

## Features

- **PEX 8609 General Features**
  - 8-lane PCI Express switch
    - Integrated 5.0 GT/s SerDes
  - Up to 8 configurable ports
  - 15 x 15mm<sup>2</sup>, 196-ball PBGA
  - Typical Power: 1.60 Watts
- **PEX 8609 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
  - **Integrated DMA Engine**
    - Four DMA Channels
    - Internal Descriptor Support
    - DMA function independent from transparent switch function
    - 64-bit Addressing
    - Prefetch Descriptor Mode
    - Up to 4.0 GB/s throughput per channel
  - **Dual-Host & Fail-Over Support**
    - Configurable Non-Transparent port (NTB)
    - Moveable upstream port
    - Crosslink port capability
  - **High Performance**
    - Cut-Thru latency: 140ns
    - 2KB max payload size
    - Read Pacing
    - Dual-Cast
  - **Flexible Configuration**
    - 8 flexible & configurable ports (x1 or x4)
    - Configurable with strapping pins, EEPROM, I<sup>2</sup>C, 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<sub>hot</sub>
  - **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 I<sup>2</sup>C (Hot-Plug Controller on every port)
    - Data path protection
    - Memory (RAM) error correction
    - Port Status bits and GPIO available
    - Per port error diagnostics
    - Performance monitoring (per port payload & header counters)



# PEX 8609

## *Flexible & Versatile 8-lane 8-port PCI Express® Switch*

The *ExpressLane*™ PEX 8609 device offers PCI Express switching capability enabling users to add scalable high bandwidth non-blocking interconnection to a wide variety of applications including **control planes, communication platforms, servers, storage systems and embedded systems**. The PEX 8609 is well suited for **fan-out, aggregation, peer-to-peer, and intelligent I/O module** applications.

### Low Packet Latency & High Performance

The PEX 8609 architecture supports packet **cut-thru with a maximum latency of 140ns**. 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 servers**. 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.

### Integrated DMA Engine

The PEX 8609 provides a versatile and powerful DMA engine built in to the device which can be used as a stand alone DMA engine. The DMA engine removes the burden resulting from having to move data between devices away from the processor. This allows the processor to perform computational tasks instead. The four DMA channels can support high data rate transfers between IO devices connected to any of the available ports in the PEX8609. Additionally, the DMA engine in the PEX 8609 can be used to complement the DMA engine in the processor by providing additional DMA channels for higher performance.

### Data Integrity

The PEX 8609 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 8609 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 8609 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 8609 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.

## Flexible Port Configurations

The PEX 8609 supports a large number of port configurations as shown in figure 1 below. Please refer to the PEX 8609 datasheet for more port configuration options.




Figure 1. PEX8609 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 8609 Hot-Plug capability feature makes it suitable for **High Availability (HA) applications**. If the PEX 8609 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 8609 is equipped with a Hot-Plug control/status register to support Hot-Plug capability through external logic via the I<sup>2</sup>C interface.

## Dual Cast

The PEX 8609 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 performances by allowing for the adequate servicing of all downstream devices by intelligent handling of read requests.

## SerDes Power and Signal Management

The PEX 8609 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 8609 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 8609 switch supports hardware fixed and Weighted Round-Robin Ingress Port Arbitration. This allows fine tuning of Quality of Service and efficient use of packet buffers for better system performance. The PEX 8609 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 applications**, PEX 8609 can be configured for a wide variety of form factors and applications.

## Fan-Out

The PEX 8609 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

In this example, the PEX 8609 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. 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. The DMA engine in the PEX8609 can alternatively be used for data transfers between the ASICs, FPGAs and the host.

The integrated DMA engine together with the PEX8609 ability to support direct peer-to-peer transfers can alternately be used to transfer data between the ASICs and FPGAs in the system.

