



PCIe X8 IP Core

IP Version: v2.2.0

User Guide

FPGA-IPUG-02243-1.3

March 2025

Disclaimers

Lattice makes no warranty, representation, or guarantee regarding the accuracy of information contained in this document or the suitability of its products for any particular purpose. All information herein is provided AS IS, with all faults, and all associated risk is the responsibility entirely of the Buyer. The information provided herein is for informational purposes only and may contain technical inaccuracies or omissions, and may be otherwise rendered inaccurate for many reasons, and Lattice assumes no obligation to update or otherwise correct or revise this information. Products sold by Lattice have been subject to limited testing and it is the Buyer's responsibility to independently determine the suitability of any products and to test and verify the same. LATTICE PRODUCTS AND SERVICES ARE NOT DESIGNED, MANUFACTURED, OR TESTED FOR USE IN LIFE OR SAFETY CRITICAL SYSTEMS, HAZARDOUS ENVIRONMENTS, OR ANY OTHER ENVIRONMENTS REQUIRING FAIL-SAFE PERFORMANCE, INCLUDING ANY APPLICATION IN WHICH THE FAILURE OF THE PRODUCT OR SERVICE COULD LEAD TO DEATH, PERSONAL INJURY, SEVERE PROPERTY DAMAGE OR ENVIRONMENTAL HARM (COLLECTIVELY, "HIGH-RISK USES"). FURTHER, BUYER MUST TAKE PRUDENT STEPS TO PROTECT AGAINST PRODUCT AND SERVICE FAILURES, INCLUDING PROVIDING APPROPRIATE REDUNDANCIES, FAIL-SAFE FEATURES, AND/OR SHUT-DOWN MECHANISMS. LATTICE EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS OF THE PRODUCTS OR SERVICES FOR HIGH-RISK USES. The information provided in this document is proprietary to Lattice Semiconductor, and Lattice reserves the right to make any changes to the information in this document or to any products at any time without notice.

Inclusive Language

This document was created consistent with Lattice Semiconductor's inclusive language policy. In some cases, the language in underlying tools and other items may not yet have been updated. Please refer to Lattice's inclusive language [FAQ 6878](#) for a cross reference of terms. Note in some cases such as register names and state names it has been necessary to continue to utilize older terminology for compatibility.

Contents

Contents.....	3
Acronyms in This Document	14
1. Introduction	15
1.1. Overview of the IP	15
1.2. Quick Facts	15
1.3. IP Support Summary.....	15
1.4. Features.....	16
1.4.1. Hard IP Limitations	17
1.4.2. Soft IP	17
1.5. Licensing and Ordering Information.....	17
1.5.1. Ordering Part Number	18
1.6. Hardware Support	18
1.7. Speed Grade Supported	18
1.8. Naming Conventions	18
1.8.1. Nomenclature	18
1.8.2. Signal Names	18
2. Functional Description	19
2.1. PCIe IP Architecture Overview	19
2.2. Clocking	21
2.2.1. Clocking Overview	21
2.3. Reset.....	23
2.3.1. Reset Overview.....	23
2.3.2. Clock and Reset Sequence	24
2.4. Protocol Layers.....	24
2.4.1. ECC and Parity Data Path Protection	25
2.4.2. Error Handling	26
2.4.3. LTSSM State	28
2.5. PHY Equalization (8 GT/s).....	32
2.5.1. Equalization Process	32
2.5.2. Equalization Time Limit	32
2.5.3. Equalization Methods.....	32
2.5.4. Equalization Quality.....	39
2.6. Multi-Function Support.....	39
2.7. Power Management.....	39
2.7.1. Power Management Supported by PCIe IP Core	39
2.7.2. Configuring Core to Support Power Management	40
2.7.3. APSM L0s	40
2.7.4. APSM L1s	41
2.8. DMA Support.....	42
2.8.1. DMA Scatter-Gather and Status Queues	42
2.8.2. DMA Channel Programming and Operation.....	45
2.8.3. DMA Performance	47
2.8.4. DMA Bypass Interface	47
2.9. Non-DMA Support	48
2.9.1. Non-DMA Overview.....	48
2.9.2. Non-DMA Write.....	48
2.9.3. Non-DMA Read.....	49
2.10. Interrupts	49
2.10.1. Generation of the Interrupts	49
2.10.2. Legacy Interrupt	50
2.10.3. MSI Interrupt	50
2.10.4. MSI-X Interrupt.....	52

2.11.	PCIe Endpoint Core Buffers	53
2.11.1.	PCI Express Credits	54
2.11.2.	Max Payload Size	54
2.12.	Hard IP Interface	55
2.12.1.	PHY Interface	55
2.12.2.	TLP TX/RX Interface	55
2.12.3.	LMMI Interface	65
2.13.	Soft IP Interface	66
2.13.1.	Data Interface Conversion	66
2.13.2.	Register Interface Conversion	69
2.14.	Multi-Protocol Support	69
3.	IP Parameter Description	70
3.1.	General	70
3.2.	Optional Port	71
3.3.	Flow Control Update	72
3.4.	Function 0	73
3.4.1.	Configuration	73
3.4.2.	Base Address Register (BAR) [0 to 5]	74
3.4.3.	Legacy Interrupt	75
3.4.4.	MSI Capability	76
3.4.5.	MSI-X Capability	76
3.4.6.	Device Serial Number Capability	77
3.4.7.	PCI Express Capability	78
3.4.8.	Advance Error Reporting Capability	78
3.4.9.	Atomic OP Capability	79
3.4.10.	Latency Tolerance Reporting Capability	80
3.4.11.	Power Budgeting Capability	80
3.4.12.	Dynamic Power Allocation Capability	80
4.	Signal Description	83
4.1.	Clock and Reset Interface	83
4.2.	PHY Interface	84
4.3.	Transaction Layer Interface	85
4.3.1.	TLP Transmit Interface	85
4.3.2.	TLP Receive Interface	88
4.4.	Lattice Memory Mapped Interface (LMMI)	92
4.5.	Power Management Interface	93
4.6.	AXI-L Configuration Interface	93
4.7.	AXI-4 Stream Data Interface	94
4.7.1.	AXI4-Stream Transmitter Interface Port Descriptions	94
4.7.2.	AXI4-Stream Receiver Interface Port Descriptions	95
4.8.	DMA Interface	96
4.9.	Unused Interface	102
5.	Register Description	103
5.1.	Hard IP Core Configuration and Status Registers	103
5.1.1.	EP Configuration Settings	104
5.1.2.	mgmt_tlb (0x4_2000)	105
5.1.3.	mgmt_ptl (0x4_3000)	170
5.1.4.	mgmt_ftl (0x4_4000)	182
5.1.5.	mgmt_ftl_mf[7:1] (0x4_5000,0x4_6000,0x4_7000,0x4_8000,0x4_9000,0x4_A000,0x4_B000)	207
5.2.	PCI Express Configuration Space Registers	213
5.2.1.	Type 00 Configuration Header	213
5.2.2.	Type 01 Configuration Header	214
5.2.3.	Capability and Extended Capability Address Locations	214
5.2.4.	Type 00 Configuration Registers	215

5.2.5.	PCI Express Capability	216
5.2.6.	Power Management Capability	222
5.2.7.	MSI-X Capability	223
5.2.8.	MSI Capability	224
5.2.9.	Advanced Error Reporting Extended Capability	225
5.2.10.	ARI Extended Capability	227
5.2.11.	Vendor-Specific Extended Capability	228
5.2.12.	Secondary PCI Express Extended Capability	229
5.2.13.	ATS Extended Capability	230
5.2.14.	DSN Extended Capability	230
5.2.15.	Resizable BAR Capability	230
5.2.16.	Power Budgeting Capability	231
5.2.17.	Dynamic Power Allocation Capability	232
5.2.18.	L1 PM Substates Extended Capability	232
5.2.19.	Latency Tolerance Reporting Capability	233
5.3.	DMA Configuration Space Registers	233
6.	Example Design	240
6.1.	Example Design Supported Configuration	240
6.2.	Overview of the Example Design and Features	241
6.3.	Example Design Components	243
6.3.1.	DMA Design	243
6.3.2.	Non-DMA Design	243
6.4.	Simulating the Example Design	246
6.4.1.	Running Functional Simulation (Non-DMA)	247
6.4.2.	Running Functional Simulation (DMA)	251
6.5.	Generating the Example Design for Hardware	255
6.5.1.	DMA Variant	255
6.5.2.	NON-DMA Variant	257
6.6.	Design Test Case Examples	258
6.6.1.	DMA Design	258
6.6.2.	Non-DMA Design	259
6.7.	Debugging Example Design Issues	259
6.7.1.	Signals to Debug	259
6.8.	Limitations of the Example Design	260
7.	Designing with the IP	261
7.1.	Generating and Instantiating the IP	261
7.1.1.	Generated Files and File Structure	268
7.1.2.	Design Implementation	268
7.1.3.	Timing Constraints	268
8.	Debugging	271
8.1.	Debug Methods	271
8.1.1.	Debug Flow Charts	271
8.1.2.	Internal Register Read for Debug	277
8.1.3.	PCIe Loopback Test	277
9.	Design Consideration	278
9.1.	DMA Based Design	278
9.2.	Non-DMA Based Design	278
Appendix A.	Resource Utilization	279
References	280
Technical Support Assistance	281
Revision History	282

Figures

Figure 2.1. Lattice PCIe x8 IP Core Block Diagram	19
Figure 2.2. Lattice 8-lane SERDES/PCS + PCIe Hard-IP	20
Figure 2.3. Lattice PCIe x8 Core Hard IP	20
Figure 2.4. PCIe IP Clock Domain Block Diagram for AXI-S Interface	21
Figure 2.5. PCIe IP Clock Domain Block Diagram for TLP Interface	22
Figure 2.6. Reset Signals in Lattice PCIe IP Core	23
Figure 2-7. Clock and Reset Sequence Diagram	24
Figure 2.8. Source Scatter-Gather Queue Elements	43
Figure 2.9. Destination Scatter-Gather Queue Elements	44
Figure 2.10. Status Queue Elements	45
Figure 2.11. Non-DMA Application Data Flow	48
Figure 2.12. Non-DMA Write Operation (TLP Data Interface)	48
Figure 2.13. Non-DMA Read Operation (TLP Data Interface)	49
Figure 2.14. MSI Capability Structure Variant	51
Figure 2.15. MSI Capability Structure Variant	51
Figure 2.16. MSI-X Capability Structure Variant	52
Figure 2.17. MSI-X Table Entries	53
Figure 2.18. Pending Bit Array	53
Figure 2.19. TLP Memory Request Header	55
Figure 2.20. TLP Memory Read Operation for Link0 (x4 Lane)	58
Figure 2.21. TLP Memory Read Operation for Link0 (x2 Lane)	58
Figure 2.22. TLP Memory Read Operation for Link0 (x1 Lane)	58
Figure 2.23. Minimum link0_tx_ready_o Timing Diagram	59
Figure 2.24. Wait State of link0_tx_ready_o Timing Diagram	59
Figure 2.25. TLP Packet formation by the Lattice PCIe IP core	61
Figure 2.26. TLP Memory Write Operation for Link0 (x4 Lane)	62
Figure 2.27. TLP Memory Write Operation for Link0 (x2 Lane)	62
Figure 2.28. TLP Memory Write Operation for Link0 (x1 Lane)	62
Figure 2.29. Minimum link0_rx_ready_i Timing Diagram	63
Figure 2.30. Wait State of link0_rx_ready_i Timing Diagram	63
Figure 2.31. LMMI Write Operation	65
Figure 2.32. LMMI Read Operation	66
Figure 2.33. AXI4-Stream Data Interface, AXI-L Register Interface	66
Figure 2.34. PCIe to AXI4-Stream Transaction for x1	67
Figure 2.35. PCIe to AXI4-Stream Transaction for x2	67
Figure 2.36. PCIe to AXI4-Stream Transaction for x8	67
Figure 2.37. AXI4-Stream to PCIe Transaction for x1	68
Figure 2.38. AXI4-Stream to PCIe Transaction for x2	68
Figure 2.39. AXI4-Stream to PCIe Transaction for x8	68
Figure 2.40. PCIe Bifurcation	69
Figure 3.1. Attributes in the General Tab	70
Figure 3.2. Attributes in the Optional Port Tab	71
Figure 3.3. Attributes in the Flow Control Update Tab	72
Figure 3.4. Attributes in Function Configuration Tab	73
Figure 3.5. Attributes in BAR Tab	74
Figure 3.6. Attributes in Legacy Interrupt	75
Figure 3.7. Attributes in MSI Capability	76
Figure 3.8. Attributes in MSI-X Capability	76
Figure 3.9. Attributes in Device Serial Number Capability	77
Figure 3.10. Attributes in PCIe Capability	78
Figure 3.11. Attributes in Advance Error Reporting Capability	78
Figure 3.12. Attributes in Atomic OP Capability	79

Figure 3.13. Attributes in Latency Tolerance Reporting Capability	80
Figure 3.14. Attributes in Power Budgeting Capability	80
Figure 3.15. Attributes in Dynamic Allocation Capability	80
Figure 6.1. PCIe x8 IP Example Design Block Diagram	242
Figure 6.2. Components within the Harden DMA Design	243
Figure 6.3. Components within NON-DMA Design	244
Figure 6.4. Non-DMA Design Data Flow	245
Figure 6.5. PCIe x8 IP Example Design Flowchart	246
Figure 6.6. Testbench Files.....	247
Figure 6.7. Project Naming	247
Figure 6.8. Simulation Top Module.....	248
Figure 6.9. Simulation Setting	248
Figure 6.10. Known Simulation Error	249
Figure 6.11 System Environment Variable for Windows	249
Figure 6.12. Command of Full License QuestaSim OEM	250
Figure 6.13. Expected Log Printing	250
Figure 6.14. Simulation Waveform	251
Figure 6.15. Testbench Files	251
Figure 6.16. Project Naming	252
Figure 6.17. Simulation Top Naming.....	252
Figure 6.18. Simulation Setting	253
Figure 6.19. Known Simulation Error	253
Figure 6.20. System Environment Variable for Windows	254
Figure 6.21. Command of Full License QuestaSim OEM	254
Figure 6.22. Expected Log Printing	254
Figure 6.23. IP Configuration for Harden DMA Example Design – General Tab	255
Figure 6.24. IP Configuration for harden DMA example design – Link 0: Funtion 0 tab (Part 1).....	256
Figure 6.25. IP Configuration for Harden DMA Example Design – Link 0: Function 0 Tab (Part 2).....	256
Figure 6.26. IP Configuration for Non-DMA Example Design – General Tab	257
Figure 6.27. WIDTH Parameter Adjustment in the example_design_top Module According to Bifurcation Selection	258
Figure 7.1. Module/IP Block Wizard	261
Figure 7.2. IP Configuration for Non-DMA.....	262
Figure 7.3. IP Configuration for Harden DMA (General Tab)	262
Figure 7.4. IP Configuration for Harden DMA (Link 0: Function 0 Tab)	263
Figure 7.5. Check Generated Result.....	264
Figure 7.6. Set IP Top Level Unit	265
Figure 7.7. Select to Add Constraint File.....	266
Figure 7.8. Select constraint.pdc from the Eval Folder	267
Figure 7.9. Project Compile Done	267
Figure 7.10. Bitstream Not Generated due to the Virtual I/O Setting	268
Figure 7.11. Timing Constraint File (.sdc) for the PCIe x8 IP	269
Figure 7.12. PLL IP Configuration for Input Clock of 125 MHz	269
Figure 7.13. Timing Constraints for PCIe x8 IP Example	270
Figure 8.1. Hardware Detection Failure Debugging Flow	271
Figure 8.2. Link Training Issue Debugging Flow	272
Figure 8.3. Checkbox to enable LTSSM Polling	272
Figure 8.4. Data Transfer Issue Debugging Flow.....	275
Figure 8.5. Debugging the FPGA Configuration Issues Flow	276

Tables

Table 1.1. Summary of the PCIe x8 IP	15
Table 1.2. Ordering Part Number	18
Table 1.3. Lattice PCIe IP Core Supported Speed Grade	18
Table 2.1. General PCI Express Error List	26
Table 2.2. Physical Layer Error List	26
Table 2.3. Data Link Layer Error List	27
Table 2.4. Transaction Layer Error List	27
Table 2.5. LTSSM State Definition	28
Table 2.6. RX L0s State Description	31
Table 2.7. Preset to Coefficient Conversion	33
Table 2.8. Simulation Data Throughput Using a Different Descriptor Size for FPGA-to-Host (F2H) Transfer	47
Table 2.9. Simulation Data Throughput Using a Different Descriptor Size for Host-to-FPGA (H2F) Transfer	47
Table 2.10. Register Access for Different Data Interfaces	48
Table 2.11. Base Address to Enable Interrupt	49
Table 2.12. Legacy Interrupt Register	50
Table 2.13. TLP Header Field	56
Table 2.14. Data Byte Order	64
Table 3.1. General Tab Attributes Description	70
Table 3.2. Optional Port Attributes	71
Table 3.3. Flow Control Attributes	72
Table 3.4. Function Configuration Tab Attributes	73
Table 3.5. BAR Tab Attributes	74
Table 3.6. Legacy Interrupt Attribute Descriptions	75
Table 3.7. MSI Capability Attributes	76
Table 3.8. MSI-X Capability Attributes	77
Table 3.9. Device Serial Number Capability Attributes	77
Table 3.10. PCIe Capability Attributes	78
Table 3.11. Advance Error Reporting Capability Attributes	79
Table 3.12. Atomic OP capability Attributes	79
Table 3.13. Latency Tolerance Reporting Capability Attributes	80
Table 3.14. Power Budgeting Capability Attributes	80
Table 3.15. Dynamic Allocation capability Attributes	81
Table 3.16. Function 1-3 Tab	82
Table 4.1. Clock and Reset Ports	83
Table 4.2. PHY Interface Descriptions	84
Table 4.3. TLP Transmit Interface Ports	85
Table 4.4. TLP Transmit Credit Interface Ports	87
Table 4.5. TLP Receive Interface Ports	88
Table 4.6. TLP Receive Credit Interface Ports	91
Table 4.7. Lattice Memory Mapped Interface Ports	92
Table 4.8. AXI-L Configuration Interface Ports	93
Table 4.9. AXI-4 Stream Transmitter Interface Ports	94
Table 4.10. AXI-4 Stream Receiver Interface Ports	95
Table 4.11. DMA AXI4 Manager Write Interface	96
Table 4.12. DMA AXI4 Manager Read Interface	98
Table 4.13. DMA AXI4 Subordinate Write Interface	100
Table 4.14. DMA AXI4 Subordinate Read Interface	101
Table 4.15. DMA AXI4 Interrupt Interface	101
Table 4.16. Unused Interface	102
Table 5.1. Register Access Abbreviations	103
Table 5.2. Hard PCIe Core Register Mapping	103
Table 5.3. CSR Values Recommended for EP Applications	104

Table 5.4. Itssm_simulation Register 0x0	105
Table 5.5. Itssm_cfg_lw_start Register 0x34	105
Table 5.6. Itssm_latch_rx Register 0x38	106
Table 5.7. Itssm_cfg Register 0x3c	106
Table 5.8. Itssm_port_type Register 0x40	108
Table 5.9. Itssm_ds_link Register 0x44	108
Table 5.10. Itssm_detect_quiet Register 0x48	109
Table 5.11. Itssm_rx_det Register 0x4c	109
Table 5.12. Itssm_nfts Register 0x50	109
Table 5.13. Itssm_ds_initial_auto Register 0x54	110
Table 5.14. Itssm_select_deemphasis Register 0x58	110
Table 5.15. Itssm_beacon Register 0x5c	111
Table 5.16. Itssm_mod_cpl Register 0x60	111
Table 5.17. Itssm_rx_elec_idle Register 0x64	111
Table 5.18. Itssm_compliance_toggle Register 0x68	112
Table 5.19. Itssm_prevent_rx_ts_entry_to Register 0x6c	113
Table 5.20. Itssm_link Register 0x80	113
Table 5.21. Itssm_Itssm Register 0x84	114
Table 5.22. Itssm_rx_I0s Register 0x88	116
Table 5.23. I0_to_rec Register 0x8c	117
Table 5.24. Itssm_rx_detect Register 0x90	118
Table 5.25. Itssm_configured Register 0x94	119
Table 5.26. Itssm_direct_to_detect Register 0x98	119
Table 5.27. Itssm_equalization Register 0x9c	119
Table 5.28. Itssm_crosslink Register 0xa0	120
Table 5.29. Physical Layer Tx Underflow Error Status Register – 0xa4	120
Table 5.30. Physical Lane Rx Status Registers	120
Table 5.31. pl_rx0 Register 0xa8 – Lane Rx Status 0 Register	120
Table 5.32. pl_rx1 Register 0xac – Lane Rx Status 1	122
Table 5.33. pl_rx2 Register 0xb0 – Lane Rx Status 2	125
Table 5.34. pl_rx3 Register 0xb4 – Lane Rx Status 3	128
Table 5.35. pl_rx4 Register 0xb8 – Lane Rx Status 4	130
Table 5.36. debugself_crosslink Register 0xc0	133
Table 5.37. debug_rx_det Register 0xc4	133
Table 5.38. debug_force_tx Register 0xc8	134
Table 5.39. debug_direct_scramble_off Register 0xcc	134
Table 5.40. debug_force_scramble_off_fast Register 0xd0	134
Table 5.41. balign Register 0xd4	135
Table 5.42. debug_pipe_rx Register 0xe0	136
Table 5.43. debug_direct_to_loopback Register 0x100	136
Table 5.44. debug_loopback_control Register 0x104	136
Table 5.45. debug_loopback_master_5g Register 0x108	137
Table 5.46. debug_loopback_slave_5g Register 0x10c	138
Table 5.47. debug_loopback_master_8g_deemph Register 0x110	138
Table 5.48. debug_loopback_slave_8g_deemph Register 0x114	139
Table 5.49. debug_direct_to_loopback_status Register 0x118	139
Table 5.50. debug_loopback_err_reset Register 0x11c	140
Table 5.51. debug_loopback_err Register 0x120	140
Table 5.52. phy_control Register 0x140	140
Table 5.53. phy_control_8g Register 0x144	140
Table 5.54. phy_eq_tx_override Register 0x148 F	141
Table 5.55. phy_eq_tx_max Register 0x14c	141
Table 5.56. phy_eq_tx_force Register 0x150	142
Table 5.57. phy_preset_to_coef_conv_control Register 0x15c	142

Table 5.58. phy_preset_conv_tab_pre Register 0x160	143
Table 5.59. phy_preset_conv_tab_post Register 0x170	143
Table 5.60. eq_control Register 0x180	143
Table 5.61. eq_ts_control Register 0x184	144
Table 5.62. eq_reduced_swing Register 0x188	146
Table 5.63. eq_method Register 0x1bc	146
Table 5.64. eq_fmerit_control Register 0x1c0	146
Table 5.65. eq_preset_method_control Register 0x1c4	147
Table 5.66. eq_alg_method_control Register 0x1c8	147
Table 5.67. eq_table_method_control Register 0x1cc	148
Table 5.68. eq_table_method_table Register 0x1d0	148
Table 5.69. eq_updn_control Register 0x240	149
Table 5.70. eq_firmware_control Register 0x280	150
Table 5.71. eq_pre_cursor Register 0x290	151
Table 5.72. eq_post_cursor Register 0x2a0	151
Table 5.73. eq_status Register 0x2c0	152
Table 5.74. eq_status_error Register 0x2c4	153
Table 5.75. eq_status_preset_coef Register 0x2c8	153
Table 5.76. eq_status_feedback_fom Register 0x2d0	153
Table 5.77. eq_status_feedback_dir Register 0x2e0	154
Table 5.78. eq_status_remote_fs Register 0x2e8	154
Table 5.79. eq_status_remote_if Register 0x2f4	154
Table 5.80. eq_status_remote_precursor Register 0x300	154
Table 5.81. eq_status_remote_postcursor Register 0x30c	155
Table 5.82. pl_rx Register 0x33c	155
Table 5.83. pl_16g Register 0x340	155
Table 5.84. pl_tx_skp Register 0x344	156
Table 5.85. pl_tx_debug Register 0x348	157
Table 5.86. pl_ctrl Register 0x34c	157
Table 5.87. pl_ts_matching Register 0x350	157
Table 5.88. dl_retry_timeout Register 0x380	158
Table 5.89. dl_ack_timeout_div Register 0x384	158
Table 5.90. dl_tx_ctrl Register 0x38c	159
Table 5.91. dl_ctrl Register 0x390	160
Table 5.92. dl_stat Register 0x394	163
Table 5.93. dl_ack_to_nak Register 0x398	166
Table 5.94. dl_inject Register 0x39c	166
Table 5.95. dllp_inject Register 0x3a0	167
Table 5.96. eq_status_table_control Register 0x3d8	168
Table 5.97. eq_status_table_info Register 0x3dc	168
Table 5.98. eq_status_table Register 0x3e0	168
Table 5.99. eq_capture_sel Register 0x3f0	169
Table 5.100. eq_capture Register 0x3f4	169
Table 5.101. phy_eq_tx_force_per_lane Register 0x400	169
Table 5.102. phy_eq_tx_force_per_lane_8g_pre Register 0x404	170
Table 5.103. phy_eq_tx_force_per_lane_8g_post Register 0x410	170
Table 5.104. Simulation Register 0x0	170
Table 5.105. pm_aspm_ios Register 0x40	170
Table 5.106. pm_aspm_l1 Register 0x50	171
Table 5.107. pm_aspm_l1_min Register 0x54	171
Table 5.108. pm_l1 Register 0x60	171
Table 5.109. pm_l1_min Register 0x64	171
Table 5.110. pm_l1pmss Register 0x68	172
Table 5.111. pm_l2 Register 0x70	172

Table 5.112. pm_pme_to_ack_ep Register 0x80.....	172
Table 5.113. pm_pme_to_ack_ds Register 0x84	172
Table 5.114. pm_pme Register 0x88	173
Table 5.115. pm_status Register 0x90.....	173
Table 5.116. vc_rx_c Register 0x108.....	174
Table 5.117. vc_rx_adv Register 0x10c.....	174
Table 5.118. vc_rx_control Register 0x110.....	174
Table 5.119. vc_rx_status Register 0x114.....	175
Table 5.120. vc_rx_credit_status_cfg Register 0x120	176
Table 5.121. vc_rx_credit_status_p Register 0x124	176
Table 5.122. vc_rx_credit_status_n Register 0x128	176
Table 5.123. vc_rx_credit_status_c Register 0x12c.....	176
Table 5.124. vc_rx_f_oc_update_timer Register 0x130	177
Table 5.125. vc_rx_p_flow_ctrl Register 0x134.....	177
Table 5.126. vc_rx_n_flow_ctrl Register 0x138.....	177
Table 5.127. vc_rx_alloc_size Register 0x140.....	178
Table 5.128. vc_rx_alloc_p Register 0x144.....	178
Table 5.129. vc_rx_alloc_n Register 0x148.....	178
Table 5.130. vc_rx_alloc_c Register 0x14c	179
Table 5.131. vc_rx_alloc_error Register 0x150.....	179
Table 5.132. vc_tx_np_fifo Register 0x180.....	180
Table 5.133. vc_tx_status Register 0x184.....	180
Table 5.134. vc_tx_credit_status_p Register 0x190	180
Table 5.135. vc_tx_credit_status_n Register 0x194	181
Table 5.136. vc_tx_credit_status_c Register 0x198.....	181
Table 5.137. vc_tx_credit_cleanup Register 0x19c.....	181
Table 5.138. tlp_tx Register 0x1c4.....	182
Table 5.139. fc_credit_init Register 0x1c8.....	182
Table 5.140. simulation Register 0x0.....	182
Table 5.141. decode Register 0x10	183
Table 5.142. decode_t1 Register 0x14.....	185
Table 5.143. tlp_processing Register 0x18	186
Table 5.144. Initial Register 0x20.....	186
Table 5.145. cfg Register 0x30	187
Table 5.146. ds_port Register 0x34	187
Table 5.147. us_port Register 0x38	188
Table 5.148. id1 Register 0x40.....	188
Table 5.149. id2 Register 0x44.....	188
Table 5.150. id3 Register 0x48.....	188
Table 5.151. Cardbus Register 0x4c	189
Table 5.152. Interrupt Register 0x50	189
Table 5.153. bar0 Register 0x60	189
Table 5.154. bar1 Register 0x64	189
Table 5.155. bar2 Register 0x68	190
Table 5.156. bar3 Register 0x6c.....	190
Table 5.157. bar4 Register 0x70	190
Table 5.158. bar5 Register 0x74	190
Table 5.159. exp_rom Register 0x78	191
Table 5.160. pcie_cap Register 0x80	191
Table 5.161. pcie_dev_cap Register 0x84.....	192
Table 5.162. pcie_link_cap Register 0x88.....	193
Table 5.163. pcie_link_stat Register 0x8c.....	194
Table 5.164. pcie_slot_cap Register 0x90.....	194
Table 5.165. pcie_dev_cap2 Register 0x98.....	196

Table 5.166. pcie_link_ctl2 Register 0xa0.....	197
Table 5.167. pm_cap Register 0xc0	198
Table 5.168. pm Register 0xc4	199
Table 5.169. pm_aux Register 0xc8	199
Table 5.170. ari_cap Register 0xe0	200
Table 5.171. aer_cap Register 0x100	200
Table 5.172. msi_cap Register 0xe8.....	201
Table 5.173. msix_cap Register 0xf0.....	202
Table 5.174. msix_table Register 0xf4	202
Table 5.175. msix_pba Register 0xf8	202
Table 5.176. vsec_cap Register 0x110	203
Table 5.177. sris_cap Register 0x120.....	203
Table 5.178. dsn_cap Register 0x130	203
Table 5.179. dsn_serial Register 0x134	204
Table 5.180. pwr_budget_cap Register 0x150	204
Table 5.181. dpa_cap Register 0x158	204
Table 5.182. dpa_xlcy Register 0x15c	205
Table 5.183. dpa_alloc Register 0x160	205
Table 5.184. ltr_cap Register 0x180.....	206
Table 5.185. l1pmss_cap Register 0x188	206
Table 5.186. atomic_op_cap Register 0x1cc.....	207
Table 5.187. Base Address for mgmt_ftl_mf	207
Table 5.188. Function Register 0x8.....	208
Table 5.189. us_port Register 0x38	208
Table 5.190. id1 Register 0x40	208
Table 5.191. id2 Register 0x44	208
Table 5.192. id3 Register 0x48	209
Table 5.193. Cardbus Register 0x4c	209
Table 5.194. Interrupt Register 0x50	209
Table 5.195. bar0 Register 0x60	209
Table 5.196. bar1 Register 0x64	210
Table 5.197. bar2 Register 0x68	210
Table 5.198. bar3 Register 0x6c.....	210
Table 5.199. bar4 Register 0x70	210
Table 5.200. bar5 Register 0x74	210
Table 5.201. exp_rom Register 0x78	211
Table 5.202. msi_cap Register 0xe8.....	211
Table 5.203. msix_cap Register 0xf0.....	211
Table 5.204. msix_table Register 0xf4	212
Table 5.205. msix_pba Register 0xf8	212
Table 5.206. dsn_cap Register 0x130	212
Table 5.207. dsn_serial Register 0x134	213
Table 5.208. Type 00 Configuration Header	213
Table 5.209. Type 01 Configuration Header	214
Table 5.210. Capability and Extended Capability Items	214
Table 5.211. Type 00 Configuration Registers	215
Table 5.212. PCI Express Capability	216
Table 5.213. Power Management Capability.....	222
Table 5.214. MSI-X Capability	223
Table 5.215. MSI Capability	224
Table 5.216. Advanced Error Reporting Extended Capability	225
Table 5.217. ARI Extended Capability	227
Table 5.218. Vendor-Specific Extended Capability	228
Table 5.219. Secondary PCI Express Extended Capability.....	229

Table 5.220. ATS Extended Capability	230
Table 5.221. DSN Extended Capability.....	230
Table 5.222. Resizable BAR Capability	230
Table 5.223. Power Budgeting Capability	231
Table 5.224. Dynamic Power Allocation (DPA) Capability	232
Table 5.225. L1 PM Substates Extended Capability	232
Table 5.226. Latency Tolerance Reporting (LTR) Capability.....	233
Table 5.227. DMA Configuration Space Registers	234
Table 6.1. PCIe x8 IP Configuration Supported by the Example Design.....	240
Table 6.2. Non-DMA Design Status Ports Description	259
Table 6.3. Non-DMA Signals to Debug Description	259
Table 7.1. Generated File List	268
Table 8.1. PCIe LTSSM State and Sub-State Definition	273
Table A.1. Lattice PCIe IP Core Resource Utilization	279

Acronyms in This Document

A list of acronyms used in this document.

Acronym	Definition
ASPM	Active State Power Management
AXI	Advanced Extensible Interface
AXI-L	Advanced Extensible Interface Lite
BAR	Base Address Register
CSR	Configuration and Status Register
DLLP	Data Link Layer Packet
DMA	Direct Memory Access
ECC	Error Correction Coding
EP	Endpoint
FIFO	First In First Out
LMMI	Lattice Memory Mapped Interface
LTSSM	Link Training and Status State Machine
MSI	Message Signaled Interrupt
RTL	Register Transfer Language
PCI	Peripheral Component Interconnect
PCIE	Peripheral Component Interconnect Express
PCS	Physical Coding Sublayer
PLL	Phase-Locked Loop
PM	Power Management
PMA	Physical Medium Attachment
RAM	Random Access Memory
RC	Root Complex
RP	Root Port
TLP	Transaction Layer Packet

1. Introduction

1.1. Overview of the IP

PCI Express® is a high performance, fully scalable, and well-defined standard for a wide variety of computing and communications platforms. As a packet-based serial technology, the PCI Express standard greatly reduces the number of required pins and simplifies board routing and manufacturing. PCI Express is a point-to-point technology, as opposed to the multi-drop bus in PCI. Each PCI Express device has the advantage of full duplex communication with its link partner to greatly increase overall system bandwidth. The basic data rate for a single lane is double that of the 32-bit/33 MHz PCI bus. A four-lane link has eight times the data rate in each direction of a conventional bus.

The Lattice PCIe x8 IP Core provides a flexible, high-performance, easy-to-use Transaction Layer Interface to the PCI Express Bus. The Lattice PCIe x8 IP Core implementation is a hardened IP with soft logic provided for interface conversion options. The hardened IP is an integration of PHY and Link Layer blocks.

The Lattice PCIe x8 IP Core is supported in the Lattice Avant™-AT-G and Lattice Avant-AT-X FPGA device family and is available in the Lattice Radiant™ software.

1.2. Quick Facts

Table 1.1. Summary of the PCIe x8 IP

IP Requirements	Supported FPGA Families	Avant-AT-G, Avant-AT-X
	IP Changes	Refer to the PCIe x8 IP Release Notes (FPGA-RN-02061) .
Resource Utilization	Targeted Devices	All Avant-AT-G and Avant-AT-X family
	Supported User Interface	LMMI, AXI-L, TLP, AXI4-Stream
Design Tool Support	Lattice Implementation	IP Core v2.2.0 – Lattice Radiant Software 2024.2.1 or later
	Synthesis	Lattice Synthesis Engine
		Synopsys® Synplify Pro® for Lattice
	Simulation	Questasim

1.3. IP Support Summary

Device Family	IP	User Interface	Gen Speed	Link Width	Data Rate (Gbps)	Radiant Timing Model	Hardware Validated
Avant-AT-G Avant-AT-X	PCIe EP	TLP	Gen 4	X8, X4, X2, X1	128, 64, 32, 16	Preliminary	No
			Gen 3	X8, X4, X2, X1	64, 32, 16, 8	Preliminary	No
			Gen 2	X8, X4, X2, X1	40, 20, 10, 5	Preliminary	No
			Gen 1	X8, X4, X2, X1	20, 10, 5, 2.5	Preliminary	No
		AXI-Stream	Gen 4	X8, X4, X2, X1	128, 64, 32, 16	Preliminary	No
			Gen 3	X8, X4, X2, X1	64, 32, 16, 8	Preliminary	No
			Gen 2	X8, X4, X2, X1	40, 20, 10, 5	Preliminary	No
			Gen 1	X8, X4, X2, X1	20, 10, 5, 2.5	Preliminary	No
	PCIe Harden DMA	AXI-MM	Gen 4	X8	128	Preliminary	No
			Gen 3	X8	64	Preliminary	No
			Gen 2	X8	40	Preliminary	No
			Gen 1	X8	20	Preliminary	No

1.4. Features

The Hard IP PHY key features include:

- Aggregation and bifurcation up to x4 lanes and x8 lanes PHY configuration
- Data rates of 2.5 Gbps, 5.0 Gbps, 8.0 Gbps, and 16.0 Gbps
- Selectable parallel data widths such as 8, 16, 32, 64
- 8b/10b encoding at 2.5 Gbps and 5 Gbps, and 128b/130b encoding at 8 Gbps and 16 Gbps
- Adaptive and configurable RX Continuous Time Linear Equalizer (CTLE) and Decision Feedback Equalizer (DFE)
- Adaptive and programmable TX equalization
- PCIe L1-substate power managements and Separate RefClk Independent SSC Architecture (SRIS)
- Extensive PMA debug capability via read/write and read-only registers in PCS
- Register-based control of all PCS-to-PMA signals
- A wide range of reference clock frequencies with optional fractional frequency correction capability
- A wide range of divided clock frequencies for external-to-PHY usage with optional spread-spectrum clock (SSC) capability
- Built-in, on-chip SSC generation and full configuration from –5000 to +5000 ppm
- Test support features such as near-end loopback, PLL bypass modes, and others.
- Protocol-compatible features such as LOS, squelch, power modes, and others.

The Hard IP Link Layer key features include:

- PCI Express Base Specification Revision 4.0 compliant including compliance with earlier PCI Express Specifications.
- Backward compatible with PCI Express 3.x, 2.x, 1.x
- x8 PCI Express Lanes with support for bifurcation
 - Supported lane configurations:
 - Radiant 2023.2 – 1 × 4, 1 × 2, 1 × 1
 - Radiant 2024.1 and future Radiant releases – 1 × 8, 1 × 4, 1 × 2, 1 × 1
- 16.0GT/s, 8.0GT/s, 5.0 GT/s, and 2.5 GT/s line rate support
- Comprehensive application support:
- Endpoint
- Root Port (future release)
- Multi-Function support per link
 - PCIe x8 Core supports 1-8 PF
 - PCIe x4 Core supports 1-8 PF
- SRIOV support per link (future release for non-DMA)
 - PCIe x8 Core supports up to 32 (PF+VF) functions
 - PCIe x4 Core supports up to 32 (PF+VF) functions
- Support for Autonomous and Software-Controlled Equalization
- Support for Figure of Merit and Up/Down PIPE PHY Equalization
- Flexible Equalization methods (Algorithm, Preset, User-Table, Adaptive-Table, Firmware-controlled)
- ECC RAM and Parity Data Path Protection
- Core Data Width
 - 64 bits for x1 lane
 - 128 bits for x2 lanes
 - 256 bits for x4 lanes
 - 512 bits for x8 lanes
- Complete error-handling support
- AER, ECRC generation/checking, recovery from Parity and ECC errors
- Supports detection of numerous optional errors, embedded simulation error checks/assertions
- Simulation and hardware error injection features enable error testing
- Flexible core options allow for design complexity/feature tradeoffs
- Configurable Receive, Transmit, and Replay Buffer sizes
- Supports Polarity Inversion, Up/Down-configure, Autonomous Link Width/Speed changes
- Power Management

- Supports L1, ASPM L0s, and ASPM L1
- L1 PM Substates with CLKREQ
- Power Budgeting
- Dynamic Power Allocation
- Latency Tolerance Reporting
- Implements Type 0 Configuration Registers in Endpoint Mode
- Implements Type 1 Configuration Registers in Root Port Mode (future release)
- Complete Root Port Configuration Register implementation (future release)
- Dual mode design supports EP or RP through register changes (future release)
- Easy to use
- Decodes received packets to provide key routing (BAR hits, Tag, and others) information
- Implements all aspects of the required PCIe Configuration Space
- Optionally consumes PCI Express Message TLPs or leaves them in band
- Interfaces have consistent timing and function over all modes of operation
- Provides a wealth of diagnostic information for superior system-level debug and link monitoring
- Implements all three PCI Express Layers (Transaction, Data Link, and Physical)
- Hardened high-performance multi-channel scatter-gather DMA controller
 - Support up to 64 DMA channels
 - Support for PCIe multi-function and SR-IOV capability
 - 64-bit address support

1.4.1. Hard IP Limitations

The following are the limitations of the Hard IP:

- Users can only access Function 0 of the PCIE Configuration Space Registers due to address bit limitation in Hard IP. This affects the usage of links other than Link 0 since it relies on register interface to get a particular function setting (for example, Max payload size and MSI/MSI-X configuration) including error reporting (AER).
- CLKREQ is not supported in MPPHY, reference clock stopping during L1 substate is not supported.
- AXI Write and Read address channel are sharing the same ports. Request will be done one at a time or sequentially – cannot be simultaneous.
- DPA and LTR interface of the Hard IP are not exposed/available to the user.
- The Link Layer port *p_clk_period_in_ps is not programmable (hard coded). This may have an effect in timeout calculations.
- Reset ports are currently not exposed/available to the user.
- DMA is supported only in x8 mode.

1.4.2. Soft IP

- TLP Data Interface
- AXI4-Stream Data Interface Option
- LMMI Register Interfaces
- AXI-L Register Interface

Radiant 2024.1 SP1 requires a patch to work with PCIe_x8 IP.

1.5. Licensing and Ordering Information

An IP specific license string is required to enable full use of the PCIe x8 IP in a complete, top-level design.

The IP can be fully evaluated through functional simulation and implementation (synthesis, map, place and route) without an IP license string. This IP supports Lattice's [Hardware Support](#) capabilities. You can create versions of the IP to operate in hardware for a limited time (approximately four hours) without requiring an IP license string. A license string is required to enable timing simulation and to generate a bitstream file that does not include the hardware evaluation timeout limitation.

For more information about pricing and availability of the PCIe x8 IP, contact your [local Lattice Sales Office](#).

1.5.1. Ordering Part Number

Table 1.2. Ordering Part Number

Device Family	Part Number	
	Multi-Site Perpetual	Single Seat Annual
Avant-AT-G	PCI-EXP8-AVG-UT	PCI-EXP8-AVG-US
	PCI-EXP4-AVG-UT	PCI-EXP4-AVG-US
	PCI-EXP2-AVG-UT	PCI-EXP2-AVG-US
Avant-AT-X	PCI-EXP8-AVX-UT	PCI-EXP8-AVX-US
	PCI-EXP4-AVX-UT	PCI-EXP4-AVG-US
	PCI-EXP2-AVX-UT	PCI-EXP2-AVG-US

1.6. Hardware Support

Refer to the [Example Design](#) section for more information on the boards used.

1.7. Speed Grade Supported

The Lattice PCIe IP core supported speed grade is provided in this section. Different configurations may be supported using different speed grade due to fabric performance requirement.

- 3 – fastest speed grade

Table 1.3. Lattice PCIe IP Core Supported Speed Grade

PCIe Core Config	Device Family	Speed Grade
Gen3x8	Avant-AT-G	1/2/3
Gen4x8	Avant-AT-X	1/2/3

1.8. Naming Conventions

1.8.1. Nomenclature

The nomenclature used in this document is based on Verilog HDL.

1.8.2. Signal Names

- `_n` are active low (asserted when value is logic 0)
- `_i` are input signals
- `_o` are output signals

2. Functional Description

2.1. PCIe IP Architecture Overview

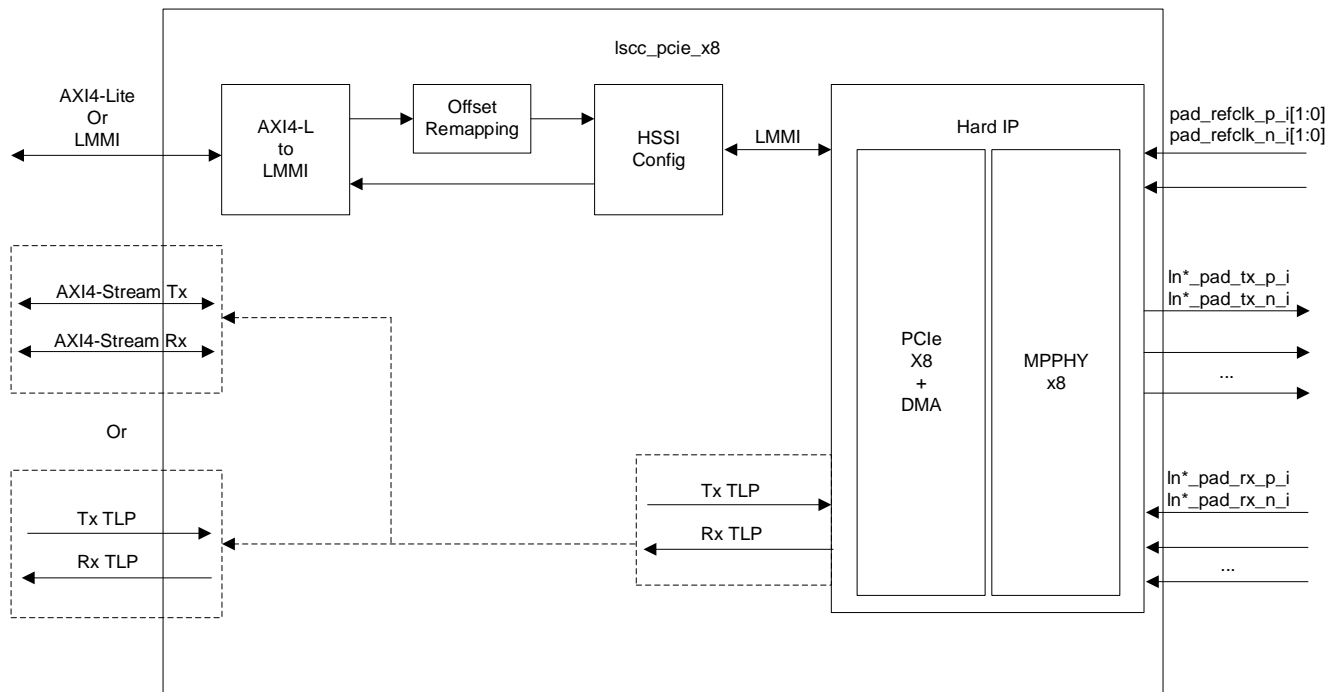


Figure 2.1. Lattice PCIe x8 IP Core Block Diagram

The Lattice PCIe x8 Core implements all three layers defined by the PCI Express specification: Physical, Data Link, and Transaction. Bifurcation is supported with link and lane configuration such as 1×8 , 1×4 , 1×2 , and 1×1 .

The Lattice PCIe x8 Core Hard IP implementation integrates the Rambus Expresso Core (Link Layer) and DMA controller, and Synopsys PMA and PCS Core (PHY). A block diagram with an 8-lane instance is shown in [Figure 2.2](#).

For convenience, the Link Layer Cores are referred to as Link 0 – PCIe x8 Core.

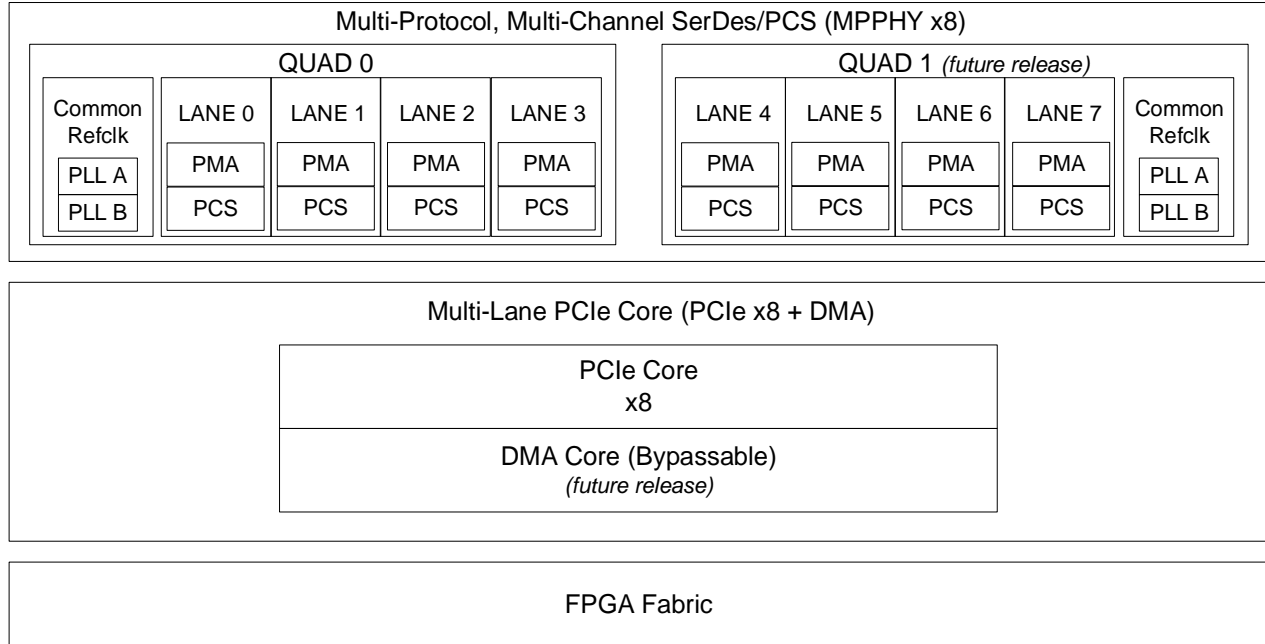


Figure 2.2. Lattice 8-lane SERDES/PCS + PCIe Hard-IP

The Lattice PCIe x8 Core Hard IP has the following interfaces as shown in Figure 2.3. The details of each interface are discussed in the following sections.

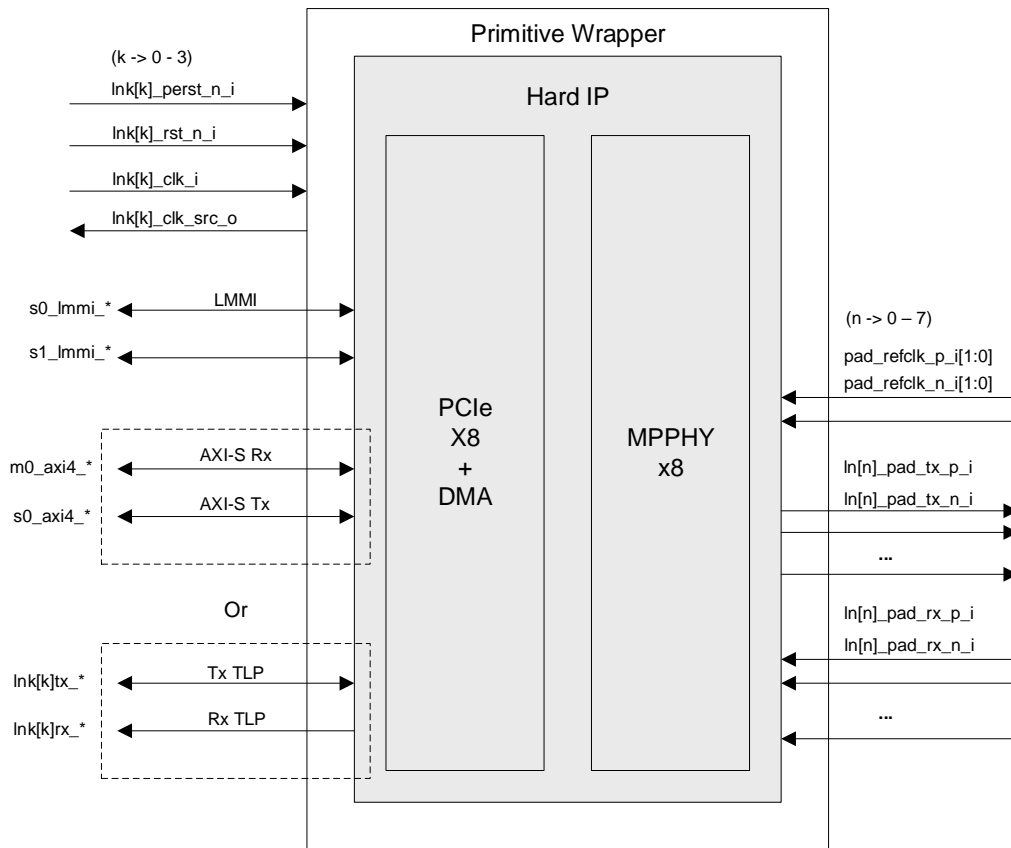


Figure 2.3. Lattice PCIe x8 Core Hard IP

- Clock and Reset Interface
 - The user domain interface can be clocked using one of several PHY clock outputs (sys_clk_i = link0_clk_src_o) or by user-generated clock. The fundamental reset (link0_perst_n_i) resets the core (PHY and Link layer blocks) except for the core Configuration Registers. Another reset (link0_rst_usr_n_i) is provided to reset only the Link layer block.
- PHY Interface
 - High Speed Serial Interface supports maximum rate of 16GT/s
 - Up to eight lanes (grouped into Two Quads)
- TLP Receive Interface
 - Receive TLPs from the PCIe link partner
 - High bandwidth interface
 - 512b data width (64b per active lane)
- TLP Transmit Interface
 - Transmit TLPs to the PCIe link partner
- High bandwidth interface
 - 512b data width (64b per active lane)
 - Power Management Interface – Not Supported
 - Ports for implementing power management capabilities
- AXI Interface
 - Available if Hard DMA Core is enabled
 - 512b data width each (TX and RX)
- AXI4-Lite or LMMI (Mutually exclusive) – Configuration and Status Register (CSR) Interface
 - This interface is used to write and read the core configuration and status registers. A typical customer application requires changing only a small number of the default values such as Vendor ID, Device ID, and BAR configuration.
 - 32b data width (AXI4-Lite) or 16b data width (LMMI)

2.2. Clocking

2.2.1. Clocking Overview

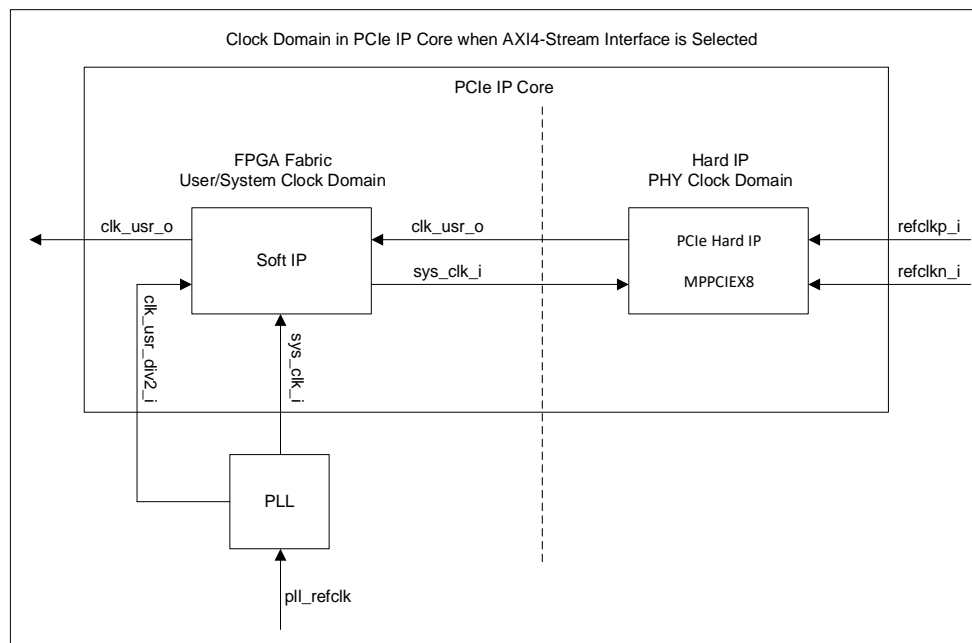


Figure 2.4. PCIe IP Clock Domain Block Diagram for AXI-S Interface




Figure 2.5. PCIe IP Clock Domain Block Diagram for TLP Interface

The PCIe x8 IP includes the following clock domains. The *sys_clk_i* and *clk_usr_div2_i* signals are generated during the PLL IP instantiation. For the TLP interface, you can connect *clk_usr_o* back to the *sys_clk_i*, as shown in [Figure 2.5](#).

- *refclkp_i/refclkn_i* are differential PHY reference clocks.
 - You can set the reference clock frequency in the PCIe IP Core.
 - There are two options available: 100 MHz or 125 MHz.
- *sys_clk_i/clk_usr_i* is the user clock domain input clock.
 - This clock is generated by the system PLL and shared to the Link layer blocks.
 - For the TLP interface variants, you can choose to connect *clk_usr_o* back to the *sys_clk_i* as shown in [Figure 2.5](#).
- *clk_usr_div2_i* is the user clock domain divided by two input clocks.
 - This clock is generated by the system PLL with simple division by two at half of the *sys_clk_i* frequency.
 - This clock does not apply to the TLP interface variants.
 -
- *clk_usr_o* is the User Clock Domain Output Clock.
 - This is the pclk output that comes from the PHY of the PCIe IP core.
 - By default, *clk_usr_o* uses the divide-by-2 version of the 125 MHz pclk from the PHY.

2.3. Reset

2.3.1. Reset Overview

There are two fundamental reset events that can occur in PCI Express:

- Cold Reset – This is a fundamental reset applied during power cycling. The signal link [LINK]_perst_n_i is asserted.
- Warm Reset – This is a fundamental reset triggered by hardware without the removal and re-application of power. The perst_n_i signal is asserted.

The fundamental reset *link [LINK]_perst_n_i* resets the core (Link Layer and PHY Layer blocks) while another reset, which is the user clock domain Link Layer reset *[LINK]_rst_usr_n_i*, is used to reset the Link Layer block only.

Depending on the PCIe IP configuration, either DMA or non-DMA, reset signal *usr_cfg_reset_n_i* is used.

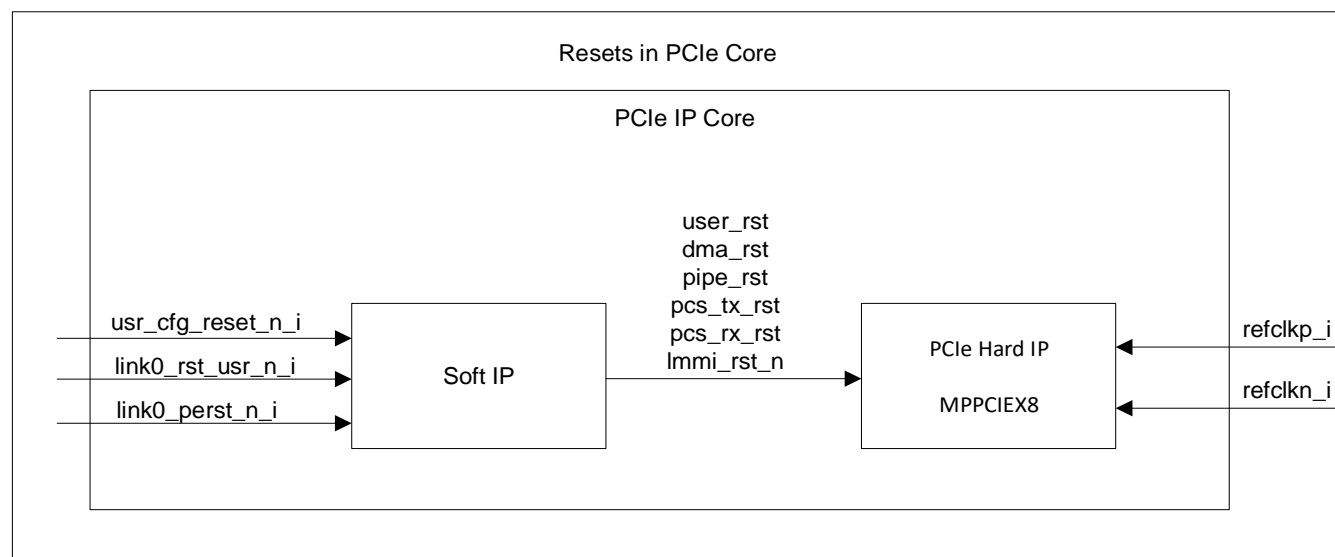


Figure 2.6. Reset Signals in Lattice PCIe IP Core

2.3.2. Clock and Reset Sequence

The PCIe IP clock and reset operation is shown in Figure 2-7.




Figure 2-7. Clock and Reset Sequence Diagram

The Lattice PCIe x8 IP Core configuration register implementation has default values that are appropriate for most applications. You can change the register configuration through the LMMI or AXI-L interface. When the LMMI or AXI-L interface is used to configure the PHY layer registers, the configuration should be done before the deassertion of the *link0_perst_n_i* signal. The PHY Layer is released from reset and is ready for operation once it is able to generate the PIPE clock output (such as the *link0_clk_usr_o* signal). The user domain reset (such as *link0_rst_usr_n_i*) can be deasserted if the Link Layer register configuration is done or skipped.

To ensure that the clock is stable before the link training, you must wait for the PLL locked status for all four channels of Tx PLL before de-asserting the user domain reset (*link0_rst_usr_n_i*). If you select the x2 or x1 link width, you may observe the two-channel Tx PLL locked status or one channel Tx PLL locked status respectively. The TX PLL status (*bit-4, offset 0x7F in PHY PMA Status register*) can be read through the LMMI or AXI-L.

2.4. Protocol Layers

There are three major classes of packets in PCIe devices: Transaction Layer Packets (TLP), Data Link Layer Packets (DLLP), and Physical Layer Packets (PLP), which is also known as ordered sets). The function of the Protocol Layer is to generate and process these packets.

- **Transaction Layer**
The Transaction Layer manages the TLPs to communicate request and completion data with other PCIe devices. The TLP packets are assembled at the *transmit side* of the link and disassembled at the *receive side* of the link. The TLP communicates through different formats either in I/O request format or in the memory request format.
- **Data Link Layer**
The Data Link Layer transfers data from the Transaction Layer to the Physical Layer. It plays an important role in assuring good reception of the TLP packets. The DLLPs are used to convey the information about the link initialization, power management, flow control, and TLP acknowledgements.

- Physical Layer

The Physical Layer converts the packets from the Data Link Layer into serialized bit streams and transfers it to the external physical link. The receive logic de-serializes the bits, reassembles the packets, and forwards it to the Data Link Layer. It conveys the communication between the Data Link Layer and the external physical link. The Physical layer is divided into the Logical sub-block and the Electrical sub-block. The Logical sub-block frames and deframes the packets and implements the LTSSM state machine. The scrambling, descrambling, and 8B/10B encoding and decoding of data are done in the logical sub-block. The Electrical sub-block provides the physical interface to the Link and contains the differential transmitters and receivers. The PLPs or ordered sets are exchanged during link training and link initialization.

2.4.1. ECC and Parity Data Path Protection

The Lattice PCIe x8 IP Core protects the TLP data path with Error Correction Coding (ECC) and Parity Protection. This is implemented in the Hard IP block.

ECC is used to protect TLP data in the following data path RAMs:

- Replay Buffer
- Receive Buffer
- Transmit Buffer

The ECC implementation enables correction for 1-bit errors and detection for 2-bit errors. The 8-bit of ECC information is included in the RAMs for each 64 (or fraction thereof) data bits.

Even (XOR) Parity ($parity[i] = \bigwedge (data[(i+1) \times 8 - 1 : i \times 8])$) is used to protect the data path. Parity provides detection for 1-bit errors (and other odd-bit errors). To enable continuous parity protection coverage, parity is passed through RAMs that are also protected by ECC.

The core includes the ability to enable/disable the reporting and handling of ECC/Parity errors. Correctable errors (ECC 1-bit errors) are fixed when correction is enabled. Uncorrectable ECC/Parity errors in the transmit data path result in the associated TLP being discarded or nullified when error handling is enabled. While error handling can be disabled, this is not recommended as passing a known TLP with bad contents can result in a more serious error condition than discarding the TLP.

2.4.1.1. Receive Data Path

For the receive data path, parity is generated for received TLPs prior to the removal and validation of the Link CRC (LCRC). Parity protection is thus overlapped with LCRC protection.

Received TLP parity is passed with the associated received TLP (header and payload) bytes through the Receive Buffer and onto the user Transaction Layer Receive interface. It is expected that parity is checked and errors are handled by the ultimate TLP consumer. Since TLP can have parity errors on any byte (toward the end of a longer TLP for instance), it is generally not possible to avoid processing the error TLP as the earlier portion of the TLP may already have been processed by the time that the error is detected.

Applications that do not want to process TLPs with errors need to store and forward the TLP for processing only after inspecting the parity of all data bytes. If the core Transaction Layer detects a parity error while it is consuming a received TLP (Type 0 Configuration Read/Write, Malformed TLP, and Message), the error is reported as Uncorrectable Error (in AER capability) and the core discards the TLP without processing it.

2.4.1.2. Transmit Data Path

For the transmit data path, parity is generated by the TLP source. For user TLPs (for example those transmitted on the core's Transaction Layer Transmit Interface), the parity is provided along with associated TLP (header and payload) bytes. The provided parity is kept with the associated data as it traverses the core. The parity is checked and discarded just after the TLP PCIe LCRC is generated.

Parity protection is thus overlapped with LCRC protection, including the associated PCIe replay mechanism. If the core detects a parity or uncorrectable ECC error during transmission of a TLP, the error is reported and the associated TLP is nullified (discarded) and not retransmitted. This is a serious error that must be handled by the software. The TLP is discarded to not propagate the error and risk potentially worse consequences in other components that receives TLPs with known bit errors.

2.4.1.3. Uncorrectable Error Recovery

PCI Express includes the ability to nullify or cancel a TLP transmission immediately after it is completed by inverting the LCRC and using End Bad (EDB) end framing instead of the normal TLP end framing. TLP can be nullified to reduce propagation, potentially multiplying the effects of the error. Nullified TLPs are not regenerated by the original TLP source as it is difficult for software to construct the missing TLP. As a result, there is a fatal system error condition regardless of whether the error TLP is nullified or not. When TLP is nullified due to errors, the core attempts to keep the transmit stream active so that the software can be notified of the error using the standard in-band mechanisms (for example, transmission of ERR_NFAT or ERR_FAT message).

TLPs are allocated a sequence number during transmission and the PCIe receiver only accepts TLPs in sequential order. When a TLP is nullified due to an uncorrectable error, the missing sequence number must be recovered before the link can continue to transmit TLPs.

TLPs are allocated Virtual Channel Flow Control Credits when they are transmitted by the Transaction Layer. The PCI Express device receiving the TLP over the PCI Express link frees the associated credits by sending Flow Control Update DLLPs. TLPs, which are nullified due to uncorrectable ECC and Parity errors, are allocated credits by the Transaction Layer, which is never freed since the TLP is nullified and not received by the Receiver. Nullified TLPs are discarded by the Receiver without affecting Flow Control Credits or Sequence Number.

Whenever a transmitted TLP is nullified due to an uncorrectable error, this causes the PCI Express link to be unable to process further TLPs. The sequence number and flow control credits that are allocated to the nullified TLP must be reclaimed before the link is repaired. The Lattice PCIe x8 IP Core contains logic to correct the link when TLPs are nullified due to uncorrectable errors.

Whenever an uncorrectable ECC or Parity error is detected, it is recommended for you to reset the link through the software to reset the link although the link is corrected for further transmission.

2.4.2. Error Handling

The Lattice PCIe x8 IP Core detects and implements the appropriate response to most error conditions without user intervention. You generally only need to detect and report errors that the core does not have enough information to detect.

2.4.2.1. PCIe-Defined Error Types

The following defines the error types in the PCIe. The *Type* column in Table 2.1 to Table 2.4 defines the PCI Express defined error severity:

- COR – Correctable
- NFAT – Uncorrectable – Non-Fatal
- FAT – Uncorrectable – Fatal

Table 2.1. General PCI Express Error List

Error	Type
Corrected Internal Error	COR
Uncorrectable Internal Error	FAT
Header Log Overflow	COR

Table 2.2. Physical Layer Error List

Error	Type
Receiver Error	COR

Table 2.3. Data Link Layer Error List

Error	Type
Bad TLP	COR
Bad DLLP	COR
Replay Timeout	COR
REPLAY_NUM Rollover	COR
Data Link Layer Protocol Error	FAT
Surprise Down	FAT

Table 2.4. Transaction Layer Error List

Error	Type
Poisoned TLP Received	NFAT
ECRC Check Failed	NFAT
Unsupported Request	NFAT
Completion Timeout	NFAT
Completer Abort	NFAT
Unexpected Completion	NFAT
ACS Violation	NFAT
MC Blocked TLP	NFAT
AtomicOp Egress Blocked	NFAT
Receiver Overflow	FAT
Flow Control Protocol Error	FAT
Malformed TLP	FAT

2.4.2.2. User Error Reporting

The User Hardware design must be able to detect and report the following errors.

- Uncorrectable Internal Error
 - Signals if AER Version 0x2 is enabled in the core and user hardware is detected and unable to correct an application-specific error that is not reported through another error mechanism.
 - If AER is supported by the core, the header of the first TLP associated with the error may optionally be logged.
- Poisoned TLP Received with Advisory Non-Fatal Severity
 - Signals if the core's default poison handling is disabled (`ignore_poison == 1`) and you receive a poisoned TLP that is considered as Advisory Non-Fatal severity. If the data payload of a poisoned packet is used or the poison can be recovered from the software or other mechanism, the poison should be treated as Advisory Non-Fatal since a non-fatal error often causes a system operation to crash.
 - If AER is supported by the core and the core is operating in Endpoint mode, an `ERR_COR` message is requested and transmitted if enabled.
 - If AER is supported by the core, the header of the poisoned packet must be logged.
- Poisoned TLP with Non-Fatal Severity
 - Signals if the core's default poison handling is disabled (`ignore_poison == 1`) and you receive a poisoned TLP that is considered as Non-Fatal severity. Handling poison as Non-Fatal severity should be avoided when possible as this is often fatal to the system operation.
 - If AER is supported by the core, the header of the poisoned packet must be logged.

- **Unsupported Request**
 - A Type0 Vendor-defined message that is received but not supported by user logic is an Unsupported Request. This is uncommon since only devices designed to receive Type0 Vendor-defined messages should receive these. However, compliance tests may require this error to be handled; hence, it is recommended to implement this check. Receiving a message with Message Code == 0x7E should cause Unsupported Request to be reported, unless the user design is designed to receive these messages.
 - Completions that are received with a Reserved Completion status must be handled as if the Completion status is an Unsupported Request.
- **Completion Timeout**
 - If you initiate a non-posted request (all reads, I/O Write, and Configuration Write), you are required to implement a completion timeout timer that fires if completions to a non-posted request are not received in the allotted time period. This error check needs to be implemented by the user design that includes initiating non-posted requests.
- **Completion Abort**
 - Signals if permanently unable to process a request due to a device-specific error condition. Generally, this error is only signaled if you choose to implement a restricted programming model (that requires the software to always perform DWORD size transactions and not support burst transactions). This is not recommended unless that the only software that can access the user design is your own software, which is designed to conform with the restricted programming model.
 - If AER is supported by the core, the header of the aborted request must be logged.
- **Unexpected Completion**
 - You must signal if a completion is received but the tag does not match any outstanding requests.
 - If the core is enabled for Target_Only mode indicating that the user design does not initiate non-posted requests, the core considers all completions as Unexpected Completions, discards them, and generates the appropriate response. In this case, you do not handle this error.
 - If AER is supported by the core, then the header of the completion must be logged.

As a minimum, it is recommended to report the following errors:

- Completion Timeout if user logic initiates non-posted requests (for example, DMA read requests)
- Unsupported Requests for the cases described above
- Unexpected Completion for the case described above
- Poison, when the core's default poison handling is disabled (ignore_poison == 1)

2.4.3. LTSSM State

2.4.3.1. Main LTSSM

The Lattice PCIe x8 IP Core follows the PCI Express specification for the Link Training and Status State Machine. However, to help hit higher frequencies, the LTSSM is split into one Major State LTSSM state machine and several separate LTSSM sub-state machines, with one sub-state state machine for each major state.

The Lattice PCIe x8 IP Core implements additional LTSSM sub-states that are necessary to meet PCIe specification LTSSM operation but are not given an explicit sub-state in the PCIe specification. [Table 2.5](#) lists each state.

Table 2.5. LTSSM State Definition

LTSSM Major State	LTSSM Sub-state	Description
0 – Detect	0 – DETECT_INACTIVE	The sub-state is <i>DETECT_INACTIVE</i> whenever the LTSSM major state is not <i>Detect</i> .
	1 – DETECT_QUIET	Detect.Quiet
	2 – DETECT_SPD_CHG0	Detect.Quiet – Sub-state to change speed change back to 2.5G if needed. Request PHY speed change.
	3 – DETECT_SPD_CHG1	Detect.Quiet – Sub-state to change speed change back to 2.5G if needed. Wait for speed change to complete.
	4 – DETECT_ACTIVE0	Detect.Active – First Rx Detection.
	5 – DETECT_ACTIVE1	Detect.Active – Wait 12 ms between Rx Detection attempts.
	6 – DETECT_ACTIVE2	Detect.Active – Second Rx Detection (if needed).

LTSSM Major State	LTSSM Sub-state	Description
	7 – DETECT_P1_TO_P0	Detect.Active – Change PHY power state from P1 to P0 (inactive to active) if needed (that is on Detect – Polling transition).
	8 – DETECT_P0_TO_P1_0	Change PHY power state from P0 to P1 (active to inactive) – Transmit Electrical Idle Ordered Sets to notify the link partner that the link is idle.
	9 – DETECT_P0_TO_P1_1	Change PHY power state from P0 to P1. Wait for TX Electrical Idle Ordered Set transit request made in DETECT_P0_TO_P1_0 to get transmitted at the output of the core.
	10 – DETECT_P0_TO_P1_2	Change PHY power state from P0 to P1. Wait for PHY to reach P1 state before continuing.
1 – Polling	0 – POLLING_INACTIVE	The sub-state is <i>POLLING_INACTIVE</i> whenever the LTSSM Major State is not Polling.
	1 – POLLING_ACTIVE_ENTRY	Polling.Active – Entry to <i>Polling.Active</i> State exists since in some cases, the LTSSM must exit Polling without Tx of TS OS.
	2 – POLLING_ACTIVE	Polling.Active
	3 – POLLING_CFG	Polling.Configuration
	4 – POLLING_COMP	Polling.Compliance – Transmitting compliance pattern.
	5 – POLLING_COMP_ENTRY	Polling.Compliance entry state – Directs a speed change through POLLING_COMP_EIOS, POLLING_COMP_EIOS_ACK, and POLLING_COMP_IDLE when necessary, before going to POLLING_COMP.
	6 – POLLING_COMP_EIOS	Polling.Compliance – Transmit Electrical Idle Ordered Sets to notify the link partner that the link is idle.
	7 – POLLING_COMP_EIOS_ACK	Polling.Compliance – Wait for the Electrical Idle Ordered Sets transmitted in POLLING_COMP_EIOS to exit the core.
	8 – POLLING_COMP_IDLE	Polling.Compliance – Perform speed change now that link is idle.
2 – Configuration	0 – CONFIGURATION_INACTIVE	The sub-state is <i>CONFIGURATION_INACTIVE</i> whenever the LTSSM Major State is not Configuration.
	1 – CONFIGURATION_US_LW_START	Acting as Upstream Port – Configuration.Linkwidth.Start
	2 – CONFIGURATION_US_LW_ACCEPT	Acting as Upstream Port – Configuration.Linkwidth.Accept
	3 – CONFIGURATION_US_LN_WAIT	Acting as Upstream Port – Configuration.Lanenum.Wait
	4 – CONFIGURATION_US_LN_ACCEPT	Acting as Upstream Port – Configuration.Lanenum.Accept
	5 – CONFIGURATION_DS_LW_START	Acting as Downstream Port – Configuration.Linkwidth.Start
	6 – CONFIGURATION_DS_LW_ACCEPT	Acting as Downstream Port – Configuration.Linkwidth.Accept
	7 – CONFIGURATION_DS_LN_WAIT	Acting as Downstream Port – Configuration.Lanenum.Wait
	8 – CONFIGURATION_DS_LN_ACCEPT	Acting as Downstream Port – Configuration.Lanenum.Accept
	9 – CONFIGURATION_COMPLETE	Configuration.Complete
	10 – CONFIGURATION_IDLE	Configuration.Idle

LTSSM Major State	LTSSM Sub-state	Description
3 – L0	0 – L0_INACTIVE	The sub-state is <i>L0_INACTIVE</i> whenever the LTSSM Major State is not L0.
	1 – L0_L0	L0 – Link is in L0.
	2 – L0_TX_EL_IDLE	Tx_L0s.Entry, L1.Entry, or L2.Entry – Transmit Electrical Idle Ordered Sets to notify the link partner that the link is idle (that is for preparing to enter low power states such as Tx_L0s, L1, and L2).
	3 – L0_TX_IDLE_MIN	Tx_L0s.Entry, L1.Entry, or L2.Entry – Guarantee the minimum Tx Elec Idle time when entering electrical idle and also require Rx EIOS to have been received when necessary.
4 – Recovery	0 – RECOVERY_INACTIVE	The sub-state is <i>RECOVERY_INACTIVE</i> whenever the LTSSM Major state is not <i>Recovery</i> .
	1 – RECOVERY_RCVR_LOCK	Recovery.RcvrLock
	2 – RECOVERY_RCVR_CFG	Recovery.RcvrCfg
	3 – RECOVERY_IDLE	Recovery.Idle
	4 – RECOVERY_SPEED0	Recovery.Speed – Transmit Electrical Idle Ordered Sets to notify the link partner that the link is idle.
	5 – RECOVERY_SPEED1	Recovery.Speed – Determine to which speed to change.
	6 – RECOVERY_SPEED2	Recovery.Speed – Wait for remote device to enter electrical idle and wait for the required minimum time.
	7 – RECOVERY_SPEED3	Recovery.Speed – Request PHY change speed and wait for PHY to finish changing speed.
	8 – RECOVERY_EQ_PH0	Recovery.Equalization – Phase 0
	9 – RECOVERY_EQ_PH1	Recovery.Equalization – Phase 1
	10 – RECOVERY_EQ_PH2	Recovery.Equalization – Phase 2
	11 – RECOVERY_EQ_PH3	Recovery.Equalization – Phase 3
5 – Disable	0 – DISABLE_INACTIVE	The sub-state is <i>DISABLE_INACTIVE</i> whenever the LTSSM Major state is not <i>Disable</i> .
	1 – DISABLE0	Disable – Transmit 16 to 32 TS1 Ordered Sets with Disable Link bit asserted.
	2 – DISABLE1	Disable – Transition to Electrical Idle.
	3 – DISABLE2	Disable – Wait to receive an Electrical Idle Ordered Set and min time of TX_IDLE_MIN afterwards.
	4 – DISABLE3	Disable – Wait until a Disable exit condition occurs.
6 – Loopback	0 – LOOPBACK_INACTIVE	The sub-state is <i>LOOPBACK_INACTIVE</i> whenever the LTSSM Major state is not <i>Loopback</i> .
	1 – LOOPBACK_ENTRY	Loopback.Entry – Loopback entry state – Loopback Leader may be required to Tx Loopback TS OS before continuing or speed may need to be changed before beginning loopback.
	2 – LOOPBACK_ENTRY_EXIT	Loopback.Entry – Prepare to enter Loopback.Active
	3 – LOOPBACK_EIOS	Loopback.Entry – Transmit Electrical Idle Ordered Sets to notify the link partner that the link is idle. (to change speed).
	4 – LOOPBACK_EIOS_ACK	Loopback.Entry – Wait for the Electrical Idle Ordered Sets transmitted in LOOPBACK_EIOS to exit the core.
	5 – LOOPBACK_IDLE	Loopback.Entry – Stay in Electrical Idle for required minimum time.
	6 – LOOPBACK_ACTIVE	Loopback.Active
	7 – LOOPBACK_EXIT0	Loopback.Exit – Tx Electrical Idle
	8 – LOOPBACK_EXIT1	Loopback.Exit – Stay in Electrical Idle for required minimum time.

LTSSM Major State	LTSSM Sub-state	Description
7 – Hot Reset	0 – HOT_RESET_INACTIVE	The sub-state is <i>HOT_RESET_INACTIVE</i> whenever the LTSSM Major state is not <i>Hot Reset</i> .
	1 – HOT_RESET_HOT_RESET	Hot Reset – as Follower
	2 – HOT_RESET_LEADER_UP	Hot Reset – as Leader with Link Up
	3 – HOT_RESET_LEADER_DOWN	Hot Reset – as Leader with Link Down
8 – TX L0s	0 – TX_L0S_INACTIVE	The sub-state is <i>TX_L0S_INACTIVE</i> whenever the LTSSM Major state is not <i>TX L0s</i> .
	1 – TX_L0S_IDLE	Tx_L0s.Idle – Idle
	2 – TX_L0S_TO_L0	Tx_L0s.Idle – Exiting TX L0s; wait for PHY to indicate exit from L0s complete
	3 – TX_L0S_FTS0	Tx_L0s.FTS – Transmit requested NFTS.
	4 – TX_L0S_FTS1	Tx_L0s.FTS – Transmit additional FTS required by Cfg Register Extended Sync.
9 – L1	0 – L1_INACTIVE	The sub-state is <i>L1_INACTIVE</i> whenever the LTSSM Major state is not <i>L1</i> .
	1 – L1_IDLE	L1.Idle
	2 – L1_SUBSTATE	L1.1 or L1.2 depending upon higher level Power Management State Machine control.
	3 – L1_TO_L0	L1.Idle – Exiting L1; wait for PHY to indicate exit from L1 complete.
10 – L2	0 – L2_INACTIVE	The sub-state is <i>L2_INACTIVE</i> whenever the LTSSM Major State is not <i>L2</i> .
	1 – L2_IDLE	L2.Idle – Idle
	2 – L2_TX_WAKE0	L2.TransmitWake – Transmit a Beacon until remote device exits electrical idle.
	3 – L2_TX_WAKE1	L2.TransmitWake – Assert Tx Electrical Idle before changing power state to P1.
	4 – L2_EXIT	L2.Idle – L2 exit; wait until PHY finishes power change out of L2.
	5 – L2_SPEED	L2.Idle – Change speed if required before going to L2.

