

PEX 8508

Features

- **PEX 8508 General Features**
 - 8-lane PCI Express switch
 - Integrated SerDes
 - Up to five configurable ports
 - 19mm x 19mm, 296-ball PBGA package
 - Typical Power: 1.6 Watts
- **PEX 8508 Key Features**
 - **Standards Compliant**
 - PCI Express Base Specification, r1.1
 - Standard SHPC Specification, r1.1 (hot-plug)
 - **High Performance**
 - Non-blocking internal architecture
 - Full line rate on all ports
 - Cut-Thru latency: 150ns
 - **Non-Transparent Bridging**
 - Configurable Non-Transparent port for Multi-Host or Intelligent I/O Support
 - **Flexible Configuration**
 - Five highly flexible & configurable ports (x1, x2, 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, L2 and L3
 - Device states: D0 and D3hot
 - Vaux, WAKE#, Beacon support
 - **Spread Spectrum Clock**
 - Dual clock domain
 - **Quality of Service (QoS)**
 - Two Virtual Channels per port
 - Eight Traffic Classes per port
 - Fixed and Round-Robin Virtual Channel Port Arbitration
 - **Reliability, Availability, Serviceability**
 - 5 Standard Hot-Plug Controllers
 - Upstream port as hot-plug client
 - Transaction Layer end-to-end CRC
 - Poison bit
 - Advanced Error Reporting
 - Lane Status bits and GPO available
 - Per port error diagnostics
 - Bad DLLPs
 - Bad TLPs
 - CRC errors
 - JTAG boundary scan
 - Fatal Error out-of-band signal option

Flexible & Versatile PCI Express[®] Switch

Multi-purpose and Feature Rich PCI Express Switch Family

The *ExpressLane*[™] PEX 8508 device offers PCI Express switching capability enabling users to add scalable high bandwidth, non-blocking interconnection to a wide variety of applications including **intelligent network interface cards, office automation products, mezzanine cards, servers, storage systems, communications platforms, laptop docking stations, and embedded-control products**. The PEX 8508 switch can be used as a **fan-out** or **aggregation** type switch, and is equally well-suited to **peer-to-peer** switching applications.

Highly Flexible Port Configurations

The PEX 8508 offers highly configurable ports. There are a maximum of five ports and eight lanes that can be configured to any legal width from x1 to x4, in any combination to support your specific bandwidth needs. The ports can be **symmetric** (each port having the same lane width) or **asymmetric** (ports having different lane widths). *If you can think of a port/lane combination, you can configure it!* Any of the ports can be designated as the upstream port, and you can even dynamically change the upstream port.

End-to-end Packet Integrity

The PEX 8508 provides end-to-end CRC protection (ECRC) and **Poison bit** support to enable designs that require **guaranteed error-free packets**. These features are optional in the PCI Express specification, but PLX provides them across its entire *ExpressLane* switch product line.

Non-Transparent “Bridging” in a PCI Express Switch

The PEX 8508 supports full non-transparent bridging (NTB) functionality to allow implementation of **multi-host systems** and **intelligent I/O modules** in applications such as **intelligent adapter cards, processor isolation, and office automation**. To ensure quick product migration, the **non-transparency** features are implemented in the same fashion as in standard PCI applications.

Non-transparent bridges allow systems to isolate 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.

Two Virtual Channels

The PEX 8508 switch supports two full-featured Virtual Channels (VCs) and eight Traffic Classes (TCs). The mapping of Traffic Classes to port-specific Virtual Channels allows for different mappings for different ports. In addition, the devices offer user-selectable Virtual Channel arbitration algorithms to enable users to fine tune the Quality of Service (QoS) required for a specific application.

Low Power with Granular SerDes Control

The PEX 8508 provides low power capability that is compliant with the PCI Express power management specification. The SerDes physical links can be turned off when unused for even lower power. Support for Vaux, Wake#, and Beacon allows D3cold operation and is crucial in such applications as docking stations.

Flexible Port Configuration

The lane width for each port can be individually configured through auto-negotiation, hardware strapping, upstream software configuration, or through an optional EEPROM.

The PEX 8508 supports a large number of port configurations. For example, if you are using the PEX 8508 in a fan-out application, you may configure the upstream port as x4 and the downstream ports as four x1 ports or two x2. In a peer-to-peer switch aggregation application you can configure four x2 ports. Figure 1 show most common port configurations.

