

Features

- **PEX 8608 General Features**
 - 8-lane PCI Express switch
 - Integrated 5.0 GT/s SerDes
 - Up to 8 configurable ports
 - 15 x 15mm², 196-ball PBGA
 - Typical Power: 1.39 Watts
- **PEX 8608 Key Features**
 - **Standards Compliant**
 - PCI Express Base Specification r2.0 (Backwards compatible with PCIe r1.0a/1.1)
 - PCI Power Management Spec r1.2
 - Microsoft Vista Compliant
 - Supports Access Control Services
 - Dynamic link-width control
 - Dynamic SerDes Speed Control
 - **High Performance**
 - Non-blocking internal architecture
 - Full line rate on all ports
 - Cut-Thru latency: 130ns
 - 2KB max payload size
 - Read Pacing (intelligent bandwidth allocation)
 - Dual Cast
 - **Dual-Host & Fail-Over Support**
 - Configurable Non-Transparent port
 - Moveable upstream port
 - Crosslink port capability
 - **Flexible Configuration**
 - 8 flexible & configurable ports (x1 or x4)
 - Configurable with strapping pins, EEPROM, I²C, or Host software
 - Lane and polarity reversal
 - **PCI Express Power Management**
 - Link power management states: L0, L0s, L1, L2/L3 Ready, and L3
 - Device states: D0 and D3_{hot}
 - **Spread Spectrum Clock Isolation**
 - Dual clock domain
 - **Quality of Service (QoS)**
 - Two Virtual Channels (VC) per port
 - Eight Traffic Classes per port
 - Weighted Round-Robin Port & VC Arbitration
 - **Reliability, Availability, Serviceability**
 - All ports Hot-Plug capable thru I²C (Hot-Plug Controller on every port)
 - ECRC & Poison bit support
 - Data path protection
 - Memory (RAM) error correction
 - Advanced Error Reporting support
 - Port Status bits and GPIO available
 - Per port error diagnostics
 - Performance monitoring (per port payload & header counters)
 - JTAG AC/DC boundary scan
 - Fatal Error (FATAL_ERR#) output signal
 - INTA# output signal

PEX 8608

PCIe Gen2, 5.0 GT/s 8-lane 8-port PCI Express® Switch

The *ExpressLane*[™] PEX 8608 device offers PCI Express switching capability enabling users to add scalable high bandwidth non-blocking interconnection to a wide variety of applications including **communications platforms, control plane applications and embedded systems**. The PEX 8608 is well suited for **fan-out, aggregation, peer-to-peer, and intelligent I/O module** applications.

Low Packet Latency & High Performance

The PEX 8608 architecture supports packet **cut-thru with a maximum latency of 130ns in x4 to x1 configuration**. This, combined with large packet memory and non-blocking internal switch architecture, provides full line rate on all ports for low-latency applications such as **communications and embedded**. The low latency enables applications to achieve high throughput and performance. In addition to low latency, the device supports a **max payload size of 2048 bytes**, enabling the user to achieve even higher throughput.

Data Integrity

The PEX 8608 provides **end-to-end CRC** protection (ECRC) and **Poison** bit support to enable designs that require **guaranteed error-free packets**. PLX also supports data path parity and memory (RAM) error correction as packets pass through the switch.

Dual-Host and Fail-Over Support

The PEX 8608 supports full non-transparent bridging (NTB) functionality to allow implementation of **multi-host systems** and **intelligent I/O modules** in applications which require redundancy support such as **communications, storage, and servers**.

Non-transparent bridges allow 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 are accessible from both address domains to allow inter-processor communication.

Interoperability

The PEX 8608 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 each port supports **auto-negotiation, lane reversal** and **polarity reversal**. Furthermore, the PEX 8608 is designed 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** to ensure compatibility with PCI Express devices in the market.

Device Operation Configuration Flexibility

The PEX 8608 provides several ways to configure its operations. The device can be configured through strapping pins, I²C interface, CPU configuration cycles and/or an optional serial EEPROM. This allows for easy debug during the development phase and functional monitoring during the operation phase.



Flexible Port Configurations

The PEX 8608 supports the port configurations as shown in figure 1 below. The figure shows an 8-port configuration, a 5-port configuration as well as a 2-port configuration.




Figure 1. Port Configurations

Hot-Plug for High Availability

Hot-Plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8608 Hot-Plug capability feature makes it suitable for **High Availability (HA) applications**. If the PEX 8608 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 PEX 8608 is equipped with a Hot-Plug control/status register to support Hot-Plug capability through external logic via the I²C interface.

Dual Cast

The PEX 8608 supports Dual Cast, a feature which allows for the copying of data (e.g. packets) from one ingress port to two egress ports allowing for higher performance in storage, security, and mirroring applications.

Read Pacing

The Read Pacing feature allows users to throttle the amount of read requests being made by downstream devices. In the case where a downstream device requests several long reads back-to-back, the Root Complex gets tied up in serving this downstream port. If this 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 performance by allowing for the adequate servicing of all downstream devices by intelligent handling of read requests.

SerDes Power and Signal Management

The PEX 8608 provides low power capability that is fully compliant with the PCI Express power management specification. In addition, the SerDes physical links can be turned off when unused for even lower power. The PEX 8608 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 debug and management of the entire system.

Port and Virtual Channel (VC) Arbitration

The PEX 8608 switch supports hardware fixed and Weighted Round-Robin (WRR) Ingress Port Arbitration. This allows fine tuning of Quality of Service and efficient use of packet buffers for better system performance. The PEX 8608 also supports WRR VC arbitration scheme between the two virtual channels.

Applications

Suitable for **fan-out, control plane applications, embedded systems** as well as **intelligent I/O and host isolation applications**, PEX 8608 can be configured for a wide variety of form factors and applications.