2.4.3.2. RX L0s State Machine

The Rx_L0s State Machine follows the L0s state of the receiver. The Rx_L0s State Machine operates independently of the main LTSSM, which controls the state of the transmitter.

Table 2.6. RX L0s State Description

LTSSM Sub-state	Description
0 – RX_L0S_L0	The sub-state is “RX_L0S_L0” whenever the receiver is in L0 (that is not en route to or in Rx L0s).
1 – RX_L0S_ENTRY	Rx_L0s.Entry
2 – RX_L0S_IDLE	Rx_L0s.Idle
3 – RX_L0S_FTS	Rx_L0s.FTS
4 – RX_L0S_REC	Rx_L0s.FTS – Wait until LTSSM Major State == Recovery due to Rx L0s exit error

2.5. PHY Equalization (8 GT/s)

Operating at 8 GT/s data rate requires an equalization process to be completed before data can be reliably transferred. The Lattice PCIe x8 IP Core supports both the autonomous equalization and software-controlled equalization methods.

When equalization is initiated either through the autonomous or software mechanism, the core Link Training and Status State Machine (LTSSM) enters recovery to perform equalization. Equalization is done for both directions of the link. The Equalization process consists of requesting several sets of transmit coefficients for the remote device to use, evaluating the quality of each set of coefficients, and then choosing the best coefficient set for operation.

The Lattice PCIe x8 IP Core supports PHY using the Figure of Merit and Up/Down Equalization feedback methods established by the PIPE Specification, as well as an option for direct firmware control/status. The Lattice PCIe x8 IP Core implements several equalization algorithms that enable users to select the method that works best for the project.

2.5.1. Equalization Process

Equalization is done for both directions of the data flow:

- Remote Transmitter to Local Receiver
 - Equalization is controlled by the local device using the mechanisms described in this section.
- Local Transmitter to Remote Receiver
 - Equalization is controlled by the remote device with the assistance of the local core LTSSM. While the Local Transmitter to Remote Receiver link is undergoing equalization, the *pipe_tx_deemph* port is periodically updated by the core based upon received requests from the remote device. Other than the PHY updating its transmitter to the new *pipe_tx_deemph* settings, no external action is required.

2.5.2. Equalization Time Limit

The equalization state machine follows the PCI Express timeout of 24 ms. You must ensure that the selected method is able to complete equalization within 24 ms, or the process is aborted and considered unsuccessful.

100 μ s (0.1 ms) should be reserved for each equalization attempt for the LTSSM to communicate the new settings to the remote device and to receive acknowledgement of those settings.

If a PHY takes 2 ms to evaluate each setting and 0.1 ms is reserved per setting for LTSSM, only 11 settings may be evaluated during equalization ($2.1\text{ ms} \times 11 = 23.1\text{ ms}$). Since a final evaluation (with the best of the trial settings) is typically required to finalize the process, a maximum of 10 different trial settings may be evaluated in this case.

The PHY vendor must specify the required time to complete an equalization evaluation. If the PHY vendor does not specify a limit, the worst case of 2 ms must be assumed and a maximum of 10 different settings may be attempted. The smaller the amount of time the PHY takes to evaluate a setting, the more settings can be attempted.

2.5.3. Equalization Methods

The Lattice PCIe x8 IP Core implements a flexible Equalization process that enables users to exercise control over the choice of equalization coefficients requested of the remote device.

Each method requires expressing Pre-Cursor and Post-Cursor Coefficients. Coefficients are expressed as a positive integer ratio $c: c[5:0]/64$. For example, $c[5:0] == 8$ yields a ratio 8/64 or 0.125.

The remote PCI Express device advertises its PHY's Full Scale (FS) and Low Frequency (LF) values. The FS and LF values and the desired coefficients are used to compute the required coefficient format that is expressed to the remote device. The Lattice PCIe x8 IP Core does not violate the following PCI Express Coefficient rules. If violated, it limits the coefficient to its maximum allowed value:

- $|C-1| \leq \text{Floor}(\text{FS}/4)$
- $C0 - |C-1| - |C+1| \geq \text{LF}$
- $|C-1| + C0 + |C+1| = \text{FS}$

The Core applies the three rules above in the order illustrated. The Pre-Cursor and Post-Cursor coefficient absolute values are used in computations. The Pre-Cursor (C-1) is limited to Floor (FS/4). Then the Post-Cursor (C+1) is limited to $((\text{FS}-\text{LF})/2) - C-1$ (solving the latter two equations for C+1). Then the Cursor (C0) is computed to be $\text{FS} - C-1 - C+1$. The Lattice PCIe x8 IP Core supports both Figure of Merit and Up/Down PIPE PHY Equalization Methods.

The available Equalization algorithms are described in the following sections.

2.5.3.1. Figure of Merit – Preset Method

TX Preset (as defined by PCIe Specification) is sweep from Presets from Preset[0] to Preset [preset_method_control_addr_limit] and Figure of Merit for each preset is determined. The PCIe Specifications defines presets 0x0 through 0xA. Although, 0xA preset is intended as a diagnostic preset and generally should not be included.

```
for (i = 0; i ≤ eq_preset_method_control_addr_limit; i = i + 1) {
    pre_preset_coef = pre-cursor coef for Preset[i] // See Preset to Coefficient Conversion
    post_preset_coef = post-cursor coef for Preset[i] // See Preset to Coefficient Conversion
    pre = (pre_preset_coef * RemoteFS) / 64
    post = (post_preset_coef * RemoteFS) / 64
    // Core requests link partner PHY Tx use {post, pre} coefficients and then core performs a
    Rx EQ Evaluation
    try {post, pre}
}
```

After all (*eq_preset_method_control_addr_limit*+1) trials are completed, an additional Rx EQ evaluation is executed using the post and pre coefficients from the trial which achieved the highest Figure of Merit. This final evaluation is skipped if the current coefficients being used are the best coefficients. If the final Figure of Merit meets or exceeds the threshold configured in *eq_fmerit_control_req_feedback*, the Rx Equalization is considered successful; otherwise, it is unsuccessful.

The preset method is optionally configured to communicate the desired preset for the remote device to select appropriate coefficients and matching presets or by calculating the coefficients equivalent to the preset number and communicating the coefficients to the remote device.

When the core is configured to communicate with the preset through the coefficients rather than preset number, the core uses the coefficients from the pre-cursor coefficient and post-cursor coefficient columns in [Table 2.7](#) to perform the calculation. The coefficient target, expressed as a real number in parenthesis, is provided along with the rounded to 1/64 coefficient value that is used. The same conversion is also used in the other Figure of Merit methods when users need to convert from a preset number into the equivalent coefficients.

Table 2.7. Preset to Coefficient Conversion

Preset Encoding	De-Emphasis dB	Pre-Shoot dB	Pre-Cursor Coefficient	Post-Cursor Coefficient
0	-6	0	0 (0)	16 (-0.25)
1	-3.5	0	0 (0)	11 (-0.167)
2	-4.5	0	0 (0)	13 (-0.2)
3	-2.5	0	0 (0)	8 (-0.125)
4	0	0	0 (0)	0 (0)
5	0	2	6 (-0.1)	0 (0)
6	0	2.5	8 (-0.125)	0 (0)
7	-6	3.5	6 (-0.1)	13 (-0.2)
8	-3.5	3.5	8 (-0.125)	8 (-0.125)
9	0	3.5	11 (-0.167)	0 (0)
A	Special case: Pre-Cursor == 0, Post-Cursor == (FS-LF)/2			

The preset method works well with PHY, which take the maximum of 2 ms to evaluate equalization settings since only the 10 presets from Preset[0] through Preset[9] are evaluated. The available channels in the PCIe Specification is expected to correspond to one of the preset method that works well with PHY, which takes the maximum of 2 ms to evaluate equalization settings since only the 10 Presets from Preset[0] through Preset[9] are evaluated. The preset method is defined by the PCIe Specification to take the maximum of 2 ms to evaluate equalization settings since only the 10 Presets from Preset[0] through Preset[9] would be evaluated. Also, all available channels in the PCIe Specification are expected to correspond to one of the PCIe preset. Therefore, this method is recommended if users are unsure which method to use to evaluate Preset[0] to Preset[9].

2.5.3.2. Figure of Merit – Algorithm Method

This method sweeps through the possible coefficient values. The Algorithm Method enables complete coefficient range coverage at the expense of longer run time. Coefficient coverage ranges from coarse to fine depending upon how many iterations can be tried before the 24 ms Equalization time out.

```
pre_cursor_limit = eq_alg_method_control_pre_cursor_limit
post_cursor_limit = eq_alg_method_control_post_cursor_limit
pre_cursor_step_size = eq_alg_method_control_pre_cursor_step_size
post_cursor_step_size = eq_alg_method_control_post_cursor_step_size
for (pre_coef = 0; pre_coef ≤ pre_cursor_limit; pre_coef = pre_coef + pre_cursor_step_size) {
    for (post_coef = 0; post_coef ≤ post_cursor_limit; post_coef = post_coef +
post_cursor_step_size) {
        pre = (pre_coef × RemoteFS) / 64
        post = (post_coef × RemoteFS) / 64
        // Core requests link partner PHY Tx use {post, pre} coefficients and then core
performs a Rx EQ Evaluation
        try {post, pre}
    }
}
```

After all trials are completed, one additional Rx EQ Evaluation is executed using the {post, pre} coefficients from the trial which achieved the highest Figure of Merit. If the final Figure of Merit meets or exceeds the threshold configured in eq_fmerit_control_req_feedback, the Rx Equalization is considered successful otherwise unsuccessful. This final evaluation is skipped if the current coefficients being used are the best coefficients.

Notes:

- Pre-Cursor coefficients (pre) from 0 to 16 (0 to 0.25) are possible.
- Post-Cursor coefficients (post) from 0 to 32 (0 to 0.5) are possible.

Stepping through all 17 (0-16) Pre-Cursor values and all 33 (0-32) Post-Cursor values takes 561 iterations. Step size is increased to walk through the values more quickly (and coarsely). Limits are lowered to exclude larger values that are less likely to produce the desired results.

Example:

- Steps of 4 for Pre-Cursor and 8 for Post Cursor with Limits == 16 and 32 respectively requires 25 iterations.
- Steps of 8 for Pre-Cursor and 16 for Post Cursor with Limits == 16 and 32 respectively requires 9 iterations.

Caution: Be careful when assigning eq_table_method_control_addr_limit not to exceed the Equalization time limit. Refer to [Equalization Time Limit](#) section for more details.

The Algorithm Method works best with PHY, which take significantly less than the maximum of 2 ms to evaluate equalization settings so that the fine step sizes can be used.

2.5.3.3. Figure of Merit – Table Method

This method selects and sweeps the preset or coefficients configured in the table.

```
for (i = 0; i ≤ eq_table_method_control_addr_limit; i = i + 1) {
    // Read the current table entry
    table_pre = eq_table_method_table_array[(i×16)+5:(i×16)+0]
    table_post = eq_table_method_table_array[(i×16)+11:(i×16)+6]
```

```

    table_interpret = eq_table_method_table_array[(i×16)+13:(i×16)+12]
    table_best      = eq_table_method_table_array[(i×16)+14]           // unused in
this method
    // Interpret the table entry
    if (table_interpret == 0) { // Current table entry specifies a preset
        pre_preset_coef = pre-cursor coef for Preset[table_pre[3:0]] // See Preset Method for
coefficients used
        post_preset_coef = post-cursor coef for Preset[table_pre[3:0]] // See Preset Method for
coefficients used
        pre = (pre_preset_coef × RemoteFS) / 64
        post = (post_preset_coef × RemoteFS) / 64
    } else { // Current table entry specifies coefficients
        pre = (table_pre[5:0] × RemoteFS) / 64
        post = (table_post[5:0] × RemoteFS) / 64
    }
    // Core requests link partner PHY Tx use {post, pre} coefficients and then core performs a
Rx EQ Evaluation
    try {post, pre}
}

```

After all (`eq_table_method_control_addr_limit+1`) trials are completed, one additional Rx EQ Evaluation is executed using the {post, pre} coefficients from the trial which achieved the highest Figure of Merit. This final evaluation is skipped if the current coefficients being used are the best coefficients. If the final Figure of Merit meets or exceeds, the threshold configured in `eq_fmerit_control_req_feedback` then Rx Equalization is considered successful otherwise unsuccessful.

Notes:

- Pre-Cursor coefficients (pre) from 0 to 16 (0 to 0.25) are possible.
- Post-Cursor coefficients (post) from 0 to 32 (0 to 0.5) are possible.

In this method, you specify up to 24 Preset/coefficients to try and may select the preset/coefficients that are most likely to work for the given PHY.

The Table Method works well for users that know the range of coefficients, which typically works well for the PHY since the table values can be concentrated on coefficient ranges that are more likely to work well. The Table Method also provides a lot of flexibility and can be configured easily.

Caution: Be careful when assigning `eq_table_method_control_addr_limit` not to exceed the Equalization time limit. Refer to [Equalization Time Limit](#) section for more details.

2.5.3.4. Figure of Merit – Adaptive Table Method

This method sweeps through the user-provided table with adaptive coefficient selection. The Adaptive Table Method is similar to the Table Method and uses the same user-configured table. Although, the Adaptive Table Method may be configured to select latter coefficients dynamically from the results of earlier equalization evaluations.

```

// Initialize the current Up/Down best coefficient pair {best_post, best_pre} to {0,0}
best_pre = 0; best_post = 0; best_fom = 0;
// Initialize the current Relative best coefficient pair {rel_best_post, rel_best_pre} to {0,0}
rel_best_pre = 0; rel_best_post = 0;
for (i = 0; i ≤ eq_table_method_control_addr_limit; i = i + 1) {
    // Read the current table entry
    table_pre      = eq_table_method_table_array[(i×16)+5:(i×16)+0];
    table_post     = eq_table_method_table_array[(i×16)+11:(i×16)+6];
    table_interpret = eq_table_method_table_array[(i×16)+13:(i×16)+12];
    table_best     = eq_table_method_table_array[(i×16)+14];
    // Interpret the current table entry
    if (table_interpret == 0) { // Current table entry specifies a preset

```

```

        pre_preset_coef = pre-cursor coef for Preset[table_pre[3:0]]; // See Preset Method for
coefficients used
        post_preset_coef = post-cursor coef for Preset[table_pre[3:0]]; // See Preset Method
for coefficients used
        pre = (pre_preset_coef × RemoteFS) / 64;
        post = (post_preset_coef × RemoteFS) / 64;
    } else if (table_interpret == 1) { // Current table entry specifies coefficients
        pre = (table_pre[5:0] × RemoteFS) / 64;
        post = (table_post[5:0] × RemoteFS) / 64;
    } else if (table_interpret == 2) { // Table entry contains relative coefficient offsets
        pre = table_pre[5] ? (rel_best_pre - table_pre[1:0]) : (rel_best_pre + table_pre[1:0]);
        post = table_post[5] ? (rel_best_post - table_post[1:0]) : (rel_best_post +
table_post[1:0]);
    } else { // Use the prior evaluation's PHY Up/Down Feedback instead of table_pre,
table_post
        if (prior trial used (table_interpret == 3)) { // Use prior trial feedback when
continuing (table_interpret==3)
            case (Prior trial's predir)
                10      : pre = best_pre - 1;
                01      : pre = best_pre + 1;
                default  : pre = best_pre;
            endcase
            case (Prior trial's postdir)
                10      : post = best_post - 1;
                01      : post = best_post + 1;
                default  : post = best_post;
            endcase
        } else { // Begin (table_interpret == 3) trials from the current best coefficient pair
            pre = best_pre; post = best_post;
        }
    }
}
// Core requests link partner PHY Tx use {post, pre} coefficients and then core performs a
Rx EQ Evaluation
{fom, postdir, predir} = try {post, pre} // fom == Figure of Merit feedback; post/predir ==
directional feedback
// Better result or, when using Up/Down feedback, the most recent trial is always
considered the best
if ((fom ≥ best_fom) | (table_interpret == 3)) {
    best_fom = fom
    best_pre = pre
    best_post = post
}
if (table_best == 1) {
    rel_best_pre = best_pre
    rel_best_post = best_post
}
if ((table_interpret == 3) and ({postdir, predir} == {0, 0}) and core configured to end on
hold) // {Hold, Hold}
    exit for i loop
}

```

After all (*eq_table_method_control_addr_limit*+1) trials are completed, one additional Rx EQ evaluation is executed using the {post, pre} coefficients from the trial, which achieved the highest Figure of Merit. If the final Figure of Merit meets or exceeds, the threshold configured in *eq_fmerit_control_req_feedback* then Rx Equalization is considered successful otherwise unsuccessful. This final evaluation is skipped if the algorithm ends using *table_interpret* == 3 (Up/Down Feedback) or the best coefficients are already in use.

Notes:

- Pre-Cursor values (pre) from 0 to 16 (0 to 0.25) are possible.
- Post-Cursor values (post) from 0 to 32 (0 to 0.5) are possible.

The Adaptive Table includes 24 table entries.

While performing the algorithm, coefficient subtractions are limited to 0, coefficient additions are limited to 63 and coefficients are limited (as in all cases), when necessary, to comply with the PCIe Specification coefficient rules.

The Adaptive Table Method can perform a coarse search to quickly identify the region with the best Figure of Merit and to select the best specific coefficients in that region with a fine-grained search. The Adaptive Table Method works well for a broad range of conditions. Although, it works best when 10 to 20 iterations are possible so that it has time to select the best coefficients.

The Adaptive Table Method is very flexible and supports many different use models. The following use models are suggested:

- Coarse Figure of Merit search followed by fine PHY-directed Up/Down search.
- Fill the first N table entries with presets or coefficients that are widely dispersed setting *interpret* == 0 (absolute Preset) or 1 (absolute coefficients) as desired. Choosing Presets/coefficients that coarsely cover the coefficient space is recommended. Set *best* == 1 on all N table entries. This results in the following Up/Down feedback evaluations starting from the preset/coefficients that received the highest Figure of Merit in the first N trials.
- Fill the next M table entries with *pre* == 0 (unused), *post* == 0 (unused), *interpret* == 3 (use the prior evaluation's PHY Up/Down Feedback), and *best*=0. This results in M trials where the PHY controls the next coefficients through its Up/Down/Hold feedback.
- The algorithm finishes the M trials and then exits Equalization. However, an early exit occurs if *eq_table_method_control_end_on_hold*==1 and all lanes provided a {Hold, Hold} response.
- Coarse Figure of Merit search followed by fine Figure of Merit search.
- Fill the first N table entries with presets or coefficients that are widely dispersed, setting *interpret* == 0 (absolute Preset) or 1 (absolute coefficients) as desired. Choosing Presets/coefficients that coarsely cover the coefficient space is recommended. Set *best* == 1 on all N table entries. This results in the following relative offset evaluations starting from the Preset/coefficients that received the highest Figure of Merit in the first N trials.
- Repeat the following sequence as many times as desired and the Equalization time limit allows.
- Fill the next four table entries with relative offsets moving up, down, right, and left by 1 and setting *best*==1 on only the fourth table entry:
 - *pre* == 01 (+1), *post* == 00 (+0), *interpret* == 2 (relative), *best*=0
 - *pre* == 20 (-1), *post* == 00 (+0), *interpret* == 2 (relative), *best*=0
 - *pre* == 00 (+0), *post* == 01 (+1), *interpret* == 2 (relative), *best*=0
 - *pre* == 00 (+0), *post* == 20 (-1), *interpret* == 2 (relative), *best*=1
- After each four-table entry iteration, the coefficient may move up to 1 pre or 1 post from the prior best coefficient pair and the new best coefficient pair is selected as the new relative starting location for subsequent iterations. After three iterations of four table entries each, it is possible for the algorithm to move pre or post up to +/- 3 from the starting best coefficient pair established by the first N trials.
- This algorithm can be modified to move in initially larger relative steps (for example, +2 instead of +1), to use different combinations of relative step directions, or to use more offsets in each trial. Additionally, other offset shapes can be utilized. For example, using a combination of X shape ({+1,+1}, {+1,-1}, {-1,-1}, {-1,+1} relative movements and + shape ({+1,0}, {-1, 0}, {0,+1}, {0,-1} relative movements enables the coefficient space to be covered more or less quickly and more or less completely.

The Adaptive Table Method is flexible enough to reproduce the other Equalization Methods. However, the other methods have been kept due to the simplicity to which they can be enabled and configured.

Caution: Be careful when assigning *eq_table_method_control_addr_limit* not to exceed the Equalization time limit. Refer to [Equalization Time Limit](#) section for more details.

2.5.3.5. Up/Down – Up/Down Convergence Method

The Up/Down convergence method supports PHYs that provide Up/Down feedback response. The Up/Down Convergence Method is described as follows:

- Request the remote device to use an initial set of coefficients. These coefficients can be programmed uniquely for each lane or alternatively, the first coefficients can be provided as a preset.
- After the remote device begins transmitting with the new coefficients, the PHY is told to evaluate the receive waveform (RxEval). The PHY then provides the up/down response for the pre-cursor and post-cursor coefficients.
- The state machine adjusts the pre-cursor and post-cursor coefficients based on the feedback from the PHY. It starts with larger adjustments (initialized with eq_updn_pre_step and eq_updn_post_step). Each time the feedback changes up/dn direction, the step size is halved until it is equal to 1. When the step size is 1 and a direction change occurs in the feedback, then the coefficient is considered converged.
- If any coefficient changes value, the alternate coefficient is allowed to continue to change, even if previously converged since the two coefficients work together.
- If a hold result is received from the PHY, then the coefficient is considered converged. In some PHYs, two hold results in a row must be seen to be considered a hold since only one coefficient is evaluated with each request (the other having a null hold status). In this case, the configuration bit eq_updn_numhold must be set to 1 so the state machine knows to look for two holds.
- Once both coefficients on all active lanes reach convergence, the equalization is complete.

2.5.3.6. Up/Down – Firmware Controlled Method

The Firmware Controlled Method is described as follows:

- The Lattice PCIe x8 IP Core is configured for Firmware to receive an interrupt at the beginning of Rx Equalization (Phase 2 entry for US port and Phase 3 entry for DS port). When an Interrupt is received by Firmware, the Firmware begins a time-critical processing loop to perform Rx Equalization in a minimum amount of time as there is a PCIe Specification-mandated 24 ms time limit to complete Rx Equalization.
- Firmware receives interrupt and begins Rx Equalization evaluation.

```
// Firmware keeps performing Equalization trials until Rx Equalization is complete
done = 0
while (done == 0) {
    Firmware calculates a coefficient pair to try
    Firmware writes the trial coefficients into the core's registers and sets advance=1 to
    begin an evaluation
    Firmware enters a polling loop to wait for the core to complete the requested
    evaluation {
        Core negotiates with the link partner to change to the desired coefficients
        Core performs an Rx Equalization evaluation
        Core updates Equalization status registers with the results of the evaluation
        Core sets Equalization status register complete == 1 to inform firmware that the
        evaluation is done
        Firmware reads complete == 1 while polling and exits the polling loop
    }
    Firmware reads Equalization status registers to obtain Rx Equalization status
    if Firmware is satisfied with the Equalization status {
        Firmware sets the core's complete==1 register to exit Rx Equalization
        done = 1
    }
}
```

- Firmware may configure the core to output interrupts instead of Firmware polling for Equalization completion. Firmware may configure the core receive an interrupt each time that the core is ready to receive a new set of coefficients. Whichever method is selected, Equalization must complete in < 24 ms to avoid the PCIe Specification-required LTSSM timeout, so it is critical that the Firmware algorithm be designed to be able to complete within the time period. Refer to [Equalization Time Limit](#) section for more details.

2.5.4. Equalization Quality

The Figure of Merit Equalization processes rely on the PHY to perform Receiver Equalization using the current remote transmitter settings and produce a quality result (`pipe_eq_rx_eval_feedback_fom`) that can be used to gauge the quality of the link. Quality is defined by the PHY. For the Figure of Merit method, the core contains a field in Management Interface to determine what PHY quality level corresponds to the necessary 10-12 Bit Error Rate (BER). For the Figure of Merit method, higher quality numbers represent better link quality (lower BER).

When using the Up/Down Convergence method, the acceptable link quality is assumed when the PHY converges on a set of coefficients or exits its algorithm due to reaching the programmed maximum iteration count.

2.6. Multi-Function Support

The Lattice PCIe x8 IP Core supports 1 to 4 functions. The Multi-Function support can only be enabled for endpoints (functions implementing Type 0 Configuration Space). See the function register 0x8 for the register configuration.

When Multiple Function support is present, each function is assigned a static function number, starting at function number 0 and incrementing upwards. For ports that communicate function-specific information, `port[0]` applies to Function[0], `port[1]` applies to Function[1]. If a function is disabled, it does not affect the function number of the other enabled functions. Function [0] is always present and cannot be disabled.

2.7. Power Management

2.7.1. Power Management Supported by PCIe IP Core

The Lattice PCIe x8 IP Core supports L0, ASPM L0s, ASPM L1, L1 PM Substates, L1, and L3 link states. L0 (fully-operational state) and L3 (off) support is always enabled. The remaining link states may be enabled/disabled through the Core Configuration ports. If ASPM L0s, ASPM L1, and L1 PM Substates, or L1 support is enabled, then the user design must configure the power management capabilities of the core and for some link states, take additional action when link states are entered or exited. This section describes the recommended actions user logic should take to control and react to power management states ASPM L0s, ASPM L1, and L1.

The PCI Express Specification defines the following link states:

- L0 – Active
 - Powered
 - Clock and PLLs active; core clock active
 - All PCI Express Transactions and operations are enabled
- ASPM L0s – Low resume latency, energy saving *standby* state
 - Powered
 - Clock and PLLs active; core clock active
 - PHY transmitter in electrical idle
 - Remote PHY receiver must re-establish symbol lock during L0s exit

- When L0s is enabled by power management software, the core autonomously enters L0s when the transmit side of the link is idle and exits L0s when there is pending information to transmit. The link management DLLPs are required to be transmitted periodically so when a link is otherwise idle, it still enters and exits L0s with regularity to transmit link management DLLPs.
- ASPM L1 – Low resume latency, energy saving *standby* state
 - Powered
 - Clock and PLLs active; core clock active
 - Significant portion of PHY powered down
 - PHY transmitter in electrical idle
 - PHY receiver in electrical idle
 - Deeper power savings but longer resume time than ASPM L0s
 - Remote and local PHY must re-establish symbol lock during L1 exit
 - When ASPM L1 is enabled by power management software, the core autonomously negotiates L1 entry with the link partner after an extended period of link inactivity. The link autonomously returns to L0 when either device in the link has TLPs to transmit.
- L1 – Higher latency, lower power *standby* state
 - Powered
 - Clock and PLLs active; core clock active
 - Significant portion of PHY powered down
 - PHY transmitter in electrical idle
 - PHY receiver in electrical idle
 - Remote and local PHY must re-establish symbol lock during L1 exit
 - The L1 state is entered both under control of power management software
 - L3 – Off
 - Main power off; auxiliary power off
 - In this state, all power is removed and the core, PHY, and user logic are all non-operational
 - All state information is lost

2.7.2. Configuring Core to Support Power Management

The Lattice PCIe IP x8 Core allows user logic to implement a wide variety of power management functionality. The design's power management capabilities are primarily advertised and controlled using core configuration ports.

2.7.3. APSM L0s

The Lattice PCIe x8 IP Core supports Active State Power Management (APSM) L0s. When L0s support is enabled, ASPM L0s TX Entry Time, the desired amount of time for TLP and DLLP transmissions to be idle before L0s TX is entered, is determined by `mgmt_ptl_pm_aspm_l0s_entry_time`. The Number of NFTS sets required by the local PHY to recover symbol lock when exiting L0s is determined by `mgmt_tlb_ltssm_nfts_nfts`. NFTS Timeout Extend, `mgmt_tlb_ltssm_nfts_to_extend` (see [Table 5.12](#)), controls how long the core waits after the expected L0s exit time before directing the link to Recovery to recover from a failed L0s exit. Due to high latencies between a PHY's Rx Electrical Idle output and the associated Rx Data it is normally necessary to choose a relatively high NFTS and NFTS Timeout Extend. See `mgmt_tlb_ltssm_nfts_to_extend` description ([Table 5.12](#)) for details.

- Configuration Register Fields
 - The PCI Express Device Capabilities configuration register has the following L0s fields:
 - Bits [8:6] – Endpoint L0s Acceptable Latency – From PCI Express Base Specification, Rev 2.1 section 7.8.3 – Acceptable total latency that an Endpoint can withstand due to the transition from L0s state to the L0 state. It is essentially an indirect measure of the Endpoint's internal buffering. Power management software uses the reported L0s Acceptable Latency number to compare against the L0s exit latencies reported by all components comprising the data path from this Endpoint to the Root Complex Root Port to determine whether ASPM L0s entry can be used with no loss of performance. Note that the amount of buffering does not refer to the Lattice PCIe x8 IP Core buffering, but rather to user application buffering. You must set this field in accordance with how long a delay is acceptable for the application.

- 000 – Maximum of 64 ns
- 001 – Maximum of 128 ns
- 010 – Maximum of 256 ns
- 011 – Maximum of 512 ns
- 100 – Maximum of 1 μ s
- 101 – Maximum of 2 μ s
- 110 – Maximum of 4 μ s
- 111 – No limit
- Non-Endpoints must hard wire this field to 000
- The PCI Express Link Capabilities configuration register has the following L0s fields:
 - Bits[14:12] – L0s Exit Latency – Length of time required to complete transition from L0s to L0:
 - 000 – Less than 64 ns
 - 001 – 64 ns to less than 128 ns
 - 010 – 128 ns to less than 256 ns
 - 011 – 256 ns to less than 512 ns
 - 100 – 512 ns to less than 1 μ s
 - 101 – 1 μ s to less than 2 μ s
 - 110 – 2 μ s-4 μ s
 - 111 – More than 4 μ s
 - Exit latencies may be significantly increased if the PCI Express reference clocks used by the two devices in the link are common or separate.
 - Bits[11:10] – Active State Power Management (ASPM) Support must be set to 01 or 11 if L0s support is enabled or 00 otherwise.

2.7.4. APSM L1s

The Lattice PCIe x8 IP Core supports both software controller L1 entry (through the Power State Configuration Register) and hardware autonomous L1 entry (Active State Power Management (APSM) L1).

- Software-controlled L1 flow for Upstream Ports (Endpoint) is as follows:
 - Software initiates changing a link to L1 by writing the core's Power Management Capability: Power State Configuration Register to a value other than 00 == D0. Note that the component's Device driver participates in this process and must ensure that all traffic is idle before permitting the system to power down to L1.
 - When the core detects a change of Power State to a non-D0 value, the core's power management state machine, which is responsible for the higher-level power management protocol, follows the following sequence:
 - Block further TLP transmissions
 - Wait for all in process TLPs to complete transmission
 - Wait for the Replay Buffer to empty (all transmitted TLPs acknowledged)
 - Core transmits PM_ENTER_L1 DLLPs until receiving a PM_REQ_ACK DLLP from remote device
 - Core directs LTSSM state machine to L1
 - When a TLP is pending or the LTSSM state machine indicates L1 state has been exited due to link partner activity, the core returns to L0.
- Configuration Register Fields:
 - The PCI Express Device Capabilities configuration register has the following L1 fields:
 - Bits [11:9] – Endpoint L1 Acceptable Latency – From PCI Express Base Specification, Rev 2.1 section 7.8.3 – This field indicates the acceptable latency that an Endpoint can withstand due to the transition from L1 state to the L0 state. It is essentially an indirect measure of the Endpoint's internal buffering. Power management software uses the reported L1 Acceptable Latency number to compare against the L1 Exit Latencies reported (see below) by all components comprising the data path from this Endpoint to the Root Complex Root Port to determine whether ASPM L1 entry can be used with no loss of performance. Note that the amount of buffering does not refer to Lattice PCIe x8 IP Core buffering, but rather to user application buffering. You must set this field in accordance with how long a delay is acceptable for the application.
 - 000 – Maximum of 1 μ s
 - 001 – Maximum of 2 μ s

- 010 – Maximum of 4 μ s
 - 011 – Maximum of 8 μ s
 - 100 – Maximum of 16 μ s
 - 101 – Maximum of 32 μ s
 - 110 – Maximum of 64 μ s
 - 111 – No limit
 - Non-Endpoints must hard wire this field to 000.
- PCI Express Link Capabilities configuration register has the following L1 fields:
 - Bits[17:15] – L1 Exit Latency – Length of time required to complete transition from L1 to L0:
 - 000 – Less than 1 μ s
 - 001 – 1 μ s to less than 2 μ s
 - 010 – 2 μ s to less than 4 μ s
 - 011 – 4 μ s to less than 8 μ s
 - 100 – 8 μ s to less than 16 μ s
 - 101 – 16 μ s to less than 32 μ s
 - 110 – 32 μ s-64 μ s
 - 111 – More than 64 μ s
 - Exit latencies may be significantly increased if the PCI Express reference clocks used by the two devices in the link are common or separate.
 - Bits[11:10] – Active State Power Management (ASPM) Support must be set to 10 or 11 if L1 support is enabled or 00 otherwise.
- Hardware-autonomous L1 (ASPM L1) entry is initiated only by Upstream Ports (Endpoint). The core ASPM L1 functionality must be enabled and advertised in PCIe Link Capabilities and software must enable ASPM L1 support in order for the hardware-autonomous L1 to be negotiated. When ASPM L1 support is present and enabled for an Upstream Port, the core requests the link to be directed to L1 using the ASPM L1 protocol, when the link is idle. The link idle refers to the no TLPs or ACK/NAL DLLPs being transmitted.

2.8. DMA Support

The Lattice PCIe x8 IP Core has a DMA Bridge Core that implements high-performance DMA and bridging between PCI Express and AXI. Following are the key features of the DMA Bridge Core.

- AXI Manager Interface
 - High-performance AXI4 Manager to complete PCIe to AXI bridge and DMA transactions.
 - 64-bit address support.
 - Includes support for up to 8 PCIe to AXI address translations.
- AXI Subordinate Interface
 - High-performance AXI4 Subordinate to complete AXI to PCIe bridge transactions.
 - 64-bit address support.
 - Includes support for up to 8 AXI to PCIe address translations.
- High-performance Multi-Channel DMA
 - Up to 64 DMA channels support
 - Each DMA channel supports 4 directions of DMA data transfer.
 - PCIe to PCIe, PCIe to AXI, AXI to PCIe, and AXI to AXI

Note: For bridge core support, user needs to ensure that an additional BAR (BAR1 – BAR5) is enabled.

2.8.1. DMA Scatter-Gather and Status Queues

The DMA Channels can perform Scatter-Gather separately for the DMA Source and DMA Destination. A Source Scatter Gather and a Destination Scatter Gather Queue is managed by software and the DMA Channel to describe the desired DMA operations. The DMA Queue Registers describe the Starting Queue Address Offsets, Queue Size, and communicate flow control ownership between software and the DMA Channel hardware.

Software provides Source and Destination Queue elements to the DMA Channel to enable the DMA Channel to perform a DMA operation. When the DMA Channel is enabled, the available Source and Destination Queue elements are fetched into a small internal buffer, on need, to complete the DMA operations.

Queue elements are a fixed size. The address of a particular element in the queue is the queue starting address plus the element index multiply the element size.

Source Scatter-Gather Queue Elements

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	Byte Offset
Source Address [31:0]																																+000h
Source Address [63:32]																																+004h
Flags																Byte Count																+008h
User ID																User Handle																+00Ch

Figure 2.8. Source Scatter-Gather Queue Elements

The Source Scatter-Gather Queue elements are 128-bits.

- [63:0] – Source Address [63:0]
- [87:64] – Byte Count [23:0]
 - 0 – 2²⁴ bytes
- [95:88] – Flags [7:0]
 - [7:4] – DMA Data Read Attributes
 - If source is AXI [7:4] is used for m0_axi4_arcache_o [3:0]
 - If source is PCIe [6:4] is used for PCIe Attributes [2:0]
 - [3] – Reserved.
 - [2] – Interrupt
 - 1 – Generate an interrupt event when the Status Queue is written with DMA Completion Status; only valid when EOP is 1; the interrupt will be generated on PCIe, AXI, or both per the DMA Channel's interrupt registers.
 - 0 – Do not generate an interrupt
 - [1] – EOP
 - 1 – End of packet; a status Queue Element containing DMA packet completion status is written whenever an End of Packet is transferred to the DMA Destination
 - 0 – Not the end of a packet; packets may span multiple Descriptors.
 - [0] – Location
 - 1 – DMA data source is AXI
 - 0 – DMA data source is PCIe

- [111:96] – User Handle [15:0]
 - User Handle is copied from Source Scatter-Gather Queue element when EOP is 1 into the corresponding Source DMA Completion Statue Queue elements' User Handle field.
- [127:112] – User ID [15:0]
 - User ID is copied from Source Scatter-Gather Queue element when EOP is 1 into the corresponding Source DMA Completion Statue Queue elements' User ID field.

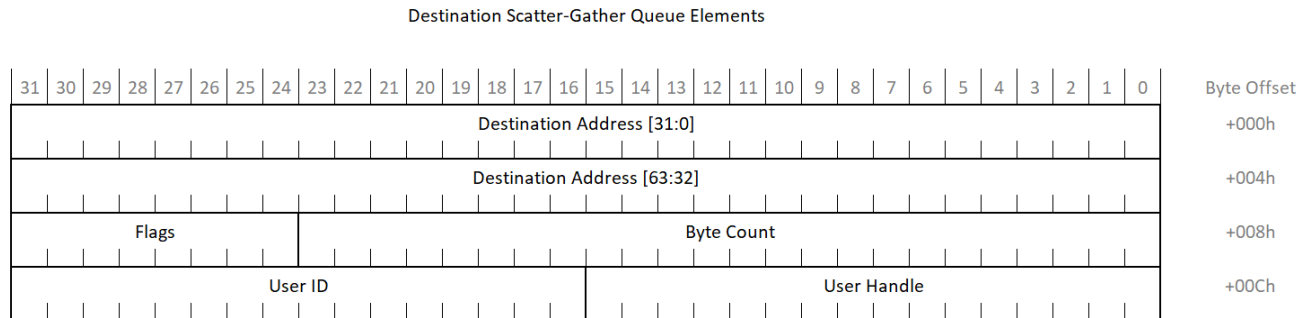


Figure 2.9. Destination Scatter-Gather Queue Elements

The Destination Scatter-Gather Queue elements are 128-bits.

- [63:0] – Destination Address [63:0]
- [87:64] – Byte Count [23:0]
 - 0 – 2²⁴ bytes
- [95:88] – Flags [7:0]
 - [7:4] – DMA Data Write Attributes
 - If source is AXI [7:4] is used for m0_axi4_arcache_o [3:0]
 - If source is PCIe [6:4] is used for PCIe Attributes [2:0]
 - [3:2] – Reserved.
 - [1] – Enable One Packet Per Destination Scatter-Gather Queue
 - 1 – Skip to next Destination Scatter-Gather element on EOP
 - Whenever an end of packet is transferred (as determined by the Source Scatter-Gather element EOP field) the remaining portion of the current Destination Scatter-Gather element (if any) is not used. The next packet starts being written into the beginning of the memory described by the following Destination Scatter-Gather element.
 - 0 – Pack Packets Back-to-Back in Destination Scatter-Gather Element
 - When a packet ends without consuming the entire Destination Scatter-Gather element, the next packet begins will be written into the remaining portion of the current Destination Scatter-Gather element location at the byte following the end of the prior packet.
- [0] – Location
 - 1 – DMA data source is AXI
 - 0 – DMA data source is PCIe

- [111:96] – User Handle [15:0]
 - User Handle is copied from the final Destination Scatter-Gather element used for a packet transfer into the corresponding Destination DMA Completion Status Queue element User Handle field.
- [127:112] – Reserved [15:0]




Figure 2.10. Status Queue Elements

The Status Queue elements are either 32-bit or 64-bit as determined by software DMA Channel register programming.

- [63:48] – User ID [15:0]
 - User ID is copied from Source Scatter-Gather element with EOP equals to 1 into the corresponding Source and Destination DMA Completion Status Queue elements' User ID fields.
 - User ID use is application specific. The DMA Channel copies but does not itself use this information.
- [47:32] – User Handle [15:0]
 - DMA Source Completion Status Queue – User Handle is copied from Source Scatter-Gather element with EOP equals to 1.
 - DMA Destination Completion Status Queue – User Handle is copied from the final Destination Scatter-Gather element used for the packet.
 - User Handle use is application specific. The DMA Channel copies but does not itself use this information. 0 – 2²⁴ bytes
- [31] – Upper Status is Non-Zero
 - For 64-bit status elements, Upper Status Is Non Zero equals to 1 when Status Queue Element bits [63:32] is non-zero. In some cases, CPUs are only capable of 32-bit atomic transactions. For such CPUs, the software reads status bits [31:0]. If bit [31] equals to 1, then software knows that Status Queue Element [63:32] is non-zero and reads status bits [63:32] until a non-0 value is returned. Software must clear all status Queue elements before giving them to the DMA Channel, so that the upper status bits are zero until updated by a DMA Completion.
 - For 32-bit status elements, this field is unused and always reads 0.
- [30:4] – Completed Byte Count [26:0]
 - The possible range is 0 to (2²⁷)-1 bytes (0 to 128MByte-1). If a DMA exceeds (2²⁷)-1 bytes, Completed Byte Count limits at its maximum value. If software needs to transfer larger packets than 128MB-1 and keep an accurate byte count, then software must split such transfers into smaller DMA transfers that are < 128 MB.
- [3] – Internal Error Detected during DMA Operation
- [2] – Destination Error Detected during DMA Operation
- [1] – Source Error Detected during DMA Operation
- [0] – Completed

2.8.2. DMA Channel Programming and Operation

Each DMA Channel implements an independent register set for controlling that DMA Channel's DMA operations. Registers are provided for Scatter-Gather and Completion Status Queue Management and DMA Control and Status. DMA registers are implemented in structures that can be implemented mostly in RAM to enable large numbers of DMA Channels to be efficiently supported.

DMA registers are described in the "DMA_CHANNEL" register section. When Multi-Function and/or SR-IOV capable PCI Express Cores are used, DMA Channels are allocated in equal portions to each function.

DMA Queues are required to be contiguous memory and address aligned to 64 bytes to enable efficient read/write of multiple queue elements in one read/write operation. This enables higher DMA bandwidth for smaller DMA transfer sizes since the overhead of queue management operations can be amortized over multiple small DMA transfers.

Source and Destination Scatter Gather Queues describe the fragmentation of the Source and Destination memory respectively. The queues are independent. The DMA Channel merges the information from the queues to perform DMA operations considering each Queue's fragmentation. This relieves software from this computation-intensive task and the necessity to pass the fragmentation information between the PCIe and AXI domains.

Prior to enabling the DMA Channel, the Source Scatter-Gather Queue, Destination Scatter-Gather Queue, DMA Source Completion Status Queue, and DMA Destination Completion Status Queue must be initialized. The following DMA Channel Enable sequence is recommended. The process can be completed in any order desired provided DMA Enable is written to 1 last:

- Verify that the DMA Channel is Idle.
 - Read DMA Running and verify it reads 0x0.
- Initialize queue base address and attributes.
 - Write SRC_Q_PTR_LO and SRC_Q_PTR_HI with the base address of the queue.
 - Write DST_Q_PTR_LO and DST_Q_PTR_HI with the base address of the queue.
 - Write STAS_Q_PTR_LO and STAS_Q_PTR_HI with the base address of the queue.
 - Write STAD_Q_PTR_LO and STAD_Q_PTR_HI with the base address of the queue.
- Initialize queue size.
 - Write SRC_Q_SIZE to the size of the queue.
 - Write DST_Q_SIZE to the size of the queue.
 - Write STAS_Q_SIZE to the size of the queue.
 - Write STAD_Q_SIZE to the size of the queue.
- Initialize queue Next pointers to the beginning of the queue.
 - Write SRC_Q_NEXT = 0x0
 - Write DST_Q_NEXT = 0x0
 - Write STAS_Q_NEXT = 0x0
 - Write STAD_Q_NEXT = 0x0
- Initialize Scatter Gather Queues to the Empty Condition (no DMA operations to execute)
 - Write SRC_Q_LIMIT = 0x0
 - Write DST_Q_LIMIT = 0x0
- Initialize Status Queues to the Fully Available Condition (all Status Queue elements except one, which is needed to preserve software flow control, are available)
 - STAS_Q_LIMIT set to (STAS_Q_SIZE-1)
 - STAD_Q_LIMIT set to (STAD_Q_SIZE-1)
- Initialize all STAS and STAD Queue Elements to 0x0.
 - DMA Completion Status Queue elements must be initialized to 0x0 so that software can tell when Status elements complete. Status elements return a non-0 value when complete.
- Optionally Initialize all SRC and DST Scatter Gather Queue Elements.
 - Source and Destination SGL Elements may be initialized or not at software discretion. Source and Destination SGL Elements will not be fetched until they have been filled in with DMA transaction instructions and the associated queue's LIMIT pointer advanced to give these elements to the DMA Channel to execute.
 - Write DMA Enable = 1 to enable the DMA Channel.

DMA transactions are given to the DMA Channel to execute by writing Source SGL and Destination SGL into the SRC/DST SGL Queues and incrementing the SRC/DST_Q_LIMIT registers.

When the DMA Channel completes a DMA operation (completes the transfer of a Source SGL that had EOP ==1), the STAS/STAD Completion Status Queues are updated with the DMA Completion Status and then a DMA Interrupt is generated (as conditioned by the DMA Channel Interrupt Coalesce Count and other registers).

Queue Management registers *_PTR_LO, *_PTR_HI, *_SIZE, and *_NEXT must not be written when the DMA Channel is enabled. The only queue registers that it is permissible to modify while the DMA Channel is enabled is the SRC/DST/STAS/STAD_Q_LIMIT registers which are incremented to provide additional elements for the DMA Channel to execute.

The minimum Queue Size is 2. This is because software must always retain ownership of one Q element for Flow Control. The DMA Channel does not execute/use the Queue element pointed to by the Queue LIMIT pointer.

The minimum Queue Size must be large enough to hold at least 1 full DMA transaction of maximum size. DMA Completion Status Queues are only written when a Source SGL Element is completed that had its EOP flag == 1 (end of a DMA transaction). If the Queue is too small to be able to place all the SGL for a single DMA transaction into the Queue, then the SGL with EOP == 1 will not be able to be added to the Queue and the DMA operation will not be able to complete. S/W will not be able to free Queue elements because no SGL with EOP can complete (because none is in the Queue) and no new SGL can be given to the DMA Channel unless Queue elements are freed.

A Queue Size of N has N Q elements: [0], [1], ..., [N-1]. For example, a Q Size of 2 has [0] and [1] elements. The Queue wraps at the N-1 element. For example, for a Queue Size of 2 the following wrap occurs: [0], [1], [0], ... The DMA Queues are intended to be setup once and re-used for multiple DMA operations. The DMA Queues are designed to enable highly overlapped transactions. Software can be setting up new DMA operations in the Queue while the DMA Channel is executing operations that software placed in the queue earlier.

The DMA Queues can also be setup for each DMA transaction if desired, although this method has lower performance, due to the need to reconfigure the queues between DMA transactions. DMA Queues may only be reconfigured when the DMA Channel is disabled. So, the process would be wait for all outstanding DMA transactions for the Channel to complete, disable the DMA Channel, reconfigure the Queues, and Re-enable the DMA Channel.

Since the DMA operations to perform are specified by the Source and Destination Scatter Gather Elements that are pointed to by the Queues, rather than being in the Queues, the Queues do not contain specific DMA transfer information. The same Queues are then readily re-used for different DMA transactions.

When a DMA transaction completes (software reads the current DMA Completion Status Queue element and reads Status == Complete), software processes the resulting DMA data and recycles the Source/Destination SGL Queue elements associated with the transfer as well as the associated DMA Source/Destination Completion Status Queue elements. Status Queue elements must be written to 0x0 when recycled since software uses a read of non-0 for a Status Queue element as indication that a DMA transfer completed. Recycled Queue elements are re-used when the queue LIMIT pointer wraps back to the position of the recycled elements.

2.8.3. DMA Performance

Table 2.8 shows the performance of the PCIe Harden DMA. Performance is measured from the start of DMA operation to until the MSI is received by the host. The values below are based on simulation using multiple descriptors of the same size. Bigger descriptors can achieve better efficiency as there are less descriptors and data transfer overhead.

2.8.3.1. FPGA-to-Host (F2H) Simulation Transfer

Table 2.8. Simulation Data Throughput Using a Different Descriptor Size for FPGA-to-Host (F2H) Transfer

Descriptor Size	Efficiency	Throughput
256 kB	77.19%	12.350 GB/s

2.8.3.2. Host-to-FPGA (H2F) Simulation Transfer

Table 2.9. Simulation Data Throughput Using a Different Descriptor Size for Host-to-FPGA (H2F) Transfer

Descriptor Size	Efficiency	Throughput
256 kB	59.19%	9.470 GB/s

Note: The throughput is limited because the hardened DMA only utilizes 16 PCIe tag for memory read. This is a hard IP limitation.

2.8.4. DMA Bypass Interface

The DMA Bypass mode allows the received MWr and MRd TLP to be converted to the AXI-MM Interface. DMA Bypass cycles can be routed through any BARs of the PCIe memory, except for BAR 0.

For read access, Read Data Channel from AXI-MM is converted to CplD TLP by the IP and transmitted to the host.

2.9. Non-DMA Support

2.9.1. Non-DMA Overview

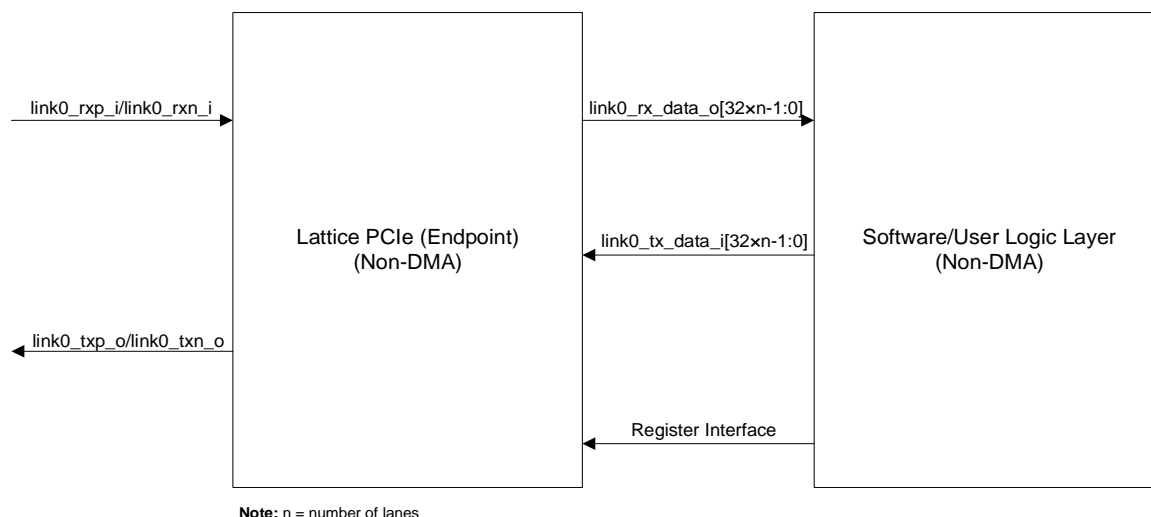


Figure 2.11. Non-DMA Application Data Flow

For the non-DMA design, the PCIe EP receives the data through the *link0_rxp_i/link0_rxn_i* serial lines from the Root Complex. The PCIe EP converts the serial data in the form of TLP packets. The TLP packets are sent to the non-DMA application layer through the *link0_rx_data_o* signal. The TLP header info is decoded and the operation is decided whether the data is written or read. For the write operation, the data is written to the RAM present in the application layer. For the read operation, the data is read from the RAM and sends the encoded data to the PCIe EP in the form of TLP packets through the *link0_tx_data_i* signal.

The register interface is enhanced as per the data interface selected in the user interface.

Table 2.10. Register Access for Different Data Interfaces

Data Interface	Register Interface
TLP	LMMI
AXI4_Stream	AXI-L

2.9.2. Non-DMA Write

The PCIe EP sends the header data to the non-DMA application layer through the *link0_rx_data_o* signal. The application layer initially verifies the operation by decoding the header information. Once the write operation is detected, the user data is received along with the valid signal from the PCIe IP. The valid data is stored in the RAM present in the application layer.




Figure 2.12. Non-DMA Write Operation (TLP Data Interface)

2.9.3. Non-DMA Read

The PCIe EP sends the header data to Non-DMA Application layer through the *link0_rx_data_o* signal. The application layer initially verifies the operation by decoding the header information. Once the operation is detected as read, the application layer waits for the ready signal sent by the PCIe EP. Based on the ready signal and header address, the user data along with the valid signal is send to PCIe EP by the RAM present in the application layer through the *link0_tx_data_i* signal.




Figure 2.13. Non-DMA Read Operation (TLP Data Interface)

For more details on the TLP write and read data transaction, refer to the [Register Interface Conversion](#) section.

2.10. Interrupts

2.10.1. Generation of the Interrupts

The Lattice PCIe Core IP supports the Legacy Interrupts, Messaged Signaled Interrupts (MSI), and MSI-X interrupts.

For each function in the PCIe IP core, the system software configures the function to use MSI-X, MSI, or Legacy Interrupt mode as part of the PCI enumeration process.

The Legacy Interrupt is supported by the PCIe Core to support the backward compatibility by enabling the INTx pins.

To minimize the pin count, the function can generate the inband interrupt message packet to indicate the assertion and de-assertion of an interrupt pin. These are the MSI and MSI-X interrupts. This interrupt mechanism is used to conserve the pins because it does not use separate wires for interrupts.

In this mechanism a single Dword provides the information about the interrupt messages MSI-X/MSI interrupts are signaled using MSI-X/MSI Message TLPs, which you can generate and transmit in the Transmit Interface.

The MSI Interrupt is a posted memory write, which is distinguished from the other memory writes by the addresses they target, which are typically reserved by the system for interrupt delivery. The MSI Capability structure is stored in the Configuration Space and is programmed using Configuration Space accesses.

The MSI-X interrupt is the extended version for the MSI interrupts, supporting a greater number of MSI Vectors and the MSI-X capability structure points to an MSI-X table structure and an MSI-X Pending Bit Array (PBA) structure, which are stored in memory.

Enabling and Disabling of interrupts can be done through PCIe IP Core user interface or through Hard IP core configuration status registers.

[Table 2.11](#) describes the register bits to enable and disable each of the interrupts.

Table 2.11. Base Address to Enable Interrupt

Base Address	Offset Address	Register Bits	Description
0x4_4000 (Function 0) 0x4_5000 (Function 1) 0x4_6000 (Function 2)	0x50	[0]	Support for Legacy Interrupts <ul style="list-style-type: none"> 0 – Enable 1 – Disable
0x4_7000 (Function 3) 0x4_8000 (Function 4) 0x4_9000 (Function 5)	0xE8	[0]	Support for MSI Interrupts <ul style="list-style-type: none"> 0 – Enable 1 – Disable

Base Address	Offset Address	Register Bits	Description
0x4_A000 (Function 6) 0x4_B000 (Function 7)	0xF0	[0]	Support for MSI-X Interrupts <ul style="list-style-type: none"> 0 – Enable 1 – Disable

2.10.2. Legacy Interrupt

When the legacy interrupts are enabled, the PCIe IP core emulates the INTx Interrupts using virtual wire. The term INTx refers to the four legacy interrupts: INTA, INTB, INTC, and INTD.

The link0_legacy_interrupt_i signal is used to generate Legacy interrupts on the PCI Express link. The link [LINK]_legacy_interrupt_i has one input for each base (physical) function. When Legacy Interrupt Mode is enabled, link0_legacy_interrupt_i implements one level-sensitive interrupt (INTA, INTB, INTC, or INTD) for each Base Function. Each functions interrupt sources must be logically ORed together and input as link0_legacy_interrupt_i [i] for a given function. Each interrupt source must continue to drive a 1 until it has been serviced and cleared by software at which time it must switch to driving 0. The core ORs together INTA/B/C/D from all functions to create an aggregated INTA/INTB/INTC/INTD. The core monitors high and low transitions on the aggregated INTA/B/C/D and sends an Interrupt assert message on each 0 to 1 transition and an Interrupt deassert message on each 1 to 0 transition of the aggregated INTA/B/C/D. Transitions, which occur too close together to be independently transmitted, are merged.

The core asserts the link0_legacy_interrupt_o, signal, which is an active high level-based interrupt signal. This interrupt is asserted by the core, whenever an interrupt is generated by the core implemented PCI Express Capability and Advanced Error Reporting Capability. The link0_legacy_interrupt_o should be merged with the link0_legacy_interrupt_i signal to produce the any user interrupt signal.

When a function has MSI-X or MSI Interrupt Mode enabled, link0_legacy_interrupt_i is not used for that function.

The selection among the four interrupts can be done through the PCIe IP core user interface or through register interface as described in [Table 2.12](#).

Table 2.12. Legacy Interrupt Register

Base Address	Offset Address	Register Bits	Description
0x4_4000 (Function 0) 0x4_5000 (Function 1) 0x4_6000 (Function 2) 0x4_7000 (Function 3) 0x4_8000 (Function 4) 0x4_9000 (Function 5) 0x4_A000 (Function 6) 0x4_B000 (Function 7)	0x50	[9:8]	Selects which Legacy Interrupt to be used: <ul style="list-style-type: none"> 0 – INTA 1 – INTB 2 – INTC 3 – INTD

2.10.3. MSI Interrupt

The Lattice PCIe IP core supports 32 MSI interrupts with a feature of enabling and disabling the vector masking. The MSI request can be either 32-bit addressable Memory Write TLP or a 64-bit addressable Memory Write TLP. There are other two registers called as Mask Bits Register and Pending Bits Register. Since there is a support for 32 interrupts, the mask bit and pending register are 32-bit length, each bit represents the masking or pending status for each interrupt. The MSI-X capability structure values are programmed through the PCI Express configuration space register.

The address is taken from the Message Address and Message Upper Address fields of the MSI Capability Structure, while the payload is taken from the Message Data field.

The type of MSI TLP sent (32-bit addressable or 64-bit addressable) depends on the value of the Upper Address field in the MSI capability structure. By default, the MSI messages are sent as 32-bit addressable Memory Write TLPs. MSI messages use 64-bit addressable Memory Write TLPs only if the system software programs a non-zero value into the Upper Address register.

The message control register in the MSI capability Structure, disables and enables the various support in the MSI Interrupt.

Figure 2.14 and Figure 2.15 show the MSI Capability Structure variant.

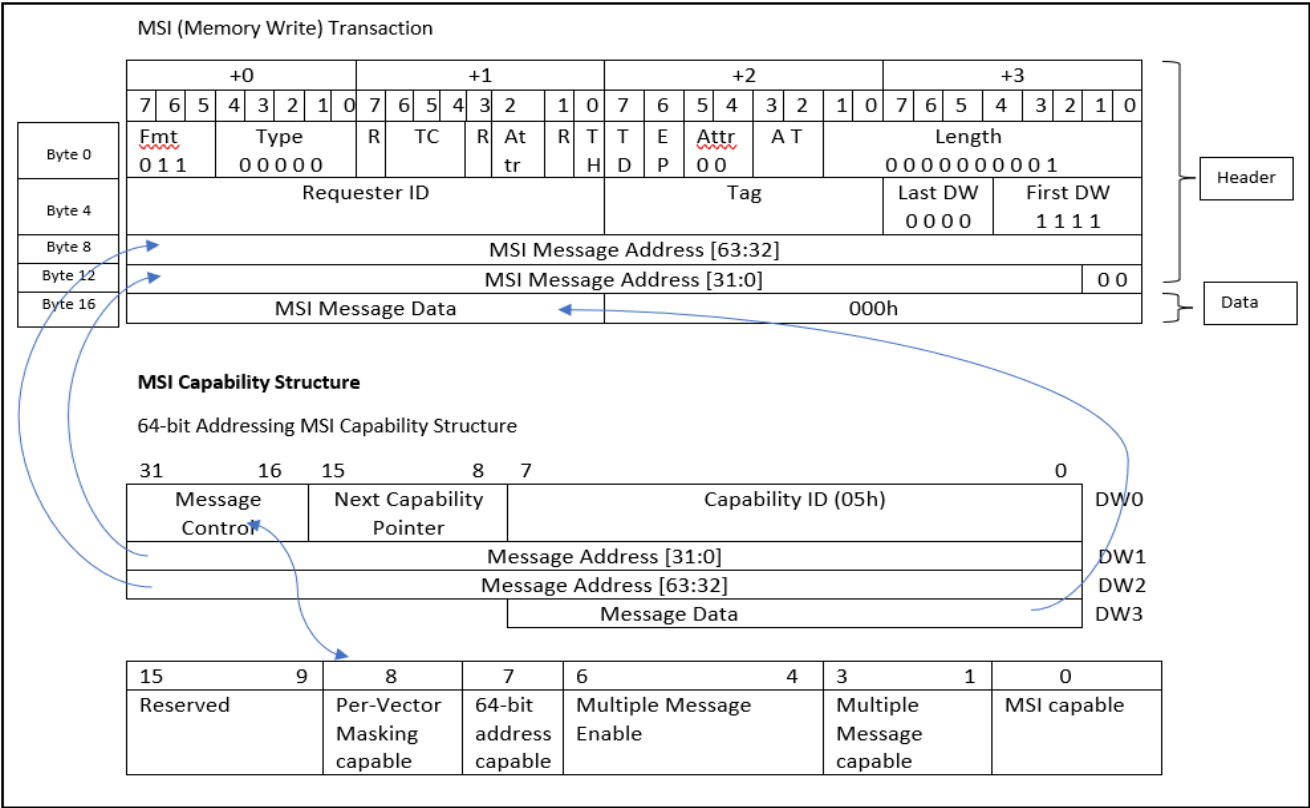


Figure 2.14. MSI Capability Structure Variant

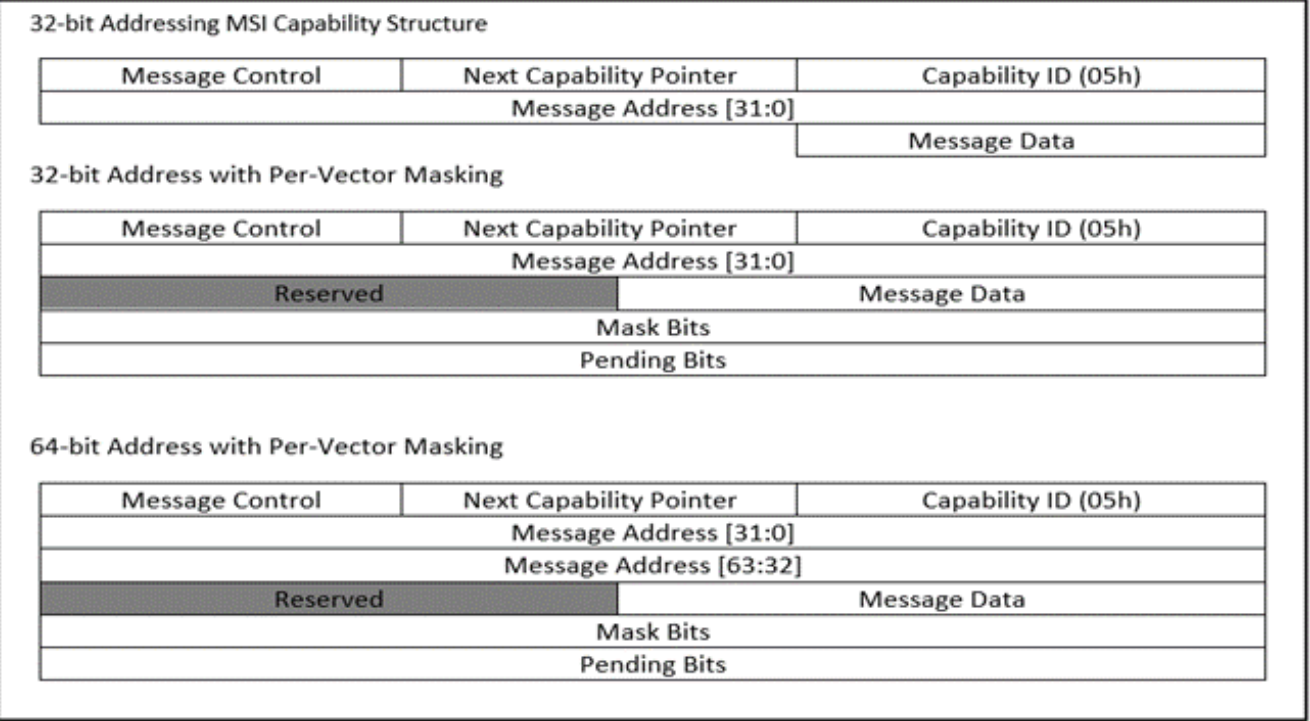


Figure 2.15. MSI Capability Structure Variant

2.10.3.1. MSI Pending Register

The MSI Pending register is used to report MSI Interrupts that are appended in the user design. MSI Pending is a PCIe Configuration Register in the MSI Capability Structure that software uses to obtain status on pending MSI Interrupt vectors. The MSI Pending register must be written whenever a MSI Interrupt Vector's pending status changes. A 1 must be written to the associated interrupt vector bit when an interrupt becomes pending and a 0 must be written to indicate that the interrupt is no longer pending.

The MSI Pending register must be updated whenever the status of your pending MSI interrupts changes. If MSI interrupts are not used, writing to the MSI Pending Register is not needed.

2.10.4. MSI-X Interrupt

2.10.4.1. MSI-X Capability Structure variant

MSI-X allows the support of large number of vectors with independent message data and address for each vector compared to the MSI Interrupts. It can support up-to 2048 vectors per function. The MSI-X Capability Structure points to an MSI-X table structure and an MSI-X Pending Bit Array (PBA) structure, which are stored in memory. In MSI-X interrupt the vector information is present in the memory location pointed by the Table Base address Indicator Register (BIR).

Figure 2.16 shows the MSI-X capability structure. The MSI-X interrupt configuration is done by the PCIe Configuration Space Registers.




Figure 2.16. MSI-X Capability Structure Variant

The description of each bit in the Message controlled are explained in the section of the PCIe Configuration Space Register configuration for MSI-X Capability Structure. The MSI-X Capability Structure variant contains information about the MSI-X Table and the PBA structure, information such as the pointers to the bases of the MSI-X Table and the PBA structure. The Table BIR in the MSI-X Capability Structure includes information about the BAR location that contains the MSI-X table.

2.10.4.2. MSI-X Table

The MSI-X table is an array of vectors and addresses. The MSI-X Table contains four Dwords. Each entry in the MSI-X table represents one vector. The DW0 and DW1 supply a unique 64-bit address for that vector and DW2 is the 32-bit data pattern for it. The DW3 is the mask bit for the vector and contains only 1 bit at a present.

DW3	DW2	DW1	DW0	
Vector Control	Message Data	Upper Address	Lower Address	Entry 0
Vector Control	Message Data	Upper Address	Lower Address	Entry 1
Vector Control	Message Data	Upper Address	Lower Address	Entry 2
....	
....	
Vector Control	Message Data	Upper Address	Lower Address	Entry N-1

Figure 2.17. MSI-X Table Entries

2.10.4.3. Pending Bit Array

The Pending Bit Array (PBA) is located within the memory address. This can use the same MSI-X Table BIR value (that is the same BAR or a different BAR). The PBA can use either qword (64-bit) or Dword (32-bit) accesses. The PBA table contains the pending bit information for each interrupt used. Same as MSI interrupts, if the event that the interrupt triggers and if its mask bit is set, the MSI-X transaction is not sent and the corresponding pending bit is set. If the interrupt vector is unmasked and if the pending bit is still set, that interrupt is generated.

DW1	DW0
Pending Bits 0 - 63	QW 0
Pending Bits 64 - 128	QW 1
Pending Bits 128 - 191	QW 2
.....	
Pending Bits	QW (N-1)/64

Figure 2.18. Pending Bit Array

2.10.4.4. MSI-X Interrupts Operation

- When the MSI-X interrupts are supported, you need to mention the size and location of the MSI-X Table and Pending Bit Array (PBA) through the PCIe CSR and the MSI-X table and the PBA structure must be implemented at the application layer.
- When the MSI-X Interrupts are generated, it uses the contents of the MSI-X Table (Address and Data) and generate a Memory Write through the TLP interface.
- The Host reads the message control register to determine the MSI-X Table size. The maximum entry in the table is 2048 entries. The BAR mentioned in the table BIR can access the MSI-X table.
- The host sets up the MSI-X table by programming the address, data, and the mask bits for each entry in the table.
- When the application generates the interrupt, it reads the MSI-X table information and generates a MWR TLP data and the corresponding bits in the PBA is set
- The generated TLP is sent to the corresponding address along with the data.
- When the MSI-X interrupt is sent, the application can clear the associated PBA bits.

2.11. PCIe Endpoint Core Buffers

The Lattice PCIe x8 IP Core contains three large RAM buffers:

- Transmit Buffer for transmitting TLPs.
- Receive Buffer for receiving TLPs.
- Replay Buffer for holding TLPs that were transmitted until positive acknowledgement of receipt is received.

The size of the Transmit Buffer, Receive Buffer, and Replay Buffer and the size of the corresponding buffers in the remote PCI Express Device have a fundamental impact on the throughput performance of the PCI Express link.

To achieve the highest throughputs, the buffers for both devices in the PCI Express link must be large enough that they can still accept more data while the oldest data begins to be freed from the buffer. If a buffer is too small, then the link stalls until the buffer has enough space to continue. The buffers must be large enough to overcome the expected latencies or the throughput is affected.

2.11.1. PCI Express Credits

The Flow Control DLLPs communicate the available buffer space in units of Header and Data Credits as defined in the PCI Express Specification. The amount of space required by a Header is 12-20 bytes or (3-5 DWORDs with 1 DWORD == 4 bytes). Each Header Credit represents the capability to store a maximum size packet header, which includes all the transaction control information (address, byte enables, and requester ID) and an optional End to End CRC (ECRC). Each Data Credit represents 16 bytes (4 DWORDs) of data payload. A transaction cannot be transmitted unless there is at least 1 header credit and enough data credits for the packet payload available in the remote device's Receive Buffer.

Credits are further divided into three categories for each of the main types of traffic:

- Posted (memory write requests and messages)
- Non-Posted (all reads, Configuration and I/O writes)
- Completion (responses to Non-Posted Requests) credit categories.

Each type of traffic must obey the PCI Express transaction ordering rules and is stored in its own buffer area. The Credit categories are annotated as:

- PH – Posted Request Header Credits
- PD – Posted Request Data Payload Credits
- NH – Non-Posted Request Header Credits
- ND – Non-Posted Request Data Payload Credits
- CH – Completion Header Credits
- CD – Completion Data Payload Credits

The PCI Express is inherently high-latency due to the serial nature of the protocol (clock rate matching and lane-lane de-skewing) and due to the latency induced by requiring packets to be fully received and robustly checked for errors before forwarding them for higher-level processing.

To achieve the best throughputs, both the Lattice PCIe x8 IP Core and the remote PCI Express device must be designed with a suitable number of credits and the capability to overlap transactions to bury the transaction latency.

The Lattice PCIe x8 IP Core Transmit, Receive, and Replay buffers are delivered with sufficient size to overcome the latencies of typical open system components.

2.11.2. Max Payload Size

The maximum payload size of any given packet is limited by the Max Payload Size field of the Device Control Configuration Register. The PCI Express Specification defines 128, 256, 512, 1024, 2048, and 4096-byte payload sizes. The maximum payload size that a device can support is limited by the size of its posted and completion TLP buffers. The Transmit Buffer and Receive Buffer Posted and Completion TLP storage and the Replay Buffer TLP storage needs to be able to hold at least four Max Payload Size TLPs to be reasonably efficient. Each device advertises the maximum payload size that it can support, and the OS/BIOS configures the devices in a link to use the lowest common maximum payload size. Thus, it is not advantageous to support a greater maximum payload size than the devices with, which one is communicating.

The higher the TLP payload size, the lower the TLP header and framing overhead is compared to the data. Above 512-byte Max Payload Size the incremental throughput benefit of higher payload sizes is small and the design area and latency for using these larger payloads is expensive. Thus, it is generally recommended to design for ≤ 512 Max Payload Size.

The Lattice PCIe x8 IP Core supports up to 512 Bytes Max Payload Size and the internal buffers can hold about 3x of the max payload size. However, given that typical PCIe devices currently available to communicate with support 256-byte maximum payloads, supporting greater than this amount is not likely to result in better performance and consumes more memory/logic resources.

2.12. Hard IP Interface

2.12.1. PHY Interface

The Link Layer is used in conjunction with a PCI Express PHY to implement a complete Lattice PCIe x8 IP Core PCI Express implementation. The PHY implements the high-speed serial and analog functions required to support PCI Express while the Link Layer implements most of the digital logic as well as the higher levels of the PCI Express protocol.

The PIPE PHY Interface that connects the Link Layer and PHY is not shown since the interface is only internal and is not visible to you.

The physical interface includes the differential receive and transmit signals along with the differential reference clock to the PCIe.

2.12.2. TLP TX/RX Interface

The Lattice PCIe core implements a complete PCI Express implementation including Physical, Data Link, and Transaction Layer functionality.

You transmit the PCI Express TLPs on the PCI Express link through the transmit interface. Also, you receive the PCI Express TLPs from the PCI Express link through the receive interface.

The PCIe core uses the Transaction Layer Interface as data interface to transmit/receive the data in the form of TLP Packets. Each TLP packet is a collection of a group of TLP frames, and each frame consists of 4DW (4X32 bit) data. A minimum of 4DW data is sent through a TLP. The Lattice PCIe core lane can access 32-bit (4 bytes) of data at a time. To transmit a single TLP frame, x1, x2, and x8 configuration takes a duration of 4, 2, and 1 clock cycles respectively.

All TLPs on the Transaction Layer Receive and Transmit Interfaces, which are processed through link0_rx_data_o/link0_tx_data_i port(s) and must be transmitted in the TLP format. The link0_rx_sel_o and link0_rx_cmd_data_o ports provide useful information about the TLP through receive interface to enable you to determine the destination of the packet (BAR and tag), traffic class, and whether it is a write or a read without having to read and parse the TLP. This allows you to optimize the code to reduce latency and relieves necessity for you to decode the TLP header to determine the packet's destination.

2.12.2.1. TLP Header Description

The Lattice PCIe uses 3DW header for memory transactions to transfer the data in the form of TLP packets. The description of each field is described as shown in [Figure 2.19](#).

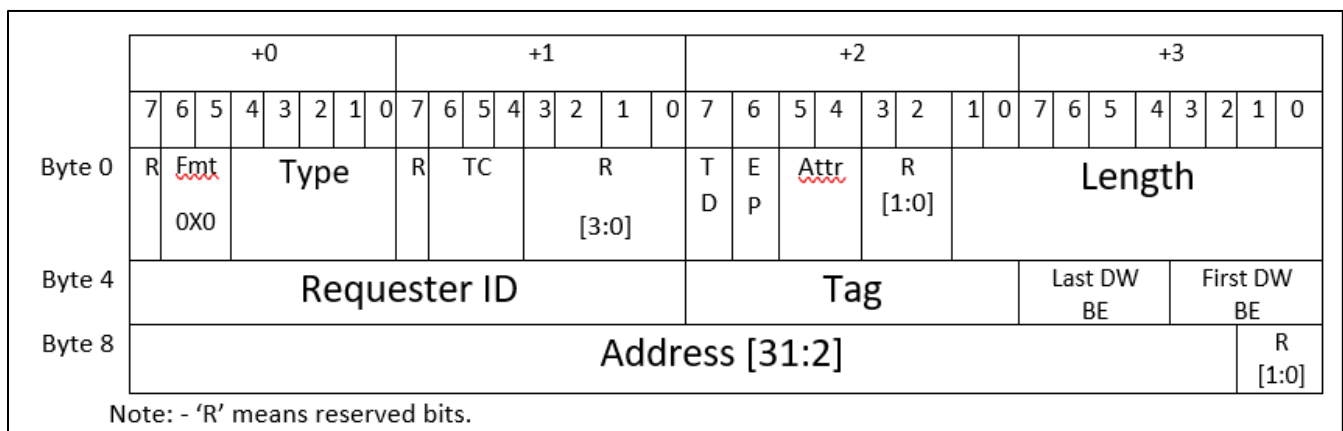


Figure 2.19. TLP Memory Request Header

Table 2.13 lists the description of each field.

Table 2.13. TLP Header Field

Field Name (with Size)	Header Byte/Bit	Function
Fmt [1:0] (Format)	Byte 0 Bit 6:5	Packet Formats: 00b = Memory Read 10b = Memory Write
Type [4:0]	Byte 0 Bit 4:0	TLP packet Type field: 00000b = Memory Read or Write 00001b = Memory Read Locked
TC [2:0] (Traffic Class)	Byte 1 Bit 6:4	These bits encode the traffic class to be applied to a Request and to any associated completion. 000b = Traffic Class 0 (Default)
TD (TLP Digest)	Byte 2 Bit 7	If 1, the optional TLP Digest field is included with this TLP.
EP (Poisoned Data)	Byte 2 Bit 6	If 1, the data accompanying this packet is considered to have an error although the transaction is allowed to complete normally.
Attr [1:0] (Attribute)	Byte 2 Bit 5:4	Bit 5 = Relaxed ordering. When set = 1, PCI-X relaxed ordering is enabled for this TLP. Otherwise, strict PCI ordering is used. Bit 4 = No Snoop. If 1, system hardware is not required to cause processor cache snoop for coherency for this TLP. Otherwise, cache snooping is required.
Length [9:0]	Byte 2 Bit 1:0 Byte 3 Bit 7:0	TLP data payload transfer size, in DW. Maximum size is 1024 DW (4 kB).
Requester ID [15:0]	Byte 4 Bit 7:0 Byte 5 Bit 7:0	Identifies a requester's return address for a completion: Byte 4, 7:0 = Bus Number Byte 5, 7:3 = Device Number Byte 5, 2:0 = Function Number
Tag [7:0]	Byte 6 Bit 7:0	These identify each outstanding request issued by the Requester. By default, only bit 4:0 is used, allowing up to 32 requests to be in progress at a time. If the Extended Tag bit in the Control Register is set, then all 8 bits may be used (256 tags).
Last DW BE [3:0] (Last DW Byte Enables)	Byte 7 Bit 7:4	These qualify bytes within the last DW of data transferred.
First DW BE [3:0] (First DW Byte Enables)	Byte 7 Bit 3:0	These qualify bytes within the first DW of the data payload.
Address [31:2]	Byte 8 Bit 7:0 Byte 9 Bit 7:0 Byte 10 Bit 7:0 Byte 11 Bit 7:2	The 32 bits start address for the memory transfer are used. The lower two bits of the address are reserved, forcing a DW-aligned start address.

2.12.2.2. TLP Transmit Interface

The Transmit Interface is the mechanism with which you transmit PCI Express TLPs over the PCI Express bus. You send a complete TLP comprised of 3DW packet header, data payload, and optionally a TLP Digest. The core Data Link Layer adds the necessary framing (STP/END/EDB), sequence number, Link CRC (LCRC), and optionally computes and appends the ECRC (TLP Digest) when ECRC is not already present in the TLP.

You transmit TLPs as completion packets in response to non-posted transaction packets sent by the Lattice PCIe IP core. If the remote device does not have sufficient space in the receive buffer for transmit TLPs, the Lattice PCIe IP core pauses the TLP transmission until space becomes available.

The Transmit Interface includes the option to nullify TLPs (instruct the Receiver to discard the TLP) to support you to cancel TLP transmissions when errors are detected after the TLP transmission has started. Nullified TLPs that target internal core resources (Root Port Configuration Registers and Power Management Messages) are discarded without affecting the internal core resources. Nullified TLPs that do not hit internal resources are discarded.

Transmit Credit Interface

The Transmit Credit Interface provides the means for flow control of non-posted transmit transactions between you and the core transmit buffer. This is important for allowing Posted and Completion TLPs to continue to make progress when non-posted TLPs are blocked (which can be necessary in some cases to avoid potential deadlock conditions). The amount of non-posted TLP storage in the core transmit buffer is communicated on the transmit credit Interface. You are expected to use this interface to limit simultaneously outstanding TLP transmission of non-posted TLPs, to the amount of non-posted TLPs that the core can absorb into the non-posted transmit buffer.

When the core Transaction Layer for Link[i] is ready to accept TLP transmissions, the core asserts `link0_tx_credit_init_o == 1` for one clock cycle and indicates the non-posted TLP Header storage capacity (NH) of the transmit buffer on `link0_tx_credit_nh_o[11:0]` on the same cycle. You are expected to keep and initialize the non-posted TLP Header capacity (NH) available transmit credit counters on `link0_tx_credit_init_o == 1`.

When a non-posted TLP is pending for transmission, you must check the currently available NH credit count for the associated link and hold the transmission until enough NH credits are available to transmit the TLP. Once the TLP is committed for transmit, the amount of NH credits required by the TLP are decremented from the NH credit count. The core forwards transmitted TLPs from the transmit buffer and thus makes room for new TLPs, the core asserts `link0_tx_credit_return_o == 1` for one clock cycle and places the number of NH credits being returned on `link0_tx_credit_nh_o[11:0]`.

In this manner, you can manage sending only enough non-posted TLPs that the core can hold in its Transmit Buffer. This allows you to know when non-posted TLPs are blocked and thus sends posted and/or completed TLPs instead. This is important for avoiding deadlocks and keeps non-posted TLP blockage from reducing posted and completion throughput. When the core receives more non-posted TLPs than the core can store in its non-posted TLP transmit storage, the core pauses the TLP transmission rather than allow an overflow to occur. Thus, if you do not wish to use the Transmit Credit Interface, you may ignore this interface provided you are willing to permit blocked non-posted TLPs from also blocking following posted and completion TLPs.

Note that core/link partner transmit TLP flow control is not managed through this interface, the core manages transmit flow control between the core and the PCIe link partner Receive Buffer without user intervention.

Transmit Interface Example Transactions

As mentioned, the TLP data interface option is made available when non-DMA support is enabled through PCIe user interface. The following are the examples of the memory read transactions that you need to send in completion to the read requests.

