



ExpressLane PEX 8509-AA 8-Lane/8-Port PCI Express Gen 1 Switch Data Book

Version 1.1

April 2009

Website www.plxtech.com

Technical Support www.plxtech.com/support

Phone 800 759-3735

408 774-9060

FAX 408 774-2169

Revision History

Version	Date	Description of Changes
1.0	November, 2007	Production Release, Silicon Revision AA.
1.1	April, 2009	<p>Production update, Silicon Revision AA.</p> <p>Reorganized chapter sequence, omitted the Software Architecture chapter, and renumbered all chapters accordingly.</p> <p>Changed register offsets 208h and 20Ch to <i>Factory Test Only</i>.</p> <p>Replaced mechanical drawing (updated Table 16-2 and Figure 16-1), and omitted references to heat spreader.</p> <p>Applied miscellaneous corrections, changes, and enhancements throughout data book.</p>

Copyright Information

Copyright © 2007 – 2009 PLX Technology, Inc. All Rights Reserved. The information in this document is proprietary and confidential to PLX Technology. No part of this document may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from PLX Technology.

PLX Technology provides this documentation without warranty, term or condition of any kind, either express or implied, including, but not limited to, express and implied warranties of merchantability, fitness for a particular purpose, and non-infringement. While the information contained herein is believed to be accurate, such information is preliminary, and no representations or warranties of accuracy or completeness are made. In no event will PLX Technology be liable for damages arising directly or indirectly from any use of or reliance upon the information contained in this document. PLX Technology may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

PLX Technology retains the right to make changes to this product at any time, without notice. Products may have minor variations to this publication, known as errata. PLX Technology assumes no liability whatsoever, including infringement of any patent or copyright, for sale and use of PLX Technology products.

PLX Technology and the PLX logo are registered trademarks and ExpressLane is a trademark of PLX Technology, Inc.

PCI Express is a trademark of the PCI Special Interest Group (PCI-SIG).

EUI-64 is a trademark of The Institute of Electrical and Electronics Engineers, Inc. (IEEE)

All product names are trademarks, registered trademarks, or service marks of their respective owners.

Document Number: 8509-AA-SIL-DB-P1-1.1

Preface

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

Audience

This data book provides functional details of PLX Technology's ExpressLane PEX 8509-AA 8-Lane/8-Port PCI Express Gen 1 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 USA
Tel: 800 759-3735 (domestic only) or 408 774-9060, Fax: 408 774-2169, www.plxtech.com
The [PLX PEX 8509 Toolbox](#) includes this data book, as well as other PEX 8509 documentation, including the Errata.
- 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 1.0a*
 - [PCI Express Architecture PCI Express Jitter and BER White Paper, Revision 1.0](#)
- The Institute of Electrical and Electronics Engineers, Inc. (IEEE)
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 ExpressCard 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 ² C Bus v2.1 I2C Bus v2.1 ^a	The I ² C-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.
BAR	Base Address register.
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.
CSRs	Configuration Space registers.
Downstream Station	A station that contains only downstream ports.
ECC	Error-Correcting Code.
GPIO	General-Purpose Input/Output.
GPU	Graphics Processing Unit.
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	A bidirectional pair of differential PCI Express I/O signals.
Local	Reference to PCI Express attributes (<i>such as</i> , credits) that belong to the PCI Express station.
LTSSM	Link Training and Status State Machine.

Terms and Abbreviations	Definitions
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 logic.
PEX	PCI Express.
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.
PRBS	Pseudo-Random Bit Sequence.
QoS	Quality of Service.
RoHS	Restrictions on the use of certain Hazardous Substances (RoHS) Directive.
RAS	Reliability, Availability, and Serviceability.
RM	Read Margin.
RR	Round-Robin scheduling.
SerDes	Serializer and De-serializer. A high-speed differential-signaling parallel-to-serial and serial-to-parallel conversion logic attached to lane pads.
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.
Upstream station	Upstream station. Contains the component's upstream port. An upstream station might contain downstream ports.
UTP	User Test Pattern.
VC	Virtual Channel.
WRR	Weighted Round-Robin scheduling.

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_XXXn[x] PEX_XXXp[x]	When the signal name appears in all CAPS, with the primary Port description listed first, field [x] indicates the number associated with the signal balls/pads assigned to a specific SerDes module/Lane. The lowercase “n” (negative) or “p” (positive) suffix indicates the differential pair of signals, which are always used together.
# = Active-Low signals	Unless specified otherwise, Active-Low signals are identified by a “#” appended to the term (<i>for example</i> , PEX_PERST#).
Program/code samples	Monospace font (<i>program or code samples</i>) 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.
<i>Parity Error Detected</i>	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 (<i>for example</i> , 01b, 010b, 1010b, and so forth). Not used with single-digit values of 0 or 1.
0 through 9	Decimal numbers, or single binary numbers.
byte	Eight bits – abbreviated to “B” (<i>for example</i> , 4B = 4 bytes)
LSB	Least-Significant Byte.
lsb	Least-significant bit.
MSB	Most-Significant Byte.
msb	Most-significant bit.
DWord	Double-Word (32 bits) is the primary register size in these devices.
QWord	Quad-Word (64 bits).
Reserved	Do not modify Reserved bits and words. Unless specified otherwise, these bits read as 0 and must be written as 0.

Contents

Chapter 1	Introduction	1
	1.1 Features	1
	1.2 Overview	3
Chapter 2	Features and Applications	5
	2.1 Flexible and Feature-Rich 8-Lane/8-Port Switch	5
	2.1.1 Highly Flexible Port Configurations	5
	2.1.2 High Performance	5
	2.1.3 End-to-End Packet Integrity	5
	2.1.4 Configuration Flexibility	5
	2.1.5 Interoperability	5
	2.1.6 Low Power with Granular SerDes Control	6
	2.1.7 Flexible Port-Width Configuration	6
	2.1.8 Hot Plug for High Availability	6
	2.1.9 Dynamic Lane Reversal	6
	2.1.10 Fully Compliant Power Management	6
	2.2 Applications	7
	2.2.1 Adapter Board Aggregation	7
	2.2.2 Control Plane Usage	8
	2.2.3 Docking Station	10
	2.3 Software Usage Model	11
	2.3.1 System Configuration	11
	2.3.2 Interrupt Sources and Events	11
Chapter 3	Signal Ball Description	13
	3.1 Introduction	13
	3.2 Abbreviations	13
	3.3 Internal Pull-Up Resistors	14
	3.4 Signal Ball Descriptions	14
	3.4.1 PCI Express Signals	15
	3.4.2 Hot Plug Signals	16
	3.4.3 Serial EEPROM Signals	20
	3.4.4 Strapping Signals	21
	3.4.5 JTAG Interface Signals	22
	3.4.6 I ² C Slave Interface Signals	23
	3.4.7 Device-Specific Signals	23
	3.4.8 No Connect Signals	25
	3.4.9 Power and Ground Signals	26
	3.5 Ball Assignments by Location	27
	3.6 Physical Layout	33
Chapter 4	Functional Overview	35
	4.1 Architecture	35
	4.1.1 Ingress and Egress Functions	36
	4.1.2 Station and Port Functions	36
	4.1.2.1 Port Combinations	36
	4.1.2.2 Port Numbering	38

4.2	PCI-Compatible Software Model	39
4.2.1	System Reset	40
4.2.2	Interrupts	40
4.2.2.1	Interrupt Sources or Events	40
4.2.2.2	INTx Switch Mapping	40
4.3	PCI Express Station Functional Description	41
4.3.1	Functional Blocks	41
4.3.1.1	Physical Layer	43
4.3.1.2	Data Link Layer	43
4.3.1.3	Transaction Layer Control	43
4.3.1.4	Non-Blocking Crossbar Switch Architecture	44
4.3.2	Cut-Thru Mode	44
Chapter 5	Reset and Initialization	45
5.1	Reset Overview	45
5.1.1	Cold Reset	45
5.1.2	Warm Reset	45
5.1.3	Hot Reset	46
5.1.3.1	Hot Reset Propagation	46
5.1.3.2	Hot Reset Disable	47
5.1.4	Secondary Bus Reset	47
5.2	Initialization Procedure	48
5.2.1	Default Port Configuration	49
5.2.2	Default Register Initialization	49
5.2.3	Device-Specific Registers	49
5.2.4	Reset and Clock Initialization Timing	50
5.2.4.1	Serial EEPROM Load Time	51
5.2.4.2	I ² C Load Time	51
Chapter 6	Serial EEPROM Controller	53
6.1	Overview	53
6.2	Serial EEPROM Data Format	55
6.3	Serial EEPROM Initialization	57
6.4	PCI Express Configuration, Control, and Status Registers	57
6.5	Serial EEPROM Registers	57
6.6	Serial EEPROM Random Write/Read Access	58
6.6.1	Writing to Serial EEPROM	58
6.6.2	Reading from Serial EEPROM	59
Chapter 7	I²C Slave Interface Operation	61
7.1	Introduction	61
7.2	I ² C Support Overview	61
7.3	I ² C Addressing – Slave Mode Access	62
7.4	Command Phase Format	62
7.5	I ² C Interface Register	64
7.6	I ² C Register Write Access	64
7.6.1	Register Write	66
7.7	I ² C Register Read Access	68
7.7.1	Register Read Address Phase and Command Packet	70
7.7.2	Register Read Data Packet	71

Chapter 8	Performance Metrics	73
	8.1 Introduction	73
	8.2 Throughput	73
	8.2.1 Shared Wire	73
	8.2.2 Unidirectional Throughput	73
	8.2.3 Ideal PCI Express Throughput	75
	8.2.4 Bidirectional PCI Express Throughput	76
	8.3 DLLP Policies	80
	8.3.1 ACK DLLP Policy	81
	8.3.2 UpdateFC DLLP Policy	82
	8.3.3 Unidirectional DLLP Policies	83
	8.4 Adjusting Ingress Resources	84
	8.4.1 Initial Credit Allocation	84
	8.4.2 Common Credit Pool	85
	8.4.3 Wait for ACK	85
	8.5 Latency	86
	8.5.1 Host-Centric Latency	87
	8.5.2 Peer-to-Peer Latency	87
	8.5.3 Other Latency Measurements	87
	8.6 Queuing Options	88
	8.6.1 Destination Queue	88
	8.6.2 Source Queue	89
Chapter 9	Device Layers	91
	9.1 Data Flow Through	91
	9.2 Physical Layer	92
	9.2.1 PHY Status and Command Registers	93
	9.2.2 Hardware Link Interface Configuration	93
	9.3 Data Link Layer	95
	9.3.1 Data Link Layer Packet	95
	9.3.1.1 DLLP Ingress	95
	9.3.1.2 DLLP Egress	96
	9.3.2 Packet Arbiter	96
	9.4 Transaction Layer	97
	9.4.1 Virtual Channel and Traffic Classes	99
	9.4.2 TL Transmit/Egress Protocol	99
	9.4.2.1 Headers	99
	9.4.2.2 Data Payloads	99
	9.4.2.3 End-to-End Cyclic Redundancy Check	99
	9.4.3 TL Receive/Ingress Protocol	100
	9.4.4 Flow Control Protocol	100
Chapter 10	Interrupts	101
	10.1 Interrupt Support	101
	10.1.1 Interrupt Handling	102
	10.2 INTx Emulation Support	102
	10.2.1 INTx-Type Interrupt Message Remapping and Collapsing	103
	10.3 Message Signaled Interrupt Support	104
	10.3.1 MSI Operation	104
	10.3.2 MSI Capability Registers	104
	10.4 PEX_INTA# Interrupts	104

Chapter 11 Hot Plug Support105

- 11.1 Hot Plug Purpose and Capability 105
 - 11.1.1 Hot Plug Controller Capabilities 105
 - 11.1.2 Hot Plug Port External Signals 105
 - 11.1.3 Hot Plug Output Signal States for Disabled Hot Plug Slots 105
- 11.2 PCI Express Capability Registers for Hot Plug 106
- 11.3 Hot Plug Interrupts 106
- 11.4 Hot Plug Controller Slot Power-Up/Down Sequence 107
 - 11.4.1 Slot Power-Up Sequence 107
 - 11.4.1.1 Configuring Slot Power-Up Sequence Features with Serial EEPROM . . . 108
 - 11.4.1.2 Slot Power-Up Sequencing when Power Controller Present Bit Is Set . . . 109
 - 11.4.1.3 HP_PERSTx# (Reset) and HP_PWRLEDx# Output Power-Up Sequencing when Power Controller Present Bit Is Clear 111
 - 11.4.1.4 Disabling Power-Up Hot Plug Output Sequencing 112
 - 11.4.2 Slot Power-Down Sequence 112
- 11.5 Hot Plug Board Insertion and Removal Process 113

Chapter 12 Power Management117

- 12.1 Overview 117
- 12.2 Features 118
- 12.3 Power Management Capability 119
 - 12.3.1 Device Power Management States 119
 - 12.3.1.1 D0 Device Power Management State 119
 - 12.3.1.2 D3hot Device Power Management State 119
 - 12.3.2 Link Power Management States 120
 - 12.3.3 PCI Express Power Management Support 121
- 12.4 Power Management Tracking 131
- 12.5 Power Management Event Handler 131

Chapter 13 Port Registers133

- 13.1 Introduction 133
- 13.2 Type 1 Port Register Map 134
- 13.3 Port Register Configuration and Map 136
- 13.4 Register Access 137
 - 13.4.1 PCI r3.0-Compatible Configuration Mechanism 137
 - 13.4.2 PCI Express Enhanced Configuration Mechanism 138
 - 13.4.3 Device-Specific Memory-Mapped Configuration Mechanism 139
- 13.5 Register Descriptions 140
- 13.6 Type 1 Configuration Header Registers 141
- 13.7 Power Management Capability Registers 156
- 13.8 Message Signaled Interrupt Capability Registers 159
- 13.9 PCI Express Capability Registers 161
- 13.10 Subsystem ID and Subsystem Vendor ID Capability Registers 180
- 13.11 Vendor-Specific Enhanced Capability Registers 181
- 13.12 Device Serial Number Extended Capability Registers 183
- 13.13 Power Budget Extended Capability Registers 185
- 13.14 Virtual Channel Extended Capability Registers 188
- 13.15 Port Arbitration Table Registers 192
- 13.16 Device-Specific Registers 197
 - 13.16.1 Device-Specific Registers – Error Checking and Debug 199
 - 13.16.2 Device-Specific Registers – Physical Layer 213
 - 13.16.3 Device-Specific Registers – I²C Interface 239
 - 13.16.4 Device-Specific Registers – Bus Number CAM 240
 - 13.16.5 Device-Specific Registers – I/O CAM 244

13.16.6	Device-Specific Registers – Address-Mapping CAM	248
13.16.7	Device-Specific Registers – Ingress Control and Port Enable	258
13.16.8	Device-Specific Registers – I/O CAM Base and Limit Upper 16 Bits	262
13.16.9	Device-Specific Registers – Base Address Shadow	264
13.16.10	Device-Specific Registers – Shadow Virtual Channel Capability	273
13.16.11	Device-Specific Registers – Ingress Credit Handler	277
13.16.12	Device-Specific Registers – Port Configuration Header	307
13.16.13	Device-Specific Registers – Source Queue Weight and Soft Error	331
13.17	Advanced Error Reporting Extended Capability Registers	335
Chapter 14	Test and Debug	343
14.1	Physical Layer Loopback Operation	343
14.1.1	Overview	343
14.1.2	Internal Loopback Mode	344
14.1.3	Analog Loopback Master Mode	345
14.1.4	Digital Loopback Master Mode	347
14.1.5	Analog Loopback Slave Mode	348
14.1.6	Digital Loopback Slave Mode	348
14.2	Using the Diagnostic Registers	349
14.3	Pseudo-Random and Bit-Pattern Generation	350
14.4	JTAG Interface	351
14.4.1	IEEE 1149.1 and IEEE 1149.6 Test Access Port	351
14.4.2	JTAG Instructions	352
14.4.3	JTAG Boundary Scan	353
14.4.4	JTAG Reset Input Signal JTAG_TRST#	353
14.5	Lane Good Status LEDs	354
Chapter 15	Electrical Specifications	355
15.1	Introduction	355
15.2	Power-Up/Power-Down Sequence	355
15.3	Absolute Maximum Ratings	355
15.4	Power Characteristics	356
15.5	Power Consumption	356
15.6	I/O Interface Signal Groupings	357
15.7	Transmit Drive Characteristics	365
15.7.1	Drive Current	365
15.7.2	Transmit Equalization	365
15.7.3	Transmit Termination Adjust	365
15.8	Receive Characteristics	365
15.8.1	Receive Equalization	365
15.8.2	Receive Termination Adjust	365
Chapter 16	Thermal and Mechanical Specifications	367
16.1	Thermal Characteristics	367
16.1.1	Sample Thermal Data	367
16.2	General Package Specifications	368
16.3	Mechanical Dimensions	369
Appendix A	General Information	371
A.1	Product Ordering Information	371
A.2	United States and International Representatives and Distributors	372
A.3	Technical Support	372

THIS PAGE INTENTIONALLY LEFT BLANK.

Registers

Type 1 Configuration Header Registers	141
13-1. 00h PCI Configuration ID	141
13-2. 04h PCI Command/Status	142
13-3. 08h Class Code and PCI Revision ID	144
13-4. 0Ch Miscellaneous Control	145
13-5. 10h Base Address 0	145
13-6. 14h Base Address 1	146
13-7. 18h Bus Number	146
13-8. 1Ch Secondary Status, I/O Limit, and I/O Base	147
13-9. 20h Memory Base and Limit	149
13-10. 24h Prefetchable Memory Base and Limit	150
13-11. 28h Prefetchable Memory Base Upper 32 Bits	151
13-12. 2Ch Prefetchable Memory Limit Upper 32 Bits	151
13-13. 30h I/O Base and Limit Upper 16 Bits	151
13-14. 34h New Capability Pointer	152
13-15. 38h Expansion ROM Base Address	152
13-16. 3Ch Bridge Control and PCI Interrupt Signal	152
Power Management Capability Registers	156
13-17. 40h Power Management Capability	156
13-18. 44h Power Management Status and Control	157
Message Signaled Interrupt Capability Registers	159
13-19. 48h MSI Control and Capability Header	159
13-20. 4Ch MSI Address	160
13-21. 50h MSI Upper Address	160
13-22. 54h MSI Data	160
13-23. 58h MSI Mask	160
13-24. 5Ch MSI Pending	160
PCI Express Capability Registers	161
13-25. 68h PCI Express Capability List and Capability	162
13-26. 6Ch Device Capability	163
13-27. 70h Device Status and Control	164
13-28. 74h Link Capability	167
13-29. 78h Link Status and Control	169
13-30. 7Ch Slot Capability	171
13-31. 80h Slot Status and Control	174
Subsystem ID and Subsystem Vendor ID Capability Registers	180
13-32. 90h Subsystem Capability	180
13-33. 94h Subsystem ID and Subsystem Vendor ID	180
Vendor-Specific Enhanced Capability Registers	181
13-34. DCh Vendor-Specific Enhanced Capability	181
13-35. E0h PLX Hardwired Configuration ID	182
13-36. E4h PLX Hardwired Revision ID	182
13-37. E8h Scratch A	182
13-38. ECh Scratch B	182
Device Serial Number Extended Capability Registers	183
13-39. 100h Device Serial Number Enhanced Capability	183
13-40. 104h Serial Number (Lower DW)	184
13-41. 108h Serial Number (Upper DW)	184

Power Budget Extended Capability Registers	185
13-42. 138h Power Budget Extended Capability Header	185
13-43. 13Ch Data Select	185
13-44. 140h Power Budget Data	186
13-45. 144h Power Budget Capability	187
Virtual Channel Extended Capability Registers	188
13-46. 148h Virtual Channel Extended Capability	188
13-47. 14Ch Port VC Capability 1	189
13-48. 154h Port VC Status and Control	189
13-49. 158h VC0 Resource Capability	190
13-50. 15Ch VC0 Resource Control	190
13-51. 160h VC0 Resource Status	191
Port Arbitration Table Registers	192
13-52. 1A8h Port Arbitration Table Phases 0 to 7	193
13-53. 1ACh Port Arbitration Table Phases 8 to 15	194
13-54. 1B0h Port Arbitration Table Phases 16 to 23	195
13-55. 1B4h Port Arbitration Table Phases 24 to 31	196
Device-Specific Registers	197
Device-Specific Registers – Error Checking and Debug	199
13-56. 1C0h Device-Specific Error Status for Egress ECC Error	200
13-57. 1C4h Device-Specific Error Mask for Egress ECC Error	201
13-58. 1C8h ECC Error Check Disable	202
13-59. 1CCh Error Handler 32-Bit Error Status	203
13-60. 1D0h Error Handler 32-Bit Error Mask	205
13-61. 1DCh Debug Control	207
13-62. 1E0h Power Management Hot Plug User Configuration	209
13-63. 1E8h Bad TLP Count	210
13-64. 1ECh Bad DLLP Count	210
13-65. 1F4h Lane Status/Software PEX_LANE_GOODx# LED Control	211
13-66. 1F8h ACK Transmission Latency Limit	212
Device-Specific Registers – Physical Layer	213
13-67. 204h Physical Layer Receiver Not Detected and Electrical Idle Detect Masks	215
13-68. 210h Physical Layer User Test Pattern 0	216
13-69. 214h Physical Layer User Test Pattern 4	216
13-70. 218h Physical Layer User Test Pattern 8	216
13-71. 21Ch Physical Layer User Test Pattern 12	216
13-72. 220h Physical Layer Command and Status	217
13-73. 224h Port Configuration	218
13-74. 228h Physical Layer Test	219
13-75. 22Ch Physical Layer	221
13-76. 230h Physical Layer Port Command	223
13-77. 234h SKIP Ordered-Set Interval and Port Control	231
13-78. 238h SerDes Quad 0 Diagnostic Data	232
13-79. 23Ch SerDes Quad 1 Diagnostic Data	233
13-80. 248h SerDes Nominal Drive Current Select	233
13-81. 24Ch SerDes Drive Current Level 1	234
13-82. 254h SerDes Drive Equalization Level Select 1	234
13-83. 260h Serial EEPROM Status and Control	235
13-84. 264h Serial EEPROM Data Buffer	238
13-85. 268h Serial EEPROM Clock Frequency	238
13-86. 26Ch Serial EEPROM 3 rd Address Byte	238
Device-Specific Registers – I²C Interface	239
13-87. 294h I ² C Configuration	239

Device-Specific Registers – Bus Number CAM	240
13-88. 2C8h Bus Number CAM 0	241
13-89. 2CCh Bus Number CAM 1	241
13-90. 2D0h Bus Number CAM 2	241
13-91. 2D4h Bus Number CAM 3	242
13-92. 2D8h Bus Number CAM 4	242
13-93. 2DCh Bus Number CAM 5	242
13-94. 2E0h Bus Number CAM 6	243
13-95. 2E4h Bus Number CAM 7	243
Device-Specific Registers – I/O CAM	244
13-96. 308h I/O CAM 0	245
13-97. 30Ah I/O CAM 1	245
13-98. 30Ch I/O CAM 2	245
13-99. 30Eh I/O CAM 3	246
13-100. 310h I/O CAM 4	246
13-101. 312h I/O CAM 5	246
13-102. 314h I/O CAM 6	247
13-103. 316h I/O CAM 7	247
Device-Specific Registers – Address-Mapping CAM	248
13-104. 348h AMCAM 0 Memory Base and Limit	250
13-105. 34Ch AMCAM 0 Prefetchable Memory Base and Limit	250
13-106. 350h AMCAM 0 Prefetchable Memory Upper Base Address	250
13-107. 354h AMCAM 0 Prefetchable Memory Upper Limit Address	250
13-108. 358h AMCAM 1 Memory Base and Limit	251
13-109. 35Ch AMCAM 1 Prefetchable Memory Base and Limit	251
13-110. 360h AMCAM 1 Prefetchable Memory Upper Base Address	251
13-111. 364h AMCAM 1 Prefetchable Memory Upper Limit Address	251
13-112. 368h AMCAM 2 Memory Base and Limit	252
13-113. 36Ch AMCAM 2 Prefetchable Memory Base and Limit	252
13-114. 370h AMCAM 2 Prefetchable Memory Upper Base Address	252
13-115. 374h AMCAM 2 Prefetchable Memory Upper Limit Address	252
13-116. 378h AMCAM 3 Memory Base and Limit	253
13-117. 37Ch AMCAM 3 Prefetchable Memory Base and Limit	253
13-118. 380h AMCAM 3 Prefetchable Memory Upper Base Address	253
13-119. 384h AMCAM 3 Prefetchable Memory Upper Limit Address	253
13-120. 388h AMCAM 4 Memory Base and Limit	254
13-121. 38Ch AMCAM 4 Prefetchable Memory Base and Limit	254
13-122. 390h AMCAM 4 Prefetchable Memory Upper Base Address	254
13-123. 394h AMCAM 4 Prefetchable Memory Upper Limit Address	254
13-124. 398h AMCAM 5 Memory Base and Limit	255
13-125. 39Ch AMCAM 5 Prefetchable Memory Base and Limit	255
13-126. 3A0h AMCAM 5 Prefetchable Memory Upper Base Address	255
13-127. 3A4h AMCAM 5 Prefetchable Memory Upper Limit Address	255
13-128. 3A8h AMCAM 6 Memory Base and Limit	256
13-129. 3ACh AMCAM 6 Prefetchable Memory Base and Limit	256
13-130. 3B0h AMCAM 6 Prefetchable Memory Upper Base Address	256
13-131. 3B4h AMCAM 6 Prefetchable Memory Upper Limit Address	256
13-132. 3B8h AMCAM 7 Memory Base and Limit	257
13-133. 3BCh AMCAM 7 Prefetchable Memory Base and Limit	257
13-134. 3C0h AMCAM 7 Prefetchable Memory Upper Base Address	257
13-135. 3C4h AMCAM 7 Prefetchable Memory Upper Limit Address	257
Device-Specific Registers – Ingress Control and Port Enable	258
13-136. 660h Ingress Control	258
13-137. 664h Ingress Control Shadow	259
13-138. 668h Ingress Port Enable	259
13-139. 66Ch Negotiated Link Width for Ports 0, 1, 2, 3, 4, 5, 6, 7	260

Device-Specific Registers – I/O CAM Base and Limit Upper 16 Bits	262
13-140. 680h I/OCAM Upper Port 0	262
13-141. 684h I/OCAM Upper Port 1	262
13-142. 688h I/OCAM Upper Port 2	262
13-143. 68Ch I/OCAM Upper Port 3	263
13-144. 690h I/OCAM Upper Port 4	263
13-145. 694h I/OCAM Upper Port 5	263
13-146. 698h I/OCAM Upper Port 6	263
13-147. 69Ch I/OCAM Upper Port 7	263
Device-Specific Registers – Base Address Shadow	264
13-148. 6C0h BAR0 Shadow for Port 0	265
13-149. 6C4h BAR1 Shadow for Port 0	265
13-150. 6C8h BAR0 Shadow for Port 1	266
13-151. 6CCh BAR1 Shadow for Port 1	266
13-152. 6D0h BAR0 Shadow for Port 2	267
13-153. 6D4h BAR1 Shadow for Port 2	267
13-154. 6D8h BAR0 Shadow for Port 3	268
13-155. 6DCh BAR1 Shadow for Port 3	268
13-156. 6E0h BAR0 Shadow for Port 4	269
13-157. 6E4h BAR1 Shadow for Port 4	269
13-158. 6E8h BAR0 Shadow for Port 5	270
13-159. 6ECh BAR1 Shadow for Port 5	270
13-160. 6F0h BAR0 Shadow for Port 6	271
13-161. 6F4h BAR1 Shadow for Port 6	271
13-162. 6F8h BAR0 Shadow for Port 7	272
13-163. 6FCh BAR1 Shadow for Port 7	272
Device-Specific Registers – Shadow Virtual Channel Capability	273
13-164. 740h VC0 Port 0 Capability	274
13-165. 748h VC0 Port 1 Capability	274
13-166. 750h VC0 Port 2 Capability	274
13-167. 758h VC0 Port 3 Capability	275
13-168. 760h VC0 Port 4 Capability	275
13-169. 768h VC0 Port 5 Capability	275
13-170. 770h VC0 Port 6 Capability	276
13-171. 778h VC0 Port 7 Capability	276
Device-Specific Registers – Ingress Credit Handler	277
13-172. 940h INCH Port Pool Setting for Ports 0, 1, 2, 3	279
13-173. 944h INCH Port Pool Setting for Ports 4, 5, 6, 7	281
13-174. A00h INCH Threshold Port 0 VC0 Posted	283
13-175. A04h INCH Threshold Port 0 VC0 Non-Posted	284
13-176. A08h INCH Threshold Port 0 VC0 Completion	285
13-177. A18h INCH Threshold Port 1 VC0 Posted	286
13-178. A1Ch INCH Threshold Port 1 VC0 Non-Posted	287
13-179. A20h INCH Threshold Port 1 VC0 Completion	288
13-180. A30h INCH Threshold Port 2 VC0 Posted	289
13-181. A34h INCH Threshold Port 2 VC0 Non-Posted	290
13-182. A38h INCH Threshold Port 2 VC0 Completion	291
13-183. A48h INCH Threshold, Port 3 VC0 Posted	292
13-184. A4Ch INCH Threshold, Port 3 VC0 Non-Posted	293
13-185. A50h INCH Threshold, Port 3 VC0 Completion	294
13-186. A60h INCH Threshold, Port 4 VC0 Posted	295
13-187. A64h INCH Threshold, Port 4 VC0 Non-Posted	296
13-188. A68h INCH Threshold, Port 4 VC0 Completion	297
13-189. A78h INCH Threshold, Port 5 VC0 Posted	298
13-190. A7Ch INCH Threshold, Port 5 VC0 Non-Posted	299

13-191.	A80h	INCH Threshold, Port 5 VC0 Completion	300
13-192.	A90h	INCH Threshold, Port 6 VC0 Posted	301
13-193.	A94h	INCH Threshold, Port 6 VC0 Non-Posted	302
13-194.	A98h	INCH Threshold, Port 6 VC0 Completion	303
13-195.	AA8h	INCH Threshold, Port 7 VC0 Posted	304
13-196.	AACH	INCH Threshold, Port 7 VC0 Non-Posted	305
13-197.	AB0h	INCH Threshold, Port 7 VC0 Completion	306
Device-Specific Registers – Port Configuration Header			307
13-198.	E00h	Command Port 0	307
13-199.	E02h	Command Port 1	308
13-200.	E04h	Command Port 2	309
13-201.	E06h	Command Port 3	310
13-202.	E08h	Command Port 4	311
13-203.	E0Ah	Command Port 5	312
13-204.	E0Ch	Command Port 6	313
13-205.	E0Eh	Command Port 7	314
13-206.	E20h	Bridge Control Port 0	315
13-207.	E22h	Bridge Control Port 1	317
13-208.	E24h	Bridge Control Port 2	319
13-209.	E26h	Bridge Control Port 3	321
13-210.	E28h	Bridge Control Port 4	323
13-211.	E2Ah	Bridge Control Port 5	325
13-212.	E2Ch	Bridge Control Port 6	327
13-213.	E2Eh	Bridge Control Port 7	329
Device-Specific Registers – Source Queue Weight and Soft Error			331
13-214.	F10h	Port Egress TLP Threshold	331
13-215.	F14h	Source Queue Weight	332
13-216.	F18h	Soft Error Counters 1	333
13-217.	F1Ch	Soft Error Counters 2	333
13-218.	F2Ch	Soft Error Counters 6	333
13-219.	F30h	Soft Error Injection	334
Advanced Error Reporting Extended Capability Registers			335
13-220.	FB4h	Advanced Error Reporting Enhanced Capability Header	335
13-221.	FB8h	Uncorrectable Error Status	336
13-222.	FBCCh	Uncorrectable Error Mask	337
13-223.	FC0h	Uncorrectable Error Severity	338
13-224.	FC4h	Correctable Error Status	339
13-225.	FC8h	Correctable Error Mask	340
13-226.	FCCh	Advanced Error Capabilities and Control	340
13-227.	FD0h	Header Log 0	341
13-228.	FD4h	Header Log 1	341
13-229.	FD8h	Header Log 2	341
13-230.	FDCh	Header Log 3	341

THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 1 Introduction

1.1 Features

PLX Technology's ExpressLane™ PEX 8509 PCI Express Switch supports the following features:

- 8-port PCI Express switch
 - Eight lanes with integrated on-chip SerDes
 - Low-power SerDes (under 90 mW per lane)
 - Fully Non-Blocking Switch architecture
 - Optional Device-Specific Relaxed Ordering
 - Port configuration
 - Eight independent ports
 - Choice of width (number of lanes) per unique link/port – x1, x2, and x4
 - Configurable with serial EEPROM or I²C
 - Designate any port as the *upstream port* (Port 0 is recommended)
 - Maximum Payload Size – 1,024 bytes
 - Dynamic Buffer Pool architecture, for faster credit updates
- Quality of Service (QoS) support
 - All ports support one, full-featured Virtual Channel (VC0)
 - All ports support eight Traffic Class (TC) mapping, independently of the other ports
 - Ingress port arbitration
- Reliability, Availability, Serviceability (RAS) features
 - PCI Express Standard Hot Plug Controller for three ports, including optional usage models for Manually operated Retention Latch, by way of MRL Sensor and Attention Button support
 - Baseline and Advanced Error Reporting capability
 - JTAG boundary scan
- INTA# ([PEX_INTA#](#)) and FATAL ERROR ([FATAL_ERR#](#)) (Conventional PCI SERR# equivalent) ball support
- Lane Status balls ([PEX_LANE_GOOD\[7:0\]#](#))
- Other PCI Express Capabilities
 - Transaction Layer Packet (TLP) Digest support
 - Poison bit
 - End-to-end Cyclic Redundancy Check (ECRC)
 - Lane reversal support
 - Polarity reversal
 - Conventional PCI-compatible Link Power Management states – L0, L0s, L1, L2/L3 Ready, and L3 (with Vaux *not supported*)
 - Conventional PCI-compatible Device Power Management states – D0 and D3hot
 - Active State Power Management (ASPM) fully supported
- Out-of-Band Initialization options
 - Serial EEPROM
 - I²C (7-bit Slave address with 100 Kbps)


- Performance
 - 40 Gbps aggregate bandwidth [2.5 Gbps/lane x 8 SerDes lanes x 2 (full duplex)]
 - Cut-Thru packet latency of 118 ns for a x4 to x1 configuration
 - Non-blocking internal crossbar supporting full wire speed
- Testability – JTAG support for DC
- 15 x 15 mm², 196-ball, Plastic Ball Grid Array (PBGA) package
- Typical power – 1.2W
- 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 ExpressCard CEM r1.1)*
 - *IEEE Standard 1149.1-1990, IEEE Standard Test Access Port and Boundary-Scan Architecture, 1990 (IEEE Standard 1149.1-1990)*
 - *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 (IEEE Standard 1149.6-2003)*
 - *The I²C-Bus Specification, Version 2.1 (I²C Bus v2.1)*

1.2 Overview

This data book describes PLX Technology’s ExpressLane PEX 8509, a fully non-blocking, low-latency, low-cost, and low-power 8-lane, 8-port PCI Express Gen1 switch. Conforming to the *PCI Express Base r1.1*, the PEX 8509 enables users to add scalable, high-bandwidth I/O to a wide variety of products, including servers, communication products, storage systems, and other embedded products. The PEX 8509’s flexible hardware configuration and software programmability allows the switch’s port configurations and QoS operating characteristics to be tailored to suit a wide variety of application requirements.

The PEX 8509 is principally aimed at fan-in/out or aggregation applications and control planes; however, it can also be used in peer-to-peer communication traffic. The PEX 8509 supports the port configurations illustrated in [Figure 1-1](#). The PEX 8509 can support x1 and x2 ports, by auto-negotiating its ports to the link width of the end-device to which it is interfacing.

Figure 1-1. PEX 8509 Port Configuration



THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 2 Features and Applications

2.1 Flexible and Feature-Rich 8-Lane/8-Port Switch

2.1.1 Highly Flexible Port Configurations

The PLX ExpressLane PEX 8509 PCI Express Switch offers flexibility in the configuration of its ports. A maximum of eight ports can be configured to the standard widths of x1, x2, and/or x4, to support specific bandwidth needs. The ports can be symmetric (each port has the same lane width) or asymmetric (ports have different lane widths). Any one port can be designated as the upstream port.

2.1.2 High Performance

The PEX 8509 architecture supports packet Cut-Thru with a latency of 118 ns (x4 to x1). This, combined with large Packet memory (256 to 1,024-byte [Maximum Payload Size](#)), and Non-Blocking Internal Switch architecture, provide full line rate on its ports.

2.1.3 End-to-End Packet Integrity

The PEX 8509 provides **End-to-end CRC** (ECRC) protection and **Poison Bit** support to enable designs that require **guaranteed error-free packets**. PLX also supports data path parity and memory (RAM) error correction as packets pass through the PEX 8509. These features are optional in the *PCI Express Base r1.1*; however, PLX provides them across its entire ExpressLane switch product line.

2.1.4 Configuration Flexibility

The PEX 8509 provides several ways to configure its operations. *For example*, the PEX 8509 can be configured through Strapping balls, the I²C interface, CPU Configuration requests, or an optional serial EEPROM. Additionally, the I²C interface allows for easy debug during the development phase, performance monitoring during the operation phase, and driver or software upgrade.

2.1.5 Interoperability

The PEX 8509 is designed to be fully compliant with the *PCI Express Base r1.1*. Additionally, the device supports **auto-negotiation**, **lane reversal**, and **polarity reversal**, for maximum board design and board placement flexibility. Furthermore, the PEX 8509 is interoperable with many popular motherboards and server boards with PCI Express connections, 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** and compliance testing at the **PCI-SIG Plug-Fest**, to ensure compatibility with PCI Express devices in the market.

2.1.6 Low Power with Granular SerDes Control

The PEX 8509 provides **low-power** capability that is fully compliant with the *PCI Express Base r1.1* Power Management specifications. Unused SerDes can be disabled to further reduce PEX 8509 power consumption.

The PEX 8509 supports **SerDes output software control**, to allow power and signal strength optimization within a system. The PLX SerDes implementation supports four power levels – *Off*, *Low*, *Typical*, and *High*. The SerDes block also supports **Loopback modes** and **Advanced Error Reporting**, which enables efficient system debug and management.

2.1.7 Flexible Port-Width Configuration

The width of each port can be individually configured to any valid width (x1, x2, or x4) depending on the application bandwidth requirements. Any of the configuration methods described in [Section 2.1.4](#) can be used for port configuration.

2.1.8 Hot Plug for High Availability

Hot Plug capability allows users to replace hardware modules and perform maintenance, without having to power down the system.

The PEX 8509 Hot Plug Capability and Advanced Error Reporting features make the switch suitable for **High-Availability (HA) applications**. Three of the eight ports include a Standard Hot Plug Controller. If the PEX 8509 is used in an application where one or more of its downstream ports connect to PCI Express slots, the ports with the Hot Plug Controller feature can be used for these slots. The PEX 8509 automatically manages Hot Plug events on its Hot Plug-capable downstream ports/slots. Additionally, the upstream port is a fully compliant Hot Plug client, and the PEX 8509 can be used on hot-pluggable adapter boards, docking stations, and line cards.

2.1.9 Dynamic Lane Reversal

The PEX 8509 supports dynamic Lane Reversal during the link training process. Lane Reversal capability allows flexibility in determining board routing, so that PCI Express components can be connected without having to crisscross wires. If the wiring of lanes to a device is reversed (on both Transmitters and Receivers), only one of the two connected devices must support Lane Reversal.

Either of the outside lanes (Transmitter and Receiver pairs) of the PEX 8509 programmed port width must be identified as being Lane 0. During link training, both devices on the link negotiate the lane numbering. During the LTSSM's Configuration state, the upstream device sends TS1 Ordered-Sets, in which each connected lane is identified by a consecutive Lane Number, starting with Lane 0 corresponding to the physical Lane Number of the port. If the PEX 8509 port receives a TS1 with a non-zero Lane Number on its Lane 0, the port reverses its lane numbers and again attempts to train. To confirm successful Lane Number negotiation, both devices exchange TS2 Ordered-Sets with identical Lane Numbers on each connected lane.

2.1.10 Fully Compliant Power Management

The PEX 8509 supports Link (L0, L0s, L1, L2/L3 Ready, and L3) and Device (D0 and D3hot) Power Management (PM) states, in compliance with the *PCI Express Base r1.1* Power Management specifications.

2.2 Applications

Suitable for aggregation as well as control plane and docking station applications, the PEX 8509 can be configured for a wide variety of form factors and applications.


2.2.1 Adapter Board Aggregation

The PEX 8509, with its versatile symmetric or asymmetric lane configuration capability, allows user-specific tuning to a variety of aggregation applications.

Figure 2-1 illustrates an example for an aggregation design, where a wide port in the PEX 8509 is used to connect to the upstream path in the host direction and multiple narrow downstream ports connect to PCI Express I/O devices.

In this example, the PEX 8509 would typically have a 4-lane upstream port, and as many as four downstream ports. The downstream ports consist of a single lane and connect to four GE MACs.

Figure 2-1. Adapter Board Aggregation



2.2.2 Control Plane Usage

Low latency and high port count make the PEX 8509 suitable for control plane applications. Figure 2-2 represents a controller card for a communications system. The PEX 8509 can be used to provide the control fabric for up to seven line cards.

The PEX 8509 can be cascaded to provide a higher port count for control applications that require connectivity to more than seven line cards, as illustrated in Figure 2-3. In this example, cascading two PEX 8509 switches provides connectivity to 13 line cards. The low latency in the PEX 8509 makes it an ideal candidate for control applications.

Figure 2-2. Control Plane


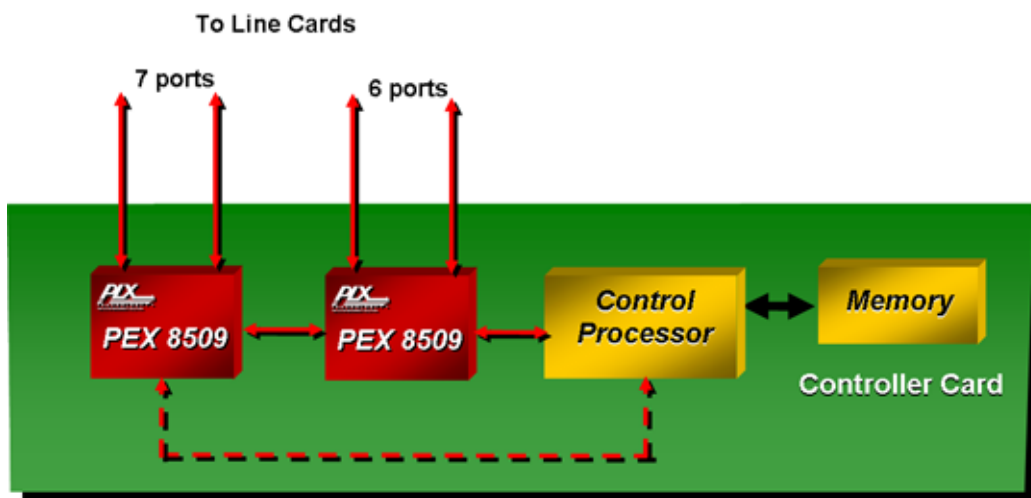


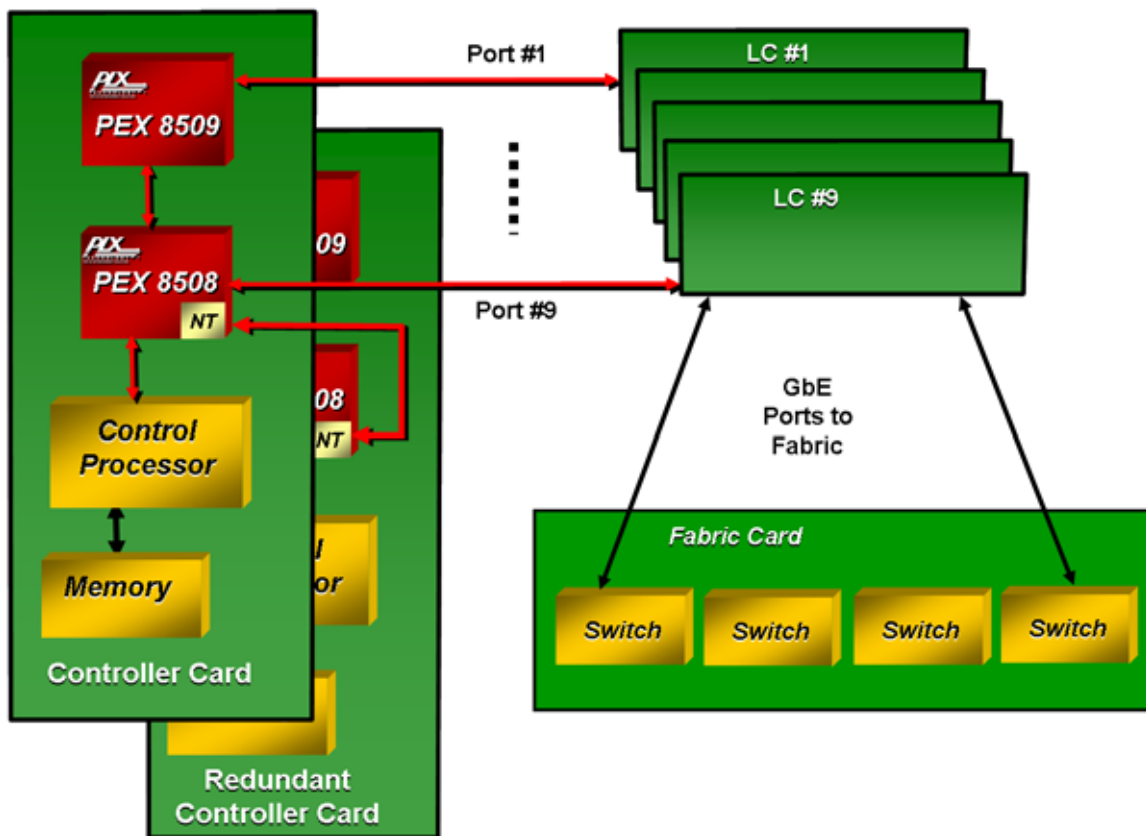
Figure 2-3. Cascaded PEX 8509



Redundancy in the control plane is something that is desirable and required in certain communications platforms. The PEX 8509 can be used in these redundant control planes to provide line card connectivity. Connectivity between the two controller cards is accomplished by using a Non-Transparent port. The Non-Transparent port guarantees that the controller cards' Memory and Address space remains isolated (not shared) between the cards. Communication between the controller cards is achieved by way of Address Translation mechanisms, as defined by the Non-Transparent port. Figure 2-3 illustrates a redundant controller card example, using PLX Technology's ExpressLane PEX 8508 PCI Express Switch to provide the Non-Transparent port between the two controller cards while the PEX 8509 provides the line card connectivity.

Note: For further details regarding Non-Transparency and the PEX 8508, refer to the [PEX 8508 Data Book](#).

Figure 2-4. Redundant Controller Card



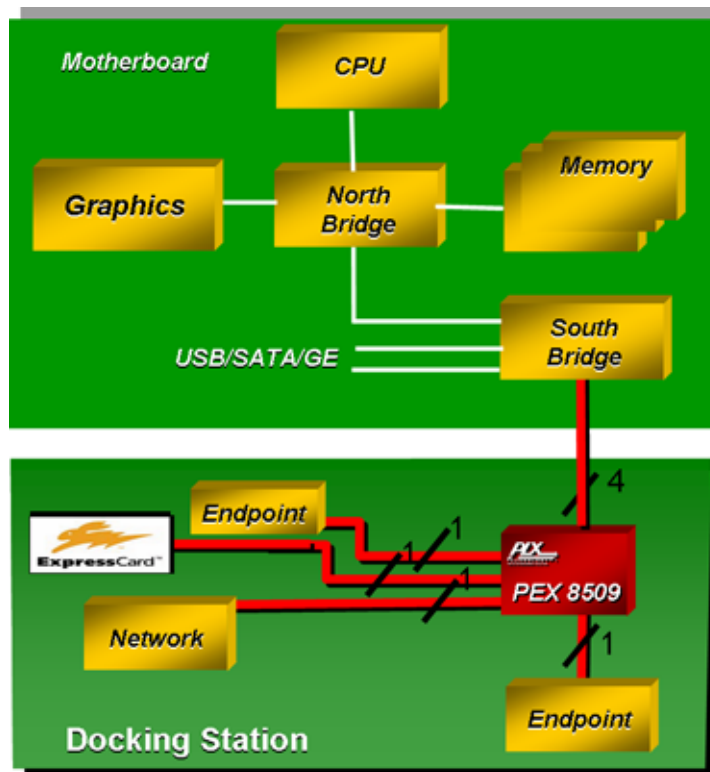
2.2.3 Docking Station

Docking station applications can take advantage of the PEX 8509’s fan-out capabilities. [Figure 2-5](#) illustrates use of the PEX 8509 in a docking station application. The upstream x4 port link connects to the Host (laptop computer) and the other downstream ports provide fan-out capability to the I/O devices located in the docking station.

For docking stations that are not bandwidth intensive, but require higher downstream connectivity, they can be re-configured to reduce the upstream port width, to provide additional downstream ports. *For example*, the PEX 8509 can be configured with an upstream port width of x2. The remaining six downstream ports provide fan-out capability to the I/O devices located in the docking station.

Three of the PEX 8509’s eight ports include a Standard Hot Plug Controller. When the PEX 8509 is used in a docking station application, one or more of its downstream ports can connect to PCI Express device(s) that are expected to be extracted and/or inserted to a live system. The ports in the PEX 8509 with the Hot Plug Controller feature can be used for the devices that require Hot Plug support. The PEX 8509 automatically manages the Hot Plug events of its Hot Plug-capable ports.

Figure 2-5. PEX 8509 PCI Express Switch in a Docking Station



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. The recommended upstream port is Port 0; however, any port can be configured as the upstream port through optional configuration, by way of a serial EEPROM, the I²C 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 8509 are compliant to the PCI and PCI Express system models. The Configuration Space registers (CSRs) in a virtual primary/secondary PCI-to-PCI bridge are accessible by Type 0 and Type 1 Configuration requests, through the virtual primary bus interface (matching Bus Number, Device Number, and Function Number).

2.3.2 Interrupt Sources and Events

The PEX 8509 supports the INTx Interrupt message type (compatible with *PCI r3.0* Interrupt signals) or Message Signaled Interrupts (MSI), when enabled. The PEX 8509 generates interrupts/messages for Hot Plug or Link State events, Device-Specific errors, and Baseline and Advanced Error Reporting.

THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 3 Signal Ball Description

3.1 Introduction

This chapter provides descriptions of the PEX 8509 signal balls. The signal name, type, location, and a brief description are provided for each signal ball. A list of signals by location and a map of the PEX 8509's physical layout are also provided.

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	Power (VDD10A) balls for SerDes Analog circuits
CMLCLK _n ^a	Differential low-voltage, high-speed, CML negative Clock inputs
CMLCLK _p ^a	Differential low-voltage, high-speed, CML positive Clock inputs
CMLR _n	Differential low-voltage, high-speed, CML negative Receiver inputs
CMLR _p	Differential low-voltage, high-speed, CML positive Receiver inputs
CMLT _n	Differential low-voltage, high-speed, CML negative Transmitter outputs
CMLT _p	Differential low-voltage, high-speed, CML positive Transmitter outputs
CPWR	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 internal pull-up resistors)
I/O	Bidirectional (Input or Output) signal (signals without internal pull-up resistors)
I/OPWR	3.3V Power (VDD33) balls for Input and Output interfaces
O	Output
OD	Open Drain output
PLL_GND	PLL Ground connection
PLL_PWR	3.3V Power (VDD33A) balls for PLL circuits
PU	Weak internal pull-up resistor
SerDes	Differential low-voltage, high-speed, I/O signal pairs (negative and positive)
SPWR	Power (VDD10S) balls for SerDes Digital circuits
STRAP	Input Strapping balls, cannot be left floating on the board

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

3.3 Internal Pull-Up Resistors

The PEX 8509 contains signals that have weak internal pull-up resistors, indicated in this chapter by PU, in the signal ball tables (**Type** column). If a signal with this notation is 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 Ω to 10K Ω resistor.

3.4 Signal Ball Descriptions

The signals are divided into the following groups:

- PCI Express Signals
- Hot Plug Signals
- Serial EEPROM Signals
- Strapping Signals
- JTAG Interface Signals
- I²C Slave Interface Signals
- Device-Specific Signals
- No Connect Signals
- Power and Ground Signals

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

3.4.1 PCI Express Signals

Table 3-2 defines the PCI Express SerDes and Control signals.

Table 3-2. PCI Express Signals – 35 Balls

Signal Name	Type	Location	Description
PEX_PERn[7:0]	CMLRn	K14, P12, N11, P8, P7, N4, P3, L1	Negative Half of PCI Express Receiver Differential Signal Pairs (8 Balls)
PEX_PERp[7:0]	CMLRp	L14, N12, P11, N8, N7, P4, N3, M1	Positive Half of PCI Express Receiver Differential Signal Pairs (8 Balls)
PEX_PERST#	I PU	G3	PCI Express Reset Used to cause a Fundamental Reset. (Refer to Chapter 5, “Reset and Initialization,” for further details.)
PEX_PETn[7:0]	CMLTn	P14, N13, P10, N9, N6, P5, N2, P1	Negative Half of PCI Express Transmitter Differential Signal Pairs (8 Balls)
PEX_PETp[7:0]	CMLTp	N14, P13, N10, P9, P6, N5, P2, N1	Positive Half of PCI Express Transmitter Differential Signal Pairs (8 Balls)
PEX_REFCLKn	CMLCLKn	J2	Negative Half of 100-MHz PCI Express Reference Clock Input Signal Pair
PEX_REFCLKp	CMLCLKp	J1	Positive Half of 100-MHz PCI Express Reference Clock Input Signal Pair

3.4.2 Hot Plug Signals

The PEX 8509 includes nine Hot Plug signals for each Hot Plug-capable downstream PCI Express port – Ports 1, 2, and 3 (3 ports x 9 signals/port = 27 total signals) – defined in Table 3-3. These signals are active only for Hot Plug-capable downstream ports configured at start-up. (Refer to Chapter 11, “Hot Plug Support,” for further details.)

