

PEX 8518

Flexible & Versatile PCI Express[®] Switch

Features

- **PEX 8518 General Features**
 - 16-lane PCI Express switch
 - Integrated SerDes
 - Up to five configurable ports
 - 23mm x 23mm, 376-ball PBGA package
 - Typical Power: 2.6 Watts
- **PEX 8518 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, x4, or x8)
 - 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

Multi-purpose, Feature Rich ExpressLane™ PCI Express Switch

The *ExpressLane* PEX 8518 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 8518 is well suited for **fan-out, aggregation, peer-to-peer, and intelligent I/O module** applications.

Highly Flexible Port Configurations

The PEX 8518 offers highly configurable ports. There are a maximum of five ports that can be configured to any legal width from x1 to x8, 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 8518 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 8518 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 8518 supports full non-transparent bridging (NTB) 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 8518 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 8518 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 8518 supports a large number of port configurations. For example, if you are using the PEX 8518 in a fan-out application, you may configure the upstream port as x8 and the downstream ports as four x2 ports; two x2 & one x4 ports; three x1 ports; or other combinations, as long as you don't run out of lanes (16) or ports (5). In a port aggregation application you can configure three x2 or x1 ports for aggregation into one x8 or x4 port. 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 8518 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 8518 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 8518 device supports 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.

SerDes Power and Signal Management

The ExpressLane PEX 8518 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 8518 switch supports **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 **intelligent I/O applications** *ExpressLane* PEX 8518 can be configured for a wide variety of form factors and applications.

Host-Centric Fan-out

The *ExpressLane* PEX 8518 switch, with its symmetric or asymmetric lane configuration capability, allows user specific tuning to a variety of **host-centric as well as peer-to-peer applications**.

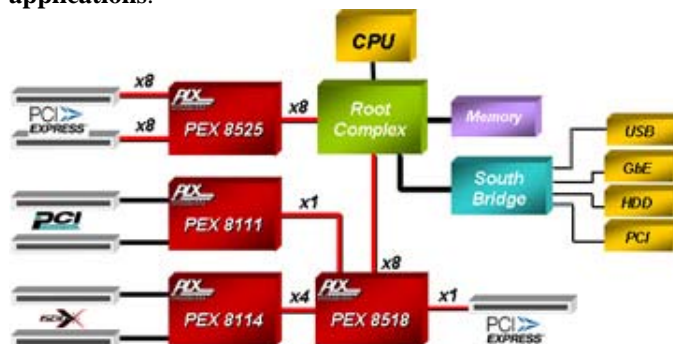


Figure 2. Fan-in/out Usage

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 8518 would typically have an 8-lane upstream port, and as many as three 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** slots through the use of the *ExpressLane* **PEX 8114 and PEX 8111** PCIe bridging devices.

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 8518 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 motherboard. Figure 3 shows the PEX 8518 in this application.




Figure 3. Aggregation Adapter Card

The adapter card in Figure 3 can be **transparent**, in which case the PCI Express I/O devices are just standard I/O endpoints such as Ethernet or Fibre Channel. Or the PEX 8518 can provide a **non-transparent** port to the system (via the card's edge connector). 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**.

Intelligent Adapter Card

The PEX 8518 supports the **non-transparency** feature. Figure 4 illustrates a host system using an intelligent adapter card.




Figure 4. Intelligent Adapter Usage

In this figure, the CPU on the adapter card is isolated from the host CPU. The PEX 8518 non-transparent port allows the two CPUs to be isolated but communicate with each other through various registers that are designed in the PEX 8518 for that purpose. The host CPU **can dynamically re-assign both the upstream port and the non-transparent port** of PEX 8518 allowing the system to be reconfigured.

Dual Host/Fabric Model

The *ExpressLane* PEX 8518 supports applications requiring **dual host, host failover, and load-sharing** applications through the **non-transparency** feature. Figure 5 illustrates a dual host system with dual switch fabric in dual-star configuration.