In case of memory read transactions, the Lattice PCIe IP core sends the header packet, which contains information about the address and size of data that you need to send in completion to the received packet. You need to send the completion packet with the header followed by the data when the `link0_tx_ready_o` signal from the PCIe is high as the data you sent is validated only when PCIe is ready. Based on the number of lanes used, the packet header and data are transmitted accordingly as shown in the below figures for four lanes, two lanes, and one lane respectively.

Note that the header packet has an unknown(trash) value in MSB, because the TLP header is of 3DW (12 bytes) whereas the TLP frame is of 4DW (32 bytes) size. To send the complete TLP, some garbage data in Dword is appended with header data, which can be ignored.

The following are the notations used in the figures:

- N – size of the data packet in Dwords
- *data* – 1 Dword of unknown data attached in case of 3DW TLP Header
- H0, H1, and H2 – Header information
- D0, D1,...,D(N-1),D(N) – User data




Figure 2.20. TLP Memory Read Operation for Link0 (x4 Lane)




Figure 2.21. TLP Memory Read Operation for Link0 (x2 Lane)

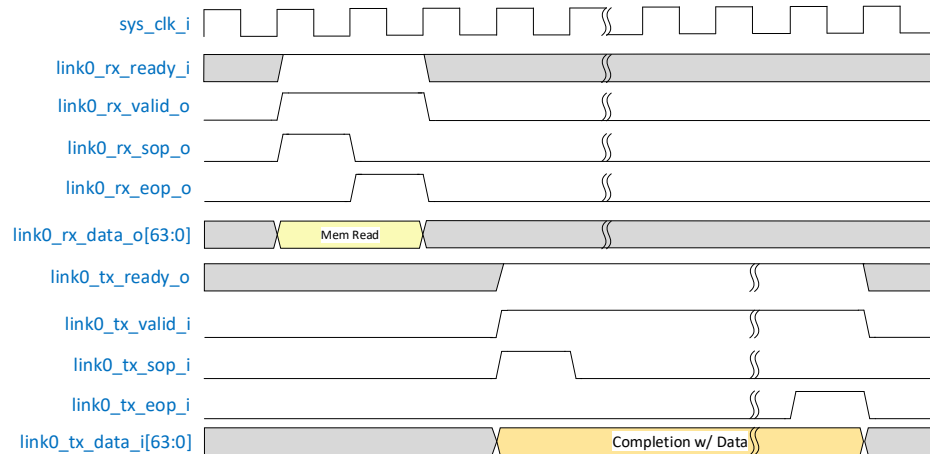


Figure 2.22. TLP Memory Read Operation for Link0 (x1 Lane)

Figure 2.23 and Figure 2.24 show the TLP transaction according to the tx_ready_o behaviour based on the minimum timing of link0_tx_ready_o:

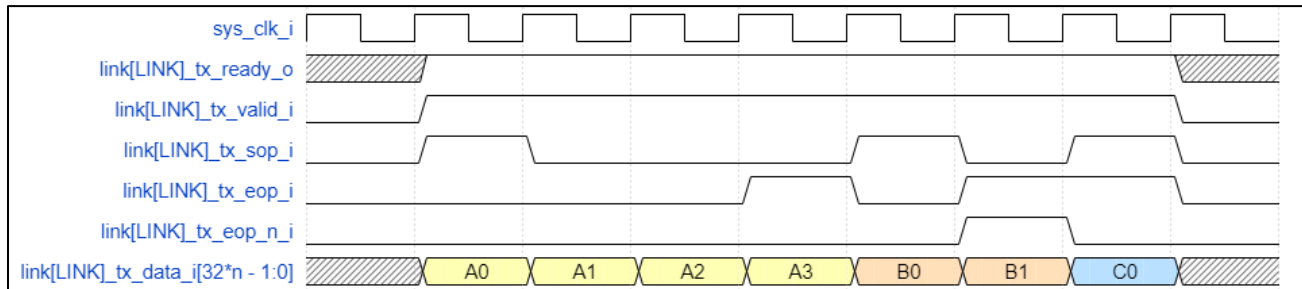


Figure 2.23. Minimum link0_tx_ready_o Timing Diagram

Transaction A begins on cycle 2 with the assertion of link0_tx_sop_i and ends on cycle 5 with the assertion of link0_tx_eop_i==link0_tx_valid_i==link0_tx_ready_o==1. The packet transfers with minimum timing with link0_tx_valid_i==link0_tx_ready_o==1 on cycles 2-5.

Transaction B begins immediately after Transaction A on cycle 6 with the assertion of link0_tx_sop_i and ends on cycle 7 with the assertion of link0_tx_eop_i==link0_tx_valid_i==link0_tx_ready_o==1. The packet transfers with minimum timing with link0_tx_valid_i==link0_tx_ready_o==1 on cycles 6 to 7. Transaction B is nullified (dropped) during packet forwarding on cycle 7 because of the following conditions: link0_tx_eop_i_n==1 happens when link0_tx_eop_i==1.

Transaction C begins immediately after Transaction B on cycle 8 with the assertion of link0_tx_sop_i and ends on the same cycle with the assertion of link0_tx_eop_i==link0_tx_valid_i==link0_tx_ready_o==1. The packet transfers with minimum timing with link0_tx_valid_i==link0_tx_ready_o==1 on cycle 8 considering the wait state timing of link0_tx_ready_o:

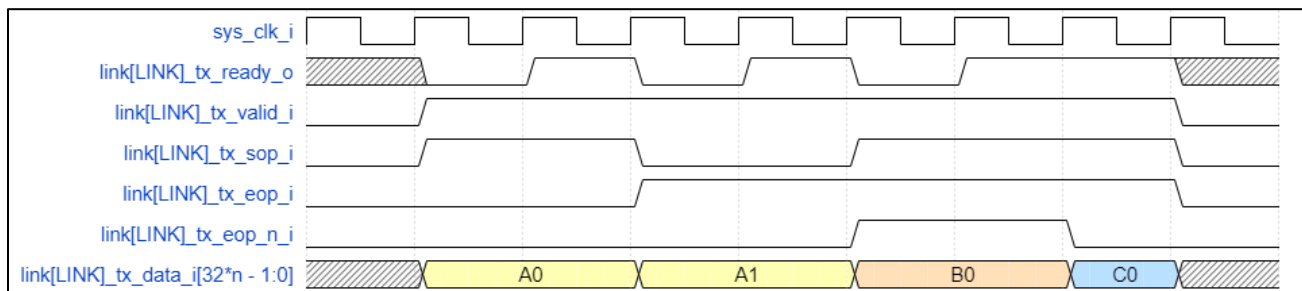


Figure 2.24. Wait State of link0_tx_ready_o Timing Diagram

Transaction A begins on cycle 2 with the assertion of link0_tx_sop_i and ends on cycle 5 with the assertion of link0_tx_eop_i==link0_tx_valid_i==link0_tx_ready_o==1. The packet transfers only on cycles 3 and 5 when link0_tx_valid_i==link0_tx_ready_o==1.

Transaction B begins immediately after Transaction A on cycle 6 with the assertion of link0_tx_sop_i and ends on cycle 7 with the assertion of link0_tx_eop_i==link0_tx_valid_i==link0_tx_ready_o==1. The packet transfers only on cycle 7 when link0_tx_valid_i==link0_tx_ready_o==1. Transaction B is nullified (dropped) during packet forwarding on cycle 7 because of the following conditions: link0_tx_eop_i_n==1 happens when link0_tx_eop_i==1.

Transaction C begins immediately after Transaction B on cycle 8 with the assertion of link0_tx_sop_i and ends on the same cycle with the assertion of link0_tx_eop_i==link0_tx_valid_i==link0_tx_ready_o==1. The packet transfers with minimum timing (no wait states) with link0_tx_valid_i==link0_tx_ready_o==1 on cycle 8.

Transmit Interface Considerations

The following considerations are provided to simplify logic using the Transmit Interface and to address common problems, which must be avoided:

- For each TLP that you transmit, the core adds a minimum of 2-bytes of STP/END/EDB framing, a 2-byte Sequence Number, and a 4-byte Link CRC for a total of 8 bytes (64-bits). These additional 8 bytes, which the core transmits but do not appear on link [LINK]_tx_data_i, allows you the flexibility of not using every clock cycle on the Transmit Interface. This flexibility can be useful to simplify user logic and improve design timing closure.
- Completions, which are transmitted in response to a previously received non-posted request, must reflect the Traffic Class, Requester ID, Tag, and Attributes of the original request. While most of these are obvious, it may not be obvious to reflect the attributes, and this is known to cause problems on some systems.
- When the link trains at less than full width or speed, link [LINK]_tx_ready_o is gaped in relation to the number of lanes being used and the number of lanes available in the core. You must remember to include a simulation case which forces the link into lower than full-width and/or speed to test that the logic properly handles the gaping of link [LINK]_tx_ready_o and the corresponding lower data transfer rate in this case.
- While TLPs are transmitted over PCI Express, these are placed into a replay buffer in case the TLPs need to be replayed due to transmission errors. The core negotiates the replay process in conjunction with the remote PCI Express Device and does not require any user intervention. You can monitor the frequency of replays, if desired, by monitoring the appropriate error status registers.
- The Lattice PCIe core interface is designed to support high throughput applications. Small interruptions in transmissions occurs, however, as the core periodically needs to transmit link management DLLPs and SKP Ordered Sets and may also need to transmit error messages, configuration write/read completions, and interrupt TLPs.
- The Lattice PCIe core handles all Data Link Layer functionality for you and handles most of the error cases for you as well. To accomplish these functions, the core occasionally delays your Local Transmit Interface requests while it completes its own TLP transmissions for these purposes. All the core's TLP transmissions are short, so it delays your request for only a few clock cycles. The core transmits DLLPs used for link maintenance, TLP messages to communicate errors, interrupt TLPs, and completions to notify the system of malformed or un-routable TLP requests.

If the user TLP transmit requests are delayed for extended periods of time, this may be due to insufficient link partner receive buffer space or local replay buffer space or due to the link having to wake from a lower power state or recover from an error before transmission can occur.

2.12.2.3. TLP Receive Interface

The Receive Interface is the mechanism with, which receives the PCI Express TLPs from the PCI Express link partner. You receive complete Transaction Layer Packets (TLPs) comprised of a three DWORD TLP header, data payload (if present), and TLP Digest (ECRC, if present).

The TLPs, which were received without errors and were not nullified, are presented on the receive interface. Therefore, the user logic only needs to handle valid received TLPs.

The PCIe core transmits the TLPs only after considering the following checks:

- The core checks received TLPs for transmission errors (Sequence Number or LCRC error) and negotiates replay of TLPs with the link partner as required.
- The core discards TLPs which are nullified by the link partner during transmission (TLP is received without transmission errors and with EDB instead of END framing).
- The core checks received TLPs which were received without transmission errors and without being nullified for Malformed TLP due to length and content errors.
 - If the core determines that a received TLP is malformed due to length (TLP length calculated from the received TLP Header Format and Type, Length, and TLP Digest does not match the received TLP length), the core discards the TLP and report the error.
 - If the TLP fails to hit an enabled resource or is malformed due to its content (invalid Traffic Class, invalid Format and Type, and invalid Byte Enables), the core discards the TLP and reports the error.




Figure 2.25. TLP Packet formation by the Lattice PCIe IP core

If the TLP passes all the above checks, it is considered a valid TLP and is forwarded to the receive interface for the user's logic to consume. The core strips the Physical Layer framing (STP/END/EDB) and Data Link Layer Sequence Number and Link CRC (LCRC) before presenting the TLPs to you on the Receive Interface. The core does not strip the received TLP ECRC (if present) as some user designs require forwarding the ECRC either to transmit the TLP out another PCIe port. The ECRC value is also checked at a later point in the user's data path in order to continue the ECRC error detection protection for a larger portion of the receive data path. If an ECRC is present in the TLP, the core checks the validity of the received ECRC and reports detected ECRC errors on the receive interface.

The core also decodes received TLPs against its Configuration Registers and provides the transaction decode information on the Receive Interface such that the TLP can be directed to the appropriate destination without the need for you to parse the TLP until its destination. For example, if the received TLP is an I/O or Memory write or read request, the Base Address Register (BAR) resource that is hit is indicated and if the TLP is a completion, the TLP's tag field is provided. The core also provides additional useful transaction attributes.

Receive Credit Interface

The Receive Credit Interface provides the means for flow control of non-posted receive transactions between the core receive buffer and user receive TLP logic. This is important for allowing Posted and Completion TLPs to continue to make progress when non-posted TLPs are blocked (which is necessary in some cases to avoid potential deadlock conditions). The amount of non-posted TLP storage in the user's design is communicated on the Receive Credit Interface. The core uses this interface to limit the simultaneously outstanding receive non-posted TLPs to the amount of non-posted TLPs that the user design advertises that it can absorb into the non-posted receive buffer.

When you are ready to accept non-posted TLP reception, assert the `link0_rx_credit_init_i == 1` (where `LINK=0` or `1`) for one clock cycle and the non-posted TLP header storage capacity of the user design is indicated through `link0_rx_credit_nh_i[11:0]` on the same clock cycle. Holding off credit initialization for an extended period can cause received non-posted TLP transactions to timeout in the source component which may cause to serious errors.

The core limits simultaneous outstanding non-posted receive TLPs on the receive interface to ensure no more than the initialized NH credits are simultaneously outstanding to user receive TLP logic.

Once the received non-posted TLPs are processed/forwarded such that more room is available to receive new non-posted TLPs, assert the `link0_rx_credit_return_i==1` for one clock cycle and place the number of NH credits being returned on `link0_rx_credit_nh_i[11:0]`. In this manner, you can limit the outstanding core receive TLPs to the user design. This permits the core to know when non-posted TLPs are blocked and thus send posted and/or completion TLPs to the user design instead. This is important for avoiding deadlocks and keeps non-posted TLP blockage from reducing posted and completion throughput.

Note that the link partner/core receive TLP flow control is not managed through this interface; the core manages receive buffer flow control between itself and the PCIe link partner transmit gating function without user intervention.

Receive Interface Example Transactions

The Lattice PCIe core sends the data in the form of TLP packets when non-DMA option is enabled. The receive interface presents the TLP data through `link0_rx_data_o` signal. The data is validated only when `link0_rx_valid_o` signal is high and you are ready (for example, `link0_rx_ready_i` must be high to access the data sent by the core). As the Lattice PCIe core transmits TLP packet, which consists of 3DW header along with data (in TLP frames), the last DW of TLP packet is sent as trash value(X) to ensure the complete TLP is transmitted.

The timing diagrams below show the receive interface behavior when the PCIe core receives a Memory Write TLP.

The following are the notations used in the figures:

- *N* – size of the data packet in Dwords
- *data* – 1 Dword of unknown data attached in case of completion of TLP packet
- H0,H1, and H2 – Header information
- D0,D1,...,D(N-1) – Write data




Figure 2.26. TLP Memory Write Operation for Link0 (x4 Lane)




Figure 2.27. TLP Memory Write Operation for Link0 (x2 Lane)




Figure 2.28. TLP Memory Write Operation for Link0 (x1 Lane)

Figure 2.28 and Figure 2.29 shows the TLP transaction according to the rx_ready_i behaviour based on the minimum timing of link0_rx_ready_i:




Figure 2.29. Minimum link0_rx_ready_i Timing Diagram

Transaction A begins on cycle 2 with the assertion of link0_rx_sop_o==link0_rx_valid_o==1 and ends on cycle 5 with the assertion of link0_rx_eop_o==link0_rx_valid_o==link0_rx_ready_i==1. The packet transfers with minimum timing since link0_rx_valid_o == link0_rx_ready_i == 1 on cycles 2-5. Data transfers on cycles 2-5.

Transaction B begins immediately after Transaction A on cycle 6 with the assertion of link0_rx_sop_o==link0_rx_valid_o==1 and ends on cycle 7 with the assertion of link0_rx_eop_o==link0_rx_valid_o==link0_rx_ready_i==1. Data transfers on cycles 6-7.

Transaction C begins immediately after Transaction B on cycle 8 with the assertion of link0_rx_sop_o==link0_rx_valid_o==1 and ends on the same cycle since link0_rx_eop_o==link0_rx_ready_i==1 are also asserted. Data transfers on cycle 8 considering the wait state timing of link0_rx_ready_i:




Figure 2.30. Wait State of link0_rx_ready_i Timing Diagram

Transaction D begins on cycle 2 with the assertion of link0_rx_sop_o==link0_rx_valid_o==1 and ends on cycle 7 with the assertion of link0_rx_eop_o==link0_rx_valid_o==link0_rx_ready_i==1. The packet transfer wait states due to link0_rx_ready_i==0 on cycles 2, 4, and 6. Data transfers on cycles 3, 5, and 7.

Transaction E begins on cycle 8 with the assertion of link0_rx_sop_o==link0_rx_valid_o==1 but is wait stated due to link0_rx_ready_i==1 on cycle 8. On cycle 9 the transaction completes with link0_rx_sop_o==link0_rx_eop_o==link0_rx_valid_o==link0_rx_ready_i==1.

Receive Interface Considerations

The following considerations are provided to simplify logic using the receive interface and to address common problems, which must be avoided:

- For each TLP that you receive, the core strips a minimum 2-bytes of STP/END/EDB framing, a 2-bytes of Sequence Number, and a 4-bytes of Link CRC, for a total of 8 bytes (64-bits). These additional 8 bytes, which the core receives but which do not appear on link0_rx_data_o, allows you the flexibility of not using every clock cycle on the receive interface. This flexibility can be useful to simplify user logic and improve design timing closure.
- TLPs that appear on the receive interface have passed the Physical Layer and Link Layer error detection and correction logic and can be assumed to be free of transmission errors. When the core receives a TLP with a STP/END/EDB framing, Sequence Number, or Link CRC error, the core coordinates re-transmission of the TLP with the remote PCI Express device and only forwards packets that pass transmission error checks onto to the receive interface.

- TLPs that are received from PCI Express are decoded for validity against the core's configuration registers and are only forwarded to the receive interface if they hit an enabled resource. Therefore, you only need to handle valid TLPs which target the user resources. TLPs, which do not hit user resources, are terminated by the core and the appropriate error message and response is handled by the core on the user's behalf.
- The Lattice PCIe core handles all Data Link Layer functionality for you and handles most of the Transaction Layer error cases as well. The core consumes Configuration Transactions, Messages, and TLPs which do not map to user resources and transmits the appropriate response. TLPs which are handled by the core do not appear on the Receive Interface.
- User logic that manages read requests (for DMA) and assigns a tag to each read request that is transmitted. The core provides the tag of each received completion on link0_rx_cmd_data_o to allow user logic to route completions from different sources to the destination without having to parse the TLP for tag information. The core does not track the outstanding tags that are in use by the user. If a completion is received with a tag that does not correspond to an outstanding user read request, then you must report the error.

Data Byte Order

The core transmits the TLP data in the following byte order:

- link0_tx_data_i[7:0], link0_tx_data_i[15:8], link0_tx_data_i[23:16],...

The core receives the TLP data in the following byte order:

- link0_rx_data_o[7:0], link0_rx_data_o[15:8], link0_rx_data_o[23:16],...

For example, you transmit, or the core receives a 32-bit Memory Read Transaction Layer Packet in the following byte order as shown in [Table 2.14](#).

Table 2.14. Data Byte Order

link0_tx_data_i/ link0_rx_data_o	First Data Word	Second Data Word	Third Data Word
[7:0]	{R, Fmt[1:0], Type[4:0]}	RequesterID[15:8]	Addr[31:24]
[15:8]	{R, TC[2:0], R[3:0]}	RequesterID[7:0]	Addr[23:16]
[23:16]	{TD, EP, Attr[1:0], Length[9:8]}	Tag[7:0]	Addr[15:8]
[31:24]	Length[7:0]	{LastDWBE[3:0], 1stDWBE[3:0]}	{Addr[7:2], R[1:0]}

2.12.2.4. Transaction Layer Interface Error Detection and Correction

The Lattice PCIe IP Core has built in error detection and correction mechanisms for both Transaction Layer Packets (TLPs) which are transferred between PCI Express and Transaction Layer Interface and Data Link Layer Packets (DLLPs) which are used by the core internally for link management.

The Lattice PCIe IP core adds the required Physical Layer framing (STP/END/EDB) and Data Link Layer error detection and correction information (Sequence Number/Link CRC) to the TLP packets transmitted on the Transmit Interface. Likewise, when TLP packets are received from PCI Express, the core validates that the packet is received correctly by checking the Physical Layer framing (STP/END/EDB) and Data Link Layer error detection and correction information (Sequence Number/Link CRC). Packets that are forwarded to you on the receive interface are sent after stripping the Physical Layer framing (STP/END/EDB) and Data Link Layer error detection and correction information (Sequence Number/Link CRC).

If transmission errors are detected in packet transmission or reception, the core coordinates with the remote PCI Express device to retry the transaction and recover from the error. This process occurs without any user intervention. The Lattice PCIe core logs both corrected and uncorrected errors. This error status information is made available through the status registers and is accessed by system software through the Configuration Registers. The core generates and transmits error Message TLPs to the remote PCI Express device in response to different types of errors detected.

ECRC (TLP Digest) generation and checking is a core option. When ECRC generation support is enabled by the software (AER Capability: ECRC Generation Enable == 1), the core generates and adds ECRC to all transmitted TLPs (except those that already contain an ECRC with TD bit set to 1). When ECRC checking support is enabled by software (AER Capability: ECRC Check Enable == 1), the ECRC fields present in received TLPs are checked for validity and any errors are noted on the Receive Interface and are reported in the AER Capability. The core does not modify the ECRC or TD (TLP Digest == ECRC indicator) fields on received TLPs and passes these fields onto the receive interface as received.

The Lattice PCIe core also handles TLPs that are Type 0 Configuration transaction requests, messages requests for link management, TLPs that don't hit an enabled resource and any requests that the core determines are malformed.

If the core found TLP having transmission errors, then that TLP is consumed by the core (and not forwarded) and then transmits any required completion packet(s), generates required error messages and logs any required errors.

The core has been designed in such a way that it is feasible to you to only consume and generate the TLPs and can make use of these TLPs for transferring data and control information between your application and the remote PCI Express devices.

2.12.3. LMMI Interface

When you select the TLP as data interface option in the PCIe IP user interface, the IP by default configures LMMI as register interface. The Core Configuration and Status Registers (CSR) are made accessible to the user design through the Lattice Memory Mapped interface (LMMI).

An example of the register configuration through the LMMI is shown below in the LMMI write and read timing diagrams.

The data transaction, through the LMMI, only starts when both `usr_lmimi_request_i` and `usr_lmimi_ready_o` are asserted high. Consecutive request must be done with at least one clock period wait cycle (for example, `usr_lmimi_request_i` should de-assert first after a successful transaction before making another request).

When both `usr_lmimi_request_i` and `usr_lmimi_ready_o` are asserted high, `usr_lmimi_wr_rdn_i`, and `usr_lmimi_offset_i` must be valid and describe the transaction to execute; if the transaction is a write as indicated by `usr_lmimi_wr_rdn_i` being asserted to high, `usr_lmimi_wdata_i` must also be valid.

Notes:

- Only one request should be active at a given time.
- Once `usr_lmimi_rst_n_i` de-asserted, you must wait 10 or more `usr_lmimi_clk_i` clocks before accessing LMMI interface. This is to allow reset to be properly propagated through the PCIe controller.

2.12.3.1. LMMI Write Operations

The data is written to PCIe registers only when PCIe IP gets a request from you by the asserting `usr_lmimi_request_i` signal to high, when the PCIe IP indicates that it is ready by asserting `usr_lmimi_ready_o` to high and when a write transaction is being indicated from your end by driving `usr_lmimi_wr_rdn_i` to high. When all three of these signals are asserted, a successful write transaction will take place.

For example, you need to write 0X01 data into 0X0A register and then 0X02 data into 0X0B register. The 0X01 data is written into 0X0A register in one transaction only as ready signal is high when request is asserted. But to write 0X02 data into 0X0B register took two transactions because ready signal is low when request is asserted in first transaction. Therefore, the data is written to the register in the second transaction only when ready signal is high as shown in [Figure 2.31](#).

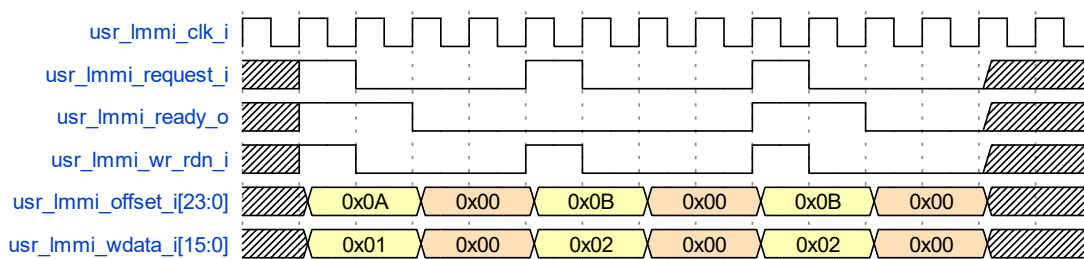


Figure 2.31. LMMI Write Operation

2.12.3.2. LMMI Read Operation

The data is read from PCIe registers only when PCIe IP gets a request from the user. For example, `usr_lmimi_request_i` must be asserted high to indicate a new request is being made and `usr_lmimi_wr_rdn_i` signal must be asserted low to indicate a read transaction. A valid offset value through `usr_lmimi_offset_i` needs to be supplied from your end during the request.

You can read the data from PCIe core registers only when the ready and read valid signals are received from PCIe IP. For example, the `usr_lmimi_ready_o` signal and the `usr_lmimi_rdata_valid_o` signal must be asserted high together for a valid read transaction to take place.

For example, you want to read the data [0X0000001d00000000] from lane 0 PMA Status register offset 0x7F. The transaction follows the steps as shown in Figure 2.32.




Figure 2.32. LMMI Read Operation

The following are the registers to be configured through the LMMI:

- Simulation Registers
 - Register address – 0x4_2000
This register is used to reduce the ltssm ts_1 and timeouts to fasten the simulation when asserted as 1.
 - Register address – 0x4_3000
This register is used to reduce the Power Management State Machine timeouts to fasten the simulation when asserted as 1.
 - Register address – 0x4_4000
This register is used to reduce the timeouts to fasten the simulation when asserted as 1.

2.13. Soft IP Interface

2.13.1. Data Interface Conversion

2.13.1.1. AXI-4 Stream Interface

This interface is available if the data interface type selected in the IP generation user interface is *AXI4_STREAM*. The data width of the transaction varies based on the lane support.

- Lane x4: 256 data width
- Lane x2: 128 data width
- Lane x1: 64 data width

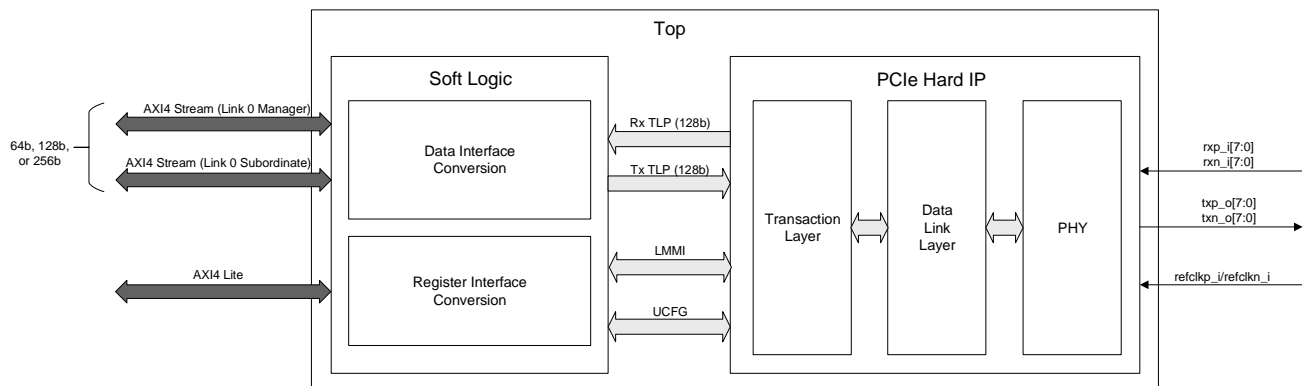


Figure 2.33. AXI4-Stream Data Interface, AXI-L Register Interface

PCIe to AXI4-Stream Transfer

For the PCIe to AXI4-Stream transfer, the PCIe sends the data to the user's application. The transaction has the header of size three double-word. After the header transaction, the actual data is transferred as shown in Figure 2.34 to Figure 2.36.

Note: The *DATA* in the transactions below are random data sent by the PCIe IP to compensate the data width.

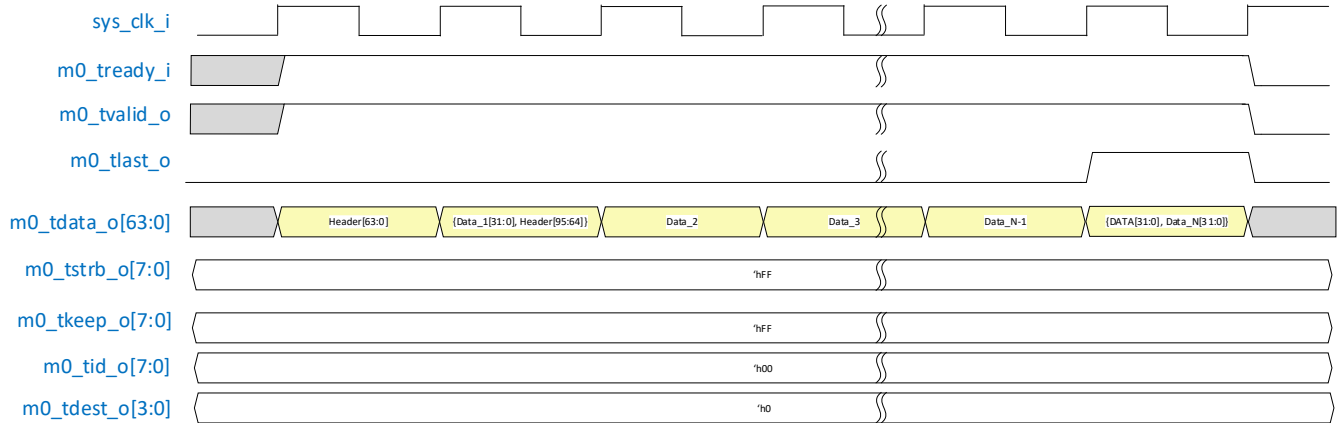


Figure 2.34. PCIe to AXI4-Stream Transaction for x1

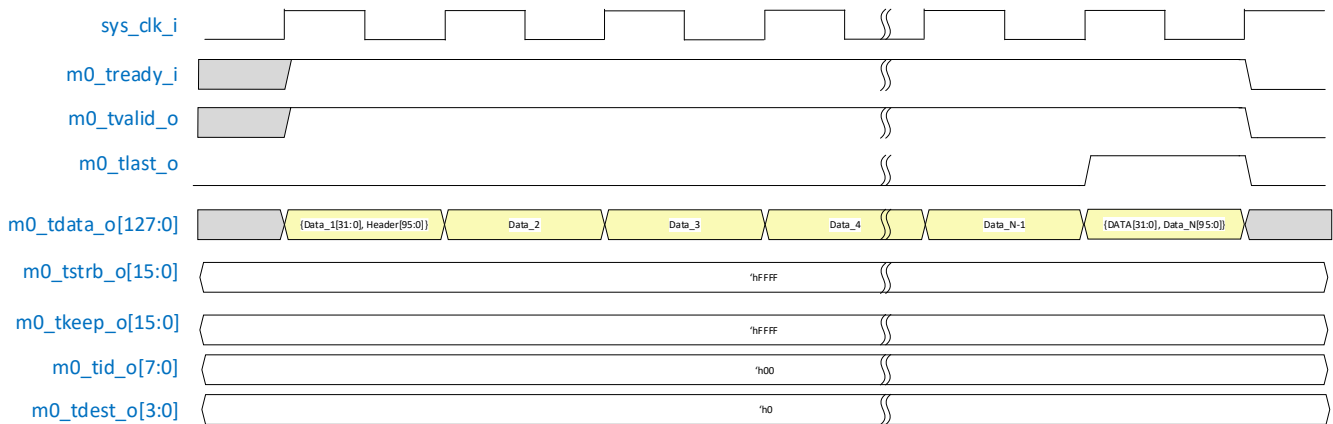


Figure 2.35. PCIe to AXI4-Stream Transaction for x2

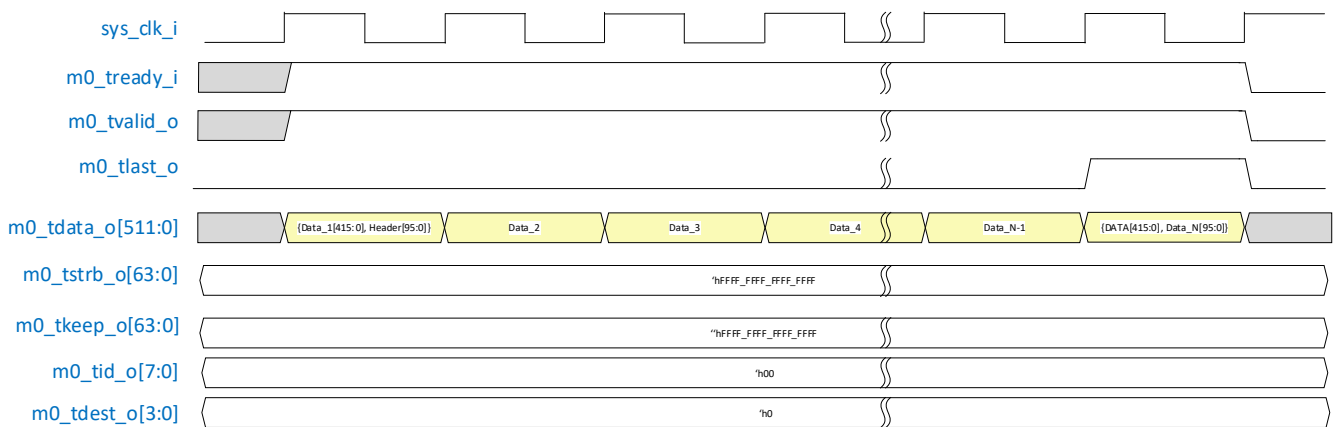


Figure 2.36. PCIe to AXI4-Stream Transaction for x8

AXI4-Stream to PCIe

For the AXI4-Stream to PCIe transaction, the user's application sends the data to the PCIe Endpoint IP. Similar to the PCIe to AXI4-Stream transfer, there is a header data of size three double word, which is transferred first followed by the actual data as shown in Figure 2.37 to Figure 2.39.

Note: The *DATA* in the transactions below are random data sent by the PCIe IP to compensate the data width.

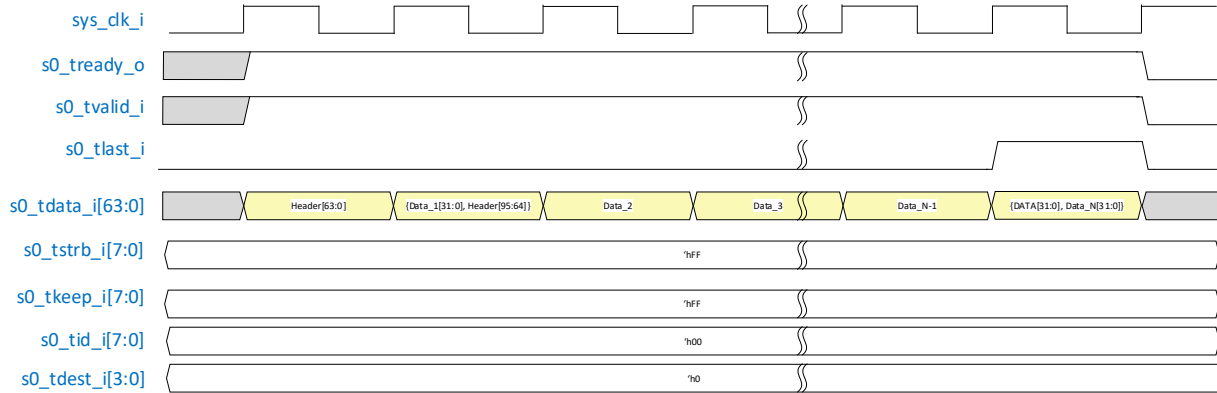


Figure 2.37. AXI4-Stream to PCIe Transaction for x1

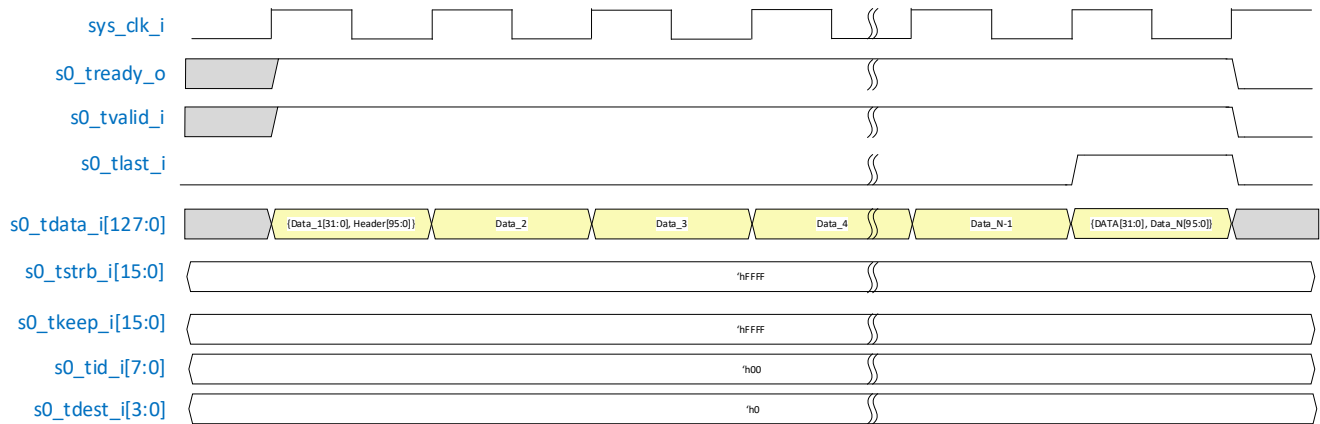


Figure 2.38. AXI4-Stream to PCIe Transaction for x2

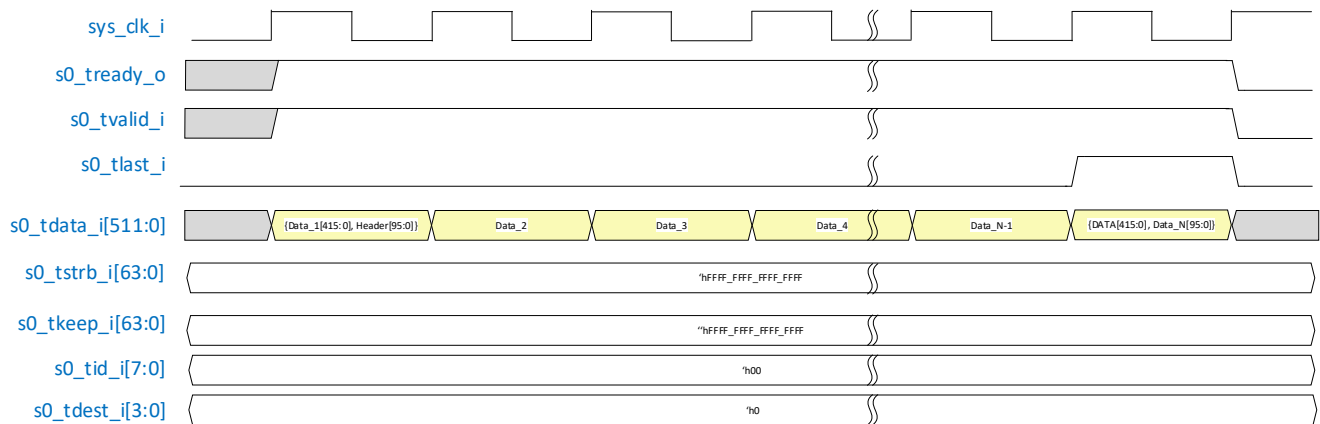


Figure 2.39. AXI4-Stream to PCIe Transaction for x8

2.13.2. Register Interface Conversion

2.13.2.1. AXI-L Interface

This interface is available if the register interface type selected in the IP generation user interface is AXI-L. For more information on the AXI-L signal description, refer to the [AXI-L Configuration Interface](#) section. The address of the registers to be accessed is available in the [Register Description](#) section.

The Core Configuration and Status Registers (CSR) are made accessible to the user design through the AXI-L. The registers that are configured through AXI-L are as follows:

- Simulation registers:
 - Register address – 0X4_2000
This register is used to reduce the Itssm ts_1 and timeouts to fasten the simulation when asserted as 1.
 - Register address – 0X4_3000
This register is used to reduce the Power Management State Machine timeouts to fasten the simulation when asserted as 1.
 - Register address – 0X4_4000
This register is used to reduce the timeouts to fasten the simulation when asserted as 1.

2.14. Multi-Protocol Support

The Lattice Avant-AT-G supports multi-protocol starting Radiant 2024.1. Quad merging is handled by the Lattice Radiant software.

PCIe Bifurcation									
Mode	PHY Lanes								
	0	1	2	3	4	5	6	7	
DMA+PCIe Mode	x8								
PCIe Mode 1x8	x8								
PCIe Mode 1x4	x4				any protocol				
PCIe Mode 1x2	x2		NA	NA	any protocol				
PCIe Mode 1x1	x1	NA	NA	NA	any protocol				
	Quad0				Quad1				
PCIe Mode 1x2 + PIPE Direct (2x1/1x2)	x2		(PIPE x1/x2)						
PCIe Mode 1x1 + PIPE Direct (2x1/1x2)	x1	NA	(PIPE x1/x2)						

Figure 2.40. PCIe Bifurcation

3. IP Parameter Description

The PCIe Endpoint Core attributes are configurable through the IP Catalog's Module/IP wizard of the Lattice Radiant Software. Refer to [Table 3.1](#) for the description of each attribute.

3.1. General

General		Flow Control	Link 0: Function 0
Property			Value
▼ General			
Bifurcation Select (Link_X_Lane)			1X8
Ref Clk Freq (MHz)			100
Simulation Reduce Timeout			0
Enabling LTSSM Polling Function			<input type="checkbox"/>
Enabling Register Autoread Function			<input type="checkbox"/>
Register Interface Type			AXI4_LITE
AXI4-Lite Data Width			32
Link 0 : Data Width			1024
Link 0 : PCIe Device Type			PCIe Endpoint
Target Link Speed			GEN4
Link 0 : Number of Physical Functions			1
AXI Harden DMA Enabled			<input type="checkbox"/>
Data Interface Type			AXI4_STREAM
▼ ASPM Capability			
Link 0 : Active State Power Management (ASPM) Support			No ASPM Support

Figure 3.1. Attributes in the General Tab

Table 3.1. General Tab Attributes Description

Attribute	Selectable Values	Description	Parameter
Bifurcation Select	1×1 1×2 1×4 1×8	Configures the number of Links and Lanes.	PCIE_BIFUR_SEL 2=1×1, 1=1×2, 5=1×4, 6=1×8
Ref Clk Freq (MHz)	100 MHz	Display only. Ref Clk from PCIe slot	—
Simulation Reduce Timeout	0, 1	0 – For hardware mode 1 – For simulation mode	—
Register Interface Type	AXI-Lite, LMMI	<ul style="list-style-type: none"> Available if default interface is not selected and mapping of register to subordinate data interface is not selected. The selected interface replaces the native Lattice Memory Mapped Interface (LMMI) of the hard IP by adding a soft logic bridge. 	USR_CFG_IF_TYPE = {"LMMI", "AXI-L"}

Attribute	Selectable Values	Description	Parameter
LMMI Data Width	32	Display only. Lattice Memory Mapped Interface (LMMI) data width	—
Link 0 Data Width	64	Display only	—
Link 0 PCIe Device Type	PCIe Endpoint	Display only. PCIe IP core supports only PCIe Endpoints	LINK0_DEVICE_TYPE = "PCIe Endpoint"
Link 0 Target Link Speed	Gen1 (2.5G) Gen2 (5.0G) Gen3 (8.0G) (2024) Gen4 (16.0G) (2024)	<ul style="list-style-type: none"> Initial value of Target Link Speed Configuration Register. Determines the maximum initial link speed which can be reached during initial training. Must be set to the lesser of the maximum speed supported by the core and the maximum speed at which you desire the core to operate. 	LINK0_FTL_INITIAL_TARGET_LINK_SPEED = {0,1,2,3}
Link 0 Number of Physical Functions	1–8	Set the number of enabled functions.	LINK0_NUM_FUNCTIONS = {1..8}
Enable DMA Support (2024)	Checked Unchecked	Enable Harden DMA	—
Data Interface Type	AXI4_STREAM, TLP	<ul style="list-style-type: none"> Available if default interface is not selected. The selected interface replaces the native TLP data interface of the hard IP by adding a soft logic bridge. 	USR_MST_IF_TYPE = {"TLP", "AXI4_STREAM"} USR_SLV_IF_TYPE = {"TLP", "AXI4_STREAM", "NONE"}

3.2. Optional Port



Figure 3.2. Attributes in the Optional Port Tab

Table 3.2. Optional Port Attributes

Link [k] (k == 0 - 1) Optional Ports			
Attribute	Selectable Values	Description	Parameter
Link [k] Enable Legacy Interrupt Ports	Checked Unchecked	<ul style="list-style-type: none"> Available if legacy interrupt is enabled. Set to add the Legacy interrupt ports. 	LINK[k]_MAIN_CTRL_4_EN_PORT_MGMT_INTERRUPT_LEG = {0,1}

3.3. Flow Control Update

Configure IP	
General	Flow Control
Property	Value
▼ Flow Control Update	
Disable FC Update Timer	<input type="checkbox"/>
FC Update Timer Divider	Use PCIe Spec recommended values
Completion Credit (CH,CD) Advertisement	Advertise [Infinite for Endpoint], [Actual values for Root Port]

Figure 3.3. Attributes in the Flow Control Update Tab

Table 3.3. Flow Control Attributes

Link [k] (k == 0 - 1) Flow Control Update			
Attribute	Selectable Values	Description	Parameter
Link [k] Disable FC Update Timer	Checked Unchecked	<ul style="list-style-type: none"> Set to disable FC Update Timer (that is, schedule a FC Update on Every Consumed RX TLP Otherwise, schedule FC Updates in accordance with PCIe Specification recommended values) 	LINK[k]_PTL_RX_CTRL_FC_UPDATE_TIMER_DISABLE = {0,1}
Link [k] FC Update Timer Divider	Use PCIe Spec recommended values, Divide by 2, Divide by 4, Divide by 8	Select the FC Update frequency of the Receive Buffer when FC update timer is enabled.	LINK[k]_PTL_RX_CTRL_FC_UPDATE_TIMER_DIV = {0,1,2,3}
Link [k] Completion Credit Advertisement	Advertise Infinite for Endpoint and Actual for Root Port, Advertise Actual, Advertise Infinite	Select the completion credit advertisement behavior.	LINK[k]_PTL_RX_CTRL_ADV_CH_CD_SEL = {0,1,2}

3.4. Function 0

3.4.1. Configuration

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	4
Property				Value			
▼ Configuration							
Link 0 : Disable Function 0				<input type="checkbox"/>			
Link 0 : Device ID (16'h)				E004			
Link 0 : Vendor ID (16'h)				19AA			
Link 0 : Subsystem ID (16'h)				E004			
Link 0 : Subsystem Vendor ID (16'h)				19AA			
Link 0 : Class Code (24'h)				118000			
Link 0 : Revision ID (8'h)				04			
Link 0 : Root Port ID (16'h)				0000			

Figure 3.4. Attributes in Function Configuration Tab

Table 3.4. Function Configuration Tab Attributes

Configuration			
Attribute	Selectable Values	Description	Parameter
Disable Function	Unchecked	<ul style="list-style-type: none"> Cannot disable function 0. Display only. 	—
Link [k] Device ID	(Hex) 0000 – FFFF	Value returned when the Device ID Configuration Register is read.	LINK[k]_FTL_ID1_DEVICE_ID = {16'h0000 – 16'hFFFF}
Link [k] Vendor ID	(Hex) 0000 – FFFF	Value returned when the Vendor ID Configuration Register is read.	LINK[k]_FTL_ID1_VENDOR_ID = {16'h0000 – 16'hFFFF}
Link [k] Subsystem ID	(Hex) 0000 – FFFF	Value returned when the Subsystem ID Configuration Register is read.	LINK[k]_FTL_ID2_SUBSYSTEM_ID = {16'h0000 – 16'hFFFF}
Link [k] Subsystem Vendor ID	(Hex) 0000 – FFFF	Value returned when the Subsystem Vendor ID Configuration Register is read.	LINK[k]_FTL_ID2_SUBSYSTEM_VENDOR_ID = {16'h0000 – 16'hFFFF}
Link [k] Class Code	(Hex) 00000 – FFFFF	Value returned when the Class Code Configuration Register is read.	LINK[k]_FTL_ID3_CLASS_CODE = {16'h0000 – 16'hFFFF}
Link [k] Revision ID	(Hex) 00 – FF	Value returned when the Revision ID Configuration Register is read.	LINK[k]_FTL_ID3_REVISION_ID = {8'h00 – 8'hFF}

3.4.2. Base Address Register (BAR) [0 to 5]

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	Link 1: Function 1
Property				Value			
▼ Base Address Register 0							
Link 0 : BAR 0 : Enable				<input type="checkbox"/>			
Link 0 : BAR 0 : Resizable				<input type="checkbox"/>			
Link 0 : BAR 0 : Address Type				Memory			
Link 0 : BAR 0 : 64 bit address				<input type="checkbox"/>			
Link 0 : BAR 0 : Prefetchable				<input type="checkbox"/>			
Link 0 : BAR 0 : Resizable BAR Supported Sizes [23:4] (20'h)				00000			
Link 0 : BAR 0 : Default Size (unit)				KiB (2 ¹⁰)			
Link 0 : BAR 0 : Default Size (value)				64			
Link 0 : BAR 0				32'h0			
Link 0 : Local Memory Base Address 0				0			
▼ Base Address Register 1							
Link 0 : BAR 1 : Enable				<input type="checkbox"/>			
Link 0 : BAR 1 : Resizable				<input type="checkbox"/>			
Link 0 : BAR 1 : Address Type				Memory			
Link 0 : BAR 1 : 64 bit address				<input type="checkbox"/>			
Link 0 : BAR 1 : Prefetchable				<input type="checkbox"/>			

Figure 3.5. Attributes in BAR Tab

Table 3.5. BAR Tab Attributes

Link [k] (k == 0 - 1) Base Address Register n (n == 0 - 5)			
Attribute	Selectable Values	Description	Parameter
Link [k] BAR n – Enable	Checked Unchecked	Set to enable the BAR.	—
Link [k] BAR n – Address Type	Memory, I/O	Select if the BAR is for Memory or I/O space.	—
Link [k] BAR n – 64-bit Address	Checked Unchecked	<ul style="list-style-type: none"> Applicable for memory space only. Set to use 64-bit address. Note that BAR n and BAR n+1 are used for the 64-bit address. 	—
Link [k] BAR n – Prefetchable	Checked Unchecked	<ul style="list-style-type: none"> Applicable for memory space only. Set to identify the memory address as prefetchable. 	—
Link [k] BAR n – Default Size (unit)	Bytes, kB (210), MB (220), GB (230), TB (240), PB (250), EB (260),	Select the size of Memory space. ¹	—

Link [k] (k == 0 - 1) Base Address Register n (n == 0 - 5)			
Attribute	Selectable Values	Description	Parameter
Link [k] BAR n – Default Size (value)	(Power of 2) 32 bits Memory Space: 16 bytes – 2 GB 64 bits Memory Space: 64 bits: 4 GB – 8 EB 32 bits I/O Space: 2 Bytes – 256 Bytes	Select the size of Memory or I/O space. ¹	—
Link [k] BAR n	32 bits: FFFF_FFF0 - 1000_0000 64 bits: FFFF_FFFF_0000_0000 - 1000_0000_0000_0000	Display Only	Function 0: LINK[k]_FTL_BAR0_CFG LINK[k]_FTL_BAR5_CFG Function m: LINK[k]_FTL_MF1_BAR0_CFG LINK[k]_FTL_MF[m]_BAR[n]_CFG
Local Memory Base Address n	(Hex, Aligned to BAR size) FFFF_FFF0 – 0000_0000	<ul style="list-style-type: none"> Applicable for memory space only. This is the base address of the local system memory that maps to the configured PCIe BAR. Must be aligned to the specified BAR size. Received Memory requests that hits the BAR are forwarded to this address. 	Function 0: F0BAR0_TO_LOCADR F0BAR5_TO_LOCADR Function m: F1BAR0_TO_LOCADR F[m]BAR[n]_TO_LOCADR

3.4.3. Legacy Interrupt

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	4
Property				Value			
▼ Legacy Interrupt							
Link 0 : Disable Legacy Interrupt				<input type="checkbox"/>			
Link 0 : Interrupt Pin				INTA			

Figure 3.6. Attributes in Legacy Interrupt

Table 3.6. Legacy Interrupt Attribute Descriptions

Link [k] (k == 0 - 1) Legacy Interrupt			
Attributes	Value	Description	Parameters
Link [k] Disable Legacy Interrupt	Checked Unchecked	<ul style="list-style-type: none"> RTL always supports legacy interrupt. The current attribute only uses for port activation. 	LINK[k]_FTL_INTERRUPT_DISABLE = {0,1}
Link [k] Interrupt Pin	INT A, INT B, INT C, INT D	Select which legacy interrupt pin is used.	LINK[k]_FTL_INTERRUPT_PIN = {0,1,2,3}

3.4.4. MSI Capability

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	◀
Property			Value				
▼ MSI Capability							
Link 0 : Disable MSI Capability			<input type="checkbox"/>				
Link 0 : Number of MSI vectors			8				
Link 0 : Enable Vector Masking			<input checked="" type="checkbox"/>				

Figure 3.7. Attributes in MSI Capability

Table 3.7. MSI Capability Attributes

Link [k] (k == 0 - 1) MSI Capability			
Attributes	Value	Description	Parameters
Link [k] Disable MSI Capability	Checked Unchecked	Set to disable the MSI Capability.	LINK[k]_FTL_MSI_CAP_DISABLE = {0,1}
Link [k] Number of MSI vectors	1 – 32	Set the number of requested MSI vectors.	LINK[k]_FTL_MSI_CAP_MULT_MESSAGE_CAPABLE = {0,1,2,3,4,5}
Link [k] Enable Vector Masking	Checked Unchecked	Set to enable vector masking capability.	LINK[k]_FTL_MSI_CAP_VEC_MASK_CAPABLE = {0,1}

3.4.5. MSI-X Capability

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	4
Property			Value				
▼ MSI-X Capability							
Link 0 : Disable MSI-X Capability			<input type="checkbox"/>				
Link 0 : MSI-X Table Size [1 - 2048]			8				
Link 0 : MSI-X Table BAR indicator			BAR 0				
Link 0 : MSI-X Table Address Offset (8bytes aligned)			6000				
Link 0 : MSI-X PBA BAR indicator			BAR 0				
Link 0 : MSI-X PBA Address Offset (8bytes aligned)			7000				

Figure 3.8. Attributes in MSI-X Capability

Table 3.8. MSI-X Capability Attributes

Link [k] (k == 0 - 1) MSI-X Capability			
Attributes	Value	Description	Parameters
Link [k] Disable MSI-X Capability	Checked Unchecked	Set to disable the MSI-X Capability.	LINK[k]_FTL_MSIX_CAP_DISABLE = {0,1}
Link [k] MSI-X Table Size	1–2048	Set the number of requested MSI-X vectors.	LINK[k]_FTL_MSIX_CAP_TABLE_SIZE = {0 – 2047}
Link [k] MSI-X Table BAR indicator	BAR 0, BAR 1, BAR 2, BAR 3, BAR 4, BAR 5	<ul style="list-style-type: none"> Select which Base Address register. Located beginning at 10h in Configuration Space, is used to map the MSI-X Table into Memory Space. 	LINK[k]_FTL_MSIX_TABLE_BIR = {0,1,2,3,4,5}
Link [k] MSI-X Table Address Offset	(Hex, 8 bytes aligned) 0000_0000 – FFFF_FFF8	Set the byte address offset (8 bytes aligned), within the BAR selected by MSI-X Table BAR indicator, at which the MSI-X table begins.	LINK[k]_FTL_MSIX_TABLE_OFFSET = {29'h00000000 – 29'h1FFFFFFF}
Link [k] MSI-X PBA BAR indicator	BAR 0, BAR 1, BAR 2, BAR 3, BAR 4, BAR 5	Select which Base Address register, located beginning at 10h in Configuration Space, is used to map the MSI-X PBA into Memory Space.	LINK[k]_FTL_MSIX_PBA_BIR = {0,1,2,3,4,5}
Link [k] MSI-X PBA Address Offset	(Hex, 8 bytes aligned) 0000_0000 – FFFF_FFF8	Set the byte address offset (8 bytes aligned), within the BAR selected by MSI-X PBA BAR indicator, at which the MSI-X PBA begins.	LINK[k]_FTL_MSIX_PBA_OFFSET = {29'h00000000 – 29'h1FFFFFFF}

3.4.6. Device Serial Number Capability

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	4
Property		Value					
▼ Device Serial Number Capability							
Link 0 : Enable DSN Capability		<input type="checkbox"/>					
Link 0 : Serial Number		0					

Figure 3.9. Attributes in Device Serial Number Capability

Table 3.9. Device Serial Number Capability Attributes

Link [k] (k == 0 - 1) Device Serial Number Capability			
Attributes	Value	Description	Parameters
Link [k] Enable DSN Capability	Checked Unchecked	Set to enable the Device Serial Number capability.	LINK[k]_FTL_DSN_CAP_ENABLE = {0,1}
Link [k] Serial Number	(Hex) 0000_0000_0000_0000 – FFFF_FFFF_FFFF_FFFF	Set the device serial number.	LINK[k]_FTL_DSN_SERIAL_NUMBER

3.4.7. PCI Express Capability

General		Flow Control	Link 0: Function 0
Property			Value
▼ PCI Express Capability			
Link 0 : Maximum Payload Size Supported			512_BYTES
Link 0 : Disable Function Level Reset (FLR)			<input checked="" type="checkbox"/>
Link 0 : Enable Extended Tag Field			<input type="checkbox"/>

Figure 3.10. Attributes in PCIe Capability

Table 3.10. PCIe Capability Attributes

Link [k] (k == 0 - 1) PCIe Device Capability			
Attributes	Value	Description	Parameters
Link [k] Maximum Payload Size Supported	128 Bytes, 256 Bytes, 512 Bytes	Select the maximum payload size supported.	LINK[k]_FTL_PCIE_DEV_CAP_MAX_PAYLOAD_SIZE_SUPPORTED = {0,1,2}
Link [k] Disable Function Level Reset (FLR)	Checked Unchecked	Set to disable Function Level Reset capability.	LINK[k]_FTL_PCIE_DEV_CAP_DISABLE_FLR_CAPABILITY = {0,1}
Link [k] Enable Extended Tag Field	Checked Unchecked	Set to enable Extended Tag Field (8-bit tag field).	LINK[k]_FTL_PCIE_DEV_CAP_EXTENDED_TAG_FIELD_SUPPORTED = {0,1}

3.4.8. Advance Error Reporting Capability

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	⌵
Property			Value				
▼ Advance Error Reporting Capability							
Link 0 : Enable ECRC Generation and Checking			<input checked="" type="checkbox"/>				
Link 0 : Enable Reporting : Correctable Internal Error			<input type="checkbox"/>				
Link 0 : Enable Reporting : Surprise Down Error			<input type="checkbox"/>				
Link 0 : Enable Reporting : Completion Timeout Error			<input checked="" type="checkbox"/>				
Link 0 : Enable Reporting : Completer Abort Error			<input type="checkbox"/>				
Link 0 : Enable Reporting : Uncorrectable Internal Error			<input type="checkbox"/>				

Figure 3.11. Attributes in Advance Error Reporting Capability

Table 3.11. Advance Error Reporting Capability Attributes

Link [k] (k == 0 - 1) Advance Error Reporting Capability			
Attributes	Value	Description	Parameters
Link [k] Enable ECRC Generation and Checking	Checked Unchecked	Set to enable ECRC generation and checking.	LINK[k]_FTL_AER_CAP_ECRC_GEN_CHK_CAPABLE = {0,1}
Link [k] Enable Reporting: Correctable Internal Error	Checked Unchecked	Set to enable reporting of correctable internal error.	LINK[k]_FTL_AER_CAP_EN_CORR_INTERNAL_ERROR = {0,1}
Link [k] Enable Reporting: Surprise Down Error	Checked Unchecked	Set to enable reporting of surprise down error.	LINK[k]_FTL_AER_CAP_EN_SURPRISE_DOWN_ERROR = {0,1}
Link [k] Enable Reporting: Completion Timeout Error	Checked Unchecked	Set to enable reporting of completion timeout error.	LINK[k]_FTL_AER_CAP_EN_COMPLETION_TIMEOUT = {0,1}
Link [k] Enable Reporting: Completer Abort Error	Checked Unchecked	Set to enable reporting of completer abort error.	LINK[k]_FTL_AER_CAP_EN_COMPLETER_ABORT = {0,1}
Link [k] Enable Reporting: Uncorrectable Internal Error	Checked Unchecked	Set to enable reporting of uncorrectable internal error.	LINK[k]_FTL_AER_CAP_EN_UCORR_INTERNAL_ERROR = {0,1}

3.4.9. Atomic OP Capability

General	Flow Control	Link 0: Function 0	Link 0: Function 1	Link 0: Function 2	Link 0: Function 3	Link 1: Function 0	4
Property				Value			
Atomic OP Capability							
Link 0 : Enable Atomic Op Capability				<input type="checkbox"/>			
Link 0 : Enable Root as Atomic Op Completer				<input type="checkbox"/>			
Link 0 : Enable Atomic Op Completer 128b Operand				<input checked="" type="checkbox"/>			
Link 0 : Enable Atomic Op Completer 64b Operand				<input checked="" type="checkbox"/>			
Link 0 : Enable Atomic Op Completer 32b Operand				<input checked="" type="checkbox"/>			
Link 0 : Enable Atomic Op Routing				<input type="checkbox"/>			

Figure 3.12. Attributes in Atomic OP Capability

Table 3.12. Atomic OP capability Attributes

Link [k] (k == 0 - 1) Atomic OP Capability			
Attributes	Value	Description	Parameters
Link [k] Enable Atomic Op Capability	Checked Unchecked	Set to enable Atomic Operations Capability.	LINK[k]_FTL_ATOMIC_OP_CAP_ENABLE = {0,1}
Link [k] Enable Root as Atomic Op Completer	Checked Unchecked	Set to enable Root as Atomic OP Completer.	LINK[k]_FTL_ATOMIC_OP_CAP_RP_COMPLETER_ENABLE = {0,1}
Link [k] Enable Atomic Op Completer 128b Operand	Checked Unchecked	Set to support Atomic Op 128b operand.	LINK[k]_FTL_ATOMIC_OP_CAP_COMPLETER_128_SUPPORTED = {0,1}
Link [k] Enable Atomic Op Completer 64b Operand	Checked Unchecked	Set to support Atomic Op 64b operand	LINK[k]_FTL_ATOMIC_OP_CAP_COMPLETER_64_SUPPORTED = {0,1}

Table 3.15. Dynamic Allocation capability Attributes

Link [k] (k == 0 - 1) Dynamic Power Allocation Capability			
Attributes	Value	Description	Parameters
Link [k] Enable DPA Capability	Checked Unchecked	Set to enable the Dynamic Power Allocation capability.	LINK[k]_FTL_DPA_CAP_ENABLE = {0,1}
Link [k] Max Substate Number	0 – 31	<ul style="list-style-type: none"> Specifies the maximum substate number. Substates from [substate_max:0] are supported. For example, substate_max==0 indicates support for 1 substate. 	LINK[k]_FTL_DPA_CAP_SUBSTATE_MAX = {0 - 31}
Link [k] Transition Latency Unit	1 ms, 10 ms, 100 ms	Specifies Transition Latency Unit.	LINK[k]_FTL_DPA_CAP_TLUNIT = {0 - 2}
Link [k] Power Allocation Scale	10.0x, 1.0x, 0.1x, 0.01x	Specifies Power Allocation Scale.	LINK[k]_FTL_DPA_CAP_PAS = {0 - 3}
Link [k] Transition Latency Value 0	0 – 255	Specifies Transition Latency Value 0.	LINK[k]_FTL_DPA_CAP_XLCY0 = {0 - 3}
Link [k] Transition Latency Value 1	0 – 255	Specifies Transition Latency Value 1.	LINK[k]_FTL_DPA_CAP_XLCY1 = {0 - 3}
Link [k] Transition Latency Indicator 32x1b	(Hex) 00000000 – FFFFFFFF	<ul style="list-style-type: none"> Specifies which Transition Latency Value applies to each substate. Each bit corresponds to a substate. 	LINK[k]_FTL_DPA_XLCY_INDICATOR = {32'h00000000 – 32'hFFFFFFFF}
Link [k] Power Allocation Array 32x8b	(Hex) {32{00}} – {32{FF}}	<ul style="list-style-type: none"> Substate Power Allocation Array. Each entry is 8b value. 	LINK[k]_FTL_DPA_ALLOC_ARRAY = { {32{8'h00}} – {32{8'hFF}} }

Table 3.16. Function 1-3 Tab

Link [k] (k == 0 - 1) Function n (n == 1 – 3)			
Configuration			
Disable Function	Checked	Available if the number of physical functions enabled is set to greater than 1. Set to disable the function. Parameter: LINK[k]_FTL_MF1_FUNCTION_DISABLE = {0,1} LINK[k]_FTL_MF2_FUNCTION_DISABLE = {0,1} LINK[k]_FTL_MF3_FUNCTION_DISABLE = {0,1}	Parameter: LINK[k]_FTL_MF1_FUNCTION_DISABLE = {0,1} LINK[k]_FTL_MF2_FUNCTION_DISABLE = {0,1} LINK[k]_FTL_MF3_FUNCTION_DISABLE = {0,1}
	Unchecked		
Device ID	Refer to Function 0 section.		—
Vendor ID			
Subsystem ID			
Subsystem Vendor ID			
Class Code			
Revision ID			
Base Address Register (see the Lattice PCIe x8 Core Configuration user interface in Function 0 section)			—
Legacy Interrupt (see the Lattice PCIe x8 Core Configuration user interface in Function 0 section)			—
MSI Capability (see the Lattice PCIe x8 Core Configuration user interface in Function 0 section)			—
MSI-X Capability (see the Lattice PCIe x8 Core Configuration user interface in Function 0 section)			—
Device Serial Number Capability (see the Lattice PCIe x8 Core Configuration user interface in Function 0 section)			—

4. Signal Description

The Lattice PCIe x8 IP Core Ports for Link 0 is defined in the following sub-sections.

4.1. Clock and Reset Interface

Table 4.1. Clock and Reset Ports

Port	Type	Description
link0_perst_n_i	Input	<ul style="list-style-type: none"> PCI Express Fundamental Reset Active-low asynchronous assert, synchronous de-assert reset to the Link Layer, PHY, and Soft Logic blocks. On link0_perst_n_i and link0_rst_usr_n_i deassertion, the core starts in the Detect.Quiet Link Training and Status State Machine (LTSSM) state with the Physical Layer down and Data Link Layer down. Normally, the PHY registers are automatically configured based on the GUI settings selected by User. However, it is still possible for the user to modify and configure the registers through the selected register interface. link0_perst_n_i must remain asserted while the PHY registers are being configured.
link0_rst_usr_n_i	Input	<ul style="list-style-type: none"> User Clock Domain Link Layer Reset (Link Layer or MAC Layer Reset) Active-low asynchronous assert, synchronous de-assert reset to the User clock domain, Link Layer and Soft Logic blocks. On link0_perst_n_i and link0_rst_usr_n_i deassertion the core starts in the Detect.Quiet Link Training and Status State Machine (LTSSM) state with the Physical Layer down and Data Link Layer down. In general, the Link Layer registers are automatically configured based on the user interface settings you selected. However, it is still possible for the user to modify and configure the registers through the selected register interface. It is recommended that link0_rst_usr_n_i remain asserted while the Link Layer core registers are being configured.
sys_clk_i	Input	<ul style="list-style-type: none"> User Clock Domain Input Clock It is recommended to use the following minimum clock frequency to achieve the maximum throughput with respect to link data rate: <ul style="list-style-type: none"> 16.0G – 250 MHz² 8.0G – 125 MHz 5.0G – 62.5 MHz 2.5G – 31.25 MHz All ports of the corresponding PCIE link are synchronized to this clock.
link0_clk_usr_o	Output	<ul style="list-style-type: none"> Optional Clock source. There are several clock sources coming from PHY that can be selected and connected to the sys_clk_i.
link0_pl_link_up_o	Output	<ul style="list-style-type: none"> Physical Layer Link Up Status – (1) Up and (0) Down link0_pl_link_up_o is used as an active-low, synchronous reset for the core's Data Link Layer. You are not expected to use this port except for status since user design does not interface directly with the Data Link Layer.

Port	Type	Description
link0_dl_link_up_o	Output	<ul style="list-style-type: none"> Data Link Layer Link Up Status – (1) Up and (0) Down link0_dl_link_up_o is used as an active-low, synchronous reset for the Transaction Layer and also indicates when TLPs can be successfully transmitted across the link. For Endpoint-only applications, users must use link0_dl_link_up_o as a synchronous reset for their RTL interfacing to the core's Transaction Layer interfaces.
link0_tl_link_up_o	Input	<ul style="list-style-type: none"> Transaction Layer Link Up Status – (1) Up and (0) Down link0_tl_up_o is an active-low, synchronous reset to the core's upper transaction layer. A Downstream Port's (Root Port) Configuration Registers, and the RTL which enables them to be accessed them from the User Transmit and Receive Interfaces, are reset by link0_tl_up_o so that the PCIe Configuration Registers may be read even when the link is down. Downstream Port applications must use link0_tl_up_o as a synchronous reset for their RTL interfacing to the core's Transaction Layer interfaces. When link0_tl_up_o == 1 the core's PCIe Configuration Registers may be read. When link0_dl_link_up_o == 1, accesses the PCIe devices behind the DS Port may be attempted. When link0_dl_link_up_o == 0 accesses to the PCIe devices behind the DS Port is not possible; if accesses are attempted, they fail with error status.
sys_rst_n_i	Input	<ul style="list-style-type: none"> System Reset (Available if application interface selected is AXI4) This is similar to link0_perst_n_i, but it is common to all links. Active-low asynchronous assert, synchronous de-assert reset to the Link Layer, PHY, and DMA Core blocks.

Notes:

1. LINK – range (0, 3)
2. Support for 250 MHz is only available in the next release (2024.1).

4.2. PHY Interface

The Link Layer is used in conjunction with a third-party PCI Express PHY to implement a complete Lattice PCIe x8 Core PCI Express implementation. The PHY implements the high-speed serial and analog functions required to support PCI Express while the Link Layer implements most of the digital logic as well as the higher levels of the PCI Express protocol.

The PIPE PHY Interface that connects the Link Layer and PHY is not shown here since the interface is only internal and is not visible to users.

Table 4.2. PHY Interface Descriptions

Port	Type	Description
link0_rxp_i	Input	<ul style="list-style-type: none"> Differential Receive Serial signal, Rx+. When PCIE Core is only enabled, the PCIE link can be configured as: <ul style="list-style-type: none"> x8 lanes – (Lane 0 – 7) or x4 lanes – (Lane 0 – 3) or x2 lanes – (Lane 0 – 1) or x1 lane – (Lane 0)
link0_rxn_i	Input	Differential Receive Serial signal, Rx- See PCIE Hard IP Core Lane assignment above.
link0_txp_o	Output	Differential Transmit Serial signal, Tx+ See PCIE Hard IP Core Lane assignment above.

Port	Type	Description
link0_txn_o	Output	Differential Transmit Serial signal, Tx- See PCIe Hard IP Core Lane assignment above.
refclkp_i	Input	Differential Reference Clock, CLK+ (100 MHz)
refclkn_i	Input	Differential Reference Clock, CLK- (100 MHz)

Note:

1. LANE – range (0, 7)

4.3. Transaction Layer Interface

4.3.1. TLP Transmit Interface

Refer the [TLP Transmit Interface](#) section for more information and timing diagrams.

4.3.1.1. TLP Transmit Interface Port Description

Table 4.3. TLP Transmit Interface Ports

Port	Clock Domain	Direction	Description
link0_tx_valid_i	sys_clk_i	Input	Source valid. (1==Valid, 0==Not valid)
link0_tx_ready_o	sys_clk_i	Output	<ul style="list-style-type: none"> • Destination ready. (1==Ready, 0==Not ready) • A transfer occurs on the transmit interface only when link0_tx_valid_i==link0_tx_ready_o==1.
link0_tx_sop_i	sys_clk_i	Input	<ul style="list-style-type: none"> • Start of packet indicator. • Set == 1 coincident with the first link0_tx_data_i word in each TLP.
link0_tx_eop_i	sys_clk_i	Input	<ul style="list-style-type: none"> • End of packet indicator. • Set == 1 coincident with the last link0_tx_data_i word in each TLP.
link0_tx_eop_n_i	sys_clk_i	Input	<ul style="list-style-type: none"> • Nullify packet indicator. • Set == 1 coincident with link0_tx_eop_i == 1 to instruct the core to nullify the current TLP (invert LCRC and use EDB framing) instead of transmitting the TLP normally.

Port	Clock Domain	Direction	Description
lnk0_tx_data_i [NUM_LANES × 64-1:0]	sys_clk_i	Input	<ul style="list-style-type: none"> TLP data to transfer link0_tx_data_i must be valid from the assertion of link0_tx_sop_i until the TLP is fully consumed with the assertion of link0_tx_eop_i. TLP data must comprise a complete Transaction Layer Packet (TLP) as defined by the PCI Express Specification including the entire 3 or 4 DWORD TLP header, data payload (if present), and optionally a TLP Digest (ECRC). The core adds the necessary STP/END/EDB framing, Sequence Number, LCRC, and add ECRC (if enabled to do so and ECRC is not already present in the transmission) as part of its Data Link Layer functionality. Transmitted TLPs are required to be formulated correctly per the PCIe Specification including filling in the Requester/Completer ID, Attributes, Traffic Class, and so on. For Multi-Function, the core uses the Requestor/Completer ID in transmitted TLPs to determine which function's Configuration Registers must be applied to determine the validity of the transmitted TLP. Data width depends on the number of lanes configured for a particular link. <ul style="list-style-type: none"> x8 – 512b x4 – 256b x2 – 128b x1 – 64b <p>Note: There is a known hardware limitation for the December 2023 release. As a workaround for this issue, application logic must have at least 20 clocks gap from tx_eop_i of a TLP to the tx_sop_i of the next TLP. In addition to performance impact, the problem that may be observed is completion timeout for continuous Memory Read TLP from Host to the FPGA.</p> <p>The hardware limitation will be solved in 2024 release.</p>
lnk0_tx_data_p_i [NUM_LANES × 8-1:0]	sys_clk_i	Input	<ul style="list-style-type: none"> Parity of associated link0_tx_data_i: <ul style="list-style-type: none"> link0_tx_data_p_i[i] == ^((link0_tx_data_i[((i+1) × 8)-1:(i×8)]) link0_tx_data_p_i must be valid for all bytes in link0_tx_data_i that contain a portion of a TLP (header, payload, and TLP digest (if present)). Parity width changes as data width, 1 parity bit per 8 bits of data.

Note:

1. LINK – values (0); NUM_LANES – range (1,8)

4.3.1.2. TLP Transmit Credit Interface Port Description

Table 4.4. TLP Transmit Credit Interface Ports

Port	Clock Domain	Direction	Description
link0_tx_credit_init_o	sys_clk_i	Output	<p>When the core Transaction Layer for Link[i] is ready to accept TLP transmissions, the core asserts link0_tx_credit_init_o == 1 for one clock cycle and on the same cycle indicates the non-posted TLP Header storage capacity of the Transmit Buffer on link0_tx_credit_nh_o[11:0]. Users are expected to keep and initialize their NH available transmit credit counters on link0_tx_credit_init_o==1.</p> <p>When a non-posted TLP is pending for transmission within user logic, user logic should check the currently available NH credit count for the associated link and hold the transmission until enough NH credits are available to transmit the TLP. Once the TLP has been committed for transmit, the amount of NH credits required by that TLP are decremented from the NH credit count. As the core forwards transmitted TLPs from the Transmit Buffer and thus makes room for new TLPs, the core asserts link0_tx_credit_return_o==1 for one clock cycle and places the number of NH credits being returned on link0_tx_credit_nh_o[11:0]. In this manner, the user can manage sending only enough non-posted TLPs that the core can hold in its Transmt Buffer. This permits the user to know when non-posted TLPs would be blocked and thus send posted and/or completion TLPs instead. This is important for avoiding deadlocks and keeps non-posted TLP blockage from reducing posted and completion throughput.</p> <p>Should the core receive more non-posted TLPs than the core can store in its non-posted TLP transmit storage, the core pauses TLP transmission rather than allow an overflow to occur. Thus, users that do not wish to use the Transmit Credit Interface may ignore this interface provided they are willing to permit blocked non-posted TLPs from also blocking following posted and completion TLPs.</p>
link0_tx_credit_return_o	sys_clk_i	Output	As the core forwards transmitted TLPs from the Transmit Buffer and thus makes room for new TLPs, the core asserts link0_tx_credit_return_o==1 for one clock cycle and places the number of NH credits being returned on link0_tx_credit_nh_o[11:0].
link0_tx_credit_nh_o[11:0]	sys_clk_i	Output	Number of NH credits to return == link0_tx_credit_nh_o[11:0].

4.3.2. TLP Receive Interface

Refer to the [TLP Receive Interface](#) section for more information and timing diagrams.

4.3.2.1. TLP Receive Interface Port Descriptions

Table 4.5. TLP Receive Interface Ports

Port	Clock Domain	Direction	Description
link0_rx_valid_o	sys_clk_i	Output	<ul style="list-style-type: none"> link0_rx_valid_o == 1 when the other link0_rx_* output ports are valid. Otherwise, value is 0. A data transfer occurs when link0_rx_valid_o == 1 and link0_rx_ready_i == 1.
link0_rx_ready_i	sys_clk_i	Input	<ul style="list-style-type: none"> Set link0_rx_ready_i == 1 whenever the user logic is ready to accept received TLP data. A data transfers occur when link0_rx_valid_o == 1 and link0_rx_ready_i == 1. For maximum throughput, user's design must consume TLPs at the rate that they are presented (link0_rx_ready_i == 1 on all clock cycles).
link0_rx_sel_o[1:0]	sys_clk_i	Output	<ul style="list-style-type: none"> Receive TLP type indicator: <ul style="list-style-type: none"> 0 == Posted Request (write request) 1 == Non-Posted Request (request requiring a completion) 2 == Completion (completion to a previous request) 3 == Reserved link0_rx_sel_o is useful for steering the TLP to the appropriate processing logic. For example, Posted Requests should be directed to receive write logic while Non-Posted Requests should be directed to receive read logic. Completions should be directed back to the original read request source using the TLP Tag information. link0_rx_sel_o is valid for the entire TLP (from link0_rx_sop_o == 1 through link0_rx_eop_o == 1)

Port	Clock Domain	Direction	Description
link0_rx_cmd_data_o[12:0]	sys_clk_i	Output	<ul style="list-style-type: none"> Received TLP Type Indicator link0_rx_cmd_data_o provides information about the received TLP to facilitate user TLP processing. This port has a different meaning in Root Port Modes of operation and Endpoint operation. The Lattice PCIe x8 Core decodes received TLPs to determine their destination. The core passes this information to the Transaction Layer Receive Interface by asserting the appropriate bits in this field. Individual bits of link0_rx_cmd_data_o[12:0] carry the following information: <ul style="list-style-type: none"> Bits[12:10] – Traffic Class[2:0] of the TLP Bit[9] – Completion/Base Address Region indicator (1) indicates the TLP is a Completion or Message routed by ID (0) indicates the TLP is a read or write request or a Message routed by address that hit an enabled Base Address Region. Bits[8:0] are decoded differently based on the value of Bit[9] Bits[8:0] – Completion or Message routed by ID (Bit[9] == 1) Bits[8] – Reserved Bits[7:0] – Tag[7:0] – The Requestor Tag contained in the TLP. Tag[7:0] is used to associate the completion with its original source request. This field is reserved if the TLP is a message rather than a completion. Bits[8:0] – Read or write request or a Message routed by address that hit an enabled Base Address Region (Bit[9] == 0) Bit[8] – When (1), the packet is a <i>write</i> transaction; when (0), the packet is a <i>read</i> transaction. Bit[7] – When (1), the packet requires one or more Completion transactions as a response; (0) otherwise. Bit[6] – (1) the TLP hit the Expansion ROM else (0) Bit[5] – (1) the TLP hit Base Address Region 5 else (0) Bit[4] – (1) the TLP hit Base Address Region 4 else (0) Bit[3] – (1) the TLP hit Base Address Region 3 else (0) Bit[2] – (1) the TLP hit Base Address Region 2 else (0) Bit[1] – (1) the TLP hit Base Address Region 1 else (0) Bit[0] – (1) the TLP hit Base Address Region 0 else (0) link0_rx_cmd_data_o is valid for the entire packet (from link0_rx_sop_o == 1 through link0_rx_eop_o == 1).
link0_rx_f_o[5:0]	sys_clk_i	Output	<ul style="list-style-type: none"> Function Hit by the Received TLP link0_rx_f_o indicates which PCIe Function received the TLP. <ul style="list-style-type: none"> link0_rx_f_o == 0 indicates Function #0. link0_rx_f_o == 1 indicates Function #1 and so on. <p>Note: This signal is only functional in 2024 release or later. Prior to 2024 release, the application logic must decode the address/ BDF of the received TLP and compare against the BAR in PCIe Configuration Registers to know which Physical Function is hit.</p>
link0_rx_sop_o	sys_clk_i	Output	<ul style="list-style-type: none"> Start of TLP indicator link0_rx_sop_o == 1 coincident with the first link0_rx_data_o word in each TLP. Otherwise, value is 0.

