

PEX 8532 & PEX 8516

Features

- **PEX 8532 General Features**
 - 32-lane PCI Express switch
 - Integrated SerDes
 - Eight configurable ports (x1, x2, x4, x8, x16)
 - 35mm x35mm, 680 pin enhanced PBGA package, 6.5 Watts
- **PEX 8516 General Features**
 - 16-lane PCI Express switch
 - Integrated SerDes
 - Four configurable ports (x1, x2, x4, x8)
 - 27mm x 27mm, 312 pin enhanced PBGA package, 3.5 Watts
- **PEX 8532 & PEX 8516 Common Features**
 - Standards Compliant
 - PCI Express Base Specification, r1.0a
 - PCI Standard SHPC Specification, r1.1
 - High Performance
 - Non-blocking Switch Fabric
 - Full Line rate
 - Configurable **Non-transparent** port for Multi-Host or Intelligent I/O Support
 - Flexible Configuration
 - Eight (four for PEX8516) highly flexible and configurable ports
 - Flexible lane width/port x2(x1), x4, x8, (x16 for PEX 8532 only)
 - Configurable with strapping pins, EEPROM, 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 D3hot
 - Quality of Service (QoS)
 - **Two Virtual Channels**/port
 - Eight Traffic Classes/port
 - Fixed and WRR Virtual Channel arbitration
 - Round robin port arbitration
 - Reliability, Availability, Serviceability (RAS)
 - Standard Hot-Plug Controller
 - Upstream port as hot-plug client
 - Transaction Layer **end-to-end CRC**
 - Poison bit
 - Error reporting in addition to advanced error reporting support of PCI Express
 - Per port performance monitoring
 - Average packet size, number of packets, CRC errors
 - JTAG boundary scan

Flexible & Versatile PCI Express™ Switches

Multi-purpose and Feature Rich PCI Express Switch Family

The **32-lane PEX 8532** and **16-lane PEX 8516** products offer PCI Express **switching** capability conforming to the latest revision of the PCI Express Base specification. These products enable users to add scalable high bandwidth, **non-blocking** interconnection to a wide variety of applications including servers, storage systems, communications platforms, blade servers, and embedded-control products. The PEX 8532/16 products can be used as **fanout, aggregation, or peer-to-peer** switches, and are equally well-suited to **fabric backplane** and **intelligent I/O** module applications.

Highly Flexible Port Configurations

The PEX 8532/16 offer highly configurable ports. There are a maximum of **8 ports** (4 for PEX 8516) that can be configured to any legal width from x1 to x16 (x8 max for PEX 8516), 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 family 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 PEX product line.

Non-Transparent “Bridging” in a PCI Express Switch

The PEX 8532/16 product family supports full non-transparent bridging functionality to allow implementation of multi-host systems and intelligent I/O modules in applications such as **communications, storage, and blade servers**. 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 8532/16 products support 2 full-featured Virtual Channels (VCs) and a full **8 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 8532/16 provide **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.



Flexible Port Width 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 8532/16 support a large number of port configurations. For example, if you are using the PEX 8532 in a fan-out application (such as in Figure 1), you may configure the upstream port as x8 and the downstream ports as six x4 ports; two x8 & two x4 ports; three x8 ports; or **any other combination** as long as you don't run out of lanes or ports. For a peer-to-peer application, you can configure all eight ports as x4 or x2, or a combination of the two. In a port aggregation application you can configure four x2 or x4 ports for aggregation into one x8 or x16 port.

Hot Plug for High Availability

Hot plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8532/16 hot plug capabilities and advanced error reporting features make them suitable for High Availability (HA) applications. Each downstream port includes a Standard Hot Plug Controller. If the PEX 8532/16 is used in an application where one or more of its downstream ports connects 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 fully compliant 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 8532/16 devices support both link (L0, L0s, L1, L2/L3 Ready, and L3) and device (D0 and D3hot) power management states, in compliance with the PCI Express power management specification.

SerDes Power and Signal Management

The PEX 8532/15 family 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 8532/16 switches support fixed **Round Robin** and **Weighted Round Robin** (32 phase) arbitration schemes for both (2) virtual channels. This allows fine tuning of Quality of Service and efficient use of the system bandwidth.

Applications

Suitable for **host-centric** as well as **peer-to-peer** traffic patterns, the PEX 8532/16 can be configured for a wide variety of form factors and applications.

Host Centric Fan-out

The PEX 8532/16 devices, with their versatile symmetric or asymmetric lane configuration capability, allow application specific tuning to a variety of host-centric applications.

