ExpressLane PEX 8508AC
5-Port/8-Lane PCI Express Switch
Data Book

Version 1.3

May 2007

Website www.plxtech.com
Technical Support www.plxtech.com/support
Phone 800 759-3735
408 774-9060
FAX 408 774-2169

Revision History

Revision	Date	Description of Changes
1.0	September, 2006	Initial Production release, Silicon Revision AA.
1.1	January, 2007	Initial Production release, Silicon Revision AB. Updated Power-related values in Chapter 23, "Electrical Specifications." Updated RDK ordering information in Appendix B, "General Information." Applied miscellaneous corrections throughout data book.
1.2	April, 2007	Initial Production release, Silicon Revision AC. Applied miscellaneous changes, corrections, and enhancements throughout data book.
1.3	May, 2007	Removed support for Silicon Revisions AA and AB. Applied miscellaneous changes, corrections, and enhancements throughout data book.

Copyright Information

Document Number: 8508AC-SIL-DB-P1-1.3

Preface

The information contained in this data book is subject to change without notice. This data book will be updated periodically as new information is made available.

This data book documents information for the PEX 8508, Silicon Revision AC.

Audience

This data book provides functional details of PLX Technology's ExpressLane PEX 8508AC 5-Port/ 8-Lane PCI Express Switch, for hardware designers and software/firmware engineers.

Supplemental Documentation

This data book assumes that the reader is familiar with the following documents:

- *PLX Technology, Inc. (PLX)*

  870 W Maude Avenue, Sunnyvale, CA 94085

  Tel: 800 759-3735 (domestic only) or 408 774-9060, Fax: 408 774-2169, www.plxtech.com

  - *PEX 85XX EEPROM – PEX 8518/8517/8508 Design Note*
  - *PEX 8508 Errata*

  The PLX PEX 8508 Toolbox includes other PEX 8508 documentation as well.

- *PCI Special Interest Group (PCI-SIG)*

  3855 SW 153rd Drive, Beaverton, OR 97006 USA

  Tel: 503 619-0569, Fax: 503 644-6708, www.pcisig.com

  - *PCI Local Bus Specification, Revision 3.0*
  - *PCI Bus Power Management Interface Specification, Revision 1.2*
  - *PCI to PCI Bridge Architecture Specification, Revision 1.2*
  - *PCI Express Base Specification, Revision 1.1*
  - *PCI Express Card Electromechanical Specification, Revision 1.1*
  - *PCI Express Architecture PCI Express Jitter and BER White Paper, Revision 1.0*

- *The Institute of Electrical and Electronics Engineers, Inc.*

  445 Hoes Lane, Piscataway, NJ 08854-4141 USA

  Tel: 800 701-4333 (domestic only) or 732 981-0060, Fax: 732 981-9667, www.ieee.org

  - *IEEE Standard 1149.1-1990, IEEE Standard Test Access Port and Boundary-Scan Architecture, 1990*
  - *IEEE Standard 1149.1a-1993, IEEE Standard Test Access Port and Boundary-Scan Architecture*
  - *IEEE Standard 1149.1-1994, Specifications for Vendor-Specific Extensions*
  - *IEEE Standard 1149.6-2003, IEEE Standard Test Access Port and Boundary-Scan Architecture Extensions*

- *NXP Semiconductors*

  www.standardics.nxp.com/

  - *The I2C-Bus Specification, Version 2.1*

Note: *In this data book, shortened titles are associated with the previously listed documents. The following table lists these abbreviations.*

Abbreviation	Document
PCI r3.0	*PCI Local Bus Specification, Revision 3.0*
PCI Power Mgmt. r1.2	*PCI Bus Power Management Interface Specification, Revision 1.2*
PCI-to-PCI Bridge r1.2	*PCI to PCI Bridge Architecture Specification, Revision 1.2*
PCI Express Base r1.1	*PCI Express Base Specification, Revision 1.1*
PCI Express CEM r1.1	*PCI Express Card Electromechanical Specification, Revision 1.1*
IEEE Standard 1149.1-1990	*IEEE Standard Test Access Port and Boundary-Scan Architecture*
IEEE Standard 1149.6-2003	*IEEE Standard Test Access Port and Boundary-Scan Architecture Extensions*
I^2C Bus v2.1 *I2C Bus v2.1*[a]	*The I^2C-Bus Specification, Version 2.1*

a. Due to formatting limitations, the specification name may appear without the superscripted "2" in its title.

Terms and Abbreviations

The following table lists common terms and abbreviations used in this data book. Terms and abbreviations defined in the *PCI Express Base r1.1* are not included in this table.

Terms and Abbreviations	Definitions
AMCAM	Address-mapping CAM that determines a memory request route. Contains mirror copies of the PCI-to-PCI bridge **Memory Base** and **Limit** registers in the switch.
BusNoCAM	Bus Number mapping CAM that determines the Completion route. Contains mirror copies of the PCI-to-PCI bridge **Secondary Bus Number** and **Subordinate Bus Number** registers in the switch.
CAM	Content Addressable Memory.
CFC	Constant Frequency Clocking.
CSRs	Configuration Space Registers.
Egress	Outgoing traffic from chip.
Egress queue	Egress queuing/scheduling mechanism.
ECRC	End-to-end Cyclic Redundancy Check (CRC).
Foreign	Reference to PCI Express attributes that belong to (off-chip) PCI Express components located on the opposite (other side) of PCI Express links.
GPIO	General-Purpose Input/Output.
HOL	Head of Line.
Ingress	Incoming traffic to chip.
Ingress queue	Ingress queuing/scheduling mechanism.
IOAMCAM	I/O Address mapping CAM that determines an I/O request route. Contains mirror copies of the PCI-to-PCI bridge **I/O Base** and **Limit** registers in the switch.
Lane	Bidirectional pair of differential PCI Express I/O signals.
Link Interface	Primary side of the NT Port, connects to external device pins. The secondary side of the NT Port is referred to as the *NT Port Virtual Interface*, and connects to the internal virtual PCI Express interface.
Local	Reference to PCI Express attributes (*such as*, credits) that belong to the PCI Express station.
MRL	Manually-operated Retention Latch.
MPS	Maximum Payload Size.
NT	Non-Transparent. A bridging technique used in the PCI Express Switch to isolate Memory spaces by presenting the processor as an endpoint rather than another memory system.
PCI Express Station	A functional unit that provides the PCI Express conforming system interface. Includes the Serializer and De-serializer (SerDes) hardware interface modules and PCI Express Interface, which provides the Physical Layer, Data Link Layer, and Transaction Layer (TL) logic.
PHY	Physical Layer.
Port	Ports are a collection of lanes configured at startup which contain the functional logic and memory resources to communicate with like resources in other PCI Express devices.
P-P	PCI-to-PCI.
PRBS	Pseudo-Random Bit Sequence.
QoS	Quality of Service.
RAS	Reliability, Availability, and Serviceability.
RoHS	Restrictions on the use of certain Hazardous Substances (RoHS) Directive.

Terms and Abbreviations	Definitions
RR	Round-Robin scheduling.
SerDes	Serializer/De-serializer. A high-speed differential-signaling parallel-to-serial and serial-to-parallel conversion logic attached to lane pads.
SSC	Spread Spectrum Clocking.
TC	Traffic Class.
TDM	Time Division Multiplexing.
TLC	Transaction Layer Control. The module performing PCI Express Transaction Layer functions.
TLP	Transaction Layer Packet. PCI Express packet formation and organization.
Transparent	Refers to standard PCI Express upstream-to-downstream routing protocol.
VC	Virtual Channel.
VC&T	Virtual Channel and Type.
Virtual Interface	Secondary side of the NT Port, connects to the internal virtual PCI Express interface.

Data Book Notations and Conventions

Notation/Convention	Description
Blue text	Indicates that the text is hyperlinked to its description elsewhere in the data book. Left-click the blue text to learn more about the hyperlinked information. This format is often used for register names, register bit and field names, register offsets, chapter and section titles, figures, and tables.
PEX_XXXp[10, 8, 6, 4:0]	When the signal name appears in all CAPS, with the primary port description listed first, field [10, 8, 6, 4:0] indicates the number of signal balls/pads assigned to that port. The lowercase "p" (positive) or "n" (negative) suffix indicates the differential pair of the signal, which are always used together.
# = Active-Low signals	Unless specified otherwise, Active-Low signals are identified by a "#" appended to the term (*for example*, PEX_PERST#).
`Program/code samples`	`Monospace font` (*`program or code samples`*) is used to identify code samples or programming references. These code samples are case sensitive, unless specified otherwise.
`command_done`	Interrupt format.
Command/Status	Register names.
Parity Error Detected	Register parameter [field] or control function.
Upper Base Address[31:16]	Specific function in 32-bit register bounded by bits [31:16].
Number multipliers	k = 1,000 (10^3)is generally used with frequency response. K = 1,024 (2^{10}) is used for memory size references. KB = 1,024 bytes. M = meg. = 1,000,000 when referring to frequency (decimal notation) = 1,048,576 when referring to memory sizes (binary notation).
1Fh	h = suffix which identifies hex values. Each prefix term is equivalent to a 4-bit binary value (nibble). Legal prefix terms are 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, F.
1010b	b = suffix which identifies binary notation (*for example*, 01b, 10b, 1010b, and so forth). Not used with single-digit values of 0 or 1.
0 through 9	Decimal numbers.
byte	Eight bits – abbreviated to "B" (*for example*, 4B = 4 bytes).
LSB	Least-Significant Byte.
lsb	Least-significant bit.
MSB	Most-Significant Byte.
msb	Most-significant bit.
DWord	DWord (32 bits) is the primary register size in these devices.
Reserved	Do not modify ***Reserved*** bits and words. Unless specified otherwise, these bits read as 0 and are to be written as 0.

THIS PAGE INTENTIONALLY LEFT BLANK.

Contents

THIS PAGE INTENTIONALLY LEFT BLANK.

Registers

Transparent Port Registers

NT Port Virtual Interface Registers

NT Port Link Interface Registers

Chapter 1 Introduction

1.1 Features

PLX Technology's ExpressLane™ PEX 8508 PCI Express Switch supports the following features:

- 5-Port PCI Express switch
  - 8 lanes with integrated on-chip SerDes
  - Low power (under 250 mW per lane)
  - Fully non-blocking switching architecture
  - Optional Device-Specific Relaxed Ordering
  - Port configuration
    - Five independent ports (Ports 0, 1, 2, 3, and 4)
    - Choice of width (# of lanes) per unique port/link – x1, x2, x4
    - Configurable with serial EEPROM or I^2C
    - Selectable Upstream port
  - Maximum payload size is 256 bytes
- Quality of Service (QoS)
  - 2 Virtual Channels (VC0 and VC1)[1]
  - All ports support eight Traffic Class (TC) mapping, independently of the other ports
  - TCs can be symmetrically or asymmetrically mapped to a VC
  - Egress port arbitration
  - VC arbitration
    - Fixed-Priority
    - Round-Robin
- Reliability, Availability, and Serviceability (RAS) features
  - Standard Hot Plug Controller for all downstream ports, including optional usage models for Manually operated Retention Latch (MRL) Sensor and Attention Button support (Electromechanical Interlock supported with the Power Enable output)
  - Baseline and Advanced Error Reporting capability
- FATAL ERROR (FATAL_ERR#) (Conventional PCI SERR# equivalent) ball support
- Optional PCI Express features implemented in the PEX 8508
  - TLP Digest support – Poison bit and end-to-end CRC (ECRC)
  - Lane reversal support
  - Link Power Management states – L0, L0s, L1, L2, L2/L3 Ready, L3
  - VAUX, WAKE#, Beacon support
  - PCI Bus Power Management states – D0, D3 (hot and cold)
  - Active State Power Management (ASPM) fully supported[2]

1. Used up to three-port configuration. For four- or five-port configuration, one VC is allowed.
2. ASPM is ***not supported*** in Cut-Thru mode.

- Non-Transparent Bridging (NTB) Port support – Any one downstream port can be selected as the NT Port
- Out-of-Band Initialization options
  - Serial EEPROM
  - I^2C (7-bit Slave address with 100 Kbps)
- Performance
  - Bandwidth of 40 Gbps [2.5 Gbps x 8 SerDes x 2 (Full Duplex)]
  - Unloaded Cut-Thru latency of 150 ns for Transparent mode in x4 to x4 mode; under 175 ns for NTB port
  - Peer-to-peer latency under 150 ns
  - Non-blocking internal crossbar supporting TLP bandwidth capacity of each x4 link
- Testability – JTAG support for DC
- Spread-Spectrum Clocking option
- 19-mm square, 296-ball PBGA-H (Plastic BGA Heat-spreader) package
- Maximum power – 2.50W
- Compliant to the following specifications:
  - *PCI Local Bus Specification, Revision 3.0* (*PCI r3.0*)
  - *PCI Bus Power Management Interface Specification, Revision 1.2* (*PCI Power Mgmt. r1.2*)
  - *PCI to PCI Bridge Architecture Specification, Revision 1.2 (PCI-to-PCI Bridge r1.2)*
  - *PCI Express Base Specification, Revision 1.1* (*PCI Express Base r1.1*)
  - *PCI Express Card Electromechanical Specification, Revision 1.1* (*PCI Express CEM r1.1*)
  - *IEEE Standard Test Access Port and Boundary-Scan Architecture (IEEE Standard 1149.1-1990)*
  - *IEEE Standard Test Access Port and Boundary-Scan Architecture Extensions (IEEE Standard 1149.6-2003)*
  - *The I^2C-Bus Specification, Version 2.1* (*I^2C Bus v2.1*)

1.2 Overview

The PEX 8508 is a fully non-blocking, low-latency, low-cost, and low-power 5-port, 8-lane PCI Express switch. Conforming to the *PCI Express Base r1.1*, the PEX 8508 enables users to add scalable, high-bandwidth I/O to a wide variety of products, including intelligent adapter boards, general-purpose servers, edge communication products, storage routers, blade servers, docking stations, printer engine controllers, and other embedded products. The PEX 8508's flexible hardware configuration and software programmability allow the switch to be tailored to suit a wide variety of applications.

The PEX 8508 is principally aimed at fan-in/fan-out or aggregation applications; however, it can also be used in peer-to-peer communication traffic patterns and as an isolation switch through its built-in Non-Transparent support. The PEX 8508 includes 8 PCI Express SerDes links that are used in a wide variety of configurations. For example, if the PEX 8508 is being used in a fan-out application, the upstream port can be configured as x4 and the downstream ports as four x1 ports; two x2 ports; one x4 port and two x2 ports; or other combinations, provided that the maximum number of lanes (8) or ports (5) is not exceeded. In a dual-Host application, the PEX 8508 can be used to isolate the two hosts, by routing traffic through the NT Port by way of two x4 links. Figure 1-1 illustrates the most common port configurations.

The PEX 8508 offers support for up to two VCs, depending upon the number of ports used. User-selectable port and Virtual Channel (VC) arbitration algorithms enable users to fine-tune QoS for the most demanding applications. The PEX 8508's Non-Transparent support, Hot Plug capability, and Advanced Error Reporting features make it suitable for High-Availability applications.

Figure 1-1. PEX 8508 Port Configuration Examples

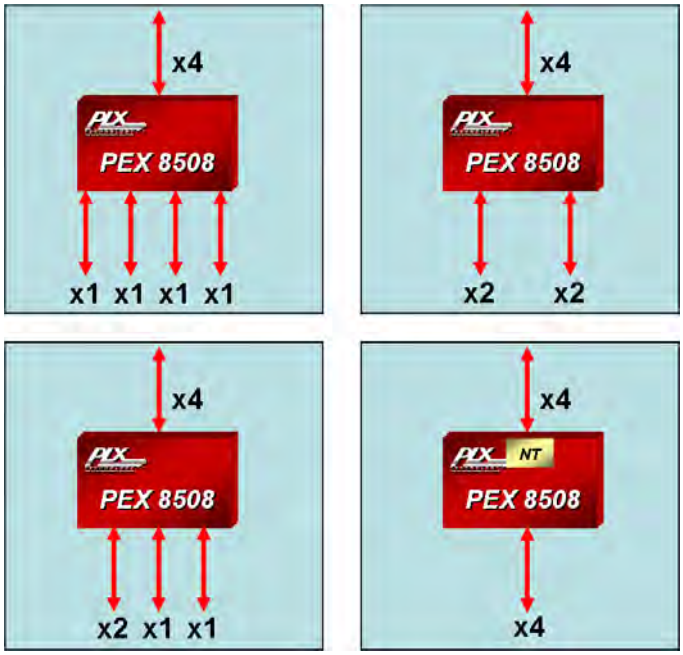

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 2 Applications

2.1 Multi-Purpose and Feature-Rich ExpressLane PEX 8508 PCI Express Switch

The PLX ExpressLane PEX 8508 is the most versatile and flexible low-port count PCI Express switch available in today's market. The PEX 8508 is used as fan-out/fan-in, processor isolation, or peer-to-peer switching. The 8-lane PEX 8508 product offers PCI Express switching capability conforming to the *PCI Express Base r1.1*. This product enables the ability to add scalable high-bandwidth, non-blocking interconnection to a wide variety of applications, including servers, adapter boards, docking stations, blade servers, and embedded-control products.

2.1.1 High Performance

The PEX 8508 architecture supports packet Cut-Thru with a latency of 150 ns for Transparent mode in x4 to x4 mode for ingress and egress ports; under 175 ns for NTB ingress and egress ports. This allows high TLP throughput on each x4 link for performance-hungry applications.

2.1.2 End-to-End Packet Integrity

The PEX 8508 provides ECRC protection and Poison-bit support to enable designs that require guaranteed error-free packets. These features are optional in the *PCI Express Base r1.1*; however, they are provided across the entire PLX PCI Express product line.

2.1.3 Interoperability

The PEX 8508 is designed to be fully compliant with the *PCI Express Base r1.1*. Additionally, the device supports auto-negotiation and lane reversal for maximum board design and board placement flexibility. Furthermore, the PEX 8508 is interoperable with many popular motherboards and server boards with PCI Express connections, and PCI Express endpoints (Ethernet, RAID Controllers), as well as PLX's family of PCI Express switches and bridges. All PLX ExpressLane devices undergo thorough interoperability testing at PLX's Interoperability Lab.

2.1.4 PCI Express Switch Non-Transparent Bridging (NTB)

The PEX 8508 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-Transparent features are implemented in the same manner as in Conventional PCI applications. NT bridges allow systems to isolate memory domains by presenting the processor subsystem as an endpoint, rather than another memory system. BARs 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.

2.1.5 Two Virtual Channels (VCs)

The PEX 8508 supports up to two full-featured VCs (Virtual Channels 0 and 1) and a full set of 8 TCs. TC mapping to port-specific VCs allows for divergent mappings for particular ports. In addition, the devices offer user-selectable VC arbitration algorithms to enable users to fine-tune the Quality of Service (QoS) required for a specific application. Use of multiple VCs is for applications requiring added QoS, where typical use is for no more than three ports. The PEX 8508 is limited to one VC in four- or five-port configurations. In this case, the PEX 8508 is normally an aggregation switch with four x1 ports and one x4 port and added VC support is not of much use. If the application is aggregating two x2 ports to one x4 port, two VCs are supported and QoS issues might become pertinent.

The PEX 8508 supports Hardware-fixed arbitration schemes for both VCs. This allows QoS fine-tuning and efficient use of system bandwidth.

2.1.6 Low Power with Granular SerDes Control

The PEX 8508 provides low-power capability that is fully compliant with the *PCI Express Base r1.1* Power Management specifications. Unused SerDes can be disabled, to reduce power consumption. All link layer and device power management schemes are supported.

The PEX 8508 supports SerDes output software control, to allow power and signal strength optimization in a system. The PLX SerDes implementation supports four power levels – *Off*, *Low*, *Typical*, and *High*. The SerDes block also supports Loop-Back modes and Advanced Error Reporting, which enables efficient system debug and management.

2.1.7 Flexible Port-Width Configuration

The PEX 8508 offers highly configurable ports. There are up to five ports, which can be configured to most legal widths (from x2 to x4), or trained down to x1, provided the total lane count does not exceed 8 in any combination, to support specific bandwidth needs. The ports are symmetric (each port has the same lane width) or asymmetric (ports have different lane widths). Any one port can be designated as, or dynamically changed to, the upstream port.

Lane width can be individually configured from each port through auto-negotiation, hardware strapping, upstream software configuration, or an optional serial EEPROM or I^2C interface.

2.1.8 Hot Plug for High-Availability Applications

Hot Plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8508's Hot Plug capabilities and Advanced Error Reporting features make it suitable for High-Availability (HA) applications. Each downstream port includes a standard Hot Plug Controller. If the PEX 8508 is used in an application where one or more of its downstream ports are connected to PCI Express slots, the Hot Plug Controller for that port is used to manage the Hot Plug event of its associated slot. Additionally, the upstream port is a fully compliant Hot Plug client, and the PEX 8508 can be used on Hot-Pluggable adapter boards, docking stations, and line cards.

2.1.9 Fully Compliant Power Management

For applications that require power management, the PEX 8508 supports Link (L0, L0s, L1, L2, L2/L3 Ready, L3) and Device (D0, D3hot, and D3cold) Power Management states, in compliance with the *PCI Express Base r1.1* Power Management specifications.

2.2 Applications

The PEX 8508 Multi-port PCI Express switch adds low-cost and low-latency aggregation and peer-to-peer switching with non-transparency for isolation. Other uses include high-volume applications in adapter add-in boards, docking station, and printer engine design applications. Other uses include isolation of processor domains in multi-Host systems with the use of non-transparency. Systems with limited endpoint quantities can also use the PEX 8508 as a peer-to-peer switch.

2.2.1 Host Centric Fan-Out

The PEX 8508 switch, with its symmetric or asymmetric lane-configuration capability, allows user-specific tuning to a variety of Host-centric applications.

Figure 2-1 illustrates a typical server-based design, where the Root Complex provides a PCI Express link that must 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 8508 would typically have a 4-lane upstream port, and as many as four downstream ports. The downstream ports can be of differing widths, if required. Figure 2-1 also illustrates how some of the ports can be bridged to provide PCI or PCI-X slots, through the use of PLX Technology's ExpressLane PEX 8114 and PEX 8111 PCI Express bridging devices.

Figure 2-1. Fan-In/Fan-Out Usage

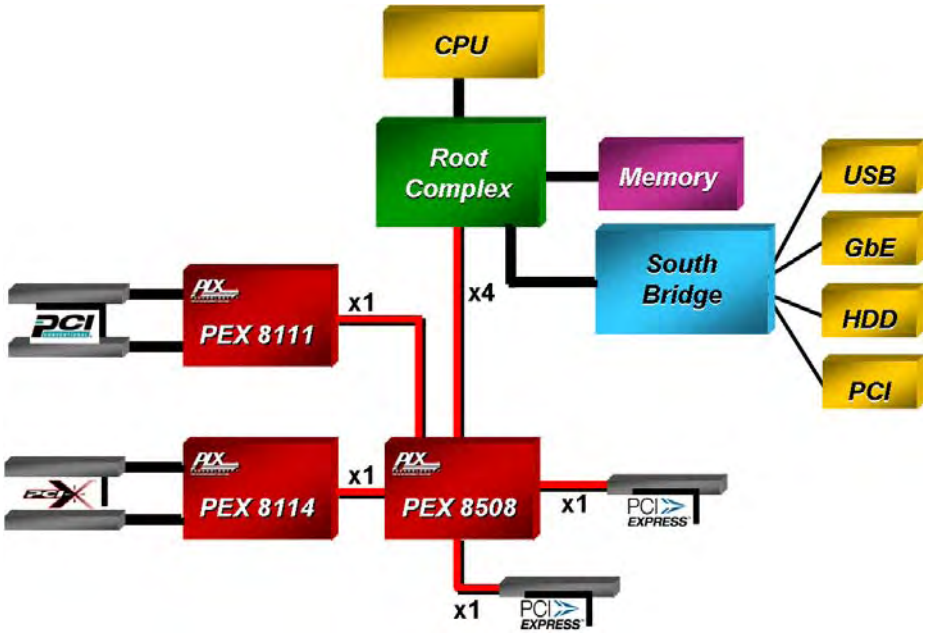

2.2.2 Dual Host/Fabric Model

The PEX 8508 supports applications requiring dual Host, Host failover, and load-sharing applications through the NTB feature. Figure 2-2 illustrates a dual Host system with dual-switch fabric in a dual-star configuration. The redundancy of the Host and fabric can be achieved through many possible configurations, using the PEX 8508's NTB function. In Figure 2-2, Host 1 controls Switch 1 and its associated I/Os, and Host 2 controls Switch 2 and its associated I/Os. The hosts and switches are isolated, using the NTB functionality 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 and data messages to all the I/Os through the surviving switch, using the NTB function.

Figure 2-2. Dual Host/Fabric Mode Embedded Systems

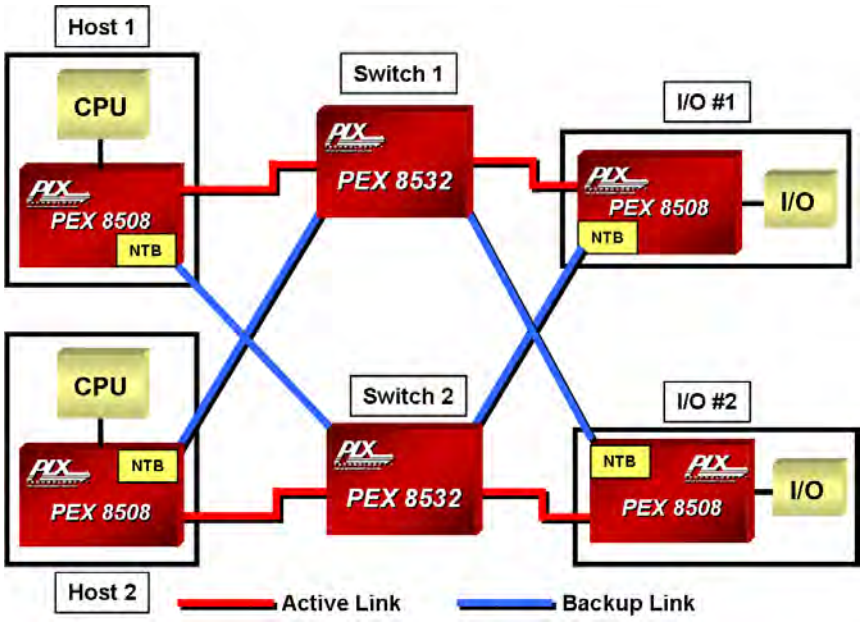

2.2.3 Adapter Board Aggregation

The number and variety of PCI Express native-mode devices is rapidly growing. As these devices become mainstream, it becomes necessary to create multi-function and multi-port adapter boards with PCI Express capability.

The PEX 8508 can be used to create an adapter or mezzanine board that aggregates PCI Express devices into a single port, that can be plugged into a backplane or motherboard. Figure 2-3 illustrates use of the PEX 8508 in this application.

The adapter board in Figure 2-3 can be Transparent, in which case the PCI Express devices are standard I/O products *such as* Ethernet, Fibre Channel, and so forth. Or, the PEX 8508 can provide a Non-Transparent port to the system (by way of the board edge). In this case, one of the PCI Express devices can be a CPU or other "intelligent" device with on-chip processing capability – thereby needing address domain isolation from the remainder of the system. This approach is commonly used in RAID Controllers.

Figure 2-3. Adapter Board Aggregation

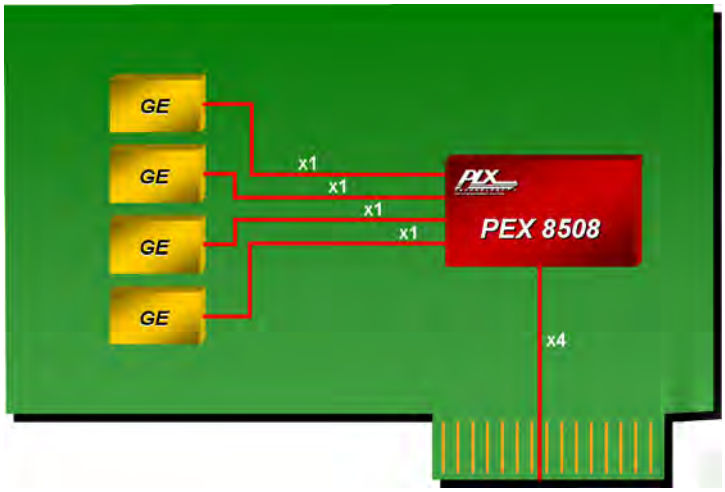

2.2.4 Intelligent Adapter Board

The PEX 8508 supports Non-Transparent Bridging (NTB). Figure 2-4 illustrates a Host system using an intelligent adapter board. In Figure 2-4, the CPU on the adapter board is isolated from the Host CPU. The PEX 8508 NT Port allows the two CPUs to be isolated, but able to communicate with one another through various PEX 8508 device-specific registers. The Host CPU can dynamically re-assign the PEX 8508 upstream port and the PEX 8508 NT port, allowing the system to be re-configured.

Figure 2-4. Intelligent Adapter Board Usage

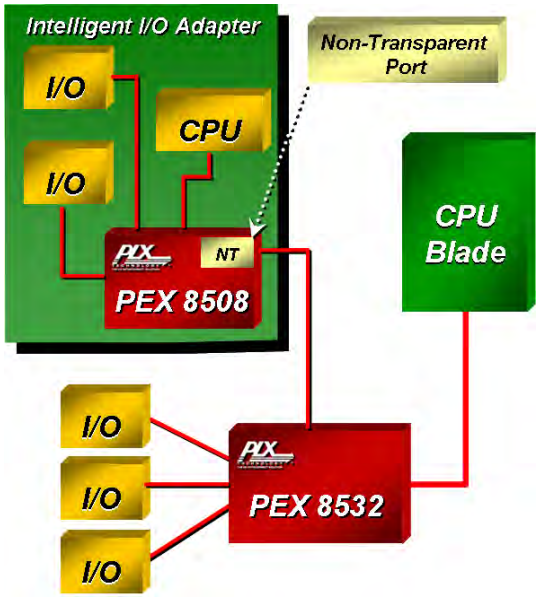

2.2.5 Docking Station Support with Low-Cost, 5-Port PCI Express Switch

The PEX 8508 includes a low-cost architectural design, along with key features useful in a docking station application. Figure 2-5 illustrates the PEX 8508 aggregating several I/O device inputs and passing data to the notebook. In this application, up to four endpoint devices are connected with a x1 link. The WAKE# and Beacon functions, along with VAUX power, are instrumental in supporting system requirements.

Figure 2-5. PCI Express Switch in a Docking Station

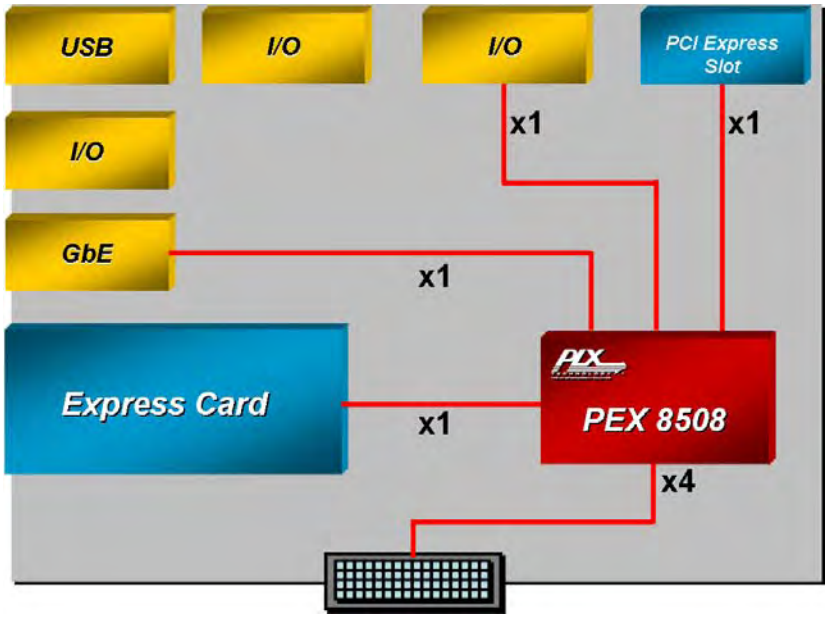

2.2.6 ATCA and XMC Mezzanine Board Applications

The PEX 8508 supports aggregation or fan-in/out, as well as peer-to-peer switching. Figure 2-6 illustrates I/O Mezzanine boards with a PEX 8508, which can aggregate data to the main baseboard or allow peer-to-peer switching between I/O endpoints. Non-transparency can isolate each of the I/O Mezzanine boards from one another.

Figure 2-6. PEX 8508 in XMC Mezzanine Board

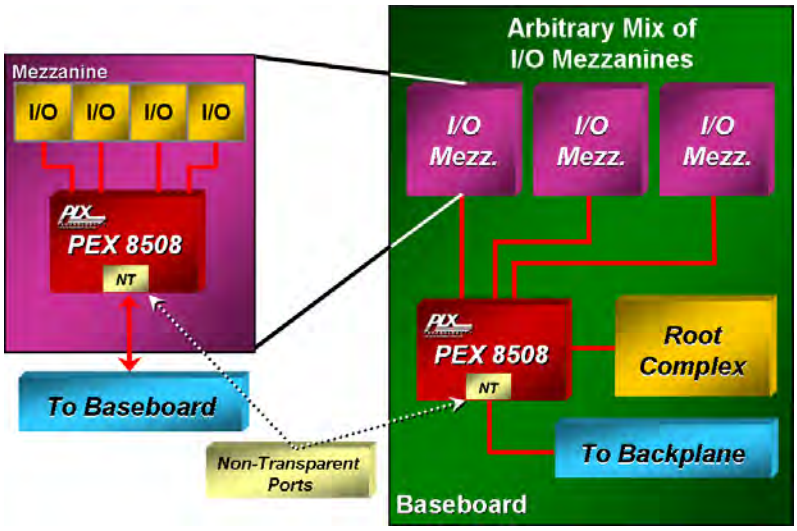

2.3 Software Usage Model

From the system model viewpoint, each PCI Express port is a virtual PCI-to-PCI bridge device with its own set of PCI Express Configuration registers. By default, the PEX 8508 uses Port 0 as the upstream port; however, any port can be configured as the upstream port through optional configuration, by way of a serial EEPROM, the I^2C interface, or Strapping balls. The BIOS or Host can configure the other ports by way of the upstream port, using Conventional PCI enumeration.

2.3.1 System Configuration

The virtual PCI-to-PCI bridges within the PEX 8508 are compliant with 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 0/1 Configuration requests, by way of the virtual primary bus interface (matching Bus, Device, and Function Numbers).

2.3.2 Interrupt Sources and Events

The PEX 8508 supports the INT*x* Interrupt message type (compatible with *PCI r3.0* Interrupt signals) or Message Signaled Interrupts (MSI), when enabled. The PEX 8508 generates messages for PCI Express Baseline and Advanced Error Reporting error reporting mechanisms. The PEX 8508 generates interrupts for Hot Plug events, and device-specific internal errors, and forwards interrupts received from downstream ports. Both forwarded and internally generated interrupts are remapped and collapsed at the upstream port.

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 3 Signal Ball Description

3.1 Introduction

This chapter provides descriptions of the 296 PEX 8508 signal balls. The signals are divided into the following groups:

- PCI Express Signals
- Hot Plug Signals
- Serial EEPROM Signals
- Strapping Signals
- JTAG Interface Signals
- I2C Interface Signals
- Fatal Error Signal
- Wake Signal
- No Connect Signals
- Power and Ground Signals

The signal name, type, location, and a brief description are provided for each signal ball.

3.2 Abbreviations

The following abbreviations are used in the signal tables provided in this chapter.

Table 3-1. Ball Assignment Abbreviations

Abbreviation	Description
#	Active-Low signal
APWR	1.0V Power (VDD10A) balls for SerDes Analog circuits
AUXPWR	3.3V Power (VDD33X) balls for Auxiliary circuits
CMLCLK_CFCn[a]	Spread-Spectrum differential low-voltage, high-speed, CML negative Clock inputs
CMLCLK_CFCp[a]	Spread-Spectrum differential low-voltage, high-speed, CML positive Clock inputs
CMLCLKn[a]	Differential low-voltage, high-speed, CML negative Clock inputs
CMLCLKp[a]	Differential low-voltage, high-speed, CML positive Clock inputs
CMLRn	Differential low-voltage, high-speed, CML negative Receiver inputs
CMLRp	Differential low-voltage, high-speed, CML positive Receiver inputs
CMLTn	Differential low-voltage, high-speed, CML negative Transmitter outputs
CMLTp	Differential low-voltage, high-speed, CML positive Transmitter outputs
CPWR	1.0V Power (VDD10) balls for low-voltage Core circuits
GND	Common Ground (VSS) for all circuits; also associated with VSS_THERMAL (thermal ground)
I	Input (signals with weak internal pull-up resistors)
I/O	Bidirectional programmable signal (input or output)
I/OPWR	3.3V Power (VDD33) balls for Input and Output interfaces
O	Output
OD	Open Drain output
PLL_GND	PLL Ground connection
PLLPWR	3.3V Power (VDD33A) balls for PLL circuits
PU	Pull-up resistor (recommended value between 3K to 10K Ohms)
SerDes	Differential low-voltage, high-speed, I/O signal pairs (negative and positive)
SPWR	1.0V Power (VDD10S) balls for SerDes Digital circuits
STRAP	Strapping balls cannot be left floating on the board
XPWR	1.0V Power (VDD10X) balls for SerDes Auxiliary circuits

a. For REFCLK input, CML source is recommended; however, LVDS source is supported.

3.3 Internal Pull-Up Resistors

The signals listed in Table 3-2 have a weak internal pull-up resistor. If a listed signal is not used and no board trace is connected to the ball, the internal resistor is normally sufficient to keep the signal from toggling. If a listed signal is not used, but is connected to a board trace, the internal resistors might not be strong enough to hold the signal in the inactive state, and therefore it is recommended that the signal be pulled High to VDD33 or Low to VSS (GND), as appropriate, through a 3K- to 10K-Ohm resistor.

Table 3-2. Balls with Internal Pull-Up Resistors

Signal Name		
EE_DO	HP_PRSNT[4:0]#	JTAG_TDI
EE_PR#	HP_PWRFLT[4:0]#	JTAG_TMS
HP_BUTTON[4:0]#	I2C_ADDR[2:0]	JTAG_TRST#
HP_MRL[4:0]#	JTAG_TCK	

3.4 Signal Ball Descriptions

3.4.1 PCI Express Signals

The PCI Express SerDes and Control signals are defined in Table 3-3.

The SerDes lanes are not numbered sequentially, because they are distributed across three SerDes Lane groups, called *quads*. The quads are defined as SerDes[0-3] (first quad), SerDes[4-7] (second quad), and SerDes[8-11] (third quad). The PEX 8508 does not use SerDes[5, 7, 9, 11]; Lanes 5, 7, 9, and 11 are ***reserved***. Therefore, Lanes 0 through 3 populate the first quad, Lanes 4 and 6 populate two SerDes in the second quad, and Lanes 8 and 10 populate two SerDes in the third quad. This arrangement affects all PCI Express lanes and the appropriate PEX_LANE_GOOD*x*# signals.

Note: *The ball numbers are ordered in sequence to follow the Signal Name sequencing [n to 0].*

Table 3-3. PCI Express Signals – 46 Balls

Signal Name	Type	Location	Description
PEX_LANE_GOOD[10, 8, 6, 4:0]#	O	E16, G17, B2, A4, C5, C6, B6, C7	**Active-Low PCI Express Lane Status Indicator Outputs for Lanes [10, 8, 6, 4-0] (8 Balls)** Directly drives common-anode LED module.
PEX_NT_RESET#	O	A13	**Active-Low Output Used to Propagate Reset in NT Mode**
PEX_PERn[10, 8, 6, 4:0]	CMLRn	P12, P10, P8, P6, P4, P2, M5, K5	**Negative Half of PCI Express Receiver Differential Signal Pairs (8 Balls)**
PEX_PERp[10, 8, 6, 4:0]	CMLRp	R12, R10, R8, R6, R4, R2, M4, K4	**Positive Half of PCI Express Receiver Differential Signal Pairs (8 Balls)**
PEX_PERST#	I	E1	**PCI Express Reset** Used to cause a Fundamental Reset. (Refer to Chapter 5, "Reset and Initialization," for further details.)
PEX_PETn[10, 8, 6, 4:0]	CMLTn	U12, U10, U8, V6, U4, U2, M2, K2	**Negative Half of PCI Express Transmitter Differential Signal Pairs (8 Balls)**
PEX_PETp[10, 8, 6, 4:0]	CMLTp	V12, V10, V8, U6, V4, V2, M1, K1	**Positive Half of PCI Express Transmitter Differential Signal Pairs (8 Balls)**

Table 3-3. PCI Express Signals – 46 Balls *(Cont.)*

Signal Name	Type	Location	Description
PEX_REFCLKn	CMLCLKn	H2	**Negative Half of 100-MHz PCI Express Reference Clock Input Signal Pair** When Dual Clocking is enabled, the PEX_REFCLKn and PEX_REFCLKp signal balls become the Spread-Spectrum Clocking (SSC) domain Clock signals. Refer to Chapter 19, "Dual Clocking Support," for information regarding the use of this signal.
PEX_REFCLKp	CMLCLKp	H1	**Positive Half of 100-MHz PCI Express Reference Clock Input Signal Pair** When Dual Clocking is enabled, the PEX_REFCLKn and PEX_REFCLKp signal balls become the SSC domain Clock signals. Refer to Chapter 19, "Dual Clocking Support," for information regarding the use of this signal.
PEX_REFCLK_CFCn	CMLCLK_CFCn	H5	**Negative Half of 100-MHz PCI Express Spread-Spectrum Reference Clock Input Signal Pair** When Dual Clocking is enabled, the PEX_REFCLK_CFCn and PEX_REFCLK_CFCp signal balls become the Constant Frequency Clocking (CFC) domain Clock signals. Refer to Chapter 19, "Dual Clocking Support," for information regarding the use of this signal.
PEX_REFCLK_CFCp	CMLCLK_CFCp	H4	**Positive Half of 100-MHz PCI Express Spread-Spectrum Reference Clock Input Signal Pair** When Dual Clocking is enabled, the PEX_REFCLK_CFCn and PEX_REFCLK_CFCp signal balls become the CFC domain Clock signals. Refer to Chapter 19, "Dual Clocking Support," for information regarding the use of this signal.

3.4.2 Hot Plug Signals

The PEX 8508 includes nine Hot Plug signals for each PCI Express port (5 ports x 9 signals/port = 45 total signals), defined in Table 3-4. These signals are active only for downstream ports configured at start-up. (Refer to Chapter 11, "Hot Plug Support," for further details.)

Table 3-4. Hot Plug Signals – 45 Balls

Signal Name	Type	Location	Description
HP_ATNLED[4:0]#	O	J16, C12, E15, C9, E4	**Hot Plug Attention LED Output (5 Balls)** Active-Low Slot Control Logic output used to drive the Attention Indicator. Output is set Low to turn On the LED. Enabled when the **Slot Capabilities** register *Attention Indicator Present* bit is set (offset 7Ch[3]=1) and controlled by the **Slot Control** register *Attention Indicator Control* field (offset 80h[7:6]). When software writes to the *Attention Indicator Control* field, a Command Completed interrupt can be generated to notify the Host that the command has been executed. Software must use a Byte or Word Write (and not a Dword Write) to control HP_ATNLED*x*#. When the following conditions exist: • **Slot Capabilities** register *Attention Indicator Present* bit is set (offset 7Ch[3]=1), and • **Slot Control** register *Command Completed Interrupt Enable* bit is not masked (offset 80h[4]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated to the Host. An external current-limiting resistor is required.
HP_BUTTON[4:0]#	I PU[a]	N18, B13, D15, A10, D5	**Hot Plug Attention Button Input (5 Balls)** Active-Low Slot Control Logic input, directly connected to the Attention Button, with input assertion status latched in the **Slot Status** register *Attention Button Pressed* field (offset 80h[16]). Enabled when the **Slot Capabilities** register *Attention Button Present* bit is set (offset 7Ch[0]=1). When the following conditions exist: • HP_BUTTON*x*# is not masked (**Slot Control** register *Attention Button Pressed Enable* bit (offset 80h[0]=1), and • **Slot Capabilities** register *Hot Plug Capable* bit is set (offset 7Ch[6]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated, to notify the Host of intended board insertion or removal. ***Note:*** *HP_BUTTONx# is internally de-bounced, but must remain stable for at least 10 ms.*

Table 3-4. Hot Plug Signals – 45 Balls *(Cont.)*

Signal Name	Type	Location	Description
HP_CLKEN[4:0]#	O	N17, D18, J17, B5, F4	**Reference Clock Enable Output (5 Balls)** Active-Low output that, when enabled, allows external REFCLK to be provided to the slot. Enabled when the **Slot Capabilities** register *Power Controller Present* bit is set (offset 7Ch[1]=1), and controlled by the **Slot Control** register *Power Controller Control* bit (offset 80h[10]). The time delay from HP_PWREN*x*# output assertion to HP_CLKEN*x*# output assertion is programmable (through serial EEPROM load) from 16 ms (default) to 128 ms, in the *HPC Tpepv Delay* field (offset 1E0h[4:3]).
HP_MRL[4:0]#	I PU[a]	H15, D12, E17, D8, B4	**Hot Plug Manually Operated Retention Latch Sensor Input (5 Balls)** Active-Low input that triggers Slot Control Logic. Directly connected to an optional MRL Sensor that is logic High when the latch is not closed. HP_MRL*x*# input assertion enables Hot Plug output sequencing to turn On the slot's power (HP_PWREN*x*# and HP_PWRLED*x*#) and clock (HP_CLKEN*x*#), and de-assert Reset (HP_PERST*x*#) after reset or under software control. Enabled when the **Slot Capabilities** register *MRL Sensor Present* bit is set (offset 7Ch[2]=1). A change in the HP_MRL*x*# Input signal state is latched in the **Slot Status** register *MRL Sensor Changed* bit (offset 80h[18]), and the state change can assert an interrupt to notify the Host of a change in the MRL Sensor state. When the following conditions exist: • HP_MRL*x*# is not masked (**Slot Control** register *MRL Sensor Changed Enable* bit, offset 80h[2]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated. If the associated downstream port connects to a PCI Express board slot that does not implement an MRL Sensor, HP_MRL*x*# is normally connected to HP_PRSNT*x*# and a pull-up resistor, with the common node connected to the PRSNT2# signal(s) at the slot. If the associated downstream port instead connects directly to a device (in which case Hot Plug is not used), pull HP_MRL*x*# Low. ***Note:*** *HP_MRLx# is internally de-bounced, but must remain stable for at least 10 ms. HP_MRLx#, if enabled, is not de-bounced when sampled immediately after reset.*
HP_PERST[4:0]#	O	L17, C14, J18, D6, D3	**Active-Low Reset Output (5 Balls)** Active-Low Hot Plug output used to reset the slot. Controlled by the **Slot Control** register *Power Controller Control* bit (offset 80h[10]).
HP_PRSNT[4:0]#	I PU[a]	U18, D13, G16, A6, B1	**Combination of Hot Plug PRSNT1# and PRSNT2# Input (5 Balls)** Active-Low input connected to the slot's PRSNT2# signal, which on the add-in board connects to the slot's PRSNT1# signal, which is normally grounded on the PRSNT2# signal at the motherboard slot. A change in the HP_PRSNT*x*# Input signal state is latched in the **Slot Status** register *Presence Detect Changed* bit (offset 80h[19]), and the state change can assert an interrupt to notify the Host of board presence or absence. When the following conditions exist: • HP_PRSNT*x*# is not masked (**Slot Control** register *Presence Detect Changed Enable* bit (offset 80h[3]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated. ***Note:*** *HP_PRSNTx# is internally de-bounced, but must remain stable for at least 10 ms.*

Table 3-4. Hot Plug Signals – 45 Balls *(Cont.)*

Signal Name	Type	Location	Description
HP_PWREN[4:0]#	O	K17, C16, G18, A8, D4	**Active-Low Hot Plug Power Enable Output (5 Balls)** Active-Low Slot Control Logic output that controls the slot power state. When this signal is Low, power is enabled to the slot. Enabled when the **Slot Capabilities** register *Power Controller Present* bit is set (offset 7Ch[1]=1). When software turns the slot's Power Controller On or Off (**Slot Control** register *Power Controller Control* bit, offset 80h[10]), a Command Completed interrupt can be generated to notify the Host that the command has been executed. When the following conditions exist: • **Slot Control** register *Command Completed Interrupt Enable* bit is not masked (offset 80h[4]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated to the Host. When HP_MRL*x*# is enabled [**Slot Capabilities** register *MRL Sensor Present* bit is set (offset 7Ch[2]=1)], HP_MRL*x*# input assertion enables Hot Plug output sequencing to turn On the slot's power, by asserting HP_PWREN*x*# after reset or under software control.
HP_PWRFLT[4:0]#	I PU[a]	T18, B18, F16, A9, C4	**Hot Plug Power Fault Input (5 Balls)** Active-Low input that indicates the slot's external Power Controller detected a power fault on one or more supply rails. Enabled when the **Slot Capabilities** register *Power Controller Present* bit is set (offset 7Ch[1]=1), and input assertion status is latched in the **Slot Status** register *Power Fault Detected* bit (offset 80h[17]). When the following conditions exist: • HP_PWRFLT*x*# is not masked (**Slot Control** register *Power Fault Detector Enable* bit (offset 80h[1]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated, to notify the Host of a power fault. ***Note:*** *If HP_PWRENx# and HP_CLKENx# are not used, HP_PWRFLTx# can be used as a general-purpose input with status reflected in the **Slot Status** register Power Fault Detected (offset 80h[17]), provided the **Slot Capabilities** register Power Controller Present bit is set (offset 7Ch[1]=1).*

Table 3-4. Hot Plug Signals – 45 Balls *(Cont.)*

Signal Name	Type	Location	Description
HP_PWRLED[4:0]#	O	T17, C15, H16, A5, E3	**Hot Plug Power LED Output (5 Balls)** Active-Low Slot Control Logic output used to drive the Power Indicator. This output is set Low to turn On the LED. Enabled when the **Slot Capabilities** register *Power Indicator Present* bit is set (offset 7Ch[4]=1), and controlled by the **Slot Status** register *Power Indicator Control* field (offset 80h[9:8]). When software writes to the *Power Indicator Control* field, a Command Completed interrupt can be generated to notify the Host that the command has been executed. Software must use a Byte or Word Write (and not a Dword Write) to control HP_PWRLED*x*#. When the following conditions exist: • **Slot Capabilities** register *Power Indicator Present* bit is set (offset 7Ch[4]=1), and • **Slot Control** register *Command Completed Interrupt Enable* bit is not masked (offset 80h[4]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated to the Host. An external current-limiting resistor is required.

a. If Hot Plug outputs (including HP_PERSTx#) are used and HP_MRLx# input is not used, pull HP_MRLx# input Low so that Hot Plug outputs (including HP_PERSTx#) will properly sequence if the serial EEPROM is blank or missing. Default register values enable HP_MRLx#, which must then be asserted to cause Hot Plug outputs to toggle (for example, to de-assert HP_PERSTx# and assert HP_PWRLEDx#).

3.4.3 Serial EEPROM Signals

The PEX 8508 includes five signals for interfacing to a serial EEPROM, defined in Table 3-5.

Table 3-5. Serial EEPROM Signals – 5 Balls

Signal Name	Type	Location	Description
EE_CS#	O	T16	**Serial EEPROM Active-Low Chip Select Output**
EE_DI	O	N16	**PEX 8508 Output to Serial EEPROM Data Input**
EE_DO	I PU	L16	**PEX 8508 Input from Serial EEPROM Data Output**
EE_PR#	I PU	R16	**Serial EEPROM Present Active-Low Input** Must be tied to VSS (GND) to indicate serial EEPROM presence.
EE_SK	O	M16	**7.8 MHz Serial EEPROM Clock Output**

3.4.4 Strapping Signals

The PEX 8508 Strapping signals, defined in Table 3-6, set the configuration of upstream and NT Port assignment, port width, spread-spectrum clocking, and various setup and test modes. These balls must be tied High to VDD33 or Low to VSS (GND). After a Fundamental Reset, the **Link Capabilities**, **Debug Control**, and **Port Configuration** registers capture ball status. Strapping ball Configuration data can be changed by writing new data to these registers from the serial EEPROM.

Table 3-6. Strapping Signals – 22 Balls

Signal Name	Type	Location	Description
STRAP_FACTORY_TEST[2:1]#	I STRAP	G4, D1	***Factory Test Only*** **(2 Balls)** Must be tied High.
STRAP_MODE_SEL[1:0]	I STRAP	A16, D10	**Mode Selects (2 Balls)** Register/Bits – **Debug Control** register *Mode Select* field (Port 0, offset 1DCh[19:18]) LL = ***Reserved*** LH = NT Intelligent Adapter mode HL = NT Dual-Host mode HH = Transparent mode
STRAP_NT_UPSTRM_PORTSEL[3:0]	I STRAP	C17, B14, C13, B16	**Select Non-Transparent Upstream Port (4 Balls)** Register/Bits – **Debug Control** register *NT Port Number* field (Port 0, offset 1DCh[27:24]) LLLL = Port 0 LLLH = Port 1 LLHL = Port 2 LLHH = Port 3 LHLL = Port 4 If NT Port is not used, set to HHHH. All other values are ***reserved***.
STRAP_PORTCFG[4:0]	I STRAP	G15, L18, P18, K16, P17	**Strapping Signals to Select Number of Lanes in Ports 0, 1, 2, 3, 4 (5 Balls)** Register/Bits – **Port Configuration** register *Port Configuration* field (Port 0, offset 224h[4:0]) LLLLL = x2, x2, x2, x2 LLLHL = x4, x4 LLLHH = x4, x2, x2 LLHLL = x4, x2, x1, x1 LLHLH = x4, x1, x1, x2 LLHHL = x4, x1, x2, x1 LHLLL = x4, x1, x1, x1, x1 LHLLH = x2, x2, x2, x1, x1 All other configurations are ***reserved*** and default to x2, x2, x2, x2.

Table 3-6. Strapping Signals – 22 Balls *(Cont.)*

Signal Name	Type	Location	Description
STRAP_SSC_XING_ENA	I STRAP	F3	**Spread Spectrum Clocking (SSC) Crossing Enable** Allows the SSC REFCLK clock source on upstream Port 0, with a x4 lane width and Constant Frequency Clocking (CFC) REFCLK source on downstream devices. (Refer to Chapter 19, "Dual Clocking Support," for further details.) L = Disabled (pulled down to VSS) H = Enabled (pulled up to VDD33A)
STRAP_TESTMODE[3:0]	I STRAP	B17, A17, C11, B12	**Test Mode Select (4 Balls)** ***Factory Test Only*** Register/Bits – **Physical Layer Test** register (Port 0, offset 228h) HHHH = Default (Test Modes are disabled)
STRAP_UPSTRM_PORTSEL[3:0]	I STRAP	D7, C8, B8, B9	**Strapping Signals to Select Upstream Port (4 Balls)** Register/Bits – **Debug Control** register *Upstream Port ID* field (Port 0, offset 1DCh[10:8]) LLLL = Port 0 LLLH = Port 1 LLHL = Port 2 LLHH = Port 3 LHLL = Port 4 All other values are ***reserved***.

3.4.5 JTAG Interface Signals

The PEX 8508 includes five interface signals for performing JTAG boundary scan, defined in Table 3-7. (Refer to Chapter 22, "Test and Debug," for further details.)

Table 3-7. JTAG Interface Signals – 5 Balls

Signal Name	Type	Location	Description
JTAG_TCK	I PU	D9	**JTAG Test Clock Input** JTAG Test Access Port (TAP) Controller clock source. Frequency can be from 0 to 10 MHz.
JTAG_TDI	I PU	A12	**JTAG Test Data Input** Serial input to the JTAG TAP Controller for test instructions and data.
JTAG_TDO	O	A14	**JTAG Test Data Output** Serial output from the JTAG TAP Controller test instructions and data.
JTAG_TMS	I PU	C10	**JTAG Test Mode Select** When High, JTAG Test mode is enabled. Input decoded by the JTAG TAP Controller, to control test operations.
JTAG_TRST#	I PU	B10	**JTAG Test Reset** Active-Low input used to reset the Test Access Port. Tie to ground through a 1.5K-Ohm resistor, to hold the JTAG TAP Controller in the *Test-Logic-Reset* state, which enables standard logic operation. When JTAG functionality is not used, the JTAG_TRST# input should be pulled or driven Low, to place the JTAG TAP Controller into the *Test-Logic-Reset* state, which disables the test logic and enables standard logic operation. Alternatively, if JTAG_TRST# input is High, the JTAG TAP Controller can be placed into the *Test-Logic-Reset* state by initializing the JTAG TAP Controller's **Instruction** register to contain the *IDCODE* instruction, or by holding the JTAG_TMS input High for at least five rising edges of the JTAG_TCK input.

3.4.6 I²C Interface Signals

The I²C interface is described Chapter 20, "I2C Interface Operation."

Table 3-8. I²C Interface Configuration Signals – 5 Balls

Signal Name	Type	Location	Description
I2C_ADDR[2:0]	I PU	P15, R14, T14	**I²C Slave Address Bits 2 through 0 (3 Balls)** Used to set the PEX 8508 Slave address on the I²C Bus. If I²C output is used, I2C_ADDR[2:0] should be strapped to a unique address, to avoid address conflict with any other I²C devices (on the same I²C Bus segment) that have the upper four bits of their 7-bit I²C Slave address also set to 1110b. Must be pulled High to VDD33 or Low to VSS (GND) through external resistors.
I2C_SCL	I/O OD	T15	**I²C Serial Clock Line** I²C Clock source.
I2C_SDA	I/O OD	P16	**I²C Serial Data Output** Transmits and receives I²C data.

3.4.7 Fatal Error Signal

Table 3-9. Fatal Error Signal – 1 Ball

Signal Name	Type	Location	Description
FATAL_ERR#	O	R15	**Fatal Error** Asserted Low when a Fatal error is detected in the PEX 8508. The **Device Control** register *Fatal Error Detected* bit (offset 70h[18]) is set.

3.4.8 Wake Signal

Table 3-10. Wake Signal – 1 Ball

Signal Name	Type	Location	Description
WAKE#	I	F1	**Wake Up** Sideband signal used to wake up the PEX 8508 when it is in the L2 (D3cold) power state. The PEX 8508 sends a Beacon signal on Lane 0 of the upstream port when WAKE# is asserted. (Refer to Chapter 21, "WAKE# and Beacon Functionality," for further details.)

3.4.9 No Connect Signals

Caution: ***Do not connect the following PEX 8508 balls to board electrical paths, because these balls are not connected within the PEX 8508.***

Table 3-11. No Connect Signals – 21 Balls

Signal Name	Type	Location	Description
N/C	***Reserved***	A1, A2, B3, C3, C18, D14, D16, D17, E18, F15, J15, K15, L15, M15, N4, U17, V17	**No Connect (17 Balls)** Do not connect these balls to board electrical paths.
NC_PROCMON	***Reserved***	D2	**No Connect** Do not connect these balls to board electrical paths.
NC_SPARE[2:0]	***Reserved***	N15, K18, D11	**No Connect (3 Balls)** Do not connect these balls to board electrical paths.

3.4.10 Power and Ground Signals

Table 3-12. Power and Ground Signals – 145 Balls

Signal Name	Type	Location	Description
VDD10	CPWR	E5, E7, E8, E11, E12, E14, F14, G5, G14, H14, K14, L14, N5, N14, P5, P9, P13, P14	**1.0V Power for Core Logic (18 Balls)**
VDD10A	APWR	J4, M3, R7, R11	**1.0V Power for SerDes Analog Circuits (4 Balls)**
VDD10S	SPWR	J3, L3, N3, T3, T5, T7, T9, T11, T13, U14, V14	**1.0V Power for SerDes Digital Circuits (11 Balls)**
VDD10X	XPWR	G3, L4, P7, P11	**1.0V Power for SerDes Auxiliary Circuits (4 Balls)**
VDD33	I/OPWR	B7, B11, B15, C2, F17, H17, M17, R17, U16	**3.3V Power for I/O Logic Functions (9 Balls)**
VDD33A	PLLPWR	F2	**3.3V Power for PLL Circuits**
VDD33X	AUXPWR	G1	**3.3V Power for Auxiliary Circuits**
VSS	GND	A3, A7, A11, A15, A18, C1, E6, E9, E10, E13, F5, F18, G2, H3, H18, J1, J2, J5, J14, K3, L2, L5, M14, M18, N1, N2, P1, P3, R1, R3, R5, R9, R13, R18, T1, T2, T4, T6, T8, T10, T12, U1, U3, U5, U7, U9, U11, U13, U15, V1, V5, V15, V16, V18	**Ground Connections (54 Balls)**
VSS_THERMAL	Thermal-GND	G7, G8, G9, G10, G11, G12, H7, H8, H9, H10, H11, H12, J7, J8, J9, J10, J11, J12, K7, K8, K9, K10, K11, K12, L7, L8, L9, L10, L11, L12, M7, M8, M9, M10, M11, M12	**Thermal-Ball Ground Connection (36 Balls)**
VSSA_PLL	PLL_GND	E2	**PLL Ground Connection**
VTT_PEX[5:0]	Supply	V13, V11, V9, V7, V3, L1	**SerDes Termination Supply (6 Balls)**[a] Tied to SerDes termination supply voltage (typically 1.5V).

a. PEX_PETn/p[x] SerDes termination supply voltage controls transmitter Common mode voltage (V_{TX-CM}) value and output voltage swing ($V_{TX-DIFFp}$), per the following formula:

$$V_{TX-CM} = V_{TT} - V_{TX-DIFFp}$$

3.5 Ball Assignments by Number

Note: *For "Type" definitions, refer to Table 3-1.*

Table 3-13. PEX 8508 Switch Ball Assignments by Number

PEX 8508 Balls				Comment
Num	**Name**	**Type**	**Signal Group**	
A1	N/C	***Reserved***	No Connect	
A2	N/C	***Reserved***	No Connect	
A3	VSS	GND	Ground	
A4	PEX_LANE_GOOD4#	O	Lane Status	
A5	HP_PWRLED1#	O	Hot Plug	
A6	HP_PRSNT1#	I, PU	Hot Plug	
A7	VSS	GND	Ground	
A8	HP_PWREN1#	O	Hot Plug	
A9	HP_PWRFLT1#	I, PU	Hot Plug	
A10	HP_BUTTON1#	I, PU	Hot Plug	
A11	VSS	GND	Ground	
A12	JTAG_TDI	I, PU	JTAG	
A13	PEX_NT_RESET#	O	PEX Control	
A14	JTAG_TDO	O	JTAG	
A15	VSS	GND	Ground	
A16	STRAP_MODE_SEL1	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_MODE_SEL[1:0]
A17	STRAP_TESTMODE2	I	STRAP	Strapping Ball – tie H
A18	VSS	GND	Ground	
B1	HP_PRSNT0#	I, PU	Hot Plug	
B2	PEX_LANE_GOOD6#	O	Lane Status	
B3	N/C	***Reserved***	No Connect	
B4	HP_MRL0#	I, PU	Hot Plug	
B5	HP_CLKEN1#	O	Hot Plug	
B6	PEX_LANE_GOOD1#	O	Lane Status	
B7	VDD33	I/OPWR	Power	
B8	STRAP_UPSTRM_PORTSEL1	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[3:0]
B9	STRAP_UPSTRM_PORTSEL0	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[3:0]
B10	JTAG_TRST#	I, PU	JTAG	
B11	VDD33	I/OPWR	Power	
B12	STRAP_TESTMODE0	I	STRAP	Strapping Ball – tie H
B13	HP_BUTTON3#	I, PU	Hot Plug	
B14	STRAP_NT_UPSTRM_PORTSEL2	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_NT_UPSTRM_PORTSEL[3:0]

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	Name	Type	Signal Group	
B15	VDD33	I/OPWR	Power	
B16	STRAP_NT_UPSTRM_PORTSEL0	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_NT_UPSTRM_PORTSEL[3:0]
B17	STRAP_TESTMODE3	I	STRAP	Strapping Ball – tie H
B18	HP_PWRFLT3#	I, PU	Hot Plug	
C1	VSS	GND	Ground	
C2	VDD33	I/OPWR	Power	
C3	N/C	***Reserved***	No Connect	
C4	HP_PWRFLT0#	I, PU	Hot Plug	
C5	PEX_LANE_GOOD3#	O	Lane Status	
C6	PEX_LANE_GOOD2#	O	Lane Status	
C7	PEX_LANE_GOOD0#	O	Lane Status	
C8	STRAP_UPSTRM_PORTSEL2	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[3:0]
C9	HP_ATNLED1#	O	Hot Plug	
C10	JTAG_TMS	I, PU	JTAG	
C11	STRAP_TESTMODE1	I	STRAP	Strapping Ball – tie H
C12	HP_ATNLED3#	O	Hot Plug	
C13	STRAP_NT_UPSTRM_PORTSEL1	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_NT_UPSTRM_PORTSEL[3:0]
C14	HP_PERST3#	O	Hot Plug	
C15	HP_PWRLED3#	O	Hot Plug	
C16	HP_PWREN3#	O	Hot Plug	
C17	STRAP_NT_UPSTRM_PORTSEL3	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_NT_UPSTRM_PORTSEL[3:0]
C18	N/C	***Reserved***	No Connect	
D1	STRAP_FACTORY_TEST1#	I	STRAP	Strapping Ball – tie H
D2	NC_PROCMON	***Reserved***	No Connect	
D3	HP_PERST0#	O	Hot Plug	
D4	HP_PWREN0#	O	Hot Plug	
D5	HP_BUTTON0#	I, PU	Hot Plug	
D6	HP_PERST1#	O	Hot Plug	
D7	STRAP_UPSTRM_PORTSEL3	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[3:0]
D8	HP_MRL1#	I, PU	Hot Plug	
D9	JTAG_TCK	I, PU	JTAG	
D10	STRAP_MODE_SEL0	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_MODE_SEL[1:0]
D11	NC_SPARE0	***Reserved***	No Connect	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	Name	Type	Signal Group	
D12	HP_MRL3#	I, PU	Hot Plug	
D13	HP_PRSNT3#	I, PU	Hot Plug	
D14	N/C	***Reserved***	No Connect	
D15	HP_BUTTON2#	I, PU	Hot Plug	
D16	N/C	***Reserved***	No Connect	
D17	N/C	***Reserved***	No Connect	
D18	HP_CLKEN3#	O	Hot Plug	
E1	PEX_PERST#	I	PEX Control	
E2	VSSA_PLL	PLL_GND	Ground	Analog GND for PLL circuit
E3	HP_PWRLED0#	O	Hot Plug	
E4	HP_ATNLED0#	O	Hot Plug	
E5	VDD10	CPWR	Power	
E6	VSS	GND	Ground	
E7	VDD10	CPWR	Power	
E8	VDD10	CPWR	Power	
E9	VSS	GND	Ground	
E10	VSS	GND	Ground	
E11	VDD10	CPWR	Power	
E12	VDD10	CPWR	Power	
E13	VSS	GND	Ground	
E14	VDD10	CPWR	Power	
E15	HP_ATNLED2#	O	Hot Plug	
E16	PEX_LANE_GOOD10#	O	Lane Status	
E17	HP_MRL2#	I, PU	Hot Plug	
E18	N/C	***Reserved***	No Connect	
F1	WAKE#	I	Wake	
F2	VDD33A	PLLPWR	Power	
F3	STRAP_SSC_XING_ENA	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_SSC_XING_ENA
F4	HP_CLKEN0#	O	Hot Plug	
F5	VSS	GND	Ground	
F14	VDD10	CPWR	Power	
F15	N/C	***Reserved***	No Connect	
F16	HP_PWRFLT2#	I, PU	Hot Plug	
F17	VDD33	I/OPWR	Power	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	Name	Type	Signal Group	
F18	VSS	GND	Ground	
G1	VDD33X	AUXPWR	Power	
G2	VSS	GND	Ground	
G3	VDD10X	XPWR	Power	
G4	STRAP_FACTORY_TEST2#	I	STRAP	Strapping Ball – tie H
G5	VDD10	CPWR	Power	
G7	VSS_THERMAL	Thermal-GND	Ground	
G8	VSS_THERMAL	Thermal-GND	Ground	
G9	VSS_THERMAL	Thermal-GND	Ground	
G10	VSS_THERMAL	Thermal-GND	Ground	
G11	VSS_THERMAL	Thermal-GND	Ground	
G12	VSS_THERMAL	Thermal-GND	Ground	
G14	VDD10	CPWR	Power	
G15	STRAP_PORTCFG4	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[4:0]
G16	HP_PRSNT2#	I, PU	Hot Plug	
G17	PEX_LANE_GOOD8#	O	Lane Status	
G18	HP_PWREN2#	O	Hot Plug	
H1	PEX_REFCLKp	CMLCLKp	SerDes / Clock	
H2	PEX_REFCLKn	CMLCLKn	SerDes / Clock	
H3	VSS	GND	Ground	
H4	PEX_REFCLK_CFCp	CMLCLK_CFCp	SerDes / Clock	
H5	PEX_REFCLK_CFCn	CMLCLK_CFCn	SerDes / Clock	
H7	VSS_THERMAL	Thermal-GND	Ground	
H8	VSS_THERMAL	Thermal-GND	Ground	
H9	VSS_THERMAL	Thermal-GND	Ground	
H10	VSS_THERMAL	Thermal-GND	Ground	
H11	VSS_THERMAL	Thermal-GND	Ground	
H12	VSS_THERMAL	Thermal-GND	Ground	
H14	VDD10	CPWR	Power	
H15	HP_MRL4#	I, PU	Hot Plug	
H16	HP_PWRLED2#	O	Hot Plug	
H17	VDD33	I/OPWR	Power	
H18	VSS	GND	Ground	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	Name	Type	Signal Group	
J1	VSS	GND	Ground	
J2	VSS	GND	Ground	
J3	VDD10S	SPWR	Power	
J4	VDD10A	APWR	Power	
J5	VSS	GND	Ground	
J7	VSS_THERMAL	Thermal-GND	Ground	
J8	VSS_THERMAL	Thermal-GND	Ground	
J9	VSS_THERMAL	Thermal-GND	Ground	
J10	VSS_THERMAL	Thermal-GND	Ground	
J11	VSS_THERMAL	Thermal-GND	Ground	
J12	VSS_THERMAL	Thermal-GND	Ground	
J14	VSS	GND	Ground	
J15	N/C	***Reserved***	No Connect	
J16	HP_ATNLED4#	O	Hot Plug	
J17	HP_CLKEN2#	O	Hot Plug	
J18	HP_PERST2#	O	Hot Plug	
K1	PEX_PETp0	CMLTp	SerDes	
K2	PEX_PETn0	CMLTn	SerDes	
K3	VSS	GND	Ground	
K4	PEX_PERp0	CMLRp	SerDes	
K5	PEX_PERn0	CMLRn	SerDes	
K7	VSS_THERMAL	Thermal-GND	Ground	
K8	VSS_THERMAL	Thermal-GND	Ground	
K9	VSS_THERMAL	Thermal-GND	Ground	
K10	VSS_THERMAL	Thermal-GND	Ground	
K11	VSS_THERMAL	Thermal-GND	Ground	
K12	VSS_THERMAL	Thermal-GND	Ground	
K14	VDD10	CPWR	Power	
K15	N/C	***Reserved***	No Connect	
K16	STRAP_PORTCFG1	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[4:0]
K17	HP_PWREN4#	O	Hot Plug	
K18	NC_SPARE1	***Reserved***	No Connect	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	**Name**	**Type**	**Signal Group**	
L1	VTT_PEX0	Supply	Power	
L2	VSS	GND	Ground	
L3	VDD10S	SPWR	Power	
L4	VDD10X	XPWR	Power	
L5	VSS	GND	Ground	
L7	VSS_THERMAL	Thermal-GND	Ground	
L8	VSS_THERMAL	Thermal-GND	Ground	
L9	VSS_THERMAL	Thermal-GND	Ground	
L10	VSS_THERMAL	Thermal-GND	Ground	
L11	VSS_THERMAL	Thermal-GND	Ground	
L12	VSS_THERMAL	Thermal-GND	Ground	
L14	VDD10	CPWR	Power	
L15	N/C	***Reserved***	No Connect	
L16	EE_DO	I, PU	Serial EEPROM	Connected to data output of serial EEPROM
L17	HP_PERST4#	O	Hot Plug	
L18	STRAP_PORTCFG3	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[4:0]
M1	PEX_PETp1	CMLTp	SerDes	
M2	PEX_PETn1	CMLTn	SerDes	
M3	VDD10A	APWR	Power	
M4	PEX_PERp1	CMLRp	SerDes	
M5	PEX_PERn1	CMLRn	SerDes	
M7	VSS_THERMAL	Thermal-GND	Ground	
M8	VSS_THERMAL	Thermal-GND	Ground	
M9	VSS_THERMAL	Thermal-GND	Ground	
M10	VSS_THERMAL	Thermal-GND	Ground	
M11	VSS_THERMAL	Thermal-GND	Ground	
M12	VSS_THERMAL	Thermal-GND	Ground	
M14	VSS	GND	Ground	
M15	N/C	***Reserved***	No Connect	
M16	EE_SK	O	Serial EEPROM	
M17	VDD33	I/OPWR	Power	
M18	VSS	GND	Ground	
N1	VSS	GND	Ground	
N2	VSS	GND	Ground	
N3	VDD10S	SPWR	Power	
N4	N/C	***Reserved***	No Connect	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	Name	Type	Signal Group	
N5	VDD10	CPWR	Power	
N14	VDD10	CPWR	Power	
N15	NC_SPARE2	***Reserved***	No Connect	
N16	EE_DI	O	Serial EEPROM	Connected to data input of serial EEPROM
N17	HP_CLKEN4#	O	Hot Plug	
N18	HP_BUTTON4#	I, PU	Hot Plug	
P1	VSS	GND	Ground	
P2	PEX_PERn2	CMLRn	SerDes	
P3	VSS	GND	Ground	
P4	PEX_PERn3	CMLRn	SerDes	
P5	VDD10	CPWR	Power	
P6	PEX_PERn4	CMLRn	SerDes	
P7	VDD10X	XPWR	Power	
P8	PEX_PERn6	CMLRn	SerDes	
P9	VDD10	CPWR	Power	
P10	PEX_PERn8	CMLRn	SerDes	
P11	VDD10X	XPWR	Power	
P12	PEX_PERn10	CMLRn	SerDes	
P13	VDD10	CPWR	Power	
P14	VDD10	CPWR	Power	
P15	I2C_ADDR2	I, PU	I^2C	
P16	I2C_SDA	I/O, OD	I^2C	
P17	STRAP_PORTCFG0	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[4:0]
P18	STRAP_PORTCFG2	I	STRAP	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[4:0]
R1	VSS	GND	Ground	
R2	PEX_PERp2	CMLRp	SerDes	
R3	VSS	GND	Ground	
R4	PEX_PERp3	CMLRp	SerDes	
R5	VSS	GND	Ground	
R6	PEX_PERp4	CMLRp	SerDes	
R7	VDD10A	APWR	Power	
R8	PEX_PERp6	CMLRp	SerDes	
R9	VSS	GND	Ground	
R10	PEX_PERp8	CMLRp	SerDes	
R11	VDD10A	APWR	Power	
R12	PEX_PERp10	CMLRp	SerDes	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	Name	Type	Signal Group	
R13	VSS	GND	Ground	
R14	I2C_ADDR1	I, PU	I^2C	
R15	FATAL_ERR#	O	Fatal Error	
R16	EE_PR#	I, PU	Serial EEPROM	
R17	VDD33	I/OPWR	Power	
R18	VSS	GND	Ground	
T1	VSS	GND	Ground	
T2	VSS	GND	Ground	
T3	VDD10S	SPWR	Power	
T4	VSS	GND	Ground	
T5	VDD10S	SPWR	Power	
T6	VSS	GND	Ground	
T7	VDD10S	SPWR	Power	
T8	VSS	GND	Ground	
T9	VDD10S	SPWR	Power	
T10	VSS	GND	Ground	
T11	VDD10S	SPWR	Power	
T12	VSS	GND	Ground	
T13	VDD10S	SPWR	Power	
T14	I2C_ADDR0	I, PU	I^2C	
T15	I2C_SCL	I/O, OD	I^2C	
T16	EE_CS#	O	Serial EEPROM	
T17	HP_PWRLED4#	O	Hot Plug	
T18	HP_PWRFLT4#	I, PU	Hot Plug	
U1	VSS	GND	Ground	
U2	PEX_PETn2	CMLTn	SerDes	
U3	VSS	GND	Ground	
U4	PEX_PETn3	CMLTn	SerDes	
U5	VSS	GND	Ground	
U6	PEX_PETp4	CMLTp	SerDes	
U7	VSS	GND	Ground	
U8	PEX_PETn6	CMLTn	SerDes	
U9	VSS	GND	Ground	
U10	PEX_PETn8	CMLTn	SerDes	
U11	VSS	GND	Ground	
U12	PEX_PETn10	CMLTn	SerDes	
U13	VSS	GND	Ground	

Table 3-13. PEX 8508 Switch Ball Assignments by Number *(Cont.)*

PEX 8508 Balls				Comment
Num	**Name**	**Type**	**Signal Group**	
U14	VDD10S	SPWR	Power	
U15	VSS	GND	Ground	
U16	VDD33	I/OPWR	Power	
U17	N/C	***Reserved***	No Connect	
U18	HP_PRSNT4#	I, PU	Hot Plug	
V1	VSS	GND	Ground	
V2	PEX_PETp2	CMLTp	SerDes	
V3	VTT_PEX1	Supply	Power	
V4	PEX_PETp3	CMLTp	SerDes	
V5	VSS	GND	Ground	
V6	PEX_PETn4	CMLTn	SerDes	
V7	VTT_PEX2	Supply	Power	
V8	PEX_PETp6	CMLTp	SerDes	
V9	VTT_PEX3	Supply	Power	
V10	PEX_PETp8	CMLTp	SerDes	
V11	VTT_PEX4	Supply	Power	
V12	PEX_PETp10	CMLTp	SerDes	
V13	VTT_PEX5	Supply	Power	
V14	VDD10S	SPWR	Power	
V15	VSS	GND	Ground	
V16	VSS	GND	Ground	
V17	N/C	***Reserved***	No Connect	
V18	VSS	GND	Ground	

3.6 PEX 8508 Physical Layout

Figure 3-1. PEX 8508 Physical Ball Assignment (See-Through Top View)

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	
A	N/C	N/C	VSS	PEX_LANE_GOOD4#	HP_PWRLED1#	HP_PRSNT1#	VSS	HP_PWREN1#	HP_PWRFLT1#	HP_BUTTON1#	VSS	JTAG_TDI	PEX_NT_RESET#	JTAG_TDO	VSS	STRAP_MODE_SEL1	STRAP_TESTMODE2	VSS	A
B	HP_PRSNT0#	PEX_LANE_GOOD6#	N/C	HP_MRL0#	HP_CLKEN1#	PEX_LANE_GOOD1#	VDD33	STRAP_UPSTRM_PORTSEL1	STRAP_UPSTRM_PORTSEL0	JTAG_TRST#	VDD33	STRAP_TESTMODE0	HP_BUTTON3#	STRAP_NT_UPSTRM_PORTSEL2	VDD33	STRAP_NT_UPSTRM_PORTSEL0	STRAP_TESTMODE3	HP_PWRFLT3#	B
C	VSS	VDD33	N/C	HP_PWRFLT0#	PEX_LANE_GOOD3#	PEX_LANE_GOOD2#	PEX_LANE_GOOD0#	STRAP_UPSTRM_PORTSEL2	HP_ATNLED1#	JTAG_TMS	STRAP_TESTMODE1	HP_ATNLED3#	STRAP_NT_UPSTRM_PORTSEL1	HP_PERST3#	HP_PWRLED3#	HP_PWREN3#	STRAP_NT_UPSTRM_PORTSEL3	N/C	C
D	STRAP_FACTORY_TEST1#	NC_PROCMON	HP_PERST0#	HP_PWREN0#	HP_BUTTON0#	HP_PERST1#	STRAP_UPSTRM_PORTSEL3	HP_MRL1#	JTAG_TCK	STRAP_MODE_SEL0	NC_SPARE0	HP_MRL3#	HP_PRSNT3#	N/C	HP_BUTTON2#	N/C	N/C	HP_CLKEN3#	D
E	PEX_PERST#	VSSA_PLL	HP_PWRLED0#	HP_ATNLED0#	VDD10	VSS	VDD10	VDD10	VSS	VSS	VDD10	VDD10	VSS	VDD10	HP_ATNLED2#	PEX_LANE_GOOD10#	HP_MRL2#	N/C	E
F	WAKE#	VDD33A	STRAP_SSC_XING_ENA	HP_CLKEN0#	VSS									VDD10	N/C	HP_PWRFLT2#	VDD33	VSS	F
G	VDD33X	VSS	VDD10X	STRAP_FACTORY_TEST2#	VDD10			VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VDD10	STRAP_PORTCFG4	HP_PRSNT2#	PEX_LANE_GOOD8#	HP_PWREN2#	G
H	PEX_REFCLKp	PEX_REFCLKn	VSS	PEX_REFCLK_CFCp	PEX_REFCLK_CFCn			VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VDD10	HP_MRL4#	HP_PWRLED2#	VDD33	VSS	H
J	VSS	VSS	VDD10S	VDD10A	VSS			VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS	N/C	HP_ATNLED4#	HP_CLKEN2#	HP_PERST2#	J
K	PEX_PETp0	PEX_PETn0	VSS	PEX_PERp0	PEX_PERn0			VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VDD10	N/C	STRAP_PORTCFG1	HP_PWREN4#	NC_SPARE1	K
L	VTT_PEX0	VSS	VDD10S	VDD10X	VSS			VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VDD10	N/C	EE_DO	HP_PERST4#	STRAP_PORTCFG3	L
M	PEX_PETp1	PEX_PETn1	VDD10A	PEX_PERp1	PEX_PERn1			VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS_THERMAL	VSS	N/C	EE_SK	VDD33	VSS	M
N	VSS	VSS	VDD10S	N/C	VDD10									VDD10	NC_SPARE2	EE_DI	HP_CLKEN4#	HP_BUTTON4#	N
P	VSS	PEX_PERn2	VSS	PEX_PERn3	VDD10	PEX_PERn4	VDD10X	PEX_PERn6	VDD10	PEX_PERn8	VDD10X	PEX_PERn10	VDD10	VDD10	I2C_ADDR2	I2C_SDA	STRAP_PORTCFG0	STRAP_PORTCFG2	P
R	VSS	PEX_PERp2	VSS	PEX_PERp3	VSS	PEX_PERp4	VDD10A	PEX_PERp6	VSS	PEX_PERp8	VDD10A	PEX_PERp10	VSS	I2C_ADDR1	FATAL_ERR#	EE_PR#	VDD33	VSS	R
T	VSS	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	I2C_ADDR0	I2C_SCL	EE_CS#	HP_PWRLED4#	HP_PWRFLT4#	T
U	VSS	PEX_PETn2	VSS	PEX_PETn3	VSS	PEX_PETp4	VSS	PEX_PETn6	VSS	PEX_PETn8	VSS	PEX_PETn10	VSS	VDD10S	VSS	VDD33	N/C	HP_PRSNT4#	U
V	VSS	PEX_PETp2	VTT_PEX1	PEX_PETp3	VSS	PEX_PETn4	VTT_PEX2	PEX_PETp6	VTT_PEX3	PEX_PETp8	VTT_PEX4	PEX_PETp10	VTT_PEX5	VDD10S	VSS	VSS	N/C	VSS	V
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 4 PEX 8508 Functional Overview

4.1 Architecture

The PEX 8508 switches are designed with a flexible and modular architecture. Figure 4-1 illustrates a block diagram of the PEX 8508.

Figure 4-1. PEX 8508 Block Diagram

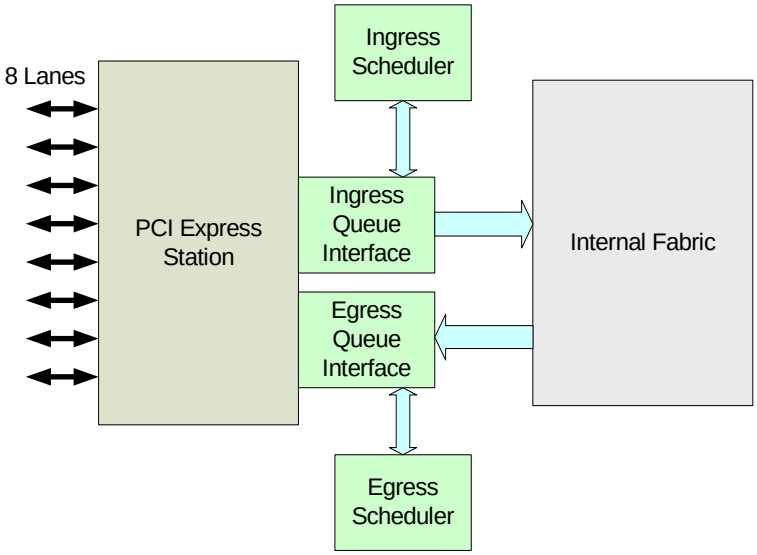

4.1.1 Ingress and Egress Functions

The Crossbar Switch Ingress Queue interfaces the PCI Express station to the Crossbar Switch Fabric (internal fabric). The queue contains a centralized packet buffer for all incoming ports, ingress port scheduler, and the Crossbar Switch Fabric scheduler.

The Crossbar Switch Egress Queue interfaces the non-blocking internal fabric to the PCI Express station. The queue contains a centralized packet buffer for all outgoing ports and the egress port scheduler.

All ingress traffic flows from the PCI Express station through the Crossbar Switch Ingress queue, the internal fabric, and the Crossbar Switch Egress queue, then returns to the PCI Express station. All ingress and egress scheduler modules include VC schedulers for each port. The ingress scheduler supports a port-width-based arbitration scheme. The egress scheduler supports a device-specific port arbitration scheme, to avoid port starvation.

4.1.2 Station and Port Functions

Each port implements the *PCI Express Base r1.1* Physical, Data Link, and Transaction layers. The PCI Express station can support up to 8 integrated Serializer/De-serializer (SerDes) modules. The SerDes modules provide the 8 PCI Express hardware interface lanes. These lanes can be combined, for a total of two to five PCI Express ports.

From the system model viewpoint, each PCI Express port is a virtual PCI-to-PCI bridge device with its own set of PCI Express Configuration registers. Any one of the ports can be the upstream port, and the remaining enabled ports are the downstream ports. The BIOS can configure the other ports, by way of the upstream port, using Conventional PCI enumeration with Configuration accesses.

The PEX 8508 port width is configurable by way of Strapped signal balls, serial EEPROM, or I^2C interface download after reset. The final port width can also be made narrower by auto-lane width negotiation, as described in the *PCI Express Base r1.1*.

4.1.2.1 PEX 8508 Port Combinations

Table 4-1 defines the PEX 8508 port and lane configuration. The equivalent system model contains one upstream port PCI-to-PCI bridge and four downstream port PCI-to-PCI bridges, as illustrated in Figure 4-2. Ports that are not configured or enabled are invisible to software.

The upstream port and downstream ports' lane widths are initially set by the Strapping balls, which are tied High to VDD33 or Low to VSS (GND). The serial EEPROM option can be used to re-configure the ports, with the options defined in Table 4-1. Serial EEPROM configuration occurs following a Fundamental Reset, and overrides the configuration set by the Strapping balls at that time. (Refer to Section 5.3.3, "Setting Port Configuration Using Serial EEPROM.") This can also be changed through the I^2C interface. The final port width can also be changed by link-width negotiation when the ports are connected to external PCI Express ports. The narrowest port on one end of the link determines the maximum link width. Additionally, if a connection is broken on one of the lanes, the training sequence removes the broken lane and negotiates to a narrower width. A x4 port can negotiate down to x2 or x1.

If the port cannot train to x1 (Lane 0 is broken), the port reverses its lanes and attempts to train again. *For example*, a x4 port that cannot train to x4 attempts to negotiate down to x2 or x1; if x1 linkup fails, the port reverses its lanes and attempts again to negotiate linkup. Either the lowest lane (Lane 0) or highest lane (if lanes are reversed) of the programmed link width must connect to the other device's Lane 0.

The PEX 8508 does not bond out SerDes Lanes in numeric order, from 0 through 7. Instead, the actual SerDes Lane number is bonded out by the device to the ports defined in Table 4-1. Lane numbers not listed do not physically exist in the device.

Table 4-1. PEX 8508 Port Configurations

Port Configuration Register Value[a] (Port 0, Offset 224h[4:0])	Lane Width [Lanes/SerDes]/Port[b,d]				
	Port 0	Port 1	Port 2	Port 3	Port 4
0h	x2 [0-1]	x2 [2-3]	x2 [4, 6]	x2 [8, 10]	–[c]
2h	x4 [0-3]	x4 [4, 6, 8, 10]	–	–	–
3h	x4 [0-3]	x2 [4, 6]	x2 [8, 10]	–	–
4h	x4 [0-3]	x2 [4, 6]	x1 [8]	x1 [10]	–
5h	x4 [0-3]	x1 [4]	x1 [8]	x2 [8, 10]	–
6h	x4 [0-3]	x1 [4]	x2 [6, 8]	x1 [10]	–
8h	x4 [0-3]	x1 [4]	x1 [6]	x1 [8]	x1 [10]
9h	x2 [0-1]	x2 [2-3]	x2 [4, 6]	x1 [8]	x1 [10]

a. All other configurations default to option 0.

b. The lanes are assigned to each enabled port in sequence, as indicated in [brackets].

c. "–" indicates that the port is not enabled for that configuration.

d. Refer to Table 4-2 for an explanation of the physical lane to SerDes quad to SerDes module relationship.

Table 4-2. Physical Lane to SerDes Quad to SerDes Module Relationship

Physical Lanes	SerDes Quad	SerDes Modules
0-3	0	0-3
4, 6	1	4, 6
8, 10	2	8, 10

Figure 4-2. Equivalent System Model with Maximum Number of Ports and Lanes

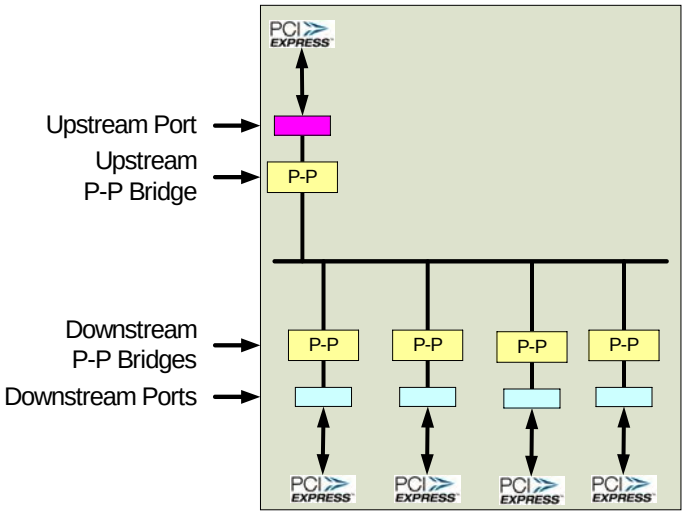

4.1.2.2 Port Numbering

The PEX 8508 Port Numbers are 0, 1, 2, 3, and 4. (Refer to Table 4-1 and Figure 4-3.) The Port Number and Device Number are the same (in the Type 1 Headers) that map to all ports, with an exception for a non-zero upstream port.

All downstream Device Numbers match their corresponding Port Number. *For example*, if Port 0 is the upstream port, Ports 1, 2, 3, and 4 are the downstream ports. The Device Numbers for the PCI-to-PCI bridges implemented on the downstream ports are 1, 2, 3, and 4, respectively. (Refer to Figure 4-3.)

Any PEX 8508 port can be configured as the upstream port. The PCI-to-PCI bridge implemented on the upstream port does not assume a Device Number – it accepts the Device Number assigned by the upstream device. Generally, the upstream device assigns Device Number 0, according to the *PCI Express Base r1.1*.

Figure 4-3. PLX Port Numbering Convention Example (When Port 0 Is Upstream Port)

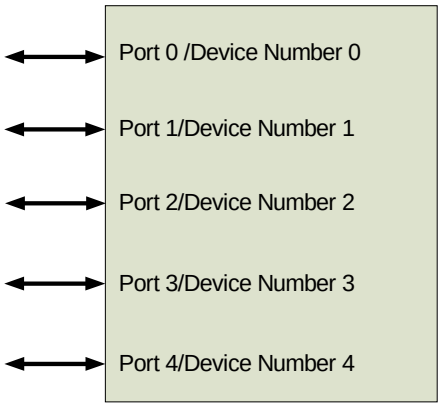

4.2 PCI-Compatible Software Model

The PEX 8508 can be thought of as a hierarchy of PCI-to-PCI bridges, with one upstream PCI-to-PCI bridge and one or more downstream PCI-to-PCI bridges connected by a virtual internal bus. (Refer to Figure 4-4.) PCI-to-PCI bridges are compliant with the PCI and PCI Express system models. Figure 4-4 illustrates the concept of hierarchical PCI-to-PCI bridges, with the bus in the middle being the virtual PCI Bus. The Configuration Space registers (CSRs) in the upstream PCI-to-PCI bridge are accessible by Type 0 Configuration requests targeting the upstream bus interface. The upstream port captures the Type 0 Configuration Write Target Bus Number and Device Number. The upstream port uses this Captured Bus Number and Captured Device Number as part of the Requester ID and Completer ID for the requests and completions generated by the upstream port.

The CSRs in the downstream port PCI-to-PCI bridges are accessible by Type 1 Configuration requests received at the upstream port that target the virtual internal bus, by having a Bus Number value that matches the upstream bridge's Secondary Bus Number value. Each downstream bridge is associated with a unique Device Number, as explained in Section 4.1.2.

The CSRs of downstream devices are hit in two ways. If the Configuration request matches the PEX 8508 downstream port Secondary Bus Number, the PEX 8508 converts the Type 1 Configuration request into a Type 0 Configuration request. However, if the Bus Number does ***not*** match the Secondary Bus Number, but falls within the Subordinate Bus Number range, the Type 1 Configuration request is forwarded out of the PEX 8508, unchanged.

After all PCI devices have been located and assigned Bus and Device Numbers, software can assign a Memory map and I/O map. Requests (Memory or I/O) go downstream if they fall within a bridge's Base and Limit range. In the PEX 8508, each downstream bridge has its own Base and Limit. The Request (Memory or I/O) goes upstream if it does not target anything within the upstream bridge's Base and Limit range.

Completions route by the Bus Number established in the Configuration registers. If the Bus Number is in the Secondary or Subordinate range, the packet goes downstream; otherwise, the packet goes upstream.

Figure 4-4. PEX 8508 System Configuration Propagation

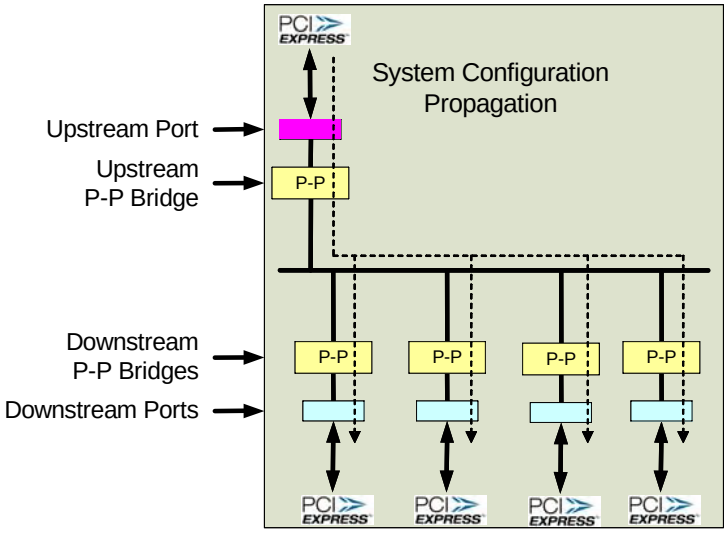

4.2.1 PEX 8508 Reset

The PEX 8508 can be reset by four different mechanisms (refer to Section 5.1, "Reset Overview," for details):

- Fundamental Reset input, through the PEX_PERST# signal
- In-band Reset propagates from upstream, through the Physical Layer mechanism, which communicates a reset through a training sequence (TS1/TS2 Ordered-Set *Hot Reset* or *Disable Link* bit is set)
- PCI Express link enters the DL_Down state on the upstream port
- Upstream port **Bridge Control** register *Secondary Bus Reset* bit is set (offset 3Ch[22]=1)

Reset is propagated from upstream to downstream. Reset is propagated to the downstream PCI Express device, through the PCI Express link by the Physical Layer mechanism (the TS1/TS2 Ordered-Set *Hot Reset* bit is set), or when the upstream port link enters the *DL_Down* state. (Refer to Section 5.1, "Reset Overview," for further details.)

Figure 4-5 illustrates an example of system reset propagation.

Figure 4-5. PEX 8508 System Reset Propagation

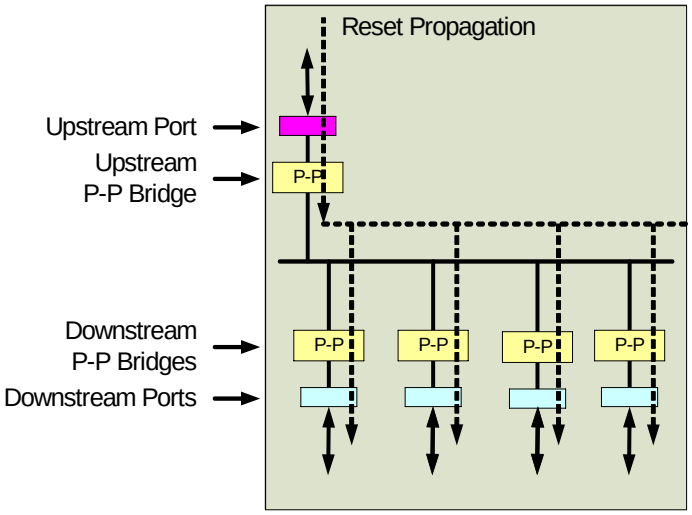

4.2.2 Interrupts

Generated interrupts are INT*x* Interrupt message-type (compatible with the *PCI r3.0*-defined Interrupt signals) or Message Signaled Interrupts (MSI), when enabled. MSI and INT*x* are mutually exclusive; either can be enabled in a system (depending upon which interrupt type the system software supports), but never both. [Refer to the **Message Signaled Interrupt Capability** register (offset 48h) and **Command** register *Interrupt Disable* bit (offset 04h[10]).] The PEX 8508 does not convert received INT*x* messages to MSI messages.

Refer to Chapter 6, "Interrupts," for complete details.

4.2.2.1 Interrupt Sources or Events

The PEX 8508-generated interrupt/message sources include:

- Hot Plug events
- Device-specific error events
- INT*x*

4.2.2.2 INT*x* Switch Mapping

The PEX 8508 remaps and collapses the INT*x* *virtual wires*, based upon the downstream port Device Number and received INT*x* message Requester ID Device Number. Each virtual PCI-to-PCI bridge of a downstream port specifies the Port Number associated with the INT*x* (Interrupt) messages received or generated, and forwards the Interrupt messages in the upstream direction. Refer to Section 6.2.1, "INTx-Type Interrupt Message Remapping and Collapsing," for interrupt routing information.

4.3 PCI Express Station Functional Description

The single-station PEX 8508 supports up to two to five ports and 8 lanes, which can be configured as defined in Table 4-1. The station forwards ingress packets to the internal fabric, to be routed to the egress station.

The station implements the PCI Express Physical (PHY) Layer and Data Link Layer (DLL) functions for each of its ports, and aggregates traffic from these ports onto a transaction-based non-blocking internal fabric. The PCI Express station also performs many Transaction Layer functions, while the packet queuing and ordering aspects of this layer are handled by the Crossbar Switch Control blocks.

During system initialization, software initiates Configuration requests that set up the PCI Express interfaces, Device Numbers, and Address maps across the various ports. These maps are used to direct traffic between ports during standard system operation. Traffic flow between the ports is supported through the central internal fabric.

4.3.1 PEX 8508 Functional Blocks

At the top level, the station has a layered organization consisting of the Physical (PHY), Data Link Layer (DLL), and Transaction Layer Control (TLC) blocks, as illustrated in Figure 4-6. The PHY and DLL blocks have port-specific data paths (one per PCI Express port) that operate independently of one another. The TLC ingress aggregates traffic for all ingress ports in the station, then sends the traffic to the internal fabric. The TLC egress accepts packets, by way of the internal fabric, from all ingress ports, and schedules them to be sent out the appropriate egress port.

Figure 4-6. PCI Express Station Block Diagram

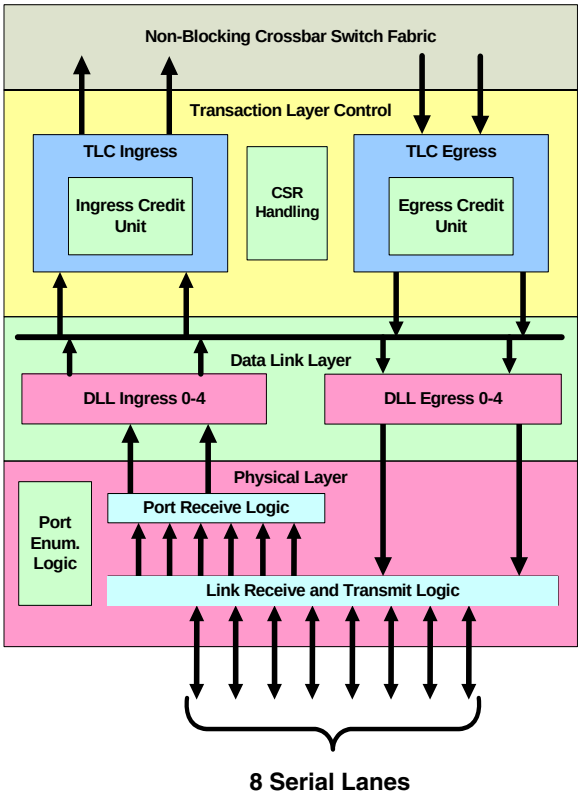

4.3.1.1 Physical Layer

The Physical (PHY) Layer module interfaces to the PCI Express lanes and implements the PHY Layer functions. The number of ports can vary from two to five, with a cumulative lane-bandwidth of x8. When there are fewer than five configured/enabled ports, the cumulative x8 lane-bandwidth is shared among the remaining active ports, as defined in Table 4-1. PHY layer functions include:

- Establishes port configurations and SerDes-to-port assignments
- Establishes internal bandwidth division among ports
- Supports cross-linked upstream and downstream ports
- 8b/10b encoding/decoding
- Data scrambling/unscrambling
- Lane reversal
- Packet framing
- Loop-back master and slave support
- Pseudo-Random Bit Sequence (PRBS) data generation and checking
- Programmable test pattern with SKIP Ordered-Set insertion and return data checking

Refer to Section 9.2, "Physical Layer," for further details.

4.3.1.2 Data Link Layer

The Data Link Layer (DLL) serves as an intermediate stage between the Transaction Layer and the Physical Layer. The primary responsibility of the Data Link Layer includes link management and data integrity, including error and error correction.

The transmission side of the Data Link Layer accepts TLPs assembled by the Transaction Layer, calculates and applies data protection code and TLP Sequence Number, and submits them to the Physical Layer for transmission across the link.

The receiving Data Link Layer is responsible for checking the integrity of received TLPs and submitting them to the Transaction Layer for further processing. On detection of TLP error(s), this layer is responsible for requesting re-transmission of TLPs until the information is correctly received, or the link is determined to have failed.

Refer to Section 9.3, "Data Link Layer (DLL)," for further details.

4.3.1.3 Transaction Layer Control

The Transaction Layer Control (TLC) module accepts Transaction Layer Packets (TLPs), then decodes and processes these TLPs per the PCI Express protocol. The TLC module serves as the intermediate pipeline, by accepting Ingress TLPs from the various ports and staging these TLPs to the source buffers of the Crossbar Switch interface. Similarly, the TLC module accepts TLPs from the Crossbar Switch destination buffers and schedules these TLPs onto Retry buffers in the various DLL Egress ports. In addition, the TLC module performs other TL functions, including:

- Credit-based Flow Control (FC)
- Data Poisoning and end-to-end data integrity detection
- Pipelined full-split transaction protocol
- TLP construction, processing, and routing
- Virtual Channel management
- PCI/PCI-X-compatible ordering
- Power management
- Error logging and reporting
- Interrupt handling
- QoS handling

The PEX 8508 does not support Locked transactions. This is consistent with limitations for Locked transaction use, as outlined in the *PCI r3.0* (Appendix F, "Exclusive Accesses"), and prevents potential deadlock, as well as serious performance degradation, that could occur with Locked transaction use. The PEX 8508 responds to "lock"-type Read Requests (MRdLk) with a Completion, having a Completion status of Unsupported Request (UR).

Refer to Section 9.4, "Transaction Layer (TL)," for further details.

4.3.1.4 Non-Blocking Crossbar Switch Architecture

The Non-Blocking Crossbar switch is an on-chip interconnect switching fabric used to link multiple stations. The Crossbar Switch architecture is built on the existing PLX switch fabric architecture technology. In addition to addressing simultaneous multiple flows, the Crossbar Switch architecture incorporates functions required to support an efficient PCI Express switch fabric, including:

- Deadlock avoidance
- Priority preemption
- Two independent VCs
- PCI Express ordering rules
- Packet fair queuing
- Oldest first scheduling

The Crossbar Switch interconnect physical topology is that of a packet beat-based internal fabric. The Crossbar Switch protocol is sufficiently flexible and robust to support a variety of embedded system requirements. The Crossbar Switch architecture basic features include:

- Multiple concurrent Data transfers
- Global ordering within the switch
- Three types of transactions (Posted, Non-Posted, and Completion) within a VC meet PCI and PCI Express Ordering and Deadlock Avoidance rules
- Internal credit updates guarantee packet forwarding progress, once scheduled
- Packet scheduling in the Crossbar switch selects the oldest eligible packet, with programmable weights for source ports, destination ports, and VC

4.3.2 Relaxed Ordering

The PEX 8508 does not support the optional Relaxed Ordering capability defined in the *PCI Express Base r1.1*. However, is does support two Device-Specific Relaxed Ordering modes:

- PEX 8508 Relaxed Ordering
- PEX 8508 Relaxed Completion Ordering

4.3.2.1 PEX 8508 Relaxed Ordering

The PEX 8508 does ***not support*** the TLP optional *Relaxed Ordering* bit, as specified in the *PCI Express Base r1.1*, Table 2-23. By default, all packets entering from a specific port are dispatched to their respective destinations, based upon strict ordering.

However, to remove unnecessary head-of-line blocking caused by PCI ordering in applications where ordering is not important, the PEX 8508 offers a Device-Specific Relaxed Ordering mode. Device-Specific Relaxed Ordering mode is enabled when any bit within a **Device-Specific Relaxed Ordering Mode** register *Enable PLX Relaxed Ordering* field is set to 1:

- Port 0 – offset BFCh[7:0]
- Port 1 – offset BFCh[15:8]
- Port 2 – offset BFCh[23:16]
- Port 3 – offset BFCh[31:24]
- Port 4 – offset BE4h[7:0]

In general, each port has 8 TCs to which it can map.

The ingress scheduler on a specific port (for a specific Traffic Class) selects packets without using ordering requirements and dispatches the packets to the egress ports. If using the Relaxed Ordering feature, ensure that it is used only for packets of a specific Traffic Class. This allows those packets to be distinguished from packets on other Traffic Classes in which the Relaxed Ordering feature is not enabled.

After the packets reach the egress ports, strict ordering is used in these queues, irrespective of the bits set on the ingress port.

