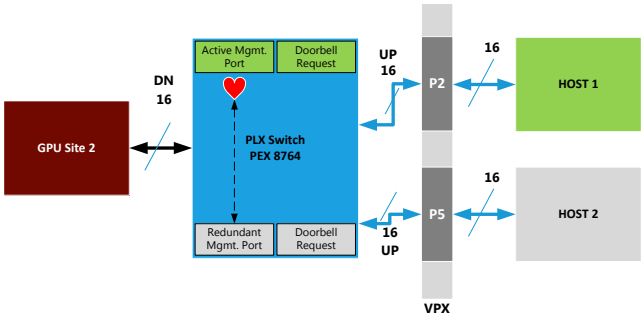

Figure 11: Virtual Switch Host Fail-Over "Before"

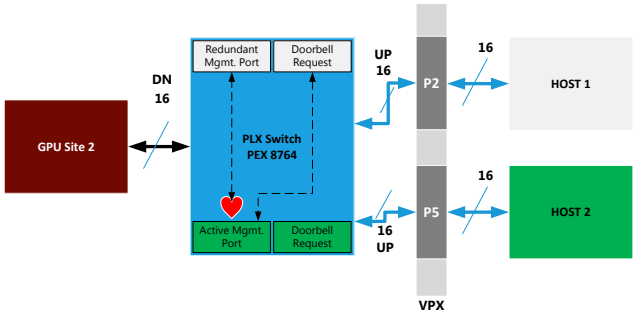

Figure 12: Virtual Switch Host Fail-Over "After"

Note that the above Fail-Over operation can also apply to the configuration where Host 2 (or Host 1) is operating in Non-Transparent Bride Mode. See *Non-Transparent Bridge* in this document.

3.4 NON-TRANSPARENT BRIDGE

The Dual Host configuration; Figure 12 supports 2 upstream X16 lane PCI Express link to 2 host CPUs supporting 2 GPU sites through dedicated X16 GEN3 lanes PCI Down Stream links from the PLX switch. In this configuration, Host 1 is managing both GPU sites with Host 2 operating in a Non-Transparent Bridge mode with no connect Downstream device. Host 2 and Host 1 have the capability to share information using the PEX 8764 on-board Doorbell and Scratchpad registers as monitored through a software-driven interface.

Figure 11 references the VPX6U implementation.

OpenVPX ANSI/VITA 65 6U Slot Profile Support: SLT6-PAY-4F1Q2U2T-10.2.1, SLT6-PAY-4F1Q2U2T-10.2.1-10.2.6, SLT6-PAY-4F2Q2U2T-10.2.7

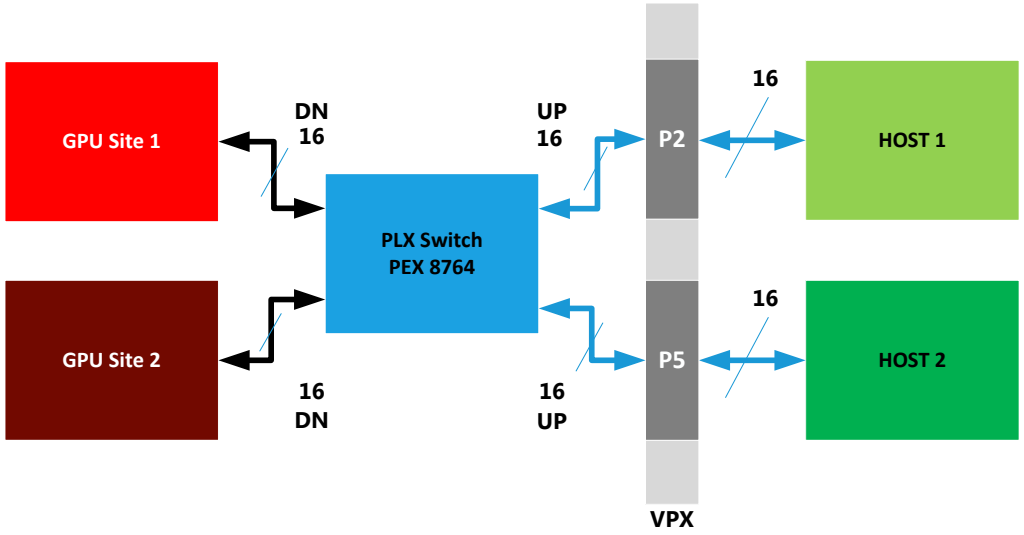

Figure 13: Non-Transparent Bridge

3.4.1 Non-Transparent Host Fail-Over Operation

Several configuration scenarios are possible for Non-Transparent Host Fail-Overs including [but not limited to] Active-Passive NT Fail-Over, Active-Active NT Fail-Over and NT Fail-Over for Virtual Switch.