Table 3-3. Hot Plug Signals – 27 Balls

Signal Name	Type	Location	Description
HP_ATNLED[3:1]#	O PU ^a	D11, C9, B1	<p>Hot Plug Attention LED Outputs for Ports 3, 2, and 1 (3 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 Capability register <i>Attention Indicator Present</i> bit is set (offset 7Ch[3]=1) and controlled by the Slot Control register <i>Attention Indicator Controls</i> field (offset 80h[7:6]). When software writes to the <i>Attention Indicator Controls</i> 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_ATNLEDx#.</p> <p>When the following conditions exist:</p> <ul style="list-style-type: none"> • Slot Capability register <i>Attention Indicator Present</i> bit is set (offset 7Ch[3]=1), and • Slot Control register <i>Command Completed Interrupt Enable</i> bit is not masked (offset 80h[4]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INTx message, or PEX_INTA# output, all mutually exclusive) can be generated to the Host.</p> <p>An external current-limiting resistor is required.</p>
HP_BUTTON[3:1]#	I PU ^a	B13, A8, B2	<p>Hot Plug Attention Button Inputs for Ports 3, 2, and 1 (3 Balls) Active-Low Slot Control Logic input, directly connected to the Attention Button, with input assertion status latched in the Slot Status register <i>Attention Button Pressed</i> bit (offset 80h[16]).</p> <p>Enabled when the Slot Capability register <i>Attention Button Present</i> bit is set (offset 7Ch[0]=1).</p> <p>When the following conditions exist:</p> <ul style="list-style-type: none"> • HP_BUTTONx# is not masked (Slot Control register <i>Attention Button Pressed Enable</i> bit, offset 80h[0]=1), and • Slot Capability register <i>Hot Plug Capable</i> bit is set (offset 7Ch[6]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INTx message, or PEX_INTA# output, all mutually exclusive) can be generated, to notify the Host of intended board insertion or removal.</p> <p><i>Note:</i> HP_BUTTONx# is internally de-bounced, but must remain stable for at least 10 ms.</p>

Table 3-3. Hot Plug Signals – 27 Balls (Cont.)

Signal Name	Type	Location	Description
HP_CLKEN[3:1]#	O PU ^a	F12, C4, F2	<p>Hot Plug Reference Clock Enable Outputs for Ports 3, 2, and 1 (3 Balls)</p> <p>Active-Low output that, when enabled, allows external REFCLK to be provided to the slot.</p> <p>Enabled when the Slot Capability register <i>Power Controller Present</i> bit is set (offset 7Ch[1]=1), and controlled by the Slot Control register <i>Power Controller Control</i> bit (offset 80h[10]).</p> <p>The time delay from HP_PWRENx# output assertion to HP_CLKENx# output assertion is programmable (through serial EEPROM load) from 16 ms (default) to 128 ms, in the <i>HPC T_{pepv} Delay</i> field (offset 1E0h[4:3]).</p>
HP_MRL[3:1]#	I PU ^a	B14, B8, C1	<p>Hot Plug Manually Operated Retention Latch Sensor Inputs for Ports 3, 2, and 1 (3 Balls)</p> <p>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_MRLx# input assertion enables Hot Plug output sequencing to turn On the slot's power (HP_PWRENx# and HP_PWRLEDx#) and clock (HP_CLKENx#), and de-assert Reset (HP_PERSTx#) after reset or under software control.</p> <p>Enabled when the Slot Capability register <i>MRL Sensor Present</i> bit is set (offset 7Ch[2]=1).</p> <p>A change in the HP_MRLx# Input signal state is latched in the Slot Status register <i>MRL Sensor Changed</i> bit (offset 80h[18]), and the state change can assert an interrupt to notify the Host of a change in the MRL Sensor state.</p> <p>When the following conditions exist:</p> <ul style="list-style-type: none"> • HP_MRLx# is not masked (Slot Control register <i>MRL Sensor Changed Enable</i> bit, offset 80h[2]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INTx message, or PEX_INTA# output, all mutually exclusive) can be generated.</p> <p>If the associated Hot Plug-capable downstream port connects to a PCI Express board slot that does not implement an MRL Sensor, HP_MRLx# is normally connected to HP_PRSNTx# and a pull-up resistor, with the common node connected to the PRSNT2# signal(s) at the slot. If the associated Hot Plug-capable downstream port instead connects directly to a device (in which case Hot Plug is not used), pull HP_MRLx# Low.</p> <p><i>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.</i></p>
HP_PERST[3:1]#	O	E12, C5, E3	<p>Hot Plug Reset Outputs for Ports 3, 2, and 1 (3 Balls)</p> <p>Active-Low Hot Plug output used to reset the slot. When the Slot Capability register <i>Power Controller Present</i> bit is set (offset 7Ch[1]=1), the HP_PERSTx# output state can be controlled by software, using the Slot Control register <i>Power Controller Control</i> bit (offset 80h[10]).</p>

Table 3-3. Hot Plug Signals – 27 Balls (Cont.)

Signal Name	Type	Location	Description
HP_PRSNT[3:1]#	I PU ^a	F11, B5, E2	<p>Hot Plug PRSNT2# Inputs for Ports 3, 2, and 1 (3 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_PRSNTx# Input signal state is latched in the Slot Status register <i>Presence Detect Changed</i> bit (offset 80h[19]), and the state change can assert an interrupt to notify the Host of board presence or absence.</p> <p>When the following conditions exist:</p> <ul style="list-style-type: none"> • HP_PRSNTx# is not masked (Slot Control register <i>Presence Detect Changed Enable</i> bit (offset 80h[3]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INTx message, or PEX_INTA# output, all mutually exclusive) can be generated.</p> <p><i>Note:</i> HP_PRSNTx# is internally de-bounced, but must remain stable for at least 10 ms.</p>
HP_PWREN[3:1]#	O	C14, C7, D1	<p>Hot Plug Power Enable Outputs for Ports 3, 2, and 1 (3 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 Capability register <i>Power Controller Present</i> bit is set (offset 7Ch[1]=1).</p> <p>When software turns the slot's Power Controller On or Off [Slot Control register <i>Power Controller Control</i> bit (offset 80h[10])], a Command Completed interrupt can be generated to notify the Host that the command has been executed.</p> <p>When the following conditions exist:</p> <ul style="list-style-type: none"> • Slot Control register <i>Command Completed Interrupt Enable</i> bit is not masked (offset 80h[4]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INTx message, or PEX_INTA# output, all mutually exclusive) can be generated to the Host.</p> <p>When HP_MRLx# is enabled [Slot Capability register <i>MRL Sensor Present</i> bit is set (offset 7Ch[2]=1)], HP_MRLx# input assertion enables Hot Plug output sequencing to turn On the slot's power, by asserting HP_PWRENx# after reset or under software control.</p>

Table 3-3. Hot Plug Signals – 27 Balls (Cont.)

Signal Name	Type	Location	Description
HP_PWRFLT[3:1]#	I PU ^a	D12, B7, D2	<p>Hot Plug Power Fault Inputs for Ports 3, 2, and 1 (3 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 Capability register <i>Power Controller Present</i> bit is set (offset 7Ch[1]=1), and input assertion status is latched in the Slot Status register <i>Power Fault Detected</i> bit (offset 80h[17]). When the following conditions exist:</p> <ul style="list-style-type: none"> • HP_PWRFLT_x# is not masked (Slot Control register <i>Power Fault Detector Enable</i> bit, offset 80h[1]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INT_x message, or PEX_INTA# output, all mutually exclusive) can be generated, to notify the Host of a power fault.</p> <p><i>Note: If HP_PWREN_x# and HP_CLKEN_x# are not used, HP_PWRFLT_x# can be used as a general-purpose input with status reflected in the Slot Status register Power Fault Detected bit (offset 80h[17]), provided the Slot Capability register Power Controller Present bit is set (offset 7Ch[1]=1).</i></p>
HP_PWRLED[3:1]#	O PU ^a	E13, A3, E1	<p>Hot Plug Power LED Outputs for Ports 3, 2, and 1 (3 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 Capability register <i>Power Indicator Present</i> bit is set (offset 7Ch[4]=1), and controlled by the Slot Control register <i>Power Indicator Control</i> field (offset 80h[9:8]). When software writes to the <i>Power Indicator Control</i> 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:</p> <ul style="list-style-type: none"> • Slot Capability register <i>Power Indicator Present</i> bit is set (offset 7Ch[4]=1), and • Slot Control register <i>Command Completed Interrupt Enable</i> bit is not masked (offset 80h[4]=1), and • Slot Control register <i>Hot Plug Interrupt Enable</i> bit is set (offset 80h[5]=1), <p>an interrupt (MSI, INT_x message, or PEX_INTA# output, all mutually exclusive) can be generated to the Host. An external current-limiting resistor is required.</p>

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

3.4.3 Serial EEPROM Signals

The PEX 8509 includes four signals for interfacing to a serial EEPROM, defined in [Table 3-4](#). For information regarding serial EEPROM use, refer to [Chapter 6](#), “Serial EEPROM Controller.”

Table 3-4. Serial EEPROM Signals – 4 Balls

Signal Name	Type	Location	Description
EE_CS#	O	F14	Serial EEPROM Active-Low Chip Select Output
EE_DI	O	H12	PEX 8509 Output to Serial EEPROM Data Input
EE_DO	I PU	G12	PEX 8509 Input from Serial EEPROM Data Output Should be pulled High to VDD33 .
EE_SK	O	G13	Serial EEPROM Clock Output Programmable [by way of the Serial EEPROM Clock Frequency register <i>EepFreq[2:0]</i> field (Port 0, offset 268h[2:0])] to the following: <ul style="list-style-type: none"> • 1 MHz (default) • 1.98 MHz • 5 MHz • 9.62 MHz • 12.5 MHz • 15.6 MHz • 17.86 MHz

3.4.4 Strapping Signals

The PEX 8509 Strapping signals, defined in [Table 3-5](#), set the configuration of upstream port assignment, port width, 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 Capability**, **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 3-5. Strapping Signals – 15 Balls

Signal Name	Type	Location	Description
STRAP_DEBUG_SEL[1:0]#	PU STRAP	A13, A12	Factory Test Only (2 Balls) Must be tied High.
STRAP_FAST_BRINGUP#	PU STRAP	G2	Factory Test Only Must be tied High.
STRAP_PLL_BYPASS#	PU STRAP	F1	Factory Test Only Must be tied High.
STRAP_PORTCFG[1:0]	STRAP	E14, F13	Strapping Signals to Select Number of Lanes in Port Configuration for Ports 0, 1, 2, 3, 4, 5, 6, and 7 (2 Balls) Register/Bits – Port Configuration register <i>Port Configuration</i> field (Port 0, offset 224h[1:0]) LL = x1, x1, x1, x1, x1, x1, x1, x1 LH = x2, x1, x1, x1, x1, x1, x1 HL = x2, x2, x1, x1, x1, x1 HH = x4, x1, x1, x1, x1
STRAP_PROBE_MODE#	PU STRAP	G14	Factory Test Only Must be tied High.
STRAP_SERDES_MODE_ENABLE#	STRAP	A2	Factory Test Only Must be tied High.
STRAP_TESTMODE[3:0]	PU STRAP	C12, A11, B11, C11	Test Mode Selects (4 Balls) Factory Test Only HHHH = Default (Test modes are disabled)
STRAP_UPSTRM_PORTSEL[2:0]	STRAP	C8, A6, A7	Strapping Signals to Select Upstream Port (3 Balls) Register/Bits – Debug Control register <i>Upstream Port ID</i> field (Port 0, offset 1DCh[11:8]) LLL = Port 0 LLH = Port 1 LHL = Port 2 LHH = Port 3 HLL = Port 4 HLH = Port 5 HHL = Port 6 HHH = Port 7

3.4.5 JTAG Interface Signals

The PEX 8509 includes five signals for performing JTAG boundary scan, defined in [Table 3-6](#). The JTAG interface is described in [Chapter 14, “Test and Debug.”](#)

Table 3-6. JTAG Interface Signals – 5 Balls

Signal Name	Type	Location	Description
JTAG_TCK	I PU	C10	JTAG Test Clock Input JTAG Test Access Port (TAP) Controller clock source. Frequency can be from 0 to 10 MHz.
JTAG_TDI	I PU	A9	JTAG Test Data Input Serial input to the JTAG TAP Controller, for test instructions and data.
JTAG_TDO	O	A10	JTAG Test Data Output Serial output from the JTAG TAP Controller test instructions and data.
JTAG_TMS	I PU	B10	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	B9	JTAG Test Reset Active-Low input used to reset the Test Access Port. Tie to ground through a 1.5K Ω resistor, to hold the JTAG TAP Controller in the <i>Test-Logic-Reset</i> 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 <i>Test-Logic-Reset</i> 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 <i>Test-Logic-Reset</i> state by initializing the JTAG TAP Controller's Instruction register to contain the <i>IDCODE</i> instruction, or by holding the JTAG_TMS input High for at least five rising edges of the JTAG_TCK input.

3.4.6 I²C Slave Interface Signals

Table 3-7 defines the I²C Slave Interface signals. For further details, refer to Chapter 7, “I²C Slave Interface Operation.”

Table 3-7. I²C Slave Interface Signals – 5 Balls

Signal Name	Type	Location	Description
I2C_ADDR[2:0]	I PU	L12, K12, K11	I²C Slave Address Bits 2 through 0 (3 Balls) Used to set the PEX 8509 Slave address on the I ² C Bus. If I ² C or PEX_INTA# output is used, I2C_ADDR[2:0] should be strapped to a unique address, to avoid address conflict with any other I ² C devices (on the same I ² C Bus segment) that have the upper four bits of their 7-bit I ² C Slave address also set to 0111b. Must be pulled High to VDD33 or Low to VSS (GND) through external resistors.
I2C_SCL	I/O OD	H14	I²C Serial Clock Line I ² C Clock source.
I2C_SDA	I/O OD	H13	I²C Serial Data Output Transmits and receives I ² C data.

3.4.7 Device-Specific Signals

Table 3-8 defines the Device-Specific signals – signals that are unique to the PEX 8509.

Table 3-8. Device-Specific Signals – 10 Balls

Signal Name	Type	Location	Description
FATAL_ERR#	O	J13	Fatal Error Asserted Low when a Fatal error is detected in the PEX 8509 and the following conditions exist (all the same conditions that are required to send a Fatal Error message to the Host): <ul style="list-style-type: none"> • Specific error is defined as Fatal in the Uncorrectable Error Severity register (offset FC0h), and • Corresponding Uncorrectable Error Mask register bit (offset FBCh) is not set, and • Device Control register <i>Fatal Error Reporting Enable</i> bit (offset 70h[2]) –or– PCI Command register <i>SERR# Enable</i> bit (offset 04h[8]) is set The Device Control register <i>Fatal Error Detected</i> bit is set (offset 70h[18]=1), and the specific error is flagged in the Uncorrectable Error Status register (offset FB8h).

Table 3-8. Device-Specific Signals – 10 Balls

Signal Name	Type	Location	Description
PEX_INTA#	OD	C3	<p>Interrupt Output PEX_INTA# Interrupt output is enabled if:</p> <ul style="list-style-type: none"> • INTx messages are enabled (PCI Command register <i>Interrupt Disable</i> bit, offset 04h[10]=0), and • MSI is disabled (MSI Control register <i>MSI Enable</i> bit, offset 48h[16]=0) • PEX_INTA# output is enabled (ECC Error Check Disable register <i>Enable PEX_INTA# Ball for Device-Specific Error</i> or <i>Enable PEX_INTA# Ball for Hot Plug or Link State Event</i> bit, Port 0, offset 1C8h[5 and/or 4]=1, respectively) <p><i>Note:</i> PEX_INTA# assertion and INTx message generation are mutually exclusive.</p> <p>PEX_INTA# assertion (Low) indicates that one or more of the following events and/or errors (if not masked) were detected:</p> <p>PCI Express Hot Plug Events (Slot Status register, offset 80h), All Downstream Ports:</p> <ul style="list-style-type: none"> • <i>Presence Detect Changed</i> (bit 19 = 1) • <i>Data Link Layer State Changed</i> (bit 24 = 1) <p><i>Note:</i> Presence is determined by the logical OR of SerDes Receiver Detect or HP_PRSENTx# input (Ports 1, 2, and 3 only) for the port.</p> <p>Hot Plug Events (Slot Status register, offset 80h), Ports 1, 2, and 3 Only:</p> <ul style="list-style-type: none"> • <i>Attention Button Pressed</i> (bit 16 = 1) • <i>Power Fault Detected</i> (bit 17 = 1) • <i>MRL Sensor Changed</i> (bit 18 = 1) • <i>Command Completed</i> (bit 20 = 1) <p>Device-Specific Error Conditions:</p> <ul style="list-style-type: none"> • Completion FIFO overflow (Error Handler 32-Bit Error Status register <i>Completion FIFO Overflow Status</i> bit (offset 1CCh[0]=1) in each port) • Internal RAM 2-bit Error-Correcting Code (ECC) error and 1-bit ECC Error Counter overflow [Device-Specific Error Status for Egress ECC Error register (Port 0, offset 1C0h[31:0]), if not masked in the Device-Specific Error Mask for Egress ECC Error register (Port 0, offset 1C4h[31:0])] • Internal RAM 2-bit ECC Error and 1-bit ECC Error Counter overflow [Error Handler 32-Bit Error Status register (offset 1CCh[26:4]), if not masked in the Error Handler 32-Bit Error Mask register (offset 1D0h[26:4])]
PEX_LANE_GOOD[7:0]#	O	C13, D13, D3, C2, B3, B4, A4, B6	<p>Active-Low PCI Express Lane Linkup Status Indicator Outputs for Lanes [7-0] –or– Programmable GPO (8 Balls)</p> <p>When the Debug Control register <i>On-Board SerDes Lane Status Control</i> bit is cleared (Port 0, offset 1DCh[30]=0), outputs indicate Lane Linkup status. These signals can directly drive common-anode LED modules (external current-limiting resistors are required).</p> <p>When the Debug Control register <i>On-Board SerDes Lane Status Control</i> bit and bit 7 are both set (Port 0, offset 1DCh[30, 7]=11b), these signals are General-Purpose outputs controlled by the Lane Status/Software PEX_LANE_GOODx# LED Control register <i>Port LED Status Control</i> field (Port 0, offset 1F4h[7:0]).</p>

3.4.8 No Connect Signals

Caution: *Do not connect these balls to board electrical paths.
These balls are internally connected to the device.*

Table 3-9. No Connect Signals – 6 Balls

Signal Name	Type	Location	Description
NC_PROCMON	<i>Reserved</i>	F3	No Connect Do not connect this ball to board electrical paths.
NC_SPARE[2:0]	PU <i>Reserved</i>	J12, D14, B12	No Connect (3 Balls) Do not connect these balls to board electrical paths.
THERMAL_DIODEn	<i>Reserved</i>	C6	No Connect Factory Test Only Do not connect this ball to board electrical paths.
THERMAL_DIODEp	<i>Reserved</i>	A5	No Connect Factory Test Only Do not connect this ball to board electrical paths.

3.4.9 Power and Ground Signals

Table 3-10. Power and Ground Signals – 89 Balls

Signal Name	Type	Location	Description
VDD10	CPWR	E5, E7, E9, F10, G5, H10, J5, K4, K6, K8, K10	1.0V Power for Core Logic (11 Balls)
VDD10A	APWR	J3, L5, L10	1.0V Power for SerDes Analog Circuits (3 Balls)
VDD10S	SPWR	K2, L3, L7, L13, M2, M4, M6, M8, M10, M12, M14	1.0V Power for SerDes Digital Circuits (11 Balls)
VDD33	I/OPWR	D4, D5, D6, D7, D8, D9, D10, E4, E11, F4, G4, G11, H4, H11, J11	3.3V Power for I/O Logic Functions (15 Balls)
VDD33A	PLLWR	H3	3.3V Power for PLL Circuits
VSS	GND	A1, A14, E6, E8, E10, F5, F6, F7, F8, F9, G6, G7, G8, G9, G10, H1, H2, H5, H6, H7, H8, H9, J4, J6, J7, J8, J9, J10, J14, K1, K3, K5, K7, K9, K13, L2, L8, M3, M5, M7, M9, M11, M13	Ground Connections (43 Balls)
VSSA_PLL	PLL_GND	G1	PLL Ground Connection
VTT_PEX[3:0]	Supply	L11, L9, L6, L4	SerDes Termination Supply^a(4 Balls) Tied to SerDes termination supply voltage (typically 1.5V).

- a. *PEX_PETn/p[x]* SerDes termination supply voltage controls the 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 Location

Table 3-11. PEX 8509 Ball Assignments by Location

Loc	Signal Name	Type	Signal Group	Comment
A1	VSS	GND	Ground	
A2	STRAP_SERDES_MODE_ENABLE#	STRAP	Strapping	Strapping Ball – tie H
A3	HP_PWRLED2#	O, PU	Hot Plug	
A4	PEX_LANE_GOOD1#	O	Device-Specific	
A5	THERMAL_DIODEp	<i>Reserved</i>	No Connect	
A6	STRAP_UPSTRM_PORTSEL1	STRAP	Strapping	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[2:0]
A7	STRAP_UPSTRM_PORTSEL0	STRAP	Strapping	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[2:0]
A8	HP_BUTTON2#	I, PU	Hot Plug	
A9	JTAG_TDI	I, PU	JTAG	
A10	JTAG_TDO	O	JTAG	
A11	STRAP_TESTMODE2	PU, STRAP	Strapping	Strapping Ball – tie H
A12	STRAP_DEBUG_SEL0#	PU, STRAP	Strapping	Strapping Ball – tie H
A13	STRAP_DEBUG_SEL1#	PU, STRAP	Strapping	Strapping Ball – tie H
A14	VSS	GND	Ground	
B1	HP_ATNLED1#	O, PU	Hot Plug	
B2	HP_BUTTON1#	I, PU	Hot Plug	
B3	PEX_LANE_GOOD3#	O	Device-Specific	
B4	PEX_LANE_GOOD2#	O	Device-Specific	
B5	HP_PRSNT2#	I, PU	Hot Plug	
B6	PEX_LANE_GOOD0#	O	Device-Specific	
B7	HP_PWRFLT2#	I, PU	Hot Plug	
B8	HP_MRL2#	I, PU	Hot Plug	
B9	JTAG_TRST#	I, PU	JTAG	
B10	JTAG_TMS	I, PU	JTAG	
B11	STRAP_TESTMODE1	PU, STRAP	Strapping	Strapping Ball – tie H
B12	NC_SPARE0	PU, <i>Reserved</i>	No Connect	
B13	HP_BUTTON3#	I, PU	Hot Plug	
B14	HP_MRL3#	I, PU	Hot Plug	
C1	HP_MRL1#	I, PU	Hot Plug	
C2	PEX_LANE_GOOD4#	O	Device-Specific	
C3	PEX_INTA#	OD	Device-Specific	
C4	HP_CLKEN2#	O, PU	Hot Plug	
C5	HP_PERST2#	O	Hot Plug	
C6	THERMAL_DIODEn	<i>Reserved</i>	No Connect	
C7	HP_PWREN2#	O	Hot Plug	

Table 3-11. PEX 8509 Ball Assignments by Location (Cont.)

Loc	Signal Name	Type	Signal Group	Comment
C8	STRAP_UPSTRM_PORTSEL2	STRAP	Strapping	Strapping Ball – tie H or L, as defined in STRAP_UPSTRM_PORTSEL[2:0]
C9	HP_ATNLED2#	O, PU	Hot Plug	
C10	JTAG_TCK	I, PU	JTAG	
C11	STRAP_TESTMODE0	PU, STRAP	Strapping	Strapping Ball – tie H
C12	STRAP_TESTMODE3	PU, STRAP	Strapping	Strapping Ball – tie H
C13	PEX_LANE_GOOD7#	O	Device-Specific	
C14	HP_PWREN3#	O	Hot Plug	
D1	HP_PWREN1#	O	Hot Plug	
D2	HP_PWRFLT1#	I, PU	Hot Plug	
D3	PEX_LANE_GOOD5#	O	Device-Specific	
D4	VDD33	I/OPWR	Power	
D5	VDD33	I/OPWR	Power	
D6	VDD33	I/OPWR	Power	
D7	VDD33	I/OPWR	Power	
D8	VDD33	I/OPWR	Power	
D9	VDD33	I/OPWR	Power	
D10	VDD33	I/OPWR	Power	
D11	HP_ATNLED3#	O, PU	Hot Plug	
D12	HP_PWRFLT3#	I, PU	Hot Plug	
D13	PEX_LANE_GOOD6#	O	Device-Specific	
D14	NC_SPARE1	PU, <i>Reserved</i>	No Connect	
E1	HP_PWRLED1#	O, PU	Hot Plug	
E2	HP_PRSNT1#	I, PU	Hot Plug	
E3	HP_PERST1#	O	Hot Plug	
E4	VDD33	I/OPWR	Power	
E5	VDD10	CPWR	Power	
E6	VSS	GND	Ground	
E7	VDD10	CPWR	Power	
E8	VSS	GND	Ground	
E9	VDD10	CPWR	Power	
E10	VSS	GND	Ground	
E11	VDD33	I/OPWR	Power	
E12	HP_PERST3#	O	Hot Plug	
E13	HP_PWRLED3#	O, PU	Hot Plug	
E14	STRAP_PORTCFG1	STRAP	Strapping	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[1:0]
F1	STRAP_PLL_BYPASS#	PU, STRAP	Strapping	Strapping Ball – tie H
F2	HP_CLKEN1#	O, PU	Hot Plug	

Table 3-11. PEX 8509 Ball Assignments by Location (Cont.)

Loc	Signal Name	Type	Signal Group	Comment
F3	NC_PROCMON	<i>Reserved</i>	No Connect	
F4	VDD33	I/OPWR	Power	
F5	VSS	GND	Ground	
F6	VSS	GND	Ground	
F7	VSS	GND	Ground	
F8	VSS	GND	Ground	
F9	VSS	GND	Ground	
F10	VDD10	CPWR	Power	
F11	HP_PRSNT3#	I, PU	Hot Plug	
F12	HP_CLKEN3#	O, PU	Hot Plug	
F13	STRAP_PORTCFG0	STRAP	Strapping	Strapping Ball – tie H or L, as defined in STRAP_PORTCFG[1:0]
F14	EE_CS#	O	Serial EEPROM	
G1	VSSA_PLL	PLL_GND	Ground	
G2	STRAP_FAST_BRINGUP#	PU, STRAP	Strapping	Strapping Ball – tie H
G3	PEX_PERST#	I, PU	PEX Control	
G4	VDD33	I/OPWR	Power	
G5	VDD10	CPWR	Power	
G6	VSS	GND	Ground	
G7	VSS	GND	Ground	
G8	VSS	GND	Ground	
G9	VSS	GND	Ground	
G10	VSS	GND	Ground	
G11	VDD33	I/OPWR	Power	
G12	EE_DO	I, PU	Serial EEPROM	Connected to data output of serial EEPROM
G13	EE_SK	O	Serial EEPROM	
G14	STRAP_PROBE_MODE#	PU, STRAP	Strapping	Strapping Ball – tie H
H1	VSS	GND	Ground	
H2	VSS	GND	Ground	
H3	VDD33A	PLL_PWR	Power	
H4	VDD33	I/OPWR	Power	
H5	VSS	GND	Ground	
H6	VSS	GND	Ground	
H7	VSS	GND	Ground	
H8	VSS	GND	Ground	
H9	VSS	GND	Ground	
H10	VDD10	CPWR	Power	
H11	VDD33	I/OPWR	Power	

Table 3-11. PEX 8509 Ball Assignments by Location (Cont.)

Loc	Signal Name	Type	Signal Group	Comment
H12	EE_DI	O	Serial EEPROM	Connected to data input of serial EEPROM
H13	I2C_SDA	I/O, OD	I ² C	
H14	I2C_SCL	I/O, OD	I ² C	
J1	PEX_REFCLKp	CMLCLKp	SerDes	
J2	PEX_REFCLKn	CMLCLKn	SerDes	
J3	VDD10A	APWR	Power	
J4	VSS	GND	Ground	
J5	VDD10	CPWR	Power	
J6	VSS	GND	Ground	
J7	VSS	GND	Ground	
J8	VSS	GND	Ground	
J9	VSS	GND	Ground	
J10	VSS	GND	Ground	
J11	VDD33	I/OPWR	Power	
J12	NC_SPARE2	PU, <i>Reserved</i>	No Connect	
J13	FATAL_ERR#	O	Device-Specific	
J14	VSS	GND	Ground	
K1	VSS	GND	Ground	
K2	VDD10S	SPWR	Power	
K3	VSS	GND	Ground	
K4	VDD10	CPWR	Power	
K5	VSS	GND	Ground	
K6	VDD10	CPWR	Power	
K7	VSS	GND	Ground	
K8	VDD10	CPWR	Power	
K9	VSS	GND	Ground	
K10	VDD10	CPWR	Power	
K11	I2C_ADDR0	I, PU	I ² C	
K12	I2C_ADDR1	I, PU	I ² C	
K13	VSS	GND	Ground	
K14	PEX_PERn7	CMLRn	SerDes	
L1	PEX_PERn0	CMLRn	SerDes	
L2	VSS	GND	Ground	
L3	VDD10S	SPWR	Power	
L4	VTT_PEX0	Supply	Power	
L5	VDD10A	APWR	Power	
L6	VTT_PEX1	Supply	Power	

Table 3-11. PEX 8509 Ball Assignments by Location (Cont.)

Loc	Signal Name	Type	Signal Group	Comment
L7	VDD10S	SPWR	Power	
L8	VSS	GND	Ground	
L9	VTT_PEX2	Supply	Power	
L10	VDD10A	APWR	Power	
L11	VTT_PEX3	Supply	Power	
L12	I2C_ADDR2	I, PU	I ² C	
L13	VDD10S	SPWR	Power	
L14	PEX_PERp7	CMLRp	SerDes	
M1	PEX_PERp0	CMLRp	SerDes	
M2	VDD10S	SPWR	Power	
M3	VSS	GND	Ground	
M4	VDD10S	SPWR	Power	
M5	VSS	GND	Ground	
M6	VDD10S	SPWR	Power	
M7	VSS	GND	Ground	
M8	VDD10S	SPWR	Power	
M9	VSS	GND	Ground	
M10	VDD10S	SPWR	Power	
M11	VSS	GND	Ground	
M12	VDD10S	SPWR	Power	
M13	VSS	GND	Ground	
M14	VDD10S	SPWR	Power	
N1	PEX_PETp0	CMLTp	SerDes	
N2	PEX_PETn1	CMLTn	SerDes	
N3	PEX_PERp1	CMLRp	SerDes	
N4	PEX_PERn2	CMLRn	SerDes	
N5	PEX_PETp2	CMLTp	SerDes	
N6	PEX_PETn3	CMLTn	SerDes	
N7	PEX_PERp3	CMLRp	SerDes	
N8	PEX_PERp4	CMLRp	SerDes	
N9	PEX_PETn4	CMLTn	SerDes	
N10	PEX_PETp5	CMLTp	SerDes	
N11	PEX_PERn5	CMLRn	SerDes	
N12	PEX_PERp6	CMLRp	SerDes	
N13	PEX_PETn6	CMLTn	SerDes	
N14	PEX_PETp7	CMLTp	SerDes	
P1	PEX_PETn0	CMLTn	SerDes	
P2	PEX_PETp1	CMLTp	SerDes	

Table 3-11. PEX 8509 Ball Assignments by Location (Cont.)

Loc	Signal Name	Type	Signal Group	Comment
P3	PEX_PERn1	CMLRn	SerDes	
P4	PEX_PERp2	CMLRp	SerDes	
P5	PEX_PETn2	CMLTn	SerDes	
P6	PEX_PETp3	CMLTp	SerDes	
P7	PEX_PERn3	CMLRn	SerDes	
P8	PEX_PERn4	CMLRn	SerDes	
P9	PEX_PETp4	CMLTp	SerDes	
P10	PEX_PETn5	CMLTn	SerDes	
P11	PEX_PERp5	CMLRp	SerDes	
P12	PEX_PERn6	CMLRn	SerDes	
P13	PEX_PETp6	CMLTp	SerDes	
P14	PEX_PETn7	CMLTn	SerDes	

3.6 Physical Layout

Figure 3-1. PEX 8509 196-Ball Physical Ball Assignment (See-Through Top View)

	1	2	3	4	5	6	7	8	9	10	11	12	13	14		
A	VSS	STRAP_S ERDES_M ODE_ENA BLE#	HP_PWRL ED2#	PEX_LAN E_GOOD1 #	THERMAL _DIODEp	STRAP_U PSTRM_P ORTSEL1	STRAP_U PSTRM_P ORTSEL0	HP_BUTT ON2#	JTAG_TDI	JTAG_TDO	STRAP_T ESTMODE 2	STRAP_D EBUG_SE L0#	STRAP_D EBUG_SE L1#	VSS	A	
B	HP_ATNL ED1#	HP_BUTT ON1#	PEX_LAN E_GOOD3 #	PEX_LAN E_GOOD2 #	HP_PRSN T2#	PEX_LAN E_GOOD0 #	HP_PWRF LT2#	HP_MRL2 #	JTAG_TRS T#	JTAG_TMS	STRAP_T ESTMODE 1	NC_SPAR E0	HP_BUTT ON3#	HP_MRL3 #	B	
C	HP_MRL1 #	PEX_LAN E_GOOD4 #	PEX_INTA #	HP_CLKE N2#	HP_PERS T2#	THERMAL _DIODEn	HP_PWRE N2#	STRAP_U PSTRM_P ORTSEL2	HP_ATNL ED2#	JTAG_TCK	STRAP_T ESTMODE 0	STRAP_T ESTMODE 3	PEX_LAN E_GOOD7 #	HP_PWRE N3#	C	
D	HP_PWRE N1#	HP_PWRF LT1#	PEX_LAN E_GOOD5 #	VDD33	VDD33	VDD33	VDD33	VDD33	VDD33	VDD33	HP_ATNL ED3#	HP_PWRF LT3#	PEX_LAN E_GOOD6 #	NC_SPAR E1	D	
E	HP_PWRL ED1#	HP_PRSN T1#	HP_PERS T1#	VDD33	VDD10 =<= Pad 1	VSS	VDD10	VSS	VDD10	VSS	VDD33	HP_PERS T3#	HP_PWRL ED3#	STRAP_P ORTCFG1	E	
F	STRAP_PL L_BYPASS #	HP_CLKE N1#	NC_PROC MON	VDD33	VSS	VSS	VSS	VSS	VSS	VSS	VDD10	HP_PRSN T3#	HP_CLKE N3#	STRAP_P ORTCFG0	EE_CS#	F
G	VSSA_PLL	STRAP_F AST_BRIN GUP#	PEX_PER ST#	VDD33	VDD10	VSS	VSS	VSS	VSS	VSS	VSS	VDD33	EE_DO	EE_SK	STRAP_P ROBE_MO DE#	G
H	VSS	VSS	VDD33A	VDD33	VSS	VSS	VSS	VSS	VSS	VSS	VDD10	VDD33	EE_DI	I2C_SDA	I2C_SCL	H
J	PEX_REF CLKp	PEX_REF CLKn	VDD10A	VSS	VDD10	VSS	VSS	VSS	VSS	VSS	VSS	VDD33	NC_SPAR E2	FATAL_ER R#	VSS	J
K	VSS	VDD10S	VSS	VDD10	VSS	VDD10	VSS	VDD10	VSS	VDD10	I2C_ADDR 0	I2C_ADDR 1	VSS	PEX_PER n7	K	
L	PEX_PER n0	VSS	VDD10S	VTT_PEX0	VDD10A	VTT_PEX1	VDD10S	VSS	VTT_PEX2	VDD10A	VTT_PEX3	I2C_ADDR 2	VDD10S	PEX_PER p7	L	
M	PEX_PER p0	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	VSS	VDD10S	M	
N	PEX_PETp 0	PEX_PETn 1	PEX_PER p1	PEX_PER n2	PEX_PETp 2	PEX_PETn 3	PEX_PER p3	PEX_PER n4	PEX_PETn 4	PEX_PETp 5	PEX_PER n5	PEX_PER p6	PEX_PETn 6	PEX_PETp 7	N	
P	PEX_PETn 0	PEX_PETp 1	PEX_PER n1	PEX_PER p2	PEX_PETn 2	PEX_PETp 3	PEX_PER n3	PEX_PER p4	PEX_PETn 4	PEX_PETp 5	PEX_PER n5	PEX_PER p6	PEX_PETn 6	PEX_PETp 7	P	

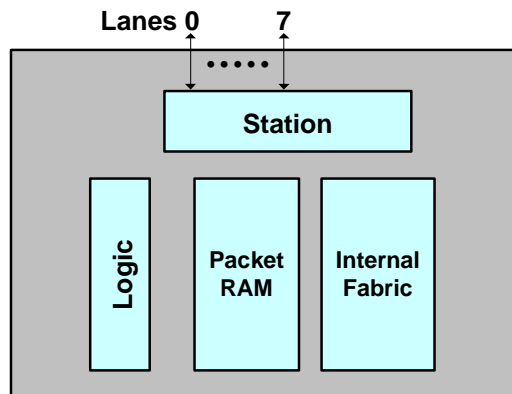
Die UP

THIS PAGE INTENTIONALLY LEFT BLANK.

4.1 Architecture

The PEX 8509 switch is designed with a flexible, modular architecture. The eight PCI Express lanes are implemented in a single station, which is connected by the internal fabric to the central RAM. [Figure 4-1](#) illustrates a block diagram of the PEX 8509.

Figure 4-1. PEX 8509 Block Diagram



4.1.1 Ingress and Egress Functions

The Switch Ingress ports interface to the Crossbar Switch fabric (internal fabric), through the PCI Express station. The PEX 8509 has a centralized packet RAM, which it uses to store all ingress traffic. Ingress traffic flows from the PCI Express station through the internal fabric, to the RAM. From the RAM, the traffic flows across the internal fabric to the Egress station, where it eventually egresses out the correct port in the station. The PEX 8509 ports support a Weighted Round-Robin-based Port Arbitration scheme, to support Quality of Service (QoS).

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 supports up to 8 integrated Serializer and De-serializer (SerDes) modules, which provide the 8 PCI Express hardware interface lanes.

The lanes within can be combined, for a total of five to eight 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. The BIOS enumerates the PEX 8509 ports, using either Conventional PCI Configuration access or PCI Express Enhanced access.

The PEX 8509 port width is configurable, by way of Strapped signal balls, serial EEPROM, or I²C download after reset. The final port width can be made narrower by auto-link width negotiation, as described in the *PCI Express Base r1.1*.

4.1.2.1 Port Combinations

[Table 4-1](#) defines the PEX 8509 port and lane configuration. Configure the Ports for one station independent of the other stations' Port configurations. Ports that are not configured nor enabled are invisible to software.

The upstream port and downstream ports' link 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. This can also be changed through the I²C interface. The final port width can be automatically negotiated down from the programmed port width, to x4, x2, or x1, through link-width negotiation for linkup to a device with fewer lanes. 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.

Table 4-1. Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[1:0])	Lanes/SerDes/Port ^{a, c}							
	Port 0	Port 1	Port 2	Port 3	Port 4	Port 5	Port 6	Port 7
00b	x1 [0]	x1 [1]	x1 [2]	x1 [3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]
01b	x2 [0-1]	x1 [2]	x1 [3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]	_b
10b	x2 [0-1]	x2 [2-3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]	-	-
11b	x4 [0-3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]	-	-	-

- a. The lanes are assigned to each enabled port in sequence, as indicated in [brackets].
- b. Configuration value and port combinations with “-” (no data) are **reserved**.
- c. Refer to [Table 4-2](#) for an explanation of the default Port/Physical Lane/SerDes Module to SerDes Quad relationship.

Table 4-2. Port/Physical Lane/SerDes Module to SerDes Quad Relationship (Port 0, Offset 224h[1:0]=0h)

Port	Physical Lanes and SerDes Modules	SerDes Quad
0	0	0
1	1	
2	2	
3	3	
4	4	1
5	5	
6	6	
7	7	


4.1.2.2 Port Numbering

The PEX 8509 Port Numbers are 0, 1, 2, 3, 4, 5, 6, and 7. (Refer to [Table 4-1](#) and [Figure 4-2](#).)

The Port Numbers have a direct relationship to the downstream ports for the PCI Device Number assigned to the internal PCI-to-PCI bridges on the internal virtual PCI Bus. *For example*, if Port 3 is a downstream port, the PCI-to-PCI bridge associated with that port is also Device Number 3. All downstream Device Numbers match their corresponding Port Number. *For example*, if Port 0 is the upstream port, Ports 1, 2, 3, 4, 5, 6, and 7 are the downstream ports. The Device Numbers for the PCI-to-PCI bridges implemented on the downstream ports are 1, 2, 3, 4, 5, 6, and 7, respectively. (Refer to [Figure 4-2](#).)

Any PEX 8509 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 4-2. PLX Port Numbering Convention Example (When Port 0 Is Upstream Port)



4.2 PCI-Compatible Software Model

The PEX 8509 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 an internal virtual bus. (Refer to [Figure 4-3.](#)) PCI-to-PCI bridges are compliant with the PCI and PCI Express system models. [Figure 4-3](#) illustrates the concept of hierarchical PCI-to-PCI bridges, with the bus in the middle being the internal 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.

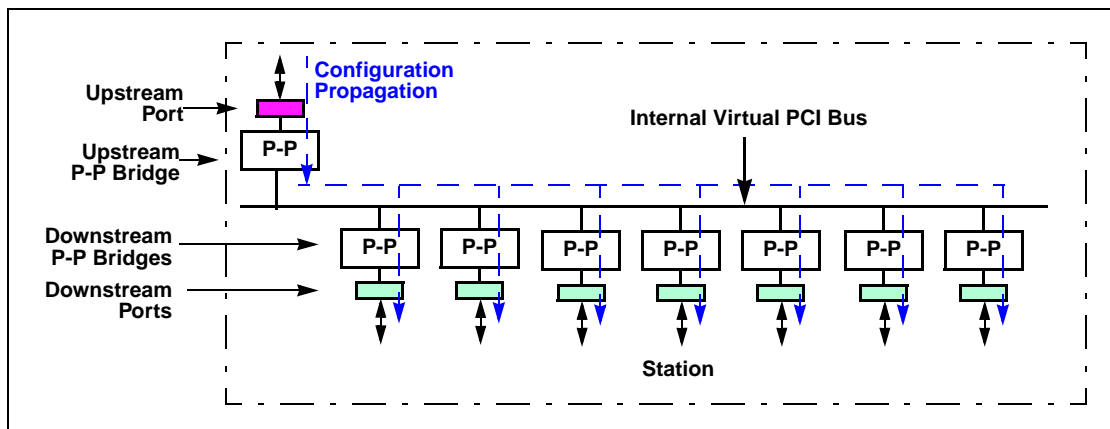
The CSRs in the downstream port PCI-to-PCI bridges are accessible by Type 1 Configuration requests received at the upstream port that target the internal virtual PCI Bus, by having a Bus Number value that matches the upstream bridge's Secondary Bus Number value. Each downstream bridge is associated with 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 8509 downstream port Secondary Bus Number, the PEX 8509 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 8509, unchanged. A Type 1 Configuration Request that targets a Bus Number that is not within range is invalid, and is terminated by the PEX 8509 upstream port as an Unsupported Request (UR).

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 8509, 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 4-3. System Configuration Propagation



4.2.1 System Reset

The PEX 8509 can be reset by four different mechanisms:

- 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.)

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 **MSI Control and Capability Header** register (offset [48h](#)) and **PCI Command** register *Interrupt Disable* bit (offset [04h\[10\]](#)).] The PEX 8509 does not convert received INT_x messages to MSI messages. (Refer to [Chapter 10, “Interrupts,”](#) for details.)

4.2.2.1 Interrupt Sources or Events

The PEX 8509-generated interrupt/message sources include:

- Hot Plug or Link Status events
- PCI Express Hot Plug events
- Device-Specific errors
- INT_x messages from downstream devices

4.2.2.2 INT_x Switch Mapping

The PEX 8509 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 10.2.1, “INT_x-Type Interrupt Message Remapping and Collapsing,”](#) for interrupt routing information.)

4.3 PCI Express Station Functional Description

The PEX 8509 groups 8 SerDes together into the station, which can be comprised of five to eight ports. (Refer to [Table 4-1](#).) The station forwards ingress packets to the internal fabric and central RAM, and the station pulls egress packets from the central RAM to send out of the PEX 8509.

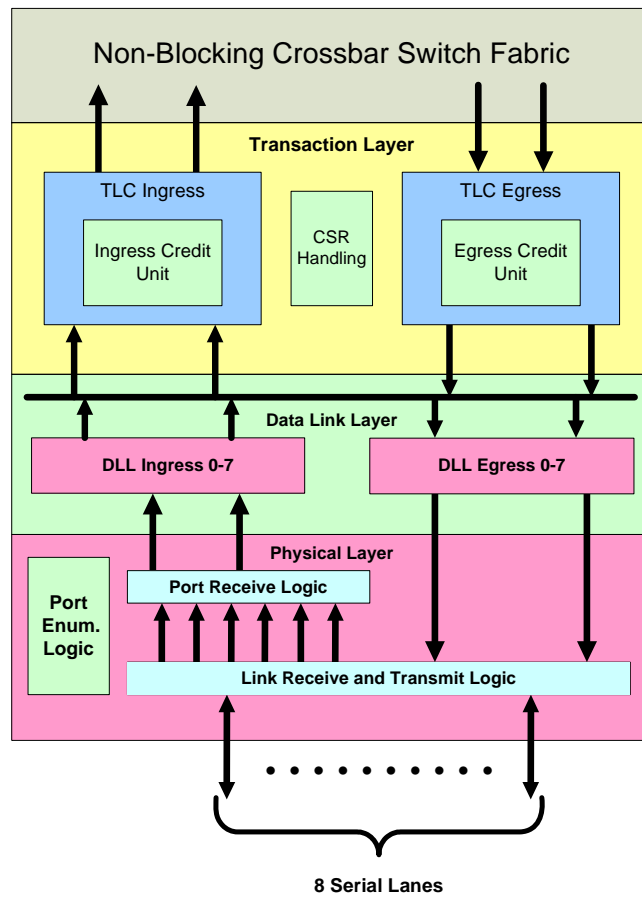
The station implements the PCI Express Physical Layer (PHY) and Data Link Layer (DLL) functions for its ports, and aggregates traffic from these ports onto a transaction-based, non-blocking internal fabric. The PCI Express station also performs many Transaction Layer functions, while the packet queuing 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. The PCI Express station can contain multiple ports (one upstream and multiple downstream). Traffic flow between the ports is supported through the central internal fabric.

4.3.1 Functional Blocks

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

Figure 4-4. PCI Express Station Block Diagram



4.3.1.1 Physical Layer

For details, refer to [Section 9.2, “Physical Layer.”](#)

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 detection and correction.

The transmission side of the Data Link Layer accepts Transaction Layer Packets (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.

For further details, refer to [Section 9.3, “Data Link Layer.”](#)

4.3.1.3 Transaction Layer Control

The upper layer of the architecture is the Transaction Layer (TL). The TL’s primary responsibility is the assembly and disassembly of TLPs. TLPs are used to communicate transactions, *such as* Read and Write, as well as certain types of events. The Transaction Layer is also responsible for managing credit-based flow control for TLPs.

Every Request packet requiring a Response packet is implemented as a Split Transaction. Each packet has a unique identifier that enables Response packets to be directed to the correct originator. The packet format supports different forms of addressing, depending upon the transaction type – *Memory, I/O, Configuration, and Message*. The packets can also have attributes, *such as No Snoop and Relaxed Ordering*.

The TL supports four Address spaces – it includes the three PCI Address spaces (Memory, I/O, and Configuration) and adds a Message space. This specification uses Message space to support all prior sideband signals, *such as* interrupts, Power Management requests, and so forth, as in-band Message transactions. PCI Express Message transactions can be thought of as *virtual wires*, because their effect is to eliminate the wide array of sideband signals currently used in a platform implementation.

The PEX 8509 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 8509 responds to “lock”-type Read Requests (MRdLk) with a Completion, having a Completion with status of Unsupported Request (UR).

For further details, refer to [Section 9.4, “Transaction Layer.”](#)

4.3.1.4 Non-Blocking Crossbar Switch Architecture

The Non-Blocking Crossbar switch is an on-chip interconnect switching fabric. The Crossbar Switch architecture is built upon 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
- PCI Express Ordering rules
- Packet fair queuing
- Oldest first scheduling

The Crossbar Switch interconnect physical topology is that of a packet-based Crossbar Switch fabric (internal fabric) designed to simultaneously connect multiple on-chip stations. The Crossbar Switch protocol is sufficiently flexible and robust to support a variety of embedded system requirements. The protocol is specifically designed to ease chip integration by strongly enforcing station boundaries and standardizing communication between stations. 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 (P, NP, and Cpl, respectively) – meet PCI and PCI Express Ordering and Deadlock Avoidance rules
- Optional weighting of source ports to support Source Port arbitration

Note: Although the internal fabric is designed to support multiple stations, the PEX 8509 has only one station.

4.3.2 Cut-Thru Mode

The PEX 8509 is designed to cut through TLPs to and from every port. By default, all ports are 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 for all ports by clearing the **Debug Control** register *Cut-Thru Enable* bit (Port 0, offset 1DCh[21]).

*Note: The **Debug Control** register *Cut-Thru Enable* bit affects the entire chip. If Cut-Thru is enabled, 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.*



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]=1).

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 8509 power-on. When the **PEX_PERST#** signal is held Low following the proper application of power to the component, a Fundamental Reset occurs.

A Fundamental Reset initializes the entire PEX 8509 device (*such as* configuration information, clocks, state machines, registers, and so forth).

When power is removed from the device, or travels outside 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 8509 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 8509 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 8509 port re-configuration. However, a Hot Reset:

- Causes all TLPs held in the PEX 8509 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 13-4, “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

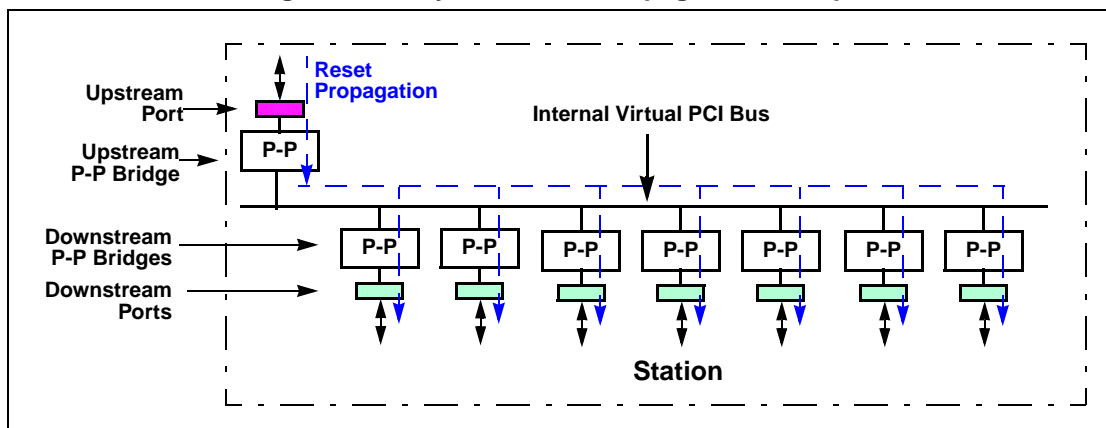
A Hot Reset is propagated to a downstream PCI Express device through the PCI Express link, using the Physical Layer Hot Reset mechanism (*that is*, a Reset bit in the Training Ordered-Set from the upstream device is set).

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

An example of reset propagation across the PEX 8509 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 for all 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. System Reset Propagation Example



5.1.3.2 Hot Reset Disable

The PEX 8509 includes a configuration option – **Debug Control** register *Upstream Port DL_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 PEX 8509 from manifesting an internal reset due to 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** port PCI-to-PCI bridge **Bridge Control** register (BCR) *Secondary Bus Reset* bit is set (offset 3Ch[22]=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** port 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 8509 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's *Secondary Bus Reset* bit. The PHY 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 8509 initialization process is started. There are two or more steps in the process, depending upon the availability of an external initialization serial EEPROM and I²C.

The initialization sequence executed is as follows:

1. PEX 8509 reads the Strapping signal balls to determine the upstream port, and lane configuration of all ports.
2. If a serial EEPROM is present, serial EEPROM data is downloaded to the PEX 8509 Configuration registers. The configuration from the Strapping signal balls can be changed by serial EEPROM data.
3. If I²C is used, it can be used to program all the registers (the same as would be done with the serial EEPROM). Because I²C is relatively slow, normally the links are already up by the time the first I²C Write occurs. The first I²C command might be to block system access while the configuration is being changed, by disabling the upstream port; ports can be disabled by setting the appropriate **Port Control** register *Disable Port x* bit (Port 0, offset 234h[23:16]).