## Active-Standby Failover Model

The PEX 8609 supports applications requiring **host failover** applications through the **non-transparency and DMA**

The diagram illustrates a dual-processor system. On the left, **Processor A** is shown in an **Active** state, connected to a **Root Complex**. Below the Root Complex is a **PEX 8609** controller with **DMA** and **NT** (Non-Transparent) blocks. On the right, **Processor B** is in a **Standby** state, also connected to its own **Root Complex** and a **PEX 8609** controller with **NT** and **DMA** blocks. A horizontal connection between the two PEX 8609 controllers is labeled **x4**. The PCI Express logo is at the bottom center.

## DMA in Control Planes

The diagram illustrates the PEX 8609 architecture. At the center is a red box labeled "PEX 8609" with "FX" and "DMA" logos. It is connected to six peripheral components: two ASICs and two FPGAs at the top, two ASICs on the left, and one FPGA and one ASIC at the bottom. All these connections are bidirectional. To the right of the PEX 8609 is a yellow box labeled "Control Processor", connected to the PEX 8609 by a single red line. Further to the right is another yellow box labeled "Memory", connected to the Control Processor by a double-headed black arrow.

## Distributed Computing


The diagram illustrates a multi-server system architecture for a data center. It shows four servers, labeled Server A, Server B, Server C, and Server D, each containing a CPU, MEM (Memory), RC (Root Complex), and a PEX8609 PCIe switch. The PEX8609 switch is connected to a common bus labeled "DMA Transfers Between Servers". A PEX8618 switch is shown at the bottom, connected to the common bus via four x4 links. The PEX8618 switch is also connected to a "vd." (Virtual Disk) component.

The PEX 8609 supports the **non-transparency** feature.

The diagram shows a Primary Host CPU Blade connected to two Intelligent I/O Adaptors. The CPU Blade contains a CPU and MEM modules. It is connected via x4 links to two PEX8609 Intelligent I/O Adaptors. Each adaptor contains a CPU, MEM, and DMA module. The PEX8609 adaptors are connected via x4 links to a PEX8614 Primary Host CPU Blade. A dashed line indicates a DMA Transfer(s) between the PEX8609 adaptors and the PEX8614 blade.

In this figure, the CPU on the adapter card is isolated from the host CPU. The PEX 8609 non-transparent port allows the two CPUs to be isolated but communicate with each other through various registers that are designed in the PEX 8609 for that purpose. Moreover, the internal DMA engine can be used to copy or mirror the data between adapter cards so that intelligent I/O adapter redundancy is achieved. This model can be expanded to an Active-Standby failover model. In this case, there is an Active host and a Standby host ready to take over the system in the event the Active host fails. The DMA engine in the PEX8609 can be used to mirror the data between hosts thus minimizing the failover time.

The DMA engine in the PEX8609 is very flexible. It implements **Four DMA channels** capable of saturating a x8 Gen2 PCI Express link. Furthermore, it supports a Descriptor Ring approach in which the descriptors can be placed in external host memory or internal to the PEX8609. The DMA engine can move data from any port to any port including the same port. That is, the source and destination for a DMA transfer can be on the same PEX8609 port. Figure 7 shows two PCIe endpoints and a PEX8609 device connected through a PCIe switch. In this example, one PCIe endpoint behind the PEX8609 takes advantage of the DMA capabilities in the PEX8609 to move data from PCIe2 to PCIe1.



