

PEX 8636, PCI Express Gen 2 Switch, 36 Lanes, 24 Ports

Highlights

■ PEX 8636 General Features

- 36-lane, 24-port PCIe Gen2 switch
 - Integrated 5.0 GT/s SerDes
- 35 x 35mm², 1156-ball FCBGA package
- Typical Power: 8.8 Watts

■ PEX 8636 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

◆ performancePAK

- ✓ Read Pacing (bandwidth throttling)
- ✓ Multicast
- ✓ Dynamic Buffer/FC Credit Pool
- Non-blocking switch fabric
- Full line rate on all ports
- Packet Cut-Thru with 200ns max packet latency (x1 to x1)
- 2KB Max Payload Size

○ Flexible Configuration

- Ports configurable as x1, x4, x8, x16
- Registers configurable with strapping pins, EEPROM, I²C, or host software
- Lane and polarity reversal
- Compatible with PCIe 1.0a PM

○ Multi-Host & Fail-Over Support

- Configurable Non-Transparent (NT) port
- Failover with NT port
- Up to Eight upstream/Host ports with 1+1 or N+1 failover to other upstream ports

○ Quality of Service (QoS)

- Eight traffic classes per port
- Weighted round-robin source port arbitration

○ Reliability, Availability, Serviceability

◆ visionPAK

- ✓ Per Port Performance Monitoring
 - Per port payload & header counters
- ✓ SerDes Eye Capture
- ✓ Error Injection and Loopback
- 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
- JTAG AC/DC boundary scan

The ExpressLane™ PEX 8636 device offers Multi-Host PCI Express switching capability enabling users to connect multiple hosts to their respective endpoints via scalable, high bandwidth, non-blocking interconnection to a wide variety of applications including communications platforms. The PEX 8636 is well suited for fan-out, aggregation, and peer-to-peer applications.

Multi-Host Architecture

The PEX 8636 employs an enhanced version of PLX's field tested PCIe switch architecture, which allows users to configure the device in legacy single-host mode or multi-host mode with up to Eight host ports capable of 1+1 (one active & one backup) or N+1 (N active & one backup) host failover. This powerful architectural enhancement enables users to build PCIe based systems to support high-availability, failover, redundant and clustered systems.

High Performance & Low Packet Latency

The PEX 8636 architecture supports packet cut-thru with a maximum latency of 200ns (x1 to x1). This, combined with large packet memory, flexible common buffer/FC credit pool 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 8636 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 circuitry throughout the internal data paths as packets pass through the switch.

Flexible Configuration

The PEX 8636's 24 Ports can be configured to lane widths of x1, 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. Figure 1 shows some of the PEX 8636's common port configurations in legacy Single-Host mode.




Figure 1. Common Port Configurations

PEX 8636, PCI Express Gen 2 Switch, 36 Lanes, 24 Ports

The PEX 8636 can also be configured in Multi-Host mode where users can choose up to Eight ports as host/upstream ports and assign a desired number of downstream ports to each host. In Multi-Host mode, a virtual switch is created for each host port and its associated downstream ports inside the device. The traffic between the ports of a virtual switch is completely isolated from the traffic in other virtual switches. Figure 2 illustrates some configurations of the PEX 8636 in Multi-Host mode where each ellipse represents a virtual switch inside the device.

The PEX 8636 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 2. Common Multi-Host Configurations

Dual-Host & Failover Support

In Single-Host mode, the PEX 8636 supports a **Non-Transparent (NT) Port**, which enables the implementation of **dual-host systems** for redundancy and host failover capability. 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 3).

