

PEX 8524

Features

- **PEX 8524 General Features**
 - 24-lane PCI Express switch
 - Integrated SerDes
 - Up to six configurable ports
 - 31mm x31mm, 644 pin PBGA package
 - Typical Power: 3.9 Watts
- **PEX 8524 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**
 - Six 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**
 - 6 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 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*TM PCI Express Switch

The *ExpressLane*TM PEX 8524 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 8524 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 8524 offers highly configurable ports. There are a maximum of 6 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 8524 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 8524 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 *ExpressLane* PEX 8524 product supports full non-transparent bridging (NTB) functionality to allow implementation of **multi-host systems and intelligent I/O modules in communications, storage, blade server, and graphics fan-out applications**. 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 8524 switch supports 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 8524 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 8524 supports a large number of port configurations. For example, if you are using the PEX 8524 in a fan-out application, you may configure the upstream port as x8 and the downstream ports as four x4 ports; two x8 ports for dual-graphics fan-out; or other combinations, as long as you don't run out of lanes (24) or ports (6). In a peer-to-peer application you can configure all six ports as x4. Figure 1 shows the 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 8524 hot plug capability and **Advanced Error Reporting** features makes it suitable for **High Availability (HA) applications**. Each downstream port includes a Standard Hot Plug Controller. If the PEX 8524 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 8524 device supports 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 8524 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 8524 switches support **hardware fixed and Round Robin arbitration schemes** for two virtual channels on each port. This allows for the fine tuning of Quality of Service for efficient use of packet buffers and system bandwidth.

Applications

Suitable for **host-centric** as well as **peer-to-peer traffic patterns**, the PEX 8524 can be configured for a wide variety of form factors and applications.

Host-Centric Fan-out

The PEX 8524, with its symmetric or asymmetric lane configuration capability, allows user specific tuning to a variety of host-centric applications.

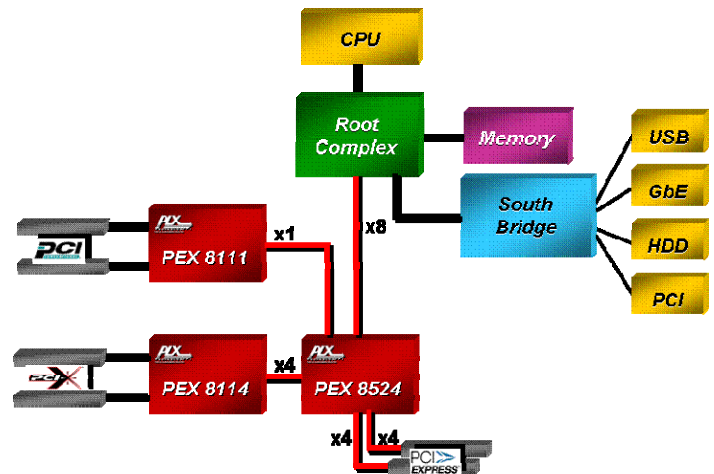


Figure 2. Fan-in/out Usage

Figure 2 shows a typical **server-based** design, where the root complex provides a PCI Express link that needs to be fanned into a larger number of smaller ports for a variety of I/O functions, each with different bandwidth requirements.

In this example, the PEX 8524 would typically have an 8-lane upstream port, and as many as 5 downstream ports (4 shown here). 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** slots through the use of the *ExpressLane* **PEX 8114 and PEX 8111** PCIe bridging devices.

Almost all (non x86 based) high-end microprocessor manufacturers are offering PCI Express interfaces. The PEX 8524 can be directly connected to a processor to fan-out its PCIe port to a larger number of ports for enhanced connectivity as illustrated in Figure 3.




Figure 3. Fan out for PowerPC/MIPS CPUs

Embedded Systems

The PEX 8524 can also be utilized in embedded applications. Figure 4 shows several independent modules connecting through the PEX 8524. The port widths for each module can be configured as required. The peer-to-peer communication feature of the PEX 8524 allows these modules to communicate with each other without any centralized control.

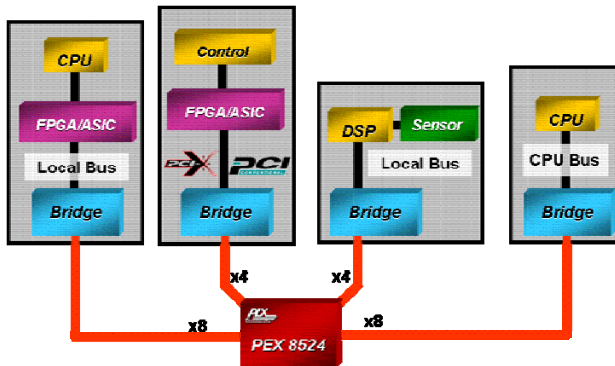


Figure 4. Embedded Systems

Peer to Peer Communication

Figure 5 represents a backplane where the ExpressLane PEX 8524 provides peer-to-peer data exchange for a large number of line cards where the CPU/Host plays the management role.




