

PEX 8747, PCI Express Gen 3 Switch, 48 Lanes, 5 Ports

Highlights

■ PEX 8747 General Features

- 48-lane, 5-port PCIe Gen 3 switch
 - Integrated 8.0 GT/s SerDes
- 27 x 27mm², 676-pin FCBGA package
- Typical Power: 8.0 Watts

■ PEX 8747 Key Features

- **Standards Compliant**
 - PCI Express Base Specification, r3.0 (compatible w/ PCIe r1.0a/1.1 & 2.0)
 - PCI Power Management Spec, r1.2
 - Microsoft Vista Compliant
 - Supports Access Control Services
 - Dynamic link-width control
 - Dynamic SerDes speed control
- **High Performance**
 - ◆ **performancePAK**
 - ✓ Read Pacing (bandwidth throttling)
 - ✓ Multicast
 - ✓ Dynamic Buffer/FC Credit Pool
 - Non-blocking switch fabric
 - Full line rate on all ports
 - Packet Cut-Thru with 100ns max packet latency (x16 to x16)
 - 2KB Max Payload Size
- **Flexible Configuration**
 - Ports configurable as x8 or x16
 - Registers configurable with strapping pins, EEPROM, I²C, or host software
 - Lane and polarity reversal
 - Compatible with PCIe 1.0a PM
- **Quality of Service (QoS)**
 - Eight traffic classes per port
 - Weighted round-robin source port arbitration
- **Reliability, Availability, Serviceability**
 - ◆ **visionPAK**
 - ✓ Per Port Performance Monitoring
 - Per port payload & header counters
 - ✓ SerDes Eye Capture
 - ✓ PCIe Packet Generator
 - ✓ Error Injection and Loopback
 - All ports hot plug capable thru I²C (Hot Plug Controller on every port)
 - ECRC and Poison bit support
 - Data Path parity
 - Memory (RAM) Error Correction
 - INTA# and FATAL_ERR# signals
 - Advanced Error Reporting
 - Port Status bits and GPIO available
 - Per port error diagnostics
 - JTAG AC/DC boundary scan

The ExpressLane™ PEX 8747 device offers Multi-Host PCI Express switching capability enabling users to connect a host to its respective endpoints via scalable, high bandwidth, non-blocking interconnection to a variety of **graphics applications**. The PEX 8747 is optimized to support high-resolution graphics while supporting **peer-to-peer** traffic and multicast for maximum performance.

High Performance & Low Packet Latency

The PEX 8747 architecture supports packet **cut-thru with a maximum latency of 100ns (x16 to x16)**. This, combined with large packet memory, flexible common buffer/FC credit pool and non-blocking internal switch architecture, provides full line rate on all ports for performance-hungry applications such as **servers** and **switch fabrics**. The low latency enables applications to achieve high throughput and performance. In addition to low latency, the device supports a packet payload size of up to 2048 bytes, enabling the user to achieve even higher throughput.

Data Integrity

The PEX 8747 provides **end-to-end CRC (ECRC)** protection and **Poison bit** support to enable designs that require **end-to-end data integrity**. PLX also supports data path parity and memory (RAM) error correction circuitry throughout the internal data paths as packets pass through the switch.

Flexible Configuration

The PEX 8747's 5 ports can be configured to lane widths of x8 or x16. Flexible buffer allocation, along with the device's **flexible packet flow control**, maximizes throughput for applications where more traffic flows in the downstream, rather than upstream, direction. Any port can be designated as the upstream port, which can be changed dynamically. Figure 1 shows some of the PEX 8747's common port configurations.




Figure 1. Common Port Configurations

SerDes Power and Signal Management

The PEX 8747 provides low power capability that is fully compliant with the PCIe power management specification and supports software control of the SerDes outputs to allow optimization of power and signal strength in a system. Furthermore, the SerDes block supports **loop-back modes** and **advanced reporting of error conditions**, which enables efficient management of the entire system.

PEX 8747, PCI Express Gen 3 Switch, 48 Lanes, 5 Ports

Interoperability

The PEX 8747 is designed to be fully compliant with the PCI Express Base Specification r2.0, and is backwards compatible to PCI Express Base Specification r1.1 and r1.0a. Additionally, it supports **auto-negotiation**, **lane reversal**, and **polarity reversal**. Furthermore, the PEX 8747 is tested for Microsoft Vista compliance. All PLX switches undergo thorough interoperability testing in PLX's **Interoperability Lab** and **compliance testing at the PCI-SIG plug-fest**.

performancePAK™

Exclusive to PLX, *performancePAK* is a suite of unique and innovative performance features which allows PLX's Gen 2 switches to be the highest performing Gen 2 switches in the market today. The *performancePAK* features consists of the Read Pacing, Multicast, and Dynamic Buffer Pool.

