

Features

- **PEX 8609 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: 1.60 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, 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_{not}
 - **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)
 - 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.

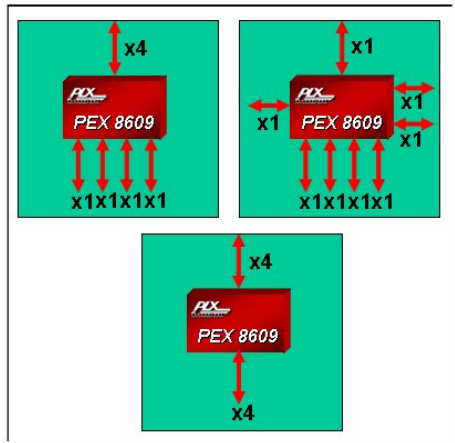


Figure 1. PEX8609 Port Configurations

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 I²C 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**.

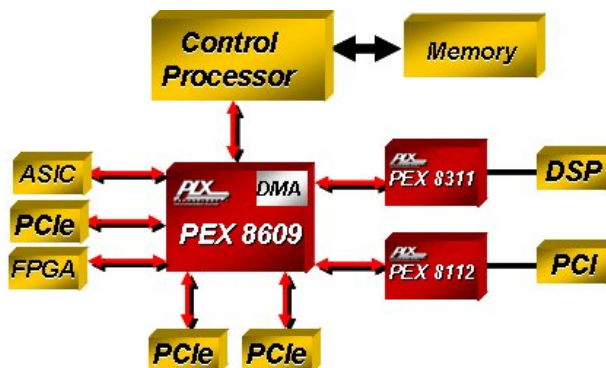


Figure 2. Fan-in/out Usage

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**

feature. Figure 3 illustrates a dual host system with an active and standby processor configuration. The redundancy of the host and the fabric can be achieved through the NTB function of PEX 8609. Additionally, the DMA engine in the PEX8609 can be used to mirror the data on both the active and standby hosts.

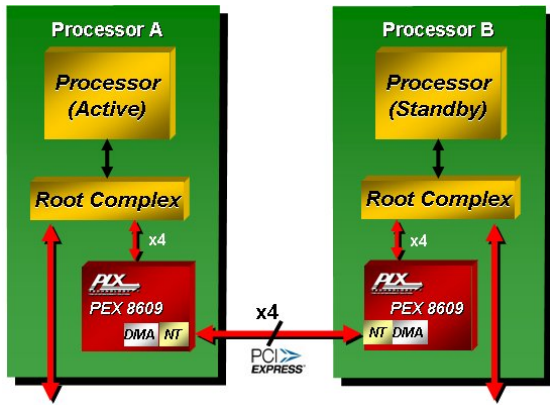


Figure 3. Active-Standby Failover

DMA in Control Planes

The DMA engine in the PEX8609 can alternatively be used to perform housekeeping tasks, such as reading device status information from the many endpoints available. This particular usage model allows the DMA engine to perform the many data transfers for endpoint status that would otherwise be handled by the processor.

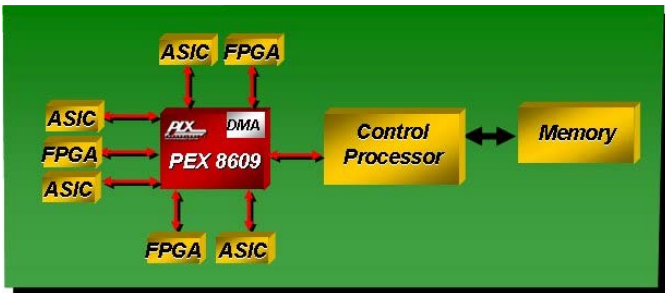


Figure 4. PEX8609 DMA in a Packet Router

Distributed Computing

The PEX 8609 is well suited for distributed computing applications. Once programmed, the DMA engine in the PEX8609 transfers the data between servers and through the NT port at a high rate. The processor can focus all of its computing cycles manipulating the large amounts of data received.

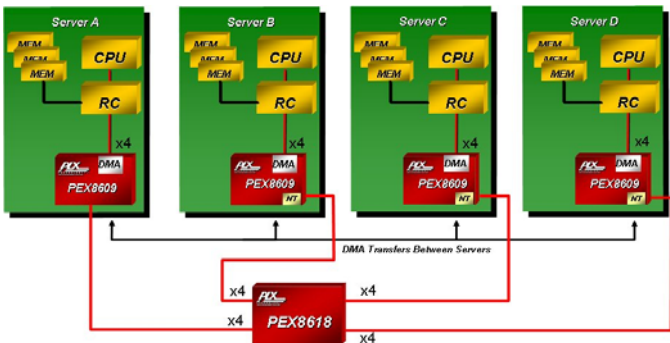


Figure 5. Distributed Computing Intelligent Adapter Card Redundancy

The PEX 8609 supports the **non-transparency** feature. Figure 6 illustrates a host system using an intelligent adapter card.

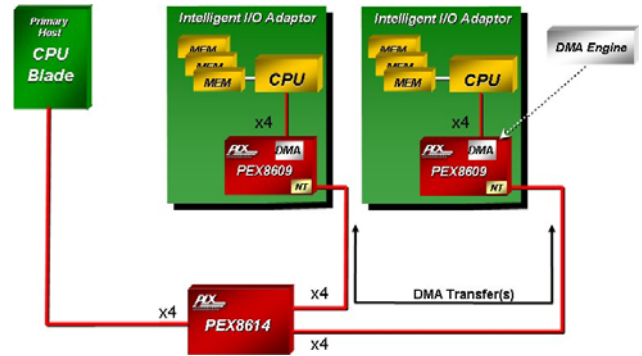


Figure 6. Intelligent Adapter Usage

In this figure, the CPU on the 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.

Generic Data Mover

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.

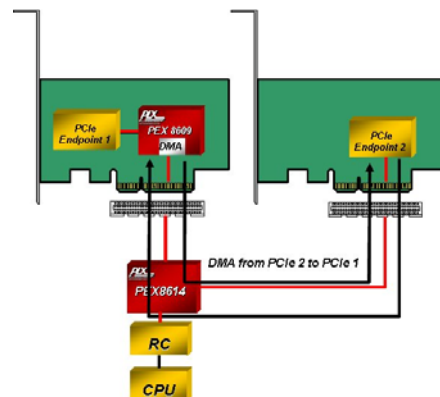


Figure 7. Generic Data Mover

Software Usage Model

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.

Interrupt Sources/Events

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.

Development Tools

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).




Figure 8. PEX 8609RDK


RDK

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.

SDK

The SDK tool set includes:

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



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
PEX8609-BA50BC	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm ² pkg
PEX8609-BA50BC G	8 Lane, 8 Port PCIe Switch, 196-ball PBGA 15x15mm ² pkg, Pb-free
PEX8609BA-AIC4U1D RDK	PEX 8609 Add-in Card Kit; x4 Upstream; x1 downstream (4)

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