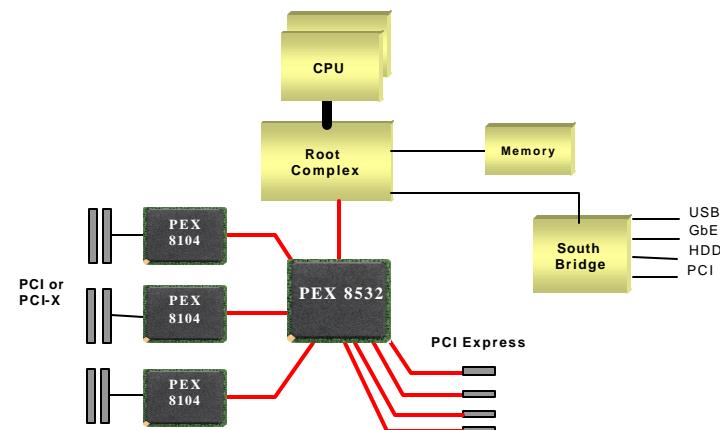


Figure 1. Fan-in/out Usage

Figure 1 shows a typical **server-based** design, where the root complex provides a PCI Express link that needs to be broken into a larger number of smaller ports for a variety of I/O functions with different bandwidth requirements.

In this example, the PEX 8532 would typically have an 8-lane upstream port, and as many as **7 downstream ports**. The downstream ports can be of differing widths if required. (perhaps several x4 ports and several x1 ports). The figure also shows how some of the ports can be bridged to provide **PCI or PCI-X** through the use of the **PEX 8104 PCI Express to PCI-X** device.

Peer-to-Peer & Backplane Usage

Figure 2 represents a backplane where the PEX 8532 provides peer-to-peer data exchange for a large number of line cards where the CPU/Host plays the management role.

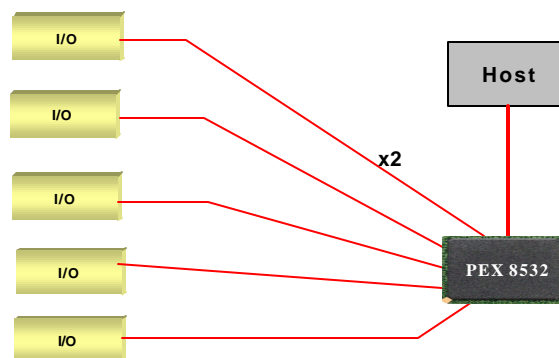


Figure 2. Peer-to-Peer/Backplane Usage

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.




Figure 3. Aggregation Adapter Card

The PEX 8516 can be used to create an adapter or mezzanine 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 8516 in this usage.

The adapter card in Figure 3 can be **transparent**, in which case the PCI Express devices are just standard I/O products such as Ethernet, Fibre Channel, etc. Or the PEX 8516 can provide a **non-transparent** port to the system (via the card edge). In this case, one of the PCI Express devices can be a CPU or other “intelligent” device with on-chip processing capability – thus needing address domain isolation from the rest of the system. This approach is commonly used in RAID controllers.

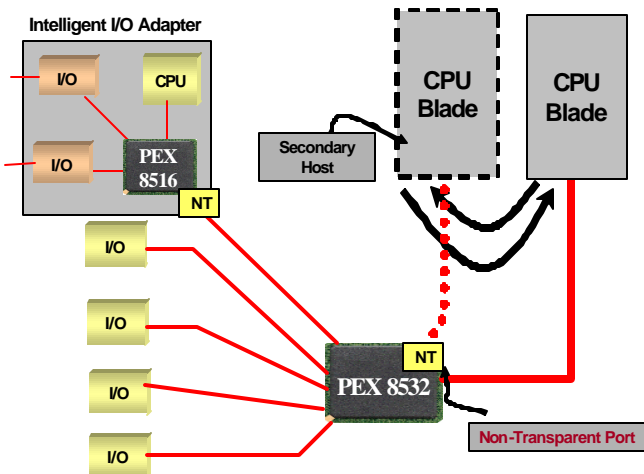


Figure 4. Dual Host Usage

Dual Host Model

The PEX 8532/16 support applications requiring **dual host**, **host failover**, and **load-sharing** applications through the **non-transparent** feature. Figure 4 illustrates a dual host system using an intelligent adapter card.

