

PEX8796, PCI Express Gen3 Switch, 96 Lanes, 24 Ports

Highlights

■ PEX8796 General Features

- 96-lane, 24-port PCIe Gen3 switch
 - Integrated 8.0 GT/s SerDes
- 35 x 35mm², 1156-ball FCBGA package
- Typical Power: 18.6 Watts

■ PEX8796 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 Windows Logo Compliant
 - Supports Access Control Services
 - Dynamic link-width control
 - Dynamic SerDes speed control
- **High Performance**
 - ◆ **performancePAK**
 - ✓ Multicast
 - ✓ Dynamic Buffer/FC Credit Pool
 - Non-blocking switch fabric
 - Full line rate on all ports
 - Cut-Thru with 150ns max packet latency
 - 2KB Max Payload Size
- **Multi-Host & Fail-Over Support**
 - 2 Configurable Non-Transparent ports
 - Failover with Non-Transparent port
 - Up to 4 upstream/Host ports with 1+1 or N+1 failover to other upstream ports
- **Quality of Service (QoS)**
 - Traffic Class Queuing
 - Eight traffic classes per port
 - Weighted round-robin source port arbitration
- **Reliability, Availability, Serviceability**
 - ◆ **visionPAK**
 - ✓ Per Port Performance Monitoring
 - ✓ SerDes Eye Capture
 - ✓ PCIe Packet Generator
 - ✓ Error Injection and Loopback
 - 6 Hot-Plug port with native HP Signals
 - All ports Hot-Plug capable thru I²C
 - SSC Isolation on up to 24 ports
 - ECRC and Poison bit support
 - Data Path parity
 - Memory (RAM) Error Correction
 - Advanced Error Reporting
 - Port Status bits and GPIO available
 - JTAG AC/DC boundary scan

The ExpressLane™ PEX8796 device offers Multi-Host PCI Express switching capability enabling users to connect multiple hosts to their respective endpoints via scalable, high bandwidth, non-blocking interconnection to a wide variety of applications including **servers, storage systems, and communications platforms**. The PEX8796 is well suited for **fan-out, aggregation, and peer-to-peer applications**.

Multi-Host Architecture

The PEX8796 employs an enhanced architecture, which allows users to configure the device in legacy single-host mode or multi-host mode with up to four host ports capable of 1+1 (one active & one backup) or N+1 (N active & one backup) host failover. This powerful architectural enhancement enables users to build PCIe based systems to support high-availability, failover, redundant and clustered systems.

High Performance & Low Packet Latency

The PEX8796 architecture supports packet **cut-thru with a maximum latency of 150ns (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 PEX8796 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 PEX8796's 24 ports can be configured to lane widths of x4, 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 PEX8796's common port configurations in legacy Single-Host mode.




Figure 1: Common Port Configurations

The PEX8796 can also be configured in Multi-Host mode where users can choose up to four ports as host/upstream ports and assign a desired number of downstream ports to each host. In Multi-Host mode, a virtual switch is created for each host port and its associated downstream ports inside the device. The traffic between the ports of a virtual switch is completely isolated from the traffic in other virtual switches. Figure 2 illustrates some configurations of the PEX8796 in Multi-Host mode where each ellipse represents a virtual switch inside the device.

PEX8796, PCI Express Gen3 Switch, 96 Lanes, 24 Ports

The PEX8796 also provides several ways to configure its registers. The device can be configured through strapping pins, I²C interface, host software, or an optional serial EEPROM. This allows for easy debug during the development phase, performance monitoring during the operation phase, and driver or software upgrade.




Figure 2: Common Multi-Host Configurations

Dual-Host & Failover Support

In Single-Host mode, the PEX8796 supports a **Non-Transparent (NT) Port**, which enables the implementation of **dual-host systems** for redundancy and host failover capability. The NT port allows systems to isolate host memory domains by presenting the processor subsystem as an endpoint rather than another memory system. Base address registers are used to translate addresses; doorbell registers are used to send interrupts between the address domains; and scratchpad registers (accessible by both CPUs) allow inter-processor communication (see Figure 3).




Figure 3: Non-Transparent Port

Multi-Host & Failover Support

In Multi-Host mode, PEX8796 can be configured with up to four upstream host ports, each with its own dedicated downstream ports. The device can be configured for 1+1 redundancy or N+1 redundancy. The PEX8796 allows the hosts to communicate their status to each other via special door-bell registers. In failover mode, if a host fails, the host designated for failover will disable the upstream port attached to the failing host and program the downstream ports of that host to its own domain. Figure 4a shows a two host system in Multi-Host mode with two virtual switches inside the device and Figure 4b shows Host 1 disabled after failure and Host 2 having taken over all of Host 1's end-points.




Figure 4a. Multi-Host




Figure 4b. Multi-Host Fail-Over

Hot-Plug for High Availability

Hot plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX8796 Hot-Plug capability feature makes it suitable for **High Availability (HA) applications**. Four downstream ports include a Standard Hot Plug Controller. If the PEX8796 is used in an application where one or more of its downstream ports connect to PCI Express slots, each port's Hot-Plug Controller can be used to manage the Hot-Plug event of its associated slot. Every port on the PEX8796 is equipped with a Hot-Plug control/status register to support Hot-Plug capability through external logic via the I²C interface.