Port Configuration (Port 0, offset 224h[1:0]) and/or upstream port designation [**Debug Control** register *Upstream Port ID* field (Port 0, offset 1DCh[11:8])] changes take effect upon subsequent Hot Reset. (Refer to Section 5.1.3.)

Note: As described in Chapter 7, “I²C Slave Interface Operation,” an external I²C Master can send the register Read/Write requests to PEX 8509 after reset. To prevent conflict, first disable the upstream port, by programming the appropriate **Port Control** register *Disable Port x* bit (Port 0, offset 234h[23:16]) to a value of 1. Restoration of the upstream port should be the last register Write of the entire I²C programming procedure.

4. After configuration from the Strapping signal balls, serial EEPROM, and/or I²C is complete, 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 link-width configuration is determined by the 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.”](#)) Some of these settings can be changed by downloading serial EEPROM data or by I²C programming.

5.2.2 Default Register Initialization

Each PEX 8509 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, “Port 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 I²C 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 8509 supports three 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”](#)

5.2.3 Device-Specific Registers

The following registers are unique to the PEX 8509, and are not referenced in PCI Express documentation. The Device-Specific registers are organized into the following sections:

- [Section 13.16.1, “Device-Specific Registers – Error Checking and Debug”](#)
- [Section 13.16.2, “Device-Specific Registers – Physical Layer”](#)
- [Section 13.16.3, “Device-Specific Registers – I²C Interface”](#)
- [Section 13.16.4, “Device-Specific Registers – Bus Number CAM”](#)
- [Section 13.16.5, “Device-Specific Registers – I/O CAM”](#)
- [Section 13.16.6, “Device-Specific Registers – Address-Mapping CAM”](#)
- [Section 13.16.7, “Device-Specific Registers – Ingress Control and Port Enable”](#)
- [Section 13.16.8, “Device-Specific Registers – I/O CAM Base and Limit Upper 16 Bits”](#)
- [Section 13.16.9, “Device-Specific Registers – Base Address Shadow”](#)
- [Section 13.16.10, “Device-Specific Registers – Shadow Virtual Channel Capability”](#)
- [Section 13.16.11, “Device-Specific Registers – Ingress Credit Handler”](#)
- [Section 13.16.12, “Device-Specific Registers – Port Configuration Header”](#)
- [Section 13.16.13, “Device-Specific Registers – Source Queue Weight and Soft Error”](#)

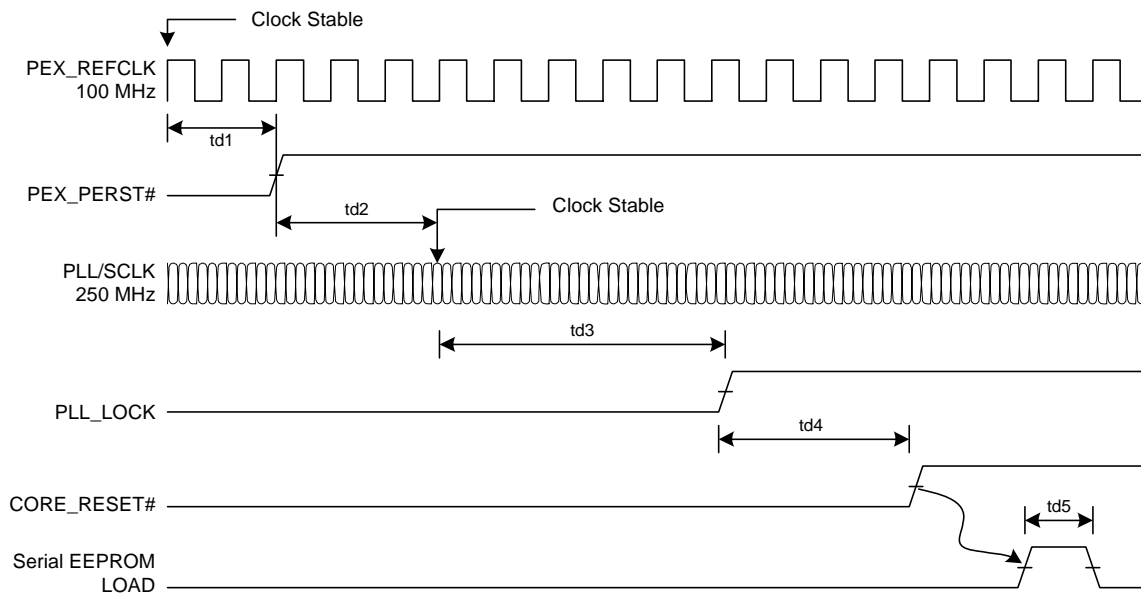
The Device-Specific registers cannot be accessed by Configuration requests; however, software can access these registers with Memory requests.

5.2.4 Reset and Clock Initialization Timing

Table 5-1. Reset and Clock Initialization Timing

Symbol	Description	Typical Delay
td1	REFCLK stable to PEX_Reset release time	100 μ s
td2	PEX_Reset release to PLL Clock Stable and Reset de-bounce	1.32 ms
td3	Clock and Reset Stable to PLL Lock	125 μ s
td4	PLL Lock to Core Reset release	33 μ s
td5	Serial EEPROM load time with no serial EEPROM present	17 μ s

Figure 5-2. Reset and Clock Initialization Timing



5.2.4.1 Serial EEPROM Load Time

Serial EEPROM initialization loads only the Configuration register data that is specifically programmed into the serial EEPROM. Registers that are not included in the serial EEPROM data are initialized to default register values.

Each register entry in the serial EEPROM consists of two Address bytes and four Data bytes (refer to [Section 6.2, “Serial EEPROM Data Format”](#)); therefore, each register entry (6 bytes, or 48 bits) requires 48 serial EEPROM clocks to download. Thus, at the serial EEPROM clock default frequency of 1 MHz, after initial overhead to read the [Serial EEPROM Status](#) register (Port 0, offset [260h](#)) (16 serial EEPROM clocks, or 16 μ s), plus another 40 serial EEPROM clocks (40 μ s) to begin reading the register data, each register entry in the serial EEPROM requires 48 μ s to download. A serial EEPROM containing 50 register entries (typical configuration, assuming the serial EEPROM is programmed only with non-default register values) and clocked at 1 MHz takes approximately 2.5 ms to load (16 + 40 + 48 * 50) μ s (5,200 μ s).

To reduce the serial EEPROM initialization time, the first register entry in the serial EEPROM could increase the clock frequency by programming the [Serial EEPROM Clock Frequency](#) register (Port 0, offset [268h](#)), to a value of 2h (5 MHz), or 3h (9.62 MHz), if the serial EEPROM supports the higher frequency at the serial EEPROM supply voltage (typically 3.3V). At 5 MHz clocking, the serial EEPROM load time for 50 register entries could be reduced to approximately 575 μ s. Because the *PCI Express Base r1.1* allows a 20-ms budget for system hardware initialization, the default 1-MHz serial EEPROM clock is often sufficient when the number of ports and registers programmed by serial EEPROM is relatively small.

5.2.4.2 I²C Load Time

Initialization using I²C is slower than serial EEPROM initialization, because the I²C interface operates at a lower clock frequency (100 KHz maximum) and the number of bits per Register access is increased (because the Device address is included in the bit stream). Writing one register using 100-KHz clocking takes approximately 830 μ s (83 clock periods).

THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 6 Serial EEPROM Controller

6.1 Overview

The PEX 8509 provides an interface to SPI (Serial Peripheral Interface)-compatible serial EEPROMs. This interface consists of a Chip Select, Clock, Write Data, and Read Data balls, and operates at a programmable frequency of up to 17.86 MHz. The PEX 8509 supports up to a 16-MB serial EEPROM, utilizing 1-, 2-, or 3-byte addressing. The PEX 8509 automatically determines the appropriate addressing mode.


The primary function of the Serial EEPROM Controller is to allow access to non-volatile memory from the PEX 8509. This is accomplished using two different methods:

- The first method of access to a serial EEPROM device is during initialization. When a serial EEPROM device is connected to the PEX 8509, during initialization, the Serial EEPROM Controller reads data from the serial EEPROM to be used to update the default values of registers within the PEX 8509.
- The second method of access is controlled by software initiating a Read or Write Request to the serial EEPROM, to store or retrieve other data.

The on-chip Serial EEPROM Controller is integrated into the PEX 8509, as illustrated in [Figure 6-1](#). The controller performs a serial EEPROM download when the following conditions exist:

- Serial EEPROM is present, and
- Validation signature (first byte read from the serial EEPROM) value is **5Ah**, and
- 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 **Debug Control** register *Disable Serial EEPROM Load on Hot Reset* bit (Port 0, offset 1DCh[17]=1)]
 - Upstream port exits a DL_Down state [downloading upon this event can be optionally disabled, by setting the **Debug Control** register *Upstream Port DL_Down Reset Propagation Disable* bit (Port 0, offset 1DCh[20]=1)]

Figure 6-1. Serial EEPROM Connections



6.2 Serial EEPROM Data Format

The data in the serial EEPROM is stored in the format defined in [Table 6-1](#). The Validation Signature byte is located in the first address. The Serial EEPROM Controller reads this byte to determine whether a valid serial EEPROM image exists versus a blank image. REG_BYTE_COUNT[15:0] contains the number of bytes of serial EEPROM data to be loaded. It is equal to the number of registers to be loaded times six (6 serial EEPROM bytes per register). For the remaining register-related locations, data is written into a 2-byte address that represents the Configuration register offset and Port Number, and the 4 bytes following are the data loaded for that Configuration register. Only Configuration register data specifically programmed into the serial EEPROM is loaded after the PEX 8509 exits reset.

Table 6-1. Serial EEPROM Data

Location	Value	Description
0h	5Ah	Validation Signature
1h	–	<i>Reserved</i>
2h	REG BYTE COUNT (LSB)	Configuration register Byte Count (LSB)
3h	REG BYTE COUNT (MSB)	Configuration register Byte Count (MSB)
4h	REGADDR (LSB)	1 st Configuration Register Address (LSB)
5h	REGADDR (MSB)	1 st Configuration Register Address (MSB)
6h	REGDATA (Byte 0)	1 st Configuration Register Data (Byte 0)
7h	REGDATA (Byte 1)	1 st Configuration Register Data (Byte 1)
8h	REGDATA (Byte 2)	1 st Configuration Register Data (Byte 2)
9h	REGDATA (Byte 3)	1 st Configuration Register Data (Byte 3)
Ah	REGADDR (LSB)	2 nd Configuration Register Address (LSB)
Bh	REGADDR (MSB)	2 nd Configuration Register Address (MSB)
Ch	REGDATA (Byte 0)	2 nd Configuration Register Data (Byte 0)
Dh	REGDATA (Byte 1)	2 nd Configuration Register Data (Byte 1)
Eh	REGDATA (Byte 2)	2 nd Configuration Register Data (Byte 2)
Fh	REGDATA (Byte 3)	2 nd Configuration Register Data (Byte 3)
.....		
FFFFh	REGDATA (Byte 3)	Last Configuration Register Data (Byte 3)

Table 6-2 defines the Configuration register Address format (REGADDR[15:0] from Table 6-1):

- Bits [9:0] represent bits [11:2] of the Register address
- Bits [15:10] represent the Port Number of the register selected to be programmed by serial EEPROM

The values stored in the serial EEPROM register Address bytes are different from the actual Configuration register address specified when the user writes a register address to the serial EEPROM. To determine what value is stored in the serial EEPROM for the Configuration register address, shift the address value by two bits to the right, then OR the remaining 10 bits with the values listed in Table 6-2, based upon the Port Number to which the address must be programmed.

For example, to load register offset 1F8h to Port 4, shift the address to the right by 2 bits (this becomes 07Eh) and concatenate 0010_00b. The resulting DWord address in the serial EEPROM will be 0010_0000_0111_1110b, which is 207Eh.

Table 6-2. Configuration Register Address Format

Port Number	REGADDR Bits [15:10] Value ^a	Port Identifier
Port 0	0000_00b	0000h
Port 1	0000_01b	0400h
Port 2	0000_10b	0800h
Port 3	0000_11b	0C00h
Port 4	0001_00b	1000h
Port 5	0001_01b	1400h
Port 6	0001_10b	1800h
Port 7	0001_11b	1C00h

a. Encodings not listed are *reserved*.

6.3 Serial EEPROM Initialization

After the device Reset is de-asserted, the serial EEPROM's internal **Status** register is read to determine whether a serial EEPROM is installed. A pull-up resistor on the **EE_DO** input produces a value of FFh if a serial EEPROM is not installed. If a serial EEPROM is detected, the first byte (validation signature) is read. If a value of 5Ah is read, it is assumed that the serial EEPROM is programmed for the PEX 8509. The serial EEPROM address width is determined while the first byte is read. If the first byte's value is not 5Ah, the serial EEPROM is blank or programmed with invalid data. In this case, no more data is read from the serial EEPROM. Also, the **Serial EEPROM Status** register *EepAddrWidth* field (Port 0, offset 260h[23:22]) reports a value of 00b (undetermined width). If the *EepAddrWidth* field reports a value of 00b, any subsequent accesses to the serial EEPROM (through the PEX 8509 Serial EEPROM registers) default to a serial EEPROM address width of 1 byte, unless the *EepAddrWidthOverride* bit is set (Port 0, offset 260h[21]=1). This allows the *EepAddrWidth* field to be manually set.

If the serial EEPROM contains valid data, the REG_BYTE_COUNT values in Bytes 2 and 3 determine the number of serial EEPROM locations that contain Configuration register addresses and data. Each Configuration register entry consists of two bytes of register Address and four bytes of register Write data. The REG_BYTE_COUNT must be a multiple of 6.

The **EE_SK** output clock frequency is determined by the **Serial EEPROM Clock Frequency** register *EepFreq[2:0]* field (Port 0, offset 268h[2:0]). The default clock frequency is 1 MHz. At this clock rate, it takes approximately 48 μ s per DWORD during Configuration register initialization. For faster loading of large serial EEPROMs that support a faster clock, the first Configuration register load from the serial EEPROM could be to the **Serial EEPROM Clock Frequency** register.

6.4 PCI Express Configuration, Control, and Status Registers

The PCI Express Configuration, Control, and Status registers that can be initialized are discussed in Chapter 13, "Port Registers."

6.5 Serial EEPROM Registers

The serial EEPROM parameters, for the registers (Port 0, offsets 260h through 26Ch) defined in Section 13.16.2, "Device-Specific Registers – Physical Layer," can be changed using the serial EEPROM. It is recommended that the first entry in the serial EEPROM change the value in the **Serial EEPROM Clock Frequency** register (Port 0, offset 268h) to increase the clock frequency, to reduce the time needed for the remainder of the serial EEPROM load. At the last entry of the serial EEPROM, the **Serial EEPROM Status and Control** register (Port 0, offset 260h) can be programmed to issue a Write Disable (WRDI) command and enable the write protection of serial EEPROM.

6.6 Serial EEPROM Random Write/Read Access

To access the serial EEPROM, a PCI Express Master uses the following registers:

- **Serial EEPROM Status and Control** (Port 0, offset 260h)
- **Serial EEPROM Data Buffer** (Port 0, offset 264h)
- **Serial EEPROM 3rd Address Byte** (Port 0, offset 26Ch)

The PCI Express Master can only access the serial EEPROM on a DWord basis (4 bytes).

6.6.1 Writing to Serial EEPROM

To write a DWord to the serial EEPROM:

1. If the 3rd Address byte (Address bits [23:16]) is needed [when the **Serial EEPROM Status** register *EepAddrWidth* field is set (Port 0, offset 260h[23:22]=11b)], write the value to the *Serial EEPROM 3rd Address Byte* field (Port 0, offset 26Ch[7:0]).
2. Write the 32-bit data into the **Serial EEPROM Data Buffer** register (Port 0, offset 264h).
3. Issue a Write Enable instruction to the serial EEPROM (Command = 110b, Set Write Enable Latch), by writing the value 0000_C000h into the **Serial EEPROM Status and Control** register (Port 0, offset 260h).
4. Calculate the combined Address and Command value to write into the **Serial EEPROM Control** register, by combining the serial EEPROM 3-bit Write Data instruction (value 010b) as bits [15:13], together with the serial EEPROM address. Serial EEPROM Address bits [14:2] must be programmed into **Serial EEPROM Control** register bits [12:0], and serial EEPROM Address bit 15 must be programmed into **Serial EEPROM Status** register bit 20 (*that is*, set bit 20 if the serial EEPROM address is in the upper 32 KB of any 64-KB address block within the serial EEPROM). The data in the **Serial EEPROM Data Buffer** register is written to the serial EEPROM when the **Serial EEPROM Status and Control** register is written.
5. The serial EEPROM Write operation is complete when a subsequent read of the **Serial EEPROM Status** register bit 18 returns 0. At this time, another serial EEPROM access can be started.
6. Issue a Write Disable instruction to the serial EEPROM (Command = 100b, Reset Write Enable Latch), by writing the value 0000_8000h to the **Serial EEPROM Status and Control** register.

Because each PEX 8509 port and register address value (REGADDR, refer to [Section 6.3](#)) and corresponding data value (REGDATA) require 6 bytes of serial EEPROM memory, and the PEX 8509 serial EEPROM interface accesses 4 bytes at a time, two serial EEPROM writes may be needed to store each set of REGADDR (one Word) and REGDATA (one Dword) entries into the serial EEPROM. To avoid overwriting a Word of another set of 6-byte REGADDR and REGDATA values, one of the two Serial EEPROM Writes might need to be a Read-Modify-Write type of operation (preserving one Word read from the serial EEPROM and writing the value back along with a new Word value).

6.6.2 Reading from Serial EEPROM

To read a DWord from the serial EEPROM:

1. If the 3rd Address byte (Address bits [23:16]) is needed [when the **Serial EEPROM Status** register *EepAddrWidth* field is set (Port 0, offset 260h[23:22]=11b)], write the value to the *Serial EEPROM 3rd Address Byte* field (Port 0, offset 26Ch[7:0]).
2. Calculate and write the combined Address and Command value to write into the **Serial EEPROM Control** register (Port 0, offset 260h), by combining the serial EEPROM 3-bit Read Data instruction (value 011b) as bits [15:13], together with the serial EEPROM address. Serial EEPROM Address bits [14:2] must be programmed into **Serial EEPROM Control** register bits [12:0], and serial EEPROM Address bit 15 must be programmed into **Serial EEPROM Status** register bit 20 (*that is*, set bit 20 if the serial EEPROM address is in the upper 32 KB of any 64-KB address block within the serial EEPROM).
3. Poll the **Serial EEPROM Status** register until the *EepCmdStatus* bit is cleared (Port 0, offset 260h[18]=0), which signals that the transaction is complete.
4. Read the four bytes of serial EEPROM data from the **Serial EEPROM Data Buffer** register (Port 0, offset 264h).

THIS PAGE INTENTIONALLY LEFT BLANK.

7.1 Introduction

This chapter describes to the I²C Slave interface, which uses the `I2C_ADDR[2:0]`, `I2C_SCL`, and `I2C_SDA` signals for PEX 8509 register access by an I²C Master.

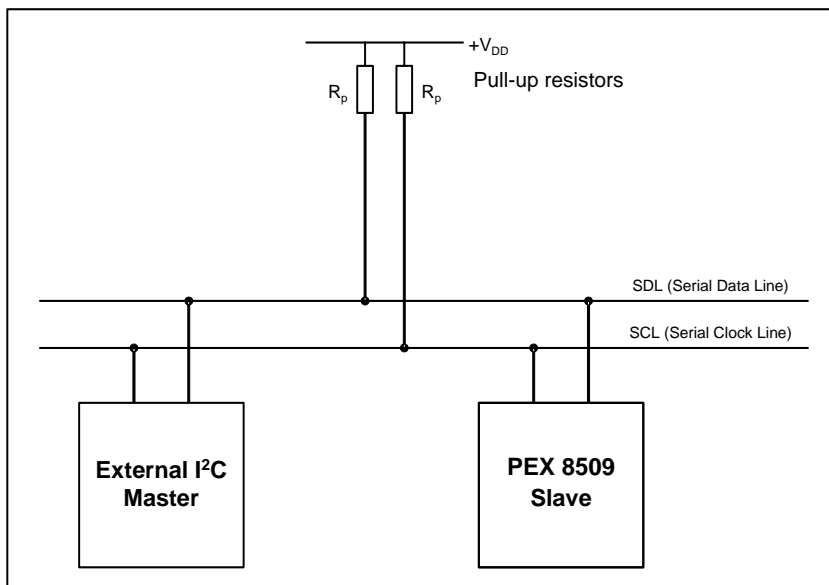
7.2 I²C Support Overview

Inter-Integrated Circuit (I²C) is a bus used to connect Integrated Circuits (ICs). Multiple ICs are connected to an I²C Bus and each IC can act as a Master by initiating a Data transfer. I²C is used for Data transfers between ICs at relatively low rates (100 Kbps) and is used in a variety of applications. For further details regarding I²C Buses, refer to the [I2C Bus, v2.1](#).

The PEX 8509 is an I²C Slave. Slave operations allow the PEX 8509 Configuration registers to be read from or written to by an I²C Master, external from the device. I²C is a sideband mechanism that allows the device Configuration registers to be programmed, read from, or written to, independent of the PCI Express upstream link.

In the past, either a serial EEPROM was required, or the PEX 8509 Configuration registers could only be accessed through a working (successful linkup) PCI Express upstream link. With I²C, users now have the option of programming the Configuration Space registers through the I²C interface. This is useful for debugging purposes, if the PEX 8509 upstream port fails to link up. I²C also provides an alternative to using a serial EEPROM.

Figure 7-1. Standard Devices to I²C Bus Connection Block Diagram



7.3 I²C Addressing – Slave Mode Access

To access the PEX 8509 Configuration registers through the I²C interface, the PEX 8509 I²C Slave address must be configured.

The PEX 8509 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 3Fh), 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 8509 devices can share the same I²C Bus segment without conflict, provided each PEX 8509 switch has its set of I2C_ADDR[2:0] inputs strapped to a unique combination.

7.4 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 7-2a](#). 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 8509 Port Number being accessed
- Byte Enable(s) of the register data being accessed

When the I²C Master is writing to the PEX 8509, 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 8509, 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 7-1 describes each Command Phase byte. Figure 7-2b illustrates the Command phase portion of an I²C Write packet.

Table 7-1. Command Phase Format

Byte	Bit(s)	Description
1 st (0)	7:3	Reserved Must be Cleared to 0000_0b.
	2:0	Command 011b = Write register 100b = Read register All other encodings are <i>reserved. Do not use.</i>
2 nd (1)	7:3	Reserved Must be Cleared to 0000_0b.
	2:0	Port Selector, Bits [3:1] 2 nd Byte, bits [2:0], and 3 rd Byte, bit 7, combine to form a 4-bit Port Selector.
3 rd (2)	7	Port Selector, Bit 0 2 nd Byte, bits [2:0], and 3 rd Byte, bit 7, combine to form a 4-bit Port Selector. <i>Port Selector, bits [3:1] (2nd Byte, bits [2:0]) select the port to access – only values 0h, 1h, 2h, 3h, 4h, 5h, 6h, and 7h are valid.</i>
	6	Reserved Must be cleared to 0.
	5:2	1 = Indicates corresponding PEX 8509 register byte is modified Bit Description 2 Byte Enable for Byte 0 (PEX 8509 register bits [7:0]) 3 Byte Enable for Byte 1 (PEX 8509 register bits [15:8]) 4 Byte Enable for Byte 2 (PEX 8509 register bits [23:16]) 5 Byte Enable for Byte 3 (PEX 8509 register bits [31:24])
	1:0	PEX 8509 Register Address [11:10]
4 th (3)	7:0	PEX 8509 Register Address [9:2] <i>Note: All register addresses are DWord-aligned. Therefore, bits [1:0] are always cleared to 00b.</i>

7.5 I²C Interface Register

The **I²C Interface** register, **I²C Configuration** (offset 294h), is described in [Section 13.16.3, “Device-Specific Registers – I²C Interface.”](#) This register is accessible only from Port 0. The default I²C Slave address can be changed in the **I²C Configuration** register to a different value, using the serial EEPROM or a Memory Write. **The I²C Slave address must not be changed by an I²C Write command.** (Refer to [Section 7.3.](#))

Other I²C interface registers exist; however, they are for *Factory Test Only*.

7.6 I²C Register Write Access

The PEX 8509 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 7-2](#) defines mapping of the I²C Data bytes to the Configuration register Data bytes.

[Figure 7-2c](#) 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 8509 register is not modified.

The PEX 8509 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 8509 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 I²C protocol, refer to the [I2C Bus, v2.1.](#))

In the packet described in [Figure 7-2](#), Command Bytes 0 through 3 follow the format specified in [Table 7-1](#).

Table 7-2. I²C Register Write Access

I ² C 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 7-2. I²C Write Packet

Figure 7-2.a I²C Write Packet Address Phase Bytes

1 st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address[7:1]	Read/Write Bit 0 = Write	A

Figure 7-2.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 7-2.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)	N	P

7.6.1 Register Write

The following tables illustrate a sample I²C packet for writing the PEX 8509 **MSI Upper Address** register (offset **50h**) for Port 7, with data 1234_5678h.

Note: The PEX 8509 has a default I²C Slave address of **3Fh**, with the **I2C_ADDR[2:0]** balls having a value of **111b**. 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 7-3. I²C Register Write Access Example – 1st Cycle

Phase	Value	Description
Address	70h	Bits [7:1] for PEX 8509 I²C Slave Address (3Fh) Last bit (bit 0) for Write = 0.

Table 7-4. I²C Register Write Access Example – Command Cycle

Byte	Value	Description
0	03h	[7:3] Reserved Must be Cleared to 0000_0b. [2:0] Command 011b = Write
1	03h	[7:3] Reserved Must be Cleared to 0000_0b. [2:0] Port Selector, Bits [3:1]
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 7-5. 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 7-3. I²C Write Command Packet Example

Figure 7-3.a I²C Write Packet Address Phase Bytes

1 st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address 0111_111b	Read/Write Bit 0 0 = Write	A

Figure 7-3.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_0011b	A	Command Byte 2 1011_1100b	A	Command Byte 3 0001_0100b	A

Figure 7-3.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	N	P

7.7 I²C Register Read Access

When the I²C Master attempts to read a PEX 8509 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 8509 re-transmits the same byte sequence, starting from Byte 3 of the internal buffer.

The 1st and 2nd I²C Read packets (illustrated in [Figure 7-4](#) and [Figure 7-5](#), 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 8509 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 7-4. I²C Read Command Packet (1st Packet)

Figure 7-4.a I²C Read Command Packet Address Phase Bytes

1 st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address[7:1]	Read/Write Bit 0 = Write	A

Figure 7-4.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 7-5. I²C Read Data Packet (2nd Packet)

Figure 7-5.a I²C Read Data Packet Address Phase Bytes

1 st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address[7:1]	Read/Write Bit, 1 = Read	A

Figure 7-5.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
Register Byte 3	A	Register Byte 2	A	Register Byte 1	A	Register Byte 0	P

7.7.1 Register Read Address Phase and Command Packet

The following is a sample I²C packet for reading the PEX 8509 **Serial EEPROM Data Buffer** register (Port 0, offset 264h) for Port 7, assuming the register value is ABCD_EF01h.

Note: The PEX 8509 has a default I²C Slave address of 3Fh, with the I2C_ADDR[2:0] bits having a value of 111b. 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 7-6. I²C Register Read Access Example – 1st Cycle

Phase	Value	Description
Address	71h	Bits [7:1] for PEX 8509 I²C Slave Address (3Fh) Last bit (bit 0) for Write = 0.

Table 7-7. I²C Register Read Access Example – Command Cycle

Byte	Value	Description
0	04h	[7:3] Reserved Must be Cleared to 0000_0b. [2:0] Command 100b = Read
1	03h	[7:3] Reserved Must be Cleared to 0000_0b. [2:0] Port Selector, Bits [3:1]
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	99h	[7:0] Register Address, Bits [9:2]

7.7.2 Register Read Data Packet

Note: The PEX 8509 has a default I²C Slave address of 3Fh, with the I2C_ADDR[2:0] bits having a value of 111b. 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 7-8. I²C Register Read Access Example – 1st Cycle

Phase	Value	Description
Address	71h	Bits [7:1] for PEX 8509 I²C Slave Address (3Fh) 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 7-6. 1st Packet – Command Phase

1 st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address 0111_111b	Read/Write Bit 0 = Write	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_0011b	A	Command Byte 2 0011_1011b	A	Command Byte 3 1001_1001b	P

Figure 7-7. 2nd Packet – Read Phase

1 st Cycle			
START	7 6 5 4 3 2 1	0	ACK/NAK
S	Slave Address[7:1] 0111_111b	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
Register Byte 3 1010_1011b	A	Register Byte 2 1100_1101b	A	Register Byte 1 1110_1111b	A	Register Byte 0 0000_0001b	P

THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 8 Performance Metrics

8.1 Introduction

This chapter discusses measures of performance, including throughput and latency. It also provides guidelines for programming on-chip registers to boost performance beyond that provided by the general-purpose default values.

8.2 Throughput

Throughput measures the amount of Payload bytes that are transferred per unit time. PCI Express has different possible throughput values, depending upon the link width, Payload size, traffic distribution, and TLP overhead, all of which are under software control. To comprehend PCI Express throughput, a basic understanding of the underlying PCI Express fundamentals is needed.

8.2.1 Shared Wire

Bytes are transmitted across PCI Express wires during each symbol time, regardless of traffic load. The bytes are classified into three wire traffic types:

- TLPs (which can carry a Payload)
- DLLPs
- PHY Ordered-Sets

Electrical Idles (including PADs) are not counted as traffic. To measure throughput and understand how the link performs, count all three wire traffic types while tracking how much time passes. PHY SKIP Ordered-Sets occur irregularly and can mostly be ignored. A fully utilized link requires 99% TLPs and DLLPs in each direction. The ratio of TLPs to DLLPs depends upon the application.

8.2.2 Unidirectional Throughput


[Figure 8-1](#) illustrates ideal PCI Express throughput in the case of unidirectional traffic. The PEX 8509 tracks to the ideal for these cases.

[Figure 8-1](#) also illustrates how the Payload size increases in tandem with the maximum throughput. As the link becomes wider, the maximum throughput increases. The *PCI Express Base r1.1* allows a default Maximum Payload Size (MPS) of 128 bytes; however, the PEX 8509 supports up to a 1,024-byte MPS.

Unidirectional PCI Express throughput has maximal TLPs on the wire going in one direction. The other direction of the bidirectional link is mostly unused. DLLPs that share the wire (per the *PCI Express Base r1.1*) are typically transmitted in response to a TLP, and therefore travel in the opposite direction of the TLPs. Thus, for unidirectional traffic, DLLP traffic does not interfere with TLP bandwidth.

It is useful to make a clarification, regarding Memory Read (MRd) requests and their corresponding Completions with Data (CPLD). The Read size can be large – the PEX 8509 supports up to 4 KB; however, the MRd TLP is only 12 to 20 bytes long. The Completion for the MRd carries the data. Typically, a Root Complex transmits multiple, partial Completions of a maximum 64-byte Payload size (endpoint devices must transmit Completions of at least 128-byte granularity). As a result, even with large Read sizes, the bandwidth expected for Completions is limited to 64-byte (or 128-byte) Payloads. A series of MRd requests in the upstream direction, each with a large Read size, results in a nearly unidirectional, 64-byte (or 128-byte, for non-Root Complex devices) CPLD data stream.

Figure 8-1. Ideal PCI Express Throughput in Unidirectional Bandwidth Case



8.2.3 Ideal PCI Express Throughput

This section discusses how to calculate ideal PCI Express throughput, as explained in [Section 8.2.2](#).

PEX 8509 signaling operates at 2,500 Mbps/lane. The PEX 8509 allows lanes to be grouped into x1, x2, and x4 widths. This bandwidth is de-rated, according to the factors described below.

PCI Express protocol has a built-in 8b/10b encoding, which immediately takes 20% off the throughput:

$$8b/10b_encoding_hit = 8/10 = 0.8$$

TLPs include overhead as part of the PCI Express protocol. Each TLP has a Header of 12 to 16 bytes (16 bytes are use only for 64-bit addressing; otherwise, all TLPs have 12-byte Headers). TLPs can also have an optional ECRC of 4 bytes. Finally, all TLPs require a Data Link Layer and Physical Layer framing symbol overhead of 8 bytes. The total TLP overhead is as follows:

$$TLP_overhead_min = 12 + 8 = 20 \text{ bytes}$$

$$TLP_overhead_max = 16 + 4 + 8 = 28 \text{ bytes}$$

The *PCI Express Base r1.1* requires that DLLPs and SKIP Ordered-Sets share the same wire as TLPs, allowing these other traffic sources to reduce TLP throughput. The best case (least impact) reduction to TLP bandwidth that corresponds to unidirectional traffic scenarios can be calculated.

The *PCI Express Base r1.1* requires both an updateFC (all three types) and ACK to be transmitted every 30 μ s, to cover lossy link behavior. Each DLLP takes 8 bytes. On a x4 link, a DLLP takes 2 symbol times; therefore, 3 DLLPs cost 6 symbol times every 7,500 symbol times. This hit is much less than 1%; however, to be complete, the throughput decrease from DLLPs in unidirectional traffic must be as follows:

$$DLLP_x1_unidirectional_derating = 7,500/7,524 \text{ symbol times}$$

$$DLLP_x2_unidirectional_derating = 7,500/7,512 \text{ symbol times}$$

$$DLLP_x4_unidirectional_derating = 7,500/7,506 \text{ symbol times}$$

A SKIP Ordered-Set can be modeled as coming once per 1,180 symbol times. The *PCI Express Base r1.1* provides a range of 1,180 to 1,538 symbol times. The value used by the PEX 8509 is once every 1,100 symbol times. A SKIP Ordered-Set requires 4 symbol times to transmit; therefore, throughput is decreased by:

$$SKIP_derating = (1,180/1,184)$$

Placing together all the overhead and throughput derating, the ideal PCI Express unidirectional bandwidth can be calculated for any Payload, as follows:

$$\text{Ideal_pcie_bandwidth} = (\text{link_rate}) * (8b10b_encoding_hit) * (\text{dllp_derating}) * (\text{skip_derating}) \\ * [\text{payload} / (\text{payload} + \text{tlpoverhead})]$$

The above formula was used to create the ideal curves in the graph illustrated in [Figure 8-1](#), using `tlp_overhead_min` and the appropriate link rate.

8.2.4 Bidirectional PCI Express Throughput

Although unidirectional flows have virtually no DLLP traffic flowing in the same direction as the TLP, to model bidirectional traffic, DLLPs require prominent consideration. Three different DLLP rates provide a useful reference – 0, 1, and 2 DLLP/TLP.

The worst case, 2 DLLP/TLP, applies wherein every TLP causes an ACK and updateFC DLLP. The ACK acknowledges that the TLP arrived, and the updateFC provides more credits, allowing more TLPs of the same type to be transmitted.

Note: Worst case is approximate. There can be an additional updateFC time every 7,500 symbol times.

The best case, 0 DLLP/TLP, is the unidirectional traffic case, because no DLLPs travel in the same direction as the TLP flow.

It is expected that a maximally and optimally used link will have throughput in the range of 0 to 2 DLLP/TLP for any traffic pattern. Because DLLP and TLP Counts are easily measured with standard PCI Express logic analyzers, if the DLLP Count is more than 2x the TLP Count, the link is probably underutilized.

The exact ratio of DLLPs to TLPs depends upon a variety of factors that, to some extent, remain outside the *PCI Express Base r1.1* guidelines. [Figure 8-2](#), [Figure 8-3](#), and [Figure 8-4](#) illustrate the measured PEX 8509 bidirectional throughput with default register values for x4, x2, and x1, respectively, compared against three different ideal DLLP policies of 0, 1, or 2 DLLP/TLP. The graphs provided in [Figure 8-2](#), [Figure 8-3](#), and [Figure 8-4](#) are based upon throughput measured for sustained back-to-back TLPs of the same size. The ideal calculated bandwidths are used for reference. The 0 DLLP/TLP graph is always the top graph (*that is*, the best performance). The middle graph is 1 DLLP/TLP, and the lower graph is 2 DLLP/TLP, which are below 0 DLLP/TLP in performance.

In general, for all link widths, the PEX 8509 (by default) runs at better than 1 DLLP/TLP for Payload sizes of 32 to 256 bytes. For larger Payload sizes, the default register values require fine-tuning, to allow for improved throughput. Tuning and consideration factors are discussed in further detail in [Section 8.3](#).

Figure 8-2. Measured PEX 8509 Bidirectional Throughput (TLP x4 Payload Bandwidth)




Figure 8-3. Measured PEX 8509 Bidirectional Throughput (TLP x2 Payload Bandwidth)





Figure 8-4. Measured PEX 8509 Bidirectional Throughput (TLP x1 Payload Bandwidth)



8.3 DLLP Policies

As previously discussed, DLLP rates can vary from 0 to 2 (or more) DLLP/TLP. The PEX 8509 allows programming to affect the DLLP rate. The graphs illustrated in [Figure 8-2](#), [Figure 8-3](#), and [Figure 8-4](#) clearly illustrate that more DLLPs can reduce the total TLP throughput; therefore, it seems that the DLLP rates should be minimized. However, transmitting fewer DLLPs can result in credit starvation or Replay buffer overflow, which can have a worse effect on TLP bandwidth. Thus, care must be taken when changing the default PEX 8509 DLLP transmission rate.

Typically, TLPs have higher transmission priority on the wire than DLLPs. However, the PEX 8509 allows for DLLPs to have higher priority under certain conditions, which means DLLPs can transmit after the current TLP completes. The decision to transmit a DLLP ahead of a TLP is referred to as *DLLP policy*.

The PEX 8509 can be programmed to alter its default DLLP policies, to emphasize improved TLP throughput, faster acknowledgement, more credit, or simplest behavior. The default policies should work for most applications. However, choices for a DLLP policy allow for further optimization to match any application.

8.3.1 ACK DLLP Policy

An ACK DLLP is a response indicating to the TLP transmitter that the receiver received a “good” copy of the TLP (it acknowledges the receipt). What an ACK means to the TLP transmitter is that it can remove any stored copy of that TLP, because it is unnecessary to resend. ACK DLLPs can be combined, such that one ACK DLLP can serve to acknowledge multiple TLPs. This collapsing of ACKs is the basis of the ACK DLLP policy choices. Less-frequent, more-collapsed ACKs have the least impact on TLP transmit bandwidth (*that is*, less-frequent ACKs result in less than 1 DLLP/TLP).

The ACK policy has two parts – a Timer and TLP Counter.

The default ACK Timer policy varies, according to the negotiated link width. [Table 8-1](#) charts an internal ACK Transmission Latency Timer (in symbol times), based upon differing link width values.

The ACK Transmission Latency Timer loads the appropriate value represented in [Table 8-1](#) when a TLP is received and known to be good (a few clocks after the END frame). The timer counts down each symbol time (every 4 ns). When the timer reaches 0, an ACK DLLP takes higher priority over new TLPs. This means that an ACK DLLP is transmitted before a new TLP is started, but also that the current TLP finished first. The ACK DLLP transmitted acknowledges all TLPs, up until the most-recently arrived good TLP.

The ACK TLP Counter policy sends a high-priority ACK after 12 TLPs have arrived, regardless of the ACK Timer policy.

If no TLP traffic is being transmitted (the Transmit link is idle), an ACK DLLP can be transmitted before the timer expires. This is an opportunistic low-priority ACK. When an ACK is transmitted, the timer resets, waiting for a new TLP to trigger.

The PEX 8509 allows a programmable override of the default Ack_Latency_Timer value, by programming the **ACK Transmission Latency Limit** register (offset 1F8h[11:0]). The value in this register is loaded when a new TLP arrives, and a high-priority ACK DLLP is attempted when the timer reaches 0. For fastest ACK response, this timer can be programmed to 0. This results in 1 DLLP ACK transmitted per each TLP received. For less impact on transmit TLP bandwidth, a larger value can be programmed, resulting in less-frequent ACKs.

In general, a slower ACK response does not impact the receive TLP stream, and aids the TLP transmit stream. In this case, the ACK becomes a performance bottleneck. The PEX 8509 can store up to 64 TLPs for Ports 0, 1, and 2 and up to 32 TLPs for Ports 3, 4, 5, 6, and 7, while waiting for an ACK.

Table 8-1. ACK Transmission Latency Timer Values

Link Width	x1	x2	x4
ACK Transmission Latency Timer (Symbol Times)	255	217	118

8.3.2 UpdateFC DLLP Policy

An UpdateFC DLLP is transmitted, in response to a received TLP, to allow the transmitter to transmit more TLPs of the same type. Each TLP that arrives consumes credit, and eventually a stream of TLPs consumes all credit, unless an updateFC DLLP allows more credit. The updateFC DLLP policy choice is the frequency of transmitting an updateFC.

There are two parts to the UpdateFC policy – credit amount and frequency of transmitting the updates. This section discusses only the frequency. Refer to [Section 8.4](#) for details regarding credit amounts.

The only reason to transmit an updateFC is to replenish credit. If the transmitter has sufficient credit, there is no reason to transmit more updateFC DLLPs. Therefore, the PEX 8509 looks at the total credit advertised when deciding when to transmit an updateFC DLLP. Four threshold options are allowed – 100%, 75%, 50%, and 25% (default). Whenever the remaining credit drops below the programmed threshold, an updateFC DLLP is given high priority (*that is*, the updateFC DLLP is transmitted before a new TLP is started). There is a separate threshold for Header and Payload credits for each TLP type (Posted, Non-Posted, and Completion) for each port.

[Table 8-2](#) illustrates the effect of the various updateFC DLLP policies for a series of 256-byte Payload TLPs (256 bytes = 16 Payload credits), as follows:

- **100% policy** – Results in a high-priority updateFC after every TLP, for a 1 DLLP per TLP ratio
- **75% policy** – Results in an updateFC every 2 TLPs
- **50% policy** – Results in an updateFC every 3 TLPs
- **25% policy** – Results in an updateFC every 4 TLPs

Once a high-priority updateFC is triggered, the running credit deficit is fully restored, if there are sufficient on-chip resources to do so. For most non-congested applications, there should be ample chip resources to fully restore credit with every updateFC. If resources are running low, however, then only part of the running credit is restored. If the threshold for transmitting an updateFC remains crossed, then, as more resources become available, a second high-priority DLLP can be transmitted.

[Table 8-2](#) charts how, for the various options, an updateFC is triggered. This example is for a traffic stream of back-to-back 256-byte Posted TLPs, where the maximum Posted Header credit is 10 and the maximum Posted Payload credit is 72. Recall that a 256-byte Payload requires 16 credits (1 credit per 16 bytes); therefore, each TLP in this case consumes 1 Header and 16 Payload credits.

Table 8-2. UpdateFC Options

Example	Running Credit Consumed/Total for Header and Payload	25% Remains Triggers at 2 Header / 18 Payload	50% Remains Triggers at 5 Header / 36 Payload	75% Remains Triggers at 7 Header / 54 Payload	100% Remains = ASAP
TLP0	9/10, 56/72	–	–	–	updateFC
TLP1	8/10, 40/72	–	–	updateFC	updateFC
TLP2	7/10, 24/72	–	updateFC	–	updateFC
TLP3	6/10, 6/72	updateFC	–	updateFC	updateFC

8.3.3 Unidirectional DLLP Policies

For unidirectional traffic, the PEX 8509 DLLP policies allow the most-frequent DLLPs, because the DLLPs do not interfere with the TLPs. Recall that DLLPs flow in the opposite direction of TLPs.

The PEX 8509 can transmit a DLLP ACK almost immediately upon receiving and verifying a TLP. A faster ACK results in fast transmitter de-allocation of the TLP, and can therefore allow a shallow TLP Replay buffer. The default values can be overwritten to increase or decrease the ACK DLLP rate. For unidirectional traffic, a small number (*such as* 1) is recommended. The ACK Latency Timer can be programmed in the **ACK Transmission Latency Limit** register (offset 1F8h[11:0]), where the value represents the number of symbol times before transmitting an ACK.

***Note:** The smallest ACK Transmission Latency Timer value programmed should be 1. Programming a value of 0 has unpredictable results.*

Similar to the ACK programmability, the PEX 8509 can immediately transmit an updateFC after receiving only the TLP Header. By transmitting an updateFC earlier, the total credit advertised can be minimized. For overly large Payloads (*such as* 1,024 bytes), the PEX 8509 resources should not be reserved more than necessary, unless actually needed. By programming fewer credits and having a fast updateFC policy, the system does not run out of credits and the PEX 8509 does not waste buffer space on reservations that do not arrive. Set the updateFC policy for unidirectional traffic to 100%. Set the credits to be sufficient to allow 3 to 4 TLPs.

8.4 Adjusting Ingress Resources

There is a finite amount of RAM storage available. The storage is used up by credit advertisement and TLP storage. The credit advertisement is programmable, although the default values should be optimal for most applications. The TLP storage is comprised of the number of TLPs that have arrived, but have not yet been forwarded to, and acknowledged by, the next PCI Express device.

RAM storage algorithms work in this way. As previously mentioned, RAM storage inside the PEX 8509 has a finite number of entries (1 TLP = 1 Entry). Additionally, the available number of entries in the RAM are shared by all available ports. Furthermore, each port can support three different types of traffic (Posted, Non-Posted, and Completions). For every port, TLP entries are used by all three types of traffic. These entries remain occupied and unavailable, until released. The TLP entries are released only after the receiver has acknowledged to the sender the arrival of the TLP. The TLP must arrive with no errors, per ACK/NAK policy. The number of entries assigned to each Traffic Class for every port is controlled with credit allocation. There is a trade-off between the number of entries that are *reserved* for a particular traffic class/port combination, perhaps more for one combination than another (initial credit), and the number of entries that are available on a first-come, first-served basis (general credit pool).

This analogy applies to the PEX 8509' RAM storage. The reservations are initial credit allocations for the three TLP types. The general tables are the common credit pool, and the time the table is occupied is the wait for ACK. Each of these topics is addressed in further detail in the sections that follow.

8.4.1 Initial Credit Allocation

The PEX 8509 default amount of advertised credit depends upon the strapped port width (not the negotiated port width). [Table 8-3](#) illustrates the initial credit for port configurations, and defines the advertised initFC DLLP values.

The values listed in [Table 8-3](#) can be changed before the initial advertisement by serial EEPROM or I²C. Take care to ensure that more credit is not advertised than there are available resources on the PEX 8509.

Per the *PCI Express Base r1.1*, the minimum initial credit must be sufficient to meet the credit requirement of a Maximum Payload Size. For the PEX 8509, the minimum initial Payload credit programmed must be at least eight more than the *PCI Express Base r1.1* minimum:

$$\text{Minimum Payload Credit} = 8 + \text{Payload Credit of Maximum Payload Size TLP}$$

Table 8-3. Initial Credit for Port Configurations

Strapped Port Width	Posted Header	Posted Payload	Non-Posted Header	Non-Posted Payload	Completion Header	Completion Payload
All	9	72	8	0 (Infinite)	8	72

8.4.2 Common Credit Pool

The PEX 8509 default credit values are optimal for most applications, to maintain back-to-back TLP traffic indefinitely, without running out of credit. After the initial credit is exhausted, more resources are automatically made available, to maintain the initial credit allotment. These additional resources are not reserved ahead of time, and can be used for any TLP type – Posted, Non-Posted, or Completion (P, NP, or Cpl, respectively). Because the Port and TLP type are not pre-specified, these extra resources are termed a *common credit pool*. There is a common credit pool for both Header and Payload credit.

The common Header pool is what remains after subtracting the advertised Header credits for each TLP type and for each port, from the total Header space of **256** credits.

The common Payload pool is what remains after subtracting the Payload credits for each type [Posted and Completion (CPLD)] for each port, from the total Payload space of **1,285** (five, six, and seven ports) or **1,159** (eight ports).

A larger common pool allows the most flexibility in handling all possible instantaneous traffic streams, without back-pressuring ingress flows. The default values of the initial credits for the different port configurations are selected to meet the criteria of sufficient credit to sustain uninterrupted traffic, while maintaining a large common pool.

8.4.3 Wait for ACK

Once a TLP arrives, it remains in the PEX 8509, until it is no longer required. The TLP can quickly egress the PEX 8509; however, until an ACK is received, stating that the TLP was correctly received, each TLP must remain in the PEX 8509 and be ready to be resent multiple times. While on the PEX 8509, the TLP continues to use the common pool resources.

The *PCI Express Base r1.1* recommends sending an ACK within the approximate time it takes to send 1.5 to 3 Maximum Payload Size TLPs. It does not suggest that smaller TLPs obtain faster ACKs. This data book describes how the PEX 8509 sends an ACK. However, the PEX 8509 has no way of knowing its link partner's ACK policy.

To minimize the amount of TLPs stored on the PEX 8509 while waiting for an ACK, follow these procedures:

- Avoid traffic patterns where a large amount of back-to-back TLP bytes go from a wide link to a single narrow link, because the narrow link can only forward the TLPs at a fraction of the ingress rate. *For example*, if a 4-KB MRd of is transmitted upstream from a x1 port and the upstream port is x4, the x4 port transmits a 4-KB CPLD four times faster than the Requester can receive the data. If the Requester transmits many of these MRd requests, overly large amounts of CPLD data that require storage on the PEX 8509 quickly accumulate.
- If there are many small TLPs, check whether the PEX 8509's ACK response time can be reduced from the *PCI Express Base r1.1* guidelines.
- Space the TLP pattern, rather than a burst of many back-to-back TLPs followed by a long stall, to even the distribution and accommodate a fixed ACK Transmission Latency Timer.


8.5 Latency

Latency is the length of time it takes to proceed from one event to another. There are numerous measures of latency. Perhaps the most common is the Start TLP-to-Start TLP (STP-to-STP) latency of TLPs, by way of the PEX 8509. Figure 8-5 illustrates the STP-to-STP latency for a x4, x2, and x1 port on the PEX 8509.

As expected with a Cut-Thru algorithm, latency is basically constant for all Payload sizes. The minimum latency for a x4 to x1 connection is 118 ns, x2 to x1 connection is 138 ns, and x1 to x1 connection is 188 ns.

The reason for the seemingly dramatic different latencies for different ingress port widths is mostly due to how quickly the TLP Header arrives to be decoded. The TLP Header, which accounts for 19 bytes, is received in 20 ns for a x4 link, 40 ns for a x2 link, and 76 ns for a x1 link.

Figure 8-5. x4, x2, and x1 Latency



8.5.1 Host-Centric Latency

Host-centric traffic flows only to or from the host. Host-centric latency depends upon the number of active streams. If there is only one active stream, or if the total host bandwidth is greater than or equal to the sum of all traffic streams, the latency graphs illustrated in [Figure 8-5](#) apply to well-balanced traffic.

However, if there is more traffic than an upstream host can sink, there can be congestion while the TLPs all try to use the limited host bandwidth. The latencies in that case depend upon the level of traffic congestion. In this case, host bandwidth is at 100%; however, the total downstream bandwidth is more than the host bandwidth, and latencies continue to increase until the congestion eases.

Another case of increased latency is if the host serially sends large chunks of traffic to one downstream port and then another downstream port. Suppose the upstream port is x2 and the downstream ports are all x1. It seems that there should not be a latency build-up because the bandwidth matches, and remains true, over a long period of time. However, unless the host can interleave the destinations, one destination must wait until the host completes transmitting traffic to the other destination, before it is allowed a turn. This can significantly increase round-trip Read latency.

For example, suppose that one downstream port transmits 16, 4-KB MRd requests upstream. Those Read requests represent 64 KB of data. If the upstream port is x2 and the downstream port is x1, then the Read Completions back up into the PEX 8509, perhaps all the way to the Root Complex. Suppose another downstream port has only one, 1-KB MRd request upstream. For many Root Complexes, this one Read request from the second device must wait for the 16 MRd requests from the first device to complete before being serviced. The PEX 8509 buffer is approximately **10 KB**; therefore, the second downstream device must wait for 64-10 KB to transmit across a x1 link before receiving its own Completion. 54 KB on x1 takes about 216 μ s, which significantly increases latency.

8.5.2 Peer-to-Peer Latency

Peer-to-peer latencies match the graphs illustrated in [Figure 8-5](#) when there is no congestion. The PEX 8509 has the same latency, regardless of whether the traffic is host-centric or peer-to-peer. Latency is constant in the non-congested case, no matter the source nor destination port, if the source port has the same or greater bandwidth than the destination port.

The discussion for host-centric traffic applies to all ports for peer-to-peer. It is recommended that there be a method outside the scope of this data book for balancing traffic flow for peer-to-peer applications.

8.5.3 Other Latency Measurements

In addition to STP-to-STP latency, there are other latencies to consider. [Table 8-4](#) illustrates various best-case latencies for a x1 link. Transmitted DLLPs can be required to wait for a TLP. DLLP policies can prevent sending a DLLP for a time period longer than the best case.

Table 8-4. Best Case Latencies for x1 Links

Latency Type	Latency (ns)
STP into updateFC Out	212
STP -> END into ACK Out	88
UpdateFC into STP Out	132

8.6 Queuing Options

On-chip queuing will not exist only in balanced bandwidth scenarios where the total ingress bandwidth is less than or equal to the egress bandwidth. In the common case, where the total ingress bandwidth is greater than the egress bandwidth, queues will develop on the PEX 8509. The PEX 8509 provides two alternatives for where to locate that queuing – Destination queue and Source queue.

8.6.1 Destination Queue

The default behavior is for all queues to develop at the destination port. If TLPs are arriving from four sources to a common destination port, the TLPs are scheduled according to First-In, First-Out (FIFO). If all four flows are equally active, the TLPs naturally interleave as:

1, 2, 3, 4, 1, 2, 3, 4

However, if a port is unusually busy while the others are not, the output might be similar to:

1, 1, 1, 1, 1, 1, 1, 1, 2, 3, 4, 1, 2, 3, 4

In this case, the other ports must wait for the earlier Port 1 traffic to complete before they can progress.