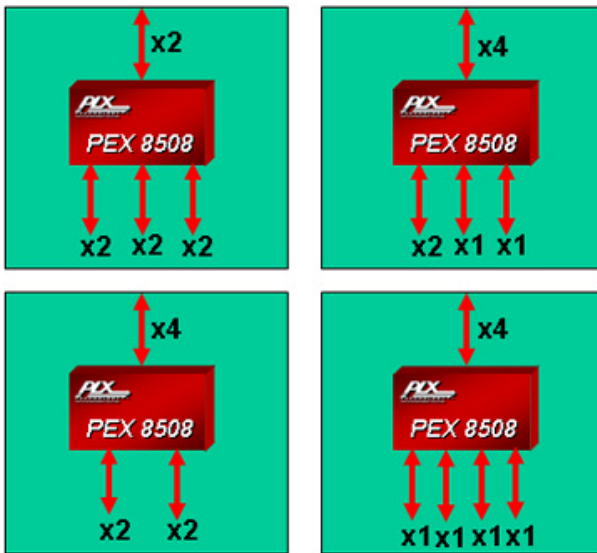


Figure 1. Common Port Configurations

Hot-Plug for High Availability

Hot plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8508 hot-plug capability and **Advanced Error Reporting** features makes it suitable for **High Availability (HA)** applications. Each port includes a Standard **Hot-Plug Controller**. If the PEX 8508 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. Furthermore, its upstream port is a **hot-plug client**, allowing it to be used on hot-pluggable **adapter cards, backplanes and fabric modules**.

Fully Compliant Power Management

For applications that require power management, the PEX 8508 devices support both link (L0, L0s, L1, L2/L3 Ready, L2 and L3) and device (D0 and D3hot) power management states, in compliance with the PCI Express power management specification. Vaux, Wake# and Beacon support the wakeup protocol from the D3cold state.

SerDes Power and Signal Management

The PEX 8508 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.

Flexible Virtual Channel Arbitration

The PEX 8508 switch supports fixed priority and Round Robin arbitration schemes for up to two virtual channels. This allows fine tuning of Quality of Service, efficient use of packet buffers and the system bandwidth.

Applications

Suitable for host-centric as well as Intelligent I/O applications the PEX 8508 can be configured for a wide variety of form factors and applications.

Processor Isolation for Redundancy

The PEX 8508 switch, with its low latency, Hot-Plug, and NTB capability, allows users to isolate processors in **multi-host or failover applications**.

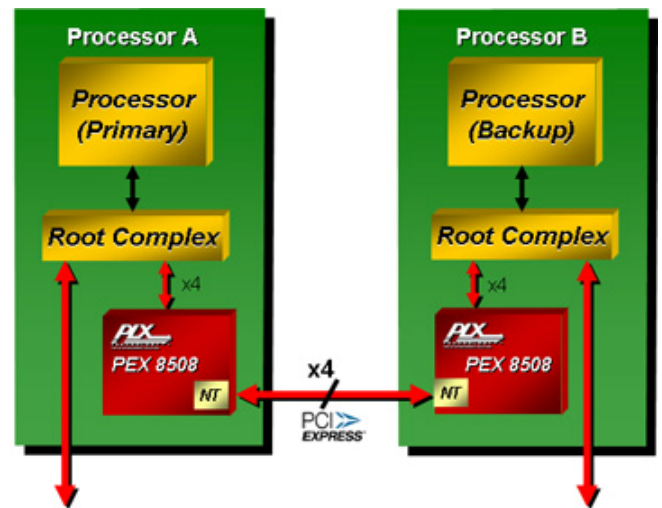


Figure 2. Processor Isolation

Figure 2 shows two server boards. The PCI Express switch is used to prevent the backup host from enumerating and configuring portions of the primary host system. In addition, non transparent ports provide doorbell and mailbox registers for inter-processor communications. The backup processor will monitor heartbeat messages from the primary system and upon failure take over as the primary host.

PEX 8508 has very **low latency**, under **150ns**, to assist in a rapid failover to the backup system. Here only one PEX 8508 switch is required with the one non transparent port but to maintain identical board layouts, each board shown has a PEX 8508 switch. This application can be generalized to where each **compute or storage or I/O board** would include a PEX 8508 to isolate each processor domain.