Refer to Section 8.3.2.3, "Device-Specific Relaxed Ordering," for further details.

4.3.2.2 PEX 8508 Relaxed Completion Ordering

The PEX 8508 provides a Relaxed Completion Ordering feature that enables Completion transactions to pass enqueued Posted transactions that are blocked. This feature is available on VC0, and is enabled by setting the following bits in the Device-Specific Configuration space:

- **Device-Specific Relaxed Ordering Enable** register *Device-Specific Relaxed Ordering Enable* bit (Ports 0, 1, 2, 3, and 4, offset 1F0h[20]) is set to 1, and
- Any bit within a **Device-Specific Relaxed Ordering Mode** register *Enable PLX Relaxed Ordering* field is set to 1:
  - Port 0 – offset BFCh[7:0]
  - Port 1 – offset BFCh[15:8]
  - Port 2 – offset BFCh[23:16]
  - Port 3 – offset BFCh[31:24]
  - Port 4 – offset BE4h[7:0]

  In general, each port has 8 TCs to which it can map.

4.3.3 No Snoop Enable

In NT mode, the PEX 8508 provides a No Snoop Disable feature to force the *No Snoop* attribute bit in the packet header to 0, for all packets transferred between the NT Link and Virtual Interfaces (across the NT boundary, in both directions). This capability can be used to handle cache coherency-related issues in a system. To enable this feature, set the **Ingress Control** register *No Snoop Disable* bit (offset 660h[24]) to 1.

4.3.4 Cut-Thru Mode

By default, the upstream port is enabled for Cut-Thru. Cut-Thru mode can reduce latency, especially for longer packets, because the entire packet does not need to be stored before being forwarded. Instead, after the header is decoded, the packet can be immediately forwarded. Cut-Thru mode can be disabled by clearing the **Debug Control** register *Cut-Thru Enable* bit (1DCh[21]).

The Cut-Thru port can be moved from the upstream port to any other port, by programming the **Debug Control** register *Cut-Thru Port Number* field (1DCh[6:4]).

Some TLPs will never use a Cut-Thru path. These include Configuration TLPs, I/O TLPs, and any TLP when the **Bridge Control** register *VGA Enable* bit (offset 3Ch[19]) is set for that port.

Cut-Thru mode, if enabled, is supported for the PEX 8508's NT Port Link Interface, if the PEX 8508 is configured for Non-Transparent mode.

Note: *The **Debug Control** register Cut-Thru Enable bit (1DCh[21]) affects the entire switch. If Cut-Thru is enabled (default), all ports use Cut-Thru. If Cut-Thru is not enabled, no ports use Cut-Thru.*

Caution: ***One of the drawbacks to using Cut-Thru mode is that the TLP is not known to be good until the last byte. If the TLP proves to be bad, the Cut-Thru packet must be discarded. If the TLP has already been forwarded to another device, that TLP will be framed with an EDB (End Data Bad), as opposed to the standard END.***

4.4 Non-Transparent Bridging Implementation

The PEX 8508 supports Non-Transparent bridge functionality, which is used to implement High-Availability systems or Intelligent I/O modules using PCI Express technology. The following discusses the basic Non-Transparent bridging concept, as it applies to a PCI Express system.

Non-Transparent (NT) bridges allow systems to isolate address spaces by appearing as an endpoint to the Host. The NT bridge exposes a Type 0 CSR header and forwards transactions from one domain to the other, using address translation. The NT bridge is used to connect two independent address/Host domains. The NT bridge includes **Doorbell** registers for transmitting interrupts from one side of the bridge to the other. The bridge also has **Scratchpad** registers, accessible from both domains for inter-host communication. The PEX 8508 PCI Express switch, with a single port configured to operate in NT Bridge mode, supports two system models:

- Intelligent Adapter Mode
- Dual-Host Mode

The PEX 8508 switch initial operating mode is determined by the STRAP_MODE_SEL[1:0] balls. The mode can later be changed by the configuration serial EEPROM, by writing to the **Debug Control** register *Mode Select* bits.

Refer to Chapter 14, "Non-Transparent Bridging (NTB)," for further details.

4.4.1 Intelligent Adapter Mode

The use of Non-Transparent bridges in PCI systems is well established for supporting intelligent adapters in enterprise and multi-Host systems. This same concept is used in PCI Express bridges and switches.

In Figure 4-7, there are two Type 0 CSR headers in the Non-Transparent PCI-to-PCI bridge. The one nearer the internal virtual PCI Bus is referred to as the *Virtual Interface*. The one nearer the PCI Express link is referred to as the *Link Interface*.

In Intelligent Adapter mode, the NT Port Link Interface is connected to the System domain. The System Host manages only the NT Port Link Interface Type 0 (PCI endpoint) function. The Local Host manages all PEX 8508 Transparent Port Type 1 (PCI-to-PCI bridge) and NT Port Virtual Interface Type 0 functions. Cross-domain traffic is routed through an Address Translation mechanism. (Refer to Section 14.1.6, "Address Translation.")

Intelligent Adapter mode does ***not support*** Host-failover applications.

Note: *P-P is used to represent PCI-to-PCI in the illustrations provided in this data book.*

Figure 4-7. PEX 8508 Intelligent Adapter Software Model

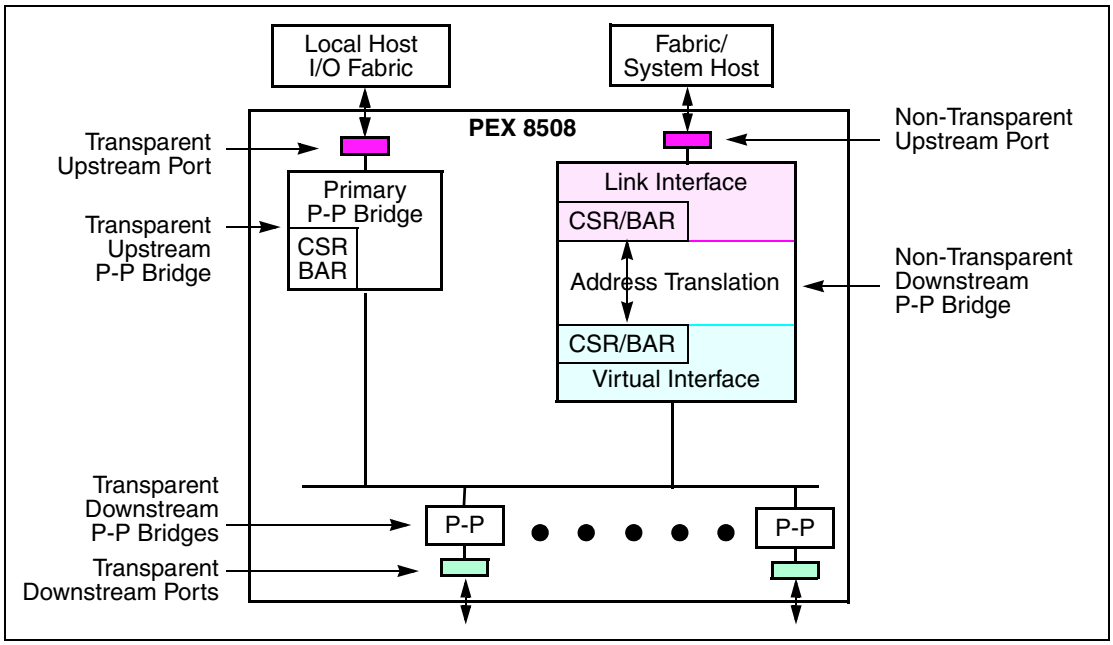

4.4.2 Dual-Host Mode

The PEX 8508 can concurrently support two hosts. The PEX 8508 Transparent upstream port is connected to the Active Host, and the NT Port Link Interface is connected to the Backup Host. (Refer to Figure 4-8.)

After power-up, the Active and Backup Hosts can enumerate their domain at the same time, which is modified using the optional serial EEPROM. The serial EEPROM is necessary to enable or resize the **BAR Setup/Limit** registers.

Dual-Host mode supports Host-Failover applications, described in the next section.

Figure 4-8. Dual-Host Model

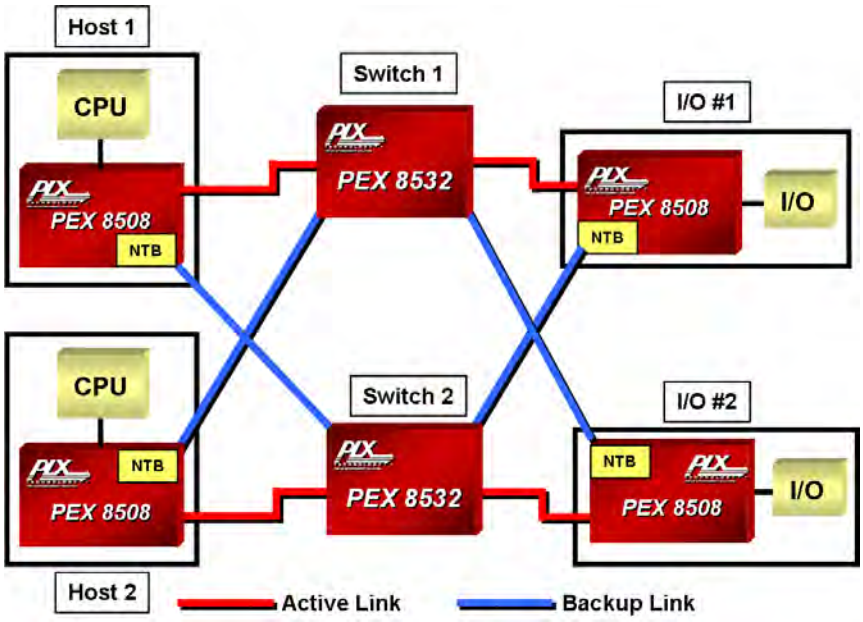

4.4.2.1 Host-Failover Application

The Host-Failover application is based upon the basic dual-Host configuration. The Active Host periodically transmits heartbeat messages, by way of the switch, to the Backup Host to indicate that it remains active. When the Backup Host fails to receive heartbeat messages before its Fail-Detect Timer expires, it starts the failover process. The Backup Host halts cross-domain traffic before it starts the failover. The Backup Host uses the Memory-Mapped access to the Configuration register to execute the failover. The Backup Host performs the following procedure to take control:

1. Failover detected.
2. Upstream port demotion.
3. NT Port (self-) promotion as a new upstream port.

Refer to Section 16.2.4, "Host-Failover Application," for further details.

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 5 Reset and Initialization

5.1 Reset Overview

Reset is a mechanism that returns a device to its initial state. Hardware or software mechanisms can trigger a reset. The re-initialized states following a reset vary, depending upon the reset type and condition.

The *PCI Express Base r1.1*, Section 6.6, defines the hardware mechanism as *Fundamental Reset*. Two actions can trigger a Fundamental Reset:

- Cold Reset
- Warm Reset

There is also a type of reset triggered by an in-band signal from an upstream PCI Express link to all its downstream ports, which is called a Hot Reset.

There is also a Secondary Bus Reset. Any PCI-to-PCI bridge can reset its downstream hierarchy by setting the **Bridge Control** register *Secondary Bus Reset* bit (offset 3Ch[22]).

Upon exit from a Cold or Warm Reset, all port configurations, port registers, and state machines are set to initial (start-up) values, as specified in Section 5.2, "Initialization Procedure."

5.1.1 Cold Reset

A Cold Reset is a Fundamental Reset that occurs following a proper PEX 8508 power-on. When the PEX_PERST# signal is held Low following the proper application of power to the device, a Fundamental Reset occurs.

A Fundamental Reset initializes the entire PEX 8508 device (*such as*, configuration information, clocks, state machines, registers, and so forth).

When power is removed from the device, or travels outside the required ranges, all settings and configuration information is lost. The device must cycle through the entire Initialization Procedure after power is accurately re-applied.

5.1.2 Warm Reset

The Fundamental Reset mechanism can also be triggered by driving the PEX 8508 hardware Reset signal (PEX_PERST#) Low, without the removal and re-application (recycling) of power. This is considered a *Warm Reset*.

PEX_PERST# can be controlled by on-board toggle switches or other external hardware resets to the device. The PEX 8508 must cycle through the entire Initialization Procedure after the PEX_PERST# Input signal is returned to High.

5.1.3 Hot Reset

A Hot Reset is equivalent to a traditional Software Reset. Triggered by an in-band signal from an upstream PCI Express link to all downstream ports, a Hot Reset causes all ports that are downstream from the initiating port to set their registers and state machines to initial values. This type of reset does not require power cycling, nor does it cause PEX 8508 port re-configuration. However, a Hot Reset:

- Causes all TLPs held in the PEX 8508 to be dropped
- Returns all State machines to their initial (default) values
- Returns all Non-Sticky register bits to their initial (default) conditions (refer to Table 13-3, "Register Types, Grouped by User Accessibility," for further details regarding Sticky register bit types)

A Hot Reset is triggered by the following actions:

- Physical Layer (at the upstream port) receives a reset through a training sequence leading to a Hot Reset
- Upstream PCI Express port enters the *DL_Inactive* state, which has the same effect as a Hot Reset

Note: *In the following sections, the terms "virtual PCI-to-PCI bridge" and "port" refer to a given Station port.*

5.1.3.1 Hot Reset Propagation

A Hot Reset is propagated to a downstream PCI Express device through the PCI Express link, using the Physical Layer Hot Reset mechanism (*that is*, a *Reset* bit in the Training Ordered-Set from the upstream device is set).

PCI Express views a *switch* as a hierarchy of virtual PCI-to-PCI bridges.

An example of reset propagation across the PEX 8508 switch is illustrated in Figure 5-1. Upon receiving a Hot Reset from the upstream PCI Express link, the virtual primary PCI-to-PCI bridge propagates the reset to virtual secondary PCI-to-PCI bridges in the upstream and downstream ports. Each virtual secondary PCI-to-PCI bridge propagates the reset to its downstream links, and initializes its internal states to initial/default conditions.

A Hot Reset does not impact Clock Logic, Port Configuration, nor Sticky register bits.

Figure 5-1. PEX 8508 System Reset Propagation Example

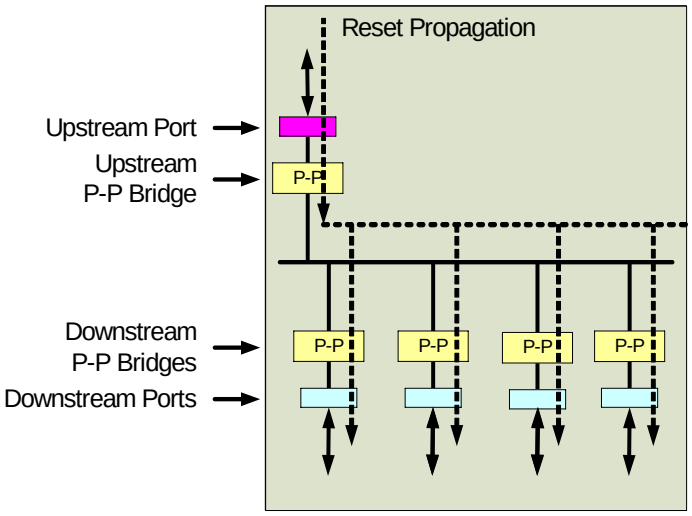

5.1.3.2 Hot Reset Disable

The PEX 8508 includes a configuration option – **Debug Control** register *Upstream Port Hot Reset and Link Down Reset Propagation Disable* bit (Port 0, offset 1DCh[20]) – to ignore the Hot Reset sequence from the upstream PCI Express link. Setting this bit enables the upstream port to ignore a Hot Reset training sequence, blocks the switch from manifesting an internal reset due a DL_Down event, and prevents the downstream ports from issuing a Hot Reset to downstream devices when either a Hot Reset or DL_Down event occurs on the upstream link.

5.1.4 Secondary Bus Reset

When the **upstream** PCI-to-PCI bridge **Bridge Control** register (**BCR**) *Secondary Bus Reset* bit (offset 3Ch[22]) is set to 1, all ports that are downstream from that port are reset to their initial/default states. The downstream ports propagate an in-band Hot Reset to their respective downstream links. In addition, the downstream ports' Configuration Space registers (CSRs) are re-initialized. The upstream PCI-to-PCI bridge (upstream port) and its CSRs are not affected; however, the queues to/from all downstream ports are drained, because their upstream-to-downstream virtual connections are re-initialized.

When the **downstream** PCI-to-PCI bridge **BCR** *Secondary Bus Reset* bit is set to 1, a Hot Reset is transmitted to its single downstream port, which resets all devices downstream from that port to their initial/default states. The reset port drops any incoming traffic. All other PEX 8508 traffic not flowing to the reset port is unaffected.

The downstream links are held in reset until software removes the condition by clearing the **BCR** *Secondary Bus Reset* bit. The PHY layer of the downstream port in question propagates the reset condition in-band to its downstream link, and remains in the Hot Reset state until the Reset condition (BCR) is cleared. The Transaction Layer draining of non-empty queues to/from the affected port(s) is handled in a manner similar to the case of that port proceeding to the *DL_Inactive* state, as defined in the *PCI Express Base r1.1*, Section 2.9.

5.2 Initialization Procedure

Upon exit from a Fundamental Reset, the PEX 8508 initialization process is started. There are two to four steps in the process, depending upon the availability of an external initialization serial EEPROM and I^2C.

The initialization sequence executed is as follows:

1. PEX 8508 reads the Strapping signal balls to determine System mode, upstream port, and lane configuration of all ports.
2. If a serial EEPROM is present (EE_PR# ball is Low), the serial EEPROM data is downloaded to the PEX 8508 Configuration registers. The configuration from the Strapping signal balls can be changed by serial EEPROM data.
3. If I^2C is used, it can be used to program all the registers (the same as would be done with the serial EEPROM).
4. After the configuration from the Strapping signal balls, serial EEPROM, and/or I^2C completes, the Physical Layer of the configured ports attempts to bring up the links. After both components on a link enter the initial Link Training state, the components proceed through Physical Layer Link initialization and then through Flow Control initialization for VC0, preparing the Data Link and Transaction Layers to use the link. Following Flow Control initialization for VC0, it is possible for VC0 Transaction Layer Packets (TLPs) and Data Link Layer Packets (DLLPs) to be transmitted across the link.

5.2.1 Default Port Configuration

The default upstream port selection and overall port lane-width configuration is determined by Strapping signal ball levels. All Strapping balls must be tied High to VDD33 or Low to VSS (GND), which sets the default device configuration. (Refer to Section 3.4.4, "Strapping Signals.") These settings can be changed by downloading serial EEPROM data or by I^2C programming.

5.2.2 Default Register Initialization

Each PEX 8508 port defined in the Port Configuration process has a set of assigned registers that control port activities and status during standard operation. These registers are set to default/initial settings, as defined in:

- Chapter 13, "Transparent Port Registers"
- Chapter 17, "NT Port Virtual Interface Registers"
- Chapter 18, "NT Port Link Interface Registers"

Following a Fundamental Reset, the basic PCI Express Support registers are initially set to the values specified in the *PCI Express Base r1.1*. The Device-Specific registers are set to the values specified in their register description tables. These registers can be changed by loading new data with the attached serial EEPROM, the I^2C interface, or by way of Transaction Layer Configuration Space register (CSR) accesses using Configuration or Memory Writes; however, registers identified as Read-Only (RO) ***cannot*** be modified by Configuration nor Memory Write requests.

The PEX 8508 supports three mechanisms for accessing registers by way of the Transaction Layer, as described in the following sections:

The Transparent Ports and NT Port support the following mechanisms for accessing registers by way of the Transaction Layer, as described in the following sections:

- Section 13.4.1, "PCI r3.0-Compatible Configuration Mechanism"
- Section 13.4.2, "PCI Express Enhanced Configuration Mechanism"
- Section 13.4.3, "Device-Specific Memory-Mapped Configuration Mechanism"

The NT Port also supports:

- Section 17.2.3 or Section 18.2.3, "Device-Specific I/O-Mapped Configuration Mechanism"
- Section 17.2.4 or Section 18.2.4, "Device-Specific Cursor Mechanism"

5.2.3 Serial EEPROM Initialization

Note: *For further details, refer to the PEX 85XX EEPROM – PEX 8518/8517/8508 Design Note.*

The on-chip Serial EEPROM Controller is integrated in the PEX 8508, as illustrated in Figure 5-2. The Controller performs a serial EEPROM download when a serial EEPROM is present, as indicated by the EE_PR# Strapping ball = Low, and when one of the following events occur:

- PEX_PERST# is returned High, following a Fundamental Reset (*such as,* a Cold or Warm Reset)
- Hot Reset is received at the upstream port (downloading upon this event can be optionally disabled, by setting the Port 0 register at offset 1DCh[16, 17, and/or 20])
- Upstream port exits a DL_Down state (downloading upon this event can be optionally disabled, by setting the Port 0 register at offset 1DCh[16, 17, and/or 20])

Figure 5-2. PEX 8508 Serial EEPROM Connections

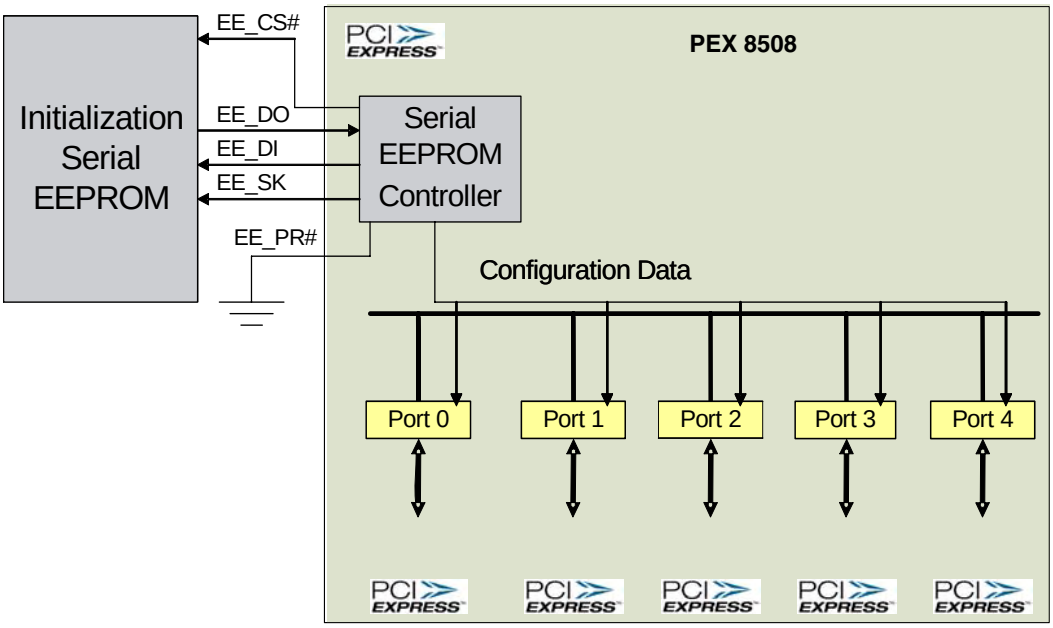

5.2.3.1 Configuration Data Download

Serial data is downloaded from the serial EEPROM and converted to parallel data, which is then routed to the PEX 8508 ports. The Serial EEPROM Controller generates a 7.8-MHz EE_SK signal and reads all DWords from the serial EEPROM, which represents the data necessary to initialize the PEX 8508 registers.

The serial EEPROM Memory map reflects the basic PEX 8508 register map, for registers discussed in the following chapters:

- Chapter 13, "Transparent Port Registers"
- Chapter 17, "NT Port Virtual Interface Registers"
- Chapter 18, "NT Port Link Interface Registers"

Because these registers are duplicated for each port, the serial EEPROM data starts with Port 0, offset 00h, and progresses through the registers for all five ports, in sequence. Certain general device registers are appended at the end of the space for Port 0, and various Address spaces are skipped due to ***unused/reserved*** register space. A detailed description of the serial EEPROM Memory map is provided in Appendix A, "Serial EEPROM Memory Map."

While downloading the data, the PEX 8508 generates a CRC value from the Read data. When the serial EEPROM download is complete, the generated CRC value is compared to a CRC value stored in the last DWord location of the serial EEPROM. If the CRC values match, the PEX 8508 sets the **Serial EEPROM Status** register *EepPrsnt[1:0]* field (Port 0, offset 260h[17:16]) to 01b (serial EEPROM download complete and serial EEPROM CRC is validated).

If the CRC fails:

- The **Serial EEPROM Status** register *EepPrsnt[1:0]* field is set to 11b to indicate that the serial EEPROM is present but a CRC error was detected, and,
- If the **Serial EEPROM Status** register *CRC Disable* bit (Port 0, offset 260h[21]) value is 0 (default), all registers that are serial EEPROM-programmable are reset/initialized to default values, or,
- If the **Serial EEPROM Status** register *CRC Disable* bit value is 1, the CRC error is ignored and all registers that are serial EEPROM-programmable are initialized with the values that were downloaded from the serial EEPROM.

Caution: ***Setting the CRC Disable bit is strongly discouraged because a corrupted serial EEPROM download could render the PEX 8508 inoperable until a Fundamental Reset is applied (by asserting PEX_PERST# input).***

During the serial EEPROM download, the Class Code 060400h is monitored. If the correct code is not found, the download is terminated.

Note: *It is the system software's responsibility to verify whether the serial EEPROM download completes without error.*

5.3 PCI Express Configuration, Control, and Status Registers

The PCI Express Configuration, Control, and Status registers that can be initialized are discussed in the following sections. However, this is not a complete list of programmable registers. For a complete list, refer to Chapter 13, "Transparent Port Registers," and Appendix A, "Serial EEPROM Memory Map."

- Section 5.3.1, "Selecting Configuration Values Using Serial EEPROM"
- Section 5.3.2, "Selecting Upstream Port Using Serial EEPROM"
- Section 5.3.3, "Setting Port Configuration Using Serial EEPROM"
- Section 5.3.4, "Power Management Parameters Using Serial EEPROM"
- Section 5.3.5, "Message Signaled Interrupt Capability Using Serial EEPROM"
- Section 5.3.6, "PCI Express Capability Using Serial EEPROM"
- Section 5.3.7, "Device Serial Number Extended Capability Using Serial EEPROM"
- Section 5.3.8, "Power Budgeting Extended Capability Using Serial EEPROM"
- Section 5.3.9, "Virtual Channel Extended Capability Using Serial EEPROM"
- Section 5.3.10, "Advanced Error Reporting Capability"

5.3.1 Selecting Configuration Values Using Serial EEPROM

Configuration values, for the registers defined in Section 13.6, "PCI-Compatible Type 1 Configuration Registers," can be changed using the serial EEPROM, with the exception of those fixed by the *PCI Express Base r1.1* or ***not supported*** by the PEX 8508.

5.3.2 Selecting Upstream Port Using Serial EEPROM

The **Debug Control** register *Upstream Port ID* field (offset 1DCh[10:8]) values defined in Table 5-1 determine the upstream port selection.

Table 5-1. Debug Control Register *Upstream Port ID* field (Offset 1DCh[10:8]) Values

Ports	
0h = Port 0	3h = Port 3
1h = Port 1	4h = Port 4
2h = Port 2	5h to 7h = ***Reserved***

5.3.3 Setting Port Configuration Using Serial EEPROM

To use the serial EEPROM to set the port configuration, program the following registers:

- **Port Configuration** register *Port Configuration* field – Sets the port configuration (Port 0, offset 224h[4:0])
- **Debug Control** register – Sets the Port Number and parameters for the upstream port

The **Port Configuration** register *Port Configuration* field value determines the number of enabled ports, as well as the maximum link widths of those ports, as defined in Table 5-2.

Table 5-2. PEX 8508 Port Configurations

Port Configuration Register Value[a] (Port 0, Offset 224h[4:0])	Lane Width [Lanes/SerDes]/Port[b,d]				
	Port 0	Port 1	Port 2	Port 3	Port 4
0h	x2 [0-1]	x2 [2-3]	x2 [4, 6]	x2 [8, 10]	–[c]
2h	x4 [0-3]	x4 [4, 6, 8, 10]	–	–	–
3h	x4 [0-3]	x2 [4, 6]	x2 [8, 10]	–	–
4h	x4 [0-3]	x2 [4, 6]	x1 [8]	x1 [10]	–
5h	x4 [0-3]	x1 [4]	x1 [8]	x2 [8, 10]	–
6h	x4 [0-3]	x1 [4]	x2 [6, 8]	x1 [10]	–
8h	x4 [0-3]	x1 [4]	x1 [6]	x1 [8]	x1 [10]
9h	x2 [0-1]	x2 [2-3]	x2 [4, 6]	x1 [8]	x1 [10]

a. All other configurations default to option 0.

b. The lanes are assigned to each enabled port in sequence, as indicated in [brackets].

c. "–" indicates that the port is not enabled for that configuration.

d. Refer to Table 5-3 for an explanation of the physical lane to SerDes quad to SerDes module relationship.

Table 5-3. Physical Lane to SerDes Quad to SerDes Module Relationship

Physical Lanes	SerDes Quad	SerDes Modules
0-3	0	0-3
4, 6	1	4, 6
8, 10	2	8, 10

5.3.4 Power Management Parameters Using Serial EEPROM

Power Management parameters, for the registers defined in Section 13.7, "Power Management Capability Registers," can be changed using the serial EEPROM, with the exception of those fixed by the *PCI Express Base r1.1* or ***not supported*** by the PEX 8508.

5.3.5 Message Signaled Interrupt Capability Using Serial EEPROM

Message Signaled Interrupt (MSI) Capability parameters, for the registers defined in Section 13.8, "Message Signaled Interrupt Capability Registers," can be changed using the serial EEPROM, with the exception of those fixed by the *PCI Express Base r1.1* or ***not supported*** by the PEX 8508.

5.3.6 PCI Express Capability Using Serial EEPROM

PCI Express Capability parameters, for the registers defined in Section 13.9, "PCI Express Capability Registers," can be changed using the serial EEPROM, with the exception of those fixed by the *PCI Express Base r1.1* or ***not supported*** by the PEX 8508.

5.3.6.1 Configuring Hot Plug Controller Slot Power-Up Sequence Features with Serial EEPROM

Refer to Section 11.4.1.1, "Configuring Hot Plug Controller Slot Power-Up Sequence Features with Serial EEPROM," for details.

5.3.7 Device Serial Number Extended Capability Using Serial EEPROM

Device Serial Number Extended Capability parameters, for the registers defined in Section 13.10, "Device Serial Extended Number Capability Registers," can be changed using the serial EEPROM.

5.3.8 Power Budgeting Extended Capability Using Serial EEPROM

Power Budgeting Extended Capability parameters, for the registers defined in Section 13.11, "Power Budgeting Extended Capability Registers," can be changed using the serial EEPROM.

5.3.9 Virtual Channel Extended Capability Using Serial EEPROM

Virtual Channel Extended Capability parameters, for the registers defined in Section 13.12, "Virtual Channel Extended Capability Registers," can be changed using the serial EEPROM.

5.3.10 Advanced Error Reporting Capability

The PEX 8508 supports Advanced Error Reporting, as defined in the *PCI Express Base r1.1*. Advanced Error Reporting Capability parameters, for the registers defined in Section 13.14, "Advanced Error Reporting Capability Registers," can be changed using the serial EEPROM.

5.3.11 Device-Specific Registers

The following registers are unique to the PEX 8508 device, and not referenced in PCI Express documentation. The Device-Specific registers are organized into the following sections:

- Section 13.13.1, "Error Checking and Debug Registers"
- Section 13.13.2, "Physical Layer Registers"
- Section 13.13.3, "I2C Interface Register"
- Section 13.13.4, "Bus Number CAM Registers"
- Section 13.13.5, "I/O CAM Registers"
- Section 13.13.6, "Address-Mapping CAM (AMCAM) Registers"
- Section 13.13.7, "Ingress Control and Port Enable Registers"
- Section 13.13.8, "I/O CAM Base and Limit Upper 16-Bit Registers"
- Section 13.13.9, "Base Address Shadow Registers (BARs)"
- Section 13.13.10, "Shadow Virtual Channel (VC) Capability Registers"
- Section 13.13.11, "Shadow Port Virtual Channel (VC) Capability Registers"
- Section 13.13.12, "Ingress Credit Handler (INCH) Registers"
- Section 13.13.13, "Relaxed Ordering and Performance Counter Registers"
- Section 13.13.14, "Internal Credit Handler (ITCH) VC&T Threshold Registers"
- Section 13.13.15, "Port Virtual Channel Queue Status Registers"

The Device-Specific registers cannot be accessed by Configuration requests; however, software can access these registers with Memory requests.

5.3.12 Reset and Clock Initialization Timing

Table 5-4. Reset and Clock Initialization Timing

Symbol	Description	Typical Delay
td1	REF Clock stable to PEX_Reset release time	100 µs
td2	PEX_Reset release to PLL Clock Stable and Reset de-bounce	1.32 ms
td3	Clock and Reset Stable to PLL Lock	125 µs
td4	PLL Lock to BIST Done time, which causes Core Reset release	4.5 ms
td5	Core Reset release to SerDes Resets active delay	10 µs
td6	SerDes Reset active time	60 µs
td7	Serial EEPROM load time	12.6 ms

Figure 5-3. Reset and Clock Initialization Timing

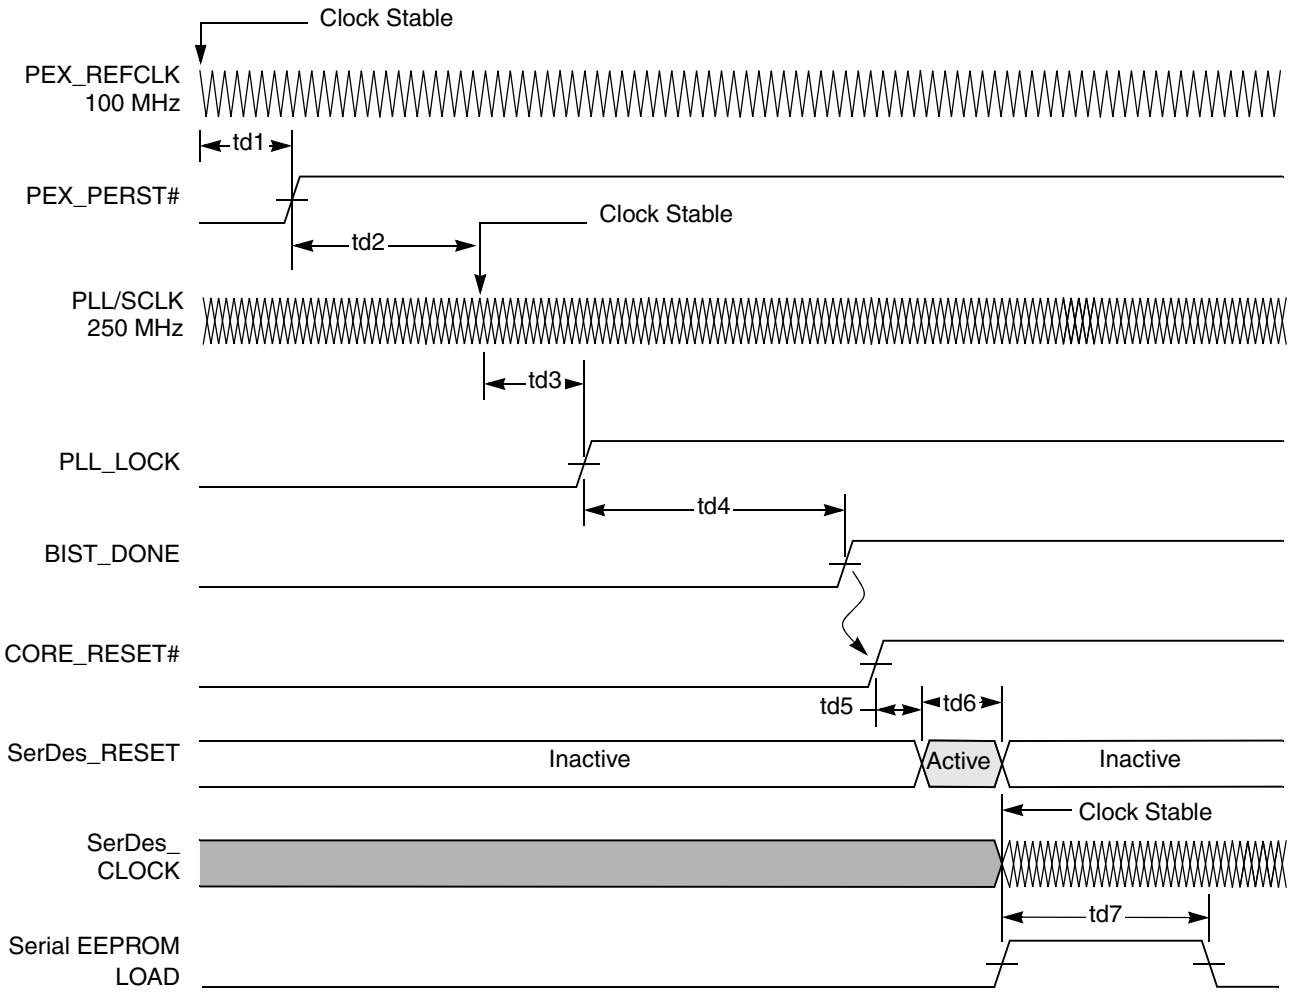

Chapter 6 Interrupts

6.1 Interrupt Support

The PEX 8508 supports the PCI Express interrupt model, which uses two mechanisms:

- INT*x* emulation
- Message Signaled Interrupt (MSI)

For Conventional PCI compatibility, the PCI INT*x* emulation mechanism is used to signal interrupts to the System Interrupt Controller. This mechanism is compatible with existing PCI software, provides the same level of service as the corresponding PCI interrupt signaling mechanism, and is independent of System Interrupt Controller specifics. The PCI INT*x* emulation mechanism virtualizes PCI physical Interrupt signals by using an in-band signaling mechanism.

In addition to PCI INT*x*-compatible interrupt emulation, the PEX 8508 supports the Message Signaled Interrupt (MSI) mechanism. The PCI Express MSI mechanism is compatible with the MSI Capability defined in the *PCI r3.0*.

The following events are supported for interrupts:

- Hot Plug
  - Attention Button Pressed
  - Power Fault Detected
  - MRL Sensor Changed
  - Command Completed
- PCI Express Hot Plug
  - Presence Detect Changed (SerDes Receiver Detect)
  - Data Link Layer State Changed
- Device-specific errors
  - ECC Error detected in internal packet RAM
  - Ingress Credit Underrun
  - Internal Error FIFO Overflow

6.1.1 PEX 8508 Interrupt Handling

The PEX 8508 provides an Interrupt Generation module with each port. The module reads the request for interrupts from different sources and generates an MSI or PCI-compatible Assert_INT*x*/ Deassert_INT*x* Interrupt message. The MSI supports a PCI Express edge-triggered interrupt, whereas Assert_INT*x* and Deassert_INT*x* Message transactions emulate PCI level-triggered interrupt signaling. The System Interrupt Controller functions include:

- Sensing Interrupt events
- Signaling the interrupt, by way of the INT*x* mechanism, and setting the Interrupt Status bit
- Signaling interrupt by way of the MSI mechanism
- Handling INT*x*-type Interrupt messages from downstream devices

6.2 INT*x* Emulation Support

The PEX 8508 supports PCI INT*x* emulation, to signal interrupts to the System Interrupt Controller. This mechanism is compatible with existing PCI software. PCI INT*x* emulation virtualizes PCI physical Interrupt signals, by using the in-band signaling mechanism.

PCI Interrupt registers (the Interrupt registers defined in the *PCI r3.0*) are supported. The *PCI r3.0* **Command** register *Interrupt Disable* and **Status** register *Interrupt Status* bits are also supported (offset 04h[10, 19], respectively).

Although the *PCI Express Base r1.1* provides INTA#, INTB#, INTC#, and INTD# for INT*x* signaling, the PEX 8508 uses only INTA# for internal Interrupt message generation, because it is a single-function device. However, incoming messages from downstream devices can be of INTA#, INTB#, INTC#, or INTD# type. Internally generated INTA# messages from the downstream port are also remapped and collapsed at the upstream port, according to the downstream port's Device Number, with its own Device Number and Received Device Number from the downstream device. When an interrupt is requested, the **Status** register *Interrupt Status* bit is set. If INT*x* interrupts are enabled [**Command** register *Interrupt Disable* and **Message Control** register *MSI Enable* bits (offsets 04h[10] and 48h[16], respectively) are cleared to 0], an Assert_INT*x* message is generated and transmitted upstream to indicate the port interrupt status. For each interrupt event, there is a corresponding Mask bit. The Interrupt request can be generated only when the Mask bit is not set. Software reads and clears the event and *Interrupt Status* bit after servicing the interrupt.

6.2.1 INT*x*-Type Interrupt Message Remapping and Collapsing

INT*x*-type Interrupt messages from downstream devices are directly forwarded to the upstream port, rather than being terminated and regenerated by the downstream port. The upstream port remaps and collapses the INT*x* message type received at the downstream port, based upon the downstream port's Device Number and Received INT*x* message Requester ID Device Number, and generates a new Interrupt message, according to the mapping defined in Table 6-1.

A downstream port transmits an Assert_INTA/Deassert_INTA message to the upstream port, due to a device-specific error.

Internally generated INT*x* messages always originate as type INTA messages, because the PEX 8508 is a single-function device. Internally generated Interrupt INTA messages from downstream ports are remapped at the upstream port to INTA, INTB, INTC, or INTD messages, according to the mapping defined in Table 6-1.

INT*x* messages from downstream devices and from internally generated Interrupt messages are ORed together to generate INTA, INTB, INTC, or INTD level-sensitive signals, and edge-detection circuitry in the upstream port generates the Assert_INT*x* and Deassert_INT*x* messages. The upstream port then forwards the new messages upstream, by way of its link.

Table 6-1. Downstream/Upstream Port INT*x* Interrupt Message Mapping

Device Number	At Downstream Port	By Upstream Port
0, 4	INTA	INTA
	INTB	INTB
	INTC	INTC
	INTD	INTD
1	INTA	INTB
	INTB	INTC
	INTC	INTD
	INTD	INTA
2	INTA	INTC
	INTB	INTD
	INTC	INTA
	INTD	INTB
3	INTA	INTD
	INTB	INTA
	INTC	INTB
	INTD	INTC

6.3 Message Signaled Interrupt (MSI) Support

One of the interrupt schemes supported by the PEX 8508 is the MSI mechanism, which is required for PCI Express devices. The MSI method uses Memory Write transactions to deliver interrupts. MSIs are edge-triggered interrupts.

Note: *MSI and INTx are mutually exclusive. These interrupt mechanisms* ***cannot*** *be simultaneously enabled.*

6.3.1 MSI Operation

At configuration time, system software traverses the function Capability list. If a Capability ID of 05h is found, the function implements MSI. System software reads the MSI Capability Structure registers to determine function capabilities.

The PEX 8508 supports only one message for MSI; therefore, the **Message Control** register *Multiple Message Enable* and *Multiple Message Capable* fields (offset 48h[22:20, 19:17], respectively) are always cleared to 000b.

The **Message Control** register *MSI 64-Bit Address Capable* bit is enabled (offset 48h[23]=1), by default.

System software initializes the MSI Capability Structure registers with a system-specified message. If the MSI function is enabled, after an Interrupt event occurs, the Interrupt Generation module generates a DWord Memory Write to the address specified by the **Message Address[31:0]** register (offset 4Ch) contents. Data written is the contents of the **Message Data** register (offset 54h) lower two bytes and zeros (0) in the upper two bytes. Because the **Message Control** register *Multiple Message Enable* field (offset 48h[22:20]) field is always cleared to 000b, the Interrupt Generation module is not allowed to change the low-order bits of Message data.

When the Hot Plug and/or device-specific error events that caused the interrupt are serviced, the device can generate a new MSI Memory Write as a result of new events. Because MSI is an edge-triggered event, a Mask bit is provided for masking the event [**Mask Bit for MSI** register Mask Bit bit (offset 58h[0])]. A new MSI can only be generated when the *Mask Bit* bit is serviced. System software should mask the *Mask Bit* bit when the MSI event is being processed.

6.3.2 MSI Capability Registers

The MSI Capability registers are described in Section 13.8, "Message Signaled Interrupt Capability Registers."

Chapter 7 Software Architecture

7.1 PEX 8508 Software Model

The PEX 8508 requires software support in the following areas:

- Switch configuration and configuration of all switch downstream links
- Moving data forward and back through the links
- Monitoring and servicing interrupts throughout the connected fabric
- Monitoring and adjusting performance-related mechanisms

The Configuration Mechanisms are straightforward and use Conventional PCI software structures and procedures to set up and identify all ports and links connected through the PEX 8508. The PEX 8508 supports an optional serial EEPROM interface or I^2C interface, to simplify downloading of Configuration data to the switch. (Refer to Section 7.3.)

After the PEX 8508 and its links are set up, data can be routed through the PEX 8508, from one port to another. Responses and other communication are returned, by way of the same links, to the Initiator. The PEX 8508 is transparent to these Data transfers. (Refer to Section 7.3.1.3.)

Errors for which the PEX 8508 is enabled to intercept cause an interrupt to be generated to the Host. It is the responsibility of the Host software to implement Interrupt Service routines to handle the problem. (Refer to Chapter 6, "Interrupts," for further details.)

7.2 Configuration Mechanisms

The PEX 8508 supports the two Configuration mechanisms described in the *PCI Express Base r1.1*:

- **PCI r3.0-Compatible Configuration Mechanism**

  This mechanism supports 100% binary compatibility with the *PCI r3.0* and its corresponding Bus Enumeration and Configuration software. The mechanism allows access to the lower 256 bytes (64 DWords) of the 4-KB Configuration space of each port. Access to the entire 4-KB Configuration Space requires 10 Address bits, which are defined in a PCI Express Configuration Request packet to include a 6-bit *Register Number* field and a 4-bit *Extended Register Number* field. The mechanism maps all six of its Address bits into the *Register Number* field, and clears the *Extended Register Number* field in the packet to 0h. Therefore, the mechanism cannot access the upper 960 DWords (PCI Express Extended Configuration space) that are implemented in each port.

- **PCI Express Enhanced Configuration Mechanism**

  This mechanism increases the size of available Configuration space and optimizes Configuration mechanisms. The mechanism allows access to the entire 4-KB Configuration space defined by the *PCI Express Base r1.1*. Registers within the PEX 8508 that are defined by the *PCI Express Base r1.1* can be accessed by this mechanism. PEX 8508 device-specific registers (which are not defined by the *PCI Express Base r1.1*) cannot be accessed by this mechanism.

The PEX 8508 also supports a third Configuration mechanism:

- **Device-Specific Memory-Mapped Configuration Mechanism**

  This mechanism supports access to all PEX 8508 registers, with the use of Memory Read and Memory Write commands and a 128-KB Address space that includes the 4-KB register sets of all ports, that is located at the Base address assigned to the upstream port's **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively).

From a software point of view, each PEX 8508 port is a PCI-to-PCI bridge. A PCI-to-PCI bridge must have a uniquely assigned Bus Number and Device Number. The upstream port has its own primary Bus Number, while all downstream ports share the same (internal) Bus Number and different Device Numbers.

For further details, refer to the following sections:

- **Transparent mode** – Section 13.4, "Register Access"
- **Non-Transparent mode, Virtual Interface** – Section 17.2, "Register Access"
- **Non-Transparent mode, Link Interface** – Section 18.2, "Register Access"

7.2.1 Software Configuration and Routing

Configuration requests must be routed from the Host, through the PEX 8508's upstream port. All Type 0 Configuration requests to the upstream port access the PEX 8508's upstream port Configuration registers.

The upstream port's PCI-to-PCI bridge forwards Type 1 Configuration requests from its upstream interface to its downstream interface. The Secondary Bus Number of the upstream port matches the Primary Bus Number of every downstream port. If the Bus Number value encoded in the Configuration Request matches the upstream port's Secondary Bus Number, the Configuration request is targeting a downstream port's registers, and therefore the upstream port converts the Type 1 Configuration request to a Type 0 Configuration request. To access the PCI-to-PCI bridge registers of a specific downstream port, the Device Number value encoded in this Configuration Request (which the upstream port converted to Type 0) must match the Device Number (Port Number) of the downstream port (the Device Number of a downstream port is always the same value as the designated Port Number). All other Device Numbers are non-existent devices, and the Configuration request terminates with an Unsupported Request (UR) Completion.

To configure additional devices in the PCI hierarchy, the switch downstream ports must have their Secondary and Subordinate Bus Numbers set. Any match on a downstream port PCI-to-PCI bridge's Secondary Bus Numbers results in the PEX 8508 converting the incoming Type 1 Configuration request to an outgoing Type 0 Configuration request, and the Device Number must be 0.

Figure 7-2 illustrates an example of PEX 8508 system configuration propagation.

Figure 7-1. PEX 8508 System Configuration Propagation

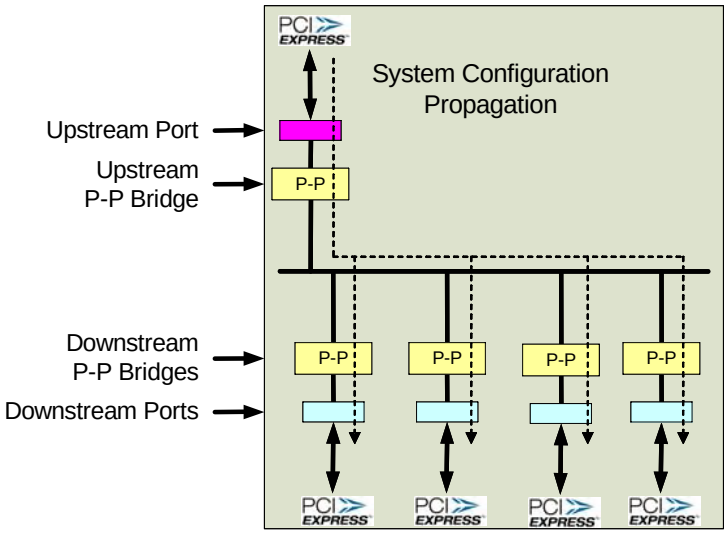

7.3 Sample Configuration Procedure

Consideration must be given to the configuration procedure when setting up and initializing a PEX 8508 switch. Certain items are processed by initial hardware configuration, connections, and operating selections. The PCI/PCI Express Configuration registers can be written from the Host, by way of the upstream port to all downstream ports and their links, or by the serial EEPROM or I^2C interface. Figure 7-2 illustrates an example of PEX 8508 system configuration.

The sequence executed to set up and initialize a PEX 8508 switch is as follows:

1. Select ports and assign lanes to each port:
   - Refer to Section 4.1.2.1, "PEX 8508 Port Combinations," for options
   - PEX 8508 must be connected to PCI Express-compatible devices
   - Strapping balls must be set to identify the selected port configuration – refer to STRAP_PORTCFG[4:0]
   - If enabled, the serial EEPROM can be used to override the Strapping ball selections
2. Select the upstream port – Set Strapping balls STRAP_UPSTRM_PORTSEL[3:0].
3. Software/serial EEPROM programs the following registers for the **upstream port**:
   - **Primary Bus Number** – Identifies the upstream link (**Bus Number** register, offset 18h[7:0])
   - **Secondary Bus Number** – Identifies the switch internal Virtual PCI Bus (**Bus Number** register, offset 18h[15:8])
   - **Subordinate Bus Number** – Must be the last (largest) Bus Number in the downstream hierarchy of the upstream port (**Bus Number** register, offset 18h[23:16])
   - Set the **Command** register *Bus Master Enable* and *Memory Access Enable* bits (offset 04h[2:1], respectively)
   - **Base** and **Limit** registers – Combines the memory of all downstream devices into one large space, with the total size given by Limit - Base, and the Start address given by Base
   - **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively) (Base address for Memory-Mapped CSR access on the PEX 8508)
4. Software/serial EEPROM programs the following registers for the **downstream ports**:
   - **Primary Bus Number** – All downstream Port Numbers are the Device Numbers on the internal virtual PCI Bus (**Bus Number** register, offset 18h[7:0])
   - **Secondary Bus Number** – Identifies the port's downstream link (**Bus Number** register, offset [15:8])
   - **Subordinate Bus Number** – Must be the last (largest) Bus Number in the downstream hierarchy of each downstream port (**Bus Number** register, offset 18h[23:16])
   - Set the **Command** register *Bus Master Enable* and *Memory Access Enable* bits (offset 04h[2:1], respectively)
   - **Base** and **Limit** registers – Combines the memory of all downstream devices into one large space, with the total size given by Limit - Base, and the Start address given by Base

Figure 7-2. PEX 8508 System Configuration Example

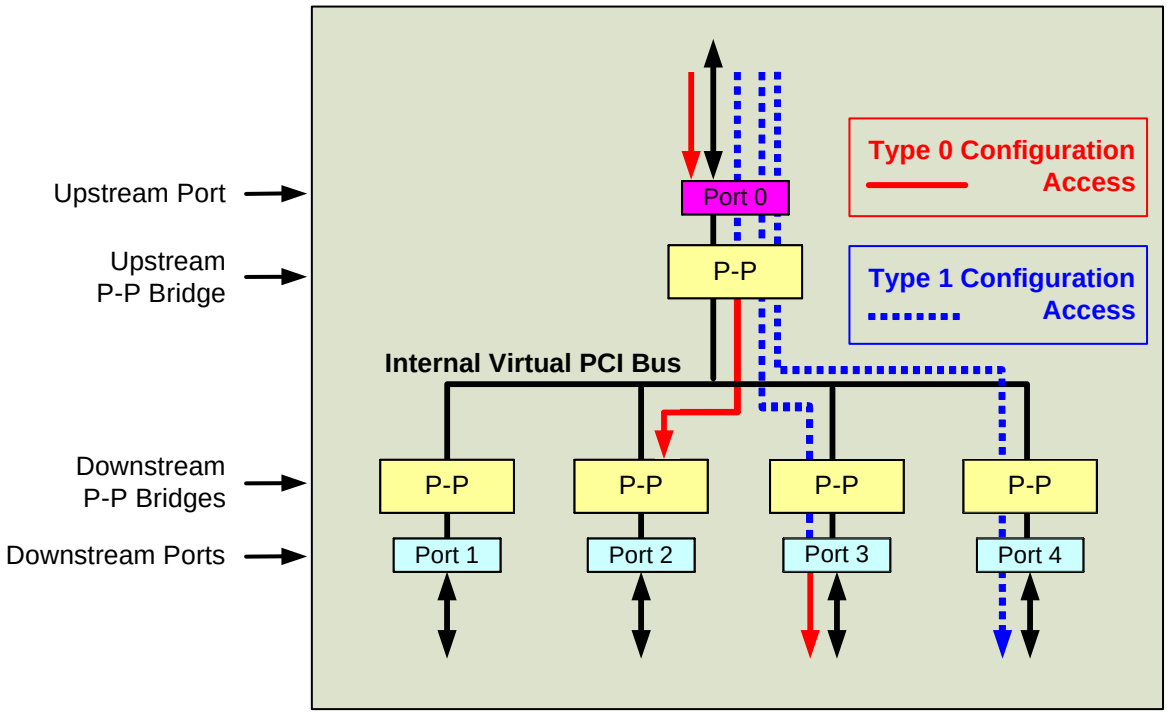

Note: *In Figure 7-2, Port 0 is designated as the upstream port; however, any port can be designated as the upstream port.*

7.3.1 Switch Device Number Assignment Example

The following is an example of how to access the PEX 8508's upstream port and downstream ports. The example assumes that Port 0 is the upstream port. A diagram of the system is illustrated in Figure 7-2.

- Initially, all ports have a Bus Number of 0, a Device Number matching the Port Number, and a Function Number fixed at 0.
- Port 0 registers are accessed with a Type 0 Configuration request, where the Primary Bus Number is copied from the request. For this example, the Primary Bus Number is 7. The Device and Function Numbers must be 0 to hit the registers.
- Port 1 registers are accessed with a Type 1 Configuration transaction to the upstream port with the Bus Number:
  - Bus Number is the internal Virtual PCI Bus (upstream port Secondary Bus Number)
  - Device Number is 01h
  - Function Number to be 000b
- Port 2 registers are accessed with a Type 1 Configuration transaction:
  - Bus Number is the internal Virtual PCI Bus
  - Device Number is 02h
  - Function Number is 000b
- Port 3 registers are accessed with a Type 1 Configuration transaction:
  - Bus Number is the internal Virtual PCI Bus (upstream port Secondary Bus Number)
  - Device Number is 03h
  - Function Number is 000b
- Port 4 registers are accessed with a Type 1 Configuration transaction
  - Bus Number is the internal Virtual PCI Bus (upstream port Secondary Bus Number)
  - Device Number is 04h
  - Function Number is 000b

7.3.1.1 Configuration Register Programming Sequence

Registers that are defined by PCI-SIG Specifications can be accessed by Configuration mechanisms or Memory command; device-specific registers can be accessed by Memory command, but not by Configuration mechanisms (except for limited, indirect access through the NT Port Virtual Interface Cursor Mechanism Control Registers, when Non-Transparent mode is enabled).

The upstream port **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively) map internal registers for Memory-Mapped I/O access. **BAR0** is a 128-KB, Non-Prefetchable BAR [the *Prefetchable* bit (bit 3) with default value 0 is ***not*** programmable]. Because **BAR0** maps to Non-Prefetchable Address space and the Address space is relatively small, it is recommended that **BAR0** be configured as a 32-bit BAR (default, with bits [2:1]=00b), rather than as a 64-bit BAR, to be mapped below the 4-GB Address Boundary space.

With **BAR0** configured as a 32-bit BAR, **BAR1** (which contains the upper 32 bits of address if **BAR0** is configured as a 64-bit BAR) must remain the default value 0h.

BAR0 and **BAR1** can be disabled by setting the **Ingress Control** register *Disable Upstream Port BAR0 and BAR1* bit (Port 0, offset 660h[26]).

Register access must be 1 DWord (Byte Enables can select individual bytes). If a Memory Mapped Read to a Configuration register requests more than 1 DWord, the PEX 8508 returns the first DWord, with a Completion status of *Completer Abort*. This error is flagged in the upstream port **Device Status** register (offset 70h) and **Uncorrectable Error Status** register *Completer Abort Status* bit (offset FB8h[15]).

To program access to internal registers:

1. Program the **Bus Number** register in the upstream port (offset 18h).
2. Program the **Bus Number** registers in all downstream ports (offset 18h).
3. Program the **Memory Base and Limit** register (offset 20h) in all downstream ports.
4. Program **BAR0/1** on the upstream port. (Optional, but necessary for Memory-Mapped access to internal registers.)
5. Program the **Memory Base and Limit** register (offset 20h) in the upstream port, ensuring the values claim all the space requested by all downstream ports.
6. Program the **Command** register *Bus Master Enable* and *Memory Access Enable* bits on all ports (offset 04h[2:1], respectively).

Figure 7-3. Programming Base and Limit Values

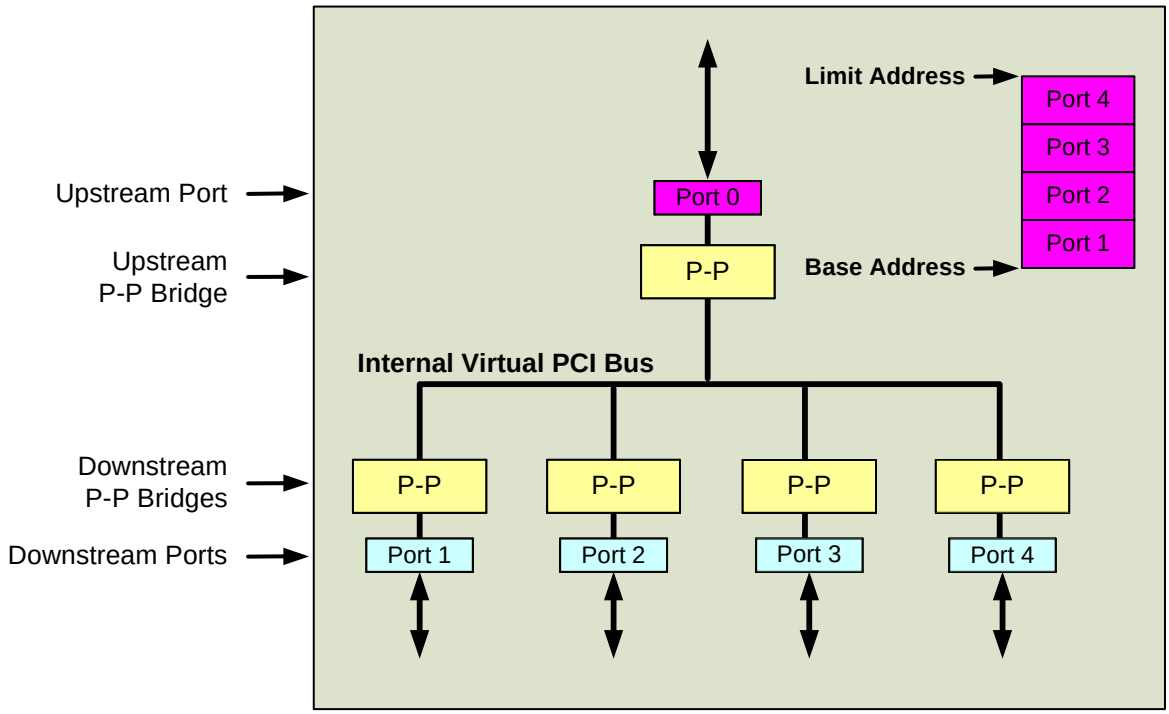

Note: *In Figure 7-3, Port 0 is designated as the upstream port; however, any port can be designated as the upstream port.*

7.3.1.2 Sample Pseudo Code

The following sample pseudo code demonstrates how to configure the PEX 8508's upstream and downstream ports after they are previously discovered by system enumeration software.

```
CFGTYPE0 Write busnum 01,devicenum 00 function 0 address 18h data 0009_0201h
    //Primary Bus Number 01, secondary Bus Number 02 and subordinate Bus Number
    09. At this step the virtual PCI bus in the PEX 8508 gets the Bus Number 02.
    After this, any access to Bus Number 02 from the upstream port would refer to
    this bus.
CFGTYPE1 Write busnum 02,devicenum 01 function 0 address 18h data 0003_0302h
    //Primary Bus Number 02,secondary Bus Number 03 and subordinate Bus Number 03.
    an endpoint is attached to Port 1
CFGTYPE1 Write busnum 02,devicenum 02 function 0 address 18h data 0004_0402h
    //Primary Bus Number 02,secondary Bus Number 04 and subordinate Bus Number 04.
    an endpoint attached to Port 2
CFGTYPE1 Write busnum 02,devicenum 03 function 0 address 18h data 0005_0502h
    //Primary Bus Number 02,secondary Bus Number 05 and subordinate Bus Number 05.
    an endpoint attached to Port 3
CFGTYPE1 Write busnum 02,devicenum 04 function 0 address 18h data 0006_0602h
    //Primary Bus Number 02,secondary Bus Number 06 and subordinate Bus Number 06.
    an endpoint attached to Port 4
CFGTYPE1 Write busnum 02,devicenum 01 function 0 address 20h data 02FF_0200h
    // need 0200_0000h to 02FF_FFFFh Memory space for Port 1.
CFGTYPE1 Write busnum 02,devicenum 02 function 0 address 20h data 03FF_0300h
    // need 0300_0000h to 03FF_FFFFh Memory space for Port 2.
CFGTYPE1 Write busnum 02,devicenum 03 function 0 address 20h data 04FF_0400h
    // need 0400_0000h to 04FF_FFFFh Memory space for Port 3.
CFGTYPE1 Write busnum 02,devicenum 04 function 0 address 20h data 05FF_0500h
    // need 0500_0000h to 05FF_FFFFh Memory space for Port 4.
CFGTYPE0 Write busnum 01,devicenum 00 function 0 address 20h data 06FF_0600h
    // the PEX 8508 will claim all Memory accesses from 0200_0000h to 06FF_FFFFh
    and would send it to the appropriate port. Any memory address not within
    any of the downstream Address spaces will go to the upstream port.
    // Now set the Bus Master Enable and Memory Access Enable bits on the
    upstream port and all downstream ports.
CFGTYPE0 Write busnum 01,devicenum 00 function 0 address 04h data 0000_0006h
CFGTYPE1 Write busnum 02,devicenum 01 function 0 address 04h data 0000_0006h
CFGTYPE1 Write busnum 02,devicenum 02 function 0 address 04h data 0000_0006h
CFGTYPE1 Write busnum 02,devicenum 03 function 0 address 04h data 0000_0006h
CFGTYPE1 Write busnum 02,devicenum 04 function 0 address 04h data 0000_0006h
    // For each port configured above, registers 24h, 28h, and 2Ch can be
    programmed to enable 64-bit Prefetchable Memory space for downstream devices.
    // Memory-Mapped access of all configuration registers listed above can also
    be performed by programming BAR0 and BAR1 (for 64-bit Memory spaces) for
    busnum 01, devicenum 00.
```

7.3.1.3 Sample Packet Transfer

When all ports are configured using the sample code provided in Section 7.3.1.2, the following occurs:

- 32-bit Memory transactions from the upstream port, destined between addresses 0200_0000h to 06FF_FFFFh, advance to the appropriate downstream port.
- 32-bit Memory transactions from a downstream port, between the addresses 0200_0000h to 06FF_FFFFh, advance to the appropriate downstream port (if the transactions are not in the Base-Limit range of that port).
- Transactions from a downstream port, outside the range of 0200_0000h to 06FF_FFFFh and outside PEX 8508 Memory-Mapped Register space (refer to Section 7.3.2 for details regarding Register space), advance to the upstream port

7.3.2 Using Base Address Registers (BARs) to Access Registers

Configuration requests can access only those registers that are defined by the *PCI Express Base r1.1*. These registers and the Device-Specific registers can all be accessed by Memory requests that target the Memory space defined by the upstream port **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively).

- Upstream port **BAR0** requests that 128-KB Memory space be set aside for internal PEX 8508 registers.
- Optionally, upstream port **BAR1** can be used to place this internal register Memory space anywhere in 64-bit System Memory space.
- After upstream port **BAR0** (and optionally, **BAR1**) is programmed, all register locations inside the PEX 8508 can be accessed from any port, using either Memory requests or Configuration requests.
- Each port consumes 4 KB of Memory space for internal registers. Port 0 is at 0000h to 0FFFh, Port 1 is at 1000h to 1FFFh, and so forth.

For example, if upstream port **BAR0** is programmed to 0100_0000h (using a Type 0 Configuration transaction) and the **Command** register *Memory Access Enable* bit is set (offset 04h[1]=1; again, programmed using a Type 0 Configuration transaction), all PEX 8508 registers can be accessed using Memory-Mapped Register accesses.

The following sections describe information specific to Transparent and Non-Transparent modes.

7.3.2.1 Transparent Mode Registers

The formula to locate register addresses for Transparent ports is as follows:

```
BAR0 + (port * 1000h) + register_offset
```

For example, to hit the **Memory Base and Limit** (offset 20h) for each port, refer to Table 7-1 and Figure 7-4. Table 7-1 defines how the registers in each port can be reached. All registers for all ports sit in the same BAR. Using the formula 1000h x Port Number provides the start address for the first register in a port. Figure 7-4 provides a graphical view of the BAR Memory space.

Table 7-1. PEX 8508 Memory-Mapped Register Access

Register	Location Address
Port 0 Base and Limit	0100_0020h
Port 1 Base and Limit	0100_1020h
Port 2 Base and Limit	0100_2020h
Port 3 Base and Limit	0100_3020h
Port 4 Base and Limit	0100_4020h

Note: *For a complete list of Memory-Mapped Register accesses, refer to Section A.1, "Serial EEPROM Memory Map."*

Figure 7-4. Using Memory-Mapped Access for PEX 8508 in Transparent Mode

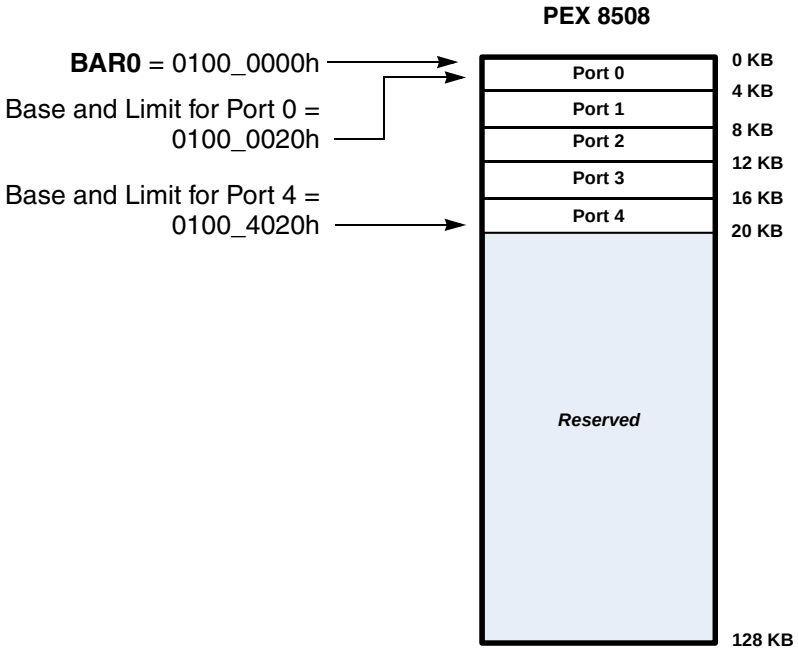

7.3.2.2 Non-Transparent Mode Registers

In NT mode, there are additional registers representing the NT Virtual and Link Interface endpoints. These registers exist at the fixed offsets of (refer to Figure 7-5):

- **Virtual Endpoint** – `BAR0 + 10000h`
- **Link Endpoint** – `BAR0 + 11000h`

Note: *For a complete listing of Memory-Mapped Register accesses, refer to Section A.1, "Serial EEPROM Memory Map."*

Figure 7-5. Using Memory-Mapped Access for PEX 8508 in Non-Transparent Mode

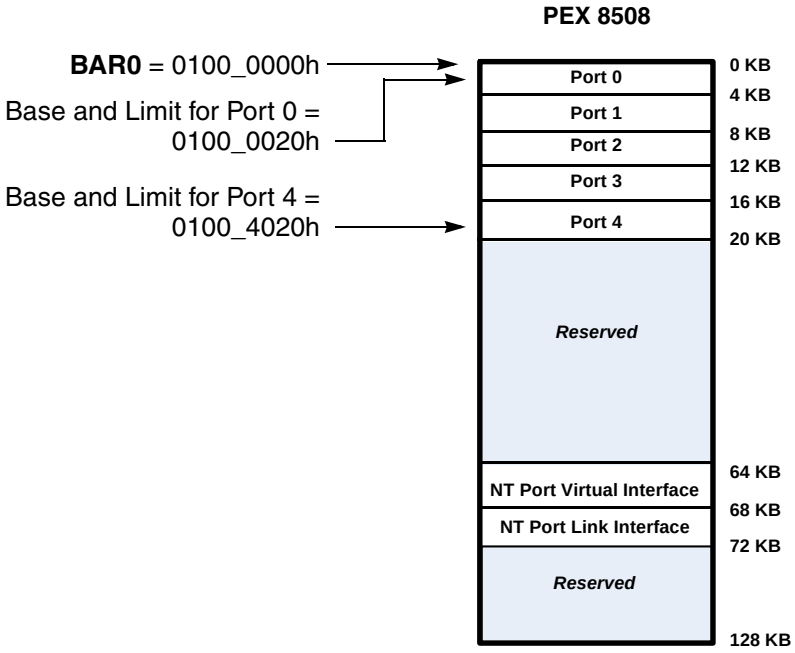

7.4 Interrupt Support

The PEX 8508 supports the PCI Express interrupt model, which uses two mechanisms:

- INT*x* Emulation
- Message Signaled Interrupt (MSI)

These interrupt mechanisms are discussed in Chapter 6, "Interrupts."

7.5 Hot Plug Support

The PEX 8508 supports the standard Hot Plug Controller (HPC) on all downstream ports. The detailed Hot Plug mechanisms are discussed in Chapter 11, "Hot Plug Support."

Note: *Refer to the PEX 85XX EEPROM – PEX 8518/8517/8508 Design Note for additional register programming information.*

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 8 Performance Metrics

8.1 Overview

The PEX 8508 includes features that optimize performance under several application scenarios. This chapter discusses the four major performance metrics:

- Internal fabric non-blocking nature
- Quality of Service (QoS)
- Sustained link throughput
- Port-to-port latency

These approaches emphasize metric optimization. In general, Host-centric applications, with transactions traveling between a wide upstream port and narrow downstream ports, are more *latency*-oriented. In comparison, peer-to-peer applications, where transactions are evenly distributed among all ports in a switch, are more *throughput*-oriented. However, achieving best performance is strongly application-dependent and the above principles are not necessarily always correct.

For example, if the PEX 8508 is linked with a graphics board in a Host-centric application, graphics port *throughput* becomes the most important performance consideration, not *latency.* Conversely, if the traffic pattern in a peer-to-peer application is lightly loaded and bursty, *latency* can overweigh *throughput*, to become the highest performance concern.

Therefore, when tuning performance, it becomes important to understand the interaction and dynamics among performance metrics. *For example*, some tunings, in which the sending of traffic from multiple ingress ports to a narrow egress port is avoided, improves overall throughput and latency by reducing hot-spot queuing within the system. Other tunings, *such as* aggregating traffic from multiple ports into a wider data path and processing them in a time-multiplexed manner, can optimize throughput at the cost of slightly longer port-to-port latency.

Once the system dynamics are understood, it is easier to exchange performance metrics against one another.

8.2 Non-Blocking Switch

In switch literature, *non-blocking* is used to indicate that a packet can be routed from an ingress port to an egress port, provided that not more than one packet is received by the same ingress port and not more than one packet is destined to the same egress port. A non-blocking switch is expected to fully route all packets for independent ingress traffic streams, with the destination uniformly distributed. The PEX 8508 is a non-blocking switch.

8.2.1 Queuing Topology

Three major queuing topologies are used in switch architecture:

- **Output Queuing (OQ)** – When a packet arrives at an ingress port, it is immediately placed into a buffer that resides in the egress port. If, in the worst case, there are *N* ingress ports simultaneously attempting to transmit packets to the same egress port, the output buffer is required to enqueue traffic *N* times faster than the egress port's dequeuing rate, requiring a Speed-up *S* of *N*.
- **Input Queuing (IQ)** – Ingress port packets have a set of *Virtual Output Queues* (VOQ). One of the packets, among all head packets in different VOQs to the same egress port, is allowed to be scheduled out of that ingress port during a given time slot. The key factor in achieving high performance using VOQ is the global scheduling algorithm, which is responsible for the selection of packets to transmit from the ingress ports to the egress ports in each time unit. The complexity of such scheduling algorithm is $O(N^2)$.
- **Input Queuing (IQ)** – Packet queueing is on the input port. A FIFO IQ is susceptible to Head of Line blocking if the head of the queue is blocked due to output contention, and other packets in the queue are destined to other contention-free output ports. Enqueue rate is the input rate, and thus the Speed-up is 1.
- **Combined Input-Output Queuing (CIOQ)** – A CIOQ switch is a non-blocking switch with *N* inputs and *N* outputs that operates *S* times as fast as the input rate. With a Speed-up of *S*, a CIOQ switch can remove up to *S* packets from each input and deliver up to *S* packets to each output within a time slot. Because $1 < S < N$, packets must be buffered before switching at the inputs, as well as after switching at the outputs. To avoid Head of Line (HOL) blocking, each input has a separate FIFO queue for each output – these are Virtual Output Queues (VOQs).

The PEX 8508 uses CIOQ as its internal switching topology to process traffic. Packets from one or more ports are aggregated to accommodate traffic from all ports within it at any time.

After extensive simulation to consider standard switching performance factors, including input traffic distribution, packet size distribution, output throughput, port-to-port latency, latency jitter, egress-to-ingress backpressure, as well as PCI Express-specific performance factors *such as* Physical Layer, Data Link Layer overhead, and packet-to-packet dependency caused by PCI ordering, PLX determined that using an internal speed-up factor of 1.25 allows the PEX 8508 to be non-blocking.

8.2.2 Port-to-Station Aggregation

As previously stated, the PEX 8508 PCI Express station is aggregated from multiple ports, provided that the combined port width is less than or equal to 8 lanes. Figure 8-1 redraws the PEX 8508 Transaction Layer architecture, by explicitly dividing a PCI Express station into individual ingress and egress ports, grouping the ingress ports into a source station and the egress ports into a destination station.

In a PCI Express source station, the write port to the ingress packet RAM is shared by up to five ingress ports in a Time Domain Multiplex (TDM) manner. The wider the port, the more TDM slots are assigned to that port. Within VOQs of a single source station, if multiple packets from different ingress ports are available to be dispatched to the same destination, a Round-Robin arbiter controls which ingress port to select next.

Moving to the internal fabric, the Read ports to ingress packet RAM and Write ports to egress packet RAM are controlled by the PLX implementation of the CIOQ scheduling algorithm, where a unit in scheduling is the station, rather than a port.

In a PCI Express destination station, the Read port to the egress packet RAM is shared by up to five egress ports in a TDM manner as well. Furthermore, there are five independent egress schedulers. Egress schedulers follow Virtual Channel arbitration, as required by the *PCI Express Base r1.1*.

Figure 8-1. PEX 8508 Queuing Data Structures

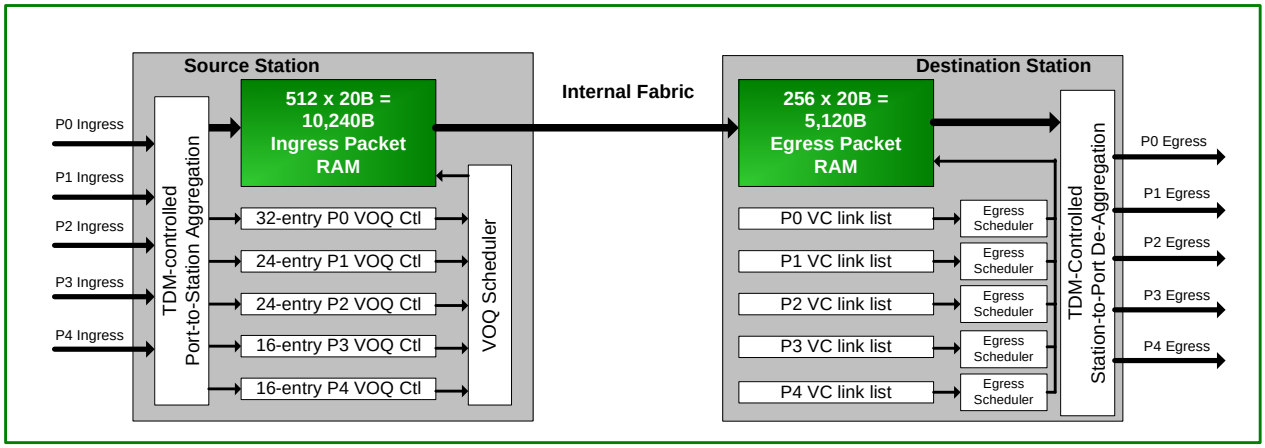

Note: *The PEX 8508 has only one station. Therefore, the Source and Destination stations are equivalent.*

8.2.3 RAM and Queue Size

Table 8-1 defines the RAM and Queue size built into the PEX 8508. The smallest unit of ingress and egress packet RAM is defined as a *beat*, which can store 20 bytes of data. In the PEX 8508, the smallest packet (12B Header) takes 14 beats to store in a packet RAM, and the largest packet (16B Header + 256B Payload + 4B digest = 276B) takes 8 beats to store in packet RAM.

In Table 8-1, the number 32 under the cell **Ingress VOQ Entries**, **Per Port** indicates the maximum number of VOQ entries allocated to an ingress port. The number of VOQ entries allocated to ingress ports are variable on the PEX 8508, and ranging from 32 for Port 0, 24 for Ports 1 and 2, and 16 for Ports 3 and 4.

Each VOQ entry holds one Transaction Layer packet. For ingress ports, the incoming packet can travel to the station, containing up to five egress ports. Also, for egress ports, up to two VCs are supported, with each VC potentially having three different packet types – Posted, Non-Posted, and Completion (P, NP, and Cpl, respectively).

Each port on the egress side can have up to six queues, for holding packets from two supported VCs and three supported packet types. A queue that stores packets of a unique VC and a unique type is referred to as a *Virtual Channel and Type (VC&T) queue*. Again, some queues can be completely empty and some queues can contain more than one packet. The maximum number of packets held by a single egress port is limited by the number of egress packet RAM beats allocated to that port.

It can be calculated from Table 8-1 that the total packet RAM size for the PEX 8508 is 15,360 bytes. Assume the maximum Transaction Layer Packet (TLP) size is 276B with a Maximum Payload Size (MPS) of 256B. Theoretically, the PEX 8508 can store up to 111 MPS packets.

Table 8-1. PEX 8508 Data Structure Size

Data Structure Name	Per Port	Overall
Ingress Packet Name	Programmable	512 beats or 10,240 bytes
Ingress VOQ Entries	32 (Port 0) 24 (Ports 1 and 2) 16 (Ports 3 and 4)	112
Egress Packet RAM	256 beats or 5,120 bytes for all five ports	256 beats or 5,120 bytes
Egress VC&T Queues	6	24

8.3 Quality of Service

Quality of Service (QoS) is a performance differentiation feature offered by PCI Express to manage multiple traffic classes of different characteristics. An application assigns a Traffic Class (TC) value to individual Transaction Layer packets, according to the QoS requested by the class to which the transaction belongs. The static TC value tagged to each packet is dynamically mapped to a VC as it passes through a system PCI Express-capable device. The TC value ultimately determines the relative priority of a single packet as it traverses the PCI Express fabric, as well as the accumulated bandwidth allocated to the packets that belong to the same class.

8.3.1 Virtual Channel

The PEX 8508 supports up to two Virtual Channels (VCs) – VC0 and VC1. Ports 0, 1, and 2 support VC0 and VC1. Ports 3 and 4 support only VC0. Each VC has its own buffer resource allocation and data path. For a single port, VC configuration and properties are determined by the **Virtual Channel Extended Capability** registers (offset 148h to 1C4h).

Registers described in the Virtual Channel Extended Capability register map cover switch egress and ingress ports. Registers related to packet arbitration are egress-specific, whereas registers defining TC/VC mapping and Low-Priority VC Count are applicable to egress and ingress ports. [Refer to Table 13-10, "PEX 8508 Virtual Channel Extended Capability Register Map (All Ports)," for further details.]

Virtual Channel and traffic labeling allow independent physical resources to handle differentiated traffic. The **VC0 Resource Control** and **VC1 Resource Control** registers (offsets 15Ch and 168h, respectively) contain bits that control TC/VC mapping.

Across various ports, the PEX 8508 supports both symmetric and asymmetric TC/VC mapping. In the latter approach, the TC/VC mapping is port-independent and configured with different values per port.

The PEX 8508 default configuration sets all TC[7:0] to VC0, as provided in the *TC/VC0 Map* bits. For applications requiring two Virtual Channels, TC[7:1] can be mapped to VC1 by removing them from the **VC0 Resource Control** register *TC/VC0 Map* bits and adding them to the **VC1 Resource Control** register *TC/VC1 Map[0]* and *TC/VC1 Map[7:1]* bits.

8.3.2 Packet Arbitration

Because of CIOQ switch architecture and multiple VC support, the PEX 8508 functions with several arbitration/scheduling points distributed in the data path from an ingress port to an egress port.

This section discusses the arbitration/scheduling/backpressure algorithm used in the Source Scheduler, internal fabric, and egress scheduler.

8.3.2.1 Source Scheduler

Source Scheduler is essentially the VOQ scheduler depicted in Figure 8-1. The Source Scheduler functions as follows:

- From all VOQ entries (each entry represents a single packet) belonging to a single ingress port, identifies one packet to be dispatched to the appropriate destination station, when egress RAM space is available
- Arbitrates among multiple-ready packets from different ingress ports with a Round-Robin mechanism
- Breaks deadlock potential by following Conventional PCI Ordering rules

Note: *Packets to different egress ports are selected with oldest first criteria on a per-queue basis. This policy offers optimum fairness and performance properties at low complexity. Packets to different VCs are selected by allowing VC1 higher priority, if configured as such. Conventional PCI Ordering rules are enforced.*

The Source Scheduler is capable of handling variable-length packets. It can schedule one packet out of a source station every Clock cycle, regardless of packet size.

The Source Scheduler has two programmable fields:

- High-Priority Virtual Channel
- Device-Specific Relaxed Ordering

8.3.2.2 High-Priority Virtual Channel

When two VCs are available, VC1 packets are, by default, given higher priority over VC0 packets. VC1 packets can be programmed to have the same priority as VC0 packets, by setting the **Port VC Capability 1** register *Low-Priority Extended VC Count* (link) bit to 1 (offset 14Ch[4]=1). The default setting can be changed only by serial EEPROM.

8.3.2.3 Device-Specific Relaxed Ordering

PLX Relaxed Ordering capability is supported to enhance the performance of "push-only" traffic (Posted packets, *such as* Memory Writes and messages) in the PEX 8508. Device-Specific Relaxed Ordering mode is enabled when any bit within a **Device-Specific Relaxed Ordering Mode** register *Enable PLX Relaxed Ordering* field is set to 1:

- Port 0 – offset BFCh[7:0]
- Port 1 – offset BFCh[15:8]
- Port 2 – offset BFCh[23:16]
- Port 3 – offset BFCh[31:24]
- Port 4 – offset BE4h[7:0]

According to PCI Ordering rules, Posted packets are not allowed to bypass previously Posted packets of the same VC&T, regardless of whether they are targeting different egress ports. In applications *such as* storage area networks or IP networks, where Posted PCI Express packets are used to transmit encapsulated data traffic through switches, unnecessary serial dependency might be created in the Source Scheduler for those Posted packets coming from the same ingress ports, but going to different egress ports, if strict PCI Ordering rules were followed. This can result in dramatically degraded overall switch throughput due to HOL blocking. For uniformly distributed traffic, throughput can be reduced to 58% of the input line rate.

The **Device-Specific Relaxed Ordering Mode** registers can be used to enable the Device-Specific Relaxed Ordering capability. There is an enable bit for TCs in each ingress port. All packets are allowed to bypass older packets from the same ingress port and TC. Packets targeting different egress ports are free to proceed without waiting for ordering dependency to be cleared. Meanwhile, packets targeting the same egress port are processed "in order," because there is no performance gain. This is how PLX allows VOQ in the PCI Ordering rules, and gets around HOL blocking.

Because the Enable bit is TC-based, taking advantage of Device-Specific Relaxed Ordering mode requires the PEX 8508 to be programmed with symmetric TC/VC mapping first.

Posted traffic benefits most from this mode. To take advantage of Device-Specific Relaxed Ordering mode, without violating other Ordering rules defined by the *PCI r3.0*, it is suggested to restrict outstanding traffic flow to be "Posted only" and shut down all Non-Posted packets.

There are two usage models:

- Restrict all Posted traffic requiring high-throughput in VC1 and program all TCs belonging to VC1 to enable Relaxed Ordering.
- Software disables Device-Specific Relaxed Ordering mode in all TCs beforehand, performing all setups that involve Non-Posted packets, and then enabling an *Enable PLX Relaxed Ordering* bit when the system enters pure Data Transfer mode. When the Data transfer completes, disable Device-Specific Relaxed Ordering mode.

Note: *A variation of this feature, Relaxed Completion Ordering, allows only Completions to bypass Posted packets, while preserving the remainder of PCI ordering. (Refer to Section 4.3.2.2, "PEX 8508 Relaxed Completion Ordering," for details.)*

8.3.2.4 Internal Fabric Backpressure

The internal fabric of the switch provides egress packet RAM space available status. The Source Scheduler never transmits a packet to overflow egress packet RAM. Moreover, to prevent packets in a particular VC&T from occupying the majority of egress packet RAM, and to speed up backpressure from Egress queues to Ingress queues and ultimately to external devices in the case of congestion avoidance, VC&T-based (VC and Type – P, NP, Cpl) packet cutoff information is passed from egress ports to the Source Scheduler.

Egress queue congestion occurs when the packet arrival rate overcomes the packet dispatching rate. There are two causes of congestion:

- Insufficient credit to transmit the packets
- Insufficient bandwidth to transmit the packets as quickly as they arrive

In either case, if the situation continues, eventually the Egress queues of the congested port will fill up.

The PEX 8508 utilizes a watermark mechanism to cut off additional packets from the ingress side when the Egress queues back up. With some egress ports cut off, Ingress queues that contain packets targeting those egress ports could then fill. A filled Ingress queue prevents additional credit to its link partner, which causes that external device to stop transmitting packets in that VC&T.

ITCH VC&T Threshold *x* registers exist for each VC&T in an egress port. All five ports share the same programmed value. Each VC&T has its own upper and lower limit. If more data than the programmed upper limit is queued, no more packets from that VC&T can be scheduled across the internal fabric, thereby cutting off that VC&T flow. After cutting off the VC&T flow, an Egress queue eventually drops below its lower limit, as packets are scheduled out of the egress port. This event turns On the internal fabric VC&T-enabling flags, which allow that VC&T to resume flow. [Refer to the **ITCH VC&T Threshold 1**, **ITCH VC&T Threshold 2**, and **ITCH VC&T Threshold 3** registers (offsets C00h, C04h, and C08h, respectively) for further details.]

Two rules are used for programming the **ITCH VC&T Threshold *x*** registers:

- The unit value for the upper and lower limits is equivalent to 8 beats. The maximum value programmable in the upper limit is a function of the port width, defined in Table 8-2.

  A x4 egress port can handle no more than 256 beats, a meaningful value for the upper and lower limits is bounded by 256/8 = 32 (or 20h). In the device, the default value used by the upper and lower limits (FFh) is larger than its maximum legal value (20h). Therefore, by default, the backpressure mechanism is not triggered.
- The upper and lower limits must be different, with the upper number being larger than the lower number by at least two units.

The main objective is to avoid clogging the egress RAM with excessive packets of the same type that might prevent packets of other types from making fast forward progress inside the switch. If necessary, program the upper and lower limits.

Table 8-2. VC&T Threshold Limits

Port Width	Maximum Upper Limit in Beats
x1	64
x2	128
x4	256

8.3.2.5 Egress Scheduler

In each egress port, the PEX 8508 strictly follows VC and port arbitration mechanisms, as defined by the *PCI Express Base r1.1*.

Virtual Channel Arbitration

In the context of scheduling traffic in VC0 and VC1, the main goal of the egress scheduler's VC arbitration is to provide differentiated services between data flows within the fabric.

The two VC arbitration choices are:

- **Strict priority** – VC1 always prevails over VC0
- **Round-Robin, or "Hardware-Fixed Arbitration" in the** *PCI Express Base r1.1* – Alternate between VC0 and VC1

The strict priority selection is made by clearing the **Port VC Capability 1** register *Low-Priority Extended VC Count* bit (offset 14Ch[4]=0). The default value is strict priority, and can only be changed by serial EEPROM initialization.

VC0 and VC1 share the low-priority pool when the *Low-Priority Extended VC Count* bit is set to 1 (by way of serial EEPROM). Within the low-priority pool, Round-Robin arbitration can be selected.

Egress Port Arbitration

Regarding egress ports, the PEX 8508 supports only one port arbitration mechanism – the non-configurable hardware-fixed arbitration scheme. The oldest Ready packet from all ingress ports arriving at the current egress port is selected first. *Ready* packet is defined as a packet with available egress credit and no ordering violations.

8.4 Throughput

8.4.1 Theoretical Upper Limit

PCI Express allows for bidirectional traffic capability and scalable widths, allowing it to closely match the needed bandwidth. As discussed in this section, compared to the 2.5-Gbps raw bandwidth provided by each SerDes lane, the achievable data Payload efficiency, assume the Maximum Payload Size (MPS) of 128B is approximately 70%.

8.4.1.1 Physical Layer Overhead

The 2.5-Gbps serial data on a SerDes is encoded with additional information for clock recovery and error detection through 8b/10b encoding. When the additional information is removed, a 2.0-Gbps data rate remains, which is 80% of the starting bandwidth.

The PHY Layer also adds a 1-byte start symbol (STP) and a 1-byte end symbol (END or EDB) to the packet size, thereby introducing 2 bytes of overhead per TLP.

Also in the PHY Layer, SKIP Ordered-Sets are used to compensate for differences in frequency between bit rates at opposite ends of a Link. The *PCI Express Base r1.1* specifies a clock frequency tolerance of 600 parts per million (ppm), which in turn requires a SKIP Ordered-Set to be transmitted within the range of 1,180 to 1,530 symbol times. This causes the achievable efficiency to drop another 4/1,180 = 0.34%.

8.4.1.2 Data Link Layer Overhead

To ensure data integrity passing over the wire, the *PCI Express Base r1.1* states that the DLL (Data Link Layer) adds a sequence number at the start of the packet and an LCRC integrity check at the end of the packet. The sequence number is 2 bytes, and the LCRC is 4 bytes, thereby introducing 6 bytes of overhead per TLP.

In addition to the overhead inherent in TLP Payload transmission, the *PCI Express Base r1.1* uses the same wire to transmit DLLPs (Data Link Layer Packets). ACK (acknowledge) and updateFC are the two most frequently used types of DLLPs during standard run time, where throughput matters. ACK is used to acknowledge TLP receipt. updateFC is used to provide additional credits, which enables additional TLP transactions.

DLLPs are structured so that a single ACK can represent receiving multiple TLPs, reducing the total number of ACKs required. Similarly, a single updateFC is structured so that credit for more than one packet can be extended at a time, reducing the number of required updateFCs per TLP. The size of a single DLLP is 8 bytes. In the worst case, two outgoing DLLPs are formed for each incoming TLP, which equals 16B in per TLP overhead. In the best case, there is zero DLLP overhead for incoming TLPs. DLLPs flow in the opposite direction of TLPs, as they are feedback mechanisms. For one-way TLP traffic, the DLLP overhead does not impact overall link utilization.

8.4.1.3 Transaction Layer Overhead

All PCI Express payloads are encapsulated in a TLP. A TLP contains a Header portion that provides the PEX 8508 with routing information for the packet. The Header can be 12 or 16 bytes.

According to the *PCI Express Base r1.1*, Maximum Payload Size (MPS) can range from 128 to 4,096 bytes. The PEX 8508 supports an MPS of up to 256 bytes.

The **Device Control** register *Maximum Payload Size* field (offset 70h[7:5]) default MPS is 128 bytes. Software can change the default value; however, the entire system must have a consistent MPS. In general, longer payloads are more efficient, but require more on-chip resources to buffer, and can cause much worse hot spot congestion in bursty traffic flows.

A TLP can also incur additional overhead when end-to-end data integrity is essential. In such cases, a 4-byte ECRC is added as another type of overhead to the end of the packet.

Note: *Refer to the **Device Control** register Maximum Payload Size field (offset 70h[7:5]) for MPS limitations.*

8.4.1.4 PCI Express Efficiency Upper Bound Summary

Table 8-3 summarizes PCI Express inherent efficiency for 0B, 4B, 8B, 40B, 128B, 256B, and 4,096B Payload sizes on various negotiated link widths. The 4,096B table row is provided for reference only, as the PEX 8508 supports MPS of up to 256 bytes.

The table columns provide three types of variations:

- Comparing to raw SerDes bandwidth of 2.5 Gbps versus 2.0 Gbps after 8b/10b decoding.
- 0% additional DLLP generation (unidirectional TLP traffic) versus 100% DLLP generation, assuming traffic is equal in both directions. 0% DLLP assumes that DLLP traffic is not injected into the TLP stream, *such as* a unidirectional traffic stream. 100% DLLP assumes that for every TLP transmitted, an additional ACK DLLP and updateFC DLLP are generated in the reverse direction for a bidirectional fully loaded traffic stream for a bidirectional fully loaded traffic stream.
- Non-Payload TLP overhead of 12B versus 20B (16B Header + ECRC).

In summary, the larger the Payload, the more efficient the PCI Express communication. For 256-byte Maximum Payload Size supported by the PEX 8508, the limit efficiency compared to the raw 2.5-Gbps bandwidth is between 68 to 73.92%.

Note: *Not all factors are reflected in this table. The SKIP Ordered-Set drops another 0.3%. Any credit shortage or transient congestion can significantly drop.*

Table 8-3. Throughput Theoretical Upper Limit

Bytes of Payload	2.5 Gbps Raw Bandwidth				2.0 Gbps Raw Bandwidth after 8b/10b Decoding			
	0% DLLP		100% DLLP		0% DLLP		100% DLLP	
	12B	20B+	12B	20B+	12B	20B+	12B	20B+ ECRC
0	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
4	13.28%	10.00%	8.00%	6.64%	16.60%	12.50%	10.00%	8.30%
8	22.80%	17.68%	14.48%	12.24%	28.50%	22.10%	18.10%	15.30%
40	53.12%	46.88%	42.00%	38.00%	66.40%	58.60%	52.50%	47.50%
128	68.96%	65.44%	62.24%	59.36%	86.20%	81.80%	77.80%	74.20%
256	73.92%	71.84%	69.92%	68.00%	92.40%	89.80%	87.40%	85.00%
4,096	79.36%	79.20%	79.04%	78.88%	99.20%	99.00%	98.80%	98.60%

8.4.2 Single-Stream Throughput

When data flows in a single-packet stream from a fixed-ingress to fixed-egress port, its throughput can be optimized on both the ingress and egress sides. The method for optimizing throughput on both sides is discussed in the following sections.

8.4.2.1 Ingress Side

Optimize Ingress Credit Allocation

A TLP cannot be transmitted to the switch without the switch providing sufficient ingress credits beforehand. When a credit is advertised, it indicates a guaranteed storage available in the credit transmitter at that time. If there is insufficient or untimely ingress credits advertised from the PEX 8508 to its link partner, the incoming TLP stream does not sustain at the highest possible rate.

Amount of Ingress Credit Required Calculation

The PEX 8508 supports up to six VC&Ts per port. The amount of ingress credits advertised in each VC&T is expected to be sufficient to cover the round-trip delay from the time the external device schedules a TLP for transmission in its Transaction Layer to the time the external device receives the replenishing credit from the PEX 8508 in the same VC&T.

To enable a burst of TLPs of the same VC&T to enter the PEX 8508 without interruption, use the following empirical equation:

```
Ingress_Credit_Advertised = (Round_trip_time_in_symbol times x link_width)
/ packet_size_in_bytes
```

Round-trip latency, which can range from 110 to 400 ns (28 to 100 symbol times), is determined by both the PEX 8508 and external device and consists of the following:

- Latency for incoming TLP to travel the entire PEX 8508 ingress data path
- Delay from writing the first byte of the packet into ingress packet RAM until writing the last byte of the packet into ingress packet RAM
- Latency for Source Scheduler to transmit the packet to egress packet RAM and free up the ingress buffers for this TLP packet
- Latency for the PEX 8508's ingress credit scheduler to generate an updateFC packet
- Latency for this updateFC DLLP to travel the PEX 8508 egress data path to SerDes
- Delay in SerDes
- Latency for this updateFC DLLP to travel the ingress data path of the external device
- Latency for external device to process the updateFC DLLP and update its Credit Limit Counter
- Latency for external device to schedule the next TLP in the same VC&T out
- Latency for the external device to move the new TLP across its egress data path to the SerDes
- Final delay in SerDes

For example, suppose a link with a 400 ns round-trip time has 8 lanes and a stream of Posted transactions is broken into packets of 64B Payload each. The amount of Posted type Header credit needed to sustain a steady incoming traffic flow is approximately 100 x 8 / (16 + 64) = 10.

The smaller the Payload size, the higher demand on the ingress Header credit to be advertised. Using the same example provided above, but changing Payload size from 64B to 4B, the number of ingress Header credits required is 40. As previously stated, the PEX 8508 has a total of up to 32, 24, or 16 VOQ entries in its ingress ports, which is insufficient to sustain 4B-Payload TLPs back-to-back. Two VOQ entries must always be used for 1 NP and 1 Cpl Header credit, the most P Header credits are the remaining VOQ entries. For Port 0, this is 30 VOQ entries. Working backward, it can be calculated that Port 0 can sustain a burst if the average Payload Size is 10B and Header is 16B.

Program Ingress Credit Threshold Rules

The PEX 8508 Ingress Credit Handler registers are used to program the ingress credit value for each VC&T in the ingress ports. Refer to Section 13.13.12, "Ingress Credit Handler (INCH) Registers," to view the registers used to program the ingress credit value. The registers range from offsets A00h to A50h and A60h to A68h.

In these registers, both Header credit and Payload credit threshold are writable. The following rules are used to program the Ingress Credit Handler registers:

- One unit of Header credit threshold represents one packet. A value of *1* allows the PEX 8508 to advertise one Header credit, *2* allows two Header credits, and so forth. Bits [13:9] allow a maximum of 31 Header credits to be programmed to VC&T.
- One unit of Payload credit threshold represents 16B of data. When the PEX 8508 Maximum Payload Size is set as 256, a value of at least *16* is required to be programmed for Posted and Completion packets. Bits [8:0] allow a maximum value of 2^9 Payload credit units to be advertised. For Posted and Completion types, bits [2:0] are ***reserved***, which forces the Payload credit threshold to be powers of 8. This effectively makes the granularity for Posted and Completion credit threshold types increase to 16B x 8 = 128B. For Non-Posted types, similar restrictions do not apply, because the Payload size is never more than 4B.
- For all ingress VC&Ts, the total Header credit cannot exceed the maximum number of Header credits available for that port, which is the number of VOQs for that port. Port 0 has 32, Ports 1 and 2 have 24, and Ports 3 and 4 have 16.
- For all VC&Ts in all five potential ports in a source station, the total Payload credit cannot exceed 1,024, which is the size of the Ingress RAM.

Ingress Credit Threshold Programming

Two methods are available for changing the ingress credit threshold – by way of a serial EEPROM or direct CSR programming. The first approach is straightforward – program the required values and the link produces the programmed values. However, the latter approach requires further explanation.

In the direct CSR programming method, if a credit threshold requires an increase, perform a CSR Write. This immediately results in a new updateFC to be transmitted, carrying the newly increased credits. In contrast, if a credit threshold requires a decrease, the *PCI Express Base r1.1* does not provide a pre-defined method for reclaiming unused credit previously advertised to the external device. The danger of programming a smaller credit threshold than the initial value is that ingress packet storage can overflow if the external device is not aware of the credit threshold, and continues sending packets according to the initial credit threshold. Use the following approach to avoid packet RAM overflow:

- Upstream port
  - Use CSR access to program the VC&Ts whose credit requires a decrease.
  - Transmit "side-impact free" traffic from Host to those VC&Ts, to deplete all credits to be reclaimed. Before the surplus credits are completely reclaimed, the PEX 8508 transmits updateFC for that VC&T. After the amount of incoming traffic attains the difference between the old and new credit thresholds, the PEX 8508 starts transmitting fresh updateFC DLLPs for incoming TLPs.
- Downstream ports
  - Program all ingress credit threshold CSRs in all VC&Ts to the needed values.
  - Execute a Secondary Bus Reset in the **Bridge Control** register.
  - Release the reset. The newly programmed values take effect afterward.

Ingress credit threshold registers are not sticky after a primary reset; therefore, this sequence requires repeating after any primary reset.

Tip: *If a serial EEPROM is available and you want to experiment with credit values, initially program all thresholds to 1. There is no impact in increasing the credit (up to the maximum of PEX 8508 resources). This allows for the most flexibility in your experiments.*

8.4.2.2 Egress Side

Provide Sufficient Egress Credit

The simplest way to improve an egress port's throughput is to provide the PEX 8508 with sufficient egress credit. In general, the number of egress credit required by the PEX 8508 follows the same equation as the ingress credit calculation:

```
Egress_Credit_Required = (Round_trip_time_in_clocks x link_width) /
packet_size_in_bytes
```

For the PEX 8508 to achieve pipelined performance, the external device is suggested to advertise at least 2 MPS worth of Payload credits. For 128B MPS, the Payload credit is greater than or equal to 16 (16 credits = 256B = 2 MPS). For 256B MPS, the Payload credit must be greater than or equal to 32 (2 MPS). The PEX 8508 scheduler requires 1 MPS of credit to send any length TLP that has a Payload. Therefore, a 2 MPS credit minimum is recommended for smaller than MPS payloads, and 3 MPS is recommended for MPS payloads.

8.4.3 Multiple Stream Throughput

8.4.3.1 Enable Device-Specific Relaxed Ordering

A burst of Posted traffic, from one ingress port targeting multiple egress ports, is subject to HOL blocking by the PCI Ordering rules.

The PEX 8508 does not support optional Relaxed Ordering bits in TLP, as specified in the *PCI Express Base r1.1*, Table 2-23. By default, all packets entering from a specific port are dispatched to their respective destinations, based upon strict ordering.

However, as described in Section 8.3.2.1, the PEX 8508 provides its own Relaxed Ordering to overcome the packet-to-packet dependency in a burst of Posted traffic from the same ingress port, but to different egress ports.

8.4.3.2 Avoid Hot Spots

A hot spot forms when multiple ingress ports attempt to transmit packets to the same egress port, and the overall influx bandwidth outweighs the efflux bandwidth. If the hot spot is not transient, the hot spot port throughput can appear High. However, eventually the Egress queues fill, backpressuring the Ingress queues. When the Ingress queues fill, ingress traffic is backpressured, potentially impacting traffic flow not targeting the congested egress port. As a result, the switch overall average throughput is dramatically reduced. PCI Express does not provide a mechanism to recognize and avoid hot spots. It is therefore left to the system designers to understand and avoid this pitfall.

8.4.4 Throughput and Packet Size Relationship

In general, sustained throughput increases as Payload Size increases due to the increased PCI Express protocol efficiency. However, the following secondary effects can also affect throughput:

- In peer-to-peer applications, longer packets can result in less interleaved or randomized egress port distribution compared to shorter packets. This increases the chance of building up transient congestion in egress ports, and can negatively impact overall throughput.
- Longer packets require fewer Header credits per unit time, and are therefore less likely to idle the link while waiting for additional Header credit.
- Longer packets burn up Payload credits faster and can stall DLLPs behind the long TLP longer, potentially causing credit starvation. If there is insufficient link credit (3 TLPs worth or more), shorter packets might provide better throughput.
- Posted packets block younger packets of other types (Non-Posted and Completions). In a system with minimal credits, Posted packets should receive the strongest consideration when allocating credits.

It is recommended to carefully compare the benefits and drawbacks of using longer packets.

8.4.5 Data Link Layer Considerations

8.4.5.1 Arbitration between DLLP and TLP

To reduce DLLP overhead on the wire, the PEX 8508 uses the following fixed-priority scheme to determine what transmits next:

1. Completion of a TLP or DLLP currently in transmission.
2. Initialization Flow Control (InitFC) DLLPs.
3. NAK DLLP.
4. ACK DLLP, due to receipt of a duplicate TLP or ACK Latency Timer expiration. *
5. Update FC DLLP, due to the FC Update Pending Timer expiration. *
6. Retry Buffer TLP, due to received NAK or Retry timeout.
7. New TLPs. *
8. Update FC DLLP, due to change in available credits. *
9. Power Management DLLP.
10. ACK DLLP for the last received TLP. *

Among these ten categories, the five most frequently detected new packets are noted with an asterisk (*). Updated-FC, InitFC, and ACK DLLPs appear twice – once as higher priority than TLPs, and again as lower priority than TLPs. A regular DLLP turns into a higher priority DLLP based upon a programmable timer. The basic idea is to reduce the number of DLLPs, the timers provide the opportunity to collapse multiple DLLPs into 1. The timers are discussed in Sections 8.4.5.2 and 8.4.5.3.

8.4.5.2 DLLP ACK Frequency Control

The **ACK Transmission Latency Limit** register (offset 1F8h) indicates a minimum amount of time (in 4-ns clocks) that the switch waits before prioritizing an ACK. By setting this register to the minimum value of 2 (refer to note below), ACKs are typically always transmitted with high priority, allowing the most DLLP traffic and the smallest possible Retry buffer in the other device on the link.

