

PEX 8532

Features

- **PEX 8532 General Features**
 - 32-lane PCI Express switch
 - Integrated SerDes
 - Up to eight configurable ports
 - (x1, x2, x4, x8, x16)
 - 35mmx35mm, 680 pin PBGA package
 - Typical Power: 5.7 Watts
- **PEX 8532 Key Features**
 - **Standard Compliant**
 - PCI Express Base Specification, r1.1
 - **High Performance**
 - Non-blocking switch fabric
 - Full line rate on all ports
 - **Non-Transparent Bridging**
 - Configurable Non-Transparent port for Multi-Host or Intelligent I/O Support
 - **Flexible Configuration**
 - Eight highly flexible & configurable ports (x1, x2, x4, x8, or x16)
 - 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 per port
 - Eight Traffic Classes per port
 - Fixed and Round-Robin Virtual Channel Port Arbitration
 - **Reliability, Availability, Serviceability (RAS)**
 - Eight Standard Hot-Plug Controllers
 - Upstream port as hot-plug client
 - Transaction Layer end-to-end CRC
 - Poison bit
 - Advanced Error Reporting
 - Per port performance monitoring
 - Average packet size
 - Number of packets
 - CRC errors and more
 - JTAG boundary scan

Flexible & Versatile PCI Express™ Switch

Multi-purpose, Feature Rich ExpressLane™ PCI Express Switch

The *ExpressLane™* PEX 8532 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, communications platforms, blade servers, and embedded-control products**. The PEX 8532 is well suited for **fan-out, aggregation, dual-graphics, peer-to-peer, and intelligent I/O module** applications.

Highly Flexible Port Configurations

The *ExpressLane* PEX 8532 offers highly configurable ports. There are a maximum of 8 ports that can be configured to any legal width from x1 to x16, in any combination to support your specific bandwidth needs. The ports can be configured for **symmetric** (each port having the same lane width and traffic load) or **asymmetric** (ports having different lane widths) traffic. In the event of asymmetric traffic, the PEX 8532 features a **flexible central packet memory** that allocates a memory buffer for each port as required by the application or endpoint. This buffer allocation along with the device's **flexible packet flow control** minimizes bottlenecks when the upstream and aggregated downstream bandwidths do not match (are asymmetric). Any of the ports can be designated as the upstream port, which can be changed dynamically.

End-to-end Packet Integrity

The PEX 8532 provides **end-to-end CRC** protection (ECRC) and **Poison bit** support to enable designs that require **end-to-end data integrity**. 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 8532 product 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 *ExpressLane* PEX 8532 switch supports 2 full-featured Virtual Channels (VCs) and 8 Traffic Classes (TCs). The mapping of Traffic Classes to port-specific Virtual Channels allows for different mappings on 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 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.



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 supports 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 (32) or ports (8). 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 ExpressLane PEX 8532 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 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 ExpressLane PEX 8532 device 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 ExpressLane PEX 8532 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 ExpressLane PEX 8532 switch supports hardware fixed and Round Robin arbitration schemes for two virtual channels. This allows fine tuning of Quality of Service, optimum use of buffers and efficient use of the system bandwidth.

Applications

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

Host Centric Fan-out

The ExpressLane PEX 8532 device, with its versatile symmetric or asymmetric lane configuration capability, allows user specific tuning to a variety of host-centric applications.

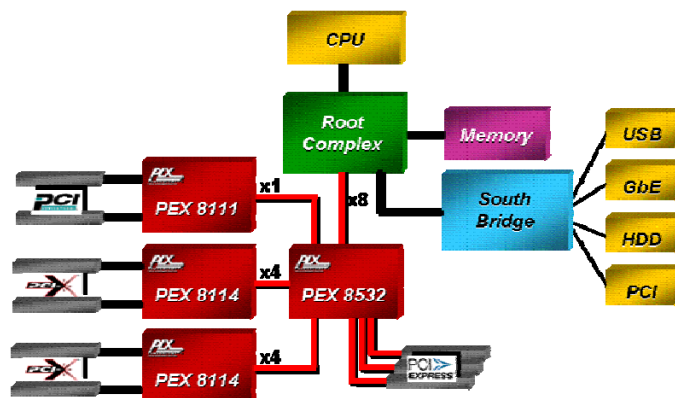


Figure 1. Fan-in/out Usage

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

The ExpressLane 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. The figure also shows how some of the ports can be bridged to provide **PCI or PCI-X** through the use of the ExpressLane PCIe bridges such as **PEX 8114 and PEX 8111**.