Adapter Card Aggregation

The number and variety of PCI Express native-mode devices is growing quickly. As these devices become mainstream it will be necessary to create multifunction and multi-port adapter cards with PCI Express capability.

The PEX 8508 can be used to create an **adapter card** that **aggregates** the PCI Express devices into a single port that can be plugged into a **backplane** or a **motherboard**. Figure 3 shows the PEX 8508 in this application.

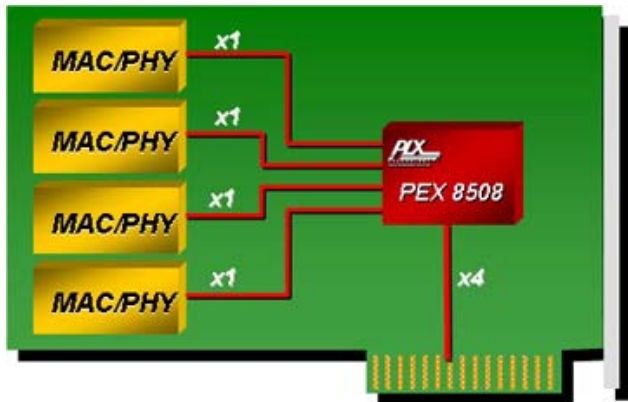


Figure 3. Aggregation Adapter Card

The adapter card in Figure 3 is transparent with no local processor present. Here four Gigabit Ethernet (GE) ports are aggregated through the switch. Recent MAC/PHY devices for GE all well as InfiniBand, Fibre Channel, etc., have native PCI Express support for direct connections to the switch. The **non-blocking** low latency design of PEX 8508 provides full throughput with **latency under 150ns**.

Intelligent Adapter Card

The PEX 8508 supports the non-transparency feature. Figure 4 illustrates an intelligent adapter card used in servers and storage hosts and is commonly used in TCP/IP Offload (TOE) engines, iSCSI processing, RDMA, or encryption/decryption.

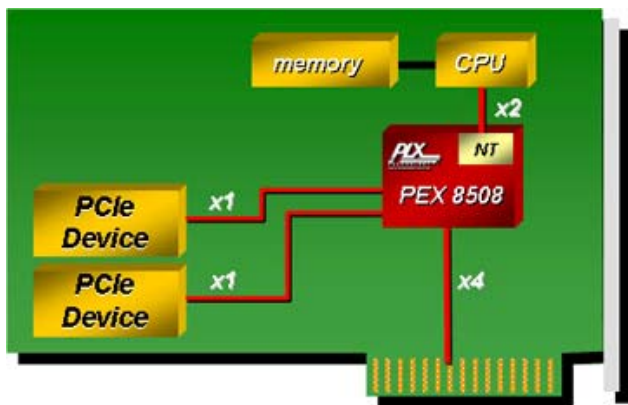


Figure 4. Intelligent Adapter

In this figure, the CPU on the adapter card is isolated from the host CPU (not shown). The PEX 8508 non-transparent port allows the two CPUs to be isolated but communicate with each other through address translation that is integrated into the PEX 8508. The upstream and downstream as well as the non-transparent port of the PEX 8508 are fully configurable. In this adapter, the CPU or other “intelligent” device with on-chip processing capability—needs to be isolated from the rest of the system to avoid memory mapping conflicts.

Office Automation Controller

A low port/lane PCI Express switch can be very useful in a print engine or similar type system controller in embedded applications. Such designs require low port/lane counts and low cost, yet the system throughput can be very high for graphics oriented data. Figure 5 shows that the PEX 8508 is used to transfer data from a scanner and other input devices to the print engine CPU and back to the marking engine. Many endpoint devices will be incorporating native PCI Express support to allow the use of all five ports of the PEX 8508. Complete power management support will reduce system power consumption. With a small 19mm x 19mm footprint, the controller board space can be minimized.