In this figure, the **primary** CPU (on the right) is the active host; it configures and enumerates the system, and handles interrupts and error conditions. If the primary host ceases proper operation, the secondary host takes over the system. The PEX 8532/16 can **dynamically re-assign** both the **upstream** port (from primary to secondary) and the **non-transparent** port (from the secondary to the primary), and allow the system to continue operation.

Dual Fabric Model

High performance **communications**, **storage**, and **blade server** systems often require more reliability than a dual host system provides. For these **high availability** systems, all single points of failure must be eliminated. The PEX 8532 can offer this additional level of redundancy by linking together 2 switch fabric cards in a **dual-star** topology.

Figure 5 shows an example of such a system. In this system, both the host and fabric are on separate cards, and both can be active simultaneously. The unique **non-transparent** feature of the PEX 8532 is used to allow one fabric/host card to be the primary and the other to be the secondary.

This approach can be used to provide more than two active processing nodes. It is straightforward to create a **generalized multiprocessor system** with smaller PEX 8516 switches on each card in non-transparent mode, and several PEX 8532 fabric cards to provide the fabric backbone.

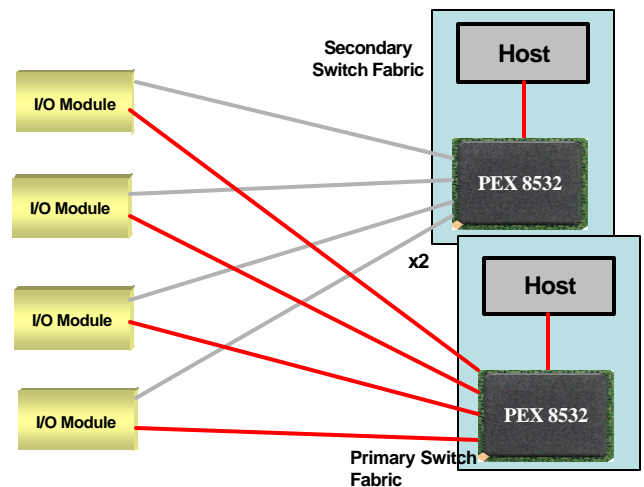


Figure 5. Dual Host & Fabric Usage

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. The PEX 8532/16 uses port 0 as the upstream port by default, but any port can be configured as the upstream port by using an optional EEPROM, or through strapping pins. It is through the upstream port that the BIOS or host can configure the other ports using standard PCI enumeration.

System Configuration

The virtual PCI to PCI bridges within the PEX 8532/16 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 8532/16 switches support 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 8532/16 for hot plug events, baseline error reporting, and advanced error reporting.

Development Tools

PLX is offering hardware and software tools (PEX 8532RDK) to enable rapid customer design activity. These tools include the hardware Rapid Development Kit (RDK) and the Software Development Kit (SDK).

PEX 8532RDK

The RDK hardware module includes the PEX 8532 with one x16 (card-edge slot) port, two x4 ports and one x8 port (slot connectors). The RDK is available with x16, x8 and x4 connectors (card-edge type).

The PEX 8532RDK can be installed in a motherboard, used as a riser card, or configured as a bench-top board.

The PEX 8532RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for the PEX 8532 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

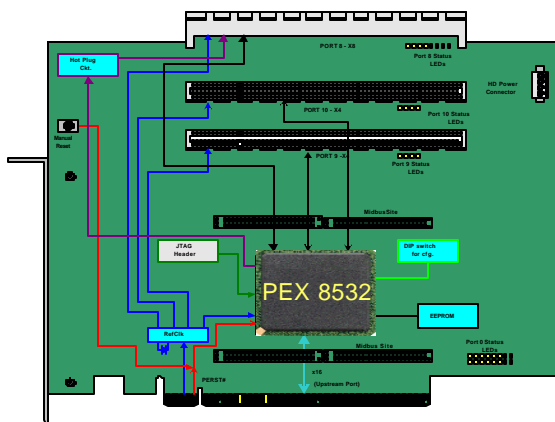


Figure 6. PEX 8532RDK



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

© 2004 PLX Technology, Inc. All rights reserved. PLX and the PLX logo are registered trademarks of PLX Technology, Inc. PowerDrive and the PowerDrive logo are trademarks 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.

Product Ordering Information

Part Number	Description
PEX 8532AA-25BI	32-Lane PCI Express Switch
PEX 8516AA-25BI	16-Lane PCI Express Switch
PEX 8532RDK	Rapid Development Kit for PEX 8532

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