In these scenarios “Active” is defined as the managing Host of specified endpoints on the PEX 8764 switch. Conversely, “Passive” defines a Host that is not managing a PEX 8764 endpoint but may be sharing information between Hosts using the PEX 8764 Doorbell and Scratchpad mechanisms.

As example, Figure 14 [before] and Figure 15 [after] describes an Active-Passive NT Fail-Over operation based on the configuration as described in *Non-Transparent Bridge* within this document.

A typical Fail-Over software-driven program flow would consist of…

1. Passive Host 2 monitors the Active Host 1 for “heartbeat” status through the Passive Management Port. Remaining operation for this step follows steps 2-4 of the *VS Host Fail-Over operation.*
2. Active Host 1 Port interface is redefined to be the NT Link Endpoint. Conversely, Passive Host 2 interface is redefined to be the NT Virtual Link Endpoint
3. The PEX 8734 is reset though the [new] Active Host 2 Upstream Port Bridge Control register.
4. Re-enumerate (optional in some systems) the Bridge hierarchy and resume operation

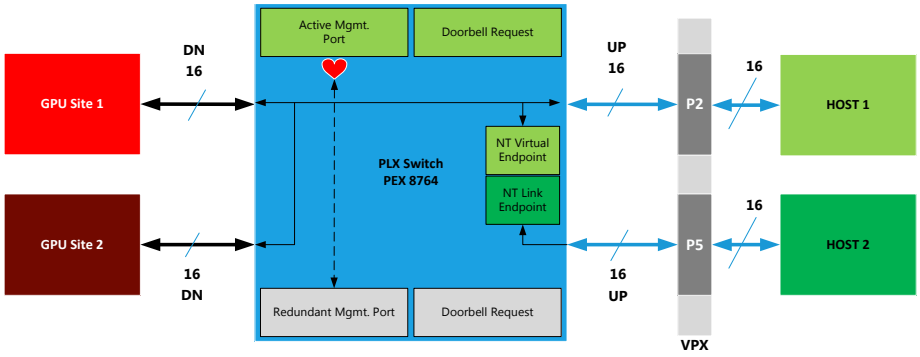

Figure 14: NT Host Fail-Over "Before"

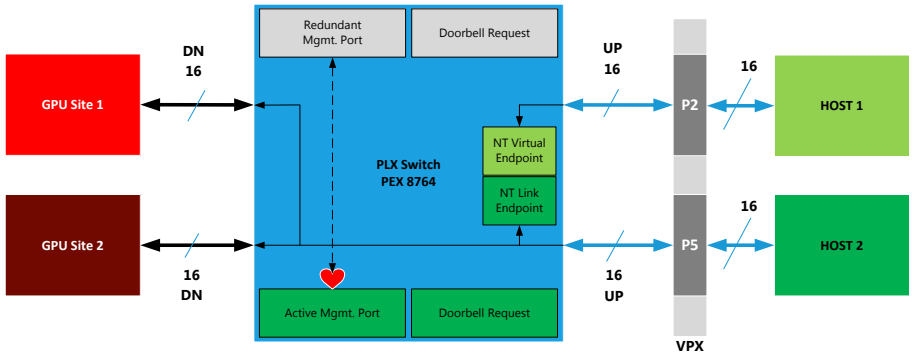

Figure 15: NT Host Fail-Over "After"

3.5 System Level Example Configuration – Virtual Switch Host Failover

This configuration as described below in Figure 15 *Dual Host 1C* highlights a system level solution that integrates 2 of WOLF's VPX6U dual GPU solutions with 2 system Hosts.

This same component combination could also be reconfigured to various other Virtual Switch, Non-Transparent Bridge and Fail-Over operating methods including…

- Dual Virtual Switch with each module of 2 GPUs managed by a dedicated Host
- Dual Virtual Switch with each module of 2 GPUs managed by a dedicated Host with Host Fail-Over
- Dual Virtual Switch with 4 GPUs (2 modules) managed by a single Host with Host Fail-Over
- Single Virtual Switch with 4 GPUs (2 modules) managed by a single Host with a 2nd Host in NTB
- Single Virtual Switch with 4 GPUs (2 modules) managed by a single Host with a 2nd Host in NTB and Fail-Over
- And more….