The redundancy of the host and the fabric can be achieved through many possible configurations using NTB function of PEX 8518. In the configuration shown below the host 1 controls the switch 1 and associated I/Os and the host 2 controls the switch 2 and associated I/Os. The hosts and switches are isolated using NTB functionality of PEX 8518 on the host boards. If one of the hosts fails the surviving host can remove the failing host from the configuration while controlling both the switches and all I/Os. Similarly, if one of the switches fails the host associated with that switch can send control messages to its I/Os through the surviving switch using NTB function.

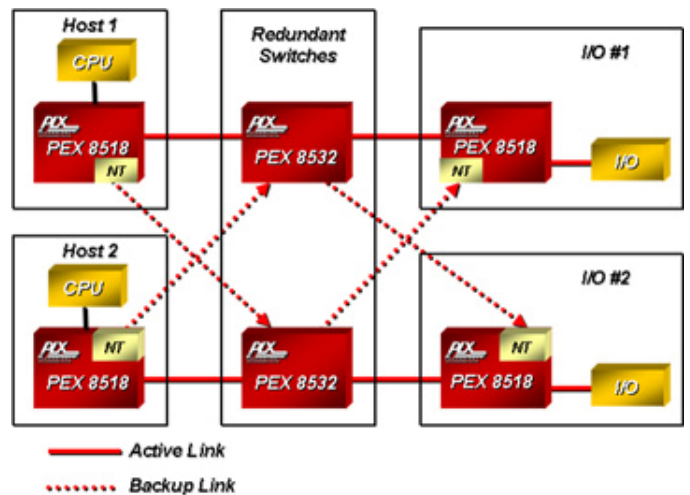


Figure 5. Dual Host/Fabric Mode

Embedded Systems

The PEX 8518 can also be utilized in embedded applications. Figure 6 shows several independent modules connecting through the PEX 8518. The port widths for each module can be configured as required. The Peer-to-peer communication feature of the PEX 8518 allows these modules to communicate with each other without any centralized control.




Figure 6. Embedded 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 8518 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 8518 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 8518 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 PEX 8518 Rapid Development Kit (RDK), hardware documentation, and a Software Development Kit (SDK).

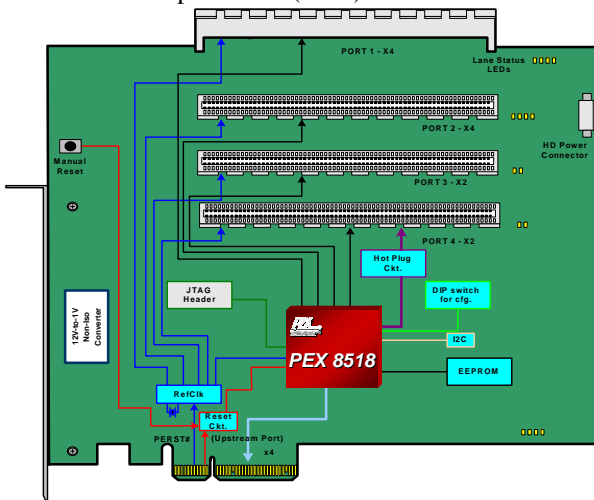


Figure 7. PEX 8518RDK


RDK

The RDK hardware module includes the PEX 8518 with one x4 (card-edge slot) port and three additional x4 ports (see Figure 7). The RDK is available with x4 card edge connector and adapters for x1 edge connectors are available to plug the RDK into smaller slots. The PEX 8518RDK hardware module can be installed in a motherboard, used as a riser card, or configured as a bench-top board. The PEX 8518RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for PEX 8518 features and benefits.

SDK

The SDK tool set includes:

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



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
PEX8518-AC25BI	16 Lane, 5 Port PCIe Switch, 376-ball PBGA 23x23mm pkg
PEX8518-AC25BI G	16 Lane, 5 Port PCIe Switch, 376-ball PBGA 23x23mm pkg, Pb-free
PEX 8518RDK	PEX 8518 Rapid Development Kit with x4 Connector

Please visit the PLX Web site at <http://www.plxtech.com/8518> or contact PLX sales at 408-774-9060 for sampling.