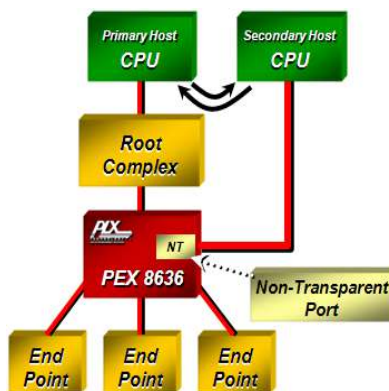


Figure 3. Non-Transparent Port

Multi-Host & Failover Support

In Multi-Host mode, PEX 8636 can be configured with up to Eight upstream host ports, each with its own dedicated downstream ports. The device can be configured for 1+1 redundancy or N+1 redundancy. The PEX 8636 allows the hosts to communicate their status to each other via special door-bell registers. In failover mode, if a host fails, the host designated for failover will disable the upstream port attached to the failing host and program the downstream ports of that host to its own domain. Figure 4a shows a two host system in Multi-Host mode with two virtual switches inside the device and Figure 4b shows Host 1 disabled after failure and Host 2 having taken over all of Host 1's end-points.




Figure 4a. Multi-host

Figure 4b. Multi-host Fail-over

Hot Plug for High Availability

Hot plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8636 hot plug capability feature makes it suitable for **High Availability (HA) applications**. Four downstream ports include a Standard Hot Plug Controller. If the PEX 8636 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 8636 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 8636 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.

PEX 8636, PCI Express Gen 2 Switch, 36 Lanes, 24 Ports

Interoperability

The PEX 8636 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** on its wide ports. Furthermore, the PEX 8636 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**.

performancePAK™

Exclusive to PLX, *performancePAK* is a suite of unique and innovative performance features which allows PLX's Gen 2 switches to be the highest performing Gen 2 switches in the market today. The *performancePAK* features consists of the Read Pacing, Multicast, and Dynamic Buffer Pool.

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 that downstream port. If that 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.

Multicast

The Multicast feature enables the copying of data (packets) from one ingress port to multiple (up to 23) egress ports in one transaction allowing for higher performance in dual-graphics, storage, security, and redundant applications, among others. Multicast relieves the CPU from having to conduct multiple redundant transactions, resulting in higher system performance.

Dynamic Buffer Pool

The PEX 8636 employs a dynamic buffer pool for Flow Control (FC) management. As opposed to a static buffer scheme which assigns fixed, static buffers to each port, PLX's dynamic buffer allocation scheme utilizes a common pool of FC Credits which are shared by other ports. This shared buffer pool is fully programmable by the user, so FC credits can be allocated among the ports as needed. Not only does this prevent wasted buffers and inappropriate buffer assignments, any unallocated buffers remain in the common buffer pool and can then be used for faster FC credit updates.

visionPAK™

Another PLX exclusive, *visionPAK* is a debug diagnostics suite of integrated hardware and software instruments that users can use to help bring their systems to market faster. *visionPAK* features consist of Performance Monitoring, SerDes Eye Capture, Error Injection, SerDes Loopback, and more.

Performance Monitoring

The PEX 8636's real time performance monitoring allows users to literally "see" ingress and egress performance on each port as traffic passes through the switch using PLX's Software Development Kit (SDK). The monitoring is completely passive and therefore has no affect on overall system performance. Internal counters provide extensive granularity down to traffic & packet type and even allows for the filtering of traffic (i.e. count only Memory Writes).

SerDes Eye Capture

Users can evaluate their system's signal integrity at the physical layer using the PEX 8636's SerDes Eye Capture feature. Using PLX's SDK, users can view the receiver eye of any lane on the switch. Users can then modify SerDes settings and see the impact on the receiver eye. Figure 5 shows a screenshot of the SerDes Eye Capture feature in the SDK.