Fan-Out

The PEX 8608 switch, with its high port count and flexible configurations, allows user specific tuning to a variety of **host-centric** as well as **peer-to-peer applications**.




Figure 2. Fan-in/out Usage

Figure 2 shows a typical **fan-out** design, where the processor provides a PCI Express link that needs to be fanned into a larger number of smaller ports for a variety of I/O functions, each with different bandwidth requirements.

In this example, the PEX 8608 would typically have a 1-lane upstream port, and as many as 7 downstream ports. The downstream ports provide x1 PCI Express connectivity to the endpoints. With its eight ports, the PEX 8608 can provide fan-out connectivity to up to seven PCI Express devices. The figure also shows how some of the ports can be bridged to provide **PCI slots or Generic devices** through the use of the **PEX 8311 and PEX 8112 PCIe bridging devices**.

Control Plane Application

The PEX 8608 is ideal for control planes in routers and other communications sub-systems to meet increased packet processing needs without compromising latency. Figure 3 shows a controller card with a PEX 8608 connecting the Control Processor to as many as seven line cards each via an x1 connection. This usage model provides connectivity to multiple line cards giving the processor control over a large number of line cards in communication platforms.




Figure 3. Control Plane Application

Control Planes for Large Systems

Two PEX 8608 devices can be used in applications where there is need for more than 7 line cards and the processor has 2 PCIe ports. An example of a control plane application with more than 7 line cards is shown in Figure 4.




Figure 4. Control Planes for Large Systems

Embedded Control Planes

The PEX 8608 is well suited for embedded control plane applications. An example of such application is a packet processing module, as shown in Figure 5. In this example, the FPGAs and ASICs implement an x1 PCI Express Interface. With its high port count, the PEX 8608 can provide connectivity to a high number of FPGAs and ASICs for a robust packet processing communications sub-system.




Figure 5. Packet Processing Module

Intelligent Adapter Card

The PEX 8608 supports the **non-transparency** feature. Figure 6 illustrates a host system using an intelligent adapter card.

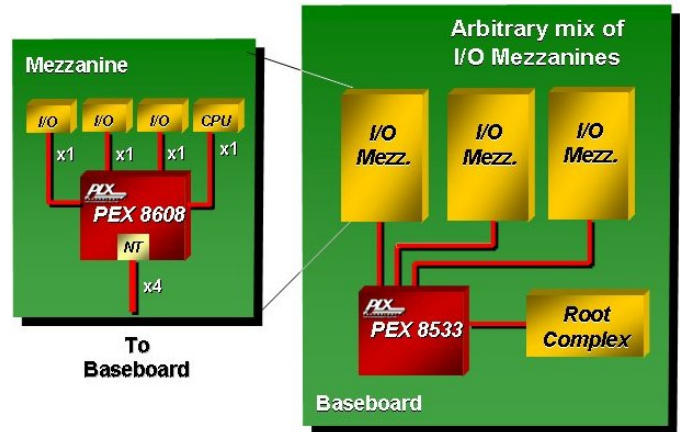


Figure 6. Intelligent Adapter Usage

In this figure, the CPU on the Mezzanine card is isolated from the host CPU. The PEX 8608 non-transparent port allows the two CPUs to be isolated but communicate with each other through various registers that are designed in the PEX 8608 for that purpose. The host CPU **can dynamically re-assign both the upstream port and the non-transparent port** of PEX 8608 allowing the system to be reconfigured.

Active-Standby Failover Model

The PEX 8608 supports applications requiring **dual host, host failover** applications through the **non-transparency** feature. Figure 7 illustrates a dual host system with an active and standby processor configuration.

The redundancy of the host and the fabric can be achieved through many possible configurations using NTB function of PEX 8608. In the configuration shown below Processor A is the primary active host of the PCI Express system. Processor B acts as the backup Host. In the case of a failure on Processor A, application software is instructed to migrate the PCI Express system to Processor B. Consequently, Processor B becomes the active host and Processor A can be replaced as the backup Host.




Figure 7. Active-Standby Host Mode

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 PEX 8608 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 8608 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 8608 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 PEX 8608 Rapid Development Kit (RDK), hardware documentation, and a Software Development Kit (also available at www.plxtech.com/sdk).

RDK

The PEX 8608RDK is a hardware module containing the PEX 8608 which plugs right into your system (Figure 8). The PEX 8608RDK hardware module can be installed in a motherboard, used as a riser card, or configured as a bench-top board. The PEX 8608RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for PEX 8608 features and benefits.

SDK

The SDK tool set includes:

- Linux & Windows drivers
- C/C++ Source code, Objects, libraries
- User's Guides & Application examples





Figure 8. PEX 8608RDK



PLX Technology, Inc.
 870 W. Maude Ave.
 Sunnyvale, CA 94085 USA
 Tel: 1-800-759-3735
 Tel: 1-408-774-9060
 Fax: 1-408-774-2169
info@plxtech.com
www.plxtech.com

Product Ordering Information

Part Number	Description
PEX8608-BA50BC	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm ² pkg
PEX8608-BA50BC G	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm ² pkg, Pb-free
PEX 8608BA-AIC4U1D RDK	PEX 8608 Rapid Development Kit – Add-in Card with x4 Upstream and Four x1 Downstream

Please visit the PLX Web site at <http://www.plxtech.com> or contact PLX sales at 408-774-9060 for sampling.

© 2009 by PLX Technology, Inc. All rights reserved. PLX and the PLX logo are registered trademarks of PLX Technology, Inc. ExpressLane is a trademark of PLX Technology, Inc., which may be registered in some jurisdiction. 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 Technology, Inc. assumes no responsibility for any errors that may appear in this material. PLX Technology, Inc. reserves the right, without notice, to make changes in product design or specification.