Note: *2 is the minimum value that has an effect; 0 or 1 wait for 255 clocks.*

The larger the number written into this CSR, the larger the chance of ACK collapse, and the more efficient the outgoing TLP throughput can be.

However, by setting the *ACK Transmission Latency Limit* field to the maximum (255), 255 symbol times (4 ns each) to occur before prioritizing an ACK. If the Retry buffer in the external device is not sufficiently deep, it can slow the incoming TLP rate. On a x4 link, 1,020B can be transmitted in 255 symbol times, which is 51 20B packets. The external device would need to have a Retry buffer that could store more than 51 TLPs, so as not to impact the back-to-back burst of incoming TLPs.

Because programming a smaller value into this CSR decreases egress TLP throughput but can increase the ingress TLP throughput, a tradeoff must be addressed.

If there is no TLP traffic, an ACK can be transmitted earlier than the timer indicates, as a low-priority DLLP.

The initial value depends upon the programmed link width. However, the value can be overwritten by serial EEPROM or a regular CSR Write.

Note: *The PEX 8508 Retry buffer is 64 entries.*

8.4.5.3 DLLP updateFC Frequency Control

The **INCH FC Update Pending Timer** (Ports 0 through 3, offset 9F4h[31:0]) and **INCH Port 4 Control** (Port 4, offset 9F0h[7:0]) registers control the length of time a port must wait before prioritizing an updateFC DLLP. Before the timer expires, TLPs have priority over updateFC DLLPs. After the timer expires, updateFC DLLPs move to higher priority. The value programmed into this CSR is a counter expiration value. All six VC&Ts in an ingress port share the same counter upper limit; however, each has its own set of counters, for counting up.

The smaller the value written into these registers, the sooner an updateFC DLLP becomes higher priority; therefore, the sooner the updateFC DLLP is transmitted. The sooner an updateFC is transmitted, the less likely the chance to collapse two VC&T updateFCs. However, even for small timer values, only one updateFC is typically sent per each incoming TLP. The updateFC is broken into multiple DLLPs for each incoming TLP only if there are insufficient resources to replace the credit.

Note: *Each VC and type has its own updateFC. Only updateFCs for the same VC&T can be collapsed.*

The *PCI Express Base r1.1* guidelines for the FC Update Pending Timer are provided in Table 8-4.

For implementation, a value of 01h or 00h into the CSR results in waiting 255 symbol times. The smallest value is 10h. The initial value of 00h is effectively 255 symbol times.

Table 8-4. FC Update Pending Timer Guidelines

Maximum Packet Size	Link Width	Recommended Timer Count
128 bytes	x1	76h
	x2	40h
	x4	24h
256 bytes	x1	D0h
	x2	6Ch
	x4	3Bh

8.5 Latency

8.5.1 Queuing Effect

In switches with large internal buffers, the latency increases once internal queuing is developed. The packet at the end of an egress VC&T queue does not transmit until all packets in front of it are transmitted. Assume the egress RAM of a x4 port is packed with packets of the same VC&T, draining the entire egress RAM takes 2,560 clocks (512 beats x 20B per beat / 4B/clock). Worst-case packet latency can be as long as 10 µs.

To overcome the queuing effect, attempt the following:

- Avoid creating hot spots. Ensure that the upstream port width in a Host-centric application matches the sum width of all active downstream ports. Otherwise, the upstream port can easily become a hot spot when all downstream ports are attempting to transmit packets to it.
- Program a small Egress queue packet upper and lower limit, to avoid packet accumulation in an egress port. Section 8.3.2.4 describes how to program these thresholds. Lower latency is achieved at the cost of reducing the PEX 8508's capability to buffer transient congestion.
- Reduce traffic load. Lighter traffic is less likely to experience congestion and can drain relatively faster, as the egress links can drain at the full link rate.

8.5.2 Time Division Multiplex Effect

As previously illustrated, the PEX 8508 source station employs port-to-station aggregation, and the destination station employs station-to-port de-aggregation. Time Division Multiplex (TDM) controls aggregation and de-aggregation. Usually, waiting for a proper TDM slot to process packet coming from or going to a particular port increases the latency. The wider the port, the more TDM slots that port owns; therefore, the less latency contributed by TDM.

Within the station, only a subset of 8 lanes are connected to SerDes. One approach to reduce latency is to strap the port as a wider port and allow it to negotiate down to the expected link width. *For example,* if there is only a x2 port owned by the station, the port can be strapped as x4, and allowed to become x2 later through the normal link training process. As a result, all TDM slots in this station are acquired by the x2 port. The worst-case TDM effect for a x1 or x2 port is 14 symbol times = 57 ns.

8.5.3 High-Priority Packets

The previous sections discussed methods for optimizing latency in a single VC system. However, a better solution for some traffic scenarios that require consistently low latency is to use a different VC.

The PEX 8508 does ***not support*** isochronous traffic that requires high-priority packets by way of a switch with a time limit. However, it does provide a high-priority packet path throughout the entire switch if there are two VCs and VC1 is configured with higher priority compared to VC0 in both ingress and egress ports.

VC1 includes independent credit, storage, and scheduling with respect to VC0. However, it shares the wires in and out of the switch. At any point where there can be congestion between the two VCs, VC1 is treated separately and preferentially to VC0. This occurs at the Ingress queues, internal fabric, and Egress queues. For contention, VC1 packets are given priority over VC0 packets.

In this case, VC0 is earmarked for slower, bulk data transfer, and VC1 processes packets with a much shorter latency if there is no over-subscription.

Two conditions are required to make the high-priority path meaningful:

- TC/VC mapping is symmetric across all ports
- All ports configure the Low-Priority Extended VC Count as 0, in the egress **Port VC Capability 1** register (offset 14Ch; default) to give VC1 the higher priority
- Certain TCs map to VC1 [**VC0 Resource Control** and **VC1 Resource Control** registers (offsets 15Ch and 168h, respectively)] and the high-priority TLPs use the TCs that map to VC1
- VC1 is enabled on ingress and egress ports

8.5.4 Cut-Thru

The PEX 8508 includes a limited Cut-Thru feature, which can dramatically reduce latency. Cut-Thru works by waiting for only enough information to route the packet before forwarding it to the destination port. The extra condition exists to ensure that the destination port does not run out of data if it immediately forwards the packet.

When egress port(s) have no stored TLPs, Cut-Thru path(s) are dynamically enabled between the associated ingress and egress port(s). If the **Debug Control** register *Data Path-Based Cut-Thru* bit is cleared (offset 1DCh[14]=0), Cut-Thru is enabled only when all egress ports have no stored TLPs in their path. This setting can be useful in special ordering cases, where Writes from one ingress port are targeting multiple egress ports.

In the case where any packet is destined to an egress port that is the same length as, or narrower than, the ingress port, implement Cut-Thru as soon as the TLP's Header is decoded, regardless of the packet length.

However, in the case where any packet that is going to a destination port is wider than its ingress port, the PEX 8508 must store a significant fraction of the packet before Cut-Thru can start. The ratio is easy to predict:

- x2 to x4 requires approximately 1/2 of the packet to be stored before being forwarded
- x1 to x4 requires 3/4 of the packet to be stored before being forwarded

8.5.5 Smaller Size Packets

If Cut-Thru is not enabled, the PEX 8508 uses a store-and-forward architecture. Without Cut-Thru, a packet must be completely written into the PEX 8508's internal packet RAM before the first byte of the packet can be transmitted out of the egress port. Cut-Thru latency is a function of the packet size; therefore, the smaller the packet size, the shorter the Cut-Thru latency. The bulk of the latency is dictated by the amount of time it takes for the packet to arrive. Narrower ingress ports contain correspondingly higher latency than wider egress ports.

8.5.6 Power Management

Saving power and optimizing latency are typically two conflicting tasks. After a chip enters Power Saving mode, the wakeup time when new Burst packets arrive always contributes to latency. For latency-sensitive applications, it is recommended to use software to turn off the Active State Power Management (ASPM) L0s entrance/exit, as well as L1 entrance/exit.

Chapter 9 Device Layers

9.1 PEX 8508 Data Flow Through

The PEX 8508 uses packets to transfer data between linked devices and to communicate information between all levels in the system, from the Root device to the Endpoint device. Packets are formed in the Transaction and Data Link Layers, to carry the information from the transmitting device to the receiving device. As the transmitted packets flow through the other layers, the packets are extended with additional information necessary to handle packets at those layers. At the receiving side, the reverse process occurs, and packets are transformed from their Physical Layer representation to the Data Link Layer representation, and finally (for TLPs) to the form that is processed by the TL of the receiving device. Figure 9-1 illustrates the conceptual flow of transaction-level packet information through the layers of each port.

The Transaction, Data Link, and Physical Layers and their functions (illustrated in Figure 9-1) are implemented by the PEX 8508, in accordance with *PCI Express Base r1.1* requirements. The details of each layer are described in the following sections.

Figure 9-1. PEX 8508 Data Formation and Flow Through

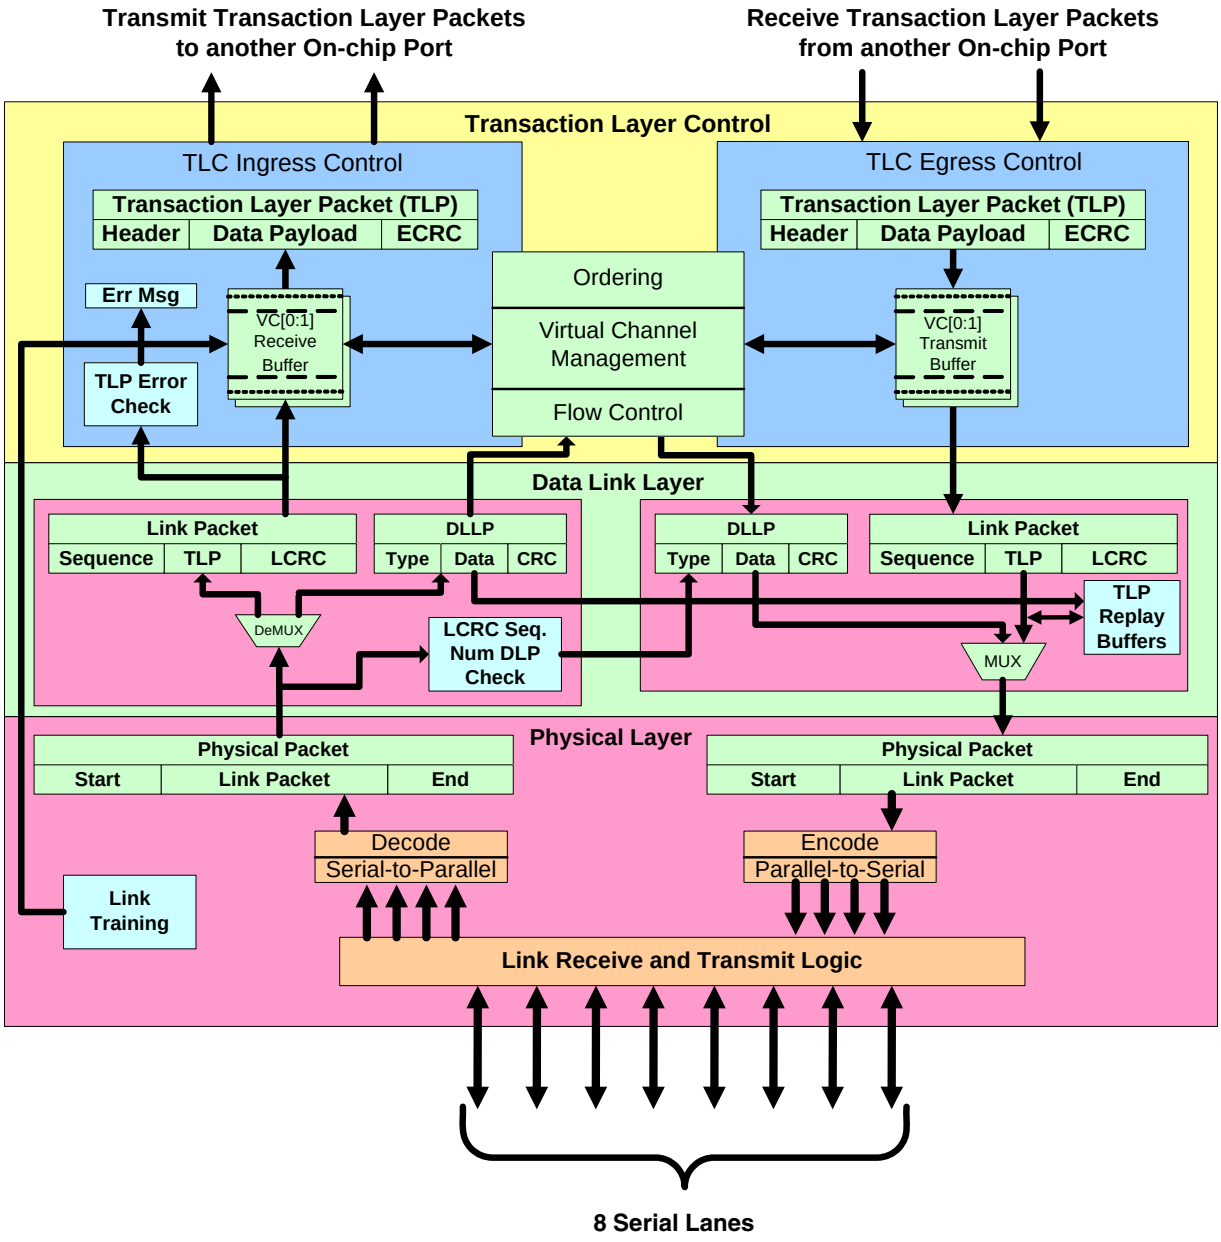

9.2 Physical Layer

The Physical Layer is responsible for converting information received from the DLL into an appropriate serialized format and transmitting it across the PCI Express Link. The Physical Layer includes all circuitry for PCI Express Link interface operation, including:

- Driver and input buffers
- Parallel-to-serial and serial-to-parallel conversion
- PLLs and clock circuitry
- Impedance matching circuitry
- Interface initialization and maintenance functions

9.2.1 Physical Layer Command and Status Registers

The Physical Layer operating conditions are defined in Section 13.13.2, "Physical Layer Registers." The system Host can track the Link operating status and re-configure Link parameters, by way of these registers.

9.2.2 Hardware Link Interface Configuration

The station's Physical Layer can include up to 8 integrated Serializer/De-serializer (SerDes) modules, that are distributed among three SerDes quads (0, 1, and 2) and provide the PCI Express hardware interface lanes. (Refer to Table 9-2, which lists the relationship of the SerDes modules and quads to the 8 lanes.) The SerDes modules also provide all physical communication controls and functions required by the *PCI Express Base r1.1*, as well as the links (clustered into ports) that connect the PEX 8508 to other PCI Express devices.

The number of ports, number of lanes per port, and the SerDes connected to those ports (the numbers within [brackets]) that the PEX 8508 supports is configurable, as defined in Table 9-1. Initial port configuration is determined by Strapped signal balls, serial EEPROM, I^2C download after reset, or auto lane-width negotiation. Link training can reduce the port width by using auto lane-width negotiation.

Table 9-1. PEX 8508 Port Configurations

Port Configuration Register Value[a] (Port 0, Offset 224h[4:0])	Lane Width [Lanes/SerDes]/Port[b,d]				
	Port 0	Port 1	Port 2	Port 3	Port 4
0h	x2 [0-1]	x2 [2-3]	x2 [4, 6]	x2 [8, 10]	–[c]
2h	x4 [0-3]	x4 [4, 6, 8, 10]	–	–	–
3h	x4 [0-3]	x2 [4, 6]	x2 [8, 10]	–	–
4h	x4 [0-3]	x2 [4, 6]	x1 [8]	x1 [10]	–
5h	x4 [0-3]	x1 [4]	x1 [8]	x2 [8, 10]	–
6h	x4 [0-3]	x1 [4]	x2 [6, 8]	x1 [10]	–
8h	x4 [0-3]	x1 [4]	x1 [6]	x1 [8]	x1 [10]
9h	x2 [0-1]	x2 [2-3]	x2 [4, 6]	x1 [8]	x1 [10]

a. All other configurations default to option 0.

b. The lanes are assigned to each enabled port in sequence, as indicated in [brackets].

c. "–" indicates that the port is not enabled for that configuration.

d. Refer to Table 9-2 for an explanation of the physical lane to SerDes quad to SerDes module relationship.

Table 9-2. Physical Lane to SerDes Quad to SerDes Module Relationship

Physical Lanes	SerDes Quad	SerDes Modules
0-3	0	0-3
4, 6	1	4, 6
8, 10	2	8, 10

9.3 Data Link Layer (DLL)

The DLL primary responsibilities include link management and link level data integrity, including error detection. The DLL handles both the transmitter (egress) and receiver (ingress) sides of a link.

The DLL transmitter side accepts a TLP assembled by the TL, prefixes a Sequence Number, and calculates and appends a 32-bit Link CRC (LCRC). The combined TLP is submitted to the Physical Layer for transmission across the link.

The transmitter also sends out DLL packets (DLLPs). ACK and NAK DLLPs are feedback recognizing received TLPs. updateFC DLLPs are feedback providing more credit for more TLPs of particular types and VC. Other DLLPs include Power Management and initFC.

9.3.1 DLLP Ingress

The DLLP Ingress module is responsible for checking the data link level integrity of received TLPs and DLLPs, by verifying the LCRC and Sequence Number (for TLPs). Upon detection of error(s) on TLPs, this layer is responsible for requesting TLP retransmission until information is correctly received, or the link is determined to have failed. LCRC errors detected on a DLLP cause the DLLP to be ignored; another DLLP will be resent eventually.

Qualified TLPs are forwarded to the TL for further processing. Additionally, the Sequence Number is sent to the egress DLL for an eventual ACK response. Non-qualified TLPs result in a NAK response.

Qualified DLLPs are forwarded to the appropriate location, depending upon the DLLP:

- ACK and NAK affect the Replay buffer in the DLL egress module
- updateFC information is sent to the TL egress flow control

Initialization Flow Control (initFC) DLLPs are decoded and used to update the Credit Counters. After the FC Initialization procedure is completed, the initial FC values are communicated to the TL. updateFC DLLPs are immediately communicated to the TL.

The DLLP ingress logic receives TLPs from the physical layer and checks the CRC and Sequence Number. If these data integrity checks pass, the TLP is stripped of the CRC and Sequence Number, aligned to the 160-bit data width, aligned with a time slot, and passed to the TL. When the DLLP egress block receives an ACK/NAK request word from the DLL's TLP ingress block, it converts the word into a properly formatted ACK or NAK DLLP, which is then conveyed to the TL. If a NAK is received, the TLP egress module schedules the specified packet for retransmission and increments the Replay Count for that port.

9.3.2 DLLP Egress

The DLLP egress module is responsible for adding Data Link Layer information to outgoing packets. For TLPs, this means that a Sequence Number is prepended and an LCRC is appended.

The source of TLPs is determined by the Replay buffer. If a replay is needed due to a received NAK, TLPs are pulled from the Replay buffer. Otherwise, TLPs originate from the TL.

The DLL egress module also generates outgoing DLLPs. If the link is busy with TLPs, the DLLPs wait. DLLPs are sent out after a DLLP has waited a sufficient length of time, or there is a break in the TLP traffic stream. A DLLP that is sent out before a TLP is a high-priority DLLP.

TLP acknowledgement (ACK or NAK) DLLPs are sent, based upon received TLPs. Before the DLLP is sent, all received TLPs are collapsed, allowing for one ACK or NAK to be sent in response to multiple TLPs received. NAK DLLPs are always high-priority DLLPs, and are sent ahead of new TLPs. ACK DLLPs are low priority at first. However, after an ACK DLLP has waited the programmed number of ACK Latency Timer clocks [refer to the **ACK Transmission Latency Limit** register (offset 1F8h)], the ACK is promoted to high priority.

updateFC DLLPs are sent indirectly in response to incoming TLPs as well. As incoming TLPs use up the previously advertised credit, additional credit is allocated with a new updateFC. As for ACKs, updateFC DLLPs start out as low priority and can be collapsed to represent multiple TLPs. However, updateFC DLLPs are specific to VC and Type, and therefore can only collapse for TLPs of the same VC and Type. The **INCH FC Update Pending Timer** (Ports 0 through 3 – offset 9F4h[31:0]) and **INCH Port 4 Control** (Port 4 – offset 9F0h[7:0]) registers determine how long an updateFC DLLP must wait before becoming high priority.

The DLL egress also sends Power Management DLLPs and the initial Flow Control (initFC) DLLPs.

9.3.3 Initial FC Credits Advertised

The PEX 8508 has programmable credits to advertise for each port, for each VC and Type. The initially advertised credits are the values of the credit registers, located at offsets A00h to A50h and A60h to A68h. A serial EEPROM can be used to override the default values of these registers.

For further details regarding programming of these registers, refer to the "Program Ingress Credit Threshold Rules" section within Section 8.4.2.1, "Ingress Side."

9.3.4 Packet Arbiter

The Packet Arbiter determines what type of packet to forward to the PHY layer, on a per port basis. The priority algorithm implemented by the Packet Arbiter is discussed in Section 8.4.5.1, "Arbitration between DLLP and TLP," and follows the recommended priority provided in the *PCI Express Base r1.1*, Section 3.5.2.1.

9.4 Transaction Layer (TL)

The Transaction Layer assembles and disassembles TLPs. TLPs are used to communicate transactions, *such as* Read and Write, as well as certain types of events. The TL is also responsible for managing credit-based FC for TLPs. The TL supports the four address spaces defined in Table 9-3.

Message space is added to PCI Express spaces, and it is used to support PCI Express sideband signals *such as* interrupts, Power Management requests, and errors. PCI Express Message transactions are considered *virtual wires* that support *virtual pins*. As virtual wires, Assert and De-assert messages are sent when a triggering event changes the state of the wire.

All Request packets requiring a Response packet are implemented as Split Transactions. Each packet has a unique identifier that enables Response packets to be directed to the correct originator. The packet format supports various forms of addressing, depending upon the transaction type – *Memory*, *I/O*, *Configuration*, or *Message*.

TL functions include:

- Decoding and checking rules for the incoming TLP
- Memory-mapped CSR access
- Checking incoming packets for Malformed or Unsupported packets
- ECRC checking of incoming packets
- Error signaling for incoming packets
- Destination lookup and TC-VC mapping
- TLP packet dispatching
- Write control to the packet RAM and packet link list RAM
- Write control to the Scoreboard Input FIFO
- Adaptive External Credit Control
- Shadow CSR registers for AMCAM/BusNoCAM/TC-VC mapping
- Power Management support
- Hot Plug support
- Message Signaled Interrupt or INT*x* generation
- Ordering
- Ingress and Egress credit management

Table 9-3. Address Spaces Support Differing Transaction Types

Address Space	Transaction Types	Transaction Functions
Configuration	Read/Write	Device configuration or setup
Input/Output		Transfers data from/to an I/O space
Memory		Transfers data from/to a memory location
Message	Baseline/Virtual Wires	General-purpose messages Event signaling (status, interrupts, and so forth)

The hardware functions provided by the PEX 8508 to implement *PCI Express Base r1.1* TL requirements are illustrated in Figure 9-2. The blocks provide a combination of Ingress and Egress control, as well as the data management at each stage in the flow sequence.

Figure 9-2. TL Controller

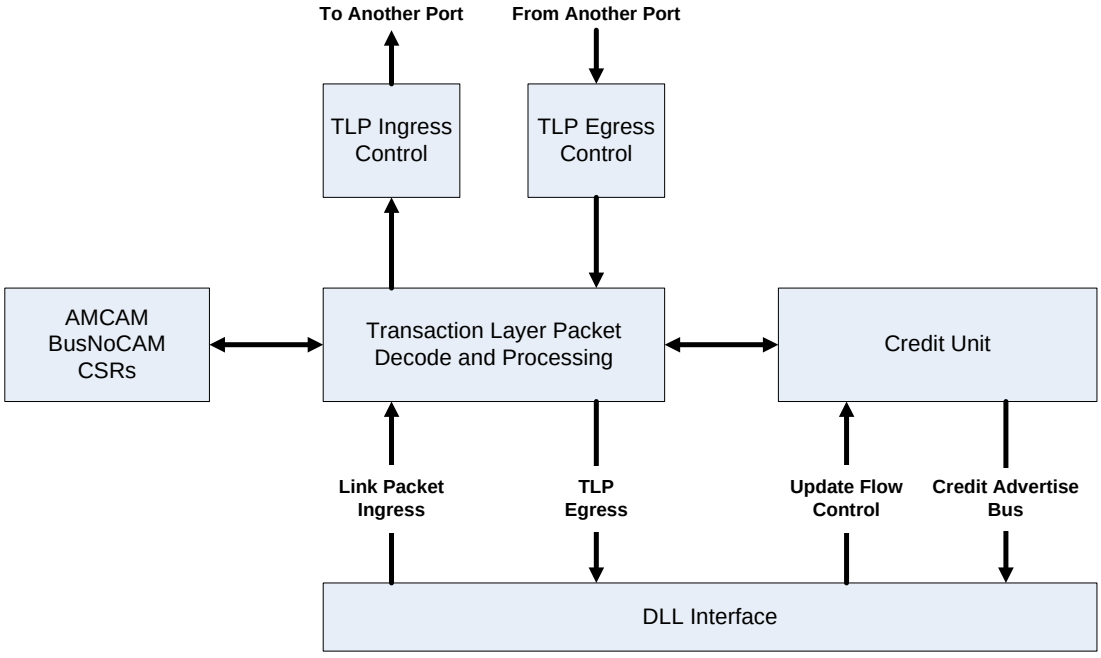

9.4.1 Transaction Layer (TL) Transmit/Egress Protocol

The TL receives information from other ports in the switch and generates outbound requests and Completion TLPs, which it stores in VC buffers. This layer assembles TLPs whose contents consist of Identification Headers, Data Payloads, and ECRC checking codes. These fields are defined in the *PCI Express Base r1.1*, Section 2.2.

9.4.1.1 Headers

The Headers contain three or four DWords, which can include the following:

- Address/Routing – 32 or 64 bits
- TLP Type
- Transfer Size
  - Write requests = Total outgoing DWords
  - Read requests = Requested DWords from Completer
- Requester ID or Completer ID
- Tag – Used to identify a Completion TLP
- Traffic Class
- Byte Enables
- Completion status
- Digest – One bit indicating ECRC presence
- Attributes

9.4.1.2 Data Payloads

The Data Payloads are variable length with a maximum of 128 bytes or 256 bytes, as defined by the *Maximum Payload Size* bits. Read requests do not include a Data Payload.

9.4.1.3 End-to-End CRC (ECRC)

ECRC is an optional 32-bit field appended to the end of the outgoing packet. ECRC is calculated over the entire packet, starting with the Header and including the Data Payload, except for the *EP* bit and bit 0 of the *Type* field, which are always considered to be a value of 1 for ECRC calculations. The *ECRC* field is transmitted, unchanged, as it moves through the fabric to the completer device. The PEX 8508 checks the ECRC on all incoming TLPs, if enabled, and can optionally report detected errors. [When the ECRC is detected, the **Uncorrectable Error Status** register *ECRC Error Status* bit (offset FB8h[19]) can be used to log ECRC errors.] Additionally, the PEX 8508 can optionally append ECRC to the end of internally generated TLPs, *such as* Interrupt and Error messages, if enabled.

9.4.2 TL Receive/Ingress Protocol

The TL Receive/Ingress logic collects and stores inbound TLP traffic in the VC buffers. The incoming data is checked for CRC errors, valid type field, and to ensure it matches the *Transfer Size* field in the Header. The *LCRC* field is stripped and the remaining Header and Data Payload are forwarded to the Source Scheduler, to be routed to the outgoing switch port. When CRC errors are detected, the packet is discarded.

9.4.3 Flow Control (FC) Protocol

The PEX 8508 implements FC protocol that ensures it does not transmit a TLP over a link to a remote receiver unless the receiving device has sufficient VC Buffer space to accommodate the packet. This FC is automatically managed by the hardware and is transparent to software. Software is used only to enable additional buffers to supplement the initial default buffer assignment.

The initial default FC DLLP buffers, which are enabled after link training, allow TLP traffic immediately following the link training. The Configuration transactions are the first transactions to use the default VC buffers to set up the initial device operating modes and capabilities.

The TL Egress Credit Unit transmits DLL packets, called *FC packets*, that update the FC to the transmitter device on a periodic basis. The updated DLLPs contain FC credit information that updates the transmitter regarding the amount of available buffer space in the receiver VC buffer. The transmitter tracks this information and transmits TLPs only when it perceives that the remote receiver contains sufficient Buffer space to accept the transmitted TLPs.

9.4.4 Virtual Channels and Traffic Classes

The PEX 8508 supports up to two Virtual Channels (VC0 and VC1), and 8 Traffic Classes (TC[7:0]). VC0 and TC0 are required by the *PCI Express Base r1.1*, and are configured at device start-up. VC1, the second Virtual Channel, is supported but not enabled for use. To use VC1, set the **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]). VC1 can be disabled by clearing the **Port VC Capability 1** register *Low-Priority Extended VC Count* and *Extended VC Count* bits (offset 14Ch[4, 0]=00b). For four- or five-port implementations, VC support is limited to two VCs (VC0 and VC1) for up to three ports and one VC (VC0) for the fourth and fifth port.

9.4.4.1 PEX 8508 PCI Express Virtual Channel Support

The PEX 8508 supports PCI Express features required for PCI Express Switch implementation. Table 9-4 provides a partial listing of supported and non-supported features and the register bits used for their configuration. For a complete listing, refer to Section 13.12, "Virtual Channel Extended Capability Registers."

Note: *The PEX 8508 supports up to two Virtual Channels (VCs) – VC0 and VC1. Ports 0, 1, and 2 support VC0 and VC1. Ports 3 and 4 support only VC0.*

Table 9-4. PCI Express Virtual Channel Capabilities Support

Register		Description	Supported	
Offset	Bit(s)		Yes	No
14Ch	0	**Extended VC Count** 0 = PEX 8508 port supports only the default Virtual Channel 0 (VC0) 1 = PEX 8508 port supports one extended virtual channel, Virtual Channel 1 (VC1) ***Note:*** *Value of 1 is **reserved** for Ports 3 and 4.*	✓	
	4	**Low-Priority Extended VC Count** For strict priority arbitration, this bit indicates the number of extended virtual channels (those in addition to the default VC0) that belong to the Low-Priority Virtual Channel group for the PEX 8508 port. PEX 8508 Serial EEPROM register initialization capability is used to change this bit to 1 to also set VC1 to the Low-Priority Virtual Channel group. 0 = For this PEX 8508 port, only the default VC0 belongs to the Low-priority Virtual Channel group 1 = For this PEX 8508 port, VC0 and VC1 belong to the Low-priority Virtual Channel group ***Note:*** *Value of 1 is **reserved** for Ports 3 and 4.*	✓	
	9:8	**Reference Clock**		✓
	11:10	**Port Arbitration Table Entry Size**		✓
150h	0	**VC Arbitration Capability** Bit 0 value of 1 indicates Round-Robin (Hardware-Fixed) Arbitration scheme is supported.	✓	
154h	1	**VC Arbitration Select** ***Not used, because the PEX 8508 supports only Hardware-Fixed arbitration. It is recommended that this bit remain programmed to 0.*** Selects the VC arbitration type for the corresponding PEX 8508 port. Indicates the bit number in the **Port VC Capability 2** register *VC Arbitration Capability* field that corresponds to the arbitration type. 0 = bit 0; Round-Robin (Hardware-Fixed) arbitration scheme 1 = ***Reserved*** Select only an arbitration type that corresponds to a bit set in the *VC Arbitration Capability* field. Cannot be modified when more than one lower priority Virtual Channel group VC is enabled.		✓

Table 9-4. PCI Express Virtual Channel Capabilities Support *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
168h[a]	0	**TC/VC1 Map[0]**		✓
	7:1	**TC/VC1 Map[7:1]** Defines Traffic Classes [7:1], respectively, and indicates which TCs are mapped into Virtual Channel 1. By default, Traffic Classes [7:1] are mapped to VC0. ***Note:*** *Bit 1 must be set to 1 for VC0.*	✓	
	16	**Load Port Arbitration Table**		✓
	19:17	**Port Arbitration Select**		✓
	24	**VC1 ID** Defines the ID code for the corresponding PEX 8508 port Virtual Channel 1 (1 is the only supported value).	✓	
	31	**VC1 Enable** 0 = Disables the corresponding PEX 8508 port Virtual Channel 1 1 = Enables the corresponding PEX 8508 port Virtual Channel 1	✓	
16Ch[a]	16	**Port Arbitration Table Status**		✓
	17	**VC1 Negotiation Pending** 0 = VC1 negotiation completed 1 = VC1 initialization or disabling is pending for the corresponding PEX 8508 port	✓	

*a. Register offsets 168h and 16Ch are **reserved** for Ports 3 and 4.*

Chapter 10 Virtual Channels and Port Arbitration

10.1 Quality of Service (QoS) Support

QoS is a feature offered by PCI Express that was not available before in PCI. *QoS* refers to being able to guarantee preferential treatment for a particular application or for traffic coming from a particular source. The PEX 8508 provides two methods for acquiring QoS:

- Virtual Channel (VC) Support
- Port Arbitration

10.2 Virtual Channel (VC) Support

The PEX 8508 supports up to two Virtual Channels (VCs) that act as two completely independent paths through the switch. Although the two VCs share the same physical wire links, the PEX 8508 has separate queues and separate credit mechanisms for each VC. At the final egress stage, the PEX 8508 has two methods for determining which TLP from which VC to send on the wire.

These features are described in the following sections.

10.2.1 Traffic Class (TC) to VC Mapping

The TC/VC mapping is port-independent and must be configured and enabled for each port. The default configuration for the PEX 8508 ports sets all TC[7:0] bits to VC0, as defined in the **VC0 Resource Control** register *TC/VC0 Map* bits. Any of the TC IDs (except TC0) can be mapped to VC1 by removing them from the *TC/VC0 Map* bits and adding them to the **VC1 Resource Control** register *TC/VC1 Map[7:1]* field. If any TCs are mapped to VC1, VC1 must be enabled by setting the **VC1 Resource Control** register *VC1 Enable* bit. Also, because the default VC1 credit is minimal, the user might want to increase the VC1 credit on a port-by-port basis. (Refer to Section 8.4.2.1, "Ingress Side," regarding ingress credit programming.)

10.2.2 Fixed-Priority Arbitration

The *PCI Express Base r1.1* defines the default hardware-fixed strict priority scheme based upon the inherent priority of the VC IDs, where VC1 has higher priority than VC0. The VC IDs alone determine the sequence in which transactions are sent. For this arbitration scheme, which is the default, the **Port VC Capability 1** register *Low-Priority Extended VC Count* bit is cleared (offset 14Ch[4]=0).

10.2.3 Round-Robin Arbitration

If the **Port VC Capability 1** register *Low-Priority Extended VC Count* bit is set (offset 14Ch[4]=1), VC0 and VC1 share the same low-priority pool. As a result, the PEX 8508 alternates selecting VC0 and VC1 in a fair Round-Robin manner.

10.3 Port Arbitration

The PEX 8508 switch implements queues at ingress and egress ports. There is arbitration capability for both Ingress and Egress queues.

By default, the PEX 8508 attempts to send every packet on the Ingress queue to the Egress queue, until the egress RAM fills. That results in optimum use of the PEX 8508 on chip buffering; however, it might create a situation in which one high-priority port takes too much latency.

The amount of queueing in the egress side is programmable per VC and Type. The **ITCH VC&T Threshold *x*** registers (refer to Section 13.13.14, "Internal Credit Handler (ITCH) VC&T Threshold Registers," offsets C00h to C08h) set a maximum for the Egress queue depth for each VC&T. By reducing the thresholds, the Egress queue size is constrained. A smaller Egress queue means that newly arriving packets from a high-priority source do not have to wait as long for any earlier arriving packets from another source, because the number of earlier arriving packets is reduced.

10.3.1 Ingress Arbitration

The PEX 8508 supports hardware-fixed Round-Robin arbitration on the ingress ports. Round-Robin servicing is performed between the five ports. When a port has no packets pending in its queue, that port is skipped, and the next port with packets waiting is serviced.

If any ingress ports are targeting a congested Egress queue (Egress queue depth is greater than its programmable threshold), that ingress port removes itself from the arbitration pool.

An Ingress queue can enable Device-Specific Relaxed Ordering, by setting any bit within a **Device-Specific Relaxed Ordering Mode** register *Enable PLX Relaxed Ordering* field to a value of 1 for the selected port:

- Port 0 – offset BFCh[7:0]
- Port 1 – offset BFCh[15:8]
- Port 2 – offset BFCh[23:16]
- Port 3 – offset BFCh[31:24]
- Port 4 – offset BE4h[7:0]

With Relaxed Ordering, packets blocked by one destination do not block later packets to another destination.

10.3.2 Egress Queue Handling

The packets in the Egress queue of each port are serviced on a first-in, first-out basis (without violating PCI Ordering rules). The size of the Egress queue in each port is programmed, using the threshold registers. The **ITCH VC&T Threshold *x*** registers allow for a maximum queueing of each VC and Type in the egress port. If that maximum is crossed, the destination backpressures any source attempting to Write to the destination queue. After the programmed minimum is crossed, the backpressure releases and the source can forward to that destination again.

Chapter 11 Hot Plug Support

11.1 Hot Plug Purpose and Capability

Hot Plug capability allows board insertion and extraction from a running system without adversely affecting the system. Boards are typically inserted or extracted to repair faulty boards or re-configure the system without system down time. Hot Plug capability allows systems to isolate faulty boards in the event of a failure. The PEX 8508 includes one Hot Plug Controller per downstream port.

11.1.1 Hot Plug Controller Capabilities

- Insertion and removal of PCI Express boards, without removing system power
- Board-present and MRL (Manually operated Retention Latch) Sensor signals supported
- Power Indicator and Attention Indicator Output signals controlled
- Attention Button monitored
- Power fault detection and Faulty board isolation
- Power switch for controlling downstream device power
- Generates PME (Power Management Event) for Hot Plug events in sleeping systems (D3hot)
- Presence detect is accomplished through either an in-band SerDes receiver detect mechanism or by using the HP_PRSNT*x*# signal
- Hot Plug interrupts can be sent in-band using INT*x* or MSI messages

11.1.2 Hot Plug Port External Signals

The PEX 8508's Hot Plug Controllers include nine Hot Plug signals for each PCI Express port (5 ports x 9 signals/port = 45 total signals), defined in Table 11-1. (Refer to Table 3-4, "Hot Plug Signals – 45 Balls," for signal-to-ball mapping.)

Table 11-1. Hot Plug Signals

Signal Name	Type	Description
HP_ATNLED*x*#	O	**Hot Plug Attention LED Output (5 Balls)** Active-Low Slot Control Logic output used to drive the Attention Indicator. Output is set Low to turn On the LED. Enabled when the **Slot Capabilities** register *Attention Indicator Present* bit is set (offset 7Ch[3]=1) and controlled by the **Slot Control** register *Attention Indicator Control* field (offset 80h[7:6]). When software writes to the *Attention Indicator Control* field, a Command Completed interrupt can be generated to notify the Host that the command has been executed. Software must use a Byte or Word Write (and not a Dword Write) to control HP_ATNLED*x*#. When the following conditions exist: • **Slot Capabilities** register *Attention Indicator Present* bit is set (offset 7Ch[3]=1), and • **Slot Control** register *Command Completed Interrupt Enable* bit is not masked (offset 80h[4]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated to the Host. An external current-limiting resistor is required.
HP_BUTTON*x*#	I PU	**Hot Plug Attention Button Input (5 Balls)** Active-Low Slot Control Logic input, directly connected to the Attention Button, with input assertion status latched in the **Slot Status** register *Attention Button Pressed* field (offset 80h[16]). Enabled when the **Slot Capabilities** register *Attention Button Present* bit is set (offset 7Ch[0]=1). When the following conditions exist: • HP_BUTTON*x*# is not masked (**Slot Control** register *Attention Button Pressed Enable* bit (offset 80h[0]=1), and • **Slot Capabilities** register *Hot Plug Capable* bit is set (offset 7Ch[6]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated, to notify the Host of intended board insertion or removal. ***Note:*** *HP_BUTTONx# is internally de-bounced, but must remain stable for at least 10 ms.*
HP_CLKEN*x*#	O	**Reference Clock Enable Output (5 Balls)** Active-Low output that, when enabled, allows external REFCLK to be provided to the slot. Enabled when the **Slot Capabilities** register *Power Controller Present* bit is set (offset 7Ch[1]=1), and controlled by the **Slot Control** register *Power Controller Control* bit (offset 80h[10]). The time delay from HP_PWREN*x*# output assertion to HP_CLKEN*x*# output assertion is programmable (through serial EEPROM load) from 16 ms (default) to 128 ms, in the *HPC Tpepv Delay* field (offset 1E0h[4:3]).

Table 11-1. Hot Plug Signals *(Cont.)*

Signal Name	Type	Description
HP_MRL*x*#	I PU	**Hot Plug Manually Operated Retention Latch Sensor Input (5 Balls)** Active-Low input that triggers Slot Control Logic. Directly connected to an optional MRL Sensor that is logic High when the latch is not closed. HP_MRL*x*# input assertion enables Hot Plug output sequencing to turn On the slot's power (HP_PWREN*x*# and HP_PWRLED*x*#) and clock (HP_CLKEN*x*#), and de-assert Reset (HP_PERST*x*#) after reset or under software control. Enabled when the **Slot Capabilities** register *MRL Sensor Present* bit is set (offset 7Ch[2]=1). A change in the HP_MRL*x*# Input signal state is latched in the **Slot Status** register *MRL Sensor Changed* bit (offset 80h[18]), and the state change can assert an interrupt to notify the Host of a change in the MRL Sensor state. When the following conditions exist: • HP_MRL*x*# is not masked (**Slot Control** register *MRL Sensor Changed Enable* bit, offset 80h[2]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated. If the associated downstream port connects to a PCI Express board slot that does not implement an MRL Sensor, HP_MRL*x*# is normally connected to HP_PRSNT*x*# and a pull-up resistor, with the common node connected to the PRSNT2# signal(s) at the slot. If the associated downstream port instead connects directly to a device (in which case Hot Plug is not used), pull HP_MRL*x*# Low. ***Note:*** *HP_MRLx# is internally de-bounced, but must remain stable for at least 10 ms. HP_MRLx#, if enabled, is not de-bounced when sampled immediately after reset.*
HP_PERST*x*#	O	**Active-Low Reset Output (5 Balls)** Active-Low Hot Plug output used to reset the slot. Controlled by the **Slot Control** register *Power Controller Control* bit (offset 80h[10]).
HP_PRSNT*x*#	I PU	**Combination of Hot Plug PRSNT1# and PRSNT2# Input (5 Balls)** Active-Low input connected to the slot's PRSNT2# signal, which on the add-in board connects to the slot's PRSNT1# signal, which is normally grounded on the PRSNT2# signal at the motherboard slot. A change in the HP_PRSNT*x*# Input signal state is latched in the **Slot Status** register *Presence Detect Changed* bit (offset 80h[19]), and the state change can assert an interrupt to notify the Host of board presence or absence. When the following conditions exist: • HP_PRSNT*x*# is not masked (**Slot Control** register *Presence Detect Changed Enable* bit (offset 80h[3]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated. ***Note:*** *HP_PRSNTx# is internally de-bounced, but must remain stable for at least 10 ms.*

Table 11-1. Hot Plug Signals *(Cont.)*

Signal Name	Type	Description
HP_PWREN*x*#	O	**Active-Low Hot Plug Power Enable Output (5 Balls)** Active-Low Slot Control Logic output that controls the slot power state. When this signal is Low, power is enabled to the slot. Enabled when the **Slot Capabilities** register *Power Controller Present* bit is set (offset 7Ch[1]=1). When software turns the slot's Power Controller On or Off (**Slot Control** register *Power Controller Control* bit, offset 80h[10]), a Command Completed interrupt can be generated to notify the Host that the command has been executed. When the following conditions exist: • **Slot Control** register *Command Completed Interrupt Enable* bit is not masked (offset 80h[4]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated to the Host. When HP_MRL*x*# is enabled [**Slot Capabilities** register *MRL Sensor Present* bit is set (offset 7Ch[2]=1)], HP_MRL*x*# input assertion enables Hot Plug output sequencing to turn On the slot's power, by asserting HP_PWREN*x*# after reset or under software control.
HP_PWRFLT*x*#	I PU	**Hot Plug Power Fault Input (5 Balls)** Active-Low input that indicates the slot's external Power Controller detected a power fault on one or more supply rails. Enabled when the **Slot Capabilities** register *Power Controller Present* bit is set (offset 7Ch[1]=1), and input assertion status is latched in the **Slot Status** register *Power Fault Detected* bit (offset 80h[17]). When the following conditions exist: • HP_PWRFLT*x*# is not masked (**Slot Control** register *Power Fault Detector Enable* bit (offset 80h[1]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated, to notify the Host of a power fault. ***Note:*** *If HP_PWRENx# and HP_CLKENx# are not used, HP_PWRFLTx# can be used as a general-purpose input with status reflected in the **Slot Status** register Power Fault Detected (offset 80h[17]), provided the **Slot Capabilities** register Power Controller Present bit is set (offset 7Ch[1]=1).*
HP_PWRLED*x*#	O	**Hot Plug Power LED Output (5 Balls)** Active-Low Slot Control Logic output used to drive the Power Indicator. This output is set Low to turn On the LED. Enabled when the **Slot Capabilities** register *Power Indicator Present* bit is set (offset 7Ch[4]=1), and controlled by the **Slot Status** register *Power Indicator Control* field (offset 80h[9:8]). When software writes to the *Power Indicator Control* field, a Command Completed interrupt can be generated to notify the Host that the command has been executed. Software must use a Byte or Word Write (and not a Dword Write) to control HP_PWRLED*x*#. When the following conditions exist: • **Slot Capabilities** register *Power Indicator Present* bit is set (offset 7Ch[4]=1), and • **Slot Control** register *Command Completed Interrupt Enable* bit is not masked (offset 80h[4]=1), and • **Slot Control** register *Hot Plug Interrupt Enable* bit is set (offset 80h[5]=1), an interrupt (MSI, or INT*x* message, both mutually exclusive) can be generated to the Host. An external current-limiting resistor is required.

Note: *If Hot Plug outputs (including HP_PERSTx#) are used and HP_MRLx# input is not used, pull HP_MRLx# input Low so that Hot Plug outputs (including HP_PERSTx#) will properly sequence if the serial EEPROM is blank or missing. Default register values enable HP_MRLx#, which must then be asserted to cause Hot Plug outputs to toggle (for example, to de-assert HP_PERSTx# and assert HP_PWRLEDx#).*

11.1.3 Hot Plug Output Signal States for Disabled Hot Plug Slots

When a Hot Plug slot is disabled, the Hot Plug Output balls for that port are in the logic states defined in Table 11-2.

Table 11-2. Hot Plug Outputs for Disabled Hot Plug Slot

Output Signal	Logic	Comments
HP_ATNLED*x*#	High	Attention LED is turned Off
HP_CLKEN*x*#	High	Reference Clock not driven to the slot
HP_PERST*x*#	Low	Slot remains in reset
HP_PWREN*x*#	High	Power Controller is turned Off
HP_PWRLED*x*#	High	Power LED is turned Off

11.2 PCI Express Capability Registers for Hot Plug

The Hot Plug Configuration, Capability, Command, Status, and Events are described in Section 13.9, "PCI Express Capability Registers." The applicable registers are as follows:

- **Slot Capabilities** (offset 7Ch)
- **Slot Status and Control** (offset 80h)

Note: *Hot Plug **Slot Status** and other Hot Plug control-related registers are "don't care" for the NT Port Virtual Interface, and should not be modified by the user.*

11.3 Hot Plug Interrupts

Each Hot Plug Controller supports Hot Plug interrupt generation on the following events:

- Attention Button Pressed
- Power Fault Detected
- MRL Sensor Changed
- Presence Detect Changed
- Command Completed
- Data Link Layer State Changed

Hot Plug interrupts can be signaled by in-band INT*x* or MSI messages. Only one interrupt mechanism can be selected, and all Hot Plug ports must use the same mechanism.

INT*x* interrupts are enabled if:

- INTA messages are enabled (**Command** register *Interrupt Disable* bit, offset 04h[10]=0) and,
- MSI is disabled (**Message Control** register *MSI Enable* bit, offset 48h[16]=0)

MSI interrupts are enabled if:

- INTA messages are disabled (**Command** register *Interrupt Disable* bit, offset 04h[10]=1) and,
- MSI is enabled (**Message Control** register *MSI Enable* bit, offset 48h[16]=1)

Depending upon the downstream port's power state, a Hot Plug event can generate a system interrupt or PME. When a PEX 8508 downstream port is in the D0 power state, Hot Plug events generate a system interrupt; when not in the D0 state, a PME interrupt message is generated by Hot Plug events. The **Slot Status** register *Command Completed* bit (offset 80h[20]) does not generate a PME interrupt message. When the system is in Sleep mode, Hot Plug operation uses PME logic to wake up the system.

11.4 Hot Plug Controller Power-Up/Down Sequence

If a Transparent downstream port is enabled, the port's Hot Plug Controller can power-up or power-down the slot. This section describes how this process occurs.

11.4.1 Slot Power-Up Sequence

If a downstream port is connected to a slot, that port's Hot Plug Controller can power up the slot, with or without an external serial EEPROM. Hot Plug Controller sequencing is determined by the states of the following bits:

- **Slot Capabilities** register *Power Controller Present* bit (offset 7Ch[1])
- **Slot Capabilities** register *MRL Sensor Present* bit (offset 7Ch[2])

  (MRL is Manually operated Retention Latch)
- **Slot Control** register *Power Controller Control* bit (offset 80h[10])

and the HP_MRL*x*# input state, if the *MRL Sensor Present* bit is set to 1. Hot Plug-configurable features are programmable only by the serial EEPROM.

11.4.1.1 Configuring Hot Plug Controller Slot Power-Up Sequence Features with Serial EEPROM

An external serial EEPROM can be used to configure the Hot Plug Controller and Hot Plug outputs. Features can be changed by using the registers defined in Table 11-3. The Hot Plug Controller outputs remain in the default state described in Table 11-2, before the serial EEPROM image is loaded into the device.

After the serial EEPROM image is loaded, the Hot Plug Controller starts a power-up sequence on each slot that has the **Slot Capabilities** register *Power Controller Present* bit set (offset 7Ch[1]=1) and the **Slot Control** register *Power Controller Control* bit cleared (offset 80h[10]=0).

Table 11-3. Configuring Power-Up Sequence Features with Serial EEPROM

Register Bit	Hot Plug Controller and Hot Plug Output Signal Configurable Features
Power Controller Present (**Slot Capabilities** register, offset 7Ch[1])	The *Power Controller Present* bit enables or disables the Hot Plug Controller on the PEX 8508 downstream ports. If the *Power Controller Present* bit is cleared to 0, the Hot Plug Controller is disabled for that slot and a power-up sequence is not executed. The slot remains in the disabled state, as defined in Table 11-2. If the *Power Controller Present* bit is enabled (set to 1), the Hot Plug Controller powers up the slot when the MRL is closed and the **Slot Control** register *Power Controller Control* bit is cleared (offset 80h[10]=0). Otherwise, if the *MRL Sensor Present* bit is disabled (cleared to 0), the MRL's position has no effect on powering up the slot.
MRL Sensor Present (**Slot Capabilities** register, offset 7Ch[2])	When enabled (set to 1), the PEX 8508 senses whether the MRL is open or closed for a slot. If this bit is set to 1, the MRL should be Low for power-on for that slot. If this bit is cleared to 0, the MRL position is "don't care" for that slot.
HPC Tpepv Delay (**Power Management Hot Plug User Configuration** register, offset 1E0h[4:3])	This field controls the delay from when HP_PWREN*x*# is asserted Low, to when power is valid at a slot. (Refer to Section 11.4.1.2.) This register is Read-Only and can be set by serial EEPROM. 00b = 16 ms (default) 01b = 32 ms 10b = 64 ms 11b = 128 ms
HPC Tpvperl Delay (**Power Management Hot Plug User Configuration** register, offset 1E0h[6])	This bit controls the delay from when Power is valid at the slot, to when HP_PERST*x*# is de-asserted High. (Refer to Section 11.4.1.2.) 0 = 20 ms 1 = 100 ms (default)
Attention Indicator Present (**Slot Capabilities** register, offset 7Ch[3])	When set to 1, this bit controls whether the HP_ATNLED*x*# output for the slot drives out Active-Low. Otherwise, this output is not functional on the slot.
Power Indicator Present (**Slot Capabilities** register, offset 7Ch[4])	When set to 1, this bit controls whether the HP_PWRLED*x*# output for the slot drives out Active-Low. Otherwise, this output is not functional on the slot.

11.4.1.2 Slot Power-Up Sequencing when *Power Controller Present* Bit Is Set

By default, the *Power Controller Present*, *MRL Sensor Present*, and *Power Controller Control* (when the MRL is open) bits are set to 1. When the serial EEPROM is not present, present but blank, or programmed with default register values, the Hot Plug Controller is initially powered up, the **PCI Express Capabilities** register *Slot Implemented* bit is set (offset 68h[24])=1), and the PEX 8508 is in the following state:

1. Hot Plug Controller is enabled for all downstream ports.
2. All slots are enabled to be powered up.
3. Attention LED (HP_ATNLED*x*#) and Power LED (HP_PWRLED*x*#) are High on the slot chassis.

Immediately after the PEX 8508 exits Reset (PEX_PERST# input goes High), if a downstream port's *MRL Sensor Present* bit is set to 1 (default), the HP_MRL*x*# input for that slot is sampled. If HP_MRL*x*# input is enabled and asserted (value of 0), the device clears the *Power Controller Control* bit to 0, to enable slot power-up. If the *Power Controller Control* bit is not cleared, either by initially enabling it (default) and asserting HP_MRL*x*#, or by programming both the *MRL Sensor Present* and *Power Controller Control* bit values to 0 in the serial EEPROM, the downstream slot is not powered up and remains in the disabled state, as defined in Table 11-2 and illustrated in Figure 11-3.

If a slot's *Power Controller Present* bit is set to 1, and the *Power Controller Control* bit is cleared to 0 (either by initially enabling and asserting HP_MRL*x*#, or by programming the *MRL Sensor Present* and *Power Controller Control* bit values to 0 in the serial EEPROM), the slot starts power-up sequencing with HP_PWREN*x*# and HP_PWRLED*x*# assertion, following PEX_PERST# input de-assertion:

- If the serial EEPROM is not present, HP_PWREN*x*# and HP_PWRLED*x*# are asserted approximately 6.1 ms after PEX_PERST# input is de-asserted
- If the serial EEPROM is present, HP_PWREN*x*# and HP_PWRLED*x*# are asserted approximately 18.7 ms after PEX_PERST# input is de-asserted

The power-up sequence is as follows:

1. The Hot Plug Controller drives HP_PWRLED*x*# Low, to turn On the Power Indicator, and drives HP_PWREN*x*# Low to turn On the external Power Controller.

2. After the programmable T_{pepv} time delay following HP_PWREN*x*# assertion, power to the slot is valid and the Hot Plug Controller drives HP_CLKEN*x*# Low. Based upon the STRAP_SSC_XING_ENA ball state, if SSC isolation is enabled, the Hot Plug Controller turns On the PEX_REFCLK_CFCn/p signal to the slot. Otherwise, if SSC isolation is disabled, the Hot Plug Controller turns On the PEX_REFCLKn/p signal to the slot. The T_{pepv} time delay is specified by setting the **Power Management Hot Plug User Configuration** register *HPC Tpepv Delay* field (offset 1E0h[4:3]) to a non-zero value. By default, this field is cleared to 00b, indicating a 16-ms time delay from the time HP_PWREN*x*# goes Low to power becoming valid at the slot.

3. After the programmable T_{pvperl} time delay following HP_CLKEN*x*# assertion, the Hot Plug Controller de-asserts HP_PERST*x*# to release slot reset. The T_{pvperl} time delay is specified in the **Power Management Hot Plug User Configuration** register *HPC Tpvperl Delay* bit (offset 1E0h[6]). By default, this bit is set to 1, indicating a 100-ms delay.

   With this default delay, if the serial EEPROM is not present, HP_PERST*x*# output is de-asserted approximately 122 ms after PEX_PERST# input is de-asserted. However, if the serial EEPROM is present, HP_PERST*x*# output is de-asserted approximately 135 ms after PEX_PERST# input is de-asserted.

   Because the *PCI Express Base r1.1* allows the Host to start Configuration accesses 100 ms after the Root Complex de-asserts its PERST# output, it is recommended that a programmed serial EEPROM be used to clear the *HPC Tpvperl Delay* bit to 0, to reduce the T_{pvperl} time delay to 20 ms, so that HP_PERST*x*# output is de-asserted approximately 55 ms after PEX_PERST# input is de-asserted.

Figure 11-1 illustrates the timing sequence with the *Power Controller Present* bit (offset 7Ch[1]) set to 1. This timing sequence occurs at system power-up, or when a slot is being powered up by the user.

If HP_MRL*x*# is enabled but not asserted to power-up the slot immediately after reset, HP_MRL*x*# can be asserted at runtime to start the slot power-up sequence, provided the *MRL Sensor Present* and *Power Controller Present* bits are set (offset 7Ch[2:1]=11b, either by default values when the serial EEPROM is not present or blank, or by programming the serial EEPROM to set these bits), and the *Power Controller Control* bit is cleared (offset 80h[10]=0, either by the programmed serial EEPROM or by software).

Power-up sequencing at runtime is controlled by software clearing the *Power Controller Control* bit in response to an interrupt caused by HP_MRL*x*# input assertion [if an MRL Sensor is present, and the **Slot Control** register *Hot Plug Interrupt Enable and MRL Sensor Changed Enable* bits are set (offset 80h[5, 2]=11b), and/or by the user pressing the Attention Button, if enabled [**Slot Control** register *Hot Plug Interrupt Enable and Attention Button Pressed Enable* bits must be set (offset 80h[5, 0]=11b)].

HP_MRL*x*# and HP_BUTTON*x*# assertion and de-assertion at runtime are not latched until the 10-ms de-bounce ensures that the state change is stable.

Figure 11-1. Slot Power-Up Timing when *Power Controller Present* Bit Is Set

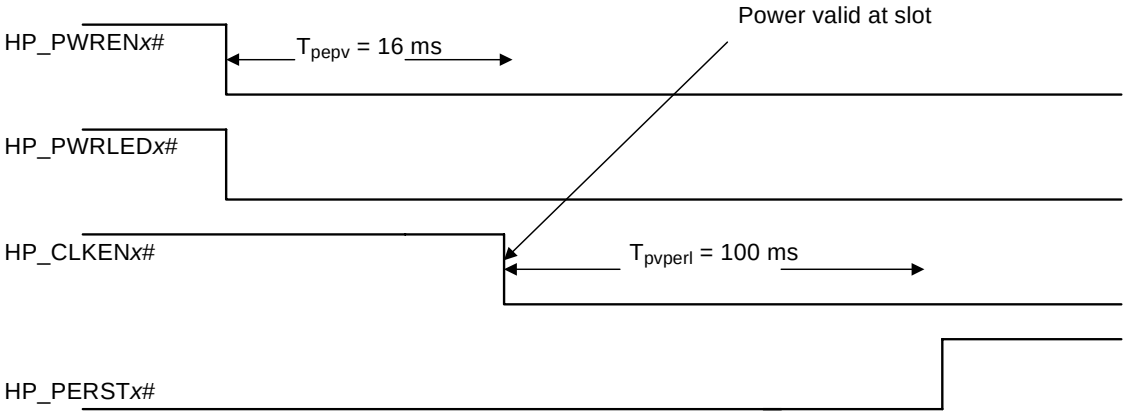

Note: *HP_PWRLEDx# is not asserted if the serial EEPROM or I^2C interface clears the Power Indicator Present bit (offset 7Ch[4]) to 0.*

11.4.1.3 HP_PERST*x*# (Reset) and HP_PWRLED*x*# Output Power-Up Sequencing when *Power Controller Present* Bit Is Clear

The HP_PERST*x*# and HP_PWRLED*x*# outputs can be used without enabling the Hot Plug Power Controller (HP_PWREN*x*# and HP_CLKEN*x*# outputs and HP_PWRFLT*x*# input). *For example*, HP_PERST*x*# can be used to reset an on-board downstream device.

If the *Power Controller Present* (offset 7Ch[1]) and *Power Controller Control* (offset 80h[10]) bits are cleared to 0 by the serial EEPROM, HP_PERST*x*# is de-asserted (High) and HP_PWRLED*x*# is asserted (Low), after the Root Complex PERST# input is de-asserted, as illustrated in Figure 11-2. However, HP_PWRLED*x*# is not asserted if the serial EEPROM also cleared the *Power Indicator Present* bit (offset 7Ch[4]) to 0.

If the serial EEPROM is initially blank, causing register default values to be loaded, HP_PERST*x*# is asserted and HP_PWRLED*x*# is not asserted unless HP_MRL*x*# is Low. Therefore, if the HP_PERST*x*# and/or HP_PWRLED*x*# outputs are used [and a Manually operated Retention Latch (MRL) is ***not*** used], pull HP_MRL*x*# Low, to allow the outputs to toggle, regardless of whether the serial EEPROM is blank.

HP_PERST*x*# can also be toggled at runtime by toggling the *Power Controller Control* bit, provided that either the *Power Controller Present* bit is cleared, or that HP_PERST*x*# is initially de-asserted during slot power-up sequencing, as described in Section 11.4.1.2. A value of 1 asserts HP_PERST*x*# (Low). A value of 0 de-asserts HP_PERST*x*# (High).

Figure 11-2. Hot Plug Outputs when *Power Controller Present* and *Power Controller Control* Bits Are Cleared

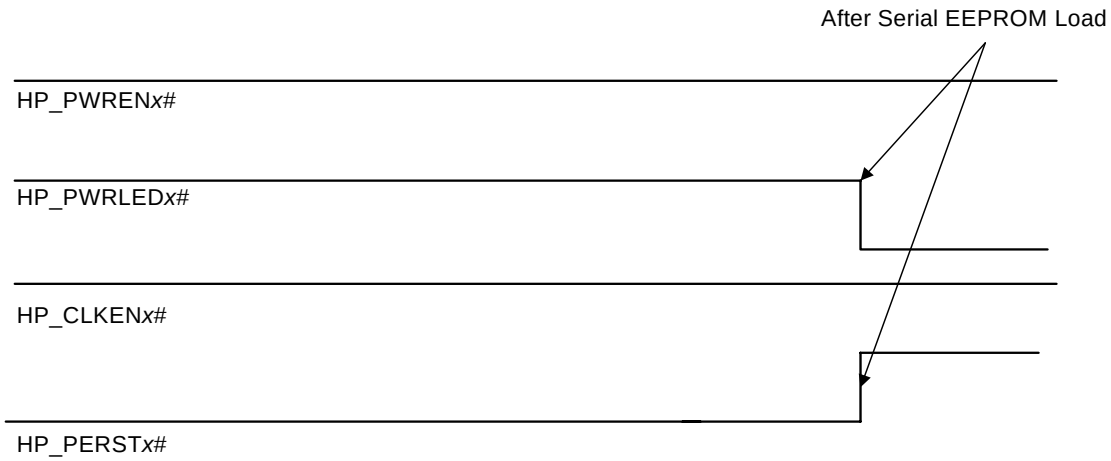

Note: *HP_PWRLEDx# is not asserted if the serial EEPROM or I^2C interface clears the Power Indicator Present bit (offset 7Ch[4]) to 0.*

11.4.1.4 Disabling Power-Up Hot Plug Output Sequencing

If the *Power Controller Control* bit is set to 1, after reset, the HP_PWREN*x*#, HP_PWRLED*x*#, and HP_CLKEN*x*# outputs remain High, and the HP_PERST*x*# output remains Low. The HP_PWREN*x*#, HP_PWRLED*x*#, and HP_CLKEN*x*# outputs also remain High if HP_MRL*x*# is not asserted in the default Hot Plug power-up sequencing described in Section 11.4.1.2. (Refer to Figure 11-3.)

Figure 11-3. Hot Plug Outputs when *Power Controller Control* Bit Is Set

11.4.2 Slot Power-Down Sequence

Software can power-down slots by setting the *Power Controller Control* bit (offset 80h[10]=1). If the *MRL Sensor Present* bit is set (offset 7Ch[2]=1), the Hot Plug Controller powers down the slot if the MRL is open. Figure 11-4 illustrates the following power-down timing sequence for either event:

1. HP_PERST*x*# to the port is asserted.
2. HP_CLKEN*x*# is de-asserted to the slot 100 µs after HP_PERST*x*# is asserted.
3. HP_PWREN*x*# is de-asserted to the slot 100 µs after HP_CLKEN*x*# is de-asserted.

Figure 11-4. Hot Plug Automatic Power-Down Sequence

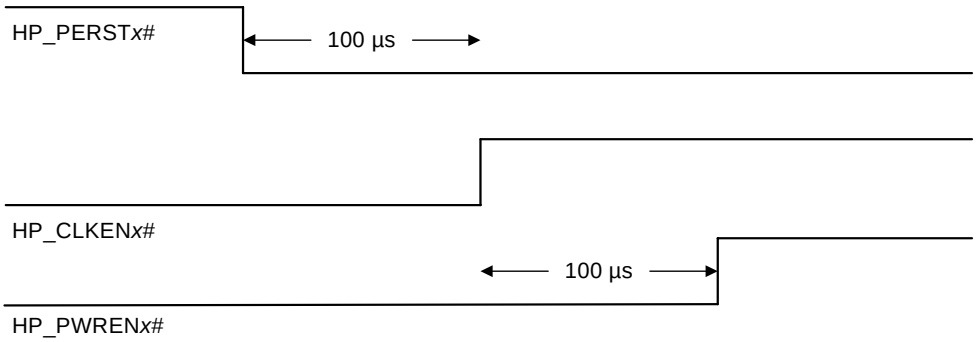

11.5 Hot Plug Board Insertion and Removal Process

Table 11-4 defines the board insertion procedure supported by the PEX 8508. Table 11-5 defines the board removal procedure.

Table 11-4. Hot Plug Board Insertion Process

Operator / Action	Hot Plug Controller	Software
A. Places board in slot.	**1.** Sets *Presence Detect State* bit to 1. **2.** Sets *Presence Detect Changed* bit to 1. **3.** Generates Interrupt message due to Presence Detect change, if enabled.	Clears *Presence Detect Changed* bit to 0.
	4. Transmits Interrupt de-assertion message, if enabled.	
B. Locks MRL.	**5.** Clears *MRL Sensor State* bit to 0. **6.** Sets *MRL Sensor Changed* bit to 1. **7.** Generates Interrupt message due to MRL Sensor state change, if enabled.	Clears *MRL Sensor Changed* bit to 0.
	8. Transmits Interrupt de-assertion message, if enabled.	
C. Presses Attention Button.	**9.** Sets *Attention Button Pressed* bit to 1. **10.** Generates Interrupt message due to Attention Button Pressed event, if enabled.	Clears *Attention Button Pressed* bit to 0.
	11. Transmits Interrupt de-assertion message, if enabled.	Writes to the **Slot Control** register *Power Indicator Control* field, to blink the Power Indicator LED, which indicates that the board is being powered up.

Continued …

Table 11-4. Hot Plug Board Insertion Process *(Cont.)*

Operator / Action	Hot Plug Controller	Software
D. Power Indicator blinks.	**12.** Sets *Power Indicator Control* field to 10b. **13.** Sets *Command Completed* bit to 1. **14.** Generates Interrupt message due to Power Indicator Blink command Completion, if enabled.	Clears *Command Completed* bit to 0.
	15. Transmits Interrupt de-assertion message, if enabled.	Clears **Slot Control** register *Power Controller Control* bit to 0, to turn On power to the port.
	16. Slot is powered up. **17.** After a T_{pepv} delay, sets *Command Completed* bit to 1. **18.** Generates Interrupt message due to Power Turn On command Completion, if enabled.	Clears *Command Completed* bit to 0.
	19. Transmits Interrupt de-assertion message, if enabled.	Writes to the **Slot Control** register *Power Indicator Control* field, to turn On the Power Indicator LED, which indicates that the slot is fully powered On.
	20. Sets *Power Indicator Control* field to 01b. **21.** Transmits Interrupt assertion message due to Power Indicator Turn On command Completion, if enabled.	Clears *Command Completed* bit to 0.
E. Power Indicator On.	**22.** Transmits Interrupt de-assertion message, if enabled.	Software can now read the **Link Status** register *Data Link Layer Link Active* bit (offset 78h[29]). A value of 1 in this bit indicates that the board is ready to be used.

Table 11-5. Hot Plug Board Removal Process

Operator / Action	Hot Plug Controller	Software
A. Presses Attention Button.	**1.** Sets *Attention Button Pressed* bit to 1. **2.** Generates Interrupt message due to Attention Button pressed, if enabled.	Clears *Attention Button Pressed* bit to 0.
	3. Transmits Interrupt de-assertion message, if enabled.	Writes to the **Slot Control** register *Power Indicator Control* field, to blink the Power Indicator LED, which indicates that the board is being powered down.
B. Power Indicator blinks.	**4.** Sets *Power Indicator Control* field to 10b. **5.** Sets *Command Completed* bit to 1. **6.** Generates Interrupt message due to Power Indicator Blink command Completion, if enabled.	Clears *Command Completed* bit to 0.
	7. Transmits Interrupt de-assertion message, if enabled.	Sets **Slot Control** register *Power Controller Control* bit to 1, to turn Off power to the port.
C. Power Indicator Off.	**8.** Slot is powered Off. **9.** After a 256-ms delay, sets the *Command Completed* bit to 1. **10.** Generates Interrupt message due to Power Turn Off command Completion, if enabled.	Clears *Command Completed* bit to 0. Clears *Power Indicator Control* field to 00b, to turn Off the Power Indicator LED, which indicates that the slot is fully powered Off and the board can be removed.
D. Power Indicator Off, board ready to be removed.	**11.** Clears *Power Indicator Control* field to 00b. **12.** Sets *Command Completed* bit to 1, due to Power Indicator Off command Completion.	Clears *Command Completed* bit to 0.
	13. Transmits Interrupt de-assertion message, if enabled.	
E. Unlocks MRL.	**14.** Sets *MRL Sensor State* bit to 1. **15.** Sets *MRL Sensor Changed* bit to 1. **16.** Generates Interrupt message due to MRL Sensor state change, if enabled.	Clears *MRL Sensor Changed* bit to 0.
	17. Transmits Interrupt de-assertion message, if enabled.	
F. Removes board from slot.	**18.** Clears *Presence Detect State* bit to 0. **19.** Sets *Presence Detect Changed* bit to 1. **20.** Generates Interrupt message due to Presence Detect change, if enabled.	Clears *Presence Detect Changed* bit to 0.
	21. Transmits Interrupt de-assertion message, if enabled.	

Chapter 12 Power Management

12.1 Power Management Capability

The Power Management (PM) module interfaces with different sections within the PEX 8508 to reduce power consumption. The PEX 8508 supports the following in PCI-compatible power management:

- Hardware-autonomous power management
- Link Power Management States (L States)
  - PCI Bus Power Management – L0, L1, L2 with VAUX support, L2/L3 Ready, and L3
  - Active State Power Management – L0s and L1
- AUX Power
- Software-driven D-State power management
- Device Power Management State (D-States) – D0 (D0_uninitialized and D0_active) and D3 (D3hot and D3cold) support
- Power Management Event (PME) support from D3hot and D3cold
- Power Management Event due to Hot Plug events
- Downstream ports generate and forward PME_Turn_Off broadcast messages
- Implements the *PCI Power Mgmt. r1.2*

The PM module interfaces with the Physical Layer electrical sub-block to transition the Link state into low-power states when the module receives a Power State Change request from an upstream component, or an internal event forces the Link State entry into low-power states in hardware-autonomous PM (ASPM) mode. PCI Express Link states are not directly visible to Conventional PCI Bus driver software; instead, these states are derived from the Power Management state of the components residing on those links. Figure 12-1 illustrates the PEX 8508 Power Management Controller functional block diagram.

Figure 12-1. Power Management Controller Functional Block Diagram

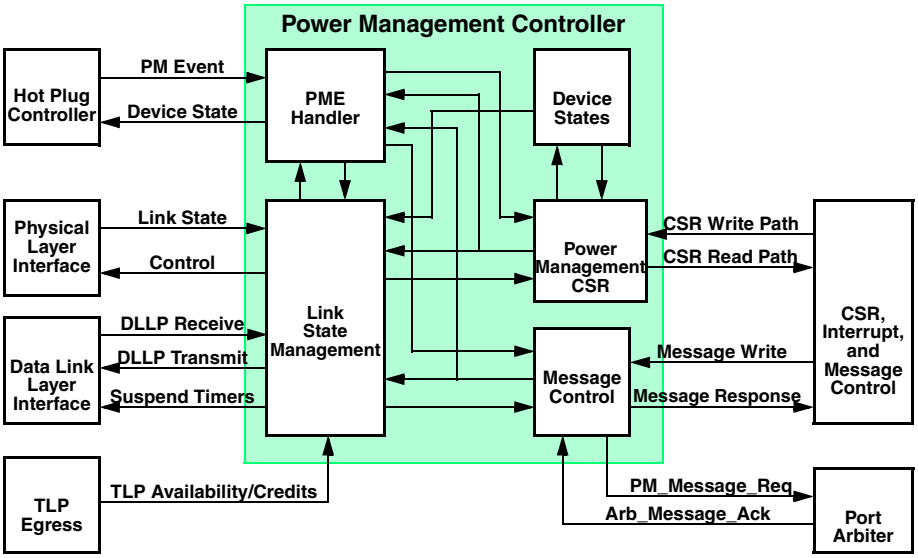

Note: *The Hot Plug Controller is available only on downstream ports.*

12.1.1 Device Power States

The PEX 8508 supports the PCI Express PCI-PM D0, D3hot, and D3cold Device Power Management states. The D1 and D2 states, which are optional in the *PCI Express Base r1.1*, are ***not supported*** by the PEX 8508.

The D3hot state can be entered from the D0 state, when system software programs the **Power Management Status and Control** register *Power State* field (offset 44h[1:0]=11b) for the appropriate port. The D0_uninitialized state can be entered from the D3hot state when the upstream and downstream links are in the L0s state and system software clears the **Power Management Status and Control** register *Power State* field (offset 44h[1:0]=00b).

12.1.1.1 D0 State

D0 is divided into two distinct sub-states – *uninitialized* and *active*. When power is initially applied to a PCI Express component, it defaults to the D0_uninitialized state. The component remains in the D0_uninitialized state until the serial EEPROM load is complete.

A device enters the *D0_active* state when:

- Any single Memory Access Enable occurs
- System software sets any combination of the **Command** register *Bus Master Enable*, *Memory Access Enable*, and/or *I/O Access Enable* bits (offset 04h[2, 1, and/or 0], respectively)

12.1.1.2 D3hot State

Once in the D3hot state, the PEX 8508 can be later transitioned into the D3cold state, by removing power from its Host component. Functions that are in the D3hot state can be transitioned, by software, to the D0_uninitialized state. When in the D3hot state, Hot Plug operations cause a PME in the PEX 8508.

Only Type 0 Configuration accesses are allowed in the D3hot state. Memory and I/O transactions result in an Unsupported Request (UR). Completions flowing in either direction are not affected.

Type 1 transactions toward a PEX 8508 port in the D3hot state are terminated as Unsupported Requests (UR). Type 0 Configuration transactions complete successfully. When the PEX 8508 upstream port is programmed to the D3hot state, the port initiates Conventional PCI-PM L1 entry.

12.1.1.3 D3cold State

The PEX 8508 transitions to the D3cold state when power is removed. Re-applying power causes the PEX 8508 to transition from the D3cold state into the D0_uninitialized state, followed by a configuration and link training sequence. The D3cold state assumes that all previous context is lost; therefore, software must save the required context while the PEX 8508 remains in the D3hot state. In the D3cold state, if auxiliary voltage (VAUX) is present, the PEX 8508 generates Beacon on its upstream port if it detects Beacon on any of the downstream ports, or its WAKE# signal is asserted.

12.1.2 Link Power Management States

The Power Management state of a link is determined by the D-state of its downstream link. The PEX 8508 switch holds its upstream link and downstream links in the L0 state when it is in standard operational mode (PCI PM state is in D0_active). Active State Power Management (ASPM) defines a mechanism for components in the D0 state to reduce link power by placing their links into a low-power state and instructs the opposite end of the link to do likewise. This capability allows hardware-autonomous, dynamic-link power reduction beyond what is achievable by Software-Only Power Management. Table 12-1 defines the relationships between the PEX 8508 power state and its downstream link.

Table 12-1. Power States of Connected Link Components

Downstream Component D State	PEX 8508 D State	Permissible Interconnect State	Power Saving Actions
D0	D0	L0	Full power.
		L0s, L1 (optional)	PHY Transmit lanes placed in high-impedance state.
D1	D0	L1	PHY Transmit lanes placed in high-impedance state.
D2			
D3hot	D0 or D3hot[a]	L1, L2/L3 Ready	PHY Transmit lanes placed in high-impedance state. FC and DLL ACK/NAK Timers suspended. PLL can be disabled.
D3cold	D0, D3hot, or D3cold	L2	Link-Off state. Draws power only from VAUX.

a. The PEX 8508 initiates a link-state transition of its upstream port to L1 when the port is programmed to D3hot.

12.1.3 PEX 8508 PCI Express Power Management Support

The PEX 8508 supports PCI Express features that are required or important for a PCI Express switch Power Management. Table 12-2 lists supported and non-supported features and the register bits/fields used for configuration or activation.

Table 12-2. PCI Express Power Management Capabilities Supported

Register		Description	Supported	
Offset	Bit(s)		Yes	No
40h		**Power Management Capability**		
	7:0	**Capability ID** Set to 01h, indicating that the data structure currently being pointed to is the Power Management data structure.	✔	
	15:8	**Next Capability Pointer** Default 48h points to the **Message Signaled Interrupt Capability** register.	✔	
	18:16	**Version** Default 011b indicates compliance with the *PCI Power Mgmt. r1.2*.	✔	
	19	**PME Clock** Does not apply to PCI Express. Returns 0.		✔
	21	**Device-Specific Initialization** Default 0 indicates that Device-Specific Initialization is ***not*** required.		✔
	24:22	**AUX Current** Value can be strapped to 111b in the serial EEPROM.	✔	
	25	**D1 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D1 Power state.		✔
	26	**D2 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D2 Power state.		✔
	31:27	**PME Support** Default value of 1100_1b indicates that the corresponding PEX 8508 port forwards PME messages in the D0, D3hot, and D3cold power states.	✔	

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
44h		**Power Management Status and Control**		
	1:0	**Power State** Used to determine the current power state of the port, and to set the port into a new power state. 00b = D0 01b = D1 – ***Not supported*** 10b = D2 – ***Not supported*** 11b = D3hot If software attempts to write an unsupported state to this field, the Write operation completes normally; however, the data is discarded and no state change occurs.	✔	
	3	**No Soft Reset**	✔	
	8	**PME Enable** 0 = Disables PME generation by the corresponding PEX 8508 port 1 = Enables PME generation by the corresponding PEX 8508 port	✔	
	12:9	**Data Select** Writable by Serial EEPROM only[a]. Bits [12:9] select the **Data** and **Data Scale** registers. 0h = D0 power consumed 3h = D3hot power consumed 4h = D0 power dissipated 7h = D3hot power dissipated All other values are ***reserved***.	✔	
		RO for hardware auto-configuration.		✔
	14:13	**Data Scale** Loaded by serial EEPROM[a]. There are four internal **Data Scale** registers (one each per **Data** register – 0, 3, 4 and 7), per port. Bits [12:9], *Data Select*, select the **Data Scale** register.	✔	
	15	**PME Status** 0 = Disables PME generation by the corresponding PEX 8508 port 1 = PME is being generated by the corresponding PEX 8508 port	✔	

*a. With no serial EEPROM, Reads return 00h for the **Data Scale** and **Data** registers (for all Data Selects).*

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
44h		**Power Management Control/Status Bridge Extensions**		
	22	**B2/B3 Support** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.		✔
	23	**Bus Power/Clock Control Enable** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.		✔
		Data		
	31:24	**Data** Loaded by serial EEPROM[a]. There are four internal **Data** registers (0, 3, 4, and 7), per port. Bits [12:9], *Data Select*, select the **Data** register.	✔	
6Ch		**Device Capabilities**		
	8:6	**Endpoint L0s Acceptable Latency** Because the PEX 8508 is a switch and not an endpoint, it does ***not support*** this feature. 000b = Disables the capability		✔
	11:9	**Endpoint L1 Acceptable Latency** Because the PEX 8508 is a switch and not an endpoint, it does ***not support*** this feature. 000b = Disables the capability		✔
	25:18	**Captured Slot Power Limit Value (Upstream Port)** For the PEX 8508 upstream port, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Scale* field. Do not change for the downstream ports.	✔	
	27:26	**Captured Slot Power Limit Scale (Upstream Port)** For the PEX 8508 upstream port, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Value* field. 00b = 1.0 01b = 0.1 10b = 0.01 11b = 0.001 Do not change for the downstream ports.	✔	

*a. With no serial EEPROM, Reads return 00h for the **Data Scale** and **Data** registers (for all Data Selects).*

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
70h		**Device Control**		
	10	**Auxiliary (AUX) Power PM Enable** Cleared to 0.	✔	
		Device Status		
	20	**Auxiliary (AUX) Power Detected** Cleared to 0.	✔	
74h		**Link Capabilities**		
	11:10	**Active State Power Management (ASPM) Support** Indicates the level of ASPM supported by the port. 01b = L0s link power state entry is supported 10b = L0s and L1 link power states are supported All other values are ***reserved***.	✔	
	14:12	**L0s Exit Latency** Indicates the L0s exit latency for the given PCI Express link. The value reported indicates the length of time the corresponding PEX 8508 port requires to complete the transition from L0s to L0. 101b = Corresponding PEX 8508 port L0s Exit Latency is between 1 and 2 µs	✔	
	17:15	**L1 Exit Latency** Indicates the L1 exit latency for the given PCI Express link. The value reported indicates the length of time the corresponding PEX 8508 port requires to complete the transition from L1 to L0. 101b = Corresponding PEX 8508 port L1 Exit Latency is between 16 and 32 µs	✔	
	18	**Clock Power Management**	✔	
78h		**Link Status and Control**		
	1:0	**Active State Power Management (ASPM) Control** 00b = Disables L0s and L1 Entry for the corresponding PEX 8508 port[b] 01b = Enables only L0s Entry 10b = Enables L1 Entry 11b = Enables both L0s and L1 Entry	✔	
	8	**Clock Power Management Enable** The PEX 8508 does ***not support*** removal of the Reference Clock in the L1 and L2/L3 Ready Link states.	✔	

b. The port receiver must be capable of entering the L0s Link PM state, regardless of whether the state is disabled

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
7Ch		**Slot Capabilities**		
	0	**Attention Button Present** 0 = Attention Button is not implemented 1 = Attention Button is implemented on the slot chassis of the corresponding PEX 8508 downstream port Do not change for the upstream port.	✔	
	1	**Power Controller Present** 0 = Power Controller is not implemented 1 = Power Controller is implemented for the slot of the corresponding PEX 8508 downstream port Do not change for the upstream port.	✔	
	2	**MRL Sensor Present** 0 = MRL Sensor is not implemented 1 = MRL Sensor is implemented on the slot chassis of the corresponding PEX 8508 downstream port Do not change for the upstream port.	✔	
	3	**Attention Indicator Present** 0 = Attention Indicator is not implemented 1 = Attention Indicator is implemented on the slot chassis of the corresponding PEX 8508 downstream port Do not change for the upstream port.	✔	

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
7Ch		**Slot Capabilities *(Cont.)***		
	4	**Power Indicator Present** 0 = Power Indicator is not implemented 1 = Power Indicator is implemented on the slot chassis of the corresponding PEX 8508 downstream port Do not change for the upstream port.	✔	
	5	**Hot Plug Surprise** 0 = No device in the corresponding PEX 8508 downstream port slot is removed from the system without prior notification 1 = Device in the corresponding PEX 8508 downstream port slot can be removed from the system without prior notification Do not change for the upstream port.	✔	
	6	**Hot Plug Capable** 0 = Corresponding PEX 8508 downstream port slot is not capable of supporting Hot Plug operations 1 = Corresponding PEX 8508 downstream port slot is capable of supporting Hot Plug operations Do not change for the upstream port.	✔	
	14:7	**Slot Power Limit Value** The maximum power supplied by the corresponding PEX 8508 downstream slot is determined by multiplying the value in this field (expressed in decimal; 25d = 19h) by the value specified by the *Slot Power Limit Scale* field. This field must be implemented if the **PCI Express Capabilities** register *Slot Implemented* bit is set (offset 68h[24])=1, default). Writes to this register also cause the downstream port to send the Set_Slot_Power_Limit message to the device connected to it, so as to convey the Limit value to the downstream device's upstream port **Device Capabilities** register *Captured Slot Power Limit Value* and *Captured Slot Power Limit Scale* fields. Do not change for the upstream port.	✔	
	16:15	**Slot Power Limit Scale** The maximum power supplied by the corresponding PEX 8508 downstream slot is determined by multiplying the value in this field by the *Slot Power Limit Value* bit. This field must be implemented if the **PCI Express Capabilities** register *Slot Implemented* bit is set (offset 68h[24])=1, default). Writes to this register also cause the downstream port to send the Set_Slot_Power_Limit message to the device connected to it, so as to convey the Limit value to the downstream device's upstream port **Device Capabilities** register *Captured Slot Power Limit Value* and *Captured Slot Power Limit Scale* fields. 00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x Do not change for the upstream port.	✔	

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
80h	**Slot Control**			
	1	**Power Fault Detector Enable** 0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Power State (**Power Management Status and Control** register *Power State* field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Power state (offset 44h[1:0]=11b), for a Power Fault Detected event on the corresponding Hot Plug-capable PEX 8508 downstream port. Do not change for the upstream port.	✔	
	9:8	**Power Indicator Control** Controls the Power Indicator on the corresponding PEX 8508 downstream port slot. 00b = ***Reserved*** – Writes are ignored 01b = Turns On indicator to constant On state 10b = Causes indicator to blink 11b = Turns Off indicator Software must use a Byte or Word Write (and not a DWord Write) to control the HP_PWRLED*x*# Output signal. Reads return the corresponding PEX 8508 downstream port Power Indicator's current state. Do not change for the upstream port.	✔	
	10	**Power Controller Control** Controls the Power Controller on the corresponding PEX 8508 downstream port slot. 0 = Turns On the Power Controller; requires some delay to be effective 1 = Turns Off the Power Controller Software must use a Byte or Word Write (and not a DWord Write) to control the Power Controller Output signals. Do not change for the upstream port.	✔	
	Slot Status			
	17	**Power Fault Detected** Set to 1 when the Power Controller of the corresponding PEX 8508 downstream port slot detects a Power Fault at the slot. Do not change for the upstream port.	✔	

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
138h		**Power Budgeting Enhanced Capability Header**		
	15:0	**PCI Express Extended Capability ID** Set to 0004h, as required by the *PCI Express Base r1.1*.	✔	
	19:16	**Capability Version** Set to 1h, as required by the *PCI Express Base r1.1*.	✔	
	31:20	**Next Capability Offset** Set to 148h, which addresses the PEX 8508 **Virtual Channel Enhanced Capability Header** registers.	✔	
13Ch		**Data Select**		
	7:0	**Data Select** Indexes the Power Budgeting Data reported, by way of eight registers per port, and selects the DWord of Power Budgeting Data that appears in each **Power Budgeting Data** register. Index values commence at 0 to select the first DWord of Power Budgeting Data; subsequent DWords of Power Budgeting Data are selected by increasing index values 1 to 7.	✔	

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
140h		**Power Budgeting Data**		
	7:0	**Base Power** Eight registers/port. Specifies (in Watts) the base power value in the operating condition. This value must be multiplied by the *Data Scale* (bits [9:8]) to produce the actual power consumption value.	✔	
	9:8	**Data Scale** Specifies the scale to apply to the Base Power value. The power consumption of the device is determined by multiplying the *Base Power* field (bits [7:0]) contents with the value corresponding to the encoding returned by this field. 00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x	✔	
	12:10	**PM Sub-State** 000b = Corresponding PEX 8508 port is in the default Power Management sub-state	✔	
	14:13	**PM State** Current power state. 00b = D0 state 11b = D3 state All other values are ***reserved***.	✔	
	17:15	**Type** Type of operating condition. 000b = PME Auxiliary 001b = Auxiliary 010b = Idle 011b = Sustained 111b = Maximum All other values are ***reserved***.	✔	
	20:18	**Power Rail** Power Rail of operating condition. 000b = Power 12V 001b = Power 3.3V 010b = Power 1.8V 111b = Thermal All other values are ***reserved***.	✔	
	Note: *There are eight registers per port that can be programmed, through the serial EEPROM. Each non-zero register value describes the power usage for a different operating condition. Each configuration is selected by writing to the* ***Data Select*** *register Data Select field (offset 13Ch[7:0]).*			

Table 12-2. PCI Express Power Management Capabilities Supported *(Cont.)*

Register		Description	Supported	
Offset	Bit(s)		Yes	No
144h		**Power Budget Capability**		
	0	**System Allocated** 1 = Power budget for the device is included within the system power budget	✔	
1E0h		**Power Management Hot Plug User Configuration**		
	0	**L0s Entry Idle Count** Time to meet to enter L0s. 0 = Idle condition lasts for 1 µs 1 = Idle condition lasts for 4 µs	✔	
	1	**L1 Upstream Port Receiver Idle Count** For active L1 entry. 0 = Upstream port receiver remains idle for 2 µs 1 = Upstream port receiver remains idle for 3 µs	✔	
	2	**HPC PME Turn-Off Enable** 1 = PME Turn-Off message is transmitted before the port is turned Off on a downstream port	✔	
	4:3	**HPC Tpepv Delay** Slot power-applied to power-valid delay time. 00b = 16 ms 01b = 32 ms 10b = 64 ms 11b = 128 ms	✔	
	5	**HPC Inband Presence Detect Enable** 0 = HP_PRSNT[4:0]# Input balls are used to detect a board present in the slot 1 = SerDes Receiver Detect mechanism is used to detect a board present in the slot	✔	
	6	**HPC Tpvperl Delay** Downstream port power valid to reset signal release time. 0 = 20 ms 1 = 100 ms (default)	✔	

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 13 Transparent Port Registers

13.1 Introduction

This chapter defines the PEX 8508 Transparent mode port registers (some of which are also used by the NT Port Virtual and Link Interfaces). The PEX 8508 ports have their own Configuration, Capability, Control, and Status register space. The register mapping is the same for each port. (Refer to Table 13-1.) This chapter also presents the PEX 8508 user-programmable registers and the order in which they appear in the register map. Register descriptions, when applicable, include details regarding their use and meaning in the upstream port and downstream ports. (Refer to Figure 13-1.)

The NT Port registers are defined in Chapter 17, "NT Port Virtual Interface Registers," and Chapter 18, "NT Port Link Interface Registers."

For further details regarding register names and descriptions, refer to the following specifications:

- *PCI r3.0*
- *PCI Power Mgmt. r1.2*
- *PCI-to-PCI Bridge r1.2*
- *PCI Express Base r1.1*
- *I^2C Bus v2.1*

Note: *For the PEX 8508 to properly route Memory and I/O requests and Completions, the station contains Content Addressable Memory (CAM) registers that hold mirror copies of certain registers in each port. Refer to the PEX 85XX EEPROM – PEX 8518/8517/8508 Design Note, Section 6.13, "Shadowed Registers." If the registers that are shadowed are programmed by the serial EEPROM to non-default values, the corresponding CAM registers must be programmed by serial EEPROM to contain the same values as the shadowed registers.*

13.2 Type 1 PEX 8508 Port Register Map

Table 13-1. Type 1 PEX 8508 Port Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
PCI-Compatible Type 1 Configuration Registers		00h …
	New Capability Pointer (40h)	34h … 3Ch
Next Capability Pointer (48h)	Capability ID (01h)	40h
Power Management Capability Registers		44h
Next Capability Pointer (68h)	Capability ID (05h)	48h
Message Signaled Interrupt Capability Registers		… 64h
Next Capability Pointer (00h)	Capability ID (10h)	68h
PCI Express Capability Registers		… 8Ch
NT Port Virtual Interface Registers (***Reserved in Transparent Mode***)		90h – FCh
Next Capability Offset (FB4h) / 1h	PCI Express Extended Capability ID (0003h)	100h
Device Serial Extended Number Capability Registers		… 134h
Next Capability Offset (148h) / 1h	PCI Express Extended Capability ID (0004h)	138h
Power Budgeting Extended Capability Registers		… 144h
Next Capability Offset (000h) / 1h	PCI Express Extended Capability ID (0002h)	148h
Virtual Channel Extended Capability Registers		… 1C4h
Device-Specific Registers		1C8h … C38h
Non-Transparent Bridging-Specific Registers (***Reserved in Transparent Mode***)		C3Ch – DF0h
Reserved		DF4h – FB0h
Next Capability Offset (138h) / 1h	PCI Express Extended Capability ID (0001h)	FB4h
Advanced Error Reporting Capability Registers		… FFCh

13.3 PEX 8508 Port Register Configuration and Map

The PEX 8508 port registers are configured similarly – not all the same. Port 0 includes more device-specific registers than the other ports. The registers for this port contain station-specific setup and control information. Also, Port 0 contains registers that are used to set up and control the PEX 8508, and serial EEPROM and/or I^2C interface logic and control.

Table 13-2 defines the register map by port. (Refer to Appendix A, "Serial EEPROM Memory Map," for a detailed description of this register map.)

Table 13-2. PEX 8508 Port Register Configuration and Map

Register Types	Port 0	Ports 1, 2, 3, 4
PCI-Compatible Type 1 Configuration Registers	00h - 3Ch	00h - 3Ch
Power Management Capability Registers	40h - 44h	40h - 44h
Message Signaled Interrupt Capability Registers	48h - 64h	48h - 64h
PCI Express Capability Registers	68h - 8Ch	68h - 8Ch
Device Serial Extended Number Capability Registers	100h - 134h	100h - 134h
Power Budgeting Extended Capability Registers	138h - 144h	138h - 144h
Virtual Channel Extended Capability Registers	148h - 1C4h	148h - 1C4h
Device-Specific Registers		
Error Checking and Debug Registers	1C8h - 1FCh	1CCh - 1D0h, 1E0h - 1F0h, 1F8h
ECC Error Check Disable	1C8h	
Error Handler	1CCh - 1D0h	1CCh - 1D0h
Debug Control	1D4h - 1DCh	
Power Management, Hot Plug, and Miscellaneous Control	1E0h - 1FCh	1E0h - 1F0h, 1F8h
Physical Layer Registers	200h - 2C4h	
SerDes	238h - 258h	
Serial EEPROM	260h - 264h	
I2C Interface Register	290h - 2C4h	
Bus Number CAM Registers	2C8h - 304h	
I/O CAM Registers	308h - 344h	
Address-Mapping CAM (AMCAM) Registers	348h - 548h	
Ingress Control and Port Enable Registers	660h - 67Ch	
I/O CAM Base and Limit Upper 16-Bit Registers	680h - 6BCh	
Base Address Shadow Registers (BARs)	6C0h - 73Ch	
Shadow Virtual Channel (VC) Capability Registers	740h - 83Ch	
Shadow Port Virtual Channel (VC) Capability Registers	840h - 9ECh	
Ingress Credit Handler (INCH) Registers	9F0h - BDCh	A00h - A50h, A60h - A68h
Relaxed Ordering and Performance Counter Registers	BE0h - BFCh	
Internal Credit Handler (ITCH) VC&T Threshold Registers	C00h - C08h	
Port Virtual Channel Queue Status Registers	C0Ch - C38h	
Advanced Error Reporting Capability Registers	FB4h - FFCh	FB4h - FFCh

13.4 Register Access

Each PEX 8508 port implements a 4-KB Configuration space. The lower 256 bytes (offsets 00h through FFh) is the PCI-compatible Configuration space, and the upper 960 Dwords (offsets 100h through FFFh) is the PCI Express Extended Configuration space. The PEX 8508 supports three mechanisms for accessing registers:

- PCI r3.0-Compatible Configuration Mechanism
- PCI Express Enhanced Configuration Mechanism
- Device-Specific Memory-Mapped Configuration Mechanism

13.4.1 *PCI r3.0-Compatible* Configuration Mechanism

The *PCI r3.0*-Compatible Configuration mechanism provides standard access to the PEX 8508 ports' first 256 bytes (the bytes at offsets 00h through FFh) of the PCI Express Configuration space. This mechanism is used to access the PEX 8508 port Type 1 (PCI-to-PCI Bridge) registers:

- PCI-Compatible Type 1 Configuration Registers
- Power Management Capability Registers
- Message Signaled Interrupt Capability Registers
- PCI Express Capability Registers

The *PCI r3.0*-Compatible Configuration mechanism uses PCI Type 0 and Type 1 Configuration transactions to access the PEX 8508 Configuration registers. The PEX 8508 upstream port captures the Bus and Device Numbers assigned by the upstream device on the PCI Express link attached to the PEX 8508 upstream port, as required by the *PCI Express Base r1.1*.

The PEX 8508 decodes all Type 1 Configuration accesses received on its upstream port, when any of the following conditions exist:

- Specified Bus Number in the Configuration access is the PEX 8508 internal virtual PCI Bus Number, the PEX 8508 automatically converts the Type 1 Configuration access into the appropriate Type 0 Configuration access for the specified device.
  - If the specified device corresponds to the PCI-to-PCI bridge in one of the PEX 8508 downstream ports, the PEX 8508 processes the Read or Write request to the specified downstream port register specified in the original Type 1 Configuration access.
  - If the specified Device Number does not correspond to any of the PEX 8508 downstream port Device Numbers, the PEX 8508 responds with an *Unsupported Request* (UR).
    - If the specified Bus Number in the Type 1 Configuration access is not the PEX 8508 internal virtual PCI Bus Number, but is the number of one of the PEX 8508 downstream port secondary/subordinate buses, the PEX 8508 passes the Configuration access on to the PCI Express link attached to that PEX 8508 downstream port.
    - If the specified Bus Number is the downstream port Secondary Bus Number, and specified Device Number is 0, the PEX 8508 converts the Type 1 Configuration access to a Type 0 Configuration access before passing it on.
    - If the specified Device Number is not 0, the downstream port drops the TLP and generates a UR.
    - If the specified Bus Number is not the downstream port Secondary Bus Number, the PEX 8508 passes along the Type 1 Configuration access, without change.

Because the *PCI r3.0*-Compatible Configuration mechanism is limited to the first 256 bytes of the PCI Express Configuration Space of the PEX 8508 ports, the PCI Express Enhanced Configuration mechanism (described in Section 13.4.2) or Device-Specific Memory-Mapped Configuration mechanism (described in Section 13.4.3) must be used to access beyond byte FFh. The PCI Express Enhanced Configuration mechanism can access the registers in the PCI-compatible region, as well as those in the PCI Express Extended Configuration space that are defined by the *PCI Express Base r1.1*; however, it generally cannot access the PEX 8508 device-specific registers above 100h. The Device-Specific Memory-Mapped Configuration mechanism can access all PEX 8508 registers.

13.4.2 PCI Express Enhanced Configuration Mechanism

The PCI Express Enhanced Configuration mechanism is implemented on all PCI Express PCs and on systems that do not implement a processor-specific firmware interface to the Configuration space, providing a Memory-Mapped Address space in the Root Complex through which the Root Complex translates a Memory access into one or more Configuration requests. Device drivers normally use an application programming interface (API) provided by the Operating System, to use the PCI Express Enhanced Configuration mechanism.

The PCI Express Enhanced Configuration mechanism is used to access the PEX 8508 port Type 1 (PCI-to-PCI Bridge) registers that are defined by the *PCI Express Base r1.1*:

- PCI-Compatible Type 1 Configuration Registers
- Power Management Capability Registers
- Message Signaled Interrupt Capability Registers
- PCI Express Capability Registers
- Device Serial Extended Number Capability Registers
- Power Budgeting Extended Capability Registers
- Virtual Channel Extended Capability Registers
- Advanced Error Reporting Capability Registers

The PEX 8508 device-specific registers that exist in the PCI Express Extended Configuration space (above 100h) generally cannot be accessed by the PCI Express Enhanced Configuration mechanism. The Device-Specific Memory-Mapped Configuration mechanism (described in Section 13.4.3) can access all PEX 8508 registers.

13.4.3 Device-Specific Memory-Mapped Configuration Mechanism

The Device-Specific Memory-Mapped Configuration mechanism provides a method to access the Configuration registers of all ports in a single Memory map, as illustrated in Figure 13-1. The registers of each port are located within a 4-KB range. The PEX 8508 supports up to five simultaneously active ports.

The PEX 8508 requires a single contiguous Memory space of 128 KB to contain all the PEX 8508 Configuration registers and sufficient Memory space to support software compatibility for future device expansion.

To use the Device-Specific Memory-Mapped Configuration mechanism, program the upstream port's Type 1 Configuration Space **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively). After the PEX 8508 upstream port Base Address registers are configured, Port 0 registers can be accessed with Memory Reads from and Writes to the first 4 KB (0000h to 0FFFh), Port 1 registers can be accessed with Memory Reads from and Writes to the second 4 KB (1000h to 1FFFh), and so forth. Within each of these 4-KB windows, individual registers are located at the DWord offsets indicated in Table 13-1.

The upstream port **BAR0** and **BAR1** are typically enumerated at boot time, by BIOS or the Operating System (OS) software. When the registers are written (by serial EEPROM, I^2C interface, or software), the PEX 8508 automatically copies the values into the **BAR0** and **BAR1 Shadow** registers that exist in Port 0, located at offsets 6C0h through 6E4h. The registers used within this block depend upon which port is the upstream port.

If the upstream port **BAR0** and **BAR1** are enumerated by serial EEPROM, rather than by BIOS/OS, the serial EEPROM must be programmed to also load the same values to the corresponding **BAR0** and **BAR1 Shadow** registers.

Figure 13-1. PEX 8508 Register Offset from Upstream Port BAR0/1 Base Address (Transparent Mode)

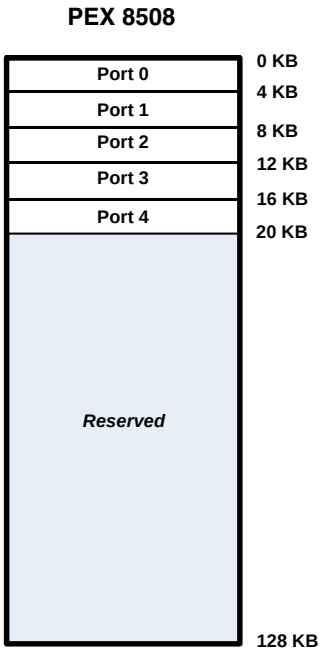

13.5 Register Descriptions

The remainder of this chapter details the PEX 8508 registers, including:

- Bit/field names
- Description of register functions for the PEX 8508 upstream port and downstream ports
- Type (*such as* RW or HwInit; refer to Table 13-3 for Type descriptions)
- Whether the power-on/reset value can be modified, by way of the PEX 8508 serial EEPROM or I^2C initialization feature
- Default power-on/reset value

Table 13-3. Register Types, Grouped by User Accessibility

Type	Description
HwInit	**Hardware-Initialized** Refers to the PEX 8508 Hardware Initialization mechanism or PEX 8508 Serial EEPROM register initialization feature. Read-Only after initialization and can only be reset with a Fundamental Reset.
RW	**Read-Write** Read/Write and is set or cleared to the needed state by software.
RW1C	**Read-Only Status, Write 1 to Clear** Write 1 to clear status register or bit. Indicates status when read. A status bit set by the system to 1 (to indicate status) is cleared by writing 1 to that bit. Writing 0 has no effect.
RW1CS	**Read-Only Status, Write 1 to Clear, Sticky** Same as RW1C, except that bits are not modified by a Hot Reset.
RW1S	**Read-Write, Write 1 to Set, Sticky** Non-Transparent ports contain these types of Device-Specific Control registers. Software writes 1 to the register to enable control and 1 to a register with RW1C privilege to clear the control. Writing 0 has no effect.
RWS	**Read-Write, Sticky** Same as RW, except that bits are not modified by a Hot Reset.
RO	**Read-Only** Read-Only and cannot be altered by software. Initialized by the PEX 8508 hardware initialization mechanism or PEX 8508 serial EEPROM register initialization feature.
ROS	**Read-Only, Sticky** Same as RO, except that bits are not initialized nor modified by a Hot Reset.
RsvdP	***Reserved* and Preserved** ***Reserved*** for future Read/Write implementations. Registers are Read-Only and must return 0 when read. Software must preserve value read for writes to bits.
RsvdZ	***Reserved* and Zero** Reserved for future RW1C implementations. Registers are Read-Only and must return 0 when read. Software must use 0 for writes to bits.
RZ	**Read Return Zero** Software Read always returns 0; however, software is allowed to write this register.

