

PEX 8648

PCIe Gen2, 5.0GT/s 48-lane, 12-port PCIe Switch

Features

- **PEX 8648 General Features**
 - 48-lane, 12-port PCIe Gen2 switch
 - Integrated 5.0 GT/s SerDes
 - 27 x 27mm², 676-pin FCBGA package
 - Typical Power: < 7.0 Watts

- **PEX 8648 Key Features**
 - **Standards Compliant**
 - PCI Express Base Specification, r2.0 (backwards compatible w/ 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 switch fabric
 - Full line rate on all ports
 - Packet Cut-Thru with 140ns max packet latency (x16 to x16)
 - 2KB Max Payload Size
 - Read Pacing (bandwidth throttling)
 - Dual-Cast
 - **Flexible Configuration**
 - Ports configurable as x1, x2, x4, x8, x16
 - Registers configurable with strapping pins, EEPROM, I²C, or host software
 - Lane and polarity reversal
 - Compatible with PCIe 1.0a PM
 - **Dual-Host & Fail-Over Support**
 - Configurable Non-Transparent port
 - Moveable upstream port
 - Crosslink port capability
 - **Quality of Service (QoS)**
 - Eight traffic classes per port
 - Weighted round-robin source port arbitration
 - **Reliability, Availability, Serviceability**
 - 3 Hot Plug Ports with native HP Signals
 - All ports hot plug capable thru I²C (Hot Plug Controller on every port)
 - ECRC and Poison bit support
 - Data Path parity
 - Memory (RAM) Error Correction
 - INTA# and FATAL_ERR# signals
 - Advanced Error Reporting
 - Port Status bits and GPIO available
 - Per port error diagnostics
 - Performance Monitoring
 - Per port payload & header counters
 - JTAG AC/DC boundary scan



The *ExpressLane*[™] PEX 8648 device offers PCI Express switching capability enabling users to add scalable high bandwidth, non-blocking interconnection to a wide variety of applications including **servers, storage systems, and communications platforms**. The PEX 8648 is well suited for **fan-out, aggregation, and peer-to-peer** applications.

High Performance & Low Packet Latency

The PEX 8648 architecture supports packet **cut-thru with a maximum latency of 140ns (x16 to x16)**. This, combined with large packet memory 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 PEX 8648 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 as packets pass through the switch.

Flexible Register & Port Configuration

The PEX 8648's 12 ports can be configured to lane widths of x1, x2, 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. The PEX 8648 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 1 shows some of the PEX 8648's common port configurations.




Figure 1. Common Port Configurations

Dual-Host & Failover Support

The PEX 8648 product supports a **Non-Transparent (NT) Port**, which enables the implementation of **multi-host systems** in communications, storage, and blade server applications. 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 2).




Figure 2. Non-Transparent Port

Dual Cast

The PEX 8648 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 dual-graphics, storage, security, and redundant applications.

Read Pacing

The Read Pacing feature allows users to throttle the amount of read requests being made by downstream devices. When 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.

Hot Plug for High Availability

Hot plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8648 hot plug capability feature makes it suitable for **High Availability (HA) applications**. Three downstream ports include a Standard Hot Plug Controller. If the PEX 8648 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 8648 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 PEX 8648 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 PEX 8648 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, it supports **auto-negotiation**, **lane reversal**, and **polarity reversal**. Furthermore, the PEX 8648 is tested 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**.