Port	Clock Domain	Direction	Description
link0_rx_eop_o	sys_clk_i	Output	<ul style="list-style-type: none"> End of TLP indicator. link0_rx_eop_o == 1 coincident with the last link0_rx_data_o word in each TLP. Otherwise, value is 0.
link0_rx_err_ecrc_o	sys_clk_i	Output	<ul style="list-style-type: none"> Received TLP ECRC error indicator link0_rx_err_ecrc_o == 1 from link0_rx_sop_o to link0_rx_eop_o inclusive for received TLPs which contain a detected ECRC error. Otherwise, value is 0. link0_rx_err_ecrc_o only reports ECRC errors when ECRC checking is enabled. ECRC checking is enabled by software through the AER Capability. TLPs with ECRC errors are presented on the Receive Interface in the same format that they are received including the TLP Digest (ECRC). ECRC errors are serious, uncorrectable errors. The user design must decide how to handle/recover from the error including whether to use the TLP with the error. ECRC errors need for higher level software to correct/handle the error. PCIe does not have a standard mechanism for retransmitting TLPs end to end as it does for a given PCIe link (through the LCRC/SequenceNumber and Replay mechanisms).
link0_rx_data_o [NUM_LANES×64-1:0]	sys_clk_i	Output	<ul style="list-style-type: none"> Received TLP Data Received TLP data comprises a complete Transaction Layer Packet (TLP) as defined by the PCI Express Specification including the entire TLP header, data payload (if present), and TLP Digest (ECRC, if present). The core strips the packet's STP/END/EDB framing, Sequence Number, and Link CRC (LCRC) prior to the TLP appearing on this interface. The core checks TLP ECRC, when present and when checking is enabled, and can be optionally enabled to remove the ECRC from the TLP. Data width depends on the number of lanes configured for a particular link. <ul style="list-style-type: none"> x8 – 512b x4 – 256b x2 – 128b x1 – 64b
link0_rx_data_p_o [NUM_LANES×8-1:0]	sys_clk_i	Output	<ul style="list-style-type: none"> Received TLP Data Parity Even parity of associated link0_rx_data_o: <ul style="list-style-type: none"> link0_rx_data_p_o[i] == ^(link0_rx_data_o[((i+1) ×8)-1:(i×8)]) link0_rx_data_p_o is valid for all bytes in link0_rx_data_o that contain a portion of a TLP (header, payload (if present), and TLP digest (ECRC, if present)). link0_rx_data_p_o is not valid for any trailing, unused bytes in the final link0_rx_data_o word in a TLP. Parity width changes as data width, 1 parity bit per 8 bits of data.

Note:

1. LINK – values (0); NUM_LANES – range (1,8)

4.3.2.2. TLP Receive Credit Interface Port Description

Table 4.6. TLP Receive Credit Interface Ports

Port	Clock Domain	Direction	Description
link0_rx_credit_init_i	sys_clk_i	Input	<p>When the user transaction layer logic is ready to accept non-posted TLP reception, assert the link0_rx_credit_init_i == 1 for one clock cycle and on the same cycle indicates the non-posted TLP header storage capacity of the user design in link0_rx_credit_nh_i[11:0].</p> <p>You must initialize link0_rx_credit_init_i shortly (within 10s of clocks) after u_tl_link_up for Root Port and shortly after u_dl_link_up for Endpoint. Holding off credit initialization for an extended period can cause received non-posted TLP transactions to timeout in the source component which may be serious errors.</p> <p>The core limits simultaneous outstanding non-posted receive TLPs on the receive interface to ensure no more than the initialized NH credits are simultaneously outstanding to user receive TLP logic.</p> <p>Once the received non-posted TLPs are processed/forwarded such that more room is available to receive new non-posted TLPs, assert link0_rx_credit_return_i==1 for one clock cycle and places the number of NH credits being returned on link0_rx_credit_nh_i[11:0]. In this manner, you can limit the outstanding core receive TLPs to the user design. This permits the core to know when non-posted TLPs would be blocked and thus send posted and/or completion TLPs to the user design instead. This is important for avoiding deadlocks and keeps non-posted TLP blockage from reducing posted and completion throughput.</p> <p>If you do not wish to implement flow control of NH credits through this interface, link0_rx_credit_init==1 and link0_rx_credit_nh_inf_i is set to 1 to advertise infinite NH credits. The NH credit flow control is not implemented for links that advertised infinite NH credits.</p>
link0_rx_credit_return_i	sys_clk_i	Input	<p>Once the received non-posted TLPs are processed/forwarded such that more room is available to receive new non-posted TLPs, assert link0_rx_credit_return_i==1 for one clock cycle and places the number of NH credits being returned on link0_rx_credit_nh_i[11:0].</p>
link0_rx_credit_nh_i[11:0]	sys_clk_i	Input	Number of NH credits to return == link0_rx_credit_nh_i[11:0].
link0_rx_credit_nh_inf_i	sys_clk_i	Input	<p>Infinite NH Credits</p> <p>link0_rx_credit_nh_inf_i: 1==Do not limit TLP reception due to NH credits 0==Limit simultaneously outstanding NH credits to the value of link0_rx_credit_nh_i[11:0] when link0_rx_credit_init was 1.</p>

4.4. Lattice Memory Mapped Interface (LMMI)

The Lattice PCIe x8 IP Core implements a bus for configuring core options and obtaining core status. The Core Configuration and Status Registers (CSR) are made accessible to the user design through the Lattice Memory Mapped interface (LMMI).

Table 4.7. Lattice Memory Mapped Interface Ports

Port	Clock Domain	Direction	Description
usr_lmml_clk_i	usr_lmml_clk_i	Input	LMMI Clock (250 MHz max). ¹
usr_lmml_rst_n_i	usr_lmml_clk_i	Input	Active low, asynchronous assert, synchronous de-assert reset Note: Once usr_lmml_rst_n_i de-asserted, you must wait for 10 or more usr_lmml_clk_i clocks before accessing LMMI interface. This is to allow reset to be properly propagated through the PCIe controller.
usr_lmml_offset_i [21:0]	usr_lmml_clk_i	Input	<ul style="list-style-type: none"> Register address. See Register Description for more details on address mapping. usr_lmml_offset_i [23:22] – Not used. Drive 0 usr_lmml_offset_i [21:20] - Quad identifier 2'b00 – Quad 0 2'b01 – Quad 1 2'b10 – Quad 2 (Future release) 2'b11 – Quad 3 (Future release) usr_lmml_offset_i [19] – Not used. Drive 0 usr_lmml_offset_i [18:16] – Core Register Block Identifier PCIe Configuration Space Registers: <ul style="list-style-type: none"> 3'b000 – 3'b001 – Hard PCIe x8 Core PF and VF registers PCIe Core Registers: <ul style="list-style-type: none"> 3'b100 – Hard PCIe x8 Core registers MPP Registers: <ul style="list-style-type: none"> 3'b101 – MPPHY, External M PLLA, External M PLLB, Aggregate Control, CMU, and HRC registers PMA Registers: <ul style="list-style-type: none"> 3'b110 – 3'b111 – PMA Lane, Raw PCS Lane, Raw PCS Always-on Logic, PMA Extra Space, Raw PCS Extra Space, Raw PCS Always-on Logic Extra Space, PMA ROM and PMA RAM Memory access registers For DW-aligned register space: <ul style="list-style-type: none"> usr_lmml_offset_i [15:2] - Register offset usr_lmml_offset_i [1:0] – reserved (tie to 0).
usr_lmml_request_i	usr_lmml_clk_i	Input	Start Transaction (1==Active; 0==Otherwise) A transaction is started when usr_lmml_request_i==usr_lmml_ready_o==1. When usr_lmml_request_i==usr_lmml_ready_o==1, usr_lmml_wr_rdn_i, and usr_lmml_offset_i must be valid and describe the transaction to execute; if the transaction is a write as indicated by usr_lmml_wr_rdn_i==1, usr_lmml_wdata_i must also be valid.
usr_lmml_wr_rdn_i	usr_lmml_clk_i	Input	Direction (1==Write, 0==Read)
usr_lmml_wdata_i [31:0]	usr_lmml_clk_i	Input	Write data
usr_lmml_rdata_o [31:0]	usr_lmml_clk_i	Output	Read data
usr_lmml_ready_o	usr_lmml_clk_i	Output	Target is ready to start a new transaction. (1==Ready; 0==Not ready)
usr_lmml_rdata_valid_o	usr_lmml_clk_i	Output	Read transaction is complete and usr_lmml_rdata_o contains valid data (1==Valid; 0==Otherwise).

Note:

- Support for 250 MHz is only available in the next release (2024.1).

4.5. Power Management Interface

The Lattice PCIe x8 IP Core supports optional capabilities such as Dynamic Power Allocation, Latency Tolerance Reporting, and Power Budgeting.

Note: Not supported in 2023.2 release. Supported in Radiant 2024.1 and later.

4.6. AXI-L Configuration Interface

The Lattice PCIe x8 Core provides an option to use AXI4-Lite interface as alternative to LMMI for register access.

The clock supplied to soft IP's user configuration path must be synchronous to LMMI clock in Hard-IP. If AXI4-Lite requires a separate clock domain, the clock crossing from AXI4-Lite clock to LMMI clock must be done externally to the soft IP.

Table 4.8. AXI-L Configuration Interface Ports

Port	Type	Function
Write Address Channel		
axi4l_awvalid_i	Input	Write address valid
axi4l_awaddr_i [23:0]	Input	Write address
axi4l_awprot_i [2:0]	Input	Protection type. Not supported (Tied to 3'b000).
axi4l_awready_o	Output	Write address ready
Write Data Channel		
axi4l_wvalid_i	Input	Write valid
axi4l_wdata_i [AXI4L_DATAWIDTH-1:0]	Input	Write data
axi4l_wstrb_i [(AXI4L_DATAWIDTH/8)-1:0]	Input	Write strobes
axi4l_wready_o	Output	Write ready
Write Response Channel		
axi4l_bready_i	Input	Response ready
axi4l_bvalid_o	Output	Write response valid
axi4l_bresp_o[1:0]	Output	Write response
Read Address Channel		
axi4l_arvalid_i	Input	Read address valid
axi4l_araddr_i [23:0]	Input	Read address
axi4l_arprot_i [2:0]	Input	Protection type. Not supported (Tied to 3'b000).
axi4l_arready_o	Output	Read address ready
Read Data Channel		
axi4l_rready_i	Input	Read ready
axi4l_rvalid_o	Output	Read valid
axi4l_rdata_o [AXI4L_DATAWIDTH-1:0]	Output	Read data
axi4l_rresp_o [1:0]	Output	Read response

4.7. AXI-4 Stream Data Interface

This interface is available if the application interface type selected in the IP generation is AXI4. The AXI4 interface is slightly different if Hard DMA Core is enabled. If you did not enable any DMA Core, the soft logic only has the AXI bridging application. If Soft DMA Core is selected, the soft logic has AXI bridging and Scatter-Gather DMA application. If Hard DMA Core is selected, the default hardware AXI interface is used (no soft logic).

4.7.1. AXI4-Stream Transmitter Interface Port Descriptions

Table 4.9. AXI-4 Stream Transmitter Interface Ports

Port	Clock Domain	Direction	Description
m0_tready_i	clk_usr_div2_i	Input	Destination ready. 1==Ready, 0==Not ready. A transfer occurs when m_tvalid_o==m_tready_i==1.
m0_tvalid_o	clk_usr_div2_i	Output	Source valid 1==Valid 0==Not valid.
m0_tdata_o [NUM_LANES*64-1:0]	clk_usr_div2_i	Output	For Link 0, data width depends on bifurcation option selected. <ul style="list-style-type: none"> 1×4 – 256b 1×2 – 128b 1×1 – 64b
m0_tstrb_o [NUM_LANES*8-1:0]	clk_usr_div2_i	Output	<ul style="list-style-type: none"> Byte qualifier that indicates whether the content of the associated byte of m_tdata_o is processed as a data byte or a position byte. For Link 0, tstrb width depends on bifurcation option selected. <ul style="list-style-type: none"> 1×4 – 256b. The value is 32'hFFFFFFF. 1×2 – 128b. The value is 16'hFFFF. 1×1 – 64b. The value is 8'hFF.
m0_tkeep_o [NUM_LANES*8-1:0]	clk_usr_div2_i	Output	<ul style="list-style-type: none"> Byte qualifier that indicates whether the content of the associated byte of m_tdata_o is processed as part of the data stream. Associated bytes that have the m_tkeep_o byte qualifier deasserted are null bytes and can be removed from the data stream. For Link 0, tkeep width depends on bifurcation option selected. <ul style="list-style-type: none"> 1×4 – 256b. The value is 32'hFFFFFFF. 1×2 – 128b. The value is 16'hFFFF. 1×1 – 64b. The value is 8'hFF.
m0_tlast_o	clk_usr_div2_i	Output	End of TLP indicator. m_tlast_o == 1 coincident with the last m_tdata_o word in each TLP. Otherwise, 0.
m0_tid_o [7:0]	clk_usr_div2_i	Output	Data stream identifier that indicates different streams of data. m0_tid_o[2:0] has the BAR number information when rx_cmd_data[9] = 0. Otherwise, 0 when rx_cmd_data[9] = 1(completion). <ul style="list-style-type: none"> m0_tid_o[3] = link0_rx_err_par m0_tid_o[6:4] = link0_rx_cmd_data[12:10] m0_tid_o[7] = link0_rx_err_ecrc

Port	Clock Domain	Direction	Description
m0_tdest_o [3:0]	clk_usr_div2_i	Output	<ul style="list-style-type: none"> m_tdest_o provides routing information for the data stream. Bits [3:2] – Function Hit by the Received TLP Bits [1:0] – Receive TLP type indicator: <ul style="list-style-type: none"> 0 == Posted Request (write request) 1 == Non-Posted Request (request requiring a completion) 2 == Completion (completion to a previous request) <p>Note: Bits [3:0] are only functional in 2024 release or later. Prior to 2024 release, the application logic must decode the address/BDF of the received TLP and compare against the BAR in PCIe Configuration Registers to know which Physical Function is hit.</p>

Note:

1. NUM_LANES – range (1,4)

4.7.2. AXI4-Stream Receiver Interface Port Descriptions

Table 4.10. AXI-4 Stream Receiver Interface Ports

Port	Clock Domain	Direction	Description
s0_tvalid_i	clk_usr_div2_i	Input	Source valid 1==Valid 0==Not valid.
s0_tdata_i [NUM_LANES*64-1:0]	clk_usr_div2_i	Input	<ul style="list-style-type: none"> TLP data to transfer For Link 0, data width depends on bifurcation option selected. <ul style="list-style-type: none"> 1×4 – 256b 1×2 – 128b 1×1 – 64b <p>Note: There is a known hardware limitation in the December 2023 release. As workaround for this issue, application logic must have at least 20 clocks gap from tx_eop_i of a TLP to the tx_sop_i of the next TLP. In addition to performance impact, another problem to look out for is the completion timeout for continuous Memory Read TLP from Host to the FPGA. This hardware limitation will be solved in the 2024 release.</p>
s0_tstrb_i [NUM_LANES*8-1:0]	clk_usr_div2_i	Input	<ul style="list-style-type: none"> Byte qualifier that indicates whether the content of the associated byte of m_tdata_o is processed as a data byte or a position byte. For Link 0, tstrb width depends on the bifurcation option selected. <ul style="list-style-type: none"> 1×4 – 256b. The value is 32'hFFFFFFF. 1×2 – 128b. The value is 16'hFFFF. 1×1 – 64b. The value is 8'hFF.

Port	Clock Domain	Direction	Description
s0_tkeep_i [NUM_LANES*8-1:0]	clk_usr_div2_i	Input	<ul style="list-style-type: none"> Byte qualifier that indicates whether the content of the associated byte of m_tdata_o is processed as part of the data stream. Associated bytes that have the m_tkeep_o byte qualifier deasserted are null bytes and can be removed from the data stream. For Link 0, tkeep width depends on the bifurcation option selected. <ul style="list-style-type: none"> 1×4 – 256b. The value is 32'hFFFFFFF. 1×2 – 128b. The value is 16'hFFFF. 1×1 – 64b. The value is 8'hFF.
s0_tlast_i	clk_usr_div2_i	Input	End of packet indicator. Set == 1 coincident with the last s[LINK]_tdata_i word in each TLP.
s0_tid_i [7:0]	clk_usr_div2_i	Input	Unused. Set to 8'h00.
s0_tdest_i [3:0]	clk_usr_div2_i	Input	Unused. Set to 4'h0.
s0_tready_o	clk_usr_div2_i	Output	Destination ready. 1==Ready 0==Not ready A transfer occurs when s_tvalid_i==s_tready_o==1.

Note:

1. NUM_LANES – range (1,4)

4.8. DMA Interface

Table 4.11. DMA AXI4 Manager Write Interface

Port	Clock Domain	Direction	Description
m0_axi_awid_o [3:0]	sys_clk_i	Output	This signal is the identification tag for the write address group of signals.
m0_axi_awaddr_o [63:0]	sys_clk_i	Output	The write address gives the address of the first transfer in a write burst transaction.
m0_axi_awlen_o [3:0]	sys_clk_i	Output	The burst length gives the exact number of transfers in a burst. This information determines the number of data transfers associated with the address.
m0_axi_awsize_o [2:0]	sys_clk_i	Output	This signal indicates the size of each transfer in the burst.
m0_axi_awprot_o [2:0]	sys_clk_i	Output	This signal is used only for Security. Only m0_axi4_awprot [1] is used; [2],[0] are unused and set to 0. <ul style="list-style-type: none"> m0_axi_awprot [1] = 0: Secure m0_axi_awprot [1] = 1: Non-Secure
m0_axi_awcache_o [3:0]	sys_clk_i	Output	This signal indicates how transactions are required to progress through a system.
m0_axi_awvalid_o	sys_clk_i	Output	This signal indicates that the channel is signaling valid write address and control information.
m0_axi_awready_i	sys_clk_i	Input	This signal indicates that the slave is ready to accept an address and associated control signals.

Port	Clock Domain	Direction	Description
m0_axi_awuser_o [47:0]	sys_clk_i	Output	<p>Provides additional useful transaction information which may be optionally used:</p> <ul style="list-style-type: none"> m0_axi_awuser_o [47:32] <ul style="list-style-type: none"> Provides the PCI Express Requester ID {Bus[7:0], Device[4:0], Function[2:0]} (or {Bus[7:0], Function[7:0]} for ARI capable sources) received with transactions that originated from PCI Express. m0_axi_awuser_o [31:24] <ul style="list-style-type: none"> Provides the function number targeted by the transaction. For transaction originating from PCIe, this is the function that has claimed the transaction. For DMA transactions, this is the function that contains the registers for the DMA channel. m0_axi_awuser_o [23] <ul style="list-style-type: none"> For DMA sources (m0_axi_awuser_o [11] == 1), it indicates when the transaction is the end of a DMA packet. This information is typically needed for FIFO DMA applications. This bit is always 1 for non-DMA sources. FIFO DMA applications need to know when the packet ends because the final packet word cannot be written to a DMA write FIFO or popped from a DMA read FIFO until the full word has been provided/consumed or the end of a packet has occurred indicating no more data will be coming/requested. m0_axi_awuser_o [22] – Reserved. m0_axi_awuser_o [21:12] <ul style="list-style-type: none"> Provides the exact byte count being requested. The requested byte count is otherwise only known to the resolution m0_axi_awsz_o of or by inspecting the byte enables. m0_axi_awuser_o [11:0] <ul style="list-style-type: none"> Encodes the transaction source. FIFO DMA applications, which may associate each hardware FIFO with a DMA Channel, can use this information to determine which FIFO is being addressed. m0_axi_awuser_o [11:0] encoding is described in the text following this table.
m0_axi_wid_o [3:0]	sys_clk_i	m0_axi_wid_o [3:0]	This signal is the ID tag of the write data transfer. Supported only in AXI3.
m0_axi_wdata_o [511:0]	sys_clk_i	m0_axi_wdata_o [511:0]	Write data.
m0_axi_wstrb_o [63:0]	sys_clk_i	m0_axi_wstrb_o [63:0]	This signal indicates which byte lanes hold valid data. There is one write strobe bit for each eight bits of the write data bus.
m0_axi_wlast_o	sys_clk_i	m0_axi_wlast_o	This signal indicates the last transfer in a write burst.
m0_axi_wvalid_o	sys_clk_i	m0_axi_wvalid_o	This signal indicates that valid write data and strobes are available.
m0_axi_wready_i	sys_clk_i	m0_axi_wready_i	This signal indicates that the subordinate can accept the write data.
m0_axi_bid_i [3:0]	sys_clk_i	m0_axi_bid_i [3:0]	This signal is the ID tag of the write response.
m0_axi_bresp_i [1:0]	sys_clk_i	m0_axi_bresp_i [1:0]	This signal indicates the status of the write transaction.

Port	Clock Domain	Direction	Description
m0_axi_bvalid_i	sys_clk_i	Input	This signal indicates that the channel is signaling a valid write response.
m0_axi_bready_o	sys_clk_i	Output	This signal indicates that the manager can accept a write response.

Table 4.12. DMA AXI4 Manager Read Interface

Port	Clock Domain	Direction	Description
m0_axi_arid_o [3:0]	sys_clk_i	Output	This signal is the identification tag for the read address group of signals.
m0_axi_araddr_o [63:0]	sys_clk_i	Output	The read address gives the address of the first transfer in a read burst transaction.
m0_axi_arlen_o [3:0]	sys_clk_i	Output	This signal indicates the exact number of transfers in a burst.
m0_axi_arsize_o [2:0]	sys_clk_i	Output	This signal indicates the size of each transfer in the burst.
m0_axi_arprot_o [2:0]	sys_clk_i	Output	This signal is used only for Security. Only m0_axi_arprot [1] is used; [2],[0] are unused and set to 0. <ul style="list-style-type: none"> m0_axi_arprot [1] = 0: Secure m0_axi_arprot [1] = 1: Non-Secure
m0_axi_arcache_o [3:0]	sys_clk_i	Output	This signal indicates how transactions are required to progress through a system.
m0_axi_arvalid_o	sys_clk_i	Output	This signal indicates that the channel is signaling valid read address and control information.
m0_axi_arready_i	sys_clk_i	Input	This signal indicates that the slave is ready to accept an address and associated control signals.

Port	Clock Domain	Direction	Description
m0_axi_aruser_o [47:0]	sys_clk_i	Input	<p>Provides additional useful transaction information which may be optionally used:</p> <ul style="list-style-type: none"> m0_axi_awuser_o [47:32] <ul style="list-style-type: none"> Provides the PCI Express Requester ID {Bus[7:0], Device[4:0], Function[2:0]} (or {Bus[7:0], Function[7:0]} for ARI capable sources) received with transactions that originated from PCI Express. m0_axi_awuser_o [31:24] <ul style="list-style-type: none"> Provides the function number targeted by the transaction. For transaction originating from PCIe, this is the function that has claimed the transaction. For DMA transactions, this is the function that contains the registers for the DMA channel. m0_axi_awuser_o [23] <ul style="list-style-type: none"> For DMA sources (m0_axi_awuser_o [11] == 1), it indicates when the transaction is the end of a DMA packet. This information is typically needed for FIFO DMA applications. This bit is always 1 for non-DMA sources. FIFO DMA applications need to know when the packet ends because the final packet word cannot be written to a DMA write FIFO or popped from a DMA read FIFO until the full word has been provided/consumed or the end of a packet has occurred indicating no more data will be coming/requested. m0_axi_awuser_o [22] – Reserved. m0_axi_awuser_o [21:12] <ul style="list-style-type: none"> The requested byte count is otherwise only known to the resolution of m0_axi_arsize_o. FIFO DMA applications, which generally cannot over-read, typically need this information. m0_axi_awuser_o [11:0] <ul style="list-style-type: none"> Encodes the transaction source. FIFO DMA applications, which may associate each hardware FIFO with a DMA Channel, can use this information to determine which FIFO is being addressed. m0_axi_awuser_o [11:0] encoding is described in the text following this table.
m0_axi_rid_i [3:0]	sys_clk_i	Input	This signal is the identification tag for the read data group of signals generated by the subordinate.
m0_axi_rdata_i [511:0]	sys_clk_i	Input	Read data.
m0_axi_rresp_i [1:0]	sys_clk_i	Input	This signal indicates the status of the read transfer.
m0_axi_rlast_i	sys_clk_i	Input	This signal indicates the last transfer in a read burst.
m0_axi_rvalid_i	sys_clk_i	Input	This signal indicates that the channel is signaling the required read data.
m0_axi_rready_o	sys_clk_i	Output	This signal indicates that the manager can accept the read data and response information.

Table 4.13. DMA AXI4 Subordinate Write Interface

Port	Clock Domain	Direction	Description
s0_axi_awid_i [15:0]	sys_clk_i	Input	This signal is the identification tag for the write address group of signals.
s0_axi_awaddr_i [63:0]	sys_clk_i	Input	The write address gives the address of the first transfer in a write burst transaction.
s0_axi_awlen_i [3:0]	sys_clk_i	Input	The burst length gives the exact number of transfers in a burst. This information determines the number of data transfers associated with the address. Max of 16 beat bursts.
s0_axi_awsz_i [2:0]	sys_clk_i	Input	This signal indicates the size of each transfer in the burst.
s0_axi_awprot_i [2:0]	sys_clk_i	Input	This signal is used only for Security. Only s0_axi_awprot [1] is used; [2],[0] are unused and set to 0. <ul style="list-style-type: none"> s0_axi_awprot [1] = 0: Secure s0_axi_awprot [1] = 1: Non-Secure
s0_axi_awcache_i [3:0]	sys_clk_i	Input	This signal indicates how transactions are required to progress through a system.
s0_axi_awvalid_i	sys_clk_i	Input	This signal indicates that the channel is signaling valid write address and control information.
s0_axi_awready_o	sys_clk_i	Output	This signal indicates that the slave is ready to accept an address and associated control signals.
s0_axi_awuser_i [7:0]	sys_clk_i	Input	Use to specify the function number to be used if this write request is destined for PCIe and the PCIe core supports multiple functions.
s0_axi_wid_i [3:0]	sys_clk_i	Input	This signal is the ID tag of the write data transfer. Supported only in AXI3. Must tie to 'b0 in AXI4.
s0_axi_wdata_i [511:0]	sys_clk_i	Input	Write data.
s0_axi_wstrb_i [63:0]	sys_clk_i	Input	This signal indicates which byte lanes hold valid data. There is one write strobe bit for each eight bits of the write data bus.
s0_axi_wlast_i	sys_clk_i	Input	This signal indicates the last transfer in a write burst.
s0_axi_wvalid_i	sys_clk_i	Input	This signal indicates that valid write data and strobes are available.
s0_axi_wready_o	sys_clk_i	Output	This signal indicates that the subordinate can accept the write data.
s0_axi_bid_o [3:0]	sys_clk_i	Output	This signal is the ID tag of the write response.
s0_axi_bresp_o [1:0]	sys_clk_i	Output	This signal indicates the status of the write transaction.
s0_axi_bvalid_o	sys_clk_i	Output	This signal indicates that the channel is signaling a valid write response.
s0_axi_bready_i	sys_clk_i	Input	This signal indicates that the manager can accept a write response.

Table 4.14. DMA AXI4 Subordinate Read Interface

Port	Clock Domain	Direction	Description
s0_axi_arid_i [15:0]	sys_clk_i	Output	This signal is the identification tag for the read address group of signals.
s0_axi_araddr_i [63:0]	sys_clk_i	Output	The read address gives the address of the first transfer in a read burst transaction.
s0_axi_arlen_i [3:0]	sys_clk_i	Output	This signal indicates the exact number of transfers in a burst.
s0_axi_arsize_i [2:0]	sys_clk_i	Output	This signal indicates the size of each transfer in the burst.
s0_axi_arprot_i [2:0]	sys_clk_i	Output	This signal is used only for Security. Only m0_axi_arprot [1] is used; [2],[0] are unused and set to 0. <ul style="list-style-type: none"> s0_axi_arprot [1] = 0: Secure s0_axi_arprot [1] = 1: Non-Secure
s0_axi_arcache_i [3:0]	sys_clk_i	Output	This signal indicates how transactions are required to progress through a system.
s0_axi_arvalid_i	sys_clk_i	Output	This signal indicates that the channel is signaling valid read address and control information.
s0_axi_arready_o	sys_clk_i	Input	This signal indicates that the slave is ready to accept an address and associated control signals.
s0_axi_aruser_i [7:0]	sys_clk_i	Input	Use to specify the function number to be used if this write request is destined for PCIe and the PCIe core supports multiple functions.
s0_axi_rid_i [3:0]	sys_clk_i	Input	This signal is the identification tag for the read data group of signals generated by the subordinate.
s0_axi_rdata_i [511:0]	sys_clk_i	Input	Read data.
s0_axi_rresp_i [1:0]	sys_clk_i	Input	This signal indicates the status of the read transfer.
s0_axi_rlast_i	sys_clk_i	Input	This signal indicates the last transfer in a read burst.
s0_axi_rvalid_i	sys_clk_i	Input	This signal indicates that the channel is signaling the required read data.
s0_axi_rready_o	sys_clk_i	Output	This signal indicates that the manager can accept the read data and response information.

Table 4.15. DMA AXI4 Interrupt Interface

Port	Clock Domain	Direction	Description
s_int_tx_i[63:0]	sys_clk_i	Input	This signal is used to cause interrupts to be generated on PCIe. It must be driven depending upon the value of s_int_tx_edge_level_n_o: 1 – Edge Interrupt mode PCIe Core is configured for MSI-X or MSI interrupts Set s_int_tx_i[i] == 1 for one clock to request an interrupt on interrupt vector[i] 0 – Level Interrupt mode PCIe Core is configured for Legacy interrupts Set s_int_tx_i[i] == 1 for as long as interrupt vector[i] continues to be pending s_int_tx[i] must be 0 except when an interrupt is desired to be generated.
s_int_tx_edge_level_n_o[63:0]	sys_clk_i	Output	s_int_tx_edge_level_n_o[i] indicates how interrupts on s_int_tx_i[i] must be signaled. For cores supporting multi-function or SR-IOV, different interrupt vectors are associated with each function and each function has independent MSI-X and MSI enables.

4.9. Unused Interface

The signals listed in the [Table 4.16](#) are unused and removed in the next release.

Table 4.16. Unused Interface

Port	Direction
link0_legacy_interrupt_i	Input
link0_legacy_interrupt_o	Output
link0_user_aux_power_detected_i	Input
link0_user_transactions_pending_i	Input
refret_i	Input
rext_i	Input

5. Register Description

Table 5.1. Register Access Abbreviations

Abbreviation	Meaning
RW	Read and Write access
RO	Read only
WO	Write only
RW1C	Read write 1 to clear

5.1. Hard IP Core Configuration and Status Registers

The Lattice PCIe x8 IP Core configuration registers have default values that are appropriate for most applications. Customers typically would only want to change a small number of values such as Vendor/Device ID and BAR configuration. Such changes can be made through LMMI writes prior to core reset release or through the IP generation user interface. The registers defined in the sections below are the same set for all links. The registers are configured through LMMI and AXI-L interface.

Table 5.2 lists the offset address for the Hard IP Core Registers.

Table 5.2. Hard PCIe Core Register Mapping

Hard IP Core	Register Block	Start Byte Offset	End Byte Offset
QUAD 0			
PCIe x8	PCIe Configuration Space Registers		
	Physical Function 0	0x0_0000	0x0_0FFF
	Physical Function 1	0x0_1000	0x0_1FFF
	Physical Function 2	0x0_2000	0x0_2FFF
	Physical Function 3	0x0_3000	0x0_3FFF
	Physical Function 4	0x0_4000	0x0_4FFF
	Physical Function 5	0x0_5000	0x0_5FFF
	Physical Function 6	0x0_6000	0x0_6FFF
	Physical Function 7	0x0_7000	0x0_7FFF
	PF0 Virtual Function 0	0x0_8000	0x0_8FFF
	PF0 Virtual Function 1	0x0_9000	0x0_9FFF
	PF0 Virtual Function 2	0x0_A000	0x0_AFFF
	PF1 Virtual Function 0	0x0_B000	0x0_BFFF
	PF1 Virtual Function 1	0x0_C000	0x0_CFFF
	PF1 Virtual Function 2	0x0_D000	0x0_DFFF
	PF2 Virtual Function 0	0x0_E000	0x0_EFFF
	PF2 Virtual Function 1	0x0_F000	0x0_FFFF
	PF2 Virtual Function 2	0x1_0000	0x1_0FFF
	PF3 Virtual Function 0	0x1_1000	0x1_1FFF
	PF3 Virtual Function 1	0x1_2000	0x1_2FFF
	PF3 Virtual Function 2	0x1_3000	0x1_3FFF
	PF4 Virtual Function 0	0x1_4000	0x1_4FFF
	PF4 Virtual Function 1	0x1_5000	0x1_5FFF
	PF4 Virtual Function 2	0x1_6000	0x1_6FFF

Hard IP Core	Register Block	Start Byte Offset	End Byte Offset
	PF5 Virtual Function 0	0x1_7000	0x1_7FFF
	PF5 Virtual Function 1	0x1_8000	0x1_8FFF
	PF5 Virtual Function 2	0x1_9000	0x1_9FFF
	PF6 Virtual Function 0	0x1_A000	0x1_AFFF
	PF6 Virtual Function 1	0x1_B000	0x1_BFFF
	PF6 Virtual Function 2	0x1_C000	0x1_CFFF
	PF7 Virtual Function 0	0x1_D000	0x1_DFFF
	PF7 Virtual Function 1	0x1_E000	0x1_EFFF
	PF7 Virtual Function 2	0x1_F000	0x1_FFFF
	PCIe x8 Core Registers		
	mgmt_tlb	0x4_2000	0x4_2717
	mgmt_ptl	0x4_3000	0x4_32E3
	mgmt_ftl	0x4_4000	0x4_4337
	mgmt_ftl_mf1	0x4_5000	0x4_5337
	mgmt_ftl_mf2	0x4_6000	0x4_6337
	mgmt_ftl_mf3	0x4_7000	0x4_7337
	mgmt_ftl_mf4	0x4_8000	0x4_8337
	mgmt_ftl_mf5	0x4_9000	0x4_9337
	mgmt_ftl_mf6	0x4_A000	0x4_A337
	mgmt_ftl_mf7	0x4_B000	0x4_B337
Quad 0 MPPHY	MPPHY 0 Registers		
	MPPHY Lane 0 – Lane 3 (PCS)	0x05_0000	0x05_0FFF
	Quad 0 PMA, Common	0x06_0000	0x06_FFFF
	Quad 0 PHY ROM	0x07_0000	0x07_3FFF
	Quad 0 PHY RAM	0x07_4000	0x07_FFFF
QUAD 1			
Quad 1 MPPHY	MPPHY 1 Registers		
	MPP Lane 4 – Lane 7 (PCS)	0x15_0000	0x15_0FFF
	Quad 1 PMA, Common	0x16_0000	0x16_FFFF
	Quad 1 PHY ROM	0x17_0000	0x17_3FFF
	Quad 1 PHY RAM	0x17_4000	0x17_FFFF

5.1.1. EP Configuration Settings

The Lattice PCIe x8 IP Core supports Endpoint (EP) operation. The current mode of operation is determined by the core CSR. The following table illustrates the CSR values that are recommended for EP and RP applications.

Table 5.3. CSR Values Recommended for EP Applications

Register Field	Offset	EndPoint
mgmt_tlb_ltssm_port_type_ds_us_n	0x4_2040	1'b0
mgmt_ftl_cfg_type1_type0_n	0x4_4030	1'b0
mgmt_ftl_decode_ignore_poison	0x4_4010	1'b0
mgmt_ftl_decode_t1_rx_bypass_msg_dec	0x4_4014	1'b0
mgmt_ftl_pcie_cap_slot_implemented	0x4_4080	1'b0
mgmt_ftl_pcie_cap_device_port_type	0x4_4080	4'h0
mgmt_ftl_id3_class_code	0x4_4048	User Application Specific

Register Field	Offset	EndPoint
mgmt_ftl_ari_cap_disable	0x4_40E0	1'b0
mgmt_ftl_msi_cap_disable	0x4_40E8	1'b0
mgmt_ftl_msi_cap_mult_message_capable	0x4_40E8	User Application Specific
mgmt_ftl_msix_cap_table_size	0x4_40F0	User Application Specific
mgmt_ftl_msix_cap_disable	0x4_40F0	1'b0 (Enabled)
mgmt_ftl_aer_cap_en_surprise_down_error	0x4_4100	1'b0

5.1.2. mgmt_tlb (0x4_2000)

The following are the register sets with the 0x4_2000 base address.

Note: Registers in the range not described in this section are considered reserved and must not be accessed. The access to this range results in unexpected behavior.

5.1.2.1. LTSSM Register Set

ltssm_simulation Register 0x0

This register set is used for LTSSM simulation speed reduction.

Table 5.4. ltssm_simulation Register 0x0

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1]	reduce_ts1	RW	1	0x0	Reduce the minimum number of TS1 transmitted in Polling.Active from 1024 to 16 to shorten simulation time. 0 – Disable 1 – Enable
[0]	reduce_timeouts	RW	1	0x0	Reduce LTSSM timeouts to shorten simulation time. When enabled, 1 ms-> 20 μ s, 2 ms->40 μ s, 12 ms->60 μ s, 24 ms->80 μ s, 32 ms->100 μ s, and 48 ms->160 μ s. 0 – Disable 1 – Enable

ltssm_cfg_lw_start Register 0x34

This register set is used for LTSSM CFG.LWSTART configuration.

Table 5.5. ltssm_cfg_lw_start Register 0x34

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1:0]	min_time	RW	2	0x0	Minimum time spent in Cfg.LW.Start before exit is permitted. 0 – 4 μ s 1 – 16 μ s 2 – 64 μ s

ltssm_latch_rx Register 0x38

This register set is used for LTSSM latch RX configuration.

Table 5.6. ltssm_latch_rx Register 0x38

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	30	0x0	—
[0]	link_lane	RW	1	0x1	Enable latching each lane's received link and lane numbers and state exit condition during LTSSM Configuration link width negotiation. 0 – Disable. The lane is included in the link if it is receiving the state exit criteria on the clock cycle that the link width and state exit transition is occurring. A received Physical Layer error occurring close to the clock cycle that the link width is being determined results in a reduction of link width even if the lane had previously recorded valid state exit criteria. 1 – Enable. The lane is included in the link if it met the state exit criteria at any time during the state. This is the recommended setting since received Physical Layer errors are less likely to result in reduced link width.

ltssm_cfg Register 0x3c

This register set is used for LTSSM configuration.

Table 5.7. ltssm_cfg Register 0x3c

Field	Name	Access	Width	Reset	Description
[31:28]	lw_start_updn_end_delay	RW	4	0x9	LTSSM CFG_[US/DS]_LW_START normal CFG_[US/DS]_LW_START TS1 transmissions and parsing of received TS OS begins (lw_start_updn_end_delay × 64) symbols after the bp_ltssm_cfg_lw_start_updn 1 to 0 transition occurs at the end of PHY adaptation. This delay is intended to flush any corrupted PHY rx data due to the PHY adaptation through the Link Layer Core before the Core begins paying attention to received data again.
[27:24]	lw_start_updn_start_delay	RW	4	0x8	LTSSM CFG_[US/DS]_LW_START bp_ltssm_cfg_lw_start_updn==1 assertion is delayed by (lw_start_updn_start_delay × 64) symbols from CFG_[US/DS]_LW_START state entry. The start delay is intended to avoid the PHY beginning adaptation, and thus corrupting the input data, before the link partner data stream has ended. When the Core reaches CFG_[US/DS]_LW_START before the link partner, the link partner may still be in Recovery.Idle with an active data stream. The start delay must be long enough to delay PHY adaptation until the receive data stream has ended or else SKP Data Parity Errors and Receiver Errors can be detected and recorded by the Core due to the PHY corrupting the receive data stream due to adaptation.
[23:12]	lw_start_updn_count	RW	12	0xfa	LTSSM CFG_[US/DS]_LW_START bp_ltssm_cfg_lw_start_updn==1 duration is set to (lw_start_updn_count × 1024) ns. 0==Disabled.

Field	Name	Access	Width	Reset	Description
[11:8]	lw_start_updn_rate_en	RW	4	0xf	LTSSM CFG_[US/DS]_LW_START bp_ltssm_cfg_lw_start_updn==1 rate enable/disable. Controls for which speeds the bp_ltssm_cfg_lw_start_updn feature is supported. One bit is provided to enable/disable each speed supported {16G, 8G, 5G, 2.5G}. Bit positions for speeds that are not supported by a given core delivery must be set to 0. 0 – Disable feature when at the associated link speed. 1 – Enable feature when at the associated link speed.
[7:6]	reserved	RO	2	0x0	—
[5]	lw_start_updn_eie_en	RW	1	0x0	<i>lw_start_updn_eie_en</i> LTSSM CFG_[US/DS]_LW_START bp_ltssm_cfg_lw_start_updn==1 EIE Tx OS enable. 0 – Disabled 1 – Enabled
[4]	lw_start_updn_en_dir_ds	RW	1	0x0	LTSSM CFG_[US/DS]_LW_START bp_ltssm_cfg_lw_start_updn==1 directed down- configure enable. 0 – Do not assert bp_ltssm_cfg_lw_start_updn==1 when the CFG_[US/DS]_LW_START entry is due to locally directed downconfigure. 1 – Assert bp_ltssm_cfg_lw_start_updn==1 when the CFG_[US/DS]_LW_START entry is due to locally directed down-configure.
[3:2]	reserved	RO	2	0x0	—
[1]	lw_start_updn_timer_en	RW	1	0x0	LTSSM CFG_[US/DS]_LW_START bp_ltssm_cfg_lw_start_updn Timer Enable. Register lw_start_updn_timer_en can be set to stay in adaptation for a fixed period instead of relying on the PHY to have a port bp_ltssm_cfg_lw_start_updn_ack that is asserted at the end of adaptation. Only one of lw_start_updn_timer_en and lw_start_updn_ack_en can be set to 1. 0 – Disabled. 1 – Deassert bp_ltssm_cfg_lw_start_updn after (lw_start_updn_count × 1024) ns has expired.

Field	Name	Access	Width	Reset	Description
[0]	lw_start_updn_ack_en	RW	1	0x0	<p>LTSSM Configuration Link Width Start</p> <p>bp_ltssm_cfg_lw_start_updn Ack Enable</p> <p>0 – Disabled. Output port bp_ltssm_cfg_lw_start_updn is held == 0 and input port bp_ltssm_cfg_lw_start_updn_ack is ignored. When CFG_[DS/US]_LW_START is entered from Recovery, the transition from CFG_[DS/US]_LW_START to CFG_[DS/US]_LW_ACCEPT occurs after a minimum of 4 μs.</p> <p>1 – Enabled. If also enabled, through mgmt_tlb_ltssm_cfg_lw_start_updn_rate_en, at the current link speed, output port bp_ltssm_cfg_lw_start_updn is set upon CFG_[DS/US]_LW_START entry from Recovery and input port bp_ltssm_cfg_lw_start_updn_ack is used. The transition from CFG_[DS/US]_LW_START to CFG_[DS/US]_LW_ACCEPT occurs only after the PHY has asserted bp_ltssm_cfg_lw_start_updn_ack == 1 and additionally a minimum of 4 μs has elapsed. bp_ltssm_cfg_lw_start_updn_ack must not be withheld so long that the state timeout of 24 ms expires or the link exits to detect, and the link goes down, which is a serious error.</p>

ltssm_port_type Register 0x40

This register set is used for the LTSSM port type configuration.

Table 5.8. ltssm_port_type Register 0x40

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	ds_us_n	RW	1	0x0	<p>Determines the PCI Express port type which affects many aspects of LTSSM training.</p> <p>0 – Upstream Port</p> <p>1 – Downstream Port</p>

ltssm_ds_link Register 0x44

This register set is used for the LTSSM downstream link configuration.

Table 5.9. ltssm_ds_link Register 0x44

Field	Name	Access	Width	Reset	Description
[31:5]	reserved	RO	25	0x0	—
[4:0]	number	RW	5	0x0	For downstream ports only, unique Link Number assigned to the link and used in TS sets during LTSSM Configuration

ltssm_detect_quiet Register 0x48

This register set is used for the LTSSM Detect.Quiet configuration.

Table 5.10. ltssm_detect_quiet Register 0x48

Field	Name	Access	Width	Reset	Description
[31:5]	reserved	RO	25	0x0	—
[4:0]	number	RW	5	0x0	For downstream ports only, unique Link Number assigned to the link and used in TS sets during LTSSM Configuration

ltssm_rx_det Register 0x4c

This register set is used for the LTSSM receiver detection configuration.

Table 5.11. ltssm_rx_det Register 0x4c

Field	Name	Access	Width	Reset	Description
[31]	override	RW	1	0x0	Lane receiver detection mask enable. 0 – Disable 1 – Enable
[30:16]	reserved	RO	15	0x0	—
[15:0]	mask	RW	16	0x0	<i>mask</i> Lane receiver detection mask. When override==1, mask determines which lanes attempt receiver detection. For each lane[i]: 0 – Skip receiver detection and exclude the lane from the link. 1 – Perform receiver detection and use result to determine whether to include/exclude the lane from the link.

ltssm_nfts Register 0x50

This register set is used for the LTSSM Nfts configuration.

Table 5.12. ltssm_nfts Register 0x50

Field	Name	Access	Width	Reset	Description
[31:16]	reserved	RO	16	0x0	—
[15:8]	to_extend	RW	8	0x7f	Number of FTS set transfer times to wait in addition to the time required to transmit the requested Nfts sets before timing out to Recovery on Rx_LOs exit.
[7:0]	nfts	RW	8	0xff	Number of FTS sets to request link partner transmit when exiting LOs. Nfts value transmitted in TS1 and TS2 Ordered Sets during training.

ltssm_ds_initial_auto Register 0x54

This register set is used for the LTSSM initial link speed configuration.

Table 5.13. ltssm_ds_initial_auto Register 0x54

Field	Name	Access	Width	Reset	Description
[31]	rate_enable	RW	1	0x0	<i>rate_enable</i> Determines whether link speed up is requested by the core after the first entry to L0 following state Detect. If neither port directs the link to a higher speed, the link remains at 2.5G unless software initiates a speed change. It is recommended to set rate_enable=1 and rate==maximum supported speed. 0 – Let the link partner or software initiate initial speed changes. 1 – Make 1 attempt to direct the link to the maximum speed specified by rate. The speed achieved is the maximum speed, less than or equal to rate, that both the core and link partner support.
[30:2]	reserved	RO	29	0x0	Number of FTS set transfer times to wait in addition to the time required to transmit the requested NFTS sets before timing out to Recovery on Rx_LOs exit.
[1:0]	rate	RW	2	0x0	<i>rate</i> When rate_enable==1, indicates the maximum rate that is attempted to negotiate on the initial link training from Detect. Only speeds supported by the core can be indicated. 0 – 2.5G 1 – 5G 2 – 8G ¹ 3 – 16G ¹

Note:

1. Available in 2024 release.

ltssm_select_deemphasis Register 0x58

This register set is used for the LTSSM 2.5/5G deemphasis configuration.

Table 5.14. ltssm_select_deemphasis Register 0x58

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	6db_3_5db_n	RW	1	0x1	For 5G capable cores only: For upstream ports only, sets the default deemphasis for 5G operation during LTSSM State Detect. 0 – -3.5dB 1 – -6dB

Itssm_beacon Register 0x5c

This register set is used for the LTSSM Beacon configuration.

Table 5.15. Itssm_beacon Register 0x5c

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	l2_d3hot_enable	RW	1	0x0	<p><i>l2_d3hot_enable</i></p> <p>L2 wake Beacon transmission control.</p> <p>0 – Disabled. The customer design must wake the link through WAKE# pin assertion. Set to 0 when using PHY which do not support Beacon transmission. Set to 0 if the core is not clocked (some PHY remove the core's clock in L2 while others supply a keep alive clock) or powered (some applications remove core power in L2 to maximize power savings) in L2, as the core is unable to initiate Beacon generation in these cases.</p> <p>1 – Transmit beacon when directed to wake the link from L2.</p>

Itssm_mod_cpl Register 0x60

This register set is used for the LTSSM Modified Compliance configuration.

Table 5.16. Itssm_mod_cpl Register 0x60

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	31	0x0	—
[1]	one_eieos	RW	1	0x1	<p>When entering Modified Compliance Pattern determines the number of EIEOS blocks to send.</p> <p>0 – Send 8 EIEOS blocks to ensure receiver lock</p> <p>1 – Send 1 EIEOS (as per Spec)</p>
[0]	exit_direct_to_detect	RW	1	0x0	<p><i>exit_direct_to_detect</i></p> <p>When transmitting Modified Compliance Pattern and <code>cfg_enter_compliance == 0</code>, determines which of the two PCIe Specification optional behaviors is selected.</p> <p>0 – Do not exit to Detect for this reason.</p> <p>1 – Exit to Detect.</p>

Itssm_rx_elec_idle Register 0x64

This register set is used for the LTSSM Rx Electrical Idle configuration.

Table 5.17. Itssm_rx_elec_idle Register 0x64

Field	Name	Access	Width	Reset	Description
[31]	rec_spd_infer_rcvr_lock	RW	1	0x0	<p>Recovery Speed successful and unsuccessful inference expand to Recovery.RcvrLock enable.</p> <p>0 – Do not include time spent in Recovery.RcvrLock when calculating successful and unsuccessful speed change electrical idle inference in Recovery.Speed.</p> <p>1 – Include time spent in Recovery.RcvrLock when calculating successful and unsuccessful speed change electrical idle inference in Recovery.Speed.</p>

Field	Name	Access	Width	Reset	Description
[30]	rec_pd_infer_rcvr_cfg	RW	1	0x0	Recovery Speed successful and unsuccessful inference expand to Recovery.RcvrCfg enable. 0 – Do not include time spent in Recovery.RcvrCfg when calculating successful and unsuccessful speed change electrical idle inference in Recovery.Speed. 1 – Include time spent in Recovery.RcvrCfg when calculating successful and unsuccessful speed change electrical idle inference in Recovery.Speed.
[29]	rec_spd_infer_eq_ph0123	RW	1	0x0	Recovery Speed successful and unsuccessful inference expand to Recovery.EqPhase0123 enable. 0 – Do not include time spent in Recovery.EqPhase0123 when calculating successful and unsuccessful speed change electrical idle inference in Recovery.Speed. 1 – Include time spent in Recovery.EqPhase0123 when calculating successful and unsuccessful speed change electrical idle inference in Recovery.Speed.
[28:4]	reserved	RO	25	0x0	—
[3:0]	filter	RW	4	0x1	After entering a LTSSM state that monitors, pipe_rx_elec_idle for exit, ignore pipe_rx_elec_idle for 128 × filter) nanoseconds to enable tolerance for pipe_rx_elec_idle not latency matched with the associaetd pipe_rx_data.

ltssm_compliance_toggle Register 0x68

This register set is used for the LTSSM Compliance Toggle configuration.

Table 5.18. ltssm_compliance_toggle Register 0x68

Field	Name	Access	Width	Reset	Description
[31:4]	reserved	RO	28	0x0	—
[3:2]	max_speed	RW	2	0x3	Maximum speed of compliance patterns that should be generated. 0 – 2.5G 1 – 5G 2 – 8G ¹ 3 – 16G ¹
[1:0]	min_speed	RW	2	0x0	Minimum speed of compliance patterns that should be generated. 0 – 2.5G 1 – 5G 2 – 8G ¹ 3 – 16G ¹

Note:

1. Available in 2024 release.

ltssm_prevent_rx_ts_entry_to Register 0x6c

This register set is used for the LTSSM State Rx TS Transition Prevention configuration.

Table 5.19. ltssm_prevent_rx_ts_entry_to Register 0x6c

Field	Name	Access	Width	Reset	Description
[31:4]	reserved	RO	28	0x0	—
[3]	compliance	RW	1	0x0	LTSSM to Polling.Compliance Rx TS state transition disable. 0 – Enabled 1 – Disabled
[2]	loopback	RW	1	0x0	LTSSM to Loopback Follower Rx TS state transition disable. 0 – Enabled 1 – Disabled
[1]	hot_reset	RW	1	0x0	LTSSM to Hot Reset Rx TS state transition disable. 0 – Enabled 1 – Disabled
[0]	disable	RW	1	0x0	LTSSM to Disable Rx TS state transition disable. 0 – Enabled 1 – Disabled

ltssm_link Register 0x80

This register is used for the Current Link Status configuration.

Table 5.20. ltssm_link Register 0x80

Field	Name	Access	Width	Reset	Description
[31]	dl_link_up	RO	1	0x0	Data Link Layer link up status. 0 – Down 1 – Up
[30]	pl_link_up	RO	1	0x0	Physical Layer link up status. 0 – Down 1 – Up
[29:20]	Reserved	RO	10	0x0	—
[19:16]	lane_rev_status	RO	4	0x0	Indicates the current lane reversal status: lane_rev_status[0], 1 == Full Reverse is in effect else 0 lane_rev_status[1], 1 == x2 Reverse is in effect (≥ 4 lane only) else 0 lane_rev_status[2], 1 == x4 Reverse is in effect (≥ 8 lane only) else 0 lane_rev_status[3], 1 == x8 Reverse is in effect (≥ 16 lane only) else 0
[15]	idle_infer_rec_rcvr_cfg	RW1C	1	0x0	Electrical Idle inference status in Recovery.RcwrCfg. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[14]	idle_infer_loopback_slave	RW1C	1	0x0	Electrical Idle inference status in Loopback.Active as a Loopback Follower. 0 – Otherwise 1 – Event occurred. Write 1 to clear
[13]	idle_infer_rec_speed2_success	RW1C	1	0x0	Electrical Idle inference status in Recovery.Speed on a successful speed negotiation. 0 – Otherwise 1 – Event occurred. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[12]	idle_infer_rec_speed2_unsuccess	RW1C	1	0x0	Electrical Idle inference status in Recovery.Speed on an unsuccessful speed negotiation. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[11]	idle_infer_l0_to_rec_rcvr_lock	RW1C	1	0x0	Electrical Idle inference status in L0 – event causes entry into Recovery.RcvrLock. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[10:9]	Reserved	RO	2	0x0	—
[8]	speed_change_fail	RW1C	1	0x0	Speed Change Failure error indicator. 0 – Otherwise 1 – Speed change failure occurred. Write 1 to clear.
[7:2]	Reserved	RO	6	0x0	—
[1:0]	speed	RO	2	0x0	Current LTSSM Link Speed. Only link speeds supported by the core is be indicated. 0 – 2.5G 1 – 5G 2 – 8G ¹ 3 – 16G ¹

Note:

1. Available in 2024 release.

ltssm_ltssm Register 0x84

This register set is used for LTSSM State Machine State configuration.

Table 5.21. ltssm_ltssm Register 0x84

Field	Name	Access	Width	Reset	Description
[31:20]	Reserved	RO	12	0x0	—
[19:16]	sub_state	RO	4	0x1	Current LTSSM Minor State. Encoding varies depending upon the Current LTSSM Major State. 0 – DETECT_INACTIVE 1 – DETECT_QUIET 2 – DETECT_SPD_CHG0 3 – DETECT_SPD_CHG1 4 – DETECT_ACTIVE0 5 – DETECT_ACTIVE1 6 – DETECT_ACTIVE2 7 – DETECT_P1_TO_P0 8 – DETECT_P0_TO_P1_0 9 – DETECT_P0_TO_P1_1 10 – DETECT_P0_TO_P1_2 0 – POLLING_INACTIVE 1 – POLLING_ACTIVE_ENTRY 2 – POLLING_ACTIVE 3 – POLLING_CFG 4 – POLLING_COMP 5 – POLLING_COMP_ENTRY 6 – POLLING_COMP_EIOS 7 – POLLING_COMP_EIOS_ACK 8 – POLLING_COMP_IDLE 0 – CONFIGURATION_INACTIVE 1 – CONFIGURATION_US_LW_START

Field	Name	Access	Width	Reset	Description
					2 – CONFIGURATION_US_LW_ACCEPT 3 – CONFIGURATION_US_LN_WAIT 4 – CONFIGURATION_US_LN_ACCEPT 5 – CONFIGURATION_DS_LW_START 6 – CONFIGURATION_DS_LW_ACCEPT 7 – CONFIGURATION_DS_LN_WAIT 8 – CONFIGURATION_DS_LN_ACCEPT 9 – CONFIGURATION_COMPLETE 10 – CONFIGURATION_IDLE 0 – L0_INACTIVE 1 – L0_L0 2 – L0_TX_EL_IDLE 3 – L0_TX_IDLE_MIN 0 – RECOVERY_INACTIVE 1 – RECOVERY_RCVR_LOCK 2 – RECOVERY_RCVR_CFG 3 – RECOVERY_IDLE 4 – RECOVERY_SPEED0 5 – RECOVERY_SPEED1 6 – RECOVERY_SPEED2 7 – RECOVERY_SPEED3 8 – RECOVERY_EQ_PH0 9 – RECOVERY_EQ_PH1 10 – RECOVERY_EQ_PH2 11 – RECOVERY_EQ_PH3 0 – DISABLED_INACTIVE 1 – DISABLED_0 2 – DISABLED_1 3 – DISABLED_2 4 – DISABLED_3 0 – LOOPBACK_INACTIVE 1 – LOOPBACK_ENTRY 2 – LOOPBACK_ENTRY_EXIT 3 – LOOPBACK_EIOS 4 – LOOPBACK_EIOS_ACK 5 – LOOPBACK_IDLE 6 – LOOPBACK_ACTIVE 7 – LOOPBACK_EXIT0 8 – LOOPBACK_EXIT1 0 – HOT_RESET_INACTIVE 1 – HOT_RESET_HOT_RESET 2 – HOT_RESET_LEADER_UP 3 – HOT_RESET_LEADER_DOWN 0 – TX_LOS_INACTIVE 1 – TX_LOS_IDLE 2 – TX_LOS_TO_L0 3 – TX_LOS_FTS0 4 – TX_LOS_FTS1 0 – L1_INACTIVE 1 – L1_IDLE 2 – L1_SUBSTATE 3 – L1_TO_L0 0 – L2_INACTIVE 1 – L2_IDLE

Field	Name	Access	Width	Reset	Description
					2 – L2_TX_WAKE0 3 – L2_TX_WAKE1 4 – L2_EXIT 5 – L2_SPEED
[15:4]	Reserved	RO	12	0x0	—
[3:0]	state	RO	4	0x0	Current LTSSM Major State 0 – DETECT 1 – POLLING 2 – CONFIGURATION 3 – L0 4 – RECOVERY 5 – DISABLED 6 – LOOPBACK 7 – HOT_RESET 8 – TX_LOS 9 – L1 10 – L2

ltssm_rx_l0s Register 0x88

This register set is used for the Rx L0s State Machine State configuration.

Table 5.22. ltssm_rx_l0s Register 0x88

Field	Name	Access	Width	Reset	Description
[31:3]	Reserved	RO	29	0x0	—
[2:0]	state	RO	3	0x0	Current LTSSM RX L0s State. 0 – RX_LOS_L0 1 – RX_LOS_ENTRY 2 – RX_LOS_IDLE 3 – RX_LOS_FTS 4 – RX_LOS_REC

IO_to_rec Register 0x8c

This register set is used to report different events causing L0 state to Recovery state transition.

Table 5.23. IO_to_rec Register 0x8c

Field	Name	Access	Width	Reset	Description
[31:15]	reserved	RO	17	0x0	—
[14]	direct_to_detect_fast	RW1C	1	0x0	Recovery is entered from L0 due to assertion of mgmt_tlb_ltssm_direct_to_detect_fast. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[13]	direct_to_recovery_ch_bond	RW1C	1	0x0	Recovery is entered from L0 due to more lane skew than the Channel Bond circuit can tolerate or is due to channel bond failing to occur within the expected timeout period. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[12]	direct_to_loopback_entry	RW1C	1	0x0	Recovery is entered from L0 due to being directed into Leader Loopback. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[11]	directed_speed_change	RW1C	1	0x0	Recovery is entered from L0 due to being directed to make a speed change. This includes the initial hardware-initiated speed change(s) which are made when first exiting Detect.Quiet to L0. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[10]	IO_to_rec_rcvr_lock_rx_ts12	RW1C	1	0x0	Recovery is entered from L0 due to receiving TS1 or TS2 ordered sets. The link partner is directing Recovery entry. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[9]	IO_to_rec_rcvr_lock_rx_8g_eie	RW1C	1	0x0	Recovery is entered from L0 due to receiving EIE ordered sets at ≥ 8G. The link partner is directing Recovery entry. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[8]	IO_to_rec_rcvr_lock_rx_infer	RW1C	1	0x0	Recovery is entered from L0 due to inferring Electrical Idle due to no SKP ordered set received in 128 μs. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[7]	direct_to_recovery_phy	RW1C	1	0x0	Recovery is entered from L0 due to receiving a burst of ~1024 clock cycles of data containing PHY errors at 2.5G or 5G. This normally occurs only when the PHY has lost lock on one or more lanes. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[6]	direct_to_recovery_frame	RW1C	1	0x0	Recovery is entered from L0 due to receiving one or more framing errors at ≥ 8G. This occurs due to Rx bit errors which are expected every few minutes at PCIe Specified BER of 10 ⁻¹² . 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[5]	direct_to_recovery_replay	RW1C	1	0x0	Recovery is entered from L0 due to the original and three replay TLP transmissions failing to receive ACK DLLP acknowledgment. 0 – Otherwise 1 – Event occurred. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[4]	direct_to_hot_reset	RW1C	1	0x0	Recovery is entered from L0 due to being directed into Hot Reset (Secondary Bus Reset Register). 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[3]	direct_to_disable	RW1C	1	0x0	Recovery is entered from L0 due to being directed into Disable (Link Disable Register). 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[2]	rx_l0s_direct_to_recovery	RW1C	1	0x0	Recovery is entered from L0 due to failing to receive the complete Rx_L0S FTS exit sequence within the PCIe Specification allowed timeout period. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[1]	autonomous_width_change	RW1C	1	0x0	Recovery is entered from L0 due to directed autonomous width change. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[0]	directed_retrain_link	RW1C	1	0x0	Recovery is entered from L0 due to directed retrain link (Retrain Link Register). 0 – Otherwise 1 – Event occurred. Write 1 to clear.

ltssm_rx_detect Register 0x90

This register set is used for the Receiver detection status.

Table 5.24. ltssm_rx_detect Register 0x90

Field	Name	Access	Width	Reset	Description
[31:16]	reserved	RO	16	0x0	—
[15:0]	lanes	RO	16	0x0	Per lane receiver detection status. For each lane: 0 – Unconnected 1 – Present

ltssm_configured Register 0x94

This register set is used for the Configured link status.

Table 5.25. ltssm_configured Register 0x94

Field	Name	Access	Width	Reset	Description
[31:25]	reserved	RO	7	0x0	—
[24:16]	link_num	RO	9	0x1ff	Link Number configured during LTSSM Training. link_num == 0x1FF on fundamental reset, changes to 0x1F7 (KPAD) when entering CFG_US_LW_START or CFG_DS_LW_START (the start of LTSSM Configuration), and then changes to the negotiated Link Number determined during LTSSM Configuration when the LTSSM changes from CFG_COMPLETE to CFG_IDLE. This field is provided for diagnostics.
[15:0]	lanes	RO	16	0x0	Per lane configured link status. Each lane status resets to 0. After Receiver Detection results are available, each lane status is updated to show which lanes detected receivers. After a link has been formed, each lane status is updated to show which lanes are part of the configured link. For each lane: 0 – Lane did not configure into the link. 1 – Lane configured into the link.

ltssm_direct_to_detect Register 0x98

This register set is used for the Rec Rcvr Lock to Detect controls.

Table 5.26. ltssm_direct_to_detect Register 0x98

Field	Name	Access	Width	Reset	Description
[31:16]	reserved	RO	16	0x0	—
[15]	fast	RW	9	0x0	A rising edge on this signal instructs the state machine to proceed from L0 or Recovery to Detect as quickly as possible.
[14:8]	Reserved	RO	7	0x0	—
[7:0]	timer	RW	8	0x0	This value determines the timeout delay for the state machine to proceed from Recovery Rcvr Lock to Detect when no TS sets are received. A value of 0 disables this timeout.

ltssm_equalization Register 0x9c

This register set is used for the for ≥ 8 G capable cores only: LTSSM equalization status.

Table 5.27. ltssm_equalization Register 0x9c

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	fail	RO	1	0x0	Equalization Failure error indicator. 0 – Otherwise 1 – Equalization failure.

Itssm_crosslink Register 0xa0

This register set is used for the TSSM crosslink status.

Table 5.28. Itssm_crosslink Register 0xa0

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1]	ds_us_n	RO	1	0x0	Crosslink port type. When active==1, indicates which personality the port assumed during crosslink negotiation. 0 – Upstream 1 – Downstream
[0]	active	RO	1	0x0	Crosslink active indicator. 0 – Otherwise 1 – Link is operating in a crosslink configuration.

5.1.2.2. Physical Layer Status Register Set

Physical Layer Tx Underflow Error Status Register – 0xa4

Table 5.29. Physical Layer Tx Underflow Error Status Register – 0xa4

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	err_tx_pipe_underflow	RW1C	1	0x0	0 – Otherwise 1 – Physical Layer Tx data needed to be forwarded to the lanes for transmission and some, but not all lanes, were ready to accept data causing some lanes to under low. This bit stays asserted once set. Write 1 to clear.

Table 5.30 illustrates the Physical Lane RX Status Register set with its offset and register address.

Table 5.30. Physical Lane Rx Status Registers

Register Name	Offset Address	Description
pl_rx0 Register	0xa8	Lane Rx Status 0 register – TS2 and TS1 OS detection [0 to 15 bits]
pl_rx1 Register	0xac	Lane Rx Status 1 – Inverted TS2 and TS1 OS detection [0 to 15 bits]
pl_rx2 Register	0xb0	Lane Rx Status 2 – FTS and SKP OS detection
pl_rx3 Register	0xb4	Lane Rx Status 3 – EIOS detection and EIE detection
pl_rx4 Register	0xb8	Lane Rx Status 4 – Data Block is received and SDS ordered set detection

pl_rx0 Register 0xa8 – Lane Rx Status 0 Register

Table 5.31. pl_rx0 Register 0xa8 – Lane Rx Status 0 Register

Field	Name	Access	Width	Reset	Description
[31]	ts2_detect15	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[30]	ts2_detect14	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[29]	ts2_detect13	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear

Field	Name	Access	Width	Reset	Description
[28]	ts2_detect12	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[27]	ts2_detect11	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[26]	ts2_detect10	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[25]	ts2_detect9	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[24]	ts2_detect8	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[23]	ts2_detect7	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[22]	ts2_detect6	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[21]	ts2_detect5	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[20]	ts2_detect4	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[19]	ts2_detect3	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[18]	ts2_detect2	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[17]	ts2_detect1	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[16]	ts2_detect0	RW1C	1	0x0	ts2_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[15]	ts1_detect15	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[14]	ts1_detect14	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[13]	ts1_detect13	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[12]	ts1_detect12	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear

Field	Name	Access	Width	Reset	Description
[11]	ts1_detect11	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[10]	ts1_detect10	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[9]	ts1_detect9	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[8]	ts1_detect8	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[7]	ts1_detect7	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[6]	ts1_detect6	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[5]	ts1_detect5	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[4]	ts1_detect4	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[3]	ts1_detect3	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[2]	ts1_detect2	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[1]	ts1_detect1	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear
[0]	ts1_detect0	RW1C	1	0x0	ts1_detect[i] is set to 1 when a TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear

pl_rx1 Register 0xac – Lane Rx Status 1

Table 5.32. pl_rx1 Register 0xac – Lane Rx Status 1

Field	Name	Access	Width	Reset	Description
[31]	ts2i_detect15	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[30]	ts2i_detect14	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[29]	ts2i_detect13	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[28]	ts2i_detect12	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[27]	ts2i_detect11	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[26]	ts2i_detect10	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[25]	ts2i_detect9	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[24]	ts2i_detect8	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[23]	ts2i_detect7	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[22]	ts2i_detect6	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[21]	ts2i_detect5	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[20]	ts2i_detect4	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[19]	ts2i_detect3	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[18]	ts2i_detect2	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[17]	ts2i_detect1	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[16]	ts2i_detect0	RW1C	1	0x0	ts2i_detect[i] is set to: 1 when an inverted TS2 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[15]	ts1i_detect15	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[14]	ts1i_detect14	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[13]	ts1i_detect13	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[12]	ts1i_detect12	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[11]	ts1i_detect11	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[10]	ts1i_detect10	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[9]	ts1i_detect9	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[8]	ts1i_detect8	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[7]	ts1i_detect7	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[6]	ts1i_detect6	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[5]	ts1i_detect5	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[4]	ts1i_detect4	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[3]	ts1i_detect3	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[2]	ts1i_detect2	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[1]	ts1i_detect1	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[0]	ts1i_detect0	RW1C	1	0x0	ts1i_detect[i] is set to: 1 when an inverted TS1 ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

pl_rx2 Register 0xb0 – Lane Rx Status 2

Table 5.33. pl_rx2 Register 0xb0 – Lane Rx Status 2

Field	Name	Access	Width	Reset	Description
[31]	fts_detect15	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[30]	fts_detect14	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[29]	fts_detect13	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[28]	fts_detect12	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[27]	fts_detect11	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[26]	fts_detect10	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[25]	fts_detect9	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[24]	fts_detect8	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[23]	fts_detect7	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[22]	fts_detect6	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[21]	fts_detect5	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[20]	fts_detect4	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[19]	fts_detect3	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[18]	fts_detect2	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[17]	fts_detect1	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[16]	fts_detect0	RW1C	1	0x0	fts_detect[i] is set to: 1 when the FTS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[15]	skp_detect15	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[14]	skp_detect14	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[13]	skp_detect13	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[12]	skp_detect12	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[11]	skp_detect11	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[10]	skp_detect10	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[9]	skp_detect9	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[8]	skp_detect8	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[7]	skp_detect7	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[6]	skp_detect6	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[5]	skp_detect5	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[4]	skp_detect4	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[3]	skp_detect3	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[2]	skp_detect2	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[1]	skp_detect1	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[0]	skp_detect0	RW1C	1	0x0	skp_detect[i] is set to: 1 when a SKP ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

pl_rx3 Register 0xb4 – Lane Rx Status 3

Table 5.34. pl_rx3 Register 0xb4 – Lane Rx Status 3

Field	Name	Access	Width	Reset	Description
[31]	eie_detect15	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[30]	eie_detect14	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[29]	eie_detect13	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[28]	eie_detect12	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[27]	eie_detect11	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[26]	eie_detect10	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[25]	eie_detect9	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[24]	eie_detect8	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[23]	eie_detect7	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[22]	eie_detect6	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[21]	eie_detect5	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[20]	eie_detect4	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[19]	eie_detect3	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[18]	eie_detect2	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[17]	eie_detect1	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[16]	eie_detect0	RW1C	1	0x0	eie_detect[i] is set to: 1 when an EIE ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[15]	eios_detect15	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[14]	eios_detect14	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[13]	eios_detect13	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[12]	eios_detect12	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[11]	eios_detect11	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[10]	eios_detect10	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[9]	eios_detect9	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[8]	eios_detect8	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[7]	eios_detect7	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[6]	eios_detect6	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[5]	eios_detect5	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[4]	eios_detect4	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[3]	eios_detect3	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[2]	eios_detect2	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[1]	eios_detect1	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[0]	eios_detect0	RW1C	1	0x0	eios_detect[i] is set to: 1 when an EIOS ordered set is received on Lane[i]. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

pl_rx4 Register 0xb8 – Lane Rx Status 4

Table 5.35. pl_rx4 Register 0xb8 – Lane Rx Status 4

Field	Name	Access	Width	Reset	Description
[31]	data_detect15	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[30]	data_detect14	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[29]	data_detect13	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[28]	data_detect12	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[27]	data_detect11	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[26]	data_detect10	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[25]	data_detect9	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[24]	data_detect8	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[23]	data_detect7	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[22]	data_detect6	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[21]	data_detect5	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[20]	data_detect4	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[19]	data_detect3	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[18]	data_detect2	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[17]	data_detect1	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[16]	data_detect0	RW1C	1	0x0	data_detect[i] is set to: 1 when a Data Block is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[15]	sds_detect15	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[14]	sds_detect14	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[13]	sds_detect13	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[12]	sds_detect12	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[11]	sds_detect11	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[10]	sds_detect10	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[9]	sds_detect9	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[8]	sds_detect8	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[7]	sds_detect7	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[6]	sds_detect6	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[5]	sds_detect5	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[4]	sds_detect4	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[3]	sds_detect3	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[2]	sds_detect2	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[1]	sds_detect1	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.
[0]	sds_detect0	RW1C	1	0x0	sds_detect[i] is set to: 1 when a SDS ordered set is received on Lane[i] at 8G. 0 otherwise. This bit stays asserted once set. Write 1 to clear.

5.1.2.3. Debug Register Set

debugself_crosslink Register 0xc0

This register set is used for debug to allow Rx detection when Tx is externally looped back to Rx.

Table 5.36. debugself_crosslink Register 0xc0

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	0 – Otherwise 1 – For debug use only, configure LTSSM so that it links with itself when core Tx is externally looped back to core Rx.

debug_rx_det Register 0xc4

This register set is used for the LTSSM receiver detection bypass configuration.

Table 5.37. debug_rx_det Register 0xc4

Field	Name	Access	Width	Reset	Description
[31:17]	reserved	RO	15	0x0	—
[16]	inhibit	RW	1	0x0	Link receiver detection inhibit. 0 – Perform receiver detection and use result to determine whether to include/exclude lanes from the link. 1 – Skip receiver detection and assume receivers are not present on all lanes.
[15:1]	reserved	RO	15	0x0	—
[0]	bypass	RW	1	0x0	Link receiver detection bypass. If both bypass and inhibit are asserted, bypass takes precedence. 0 – Perform receiver detection and use result to determine whether to include/exclude lanes from the link. 1 – Skip receiver detection and assume receivers are present on all lanes.

debug_force_tx Register 0xc8

This register set is used for debug using TX PIPE signals.

Table 5.38. debug_force_tx Register 0xc8

Field	Name	Access	Width	Reset	Description
[31:10]	reserved	RO	22	0x0	—
[9]	deemph_5g_enable	RW	1	0x0	For 5G capable cores only: Force pipe_tx_deemph at 5G enable. 0 – Disable. 1 – Enable. Force phy_tx_deemph at 5G speed to the value specified by deemph_5g_6db_3_5db_n. The force is applied at 5G speed except during Polling.Compliance, where for compatibility with PCI SIG Workshop Electrical Testing, the force is not applied.
[8]	deemph_5g_3_5db_6db_n	RW	1	0x0	For 5G capable cores only: Force pipe_tx_deemph at 5G value. 0 – -6dB 1 – -3.5dB
[7:4]	reserved	RO	4	0x0	—
[3]	margin_enable	RW	1	0x0	Force pipe_tx_margin enable. 0 – Drive pipe_tx_margin per PCIe Specification. 1 – Drive pipe_tx_margin to value.
[2:0]	margin_value	RW	3	0x0	Force pipe_tx_margin Value.

debug_direct_scramble_off Register 0xcc

This register set is used for scrambling disable control for 2.5G and 5G data rate.

Table 5.39. debug_direct_scramble_off Register 0xcc

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	LTSSM direct scrambling disabled at 2.5G and 5G. 0 – Otherwise 1 – Direct to disable scrambling at 2.5G and 5G during Configuration.Complete.

debug_force_scramble_off_fast Register 0xd0

This register set is used for scrambling disable control for 8G data rate.

Table 5.40. debug_force_scramble_off_fast Register 0xd0

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	LTSSM force scrambling disabled at ≥ 8G. 0 – Otherwise 1 – Disable scrambling at ≥ 8G. Only works for simulation when link partner is also disabling scrambling. Cannot be set for hardware because disabling scrambling at ≥ 8G is not permitted per PCIe Specification.

balign Register 0xd4

This register set is used for pipe_block_align_control generation options for 8G data rate. It is not recommended to change the default values of this register.

Table 5.41. balign Register 0xd4

Field	Name	Access	Width	Reset	Description
[31]	state_data_n	RW	1	0x0	When generating pipe_block_align_control (which may be used by some PHY to aid in acquiring $\geq 8G$ block alignment), select between the LTSSM State Algorithm and the Rx Data Observation Algorithm. 0 – Use Rx Data Observation Algorithm 1 – Use the LTSSM State Algorithm
[30:6]	reserved	RO	25	0x0	—
[5]	exclude_loopback_master	RW	1	0x0	When generating pipe_block_align_control through the LTSSM State algorithm, exclude/include the Loopback Leader state in driving pipe_block_align_control to 0. 0 – Include 1 – Exclude
[4]	exclude_cfg_complete	RW	1	0x0	When generating pipe_block_align_control through the LTSSM State algorithm, exclude/include the CFG_COMPLETE LTSSM state in driving pipe_block_align_control to 0. pipe_block_align_control is only driven to 0 during CFG_COMPLETE after receiving the required Rx exit criteria to state CFG_IDLE. Due to this requirement, the core may need to stay in CFG_COMPLETE for a while after receiving the Rx exit criteria to also meet the required Tx exit criteria. 0 – Include 1 – Exclude
[3]	exclude_cfg_idle	RW	1	0x0	When generating pipe_block_align_control through the LTSSM State algorithm, exclude/include the CFG_IDLE LTSSM state in driving pipe_block_align_control to 0. 0 – Include 1 – Exclude
[2]	exclude_rec_rcvr_cfg	RW	1	0x0	When generating pipe_block_align_control through the LTSSM State algorithm, exclude/include the REC_RCVR_CFG LTSSM state in driving pipe_block_align_control to 0. pipe_block_align_control is only driven to 0 during REC_RCVR_CFG after receiving the required Rx exit criteria to state REC_IDLE. Due to the PCIe Specification, the core may need to stay in REC_RCVR_CFG for a while after receiving the Rx exit criteria to also meet the required Tx exit criteria. 0 – Include 1 – Exclude
[1]	exclude_rec_idle	RW	1	0x0	When generating pipe_block_align_control through the LTSSM State algorithm, exclude/include the REC_IDLE LTSSM state in driving pipe_block_align_control to 0. 0 – Include 1 – Exclude
[0]	exclude_l0	RW	1	0x0	When generating pipe_block_align_control through the LTSSM State algorithm, exclude/include the L0 and TX_L0s LTSSM states in driving pipe_block_align_control to 0. 0 – Include 1 – Exclude

debug_pipe_rx Register 0xe0

This register set is used for the PIPE Interface Debug status.

Table 5.42. debug_pipe_rx Register 0xe0

Field	Name	Access	Width	Reset	Description
[31:16]	polarity	RO	16	0x0	PHY PIPE Interface pipe_rx_polarity current value. For each lane: 0 – Otherwise 1 – PHY lane has been instructed to invert its receiver polarity to compensate for serial rx_p and rx_n being swapped.
[15:0]	valid	RO	16	0x0	PHY PIPE Interface pipe_rx_valid current value. For each lane: 0 – Otherwise 1 – PHY lane is locked to data stream

debug_direct_to_loopback Register 0x100

This register set is used to enable LTSSM Leader loopback.

Table 5.43. debug_direct_to_loopback Register 0x100

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	LTSSM leader loopback enable. 0 – Otherwise. 1 – Direct LTSSM to Loopback.Leader. Before this field is set to 1, all relevant registers containing Leader Loopback control options must be set to the desired values. When mgmt_tlb_debug_direct_to_loopback == 1 no Leader Loopback control options may be changed.

debug_loopback_control Register 0x104

This register set is used enable different control features related to loopback for 2.5G and 5G data rates.

Table 5.44. debug_loopback_control Register 0x104

Field	Name	Access	Width	Reset	Description
[31:28]	inject_err_lane_select	RW	4	0x0	Lane selection to inject error in Loopback. Only lanes configured by the core may be programmed. 0 = Lane 0. 15 = Lane 15.
[27]	inject_rx_2bit_data_err	RW	1	0x0	When enabled during loopback, the rising edge of inject_rx_2bit_data_err bit injects a back-to-back error on the received loopback data. This simulates the PHY losing lock and increments the counter by one, and the Loopback Leader restarts the loopback pattern so that the PHY can recover symbol lock.
[26]	inject_rx_1bit_data_err	RW	1	0x0	When enabled during loopback, the rising edge of inject_rx_1bit_data_err bit injects a single clk error on the received loopback data. This causes the error count to increment by 1 for each received data byte.
[25]	inject_rx_valid_err	RW	1	0x0	When enabled during loopback, the rising edge of inject_rx_valid_err bit injects a single clk error on the received PIPE PHY interface phy_rx_valid signal. This simulates the PHY losing lock during Loopback Leader operation which causes the error count to increment by one, and the Loopback Leader restarts the loopback pattern so that the PHY can recover symbol lock.

Field	Name	Access	Width	Reset	Description
[24]	inject_rx_skp_err	RW	1	0x0	When enabled during loopback, the rising edge of inject_rx_skp_err bit injects a single clk error on the next received SKP Ordered Set. When a SKP Ordered Set is corrupted, the lane's RX descrambling LFSR goes out of sync with the transmitter lane's scrambling LFSR causing all the subsequent data checks to fail. This simulates the PHY losing lock and increments the counter by one, and the Loopback Leader restarts the loopback pattern so that the PHY can recover symbol lock.
[23:19]	reserved	RO	5	0x0	—
[18:16]	pattern	RW	3	0x0	Loopback data pattern. 0 – Unscrambled PRBS31 Polynomial Pattern using Galois implementation with non-inverted output. The polynomial representation is $G(x) = X^{31} + X^{28} + 1$.
[15:9]	reserved	RO	7	0x0	—
[8]	tx_comp_receive	RW	1	0x0	Loopback compliance receive behavior. 0 – Loopback Leader does not assert Compliance Receive (recommended default) 1 – Loopback Leader asserts Compliance Receive in TS sets transmitted during Loopback Entry
[7:2]	reserved	RO	6	0x0	—
[1:0]	speed	RW	2	0x0	Desired speed in loopback. Only speeds supported by the core may be programmed. A speed change is only implemented if Loopback is entered from Configuration; if entered from Recovery, the speed is not changed. 0 – 2.5G 1 – 5G 2 – 8G ¹ 3 – 16G ¹

Note:

1. Available in 2024 release.

debug_loopback_master_5g Register 0x108

This register set is used to select Deemphasis values by loopback Leader for 2.5G and 5G data rates.