13.6 PCI-Compatible Type 1 Configuration Registers

Table 13-4. PCI-Compatible Type 1 Configuration Register Map (All Ports)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Device ID		Vendor ID		00h
Status		Command		04h
Class Code			Revision ID	08h
BIST *(Not Supported)*	Header Type	Master Latency Timer	Cache Line Size	0Ch
Base Address 0				10h
Base Address 1				14h
Secondary Latency Timer *(Not Supported)*	Subordinate Bus Number	Secondary Bus Number	Primary Bus Number	18h
Secondary Status		I/O Limit	I/O Base	1Ch
Memory Limit		Memory Base		20h
Prefetchable Memory Limit		Prefetchable Memory Base		24h
Prefetchable Memory Upper Base Address				28h
Prefetchable Memory Upper Limit Address				2Ch
I/O Limit Upper 16 Bits		I/O Base Upper 16 Bits		30h
Reserved			New Capability Pointer (40h)	34h
Expansion ROM Base Address ***(Not Supported)***				38h
Bridge Control		Interrupt Pin	Interrupt Line	3Ch

Register 13-1. 00h Configuration ID (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Vendor ID** Unless overwritten by the serial EEPROM, returns the PLX PCI-SIG-assigned Vendor ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Vendor ID with another Vendor ID.	HwInit	Yes	10B5h
31:16	**Device ID** Unless overwritten by the serial EEPROM, the PEX 8508 returns 8508h, the PLX-assigned Device ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX-assigned Device ID with another Device ID.	HwInit	Yes	8508h

Register 13-2. 04h Command/Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Command			
0	**I/O Access Enable** 0 = PEX 8508 ignores I/O accesses on the corresponding port's primary interface 1 = PEX 8508 responds to I/O accesses on the corresponding port's primary interface	RW	Yes	0
1	**Memory Access Enable** 0 = PEX 8508 ignores Memory accesses on the corresponding port's primary interface 1 = PEX 8508 responds to Memory accesses on the corresponding port's primary interface	RW	Yes	0
2	**Bus Master Enable** Controls the PEX 8508 Memory and I/O request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8508 handles Memory and I/O requests received on the corresponding port downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8508 returns a Completion with UR Completion status 1 = PEX 8508 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
3	**Special Cycle Enable** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
4	**Memory Write and Invalidate Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
5	**VGA Palette Snoop** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
6	**Parity Error Response Enable** Controls bit 24 (*Master Data Parity Error* bit).	RW	Yes	0
7	**IDSEL Stepping/Wait Cycle Control** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
8	**SERR# Enable** Controls bit 30 (*Signaled System Error* bit). 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
9	**Fast Back-to-Back Transactions Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
10	**Interrupt Disable** 0 = Corresponding PEX 8508 port is enabled to generate INT*x* Interrupt messages 1 = Corresponding PEX 8508 port is prevented from generating INT*x* Interrupt messages	RW	Yes	0
15:11	***Reserved***	RsvdP	No	00h

Register 13-2. 04h Command/Status (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Status			
18:16	***Reserved***	RsvdP	No	000b
19	**Interrupt Status** 0 = No INT*x* Interrupt message is pending 1 = INT*x* Interrupt message is pending internally to the corresponding PEX 8508 port	RO	Yes	0
20	**Capability List** Required by the *PCI Express Base r1.1* to be 1 at all times.	RO	Yes	1
21	**66 MHz Capable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
22	***Reserved***	RsvdP	No	0
23	**Fast Back-to-Back Transactions Capable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
24	**Master Data Parity Error** When the *Parity Error Response Enable* bit is set to 1, the corresponding PEX 8508 port sets this bit to 1 when the port: • Forwards the poisoned TLP Write request from the secondary to the primary interface, **or** • Receives a Completion marked as poisoned on the primary interface When the *Parity Error Response Enable* bit is cleared to 0, the PEX 8508 never sets this bit. This error is natively reported by the **Uncorrectable Error Status** register *Poisoned TLP Status* bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0
26:25	**DEVSEL# Timing** Cleared to 00b, as required by the *PCI Express Base r1.1*.	RsvdP	No	00b
27	**Signaled Target Abort** When a Memory-Mapped access payload length is greater than one DWord, the PEX 8508 upstream port sets this bit to 1. This error is natively reported by the **Uncorrectable Error Status** register *Completer Abort Status* bit (offset FB8h[15]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0

Register 13-2. 04h Command/Status (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
28	**Received Target Abort** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
29	**Received Master Abort** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
30	**Signaled System Error** When the *SERR# Enable* bit is set to 1, the corresponding PEX 8508 port sets this bit to 1 when transmitting an ERR_FATAL or ERR_NONFATAL message to its upstream device. This error is natively reported by the **Device Status** register *Fatal Error Detected* and *Non-Fatal Error Detected* bits (offset 70h[18:17], respectively), which are mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0
31	**Detected Parity Error** Set to 1 when the corresponding port receives a Poisoned TLP on its primary side, regardless of the *Parity Error Response Enable* bit state. This error is natively reported by the **Uncorrectable Error Status** register *Poisoned TLP Status* bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0

Register 13-3. 08h Class Code and Revision ID (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Revision ID** Unless overwritten by the serial EEPROM, returns ACh, the PLX-assigned Revision ID for this version of the PEX 8508. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Revision ID with another Revision ID. ***Note:*** *When a serial EEPROM is not used, the default values for this field are loaded. When using a serial EEPROM, refer to the following table for the appropriate configuration settings for this field.*	RO	Yes	ACh

	Serial EEPROM Value	Results in Actual Register Value
Revision ID	ACh	AAh
	AAh	ACh

Class Code				060400h
15:8	**Register-Level Programming Interface** PEX 8508 ports support the *PCI-to-PCI Bridge r1.2* requirements, but not subtractive decoding, on its upstream interface.	RO	Yes	00h
23:16	**Sub-Class Code** PCI-to-PCI bridge.	RO	Yes	04h
31:24	**Base Class Code** Bridge device.	RO	Yes	06h

Register 13-4. 0Ch Miscellaneous Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Cache Line Size** System Cache Line Size. Implemented as a Read-Write field for Conventional PCI compatibility purposes and does not impact PEX 8508 functionality.	RW	Yes	00h
15:8	**Master Latency Timer** ***Not supported*** Cleared to 00h, as required by the *PCI Express Base r1.1*.	RsvdP	No	00h
22:16	**Configuration Layout Type** The Corresponding PEX 8508 port Configuration Space Header adheres to the Type 1 PCI-to-PCI Bridge Configuration Space layout defined by the *PCI-to-PCI Bridge r1.2*.	RO	Yes	01h
23	**Header Type** Always 0, because the PEX 8508 is a single-function device.	RO	Yes	0
31:24	**BIST** ***Not supported*** Built-In Self-Test (BIST) Pass/Fail.	RsvdP	No	00h

Register 13-5. 10h Base Address 0 [Upstream Port Only; *Reserved* (RsvdP) for Downstream Ports]

Bit(s)	Description	Ports	Type	Serial EEPROM and I^2C	Default
0	**Memory Space Indicator** When enabled, the Base Address register maps the corresponding PEX 8508 port Configuration registers into Memory space. ***Note:*** *The upstream port is hardwired to 0.*	Upstream	RO	No	0
	Reserved for the downstream ports.	Downstream	RsvdP	No	0
2:1	**Memory Map Type** 00b = Corresponding PEX 8508 port Configuration registers can be mapped anywhere in 32-bit Memory Address space 10b = Corresponding PEX 8508 port Configuration registers can be mapped anywhere in 64-bit Memory Address space 01b, 11b = ***Reserved***	Upstream	RO	Yes	00b
	Reserved for the downstream ports.	Downstream	RsvdP	No	00b
3	**Prefetchable** The Base Address register maps the corresponding PEX 8508 port Configuration registers into Non-Prefetchable Memory space by default. ***Notes:*** *The upstream port is hardwired to 0.* *If an application (System) needs to map this BAR into Prefetchable Address space, the System must use a Configuration TLP to access the Performance Counter. Otherwise, the Prefetchable CSR Read clears the Performance Counter, irrespective of Byte Enable bits.*	Upstream	RO	No	0
	Reserved for the downstream ports.	Downstream	RsvdP	No	0
16:4	***Reserved***		RsvdP	No	0-0h
31:17	**Base Address** Base Address for Device-Specific Memory-Mapped Configuration Mechanism.	Upstream	RW	Yes	0000h
	Reserved for the downstream ports.	Downstream	RsvdP	No	0000h

Register 13-6. 14h Base Address 1 [Upstream Port Only; *Reserved* (RsvdP) for Downstream Ports]

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 1** For 64-bit addressing [**Base Address 0** register *Memory Map Type* field (offset 10h[2:1]) is set to 10b], Base Address 1 extends Base Address 0 to provide the upper 32 Address bits.	Upstream	RW	Yes	0000_0000h
	Read-Only when the **Base Address 0** register *Memory Map Type* field indicates 32-bit memory addressing (offset 10h[2:1]=00b).		RO	No	0000_0000h
	Reserved for the downstream ports.	Downstream	RsvdP	No	0000_0000h

Register 13-7. 18h Bus Number (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Primary Bus Number** Records the Bus Number of the PCI Bus segment to which the primary interface of this port is connected. Set by Configuration software.	RW	Yes	00h
15:8	**Secondary Bus Number** Records the Bus Number of the PCI Bus segment that is the secondary interface of this port. Set by Configuration software.	RW	Yes	00h
23:16	**Subordinate Bus Number** Records the Bus Number of the highest numbered PCI Bus segment that is subordinate to this port. Set by Configuration software.	RW	Yes	00h
31:24	**Secondary Latency Timer** Cleared to 00h, as required by the *PCI Express Base r1.1*.	RsvdP	No	00h

Register 13-8. 1Ch Secondary Status, I/O Limit, and I/O Base (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	I/O Base			
3:0	**I/O Base Addressing Capability** 1h = 32-bit I/O address decoding is supported Other values are ***reserved***.	RO	Yes	1h
7:4	**I/O Base Address[15:12]** The PEX 8508 ports use their **I/O Base** and **I/O Limit** registers to determine the Address range of I/O transactions to forward from one interface to the other. I/O Base Address[15:12] bits specify the corresponding PEX 8508 port I/O Base Address[15:12]. The PEX 8508 assumes I/O Base Address[11:0]=000h. For 16-bit I/O addressing, the PEX 8508 assumes Address[31:16]=0000h. For 32-bit addressing, the PEX 8508 decodes Address[31:0], and uses the **I/O Upper Base and Limit Address** register *I/O Base Upper 16 Bits* and *I/O Limit Upper 16 Bits* fields (offset 30h[15:0 and 31:16], respectively).	RW	Yes	Fh
	I/O Limit			
11:8	**I/O Limit Addressing Capability** 1h = 32-bit I/O address decoding is supported Other values are ***reserved***.	RO	Yes	1h
15:12	**I/O Limit Address[15:12]** The PEX 8508 ports use their **I/O Base** and **I/O Limit** registers to determine the Address range of I/O transactions to forward from one interface to the other. I/O Limit Address[15:12] specify the corresponding PEX 8508 port I/O Limit Address[15:12]. The PEX 8508 assumes Address bits [11:0] of the I/O Limit Address are FFFh. For 16-bit I/O addressing, the PEX 8508 decodes Address bits [15:0] and assumes Address bits [31:16] of the I/O Limit Address are 0000h. For 32-bit addressing, the PEX 8508 decodes Address bits [31:0], and uses the **I/O Upper Base and Limit Address** register *I/O Base Upper 16 Bits* and *I/O Limit Upper 16 Bits* fields (offset 30h[15:0 and 31:16], respectively). When the I/O Limit Address is less than the I/O Base Address, the PEX 8508 does not forward I/O transactions from the corresponding port primary/upstream bus to its secondary/downstream bus. However, the PEX 8508 forwards all I/O transactions from the secondary bus of the corresponding port to its primary bus.	RW	Yes	0h

Register 13-8. 1Ch Secondary Status, I/O Limit, and I/O Base (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Secondary Status			
20:16	***Reserved***	RsvdP	No	0-0h
21	**66 MHz Capable** ***Not supported*** 0 = Not enabled, because PCI Express does ***not support*** 66 MHz	RsvdP	No	0
22	***Reserved***	RsvdP	No	0
23	**Fast Back-to-Back Transactions Capable** ***Reserved*** Not enabled, because PCI Express does ***not support*** this function.	RsvdP	No	0
24	**Master Data Parity Error** If the *Parity Error Response Enable* bit is set to 1, the corresponding PEX 8508 port sets this bit to 1 when transmitting or receiving a TLP on its downstream side, and when either of the following two conditions occur: • Port receives Completion marked poisoned • Port forwards poisoned TLP write request When the *Parity Error Response Enable* bit = 0, the PEX 8508 never sets this bit.	RW1C	Yes	0
26:25	**DEVSEL# Timing** ***Not supported*** Cleared to 00b, as required by the *PCI Express Base r1.1*.	RsvdP	No	00b
27	**Signaled Target Abort** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
28	**Received Target Abort** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
29	**Received Master Abort** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
30	**Received System Error** Set to 1 when a port receives an ERR_FATAL or ERR_NONFATAL message on its secondary interface.	RW1C	Yes	0
31	**Detected Parity Error** Set to 1 by the secondary side of a Type 1 Configuration Space Header device when receiving a poisoned TLP, regardless of the *Parity Error Response Enable* bit state.	RW1C	Yes	0

Register 13-9. 20h Memory Base and Limit (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Memory Base				
3:0	***Reserved***	RsvdP	No	0h
15:4	**Memory Base Address[31:20]** Specifies the corresponding PEX 8508 port Memory Base Address[31:20]. The PEX 8508 assumes Memory Base Address[19:0]=00000h.	RW	Yes	FFFh
Memory Limit				
19:16	***Reserved***	RsvdP	No	0h
31:20	**Memory Limit Address[31:20]** Specifies the corresponding PEX 8508 port non-prefetchable Memory Limit Address[31:20]. The PEX 8508 assumes Memory Limit Address[19:0]=FFFFFh.	RW	Yes	000h

Note: *The PEX 8508 port forwards Memory transactions from its primary interface to its secondary interface (downstream) if a Memory address is within the range defined by the* ***Memory Base*** *and* ***Memory Limit*** *registers (when the Base is less than or equal to the Limit).*

Conversely, the PEX 8508 port forwards Memory transactions from its secondary interface to its primary interface (upstream) if a Memory address is outside this Address range [provided the address is not within the range defined by the ***Prefetchable Memory Base*** *(offsets 28h + 24h[15:0]) and* ***Prefetchable Memory Limit*** *(offsets 2Ch + 24h[31:16]) registers].*

Register 13-10. 24h Prefetchable Memory Base and Limit (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Prefetchable Memory Base				
3:0	**Prefetchable Memory Base Capability** 0h = Corresponding PEX 8508 port supports 32-bit Prefetchable Memory Addressing 1h = Corresponding PEX 8508 port defaults to 64-bit Prefetchable Memory Addressing support, as required by the *PCI Express Base r1.1* ***Note:*** *If the application needs 32-bit only Prefetchable space, the serial EEPROM or I^2C must clear both this field and field [19:16] (**Prefetchable Memory Limit** register Prefetchable Memory Limit Capability field).*	RO	Yes	1h
15:4	**Prefetchable Memory Base Address[31:20]** Specifies the corresponding PEX 8508 port Prefetchable Memory Base Address[31:20]. The PEX 8508 assumes Prefetchable Memory Base Address[19:0]=00000h.	RW	Yes	FFFh
Prefetchable Memory Limit				
19:16	**Prefetchable Memory Limit Capability** 0h = Corresponding PEX 8508 port supports 32-bit Prefetchable Memory Addressing 1h = Corresponding PEX 8508 port defaults to 64-bit Prefetchable Memory Addressing support, as required by the *PCI Express Base r1.1*	RO	Yes	1h
31:20	**Prefetchable Memory Limit Address[31:20]** Specifies the corresponding PEX 8508 port Prefetchable Memory Limit Address[31:20]. The PEX 8508 assumes Prefetchable Memory Limit Address[19:0]=FFFFFh.	RW	Yes	000h

Note: *The PEX 8508 port forwards Memory transactions from its primary interface to its secondary interface (downstream) if a Memory address is within the range defined by the **Prefetchable Memory Base** (offsets 28h + 24h[15:0]) and **Prefetchable Memory Limit** (offsets 2Ch + 24h[31:16]) registers (when the Base is less than or equal to the Limit).*

*Conversely, the PEX 8508 port forwards Memory transactions from its secondary interface to its primary interface (upstream) if a Memory address is outside this Address range [provided the address is not within the range defined by the **Memory Base** and **Memory Limit** registers (offset 20h)].*

Register 13-11. 28h Prefetchable Memory Upper Base Address (All Ports)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
31:0	**Prefetchable Memory Base Address[63:32]** The PEX 8508 uses this register for Prefetchable Memory Upper Base Address[63:32]. When the **Prefetchable Memory Base** register *Prefetchable Memory Base Capability* field indicates 32-bit addressing, this register is Read-Only and returns 0000_0000h.	When offset 24h[3:0]=1h	RW	Yes	0000_0000h
		When offset 24h[3:0]=0h	RO	No	0000_0000h

Register 13-12. 2Ch Prefetchable Memory Upper Limit Address (All Ports)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
31:0	**Prefetchable Memory Limit Address[63:32]** The PEX 8508 uses this register for Prefetchable Memory Upper Limit Address[63:32]. When the **Prefetchable Memory Limit** register *Prefetchable Memory Limit Capability* field indicates 32-bit addressing, this register is Read-Only and returns 0000_0000h.	When offset 24h[19:16]=1h	RW	Yes	0000_0000h
		When offset 24h[19:16]=0h	RO	No	0000_0000h

Register 13-13. 30h I/O Upper Base and Limit Address (All Ports)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
15:0	**I/O Base Upper 16 Bits** The PEX 8508 uses this register for I/O Base Address[31:16]. When the **I/O Base** register *I/O Base Addressing Capability* field indicates 16-bit addressing, this register is Read-Only and returns 0000h.	When offset 1Ch[3:0]=1h	RW	Yes	FFFFh
		When offset 1Ch[3:0]=0h	RO	No	0000h
31:16	**I/O Limit Upper 16 Bits** The PEX 8508 uses this register for I/O Limit Address[31:16]. When the **I/O Limit** register *I/O Limit Addressing Capability* field indicates 16-bit addressing, this register is Read-Only and returns 0000h.	When offset 1Ch[11:8]=1h	RW	Yes	0000h
		When offset 1Ch[11:8]=0h	RO	No	0000h

Register 13-14. 34h New Capability Pointer (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**New Capability Pointer** Default 40h points to the **Power Management Capability** register.	RO	Yes	40h
31:8	***Reserved***	RsvdP	No	0000_00h

Register 13-15. 38h Expansion ROM Base Address (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Expansion ROM Base Address** ***Reserved*** Cleared to 0000_0000h.	RsvdP	No	0000_0000h

Register 13-16. 3Ch Bridge Control and Interrupt Signal (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Interrupt Signal			
7:0	**Interrupt Line** Interrupt Line Routing value. The PEX 8508 does ***not*** use this register in Transparent mode; however, the register is included for operating system and device driver use.	RW	Yes	00h
15:8	**Interrupt Pin** Identifies the Conventional PCI Interrupt message(s) that the device (or device function) uses. Only value 00h or 01h is allowed in the PEX 8508. 00h = Indicates that the device does not use Conventional PCI Interrupt message(s) 01h, 02h, 03h, and 04h = Maps to Conventional PCI Interrupt messages for INTA#, INTB#, INTC#, and INTD#, respectively	RO	Yes	01h

Register 13-16. 3Ch Bridge Control and Interrupt Signal (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Bridge Control			
16	**Parity Error Response Enable** Controls the response to Poisoned TLPs. 1 = Enables the secondary *Master Data Parity Error* bit	RW	Yes	0
17	**SERR# Enable** Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface. When set to 1, and the **Command** register *SERR# Enable* bit is set to 1, enables the *Signaled System Error* bit.	RW	Yes	0
18	**ISA Enable** Modifies the bridge response to ISA I/O addresses. Applies only to I/O addresses, enabled by the **I/O Base** and **I/O Limit** registers, that are within the first 64 KB of PCI I/O address space (0000_0000h to 0000_FFFFh). When set to 1, the bridge blocks forwarding of I/O transactions from the primary to secondary interface that address the last 768 bytes in each 1-KB block. In the opposite direction (secondary to primary interface), I/O transactions are forwarded if they address the last 768 bytes in each 1-KB block. The default state of this bit after reset must be 0. 0 = Forward downstream all I/O addresses in the Address range defined by the **I/O Base** and **I/O Limit** registers 1 = Forward upstream ISA I/O addresses in the Address range defined by the **I/O Base** and **I/O Limit** registers that are within the first 64 KB of PCI I/O Address space (top 768 bytes of each 1-KB block)	RW	Yes	0
19	**VGA Enable** Modifies the bridge response to VGA-compatible addresses. When set to 1, the bridge positively decodes and forwards the following addresses on the primary interface to the secondary interface (and, conversely, blocks forwarding of these addresses from the secondary interface to the primary interface): • Memory addresses within the range 000A_0000h to 000B_FFFFh • I/O addresses in the first 64 KB of the I/O Address space (AD[31:16] is 0000h), where AD[9:0] is within the ranges 3B0h to 3BBh and 3C0h to 3DFh (inclusive of ISA address aliases – AD[15:10] is not decoded) When set to 1, forwarding of these addresses is independent of the: • Memory and I/O Address ranges defined by the bridge **I/O Base**, **I/O Limit**, **Memory Base**, **Memory Limit**, **Prefetchable Memory Base**, and **Prefetchable Memory Limit** registers • Bit 18 (*ISA Enable* bit) or **Command** register *VGA Palette Snoop* bit settings Forwarding of these addresses is qualified by the **Command** register *I/O Access Enable* and *Memory Access Enable* bits. The default state of this bit after reset must be 0. 0 = Do not forward VGA-compatible Memory and I/O addresses from the primary to the secondary interface (addresses defined above) unless they are enabled for forwarding by the defined Memory and I/O Address ranges 1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the *I/O Access Enable* and *Memory Access Enable* bits are set), independent of the Memory and I/O Address ranges and independent of the *ISA Enable* bit ***Note:*** *When set in an egress port, configured as a non-Cut-Thru path. (Refer to Section 4.3.4, "Cut-Thru Mode," for further details.)*	RW	Yes	0

Register 13-16. 3Ch Bridge Control and Interrupt Signal (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
20	**VGA 16-Bit Enable** Used only when bit 19 (*VGA Enable*) or the **Command** register *VGA Palette Snoop* bit is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0. Enables system configuration software to select between 10- and 16-bit I/O address decoding, for VGA I/O register accesses forwarded from the primary interface to the secondary interface. 0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses	RW	Yes	0
21	**Master Abort Mode** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
22	**Secondary Bus Reset** 1 = Causes a Hot Reset on the corresponding PEX 8508 port secondary/downstream PCI Bus	RW	Yes	0
23	**Fast Back-to-Back Transactions Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
24	**Primary Discard Timer** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
25	**Secondary Discard Timer** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
26	**Discard Timer Status** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
27	**Discard Timer SERR# Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
31:28	***Reserved***	RsvdP	No	0h

13.7 Power Management Capability Registers

This section details the PEX 8508 Power Management Capability registers. Table 13-5 defines the register map.

Table 13-5. Power Management Capability Register Map (All Ports)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Power Management Capability		Next Capability Pointer (48h)	Capability ID (01h)	40h
Data	Power Management Control/ Status Bridge Extensions	Power Management Status and Control		44h

Register 13-17. 40h Power Management Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Capability ID** Set to 01h, indicating that the data structure currently being pointed to is the Power Management data structure.	RO	Yes	01h
15:8	**Next Capability Pointer** Default 48h points to the **Message Signaled Interrupt Capability** register.	RO	Yes	48h
18:16	**Version** Default 011b indicates compliance with the *PCI Power Mgmt. r1.2*.	RO	Yes	011b
19	**PME Clock** Does not apply to PCI Express. Returns 0.	RsvdP	No	0
20	***Reserved***	RsvdP	No	0
21	**Device-Specific Initialization** Default 0 indicates that Device-Specific Initialization is ***not*** required.	RO	Yes	0
24:22	**AUX Current** Value can be strapped to 111b in the serial EEPROM.	RO	Yes	000b
25	**D1 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D1 Power state.	RsvdP	No	0
26	**D2 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D2 Power state.	RsvdP	No	0
31:27	**PME Support** Default value of 1100_1b indicates that the corresponding PEX 8508 port forwards PME messages in the D0, D3hot, and D3cold power states.	RO	Yes	1100_1b

Register 13-18. 44h Power Management Status and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Power Management Status and Control				
1:0	**Power State** Used to determine the current power state of the port, and to set the port into a new power state. 00b = D0 01b = D1 – ***Not supported*** 10b = D2 – ***Not supported*** 11b = D3hot If software attempts to write an unsupported state to this field, the Write operation completes normally; however, the data is discarded and no state change occurs.	RW	Yes	00b
2	***Reserved***	RsvdP	No	0
3	**No Soft Reset**	RO	Yes	1
7:4	***Reserved***	RsvdP	No	0h
8	**PME Enable** 0 = Disables PME generation by the corresponding PEX 8508 port 1 = Enables PME generation by the corresponding PEX 8508 port	RWS	No	0
12:9	**Data Select** Writable by Serial EEPROM only[a]. Bits [12:9] select the **Data** and **Data Scale** registers. 0h = D0 power consumed 3h = D3hot power consumed 4h = D0 power dissipated 7h = D3hot power dissipated All other values are ***reserved***.	RO	Yes	0h
	Not supported RO for hardware auto-configuration.	RO	No	0h
14:13	**Data Scale** Loaded by serial EEPROM[a]. There are four internal **Data Scale** registers (one each per **Data** register – 0, 3, 4 and 7), per port. Bits [12:9], *Data Select*, select the **Data Scale** register.	RO	Yes	00b
15	**PME Status** 0 = PME is not generated by the corresponding PEX 8508 port 1 = PME is being generated by the corresponding PEX 8508 port	RW1CS	No	0

Register 13-18. 44h Power Management Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Power Management Control/Status Bridge Extensions				
21:16	***Reserved***	RsvdP	No	0-0h
22	**B2/B3 Support** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.	RsvdP	No	0
23	**Bus Power/Clock Control Enable** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.	RsvdP	No	0
Data				
31:24	**Data** Loaded by serial EEPROM[a]. There are four internal **Data** registers (0, 3, 4, and 7), per port. Bits [12:9], *Data Select*, select the **Data** register.	RO	Yes	00h

*a. With no serial EEPROM, Reads return 00h for the **Data Scale** and **Data** registers (for all Data Selects).*

13.8 Message Signaled Interrupt Capability Registers

This section details the PEX 8508 Message Signaled Interrupt (MSI) Capability registers. Table 13-6 defines the register map.

Table 13-6. Message Signaled Interrupt Capability Register Map (All Ports)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Reserved	Message Control	Next Capability Pointer (68h)	Capability ID (05h)	48h
Message Address[31:0]				4Ch
Message Upper Address[63:32]				50h
Reserved		Message Data		54h
Mask Bit for MSI				58h
Reserved			5Ch –	64h

Register 13-19. 48h Message Signaled Interrupt Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
MSI Capability Header				
7:0	**Capability ID** Set to 05h, as required by the *PCI r3.0*.	RO	Yes	05h
15:8	**Next Capability Pointer** Set to 68h to point to the PEX 8508 PCI Express Capability registers.	RO	Yes	68h
Message Control				
16	**MSI Enable** 0 = Message Signaled Interrupts for the corresponding port are disabled 1 = Message Signaled Interrupts for the corresponding port are enabled	RW	Yes	0
19:17	**Multiple Message Capable** 000b = PEX 8508 port is requesting one message – the only value supported	RO	Yes	000b
22:20	**Multiple Message Enable** 000b = PEX 8508 port contains only one allocated message – the only value supported	RW	Yes	000b
23	**MSI 64-Bit Address Capable** 1 = PEX 8508 is capable of generating 64-bit Message Signaled Interrupt addresses	RO	Yes	1
31:24	***Reserved***	RsvdP	No	00h

Register 13-20. 4Ch Message Address[31:0] (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
1:0	***Reserved***	RsvdP	No	00b
31:2	**Message Address[31:2]** MSI Write transaction lower address[31:2]. ***Note:*** *Refer to offset 50h for Message Upper Address[63:32].*	RW	Yes	0000_000h

Register 13-21. 50h Message Upper Address[63:32] (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Message Upper Address[63:32]** MSI Write transaction upper address[63:32]. ***Note:*** *Refer to offset 4Ch for Message Address[31:0].*	RW	Yes	0000_0000h

Register 13-22. 54h Message Data (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Message Data** MSI Write transaction TLP payload.	RW	Yes	0000h
31:16	***Reserved***	RsvdP	No	0000h

Register 13-23. 58h Mask Bit for MSI (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Mask Bit** MSI mask for Hot Plug triggered event.	RW	Yes	0
31:1	***Reserved***	RsvdP	No	0000_000h

13.9 PCI Express Capability Registers

This section details the PEX 8508 PCI Express Capability registers. Hot Plug Capability, Command, Status, and Events are included in these registers. Table 13-7 defines the register map.

Table 13-7. PCI Express Capability Register Map (All Ports)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
PCI Express Capabilities		Next Capability Pointer (00h)	Capability ID (10h)	68h
Device Capabilities				6Ch
Device Status		Device Control		70h
Link Capabilities				74h
Link Status		Link Control		78h
Slot Capabilities				7Ch
Slot Status		Slot Control		80h
Reserved				84h – 8Ch

Register 13-24. 68h PCI Express Capability List and Capabilities (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I^2C	Default
PCI Express Capability List					
7:0	**Capability ID** Set to 10h, as required by the *PCI Express Base r1.1*.		RO	Yes	10h
15:8	**Next Capability Pointer** 00h = PCI Express Capability is the last capability in the PEX 8508 port Capabilities list The PEX 8508 port Extended Capabilities list starts at offset 100h.		RO	Yes	00h
PCI Express Capabilities					
19:16	**Capability Version** The PEX 8508 ports set this field to 1h, as required by the *PCI Express Base r1.1*.		RO	Yes	1h
23:20	**Device/Port Type** Set at reset, as required by the *PCI Express Base r1.1*.	Upstream	RO	Yes	5h
		Downstream	RO	Yes	6h
24	**Slot Implemented** Not valid for the upstream port.	Upstream	RsvdP	No	0
	0 = Disables or connects to an integrated component[a] 1 = Indicates that the downstream port connects to a slot, as opposed to being connected to an integrated component or being disabled	Downstream	RO	Yes	1
29:25	**Interrupt Message Number** The serial EEPROM writes 00_000b, because the Base message and MSI messages are the same.		RO	Yes	00_000b
31:30	***Reserved***		RsvdP	No	00b

a. The PEX 8508 Serial EEPROM register initialization capability is used to change this value to 0h, indicating that the corresponding PEX 8508 downstream port connects to an integrated component or is disabled.

Register 13-25. 6Ch Device Capabilities (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
2:0	**Maximum Payload Size Supported** 000b = PEX 8508 port supports 128-byte maximum payload 001b = PEX 8508 port supports 256-byte maximum payload No other values are supported.		RO	Yes	001b
4:3	**Phantom Functions Supported** ***Not supported*** Cleared to 00b.		RO	Yes	00b
5	**Extended Tag Field Supported** ***Not supported*** 0 = Maximum *Tag* field is 5 bits 1 = Maximum *Tag* field is 8 bits		RO	Yes	0
8:6	**Endpoint L0s Acceptable Latency** ***Not supported*** Because the PEX 8508 is a switch and not an endpoint, it does ***not support*** this feature. 000b = Disables the capability		RO	Yes	000b
11:9	**Endpoint L1 Acceptable Latency** ***Not supported*** Because the PEX 8508 is a switch and not an endpoint, it does ***not support*** this feature. 000b = Disables the capability		RO	Yes	000b
17:12	***Reserved***, as required by the *PCI Express Base r1.1*.		RsvdP	No	00_0000b
25:18	**Captured Slot Power Limit Value** For the PEX 8508 upstream port, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Scale* field.	Upstream	RO	Yes	00h
	Not valid for the downstream ports.	Downstream	RsvdP	No	00h
27:26	**Captured Slot Power Limit Scale** For the PEX 8508 upstream port, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Value* field. 00b = 1.0 01b = 0.1 10b = 0.01 11b = 0.001	Upstream	RO	Yes	00b
	Not valid for the downstream ports.	Downstream	RsvdP	No	00b
31:28	***Reserved***		RsvdP	No	0h

Register 13-26. 70h Device Status and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Device Control			
0	**Correctable Error Reporting Enable** 0 = Disables 1 = Enables the corresponding PEX 8508 port to report Correctable errors	RW	Yes	0
1	**Non-Fatal Error Reporting Enable** 0 = Disables 1 = Enables the corresponding PEX 8508 port to report Non-Fatal errors	RW	Yes	0
2	**Fatal Error Reporting Enable** 0 = Disables 1 = Enables the corresponding PEX 8508 port to report Fatal errors	RW	Yes	0
3	**Unsupported Request Reporting Enable** 0 = Disables 1 = Enables the corresponding PEX 8508 port to report Unsupported Request errors	RW	Yes	0
4	**Enable Relaxed Ordering** ***Not supported*** Cleared to 0.	RsvdP	No	0
7:5	**Maximum Payload Size** Software can change this field to configure the PEX 8508 ports to support other Payload sizes; however, software cannot change this field to a value larger than that indicated by the **Device Capabilities** register *Maximum Payload Size Supported* field value (offset 6Ch[2:0]). 000b = Indicates that initially the PEX 8508 port is configured to support a Maximum Payload Size of 128 bytes 001b = Indicates that initially the PEX 8508 port is configured to support a Maximum Payload Size of 256 bytes No other values are supported.	RW	Yes	000b
8	**Extended Tag Field Enable** ***Not supported*** Cleared to 0.	RsvdP	No	0
9	**Phantom Functions Enable** ***Not supported*** Cleared to 0.	RsvdP	No	0
10	**Auxiliary (AUX) Power PM Enable** Cleared to 0.	RO	No	0
11	**Enable No Snoop** ***Not supported*** Cleared to 0.	RO	No	0
14:12	**Maximum Read Request Size** ***Not supported*** Cleared to 000b.	RsvdP	No	000b
15	***Reserved***	RsvdP	No	0

Register 13-26. 70h Device Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Device Status			
16	**Correctable Error Detected** Set when the corresponding port detects a Correctable error, regardless of the bit 0 (*Correctable Error Reporting Enable* bit) state. 0 = Corresponding PEX 8508 port did not detect a Correctable error 1 = Corresponding PEX 8508 port detected a Correctable error	RW1C	Yes	0
17	**Non-Fatal Error Detected** Set when the corresponding port detects a Non-Fatal error, regardless of the bit 1 (*Non-Fatal Error Reporting Enable* bit) state. 0 = Corresponding PEX 8508 port did not detect a Non-Fatal error 1 = Corresponding PEX 8508 port detected a Non-Fatal error	RW1C	Yes	0
18	**Fatal Error Detected** Set when the corresponding port detects a Fatal error, regardless of the bit 2 (*Fatal Error Reporting Enable* bit) state. 0 = Corresponding PEX 8508 port did not detect a Fatal error 1 = Corresponding PEX 8508 port detected a Fatal error	RW1C	Yes	0
19	**Unsupported Request Detected** Set when the corresponding port detects an Unsupported Request, regardless of the bit 3 (*Unsupported Request Reporting Enable* bit) state. 0 = Corresponding PEX 8508 port did not detect an Unsupported Request 1 = Corresponding PEX 8508 port detected an Unsupported Request	RW1C	Yes	0
20	**Auxiliary (AUX) Power Detected** ***Not supported*** Cleared to 0.	RO	No	0
21	**Transactions Pending** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
31:22	***Reserved***	RsvdP	No	000h

Register 13-27. 74h Link Capabilities (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
3:0	**Maximum Link Speed** Set to 0001b, as required by the *PCI Express Base r1.1*.		RO	Yes	0001b
9:4	**Maximum Link Width** Actual link width is set by STRAP_PORTCFG[4:0]. The PEX 8508 Maximum Link Width is x4 = 00_0100b.		RO	No	Set by Strapping ball levels
11:10	**Active State Power Management (ASPM) Support** Indicates the level of ASPM supported by the port. 01b = L0s link power state entry is supported 10b = L0s and L1 link power states are supported All other values are ***reserved***.		RO	Yes	11b
14:12	**L0s Exit Latency** Indicates the L0s exit latency for the given PCI Express link. The value reported indicates the length of time the corresponding PEX 8508 port requires to complete the transition from L0s to L0. 101b = Corresponding PEX 8508 port L0s Exit Latency is between 1 and 2 µs		RO	No	101b
17:15	**L1 Exit Latency** Indicates the L1 exit latency for the given PCI Express link. The value reported indicates the length of time the corresponding PEX 8508 port requires to complete the transition from L1 to L0. 101b = Corresponding PEX 8508 port L1 Exit Latency is between 16 and 32 µs		RO	Yes	101b
18	**Clock Power Management**		RO	Yes	0
19	**Surprise Down Error Reporting Capable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Valid for the downstream ports only.	Downstream	RO	Yes	1
20	**Data Link Layer Link Active Reporting Capable** Not valid for the upstream port.	Upstream	RsvdP	No	0
	Valid for the downstream ports only.	Downstream	RO	Yes	1
23:21	***Reserved***		RsvdP	No	000b
31:24	**Port Number** The Port Number is set by signal ball Strapping options: STRAP_PORTCFG[4:0] – Ports 0, 1, 2, 3, 4		HwInit	No	Set by Strapping ball levels

Register 13-28. 78h Link Status and Control (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
Link Control					
1:0	**Active State Power Management (ASPM) Control** 00b = Disables L0s and L1 Entry for the corresponding PEX 8508 port[a] 01b = Enables only L0s Entry 10b = Enables L1 Entry 11b = Enables both L0s and L1 Entry		RW	Yes	00b
2	***Reserved***		RsvdP	No	0
3	**Read Completion Boundary (RCB)** Cleared to 0, as required by the *PCI Express Base r1.1*.		RO	Yes	0
4	**Link Disable** Not valid for the upstream port.	Upstream	RsvdP	No	0
	Setting to 1 places the link on the corresponding PEX 8508 downstream port to the Disabled Link Training state.	Downstream	RW	Yes	0
5	**Retrain Link** Not valid for the upstream port.	Upstream	RsvdP	No	0
	For PEX 8508 ports, always returns 0 when read. Writing 1 to this bit causes the corresponding PEX 8508 downstream port to initiate retraining of its PCI Express link.	Downstream	RZ	Yes	0
6	**Common Clock Configuration** 0 = Corresponding PEX 8508 port and the device at the opposite end of the corresponding port's PCI Express link are operating with an **asynchronous** Reference Clock 1 = Corresponding PEX 8508 port and the device at the opposite end of the corresponding port's PCI Express link are operating with a **distributed** common Reference Clock		RW	Yes	0
7	**Extended Sync** Set to 1 causes the corresponding PEX 8508 port to transmit: • 4,096 FTS Ordered-Sets in the L0s state, • Followed by a single SKIP Ordered-Set prior to entering the L0 state, • Finally, transmission of 1,024 TS1 Ordered-Sets in the Recovery state.		RW	Yes	0
8	**Clock Power Management Enable** The PEX 8508 does ***not support*** removal of the Reference Clock in the L1 and L2/L3 Ready Link states.		RW	Yes	0
15:9	***Reserved***		RsvdP	No	00h

Register 13-28. 78h Link Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
	Link Status				
19:16	**Link Speed** Set to 0001b, as required by the *PCI Express Base r1.1* for a 2.5 Gbps PCI Express link.		RO	Yes	0001b
25:20	**Negotiated Link Width** Link width is determined by negotiated value with attached port/lane. 00_0001b = x1 00_0010b = x2 00_0100b = x4 All other values are ***not supported***. The value in this field is undefined when the link is not up.		RO	Yes	00_0001b
26	***Undefined***		RsvdP	No	0
27	**Link Training** ***Reserved*** for the upstream ports.	Upstream	RsvdP	No	0
	When set to 1, indicates that the corresponding PEX 8508 downstream port requested link training and either the link training is in process or about to start.	Downstream	RO	No	0
28	**Slot Clock Configuration** Set by the upstream port. 0 = Indicates that the PEX 8508 uses an independent clock 1 = Indicates that the PEX 8508 uses the same physical Reference Clock that the platform provides on the connector		HwInit	Yes	0
29	**Data Link Layer Link Active** When set to 1, indicates the following: • Data Link Layer is in the DL_Active state • Link is operational • Flow Control Initialization has successfully completed		RO	No	0
31:30	***Reserved***		RsvdP	No	00b

a. The port receiver must be capable of entering the L0s Link PM state, regardless of whether the state is disabled.

Register 13-29. 7Ch Slot Capabilities (Only Downstream Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I^2C	Default
0	**Attention Button Present** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Attention Button is not implemented 1 = Attention Button is implemented on the slot chassis of the corresponding PEX 8508 downstream port	Downstream	HwInit	Yes	1
1	**Power Controller Present** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Power Controller is not implemented 1 = Power Controller is implemented for the slot of the corresponding PEX 8508 downstream port	Downstream	HwInit	Yes	1
2	**MRL Sensor Present** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = MRL Sensor is not implemented 1 = MRL Sensor is implemented on the slot chassis of the corresponding PEX 8508 downstream port	Downstream	HwInit	Yes	1
3	**Attention Indicator Present** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Attention Indicator is not implemented 1 = Attention Indicator is implemented on the slot chassis of the corresponding PEX 8508 downstream port	Downstream	HwInit	Yes	1
4	**Power Indicator Present** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Power Indicator i*s* not implemented 1 = Power Indicator is implemented on the slot chassis of the corresponding PEX 8508 downstream port	Downstream	HwInit	Yes	1
5	**Hot Plug Surprise** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = No device in the corresponding PEX 8508 downstream port slot is removed from the system without prior notification 1 = Device in the corresponding PEX 8508 downstream port slot can be removed from the system without prior notification	Downstream	HwInit	Yes	0
6	**Hot Plug Capable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Corresponding PEX 8508 downstream port slot is not capable of supporting Hot Plug operations 1 = Corresponding PEX 8508 downstream port slot is capable of supporting Hot Plug operations	Downstream	HwInit	Yes	1

Register 13-29. 7Ch Slot Capabilities (Only Downstream Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
14:7	**Slot Power Limit Value** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	00h
	The maximum power supplied by the corresponding PEX 8508 downstream slot is determined by multiplying the value in this field (expressed in decimal; 25d = 19h) by the value specified by the *Slot Power Limit Scale* field. This field must be implemented if the **PCI Express Capabilities** register *Slot Implemented* bit is set (offset 68h[24])=1, default). Writes to this register also cause the downstream port to send the Set_Slot_Power_Limit message to the device connected to it, so as to convey the Limit value to the downstream device's upstream port **Device Capabilities** register *Captured Slot Power Limit Value* and *Captured Slot Power Limit Scale* fields.	Downstream	HwInit	Yes	19h
16:15	**Slot Power Limit Scale** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	00b
	The maximum power supplied by the corresponding PEX 8508 downstream slot is determined by multiplying the value in this field by the *Slot Power Limit Value* bit. This field must be implemented if the **PCI Express Capabilities** register *Slot Implemented* bit is set (offset 68h[24])=1, default). Writes to this register also cause the downstream port to send the Set_Slot_Power_Limit message to the device connected to it, so as to convey the Limit value to the downstream device's upstream port **Device Capabilities** register *Captured Slot Power Limit Value* and *Captured Slot Power Limit Scale* fields. 00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x	Downstream	HwInit	Yes	00b
17	**Electromechanical Interlock Present** When set to 1, indicates that an Electromechanical Interlock is implemented on the chassis for this slot.		RsvdP	No	0
18	**No Command Completed Support** When set to 1, indicates that this slot does not generate software notification when an issued command is completed by the Hot Plug Controller. Allowed to be set to 1 only when the port is able to accept Writes to all fields of the **Slot Control** register, without delay between successive Writes.		RsvdP	No	0
31:19	**Physical Slot Number** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0-0h
	Indicates the physical slot number attached to this port. If the **PCI Express Capabilities** register *Slot Implemented* bit is set (offset 68h[24])=1, default), this field must be hardware-initialized to a value that assigns a slot number that is unique within the chassis, regardless of the form factor associated with the slot. Must be initialized to 0h for ports connected to devices that are integrated on the system board.	Downstream	HwInit	Yes	0-0h

Register 13-30. 80h Slot Status and Control (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I^2C	Default
Slot Control					
0	**Attention Button Pressed Enable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Power state (**Power Management Status and Control** register *Power State* field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Power state (offset 44h[1:0]=11b), for an Attention Button Pressed event on the corresponding Hot Plug-capable PEX 8508 downstream port	Downstream	RW	Yes	0
1	**Power Fault Detector Enable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Power state (**Power Management Status and Control** register *Power State* field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Power state (offset 44h[1:0]=11b), for a Power Fault Detected event on the corresponding Hot Plug-capable PEX 8508 downstream port	Downstream	RW	Yes	0

Register 13-30. 80h Slot Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
2	**MRL Sensor Changed Enable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Power state (**Power Management Status and Control** register *Power State* field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Power state (offset 44h[1:0]=11b), for an MRL Sensor Changed event on the corresponding Hot Plug-capable PEX 8508 downstream port	Downstream	RW	Yes	0
3	**Presence Detect Changed Enable** Not valid for the upstream port.	Upstream	RsvdP	No	0
	0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Power state (**Power Management Status and Control** register *Power State* field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Power state (offset 44h[1:0]=11b), for a Presence Detect Changed event on the corresponding Hot Plug-capable PEX 8508 downstream port A Presence Detect Changed event is triggered from one of two sources, depending upon the state of the *HPC Inband Presence Detect Enable* bit (offset 1E0h[5]): • If the *HPC Inband Presence Detect Enable* bit is cleared (offset 1E0h[5]=0, default), Presence Detect is input from the HP_PRSNT*x*# signal on the corresponding PEX 8508 Hot Plug-capable downstream port • If the *HPC Inband Presence Detect Enable* bit is set (offset 1E0h[5]=1), Presence Detect is input from the SerDes Receiver Detect on the corresponding PEX 8508 downstream port	Downstream	RW	Yes	0

Register 13-30. 80h Slot Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I^2C	Default
4	**Command Completed Interrupt Enable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt when a command is completed by the Hot Plug Controller on the corresponding Hot Plug-capable PEX 8508 downstream port	Downstream	RW	Yes	0
5	**Hot Plug Interrupt Enable** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	0 = Function is disabled 1 = Enables a Hot Plug Interrupt on enabled Hot Plug events for the corresponding PEX 8508 downstream port	Downstream	RW	Yes	0
7:6	**Attention Indicator Control** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	00b
	Controls the Attention Indicator on the corresponding PEX 8508 downstream port slot. 00b = ***Reserved*** – Writes are ignored 01b = Turns On indicator to constant On state 10b = Causes indicator to blink 11b = Turns Off indicator Software must use a Byte or Word Write (and not a DWord Write) to control the HP_ATNLED*x*# Output signal. Reads return the corresponding PEX 8508 downstream port Attention Indicator's current state.	Downstream	RW	Yes	11b

Register 13-30. 80h Slot Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
9:8	**Power Indicator Control** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	00b
	Controls the Power Indicator on the corresponding PEX 8508 downstream port slot. 00b = ***Reserved*** – Writes are ignored 01b = Turns On indicator to constant On state 10b = Causes indicator to blink 11b = Turns Off indicator Software must use a Byte or Word Write (and not a DWord Write) to control the HP_PWRLED*x*# Output signal. Reads return the corresponding PEX 8508 downstream port Power Indicator's current state.	Downstream	RW	Yes	11b (MRL open) 01b (MRL closed)
10	**Power Controller Control** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Controls the Power Controller on the corresponding PEX 8508 downstream port slot. 0 = Turns On the Power Controller; requires some delay to be effective 1 = Turns Off the Power Controller Software must use a Byte or Word Write (and not a DWord Write) to control the Power Controller Output signals.	Downstream	RW	Yes	1 (MRL open) 0 (MRL closed)
11	**Electromechanical Interlock Control** When an Electromechanical Interlock is implemented, this bit indicates the current status of the Electromechanical Interlock. 0 = Electromechanical Interlock is disengaged 1 = Electromechanical Interlock is engaged		RsvdP	No	0
12	**Data Link Layer State Changed Enable** Not valid for the upstream port.	Upstream	RsvdP	No	0
	Enables software notification with a Hot Plug interrupt if the port is in the D0 Power state (**Power Management Status and Control** register *Power State* field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Power state (offset 44h[1:0]=11b), when the **Link Status** register *Data Link Layer Link Active* bit (offset 78h[29]) is changed.	Downstream	RW	Yes	0
15:13	***Reserved***		RsvdP	No	000b

Register 13-30. 80h Slot Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I^2C	Default
Slot Status					
16	**Attention Button Pressed** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Set to 1 when the Attention Button of the corresponding PEX 8508 downstream port slot is pressed.	Downstream	RW1C	Yes	0
17	**Power Fault Detected** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Set to 1 when the Power Controller of the corresponding PEX 8508 downstream port slot detects a Power Fault at the slot.	Downstream	RW1C	Yes	0
18	**MRL Sensor Changed** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Set to 1 when an MRL state change is detected on the corresponding PEX 8508 downstream port slot.	Downstream	RW1C	Yes	0
19	**Presence Detect Changed** Not valid for the upstream port.	Upstream	RsvdP	No	0
	Set to 1 when a Presence Detect Change is detected on the corresponding PEX 8508 downstream port slot. A Presence Detect Changed event is triggered from one of two sources, depending upon the state of the *HPC Inband Presence Detect Enable* bit (offset 1E0h[5]): • If the *HPC Inband Presence Detect Enable* bit is cleared (offset 1E0h[5]=0, default), Presence Detect is input from the HP_PRSNT*x*# signal on the corresponding Hot Plug-capable PEX 8508 downstream port • If the *HPC Inband Presence Detect Enable* bit is set (offset 1E0h[5]=1), Presence Detect is input from the SerDes Receiver Detect on the corresponding PEX 8508 downstream port	Downstream	RW1C	Yes	0

Register 13-30. 80h Slot Status and Control (All Ports) *(Cont.)*

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
20	**Command Completed** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Set to 1 when the Hot Plug Controller on the corresponding PEX 8508 downstream port slot completes an issued command. ***Note:*** *RW1C functionality is dependent on bits [10:6] of this register.*	Downstream	RW1C	Yes	0
21	**MRL Sensor State** ***Reserved*** for the upstream port.	Upstream	RsvdP	No	0
	Reveals the corresponding PEX 8508 downstream port MRL Sensor's current state. 0 = MRL Sensor is closed 1 = MRL Sensor is open	Downstream	RO	Yes	0
22	**Presence Detect State** Not valid for the upstream port.	Upstream	RsvdP	No	0
	Reveals the corresponding PEX 8508 downstream port's current Presence state. Presence is determined from one of two sources, depending upon the state of the *HPC Inband Presence Detect Enable* bit (offset 1E0h[5]): • If the *HPC Inband Presence Detect Enable* bit is cleared (offset 1E0h[5]=0, default), Presence Detect is input from the HP_PRSNT*x*# signal on the corresponding Hot Plug-capable PEX 8508 downstream port • If the *HPC Inband Presence Detect Enable* bit is set (offset 1E0h[5]=1), Presence Detect is input from the SerDes Receiver Detect on the corresponding PEX 8508 downstream port 0 = Slot is empty, or device is not present 1 = Slot is occupied, or device is present	Downstream	RO	Yes	0
23	**Electromechanical Interlock Status** When an Electromechanical Interlock is implemented, indicates the Electromechanical Interlock's current status. 0 = Electromechanical Interlock is disengaged 1 = Electromechanical Interlock is engaged		RsvdZ	No	0
24	**Data Link Layer State Changed** Not valid for the upstream port.	Upstream	RsvdP	Yes	0
	Set when the value reported in the **Link Status** register *Data Link Layer Link Active* bit changes. In response to a Data Link Layer State Changed event, software must read the **Link Status** register *Data Link Layer Link Active* bit, to determine whether the link is active before initiating Configuration requests to the Hot Plugged device.	Downstream	RW1C	Yes	0
31:25	***Reserved***		RsvdZ	No	0-0h

13.10 Device Serial Extended Number Capability Registers

This section details the PEX 8508 Device Serial Number Extended Capability registers. Table 13-8 defines the register map.

Table 13-8. PEX 8508 Device Serial Number Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (FB4h)	Capability Version (1h) / PCI Express Extended Capability ID (0003h)	100h
Serial Number (Lower DW)		104h
Serial Number (Upper DW)		108h
Reserved	10Ch –	134h

Register 13-31. 100h Device Serial Number Enhanced Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**PCI Express Extended Capability ID** Set to 0003h, as required by the *PCI Express Base r1.1*.	RO	Yes	0003h
19:16	**Capability Version** Set to 1h, as required by the *PCI Express Base r1.1*.	RO	Yes	1h
31:20	**Next Capability Offset** Set to FB4h, which is the **Advanced Error Reporting Enhanced Capability Header** registers.	RO	Yes	FB4h

Register 13-32. 104h Serial Number (Lower DW) (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**PCI Express Device Serial Number (1st DW)** Lower half of a 64-bit register. Value set by Serial EEPROM register initialization.	RO	Yes	0000_0EDFh

Register 13-33. 108h Serial Number (Upper DW) (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**PCI Express Device Serial Number (2nd DW)** Upper half of a 64-bit register. Value set by Serial EEPROM register initialization.	RO	Yes	0000_0001h

13.11 Power Budgeting Extended Capability Registers

This section details the PEX 8508 Power Budgeting Extended Capability registers. Table 13-9 defines the register map.

Table 13-9. PEX 8508 Power Budgeting Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (148h) / Capability Version (1h)	PCI Express Extended Capability ID (0004h)	138h
Reserved	Data Select	13Ch
Power Budgeting Data		140h
Power Budget Capability		144h

Register 13-34. 138h Power Budgeting Enhanced Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**PCI Express Extended Capability ID** Set to 0004h, as required by the *PCI Express Base r1.1*.	RO	Yes	0004h
19:16	**Capability Version** Set to 1h, as required by the *PCI Express Base r1.1*.	RO	Yes	1h
31:20	**Next Capability Offset** Set to 148h, which addresses the PEX 8508 **Virtual Channel Enhanced Capability Header** registers.	RO	Yes	148h

Register 13-35. 13Ch Data Select (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Data Select** Indexes the Power Budgeting Data reported, by way of eight registers per port, and selects the DWord of Power Budgeting Data that appears in each **Power Budgeting Data** register. Index values commence at 0 to select the first DWord of Power Budgeting Data; subsequent DWords of Power Budgeting Data are selected by increasing index values 1 to 7.	RW	Yes	00h
31:8	***Reserved***	RsvdP	No	0-0h

Register 13-36. 140h Power Budgeting Data (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Base Power** Eight registers/port. Specifies (in Watts) the base power value in the operating condition. This value must be multiplied by the *Data Scale* (bits [9:8]) to produce the actual power consumption value.	RO	Yes	00h
9:8	**Data Scale** Specifies the scale to apply to the Base Power value. The power consumption of the device is determined by multiplying the *Base Power* field (bits [7:0]) contents with the value corresponding to the encoding returned by this field. 00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x	RO	Yes	00b
12:10	**PM Sub-State** 000b = Corresponding PEX 8508 port is in the default Power Management sub-state	RO	Yes	000b
14:13	**PM State** Current power state. 00b = D0 state 11b = D3 state All other values are ***reserved***.	RO	Yes	00b
17:15	**Type** Type of operating condition. 000b = PME Auxiliary 001b = Auxiliary 010b = Idle 011b = Sustained 111b = Maximum All other values are ***reserved***.	RO	Yes	000b
20:18	**Power Rail** Power Rail of operating condition. 000b = Power 12V 001b = Power 3.3V 010b = Power 1.8V 111b = Thermal All other values are ***reserved***.	RO	Yes	000b
31:21	***Reserved***	RsvdP	No	0-0h

Note: *There are eight registers per port that can be programmed, through the serial EEPROM. Each non-zero register value describes the power usage for a different operating condition. Each configuration is selected by writing to the* ***Data Select*** *register Data Select field (offset 13Ch[7:0]).*

Register 13-37. 144h Power Budget Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**System Allocated** 1 = Power budget for the device is included within the system power budget	HwInit	Yes	1
31:1	***Reserved***	RsvdP	No	0-0h

13.12 Virtual Channel Extended Capability Registers

This section details the PEX 8508 Virtual Channel Extended Capability registers. These registers are duplicated for each port. Table 13-10 defines the register map for one port.

Table 13-10. PEX 8508 Virtual Channel Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (000h) / Capability Version (1h)	PCI Express Extended Capability ID (0002h)	148h
Port VC Capability 1		14Ch
Port VC Capability 2		150h
Port VC Status ***(Reserved)***	Port VC Control	154h
VC0 Resource Capability		158h
VC0 Resource Control		15Ch
VC0 Resource Status	***Reserved***	160h
VC1 Resource Capability (Only Ports 0, 1, and 2) ***Reserved*** (Only Ports 3 and 4)		164h
VC1 Resource Control (Only Ports 0, 1, and 2) ***Reserved*** (Only Ports 3 and 4)		168h
VC1 Resource Status (Only Ports 0, 1, and 2) ***Reserved*** (Only Ports 3 and 4)	***Reserved***	16Ch
Reserved		170h – 1C4h

Register 13-38. 148h Virtual Channel Enhanced Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
15:0	**PCI Express Extended Capability ID** Set to 0002h, as required by the *PCI Express Base r1.1*.	RO	Yes	0002h
19:16	**Capability Version** Set to 1h, as required by the *PCI Express Base r1.1*.	RO	Yes	1h
31:20	**Next Capability Offset** Cleared to 000h, indicating that the Virtual Channel Extended Capability is the last extended capability in the port Extended Capability list.	RO	Yes	000h

Register 13-39. 14Ch Port VC Capability 1 (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
0	**Extended VC Count** 0 = PEX 8508 port supports only the default Virtual Channel 0 (VC0) 1 = PEX 8508 port supports one extended virtual channel, Virtual Channel 1 (VC1)	0, 1, 2	RO	Yes	1
	0 = PEX 8508 port supports only the default Virtual Channel 0 (VC0) 1 = ***Reserved***	3, 4	RO	Yes	1
3:1	***Reserved***		RsvdP	No	000b
4	**Low-Priority Extended VC Count** For strict priority arbitration, this bit indicates the number of extended virtual channels (those in addition to the default VC0) that belong to the Low-Priority Virtual Channel group for the PEX 8508 port. PEX 8508 Serial EEPROM register initialization capability is used to change this bit to 1 to also set VC1 to the Low-Priority Virtual Channel group. 0 = For this PEX 8508 port, only the default VC0 belongs to the Low-priority Virtual Channel group 1 = For this PEX 8508 port, VC0 and VC1 belong to the Low-priority Virtual Channel group	0, 1, 2	RO	Yes	0
	For strict priority arbitration, this bit indicates the number of extended virtual channels (those in addition to the default VC0) that belong to the Low-Priority Virtual Channel group for the PEX 8508 port. 0 = For this PEX 8508 port, VC0 belongs to the Low-priority Virtual Channel group 1 = ***Reserved***	3, 4	RO	Yes	0
7:5	***Reserved***		RsvdP	No	000b
9:8	**Reference Clock** Cleared to 00b.	0, 1, 2, 3, 4	RsvdP	No	00b
11:10	**Port Arbitration Table Entry Size** Cleared to 00b, as the size of the Port Arbitration Table is 1 bit.	0, 1, 2, 3, 4	RsvdP	No	00b
31:12	***Reserved***		RsvdP	No	0-0h

Register 13-40. 150h Port VC Capability 2 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**VC Arbitration Capability** Bit 0 value of 1 indicates Round-Robin (Hardware-Fixed) Arbitration scheme is supported.	RO	Yes	1
31:1	***Reserved***	RsvdP	No	0-0h

Register 13-41. 154h Port VC Status and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Port VC Control				
0	***Reserved***	RsvdP	No	0
1	**VC Arbitration Select** ***Not used, because the PEX 8508 supports only Hardware-Fixed arbitration. It is recommended that this bit remain programmed to 0.*** Selects the VC arbitration type for the corresponding PEX 8508 port. Indicates the bit number in the **Port VC Capability 2** register *VC Arbitration Capability* field that corresponds to the arbitration type. 0 = bit 0; Round-Robin (Hardware-Fixed) arbitration scheme 1 = ***Reserved*** Select only an arbitration type that corresponds to a bit set in the *VC Arbitration Capability* field. Cannot be modified when more than one lower priority Virtual Channel group VC is enabled.	RW	Yes	0
15:2	***Reserved***	RsvdP	No	0-0h
Port VC Status				
31:16	***Reserved***	RsvdP	No	0-0h

Register 13-42. 158h VC0 Resource Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Port Arbitration Capability** 1 = Non-configurable hardware-fixed port arbitration Non-configurable hardware-fixed port arbitration is the only configuration supported by the PEX 8508.	RO	Yes	1
14:1	***Reserved***	RsvdP	No	0-0h
15	**Reject Snoop Transactions** Not a PCI Express switch feature; therefore, cleared to 0.	RsvdP	No	0
22:16	**Maximum Time Slots** ***Not supported*** Cleared to 000_0000b.	RsvdP	No	000_0000b
23	***Reserved***	RsvdP	No	0
31:24	**Port Arbitration Table Offset** Cleared to 00h, as the Port Arbitration Table is not present.	RsvdP	No	00h

Register 13-43. 15Ch VC0 Resource Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC/VC0 Map** Defines Traffic Classes [7:0], respectively, and indicates which TCs are mapped into Virtual Channel 0. Traffic Class 0 (TC0) must be mapped to Virtual Channel 0. By default, Traffic Classes [7:1] are mapped to VC0. ***Note:*** *Bit 0 must be set to 1 for VC0.*	RO	Yes	1
7:1		RW	Yes	7Fh
15:8	***Reserved***	RsvdP	No	00h
16	**Load Port Arbitration Table** Cleared to 0 because the Port Arbitration Table is not used to select a Port Arbitration scheme.	RsvdP	No	0
19:17	**Port Arbitration Select** Cleared to 000b because the PEX 8508 uses a non-configurable Hardware Fixed-Port Arbitration.	RsvdP	No	000b
23:20	***Reserved***	RsvdP	No	0-0h
24	**VC0 ID** Defines the corresponding PEX 8508 port Virtual Channel 0 ID code. Because this is the default VC, it is cleared to 0.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	0-0h
31	**VC0 Enable** 0 = Not allowed 1 = Enables the corresponding PEX 8508 port default Virtual Channel 0	RO	Yes	1

Register 13-44. 160h VC0 Resource Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	***Reserved***	RsvdP	No	00000h
16	**Port Arbitration Table Status** By default, cleared to 0 because the PEX 8508 uses a non-configurable Hardware Fixed-Port Arbitration.	RsvdP	No	0
17	**VC0 Negotiation Pending** 0 = VC0 negotiation completed 1 = VC0 initialization is not complete for the corresponding PEX 8508 port	RO	Yes	1
31:18	***Reserved***	RsvdP	No	0-0h

Register 13-45. 164h VC1 Resource Capability (Only Ports 0, 1, and 2)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Port Arbitration Capability** 1 = Non-configurable hardware-fixed port arbitration The only configuration supported by PEX 8508.	RO	Yes	1
14:1	***Reserved***	RsvdP	No	0-0h
15	**Reject Snoop Transactions** Not a PCI Express switch feature; therefore, this bit is cleared to 0.	RsvdP	No	0
22:16	**Maximum Time Slots** ***Not supported*** Cleared to 000_0000b.	RsvdP	No	000_0000b
23	***Reserved***	RsvdP	No	0
31:24	**Port Arbitration Table Offset** Cleared to 0, the Port Arbitration Table is not present.	RsvdP	No	00h

Register 13-46. 168h VC1 Resource Control (Only Ports 0, 1, and 2)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC/VC1 Map[0]** ***Reserved***	RsvdP	No	0
7:1	**TC/VC1 Map[7:1]** Defines Traffic Classes [7:1], respectively, and indicates which TCs are mapped into Virtual Channel 1. By default, Traffic Classes [7:1] are mapped to VC0. ***Note:*** *Bit 1 must be set to 1 for VC0.*	RW	Yes	00h
15:8	***Reserved***	RsvdP	No	0-0h
16	**Load Port Arbitration Table** Cleared to 0 because the Port Arbitration Table is not used to select a Port Arbitration scheme.	RsvdP	No	0
19:17	**Port Arbitration Select** Cleared to 0 because the PEX 8508 uses a non-configurable Hardware Fixed-Port Arbitration.	RsvdP	No	000b
23:20	***Reserved***	RsvdP	No	0-0h
24	**VC1 ID** Defines the ID code for the corresponding PEX 8508 port Virtual Channel 1 (1 is the only supported value).	RW	Yes	1
30:25	***Reserved***	RsvdP	No	00_0000b
31	**VC1 Enable** 0 = Disables the corresponding PEX 8508 port Virtual Channel 1 1 = Enables the corresponding PEX 8508 port Virtual Channel 1	RW	Yes	0

Register 13-47. 16Ch VC1 Resource Status (Only Ports 0, 1, and 2)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
15:0	***Reserved***	RsvdP	No	0000h
16	**Port Arbitration Table Status** By default, cleared to 0 because the PEX 8508 uses a non-configurable Hardware Fixed-Port Arbitration.	RsvdP	No	0
17	**VC1 Negotiation Pending** 0 = VC1 negotiation completed 1 = VC1 initialization or disabling is pending for the corresponding PEX 8508 port	RO	Yes	0
31:18	***Reserved***	RsvdP	No	0-0h

13.13 Device-Specific Registers

Table 13-11 defines the Device-Specific register set – registers that are unique to the PEX 8508 and not referenced in the *PCI Express Base r1.1*.

Note: *This register group is accessed using a Memory-Mapped cycle.* ***It is recommended that these register values not be changed.***

Table 13-11. Device-Specific Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Error Checking and Debug Registers	1C8h … 1FCh
Physical Layer Registers	200h … 28Ch
I2C Interface Register	290h … 2C4h
Bus Number CAM Registers	2C8h … 304h
I/O CAM Registers	308h … 344h
Address-Mapping CAM (AMCAM) Registers	348h … 548h
Reserved	54Ch – 65Ch
Ingress Control and Port Enable Registers	660h … 67Ch
I/O CAM Base and Limit Upper 16-Bit Registers	680h … 6BCh
Base Address Shadow Registers (BARs)	6C0h … 73Ch
Shadow Virtual Channel (VC) Capability Registers	740h … 83Ch

Table 13-11. Device-Specific Register Map *(Cont.)*

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Shadow Port Virtual Channel (VC) Capability Registers				840h … 9ECh
Ingress Credit Handler (INCH) Registers				9F0h … BDCh
Relaxed Ordering and Performance Counter Registers				BE0h … BFCh
Internal Credit Handler (ITCH) VC&T Threshold Registers				C00h … C08h
Port Virtual Channel Queue Status Registers				C0Ch … C38h

13.13.1 Error Checking and Debug Registers

Table 13-12. Device-Specific Error Checking and Debug Register Map (Ports[a])

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
ECC Error Check Disable				1C8h
Error Handler 32-Bit Error Status ***(Factory Test Only)***				1CCh
Error Handler 32-Bit Error Mask ***(Factory Test Only)***				1D0h
Factory Test Only				1D4h – 1D8h
Debug Control				1DCh
Power Management Hot Plug User Configuration				1E0h
Egress Control and Status				1E4h
Reserved			Bad TLP Count	1E8h
Reserved			Bad DLLP Count	1ECh
Device-Specific Relaxed Ordering Enable				1F0h
Reserved		Software-Controlled Lane Status		1F4h
Reserved	ACK Transmission Latency Limit			1F8h
Reserved				1FCh

a. Certain registers are port-specific, while others are device-specific.

Register 13-48. 1C8h ECC Error Check Disable (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**ECC 1-Bit Error Check Disable** 0 = RAM 1-Bit Soft Error Check is enabled 1 = Disables RAM 1-Bit Soft Error Check	RW	Yes	0
1	**ECC 2-Bit Error Check Disable** 0 = RAM 2-Bit Soft Error Check is enabled 1 = Disables RAM 2-Bit Soft Error Check	RW	Yes	0
31:2	***Reserved***	RsvdP	No	0-0h

Note: *All errors in register offset 1CCh generate MSI/INTx interrupts, when enabled.*

Register 13-49. 1CCh Error Handler 32-Bit Error Status (*Factory Test Only*)

Bit(s)	Description	Bit Exists Only on Port(s)	Type	Serial EEPROM and I^2C	Default
0	**Completion FIFO Overflow Status** 0 = No overflow detected 1 = Completion FIFO Overflow detected when the 4-deep Completion FIFO (ingress) or 2-deep Completion FIFO (egress) overflows	0, 1, 2, 3, 4	RW1CS	Yes	0
1	**Egress PRAM Soft Error Overflow** Egress Packet RAM 1-bit Soft Error Counter Overflow. 0 = No error detected 1 = Egress PRAM 1-bit Soft-Error (8-bit counter) overflow; when destination packet RAM 1-bit Soft Error Count is greater than or equal to 256, generates an MSI/INT*x* interrupt, when enabled	0	RW1CS	Yes	0
2	***Reserved***		RsvdP	No	0
3	**Egress PRAM ECC Error** Egress Packet RAM 2-bit error detection. 0 = No error detected 1 = Egress PRAM 2-bit ECC error detected	0	RW1CS	Yes	0
4	***Reserved***		RsvdP	No	0
5	**Ingress RAM 1-Bit ECC Error** Source Packet RAM 1-bit soft error detection. 0 = No error detected 1 = Ingress RAM 1-BIT ECC error detected	0	RW1CS	Yes	0
7:6	***Reserved***		RsvdP	No	00b
8	**Ingress RAM Uncorrectable ECC Error** Ingress Packet RAM 2-bit Error detection. 0 = No 2-bit error detected 1 = Packet RAM Uncorrectable ECC error detected	0	RW1CS	Yes	0
9	**Ingress LLIST 1-Bit ECC Error** Ingress Link-List RAM 1-bit soft error detection. 0 = No error detected 1 = 1-bit ECC error detected	0	RW1CS	Yes	0
10	**Ingress LLIST Uncorrectable ECC Error** Ingress packet Link-List RAM 2-bit error detection. 0 = No 2-bit error detected 1 = Ingress Link-List Uncorrectable ECC Error detected	0	RW1CS	Yes	0
11	**Credit Update Timeout Status** No useful credit update to make forward progress for 512 ms or 1s (disabled by default). 0 = No Credit Update Timeout detected 1 = Credit Update Timeout completed	0, 1, 2, 3, 4	RW1CS	Yes	0

Register 13-49. 1CCh Error Handler 32-Bit Error Status (*Factory Test Only*) *(Cont.)*

Bit(s)	Description	Bit Exists Only on Port(s)	Type	Serial EEPROM and I²C	Default
12	**INCH Underrun Error** Ingress Credit Underrun. 0 = No error detected 1 = Credit Underrun error detected	0, 1, 2, 3, 4	RW1CS	Yes	0
13	**Ingress Memory Allocation Unit 1-Bit Soft Error Counter Overflow** Ingress Memory Allocation/De-allocation RAM 1-bit Soft Error Count is greater than or equal to 8. 0 = No error detected 1 = 1-bit Soft Error Counter is ≥ 8	0	RW1CS	Yes	0
14	**Ingress Memory Allocation Unit 2-Bit Soft Error** Ingress Memory Allocation/De-allocation RAM 2-bit error detection for Transaction Layer Ingress Memory Allocation/De-allocation unit. 0 = No error detected 1 = 2-bit Soft error detected	0	RW1CS	Yes	0
15	**NT Port DL_Down Virtual Interface Status** 0 = No DL_Down event has occurred 1 = DL_Down event occurred on the NT Port Link Interface ***Note:*** *Not valid in Transparent mode.*	NT Port Virtual Interface	RW1CS	Yes	0
17:16	***Factory Test Only***		RW1CS	Yes	11b
31:18	***Reserved***		RsvdP	No	0-0h

Note: *Error logging is enabled in register offset 1D0h by default.*

Register 13-50. 1D0h Error Handler 32-Bit Error Mask *(Factory Test Only)*

Bit(s)	Description	Bit Exists Only on Port(s)	Type	Serial EEPROM and I²C	Default
0	**Completion FIFO Overflow Mask** 0 = When enabled, error generates MSI/INT*x* interrupt 1 = *Completion FIFO Overflow Status* bit is masked/disabled	0, 1, 2, 3, 4	RWS	Yes	1
1	**Egress PRAM Soft Error Overflow Mask** 0 = No effect on reporting activity 1 = *Egress PRAM Soft Error Overflow* bit is masked/disabled	0	RWS	Yes	1
2	***Reserved***		RsvdP	No	0
3	**Egress PRAM ECC Error Mask** 0 = No effect on reporting activity 1 = *Egress PRAM ECC Error* bit is masked/disabled	0	RWS	Yes	1
4	***Reserved***		RsvdP	No	0
5	**Ingress RAM 1-Bit ECC Error Mask** 0 = No effect on reporting activity 1 = *Ingress RAM 1-Bit ECC Error* bit is masked/disabled	0	RWS	Yes	1
7:6	***Reserved***		RsvdP	No	00b
8	**Ingress RAM Uncorrectable ECC Error Mask** 0 = No effect on reporting activity 1 = *Ingress RAM Uncorrectable ECC Error* bit is masked/disabled	0	RWS	Yes	1
9	**Ingress LLIST 1-Bit ECC Error Mask** 0 = No effect on reporting activity 1 = *Ingress RAM 1-Bit ECC Error* bit is masked/disabled	0	RWS	Yes	1
10	**Ingress LLIST Uncorrectable ECC Error Mask** 0 = No effect on reporting activity 1 = *Ingress LLIST Uncorrectable ECC Error* bit is masked/disabled	0	RWS	Yes	1
11	**Credit Update Timeout Status Mask** 0 = No effect on reporting activity 1 = *Credit Update Timeout Status* bit is masked/disabled	0, 1, 2, 3, 4	RWS	Yes	1

Register 13-50. 1D0h Error Handler 32-Bit Error Mask *(Factory Test Only) (Cont.)*

Bit(s)	Description	Bit Exists Only on Port(s)	Type	Serial EEPROM and I^2C	Default
12	**INCH Underrun Error Mask** 0 = No effect on reporting activity 1 = *INCH Underrun Error* bit is masked/disabled	0, 1, 2, 3, 4	RWS	Yes	1
13	**Ingress Memory Allocation Unit 1-Bit Soft Error Counter Overflow Mask** 0 = No effect on reporting activity 1 = *Ingress Memory Allocation Unit 1-Bit Soft Error Counter Overflow* bit is masked/disabled	0	RWS	Yes	1
14	**Ingress Memory Allocation Unit 2-Bit Soft Error Mask** 0 = Error reporting is enabled using interrupts 1 = *Ingress Memory Allocation Unit 2-Bit Soft Error* bit is masked/disabled	0	RWS	Yes	1
15	**NT Port DL_Down Virtual Interface Interrupt Mask** 0 = Enable Interrupt Generation from the NT Port Virtual Interface for a DL_Down event occurring on the NT Port Link Interface 1 = Mask Interrupt Generation from the NT Port Virtual Interface for a DL_Down event occurring on the NT Port Link Interface ***Note:*** *Not valid in Transparent mode. Writable, but reads 0 in Transparent mode.*	NT Port Virtual Interface	RWS	Yes	1
17:16	***Factory Test Only***		RWS	Yes	11b
31:18	***Reserved***		RsvdP	No	0-0h

Note: *For register offset 1DCh, if Port 0 is the NT Port:*

- *This register is loaded from the NT Port Virtual Interface register offset 1DCh location in the serial EEPROM*
- *The serial EEPROM must be programmed such that both the Port 0 and NT Port Virtual Interface register offset 1DCh locations contain the same value for this register*

Register 13-51. 1DCh Debug Control (Port 0, and also NT Port Virtual Interface if Port 0 is the NT Port)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
3:0	***Factory Test Only***	RsvdP	No	0h
6:4	**Cut-Thru Port Number** Selects the Port Number to be used for Cut-Thru. Software can change this field when the Cut-Thru port is not the upstream port. 000b = Port 0 001b = Port 1 010b = Port 2 011b = Port 3 100b = Port 4 All other values are ***reserved***. If the upstream port is intended to be a Cut-Thru port, the Cut-Thru Port Number is strapped by the STRAP_UPSTRM_PORTSEL[3:0] balls.	RW	Yes	Set by Strapping ball levels
7	***Factory Test Only***	RsvdP	No	0
10:8	**Upstream Port ID** When bit 15 (*Hardware/Software Configuration Mode Control* bit) is cleared to 0, Upstream Port Number – Reads External Strap value on the STRAP_UPSTRM_PORTSEL[3:0] balls when Reset is de-asserted. Software is not allowed to change this value.	HwInit	Yes	Set by Strapping ball levels
	When bit 15 (*Hardware/Software Configuration Mode Control* bit) is set to 1, Upstream Port Number is set by software. 0h = Port 0 1h = Port 1 2h = Port 2 3h = Port 3 4h = Port 4 5h to 7h = ***Reserved***	RW[a]	Yes	
11	***Factory Test Only***	RW	Yes	1

Register 13-51. 1DCh Debug Control (Port 0, and also NT Port Virtual Interface if Port 0 is the NT Port) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
13:12	***Reserved***	RsvdP	No	00b
14	**Data Path-Based Cut-Thru** Used for downstream traffic. 0 = Cut-Thru between the Cut-Thru port and another specific port is not allowed unless the paths between **all** ports and the Cut-Thru port are available for Cut-Thru (*that is*, there can be no path where TLPs are stored in the non-Cut-Thru path). Select this option to satisfy special ordering cases where, *for example*, Writes from the Cut-Thru port, to multiple other ports, must be performed in order. 1 = Cut-Thru is enabled on each path between the Cut-Thru port and other ports any time the path in question has no stored TLPs in the path.	RW	Yes	1
15	**Hardware/Software Configuration Mode Control** 0 = Strapping balls control the upstream port, NT Port, and bits [19:18] (*Mode Select* field); the upstream port is controlled by STRAP_UPSTRM_PORTSEL[3:0], and the NT Port is controlled by STRAP_NT_UPSTRM_PORTSEL[3:0] 1 = Software controls the upstream port, NT Port, and *Mode Select*	RW	Yes	0
16	**Upstream Hot Reset Control** 0 = Resets NT Port data path when the upstream port receives a Hot Reset or DL_Down state 1 = Does not reset NT Port data path when the upstream port receives a Hot Reset or DL_Down state ***Note:*** *In NT mode, value of 1 is automatically set as default.*	RW	Yes	0
17	**Disable Serial EEPROM Load Disable on Level 1 Reset** Serial EEPROM load is disabled only for the Hot Reset or DL_Down state; does not affect Fundamental Reset. 0 = Enables Serial EEPROM load upon a Hot Reset 1 = Disables Serial EEPROM load upon a Hot Reset	RW	Yes	0
19:18	**Mode Select** Mode is selected by STRAP_MODE_SEL[1:0] Strapping balls, and overridden by the serial EEPROM. Software is not allowed to change this value. 00b = ***Reserved*** 01b = NT Intelligent Adapter mode 10b = NT Dual-Host mode 11b = Transparent mode	HwInit RW[a]	Yes Yes	Set by Strapping ball levels
20	**Upstream Port Hot Reset and Link Down Reset Propagation Disable** 0 = Internal Reset and Hot Reset propagation are enabled 1 = Internal Reset and Hot Reset propagation are disabled Set to 1 for NT Dual-Host mode.	RW	Yes	0
21	**Cut-Thru Enable** Enables Cut-Thru support.	RW	Yes	1
22	***Reserved***	RsvdP	No	0
23	***Factory Test Only***	RW	Yes	0

Register 13-51. 1DCh Debug Control (Port 0, and also NT Port Virtual Interface if Port 0 is the NT Port) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
27:24	**NT Port Number** When field [19:18] (*Mode Select* field) is set to 01b or 10b and bit 15 (*Hardware/Software Configuration Mode Control* bit) is cleared to 0, the NT Port Number is set by the STRAP_NT_UPSTRM_PORTSEL[3:0] Strapping balls. Otherwise, this field is Read-Only. This field is "don't care" for Transparent mode. Software is not allowed to change this value.	HwInit	Yes	Set by Strapping ball levels
	When field [19:18] (*Mode Select* field) is set to 01b or 10b and bit 15 (*Hardware/Software Configuration Mode Control* bit) is set to 1, the NT Port Number is selected by this field, as set by software. 0h = Port 0 1h = Port 1 2h = Port 2 3h = Port 3 4h = Port 4 5h to 7h = ***Reserved***	RW[a]	Yes	
28	**Virtual Interface Access Enable** Used only in NT mode. When the serial EEPROM is not present, the default value is 1; otherwise, the default value is 0. 0 = Retries Type 0 Configuration TLP received on the NT Port Virtual Interface 1 = Accepts Type 0 Configuration TLP on the NT Port Virtual Interface ***Note:*** *This bit does not affect the PEX 8508 in Transparent mode.* *Set this bit to enable Configuration access to the NT Port Virtual Interface.*	RW	Yes	1
29	**Link Interface Access Enable** Used only in NT mode. Default value is 1 when bits [19:18] (*Mode Select* field) are set to 10b (Dual Host mode). Otherwise, the default value is 0. 0 = Retries Type 0 Configuration TLP received on the NT Port Link Interface 1 = Accepts Type 0 Configuration TLP on the NT Port Link Interface ***Note:*** *This bit does not affect the PEX 8508 in Transparent mode.* *Set this bit to enable Configuration access to the NT Port Link Interface.*	RW	Yes	0
30	**On-Board SerDes Lane Status Control** 0 = Physical Layer Lane Up status controls the on-board SerDes PEX_LANE_GOOD[10, 8, 6, 4:0]# outputs 1 = Software-driven value to the **Software-Controlled Lane Status** register (offset 1F4h) controls the on-board SerDes PEX_LANE_GOOD[10, 8, 6, 4:0]# outputs	RW	Yes	0
31	**Power-Up RAM BIST Status** Reports the power-up RAM BIST result. 0 = No power-up RAM BIST error 1 = Non-recoverable Fatal RAM BIST error detected, must replace the PEX 8508 device prior to bring-up	RO	No	0

a. Although these bits are RW, do not change by software.

Register 13-52. 1E0h Power Management Hot Plug User Configuration (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**L0s Entry Idle Count** Time to meet to enter L0s. 0 = Idle condition lasts for 1 µs 1 = Idle condition lasts for 4 µs	RW	Yes	0
1	**L1 Upstream Port Receiver Idle Count** For active L1 entry. 0 = Upstream port receiver remains idle for 2 µs 1 = Upstream port receiver remains idle for 3 µs	RW	Yes	0
2	**HPC PME Turn-Off Enable** 1 = PME Turn-Off message is transmitted before the port is turned Off on a downstream port	RW	Yes	0
4:3	**HPC T_{pepv} Delay** Slot power-applied to power-valid delay time. 00b = 16 ms 01b = 32 ms 10b = 64 ms 11b = 128 ms	RO	Yes	00b
5	**HPC Inband Presence Detect Enable** 0 = HP_PRSNT[4:0]# Input balls are used to detect a board present in the slot 1 = SerDes Receiver Detect mechanism is used to detect a board present in the slot	RO	Yes	0
6	**HPC T_{pvperl} Delay** Downstream port power valid to reset signal release time. 0 = 20 ms 1 = 100 ms (default)	RO	Yes	1
12:7	***Factory Test Only***	RW	Yes	0_0000_0b
31:13	***Reserved***	RsvdP	No	0-0h

Register 13-53. 1E4h Egress Control and Status (All Ports)

Bit(s)	Description	Bit Exists Only on Port(s)	Type	Serial EEPROM and I²C	Default
0	**Egress Credit Update Timer Enable** In this mode, when the port is not receiving credits to make forward progress and the Egress Timeout Timer times out, the downstream link is brought down. 0 = Egress Credit Update Timer is disabled 1 = Egress Credit Update Timer is enabled	0, 1, 2, 3, 4	RW	Yes	0
1	**Egress Timeout Value** 0 = Minimum 512 ns (Maximum 768 ns) 1 = Minimum 1,024 ns (Maximum 1,280 ns)	0, 1, 2, 3, 4	RW	Yes	0
8:2	***Reserved***		RW	Yes	0000_000b
10:9	***Factory Test Only***		RW	Yes	00b
15:11	***Factory Test Only*** Testing bits – must be 0000_0b.		RW	Yes	0000_0b
19:16	**VC&T Encountered Timeout** 0h = VC0 Posted 1h = VC0 Non-Posted 2h = VC0 Completion 3h = VC1 Posted 4h = VC1 Non-Posted 5h = VC1 Completion All other values are ***reserved***.	0, 1, 2	RO	Yes	0h
	0h = VC0 Posted 1h = VC0 Non-Posted 2h = VC0 Completion All other values are ***reserved***.	3, 4	RO	Yes	0h
23:20	***Reserved***		RsvdP	No	0h
31:24	**Packet RAM 8-Bit Soft Error Counter Value** Hardware increments this counter when a 1-bit Soft error in packet RAM Read is detected.	0	RO	No	00h

Register 13-54. 1E8h Bad TLP Count (All Ports Except NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Bad TLP Count** Counts the number of TLPs with bad LCRC, or number of TLPs with a sequence number mismatch error. The maximum value is FFh. The Counter saturates at FFh and does not roll over to 00h.	RW	Yes	00h
31:8	***Reserved***	RsvdP	No	0h

Register 13-55. 1ECh Bad DLLP Count (All Ports Except NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Bad DLLP Count** Counts the number of DLLPs with bad LCRC, or number of DLLPs with a sequence number mismatch error. The maximum value is FFh. The Counter saturates at FFh and does not roll over to 00h.	RW	Yes	00h
31:8	***Reserved***	RsvdP	No	0h

Register 13-56. 1F0h Device-Specific Relaxed Ordering Enable (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
19:0	***Reserved***	RsvdP	No	0h
20	**Device-Specific Relaxed Ordering Enable** Enables VC0 Completions to pass VC0 Posted packets. 0 = VC0 Completion TLP blocked by older VC0 Posted TLP 1 = VC0 Completion TLP bypasses the older VC0 Posted TLP	RW	Yes	0
31:21	***Reserved***	RsvdP	No	0h

Note: *For register offset 1F4h, if the following conditions are met:*

- *Serial EEPROM programs Port 0 to be the NT Port*
- *Debug Control register On-Board SerDes Lane Status Control bit is set (offset 1DCh[30]=1) to enable software or serial EEPROM control of the PEX_LANE_GOOD[10, 8, 6, 4:0]# outputs*

the serial EEPROM must be programmed such that both Port 0 and the NT Port Virtual Interface register offset 1F4h locations contain the same value for this register.

Register 13-57. 1F4h Software-Controlled Lane Status [Port 0, and also NT Port Virtual Interface if Port 0 is the NT Port)

Bit(s)	Description		Type	Serial EEPROM and I^2C	Default
10, 8, 6, 4:0	**Software-Controlled Lane Status** Echoes the state of the PEX_LANE_GOOD[10, 8, 6, 4:0]# balls when the **Debug Control** register *On-Board SerDes Lane Status Control* bit (offset 1DCh[30]) is cleared to 0; otherwise, these bits are Read-Write when the *On-Board SerDes Lane Status Control* bit is set to 1. Bit 0 controls SerDes Lane 0. Bit 4 controls SerDes Lane 4. Bit 6 controls SerDes Lane 6. Bit 8 controls SerDes Lane 8. Bit 10 controls SerDes Lane 10.	If **Debug Control** register *On-Board SerDes Lane Status Control* bit is cleared (offset 1DCh[30]=0, default)	RO	Yes	0-0h
		If **Debug Control** register *On-Board SerDes Lane Status Control* bit is set (offset 1DCh[30]=1)	RW	Yes	0-0h
31:16, 15:11, 9, 7, 5	***Reserved***		RsvdP	No	0-0h

Register 13-58. 1F8h ACK Transmission Latency Limit (All Ports)

<table>
<tr><th>Bit(s)</th><th>Description</th><th>Type</th><th>Serial EEPROM and I²C</th><th>Default</th></tr>
<tr><td>7:0</td><td>ACK Transmission Latency Limit
If the serial EEPROM is not present, the value of this register changes based upon the negotiated link width after the link is up. This value assumes that Maximum Payload Size is 256 bytes.
If the serial EEPROM is present, program the serial EEPROM to load the value based upon the programmed port width and Device Control register Maximum Payload Size field value (offset 70h[7:5]):
<table><tr><th rowspan="2">Maximum Payload Size</th><th colspan="3">Port Width</th></tr><tr><th>x1</th><th>x2</th><th>x4</th></tr><tr><td>128B</td><td>FAh</td><td>80h</td><td>49h</td></tr><tr><td>256B</td><td>FFh</td><td>D9h</td><td>76h</td></tr></table>Note: The value of this field is valid after link negotiation completes.</td><td>RWS</td><td>Yes</td><td>FFh</td></tr>
<tr><td>15:8</td><td>Factory Test Only
Testing bits – must be 00h.</td><td>RWS</td><td>Yes</td><td>00h</td></tr>
<tr><td>23:16</td><td>Upper 8 Bits of Replay Timer Limit
If the serial EEPROM is not present, the value of this register changes based upon the negotiated link width after the link is up.
If the serial EEPROM is present, program the serial EEPROM to load the value based upon the programmed port width.
Value is based upon negotiated link width encoding:
<table><tr><th>Link Width</th><th>Register Value</th></tr><tr><td>x1</td><td>05h</td></tr><tr><td>x2</td><td>03h</td></tr><tr><td>x4</td><td>02h</td></tr></table></td><td>RWS</td><td>Yes</td><td>FFh</td></tr>
<tr><td>31:24</td><td>Reserved</td><td>RsvdP</td><td>No</td><td>00h</td></tr>
</table>

13.13.2 Physical Layer Registers

Note: *In this section, the term "SerDes quad" or "quad" refers to assembling SerDes lanes into groups of four lanes for testing purposes. The quads are defined as SerDes[0-3] (quad 0), SerDes[4-7] (quad 1), and SerDes[8-11] (quad 2).* ***The*** *PEX 8508* ***does not use SerDes[5, 7, 9, 11]; these lanes are reserved.***

Table 13-13. Device-Specific Physical Layer Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Factory Test Only		200h
Physical Receiver Not Detected Mask	Physical Electrical Idle Detect Mask	204h
Physical Deskew Level Low		208h
Physical Deskew Level High		20Ch
Phy User Test Pattern 0		210h
Phy User Test Pattern 4		214h
Phy User Test Pattern 8		218h
Phy User Test Pattern 12		21Ch
Physical Layer Status	Physical Layer Command	220h
Port Configuration		224h
Physical Layer Test		228h
Physical Layer		22Ch
Physical Layer Port Command		230h
Port 4 Receive Error Count / Port Control	SKIP Ordered-Set Interval	234h
Quad 0 SerDes Diagnostic Data		238h
Quad 1 SerDes Diagnostic Data		23Ch
Quad 2 SerDes Diagnostic Data		240h
Reserved		244h
SerDes Nominal Drive Current Select		248h
SerDes Drive Current Level Select 1		24Ch
SerDes Drive Current Level Select 2		250h
SerDes Drive Equalization Level Select 1		254h
SerDes Drive Equalization Level Select 2		258h
Physical Receive Error Count		25Ch
Status Data from Serial EEPROM / Serial EEPROM Status	Serial EEPROM Control	260h
Serial EEPROM Data Buffer		264h
Reserved		268h – 28Ch

Note: *For configuration purposes, cross-link connections are* ***not supported****.*

Register 13-59. 204h Physical Receiver Not Detected and Electrical Idle Detect Masks (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Physical Electrical Idle Detect Mask				
10, 8, 6, 4:0	**SerDes Mask Electrical Idle Detect** Bit 0 controls SerDes Lane 0. … Bit 4 controls SerDes Lane 4. Bit 6 controls SerDes Lane 6. Bit 8 controls SerDes Lane 8. Bit 10 controls SerDes Lane 10.	RWS	Yes	0-0h
15:11, 9, 7, 5	***Reserved***	RsvdP	No	0-0h
Physical Receiver Not Detected Mask				
26, 24, 22, 20:16	**SerDes Mask Receiver Not Detected** Bit 16 controls SerDes Lane 0. … Bit 20 controls SerDes Lane 4. Bit 22 controls SerDes Lane 6. Bit 24 controls SerDes Lane 8. Bit 26 controls SerDes Lane 10.	RWS	Yes	0-0h
31:27, 25, 23, 21	***Reserved***	RsvdP	No	0-0h

Register 13-60. 208h Physical Deskew Level Low[a] (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
2:0	**Lane 0 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 0. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
3	**Lane 0 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 0 delay. 1 = Enables the Lane 0 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
6:4	**Lane 1 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 1. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
7	**Lane 1 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 1 delay. 1 = Enables the Lane 1 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
10:8	**Lane 2 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 2. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
11	**Lane 2 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 2 delay. 1 = Enables the Lane 2 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
14:12	**Lane 3 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 3. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
15	**Lane 3 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 3 delay. 1 = Enables the Lane 3 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0

Register 13-60. 208h Physical Deskew Level Low[a] (Only Port 0) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
18:16	**Lane 4 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 4. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
19	**Lane 4 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 4 delay. 1 = Enables the Lane 4 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
23:20	***Reserved***	RsvdP	No	0h
26:24	**Lane 6 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 6. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
27	**Lane 6 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 6 delay. 1 = Enables the Lane 6 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
31:28	***Reserved***	RsvdP	No	0h

a. It is not recommended to change the Lane Deskew level from Auto mode to Manual mode.

b. Refer to Table 4-1, "PEX 8508 Port Configurations," for Port/Lane associations.

Register 13-61. 20Ch Physical Deskew Level High[a] (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
2:0	**Lane 8 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 8. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
3	**Lane 8 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 8 delay. 1 = Enables the Lane 8 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
7:4	***Reserved***	RsvdP	No	0h
10:8	**Lane 10 Deskew Level** When written, controls the Deskew Delay (in symbol times) for Lane 10. When read, reads the Auto Deskew Delay setting or manually enabled value.	RWS	Yes	000b
11	**Lane 10 Manual Deskew Level Enable** 0 = PEX 8508 performs an Auto Deskew function, to balance the Lane 10 delay. 1 = Enables the Lane 10 manual Lane Deskew option. The Lane Deskew value (in symbol times) must be set in this case.	RWS	Yes	0
31:12	***Reserved***	RsvdP	No	0000_0h

a. It is not recommended to change the Lane Deskew level from Auto mode to Manual mode.

b. Refer to Table 4-1, "PEX 8508 Port Configurations," for Port/Lane associations.

Register 13-62. 210h Phy User Test Pattern 0 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Test Pattern 0** Test pattern bytes 0-3. Used for Digital Far-End Loop-Back testing.	RW	Yes	0-0h

Register 13-63. 214h Phy User Test Pattern 4 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Test Pattern 4** Test pattern bytes 4-7. Used for Digital Far-End Loop-Back testing.	RW	Yes	0-0h

Register 13-64. 218h Phy User Test Pattern 8 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Test Pattern 8** Test pattern bytes 8-11. Used for Digital Far-End Loop-Back testing.	RW	Yes	0-0h

Register 13-65. 21Ch Phy User Test Pattern 12 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Test Pattern 12** Test pattern bytes 12-15. Used for Digital Far-End Loop-Back testing.	RW	Yes	0-0h

Register 13-66. 220h Physical Layer Command and Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Physical Layer Command				
0	**Port Enumerator Enable** 0 = Enumerate is not enabled 1 = Enumerate is enabled ***Note:*** *Value is 1 after ee_csrloaded asserts.*	HwInit	Yes	0
1	**TDM Enable** 0 = TDM is not enabled 1 = TDM is enabled ***Note:*** *Value is 1 after ee_csrloaded asserts.*	HwInit	Yes	0
2	***Reserved***	RWS	Yes	0
3	**Upstream Port as Configuration Master Enable** 0 = Upstream Port Cross-link is not supported 1 = Upstream Port Cross-link is supported	RWS	Yes	0
4	**Downstream Port as Configuration Slave Enable** 0 = Downstream Port Cross-link is not supported 1 = Downstream Port Cross-link is supported	RWS	Yes	0
5	**Lane Reversal Disable** Provides ability to enable or disable lane reversal. 0 = Lane reversal is supported 1 = Lane reversal is not supported ***Note:*** *Lane reversal is not supported on Port 0.*	RWS	Yes	0
6	**PLL Turn-Off Enable**	RWS	Yes	0
7	***Factory Test Only***	RWS	Yes	0
15:8	**N_FTS Value** Number of Fast Training Sets (N_FTS) value to transmit in training sets.	RWS	Yes	40h
Physical Layer Status				
19:16	***Reserved***	RsvdP	No	0h
22:20	**Number of Ports Enumerated** Number of ports in current configuration.	HwInit	Yes	000b
23	***Reserved***	RsvdP	No	0
24	**Port 0 Deskew Buffer Error Status** 1 = Port 0 Deskew Buffer overflow or underflow	RW1C	Yes	0
25	**Port 1 Deskew Buffer Error Status** 1 = Port 1 Deskew Buffer overflow or underflow	RW1C	Yes	0
26	**Port 2 Deskew Buffer Error Status** 1 = Port 2 Deskew Buffer overflow or underflow	RW1C	Yes	0
27	**Port 3 Deskew Buffer Error Status** 1 = Port 3 Deskew Buffer overflow or underflow	RW1C	Yes	0
31:28	***Reserved***	RsvdP	No	0h

Register 13-67. 224h Port Configuration (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
4:0	**Port Configuration** The serial EEPROM bit values always override the values of the STRAP_PORTCFG[4:0] Strapping signals (assuming the serial EEPROM values are loaded; refer to Table 13-14). Bits [4:3] must always be programmed to 00b. This register is reset only by a Fundamental Reset (PEX_PERST# assertion). The serial EEPROM is used to optionally re-configure the ports from the value set by the STRAP_PORTCFG[4:0] inputs, using the values defined in Table 13-14. ***Note:*** *All other configurations default to option 0.*	HwInit	Yes	0_0000b
31:5	***Reserved***	RsvdP	No	0-0h

Table 13-14. PEX 8508 Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[4:0])	Lane Width Per Port				
	0	1	2	3	4
0h	x2	x2	x2	x2	
2h	x4	x4			
3h	x4	x2	x2		
4h	x4	x2	x1	x1	
5h	x4	x1	x1	x2	
6h	x4	x1	x2	x1	
8h	x4	x1	x1	x1	x1
9h	x2	x2	x2	x1	x1

Register 13-68. 228h Physical Layer Test (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Timer Test Mode Enable** 0 = Normal Physical Layer Timer parameters used 1 = Shortens Timer scale from milliseconds to microseconds	RW	Yes	0
1	**SKIP Timer Test Mode Enable** 0 = Disables SKIP Timer Test mode 1 = Enables SKIP Timer Test mode	RW	Yes	0
2	**Port 0 x1 Only** 1 = Port 0 is configured as x1 only	RW	Yes	0
3	**TCB Capture Disable** 0 = Training Control Bit (TCB) Capture is enabled 1 = Disables TCB Capture	RW	Yes	0
4	**Analog Loop-Back Enable** 0 = PEX 8508 enters Digital Loop-Back Slave mode if an external device sends at least two consecutive TS1 Ordered-Sets that have the *Loop-Back* bit exclusively set in the TS1 Training Control symbol. The PEX 8508 then loops back data through the elastic buffer, 8b/10b decoder, and 8b/10b encoder. 1 = PEX 8508 enters Analog Loop-Back Slave mode if an external device sends at least two consecutive TS1 Ordered-Sets that have the *Loop-Back* bit exclusively set in the TS1 Training Control symbol. The PEX 8508 then loops back the symbol stream from the 10-bit Receive interface (before the elastic buffer) to the 10-bit Transmit interface.	RW	Yes	0
5	**Port/SerDes Test Pattern Enable Select** 1 = Bits [30:28] (*User Test Pattern Enable* bits) select ports rather than SerDes quads	RW	Yes	0
6	***Reserved***	RsvdP	No	0
7	**SerDes BIST Enable** When programmed to 1 by serial EEPROM, enables SerDes internal loop-back Pseudo-Random Bit Sequence (PRBS) test for 512 µs before starting link initialization at power-on.	RO	Yes	0
9:8	**PRBS Diagnostic Data Select** Selects the SerDes within the quad for PRBS generation/checking. **Value** **Port 0, SerDes** 00b = [0, 4, 8] 01b = [1] 10b = [2, 6, 10] 11b = [3]	RW	Yes	00b
15:10	***Reserved***	RsvdP	No	0-0h

Register 13-68. 228h Physical Layer Test (Only Port 0 and NT Port Link Interface) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
18:16	**PRBS Enable** When set to 1, enables PRBS sequence generation/checking on the SerDes quads. **Bit** **Port 0, SerDes** 16 [0-3] 17 [4, 6] 18 [8, 10] ***Note: PRBS Enable and User Test Pattern Enable (bits [30:28]) are mutually exclusive functions and must not be enabled together for the same SerDes quad. In the Port 0 register, the logical result of bits [18:16] ANDed with bits [30:28] must be 000b.***	RW	Yes	000b
19	***Reserved***	RsvdP	No	0
22:20	**PRBS External Loop-Back** 0 = SerDes quad establishes Internal Analog Loop-Back mode when the corresponding *PRBS Enable* bit (bit 18, 17, or 16) is set to 1 1 = SerDes quad establishes Analog Loop-Back Master mode when the corresponding *PRBS Enable* bit (bit 18, 17, or 16) is set to 1 The following bit commands are valid when the **Physical Layer Port Command** register *Port x Loop-Back Command* bit (Port 0, offset 230h[0, 4, 8, 12, and/or 16]) is set for the associated port. **Bit** **Port 0, SerDes** 20 [0-3] 21 [4, 6] 22 [8, 10]	RW	Yes	000b
23	***Reserved***	RsvdP	No	0
26:24	**PRBS Error Count Reset** When set to 1, resets the PRBS Error Counter (**Quad *x* SerDes Diagnostic Data** register *PRBS Error Count* field). The bits in this field are self-clearing. **Bit** **Port 0** 24 Offset 238h[31:24] 25 Offset 23Ch[31:24] 26 Offset 240h[31:24]	RO	Yes	000b
27	***Reserved***	RsvdP	No	0b

Register 13-68. 228h Physical Layer Test (Only Port 0 and NT Port Link Interface) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
30:28	**User Test Pattern Enable** 0 = Disables transmission of the 128-bit test pattern 1 = Enables transmission of the 128-bit test pattern [**Phy User Test Pattern *x*** registers (Port 0, offsets 210h through 21Ch)] on the SerDes quads in Digital Far-End Loop-Back Master mode **Bit** **Port 0, SerDes** 28 [0-3] 29 [4, 6] 30 [8, 10] ***Note: User Test Pattern Enable and PRBS Enable (bits [18:16]) are mutually exclusive functions and must not be enabled together for the same SerDes quad. In the Port 0 register, the logical result of bits [18:16] ANDed with bits [30:28] must be 000b.***	RW	Yes	000b
31	***Reserved***	RsvdP	No	0

Note: *Port 0 parameters apply to SerDes[0-3], SerDes[4, 6], and SerDes[8, 10]. SerDes[5, 7] and SerDes[9, 11] are* ***reserved.***

Register 13-69. 22Ch Physical Layer (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
5:0	***Factory Test Only***	RW1S	Yes	00_0000b
7:6	***Reserved***	RW1S	Yes	00b
9:8	**SerDes Quad 0 TxTermAdjust** SerDes Quad 0 TxTermAdj[1:0]. Control bus to adjust Transmit termination values above or below the nominal 50 Ohms for physical Lanes [0-3]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50 Ohms) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RW1S	Yes	00b
11:10	**SerDes Quad 1 TxTermAdjust** SerDes Quad 1 TxTermAdj[1:0]. Control bus to adjust Transmit termination values above or below the nominal 50 Ohms for physical Lanes [4, 6]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50 Ohms) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RW1S	Yes	00b
13:12	**SerDes Quad 2 TxTermAdjust** SerDes Quad 2 TxTermAdj[1:0]. Control bus to adjust Transmit termination values above or below the nominal 50 Ohms for physical Lanes [8, 10]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50 Ohms) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RW1S	Yes	00b
15:14	***Reserved***	RsvdP	No	00b

Register 13-69. 22Ch Physical Layer (Only Port 0 and NT Port Link Interface) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
17:16	**SerDes Quad 0 RxTermAdjust** SerDes Quad 0 RxTermAdj[1:0]. Control bus to adjust Receive termination values above or below the nominal 50 Ohms for physical Lanes [0-3]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50 Ohms) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RW1S	Yes	00b
19:18	**SerDes Quad 1 RxTermAdjust** SerDes Quad 1 RxTermAdj[1:0]. Control bus to adjust Receive termination values above or below the nominal 50 Ohms for physical Lanes [4, 6]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50 Ohms) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RW1S	Yes	00b
21:20	**SerDes Quad 2 RxTermAdjust** SerDes Quad 2 RxTermAdj[1:0]. Control bus to adjust Receive termination values above or below the nominal 50 Ohms for physical Lanes [8, 10]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50 Ohms) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RW1S	Yes	00b
23:22	***Reserved***	RsvdP	No	00b
25:24	**SerDes Quad 0 RxEqCtl** SerDes Quad 0 RxEqCtl[1:0]. Control bus to adjust the Receiver equalization, globally for physical Lanes [0-3]. For further details, refer to the expanded description that follows this register table. Table 13-15 defines the field decode.	RW1S	Yes	00b
27:26	**SerDes Quad 1 RxEqCtl** SerDes Quad 1 RxEqCtl[1:0]. Control bus to adjust the Receiver equalization, globally for physical Lanes [4, 6]. For further details, refer to the expanded description that follows this register table. Table 13-15 defines the field decode.	RW1S	Yes	00b
29:28	**SerDes Quad 2 RxEqCtl** SerDes Quad 2 RxEqCtl[1:0].Control bus to adjust the Receiver equalization, globally for physical Lanes [8, 10]. For further details, refer to the expanded description that follows this register table. Table 13-15 defines the field decode.	RW1S	Yes	00b
31:30	***Reserved***	RsvdP	No	00b

SerDes Quad *x* RxEqCtl Expanded Description. At high speeds, the channel between a PCI Express Transmitter and Receiver exhibits frequency-dependent losses (*such as* due to PCB dielectric and conductor skin-effect). The channel acts as a low-pass filter, attenuating the high-frequency components of a signal passing through it. This distortion results in Inter Symbol Interference (ISI). ISI is a form of deterministic jitter that can easily close the received data "eye," reducing the ability to reliably recover a data stream across the channel. To mitigate the effects of ISI, the receiver at each lane includes a receive equalizer. The receive equalizer is implemented as a selectable, high-pass filter at the receiver input pad and is capable of removing as much as 0.4 UI of ISI-related jitter. SerDes Quad *x* RxEqCtl decodes as defined in Table 13-15.

The Channel Length assumes standard FR4 material. The Rx Equalizer settings should be chosen based upon the amount of deterministic jitter induced by the channel. The channel lengths listed in the table above are included as a general guideline, not as an absolute reference. Deterministic jitter as a function of channel length can vary with PCB layer stackup, PCB material, and the type of connector(s) used.

