

PEX 8647

Features

- **PEX 8647 Vitals**
 - 48-lane, 3-port PCIe Gen2 switch
 - Integrated 5.0 GT/s SerDes
- 27 x 27mm², 676-pin FCBGA package
- Typical Power: 2.8 Watts

- **PEX 8647 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**
 - Non-blocking switch fabric
 - Full line rate on all ports
 - Packet Cut-Thru with 140ns max packet latency (x16 to x16)
 - 2KB Max Payload Size
 - Read Pacing (bandwidth throttling)
 - Dual Cast
 - **Flexible Configuration**
 - Registers configurable with strapping pins, EEPROM, I²C, or host software
 - Lane and polarity reversal
 - Selectable upstream port
 - Compatible with PCIe 1.0a PM
 - **Quality of Service (QoS)**
 - Eight traffic classes per port
 - Weighted round-robin source port arbitration
 - **Reliability, Availability, Serviceability**
 - 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
 - Performance Monitoring
 - Per port payload & header counters

PCIe Gen2, 5.0GT/s 48-lane, 3-port PCIe Switch

The *ExpressLane*TM PEX 8647 device offers PCI Express switching capability enabling users to add scalable high bandwidth, non-blocking interconnection to **high-end graphics applications**. The PEX 8647 is optimized to support high-resolution graphics while supporting **peer-to-peer traffic** and dual-cast for maximum performance.

High Performance & Low Packet Latency

The PEX 8647 architecture supports packet **cut-thru with a maximum latency of 140ns (x16 to x16)**. This, combined with large packet memory and non-blocking internal switch architecture, provides full line rate on all ports for performance-hungry graphics applications. 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. Flexible buffer allocation, along with the device's **flexible packet flow control**, maximizes throughput for graphics applications where more traffic flows in the downstream, rather than upstream, direction.

Data Integrity

The PEX 8647 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 as packets pass through the switch.

Register Configuration Flexibility

The PEX 8647 provides several ways to configure its operations. The device can be configured through strapping pins, **I²C interface**, CPU configuration cycles, 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.

Dual Cast

The PEX 8647 supports Dual Cast, a feature which allows for the copying of data (e.g. packets) from one ingress port to two egress ports allowing for higher performance in dual-graphics applications.

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 this downstream port. If this 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, under user control, system performance by allowing for the adequate servicing of all downstream devices.



SerDes Power and Signal Management

The PEX 8647 supports software control of the SerDes outputs to allow for the 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.

Interoperability

The PEX 8647 is designed to be fully compliant with the PCI Express Base Specification r2.0. Additionally, it supports **auto-negotiation**, **lane reversal**, and **polarity reversal**. Furthermore, the PEX 8647 is tested for Microsoft Vista compliance testing. All PLX switches undergo thorough interoperability testing in PLX's **Interoperability Lab** and **compliance testing at the PCI-SIG plug-fest**.

Applications

Targeted at high-end graphics applications, the PEX 8647 supports **host-centric** as well as **peer-to-peer** traffic patterns.

Graphics Fan-out

In a graphics fan-out application (see Figure 1), the PEX 8647 drives a dual-output display. The PEX 8647 will fan out to two Graphics Modules (shown as GPUs in the Figures) via the two x16 downstream ports while the x16 upstream port links to the Root Complex. Each graphics module drives its own monitor. Increasing memory and bandwidth requirements have put a strain on local GPU memory. The PEX 8647 allows for highly efficient data transfers over the PCI Express bus, allowing the Graphics Module to utilize the system memory and render it as if it were local graphics memory. In a fan-out application such as this one, each Graphics Module drives its own output.