Table 5.45. debug_loopback_master_5g Register 0x108

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	deemph	RW	1	0x0	Select Deemphasis value used by Loopback Leader when Loopback.Active occurs at 5G data rate. 0 – -6.0dB 1 – -3.5dB

debug_loopback_slave_5g Register 0x10c

This register set is used to select Deemphasis value transmitted in TS sets during loopback for 2.5G and 5G data rates.

Table 5.46. debug_loopback_slave_5g Register 0x10c

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	deemph	RW	1	0x0	Select Deemphasis value transmitted in TS sets for the Follower to use when Loopback.Active occurs at 5G data rate. 0 – -6.0dB 1 – -3.5dB

debug_loopback_master_8g_deemph Register 0x110

This register set is used to select TX Preset related coefficients for 8G data rate for loopback Leader.

Table 5.47. debug_loopback_master_8g_deemph Register 0x110

Field	Name	Access	Width	Reset	Description
[31]	coef_en	RW	1	0x0	Leader coefficient enable. 0 – Otherwise 1 – Direct local transmitter to use coef == mgmt_tlb_debug_loopback_master_8g_deemph_coef when acting as a Loopback Leader in Loopback.Active.
[30:26]	reserved	RO	5	0x0	—
[25:8]	coef	RW	18	0x0	Coefficients to use when mgmt_tlb_debug_loopback_master_8g_deemph_coef_en==1. The coefficients must be a valid set of coefficients, considering the leader's FS and LF values and PCI Express coefficient rules, or the results are undefined. Coefficient mapping: [17:12]==Post-Cursor [11:6]==Cursor [5:0]==Pre-Cursor
[7]	preset_en	RW	1	0x0	Leader preset enable. 0 – Otherwise 1 – Direct local transmitter to use preset == mgmt_tlb_debug_loopback_master_8g_deemph_preset when acting as a Loopback Leader in Loopback.Active
[6:4]	reserved	RO	3	0x0	—
[3:0]	preset	RW	4	0x0	Preset to use when mgmt_tlb_debug_loopback_master_8g_deemph_preset_en==1. Must be a valid preset in range 0x0 to 0xa or the result is undefined.

debug_loopback_slave_8g_deemph Register 0x114

This register set is used to select TX Preset related coefficients for 8G data rate for loopback Follower.

Table 5.48. debug_loopback_slave_8g_deemph Register 0x114

Field	Name	Access	Width	Reset	Description
[31]	coef_en	RW	1	0x0	Follower coefficient enable. 0 – Otherwise 1 – Direct follower's transmitter to use coef == mgmt_tlb_debug_loopback_slave_8g_deemph_coef through TS sets transmitted while directing slave to Loopback.
[30:26]	reserved	RO	5	0x0	—
[25:8]	coef	RW	18	0x0	Coefficients to use when mgmt_tlb_debug_loopback_slave_8g_deemph_coef_en==1. The coefficients must be a valid set of coefficients, considering the follower's FS and LF values and PCI Express coefficient rules, or the follower rejects them. Coefficient mapping: [17:12]==Post-Cursor [11:6]==Cursor [5:0]==Pre-Cursor.
[7]	preset_en	RW	1	0x0	Follower preset enable. 0 – Otherwise 1 – Direct follower's transmitter to use preset == mgmt_tlb_debug_loopback_slave_8g_deemph_preset through TS sets transmitted while directing follower to Loopback.
[6:4]	hint	RW	3	0x0	Follower Rx Hint transmitted in EQ TS1 sets when mgmt_tlb_debug_loopback_slave_8g_deemph_preset_en == 1. PCIe Specification does not indicate what to transmit for RxHint when requesting a Preset through 2.5/5G EQTS1 sets, the follower likely ignores whatever is transmitted in this field.
[3:0]	preset	RW	4	0x0	Follower preset enable. 0 – Otherwise 1 – Direct follower's transmitter to use preset == mgmt_tlb_debug_loopback_slave_8g_deemph_preset through TS sets transmitted while directing follower to Loopback.

debug_direct_to_loopback_status Register 0x118

This register set is used for the Leader loopback status.

Table 5.49. debug_direct_to_loopback_status Register 0x118

Field	Name	Access	Width	Reset	Description
[31:16]	sync	RO	16	0x0	Loopback per lane sync to data pattern indicator. For each lane: 0 – Not locked to loopback pattern. 1 – Locked to loopback pattern.
[15:1]	reserved	RO	15	0x0	—
[0]	cfg_entry	RO	1	0x0	Loopback entered from Configuration or Recovery indicator. 0 – Loopback entry is from Recovery. 1 – Loopback entry is from Configuration.

debug_loopback_err_reset Register 0x11c

This register set is used for the Leader loopback error reset.

Table 5.50. debug_loopback_err_reset Register 0x11c

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	Loopback error counter reset. 0 – Leader Loopback error count increments as errors are detected during Leader Loopback – saturating at maximum value. 1 – Reset the leader loopback error count on all lanes to 0x0. The reset stays in force for as long as mgmt_tlb_debug_loopback_err_reset_enable remains at 1.

debug_loopback_err Register 0x120

This register set is used for the Leader loopback error count.

Table 5.51. debug_loopback_err Register 0x120

Field	Name	Access	Width	Reset	Description
[255:0]	count	RO	256	0x0	Loopback per lane error count – 16 bits per lane. Errors are counted only after the lane is locked to the loopback pattern.

5.1.2.4. Physical control Register Set

phy_control Register 0x140

This register set is used for LTSSM PIPE Interface configuration for 2.5G and 5G data rates.

Table 5.52. phy_control Register 0x140

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	pipe_tx_swing	RW	1	0x0	Directly controls the value of pipe_tx_swing which sets PHY 2.5G/5G Transmitter Amplitude. 0 – Full Swing Full Swing is required for most applications. 1 – Reduced Swing Reduced Swing is useful to support low power form factors which encourage or require reduced transmitter amplitudes.

phy_control_8g Register 0x144

This register set is used for LTSSM 8G PIPE Interface configuration for 8G data rate.

Table 5.53. phy_control_8g Register 0x144

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1]	no_tx_idle_delay	RW	1	0x0	Controls the deassertion of pipe_tx_elec_idle to the PHY when operating with 128b130b encoding. 0 – Deassert pipe_tx_elec_idle at the next data_valid gap. 1 – Deassert pipe_tx_elec_idle at the next ordered set boundary.

Field	Name	Access	Width	Reset	Description
[0]	double_tx_data_valid	RW	1	0x0	Controls the number of consecutive 8G pipe_tx_data_valid deassertions used when compensating for 128b130b encoding differences. 0 – Deassert pipe_tx_data_valid for 1 clock every 64 clocks – the required value for the majority of PHY. 1 – Deassert pipe_tx_data_valid for 2 back-back clocks every 128 clocks, simplifies connecting 8G PHY which have double the per lane width of the controller.

phy_eq_tx_override Register 0x148 F

This register set is used for local PHY transmitter FS/LF override for 8G data rate.

Table 5.54. phy_eq_tx_override Register 0x148 F

Field	Name	Access	Width	Reset	Description
[31:30]	reserved	RO	2	0x0	—
[29:24]	fs	RW	6	0x30	Local PHY Transmitter: Full Scale Value. When 141gmt._tlb_phy_eq_tx_override_enable == 1, 141gmt._tlb_phy_eq_tx_override_fs is used for the local PHY Full Scale (FS) value, otherwise PIPE PHY interface port pipe_local_fs is used.
[23:22]	reserved	RO	2	0x0	—
[21:16]	lf	RW	6	0x8	Local PHY Transmitter: Low Frequency Value. When 141gmt._tlb_phy_eq_tx_override_enable == 1, 141gmt._tlb_phy_eq_tx_override_lf is used for the local PHY Low Frequency (LF) value, otherwise PIPE PHY interface port pipe_local_fs is used.
[15:1]	reserved	RO	15	0x0	—
[0]	enable	RW	1	0x0	Controls whether 141gmt._tlb_phy_eq_tx_override_fs and 141gmt._tlb_phy_eq_tx_override_lf or pipe_local_fs and pipe_local_lf are used to determine the FS and LF values of the local PHY. 0 – Use pipe_local_fs and pipe_local_lf. 1 – Use 141gmt._tlb_phy_eq_tx_override_fs and 141gmt._tlb_phy_eq_tx_override_lf.

phy_eq_tx_max Register 0x14c

This register set is used to specify local PHY maximum allowed coefficient values for 8G data rate.

Table 5.55. phy_eq_tx_max Register 0x14c

Field	Name	Access	Width	Reset	Description
[31:30]	reserved	RO	2	0x0	—
[29:24]	pre	RW	6	0x0	Local PHY transmitter maximum pre-cursor[5:0] coefficient value. If a coefficient request exceeds mgmt_tlb_phy_eq_tx_max_pre, the coefficient is limited to mgmt_tlb_phy_eq_tx_max_pre before being passed to the PHY.
[23:22]	reserved	RO	2	0x0	—

Field	Name	Access	Width	Reset	Description
[21:16]	post	RW	6	0x0	Local PHY transmitter maximum post-cursor [5:0] coefficient value. If a coefficient request exceeds mgmt_tlb_phy_eq_tx_max_post, then the coefficient is limited to mgmt_tlb_phy_eq_tx_max_post before being passed to the PHY.
[15:0]	reserved	RO	16	0x0	—

phy_eq_tx_force Register 0x150

This register set is used to force PHY TX Deemphasis configuration for 8G data rate. This register forces the local PHY TX Deemphasis instead of allowing the link partner to determine the TX Deemphasis during Equalization.

Table 5.56. phy_eq_tx_force Register 0x150

Field	Name	Access	Width	Reset	Description
[31:26]	reserved	RO	6	0x0	—
[25:8]	coef	RW	18	0x0	Coefficients to use when coef_enable==1
[7:4]	preset	RW	4	0x0	Preset to use when preset_enable==1
[3:2]	reserved	RO	2	0x0	—
[1]	coef_enable	RW	1	0x0	Force local PHY pipe_tx_deemph Tx De-Emphasis port. 0 – Determine local PHY ≥ 8G pipe_tx_deemph per PCIe Specification. 1 – Force all local PHY lanes' pipe_tx_deemph output port at ≥ 8G to the coefficients specified by coef. This setting is not PCIe compliant and is intended for debug only.
[0]	preset_enable	RW	1	0x0	Force remote PHY transmitter Tx De-Emphasis to the specified Preset during Equalization Phase 2/3. 0 – Normal operation. 1 – For Figure of Merit Equalization, override the standard 8G Equalization coefficient selection methods and force the core to use the Preset Equalization method with only one Preset == preset. Two total Rx Eq Evaluations are performed, 1 Trial + 1 Final, both using Preset == preset. For Up/Down Equalization, force Equalization to start requesting the link partner change to Preset = preset.

phy_preset_to_coef_conv_control Register 0x15c

This register set is used for local PHY TX preset to coefficient conversion configuration for 8G data rate.

Table 5.57. phy_preset_to_coef_conv_control Register 0x15c

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1:0]	conv_method	RW	2	0x0	Local PHY Transmitter Preset Conversion Method. 0 – Compute using the coefficients from PCIe Specification. PCIe Preset Table 4.3.5.2.2. Tx Equalization Presets. 1 – Lookup table specified by mgmt_tlb_phy_preset_conv_tab_post and mgmt_tlb_phy_preset_conv_tab_pre. 2 – Lookup table obtained from the PIPE PHY using the pipe_local_get_* PIPE interface ports.

phy_preset_conv_tab_pre Register 0x160

This register set is used for pre-cursor coefficients configuration for 8G data rate.

Table 5.58. phy_preset_conv_tab_pre Register 0x160

Field	Name	Access	Width	Reset	Description
[127:66]	reserved	RO	62	0x0	—
[65:0]	coef	RW	66	0xc20c2040000000	For ≥ 8G capable cores only: Pre-cursor table – 6-bit pre-cursor coefficients for each of the 11 possible preset to coefficient conversions packed back-back. A particular preset pre-cursor entry[i] is accessed as [(i*6)+5:(i*6)].

phy_preset_conv_tab_post Register 0x170

This register set is used for post-cursor coefficients configuration for 8G data rate.

Table 5.59. phy_preset_conv_tab_post Register 0x170

Field	Name	Access	Width	Reset	Description
[127:66]	reserved	RO	62	0x0	—
[65:0]	coef	RW	66	0xc20c2040000000	For ≥ 8G capable cores only: Post-cursor table – 6-bit post-cursor coefficients for each of the 11 possible preset to coefficient conversions packed back-back. A particular preset post-cursor entry[i] is accessed as [(i*6)+5:(i*6)].

5.1.2.5. Equalization Configuration Register Set

eq_control Register 0x180

This register set is used for upstream and downstream port preset configuration for 8G data rate.

Table 5.60. eq_control Register 0x180

Field	Name	Access	Width	Reset	Description
[31]	reserved	RO	1	0x0	—
[30:28]	us_port_rx_preset_hint	RW	3	0x2	Rx Preset Hint value that is requested for the Upstream Port to use during the initial stage of Equalization (value transmitted by Downstream Ports in EQ TS2 Ordered Sets) Upstream Ports (Endpoints) do not use this field.
[27:24]	us_port_tx_preset	RW	4	0x4	Tx Preset value that is requested for the Upstream Port to use during the initial stage of Equalization (value transmitted by Downstream Ports in EQ TS2 Ordered Sets) Upstream Ports (Endpoints) do not use this field.
[23]	reserved	RO	1	0x0	—
[22:20]	ds_port_rx_preset_hint	RW	3	0x1	Downstream Port Rx Preset Hint used during initial stage of Equalization. Upstream Ports (Endpoints) uses the value requested by the link partner. If no value is provided by the link partner, Upstream Ports (Endpoints) use this value.

Field	Name	Access	Width	Reset	Description
[19:16]	ds_port_tx_preset	RW	4	0x3	Downstream Port Tx Preset used during initial stage of Equalization. Upstream Ports (Endpoints) uses the value requested by the link partner. If the value requested by the link partner is illegal, or no value is requested by the link partner, Upstream Ports (Endpoints) use this value.
[15:2]	reserved	RO	14	0x0	—
[1]	reset_eieos_interval_count	RW	1	0x0	Reset_EIEOS_Interval_Count :- value transmitted for Reset EIEOS Interval Count in TS1/2 Ordered Set transmissions during appropriate Recovery.Equalization states.
[0]	downstream_eq_skip_phase_2_3	RW	1	0x0	Downstream_Eq_Skip_Phase_2_3 may be set in simulation for Root Port cores to speed simulation since the time-consuming portions of Equalization (Phase 2 and Phase 3) are skipped. 0 – Normal operation (perform all 4 equalization phases) 1 – Skip Equalization Phase 2 and Phase 3 (it is known that full equalization is unnecessary)

eq_ts_control Register 0x184

This register set is used for the Equalization reduced swing configuration for 8G data rate.

Table 5.61. eq_ts_control Register 0x184

Field	Name	Access	Width	Reset	Description
[31:24]	reserved	RO	8	0x0	—
[23:16]	rx_eq_resp_wait	RW	8	0x2	<p>This register determines the number of microseconds the core waits after making an equalization remote PHY Tx coefficient change request before timing out and giving up on getting a TS response from the link partner for that request.</p> <p>This timeout is only used if the link partner fails to acknowledge the request within the timeout window.</p> <p>A timeout may occur, for example, if the local PHY receiver is unable to recover the receive data stream after the link partner changes to the new remote Tx coefficients that were requested. The timeout value is set to $(1.024 \times \text{rx_eq_resp_wait}) \mu\text{s}$. 0 is a special value that selects $(1.024 \times 8) \mu\text{s}$. PHY use EIEOS Rx reception at $\geq 8\text{G}$ and COM symbol reception at $\leq 5\text{G}$ to lock to the data stream.</p> <p>For the PHY to have a good opportunity to recover the data stream, the timeout value chosen must be large enough to include at least several of the Rx OS that the PHY needs to lock to the data stream.</p> <p>At $\geq 8\text{G}$, 1 EIE OS is received every 32 TS OS so 1 EIE OS is received every 33 OS with OS size == 16 symbols.</p> <p>At $\leq 5\text{G}$ a COM symbol is received at the beginning of every OS. rx_eq_resp_wait==8 (8.192 μs) is recommended for most PHY.</p> <p>At 16G, 8.192 μs allows for ~31 EIE OS (8.192 μs / 264 ns) to be received before the timeout occurs.</p> <p>At 8G, 8.192 μs allows for ~15.5 EIE OS (8.192 μs / 528 ns per 1EIE+32TS) to be received before the timeout occurs.</p> <p>At 5G, 8.192 μs allows for ~256 COM symbols (8.192 μs / 32 ns per TS OS) to be received before the timeout occurs.</p> <p>At 2.5G, 8.192 μs allows for ~128 COM symbols (8.192 μs / 64 ns per TS OS) to be received before the timeout occurs.</p>

Field	Name	Access	Width	Reset	Description
[15:8]	ts1_ack_delay	RW	8	0x1f	Defines how long the upstream port (Phase 2) or downstream port (Phase 3) waits after requesting new coefficients/presets before looking for incoming EQ TS1 sets from the remote link partner. This delay by specification should be set to the round-trip delay to the remote link partner (including logic delays in the requesting port) + 500 ns. The delay value used = (eq_ts1_ack_delay [7:0] × 16) + 500 ns. However, 0 is a special value that selects 4.596 microseconds.
[7:6]	request_eq_max_count	RW	2	0x2	Maximum times Request Equalization bit is set in Recovery.RcvrCfg. When set to 2'b00, selects infinite times.
[5]	tx_eq_eval_cnt_sel	RW	1	0x0	Determines the number of clock cycles to wait after the first lane receives an Equalization Tx De-emphasis change request from the link partner until all lanes transmit a change response. The wait time must include the worst-case lane skew that can exist on the lanes as well as time to complete the coefficient computations for the new request. 1 == Wait 127 clocks (conservative). 0 == Wait 8 clocks plus 64 symbols which is: 72 clocks for 8-bit per lane PHY, 40 clocks for 16-bit per lane PHY, or 24 clocks for 32-bit per lane PHY.
[4]	skip_final_coef_check	RW	1	0x0	When set to 1, the Upstream Port skips the check in the Recovery.RcvrLock state after RX Equalization, which compares the TS sets coefficient/preset data with the last requested coefficients/preset from Phase 2 of Equalization. This exception is required for some non-compliant Downstream Port devices.
[3]	ts1_ack_block_use_preset	RW	1	0x1	When set to 1, the <i>use preset</i> bit is always forced to 0 in all EQ TS1 ordered sets transmitted in RX Equalization.
[2]	ts1_ack_mask_use_preset	RW	1	0x1	When set to 1, ignores the state of the use preset bit in incoming EQ TS1 sets when in Phase 2 (upstream port) or Phase 3 (downstream port). When set to 0, the use_preset bit is checked to make sure it matches the setting used in the EQ TS1 sets sent to request remote transmitter settings. In either setting, the preset value or coefficient values are compared as required by the PCIe 3.0 Specification. The value of this bit MUST be 1 for proper PCIe operation.
[1]	early_rx_eval	RW	1	0x0	When set to 1, the local PHY is told to evaluate the RX serial signals (using RxEval) BEFORE checking incoming TS1 sets for acknowledgment after each new request. When 0, the local PHY is told to evaluate the RX serial signals (using RxEval) AFTER checking incoming TS1 sets for acknowledgment after each new request. The 0 state is the default recommended state.
[0]	no_remote_change	RW	1	0x0	When set to 1, the equalization algorithm monitors the advertised coefficients from the Link Partner before equalization starts, and requests the same coefficients when performing equalization, so that the link partner does not change the TX coefficients during equalization. This is primarily for debugging purposes.

eq_reduced_swing Register 0x188

This register set is used for the Equalization reduced swing configuration for 8G data rate.

Table 5.62. eq_reduced_swing Register 0x188

Field	Name	Access	Width	Reset	Description
[31:19]	reserved	RO	13	0x0	—
[18:8]	preset_reject	RW	11	0x0	Specifies which presets (if any) are rejected if requested by the remote link partner during 8G RX Equalization. A request to use preset[i] is rejected if mgmt_tlb_eq_reduced_swing_en==1 and mgmt_tlb_eq_reduced_swing_preset_reject[i]==1 and is otherwise accepted. Per PCIe Specification, all preset requests with a valid preset value (0x0 to 0xA) must be accepted with Full Swing, but selected presets may be rejected when implementing reduced swing.
[7:1]	reserved	RO	7	0x0	—
[0]	en	RW	1	0x0	Reduced swing support enable. When mgmt_tlb_eq_reduced_swing_en==1, a request to transmit with preset[i] is rejected as illegal if mgmt_tlb_eq_reduced_swing_preset_reject[i] == 1. Only affects the acceptance/rejection of preset requests. It is also necessary to use mgmt_tlb_phy_control_pipe_tx_swing to configure the PHY transmitter for reduced swing.

eq_method Register 0x1bc

This register set is used to select Equalization method for 8G data rate.

Table 5.63. eq_method Register 0x1bc

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	select_dir_fom_n	RW	1	0x0	The core implements two primary methods: Figure of Merit and Up/Down. This register chooses which major algorithm is used. Each major algorithm has its own CSR registers for control/status Primary Equalization Method Selection. 0 – Use Figure of Merit Equalization Methods 1 – Use Up/Down Equalization Methods

eq_fmerit_control Register 0x1c0

This register set is used for Equalization Figure of Merit Method configuration for 8G data rate.

Table 5.64. eq_fmerit_control Register 0x1c0

Field	Name	Access	Width	Reset	Description
[31:24]	req_feedback	RW	8	0x80	When phy_eq_rx_eval_f_merit ≥ eq_req_feedback, the link is BER 10 ⁻¹² or better.
[23:2]	reserved	RO	22	0x0	—

Field	Name	Access	Width	Reset	Description
[1:0]	method	RW	2	0x3	Equalization Method Selection. 0 – Step through PCIe-defined Tx Presets 1 – Evenly step through the coefficients range 2 – Step through the user-provided coefficient table 3 – Step through the user-provided coefficient table with adaptive coefficient selection

eq_preset_method_control Register 0x1c4

This register set is used for Equalization Figure of Merit Preset Method configuration for 8G data rate.

Table 5.65. eq_preset_method_control Register 0x1c4

Field	Name	Access	Width	Reset	Description
[31:12]	reserved	RO	20	0x0	—
[11:8]	addr_limit	RW	4	0x4	Last preset to use. The Preset Algorithm steps through PCIe-Specification-defined Tx Equalization Presets from 0 to eq_preset_addr_limit; eq_preset_addr_limit has a maximum value of 9 (step through presets 0-9).
[7:1]	reserved	RO	7	0x0	—
[0]	use_coef	RW	1	0x0	Controls whether the presets are communicated to the remote device using (1) the associated coefficient values or (0) the Preset value. 1 – The LTSSM requests coefficient values that are equivalent to the desired presets in the C-1, C0, and C+1 fields of TS1 Ordered Sets. 0 – LTSSM requests the desired presets using the Use Preset and Tx Preset fields of TS1 Ordered Sets.

eq_alg_method_control Register 0x1c8

This register set is used for Equalization Figure of Merit Algorithm Method configuration for 8G data rate.

Table 5.66. eq_alg_method_control Register 0x1c8

Field	Name	Access	Width	Reset	Description
[31:30]	reserved	RO	2	0x0	—
[29:24]	post_cursor_step_size	RW	6	0x8	Step size to use when walking through Post-Cursor coefficient values.
[23:22]	reserved	RO	2	0x0	—
[21:16]	pre_cursor_step_size	RW	6	0x4	Step size to use when walking through Pre-Cursor coefficient values.
[15:14]	reserved	RO	2	0x0	—
[13:8]	post_cursor_limit	RW	6	0x20	Upper bound on the Post-Cursor coefficient values to try. Permissible values are 0-32 (0 to 0.5).
[7:6]	reserved	RO	2	0x0	—
[5:0]	pre_cursor_limit	RW	6	0x10	Upper bound on the Pre-cursor coefficient values to try. Permissible values are 0-16 (0 to 0.25).

eq_table_method_control Register 0x1cc

This register set is used for Equalization Figure of Merit Table Method and Adaptive Table Method configuration for 8G data rate.

Table 5.67. eq_table_method_control Register 0x1cc

Field	Name	Access	Width	Reset	Description
[31:9]	reserved	RO	23	0x0	—
[8]	end_on_hold	RW	1	0x1	When the Adaptive Table Method is selected, determines whether to exit early if a Prior PHY Up/Down Feedback response == {HOLD, HOLD} is received on all lanes while processing a Table Entry with interpret==3. 0 – NoExitOnHold 1 – YesExitOnHold
[7:5]	reserved	RO	3	0x0	—
[4:0]	addr_limit	RW	5	0x11	Last table entry to use for the Table Method and Adaptive Table Method. The Table method walks through table entries from $i == 0$ to $eq_table_addr_limit$ selecting $eq_table_method_pre_cursor[(i*6)+5:(i*6)]$ as the pre-cursor coefficient and $eq_table_method_post_cursor[(i*6)+5:(i*6)]$ as the post-cursor coefficient for each Equalization evaluation trial.

eq_table_method_table Register 0x1d0

This register set is used to set the table array for the Equalization Figure of Merit Table Method and Adaptive Table Method for 8G data rate.

Table 5.68. eq_table_method_table Register 0x1d0

Field	Name	Access	Width	Reset	Description
[383:0]	array	RW	384	0x300030003000300030003000300030003000300030003000400700080009000600050004	For $\geq 8G$ capable cores only. The array consists of 16-bit table entries, packed back-back, for each of the 24 implemented table entries. Table entries are used to configure the Rx Equalization algorithms Table Method and Adaptive Table Method.

Each 16-bit table entry [i] is in the following format:

- $array[(i*16)+15]$ – reserved
- $array[(i*16)+14]$ – best
- $array[(i*16)+13:(i*16)+12]$ – Interpret[1:0]
- $array[(i*16)+11:(i*16)+6]$ – post[5:0]
- $array[(i*16)+5:(i*16)+0]$ – pre [5:0]

When the Table Method is selected:

- When $interpret[1:0] == 00$, use the coefficients corresponding to $Preset[pre[3:0]]$.
- When $interpret[1:0] == 01$, $pre[5:0]$ is the desired pre-cursor coefficient and $post[5:0]$ is the desired post-cursor coefficient. best is unused

When the Adaptive Table Method is selected:

- When $interpret[1:0] == 00$, use the coefficients corresponding to $Preset[pre[3:0]]$.
- When $interpret[1:0] == 01$, $pre[5:0]$ is the desired pre-cursor coefficient and $post[5:0]$ is the desired post-cursor coefficient.
- When $interpret[1:0] == 10$, $pre[5:0]$ is the relative pre-cursor offset from the current relative best coefficient pair {rel_best_post, rel_best_pre} and $post[5:0]$ is the relative post-cursor offset from the current relative best coefficient pair {rel_best_post, rel_best_pre}.

- When `interpret[1:0] == 11`, apply the prior PHY Up/Down Feedback result to the current best coefficient pair {`best_post`, `best_pre`} and make the result the new best coefficient pair. `best==1` instructs the core to set {`rel_best_post`, `rel_best_pre`} to the coefficient pair that returned the highest Figure of Merit from among all the coefficient pairs tried since the beginning of Rx Equalization (that is, include the currently executing table entry through table entry 0 inclusive).

eq_updn_control Register 0x240

This register set is used for the Equalization Up/Down Convergence Method configuration for 8G data rate.

Table 5.69. eq_updn_control Register 0x240

Field	Name	Access	Width	Reset	Description
[31:26]	iteration_max	RW	6	0x0	The iteration count at which Up/Dn Equalization should assume convergence if not yet converged. A 0 value indicates no limit.
[25]	start_remote_adv	RW	1	0x0	If (1), then the initial pre-cursor and post-cursor coefficients requested from the remote link partner are initialized by the coefficients advertised by the link partner in received TS sets. This option is available to simplify the selection of the initial coefficient values by using the initial coefficients decided determined to be an optimal starting case for the link partner.
[24]	fail_limit_err	RW	1	0x0	If (1), then if an <i>up</i> response is received while a coefficient is at its maximum legal value, this is considered an error which causes RX Equalization to fail. If (0), it does not cause RX Equalization to fail, but the coefficient value is limited by the maximum value. This is also true for cases where a <i>down</i> response is received for a coefficient set to the minimum value (0).
[23:18]	reserved	RO	6	0x0	—
[17:16]	post_step	RW	2	0x2	Initial step size for changing the post-cursor coefficient values. 0 – Step Size 1 1 – Step Size 2 2 – Step Size 4 3 – Step Size 8
[15:10]	reserved	RO	6	0x0	—
[9:8]	pre_step	RW	2	0x1	This field defines the initial step size for changing the pre-cursor coefficient values based on the PHY up/dn response. 0 – Step Size 1 1 – Step Size 2 2 – Step Size 4 3 – Step Size 8
[7:3]	reserved	RO	5	0x0	—
[2]	use_coef	RW	1	0x0	Controls whether the presets are communicated to the remote device using (1) the associated coefficient values or (0) the Preset value. When (1) the LTSSM requests coefficient values that are equivalent to the desired presets in the C-1, C0, and C+1 fields of TS1 Ordered Sets. When (0) the LTSSM requests the desired presets using the Use Preset and Tx Preset fields of TS1 Ordered Sets.
[1]	start_preset	RW	1	0x1	If set to (1) then the initial pre-cursor and post-cursor coefficients requested from the remote link partner are controlled as a preset value, which is loaded into <code>eq_pre_cursor_laneX[3:0]</code> . This option is available to simplify the selection of the initial coefficient values by using the preset settings specified in the PCIe Specification.

Field	Name	Access	Width	Reset	Description
[0]	numhold	RW	1	0x0	If set to (1), then two consecutive hold status responses are needed for each coefficient to indicate that the coefficient is at an optimal setting. This is required for PHYs which update only one coefficient on each response and set to the other coefficient response to hold. If set to (0), then a single hold response is sufficient to consider a coefficient at an optimal setting.

eq_firmware_control Register 0x280

This register set is used for Equalization Up/Down Firmware Controlled Method configuration for 8G data rate.

Table 5.70. eq_firmware_control Register 0x280

Field	Name	Access	Width	Reset	Description
[31:7]	reserved	RO	25	0x0	—
[6]	tx_quiesce	RW	1	0x0	Transmit Quiesce Guarantee Control Value of Quiesce Guarantee in transmitted TS2 sets in Recovery.RcvrCfg when in the DL_ACTIVE Data Link Layer state (which is L0 and having exchanged Flow Control DLLPs successfully). Quiesce Guarantee is set to 1 in Detect.Quiet and is set to tx_quiesce when in DL_ACTIVE. Quiesce Guarantee holds the last value to which it is set. When the Upstream Port is requesting Equalization be re-run, the Upstream Port is permitted, but not required, to set Quiesce Guarantee == 1 to inform the Downstream Port that an equalization process initiated within 1 ms does not cause any side-effects to its operation.
[5]	tx_req_eq	RW	1	0x0	Transmit Link Equalization Request Control. Set to 1 to set Link Equalization Request == 1 in transmitted TS2 sets in Recovery.RcvrCfg. May be set to 1 by an Upstream Port to inform the Downstream Port link partner (which is the only port that is able to re-run Equalization) that the Upstream Port wants to re-run Equalization. This is a hint to the Downstream Port and the Downstream Port may or may not re-run Equalization.
[4]	int_clr	RW	1	0x0	Writing a 1 to this register clears any pending Equalization interrupts in eq_status – rx_quiesce_hold eq_status – rx_req_eq_hold eq_status – int
[3]	int_en	RW	1	0x0	Local Rx Equalization Interrupt Enable. Enables interrupts to be generated for Rx Equalization events. An interrupt is generated on Entry to Local Rx Equalization Phase 2 (US Port) or Phase 3 (DS Port) and after the PHY Response is received for each Local Rx Equalization trial.
[2]	complete	RW	1	0x0	Firmware Controlled Equalization – Complete. Set to 1 to instruct the Local Rx Firmware Controlled Equalization algorithm to consider Rx Equalization complete. Complete is set to 1, instead of setting advance to 1, when no new equalization step is needed, and Equalization can be considered complete. 0 – NotComplete 1 – Complete

Field	Name	Access	Width	Reset	Description
[1]	advance	RW	1	0x0	Firmware Controlled Equalization – Advance. Set to 1 to instruct the Local Rx Firmware Controlled Equalization algorithm to try the new coefficients specified in the eq_pre_cursor and eq_post_cursor registers. When the core completes a trial, the core stores the resulting PHY Figure of Merit and Directional feedback and waits until new coefficients are loaded by Firmware and Firmware sets the advance register to start another trial or else Firmware sets the complete register to end Equalization. 0 – DoNotAdvance 1 – Advance
[0]	ext_control	RW	1	0x0	Specifies whether the local PHY RX Equalization is under Firmware or local hardware control. This option is only supported when the Equalization Up/Down method is selected (eq_method:select_dir_fom_n == 1). 0 – HardwareControl 1 – FirmwareControl

eq_pre_cursor Register 0x290

This register set is used to set pre-cursor coefficients for the Equalization Up/Down Firmware Controlled Method for 8G data rate.

Table 5.71. eq_pre_cursor Register 0x290

Field	Name	Access	Width	Reset	Description
[95:0]	coef	RW	96	0x0	For ≥ 8G capable cores only, which are delivered with PHY implementing Up/Down Equalization Feedback -Pre-cursor table – 6-bit pre-cursor coefficients for each of the 16 possible lanes packed back-back. A particular lane entry[i] is accessed as [(i*6)+5:(i*6)].

eq_post_cursor Register 0x2a0

This register set is used to set post-cursor coefficients for the Equalization Up/Down Firmware Controlled Method for 8G data rate.

Table 5.72. eq_post_cursor Register 0x2a0

Field	Name	Access	Width	Reset	Description
[95:0]	coef	RW	96	0x0	For ≥ 8G capable cores only, which are delivered with PHY implementing Up/Down Equalization Feedback – Post-cursor table – 6-bit post-cursor coefficients for each of the 16 possible lanes packed back-back. A particular lane entry[i] is accessed as [(i*6)+5:(i*6)].

eq_status Register 0x2c0

This register set is used to read the status registers for the Equalization Up/Down Firmware Controlled Method for 8G data rate.

Table 5.73. eq_status Register 0x2c0

Field	Name	Access	Width	Reset	Description
[31:7]	reserved	RO	25	0x0	—
[6]	rx_quiesce_hold	RO	1	0x0	<p>When the Request Equalization bit is set on ingress TS2 sets at 8G speed, mgmt_tlb_eq_status_rx_quiesce_hold is set to the value of the Quiesce Guarantee bit in those same TS2 sets.</p> <p>The Quiesce Guarantee bit indicates that it is safe for the remote link partner's application to re-run RX Equalization, which can take tens of milliseconds.</p> <p>mgmt_tlb_eq_status_rx_quiesce_hold is cleared when mgmt_tlb_eq_firmware_control_int_clr==1.</p> <p>Refer to the PCI Express 3.0 Spec, Section 4.2.3 for details on how to use this information.</p>
[5]	rx_req_eq_hold	RO	1	0x0	<p>When the Request Equalization bit is set on ingress TS2 sets at 8G speed, mgmt_tlb_eq_status_rx_req_eq_hold is set to 1.</p> <p>This indicates that the remote link partner is requesting RX Equalization be re-run.</p> <p>mgmt_tlb_eq_status_rx_req_eq_hold is cleared when mgmt_tlb_eq_firmware_control_int_clr==1.</p> <p>See the PCI Express 3.0 Spec, Section 4.2.3 for details on how to use this information.</p>
[4]	int_edge	RO	1	0x0	<p>Pulsed interrupt signal. Indicates an interrupt is signaled with a 1 clock wide pulse.</p> <p>Used for state machine logic, not applicable for CSR register reads or polling.</p>
[3]	err	RO	1	0x0	Indicates that an error occurred during RX Equalization.
[2]	exit	RO	1	0x1	Indicates that RX Equalization has not started or has completed.
[1]	ready	RO	1	0x0	Indicates the RX Equalization process requires a response from the external control process.
[0]	int	RO	1	0x0	<p>Rx Equalization Interrupt Status.</p> <p>int==1 indicates an interrupt is active.</p> <p>int is set to 1 when mgmt_tlb_eq_status_ready is set to 1 indicating that the Firmware Controlled Equalization algorithm is ready for new coefficients to try.</p> <p>int is also set to 1 when mgmt_tlb_eq_status_rx_req_eq_hold is set to 1 indicating that the link partner set Link Equalization Request==1 in its transmitted TS2 OS in Recovery.RcvrCfg to request that Equalization be re-run.</p> <p>int=1 is cleared to 0 when eq_firmware_control:advance is set to 1 indicating that new coefficients have been provided or when eq_firmware_control:complete is set to 1 indicating that Equalization is complete, or when eq_firmware_control:int_clr is written to 1.</p>

eq_status_error Register 0x2c4

This register set is used to read the error status registers for the Equalization Up/Down Firmware Controlled Method for 8G data rate.

Table 5.74. eq_status_error Register 0x2c4

Field	Name	Access	Width	Reset	Description
[31:16]	lane_error	read-only	16	0x0	Per-lane equalization error status. Errors can be caused either by the coefficients being rejected by the remote link partner (which is only permitted when certain presets are used in reduced-swing mode) or if a response is not detected by the remote link partner (which might be caused by extremely low link signal quality). For each lane: 0 – No Error 1 – Error
[15:0]	lane_active	read-only	16	0x0	Per-lane active indicator. Only lanes that are active (trained to be part of the link) can be expected to participate in equalization. For each lane: 0 – Inactive 1 – Active

eq_status_preset_coef Register 0x2c8

This register set is used to read the preset coefficient registers for the Equalization Up/Down Firmware Controlled Method for 8G data rate.

Table 5.75. eq_status_preset_coef Register 0x2c8

Field	Name	Access	Width	Reset	Description
[31:16]	lane_match	RO	16	0x0	Per-lane Coefficients or Preset Match Last Request. Indicates that in TS sets received in Recovery.RcvrLock, coefficients or preset fields matched those of the last coefficient or preset request. 0 – No Error 1 – Error
[15:0]	lane_match_valid	RO	16	0x0	Per-lane Remote PHY Match Field Valid. Indicates that the last time through Recovery.RcvrLock, RX Equalization phases 2 and 3 had completed, and this lane received 8 consecutive 8G EC00 TS Sets. 0 – Inactive 1 – Active

eq_status_feedback_fom Register 0x2d0

This register set is used to read the per lane FOM from local PHY for the Equalization Figure of Merit Method for 8G data rate.

Table 5.76. eq_status_feedback_fom Register 0x2d0

Field	Name	Access	Width	Reset	Description
[127:0]	value	RO	128	0x0	Per-lane Figure of Merit Equalization feedback received from the local PHY – 8-bits per lane. value[(i*8)+7:(i*8)] is the measure of Equalization quality for Lane[i] with higher values indicating better BER.

eq_status_feedback_dir Register 0x2e0

This register set is used to read the per lane pre-cursor and post-cursor values from local PHY for the Equalization Up/Down Feedback Method for 8G data rate.

Table 5.77. eq_status_feedback_dir Register 0x2e0

Field	Name	Access	Width	Reset	Description
[63:0]	value	RO	164	0x0	Per-lane Up/Down Equalization feedback received from the local PHY – 4-bits per lane. value[(i*4)+3:(i*4)+2] is the post-cursor feedback and [(i*4)+1:(i*4)+0] is the pre-cursor feedback for Lane[i] with feedback encoded as:- 00==NoChange/Hold, 01==Increment, 10==Decrement, and 11==Reserved.

eq_status_remote_fs Register 0x2e8

This register set is used to read the per lane FS value of remote link partner during RX Equalization for 8G data rate.

Table 5.78. eq_status_remote_fs Register 0x2e8

Field	Name	Access	Width	Reset	Description
[95:0]	value	RO	96	0x0	Per-lane FS value advertised by the remote link partner in TS Sets received in Recovery.Equalization.Phase1. The FS value for lane[i] is value[(i*6)+5:(i*6)].

eq_status_remote_lf Register 0x2f4

This register set is used to read the per lane LF value of remote link partner during RX Equalization for 8G data rate.

Table 5.79. eq_status_remote_lf Register 0x2f4

Field	Name	Access	Width	Reset	Description
[95:0]	value	RO	96	0x0	Per-lane LF value advertised by the remote link partner in TS Sets received in Recovery.Equalization.Phase1. The LF value for lane[i] is value[(i*6)+5:(i*6)].

eq_status_remote_precursor Register 0x300

This register set is used to read the per lane PHY pre-cursor coefficient of remote link partner during RX Equalization for 8G data rate.

Table 5.80. eq_status_remote_precursor Register 0x300

Field	Name	Access	Width	Reset	Description
[95:0]	value	RO	96	0x0	Per-lane Remote PHY Pre-Cursor Coefficient. Indicates the remote device pre-cursor coefficient advertised in 8G TS Sets received in the last time through Recovery.RcvrLock. The precursor for lane[i] is value[(i*6)+5:(i*6)].

eq_status_remote_postcursor Register 0x30c

This register set is used to read the per lane PHY post-cursor coefficient of remote link partner during RX Equalization for 8G data rate.

Table 5.81. eq_status_remote_postcursor Register 0x30c

Field	Name	Access	Width	Reset	Description
[95:0]	value	RO	96	0x0	Per-lane Remote PHY Post-Cursor Coefficient. Indicates the remote device post-cursor coefficient advertised in 8G TS Sets received in the last time through Recovery.RcvrLock. The postcursor for lane[i] is value[(i*6)+5:(i*6)].

5.1.2.6. Physical Layer Register Set

pl_rx Register 0x33c

This register set is used for the Physical Layer Per Lane Receiver Control.

Table 5.82. pl_rx Register 0x33c

Field	Name	Access	Width	Reset	Description
[31:17]	reserved	RO	15	0x0	—
[16]	inject_data_error_en	RW	1	0x0	pipe_rx_data error injection for test purposes. inject_data_error_mask[NUM_LANES-1:0] controls which lane(s) is used for error injection. For example, to enable a bit error on lane 2, first write 0x04 to inject_data_error_mask, next write a 1 to inject_data_error_en, then finally write 0 to inject_data_error_en. This is a test only feature which is not used for normal operation. inject_data_error_en must be set to 0 whenever error injections are not desired. The write to this field must occur after writes to the other fields in this register have established the desired parameters of the injection.
[15:0]	inject_data_error_mask	RW	16	0x0	Lane select for data error injection. Bit 0 corresponds to lane 0, Bit 1 corresponds to lane 1. Setting the bit for a corresponding lane to 1 result in an error being injected on that lane when inject_data_error_en changes from 0 to 1. Always setup inject_data_error_mask before setting inject_data_error_en to 1. This field may not be changed while inject_data_error_en==1.

pl_16g Register 0x340

This register set is used for the Physical Layer 16G Optional Behavior Enable.

Table 5.83. pl_16g Register 0x340

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable_16g_eie_same_as_8g	RW	1	0x0	Determines whether 16G EIEOS use Draft 0.7 and higher OS definition or Draft 0.5 and lower OS definition (which is the same as 8G). 0 – Use PCIe 4.0 Draft 0.7 and higher 16G EIEOS definition – 0xFFFF0000FFFF0000FFFF0000FFFF0000 1 – Use PCIe 4.0 Draft 0.5 and lower 16G EIEOS definition – 0xFF00FF00FF00FF00FF00FF00FF00FF00

pl_tx_skp Register 0x344

This register set is used for the Physical Layer Transmit SKP Period Control.

Table 5.84. pl_tx_skp Register 0x344

Field	Name	Access	Width	Reset	Description
[31:30]	reserved	RO	2	0x0	—
[29:24]	period_sris_128b130b	RW	6	0x0	The transmit SKP period used when operating at $\geq 8G$ with the SRIS capability enabled and configured for $SRIS = period_sris_128b130b$ Blocks. PCIe Specification is < 38 blocks. 0 is a special case that selects 36 Blocks. This register must be configured for a PCIe Specification compliant value.
[23:16]	period_srns_128b130b	RW	8	0x0	The transmit SKP period used when operating at $\geq 8G$ with the SRIS capability disabled or with SRIS enabled but configured for $SRNS = 256 + period_srns_128b130b$ Blocks. PCIe Specification is 370-375 blocks. 0 is a special case that selects 116 == 372 Blocks. This register must be configured for a PCIe Specification compliant value.
[15:8]	period_sris_8b10b	RW	8	0x0	The transmit SKP period used when operating at $\leq 5G$ with the SRIS capability enabled and configured for $SRIS = period_sris_8b10b$ Symbol Times. PCIe Specification is < 154 Symbol Times. 0 is a special case that selects 146 Symbol Times. This register must be configured for a PCIe Specification compliant value. The number of symbol times selected must be a multiple of the PHY per lane symbol data width or the lower bits are truncated. For example: For 16-bit per lane PHY, $period_sris_8b10b[0]$ is always treated as 0. For 32-bit per lane PHY, $period_sris_8b10b[1:0]$ are always treated as 00. For 64-bit per lane PHY, $period_sris_8b10b[2:0]$ are always treated as 000.
[7:0]	period_srns_8b10b	RW	8	0x0	The transmit SKP period used when operating at $\leq 5G$ with the SRIS capability disabled or with SRIS enabled but configured for $SRNS = (256 + period_srns_8b10b) * 4$ Symbol Times. PCIe Specification is 1180-1538 Symbol Times. 0 is a special case that selects 44 == 1200 Symbol Times. This register must be configured for a PCIe Specification compliant value. The number of symbol times selected must be a multiple of the PHY per lane symbol data width or the lower bits are truncated. For example: For 64-bit per lane PHY $period_srns_8b10b[0]$ is always treated as 0. For 32, 16, and 8-bit per lane PHY, all of $period_srns_8b10b$ is relevant.

pl_tx_debug Register 0x348

This register set is used for the Physical Layer Debug Control.

Table 5.85. pl_tx_debug Register 0x348

Field	Name	Access	Width	Reset	Description
[31:3]	reserved	RO	29	0x0	—
[2]	inject_margin_crc_error	RW	1	0x0	Setting this to 1 injects errors into the margin crc value of the control skp ordered set on all lanes
[1]	inject_margin_parity_error	RW	1	0x0	Setting this to 1 injects errors into the margin parity bit of the control skp ordered set on all lanes
[0]	inject_data_parity_error	RW	1	0x0	Setting this to 1 injects errors into the data parity bit of the control skp ordered set on all lanes

pl_ctrl Register 0x34c

This register set is used for the Physical Layer Control.

Table 5.86. pl_ctrl Register 0x34c

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	8b10b_err_rec_entry_sel	RW	1	0x0	Selects the Physical Layer error threshold required to be received in L0, when operating with 8b10b encoding (2.5G/5G speed), before the link is directed to Recovery. 0 – When in L0 and using 8b10b encoding, direct the link to Recovery after receiving a single Physical Layer Error. This is the more conservative setting but has the disadvantage of causing Recovery entry on all L0 Physical Layer errors – even those errors that the link would be able to recover from on its own without having to go through Recovery. 1 – When in L0 and using 8b10b encoding, direct the link to Recovery only after receiving a massive burst of errors (which typically is an indication that there is a persistent problem, such as PHY loss of lock, for which Recovery entry is required to fix). The core implements an error counter. For each enabled PHY Rx clock cycle, the error counter is incremented when a clock cycle contains a Physical Layer error, and the error counter is decremented when a clock cycle contains no Physical Layer errors. If the counter reaches 1023, the link is directed to Recovery.

pl_ts_matching Register 0x350

This register set is used for the Physical Layer TS Match Control.

Table 5.87. pl_ts_matching Register 0x350

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	legacy_mode	RW	1	0x1	Setting this to 1 compares all symbols when matching TS sets (legacy behavior). Setting this to 0 compares only the symbols required to meet specifications.

5.1.2.7. Data Link Layer Control Register Set

dl_retry_timeout Register 0x380

This register set is used for the Replay Timeout Control.

Table 5.88. dl_retry_timeout Register 0x380

Field	Name	Access	Width	Reset	Description
[31:24]	pcie4_symt_sync	RW	8	0x0	Replay Timeout Timer PCIe 4.0 Simplified REPLAY_TIMER Limits Extended Sync==1 Value. {pcie4_symt_sync, 10'h0} = Symbol times to use for Replay Timer when pcie4_enable==1 and Extended Synch==1. 0 is a special case selecting 8'd80.
[23:16]	pcie4_symt_sync_n	RW	8	0x0	Replay Timeout Timer PCIe 4.0 Simplified REPLAY_TIMER Limits Extended Sync==0 Value. {pcie4_symt_sync_n, 10'h0} = Symbol times to use for Replay Timer when pcie4_enable==1 and Extended Synch==0. 0 is a special case selecting 8'd24.
[15]	pcie4_enable	RW	1	0x1	Replay Timeout Timer PCIe 4.0 Simplified REPLAY_TIMER Limits Enable. 0 – Use PCIe 3.0 Specification and prior REPLAY_TIMER Limits from UNADJUSTED REPLAY_TIMER LIMITS tables in the PCIe Specification. 1 – Use PCIe 4.0 Specification Simplified REPLAY_TIMER Limits.
[14:1]	l0s_adj	RW	14	0x180	Replay Timeout L0s Adjustment. The number of symbol times to add to the recommended PCIe Replay Timer timeout period to compensate for the remote link having to exit L0s before it can send an ACK/NAK DLLP. l0s_adj is added to the Replay Timer timeout period after the optional doubling controlled by mult_enable is applied. Not applicable when pcie4_enable==1.
[0]	mult_enable	RW	1	0x0	Replay Timeout Timer Multiplier Enable. Not applicable when pcie4_enable==1. 0 – The Replay Timer timeout period implemented is the recommended (-0%) value in the PCIe Specification UNADJUSTED REPLAY_TIMER LIMITS FOR 2.5/5.0/8.0 GT/S MODE OPERATION BY LINK WIDTH AND MAX_PAYLOAD_SIZE tables. 1 – The Replay Timer timeout period implemented is 2 times the recommended (-0%) value in the PCIe Specification UNADJUSTED REPLAY_TIMER LIMITS FOR 2.5/5.0/8.0 GT/S MODE OPERATION BY LINK WIDTH AND MAX_PAYLOAD_SIZE tables.

dl_ack_timeout_div Register 0x384

This register set is used for the ACK Timer Control.

Table 5.89. dl_ack_timeout_div Register 0x384

Field	Name	Access	Width	Reset	Descriptions
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	ACK Timer Control. 0 – Ack according to specifications. 1 – Ack twice as often as recommended by specifications.

dl_tx_ctrl Register 0x38c

This register set is used for the Data Link Layer TX Control.

Table 5.90. dl_tx_ctrl Register 0x38c

Field	Name	Access	Width	Reset	Descriptions
[31]	stp_override_en	RW	1	0x0	<p>8G STP Symbol Debug Length Override Enable.</p> <p>Enables the injection of 8G STP symbol length errors (for debug only).</p> <p>On a rising edge of stp_override_en, a request to inject an 8G STP length error is stored until it can be acted upon.</p> <p>The next TLP with a computed 8G STP length == stp_override_len has its STP token length field replaced by stp_override_new_len instead of using the computed value.</p> <p>Per PCIe Specification, the 8G STP token length must consider the full framed TLP length including 8G STP Token, TLP Header, TLP Payload, ECRC (if present), and LCRC, but must not include the EDB symbol (if the TLP is being nullified).</p> <p>These fields enable you to transmit an 8G TLP of incorrect length by placing the EOP at the end of the TLP data to transmit and then writing these fields to substitute the desired STP length matching the TLP being transmitted for the actual length of the TLP that would be computed from the TLP header.</p> <p>At 2.5 and 5G, these fields are unused as the full framed TLP length is not included in the STP token at these speeds and a malformed length TLP can be transmitted just by placing EOP at the incorrect location.</p> <p>0 – Disabled 1 – Enabled</p>
[30:27]	reserved	RO	4	0x0	—
[26:16]	stp_override_new_len	RW	11	0x0	8G STP Symbol Debug Length Replacement Value.
[15:11]	reserved	RO	5	0x0	—
[10:0]	stp_override_len	RW	11	0x0	8G STP Symbol Debug Length Replacement Match.

dl_ctrl Register 0x390

This register set is used for the Data Link Layer Control.

Table 5.91. dl_ctrl Register 0x390

Field	Name	Access	Width	Reset	Descriptions
[31:26]	reserved	RO	6	0x0	—
[25]	tx_pfx_par_inject_en	RW	1	0x0	Transmit Data Link Layer Prefix Parity Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single prefix parity error injection, applied prior to assigning the TLP sequence number, is scheduled and injected at the next opportunity (TLP transmit). Tx prefix parity error Handling and Reporting are governed by tx_par2_handle_disable and tx_par2_report_disable.
[24]	rx_early_forward_disable	RW	1	0x0	Receive Data Link Layer down-trained early forwarding disable. 0 – When down-trained, forward Rx Data Link Layer data for processing whenever a TLP/DLLP end occurs without a following TLP/DLLP start the same clock cycle. This setting results in lower Rx TLP/DLLP latency. 1 – When down-trained, always aggregate Rx Data Link Layer data to full width before forwarding the data. For example, a x16 core operating at x1 receives and aggregate 16 clock cycles of 1 lane data before outputting one clock cycle of 16 lane data for further processing.
[23]	reserved	RO	1	0x0	—
[22]	tx_gap_inject_en	RW	1	0x0	Transmit Data Link Layer TX Valid Gap Injection Enable. 0 – Do not inject gap. 1 – On the rising edge, a single clock bp_tx_valid gap is scheduled and is injected at the next opportunity (within a TLP). This gap in the bp_tx_valid can cause a data underflow at the Physical Layer.
[21]	rx_malf_inject_en	RW	1	0x0	Receive Data Link Layer Malformed Length TLP Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single malformed TLP error injection is scheduled and is injected at the next opportunity (Tx TLP EOP). The TLP is malformed by deleting its end of TLP indicator causing the TLP to end at the incorrect location.
[20]	rx_lcrc_inject_en	RW	1	0x0	Receive Data Link Layer LCRC Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single LCRC TLP error injection is scheduled and injected at the next opportunity (Tx TLP EOP). The LCRC is corrupted by inverting LCRC bit 0 of the received TLP.
[19]	reserved	RO	1	0x0	—

Field	Name	Access	Width	Reset	Descriptions
[18]	rx_dl_active_disable	RW	1	0x0	Control the use of DL_Active to block reception of TLPs. 0 – Block reception of TLPs when dl_active is low. 1 – Do not block TLP reception based on dl_active.
[17]	rx_inhibit_tlp	RW	1	0x0	Receive Data Link Layer TLP Rx Inhibit Enable. 0 – Process received TLPs per PCIe Specification Required setting for compliant PCIe operation. 1 – For test purposes only, discard and do not accept received TLPs. Received TLPs are processed as if the Sequence Number is one greater than received. This prevents the TLP with the current expected Sequence Number and all following TLPs from being received. This causes the link partner to do TLP replays as received TLPs are incorrect due to perceived Sequence Number errors.
[16]	rx_inhibit_ack_nak	RW	1	0x0	Receive Data Link Layer ACK/NAK Inhibit Enable. 0 – Process received ACK and NAK DLLPs per PCIe Specification. Required setting for compliant PCIe operation. 1 – For test purposes only, discard and do not process received ACK and NAK DLLPs. This causes the core to do TLP replays because TLP acknowledgements do not received.
[15]	reserved	RO	1	0x0	—
[14]	tx_par2_report_disable	RW	1	0x0	Transmit Data Link Layer Parity 2 Error Reporting Disable. 0 – Enable reporting. 1 – Disable reporting of Data Link Layer transmit parity errors.
[13]	tx_par2_handle_disable	RW	1	0x0	Transmit Data Link Layer Parity 2 Error Handling Disable. 0 – Enable handling. TLPs with errors are nullified and not retransmitted. 1 – Disable handling of Data Link Layer transmit parity errors. When error handling is disabled, TLPs with parity errors continue to be transmitted.
[12]	tx_par2_inject_en	RW	1	0x0	Transmit Data Link Layer Parity 2 Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single parity error injection, applied after assigning the TLP sequence number, is scheduled and is injected at the next opportunity (TLP transmit).
[11:9]	reserved	RO	3	0x0	—

Field	Name	Access	Width	Reset	Descriptions
[8]	tx_par1_inject_en	RW	1	0x0	Transmit Data Link Layer Parity 1 Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single parity error injection, applied prior to assigning the TLP sequence number, is scheduled and is injected at the next opportunity (TLP transmit).
[7]	reserved	RO	1	0x0	—
[6]	tx_replay_ecc2_report_disable	RW	1	0x0	Transmit Replay Buffer ECC 2-bit Error Reporting Disable. 0 – Enable reporting. 1 – Disable reporting of ECC 2-bit errors.
[5]	tx_replay_ecc2_handle_disable	RW	1	0x0	Transmit Replay Buffer ECC 2-bit Error Handling Disable. 0 – Enable handling. TLPs with errors are nullified and not retransmitted. 1 – Disable handling of ECC 2-bit errors. When error handling is disabled, TLPs with ECC 2-bit errors continue to be transmitted.
[4]	tx_replay_ecc2_inject_en	RW	1	0x0	Transmit Replay Buffer ECC 2-bit Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single ECC 2-bit error injection is scheduled and is injected at the next opportunity (Replay Buffer RAM write). The error is only seen if a Replay occurs of the TLP receiving the error injection.
[3]	reserved	RO	1	0x0	—
[2]	tx_replay_ecc1_report_disable	RW	1	0x0	Transmit Replay Buffer ECC 1-bit Error Reporting Disable. 0 – Enable reporting. 1 – Disable reporting of ECC 1-bit errors.
[1]	tx_replay_ecc1_handle_disable	RW	1	0x0	Transmit Replay Buffer ECC 1-bit Error Handling Disable. 0 – Enable correction. 1 – Disable correction of ECC 1-bit errors. When error correction is disabled, ECC 1-bit errors are treated the same as uncorrectable ECC 2-bit errors.
[0]	tx_replay_ecc1_inject_en	RW	1	0x0	Transmit Replay Buffer ECC 1-bit Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single ECC 1-bit error injection is scheduled and is injected at the next opportunity (Replay Buffer RAM write). The error is only seen if a Replay occurs of the TLP receiving the error injection.

dl_stat Register 0x394

This register set is used for the Data Link Layer Status.

Table 5.92. dl_stat Register 0x394

Field	Name	Access	Width	Reset	Description
[31]	info_bad_tlp_null_err	RW, W1C	1	0x0	Nullified TLP Received Status. This is not a reported error but is useful information to store for debug. This is a sub-class of err_aer_bad_tlp that provides more information as to why the TLP is bad. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[30]	info_bad_tlp_phy_err	RW, W1C	1	0x0	TLP PHY Error Status. This is not a reported error but is useful information to store for debug. This is a sub-class of err_aer_bad_tlp that provides more information as to why the TLP is bad. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[29]	info_bad_tlp_malf_err	RW, W1C	1	0x0	Malformed TLP Status. This is not a reported error but is useful information to store for debug. This is a sub-class of err_aer_bad_tlp that provides more information as to why the TLP is bad. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[28]	info_bad_tlp_ecrc_err	RW, W1C	1	0x0	TLP ECRC Mismatch Status. This is not a reported error but is useful information to store for debug. This is a sub-class of err_aer_bad_tlp that provides more information as to why the TLP is bad. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[27]	info_schedule_dupl_ack	RW, W1C	1	0x0	Duplicate TLP Received Status. This is not a reported error but is useful information to store for debug. Duplicate TLPs are received during TLP Replay. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[26]	info_bad_tlp_seq_err	RW, W1C	1	0x0	TLP Sequence Number Mismatch Status. This is not a reported error but is useful information to store for debug. This is a sub-class of err_aer_bad_tlp that provides more information as to why the TLP is bad. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[25]	info_bad_tlp_crc_err	RW, W1C	1	0x0	TLP LCRC Mismatch Status. This is not a reported error but is useful information to store for debug. This is a sub-class of err_aer_bad_tlp that provides more information as to why the TLP is bad. 0 – Otherwise 1 – Event occurred. Write 1 to clear.

Field	Name	Access	Width	Reset	Description
[24]	info_nak_received	RW, W1C	1	0x0	NAK Received Status. This is not a reported error but is useful information to store for debug. Receiving a NAK indicates that the link partner requested a Replay. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[23]	info_deskew_overflow_error	RW, W1C	1	0x0	Rx Deskew FIFO Overflow Error Status. The lane-lane skew of Rx data on one or more lanes are latent from the other lanes that the deskew range of the Rx Deskew FIFO is exceeded. This is a correctable error since the core drives the link to Recovery to fix this issue. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[22]	info_tx_data_underflow	RW, W1C	1	0x0	Physical Layer TLP Transmit Underflow Error Status. A TLP is transmitted by the physical layer and more data is needed to continue the transmission, but no data is provided. This error is normally caused by the Transaction Layer failing to provide TLP data at \geq PCIe Line Rate. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[21]	info_replay_started	RW, W1C	1	0x0	A Replay is started. This is not a reported error but is useful information to store for debug. Indicates a Replay occurred on local TX interface due to either Ack Timeout or Nack Reception. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[20]	reserved	RO	1	0x0	—
[19]	err_aer_tx_par2	RW, W1C	1	0x0	Transmit Data Link Layer Parity 2 Error Status. Indicates that a Data Link Layer transmit parity error is detected after sequence number application. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[18]	reserved	RO	1	0x0	—
[17]	err_aer_tx_replay_ecc2	RW, W1C	1	0x0	Transmit Replay Buffer ECC 2-bit Error Status. Indicates that an uncorrectable ECC error occurred during Replay. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[16]	err_aer_tx_replay_ecc1	RW, W1C	1	0x0	Transmit Replay Buffer ECC 1-bit Error Status. Indicates that a correctable ECC error occurred during Replay. 0 – Otherwise 1 – Event occurred. Write 1 to clear.
[15:8]	reserved	RO	8	0x0	—
[7]	reserved	RO	1	0x0	—

Field	Name	Access	Width	Reset	Description
[6]	err_aer_surprise_down	RW, W1C	1	0x0	Surprise Down Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.
[5]	err_aer_dl_protocol_error	RW, W1C	1	0x0	DL Protocol Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.
[4]	err_aer_replay_timer_timeout	RW, W1C	1	0x0	Replay Timer Timeout Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.
[3]	err_aer_replay_num_rollover	RW, W1C	1	0x0	Replay Num Rollover Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.
[2]	err_aer_bad_dlp	RW, W1C	1	0x0	Bad DLLP Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.
[1]	err_aer_bad_tlp	RW, W1C	1	0x0	Bad TLP Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.
[0]	err_aer_receiver_error	RW, W1C	1	0x0	Receiver Error Status. 0 – Otherwise 1 – Error occurred. Errors of this type must be logged in the Transaction Layer AER Capability. Write 1 to clear.

dl_ack_to_nak Register 0x398

This register set is used for the ACK-to-NAK error injection controls.

Table 5.93. dl_ack_to_nak Register 0x398

Field	Name	Access	Width	Reset	Description
[31]	enable	RW	1	0x0	Enable ACK-to-NAK injection. The write to this field must occur after writes to the other fields in this register have established the desired parameters of the injection. 0 – Do nothing 1 – Enable injection and load parameters into the ACK-to-NAK injector.
[30:24]	reserved	RO	7	0x0	—
[23:16]	count	RW	8	0x0	Number of times to replace the ACK with a NAK.
[15:12]	reserved	RO	4	0x0	—
[11:0]	seq_num	RW	12	0x0	Sequence Number of ACK to be changed to a NAK.

dl_inject Register 0x39c

This register set is used for the DLLP CRC/TLP ECRC error injection controls.

Table 5.94. dl_inject Register 0x39c

Field	Name	Access	Width	Reset	Description
[31]	dllp_crc_err_enable	RW	1	0x0	Enable DLLP CRC error injection. The write to this field must occur after writes to the other fields in this register have established the desired parameters of the injection. 0 – Disable the DLLP CRC Error injector. 1 – Enable DLLP CRC Error injector.
[30:28]	reserved	RO	3	0x0	—
[27:16]	dllp_crc_err_rate	RW	12	0x0	Rate at which DLLP CRC errors are to be injected. A value of 0 injects a single DLLP CRC error. A non-zero value injects errors at intervals of $\text{Rate} \times 256 \times \text{clk_period}$. This field may not be changed while <code>dllp_crc_err_enable == 1</code> .
[15:13]	reserved	RO	3	0x0	—
[12]	dllp_inject_enable	RW	1	0x0	Inject a DLLP (transmit) using the data in the <code>dllp_inject_data</code> register. A single DLLP is injected after each rising edge of this signal.
[11:9]	reserved	RO	3	0x0	—
[8]	tlp_seq_err_enable	RW	1	0x0	Modify the sequence number in the next transmitted TLP to an invalid value (bad sequence number error). A single TLP is altered after each rising edge of this signal.
[7:4]	reserved	RO	4	0x0	—
[3]	tlp_lcrc_err_enable	RW	1	0x0	Enable TLP LCRC error injection. The write to this field must occur after writes to the other fields in this register have established the desired parameters of the injection. 0 – Disable the TLP LCRC Error injector 1 – Enable TLP LCRC Error injector

Field	Name	Access	Width	Reset	Description
[2:0]	tlp_lcrc_err_rate	RW	3	0x0	Rate at which TLP LCRC errors are to be injected. A value of 0 injects LCRC errors into all TLPs. A non-zero value injects an error into a TLP and then pass Rate TLPs without error, and then repeat. This field may not be changed while tlp_lcrc_err_enable==1.

dllp_inject Register 0x3a0

This register set is used for the DLLP Injector Data.

Table 5.95. dllp_inject Register 0x3a0

Field	Name	Access	Width	Reset	Description
[31:0]	data	RW	32	0x0	Data to include in injected DLLP. This field may not be changed while dl_inject_dllp_inject_enable==1.

5.1.2.8. LTSSM Equalization Status Control

eq_status_table_control Register 0x3d8

This register set is used for LTSSM Equalization Status Control for 8G data rate.

The LTSSM Equalization Status is a debug feature that captures the status of the most recent Equalization execution from a single selected lane and allows software to read the status.

To use this feature:

1. Select a lane to monitor by writing lane_select.
2. Cause Equalization to be executed; Equalization can be executed by writing the following values in sequence to the Downstream Port PCIe Cfg Registers (see PCIe Specification for details):
 - a. Perform Equalization = 1
 - b. Target Link Speed = Desired Link Speed (1 = 2.5G, 2 = 5G, 3 = 8G¹, 4 = 16G¹)
 - c. Retrain Link = 1 (this must be written last)
3. Wait for Equalization to complete by reading eq_status_table_info_done until it indicates done.
4. The speed at which EQ is executed may be read through eq_status_table_info_16g_speed.
5. For (i=0; i<25; i=i+1).

```
begin
Write eq_status_table_control_step_select = i
Read eq_status_table_data == Equalization Status of iteration[i]
End
```

Note:

1. Available in 2024 release.

Table 5.96. eq_status_table_control Register 0x3d8

Field	Name	Access	Width	Reset	Description
[31:13]	reserved	RO	19	0x0	—
[12:8]	step_select	RW	5	0x0	Used to select which step (which EQ trial iteration) of the most recently completed EQ process is accessed at the eq_status_table_data. After Equalization has completed, step_select is intended to be written from 0 to 24 to read the Equalization status for all trials.
[7:4]	reserved	RO	4	0x0	—
[3:0]	lane_select	RW	4	0x0	Used to select which lane's data is collected during the next Equalization. This value of this register is captured at the start of Equalization and Equalization results are then saved for the selected lane.

eq_status_table_info Register 0x3dc

This register set is used to read whether the RX Equalization status registers are valid or not at 8G data rate.

Table 5.97. eq_status_table_info Register 0x3dc

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1]	16g_speed	RO	1	0x0	Indicates speed at which Equalization is executed. This register is reset only by fundamental reset so that the status survives link down and other soft reset conditions. 0 – Equalization run at 8G speed 1 – Equalization run at 16G speed
[0]	done	RO	1	0x0	Indicates that Equalization has been completed, and the table results are valid. This register is reset only by fundamental reset so that the status survives link down and other soft reset conditions.

eq_status_table Register 0x3e0

This register set is used to read RX Equalization status registers for 8G data rate.

Table 5.98. eq_status_table Register 0x3e0

Field	Name	Access	Width	Reset	Description
[31:0]	data	RO	32	0x0	<ul style="list-style-type: none"> Equalization Status for the selected lane and step (iteration) recorded as follows: <ul style="list-style-type: none"> [1:0] – EQ Pre-Cursor Up/Down Feedback[1:0] (00=Hold,01=Inc,10=Dec,11=Rsvd) [3:2] – EQ Post-Cursor Up/Down Feedback[1:0] (00=Hold,01=Inc,10=Dec,11=Rsvd) [11: 4] – EQ Figure of Merit Feedback[7:0] [17:12] – Remote PHY Tx Post Cursor[5:0] [23:18] – Remote PHY Tx Pre Cursor[5:0] [24] – Error Status: 1==Error, 0==No Error [25] – Active Status: 1==Lane is part of the link, 0==Lane is not part of the link [31:16] – Reserved Equalization Status is reset to 0x0 for all iterations when Equalization begins so iterations that were not executed reads 0x0. Equalization Status is re-captured on every Equalization, so Equalization Status may only be read between Equalization attempts or it is unknown from which Equalization run the status is captured. Captured Equalization Status is reset only by fundamental reset so that the status survives link down and other soft reset conditions.

eq_capture_sel Register 0x3f0

This register set is used for selection of lane, TX or RX for RX Equalization results for 8G data rate.

Table 5.99. eq_capture_sel Register 0x3f0

Field	Name	Access	Width	Reset	Description
[31:6]	reserved	RO	26	0x0	—
[5]	dir	RW	1	0x0	Used to select TX (dir == 1) or RX (dir == 0) for equalization results read-back.
[4]	speed	RW	1	0x0	Used to select 16G (speed == 1) or 8G (speed == 0) for equalization results read-back.
[3:0]	lane	RW	4	0x0	Used to select the lane number (0 to 15) for equalization results read-back.

eq_capture Register 0x3f4

This register set is used to read out RX Equalization results for 8G data rate.

Table 5.100. eq_capture Register 0x3f4

Field	Name	Access	Width	Reset	Description
[31:22]	reserved	RO	10	0x0	—
[21:0]	result	RO	22	0x0	<ul style="list-style-type: none"> Equalization results from LTSSM recorded as follows: <ul style="list-style-type: none"> [5: 0] – EQ Post-Cursor[5:0] [11: 6] – EQ Cursor[5:0] [17:12] – EQ Pre-Cursor[5:0] [21:18] – EQ Preset <p>Note: The Pre-Cursor, Cursor, and Post-Cursor values is 0x0 if Equalization Presets are used. Similarly, Preset is 0x0 if Equalization Coefficients are used.</p> Equalization results are set to 0x0 for all lanes, speed, and directions by a fundamental reset. The values are updated at each step of the equalization procedure, and then held until the next equalization or fundamental reset.

phy_eq_tx_force_per_lane Register 0x400

This register set is used to enable per-Lane Hardcoded Preset/Coefficient configuration for 8G data rate.

Table 5.101. phy_eq_tx_force_per_lane Register 0x400

Field	Name	Access	Width	Reset	Description
[31:4]	reserved	RO	28	0x0	—
[3]	16g_coef_enable	RW	1	0x0	Used to enable per-lane hardcoded coefficients at 16G speed.
[2]	16g_preset_enable	RW	1	0x0	Used to enable per-lane hardcoded presets at 16G speed
[1]	8g_coef_enable	RW	1	0x0	Used to enable per-lane hardcoded coefficients at 8G speed.
[0]	8g_preset_enable	RW	1	0x0	Used to enable per-lane hardcoded presets at 8G speed

phy_eq_tx_force_per_lane_8g_pre Register 0x404

This register set is used for per-Lane Hardcoded Preset/Pre-cursor values configuration for 8G data rate.

Table 5.102. phy_eq_tx_force_per_lane_8g_pre Register 0x404

Field	Name	Access	Width	Reset	Description
[95:0]	value	RW	96	0x0	Per-lane preset/pre-cursor values to use in hardcoded per-lane mode at 8G speed.

phy_eq_tx_force_per_lane_8g_post Register 0x410

This register set is used for per-Lane Hardcoded Preset/Post-cursor values configuration for 8G data rate.

Table 5.103. phy_eq_tx_force_per_lane_8g_post Register 0x410

Field	Name	Access	Width	Reset	Description
[95:0]	value	RW	96	0x0	Per-lane post-cursor values to use in hardcoded per-lane mode at 8G speed.

5.1.3. mgmt_ptl (0x4_3000)

The following are the register sets with the 0x3000 base address.

5.1.3.1. Simulation Register

Simulation Register 0x0

This register set is used for the Partial Transaction Layer simulation speed reduction.

Table 5.104. Simulation Register 0x0

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	pm_reduce_timeouts	RW	1	0x0	Reduce Power Management State Machine timeouts from their value in ms to their value in μ s to shorten simulation time. 0 – Disable 1 – Enable

5.1.3.2. Power Management State Machine Register Set

pm_aspm_l0s Register 0x40

This register set is used for the Power Management State Machine ASPM L0s entry control.

Table 5.105. pm_aspm_l0s Register 0x40

Field	Name	Access	Width	Reset	Description
[31:16]	entry_time	RW	16	0x0	ASPM L0s TX Entry Time in μ s. 0 is a special case == 6.9 μ s.
[15:1]	reserved	RO	15	0x0	—
[0]	enable	RW	1	0x1	Enable Power Management State Machine to enter ASPM L0s. 0 – Disable 1 – Enable

pm_aspm_l1 Register 0x50

This register set is used for the Power Management State Machine ASPM L1 entry control.

Table 5.106. pm_aspm_l1 Register 0x50

Field	Name	Access	Width	Reset	Description
[31:16]	entry_time	RW	16	0x0	ASPM L1 TX Entry Time in ms. 0 is a special case == 1000 μ s
[15:1]	reserved	RO	15	0x0	—
[0]	enable	RW	1	0x1	Enable Power Management State Machine to enter ASPM L1. 0 – Disable 1 – Enable

pm_aspm_l1_min Register 0x54

This register set is used for the Power Management State Machine ASPM L1 re-entry control.

Table 5.107. pm_aspm_l1_min Register 0x54

Field	Name	Access	Width	Reset	Description
[31:30]	reserved	RO	2	0x0	—
[29:16]	reentry_time	RW	14	0x0	When reentry_disable==0, specifies the minimum time between ASPM L1 requests in ns. 0 is a special case == 9500 ns (PCIe Specification value).
[15:1]	reserved	RO	15	0x0	—
[0]	reentry_disable	RW	1	0x0	Disable enforcing a minimum time between ASPM L1 requests. 0 – Enable 1 – Disable

pm_l1 Register 0x60

This register set is used for the Power Management State Machine L1 entry control.

Table 5.108. pm_l1 Register 0x60

Field	Name	Access	Width	Reset	Description
[31:16]	us_port_ps_entry_time	RW	16	0x0	Upstream Ports only: Number of μ s to wait for the transmission of the completion to the PowerState Cfg Write that initiated L1 entry, before beginning to block TLPs and enter L1. 0 is a special case == 4 μ s.
[15:1]	reserved	RO	15	0x0	—
[0]	enable	RW	1	0x1	Enable Power Management State Machine to enter L1. 0 – Disable 1 – Enable

pm_l1_min Register 0x64

This register set is used for the Power Management State Machine L1 re-entry control.

Table 5.109. pm_l1_min Register 0x64

Field	Name	Access	Width	Reset	Description
[31:24]	reserved	RO	8	0x0	—
[23:16]	ps_reentry_time	RW	8	0x0	Minimum number of μ s to wait following an L1 exit when Power State != D0, before re-entering L1 due to Power State != D0. A wait time is needed to give the transaction layer time to process a Power State Cfg Write to D0 that caused L1 exit. 0 is a special case == 50 μ s.
[15:0]	reserved	RO	16	0x0	—

pm_l1pmss Register 0x68

This register set is used for the Power Management State Machine L1PMSS control.

Table 5.110. pm_l1pmss Register 0x68

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	ds_drive_clkreq	RW	1	0x0	Enable driving the clkreq_n signal when operating as a downstream port. 0 – Disable 1 – Enable

pm_l2 Register 0x70

This register set is used for the Power Management State Machine L2 entry control.

Table 5.111. pm_l2 Register 0x70

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x1	Enable Power Management State Machine to enter L2. 0 – Disable 1 – Enable

pm_pme_to_ack_ep Register 0x80

This register set is used for the Power Management State Machine Endpoint PME_TO_Ack control.

Table 5.112. pm_pme_to_ack_ep Register 0x80

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	user_auto_n	RW	1	0x0	For Endpoints only: PME_TO_Ack message transmission scheduling method. Endpoints are required to respond to a PME_Turn_Off message with a PME_TO_Ack message when they are ready to allow power down. 0 – Schedule PME_TO_Ack message automatically on reception of PME_Turn_Off message. 1 – Schedule PME_TO_Ack message under user control through the pm_l2_enter_ack rising edge.

pm_pme_to_ack_ds Register 0x84

This register set is used for the Power Management State Machine Downstream Port PME_TO_Ack control.

Table 5.113. pm_pme_to_ack_ds Register 0x84

Field	Name	Access	Width	Reset	Description
[31:8]	reserved	RO	24	0x0	—
[7:0]	timeout_threshold	RW	8	0x0	For Root Port only: ms to wait for a transmitted PME_Turn_Off Message to be acknowledged by receipt of PME_TO_Ack message before continuing with L2/L3 entry. 0xFF is a special case that disables the timeout mechanism. 0x00 is a special case == 10 ms.

pm_pme Register 0x88

This register set is used for the Power Management State Machine PM_PME control.

Table 5.114. pm_pme Register 0x88

Field	Name	Access	Width	Reset	Description
[31:12]	reserved	RO	20	0x0	—
[11:0]	timeout_threshold	RW	12	0x0	ms to wait for a transmitted PM_PME Message to be acknowledged, by clearing of the PME_Status register, before reissuing the PM_PME message. 0xFFFF is a special case that disables the timeout mechanism. 0x000 is a special case == 100 ms.

pm_status Register 0x90

This register set is used for the Power Management State Machine Status.

Table 5.115. pm_status Register 0x90

Field	Name	Access	Width	Reset	Description
[31:5]	reserved	RO	27	0x0	—
[4:0]	state	RO	5	0x0	Power Management State Machine State. 0 – IDLE 1 – L1_WAIT_IDLE 2 – L1_WAIT_REPLAY 3 – L1_READY 4 – L1_STOP_DLLP 5 – L1 6 – L1_1 7 – L1_2_ENTRY 8 – L1_2_IDLE 9 – L1_2_EXIT 10 – L1_EXIT 11 – L2_WAIT_IDLE 12 – L2_WAIT_REPLAY 13 – L23_READY 14 – L2_STOP_DLLP 15 – L2 16 – LOS

5.1.3.3. Receive Buffer

vc_rx_c Register 0x108

This register set is used for the Receive Buffer completion handling configuration.

Table 5.116. vc_rx_c Register 0x108

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1]	force_ro	RW	1	0x0	Force completion relaxed ordering (RO==1) behavior for all completion TLPs, even those with RO==0. Note that setting this register to 1 is not PCIe Specification compliant but this may be fine for some designs since it is acceptable in many designs for C without RO==1 to pass prior P. 0 – Disable. Received completions are handled using the received RO attribute. 1 – Enable. All received completions are handled as if the RO attribute is 1.
[0]	priority	RW	1	0x0	Completion priority enable. 0 – Disable. Arbitration between Posted, Non-Posted, and Completion TLPs is round robin. 1 – Enable. While arbitrating between putting pending received Posted, Non-Posted, and Completion TLPs on the user received TLP interface, completions are given highest priority. Posted and non-posted requests transact only when a completion is not pending or as needed to prevent starving.

vc_rx_adv Register 0x10c

This register set is used for the Receive Buffer completion credit advertisement configuration.

Table 5.117. vc_rx_adv Register 0x10c

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1:0]	ch_cd_sel	RW	2	0x0	PCIe Specification requires CH and CD credit advertisements to be infinite for Endpoints and the finite (actual credit values) for Root Port. ch_cd_sel may be configured to over-ride the default PCIe Specification expected behavior. 0 – Implement CH, CD credit advertisements per port type: Endpoints == infinite, Root Port == actual. 1 – Advertise actual CH, CD credits. 2 – Advertise Infinite CH, CD credits

vc_rx_control Register 0x110

This register set is used for the Receive Buffer Parity/ECC control.

Table 5.118. vc_rx_control Register 0x110

Field	Name	Access	Width	Reset	Description
[31:20]	reserved	RO	12	0x0	—
[19]	reserved	RO	1	0x0	—
[18]	par1_report_disable	RW	1	0x0	Receive Buffer Prefix Parity Error and Parity 1 Error Reporting Disable. 0 – Enable reporting. 1 – Disable reporting of detected parity errors.
[17]	pfx_par_inject_en	RW	1	0x0	Receive Buffer Prefix Parity Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single prefix parity error injection is scheduled and is injected at the next opportunity (TLP receipt).