Table 13-15. RxEqCtl[1:0] Decode for Register Offset 224h[29:24]

RxEqCtl[1:0]	RX Eq Setting	Input Jitter	Channel Length
00b	Maximum Rx Eq	0.25 UI	50.8 cm (20 inches) and two or more connectors
01b	Minimum Rx Eq	Between 0.1 and 0.25 UI	Between 20.32 and 50.8 cm (8 and 20 inches) and up to two connectors
10b, 11b	Rx Eq OFF	< 0.1 UI	20.32 cm (8 inches) or less, up to one connector

Register 13-70. 230h Physical Layer Port Command (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Port 0 Loop-Back Command** 0 = Port 0 is not enabled to go to the Loop-Back Master state 1 = Port 0 is enabled to go to the Loop-Back Master state	RW	Yes	0
1	**Port 0 Scramble Disable Command** When the serial EEPROM load sets this bit, the scrambler is disabled in the Configuration-Complete state. When software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing Link Training protocol. The upstream/downstream device scrambler will not be disabled. 0 = Port 0 scrambler is enabled 1 = Port 0 scrambler is disabled	RW	Yes	0
2	**Port 0 Rx L1 Only Command** Port 0 Receiver enters into the ASPM L1 Link PM state. 0 = Port 0 receiver is allowed to go to the ASPM L0s or L1 Link PM state when an Electrical Idle Ordered-Set in the L0 Link PM state is detected 1 = Port 0 receiver is allowed to go to the ASPM L1 Link PM state only when an Electrical Idle Ordered-Set in the L0 Link PM state is detected	RW	Yes	0
3	**Port 0 Ready as Loop-Back Master** Port 0 LTSSM established Loop-Back as a Master. 0 = Port 0 is not in Loop-Back Master mode 1 = Port 0 is in Loop-Back Master mode	RO	No	0

Register 13-70. 230h Physical Layer Port Command (Only Port 0 and NT Port Link Interface) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
4	**Port 1 Loop-Back Command** 0 = Port 1 is not enabled to go to the Loop-Back Master state 1 = Port 1 is enabled to go to the Loop-Back Master state	RW	Yes	0
5	**Port 1 Scramble Disable Command** When the serial EEPROM load sets this bit, the scrambler is disabled in the Configuration-Complete state. When software sets this bit when the Link is in the up state, hardware immediately disables its scrambler without executing Link Training protocol. The upstream/downstream device scrambler is not disabled. 0 = Port 1 scrambler is enabled 1 = Port 1 scrambler is disabled	RW	Yes	0
6	**Port 1 Rx L1 Only Command** Port 1 Receiver enters into the ASPM L1 Link PM state. 0 = Port 1 receiver is allowed to go to the ASPM L0s or L1 Link PM state when an Electrical Idle Ordered-Set in the L0 Link PM state is detected 1 = Port 1 receiver is allowed to go to the ASPM L1 Link PM state only when an Electrical Idle Ordered-Set in the L0 Link PM state is detected	RW	Yes	0
7	**Port 1 Ready as Loop-Back Master** Port 1 LTSSM established Loop-Back as a Master. 0 = Port 1 is not in Loop-Back Master mode 1 = Port 1 is in Loop-Back Master mode	RO	No	0
8	**Port 2 Loop-Back Command** 0 = Port 2 is not enabled to go to the Loop-Back Master state 1 = Port 2 is enabled to go to the Loop-Back Master state	RW	Yes	0
9	**Port 2 Scramble Disable Command** When the serial EEPROM load sets this bit, the scrambler is disabled in the Configuration-Complete state. When software sets this bit when the Link is in the up state, hardware immediately disables its scrambler without executing Link Training protocol. The upstream/downstream device scrambler is not disabled. 0 = Port 2 scrambler is enabled 1 = Port 2 scrambler is disabled	RW	Yes	0
10	**Port 2 Rx L1 Only Command** Port 2 Receiver enters into the ASPM L1 Link PM state. 0 = Port 2 receiver is allowed to go to the ASPM L0s or L1 Link PM state when an Electrical Idle Ordered-Set in the L0 Link PM state is detected 1 = Port 2 receiver is allowed to go to the ASPM L1 Link PM state only when an Electrical Idle Ordered-Set in the L0 Link PM state is detected	RW	Yes	0
11	**Port 2 Ready as Loop-Back Master** Port 2 LTSSM established Loop-Back as a Master. 0 = Port 2 is not in Loop-Back Master mode 1 = Port 2 is in Loop-Back Master mode	RO	No	0

Register 13-70. 230h Physical Layer Port Command (Only Port 0 and NT Port Link Interface) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
12	**Port 3 Loop-Back Command** 0 = Port 3 is not enabled to go to the Loop-Back Master state 1 = Port 3 is enabled to go to the Loop-Back Master state	RW	Yes	0
13	**Port 3 Scramble Disable Command** When the serial EEPROM load sets this bit, the scrambler is disabled in the Configuration-Complete state. When software sets this bit when the Link is in the up state, hardware immediately disables its scrambler without executing Link Training protocol. The upstream/downstream device scrambler is not disabled. 0 = Port 3 scrambler is enabled 1 = Port 3 scrambler is disabled	RW	Yes	0
14	**Port 3 Rx L1 Only Command** Port 3 Receiver enters into the ASPM L1 Link PM state. 0 = Port 3 receiver is allowed to go to the ASPM L0s or L1 Link PM state when an Electrical Idle Ordered-Set in the L0 Link PM state is detected 1 = Port 3 receiver is allowed to go to the ASPM L1 Link PM state only when an Electrical Idle Ordered-Set in the L0 Link PM state is detected	RW	Yes	0
15	**Port 3 Ready as Loop-Back Master** Port 3 LTSSM established Loop-Back as a Master. 0 = Port 3 is not in Loop-Back Master mode 1 = Port 3 is in Loop-Back Master mode	RO	No	0

Register 13-70. 230h Physical Layer Port Command (Only Port 0 and NT Port Link Interface) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
16	**Port 4 Loop-Back Command** 0 = Port 4 is not enabled to go to the Loop-Back Master state 1 = Port 4 is enabled to go to the Loop-Back Master state	RW	Yes	0
17	**Port 4 Scramble Disable Command** When the serial EEPROM load sets this bit, the scrambler is disabled in the Configuration-Complete state. When software sets this bit when the Link is in the up state, hardware immediately disables its scrambler without executing Link Training protocol. The upstream/downstream device scrambler is not disabled. 0 = Port 4 scrambler is enabled 1 = Port 4 scrambler is disabled	RW	Yes	0
18	**Port 4 Rx L1 Only Command** Port 4 Receiver enters into the ASPM L1 Link PM state. 0 = Port 4 receiver is allowed to go to the ASPM L0s or L1 Link PM state when an Electrical Idle Ordered-Set in the L0 Link PM state is detected 1 = Port 4 receiver is allowed to go to the ASPM L1 Link PM state only when an Electrical Idle Ordered-Set in the L0 Link PM state is detected	RW	Yes	0
19	**Port 4 Ready as Loop-Back Master** Port 4 LTSSM established Loop-Back as a Master. 0 = Port 4 is not in Loop-Back Master mode 1 = Port 4 is in Loop-Back Master mode	RO	No	0
31:20	***Reserved***	RsvdP	No	0-0h

Register 13-71. 234h Port 4 Receive Error Count; Port Control and SKIP Ordered-Set Interval (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
SKIP Ordered-Set Interval				
11:0	**SKIP Ordered-Set Interval** SKIP Ordered-Set interval (symbol times). 49Ch = Minimum interval (1,180 symbol times) 602h = Maximum interval (1,538 symbol times)	RWS	Yes	49Ch
15:12	***Reserved***	RsvdP	No	0h
Port Control				
16	**Disable Port 0** 0 = Enables Link Training operation on Port 0 1 = Disables Link Training operation on Port 0	RWS	Yes	0
17	**Disable Port 1** 0 = Enables Link Training operation on Port 1 1 = Disables Link Training operation on Port 1	RWS	Yes	0
18	**Disable Port 2** 0 = Enables Link Training operation on Port 2 1 = Disables Link Training operation on Port 2	RWS	Yes	0
19	**Disable Port 3** 0 = Enables Link Training operation on Port 3 1 = Disables Link Training operation on Port 3	RWS	Yes	0
20	**Disable Port 4** 0 = Enables Link Training operation on Port 4 1 = Disables Link Training operation on Port 4	RWS	Yes	0
23:21	***Reserved***	RsvdP	No	000b
Port 4 Receive Error Count				
31:24	**Port 4 Receive Error Count** Reports the number of PHY packets received on Port 4 with CRC errors.	RO	No	00h

Register 13-72. 238h Quad 0 SerDes Diagnostic Data (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**UTP Expected Data** Expected User Test Pattern (UTP) SerDes[0-3] Diagnostic data when UTP is enabled [**Physical Layer Test** register *User Test Pattern Enable* bit for SerDes Quad 0 (Port 0, offset 228h[28]) is set to 1].	RO	No	00h
15:8	**Actual Data** Actual UTP SerDes[0-3] Diagnostic data when UTP is enabled [**Physical Layer Test** register *User Test Pattern Enable* bit for SerDes Quad 0 (Port 0, offset 228h[28]) is set to 1].	RO	No	00h
19:16	**Receiver Detected SerDes[0-3]**	RO	No	0h
23:20	***Reserved***	RO	No	0h
31:24	**PRBS Error Count** PRBS SerDes[0-3] Error Count (0 to 255).	RO	No	00h

Register 13-73. 23Ch Quad 1 SerDes Diagnostic Data (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**UTP Expected Data** Expected User Test Pattern (UTP) SerDes[4, 6] Diagnostic data when UTP is enabled [**Physical Layer Test** register *User Test Pattern Enable* bit for SerDes Quad 1 (Port 0, offset 228h[29]) is set to 1].	RO	No	00h
15:8	**Actual Data** Actual UTP SerDes[4, 6] Diagnostic data when UTP is enabled [**Physical Layer Test** register *User Test Pattern Enable* bit for SerDes Quad 1 (Port 0, offset 228h[29]) is set to 1].	RO	No	00h
19:16	**Receiver Detected SerDes[4, 6]**	RO	No	0h
23:20	***Reserved***	RO	No	0h
31:24	**PRBS Error Count** PRBS SerDes[4, 6] Error Count (0 to 255).	RO	No	00h

Register 13-74. 240h Quad 2 SerDes Diagnostic Data (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**UTP Expected Data** Expected User Test Pattern (UTP) SerDes[8, 10] Diagnostic data when UTP is enabled [**Physical Layer Test** register *User Test Pattern Enable* bit for SerDes Quad 2 (Port 0, offset 228h[30]) is set to 1].	RO	No	00h
15:8	**Actual Data** Actual UTP SerDes[8, 10] Diagnostic data when UTP is enabled [**Physical Layer Test** register *User Test Pattern Enable* bit for SerDes Quad 2 (Port 0, offset 228h[30]) is set to 1].	RO	No	00h
19:16	**Receiver Detected SerDes[8, 10]**	RO	No	0h
23:20	***Reserved***	RO	No	0h
31:24	**PRBS Error Count** PRBS SerDes[8, 10] Error Count (0 to 255).	RO	No	00h

Register 13-75. 248h SerDes Nominal Drive Current Select (Only Port 0 and NT Port Link Interface)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
1:0	**SerDes_0 Nominal Drive Current**	The following values for Nominal Current apply to each drive: • 00b = 20 mA • 01b = 10 mA • 10b = 28 mA • 11b = 20 mA	RWS	Yes	00b
3:2	**SerDes_1 Nominal Drive Current**		RWS	Yes	00b
5:4	**SerDes_2 Nominal Drive Current**		RWS	Yes	00b
7:6	**SerDes_3 Nominal Drive Current**		RWS	Yes	00b
9:8	**SerDes_4 Nominal Drive Current**		RWS	Yes	00b
11:10	***Reserved***		RsvdP	No	00b
13:12	**SerDes_6 Nominal Drive Current**		RWS	Yes	00b
15:14	***Reserved***		RsvdP	No	00b
17:16	**SerDes_8 Nominal Drive Current**		RWS	Yes	00b
19:18	***Reserved***		RsvdP	No	00b
21:20	**SerDes_10 Nominal Drive Current**		RWS	Yes	00b
31:22	***Reserved***		RsvdP	No	0-0h

Note: *Port 0 parameters apply to SerDes[0-3], SerDes[4, 6], and SerDes[8, 10]. SerDes[5, 7] and SerDes[9, 11] are* ***reserved****.*

Register 13-76. 24Ch SerDes Drive Current Level Select 1 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
3:0	**SerDes 0 Drive Current Level**	The following values represent the ratio of Actual Current/Nominal Current (selected in **SerDes Nominal Drive Current Select** register) and apply to each drive: 0h = 1.00 8h = 0.60 1h = 1.05 9h = 0.65 2h = 1.10 Ah = 0.70 3h = 1.15 Bh = 0.75 4h = 1.20 Ch = 0.80 5h = 1.25 Dh = 0.85 6h = 1.30 Eh = 0.90 7h = 1.35 Fh = 0.95	RWS	Yes	0h
7:4	**SerDes 1 Drive Current Level**		RWS	Yes	0h
11:8	**SerDes 2 Drive Current Level**		RWS	Yes	0h
15:12	**SerDes 3 Drive Current Level**		RWS	Yes	0h
19:16	**SerDes 4 Drive Current Level**		RWS	Yes	0h
23:20	***Reserved***		RsvdP	No	0h
27:24	**SerDes 6 Drive Current Level**		RWS	Yes	0h
31:28	***Reserved***		RsvdP	No	0h

Note: *Port 0 parameters apply to SerDes[0-4, 6]. SerDes[5, 7] are* ***reserved****.*

Register 13-77. 250h SerDes Drive Current Level Select 2 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
3:0	**SerDes_8 Drive Current Level**	The following values represent the ratio of Actual Current/Nominal Current (selected in **SerDes Nominal Drive Current Select** register) and apply to each drive: 0h = 1.00 8h = 0.60 1h = 1.05 9h = 0.65 2h = 1.10 Ah = 0.70 3h = 1.15 Bh = 0.75 4h = 1.20 Ch = 0.80 5h = 1.25 Dh = 0.85 6h = 1.30 Eh = 0.90 7h = 1.35 Fh = 0.95	RWS	Yes	0h
7:4	***Reserved***		RsvdP	No	0h
11:8	**SerDes_10 Drive Current Level**		RWS	Yes	0h
31:12	***Reserved***		RsvdP	No	0h

Note: *Port 0 parameters apply to SerDes[8, 10]. SerDes[9, 11] are* ***reserved****.*

Register 13-78. 254h SerDes Drive Equalization Level Select 1 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description		Type	Serial EEPROM and I^2C	Default
3:0	**SerDes 0 Drive Equalization Level**	The following values represent the percentage of Drive Current attributable to Equalization Current and apply to each drive:	RWS	Yes	8h
7:4	**SerDes 1 Drive Equalization Level**		RWS	Yes	8h
11:8	**SerDes 2 Drive Equalization Level**		RWS	Yes	8h
15:12	**SerDes 3 Drive Equalization Level**		RWS	Yes	8h
19:16	**SerDes 4 Drive Equalization Level**		RWS	Yes	8h
23:20	***Reserved***		RsvdP	No	0h
27:24	**SerDes 6 Drive Equalization Level**		RWS	Yes	8h
31:28	***Reserved***		RsvdP	No	0h

I_{EQ}/I_{DR}	De-Emphasis (dB)
0h = 0.00	0.00
1h = 0.04	-0.35
2h = 0.08	-0.72
3h = 0.12	-1.11
4h = 0.16	-1.51
5h = 0.20	-1.94
6h = 0.24	-2.38
7h = 0.28	-2.85
8h = 0.32	-3.35
9h = 0.36	-3.88
Ah = 0.40	-4.44
Bh = 0.44	-5.04
Ch = 0.48	-5.68
Dh = 0.52	-6.38
Eh = 0.56	-7.13
Fh = 0.60	-7.96

Note: *Port 0 parameters apply to SerDes[0-4, 6]. SerDes[5, 7] are* ***reserved.***

Register 13-79. 258h SerDes Drive Equalization Level Select 2 (Only Port 0 and NT Port Link Interface)

Bit(s)	Description		Type	Serial EEPROM and I^2C	Default
3:0	**SerDes_8 Drive Equalization Level**	The following values represent the percentage of Drive Current attributable to Equalization Current and apply to each drive:	RWS	Yes	8h
7:4	***Reserved***		RsvdP	No	0h
11:8	**SerDes_10 Drive Equalization Level**		RWS	Yes	8h
31:12	***Reserved***		RsvdP	No	0h

I_{EQ}/I_{DR}	De-Emphasis (dB)
0h = 0.00	0.00
1h = 0.04	-0.35
2h = 0.08	-0.72
3h = 0.12	-1.11
4h = 0.16	-1.51
5h = 0.20	-1.94
6h = 0.24	-2.38
7h = 0.28	-2.85
8h = 0.32	-3.35
9h = 0.36	-3.88
Ah = 0.40	-4.44
Bh = 0.44	-5.04
Ch = 0.48	-5.68
Dh = 0.52	-6.38
Eh = 0.56	-7.13
Fh = 0.60	-7.96

Note: *Port 0 parameters apply to SerDes[8, 10]. SerDes[9, 11] are* ***reserved****.*

Register 13-80. 25Ch Physical Receive Error Count (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Port 0 Receive Error Count** Reports the number of PHY packets received on Port 0 with CRC errors. Clears on Read.	RO	No	00h
15:8	**Port 1 Receive Error Count** Reports the number of PHY packets received on Port 1 with CRC errors. Clears on Read.	RO	No	00h
23:16	**Port 2 Receive Error Count** Reports the number of PHY packets received on Port 2 with CRC errors. Clears on Read.	RO	No	00h
31:24	**Port 3 Receive Error Count** Reports the number of PHY packets received on Port 3 with CRC errors. Clears on Read.	RO	No	00h

Register 13-81. 260h Serial EEPROM Status and Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Serial EEPROM Control				
12:0	**EepBlkAddr** Serial EEPROM Block Address for 32 KB.	RW	Yes	0000h
15:13	**EepCmd[2:0]** Commands to the Serial EEPROM Controller. 000b = ***Reserved*** 001b = Data from bits [31:24] (**Status Data from Serial EEPROM** register) are written to the Serial EEPROM internal **Status** register 010b = Write four bytes of data from the *EepBuf* into the memory location pointed to by the *EepBlkAddr* field 011b = Read four bytes of data from the memory location pointed to by the *EepBlkAddr* field into the *EepBuf* 100b = Reset Write Enable latch 101b = Data from Serial EEPROM internal **Status** register written to bits [31:24] (**Status Data from Serial EEPROM** register) 110b = Set Write Enable latch 111b = ***Reserved*** ***Note:*** *For value of 001b, only bits [31, 27:26] can be written into the serial EEPROM's internal* ***Status*** *register.*	RW	Yes	000b

Register 13-81. 260h Serial EEPROM Status and Control (Only Port 0) *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Serial EEPROM Status			
17:16	**EepPrsnt[1:0]** Serial EEPROM Present status, unless bit 21 (*CRC Disable* bit) is set to ignore CRC checking ***(not recommended)***. 00b = Not present 01b = Serial EEPROM is present – no CRC error 10b = ***Reserved*** 11b = Serial EEPROM is present, but with CRC error – unless bit 21 (*CRC Disable* bit) is set to ignore CRC checking ***(not recommended)***	RO	Yes	00b
19:18	**EepCmdStatus** Serial EEPROM Command status. 00b = Serial EEPROM Command is complete 01b = Serial EEPROM Command is not complete 10b = Serial EEPROM Command is complete, with CRC error 11b = ***Reserved***	RO	Yes	00b
20	**EepBlkAddrUp** Serial EEPROM Block Address upper bit 13. Extends the serial EEPROM to 64 KB.	RW	Yes	0
21	**CRC Disable** 0 = Serial EEPROM input data uses CRC 1 = Serial EEPROM input data CRC is disabled ***(not recommended)***	RW	Yes	0
23:22	***Reserved***	RsvdP	Yes	00b

Register 13-81. 260h Serial EEPROM Status and Control (Only Port 0) *(Cont.)*

<table>
<tr><th>Bit(s)</th><th>Description</th><th>Type</th><th>Serial EEPROM and I²C</th><th>Default</th></tr>
<tr><td colspan="5">Status Data from Serial EEPROM[a]</td></tr>
<tr><td>24</td><td>EepRdy
Serial EEPROM RDY#.
0 = Serial EEPROM is ready to transmit data
1 = Write cycle is in progress</td><td>RW</td><td>Yes</td><td>0</td></tr>
<tr><td>25</td><td>EepWen
Serial EEPROM Write Enable.
0 = Serial EEPROM Write is disabled
1 = Serial EEPROM Write is enabled</td><td>RW</td><td>Yes</td><td>0</td></tr>
<tr><td>27:26</td><td>EepBp[1:0]
Serial EEPROM Block-Write Protect bits. Block Protection options protect the top ¼, top ½, or the entire serial EEPROM. PEX 8508 Configuration data is stored in the lower addresses; therefore, when using Block Protection, the entire serial EEPROM should be protected with BP[1:0]=11b:
<table>
<tr><th rowspan="2">BP[1:0]</th><th rowspan="2">Level</th><th colspan="4">Array Addresses Protected</th></tr>
<tr><th>8-KB Device</th><th>16-KB Device</th><th>32-KB Device</th><th>64-KB Device</th></tr>
<tr><td>00b</td><td>0</td><td>None</td><td>None</td><td>None</td><td>None</td></tr>
<tr><td>01b</td><td>1 (top 1/4)</td><td>1800h - 1FFFh</td><td>3000h - 3FFFh</td><td>6000h - 7FFFh</td><td>–</td></tr>
<tr><td>10b</td><td>2 (top 1/2)</td><td>1000h - 1FFFh</td><td>2000h - 3FFFh</td><td>4000h - 7FFFh</td><td>–</td></tr>
<tr><td>11b</td><td>3 (All)</td><td>0000h - 1FFFh</td><td>0000h - 3FFFh</td><td>0000h - 7FFFh</td><td>–</td></tr>
</table></td><td>RW</td><td>Yes</td><td>00b</td></tr>
<tr><td>30:28</td><td>EepWrStatus
Serial EEPROM Write status. Value is 000b when the serial EEPROM is not in an internal Write cycle.
Note: Definition of this field varies among serial EEPROM manufacturers. Reads of the serial EEPROM internal Status register can return 000b or 111b, depending upon the serial EEPROM that is used.</td><td>RO</td><td>Yes</td><td>000b</td></tr>
<tr><td>31</td><td>EepWpen
Serial EEPROM Write Protect Enable. Overrides the internal serial EEPROM Write Protect WP# input and enables/disables Writes to the Serial EEPROM Status register:
• When WP# is High or EepWpen = 0, and EepWen = 1, the Serial EEPROM Status register is writable
• When WP# is Low and EepWpen = 1, or EepWen = 0, the Serial EEPROM Status register is protected
Notes: If the internal serial EEPROM Write Protect WP# input is Low, after software sets the EepWen bit to write-protect the Serial EEPROM Status register, the EepWen value cannot be changed to 0, nor can the EepBp[1:0] bits be cleared to disable Block Protection, until WP# is high.
This bit is not implemented in certain serial EEPROMs. Refer to the serial EEPROM manufacturer's data sheet.</td><td>RW</td><td>Yes</td><td>0</td></tr>
</table>

*a. Within the serial EEPROM's internal **Status** register, only bits [31, 27:26] can be written.*

Register 13-82. 264h Serial EEPROM Data Buffer (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
31:0	**EepBuf** Serial EEPROM R/W buffer. Read/Write command to the corresponding **Serial EEPROM Control** register results in a 4-byte Read/Write to or from the serial EEPROM device.	RW	Yes	0-0h

13.13.3 I²C Interface Register

Table 13-16 defines the I²C interface register map.

Table 13-16. I²C Interface Register Map (Only Port 0)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0

Factory Test Only	290h
I2C Configuration	294h
Factory Test Only	298h – 2A8h
Reserved	2ACh – 2C4h

Register 13-83. 294h I²C Configuration (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default	
2:0	**Slave Address** Bits [6:0] comprise the I²C Slave address, 70h – bits [2:0] are strapped, using the I2C_ADDR[2:0] balls. ***Note: The I²C Slave address must not be changed by an I²C Write command.***	HwInit	Yes	111b	70h
6:3	**Slave Address** Bits [6:0] comprise the I²C Slave address, 70h – bits [6:3] default to 1110b. ***Note: The I²C Slave address must not be changed by an I²C Write command.***	RW	Yes	1110b	
9:7	***Reserved***	RsvdP	No	000b	
10	***Factory Test Only***	RW	Yes	0	
31:11	***Reserved***	RW	Yes	0000_00h	

13.13.4 Bus Number CAM Registers

The **Bus Number Content-Addressable Memory (CAM)** registers, defined in Table 13-17, are used to determine the Configuration TLP Completion route. These registers contain mirror copies of the **Primary Bus Number**, **Secondary Bus Number**, and **Subordinate Bus Number** registers of each PEX 8508 port.

These registers are automatically updated by hardware. Modifying these registers by writing to the addresses listed here is not recommended.

Table 13-17. Device-Specific Bus Number CAM Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Reserved	Bus Number CAM 0			2C8h
Reserved	Bus Number CAM 1			2CCh
Reserved	Bus Number CAM 2			2D0h
Reserved	Bus Number CAM 3			2D4h
Reserved	Bus Number CAM 4			2D8h
Reserved			2DCh –	304h

Register 13-84. 2C8h Bus Number CAM 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Primary Bus Number** Mirror copy of Port 0 Primary Bus Number.	RW	Yes	00h
15:8	**Secondary Bus Number** Mirror copy of Port 0 Secondary Bus Number.	RW	Yes	FFh
23:16	**Subordinate Bus Number** Mirror copy of Port 0 Subordinate Bus Number.	RW	Yes	00h
31:24	***Reserved***	RsvdP	No	00h

Register 13-85. 2CCh Bus Number CAM 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Primary Bus Number** Mirror copy of Port 1 Primary Bus Number.	RW	Yes	00h
15:8	**Secondary Bus Number** Mirror copy of Port 1 Secondary Bus Number.	RW	Yes	FFh
23:16	**Subordinate Bus Number** Mirror copy of Port 1 Subordinate Bus Number.	RW	Yes	00h
31:24	***Reserved***	RsvdP	No	00h

Register 13-86. 2D0h Bus Number CAM 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Primary Bus Number** Mirror copy of Port 2 Primary Bus Number.	RW	Yes	00h
15:8	**Secondary Bus Number** Mirror copy of Port 2 Secondary Bus Number.	RW	Yes	FFh
23:16	**Subordinate Bus Number** Mirror copy of Port 2 Subordinate Bus Number.	RW	Yes	00h
31:24	***Reserved***	RsvdP	No	00h

Register 13-87. 2D4h Bus Number CAM 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Primary Bus Number** Mirror copy of Port 3 Primary Bus Number.	RW	Yes	00h
15:8	**Secondary Bus Number** Mirror copy of Port 3 Secondary Bus Number.	RW	Yes	FFh
23:16	**Subordinate Bus Number** Mirror copy of Port 3 Subordinate Bus Number.	RW	Yes	00h
31:24	***Reserved***	RsvdP	No	00h

Register 13-88. 2D8h Bus Number CAM 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Primary Bus Number** Mirror copy of Port 4 Primary Bus Number.	RW	Yes	00h
15:8	**Secondary Bus Number** Mirror copy of Port 4 Secondary Bus Number.	RW	Yes	FFh
23:16	**Subordinate Bus Number** Mirror copy of Port 4 Subordinate Bus Number.	RW	Yes	00h
31:24	***Reserved***	RsvdP	No	00h

13.13.5 I/O CAM Registers

The **I/O CAM** registers, defined in Table 13-18, are used to determine I/O request routing. These registers contain mirror copies of the **I/O Base** and **I/O Limit** registers of each PEX 8508 port.

These registers are automatically updated by hardware. Modifying these registers by writing to the addresses listed here is not recommended.

Table 13-18. Device-Specific I/O CAM Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
I/O CAM 1		I/O CAM 0		308h
I/O CAM 3		I/O CAM 2		30Ch
Reserved		I/O CAM 4		310h
Reserved				314h – 344h

Register 13-89. 308h I/O CAM 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	1h
7:4	**I/O Base** Mirror copy of Port 0 I/O Base value.	RW	Yes	Fh
11:8	***Reserved***	RsvdP	No	1h
15:12	**I/O Limit** Mirror copy of Port 0 I/O Limit value.	RW	Yes	0h

Register 13-90. 30Ah I/O CAM 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	1h
7:4	**I/O Base** Mirror copy of Port 1 I/O Base value.	RW	Yes	Fh
11:8	***Reserved***	RsvdP	No	1h
15:12	**I/O Limit** Mirror copy of Port 1 I/O Limit value.	RW	Yes	0h

Register 13-91. 30Ch I/O CAM 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	1h
7:4	**I/O Base** Mirror copy of Port 2 I/O Base value.	RW	Yes	Fh
11:8	***Reserved***	RsvdP	No	1h
15:12	**I/O Limit** Mirror copy of Port 2 I/O Limit value.	RW	Yes	0h

Register 13-92. 30Eh I/O CAM 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	1h
7:4	**I/O Base** Mirror copy of Port 3 I/O Base value.	RW	Yes	Fh
11:8	***Reserved***	RsvdP	No	1h
15:12	**I/O Limit** Mirror copy of Port 3 I/O Limit value.	RW	Yes	0h

Register 13-93. 310h I/O CAM 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	1h
7:4	**I/O Base** Mirror copy of Port 4 I/O Base value.	RW	Yes	Fh
11:8	***Reserved***	RsvdP	No	1h
15:12	**I/O Limit** Mirror copy of Port 4 I/O Limit value.	RW	Yes	0h

13.13.6 Address-Mapping CAM (AMCAM) Registers

The **AMCAM** registers, defined in Table 13-19, are used to used to determine Memory request routing. These registers contain mirror copies of the **Memory Base and Limit**. **Prefetchable Memory Base and Limit**, **Prefetchable Memory Upper Base Address**, and **Prefetchable Memory Upper Limit Address** registers of each PEX 8508 port.

These registers are automatically updated by hardware. Modifying these registers by writing to the addresses listed here is not recommended.

Table 13-19. Device-Specific AMCAM Register Map (Only Port 0)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
AMCAM 0 Memory Limit	AMCAM 0 Memory Base	348h
AMCAM 0 Prefetchable Memory Limit	AMCAM 0 Prefetchable Memory Base	34Ch
AMCAM 0 Prefetchable Memory Upper Base Address		350h
AMCAM 0 Prefetchable Memory Upper Limit Address		354h
AMCAM 1 Memory Limit	AMCAM 1 Memory Base	358h
AMCAM 1 Prefetchable Memory Limit	AMCAM 1 Prefetchable Memory Base	35Ch
AMCAM 1 Prefetchable Memory Upper Base Address		360h
AMCAM 1 Prefetchable Memory Upper Limit Address		364h
AMCAM 2 Memory Limit	AMCAM 2 Memory Base	368h
AMCAM 2 Prefetchable Memory Limit	AMCAM 2 Prefetchable Memory Base	36Ch
AMCAM 2 Prefetchable Memory Upper Base Address		370h
AMCAM 2 Prefetchable Memory Upper Limit Address		374h
AMCAM 3 Memory Limit	AMCAM 3 Memory Base	378h
AMCAM 3 Prefetchable Memory Limit	AMCAM 3 Prefetchable Memory Base	37Ch
AMCAM 3 Prefetchable Memory Upper Base Address		380h
AMCAM 3 Prefetchable Memory Upper Limit Address		384h
AMCAM 4 Memory Limit	AMCAM 4 Memory Base	388h
AMCAM 4 Prefetchable Memory Limit	AMCAM 4 Prefetchable Memory Base	38Ch
AMCAM 4 Prefetchable Memory Upper Base Address		390h
AMCAM 4 Prefetchable Memory Upper Limit Address		394h
Reserved		398h – 544h
Ingress Performance Counter		548h

Register 13-94. 348h AMCAM 0 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 0 Memory Base				
3:0	***Reserved***	RsvdP	No	0h
15:4	**AMCAM 0 Memory Base** Mirror copy of Port 0 Memory Base value.	RW	Yes	FFFh
AMCAM 0 Memory Limit				
19:16	***Reserved***	RsvdP	No	0h
31:20	**AMCAM 0 Memory Limit** Mirror copy of Port 0 Memory Limit value.	RW	Yes	000h

Register 13-95. 34Ch AMCAM 0 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 0 Prefetchable Memory Base				
3:0	**AMCAM 0 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	**AMCAM 0 Prefetchable Memory Base** AMCAM 0 Port 0 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 0 Prefetchable Memory Limit				
19:16	**AMCAM 0 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	**AMCAM 0 Prefetchable Memory Limit** AMCAM 0 Port 0 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-96. 350h AMCAM 0 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 0 Prefetchable Memory Base[63:32]** AMCAM 0 Port 0 Prefetchable Memory Base[63:32].	RW	Yes	0h

Register 13-97. 354h AMCAM 0 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 0 Prefetchable Memory Limit[63:32]** AMCAM 0 Port 0 Prefetchable Memory Limit[63:32].	RW	Yes	0h

Register 13-98. 358h AMCAM 1 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 1 Memory Base				
3:0	***Reserved***	RsvdP	No	0h
15:4	**AMCAM 1 Memory Base** Mirror copy of Port 1 Memory Base value.	RW	Yes	FFFh
AMCAM 1 Memory Limit				
19:16	***Reserved***	RsvdP	No	0h
31:20	**AMCAM 1 Memory Limit** Mirror copy of Port 1 Memory Limit value.	RW	Yes	000h

Register 13-99. 35Ch AMCAM 1 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 1 Prefetchable Memory Base				
3:0	**AMCAM 1 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	**AMCAM 1 Prefetchable Memory Base** AMCAM 1 Port 1 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 1 Prefetchable Memory Limit				
19:16	**AMCAM 1 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	**AMCAM 1 Prefetchable Memory Limit** AMCAM 1 Port 1 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-100. 360h AMCAM 1 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 1 Prefetchable Memory Base[63:32]** AMCAM 1 Port 1 Prefetchable Memory Base[63:32].	RW	Yes	0h

Register 13-101. 364h AMCAM 1 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 1 Prefetchable Memory Limit[63:32]** AMCAM 1 Port 1 Prefetchable Memory Limit[63:32].	RW	Yes	0h

Register 13-102. 368h AMCAM 2 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 2 Memory Base				
3:0	***Reserved***	RsvdP	No	0h
15:4	**AMCAM 2 Memory Base** Mirror copy of Port 2 Memory Base value.	RW	Yes	FFFh
AMCAM 2 Memory Limit				
19:16	***Reserved***	RsvdP	No	0h
31:20	**AMCAM 2 Memory Limit** Mirror copy of Port 2 Memory Limit value.	RW	Yes	000h

Register 13-103. 36Ch AMCAM 2 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 2 Prefetchable Memory Base				
3:0	**AMCAM 2 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	**AMCAM 2 Prefetchable Memory Base** AMCAM 2 Port 2 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 2 Prefetchable Memory Limit				
19:16	**AMCAM 2 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	**AMCAM 2 Prefetchable Memory Limit** AMCAM 2 Port 2 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-104. 370h AMCAM 2 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 2 Prefetchable Memory Base[63:32]** AMCAM 2 Port 2 Prefetchable Memory Base[63:32].	RW	Yes	0h

Register 13-105. 374h AMCAM 2 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 2 Prefetchable Memory Limit[63:32]** AMCAM 2 Port 2 Prefetchable Memory Limit[63:32].	RW	Yes	0h

Register 13-106. 378h AMCAM 3 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 3 Memory Base				
3:0	***Reserved***	RsvdP	No	0h
15:4	**AMCAM 3 Memory Base** Mirror copy of Port 3 Memory Base value.	RW	Yes	FFFh
AMCAM 3 Memory Limit				
19:16	***Reserved***	RsvdP	No	0h
31:20	**AMCAM 3 Memory Limit** Mirror copy of Port 3 Memory Limit value.	RW	Yes	000h

Register 13-107. 37Ch AMCAM 3 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 3 Prefetchable Memory Base				
3:0	**AMCAM 3 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	**AMCAM 3 Prefetchable Memory Base** AMCAM 3 Port 3 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 3 Prefetchable Memory Limit				
19:16	**AMCAM 3 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	**AMCAM 3 Prefetchable Memory Limit** AMCAM 3 Port 3 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-108. 380h AMCAM 3 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 3 Prefetchable Memory Base[63:32]** AMCAM 3 Port 3 Prefetchable Memory Base[63:32].	RW	Yes	0h

Register 13-109. 384h AMCAM 3 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 3 Prefetchable Memory Limit[63:32]** AMCAM 3 Port 3 Prefetchable Memory Limit[63:32].	RW	Yes	0h

Register 13-110. 388h AMCAM 4 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 4 Memory Base				
3:0	***Reserved***	RsvdP	No	0h
15:4	**AMCAM 4 Memory Base** Mirror copy of Port 4 Memory Base value.	RW	Yes	FFFh
AMCAM 4 Memory Limit				
19:16	***Reserved***	RsvdP	No	0h
31:20	**AMCAM 4 Memory Limit** Mirror copy of Port 4 Memory Limit value.	RW	Yes	000h

Register 13-111. 38Ch AMCAM 4 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
AMCAM 4 Prefetchable Memory Base				
3:0	**AMCAM 4 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	**AMCAM 2 Prefetchable Memory Base** AMCAM 4 Port 4 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 4 Prefetchable Memory Limit				
19:16	**AMCAM 4 Addressing Support** 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	**AMCAM 4 Prefetchable Memory Limit** AMCAM 4 Port 4 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-112. 390h AMCAM 4 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 4 Prefetchable Memory Base[63:32]** AMCAM 4 Port 4 Prefetchable Memory Base[63:32].	RW	Yes	0h

Register 13-113. 394h AMCAM 4 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**AMCAM 4 Prefetchable Memory Limit[63:32]** AMCAM 4 Port 4 Prefetchable Memory Limit[63:32].	RW	Yes	0h

Register 13-114. 548h Ingress Performance Counter (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Port 0 Ingress Queue Depth** When written, allows the maximum depth of the Ingress queue (in TLP packets) to be set for Port 0 before an overflow can occur. When read, returns the Port 0 Ingress Queue depth.	RW	Yes	00h
15:8	**Port 1 Ingress Queue Depth** When written, allows the maximum depth of the Ingress queue (in TLP packets) to be set for Port 1 before an overflow can occur. When read, returns the Port 1 Ingress Queue depth.	RW	Yes	00h
23:16	**Port 2 Ingress Queue Depth** When written, allows the maximum depth of the Ingress queue (in TLP packets) to be set for Port 2 before an overflow can occur. When read, returns the Port 2 Ingress Queue depth.	RW	Yes	00h
31:24	**Port 3 Ingress Queue Depth** When written, allows the maximum depth of the Ingress queue (in TLP packets) to be set for Port 3 before an overflow can occur. When read, returns the Port 3 Ingress Queue depth.	RW	Yes	00h

13.13.7 Ingress Control and Port Enable Registers

Table 13-20. Device-Specific Ingress Control and Port Enable Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Ingress Control				660h
Reserved				664h
Ingress Port Enable				668h
Reserved				66Ch – 67Ch

Register 13-115. 660h Ingress Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Enable CSR Access by Downstream Devices** Enables acceptance of both Configuration and Memory requests from a Requester that is downstream from a Transparent port, targeting any PEX 8508 registers. 0 = Configuration requests from a downstream device that is targeting PEX 8508 registers are ***not supported***; the downstream port flags an Uncorrectable Error, and, returns a Completion with Unsupported Request (UR) specified in the *Completion Status* field, to the downstream Requester. Only this mode is *PCI Express Base r1.1*-compliant. 1 = Configuration and Memory requests from downstream Requesters, targeting any PEX 8508 registers in any port, are allowed. ***Note:*** *This bit can be initially set only through the upstream port, the NT Port Link Interface, the I²C interface, or by the serial EEPROM, to enable register access through the downstream Transparent ports; a Requester downstream from a Transparent port cannot set the bit to grant itself (or peers) access to PEX 8508 registers. Configuration requests can access those registers that are defined by PCI-SIG specifications, and generally cannot access device-specific registers other than the NT Port* ***Cursor Mechanism*** *registers. Memory requests can access all PEX 8508 registers.*	RW	Yes	0
1	**Configuration Write to Device-Specific Register without Unsupported Request** When set 1, disables completions with an Unsupported Request status from being returned when Configuration Writes are attempted on Device-Specific registers.	RW	Yes	0
25:2	***Factory Test Only***	RW	Yes	0-0h
26	**Disable Upstream Port BAR0 and BAR1** 0 = Enables upstream **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively) 1 = Disables upstream **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively)	RW	Yes	0

Register 13-115. 660h Ingress Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
28:27	***Factory Test Only***	RW	Yes	00b
29	**Disable Configuration Access from Upstream Port** 0 = Enables Configuration access from the upstream port 1 = Disables Configuration access from the upstream port	RW	Yes	0
30	***Factory Test Only***	RW	Yes	0
31	**Memory-Mapped Configuration Space Register Access Enable** Affects all ports. 0 = Disabled 1 = Enabled	RW	Yes	0

Register 13-116. 668h Ingress Port Enable (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Ingress Port Enable** The value of this register depends upon the number of ports used, which is dependent on the **Port Configuration** register *Port Configuration* field (offset 224h[4:0]) value or STRAP_PORTCFG[4:0] balls. Set the upper 16 bits to FFFFh. For the lower 16 bits, set each bit position that corresponds to an enabled port. Table 13-21 provides a sample mapping, based upon the Port Configuration value.	RW	Yes	FFFF_000Fh

Table 13-21. PEX 8508 Ingress Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[4:0])	Lane Width/Port					Ingress Port Enable Register Value (Port 0, Offset 668h)
	0	1	2	3	4	
0h	x2	x2	x2	x2		FFFF_000Fh
2h	x4	x4				FFFF_0003h
3h	x4	x2	x2			FFFF_0007h
4h	x4	x2	x1	x1		FFFF_000Fh
5h	x4	x1	x1	x2		FFFF_000Fh
6h	x4	x1	x2	x1		FFFF_000Fh
8h	x4	x1	x1	x1	x1	FFFF_001Fh
9h	x2	x2	x2	x1	x1	FFFF_001Fh

13.13.8 I/O CAM Base and Limit Upper 16-Bit Registers

Table 13-22. Device-Specific I/O CAM Base and Limit Upper 16-Bit Register Map (Only Port 0)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
I/O CAM Limit[31:16] Upper Port 0	I/O CAM Base[31:16] Upper Port 0	680h
I/O CAM Limit[31:16] Upper Port 1	I/O CAM Base[31:16] Upper Port 1	684h
I/O CAM Limit[31:16] Upper Port 2	I/O CAM Base[31:16] Upper Port 2	688h
I/O CAM Limit[31:16] Upper Port 3	I/O CAM Base[31:16] Upper Port 3	68Ch
I/O CAM Limit[31:16] Upper Port 4	I/O CAM Base[31:16] Upper Port 4	690h
Reserved		694h – 6BCh

Register 13-117. 680h I/O CAM Upper Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**I/O CAM Base[31:16]** I/O Base Upper 16 bits.	RW	Yes	FFFFh
31:16	**I/O CAM Limit[31:16]** I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-118. 684h I/O CAM Upper Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**I/O CAM Base[31:16]** I/O Base Upper 16 bits.	RW	Yes	FFFFh
31:16	**I/O CAM Limit[31:16]** I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-119. 688h I/O CAM Upper Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**I/O CAM Base[31:16]** I/O Base Upper 16 bits.	RW	Yes	FFFFh
31:16	**I/O CAM Limit[31:16]** I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-120. 68Ch I/O CAM Upper Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**I/O CAM Base[31:16]** I/O Base Upper 16 bits.	RW	Yes	FFFFh
31:16	**I/O CAM Limit[31:16]** I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-121. 690h I/O CAM Upper Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**I/O CAM Base[31:16]** I/O Base Upper 16 bits.	RW	Yes	FFFFh
31:16	**I/O CAM Limit[31:16]** I/O Limit Upper 16 bits.	RW	Yes	0000h

13.13.9 Base Address Shadow Registers (BARs)

The registers defined in Table 13-23 contain a shadow copy of the two Type 1 Configuration Base Address registers for each PEX 8508 port.

Table 13-23. Device-Specific BARs for Shadow Registers Register Map (Only Port 0)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
BAR0 Shadow for Port 0	6C0h
BAR1 Shadow for Port 0	6C4h
BAR0 Shadow for Port 1	6C8h
BAR1 Shadow for Port 1	6CCh
BAR0 Shadow for Port 2	6D0h
BAR1 Shadow for Port 2	6D4h
BAR0 Shadow for Port 3	6D8h
BAR1 Shadow for Port 3	6DCh
BAR0 Shadow for Port 4	6E0h
BAR1 Shadow for Port 4	6E4h
Reserved	6E8h – 73Ch

Register 13-122. 6C0h BAR0 Shadow for Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR 1 = I/O BAR	RO	No	0
2:1	**Memory Map Type** Memory Mapping for Port 0. 00b = 32-bit mode 10b = 64-bit mode 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Not Prefetchable 1 = Prefetchable	RO	Yes	0
16:4	***Reserved***	RsvdP	No	000h
31:17	**Base Address** Shadow copy of Port 0 Base Address 0.	RW	Yes	0000h

Register 13-123. 6C4h BAR1 Shadow for Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 1[63:32]** When the **BAR0 Shadow for Port 0** register *Memory Map Type* field (offset 6C0h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 0 Base Address 1[63:32].	RW	Yes	0000_0000h
	Reserved when the **BAR0 Shadow for Port 0** register *Memory Map Type* field (offset 6C0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-124. 6C8h BAR0 Shadow for Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR 1 = I/O BAR	RO	No	0
2:1	**Memory Map Type** Memory Mapping for Port 1. 00b = 32-bit mode 10b = 64-bit mode 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Not Prefetchable 1 = Prefetchable	RO	Yes	0
16:4	***Reserved***	RsvdP	No	000h
31:17	**Base Address** Shadow copy of Port 1 Base Address 0.	RW	Yes	0000h

Register 13-125. 6CCh BAR1 Shadow for Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 1[63:32]** When the **BAR0 Shadow for Port 1** register *Memory Map Type* field (offset 6C8h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 1 Base Address 1[63:32].	RW	Yes	0000_0000h
	Reserved when the **BAR0 Shadow for Port 1** register *Memory Map Type* field (offset 6C8h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-126. 6D0h BAR0 Shadow for Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR 1 = I/O BAR	RO	No	0
2:1	**Memory Map Type** Memory Mapping for Port 2. 00b = 32-bit mode 10b = 64-bit mode 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Not Prefetchable 1 = Prefetchable	RO	Yes	0
16:4	***Reserved***	RsvdP	No	000h
31:17	**Base Address** Shadow copy of Port 2 Base Address 0.	RW	Yes	0000h

Register 13-127. 6D4h BAR1 Shadow for Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 1[63:32]** When the **BAR0 Shadow for Port 2** register *Memory Map Type* field (offset 6D0h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 2 Base Address 1[63:32].	RW	Yes	0000_0000h
	Reserved when the **BAR0 Shadow for Port 2** register *Memory Map Type* field (offset 6D0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-128. 6D8h BAR0 Shadow for Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR 1 = I/O BAR	RO	No	0
2:1	**Memory Map Type** Memory Mapping for Port 3. 00b = 32-bit mode 10b = 64-bit mode 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Not Prefetchable 1 = Prefetchable	RO	Yes	0
16:4	***Reserved***	RsvdP	No	000h
31:17	**Base Address** Shadow copy of Port 3 Base Address 0.	RW	Yes	0000h

Register 13-129. 6DCh BAR1 Shadow for Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 1[63:32]** When the **BAR0 Shadow for Port 3** register *Memory Map Type* field (offset 6D8h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 3 Base Address 1[63:32].	RW	Yes	0000_0000h
	Reserved when the **BAR0 Shadow for Port 3** register *Memory Map Type* field (offset 6D8h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-130. 6E0h BAR0 Shadow for Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR 1 = I/O BAR	RO	No	0
2:1	**Memory Map Type** Memory Mapping for Port 4. 00b = 32-bit mode 10b = 64-bit mode 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Not Prefetchable 1 = Prefetchable	RO	Yes	0
16:4	***Reserved***	RsvdP	No	000h
31:17	**Base Address** Shadow copy of Port 4 Base Address 0.	RW	Yes	0000h

Register 13-131. 6E4h BAR1 Shadow for Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 1[63:32]** When the **BAR0 Shadow for Port 4** register *Memory Map Type* field (offset 6E0h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 4 Base Address 1[63:32].	RW	Yes	0000_0000h
	Reserved when the **BAR0 Shadow for Port 4** register *Memory Map Type* field (offset 6E0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

13.13.10 Shadow Virtual Channel (VC) Capability Registers

Table 13-24. Device-Specific Shadow Virtual Channel (VC) Capability Register Map (Only Port 0)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
VC0 Capability Port 0	740h
VC1 Capability Port 0	744h
VC0 Capability Port 1	748h
VC1 Capability Port 1	74Ch
VC0 Capability Port 2	750h
VC1 Capability Port 2	754h
VC0 Capability Port 3	758h
Reserved	75Ch
VC0 Capability Port 4	760h
Reserved	764h – 83Ch

Register 13-132. 740h VC0 Capability Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC VC0 Port 0 Map[0]** Always mapped to Virtual Channel 0. Tied to 1.	RO	Yes	1
7:1	**TC VC0 Port 0 Map[7:1]** Mapped to Virtual Channel 0 by default. Software can change this field during enumeration or when assigning traffic to the traffic class.	RW	Yes	7Fh
23:8	***Reserved***	RsvdP	No	000h
24	**VC0 Port 0 ID** Port 0 Virtual Channel 0 ID.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	00h
31	**VC0 Port 0 Enable** Port 0 Virtual Channel 0 Enable.	RO	Yes	1

Register 13-133. 744h VC1 Capability Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
7:1	**TC VC1 Port 0 Map[7:1]** Mapped to Virtual Channel 1 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	00h
23:8	***Reserved***	RsvdP	No	000h
24	**VC1 Port 0 ID** Port 0 Virtual Channel 1 ID.	RW	Yes	1
30:25	***Reserved***	RsvdP	No	00h
31	**VC1 Port 0 Enable** Port 0 Virtual Channel 1 Enable.	RW	Yes	0

Register 13-134. 748h VC0 Capability Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC VC0 Port 1 Map[0]** Always mapped to Virtual Channel 0. Tied to 1.	RO	Yes	1
7:1	**TC VC0 Port 1 Map[7:1]** Mapped to Virtual Channel 0 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	7Fh
23:8	***Reserved***	RsvdP	No	000h
24	**VC0 Port 1 ID** Port 1 Virtual Channel 0 ID.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	00h
31	**VC0 Port 1 Enable** Port 1 Virtual Channel 0 Enable.	RO	Yes	1

Register 13-135. 74Ch VC1 Capability Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
7:1	**TC VC1 Port 1 Map[7:1]** Mapped to Virtual Channel 1 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	00h
23:8	***Reserved***	RsvdP	No	000h
24	**VC1 Port 1 ID** Port 1 Virtual Channel 1 ID.	RW	Yes	1
30:25	***Reserved***	RsvdP	No	00h
31	**VC1 Port 1 Enable** Port 1 Virtual Channel 1 Enable.	RW	Yes	0

Register 13-136. 750h VC0 Capability Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC VC0 Port 2 Map[0]** Always mapped to Virtual Channel 0. Tied to 1.	RO	Yes	1
7:1	**TC VC0 Port 2 Map[7:1]** Mapped to Virtual Channel 0 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	7Fh
23:8	***Reserved***	RsvdP	No	000h
24	**VC0 Port 2 ID** Port 2 Virtual Channel 0 ID.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	00h
31	**VC0 Port 2 Enable** Port 2 Virtual Channel 0 Enable.	RO	Yes	1

Register 13-137. 754h VC1 Capability Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
7:1	**TC VC1 Port 2 Map[7:1]** Mapped to Virtual Channel 1 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	00h
23:8	***Reserved***	RsvdP	No	000h
24	**VC1 Port 2 ID** Port 2 Virtual Channel 1 ID.	RW	Yes	1
30:25	***Reserved***	RsvdP	No	00h
31	**VC1 Port 2 Enable** Port 2 Virtual Channel 1 Enable.	RW	Yes	0

Register 13-138. 758h VC0 Capability Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC VC0 Port 3 Map[0]** Always mapped to Virtual Channel 0. Tied to 1.	RO	Yes	1
7:1	**TC VC0 Port 3 Map[7:1]** Mapped to Virtual Channel 0 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	7Fh
23:8	***Reserved***	RsvdP	No	000h
24	**VC0 Port 3 ID** Port 3 Virtual Channel 0 ID.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	00h
31	**VC0 Port 3 Enable** Port 3 Virtual Channel 0 Enable.	RO	Yes	1

Register 13-139. 760h VC0 Capability Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC VC0 Port 4 Map[0]** Always mapped to Virtual Channel 0. Tied to 1.	RO	Yes	1
7:1	**TC VC0 Port 4 Map[7:1]** Mapped to Virtual Channel 0 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	7Fh
23:8	***Reserved***	RsvdP	No	000h
24	**VC0 Port 4 ID** Port 4 Virtual Channel 0 ID.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	00h
31	**VC0 Port 4 Enable** Port 4 Virtual Channel 0 Enable.	RO	Yes	1

13.13.11 Shadow Port Virtual Channel (VC) Capability Registers

Table 13-25. Device-Specific Shadow Port Virtual Channel (VC) Capability Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Port 0 VC Capability 1				840h
Port 1 VC Capability 1				844h
Port 2 VC Capability 1				848h
Port 3 VC Capability 1				84Ch
Port 4 VC Capability 1				850h
Reserved				854h – 9ECh

Register 13-140. 840h Port 0 VC Capability 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Low-Priority Extended VC Count** Low-Priority Virtual Channel Count.	RW	Yes	0
31:5	***Reserved***	RsvdP	No	0-0h

Register 13-141. 844h Port 1 VC Capability 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Low-Priority Extended VC Count** Low-Priority Virtual Channel Count.	RW	Yes	0
31:5	***Reserved***	RsvdP	No	0-0h

Register 13-142. 848h Port 2 VC Capability 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Low-Priority Extended VC Count** Low-Priority Virtual Channel Count.	RW	Yes	0
31:5	***Reserved***	RsvdP	No	0-0h

Register 13-143. 84Ch Port 3 VC Capability 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Low-Priority Extended VC Count** Low-Priority Virtual Channel Count.	RW	Yes	0
31:5	***Reserved***	RsvdP	No	0-0h

Register 13-144. 850h Port 4 VC Capability 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Low-Priority Extended VC Count** Low-Priority Virtual Channel Count.	RW	Yes	0
31:5	***Reserved***	RsvdP	No	0-0h

13.13.12 Ingress Credit Handler (INCH) Registers

Changing credits values from default register values must be done carefully. Credits must be programmed properly, otherwise the device will not function as expected. Also, there are minimal required Header credits for all the flows, which are required to achieve reasonable performance. (Refer to the *PEX 85XX EEPROM – PEX 8518/8517/8508 Design Note*, Section 6.9, "INCH Threshold Port Virtual Channel Registers," for details.) The minimum initial Payload credits for Posted and Completions must exceed the required credits for a Maximum Payload Size TLP by 8.

Table 13-26. Device-Specific Ingress Credit Handler (INCH) Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
INCH Port 4 Control				9F0h
INCH FC Update Pending Timer				9F4h
Reserved				9F8h
Reserved		INCH Mode		9FCh
INCH Threshold Port 0 VC0 Posted				A00h
INCH Threshold Port 0 VC0 Non-Posted				A04h
INCH Threshold Port 0 VC0 Completion				A08h
INCH Threshold Port 0 VC1 Posted				A0Ch
INCH Threshold Port 0 VC1 Non-Posted				A10h
INCH Threshold Port 0 VC1 Completion				A14h
INCH Threshold Port 1 VC0 Posted				A18h
INCH Threshold Port 1 VC0 Non-Posted				A1Ch
INCH Threshold Port 1 VC0 Completion				A20h
INCH Threshold Port 1 VC1 Posted				A24h
INCH Threshold Port 1 VC1 Non-Posted				A28h
INCH Threshold Port 1 VC1 Completion				A2Ch
INCH Threshold Port 2 VC0 Posted				A30h
INCH Threshold Port 2 VC0 Non-Posted				A34h
INCH Threshold Port 2 VC0 Completion				A38h
INCH Threshold Port 2 VC1 Posted				A3Ch
INCH Threshold Port 2 VC1 Non-Posted				A40h
INCH Threshold Port 2 VC1 Completion				A44h
INCH Threshold Port 3 VC0 Posted				A48h
INCH Threshold Port 3 VC0 Non-Posted				A4Ch
INCH Threshold Port 3 VC0 Completion				A50h
Reserved				A54h – A5Ch
INCH Threshold Port 4 VC0 Posted				A60h
INCH Threshold Port 4 VC0 Non-Posted				A64h
INCH Threshold Port 4 VC0 Completion				A68h
Reserved				A6Ch – BCCh

Register 13-145. 9F0h INCH Port 4 Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**INCH FC Update Pending Timer** Update Pending Timer for Port 4. (Refer to Table 13-27.) For implementation, a value of 01h or 00h into the CSR results in waiting 255 symbol times. Values are automatically set at reset according to the **Device Control** register *Maximum Payload Size* (offset 70h[7:5]) and **Link Status** register *Negotiated Link Width* (offset 78h[25:20]) fields. Programmable by serial EEPROM if bit 8 (*INCH FC Update Pending Timer Uses Serial EEPROM Value* bit) is set.	RWS	Yes	00h
8	**INCH FC Update Pending Timer Uses Serial EEPROM Value** Update Pending Timer for Port 4, using the value programmed in the serial EEPROM. 0 = Use default values (refer to Table 13-27) 1 = Use serial EEPROM value	RWS	Yes	0
31:9	***Reserved***	RsvdP	No	0000_00h

Register 13-146. 9F4h INCH FC Update Pending Timer (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Port 0 Update Timer** Update Pending Timer for Port 0. (Refer to Table 13-27.) For implementation, a value of 01h or 00h into the CSR results in waiting 255 symbol times. Values are automatically set at reset according to the **Device Control** register *Maximum Payload Size* (offset 70h[7:5]) and **Link Status** register *Negotiated Link Width* (offset 78h[25:20]) fields. Programmable by serial EEPROM if the associated **INCH Mode** register *INCH FC Update Pending Timer Uses Serial EEPROM Value* bit (offset 9FCh[20]) is set.	RWS	Yes	00h
15:8	**Port 1 Update Timer** Update Pending Timer for Port 1. (Refer to Table 13-27.) For implementation, a value of 01h or 00h into the CSR results in waiting 255 symbol times. Values are automatically set at reset according to the **Device Control** register *Maximum Payload Size* (offset 70h[7:5]) and **Link Status** register *Negotiated Link Width* (offset 78h[25:20]) fields. Programmable by serial EEPROM if the associated **INCH Mode** register *INCH FC Update Pending Timer Uses Serial EEPROM Value* bit (offset 9FCh[21]) is set.	RWS	Yes	00h
23:16	**Port 2 Update Timer** Update Pending Timer for Port 2. (Refer to Table 13-27.) For implementation, a value of 01h or 00h into the CSR results in waiting 255 symbol times. Values are automatically set at reset according to the **Device Control** register *Maximum Payload Size* (offset 70h[7:5]) and **Link Status** register *Negotiated Link Width* (offset 78h[25:20]) fields. Programmable by serial EEPROM if the associated **INCH Mode** register *INCH FC Update Pending Timer Uses Serial EEPROM Value* bit (offset 9FCh[22]) is set.	RWS	Yes	00h
31:24	**Port 3 Update Timer** Update Pending Timer for Port 3. (Refer to Table 13-27.) For implementation, a value of 01h or 00h into the CSR results in waiting 255 symbol times. Values are automatically set at reset according to the **Device Control** register *Maximum Payload Size* (offset 70h[7:5]) and **Link Status** register *Negotiated Link Width* (offset 78h[25:20]) fields. Programmable by serial EEPROM if the associated **INCH Mode** register *INCH FC Update Pending Timer Uses Serial EEPROM Value* bit (offset 9FCh[23]) is set.	RWS	Yes	00h

Table 13-27. FC Update Pending Timer Guidelines

Maximum Packet Size	Link Width	Recommended Timer Count
128 bytes	x1	76h
	x2	40h
	x4	24h
256 bytes	x1	D0h
	x2	6Ch
	x4	3Bh

Register 13-147. 9FCh INCH Mode (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Maximum Mode Enable** ***Factory Test Only***	RO	No	FFh
15:8	***Reserved***	RsvdP	No	00h
19:16	***Factory Test Only***	RWS	Yes	0h
23:20	**INCH FC Update Pending Timer Uses Serial EEPROM Value** Update Pending Timer for Ports 0, 1, 2, or 3, using the value programmed in the serial EEPROM. 0 = Use default values (refer to Table 13-27) 1 = Use serial EEPROM value **Bit** **Controls** 20 Port 0 Pending Timer Source 21 Port 1 Pending Timer Source 22 Port 2 Pending Timer Source 23 Port 3 Pending Timer Source	RWS	Yes	0h
31:24	***Reserved***	RsvdP	No	00h

Register 13-148. A00h INCH Threshold Port 0 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 0 (VC0) Memory Write and Message transactions.				
8:0	**Payload** Payload = 00_1111b = 120d = 78h.	RWS	Yes	1878h
13:9	**Header** Header = 0_1100b = 12d = Ch.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-149. A04h INCH Threshold Port 0 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	**Payload** Payload = 0_0000_0111b = 7d = 7h.	RWS	Yes	0E07h
13:9	**Header** Header = 0_0111b = 7d = 7h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-150. A08h INCH Threshold Port 0 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
8:0	**Payload** Payload = 00_1101b = 104d = 68h.	RWS	Yes	1468h
13:9	**Header** Header = 0_1010b = 10d = Ah.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-151. A0Ch INCH Threshold Port 0 VC1 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 1 (VC1) Memory Write and Message transactions. Port 0 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 0 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 00_0010b = 16d = 10h.	RWS	Yes	0210h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-152. A10h INCH Threshold Port 0 VC1 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 1 (VC1) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions. Port 0 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 0 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 0_0000_0001b = 1d = 1h.	RWS	Yes	0201h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-153. A14h INCH Threshold Port 0 VC1 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 1 (VC1) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions. Port 0 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 0 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 00_0010b = 16d = 10h.	RWS	Yes	0210h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-154. A18h INCH Threshold Port 1 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 0 (VC0) Memory Write and Message transactions.				
8:0	**Payload** Payload = 00_1011b = 88d = 58h.	RWS	Yes	1258h
13:9	**Header** Header = 0_1001b = 9d = 9h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-155. A1Ch INCH Threshold Port 1 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	**Payload** Payload = 0_0000_0101b = 5d = 5h.	RWS	Yes	0A05h
13:9	**Header** Header = 0_0101b = 5d = 5h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-156. A20h INCH Threshold Port 1 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
8:0	**Payload** Payload = 00_1001b = 72d = 48h.	RWS	Yes	0E48h
13:9	**Header** Header = 0_0111b = 7d = 7h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-157. A24h INCH Threshold Port 1 VC1 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 1 (VC1) Memory Write and Message transactions. Port 1 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 1 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 00_0010b = 16d = 10h.	RWS	Yes	0210h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-158. A28h INCH Threshold Port 1 VC1 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 1 (VC1) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions. Port 1 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 1 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 0_0000_0001b = 1d = 1h.	RWS	Yes	0201h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-159. A2Ch INCH Threshold Port 1 VC1 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 1 (VC1) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions. Port 1 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 1 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 00_0010b = 16d = 10h.	RWS	Yes	0210h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-160. A30h INCH Threshold Port 2 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 0 (VC0) Memory Write and Message transactions.				
8:0	**Payload** Payload = 00_1011b = 88d = 58h.	RWS	Yes	1258h
13:9	**Header** Header = 0_1001b = 9d = 9h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-161. A34h INCH Threshold Port 2 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	**Payload** Payload = 0_0000_0101b = 5d = 5h.	RWS	Yes	0A05h
13:9	**Header** Header = 0_0101b = 5d = 5h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-162. A38h INCH Threshold Port 2 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
8:0	**Payload** Payload = 00_1001b = 72d = 48h.	RWS	Yes	0E48h
13:9	**Header** Header = 0_0111b = 7d = 7h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-163. A3Ch INCH Threshold Port 2 VC1 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 1 (VC1) Memory Write and Message transactions. Port 2 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 2 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 00_0010b = 16d = 10h.	RWS	Yes	0210h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-164. A40h INCH Threshold Port 2 VC1 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 1 (VC1) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions. Port 2 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 2 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 0_0000_0001b = 1d = 1h.	RWS	Yes	0201h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-165. A44h INCH Threshold Port 2 VC1 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 1 (VC1) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions. Port 2 does not advertise VC1 credits unless its VC1 is enabled, by setting the Port 2 **VC1 Resource Control** register *VC1 Enable* bit (offset 168h[31]=1).				
8:0	**Payload** Payload = 00_0010b = 16d = 10h.	RWS	Yes	0210h
13:9	**Header** Header = 0_0001b = 1d = 1h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-166. A48h INCH Threshold Port 3 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 0 (VC0) Memory Write and Message transactions.				
8:0	**Payload** Payload = 00_1000b = 64d = 40h.	RWS	Yes	0C40h
13:9	**Header** Header = 0_0110b = 6d = 6h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-167. A4Ch INCH Threshold Port 3 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	**Payload** Payload = 0_0000_0100b = 4d = 4h.	RWS	Yes	0804h
13:9	**Header** Header = 0_0100b = 4d = 4h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-168. A50h INCH Threshold Port 3 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
8:0	**Payload** Payload = 00_0111b = 56d = 38h.	RWS	Yes	0C38h
13:9	**Header** Header = 0_0110b = 6d = 6h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-169. A60h INCH Threshold Port 4 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Posted credits are used for Virtual Channel 0 (VC0) Memory Write and Message transactions.				
8:0	**Payload** Payload = 00_1000b = 64d = 40h.	RWS	Yes	0C40h
13:9	**Header** Header = 0_0110b = 6d = 6h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-170. A64h INCH Threshold Port 4 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Non-Posted credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	**Payload** Payload = 0_0000_0100b = 4d = 4h.	RWS	Yes	0804h
13:9	**Header** Header = 0_0100b = 4d = 4h.			
31:14	***Reserved***	RsvdP	No	0-0h

Register 13-171. A68h INCH Threshold Port 4 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Completion credits are used for Virtual Channel 0 (VC0) Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
8:0	**Payload** Payload = 00_0111b = 56d = 38h.	RWS	Yes	0C38h
13:9	**Header** Header = 0_0110b = 6d = 6h.			
31:14	***Reserved***	RsvdP	No	0-0h

13.13.13 Relaxed Ordering and Performance Counter Registers

Table 13-28. Device-Specific Relaxed Ordering and Performance Counter Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Reserved			***Factory Test Only***	BE0h
Reserved			Device-Specific Relaxed Ordering Mode Port 4	BE4h
One-Bit ECC Error Count		***Reserved***		BE8h
Factory Test Only				BECh – BF8h
Device-Specific Relaxed Ordering Mode				BFCh

Register 13-172. BE4h Device-Specific Relaxed Ordering Mode Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Enable PLX Relaxed Ordering on Port 4** When any of these bits are set, the corresponding Traffic Class for this port allows Device-Specific Relaxed Ordering on Port 4.	RW	Yes	00h
31:8	***Reserved***	RW	Yes	0000_00h

Register 13-173. BE8h One-Bit ECC Error Count (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	***Reserved***	RsvdP	No	0000h
23:16	**Ingress Packet RAM 1-Bit ECC Count** Record number of 1-bit Correctable errors that occurred in Ingress RAM.	RO	No	00h
31:24	**Ingress Pointer Linked List RAM 1-Bit ECC Count** Record number of 1-bit Correctable errors that occurred in the Ingress Pointer Linked List RAM.	RO	No	00h

Register 13-174. BFCh Device-Specific Relaxed Ordering Mode (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Enable PLX Relaxed Ordering on Port 0** When any of these bits are set, the corresponding Traffic Class for this port allows Device-Specific Relaxed Ordering on Port 0.	RW	Yes	00h
15:8	**Enable PLX Relaxed Ordering on Port 1** When any of these bits are set, the corresponding Traffic Class for this port allows Device-Specific Relaxed Ordering on Port 1.	RW	Yes	00h
23:16	**Enable PLX Relaxed Ordering on Port 2** When any of these bits are set, the corresponding Traffic Class for this port allows Device-Specific Relaxed Ordering on Port 2.	RW	Yes	00h
31:24	**Enable PLX Relaxed Ordering on Port 3** When any of these bits are set, the corresponding Traffic Class for this port allows Device-Specific Relaxed Ordering on Port 3.	RW	Yes	00h

13.13.14 Internal Credit Handler (ITCH) VC&T Threshold Registers

Note: *Previously scheduled packets arrive in their entirety, completely unaffected by the cutoff signal.*

The registers defined in Table 13-29 control internal traffic from the Link Interface to the Virtual Interface. The threshold (Packet Count) units are equivalent to 8 beats, where each beat can be up to 20 bytes. Therefore, a programmed value of 1 represents 160 bytes, 2 represents 320 bytes, and so forth. The entire TLP (header, payload, and ECRC, if any) is used to determine a total byte size, and the total byte size is divided by 20 and rounded up to the nearest integer to ascertain the number of beats. Every 8 beats counts as 1 threshold unit.

The Upper Packet Count is the high threshold. If more units than the programmed upper count are queued, no more packets of that type can be scheduled across the internal fabric.

The Lower Packet Count is the low threshold. After cutting off a VC&T due to the high threshold, that VC&T is turned On again after the count returns below the low threshold.

The upper and lower counts must be different, and the upper number must be at least two units larger than the lower number.

Table 13-29. Device-Specific Internal Credit Handler (ITCH) VC&T Threshold Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
ITCH VC&T Threshold 1				C00h
ITCH VC&T Threshold 2				C04h
ITCH VC&T Threshold 3				C08h

The **ITCH VC&T Threshold *x*** registers define internal credits between the Egress queue of one port and the Ingress queues of ingress ports. The **ITCH VC&T Threshold *x*** registers are programmed in the serial EEPROM with values to prevent filling of the Egress RAM with Posted packets, which might prevent Completion packets from making forward progress inside the PEX 8508. The programmed values apply to all ports within the station. While separate registers exist for Posted, Non-Posted, and Completion Packet Counts, programming only the registers for Posted Counts prevents clogging of the Egress RAM for any combination of packets.

Recommended register values for Threshold mode are based upon the number of ports enabled in the station (as configured in the **Port Configuration** register, offset 224h). A port is considered enabled, regardless of whether it is used. The number of enabled ports is usually the same as the **Physical Layer Status** register *Number of Ports Enumerated* field (offset 220h[22:20]), which the PEX 8508 automatically updates after a reset.

Table 13-30 lists the recommended VC0 Posted Lower and Upper Packet Count values, based upon the number of enabled ports and only VC0 enabled.

Table 13-30. ITCH VC&T VC0 Posted Packet Count, for One VC Enabled (VC0 Only; All Ports)

Number of Enabled Ports	Port Configuration Register Value (Port 0, Offset 224h[4:0])	Destination RAM Entries	VC0 Posted Lower Packet Count Value (C00h[15:8])	VC0 Posted Upper Packet Count Value (C00h[7:0])
2	2h	256-256	54h	8Ch
3	3h	250-131-131	1Ch	3Fh
4	0h, 4h, 5h, 6h	152-120-120-120	1Ch	38h
5	8h, 9h	120-100-100-96-96	1Ch	2Ah

Table 13-31 and Table 13-32 list the recommended VC0 and VC1 Posted Lower and Upper Packet Count values, respectively, based upon the number of enabled ports, with both VC0 and VC1 enabled.

Caution: ***It is not recommended to use two VCs with five ports enabled, due to inefficient RAM utilization in this configuration.***

Table 13-31. ITCH VC&T VC0 Posted Packet Count, for Two VCs Enabled (Only Ports 0, 1, and 2)

Number of Enabled Ports	Port Configuration Register Value (Port 0, Offset 224h[4:0])	Destination RAM Entries	VC0 Posted Lower Packet Count Value (C00h[15:8])	VC0 Posted Upper Packet Count Value (C00h[7:0])
2	2h	256-256	54h	8Ch
3	3h	250-131-131	1Ch	3Fh
4	0h, 4h, 5h, 6h	152-120-120-120	1Ch	38h
5	8h, 9h	120-100-100-96-96	1Ch	2Ah

Table 13-32. ITCH VC&T VC1 Posted Packet Count, for Two VCs Enabled (Only Ports 0, 1, and 2)

Number of Enabled Ports	Port Configuration Register Value (Port 0, Offset 224h[4:0])	Destination RAM Entries	VC1 Posted Lower Packet Count Value (C04h[31:24])	VC1 Posted Upper Packet Count Value (C04h[23:16])
2	2h	256-256	2Ah	46h
3	3h	250-131-131	14h	2Ah
4	0h, 4h, 5h, 6h	152-120-120-120	14h	2Ah
5	8h, 9h	120-100-100-96-96	12h	22h

Register 13-175. C00h ITCH VC&T Threshold 1 (Only Port 0)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
7:0	**VC0 Posted Upper Packet Count** VC0 Posted upper packet beat limit.	0, 1, 2, 3, 4	RW	Yes	FFh
15:8	**VC0 Posted Lower Packet Count** VC0 Posted lower packet beat limit.	0, 1, 2, 3, 4	RW	Yes	7Fh
23:16	**VC0 Non-Posted Upper Packet Count** VC0 Non-Posted upper packet beat limit.	0, 1, 2, 3, 4	RW	Yes	FFh
31:24	**VC0 Non-Posted Lower Packet Count** VC0 Non-Posted lower packet beat limit.	0, 1, 2, 3, 4	RW	Yes	7Fh

Register 13-176. C04h ITCH VC&T Threshold 2 (Only Port 0)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
7:0	**VC0 Completion Upper Packet Count** VC0 Completion upper packet beat limit.	0, 1, 2, 3, 4	RW	Yes	FFh
15:8	**VC0 Completion Lower Packet Count** VC0 Completion lower packet beat limit.	0, 1, 2, 3, 4	RW	Yes	7Fh
23:16	**VC1 Posted Upper Packet Count** VC1 Posted upper packet beat limit.	0, 1, 2	RW	Yes	FFh
31:24	**VC1 Posted Lower Packet Count** VC1 Posted lower packet beat limit.	0, 1, 2	RW	Yes	7Fh

Register 13-177. C08h ITCH VC&T Threshold 3 (Only Port 0)

Bit(s)	Description	Ports	Type	Serial EEPROM and I²C	Default
7:0	**VC1 Non-Posted Upper Packet Count** VC1 Non-Posted upper packet beat limit.	0, 1, 2	RW	Yes	FFh
15:8	**VC1 Non-Posted Lower Packet Count** VC1 Non-Posted lower packet beat limit.	0, 1, 2	RW	Yes	7Fh
23:16	**VC1 Completion Upper Packet Count** VC1 Completion upper packet beat limit.	0, 1, 2	RW	Yes	FFh
31:24	**VC1 Completion Lower Packet Count** VC1 Completion lower packet beat limit.	0, 1, 2	RW	Yes	7Fh

13.13.15 Port Virtual Channel Queue Status Registers

Table 13-33. Device-Specific Port Virtual Channel Queue Status Register Map (Only Port 0 and NT Port Link Interface)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Port 0 VC0 Posted and Non-Posted Queue Status	C0Ch
Port 0 VC0 Completion and VC1 Posted Queue Status	C10h
Port 0 VC1 Non-Posted and Completion Queue Status	C14h
Port 1 VC0 Posted and Non-Posted Queue Status	C18h
Port 1 VC0 Completion and VC1 Posted Queue Status	C1Ch
Port 1 VC1 Non-Posted and Completion Queue Status	C20h
Port 2 VC0 Posted and Non-Posted Queue Status	C24h
Port 2 VC0 Completion and VC1 Posted Queue Status	C28h
Port 2 VC1 Non-Posted and Completion Queue Status	C2Ch
Port 3 VC0 Posted and Non-Posted Queue Status	C30h
Port 3 VC0 Completion Queue Status	C34h
Reserved	C38h

Register 13-178. C0Ch Port 0 VC0 Posted and Non-Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 0 VC0 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 0 VC0 queue,	RO	Yes	000h
21:11	**Port 0 VC0 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 0 VC0 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 0 ITCH Disabled VC0 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 0 VC0 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 0 ITCH Disabled VC0 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 0 VC0 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-179. C10h Port 0 VC0 Completion and VC1 Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 0 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 0 VC0 queue.	RO	Yes	000h
21:11	**Port 0 VC1 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 0 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 0 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 0 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 0 ITCH Disabled VC1 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 0 VC1 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-180. C14h Port 0 VC1 Non-Posted and Completion Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 0 VC1 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 0 VC1 queue.	RO	Yes	000h
21:11	**Port 0 VC1 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 0 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 0 ITCH Disabled VC1 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 0 VC1 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 0 ITCH Disabled VC1 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 0 VC1 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-181. C18h Port 1 VC0 Posted and Non-Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 1 VC0 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 1 VC0 queue,	RO	Yes	000h
21:11	**Port 1 VC0 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 1 VC0 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 1 ITCH Disabled VC0 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC0 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC0 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC0 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-182. C1Ch Port 1 VC0 Completion and VC1 Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 1 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 1 VC0 queue.	RO	Yes	000h
21:11	**Port 1 VC1 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 1 VC1 queue.	RO	Yes	000h
23:22	***Factory Test Only***	RW	Yes	00b
27:24	**Port 1 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC1 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-183. C20h Port 1 VC1 Non-Posted and Completion Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 1 VC1 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 1 VC1 queue.	RO	Yes	000h
21:11	**Port 1 VC1 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 1 VC1 queue.	RO	Yes	000h
23:22	***Factory Test Only***	RW	Yes	00b
27:24	**Port 1 ITCH Disabled VC1 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC1 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-184. C24h Port 2 VC0 Posted and Non-Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 2 VC0 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 2 VC0 queue,	RO	Yes	000h
21:11	**Port 2 VC0 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 2 VC0 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 2 ITCH Disabled VC0 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC0 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 2 ITCH Disabled VC0 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC0 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-185. C28h Port 2 VC0 Completion and VC1 Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 2 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 2 VC0 queue.	RO	Yes	000h
21:11	**Port 2 VC1 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 2 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 2 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 2 ITCH Disabled VC1 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC1 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-186. C2Ch Port 2 VC1 Non-Posted and Completion Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 2 VC1 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 2 VC1 queue.	RO	Yes	000h
21:11	**Port 2 VC1 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 2 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 2 ITCH Disabled VC1 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC1 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 2 ITCH Disabled VC1 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC1 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-187. C30h Port 3 VC0 Posted and Non-Posted Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 3 VC0 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 3 VC0 queue,	RO	Yes	000h
21:11	**Port 3 VC0 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 3 VC0 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 3 ITCH Disabled VC0 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 3 VC0 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 3 ITCH Disabled VC0 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 3 VC0 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 13-188. C34h Port 3 VC0 Completion Queue Status (Only Port 0 and NT Port Link Interface)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 3 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 3 VC0 queue.	RO	Yes	000h
23:11	***Reserved***	RsvdP	No	0000h
27:24	**Port 3 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 3 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	***Reserved***	RsvdP	No	0h

13.14 Advanced Error Reporting Capability Registers

Table 13-34. Advanced Error Reporting Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (138h)	Capability Version (1h) / PCI Express Extended Capability ID (0001h)	FB4h
Uncorrectable Error Status[a]		FB8h
Uncorrectable Error Mask[a]		FBCh
Uncorrectable Error Severity[a]		FC0h
Correctable Error Status		FC4h
Correctable Error Mask		FC8h
Advanced Error Capabilities and Control		FCCh
Header Log 0		FD0h
Header Log 1		FD4h
Header Log 2		FD8h
Header Log 3		FDCh
Reserved		FE0h – FFCh

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 13-189. FB4h Advanced Error Reporting Enhanced Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**PCI Express Extended Capability ID**	RO	Yes	0001h
19:16	**Capability Version**	RO	Yes	1h
31:20	**Next Capability Offset** Set to 138h, which addresses the PEX 8508 **Power Budgeting Enhanced Capability Header** registers.	RO	Yes	138h

Register 13-190. FB8h Uncorrectable Error Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
5	**Surprise Down Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
11:6	***Reserved***	RsvdP	No	0000_00b
12	**Poisoned TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
13	**Flow Control Protocol Error Status**	RsvdP[a]	Yes	0
14	**Completion Timeout Status**	RsvdP[a]	No	0
15	**Completer Abort Status**	RsvdP[a]	Yes	0
16	**Unexpected Completion Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
17	**Receiver Overflow Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
18	**Malformed TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
19	**ECRC Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
20	**Unsupported Request Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 13-191. FBCh Uncorrectable Error Mask (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
5	**Surprise Down Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
11:6	***Reserved***	RsvdP	No	0000_00b
12	**Poisoned TLP Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
13	**Flow Control Protocol Error Mask**	RsvdP[a]	Yes	0
14	**Completion Timeout Mask**	RsvdP[a]	No	0
15	**Completer Abort Mask**	RsvdP[a]	Yes	0
16	**Unexpected Completion Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
17	**Receiver Overflow Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
18	**Malformed TLP Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
19	**ECRC Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
20	**Unsupported Request Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 13-192. FC0h Uncorrectable Error Severity (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
5	**Surprise Down Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
11:6	***Reserved***	RsvdP	No	0-0h
12	**Poisoned TLP Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
13	**Flow Control Protocol Error Severity**	RsvdP[a]	Yes	1
14	**Completion Timeout Severity**	RsvdP[a]	No	0
15	**Completer Abort Severity**	RsvdP[a]	Yes	0
16	**Unexpected Completion Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
17	**Receiver Overflow Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
18	**Malformed TLP Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
19	**ECRC Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
20	**Unsupported Request Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 13-193. FC4h Correctable Error Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Receiver Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
5:1	***Reserved***	RsvdP	No	0-0h
6	**Bad TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
7	**Bad DLLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
8	**Replay Number Rollover Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
11:9	***Reserved***	RsvdP	No	000b
12	**Replay Timer Timeout Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
31:13	***Reserved***	RsvdP	No	0-0h

Register 13-194. FC8h Correctable Error Mask (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Receiver Error Mask** 0 = Error reporting not masked 1 = Error reporting masked	RWS	Yes	0
5:1	***Reserved***	RsvdP	No	0-0h
6	**Bad TLP Mask** 0 = Error reporting not masked 1 = Error reporting masked	RWS	Yes	0
7	**Bad DLLP Mask** 0 = Error reporting not masked 1 = Error reporting masked	RWS	Yes	0
8	**Replay Number Rollover Mask** 0 = Error reporting not masked 1 = Error reporting masked	RWS	Yes	0
11:9	***Reserved***	RsvdP	No	000b
12	**Replay Timer Timeout Mask** 0 = Error reporting not masked 1 = Error reporting masked	RWS	Yes	0
31:13	***Reserved***	RsvdP	No	0-0h

Register 13-195. FCCh Advanced Error Capabilities and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
4:0	**First Error Pointer** Identifies the bit position of the first error reported in the **Uncorrectable Error Status** register.	ROS	Yes	1_1111b
5	**ECRC Generation Capable** 0 = ECRC generation is not supported 1 = ECRC generation is supported, but must be enabled	RO	Yes	1
6	**ECRC Generation Enable** 0 = ECRC generation is disabled 1 = ECRC generation is enabled	RWS	Yes	0
7	**ECRC Checking Capable** 0 = ECRC checking is not supported 1 = ECRC checking is supported, but must be enabled	RO	Yes	1
8	**ECRC Checking Enable** 0 = ECRC checking is disabled 1 = ECRC checking is enabled	RWS	Yes	0
31:9	***Reserved***	RsvdP	No	0-0h

Register 13-196. FD0h Header Log 0 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**TLP Header 0** First DWord header. TLP header associated with error.	ROS	Yes	0-0h

Register 13-197. FD4h Header Log 1 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**TLP Header 1** Second DWord header. TLP header associated with error.	ROS	Yes	0-0h

Register 13-198. FD8h Header Log 2 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**TLP Header 2** Third DWord header. TLP header associated with error.	ROS	Yes	0-0h

Register 13-199. FDCh Header Log 3 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**TLP Header 3** Fourth DWord header. TLP header associated with error.	ROS	Yes	0-0h

Chapter 14 Non-Transparent Bridging (NTB)

14.1 Introduction

A Non-Transparent (NT) bridge is functionally similar to a Transparent bridge in that both provide a path between two independent PCI Express interfaces. The key difference is that when an NT bridge is used, devices on the downstream side (relative to the system Host) of the bridge are not visible from the upstream side. This allows an intelligent controller on the downstream side to manage devices there, making them appear as a single controller to the system Host.

The path between the two interfaces allows the devices on the downstream side to directly transfer data to the upstream side of the bus, without directly involving the intelligent controller in the data move. Thus, transactions are forwarded across the interface unrestricted, just as in a PCI-to-PCI bridge, but the resources responsible are hidden from the Host, which detects a single device. An NT bridge can also be used to link a secondary Host with the hierarchy of a primary Host. It provides isolation while allowing communication between the two systems.

The PEX 8508 includes **Doorbell** registers to send interrupts from one side of the bridge to the other, and **Scratchpad** registers accessible from both sides for inter-processor communication. Upon primary Host failure, the NT bridge resources allow access by the secondary Host to re-configure the system so that it can take over as Host.

Any one of the PEX 8508 downstream ports can be configured as an NT Port, while the other downstream ports remain Transparent. The primary side of the NT Port is referred to as the *NT Port Link Interface*, and connects to external device pins. The secondary side of the NT Port is referred to as the *NT Port Virtual Interface*, and connects to the internal virtual PCI Express interface.

After power-up, the PEX 8508 NT Type 1 ports, including the NT Port Virtual Interface, are enumerated by the Local Host connected to the PEX 8508 upstream port. The Local Host enables/resizes the Base Address registers (BARs) by programming the NT Port Link Interface BAR Setup/Limit registers, before the System Host assigns resources for these BARs. This behavior is changed with serial EEPROM initialization.

After the Local Host finishes its enumeration, it enables the NT Port Link Interface to be enumerated by the System Host connected to the NT Port Link Interface. The NT Port Link Interface Retries the System Host Configuration transaction until the Local Host enables the NT Port Link Interface to process the System Host Configuration transaction. The **Debug Control** register *Virtual Interface Access Enable* bit (offset 1DCh[28]) enables access to the Virtual Interface Configuration registers. The *Link Interface Access Enable* bit (offset 1DCh[29]) enables access to the Link Interface Configuration registers. These bits do not affect normal Memory, Memory-Mapped CSRs, nor I/O-Mapped CSR transactions.

The following are key elements of PEX 8508 NTB:

- Device Type Identification
- Non-Transparent Port (NT Port) Reset
- Scratchpad Registers
- Doorbell Registers
- BAR Setup Registers
- Address Translation

14.1.1 Device Type Identification

Devices identify themselves by way of the Conventional PCI CSR header **Class Code** register. A Transparent PCI-to-PCI bridge identifies itself as a Class Code 060400h. The NT bridge identifies itself as "other," 068000h, with a Type 0 Header. This identification is consistent with the use of other Non-Transparent bridges available in the industry.

The **PCI Express Capabilities** register includes a *Device/Port Type* field (offset 68h[23:20]). In this register, a Transparent bridge/switch port identifies itself as an upstream or downstream port, while the NT bridge/switch NT Port identifies itself as a PCI Express endpoint.

14.1.2 Non-Transparent Port (NT Port) Reset

The section discusses Non-Transparent mode exceptions and enhancements to Transparent mode PCI Express (standard) reset behavior.

14.1.2.1 Fundamental Reset (PEX_PERST#)

PEX_PERST# resets all PEX 8508 states, including NT Port states. All Sticky bits and Configuration registers in Virtual and Link spaces are initialized to default values by this reset.

14.1.2.2 Intelligent Adapter Mode NT Port Reset

In Intelligent Adapter mode, when a Hot Reset is received by the Transparent-side upstream port, the PEX 8508 propagates the reset to all Transparent downstream ports to reset them, resets the internal fabric and NT Port Virtual Interface states. There is no reset propagation to the NT Port and its Link side remains intact. The PEX 8508 supports an option that allows these Hot Reset conditions at its Transparent upstream port to be masked (disabled), by setting the **Debug Control** register *Upstream Port Hot Reset and Link Down Reset Propagation Disable* bit (Port 0, offset 1DCh[20]).

When the NT Port Link Interface receives a Hot Reset, the NT Port Link Interface registers are reset. This reset type does not reset the Transparent ports nor NT Port Virtual Interface. Instead, when the NT Port Link Interface receives a Hot Reset (or enters the DL_Down state), the PEX_NT_RESET# signal is asserted Low for 1 µs. The system can use this signal to trigger a reset of the entire Local subsystem (Transparent side).

When software writes to the PEX 8508 Transparent-side upstream port **Bridge Control** register *Secondary Bus Reset* bit (offset 3Ch[22]), the resulting secondary bus reset is (as above) propagated to all PEX 8508 Transparent downstream ports, and the port states and NT Port Virtual Interface states are reset.

14.1.2.3 Dual-Host Mode NT Port Reset

Dual-Host mode reset behavior is the same as in Intelligent Adapter mode, with the following exception – in Dual-Host mode, a Hot Reset received from the Active Host as detected at the PEX 8508 Transparent upstream port (or DL_Down state) does ***not*** reset the Transparent ports, the internal switch-fabric nor the NT Port Virtual Interface. The queues' internal operation and downstream ports remain intact, allowing the Backup Host to take over (as described in Section 4.4.2, "Dual-Host Mode").

There is no reset propagation onto the NT Port.

14.1.2.4 Reset Propagation

Reset propagation, during a Hot Reset or by way of the **Bridge Control** register *Secondary Bus Reset* mechanism is limited to Transparent downstream ports. In the NT bridge, this reset cannot be propagated across the bridge (across the NT Port). (Refer to Chapter 5, "Reset and Initialization," for details regarding PEX 8508 Transparent mode reset behavior.)

14.1.3 Scratchpad Registers

Scratchpad registers are readable and writable from both sides of the NT bridge, providing a generic means for inter-host communication. A block of eight registers are provided, accessible in Memory or I/O space from the NT Port Virtual and Link Interfaces. These registers pass Control and Status information between Virtual and Link Interface devices or they can be generic RW registers. Reading or writing Scratchpad registers does not cause interrupts to assert – **Doorbell** registers are used for that purpose.

14.1.4 Doorbell Registers

Doorbell registers are used to transmit interrupts from one side of the NT bridge to the other. This section describes a typical set of Doorbell Control registers.

A 16-bit software-controlled Interrupt Request register and associated 16-bit Mask register are implemented for the NT Port Virtual and Link Interfaces. These registers can be accessed from the NT Port Virtual and Link Interfaces, in Memory or I/O space. The Doorbell mechanisms consist of the following registers:

- **Set Virtual Interface IRQ**
- **Clear Virtual Interface IRQ**
- **Set Virtual Interface IRQ Mask**
- **Clear Virtual Interface IRQ Mask**
- **Set Link Interface IRQ**
- **Clear Link Interface IRQ**
- **Set Link Interface IRQ Mask**
- **Clear Link Interface IRQ Mask**

An interrupt is asserted on the NT Port Virtual Interface when one or more of the **IRQ Set** register bits are set to 1 and their corresponding Mask bits are cleared to 0. An interrupt is de-asserted on the NT Port Virtual Interface when one or more of the **IRQ Clear** register bits are set to 1 and their corresponding Mask bits are cleared to 0. The NT Port Link Interface operates in the same manner.

In a PCI Express switch, interrupt state transitions (from setting to clearing, or vice versa) result in packets being transmitted upstream on the appropriate side of the bridge when INT*x* is enabled. Standard PCI Express Capability structures allow these interrupts to be configured as MSI or INT*x*. When MSI is enabled, packets are transmitted only when interrupts transition from Clear IRQ to Set IRQ.

The Set IRQ and Clear IRQ are two interfaces to the same register. To the user, the Set IRQ register must be accessed to set the bit while the Clear IRQ register is used to clear the bits. The status of this register can be read from either interface.

The PEX 8508 Virtual interrupts are also disabled/removed when the link to the other device is down.

14.1.5 BAR Setup Registers

All NT Port Virtual and Link Interface BARs include programmable window sizes, with the exception of **BAR0** and **BAR1** (on both interfaces), which provide Memory and/or I/O-Mapped access to the CSRs. The **BAR Setup** registers are used to program the window size of each BAR. A detailed description of the NT Port Virtual and Link Interface BARs follows.

14.1.5.1 NT Port Virtual Interface BARs

Table 14-1 defines the NT Port Virtual Interface BARs.

Table 14-1. NT Port Virtual Interface BARs

BAR	Description
BAR0	***Reserved*** All PEX 8508 Port Configuration registers are mapped into Memory space using Transparent upstream port Type 1 space **BAR0** and **BAR1**. The Local Host, connected to the Transparent ports, can use the Transparent upstream port BARs to access the PEX 8508 Port Configuration registers.
BAR1	**BAR1** is ***reserved***.
BAR2	Configured by the **NT Port Virtual Interface Memory BAR2 Setup** register. **BAR2** is always a 32-bit BAR and uses Direct Address Translation.
BAR3	Configured by the **NT Port Virtual Interface Memory BAR3 Setup** register. **BAR3** is always a 32-bit BAR and uses Lookup Table-Based Address Translation.
BAR4	Configured by the **NT Port Virtual Interface Memory BAR4/5 Setup** register. **BAR4** can be implemented as a 32-bit BAR or lower half of a 64-bit BAR by combining it with **BAR5**. **BAR4** uses Direct Address Translation.
BAR5	Enabled only when **BAR4** is configured as a 64-bit BAR. Holds the upper 32-bit Base address of the 64-bit Memory Address range. The **NT Port Virtual Interface Memory BAR4/5 Setup** register defines the size. **BAR5** is not implemented as a 32-bit only BAR. **BAR5** uses Direct Address Translation.

14.1.5.2 NT Port Link Interface BARs

Table 14-2 defines the NT Port Link Interface BARs.

Table 14-2. NT Port Link Interface BARs

BAR	Description
BAR0	Maps all PEX 8508 Port Configuration registers into System Host Memory space. **BAR0** is always enabled.
BAR1	Maps only NT Port Virtual Interface and Link Interface Configuration registers into System Host I/O space. BARs can be disabled (enabled by default) by the **NT Port Link Interface BAR0/BAR1 Setup** register.
BAR2	Configured by the **NT Port Link Interface Memory BAR2/3 Setup** register. **BAR2** can be implemented as a 32-bit BAR or lower half of a 64-bit BAR by combining it with **BAR3**. **BAR2** uses Direct Address Translation.
BAR3	Enabled only when **BAR2** is configured as a 64-bit BAR. Holds the upper 32-bit Base address of the 64-bit Memory Address range. The **NT Port Link Interface Memory BAR3 Setup** register defines the size. **BAR3** is not implemented as a 32-bit only BAR. **BAR3** uses Direct Address Translation.
BAR4	Configured by the **NT Port Link Interface Memory BAR4/5 Setup** register. **BAR4** can be implemented as a 32-bit BAR or lower half of a 64-bit BAR by combining it with **BAR5**. **BAR4** uses Direct Address Translation.
BAR5	Enabled only when **BAR4** is configured as a 64-bit BAR. Holds the upper 32-bit Base address of the 64-bit Memory Address range. The **NT Port Link Interface Memory BAR5 Setup** register defines the size. **BAR5** is not implemented as a 32-bit only BAR. **BAR5** uses Direct Address Translation.

14.1.5.3 BAR Limit Registers

The Base Address register (BAR) Address Range size is a power of two (BARs can be assigned only Memory resources in "power of two" granularity).

A Limit register is used to reduce the Address Range size. When a Limit register is implemented, the Address range extends from the Base register to the Limit register, instead of to the Base register plus an offset of $2^N - 1$, where N is the number of bits in the offset (as defined by the Setup register).

The NT Port forwards the transaction when the transaction address falls within the range of the BAR lower address limit value to the BAR Limit register value. If the transaction address is outside this range, the NT Port reports an "Unsupported Request" and discards the transaction.

Figure 14-1. Limit Register Application

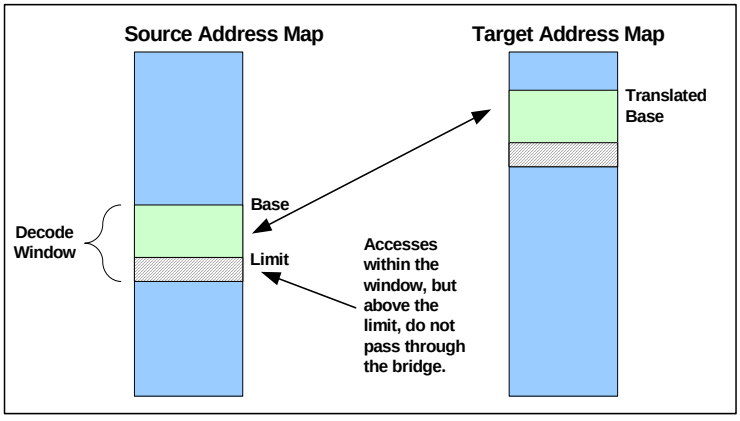

14.1.6 Address Translation

The Transparent bridge uses Base and Limit registers in I/O space, Non-Prefetchable Memory space, and Prefetchable Memory space to map transactions in the downstream direction across the bridge. All downstream devices must be mapped in contiguous address regions, such that a single Address range in each space is sufficient. Upstream mapping is done by way of inverse decode, relative to the same registers. A Transparent bridge does not translate the addresses of forwarded transactions/packets.

Address domain is unique per Host. If a transaction originates in one Host domain and targets a device in another Host domain, it must travel through the NT Port. If the NT Port does not process address translation, the transaction travels to a non-targeted device on a second Host domain, or is rejected by the NT Port upstream bridge. Transactions crossing the address domain must be address-translated by the NT Port before transmitting the transaction to the target Host domain.

The NT bridge uses the Conventional PCI set of BARs in its Type 0 CSR header to define Address ranges into the Memory space on the other side of the bridge. BARs define resource Address ranges that allow transaction forwarding to the opposite (other side) interface.

There are two sets of BARs – one each for the Virtual and Link Interfaces. BARs contain a Setup and Address Translation register:

- **BAR Setup** registers enable/disable the BAR and define the Address Range size and type. Certain BARs contain a Limit register, which is used to restrict its Address Range size to less than a power of two. BAR Setup registers must be programmed prior to allowing configuration software to assign a resource for these BARs.
- **BAR Address Translation** registers must be programmed before generating traffic across the NT Port. This programming is typically done by information downloaded from the serial EEPROM, by software, or through the I^2C interface.

The PEX 8508 NT Port Virtual Interface supports two types of address translation:

- Direct Address Translation
- Lookup Table-Based Address Translation

The PEX 8508 NT Port Link Interface supports only Direct Address Translation.

14.1.6.1 Direct Address Translation

The BAR Setup registers define a mask that splits the address into an upper *base* field and a lower *offset* field. Translation then consists of replacing, under the maskable portion of the Setup register, the Address Base bits with the corresponding Address Translation register bits. Figure 14-2 illustrates Direct Address Translation.

The device(s) on the originating Host domain can communicate to a single device or multiple devices mapped to consecutive Memory Address space on the Target Host domain, by using the Direct Address Translation mechanism. Figure 14-3 illustrates the entire address map, claimed by the NT Port, mapped into the single target device. Figure 14-4 illustrates the entire address map claimed by the NT Port, mapped into multiple target devices. Multiple devices must be in contiguous Memory ranges.

Figure 14-2. Direct Address Translation

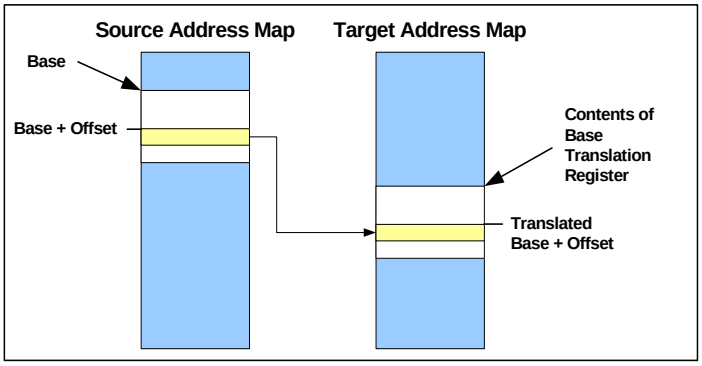

Figure 14-3. NT Port Mapped into Single Target Device

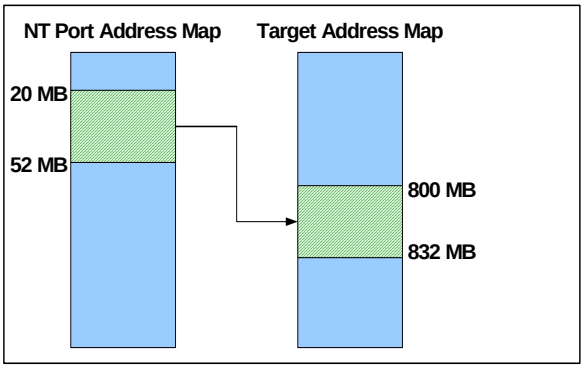

Figure 14-4. NT Port Mapped into Multiple Target Devices

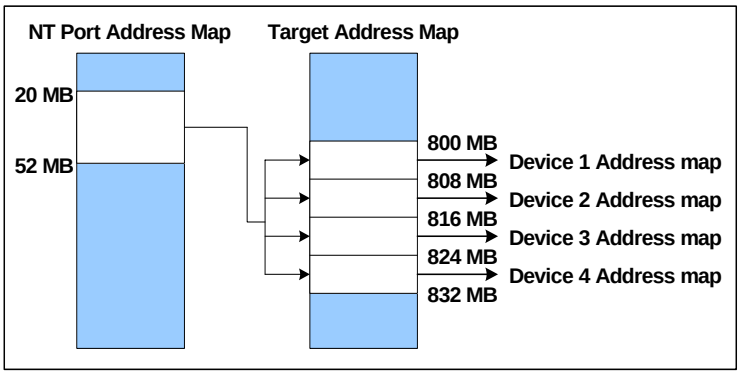

Address Translation Example

Assume the following:

1. NT Port Virtual Interface **BAR2** claims 128-KB Memory space.
2. Configuration software assigns the 5F00_0000h address value to NT Port Virtual Interface **BAR2** and it is within the Transparent upstream port Memory window.
3. Device driver software programs the **BAR2** Address Translation register to 2754_0000h.

The PEX 8508 receives a transaction to the NT Port Virtual Interface with address 5F00_0080h. The received transaction address is attaining the NT Port Virtual Interface **BAR2**. The PEX 8508 claims the transaction and executes the address translation described in Figure 14-5.

Figure 14-5. Address Translation Example

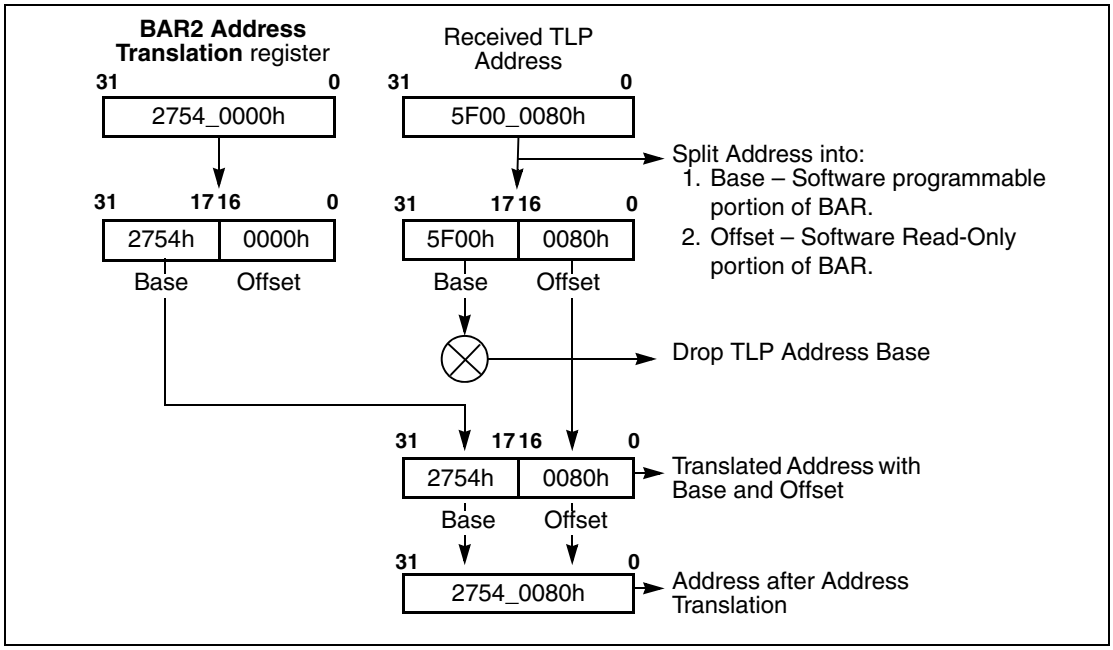

Note: *Nibble boundary-aligned hex address, bar_addr_xlation_reg[31:16], is 0010_0111_0101_0100b (2754h).*

14.1.6.2 Lookup Table-Based Address Translation

On the NT Port Virtual Interface, **BAR3** uses a special Lookup Table-Based Address translation for transactions that fall within its window.

The **NT Port Virtual Interface BAR3 Setup** register defines the lookup table (LUT) entry/page size. The **NT Port Virtual Interface BAR3** size is determined by multiplying the page size by 64.

This **BAR3 Setup** register defines a mask that splits the transaction address into upper and lower/offset fields. The upper field is further divided into two portions – the upper portion is referred to as the *Base address* and lower portion as the *Lookup Table Index*. The index field location of the received TLP address varies, based upon the page size selection. Table 14-3 defines the LUT entry/page size, corresponding BAR size, and bit position of individual fields in the received transaction address.

Figure 14-6 describes the Lookup Table-Based Address Translation scheme. The received transaction address is divided into three parts, based upon the **BAR Setup** register. The *Base Address* and *LUT Index* fields are compared against the BAR. If the transaction address attains the BAR, the PEX 8508 uses the LUT Index to select the LUT Entry. The PEX 8508 replaces the Base address and LUT Index with the selected Lookup Table Entry value, if the entry is valid. Otherwise, the NT Port Virtual Interface returns an *Unsupported Request* (UR) error condition. The PEX 8508 passes the received transaction address offset into translated address offset without modification.

Applications can use the Lookup Table-Based Address translation when the target device's Address range is scattered over 32-bit Memory space.

The NT Port Lookup table descriptions are discussed in Section 17.13.2, "NT Port Virtual Interface Lookup Table-Based Address Translation Registers."

Table 14-3. Received Transaction-Address Breakdown

Page Size (Bytes)	Window Size (Bytes)	Base Address (Bits)	LUT Index (Bits)	Offset (Bits)
4K	256K	[31:18]	[17:12]	[11:0]
8K	512K	[31:19]	[18:13]	[12:0]
16K	1M	[31:20]	[19:14]	[13:0]
32K	2M	[31:21]	[20:15]	[14:0]
64K	4M	[31:22]	[21:16]	[15:0]
128K	8M	[31:23]	[22:17]	[16:0]
256K	16M	[31:24]	[23:18]	[17:0]
512K	32M	[31:25]	[24:19]	[18:0]
1M	64M	[31:26]	[25:20]	[19:0]
2M	128M	[31:27]	[26:21]	[20:0]
4M	256M	[31:28]	[27:22]	[21:0]
8M	12M	[31:29]	[28:23]	[22:0]
16M	1G	[31:30]	[29:24]	[23:0]
32M	2G	31	[30:25]	[24:0]

Figure 14-6. Lookup Table-Based Address Translation

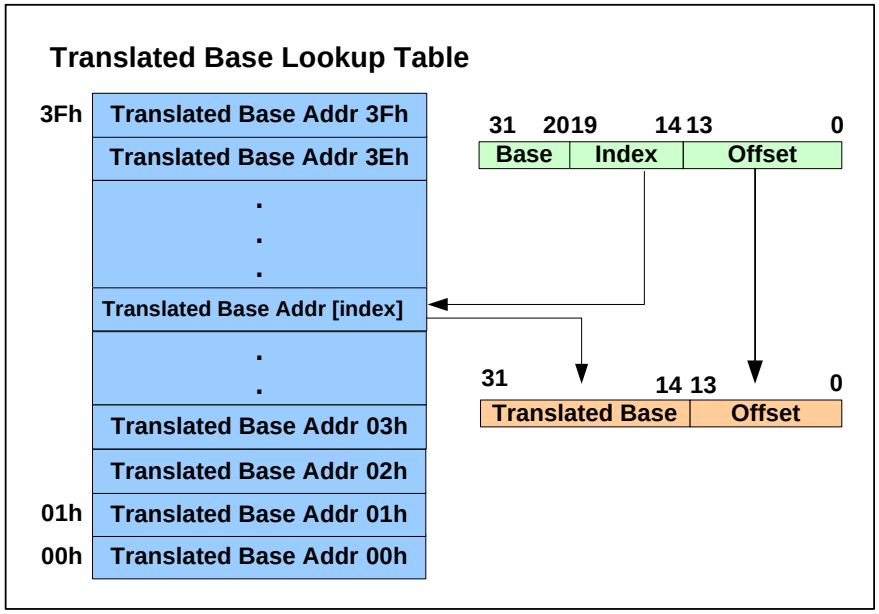

14.2 Requester ID Translation

Configuration, Message, and Completion transactions are ID-routed instead of address-routed. Of these, the NT Port forwards only the Completion transaction between the two Host domains. PCI Express switches and bridges use the Requester ID [defined in Completion Transaction Layer Packet (TLP) header] to route these packets.

The Requester ID consists of the following:

- Requester's PCI Bus Number
- Device Number
- Function Number

The Completer ID consists of the following:

- Completer's PCI Bus Number
- Device Number
- Function Number

Note: *The PCI Bus Number is unique for each Host domain.*

Figure 14-7 illustrates the Memory Request TLP header format. Figure 14-8 illustrates the Completion TLP Header format.