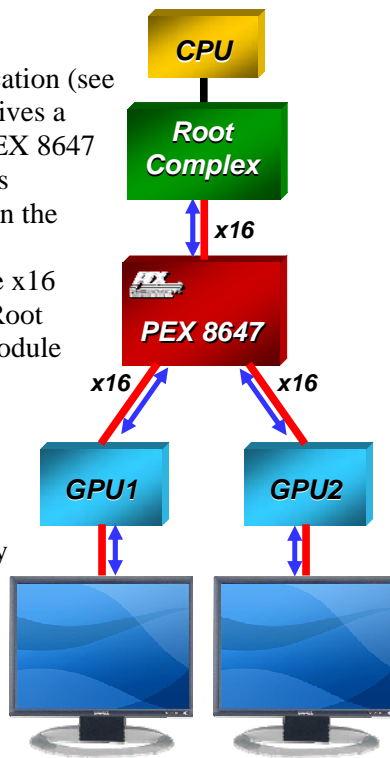


Figure 1. Dual Graphics Fan-Out

Dual Graphics with Peer-to-Peer Communication

High resolution 3D graphics applications can take full advantage of the PEX 8647 three port configuration. Applications such as high-resolution gaming, high resolution scientific use, and image processing can benefit from the performance of the PEX 8647 switch.

Figure 2 illustrates the use of the device in a high resolution gaming application where two Graphics Modules drive a single monitor for the ultimate gaming experience. The upstream x16 port links to the Root Complex and the two downstream ports connect to the Graphics Modules. The peer-to-peer support of the PEX 8647 allows the two Graphics Modules to communicate with each other for maximum performance.

In this example, the two Graphics Modules divide the screen into a checkerboard pattern. In Figure 2, the screen is divided into white frames and blue frames, with one GPU managing the white frames and the other managing the blue frames. This mode of operation is referred to as **Supertiling**, and is generally the most efficient because it evenly divides the processing and graphics rendering workload across the two Graphics Modules.

This usage model calls for heavy peer-to-peer communication between the two Graphics Modules. The PEX 8647 can also support dual-graphics solutions running in scissor, or alternate frame-rate modes. In each of these modes, the processing and graphics rendering workload is shared by the Graphics Modules, and therefore requires a great amount of peer-to-peer communication between the Graphics Modules to monitor each other's progress and execution.

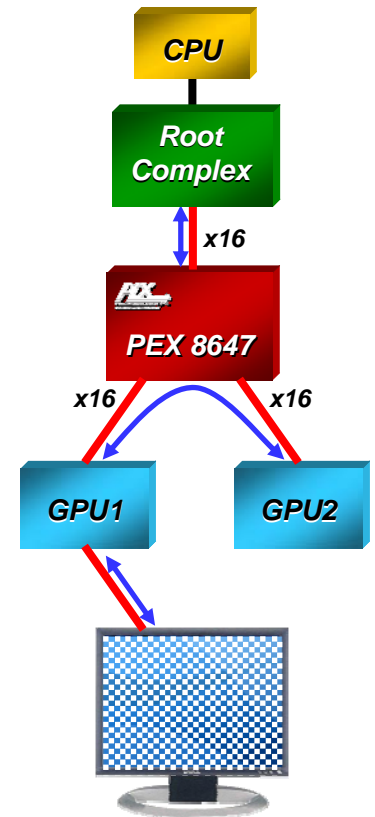


Figure 2. Dual Graphics with Peer-to-Peer Communication

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 8647 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 8647 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 8647 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 8647RDK), hardware documentation (available at www.plxtech.com), and a Software Development Kit (also available at www.plxtech.com).

ExpressLane PEX 8647RDK

The PEX 8647RDK is a hardware module containing the PEX 8647 which plugs right into your system. The PEX 8647RDK can be used to test and validate customer software, or used as an evaluation vehicle for PEX 8647 features and benefits. The PEX 8647RDK 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 8647.



PLX Technology, Inc.
870 Maude Ave.
Sunnyvale, CA 94085 USA
info@plxtech.com
www.plxtech.com

Product Ordering Information

Part Number	Description
PEX8647-BB50BC F	48-Lane, 3-Port PCI Express Switch, Pb-Free (27x27mm ²)
PEX8647-BB RDK	PEX 8647 Rapid Development Kit

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

© 2009 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.