Note: For the queuing examples provided in this chapter, “Port 1” indicates “first port, not the port physically identified as Port 1.”

8.6.2 Source Queue

There may be applications that require deterministic bounded latency for a few ports, while the latency for other ports is not as important. For those applications, source queuing can be enabled.

Source queuing limits the Destination queue depth. When the Destination queue reaches the maximum depth, any subsequent TLPs targeting that port are not forwarded; rather they queue up in a per-Source-Port-based queue. The Source-Port queue does not forward TLPs until the Destination queue drops to a low watermark threshold, upon which TLP forwarding is re-enabled.

Note: *A Source Port queue, that cannot forward to a Destination queue, blocks all subsequent TLPs arriving in that same source port, although the target port is a different destination.*

The **Port Egress TLP Threshold** register (offset **F10h**) controls the maximum and minimum queue depths. (Refer to [Table 8-5](#).) The Upper TLP Count is the number of TLPs that can be queued in the Destination queue. The Lower TLP Count is the number of TLPs to which the destination queue must drop after reaching saturation, before re-enabling TLP forwarding.

In the example provided in [Section 8.6.1](#), with Port 1 turning On earlier than the other ports, by setting the destination Upper TLP Count to 3 and the Lower TLP Count to 1, the worst case that Port 2, 3, or 4 would be required to wait is for six Port 1 TLPs before being forwarded.

To avoid unnecessary idles on the destination link, the smallest value for the Upper TLP Count must be 2 for the PEX 8509. The corresponding value for the Lower TLP Count must be 1.

Table 8-5. Port Egress TLP Threshold Register Port Lower and Upper TLP Counts (Offset F10h)

Bit(s)	Name	Description
7:0	Port Lower TLP Count	When Source Scheduling is disabled due to threshold, Source Scheduling is re-enabled when the Port TLP Count goes below this threshold.
23:16	Port Upper TLP Count	When the Port TLP Count is greater than or equal to this value, the Source Scheduler disables TLP Scheduling to this egress port.

Note: *Bits not identified in [Table 8-5](#) are reserved.*

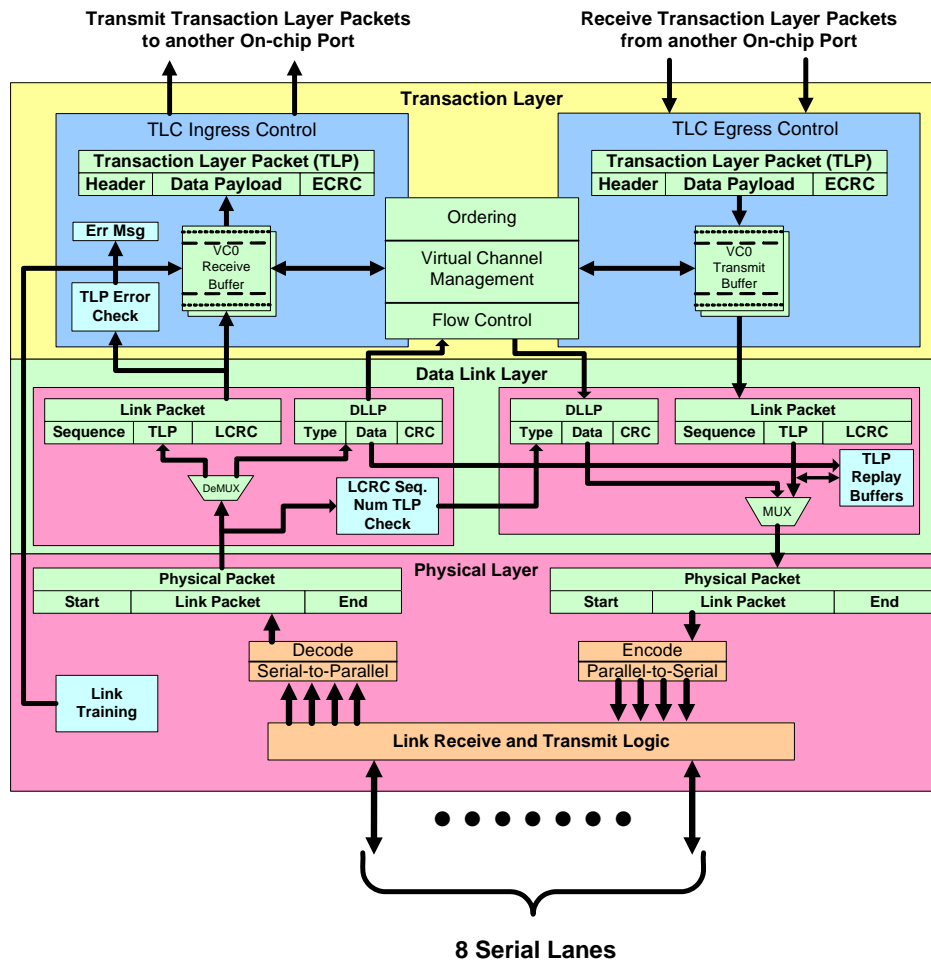
THIS PAGE INTENTIONALLY LEFT BLANK.

9.1 Data Flow Through

The PEX 8509 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 8509, in accordance with *PCI Express Base r1.1* requirements. The details of each layer are described in the following sections.

Figure 9-1. Data Formation and Flow Through



9.2 Physical Layer

The Physical Layer (PHY) is responsible for converting information received from the DLL into an appropriate serialized format and transmitting it across the PCI Express link. The PHY also receives the serialized input from the SerDes, converts it to parallel data (internal Data Bus), then writes it to the TLC Ingress buffer.

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

The PHY module interfaces to the PCI Express lanes and implements the PHY functions. The number of ports can vary from five to eight, with a cumulative lane bandwidth of x8. PHY functions include:

- SerDes modules, which provide all functions required by the *PCI Express Base r1.1*
- User-configurable port division
- x1, x2, x4 link widths supported
- Hardware link training and initialization
- Hardware detection of polarity reversal
- Hardware detection of lane reversal
- Data scrambling/de-scrambling and 8b/10b encode/decode
- Packet framing
- Loopback Master and Slave support
- Programmable test pattern with SKIP Ordered-Set insertion and return data checking
- Receiver error checking (packet framing, disparity, and symbol encoding)
- Run-on TLP check – Receive error flagged if number of beats received exceed the programmed MPS
- Checks and removes DLLP and TLP LCRC
- Link state Power Management – Supports L0, L0s, L1, L2/L3 Ready, and L3
- Supports cross-linked upstream port and downstream ports

9.2.1 PHY Status and Command Registers

The PHY operating conditions are defined in [Section 13.16.2, “Device-Specific Registers – Physical Layer.”](#) 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 PHY can include up to 8 integrated Serializer and De-serializer (SerDes) modules, which are distributed among two quads (0 and 1) 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 8509 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 8509 supports is configurable, as defined in [Table 9-1](#). Initial port configuration is determined by Strapped signal balls, serial EEPROM, or auto link-width negotiation.

Table 9-1. Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[1:0])	Lanes/SerDes/Port ^{a, c}							
	Port 0	Port 1	Port 2	Port 3	Port 4	Port 5	Port 6	Port 7
00b	x1 [0]	x1 [1]	x1 [2]	x1 [3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]
01b	x2 [0-1]	x1 [2]	x1 [3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]	_b
10b	x2 [0-1]	x2 [2-3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]	-	-
11b	x4 [0-3]	x1 [4]	x1 [5]	x1 [6]	x1 [7]	-	-	-

- a. The lanes are assigned to each enabled port in sequence, as indicated in [brackets].
- b. Configuration value and port combinations with “-” (no data) are **reserved**.
- c. Refer to [Table 9-2](#) for an explanation of the default Port/Physical Lane/SerDes Module to SerDes Quad relationship.

**Table 9-2. Port/Physical Lane/SerDes Module to SerDes Quad Relationship
(Port 0, Offset 224h[1:0]=0h)**

Port	Physical Lanes and SerDes Modules	SerDes Quad
0	0	0
1	1	
2	2	
3	3	
4	4	1
5	5	
6	6	
7	7	

9.3 Data Link Layer

The Data Link Layer (DLL) primary responsibilities include link management, DLLP decoding and generation, Sequence Number checking, and ACK/NAK generation. Upon detection of TLP error(s) by the PHY, the DLL is responsible for requesting TLP re-transmission until information is correctly received, or the link is determined to have failed.

9.3.1 Data Link Layer Packet

The Data Link Layer also generates and consumes packets used for Link management functions. To differentiate these packets from the TLPs used by the Transaction Layer, the term *Data Link Layer Packet (DLLP)* is used when referring to packets generated and consumed at the Data Link Layer. The rules governing the identification and formation of these packets are defined in the *PCI Express Base r1.1*, Section 3.4.1.

9.3.1.1 DLLP Ingress

The DLLP ingress module is responsible for receiving DLLPs from the Physical Layer, decoding the packet, and responding appropriately. If the DLLP is an FC Init/Update packet, the 32-bit data is written to the Egress Credit Handler. Other DLLPs are handled within the DLL logic. Certain DLLPs (*such as* Power Management) cause interaction with logic outside the DLL. A 16-bit CRC is performed on all received DLLPs. DLLPs failing CRC checking are ignored.

The DLLP ingress logic receives the sequence number for each incoming TLP from the Physical Layer and checks the sequence for correctness. The DLLP Ingress module sends requests for ACK/NAKs to the DLL egress block, where it is formatted and sent to the TL. If the DLL receives a NAK, the DLL transmits the Sequence Number and a NAK indicator to the TL egress module for re-transmission, then increments the Replay Count for that port.

9.3.1.2 DLLP Egress

The egress side of the DLL formats DLLPs for transmission by the TL Layer. These DLLPs include FC Init/Update packets from the INCH, ACK or NAK DLLPs and DLLPs from the Power Management module.

For each DLLP, the DLL egress generates a 16-bit CRC value that is included in the DLLP.

Normally, a TLP has higher priority than a ACK DLLP. The DLL has two ways to force a pending ACK to be sent at the next TLP boundary – the ACK Latency Timer and TLP Counter. If an ACK is waiting to be sent and the ACK Latency Timer reaches its maximum value, the pending ACK is given high priority for transmission. Another way to set the ACK to high priority is with the TLP Counter. If the pending ACK has waited while *N* number of TLPs are sent, the ACK's priority is set to High.

9.3.2 Packet Arbiter

The Packet Arbiter logic determines what type of packet to transmit, on a per-port basis, and arbitrates between the DLLPs and TLPs. The TLP request is a single signal that asserts when a TLP is scheduled for transmission. The DLLP request is a 4-bit bus that asserts the msb when a DLLP is scheduled for transmission. The lower three bits encode the type of DLLP, as defined in [Table 9-3](#).

The priority implemented by the Packet Arbiter is specified in the *PCI Express Base r1.1*, Section 3.5.2.1.

Table 9-3. DLLP Type Encoding

DLLP Encoding	DLLP Type
000b	NAK
001b	ACK
010b	Init FC
011b	FC Update
100b	Power Management
101b	Vendor-Specific
110b, 111b	<i>Reserved</i>

9.4 Transaction Layer

The Transaction Layer (TL) 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 9-4](#).

Table 9-4. 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)

Message space is added to PCI Express, and is used to support PCI Express sideband signals *such as* interrupts, Power Management requests, and so forth. 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
- Data Poisoning and end-to-end data integrity detection
- ECRC checking of incoming packets
- Error logging and reporting for incoming packets
- TLP packet dispatching
- Write control to the packet RAM and packet link list RAM
- Destination lookup and TC-VC mapping
- Shadow CSRs for BusNoCAM/IOCAM/AMCAM/TC-VC mapping
- Message Signaled Interrupt or INT_x generation
- Credit-based scheduling
- Pipelined full Split Transaction protocol
- PCI/PCI-X-compatible ordering
- Interrupt handling (INT_x or Message Signal Interrupt)
- Power Management support
- Hot Plug and PCI Express Hot Plug support
- Link State event support
- QoS support
- Ordering
- Ingress and Egress credit management

The hardware functions provided by the PEX 8509 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 Virtual Channel and Traffic Classes

The PEX 8509 supports one Virtual Channel (VC0) and eight Traffic Classes (TC[7:0]). VC0 and TC0 are required by the *PCI Express Base r1.1*, and configured at device start-up.

9.4.2 TL Transmit/Egress Protocol

The egress side TL receives TLP information from the internal fabric and makes a decision, based upon credit and ordering, regarding which TLP to send next from an Egress port.

The PEX 8509 implements the *PCI Express Base r1.1*-specified Flow Control (FC) protocol, which ensures that it does not transmit a TLP over a link to a remote receiver unless the receiving device contains sufficient Buffer space to accommodate the packet. This flow control 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.

9.4.2.1 Headers

The Headers contain three or four DWords that 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

9.4.2.2 Data Payloads

The Data Payloads are variable length with a maximum of 1,024 bytes, as defined by the *Maximum Payload Size* field (available sizes are 128, 256, 512, and 1,024, depending upon the number of ports used). Read requests do not include a Data Payload.

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

9.4.2.3 End-to-End Cyclic Redundancy Check

End-to-end Cyclic Redundancy Check (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 8509 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 8509 can optionally append ECRC to the end of internally generated TLPs, *such as* Interrupt and Error messages, if enabled.

9.4.3 TL Receive/Ingress Protocol

The ingress side TL collects and stores inbound TLP traffic in the packet RAM. The incoming data is checked for ECRC errors, valid type field, length matching the Header *Transfer Size* field, and other TLP-specific errors defined by the *PCI Express Base r1.1*.

Header and Data Payload information is forwarded to the Source Scheduler, to be routed across the internal fabric, to the Egress port. When CRC errors are detected, the packet is discarded.

9.4.4 Flow Control Protocol

The initial number of VC0 Flow Control (FC) credits are advertised as programmed in the Threshold registers, for each type of Header and Payload. The FC initial credits received are sent to the Egress buffer. After FC initialization is complete, the FC update credits received are transferred to the Egress buffer. For FC Credit updates, the Ingress buffer sends update requests to the DLL for DLLP for transmission, to increase the number of advertised credits in the PEX 8509.



Chapter 10 Interrupts

10.1 Interrupt Support

The PEX 8509 supports the PCI Express interrupt model, which uses two mechanisms:

- INTx emulation
- Message Signaled Interrupt (MSI)

For Conventional PCI compatibility, the PCI INTx 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 INTx emulation mechanism virtualizes PCI physical Interrupt signals by using an in-band signaling mechanism.

In addition to PCI INTx-compatible interrupt emulation, the PEX 8509 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
 - Presence Detect Changed (HP_PRSENTx# Input signal)
 - 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
 - Error-Correcting Code (ECC) error detected in the internal packet RAM
 - Internal Error FIFO overflow

The PEX 8509's external Interrupt ball, **PEX_INTA#**, indicates the assertion and/or de-assertion of the internally generated INTx signal:

- For Hot Plug or Link State triggered INTx events, PEX_INTA# assertion is controlled by the **ECC Error Check Disable** register *Enable PEX_INTA# Ball for Hot Plug or Link State Event* bit (Port 0, offset 1C8h[4]). When this bit is set, Hot Plug or Link State events trigger PEX_INTA# assertion; however, an INTx message is not generated in this case. PEX_INTA# assertion and INTx message generation for Hot Plug or Link State cases are mutually exclusive.
- PEX_INTA# assertion is controlled by the **ECC Error Check Disable** register *Enable PEX_INTA# Ball for Device-Specific Error* bit (Port 0, offset 1C8h[5]). When this bit is set, Device-Specific errors trigger PEX_INTA# assertion; however, PEX_INTA# assertion and INTx message generation are mutually exclusive.

10.1.1 Interrupt Handling

The PEX 8509 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 the interrupt, by way of the MSI mechanism
- Handling INT_x-type Interrupt messages from downstream devices

10.2 INT_x Emulation Support

The PEX 8509 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* **PCI Command** register *Interrupt Disable* and **PCI 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 8509 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 **PCI Status** register *Interrupt Status* bit is set. If INT_x interrupts are enabled (**PCI Command** register *Interrupt Disable* and **MSI Control** register *MSI Enable* bits, offsets 04h[10]=0 and 48h[16]=0, respectively), 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.

10.2.1 INTx-Type Interrupt Message Remapping and Collapsing

INTx-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 INTx message type received at the downstream port, based upon the downstream port's Device Number and Received INTx message Requester ID Device Number, and generates a new Interrupt message, according to the mapping defined in [Table 10-1](#).

A downstream Port transmits an Assert_INTA/Deassert_INTA message to the upstream port, due to a Hot Plug or Link State event or Device-Specific error.

Internally generated INTx messages always originate as type INTA messages, because the PEX 8509 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 10-1](#).

INTx 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_INTx and Deassert_INTx messages. The upstream port then forwards the new messages upstream, by way of its link.

Table 10-1. Downstream/Upstream Port INTx Interrupt Message Mapping

Device Number	At Downstream Port	By Upstream Port
0, 4	INTA	INTA
	INTB	INTB
	INTC	INTC
	INTD	INTD
1, 5	INTA	INTB
	INTB	INTC
	INTC	INTD
	INTD	INTA
2, 6	INTA	INTC
	INTB	INTD
	INTC	INTA
	INTD	INTB
3, 7	INTA	INTD
	INTB	INTA
	INTC	INTB
	INTD	INTC

10.3 Message Signaled Interrupt Support

One of the interrupt schemes supported by the PEX 8509 is the Message Signaled Interrupt (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.*

10.3.1 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 8509 supports two messages for MSI – one for Hot Plug or Link State events, the other for Device-Specific error events. Therefore, the **MSI Control** register *Multiple Message Capable* field (offset 48h[19:17]) is always set to 001b. When the register's *Multiple Message Enable* field (offset 48h[22:20]) is cleared to 000b (default), the PEX 8509 can generate only one message for Hot Plug or Link State events and Device-Specific error events. When a non-zero value is written to the *Multiple Message Enable* field, two-message support is enabled.

The **MSI 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 **MSI Address** register (offset 4Ch) contents. Data written is the contents of the **MSI Data** register (offset 54h) lower two bytes and zeros (0) in the upper two bytes.

When the Hot Plug or Link State event or Device-Specific error that caused the interrupt is serviced, the device can generate a new MSI Memory Write as a result of new events. Because MSI is an edge-triggered event, two bits are provided for masking the events [**MSI Mask** register *MSI Mask for Device-Specific Interrupts* and *MSI Mask for Hot Plug or Link State Interrupts* bits (offset 58h[1:0]), respectively]. A new MSI can be generated only after the Mask bits are serviced. System software should mask the Mask bits when the MSI event is being processed.

10.3.2 MSI Capability Registers

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

10.4 PEX_INTA# Interrupts

PEX_INTA# Interrupt output is enabled when the following conditions exist:

- INTA messages are enabled (**PCI Command** register *Interrupt Disable* bit, offset 04h[10]=0) and MSI is disabled (**MSI Control** register *MSI Enable* bit, offset 48h[16]=0)
- PEX_INTA# output is enabled [**ECC Error Check Disable** register *Enable PEX_INTA# Ball for Device-Specific Error* and/or *Enable PEX_INTA# Ball for Hot Plug or Link State Event* bits (Port 0, offset 1C8h[5 and/or 4], respectively) are set to 1]

Note: PEX_INTA# assertion and INTx messaging are mutually exclusive for a given interrupt event. When MSI is enabled (offset 48h[16]=1), both PEX_INTA# and INTx are disabled for PEX 8509 internally generated interrupts. The forwarding of external INTx messages received from a downstream port to the upstream port is always enabled.



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 8509 includes one Hot Plug Controller per Hot Plug-capable downstream port (Ports 1, 2, and 3).

Note: Ports 0, 4, 5, 6, and 7 do not have a Hot Plug Controller because these ports are not Hot Plug-capable.

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 Device PM state)
- Presence detect is accomplished through either an in-band SerDes receiver detect mechanism or by using the HP_PRSNTx# signal
- Hot Plug interrupts can be sent in-band using INTx or MSI messages, or signaled externally using PEX_INTA#

11.1.2 Hot Plug Port External Signals

The signals for Hot Plug support are defined in [Section 3.4.2, “Hot Plug Signals.”](#)

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 11-1](#).

Table 11-1. Hot Plug Outputs for Disabled Hot Plug Slot

Output Signal	Logic	Comments
HP_ATNLEDx#	High	Attention LED is turned Off
HP_CLKENx#	High	Reference Clock is not driven to the slot
HP_PERSTx#	Low	Slot remains in reset
HP_PWRENx#	High	Power Controller is turned Off
HP_PWRLEDx#	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 Capability** (offset 7Ch)
- **Slot Status and Control** (offset 80h)

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 INTx or MSI messages, or by the side-band PEX_INTA# output. Only one interrupt mechanism can be selected, and all Hot Plug ports must use the same mechanism.

INTx interrupts are enabled if:

- INTx messages are enabled (**PCI Command** register *Interrupt Disable* bit, offset 04h[10]=0) and MSI is disabled (**MSI Control** register *MSI Enable* bit, offset 48h[16]=0)
- PEX_INTA# output is disabled [**ECC Error Check Disable** register *Enable PEX_INTA# Ball for Hot Plug or Link State Event* bit is cleared (Port 0, offset 1C8h[4]=0)]

PEX_INTA# interrupts are enabled if:

- INTx messages are enabled (**PCI Command** register *Interrupt Disable* bit, offset 04h[10]=0) and MSI is disabled (**MSI Control** register *MSI Enable* bit, offset 48h[16]=0)
- PEX_INTA# output is enabled [**ECC Error Check Disable** register *Enable PEX_INTA# Ball for Hot Plug or Link State Event* bit is set (Port 0, offset 1C8h[4]=1)]

MSI interrupts are enabled if:

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

Depending upon the downstream Port 1, 2, and/or 3 Device PM state, a Hot Plug event can generate a system interrupt or PME. When a PEX 8509 Hot Plug-capable downstream port is in the D0 Device PM state, Hot Plug events generate a system interrupt; when not in the D0 Device PM 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 Slot Power-Up/Down Sequence

If a Hot Plug-capable downstream port (Port 1, 2, and/or 3) 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 Port 1, 2, and/or 3 is connected to a slot, its associated Hot Plug Controller can power up that slot, with or without an external serial EEPROM. Hot Plug Controller sequencing is determined by the states of the following bits:

- **Slot Capability** register *Power Controller Present* bit (offset 7Ch[1])
- **Slot Capability** 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 λ # 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 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 11-2](#). The Hot Plug Controller outputs remain in the default state described in [Table 11-1](#), 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 Capability** 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 11-2. Configuring Power-Up Sequence Features with Serial EEPROM

Register Bit	Hot Plug Controller and Hot Plug Output Signal Configurable Features
<i>Power Controller Present</i> (Slot Capability register, offset 7Ch[1])	The <i>Power Controller Present</i> bit enables or disables the Hot Plug Controller on the PEX 8509 Hot Plug-capable downstream ports. If the <i>Power Controller Present</i> 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 11-1 . If the <i>Power Controller Present</i> bit is enabled (set to 1), the Hot Plug Controller powers up the slot when the MRL is closed and the Slot Control register <i>Power Controller Control</i> bit is cleared (offset 80h[10]=0). Otherwise, if the <i>MRL Sensor Present</i> bit is disabled (cleared to 0), the MRL's position has no effect on powering up the slot.
<i>MRL Sensor Present</i> (Slot Capability register, offset 7Ch[2])	When enabled (set to 1), the PEX 8509 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.
<i>HPC T_{pepv} Delay</i> (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
<i>HPC T_{pperl} Delay</i> (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 # are de-asserted High. (Refer to Section 11.4.1.2 .) 0 = 20 ms 1 = 100 ms (default)
<i>Attention Indicator Present</i> (Slot Capability 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.
<i>Power Indicator Present</i> (Slot Capability 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 Capability** register *Slot Implemented* bit is set (offset 68h[24]=1), and the PEX 8509 is in the following state:

1. Hot Plug Controller is enabled for Ports 1, 2, and/or 3.
2. Slots 1, 2, and/or 3 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 8509 exits Reset (**PEX_PERST#** input goes High), if the downstream Port 1, 2, and/or 3 *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 11-1](#).

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 and serial EEPROM initialization. The serial EEPROM initialization delay is determined by the following:

- Serial EEPROM clock (**EE_SK**) frequency, programmable through the **Serial EEPROM Clock Frequency** register *EepFreq[2:0]* field (Port 0, offset 268h[2:0])
- Number of registers that are programmed to be initialized by the serial EEPROM

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 to turn On the Reference Clock (PEX_REFCLK_{n/p}) to the slot. The T_{pepv} time delay is specified by setting the **Power Management Hot Plug User Configuration** register *HPC T_{pepv} 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 T_{pvperl} Delay* bit (offset 1E0h[6]). By default, this bit is set to 1, indicating a 100-ms delay.

Consideration should be given to the combination of the serial EEPROM clock (**EE_SK**) frequency [programmable in the **Serial EEPROM Clock Frequency** register *EepFreq[2:0]* field (Port 0, offset 268h[2:0])], along with the number of registers to be initialized by serial EEPROM, as well as any delay for cascaded resets through multiple devices, and allow sufficient margin for devices to be ready for enumeration by the Host.


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 using software control.

If HP_MRLx# is enabled but not asserted to power-up the slot immediately after reset, HP_MRLx# 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_MRLx# 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_MRLx# and HP_BUTTONx# 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 I²C interface clears the *Power Indicator Present* bit (offset 7Ch[4]) to 0.

11.4.1.3 HP_PERSTx# (Reset) and HP_PWRLEDx# Output Power-Up Sequencing when *Power Controller Present* Bit Is Clear


The HP_PERSTx# and HP_PWRLEDx# outputs can be used without enabling the Hot Plug Power Controller (HP_PWRENx# and HP_CLKENx# outputs and HP_PWRFLTx# input). For example, HP_PERSTx# 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_PERSTx# is de-asserted (High) and HP_PWRLEDx# is asserted (Low), after the Root Complex PERST# input is de-asserted, as illustrated in Figure 11-2. However, HP_PWRLEDx# 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_PERSTx# is asserted and HP_PWRLEDx# is not asserted unless HP_MRLx# is Low. Therefore, if the HP_PERSTx# and/or HP_PWRLEDx# outputs are used [and a Manually operated Retention Latch (MRL) is *not* used], pull HP_MRLx# Low, to allow the outputs to toggle, regardless of whether the serial EEPROM is blank.

HP_PERSTx# can also be toggled at runtime by toggling the *Power Controller Control* bit, provided the *Power Controller Present* bit is set (offset 7Ch[1]=1), and that either the *Power Controller Present* bit is cleared (offset 7Ch[1]=0) or HP_PERSTx# is initially de-asserted during slot power-up sequencing, as described in Section 11.4.1.2. A value of 1 asserts HP_PERSTx# (Low). A value of 0 de-asserts HP_PERSTx# (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 I²C 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](#).

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 automatically powers down the slot if the MRL is open. [Figure 11-3](#) 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 11-3. Hot Plug Automatic Power-Down Sequence



11.5 Hot Plug Board Insertion and Removal Process

Table 11-3 defines the board insertion procedure supported by the PEX 8509. Table 11-4 defines the board removal procedure.

Table 11-3. Hot Plug Board Insertion Process

Operator / Action	Hot Plug Controller	Software
A. Places board in slot.	<ol style="list-style-type: none"> 1. Sets <i>Presence Detect State</i> bit to 1. 2. Sets <i>Presence Detect Changed</i> bit to 1. 3. Generates Interrupt message due to Presence Detect change, if enabled. 	Clears <i>Presence Detect Changed</i> bit to 0.
	<ol style="list-style-type: none"> 4. Transmits Interrupt de-assertion message, if enabled. 	
B. Locks MRL.	<ol style="list-style-type: none"> 5. Clears <i>MRL Sensor State</i> bit to 0. 6. Sets <i>MRL Sensor Changed</i> bit to 1. 7. Generates Interrupt message due to MRL Sensor state change, if enabled. 	Clears <i>MRL Sensor Changed</i> bit to 0.
	<ol style="list-style-type: none"> 8. Transmits Interrupt de-assertion message, if enabled. 	
C. Presses Attention Button.	<ol style="list-style-type: none"> 9. Sets <i>Attention Button Pressed</i> bit to 1. 10. Generates Interrupt message due to Attention Button Pressed event, if enabled. 	Clears <i>Attention Button Pressed</i> bit to 0.
	<ol style="list-style-type: none"> 11. Transmits Interrupt de-assertion message, if enabled. 	
Continued ...		

Table 11-3. Hot Plug Board Insertion Process (Cont.)

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

Table 11-4. Hot Plug Board Removal Process

Operator / Action	Hot Plug Controller	Software
A. Presses Attention Button.	<ol style="list-style-type: none"> 1. Sets <i>Attention Button Pressed</i> bit to 1. 2. Generates Interrupt message due to Attention Button pressed, if enabled. 	Clears <i>Attention Button Pressed</i> bit to 0.
	<ol style="list-style-type: none"> 3. Transmits Interrupt de-assertion message, if enabled. 	Writes to the Slot Control register <i>Power Indicator Control</i> field, to blink the Power Indicator LED, which indicates that the board is being powered down.
B. Power Indicator blinks.	<ol style="list-style-type: none"> 4. Sets <i>Power Indicator Control</i> field to 10b. 5. Sets <i>Command Completed</i> bit to 1. 6. Generates Interrupt message due to Power Indicator Blink command completion, if enabled. 	Clears <i>Command Completed</i> bit to 0.
	<ol style="list-style-type: none"> 7. Transmits Interrupt de-assertion message, if enabled. 	Sets Slot Control register <i>Power Controller Control</i> bit to 1, to turn Off power to the port.
C. Power Indicator Off.	<ol style="list-style-type: none"> 8. Slot is powered Off. 9. After a 256-ms delay, sets the <i>Command Completed</i> bit to 1. 10. Generates Interrupt message due to Power Turn Off command completion, if enabled. 	Clears <i>Command Completed</i> bit to 0. Clears <i>Power Indicator Control</i> 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.	<ol style="list-style-type: none"> 11. Clears <i>Power Indicator Control</i> field to 00b. 12. Sets <i>Command Completed</i> bit to 1, due to Power Indicator Off command completion. 	Clears <i>Command Completed</i> bit to 0.
	<ol style="list-style-type: none"> 13. Transmits Interrupt de-assertion message, if enabled. 	
E. Unlocks MRL.	<ol style="list-style-type: none"> 14. Sets <i>MRL Sensor State</i> bit to 1. 15. Sets <i>MRL Sensor Changed</i> bit to 1. 16. Generates Interrupt message due to MRL Sensor state change, if enabled. 	Clears <i>MRL Sensor Changed</i> bit to 0.
	<ol style="list-style-type: none"> 17. Transmits Interrupt de-assertion message, if enabled. 	
F. Removes board from slot.	<ol style="list-style-type: none"> 18. Clears <i>Presence Detect State</i> bit to 0. 19. Sets <i>Presence Detect Changed</i> bit to 1. 20. Generates Interrupt message due to Presence Detect change, if enabled. 	Clears <i>Presence Detect Changed</i> bit to 0.
	<ol style="list-style-type: none"> 21. Transmits Interrupt de-assertion message, if enabled. 	

THIS PAGE INTENTIONALLY LEFT BLANK.

12.1 Overview

The PEX 8509 Power Management features provide the following services:

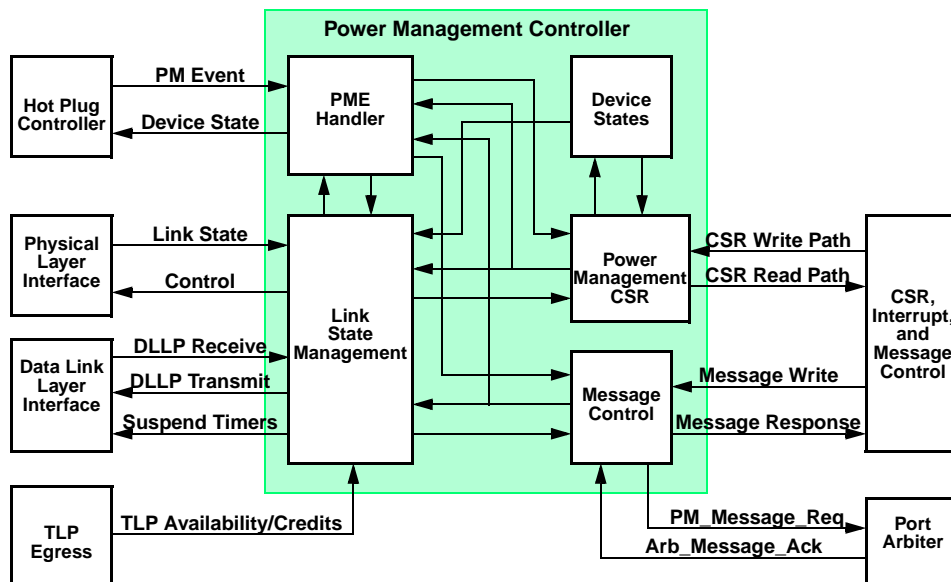
- Mechanisms to identify power management capabilities
- Ability to transition into certain power management states
- Notification of the current power management state of each port
- Support for the option to wakeup the system upon a specific event

The PEX 8509 supports hardware-autonomous power management and software-driven D-State power management. The switch also supports the L0s and L1 Link PM states in hardware-autonomous Active State Power Management (ASPM), as well as the L1, L2/L3 Ready, and L3 Link PM states in Conventional PCI-compatible Power Management. D0, D3hot, and D3cold Device PM states are supported in Conventional PCI-compatible Power Management. Because the PEX 8509 does not support Vaux, PME generation from the D3cold Device PM state is **not supported**.

The Power Management module interfaces with a Physical Layer electrical sub-block, to transition the Link state into a low-power state, when the module receives a Power State Change request from a downstream component, or an internal event forces the Link state entry into low-power states in hardware-autonomous ASPM mode. PCI Express Link states are not directly visible to Conventional PCI Bus driver software; however, they are derived from the Power Management state of the components residing on those links.

Figure 12-1 illustrates the PEX 8509 Power Management Controller functional block diagram.

Figure 12-1. Power Management Controller Functional Block Diagram



Note: The PEX 8509 Standard Hot Plug Controllers are available on Ports 1, 2, and 3, when configured as downstream ports.

12.2 Features

- Compliant to the following specifications:
 - *PCI Express Base r1.1*
 - *PCI Power Mgmt. r1.2*
- Link Power Management State (L-States)
 - Conventional PCI-compatible Link Power Management – L1, L2/L3Ready and L3 (with Vaux **not supported**)
 - Active State Power Management (ASPM) – L0s and L1
- Conventional PCI-compatible Device Power Management States (D-States) – D0 (D0_uninitialized and D0_active) and D3hot
- Power Management Event (PME) support from D3hot Device PM state
- PME for Hot Plug events
- Forwarding of PME_Turn_Off broadcast messages
- *PCI Express Base r1.1*-specific Control and Status registers and interrupts

12.3 Power Management Capability

12.3.1 Device Power Management States

The PEX 8509 supports the PCI Express PCI-PM D0 and D3hot Device PM states. The D1 and D2 Device PM states, which are optional in the *PCI Express Base r1.1*, are **not supported** by the PEX 8509.

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

12.3.1.1 D0 Device Power Management State

The D0 Device PM state 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 Device PM state. The component remains in the D0_uninitialized Device PM state until the serial EEPROM load and initial link training completes.

A device enters the D0_active Device PM state when:

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

12.3.1.2 D3hot Device Power Management State

Once in the D3hot Device PM state, the PEX 8509 can later be transitioned into the D3cold Device PM state, by removing power from its host component. Functions that are in the D3hot Device PM state can be transitioned, by software, to the D0_uninitialized Device PM state. When in the D3hot Device PM state, Link State operations cause a PME in the PEX 8509.

Only Type 0 Configuration accesses are allowed in the D3hot Device PM 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 8509 port in the D3hot Device PM state are terminated as Unsupported Requests (UR). Type 0 Configuration transactions complete successfully. When the PEX 8509 upstream port is programmed to the D3hot Device PM state, the port initiates Conventional PCI-PM L1 Link PM state entry.

12.3.2 Link Power Management States

PEX 8509 components hold their upstream and downstream links in the L0 Link PM state when they are in the standard operational state (Conventional PCI-PM state is in the D0_active Device PM state). ASPM defines a mechanism for components in the D0 Device PM state, to reduce link power by placing their links into a low-power state and instructs the other end of the link to do likewise. This allows hardware-autonomous, dynamic link power reduction beyond what is achievable by software-only-controlled power management. Table 12-1 defines the relationships between the Power state of a component and its upstream link. Table 12-2 defines the relationships between Link Power Management states and power saving actions.

Conventional PCI Power Management, and the L1 and L2/L3 Ready Link PM states are controlled by system software programming the PEX 8509 into the D3hot Device PM state, and subsequently causing the Root Complex to broadcast the PME_Turn_Off message to the downstream hierarchy.

Table 12-1. Relationships between Component Power State and Upstream Link

Downstream Component Device PM State	Permissible Upstream Component Device PM State	Permissible Interconnect Link PM State
D0	D0	L0, L0s, L1 (optional) – Active State Power Management (ASPM).
D3hot	D0 to D3hot	L1, L2/L3 Ready.
D3cold (no Vaux)	D0 to D3cold	L3 (off). Zero power.

Table 12-2. Relationships between Link PM States and Power-Saving Actions

Link PM State	Power-Saving Actions
Tx L0s	PHY Tx Lanes are in a high-impedance state.
Rx L0s	PHY Rx Lanes in a low-power state.
L1	PHY Tx and Rx Lanes are in a low-power state. FC timers are suspended.
L2/L3 Ready	PHY Tx and Rx Lanes are in a low-power state. FC timers are suspended.
L3 (D3cold)	Component is fully powered Off.

12.3.3 PCI Express Power Management Support

The PEX 8509 supports PCI Express features that are required or important for PCI Express switch Power Management. Table 12-3 lists supported and non-supported features and the register bits/fields used for configuration or activation.

Table 12-3. Supported PCI Express Power Management Capabilities

Register		Description	Supported	
Offset	Bit(s)		Yes	No
40h	Power Management Capability (All Ports)			
	7:0	Capability ID Set to 01h, indicating that the Capability structure is the PCI Power Management Capability structure.	✓	
	15:8	Next Capability Pointer Default 48h points to the Message Signaled Interrupt Capability structure.	✓	
	18:16	Version Default 011b indicates compliance with the <i>PCI Power Mgmt. r1.2</i> .	✓	
	19	PME Clock Does not apply to PCI Express. Returns 0.		✓
	21	Device-Specific Initialization Default 0 indicates that Device-Specific Initialization is <i>not</i> required.	✓	
	24:22	AUX Current The PEX 8509 does <i>not support</i> PME generation from D3cold; therefore, the serial EEPROM value for this field should be 000b.		✓
	25	D1 Support Default value of 0 indicates that the PEX 8509 does <i>not support</i> the D1 Device PM state.		✓
	26	D2 Support Default value of 0 indicates that the PEX 8509 does <i>not support</i> the D2 Device PM state.		✓
	31:27	PME Support Bits [31, 30, and 27] must be set to 1, to indicate that the PEX 8509 will forward PME messages, as required by the <i>PCI Express Base r1.1</i> .	✓	

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
Power Management Status and Control (All Ports)				
44h	1:0	Power State Used to determine the current Device PM state of the port, and to set the port into a new Device PM state. 00b = D0 01b = D1 – <i>Not supported</i> 10b = D2 – <i>Not supported</i> 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 8509 port ^a 1 = Enables PME generation by the corresponding PEX 8509 port	✓	
	12:9	Data Select Initially writable by serial EEPROM and I ² C only ^b . After a Serial EEPROM or I ² C Write occurs to this register, RW for all CSR accesses. 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	✓	
		RO for hardware auto-configuration.		✓
	14:13	Data Scale Writable by serial EEPROM and I ² C only ^b . There are four internal Data Scale registers (one each per Data register – 0, 3, 4 and 7), per port. Bits [12:9], <i>Data Select</i> , select the Data Scale register.	✓	
	15	PME Status 0 = PME is not generated by the corresponding PEX 8509 port ^{aa} 1 = PME is being generated by the corresponding PEX 8509 port	✓	

- a. Because the PEX 8509 does not consume auxiliary power, this bit is not sticky, and is always cleared to 0 at power-on reset.
- b. With no serial EEPROM nor previous I²C programming, Reads return 00h for the **Data Scale** and **Data** registers (for all Data Selects).

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
44h	Power Management Control/Status Bridge Extensions (All Ports)			
	22	B2/B3 Support <i>Reserved</i> Cleared to 0, as required by the <i>PCI Power Mgmt. r1.2</i> .		✓
	23	Bus Power/Clock Control Enable <i>Reserved</i> Cleared to 0, as required by the <i>PCI Power Mgmt. r1.2</i> .		✓
	Power Management Data (All Ports)			
	31:24	Data Writable by serial EEPROM and I ² C only ^b . There are four internal Data registers per port. Bits [12:9], <i>Data Select</i> , select the Data register.	✓	
6Ch	Device Capability (All Ports)			
	8:6	Endpoint L0s Acceptable Latency Because the PEX 8509 is a switch and not an endpoint, the PEX 8509 does <i>not support</i> this feature. 000b = Disables the capability		✓
	11:9	Endpoint L1 Acceptable Latency Because the PEX 8509 is a switch and not an endpoint, the PEX 8509 does <i>not support</i> this feature. 000b = Disables the capability		✓
	25:18	Captured Slot Power Limit Value (Upstream Port) For the PEX 8509 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 <i>Captured Slot Power Limit Scale</i> field. Do not change for the downstream ports.	✓	
	27:26	Captured Slot Power Limit Scale (Upstream Port) For the PEX 8509 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 <i>Captured Slot Power Limit Value</i> field. 00b = 1.0 01b = 0.1 10b = 0.01 11b = 0.001 Do not change for the downstream ports.	✓	

b. With no serial EEPROM nor previous I²C programming, Reads return 00h for the **Data Scale** and **Data** registers (for all *Data Selects*).

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
70h	Device Control (All Ports)			
	10	AUX Power PM Enable Cleared to 0.		✓
	Device Status (All Ports)			
	20	AUX Power Detected Cleared to 0.		✓
74h	Link Capability (All Ports)			
	11:10	Active State Power Management (ASPM) Support Indicates the level of ASPM supported by the port. 01b = L0s Link PM state entry is supported 11b = L0s and L1 Link PM states are supported All other encodings are <i>reserved</i> .	✓	
	14:12	L0s Exit Latency Indicates the L0s Link PM state exit latency for the given PCI Express link. Value depends upon the Physical Layer Command and Status register N_FTS Value field (offset 220h[15:8]) value. 101b = Corresponding PEX 8509 port L0s Link PM state Exit Latency is 1 μ s to less than 2 μ s	✓	
	17:15	L1 Exit Latency Indicates the L1 Link PM state exit latency for the given PCI Express link. The value reported indicates the length of time that the corresponding PEX 8509 port requires to complete the transition from the L1 to L0 Link PM state. 101b = Corresponding PEX 8509 port L1 Link PM state Exit Latency is 16 μ s to less than 32 μ s	✓	
	18	Clock Power Management	✓	
78h	Link Control (All Ports)			
	1:0	Active State Power Management (ASPM) 00b = Disables L0s and L1 Link PM state Entries for the corresponding PEX 8509 port ^b 01b = Enables only L0s Link PM state Entry 10b = Enables only L1 Link PM state Entry 11b = Enables both L0s and L1 Link PM state Entries	✓	
	8	Clock Power Management Enable The PEX 8509 does <i>not support</i> removal of the Reference Clock in the L1 and L2/L3 Ready Link PM states.		✓

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

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
7Ch	Slot Capability (All Downstream Ports; Upstream Port Always Reads 0)			
	0	Attention Button Present 0 = Attention Button is not implemented 1 = Attention Button is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port <i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.	✓	
	1	Power Controller Present 0 = Power Controller is not implemented 1 = Power Controller is implemented for the slot of the corresponding PEX 8509 Hot Plug-capable downstream port <i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.	✓	
	2	MRL Sensor Present 0 = MRL Sensor is not implemented 1 = MRL Sensor is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port <i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.	✓	
	3	Attention Indicator Present 0 = Attention Indicator is not implemented 1 = Attention Indicator is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port <i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.	✓	

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
80h	Slot Control (All Ports)			
	1	<p>Power Fault Detector Enable</p> <p>0 = Function is disabled</p> <p>1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Device PM state (Power Management Status and Control register <i>Power State</i> field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Device PM state (offset 44h[1:0]=11b), for a Power Fault Detected event on the corresponding PEX 8509 Hot Plug-capable downstream port</p> <p><i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.</p>	✓	
	9:8	<p>Power Indicator Control</p> <p>Controls the Power Indicator on the corresponding PEX 8509 downstream port slot.</p> <p>00b = <i>Reserved</i> – Writes are ignored</p> <p>01b = Turns On indicator to constant On state</p> <p>10b = Causes indicator to blink</p> <p>11b = Turns Off indicator</p> <p>Software must use a Byte or Word Write (and not a DWord Write) to control the HP_PWRLEDx# Output signal.</p> <p>Reads return the corresponding PEX 8509 Hot Plug-capable downstream port Power Indicator's current state.</p> <p><i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.</p>	✓	
	10	<p>Power Controller Control</p> <p>Controls the Power Controller on the corresponding PEX 8509 Hot Plug-capable downstream port slot.</p> <p>0 = Turns On the Power Controller; requires some delay to be effective</p> <p>1 = Turns Off the Power Controller</p> <p>Software must use a Byte or Word Write (and not a DWord Write) to control the Power Controller Output signals.</p> <p><i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.</p>	✓	
	Slot Status (All Ports)			
17	<p>Power Fault Detected</p> <p>Set to 1 when the Power Controller of the corresponding PEX 8509 Hot Plug-capable downstream port slot detects a Power Fault at the slot.</p> <p><i>Reserved</i> for the upstream port and non-Hot Plug-capable downstream ports.</p>	✓		

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
Power Budget Extended Capability Header (All Ports)				
138h	15:0	PCI Express Extended Capability ID Set to 0004h, as required by the <i>PCI Express Base r1.1</i> .	✓	
	19:16	Capability Version Set to 1h, as required by the <i>PCI Express Base r1.1</i> .	✓	
	31:20	Next Capability Offset Set to 148h, which addresses the PEX 8509 Virtual Channel Extended Capability structure.	✓	
Data Select (All Ports)				
13Ch	7:0	Data Select Indexes the Power Budget Data reported, by way of eight Power Budget Data registers, and selects the DWord of Power Budget data that appears in each Power Budget Data register. Index values start at 0, to select the first DWord of Power Budget data; subsequent DWords of Power Budget data are selected by increasing index values 1 to 7.	✓	

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
140h	Power Budget Data (All Ports)			
	7:0	<p>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, to produce the actual power consumption value.</p>	✓	
	9:8	<p>Data Scale Specifies the scale to apply to the Base Power value. The device power consumption is determined by multiplying the <i>Base Power</i> field contents with the value corresponding to the encoding returned by this field.</p> <p>00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x</p>	✓	
	12:10	<p>PM Sub-State 000b = Corresponding PEX 8509 port is in the default Power Management sub-state</p>	✓	
	14:13	<p>PM State Current Device Power Management (PM) state.</p> <p>00b = D0 Device PM state 11b = D3 Device PM state All other encodings are <i>reserved</i>.</p>	✓	
	17:15	<p>Type Type of operating condition.</p> <p>000b = PME Auxiliary 001b = Auxiliary 010b = Idle 011b = Sustained 111b = Maximum All other encodings are <i>reserved</i>.</p>	✓	
	20:18	<p>Power Rail Power Rail of operating condition.</p> <p>000b = Power 12V 001b = Power 3.3V 010b = Power 1.8V 111b = Thermal All other encodings are <i>reserved</i>.</p>	✓	
<p>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]) contents.</p>				

Table 12-3. Supported PCI Express Power Management Capabilities (Cont.)

Register		Description	Supported	
Offset	Bit(s)		Yes	No
144h	Power Budget Capability (All Ports)			
	0	System Allocated 1 = Power budget for the device is included within the system power budget	✓	
1E0h	Power Management Hot Plug User Configuration (All Ports)^c			
	0	L0s Entry Idle Count Time to meet to enter the L0s Link PM state. 0 = Idle condition lasts for 1 μ s 1 = Idle condition lasts for 4 μ s	✓	
	1	L1 Upstream Port Receiver Idle Count For active L1 Link PM state entry. 0 = Upstream port receiver remains idle for 2 μ s 1 = Upstream port receiver remains idle for 3 μ s	✓	
	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 T_{pepv} Delay Slot power-applied to power-valid delay time. 00b = 16 ms (default) 01b = 32 ms 10b = 64 ms 11b = 128 ms	✓	
	6	HPC T_{ppervl} Delay Downstream port power-valid to reset signal release time. 0 = 20 ms 1 = 100 ms (default)	✓	
	7	Disable PCI PM L1 Entry 0 = Enables upstream port entry into the L1 Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field (offset 44h[1:0]) is set to 11b] 1 = Disables upstream port entry into the L1 Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field (offset 44h[1:0]) is set to 11b]	✓	
	10	L0s Entry Disable 0 = Enables upstream port entry into the L0s Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field (offset 44h[1:0]) is set to 11b] 1 = Disables upstream port entry into the L0s Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field (offset 44h[1:0]) is set to 11b]	✓	

c. Hot Plug-related bits/fields are valid only for Hot Plug-capable downstream Ports 1, 2, and 3. Ports 0, 4, 5, 6, and 7 are reserved for Hot Plug-related bits/fields.

12.4 Power Management Tracking

Upstream port logic tracks the link status of each downstream and upstream port link, to derive the following conditions:

1. Upstream port enters the L0s Link PM state when all enabled downstream Receivers are in the L0s Link PM state or deeper, or in a Link Down state.
2. Upstream port enters the active L1 Link PM state, only when all downstream ports are in the active L1 Link PM state or deeper, or the link is down.
3. When a downstream port is in the active L1 Link PM state and an ASPM L1 Link PM state exit is occurring in the downstream port, the upstream port exits the L1 Link PM state.
4. When the upstream port is in the active L1 Link PM state and an active L1 Link PM state exit is occurring, due to Receiver Electrical Idle exit, the downstream port exits the L1 Link PM state.
5. When a PME_TO_Ack message is received only on all active (not in Link Down) downstream ports, a PME_TO_Ack message is issued toward the upstream port.
6. When all downstream ports are in the L2/L3 Ready Link PM or Link Down state, the upstream port transmits only a PM_ENTER_L23 DLLP toward the Root Complex.

12.5 Power Management Event Handler

PM_PME messages are Posted TLPs that inform the Power Management software which agent within the PCI Express hierarchy has requested a PM-state change. PM_PME messages are always routed toward the Root Complex.

PCI Express components are permitted to wake the system from any supported Power Management state, through the request of a Power Management Event (PME).

When a PEX 8509 Hot Plug-capable downstream port is in the D3hot Device PM state, the following Hot Plug events cause the **Power Management Status and Control** register *PME Status* bit to be set (offset 44h[15]=1):

- Hot Plug
 - Presence Detect Changed (HP_PRSENT_x# input signal)
 - 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

This causes the downstream port to generate a PM_PME message, if the **Power Management Status and Control** register *PME Enable* bit is set (offset 44h[8]=1).

THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 13 Port Registers

13.1 Introduction

This chapter defines the PEX 8509 port registers. The PEX 8509 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 8509 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 [Table 13-3](#).)

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*
- *I²C Bus v2.1*

13.2 Type 1 Port Register Map

Table 13-1. Type 1 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	
Type 1 Configuration Header Registers			00h
			...
			34h
			...
			3Ch
			40h
Power Management Capability Registers			44h
			48h
			...
Message Signaled Interrupt Capability Registers			64h
			68h
			8Ch
			90h
Subsystem ID and Subsystem Vendor ID Capability Registers			94h
<i>Reserved</i>			98h –
			DCh
			...
Vendor-Specific Enhanced Capability Registers			FCh
			100h
Next Capability Offset (FB4h)	1h	PCI Express Extended Capability ID (0003h)	100h
Device Serial Number Extended Capability Registers			134h
			138h
			...
Power Budget Extended Capability Registers			144h
			148h
			...
Next Capability Offset (000h)	1h	PCI Express Extended Capability ID (0002h)	148h
Virtual Channel Extended Capability Registers			1A4h
Port Arbitration Table Registers			1A8h
			1BCh

Table 13-1. Type 1 Port 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																																
Device-Specific Registers																																1C0h																
																																...																
																																FB0h																
Next Capability Offset (138h)																1h																PCI Express Extended Capability ID (0001h)																FB4h
Advanced Error Reporting Extended Capability Registers																																...																
																																FFCh																

13.3 Port Register Configuration and Map

The PEX 8509 port registers are configured similarly – not all the same. Port 0 includes more Device-Specific registers than the other ports. These Device-Specific registers contain setup and control information specific to the station. Port 0 also contains registers that are used to set up and control the PEX 8509, as well as serial EEPROM interface and I²C interface logic and control. Table 13-2 defines the port register configuration and map.