Figure 14-7. Memory Request TLP Header Format

	Byte 0								Byte 1								Byte 2								Byte 3							
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Bytes 0-3	R	Fmtx1		Type					R	TC			R				TD	EP	Attr		R				Length							
Bytes 4-7	Requester ID																Tag								Last DW BE				1st DW BE			
Bytes 8-11	Address[63:32]																															
Bytes 12-15	Address[31:0]																								R							

Figure 14-8. Completion TLP Header Format

	Byte 0								Byte 1								Byte 2								Byte 3							
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Bytes 0-3	R	Fmt		Type					R	TC			R				TD	EP	Attr		R				Length							
Bytes 4-7	Completer ID																Completer			BCM	Byte Count											
Bytes 8-11	Requester ID																Tag								R	Lower Address						

14.2.1 Transaction Sequence

To implement a transaction sequence:

1. Requester inserts ID information into the Memory Read TLP that it generates on the initiating Host domain.
2. Switches and bridges between the transaction initiator and PEX 8508 NT Port route this Memory Read TLP based upon the address.
3. NT Port replaces the Memory Read TLP Requester ID with its ID, and conducts the address translation before it forwards this Requester ID-translated TLP to the Target Host domain, because the NT Port is the transaction initiator in the target Host domain.
4. Switches and bridges between the PEX 8508 NT Port and target device route this Memory Read TLP, based upon the address.
5. When the target device generates the Completion TLP, it copies the Memory Read TLP Requester ID into the corresponding Completion *TLP Requester* ID field and inserts its ID into the *TLP Completer ID* field.
6. Switches and bridges between the target device and PEX 8508 NT Port route the Completion TLP, based upon Requester ID information.
7. NT Port restores the original Requester ID value from the Configuration register and implements another Requester ID and Completer ID translation for the Completion TLP before it forwards the Completion TLP to the Requester Host domain.
8. Switches and bridges between the PEX 8508 NT Port and Requester route the Completion TLP, based upon the Requester ID.
9. Requester accepts the Completion TLP and processes it.

14.2.2 Transaction Originating in Local Host Domain

The translation of outgoing requests from the NT Port Virtual Interface to the NT Port Link Interface uses an 8-entry LUT, as discussed in Section 17.13.5, "NT Port Virtual Interface Send Lookup Table Entry Registers." Each LUT entry supports all outgoing requests and any number of outstanding requests made by a single device or function. If a device uses Phantom Function Numbers to increase the maximum number of outstanding transactions, each phantom function consumes a LUT entry. The LUT must be configured, by a serial EEPROM or local firmware, before it is possible to transmit requests to the system domain, which provides a measure of security and protection.

When a Memory request arrives at the NT Port Virtual Interface, the packet Requester ID is associated with this LUT. If it attains one of the enabled LUT entries, the corresponding entry address (TxIndex) is inserted into the *Function Number* field of the packet's Requester ID. Conversely, if it does not attain one of the enabled LUT entries, an Unsupported Request Completion is returned.

At the same time, the contents of the NT Port Link Interface Bus Number and Device Number Capture registers (the values used during the last CSR write to the port) are copied into the packet Requester ID's *Bus Number* and *Device Number* fields.

A Completion, with translated Requester ID, returned from the system domain to the PEX 8508, is recognized when its Requester ID Bus and Device Numbers match the NT Port Link Interface captured Bus and Device Numbers. (Refer to Figure 14-9.)

When the original Requester ID is restored, the following occurs:

1. TxIndex is retrieved from the *Function Number* field of the Completion TLP Requester ID.
2. TxIndex is used to look up the same 8-entry LUT, to restore the original Requester ID.
3. If the selected entry is valid, the restored Requester ID is placed into the Completion *TLP Requester* field; otherwise, an Unexpected Completion is returned.
4. Completion *TLP Completer ID* field is replaced by the NT Port Virtual Interface captured Bus, Device, and Function Numbers.
5. Translated Completion TLP is forwarded to the original Requester in the local domain.

Figure 14-9. Requester ID Translation for Request Originating in Local Domain

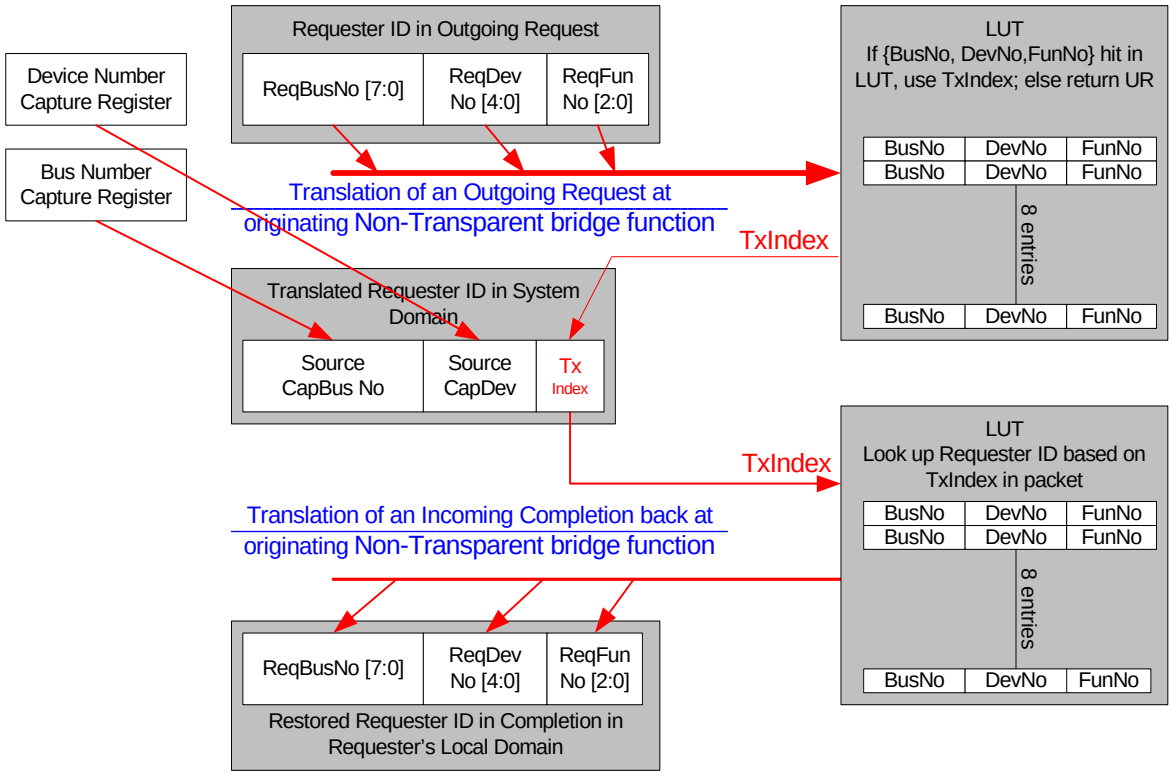

14.2.3 System Host Domain Transaction Originating

Transactions originating in System Host domain use a second LUT, with 32 entries, as illustrated in Figure 14-10. This data structure supports up to 32 devices (elsewhere in the system domain) that are transmitting requests through the associated NT Port. Because the Function Number is not used in the LUT association, a separate LUT entry is not required for each requesting or phantom function device. The LUT must be configured before transmitting requests through the NT Port. This Requester registration process, which cannot be accomplished by a peer, is an effective security and protection mechanism.

When a request is received from the system domain and routed to the NT Port, its Requester ID is again translated – Bus and Device Numbers, but not Function Number. The received Memory request TLP Requester ID is associated with this LUT, and the address (RxIndex) of the corresponding matching entry is substituted into the *Device Number* field of the Memory request's TLP *Requester ID* field.

If no match is found, or the matched entry is not enabled, the request receives a UR response.

If a match is found, and matched entry is enabled, the PEX 8508 internal virtual PCI Bus Number is copied into the packet Requester ID's *Bus Number* field. The translated Memory request TLP is address-translated and forwarded into the Local domain.

The PEX 8508 internal virtual PCI Bus Number is sufficient to route the Completion from the completer back to the NT Port in the completer's domain, because the NT Port is the only possible Requester on the switch internal virtual PCI Bus. Elsewhere in the PCI Express hierarchy, the Bus Number is sufficient to route the Completion back into the switch containing the NT Port.

The inverse translation occurs when a Completion passes through the NT bridge from the local domain to the system domain. The RxIndex is retrieved from the *Device Number* field of the received Completion *TLP Requester ID* header field, and used to look up the 32-entry LUT. The Completion TLP *Requester ID*, *Bus Number*, and *Device Number* fields are replaced by the decoded LUT-entry Bus Number and Device Number values if the entry is valid; otherwise, an Unexpected Completion is returned.

The Completion TLP Completer ID is replaced by the NT Port Link Interface captured Bus Number, captured Device Number and Function Number values before forwarding the Completion TLP to the system domain.

Figure 14-10. Requester ID Translation for Request Originating in System Domain

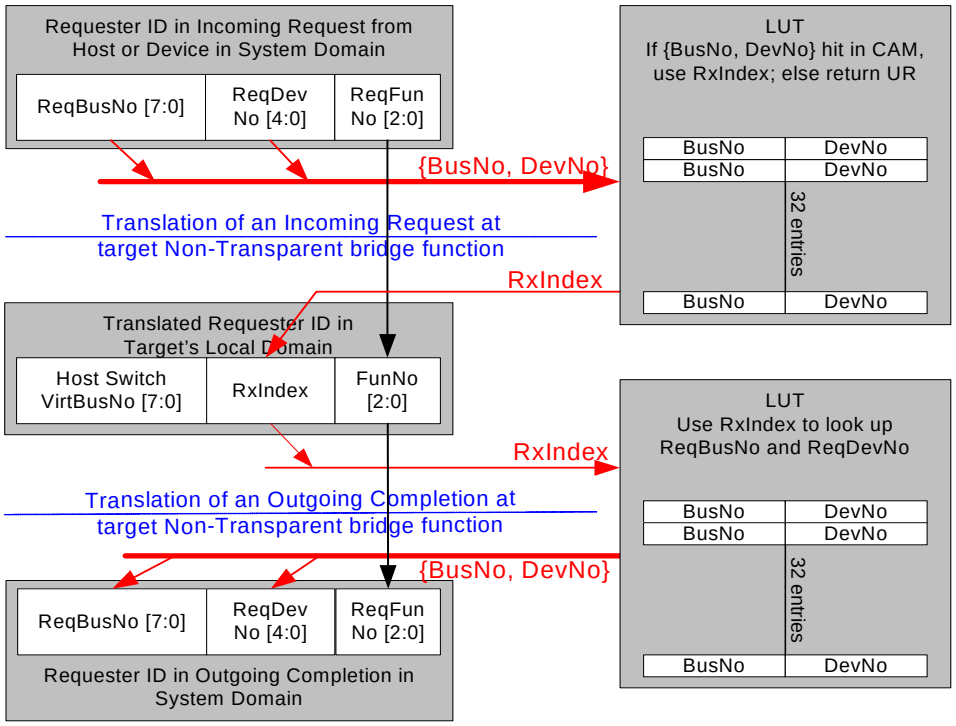

14.3 NT Port Power Management Handling

14.3.1 Active-State Power Management (ASPM)

When NT mode is enabled (Intelligent Adapter or Dual-Host mode), the PEX 8508 does not enter into active Power Management states L0s and L1 on any of its ports, although the PEX 8508 NT Link Interface Type 0 Endpoint, NT Virtual Interface Type 0 Endpoint, and Transparent Type 1 ports are enabled for ASPM by programming the **Link Control** register *Active State Power Management (ASPM) Control* field (offset 78h[1:0]). The PEX 8508 NT Link Interface Type 0 Endpoint and Transparent upstream port do not enter L0s nor request an ASPM L1 entry on its transmit direction. Similarly, the PEX 8508 Transparent downstream ports do not enter L0s on their transmit direction. If an ASPM L1 request is received on a PEX 8508 Transparent downstream port, a PM_Active_State_Nak message is transmitted downstream, irrespective of the *Active State Power Management (ASPM) Control* field value. The PEX 8508 allows all ports to receive a lane entry to the L0s state.

14.3.2 PCI-PM and PME Turn Off Support

When NT mode is enabled, the NT Port Link Interface Type 0 Endpoint behaves as other endpoints in the D3hot PCI-PM power states. Once in the D3hot Power state, the PEX 8508 NT Port Link Interface Type 0 Endpoint requests PCI_PM L1 entry and finally settles in the L1 link state. Only Configuration accesses and messages to the NT Port Link Interface Type 0 Endpoint are supported in the D3hot Power state. NT Host software can transmit PME_Turn_Off messages when the NT Host decides to turn Off the main power and main Reference Clock. The PEX 8508 NT Link Interface Type 0 Endpoint indicates its readiness to lose power by transmitting a PME_TO_Ack message toward the upstream device. The PME_TO_Ack message is transmitted when there is no pending TLP waiting to transmit in the PEX 8508 NT Port Link Interface upstream direction. The port requests the L2/L3 Ready state, by transmitting PM_Enter_L23 DLLP to the upstream device after transmitting a PME_TO_Ack TLP. The port settles into the L3 link state when the Power Controller removes the main power and main Reference Clock.

When the PME_Turn_Off message is received on the PEX 8508 Transparent upstream port, the port broadcasts this message to all PEX 8508 downstream devices, including the NT Port Virtual Interface Type 0 Endpoint. After the PME_TO_Ack message is received from all downstream devices and from the PEX 8508 NT Port Virtual Interface Type 0 Endpoint, the PEX 8508 Transparent upstream port transmits an aggregated PME_TO_Ack message to the upstream component after it finishes transmitting all pending TLPs to the upstream component. When NT mode is enabled, the PEX 8508 Transparent downstream ports allow the attached devices to enter the PCI-PM-compatible L1 state. The PEX 8508 NT Port Virtual Interface Type 0 Endpoint never enters the PCI-PM L1 state.

14.3.3 Message Generation

The PEX 8508 NT Port Link Interface Type 0 Endpoint never generates PM_PME messages. The PEX 8508 NT Port Virtual Interface Type 0 Endpoint never receives Set_Slot_Power_Limit messages and never generates PM_PME messages.

14.4 NT Hot Plug Support

The PEX 8508 Transparent downstream ports, with NT mode enabled, behave in the same way when NT mode is disabled. The PEX 8508 NT Port Virtual Interface Type 0 Endpoint never receives nor generates Hot Plug messages. The PEX 8508 NT Port Link Interface Type 0 Endpoint generates and receives Hot Plug messages that an Endpoint (or switch upstream port) receives/generates. The PEX 8508 NT Port Link Interface Type 0 Endpoint and Transparent upstream port implements Hot Plug client features. The *Attention Button Present*, *Attention Indicator Present*, and *Power Indicator Present* bits and their functionality are implemented on the PEX 8508 NT Port Link Interface Type 0 Endpoint and PEX 8508 downstream ports. The PEX 8508 NT Port Virtual Interface Type 0 Endpoint does not implement Hot Plug Control nor Hot Plug Client functionality, because it is an Endpoint device not connected to a Physical Link.

Figure 14-11. Sample – Intelligent Adapter Mode Hot Plug View

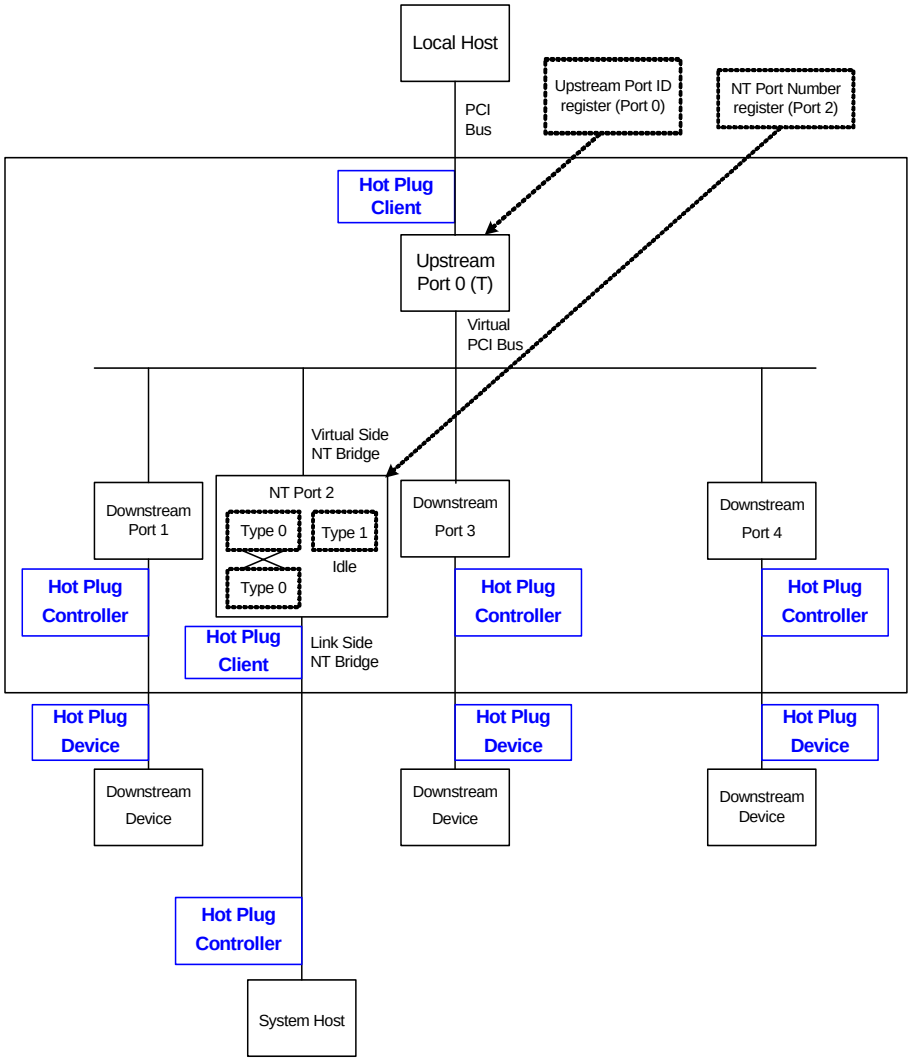

Figure 14-12. Sample Dual-Host Mode Hot Plug View

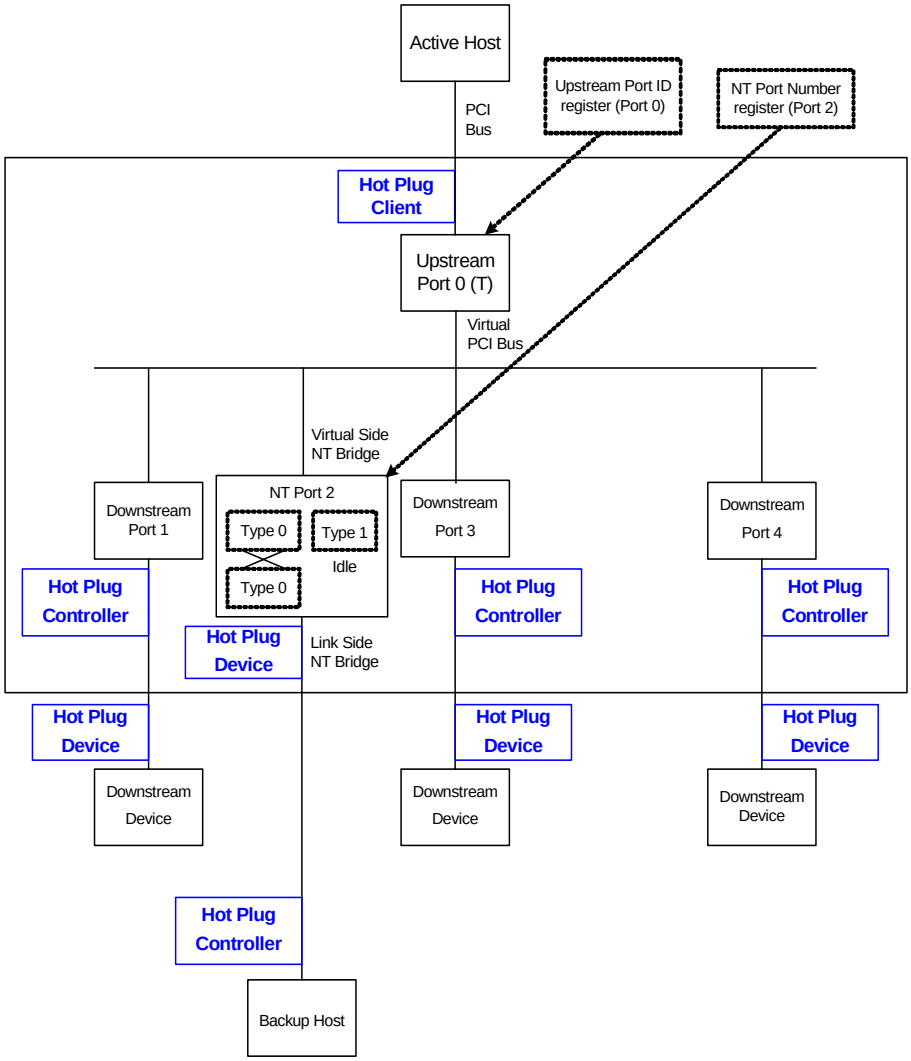

14.4.1 Hot Plug Sequence during Host-Failover

Hot Plugging the complete Active Host domain into the Backup Host is similar to Hot Plugging the endpoint device into the downstream port of a Switch or Root Port. If the Backup Host is dead, it is not a problem. If the Active Host is dead, the Backup Host first completes the Hot Plug insertion sequence before it starts the failover sequence.

The Active Host can service none or a portion of the Hot Plug sequence on a Transparent downstream port and dies before servicing the remaining sequence. The downstream port Hot Plug Controller module previously transmitted an interrupt for the next Hot Plug sequence to the failed Transparent upstream port and it did not receive service from failed active Host. The Backup Host disables "MSI and INT*x* interrupt generation" before it starts the failover sequence. After the Backup Host finishes the failover sequence, it enables "MSI and/or INT*x* interrupt generation" for the Transparent downstream port. This interrupt generation re-enabling generates Interrupt Assertion messages (or MSI) to the self-promoted Backup (Active) Host. The self-promoted Backup (Active) Host continues the remaining Hot Plug insertion/removal sequence from the failed Active Host.

Chapter 15 NT Port Interrupts

15.1 Introduction

The NT Port Virtual and Link sides can generate MSI or INT*x* interrupts, which are enabled by the **Message Control** register *MSI Enable* bit (offset 48h[16]) or **Command** register *Interrupt Disable* bit (offset 04h[10]), respectively.

Because they are endpoints, the NT Port Virtual and Link sides cannot receive Interrupt messages; therefore, if an interrupt message is received, it is reported as an error condition.

The NT Port Virtual side generates interrupts to the Local Host/Active Host for device-specific errors reported by NT Port Egress modules or Doorbell interrupts.

The NT Port Link side generates interrupts to the System/Inactive Secondary Host when device-specific errors are reported by NT Port Ingress modules.

15.2 Doorbell Interrupts

By default, all interrupt sources are masked. If software processes an interrupt, it first clears the Interrupt Mask register for the interrupt source.

The asserted INT*x* virtual wires are de-asserted when the software clears the event Status bit that caused the assertion.

The Interrupt handler has two set of registers for the NT Port – one set for Virtual Side Type 0 Configuration Space and another set for Link Side Type 0 Configuration Space.

Further details regarding MSI and INT*x* interrupts are provided in Chapter 6, "Interrupts."

15.3 Doorbell Registers

A 16-bit software-controlled Interrupt Request register and associated 16-bit Mask register are implemented for the NT Port Virtual and Link Interfaces. These registers can be accessed from the NT Port Virtual or Link Interface, in Memory or I/O space. The Doorbell mechanisms consist of the following registers:

- **Set Virtual Interface IRQ**
- **Clear Virtual Interface IRQ**
- **Set Virtual Interface IRQ Mask**
- **Clear Virtual Interface IRQ Mask**
- **Set Link Interface IRQ**
- **Clear Link Interface IRQ**
- **Set Link Interface IRQ Mask**
- **Clear Link Interface IRQ Mask**

An interrupt is asserted on the NT Port Virtual Interface when one or more of the **IRQ Set** register bits are set to 1 and their corresponding Mask bits are cleared to 0. An interrupt is de-asserted on the NT Port Virtual Interface when one or more of the **IRQ Clear** register bits are set to 1 and their corresponding Mask bits are cleared to 0. The NT Port Link Interface operates in the same manner.

The Set/Clear registers are internally the same physical Interrupt Request register, which includes two separate DWords – one DWord is used to set bits, the other is used to clear bits. The status can be read from either register. An interrupt on the NT Port Virtual Interface is asserted/de-asserted each time the **IRQ Set/Clear** registers are written to, regardless of whether the registers have already been set. After the registers are initially set, the register bits remain set, unless cleared by serial EEPROM load or a Hot or Cold Reset.

The **Set/Clear Mask** registers are also internally the same register – one interface is used to set a Mask bit, the other is used to clear a Mask bit.

Chapter 16 Non-Transparent Bridging Software Architecture

16.1 Introduction

This chapter focuses on system configuration and data transfer through the NT Port. The PEX 8508 supports two types of NT modes:

- Intelligent Adapter
- Dual-Host

The NT Port and NT-mode type are described in Chapter 14, "Non-Transparent Bridging (NTB)."

The PEX 8508 NT feature and mode (Intelligent Adapter or Dual-Host) are enabled using board-level Strapping balls. The PEX 8508 requires software support for the following:

- System Configuration
- Data transfer through NT Port
- Quality of Service (QoS) management in a switch
- Performance tuning in a switch
- Interrupt Service routine
- Hot Plug routine
- Power Management routine
- Error Handling routine

16.2 System Configuration

The PCI Express Configuration model supports two Configuration mechanisms:

- PCI-Compatible Configuration
- PCI Express Enhanced Configuration

The PCI-Compatible Configuration mechanism supports 100% binary compatibility with the *PCI r3.0* or later operating systems and corresponding bus enumeration and configuration software.

The PCI Express Enhanced Configuration mechanism is provided to increase the size of available Configuration space and optimize access mechanisms.

16.2.1 PEX 8508 Intelligent Adapter Mode

Figure 16-1 describes a sample system view with PEX 8508 NT Intelligent Adapter mode enabled.

The PEX 8508 Transparent ports are PCI-to-PCI bridges, and the PEX 8508 NT Port is two Type 0 Endpoint devices connected back-to-back. PCI Express devices must include an assigned unique Requester ID (Bus, Device, and Function Numbers).

Each PEX 8508 Transparent port has its own 4-KB PCI Express Configuration registers and the NT Port includes an 8-KB Configuration space – 4 KB for the NT Port Virtual Interface Type 0 Endpoint and another 4 KB for the NT Port Link Interface Type 0 Endpoint.

At power-up, one of PEX 8508 ports is selected as the upstream port and one of PEX 8508 downstream ports is selected as the NT Port, using board-level Strapping balls, a serial EEPROM, or through the I^2C interface.

The BIOS running in the Local Host configures the PEX 8508 upstream port, downstream ports, and NT Port Virtual Interface Type 0 Endpoint. The BIOS running in the System Host configures the NT Port Link Interface Type 0 Endpoint.

The Local Host-connected Root Complex initiates a Type 0 Configuration request to configure only the PEX 8508 upstream port and initiates Type 1 Configuration requests to configure the PCI Express hierarchy behind the PEX 8508 upstream port, including the PEX 8508 NT Port Virtual Interface Type 0 Endpoint. The Local Host is not allowed to configure the NT Port Link Interface Type 0 Endpoint using Type 0/Type 1 Configuration requests.

The System Host-connected Root Complex initiates a Type 0 Configuration request to configure only the PEX 8508 NT Port Link Interface Type 0 Endpoint. If this Root Complex initiates Type 1 Configuration requests, the NT Port Link Interface Type 0 Endpoint rejects the cycles as Unsupported Request (UR) errors.

Figure 16-1. Sample System Configuration with Non-Transparent PEX 8508

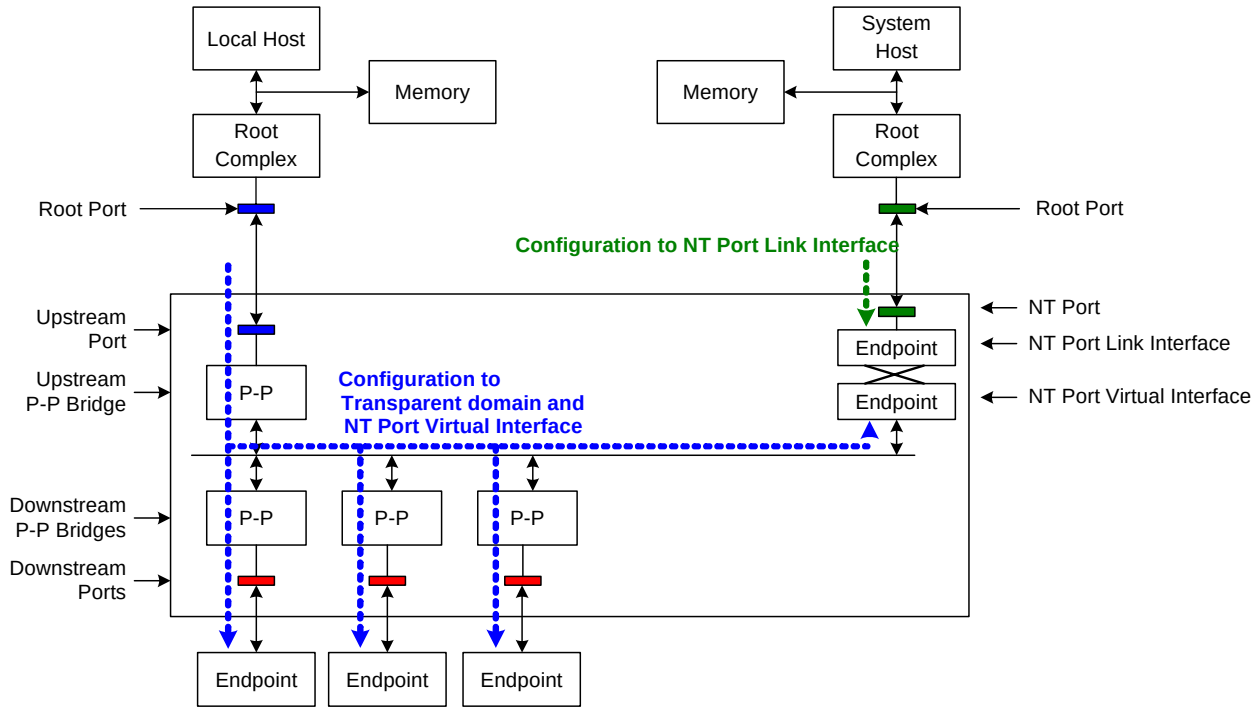

NT Port Virtual and Link Interface Type 0 Endpoint **BAR2** through **BAR5** are disabled, by default. Each BAR includes a corresponding BAR Setup register. This BAR Setup register controls the corresponding BAR, in the following manner:

1. Enables/disables BAR.
2. Programs BAR size.
3. Maps BAR group into 32- or 64-bit Address space.
4. Maps BAR into Prefetchable or Non-Prefetchable Memory space.

Refer Section 14.1.5, "BAR Setup Registers," for a detailed description of BARs and BAR Setup registers.

Use a serial EEPROM to enable the necessary BARs for inter-host domain traffic by default; otherwise, System software/BIOS and Device Driver software must work together to enable the NT Port Type 0 Endpoint BARs before assigning resources to the BARs. Alternatively, the system can enable the NT Port Type 0 Endpoint BARs through the I^2C interface before assigning resources to the BARs.

The PEX 8508 implements the following two configuration register bits to take advantage of software layer resource assignment for the NT Port. The NT Port includes **Debug Control** register *Link Interface Access Enable* and *Virtual Interface Access Enable* device-specific Configuration register bits (Port 0, offset 1DCh[29:28], respectively).

By default, the *Virtual Interface Access Enable* configuration bit is set to 1, and the *Link Interface Access Enable* configuration bit is cleared to 0. The serial EEPROM overrides these default values.

If the *Virtual Interface Access Enable* configuration bit is cleared to 0, the NT Port Virtual Interface Type 0 Endpoint returns a "Configuration Retry Status (CRS)" response (Completion with CRS status) for the received Configuration request from the Local Host. Otherwise, it accesses the corresponding Configuration registers.

If the *Link Interface Access Enable* configuration bit is cleared to 0, the NT Port Link Interface Type 0 Endpoint returns a "Configuration Retry Status (CRS)" response (Completion with CRS status) for the received Configuration request from the System Host. Otherwise, it accesses the corresponding Configuration registers.

16.2.2 Sample PEX 8508 Configuration Steps

The PEX 8508 can be configured using a serial EEPROM, software, or through the I^2C interface.

To configure the PEX 8508:

1. Select PEX 8508 operating mode:
   a. Use STRAP_MODE_SEL[1:0] balls to select the PEX 8508 operating mode. (Refer to Table 3-6, "Strapping Signals – 22 Balls," for the operating mode encoding values.)
   b. Serial EEPROM is used to override the Strapping ball selection.
2. Select port configuration:
   a. Select STRAP_PORTCFG[4:0] for PEX 8508 port configurations. (Refer to Table 4-1, "PEX 8508 Port Configurations," for the Strapping ball encoding values.)
   b. Serial EEPROM is used to override the Strapping ball selection.
3. Select upstream port:
   a. STRAP_UPSTRM_PORTSEL[3:0] selects one of PEX 8508 ports as the upstream port.
   b. Serial EEPROM is used to override the Strapping ball selection.
4. NT Port selection:
   a. STRAP_NT_UPSTRM_PORTSEL[3:0] selects one of the PEX 8508 downstream ports as the NT Port.
   b. Serial EEPROM is used to override the Strapping ball selection.
5. Power-up the system.
6. Software enumeration directives:
   a. Locating Root Port devices and Root Complex integrated Endpoint devices within a Root Complex are implementation-specific.
   b. Root Complex is allowed more than one Root Port.
   c. PCI Express hierarchy starts from Root Complex Root Port.
   d. Local Host BIOS and System Host BIOS scans the device presence behind Root Port, using a Type 0 Configuration request.
   e. BIOS reads key configuration registers (*for example*, *Header Type, Class Code*, PCI Express *Device/Port Type* field, and so forth) to locate the Device and Header types.
   f. PEX 8508 NT Port Link Interface Type 0 Endpoint responds to the Configuration access with CRS, as its "Link Interface Access Enable," is disabled by default.
   g. System Host connected Root Complex later Retries this Configuration request to PEX 8508 NT Port Link Interface Type 0 Endpoint.
   h. PEX 8508 NT Port Link Interface Type 0 Endpoint continues to Retry the Configuration request until it comprehends that "Link Interface Access Enable" is enabled.
   i. PEX 8508 upstream port accesses the corresponding Configuration registers and returns a successful Completion.
   j. BIOS detects a PCI-to-PCI bridge device behind Local Host connected Root Complex and programs the PEX 8508 upstream port **Primary Bus Number**, **Secondary Bus Number**, and **Subordinate Bus Number** registers.
   k. BIOS commences scanning devices behind the PEX 8508 upstream port, using Type 1 Configuration request.

l. If the received Type 1 Configuration request Bus Number is equal to the PEX 8508 upstream port Secondary Bus Number register value, and the Type 1 Configuration request Device Number is equal to the downstream Port Number or NT Port Number, PEX 8508 accesses the corresponding port configuration register. (Downstream ports "Port Number" and NT Port "Port Number" are equal to Device Number on the PEX 8508 Internal PCI Bus. NT Port Virtual Interface Type 0 Endpoint Device Number is the NT Port "Port Number" and NT Port Link Interface Type 0 Endpoint Device Number is the Device Number assigned by the upstream device.)

m. If the BIOS locates a PCI-to-PCI device (downstream port) on the PEX 8508 internal PCI Bus, it commences the depth-first device scan behind PEX 8508 downstream port, by programming PEX 8508 downstream port's **Primary Bus Number**, **Secondary Bus Number**, and **Subordinate Bus Number** registers.

n. PEX 8508 upstream port routes received Type 1 Configuration request to a PEX 8508 downstream port, if the received Type 1 Configuration request Bus Number is greater than the PEX 8508 upstream port Secondary Bus Number and less than or equal to PEX 8508 upstream port Subordinate Bus Number, as well as within the PEX 8508 downstream port Secondary Bus Number and Subordinate Bus Number window.

o. If the received Type 1 Configuration request Bus Number is equal to PEX 8508 downstream port Secondary Bus Number, the PEX 8508 downstream converts this Type 1 Configuration access to a Type 0 Configuration access, if the Type 1 Configuration request Device Number is 0. If the Type 1 Configuration request Device Number is non-zero, the PEX 8508 downstream port returns an Unsupported Request (UR) error.

p. If the received Type 1 Configuration request Bus Number is greater than the PEX 8508 downstream port Secondary Bus Number and less than or equal to the PEX 8508 downstream port Subordinate Bus Number, the PEX 8508 downstream port forwards this Type 1 Configuration request to the downstream device, unmodified.

q. If BIOS locates an NT Port Virtual Interface Type 0 Endpoint of the PEX 8508 internal PCI Bus, it stops scanning behind the NT Port, because it is an Endpoint and the PCI Express hierarchy ends in an Endpoint.

r. After locating all devices within the PCI Express hierarchy, BIOS commences resource assignment (*for example*, Memory, I/O, and/or Interrupt resource).

s. If the application is not using a serial EEPROM, system software/BIOS requires modification to implement the following:

- Assign a Memory, I/O, and/or Interrupt resource to the PEX 8508 upstream port **BAR0** (Memory BAR) or PEX 8508 NT Port Virtual Interface **BAR1** (I/O BAR) before assigning Memory resources to PEX 8508 NT Port Virtual Interface Type 0 Endpoint **BAR2** to **BAR5**
- Use a Memory/I/O-Mapped cycle (refer to Section 17.2, "Register Access," for details) to program the PEX 8508 NT Port Virtual Interface Type 0 Endpoint **BAR Setup** register
- Assign a resource to the PEX 8508 NT Port Virtual Interface Type 0 Endpoint

t. Use a Memory/I/O-Mapped cycle to program the NT Port Link Interface Type 0 Endpoint **BAR Setup** register.

u. Use a Memory-Mapped cycle to enable the *Link Interface Access Enable* configuration bit.

v. After enabling the PEX 8508 NT Port Link Interface Type 0 Endpoint, System Host BIOS is allowed to enumerate and assign a resource to the PEX 8508 NT Port Link Interface Type 0 Endpoint.

w. Use a Memory/I/O-Mapped cycle to program the PEX 8508 NT Port Type 0 Endpoint Address Translation registers. (Refer to Chapter 17, "NT Port Virtual Interface Registers," and Chapter 18, "NT Port Link Interface Registers," for details.)

x. Use a Memory/I/O-Mapped cycle to program the PEX 8508 NT Port Type 0 Endpoint BAR Limit registers if application is to efficiently use the memory resource. (Refer to Chapter 17, "NT Port Virtual Interface Registers," and Chapter 18, "NT Port Link Interface Registers," for details.)

y. Use a Memory/I/O-Mapped cycle to program the PEX 8508 NT Port Type 0 Endpoint "Send LUT Entry" and "Receive LUT Entry" registers. (Refer to Chapter 17, "NT Port Virtual Interface Registers," and Chapter 18, "NT Port Link Interface Registers," for details.)

16.2.3 PEX 8508 Dual-Host Mode

PEX 8508 NT Dual-Host mode is similar to PEX 8508 NT Intelligent Adapter mode, except the default values of the NT Port *Link Interface Access Enable* and *Virtual Interface Access Enable* bits (offset 1DCh[29:28], respectively) are set to 1. Both hosts connected to the PEX 8508 upstream port and NT Port Link Interface Type 0 Endpoint can concurrently enumerate the devices. The PEX 8508 does not generate a CRS response in NT Dual-Host mode.

16.2.4 Host-Failover Application

The Host-Failover application is based upon the basic Dual-Host configuration, and dynamic swapping of the upstream port and NT Port is supported on Port 0. The Active Host periodically transmits heartbeat messages, by way of the PEX 8508 to the Backup Host, to indicate that it remains active. When the Backup Host fails to receive heartbeat messages before its Fail Detect Timer expires, it starts the Failover process. The Backup Host halts cross-domain traffic before it starts the failover. The Backup Host uses the Memory-Mapped access to the **PCI Express Capability List and Capabilities** register (offset 68h) to execute the failover. The Backup Host follows the ensuing procedure to take control:

1. Failover Detected:
   a. Backup Host detects the Active Host's failure condition, then starts the Failover process (*such as*, Heartbeat Message Reception timeout).
   b. Upstream port remains active in this state.

2. Upstream Port Demotion:
   a. PEX 8508 is not in Reset when the Failover process starts.
   b. **Debug Control** register *Upstream Port Hot Reset and Link Down Reset Propagation Disable* bit (Port 0, offset 1DCh[20]) can disable Reset generation due to a Hot Reset and the upstream port DL_Down state. The bit is asserted, by default, for Dual-Host mode.
   c. As one of the first steps in the Failover process, the Backup Host demotes the upstream port by writing 0000b into the **PCI Express Capabilities** register (PCI Express Endpoint) *Device/Port Type* field (offset 68h[23:20]).
   d. Transaction Layer Ingress snoops this access and informs Event *c* detection to the TLP-destined Transparent upstream port.
   e. Event *d* causes the upstream port Physical Layer to bring down its upstream link, which generates the upstream port DL_Down state.
   f. Event *d* also causes the upstream port to change its *Device/Port Type* field to "PCI Express Endpoint", and changes the *PCI Class Code* to "other bridged devices". The Transparent upstream port becomes a PCI Express endpoint.
   g. Upstream port Transaction Layer Egress module drops all outgoing packets to the upstream device when it comprehends the upstream port DL_Down state.

h. **Debug Control** register *Upstream Port ID* and *NT Port Number* fields (offset 1DCh[10:8 and 27:24], respectively) remain unchanged in this state.

i. Downstream port internal modules forward/generate the packet for the upstream device, which is transmitted to the previous upstream port, and the packets are dropped by the demoted upstream port Transaction Layer Egress module due to this DL_Down state.

j. When a Transaction Layer Ingress module receives TLPs from the demoted upstream connected device, the module processes the normal upstream port type of address decoding and forwards the packet to the destination port, based upon AMCAM, IOCAM, or BusNoCAM lookup.

3. NT Port (Self) Promotion as a New Upstream Port:

a. Software can disable the *I/O Access Enable*, *Memory Access Enable*, and *Bus Master Enable*, Interrupt generation on the NT Port Virtual Interface, and Error message generation CSR bits, if it does not want to receive spurious traffic immediately after self-promotion.

b. Refer to the *PCI-to-PCI Bridge r1.2* **Command** register for these CSR descriptions and the **Device Control** register for Error message generation Enable CSR bits.

c. Backup Host promotes itself as an Active Host by writing 0101b into the NT Port Virtual Interface Type 0 Configuration **PCI Express Capabilities** register (Transparent upstream port) *Device/Port Type* field (offset 68h[23:20]).

d. PEX 8508 swaps the **Debug Control** register *Upstream Port ID* and *NT Port Number* field (offset 1DCh[10:8 and 27:24], respectively) values when a Transaction Layer ingress informs Event *c* of this condition.

e. PEX 8508 processes the port transition. PEX 8508 converts the demoted Transparent upstream port to the NT Port, and the previous NT Port to the Transparent upstream port.

f. New upstream port retrieves previously programmed NT Port Virtual Interface Type 0 CSR values.

g. PEX 8508 does not swap the Configuration space value from previous upstream port to new upstream port by itself after failover. Software running in Promoted Active Host should follow the ensuing procedure to bring the system into a communicating state:

- Copy the previous upstream port Configuration space value to a new upstream port Configuration space value.
- Copy the previous NT Port Virtual Interface Type 0 endpoint Configuration space value to a new NT Port Virtual Interface Type 0 Endpoint Configuration space value.
- If possible, reset the entire hierarchy and restart the system, using full software re-enumeration.

16.3 Data Transfer through NT Port

The following discusses the configuration registers mainly programmed for data transfer through the NT Port.

To transfer data from NT Port Virtual Interface to NT Port Link Interface direction:

1. Assign Memory space to NT Port Virtual Interface Type 0 Endpoint BARs.
2. Enable NT Port Virtual Interface Type 0 Endpoint *Memory Access Enable* bit (offset 04h[1]).
3. Enable NT Port Link Interface Type 0 Endpoint *Bus Master Enable* bit (offset 04h[2]).
4. Program NT Port Virtual Interface Type 0 Endpoint Address Translation registers with transaction completer (target) BAR value. If the application is using Lookup Table-Based Address translation, it must enable the corresponding LUT Entry as well. Address translation register values can be dynamically changed by a device driver, depending upon where the Requester is located. Before changing the Address Translation register values, the device driver must ensure an outstanding request is not pending to the NT Port.
5. Enable and program NT Port Virtual Interface "Send LUT Entry" registers with Requester ID (Bus, Device, and Function Numbers). "Send LUT Entry" register values can be dynamically changed by a device driver, depending upon the request enabled to communicate through the PEX 8508 NT Port. Before changing the "Send LUT Entry," the device driver must ensure an outstanding request is not pending for that Requester.

To transfer data from NT Port Link Interface to NT Port Virtual Interface direction:

1. Assign Memory space to NT Port Link Interface Type 0 Endpoint BARs.
2. Enable NT Port Link Interface Type 0 Endpoint *Memory Access Enable* bit (offset 04h[1]).
3. Enable NT Port Virtual Interface Type 0 Endpoint *Bus Master Enable* bit (offset 04h[2]).
4. Program NT Port Link Interface Type 0 Endpoint Address Translation registers with transaction completer (target) BAR value. Address Translation register values can be dynamically changed by a device driver, resident on the Host that interfaces to the Virtual Link Side. Before changing the Address translation register values, the device driver must ensure an outstanding request is not pending to the NT Port.
5. Enable and program NT Port Link Interface "Receive LUT Entry" registers with Requester Bus Number and Device Number values. "Receive LUT Entry" register values can be dynamically changed by a device driver, depending upon the request enabled to communicate through the PEX 8508 NT Port. Before changing the "Receive LUT Entry," the device driver must ensure an outstanding request is not pending for that Requester.

Chapter 17 NT Port Virtual Interface Registers

17.1 Introduction

This chapter defines the registers for the PEX 8508 Non-Transparent (NT) Port Virtual Interface (interface) registers. The NT Port includes two sets of configuration-capability, control, and status registers to support the Virtual and Link Interfaces. Table 17-1 defines the NT Port Virtual Interface register mapping.

The NT Port Link Interface registers are defined in Chapter 18, "NT Port Link Interface Registers." The Transparent mode registers are defined in Chapter 13, "Transparent Port Registers."

For further details regarding register names and descriptions, refer to the following specifications:

- *PCI r3.0*
- *PCI Power Mgmt. r1.2*
- *PCI Express Base r1.1*

Table 17-1. NT Port Virtual Interface Type 0 Register Map

31 30 29 28 27 26 25 24 23 22 21 20	19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
PCI-Compatible Type 0 Configuration Registers			Capability Pointer (40h)	00h … 34h … 3Ch
		Next Capability Pointer (48h)	Capability ID (01h)	40h
Power Management Capability Registers				44h
		Next Capability Pointer (68h)	Capability ID (05h)	48h
Message Signaled Interrupt Capability Registers				… 64h
		Next Capability Pointer (00h)	Capability ID (10h)	68h
PCI Express Capability Registers				… 8Ch
NT Port Virtual Interface Registers				90h … FCh
Next Capability Offset (FB4h)	1h	PCI Express Extended Capability ID (0003h)		100h
Device Serial Number Extended Capability Registers				104h 108h
Reserved				10Ch – 134h
Next Capability Offset (148h)	1h	PCI Express Extended Capability ID (0004h)		138h
Power Budgeting Extended Capability Registers				… 144h
Next Capability Offset (000h)	1h	PCI Express Extended Capability ID (0002h)		148h
Virtual Channel Extended Capability Registers				… 1C4h
Device-Specific Registers				1C8h … C38h
Non-Transparent Bridging-Specific Registers				C3Ch … DF0h
Reserved				DF4h – FB0h
Next Capability Offset (138h)	1h	PCI Express Extended Capability ID (0001h)		FB4h
Advanced Error Reporting Capability Registers				… FFCh

17.2 Register Access

The PEX 8508 NT Port Virtual Interface implements a 4-KB Configuration space. The lower 256 bytes (offsets 00h through FFh) is the PCI-compatible Configuration space, and the upper 960 Dwords (offsets 100h through FFFh) is the PCI Express Extended Configuration space. The PEX 8508 supports four mechanisms for accessing NT Port Virtual Interface registers:

- PCI Express Base r1.1 Configuration Mechanism
- Device-Specific Memory-Mapped Configuration Mechanism
- Device-Specific I/O-Mapped Configuration Mechanism
- Device-Specific Cursor Mechanism

17.2.1 *PCI Express Base r1.1* Configuration Mechanism

The PCI Express Configuration Mechanism is divided into two mechanisms:

- *PCI r3.0*-Compatible Configuration
- PCI Express Enhanced Configuration

The *PCI r3.0*-compatible Configuration mechanism provides Conventional PCI access to the first 256 bytes (the bytes at offsets 00h through FFh) of the NT Port Virtual Interface Configuration register space. The PCI Express Enhanced Configuration mechanism provides access to the remaining 4 KB (offsets 100h through FFFh).

17.2.1.1 *PCI r3.0*-Compatible Configuration Mechanism

The *PCI r3.0*-Compatible Configuration mechanism provides standard access to the PEX 8508 NT Port Virtual Interface's first 256 bytes (the bytes at offsets 00h through FFh) of the PCI Express Configuration space. (Refer to Figure 17-1.) This mechanism is used to access the PEX 8508 NT Port Virtual Interface Type 0 (PCI endpoint) registers:

- PCI-Compatible Type 0 Configuration Registers
- Power Management Capability Registers
- Message Signaled Interrupt Capability Registers
- PCI Express Capability Registers

The *PCI r3.0*-Compatible Configuration mechanism uses PCI Type 0 and Type 1 Configuration transactions to access the PEX 8508 Configuration registers. The PEX 8508 upstream port captures the Bus and Device Numbers assigned by the upstream device on the PCI Express link attached to the PEX 8508 upstream port, as required by the *PCI Express Base r1.1*.

The PEX 8508 decodes all Type 1 Configuration accesses received on its upstream port, when any of the following conditions exist:

- If the Bus Number specified in the Configuration access is the number of the PEX 8508 internal virtual PCI Bus, the PEX 8508 automatically converts the Type 1 Configuration access into the appropriate Type 0 Configuration access for the specified device.
  - If the specified device corresponds to the NT Port Virtual Interface (or to the PCI-to-PCI bridge in one of the PEX 8508 downstream Transparent ports), the PEX 8508 processes the Read or Write request to the downstream port register specified in the original Type 1 Configuration access.
  - If the specified Device Number does not correspond to any of the PEX 8508 downstream port Device Numbers, the PEX 8508 responds with an *Unsupported Request* (UR).

Because the *PCI r3.0*-compatible Configuration mechanism is limited to the first 256 bytes of the NT Port Virtual Interface Configuration register space, one of the following must be used to access beyond byte FFh:

- PCI Express Enhanced Configuration Mechanism
- Device-Specific Memory-Mapped Configuration Mechanism
- Device-Specific Cursor Mechanism

The *PCI r3.0*-compatible Configuration mechanism uses the same request format as the PCI Express Enhanced Configuration Mechanism. For PCI-compatible Configuration requests, the Extended Register Address field must be all zeros (0).

Do not use this mechanism to access the PEX 8508 Device-Specific Configuration registers.

17.2.1.2 PCI Express Enhanced Configuration Mechanism

The PCI Express Enhanced Configuration mechanism uses a flat, Root Complex Memory-Mapped Address space to access device Configuration registers. In this case, the Memory address determines the Configuration register accessed, and the Memory data returns the addressed register's contents. The Root Complex converts the Memory transaction into a Configuration transaction.

This mechanism is used to access the NT Port Virtual Interface Type 0 registers:

- PCI-Compatible Type 0 Configuration Registers
- Power Management Capability Registers
- Message Signaled Interrupt Capability Registers
- PCI Express Capability Registers
- Device Serial Number Extended Capability Registers
- Power Budgeting Extended Capability Registers
- Virtual Channel Extended Capability Registers
- Advanced Error Reporting Capability Registers

Do not use this mechanism to access the PEX 8508 Device-Specific Configuration registers.

17.2.2 Device-Specific Memory-Mapped Configuration Mechanism

The Device-Specific Memory-Mapped Configuration mechanism provides a method to access the PEX 8508 Port Configuration registers of all ports in a single Memory map, as illustrated in Figure 17-1. The registers of each port are located within a 4-KB range.

When the NT Port is enabled at Fundamental Reset, the NT Port Virtual Interface and Link Interface Configuration registers are used in place of the Type 1 Configuration registers for that port.

To utilize the Device-Specific Memory-Mapped Configuration mechanism, use the *PCI r3.0*-Compatible Configuration Mechanism to program the PEX 8508 upstream port **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively). After the PEX 8508 upstream port Memory-Mapped register Base address is set, the upstream port register is accessed with Memory Reads from and Writes to the Configuration Space registers. The NT Port registers are accessed with Memory Reads from and Writes to the 4-KB range, starting at offset 64 KB for the Virtual Interface registers and offset 68 KB for the Link Interface registers.

This mechanism is used to access all PEX 8508 Configuration registers.

Figure 17-1. PEX 8508 Register Offset from Upstream Port BAR0/1 Base Address (Non-Transparent Mode)

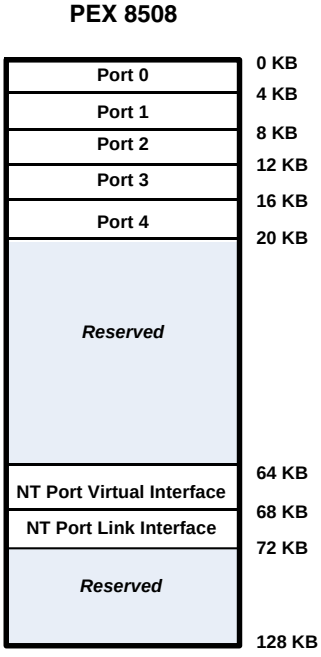

17.2.3 Device-Specific I/O-Mapped Configuration Mechanism

The first 256 bytes of NT Port Virtual Interface Configuration Space registers are directly accessed by an I/O transaction. The NT Port Virtual Interface **Base Address 1** register (**BAR1**. offset 14h) is used for I/O-mapped access. (Refer to Figure 17-2.)

Extended Configuration Space registers are accessed by using the Cursor mechanism in I/O space.

Figure 17-2. I/O-Mapped Configuration Space View

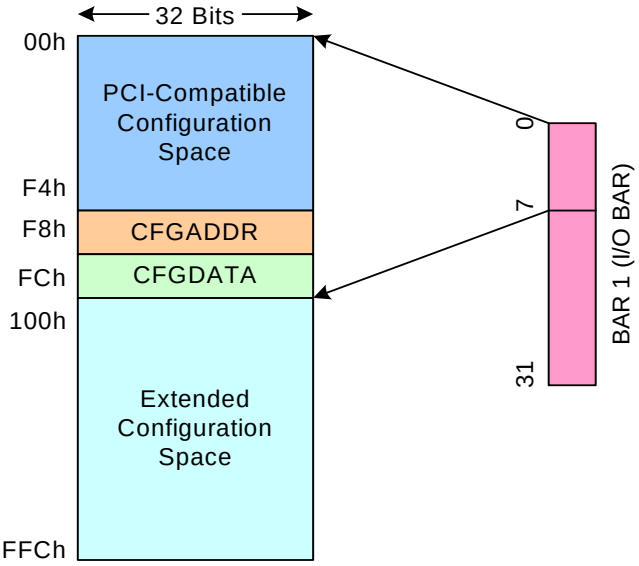

17.2.4 Device-Specific Cursor Mechanism

In Figure 17-2, the software uses the **Configuration Address Window** (CFGADDR) register (offset F8h) to point to the NT Port Virtual or Link Interface Configuration Space registers, including the Extended Space register.

Software uses the **Configuration Data Window** (CFGDATA) register (offset FCh) to write to or read from the selected Configuration Space registers.

Refer to Section 17.8.4, "NT Port Virtual Interface Cursor Mechanism Control Registers," for the register descriptions.

17.3 Register Descriptions

The remainder of this chapter details the PEX 8508 NT Port Virtual Interface registers, including:

- Bit/field names
- Description of register functions in the PEX 8508 NT Port Virtual and Link Interfaces
- Type (*such as* RW or HwInit; refer to Table 13-3, "Register Types, Grouped by User Accessibility." for Type descriptions)
- Whether the power-on/reset value can be modified, by way of the PEX 8508 serial EEPROM or I^2C initialization feature
- Default power-on/reset value

17.4 PCI-Compatible Type 0 Configuration Registers

Table 17-2. PCI-Compatible Type 0 Configuration Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Device ID		Vendor ID		00h
Status		Command		04h
Class Code			Revision ID	08h
BIST ***(Not Supported)***	Header Type	Master Latency Timer	Cache Line Size	0Ch
Reserved				10h
Base Address 1				14h
Base Address 2				18h
Base Address 3				1Ch
Base Address 4				20h
Base Address 5				24h
Reserved				28h
Subsystem ID		Subsystem Vendor ID		2Ch
Expansion ROM Base Address				30h
Reserved			Capability Pointer (40h)	34h
Reserved				38h
Minimum Latency ***(Reserved)***	Minimum Grant ***(Reserved)***	Interrupt Pin	Interrupt Line	3Ch

Register 17-1. 00h Configuration ID

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Vendor ID** Unless overwritten by the serial EEPROM, returns the PLX PCI-SIG-assigned Vendor ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Vendor ID with another Vendor ID.	HwInit	Yes	10B5h
31:16	**Device ID** Unless overwritten by the serial EEPROM, the PEX 8508 returns 8508h, the PLX-assigned Device ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX-assigned Device ID with another Device ID.	HwInit	Yes	8508h

Register 17-2. 04h Command/Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Command			
0	**I/O Access Enable** 0 = PEX 8508 ignores I/O requests on the NT Port Virtual Interface 1 = PEX 8508 accepts I/O requests received on the NT Port Virtual Interface	RW	Yes	0
1	**Memory Access Enable** 0 = PEX 8508 ignores Memory requests on the NT Port Virtual Interface 1 = PEX 8508 accepts Memory requests received on the NT Port Virtual Interface	RW	Yes	0
2	**Bus Master Enable** Controls the PEX 8508 Memory request forwarding in the upstream direction. Does not affect message forwarding nor Completions in the upstream direction. 0 = PEX 8508 handles Memory requests received on the NT Port Link Interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8508 returns a Completion with UR Completion status 1 = PEX 8508 forwards Memory requests in the upstream direction	RW	Yes	0
3	**Special Cycle Enable** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
4	**Memory Write and Invalidate Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
5	**VGA Palette Snoop** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
6	**Parity Error Response Enable** Controls bit 24 (*Master Data Parity Error* bit).	RW	Yes	0
7	**IDSEL Stepping/Wait Cycle Control** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0

Register 17-2. 04h Command/Status *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
8	**SERR# Enable** Controls bit 30 (*Signaled System Error* bit). 1 = Enables reporting of Fatal and Non-Fatal errors detected by the NT Port Virtual Interface to the Root Complex	RW	Yes	0
9	**Fast Back-to-Back Transactions Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
10	**Interrupt Disable** 0 = NT Port Virtual Interface is enabled to generate INT*x* Interrupt messages 1 = NT Port Virtual Interface is prevented from generating INT*x* Interrupt messages	RW	Yes	0
15:11	***Reserved***	RsvdP	No	00h
Status				
18:16	***Reserved***	RsvdP	No	000b
19	**Interrupt Status** 0 = No INT*x* Interrupt pending 1 = INT*x* Interrupt pending internally to the NT Port Virtual Interface	RO	Yes	0
20	**Capability List** Set to 1, as required by the *PCI Express Base r1.1*.	RO	Yes	1
21	**66 MHz Capable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
22	***Reserved***	RsvdP	No	0
23	**Fast Back-to-Back Transactions Capable** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0

Register 17-2. 04h Command/Status *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
24	**Master Data Parity Error** When the *Parity Error Response Enable* bit is set to 1, the NT Port Virtual Interface sets this bit to 1 when the NT Port: • Forwards the poisoned TLP Write request from the Link Interface to the Virtual Interface, **or** • Receives a Completion marked as poisoned on the Virtual Interface When the *Parity Error Response Enable* bit is cleared to 0, the PEX 8508 never sets this bit. This error is natively reported by the **Uncorrectable Error Status** register *Poisoned TLP Status* bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0
26:25	**DEVSEL# Timing** ***Not supported*** Always cleared to 00b.	RsvdP	No	00b
27	**Signaled Target Abort** Set to 1, when the NT Port forwards a Completion with Completer Abort (CA) status from the Link Interface to the Virtual Interface. ***Note:*** *When set during a forwarded Completion, the **Uncorrectable Error Status** register Completer Abort Status bit is not set.*	RW1C	Yes	0
28	**Received Target Abort** Defaults to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
29	**Received Master Abort** Defaults to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
30	**Signaled System Error** When the *SERR# Enable* bit is set to 1, the NT Port Virtual Interface sets this bit to 1 when transmitting an ERR_FATAL or ERR_NONFATAL message to its upstream port. ***Note:*** *When set during a forwarded Completion, the **Uncorrectable Error Status** register Completer Abort Status bit (offset FB8h[15]) is not updated, because the NT Port does not log the requests that it forwards.*	RW1C	Yes	0
31	**Detected Parity Error** The NT Port Virtual Interface sets this bit to 1 when receiving a Poisoned TLP, regardless of the *Parity Error Response Enable* bit state. This error is natively reported by the **Uncorrectable Error Status** register *Poisoned TLP Status* bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0

Register 17-3. 08h Class Code and Revision ID

<table>
<tr><th>Bit(s)</th><th>Description</th><th>Type</th><th>Serial EEPROM and I²C</th><th>Default</th></tr>
<tr><td>7:0</td><td>Revision ID
Unless overwritten by the serial EEPROM, returns ACh, the PLX-assigned Revision ID for this version of the PEX 8508. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Revision ID with another Revision ID.
Note: When a serial EEPROM is not used, the default values for this field are loaded. When using a serial EEPROM, refer to the following table for the appropriate configuration settings for this field.
<table>
<tr><th></th><th>Serial EEPROM Value</th><th>Results in Actual Register Value</th></tr>
<tr><td rowspan="2">Revision ID</td><td>ACh</td><td>AAh</td></tr>
<tr><td>AAh</td><td>ACh</td></tr>
</table></td><td>RO</td><td>Yes</td><td>ACh</td></tr>
<tr><td colspan="3">Class Code</td><td>068000h</td></tr>
<tr><td>15:8</td><td>Register-Level Programming Interface
Reserved, as required by the PCI r3.0.</td><td>RO</td><td>Yes</td><td>00h</td></tr>
<tr><td>23:16</td><td>Sub-Class Code
Other bridge devices.</td><td>RO</td><td>Yes</td><td>80h</td></tr>
<tr><td>31:24</td><td>Base Class Code
Bridge devices.</td><td>RO</td><td>Yes</td><td>06h</td></tr>
</table>

Register 17-4. 0Ch Miscellaneous Control

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Cache Line Size** System Cache Line Size. Implemented as a Read-Write field for Conventional PCI compatibility purposes and does not impact PEX 8508 functionality.	RW	Yes	00h
15:8	**Master Latency Timer** ***Not supported*** Cleared to 00h, as required by the *PCI Express Base r1.1*.	RsvdP	No	00h
22:16	**Configuration Layout Type** Type 0 Configuration Header for the NT Port.	RO	Yes	00h
23	**Header Type** 0 = PEX 8508 is a single-function device	RO	Yes	0
31:24	**BIST** ***Not supported*** BIST Pass/Fail.	RsvdP	No	00h

Register 17-5. 14h Base Address 1 (NT Port Virtual Interface I/O Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**I/O Space Indicator** I/O BAR when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is set to 11b.	RO	Yes	1
	Reserved when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is not set to 11b.	RsvdP	No	0
7:1	***Reserved***	RsvdP	No	0000_000b
31:8	**I/O Base Address** 256-byte I/O Space Base address when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is set to 11b.	RW	Yes	0000_00h
	Reserved when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is not set to 11b.	RsvdP	No	0000_00h

Note: *When software writes to the NT Port Virtual Interface **Base Address 1 (BAR1)** register, the value is automatically copied to the **NT Port Virtual Interface BAR1 (Shadow Copy)** register (offset D6Ch).*

Register 17-6. 18h Base Address 2 (NT Port Virtual Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR – only value supported	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 01b, 10b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Address 2** Contains the software-assigned Memory Space Base address: • Enabled and sized by the **NT Port Virtual Interface Memory BAR2 Setup** register (offset D4h) • Used for Memory transactions crossing the NT Port • Minimum Address range requested is 4 KB • Uses direct address translation • Includes a **Limit** register	RW	Yes	0000_0h
	Reserved when the **NT Port Virtual Interface Memory BAR2 Setup** register *BAR2 Type* field (offset D4h[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Note: *When software writes to the NT Port Virtual Interface **Base Address 2 (BAR2)** register, the value is automatically copied to the **NT Port Virtual Interface BAR2 (Shadow Copy)** register (offset D70h).*

Register 17-7. 1Ch Base Address 3 (NT Port Virtual Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Memory Space Indicator** 0 = Memory BAR – only value supported	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory Address space 01b, 10b, 11b = ***Not allowed***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
17:4	***Reserved***	RsvdP	No	000h
31:18	**Base Address 3** Contains the software-assigned Memory Space Base address: • Enabled and sized by the **NT Port Virtual Interface Memory BAR3 Setup** register (offset D8h) • No **Limit** register • Used for Memory transactions crossing the NT Port • Minimum Address range requested is 256 KB • Uses LUT address translation	RW	Yes	0000h
	Reserved when the **NT Port Virtual Interface Memory BAR2 Setup** register *BAR2 Type* field (offset D4h[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Note: *When software writes to the NT Port Virtual Interface **Base Address 3** register (**BAR3**), the value is automatically copied to **NT Port Virtual Interface BAR3 (Shadow Copy)** (offset D74h).*

Register 17-8. 20h Base Address 4 (NT Port Virtual Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR – only value supported	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 10b = Mappable anywhere in 64-bit Memory space 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Address 4** Contains the software-assigned Memory Space Base address: • Enabled and sized by the **NT Port Virtual Interface Memory BAR4/5 Setup** register (offset DCh) • Used for Memory transactions crossing the NT Port • Minimum Address range requested is 4 KB • Uses direct address translation	RW	Yes	0000_0h
	Reserved when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Note: *When software writes to the NT Port Virtual Interface **Base Address 4** register (**BAR4**), the value is automatically copied to **NT Port Virtual Interface BAR4 (Shadow Copy)** (offset D78h).*

Register 17-9. 24h Base Address 5 (NT Port Virtual Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 5** NT Port Virtual Interface upper 32-bit address when **BAR4/5** is implemented as a 64-bit BAR [**NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is set to 10b]. RW, based upon the **NT Port Virtual Interface Memory BAR4/5 Setup** and **NT Port Virtual Interface Memory BAR5 Setup** registers (offsets DCh and E0h, respectively). The **BAR4/5** group uses direct address translation. Contains a **Limit** register.	RW	Yes	0-0h
	Reserved when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Note: *When software writes to the NT Port Virtual Interface **Base Address 5** register (**BAR5**), the value is automatically copied to **NT Port Virtual Interface BAR5 (Shadow Copy)** (offset D7Ch).*

Register 17-10. 2Ch Subsystem ID and Subsystem Vendor ID

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Subsystem Vendor ID** Unless overwritten by the serial EEPROM, returns the PLX PCI-SIG-assigned Vendor ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Vendor ID with another Vendor ID.	HwInit	Yes	10B5h
31:16	**Subsystem ID** Unless overwritten by the serial EEPROM, the PEX 8508 returns 8508h, the PLX-assigned Device ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX-assigned Device ID with another Device ID.	HwInit	Yes	8508h

Register 17-11. 30h Expansion ROM Base Address

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
0	**Expansion ROM Enable** 0 = NT Port Virtual Interface Expansion ROM is disabled		RsvdP	No	0
	1 = NT Port Virtual Interface Expansion ROM is disabled, and NT Port Link Interface Expansion ROM is enabled		RO	Yes	0
10:1	***Reserved***		RsvdP	No	0-0h
31:11	**Expansion ROM Base Address** Expansion ROM minimum size is 2 KB (program register initially to FFFF_F801h). Expansion ROM maximum size is 32 KB (program register initially to FFFF_8001h). ***Note: This BAR must not be programmed to enable more than 32 KB.*** *Expansion ROM is limited to 32 KB, because the largest serial EEPROM that can be used is 64 KB (limit of 16-bit addressing) and the Expansion ROM image is stored in the serial EEPROM, beginning at address 32 KB.*	When Bit 0 = 0	RsvdP	No	0-0h
		When Bit 0 = 1	RW	Yes	0-0h

Note: Expansion ROM can be enabled in either the NT Port Virtual Interface or NT Port Link Interface, but not both simultaneously*. The default values of the NT Port Virtual Interface Expansion ROM Enable (NT Virtual Port, offset 30h[0]=0) and NT Port Link Interface Expansion ROM Enable (NT Link Port, offset 30h[0]=1) bits enable the NT Port Link Interface Expansion ROM, and the NT Port Virtual Interface* ***Expansion ROM Base Address*** *register (NT Virtual Port, offset 30h) is cleared and is not programmable.*

If the NT Port Virtual Interface Expansion ROM is enabled (NT Virtual Port, offset 30h[0]=1), the NT Port Link Interface ***Expansion ROM Base Address*** *register (NT Link Port, offset 30h) is cleared and is not programmable.*

Register 17-12. 34h Capability Pointer

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Capability Pointer** Default 40h points to the **Power Management Capability** register.	RO	Yes	40h
31:8	***Reserved***	RsvdP	No	0000_00h

Register 17-13. 3Ch Interrupt

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Interrupt Line** The Interrupt Line Routing value communicates interrupt line routing information. Values in this register are programmed by system software and are system architecture-specific. The value is used by device drivers and operating systems.	RW	Yes	00h
15:8	**Interrupt Pin** Identifies the Conventional PCI interrupt message(s) the device (or device function) uses. Only value 00h or 01h is allowed in the PEX 8508. 00h = Indicates that the device does not use Conventional PCI Interrupt message(s) 01h, 02h, 03h, and 04h = Maps to Conventional PCI Interrupt messages for INTA#, INTB#, INTC#, and INTD#, respectively	RO	Yes	01h
23:16	**Minimum Grant** ***Reserved*** Does not apply to PCI Express.	RsvdP	No	00h
31:24	**Minimum Latency** ***Reserved*** Does not apply to PCI Express.	RsvdP	No	00h

17.5 Power Management Capability Registers

This section details the NT Port Virtual Interface Power Management Capability registers. Table 17-3 defines the register map.

Table 17-3. Power Management Capability Register Map (All Ports)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Power Management Capability		Next Capability Pointer (48h)	Capability ID (01h)	40h
Data	Power Management Control/ Status Bridge Extensions	Power Management Status and Control		44h

Register 17-14. 40h Power Management Capability

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
7:0	**Capability ID** Default = 01h – only value allowed.	RO	Yes	01h
15:8	**Next Capability Pointer** Default 48h points to the **Message Signaled Interrupt Capability** register.	RO	Yes	48h
18:16	**Version** Default = 011b – only value allowed.	RO	Yes	011b
19	**PME Clock** Does not apply to PCI Express. Returns 0.	RsvdP	No	0
20	***Reserved***	RsvdP	No	0
21	**Device-Specific Initialization** Default 0 indicates that Device-Specific Initialization is ***not*** required.	RO	Yes	0
24:22	**AUX Current** Value can be strapped to 111b in the serial EEPROM.	RO	Yes	000b
25	**D1 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D1 Power state.	RsvdP	No	0
26	**D2 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D2 Power state.	RsvdP	No	0
31:27	**PME Support** PME messages are disabled by default.	RO	Yes	0000_0b

Register 17-15. 44h Power Management Status and Control

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
Power Management Status and Control				
1:0	**Power State** This field is used to determine the current power state of the port, and to set the port into a new power state. 00b = D0 01b = D1 – ***Not supported*** 10b = D2 – ***Not supported*** 11b = D3hot If software attempts to write an unsupported state to this field, the Write operation completes normally; however, the data is discarded and no state change occurs.	RW	Yes	00b
2	***Reserved***	RsvdP	No	0
3	**No Soft Reset**	RO	Yes	1
7:4	***Reserved***	RsvdP	No	0h
8	**PME Enable** Tied to 0, because the PEX 8508 does not generate PME in PCI Express mode.	RsvdP	No	0
12:9	**Data Select** Writable by Serial EEPROM only[a]. Bits [12:9] select the **Data** and **Data Scale** registers. 0h = D0 power consumed 3h = D3hot power consumed 4h = D0 power dissipated 7h = D3hot power dissipated All other values are ***reserved***.	RO	Yes	0h
	Not supported RO for hardware auto-configuration.	RO	No	0h
14:13	**Data Scale** Loaded by serial EEPROM[a]. There are four internal **Data Scale** registers (one each per **Data** register – 0, 3, 4 and 7), per port. Bits [12:9], *Data Select*, select the **Data Scale** register.	RO	Yes	00b
15	**PME Status** 0 = PME is not being generated by the NT Port	RsvdP	No	0

Register 17-15. 44h Power Management Status and Control *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Power Management Control/Status Bridge Extensions				
21:16	***Reserved***	RsvdP	No	0-0h
22	**B2/B3 Support** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.	RsvdP	No	0
23	**Bus Power/Clock Control Enable** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.	RsvdP	No	0
Data				
31:24	**Data** Loaded by serial EEPROM[a]. There are four internal **Data** registers (0, 3, 4, and 7), per port. Bits [12:9], *Data Select*, select the **Data** register.	RO	Yes	00h

*a. With no serial EEPROM, Reads return 00h for the **Data Scale** and **Data** registers (for all Data Selects).*

17.6 Message Signaled Interrupt Capability Registers

The registers defined in Section 13.8, "Message Signaled Interrupt Capability Registers," are also applicable to the NT Port Virtual Interface. Table 17-4 defines the register map used by all ports.

Table 17-4. Message Signaled Interrupt Capability Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Reserved	Message Control	Next Capability Pointer (68h)	Capability ID (05h)	48h
Message Address[31:0]				4Ch
Message Upper Address[63:32]				50h
Reserved		Message Data		54h
Mask Bit for MSI				58h
Reserved				5Ch – 64h

17.7 PCI Express Capability Registers

This section details the PEX 8508 PCI Express Capability registers. Table 17-5 defines the register map.

Table 17-5. PCI Express Capability Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
PCI Express Capabilities		Next Capability Pointer (00h)	Capability ID (10h)	68h
Device Capabilities				6Ch
Device Status		Device Control		70h
Link Capabilities				74h
Link Status		Link Control		78h
Reserved				7Ch – 8Ch

Register 17-16. 68h PCI Express Capability List and Capabilities

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
PCI Express Capability List				
7:0	**Capability ID** Set to 10h by default, as required by the *PCI Express Base r1.1*.	RO	Yes	10h
15:8	**Next Capability Pointer** 00h = PCI Express Capability is the last capability in the first 256-byte Configuration space of the PEX 8508 NT Port Virtual Interface capability list The PEX 8508 NT Port Virtual Interface Extended Capabilities list starts at 100h.	RO	Yes	00h
PCI Express Capabilities				
19:16	**Capability Version** Set to 1h.	RO	Yes	1h
23:20	**Device/Port Type** Default = PCI Express Endpoint device.	RO	Yes	0h
24	**Slot Implemented** ***Not valid for PCI Express Endpoint devices.***	RsvdP	No	0
29:25	**Interrupt Message Number** The serial EEPROM writes 00_000b, because the Base message and MSI messages are the same.	RO	Yes	00_000b
31:30	***Reserved***	RsvdP	No	00b

Register 17-17. 6Ch Device Capabilities

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
2:0	**Maximum Payload Size Supported** 000b = NT Port Virtual Interface supports 128-byte maximum payload 001b = NT Port Virtual Interface supports 256-byte maximum payload No other values are supported.	RO	Yes	001b
4:3	**Phantom Functions Supported** ***Not supported*** Cleared to 00b.	RO	Yes	00b
5	**Extended Tag Field Supported** ***Not supported*** 0 = Maximum *Tag* field is 5 bits 1 = Maximum *Tag* field is 8 bits	RO	Yes	0
8:6	**Endpoint L0s Acceptable Latency**	RO	Yes	000b
11:9	**Endpoint L1 Acceptable Latency**	RO	Yes	000b
12	**Attention Button Present** No Attention Button is present for the NT Virtual Interface.	RsvdP	No	0
13	**Attention Indicator Present** No Attention Indicator is present for the NT Virtual Interface.	RsvdP	No	0
14	**Power Indicator Present** No Power Indicator is present for the NT Virtual Interface.	RsvdP	No	0
17:15	***Reserved***	RsvdP	No	000b
25:18	**Captured Slot Power Limit Value** For the NT Port Virtual Interface register, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Scale* field.	RO	Yes	00h
27:26	**Captured Slot Power Limit Scale** For the NT Port Virtual Interface register, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Value* field. 00b = 1.0 01b = 0.1 10b = 0.01 11b = 0.001	RO	Yes	00b
31:28	***Reserved***	RsvdP	No	0h

Register 17-18. 70h Device Status and Control

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Device Control			
0	**Correctable Error Reporting Enable** 0 = Disables 1 = Enables NT Port Virtual Interface to report Correctable errors to the local Host	RW	Yes	0
1	**Non-Fatal Error Reporting Enable** 0 = Disables 1 = Enables NT Port Virtual Interface to report Non-Fatal errors to the local Host	RW	Yes	0
2	**Fatal Error Reporting Enable** 0 = Disables 1 = Enables NT Port Virtual Interface to report Fatal errors to the local Host	RW	Yes	0
3	**Unsupported Request Reporting Enable** 0 = Disables 1 = Enables NT Port Virtual Interface to report Unsupported Request errors to the local Host	RW	Yes	0
4	**Enable Relaxed Ordering** ***Not supported*** Cleared to 0.	RsvdP	No	0
7:5	**Maximum Payload Size** The NT Port Virtual Interface power-on/reset value is 000b, to support a Maximum Payload Size of 128 bytes. Software can change this field to configure the NT Port Virtual Interface to support other Payload sizes; however, software cannot change this field to a value larger than that indicated by the **Device Capabilities** register *Maximum Payload Size Supported* field (offset 6Ch[2:0]). for the Virtual and Link Interfaces. (Requester and Completer domains must possess the same Maximum Payload Size.) 000b = Indicates that initially the PEX 8508 port is configured to support a Maximum Payload Size of 128 bytes 001b = Indicates that initially the PEX 8508 port is configured to support a Maximum Payload Size of 256 bytes No other values are supported. ***Note:*** *Software must halt all transactions through the NT Port before changing this field.*	RW	Yes	000b
8	**Extended Tag Field Enable** ***Not supported*** Cleared to 0.	RsvdP	No	0
9	**Phantom Functions Enable** ***Not supported*** Cleared to 0.	RsvdP	No	0
10	**Auxiliary (AUX) Power PM Enable** Cleared to 0.	RO	No	0
11	**Enable No Snoop** ***Not supported*** Cleared to 0.	RsvdP	No	0

Register 17-18. 70h Device Status and Control *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
14:12	**Maximum Read Request Size** ***Not supported*** Cleared to 000b.	RsvdP	No	000b
15	***Reserved***	RsvdP	No	0
	Device Status			
16	**Correctable Error Detected** Set when the NT Port Virtual Interface detects a Correctable error, regardless of the bit 0 (*Correctable Error Reporting Enable* bit) state. 0 = NT Port Virtual Interface did not detect a Correctable error 1 = NT Port Virtual Interface detected a Correctable error	RW1C	Yes	0
17	**Non-Fatal Error Detected** Set when the NT Port Virtual Interface detects a Non-Fatal error, regardless of the bit 1 (*Non-Fatal Error Reporting Enable* bit) state. 0 = NT Port Virtual Interface did not detect a Non-Fatal error 1 = NT Port Virtual Interface detected a Non-Fatal error	RW1C	Yes	0
18	**Fatal Error Detected** Set when the NT Port Virtual Interface detects a Fatal error, regardless of the bit 2 (*Fatal Error Reporting Enable* bit) state. 0 = NT Port Virtual Interface did not detect a Fatal error 1 = NT Port Virtual Interface detected a Fatal error	RW1C	Yes	0
19	**Unsupported Request Detected** Set when the NT Port Virtual Interface detects an Unsupported Request, regardless of the bit 3 (*Unsupported Request Reporting Enable* bit) state. 0 = NT Port Virtual Interface did not detect an Unsupported Request 1 = NT Port Virtual Interface detected an Unsupported Request	RW1C	Yes	0
20	**Auxiliary (AUX) Power Detected** ***Not supported*** Cleared to 0.	RO	No	0
21	**Transactions Pending** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
31:22	***Reserved***	RsvdP	No	000h

Register 17-19. 74h Link Capabilities

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	**Maximum Link Speed** Set to 0001b for 2.5 Gbps.	RO	Yes	0001b
9:4	**Maximum Link Width** Actual link width is set by STRAP_PORTCFG[4:0]. The PEX 8508 Maximum Link Width is x4 = 00_0100b.	RO	No	Set by Strapping ball levels
11:10	**Active State Power Management (ASPM) Support** Indicates the level of ASPM supported by the port. 01b = L0s link power state entry is supported All other values are ***reserved***.	RO	Yes	01b
14:12	**L0s Exit Latency** 101b = NT Port Virtual Interface L0s Exit Latency is between 1 and 2 µs	RO	No	101b
17:15	**L1 Exit Latency** 101b = NT Port Virtual Interface L1 Exit Latency is between 16 and 32 µs	RO	Yes	101b
23:18	***Reserved***	RsvdP	No	0-0h
31:24	**Port Number** The NT Port Number is selected by signal ball strapping options. Refer to STRAP_NT_UPSTRM_PORTSEL[3:0] for details.	HwInit	No	Set by Strapping ball levels

Register 17-20. 78h Link Status and Control

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Link Control			
1:0	**Active State Power Management (ASPM) Control** Not applicable to the NT Port Virtual Interface, because no external port connection exists.	RW	Yes	00b
2	***Reserved***	RsvdP	No	0
3	**Read Completion Boundary (RCB)** ***Not supported*** Cleared to 0.	RO	Yes	0
4	**Link Disable** ***Reserved*** for NT Port Virtual Interface.	RsvdP	No	0
5	**Retrain Link** ***Reserved*** for NT Port Virtual Interface.	RsvdP	No	0
6	**Common Clock Configuration** Not applicable to the NT Port Virtual Interface, because no external port connection exists.	RW	Yes	0
7	**Extended SYNC** Not applicable to the NT Port Virtual Interface, because no external port connection exists.	RW	Yes	0
15:8	***Reserved***	RsvdP	No	00h
	Link Status			
19:16	**Link Speed** The NT Port Virtual Interface is set to 0001b for 2.5 Gbps.	RO	Yes	0001b
25:20	**Negotiated Link Width** Not applicable to the NT Port Virtual Interface, because no external port connection exists.	RO	Yes	00_0000b
26	**Training Error** ***Reserved*** for NT Port Virtual Interface.	RsvdP	No	0
27	**Link Training** ***Reserved*** for NT Port Virtual Interface.	RsvdP	No	0
28	**Slot Clock Configuration** Because there is no external connection to the NT Port Virtual Interface, this bit is always cleared to 0, which indicates that the PEX 8508 uses an independent clock.	HwInit	Yes	0
31:29	***Reserved***	RsvdP	No	000b

17.8 NT Port Virtual Interface Registers

Table 17-6. NT Port Virtual Interface Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
NT Port Virtual Interface IRQ Doorbell Registers		90h … ACh
NT Port Scratchpad (Mailbox) Registers		B4h … CCh
NT Port Virtual Interface BAR Setup Registers		D0h … F4h
NT Port Virtual Interface Cursor Mechanism Control Registers		F8h FCh

17.8.1 NT Port Virtual Interface IRQ Doorbell Registers

This section details the PEX 8508 NT Port Virtual Interface Interrupt Request (IRQ) Doorbell registers. Table 17-7 defines the register map.

Table 17-7. NT Port Virtual Interface Interrupt Request (IRQ) Doorbell Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Reserved	Set Virtual Interface IRQ	90h
Reserved	Clear Virtual Interface IRQ	94h
Reserved	Set Virtual Interface IRQ Mask	98h
Reserved	Clear Virtual Interface IRQ Mask	9Ch
Reserved	Set Link Interface IRQ	A0h
Reserved	Clear Link Interface IRQ	A4h
Reserved	Set Link Interface IRQ Mask	A8h
Reserved	Clear Link Interface IRQ Mask	ACh

Register 17-21. 90h Set Virtual Interface IRQ

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Set Virtual IRQ** Controls the state of the Virtual Interface Doorbell Interrupt request. Reading returns the status of the bits. Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register sets the corresponding Interrupt request. The Virtual Interface interrupt is asserted if the following conditions exist: • This register (offset 90h or 94h) value is non-zero, and, • The corresponding mask bit in the **Set Virtual Interface IRQ Mask** or **Clear Virtual Interface IRQ Mask** register (offset 98h or 9Ch, respectively) is not set, and, • Interrupts (either INT*x* or MSI) are enabled	RW1S	Yes	0000h
31:16	***Reserved***	RsvdP	No	0000h

Register 17-22. 94h Clear Virtual Interface IRQ

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Clear Virtual IRQ** Controls the state of the Virtual Interface Doorbell Interrupt request. Reading returns the status of the bits. Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register clears the corresponding Interrupt request. The Virtual Interface interrupt is asserted if the following conditions exist: • This register (offset 94h or 90h) value is non-zero, and, • The corresponding mask bit in the **Set Virtual Interface IRQ Mask** or **Clear Virtual Interface IRQ Mask** register (offset 98h or 9Ch, respectively) is not set, and, • Interrupts (either INT*x* or MSI) are enabled	RW1C	Yes	0000h
31:16	***Reserved***	RsvdP	No	0000h

Register 17-23. 98h Set Virtual Interface IRQ Mask

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Set Virtual IRQ Mask** Virtual Interface interrupt IRQ Mask Set. Reading returns the state of the mask bits. 0 = Corresponding interrupt request bit in the **Set Virtual Interface IRQ** register (offset 90h) is unmasked 1 = Corresponding interrupt request bit in the **Set Virtual Interface IRQ** register (offset 90h) is masked/disabled Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register clears the corresponding interrupt mask bit.	RW1S	Yes	FFFFh
31:16	***Reserved***	RsvdP	No	0000h

Register 17-24. 9Ch Clear Virtual Interface IRQ Mask

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Clear Virtual IRQ Mask** Controls the state of the Virtual Interface Interrupt Request bits. Reading returns the status of the bits. 0 = Corresponding interrupt request bit in the **Clear Virtual Interface IRQ** register (offset 94h) is unmasked 1 = Corresponding interrupt request bit in the **Clear Virtual Interface IRQ** register (offset 94h) is masked/disabled Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register clears the corresponding interrupt mask bit.	RW1C	Yes	FFFFh
31:16	***Reserved***	RsvdP	No	0000h

Register 17-25. A0h Set Link Interface IRQ

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Set Link IRQ** Controls the state of the Link Interface Doorbell Interrupt request. Reading returns the status of the bits. Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register sets the corresponding Interrupt request. The Link Interface interrupt is asserted if the following conditions exist: • This register (offset A0h or A4h) value is non-zero, and, • The corresponding mask bit in the **Set Link Interface IRQ Mask** or **Clear Link Interface IRQ Mask** register (offset A8h or ACh, respectively) is not set, and, • Interrupts (either INT*x* or MSI) are enabled	RW1S	Yes	0000h
31:16	***Reserved***	RsvdP	No	0000h

Register 17-26. A4h Clear Link Interface IRQ

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Clear Link IRQ** Controls the state of the Link Interface Doorbell Interrupt request. Reading returns the status of the bits. Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register clears the corresponding Interrupt request. The Link Interface interrupt is asserted if the following conditions exist: • This register (offset A4h or A0h) value is non-zero, and, • The corresponding mask bit in the **Set Link Interface IRQ Mask** or **Clear Link Interface IRQ Mask** register (offset A8h or ACh, respectively) is not set, and, • Interrupts (either INT*x* or MSI) are enabled	RW1C	Yes	0000h
31:16	***Reserved***	RsvdP	No	0000h

Register 17-27. A8h Set Link Interface IRQ Mask

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Set Link IRQ Mask** Link Interface Interrupt IRQ Mask Set. Reading returns the state of the IRQ mask. 0 = Corresponding interrupt request bit in the **Set Link Interface IRQ** register (offset A0h) is unmasked 1 = Corresponding interrupt request bit in the **Set Link Interface IRQ** register (offset A0h) is masked/disabled Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register sets the corresponding interrupt mask bit.	RW1S	Yes	FFFFh
31:16	***Reserved***	RsvdP	No	0000h

Register 17-28. ACh Clear Link Interface IRQ Mask

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Clear Link IRQ Mask** Link Interface Interrupt IRQ Mask Clear. Reading returns the state of the IRQ mask. 0 = Corresponding interrupt request bit in the **Clear Link Interface IRQ** register (offset A4h) is unmasked 1 = Corresponding interrupt request bit in the **Clear Link Interface IRQ** register (offset A4h) is masked/disabled Writing 0 to a bit in the register has no effect. Writing 1 to a bit in the register clears the corresponding interrupt mask bit.	RW1C	Yes	FFFFh
31:16	***Reserved***	RsvdP	No	0000h

17.8.2 NT Port Scratchpad (Mailbox) Registers

The PEX 8508 NT Port Scratchpad (Mailbox) registers are defined in Section 18.8.2, "NT Port Scratchpad (Mailbox) Registers." Table 17-8 defines the register map used by the NT Port Virtual Interface.

Table 17-8. PEX 8508 NT Port Scratchpad (Mailbox) Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
NT Port Scratchpad_0				B0h
NT Port Scratchpad_1				B4h
NT Port Scratchpad_2				B8h
NT Port Scratchpad_3				BCh
NT Port Scratchpad_4				C0h
NT Port Scratchpad_5				C4h
NT Port Scratchpad_6				C8h
NT Port Scratchpad_7				CCh

17.8.3 NT Port Virtual Interface BAR Setup Registers

This section details the NT Port Virtual Interface BAR Setup registers. Table 17-9 defines the register map used by the NT Port Virtual Interface.

The **NT Port Virtual Interface BAR*x* Setup** (offsets D0h through E0h) register values are shadowed in the corresponding **NT Port Virtual Interface BAR*x* Setup Shadow** registers (offsets D80h through D90h, respectively). When software writes to an **NT Port Virtual Interface BAR*x* Setup** register, the value is automatically copied to the corresponding **NT Port Virtual Interface BAR*x* Setup (Shadow Copy)** register. If the **NT Port Virtual Interface BAR*x* Setup** registers are programmed by serial EEPROM, the **NT Port Virtual Interface BAR*x* Setup (Shadow Copy)** registers must also be programmed by serial EEPROM, to the same respective register values.

Table 17-9. PEX 8508 NT Port Virtual Interface BAR Setup Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
NT Port Virtual Interface BAR1 Setup				D0h
NT Port Virtual Interface Memory BAR2 Setup				D4h
NT Port Virtual Interface Memory BAR3 Setup				D8h
NT Port Virtual Interface Memory BAR4/5 Setup				DCh
NT Port Virtual Interface Memory BAR5 Setup				E0h
Reserved				E4h – F4h

Register 17-29. D0h NT Port Virtual Interface BAR1 Setup

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
1:0	**I/O BAR1 Enable** 11b = Enables Virtual Interface **BAR1** as an I/O BAR All other values disable **BAR1**.	RW	Yes	11b
31:2	***Reserved***	RsvdP	No	0-0h

Register 17-30. D4h NT Port Virtual Interface Memory BAR2 Setup

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	***Reserved***	RsvdP	No	0
2:1	**BAR2 Type** 00b = Selects 32-bit Memory BAR No other values are allowed.	RW	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
30:12	**BAR2 Size** Specifies the Address Range size requested by **BAR2**. 0 = Corresponding bits in **BAR2** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR2** are RW bits	RW	Yes	0000_0h
31	**BAR2 Enable** 0 = **BAR2** is disabled, all bits in **BAR2** read 0 1 = **BAR2** is enabled, Size and Type specified in this register	RW	Yes	0

Register 17-31. D8h NT Port Virtual Interface Memory BAR3 Setup

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
15:1	***Reserved***	RsvdP	No	0-0h
19:16	**BAR3 LUT Page Size** Selects the page size of the Virtual Interface **BAR3** Address range. The total size of this range is dependent on the page size. When the *BAR3 LUT Page Size Extension* bit value is 0, the encodings are as follows: 0h = Disables **BAR3** 1h to 4h = ***Reserved*** 5h = 4 KB 6h = 8 KB 7h = 16 KB 8h = 32 KB 9h = 64 KB Ah = 128 KB Bh = 256 KB Ch = 512 KB Dh = 1 MB Eh = 2 MB Fh = 4 MB When the *BAR3 LUT Page Size Extension* bit value is 1, the encodings are as follows: 0h = 8 MB 1h = 16 MB 2h = 32 MB 3h to Fh = ***Reserved***	RW	Yes	0h
20	**BAR3 LUT Page Size Extension** Allows selection of larger page sizes when programming Page Size[19:16]. 0 = Page sizes 4 KB through 4 MB are available in the Page Size[19:16] 1 = Page sizes 8 through 32 MB are available in the Page Size[19:16]	RW	Yes	0
31:21	***Reserved***	RsvdP	0	000h

Register 17-32. DCh NT Port Virtual Interface Memory BAR4/5 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
2:1	**BAR4 Type** 00b = **BAR4** is implemented as a 32-bit Memory BAR 10b = **BAR4/5** is implemented as a 64-bit Memory BAR ***Note:*** *It is illegal to program 10b and clear the* ***NT Port Virtual Interface Memory BAR5 Setup*** *register BAR5 Enable bit (offset E0h[31]).*	RW	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
30:12	**BAR4 Size** Specifies the Address Range size requested by **BAR4**. 0 = Corresponding bits in **BAR4** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR4** are RW bits	RW	Yes	0000_0h
31	**BAR4 Enable** When bits [2:1] = 00b, enables **BAR4**; otherwise, belongs to the *BAR4 Size*[30:12] field. 0 = **BAR4** is disabled, all bits in **BAR4** read 0 1 = **BAR4** is enabled, Size and Type specified in this register	RW	Yes	0

Register 17-33. E0h NT Port Virtual Interface Memory BAR5 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
30:0	**BAR5 Size** Together with the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Size* field (offset DCh[31:12]), specifies the Address Range size requested by **BAR4/5** in 64-bit mode when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is set to 10b. 0 = Corresponding bits in **BAR5** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR5** are RW bits	RW	Yes	0-0h
	Reserved when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is cleared to 00b.	RsvdP	No	0-0h
31	**BAR5 Enable** 0 = **BAR5** is disabled 1 = **BAR5** is enabled when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is set to 10b ***Note:*** *It is illegal to program the* ***NT Port Virtual Interface Memory BAR4/5 Setup*** *register BAR4 Type* field *(offset DCh[2:1]) to 10b and clear this bit.*	RW	Yes	0
	Reserved when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is cleared to 00b.	RsvdP	No	0

17.8.4 NT Port Virtual Interface Cursor Mechanism Control Registers

This section details the NT Port Virtual Interface Cursor Mechanism Control registers. Table 17-10 defines the register map.

The Cursor Mechanism registers at offsets F8h/FCh provide a means for accessing PCI Express Extended Configuration Space registers (offsets 100h through FFFh) within the NT Port Virtual and Link Interfaces, when only standard PCI Configuration transactions (that do not support Extended Register Number), and/or I/O Request transactions (using the NT Port **BAR1** address, if enabled) are available. The Cursor Mechanism can generally access only those registers that are defined by the *PCI Express Base r1.1*, and not the device-specific registers. However if Port 0 is the NT Port, the Cursor Mechanism in the NT Port Virtual Interface registers can also access the device-specific registers.

Table 17-10. NT Port Virtual Interface Cursor Mechanism Control Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Configuration Address Window		***Reserved***		F8h
Configuration Data Window				FCh

Register 17-34. F8h Configuration Address Window

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	***Reserved***	RsvdP	No	0000h
25:16	**Register Offset**	RW	Yes	000h
30:26	***Reserved***	RsvdP	No	00h
31	**Interface Select** 0 = Access to NT Port Link Interface Type 0 Configuration Space register 1 = Access to NT Port Virtual Interface Type 0 Configuration Space register	RW	Yes	0

Register 17-35. FCh Configuration Data Window

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Data Window** Software selects a register by writing into the NT Port Virtual Interface Configuration Address window, then reads from or writes to that register using this register.	RW	Yes	0-0h

17.9 Device Serial Number Extended Capability Registers

The registers defined in Section 13.10, "Device Serial Extended Number Capability Registers," are also applicable to the NT Port Virtual Interface. Table 17-11 defines the register map used by all ports.

Table 17-11. PEX 8508 Device Serial Number Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (FB4h) / Capability Version (1h)	PCI Express Extended Capability ID (0003h)	100h
Serial Number (Lower DW)		104h
Serial Number (Upper DW)		108h

17.10 Power Budgeting Extended Capability Registers

The registers defined in Section 13.11, "Power Budgeting Extended Capability Registers," are also applicable to the NT Port Virtual Interface. Table 17-12 defines the register map used by all ports.

Table 17-12. PEX 8508 Power Budgeting Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (148h) / Capability Version (1h)	PCI Express Extended Capability ID (0004h)	138h
Reserved	Data Select	13Ch
Power Budgeting Data		140h
Power Budget Capability		144h

17.11 Virtual Channel Extended Capability Registers

The registers defined in Section 13.12, "Virtual Channel Extended Capability Registers," are also applicable to the NT Port Virtual Interface. Table 17-13 defines the register map used by all ports.

Table 17-13. PEX 8508 Virtual Channel Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16		15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (000h)	Capability Version (1h)	PCI Express Extended Capability ID (0002h)	148h
Port VC Capability 1			14Ch
Port VC Capability 2			150h
Port VC Status ***(Reserved)***		Port VC Control	154h
VC0 Resource Capability			158h
VC0 Resource Control			15Ch
VC0 Resource Status		***Reserved***	160h
VC1 Resource Capability			164h
VC1 Resource Control			168h
VC1 Resource Status		***Reserved***	16Ch
Reserved			170h – 1C4h

17.12 Device-Specific Registers

The registers defined in Section 13.13, "Device-Specific Registers," are unique to the PEX 8508 device and not referenced in the *PCI Express Base r1.1*. These registers are also applicable to the NT Port Virtual Interface, except as defined in Table 17-14 through Table 17-16 and Register 17-36 through Register 17-38. Table 17-14 defines the register map used by the NT Port Virtual Interface.

Note: *This register group is accessed using a Memory-Mapped cycle.* ***It is recommended that these register values not be changed.***

Table 17-14. Device-Specific NT Port Virtual Interface Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Reserved	1C8h – 2C4h
Bus Number CAM Registers	2C8h … 304h
I/O CAM Registers	308h … 344h
Address-Mapping CAM (AMCAM) Registers	348h … 548h
Reserved	54Ch – 65Ch
NT Port Virtual Interface Ingress Control Register	660h … 67Ch
I/O CAM Base and Limit Upper 16-Bit Registers	680h … 6BCh
Base Address Shadow Registers (BARs)	6C0h … 73Ch
Shadow Virtual Channel (VC) Capability Registers	740h … 83Ch
Shadow Port Virtual Channel (VC) Capability Registers	840h … 9ECh

Table 17-14. Device-Specific NT Port Virtual Interface Register Map *(Cont.)*

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Ingress Credit Handler (INCH) Registers				9F0h … BDCh
Reserved				BE0h – BFCh
Internal Credit Handler (ITCH) VC&T Threshold Registers				C00h … C08h
NT Port Virtual Interface Virtual Channel Queue Status Registers				C0Ch … C38h

17.12.1 NT Port Virtual Interface Ingress Control Register

The NT Port Virtual Interface **Ingress Control** register is defined in Section 13.13.7, "Ingress Control and Port Enable Registers," with the addition of the *No Snoop Disable* bit. Table 17-15 defines the register map.

Table 17-15. Device-Specific NT Port Virtual Interface Ingress Control Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Ingress Control				660h
Reserved				664h – 67Ch

Register 17-36. 660h Ingress Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Enable CSR Access by Downstream Devices** Enables acceptance of both Configuration and Memory requests from a Requester that is downstream from a Transparent port, targeting any PEX 8508 registers. 0 = Configuration requests from a downstream device that is targeting PEX 8508 registers are ***not supported***; the downstream port flags an Uncorrectable Error, and, returns a Completion with Unsupported Request (UR) specified in the *Completion Status* field, to the downstream Requester. Only this mode is *PCI Express Base r1.1*-compliant. 1 = Configuration and Memory requests from downstream Requesters, targeting any PEX 8508 registers in any port, are allowed. ***Notes:*** *This bit can be initially set only through the upstream port, the NT Port Link Interface, the I^2C interface, or by the serial EEPROM, to enable register access through downstream Transparent ports; a Requester downstream from a Transparent port cannot set the bit to grant itself (or peers) access to PEX 8508 registers. Configuration requests can access those registers that are defined by PCI-SIG specifications, and generally cannot access device-specific registers other than the NT Port* ***Cursor Mechanism*** *registers (offsets F8h/FCh). Memory requests can access all PEX 8508 registers.* *Configuration requests can access the NT Port to provide indirect access to NT Port offsets above 100h, for Conventional PCI Requesters such as a PCI Master connected to a PCI Express-to-PCI bridge, that cannot generate Configuration requests containing an Extended Register Number. The NT Port Virtual Interface Cursor Mechanism (but not the NT Port Link Interface Cursor Mechanism) can access device-specific registers that exist in the NT Port; if Port 0 is the NT Port, the NT Port Virtual Interface Cursor Mechanism can access the device-specific registers (which exist only in Port 0) and the Station registers (which exist in Port 0 for all enabled ports).*	RW	Yes	0
1	**Disable Unsupported Request Response for *Reserved* Configuration Registers**	RW	Yes	0
23:2	**Ingress Control** ***Factory Test Only***	RW	Yes	0-0h
24	**No Snoop Disable** Forces the packet header No Snoop attribute bit to 0, for all packets transferred between the NT Port Link and Virtual Interfaces (across the NT boundary, in both directions). Can be used to handle cache coherency-related issues in a system. 0 = Disables the No Snoop Disable feature 1 = Enables the No Snoop Disable feature	RW	Yes	0
31:25	**Ingress Control** ***Factory Test Only***	RW	Yes	0-0h

17.12.2 NT Port Virtual Interface Virtual Channel Queue Status Registers

The registers defined in Section 13.13.15, "Port Virtual Channel Queue Status Registers," are also applicable to the NT Port Virtual Interface, except as defined in Register 17-37 and Register 17-38. Table 17-16 defines the register map used by the NT Port Virtual Interface.

Table 17-16. Device-Specific Port Virtual Channel Queue Status Register Map (Only Port 0 and NT Link Interface)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Port 0 VC0 Posted and Non-Posted Queue Status	C0Ch
Port 0 VC0 Completion and VC1 Posted Queue Status	C10h
Port 0 VC1 Non-Posted and Completion Queue Status	C14h
Port 1 VC0 Posted and Non-Posted Queue Status	C18h
Port 1 VC0 Completion and VC1 Posted Queue Status	C1Ch
Port 1 VC1 Non-Posted and Completion Queue Status	C20h
Port 2 VC0 Posted and Non-Posted Queue Status	C24h
Port 2 VC0 Completion and VC1 Posted Queue Status	C28h
Port 2 VC1 Non-Posted and Completion Queue Status	C2Ch
Port 3 VC0 Posted and Non-Posted Queue Status	C30h
Port 3 VC0 Completion Queue Status	C34h
Reserved	C38h

Register 17-37. C1Ch Port 1 VC0 Completion and VC1 Posted Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 1 VC0 Completion Packets Status**	RO	Yes	000h
21:11	**Port 1 VC1 Posted Packets Status**	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 1 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC1 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 17-38. C20h Port 1 VC1 Non-Posted and Completion Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 1 VC1 Non-Posted Packets Status**	RO	Yes	000h
21:11	**Port 1 VC1 Completion Packets Status**	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 1 ITCH Disabled VC1 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC1 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h

17.13 Non-Transparent Bridging-Specific Registers

Table 17-17 defines the register map of the registers implemented to support the PEX 8508 NT Port Virtual Interface Non-Transparent Bridging-Specific registers. These registers are accessed by Memory-Mapped access to Port 0 or the NT Port.

Table 17-17. NT Port Virtual Interface NT Bridging-Specific Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
NT Port Virtual Interface Memory Address Translation and BAR Limit Registers				C3Ch … C58h
NT Port Virtual Interface Lookup Table-Based Address Translation Registers				C5Ch … D58h
NT Port Link Interface VC Shadow Registers (Shadow Copy)				D5Ch … D64h
NT Port Virtual Interface Base Address Registers (BARs) and BAR Setup Registers (Shadow Copy)				D68h … D90h
NT Port Virtual Interface Send Lookup Table Entry Registers				D94h … DB0h
Reserved				DB4h – DF0h

17.13.1 NT Port Virtual Interface Memory Address Translation and BAR Limit Registers

Table 17-18 defines the register map of the registers implemented to support the PEX 8508 NT Port Virtual Interface Memory Address Translation and BAR Limit registers. These registers are accessed by Memory-Mapped access to Port 0 or the NT Port.

Table 17-18. NT Port Virtual Interface Memory Address Translation and BAR Limit Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Memory BAR2 Address Translation		C3Ch
Reserved		C40h
Memory BAR4/5 Address Translation Lower		C44h
Memory BAR4/5 Address Translation Upper		C48h
Memory BAR2 Limit Address		C4Ch
Reserved		C50h
Memory BAR4/5 Limit Lower Address		C54h
Memory BAR4/5 Limit Upper Address		C58h

Register 17-39. C3Ch Memory BAR2 Address Translation

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR2 Base Translation Address[31:12]** NT Port Virtual Interface Base Translation address when the **NT Port Virtual Interface Memory BAR2 Setup** register *BAR2 Enable* bit (offset D4h[31]) is set to 1.	RW	Yes	0-0h

Register 17-40. C44h Memory BAR4/5 Address Translation Lower

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR4 Base Translation Address[31:12]** NT Port Virtual Interface Base Translation address when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Enable* bit (offset DCh[31]) is set to 1.	RW	Yes	0-0h

Register 17-41. C48h Memory BAR4/5 Address Translation Upper

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**BAR4/5 Base Translation Address[63:32]** NT Port Virtual Interface Base Translation upper address when **BAR4/5** is enabled as a 64-bit BAR [**NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* (offset DCh[2:1]) and **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* (offset D8Ch[2:1]) fields are both set to 10b].	RW	Yes	0-0h
	Read-Only when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* (offset DCh[2:1]) and **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* (offset D8Ch[2:1]) fields are both cleared to 00b.	RO	No	0-0h

Register 17-42. C4Ch Memory BAR2 Limit Address

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR2 Limit[31:0]** Contains the address of the memory window upper limit defined in the **Memory BAR2 Setup (Shadow Copy)** register (offset D84h). 1 MB granularity. When the limit is greater than the window size, the limit is ignored.	RW	Yes	000h

Register 17-43. C54h Memory BAR4/5 Limit Lower Address

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR4/5 Limit[31:0]** Contains the address of the memory window lower limit defined in the **Memory BAR4 Setup (Shadow Copy)** register (offset D8Ch). 1 MB granularity. When the limit is greater than the window size, the limit is ignored.	RW	Yes	0-0h

Register 17-44. C58h Memory BAR4/5 Limit Upper Address

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
31:0	**BAR5 Limit[63:32]** Contains the address of the memory window upper limit defined in the **Memory BAR5 Setup (Shadow Copy)** register (offset D90h) when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* (offset DCh[2:1]) and **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* (offset D8Ch[2:1]) fields are both set to 10b. 1 MB granularity. When the limit is greater than the window size, the limit is ignored.	RW	Yes	0-0h
	Read-Only when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* (offset DCh[2:1]) and **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* (offset D8Ch[2:1]) fields are both cleared to 00b.	RO	No	0-0h

17.13.2 NT Port Virtual Interface Lookup Table-Based Address Translation Registers

There are 64 **Base-Translation Lookup Table (LUT) Entry** registers to support the LUT-based address translation. These registers are accessed using Port 0 and the NT Port Virtual Interface Memory-Mapped or Cursor mechanism. Table 17-19 defines the register and address locations. Register 17-45 defines the bit definitions that apply to all 64 registers.