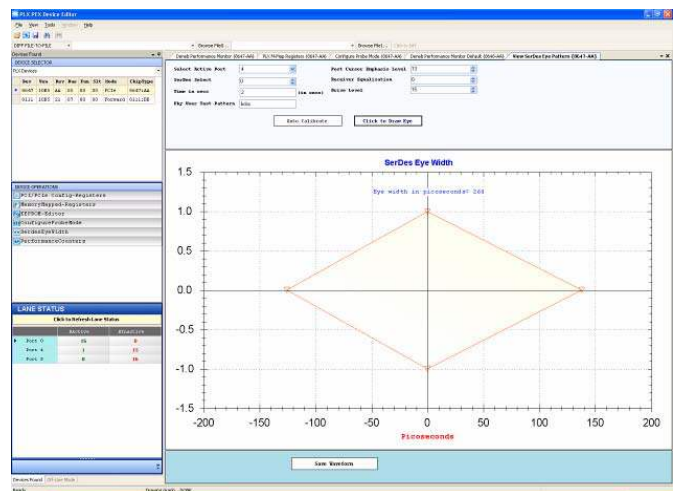


Figure 5. SerDes Eye Capture

Error Injection & SerDes Loopback

Using the PEX 8636's Error Injection feature, users can inject malformed packets and/or fatal errors into their system and evaluate a system's ability to detect and recover from such errors. The PEX 8636 also supports Internal Tx, External Tx, Recovered Clock, and Recovered Data Loopback modes.

PEX 8636, PCI Express Gen 2 Switch, 36 Lanes, 24 Ports

Applications

Suitable for **host-centric** as well as **peer-to-peer traffic patterns**, the PEX 8636 can be configured for a wide variety of form factors and applications.

Host Centric Fan-out in Control Plane

The PEX 8636, with its high port count, provides adequate connectivity when used in host-centric applications. Figure 6 shows a **control plane** design where a single processor is used to control a large number of proprietary ASICs, FPGAs and/or other fabric interface chips. The upstream port in this case can be programmed with a fat pipe, up to x16, to the processor in order to sustain bandwidth requirements from the downstream devices. Moreover, each port in the PEX 8636 can independently support 2.5GT/s and/or 5.0GT/s devices.

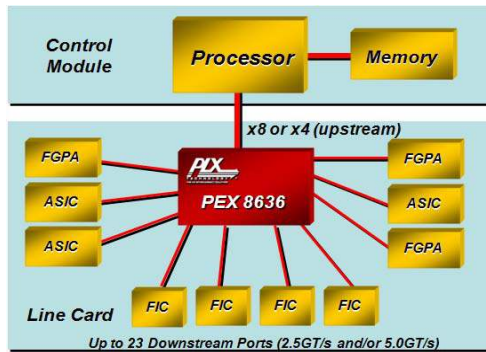


Figure 6. Host Centric Application

Multicast in Control Plane

The PEX 8636 implements Multicast as defined in the PCIe Specification. It provides support for 64 multicast groups and any port in the Switch can be the source port for a multicast transaction. In a control plane application, the multicast feature provides processor relief during configuration and run-time by requiring the processor to issue a single transaction and allowing the switch to replicate that particular transaction to multiple destinations as defined in the multicast group.

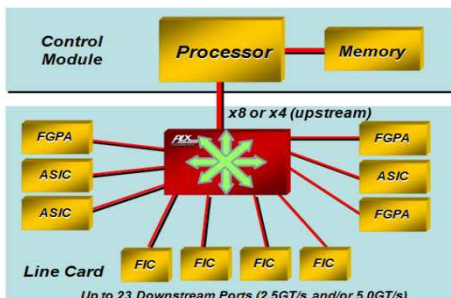


Figure 7. Multicast on PEX 8636

Host Failover in Comms Control Plane

The PEX 8636 can also be utilized in applications where host failover is required. In the control plane application below (Figure 8), an Active-Standby failover example is shown. In this case, there is one active host which is managing the endpoints during run-time. The PEX 8636 provides Non-Transparent functionality which can be enabled on a single port. The Non-Transparent port allows the Standby processor to communicate and exchange status information with the primary processor without interfering with the endpoints. When failover needs to occur, system software configures the PEX 8636 so that the upstream port is connected to the standby processor.