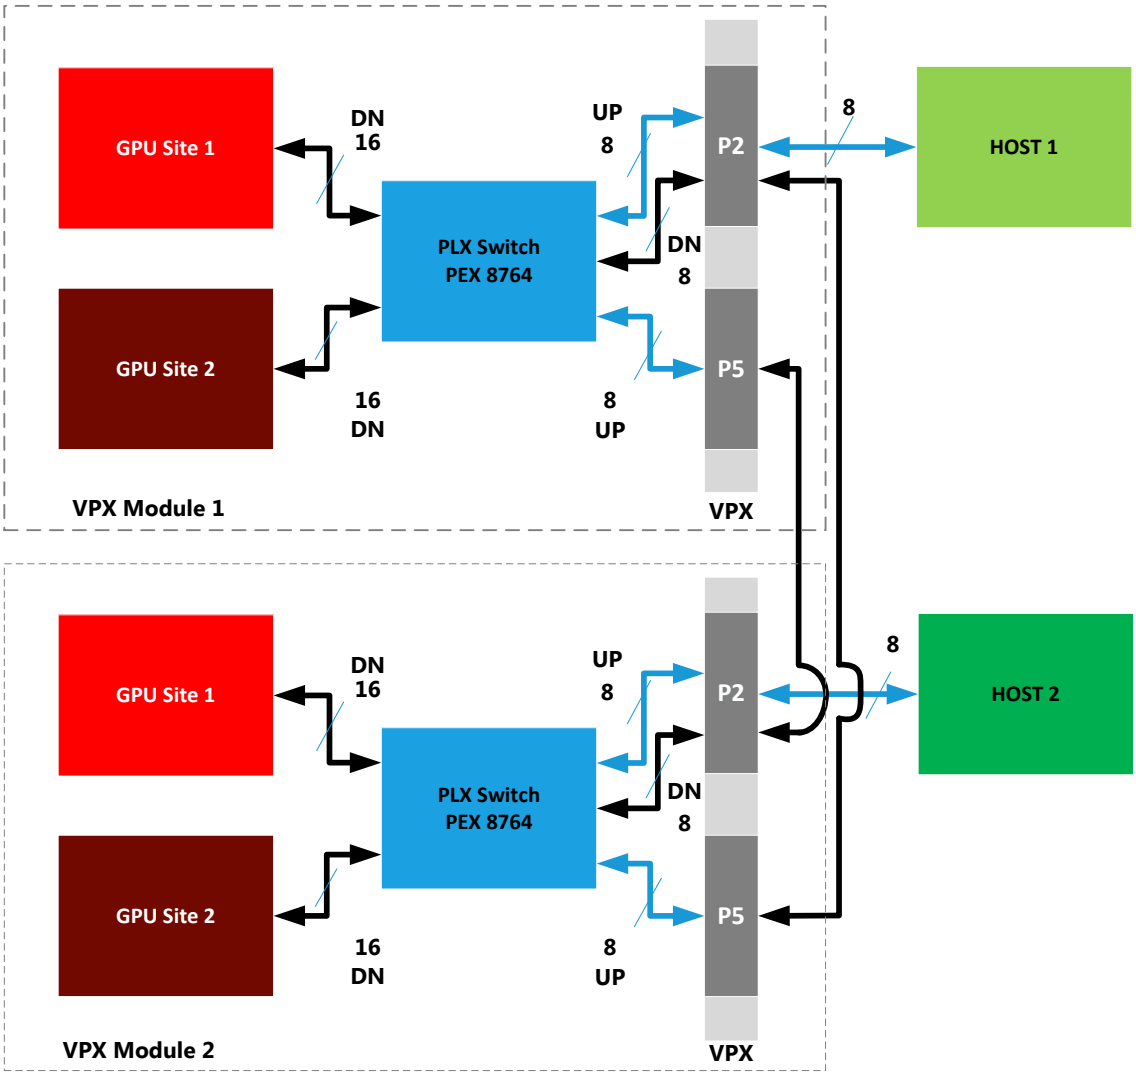

Figure 16: Dual Host 1C

4 APPENDIX

4.1 VPX6U P2

	Row G	Row F	Row E	Row D	Row C	Row B	Row A
1	PEX_RST1_L	GND	PEX_TX0_N	PEX_TX0_P	GND	PEX_RX0_N	PEX_RX0_P
2	GND	PEX_TX1_N	PEX_TX1_P	GND	PEX_RX1_N	PEX_RX1_P	GND
3	PEX_RST2_L	GND	PEX_TX2_N	PEX_TX2_P	GND	PEX_RX2_N	PEX_RX2_P
4	GND	PEX_TX3_N	PEX_TX3_P	GND	PEX_RX3_N	PEX_RX3_P	GND
5	PEX_CLK1_N	GND	PEX_TX4_N	PEX_TX4_P	GND	PEX_RX4_N	PEX_RX4_P
6	GND	PEX_TX5_N	PEX_TX5_P	GND	PEX_RX5_N	PEX_RX5_P	GND
7	PEX_CLK1_P	GND	PEX_TX6_N	PEX_TX6_P	GND	PEX_RX6_N	PEX_RX6_P
8	GND	PEX_TX7_N	PEX_TX7_P	GND	PEX_RX7_N	PEX_RX7_P	GND
9	PEX_CLK2_N	GND	PEX_TX8_N	PEX_TX8_P	GND	PEX_RX8_N	PEX_RX8_P
10	GND	PEX_TX9_N	PEX_TX9_P	GND	PEX_RX9_N	PEX_RX9_P	GND
11	PEX_CLK2_P	GND	PEX_TX10_N	PEX_TX10_P	GND	PEX_RX10_N	PEX_RX10_P
12	GND	PEX_TX11_N	PEX_TX11_P	GND	PEX_RX11_N	PEX_RX11_P	GND
13	PEX_REFCLK_SEL	GND	PEX_TX12_N	PEX_TX12_P	GND	PEX_RX12_N	PEX_RX12_P
14	GND	PEX_TX13_N	PEX_TX13_P	GND	PEX_RX13_N	PEX_RX13_P	GND
15	VPX_PEX_REFCLK_SEL	GND	PEX_TX14_N	PEX_TX14_P	GND	PEX_RX14_N	PEX_RX14_P
16	GND	PEX_TX15_N	PEX_TX15_P	GND	PEX_RX15_N	PEX_RX15_P	GND