Field	Name	Access	Width	Reset	Description
[16]	par1_inject_en	RW	1	0x0	Receive Buffer Parity 1 Error Injection Enable. 0 – Do not inject error. 1 – On the rising edge, a single parity error injection is scheduled and is injected at the next opportunity (TLP receipt).
[15:13]	reserved	RO	3	0x0	—
[12]	ecc2_report_disable	RW	1	0x0	Receive Buffer ECC 2-bit Error Reporting Disable. 0 – Enable reporting. 1 – Disable reporting of ECC 2-bit errors.
[11]	ecc2_handle_disable	RW	1	0x0	Receive Buffer ECC 2-bit Error Handling Disable. 0 – Enable handling. 1 – Disable handling of ECC 2-bit errors.
[10:9]	reserved	RO	2	0x0	—
[8]	ecc2_inject_en	RW	1	0x0	Receive Buffer ECC 2-bit Error Injection Enable. 1 – On the rising edge, a single ECC 2-bit error injection is scheduled and is injected at the next opportunity (Receive Buffer RAM read).
[7:5]	reserved	RO	3	0x0	—
[4]	ecc1_report_disable	RW	1	0x0	Receive Buffer ECC 1-bit Error Reporting Disable. 0 – Enable reporting. 1 – Disable reporting of ECC 1-bit errors.
[3]	ecc1_handle_disable	RW	1	0x0	Receive Buffer ECC 1-bit Error Handling Disable. 0 – Enable correction. 1 – Disable correction of ECC 1-bit errors. When error correction is disabled, ECC 1-bit errors are treated the same as uncorrectable ECC 2-bit errors.
[2:1]	reserved	RO	2	0x0	—
[0]	ecc1_inject_en	RW	1	0x0	Receive Buffer ECC 1-bit Error Injection Enable. 1 – On the rising edge, a single ECC 1-bit error injection is scheduled and is injected at the next opportunity (Receive Buffer RAM read).

vc_rx_status Register 0x114

This register is used for the Receive Buffer Parity/ECC status.

Table 5.119. vc_rx_status Register 0x114

Field	Name	Access	Width	Reset	Description
[31:4]	reserved	RO	28	0x0	—
[3]	err_pfx_par	RW, wr:oneToClear	1	0x0	Receive Buffer Prefix Parity Error Detection Status. 0 – Otherwise 1 – Error occurred.
[2]	err_par1	RW, wr:oneToClear	1	0x0	Receive Buffer Parity 1 Error Detection Status. 0 – Otherwise 1 – Error occurred.
[1]	err_ecc2	RW, wr:oneToClear	1	0x0	Receive Buffer ECC 2-bit Error Detection Status. 0 – Otherwise 1 – Error occurred.
[0]	err_ecc1	RW, wr:oneToClear	1	0x0	Receive Buffer ECC 1-bit Error Detection Status. 0 – Otherwise 1 – Error occurred.

vc_rx_credit_status_cfg Register 0x120

This register set is used for the Receive Buffer credit status configuration.

Table 5.120. vc_rx_credit_status_cfg Register 0x120

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1:0]	sel	RW	2	0x0	Receive Buffer credit status configuration selection. Configures which credit values are displayed in the remaining vc_rx_credit_status registers. 0 – Display static initial credit advertisements. 1 – Display dynamic current credits available. 2 – Display static number of credits implemented by the receiver buffer. This may be larger than the initial credit advertisement values when number of credits implemented exceeds the PCIe Specification advertised maximum values of 127 H and 2047 D. 3 – Reserved.

vc_rx_credit_status_p Register 0x124

This register set is used for the Receive Buffer Posted credit status.

Table 5.121. vc_rx_credit_status_p Register 0x124

Field	Name	Access	Width	Reset	Description
[31:16]	d	RO	16	0x0	Receive Buffer Posted Data credit status. Content dependent on the value of vc_rx_credit_status_cfg_sel.
[15:0]	h	RO	16	0x0	Receive Buffer Posted Header credit status. Content dependent on the value of vc_rx_credit_status_cfg_sel.

vc_rx_credit_status_n Register 0x128

This register set is used for the Receive Buffer Non-Posted credit status.

Table 5.122. vc_rx_credit_status_n Register 0x128

Field	Name	Access	Width	Reset	Description
[31:16]	d	RO	16	0x0	Receive Buffer Non-Posted Data credit status. Content dependent on the value of vc_rx_credit_status_cfg_sel.
[15:0]	h	RO	16	0x0	Receive Buffer Non-Posted Header credit status. Content dependent on the value of vc_rx_credit_status_cfg_sel.

vc_rx_credit_status_c Register 0x12c

This register set is used for the Receive Buffer Completion credit status.

Table 5.123. vc_rx_credit_status_c Register 0x12c

Field	Name	Access	Width	Reset	Description
[31:16]	d	RO	16	0x0	Receive Buffer Completion Data credit status. Content dependent on the value of vc_rx_credit_status_cfg_sel.
[15:0]	h	RO	16	0x0	Receive Buffer Completion Header credit status. Content dependent on the value of vc_rx_credit_status_cfg_sel.

vc_rx_f_oc_update_timer Register 0x130

This register set is used for the Receive Buffer FC Update Timer Control.

Table 5.124. vc_rx_f_oc_update_timer Register 0x130

Field	Name	Access	Width	Reset	Description
[31:3]	reserved	RO	29	0x0	—
[2:1]	div	RW	2	0x0	Divider Control to Reduce Period of FC Updates for Highly Latent Systems 0 – No adjustment 1 – Divide the PCIe Spec Guideline Value by 2 2 – Divide the PCIe Spec Guideline Value by 4 3 – Divide the PCIe Spec Guideline Value by 8
[0]	disable	RW	1	0x0	Disable Control for the FC Update Timer 0 – Enable the FC Update Timer. 1 – Disable FC Update Timer, sending FC Updates on Every Consumed RX Packet

vc_rx_p_flow_ctrl Register 0x134

Receive Buffer Posted TLP Flow Control.

Table 5.125. vc_rx_p_flow_ctrl Register 0x134

Field	Name	Access	Width	Reset	Description
[31:26]	reserved	RO	6	0x0	—
[25:16]	thresh	RW	10	0x10	Receive Buffer Posted TLP Flow Control Threshold Enable. Threshold to use when thresh_en == 1.
[15:9]	reserved	RO	7	0x0	—
[8]	thresh_en	RW	1	0x0	Receive Buffer Posted TLP Flow Control Threshold. 0 – Use threshold provided on vc_rx_p_thresh/ptl_rx_p_thresh input port. 1 – Use threshold provided by the thresh register
[7:1]	reserved	RO	7	0x0	—
[0]	disable	RW	1	0x0	Receive Buffer Posted TLP Flow Control Disable. 0 – Enable Posted TLP Flow Control. 1 – Disable Posted TLP Flow Control

vc_rx_n_flow_ctrl Register 0x138

This register set is used for the Receive Buffer Non-Posted TLP Flow Control.

Table 5.126. vc_rx_n_flow_ctrl Register 0x138

Field	Name	Access	Width	Reset	Description
[31:26]	reserved	RO	6	0x0	—
[25:16]	thresh	RW	10	0x10	Receive Buffer Non-Posted TLP Flow Control Threshold. 0 – Use threshold provided on vc_rx_np_thresh/ptl_rx_np_thresh input port. 1 – Use threshold provided by the thresh register.
[15:9]	reserved	RO	7	0x0	—
[8]	thresh_en	RW	1	0x0	Receive Buffer Non-Posted TLP Flow Control Disable. 0 – Enable Posted TLP Flow Control. 1 – Disable Posted TLP Flow Control.
[7:1]	reserved	RO	7	0x0	—
[0]	disable	RW	1	0x0	—

vc_rx_alloc_size Register 0x140

This register set is used for the Receive Buffer Size Information.

Table 5.127. vc_rx_alloc_size Register 0x140

Field	Name	Access	Width	Reset	Description
[31:24]	c_hdr	RO	8	0x0	Number of Receive Buffer Data storage bytes required for each unit of alloc_ch.
[23:16]	pn_hdr	RO	8	0x0	Number of Receive Buffer Data storage bytes required for each unit of alloc_ph and alloc_nh.
[15:8]	storage_data	RO	8	0x0	Receive Buffer Data Storage size in bytes == $2^{\text{storage_data}}$. The storage space required for alloc_ph, alloc_nh, alloc_ch, alloc_pd, alloc_nd, and alloc_cd must be $\leq (2^{\text{storage_data}})$ bytes. The space required for each unit of alloc_ph, alloc_nh, and alloc_cd is indicated in the fields pn_hdr and c_hdr. 16 bytes of space is required for each unit of alloc_pd, alloc_nd, and alloc_cd. Credit allocations must fit within the Receive Buffer Data Storage, or the Receive Buffer is unable to function correctly, which leads to serious errors.
[7:0]	storage_hdr	RO	8	0x0	Receive Buffer Header Storage size in number of TLPs == $2^{\text{storage_hdr}}$. alloc_ph + alloc_nh + alloc_ch must be $\leq 2^{\text{storage_hdr}}$. Credit allocations must fit within the Receive Buffer Header Storage, or the Receive Buffer is unable to function correctly, which leads to serious errors.

vc_rx_alloc_p Register 0x144

This register set is used for the Receive Buffer Posted TLP Credit Allocation.

Table 5.128. vc_rx_alloc_p Register 0x144

Field	Name	Access	Width	Reset	Description
[31:16]	d	RW	16	0x440	Number of PCI Express PD credits to allocate in the Receive Buffer. The minimum value is (Max Payload Size Supported in Bytes/16). Number of PCI Express CD credits to allocate in the Receive Buffer. The minimum value is (Max Payload Size Supported in Bytes/16). If more than 2047 credits are allocated, the PCIe Core reserves the requested amount of space in the Receive Buffer, but it advertises the maximum allowed PCIe Specification value of 2047.
[15:12]	reserved	RO	4	0x0	—
[11:0]	h	RW	12	0x20	Number of PCI Express PH credits to allocate in the Receive Buffer. The minimum value is 1. If more than 127 credits are allocated, the PCIe Core reserves the requested amount of space in the Receive Buffer but it advertises the maximum allowed PCIe Specification value of 127.

vc_rx_alloc_n Register 0x148

This register set is used for the Receive Buffer Non-Posted TLP Credit Allocation.

Table 5.129. vc_rx_alloc_n Register 0x148

Field	Name	Access	Width	Reset	Description
[31:16]	d	RW	16	0x40	Number of PCI Express ND credits to allocate in the Receive Buffer. The minimum value is 2.
[15:12]	reserved	RO	4	0x0	—
[11:0]	h	RW	12	0x20	Number of PCI Express NH credits to allocate in the Receive Buffer. The minimum value is 1. If more than 127 credits are allocated, the PCIe Core reserves the requested amount of space in the Receive Buffer but it advertises the maximum allowed PCIe Specification value of 127.

vc_rx_alloc_c Register 0x14c

This register set is used for the Receive Buffer Completion TLP Credit Allocation.

Table 5.130. vc_rx_alloc_c Register 0x14c

Field	Name	Access	Width	Reset	Description
[31:16]	d	RW	16	0x700	Number of PCI Express CD credits to allocate in the Receive Buffer. The minimum value is (Max Payload Size Supported in Bytes/16). If more than 2047 credits are allocated, the PCIe Core reserves the requested amount of space in the Receive Buffer, but it advertises the maximum allowed PCIe Specification value of 2047 (or infinite credits when operating as an Endpoint as required by PCIe Specification.).
[15:12]	reserved	RO	4	0x0	—
[11:0]	h	RW	12	0x1c0	Number of PCI Express CH credits to allocate in the Receive Buffer. The minimum value is 1. If more than 127 credits are allocated, the PCIe Core reserves the requested amount of space in the Receive Buffer but it advertises the maximum allowed PCIe Specification value of 127.

vc_rx_alloc_error Register 0x150

This register set is used for the Receive Buffer Allocation Error Status.

Table 5.131. vc_rx_alloc_error Register 0x150

Field	Name	Access	Width	Reset	Description
[31:4]	reserved	RO	28	0x0	—
[3]	d_sum	RO	1	0x0	Receive Buffer Data Storage Allocation Error. The storage space required for alloc_ph, alloc_nh, alloc_ch, alloc_pd, alloc_nd, and alloc_cd must be $\leq (2^{\text{storage_data}})$ bytes. The space required for each unit of alloc_ph, alloc_nh, and alloc_cd is indicated in the fields pn_hdr and c_hdr. 16 bytes of space is required for each unit of alloc_pd, alloc_nd, and alloc_cd. Credit allocations must fit within the Receive Buffer Data Storage, or the Receive Buffer is unable to function correctly, which leads to serious errors. 0 – No Error 1 – Error
[2]	h_sum	RO	1	0x0	Receive Buffer Header Storage Allocation Error. alloc_ph + alloc_nh + alloc_ch must be $\leq 2^{\text{storage_hdr}}$. Credit allocations must fit within the Receive Buffer Header Storage, or the Receive Buffer is unable to function correctly, which leads to serious errors. 0 – No Error 1 – Error
[1]	min_d	RO	1	0x0	Receive Buffer Data Credit Allocation Minimum Error. The alloc_pd and alloc_cd minimum value is Max Payload Size Supported. alloc_nd minimum value is 2. alloc_pd, alloc_nd, and alloc_cd must be a multiple of 2 for cores with DATA_WIDTH=256. 0 – No Error 1 – Error

Field	Name	Access	Width	Reset	Description
[0]	min_h	RO	1	0x0	Receive Buffer Header Credit Allocation Minimum Error. alloc_ph, alloc_nh, and alloc_ch must be > 0. 0 – No Error 1 – Error

vc_tx_np_fifo Register 0x180

This register set is used for the Transmit buffer non-posted TLP FIFO configuration.

Table 5.132. vc_tx_np_fifo Register 0x180

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	disable	RW	1	0x0	Transmit buffer non-posted TLP FIFO disable. The transmit non-posted TLP FIFO is present to allow Posted and Completion TLPs to continue to make progress, to avoid deadlock conditions, when Non-Posted TLP transmission is blocked by available credits in the link partner receive buffer. 0 – Enable non-posted TLP FIFO. PCIe Specification required value. 1 – For debug only: Disable non-posted TLP FIFO in which case non-posted TLPs are carried in the same data path as posted and completion TLPs.

vc_tx_status Register 0x184

This register set is used for the Transmit buffer status.

Table 5.133. vc_tx_status Register 0x184

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	remote_credit_block	RO	1	0x0	Set to one whenever a transmit TLP transmission is blocked by insufficient credits in the remote device's receive buffer.

vc_tx_credit_status_p Register 0x190

This register set is used for the TLP transmit Posted credits currently available in link partner receive buffer.

Table 5.134. vc_tx_credit_status_p Register 0x190

Field	Name	Access	Width	Reset	Description
[31:28]	reserved	RO	4	0x0	—
[27:16]	d	RO	12	0x0	Link partner current Receive Buffer Posted Data credits available.
[15:8]	reserved	RO	8	0x0	—
[7:0]	h	RO	8	0x0	Link partner current Receive Buffer Posted Header credits available.

vc_tx_credit_status_n Register 0x194

This register set is used for the TLP transmit Non-Posted credits currently available in link partner receive buffer.

Table 5.135. vc_tx_credit_status_n Register 0x194

Field	Name	Access	Width	Reset	Description
[31:28]	reserved	RO	4	0x0	—
[27:16]	d	RO	12	0x0	Link partner current Receive Buffer Non-Posted Data credits available.
[15:8]	reserved	RO	8	0x0	—
[7:0]	h	RO	8	0x0	Link partner current Receive Buffer Non-Posted Header credits available.

vc_tx_credit_status_c Register 0x198

This register set is used for the TLP transmit Completion credits currently available in link partner receive buffer.

Table 5.136. vc_tx_credit_status_c Register 0x198

Field	Name	Access	Width	Reset	Description
[31:28]	reserved	RO	4	0x0	—
[27:16]	d	RO	12	0x0	Link partner current Receive Buffer Completion Data credits available.
[15:8]	reserved	RO	8	0x0	—
[7:0]	h	RO	8	0x0	Link partner current Receive Buffer Completion Header credits available.

vc_tx_credit_cleanup Register 0x19c

This register set is used for the TLP transmit error credit cleanup control.

Table 5.137. vc_tx_credit_cleanup Register 0x19c

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	method	RW	1	0x0	TLP Transmit Credit Cleanup Method. 0 – Use the headers of the cleaned-up TLPs to recover the credits. The credits in TLPs with corrupted headers are not recovered. 1 – Use a credit lookup table based on the ID assigned to the TLP. This table is implemented in pcie_user_if.

5.1.3.4. TLP Transmit Control

tlp_tx Register 0x1c4

This register set is used to enable TD bit.

Table 5.138. tlp_tx Register 0x1c4

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	td1_means_add_has_n	RW	1	0x0	TLP Transmit TD==1 Header Field Interpretation. 0 – When a TLP is transmitted with TLP header bit TD==1, this means that the TLP already contains an ECRC. The core transmits the TLP with the TLP's existing ECRC and does not attempt to generate/append a new ECRC. 1 – Not supported for Full Transaction Layer cores like the CrossLink™-NX cores. td1_means_add_has_n must be set to 0.

5.1.3.5. FC Credit Init Control.

fc_credit_init Register 0x1c8

This register set is used to force the core to reperform FC credit initialization.

Table 5.139. fc_credit_init Register 0x1c8

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	redo	RW	1	0x0	Force the core to redo FC Credit Initialization without taking the link down. This is only possible if both ends of the link are instructed to redo the initialization.

5.1.4. mgmt_ftl (0x4_4000)

5.1.4.1. Simulation Register

simulation Register 0x0

This register set is used for simulation only such as Full Transaction Layer simulation speed reduction.

Table 5.140. simulation Register 0x0

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	reduce_timeouts	RW	1	0x0	Reduce timeouts to shorten simulation time. When enabled ms timeouts are shortened to the value in μ s. 0 – Disable 1 – Enable

5.1.4.2. Transaction Layer Decode Configuration Register

decode Register 0x10

This register set is used for the Transaction Layer Decode configuration.

Table 5.141. decode Register 0x10

Field	Name	Access	Width	Reset	Description
[31:24]	reserved	RO	8	0x0	—
[23:18]	reserved	RO	6	0x0	—
[17]	tx_bypass_decode_en	RW	1	0x0	Bypass the TLP decode block in the Transmit path. 0 – Decode_in_path module is enabled. 1 – Decode_in_path module is bypassed.
[16]	rx_bypass_decode_en	RW	1	0x0	Bypass the TLP decode block in the Receive path. 0 – Decode_in_path module is enabled. 1 – Decode_in_path module is bypassed.
[15:11]	reserved	RO	5	0x0	—
[10]	tx_convert_ur_to_ca	RW	1	0x0	When decoding TX packets convert Unsupported Request (UR) packets to Completer Abort (CA). 0 – Normal Operation. 1 – Convert UR to CA.
[9]	rx_convert_ur_to_ca	RW	1	0x0	When decoding RX packets convert Unsupported Request (UR) packets to Completer Abort (CA). 0 – Normal Operation. 1 – Convert UR to CA.
[8]	t0_rx_bypass_msg_dec	RW	1	0x0	When implementing Type 0 Configuration Space (Endpoint) – Bypass RX Message TLP Decode Enable. 0 – Normal operation. The core claims and does not forward Message TLPs to the TLP Receive Interface. 1 – All valid Msg TLPs received on PCIe (except Routed by ID and Routed by Address which are routed according to the routing type) are forwarded to the TLP Receive Interface.
[7:3]	reserved	RO	5	0x0	—
[2]	vendor0_ur	RW	1	0x1	Vendor Type 0 Messages received from PCIe are reported as UR. 0 – Do not report received Vendor Type 0 Messages as Unsupported Request (UR). 1 – Report received Vendor Type 0 Messages as Unsupported Request (UR).
[1]	target_only	RW	1	0x0	Target Only. Enable for user designs that implement purely target-only functionality. When enabled all received completions are considered Unexpected Completions and are not forwarded to the TLP Receive Interface. 0 – Disable 1 – Enable

Field	Name	Access	Width	Reset	Description
[0]	ignore_poison	RW	1	0x1	<p>Ignore Poison – Set to 1 to have the core ignore the EP poison indicator for received TLPs with data payload that do not terminate in the core. When set to 1, the core passes all poisoned TLPs to you the same way it would pass the TLP if the TLP is not poisoned. Note that the Ignore Poison control is forced to 1 by the core when the core is configured as a Root-Port.</p> <p>Note that the following TLP types ignore the setting of this bit.</p> <p>Poisoned Configuration Type 0 writes is terminated in the core in all cases, independent of the Ignore Poison bit setting. A completion with UR status is generated and the appropriate error message, ERR COR or ERR FAT, is generated if not masked. Note that Poisoned Configuration Type 0 reads are always treated as if they were not poisoned. The read completes with successful completion status and an optional Advisory Non-Fatal Error status is set provided the severity level is set to NON-FATAL.</p> <p>Poisoned packets without data payload is passed to you in all cases since EP should not be set on packets without data payload and these packets should generally be handled as if they were not poisoned or alternatively handled as Advisory Non-Fatal Errors by user logic.</p> <p>Poisoned Vendor-defined Type 1 messages with data payload are always passed to you and, if ignore poison is 0, additionally an Advisory Non-Fatal Error status is set provided the severity level is set to NON-FATAL.</p> <p>When Ignore Poison is set to 0, the core handles the remaining poisoned TLPs with data payload as follows.</p> <p>Poisoned Write request and poisoned read completions with data TLPs are consumed by the core and handled as TLP Poisoned errors that generate the appropriate poison, ERR NON-FATAL or ERR FATAL, depending upon the error severity register error message. Poisoned Message with data payload (other than vendor-defined type 1) are consumed by the core and handled as TLP Poisoned errors that generate the appropriate poison, ERR NON-FATAL or ERR FATAL depending upon the error severity register, error message.</p> <p>The recommended default for target-only endpoints is to set Ignore Poison == 0 and to have user logic ignore the EP header bit on TLPs that it receives. In this case poisoned TLPs with data payload (other than config 0 writes and vendor-defined type 1 messages) generates a NON-FATAL error message and is discarded by the core. Poisoned TLPs without data payload (for which EP does not apply) is processed as if they were not poisoned.</p> <p>0 – Disable 1 – Enable</p>

decode_t1 Register 0x14

This register set is used for the Type 1 Configuration Space Transaction Layer Decode configuration.

Table 5.142. decode_t1 Register 0x14

Field	Name	Access	Width	Reset	Description
[31:24]	reserved	RO	8	0x0	—
[23:16]	reserved	RO	8	0x0	—
[15:11]	reserved	RO	5	0x0	—
[10]	bypass_addr_dec	RW	1	0x0	<p>When implementing Type 1 Configuration Space (Root Port): Bypass TLP Address Decode Enable. <code>bypass_addr_dec</code> controls the decoding of received Memory and I/O Request TLPs on both the Primary and Secondary sides of the core.</p> <p>0 – Normal PCI Express compliant address validity checks are performed. Transactions must target a valid region to be forwarded. Transactions with an invalid address cause an error response to be generated. Recommended value unless the customer application requires a value of 1 to be used.)</p> <p>1 – The address validity checks for I/O, Memory, and Prefetchable Memory regions are bypassed. Memory Requests, I/O Requests, and Messages Routed by Address is considered valid regardless of address. <code>bypass_addr_dec</code> does not affect the other validity checks (Memory and I/O regions enabled, in a Power State where transactions can be accepted, link is up, and not poisoned). Transactions, which fail these other validity checks, still causes an error response to be generated. When <code>bypass_addr_dec == 1</code>, Memory Requests which appear on the Receive Interface indicates a hit to the Prefetchable Memory Window and I/O Requests which appear on the Receive Interface indicates a hit to the I/O Window.</p>
[9]	tx_bypass_msg_dec	RW	1	0x0	<p>When implementing Type 1 Configuration Space (Root Port): Bypass TX Message TLP Decode Enable.</p> <p>0 – Normal operation. Claim and do not forward Local Routing messages and Error Messages that are transmitted without error propagation being enabled. Other message types are left in the stream.</p> <p>1 – All valid Message TLPs transmitted towards PCIe (except Routed by ID and Routed by Address, which are routed according to the routing type) are forwarded to PCIe.</p>
[8]	rx_bypass_msg_dec	RW	1	0x0	<p>When implementing Type 1 Configuration Space (Root Port): Bypass RX Message TLP Decode Enable.</p> <p>0 – Normal operation. Claim and do not forward Local Routing messages and Error Messages that are transmitted without error propagation being enabled. Other message types are left in the stream.</p> <p>1 – All valid Message TLPs received from PCIe (except Routed by ID and Routed by Address which are routed according to the routing type) are forwarded to the TLP Receive Interface.</p>
[7:0]	reserved	RO	8	0x0	—

5.1.4.3. Transaction Layer TLP Processing Configuration Register

tlp_processing Register 0x18

This register set is used for the Transaction Layer TLP Processing configuration.

Table 5.143. tlp_processing Register 0x18

Field	Name	Access	Width	Reset	Description
[31:24]	reserved	RO	8	0x0	—
[23:16]	reserved	RO	8	0x0	—
[15:8]	reserved	RO	8	0x0	—
[7:2]	reserved	RO	6	0x0	—
[1]	ignore_ecrc	RW	1	0x0	Ignore ECRC Error Enable. When enabled ECRC errors are ignored for TLPs passed to you in the TLP Receive Interface. 0 – Disable 1 – Enable
[0]	crs_enable	RW	1	0x0	Configuration Request Retry Status Enable. 0 – Disable. Type 0 Configuration Writes and Reads are performed normally. 1 – Enable. Type 0 Configuration Writes and Reads return Configuration Request Retry Status.

5.1.4.4. Initial Register

Initial Register 0x20

This register set is used for the initial speed and width configuration.

Table 5.144. Initial Register 0x20

Field	Name	Access	Width	Reset	Description
[31:19]	reserved	RO	13	0x0	—
[18:16]	max_link_width	RW	3	0x5	Max Link Width Override. This setting, if different from zero, overrides the value of Maximum Link Width in the PCIe Link Capabilities register. 0 – Maximum core lane width 1 – 1 lane 2 – 2 lanes 3 – 4 lanes 4 – 8 lanes 5 – 16 lanes 6 – Reserved6 7 – Reserved7
[15:2]	reserved	RO	14	0x0	—

Field	Name	Access	Width	Reset	Description
[1:0]	target_link_speed	RW	2	0x3	Initial value of Target Link Speed Configuration Register. Determines the maximum initial link speed which can be reached during initial training. Must be set to the lesser of the maximum speed supported by the core and the maximum speed at which you desired the core to operate. 0 – 2.5G 1 – 5.0G 2 – 8.0G 3 – 16.0G

5.1.4.5. Configuration Register type

cfg Register 0x30

This register set is used for the Configuration Register type.

Table 5.145. cfg Register 0x30

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	type1_type0_n	RW	1	0x0	Determines the type of Configuration Registers implemented by the core. 0 – Type 0 – Endpoint 1 – Type 1 – Root

5.1.4.6. Downstream Port Configuration

ds_port Register 0x34

This register set is used for the Downstream Port configuration.

Table 5.146. ds_port Register 0x34

Field	Name	Access	Width	Reset	Description
[31:17]	reserved	RO	15	0x0	—
[16]	rcb	RW	1	0x0	Read Completion Boundary (RCB). RCB value advertised when the core is operating as a Root Port.
[15:0]	id	RW	16	0x0	Root Port ID. This 16-bit field is used to define the ID used for PCIe Requester ID and Completer ID when the core is operating as a Root Port.

5.1.4.7. Upstream Port Configuration

us_port Register 0x38

This register set is used for the Upstream Port configuration.

Table 5.147. us_port Register 0x38

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	adv_target_link_speed	RW	1	0x0	For an upstream port, advertise the link speeds specified by the target_link_speed field rather than the maximum supported speed.

5.1.4.8. Device ID Configuration

id1 Register 0x40

This register set is used for the ID1 configuration.

Table 5.148. id1 Register 0x40

Field	Name	Access	Width	Reset	Description
[31:16]	device_id	RW	16	0xe004	Value returned when the Device ID Configuration Register is read.
[15:0]	vendor_id	RW	16	0x19aa	Value returned when the Vendor ID Configuration Register is read.

id2 Register 0x44

This register set is used for the ID2 configuration.

Table 5.149. id2 Register 0x44

Field	Name	Access	Width	Reset	Description
[31:16]	subsystem_id	RW	16	0xe004	Value returned when the Subsystem ID Configuration Register is read.
[15:0]	subsystem_vendor_id	RW	16	0x19aa	Value returned when the Subsystem Vendor ID Configuration Register is read.

id3 Register 0x48

This register set is used for the ID3 configuration.

Table 5.150. id3 Register 0x48

Field	Name	Access	Width	Reset	Description
[31:8]	class_code	RW	24	0x118000	Value returned when the Class Code Configuration Register is read. Must be set to the correct value for the type of device being implemented; see PCI Local Bus Specification Revision 2.3 Appendix D for details on setting Class Code.
[7:0]	revision_id	RW	8	0x4	Value returned when the Revision ID Configuration Register is read.

5.1.4.9. Cardbus Configuration

Cardbus Register 0x4c

This register set is used for the Cardbus configuration.

Table 5.151. Cardbus Register 0x4c

Field	Name	Access	Width	Reset	Description
[31:0]	cis_pointer	RW	32	0x0	Value returned when the Cardbus CIS Pointer Configuration Register is read. Set to 0x00000000 unless a Cardbus CIS structure is implemented in memory (which is rare), in which case set to the address of the CIS Structure.

5.1.4.10. Interrupt Configuration

Interrupt Register 0x50

This register set is used for the Interrupt configuration.

Table 5.152. Interrupt Register 0x50

Field	Name	Access	Width	Reset	Description
[31:10]	reserved	RO	22	0x0	—
[9:8]	pin	RW	2	0x0	Selects which legacy interrupt is used. 0 – INTA 1 – INTB 2 – INTC 3 – INTD
[7:1]	reserved	RO	7	0x0	—
[0]	disable	RW	1	0x0	Disable support for interrupts. 0 – Enable 1 – Disable

5.1.4.11. BAR Configuration

bar0 Register 0x60

This register set is used for the BAR0 configuration.

Table 5.153. bar0 Register 0x60

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffff000c	Configuration of BAR0 (Cfg address 0x10). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

bar1 Register 0x64

This register set is used for the BAR1 configuration.

Table 5.154. bar1 Register 0x64

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffffffff	Configuration of bar1 (Cfg address 0x14). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

bar2 Register 0x68

This register set is used for the BAR2 configuration.

Table 5.155. bar2 Register 0x68

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xfffffe00c	Configuration of bar2 (Cfg address 0x18). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

bar3 Register 0x6c

This register set is used for the BAR3 configuration.

Table 5.156. bar3 Register 0x6c

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xfffffffff	Configuration of bar3 (Cfg address 0x1C). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

bar4 Register 0x70

This register set is used for the BAR4 configuration.

Table 5.157. bar4 Register 0x70

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xfffffe00c	Configuration of bar4 (Cfg address 0x20). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

bar5 Register 0x74

This register set is used for the BAR5 configuration.

Table 5.158. bar5 Register 0x74

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xfffffffff	Configuration of bar5 (Cfg address 0x24). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.4.12. Expansion ROM configuration

exp_rom Register 0x78

This register set is used for the Expansion ROM configuration.

Table 5.159. exp_rom Register 0x78

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0x0	Configuration of exp_rom. Use to define a 32-bit Memory Expansion ROM region. If an Expansion ROM region is defined, the region must map to PCIe-compliant Expansion ROM code, or the device may fail to boot.

5.1.4.13. PCI Express configuration

pcie_cap Register 0x80

This register set is used for the PCI Express Capabilities configuration.

Table 5.160. pcie_cap Register 0x80

Field	Name	Access	Width	Reset	Description
[31:14]	reserved	RO	18	0x0	—
[13:9]	interrupt_message_number	RW	5	0x0	MSI/MSI-X vector is used for the interrupt message generated in association with any of the status bits of the PCI Express Capability structure.
[8]	slot_implemented	RW	1	0x0	Indicates that the Link associated with this Port is connected to a slot. This field is valid for Downstream Ports only.
[7:4]	device_port_type	RW	4	0x0	Indicates the specific type of this PCI Express Function. 0 – PCI Express Endpoint 1 – Legacy PCI Express Endpoint 2 – Reserved 3 – Reserved 4 – Root Port of PCI Express Root Complex 5 – Upstream Port of PCI Express Switch 6 – Downstream Port of PCI Express Switch 7 – PCI Express to PCI/PCI-X Bridge 8 – PCI/PCI-X to PCI Express Bridge 9 – Root Complex Integrated Endpoint 10 – Root Complex Event Collector 11 – Reserved 12 – Reserved 13 – Reserved 14 – Reserved 15 – Reserved
[3:0]	capability_version	RW	4	0x2	Indicates PCI-SIG defined PCI Express Capability structure version number. Must be set to 0x2.

pcie_dev_cap Register 0x84

This PCI Express Device Capabilities configuration.

Table 5.161. pcie_dev_cap Register 0x84

Field	Name	Access	Width	Reset	Description
[31:29]	reserved	RO	3	0x0	—
[28]	disable_flr_capability	RW	1	0x0	Function Level Reset Capability 0 – Enable 1 – Disable
[27:26]	reserved	RO	2	0x0	—
[25:18]	reserved	RO	8	0x0	—
[17:16]	reserved	RO	2	0x0	—
[15]	reserved	RO	1	0x0	—
[14:13]	reserved	RO	2	0x0	—
[12]	extended_tag_field_en_default	RW	1	0x1	Extended Tag Field Enable Default Value. PCIe Specification allows the Extended Tag Field Enable register to reset to either 1 or 0. This register determines the reset value. 0 – 5-bit Tag field enabled on reset 1 – 8-bit Tag field enabled on reset
[11:9]	endpoint_l1_acceptable_latency	RW	3	0x0	Endpoint L1 Acceptable Latency 0 – Maximum of 1 μ s. Must be 0 when not an Endpoint. 1 – Maximum of 2 μ s 2 – Maximum of 4 μ s 3 – Maximum of 8 μ s 4 – Maximum of 16 μ s 5 – Maximum of 32 μ s 6 – Maximum of 64 μ s 7 – No limit
[8:6]	endpoint_l0s_acceptable_latency	RW	3	0x0	Endpoint L0s Acceptable Latency 0 – Maximum of 64 ns. Must be 0 when not an Endpoint. 1 – Maximum of 128 ns 2 – Maximum of 256 ns 3 – Maximum of 512 ns 4 – Maximum of 1 μ s 5 – Maximum of 2 μ s 6 – Maximum of 4 μ s 7 – No limit
[5]	extended_tag_field_supported	RW	1	0x1	Extended Tag Field Supported 0 – 5-bit Tag field supported 1 – 8-bit Tag field supported
[4:3]	phantom_functions_supported	RW	2	0x0	Phantom Functions Supported 0 – No Function Number bits are used for Phantom Functions 1 – The most significant bit of the Function number in Requester ID is used for Phantom Functions 2 – The two most significant bits of Function Number in Requester ID are used for Phantom Functions 3 – All 3 bits of Function Number in Requester ID used for Phantom Functions.

Field	Name	Access	Width	Reset	Description
[2:0]	max_payload_size_supported	RW	3	0x2	Max Payload Size Supported 0 – 128 Bytes 1 – 256 Bytes 2 – 512 Bytes 3 – 1024 Bytes 4 – 2048 Bytes 5 – 4096 Bytes 6 – Reserved 7 – Reserved

pcie_link_cap Register 0x88

This register set is used for the PCI Express Link Capabilities configuration.

Table 5.162. pcie_link_cap Register 0x88

Field	Name	Access	Width	Reset	Description
[31:24]	port_number	RW	8	0x0	Indicates the PCI Express Port number for the PCI Express Link.
[23:18]	reserved	RO	6	0x0	—
[17:15]	l1_exit_latency	RW	3	0x7	L1 Exit Latency. The value reported indicates the length of time this Port requires to complete transition from ASPM L1 to L0. 0 – Less than 1 μ s 1 – 1 μ s to less than 2 μ s 2 – 2 μ s to less than 4 μ s 3 – 4 μ s to less than 8 μ s 4 – 8 μ s to less than 16 μ s 5 – 16 μ s to less than 32 μ s 6 – 32 μ s to 64 μ s 7 – More than 64 μ s
[14:12]	l0s_exit_latency	RW	3	0x7	L0s Exit Latency. The value reported indicates the length of time this Port requires to complete transition from ASPM L0s to L0. 0 – Less than 64 ns 1 – 64 ns to less than 128 ns 2 – 128 ns to less than 256 ns 3 – 256 ns to less than 512 ns 4 – 512 ns to less than 1 μ s 5 – 1 μ s to less than 2 μ s 6 – 2 μ s to 4 μ s 7 – More than 4 μ s
[11:10]	aspm_support	RW	2	0x3	Active State Power Management (ASPM) Support 0 – No ASPM Support 1 – L0s Supported 2 – L1 Supported 3 – L0s and L1 Supported
[9:0]	reserved	RO	10	0x0	—

pcie_link_stat Register 0x8c

This register set is used for the PCI Express Link Status configuration.

Table 5.163. pcie_link_stat Register 0x8c

Field	Name	Access	Width	Reset	Description
[31:13]	reserved	RO	19	0x0	—
[12]	slot_clock_configuration	RW	1	0x1	Indicates whether the component uses the physical reference clock that the platform provides on the connector. 0 – Using independent reference clock. 1 – Using reference clock provided by slot.
[11:0]	reserved	RO	12	0x0	—

pcie_slot_cap Register 0x90

This register set is used for the PCI Express Slot Capabilities configuration.

Table 5.164. pcie_slot_cap Register 0x90

Field	Name	Access	Width	Reset	Description
[31:19]	physical_slot_number	RW	13	0x1	Indicates whether the physical slot number attached to this Port. 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. This field must be initialized to zero for Ports connected to devices that are either integrated on the system board or integrated within the same silicon as the Root Port.
[18]	no_command_completed_support	RW	1	0x0	Indicates whether the slot generates software notification when an issued command is completed by the Hot-Plug Controller. This bit is only permitted to be 1 if the hot-plug capable Port can accept writes to all fields of the Slot Control register without delay between successive writes. 0 – Software notification provided. 1 – Software notification not provided.
[17]	em_interlock_present	RW	1	0x0	Indicates whether an Electromechanical Interlock is implemented on the chassis for this slot. 0 – Not Supported 1 – Supported
[16:15]	slot_power_limit_scale	RW	2	0x0	Slot Power Limit Scale. In combination with the Slot Power Limit Value, specifies the upper limit on power supplied by the slot or by other means to the adapter. Refer PCIe Specification section for details.

Field	Name	Access	Width	Reset	Description
[14:7]	slot_power_limit_value	RW	8	0xa	Slot Power Limit Value. In combination with the Slot Power Limit Scale, specifies the upper limit on power supplied by the slot or by other means to the adapter. Refer PCIe Specification section for details.
[6]	hot_plug_capable	RW	1	0x0	Indicates whether this slot can support hot-plug operations. 0 – Not Supported 1 – Supported
[5]	hot_plug_surprise	RW	1	0x0	Indicates whether an adapter present in this slot might be removed from the system without any prior notification. This is a form factor specific capability. This bit is an indication to the operating system to allow for such removal without impacting continued software operation. 0 – Hot Plug Surprise not possible 1 – Hot Plug Surprise possible
[4]	power_indicator_present	RW	1	0x0	Indicates whether a Power Indicator is electrically controlled by the chassis for this slot. 0 – Not Supported 1 – Supported
[3]	attention_indicator_present	RW	1	0x0	Indicates whether an Attention Indicator is electrically controlled by the chassis. 0 – Not Supported 1 – Supported
[2]	mrl_sensor_present	RW	1	0x0	Indicates whether a MRL Sensor is implemented on the chassis for this slot. 0 – Not Supported 1 – Supported
[1]	power_controller_present	RW	1	0x0	Indicates whether a software programmable Power Controller is implemented for this slot/adapter. 0 – Not Supported 1 – Supported
[0]	attention_button_present	RW	1	0x0	Indicates whether an Attention Button for this slot is electrically controlled by the chassis. 0 – Not Supported 1 – Supported

pcie_dev_cap2 Register 0x98

This register set is used for the PCI Express Device Capabilities 2 configuration.

Table 5.165. pcie_dev_cap2 Register 0x98

Field	Name	Access	Width	Reset	Description
[31:22]	reserved	RO	10	0x0	—
[21]	end_end_prefixes_supported	RW	1	0x0	End-End TLP Prefix Supported 0 – Not Supported 1 – Supported
[20:19]	reserved	RO	2	0x0	—
[18]	obff_supported	RW	1	0x0	OBFF Supported 0 – OBFF Not Supported 1 – OBFF supported using Message signaling only
[17:16]	reserved	RO	2	0x0	—
[15:8]	reserved	RO	8	0x0	—
[7:5]	reserved	RO	3	0x0	—
[4]	cpl_timeout_disable_supported	RW	1	0x1	Completion Timeout Disable Supported. Completion timeout is not implemented by the core, so the advertised value must match the capabilities of the connected design which is implementing completion timeouts. 0 – Not Supported 1 – Supported

Field	Name	Access	Width	Reset	Description
[3:0]	cpl_timeout_ranges_supported	RW	4	0x0	Completion Timeout Ranges Supported advertised value. Completion timeout is not implemented by the core, so the advertised value must match the capabilities of the connected design which is implementing completion timeouts. 0 – Completion Timeout programming not supported. Timeout value in the range 50 μ s to 50 ms is used. 1 – Range A (50 μ s to 10 ms) 2 – Range B (10 ms to 250 ms) 3 – Range A (50 μ s to 10 ms) and B (10 ms to 250 ms) 4 – Range B (10 ms to 250 ms) and C (250 ms to 4 s) 5 – Range A (50 μ s to 10 ms) and B (10 ms to 250 ms) and C (250 ms to 4 s) 6 – Range B (10 ms to 250 ms) and C (250 ms to 4 s) and D (4 s to 64 s) 7 – Range A (50 μ s to 10 ms) and B (10 ms to 250 ms) and C (250 ms to 4 s) and D (4 s to 64 s) 8 – Reserved 9 – Reserved 10 – Reserved 11 – Reserved 12 – Reserved 13 – Reserved 14 – Reserved 15 – Reserved

pcie_link_ctl2 Register 0xa0

This register set is used for the PCI Express Link Control 2 configuration.

Table 5.166. pcie_link_ctl2 Register 0xa0

Field	Name	Access	Width	Reset	Description
[31:7]	reserved	RO	25	0x0	—
[6]	selectable_deemphasis	RW	1	0x0	Selectable Deemphasis setting for Root Port only: When the Link is operating at 5.0 GT/s speed, this bit is used to control the transmit deemphasis of the link. 0 – -6 dB 1 – -3.5 dB
[5:0]	reserved	RO	6	0x0	—

5.1.4.14. Power Management configuration

pm_cap Register 0xc0

This register set is used for the Power Management Capabilities configuration.

Table 5.167. pm_cap Register 0xc0

Field	Name	Access	Width	Reset	Description
[31:16]	reserved	RO	16	0x0	—
[15:11]	pme_support	RW	5	0x1f	PME Support. Indicates the power states from which the function may generate a PME. For each power state {D3Cold, D3hot, D2, D1, D0}: 0 – PME# not supported 1 – PME# supported
[10]	d2_support	RW	1	0x1	D2 Power Management State support. 0 – Not supported 1 – Supported
[9]	d1_support	RW	1	0x1	D1 Power Management State support. 0 – Not supported 1 – Supported
[8:6]	aux_current	RW	3	0x0	Aux Current. Reports the 3.3Vaux auxiliary current requirements for the PCI function. See PCIe Specification for details. 0 – Self-powered 1 – 55 mA 2 – 100 mA 3 – 160 mA 4 – 220 mA 5 – 270 mA 6 – 320 mA 7 – 375 mA
[5]	dsi	RW	1	0x0	Device Specific Initialization. Indicates whether special initialization of this function is required (beyond the standard PCI configuration header) before the generic class device driver can use it. 0 – No Device Specific Initialization necessary. 1 – Function requires a device specific initialization sequence following transition to the D0 uninitialized state.
[4]	reserved	RO	1	0x0	—
[3]	pme_clock	RW	1	0x0	PME Clock. Does not apply to PCI Express and must be 0.
[2:0]	version	RW	3	0x3	PCI Power Management Interface Specification Version. Must be set to 0x3 to indicate revision 1.2 of the PCI Power Management Interface Specification.

pm Register 0xc4

This register set is used for the Power Management Control/Status configuration.

Table 5.168. pm Register 0xc4

Field	Name	Access	Width	Reset	Description
[31:24]	data	RW	8	0x0	—
[23]	pmcsr_bus_p_c_en	RW	1	0x0	—
[22]	pmcsr_b2_b3_support	RW	1	0x0	—
[21:16]	reserved	RO	6	0x0	—
[15]	reserved	RO	1	0x0	—
[14:13]	cstat_data_scale	RW	2	0x0	0 – Unknown scale 1 – power = data × 0.1 Watts 2 – power = data × 0.01 Watts 3 – power = data × 0.001 Watts
[12:9]	cstat_data_select	RW	4	0x0	0 – D0 Power Consumed 1 – D1 Power Consumed 2 – D2 Power Consumed 3 – D3 Power Consumed 4 – D0 Power Dissipated 5 – D1 Power Dissipated 6 – D2 Power Dissipated 7 – D3 Power Dissipated 8 – Common logic power consumption. For multifunction devices, reported in Function 0 only. 9 – Reserved 10 – Reserved 11 – Reserved 12 – Reserved 13 – Reserved 14 – Reserved 15 – Reserved
[8:0]	reserved	RO	9	0x0	—

pm_aux Register 0xc8

This register set is used for the Power Management Auxiliary Power configuration.

Table 5.169. pm_aux Register 0xc8

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	power_required	RW	1	0x0	<ul style="list-style-type: none"> Identifies whether the design requires auxiliary power. <ul style="list-style-type: none"> 0 – Aux Power is not required. 1 – Aux Power is required. If Aux Power is required, PME is advertised supported from D3 Cold, or advertised. aux_current != 0, then the value of Aux Power PM Enable is sticky and preserved through conventional reset when Aux Power is provided.

5.1.4.15. ARI Capability configuration

ari_cap Register 0xe0

This register set is used for the ARI Capability configuration.

Table 5.170. ari_cap Register 0xe0

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	disable	RW	1	0x0	ARI Capability Disable When disabled, the ARI Capability does not appear in PCIe Configuration Space. This must be enabled when SR-IOV is enabled and must be disabled for downstream ports, Root Complex Integrated Endpoints, and Root Complex Event Collectors.

aer_cap Register 0x100

This register set is used for the AER Capability configuration.

Table 5.171. aer_cap Register 0x100

Field	Name	Access	Width	Reset	Description
[31]	en_tlp_prefix_blocked	RW	1	0x0	Enable TLP Prefix Blocked error reporting. 0 – Disable 1 – Enable
[30]	en_atomicop_egress_blocked	RW	1	0x0	Enable AtomicOp Egress Blocked error reporting. 0 – Disable 1 – Enable
[29]	en_mc_blocked_tlp	RW	1	0x0	Enable MC Blocked TLP error reporting. Not supported by core, so must be 0. 0 – Disable 1 – Enable
[28]	en_uccorr_internal_error	RW	1	0x0	Enable Uncorrectable Internal Error. 0 – Disable 1 – Enable
[27]	en_acs_violation	RW	1	0x0	Enable ACS Violation error reporting. Not supported by core, so must be 0. 0 – Disable 1 – Enable

Field	Name	Access	Width	Reset	Description
[26]	en_receiver_overflow	RW	1	0x0	Enable Receiver Overflow error reporting. Not supported by core, so must be 0. 0 – Disable 1 – Enable
[25]	en_completer_abort	RW	1	0x0	Enable Completer Abort error reporting. 0 – Disable 1 – Enable
[24]	en_completion_timeout	RW	1	0x1	Enable Completion Timeout error reporting. 0 – Disable 1 – Enable
[23]	en_surprise_down_error	RW	1	0x0	Enable Surprise Down Error reporting. 0 – Disable 1 – Enable
[22]	en_corr_internal_error	RW	1	0x0	Enable Correctable Internal Error reporting. 0 – Disable 1 – Enable
[21:16]	reserved	RO	6	0x0	—
[15:2]	reserved	RO	14	0x0	—
[1]	ecrc_gen_chk_capable	RW	1	0x1	ECRC Generation/Checking Capable. 0 – Not supported 1 – Supported
[0]	version	RW	1	0x0	AER Capability Version. 0 – Version 0x1 1 – Version 0x2

5.1.4.16. MSI Capability configuration

msi_cap Register 0xe8

This register set is used for the MSI Capability configuration.

Table 5.172. msi_cap Register 0xe8

Field	Name	Access	Width	Reset	Description
[31:8]	reserved	RO	24	0x0	—
[7]	reserved	RO	1	0x0	—
[6:4]	mult_message_capable	RW	3	0x5	Number of requested MSI vectors. 0 – 1 1 – 2 2 – 4 3 – 8 4 – 16 5 – 32 6 – Reserved 7 – Reserved
[3:2]	reserved	RO	2	0x0	—
[1]	vec_mask_capable	RW	1	0x1	MSI Capability Per Vector Mask Capable. 0 – Disable 1 – Enable
[0]	disable	RW	1	0x0	MSI Capability Disable. When disabled, the MSI Capability does not appear in PCIe Configuration Space. 0 – Enable 1 – Disable

msix_cap Register 0xf0

This register set is used for the MSI-X Capability configuration.

Table 5.173. msix_cap Register 0xf0

Field	Name	Access	Width	Reset	Description
[31:27]	reserved	RO	5	0x0	—
[26:16]	table_size	RW	11	0x1f	Number of requested MSI-X vectors == (table_size+1).
[15:1]	reserved	RO	15	0x0	—
[0]	disable	RW	1	0x0	MSI-X Capability Disable. When disabled, the MSI-X Capability does not appear in PCIe Configuration Space. 0 – Enable 1 – Disable

msix_table Register 0xf4

This register set is used for the MSI-X Capability – MSI-X Table configuration.

Table 5.174. msix_table Register 0xf4

Field	Name	Access	Width	Reset	Description
[31:3]	offset	RW	29	0xc00	{offset, 3'b000} == byte address offset, within the BAR selected by bir, at which the MSI-X Table begins.
[2:0]	bir	RW	3	0x0	Indicates which Base Address register, located beginning at 10h in Configuration Space, is used to map the MSI-X Table into Memory Space. 0 – 0x10 (BAR0) 1 – 0x14 (BAR1) 2 – 0x18 (BAR2) 3 – 0x1C (BAR3) 4 – 0x20 (BAR4) 5 – 0x24 (BAR5) 6 – Reserved 7 – Reserved

msix_pba Register 0xf8

This register set is used for the MSI-X Capability – MSI-X PBA configuration.

Table 5.175. msix_pba Register 0xf8

Field	Name	Access	Width	Reset	Description
[31:3]	offset	RW	29	0xe00	{offset, 3'b000} == byte address offset, within the BAR selected by bir, at which the MSI-X PBA begins
[2:0]	bir	RW	3	0x0	Indicates which Base Address register, located beginning at 10h in Configuration Space, is used to map the MSI-X PBA into Memory Space. 0 – 0x10 (BAR0) 1 – 0x14 (BAR1) 2 – 0x18 (BAR2) 3 – 0x1C (BAR3) 4 – 0x20 (BAR4) 5 – 0x24 (BAR5) 6 – Reserved 7 – Reserved

5.1.4.17. Vendor-Specific Capability Configuration

vsec_cap Register 0x110

This register set is used for the Vendor-Specific Capability configuration.

Table 5.176. vsec_cap Register 0x110

Field	Name	Access	Width	Reset	Description
[31:16]	id	RW	16	0x1	Vendor-Specific Capability ID.
[15:1]	reserved	RO	15	0x0	—
[0]	enable	RW	1	0x1	Vendor-Specific Capability Enable. When disabled, the VSEC Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.4.18. SRIS Capability configuration

sris_cap Register 0x120

This register set is used for the SRIS Capability configuration.

Table 5.177. sris_cap Register 0x120

Field	Name	Access	Width	Reset	Description
[31:16]	reserved	RO	16	0x0	—
[15:12]	low_skp_generation_speeds	RW	4	0x0	SRIS Lower SKP OS Generation Supported Speeds Vector advertisement
[11:8]	low_skp_reception_speeds	RW	4	0x0	SRIS Lower SKP OS Reception Supported Speeds Vector advertisement
[7:1]	reserved	RO	7	0x0	—
[0]	enable	RW	1	0x0	SRIS Capability Enable. When disabled, the SRIS Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.4.19. Device Serial Number

dsn_cap Register 0x130

This register set is used for the DSN capable cores only such as Device Serial Number Capability configuration.

Table 5.178. dsn_cap Register 0x130

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	Device Serial Number Capability Enable. When disabled, the Device Serial Number Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

dsn_serial Register 0x134

This register set is used for the Device Serial Number Capability – Serial Number.

Table 5.179. dsn_serial Register 0x134

Field	Name	Access	Width	Reset	Description
[63:0]	number	RW	64	0x0	Device Serial Number.

5.1.4.20. Power Budgeting Capability configuration

pwr_budget_cap Register 0x150

This register set is used for the Power Budgeting Capability configuration.

Table 5.180. pwr_budget_cap Register 0x150

Field	Name	Access	Width	Reset	Description
[31:2]	reserved	RO	30	0x0	—
[1]	sys_alloc	RW	1	0x0	Power Budgeting System Allocated. 0 – Power Budget should use Power Budgeting Capability Values 1 – Power Budget is System Allocated
[0]	enable	RW	1	0x0	Power Budgeting Capability Enable. When disabled, the Power Budgeting Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.4.21. Dynamic Power Allocation Configuration

dpa_cap Register 0x158

This register set is used for the Dynamic Power Allocation Capability configuration.

Table 5.181. dpa_cap Register 0x158

Field	Name	Access	Width	Reset	Description
[31:24]	xlcy1	RW	8	0x0	Transition Latency Value 1. When the Transition Latency Indicator for a substate is 1, this value is multiplied by the Transition Latency Unit to determine the maximum Transition Latency for the substate.
[23:16]	xlcy0	RW	8	0x0	Transition Latency Value 0. When the Transition Latency Indicator for a substate is 0, this value is multiplied by the Transition Latency Unit to determine the maximum Transition Latency for the substate.
[15:14]	reserved	RO	2	0x0	—
[13:12]	pas	RW	2	0x0	Power Allocation Scale. The value of the substate Power Allocation Register is multiplied by the decoded value of this field to determine the power allocation of the substate. 0 – 10x 1 – 1x 2 – 0.1x 3 – 0.01x
[11:10]	reserved	RO	2	0x0	—

Field	Name	Access	Width	Reset	Description
[9:8]	tlunit	RW	2	0x0	Transition Latency Unit. The substate Transition Latency Value is multiplied by the decoded Transition Latency Unit to Determine the maximum Transition Latency for the substate. 0 – 1 ms 1 – 10 ms 2 – 100 ms 3 – Reserved
[7:3]	substate_max	RW	5	0x0	Substate_Max. Specifies the maximum substate number. Substates from [substate_max:0] are supported. For example, substate_max==0 indicates support for 1 substate.
[2:1]	reserved	RO	2	0x0	—
[0]	enable	RW	1	0x0	Dynamic Power Allocation (DPA) Capability Enable. When disabled, the Dynamic Power Allocation Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

dpa_xlcy Register 0x15c

This register set is used for the Dynamic Power Allocation – Transition Latency.

Table 5.182. dpa_xlcy Register 0x15c

Field	Name	Access	Width	Reset	Description
[31:0]	indicator	RW	32	0x0	Transition Latency Indicator. Indicates which Transition Latency Value applies to each substate. For each substate[i], indicator[i] indicates which Transition Latency Value applies: 0 – Use Transition Latency Value 0 1 – Use Transition Latency Value 1

dpa_alloc Register 0x160

This register set is used for the Dynamic Power Allocation Capability – Dynamic Power Allocation Array.

Table 5.183. dpa_alloc Register 0x160

Field	Name	Access	Width	Reset	Description
[255:0]	array	RW	256	0x0	Substate Power Allocation Array. For each substate[i], multiply array[(i*8)+7i*8] times the Power Allocation Scale to determine the power allocation in Watts for the associated substate.

5.1.4.22. Latency Tolerance Reporting Capability configuration.

ltr_cap Register 0x180

This register set is used for the Latency Tolerance Reporting Capability configuration.

Table 5.184. ltr_cap Register 0x180

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	Latency Tolerance Reporting Capability Enable. When disabled, the Latency Tolerance Reporting Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.4.23. L1 PM Substates Capability configuration

l1pmss_cap Register 0x188

This register set is used for the L1 PM Substates Capability configuration.

Table 5.185. l1pmss_cap Register 0x188

Field	Name	Access	Width	Reset	Description
[31:24]	cm_restore_time	RW	8	0x0	Default Common Mode Restore Time. Default time, in microseconds, used by the Downstream Port for timing the re-establishment of common mode. See the L1 PM Substates ECN for further details.
[23:16]	port_cm_restore_time	RW	8	0x0	Port Common Mode Restore Time. Time, in microseconds, required for this port to re-establish common mode. See the L1 PM Substates ECN for further details
[15:11]	port_tpower_on_value	RW	5	0x0	Port TPOWER_ON Value. Required for ports supporting PCI-PM L1.2 or ASPM L1.2. The value of TPOWER_ON is calculated by multiplying the value in this field by the decoded TPOWER_ON Scale field.
[10]	reserved	RO	1	0x0	—
[9:8]	port_tpower_on_scale	RW	2	0x0	Port TPOWER_ON Scale. Required for ports supporting PCI-PM L1.2 or ASPM L1.2. 0 – 2 μ s 1 – 10 μ s 2 – 100 μ s 3 – Reserved
[7]	pcipm_l1_1_supported	RW	1	0x1	PCI-PM L1.1 Substate Supported. Must be set to 1 for all ports supporting L1 PM Substates. 0 – Not supported 1 – Supported
[6]	pcipm_l1_2_supported	RW	1	0x1	PCI-PM L1.2 Substate Supported. 0 – Not supported 1 – Supported
[5]	aspm_l1_1_supported	RW	1	0x1	ASPM L1.1 Substate Supported. 0 – Not supported 1 – Supported
[4]	aspm_l1_2_supported	RW	1	0x1	ASPM L1.2 Substate Supported. 0 – Not supported 1 – Supported
[3]	l1pm_supported	RW	1	0x1	L1 PM Substates Supported. 0 – Not supported 1 – Supported

Field	Name	Access	Width	Reset	Description
[2:1]	reserved	RO	2	0x0	—
[0]	enable	RW	1	0x0	L1 PM Substates Capability Enable. When disabled, the L1 PM Substates Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.4.24. Atomic Op Capability configuration.

atomic_op_cap Register 0x1cc

This register set is used for the Atomic Op Capability configuration.

Table 5.186. atomic_op_cap Register 0x1cc

Field	Name	Access	Width	Reset	Description
[31:6]	reserved	RO	26	0x0	—
[5]	rp_completer_enable	RW	1	0x0	Enable Root Port to be an Atomic Op Completer which means that the Root Port completes rather than forwards Atomic Op TLPs. 0 – Disable 1 – Enable
[4]	completer_128_supported	RW	1	0x0	Atomic Op Completer 128-bit Operand Support. 0 – Not Supported 1 – Supported
[3]	completer_64_supported	RW	1	0x0	Atomic Op Completer 64-bit Operand Support. 0 – Not Supported 1 – Supported
[2]	completer_32_supported	RW	1	0x0	Atomic Op Completer 32-bit Operand Support. 0 – Not Supported 1 – Supported
[1]	routing_supported	RW	1	0x0	Atomic Op Routing Supported. 0 – Not Supported 1 – Supported
[0]	enable	RW	1	0x0	Atomic Op Capability Enable. When disabled, the Atomic Op Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.5. mgmt_ftl_mf[7:1] (0x4_5000,0x4_6000,0x4_7000,0x4_8000,0x4_9000,0x4_A000,0x4_B000)

The base address for mgmt._ftl_mf is shown in [Table 5.187](#).

Table 5.187. Base Address for mgmt_ftl_mf

Port	Base Address
mgmt_ftl_mf1_BASE	0x4_5000
mgmt_ftl_mf2_BASE	0x4_6000
mgmt_ftl_mf3_BASE	0x4_7000
mgmt_ftl_mf4_BASE	0x4_8000
mgmt_ftl_mf5_BASE	0x4_9000
mgmt_ftl_mf6_BASE	0x4_A000

Port	Base Address
mgmt_ftl_mf7_BASE	0x4_B000

5.1.5.1. Function Register 0x8

This register set is used for the Function disable for Functions[3:1]. Function[0] may not be disabled.

Table 5.188. Function Register 0x8

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	disable	RW	1	0x0	Function disable for Functions[7:1]. Function[0] may not be disabled. 0 – Enable 1 – Disable

5.1.5.2. us_port Register 0x38

This register set is used for the Upstream Port configuration.

Table 5.189. us_port Register 0x38

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	adv_target_link_speed	RW	1	0x0	For an upstream port, advertise the link speeds specified by the target_link_speed field rather than the maximum supported speed.

5.1.5.3. id1 Register 0x40

This register set is used for the ID1 configuration.

Table 5.190. id1 Register 0x40

Field	Name	Access	Width	Reset	Description
[31:16]	device_id	RW	16	0xe004	Value returned when the Device ID Configuration Register is read
[15:0]	vendor_id	RW	16	0x19aa	Value returned when the Vendor ID Configuration Register is read

5.1.5.4. id2 Register 0x44

This register set is used for the ID2 configuration.

Table 5.191. id2 Register 0x44

Field	Name	Access	Width	Reset	Description
[31:16]	subsystem_id	RW	16	0xe004	Value returned when the Subsystem ID Configuration Register is read.
[15:0]	subsystem_vendor_id	RW	16	0x19aa	Value returned when the Subsystem Vendor ID Configuration Register is read.

5.1.5.5. id3 Register 0x48

This register set is used for the ID3 configuration.

Table 5.192. id3 Register 0x48

Field	Name	Access	Width	Reset	Description
[31:8]	class_code	RW	24	0x118000	Value returned when the Class Code Configuration Register is read. Must be set to the correct value for the type of device being implemented. Refer to PCI Local Bus Specification Revision 2.3 Appendix D for details on setting Class Code.
[7:0]	revision_id	RW	8	0x4	Value returned when the Revision ID Configuration Register is read.

5.1.5.6. Cardbus Register 0x4c

This register set is used for the Cardbus configuration.

Table 5.193. Cardbus Register 0x4c

Field	Name	Access	Width	Reset	Description
[31:0]	cis_pointer	RW	32	0x0	Value returned when the Cardbus CIS Pointer Configuration Register is read. Set to 0x00000000 unless a Cardbus CIS structure is implemented in memory (which is rare), in which case set to the address of the CIS Structure.

5.1.5.7. Interrupt Register 0x50

This register set is used for the Interrupt configuration.

Table 5.194. Interrupt Register 0x50

Field	Name	Access	Width	Reset	Description
[31:10]	reserved	RO	22	0x0	—
[9:8]	pin	RW	2	0x0	Selects which legacy interrupt is used. 0 – INTA 1 – INTB 2 – INTC 3 – INTD
[7:1]	reserved	RO	7	0x0	—
[0]	disable	RW	1	0x0	Disable support for interrupts. 0 – Enable 1 – Disable

5.1.5.8. bar0 Register 0x60

This register set is used for the BAR0 configuration.

Table 5.195. bar0 Register 0x60

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffff000c	Configuration of BAR0 (Cfg address 0x10). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.5.9. bar1 Register 0x64

This register set is used for the BAR1 configuration.

Table 5.196. bar1 Register 0x64

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffffffff	Configuration of bar1 (Cfg address 0x14). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.5.10. bar2 Register 0x68

This register set is used for the BAR2 configuration.

Table 5.197. bar2 Register 0x68

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffffe00c	Configuration of bar2 (Cfg address 0x18). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.5.11. bar3 Register 0x6c

This register set is used for the BAR3 configuration.

Table 5.198. bar3 Register 0x6c

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffffffff	Configuration of bar3 (Cfg address 0x1C). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.5.12. bar4 Register 0x70

This register set is used for the BAR4 configuration.

Table 5.199. bar4 Register 0x70

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffffe00c	Configuration of bar4 (Cfg address 0x20). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.5.13. bar5 Register 0x74

This register set is used for the BAR5 configuration.

Table 5.200. bar5 Register 0x74

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0xffffffff	Configuration of bar5 (Cfg address 0x24). Use to define a 32-bit Memory or I/O region or combine with an adjacent BAR to define a 64-bit Memory region.

5.1.5.14. exp_rom Register 0x78

This register set is used for the Expansion ROM configuration.

Table 5.201. exp_rom Register 0x78

Field	Name	Access	Width	Reset	Description
[31:0]	cfg	RW	32	0x0	Configuration of exp_rom. Use to define a 32-bit Memory Expansion ROM region. If an Expansion ROM region is defined, the region must map to PCIe-compliant Expansion ROM code or the device may fail to boot.

5.1.5.15. msi_cap Register 0xe8

This register set is used for the MSI Capability configuration.

Table 5.202. msi_cap Register 0xe8

Field	Name	Access	Width	Reset	Description
[31:8]	reserved	RO	24	0x0	—
[7]	reserved	RO	1	0x0	—
[6:4]	mult_message_capable	RW	3	0x5	Number of requested MSI vectors. 0 – 1 1 – 2 2 – 4 3 – 8 4 – 16 5 – 32 6 – Reserved 7 – Reserved
[3:2]	reserved	RO	2	0x0	—
[1]	vec_mask_capable	RW	1	0x1	MSI Capability Per Vector Mask Capable. 0 – Disable 1 – Enable
[0]	disable	RW	1	0x0	MSI Capability Disable. When disabled, the MSI Capability does not appear in PCIe Configuration Space. 0 – Enable 1 – Disable

5.1.5.16. msix_cap Register 0xf0

This register set is used for the MSI-X Capability configuration.

Table 5.203. msix_cap Register 0xf0

Field	Name	Access	Width	Reset	Description
[31:27]	reserved	RO	5	0x0	—
[26:16]	table_size	RW	11	0x1f	Number of requested MSI-X vectors == (table_size+1)
[15:1]	reserved	RO	15	0x0	—
[0]	disable	RW	1	0x0	MSI-X Capability Disable. When disabled, the MSI-X Capability does not appear in PCIe Configuration Space. 0 – Enable 1 – Disable

5.1.5.17. msix_table Register 0xf4

This register set is used for the MSI-X Capability – MSI-X Table configuration.

Table 5.204. msix_table Register 0xf4

Field	Name	Access	Width	Reset	Description
[31:3]	offset	RW	29	0xc00	{offset, 3'b000} == byte address offset, within the BAR selected by bir, at which the MSI-X Table begins
[2:0]	bir	RW	3	0x0	Indicates which Base Address register, located beginning at 10h in Configuration Space, is used to map the MSI-X Table into Memory Space. 0 – 0x10 (BAR0) 1 – 0x14 (BAR1) 2 – 0x18 (BAR2) 3 – 0x1C (BAR3) 4 – 0x20 (BAR4) 5 – 0x24 (BAR5) 6 – Reserved 7 – Reserved

5.1.5.18. msix_pba Register 0xf8

This register set is used for the MSI-X Capability – MSI-X PBA configuration.

Table 5.205. msix_pba Register 0xf8

Field	Name	Access	Width	Reset	Description
[31:3]	offset	RW	29	0xe00	{offset, 3'b000} == byte address offset, within the BAR selected by bir, at which the MSI-X PBA begins.
[2:0]	bir	RW	3	0x0	Indicates which Base Address register, located beginning at 10h in Configuration Space, is used to map the MSI-X PBA into Memory Space. 0 – 0x10 (BAR0) 1 – 0x14 (BAR1) 2 – 0x18 (BAR2) 3 – 0x1C (BAR3) 4 – 0x20 (BAR4) 5 – 0x24 (BAR5) 6 – Reserved 7 – Reserved

5.1.5.19. dsn_cap Register 0x130

This register set is used for the DSN capable cores only such as Device Serial Number Capability configuration.

Table 5.206. dsn_cap Register 0x130

Field	Name	Access	Width	Reset	Description
[31:1]	reserved	RO	31	0x0	—
[0]	enable	RW	1	0x0	Device Serial Number Capability Enable. When disabled, the Device Serial Number Capability does not appear in PCIe Configuration Space. 0 – Disable 1 – Enable

5.1.5.20. dsn_serial Register 0x134

This register set is used for the Device Serial Number Capability – Serial Number.

Table 5.207. dsn_serial Register 0x134

Field	Name	Access	Width	Reset	Description
[63:0]	number	RW	64	0x0	Device Serial Number

5.2. PCI Express Configuration Space Registers

The Lattice PCIe x8 IP Core implements Header Type 00 and Header Type 01 Configuration Registers, including Capability and Extended Capability Items, as detailed in the PCI Express Base Specification, Rev 3.0, PCI Local Interface Specification Revision 3.0, and PCI Bus Power Management Interface Specification Revision 1.2.

Type 00 and Type 01 Configuration Registers implement the first 64 bytes of Configuration Space differently:

- Type 00 – Implemented by Endpoints; refer to [Table 5.208](#).
- Type 01 – Implemented by Root Ports; refer to [Table 5.209](#).

Capability and Extended Capability Items are located at the same addresses regardless of which the header type is implemented, see [Table 5.210](#) for details.

[Table 5.208](#), [Table 5.209](#), and [Table 5.210](#) illustrate the core's PCIe Configuration Register map.

The core's Configuration Registers are highly configurable. In dual-mode (Root-Port/Endpoint) applications, the registers configure themselves according to the mode of operation, changing between Type 00 and Type 01 for instance when changing between an Endpoint and Root Port design.

The Configuration Registers provide the ability for standard PCI/PCIe BIOS/OS software to discover the device, determine its capabilities, and configure the core's features. Since there are a tremendous variety of applications, the core's Configuration Registers are highly configurable.

5.2.1. Type 00 Configuration Header

Table 5.208. Type 00 Configuration Header

Addr	Byte3	Byte2	Byte1	Byte0
00	Device ID		Vendor ID	
04	Status		Command	
08	Class Code			Revision ID
0C	BIST	Header Type	Latency Timer	Cache Line Size
10	Base Address Register 0			
14	Base Address Register 1			
18	Base Address Register 2			
1C	Base Address Register 3			
20	Base Address Register 4			
24	Base Address Register 5			
28	Cardbus CIS Pointer			
2C	Subsystem ID		Subsystem Vendor ID	
30	Expansion ROM Base Address			
34	Reserved			Capabilities Pointer
38	Reserved			
3C	Max Latency	Min Grant	Interrupt Pin	Interrupt Line

5.2.2. Type 01 Configuration Header

Table 5.209. Type 01 Configuration Header

Addr	Byte3	Byte2	Byte1	Byte0
00	Device ID		Vendor ID	
04	Status		Command	
08	Class Code			Revision ID
0C	BIST	Header Type	Primary Latency Timer	Cache Line Size
10	Base Address Register 0			
14	Base Address Register 1			
18	Secondary Latency Timer	Subordinate Bus Number	Secondary Bus Number	Primary Bus Number
1C	Secondary Status		I/O Limit	I/O Base
20	Memory Limit		Memory Base	
24	Prefetchable Memory Limit		Prefetchable Memory Base	
28	Prefetchable Base Upper 32 Bits			
2C	Prefetchable Limit Upper 32 Bits			
30	I/O Limit Upper 16 Bits		I/O Base Upper 16 Bits	
34	Reserved			Capability Pointer
38	Expansion ROM Base Address			
3C	Bridge Control		Interrupt Pin	Interrupt Line

5.2.3. Capability and Extended Capability Address Locations

Table 5.210. Capability and Extended Capability Items

Addr	Byte3	Byte2	Byte1	Byte0
7B-40	PCI Express Capability			
7F-7C	Reserved			
87-80	Power Management Capability			
8F-88	Reserved			
9B-90	MSI-X Capability			
9F-9C	Reserved			
B7-A0	MSI Capability			
FF-B8	Reserved			
147-100	Advanced Error Reporting Capability			
14F-148	ARI Capability			
17F-150	Vendor-Specific Extended Capability			
1AB-180	Secondary PCI Express Extended Capability			
1FF-1AC	Reserved			
207-200	Reserved			
20F-208	Reserved			
21B-210	DSN Capability			
26B-240	Reserved			
2BF-280	Reserved			
38F-2C0	Reserved			
39F-390	Power Budgeting Capability			
3CF-3A0	Dynamic Power Allocation (DPA) Capability			
3DF-3D0	L1 PM Substates Extended Capability			
3E7-3E0	Latency Tolerance Reporting (LTR) Capability			
FFF-3E8	Reserved			

5.2.4. Type 00 Configuration Registers

Table 5.211. Type 00 Configuration Registers

Addr	Config Register	Register Description
01–00	Vendor ID	Read Only: This field identifies the manufacturer of the device.
03–02	Device ID	Read Only: This field identifies the device.
05–04	Command Register	<p>Command Register Bits: Bits 10, 8, 6, and 2..0 are Read/Write.</p> <p>Bits[15:11] = 00000. Not implemented.</p> <p>Bit[10] – Interrupt Disable – If set, interrupts are disabled and cannot be generated; if clear interrupts are enabled</p> <p>Bit[9] = 0. Not implemented.</p> <p>Bit[8] – SERR Enable – When set enables the reporting of fatal and non-fatal errors detected by the device to the root complex (not supported).</p> <p>Bit[7] = 0. Not implemented.</p> <p>Bit[6] – Parity Error Enable – Affects the mapping of PCI Express errors to legacy PCI errors. See <i>PCI Express Base Specification Rev1.1</i>, Section 6.2 for details.</p> <p>Bit[5] = 0. Not implemented.</p> <p>Bit[4] = 0. Not implemented.</p> <p>Bit[3] = 0. Not implemented.</p> <p>Bit[2] – Bus Leader Enable – Memory and I/O Requests can only be generated on the Transaction Layer Interface if this bit is set.</p> <p>Bit[1] – Memory Space Enable – If set, the core decodes the packets to determine memory BAR hits; if clear, memory BARs are disabled.</p> <p>Bit[0] – I/O Space Enable – If set, the core decodes the packets to determine I/O BAR hits; if clear, I/O BARs are disabled.</p>
07–06	Status Register	<p>Status Register Bits: Bits 15...11 and 8 are Read/Write. Writing a 1 to a bit location clears that bit. Writing a 0 to a bit location has no affect.</p> <p>Bit[15] – Set by a device whenever it receives a Poisoned TLP.</p> <p>Bit[14] – Set when a device sends an ERR_FATAL or ERR_NONFATAL Message and the SERR Enable bit in the Command Register is set.</p> <p>Bit[13] – Set when a requestor receives a completion with Unrecognized Request Completion Status</p> <p>Bit[12] – Set when a requestor receives a completion with Completer Abort Completion Status</p> <p>Bit[11] – Set when a device completes a request using Completer Abort Completion Status</p> <p>Bits[10:9] = 00. Not implemented.</p> <p>Bit[8] – Leader Data Parity Error – This bit is set by a Requestor if its Parity Error Enable bit is set and either a Completion is received that is marked poisoned or the requestor poisons a write request.</p> <p>Bits[7:5] = 000. Not implemented.</p> <p>Bit[4] = 1 to indicate the presence of a Capabilities List.</p> <p>Bit[3] – Interrupt Status – Reflects the value of mgmt_interrupt.</p> <p>Bits[2:0] = 000. Reserved.</p>
08	Revision ID	Read Only: This register specifies the device specific revision identifier.
0B–09	Class Code	Read Only: The Class Code identifies the generic function of the device.
0C	0x0C: Cache Line Size	Read/Write: Cache Line Size is not used with PCI Express but is still implemented as read/write register for legacy compatibility purposes.
0D	0x0D: Latency Timer	Read Only returning 0x00.
0E	0x0E: Header Type	Read Only: This register reads 0x00 to indicate that the core complies to the standard PCI configuration register mapping and that it is a single function device.
0F	0x0F: BIST	Not implemented. Reads return 0x00.
13–10	Base Address Register 0	Read/Write: Base Address Register0, Base Address Register1, Base Address Register2, Base Address Register3, Base Address Register4, and Base Address Register5 inform system software of the device's resource requirements and are subsequently programmed to allocate memory and I/O resources to the device.

Addr	Config Register	Register Description
17–14	Base Address Register 1	See Base Address Register 0 description.
1B–18	Base Address Register 2	See Base Address Register 0 description.
1F–1C	Base Address Register 3	See Base Address Register 0 description.
23–20	Base Address Register 4	See Base Address Register 0 description.
27–24	Base Address Register 5	See Base Address Register 0 description.
2B–28	Card Bus CIS Pointer	Read Only: Reads return the value of the Cardbus CIS Pointer.
2D–2C	Subsystem Vendor ID	Read Only: Additional vendor information. Reads return the value of the Subsystem Vendor ID.
2F–2E	Subsystem ID	Read Only: Additional device information. Reads return the value of the Subsystem ID.
33–30	Expansion ROM Base Addr. Reg.	<p>Informs system software of the device's Expansion ROM resource requirements and is subsequently programmed to allocate memory resources to the device.</p> <p>Read/Write: Expansion ROM Base Address Register</p> <p>Bits[31:11] – Written to specify where to locate this region in memory space</p> <p>Bits[10:1] = 0..0 Reserved</p> <p>Bit[0] = Set by S/W to enable decoding the Expansion ROM and clear to disable</p>
34	Capabilities Pointer	Read Only: Reads return 0x40 which is the beginning address of the PCI Express Capabilities Item.
37–35	Reserved	Not implemented. Reads return 0x000000.
3B–38	Reserved	Not implemented. Reads return 0x00000000.
3C	Interrupt Line	Legacy interrupt is always ENABLED.
3D	Interrupt Pin	Interrupt support is enabled/disabled by CSR register. When interrupts are enabled, Interrupt Pin returns 0x01 indicating the core implements INTA# and when interrupts are disabled, Interrupt Pin returns 0x00 indicating no interrupts are used.
3E	Minimum Grant	Read Only: Returns 0x00.
3F	Maximum Latency	Read Only: Returns 0x00.