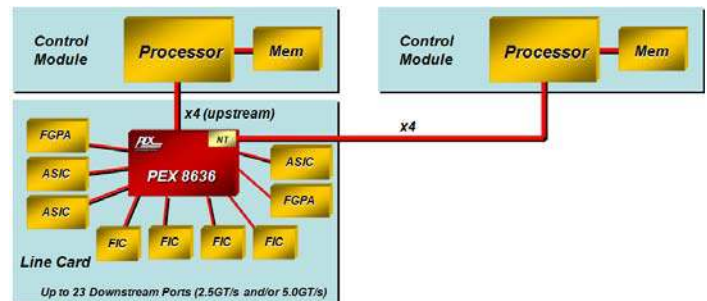


Figure 8. Failover using Non-Transparency

A second method for implementing failover is through the Virtual Switch mode function. Figure 9 shows a control plane application using the Virtual Switch mode. In these examples, two hosts may be active simultaneously and controlling their own domains while exchanging status information through doorbell registers or I²C interface. The devices can be programmed to trigger fail-over if the heartbeat information is not provided. In the event of a failure, the surviving device will reset the endpoints connected to the failing processor and enumerate them in its own domain without impacting the operation of endpoints already in its domain.

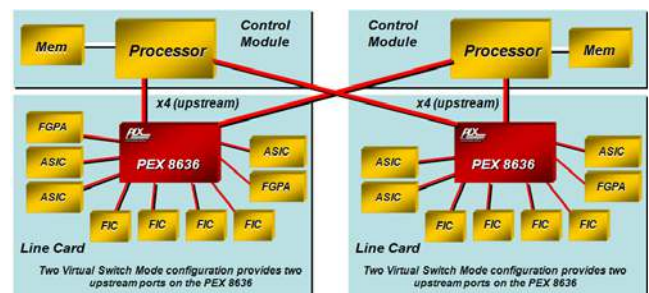


Figure 8. Failover using Virtual Switch function

PEX 8636, PCI Express Gen 2 Switch, 36 Lanes, 24 Ports

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 8636 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 8636 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 8636 for hot plug events, doorbell interrupts, baseline error reporting, and advanced error reporting.

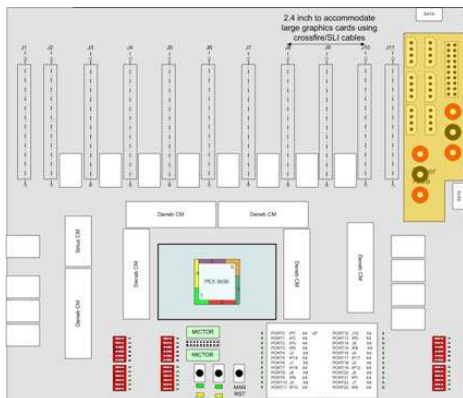


Figure 10. PEX8636-16U1D AIC RDK

Development Tools

PLX offers hardware and software tools to enable rapid customer design activity. These tools consist of a hardware module (PEX 8636 RDK), hardware documentation (available at www.plxtech.com), and a Software Development Kit (also available at www.plxtech.com).

ExpressLane PEX 8636 RDK

The PEX 8636 RDK (see Figure 10) is a hardware module containing the PEX 8636 which plugs right into your system. The PEX 8636 RDK can be used to test and validate customer software, or used as an evaluation vehicle for PEX 8636 features and benefits. The PEX 8636

RDK provides everything that a user needs to get their hardware and software development started.

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 8636. For more information, please refer to the PEX 8636 RDK Product Brief.

Both *performancePAK* and *visionPAK* are supported by PLX's RDK and SDK, the industry's most advanced hardware- and software-development kits.

Product Ordering Information

Part Number	Description
PEX8636-AA50BC F	36-lane, 24-port PCI Express Switch, Pb-Free (35x35mm ²)
PEX8636-16U1D BB RDK	PEX 8636 Rapid Development Kit with x16 Upstream and 20 x1 Downstream

PLX Technology, Inc. All rights reserved. PLX, the PLX logo, ExpressLane, Read Pacing and Dual Cast are trademarks of PLX Technology, Inc. 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 assumes no responsibility for any errors that may appear in this material. PLX reserves the right, without notice, to make changes in product design or specification.

Visit www.plxtech.com for more information.