Table 2: VPX6U P2 Pin Mapping

VPX6U P5

P5 - Differential wafers							
	Row G	Row F	Row E	Row D	Row C	Row B	Row A
1	NC	GND	PEX_TX16_N	PEX_TX16_P	GND	PEX_RX16_N	PEX_RX16_P
2	GND	PEX_TX17_N	PEX_TX17_P	GND	PEX_RX17_N	PEX_RX17_P	GND
3	NC	GND	PEX_TX18_N	PEX_TX18_P	GND	PEX_RX18_N	PEX_RX18_P
4	GND	PEX_TX19_N	PEX_TX19_P	GND	PEX_RX19_N	PEX_RX19_P	GND
5	NC	GND	PEX_TX20_N	PEX_TX20_P	GND	PEX_RX20_N	PEX_RX20_P
6	GND	PEX_TX21_N	PEX_TX21_P	GND	PEX_RX21_N	PEX_RX21_P	GND
7	NC	GND	PEX_TX22_N	PEX_TX22_P	GND	PEX_RX22_N	PEX_RX22_P
8	GND	PEX_TX23_N	PEX_TX23_P	GND	PEX_RX23_N	PEX_RX23_P	GND
9	NC	GND	PEX_TX24_N	PEX_TX24_P	GND	PEX_RX24_N	PEX_RX24_P
10	GND	PEX_TX25_N	PEX_TX25_P	GND	PEX_RX25_N	PEX_RX25_P	GND
11	NC	GND	PEX_TX26_N	PEX_TX26_P	GND	PEX_RX26_N	PEX_RX26_P
12	GND	PEX_TX27_N	PEX_TX27_P	GND	PEX_RX27_N	PEX_RX27_P	GND
13	NC	GND	PEX_TX28_N	PEX_TX28_P	GND	PEX_RX28_N	PEX_RX28_P
14	GND	PEX_TX29_N	PEX_TX29_P	GND	PEX_RX29_N	PEX_RX29_P	GND
15	NC	GND	PEX_TX30_N	PEX_TX30_P	GND	PEX_RX30_N	PEX_RX30_P
16	GND	PEX_TX31_N	PEX_TX31_P	GND	PEX_RX31_N	PEX_RX31_P	GND

Table 3: VPX6U P5 Pin Mapping