SerDes Power and Signal Management

The PEX8796 supports software control of the SerDes outputs to allow optimization of power and signal strength in a system. The PLX SerDes implementation supports four levels of power – off, low, typical, and high. The SerDes block also supports **loop-back modes** and **advanced reporting of error conditions**, which enables efficient management of the entire system.

Interoperability

The PEX8796 is designed to be fully compliant with the PCI Express Base Specification r3.0, 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 PEX8796 is tested for Microsoft Windows Logo compliance. All PLX switches undergo thorough interoperability testing in PLX's **Interoperability Lab** and **compliance testing at the PCI-SIG plug-fest**.

PEX8796, PCI Express Gen3 Switch, 96 Lanes, 24 Ports

performancePAK™

Exclusive to PLX, *performancePAK* is a suite of unique and innovative performance features which allows PLX's Gen 3 switches to be the highest performing Gen 3 switches in the market today. The *performancePAK* features consists of Multicast, and Dynamic Buffer Pool.

SSC Isolation

The PEX8796 employs a multi-clock domain which allows the user to terminate Spread Spectrum Clock enabled domains. Terminating the SSC clock removes the need to pass a common clock across a backplane. The PEX8796 supports SSC isolation on all of its ports.

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 PEX8796'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 PEX8796'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 5 shows a screenshot of the SerDes Eye Capture feature in the SDK.

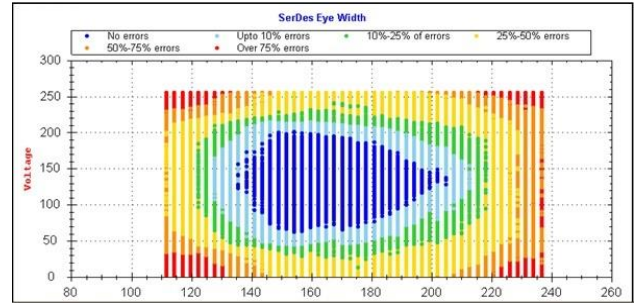


Figure 5. SerDes Eye Capture

Error Injection & SerDes Loopback

Using the PEX8796'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 PEX8796 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 PEX8796 can be configured for a wide variety of form factors and applications.

Host Centric Fan-out

The PEX8796, with its symmetric or asymmetric lane configuration capability, allows user-specific tuning to a variety of host-centric applications. Figure 6 shows a **server** design where, in a quad or multi processor system, users can assign endpoints/slots to CPU cores to distribute the system load. The packets directed to different CPUs will go to different (user assigned) PEX8796 upstream ports, allowing better queuing and load balancing capability for higher performance.




Figure 6. Host Centric Dual Upstream

PEX8796, PCI Express Gen3 Switch, 96 Lanes, 24 Ports

Multi-Host Systems

In multi-host mode, the PEX8796 can support up to four hosts at once. By creating four virtual switches, the PEX8796 allows four hosts to fan-out to their respective endpoints. This reduces the number of switches required for fan-out, saving precious board space and power. In Figure 7, the PEX8796 is being shared by four different CPU cores (hosts) on three servers, with each CPU core running its own applications (I/Os). The PEX8796 assigns the endpoints to the appropriate host and isolates them from the other hosts. In Figure 7, the endpoints are assigned to the CPU core of the same color.




Figure 7. Multi-Host System

High-End Graphics Systems

The PEX8796's 96 lanes can come in handy for high-end graphics systems, such as high-resolution gaming and medical imaging applications, employing two or more GPUs alongside multiple endpoints. Figure 8a shows a dual-graphics system where the PEX8796 is being used to fan-out to two GPUs using x16 links and numerous (up to 16) endpoints using x4 and x8 links. Figure 8b, shows an ultra-gaming system using the PEX8796 to fan-out to four GPUs via x16 links. In both examples, two x16 upstream ports are being used. Using PEX8796's multi-host capabilities, two separate virtual hierarchies can be created, assigning GPUs & endpoints to a particular CPU core via the x16 upstream ports.




Figure 8a. Dual-Graphics System




Figure 8b. Ultra-Gaming System

Host Failover

The PEX8796 can also be utilized in applications where host failover is required. In the below application (Figure 9), two hosts may be active simultaneously and controlling their own domains while exchange status information through doorbell registers or I²C interface. The devices can be programmed to trigger fail-over if the heartbeat information is not provided. In the event of a failure, the surviving device will reset the endpoints connected to the failing CPU and enumerate them in its own domain without impacting the operation of endpoints already in its domain.




Figure 9. Host Fail-Over

PEX8796, PCI Express Gen3 Switch, 96 Lanes, 24 Ports

Software Usage 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 PEX8796 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 PEX8796 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 PEX8796 for hot plug events, doorbell interrupts, baseline error reporting, and advanced error reporting.

Development Tools

PLX offers hardware and software tools to enable rapid customer design activity. These tools consist of a hardware module (PEX8796 RDK), hardware documentation (available at www.plxtech.com), and a Software Development Kit (also available at www.plxtech.com).

ExpressLane PEX8796 RDK

The PEX8796 RDK is a baseboard RDK containing the PEX8796 and multiple PCIe slots for connecting endpoints. The PEX8796 RDK can be used to test and

validate customer software, or used as an evaluation vehicle for PEX8796 features and benefits. The PEX8796 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 PEX8796.

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
PEX8796-AB80BI G	96-Lane, 24-Port PCI Express Gen 3 Switch, Pb-Free (35x35mm ²)
PEX8796-AA RDK	PEX8796 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.