Figure 5. Peer-to-Peer Usage

Graphics Fan-out Switch

As PCIe based graphics cards become more mainstream, it will be necessary to take a x8 port on the root complex device and fan it out to two x8 ports for dual graphics applications. Root Complex (Northbridge) devices are available with multiple PCIe ports. These ports can be further expanded to connect to a larger number of I/Os or to support dual-graphics using the PEX 8524 as shown in Figure 6.

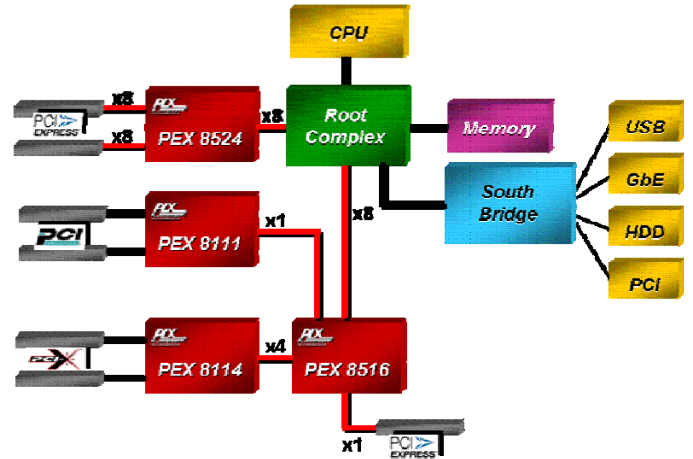


Figure 6. Graphics Fan-out

PCI Express Port Expansion

The PEX 8524 enables designers to take, for example, two x8 PCIe ports and expand them into ten ports. Some of these PCIe ports can be bridged to PCI or PCI-X using bridging products from PLX. Figure 7 illustrates one of the many configurations the PEX 8524 can support.

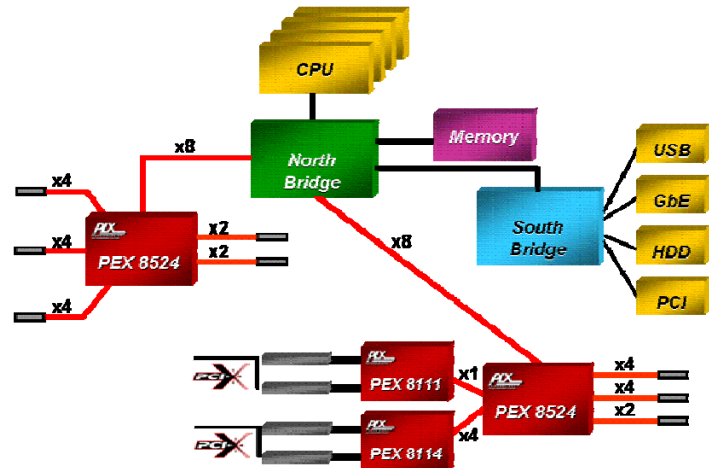


Figure 7. PCIe Port Expansion

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 8524 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 8524 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 8524 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 8524RDK), hardware documentation, and a Software Development Kit (SDK).




Figure 8. PEX 8524RDK


ExpressLane PEX 8524RDK

The RDK hardware module includes the PEX 8524 with one x8 upstream port, one x8 downstream port, and two x4 downstream ports. The upstream port uses a x8 PCI Express edge connector. PLX offers adapters (x4 and x1) which can be used to plug the RDK in smaller slots. The PEX 8524RDK hardware module can be installed in a motherboard, used as a riser card, or configured as a bench-top board. The PEX 8524RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for PEX 8524 features and benefits.

SDK

The SDK tool set includes:

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



Product Ordering Information

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

Part Number	Description
PEX8524-BC25BI	24-Lane, 6-Port PCIe Switch, 644-ball PBGA 31x31mm ² package
PEX8524-BC25BI G	24-Lane, 6-Port PCIe Switch, 644-ball PBGA 31x31mm ² package, Pb-free
PEX 8524-BC RDK	PEX 8524 Rapid Development Kit with x8 Edge Connector
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.