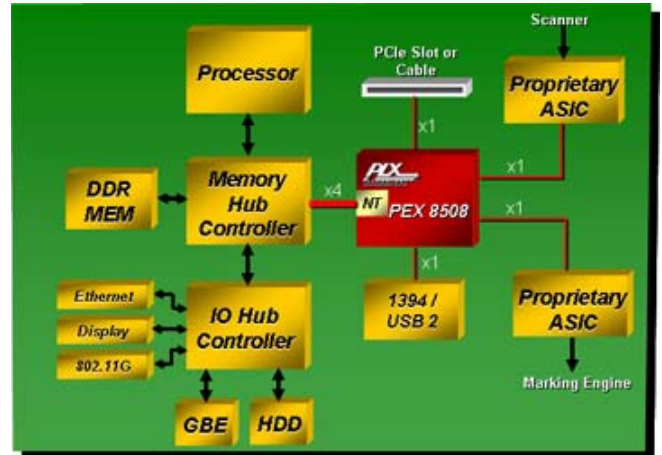


Figure 5. Printer Controller Engine

ATCA and XMC Mezzanine Card Applications

The PEX 8508 supports aggregation or fan-in/out as well as peer-to-peer switching. Figure 6 shows each I/O Mezzanine card with a PEX 8508 which can aggregate data to the main baseboard or allow peer-to-peer switching between I/O endpoints. Non transparency can isolate each of the I/O Mezzanine cards from each other.

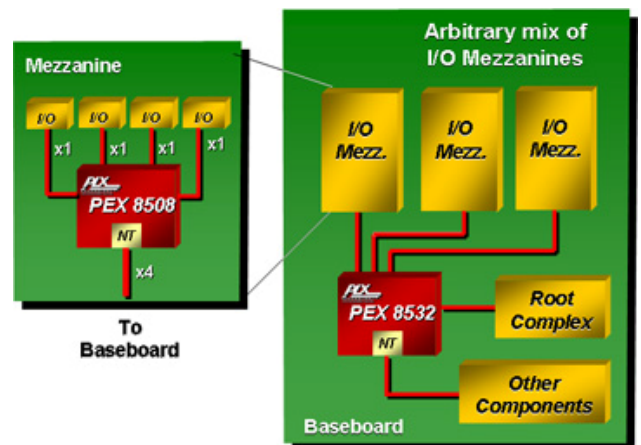


Figure 6. XMC Mezzanine Card

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 8508 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 8508 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 8508 for hot-plug events, baseline error reporting, and advanced error reporting.

Development Tools

PLX is offering hardware and software tools (PEX 8508RDK) to enable rapid customer design activity. These tools include the hardware Rapid Development Kit (RDK) and the Software Development Kit (SDK).

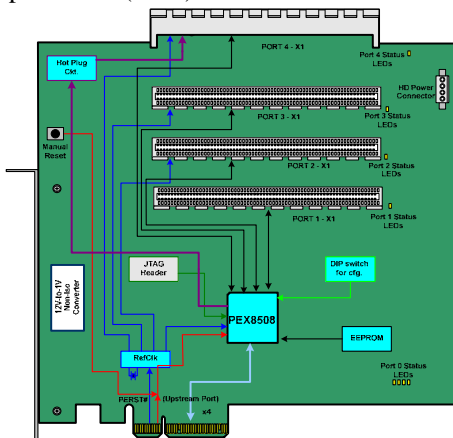


Figure 7. PEX 8508RDK

RDK

The RDK hardware module includes the PEX 8508 with by default one upstream x4 port connected to a male-edge connector, and four downstream x1 ports each connected to a x16 female slot connector. Third party adapters from x16 to x1 allow the x4 connector to plug into narrower slots. Cards with larger than x1 connectors can plug into the x16 female connectors and will auto-negotiate down to x1.

The PEX 8508RDK can be installed in a motherboard, used as a riser card, or configured as a bench-top board.

The PEX 8508RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for the PEX 8508 features and benefits.

SDK

The SDK tool set includes:

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



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
 Email: info@plxtech.com
 Web: www.plxtech.com

Product Ordering Information

Part Number	Description
PEX 8508-AC25BI	8-Lane, 5 Port PCIe Switch, 296-ball PBGA 19x19 pkg
PEX 8508-AC25BI G	8-Lane, 5 Port PCIe Switch, 296-ball PBGA 19x19 pkg, Pb-free
PEX 8508RDK	Rapid Development Kit for PEX 8508

Please visit the PLX Web site at <http://www.plxtech.com/8508> or contact PLX sales at 408-774-9060 for sampling.

© 2006 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.