Table 17-19. Base-Translation Lookup Table Entry_*n* Register Locations

ADDR Location	Lookup Table Entry_*n*	ADDR Location	Lookup Table Entry_*n*	ADDR Location	Lookup Table Entry_*n*	ADDR Location	Lookup Table Entry_*n*
C5Ch	0	C60h	1	C64h	2	C68h	3
C6Ch	4	C70h	5	C74h	6	C78h	7
C7Ch	8	C80h	9	C84h	10	C88h	11
C8Ch	12	C90h	13	C94h	14	C98h	15
C9Ch	17	CA0h	17	CA4h	18	CA8h	19
CACh	20	CB0h	21	CB4h	22	CB8h	23
CBCh	24	CC0h	25	CC4h	26	CC8h	27
CCCh	28	CD0h	29	CD4h	30	CD8h	31
CDCh	32	CE0h	33	CE4h	34	CE8h	35
CECh	36	CF0h	37	CF4h	38	CF8h	39
CFCh	40	D00h	41	D04h	42	D08h	43
D0Ch	44	D10h	45	D14h	46	D18h	47
D1Ch	48	D20h	49	D24h	50	D28h	51
D2Ch	52	D30h	53	D34h	54	D38h	55
D3Ch	56	D40h	57	D44h	58	D48h	59
D4Ch	60	D50h	61	D54h	62	D58h	63

Register 17-45. C5Ch - D58h Base-Translation Lookup Table Entry_*n* (where *n* = 0 through 63)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Entry Status** 0 = Invalid 1 = Valid	RW	Yes	0
2:1	***Reserved***	RsvdP	No	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Translation** Base Translation address value.	RW	Yes	0-0h

17.13.3 NT Port Link Interface VC Shadow Registers

The registers in this section are shadow copies and valid only for Port 0. If Port 0 is the NT Port, the register is in Virtual Interface Configuration space. ***It is recommended that these register values not be changed.*** Table 17-20 defines the register map.

Table 17-20. NT Port Link Interface VC Shadow Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
NT Link Interface VC0 Resource Control (Shadow Copy)				D5Ch
NT Link Interface VC1 Resource Control (Shadow Copy)				D60h
NT Link Interface VC Capability 1 (Shadow Copy)				D64h

Register 17-46. D5Ch NT Link Interface VC0 Resource Control (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**TC VC0 Map[0]** Always mapped to Virtual Channel 0.	RO	Yes	1
7:1	**TC VC0 Map[7:1]** Mapped to Virtual Channel 0 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	7Fh
23:8	***Reserved***	RsvdP	No	0000h
24	**VC0 ID** Virtual Channel 0 identification number.	RO	Yes	0
30:25	***Reserved***	RsvdP	No	00h
31	**VC0 Enable** Virtual Channel 0 enable.	RO	Yes	1

Register 17-47. D60h NT Link Interface VC1 Resource Control (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
7:1	**TC VC1 Map[7:1]** Mapped to Virtual Channel 1 by default. Software can change this field during enumeration or when assigning the traffic to the traffic class.	RW	Yes	00h
23:8	***Reserved***	RsvdP	No	0000h
24	**VC1 ID** Virtual Channel 1 identification number.	RW	Yes	1
30:25	***Reserved***	RsvdP	No	00h
31	**VC1 Enable** Virtual Channel 1 enable.	RW	Yes	0

Register 17-48. D64h NT Link Interface VC Capability 1 (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
6:4	**Low-Priority Virtual Channel Count** Indicates the number of Virtual Channels mapped to the low-priority group.	RO	Yes	000b
31:7	***Reserved***	RsvdP	No	0000_00h

17.13.4 NT Port Virtual Interface Base Address Registers (BARs) and BAR Setup Registers

The registers in this section are shadow copies and valid only for Port 0. If Port 0 is the NT Port, the register is in Virtual Interface Configuration space. ***It is recommended that these register values not be changed.*** Table 17-21 defines the register map.

Table 17-21. NT Port Virtual Interface Base Address Register (BAR) and BAR Setup Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
BAR0 = ***Reserved***	D68h
BAR1 (Shadow Copy)	D6Ch
BAR2 (Shadow Copy)	D70h
BAR3 (Shadow Copy)	D74h
BAR4 (Shadow Copy)	D78h
BAR5 (Shadow Copy)	D7Ch
BAR1 Setup (Shadow Copy)	D80h
Memory BAR2 Setup (Shadow Copy)	D84h
Memory BAR3 Setup (Shadow Copy)	D88h
Memory BAR4 Setup (Shadow Copy)	D8Ch
Memory BAR5 Setup (Shadow Copy)	D90h

Register 17-49. D6Ch BAR1 (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**I/O Space Indicator** I/O BAR when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is set to 11b.	RO	Yes	1
	Reserved when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is not set to 11b.	RsvdP	No	0
7:1	***Reserved***	RsvdP	No	0000_000b
31:8	**I/O Base Address** 256-byte I/O Space Base address when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is set to 11b.	RW	Yes	0000_00h
	Reserved when the **NT Port Virtual Interface BAR1 Setup** register *I/O BAR1 Enable* field (offset D0h[1:0]) is not set to 11b.	RsvdP	No	00000_00h

Register 17-50. D70h BAR2 (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR – only value supported	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 01b, 10b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Address 2** Contains the software-assigned Memory Space Base Address: • Enabled and sized by the **Memory BAR2 Setup (Shadow Copy)** register (offset D84h) • Used for Memory transactions crossing the NT Port • Minimum Address range requested is 4 KB • Uses direct address translation	RW	Yes	0000_0h

Register 17-51. D74h BAR3 (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR – only value supported	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 01b, 10b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
17:4	***Reserved***	RsvdP	No	0000h
31:18	**Base Address 3** Contains the software-assigned Memory Space Base Address: • Enabled and sized by **Memory BAR3 Setup (Shadow Copy)** register (offset D88h) • No **Limit** register • Used for Memory transactions crossing the NT Port • Minimum Address range requested is 256 KB • Uses LUT address translation	RW	Yes	0000h

Register 17-52. D78h BAR4 (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Memory BAR – only value supported	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 10b = Mappable anywhere in 64-bit Memory space 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Address 4** Contains the software-assigned Memory Space Base Address: • Enabled and sized by **Memory BAR4 Setup (Shadow Copy)** register (offset D8Ch) • Used for Memory transactions crossing the NT Port • Minimum Address range requested is 4 KB • Uses direct address translation	RW	Yes	0000_0h

Register 17-53. D7Ch BAR5 (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 5** NT Port Virtual Interface upper 32-bit address when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is set to 10b. RW, based upon **Memory BAR5 Setup (Shadow Copy)** register (offset D90h). The **BAR4/5** group uses direct address translation.	RW	Yes	0-0h
	Reserved when the **NT Port Virtual Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset DCh[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Register 17-54. D80h BAR1 Setup (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
1:0	**I/O BAR1 Enable** 11b = Enables Virtual Interface **BAR1** as an I/O BAR All other values disable **BAR1**.	RW	Yes	11b
31:2	***Reserved***	RsvdP	No	0-0h

Register 17-55. D84h Memory BAR2 Setup (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
2:1	**BAR2 Type** 00b = Selects 32-bit memory BAR No other values are allowed.	RW	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
30:12	**BAR2 Size** Specifies the Address Range size requested by **BAR2**. 0 = Corresponding bits in **BAR2** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR2** are RW bits	RW	Yes	0000_0h
31	**BAR2 Enable** 0 = **BAR2** is disabled, all bits in **BAR2** read 0 1 = **BAR2** is enabled, Size and Type specified in this register	RW	Yes	0

Note: *Register offset D84h must be programmed with the same value as the **NT Port Virtual Interface Memory BAR2 Setup** register (offset D4h). This requirement applies only to the NT Virtual Interface.*

Register 17-56. D88h Memory BAR3 Setup (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
15:1	***Reserved***	RsvdP	No	0-0h
19:16	**LUT Page Size** Virtual Interface **BAR3** Address range. The total size of this range is dependent upon the page size. When the *BAR3 LUT Page Size Extension* bit value is 0, the encodings are as follows: 0h = Disables **BAR3** 1h to 4h = ***Reserved*** 5h = 4 KB 6h = 8 KB 7h = 16 KB 8h = 32 KB 9h = 64 KB Ah = 128 KB Bh = 256 KB Ch = 512 KB Dh = 1 MB Eh = 2 MB Fh = 4 MB When the *BAR3 LUT Page Size Extension* bit value is 1, the encodings are as follows: 0h = 8 MB 1h = 16 MB 2h = 32 MB 3h to Fh = ***Reserved***	RW	Yes	0h
20	**LUT Page Size Extension** Allows selection of larger page sizes when programming Page Size[19:16]. 0 = Page sizes 4 KB through 4 MB are available in Page Size[19:16] 1 = Page sizes 8 through 32 MB are available in Page Size[19:16]	RW	Yes	0
31:21	***Reserved***	RsvdP	No	000h

Note: *Register offset D88h must be programmed with the same value as the **NT Port Virtual Interface Memory BAR3 Setup** register (offset D8h). This requirement applies only to the NT Virtual Interface.*

Register 17-57. D8Ch Memory BAR4 Setup (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
2:1	**BAR4 Type** 00b = **BAR4** is implemented as a 32-bit Memory BAR 10b = **BAR4/5** is implemented as a 64-bit Memory BAR ***Note:*** *It is illegal to program 10b and clear the* ***Memory BAR5 Setup (Shadow Copy)*** *register BAR5 Enable bit (offset D90h[31]).*	RW	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
30:12	**BAR4 Size** Specifies the Address Range size requested by **BAR4**. 0 = Corresponding bits in **BAR4** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR4** are RW bits	RW	Yes	0000_0h
31	**BAR4 Enable** When bits [2:1] = 00b, enables **BAR4**; otherwise, belongs to the *BAR4 Size*[30:12] field. 0 = **BAR4** is disabled, all bits in **BAR4** read 0 1 = **BAR4** is enabled, Size and Type specified in this register	RW	Yes	0

Note: *Register offset D8Ch must be programmed with the same value as the* ***NT Port Virtual Interface Memory BAR4/5 Setup*** *register (offset DCh). This requirement applies only to the NT Virtual Interface.*

Register 17-58. D90h Memory BAR5 Setup (Shadow Copy)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
30:0	**BAR5 Size** Together with the **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Size* field (offset D8Ch[31:12]), specifies the Address Range size requested by **BAR4/5** in 64-bit mode when the **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* field (offset D8Ch[2:1]) is set to 10b. 0 = Corresponding bits in **BAR5** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR5** are RW bits	RW	Yes	0-0h
	Reserved when the **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* field (offset D8Ch[2:1]) is cleared to 00b.	RsvdP	No	0-0h
31	**BAR5 Enable** 0 = **BAR5** is disabled 1 = **BAR5** is enabled when the **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* field (offset D8Ch[2:1]) is set to 10b ***Note:*** *It is illegal to program the **Memory BAR4 Setup (Shadow Copy)** register BAR4 Type field (offset D8Ch[2:1]) to 10b and clear this bit.*	RW	Yes	0
	Reserved when the **Memory BAR4 Setup (Shadow Copy)** register *BAR4 Type* field (offset D8Ch[2:1]) is cleared to 00b.	RsvdP	No	0

Note: *Register offset D90h must be programmed with the same value as the **NT Port Virtual Interface Memory BAR5 Setup** register (offset E0h). This requirement applies only to the NT Virtual Interface.*

17.13.5 NT Port Virtual Interface Send Lookup Table Entry Registers

This section describes the NT Port Virtual Interface Send (Requester ID Translation) Lookup Table (LUT) Entry registers. NT Port uses these registers for Requester ID translation when it forwards:

- Memory requests from NT Port Virtual Interface to the NT Port Link Interface, or
- Completion TLP from NT Port Link Interface to the NT Port Virtual Interface

When non-posted traffic must be sent through the NT Virtual Interface, program the registers listed in this group with the upstream Requester ID and set the *LUT Entry_n Enable* bit (bit 31) of each register accordingly.

Table 17-22 defines the register and address locations. The register descriptions that follow the table define the bit definitions that apply to the two register types.

Table 17-22. NT Port Virtual Interface Send Lookup Table Entry_*n* Register Locations

ADDR Location	Lookup Table Entry_*n*	ADDR Location	Lookup Table Entry_*n*
D94h	0	DA4h	4
D98h	1	DA8h	5
D9Ch	2	DACh	6
DA0h	3	DB0h	7

Register 17-59. D94h - DB0h Virtual Interface Send Lookup Table Entry_*n* (where *n* = 0 through 7)

Bit(s)	Description		Type	Serial EEPROM and I^2C	Default
2:0	**Requester ID**	**Function Number** LUT Entry_*n* Requester Function Number.	RW	Yes	000b
7:3		**Device Number** LUT Entry_*n* Requester Device Number.	RW	Yes	0000_0b
15:8		**Bus Number** LUT Entry_*n* Requester Bus Number.	RW	Yes	00h
30:16	***Reserved***		RsvdP	No	0-0h
31	**LUT Entry_*n* Enable** 0 = Disables 1 = Enables		RW	Yes	0

17.14 Advanced Error Reporting Capability Registers

The registers defined in Section 13.14, "Advanced Error Reporting Capability Registers," are also applicable to the NT Port Virtual Interface, except as defined in Register 17-60 through Register 17-62. Table 17-23 defines the register map used by the NT Port Virtual and Link Interfaces.

Table 17-23. Advanced Error Reporting Capability Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (138h) / Capability Version (1h)	PCI Express Extended Capability ID (0001h)	FB4h
Uncorrectable Error Status[a]		FB8h
Uncorrectable Error Mask[a]		FBCh
Uncorrectable Error Severity[a]		FC0h
Correctable Error Status		FC4h
Correctable Error Mask		FC8h
Advanced Error Capabilities and Control		FCCh
Header Log 0		FD0h
Header Log 1		FD4h
Header Log 2		FD8h
Header Log 3		FDCh
Reserved		FE0h – FFCh

a. *Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 17-60. FB8h Uncorrectable Error Status

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
5	**Surprise Down Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
11:6	***Reserved***	RsvdP	No	0000_00b
12	**Poisoned TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
13	**Flow Control Protocol Error Status**	RsvdP[a]	No	0
14	**Completion Timeout Status**	RsvdP[a]	No	0
15	**Completer Abort Status**	RsvdP[a]	No	0
16	**Unexpected Completion Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
17	**Receiver Overflow Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
18	**Malformed TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
19	**ECRC Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
20	**Unsupported Request Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 17-61. FBCh Uncorrectable Error Mask

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
5	**Surprise Down Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
11:6	***Reserved***	RsvdP	No	0000_00b
12	**Poisoned TLP Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
13	**Flow Control Protocol Error Mask**	RsvdP[a]	No	0
14	**Completion Timeout Mask**	RsvdP[a]	No	0
15	**Completer Abort Mask**	RsvdP[a]	No	0
16	**Unexpected Completion Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
17	**Receiver Overflow Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
18	**Malformed TLP Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
19	**ECRC Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
20	**Unsupported Request Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 17-62. FC0h Uncorrectable Error Severity

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
5	**Surprise Down Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
11:6	***Reserved***	RsvdP	No	0-0h
12	**Poisoned TLP Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
13	**Flow Control Protocol Error Severity**	RsvdP[a]	No	0
14	**Completion Timeout Severity**	RsvdP[a]	No	0
15	**Completer Abort Severity**	RsvdP[a]	No	0
16	**Unexpected Completion Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
17	**Receiver Overflow Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
18	**Malformed TLP Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
19	**ECRC Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
20	**Unsupported Request Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Chapter 18 NT Port Link Interface Registers

18.1 Introduction

This chapter defines the PEX 8508 Non-Transparent (NT) Port Link Interface (interface) registers. The NT Port includes two sets of Configuration Capability, Control, and Status registers to support the Virtual and Link Interfaces. Table 18-1 defines the NT Port Link Interface register mapping.

The NT Port Virtual Interface registers are defined in Chapter 17, "NT Port Virtual Interface Registers." The Transparent mode registers are defined in Chapter 13, "Transparent Port Registers."

For further details regarding register names and descriptions, refer to the following specifications:

- *PCI r3.0*
- *PCI Power Mgmt. r1.2*
- *PCI Express Base r1.1*

Table 18-1. NT Port Link Interface Type 0 Register Map

31 30 29 28 27 26 25 24 23 22 21 20	19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
PCI-Compatible Type 0 Configuration Registers			Capability Pointer (40h)	00h … 34h … 3Ch
		Next Capability Pointer (48h)	Capability ID (01h)	40h
Power Management Capability Registers				44h
		Next Capability Pointer (68h)	Capability ID (05h)	48h
Message Signaled Interrupt Capability Registers				… 64h
		Next Capability Pointer (00h)	Capability ID (10h)	68h
PCI Express Capability Registers				… 8Ch
NT Port Link Interface Registers				90h … FCh
Next Capability Offset (FB4h)	1h	PCI Express Extended Capability ID (0003h)		100h
Device Serial Number Extended Capability Registers				104h 108h
Reserved				10Ch – 134h
Next Capability Offset (148h)	1h	PCI Express Extended Capability ID (0004h)		138h
Power Budgeting Extended Capability Registers				… 144h
Next Capability Offset (000h)	1h	PCI Express Extended Capability ID (0002h)		148h
Virtual Channel Extended Capability Registers				… 1C4h
Device-Specific Registers				1C8h … C08h
Non-Transparent Bridging-Specific Registers				C3Ch … FB0h
Next Capability Offset (138h)	1h	PCI Express Extended Capability ID (0001h)		FB4h
Advanced Error Reporting Capability Registers				… FFCh

18.2 Register Access

The PEX 8508 NT Port Link Interface implements a 4-KB Configuration space. The lower 256 bytes (offsets 00h through FFh) is the PCI-compatible Configuration space, and the upper 960 Dwords (offsets 100h through FFFh) is the PCI Express Extended Configuration space. The PEX 8508 supports four mechanisms for accessing NT Port Link Interface registers:

- PCI Express Base r1.1 Configuration Mechanism
- Device-Specific Memory-Mapped Configuration Mechanism
- Device-Specific I/O-Mapped Configuration Mechanism
- Device-Specific Cursor Mechanism

18.2.1 *PCI Express Base r1.1* Configuration Mechanism

The PCI Express Configuration Mechanism is divided into two mechanisms:

- *PCI r3.0*-Compatible Configuration
- PCI Express Enhanced Configuration

The PCI r3.0-Compatible Configuration Mechanism provides Conventional PCI access to the first 256 bytes (the bytes at offsets 00h through FFh) of the NT Port Link Interface Configuration Register space. The PCI Express Enhanced Configuration Mechanism provides access to the remaining 4 KB (offsets 100h through FFFh).

The PEX 8508 decodes Type 0 Configuration transactions received on its NT Port Link Interface. The PEX 8508 reads from or writes to the NT Port Link Interface register, as specified in the original Type 0 Configuration access.

18.2.1.1 *PCI r3.0*-Compatible Configuration Mechanism

The *PCI r3.0*-Compatible Configuration space consists of the first 256 bytes of the NT Port Link Interface Configuration space. (Refer to Figure 18-1.) The *PCI r2.3*-Compatible Configuration mechanism provides standard access to the PEX 8508 NT Port Link Interface's first 256 bytes (the bytes at offsets 00h through FFh) of the PCI Express Configuration space. This mechanism is used to access the PEX 8508 NT Port Link Interface Type 0 (PCI endpoint) registers:

- PCI-Compatible Type 0 Configuration Registers
- Power Management Capability Registers
- Message Signaled Interrupt Capability Registers
- PCI Express Capability Registers

Because the *PCI r3.0*-Compatible Configuration mechanism is limited to the first 256 bytes of the NT Port Link Interface Configuration Register space, one of the following must be used to access beyond byte FFh:

- PCI Express Enhanced Configuration Mechanism
- Device-Specific Memory-Mapped Configuration Mechanism
- Device-Specific Cursor Mechanism

The *PCI r3.0*-Compatible Configuration mechanism uses the same request format as the Extended PCI Express Mechanism. For PCI-compatible Configuration requests, the Extended Register Address field must be all zeros (0).

Do not use this mechanism to access the PEX 8508 Device-Specific Configuration registers.

18.2.1.2 PCI Express Enhanced Configuration Mechanism

The PCI Express Enhanced Configuration mechanism uses a flat, Root Complex Memory-Mapped Address space to access Device Configuration registers. In this case, the Memory address determines the Configuration register accessed, and Memory data returns the addressed register contents. The Root Complex converts the Memory transaction into a Configuration transaction before transmitting this access to the downstream devices.

This mechanism is used to access the NT Port Link Interface Type 0 registers:

- PCI-Compatible Type 0 Configuration Registers
- Power Management Capability Registers
- Message Signaled Interrupt Capability Registers
- PCI Express Capability Registers
- Device Serial Number Extended Capability Registers
- Power Budgeting Extended Capability Registers
- Virtual Channel Extended Capability Registers
- Advanced Error Reporting Capability Registers

18.2.2 Device-Specific Memory-Mapped Configuration Mechanism

The Device-Specific Memory-Mapped Configuration Mechanism provides a method to access the PEX 8508 Port Configuration registers of all ports in a single Memory map, as illustrated in Figure 18-1. The registers of each port are located within a 4-KB range.

When the NT Port is enabled at Fundamental Reset, the NT Port Virtual and Link Interface Configuration registers are used in place of the Type 1 Configuration registers for that port.

To utilize the Device-Specific Memory-Mapped Configuration mechanism, use the *PCI r3.0*-Compatible Configuration Mechanism to program the PEX 8508 upstream port **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets 10h and 14h, respectively). After the PEX 8508 upstream port Memory-Mapped register Base address is set, the upstream port register is accessed with Memory Reads from and Writes to the Configuration Space registers. The NT Port registers are accessed with Memory Reads from and Writes to the 4-KB range, starting at offset 64 KB for the Virtual Interface registers and offset 68 KB for the Link Interface registers.

This mechanism is used to access all PEX 8508 Configuration registers.

Figure 18-1. PEX 8508 Register Offset from Upstream Port BAR0/1 Base Address (Non-Transparent Mode)

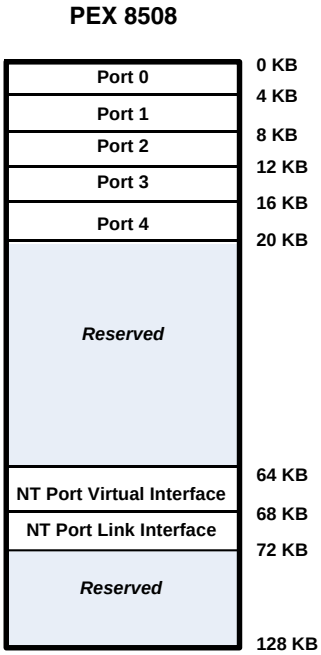

18.2.3 Device-Specific I/O-Mapped Configuration Mechanism

The first 256 bytes of NT Port Link Interface Configuration Space registers are directly accessed by an I/O transaction. The NT Port Link Interface **Base Address 1** register (**BAR1**. offset 14h) is used for I/O-Mapped access. (Refer to Figure 18-2.)

Extended Configuration Space registers are accessed by using the Cursor mechanism in I/O space.

Figure 18-2. I/O-Mapped Configuration Space View

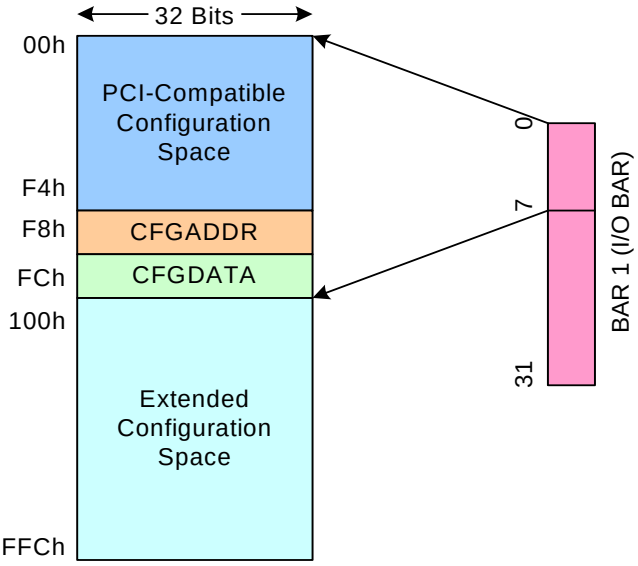

18.2.4 Device-Specific Cursor Mechanism

In Figure 18-2, the software uses the **Configuration Address Window** (CFGADDR) register (offset F8h) to select the NT Port Virtual or Link Interface Configuration Space registers, including the Extended Configuration Space register.

Software uses the **Configuration Data Window** (CFGDATA) register (offset FCh) to read or write to the selected Configuration Space registers.

Refer to Section 18.8.4, "NT Port Link Interface Cursor Mechanism Control Registers," for the register descriptions.

18.3 Register Descriptions

The remainder of this chapter details the PEX 8508 NT Port Link Interface registers, including:

- Bit/field names
- Description of register functions in the PEX 8508 NT Port Link and Virtual Interfaces
- Type (*such as* RW or HwInit; refer to Table 13-3, "Register Types, Grouped by User Accessibility." for Type descriptions)
- Whether the power-on/reset value can be modified, by way of the PEX 8508 serial EEPROM or I^2C initialization feature
- Default power-on/reset value

18.4 PCI-Compatible Type 0 Configuration Registers

Table 18-2. PCI-Compatible Type 0 Configuration Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Device ID		Vendor ID		00h
Status		Command		04h
Class Code			Revision ID	08h
BIST ***(Not Supported)***	Header Type	Master Latency Timer	Cache Line Size	0Ch
Base Address 0				10h
Base Address 1				14h
Base Address 2				18h
Base Address 3				1Ch
Base Address 4				20h
Base Address 5				24h
Reserved				28h
Subsystem ID		Subsystem Vendor ID		2Ch
Expansion ROM Base Address				30h
Reserved			Capability Pointer	34h
Reserved				38h
Minimum Latency ***(Reserved)***	Minimum Grant ***(Reserved)***	Interrupt Pin	Interrupt Line	3Ch

Register 18-1. 00h Configuration ID

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Vendor ID** Unless overwritten by the serial EEPROM, returns the PLX PCI-SIG-assigned Vendor ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Vendor ID with another Vendor ID.	HwInit	Yes	10B5h
31:16	**Device ID** Unless overwritten by the serial EEPROM, the PEX 8508 returns 8508h, the PLX-assigned Device ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX-assigned Device ID with another Device ID.	HwInit	Yes	8508h

Register 18-2. 04h Command/Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Command			
0	**I/O Access Enable** 0 = PEX 8508 ignores I/O requests received on the NT Port Link Interface 1 = PEX 8508 accepts I/O requests received on the NT Port Link Interface	RW	Yes	0
1	**Memory Access Enable** 0 = PEX 8508 ignores Memory requests received on the NT Port Link Interface 1 = PEX 8508 accepts Memory requests received on the NT Port Link Interface	RW	Yes	0
2	**Bus Master Enable** Controls the PEX 8508 Memory request forwarding in the upstream direction. Does not affect message forwarding nor Completions in the upstream direction. 0 = PEX 8508 handles Memory requests received on the NT Port Virtual Interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8508 returns a Completion with UR Completion status 1 = PEX 8508 forwards Memory requests from the NT Port Virtual Interface to the Link Interface in the upstream direction	RW	Yes	0
3	**Special Cycle Enable** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
4	**Memory Write and Invalidate Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
5	**VGA Palette Snoop** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
6	**Parity Error Response Enable** Controls bit 24 (*Master Data Parity Error* bit).	RW	Yes	0
7	**IDSEL Stepping/Wait Cycle Control** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0

Register 18-2. 04h Command/Status *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
8	**SERR# Enable** Controls bit 30 (*Signaled System Error* bit). 1 = Enables reporting of Fatal and Non-Fatal errors detected by the NT Port Link Interface to the Root Complex	RW	Yes	0
9	**Fast Back-to-Back Transactions Enable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
10	**Interrupt Disable** 0 = NT Port Link Interface is enabled to generate INT*x* Interrupt messages 1 = NT Port Link Interface is prevented from generating INT*x* Interrupt messages	RW	Yes	0
15:11	***Reserved***	RsvdP	No	00h

Register 18-2. 04h Command/Status *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Status			
18:16	***Reserved***	RsvdP	No	000b
19	**Interrupt Status** 0 = No INT*x* Interrupt pending 1 = INT*x* Interrupt pending internally to the NT Port Link Interface	RO	Yes	0
20	**Capability List** Set to 1, as required by the *PCI Express Base r1.1*.	RO	Yes	1
21	**66 MHz Capable** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
22	***Reserved***	RsvdP	No	0
23	**Fast Back-to-Back Transactions Capable** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
24	**Master Data Parity Error** When the *Parity Error Response Enable* bit = 1, the NT Port Link Interface sets this bit to 1 when the NT Port: • Forwards the poisoned TLP Write request from the Virtual Interface to the Link Interface, **or** • Receives a Completion marked as poisoned on the Link Interface When the *Parity Error Response Enable* bit = 0, the PEX 8508 never sets this bit. This error is natively reported by the **Uncorrectable Error Status** register *Poisoned TLP Status* bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0
26:25	**DEVSEL# Timing** ***Not supported*** Always cleared to 00b.	RsvdP	No	00b
27	**Signaled Target Abort** When a Memory-Mapped access payload length is greater than one DWord, the NT Port Link Interface sets this bit to 1. This bit is also set to 1 when the NT Port forwards a Completion with Completer Abort (CA) status from the Virtual Interface to the Link Interface. ***Note:*** *When set during a forwarded Completion, the* ***Uncorrectable Error Status*** *register Completer Abort Status bit (offset FB8h[15]) is not updated, because the NT Port does not log the requests that it forwards.*	RW1C	Yes	0

Register 18-2. 04h Command/Status *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
28	**Received Target Abort** Defaults to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
29	**Received Master Abort** Defaults to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
30	**Signaled System Error** When the *SERR# Enable* bit is set to 1, the NT Port Link Interface sets this bit to 1 when transmitting an ERR_FATAL or ERR_NONFATAL message to its upstream device. This error is natively reported by the **Device Status** register *Fatal Error Detected* and *Non-Fatal Error Detected* bits (offset 70h[18:17], respectively), which are mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0
31	**Detected Parity Error** The NT Port Link Interface sets this bit to 1 when receiving a Poisoned TLP, regardless of the *Parity Error Response Enable* bit state. This error is natively reported by the **Uncorrectable Error Status** register *Poisoned TLP Status* bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RW1C	Yes	0

Register 18-3. 08h Class Code and Revision ID

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Revision ID** Unless overwritten by the serial EEPROM, returns ACh, PLX-assigned Revision ID for this version of the PEX 8508. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Revision ID with another Revision ID. ***Note:*** *When a serial EEPROM is not used, the default values for this field are loaded. When using a serial EEPROM, refer to the following table for the appropriate configuration settings for this field.*	RO	Yes	ACh

	Serial EEPROM Value	Results in Actual Register Value
Revision ID	ACh	AAh
	AAh	ACh

Class Code				068000h
15:8	**Register-Level Programming Interface** Cleared to 00h, as required by the *PCI r3.0* for other bridge devices.	RO	Yes	00h
23:16	**Sub-Class Code** Other bridge devices.	RO	Yes	80h
31:24	**Base Class Code** Bridge devices.	RO	Yes	06h

Register 18-4. 0Ch Miscellaneous Control

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Cache Line Size** System Cache Line Size. Implemented as a Read-Write field for Conventional PCI compatibility purposes and does not impact PEX 8508 functionality.	RW	Yes	00h
15:8	**Master Latency Timer** ***Not supported*** Cleared to 00h, as required by the *PCI Express Base r1.1*.	RsvdP	No	00h
22:16	**PCI Header** Type 0 Configuration Header for the NT Port.	RO	Yes	00h
23	**Header Type** 0 = PEX 8508 is a single-function device	RO	Yes	0
31:24	**BIST** ***Not supported*** BIST Pass/Fail.	RO	No	00h

Register 18-5. 10h Base Address 0 (NT Port Link Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** When enabled, the Base Address register maps PEX 8508 port Configuration registers into Memory space. NT Port Link **BAR0** is configured by the serial EEPROM and Local Host. By default, the BAR Setup register selects 32-bit Memory **BAR0** and 32-bit I/O **BAR1** for CSR mapping. ***Note:*** *Hardwired to 0.*	RO	No	0
2:1	**Memory Map Type** 00b = PEX 8508 Configuration registers are mapped anywhere in 32-bit Memory Address space only	RO	Yes	00b
3	**Prefetchable** The Base Address register maps PEX 8508 Configuration registers into Non-Prefetchable Memory space by default. ***Note:*** *Hardwired to 0.*	RO	No	0
16:4	***Reserved***	RsvdP	No	0-0h
31:17	**Base Address 0** 128-KB Base address in which to map the PEX 8508 Configuration Space registers into Memory space. ***Note:*** *The upstream port is hardwired to 0.*	RW	Yes	0000h

Register 18-6. 14h Base Address 1 (NT Port Link Interface I/O Space)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**I/O Space Indicator** I/O BAR when the **NT Port Link Interface BAR0/BAR1 Setup** register *BAR0/1 Type Selector* field (offset E4h[1:0]) is set to 11b.	RO	Yes	1
	Reserved when the **NT Port Link Interface BAR0/BAR1 Setup** register *BAR0/1 Type Selector* field (offset E4h[1:0]) is cleared to 00b.	RsvdP	No	0
7:1	***Reserved***	RsvdP	No	0000_000b
31:8	**I/O Base Address** 256-byte I/O Space Base address when the **NT Port Link Interface BAR0/ BAR1 Setup** register *BAR0/1 Type Selector* field (offset E4h[1:0]) is set to 11b.	RW	Yes	0000_00h
	Reserved when the **NT Port Link Interface BAR0/BAR1 Setup** register *BAR0/1 Type Selector* field (offset E4h[1:0]) is cleared to 00b.	RsvdP	No	0000_00h

Register 18-7. 18h Base Address 2 (NT Port Link Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Memory Space Indicator** 0 = Implemented as a Memory BAR; otherwise, ***reserved***	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 10b = Mappable anywhere in 64-bit Memory space 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Address 2** Base Address is enabled and sized by the **NT Port Link Interface Memory BAR2/3 Setup** register (offset E8h). This **BAR2/3** group uses direct address translation. The minimum BAR size is programmed to 4 KB.	RW	Yes	0000_0h
	Reserved when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Register 18-8. 1Ch Base Address 3 (NT Port Link Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Upper 32-Bit BAR Address** NT Port Link Interface upper 32-bit address when **BAR2/3** is implemented as a 64-bit BAR [**NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is set to 11b]. RW, based upon the **NT Port Link Interface Memory BAR2/3 Setup** and **NT Port Link Interface Memory BAR3 Setup** registers (offsets E8h and ECh, respectively). The **BAR2/3** group uses direct address translation.	RW	Yes	0-0h
	Reserved when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Register 18-9. 20h Base Address 4 (NT Port Link Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Memory Space Indicator** 0 = Implemented as a Memory BAR; otherwise, ***reserved***	RO	Yes	0
2:1	**Memory Map Type** 00b = Mappable anywhere in 32-bit Memory space 10b = Mappable anywhere in 64-bit Memory space 01b, 11b = ***Reserved***	RO	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RO	Yes	0
11:4	***Reserved***	RsvdP	No	00h
31:12	**Base Address 4** Base Address size is enabled and sized by the **NT Port Link Interface Memory BAR4/5 Setup** register (offset F0h). The **BAR4/5** group uses direct address translation. The minimum BAR size is programmed to 4 KB.	RW	Yes	0000_0h
	Reserved when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Register 18-10. 24h Base Address 5 (NT Port Link Interface Memory Space)

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Base Address 5** NT Port Link Interface upper 32-bit address when **BAR4/5** is implemented as a 64-bit BAR [**NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is set to 10b]. RW, based upon the **NT Port Link Interface Memory BAR4/5 Setup** and **NT Port Link Interface Memory BAR5 Setup** registers (offsets F0h and F4h, respectively). The **BAR4/5** group uses direct address translation.	RW	Yes	0-0h
	Reserved when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is cleared to 00b.	RsvdP	No	0-0h

Register 18-11. 2Ch Subsystem ID and Subsystem Vendor ID

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
15:0	**Subsystem Vendor ID** Unless overwritten by the serial EEPROM, returns the PLX PCI-SIG-assigned Vendor ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX Vendor ID with another Vendor ID.	HwInit	Yes	10B5h
31:16	**Subsystem ID** Unless overwritten by the serial EEPROM, the PEX 8508 returns 8508h, the PLX-assigned Device ID. The PEX 8508 Serial EEPROM register initialization capability is used to replace the PLX-assigned Device ID with another Device ID.	HwInit	Yes	8508h

Register 18-12. 30h Expansion ROM Base Address

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
0	**Expansion ROM Enable** 0 = Expansion ROM is disabled	NT Virtual Port, offset 30h[0]=1	RsvdP	No	0
	1 = NT Port Virtual Interface Expansion ROM is enabled, and NT Port Link Interface Expansion ROM is disabled	NT Virtual Port, offset 30h[0]=0	RO	Yes	0
		NT Virtual Port, offset 30h[0]=1	RsvdP	No	0
10:1	***Reserved***		RsvdP	No	0-0h
31:11	**Expansion ROM Base Address** Expansion ROM minimum size is 2 KB (program register initially to FFFF_F801h). Expansion ROM maximum size is 32 KB (program register initially to FFFF_8001h). ***Note: This BAR must not be programmed to enable more than 32 KB.*** *Expansion ROM is limited to 32 KB, because the largest serial EEPROM that can be used is 64 KB (limit of 16-bit addressing) and the Expansion ROM image is stored in the serial EEPROM, beginning at address 32 KB.*	NT Virtual Port, offset 30h[0]=0	RsvdP	No	0-0h
		NT Virtual Port, offset 30h[0]=1	RW	Yes	0-0h

Note: ***Expansion ROM can be enabled in either the NT Port Link Interface or NT Port Virtual Interface, but not both simultaneously****. The default values of the NT Port Link Interface Expansion ROM Enable (NT Link Port, offset 30h[0]=1) and NT Port Virtual Interface Expansion ROM Enable (NT Virtual Port, offset 30h[0]=0) bits enable the NT Port Link Interface Expansion ROM, and disable the NT Port Virtual Interface Expansion ROM.*

If the NT Port Virtual Interface Expansion ROM is enabled (NT Virtual Port, offset 30h[0]=1), the NT Port Link Interface ***Expansion ROM Base Address*** *register (NT Link Port, offset 30h) is cleared and is not programmable.*

Register 18-13. 34h Capability Pointer

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Capability Pointer** Default 40h points to the **Power Management Capability** register.	RO	Yes	40h
31:8	***Reserved***	RsvdP	No	0000_00h

Register 18-14. 3Ch Interrupt

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Interrupt Line** The Interrupt Line Routing Value communicates interrupt line routing information. Values in this register are programmed by system software and are system architecture-specific. The value is used by device drivers and operating systems.	RW	Yes	00h
15:8	**Interrupt Pin** Identifies the Conventional PCI interrupt message(s) the device (or device function) uses. Only value 00h or 01h is allowed in the PEX 8508. 00h = Indicates that the device does not use Conventional PCI Interrupt message(s) 01h, 02h, 03h, and 04h = Maps to Conventional PCI Interrupt messages for INTA#, INTB#, INTC#, and INTD#, respectively	RO	Yes	01h
23:16	**Minimum Grant** ***Reserved*** Does not apply to PCI Express.	RsvdP	No	00h
31:24	**Minimum Latency** ***Reserved*** Does not apply to PCI Express.	RsvdP	No	00h

18.5 Power Management Capability Registers

This section details the NT Port Link Interface Power Management Capability registers. Table 18-3 defines the register map.

Table 18-3. Power Management Capability Register Map (All Ports)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Power Management Capability		Next Capability Pointer (48h)	Capability ID (01h)	40h
Data	Power Management Control/ Status Bridge Extensions	Power Management Status and Control		44h

Register 18-15. 40h Power Management Capability

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**Capability ID** Default = 01h – only value allowed.	RO	Yes	01h
15:8	**Next Capability Pointer** Default 48h points to the **Message Signaled Interrupt Capability** register.	RO	Yes	48h
18:16	**Version** Default = 011b – only value allowed.	RO	Yes	011b
19	**PME Clock** Does not apply to PCI Express. Returns 0.	RsvdP	No	0
20	***Reserved***	RsvdP	No	0
21	**Device-Specific Initialization** Default 0 indicates that Device-Specific Initialization is ***not*** required.	RO	Yes	0
24:22	**AUX Current** Value can be strapped to 111b in the serial EEPROM.	RO	Yes	000b
25	**D1 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D1 Power state.	RsvdP	No	0
26	**D2 Support** Default value of 0 indicates that the PEX 8508 does ***not support*** the D2 Power state.	RsvdP	No	0
31:27	**PME Support** PME messages are disabled by default.	RO	Yes	0000_0b

Register 18-16. 44h Power Management Status and Control

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Power Management Status and Control				
1:0	**Power State** This field is used to determine the current power state of the port, and to set the port into a new power state. 00b = D0 01b = D1 – ***Not supported*** 10b = D2 – ***Not supported*** 11b = D3hot If software attempts to write an unsupported state to this field, the Write operation completes normally; however, the data is discarded and no state change occurs.	RW	Yes	00b
2	***Reserved***	RsvdP	No	0
3	**No Soft Reset**	RO	Yes	1
7:4	***Reserved***	RsvdP	No	0h
8	**PME Enable** Default value of 0 indicates that PME generation is disabled.	RsvdP	No	0
12:9	**Data Select** Writable by Serial EEPROM only[a]. Bits [12:9] select the **Data** and **Data Scale** registers. 0h = D0 power consumed 3h = D3hot power consumed 4h = D0 power dissipated 7h = D3hot power dissipated All other values are ***reserved***.	RO	Yes	0h
	Not supported RO for hardware auto-configuration.	RO	No	0h
14:13	**Data Scale** Loaded by serial EEPROM[a]. There are four internal **Data Scale** registers (one each per **Data** register – 0, 3, 4 and 7), per port. Bits [12:9], *Data Select*, select the **Data Scale** register.	RO	Yes	00b
15	**PME Status** 0 = PME is not being generated by the NT Port	RsvdP	No	0

Register 18-16. 44h Power Management Status and Control *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Power Management Control/Status Bridge Extensions				
21:16	***Reserved***	RsvdP	No	0-0h
22	**B2/B3 Support** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.	RsvdP	No	0
23	**Bus Power/Clock Control Enable** ***Reserved*** Cleared to 0, as required by the *PCI Power Mgmt. r1.2*.	RsvdP	No	0
Data				
31:24	**Data** Loaded by serial EEPROM[a]. There are four internal **Data** registers (0, 3, 4, and 7), per port. Bits [12:9], *Data Select*, select the **Data** register.	RO	Yes	00h

*a. With no serial EEPROM, Reads return 00h for the **Data Scale** and **Data** registers (for all Data Selects).*

18.6 Message Signaled Interrupt Capability Registers

The Message Signaled Interrupt (MSI) Capability registers are defined in Section 13.8, "Message Signaled Interrupt Capability Registers." Table 18-4 defines the register map used by the NT Port Link Interface.

Table 18-4. Message Signaled Interrupt Capability Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Reserved	Message Control	Next Capability Pointer (68h)	Capability ID (05h)	48h
Message Address[31:0]				4Ch
Message Upper Address[63:32]				50h
Reserved		Message Data		54h
Mask Bit for MSI				58h
Reserved				5Ch – 64h

18.7 PCI Express Capability Registers

This section details the PEX 8508 PCI Express Capability registers. Table 18-5 defines the register map.

Table 18-5. PCI Express Capability Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
PCI Express Capabilities		Next Capability Pointer (00h)	Capability ID (10h)	68h
Device Capabilities				6Ch
Device Status		Device Control		70h
Link Capabilities				74h
Link Status		Link Control		78h
Reserved			7Ch –	8Ch

Register 18-17. 68h PCI Express Capability List and Capabilities

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	PCI Express Capability List			
7:0	**Capability ID** Set to 10h by default.	RO	Yes	10h
15:8	**Next Capability Pointer** 00h = PCI Express Capability is the last capability in the first 256-byte Configuration space of the PEX 8508 NT Port Link Interface capability list The PEX 8508 NT Port Link Interface Capabilities list starts at 100h.	RO	Yes	00h
	PCI Express Capabilities			
19:16	**Capability Version** The PEX 8508 NT Port Link Interface sets this field to 1h, as required by the *PCI Express Base r1.1*.	RO	Yes	1h
23:20	**Device/Port Type** Default = PCI Express Endpoint device.	RO	Yes	0h
24	**Slot Implemented** ***Not valid for PCI Express Endpoint devices.***	RsvdP	No	0
29:25	**Interrupt Message Number** The serial EEPROM writes 00_000b, because the Base message and MSI messages are the same.	RO	Yes	00_000b
31:30	***Reserved***	RsvdP	No	00b

Register 18-18. 6Ch Device Capabilities

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
2:0	**Maximum Payload Size Supported** 000b = NT Port Link Interface supports 128-byte maximum payload 001b = NT Port Link Interface supports 256-byte maximum payload No other values are supported.	RO	Yes	001b
4:3	**Phantom Functions Supported** ***Not supported*** Cleared to 00b.	RO	Yes	00b
5	**Extended Tag Field Supported** ***Not supported*** 0 = Maximum *Tag* field is 5 bits 1 = Maximum *Tag* field is 8 bits	RO	Yes	0
8:6	**Endpoint L0s Acceptable Latency**	RO	Yes	000b
11:9	**Endpoint L1 Acceptable Latency**	RO	Yes	000b
17:12	***Reserved***	RsvdP	No	00h
25:18	**Captured Slot Power Limit Value** For the NT Port Link Interface, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Scale* field.	RO	Yes	00h
27:26	**Captured Slot Power Limit Scale** For the NT Port Link Interface, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in the *Captured Slot Power Limit Value* field. 00b = 1.0 01b = 0.1 10b = 0.01 11b = 0.001	RO	Yes	00b
31:28	***Reserved***	RsvdP	No	0h

Register 18-19. 70h Device Status and Control

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Device Control			
0	**Correctable Error Reporting Enable** 0 = Disables 1 = Enables NT Port Link Interface to report Correctable errors to the system Host	RW	Yes	0
1	**Non-Fatal Error Reporting Enable** 0 = Disables 1 = Enables NT Port Link Interface to report Non-Fatal errors to the system Host	RW	Yes	0
2	**Fatal Error Reporting Enable** 0 = Disables 1 = Enables NT Port Link Interface to report Fatal errors to the system Host	RW	Yes	0
3	**Unsupported Request Reporting Enable** 0 = Disables 1 = Enables NT Port Link Interface to report Unsupported Request errors as an error message with a programmed uncorrectable error severity	RW	Yes	0
4	**Enable Relaxed Ordering** ***Not supported*** Cleared to 0.	RsvdP	No	0
7:5	**Maximum Payload Size** The NT Port Link Interface power-on/reset value is 000b, to support a Maximum Payload Size of 128 bytes. Software can change this field to configure the NT Port Link Interface to support other Payload sizes; however, software cannot change this field to a value larger than that indicated by the **Device Capabilities** register *Maximum Payload Size Supported* field (offset 6Ch[2:0]). for the Virtual and Link Interfaces. (Requester and Completer domains must possess the same Maximum Payload Size.) 000b = Indicates that initially the PEX 8508 port is configured to support a Maximum Payload Size of 128 bytes 001b = Indicates that initially the PEX 8508 port is configured to support a Maximum Payload Size of 256 bytes No other values are supported. ***Note:*** *Software must halt all transactions through the NT Port before changing this field.*	RW	Yes	000b
8	**Extended Tag Field Enable** ***Not supported*** Cleared to 0.	RsvdP	No	0
9	**Phantom Functions Enable** ***Not supported*** Cleared to 0.	RsvdP	No	0
10	**Auxiliary (AUX) Power PM Enable** Cleared to 0.	RO	No	0
11	**Enable No Snoop** ***Not supported*** Cleared to 0.	RsvdP	No	0

Register 18-19. 70h Device Status and Control *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
14:12	**Maximum Read Request Size** ***Not supported*** Cleared to 000b.	RsvdP	No	000b
15	***Reserved***	RsvdP	No	0
	Device Status			
16	**Correctable Error Detected** Set when the NT Port Link Interface detects a Correctable error, regardless of the bit 0 (*Correctable Error Reporting Enable* bit) state. 0 = NT Port Link Interface did not detect a Correctable error 1 = NT Port Link Interface detected a Correctable error	RW1C	Yes	0
17	**Non-Fatal Error Detected** Set when the NT Port Link Interface detects a Non-Fatal error, regardless of the bit 1 (*Non-Fatal Error Reporting Enable* bit) state. 0 = NT Port Link Interface did not detect a Non-Fatal error 1 = NT Port Link Interface detected a Non-Fatal error	RW1C	Yes	0
18	**Fatal Error Detected** Set when the NT Port Link Interface detects a Fatal error, regardless of the bit 2 (*Fatal Error Reporting Enable* bit) state. 0 = NT Port Link Interface did not detect a Fatal error 1 = NT Port Link Interface detected a Fatal error	RW1C	Yes	0
19	**Unsupported Request Detected** Set when the NT Port Link Interface detects an Unsupported Request, regardless of the bit 3 (*Unsupported Request Reporting Enable* bit) state. 0 = NT Port Link Interface did not detect an Unsupported Request 1 = NT Port Link Interface detected an Unsupported Request	RW1C	Yes	0
20	**Auxiliary (AUX) Power Detected** ***Not supported*** Cleared to 0.	RO	No	0
21	**Transactions Pending** ***Not supported*** Cleared to 0, as required by the *PCI Express Base r1.1*.	RsvdP	No	0
31:22	***Reserved***	RsvdP	No	000h

Register 18-20. 74h Link Capabilities

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	**Maximum Link Speed** Set to 0001b for 2.5 Gbps.	RO	Yes	0001b
9:4	**Maximum Link Width** Actual link width is set by STRAP_PORTCFG[4:0]. The PEX 8508 Maximum Link Width is x4 = 00_0100b.	RO	No	Set by Strapping ball levels
11:10	**Active State Power Management (ASPM) Support** Indicates the level of ASPM supported by the port. 01b = L0s link power state entry is supported All other values are ***reserved***.	RO	Yes	01b
14:12	**L0s Exit Latency** 101b = NT Port Link Interface L0s Exit Latency is between 1 and 2 µs	RO	No	101b
17:15	**L1 Exit Latency** 101b = NT Port Link Interface L1 Exit Latency is between 16 and 32 µs	RO	Yes	101b
23:18	***Reserved***	RsvdP	No	0-0h
31:24	**Port Number** The NT Port Number is selected by signal ball strapping options. Refer to STRAP_NT_UPSTRM_PORTSEL[3:0] for details.	HwInit	No	Set by Strapping ball levels

Register 18-21. 78h Link Status and Control

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
	Link Control			
1:0	**Active State Power Management (ASPM) Control** 00b = Disables L0s Link PM State Entry for NT Port 01b = Enables L0s Link PM State Entry Values of 10b and 11b are ***not allowed***.	RW	Yes	00b
2	***Reserved***	RsvdP	No	0
3	**Read Completion Boundary (RCB)** ***Not supported*** Cleared to 0.	RO	Yes	0
4	**Link Disable** ***Reserved*** for NT Port Link Interface.	RsvdP	No	0
5	**Retrain Link** ***Reserved*** for NT Port Link Interface.	RsvdP	No	0
6	**Common Clock Configuration** 0 = NT Port Link Interface and the device at the opposite end of the PCI Express link are operating with an **asynchronous** Reference Clock 1 = NT Port Link Interface and the device at the opposite end of the PCI Express link are operating with a **distributed** common Reference Clock	RW	Yes	0
7	**Extended Sync** Set to 1 causes the NT Port Link Interface to transmit: • 4,096 FTS Ordered-Sets in the L0s state, • Followed by a single SKIP Ordered-Set prior to entering the L0 state, • Finally, transmission of 1,024 TS1 Ordered-Sets in the Recovery state.	RW	Yes	0
15:8	***Reserved***	RsvdP	No	00h

Register 18-21. 78h Link Status and Control *(Cont.)*

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
	Link Status			
19:16	**Link Speed** The NT Port Link Interface is set to 0001b for 2.5 Gbps.	RO	Yes	0001b
25:20	**Negotiated Link Width** Indicates the negotiated width of the PCI Express link. 00_0001b = x1 00_0010b = x2 00_0100b = x4 All other values are ***not supported***. The value in this field is undefined when the link is not up.	RO	Yes	00_0001b
26	**Training Error** ***Reserved*** for NT Port Link Interface.	RsvdP	No	0
27	**Link Training** ***Reserved*** for NT Port Link Interface.	RsvdP	No	0
28	**Slot Clock Configuration** Set by the upstream port or NT Port Link Interface, but not both. 0 = Indicates that the PEX 8508 uses an independent clock 1 = Indicates that the PEX 8508 uses the same physical Reference Clock that the platform provides on the connector	HwInit	Yes	0
31:29	***Reserved***	RsvdP	No	000b

18.8 NT Port Link Interface Registers

Table 18-6. NT Port Link Interface Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
NT Port Link Interface IRQ Doorbell Registers		90h … ACh
NT Port Scratchpad (Mailbox) Registers		B4h … CCh
NT Port Link Interface BAR Setup Registers		D0h … F4h
NT Port Link Interface Cursor Mechanism Control Registers		F8h FCh

18.8.1 NT Port Link Interface IRQ Doorbell Registers

The PEX 8508 NT Port Link Interface Interrupt Request (IRQ) Doorbell registers are defined in Section 17.8.1, "NT Port Virtual Interface IRQ Doorbell Registers." Table 18-7 defines the register map used by the NT Port Link Interface.

Table 18-7. NT Port Link Interface Interrupt Request (IRQ) Doorbell Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Reserved	Set Virtual Interface IRQ	90h
Reserved	Clear Virtual Interface IRQ	94h
Reserved	Set Virtual Interface IRQ Mask	98h
Reserved	Clear Virtual Interface IRQ Mask	9Ch
Reserved	Set Link Interface IRQ	A0h
Reserved	Clear Link Interface IRQ	A4h
Reserved	Set Link Interface IRQ Mask	A8h
Reserved	Clear Link Interface IRQ Mask	ACh

18.8.2 NT Port Scratchpad (Mailbox) Registers

This section details the PEX 8508 NT Port Scratchpad (Mailbox) registers. Table 18-8 defines the register map.

Table 18-8. PEX 8508 NT Port Scratchpad (Mailbox) Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
NT Port Scratchpad_0				B0h
NT Port Scratchpad_1				B4h
NT Port Scratchpad_2				B8h
NT Port Scratchpad_3				BCh
NT Port Scratchpad_4				C0h
NT Port Scratchpad_5				C4h
NT Port Scratchpad_6				C8h
NT Port Scratchpad_7				CCh

Register 18-22. B0h NT Port Scratchpad_0

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_0** 32-bit Scratchpad_0 register.	RW	Yes	0-0h

Register 18-23. B4h NT Port Scratchpad_1

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_1** 32-bit Scratchpad_1 register.	RW	Yes	0-0h

Register 18-24. B8h NT Port Scratchpad_2

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_2** 32-bit Scratchpad_2 register.	RW	Yes	0-0h

Register 18-25. BCh NT Port Scratchpad_3

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_3** 32-bit Scratchpad_3 register.	RW	Yes	0-0h

Register 18-26. C0h NT Port Scratchpad_4

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_4** 32-bit Scratchpad_4 register.	RW	Yes	0-0h

Register 18-27. C4h NT Port Scratchpad_5

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_5** 32-bit Scratchpad_5 register.	RW	Yes	0-0h

Register 18-28. C8h NT Port Scratchpad_6

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_6** 32-bit Scratchpad_6 register.	RW	Yes	0-0h

Register 18-29. CCh NT Port Scratchpad_7

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**Scratchpad_7** 32-bit Scratchpad_7 register.	RW	Yes	0-0h

18.8.3 NT Port Link Interface BAR Setup Registers

This section details the NT Port Link Interface BAR Setup registers. Table 18-9 defines the register map.

Table 18-9. PEX 8508 NT Port Link Interface BAR Setup Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Reserved	D0h – E0h
NT Port Link Interface BAR0/BAR1 Setup	E4h
NT Port Link Interface Memory BAR2/3 Setup	E8h
NT Port Link Interface Memory BAR3 Setup	ECh
NT Port Link Interface Memory BAR4/5 Setup	F0h
NT Port Link Interface Memory BAR5 Setup	F4h

Register 18-30. E4h NT Port Link Interface BAR0/BAR1 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
1:0	**BAR0/1 Type Selector** 00b = 32-bit Memory **BAR0** 11b = Link Interface **BAR0** is a 32-bit Memory BAR and Link Interface **BAR1** is an I/O BAR All other codes disable **BAR1** implementation.	RW	Yes	11b
31:2	***Reserved***	RsvdP	No	0-0h

Register 18-31. E8h NT Port Link Interface Memory BAR2/3 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
2:1	**BAR2 Type** 00b = **BAR2** is implemented as a 32-bit Memory BAR 10b = **BAR2/3** is implemented as a 64-bit Memory BAR ***Note:*** *It is illegal to program 10b and clear the* ***NT Port Link Interface Memory BAR3 Setup*** *register BAR3 Enable bit (offset ECh[31]).*	RW	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
30:12	**BAR2 Size** Specifies the Address Range size requested by **BAR2**. 0 = Corresponding bits in **BAR2** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR2** are RW bits	RW	Yes	0000_0h
31	**BAR2 Enable** 0 = **BAR2** is disabled (bits [2:1] = 00b), all bits in **BAR2** read 0 1 = **BAR2** is enabled, Size and Type specified in this register	RW	Yes	0

Register 18-32. ECh NT Port Link Interface Memory BAR3 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
30:0	**BAR3 Size** Specifies the Address Range size requested by **BAR2/3** in 64-bit mode when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is set to 10b. 0 = Read-Only bits that always return 0, writes are ignored 1 = Corresponding bits are RW bits	RW	Yes	0-0h
	Reserved when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is cleared to 00b.	RsvdP	No	0
31	**BAR3 Enable** 0 = **BAR3** is disabled 1 = **BAR2/3** is enabled when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is set to 10b ***Note:*** *It is illegal to program the* ***NT Port Link Interface Memory BAR2/3 Setup*** *register BAR2 Type field (offset E8h[2:1]) to 10b and clear this bit.*	RW	Yes	0
	Reserved when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is cleared to 00b.	RsvdP	No	0

Register 18-33. F0h NT Port Link Interface Memory BAR4/5 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	***Reserved***	RsvdP	No	0
2:1	**BAR4 Type** 00b = **BAR4** is implemented as a 32-bit BAR 10b = **BAR4/5** is implemented as a 64-bit Memory BAR ***Note:*** *It is illegal to program 10b and clear the* ***NT Port Link Interface Memory BAR5 Setup*** *register BAR5 Enable bit (offset F4h[31]).*	RW	Yes	00b
3	**Prefetchable** 0 = Non-Prefetchable 1 = Prefetchable	RW	Yes	0
11:4	***Reserved***	RsvdP	No	00h
30:12	**BAR4 Size** Specifies the Address Range size requested by **BAR4**. 0 = Corresponding bits in **BAR4** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR4** are RW bits	RW	Yes	0000_0h
31	**BAR4 Enable** When bits [2:1] = 00b, enables **BAR4**; otherwise, belongs to the *BAR4 Size*[30:12] field. 0 = **BAR4** is disabled, all bits in **BAR4** read 0 1 = **BAR4** is enabled, Size and Type specified in this register	RW	Yes	0

Register 18-34. F4h NT Port Link Interface Memory BAR5 Setup

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
30:0	**BAR5 Size** Together with the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Size* field (offset F0h[31:12]), specifies the Address Range size requested by **BAR4/5** in 64-bit mode when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is set to 10b. 0 = Corresponding bits in **BAR5** are Read-Only bits that always return 0, and Writes are ignored 1 = Corresponding bits in **BAR5** are RW bits	RW	Yes	0-0h
	Reserved when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is cleared to 00b.	RsvdP	No	0-0h
31	**BAR5 Enable** 0 = **BAR5** is disabled 1 = **BAR4/5** is enabled when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is set to 10b ***Note:*** *It is illegal to program the* ***NT Port Link Interface Memory BAR4/5 Setup*** *register BAR4 Type field (offset F0h[2:1]) to 10b and clear this bit.*	RW	Yes	0
	Reserved when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is cleared to 00b.	RsvdP	No	0

18.8.4 NT Port Link Interface Cursor Mechanism Control Registers

This section details the NT Port Link Interface Cursor Mechanism Control registers. Table 18-10 defines the register map.

The Cursor Mechanism registers at offsets F8h/FCh provide a means for accessing PCI Express Extended Configuration Space registers (offsets 100h through FFFh) within the NT Port Virtual and Link Interfaces, when only standard PCI Configuration transactions (that do not support Extended Register Number), and/or I/O Request transactions (using the NT Port **BAR1** address, if enabled) are available. The Cursor Mechanism can generally access only those registers that are defined by the *PCI Express Base r1.1*, and not the device-specific registers.

Table 18-10. NT Port Link Interface Cursor Mechanism Control Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Configuration Address Window		***Reserved***		F8h
Configuration Data Window				FCh

Register 18-35. F8h Configuration Address Window

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
15:0	***Reserved***	RsvdP	No	0000h
25:16	**Register Offset**	RW	Yes	000h
30:26	***Reserved***	RsvdP	No	00h
31	**Interface Select** 0 = Access to NT Port Link Interface Type 0 Configuration Space register 1 = Access to NT Port Virtual Interface Type 0 Configuration Space register	RW	Yes	0

Register 18-36. FCh Configuration Data Window

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
31:0	**Data Window** Software selects a register by writing into the NT Port Link Interface Configuration Address window, then reads from or writes to that register using this register.	RW	Yes	0-0h

18.9 Device Serial Number Extended Capability Registers

The registers defined in Section 13.10, "Device Serial Extended Number Capability Registers," are also applicable to the NT Port Link Interface. Table 18-11 defines the register map.

Table 18-11. PEX 8508 Device Serial Number Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (FB4h) / Capability Version (1h)	PCI Express Extended Capability ID (0003h)	100h
Serial Number (Lower DW)		104h
Serial Number (Upper DW)		108h

18.10 Power Budgeting Extended Capability Registers

The registers defined in Section 13.11, "Power Budgeting Extended Capability Registers," are also applicable to the NT Port Link Interface. Table 18-12 defines the register map.

Table 18-12. PEX 8508 Power Budgeting Extended Capability Register Map (All Ports)

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (148h) / Capability Version (1h)	PCI Express Extended Capability ID (0004h)	138h
Reserved	Data Select	13Ch
Power Budgeting Data		140h
Power Budget Capability		144h

18.11 Virtual Channel Extended Capability Registers

The registers defined in Section 13.12, "Virtual Channel Extended Capability Registers," are also applicable to the NT Port Link Interface. Table 18-13 defines the register map.

The **Port VC Capability 1**, **VC0 Resource Control**, and **VC1 Resource Control** register (offsets 14Ch, 15Ch, and 16Ch, respectively) values are shadowed in the **NT Link Interface VC Capability 1 (Shadow Copy)**, **NT Link Interface VC0 Resource Control (Shadow Copy)**, and **NT Link Interface VC1 Resource Control (Shadow Copy)** registers (offsets D64h, D5Ch, and D60h, respectively).

Note: *These registers are not automatically shadowed when programmed by serial EEPROM; therefore, if the serial EEPROM programs these registers to non-default values, software must write the values to the registers (such as by reading the register and writing back the value), to cause the corresponding Shadow registers in the NT Port Virtual Interface to update.*

Table 18-13. NT Port Link Interface Virtual Channel Extended Capability Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (000h) / Capability Version (1h)	PCI Express Extended Capability ID (0002h)	148h
Port VC Capability 1		14Ch
Port VC Capability 2		150h
Port VC Status ***(Reserved)***	Port VC Control	154h
VC0 Resource Capability		158h
VC0 Resource Control		15Ch
VC0 Resource Status	***Reserved***	160h
VC1 Resource Capability		164h
VC1 Resource Control		168h
VC1 Resource Status	***Reserved***	16Ch
Reserved	170h –	1C4h

18.12 Device-Specific Registers

The registers defined in Section 13.13, "Device-Specific Registers," are unique to the PEX 8508 device and not referenced in the *PCI Express Base r1.1*. These registers are also applicable to the NT Port Link Interface, except as defined in Table 18-14 through Table 18-18 and Register 18-37 through Register 18-49. Table 18-14 defines the register map used by the NT Port Link Interface.

Note: *This register group is accessed using a Memory-Mapped cycle.* ***It is recommended that these register values not be changed.***

Table 18-14. Device-Specific NT Port Link Interface Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
NT Port Link Interface Error Checking and Debug Registers	1C8h … 1FCh
NT Port Link Interface Physical Layer Registers	200h … 28Ch
Reserved	290h – 65Ch
NT Port Link Interface Ingress Control Register	660h … 67Ch
Reserved	680h – BFCh
Internal Credit Handler (ITCH) VC&T Threshold Registers	C00h … C08h
NT Port Link Interface Virtual Channel Queue Status Registers	C0Ch … C38h

18.12.1 NT Port Link Interface Error Checking and Debug Registers

The registers defined in Section 13.13.1, "Error Checking and Debug Registers," are also applicable to the NT Port Link Interface, except as listed in Table 18-15 (offsets 1C8h, 1D4h, 1D8h, 1E0h, and 1E8h to 1F4h are ***reserved***), which defines the register map used by the NT Port Link Interface, and Register 18-37 through Register 18-40.

Table 18-15. Device-Specific Error Checking and Debug Register Map (Ports[a])

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Reserved				1C8h
Error Handler 32-Bit Error Status ***(Factory Test Only)***				1CCh
Error Handler 32-Bit Error Mask ***(Factory Test Only)***				1D0h
Reserved				1D4h – 1D8h
Debug Control				1DCh
Reserved				1E0h
Egress NT Port Link Interface Control and Status				1E4h
Reserved				1E8h – 1F4h
Reserved		ACK Transmission Latency Limit		1F8h
Reserved				1FCh

a. Certain registers are port-specific, while others are device-specific.

Note: *All errors in register offset 1CCh generate MSI/INTx interrupts, when enabled.*

Register 18-37. 1CCh Error Handler 32-Bit Error Status *(Factory Test Only)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Completion FIFO Overflow Status** 0 = No overflow detected 1 = Completion FIFO Overflow detected when the 4-deep Completion FIFO (ingress) or 2-deep Completion FIFO (egress) overflows	RW1CS	Yes	0
10:1	***Reserved***	RsvdP	No	0h
11	**Credit Update Timeout Status** No useful credit update to make forward progress for 512 ms or 1s (disabled by default). 0 = No Credit Update Timeout detected 1 = Credit Update Timeout completed	RW1CS	Yes	0
12	**INCH Underrun Error** Ingress Credit Underrun. 0 = No error detected 1 = Credit underrun error detected	RW1CS	Yes	0
31:13	***Reserved***	RsvdP	No	0h

Note: *Error logging is enabled in register offset 1D0h by default.*

Register 18-38. 1D0h Error Handler 32-Bit Error Mask *(Factory Test Only)*

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Completion FIFO Overflow Mask** 0 = No effect on reporting activity 1 = *Completion FIFO Overflow Status* bit is masked/disabled	RWS	Yes	1
10:1	***Reserved***	RsvdP	No	0h
11	**Credit Update Timeout Status Mask** 0 = No effect on reporting activity 1 = *Credit Update Timeout Status* bit is masked/disabled	RWS	Yes	1
12	**INCH Underrun Error Mask** 0 = No effect on reporting activity 1 = *INCH Underrun Error* bit is masked/disabled	RWS	Yes	1
31:13	***Reserved***	RsvdP	No	0h

Register 18-39. 1E4h Egress NT Port Link Interface Control and Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
0	**Egress Credit Update Timer Enable** 0 = Disables Egress Credit Update Timer 1 = Enables Egress Credit Update Timer	RW	Yes	0
1	**Egress Credit Timeout Value** 0 = Minimum 512 ms (Maximum 768 ms) 1 = Minimum 1,024 ms (Maximum 1,280 ms)	RW	Yes	0
2	***Reserved***	RW	Yes	0
15:3	***Reserved***	RsvdP	No	0-0h
19:16	**VC&T Encountered Timeout** 0h = VC0 Posted 1h = VC0 Non-Posted 2h = VC0 Completion 3h = VC1 Posted 4h = VC1 Non-Posted 5h = VC1 Completion All other values are ***reserved***.	RO	Yes	0h
31:20	***Reserved***	RsvdP	No	000h

Register 18-40. 1F8h ACK Transmission Latency Limit

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
7:0	**ACK Transmission Latency Limit** ***Note:*** *The value of this register remains 00h.*	RW	Yes	00h
15:8	***Factory Test Only*** Testing bits – must be 00h.	RW	Yes	00h
23:16	**Upper 8 Bits of Replay Timer Limit** Refer to Transparent Mode register, offset 1F8h. ***Note:*** *The value of this register remains FFh.*	RW	Yes	FFh
31:24	***Reserved***	RsvdP	No	00h

18.12.2 NT Port Link Interface Physical Layer Registers

The registers defined in Section 13.13.2, "Physical Layer Registers," are also applicable to the NT Port Link Interface, except as defined in Table 18-16 (offsets 204h through 20Ch and 260h through 28Ch are ***reserved***) and Register 18-41.

Table 18-16. Device-Specific NT Port Link Interface Physical Layer Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Factory Test Only				200h
Reserved				204h – 20Ch
Phy User Test Pattern 0				210h
Phy User Test Pattern 4				214h
Phy User Test Pattern 8				218h
Phy User Test Pattern 12				21Ch
Physical Layer Status		Physical Layer Command		220h
Port Configuration				224h
Physical Layer Test				228h
Physical Layer				22Ch
Physical Layer Port Command				230h
Port 4 Receive Error Count	Port Control	SKIP Ordered-Set Interval		234h
Quad 0 SerDes Diagnostic Data				238h
Quad 1 SerDes Diagnostic Data				23Ch
Quad 2 SerDes Diagnostic Data				240h
Reserved				244h
SerDes Nominal Drive Current Select				248h
SerDes Drive Current Level Select 1				24Ch
SerDes Drive Current Level Select 2				250h
SerDes Drive Equalization Level Select 1				254h
SerDes Drive Equalization Level Select 2				258h
Physical Receive Error Count				25Ch
Reserved				260h – 28Ch

Note: *For configuration purposes, cross-link connections are* ***not supported****.*

Register 18-41. 220h Physical Layer Command and Status

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
Physical Layer Command				
0	**Port Enumerator Enable** 0 = Enumerate is not enabled 1 = Enumerate is enabled	HwInit	No	0
1	**TDM Enable** 0 = TDM is not enabled 1 = TDM is enabled	HwInit	No	0
2	***Reserved***	RW	No	0
3	**Upstream Port as Configuration Master Enable** 0 = Upstream Port Cross-link is not supported 1 = Upstream Port Cross-link is supported	RW	No	0
4	**Downstream Port as Configuration Slave Enable** 0 = Downstream Port Cross-link is not supported 1 = Downstream Port Cross-link is supported	RW	No	0
5	**Lane Reversal Disable** Provides ability to enable or disable lane reversal. 0 = Lane reversal is supported 1 = Lane reversal is not supported ***Note:*** *Lane reversal is not supported on Port 0.*	RW	No	0
6	**PLL Turn-Off Enable**	RWS	No	0
7	***Reserved***	RsvdP	No	0
15:8	**N_FTS Value** Number of Fast Training Sets (N_FTS) value to transmit in training sets.	RW	No	40h
Physical Layer Status				
19:16	***Reserved***	RsvdP	No	0h
22:20	**Number of Ports Enumerated** Number of ports in current configuration.	HwInit	No	000b
23	***Reserved***	RsvdP	No	0
24	**Port 0 or 4 Deskew Buffer Error Status** 1 = Port 0 or 4 Deskew Buffer overflow or underflow	RW1C	No	0
25	**Port 1 Deskew Buffer Error Status** 1 = Port 1 Deskew Buffer overflow or underflow	RW1C	No	0
26	**Port 2 Deskew Buffer Error Status** 1 = Port 2 Deskew Buffer overflow or underflow	RW1C	No	0
27	**Port 3 Deskew Buffer Error Status** 1 = Port 3 Deskew Buffer overflow or underflow	RW1C	No	0
31:28	***Reserved***	RsvdP	No	0h

18.12.3 NT Port Link Interface Ingress Control Register

The NT Port Link Interface **Ingress Control** register is defined in Section 13.13.7, "Ingress Control and Port Enable Registers," with the addition of the *No Snoop Disable* bit. Table 18-17 defines the register map.

Table 18-17. Device-Specific NT Port Link Interface Ingress Control Register Map (Only Port 0)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Ingress Control				660h
Reserved				664h – 67Ch

Register 18-42. 660h Ingress Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
0	**Enable CSR Access by Downstream Devices** Enables acceptance of both Configuration and Memory requests from a Requester that is downstream from a Transparent port, targeting any PEX 8508 registers. 0 = Configuration requests from a downstream device that is targeting PEX 8508 registers are ***not supported***; the downstream port flags an Uncorrectable Error, and, returns a Completion with Unsupported Request (UR) specified in the *Completion Status* field, to the downstream Requester. Only this mode is *PCI Express Base r1.1*-compliant. 1 = Configuration and Memory requests from downstream Requesters, targeting any PEX 8508 registers in any port, are allowed. ***Notes:*** *This bit can be initially set only through the upstream port, the NT Port Link Interface, the I^2C interface, or by the serial EEPROM, to enable register access through downstream Transparent ports; a Requester downstream from a Transparent port cannot set the bit to grant itself (or peers) access to PEX 8508 registers. Configuration requests can access those registers that are defined by PCI-SIG specifications, and generally cannot access device-specific registers other than the NT Port* ***Cursor Mechanism*** *registers (offsets F8h/FCh). Memory requests can access all PEX 8508 registers.* *Configuration requests can access the NT Port to provide indirect access to NT Port offsets above 100h, for Conventional PCI Requesters such as a PCI Master connected to a PCI Express-to-PCI bridge, that cannot generate Configuration requests containing an Extended Register Number. The NT Port Virtual Interface Cursor Mechanism (but not the NT Port Link Interface Cursor Mechanism) can access device-specific registers that exist in the NT Port; if Port 0 is the NT Port, the NT Port Virtual Interface Cursor Mechanism can access the device-specific registers (which exist only in Port 0) and the Station registers (which exist in Port 0 for all enabled ports).*	RW	Yes	0
1	**Disable Unsupported Request Response for *Reserved* Configuration Registers**	RW	Yes	0
23:2	**Ingress Control** ***Factory Test Only***	RW	Yes	0-0h
24	**No Snoop Disable** Forces the packet header No Snoop attribute bit to 0, for all packets transferred between the NT Port Link and Virtual Interfaces (across the NT boundary, in both directions). Can be used to handle cache coherency-related issues in a system. 0 = Disables the No Snoop Disable feature 1 = Enables the No Snoop Disable feature	RW	Yes	0
31:25	**Ingress Control** ***Factory Test Only***	RW	Yes	0-0h

18.12.4 NT Port Link Interface Virtual Channel Queue Status Registers

The registers defined in Section 13.13.15, "Port Virtual Channel Queue Status Registers," are also applicable to the NT Port Link Interface, except as defined in Register 18-43 through Register 18-49. Table 18-18 defines the register map used by the NT Port Link Interface.

Table 18-18. Device-Specific Port Virtual Channel Queue Status Register Map (Only Port 0 and NT Link Interface)

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
Port 0 VC0 Posted and Non-Posted Queue Status				C0Ch
Port 0 VC0 Completion and VC1 Posted Queue Status				C10h
Port 0 VC1 Non-Posted and Completion Queue Status				C14h
Port 1 VC0 Posted and Non-Posted Queue Status				C18h
Port 1 VC0 Completion and VC1 Posted Queue Status				C1Ch
Port 1 VC1 Non-Posted and Completion Queue Status				C20h
Port 2 VC0 Posted and Non-Posted Queue Status				C24h
Port 2 VC0 Completion and VC1 Posted Queue Status				C28h
Port 2 VC1 Non-Posted and Completion Queue Status				C2Ch
Port 3 VC0 Posted and Non-Posted Queue Status				C30h
Port 3 VC0 Completion Queue Status				C34h
Reserved				C38h

Register 18-43. C1Ch Port 1 VC0 Completion and VC1 Posted Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 1 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 1 VC0 queue.	RO	Yes	000h
21:11	**Port 1 VC1 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 1 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 1 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC1 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 18-44. C20h Port 1 VC1 Non-Posted and Completion Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 1 VC1 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 1 VC1 queue.	RO	Yes	000h
21:11	**Port 1 VC1 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 1 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 1 ITCH Disabled VC1 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 1 ITCH Disabled VC1 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 1 VC1 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 18-45. C24h Port 2 VC0 Posted and Non-Posted Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 2 VC0 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 2 VC0 queue,	RO	Yes	000h
21:11	**Port 2 VC0 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 2 VC0 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 2 ITCH Disabled VC0 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC0 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 2 ITCH Disabled VC0 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC0 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 18-46. C28h Port 2 VC0 Completion and VC1 Posted Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 2 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 2 VC0 queue.	RO	Yes	000h
21:11	**Port 2 VC1 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 2 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 2 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 2 ITCH Disabled VC1 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC1 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 18-47. C2Ch Port 2 VC1 Non-Posted and Completion Queue Status

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
10:0	**Port 2 VC1 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 2 VC1 queue.	RO	Yes	000h
21:11	**Port 2 VC1 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 2 VC1 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 2 ITCH Disabled VC1 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC1 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 2 ITCH Disabled VC1 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 2 VC1 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 18-48. C30h Port 3 VC0 Posted and Non-Posted Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 3 VC0 Posted Packets Status** Defines the number of Posted packets waiting to be sent in the Port 3 VC0 queue,	RO	Yes	000h
21:11	**Port 3 VC0 Non-Posted Packets Status** Defines the number of Non-Posted packets waiting to be sent in the Port 3 VC0 queue.	RO	Yes	000h
23:22	***Reserved***	RsvdP	No	00b
27:24	**Port 3 ITCH Disabled VC0 Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 3 VC0 Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	**Port 3 ITCH Disabled VC0 Non-Posted Traffic Counter** Indicates the number of times that ITCH disabled Port 3 VC0 Non-Posted traffic. Fh = Counter is saturated	RW1C	Yes	0h

Register 18-49. C34h Port 3 VC0 Completion Queue Status

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
10:0	**Port 3 VC0 Completion Packets Status** Defines the number of Completion packets waiting to be sent in the Port 3 VC0 queue.	RO	Yes	000h
23:11	***Reserved***	RsvdP	No	0000h
27:24	**Port 3 ITCH Disabled VC0 Completion Traffic Counter** Indicates the number of times that ITCH disabled Port 3 VC0 Completion traffic. Fh = Counter is saturated	RW1C	Yes	0h
31:28	***Reserved***	RsvdP	No	0h

18.13 Non-Transparent Bridging-Specific Registers

Table 18-19 defines the register map of the registers implemented to support the PEX 8508 NT Port Link Interface Non-Transparent Bridging-Specific registers. These registers are accessed by Memory-Mapped access to Port 0 or the NT Port.

Table 18-19. NT Port Link Interface NT Bridging-Specific Register Map

31 30 29 28 27 26 25 24	23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8	7 6 5 4 3 2 1 0	
NT Port Link Interface Memory Address Translation and BAR Limit Registers				C3Ch … C58h
Reserved				C5Ch – DB0h
NT Port Link Interface Receive Lookup Table Entry Registers				DB4h … DF0h
Reserved				DF4h – FB0h

18.13.1 NT Port Link Interface Memory Address Translation and BAR Limit Registers

Table 18-20 defines the register map of the registers implemented to support the PEX 8508 NT Port Link Interface Memory Address Translation and BAR Limit registers. These registers are accessed by Memory-Mapped access to Port 0 or the NT Port.

Table 18-20. NT Port Link Interface Memory Address Translation and BAR Limit Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Memory BAR2/3 Address Translation Lower	C3Ch
Memory BAR2/3 Address Translation Upper	C40h
Memory BAR4/5 Address Translation Lower	C44h
Memory BAR4/5 Address Translation Upper	C48h
Memory BAR2/3 Limit Lower Address	C4Ch
Memory BAR2/3 Limit Upper Address	C50h
Memory BAR4/5 Limit Lower Address	C54h
Memory BAR4/5 Limit Upper Address	C58h

Register 18-50. C3Ch Memory BAR2/3 Address Translation Lower

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR2/3 Base Translation Address[31:12]** NT Port Link Interface Base Translation address when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Enable* bit (offset E8h[31]) is set to 1.	RW	Yes	0-0h

Register 18-51. C40h Memory BAR2/3 Address Translation Upper

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**BAR2/3 Base Translation Address[63:32]** NT Port Link Interface Base Translation upper address when **BAR2/3** is enabled as a 64-bit BAR [**NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is set to 10b].	RW	Yes	0-0h
	Read-Only when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is cleared to 00b.	RO	No	0-0h

Register 18-52. C44h Memory BAR4/5 Address Translation Lower

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR4/5 Base Translation Address[31:12]** NT Port Link Interface Base Translation address when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Enable* bit (offset F0h[31]) is set to 1.	RW	Yes	0-0h

Register 18-53. C48h Memory BAR4/5 Address Translation Upper

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**BAR4/5 Base Translation Address[63:32]** NT Port Link Interface Base Translation upper address **BAR4/5** is enabled as a 64-bit BAR [**NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is set to 10b].	RW	Yes	0-0h
	Read-Only when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is cleared to 00b.	RO	No	0-0h

Register 18-54. C4Ch Memory BAR2/3 Limit Lower Address

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR2/3 Limit[31:0]** Contains the address of the memory window lower limit defined in the **NT Port Link Interface Memory BAR2/3 Setup** register (offset E8h). 1 MB granularity. When the limit is greater than the window size, the limit is ignored.	RW	Yes	000h

Register 18-55. C50h Memory BAR2/3 Limit Upper Address

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**BAR2/3 Limit[63:32]** Contains the address of the memory window upper limit defined in the **NT Port Link Interface Memory BAR3 Setup** register (offset ECh), when the following conditions exist: • **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is set to 10b, and • **NT Port Link Interface Memory BAR3 Setup** register *BAR3 Enable* bit (offset ECh[31]) is set to 1 When the limit is greater than the window size, the limit is ignored.	RW	Yes	0-0h
	Read-Only when the **NT Port Link Interface Memory BAR2/3 Setup** register *BAR2 Type* field (offset E8h[2:1]) is cleared to 00b.	RO	No	0-0h

Register 18-56. C54h Memory BAR4/5 Limit Lower Address

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
11:0	***Reserved***	RsvdP	No	0-0h
31:12	**BAR4/5 Limit[31:0]** Contains the address of the memory window lower limit defined in the **NT Port Link Interface Memory BAR4/5 Setup** register (offset F0h). 1 MB granularity. When the limit is greater than the window size, the limit is ignored.	RW	Yes	0-0h

Register 18-57. C58h Memory BAR4/5 Limit Upper Address

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
31:0	**BAR4/5 Limit[63:32]** Contains the address of the memory window upper limit defined in the **NT Port Link Interface Memory BAR4/5 Setup** register (offset F0h), when the following conditions exist: • **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is set to 10b, and • **NT Port Link Interface Memory BAR5 Setup** register *BAR5 Enable* bit (offset F4h[31]) is set to 1 When the limit is greater than the window size, the limit is ignored.	RW	Yes	0-0h
	Read-Only when the **NT Port Link Interface Memory BAR4/5 Setup** register *BAR4 Type* field (offset F0h[2:1]) is cleared to 00b.	RO	No	0-0h

18.13.2 NT Port Link Interface Receive Lookup Table Entry Registers

This section describes the NT Port Link Interface Receive (Requester ID Translation) Lookup Table (LUT) Entry registers. The NT Port uses these registers for Requester ID translation when it forwards:

- TLP request from NT Port Link Interface to the Virtual Interface, or
- Completion TLP from NT Port Virtual Interface to the Link Interface

When Non-Posted traffic must be sent through the NT Link Interface, program the registers listed in this group with the upstream Requester ID and set the *LUT Entry Enable* bits (bits 0 and 16) for each LUT entry, as needed.

Table 18-21 defines the register and address locations. The register descriptions that follow the table define the bit definitions that apply to the registers.

Note: *Writes to NT Port **Link Interface Receive Lookup Table Entry** registers (offsets DB4h through DF0h) are automatically copied to the same offsets in the NT Port Virtual Interface. If these NT Port Link Interface registers are programmed by serial EEPROM, the same data values must be programmed into the same register offsets in the NT Port Virtual Interface (that is, the values in serial EEPROM DWord addresses 61Fh through 62Eh must also be programmed into serial EEPROM DWord addresses 717h through 726h, respectively).*

Table 18-21. Link Interface Receive Lookup Table Entry_*n_m* Register Locations

ADDR Location	Lookup Table Entry_*n_m*	ADDR Location	Lookup Table Entry_*n_m*
DB4h	0_1	DD4h	16_17
DB8h	2_3	DD8h	18_19
DBCh	4_5	DDCh	20_21
DC0h	6_7	DE0h	22_23
DC4h	8_9	DE4h	24_25
DC8h	10_11	DE8h	26_27
DCCh	12_13	DECh	28_29
DD0h	14_15	DF0h	30_31

Register 18-58. DB4h - DF0h Link Interface Receive Lookup Table Entry_*n_m*
(where *n_m* = 0_1 through 30_31)

Bit(s)	Description		Type	Serial EEPROM and I²C	Default
0	**LUT Entry_*n* Enable** 0 = Disables 1 = Enables		RW	Yes	0
2:1	***Reserved***		RsvdP	No	00b
7:3	**Requester ID**	**LUT Entry_*n* Device Number**	RW	Yes	0000_0b
15:8		**LUT Entry_*n* Bus Number**	RW	Yes	00h
16	**LUT Entry_*m* Enable** 0 = Disables 1 = Enables		RW	Yes	0
18:17	***Reserved***		RsvdP	No	00b
23:19	**LUT Entry_*m* Device Number**		RW	Yes	0000_0b
31:24	**LUT Entry_*m* Bus Number**		RW	Yes	00h

18.14 Advanced Error Reporting Capability Registers

The registers defined in Section 13.14, "Advanced Error Reporting Capability Registers," are also applicable to the NT Port Link Interface, except as defined in Register 18-59 through Register 18-61. Table 18-22 defines the register map used by the NT Port Virtual and Link Interfaces.

Table 18-22. Advanced Error Reporting Capability Register Map

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16	15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	
Next Capability Offset (138h)	Capability Version (1h)	PCI Express Extended Capability ID (0001h)	FB4h
Uncorrectable Error Status[a]		FB8h
Uncorrectable Error Mask[a]		FBCh
Uncorrectable Error Severity[a]		FC0h
Correctable Error Status		FC4h
Correctable Error Mask		FC8h
Advanced Error Capabilities and Control		FCCh
Header Log 0		FD0h
Header Log 1		FD4h
Header Log 2		FD8h
Header Log 3		FDCh
Reserved		FE0h – FFCh

a. *Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 18-59. FB8h Uncorrectable Error Status

Bit(s)	Description	Type	Serial EEPROM and I^2C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
5	**Surprise Down Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
11:6	***Reserved***	RsvdP	No	0000_00b
12	**Poisoned TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
13	**Flow Control Protocol Error Status**	RsvdP[a]	No	0
14	**Completion Timeout Status**	RsvdP[a]	No	0
15	**Completer Abort Status**	RsvdP[a]	No	0
16	**Unexpected Completion Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
17	**Receiver Overflow Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
18	**Malformed TLP Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
19	**ECRC Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
20	**Unsupported Request Error Status** 0 = No error detected 1 = Error detected	RW1CS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 18-60. FBCh Uncorrectable Error Mask

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
5	**Surprise Down Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
11:6	***Reserved***	RsvdP	No	0000_00b
12	**Poisoned TLP Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
13	**Flow Control Protocol Error Mask**	RsvdP[a]	No	0
14	**Completion Timeout Mask**	RsvdP[a]	No	0
15	**Completer Abort Mask**	RsvdP[a]	No	0
16	**Unexpected Completion Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
17	**Receiver Overflow Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
18	**Malformed TLP Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
19	**ECRC Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
20	**Unsupported Request Error Mask** 0 = No mask is set 1 = Error reporting, first error update, and header logging are masked for this error	RWS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Register 18-61. FC0h Uncorrectable Error Severity

Bit(s)	Description	Type	Serial EEPROM and I²C	Default
3:0	***Reserved***	RsvdP	No	0h
4	**Data Link Protocol Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
5	**Surprise Down Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
11:6	***Reserved***	RsvdP	No	0-0h
12	**Poisoned TLP Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
13	**Flow Control Protocol Error Severity**	RsvdP[a]	No	0
14	**Completion Timeout Severity**	RsvdP[a]	No	0
15	**Completer Abort Severity**	RsvdP[a]	No	0
16	**Unexpected Completion Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
17	**Receiver Overflow Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
18	**Malformed TLP Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	1
19	**ECRC Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
20	**Unsupported Request Error Severity** 0 = Error reported as non-fatal 1 = Error reported as fatal	RWS	Yes	0
31:21	***Reserved***	RsvdP	No	0-0h

*a. Bits [15, 13] are **reserved** or **not supported**. Additionally, bit 14 is not applicable to switches; therefore, the PCI Express Base r1.1 provides exemption from supporting bit 14.*

Chapter 19 Dual Clocking Support

19.1 Introduction

The Dual Clocking feature allows the PEX 8508 to simultaneously support Spread Spectrum Clocking (SSC) and Constant Frequency Clocking (CFC). This capability is necessary for systems in which the upstream port connects to the SSC source, and the downstream ports connect to links running on CFC domains. The upstream port (default is Port 0) must be connected to an SSC domain and downstream ports must be connected to a CFC domain, as illustrated in Figure 19-1.

Most systems support SSC; however, the *PCI Express Base r1.1* specifies that REFCLKs of two interconnected devices must be within a ±300 ppm frequency tolerance of one another. Therefore, a device on a system that supports SSC clocking cannot be connected to a device on another system that does not use SSC. If a system supports SSC, all devices in the PCI Express interconnect must have the same Reference Clock source, to allow traffic to communicate between them. However, this is not practical in the case of chassis-to-chassis or blade-to-blade PCI Express traffic. If each chassis or blade has its own SSC source, special clock isolation circuitry is required to enable such configurations to communicate between themselves. For this reason, the PEX 8508 supports two REFCLK inputs that allow traffic to move from the SSC clock domain on the upstream port to a CFC clock domain on the downstream ports, thereby allowing two different chassis or blades to communicate with one another.

The SSC domain runs at a frequency between 99.5 and 100 MHz. The CFC domain runs at a frequency of 100 MHz. Because the SSC and CFC domain average frequencies are not identical, the PEX 8508 provides buffering and flow control to ensure that data does not overflow nor underflow as it passes from one clock domain to the other. Dual Clocking support is provided in fan-in/fan-out configurations when the upstream port is set to a x4 port width.

Figure 19-1. PEX 8508 in a Dual Clock System

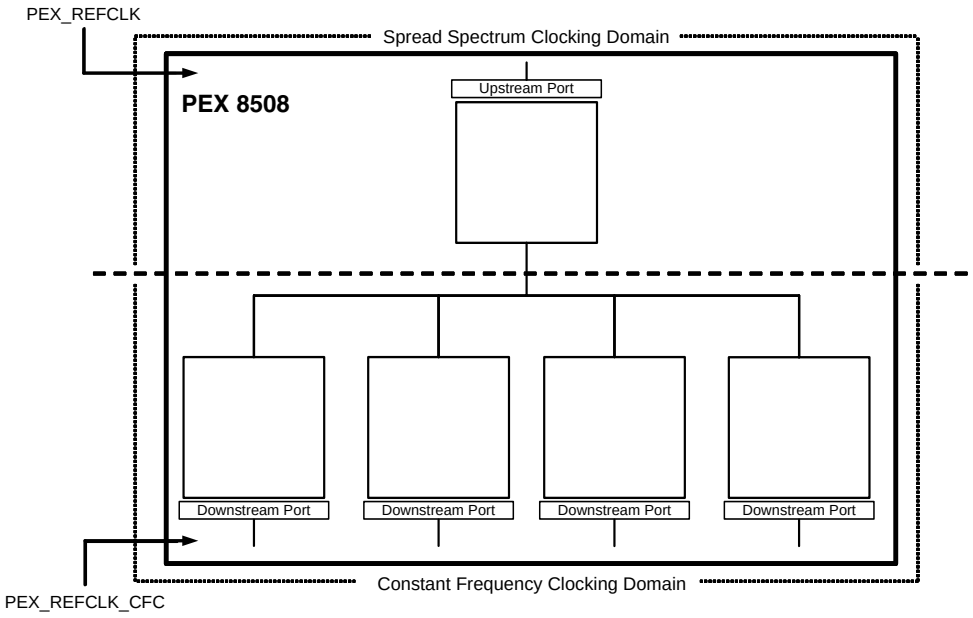

19.2 Dual Clocking Operation

Dual Clocking is enabled when the following conditions are met:

- Strap ball STRAP_SSC_XING_ENA is pulled High to VDD33A
- Selected port configuration assigns Port 0 as the upstream port with a x4 lane width.

When Dual Clocking is enabled, the REFCLK balls are mapped as follows:

- PEX_REFCLKn and PEX_REFCLK_CFCp signal balls become the SSC domain Clock signals
- PEX_REFCLK_CFCn and PEX_REFCLK_CFCp signal balls become the CFC domain Clock signals

If the system design does not utilize Dual Clocking, the STRAP_SSC_XING_ENA ball must be pulled Low to VSS and the PEX_REFCLK_CFCn and PEX_REFCLK_CFCp signals must remain floating. If the PEX 8508 utilizes SSC on its upstream port, or the upstream port is not configured to Port 0, the PEX_REFCLKn and PEX_REFCLK_CFCp signals must be shared by all downstream devices. If the PEX 8508 utilizes a Constant Frequency source on its upstream port, each downstream device can use an independent CFC REFCLK source, provided that it meets the *PCI Express Base r1.1*-defined PCI Express frequency tolerance of ±300 ppm.

Chapter 20 I²C Interface Operation

20.1 I²C Support Overview

Inter-Integrated Circuit (I^2C) is a bus used to connect Integrated Circuits (ICs). Multiple ICs are connected to an I^2C Bus and each IC can act as a Master by initiating a Data transfer. I^2C is used for Data transfers between ICs at relatively low rates (100 Kbps to 3.4 Mbps) and is used in a variety of applications. For further details regarding I^2C Buses, refer to the *I2C Bus, v2.1*.

The PEX 8508 is an I^2C Slave. Slave operations allow the PEX 8508 Configuration registers to be read from or written to by an I^2C Master, external from the device. I^2C is a sideband mechanism that allows the device Configuration registers to be programmed, read from, or written to, independently of the PCI Express upstream link.

In the past, either a serial EEPROM was required, or the PEX 8508 Configuration registers could only be accessed through a working (successful linkup) PCI Express upstream link. With I^2C, users now have the option of programming the Configuration Space registers through the I^2C interface. This is useful for debugging purposes, if the PEX 8508 upstream port fails to link up. I^2C also provides an alternative to using a serial EEPROM.

I^2C operation is supported in Transparent and Non-Transparent modes. Registers for Transparent and Non-Transparent ports can be accessed through I^2C operations.

Figure 20-1. Standard Devices to I²C Bus Connection Block Diagram

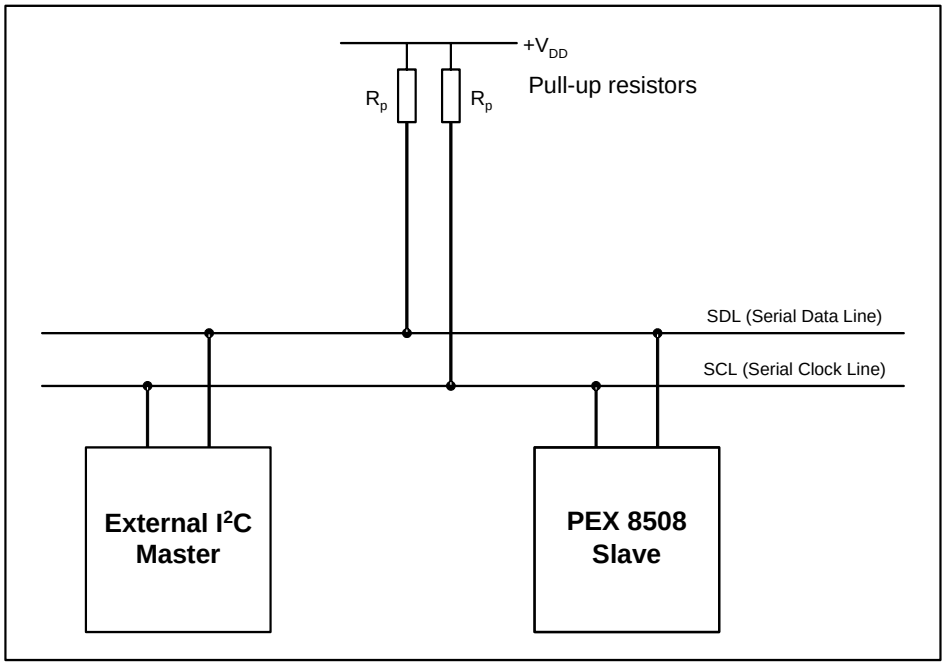

20.2 I2C Addressing – Slave Mode Access

To access the PEX 8508 Configuration registers through the I2C interface, the PEX 8508 I2C Slave address must be configured.

The PEX 8508 supports a 7-bit I2C Slave address. The 7-bit I2C Address bits can be configured from the PCI Express side, in the **I2C Configuration** register (offset 294h, default value 70h), with the lower three bits of the address strapped through the I2C_ADDR[2:0] balls.

The I2C_ADDR[2:0] balls can be pulled High or Low to select a different Slave address. Up to eight PEX 8508 devices can share the same I2C Bus segment without conflict, provided each PEX 8508 switch has its set of I2C_ADDR[2:0] inputs strapped to a unique combination.

20.3 Command Phase Format

An I2C transfer starts as a packet with Address Phase bytes, followed by four Command Phase bytes, and one or more Data Phase bytes. The I2C packet Address Phase Byte format is illustrated in Figure 20-2a. The Command Phase portion must contain 4 bytes of data. The Command phase bytes contain the following:

- I2C Transfer type (Read/Write)
- PCI Express Configuration Register address
- PEX 8508 Port Number being accessed
- Byte Enable(s) of the register data being accessed

When the I2C Master is writing to the PEX 8508, the I2C Master must transmit the Data bytes to be written to that register within the same packet that contains the Command bytes.

When the I2C Master is reading from the PEX 8508, the I2C Master must separately transmit a Command Phase packet and Data Phase packet.

Each I2C packet must contain 4 bytes of data. Pad unused packet Data bytes with zeros (0) to meet this requirement.

Table 20-1 describes each Command Phase byte. Figure 20-2b illustrates the Command phase portion of an I2C Write packet.

Table 20-1. Command Phase Format

Byte	Bit(s)	Description
1st (0)	7:3	***Reserved*** Must be cleared to 0.
	2:0	**Command** 011b = Write register 100b = Read register All other values are ***reserved***. ***Do not use.***
2nd (1)	7:3	***Reserved*** Must be cleared to 0.
	2	1 = Virtual Interface CSR access (valid only in NT mode)
	1:0	**Port Selector, Bits [2:1]** 2nd Byte, bits [1:0], and 3rd Byte, bit 7, combine to form a 3-bit Port Selector.
3rd (2)	7	**Port Selector, Bit 0** 2nd Byte, bits [1:0], and 3rd Byte, bit 7, combine to form a 3-bit Port Selector. *Port Selector, bits [2:1]* (2nd Byte, bits [1:0]) select the port to access – only values 0h, 1h, 2h, 3h, and 4h are valid.
	6	***Reserved*** Must be cleared to 0.
	5:2	1 = Indicates corresponding PEX 8508 register byte is modified **Bit** **Description** 2 Byte Enable for Byte 0 (PEX 8508 register bits [7:0]) 3 Byte Enable for Byte 1 (PEX 8508 register bits [15:8]) 4 Byte Enable for Byte 2 (PEX 8508 register bits [23:16]) 5 Byte Enable for Byte 3 (PEX 8508 register bits [31:24])
	1:0	**PEX 8508 Register Address [11:10]**
4th (3)	7:0	**PEX 8508 Register Address [9:2]** ***Note:*** *All register addresses are DWord-aligned. Therefore, bits [1:0] are always cleared to 00b.*

20.4 I2C Interface Register

The **I2C Interface** register, **I2C Configuration**, is defined in Section 13.13.3, "I2C Interface Register." This register is accessible only from Port 0. The default I2C Slave address can be changed in the **I2C Configuration** register (offset 294h) to a different value. Changes to the I2C Slave address should be done using the serial EEPROM or a Memory Write. **The I2C Slave address must not be changed by an I2C Write command.** (Refer to Section 20.2.)

Other I2C interface registers exist; however, they are for ***Factory Test Only***.

20.5 PEX 8508 I2C Register Write Access

The PEX 8508 Configuration registers can be read from and written to, based upon I2C register Read and Write operations, respectively. An I2C Write packet consists of Address Phase bytes and Command Phase bytes, followed by one to four additional I2C Data bytes. Table 20-2 defines mapping of the I2C Data bytes to the Configuration register Data bytes.

Figure 20-2c illustrates the I2C Data byte format. The I2C packet starts with the "S" (START condition) bit. Data bytes are separated by the "A" (ACKNOWLEDGE) or "N" (NOT ACKNOWLEDGE) bit. The packet ends with the "P" (STOP condition) bit.

If the Master generates an invalid command, the targeted PEX 8508 register is not modified.

The PEX 8508 considers the 1st Data byte (register Byte 3) of the Data packet, after the four Command bytes in the Command phase are transmitted. This is independent of the Byte Enable settings in the Command phase. If additional I2C bytes are present in the Data phase, these bytes access register Bytes 2 through 0, respectively, regardless of the Byte Enable settings in the Command phase. After the 8th byte of an I2C Write transfer, the PEX 8508 generates a NAK (Not Acknowledge). The Master can then generate either a STOP condition (to finish the transfer) or a repeated START condition (to start a new transfer). (For further details regarding ACK/NAK protocol, refer to the *I2C Bus, v2.1*.)

In the packet described in Figure 20-2, Command Bytes 0 through 3 follow the format specified in Table 20-1.

Table 20-2. I2C Register Write Access

I2C Data Byte Order	PCI Express Configuration Register Bytes
0	Written to register Byte 3
1	Written to register Byte 2
2	Written to register Byte 1
3	Written to register Byte 0

Figure 20-2. I2C Write Packet

Figure 20-2.a I2C Write Packet Address Phase Bytes

1st Cycle					
START	7 6 5 4 3 2 1	0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK
S	Slave Address[7:1]	Read/Write Bit 0 = Write	A	*Reserved*	A

Figure 20-2.b I2C Write Packet Command Phase Bytes

Command Cycle							
7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK
Command Byte 0	A	Command Byte 1	A	Command Byte 2	A	Command Byte 3	A

Figure 20-2.c I2C Write Packet Data Phase Bytes

Write Cycle								
7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	STOP
Data Byte 0 (to selected register Byte 3)	A	Data Byte 1 (to selected register Byte 2)	A	Data Byte 2 (to selected register Byte 1)	A	Data Byte 3 (to selected register Byte 0)	A	P

20.5.1 Register Write

The following tables illustrate a sample I2C packet for writing the PEX 8508 **Message Upper Address[63:32]** register (offset 50h) for Port 3, with data 1234_5678h.

Note: *Assume that the PEX 8508 has a default I2C Slave address of 70h, with the I2C_ADDR[2:0] balls all having a value of 0. The byte sequence on the I2C Bus, as listed in the following tables, occurs after the START and before the STOP bits are set in the packet.*

Table 20-3. I2C Register Write Access Example – 1st Cycle

Phase	Value	Description
Address	E0h	**Bits [7:1] for PEX 8508 I2C Slave Address (70h)** Last bit (bit 0) for Write = 0.

Table 20-4. I2C Register Write Access Example – Command Cycle

Byte	Value	Description
0	03h	[7:3] ***Reserved*** Must be cleared to 0. [2:0] **Register Selector Bits** 011b = Write
1	01h	[7:3] ***Reserved*** Must be cleared to 0. 2 1 = Virtual Interface CSR access (valid in NT mode only). [1:0] **Port Selector, Bits [2:1]**
2	BCh	7 **Port Selector, Bit 0** 6 ***Reserved*** Must be cleared to 0. [5:2] **Byte Enables** All active. [1:0] **Register Address, Bits [11:10]**
3	14h	[7:0] **Register Address, Bits [9:2]**

Table 20-5. I2C Register Write Access Example – Write Cycle

Byte	Value	Description
0	12h	**Data to Write for Byte 3**
1	34h	**Data to Write for Byte 2**
2	56h	**Data to Write for Byte 1**
3	78h	**Data to Write for Byte 0**

Figure 20-3. I²C Write Command Packet Example

Figure 20-3.a I²C Write Packet Address Phase Bytes

1st Cycle					
START	**7 6 5 4 3 2 1**	**0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**
S	Slave Address 1110_000b	Read/Write Bit 0 0 = Write	A	***Reserved*** 0000_0000b	A

Figure 20-3.b I²C Write Packet Command Phase Bytes

Command Cycle							
7 6 5 4 3 2 1 0	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**
Command Byte 0 0000_0011b	A	Command Byte 1 0000_0001b	A	Command Byte 2 1011_1100b	A	Command Byte 3 0001_0100b	A

Figure 20-3.c I²C Write Packet Data Phase Bytes

Write Cycle								
7 6 5 4 3 2 1 0	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**STOP**
Data Byte 0 0001_0010b	A	Data Byte 1 0011_0100b	A	Data Byte 2 0101_0110b	A	Data Byte 3 0111_1000b	A	P

20.6 PEX 8508 I2C Register Read Access

When the I2C Master attempts to read a PEX 8508 register, two packets are transmitted. The 1st packet consists of Address and Command Phase bytes to the Slave. The 2nd packet consists of Address and Data Phase bytes.

According to the *I2C Bus v2.1*, a Read cycle is triggered when the Read/Write bit (bit 0) of the 1st cycle is set to 1. The Command phase reads the requested register content into the internal buffer. When the I2C Read access occurs, the internal buffer value is transferred on to the I2C Bus, starting from Byte 3 (bits [31:24]), followed by the subsequent bytes, with Byte 0 (bits [7:0]) being transferred last. If the I2C Master requests more than four bytes, the PEX 8508 re-transmits the same byte sequence, starting from Byte 3 of the internal buffer.