5.2.5. PCI Express Capability

Table 5.212. PCI Express Capability

Addr	Config Register	Register Description
40	PCI Express Capability ID	Read Only = 0x10 (Beginning of PCI Express Capability Item)
41	Next Capability Pointer	Read Only = 0x80 (Pointer to beginning of Power Management Capability)
43-42	PCI Express Capabilities	<p>Read Only</p> <ul style="list-style-type: none"> • Bits[15:14] – Reserved = 00 • Bits[13:9] – Interrupt Message Number[4:0]; MSI/MSI-X interrupt vector associated with interrupts generated by Configuration Register events (change in link bandwidth and root port error) • Bit[8] – Slot Implemented; Downstream Switch/Root Port only • Bits[7:4] – Device/Port Type – Must match the core application since the value programmed enables/hides Configuration Registers and functionality that is only applicable to some Device/Port types: <ul style="list-style-type: none"> • 0000 – PCI Express Endpoint • Required for Endpoint applications

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> 0001 – Legacy PCI Express Endpoint 0100 – Root Port of PCI Express Root Complex (future release) Required for Root Port applications 0101 – Upstream Port of PCI Express Switch Required for Upstream Switch Ports 0110 – Downstream Port of PCI Express Switch Required for Downstream Switch Ports 0111 – PCI Express to PCI/PCI-X Bridge 1000 – PCI/PCI-X to PCI Express Bridge 1001 – Root Complex Integrated Endpoint (future release) 1010 – Root Complex Event Collector (future release) Bits[3:0] – Capability Version – Must be 0x2 for PCIe 3.0
47-44	Device Capabilities Register	<p>Read Only</p> <ul style="list-style-type: none"> Bits[31:29] – Reserved. Bit[28] – Function Level Reset Capability <ul style="list-style-type: none"> 1 – Capability Present 0 – Capability Not Present Bits[27:26] – Captured Slot Power Limit Scale Bits[25:18] – Captured Slot Power Limit Value Bits[17:16] = 00. Reserved. Bit[15] = 1. Role-based Error Reporting Bit[14] = 0 – Reserved Bit[13] = 0 – Reserved Bit[12] = 0 – Reserved Bits[11:9] – Endpoint L1 Acceptable Latency Bit[8:6] – Endpoint L0s Acceptable Latency Bit[5] – Extended Tag Field Supported Bits[4:3] – Phantom Functions Supported Bits[2:0] – Max Payload Size Supported <ul style="list-style-type: none"> 000 – 128 bytes max payload size 001 – 256 bytes max payload size 010 – 512 bytes max payload size 011 – 1024 bytes max payload size 100 – 2048 bytes max payload size 101 – 4096 bytes max payload size 110 – Reserved 111 – Reserved
49-48	Device Control Register	<p>Read/Write</p> <ul style="list-style-type: none"> Bit[15] – Bridge Configuration Retry Enable/Initiate Function Level Reset Bits[14:12] – Max Read Request Size; the Transmit Interface may not transmit a read request TLP with a length larger than the size indicated by Max Read Request Size: <ul style="list-style-type: none"> 000 == 128 bytes 001 == 256 bytes 010 == 512 bytes 011 == 1024 bytes 100 == 2048 bytes 101 == 4096 bytes 110 == Reserved 111 == Reserved Bit[11] – Enable No Snoop

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> • Bit[10] – Aux Power PM Enable • Bit[9] – Phantom Functions Enable • Bit[8] – Extended Tag Field Enable • Bits[7:5] – Max Payload Size; the Transmit Interface may not transmit a TLP with a payload larger than the size indicated by Max Payload Size: <ul style="list-style-type: none"> • 000 == 128 bytes • 001 == 256 bytes • 010 == 512 bytes • 011 == Reserved • 100 == Reserved • 101 == Reserved • 110 == Reserved • 111 == Reserved • Bit[4] – Enable Relaxed Ordering • Bit[3] – Unsupported Request Reporting Enable • Bit[2] – Fatal Error Reporting Enable • Bit[1] – Non-Fatal Error Reporting Enable • Bit[0] – Correctable Error Reporting Enable
4B-4A	Device Status Register	<p>Bits[15:4] are Read Only. Bits[3:0] are cleared by writing a 1 to the corresponding bit location.</p> <ul style="list-style-type: none"> • Bits[15:6] = 0000000000. Reserved • Bit[5] – Transactions Pending • Bit[4] – AUX Power Detected • Bit[3] – Unsupported Request Detected • Bit[2] – Fatal Error Detected • Bit[1] – Non-Fatal Error Detected • Bit[0] – Correctable Error Detected
4F-4C	Link Capabilities Register	<p>Read Only.</p> <ul style="list-style-type: none"> • Bits[31:24] – Port Number • Bits[22] = 1. ASPM Optional Compliance • Bit[21] – Link Bandwidth Notification Capability <ul style="list-style-type: none"> • == 1 when operating as a Downstream Port; else 0 • Bit[20] – Data Link Layer Active Reporting Capable <ul style="list-style-type: none"> • == 1 when operating as a Downstream Port; else 0 • Bit[19] – Surprise Down Error Reporting Capable <ul style="list-style-type: none"> • == 1 when operating as a Downstream Port; else 0 • Bit[18] = 0. Clock Power Management • Bits[17:15] – L1 Exit Latency • Bits[14:12] – L0s Exit Latency • Bits[11:10] – Active State Power Management (ASPM) Support <ul style="list-style-type: none"> • 00 – No ASMP Support • 01 – L0s Supported • 10 – L1 Supported • 11 – L0s and L1 Supported • Bits[9:4] – Maximum Link Width <ul style="list-style-type: none"> • 000001 – x1 • 000010 – x2 • 000100 – x4 • 001000 – x8 • 010000 – x16 • Bits[3:0] – Maximum Link Speed

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> 0001 (2.5GT/s) 0010 (5GT/s) 0011 (8GT/s) (2024) 0100 (16GT/s) (2024)
51-50	Link Control Register	<p>Read/Write</p> <ul style="list-style-type: none"> Bits[15:12] = 0000. Reserved. Bit[11] – Link Autonomous Bandwidth Interrupt Enable Bit[10] – Link Bandwidth Management Interrupt Enable Bit[9] – Hardware Autonomous Width Disable Bit[8] = 0. Enable Clock Power Management Bit[7] – Extended Sync Bit[6] – Common Clock Configuration Bit[5] – Retrain Link Bit[4] – Link Disable Bit[3] – Read Completion Boundary (RCB) <ul style="list-style-type: none"> 0 == 64 bytes 1 == 128 bytes Bit[2] = 0. Reserved. Bits[1:0] – Active State Power Management (ASPM) Control <ul style="list-style-type: none"> 00 – Disabled 01 – L0s Enabled 10 – L1 Enabled 11 – L0s and L1 Enabled
53-52	Link Status Register	<p>Read Only</p> <ul style="list-style-type: none"> Bit[15] – Link Autonomous Bandwidth Status Bit[14] – Link Bandwidth Management Status Bit[13] – Data Link Layer Active Bit[12] – Slot Clock Configuration Bit[11] – Link Training Bit[10] = 0. Reserved. Bits[9:4] Negotiated Link Width – indicates the number of lanes currently in use <ul style="list-style-type: none"> 010000 = x16 001000 = x8 000100 = x4 000010 = x2 000001 = x1 Bits[3:0] Link Speed <ul style="list-style-type: none"> 0001 (2.5 GT/s) 0010 (5.0 GT/s) 0011 (8.0 GT/s)
57-54	Slot Capabilities Root Port/Switch Only	<p>Normally Read Only; Writable when HW.Init Write Enable == 1 (see Table 5.218).</p> <ul style="list-style-type: none"> Bits[31:19] – Physical Slot Number Bit[18] – No Command Completed Support Bit[17] – Electromechanical Interlock Present Bits[16:15] – Slot Power Limit Scale[1:0] Bits[14:7] – Slot Power Limit Value[7:0] Bit[6] – Hot-Plug Capable Bit[5] – Hot-Plug Surprise Bit[4] – Power Indicator Present Bit[3] – Attention Indicator Present

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> • Bit[2] – MRL Sensor Present • Bit[1] – Power Controller Present • Bit[0] – Attention Button Present
59-58	Slot Control Root Port/Switch Only	Read Only <ul style="list-style-type: none"> • Bits[15:13] = 0. Reserved. • Bit[12] – Data Link Layer State Changed Enable • Bit[11] = 0. Electromechanical Interlock Control • Bit[10] – Power Controller Control • Bit[9:8] – Power Indicator Control • Bit[7:6] – Attention Indicator Control • Bit[5] – Hot-Plug Interrupt Enable • Bit[4] – Command Completed Interrupt Enable • Bit[3] – Presence Detect Changed Enable • Bit[2] – MRL Sensor Changed Enable • Bit[1] – Power Fault Detected Enable • Bit[0] – Attention Button Pressed Enable
5b-5a	Slot Status Root Port/Switch Only	Read Only <ul style="list-style-type: none"> • Bits[15:9] = 0. Reserved. • Bit[8] – Data Link Layer State Changed • Bit[7] – Electromechanical Interlock Status • Bit[6] – Presence Detect State • Bit[5] – MRL Sensor State • Bit[4] – Command Completed • Bit[3] – Presence Detect Changed • Bit[2] – MRL Sensor Changed • Bit[1] – Power Fault Detected • Bit[0] – Attention Button Pressed
5d-5c	Root Control Root Port Only	Read/Write <ul style="list-style-type: none"> • Bits[15:5] = 0. Reserved. • Bit[4] = CRS Software Visibility Enable • Bit[3] – PME Interrupt Enable • Bit[2] – System Error on Fatal Error Enable • Bit[1] – System Error on Non-Fatal Error Enable • Bit[0] – System Error on Correctable Error Enable
5f-5e	Root Capabilities Root Port Only	Read Only; Bit[16] – Write 1 to clear. <ul style="list-style-type: none"> • Bits[15:1] = 0. Reserved • Bit[0] = 1. CRS Software Visibility supported.
63-60	Root Status Root Port Only	Read Only; Bit[16] – Write 1 to clear. <ul style="list-style-type: none"> • Bits[31:18] = 0. Reserved • Bit[17] – PME Pending • Bit[16] – PME Status • Bits[15:0] – PME Requester ID
67-64	Device Capabilities 2	Read Only <ul style="list-style-type: none"> • Bits[31:24] = 0. Reserved • Bits[23:22] = 00. Max End-End TLP Prefixes • Bit[21] = 0. End-End TLP Prefix Supported • Bit[20] = 0. Extended Fmt Field Supported • Bit[19:18] = 00. OBFF Supported • Bits[17:14] = 0000. Reserved • Bits[13:12] = 00. TPH Completer Supported • Bit[11] = LTR Mechanism Supported

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> • Bit[10] = 0. No RO-enabled PR-PR Passing • Bit[9] = 0. 128-bit CAS Completer Supported • Bit[8] = 0. 64-bit AtomicOp Completer Supported • Bit[7] = 0. 32-bit AtomicOp Completer Supported • Bit[6] = 0. AtomicOp Routing Supported • Bit[5] = 0. ARI Forwarding Supported • Bit[4] – Completion Timeout Disable Supported • Bits[3:0] – Completion Timeout Ranges Supported
69-68	Device Control 2	<p>Read/Write</p> <ul style="list-style-type: none"> • Bit[15] – End-End TLP Prefix Blocking • Bits[14:13] – OBFF Enable; not supported • Bits[12:11] = 00. Reserved. • Bit[10] – LTR Mechanism Enable • Bit[9] – IDO Completion Enable • Bit[8] – IDO Request Enable • Bit[7] – AtomicOp Egress Blocking • Bit[6] – AtomicOp Request Enable • Bit[5] – ARI Forwarding Enable • Bit[4] – Completion Timeout Disable – Set by system software to disable this device from generating completion timeouts. You must disable completion timeout error generation when this bit is set. • Bits[3:0] – Completion Timeout Value – Set by system software to select the completion timeout range which must be used by users which are implementing completion timeouts. See PCI Express Specification Table 7.24 for details.
6B-6A	Device Status 2	Reserved by PCI SIG for future use. Reads return 0x00000000.
6F-6C	Link Capabilities 2	<p>Read Only</p> <ul style="list-style-type: none"> • Bits[31:23] – Reserved • Bits[22:16] – Lower SKP OS Reception Supported Speeds Vector • Bits[15:9] – Lower SKP OS Generation Supported Speeds Vector • Bit[8] – Crosslink Supported • Bits[7:1] – Supported Link Speeds Vector • Bit[0] = 0. Reserved
71-70	LinkControl 2	<p>Read/Write</p> <ul style="list-style-type: none"> • Bit[15:12] – Compliance Preset/De-emphasis[3:0] • Bit[11] – Compliance SOS • Bit[10] – Enter Modified Compliance • Bits[9:7] – Transmit Margin • Bit[6] – Selectable De-emphasis • Bit[5] – Hardware Autonomous Speed Disable • Bit[4] – Enter Compliance • Bits[3:0] – Target Link Speed[3:0] <ul style="list-style-type: none"> • 0001 (2.5 GT/s) • 0010 (5.0 GT/s) • 0011 (8.0 GT/s)

Addr	Config Register	Register Description
73-72	Link Status 2	Read Only; Bit[5] – write 1 to clear: <ul style="list-style-type: none"> Bits[15:6] = 0000000000. Reserved. Bit[5] – Link Equalization Reset Bit[4] – Equalization Phase 3 Successful Bit[3] – Equalization Phase 2 Successful Bit[2] – Equalization Phase 1 Successful Bit[1] – Equalization Complete Bit[0] – Current De-emphasis Level <ul style="list-style-type: none"> 1== -3.5 dB 0== -6 dB
77-74	Slot Capabilities 2 Root Port/Switch Only	Reserved by PCI SIG for future use. Reads return 0x00000000.
79-78	Slot Control 2 Root Port/Switch Only	Reserved by PCI SIG for future use. Reads return 0x00000000.
7b-7a	Slot Status 2 Root Port/Switch Only	Reserved by PCI SIG for future use. Reads return 0x00000000.
7F-7C	Reserved	Reads return 0x00000000.

5.2.6. Power Management Capability

Table 5.213. Power Management Capability

Addr	Config Register	Register Description
80	Power Management Capability ID	Read Only = 0x01 (Beginning of Power Management Capability Item)
81	Next Capability Pointer	Read Only. Pointer to next Capability Item in the list.
83-82	Power Management Capabilities	Read Only. <ul style="list-style-type: none"> Bits[15:11] – PME Support; recommended default == 0. Bits[10] – D2 Support (1) Yes (0) No; this bit must be set for the core to allow Power State to be written to D2; recommended default == 0. Bit[9] – D1 Support (1) Yes (0) No; this bit must be set for the core to allow Power State to be written to D1; recommended default == 0. Bit[8:6] – Aux Current; recommended default = 0. Bit[5] – Device Specific Initialization(DSI); recommended default = 0. Bit[4] – Reserved; set to 0. Bit[3] – PME Clock; recommended default = 0. Bits[2:0] – Version; set to 011 (complies with revision 1.2 of the PCI Power Management Interface Specification). Refer Error Handling for additional detail.

Addr	Config Register	Register Description
85-84	Power Management Control/Status	<p>Read/Write.</p> <ul style="list-style-type: none"> Bit[15] – PME Status; if Power Management Capabilities[15] == 1 indicating that PME is generated from D3cold, then PME_Status is implemented by the core; otherwise PME_Status == 0. Bits[14:13] – Data Scale; recommend == 0 (Data not implemented) Bits[12:9] – Data Select; recommend == 0 (Data not implemented) Bit[8] – PME En – ; if Power Management Capabilities[15:11] == 0 indicating that PME is not generated from any power state then PME_En == 0; is implemented by the core and written by system software to enable PME generation from D3cold; otherwise PME_En == 0. Bits[7:4] – Reserved – set to 0 Bit[3] – No Soft Reset – Core sets to 1 since the core is not reset when transitioning from D3hot to D0 purely due to power state changes. This bit is used by system software to know whether the device needs to be reinitialized when transitioning between D3hot and D0. Bit[2] – Reserved; set to 0 Bits[1:0] – Power State; software writes this field to transition a device into a different power state; increasing Dx numbers represent increasingly lower power states. <ul style="list-style-type: none"> 00 – D0; normal operation 01 – D1; not allowed to be written unless D1 Support == 1 10 – D2; not allowed to be written unless D2 Support == 1 11 – D3hot; “off” <p>Refer Error Handling for additional detail.</p>
86	PMCSR PCI to PCI Bridge Support	<p>Read Only.</p> <ul style="list-style-type: none"> Bit[7] – Bus Power/Clock Control Enable; set to 0 Bit[6] – B2/B3 Support for D3bat; set to 0 Bits[5:0] – Reserved; set to 0
87	Data	<p>Read Only.</p> <ul style="list-style-type: none"> Bits[7:0] – Data; recommended default = 0; not implemented
8F-88	Reserved	Reads return 0x00000000.

5.2.7. MSI-X Capability

Table 5.214. MSI-X Capability

Addr	Config Register	Register Description
90	MSI-X Capability ID	<p>Read Only = 0x11 (Beginning of MSI-X Capability Item)</p> <p>MSI-X Support may be enabled/disabled through the CSR registers. If present, its capability is defined as follows otherwise all the following registers reads 0x0.</p>
91	Next Capability Pointer	Read Only. Pointer to next Capability Item in the list.
93-92	Message Control	<p>Only bits[15:14] are Read/Write.</p> <ul style="list-style-type: none"> Bit[15] – MSI-X Enable (Read/Write) Bit[14] – Function Mask (Read/Write) Bits[13:11] – Reserved – 000 (Read Only) Bit[10:0] – Table Size[10:0] (Read Only) <ul style="list-style-type: none"> The number of MSI-X vectors requested/supported by the user’s design is Table Size + 1.
97-94	Table_Offset, Table_BIR	<p>Read Only.</p> <p>Bits[31:3] – Table_Offset[31:3]</p> <ul style="list-style-type: none"> {Table_Offset[31:3], 000} is the offset into the BAR indicated by Table_BIR where the MSI-X Table begins. Bits[2:0] – Table BIR[2:0]

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> Indicates which BAR location contains the MSI-X Table. In the case of a 64-bit BAR Table BIR indicates the BAR that contains the lower 32-bit address: <ul style="list-style-type: none"> 000 – BAR0 001 – BAR1 010 – BAR2 011 – BAR3 100 – BAR4 101 – BAR5 110, 111 – Reserved
9B-98	PBA_Offset, PBA_BIR	<p>Read Only.</p> <p>Bits[31:3] – PBA_Offset[31:3]</p> <ul style="list-style-type: none"> Same as Table Offset above but indicates the location of the PBA (Pending Bit Array). <ul style="list-style-type: none"> Bits[2:0] – PBA BIR[2:0] Same as Table BIR above, but indicates the location of the PBA.

5.2.8. MSI Capability

Table 5.215. MSI Capability

Addr	Config Register	Register Description
9F-9C	Reserved	Reads return 0x00000000.
A0	Message Capability ID	Read Only = 0x05 (Beginning of Message Capability Item); MSI Support is enabled/disabled by CSR registers. If present, its capability is defined as follows otherwise all the following registers read 0x0.
A1	Next Capability Pointer	Read Only. Pointer to next Capability Item in the list.
A3-A2	Message Control	<ul style="list-style-type: none"> Bits[6:4] and Bit[0] are Read/Write; remainder are Read Only. Bits[15:9] = 0x00. Reserved. Bit[8] = 0. Note per vector masking capable. Bit[7] – 64-bit Address Capable = 1 (Capable of generating 64-bit messages). Bits[6:4] – Multiple Message Enable – system software writes the number of allocated messages; 000==1, 001==2, 010==4, 011==8, 100==16, 101==32, 110 Reserved, 111 Reserved. Bits[3:1] – Multiple Message Capable – Number of messages requested by the device == 000 (1 Message). Bit[0] – MSI Enable – System software sets this bit to enable MSI. When set, the core uses the MSI mechanism instead of the legacy interrupt mechanism to forward user interrupts on mgmt_interrupt to PCI Express.
A7-A4	Message Address	<p>Bits[31:2] are Read/Write; Bits[1:0] are Read Only.</p> <ul style="list-style-type: none"> Bits[31:2] Message Address[31:2] Bits[1:0] – Reserved – Message Address[1:0] is always 00.
AB-A8	Message Upper Address	<p>Read/Write</p> <ul style="list-style-type: none"> Bits[31:0] Message Address[63:32]; if Message Address[63:32] == 0, then the core uses only Message Address[31:0] and does 32-bit address MSI writes. If Message Address[63:32] != 0, the core uses Message Address[63:0] and does 64-bit address MSI writes.
AD-AC	Message Data	<p>Read/Write</p> <ul style="list-style-type: none"> Bits[15:0] Message Data[15:0] – An MSI Message is sent by writing Message Data to Message Address.

5.2.9. Advanced Error Reporting Extended Capability

Table 5.216. Advanced Error Reporting Extended Capability

Addr	Config Register	Register Description
103-100	Advanced Error Reporting Enhanced Capability Header	<p>Beginning of Advanced Error Reporting (AER) Capability; the AER capability is only present if AER support is enabled in the design, however, AER support is a standard core feature that is present unless AER removal has been specifically requested to be excluded at core deliver time (which is unusual).</p> <ul style="list-style-type: none"> • Bits[15:0] – Read Only = 0x0001 == AER Capability ID • Bits[19:16] – Read Only = 0x01 == AER Capability Version (PCIe 2.0/1.1) • Bits[31:20] – Read Only. Pointer to next Enhanced/Extended Capability Item in the list.
107-104	Uncorrectable Error Status	<ul style="list-style-type: none"> • Read/Write: Bit set when corresponding error event occurs, and the error is not masked by the Uncorrectable Error Mask register; clear set bits by writing a 1: • Bits[3:0] – Reserved == 0 • Bit[4] – DataLink_Protocol_Error_Status • Bit[5] – Surprise_Down_Error_Status • Bits[11:6] – Reserved == 0 • Bit[12] – Poisoned_TLP_Status • Bit[13] – Flow_Control_Protocol_Error_Status • Bit[14] – Completion_Timeout_Status • Bit[15] – Completer_Abort_Status • Bit[16] – Unexpected_Completion_Status • Bit[17] – Receiver_Overflow_Status • Bit[18] – Malformed_TLP_Status • Bit[19] – ECRC_Error_Status • Bit[20] – Unsupported_Request_Error_Status • Bit[21] – Reserved = 0 • Bit[22] – Uncorrectable Internal Error Status • Bits[31:23] – Reserved == 0
10B-108	Uncorrectable Error Mask	<p>Read/Write: Set corresponding bit to mask (not report) selected error events; clear to unmask (report):</p> <ul style="list-style-type: none"> • Bits[3:0] – Reserved == 0 • Bit[4] – DataLink_Protocol_Error_Mask • Bit[5] – Surprise_Down_Error_Mask • Bits[11:6] – Reserved == 0 • Bit[12] – Poisoned_TLP_Mask • Bit[13] – Flow_Control_Protocol_Error_Mask • Bit[14] – Completion_Timeout_Mask • Bit[15] – Completer_Abort_Mask • Bit[16] – Unexpected_Completion_Mask • Bit[17] – Receiver_Overflow_Mask • Bit[18] – Malformed_TLP_Mask • Bit[19] – ECRC_Error_Mask • Bit[20] – Unsupported_Request_Error_Mask • Bit[21] – Reserved = 0 • Bit[22] – Uncorrectable Internal Error Mask • Bits[31:23] – Reserved == 0
10F-10C	Uncorrectable Error Severity	<p>Read/Write: Set corresponding bit to mark selected error events as FATAL errors; clear to mark selected error events as NON-FATAL errors:</p> <ul style="list-style-type: none"> • Bits[3:0] – Reserved == 0 • Bit[4] – DataLink_Protocol_Error_Severity • Bit[5] – Surprise_Down_Error_Severity • Bits[11:6] – Reserved == 0

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> • Bit[12] – Poisoned_TLP_Severity • Bit[13] – Flow_Control_Protocol_Error_Severity • Bit[14] – Completion_Timeout_Severity • Bit[15] – Completer_Abort_Severity • Bit[16] – Unexpected_Completion_Severity • Bit[17] – Receiver_Overflow_Severity • Bit[18] – Malformed_TLP_Severity • Bit[19] – ECRC_Error_Severity • Bit[20] – Unsupported_Request_Error_Severity • Bit[21] – Reserved = 0 • Bit[22] – Uncorrectable Internal Error Severity • Bits[31:23] – Reserved == 0
113-110	Correctable Error Status	<p>Read/Write: Bit set when corresponding error event occurs, and the error is not masked by the Correctable Error Mask register; clear set bits by writing a 1:</p> <ul style="list-style-type: none"> • Bit[0] – Receiver_Error_Status • Bits[5:1] – Reserved == 0 • Bit[6] – Bad_TLP_Status • Bit[7] – Bad_DLLP_Status • Bit[8] – Replay_Num_Rollover_Status • Bits[11:9] – Reserved == 000 • Bit[12] – Replay_Timer_Timeout_Status • Bit[13] – Advisory_Non_Fatal_Error_Status • Bit[14] – Corrected Internal Error Status • Bit[15] – Header Log Overflow Status • Bits[31:16] – Reserved == 0
117-114	Correctable Error Mask	<p>Read/Write: Set corresponding bit to mask (not report) selected error events; clear to unmask (report):</p> <ul style="list-style-type: none"> • Bit[0] – Receiver_Error_Mask • Bits[5:1] – Reserved == 0 • Bit[6] – Bad_TLP_Mask • Bit[7] – Bad_DLLP_Mask • Bit[8] – Replay_Num_Rollover_Mask • Bits[11:9] – Reserved == 000 • Bit[12] = Replay_Timer_Timeout_Mask • Bit[13] = Advisory_Non_Fatal_Error_Mask • Bit[14] – Corrected Internal Error Mask • Bit[15] – Header Log Overflow Mask • Bits[31:16] – Reserved == 0
11B-118	Advanced Error Capabilities and Control	<p>Read/Write: Misc Capabilities and Control</p> <ul style="list-style-type: none"> • Bits[4:0] = Read Only – First_Error_Pointer[4:0] • Bit[5] = Read Only – ECRC_Generation_Capable <ul style="list-style-type: none"> • 1==Device can generate ECRC; set if the core includes ECRC generation logic (non-standard core option). • 0 == Device is not capable of generating ECRC. • Bit[6] = Read/Write – ECRC_Generation_Enable <ul style="list-style-type: none"> • Software sets to control whether ECRCs are generated and inserted for TLPs transmitted by the core; if ECRC support is not implemented in the core, this bit is Read Only == 0. • 1== Generate and insert ECRC for TLPs transmitted by the core. • 0 == Do not generate and insert ECRC.

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> Bit[7] = Read Only – ECRC_Check_Capable <ul style="list-style-type: none"> 1==Device is capable of checking ECRC; set if the core includes ECRC generation logic (non-standard core option). 0 == Device is not capable of checking ECRC. Bit[8] = Read/Write – ECRC_Check_Enable <ul style="list-style-type: none"> Software sets to control whether ECRCs are checked for TLPs received by the core; if ECRC support is not implemented in the core, this bit is Read Only == 0. 1== Check ECRC for all TLPs with ECRC received by the core. 0 == Do not check ECRC. Bits[31:9] – Reserved = 0
12B-11C	Header Log	Header[127:0] of the TLP associated with the error. TLP format is in same order as illustrated in PCIe Specification: <ul style="list-style-type: none"> 0x11F-11C – {Byte0, Byte1, Byte2, Byte3} 0x123-120 – {Byte4, Byte5, Byte6, Byte7} 0x127-124 – {Byte8, Byte9, Byte10, Byte11} 0x12B-0x128 – {Byte12, Byte13, Byte14, Byte15}
137-12C	Reserved	Only implemented by AER Root Ports. Reads return 0x00000000.
147-138	Reserved	TLP Prefix Log Register

5.2.10. ARI Extended Capability

ARI is located at offset 0x148 unless AER is not present in which case it is moved to 0x100.

Table 5.217. ARI Extended Capability

Addr	Config Register	Register Description
14B-148 or 103-100	ARI Capability Extended Capability Header	Beginning of ARI Extended Capability – Read Only <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x000E == Capability ID
14D-14C or 105-104	ARI Capability Register	Read Only <ul style="list-style-type: none"> Bits[15:8] – Next Function Number = 0 (not implemented) Bits[7:2] – Reserved = 0 Bit[1] – ACS Function Groups Capability = 0 (not implemented) Bit[0] – MFVC Function Groups Capability = 0 (not implemented)
14F-14E or 107-106	ARI Control Register	Read Only <ul style="list-style-type: none"> Bits[15:7] – Reserved Bit[6:4] – Function Group = 0 (not implemented) Bits[3:2] – Reserved = 0 Bit[1] – ACS Function Groups Enable = 0 (not implemented) Bit[0] – MFVC Function Groups Enable = 0 (not implemented)

5.2.11. Vendor-Specific Extended Capability

Table 5.218. Vendor-Specific Extended Capability

Addr	Config Register	Register Description
153-150	Vendor-Specific PCI Express Extended Capability Header	Beginning of Vendor-Specific Extended Capability (VSEC) <ul style="list-style-type: none"> • Bits[31:20] – Read Only. Pointer to next Enhanced/Extended Capability Item in the list. • Bits[19:16] – Read Only = 0x1 == Capability Version • Bits[15:0] – Read Only = 0x000B == Capability ID
157-154	Vendor-Specific Header	Read Only <ul style="list-style-type: none"> • Bits[31:20] – VSEC Length = 0x24 (36 bytes) • Bits[19:16] – VSEC Rev = 0x1 • Bits[15:0] – VSEC ID
15B-158	HW.Init	Read/Write <ul style="list-style-type: none"> • Bits[31:1] = 0. Reserved • Bit[0] – HW.Init Write Enable – Used to allow software to write some Configuration Registers which are type <i>HW.Init</i> when they would otherwise not be writable; default value == 0 <ul style="list-style-type: none"> • 1 – HW.Init Write Enabled – Allow specific HW.Init fields to be written by software; only relevant for Configuration Registers in this document which specifically state they are writable when HW.Init Write Enable == 1 (for example, PCI Express Capability: Slot Capabilities). • 0 – HW.Init Write Disabled
15F-15C	Link Power Down Root Port / Downstream Switch Port Only	Read/Write – Used by system software in a Root Port or Downstream Switch Port application to cause a PME_Turn_Off Message to be transmitted on PCI Express to request that the downstream PCI Express hierarchy prepare for Power Down. <ul style="list-style-type: none"> • Bits[31:3] = 0. Reserved. • Bit[2] – L2 Request Timeout; indicates when an L2 Request completed due to a timeout; L2 Request Timeout is cleared (0) when L2_Request is written to 1 or when 1 is written to this register; L2 Request Timeout is set (1) when a PME_TO_ACK message is not received in response to a transmitted PME_Turn_Off within the expected 100 ms (100 μs for simulation when mgmt_short_sim == 1) timeout window. • Bit[1] – L2 Request Status; indicates when an L2 Request has completed either due to receiving the expected PME_TO_Ack message response or due to timeout; L2 Request Status is cleared (0) when L2_Request is written to 1 or when 1 is written to this register; L2 Request Status is set (1) when a PME_TO_ACK message is received or a timeout occurs • Bit[0] – L2 Request; write to 1 to cause a PME_Turn_Off Message to be transmitted downstream to the PCI Express hierarchy; after all downstream devices have prepared for power-down the core should receive a PME_TO_Ack message in response indicating the downstream PCIe hierarchy is ready for removal of power; L2 Request stays set until a PME_TO_ACK message is received or a timeout occurs.

Addr	Config Register	Register Description
163-160	Autonomous Recovery, Speed, and Width	<p>Read/Write – Used by system software in an US Port to perform autonomous speed change, width change, or entry to recovery.</p> <p>Bits[31:16] – Lane Width Mask. A 1 indicates that the lane can be used. [16] = Lane 0 .. [31] = Lane 15.</p> <p>Bits[15:12] – Reserved</p> <p>Bits[11:8] – Target Speed. (1=2.5G, 2=5G, 3=8G, 4=16G).</p> <p>Bits[7:3] – Reserved</p> <p>Bit[2] – Autonomous Entry to Recovery command. When Set to 1, Bits[1] and [0] must both be set to 0. Setting this bit to 1 causes the Link to immediately transition to recovery.</p> <p>Bit[1] – Autonomous Width Change command. When set to 1, the Link transitions to recovery to perform a link width change, using the Lane Width Mask field. This bit is ignored if HW Autonomous Width Disable has been set in the Link Control register.</p> <p>Bit[0] – Autonomous Speed Change command. When set to 1, the Link transitions to Recovery to perform a speed change, using the Target Speed field. This bit is ignored if HW Autonomous Speed Disable has been set in the Link Control 2 register.</p> <p>Speed and width changes can be signaled together. However, entry to recovery must be signaled independently from speed or width changes.</p>
173-164	Reserved	Reserved

5.2.12. Secondary PCI Express Extended Capability

Table 5.219. Secondary PCI Express Extended Capability

Addr	Config Register	Register Description
183-180	Secondary PCI Express Extended Capability Header	<p>Beginning of Secondary PCI Express Extended Capability; this capability is only present if the PCIe 3.0 support is enabled in the design. If the AER capability is not present, this capability is located at offset 0x100 instead.</p> <p>Bits[31:20] – Read Only. Pointer to next Enhanced/Extended Capability Item in the list.</p> <ul style="list-style-type: none"> • Bits[19:16] – Read Only = 0x1 == Capability Version • Bits[15:0] – Read Only = 0x0019 == Capability ID
187-184	Link Control 3	<p>Read/Write</p> <ul style="list-style-type: none"> • Bits[31:16] – Reserved = 0 • Bits[15:9] – Enable Lower SKP OS Generation Vector • Bits[8:2] – Reserved = 0 • Bit[1] – Link Equalization Request Interrupt Enable • Bit[0] – Perform Equalization
18B-188	Lane Error Status	<p>Read Only: Indicates lane-specific error status.</p> <ul style="list-style-type: none"> • Bits[31:NUM_LANES] – Reserved = 0 • Bit[Lane#] – 1 == Error detected on lane[[Lane#]]; 0 == no error
1AB-18C	Lane Equalization Control Register	<p>Read Only: Control and status fields for link equalization; 16-bits per lane starting with Lane[0] with higher lane #s at higher addresses.</p> <p>Per lane format is as follows:</p> <ul style="list-style-type: none"> • Bit[15] – Reserved = 0 • Bits[14:12] – Upstream Port Receiver Preset Hint • Bits[11:8] – Upstream Port Transmitter Preset • Bit[7] – Reserved = 0 • Bits[6:4] – Downstream Port Receiver Preset Hint • Bits[3:0] – Downstream Port Transmitter Preset

5.2.13. ATS Extended Capability

Table 5.220. ATS Extended Capability

Addr	Config Register	Register Description
203-200	ATS Capability Extended Capability Header	Beginning of ATS Extended Capability – Read Only <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x000F == Capability ID
205-204	ATS Capability Register	Read Only <ul style="list-style-type: none"> Bits[4:0] – Invalidate Queue Depth Bit[5] Page Aligned Request Bits[15:6] – Reserved
207-206	ATS Control Register	Read/Write <ul style="list-style-type: none"> Bits[4:0] – Smallest Translation Unit Bits[14:5] – Reserved Bit[15] – Enable

5.2.14. DSN Extended Capability

Table 5.221. DSN Extended Capability

Addr	Config Register	Register Description
213-210	DSN Capability Extended Capability Header	Beginning of DSN Extended Capability – Read Only <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x0003 == Capability ID
21B-214	DSN Serial Number	Read Only <ul style="list-style-type: none"> Bits[63:0] – DSN Serial Number

5.2.15. Resizable BAR Capability

Table 5.222. Resizable BAR Capability

Addr	Config Register	Register Description
283-280	Resizable BAR Extended Capability Header	Resizable BAR Capability Header – Read Only <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x0015 == Capability ID
287-284	Resizable BAR Capability(0)	Read Only <ul style="list-style-type: none"> Bits[31:24] – Reserved Bits[23:4] – When Bit n is Set, The BAR Indicated by the BAR Index in the Control Register operates with BAR sized to $2^{(n+16)}$ Bytes. For example, bit[4] = 2^{20} Bytes = 1 MB Bits[3:0] – Reserved

Addr	Config Register	Register Description
28B-288	Resizable Bar Control Register(0)	<ul style="list-style-type: none"> Read Only <ul style="list-style-type: none"> Bits[31:13] – Reserved Read/Write <ul style="list-style-type: none"> Bits[12:8] – BAR Size. Encoded Value for the Size this BAR should use. Read Only <ul style="list-style-type: none"> Bits[7:5] – Number of Resizable BARs. Value must be between 1 and 6. These bits are only valid in the Resizable BAR Control Register (0). In Control Registers (1) or higher, these bits are Reserved. Bits[2:0] – BAR Index for this BAR: <ul style="list-style-type: none"> 0 = BAR located at offset 0x10 1 = BAR located at offset 0x14 2 = BAR located at offset 0x18 3 = BAR located at offset 0x1C 4 = BAR located at offset 0x20 5 = BAR located at offset 0x24 Other values reserved. For a 64-bit BAR, this index should point to the lower DWORD.
28F-28C	Resizable BAR Capability and Control Registers (1..6)	<p>See Resizable BAR Capability (0)</p> <p>See Resizable BAR Control Register(0)</p> <p>The number of Implemented BAR Capability and Control Registers depends on the setting of “Number of Resizable BARs” Control Register (0).</p>

5.2.16. Power Budgeting Capability

Table 5.223. Power Budgeting Capability

Addr	Config Register	Register Description
393-390	Power Budgeting Capability Extended Capability Header	<p>Beginning of Power Budgeting Extended Capability – Read Only</p> <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x0004 == Capability ID
394	Data Select Register	<p>Read/Write</p> <ul style="list-style-type: none"> Bits[7:0] – Data Select Register
397-395	Reserved	Reserved
39B-398	Data Register	<p>Read Only</p> <ul style="list-style-type: none"> Bits[31:21] – Reserved Bits[20:18] – Power Rail (0:12V,1:3.3V,2:1.5/1.8V,7:Thermal) Bits[17:15] – Type (0:PME Aux,1:Aux,2:Idle,3:Sustained,7:Max) Bits[14:13] – PM State (0:D0,1:D1,2:D2,3:D3) Bits[12:10] – PM Sub State (0:Default,others: Device Specific) Bits[9:8] – Data Scale (0:1x,1:0.1x,2:0.01x,3:0.001x) Bits[7:0] – Base Power
39C	Capabilities Register	<p>Read Only</p> <ul style="list-style-type: none"> Bits[7:1] – Reserved Bit[0] – System Allocated – Set to 1 to indicate that the Power Budget Should be System Allocated, and the values from the Data Register should NOT be used for System Power Budgeting. Set to 0 to indicate that the values provided in the Data Register should be used for System Power Budgeting.

5.2.17. Dynamic Power Allocation Capability

Table 5.224. Dynamic Power Allocation (DPA) Capability

Addr	Config Register	Register Description
3A3-3A0	DPA Capability Extended Capability Header	Beginning of DPA Extended Capability – Read Only <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x0016 == Capability ID
3A7-3A4	DPA Capability Register	Read Only <ul style="list-style-type: none"> Bits[31:24] – Transition Latency Value1 (xlcy1) Bits[23:16] – Transition Latency Value0 (xlcy0) Bits[15:14] – Reserved Bits[13:12] – Power Allocation Scale (PAS) Bits[11:10] – Reserved Bits[9:8] – Transition Latency Unit (tlunit) Bits[7:5] –Reserved Bits[4:0] – Substate_Max
3AB-3A8	DPA Latency Indicator Register	Read Only <ul style="list-style-type: none"> Bits[31:Substate_Max+1] – Reserved Bits[Substate_Max:0] – Transition Latency Indicator Bits
3AD-3AC	DPA Status Register	Read Only <ul style="list-style-type: none"> Bits[15:9] – Reserved Read, Write 1 to Clear <ul style="list-style-type: none"> Bits[8] – Substate Control Enabled Read Only <ul style="list-style-type: none"> Bits[7:0] – Substate Status
3EF-3AE	DPA Control Register	Read Only <ul style="list-style-type: none"> Bits[15:5] – Reserved Read/Write <ul style="list-style-type: none"> Bits[4:0] – Substate Control
3CF-3B0	DPA Power Allocation Array	Read Only <ul style="list-style-type: none"> Bits[7:0] – Substate Power Allocation Register Address 3B0 is for Substate 0 Address 3B1 is for Substate 1, up to Substate Substate_Max

5.2.18. L1 PM Substates Extended Capability

Table 5.225. L1 PM Substates Extended Capability

Addr	Config Register	Register Description
3D3-3D0	L1 PM Substates Capability Extended Capability Header	Beginning of L1 PM Substates Extended Capability – Read Only <ul style="list-style-type: none"> Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. Bits[19:16] = 0x1 == Capability Version Bits[15:0] = 0x001E == Capability ID
3D7-3D4	L1 PM Substates Capabilities Register	HwInit <ul style="list-style-type: none"> Bits[31:24] – Reserved Bits[23:19] – Port TPOWER_ON Value Bits [18] – Reserved Bits[17:16] – Port TPOWER_ON Scale Bits[15:8] – Port Common_Mode_Restore_Time (in μs) Bits[7:5] – Reserved Bit[4] – L1 PM Substates Supported Bit[3] – ASPM L1.1 Supported

Addr	Config Register	Register Description
		<ul style="list-style-type: none"> • Bit[2] – ASPM L1.2 Supported • Bit[1] – PCI-PM L1.1 Supported • Bit[0] – PCI-PM L1.2 Supported
3DB-3D8	L1 PM Substates Control 1 Register	Read/Write <ul style="list-style-type: none"> • Bits[31:29] – LTR_L1.2_THRESHOLD_Scale • Bits[28:26] – Reserved • Bits[25:16] – LTR_L1.2_THRESHOLD_Value • Bits[15:8] – Common_Mode_Restore_Time • Bits[7:4] – Reserved • Bit[3] – ASPM L1.1 Enable • Bit[2] – ASPM L1.2 Enable • Bit[1] – PCI-PM L1.1 Enable • Bit[0] – PCI-PM L1.2 Enable
3DF-3DC	L1 PM Substates Control 2 Register	Read/Write <ul style="list-style-type: none"> • Bits[31:8] – Reserved • Bits[7:3] – TPOWER_ON Value • Bit[2] – Reserved • Bits[1:0] – T_{POWER_ON} Scale

5.2.19. Latency Tolerance Reporting Capability

Table 5.226. Latency Tolerance Reporting (LTR) Capability

Addr	Config Register	Register Description
3E3-3E0	LTR Capability Extended Capability Header	Beginning of LTR Extended Capability – Read Only <ul style="list-style-type: none"> • Bits[31:20] – Pointer to next Enhanced/Extended Capability Item in the list. • Bits[19:16] = 0x1 == Capability Version • Bits[15:0] = 0x0018 == Capability ID
3E5-3E4	Max Snoop Latency Register	Read/Write <ul style="list-style-type: none"> • Bits[15:13] – Reserved • Bits[12:10] – Max Snoop LatencyScale • Bits[9:0] – Max Snoop LatencyValue
3E7-3E6	Max No-Snoop Latency Register	Read/Write <ul style="list-style-type: none"> • Bits[15:13] – Reserved • Bits[12:10] – Max No-Snoop LatencyScale • Bits[9:0] – Max No-Snoop LatencyValue

5.3. DMA Configuration Space Registers

Each DMA Channel implements 128 (0x80) bytes of DMA Registers. DMA Channel Registers are accessed via PCI Express through the PCI Express Base Address Register associated with DMA Channel Registers. When multiple PCIe functions are present, each function is assigned 2 DMA Channels and each function's share of DMA Channels is accessed through the function's associated DMA Register BAR. DMA Channel Registers within the same function are packed back-back.

To discover all 64 DMA Channels, 8 physical functions and SR-IOV need to be enabled. This will give a total of 32 functions available to the PCIe Endpoint.

Table 5.227. DMA Configuration Space Registers

Register Offset	Config Register	Register Description
0x0	SRC_Q_PTR_LO	<ul style="list-style-type: none"> Bits[31:6] – Queue Base Address [31:6]. Queues are required to be 64-byte aligned. Bits[5:2] – Queue Read Request Attributes. Transaction attributes used by Queue read requests when reading SGL elements. If Queue Location == AXI, read_attr[3:0] is used for m_arsache[3:0]. If Queue Location == PCIe, read_attr[2:0] is used for PCIe Attr[2:0]. Bits[1] – Queue Enable <ul style="list-style-type: none"> 0 – Disabled. The DMA Channel will not read from the queue. 1 – Enabled. The DMA Channel will fetch queue elements as needed from queue as long as Q_LIMIT != Q_NEXT Bits[0] – Queue Location <ul style="list-style-type: none"> 0 – Queue is located in PCIe Memory 1 – Queue is located in AXI Memory
0x4	SRC_Q_PTR_HI	<ul style="list-style-type: none"> Bits[31:0] – Queue Base Address [63:32]. Must be set to 0x0 if the Queue is located in 32-bit address space.
0x8	SRC_Q_SIZE	<ul style="list-style-type: none"> Bits[31:0] – Queue Size. Number of Elements in the Queue. queue_size must be >= 2. A minimum of 2 elements is required to support software/hardware queue flow control ownership. queue_size is used to identify the wrap boundary of the Queue.
0xC	SRC_Q_LIMIT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Limit Pointer. Index of the first Queue element still “owned” by software. Incremented by software to give the DMA Channel additional elements to execute. DMA Channel hardware will pause and not utilize queue elements when Q_LIMIT is reached until Q_LIMIT is advanced to provide additional elements to execute.
0x10	DST_Q_PTR_LO	<ul style="list-style-type: none"> Bits[31:6] – Queue Base Address [31:6]. Queues are required to be 64-byte aligned. Bits[5:2] – Queue Read Request Attributes. Transaction attributes used by Queue read requests when reading SGL elements. If Queue Location == AXI, read_attr[3:0] is used for m_arsache[3:0]. If Queue Location == PCIe, read_attr[2:0] is used for PCIe Attr[2:0]. Bits[1] – Queue Enable <ul style="list-style-type: none"> 0 – Disabled. The DMA Channel will not read from the queue. 1 – Enabled. The DMA Channel will fetch queue elements as needed from queue as long as Q_LIMIT != Q_NEXT Bits[0] – Queue Location <ul style="list-style-type: none"> 0 – Queue is located in PCIe Memory 1 – Queue is located in AXI Memory
0x14	DST_Q_PTR_HI	<ul style="list-style-type: none"> Bits[31:0] – Queue Base Address [63:32]. Must be set to 0x0 if the Queue is located in 32-bit address space.
0x18	DST_Q_SIZE	<ul style="list-style-type: none"> Bits[31:0] – Queue Size. Number of Elements in the Queue. queue_size must be >= 2. A minimum of 2 elements is required to support software/hardware queue flow control ownership. queue_size is used to identify the wrap boundary of the Queue.
0x1C	DST_Q_LIMIT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Limit Pointer. Index of the first Queue element still “owned” by software. Incremented by software to give the DMA Channel additional elements to execute. DMA Channel hardware will pause and not utilize queue elements when Q_LIMIT is reached until Q_LIMIT is advanced to provide additional elements to execute.
0x20	STAS_Q_PTR_LO	<ul style="list-style-type: none"> Bits[31:6] – Queue Base Address [31:6]. Queues are required to be 64-byte aligned. Bits[5:2] – Queue Read Request Attributes. Transaction attributes used by Queue read requests when reading SGL elements. If Queue Location == AXI, read_attr[3:0] is used for m_arsache[3:0]. If Queue Location == PCIe, read_attr[2:0] is used for PCIe Attr[2:0]. Bits[1] – Queue Enable <ul style="list-style-type: none"> 0 – Disabled. The DMA Channel will not read from the queue. 1 – Enabled. The DMA Channel will fetch queue elements as needed from queue as long as Q_LIMIT != Q_NEXT Bits[0] – Queue Location

Register Offset	Config Register	Register Description
		<ul style="list-style-type: none"> 0 – Queue is located in PCIe Memory 1 – Queue is located in AXI Memory
0x24	STAS_Q_PTR_HI	<ul style="list-style-type: none"> Bits[31:0] – Queue Base Address [63:32]. Must be set to 0x0 if the Queue is located in 32-bit address space.
0x28	STAS_Q_SIZE	<ul style="list-style-type: none"> Bits[31:0] – Queue Size. Number of Elements in the Queue. queue_size must be ≥ 2. A minimum of 2 elements is required to support software/hardware queue flow control ownership. queue_size is used to identify the wrap boundary of the Queue.
0x2C	STAS_Q_LIMIT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Limit Pointer. Index of the first Queue element still “owned” by software. Incremented by software to give the DMA Channel additional elements to execute. DMA Channel hardware will pause and not utilize queue elements when Q_LIMIT is reached until Q_LIMIT is advanced to provide additional elements to execute.
0x30	STAD_Q_PTR_LO	<ul style="list-style-type: none"> Bits[31:6] – Queue Base Address [31:6]. Queues are required to be 64-byte aligned. Bits[5:2] – Queue Read Request Attributes. Transaction attributes used by Queue read requests when reading SGL elements. If Queue Location == AXI, read_attr[3:0] is used for m_arcache[3:0]. If Queue Location == PCIe, read_attr[2:0] is used for PCIe Attr[2:0]. Bits[1] – Queue Enable <ul style="list-style-type: none"> 0 – Disabled. The DMA Channel will not read from the queue. 1 – Enabled. The DMA Channel will fetch queue elements as needed from queue as long as Q_LIMIT \neq Q_NEXT Bits[0] – Queue Location <ul style="list-style-type: none"> 0 – Queue is located in PCIe Memory 1 – Queue is located in AXI Memory
0x34	STAD_Q_PTR_HI	<ul style="list-style-type: none"> Bits[31:0] – Queue Base Address [63:32]. Must be set to 0x0 if the Queue is located in 32-bit address space.
0x38	STAD_Q_SIZE	<ul style="list-style-type: none"> Bits[31:0] – Queue Size. Number of Elements in the Queue. queue_size must be ≥ 2. A minimum of 2 elements is required to support software/hardware queue flow control ownership. queue_size is used to identify the wrap boundary of the Queue.
0x3C	STAD_Q_LIMIT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Limit Pointer. Index of the first Queue element still “owned” by software. Incremented by software to give the DMA Channel additional elements to execute. DMA Channel hardware will pause and not utilize queue elements when Q_LIMIT is reached until Q_LIMIT is advanced to provide additional elements to execute.
0x40	SRC_Q_NEXT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Next Pointer. Index of the next Queue element that will be read by DMA Channel hardware. Incremented by DMA Channel hardware as queue read requests are generated. The number of queue elements available for the DMA Channel == SRC_Q_LIMIT - SRC_Q_NEXT (taking into account wrapping). SRC_Q_NEXT does not indicate that Queue Elements have been completed, only that the DMA Channel has started processing for the queue elements. Software must write this register to 0x0 to initialize the queue prior to enabling the DMA Channel. Software is prohibited from writing this register while the DMA Channel is enabled. SRC_Q_NEXT is utilized by DMA Channel hardware to track its location in the queue and should not be used by DMA software.
0x44	DST_Q_NEXT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Next Pointer. Index of the next Queue element that will be read by DMA Channel hardware. Incremented by DMA Channel hardware as queue read requests are generated. The number of queue elements available for the DMA Channel == SRC_Q_LIMIT - SRC_Q_NEXT (taking into account wrapping). SRC_Q_NEXT does not indicate that Queue Elements have been completed, only that the DMA Channel has started processing for the queue elements. Software must write this register to 0x0 to initialize the queue prior to enabling the DMA Channel. Software is prohibited from writing this register while the DMA Channel is enabled. SRC_Q_NEXT is utilized by DMA Channel hardware to track its location in the queue and should not be used by DMA software.

Register Offset	Config Register	Register Description
0x48	STAS_Q_NEXT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Next Pointer. Index of the next Queue element that will be read by DMA Channel hardware. Incremented by DMA Channel hardware as queue read requests are generated. The number of queue elements available for the DMA Channel == SRC_Q_LIMIT - SRC_Q_NEXT (taking into account wrapping). SRC_Q_NEXT does not indicate that Queue Elements have been completed, only that the DMA Channel has started processing for the queue elements. Software must write this register to 0x0 to initialize the queue prior to enabling the DMA Channel. Software is prohibited from writing this register while the DMA Channel is enabled. SRC_Q_NEXT is utilized by DMA Channel hardware to track its location in the queue and should not be used by DMA software.
0x4C	STAD_Q_NEXT	<ul style="list-style-type: none"> Bits[31:0] – Queue Flow Control - Next Pointer. Index of the next Queue element that will be read by DMA Channel hardware. Incremented by DMA Channel hardware as queue read requests are generated. The number of queue elements available for the DMA Channel == SRC_Q_LIMIT - SRC_Q_NEXT (taking into account wrapping). SRC_Q_NEXT does not indicate that Queue Elements have been completed, only that the DMA Channel has started processing for the queue elements. Software must write this register to 0x0 to initialize the queue prior to enabling the DMA Channel. Software is prohibited from writing this register while the DMA Channel is enabled. SRC_Q_NEXT is utilized by DMA Channel hardware to track its location in the queue and should not be used by DMA software.
0x50	SCRATCH0	<ul style="list-style-type: none"> Bits[31:0] – Scratchpad Register. Intended to enable information to be passed between software. For example, applications with both an AXI CPU and an PCIe CPU may use this register to pass information between CPUs. The DMA Channel implementation does not use or alter this information.
0x54	SCRATCH1	<ul style="list-style-type: none"> Bits[31:0] – Scratchpad Register. Intended to enable information to be passed between software. For example, applications with both an AXI CPU and an PCIe CPU may use this register to pass information between CPUs. The DMA Channel implementation does not use or alter this information.
0x58	SCRATCH2	<ul style="list-style-type: none"> Bits[31:0] – Scratchpad Register. Intended to enable information to be passed between software. For example, applications with both an AXI CPU and an PCIe CPU may use this register to pass information between CPUs. The DMA Channel implementation does not use or alter this information.
0x5C	SCRATCH3	<ul style="list-style-type: none"> Bits[31:0] – Scratchpad Register. Intended to enable information to be passed between software. For example, applications with both an AXI CPU and an PCIe CPU may use this register to pass information between CPUs. The DMA Channel implementation does not use or alter this information.
0x60	PCIE_INTERRUPT_CONTROL	<ul style="list-style-type: none"> Bits[31:24] – Reserved. Bits[23:16] – PCIe DMA SGL Interrupt Coalesce Count. Controls the frequency at which a DMA SGL Interrupt Event, completion of a source SGL Element that had both the "EOP" and "Interrupt" SGL control bits set to 1, causes PCIe Interrupts. An internal DMA Channel-specific PCIe DMA SGL Coalesce Counter is maintained. For each DMA SGL Interrupt Event, if the current PCIe DMA SGL Coalesce Counter is equal to PCIe DMA SGL Interrupt Coalesce Count, then a PCIe interrupt is generated and the PCIe DMA SGL Coalesce Counter is cleared otherwise the PCIe DMA SGL Coalesce Counter is incremented. The PCIe DMA SGL Coalesce Counter is also cleared when DMA Enable == 0. This mechanism allows software to configure the DMA Channel to interrupt PCIe once every 1-256 DMA SGL Interrupt Events. When a non-0 value is programmed into PCIe DMA SGL Interrupt Coalesce Count, software must anticipate that at the end of all DMA packet transfers, interrupts may be pending in the PCIe DMA SGL Coalesce Counter that will not cause an interrupt because no more interrupts will arrive for PCIe DMA SGL Coalesce Counter to reach the PCIe DMA SGL Interrupt Coalesce Count threshold. Bits[15:3] – Reserved. Bits[2] – PCIe Enable DMA SGL Interrupt Event

Register Offset	Config Register	Register Description
		<ul style="list-style-type: none"> 0 – Disabled. DMA SGL Interrupt Events are ignored and will not cause PCI Express interrupts to be generated. 1 – Enabled. DMA SGL Interrupt Event interrupts are enabled for this DMA Channel. Bits[1] – PCIe Enable DMA Error Interrupt Event. <ul style="list-style-type: none"> 0 – Disabled. DMA Error Interrupt Events are ignored and will not cause PCI Express interrupts to be generated. 1 – Enabled. DMA Error Interrupt Event interrupts are enabled for this DMA Channel. Bits[0] – PCIe Interrupt Enable <ul style="list-style-type: none"> 0 – Disable Interrupts. Hold PCIe interrupts which occur for this DMA Channel and assert them when PCIe Interrupt Enable is set back to 1. Pending interrupts are cleared when DMA_Enable == 0. 1 – Enable Interrupts. Allow PCIe interrupts to be asserted by this DMA Channel.
0x64	PCIE_INTERRUPT_STATUS	<ul style="list-style-type: none"> Bits[31:4] – Reserved. Bits[3] – PCIe Software Interrupt Status. <ul style="list-style-type: none"> 0 – No software interrupt pending 1 – A Software Interrupt was generated Bits[2] – PCIe DMA SGL Interrupt Status. <ul style="list-style-type: none"> 0 – No DMA interrupt pending 1 – DMA SGL Interrupt Events caused an interrupt to be generated Bits[1] – PCIe DMA Error Interrupt Status. <ul style="list-style-type: none"> 0 – No DMA interrupt pending 1 – DMA Error Interrupt Events caused an interrupt to be generated Bits[0] – Reserved.
0x68	AXI_INTERRUPT_CONTROL	<ul style="list-style-type: none"> Bits[31:24] – Reserved. Bits[23:16] – AXI DMA SGL Interrupt Coalesce Count. Controls the frequency at which a DMA SGL Interrupt Event (completion of a source SGL Element that had both the "EOP" and "Interrupt" SGL control bits set to 1) causes AXI Interrupts. An internal DMA Channel-specific AXI DMA SGL Coalesce Counter is maintained. For each DMA SGL Interrupt Event, if the current AXI DMA SGL Coalesce Counter is equal to AXI DMA SGL Interrupt Coalesce Count, then a AXI interrupt is generated and the AXI DMA SGL Coalesce Counter is cleared otherwise the AXI DMA SGL Coalesce Counter is incremented. The AXI DMA SGL Coalesce Counter is also cleared when DMA_Enable == 0. This mechanism allows software to configure the DMA Channel to interrupt AXI once every 1-256 DMA SGL Interrupt Events. When a non-0 value is programmed into AXI DMA SGL Interrupt Coalesce Count, software must anticipate that at the end of all DMA packet transfers, interrupts may be pending in the AXI DMA SGL Coalesce Counter that will not cause an interrupt because no more interrupts will arrive for AXI DMA SGL Coalesce Counter to reach the AXI DMA SGL Interrupt Coalesce Count threshold. Bits[15:3] – Reserved. Bits[2] – AXI Enable DMA SGL Interrupt Event <ul style="list-style-type: none"> 0 – Disabled. DMA SGL Interrupt Events are ignored and will not cause PCI Express interrupts to be generated. 1 – Enabled. DMA SGL Interrupt Event interrupts are enabled for this DMA Channel. Bits[1] – AXI Enable DMA Error Interrupt Event. <ul style="list-style-type: none"> 0 – Disabled. DMA Error Interrupt Events are ignored and will not cause PCI Express interrupts to be generated. 1 – Enabled. DMA Error Interrupt Event interrupts are enabled for this DMA Channel.

Register Offset	Config Register	Register Description
		<ul style="list-style-type: none"> Bits[0] – AXI Interrupt Enable <ul style="list-style-type: none"> 0 – Disable Interrupts. Hold PCIe interrupts which occur for this DMA Channel and assert them when PCIe Interrupt Enable is set back to 1. Pending interrupts are cleared when DMA_Enable == 0. 1 – Enable Interrupts. Allow PCIe interrupts to be asserted by this DMA Channel.
0x6C	AXI_INTERRUPT_STATUS	<ul style="list-style-type: none"> Bits[31:4] – Reserved. Bits[3] – AXI Software Interrupt Status. <ul style="list-style-type: none"> 0 – No software interrupt pending 1 – A Software Interrupt was generated Bits[2] – AXI DMA SGL Interrupt Status. <ul style="list-style-type: none"> 0 – No DMA interrupt pending 1 – DMA SGL Interrupt Events caused an interrupt to be generated Bits[1] – AXI DMA Error Interrupt Status. <ul style="list-style-type: none"> 0 – No DMA interrupt pending 1 – DMA Error Interrupt Events caused an interrupt to be generated Bits[0] – Reserved.
0x70	PCIE_INTERRUPT_ASSERT	<ul style="list-style-type: none"> Bits[31:4] – Reserved. Bits[3] – PCIe Software Interrupt. Write a 1 to this register to generate a PCIe Software Interrupt. An interrupt is generated and propagated through the PCI Express Core to PCIe in the same manner as a DMA Channel interrupt (using the same Interrupt Vector). Interrupts are generated by writes to this register independent of whether the DMA Channel is enabled. Bits[2:0] – Reserved.
0x74	AXI_INTERRUPT_ASSERT	<ul style="list-style-type: none"> Bits[31:4] – Reserved. Bits[3] – AXI Software Interrupt. Write a 1 to this register to generate a AXI Software Interrupt. An interrupt is generated and propagated through to AXI in the same manner as a DMA Channel interrupt (using the same Interrupt Vector). Interrupts are generated by writes to this register independent of whether the DMA Channel is enabled. Bits[2:0] – Reserved.
0x78	DMA_CONTROL	<ul style="list-style-type: none"> Bits[31:3] – Reserved. Bits[2] – DMA Channel Completion Status Queue Element Size. If User ID and/or User Handle information is needed for an application, the application must setup the DMA Channel with 64-bit Status Queue Elements. If these features are not needed, then the DMA Channel may be setup with 32-bit Status Queue Elements to reduce the bus utilization required to write status Queue Elements. <ul style="list-style-type: none"> 0 – 32-bit - Status Queue elements contain only essential DMA information 1 – 64-bit - Status Queue elements include additional User ID and/or User Handle information that is useful for some applications Bits[1] – DMA Channel Reset. Each DMA Channel has a small Source SGL FIFO, Destination SGL FIFO, Source DMA Completion Status FIFO, and Destination DMA Completion Status FIFO to enable overlapping of DMA transactions for higher throughput. When a DMA Channel is disabled, these FIFOs may not empty fully and may need to be flushed before the DMA Channel can be reused for a new operation. <ul style="list-style-type: none"> 0 – Normal operation. 1 – Reset DMA Channel Source SGL FIFO, Destination SGL FIFO, Source DMA Completion Status FIFO, and Destination DMA Completion Status FIFO Bits[0] – DMA Channel Enable <ul style="list-style-type: none"> 0 – Disable DMA Channel. DMA Channel Queues must be configured while the DMA Channel is disabled. After disabling a DMA Channel, software must wait for all outstanding DMA transactions to complete before re-enabling the DMA Channel or software can cause the prior disabled, but not yet quiescent DMA

Register Offset	Config Register	Register Description
		<p>operation, to fail potentially causing a fatal error. DMA Running can be read to determine when a prior DMA operation has completed.</p> <ul style="list-style-type: none"> 1 – Enable DMA Channel. DMA Channel fetches Source and Destination Queue Scatter Gather Elements, executes them, and writes DMA Packet Completion status to the Source and Destination DMA Completion Status Queues. The DMA Channel pauses DMA data transactions whenever it does not have at least 1 queue element available for execution in all queues. Queue LIMIT registers are advanced to provide additional elements to execute.
0x7C	DMA_STATUS	<ul style="list-style-type: none"> Bits[31:16] – Reserved. Bits[15] – DMA Channel Present. During initialization, the DMA Driver can read this register at all possible DMA Channel Register locations to determine how many DMA Channels are implemented. <ul style="list-style-type: none"> 0 – Otherwise 1 – A DMA Channel Register set is present at this location Bits[14] – Reserved. Bits[13:4] – DMA Channel Number. Unique DMA Channel Number assigned to this DMA Channel. The DMA Channel Number register is for informational purposes and is not needed for DMA operation. DMA Channel number is unique for each DMA Channel even for multi-function and SR-IOV applications. Bits[3:1] – Reserved. Bits[0] – DMA Running. Prior to transitioning DMA Enable from 0 to 1 or modifying the contents of the DMA Channel's Queue Management Registers, software must read DMA Running == 0 to verify that the prior DMA operation completed and that it is safe to re-initialize and re-start the DMA Channel. The DMA Channel Source and Destination Queues can be re-used for new DMA operations (by changing the Queue Scatter-Gather List contents) without having to disable the DMA Channel. In general software should setup the Source SGL, Destination SGL, Source DMA Completion Status, and Destination DMA Completion Status Queues once at driver initialization and keep the DMA Channel Enabled for the entire time the driver is loaded. When there is new work to do, the Driver makes the associated queue elements available to be executed by updating the Queue LIMIT pointers. <ul style="list-style-type: none"> 0 – DMA Channel is idle. 1 – DMA Channel is busy processing.

6. Example Design

The PCIe x8 IP example design is only available for the simulation purposes in this IP version. The steps to run the functional simulation is described in the [Running Functional Simulation](#) section. The PCIe x8 IP generates two types of example designs:

- DMA Design (2024.2)
- Non-DMA Design

6.1. Example Design Supported Configuration

The Example Design Supported Configuration is shown in [Table 6.1](#).

Table 6.1. PCIe x8 IP Configuration Supported by the Example Design

PCIe x8 IP User Interface Parameter	PCIe x8 IP Configuration Supported in the Example Demo Design	
	DMA Design	Non-DMA Design
Bifurcation selection	1 × 8	1 × 1, 1 × 2, 1 × 4, 1 × 8
Target Link Speed	5G, 2.5G, 8G ² , 16G ²	5G, 2.5G, 8G ² , 16G ²
Data Interface Type	AXI-MM	TLP
Number of Physical Function	1 is supported (Function 0)	1 is supported (Function 0) Multi-Function ²
Enable Descriptor ID	N/A	N/A
PCIe CSR Base Address (512 kB aligned)	N/A	0xC5200000
Rx TLP Base Address	N/A	Posted TLP Base address – 0xFFFF0000 Non-Posted TLP Base Address – 0xFFFF1000 Completion TLP Base Address – 0xFFFF2000
Optional Ports :- Enable Clkreq port Enable LTSSM disable port	N/A	N/A
Flow Control Tab	N/A	Refer to Flow Control Update section for the configuration performed in this tab.
Descriptor Queue Base address	N/A	N/A
Descriptor Queue Base Depth	N/A	N/A
Status Queue Base Address	N/A	N/A
Status Queue Base Depth	N/A	N/A
Configuration Device ID and Vendor ID Subsystem ID Subsystem Vendor ID Class Code and Revision ID	N/A	Device ID and Vendor ID is Configured from AXI-L and rest are default

PCIe x8 IP User Interface Parameter	PCIe x8 IP Configuration Supported in the Example Demo Design	
	DMA Design	Non-DMA Design
BAR 0 Enable	✓ ¹	✓ ¹
BAR 1 Enable	X ¹	X ¹
BAR 3, BAR 4, BAR 5	X ¹	X ¹
Disable Legacy Interrupt	X ¹	X ¹
Disable MSI Capability	X ¹	✓ ¹
Disable MSI-X capability	✓ ¹	✓ ¹
Enable DSN Capability	X ¹	X ¹
Maximum Payload Size Supported	512 bytes	128 bytes, 256 bytes, or 512 bytes
Disable Function Level Reset	✓ ¹	✓ ¹
Enable Extended Tag Filed	X ¹	✓ ¹
Advance Error Reporting Capability	Not supported in 2024.2 release	Refer to the Advance Error Reporting Capability section for the configuration done in this tab.

Notes:

- ✓ refers to a checked option in the PCIe x8 IP example design and X refers to an unchecked option or a non-applicable option in the PCIe x8 IP example design.
- Available in 2024 release.

6.2. Overview of the Example Design and Features

The Example Design contains the PCIe DMA design¹ and PCIe non-DMA design. Using the graphical user interface, you can test the PCIe in supported configurations. The testbench adapts and generates the testcases based on the configuration. You can configure the parameters like the PCIe generation (Gen1, 2, 3, or 4¹), PCIe lane width (x1, x2, x4, or x8¹), PCIe DMA enabled¹ or disable, and data interface to be used.

To perform a Non-DMA write or read process, you must select the following parameters:

- PCIe Gen speed
- Interface used for data transfer, TLP
- DMA support enable¹/disable

Based on these instructions, the BFM selects the type of testcase that needs to be implemented. If a 2 × 4 PCIe with non-DMA with TLP interface is selected, the BFM sends a testcase that is compatible with the DUT (PCIe Endpoint).

Note:

- Available in 2024 release.

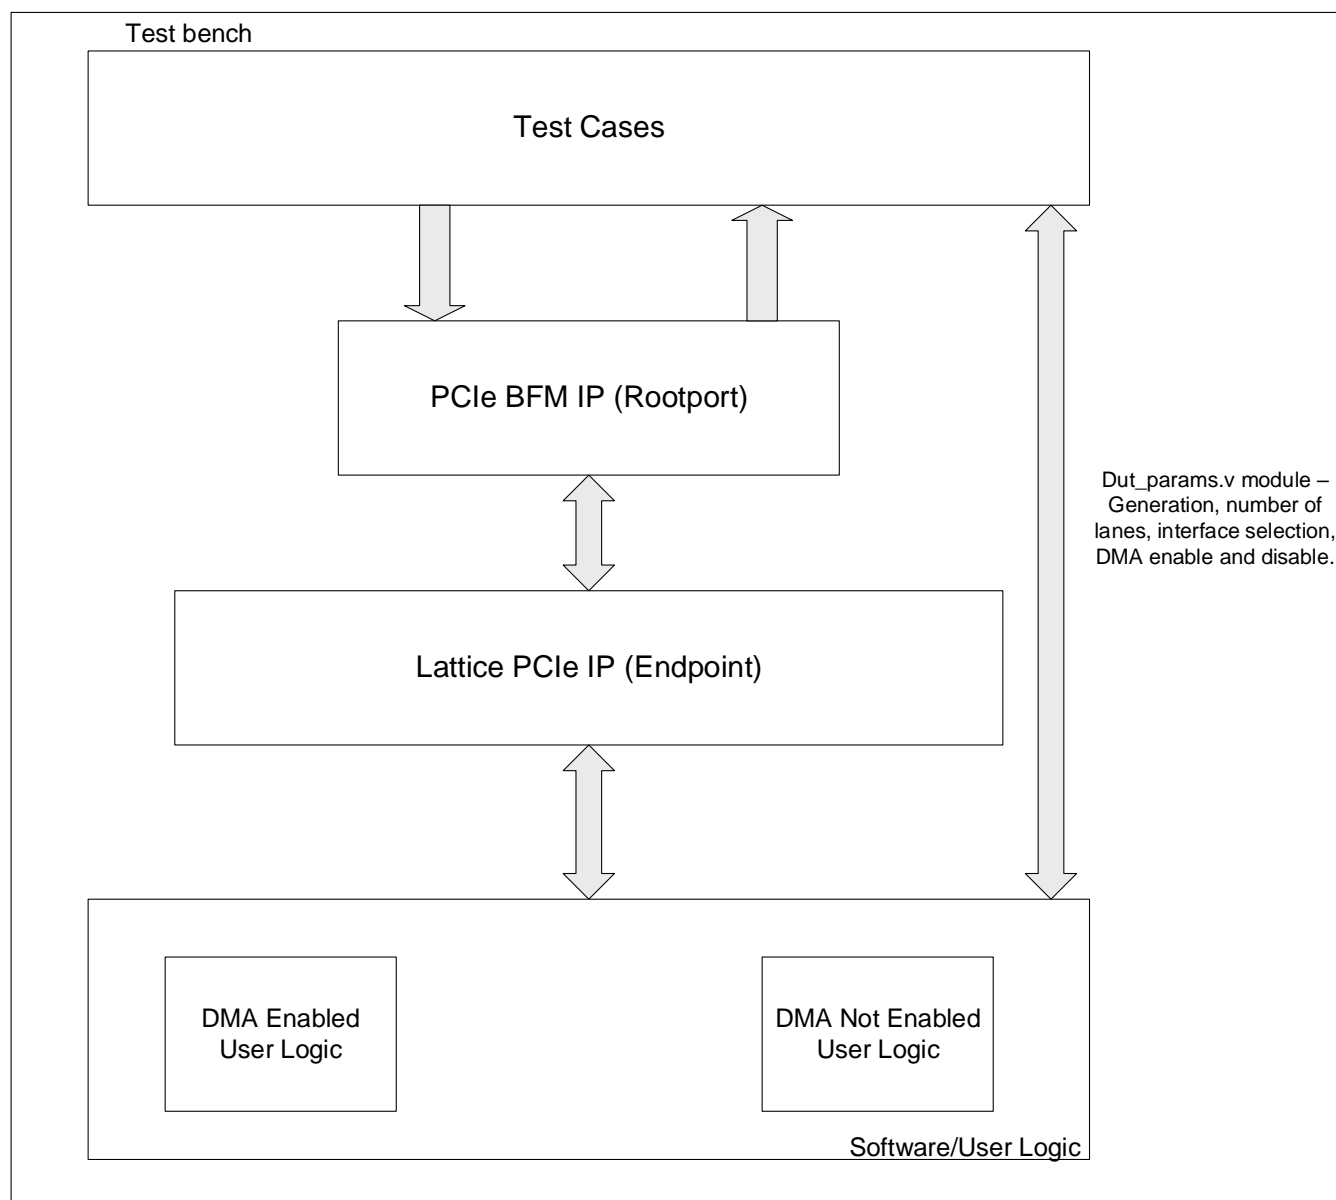


Figure 6.1. PCIe x8 IP Example Design Block Diagram

6.3. Example Design Components

6.3.1. DMA Design

The PCIe x8 harden DMA is supported in IP version 2.1.0 and onwards. The PCIe x8 implements the harden DMA example design with the following components:

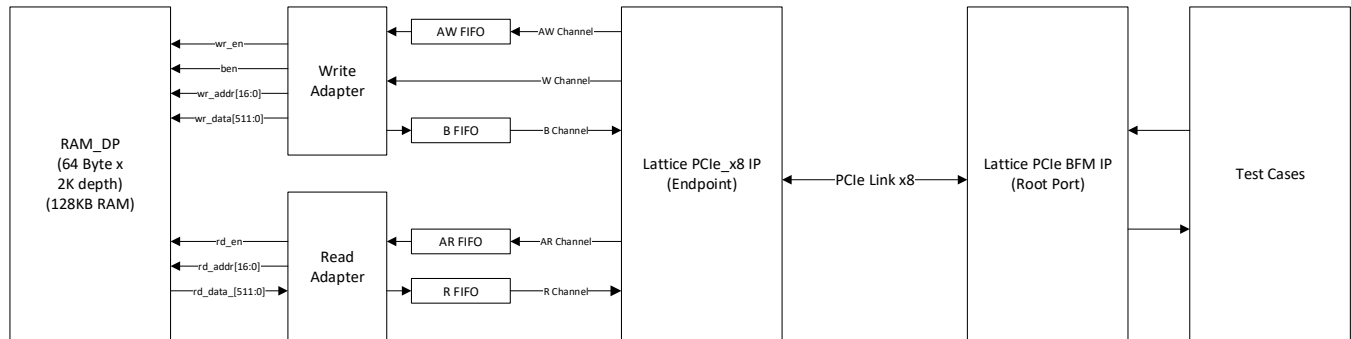


Figure 6.2. Components within the Harden DMA Design

The DMA Example Design supports both DMA mode and non-DMA mode (DMA Bypass mode). In DMA mode, the DUT fetches data from RAM_DP to Root Port or vice versa depending on DMA Descriptors. In DMA Bypass mode, the read/write operation from Root Port bypasses the DMA engine and is forwarded out to RAM_DP through the same AXI interface. The DMA Bypass mode is accessed through BAR1 while all DMA traffic uses BAR0.

Note: Both DMA and non-DMA traffic in the Example Design targets RAM_DP. This is done by ignoring the upper bits of BAR0 and BAR1 address (bit [63:17]) on the AXI address bus.

6.3.2. Non-DMA Design

The PCIe x8 Example Design implements the Non-DMA Design with the following components:

- **PCIe_rx_engine** – The received TLPs on the Rx TLP Interface is decoded in this block. For write(posted) operations, the received TLP data is sent to the *pcie_ep_mem* block to store this data into a RAM. For read(non-posted) operations, the received TLP header information is sent to the *pcie_tx_engine* block for the completion TLP.
- **PCIe_tx_engine** – The transmitted TLPs on Tx TLP Interface managed by this block. This block sends out the completion packets in response to the received non-posted TLP packets to meet the PCIe specification requirements. For example, in case of memory read TLP type packet. The required header information for the completion packet is received from the *pcie_rx_engine*. The data payload is read from the *pcie_ep_mem* block for transmitting along with the completion header.

- **PCIe_ep_mem** – The *PCIe_ep_mem* module receives the instructions from *rx_engine*, whether the data is written or read. This module consists of a RAM, which is used to store the received data. The design consists of two bars (BAR 0 and BAR 1) enabled in the PCIe endpoint. Based on the address of the BAR, the RAM writes/reads the data from/to the bar specified.

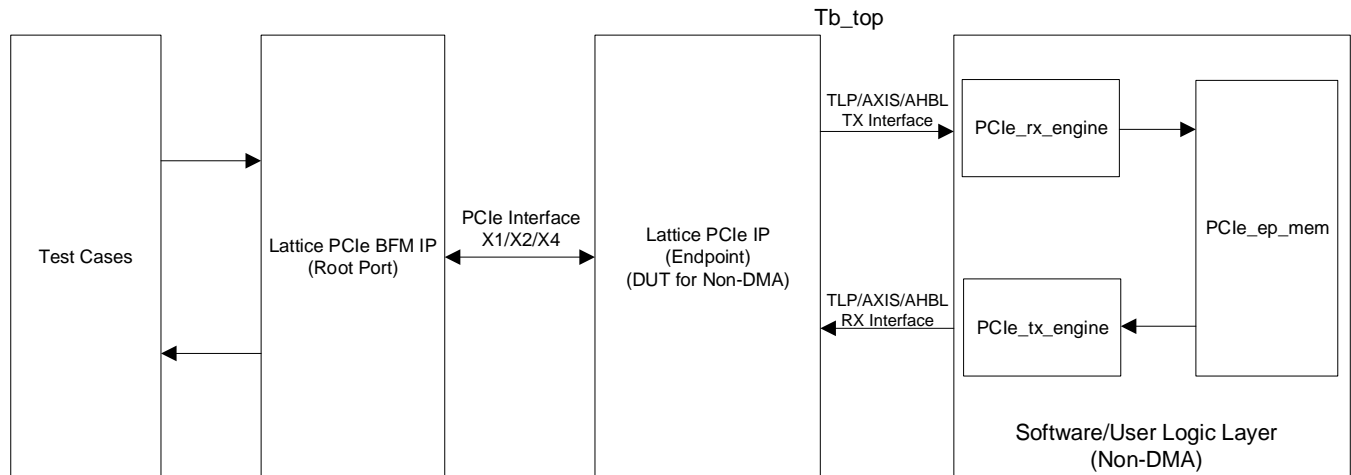


Figure 6.3. Components within NON-DMA Design

Additional Lattice IPs are used to enable the components required for the Non-DMA design as specified below:

- **Avant OSC Module – Lattice Radiant Software (FPGA-IPUG-02184)** – An internal oscillator is implemented to generate a clock to drive the LMMI clock.
- **Memory Modules User Guide (FPGA-IPUG-02033)** – A RAM_DQ of 32-bit, 64-bit, and 128-bit are used in the *pcie_ep_mem* module to store the data and can read or write the data from/to this RAM.

The following shows the Non-DMA data flow:

- Reading the configuration of Lattice PCIe x8 IP
- BFM performs the linkup with PCIe IP.
- Once linkup is done, the BFM sends the header info of the data packet to the PCIe whether to write/read into/from the preferred BAR location of the PCIe.
- PCIe sends the packet information to application layer, which the *pcie_rx_engine* decodes the header data and performs read/write operation accordingly.
- For write operations, the data is written into RAM which stores the data received.
- For read operations, the data is read from RAM and sent to the PCIe along with the header information of the packet.

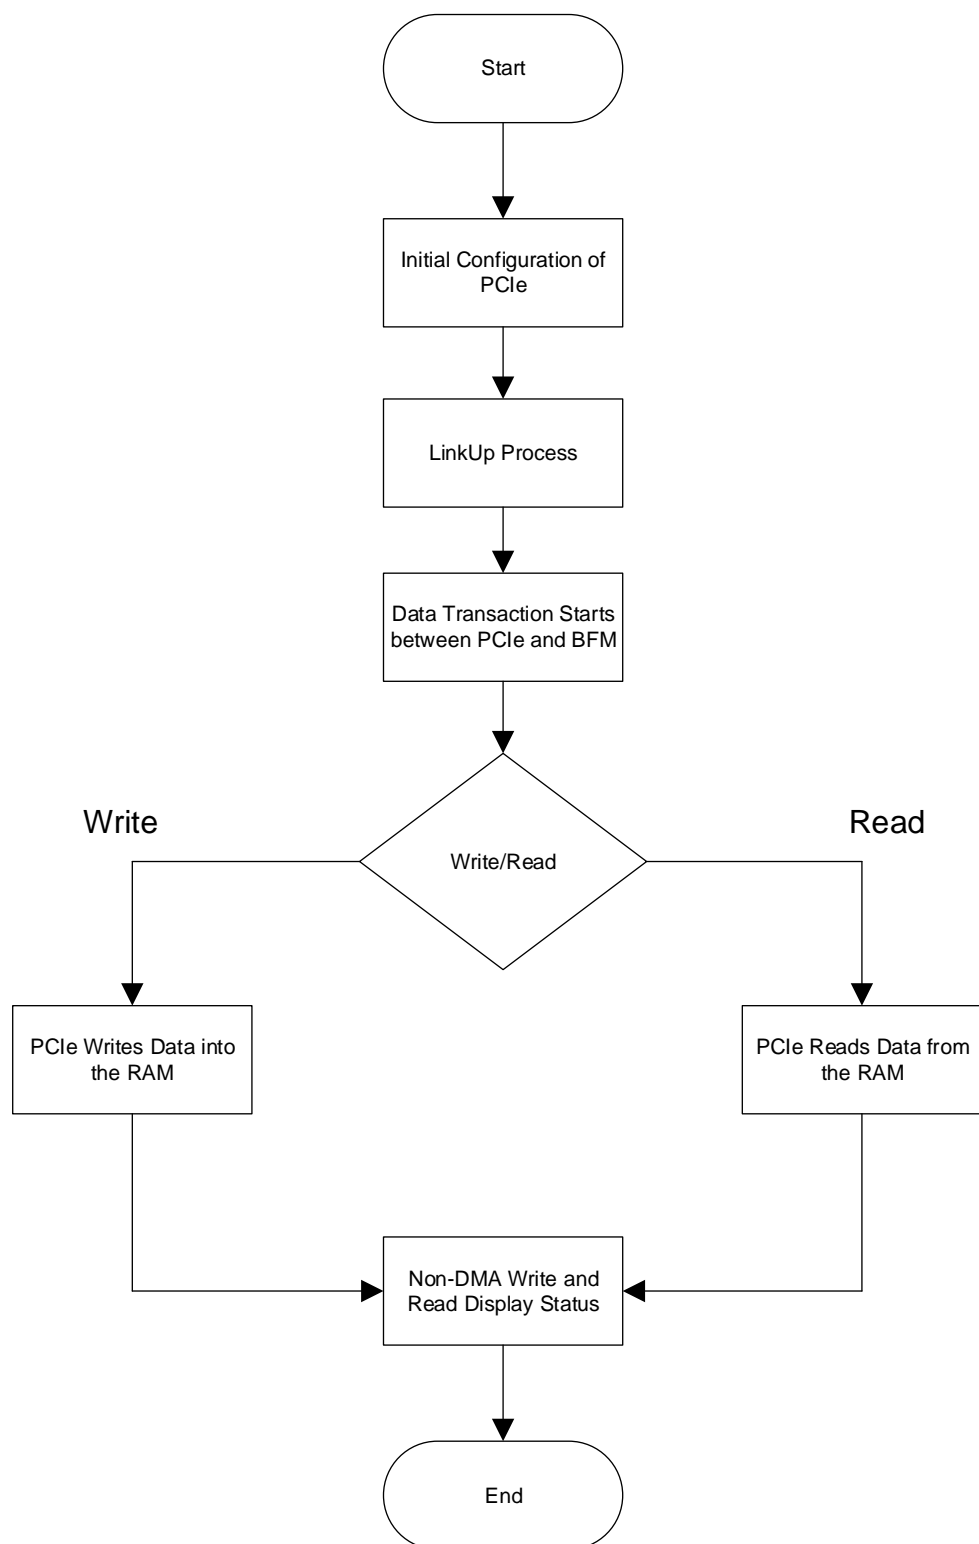


Figure 6.4. Non-DMA Design Data Flow

6.4. Simulating the Example Design




Figure 6.5. PCIe x8 IP Example Design Flowchart

The Example Design can run in simulation as follows:

1. Generate the PCIe x8 IP with the required configuration. Some of the configurations of PCIe can be done through the LMMI. The Testbench then waits for the linkup to occur.
2. Enumeration is started and wait for completion.
3. The BFM waits for the PCIe to link up.
4. The BFM starts sending the testcase based on the user requirement.
5. The status of the testcase is displayed as PASS or FAIL.

6.4.1. Running Functional Simulation (Non-DMA)

Functional Simulation can be performed after the IP is generated through the Example Design testbench. For more details on the Example Design configuration and test cases, refer to the [Example Design Supported Configuration](#) section. The limitation of functional simulation is described in the [Limitations of the Example Design](#) section.