Table 13-2. Port Register Configuration and Map

Register Types	Port 0	Ports 1, 2, 3, 4, 5, 6, 7
Type 1 Configuration Header 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
Subsystem ID and Subsystem Vendor ID Capability Registers	90h - 94h	90h - 94h
Vendor-Specific Enhanced Capability Registers	DCh - FCh	DCh - FCh
Device Serial Number Extended Capability Registers	100h - 134h	100h - 134h
Power Budget Extended Capability Registers	138h - 144h	138h - 144h
Virtual Channel Extended Capability Registers	148h - 1A4h	148h - 1A4h
Port Arbitration Table Registers	1A8h - 1BCh	1A8h - 1BCh
Device-Specific Registers		
Error Checking	1C0h - 1D0h	1CCh, 1D0h
Debug Control	1D4h - 1DCh	
Power Management, Hot Plug ^a , and Miscellaneous Control	1E0h - 1FCh	1E0h - 1ECh, 1F8h, 1FCh
Physical Layer (all except serial EEPROM-related)	200h - 25Ch, 270h - 28Ch	
Serial EEPROM	260h - 26Ch	
Device-Specific Registers – I ² C Interface	290h - 2C4h	
Device-Specific Registers – Bus Number CAM	2C8h - 304h	
Device-Specific Registers – I/O CAM	308h - 344h	
Device-Specific Registers – Address-Mapping CAM	348h - 548h	
Device-Specific Registers – Ingress Control and Port Enable	660h - 67Ch	
Device-Specific Registers – I/O CAM Base and Limit Upper 16 Bits	680h - 6BCh	
Device-Specific Registers – Base Address Shadow	6C0h - 73Ch	
Device-Specific Registers – Shadow Virtual Channel Capability	740h - 83Ch	
Device-Specific Registers – Ingress Credit Handler	940h - B7Ch	A24h - B7Ch
Device-Specific Registers – Port Configuration Header	E00h - E3Ch	
Device-Specific Registers – Source Queue Weight and Soft Error	F10h - FB0h	F10h
Advanced Error Reporting Extended Capability Registers	FB4h - FFCh	FB4h - FFCh

a. Hot Plug is supported only on downstream Ports 1, 2, and 3.

13.4 Register Access

Each PEX 8509 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 8509 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 8509 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 8509 port Type 1 (PCI-to-PCI Bridge) registers:

- [Type 1 Configuration Header Registers](#)
- [Power Management Capability Registers](#)
- [Message Signaled Interrupt Capability Registers](#)
- [PCI Express Capability Registers](#)
- [Subsystem ID and Subsystem Vendor ID Capability Registers](#)
- [Vendor-Specific Enhanced Capability Registers](#)

The *PCI r3.0*-Compatible Configuration mechanism uses PCI Type 0 and Type 1 Configuration transactions to access the PEX 8509 Configuration registers. The PEX 8509 upstream port captures the Bus and Device Numbers assigned by the upstream device on the PCI Express link attached to the PEX 8509 upstream port, as required by the *PCI Express Base r1.1*.

The PEX 8509 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 8509 internal virtual PCI Bus Number, the PEX 8509 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 8509 downstream ports, the PEX 8509 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 8509 downstream port Device Numbers, the PEX 8509 responds with an *Unsupported Request* (UR).
 - If the specified Bus Number in the Type 1 Configuration access is not the PEX 8509 internal virtual PCI Bus Number, but is the number of one of the PEX 8509 downstream port secondary/subordinate buses, the PEX 8509 passes the Configuration access on to the PCI Express link attached to that PEX 8509 downstream port.
 - If the specified Bus Number is the downstream port Secondary Bus Number, and specified Device Number is 0, the PEX 8509 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 8509 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 8509 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 8509 Device-Specific registers above 100h. The Device-Specific Memory-Mapped Configuration mechanism can access all PEX 8509 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 8509 port Type 1 (PCI-to-PCI Bridge) registers that are defined the *PCI Express Base r1.1*:

- [Type 1 Configuration Header Registers](#)
- [Power Management Capability Registers](#)
- [Message Signaled Interrupt Capability Registers](#)
- [PCI Express Capability Registers](#)
- [Subsystem ID and Subsystem Vendor ID Capability Registers](#)
- [Vendor-Specific Enhanced Capability Registers](#)
- [Device Serial Number Extended Capability Registers](#)
- [Power Budget Extended Capability Registers](#)
- [Virtual Channel Extended Capability Registers](#)
- [Advanced Error Reporting Extended Capability Registers](#)

The PEX 8509 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 8509 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 listed in [Table 13-3](#). The registers of each port are contained within a 4-KB range. The PEX 8509 supports up to eight simultaneously active ports.

The PEX 8509 requires a single contiguous Memory space of 128 KB to contain all the PEX 8509 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 Type 1 Configuration Space **Base Address 0** and **Base Address 1** registers (**BAR0** and **BAR1**, offsets **10h** and **14h**, respectively). After the PEX 8509 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. (Refer to [Table 13-3](#).) Within each of these 4-KB windows, individual registers are located at the DWord offsets indicated in [Table 13-1](#).

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, I²C interface, or software), the PEX 8509 automatically copies the values into the **BAR0** and **BAR1 Shadow** registers that exist in Ports 0. (Refer to [Table 13-29](#) for register mapping.) The particular 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.

Table 13-3. PEX 8509 Register Offsets from Upstream Port BAR0/1 Base Address

Port Number	Internal Register 4-KB Memory Space Range	Location Range
Port 0	0000h to 0FFFh	0 to 4 KB
Port 1	1000h to 1FFFh	4 to 8 KB
Port 2	2000h to 2FFFh	8 to 12 KB
Port 3	3000h to 3FFFh	12 to 16 KB
Port 4	4000h to 4FFFh	16 to 20 KB
Port 5	5000h to 5FFFh	20 to 24 KB
Port 6	6000h to 6FFFh	24 to 28 KB
Port 7	7000h to 7FFFh	28 to 32 KB

13.5 Register Descriptions

The remainder of this chapter details the PEX 8509 registers, including:

- Bit/field names
- Description of register functions for the PEX 8509 upstream port and downstream ports
- Type (*such as* RW or HwInit; refer to [Table 13-4](#) for Type descriptions)
- Whether the power-on/reset value can be modified, by way of the PEX 8509 serial EEPROM or I²C Initialization feature
- Default power-on/reset value

Table 13-4. Register Types, Grouped by User Accessibility

Type	Description
HwInit	Hardware-Initialized Refers to the PEX 8509 Hardware-Initialization mechanism or PEX 8509 Serial EEPROM and I ² C register Initialization features. Read-Only after initialization and can only be reset with a Fundamental Reset.
RO	Read-Only Read-Only and cannot be altered by software. Initialized by the PEX 8509 Hardware-Initialization mechanism or PEX 8509 serial EEPROM and/or I ² C register Initialization features.
ROS	Read-Only, Sticky Same as RO, except that bits are not initialized nor modified by a Hot Reset.
RsvdP	Reserved and Preserved <i>Reserved</i> for future RW implementations. Registers are Read-Only and must return 0 when read. Software must preserve value read for writes to bits.
RsvdZ	Reserved and Zero <i>Reserved</i> for future RWC implementations. Registers are Read-Only and must return 0 when read. Software must use 0 for writes to bits.
RW	Read-Write Read/Write and is set or cleared to the needed state by software.
RWC	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.
RWCS	Read-Only Status, Write 1 to Clear, Sticky Same as RWC, except that bits are not modified by a Hot Reset.
RWS	Read-Write, Sticky Same as RW, except that bits are not modified by a Hot Reset.
RZ	Software Read Zero Software Read always return 0; however, software is allowed to write this register.

13.6 Type 1 Configuration Header Registers

This section details the PEX 8509 Type 1 Configuration Header registers. [Table 13-5](#) defines the register map.

Table 13-5. Type 1 Configuration Header 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
PCI Status								PCI Command								04h
PCI Class Code										Revision ID						08h
BIST (Not Supported)				Header Type				Master Latency Timer (Not Supported)				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 Base Upper 32 Bits																28h
Prefetchable Memory Limit Upper 32 Bits																2Ch
I/O Limit Upper 16 Bits								I/O Base Upper 16 Bits								30h
<i>Reserved</i>										New Capability Pointer (40h)						34h
Expansion ROM Base Address (<i>Reserved</i>)																38h
Bridge Control								PCI Interrupt Pin				PCI Interrupt Line				3Ch

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	Vendor ID Identifies the device manufacturer. Defaults to the PCI-SIG-issued Vendor ID of PLX (10B5h) if not overwritten by serial EEPROM or I ² C.	RO	Yes	10B5h
31:16	Device ID Identifies the particular device. Defaults to the PLX part number for the PEX 8509, if not overwritten by serial EEPROM or I ² C.	RO	Yes	8509h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
PCI Command				
0	I/O Access Enable 0 = PEX 8509 ignores I/O Space accesses on the corresponding port's primary interface 1 = PEX 8509 responds to I/O Space accesses on the corresponding port's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory Space accesses on the corresponding port's primary interface 1 = PEX 8509 responds to Memory Space accesses on the corresponding port's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 forwarding of Memory or I/O requests in the upstream direction. Does not affect forwarding of Completions in the upstream nor downstream direction, nor forwarding of messages (including INTx Interrupt messages). 0 = PEX 8509 handles Memory and I/O requests received on the corresponding port downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status. Because MSI Interrupt messages are in-band Memory Writes, disables MSI Interrupt messages as well. 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction.	RW	Yes	0
3	Special Cycle Enable <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
4	Memory Write and Invalidate <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
5	VGA Palette Snoop <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
6	Parity Error Response Enable Controls bit 24 (<i>Master Data Parity Error</i>).	RW	Yes	0
7	IDSEL Stepping/Wait Cycle Control <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
8	SERR# Enable Controls bit 30 (<i>Signaled System Error</i>). 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 <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
10	Interrupt Disable 0 = Corresponding PEX 8509 port is enabled to generate INT _x Interrupt messages and assert PEX_INTA# output 1 = Corresponding PEX 8509 port is prevented from generating INT _x Interrupt messages and asserting PEX_INTA# output	RW	Yes	0
15:11	<i>Reserved</i>	RsvdP	No	00h
PCI Status				
18:16	<i>Reserved</i>	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 8509 port and PEX_INTA# is asserted	RO	Yes	0
20	Capability List New Capability function is supported. Required by the <i>PCI Express Base r1.1</i> to be 1 at all times.	RO	Yes	1
21	66 MHz Capable Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
22	<i>Reserved</i>	RsvdP	No	0
23	Fast Back-to-Back Transactions Capable <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
24	Master Data Parity Error If bit 6 (<i>Parity Error Response Enable</i>) is set to 1, the corresponding PEX 8509 port sets this bit to 1 when the port: <ul style="list-style-type: none"> 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 <i>Parity Error Response Enable</i> bit is cleared to 0, the PEX 8509 never sets this bit. This error is natively reported by the Uncorrectable Error Status register <i>Poisoned TLP Status</i> bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RWC	Yes	0
26:25	DEVSEL# Timing <i>Not supported</i> Cleared to 00b, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	00b
27	Signaled Target Abort This bit is set if a Memory request targets a non-existent port.	RWC	Yes	0

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
28	Received Target Abort Cleared to 0. Never set to 1.	RsvdP	No	0
29	Received Master Abort Cleared to 0. Never set to 1.	RsvdP	No	0
30	Signaled System Error If bit 8 (<i>SERR# Enable</i>) is set to 1, the corresponding PEX 8509 port sets this bit to 1 when it transmits or forwards an ERR_FATAL or ERR_NONFATAL message upstream. This error is natively reported by the Device Status register <i>Fatal Error Detected</i> and <i>Non-Fatal Error Detected</i> bits (offset 70h[18:17], respectively), which are mapped to this bit for Conventional PCI backward compatibility.	RWC	Yes	0
31	Detected Parity Error Set to 1 when the corresponding port receives a Poisoned TLP on its primary side, regardless of the bit 6 (<i>Parity Error Response Enable</i>) state. This error is natively reported by the Uncorrectable Error Status register <i>Poisoned TLP Status</i> bit (offset FB8h[12]), which is mapped to this bit for Conventional PCI backward compatibility.	RWC	Yes	0

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Revision ID Unless overwritten by the serial EEPROM, returns the Silicon Revision (AAh), the PLX-assigned Revision ID for this version of the PEX 8509. The PEX 8509 Serial EEPROM register Initialization capability is used to replace the PLX Revision ID with another Revision ID.	RO	Yes	AAh
PCI Class Code				060400h
15:8	Register-Level Programming Interface The PEX 8509 ports support the <i>PCI-to-PCI Bridge v1.2</i> requirements, but not subtractive decoding, on their 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 ² C	Default
7:0	Cache Line Size System Cache Line Size. Implemented as a RW field for Conventional PCI compatibility purposes and does not impact PEX 8509 functionality.	RW	Yes	00h
15:8	Master Latency Timer <i>Not supported</i> Cleared to 00h, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	00h
22:16	Configuration Layout Type The corresponding PEX 8509 port Configuration Space Header adheres to the Type 1 PCI-to-PCI Bridge Configuration Space layout defined by the <i>PCI-to-PCI Bridge r1.2</i> .	RO	Yes	01h
23	Header Type Always 0, because the PEX 8509 is a single-function device.	RO	Yes	0
31:24	BIST <i>Not supported</i> 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 ² C	Default
0	Memory Space Indicator 0 = Base Address register maps the PEX 8509 Configuration registers into Memory space <i>Note: Upstream port is hardwired to 0.</i>	Upstream	RO	No	0
	<i>Reserved</i>	Downstream	RsvdP	No	0
2:1	Memory Map Type 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	Upstream	RO	Yes	00b
	<i>Reserved</i>	Downstream	RsvdP	No	00b
3	Prefetchable 0 = Base Address register maps the PEX 8509 Configuration registers into Non-Prefetchable Memory space <i>Note: Upstream port is hardwired to 0.</i>	Upstream	RO	Yes	0
	<i>Reserved</i>	Downstream	RsvdP	No	0
16:4	<i>Reserved</i>		RsvdP	No	0-0h
31:17	Base Address 0 Base Address (BAR0) for Device-Specific Memory-Mapped Configuration mechanism.	Upstream	RW	Yes	0000h
	<i>Reserved</i>	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 I ² C	Default
31:0	Base Address 1 For 64-bit addressing, Base Address 1 (BAR1) extends Base Address 0 to provide the upper 32 Address bits when the Base Address 0 register <i>Memory Map Type</i> field (offset 10h[2:1]) is set to 10b.	Upstream	RW	Yes	0000_0000h
	Read-Only when the Base Address 0 register is not enabled as a 64-bit BAR [<i>Memory Map Type</i> field (offset 10h[2:1]) is not equal to 10b].		RO	Yes	0000_0000h
	<i>Reserved</i>	Downstream	RsvdP	No	0000_0000h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 <i>Not supported</i> Cleared to 00h, as required by the <i>PCI Express Base r1.1</i> .	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 ² C	Default
<p>Note: If ISA Addressing mode is enabled [PCI Command register I/O Access Enable bit (offset 04h[0]=1)], the PEX 8509 port forwards I/O transactions from its primary interface to its secondary interface (downstream) if an I/O address is within the range defined by the I/O Base and I/O Limit registers when the Base is less than or equal to the Limit.</p> <p>Conversely, the PEX 8509 port forwards I/O transactions from its secondary interface to its primary interface (upstream) if an I/O address is outside this Address range. If the PEX 8509 port does not implement an I/O Address range, the port forwards all I/O transactions on its secondary interface upstream, to its primary interface.</p>				
I/O Base				
3:0	<p>I/O Base Addressing Capability</p> <p>1h = 32-bit Address decoding is supported</p> <p>All other encodings are <i>reserved</i>.</p>	RO	Yes	1h
7:4	<p>IO_BAR</p> <p>I/O Base Address[15:12]. The PEX 8509 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.</p> <p>I/O Base Address[15:12] bits specify the corresponding PEX 8509 port I/O Base Address[15:12]. The PEX 8509 assumes I/O Base Address[11:0]=000h.</p> <p>For 16-bit I/O addressing, the PEX 8509 assumes Address[31:16]=0000h.</p> <p>For 32-bit addressing, the PEX 8509 decodes Address[31:0], and uses the I/O Base and Limit Upper 16 Bits register <i>I/O Base Upper 16 Bits</i> and <i>I/O Limit Upper 16 Bits</i> fields (offset 30h[15:0 and 31:16], respectively).</p>	RW	Yes	Fh
I/O Limit				
11:8	<p>I/O Limit Addressing Capability</p> <p>1h = 32-bit Address decoding is supported</p> <p>All other encodings are <i>reserved</i>.</p>	RO	Yes	1h
15:12	<p>IO_Limit</p> <p>I/O Limit Address[15:12]. The PEX 8509 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.</p> <p>I/O Limit Address[15:12] bits specify the corresponding PEX 8509 port I/O Limit Address[15:12]. The PEX 8509 assumes I/O Limit Address[11:0]=FFFh.</p> <p>For 16-bit I/O addressing, the PEX 8509 decodes Address bits [15:0] and assumes I/O Limit Address[31:16]=0000h.</p> <p>For 32-bit addressing, the PEX 8509 decodes Address bits [31:0], and uses the I/O Base and Limit Upper 16 Bits register <i>I/O Base Upper 16 Bits</i> and <i>I/O Limit Upper 16 Bits</i> fields (offset 30h[15:0 and 31:16], respectively).</p> <p>If the I/O Limit Address is less than the I/O Base Address, the PEX 8509 does not forward I/O transactions from the corresponding port primary/upstream bus to its secondary/downstream bus. However, the PEX 8509 forwards all I/O transactions from the secondary bus of the corresponding port to its primary bus.</p>	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 I ² C	Default
Secondary Status				
20:16	<i>Reserved</i>	RsvdP	No	0-0h
21	66 MHz Capable <i>Not supported</i> 0 = Not enabled, because PCI Express does not support 66 MHz	RsvdP	No	0
22	<i>Reserved</i>	RsvdP	No	0
23	Fast Back-to-Back Transactions Capable <i>Reserved</i> Not enabled, because PCI Express does <i>not support</i> this function.	RsvdP	No	0
24	Master Data Parity Error If the Bridge Control register <i>Parity Error Response Enable</i> bit (offset 3Ch[16]) is set to 1, the corresponding PEX 8509 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: <ul style="list-style-type: none"> Port receives Completion marked poisoned Port forwards poisoned TLP Write request When the <i>Parity Error Response Enable</i> bit is cleared to 0, the PEX 8509 never sets this bit. These errors are reported by the port's Uncorrectable Error Status register <i>Poisoned TLP Status</i> bit (offset FB8h[12]), and mirrored to this bit for Conventional PCI backward compatibility.	RWC	Yes	0
26:25	DEVSEL# Timing <i>Not supported</i> Cleared to 00b, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	00b
27	Signaled Target Abort Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
28	Received Target Abort Cleared to 0, as required by the <i>PCI Express Base r1.1</i> , because the PEX 8509 never initiates a request itself.	RsvdP	No	0
29	Received Master Abort Cleared to 0, as required by the <i>PCI Express Base r1.1</i> , because the PEX 8509 never initiates a request itself.	RsvdP	No	0
30	Received System Error Set to 1 when a downstream port receives an ERR_FATAL or ERR_NONFATAL message on its secondary interface from a downstream device.	RWC	Yes	0
31	Detected Parity Error Set to 1 by a downstream port when receiving a poisoned TLP from a downstream device, regardless of the Bridge Control register <i>Parity Error Response Enable</i> bit (offset 3Ch[16]) state.	RWC	Yes	0

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<p><i>Note:</i> The PEX 8509 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).</p> <p>Conversely, the PEX 8509 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].</p>				
Memory Base				
3:0	<i>Reserved</i>	RsvdP	No	0h
15:4	MEM_BAR[31:20] Specifies the corresponding PEX 8509 port Non-Prefetchable Memory Base Address[31:20]. The PEX 8509 assumes Memory Base Address[19:0]=0_0000h.	RW	Yes	FFFh
Memory Limit				
19:16	<i>Reserved</i>	RsvdP	No	0h
31:20	MEM_Limit[31:20] Specifies the corresponding PEX 8509 port Non-Prefetchable Memory Limit Address[31:20]. The PEX 8509 assumes Memory Limit Address[19:0]=F_FFFFh.	RW	Yes	000h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<p>Note: The PEX 8509 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).</p> <p>Conversely, the PEX 8509 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)].</p>				
Prefetchable Memory Base				
3:0	<p>Prefetchable Memory Base Capability</p> <p>0h = Corresponding PEX 8509 port supports 32-bit Prefetchable Memory Addressing</p> <p>1h = Corresponding PEX 8509 port defaults to 64-bit Prefetchable Memory Addressing support, as required by the <i>PCI Express Base r1.1</i></p> <p>Note: If the application needs 32-bit-only Prefetchable space, the serial EEPROM or I²C must clear both this field and field [19:16] (Prefetchable Memory Limit register Prefetchable Memory Limit Capability).</p>	RO	Yes	1h
15:4	<p>PMEM_BAR[31:20]</p> <p>Specifies the corresponding PEX 8509 port Prefetchable Memory Base Address[31:20].</p> <p>The PEX 8509 assumes Prefetchable Memory Base Address[19:0]=0_0000h.</p>	RW	Yes	FFFh
Prefetchable Memory Limit				
19:16	<p>Prefetchable Memory Limit Capability</p> <p>0h = Corresponding PEX 8509 port supports 32-bit Prefetchable Memory Addressing</p> <p>1h = Corresponding PEX 8509 port defaults to 64-bit Prefetchable Memory Addressing support, as required by the <i>PCI Express Base r1.1</i></p>	RO	Yes	1h
31:20	<p>PMEM_Limit[31:20]</p> <p>Specifies the corresponding PEX 8509 port Prefetchable Memory Limit Address[31:20].</p> <p>The PEX 8509 assumes Prefetchable Memory Limit Address[19:0]=F_FFFFh.</p>	RW	Yes	000h

Register 13-11. 28h Prefetchable Memory Base Upper 32 Bits (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default	
31:0	PBUP[63:32] The PEX 8509 uses this register for Prefetchable Memory Upper Base Address[63:32]. When the Prefetchable Memory Base register <i>Prefetchable Memory Base Capability</i> field indicates 32-bit addressing, this register is RO 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 Limit Upper 32 Bits (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default	
31:0	PLIMUP[63:32] The PEX 8509 uses this register for Prefetchable Memory Upper Limit Address[63:32]. When the Prefetchable Memory Limit register <i>Prefetchable Memory Limit Capability</i> field indicates 32-bit addressing, this register is RO 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 Base and Limit Upper 16 Bits (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default	
15:0	I/O Base Upper 16 Bits The PEX 8509 uses this register for I/O Base Address[31:16]. When the I/O Base register <i>I/O Base Addressing Capability</i> field indicates 16-bit addressing, this register is Read-Only and returns 0000h.	When offset 1Ch[3:0]=1h	RW	Yes	0000h
		When offset 1Ch[3:0]=0h	RO	No	0000h
31:16	I/O Limit Upper 16 Bits The PEX 8509 uses this register for I/O Limit Address[31:16]. When the I/O Limit register <i>I/O Limit Addressing Capability</i> 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 I ² C	Default
7:0	New Capability Pointer Default 40h points to the Power Management Capability structure.	RO	Yes	40h
31:8	<i>Reserved</i>	RsvdP	No	0000_00h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Expansion ROM Base Address <i>Reserved</i> Cleared to 0000_0000h.	RsvdP	No	0000_0000h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
PCI Interrupt Signal				
7:0	PCI Interrupt Line The PEX 8509 does <i>not</i> use this register; however, the register is included for operating system and device driver use.	RW	Yes	00h
15:8	PCI 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 8509. 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 PCI Interrupt Signal (All Ports) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Bridge Control				
16	<p>Parity Error Response Enable Controls the response to Poisoned TLPs.</p> <p>0 = Disables the Secondary Status register <i>Master Data Parity Error</i> bit (offset 1Ch[24]) 1 = Enables the Secondary Status register <i>Master Data Parity Error</i> bit (offset 1Ch[24])</p>	RW	Yes	0
17	<p>SERR# Enable Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the PCI Command register <i>SERR# Enable</i> bit (offset 04h[8]) is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit (offset 04h[30]).</p>	RW	Yes	0
18	<p>ISA Enable Modifies the PEX 8509's response to ISA I/O addresses enabled by the I/O Base and I/O Limit registers (offset 1Ch[15:8] and [7:0], respectively) and located in the first 64 KB of the PCI I/O Address space (0000_0000h to 0000_FFFFh).</p> <p>0 = If ISA Addressing mode is enabled [PCI Command register <i>I/O Access Enable</i> bit (offset 04h[0]) is set to 1], the PEX 8509 port forwards I/O requests within the Address range defined by the I/O Base and I/O Limit registers.</p> <p>1 = PEX 8509 blocks forwarding from the primary to secondary interface, of I/O transactions addressing the last 768 bytes in each 1-KB block of the port's I/O Address range. In the opposite direction (secondary to primary), if I/O addressing mode is enabled [PCI Command register <i>I/O Access Enable</i> bit (offset 04h[0]) is set to 1], the PEX 8509 port forwards I/O transactions that address the last 768 bytes in each 1-KB block of the port's I/O Address range.</p> <p><i>Note: Refer also to the Ingress Control register <i>Disable VGA BIOS Memory Access Decoding</i> bit (Port 0, offset 660h[28]).</i></p>	RW	Yes	0

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
19	<p>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):</p> <ul style="list-style-type: none"> • 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> • 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 (<i>ISA Enable</i>) or PCI Command register <i>VGA Palette Snoop</i> bit (offset 04h[5]) settings <p>VGA address forwarding is qualified by the PCI Command register <i>Memory Access Enable</i> and <i>I/O Access Enable</i> bits (offset 04h[1:0], respectively). The default state of this bit after reset must be 0.</p> <p>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>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p> <p><i>Notes: When set in an egress port, the port is configured as a non-Cut-Thru path. (Refer to Section 4.3.2, “Cut-Thru Mode,” for further details.)</i></p> <p><i>Refer also to the Ingress Control register <i>Disable VGA BIOS Memory Access Decoding</i> bit (Port 0, offset 660h[28]).</i></p> <p><i>For Conventional PCI VGA support, to avoid potential I/O address conflicts, if the VGA Enable bit is set in the upstream port and a downstream port, the PCI Command register <i>I/O Access Enable</i> bit (offset 04h[0]) should be set to 1 in all other PEX 8509 downstream ports, unless those downstream ports are configured to use default 32-bit address decoding and their I/O Address range is set above 1000h.</i></p>	RW	Yes	0

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
20	<p>VGA 16-Bit Decode Enable</p> <p>Used only when bit 19 (<i>VGA Enable</i>) or the PCI Command register <i>VGA Palette Snoop</i> bit (offset 04h[5]) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p> <p><i>Note:</i> Refer also to the Ingress Control register <i>Disable VGA BIOS Memory Access Decoding bit</i> (Port 0, offset 660h[28]).</p>	RW	Yes	0
21	<p>Master Abort Mode</p> <p><i>Not supported</i></p> <p>Cleared to 0, as required by the <i>PCI Express Base r1.1</i>.</p>	RsvdP	No	0
22	<p>Secondary Bus Reset</p> <p>1 = Causes a Hot Reset on the corresponding PEX 8509 port downstream link</p>	RW	Yes	0
23	<p>Fast Back-to-Back Transactions Enable</p> <p><i>Not supported</i></p> <p>Cleared to 0, as required by the <i>PCI Express Base r1.1</i>.</p>	RsvdP	No	0
24	<p>Primary Discard Timer</p> <p><i>Not supported</i></p> <p>Cleared to 0, as required by the <i>PCI Express Base r1.1</i>.</p>	RsvdP	No	0
25	<p>Secondary Discard Timer</p> <p><i>Not supported</i></p> <p>Cleared to 0, as required by the <i>PCI Express Base r1.1</i>.</p>	RsvdP	No	0
26	<p>Discard Timer Status</p> <p><i>Not supported</i></p> <p>Cleared to 0, as required by the <i>PCI Express Base r1.1</i>.</p>	RsvdP	No	0
27	<p>Discard Timer SERR# Enable</p> <p><i>Not supported</i></p> <p>Cleared to 0, as required by the <i>PCI Express Base r1.1</i>.</p>	RsvdP	No	0
31:28	Reserved	RsvdP	No	0h

13.7 Power Management Capability Registers

This section details the PEX 8509 Power Management Capability registers. Table 13-6 defines the register map.

Table 13-6. 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
Power Management Data	Power Management Control/ Status Bridge Extensions <i>(Reserved)</i>	Power Management Status and Control		44h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Capability ID Set to 01h, indicating that the Capability structure is the PCI Power Management Capability structure.	RO	Yes	01h
15:8	Next Capability Pointer Default 48h points to the Message Signaled Interrupt Capability structure.	RO	Yes	48h
18:16	Version Default 011b indicates compliance with the <i>PCI Power Mgmt. r1.2</i> .	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 <i>not</i> required.	RO	Yes	0
24:22	AUX Current The PEX 8509 does <i>not support</i> PME generation from the D3cold Device PM state; therefore, the serial EEPROM value for this field should be 000b.	RO	Yes	000b
25	D1 Support <i>Not supported</i> Default value of 0 indicates that the PEX 8509 does <i>not support</i> the D1 Device PM state.	RsvdP	No	0
26	D2 Support <i>Not supported</i> Default value of 0 indicates that the PEX 8509 does <i>not support</i> the D2 Device PM state.	RsvdP	No	0
31:27	PME Support Bits [31, 30, and 27] must be set to 1, to indicate that the PEX 8509 will forward PME messages, as required by the <i>PCI Express Base r1.1</i> .	RO	Yes	1100_1b

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Power Management Status and Control				
1:0	<p>Power State Used to determine the current Device PM state of the port, and to set the port into a new Device PM state.</p> <p>00b = D0 01b = D1 – <i>Not supported</i> 10b = D2 – <i>Not supported</i> 11b = D3hot</p> <p>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.</p>	RW	Yes	00b
2	Reserved	RsvdP	No	0
3	<p>No Soft Reset 1 = Devices transitioning from the D3hot to D0 Device PM state, because of Power State commands, do not perform an internal reset</p>	RO	Yes	1
7:4	Reserved	RsvdP	No	0h
8	<p>PME Enable 0 = Disables PME generation by the corresponding PEX 8509 port^a 1 = Enables PME generation by the corresponding PEX 8509 port</p>	RWS	No	0
12:9	<p>Data Select Initially writable by serial EEPROM and I²C only^b. After a Serial EEPROM or I²C Write occurs to this register, RW for all CSR accesses. Bits [12:9] select the Data and Data Scale registers.</p> <p>0h = D0 power consumed 3h = D3hot power consumed 4h = D0 power dissipated 7h = D3hot power dissipated</p>	RO	Yes	0h
	<p>Not supported RO for hardware auto-configuration.</p>	RO	No	0h
14:13	<p>Data Scale Writable by serial EEPROM and I²C only^b. There are four internal Data Scale registers (one each per Data register – 0, 3, 4 and 7), per port. Bits [12:9], <i>Data Select</i>, select the Data Scale register.</p>	RO	Yes	00b
15	<p>PME Status 0 = PME is not generated by the corresponding PEX 8509 port^a 1 = PME is being generated by the corresponding PEX 8509 port</p>	RWCS	No	0

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

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

- a. Because the PEX 8509 does not consume auxiliary power, this bit is not sticky, and is always cleared to 0 at power-on reset.
- b. With no serial EEPROM nor previous I²C programming, 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 8509 Message Signaled Interrupt (MSI) Capability registers. Table 13-7 defines the register map.

Table 13-7. 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	
MSI Control		Next Capability Pointer (68h)		Capability ID (05h)
MSI Address				
MSI Upper Address				
<i>Reserved</i>		MSI Data		
MSI Mask				
MSI Pending				
<i>Reserved</i>				60h – 64h

Register 13-19. 48h MSI Control and Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
MSI Capability Header				
7:0	Capability ID Set to 05h, as required by the <i>PCI r3.0</i> .	RO	Yes	05h
15:8	Next Capability Pointer Set to 68h, to point to the PCI Express Capability structure.	RO	Yes	68h
MSI Control				
16	MSI Enable 0 = Message Signaled interrupts for the corresponding port are disabled 1 = Message Signaled interrupts for the corresponding port are enabled, and INT _x Interrupt messages and PEX_INTA# output assertion are disabled	RW	Yes	0
19:17	Multiple Message Capable 000b = PEX 8509 port is requesting one message 001b = PEX 8509 port is requesting two messages All other encodings are <i>reserved</i> .	RO	Yes	001b
22:20	Multiple Message Enable 000b = PEX 8509 port contains only one allocated message – the only value supported	RW	Yes	000b
23	MSI 64-Bit Address Capable 1 = PEX 8509 is capable of generating 64-bit Message Signaled Interrupt addresses	RO	Yes	1
24	Per Vector Masking Capable 0 = PEX 8509 does not have Per Vector Masking capability 1 = PEX 8509 has Per Vector Masking capability	RO	Yes	1
31:25	<i>Reserved</i>	RsvdP	No	00h

Register 13-20. 4Ch MSI Address (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
1:0	<i>Reserved</i>	RsvdP	No	00b
31:2	Message Address <i>Note: Refer to register offset 50h for MSI Upper Address.</i>	RW	Yes	0000_000h

Register 13-21. 50h MSI Upper Address (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Message Upper Address MSI Write transaction upper address[63:32]. <i>Note: Refer to register offset 4Ch for MSI Address.</i>	RW	Yes	0000_0000h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	Message Data MSI Write transaction TLP payload.	RW	Yes	0000h
31:16	<i>Reserved</i>	RsvdP	No	0000h

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

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	MSI Mask for Hot Plug or Link State Interrupts	RW	Yes	0
1	MSI Mask for Device-Specific Interrupts	RW	Yes	0
31:2	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-24. 5Ch MSI Pending (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	MSI Pending Status for Hot Plug or Link State Interrupts	RO	No	0
1	MSI Pending Status for Device-Specific Interrupts	RO	No	0
31:2	<i>Reserved</i>	RsvdP	No	0-0h

13.9 PCI Express Capability Registers

This section details the PEX 8509 PCI Express Capability registers. Hot Plug Capability, Command, Status, and Events are included in these registers. [Table 13-8](#) defines the register map.

Table 13-8. 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 Capability		Next Capability Pointer (90h)		Capability ID (10h) 68h
Device Capability				6Ch
Device Status		<i>Not Supported/Reserved</i>		Device Control 70h
Link Capability				74h
Link Status		Link Control		78h
Slot Capability				7Ch
Slot Status		Slot Control		80h
<i>Reserved</i>				84h – 8Ch

Register 13-25. 68h PCI Express Capability List and Capability (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
PCI Express Capability List					
7:0	Capability ID Set to 10h, as required by the <i>PCI Express Base r1.1</i> .		RO	Yes	10h
15:8	Next Capability Pointer Set to 90h, to point to the Subsystem Capability structure.		RO	Yes	90h
PCI Express Capability					
19:16	Capability Version The PEX 8509 ports set this field to 1h, as required by the <i>PCI Express Base r1.1</i> .		RO	Yes	1h
23:20	Device/Port Type Set at reset, as required by the <i>PCI Express Base r1.1</i> .	Upstream	RO	No	5h
		Downstream	RO	No	6h
24	Slot Implemented 0 = Disables or connects to an 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 8509 Serial EEPROM register Initialization capability is used to change this value to 0h, indicating that the corresponding PEX 8509 downstream port connects to an integrated component or is disabled.

Register 13-26. 6Ch Device Capability (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
2:0	Maximum Payload Size Supported 000b = PEX 8509 port supports a 128-byte maximum payload 001b = PEX 8509 port supports a 256-byte maximum payload 010b = PEX 8509 port supports a 512-byte maximum payload 011b = PEX 8509 port supports a 1,024-byte maximum payload No other encodings are supported.		RO	Yes	011b
4:3	Phantom Functions Supported <i>Not supported</i> Cleared to 00b.		RO	Yes	00b
5	Extended Tag Field Supported 0 = Maximum <i>Tag</i> field is 5 bits 1 = Maximum <i>Tag</i> field is 8 bits		RO	Yes	0
8:6	Endpoint L0s Acceptable Latency <i>Not supported</i> Because the PEX 8509 is a switch and not an endpoint, the PEX 8509 does <i>not support</i> this feature. 000b = Disables the capability		RO	Yes	000b
11:9	Endpoint L1 Acceptable Latency <i>Not supported</i> Because the PEX 8509 is a switch and not an endpoint, the PEX 8509 does <i>not support</i> this feature. 000b = Disables the capability		RO	Yes	000b
14:12	<i>Reserved</i> , as required by the <i>PCI Express Base r1.1</i> .		RsvdP	No	000b
15	Role-Based Error Reporting		RO	Yes	1
17:16	<i>Reserved</i>		RsvdP	No	00b
25:18	Captured Slot Power Limit Value For the PEX 8509 upstream port, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in field [27:26] (<i>Captured Slot Power Limit Scale</i>).	Upstream	RO	Yes	00h
	Not valid for the downstream ports.	Downstream	RsvdP	No	00h
27:26	Captured Slot Power Limit Scale For the PEX 8509 upstream port, the upper limit on power supplied by the slot is determined by multiplying the value in this field by the value in field [25:18] (<i>Captured Slot Power Limit Value</i>). 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	<i>Reserved</i>		RsvdP	No	0h

Register 13-27. 70h Device Status and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Device Control				
0	Correctable Error Reporting Enable 0 = Disables 1 = Enables the corresponding PEX 8509 port to report Correctable errors	RW	Yes	0
1	Non-Fatal Error Reporting Enable 0 = Disables 1 = Enables the corresponding PEX 8509 port to report Non-Fatal errors	RW	Yes	0
2	Fatal Error Reporting Enable 0 = Disables 1 = Enables the corresponding PEX 8509 port to report Fatal errors	RW	Yes	0
3	Unsupported Request Reporting Enable 0 = Disables 1 = Enables the corresponding PEX 8509 port to report Unsupported Request errors	RW	Yes	0
4	PCI Express Relaxed Ordering Enable <i>Not supported</i> Cleared to 0.	RsvdP	No	0
7:5	Maximum Payload Size Software can change this field to configure the PEX 8509 ports to support other Payload Sizes; however, software cannot change this field to a value larger than that indicated by the Device Capability register <i>Maximum Payload Size Supported</i> field (offset 6Ch[2:0]). Maximum Payload Size port limitations are as follows: <ul style="list-style-type: none"> • 1,024 if the number of ports is ≤ 3 • 512 if the number of ports is ≤ 5 • 256 if the number of ports is ≤ 8 000b = Indicates that the PEX 8509 port is configured to support a Maximum Payload Size of 128 bytes 001b = Indicates that the PEX 8509 port is configured to support a Maximum Payload Size of 256 bytes 010b = Indicates that the PEX 8509 port is configured to support a Maximum Payload Size of 512 bytes 011b = Indicates that the PEX 8509 port is configured to support a Maximum Payload Size of 1,024 bytes	RW	Yes	000b

Register 13-27. 70h Device Status and Control (All Ports) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
8	Extended Tag Field Enable <i>Not supported</i> Cleared to 0.	RsvdP	No	0
9	Phantom Functions Enable <i>Not supported</i> Cleared to 0.	RsvdP	No	0
10	AUX Power PM Enable <i>Not supported</i> Cleared to 0.	RsvdP	No	0
11	Enable No Snoop <i>Not supported</i> Cleared to 0.	RsvdP	No	0
14:12	Maximum Read Request Size <i>Not supported</i> Cleared to 000b.	RsvdP	No	000b
15	Reserved Hardwired to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0

Register 13-27. 70h Device Status and Control (All Ports) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Device Status				
16	Correctable Error Detected Set when the corresponding port detects a Correctable error, regardless of the bit 0 (<i>Correctable Error Reporting Enable</i>) state. 0 = Corresponding PEX 8509 port did not detect a Correctable error 1 = Corresponding PEX 8509 port detected a Correctable error	RWC	Yes	0
17	Non-Fatal Error Detected Set when the corresponding port detects a Non-Fatal error, regardless of the bit 1 (<i>Non-Fatal Error Reporting Enable</i>) state. 0 = Corresponding PEX 8509 port did not detect a Non-Fatal error 1 = Corresponding PEX 8509 port detected a Non-Fatal error	RWC	Yes	0
18	Fatal Error Detected Set when the corresponding port detects a Fatal error, regardless of the bit 2 (<i>Fatal Error Reporting Enable</i>) state. 0 = Corresponding PEX 8509 port did not detect a Fatal error 1 = Corresponding PEX 8509 port detected a Fatal error	RWC	Yes	0
19	Unsupported Request Detected Set when the corresponding port detects an Unsupported Request, regardless of the bit 3 (<i>Unsupported Request Reporting Enable</i>) state. 0 = Corresponding PEX 8509 port did not detect an Unsupported Request 1 = Corresponding PEX 8509 port detected an Unsupported Request	RWC	Yes	0
20	AUX Power Detected <i>Not supported</i> Cleared to 0.	RsvdP	No	0
21	Transactions Pending <i>Not supported</i> Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .	RsvdP	No	0
31:22	Reserved	RsvdP	No	000h

Register 13-28. 74h Link Capability (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
3:0	Maximum Link Speed Set to 0001b, as required by the <i>PCI Express Base r1.1</i> for a 2.5 Gbps PCI Express link.		RO	Yes	0001b
9:4	Maximum Link Width Actual link width is set by STRAP_PORTCFG[1:0] . (Refer to Table 13-9 .) The PEX 8509 Maximum Link Width is x4 = 00_0100b. Valid widths are x1, x2, and x4.		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 PM state entry is supported 11b = L0s and L1 Link PM states are supported All other encodings are <i>reserved</i> .		RO	Yes	11b
14:12	L0s Exit Latency Indicates the L0s Link PM state exit latency for the given PCI Express link. Value depends upon the Physical Layer Command and Status register <i>N_FTS Value</i> field (offset 220h[15:8]) value. 101b = Corresponding PEX 8509 port L0s Link PM state Exit Latency is 1 μ s to less than 2 μ s		RO	No	101b
17:15	L1 Exit Latency Indicates the L1 Link PM state exit latency for the given PCI Express link. The value reported indicates the length of time that the corresponding PEX 8509 port requires to complete the transition from the L1 to L0 Link PM state. 101b = Corresponding PEX 8509 port L1 Link PM state Exit Latency is 16 μ s to less than 32 μ s		RO	Yes	101b
18	Clock Power Management		RO	Yes	0
19	<i>Reserved</i>	Upstream	RsvdP	No	0
	Surprise Down Error Reporting Capable Valid for the downstream ports only.	Downstream	RO	Yes	1
20	<i>Reserved</i>	Upstream	RsvdP	No	0
	Data Link Layer Link Active Reporting Capable	Downstream	RO	Yes	1
23:21	<i>Reserved</i>		RsvdP	No	000b
31:24	Port Number The Port Number is set by signal ball Strapping options. STRAP_PORTCFG[1:0] – Ports 0, 1, 2, 3, 4, 5, 6, 7		HwInit	No	Set by Strapping ball levels

Table 13-9. Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[1:0])	Link Width/Port							
	Port 0	Port 1	Port 2	Port 3	Port 4	Port 5	Port 6	Port 7
00b	x1	x1	x1	x1	x1	x1	x1	x1
01b	x2	x1	x1	x1	x1	x1	x1	
10b	x2	x2	x1	x1	x1	x1		
11b	x4	x1	x1	x1	x1			

Register 13-29. 78h Link Status and Control (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
Link Control					
1:0	Active State Power Management (ASPM) 00b = Disables L0s and L1 Link PM state Entries for the corresponding PEX 8509 port ^a 01b = Enables only L0s Link PM state Entry 10b = Enables only L1 Link PM state Entry 11b = Enables both L0s and L1 Link PM state Entries		RW	Yes	00b
2	<i>Reserved</i>		RsvdP	No	0
3	Read Request Return Parameter Control Read Request Return Parameter “R” control. Read Completion Boundary (RCB). Cleared to 0, as required by the <i>PCI Express Base r1.1</i> .		RO	Yes	0
4	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Link Disable Setting to 1 places the link on the corresponding PEX 8509 downstream port to the Disabled Link Training state.	Downstream	RW	Yes	0
5	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Retrain Link For PEX 8509 ports, always returns 0 when read. Writing 1 to this bit causes the corresponding PEX 8509 downstream port to initiate retraining of its PCI Express link.	Downstream	RZ	Yes	0
6	Common Clock Configuration 0 = Corresponding PEX 8509 port and the device at the other end of the corresponding port’s PCI Express link are operating with an asynchronous Reference Clock 1 = Corresponding PEX 8509 port and the device at the other 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 8509 port to transmit: <ul style="list-style-type: none"> • 4,096 FTS Ordered-Sets in the L0s Link PM state, • Followed by a single SKIP Ordered-Set prior to entering the L0 Link PM state, • Finally, transmission of 1,024 TS1 Ordered-Sets in the Recovery state. 		RW	Yes	0
8	Clock Power Management Enable The PEX 8509 does <i>not support</i> removal of the Reference Clock in the L1 and L2/L3 Ready Link PM states.		RsvdP	No	0
15:9	<i>Reserved</i>		RsvdP	No	00h

Register 13-29. 78h Link Status and Control (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
Link Status					
19:16	Link Speed Set to 0001b, as required by the <i>PCI Express Base r1.1</i> for a 2.5 Gbps PCI Express link.		RO	Yes	0001b
25:20	Negotiated Link Width Link width is determined by the negotiated value with the attached lane/port. 00_0000b = Link is down 00_0001b = x1 00_0010b = x2 00_0100b = x4 All other encodings are <i>not supported</i> .		RO	No	00_0000b
26	Reserved		RsvdP	No	0
27	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Link Training 1 = Indicates that the corresponding PEX 8509 downstream port requested link training, and the link training is in-progress or about to start	Downstream	RO	No	0
28	Slot Clock Configuration 0 = Indicates that the PEX 8509 uses an independent clock 1 = Indicates that the PEX 8509 uses the same physical Reference Clock that the platform provides on the connector		HwInit	Yes	0
29	Not valid	Upstream	RsvdP	No	0
	Data Link Layer Link Active When set to 1, and the Link Capability register <i>Data Link Layer Link Active Reporting Capable</i> bit is also set (offset 74h[20]=1), indicates the following: <ul style="list-style-type: none"> Data Link Layer is in the <i>DL_Active</i> state Link is operational Flow Control Initialization has successfully completed 	Downstream	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 13-30. 7Ch Slot Capability (All Downstream Ports; Upstream Port Always Reads 0)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
<i>Note: Hot Plug-capable ports are downstream Ports 1, 2, and 3.</i>					
0	Reserved	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Attention Button Present 0 = Attention Button is not implemented 1 = Attention Button is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	HwInit	Yes	1
1	Reserved	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Power Controller Present 0 = Power Controller is not implemented 1 = Power Controller is implemented for the slot of the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	HwInit	Yes	1
2	Reserved	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	MRL Sensor Present 0 = MRL Sensor is not implemented 1 = MRL Sensor is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	HwInit	Yes	1
3	Reserved	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Attention Indicator Present 0 = Attention Indicator is not implemented 1 = Attention Indicator is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	HwInit	Yes	1

Register 13-30. 7Ch Slot Capability (All Downstream Ports; Upstream Port Always Reads 0) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
4	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Power Indicator Present 0 = Power Indicator is not implemented 1 = Power Indicator is implemented on the slot chassis of the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	HwInit	Yes	1
5	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Hot Plug Surprise 0 = No device in the corresponding PEX 8509 downstream port slot is removed from the system without prior notification 1 = Device in the corresponding PEX 8509 downstream port slot can be removed from the system without prior notification	Downstream Hot Plug-capable	HwInit	Yes	0
6	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Hot Plug Capable 0 = Corresponding PEX 8509 downstream port slot is not capable of supporting Hot Plug operations 1 = Corresponding PEX 8509 downstream port slot is capable of supporting Hot Plug operations	Downstream Hot Plug-capable	HwInit	Yes	1
14:7	<i>Reserved</i>	Upstream	RsvdP	No	00h
	Slot Power Limit Value The maximum power supplied by the corresponding PEX 8509 downstream slot is determined by multiplying this field's value (expressed in decimal; 25d = 19h) by the value specified in the field [16:15] (<i>Slot Power Limit Scale</i>). This field must be implemented if the PCI Express Capability register <i>Slot Implemented</i> bit is set (offset 68h[24]=1, default). Serial EEPROM or I ² C Writes to this register or a Data Link Layer Up event 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 Capability register <i>Captured Slot Power Limit Value</i> and <i>Captured Slot Power Limit Scale</i> fields.	Downstream	HwInit	Yes	19h

Register 13-30. 7Ch Slot Capability (All Downstream Ports; Upstream Port Always Reads 0) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
16:15	<i>Reserved</i>	Upstream	RsvdP	No	00b
	<p>Slot Power Limit Scale</p> <p>The maximum power supplied by the corresponding PEX 8509 downstream slot is determined by multiplying this field's value by the value specified in field [14:7] (<i>Slot Power Limit Value</i>).</p> <p>This field must be implemented if the PCI Express Capability register <i>Slot Implemented</i> bit is set (offset 68h[24]=1, default).</p> <p>Serial EEPROM or I²C Writes to this register or a Data Link Layer Up event 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 Capability register <i>Captured Slot Power Limit Value</i> and <i>Captured Slot Power Limit Scale</i> fields.</p> <p>00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x</p>	Downstream	HwInit	Yes	00b
17	Electromechanical Interlock Present <i>Not supported</i>		RsvdP	No	0
18	No Command Completed Support <i>Not supported</i>		RsvdP	No	0
31:19	<i>Reserved</i>	Upstream	RsvdP	No	0-0h
	<p>Physical Slot Number</p> <p>Indicates the physical Slot Number attached to this port. If the PCI Express Capability register <i>Slot Implemented</i> 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.</p>	Downstream	HwInit	Yes	0-0h

Register 13-31. 80h Slot Status and Control (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
<i>Note:</i> Hot Plug-capable ports are Ports 1, 2, and 3.					
Slot Control					
0	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Attention Button Pressed Enable 0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Device PM state (Power Management Status and Control register <i>Power State</i> field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Device PM state (offset 44h[1:0]=11b), for an Attention Button Pressed event on the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	RW	Yes	0
1	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Power Fault Detector Enable 0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Device PM state (Power Management Status and Control register <i>Power State</i> field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Device PM state (offset 44h[1:0]=11b), for a Power Fault Detected event on the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	RW	Yes	0

Register 13-31. 80h Slot Status and Control (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
2	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	MRL Sensor Changed Enable 0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Device PM state (Power Management Status and Control register <i>Power State</i> field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Device PM state (offset 44h[1:0]=11b), for an MRL Sensor Changed event on the corresponding PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	RW	Yes	0
3	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Presence Detect Changed Enable 0 = Function is disabled 1 = Enables software notification with a Hot Plug interrupt if the port is in the D0 Device PM state (Power Management Status and Control register <i>Power State</i> field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Device PM state (offset 44h[1:0]=11b), for a Presence Detect Changed event on the corresponding downstream port A Presence Detect Changed event is triggered by either the SerDes Receiver Detect on the corresponding PEX 8509 downstream port, or by HP_PRSENT _x # input on the corresponding PEX 8509 Hot Plug-capable downstream port.	Downstream	RW	Yes	0

Register 13-31. 80h Slot Status and Control (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
4	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Command Completed Interrupt Enable 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 PEX 8509 Hot Plug-capable downstream port	Downstream Hot Plug-capable	RW	Yes	0
5	<i>Reserved</i>	Upstream	RsvdP	No	0
	Hot Plug Interrupt Enable 0 = Function is disabled 1 = Enables a Hot Plug interrupt on enabled Hot Plug/Link State events for the corresponding PEX 8509 downstream port	Downstream	RW	Yes	0
7:6	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	00b (Upstream) 11b (Downstream)
	Attention Indicator Controls Controls the Attention Indicator on the corresponding PEX 8509 downstream port slot. 00b = <i>Reserved</i> – 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_ATNLEDx# Output signal. Reads return the corresponding PEX 8509 Hot Plug-capable downstream port Attention Indicator's current state.	Downstream Hot Plug-capable	RW	Yes	11b

Register 13-31. 80h Slot Status and Control (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
9:8	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	00b (Upstream) 11b (Downstream)
	<p>Power Indicator Control Controls the Power Indicator on the corresponding PEX 8509 downstream port slot.</p> <p>00b = <i>Reserved</i> – Writes are ignored 01b = Turns On indicator to constant On state 10b = Causes indicator to blink 11b = Turns Off indicator</p> <p>Software must use a Byte or Word Write (and not a DWord Write) to control the HP_PWRLEDx# Output signal.</p> <p>Reads return the corresponding PEX 8509 Hot Plug-capable downstream port Power Indicator's current state.</p>	Downstream Hot Plug-capable	RW	Yes	11b (MRL open) 01b (MRL closed)
10	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	<p>Power Controller Control Controls the Power Controller on the corresponding PEX 8509 Hot Plug-capable downstream port slot.</p> <p>0 = Turns On the Power Controller; requires some delay to be effective 1 = Turns Off the Power Controller</p> <p>Software must use a Byte or Word Write (and not a DWord Write) to control the Power Controller Output signals.</p>	Downstream Hot Plug-capable	RW	Yes	1 (MRL open) 0 (MRL closed)
11	Electromechanical Interlock Control <i>Not supported</i>		RsvdP	No	0
12	Not valid for the upstream port.	Upstream	RsvdP	No	0
	<p>Data Link Layer State Changed Enable Enables software notification with a Hot Plug interrupt if the port is in the D0 Device PM state (Power Management Status and Control register <i>Power State</i> field, offset 44h[1:0]=00b), or with a PME message if the port is in the D3hot Device PM state (offset 44h[1:0]=11b), when the Link Status register <i>Data Link Layer Link Active</i> bit (offset 78h[29]) is changed.</p>	Downstream	RW	Yes	0
15:13	<i>Reserved</i>		RsvdP	No	000b

Register 13-31. 80h Slot Status and Control (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
Slot Status					
16	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Attention Button Pressed Set to 1 when the Attention Button of the corresponding PEX 8509 Hot Plug-capable downstream port slot is pressed.	Downstream Hot Plug-capable	RWC	Yes	0
17	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Power Fault Detected Set to 1 when the Power Controller of the corresponding PEX 8509 Hot Plug-capable downstream port slot detects a Power Fault at the slot.	Downstream Hot Plug-capable	RWC	Yes	0
18	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	MRL Sensor Changed Set to 1 when an MRL Sensor state change is detected on the corresponding PEX 8509 Hot Plug-capable downstream port slot.	Downstream Hot Plug-capable	RWC	Yes	0
19	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Presence Detect Changed Set to 1 when the value reported in bit 22 (<i>Presence Detect State</i>) changes. Write 1 to clear. A Presence Detect Changed event is triggered by either the SerDes Receiver Detect on the corresponding PEX 8509 downstream port, or by HP_PRSENT _x # input on the corresponding PEX 8509 Hot Plug-capable downstream port.	Downstream	RWC	Yes	0

Register 13-31. 80h Slot Status and Control (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
20	<i>Reserved</i>	Upstream; Downstream Non-Hot Plug-capable	RsvdP	No	0
	Command Completed Set to 1 when the Hot Plug Controller on the corresponding PEX 8509 Hot Plug-capable downstream port slot completes an issued command to: <ul style="list-style-type: none"> • <i>Attention Indicator Controls</i> (field [7:6]) • <i>Power Indicator Control</i> (field [9:8]) • <i>Power Controller Control</i> (bit 10) 	Downstream Hot Plug-capable	RWC	Yes	0
21	<i>Reserved</i>	0, 4, 5, 6, 7	RsvdP	No	0
	MRL Sensor State Reveals the corresponding PEX 8509 Hot Plug-capable downstream port MRL Sensor's current state. 0 = MRL Sensor is closed 1 = MRL Sensor is open	1, 2, 3	RO	No	0
22	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Presence Detect State Indicates the presence of a downstream device, reflected by the logical OR of the corresponding downstream port's SerDes Receiver Detect, and, if present, the port's HP_PRSN _{Tx} # input (debounced). 0 = Slot is empty, or device is not present 1 = Slot is occupied, or device is present	Downstream	RO	No	0
23	Electromechanical Interlock Status <i>Not supported</i>		RsvdZ	No	0
24	Not valid for the upstream port.	Upstream	RsvdP	No	0
	Data Link Layer State Changed Set when the value reported in the Link Status register <i>Data Link Layer Link Active</i> bit changes. In response to a Data Link Layer State Changed event, software must read the Link Status register <i>Data Link Layer Link Active</i> bit to determine whether the link is active before initiating Configuration requests to the device.	Downstream	RWC	Yes	0
31:25	<i>Reserved</i>		RsvdZ	No	0-0h

13.10 Subsystem ID and Subsystem Vendor ID Capability Registers

This section details the PEX 8509 Subsystem ID and Subsystem Vendor ID Capability registers. Table 13-10 defines the register map.