The 1st and 2nd I2C Read packets (illustrated in Figure 20-4 and Figure 20-5, respectively) perform the following functions:

- **1st packet** – Selects the register to read
- **2nd packet** – Reads the register (sample 2nd packet provided is for a 7-bit PEX 8508 I2C Slave address)

Although two packets are shown for the I2C Read, the I2C Master can merge the two packets together into a single packet, by not generating the STOP at the end of the first packet (Master does not relinquish the bus) and generating REPEAT START.

Figure 20-4. I2C Read Command Packet (1st Packet)

Figure 20-4.a I2C Read Command Packet Address Phase Bytes

1st Cycle					
START	7 6 5 4 3 2 1	0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK
S	Slave Address[7:1]	Read/Write Bit Write = 0	A	*Reserved*	A

Figure 20-4.b I2C Read Command Packet Command Phase Bytes

Command Cycle							
7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	STOP
Command Byte 0	A	Command Byte 1	A	Command Byte 2	A	Command Byte 3	P

Figure 20-5. I2C Read Data Packet (2nd Packet)

Figure 20-5.a I2C Read Data Packet Address Phase Bytes

1st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address[7:1]	Read/Write Bit, 1 = Read	A

Figure 20-5.b I2C Read Data Packet Data Phase Bytes

Read Cycle							
7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	ACK/NAK	7 6 5 4 3 2 1 0	STOP
Buffer Byte 3	A	Buffer Byte 2	A	Buffer Byte 1	A	Buffer Byte 0	P

20.6.1 Register Read Address Phase and Command Packet

The following is a sample I2C packet for reading the PEX 8508 **Serial EEPROM Data Buffer** register (Port 0, offset 264h) for Port 4, assuming the register value is ABCD_EF01h.

Note: *Assume that the PEX 8508 has a default I2C Slave address of 70h, with the I2C_ADDR[2:0] balls all having a value of 0. The byte sequence on the I2C Bus, as listed in the following tables, occurs after the START and before the STOP bits are set in the packet.*

Table 20-6. I2C Register Read Access Example – 1st Cycle

Phase	Value	Description
Address	E0h	**Bits [7:1] for PEX 8508 I2C Slave Address (70h)** Last bit (bit 0) for Write = 0.

Table 20-7. I2C Register Read Access Example – Command Cycle

Byte	Value	Description
0	04h	[7:3] ***Reserved*** Must be cleared to 0. [2:0] **Command Bits** 100b = Read
1	02h	[7:3] ***Reserved*** Must be cleared to 0. 2 1 = Virtual Interface CSR access (valid in NT mode only). [1:0] **Port Selector, Bits [2:1]**
2	3Ch	7 **Port Selector, Bit 0** 6 ***Reserved*** Must be cleared to 0. [5:2] **Byte Enables** All active. [1:0] **Register Address, Bits [11:10]**
3	99h	[7:0] **Register Address, Bits [9:2]**

20.6.2 Register Read Data Packet

Note: *Assume that the PEX 8508 has a default I²C Slave address of 70h, with the I2C_ADDR[2:0] balls all having a value of 0. The byte sequence on the I²C Bus, as listed in the following tables, occurs after the START and before the STOP bits are set in the packet.*

Table 20-8. I²C Register Read Access Example – 1st Cycle

Phase	Value	Description
Address	E1h	**Bits [7:1] for PEX 8508 I²C Slave Address (70h)** Last bit (bit 0) for Read = 1.
Read	ABh	**Byte 3 of Register Read**
	CDh	**Byte 2 of Register Read**
	EFh	**Byte 1 of Register Read**
	01h	**Byte 0 of Register Read**

Figure 20-6. 1st Packet – Command Phase

1st Cycle					
START	**7 6 5 4 3 2 1**	**0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**
S	Slave Address 1110_000b	Read/Write Bit 0 = Write	A	***Reserved*** 0000_0000b	A

Command Cycle							
7 6 5 4 3 2 1 0	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**STOP**
Command Byte 0 0000_0100b	A	Command Byte 1 0000_0010b	A	Command Byte 2 0011_1110b	A	Command Byte 3 1001_1001b	P

Figure 20-7. 2nd Packet – Read Phase

1st Cycle			
START	**7 6 5 4 3 2 1**	**0**	**ACK/NAK**
S	Slave Address[7:1] 1110_000b	Read/Write Bit 1 = Read	A

Read Cycle							
7 6 5 4 3 2 1 0	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**ACK/NAK**	**7 6 5 4 3 2 1 0**	**STOP**
Buffer Byte 3 1010_1011b	A	Buffer Byte 2 1100_1101b	A	Buffer Byte 1 1110_1111b	A	Buffer Byte 0 0000_0001b	P

THIS PAGE INTENTIONALLY LEFT BLANK.

Chapter 21 WAKE# and Beacon Functionality

21.1 Overview

The PEX 8508 supports auxiliary power (Aux) during the L2 (D3cold) power state. During the L2 power state with Aux power applied, the following conditions exist:

- Main power is not present
- Reference Clock is not present
- Internal PLL is powered down
- No internal clocks are present

For a PCI Express system, two mechanisms are used to wake up the system from the L2 power state. The first is out-of-band WAKE# signal assertion. The second is an in-bank low-frequency Beacon transmitted upstream.

For the PEX 8508, WAKE# signal assertion causes a Beacon signal to transmit on Lane 0 of the upstream port. The Beacon signal is generated only while the WAKE# input remains asserted; otherwise, the signal state is not internally latched. Beacon signal detection on a downstream port also causes Beacon signal transmission on Lane 0 of the upstream port.

For further details, refer to the *PCI Express Base r1.1*.

21.2 WAKE#-to-Beacon Translation

When the PEX 8508 WAKE# signal is asserted, the following sequence of events occur:

1. PEX 8508 detects that the WAKE# signal is asserted.
2. Upstream port SerDes Lane 0 transmits the Beacon signal upstream.
3. System restores main power to the PEX 8508.
4. System de-asserts the PEX_PERST# signal.
5. After the links are up, system software restores the context saved before entering the D3cold state.

The PEX 8508 supports full power management, including support for D3cold (L2 Link PM state). From a system design perspective, three elements must be considered when implementing PEX 8508 support for D3cold – auxiliary power, WAKE#, and Beacon.

Prior to the removal of main power and the Reference Clock, software must perform the necessary protocol to place the PEX 8508 into the L2/L3 Ready state. The system is expected to assert PEX_PERST# prior to removing the Reference Clock and shutting off main power. Three PEX 8508 supply rails must remain powered in the L2 link state – VDD10X, VDD33X, and VTT_PEX.

For add-in board designs, while main power is shut off, all three supplies are powered from the board-edge connector's +3.3VAUX supply pin. The VDD10X and VDD33X supplies utilize small amounts of power in all links states (L0, L0s, L1, and L2); Therefore, they can be permanently powered by the +3.3 VAUX supply rail. The VTT_PEX supply provides a termination voltage for the PCI Express transmitters and therefore utilizes larger amounts of power in the L0 state, as compared to the L2 state. It is recommended to use a smart-OR power switch, *such as* the California Micro Devices CMPWR025, to switch to VTT_PEX power when main power is shut down. All three supplies can be implemented in the same manner.

The PEX 8508 supports link wakeup by WAKE# assertion or Beacon propagation. The PEX 8508 generates a Beacon signal from Lane 0 of its upstream port when WAKE# input is asserted or a Beacon signal is received on any downstream port lane. In the former case, Beacon is generated only while the WAKE# input remains asserted; therefore, the signal is not internally latched.

21.3 Beacon Detection

When the PEX 8508 detects a Beacon signal on a downstream port, the following sequence of events occur:

1. PEX 8508 detects the Beacon signal.
2. Upstream port SerDes Lane 0 transmits the Beacon signal upstream.
3. System restores main power to the PEX 8508.
4. System de-asserts the PEX_PERST# signal.
5. After the links are up, system software restores the context saved before entering the D3cold state.

Chapter 22 Test and Debug

22.1 Physical Layer Loop-Back Operation

22.1.1 Overview

Physical layer loop-back functions are used to test SerDes in the PEX 8508, connections between devices, and SerDes of external devices, as well as various PEX 8508 and external digital logic. The PEX 8508 supports five types of loop-back operations, as described in Table 22-1. Additional information regarding each type is provided in the sections that follow.

Table 22-1. Loop-Back Operations

Operation	Description
Internal Loop-Back Mode	Internal Loop-Back mode connects SerDes serial Tx output to serial Rx input. The Pseudo-Random Bit Sequence (PRBS) generator is used to create a pseudo-random data pattern that is transmitted and returned to the PRBS checker.
Analog Loop-Back Master Mode	Analog Loop-Back Master mode depends upon an external device or dumb connection (*such as* a cable) to loop back the transmitted data to the PEX 8508. If an external device is used, it must not include its elastic buffer in the loop-back data path, because no SKIP Ordered-Sets are transmitted. Use the PRBS generator and checker to create and check the data pattern. The PEX 8508 enters Analog Loop-Back Master mode when the **Physical Layer Port Command** register *Port x Loop-Back Command* bit (Port 0, offset 230h[0, 4, 8, 12, and/or 16]) is set.
Digital Loop-Back Master Mode	As with the Analog Loop-Back Master mode, Digital Loop-Back Master mode depends upon an external device to loop back the transmitted data. This method is best utilized with an external device that includes at least its elastic buffer in the loop-back data path. The PEX 8508 provides a programmable data pattern generator and checker that inserts the SKIP Ordered-Set at the proper intervals. The PEX 8508 enters Digital Loop-Back Master mode when the **Physical Layer Port Command** register *Port x Loop-Back Command* bit (Port 0, offset 230h[0, 4, 8, 12, and/or 16]) is set.
Analog Loop-Back Slave Mode	The PEX 8508 enters Analog Loop-Back Slave mode when an external device transmits training sets with the *Loop-Back Training Control* bit set and the **Physical Layer Test** register *Analog Loop-Back Enable* bit (Port 0, offset 228h[4]) is set. The received data is looped back from the SerDes 10-bit receive interface to the 10-bit transmit interface. All digital logic is excluded from the loop-back data path.
Digital Loop-Back Slave Mode	The PEX 8508 enters Digital Loop-Back Slave mode when an external device transmits training sets with the *Loop-Back Training Control* bit set and the **Physical Layer Test** register *Analog Loop-Back Enable* bit (Port 0, offset 228h[4]) is cleared. In this mode, the data is looped back at an 8-bit level, which includes the PEX 8508's elastic buffer, 8b/10b decoder, and 8b/10b encoder in the loop-back data path.

22.1.2 Internal Loop-Back Mode

Figure 22-1 illustrates the Loop-Back data path when Internal Loop-Back mode is enabled. The only items in the data path are the serializer and de-serializer. Internal Loop-Back mode is used when SerDes Built-In Self-Test (BIST) is enabled (**Physical Layer Test** register *SerDes BIST Enable* bit (Port 0, offset 228h[7]=1).

The SerDes BIST is intended to overlap with the serial EEPROM load operation. To achieve this overlap, the *SerDes BIST Enable* bit is written early in the serial EEPROM load operation. After the *SerDes BIST Enable* bit is set, the SerDes is placed in Internal Loop-Back mode and the PRBS generator is started. The BIST is run for 15 ms; if an error is detected on a SerDes, the BIST_ERROR bit associated with the station that includes the SerDes in error is asserted. While the SerDes BIST is in progress, the PRBS test data is present on the external TxP and TxN balls. The continuing Serial EEPROM register load has no effect on the SerDes BIST.

Figure 22-1. Internal Loop-Back Mode (Analog near End) Data Path

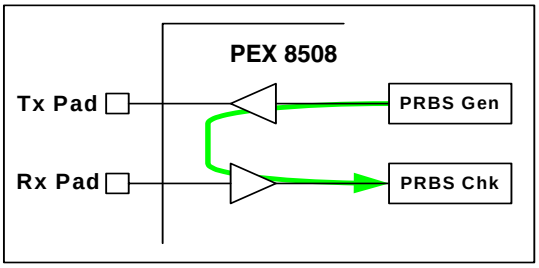

22.1.3 Analog Loop-Back Master Mode

Analog Loop-Back Master mode is typically used for Analog Far-End testing, as illustrated in Figure 22-2.

The mode can also be used to re-create the previously described BIST, by looping back the data with a cable. Looping back with a cable includes the internal bond, external balls, board trace, and connectors in the test data path, as illustrated in Figure 22-3.

Figure 22-2. Analog Far-End Loop-Back

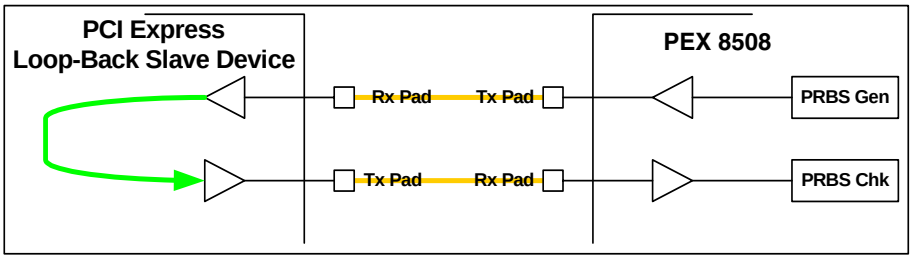

Figure 22-3. Cable Loop-Back

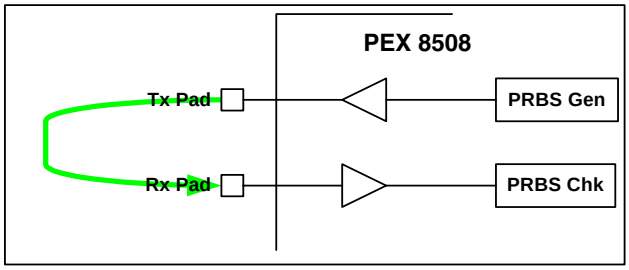

To cause a PEX 8508 port to request to become a Loop-Back Master:

1. After the link is up, a Configuration Write to the appropriate **Physical Layer Port Command** register *Port Loop-Back Command* bit (Port 0, offset 230h[0, 4, 8, 12, and/or 16]), causes the port to transition from the L0 state to Recovery, and then to the Loop-Back state:
   - If a cable is used for a loop-back, the port transitions from the Configuration state to the Loop-Back state. Connect the cable only after the upstream link is up and Configuration Writes are possible.
   - If the cable is connected before the upstream device is able to set the **Physical Layer Test** register *PRBS External Loop-Back* bit (Port 0, offset 228h[22:20]), the link with the cable can reach the L0 state and not go to the Loop-Back state.
   - Cable length is limited only by the PCI Express drivers and cable properties.
2. After the port is in the Loop-Back state, the corresponding **Physical Layer Port Command** register *Port x Ready as Loop-Back Master* bit (Port 0, offset 230h[3, 7, 11, 15, and/or 19]) is set:
   - At this time, the PRBS engine can be enabled by setting the **Physical Layer Test** register *PRBS Enable* bit (Port 0, offset 228h[18:16]) associated with the SerDes assigned to the port being tested.
   - The PRBS checker checks the returned PRBS data. Any errors are logged in the **Quad *x* SerDes Diagnostic Data** register (Port 0, offsets 238h through 240h) that corresponds to the SerDes quad being tested.

22.1.4 Digital Loop-Back Master Mode

The only difference between Analog and Digital Loop-Back Master modes is that the external device is assumed to retain, to some extent, its digital logic in the loop-back data path. Because this includes the elastic buffer, SKIP Ordered-Sets must be included in the test pattern. For the PEX 8508, this precludes PRBS engine use, because the PRBS generator does not generate SKIP Ordered-Sets.

The PEX 8508 provides the programmable test pattern transmitter for Digital Far-End Loop-Back testing, as illustrated in Figure 22-4. After Digital Loop-Back Master mode is established, Configuration Writes are used to fill the **Phy User Test Pattern *x*** registers (Port 0, offsets 210h through 21Ch). The corresponding **Physical Layer Test** register *User Test Pattern Enable* bit(s) (Port 0, offset 228h[30:28]) is set, which starts the transmission of the test pattern on all lanes. If one or more of the **Physical Layer Test** register *PRBS Enable* bits (Port 0, offset 228h[18:16]) are also set, the test pattern is transmitted on all lanes of the corresponding port, regardless of the port's width. However, if the *PRBS Enable* bit is cleared, the test pattern is transmitted only on the corresponding SerDes quad lanes.

SKIP Ordered-Sets are inserted at an interval determined by the value in the **SKIP Ordered-Set Interval** register *SKIP Ordered-Set Interval* field (Port 0, offset 234h[11:0]) (default value is 1,180 symbol times), at the nearest data pattern boundary.

The Test Pattern checker ignores SKIP Ordered-Sets returned by the Loop-Back Slave, because the number of SKIP symbols received can differ from the number transmitted. All other data is compared to the transmitted data, and errors are logged in the **Quad *x* SerDes Diagnostic Data** register(s) (Port 0, offsets 238h through 240h).

Figure 22-4. Digital Far-End Loop-Back

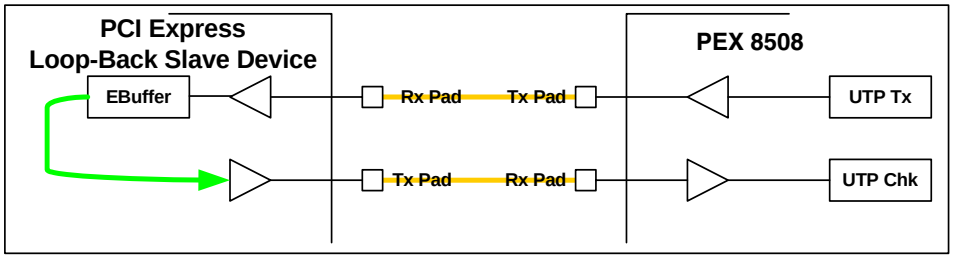

22.1.5 Analog Loop-Back Slave Mode

The PEX 8508 becomes an Analog Loop-Back Slave if it receives training sets with the *Loop-Back Training Control* bit set while the **Physical Layer Test** register *Analog Loop-Back Enable* bit (Port 0, offset 228h[4]) is set. While it is an Analog Loop-Back Slave, the PEX 8508 includes only the serializer and de-serializer in the loop-back data path. The Loop-Back Master must provide the test pattern and data pattern checking. It is unnecessary for the Loop-Back Master to include SKIP Ordered-Sets in the data pattern.

Figure 22-5 illustrates the loop-back data path when Analog Loop-Back Slave mode is enabled.

Figure 22-5. Analog Loop-Back Slave Mode

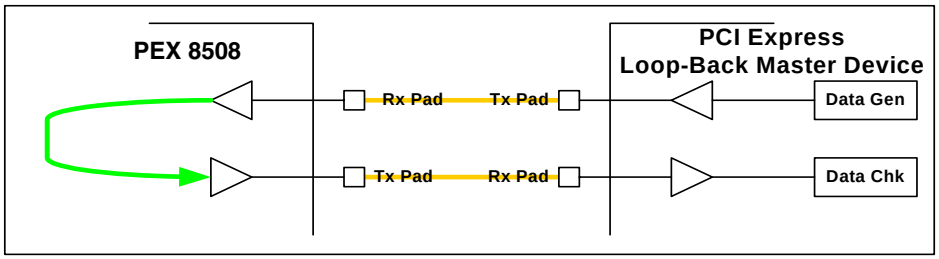

22.1.6 Digital Loop-Back Slave Mode

The PEX 8508 becomes a Digital Loop-Back Slave if it receives training sets with the *Loop-Back Training Control* bit set while the **Physical Layer Test** register *Analog Loop-Back Enable* bit (Port 0, offset 228h[4]) is cleared.

When a PEX 8508 port is a Digital Loop-Back Slave, the port includes the elastic buffer and 8b/10b decoder and encoder in the loop-back data path. The Loop-Back Master must provide the test pattern and data pattern checker. The Loop-Back Master must also transmit SKIP Ordered-Sets with the data pattern. Because the PEX 8508 can return more or fewer SKIP symbols than it receives, the data checker must make provisions for this possibility.

Figure 22-6. Digital Loop-Back Slave Mode

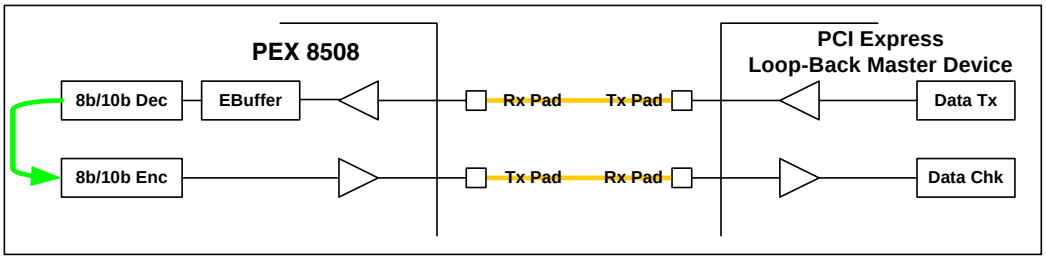

22.1.7 Using the Diagnostic Registers

There are three Diagnostic registers, one for each SerDes quad. The **Quad *x* SerDes Diagnostic Data** register (Port 0, offsets 238h through 240h) contents reflect the performance of the SerDes selected by the register's *SerDes Diagnostic Data Select* field (bits [25:24]). *For example*, if register offset 238h[25:24] is set to 10b, the information in that diagnostic register is for quad 0, SerDes 2. Table 22-2 further illustrates this example.

Table 22-2. SerDes Register Contents When *PRBS Diagnostic Data Select* Field = 10b

Port 0 Register Offset	Register	SerDes
238h	**Quad 0 SerDes Diagnostic Data**	2
23Ch	**Quad 1 SerDes Diagnostic Data**	6
240h	**Quad 2 SerDes Diagnostic Data**	10

22.2 Pseudo-Random and Bit-Pattern Generation

Each SerDes quad has an associated PRBS generator and checker. The PRBS generator is based upon a 7-bit **Linear Feedback Shift** register (LFSR), which can generate up to (2^7 – 1) unique patterns. The PRBS logic is assigned to a SerDes in the quad by manipulating the *SerDes Diagnostic Data Select* bits in the appropriate **Quad *x* SerDes Diagnostic Data** register (Port 0, offsets 238h through 240h[25:24]). The PRBS bit stream is used for internal SerDes or Analog Far-End Loop-Back testing.

The PEX 8508 also provides a method of creating a repeating programmable bit pattern. Each of the four 32-bit **Phy User Test Pattern *x*** registers (Port 0, offsets 210h through 21Ch) are loaded with a 32-bit data pattern. After a port is established as a Loop-Back Master, set the **Physical Layer Test** register *User Test Pattern Enable* bit(s) (Port 0, offset 228h[30:28]) to 1, for the SerDes quad(s) associated with that port. The PEX 8508 proceeds to transmit the data pattern on all lanes, starting with Byte 0 of the **Phy User Test Pattern 0** register and continuing, in sequence, through Byte 3 of the **Phy User Test Pattern 12** register. SKIP Ordered-Sets are inserted at the proper intervals, which makes this method appropriate for Digital Far-End Loop-Back testing. The received pattern is compared to the transmitted pattern. Any errors are logged and retrieved, by reading the **Quad *x* SerDes Diagnostic Data** registers (Port 0, offsets 238h through 240h).

To produce a pseudo-clock bitstream in Analog Loop-Back mode, set the registers as follows:

1. In the Slave device, enable Analog Loop-Back by setting the **Physical Layer Test** register *Analog Loop-Back Enable* bit (Port 0, offset 228h[4] in PLX switches).
2. In the PEX 8508 Loop-Back Master device:
   a. Write the value 4A4A_4A4Ah into each of the **Phy User Test Pattern *x*** registers (Port 0, offsets 210h through 21Ch).
   b. Set the *Port x Loop-Back Command* bit for the specific port in the **Physical Layer Port Command** register (Port 0, offset 230h[0, 4, 8, 12, and/or 16]).
   c. To check whether loop-back is successful, read the corresponding **Physical Layer Port Command** register *Port x Ready as Loop-Back Master* bit (Port 0, offset 230h[3, 7, 11, 15, and/or 19]) in the same Nibble that was set in step a. The Nibble value will be 9h if loop-back is successful.
   d. Set the **Physical Layer Test** register *User Test Pattern Enable* bit(s) (Port 0, offset 228h[30:28]) for the SerDes quad(s) used by the port selected in step b.
   e. The interval between SKIP Ordered-Sets can be programmed in the **SKIP Ordered-Set Interval** register *SKIP Ordered-Set Interval* field (Port 0, offset 234h[11:0]).

   Note: *A high value (such as FFFh) can cause the link to fail.*
3. Exit Analog Loop-Back mode by clearing the **Physical Layer Port Command** register (Port 0, offset 230h), and then the **Physical Layer Test** register (Port 0, offset 228h). The link will re-establish itself.

22.3 JTAG Interface

The PEX 8508 provides a JTAG Boundary Scan interface, which is utilized to debug board connectivity for each ball.

22.3.1 *IEEE 1149.1* and *1149.6* Test Access Port

The *IEEE 1149.1* Test Access Port (TAP), commonly called the JTAG (Joint Test Action Group) debug port, is an architectural standard described in the *IEEE Standard 1149.1-1990*. The *IEEE Standard 1149.6-2003*, *IEEE Standard Test Access Port and Boundary-Scan Architecture Extensions,* defines extensions to *1149.1* to support PCI Express SerDes testing. These standards describe methods for accessing internal chip facilities, using a four- or five-signal interface.

The JTAG debug port, originally designed to support scan-based board testing, is enhanced to support the attachment of debug tools. The enhancements, which comply with *IEEE Standard 1149.1-1994 Specifications for Vendor-Specific Extensions*, are compatible with standard JTAG hardware for boundary-scan system testing.

- **JTAG Signals** – JTAG debug port implements the four required JTAG signals – JTAG_TCK, JTAG_TDI, JTAG_TDO, JTAG_TMS – and optional JTAG_TRST# signal
- **Clock Requirements** – JTAG_TCK signal frequency ranges from DC to 10 MHz
- **JTAG Reset Requirements** – Section 22.3.4, "JTAG Reset Input Signal JTAG_TRST#"

22.3.2 JTAG Instructions

The JTAG debug port provides the *IEEE Standard 1149.1-1990* EXTEST, SAMPLE/PRELOAD, BYPASS, and IDCODE instructions. *IEEE Standard 1149.6-2003* EXTEST_PULSE and EXTEST_TRAIN instructions are also supported. ***PRIVATE instructions are for PLX use only.*** Invalid instructions behave as BYPASS instructions. Table 22-3 lists the JTAG instructions, along with their input codes.

Table 22-3. JTAG Instructions

Instruction	Input Code	Comments
EXTEST	00000b	*IEEE Standard 1149.1-1990*
IDCODE	00001b	
SAMPLE/PRELOAD	00010b	
BYPASS	11111b	
EXTEST_PULSE	00011b	*IEEE Standard 1149.6-2003*
EXTEST_TRAIN	00100b	
PRIVATE[a]	00101b	
	00110b	
	00111b	
	01000b	
	01001b	
	01010b	
	01011b	
	01100b	
	01101b	
	01110b	
	01111b	

*a. **Warning: Non-PLX use of PRIVATE instructions can cause a component to operate in a hazardous manner.***

Table 22-4 defines the JTAG IDCODE values returned by the PEX 8508.

Table 22-4. PEX 8508 JTAG IDCODE Values

Unit of Measure	Version	Part Number	PLX Manufacturer Identity	Least Significant Bit
Bits	0100b	0010_0001_0011_1100b	001_1100_1101b	1
Hex	4h	213Ch	1CDh	1h
Decimal	4	8508	461	1

22.3.3 JTAG Boundary Scan

Scan Description Language (BSDL), IEEE 1149.1-1994, is a supplement to IEEE Standard *1149.1-1990* and *IEEE Standard 1149.1a-1993, IEEE Standard Test Access Port and Boundary-Scan Architecture*. BSDL, a subset of the *IEEE 1076-1993 Standard VHSIC Hardware Description Language (VHDL)*, allows a rigorous description of testability features in components which comply with the standard. This standard is used by automated test pattern generation tools for package interconnect tests and Electronic Design Automation (EDA) tools for synthesized test logic and verification. BSDL supports robust extensions that can be used for internal test generation and to write software for hardware debug and diagnostics.

The primary components of BSDL include the logical port description, physical ball map, instruction set, and **Boundary** register description.

The logical port description assigns symbolic names to the chip balls. Each ball includes a logical type of *in*, *out*, *in out*, *buffer*, or *linkage* that defines the logical direction of signal flow.

The physical ball map correlates the chip logical ports to the physical balls of a specific package. A BSDL description can include several physical ball maps, and maps are provided with a unique name.

Instruction set statements describe the bit patterns that must be shifted into the **Instruction** register to place the chip in the various test modes defined by the standard. Instruction set statements also support descriptions of instructions that are unique to the chip.

The **Boundary** register description lists each cell or shift stage of the **Boundary** register. Each cell has a unique number, the cell numbered 0 is the closest to the Test Data Out (JTAG_TDO) ball and the cell with the highest number is closest to the Test Data In (JTAG_TDI) ball. Each cell contains additional information, including:

- Cell type
- Logical port associated with the cell
- Logical function of the cell
- Safe value
- Control cell number
- Disable value
- Result value

22.3.4 JTAG Reset Input Signal JTAG_TRST#

The JTAG_TRST# Input ball is the asynchronous JTAG logic reset. When JTAG_TRST# is set Low, it causes the PEX 8508's JTAG TAP Controller to initialize. In addition, when the JTAG TAP Controller is initialized, it selects the PEX 8508 standard logic path (core-to-I/O). It is recommended to take the following into consideration when implementing the asynchronous JTAG logic reset on a board:

- If JTAG functionality is required, consider one of the following:
  - JTAG_TRST# Input signal to use a Low-to-High transition once during PEX 8508 boot-up, along with the system PEX_PERST# signal
  - Hold the JTAG_TMS ball High while clocking the JTAG_TCK ball five times
- If JTAG functionality is not required, the JTAG_TRST# signal must be directly connected to VSS, to hold the JTAG TAP Controller inactive
- If the PEX 8508's JTAG TAP Controller is not intended to be used by the design, it is recommended that a 1.5K-Ohm pull-down resistor be connected to the JTAG_TRST# ball, to hold the JTAG TAP Controller in the *Test-Logic-Reset* state, which enables standard logic operation

22.4 Lane Good Status LEDs

The PEX 8508 provides Lane Good outputs, PEX_LANE_GOOD[10, 8, 6, 4:0]#, that can directly drive external common anode LED modules to provide visual indication that the Physical Layer of the link for each lane has trained to at least x1 width.

Software can determine:

- Which lanes have completed Physical Layer linkup, by performing a Memory Read of the **Software-Controlled Lane Status** register *Software-Controlled Lane Status* bits (Port 0, offset 1F4h[10, 8, 6, 4:0]). Bits [10, 8, 6, 4:0] correspond to Lanes [10, 8, 6, 4-0], respectively.
- Whether the link for each port has trained, by reading the **VC0 Resource Status** register *VC0 Negotiation Pending* and **VC1 Resource Status** register *VC1 Negotiation Pending* bits (offsets 160h[17] and 16Ch[17], respectively) in each port. If the bit value is 0, the link has completed Flow Control initialization. These registers can be read by either a PCI Express Configuration request or Memory Read.
- The negotiated link width of each port, by reading the **Link Status** register *Negotiated Link Width* field (offset 78h[25:20]) in each port. This register can be read by either a Configuration request or Memory Read.

Chapter 23 Electrical Specifications

23.1 Introduction

This chapter contains the PEX 8508 power-up/power-down sequencing rules and electrical specifications.

23.2 Power-Up/Power-Down Sequence

When the AUX D3cold state is supported:

- For reliable operation, VDD10, VDD10A, VDD10S, and VDD10X must power-up first and power-down last. No specific sequence is required between the VTT_PEX, VDD33, and VDD33A supplies. All supply rails should power-up within 50 ms of one another.
- VDD10X, VDD33X, and VTT_PEX must remain on after main power (VDD10, VDD10A, VDD10S, and VDD33) is removed.
- When AUX power is applied to the PEX 8508, VDD10X must power-up before VDD33X and VTT_PEX.

When the AUX power D3cold state is not supported:

- For reliable operation, VDD10, VDD10A, VDD10S, and VDD10X should power-up first and power-down last. No specific sequence is required between the VTT_PEX, VDD33, VDD33A, and VDD33X supplies. All supply rails should power-up within 50 ms of one another.

23.3 Absolute Maximum Ratings

Warning: *Maximum limits indicate the temperatures and voltages above which permanent damage can occur. Proper operation at these conditions is not guaranteed, and continuous operation of the device at these limits is not recommended.*

Table 23-1. Absolute Maximum Rating (All Voltages Referenced to VSS System Ground)

Item	Symbol	Absolute Maximum Rating	Units
I/O Interface Supply Voltage	VDD33	-0.5 to +4.6	V
PLL Supply Voltage	VDD33A	-0.5 to +4.6	V
Auxiliary Voltage	VDD33X	-0.5 to +4.6	V
Core (Logic) Supply Voltage	VDD10	-0.3 to +1.65	V
SerDes Analog Supply Voltage	VDD10A	-0.3 to +1.65[a]	V
SerDes Digital Supply Voltage	VDD10S	-0.3 to +1.65[a]	V
SerDes Auxiliary Supply Voltage	VDD10X	-0.3 to +1.65[a]	V
SerDes Termination Supply Voltage	VTT_PEX	-0.3 to +2.5	V
Input Voltage (3.3V Interface)	V_I	-0.3 to +4.6	V
Operating Ambient Temperature (Industrial)	T_A	-40 to +85	°C
Storage Temperature	T_{STG}	-55 to +150	°C

a. The SerDes Analog and Digital power supplies should track within 0.01V of one another.

23.4 Power Characteristics

Table 23-2. Operating Condition Power Supply Rails

Symbol	Parameter	Min	Typ	Max	Units
VDD33	I/O Supply {3.3V ±10%}	3.0	3.3	3.6	V
VDD33A	PLL Supply {3.3V ±10%}	3.0	3.3	3.6	V
VDD33X	Auxiliary I/O Supply {3.3V ±10%}	3.0	3.3	3.6	V
VDD10	Digital Core Supply {1.0V ±10%}	0.9	1.0	1.1	V
VDD10A	Analog SerDes Supply {1.0V ±10%}	0.9	1.0	1.1	V
VDD10S	Digital SerDes Supply {1.0V ±10%}	0.9	1.0	1.1	V
VDD10X	Auxiliary Core Supply {1.0V ±10%}	0.9	1.0	1.1	V
VTT_PEX	SerDes Termination Supply Voltage	1.35	1.5	1.8	V

23.5 Power Consumption

Table 23-3. PEX 8508 Power Estimates (Watts) (8 Lanes)

Traffic Conditions	Core/ SerDes Digital (VDD10)		PCI Express Digital (VDD10S)		PCI Express Auxiliary (VDD10X)		PCI Express Analog (VDD10A)		SerDes Termination (VTT_PEX)		PLL (VDD33A)		GPIO (VDD33/X)		Total	
	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max
A. Heavy	1.06	1.57	0.3	0.34	0.08	0.09	0.11	0.12	0.26	0.32	0.025	0.04	0.02	0.03	**1.84**	**2.50**
B. Medium	0.81	1.25	0.3	0.34	0.08	0.09	0.11	0.12	0.26	0.32	0.025	0.04	0.02	0.03	**1.60**	**2.18**
C. Light	0.64	0.72	0.3	0.34	0.08	0.09	0.11	0.12	0.26	0.32	0.025	0.04	0.02	0.03	**1.43**	**1.65**

A. 85% lane bandwidth utilization. All 8 lanes in the active L0 Link PM state.

B. 35% lane bandwidth utilization. All 8 lanes in the active L0 Link PM state.

C. 10% lane bandwidth utilization. All 8 lanes in the active L0 Link PM state.

Typical Condition – Typical silicon process, 25 °C, nominal supply voltage.
Maximum Condition – Fast silicon process, 0 °C, +10% supply voltage.

23.6 I/O Interface Signal Groupings

Table 23-4. Signal Group PCI Express Analog Interface

Signal Group	Signal Type	Signals	Notes
(a)	PCI Express Output (Transmit)	PEX_PETn[10, 8, 6, 4:0], PEX_PETp[10, 8, 6, 4:0]	Refer to Table 23-7
(b)	PCI Express Input (Receive)	PEX_PERn[10, 8, 6, 4:0], PEX_PERp[10, 8, 6, 4:0]	Refer to Table 23-8
(c)	PCI Express Differential Clock Input	PEX_REFCLKn, PEX_REFCLKp, PEX_REFCLK_CFCn, PEX_REFCLK_CFCp	Refer to Table 23-9

Table 23-5. Signal Group Digital Interface

Signal Group	Signal Type	Signals	Notes
(d)	Digital Output	EE_CS#, EE_DI, EE_SK, FATAL_ERR#, HP_ATNLED[4:0]#, HP_CLKEN[4:0]#, HP_PERST[4:0]#, HP_PWREN[4:0]#, HP_PWRLED[4:0]#, JTAG_TDO, PEX_LANE_GOOD[10, 8, 6, 4:0]#, PEX_NT_RESET#	Refer to Table 23-6
(e)	Digital Input[a]	PEX_PERST#, STRAP_FACTORY_TEST[2:1]#, STRAP_MODE_SEL[1:0], STRAP_NT_UPSTRM_PORTSEL[3:0], STRAP_PORTCFG[4:0], STRAP_SSC_XING_ENA, STRAP_TESTMODE[3:0], STRAP_UPSTRM_PORTSEL[3:0], WAKE#	
(f)	Digital Input with Internal Pull-up Resistor	EE_DO, EE_PR#, HP_BUTTON[4:0]#, HP_MRL[4:0]#, HP_PRSNT[4:0]#, HP_PWRFLT[4:0]#, I2C_ADDR[2:0], JTAG_TCK, JTAG_TDI, JTAG_TMS, JTAG_TRST#	
(g)	Bidirectional (Open Drain)	I2C_SCL, I2C_SDA	Refer to Table 23-6

a. STRAP_ signals must be tied High to VDD33 or Low to VSS (GND).

Table 23-6. DC Electrical Characteristics – Digital Interface

Symbol	Signal Group	Parameter	Min	Typ	Max	Unit	Conditions
I_{OL}	(d) (g)	Output Low Current	8			mA	V_{OL} = 0.4V
I_{OH}	(d)	Output High Current	8			mA	V_{OH} = 2.4V
V_{IL}	(e) (f)	Input Low Voltage			0.8	V	
V_{IH}	(e) (f)	Input High Voltage	2.0			V	
C_{PIN}	(a) (b) (d)	Ball Capacitance			6	pF	
	(c) (e) (f)	Ball Capacitance			5	pF	
	(g)	Ball Capacitance			10	pF	
$I_{LEAKAGE}$	(d)	Three-state Leakage			±500	nA	
	(e)	Input Leakage			±50	nA	
	(f)	Pull-Up Leakage	+0.1/-8		+0.1/-20	µA	
R_{PU}	(f)	Pull-Up Impedance	200K			Ohm	
V_{HYS}	(g)	Input Hysteresis	150			mV	

Table 23-7. PCI Express Transmit (Signal Group a) AC and DC Characteristics

Symbol	Parameter	Min	Typ	Max	Units	Comments
UI	Unit Interval	399.88	400	400.12	ps	
$V_{TX-DIFFp-p}$	Differential Peak-to-Peak Output Voltage	0.800		1.2	V	$V_{TX-DIFFp-p} = 2 * \lvert V_{TX-D+} - V_{TX-D-} \rvert$
$V_{TX-DE-RATIO}$	De-Emphasized Differential Output Voltage (Ratio)	-3.0	-3.5	-4.0	dB	Ratio of the $V_{TX-DIFFp-p}$ of the 2nd and following bits after a transition divided by the $V_{TX-DIFFp-p}$ of the 1st bit after a transition. Refer to Note 1.
T_{TX-EYE}	Minimum Tx Eye Width	0.75			UI	The maximum Transmitter jitter can be derived as: $T_{TX-MAX-JITTER} = 1 - T_{TX-EYE} = 0.25$ UI This parameter is measured with the equivalent of a zero-jitter Reference Clock. Refer to Notes 1 and 2.
$T_{TX-EYE-MEDIAN-to-MAX-JITTER}$	Maximum time between the Jitter Median and Maximum Deviation from the Median			0.125	UI	Jitter is defined as the measurement variation of the crossing points ($V_{TX-DIFF}$ = 0V) in relation to recovered Tx UI. Refer to Notes 1 and 2.
$T_{TX-RISE}$, $T_{TX-FALL}$	D+/D- Tx Output Rise/Fall Time	0.125			UI	Refer to Notes 1 and 4.
$V_{TX-CM-ACp}$	RMS AC Peak Common Mode Output Voltage			20	mV	$V_{TX-CM-ACp} = RMS(\lvert V_{TX-D+} + V_{TX-D-} \rvert/2 - V_{TX-CM-DC})$ $V_{TX-CM-DC} = DC_{(avg)}$ of $\lvert V_{TX-D+} + V_{TX-D-} \rvert/2$ Refer to Note 1.
$V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$	Absolute Delta of DC Common Mode Voltage during L0 and Electrical Idle	0		100	mV	$\lvert V_{TX-CM-DC}$ [during L0] - $V_{TX-CM-Idle-DC}$ [during Electrical Idle]$\rvert \leq 100$mV $V_{TX-CM-DC} = DC_{(avg)}$ of $\lvert V_{TX-D+} + V_{TX-D-} \rvert/2$ [L0] $V_{TX-CM-Idle-DC} = DC_{(avg)}$ of $\lvert V_{TX-D+} + V_{TX-D-} \rvert/2$ [Electrical Idle] Refer to Note 1.
$V_{TX-CM-DC-LINE-DELTA}$	Absolute Delta of DC Common Mode Voltage between D+ and D-	0		25	mV	$\lvert V_{TX-CM-DC-D+} - V_{TX-CM-DC-D-} \rvert \leq 25$mV $V_{TX-CM-DC-D+} = DC_{(avg)}$ of $\lvert V_{TX-D+} \rvert$ $V_{TX-CM-DC-D-} = DC_{(avg)}$ of $\lvert V_{TX-D-} \rvert$ Refer to Note 1.
$V_{TX-IDLE-DIFFp}$	Electrical Idle Differential Peak Output Voltage	0		20	mV	$V_{TX-IDLE-DIFFp} = \lvert V_{TX-Idle-D+} - V_{TX-Idle-D-} \rvert \leq 20$mV Refer to Note 1.
$V_{TX-RCV-DETECT}$	Amount of Voltage Change Allowed during Receiver Detection			600	mV	The total amount of voltage change that a Transmitter can apply to sense whether a low-impedance Receiver is present.
$V_{TX-DC-CM}$	Tx DC Common Mode Voltage	0		3.6	V	The allowed DC Common Mode voltage under any condition.

Table 23-7. PCI Express Transmit (Signal Group a) AC and DC Characteristics *(Cont.)*

Symbol	Parameter	Min	Typ	Max	Units	Comments
$I_{TX-SHORT}$	Tx Short Circuit Current Limit			90	mA	The total current the Transmitter can provide when shorted to its ground.
$T_{TX-IDLE-MIN}$	Minimum Time Spent in Electrical Idle	50			UI	Minimum time a Transmitter must be in Electrical Idle. Utilized by the Receiver to start looking for an Electrical Idle Exit after successfully receiving an Electrical Idle Ordered-Set.
$T_{TX-IDLE-SET-TO-IDLE}$	Maximum Time to Transition to a Valid Electrical Idle after Sending an Electrical Idle Ordered-Set			20	UI	After sending an Electrical Idle Ordered-Set, the Transmitter must meet all Electrical Idle specifications within this time. A de-bounce time for the Transmitter to meet Electrical Idle after transitioning from L0.
$RL_{TX-DIFF}$	Differential Return Loss	10			dB	Measured over 50 MHz to 1.25 GHz.
RL_{TX-CM}	Common Mode Return Loss	6			dB	Measured over 50 MHz to 1.25 GHz.
$Z_{TX-DIFF-DC}$	DC Differential Tx Impedance	80	100	120	Ohm	Tx DC Differential mode low impedance. Refer to Note 5.
$L_{TX-SKEW}$	Lane-to-Lane Output Skew			500 + 2 UI	ps	Static skew between any two Transmitter lanes within a single link.

Notes:

1. *Specified at the measurement point into a timing and voltage compliance test load, as illustrated in Figure 23-1.*

Figure 23-1. Compliance Test/Measurement Load

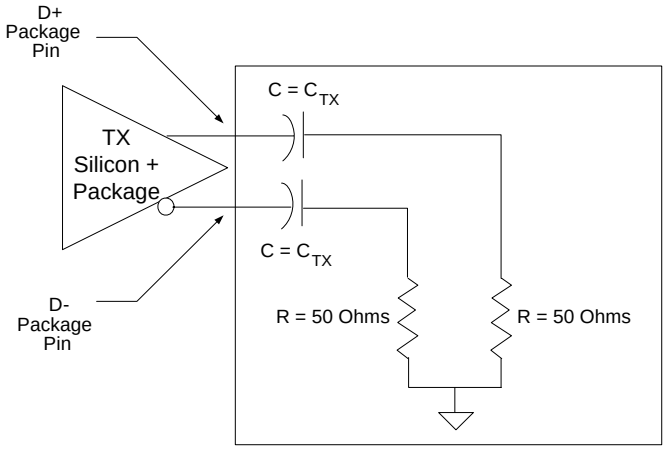

2. *At T_{TX-EYE} = 0.75 UI provides for a total sum of deterministic and random jitter budget of $T_{TX-JITTER-MAX}$ = 0.25 UI for the Transmitter. The $T_{TX-EYE-MEDIAN-to-MAX-JITTER}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median is less than half the total Tx jitter budget. (**Note:** The median is not the same as the mean.) The jitter median describes the point in time where the number of jitter points on either side is approximately equal as opposed to the averaged time value. This parameter is measured with the equivalent of a zero-jitter Reference Clock. The T_{TX-EYE} measurement is to be met at the target bit error rate. The $T_{TX-EYE-MEDIAN-to-MAX-JITTER}$ specification is to be met using the compliance pattern at a sample size of 1,000,000 UI.*

3. *The Transmitter input impedance shall result in a differential return loss, greater than or equal to 10 dB, with a Differential Test Input signal no less than 200 mV (peak value, 400 mV differential peak-to-peak) swing around ground, applied to D+ and D- lines and a common mode return loss greater than or equal to 6 dB over a frequency range of 50 MHz to 1.25 GHz. This input impedance requirement applies to all valid input levels. The reference impedance for return loss measurements is 50 Ohms to ground for both the D+ and D- lines. The series capacitance C_{TX} is optional for the return loss measurement.*

4. *Measured between 20 to 80% at Transmitter package balls into a test load, as illustrated in Figure 23-1, both V_{TX-D+} and V_{TX-D-}.*

5. *$Z_{TX-DIFF-DC}$ is the small signal resistance of the transmitter measured at a DC operating point that is equivalent to that established by connecting a 100 Ohm resistor from D+ and D- while the Tx is driving a static logic 1 or logic 0. Equivalently, this parameter can be derived by measuring the RMS voltage of the Tx while transmitting a test pattern into two different differential terminations that are near 100 Ohms.*

   Small signal resistance is measured by forcing a small change in differential voltage and dividing this by the corresponding change in current.

Table 23-8. PCI Express Receive (Signal Group b) AC and DC Characteristics

Symbol	Parameter	Min	Typ	Max	Units	Comments
UI	Unit Interval	399.88	400	400.12	ps	The UI is 400 ps ±300 ppm. UI does not account for SSC-dictated variations.
$V_{RX-DIFFp-p}$	Differential Input Peak-to-Peak Voltage	0.175		1.200	V	$V_{RX-DIFFp-p} = 2*\|V_{RX-D+} - V_{RX-D-}\|$
T_{RX-EYE}	Minimum Receiver Eye Width	0.4			UI	The maximum interconnect media and Transmitter jitter that can be tolerated by the Receiver can be derived as: $T_{RX-MAX-JITTER} = 1 - T_{RX-EYE} = 0.6$ UI Refer to Notes 6, 7, and 8.
$T_{RX-EYE-MEDIAN-to-MAX-JITTER}$	Maximum Time between the Jitter Median and Maximum Deviation from the Median			0.3	UI	Jitter is defined as the measurement variation of the crossing points ($V_{TX-DIFF}$ = 0V) in relation to recovered Tx UI. Refer to Notes 6 and 7.
$V_{RX-CM-ACp}$	AC Peak Common Mode Input Voltage			150	mV	$V_{RX-CM-ACp} = \|V_{RX-D+} + V_{RX-D-}\| /2 - V_{RX-CM-DC})$ $V_{RX-CM-DC} = DC_{(avg)}$ of $\|V_{RX-D+}\|$ Refer to Note 6.
$RL_{RX-DIFF}$	Differential Return Loss	10			dB	Measured over 50 MHz to 1.25 GHz. Refer to Note 9.
RL_{RX-CM}	Common Mode Return Loss	6			dB	Measured over 50 MHz to 1.25 GHz. Refer to Note 9.
$Z_{RX-DIFF-DC}$	DC Differential Tx Impedance	80	100	120	Ohm	Rx DC Differential mode impedance.
Z_{RX-DC}	DC Input Impedance	40	50	60	Ohm	Required Rx D+ and D- DC impedance (50 Ohms ±20% tolerance). Refer to Note 6.
$V_{RX-IDLE-DET-DIFFp-p}$	Electrical Idle Detect Threshold	65		175	mV	$V_{RX-IDLE-DET-DIFFp-p} = 2* \|V_{RX-D+} + V_{RX-D-}\|$ Measured at the package balls of the Receiver.
$T_{RX-IDLE-DET-DIFF-ENTERTIME}$	Unexpected Electrical Idle Enter Idle Detect Threshold Integration Time			10	ms	An un-expected Electrical Idle ($V_{RX-DIFFp-p} < V_{RX-IDLE-DET-DIFFp-p}$) must be recognized no longer than $T_{RX-IDLE-DET-DIFF-ENTERTIME}$ to signal an unexpected idle condition.
$L_{RX-SKEW}$	Total Skew			20	ns	Skew across all lanes in a link.

Notes:

6. *The test load in Figure 23-1 should be used as the Rx device when taking measurements.*
7. *The $T_{RX\text{-}EYE\text{-}MEDIAN\text{-}to\text{-}MAX\text{-}JITTER}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median is less than half the total, 0.64. (**Note:** The median is not the same as the mean.) The jitter median describes the point in time where the number of jitter points on either side is approximately equal as opposed to the averaged time value. The $T_{RX\text{-}EYE}$ measurement is to be met at the target bit error rate. The $T_{RX\text{-}EYE\text{-}MEDIAN\text{-}to\text{-}MAX\text{-}JITTER}$ specification is to be met using the compliance pattern at the sample size of 1,000,000 UI.*
8. *Refer to the PCI Express Jitter and BER White Paper for details regarding the Rx-Eye measurement.*
9. *The Receiver input impedance shall result in a differential return loss, greater than or equal to 10 dB, with a Differential Test Input signal of no less than 200 mV (peak value, 400 mV differential peak-to-peak) swing around ground, applied to D+ and D- lines and a common mode return loss greater than or equal to 6 dB (no bias required) over a frequency range of 50 MHz to 1.25 GHz. This input impedance requirement applies to all valid input levels. The reference impedance for return loss measurements is 50 Ohms to ground for both the D+ and D- lines. The series capacitance C_{TX} is optional for the return loss measurement.*

Table 23-9. PCI Express Differential Clock (Signal Group c) AC and DC Characteristics

Symbol	Parameter	Min	Typ	Max	Unit	Notes
F_{REFCLK}	Reference Clock Frequency		100		MHz	
V_{CM}	Input Common Mode Voltage	0.6	0.65	0.7	V	1
V_{SW}	Differential Voltage Swing (0-to-peak)	125		800	mV	
	Differential Voltage Swing (peak-to-peak)	250		1,600	mV	
T_R/T_F	Clock Input Rise/Fall Time			1.5	ns	2
DC_{REFCLK}	Input Clock Duty Cycle	40	50	60	%	
R_{TERM}	Input Parallel Termination (Single-ended)		55		Ohm	
	Input Parallel Termination (Differential)		110		Ohm	
PPM	Reference Clock Tolerance	-300		+300	ppm	

Notes:

1. *PEX_REFCLKn/p must be AC-coupled. Use a 0.01 to 0.1 µF capacitor.*
2. *Specified at 20 to 80% points at the package balls.*

23.7 Transmit Drive Characteristics

The drive current and the transmit equalization function is programmable, to allow for optimization of different backplane lengths and materials.

23.7.1 Drive Current

The nominal drive current is programmable (2-bit) within the range of 10 to 28 mA. [Refer to the **SerDes Nominal Drive Current Select** register (offset 248h) for details.]

The nominal drive current can be further programmed (4-bit) with finer granularity, within the range of 0.65X to 1.35X. [Refer to the **SerDes Drive Current Level Select 1** and **SerDes Drive Current Level Select 2** registers (offsets 24Ch and 250h, respectively) for details.]

23.7.2 Transmit Equalization

The Transmitter incorporates programmable (4-bit) first-order equalization, within the range of 0 to -7.96 dB. [Refer to the **SerDes Drive Equalization Level Select 1** and **SerDes Drive Equalization Level Select 2** registers (offsets 254h and 258h, respectively) for details.]

23.7.3 Transmit Termination Adjust

The *PCI Express Base r1.1* specifies termination (50 Ohms nominal) at the Transmit side to VTT. The Transmit driver incorporates a 2-bit register (per SerDes quad), which allows for a ±20% termination adjustment to mitigate stub effects and other non-idealities in the PCB channel. Refer to the **Physical Layer** register *SerDes Quad x TxTermAdjust* fields (offset 22Ch[13:8]) for details.

23.8 Receive Characteristics

The following programmable bits control the electrical characteristics of the Receiver circuit, to mitigate the effects of signal loss and distortion across the PCB channel.

23.8.1 Receive Equalization

The Receiver incorporates a programmable 2-bit register (per SerDes quad) to modify the high-pass filter within the circuit, which serves to mitigate the effects of Inter Symbol Interference due to frequency-dependent losses across the PCB material. Refer to **Physical Layer** register *SerDes Quad x RxEqCtl* fields (offset 22Ch[29:24]) for details.

23.8.2 Receive Termination Adjust

The *PCI Express Base r1.1* specifies termination (50 Ohms nominal) at the Receive side to ground. The Receiver input incorporates a 2-bit register (per SerDes quad), which allows for a ±20% termination adjustment to mitigate stub effects and other non-idealities in the PCB channel. Refer to the **Physical Layer** register *SerDes Quad x RxTermAdjust* fields (offset 22Ch[21:16]) for details.

Chapter 24 Thermal and Mechanical Specifications

24.1 Thermal Characteristics

The PEX 8508 does not include a heat sink. The information described in this section is based upon sample thermal performance when a heat sink is used with the PEX 8508, and is provided for reference only.

24.1.1 Sample Thermal Data

- **Heat Sink** – Wakefield 624-45AB, 21 x 21 x 2.0 mm base, 11.4 mm tall, 36 pins
- **Heat Sink Interface** – Chomerics T410, 0.18 mm thick, 7.1 C-cm2/W
- **Maximum Junction Temperature** – 125 °C

Table 24-1. Sample PEX 8508 Thermal Data with Heat Spreader – PBGA-H 19 x 19 mm Package

New Power (W)	Heat Sink	Airflow (m/s)	Θ_{JA} (°C/W)	Ψ_{JT} (°C/W)	Ψ_{JB} (°C/W)	Θ_{JC} (°C/W)	Θ_{JB} (°C/W)
2.10	No	0.00	17.20	3.63	7.86	5.74	9.23
		1.02	14.90	3.67	7.50		
		2.04	N/A	N/A	N/A		
	Yes	0.00	15.50	4.04	7.09		
		1.02	11.60	4.44	6.35		
		2.04	N/A	N/A	N/A		

24.1.2 Sample Heat Sink Specifications

Figure 24-1. Sample PEX 8508 Heat Sink (21 x 21 x 11.4 mm)

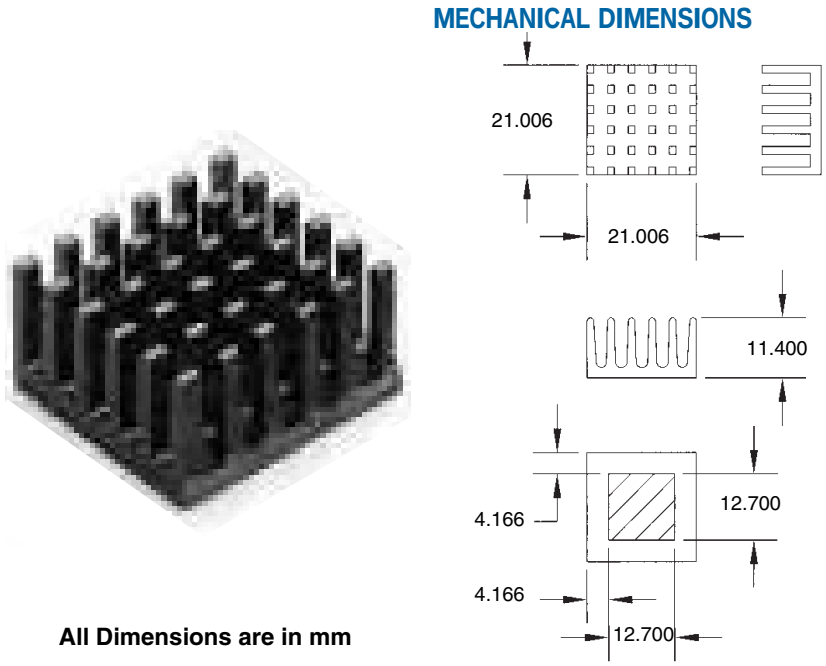

All Dimensions are in mm

24.2 Package Specifications

The PEX 8508 is offered in a 19-mm square PBGA-H (Plastic BGA Heat-spreader) package. (Refer to Table 24-2.)

Unpopulated BGA balls allow board design and placement of board-level de-coupling capacitors between VDD10, VDD10A, VDD10S, VDD10X, VDD33, VDD33A, VDD33X, and VSS/Ground.

Table 24-2. PEX 8508 Package Specifications

Parameter	Specification
Package Type	Plastic Ball Grid Array Heat-spreader (PBGA-H)
Number of Balls	296
Package Dimensions	19 x 19 mm (approximately 2.5 mm high)
Ball Matrix Pattern	18 x 18 mm (6 x 6 mm center area ***reserved*** for Ground)
Ball Pitch	1.00 mm
Ball Diameter	0.60 ±0.15 mm
Ball Spacing	0.40 mm

24.3 Mechanical Dimensions

Figure 24-2. PEX 8508 296-Ball PBGA-H Mechanical Dimensions

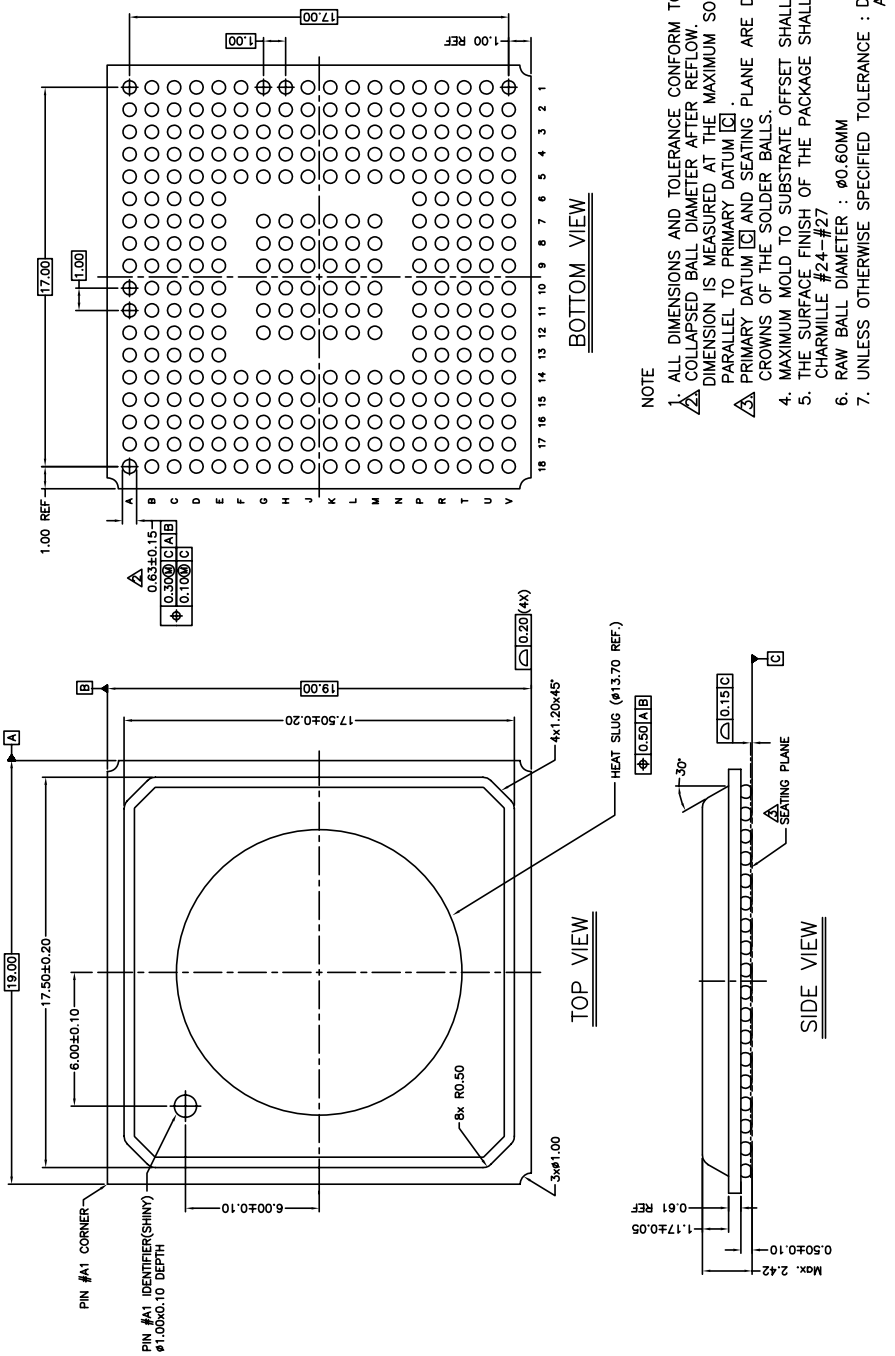

Appendix A Serial EEPROM Memory Map

A.1 Serial EEPROM Memory Map

In Table A-1, all register offsets are Byte addresses, and all serial EEPROM addresses are DWord addresses (Byte addresses shifted right 2 bits). The serial EEPROM DWord addresses are used directly for the **Serial EEPROM Control** register *EepBlkAddr* field (Port 0, offset 260h[12:0]).

Table A-1. PEX 8508 Serial EEPROM Memory Map

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
000h	Configuration ID	000h	31Ah	3A4h	42Eh	4B8h	Configuration ID	542h	63Ah
004h	Command/Status	001h	31Bh	3A5h	42Fh	4B9h	Command/Status	543h	63Bh
008h	Class Code and Revision ID	002h	31Ch	3A6h	430h	4BAh	Class Code and Revision ID	544h	63Ch
00Ch	Miscellaneous Control	003h	31Dh	3A7h	431h	4BBh	Miscellaneous Control	545h	63Dh
010h	Base Address 0	004h	31Eh	3A8h	432h	4BCh	Base Address 0 (Link) ***Reserved*** (Virtual)	546h	63Eh
014h	Base Address 1	005h	31Fh	3A9h	433h	4BDh	Base Address 1	547h	63Fh
018h	Bus Number	006h	320h	3AAh	434h	4BEh	Base Address 2	548h	640h
01Ch	Secondary Status, I/O Limit, and I/O Base	007h	321h	3ABh	435h	4BFh	Base Address 3	549h	641h
020h	Memory Base and Limit	008h	322h	3ACh	436h	4C0h	Base Address 4	54Ah	642h
024h	Prefetchable Memory Base and Limit	009h	323h	3ADh	437h	4C1h	Base Address 5	54Bh	643h
028h	Prefetchable Memory Upper Base Address	00Ah	324h	3AEh	438h	4C2h	***Reserved***	54Ch	644h
02Ch	Prefetchable Memory Upper Limit Address	00Bh	325h	3AFh	439h	4C3h	Subsystem ID and Subsystem Vendor ID	54Dh	645h
030h	I/O Upper Base and Limit Address	00Ch	326h	3B0h	43Ah	4C4h	Expansion ROM Base Address (Link) Expansion ROM Base Address (***Reserved***) (Virtual)	54Eh	646h
034h	New Capability Pointer	00Dh	327h	3B1h	43Bh	4C5h	Capability Pointer	54Fh	647h
038h	Expansion ROM Base Address ***(Reserved)***	00Eh	328h	3B2h	43Ch	4C6h	***Reserved***	550h	648h
03Ch	Bridge Control and Interrupt Signal	00Fh	329h	3B3h	43Dh	4C7h	Interrupt	551h	649h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
040h	Power Management Capability	010h	32Ah	3B4h	43Eh	4C8h	Power Management Capability	552h	64Ah
044h	Power Management Status and Control	011h	32Bh	3B5h	43Fh	4C9h	Power Management Status and Control	553h	64Bh
	PM Data[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 0h	74Eh	75Bh	768h	775h	782h	*PM Data*[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 0h	734h	741h
	PM Data[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 3h	74Fh	75Ch	769h	776h	783h	*PM Data*[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 3h	735h	742h
	PM Data[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 4h	750h	75Dh	76Ah	777h	784h	*PM Data*[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 4h	736h	743h
	PM Data[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 7h	751h	75Eh	76Bh	778h	785h	*PM Data*[31:24] and *Data Scale*[14:13], for *Data Select*[12:9] = 7h	737h	744h
048h	Message Signaled Interrupt Capability	012h	32Ch	3B6h	440h	4CAh	Message Signaled Interrupt Capability	554h	64Ch
04Ch	Message Address[31:0]	013h	32Dh	3B7h	441h	4CBh	Message Address[31:0]	555h	64Dh
050h	Message Upper Address[63:32]	014h	32Eh	3B8h	442h	4CCh	Message Upper Address[63:32]	556h	64Eh
054h	Message Data	015h	32Fh	3B9h	443h	4CDh	Message Data	557h	64Fh
058h	Mask Bit for MSI	016h	330h	3BAh	444h	4CEh	Mask Bit for MSI	558h	650h
058h	***Reserved***	017h	331h	3BBh	445h	4CFh	***Reserved***	559h	651h
…	…	…	…	…	…	…	…	…	…
064h	***Reserved***	019h	333h	3BDh	447h	4D1h	***Reserved***	55Bh	653h
068h	PCI Express Capability List and Capabilities	01Ah	334h	3BEh	448h	4D2h	PCI Express Capability List and Capabilities	55Ch	654h
06Ch	Device Capabilities	01Bh	335h	3BFh	449h	4D3h	Device Capabilities	55Dh	655h
070h	Device Status and Control	01Ch	336h	3C0h	44Ah	4D4h	Device Status and Control	55Eh	656h
074h	Link Capabilities	01Dh	337h	3C1h	44Bh	4D5h	Link Capabilities	55Fh	657h
078h	Link Status and Control	01Eh	338h	3C2h	44Ch	4D6h	Link Status and Control	560h	658h
07Ch	Slot Capabilities	01Fh	339h	3C3h	44Dh	4D7h	***Reserved***	561h	659h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
080h	Slot Status and Control	020h	33Ah	3C4h	44Eh	4D8h	***Reserved***	562h	65Ah
084h	***Reserved***	021h	33Bh	3C5h	44Fh	4D9h	***Reserved***	563h	65Bh
088h	***Reserved***	022h	33Ch	3C6h	450h	4DAh	***Reserved***	564h	65Ch
08Ch	***Reserved***	023h	33Dh	3C7h	451h	4DBh	***Reserved***	565h	65Dh
090h	***Reserved***	024h	33Eh	3C8h	452h	4DCh	Set Virtual Interface IRQ	566h	65Eh
094h	***Reserved***	025h	33Fh	3C9h	453h	4DDh	Clear Virtual Interface IRQ	567h	65Fh
098h	***Reserved***	026h	340h	3CAh	454h	4DEh	Set Virtual Interface IRQ Mask	568h	660h
09Ch	***Reserved***	027h	341h	3CBh	455h	4DFh	Clear Virtual Interface IRQ Mask	569h	661h
0A0h	***Reserved***	028h	342h	3CCh	456h	4E0h	Set Link Interface IRQ	56Ah	662h
0A4h	***Reserved***	029h	343h	3CDh	457h	4E1h	Clear Link Interface IRQ	56Bh	663h
0A8h	***Reserved***	02Ah	344h	3CEh	458h	4E2h	Set Link Interface IRQ Mask	56Ch	664h
0ACh	***Reserved***	02Bh	345h	3CFh	459h	4E3h	Clear Link Interface IRQ Mask	56Dh	665h
0B0h	***Reserved***	02Ch	346h	3D0h	45Ah	4E4h	NT Port Scratchpad_0	56Eh	666h
0B4h	***Reserved***	02Dh	347h	3D1h	45Bh	4E5h	NT Port Scratchpad_1	56Fh	667h
0B8h	***Reserved***	02Eh	348h	3D2h	45Ch	4E6h	NT Port Scratchpad_2	570h	668h
0BCh	***Reserved***	02Fh	349h	3D3h	45Dh	4E7h	NT Port Scratchpad_3	571h	669h
0C0h	***Reserved***	030h	34Ah	3D4h	45Eh	4E8h	NT Port Scratchpad_4	572h	66Ah
0C4h	***Reserved***	031h	34Bh	3D5h	45Fh	4E9h	NT Port Scratchpad_5	573h	66Bh
0C8h	***Reserved***	032h	34Ch	3D6h	460h	4EAh	NT Port Scratchpad_6	574h	66Ch
0CCh	***Reserved***	033h	34Dh	3D7h	461h	4EBh	NT Port Scratchpad_7	575h	66Dh
0D0h	***Reserved***	034h	34Eh	3D8h	462h	4ECh	NT Port Virtual Interface BAR1 Setup	576h	66Eh
0D4h	***Reserved***	035h	34Fh	3D9h	463h	4EDh	NT Port Virtual Interface Memory BAR2 Setup	577h	66Fh
0D8h	***Reserved***	036h	350h	3DAh	464h	4EEh	NT Port Virtual Interface Memory BAR3 Setup	578h	670h
0DCh	***Reserved***	037h	351h	3DBh	465h	4EFh	NT Port Virtual Interface Memory BAR4/5 Setup	579h	671h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
0E0h	***Reserved***	038h	352h	3DCh	466h	4F0h	NT Port Virtual Interface Memory BAR5 Setup	57Ah	672h
0E4h	***Reserved***	039h	353h	3DDh	467h	4F1h	NT Port Link Interface BAR0/BAR1 Setup (Link) ***Reserved*** (Virtual)	57Bh	673h
0E8h	***Reserved***	03Ah	354h	3DEh	468h	4F2h	NT Port Link Interface Memory BAR2/3 Setup (Link) ***Reserved*** (Virtual)	57Ch	674h
0ECh	***Reserved***	03Bh	355h	3DFh	469h	4F3h	NT Port Link Interface Memory BAR3 Setup (Link) ***Reserved*** (Virtual)	57Dh	675h
0F0h	***Reserved***	03Ch	356h	3E0h	46Ah	4F4h	NT Port Link Interface Memory BAR4/5 Setup (Link) ***Reserved*** (Virtual)	57Eh	676h
0F4h	***Reserved***	03Dh	357h	3E1h	46Bh	4F5h	NT Port Link Interface Memory BAR5 Setup (Link) ***Reserved*** (Virtual)	57Fh	677h
0F8h	***Reserved***	03Eh	358h	3E2h	46Ch	4F6h	Configuration Address Window	580h	678h
0FCh	***Reserved***	03Fh	359h	3E3h	46Dh	4F7h	Configuration Data Window	581h	679h
100h	Device Serial Number Enhanced Capability Header	040h	35Ah	3E4h	46Eh	4F8h	Device Serial Number Enhanced Capability Header	582h	67Ah
104h	Serial Number (Lower DW)	041h	35Bh	3E5h	46Fh	4F9h	Serial Number (Lower DW)	583h	67Bh
108h	Serial Number (Upper DW)	042h	35Ch	3E6h	470h	4FAh	Serial Number (Upper DW)	584h	67Ch
10Ch	***Reserved***	043h	35Dh	3E7h	471h	4FBh	***Reserved***	585h	67Dh
…	…	…	…	…	…	…	…	…	…
134h	***Reserved***	04Dh	367h	3F2h	47Bh	505h	***Reserved***	58Fh	687h
138h	Power Budgeting Enhanced Capability Header	04Eh	368h	3F2h	47Ch	506h	Power Budgeting Enhanced Capability Header	590h	688h
13Ch	Data Select	04Fh	369h	3F3h	47Dh	507h	Data Select	591h	689h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
140h	Power Budgeting Data (next row overwrites this row)	050h	36Ah	3F4h	47Eh	508h	Power Budgeting Data (next row overwrites this row)	592h	68Ah
	Power Budgeting Data, Data Select (offset 13Ch) = 00h	753h	760h	76Dh	77Ah	787h	Power Budgeting Data, Data Select (offset 13Ch) = 00h	739h	746h
	Power Budgeting Data, Data Select (offset 13Ch) = 01h	754h	761h	76Eh	77Bh	788h	Power Budgeting Data, Data Select (offset 13Ch) = 01h	73Ah	747h
	Power Budgeting Data, Data Select (offset 13Ch) = 02h	755h	762h	76Fh	77Ch	789h	Power Budgeting Data, Data Select (offset 13Ch) = 02h	73Bh	748h
	Power Budgeting Data, Data Select (offset 13Ch) = 03h	756h	763h	770h	77Dh	78Ah	Power Budgeting Data, Data Select (offset 13Ch) = 03h	73Ch	749h
	Power Budgeting Data, Data Select (offset 13Ch) = 04h	757h	764h	771h	77Eh	78Bh	Power Budgeting Data, Data Select (offset 13Ch) = 04h	73Dh	74Ah
	Power Budgeting Data, Data Select (offset 13Ch) = 05h	758h	765h	772h	77Fh	78Ch	Power Budgeting Data, Data Select (offset 13Ch) = 05h	73Dh	74Bh
	Power Budgeting Data, Data Select (offset 13Ch) = 06h	759h	766h	773h	780h	78Dh	Power Budgeting Data, Data Select (offset 13Ch) = 06h	73Fh	74Ch
	Power Budgeting Data, Data Select (offset 13Ch) = 07h	75Ah	767h	774h	781h	78Eh	Power Budgeting Data, Data Select (offset 13Ch) = 07h	740h	74Dh
144h	Power Budget Capability	051h	36Bh	3F5h	47Fh	509h	Power Budget Capability	593h	68Bh
148h	Virtual Channel Enhanced Capability Header	052h	36Ch	3F6h	480h	50Ah	Virtual Channel Enhanced Capability Header	594h	68Ch
14Ch	Port VC Capability 1	053h	36Dh	3F7h	481h	50Bh	Port VC Capability 1	595h	68Dh
150h	Port VC Capability 2	054h	36Eh	3F8h	482h	50Ch	Port VC Capability 2	596h	68Eh
154h	Port VC Status and Control (#1)	055h	36Fh	3F9h	483h	50Dh	Port VC Status and Control (#1)	597h	68Fh
	Port VC Status and Control (#2)	752h	75Fh	76Ch	779h	786h	Port VC Status and Control (#2)	738h	745h
158h	VC0 Resource Capability	056h	370h	3FAh	484h	50Eh	VC0 Resource Capability	598h	690h
15Ch	VC0 Resource Control	057h	371h	3FBh	485h	50Fh	VC0 Resource Control	599h	691h
160h	VC0 Resource Status	058h	372h	3FCh	486h	510h	VC0 Resource Status	59Ah	692h
164h	VC1 Resource Capability	059h	373h	3FDh	487h	511h	VC1 Resource Capability	59Bh	693h
168h	VC1 Resource Control	05Ah	374h	3FEh	488h	512h	VC1 Resource Control	59Ch	694h
16Ch	VC1 Resource Status	05Bh	375h	3FFh	489h	513h	VC1 Resource Status	59Dh	695h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
170h	***Reserved***	05Ch	376h	400h	48Ah	514h	***Reserved***	59Eh	696h
…	…	…	…	…	…	…	…	…	…
1C4h	***Reserved***	071h	38Bh	415h	49Fh	529h	***Reserved***	5B3h	6ABh
1C8h	ECC Error Check Disable	072h	–	–	–	–	***Reserved***	5B4h	6ACh
1CCh	Error Handler 32-Bit Error Status ***(Factory Test Only)***	073h	–	–	–	–	Error Handler 32-Bit Error Status ***(Factory Test Only)***	5B5h	6ADh
1D0h	Error Handler 32-Bit Error Mask ***(Factory Test Only)***	074h	–	–	–	–	Error Handler 32-Bit Error Status ***(Factory Test Only)***	5B6h	6AEh
1D4h	***Factory Test Only***	075h	–	–	–	–	***Reserved***	–	–
1D8h	***Factory Test Only***	076h	–	–	–	–	***Reserved***	–	–
1DCh	Debug Control	077h	–	–	–	–	Debug Control	5B9h	6B1h
1E0h	Power Management Hot Plug User Configuration	078h	392h	41Ch	4A6h	530h	***Reserved***	5BAh	6B2h
1E4h	Egress Control and Status	079h	393h	41Dh	4A7h	531h	Egress NT Port Link Interface Control and Status (Link) ***Reserved*** (Virtual)	5BBh	6B3h
1E8h	Bad TLP Count	07Ah	394h	41Eh	4A8h	531h	***Reserved***	5BCh	6B4h
1ECh	Bad DLLP Count	07Bh	395h	41Fh	4A9h	532h	***Reserved***	5BDh	6B5h
1F0h	Device-Specific Relaxed Ordering Enable	07Ch	396h	420h	4AAh	533h	***Reserved***	5BEh	6B6h
1F4h	Software-Controlled Lane Status	07Dh	–	–	–	–	***Reserved***	5BFh	6B7h
1F8h	ACK Transmission Latency Limit	07Eh	398h	422h	4ACh	535h	ACK Transmission Latency Limit (Link) ***Reserved*** (Virtual)	5C0h	6B8h
1FCh	***Reserved***	07Fh	–	–	–	–	***Reserved***	–	–
200h	***Factory Test Only***	080h	–	–	–	–	***Reserved***	–	–
204h	Physical Receiver Not Detected and Electrical Idle Detect Masks	081h	–	–	–	–	***Reserved***	–	–
208h	Physical Deskew Level Low	082h	–	–	–	–	***Reserved***	–	–
20Ch	Physical Deskew Level High	083h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
210h	Phy User Test Pattern 0	084h	–	–	–	–	***Reserved***	–	–
214h	Phy User Test Pattern 4	085h	–	–	–	–	***Reserved***	–	–
218h	Phy User Test Pattern 8	086h	–	–	–	–	***Reserved***	–	–
21Ch	Phy User Test Pattern 12	087h	–	–	–	–	***Reserved***	–	–
220h	Physical Layer Command and Status	088h	–	–	–	–	***Reserved***	–	–
224h	Port Configuration	089h	–	–	–	–	***Reserved***	–	–
228h	Physical Layer Test	08Ah	–	–	–	–	***Reserved***	–	–
22Ch	Physical Layer ***(Factory Test Only)***	08Bh	–	–	–	–	***Reserved***	–	–
230h	Physical Layer Port Command	08Ch	–	–	–	–	***Reserved***	–	–
234h	Port 4 Receive Error Count; Port Control and SKIP Ordered-Set Interval	08Dh	–	–	–	–	***Reserved***	–	–
238h	Quad 0 SerDes Diagnostic Data	08Eh	–	–	–	–	***Reserved***	–	–
23Ch	Quad 1 SerDes Diagnostic Data	08Fh	–	–	–	–	***Reserved***	–	–
240h	Quad 2 SerDes Diagnostic Data	090h	–	–	–	–	***Reserved***	–	–
244h	***Reserved***	091h	–	–	–	–	***Reserved***	–	–
248h	SerDes Nominal Drive Current Select	092h	–	–	–	–	***Reserved***	–	–
24Ch	SerDes Drive Current Level Select 1	093h	–	–	–	–	***Reserved***	–	–
250h	SerDes Drive Current Level Select 2	094h	–	–	–	–	***Reserved***	–	–
254h	SerDes Drive Equalization Level Select 1	095h	–	–	–	–	***Reserved***	–	–
258h	SerDes Drive Equalization Level Select 2	096h	–	–	–	–	***Reserved***	–	–
25Ch	Physical Receive Error Count	097h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
260h	Serial EEPROM Status and Control	098h	–	–	–	–	***Reserved***	–	–
264h	Serial EEPROM Data Buffer	099h	–	–	–	–	***Reserved***	–	–
268h	***Reserved***	09Ah	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
28Ch	***Reserved***	0A3h	–	–	–	–	***Reserved***	–	–
290h	***Factory Test Only***	0A4h	–	–	–	–	***Reserved***	–	–
294h	I2C Configuration	0A5h	–	–	–	–	***Reserved***	–	–
298h	***Factory Test Only***	0A6h	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
2A8h	***Factory Test Only***	0AAh	–	–	–	–	***Reserved***	–	–
2ACh	***Reserved***	0ABh	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
2C4h	***Reserved***	0B1h	–	–	–	–	***Reserved***	–	–
2C8h	Bus Number CAM 0	0B2h	–	–	–	–	***Reserved***	–	–
2CCh	Bus Number CAM 1	0B3h	–	–	–	–	***Reserved***	–	–
2D0h	Bus Number CAM 2	0B4h	–	–	–	–	***Reserved***	–	–
2D4h	Bus Number CAM 3	0B5h	–	–	–	–	***Reserved***	–	–
2D8h	Bus Number CAM 4	0B6h	–	–	–	–	***Reserved***	–	–
2DCh	***Reserved***	0B7h	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
304h	***Reserved***	0C1h	–	–	–	–	***Reserved***	–	–
308h	I/O CAM 0 and I/O CAM 1	0C2h	–	–	–	–	***Reserved***	–	–
30Ch	I/O CAM 2 and I/O CAM 3	0C3h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
310h	I/O CAM 4	0C4h	–	–	–	–	***Reserved***	–	–
314h	***Reserved***	0C5h	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
344h	***Reserved***	0D1h	–	–	–	–	***Reserved***	–	–
348h	AMCAM 0 Memory Base and Limit	0D2h	–	–	–	–	***Reserved***	–	–
34Ch	AMCAM 0 Prefetchable Memory Base and Limit	0D3h	–	–	–	–	***Reserved***	–	–
350h	AMCAM 0 Prefetchable Memory Upper Base Address	0D4h	–	–	–	–	***Reserved***	–	–
354h	AMCAM 0 Prefetchable Memory Upper Limit Address	0D5h	–	–	–	–	***Reserved***	–	–
358h	AMCAM 1 Memory Base and Limit	0D6h	–	–	–	–	***Reserved***	–	–
35Ch	AMCAM 1 Prefetchable Memory Base and Limit	0D7h	–	–	–	–	***Reserved***	–	–
360h	AMCAM 1 Prefetchable Memory Upper Base Address	0D8h	–	–	–	–	***Reserved***	–	–
364h	AMCAM 1 Prefetchable Memory Upper Limit Address	0D9h	–	–	–	–	***Reserved***	–	–
368h	AMCAM 2 Memory Base and Limit	0DAh	–	–	–	–	***Reserved***	–	–
36Ch	AMCAM 2 Prefetchable Memory Base and Limit	0DBh	–	–	–	–	***Reserved***	–	–
370h	AMCAM 2 Prefetchable Memory Upper Base Address	0DCh	–	–	–	–	***Reserved***	–	–
374h	AMCAM 2 Prefetchable Memory Upper Limit Address	0DDh	–	–	–	–	***Reserved***	–	–
378h	AMCAM 3 Memory Base and Limit	0DEh	–	–	–	–	***Reserved***	–	–
37Ch	AMCAM 3 Prefetchable Memory Base and Limit	0DFh	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
380h	AMCAM 3 Prefetchable Memory Upper Base Address	0E0h	–	–	–	–	***Reserved***	–	–
384h	AMCAM 3 Prefetchable Memory Upper Limit Address	0E1h	–	–	–	–	***Reserved***	–	–
388h	AMCAM 4 Memory Base and Limit	0E2h	–	–	–	–	***Reserved***	–	–
38Ch	AMCAM 4 Prefetchable Memory Base and Limit	0E3h	–	–	–	–	***Reserved***	–	–
390h	AMCAM 4 Prefetchable Memory Upper Base Address	0E4h	–	–	–	–	***Reserved***	–	–
394h	AMCAM 4 Prefetchable Memory Upper Limit Address	0E5h	–	–	–	–	***Reserved***	–	–
398h	***Reserved***	0E6h	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
544h	***Reserved***	151h	–	–	–	–	***Reserved***	–	–
548h	Ingress Performance Counter	152h	–	–	–	–	***Reserved***	–	–
54Ch	***Reserved***	153h	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
65Ch	***Reserved***	197h	–	–	–	–	***Reserved***	–	–
660h	Ingress Control	198h	–	–	–	–	***Reserved***	–	–
664h	***Reserved***	199h	–	–	–	–	***Reserved***	–	–
668h	Ingress Port Enable	19Ah	–	–	–	–	***Reserved***	–	–
66Ch	***Reserved***	19Bh	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
67Ch	***Reserved***	19Fh	–	–	–	–	***Reserved***	–	–
680h	I/O CAM Upper Port 0	1A0h	–	–	–	–	***Reserved***	–	–
684h	I/O CAM Upper Port 1	1A1h	–	–	–	–	***Reserved***	–	–
688h	I/O CAM Upper Port 2	1A2h	–	–	–	–	***Reserved***	–	–
68Ch	I/O CAM Upper Port 3	1A3h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
690h	I/O CAM Upper Port 4	1A4h	–	–	–	–	***Reserved***	–	–
694h	***Reserved***	1A5h	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
6BCh	***Reserved***	1AFh	–	–	–	–	***Reserved***	–	–
6C0h	BAR0 Shadow for Port 0	1B0h	–	–	–	–	***Reserved***	–	–
6C4h	BAR1 Shadow for Port 0	1B1h	–	–	–	–	***Reserved***	–	–
6C8h	BAR0 Shadow for Port 1	1B2h	–	–	–	–	***Reserved***	–	–
6CCh	BAR1 Shadow for Port 1	1B3h	–	–	–	–	***Reserved***	–	–
6D0h	BAR0 Shadow for Port 2	1B4h	–	–	–	–	***Reserved***	–	–
6D4h	BAR1 Shadow for Port 2	1B5h	–	–	–	–	***Reserved***	–	–
6D8h	BAR0 Shadow for Port 3	1B6h	–	–	–	–	***Reserved***	–	–
6DCh	BAR1 Shadow for Port 3	1B7h	–	–	–	–	***Reserved***	–	–
6E0h	BAR0 Shadow for Port 4	1B8h	–	–	–	–	***Reserved***	–	–
6E4h	BAR1 Shadow for Port 4	1B9h	–	–	–	–	***Reserved***	–	–
6E8h	***Reserved***	1BAh	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
73Ch	***Reserved***	1CFh	–	–	–	–	***Reserved***	–	–
740h	VC0 Capability Port 0	1D0h	–	–	–	–	***Reserved***	–	–
744h	VC1 Capability Port 0	1D1h	–	–	–	–	***Reserved***	–	–
748h	VC0 Capability Port 1	1D2h	–	–	–	–	***Reserved***	–	–
74Ch	VC1 Capability Port 1	1D3h	–	–	–	–	***Reserved***	–	–
750h	VC0 Capability Port 2	1D4h	–	–	–	–	***Reserved***	–	–
754h	VC1 Capability Port 2	1D5h	–	–	–	–	***Reserved***	–	–
758h	VC0 Capability Port 3	1D6h	–	–	–	–	***Reserved***	–	–
75Ch	***Reserved***	1D7h	–	–	–	–	***Reserved***	–	–
760h	VC0 Capability Port 4	1D8h	–	–	–	–	***Reserved***	–	–
764h	***Reserved***	1D9h	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
83Ch	***Reserved***	20Fh	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
840h	Port 0 VC Capability 1	210h	–	–	–	–	***Reserved***	–	–
844h	Port 1 VC Capability 1	211h	–	–	–	–	***Reserved***	–	–
848h	Port 2 VC Capability 1	212h	–	–	–	–	***Reserved***	–	–
84Ch	Port 3 VC Capability 1	213h	–	–	–	–	***Reserved***	–	–
850h	Port 4 VC Capability 1	214h	–	–	–	–	***Reserved***	–	–
854h	***Reserved***	215h	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
9ECh	***Reserved***	27Bh	–	–	–	–	***Reserved***	–	–
9F0h	INCH Port 4 Control	27Ch	–	–	–	–	***Reserved***	–	–
9F4h	INCH FC Update Pending Timer	27Dh	–	–	–	–	***Reserved***	–	–
9F8h	***Reserved***	27Eh	–	–	–	–	***Reserved***	–	–
9FCh	INCH Mode	27Fh	–	–	–	–	***Reserved***	–	–
A00h	INCH Threshold Port 0 VC0 Posted	280h	–	–	–	–	***Reserved***	–	–
A04h	INCH Threshold Port 0 VC0 Non-Posted	281h	–	–	–	–	***Reserved***	–	–
A08h	INCH Threshold Port 0 VC0 Completion	282h	–	–	–	–	***Reserved***	–	–
A0Ch	INCH Threshold Port 0 VC1 Posted	283h	–	–	–	–	***Reserved***	–	–
A10h	INCH Threshold Port 0 VC1 Non-Posted	284h	–	–	–	–	***Reserved***	–	–
A14h	INCH Threshold Port 0 VC1 Completion	285h	–	–	–	–	***Reserved***	–	–
A18h	INCH Threshold Port 1 VC0 Posted	286h	–	–	–	–	***Reserved***	–	–
A1Ch	INCH Threshold Port 1 VC0 Non-Posted	287h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
A20h	INCH Threshold Port 1 VC0 Completion	288h	–	–	–	–	***Reserved***	–	–
A24h	INCH Threshold Port 1 VC1 Posted	289h	–	–	–	–	***Reserved***	–	–
A28h	INCH Threshold Port 1 VC1 Non-Posted	28Ah	–	–	–	–	***Reserved***	–	–
A2Ch	INCH Threshold Port 1 VC1 Completion	28Bh	–	–	–	–	***Reserved***	–	–
A30h	INCH Threshold Port 2 VC0 Posted	28Ch	–	–	–	–	***Reserved***	–	–
A34h	INCH Threshold Port 2 VC0 Non-Posted	28Dh	–	–	–	–	***Reserved***	–	–
A38h	INCH Threshold Port 2 VC0 Completion	28Eh	–	–	–	–	***Reserved***	–	–
A3Ch	INCH Threshold Port 2 VC1 Posted	28Fh	–	–	–	–	***Reserved***	–	–
A40h	INCH Threshold Port 2 VC1 Non-Posted	290h	–	–	–	–	***Reserved***	–	–
A44h	INCH Threshold Port 2 VC1 Completion	291h	–	–	–	–	***Reserved***	–	–
A48h	INCH Threshold Port 3 VC0 Posted	292h	–	–	–	–	***Reserved***	–	–
A4Ch	INCH Threshold Port 3 VC0 Non-Posted	293h	–	–	–	–	***Reserved***	–	–
A50h	INCH Threshold Port 3 VC0 Completion	294h	–	–	–	–	***Reserved***	–	–
A54h	***Reserved***	295h	–	–	–	–	***Reserved***	–	–
A58h	***Reserved***	296h	–	–	–	–	***Reserved***	–	–
A5Ch	***Reserved***	297h	–	–	–	–	***Reserved***	–	–
A60h	INCH Threshold Port 4 VC0 Posted	298h	–	–	–	–	***Reserved***	–	–
A64h	INCH Threshold Port 4 VC0 Non-Posted	299h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
A68h	INCH Threshold Port 4 VC0 Completion	29Ah	–	–	–	–	***Reserved***	–	–
A6Ch	***Reserved***	29Bh	–	–	–	–	***Reserved***	–	–
…	…	…	–	–	–	–	…	–	–
BDCh	***Reserved***	2F7h	–	–	–	–	***Reserved***	–	–
BE0h	***Factory Test Only / Reserved***	2F8h	–	–	–	–	***Reserved***	–	–
BE4h	Device-Specific Relaxed Ordering Mode Port 4	2F9h	–	–	–	–	***Reserved***	–	–
BE8h	One-Bit ECC Error Count	2FAh	–	–	–	–	***Reserved***	–	–
BECh	***Factory Test Only***	2FBh	–	–	–	–	***Reserved***	–	–
BF0h	***Factory Test Only***	2FCh	–	–	–	–	***Reserved***	–	–
BF4h	***Factory Test Only***	2FDh	–	–	–	–	***Reserved***	–	–
BF8h	***Factory Test Only***	2FEh	–	–	–	–	***Reserved***	–	–
BFCh	Device-Specific Relaxed Ordering Mode	2FFh	–	–	–	–	***Reserved***	–	–
C00h	ITCH VC&T Threshold 1	300h	–	–	–	–	***Reserved***	–	–
C04h	ITCH VC&T Threshold 2	301h	–	–	–	–	***Reserved***	–	–
C08h	ITCH VC&T Threshold 3	302h	–	–	–	–	***Reserved***	–	–
C0Ch	Port 0 VC0 Posted and Non-Posted Queue Status	303h	–	–	–	–	***Reserved***	–	–
C10h	Port 0 VC0 Completion and VC1 Posted Queue Status	304h	–	–	–	–	***Reserved***	–	–
C14h	Port 0 VC1 Non-Posted and Completion Queue Status	305h	–	–	–	–	***Reserved***	–	–
C18h	Port 1 VC0 Posted and Non-Posted Queue Status	306h	–	–	–	–	***Reserved***	–	–
C1Ch	Port 1 VC0 Completion and VC1 Posted Queue Status	307h	–	–	–	–	***Reserved***	–	–

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
C20h	Port 1 VC1 Non-Posted and Completion Queue Status	308h	–	–	–	–	***Reserved***	–	–
C24h	Port 2 VC0 Posted and Non-Posted Queue Status	309h	–	–	–	–	***Reserved***	–	–
C28h	Port 2 VC0 Completion and VC1 Posted Queue Status	30Ah	–	–	–	–	***Reserved***	–	–
C2Ch	Port 2 VC1 Non-Posted and Completion Queue Status	30Bh	–	–	–	–	***Reserved***	–	–
C30h	Port 3 VC0 Posted and Non-Posted Queue Status	30Ch	–	–	–	–	***Reserved***	–	–
C34h	Port 3 VC0 Completion Queue Status	30Dh	–	–	–	–	***Reserved***	–	–
C38h	***Reserved***	30Eh	–	–	–	–	***Reserved***	–	–
C3Ch	***Reserved***	–	–	–	–	–	Memory BAR2/3 Address Translation Lower (Link) Memory BAR2 Address Translation (Virtual)	5C1h	6B9h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
C40h	***Reserved***	–	–	–	–	–	Memory BAR2/3 Address Translation Upper (Link) ***Reserved*** (Virtual)	5C2h	6BAh
C44h	***Reserved***	–	–	–	–	–	Memory BAR4/5 Address Translation Lower	5C3h	6BBh
C48h	***Reserved***	–	–	–	–	–	Memory BAR4/5 Address Translation Upper	5C4h	6BCh
C4Ch	***Reserved***	–	–	–	–	–	Memory BAR2/3 Limit Lower Address (Link) Memory BAR2 Limit Address (Virtual)	5C5h	6BDh
C50h	***Reserved***	–	–	–	–	–	Memory BAR2/3 Limit Upper Address (Link) ***Reserved*** (Virtual)	5C6h	6BEh
C54h	***Reserved***	–	–	–	–	–	Memory BAR4/5 Limit Lower Address	5C7h	6BFh
C58h	***Reserved***	–	–	–	–	–	Memory BAR4/5 Limit Upper Address	5C8h	6C0h
C5Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 0 (Virtual)	5C9h	6C1h
C60h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 1 (Virtual)	5CAh	6C2h
C64h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 2 (Virtual)	5CBh	6C3h
C68h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 3 (Virtual)	5CCh	6C4h
C6Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 4 (Virtual)	5CDh	6C5h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
C70h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 5 (Virtual)	5CEh	6C6h
C74h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 6 (Virtual)	5CFh	6C7h
C78h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 7 (Virtual)	5D0h	6C8h
C7Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 8 (Virtual)	5D1h	6C9h
C80h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 9 (Virtual)	5D2h	6CAh
C84h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 10 (Virtual)	5D3h	6CBh
C88h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 11 (Virtual)	5D4h	6CCh
C8Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 12 (Virtual)	5D5h	6CDh
C90h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 13 (Virtual)	5D6h	6CEh
C94h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 14 (Virtual)	5D7h	6CFh
C98h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 15 (Virtual)	5D8h	6D0h
C9Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 16 (Virtual)	5D9h	6D1h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
CA0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 17 (Virtual)	5DAh	6D2h
CA4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 18 (Virtual)	5DBh	6D3h
CA8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 19 (Virtual)	5DCh	6D4h
CACh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 20 (Virtual)	5DDh	6D5h
CB0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 21 (Virtual)	5DEh	6D6h
CB4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 22 (Virtual)	5DFh	6D7h
CB8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 23 (Virtual)	5E0h	6D8h
CBCh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 24 (Virtual)	5E1h	6D9h
CC0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 25 (Virtual)	5E2h	6DAh
CC4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 26 (Virtual)	5E3h	6DBh
CC8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 27 (Virtual)	5E4h	6DCh
CCCh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 28 (Virtual)	5E5h	6DDh

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
CD0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 29 (Virtual)	5E6h	6DEh
CD4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 30 (Virtual)	5E7h	6DFh
CD8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 31 (Virtual)	5E8h	6E0h
CDCh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 32 (Virtual)	5E9h	6E1h
CE0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 33 (Virtual)	5EAh	6E2h
CE4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 34 (Virtual)	5EBh	6E3h
CE8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 35 (Virtual)	5ECh	6E4h
CECh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 36 (Virtual)	5EDh	6E5h
CF0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 37 (Virtual)	5EEh	6E6h
CF4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 38 (Virtual)	5EFh	6E7h
CF8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 39 (Virtual)	5F0h	6E8h
CFCh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 40 (Virtual)	5F1h	6E9h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
D00h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 41 (Virtual)	5F2h	6EAh
D04h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 42 (Virtual)	5F3h	6EBh
D08h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 43 (Virtual)	5F4h	6ECh
D0Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 44 (Virtual)	5F5h	6EDh
D10h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 45 (Virtual)	5F6h	6EEh
D14h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 46 (Virtual)	5F7h	6EFh
D18h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 47 (Virtual)	5F8h	6F0h
D1Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 48 (Virtual)	5F9h	6F1h
D20h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 49 (Virtual)	5FAh	6F2h
D24h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 50 (Virtual)	5FBh	6F3h
D28h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 51 (Virtual)	5FCh	6F4h
D2Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 52 (Virtual)	5FDh	6F5h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
D30h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 53 (Virtual)	5FEh	6F6h
D34h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 54 (Virtual)	5FFh	6F7h
D38h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 55 (Virtual)	600h	6F8h
D3Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 56 (Virtual)	601h	6F9h
D40h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 57 (Virtual)	602h	6FAh
D44h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 58 (Virtual)	603h	6FBh
D48h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 59 (Virtual)	604h	6FCh
D4Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 60 (Virtual)	605h	6FDh
D50h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 61 (Virtual)	606h	6FEh
D54h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 62 (Virtual)	607h	6FFh
D58h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Lookup Table Entry 63 (Virtual)	608h	700h
D5Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) NT Link Interface VC0 Resource Control (Shadow Copy) (Virtual)	609h	701h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
D60h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) NT Link Interface VC1 Resource Control (Shadow Copy) (Virtual)	60Ah	702h
D64h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) NT Link Interface VC Capability 1 (Shadow Copy) (Virtual)	60Bh	703h
D68h	***Reserved***	–	–	–	–	–	BAR0 (***Reserved***)	60Ch	704h
D6Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) BAR1 (Shadow Copy) (Virtual)	60Dh	705h
D70h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) BAR2 (Shadow Copy) (Virtual)	60Eh	706h
D74h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) BAR3 (Shadow Copy) (Virtual)	60Fh	707h
D78h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) BAR4 (Shadow Copy) (Virtual)	610h	708h
D7Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) BAR5 (Shadow Copy) (Virtual)	611h	709h
D80h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) BAR1 Setup (Shadow Copy) (Virtual)	612h	70Ah
D84h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Memory BAR2 Setup (Shadow Copy) (Virtual)	613h	70Bh
D88h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Memory BAR3 Setup (Shadow Copy) (Virtual)	614h	70Ch
D8Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Memory BAR4 Setup (Shadow Copy) (Virtual)	615h	70Dh

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
D90h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Memory BAR5 Setup (Shadow Copy) (Virtual)	616h	70Eh
D94h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 0 (Virtual)	617h	70Fh
D98h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 1 (Virtual)	618h	710h
D9Ch	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 2 (Virtual)	619h	711h
DA0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 3 (Virtual)	61Ah	712h
DA4h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 4 (Virtual)	61Bh	713h
DA8h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 5 (Virtual)	61Ch	714h
DACh	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 6 (Virtual)	61Dh	715h
DB0h	***Reserved***	–	–	–	–	–	***Reserved*** (Link) Requester ID Translation LUT Entry 7 (Virtual)	61Eh	716h
DB4h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 0_1 (Link)	61Fh	717h
DB8h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 2_3 (Link)	620h	718h
DBCh	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 4_5 (Link)	621h	719h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
DC0h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 6_7 (Link)	622h	71Ah
DC4h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 8_9 (Link)	623h	71Bh
DC8h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 10_11 (Link)	624h	71Ch
DCCh	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 12_13 (Link)	625h	71Dh
DD0h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 14_15 (Link)	626h	71Eh
DD4h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 16_17 (Link)	627h	71Fh
DD8h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 18_19 (Link)	628h	720h
DDCh	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 20_21 (Link)	629h	721h
DE0h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 22_23 (Link)	62Ah	722h
DE4h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 24_25 (Link)	62Bh	723h
DE8h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 26_27 (Link)	62Ch	724h
DECh	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 28_29 (Link)	62Dh	725h

Table A-1. PEX 8508 Serial EEPROM Memory Map *(Cont.)*

Register Offset	Transparent Mode						Non-Transparent Mode		
	Register Name	Port Register Loaded from Listed Serial EEPROM Address					Register Name	Port Register Loaded from Listed Serial EEPROM Address	
		Port 0	Port 1	Port 2	Port 3	Port 4		Link	Virtual
DF0h	***Reserved***	–	–	–	–	–	Requester ID Translation LUT Entry 30_31 (Link)	62Eh	726h
DF4h	***Reserved***	–	–	–	–	–	***Reserved***	–	727h
DF8h	***Reserved***	–	–	–	–	–	***Reserved***	–	728h
DFCh	***Reserved***	–	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
FB0h	***Reserved***	–	–	–	–	–	***Reserved***	–	–
FB4h	Advanced Error Reporting Enhanced Capability Header	30Fh	399h	423h	4ADh	537h	Advanced Error Reporting Enhanced Capability Header	62Fh	729h
FB8h	Uncorrectable Error Status	310h	39Ah	424h	4AEh	538h	Uncorrectable Error Status	630h	72Ah
FBCh	Uncorrectable Error Mask	311h	39Bh	425h	4AFh	539h	Uncorrectable Error Mask	631h	72Bh
FC0h	Uncorrectable Error Severity	312h	39Ch	426h	4B0h	53Ah	Uncorrectable Error Severity	632h	72Ch
FC4h	Correctable Error Status	313h	39Dh	427h	4B1h	53Bh	Correctable Error Status	633h	72Dh
FC8h	Correctable Error Mask	314h	39Eh	428h	4B2h	53Ch	Correctable Error Mask	634h	72Eh
FCCh	Advanced Error Capabilities and Control	315h	39Fh	429h	4B3h	53Dh	Advanced Error Capabilities and Control	635h	72Fh
FD0h	Header Log 0	316h	3A0h	42Ah	4B4h	53Eh	Header Log 0	636h	730h
FD4h	Header Log 1	317h	3A1h	42Bh	4B5h	53Fh	Header Log 1	637h	731h
FD8h	Header Log 2	318h	3A2h	42Ch	4B6h	540h	Header Log 2	638h	732h
FDCh	Header Log 3	319h	3A3h	42Dh	4B7h	541h	Header Log 3	639h	733h
FE0h	***Reserved***	–	–	–	–	–	***Reserved***	–	–
…	…	…	…	…	…	…	…	…	…
FFCh	***Reserved***	–	–	–	–	–	***Reserved***	–	–
N/A	CRC Value								1E3Ch

Appendix B General Information

B.1 Product Ordering Information

Contact your local PLX Sales Representative for ordering information.

Table B-1. Product Ordering Information

Part Number	Description
PEX8508-AC25BI	PEX 8508AC 5-port, 8-Lane PCI Express Switch Plastic BGA-H (19-mm square, 296-ball) package
PEX8508-AC25BI G	PEX 8508AC 5-port, 8-Lane PCI Express Switch Plastic BGA-H (19-mm square, 296-ball), Lead-Free RoHS Green package
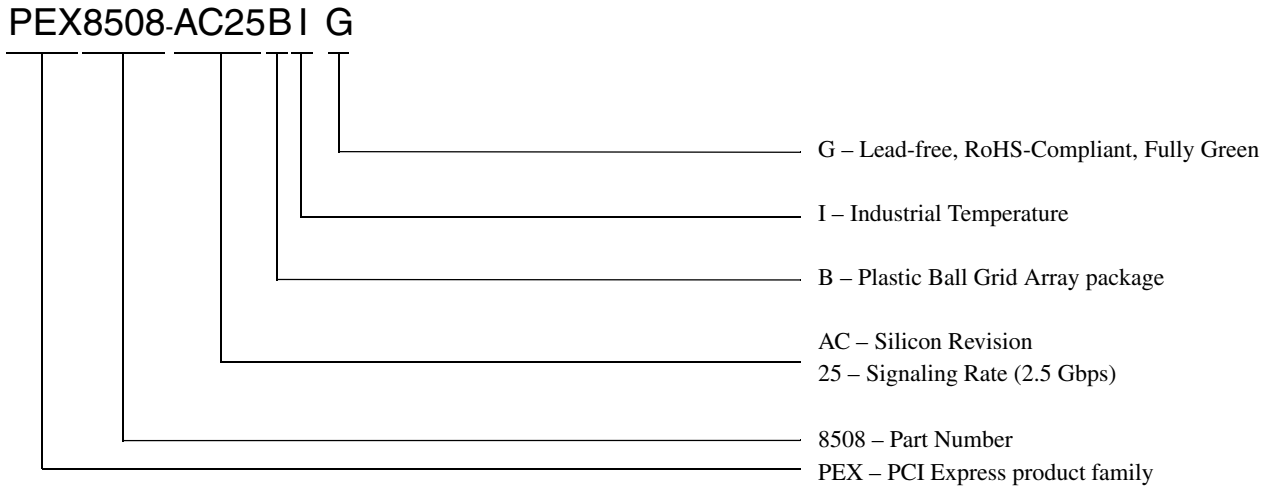	
PEX 8508-AC RDK	PEX 8508AC Rapid Development Kit with x4 edge connector
x1 Adapter	PCI Express x4 to x1 Adapter

B.2 United States and International Representatives and Distributors

PLX Technology, Inc., representatives and distributors are listed at www.plxtech.com.

B.3 Technical Support

PLX Technology, Inc., technical support information is listed at www.plxtech.com/support/, or call 800 759-3735 (domestic only) or 408 774-9060.