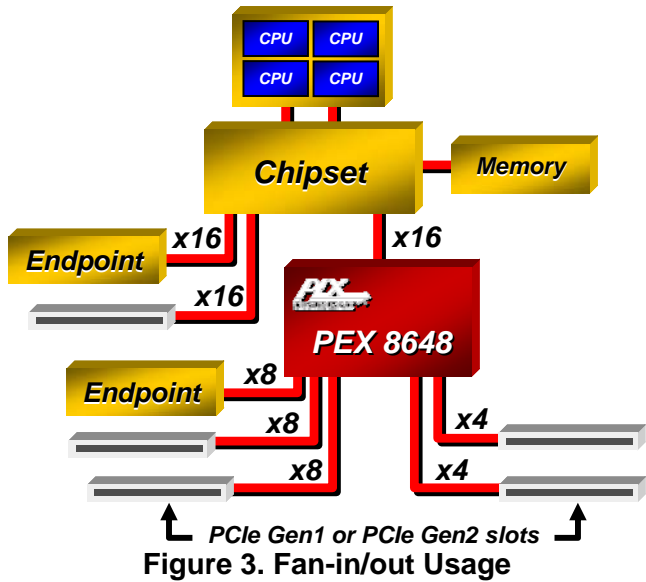
Applications

Suitable for **host-centric** as well as **peer-to-peer traffic patterns**, the PEX 8648 can be configured for a wide variety of form factors and applications.

Host Centric Fan-out

The PEX 8648, with its symmetric or asymmetric lane configuration capability, allows user-specific tuning to a variety of host-centric applications. Figure 3 shows a typical **server** design where the root complex provides a PCI Express link that needs to be expanded to a larger number of smaller ports for a variety of I/O functions. In this example, the PEX 8648 has a 16-lane upstream port, and five downstream ports using x8 and x4 links.

The PEX 8648 can also be used to create PCIe Gen1 (2.5 Gbps) ports. The PEX 8648 is backwards compatible with PCIe Gen1 devices. Therefore, the PEX 8648 enables a Gen 2 native Chip Set to fan-out to Gen 1 endpoints. In Figure 3, the PCIe slots connected to the PEX 8648's downstream ports can be populated with either PCIe Gen1 or PCIe Gen 2 devices. Conversely, the PEX 8648 can also be used to create Gen 2 ports on a Gen 1 native Chip Set in the same fashion.



PEX 8648 utilizes its NT port to isolate the hosts on the switch fabrics and the remaining 10 downstream ports on each switch are used to fan-out to the 20 AMC blades (or 10 AMC Carrier Modules).

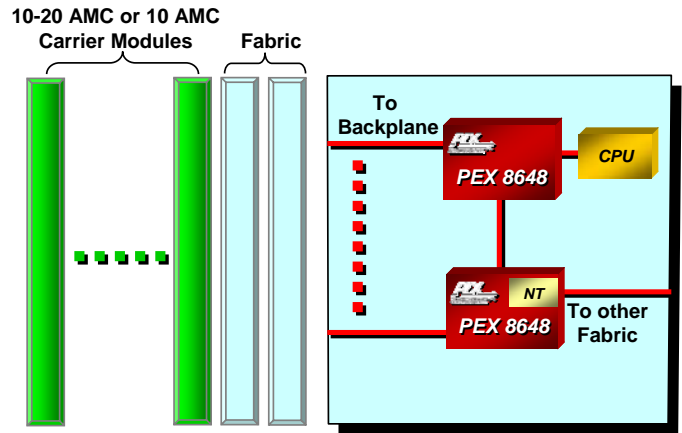


Figure 5. ATCA/MicroTCA Backplane Usage

Communications Systems

The PEX 8648 can also be utilized in communications applications. Figure 4 shows a PEX 8648 being used in a router to connect multiple endpoints in a line card to the control module using PowerPC. The link-widths for each module can be configured as required. The peer-to-peer communication feature of the PEX 8648 allows the endpoints to communicate with each other without any intervention or management by the host.