Table 13-10. PEX 8509 Subsystem ID and Subsystem Vendor ID 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	
<i>Reserved</i>	Next Capability Pointer (DCh)	SSID/SSVID Capability ID (0Dh)	90h
Subsystem ID	Subsystem Vendor ID		94h

Register 13-32. 90h Subsystem Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	SSID/SSVID Capability ID Detects the SSID/SSVID registers for the PCI-to-PCI bridge. Set to 0Dh, as required by the <i>PCI-to-PCI Bridge r1.2</i> .	RO	Yes	0Dh
15:8	Next Capability Pointer Set to DCh, to point to the Vendor-Specific Enhanced Capability register.	RO	Yes	DCh
31:16	<i>Reserved</i>	RsvdP	No	0000h

Register 13-33. 94h Subsystem ID and Subsystem Vendor ID (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	Subsystem Vendor ID The Vendor ID (offset 00h[15:0]) identifies the manufacturer of the PEX 8509 switch, and the Subsystem Vendor ID optionally identifies the board or system vendor. As with the Vendor ID value, the <i>Subsystem Vendor ID</i> value must be a valid PCI-SIG-assigned Vendor ID.	RO	Yes	10B5h
31:16	Subsystem ID The Device ID (offset 00h[31:16]) identifies the PEX 8509 switch, and optionally the Subsystem ID in combination with the Subsystem Vendor ID, uniquely identifies the board or system. The <i>Subsystem ID</i> value is chosen or assigned only by the “owner” of the valid Vendor ID value used for the Subsystem Vendor ID. If the board or system vendor is not a PCI-SIG member, PLX can assign, free of charge, a unique <i>Subsystem ID</i> value, in which case the <i>Subsystem Vendor ID</i> remains the PLX default value, 10B5h. The <i>Subsystem Vendor ID</i> and <i>Subsystem ID</i> values are usually identical for all PEX 8509 ports.	RO	Yes	8509h

13.11 Vendor-Specific Enhanced Capability Registers

This section details the PEX 8509 Vendor-Specific Enhanced Capability registers. Table 13-11 defines the register map.

Table 13-11. Vendor-Specific Enhanced 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	
<i>Reserved</i>	Vendor-Specific Enhanced Capability	Next Capability Pointer (00h)	SSID/SSVID Capability ID (09h)
Hardwired Device ID		Hardwired Vendor ID	
<i>Reserved</i>		PLX Hardwired Revision ID	
Scratch A			
Scratch B			
<i>Reserved</i>		F0h –	
			DCh E0h E4h E8h ECh FCh

Register 13-34. DCh Vendor-Specific Enhanced Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Vendor-Specific Capability ID Set to 09h, indicating that the Capability structure is the Vendor-Specific Enhanced Capability structure.	RO	Yes	09h
15:8	Next Capability Pointer 00h = Vendor-Specific Enhanced Capability is the last capability in the PEX 8509 port Capabilities list The PEX 8509 port Extended Capabilities list starts at offset 100h.	RO	Yes	00h
23:16	Number of Bytes in this Capability	RO	Yes	14h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-35. E0h PLX Hardwired Configuration ID (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	Hardwired Vendor ID Always returns the PLX PCI-SIG-assigned Vendor ID value, 10B5h.	RO	No	10B5h
31:16	Hardwired Device ID Always returns the PEX 8509 default Device ID value, 8509h.	RO	No	8509h

Register 13-36. E4h PLX Hardwired Revision ID (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Hardwired Revision ID Always returns the PEX 8509 default Revision ID value, AAh.	RO	No	Current Rev # (AAh)
31:8	<i>Reserved</i>	RsvdP	No	0000_00h

Register 13-37. E8h Scratch A (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Scratch A General-purpose register.	RW	Yes	0000_0000h

Register 13-38. ECh Scratch B (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Scratch B General-purpose register.	RW	Yes	0000_0000h

13.12 Device Serial Number Extended Capability Registers

This section details the PEX 8509 Device Serial Number Extended Capability registers. [Table 13-12](#) defines the register map.

Table 13-12. PEX 8509 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
<i>Reserved</i>			10Ch – 134h

Register 13-39. 100h Device Serial Number Enhanced Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	PCI Express Extended Capability ID Set to 0003h, as required by the <i>PCI Express Base r1.1</i> .	RO	Yes	0003h
19:16	Capability Version Set to 1h, as required by the <i>PCI Express Base r1.1</i> .	RO	Yes	1h
31:20	Next Capability Offset Set to FB4h, which is the Advanced Error Reporting Enhanced Capability structure.	RO	Yes	FB4h

Register 13-40. 104h Serial Number (Lower DW) (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	<p>PCI Express Device Serial Number (1st DW)</p> <p>Lower half of a 64-bit register. Value set by Serial EEPROM register initialization. Per the <i>PCI Express Base r1.1</i>, all switch ports must contain the same value.</p> <p>The Serial Number registers contain the IEEE-defined 64-bit Extended Unique Identifier (EUI-64™). The lower 24 bits are the Company ID value assigned by the IEEE registration authority, and the upper 40 bits are the Extension ID assigned by the identified Company.</p>	RO	Yes	B5DF_0E00h

Register 13-41. 108h Serial Number (Upper DW) (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	<p>PCI Express Device Serial Number (2nd DW)</p> <p>Upper half of a 64-bit register. Value set by Serial EEPROM register initialization. Per the <i>PCI Express Base r1.1</i>, all switch ports must contain the same value.</p> <p>The Serial Number registers contain the IEEE-defined 64-bit Extended Unique Identifier (EUI-64™). The lower 24 bits are the Company ID value assigned by the IEEE registration authority, and the upper 40 bits are the Extension ID assigned by the identified Company.</p>	RO	Yes	AA_8509_10h

13.13 Power Budget Extended Capability Registers

This section details the PEX 8509 Power Budget Extended Capability registers. These registers work a bit differently than all the others, especially with respect to serial EEPROM Writes. When read, the register to be used is selected by writing to the **Data Select** register *Data Select* field (offset 13Ch[7:0]), then reading the indexed register in the **Power Budget Data** register. If these registers are to be used, the serial EEPROM must write values to them. Instead of the serial EEPROM writing the index and then the data sequence, the index and data sequence are merged into one Write at the **Data Select** register, with select in bits [7:0] and the actual data in bits [28:8].

Table 13-13 defines the register map.

Table 13-13. Power Budget 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
<i>Reserved</i>		Data Select	13Ch
Power Budget Data			140h
Power Budget Capability			144h

Register 13-42. 138h Power Budget Extended Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	PCI Express Extended Capability ID Set to 0004h, as required by the <i>PCI Express Base r1.1</i> .	RO	Yes	0004h
19:16	Capability Version Set to 1h, as required by the <i>PCI Express Base r1.1</i> .	RO	Yes	1h
31:20	Next Capability Offset Set to 148h, which addresses the PEX 8509 Virtual Channel Extended Capability structure.	RO	Yes	148h

Register 13-43. 13Ch Data Select (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Data Select Indexes the Power Budget Data reported, by way of eight Power Budget Data registers per port, and selects the DWord of Power Budget data that appears in each Power Budget Data register. Index values start at 0, to select the first DWord of Power Budget data; subsequent DWords of Power Budget data are selected by increasing index values 1 to 7.	RW	Yes	00h
31:8	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-44. 140h Power Budget Data (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<p>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]) contents.</p>				
7:0	<p>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, to produce the actual power consumption value.</p>	RO	Yes	00h
9:8	<p>Data Scale Specifies the scale to apply to the Base Power value. The device power consumption is determined by multiplying the <i>Base Power</i> field contents with the value corresponding to the encoding returned by this field.</p> <p>00b = 1.0x 01b = 0.1x 10b = 0.01x 11b = 0.001x</p>	RO	Yes	00b
12:10	<p>PM Sub-State 000b = Corresponding PEX 8509 port is in the default Power Management sub-state</p>	RO	Yes	000b
14:13	<p>PM State Current Device Power Management (PM) state.</p> <p>00b = D0 Device PM state 11b = D3 Device PM state All other encodings are <i>reserved</i>.</p>	RO	Yes	00b
17:15	<p>Type Type of operating condition.</p> <p>000b = PME Auxiliary 001b = Auxiliary 010b = Idle 011b = Sustained 111b = Maximum All other encodings are <i>reserved</i>.</p>	RO	Yes	000b
20:18	<p>Power Rail Power Rail of operating condition.</p> <p>000b = Power 12V 001b = Power 3.3V 010b = Power 1.8V 111b = Thermal All other encodings are <i>reserved</i>.</p>	RO	Yes	000b
31:21	Reserved	RsvdP	No	0-0h

Register 13-45. 144h Power Budget Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	System Allocated 1 = Power budget for the device is included within the system power budget	HwInit	Yes	1
31:1	<i>Reserved</i>	RsvdP	No	0-0h

13.14 Virtual Channel Extended Capability Registers

This section details the PEX 8509 Virtual Channel Extended Capability registers. These registers are duplicated for each port. Table 13-14 defines the register map for one port.

Table 13-14. PEX 8509 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
<i>Reserved</i>			150h
Port VC Status (<i>Reserved</i>)		Port VC Control	154h
VC0 Resource Capability			158h
VC0 Resource Control			15Ch
VC0 Resource Status		<i>Reserved</i>	160h
<i>Reserved</i>			164h – 1A4h

Register 13-46. 148h Virtual Channel Extended Capability (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	PCI Express Extended Capability ID Set to 0002h, as required by the <i>PCI Express Base r1.1</i> .	RO	No	0002h
19:16	Capability Version Set to 1h, as required by the <i>PCI Express Base r1.1</i> .	RO	No	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	No	000h

Register 13-47. 14Ch Port VC Capability 1 (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
0	Extended VC Count		RsvdP	No	0
3:1	<i>Reserved</i>		RsvdP	No	000b
4	Low-Priority Extended VC Count		RsvdP	No	0
7:5	<i>Reserved</i>		RsvdP	No	000b
9:8	Reference Clock Cleared to 00b.		RsvdP	No	00b
11:10	Port Arbitration Table Entry Size 10b = Port Arbitration Table entry size is 4 bits	Upstream	RO	Yes	10b
	<i>Reserved</i>	Downstream	RsvdP	No	00b
31:12	<i>Reserved</i>		RsvdP	No	0-0h

Register 13-48. 154h Port VC Status and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Port VC Control				
0	Load VC Arbitration Table Writing 1 updates the VC Arbitration Table for the corresponding PEX 8509 port. Reads always return 0.	RsvdP	No	0
3:1	VC Arbitration Select 000b = Bit 0; Round-Robin (Hardware-Fixed) arbitration scheme All other encodings are <i>reserved</i>	RW	Yes	000b
15:4	<i>Reserved</i>	RsvdP	No	0-0h
Port VC Status				
16	VC Arbitration Table Status <i>Reserved</i>	RsvdP	No	0
31:17	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-49. 158h VC0 Resource Capability (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
0	Port Arbitration Capability 1 = Non-configurable Round-Robin (Hardware-Fixed) Arbitration		RO	Yes	1
1	Port Arbitration Capability 1 = Weighted Round-Robin (WRR) arbitration with 32 phases	Upstream	RO	Yes	1
	Not valid for the downstream ports.	Downstream	RsvdP	No	0
13:2	Reserved		RsvdP	No	0-0h
14	Advanced Packet Switching		RsvdP	No	0
15	Reject Snoop Transactions Not a PCI Express switch feature; therefore, cleared to 0.		RsvdP	No	0
22:16	Maximum Time Slots Cleared to 000_0000b.		RsvdP	No	000_0000b
23	Reserved		RsvdP	No	0
31:24	Port Arbitration Table Offset Offset of the Port Arbitration Table, as the number of DQWords from the Base address of the Virtual Channel Extended Capability structure. 00h = Port Arbitration Table is not present 06h = Port Arbitration Table is located at register offset 1A8h (148h + 6 * 4 DWords) <i>Note: For further details, refer to Section 13.15.</i>	Upstream	RO	Yes	06h
	Not valid for the downstream ports.	Downstream	RsvdP	No	00h

Register 13-50. 15Ch VC0 Resource Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC/VC0 Map Defines Traffic Classes [7:0], respectively, and indicates which TCs are mapped into Virtual Channel 0.	RO	No	1
7:1	Traffic Class 0 (TC0) must be mapped to Virtual Channel 0. By default, Traffic Classes [7:1] are mapped to VC0.	RW	Yes	7Fh
15:8	Reserved	RsvdP	No	00h
16	Load Port Arbitration Table Read always returns 0.	RW	Yes	0

Register 13-50. 15Ch VC0 Resource Control (All Ports) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
19:17	Port Arbitration Select Selects the Port Arbitration type for the corresponding PEX 8509 port. Indicates the bit number in the VC0 Resource Capability register <i>Port Arbitration Capability</i> field (offset 158h[1:0]) that corresponds to the arbitration type. 0 = Round-Robin (Hardware-Fixed) arbitration scheme 1 = Weighted Round-Robin with 32 Phases	RW	Yes	000b
23:20	<i>Reserved</i>	RsvdP	No	0-0h
24	VC0 ID Defines the corresponding PEX 8509 port Virtual Channel 0 ID code. Cleared to 0, because there is only one Virtual Channel.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	0-0h
31	VC0 Enable 0 = Not allowed 1 = Enables the corresponding PEX 8509 port Virtual Channel 0	RO	No	1

Register 13-51. 160h VC0 Resource Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	<i>Reserved</i>	RsvdP	No	0000h
16	Port Arbitration Table Status 0 = Hardware has finished loading values stored in the Port Arbitration Table after software sets the VC0 Resource Control register <i>Load Port Arbitration Table</i> bit (offset 15Ch[16]=1) 1 = An entry of the Port Arbitration Table was written to by software	RO	No	0
17	VC0 Negotiation Pending 0 = VC0 negotiation is complete 1 = VC0 initialization is not complete for the corresponding PEX 8509 port	RO	Yes	1
31:18	<i>Reserved</i>	RsvdP	No	0-0h

13.15 Port Arbitration Table Registers

This section details the PEX 8509 Port Arbitration Table registers. Port Arbitration Table phases are used to determine port weighting during “Weighted Round-Robin with 32 Phases” port arbitration. [Table 13-15](#) defines the register map.

Note: The Port Arbitration Table is used only when Weighted Round-Robin with 32-phase Port Arbitration is selected, by way of the VC0 Resource Control register Port Arbitration Select field (offset 15Ch[19:17]).

Table 13-15. Port Arbitration Table Register Map (Only Upstream Port)

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		
Phase 7	Phase 6	Phase 5	Phase 4	Phase 3	Phase 2	Phase 1	Phase 0																										1A8h
Phase 15	Phase 14	Phase 13	Phase 12	Phase 11	Phase 10	Phase 9	Phase 8																										1ACh
Phase 23	Phase 22	Phase 21	Phase 20	Phase 19	Phase 18	Phase 17	Phase 16																										1B0h
Phase 31	Phase 30	Phase 29	Phase 28	Phase 27	Phase 26	Phase 25	Phase 24																										1B4h
<i>Reserved</i>																												1B8h –	1BCh				

**Register 13-52. 1A8h Port Arbitration Table Phases 0 to 7
(Only Upstream Port, Reserved (RsvdP) for All Other Ports)**

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
<i>Note:</i> The fields within this register are valid only on the upstream port, and reserved on the downstream ports.					
3:0	Port Arbitration Table Phase 0	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
7:4	Port Arbitration Table Phase 1	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
11:8	Port Arbitration Table Phase 2	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
15:12	Port Arbitration Table Phase 3	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
19:16	Port Arbitration Table Phase 4	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
23:20	Port Arbitration Table Phase 5	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
27:24	Port Arbitration Table Phase 6	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
31:28	Port Arbitration Table Phase 7	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h

**Register 13-53. 1ACh Port Arbitration Table Phases 8 to 15
(Only Upstream Port, Reserved (RsvdP) for All Other Ports)**

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
<i>Note: The fields within this register are valid only on the upstream port, and reserved on the downstream ports.</i>					
3:0	Port Arbitration Table Phase 8	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
7:4	Port Arbitration Table Phase 9	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
11:8	Port Arbitration Table Phase 10	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
15:12	Port Arbitration Table Phase 11	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
19:16	Port Arbitration Table Phase 12	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
23:20	Port Arbitration Table Phase 13	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
27:24	Port Arbitration Table Phase 14	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
31:28	Port Arbitration Table Phase 15	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h

**Register 13-54. 1B0h Port Arbitration Table Phases 16 to 23
(Only Upstream Port, Reserved (RsvdP) for All Other Ports)**

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
<i>Note:</i> The fields within this register are valid only on the upstream port, and reserved on the downstream ports.					
3:0	Port Arbitration Table Phase 16	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
7:4	Port Arbitration Table Phase 17	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
11:8	Port Arbitration Table Phase 18	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
15:12	Port Arbitration Table Phase 19	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
19:16	Port Arbitration Table Phase 20	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
23:20	Port Arbitration Table Phase 21	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
27:24	Port Arbitration Table Phase 22	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h
31:28	Port Arbitration Table Phase 23	Upstream (refer to Note)	RW	Yes	0h
	Reserved	Downstream	RsvdP	No	0h

**Register 13-55. 1B4h Port Arbitration Table Phases 24 to 31
(Only Upstream Port, Reserved (RsvdP) for All Other Ports)**

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
<i>Note: The fields within this register are valid only on the upstream port, and reserved on the downstream ports.</i>					
3:0	Port Arbitration Table Phase 24	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
7:4	Port Arbitration Table Phase 25	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
11:8	Port Arbitration Table Phase 26	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
15:12	Port Arbitration Table Phase 27	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
19:16	Port Arbitration Table Phase 28	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
23:20	Port Arbitration Table Phase 29	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
27:24	Port Arbitration Table Phase 30	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h
31:28	Port Arbitration Table Phase 31	Upstream (refer to Note)	RW	Yes	0h
	<i>Reserved</i>	Downstream	RsvdP	No	0h

13.16 Device-Specific Registers

Table 13-16 defines Device-Specific registers – registers that are unique to the PEX 8509 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 13-16. 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	
Device-Specific Registers – Error Checking and Debug				1C0h ... 1FCh
Device-Specific Registers – Physical Layer				200h ... 28Ch
Device-Specific Registers – I ² C Interface				290h ... 2C4h
Device-Specific Registers – Bus Number CAM				2C8h ... 304h
Device-Specific Registers – I/O CAM				308h ... 344h
Device-Specific Registers – Address-Mapping CAM				348h ... 548h
<i>Reserved</i>				54Ch – 63Ch
<i>Factory Test Only</i>				640h – 65Ch
Device-Specific Registers – Ingress Control and Port Enable				660h ... 67Ch
Device-Specific Registers – I/O CAM Base and Limit Upper 16 Bits				680h ... 6BCh
Device-Specific Registers – Base Address Shadow				6C0h ... 73Ch

Table 13-16. 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	
Device-Specific Registers – Shadow Virtual Channel Capability				740h ...
<i>Reserved</i>				83Ch 840h – 93Ch
Device-Specific Registers – Ingress Credit Handler				940h ...
<i>Reserved</i>				B80h B84h – DFCh
Device-Specific Registers – Port Configuration Header				E00h ...
Device-Specific Registers – Source Queue Weight and Soft Error				E3Ch F10h ...
				FB0h

13.16.1 Device-Specific Registers – Error Checking and Debug

Table 13-17. 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	
Device-Specific Error Status for Egress ECC Error				1C0h
Device-Specific Error Mask for Egress ECC Error				1C4h
ECC Error Check Disable				1C8h
Error Handler 32-Bit Error Status (<i>Factory Test Only</i>)				1CCh
Error Handler 32-Bit Error Mask (<i>Factory Test Only</i>)				1D0h
<i>Factory Test Only</i>			1D4h –	1D8h
Debug Control				1DCh
<i>Reserved</i>		Power Management Hot Plug User Configuration		1E0h
<i>Factory Test Only</i>				1E4h
Bad TLP Count				1E8h
Bad DLLP Count				1ECh
<i>Reserved</i>				1F0h
<i>Reserved</i>			Lane Status/Software PEX_LANE_GOODx# LED Control	1F4h
ACK Transmission Latency Limit				1F8h
<i>Factory Test Only</i>				1FCh

a. Certain registers are port-specific, some are station-specific, and some are chip-specific; all are device-specific.

Register 13-56. 1C0h Device-Specific Error Status for Egress ECC Error (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note: The bits in this register can be masked by their respective bits in the Device-Specific Error Mask for Egress ECC Error register (Port 0, offset 1C4h).</i>				
0	Packet Link List RAM 1-Bit Soft Error Tracking Counter Overflow Detected 0 = No overflow detected 1 = 1-bit Soft Error Tracking Counter overflow detected	RWCS	Yes	0
5:1	<i>Reserved</i>	RsvdP	No	00h
6	Packet Link List RAM Read Detected 2-Bit Soft Error 0 = No error detected 1 = Read detected 2-Bit Soft error	RWCS	Yes	0
11:7	<i>Reserved</i>	RsvdP	No	00h
12	Packet RAM 1-Bit Soft Error Tracking Counter Overflow Detected 0 = No overflow detected 1 = 1-bit Soft Error Tracking Counter overflow detected	RWCS	Yes	0
20:13	<i>Reserved</i>	RsvdP	No	00h
21	Packet RAM Read Detected 2-Bit Soft Error 0 = No error detected 1 = Read detected 2-Bit Soft error	RWCS	Yes	0
31:22	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-57. 1C4h Device-Specific Error Mask for Egress ECC Error (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note:</i> The bits in this register can be used to mask their respective bits in the <i>Device-Specific Error Status for Egress ECC Error</i> register (Port 0, offset 1C0h).				
0	Packet Link List RAM 1-Bit Soft Error Tracking Counter Overflow Mask 0 = No effect on reporting activity 1 = <i>Packet Link List RAM 1-Bit Soft Error Tracking Counter Overflow Detected</i> bit is masked/disabled	RWS	Yes	1
5:1	<i>Reserved</i>	RsvdP	No	
6	Packet Link List RAM Read Detected 2-Bit Soft Error Mask 0 = No effect on reporting activity 1 = <i>Packet Link List RAM Read Detected 2-Bit Soft Error</i> bit is masked/disabled	RWS	Yes	1
11:7	<i>Reserved</i>	RsvdP	No	00h
12	Packet RAM 1-Bit Soft Error Tracking Counter Overflow Mask 0 = No effect on reporting activity 1 = <i>Packet RAM 1-Bit Soft Error Tracking Counter Overflow Detected</i> bit is masked/disabled	RWS	Yes	1
20:13	<i>Reserved</i>	RsvdP	No	00h
21	Packet RAM Read Detected 2-Bit Soft Error Mask 0 = No effect on reporting activity 1 = <i>Packet RAM Read Detected 2-Bit Soft Error</i> bit is masked/disabled	RWS	Yes	1
31:22	<i>Reserved</i>	RsvdP	No	000h

Register 13-58. 1C8h ECC Error Check Disable (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	ECC 1-Bit Error Check Disable 0 = RAM 1-Bit Soft Error Check enabled 1 = Disables RAM 1-Bit Soft Error Check	RWS	Yes	0
1	ECC 2-Bit Error Check Disable 0 = RAM 2-Bit Soft Error Check enabled 1 = Disables RAM 2-Bit Soft Error Check	RWS	Yes	0
2	Software Force Error Enable 1 = Correctable Error Status and Uncorrectable Error Status registers (offsets FC4h and FB8h , respectively) change from RWCS to RW	RWS	Yes	0
3	Software Force Non-Posted Request 1 = Enables handling of errors associated with Posted TLPs as if those errors are associated with Non-Posted TLPs	RWS	Yes	0
4	Enable PEX_INTA# Ball for Hot Plug or Link State Event 0 = Hot Plug or Link State Event Interrupt requests send an INT _x message (and do not assert PEX_INTA#) 1 = Hot Plug or Link State Event Interrupt requests assert PEX_INTA# (and do not send an INT _x message)	RWS	Yes	0
5	Enable PEX_INTA# Ball for Device-Specific Error 0 = Device-Specific Error Interrupt requests send an INT _x message (and do not assert PEX_INTA#) 1 = Device-Specific Error Interrupt requests assert PEX_INTA# (and do not send an INT _x message)	RWS	Yes	0
31:6	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-59. 1CCh Error Handler 32-Bit Error Status (Factory Test Only)

Bit(s)	Description	Bit Exists Only on Ports	Type	Serial EEPROM and I ² C	Default
<i>Notes: All errors in this register generate MSI/INTx interrupts, if enabled.</i>					
<i>The bits in this register can be masked by their respective bits in the Error Handler 32-Bit Error Mask register (offset 1D0h).</i>					
0	Completion FIFO Overflow Status 0 = No overflow is detected 1 = Completion FIFO overflow is detected when 4-deep Completion FIFO for ingress, or 1-deep Completion FIFO for egress, overflows	0, 1, 2, 3, 4, 5, 6, 7	RWCS	Yes	0
1	Reserved		RsvdP	No	0
2	Factory Test Only		RWCS	No	0
3	Reserved		RsvdP	No	0
4	Destination Queue Link List RAM 2-Bit Error 0 = No error is detected 1 = Destination Queue Link List RAM 2-bit error is detected	0	RWCS	Yes	0
5	Destination Queue Link List RAM 1-Bit Error Counter Overflow Detected 0 = No error is detected 1 = Destination Queue Link List RAM 1-bit error is detected	0	RWCS	Yes	0
6	64 Entry Retry Buffer RAM 1-Bit Error Counter Overflow Detected 0 = No error is detected 1 = 64 Entry Retry Buffer RAM 1-bit error is detected	0	RWCS	Yes	0
7	64 Entry Retry Buffer RAM 2-Bit Error 0 = No error is detected 1 = 64 Entry Retry Buffer RAM 2-bit error is detected	0	RWCS	Yes	0
8	Source Queue Link List RAM 1-Bit Error Counter Overflow 0 = No error is detected 1 = Source Queue Link List RAM 1-bit error is detected	0	RWCS	Yes	0
9	Source Queue Link List RAM 2-Bit Error 0 = No error is detected 1 = Source Queue Link List RAM 2-bit error is detected	0	RWCS	Yes	0
10	32 Entry Retry Buffer 1-Bit Error Counter Overflow 0 = No error is detected 1 = 32 Entry Retry Buffer 1-bit error is detected	0	RWCS	Yes	0
11	32 Entry Retry Buffer 2-Bit ECC Error 0 = No error is detected 1 = 32 Entry Retry Buffer 2-bit error is detected	0	RWCS	Yes	0
12	Scheduler RAM 2-Bit ECC Error 0 = No error is detected 1 = Scheduler RAM 2-bit ECC error is detected	0	RWCS	Yes	0
15:13	Reserved		RsvdP	No	000b

Register 13-59. 1CCh Error Handler 32-Bit Error Status (Factory Test Only) (Cont.)

Bit(s)	Description	Bit Exists Only on Ports	Type	Serial EEPROM and I ² C	Default
16	Scheduler RAM 1-Bit ECC Error Detected 0 = No error is detected 1 = Scheduler RAM 1-bit ECC error is detected	0	RWCS	Yes	0
19:17	<i>Reserved</i>		RsvdP	No	000b
20	TLP ID RAM 2-Bit ECC Error 0 = No error is detected 1 = TLP ID RAM 2-bit ECC error is detected	0	RWCS	Yes	0
23:21	<i>Reserved</i>		RsvdP	No	000b
24	TLP ID RAM 1-Bit ECC Error Counter Overflow 0 = No error is detected 1 = TLP ID RAM 1-bit ECC Error Counter overflow is detected	0	RWCS	Yes	0
31:25	<i>Reserved</i>		RsvdP	No	0-0h

Register 13-60. 1D0h Error Handler 32-Bit Error Mask (Factory Test Only)

Bit(s)	Description	Bit Exists Only on Ports	Type	Serial EEPROM and I ² C	Default
<i>Notes: Error logging is enabled in this register, by default.</i>					
<i>The bits in this register can be used to mask their respective bits in the Error Handler 32-Bit Error Status register (offset 1CCh).</i>					
0	Completion FIFO Overflow Mask 0 = If enabled, error generates MSI/INTx interrupt 1 = <i>Completion FIFO Overflow Status</i> bit is masked/disabled	0, 1, 2, 3, 4, 5, 6, 7	RWS	Yes	1
1	Reserved		RWS	Yes	1
2	Factory Test Only		RWS	Yes	1
3	Reserved		RsvdP	No	0
4	Destination Queue Link List RAM 2-Bit Error Mask 0 = No effect on reporting activity 1 = <i>Destination Queue Link List RAM 2-Bit Error</i> bit is masked/disabled	0	RWS	Yes	1
5	Destination Queue Link List RAM 1-Bit Error Counter Overflow Detected Mask 0 = No effect on reporting activity 1 = <i>Destination Queue Link List RAM 1-Bit Error Counter Overflow Detected</i> bit is masked/disabled	0	RWS	Yes	1
6	64 Entry Retry Buffer RAM 1-Bit Error Counter Overflow Detected Mask 0 = No effect on reporting activity 1 = <i>64 Entry Retry Buffer RAM 1-Bit Error Counter Overflow Detected</i> bit is masked/disabled	0	RWS	Yes	1
7	64 Entry Retry Buffer RAM 2-Bit Error Mask 0 = No effect on reporting activity 1 = <i>64 Entry Retry Buffer RAM 2-Bit Error</i> bit is masked/disabled	0	RWS	Yes	1

Register 13-60. 1D0h Error Handler 32-Bit Error Mask (Factory Test Only) (Cont.)

Bit(s)	Description	Bit Exists Only on Ports	Type	Serial EEPROM and I ² C	Default
8	Source Queue Link List RAM 1-Bit Error Counter Overflow Mask 0 = No effect on reporting activity 1 = <i>Source Queue Link List RAM 1-Bit Error Counter Overflow</i> bit is masked/disabled	0	RWS	Yes	1
9	Source Queue Link List RAM 2-Bit Error Mask 0 = No effect on reporting activity 1 = <i>Source Queue Link List RAM 2-Bit Error</i> bit is masked/disabled	0	RWS	Yes	1
10	32 Entry Retry Buffer 1-Bit Error Counter Overflow Mask 0 = No effect on reporting activity 1 = <i>32 Entry Retry Buffer 1-Bit Error Counter Overflow</i> bit is masked/disabled	0	RWS	Yes	1
11	32 Entry Retry Buffer 2-Bit ECC Error Mask 0 = No effect on reporting activity 1 = <i>32 Entry Retry Buffer 2-Bit ECC Error</i> bit is masked/disabled	0	RWS	Yes	1
12	Scheduler RAM 2-Bit ECC Error Mask 0 = No effect on reporting activity 1 = <i>Scheduler RAM 2-Bit ECC Error</i> bit is masked/disabled	0	RWS	Yes	1
15:13	<i>Reserved</i>		RsvdP	No	000b
16	Scheduler RAM 1-Bit ECC Error Counter Overflow Mask 0 = No effect on reporting activity 1 = <i>Scheduler RAM 1-Bit ECC Error Detected</i> bit is masked/disabled	0	RWS	Yes	1
19:17	<i>Reserved</i>		RsvdP	No	000b
20	TLP ID RAM 2-Bit ECC Error Mask 0 = No effect on reporting activity 1 = <i>TLP ID RAM 2-Bit ECC Error</i> bit is masked/disabled	0	RWS	Yes	1
23:21	<i>Reserved</i>		RsvdP	No	000b
24	TLP ID RAM 1-Bit ECC Error Counter Overflow Mask 0 = No effect on reporting activity 1 = <i>TLP ID RAM 1-Bit ECC Error Counter Overflow</i> bit is masked/disabled	0	RWS	Yes	1
31:25	<i>Reserved</i>		RsvdP	No	0-0h

Register 13-61. 1DCh Debug Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	<i>Factory Test Only</i>	RO	No	1111b
6:4	<i>Reserved</i>	RsvdP	No	000b
7	<i>Factory Test Only</i>	RWS	Yes	0
11:8	Upstream Port ID Upstream Port Number – Reads the external Strap value on the STRAP_UPSTRM_PORTSEL[2:0] balls, at Reset de-assertion. When bit 15 (<i>Software Configuration Control</i>) is cleared to 0, software is not allowed to change this value.	RO	Yes	Set by Strapping ball levels
	When bit 15 (<i>Software Configuration Control</i>) is set to 1, Upstream Port Number can be set by software. 0h = Port 0 (recommended) 1h = Port 1 2h = Port 2 3h = Port 3 4h = Port 4 5h = Port 5 6h = Port 6 7h = Port 7 All other encodings are <i>reserved</i> .	RW	Yes	
13:12	<i>Reserved</i>	RsvdP	No	00b
14	<i>Factory Test Only</i>	RWS	Yes	0
15	Software Configuration Control 0 = External Straps control the upstream port 1 = Software can control the upstream port	RWS	Yes	0

Register 13-61. 1DCh Debug Control (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
16	<p>Upstream Hot Reset Control</p> <p>0 = Assert Hot Reset 1 = Assert Secondary Bus Reset</p> <p><i>Note: Only a Fundamental Reset serial EEPROM load affects this bit.</i></p>	RWS	Yes	0
17	<p>Disable Serial EEPROM Load on Hot Reset</p> <p>0 = Enables serial EEPROM load upon upstream port Hot Reset or <i>DL_Down</i> state 1 = Disables serial EEPROM load upon upstream port Hot Reset or <i>DL_Down</i> state</p>	RWS	Yes	0
19:18	Reserved	RO	No	11b
20	<p>Upstream Port DL_Down Reset Propagation Disable</p> <p>Setting this bit enables the upstream port to ignore a Hot Reset training sequence, blocks the PEX 8509 from manifesting an internal reset due to a <i>DL_Down</i> event, and prevents the downstream ports from issuing a Hot Reset to downstream devices when either a Hot Reset or <i>DL_Down</i> event occurs on the upstream link.</p>	RWS	Yes	0
21	<p>Cut-Thru Enable</p> <p>0 = Disables Cut-Thru support 1 = Enables Cut-Thru support</p>	RWS	Yes	1
23:22	Factory Test Only	RWS	Yes	00b
29:24	Reserved	RsvdP	No	10h
30	<p>On-Board SerDes Lane Status Control</p> <p>0 = Physical Layer controls on-board SerDes Lane Status LED 1 = When bit 7 (Factory Test Only) is also set to 1, software-driven value to the Lane Status/Software PEX_LANE_GOODx# LED Control register <i>Port LED Status Control</i> field (Port 0, offset 1F4h[7:0]) controls the PEX_LANE_GOOD[7:0]# output states</p>	RWS	Yes	0
31	Reserved	RsvdP	No	0

Register 13-62. 1E0h Power Management Hot Plug User Configuration (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
0	L0s Entry Idle Count Time to meet to enter the L0s Link PM state. 0 = Idle condition lasts for 1 μ s 1 = Idle condition lasts for 4 μ s		RW	Yes	0
1	L1 Upstream Port Receiver Idle Count For active L1 Link PM state entry. 0 = Upstream port receiver remains idle for 2 μ s 1 = Upstream port receiver remains idle for 3 μ s		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 (default) 01b = 32 ms 10b = 64 ms 11b = 128 ms		RO	Yes	00b
5	Factory Test Only		RO	Yes	0
6	Reserved	Upstream	RsvdP	No	0
	HPC T_{ppervl} Delay Downstream port power-valid to reset signal release time. 0 = 20 ms 1 = 100 ms (default)	Downstream	RO	Yes	1
7	Disable PCI PM L1 Entry 0 = Enables upstream port entry into the L1 Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field is set (offset 44h[1:0]=11b)] 1 = Disables upstream port entry into the L1 Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field is set (offset 44h[1:0]=11b)]	Upstream	RW	Yes	0
	Reserved	Downstream	RsvdP	No	0

Register 13-62. 1E0h Power Management Hot Plug User Configuration (All Ports) (Cont.)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
8	<i>Reserved</i>	Upstream	RsvdP	No	0
	DLLP Timeout Link Retrain Disable 0 = Enables link retraining when no DLLPs are received for more than 200 μs (default) 1 = DLLP Timeout is disabled	Downstream	RW	Yes	0
9	Factory Test Only		RW	Yes	0
10	L0s Entry Disable 0 = Enables upstream port entry into the L0s Link PM state when the upstream port is placed in the D3hot Device PM state [Power Management Status and Control register <i>Power State</i> field is set (offset 44h[1:0]=11b)] 1 = Disables upstream port entry into the L0s Link PM state when the upstream port is placed in the D3hot Device PM state [[Power Management Status and Control register <i>Power State</i> field is set (offset 44h[1:0]=11b)]		RW	Yes	0
11	Enable Software-Controlled PERST RW for all ports. Functionally implemented only for Hot Plug-capable ports.		RW	Yes	0
12	Port Slot PERST State RW for all ports. Functionally implemented only for Hot Plug-capable ports.		RW	Yes	0
15:13	Factory Test Only		RW	Yes	000b
31:16	<i>Reserved</i>		RsvdP	No	0000h

Register 13-63. 1E8h Bad TLP Count (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Bad TLP Count Counts the number of TLPs with bad LCRC, or number of TLPs with a Sequence Number Mismatch error. The Counter saturates at FFFF_FFFFh and does not roll over to 0000_0000h.	RWS	Yes	0000_0000h

Register 13-64. 1ECh Bad DLLP Count (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Bad DLLP Count Counts the number of DLLPs with bad LCRC, or number of DLLPs with a Sequence Number Mismatch error. The Counter saturates at FFFF_FFFFh and does not roll over to 0000_0000h.	RWS	Yes	0000_0000h

Register 13-65. 1F4h Lane Status/Software PEX_LANE_GOODx# LED Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default																		
	<p>Lane Up Status</p> <p>When read, and the Debug Control register <i>On-Board SerDes Lane Status Control</i> bit is cleared (Port 0, offset 1DCh[30]=0), bits [7:0] indicate the physical linkup status of Lanes [7-0], respectively.</p> <p>0 = Lane is down 1 = Lane is up</p> <p><i>Note:</i> <i>DLL Layer linkup (Flow Control initialization) status is indicated by the VC0 Resource Status register VC0 Negotiation Pending bit (offset 160h[17]).</i></p>																					
	<p>Debug Control register <i>On-Board SerDes Lane Status Control</i> bit is cleared (Port 0, offset 1DCh[30]=0, default)</p>	RO	No	FFh																		
7:0	<p>Port LED Status Control</p> <p>When written, and the Debug Control register <i>On-Board SerDes Lane Status Control</i> bit and bit 7 are both set (Port 0, offset 1DCh[30, 7]=11b), controls the PEX_LANE_GOOD[7:0]# outputs.</p> <p><i>Note:</i> <i>Reads of this register always return Lane Up Status.</i></p> <p>Writes:</p> <p>0 = PEX_LANE_GOOD[x]# output is driven Low 1 = PEX_LANE_GOOD[x]# output is driven High</p> <table border="0"> <tr> <td>Bit</td> <td>Controls LED Status for Port</td> </tr> <tr> <td>0</td> <td>0</td> </tr> <tr> <td>1</td> <td>1</td> </tr> <tr> <td>2</td> <td>2</td> </tr> <tr> <td>3</td> <td>3</td> </tr> <tr> <td>4</td> <td>4</td> </tr> <tr> <td>5</td> <td>5</td> </tr> <tr> <td>6</td> <td>6</td> </tr> <tr> <td>7</td> <td>7</td> </tr> </table> <p>Reads:</p> <p>Bits [7:0] indicate the physical linkup status of Lanes [7-0], respectively.</p> <p>0 = Lane is down 1 = Lane is up</p>	Bit	Controls LED Status for Port	0	0	1	1	2	2	3	3	4	4	5	5	6	6	7	7			
Bit	Controls LED Status for Port																					
0	0																					
1	1																					
2	2																					
3	3																					
4	4																					
5	5																					
6	6																					
7	7																					
	<p>Debug Control register <i>On-Board SerDes Lane Status Control</i> bit and bit 7 are both set (Port 0, offset 1DCh[30, 7]=11b)</p>	RW	No	FFh																		
31:8	Reserved	RsvdP	No	0000_00h																		

Register 13-66. 1F8h ACK Transmission Latency Limit (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default																							
11:0	<p>ACK Transmission Latency Limit Acknowledge Control Packet Transmission Latency Limit. If the serial EEPROM is not present, the value of this field changes based upon a Maximum Payload Size of 256 bytes and the Negotiated Link Width Negotiated Link Width (offset 78h[25:20]) encoding, after that link is up.</p> <table border="1"> <thead> <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> </thead> <tbody> <tr> <td>128B</td> <td>237d</td> <td>128d</td> <td>73d</td> </tr> <tr> <td>256B</td> <td>416d</td> <td>217d</td> <td>118d</td> </tr> <tr> <td>512B</td> <td>559d</td> <td>289d</td> <td>154d</td> </tr> <tr> <td>1 KB</td> <td>1,071d</td> <td>545d</td> <td>282d</td> </tr> </tbody> </table>	Maximum Payload Size	Port Width			x1	x2	x4	128B	237d	128d	73d	256B	416d	217d	118d	512B	559d	289d	154d	1 KB	1,071d	545d	282d	RWS	Yes	0EDh
Maximum Payload Size	Port Width																										
	x1	x2	x4																								
128B	237d	128d	73d																								
256B	416d	217d	118d																								
512B	559d	289d	154d																								
1 KB	1,071d	545d	282d																								
15:12	Factory Test Only	RWS	Yes	0h																							
23:16	<p>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 (offset 78h[25:20]) encoding, after the Link is up. The value in this register is a multiplier of the default internal timer values that are compliant to <i>PCI Express Base r1.1</i>, and should be valid after link negotiation is complete. This field should normally remain the default value 00h.</p>	RWS	Yes	00h																							
30:24	Reserved	RsvdP	No	00h																							
31	<p>Status of ACK Transmission Latency Update Bit is set when bits [9:0] are written by serial EEPROM, I²C, or CSR access.</p>	RO	No	0																							

13.16.2 Device-Specific Registers – Physical Layer

Table 13-18. 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				
<i>Factory Test Only</i>								200h
<i>Reserved</i>		Physical Layer Receiver Not Detected Mask		<i>Reserved</i>		Physical Layer Electrical Idle Detect Mask		204h
<i>Factory Test Only</i>								208h
<i>Reserved</i>								20Ch
Physical Layer User Test Pattern 0								210h
Physical Layer User Test Pattern 4								214h
Physical Layer User Test Pattern 8								218h
Physical Layer User Test Pattern 12								21Ch
<i>Reserved</i>		Physical Layer Command and Status						220h
Port Configuration								224h
Physical Layer Test								228h
Physical Layer								22Ch
Physical Layer Port Command								230h
<i>Reserved</i>		Port Control		SKIP Ordered-Set Interval				234h
SerDes Quad 0 Diagnostic Data								238h
SerDes Quad 1 Diagnostic Data								23Ch
<i>Reserved</i>							240h –	244h
SerDes Nominal Drive Current Select								248h
SerDes Drive Current Level 1								24Ch
<i>Reserved</i>								250h
SerDes Drive Equalization Level Select 1								254h
<i>Reserved</i>								258h
<i>Factory Test Only</i>								25Ch
Status Data from Serial EEPROM		Serial EEPROM Status		Serial EEPROM Control				260h
Serial EEPROM Data Buffer								264h
Serial EEPROM Clock Frequency								268h
<i>Reserved</i>						Serial EEPROM 3 rd Address Byte		26Ch
<i>Reserved</i>								270h
<i>Factory Test Only</i>							274h –	284h
<i>Reserved</i>								288h
<i>Factory Test Only</i>								28Ch

Notes: In this section, the term “SerDes quad” or “quad” refers to assembling SerDes lanes into groups of four contiguous lanes for testing purposes.

Table 13-19 defines the default Port/Physical Lane/SerDes Module to SerDes Quad relationship.

**Table 13-19. Port/Physical Lane/SerDes Module to SerDes Quad Relationship
(Port 0, Offset 224h[1:0]=0h)**

Port	Physical Lanes and SerDes Modules	SerDes Quad
0	0	0
1	1	
2	2	
3	3	
4	4	1
5	5	
6	6	
7	7	

Register 13-67. 204h Physical Layer Receiver Not Detected and Electrical Idle Detect Masks (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Physical Layer Electrical Idle Detect Mask				
7:0	SerDes Mask Electrical Idle Detect 1 = Masks the Electrical Idle Detect for the SerDes lanes	RWS	Yes	0-0h
	Bit Port 0			
	0 SerDes 0			
	1 SerDes 1			
	2 SerDes 2			
	3 SerDes 3			
	4 SerDes 4			
	5 SerDes 5			
6 SerDes 6				
7 SerDes 7				
15:8	<i>Reserved</i>	RsvdP	No	00h
Physical Layer Receiver Not Detected Mask				
23:16	SerDes Mask Receiver Not Detected 1 = Masks the Receiver Not Detected for the SerDes lanes	RWS	Yes	00h
	Bit Port 0			
	16 SerDes 0			
	17 SerDes 1			
	18 SerDes 2			
	19 SerDes 3			
	20 SerDes 4			
	21 SerDes 5			
22 SerDes 6				
23 SerDes 7				
31:24	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-68. 210h Physical Layer User Test Pattern 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Test Pattern 0 Test pattern Bytes 0 through 3. Used for Digital Far-End Loopback testing.	RW	Yes	0-0h

Register 13-69. 214h Physical Layer User Test Pattern 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Test Pattern 4 Test pattern Bytes 4 through 7. Used for Digital Far-End Loopback testing.	RW	Yes	0-0h

Register 13-70. 218h Physical Layer User Test Pattern 8 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Test Pattern 8 Test pattern Bytes 8 through 11. Used for Digital Far-End Loopback testing.	RW	Yes	0-0h

Register 13-71. 21Ch Physical Layer User Test Pattern 12 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Test Pattern 12 Test pattern Bytes 12 through 15. Used for Digital Far-End Loopback testing.	RW	Yes	0-0h

Register 13-72. 220h Physical Layer Command and Status (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	<i>Reserved</i>	RsvdP	No	000b
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 the ability to enable or disable lane reversal. 0 = Lane reversal is supported 1 = Lane reversal is not supported	RWS	Yes	0
6	<i>Reserved</i>	RsvdP	No	0
7	Elastic Buffer Low-Latency Mode Disable 0 = Enables Elastic Buffer Low-Latency mode 1 = Disables Elastic Buffer Low-Latency mode	RWS	Yes	0
15:8	N_FTS Value Number of Fast Training Sets (N_FTS) value to transmit in training sets.	RWS	Yes	40h
19:16	<i>Reserved</i>	RsvdP	No	0h
23:20	Number of Ports Enumerated Number of ports in current configuration.	HwInit	Yes	0h
31:24	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-73. 224h Port Configuration (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
1:0	<p>Port Configuration</p> <p>The serial EEPROM bit values always override the STRAP_PORTCFG[1:0] Strapping signal values (assuming the serial EEPROM values are loaded; refer to Table 13-20).</p> <p><i>Note:</i> All other configurations default to option 0h.</p> <p>Bits [4:3] must always be programmed to 00b.</p> <p>This register is reset only by a Fundamental Reset (PEX_PERST# assertion).</p>	HwInit	Yes	00b
7:2	<i>Reserved</i>	RsvdP	No	00h
10:8	<p>x1 Only</p> <p>Bit 8 value of 1 forces Port 0 to be x1.</p> <p>Bit 9 value of 1 forces Port 1 to be x1.</p> <p>Bit 10 value is “don’t care.”</p>	RWS	Yes	000b
31:11	<i>Reserved</i>	RsvdP	No	0-0h

Table 13-20. Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[1:0])	Link Width/Port							
	Port 0	Port 1	Port 2	Port 3	Port 4	Port 5	Port 6	Port 7
00b	x1	x1	x1	x1	x1	x1	x1	x1
01b	x2	x1	x1	x1	x1	x1	x1	
10b	x2	x2	x1	x1	x1	x1		
11b	x4	x1	x1	x1	x1			

Register 13-74. 228h Physical Layer Test (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note: Port 0 parameters apply to SerDes[0-3] and SerDes[4-7], which map to Lanes [0-7], respectively.</i>				
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	Reserved	RW	Yes	0
3	TCB Capture Disable 0 = Training Control Bit (TCB) Capture is enabled 1 = Disables TCB Capture	RW	Yes	0
4	Analog Loopback Enable 0 = PEX 8509 enters Digital Loopback Slave mode if an external device sends at least two consecutive TS1 Ordered-Sets that have the <i>Loopback</i> bit exclusively set in the TS1 Training Control symbol. The PEX 8509 then loops back data through the elastic buffer, 8b/10b decoder, and 8b/10b encoder. 1 = PEX 8509 enters Analog Loopback Slave mode if an external device sends at least two consecutive TS1 Ordered-Sets that have the <i>Loopback</i> bit exclusively set in the TS1 Training Control symbol. The PEX 8509 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
6:5	Reserved	RsvdP	No	00b
7	PHY BIST Enable Physical Layer SerDes Built-In Self-Test Enable. When programmed to 1 by serial EEPROM, enables SerDes internal loopback Pseudo-Random Bit Sequence (PRBS) test for 16 ms before starting link initialization. Reads back as 0, but functions as BIST Enable in the logic.	RW	Yes	0
15:8	Reserved	RsvdP	No	00h

Register 13-74. 228h Physical Layer Test (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default						
17:16	<p>PRBS Enable When set to 1, enables PRBS sequence generation/checking on the SerDes quads.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Port 0, SerDes</th> </tr> </thead> <tbody> <tr> <td>16</td> <td>[0-3]</td> </tr> <tr> <td>17</td> <td>[4-7]</td> </tr> </tbody> </table> <p><i>Notes: When either bit in this field is set to 1, it must always be used with corresponding bits [21 and/or 20] (PRBS External Loopback) also set to 1. PRBS Enable and User Test Pattern Enable (bits [29:28]) 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 [17:16] ANDed with bits [29:28] must be 00b.</i></p>	Bit	Port 0, SerDes	16	[0-3]	17	[4-7]	RW	Yes	00b
Bit	Port 0, SerDes									
16	[0-3]									
17	[4-7]									
19:18	Reserved	RsvdP	No	00b						
21:20	<p>PRBS External Loopback 0 = SerDes quad establishes Internal Analog Loopback mode when the corresponding <i>PRBS Enable</i> bit (bit 17 or 16) is set to 1 1 = SerDes quad establishes Analog Loopback Master mode when the corresponding <i>PRBS Enable</i> bit (bit 17 or 16) is set to 1</p> <p><i>Note: This field does not enable Internal Analog Loopback mode when its bits are at a state of logic zero (0).</i></p> <p>The following bit commands are valid when the Physical Layer Port Command register <i>Port Loopback Command</i> bit (Port 0, offset 230h[0, 4, 8, 12, 16, 20, 24, or 28]) is set for the associated port.</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Port 0, SerDes</th> </tr> </thead> <tbody> <tr> <td>20</td> <td>[0-3]</td> </tr> <tr> <td>21</td> <td>[4-7]</td> </tr> </tbody> </table>	Bit	Port 0, SerDes	20	[0-3]	21	[4-7]	RW	Yes	00b
Bit	Port 0, SerDes									
20	[0-3]									
21	[4-7]									
23:22	Reserved	RsvdP	No	00b						
27:24	Reserved	RO	Yes	0h						
29:28	<p>User Test Pattern Enable 0 = Disables transmission of the 128-bit test pattern 1 = Enables transmission of the 128-bit test pattern [Physical Layer User Test Pattern x registers (Port 0, offsets 210h through 21Ch)] on the SerDes quads in Digital Far-End Loopback Master mode</p> <table border="1"> <thead> <tr> <th>Bit</th> <th>Port 0, SerDes</th> </tr> </thead> <tbody> <tr> <td>28</td> <td>[0-3]</td> </tr> <tr> <td>29</td> <td>[4-7]</td> </tr> </tbody> </table> <p><i>Note: User Test Pattern Enable and PRBS Enable (bits [17:16]) 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 [17:16] ANDed with bits [29:28] must be 00b.</i></p>	Bit	Port 0, SerDes	28	[0-3]	29	[4-7]	RW	Yes	00b
Bit	Port 0, SerDes									
28	[0-3]									
29	[4-7]									
31:30	Reserved	RsvdP	No	00b						

Register 13-75. 22Ch Physical Layer (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note: Port 0 parameters apply to SerDes[0-3] and SerDes[4-7], which map to Lanes [0-7], respectively.</i>				
5:0	Factory Test Only	RWS	Yes	00_0000b
7:6	Reserved	RWS	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 50Ω for physical Lanes [0-3]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50Ω) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RWS	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 50Ω for physical Lanes [4-7]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Tx termination to nominal (approximately 50Ω) 01b = Sets Tx termination to (nominal -17%) 10b = Sets Tx termination to (nominal +10%) 11b = Sets Tx termination to (nominal -15%)	RWS	Yes	00b
15:12	Reserved	RsvdP	No	0h
17:16	SerDes Quad 0 RxTermAdjust SerDes Quad 0 RxTermAdj[1:0]. Control bus to adjust Receive termination values above or below the nominal 50Ω for physical Lanes [0-3]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Rx termination to nominal (approximately 50Ω) 01b = Sets Rx termination to (nominal -17%) 10b = Sets Rx termination to (nominal +10%) 11b = Sets Rx termination to (nominal -15%)	RWS	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 50Ω for physical Lanes [4-7]. This allows precise matching to compensate for package or board impedance mismatch. 00b = Sets Rx termination to nominal (approximately 50Ω) 01b = Sets Rx termination to (nominal -17%) 10b = Sets Rx termination to (nominal +10%) 11b = Sets Rx termination to (nominal -15%)	RWS	Yes	00b
23:20	Reserved	RsvdP	No	0h

Register 13-75. 22Ch Physical Layer (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
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 13-21 defines the field decode.	RWS	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 [4-7]. For further details, refer to the expanded description that follows this register table. Table 13-21 defines the field decode.	RWS	Yes	00b
31:28	Reserved	RsvdP	No	0h

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 13-21](#).