Read Pacing

The Read Pacing feature allows users to throttle the amount of read requests being made by downstream devices. When a downstream device requests several long reads back-to-back, the Root Complex gets tied up in serving that downstream port. If that port has a narrow link and is therefore slow in receiving these read packets from the Root Complex, then other downstream ports may become starved – thus, impacting performance. The Read Pacing feature enhances performances by allowing for the adequate servicing of all downstream devices.

Multicast

The Multicast feature enables the copying of data (packets) from one ingress port to multiple (up to 4) egress ports in one transaction allowing for higher performance in dual-graphics, storage, security, and redundant applications, among others. Multicast relieves the CPU from having to conduct multiple redundant transactions, resulting in higher system performance.

Dynamic Buffer Pool

The PEX 8747 employs a dynamic buffer pool for Flow Control (FC) management. As opposed to a static buffer scheme which assigns fixed, static buffers to each port, PLX's dynamic buffer allocation scheme utilizes a common pool of FC Credits which are shared by other ports. This shared buffer pool is fully programmable by the user, so FC credits can be allocated among the ports as needed. Not only does this prevent wasted buffers and inappropriate buffer assignments, any unallocated buffers

remain in the common buffer pool and can then be used for faster FC credit updates.

visionPAK™

Another PLX exclusive, *visionPAK* is a debug diagnostics suite of integrated hardware and software instruments that users can use to help bring their systems to market faster. *visionPAK* features consist of Performance Monitoring, SerDes Eye Capture, Error Injection, SerDes Loopback, and more.

Performance Monitoring

The PEX 8747's real time performance monitoring allows users to literally "see" ingress and egress performance on each port as traffic passes through the switch using PLX's Software Development Kit (SDK). The monitoring is completely passive and therefore has no affect on overall system performance. Internal counters provide extensive granularity down to traffic & packet type and even allows for the filtering of traffic (i.e. count only Memory Writes).

SerDes Eye Capture

Users can evaluate their system's signal integrity at the physical layer using the PEX 8747's SerDes Eye Capture feature. Using PLX's SDK, users can view the receiver eye of any lane on the switch. Users can then modify SerDes settings and see the impact on the receiver eye. Figure 2 shows a screenshot of the SerDes Eye Capture feature in the SDK.

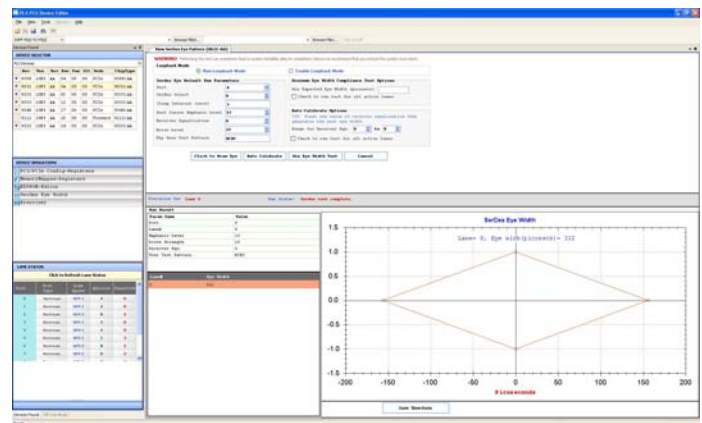


Figure 2. SerDes Eye Capture

PCIe Packet Generator

The PEX 8747 features a full-fledged PCIe Packet Generator capable of creating programmable PCIe traffic running at up to Gen 3 speeds and capable of saturating a x16 link. Using PLX's Software Development Kit (www.plxtech.com/sdk), designers can create custom

PEX 8747, PCI Express Gen 3 Switch, 48 Lanes, 5 Ports

traffic scripts for system bring-up and debug. Fully integrated into the PEX 8747, the Packet Generator proves to be a very convenient on-chip debug tool. Furthermore, the Packet Generator can be used to create PCIe traffic to test and debug other devices on the system.

Error Injection & SerDes Loopback

Using the PEX 8747's Error Injection feature, users can inject malformed packets and/or fatal errors into their system and evaluate a system's ability to detect and recover from such errors. The PEX 8747 also supports Internal Tx, External Tx, Recovered Clock, and Recovered Data Loopback modes.

Applications

Suitable for **host-centric** as well as **peer-to-peer traffic patterns**, the PEX 8747 can be configured for a variety of graphics applications.

Dual-GPU Fan-out

In a graphics fan-out application (see Figure 3), the PEX 8747 drives a dual-output display. The PEX 8747 will fan out to two Graphics Modules (shown as GPUs in the Figures) via the two x16 downstream ports while the x16 upstream port links to the Root Complex. Each graphics module drives its own monitor. Increasing memory and bandwidth requirements have put a strain on local GPU memory. The PEX 8747 allows for highly efficient data transfers over the PCI Express bus, allowing the Graphics Module to utilize the system memory and render it as if it were local graphics memory. In a fan-out application such as this one, each Graphics Module drives its own output.