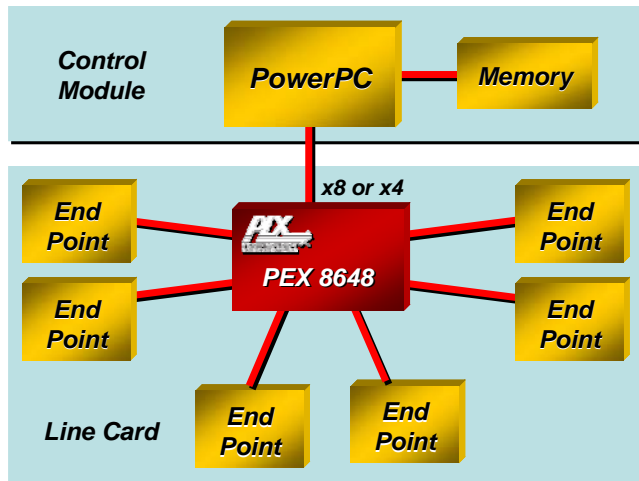


Figure 4. Router Usage

Backplane Communication

With 48 lanes and 12 ports, the PEX 8648 is well suited for backplane applications requiring high connectivity (a large number of ports). Figure 5 represents an ATCA or MicroTCA backplane application with two switch fabric blades and multiple AMC blades. In this example, two PEX 8648s provide peer-to-peer data exchange for up to 20 AMC blades connecting to the switch fabric. The

Fail-Over Storage Systems with Dual Cast

The PEX 8648's Dual Cast feature proves to be very useful in storage systems. In the example shown in Figure 6, the Dual Cast feature enables the PEX 8648 to copy data coming from the host to two downstream ports (see yellow traffic patterns) in one transaction as opposed to having to execute two separate transactions to send data to the redundant chassis. By offloading the task of backing up data onto the secondary system, processor and system performance is enhanced. The PEX 8648's NT port is used to isolate the backup system from the primary system.

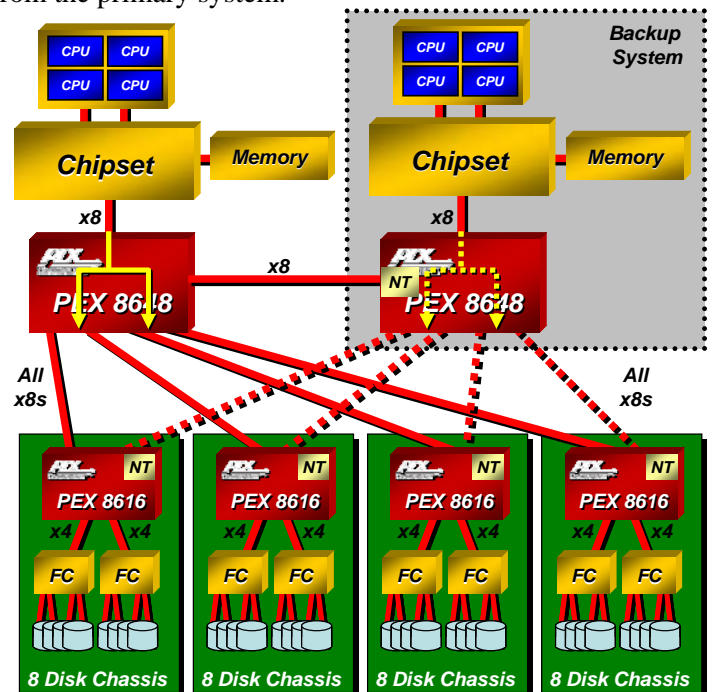


Figure 6. Dual Cast in Storage Systems

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 8648 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 8648 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 8648 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 (PEX 8648 RDK), hardware documentation (available at www.plxtech.com), and a Software Development Kit (also available at www.plxtech.com).

ExpressLane PEX 8648 RDK

The PEX 8648 RDK is a hardware module containing the PEX 8648 which plugs right into your system. The PEX 8648 RDK can be used to test and validate customer software, or used as an evaluation vehicle for PEX 8648 features and benefits. The PEX 8648 RDK provides everything that a user needs to get their hardware and software development started. For more information, please refer to the PEX 8648 RDK Product Brief.

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 PEX 8648. For more information, please refer to the PEX 8648 RDK Product Brief.



PLX Technology, Inc.
870 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 Site: www.plxtech.com

Product Ordering Information

Part Number	Description
PEX8648-AA50BC G	48-Lane, 12-Port PCI Express Switch, Pb-Free (27x27mm ²)
PEX8648-AA RDK	PEX 8648 Rapid Development Kit

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

© 2007 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.

PEX8648-SIL-PB-P1-0.8

09/07