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 13-21. RxEqCtl[1:0] Decode for Register Offset 22Ch[27:24] (Only Port 0)

RxEqCtl[1:0]	RX Eq Setting	Input Jitter	Channel Length
00b	Maximum Rx Eq	0.25 UI	50.8 cm (20 in.) and two or more connectors
10b	Minimum Rx Eq	Between 0.1 and 0.25 UI	Between 20.32 and 50.8 cm (8 and 20 in.) and up to two connectors
01b, 11b	Rx Eq OFF	< 0.1 UI	20.32 cm (8 in.) or less, up to one connector

Register 13-76. 230h Physical Layer Port Command (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<p>Port 0 Loopback Command</p> <p>0 = Port 0 is not enabled to go to Loopback Master state 1 = Port 0 is enabled to go to Loopback Master state</p>	RW	Yes	0
1	<p>Port 0 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 0 scrambler is enabled 1 = Port 0 scrambler is disabled</p>	RW	Yes	0
2	<p>Port 0 Rx L1 Only</p> <p>Port 0 Receiver enters into the ASPM L1 Link PM state.</p> <p>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</p>	RW	Yes	0
3	<p>Port 0 Ready as Loopback Master</p> <p>Port 0 Link Training and Status State Machine (LTSSM) established Loopback as a Master.</p> <p>0 = Port 0 is not in Loopback Master mode 1 = Port 0 is in Loopback Master mode</p>	RO	No	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>Port 1 Loopback Command</p> <p>0 = Port 1 is not enabled to go to Loopback Master state 1 = Port 1 is enabled to go to Loopback Master state</p>	RW	Yes	0
5	<p>Port 1 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 1 scrambler is enabled 1 = Port 1 scrambler is disabled</p>	RW	Yes	0
6	<p>Port 1 Rx L1 Only</p> <p>Port 1 Receiver enters into the ASPM L1 Link PM state.</p> <p>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</p>	RW	Yes	0
7	<p>Port 1 Ready as Loopback Master</p> <p>Port 1 LTSSM established Loopback as a Master.</p> <p>0 = Port 1 is not in Loopback Master mode 1 = Port 1 is in Loopback Master mode</p>	RO	No	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
8	<p>Port 2 Loopback Command</p> <p>0 = Port 2 is not enabled to go to Loopback Master state 1 = Port 2 is enabled to go to Loopback Master state</p>	RW	Yes	0
9	<p>Port 2 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 2 scrambler is enabled 1 = Port 2 scrambler is disabled</p>	RW	Yes	0
10	<p>Port 2 Rx L1 Only</p> <p>Port 2 Receiver enters into the ASPM L1 Link PM state.</p> <p>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</p>	RW	Yes	0
11	<p>Port 2 Ready as Loopback Master</p> <p>Port 2 LTSSM established Loopback as a Master.</p> <p>0 = Port 2 is not in Loopback Master mode 1 = Port 2 is in Loopback Master mode</p>	RO	No	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
12	<p>Port 3 Loopback Command</p> <p>0 = Port 3 is not enabled to go to Loopback Master state 1 = Port 3 is enabled to go to Loopback Master state</p>	RW	Yes	0
13	<p>Port 3 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 3 scrambler is enabled 1 = Port 3 scrambler is disabled</p>	RW	Yes	0
14	<p>Port 3 Rx L1 Only</p> <p>Port 3 Receiver enters into the ASPM L1 Link PM state.</p> <p>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</p>	RW	Yes	0
15	<p>Port 3 Ready as Loopback Master</p> <p>Port 3 LTSSM established Loopback as a Master.</p> <p>0 = Port 3 not in Loopback Master mode 1 = Port 3 in Loopback Master mode</p>	RO	No	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
16	<p>Port 4 Loopback Command</p> <p>0 = Port 4 is not enabled to go to Loopback Master state 1 = Port 4 is enabled to go to Loopback Master state</p>	RW	Yes	0
17	<p>Port 4 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 4 scrambler is enabled 1 = Port 4 scrambler is disabled</p>	RW	Yes	0
18	<p>Port 4 Rx L1 Only</p> <p>Port 4 Receiver enters into the ASPM L1 Link PM state.</p> <p>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</p>	RW	Yes	0
19	<p>Port 4 Ready as Loopback Master</p> <p>Port 4 LTSSM established Loopback as a Master.</p> <p>0 = Port 4 is not in Loopback Master mode 1 = Port 4 is in Loopback Master mode</p>	RO	No	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
20	<p>Port 5 Loopback Command</p> <p>0 = Port 5 is not enabled to go to Loopback Master state 1 = Port 5 is enabled to go to Loopback Master state</p>	RW	Yes	0
21	<p>Port 5 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 5 scrambler is enabled 1 = Port 5 scrambler is disabled</p>	RW	Yes	0
22	<p>Port 5 Rx L1 Only</p> <p>Port 5 Receiver enters into the ASPM L1 Link PM state.</p> <p>0 = Port 5 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 5 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</p>	RW	Yes	0
23	<p>Port 5 Ready as Loopback Master</p> <p>Port 5 LTSSM established Loopback as a Master.</p> <p>0 = Port 5 is not in Loopback Master mode 1 = Port 5 is in Loopback Master mode</p>	RW	Yes	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
24	<p>Port 6 Loopback Command</p> <p>0 = Port 6 is not enabled to go to Loopback Master state 1 = Port 6 is enabled to go to Loopback Master state</p>	RW	Yes	0
25	<p>Port 6 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 6 scrambler is enabled 1 = Port 6 scrambler is disabled</p>	RW	Yes	0
26	<p>Port 6 Rx L1 Only</p> <p>Port 6 Receiver enters into the ASPM L1 Link PM state.</p> <p>0 = Port 6 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 6 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</p>	RW	Yes	0
27	<p>Port 6 Ready as Loopback Master</p> <p>Port 6 LTSSM established Loopback as a Master.</p> <p>0 = Port 6 is not in Loopback Master mode 1 = Port 6 is in Loopback Master mode</p>	RO	No	0

Register 13-76. 230h Physical Layer Port Command (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
28	<p>Port 7 Loopback Command</p> <p>0 = Port 7 is not enabled to go to Loopback Master state 1 = Port 7 is enabled to go to Loopback Master state</p>	RW	Yes	0
29	<p>Port 7 Scrambler Disable</p> <p>If a serial EEPROM load sets this bit, the scrambler is disabled in a Configuration Complete state.</p> <p>If software sets this bit when the link is in the Up state, hardware immediately disables its scrambler without executing the Link Training protocol. The upstream/downstream device scrambler will not be disabled.</p> <p>0 = Port 7 scrambler is enabled 1 = Port 7 scrambler is disabled</p>	RW	Yes	0
30	<p>Port 7 Rx L1 Only</p> <p>Port 7 Receiver enters into the ASPM L1 Link PM state.</p> <p>0 = Port 7 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 7 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</p>	RW	Yes	0
31	<p>Port 7 Ready as Loopback Master</p> <p>Port 7 LTSSM established Loopback as a Master.</p> <p>0 = Port 7 is not in Loopback Master mode 1 = Port 7 is in Loopback Master mode</p>	RO	No	0

Register 13-77. 234h SKIP Ordered-Set Interval and Port Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
SKIP Ordered-Set Interval				
11:0	SKIP Ordered-Set Interval SKIP Ordered-Set interval (in symbol times). 49Ch = Minimum interval (1,180 symbol times) 602h = Maximum interval (1,538 symbol times)	RWS	Yes	49Ch
15:12	<i>Reserved</i>	RsvdP	No	0h
Port Control				
16	Disable Port 0 0 = Enables Link Training operation on Port 0 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 0	RWS	Yes	0
17	Disable Port 1 0 = Enables Link Training operation on Port 1 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 1	RWS	Yes	0
18	Disable Port 2 0 = Enables Link Training operation on Port 2 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 2	RWS	Yes	0
19	Disable Port 3 0 = Enables Link Training operation on Port 3 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 3	RWS	Yes	0
20	Disable Port 4 0 = Enables Link Training operation on Port 4 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 4	RWS	Yes	0
21	Disable Port 5 0 = Enables Link Training operation on Port 5 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 5	RWS	Yes	0
22	Disable Port 6 0 = Enables Link Training operation on Port 6 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 6	RWS	Yes	0
23	Disable Port 7 0 = Enables Link Training operation on Port 7 1 = LTSSM remains in <i>Detect.Quiet</i> state on Port 7	RWS	Yes	0
31:24	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-78. 238h SerDes Quad 0 Diagnostic Data (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note: Port 0 parameters apply to SerDes[0-3], which map to Lanes [0-3], respectively.</i>				
7:0	UTP Expected Data Expected User Test Pattern (UTP) SerDes[0-3] Diagnostic data when UTP is enabled [Physical Layer Test register <i>User Test Pattern Enable</i> bit for SerDes Quad 0 (Port 0, offset 228h[28]=1)].	RO	No	00h
15:8	Actual Data Actual UTP SerDes[0-3] Diagnostic data when UTP is enabled [Physical Layer Test register <i>User Test Pattern Enable</i> bit for SerDes Quad 0 (Port 0, offset 228h[28]=1)].	RO	No	00h
23:16	UTP/PRBS Error Count SerDes Quad 0 detected flags. UTP/PRBS SerDes[0-3] Error Count (0 to 255).	RO	Yes	00h
25:24	SerDes Diagnostic Data Select SerDes[0-3] Diagnostic Data Select. Status selection code for the fields representing RO bits [23:0] of this register. The binary code represents a status selection for one of the four Lanes of SerDes Quad 0. The test results for physical device Lanes [0-3] are selected with corresponding binary codes from 0-3.	RW	Yes	00b
29:26	Reserved	RO	No	0h
30	PRBS Count/-UTP Count 0 = Indicates that field [23:16] (<i>UTP/PRBS Error Count</i>) is the UTP Error Count 1 = Indicates that field [23:16] (<i>UTP/PRBS Error Count</i>) is the PRBS Error Count	RO	No	0
31	BIST Done	RO	No	0

Register 13-79. 23Ch SerDes Quad 1 Diagnostic Data (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note: Port 0 parameters apply to SerDes[4-7], which map to Lanes [4-7], respectively.</i>				
7:0	UTP Expected Data Expected UTP SerDes[4-7] Diagnostic data when UTP is enabled [Physical Layer Test register <i>User Test Pattern Enable</i> bit for SerDes Quad 1 (Port 0, offset 228h[29]=1)].	RO	No	00h
15:8	Actual Data Actual UTP SerDes[4-7] Diagnostic data when UTP is enabled [Physical Layer Test register <i>User Test Pattern Enable</i> bit for SerDes Quad 1 (Port 0, offset 228h[29]=1)].	RO	No	00h
23:16	UTP/PRBS Error Count SerDes Quad 1 detected flags. UTP/PRBS SerDes[4-7] Error Count (0 to 255).	RO	Yes	00h
25:24	SerDes Diagnostic Data Select SerDes[4-7] Diagnostic Data Select. Status selection code for the fields representing RO bits [23:0] of this register. The binary code represents a status selection for one of the four Lanes of SerDes Quad 1. The test results for physical device Lanes [4-7] are selected with corresponding binary codes from 0-3.	RW	Yes	00b
29:26	Reserved	RO	No	0h
30	PRBS Count/-UTP Count 0 = Indicates that field [23:16] (<i>UTP/PRBS Error Count</i>) is the UTP Error Count 1 = Indicates that field [23:16] (<i>UTP/PRBS Error Count</i>) is the PRBS Error Count	RO	No	0
31	BIST Done	RO	No	0

Register 13-80. 248h SerDes Nominal Drive Current Select (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default	
<i>Note: Port 0 parameters apply to SerDes[0-3] and SerDes[4-7], which map to Lanes [0-7], respectively.</i>					
1:0	SerDes 0 Nominal Drive Current	The following values for Nominal Current apply to each drive: <ul style="list-style-type: none"> • 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	SerDes 5 Nominal Drive Current		RWS	Yes	00b
13:12	SerDes 6 Nominal Drive Current		RWS	Yes	00b
15:14	SerDes 7 Nominal Drive Current		RWS	Yes	00b
31:16	Reserved	RsvdP	No	0000h	

Register 13-81. 24Ch SerDes Drive Current Level 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
<i>Note: Port 0 parameters apply to SerDes[0-3] and SerDes[4-7], which map to Lanes [0-7], respectively.</i>				
3:0	SerDes 0 Drive Current Level	RWS	Yes	0h
7:4	SerDes 1 Drive Current Level			
11:8	SerDes 2 Drive Current Level			
15:12	SerDes 3 Drive Current Level			
19:16	SerDes 4 Drive Current Level			
23:20	SerDes 5 Drive Current Level			
27:24	SerDes 6 Drive Current Level			
31:28	SerDes 7 Drive Current Level			
		The following values represent the ratio of Actual Current/Nominal Current (selected in the 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		

Register 13-82. 254h SerDes Drive Equalization Level Select 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default																																		
<i>Note: Port 0 parameters apply to SerDes[0-3] and SerDes[4-7], which map to Lanes [0-7], respectively.</i>																																						
3:0	SerDes 0 Drive Equalization Level	RWS	Yes	8h																																		
7:4	SerDes 1 Drive Equalization Level																																					
11:8	SerDes 2 Drive Equalization Level																																					
15:12	SerDes 3 Drive Equalization Level																																					
19:16	SerDes 4 Drive Equalization Level																																					
23:20	SerDes 5 Drive Equalization Level																																					
27:24	SerDes 6 Drive Equalization Level																																					
31:28	SerDes 7 Drive Equalization Level																																					
		The following values represent the percentage of Drive Current attributable to Equalization Current and apply to each drive: <table border="1"> <thead> <tr> <th>I_{EQ} / I_{DR}</th> <th>De-Emphasis (dB)</th> </tr> </thead> <tbody> <tr><td>0h = 0.00</td><td>0.00</td></tr> <tr><td>1h = 0.04</td><td>-0.35</td></tr> <tr><td>2h = 0.08</td><td>-0.72</td></tr> <tr><td>3h = 0.12</td><td>-1.11</td></tr> <tr><td>4h = 0.16</td><td>-1.51</td></tr> <tr><td>5h = 0.20</td><td>-1.94</td></tr> <tr><td>6h = 0.24</td><td>-2.38</td></tr> <tr><td>7h = 0.28</td><td>-2.85</td></tr> <tr><td>8h = 0.32</td><td>-3.35</td></tr> <tr><td>9h = 0.36</td><td>-3.88</td></tr> <tr><td>Ah = 0.40</td><td>-4.44</td></tr> <tr><td>Bh = 0.44</td><td>-5.04</td></tr> <tr><td>Ch = 0.48</td><td>-5.68</td></tr> <tr><td>Dh = 0.52</td><td>-6.38</td></tr> <tr><td>Eh = 0.56</td><td>-7.13</td></tr> <tr><td>Fh = 0.60</td><td>-7.96</td></tr> </tbody> </table>			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
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																																					

Register 13-83. 260h Serial EEPROM Status and Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Serial EEPROM Control				
12:0	EepBlkAddr Serial EEPROM Block Address for 32 KB.	RW	Yes	000h
15:13	EepCmd[2:0] Commands to the Serial EEPROM Controller. 000b = <i>Reserved</i> 001b = Data from bits [31:24] (Status Data from Serial EEPROM register) is written to the serial EEPROM's internal Status register 010b = Write four bytes of data from the <i>EepBuf</i> into the memory location pointed to by the <i>EepBlkAddr</i> field 011b = Read four bytes of data from the memory location pointed to by the <i>EepBlkAddr</i> field into the <i>EepBuf</i> 100b = Reset Write Enable latch 101b = Data from the serial EEPROM's internal Status register is written to bits [31:24] (Status Data from Serial EEPROM register) 110b = Set Write Enable latch 111b = <i>Reserved</i> <i>Note: For value of 001b, only bits [31, 27:26] can be written into the serial EEPROM's internal Status register.</i>	RW	Yes	000b

Register 13-83. 260h Serial EEPROM Status and Control (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Serial EEPROM Status				
17:16	EepPrsnt[1:0] Serial EEPROM Present status. 00b = Not present 01b = Serial EEPROM is present – validation signature verified 10b = <i>Reserved</i> 11b = Serial EEPROM is present – validation signature not verified	RO	No	–
18	EepCmdStatus Serial EEPROM Command status. 0 = Serial EEPROM Command is complete 1 = Serial EEPROM Command is not complete	RO	No	0
19	<i>Reserved</i>	RsvdP	No	0
20	EepBlkAddrUp Serial EEPROM Block Address upper bit 13. Extends the serial EEPROM to 64 KB.	RW	Yes	0
21	EepAddrWidthOverride 0 = <i>EepAddrWidth</i> is Read-Only 1 = <i>EepAddrWidth</i> is software writable	RW	Yes	0
23:22	EepAddrWidth Serial EEPROM Address width. If the addressing width cannot be determined, 00b is returned. A non-zero value is reported only if the validation signature (5Ah) is successfully read from the first serial EEPROM location. 00b = Undetermined 01b = 1 byte 10b = 2 bytes 11b = 3 bytes	RO	No	–

Register 13-83. 260h Serial EEPROM Status and Control (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default																																		
Status Data from Serial EEPROM^a																																						
24	<p>EepRdy Serial EEPROM RDY#.</p> <p>0 = Serial EEPROM is ready to transmit data 1 = Write cycle is in progress</p>	RW	Yes	0																																		
25	<p>EepWen Serial EEPROM Write Enable.</p> <p>0 = Serial EEPROM Write is disabled 1 = Serial EEPROM Write is enabled</p>	RW	Yes	0																																		
27:26	<p>EepBp[1:0] Serial EEPROM Block-Write Protect bits. Block Protection options protect the top ¼, top ½, or the entire serial EEPROM. PEX 8509 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.</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <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> </thead> <tbody> <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 ¼)</td> <td>1800h – 1FFFh</td> <td>3000h – 3FFFh</td> <td>6000h – 7FFFh</td> <td>–</td> </tr> <tr> <td>10b</td> <td>2 (top ½)</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> </tbody> </table>	BP[1:0]	Level	Array Addresses Protected				8-KB Device	16-KB Device	32-KB Device	64-KB Device	00b	0	None	None	None	None	01b	1 (top ¼)	1800h – 1FFFh	3000h – 3FFFh	6000h – 7FFFh	–	10b	2 (top ½)	1000h – 1FFFh	2000h – 3FFFh	4000h – 7FFFh	–	11b	3 (All)	0000h – 1FFFh	0000h – 3FFFh	0000h – 7FFFh	–	RW	Yes	00b
BP[1:0]	Level			Array Addresses Protected																																		
		8-KB Device	16-KB Device	32-KB Device	64-KB Device																																	
00b	0	None	None	None	None																																	
01b	1 (top ¼)	1800h – 1FFFh	3000h – 3FFFh	6000h – 7FFFh	–																																	
10b	2 (top ½)	1000h – 1FFFh	2000h – 3FFFh	4000h – 7FFFh	–																																	
11b	3 (All)	0000h – 1FFFh	0000h – 3FFFh	0000h – 7FFFh	–																																	
30:28	<p>EepWrStatus Serial EEPROM Write status. Value is 000b when the serial EEPROM is not in an internal Write cycle.</p> <p><i>Note: Definition of this field varies among serial EEPROM manufacturers. Reads of the serial EEPROM's internal Status register can return 000b or 111b, depending upon the serial EEPROM that is used.</i></p>	RO	Yes	000b																																		
31	<p>EepWpen Serial EEPROM Write Protect Enable.</p> <p>Overrides the internal serial EEPROM Write Protect WP# input and enables/disables Writes to the Serial EEPROM Status register:</p> <ul style="list-style-type: none"> • When WP# is high or <i>EepWpen</i> = 0, and <i>EepWen</i> = 1, the Serial EEPROM Status register is writable • When WP# is low and <i>EepWpen</i> = 1, or <i>EepWen</i> = 0, the Serial EEPROM Status register is protected <p><i>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] field be cleared to disable Block Protection, until WP# is High.</i></p> <p><i>This bit is not implemented in certain serial EEPROMs. Refer to the serial EEPROM manufacturer's data sheet.</i></p>	RW	Yes	0																																		

a. Within the serial EEPROM's internal Status register, only bits [31, 27:26] can be written.

Register 13-84. 264h Serial EEPROM Data Buffer (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	EepBuf Serial EEPROM RW 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	0000_0000h

Register 13-85. 268h Serial EEPROM Clock Frequency (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	EepFreq[2:0] Serial EEPROM clock (EE_SK) frequency control. 000b = 1 MHz (default) 001b = 1.98 MHz 010b = 5 MHz 011b = 9.62 MHz 100b = 12.5 MHz 101b = 15.6 MHz 110b = 17.86 MHz 111b = <i>Reserved</i>	RW	Yes	000b
7:3	<i>Reserved</i>	RsvdP	No	0-0h
10:8	EepCsStHld[2:0] Number of 1/2 EE_SK Clock cycles. 001b = Setup and hold of EE_CS# active to EE_SK active or EE_SK inactive to EE_CS# inactive	RW	Yes	000b
31:11	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-86. 26Ch Serial EEPROM 3rd Address Byte (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Serial EEPROM 3rd Address Byte	RW	Yes	00h
31:8	<i>Reserved</i>	RsvdP	No	0000_00h

13.16.3 Device-Specific Registers – I²C Interface

Table 13-22 defines the I²C interface register map.

Table 13-22. Device-Specific 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	<i>Factory Test Only</i>	290h
<i>I²C Configuration</i>			294h
<i>Factory Test Only</i>			298h – 2A8h
<i>Reserved</i>			2ACh – 2C4h

Register 13-87. 294h I²C Configuration (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	<p>Slave Address</p> <p>Bits [6:0] comprise the I²C Slave address, 3Fh – the value of bits [2:0] reflects the I2C_ADDR[2:0] ball state, which default to 111b, by virtue of weak internal pull-up resistors.</p> <p><i>Note: The I²C Slave address must not be changed by an I²C Write command.</i></p>	HwInit	Yes	111b
6:3	<p>Slave Address</p> <p>Bits [6:0] comprise the I²C Slave address, 3Fh – bits [6:3] default to 0111b.</p> <p><i>Note: The I²C Slave address must not be changed by an I²C Write command.</i></p>	RW	Yes	0111b
9:7	<i>Reserved</i>	RsvdP	No	000b
10	<i>Factory Test Only</i>	RW	Yes	0
31:11	<i>Reserved</i>	RW	Yes	0000_00h

13.16.4 Device-Specific Registers – Bus Number CAM

The **Bus Number Content-Addressable Memory (BusNoCAM)** 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 8509 port.

These registers are automatically updated by hardware. *Modifying these registers by writing to the addresses listed here is not recommended.*

Table 13-23. 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	
<i>Reserved</i>	Bus Number CAM 0			2C8h
<i>Reserved</i>	Bus Number CAM 1			2CCh
<i>Reserved</i>	Bus Number CAM 2			2D0h
<i>Reserved</i>	Bus Number CAM 3			2D4h
<i>Reserved</i>	Bus Number CAM 4			2D8h
<i>Reserved</i>	Bus Number CAM 5			2DCh
<i>Reserved</i>	Bus Number CAM 6			2E0h
<i>Reserved</i>	Bus Number CAM 7			2E4h
<i>Reserved</i>				2E8h – 304h

Register 13-88. 2C8h Bus Number CAM 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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	00h
23:16	Subordinate Bus Number Mirror copy of Port 0 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-89. 2CCh Bus Number CAM 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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	00h
23:16	Subordinate Bus Number Mirror copy of Port 1 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-90. 2D0h Bus Number CAM 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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	00h
23:16	Subordinate Bus Number Mirror copy of Port 2 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-91. 2D4h Bus Number CAM 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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	00h
23:16	Subordinate Bus Number Mirror copy of Port 3 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-92. 2D8h Bus Number CAM 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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	00h
23:16	Subordinate Bus Number Mirror copy of Port 4 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-93. 2DCh Bus Number CAM 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Primary Bus Number Mirror copy of Port 5 Primary Bus Number.	RW	Yes	00h
15:8	Secondary Bus Number Mirror copy of Port 5 Secondary Bus Number.	RW	Yes	00h
23:16	Subordinate Bus Number Mirror copy of Port 5 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-94. 2E0h Bus Number CAM 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Primary Bus Number Mirror copy of Port 6 Primary Bus Number.	RW	Yes	00h
15:8	Secondary Bus Number Mirror copy of Port 6 Secondary Bus Number.	RW	Yes	00h
23:16	Subordinate Bus Number Mirror copy of Port 6 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

Register 13-95. 2E4h Bus Number CAM 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Primary Bus Number Mirror copy of Port 7 Primary Bus Number.	RW	Yes	00h
15:8	Secondary Bus Number Mirror copy of Port 7 Secondary Bus Number.	RW	Yes	00h
23:16	Subordinate Bus Number Mirror copy of Port 7 Subordinate Bus Number.	RW	Yes	00h
31:24	<i>Reserved</i>	RsvdP	No	00h

13.16.5 Device-Specific Registers – I/O CAM

The **I/O Content-Addressable Memory (IOCAM)** registers 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 8509 port.

These registers are automatically updated by hardware. *Modifying these registers by writing to the addresses listed here is not recommended.*

Table 13-24. 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
I/O CAM 5		I/O CAM 4		310h
I/O CAM 7		I/O CAM 6		314h
<i>Reserved</i>				318h – 344h

Register 13-96. 308h I/O CAM 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 0 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 0 I/O Limit value.	RW	Yes	0h

Register 13-97. 30Ah I/O CAM 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 1 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 1 I/O Limit value.	RW	Yes	0h

Register 13-98. 30Ch I/O CAM 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 2 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 2 I/O Limit value.	RW	Yes	0h

Register 13-99. 30Eh I/O CAM 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 3 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 3 I/O Limit value.	RW	Yes	0h

Register 13-100. 310h I/O CAM 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 4 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 4 I/O Limit value.	RW	Yes	0h

Register 13-101. 312h I/O CAM 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 5 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 5 I/O Limit value.	RW	Yes	0h

Register 13-102. 314h I/O CAM 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 6 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 6 I/O Limit value.	RW	Yes	0h

Register 13-103. 316h I/O CAM 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
7:4	I/O Base Mirror copy of Port 7 I/O Base value.	RW	Yes	Fh
11:8	I/O Addressing Capability 0h = 16-bit I/O addressing 1h = 32-bit I/O addressing	RO	Yes	1h
15:12	I/O Limit Mirror copy of Port 7 I/O Limit value.	RW	Yes	0h

13.16.6 Device-Specific Registers – Address-Mapping CAM

The **Address-Mapping Content-Addressable Memory (AMCAM)** registers are used to determine Memory Request routing. These registers contain mirror copies of the **Memory Base and Limit, Prefetchable Memory Base and Limit, Prefetchable Memory Base Upper 32 Bits, and Prefetchable Memory Limit Upper 32 Bits** registers of each PEX 8509 port.

These registers are automatically updated by hardware. *Modifying these registers by writing to the addresses listed here is not recommended.*

Table 13-25. 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
AMCAM 5 Memory Limit		AMCAM 5 Memory Base		398h
AMCAM 5 Prefetchable Memory Limit		AMCAM 5 Prefetchable Memory Base		39Ch

Table 13-25. Device-Specific AMCAM Register Map (Only Port 0) (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	
AMCAM 5 Prefetchable Memory Upper Base Address				3A0h
AMCAM 5 Prefetchable Memory Upper Limit Address				3A4h
AMCAM 6 Memory Limit		AMCAM 6 Memory Base		3A8h
AMCAM 6 Prefetchable Memory Limit		AMCAM 6 Prefetchable Memory Base		3ACh
AMCAM 6 Prefetchable Memory Upper Base Address				3B0h
AMCAM 6 Prefetchable Memory Upper Limit Address				3B4h
AMCAM 7 Memory Limit		AMCAM 7 Memory Base		3B8h
AMCAM 7 Prefetchable Memory Limit		AMCAM 7 Prefetchable Memory Base		3BCh
AMCAM 7 Prefetchable Memory Upper Base Address				3C0h
AMCAM 7 Prefetchable Memory Upper Limit Address				3C4h
<i>Reserved</i>				3C8h – 548h

Register 13-104. 348h AMCAM 0 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 0 Memory Base				
3:0	<i>Reserved</i>	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	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 0 Memory Limit Mirror copy of Port 0 Memory Limit value.	RW	Yes	000h

Register 13-105. 34Ch AMCAM 0 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 13-106. 350h AMCAM 0 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 0 Prefetchable Memory Base[63:32] AMCAM 0 Port 0 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-107. 354h AMCAM 0 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 0 Prefetchable Memory Limit[63:32] AMCAM 0 Port 0 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-108. 358h AMCAM 1 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 1 Memory Base				
3:0	<i>Reserved</i>	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	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 1 Memory Limit Mirror copy of Port 1 Memory Limit value.	RW	Yes	000h

Register 13-109. 35Ch AMCAM 1 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 13-110. 360h AMCAM 1 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 1 Prefetchable Memory Base[63:32] AMCAM 1 Port 1 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-111. 364h AMCAM 1 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 1 Prefetchable Memory Limit[63:32] AMCAM 1 Port 1 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-112. 368h AMCAM 2 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 2 Memory Base				
3:0	<i>Reserved</i>	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	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 2 Memory Limit Mirror copy of Port 2 Memory Limit value.	RW	Yes	000h

Register 13-113. 36Ch AMCAM 2 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 13-114. 370h AMCAM 2 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 2 Prefetchable Memory Base[63:32] AMCAM 2 Port 2 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-115. 374h AMCAM 2 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 2 Prefetchable Memory Limit[63:32] AMCAM 2 Port 2 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-116. 378h AMCAM 3 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 3 Memory Base				
3:0	<i>Reserved</i>	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	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 3 Memory Limit Mirror copy of Port 3 Memory Limit value.	RW	Yes	000h

Register 13-117. 37Ch AMCAM 3 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 13-118. 380h AMCAM 3 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 3 Prefetchable Memory Base[63:32] AMCAM 3 Port 3 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-119. 384h AMCAM 3 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 3 Prefetchable Memory Limit[63:32] AMCAM 3 Port 3 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-120. 388h AMCAM 4 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 4 Memory Base				
3:0	<i>Reserved</i>	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	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 4 Memory Limit Mirror copy of Port 4 Memory Limit value.	RW	Yes	000h

Register 13-121. 38Ch AMCAM 4 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 4 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 13-122. 390h AMCAM 4 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 4 Prefetchable Memory Base[63:32] AMCAM 4 Port 4 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-123. 394h AMCAM 4 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 4 Prefetchable Memory Limit[63:32] AMCAM 4 Port 4 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-124. 398h AMCAM 5 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 5 Memory Base				
3:0	<i>Reserved</i>	RsvdP	No	0h
15:4	AMCAM 5 Memory Base Mirror copy of Port 5 Memory Base value.	RW	Yes	FFFh
AMCAM 5 Memory Limit				
19:16	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 5 Memory Limit Mirror copy of Port 5 Memory Limit value.	RW	Yes	000h

Register 13-125. 39Ch AMCAM 5 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 5 Prefetchable Memory Base				
3:0	AMCAM 5 Addressing Support 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	AMCAM 5 Prefetchable Memory Base AMCAM 5 Port 5 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 5 Prefetchable Memory Limit				
19:16	AMCAM 5 Addressing Support 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	AMCAM 5 Prefetchable Memory Limit AMCAM 5 Port 5 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-126. 3A0h AMCAM 5 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 5 Prefetchable Memory Base[63:32] AMCAM 5 Port 5 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-127. 3A4h AMCAM 5 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 5 Prefetchable Memory Limit[63:32] AMCAM 5 Port 5 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-128. 3A8h AMCAM 6 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 6 Memory Base				
3:0	<i>Reserved</i>	RsvdP	No	0h
15:4	AMCAM 6 Memory Base Mirror copy of Port 6 Memory Base value.	RW	Yes	FFFh
AMCAM 6 Memory Limit				
19:16	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 6 Memory Limit Mirror copy of Port 6 Memory Limit value.	RW	Yes	000h

Register 13-129. 3ACh AMCAM 6 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 6 Prefetchable Memory Base				
3:0	AMCAM 6 Addressing Support 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	AMCAM 6 Prefetchable Memory Base AMCAM 6 Port 6 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 6 Prefetchable Memory Limit				
19:16	AMCAM 6 Addressing Support 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	AMCAM 6 Prefetchable Memory Limit AMCAM 6 Port 6 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-130. 3B0h AMCAM 6 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 6 Prefetchable Memory Base[63:32] AMCAM 6 Port 6 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-131. 3B4h AMCAM 6 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 6 Prefetchable Memory Limit[63:32] AMCAM 6 Port 6 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

Register 13-132. 3B8h AMCAM 7 Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 7 Memory Base				
3:0	<i>Reserved</i>	RsvdP	No	0h
15:4	AMCAM 7 Memory Base Mirror copy of Port 7 Memory Base value.	RW	Yes	FFFh
AMCAM 7 Memory Limit				
19:16	<i>Reserved</i>	RsvdP	No	0h
31:20	AMCAM 7 Memory Limit Mirror copy of Port 7 Memory Limit value.	RW	Yes	000h

Register 13-133. 3BCh AMCAM 7 Prefetchable Memory Base and Limit (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
AMCAM 7 Prefetchable Memory Base				
3:0	AMCAM 7 Addressing Support 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
15:4	AMCAM 7 Prefetchable Memory Base AMCAM 7 Port 7 Prefetchable Memory Base[31:20].	RW	Yes	FFFh
AMCAM 7 Prefetchable Memory Limit				
19:16	AMCAM 7 Addressing Support 0h = 32-bit addressing is supported 1h = 64-bit addressing is supported	RO	Yes	1h
31:20	AMCAM 7 Prefetchable Memory Limit AMCAM 7 Port 7 Prefetchable Memory Limit[31:20].	RW	Yes	000h

Register 13-134. 3C0h AMCAM 7 Prefetchable Memory Upper Base Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 7 Prefetchable Memory Base[63:32] AMCAM 7 Port 7 Prefetchable Memory Base[63:32].	RW	Yes	0000_0000h

Register 13-135. 3C4h AMCAM 7 Prefetchable Memory Upper Limit Address (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	AMCAM 7 Prefetchable Memory Limit[63:32] AMCAM 7 Port 7 Prefetchable Memory Limit[63:32].	RW	Yes	0000_0000h

13.16.7 Device-Specific Registers – Ingress Control and Port Enable

Table 13-26. 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
<i>Not Used</i>																Ingress Control Shadow																664h
Ingress Port Enable																																668h
Negotiated Link Width for Ports 0, 1, 2, 3, 4, 5, 6, 7																																66Ch
<i>Reserved</i>																																670h – 67Ch

Register 13-136. 660h Ingress Control (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Factory Test Only</i>	RWS	Yes	0
1	Configuration Write to Device-Specific Register without Unsupported Request When set to 1, disables completions with an Unsupported Request status from being returned when Configuration Writes are attempted on Device-Specific registers.	RWS	Yes	0
21:2	<i>Factory Test Only</i>	RWS	Yes	0-0h
24:22	<i>Not used</i>	RWS	Yes	000b
25	<i>Factory Test Only</i>	RWS	Yes	0
26	Disable Upstream Port BAR0 and BAR1 Registers 0 = Enables the upstream port Base Address 0 and Base Address 1 registers (BAR0 and BAR1 , offsets 10h and 14h, respectively) 1 = Disables the upstream port Base Address 0 and Base Address 1 registers (BAR0 and BAR1 , offsets 10h and 14h, respectively)	RWS	Yes	0
27	<i>Not used</i>	RWS	Yes	0
28	Disable VGA BIOS Memory Access Decoding 0 = Enables the Bridge Control register <i>VGA 16-Bit Decode Enable</i> , <i>VGA Enable</i> , and <i>ISA Enable</i> bits (offset 3Ch[20:18], respectively), and enables decoding of PC ROM shadow addresses C0000h to CFFFFh (packets destined to these addresses are blocked) 1 = Disables the Bridge Control register <i>VGA 16-Bit Decode Enable</i> , <i>VGA Enable</i> , and <i>ISA Enable</i> bits (offset 3Ch[20:18], respectively), and disables decoding of PC ROM shadow addresses C0000h to CFFFFh (packets destined to these addresses are <i>not</i> blocked)	RWS	Yes	0
29	Disable Configuration Access from Upstream Port 0 = Enables Configuration access from the upstream port 1 = Disables Configuration access from the upstream port	RWS	Yes	0
30	<i>Factory Test Only</i>	RWS	Yes	0
31	<i>Not used</i>	RWS	Yes	0

Register 13-137. 664h Ingress Control Shadow (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RWS	Yes	0
1	Use Serial EEPROM Values for Ingress Credit Initialization 0 = Use default values for ingress credit initialization 1 = Use serial EEPROM values for ingress credit initialization	RWS	Yes	0
7:2	<i>Reserved</i>	RWS	Yes	0-0h
31:8	<i>Not used</i>	RWS	Yes	0000_00h

Register 13-138. 668h Ingress Port Enable (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	<p>Ingress Port Enable</p> <p>The value of this register depends upon the number of ports used, which is dependent upon the Port Configuration register <i>Port Configuration</i> field (Port 0, offset 224h[1:0]) value or STRAP_PORTCFG[1:0] ball strapping. Set the upper 24 bits to FFFF_FFh. For the lower 8 bits, set each bit position that corresponds to an enabled port.</p> <p>When a port is enabled in the corresponding station configuration, the bit is set to 1; otherwise, the bit is cleared to 0. <i>For example:</i></p> <ul style="list-style-type: none"> • When configured as x1x1x1x1x1x1x1x1, bits [7:0]=1111_1111b • When configured as x2x1x1x1x1x1x1x1, bits [7:0]=0111_1111b • When configured as x2x2x1x1x1x1x1x1, bits [7:0]=0011_1111b • When configured as x4x1x1x1x1x1x1x1, bits [7:0]=0001_1111b <p>Table 13-27 provides a sample mapping, based upon the Port Configuration value.</p>	RO	Yes	0000_0000h

Table 13-27. Ingress Port Configurations

Port Configuration Register Value (Port 0, Offset 224h[1:0])	Link Width/Port								Ingress Port Enable Register Value (Port 0, Offset 668h)
	0	1	2	3	4	5	6	7	
00b	x1	x1	x1	x1	x1	x1	x1	x1	Set the upper 16 bits to FFFFh. To determine the value for the lower 16 bits, start with 0h and set only those bit positions that correspond to the enabled ports. Bits [15:8] must be cleared to 0.
01b	x2	x1	x1	x1	x1	x1	x1		
10b	x2	x2	x1	x1	x1	x1			
11b	x4	x1	x1	x1	x1				

Register 13-139. 66Ch Negotiated Link Width for Ports 0, 1, 2, 3, 4, 5, 6, 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	Negotiated Link Width for Port 0 000b = x1 001b = x2 010b = x4 All other encodings are <i>reserved</i> .	RO	No	000b
3	Valid Negotiated Link Width for Port 0 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0
6:4	Negotiated Link Width for Port 1 000b = x1 001b = x2 All other encodings are <i>reserved</i> .	RO	No	000b
7	Valid Negotiated Link Width for Port 1 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0
10:8	Negotiated Link Width for Port 2 000b = x1 All other encodings are <i>reserved</i> .	RO	No	000b
11	Valid Negotiated Link Width for Port 2 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0
14:12	Negotiated Link Width for Port 3 000b = x1 All other encodings are <i>reserved</i> .	RO	No	000b
15	Valid Negotiated Link Width for Port 3 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0

Register 13-139. 66Ch Negotiated Link Width for Ports 0, 1, 2, 3, 4, 5, 6, 7 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
18:16	Negotiated Link Width for Port 4 000b = x1 All other encodings are <i>reserved</i> .	RO	No	000b
19	Valid Negotiated Link Width for Port 4 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0
22:20	Negotiated Link Width for Port 5 000b = x1 All other encodings are <i>reserved</i> .	RO	No	000b
23	Valid Negotiated Link Width for Port 5 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0
26:24	Negotiated Link Width for Port 6 000b = x1 All other encodings are <i>reserved</i> .	RO	No	000b
27	Valid Negotiated Link Width for Port 6 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0
30:28	Negotiated Link Width for Port 7 000b = x1 All other encodings are <i>reserved</i> .	RO	No	000b
31	Negotiated Link Width for Port 7 1 = Valid negotiated width allowing Cut-Thru from/to port.	RO	No	0

13.16.8 Device-Specific Registers – I/O CAM Base and Limit Upper 16 Bits

Table 13-28. Device-Specific I/O CAM Base and Limit Upper 16 Bits 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_0 Limit[31:16] Upper Port 0		I/O CAM_0 Base[31:16] Upper Port 0		680h
I/O CAM_1 Limit[31:16] Upper Port 1		I/O CAM_1 Base[31:16] Upper Port 1		684h
I/O CAM_2 Limit[31:16] Upper Port 2		I/O CAM_2 Base[31:16] Upper Port 2		688h
I/O CAM_3 Limit[31:16] Upper Port 3		I/O CAM_3 Base[31:16] Upper Port 3		68Ch
I/O CAM_4 Limit[31:16] Upper Port 4		I/O CAM_4 Base[31:16] Upper Port 4		690h
I/O CAM_5 Limit[31:16] Upper Port 5		I/O CAM_5 Base[31:16] Upper Port 5		694h
I/O CAM_6 Limit[31:16] Upper Port 6		I/O CAM_6 Base[31:16] Upper Port 6		698h
I/O CAM_7 Limit[31:16] Upper Port 7		I/O CAM_7 Base[31:16] Upper Port 7		69Ch
<i>Reserved</i>				6A0h – 6BCh

Register 13-140. 680h I/OCAM Upper Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-141. 684h I/OCAM Upper Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-142. 688h I/OCAM Upper Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-143. 68Ch I/OCAM Upper Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-144. 690h I/OCAM Upper Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-145. 694h I/OCAM Upper Port 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-146. 698h I/OCAM Upper Port 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

Register 13-147. 69Ch I/OCAM Upper Port 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	I/OCAM Base[31:16] I/O Base Upper 16 bits.	RW	Yes	0000h
31:16	I/OCAM Limit[31:16] I/O Limit Upper 16 bits.	RW	Yes	0000h

13.16.9 Device-Specific Registers – Base Address Shadow

The registers defined in Table 13-29 contain a shadow copy of the two Type 1 Configuration Base Address registers (**BAR0** and **BAR1**) for each PEX 8509 port.

Table 13-29. Device-Specific BAR Shadow 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
BAR0 Shadow for Port 5				6E8h
BAR1 Shadow for Port 5				6ECh
BAR0 Shadow for Port 6				6F0h
BAR1 Shadow for Port 6				6F4h
BAR0 Shadow for Port 7				6F8h
BAR1 Shadow for Port 7				6FCh
<i>Reserved</i>				700h – 73Ch

Register 13-148. 6C0h BAR0 Shadow for Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 0. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	<i>Reserved</i>	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 0 Base Address 0.	RW	Yes	0000h

Register 13-149. 6C4h BAR1 Shadow for Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 0 register <i>Memory Map Type</i> 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
	<i>Reserved</i> when the BAR0 Shadow for Port 0 register <i>Memory Map Type</i> field (offset 6C0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-150. 6C8h BAR0 Shadow for Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 1. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	<i>Reserved</i>	RsvdP	No	000h
31:17	Base Address 1 Shadow copy of Port 1 Base Address 0.	RW	Yes	0000h

Register 13-151. 6CCh BAR1 Shadow for Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 1 register <i>Memory Map Type</i> 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
	<i>Reserved</i> when the BAR0 Shadow for Port 1 register <i>Memory Map Type</i> field (offset 6C8h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-152. 6D0h BAR0 Shadow for Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 2. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	Reserved	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 2 Base Address 0.	RW	Yes	0000h

Register 13-153. 6D4h BAR1 Shadow for Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 2 register <i>Memory Map Type</i> 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 <i>Memory Map Type</i> field (offset 6D0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-154. 6D8h BAR0 Shadow for Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 3. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	<i>Reserved</i>	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 3 Base Address 0.	RW	Yes	0000h

Register 13-155. 6DCh BAR1 Shadow for Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 3 register <i>Memory Map Type</i> 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
	<i>Reserved</i> when the BAR0 Shadow for Port 3 register <i>Memory Map Type</i> field (offset 6D8h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-156. 6E0h BAR0 Shadow for Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 4. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	Reserved	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 4 Base Address 0.	RW	Yes	0000h

Register 13-157. 6E4h BAR1 Shadow for Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 4 register <i>Memory Map Type</i> 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 <i>Memory Map Type</i> field (offset 6E0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-158. 6E8h BAR0 Shadow for Port 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 5. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	<i>Reserved</i>	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 5 Base Address 0.	RW	Yes	0000h

Register 13-159. 6ECh BAR1 Shadow for Port 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 5 register <i>Memory Map Type</i> field (offset 6E8h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 5 Base Address 1[63:32].	RW	Yes	0000_0000h
	<i>Reserved</i> when the BAR0 Shadow for Port 5 register <i>Memory Map Type</i> field (offset 6E8h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-160. 6F0h BAR0 Shadow for Port 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 6. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	Reserved	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 6 Base Address 0.	RW	Yes	0000h

Register 13-161. 6F4h BAR1 Shadow for Port 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 6 register <i>Memory Map Type</i> field (offset 6F0h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 6 Base Address 1[63:32].	RW	Yes	0000_0000h
	Reserved when the BAR0 Shadow for Port 6 register <i>Memory Map Type</i> field (offset 6F0h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

Register 13-162. 6F8h BAR0 Shadow for Port 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Memory Space Indicator 0 = Memory BAR 1 = I/O BAR <i>Note: Hardwired to 0.</i>	RO	No	0
2:1	Memory Map Type Memory Mapping for Port 7. 00b = Base Address register is 32 bits wide and can be mapped anywhere in the 32-bit Memory space 10b = Base Address register is 64 bits wide and can be mapped anywhere in the 64-bit Address space 01b, 11b = <i>Reserved</i>	RO	Yes	00b
3	Prefetchable 0 = Not Prefetchable 1 = Prefetchable <i>Note: Hardwired to 0.</i>	RO	Yes	0
16:4	<i>Reserved</i>	RsvdP	No	000h
31:17	Base Address 0 Shadow copy of Port 7 Base Address 0.	RW	Yes	0000h

Register 13-163. 6FCh BAR1 Shadow for Port 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	Base Address 1[63:32] When the BAR0 Shadow for Port 7 register <i>Memory Map Type</i> field (offset 6F8h[2:1]) is set to 10b, 64-bit mode is enabled and this register becomes a Shadow copy of Port 7 Base Address 1[63:32].	RW	Yes	0000_0000h
	<i>Reserved</i> when the BAR0 Shadow for Port 7 register <i>Memory Map Type</i> field (offset 6F8h[2:1]) is not set to 10b.	RsvdP	Yes	0000_0000h

13.16.10 Device-Specific Registers – Shadow Virtual Channel Capability

Table 13-30. 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 Port 0 Capability		740h
<i>Reserved</i>		744h
VC0 Port 1 Capability		748h
<i>Reserved</i>		74Ch
VC0 Port 2 Capability		750h
<i>Reserved</i>		754h
VC0 Port 3 Capability		758h
<i>Reserved</i>		75Ch
VC0 Port 4 Capability		760h
<i>Reserved</i>		764h
VC0 Port 5 Capability		768h
<i>Reserved</i>		76Ch
VC0 Port 6 Capability		770h
<i>Reserved</i>		774h
VC0 Port 7 Capability		778h
<i>Reserved</i>		77Ch – 83Ch

Register 13-164. 740h VC0 Port 0 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_0 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_0 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_0 ID Port 0 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_0 Enable Port 0 Virtual Channel 0 Enable.	RO	No	1

Register 13-165. 748h VC0 Port 1 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_1 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_1 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_1 ID Port 1 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_1 Enable Port 1 Virtual Channel 0 Enable.	RO	No	1

Register 13-166. 750h VC0 Port 2 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_2 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_2 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_2 ID Port 2 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_2 Enable Port 2 Virtual Channel 0 Enable.	RO	No	1

Register 13-167. 758h VC0 Port 3 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_3 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_3 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_3 ID Port 3 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_3 Enable Port 3 Virtual Channel 0 Enable.	RO	No	1

Register 13-168. 760h VC0 Port 4 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_4 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_4 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_4 ID Port 4 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_4 Enable Port 4 Virtual Channel 0 Enable.	RO	No	1

Register 13-169. 768h VC0 Port 5 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_5 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_5 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_5 ID Port 5 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_5 Enable Port 5 Virtual Channel 0 Enable.	RO	No	1

Register 13-170. 770h VC0 Port 6 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_6 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_6 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_6 ID Port 6 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_6 Enable Port 6 Virtual Channel 0 Enable.	RO	No	1

Register 13-171. 778h VC0 Port 7 Capability (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	TC_VC0_7 Map[0] Always mapped to Virtual Channel 0. Tied to 1.	RO	No	1
7:1	TC_VC0_7 Map[7:1] Always mapped to Virtual Channel 0.	RW	Yes	7Fh
23:8	<i>Reserved</i>	RsvdP	No	000h
24	VC0_7 ID Port 7 Virtual Channel 0 ID.	RO	No	0
30:25	<i>Reserved</i>	RsvdP	No	00h
31	VC0_7 Enable Port 7 Virtual Channel 0 Enable.	RO	No	1

13.16.11 Device-Specific Registers – Ingress Credit Handler

Changing credit values from default register values must be done carefully. The total sum of the Header for all flows of all ports within the station must be $\leq 245d$, and the sum of all Payload credits must be $\leq 1,514d$; otherwise, the device will not function as expected. Also, there are minimal required Header credits for all the flows, that are required to achieve reasonable performance. The minimum initial Payload credits for Posted and Completions must exceed the required credits for a Maximum Payload Size TLP by 8.