Figure 3. Dual Graphics Fan-Out

Quad-GPU Fan-Out

For Quad-GPU applications using four GPUs to drive four separate displays, users can take advantage of the PEX 8747's four x8 downstream ports. Mid-range to low-end systems that may not require the bandwidth provided by a x16 Link can instead take advantage of the PEX 8747's x8 Links and fan-out to up to four display units. Figure 4 illustrates a graphics system which uses four x8 links to fan-out to four display units.




Figure 4. Quad Graphics Fan-Out

Dual-GPU w/ Peer-to-Peer Communication

High resolution 3D graphics applications can take full advantage of the PEX 8747. Applications such as high-resolution gaming, high resolution scientific use, and image processing can benefit from the performance of the PEX 8747 switch.

Figure 2 illustrates the use of the device in a high resolution gaming application where two Graphics Modules (GPUs) drive a single monitor for the ultimate gaming experience. The upstream x16 port links to the Root Complex and the two downstream ports connect to the GPUs. The peer-to-peer support of the PEX 8747 allows the two GPUs to communicate with each other for maximum performance.




Figure 5. Dual Graphics with Peer-to-Peer Communication

In this example, the two GPUs divide the screen into a checkerboard pattern. In Figure 5, the screen is divided into white frames and blue frames, with one GPU managing the white frames and the other managing the blue frames. This mode of operation is referred to as Supertiling, and is generally the most efficient because it evenly divides the processing and graphics rendering workload across the two GPUs. This usage model calls for heavy peer-to-peer communication between the two GPUs. The PEX 8747 can also support dual-graphics solutions running in Scissor, or Alternate Frame-Rate (AFR) modes. In each of these modes, the processing and graphics rendering workload is shared by the GPUs, and therefore requires a great amount of peer-to-peer communication between the GPUs to monitor each other's progress and execution.

PEX 8747, PCI Express Gen 3 Switch, 48 Lanes, 5 Ports

Software Model

From a system model viewpoint, each PCI Express port is a virtual PCI to PCI bridge device and has its own set of PCI Express configuration registers. It is through the upstream port that the BIOS or host can configure the other ports using standard PCI enumeration. The virtual PCI to PCI bridges within the PEX 8747 are compliant to the PCI and PCI Express system models. The Configuration Space Registers (CSRs) in a virtual primary/secondary PCI to PCI bridge are accessible by type 0 configuration cycles through the virtual primary bus interface (matching bus number, device number, and function number).

Interrupt Sources/Events

The PEX 8747 switch supports the INTx interrupt message type (compatible with PCI 2.3 Interrupt signals) or Message Signaled Interrupts (MSI) when enabled. Interrupts/messages are generated by PEX 8747 for doorbell interrupts, baseline error reporting, and advanced error reporting.




Figure 6. PEX 8747 RDK

Development Tools

PLX offers hardware and software tools to enable rapid customer design activity. These tools consist of a hardware module (PEX 8747 RDK), hardware documentation (available at www.plxtech.com), and a Software Development Kit (also available at www.plxtech.com).

ExpressLane PEX 8747 RDK

The PEX 8747 RDK (see Figure 6) is a hardware module containing the PEX 8747 which plugs right into your system. The PEX 8747 RDK can be used to test and validate customer software, or used as an evaluation vehicle for PEX 8747 features and benefits. The PEX 8747 RDK provides everything that a user needs to get their hardware and software development started.

Software Development Kit (SDK)

PLX's Software Development Kit is available for download at www.plxtech.com/sdk. The software development kit includes drivers, source code, and GUI interfaces to aid in configuring and debugging the PEX 8747.

Both *performancePAK* and *visionPAK* are supported by PLX's RDK and SDK, the industry's most advanced hardware- and software-development kits.

Product Ordering Information

Part Number	Description
PEX8747-AA80BC G	48-Lane, 5-Port PCI Express Switch, Pb-Free (27x27mm ²)
PEX8747-AA RDK	PEX 8747 Rapid Development Kit

PLX Technology, Inc. All rights reserved. PLX, the PLX logo, ExpressLane, Read Pacing and Dual Cast are trademarks of PLX Technology, Inc. All other product names that appear in this material are for identification purposes only and are acknowledged to be trademarks or registered trademarks of their respective companies. Information supplied by PLX is believed to be accurate and reliable, but PLX assumes no responsibility for any errors that may appear in this material. PLX reserves the right, without notice, to make changes in product design or specification.

Visit www.plxtech.com for more information.