Peer-to-Peer & Backplane Usage

Figure 2 represents a backplane where the ExpressLane 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

Graphics Fan-out Switch

The number and variety of PCI Express native-mode devices is growing quickly. The devices such as graphics cards are expected to become mainstream very rapidly. As that happens, it will be necessary to take a wide (x16) graphics port on the root complex device and fan it out to four x4 or two x8 ports.

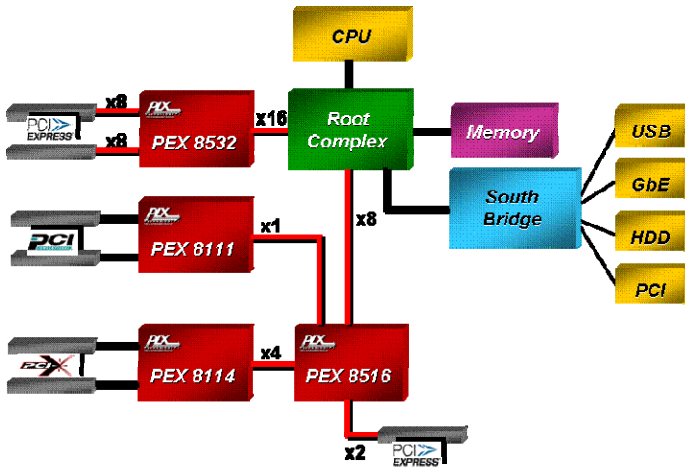


Figure 3. Graphics Fan-out

Dual Host Model

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




Figure 4. Dual Host Usage

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 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 single host system provides. For these **high availability** systems, all single points of failure must be eliminated. The ExpressLane 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-transparency** 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 ExpressLane 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

Switch Fabric Module

The ExpressLane PEX 8532 can also be used in Blade Server switch fabric applications. Figure 6 illustrate the use of PEX 8532 in a Benes Switch configuration. The non-transparency and peer-to-peer functions of PEX 8532 enable use of fewer total switching elements and less hops from source to destination.




Figure 6. Benes Switch Fabric

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 8532 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 ExpressLane PEX 8532 switch 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 for hot plug events, baseline error reporting, and advanced error reporting.

Development Tools

PLX is offering hardware and software tools (PEX 8532 RDK) to enable rapid customer design activity. These tools are bundled in a Rapid Development Kit (RDK). The RDK consists of hardware, hardware documentation and Software Development Kit (SDK).

ExpressLane PEX 8532 RDK

The RDK hardware module includes the PEX 8532 with a x16 upstream port, a x8 downstream port and two x4 downstream ports. The upstream port is a x16 PCI Express edge connector. PLX offers adapters (x8, x4, and x1) which can be used to plug the RDK in smaller slots.

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

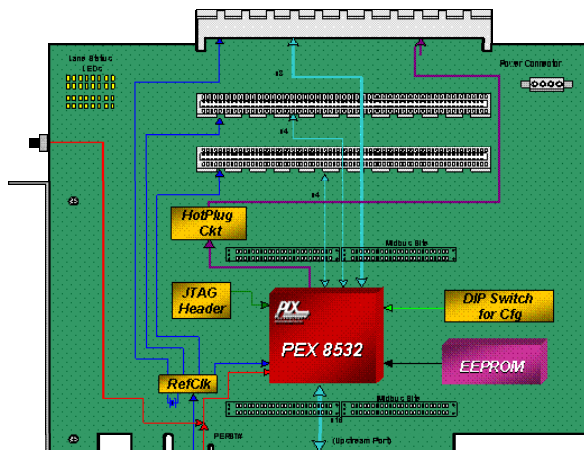


Figure 7. PEX 8532 RDK



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
PEX8532-BC25BI	32 Lane, 8 Port PCIe Switch, 680-ball PBGA 35x35mm pkg
PEX8532-BC25BI G	32 Lane, 8 Port PCIe Switch, 680-ball PBGA 35x35mm pkg, Pb-free
PEX 8532-BC RDK	PEX 8532 Rapid Development Kit
x8 Adapter	PCI Express x16 to x8 Adapter
x4 Adapter	PCI Express x16 to x4 Adapter
x1 Adapter	PCI Express x16 to x1 Adapter

Please visit the PLX Web site at <http://www.plxtech.com> 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.