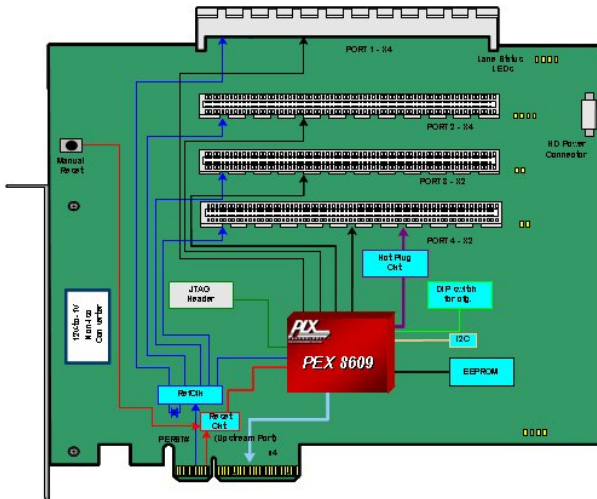
### Figure 7. Generic Data Mover



The PEX8609 is a multi-function device. Function 0 implements the function of a transparent PCI Express switch while Function 1 implements the endpoint functionality of the DMA. 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 DMA function of the PEX8609 implements a Type 0 header, same as an endpoint, and requires a driver. The driver for the DMA function is available from PLX and included in the SDK.


The PEX 8609 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 8609 for Hot-Plug events, doorbell interrupts, baseline error reporting, and advanced error reporting.

PLX offers hardware and software tools to enable rapid customer design activity. These tools consist of a PEX 8609 Rapid Development Kit (RDK), hardware documentation, and a Software Development Kit (also available at [www.plxtech.com/sdk](http://www.plxtech.com/sdk)).



The PEX 8609RDK is a hardware module containing the PEX 8609 which plugs right into your system (Figure 8). The PEX 8609RDK hardware module can be installed in a motherboard, used as a riser card, or configured as a bench-top board. The PEX 8609RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for PEX 8609 features and benefits.

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



Part Number	Description
PEX8609-BA50BC	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm <sup>2</sup> pkg
PEX8609-BA50BC G	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm <sup>2</sup> pkg, Pb-free
PEX8609-BA50BI G	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm <sup>2</sup> pkg, Pb-free Industrial Temperature
PEX8609BA-AIC4U1D RDK	PEX 8609 Add-in Card Kit; x4 Upstream; x1 downstream (4)

4/10

# Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[PLX Technology:](#)

[PEX8609-BA50BIG](#) [PEX8609BA-AIC4U1DRDK](#) [PEX 8609-BA50BC](#) [PEX 8609-BA50BC G](#)



Компания «ЭлектроПласт» предлагает заключение долгосрочных отношений при поставках импортных электронных компонентов на взаимовыгодных условиях!

Наши преимущества:

- Оперативные поставки широкого спектра электронных компонентов отечественного и импортного производства напрямую от производителей и с крупнейших мировых складов;
- Поставка более 17-ти миллионов наименований электронных компонентов;
- Поставка сложных, дефицитных, либо снятых с производства позиций;
- Оперативные сроки поставки под заказ (от 5 рабочих дней);
- Экспресс доставка в любую точку России;
- Техническая поддержка проекта, помощь в подборе аналогов, поставка прототипов;
- Система менеджмента качества сертифицирована по Международному стандарту ISO 9001;
- Лицензия ФСБ на осуществление работ с использованием сведений, составляющих государственную тайну;
- Поставка специализированных компонентов (Xilinx, Altera, Analog Devices, Intersil, Interpoint, Microsemi, Aeroflex, Peregrine, Syfer, Eurofarad, Texas Instrument, Miteq, Cobham, E2V, MA-COM, Hittite, Mini-Circuits, General Dynamics и др.);

Помимо этого, одним из направлений компании «ЭлектроПласт» является направление «Источники питания». Мы предлагаем Вам помощь Конструкторского отдела:

- Подбор оптимального решения, техническое обоснование при выборе компонента;
- Подбор аналогов;
- Консультации по применению компонента;
- Поставка образцов и прототипов;
- Техническая поддержка проекта;
- Защита от снятия компонента с производства.



#### Как с нами связаться

**Телефон:** 8 (812) 309 58 32 (многоканальный)

**Факс:** 8 (812) 320-02-42

**Электронная почта:** [org@eplast1.ru](mailto:org@eplast1.ru)

**Адрес:** 198099, г. Санкт-Петербург, ул. Калинина, дом 2, корпус 4, литера А.