Table 13-31. 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	
<i>Reserved</i>								INCH Port Pool Setting for Ports 0, 1, 2, 3																940h								
																INCH Port Pool Setting for Ports 4, 5, 6, 7																944h
<i>Reserved</i>																								948h –	9FCh							
																INCH Threshold Port 0 VC0 Posted																A00h
																INCH Threshold Port 0 VC0 Non-Posted																A04h
																INCH Threshold Port 0 VC0 Completion																A08h
<i>Reserved</i>																								A0Ch –	A14h							
																INCH Threshold Port 1 VC0 Posted																A18h
																INCH Threshold Port 1 VC0 Non-Posted																A1Ch
																INCH Threshold Port 1 VC0 Completion																A20h
<i>Reserved</i>																								A24h –	A2Ch							
																INCH Threshold Port 2 VC0 Posted																A30h
																INCH Threshold Port 2 VC0 Non-Posted																A34h
																INCH Threshold Port 2 VC0 Completion																A38h
<i>Reserved</i>																								A3Ch –	A44h							
																INCH Threshold, Port 3 VC0 Posted																A48h
																INCH Threshold, Port 3 VC0 Non-Posted																A4Ch
																INCH Threshold, Port 3 VC0 Completion																A50h
<i>Reserved</i>																								A54h –	A5Ch							
																INCH Threshold, Port 4 VC0 Posted																A60h
																INCH Threshold, Port 4 VC0 Non-Posted																A64h
																INCH Threshold, Port 4 VC0 Completion																A68h
<i>Reserved</i>																								A6Ch –	A74h							
																INCH Threshold, Port 5 VC0 Posted																A78h
																INCH Threshold, Port 5 VC0 Non-Posted																A7Ch
																INCH Threshold, Port 5 VC0 Completion																A80h
<i>Reserved</i>																								A84h –	A8Ch							

Table 13-31. Device-Specific Ingress Credit Handler (INCH) 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	
INCH Threshold, Port 6 VC0 Posted				A90h
INCH Threshold, Port 6 VC0 Non-Posted				A94h
INCH Threshold, Port 6 VC0 Completion				A98h
<i>Reserved</i>				A9Ch – AA4h
INCH Threshold, Port 7 VC0 Posted				AA8h
INCH Threshold, Port 7 VC0 Non-Posted				AACH
INCH Threshold, Port 7 VC0 Completion				AB0h
<i>Reserved</i>				AB4h – B7Ch
<i>Factory Test Only</i>				B80h

Register 13-172. 940h INCH Port Pool Setting for Ports 0, 1, 2, 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	Port 0 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 0. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
3	Unused 0 <i>Reserved</i> Program to 0. Additional bit for Port 0 Payload Pool.	RWS	Yes	0
6:4	Port 0 Header Pool Combined header credits (other than the initial credits) dedicated for Port 0. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
7	Unused 1 <i>Reserved</i> Should be kept at 0. Additional bit for Port 0 Header Pool.	RWS	Yes	0
10:8	Port 1 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 1. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
11	Unused 2 <i>Reserved</i> Program to 0. Additional bit for Port 1 Payload Pool.	RWS	Yes	0
14:12	Port 1 Header Pool Combined header credits (other than the initial credits) dedicated for Port 1. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
15	Unused 3 <i>Reserved</i> Should be kept at 0. Additional bit for Port 1 Header Pool.	RWS	Yes	0

Register 13-172. 940h INCH Port Pool Setting for Ports 0, 1, 2, 3 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
18:16	Port 2 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 2. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
19	Unused 4 <i>Reserved</i> Program to 0. Additional bit for Port 2 Payload Pool.	RWS	Yes	0
22:20	Port 2 Header Pool Combined header credits (other than the initial credits) dedicated for Port 2. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
23	Unused 5 <i>Reserved</i> Should be kept at 0. Additional bit for Port 2 Header Pool.	RWS	Yes	0
26:24	Port 3 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 3. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
27	Unused 6 <i>Reserved</i> Program to 0. Additional bit for Port 3 Payload Pool.	RWS	Yes	0
30:28	Port 3 Header Pool Combined header credits (other than the initial credits) dedicated for Port 3. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
31	Unused 7 <i>Reserved</i> Should be kept at 0. Additional bit for Port 3 Header Pool.	RWS	Yes	0

Register 13-173. 944h INCH Port Pool Setting for Ports 4, 5, 6, 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	Port 4 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 4. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
3	Unused 8 <i>Reserved</i> Program to 0. Additional bit for Port 4 Payload Pool.	RWS	Yes	0
6:4	Port 4 Header Pool Combined header credits (other than the initial credits) dedicated for Port 4. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
7	Unused 9 <i>Reserved</i> Should be kept at 0. Additional bit for Port 4 Header Pool.	RWS	Yes	0
10:8	Port 5 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 5. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
11	Unused 10 <i>Reserved</i> Program to 0. Additional bit for Port 5 Payload Pool.	RWS	Yes	0
14:12	Port 5 Header Pool Combined header credits (other than the initial credits) dedicated for Port 5. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
15	Unused 11 <i>Reserved</i> Should be kept at 0. Additional bit for Port 5 Header Pool.	RWS	Yes	0

Register 13-173. 944h INCH Port Pool Setting for Ports 4, 5, 6, 7 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
18:16	Port 6 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 6. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
19	Unused 12 <i>Reserved</i> Program to 0. Additional bit for Port 6 Payload Pool.	RWS	Yes	0
22:20	Port 6 Header Pool Combined header credits (other than the initial credits) dedicated for Port 6. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
23	Unused 13 <i>Reserved</i> Should be kept at 0. Additional bit for Port 6 Header Pool.	RWS	Yes	0
26:24	Port 7 Payload Pool Payload credits (other than the initial credits) for Posted/Completion TLPs dedicated to Port 7. 000b = 0 100b = 128 001b = 32 101b = 192 010b = 64 110b, 111b = 256 011b = 96	RWS	Yes	000b
27	Unused 14 <i>Reserved</i> Program to 0. Additional bit for Port 7 Payload Pool.	RWS	Yes	0
30:28	Port 7 Header Pool Combined header credits (other than the initial credits) dedicated for Port 7. 000b = 0 TLP 100b = 32 TLPs 001b = 4 TLPs 101b = 48 TLPs 010b = 8 TLPs 110b, 111b = 64 TLPs 011b = 16 TLPs	RWS	Yes	000b
31	Unused 15 <i>Reserved</i> Should be kept at 0. Additional bit for Port 7 Header Pool.	RWS	Yes	0

Register 13-174. A00h INCH Threshold Port 0 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 73d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-175. A04h INCH Threshold Port 0 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-176. A08h INCH Threshold Port 0 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-177. A18h INCH Threshold Port 1 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-178. A1Ch INCH Threshold Port 1 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-179. A20h INCH Threshold Port 1 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-180. A30h INCH Threshold Port 2 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-181. A34h INCH Threshold Port 2 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-182. A38h INCH Threshold Port 2 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-183. A48h INCH Threshold, Port 3 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-184. A4Ch INCH Threshold, Port 3 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-185. A50h INCH Threshold, Port 3 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-186. A60h INCH Threshold, Port 4 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-187. A64h INCH Threshold, Port 4 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-188. A68h INCH Threshold, Port 4 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-189. A78h INCH Threshold, Port 5 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-190. A7Ch INCH Threshold, Port 5 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-191. A80h INCH Threshold, Port 5 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-192. A90h INCH Threshold, Port 6 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-193. A94h INCH Threshold, Port 6 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-194. A98h INCH Threshold, Port 6 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-195. AA8h INCH Threshold, Port 7 VC0 Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Posted credits are used for VC0 Memory Write and Message transactions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 9d.	RWS	Yes	9h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-196. ACh INCH Threshold, Port 7 VC0 Non-Posted (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Non-Posted credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transactions.				
8:0	Payload <i>Reserved</i> Value of 000h indicates infinite credit.	RsvdP	Yes	000h
13:9	Header Header Initial Credit.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

Register 13-197. AB0h INCH Threshold, Port 7 VC0 Completion (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
Completion credits are used for VC0 Memory Read, I/O Read, I/O Write, Configuration Read, and Configuration Write transaction Completions.				
2:0	<i>Reserved</i>	RsvdP	No	000b
8:3	Payload Payload Credit/8 = 72d (8 * 9h) payload.	RWS	Yes	9h
13:9	Header Header Credit = 8d.	RWS	Yes	8h
15:14	FC Update High-Priority Threshold for Payload Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
17:16	FC Update High-Priority Threshold for Header Credit 00b = 75% 01b = 50% 10b = 25% 11b = 100%	RWS	Yes	00b
20:18	Payload Pool TLP payload pool reserved for this flow (in Payload credits). 000b = 0 001b = 16 010b = 32 011b = 48 100b, 101b, 110b, 111b = 64	RWS	Yes	000b
31:21	<i>Reserved</i>	RsvdP	No	000h

13.16.12 Device-Specific Registers – Port Configuration Header

Table 13-32. Device-Specific Port Configuration Header 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	
Command Port 1		Command Port 0		E00h
Command Port 3		Command Port 2		E04h
Command Port 5		Command Port 4		E08h
Command Port 7		Command Port 6		E0Ch
<i>Reserved</i>				E10h – E1Ch
Bridge Control Port 1		Bridge Control Port 0		E20h
Bridge Control Port 3		Bridge Control Port 2		E24h
Bridge Control Port 5		Bridge Control Port 4		E28h
Bridge Control Port 7		Bridge Control Port 6		E2Ch
<i>Reserved</i>				E30h – E3Ch

Register 13-198. E00h Command Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<p>I/O Access Enable</p> <p>0 = PEX 8509 ignores I/O accesses on Port 0’s primary interface 1 = PEX 8509 responds to I/O accesses on Port 0’s primary interface</p>	RW	Yes	0
1	<p>Memory Access Enable</p> <p>0 = PEX 8509 ignores Memory accesses on Port 0’s primary interface 1 = PEX 8509 responds to Memory accesses on Port 0’s primary interface</p>	RW	Yes	0
2	<p>Bus Master Enable</p> <p>Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction.</p> <p>0 = PEX 8509 handles Memory and I/O requests received on Port 0’s downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction</p>	RW	Yes	0
7:3	<i>Reserved</i>	RsvdP	No	0h
8	<p>SERR# Enable</p> <p>Controls the PCI Status register <i>Signaled System Error</i> bit.</p> <p>1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex</p>	RW	Yes	0
15:9	<i>Reserved</i>	RsvdP	No	00h

Register 13-199. E02h Command Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 1's primary interface 1 = PEX 8509 responds to I/O accesses on Port 1's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 1's primary interface 1 = PEX 8509 responds to Memory accesses on Port 1's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 1's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	Reserved	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	Reserved	RsvdP	No	00h

Register 13-200. E04h Command Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 2's primary interface 1 = PEX 8509 responds to I/O accesses on Port 2's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 2's primary interface 1 = PEX 8509 responds to Memory accesses on Port 2's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 2's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	<i>Reserved</i>	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	<i>Reserved</i>	RsvdP	No	00h

Register 13-201. E06h Command Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 3's primary interface 1 = PEX 8509 responds to I/O accesses on Port 3's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 3's primary interface 1 = PEX 8509 responds to Memory accesses on Port 3's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 3's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	Reserved	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	Reserved	RsvdP	No	00h

Register 13-202. E08h Command Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 4's primary interface 1 = PEX 8509 responds to I/O accesses on Port 4's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 4's primary interface 1 = PEX 8509 responds to Memory accesses on Port 4's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 4's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	<i>Reserved</i>	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	<i>Reserved</i>	RsvdP	No	00h

Register 13-203. E0Ah Command Port 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 5's primary interface 1 = PEX 8509 responds to I/O accesses on Port 5's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 5's primary interface 1 = PEX 8509 responds to Memory accesses on Port 5's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 5's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	Reserved	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	Reserved	RsvdP	No	00h

Register 13-204. E0Ch Command Port 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 6's primary interface 1 = PEX 8509 responds to I/O accesses on Port 6's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 6's primary interface 1 = PEX 8509 responds to Memory accesses on Port 6's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 6's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	<i>Reserved</i>	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	<i>Reserved</i>	RsvdP	No	00h

Register 13-205. E0Eh Command Port 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	I/O Access Enable 0 = PEX 8509 ignores I/O accesses on Port 7's primary interface 1 = PEX 8509 responds to I/O accesses on Port 7's primary interface	RW	Yes	0
1	Memory Access Enable 0 = PEX 8509 ignores Memory accesses on Port 7's primary interface 1 = PEX 8509 responds to Memory accesses on Port 7's primary interface	RW	Yes	0
2	Bus Master Enable Controls PEX 8509 Memory and I/O Request forwarding in the upstream direction. Neither affect message forwarding nor Completions in the upstream or downstream direction. 0 = PEX 8509 handles Memory and I/O requests received on Port 7's downstream/secondary interface as Unsupported Requests (UR); for Non-Posted requests, the PEX 8509 returns a Completion with UR completion status 1 = PEX 8509 forwards Memory and I/O requests in the upstream direction	RW	Yes	0
7:3	Reserved	RsvdP	No	0h
8	SERR# Enable Controls the PCI Status register <i>Signaled System Error</i> bit. 1 = Enables reporting of Fatal and Non-Fatal errors detected by the device to the Root Complex	RW	Yes	0
15:9	Reserved	RsvdP	No	00h

Register 13-206. E20h Bridge Control Port 0 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>SERR# Enable</p> <p>Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the Command Port 0 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>ISA Enable</p> <p>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).</p> <p>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.</p> <p>0 = Forward downstream all I/O addresses in the Address range defined by the I/O Base and I/O Limit registers</p> <p>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)</p>	RW	Yes	0
3	<p>VGA Enable</p> <p>Modifies the bridge response to VGA-compatible addresses.</p> <p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 0 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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</p> <p>1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-206. E20h Bridge Control Port 0 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-207. E22h Bridge Control Port 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>SERR# Enable</p> <p>Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the Command Port 1 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>ISA Enable</p> <p>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).</p> <p>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.</p> <p>0 = Forward downstream all I/O addresses in the Address range defined by the I/O Base and I/O Limit registers</p> <p>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)</p>	RW	Yes	0
3	<p>VGA Enable</p> <p>Modifies the bridge response to VGA-compatible addresses.</p> <p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 1 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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</p> <p>1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-207. E22h Bridge Control Port 1 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-208. E24h Bridge Control Port 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>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 Port 2 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>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.</p> <p>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)</p>	RW	Yes	0
3	<p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 2 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-208. E24h Bridge Control Port 2 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-209. E26h Bridge Control Port 3 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>SERR# Enable</p> <p>Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the Command Port 3 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>ISA Enable</p> <p>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).</p> <p>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.</p> <p>0 = Forward downstream all I/O addresses in the Address range defined by the I/O Base and I/O Limit registers</p> <p>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)</p>	RW	Yes	0
3	<p>VGA Enable</p> <p>Modifies the bridge response to VGA-compatible addresses.</p> <p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 3 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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</p> <p>1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-209. E26h Bridge Control Port 3 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-210. E28h Bridge Control Port 4 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>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 Port 4 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>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)</p>	RW	Yes	0
3	<p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 4 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> 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>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-210. E28h Bridge Control Port 4 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-211. E2Ah Bridge Control Port 5 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>SERR# Enable</p> <p>Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the Command Port 5 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>ISA Enable</p> <p>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).</p> <p>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.</p> <p>0 = Forward downstream all I/O addresses in the Address range defined by the I/O Base and I/O Limit registers</p> <p>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)</p>	RW	Yes	0
3	<p>VGA Enable</p> <p>Modifies the bridge response to VGA-compatible addresses.</p> <p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 5 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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</p> <p>1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-211. E2Ah Bridge Control Port 5 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-212. E2Ch Bridge Control Port 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>SERR# Enable</p> <p>Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the Command Port 6 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>ISA Enable</p> <p>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).</p> <p>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.</p> <p>0 = Forward downstream all I/O addresses in the Address range defined by the I/O Base and I/O Limit registers</p> <p>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)</p>	RW	Yes	0
3	<p>VGA Enable</p> <p>Modifies the bridge response to VGA-compatible addresses.</p> <p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 6 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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</p> <p>1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-212. E2Ch Bridge Control Port 6 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-213. E2Eh Bridge Control Port 7 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	<i>Reserved</i>	RsvdP	No	0
1	<p>SERR# Enable</p> <p>Controls forwarding of ERR_COR, ERR_FATAL, and ERR_NONFATAL from the secondary interface to the primary interface.</p> <p>When set to 1, and the Command Port 7 register <i>SERR# Enable</i> bit is set to 1, enables the PCI Status register <i>Signaled System Error</i> bit.</p>	RW	Yes	0
2	<p>ISA Enable</p> <p>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).</p> <p>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.</p> <p>0 = Forward downstream all I/O addresses in the Address range defined by the I/O Base and I/O Limit registers</p> <p>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)</p>	RW	Yes	0
3	<p>VGA Enable</p> <p>Modifies the bridge response to VGA-compatible addresses.</p> <p>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):</p> <ul style="list-style-type: none"> 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) <p>When set to 1, forwarding of these addresses is independent of the:</p> <ul style="list-style-type: none"> 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 2 (<i>ISA Enable</i>) setting <p>Forwarding of these addresses is qualified by the Command Port 7 register <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits. The default state of this bit after reset must be 0.</p> <p>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</p> <p>1 = Forward VGA-compatible Memory and I/O addresses (addresses defined above) from the primary interface to the secondary interface (when the <i>I/O Access Enable</i> and <i>Memory Access Enable</i> bits are set), independent of the Memory and I/O Address ranges and independent of the <i>ISA Enable</i> bit</p>	RW	Yes	0

Register 13-213. E2Eh Bridge Control Port 7 (Only Port 0) (Cont.)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4	<p>VGA 16-Bit Enable</p> <p>Used only when bit 3 (<i>VGA Enable</i>) is also set to 1, enabling VGA I/O decoding and forwarding by the bridge. Status after reset is 0.</p> <p>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.</p> <p>0 = Execute 10-bit address decodes on VGA I/O accesses 1 = Execute 16-bit address decodes on VGA I/O accesses</p>	RW	Yes	0
15:5	<i>Reserved</i>	RsvdP	No	0-0h

13.16.13 Device-Specific Registers – Source Queue Weight and Soft Error

Table 13-33. Device-Specific Source Queue Weight and Soft Error 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	
<i>Reserved</i>		Port Egress TLP Threshold		F10h
Source Queue Weight				F14h
Soft Error Counters 1				F18h
<i>Reserved</i>		Soft Error Counters 2		F1Ch
<i>Reserved</i>				F20h – F28h
Soft Error Counters 6		<i>Reserved</i>		F2Ch
Soft Error Injection				F30h
<i>Reserved</i>				F34h – FB0h

Register 13-214. F10h Port Egress TLP Threshold (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
7:0	Port Lower TLP Count When Source Scheduling is disabled due to threshold, Source Scheduling is re-enabled when the Port TLP Count goes below this threshold.	Upstream	RWS	Yes	FFh
	<i>Reserved</i>	Downstream	RsvdP	No	00h
15:8	<i>Reserved</i>		RsvdP	No	00h
23:16	Port Upper TLP Count When the Port TLP Count is greater than or equal to this value, the Source Scheduler disables TLP Scheduling to this egress port.	Upstream	RWS	Yes	FFh
	<i>Reserved</i>	Downstream	RsvdP	No	00h
31:24	<i>Reserved</i>		RsvdP	No	00h

Register offset **F14h** defines the Source Queue weight for each downstream port, where **a, b, c, d, e, f, g,** and/or **h** does not represent the current destination port source queue, because the same source port does not transmit TLPs to the same destination port. These registers are also loaded from the Weighted Round-Robin Port Arbitration table [refer to [Table 13-15, “Port Arbitration Table Register Map \(Only Upstream Port\)”](#)], when software writes to load the Port Arbitration Table registers (**VC0 Resource Control** register *Load Port Arbitration Table* bit (offset **15Ch**[16])).

Register 13-215. F14h Source Queue Weight (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
2:0	src_que_a_weight Source Scheduler A schedules src_que_a_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
3	<i>Reserved</i>	RsvdP	No	0
6:4	src_que_b_weight Source Scheduler B schedules src_que_b_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
7	<i>Reserved</i>	RsvdP	No	0
10:8	src_que_c_weight Source Scheduler C schedules src_que_c_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
11	<i>Reserved</i>	RsvdP	No	0
14:12	src_que_d_weight Source Scheduler D schedules src_que_d_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
15	<i>Reserved</i>	RsvdP	No	0
18:16	src_que_e_weight Source Scheduler E schedules src_que_e_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
19	<i>Reserved</i>	RsvdP	No	0
22:20	src_que_f_weight Source Scheduler F schedules src_que_f_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
23	<i>Reserved</i>	RsvdP	No	0
26:24	src_que_g_weight Source Scheduler G schedules src_que_g_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
27	<i>Reserved</i>	RsvdP	No	0
30:28	src_que_h_weight Source Scheduler H schedules src_que_h_weight (range 1 to 8) TLPs back-to-back when another source queue(s) to the same downstream port is not empty.	RWS	Yes	000b
31	<i>Reserved</i>	RsvdP	No	0

Register 13-216. F18h Soft Error Counters 1 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	Packet RAM 1-Bit Soft Error Counter	RO	No	00h
15:8	PLL RAM 1-Bit Soft Error Counter	RO	No	00h
23:16	TLP ID RAM 1-Bit Soft Error Counter	RO	No	00h
31:24	Scheduler RAM 1-Bit Soft Error Counter	RO	No	00h

Register 13-217. F1Ch Soft Error Counters 2 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
7:0	DQLL RAM 1-Bit Soft Error Counter	RO	No	00h
15:8	SQLL RAM 1-Bit Soft Error Counter Station 1 Packet RAM Instance 2 1-Bit Soft Error Counter	RO	No	00h
31:16	<i>Reserved</i>	RsvdP	No	0000h

Register 13-218. F2Ch Soft Error Counters 6 (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
15:0	<i>Reserved</i>	RsvdP	No	0000h
23:16	64 Entry Retry Buffer 1-Bit Soft Error Counter	RO	No	00h
31:24	32 Entry Retry Buffer 1-Bit Soft Error Counter	RO	No	00h

Register 13-219. F30h Soft Error Injection (Only Port 0)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Destination Queue Link List RAM Instance 0 1-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
1	Destination Queue Link List RAM Instance 0 2-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
2	Destination Queue Link List RAM Instance 0 Error Injection Field Select 0 = Inject Soft error in ECC code 1 = Inject Soft error in data	RWS	Yes	0
3	Destination Queue Link List RAM Instance 1 1-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
4	Destination Queue Link List RAM Instance 1 2-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
5	Destination Queue Link List RAM Instance 1 Error Injection Field Select 0 = Inject Soft error in ECC code 1 = Inject Soft error in data	RWS	Yes	0
6	64 Entry Retry Buffer 1-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
7	64 Entry Retry Buffer 2-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
8	64 Entry Retry Buffer Soft Error Injection Field Select 0 = Inject Soft error in ECC code 1 = Inject Soft error in data	RWS	Yes	0
9	32 Entry Retry Buffer 1-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
10	32 Entry Retry Buffer 2-Bit Soft Error Injection Toggle Every toggle injects an error.	RWS	Yes	0
11	32 Entry Retry Buffer Soft Error Injection Field Select 0 = Inject Soft error in ECC code 1 = Inject Soft error in data	RWS	Yes	0
31:12	<i>Reserved</i>	RsvdP	No	0-0h

13.17 Advanced Error Reporting Extended Capability Registers

Table 13-34. Advanced Error Reporting 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 (138h)	Capability Version (1h)	PCI Express Extended Capability ID (0001h)	FB4h
Uncorrectable Error Status			FB8h
Uncorrectable Error Mask			FBCh
Uncorrectable Error Severity			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
<i>Reserved</i>			FE0h – FFCh

Register 13-220. FB4h Advanced Error Reporting Enhanced Capability Header (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	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 is the Power Budget Extended Capability structure.	RO	Yes	138h

Register 13-221. FB8h Uncorrectable Error Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	<i>Reserved</i>	RsvdP	No	0h
4	Data Link Protocol Error Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
5	Surprise Down Error Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
11:6	<i>Reserved</i>	RsvdP	No	0000_00b
12	Poisoned TLP Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
13	Flow Control Protocol Error Status	RsvdP ^b	No	0
14	Completion Timeout Status	RsvdP ^b	No	0
15	Completer Abort Status	RWCS ^a	Yes	0
16	Unexpected Completion Status 0 = No error is detected 1 = Error is detected	RsvdP ^b	No	0
17	Receiver Overflow Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
18	Malformed TLP Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
19	ECRC Error Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
20	Unsupported Request Error Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
31:21	<i>Reserved</i>	RsvdP	No	0-0h

a. When the **ECC Error Check Disable** register **Software Force Error Enable** bit is set (Port 0, offset 1C8h[2]=1), Type changes from RWCS to RW.

b. Bits [16, 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 13-222. FBCh Uncorrectable Error Mask (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
3:0	<i>Reserved</i>	RsvdP	No	0h
4	Data Link Protocol Error Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
5	Surprise Down Error Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
11:6	<i>Reserved</i>	RsvdP	No	0000_00b
12	Poisoned TLP Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
13	Flow Control Protocol Error Mask	RWS	Yes	0
14	Completion Timeout Mask	RWS	No	0
15	Completer Abort Mask	RWS	Yes	0
16	Unexpected Completion Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
17	Receiver Overflow Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
18	Malformed TLP Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
19	ECRC Error Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
20	Unsupported Request Error Mask 0 = No mask is set 1 = Masks error reporting, first error update, and header logging for this error	RWS	Yes	0
31:21	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-223. FC0h Uncorrectable Error Severity (All Ports)

Bit(s)	Description	Ports	Type	Serial EEPROM and I ² C	Default
3:0	<i>Reserved</i>		RsvdP	No	0h
4	Data Link Protocol Error Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	1
5	<i>Reserved</i>	Upstream	RsvdP	No	1
	Surprise Down Error Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal	Downstream	RWS	Yes	1
11:6	<i>Reserved</i>		RsvdP	No	0000_00b
12	Poisoned TLP Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	0
13	Flow Control Protocol Error Severity		RWS	Yes	1
14	Completion Timeout Severity		RWS	No	0
15	Completer Abort Severity		RWS	Yes	0
16	Unexpected Completion Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	0
17	Receiver Overflow Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	1
18	Malformed TLP Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	1
19	ECRC Error Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	0
20	Unsupported Request Error Severity 0 = Error is reported as non-fatal 1 = Error is reported as fatal		RWS	Yes	0
31:21	<i>Reserved</i>		RsvdP	No	0-0h

Register 13-224. FC4h Correctable Error Status (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Receiver Error Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
5:1	Reserved	RsvdP	No	0-0h
6	Bad TLP Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
7	Bad DLLP Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
8	Replay Number Rollover Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
11:9	Reserved	RsvdP	No	000b
12	Replay Timer Timeout Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
13	Advisory Non-Fatal Error Status 0 = No error is detected 1 = Error is detected	RWCS ^a	Yes	0
31:14	Reserved	RsvdP	No	0-0h

- a. When the **ECC Error Check Disable** register *Software Force Error Enable* bit is set (Port 0, offset 1C8h[2]=1), Type changes from RWCS to RW.

Register 13-225. FC8h Correctable Error Mask (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
0	Receiver Error Mask 0 = Error reporting is not masked 1 = Error reporting is masked	RWS	Yes	0
5:1	<i>Reserved</i>	RsvdP	No	0-0h
6	Bad TLP Mask 0 = Error reporting is not masked 1 = Error reporting is masked	RWS	Yes	0
7	Bad DLLP Mask 0 = Error reporting is not masked 1 = Error reporting is masked	RWS	Yes	0
8	Replay Number Rollover Mask 0 = Error reporting is not masked 1 = Error reporting is masked	RWS	Yes	0
11:9	<i>Reserved</i>	RsvdP	No	000b
12	Replay Timer Timeout Mask 0 = Error reporting is not masked 1 = Error reporting is masked	RWS	Yes	0
13	Advisory Non-Fatal Error Mask 0 = Error reporting is not masked 1 = Error reporting is masked	RWS	Yes	1
31:14	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-226. FCCh Advanced Error Capabilities and Control (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
4:0	First Error Pointer Identifies the bit position of the first error reported in the Uncorrectable Error Status register (offset FB8h).	ROS	No	1Fh
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	<i>Reserved</i>	RsvdP	No	0-0h

Register 13-227. FD0h Header Log 0 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	TLP Header 0 First DWord header. TLP header associated with error.	ROS	Yes	0-0h

Register 13-228. FD4h Header Log 1 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	TLP Header 1 Second DWord header. TLP header associated with error.	ROS	Yes	0-0h

Register 13-229. FD8h Header Log 2 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	TLP Header 2 Third DWord header. TLP header associated with error.	ROS	Yes	0-0h

Register 13-230. FDCh Header Log 3 (All Ports)

Bit(s)	Description	Type	Serial EEPROM and I ² C	Default
31:0	TLP Header 3 Fourth DWord header. TLP header associated with error.	ROS	Yes	0-0h

THIS PAGE INTENTIONALLY LEFT BLANK.



14.1 Physical Layer Loopback Operation

14.1.1 Overview

Physical Layer loopback functions are used to test SerDes in the PEX 8509, connections between devices, and SerDes of external devices, as well as various PEX 8509 and external digital logic. The PEX 8509 supports five types of loopback operations, as described in [Table 14-1](#). Additional information regarding each type is provided in the sections that follow.

Table 14-1. Loopback Operations


Operation	Description
Internal Loopback Mode	Internal Loopback 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 Loopback Master Mode	Analog Loopback Master mode depends upon an external device or dumb connection (<i>such as</i> a cable) to loopback the transmitted data to the PEX 8509. If an external device is used, it must not include its elastic buffer in the loopback data path, because no SKIP Ordered-Sets are transmitted. Use the PRBS generator and checker to create and check the data pattern. The PEX 8509 enters Analog Loopback Master mode when the Physical Layer Port Command register <i>Port x Loopback Command</i> bit (Port 0, offset 230h[0, 4, 8, 12, 16, 20, 24, or 28]) is set.
Digital Loopback Master Mode	As with the Analog Loopback Master mode, Digital Loopback Master mode depends upon an external device to loopback the transmitted data. This method is best utilized with an external device that includes at least its elastic buffer in the loopback data path. The PEX 8509 provides a programmable data pattern generator and checker that inserts the SKIP Ordered-Set at the proper intervals. The PEX 8509 enters Digital Loopback Master mode when the Physical Layer Port Command register <i>Port x Loopback Command</i> bit (Port 0, offset 230h[0, 4, 8, 12, 16, 20, 24, or 28]) is set.
Analog Loopback Slave Mode	The PEX 8509 enters Analog Loopback Slave mode when an external device transmits training sets with the <i>Loopback Training Control</i> bit set and the Physical Layer Test register <i>Analog Loopback Enable</i> 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. <i>Note:</i> There are no serializers nor de-serializers in the loopback data path.
Digital Loopback Slave Mode	The PEX 8509 enters Digital Loopback Slave mode when an external device transmits training sets with the <i>Loopback Training Control</i> bit set and the Physical Layer Test register <i>Analog Loopback Enable</i> 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 8509's elastic buffer, 8b/10b decoder, and 8b/10b encoder in the loopback data path. <i>Note:</i> There are no serializers nor de-serializers in the loopback data path.

14.1.2 Internal Loopback Mode

Figure 14-1 illustrates the loopback data path when Internal Loopback mode is enabled. The only items in the data path are the SerDes. Internal Loopback mode is used only when SerDes Built-In Self-Test (BIST) is enabled (**Physical Layer Test** register *PHY BIST Enable* bit, Port 0, offset 228h[7]=1).

The SerDes BIST is intended to be overlapped with the serial EEPROM load operation. To achieve this overlap, the *PHY BIST Enable* bit is written early in the serial EEPROM load operation. After the *PHY BIST Enable* bit is set, the SerDes is placed in Internal Loopback 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 14-1. Internal Loopback Mode (Analog near End) Data Path



14.1.3 Analog Loopback Master Mode

Analog Loopback Master mode is typically used for Analog Far-End testing, as illustrated in Figure 14-2.

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 14-3.

Figure 14-2. Analog Far-End Loopback


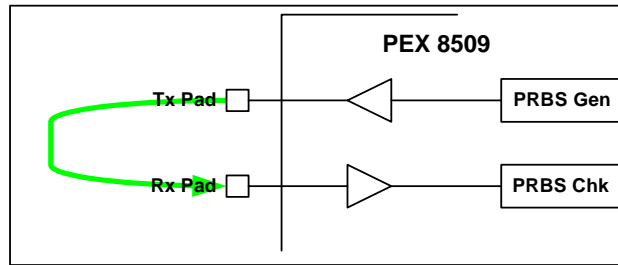


Figure 14-3. Cable Loopback



To cause a PEX 8509 port to request to become a Loopback Master:

1. After the link is up, a Configuration Write to the appropriate **Physical Layer Port Command** register *Port x Loopback Command* bit (Port 0, offset 230h[0, 4, 8, 12, 16, 20, 24, or 28]), causes the port to transition from the L0 Link PM state to Recovery, and then to the Loopback state:
 - If a cable is used for the loopback, the port transitions from the Configuration state to the Loopback 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 Loopback* bit (Port 0, offset 228h[21:20]), the link with the cable can reach the L0 Link PM state and not go to the Loopback state.
 - Cable length is limited only by the PCI Express drivers and cable properties.
2. After the port is in the Loopback state, the corresponding **Physical Layer Port Command** register *Port x Ready as Loopback Master* bit (Port 0, offset 230h[3, 7, 11, 15, 19, 23, 27, or 31]) 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[17:16]) associated with the SerDes assigned to the port being tested.
 - The PRBS checker checks the returned PRBS data. Any errors are logged in the **SerDes Quad x Diagnostic Data** register(s) (Port 0, offsets 238h and 23Ch) that corresponds to the SerDes quad being tested.

14.1.4 Digital Loopback Master Mode


The only difference between Analog and Digital Loopback Master modes is that the external device is assumed to retain, to some extent, its digital logic in the loopback data path. Because this includes the elastic buffer, SKIP Ordered-Sets must be included in the test pattern. For the PEX 8509, this precludes PRBS engine use, because the PRBS generator does not generate SKIP Ordered-Sets.

The PEX 8509 provides the programmable test pattern transmitter for Digital Far-End Loopback testing, as illustrated in Figure 14-4. After Digital Loopback Master mode is established, Configuration Writes are used to fill the **Physical Layer 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[29:28]) 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[17:16]) 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 and Port Control** 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 Loopback 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 **SerDes Quad x Diagnostic Data** register(s) (Port 0, offsets 238h and 23Ch).

Figure 14-4. Digital Far-End Loopback




14.1.5 Analog Loopback Slave Mode

The PEX 8509 becomes an Analog Loopback Slave if it receives training sets with the *Loopback Training Control* bit set while the **Physical Layer Test** register *Analog Loopback Enable* bit (Port 0, offset 228h[4]) is set. While it is an Analog Loopback Slave, the PEX 8509 includes only the SerDes in the loopback data path. The Loopback Master must provide the test pattern and data pattern checking. It is unnecessary for the Loopback Master to include SKIP Ordered-Sets in the data pattern.

Figure 14-5 illustrates the loopback data path when Analog Loopback Slave mode is enabled.

Note: There is no scrambling nor de-scrambling logic in the Slave analog loopback data path.

Figure 14-5. Analog Loopback Slave Mode




14.1.6 Digital Loopback Slave Mode

The PEX 8509 becomes a Digital Loopback Slave if it receives training sets with the *Loopback Training Control* bit set while the **Physical Layer Test** register *Analog Loopback Enable* bit (Port 0, offset 228h[4]) is cleared.

When a PEX 8509 port is a Digital Loopback Slave, the port includes the elastic buffer and 8b/10b decoder and encoder in the loopback data path. The Loopback Master must provide the test pattern and data pattern checker. The Loopback Master must also transmit SKIP Ordered-Sets with the data pattern. Because the PEX 8509 can return more or fewer SKIP symbols than it receives, the data checker must make provisions for this possibility.

Note: There is no scrambling nor de-scrambling logic in the Slave digital loopback data path.

Figure 14-6. Digital Loopback Slave Mode



14.2 Using the Diagnostic Registers

There are four diagnostic registers per station, one for each SerDes quad. The **SerDes Quad *x* Diagnostic Data** register(s) (Port 0, offsets [238h](#) and [23Ch](#)) 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 cleared to 00b, the information in that Diagnostic register is for Quad 0, SerDes 0 for Port 0. [Table 14-2](#) further illustrates this example.

Table 14-2. SerDes Register Contents When *SerDes Diagnostic Data Select* Field = 00b

Port 0 Register Offset	Register	Port 0
238h	SerDes Quad 0 Diagnostic Data	SerDes 0
23Ch	SerDes Quad 1 Diagnostic Data	SerDes 4

14.3 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 bit stream is used for internal SerDes or Analog Far-End Loopback testing.

The PEX 8509 also provides a method of creating a repeating programmable bit pattern. Each of the four 32-bit **Physical Layer 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 Loopback Master, set the **Physical Layer Test** register *User Test Pattern Enable* bit(s) (Port 0, offset 228h[29:28]) to 1, for the SerDes quad(s) associated with that port. The PEX 8509 proceeds to transmit the data pattern on all lanes of the quad, starting with Byte 0 of the **Physical Layer User Test Pattern 0** register and continuing, in sequence, through Byte 3 of the **Physical Layer User Test Pattern 12** register. SKIP Ordered-Sets are inserted at the proper intervals, which makes this method appropriate for Digital Far-End Loopback testing. The received pattern is compared to the transmitted pattern. Any errors are logged and can be retrieved, by reading the **SerDes Quad x Diagnostic Data** register(s) (Port 0, offsets 238h and 23Ch, RO bits [30, 23:0]).

The error status for the PRBS or User Pattern method is selected separately for each lane of the SerDes quad. SerDes quad lane status is selected by manipulating the *SerDes Diagnostic Data Select* bits in the appropriate **SerDes Quad x Diagnostic Data** register(s) (Port 0, offsets 238h and 23Ch[25:24]).

The following example of User Test Pattern generation transmits and compares a clock-like (0-1-0) data pattern in Analog Loopback mode:

To produce a pseudo-clock bitstream in Analog Loopback mode, set the registers as follows:

1. In the Slave device, enable Analog Loopback by setting the **Physical Layer Test** register *Analog Loopback Enable* bit (Port 0, offset 228h[4]).
2. In the PEX 8509 Loopback Master device:
 - a. Write the value 4A4A_4A4Ah into each of the **Physical Layer User Test Pattern x** registers (Port 0, offsets 210h through 21Ch).
 - b. Set the *Port x Loopback Command* bit for the specific port in the **Physical Layer Port Command** register (Port 0, offset 230h[0, 4, 8, 12, 16, 20, 24, or 28]).
 - c. To check whether loopback is successful, read the corresponding **Physical Layer Port Command** register *Port x Ready as Loopback Master* bit (Port 0, offset 230h[3, 7, 11, 15, 19, 23, 27, or 31]) in the same Nibble that was set in step a. The Nibble value will be 9h if loopback is successful.
 - d. Set the **Physical Layer Test** register *User Test Pattern Enable* bit(s) (Port 0, offset 228h[29:28]) 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 Loopback 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.

14.4 JTAG Interface

The PEX 8509 provides a Joint Test Action Group (JTAG) Boundary Scan interface, which is utilized to debug board connectivity for each ball.

14.4.1 *IEEE 1149.1* and *IEEE 1149.6* Test Access Port

The *IEEE 1149.1* Test Access Port (TAP), commonly called the *JTAG Debug port*, is an architectural standard described in the *IEEE Standard 1149.1-1990*. The *IEEE Standard 1149.6-2003* defines extensions to *IEEE 1149.1* to support PCI Express SerDes testing. These standards describe methods for accessing internal device 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 the *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 14.4.4](#)

14.4.2 JTAG Instructions

The JTAG Debug port provides the *IEEE Standard 1149.1-1990* EXTEST, SAMPLE/PRELOAD, BYPASS, CLAMP, and IDCODE instructions. *IEEE Standard 1149.6-2003* EXTEST_PULSE and EXTEST_TRAIN instructions are also supported. [Table 14-3](#) lists the JTAG instructions, along with their input codes.

The PEX 8509 returns the JTAG IDCODE values listed in [Table 14-4](#).

Table 14-3. JTAG Instructions

Instruction	Input Code	Comments
BYPASS	1111_1111_1111_1111_1111_1111_1111b	<i>IEEE Standard 1149.1-1990</i>
EXTEST	1111_1111_1111_1111_1111_1110_1000b	
SAMPLE	1111_1111_1111_1111_1111_1111_1000b	
PRELOAD	1111_1111_1111_1111_1111_1111_1000b	
EXTEST_PULSE	1111_1111_1011_1111_1111_1110_1000b	<i>IEEE Standard 1149.6-2003</i>
EXTEST_TRAIN	1111_1110_1001_1111_1111_1110_1000b	
CLAMP	1111_1111_1111_1111_1111_1110_1111b	<i>IEEE Standard 1149.1-1990</i>
IDCODE	1111_1111_1111_1111_1111_1111_1110b	

Table 14-4. JTAG IDCODE Values

Unit of Measure	Version	Part Number	PLX Manufacturer Identity	Least Significant Bit
Bits	0001b	1000_0101_0000_1001b	001_1100_1101b	1
Hex	1h	8509h	1CDh	1h
Decimal	1	34057	461	1

14.4.3 JTAG Boundary Scan

Scan Description Language (BSDL), IEEE 1149.1-1994, is a supplement to the 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 device's signal 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 device's 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 device in the various test modes defined by the standard. Instruction Set statements also support descriptions of instructions that are unique to the PEX 8509.

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 includes 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

14.4.4 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 8509's JTAG TAP Controller to initialize. In addition, when the JTAG TAP Controller is initialized, it selects the PEX 8509 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 8509 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 8509's JTAG TAP Controller is not intended to be used by the design, it is recommended that a 1.5K Ω 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

14.5 Lane Good Status LEDs

The PEX 8509 provides Lane Good outputs, `PEX_LANE_GOOD[7:0]#`, that can directly drive external common anode LED modules to provide visual indication that the Physical Layer of the link for each lane is trained to at least x1 width.

Software can determine:

- Which lanes have completed Physical Layer linkup, by performing a Memory Read of the *Lane Up Status* bits in Port 0:
 - Offset `1F4h[7:0]` corresponds to Lanes [7-0], respectively
- Whether the link for each port has trained, by reading the **VC0 Resource Status** register *VC0 Negotiation Pending* bit (offset `160h[17]`) in each port. If the register value is 0, the link has completed Flow Control initialization. This register 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 15 Electrical Specifications

15.1 Introduction

This chapter contains the PEX 8509 power-up/power-down sequencing rules and electrical specifications.

15.2 Power-Up/Power-Down Sequence

For reliable operation, the **VDD10**, **VDD10S**, and **VDD10A** should 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.

15.3 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 PEX 8509 at these limits is not recommended.*

Table 15-1. 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
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 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}	-65 to +150	°C

a. The SerDes Analog and Digital power supplies should track within 0.01V of one another.

15.4 Power Characteristics

Table 15-2. Operating Condition Power Supply Rails

Symbol	Parameter	Min	Typ	Max	Units
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
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
VTT_PEX	SerDes Termination Supply Voltage	1.35	1.5	1.8	V

15.5 Power Consumption

Table 15-3. PEX 8509 Power Estimates (Watts) (4 to 8 Lanes)

Traffic Conditions	Core/SerDes Digital (VDD10)		PCI Express Digital (VDD10S)		PCI Express Analog (VDD10A)		SerDes Termination (VTT_PEX)		PLL (VDD33A)		I/O (VDD33)		Total (Watts)	
	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max	Typ	Max
A. Heavy	0.60	0.89	0.38	0.43	0.11	0.12	0.26	0.32	0.02	0.03	0.02	0.03	1.38	1.81
B. Medium	0.50	0.74	0.32	0.36	0.09	0.10	0.22	0.26	0.02	0.03	0.02	0.03	1.16	1.53
C. Light	0.40	0.59	0.27	0.30	0.07	0.08	0.18	0.22	0.02	0.03	0.02	0.03	0.96	1.26

A. 85% lane bandwidth utilization. All 8 lanes in active L0 Link PM state.

B. 35% lane bandwidth utilization. All 8 lanes in active L0 Link PM state.

C. 10% lane bandwidth utilization. All 8 lanes in active L0 Link PM state.

15.6 I/O Interface Signal Groupings

Table 15-4. Signal Group PCI Express Analog Interface

Signal Group	Signal Type	Signals	Notes
(a)	PCI Express Output (Transmit)	PEX_PETn[7:0], PEX_PETp[7:0]	Refer to Table 15-7
(b)	PCI Express Input (Receive)	PEX_PERn[7:0], PEX_PERp[7:0]	Refer to Table 15-8
(c)	PCI Express Differential Clock Input	PEX_REFCLKn, PEX_REFCLKp	Refer to Table 15-9

Table 15-5. Signal Group Digital Interface

Signal Group	Signal Type	Signals	Note
(d)	Digital Output	EE_CS#, EE_DI, EE_SK, FATAL_ERR#, HP_ATNLED[3:1]#, HP_CLKEN[3:1]#, HP_PERST[3:1]#, HP_PWREN[3:1]#, HP_PWRLED[3:1]#, JTAG_TDO, PEX_LANE_GOOD[7:0]#	Refer to Table 15-6
(e)	Digital Input ^a	PEX_PERST#, STRAP_DEBUG_SEL[1:0]#, STRAP_FAST_BRINGUP#, STRAP_PLL_BYPASS#, STRAP_PROBE_MODE#, STRAP_SERDES_MODE_ENABLE#, STRAP_PORTCFG[1:0], STRAP_TESTMODE[3:0], STRAP_UPSTRM_PORTSEL[2:0]	
(f)	Digital Input with Internal Pull-up Resistor	EE_DO, HP_BUTTON[3:1]#, HP_MRL[3:1]#, HP_PRSNT[3:1]#, HP_PWRFLT[3:1]#, I2C_ADDR[2:0], JTAG_TCK, JTAG_TDI, JTAG_TMS, JTAG_TRST#	
(g)	Bidirectional (Open Drain)	I2C_SCL, I2C_SDA, PEX_INTA#	

a. *STRAP_* signals must be tied High to VDD33 or Low to VSS (GND). Refer to [Section 3.4.4](#), “Strapping Signals.” for details.

Table 15-6. 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}	(d) (e) (f) (g)	Ball Capacitance		5	8	pF	
$I_{LEAKAGE}$	(d)	Three-state Leakage			± 500	nA	
	(e) (f)	Input Leakage			± 50	nA	
	(f)	Pull-Up Leakage	+0.1/-8		+0.1/-20	μA	
R_{PU}	(f)	Pull-Up Impedance	200K			Ω	
V_{HYS}	(g)	Input Hysteresis	150			mV	

Table 15-7. 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 * V_{TX-D+} - V_{TX-D-} $
$V_{TX-DE-RATIO}$	De-Emphasized Differential Output Voltage (Ratio)	-3.0	-3.5	-4.0	dB	Ratio of the $V_{TX-DIFFp-p}$ of the 2 nd and following bits after a transition divided by the $V_{TX-DIFFp-p}$ of the 1 st 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: $T_{TX-MAX-JITTER} = 1 - T_{TX-EYE} = 0.25 UI$ 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	$V_{TX-CM-ACp} = \text{RMS}(V_{TX-D+} + V_{TX-D-} / 2 - V_{TX-CM-DC})$ $V_{TX-CM-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} / 2$ Refer to Note 1.
$V_{TX-CM-DC-ACTIVE-IDLE-DELTA}$	Absolute Delta of DC Common Mode Voltage during L0 Link PM state and Electrical Idle	0		100	mV	$ V_{TX-CM-DC} [\text{during L0}] - V_{TX-CM-Idle-DC} [\text{during Electrical Idle}] \leq 100mV$ $V_{TX-CM-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} / 2 [L0]$ $V_{TX-CM-Idle-DC} = DC_{(avg)}$ of $ V_{TX-D+} + V_{TX-D-} / 2 [\text{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	$ V_{TX-CM-DC-D+} - V_{TX-CM-DC-D-} \leq 25mV$ $V_{TX-CM-DC-D+} = DC_{(avg)}$ of $ V_{TX-D+} $ $V_{TX-CM-DC-D-} = DC_{(avg)}$ of $ V_{TX-D-} $ Refer to Note 1.
$V_{TX-IDLE-DIFFp}$	Electrical Idle Differential Peak Output Voltage	0		20	mV	$V_{TX-IDLE-DIFFp} = V_{TX-Idle-D+} - V_{TX-Idle-D-} \leq 20mV$ 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 15-7. 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. This is a de-bounce time for the Transmitter to meet Electrical Idle after transitioning from the L0 Link PM state.
$RL_{TX-DIFF}$	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	Ω	Tx DC Differential mode low impedance. Refer to Note 5.
$L_{TX-SKEW}$	Lane-to-Lane Output Skew			$500 + 2 \text{ UI}$	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 15-1](#).

Figure 15-1. 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 are 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 50Ω 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 15-1](#), 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 100Ω 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 100Ω

Small signal resistance is measured by forcing a small change in differential voltage and dividing this by the corresponding change in current.

Table 15-8. 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 \pm 300 ppm.
$V_{RX-DIFFp-p}$	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)} \text{ 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 Rx Impedance	80	100	120	Ω	Rx DC Differential mode impedance.
Z_{RX-DC}	DC Input Impedance	40	50	60	Ω	Required Rx D+ and D- DC impedance (50 Ω \pm 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 Receiver's package pins.
$T_{RX-IDLE-DET-DIFF-ENTERTIME}$	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 15-1](#) should be used as the Rx device when taking measurements.
7. The $T_{RX-EYE-MEDIAN-10-MAX-JITTER}$ specification ensures a jitter distribution in which the median and the maximum deviation from the median are 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-EYE} measurement is to be met at the target bit error rate. The $T_{RX-EYE-MEDIAN-10-MAX-JITTER}$ 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 Ω to ground for both the D+ and D- lines. The series capacitance C_{TX} is optional for the return loss measurement.

Table 15-9. 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		Ω	
	Input Parallel Termination (Differential)		110		Ω	
PPM	Reference Clock Tolerance	-300		+300	ppm	

Notes:

1. *PEX_REFCLKn/p must be AC-coupled. Use a 0.01 to 0.1 μ F capacitor.*
2. *Specified at 20 to 80% points at the package balls.*

15.7 Transmit Drive Characteristics

The drive current and transmit equalization function are programmable, to allow for optimization of different backplane lengths and materials.

15.7.1 Drive Current

The nominal drive current is programmable (2-bit, per SerDes/lane) 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 1](#) register (offset 24Ch) for details.]

15.7.2 Transmit Equalization

The Transmitter incorporates programmable (4-bit, per SerDes/lane) first-order equalization, within the range of 0 to -7.96 dB. [Refer to the [SerDes Drive Equalization Level Select 1](#) register (offset 254h) for details.]

15.7.3 Transmit Termination Adjust

The *PCI Express Base r1.1* specifies termination (50Ω nominal) at the Transmit side to VTT. The Transmit driver incorporates a 2-bit register (per SerDes quad), which allows for a $\pm 20\%$ 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[11:8]) for details.

15.8 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.

15.8.1 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 (ISI) due to frequency-dependent losses across the PCB material. Refer to [Physical Layer](#) register *SerDes Quad x RxEqCtl* fields (offset 22Ch[27:24]) for details.

15.8.2 Receive Termination Adjust

The *PCI Express Base r1.1* specifies termination (50Ω nominal) at the Receive side to ground. The Receiver input incorporates a 2-bit register (per SerDes quad), which allows for a $\pm 20\%$ 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[19:16]) for details.

THIS PAGE INTENTIONALLY LEFT BLANK.



Chapter 16 Thermal and Mechanical Specifications

16.1 Thermal Characteristics

The PEX 8509 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 8509, and is provided for reference only.

16.1.1 Sample Thermal Data

- Maximum Junction Temperature – 125°C

Table 16-1. Sample Thermal Data – 15 x 15 mm² PBGA Package

Heat Sink	Airflow (m/s)	Θ_{JA} (°C/W)	Θ_{JC} (°C/W)	Θ_{JB} (°C/W)
No	0.00	29.3	11.1	18.4
	1.02	26.7	–	–
	2.04	25.7	–	–
Yes	0.00	24.5	–	–
	1.02	18.0	–	–
	2.04	16.7	–	–

16.2 General Package Specifications

Table 16-2 defines the general package specifications. For a more complete list, refer to Figure 16-1.


Unpopulated BGA balls allow board design and placement of board-level de-coupling capacitors between VDD10, VDD10A, VDD10S, VDD33, and VSS/Ground.

Table 16-2. General Package Specifications

Parameter	Specification
Package Type	Plastic Ball Grid Array (PBGA)
Number of Balls	196
Package Dimensions	15 x 15 mm ² (approximately 1.81 ±0.19 mm high)
Ball Matrix Pattern	14 x 14
Ball Pitch	1.00 mm
Ball Diameter	0.50 ±0.10 mm
Ball Spacing	0.50 mm

16.3 Mechanical Dimensions

Figure 16-1. Mechanical Dimensions (196-Ball PBGA Package)



THIS PAGE INTENTIONALLY LEFT BLANK.



Appendix A General Information

A.1 Product Ordering Information

Contact your local [PLX Sales Representative](#) for ordering information.

Table A-1. Product Ordering Information

Part Numbers	Description
PEX8509-AA25BI	PEX 8509 8-Lane, 8-port PCI Express Gen 1 Switch Plastic BGA (15 x 15 mm ² , 196-ball) Leaded Package
PEX8509-AA25BI G	PEX 8509 8-Lane, 8-port PCI Express Gen 1 Switch Plastic BGA (15 x 15 mm ² , 196-ball) Lead-Free RoHS Green Package
<p>PEX8509-AA25BI G</p> <ul style="list-style-type: none"> G – Lead-Free, RoHS-Compliant, Fully Green I – Industrial Temperature B – Ball Grid Array Package AA – Silicon Revision 25 – Signaling Rate (2.5 Gbps) 8509 – Part Number PEX – PCI Express Product Family 	
PEX 8509-AA RDK	PEX 8509 Rapid Development Kit with x4 Edge Connector
x1 Adapter	PCI Express x4 to x1 Adapter

A.2 United States and International Representatives and Distributors

PLX Technology, Inc., representatives and distributors are listed at www.plxtech.com.

A.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.