To run the functional simulation:

1. Make sure that the testbench files are generated during PCIe x8 IP generation.

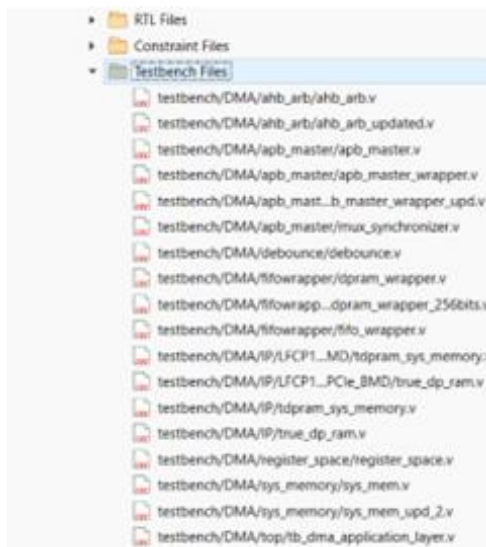



Figure 6.6. Testbench Files

2. Click the  icon to initiate the Simulation Wizard and create a new simulation project.
3. Name the project.

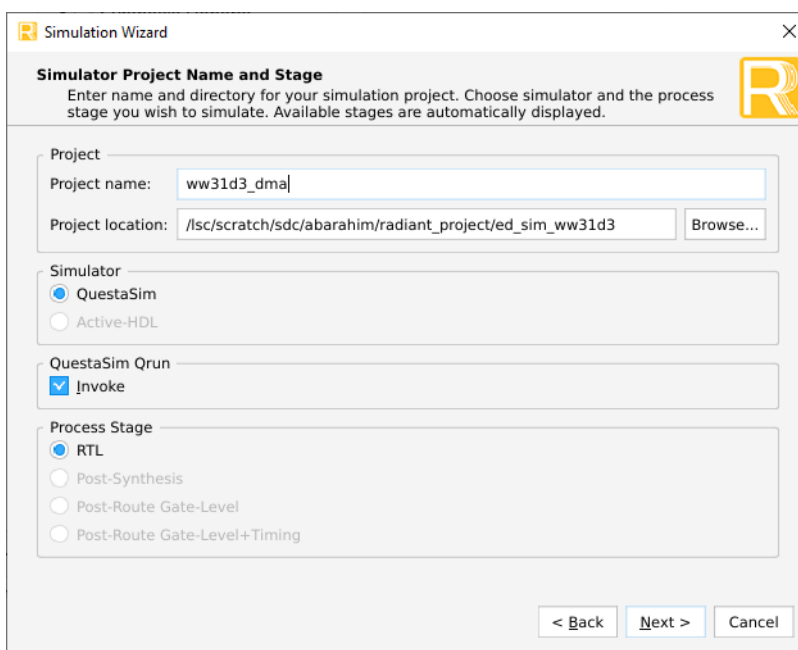


Figure 6.7. Project Naming

4. Select **tb_top** as *Simulation Top Module*.

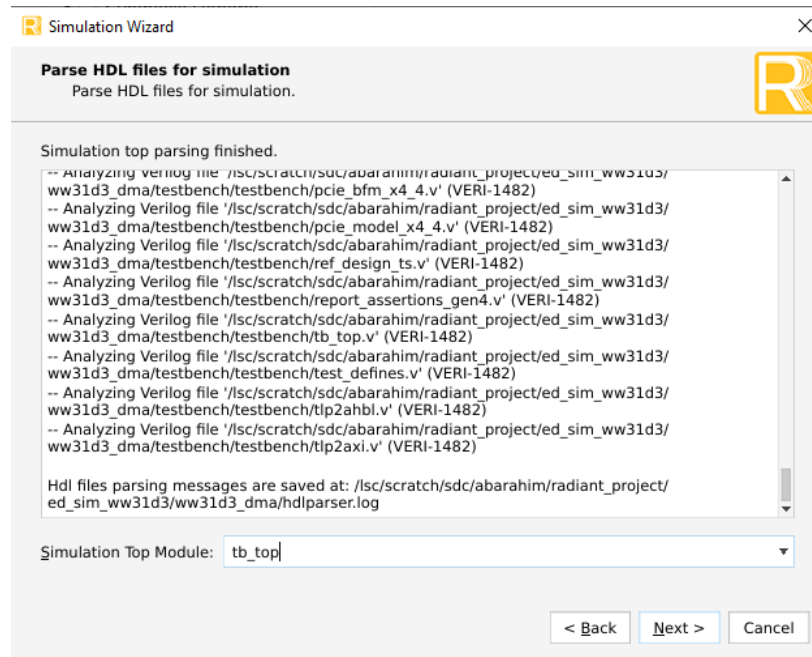


Figure 6.8. Simulation Top Module

- Use the following simulation settings. Default Run set to 0 is required.

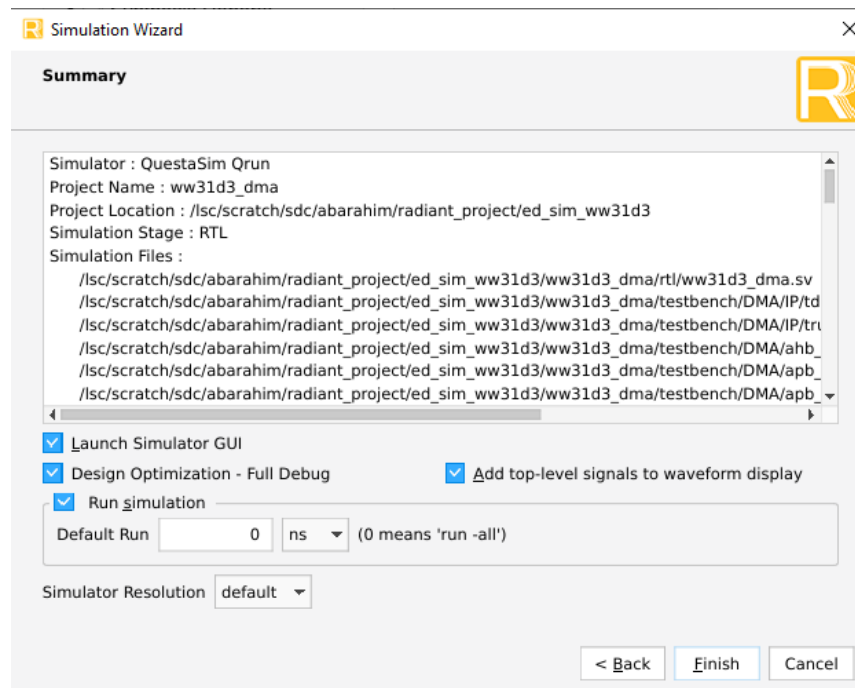


Figure 6.9. Simulation Setting

- QuestaSim Lattice-Edition is launched but it fails as you need QuestaSim OEM-Edition full license version (due to a known issue). Proceed to close the current QuestaSim window.


```
# End time: 01:08:33 on Jul 31,2024, Elapsed time: 0:01:02
# Errors: 1, Warnings: 70
# *** Summary *****
#   qrun: Errors: 0, Warnings: 0
#   vlog: Errors: 0, Warnings: 42
#   vopt: Errors: 1, Warnings: 70
# Totals: Errors: 1, Warnings: 112
# /home/rel/ng2024_1p.34/cd-LIN/diamondng/questasim_tool/questasim/linux_x86_64/qrun failed.
```

Figure 6.10. Known Simulation Error

7. Open the <project>.f and then update the -reflib to following:

For Linux:

```
-reflib /home/rel/ng2024_1p.34/cd-
LIN/diamondng/questasim_tool/cae_library/simulation/libs/pmi_work
-reflib /home/rel/ng2024_1p.34/cd-
LIN/diamondng/questasim_tool/cae_library/simulation/libs/ovi_ap6a00c
```

For Windows:

```
-reflib C:/lsc/radiant/2024.1/cae_library/simulation/libs/pmi_work
-reflib C:/lsc/radiant/2024.1/cae_library/simulation/libs/ovi_ap6a00c
```

8. In the <project>.f file, append vopt-7070 to -vopt.options.

```
-vopt.options
  -suppress vopt-7033,vopt-7070
-end
```

9. Set up the environment variable for FOUNDRY before launching the simulator.

For Linux:

- a. In the terminal used to launch the simulator, set the environment variable as follow:

```
setenv FOUNDRY <Radiant_installation_directory>/rtf/cae_library
```

Example: setenv FOUNDRY /home/rel/ng2024_1p.34/rtf/cae_library

For Windows:

- a. Run the Command Prompt as Admin, then set the FOUNDRY as shown below.
- b. After that, run the QuestaSim OEM full version.

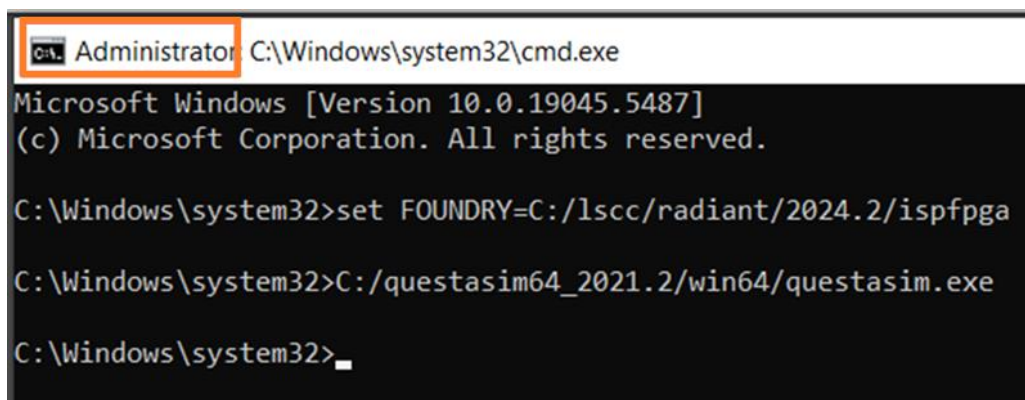


Figure 6.11 System Environment Variable for Windows

10. Launch full license QuestaSim OEM. Make sure QuestaSim is at the correct project directory and run following qrun command.

```
QuestaSim> pwd
# /lsc/scratch/sdc/abraham/radiant_project/ed_sim_ww31d3/ww31d3_dma
QuestaSim> grun -clean -f ww31d3_dma.f |
```

Figure 6.12. Command of Full License QuestaSim OEM

11. Once simulation is completed, the log printing below must be observed.

[illegible]

Figure 6.13. Expected Log Printing

12. For the Non-DMA example design, you can change the below parameters in *tb_top* module, under Non-DMA Design parameters section only:

- NON_DMA_TESTCASE_TYPE
 - 3'd0 – SINGLE BAR MEM_WRITE
 - 3'd1 – SINGLE BAR MEM_READ
 - 3'd2 – SINGLE BAR MEM_WRITE AND MEM_READ
 - 3'd3 – TWO BAR CONCURRENT MEM_WRITE AND READ
 - 3'd4 – CONFIGURE WRITE
 - 3'd5 – CONFIGURE READ
 - 3'd6 – Single BAR Multiple Write and Multiple Read transactions to incremental address with incremental data.
 - 3'd7 – Two BAR Multiple Write and Multiple Read transactions to incremental address with incremental data.
 - 3'd8 – Single BAR Multiple Write and Multiple Read transactions to same address with similar data.
 - 3'd9 – Two BAR Multiple Write and Multiple Read transactions to same address
- NON_DMA_SINGLE_BAR_SEL – This is selected only when single BAR operations are performed.
 - 1'b0 – BAR0
 - 1'b1 – BAR1
- NON_DMA_SERIES_PATTERN:
 - 1'b1 – Incremental pattern
 - 1'b0 – Fixed pattern
- NON_DMA_FIXED_DATA – You can give a fixed 32-bit data when fixed pattern is selected.
- NON_DMA_NUM_DWORD – Number of DWORDS (32-bit data) to be written in each packet.
- NON_DMA_NUM_PACKETS – Number of packets used in data transaction.
- NUM_WRITES – Number of write transactions in multiple transactions.
- NUM_READS – Number of read transactions in multiple transactions.

Note: The number of bytes in a packet in a Non-DMA design must be less than the Maximum Payload Size. Otherwise, the maximum payload size is taken as the number of bytes in a packet.

- The simulation run completion is indicated by *SIMULATION STATUS* of *ALL TESTS PASS* or *ERRORS DETECTED* in the QuestaSim transcript window. [Figure 6.14](#) shows the example of a simulation waveform.

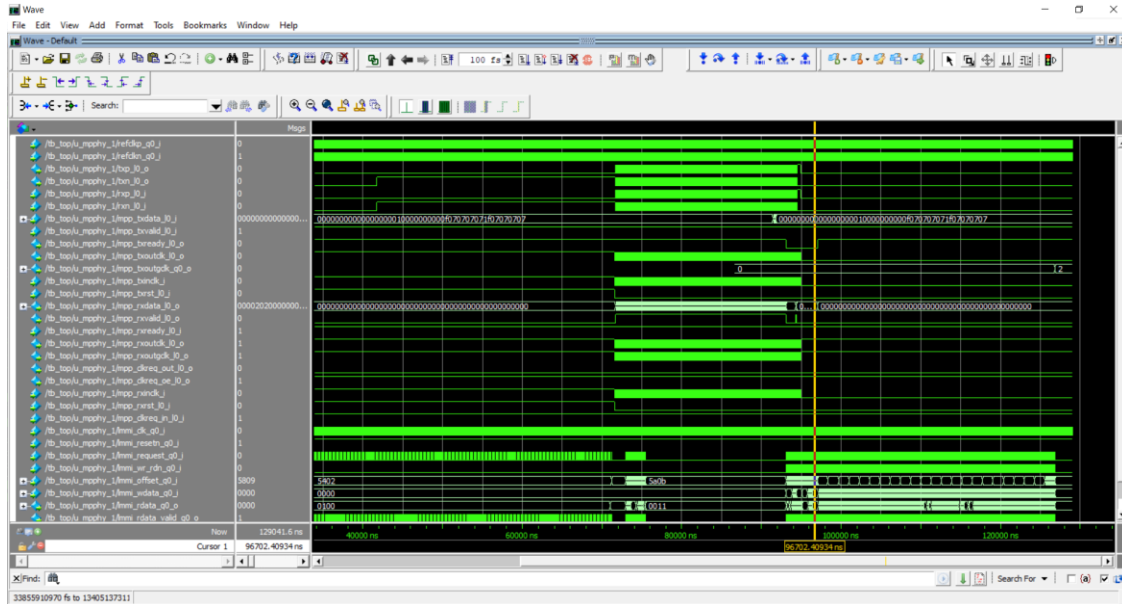


Figure 6.14. Simulation Waveform

6.4.2. Running Functional Simulation (DMA)

Functional Simulation can be performed after the IP is generated through the Example Design testbench. For more details on the Example Design configuration and test cases, refer to the [Example Design Supported Configuration](#) section. The limitation of functional simulation is described in the [Limitations of the Example Design](#) section.

To run the functional simulation:

1. Make sure that the testbench files are generated during PCIe x8 IP generation.

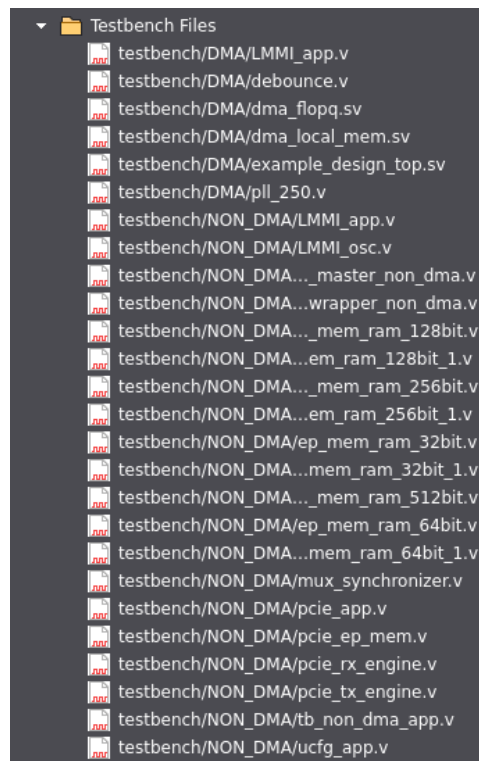



Figure 6.15. Testbench Files

- Click the  icon to initiate the Simulation Wizard and create a new simulation project.
- Name the project. Click **Next**.

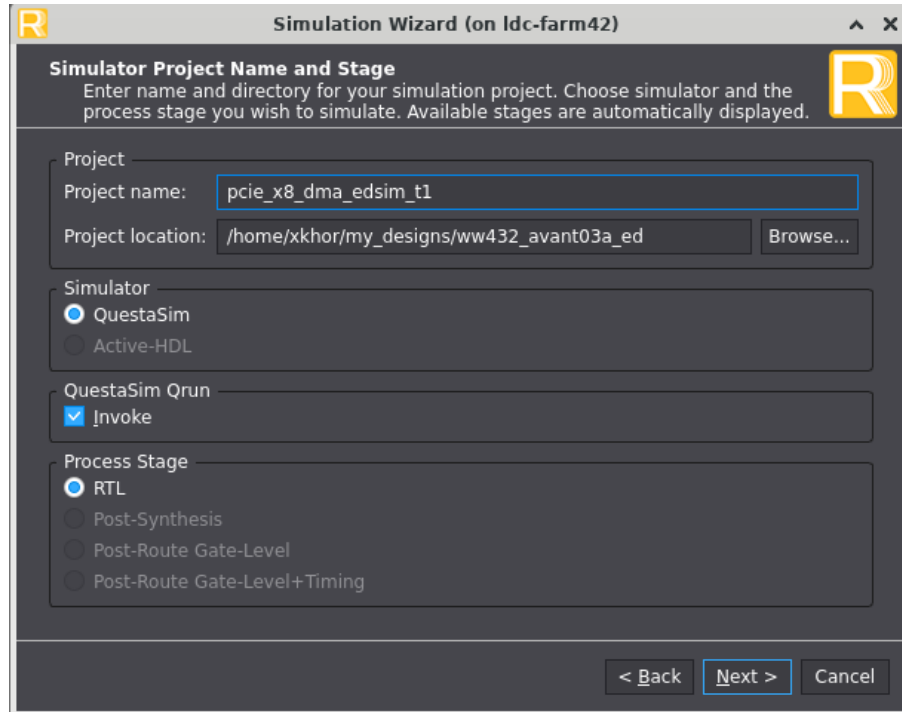


Figure 6.16. Project Naming

- Click **Next** again when the *Add and Reorder Source* window appears.
- Select **tb_top** as *Simulation Top Module*.

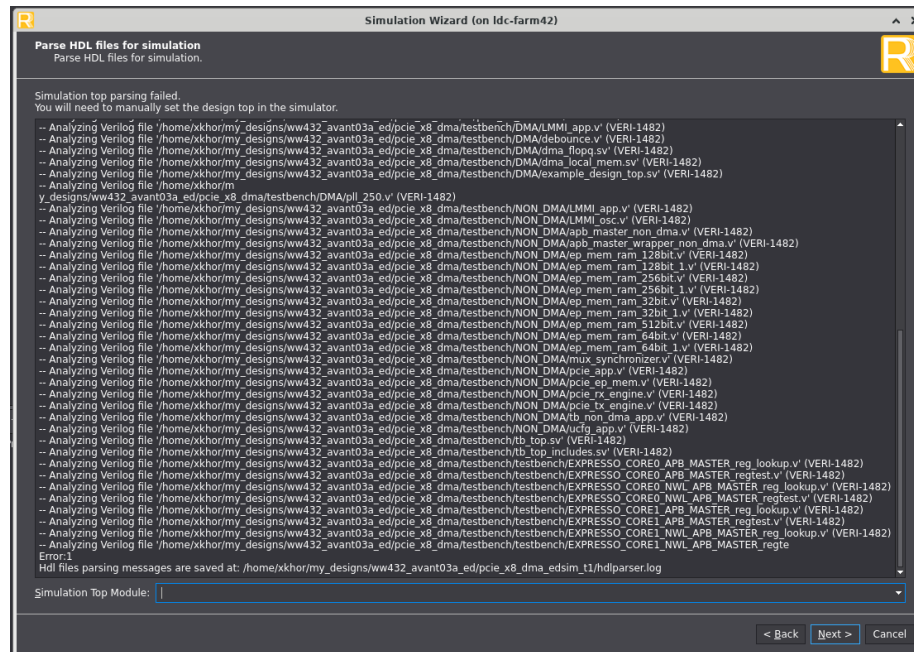


Figure 6.17. Simulation Top Naming

- Use the following simulation settings. Set *Default Run* to **0**. Click **Finish**.




Figure 6.18. Simulation Setting

7. QuestaSim Lattice Edition is launched but it fails as you need the QuestaSim OEM-Edition full license version (due to a known issue). Proceed to close the current QuestaSim window.



Figure 6.19. Known Simulation Error

8. Open the `<project>.f` file and then update the `-reflib` to following:

For Linux:

```
-reflib /home/rel/ng2024_1p.34/cd-  
LIN/diamondng/questasim_tool/cae_library/simulation/libs/pmi_work  
-reflib /home/rel/ng2024_1p.34/cd-  
LIN/diamondng/questasim_tool/cae_library/simulation/libs/ovi_ap6a00c
```

For Windows:

```
-reflib C:/lsc/radiant/2024.1/cae_library/simulation/libs/pmi_work  
-reflib C:/lsc/radiant/2024.1/cae_library/simulation/libs/ovi_ap6a00c
```

9. In the `<project>.f` file, append `vopt-7070` to `-vopt.options`:

```
-vopt.options  
-suppress vopt-7033,vopt-7070  
-end
```


6.5. Generating the Example Design for Hardware

The Example Design can also run on hardware from the IPK version 2.2.0 onwards.

6.5.1. DMA Variant

To successfully generate the DMA example design for hardware using the Lattice Radiant software, perform the following steps:

1. Create a Lattice Radiant software project. Double-click the **PCI Express for Avant** in the IP Catalog, configure the IP as required for your specific application. Be aware of the supported configurations when determining the configuration of the IP. The limitations are stipulated in the [Limitations of the Example Design](#) section.

The schematic diagram on the left shows the HardenDMA IP block with various input and output signals. The configuration table on the right is divided into two tabs: General and Flow Control. The General tab is currently selected, showing the following properties and values:

Property	Value
General	
Bifurcation Select (Link_X_Lane)	1X8
Ref Clk Freq (MHz)	100
Simulation Reduce Timeout	0
Register Interface Type	LMMI
LMMI Data Width	16
Link 0 : Data Width	1024
Link 0 : PCIe Device Type	PCIe Endpoint
Target Link Speed	GEN4
Link 0 : Number of Physical Functions	1
AXI Harden DMA Enabled	<input checked="" type="checkbox"/>
Data Interface Type	AXI_MM
ASPM Capability	
Link 0 : Active State Power Management (ASPM) Support	No ASPM Support
Virtual Function Support	
PF0 : Virtual Function Enabled	<input type="checkbox"/>
PF1 : Virtual Function Enabled	<input type="checkbox"/>
PF2 : Virtual Function Enabled	<input type="checkbox"/>
PF3 : Virtual Function Enabled	<input type="checkbox"/>

Figure 6.23. IP Configuration for Harden DMA Example Design – General Tab

General	Flow Control	Link 0: Function 0
Property	Value	
▼ Configuration		
Link 0 : Device ID (16'h)	9C25	
Link 0 : Vendor ID (16'h)	1204	
Link 0 : Subsystem ID (16'h)	E004	
Link 0 : Subsystem Vendor ID (16'h)	19AA	
Link 0 : Class Code (24'h)	118000	
Link 0 : Revision ID (8'h)	04	
▼ Base Address Register 0		
Link 0 : BAR 0 : Enable	<input checked="" type="checkbox"/>	
Link 0 : BAR 0 : Address Type	Memory	
Link 0 : BAR 0 : 64 bit address	<input type="checkbox"/>	
Link 0 : BAR 0 : Prefetchable	<input type="checkbox"/>	
Link 0 : BAR 0 : Default Size (unit)	KiB (2^10)	
Link 0 : BAR 0 : Default Size (value)	256	
Link 0 : BAR 0	32'hfffc0000	
▼ Base Address Register 1		
Link 0 : BAR 1 : Enable	<input checked="" type="checkbox"/>	
Link 0 : BAR 1 : Address Type	Memory	
Link 0 : BAR 1 : Prefetchable	<input type="checkbox"/>	
Link 0 : BAR 1 : Default Size (unit)	KiB (2^10)	
Link 0 : BAR 1 : Default Size (value)	256	
Link 0 : BAR 1	32'hfffc0000	

Figure 6.24. IP Configuration for harden DMA example design – Link 0: Function 0 tab (Part 1)

General		Flow Control	Link 0: Function 0
Property		Value	
▼ Legacy Interrupt			
Link 0 : Disable Legacy Interrupt		<input type="checkbox"/>	
Link 0 : Interrupt Pin		INTA	
▼ MSI Capability			
Link 0 : Disable MSI Capability		<input type="checkbox"/>	
Link 0 : Number of MSI vectors		2	
Link 0 : Enable Vector Masking		<input checked="" type="checkbox"/>	
▼ MSI-X Capability			
Link 0 : Disable MSI-X Capability		<input checked="" type="checkbox"/>	
▼ Device Serial Number Capability			
Link 0 : Enable DSN Capability		<input type="checkbox"/>	
▼ PCI Express Capability			
Link 0 : Maximum Payload Size Supported		512_BYTES	
Link 0 : Disable Function Level Reset (FLR)		<input checked="" type="checkbox"/>	
Link 0 : Enable Extended Tag Field		<input type="checkbox"/>	
▼ Advance Error Reporting Capability			
Link 0 : Enable ECRC Generation and Checking		<input checked="" type="checkbox"/>	
Link 0 : Enable Reporting : Correctable Internal Error		<input type="checkbox"/>	
Link 0 : Enable Reporting : Surprise Down Error		<input type="checkbox"/>	
Link 0 : Enable Reporting : TLP Prefix Blocked error		<input type="checkbox"/>	
Link 0 : Enable Reporting : Completion Timeout Error		<input checked="" type="checkbox"/>	
Link 0 : Enable Reporting : Completer Abort Error		<input type="checkbox"/>	
Link 0 : Enable Reporting : Uncorrectable Internal Error		<input type="checkbox"/>	
▼ ATS Capability			

Figure 6.25. IP Configuration for Harden DMA Example Design – Link 0: Function 0 Tab (Part 2)

2. Add the DMA files. Right-click on the Input Files and select **Add > Existing Files**. Manually add all the DMA Example Design files located in the folder of the generated IP under the *<instance name>/testbench/DMA* directory. The files include **debounce.v**, **dma_flopq.sv**, **dma_local_mem.sv**, **example_design_top.sv**, and **pll_250.v**.
3. Set **example_design_top.sv** as the top level of the design.
4. Select the PDC file: *<instance name>/eval/constraint_HDMA_ED.pdc*.
5. Once you have completed all the steps, proceed with generating the bitstream and programming it into the Avant Versa Board. Note of the device OPN stated in the Avant Versa Board User Guide.

Refer to the document linked in [Avant-G/X PCIe Host DMA Driver Software User Guide \(FPGA-TN-02405\)](#) for the steps on how to set up the driver for the DMA version of the IP as well so that you are able to interact with the IP and perform required transactions.

6.5.2. NON-DMA Variant

To successfully generate the NON-DMA example design on the hardware using the Lattice Radiant software, perform the following steps:

1. Create a Lattice Radiant software project. Double-click the **PCI Express for Avant** in the IP Catalog, configure the IP as required for your specific application. Be aware of the supported configurations when determining the configuration of the IP. The limitations are stipulated in the [Limitations of the Example Design](#) section.

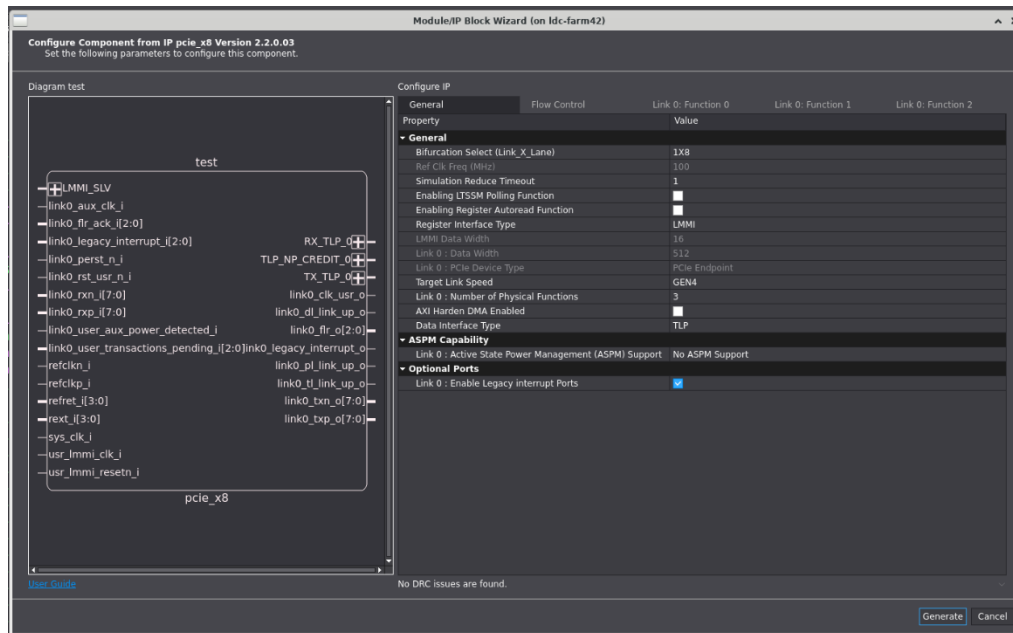
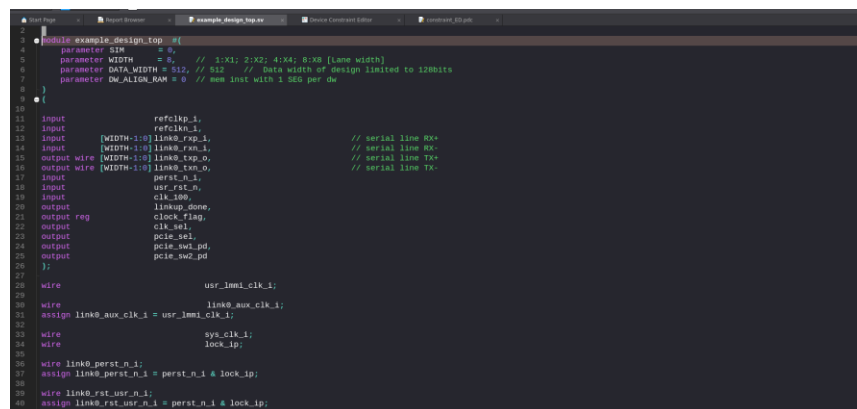


Figure 6.26. IP Configuration for Non-DMA Example Design – General Tab

2. Add the NON-DMA files. Right-click on the input files and select **Add > Existing Files**. Manually add all the Non-DMA Example Design files located in the folder of the generated IP under the *<instance name>/testbench/NON_DMA* directory. The following shows the list of files in the folder:

- axi2tlp.v
 - debounce.v
 - display_controller.v
 - ep_mem_arm_64bit.v
 - ep_mem_arm_128bit.v
 - ep_mem_arm_256bit.v
 - ep_mem_arm_512bit.v
 - example_design_top.sv
 - LMMI_app.v
 - osc_ip.v
 - pcie_app.v
 - pcie_ep_mem.sv
 - pcie_flopq.sv
 - pci_tx_engine.sv
 - pci_rx_engine.sv
 - pll_clk
 - tp2axi.v
3. Set **example_design_top.sv** as the top level of the design.
 4. Modify the parameters. Adjust the WIDTH parameter based on the selected bifurcation.



```

1  module example_design_top #(
2      parameter SIM = 0,
3      parameter WIDTH = 8, // 1x1; 2x2; 4x4; 8x8 (lane width)
4      parameter DATA_WIDTH = 512, // 512 // data width of design limited to 128bits
5      parameter DM_ALIGN_RAM = 0 // mem init with 1.5GB per dw
6  )
7  {
8      input refclkp_i;
9      input refclk_n_i;
10     input [WIDTH-1:0] link0_rxp_i; // serial line Rx+
11     input [WIDTH-1:0] link0_rxp_i; // serial line Rx-
12     output wire [WIDTH-1:0] link0_txp_o; // serial line Tx+
13     output wire [WIDTH-1:0] link0_txo_o; // serial line Tx-
14     input wire [WIDTH-1:0] link0_txw_o;
15     input perst_n_i;
16     input usr_rst_n_i;
17     input clk_500;
18     output linkup_done;
19     output req;
20     output clock_flag;
21     output clk_sel;
22     output pcie_sel;
23     output pcie_sel_pd;
24     output pcie_sw2_pd;
25 }
26
27 wire usr_lmli_clk_i;
28 wire link0_aux_clk_i;
29 assign link0_aux_clk_i = usr_lmli_clk_i;
30
31 wire sys_clk_i;
32 wire lock_ip;
33
34 wire link0_perst_n_i;
35 assign link0_perst_n_i = perst_n_i & lock_ip;
36
37 wire link0_rst_usr_n_i;
38 assign link0_rst_usr_n_i = perst_n_i & lock_ip;

```

Figure 6.27. WIDTH Parameter Adjustment in the example_design_top Module According to Bifurcation Selection

5. Select the PDC file: <instance name>/eval/constraint_ED.pdc.
6. Once you have done all of these steps, proceed to generating the bitstream and programming it into an Avant Versa Board. Note of the Device OPN stated in the Avant Versa Board User Guide.
7. Refer to the document linked on [CertusPro-NX and Avant-G/X PCIe Basic Memory-Mapped Host Driver \(Non-DMA\) User Guide \(FPGA-TN-02387\)](#) for steps on how to set up the Driver for the NON-DMA version of the IP as well so that you can interact with the IP and perform required transactions

6.6. Design Test Case Examples

6.6.1. DMA Design

The testcases are implemented using the task included in the ref_design_ts module. In this example design, the host is the Lattice PCIe BFM and the FPGA will be the RAM component. In the DMA design, the tasks performed are:

- `setup_pcie_axi_hdma_h2f_scatter_gather_queue`
 - This task is used to setup the scatter gather queue for the host-to-FPGA (H2F) DMA transaction. For the host-to-FPGA data transfer, the source and destination addresses are configured to the address space of Lattice PCIe BFM and RAM respectively. The DMA data transfer flags for the host-to-FPGA transaction are setup in this task also.
- `setup_pcie_axi_hdma_regs`
 - This task is used to setup the scatter gather queue registers for the DMA transfer. The source and destination queue pointers, source and destination queue size, source and destination next queue pointer and source and destination queue limit are configured. Finally, the interrupt generation for the respective DMA channel is configured also.
- `enable_dma_channel`
 - This task is used to enable the DMA channel to start the DMA transfer.
- `disable_dma_channel`
 - This task is used to disable the DMA channel to stop the DMA transfer.
- `setup_pcie_axi_hdma_f2h_scatter_gather_queue`
 - This task is used to setup the scatter gather queue for the FPGA-to-host (F2H) DMA transaction. For the FPGA-to-host data transfer, the source and destination addresses are configured to the address space of RAM and Lattice PCIe BFM respectively. The DMA data transfer flags for the FPGA-to-host transaction are setup in this task also.
- `reinit_pcie_axi_hdma_regs`
 - This task is used to re-initialize the the scatter gather queue registers for the DMA transfer if the scatter gather queue is being setup again. The source and destination next queue pointer and source and destination queue limit are re-configured.
- `dma_data_check`
 - This task is used to do the DMA data check between the data being transferred through the H2F and F2H DMA transactions.

6.6.2. Non-DMA Design

The testcases are implemented using the task included in the `ref_design_ts` module. For the non-DMA design, the task performed is:

- `generic_testcase_for_non_dma`
 - This task is used to write data patterns for single transactions and multiple transactions. The task also reads data from PCIe IP and implements a status check on the read data.

Table 6.2 lists the status ports related to the non-DMA design operation.

Table 6.2. Non-DMA Design Status Ports Description

Ports	Description
<code>wr_vector</code>	The array used to generate write data pattern.
<code>mul_wr_vector</code>	The array used to generate write data pattern in case of multiple transactions.
<code>address</code>	The value corresponds to address of the DWORD (32-bit data).
<code>read_data</code>	The read data from the PCIe IP when read test case is implemented.
<code>READ_DATA_CHECK</code>	This signal indicates the status of the read data.

6.7. Debugging Example Design Issues

6.7.1. Signals to Debug

6.7.1.1. Simulation Debug for Non-DMA Design

Table 6.3. Non-DMA Signals to Debug Description

Module Name	Signal Name	Description
<code>tb_top</code>	<code>lmmi_offset</code>	Lower 17-bit address of LMMI interface registers

Module Name	Signal Name	Description
tb_top	lmmi_wdata	Data written into PCIe IP through LMMI interface
tb_top	lmmi_rdata	Data read from PCIe IP through LMMI interface
tb_top	link0_pl_link_up	PCIe IP physical layer linkup
tb_top	link0_dl_link_up	PCIe IP data layer linkup
tb_top	link0_tl_link_up	PCIe IP transaction layer linkup

The following are the steps to debug the non-DMA design:

- You must check whether the initial configuration is performed properly. You can check the *lmmi_offset*, *lmmi_wdata*, and *lmmi_rdata* signals to verify. Note that in the actual application, the register configuration may not be necessary if the corresponding register is configured through the IP Catalog's Module/IP wizard.
- The linkup signals such as *link0_pl_link_up*, *link0_dl_link_up*, and *link0_tl_link_up* must be asserted.

6.8. Limitations of the Example Design

For non-DMA variant supported in IP version 1.0.0 up to version 2.1.0, the functional simulation with the example design testbench can only support the following:

- Gen 1 (2.5G)/Gen2 (5G) for non-DMA mode
- TLP interface for Data, LMMI interface for Config
- Bifurcation 1 × 1, 1 × 2, and 1 × 4
- MSI-X capability disabled
- The Example Design simulation is not supported in ModelSim OEM-Edition, ModelSim Lattice-Edition and QuestaSim Lattice-Edition. Use the QuestaSim OEM-Edition for simulation.

For non-DMA variant supported in IP version 2.2.0 and onwards, the functional simulation with the example design testbench can only support the following:

- Gen 1 (2.5G)/Gen2 (5G)/Gen3 (8G)/Gen 4(16G) for non-DMA mode
- TLP interface for Data, LMMI interface for Config
- Bifurcation 1 × 1, 1 × 2, 1 × 4 and 1 × 8
- MSI-X capability disabled
- The Example Design simulation is not supported in ModelSim OEM-Edition, ModelSim Lattice-Edition and QuestaSim Lattice-Edition. Use the QuestaSim OEM-Edition for simulation.

For Harden DMA variant supported in IP version 2.1.0 and onwards, the functional simulation with the example design testbench can only support the following IP configuration:

- Gen 4 (16G)
- AXI-MM interface for Data, LMMI interface for Config
- Bifurcation 1 × 8
- No Virtual Function Supported
- MSI-X capability disabled

7. Designing with the IP

This section provides information on how to generate the IP Core using the Lattice Radiant software and how to run simulation and synthesis. For more details on the Lattice Radiant software, refer to the Lattice Radiant software user guide.

7.1. Generating and Instantiating the IP

You can use the Lattice Radiant software to generate IP modules and integrate them into the device's architecture. The steps below describe how to generate the PCIe x8 IP in the Lattice Radiant software.

To generate the PCIe x8 IP:

1. Create a new Lattice Radiant software project or open an existing project.
2. In the **IP Catalog** tab, double-click **PCI Express for Avant-G/X** under **IP, Connectivity** category. The **Module/IP Block Wizard** opens as shown in [Figure 7.1](#). Enter values in the **Component name** and the **Create in** fields and click **Next**.

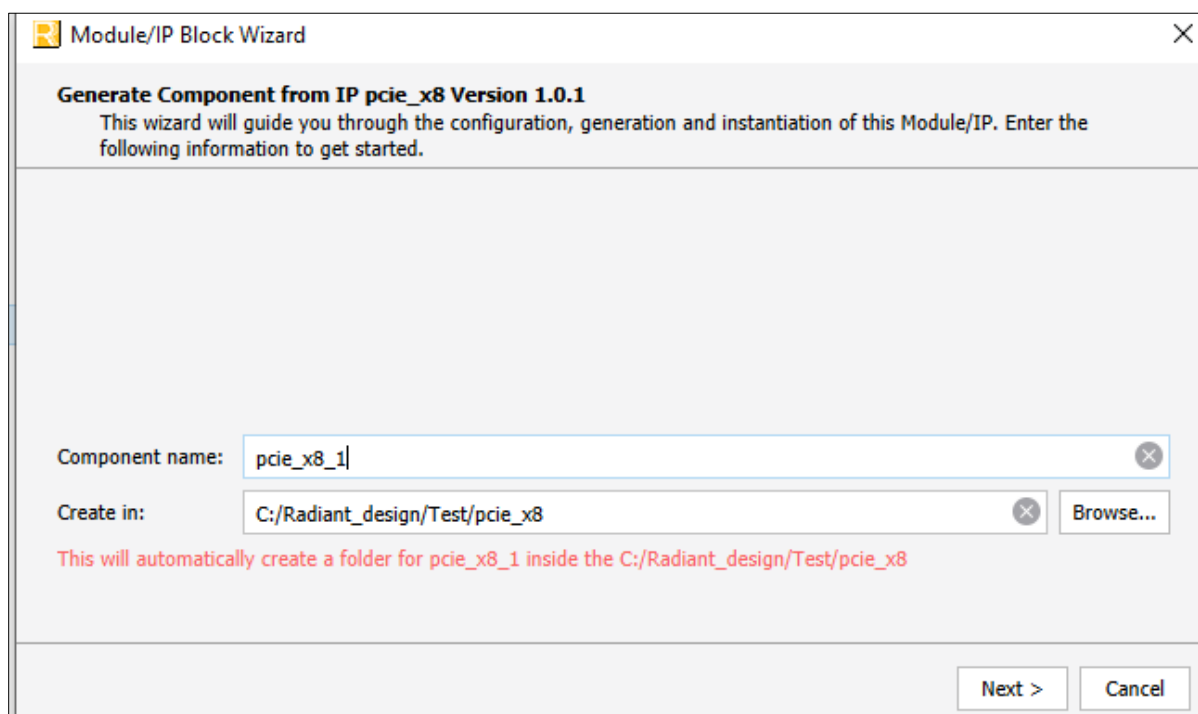


Figure 7.1. Module/IP Block Wizard

3. In the next **Module/IP Block Wizard** window, customize the selected PCIe x8 IP using drop-down lists and check boxes. [Figure 7.2](#) shows an example configuration of the non-DMA PCIe x8 IP. For the harden DMA example design support, [Figure 7.3](#) and [Figure 7.4](#) show the supported configuration of the harden DMA PCIe x8 IP. For details on the configuration options, refer to the [IP Parameter Description](#) section.

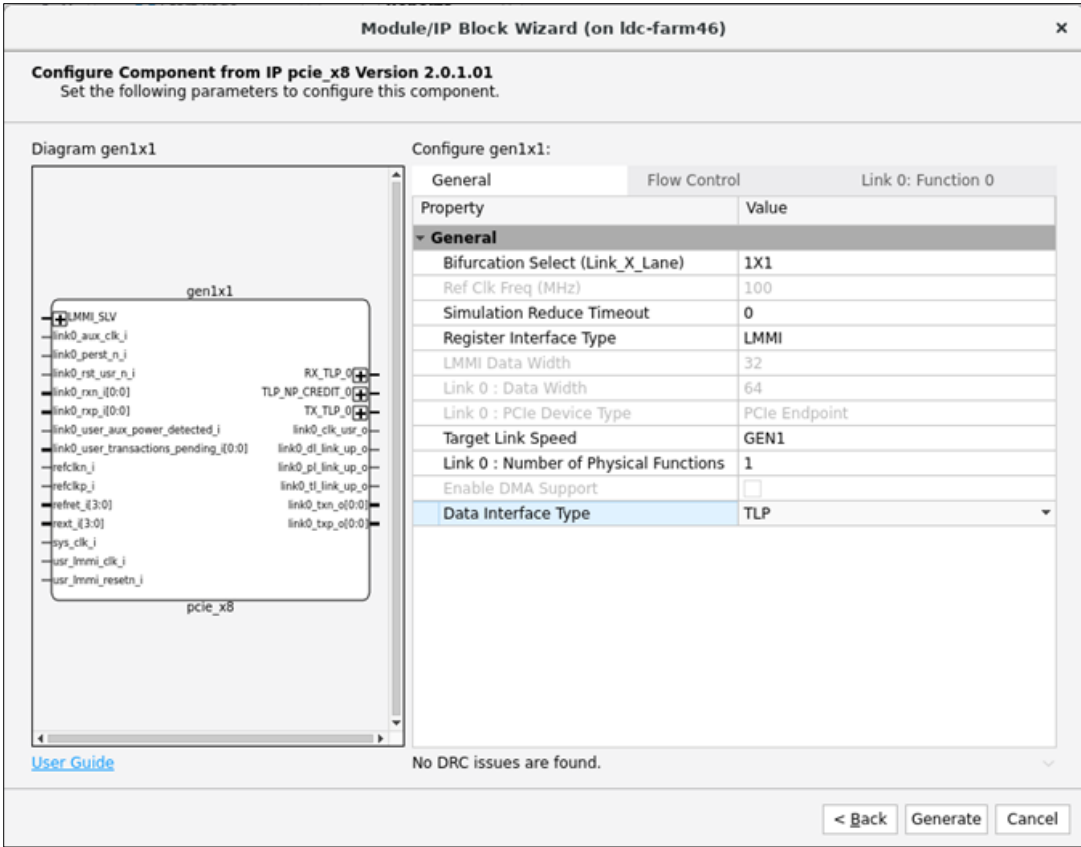


Figure 7.2. IP Configuration for Non-DMA

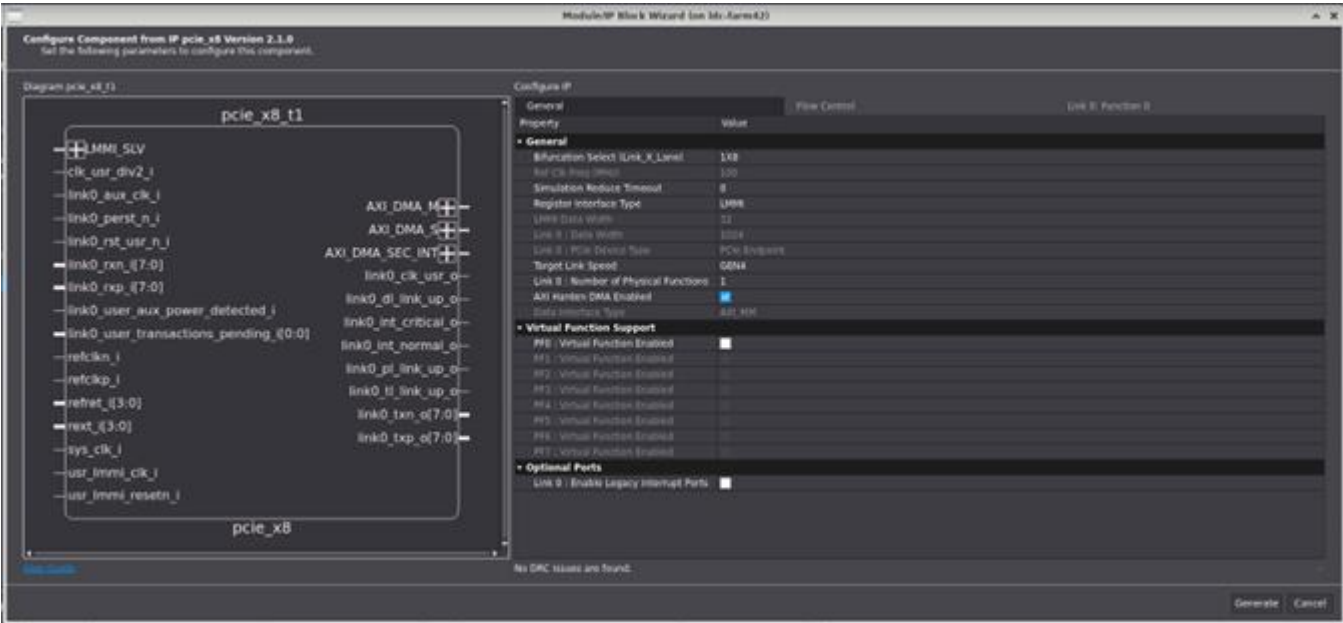


Figure 7.3. IP Configuration for Hardwired DMA (General Tab)

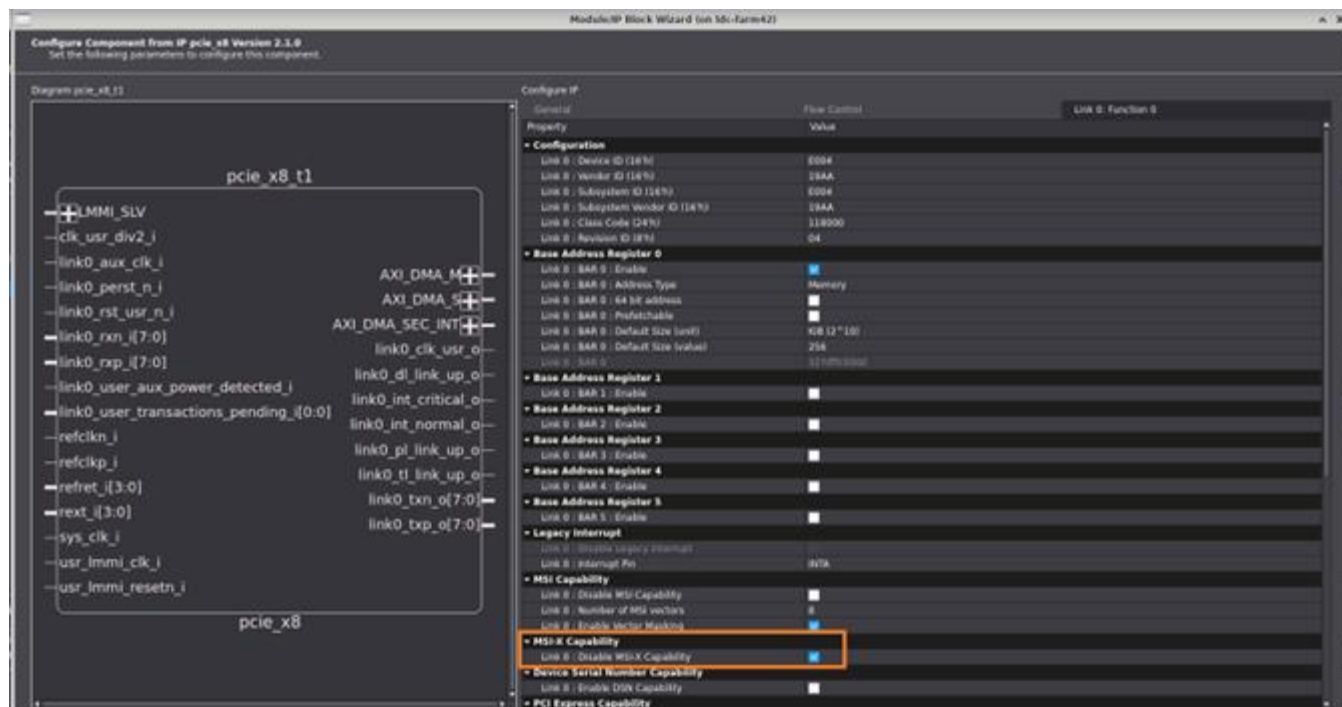


Figure 7.4. IP Configuration for Harden DMA (Link 0: Function 0 Tab)

- Click **Generate**. The **Check Generating Result** dialog box opens, showing design block messages and results as shown in Figure 7.5.

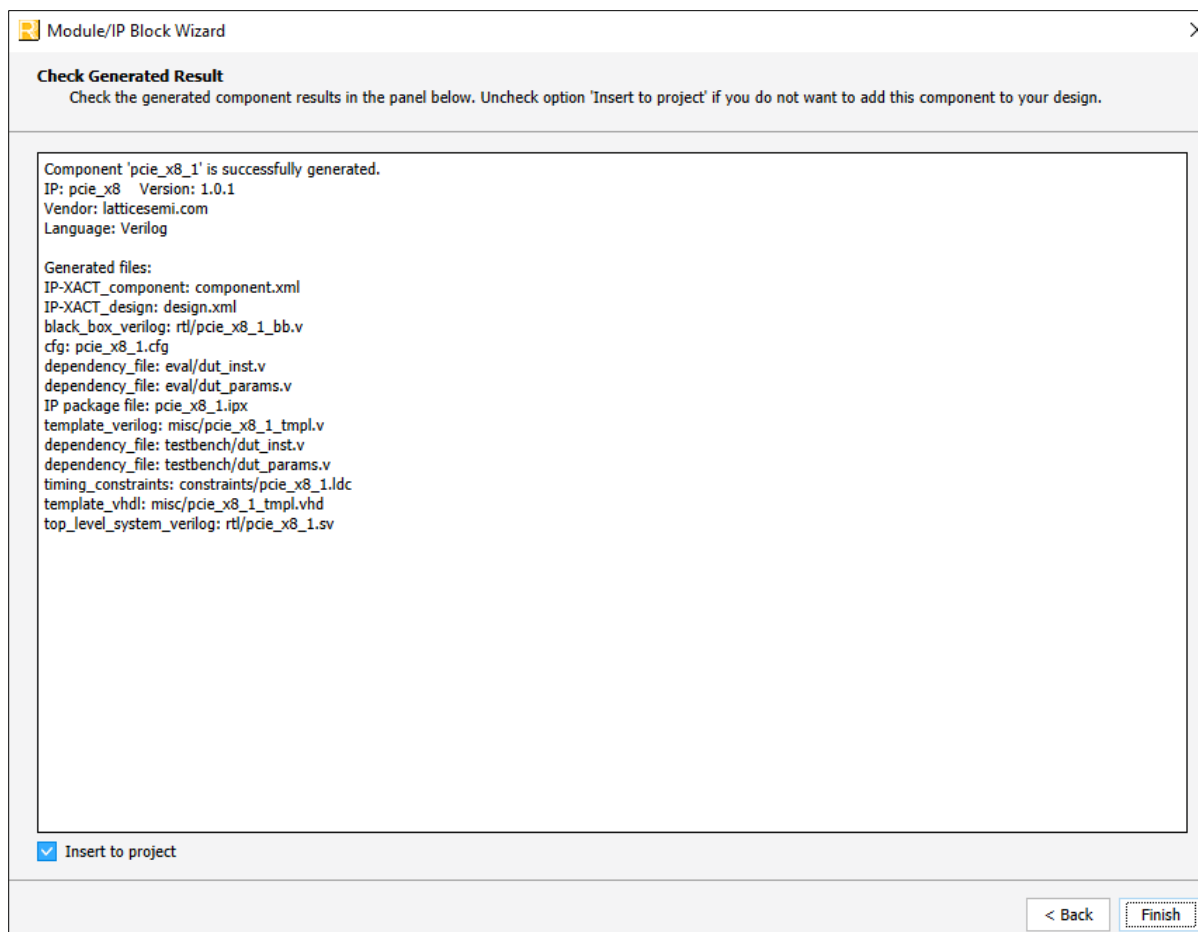


Figure 7.5. Check Generated Result

5. Click **Finish**. All the generated files are placed under the directory paths in the **Create in** and the **Component name** fields shown in [Figure 7.1](#).
6. IP version 2.1.0 or below requires an additional step to set the IP top level. Right click on the IP name and select *Set as Top-Level Unit* shown in [Figure 7.6](#).

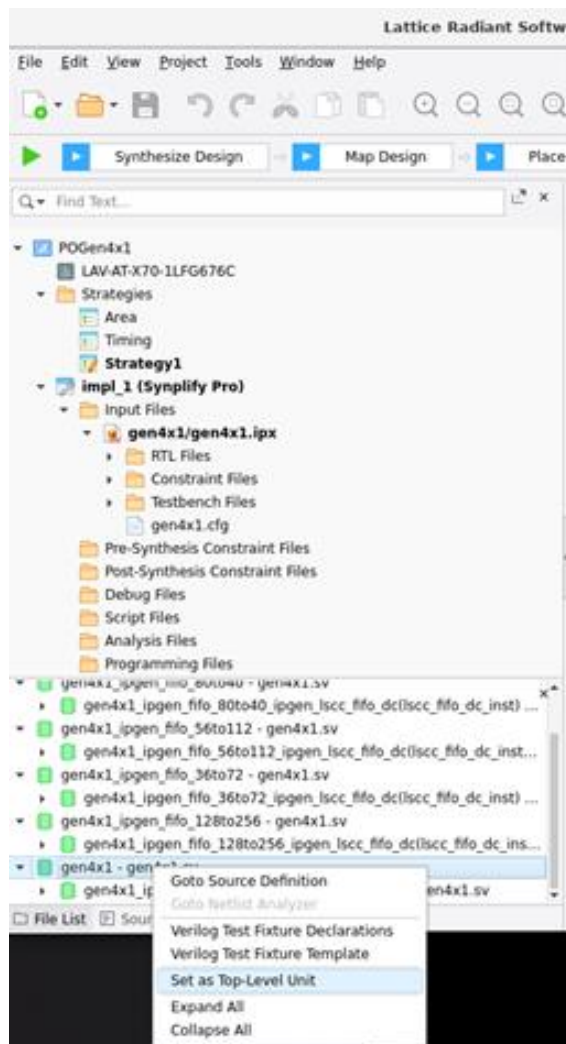


Figure 7.6. Set IP Top Level Unit

7. As this is an IP level compilation, you need to set all TLP/AXI-S/LMMI/AXI-L signal as Virtual I/O through a constraint file. In the system level design where these I/O are connected properly to the user bus or config bus, these constraints are not required.

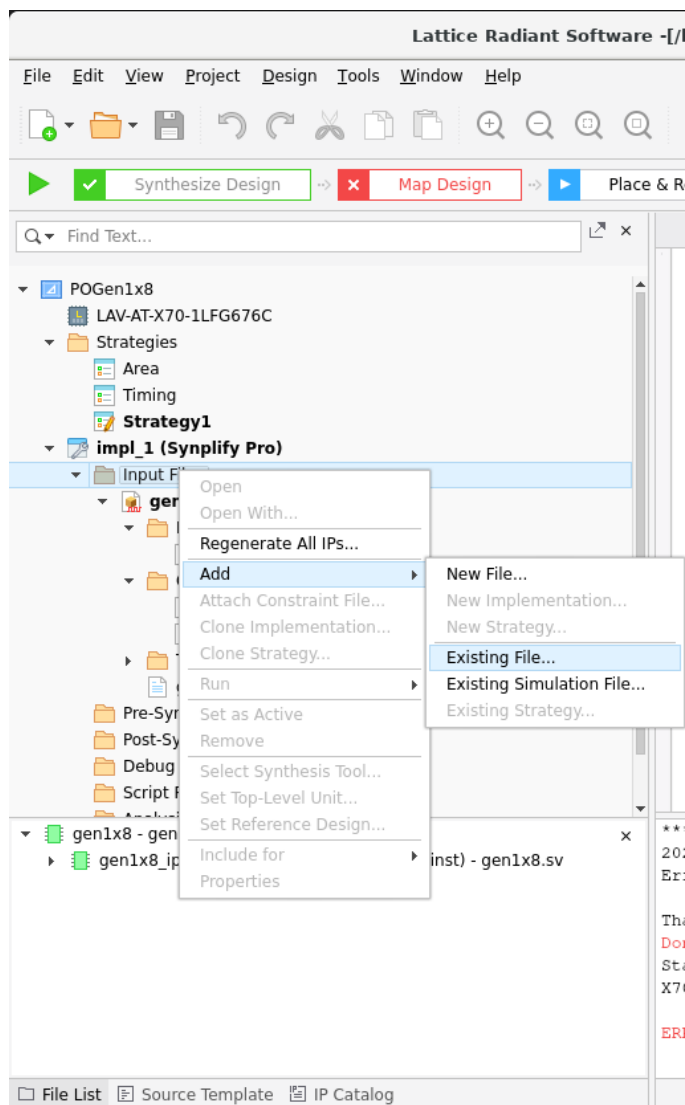


Figure 7.7. Select to Add Constraint File

8. Browse to `<project_folder>/<project_name>/<ip_name>/eval/` and select **constraint.pdc**.

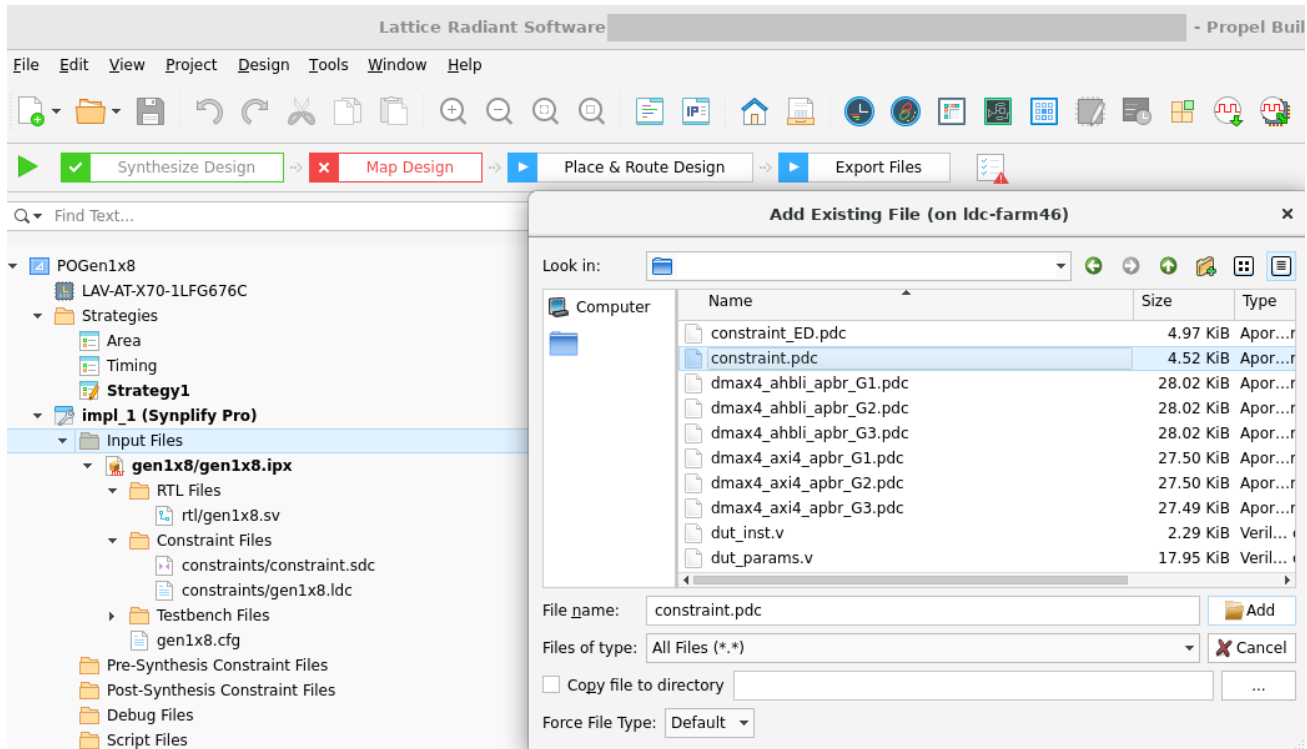


Figure 7.8. Select constraint.pdc from the Eval Folder

- Click **Run All** to compile the IP. Compilation successful is shown in Figure 7.9.

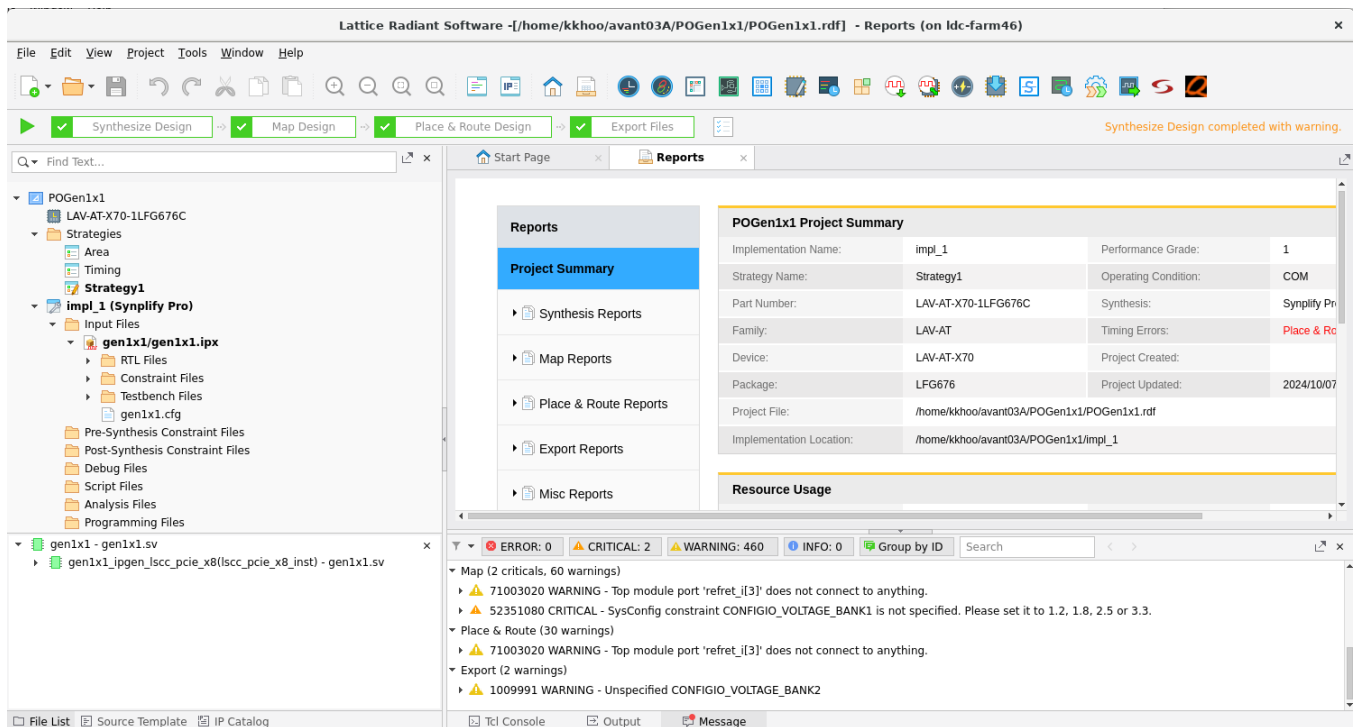


Figure 7.9. Project Compile Done

- As the IP level compile involves the virtual I/O, the bitstream is not generated. For bitstream generation, refer to PCIe Reference design on Lattice website

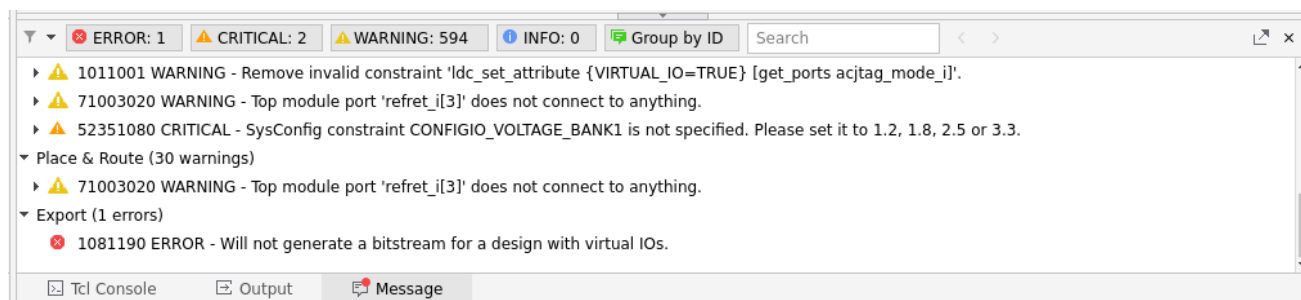


Figure 7.10. Bitstream Not Generated due to the Virtual I/O Setting

7.1.1. Generated Files and File Structure

The generated PCIe x8 module package includes the black box (<Component name>_bb.v) and instance templates (<Component name>_tmpl.v/vhd) that can be used to instantiate the core in a top-level design. An example RTL top-level reference source file (<Component name>.v) that can be used as an instantiation template for the module is also provided. You may also use this top-level reference as the starting template for the top-level for the complete design. The generated files are listed in [Table 7.1](#).

Table 7.1. Generated File List

Attribute	Description
<Component name>.ipx	This file contains the information on the files associated to the generated IP.
<Component name>.cfg	This file contains the parameter values used in IP configuration.
component.xml	Contains the ipxact:component information of the IP.
design.xml	Documents the configuration parameters of the IP in IP-XACT 2014 format.
rtl/<Component name>.v	This file provides an example RTL top file that instantiates the module.
rtl/<Component name>_bb.v	This file provides the synthesis black box.
misc/<Component name>_tmpl.v misc/<Component name>_tmpl.vhd	These files provide instance templates for the module.

7.1.2. Design Implementation

Completing the design includes additional steps to specify analog properties, pin assignments, and timing constraints. You can add and edit the constraints using the Device Constraint Editor or by manually creating a PDC File.

7.1.2.1. Device Constraint Editor

Refer to the [Lattice Radiant Software 2024.1 User Guide](#) for more information on how to use the device constraint editor.

7.1.2.2. Manual PDC File Creation

To create the manual PDC file, add the .pdc (post synthesis constraint file) file in the Lattice Radiant software and define the I/O pins according to the schematic design for ports defined in your design. You can define different types of constraints such as pins, clocks, and other timing paths. Post-Synthesis constraint files (.pdc) contain both timing and non-timing constraint .pdc source files for storing logical timing/physical constraints.

Refer to the relevant sections in the Lattice Radiant Software User Guide for more information on how to create or edit constraints.

7.1.3. Timing Constraints

The timing constraints are based on the clock frequency used. The timing constraints for the IP are defined in relevant constraint files. The example below shows the IP timing constraints generated for the PCIe x8 IP.

```
create_clock -name {sys_clk_i} -period 8 -waveform {0 4} [get_ports sys_clk_i]

if {$SUSE_DEFAULT_IF == 1} {
  create_clock -name {usr_lmml_clk_i} -period 8 -waveform {0 4} [get_ports usr_lmml_clk_i]
} else {
  if {$USR_CFG_IF_TYPE == "APB"} {
    create_clock -name {c_apb_pclk_i} -period 8 -waveform {0 4} [get_ports c_apb_pclk_i]
  }
}
```

Figure 7.11. Timing Constraint File (.sdc) for the PCIe x8 IP

- Add the timing constraints shown in [Figure 7.13](#) in the design's .pdc or constraint file. Refer to the [Lattice Radiant Software 2024.1 User Guide](#) to learn more about the .pdc file.
- For sys_clk_i and clk_usr_div2, refer to [Table 2.1](#) on selecting the frequencies for Gen1, Gen2, Gen3, or Gen4 data rates. You can use a PLL IP to create these clocks. Refer to the [PLL Module IP User Guide \(FPGA-IPUG-02063\)](#) for instantiation and generation of PLL IP. [Figure 7.12](#) shows the IP configuration for Gen 3 rates if using an input clock of 125 MHz.

Module/IP Block Wizard

Configure Component from Module pll Version 1.7.0
Set the following parameters to configure this component.

Diagram pll_clk2

Configure IP

Property	Value
Reference Clock	
CLKI: Frequency (MHz) [18 - 800]	125
CLKI: Divider Actual Value [1 - 44]	1
Phase Detector Frequency (MHz) [18 - 500]	125
Enable Reference Clock Monitor	<input type="checkbox"/>
Feedback	
CLKFB: Feedback Mode	CLKOP
CLKFB: FBK Divider Actual Value (Integer) [1 - 128]	2
Primary Clock Output	
CLKOP: Frequency Desired Value (MHz) [10 - 800]	250
CLKOP: Divider Actual Value [1 - 128]	4
CLKOP Tolerance (%)	0.0
CLKOP: ERROR (PPM)	0
CLKOP: Enable Trim for CLKOP	<input type="checkbox"/>
Secondary Clock Output	
CLKOS: Enable	<input checked="" type="checkbox"/>
CLKOS: Bypass	<input type="checkbox"/>
CLKOS: Frequency Desired Value (MHz) [6.25 - 800]	125
CLKOS: Divider Actual Value [1 - 128]	8
CLKOS Tolerance (%)	0.0
CLKOS: ERROR (PPM)	0
CLKOS: Static Phase Shift (Degrees)	0
CLKOS: Enable Trim for CLKOS	<input type="checkbox"/>
Secondary Clock Output (2)	
CLKOS2: Enable	<input checked="" type="checkbox"/>
CLKOS2: Bypass	<input type="checkbox"/>
CLKOS2: Frequency Desired Value (MHz) [6.25 - 800]	250
CLKOS2: Divider Actual Value [1 - 128]	4
CLKOS2 Tolerance (%)	0.0
CLKOS2: ERROR (PPM)	0
CLKOS2: Static Phase Shift (Degrees)	90

Figure 7.12. PLL IP Configuration for Input Clock of 125 MHz

- Define input reference clock of PLL in the timing constraints using *create_clock -name* As shown in Figure 7.13, the input reference clock of 125 MHz is sourced from an I/O pin and named *clk_125* in the design.

```
create_clock -name {clk_125} -period 8 [get_ports clk_125]
set_clock_groups -group [get_clocks clk_usr_i] -group [get_clocks c_apb_pclk_i] -asynchronous
set_clock_groups -group [get_clocks clk_usr_div2_i] -group [get_clocks c_apb_pclk_i] -asynchronous
set_clock_groups -group [get_clocks clk_usr_i] -group [get_clocks clk_usr_div2_i] -asynchronous
set_clock_groups -group [get_clocks clk_usr_ps90_i] -group [get_clocks clk_usr_div2_i] -asynchronous
set_multicycle_path -setup 2 -from [get_clocks clk_usr_i] -to [get_clocks clk_usr_ps90]
```

Figure 7.13. Timing Constraints for PCIe x8 IP Example

8. Debugging

The PCIe protocol involves the interface between a root port and endpoint with both sides being linked up. Hence, PCIe issues can range from device recognition issues, link training issue, flow control errors, enumeration issues, link down due to fatal errors, and others. This section provides the debug flow diagrams for some of the most common issues when using the PCIe x8 IP. Several debug flow charts are introduced with additional information on critical debug registers to refer to and loopback diagnostic features. This section also provides a short description on signals to be used for debugging simulation.

8.1. Debug Methods

8.1.1. Debug Flow Charts

One debugging method is to identify the type of PCIe issue. The following sections show the steps to debug various issues.

8.1.1.1. Hardware Detection Failure

Follow the steps shown in the flow diagram below if the system is not detecting the hardware.

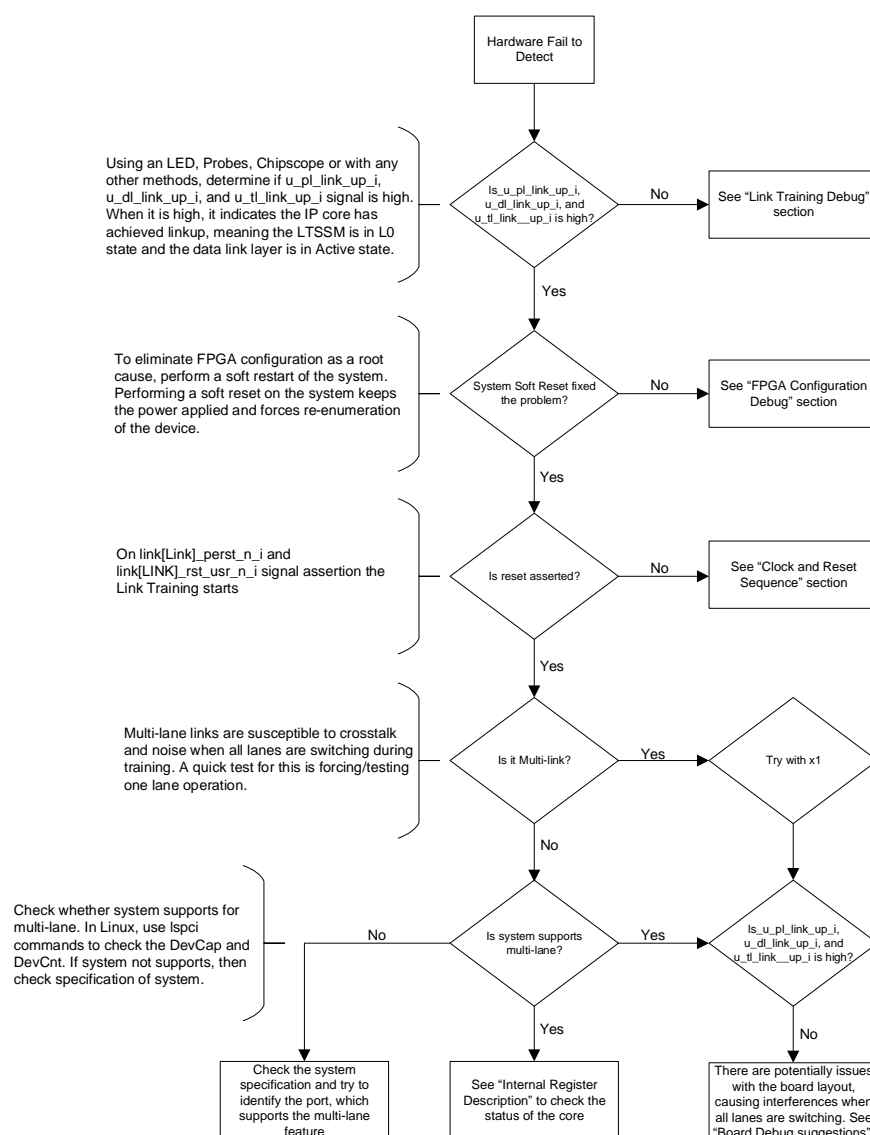


Figure 8.1. Hardware Detection Failure Debugging Flow

8.1.1.2. Link Training Debug

For the link training debug, refer to the flow chart as shown [Figure 8.2](#).

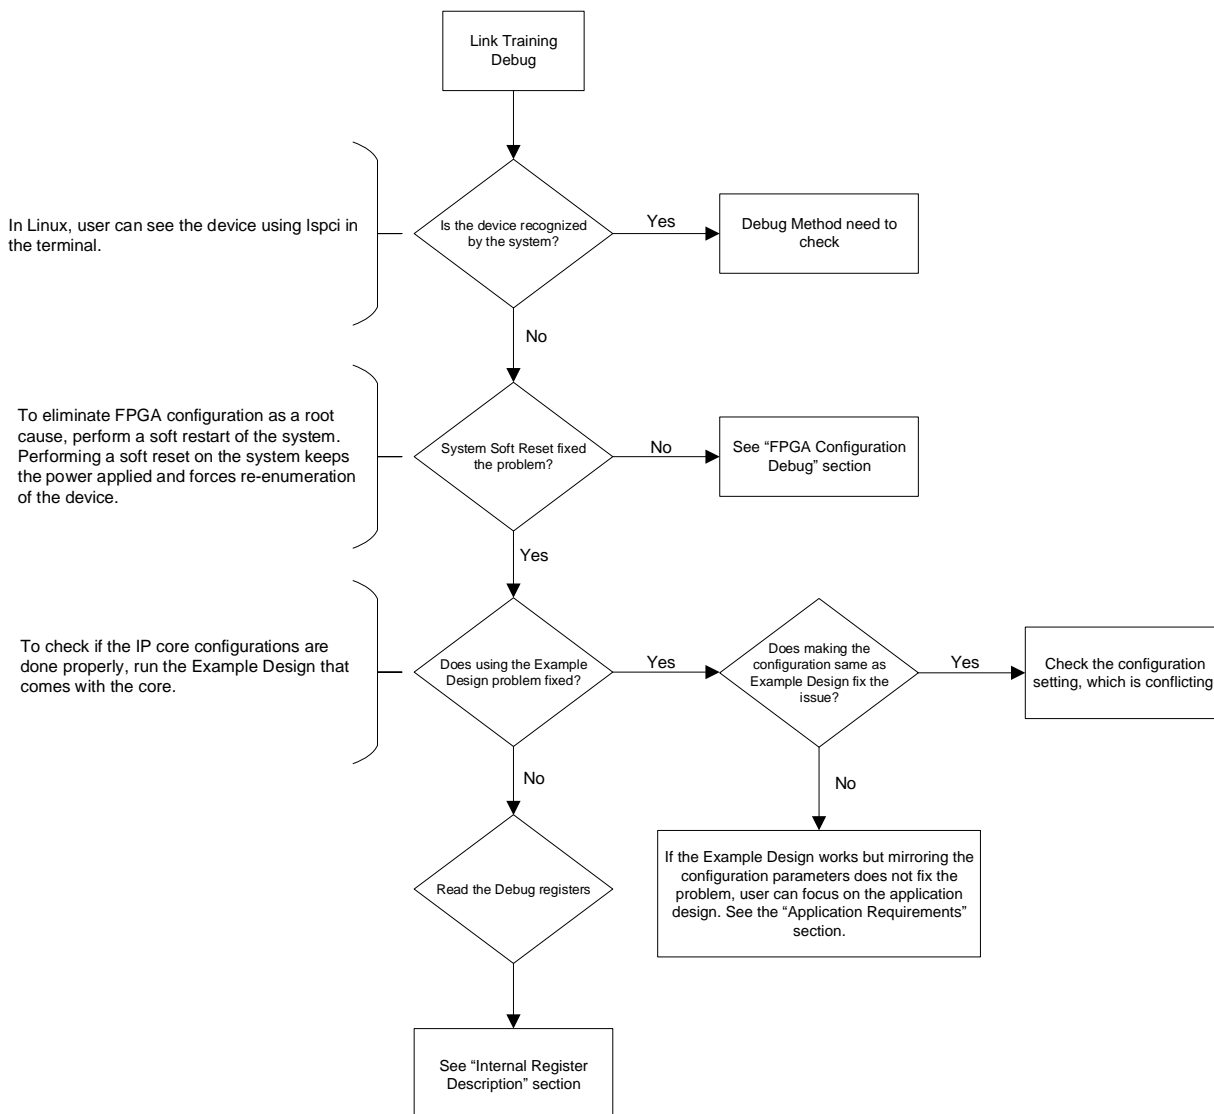


Figure 8.2. Link Training Issue Debugging Flow

You can enable the LTSSM polling feature to assist the debugging. This feature is disabled by default and must only be enabled for debugging and disabled for production. You can drive the `ltssm_polling_en` signal to 1 to enable. When enabled, the `ltssm_state` and `ltssm_sub_state` signals indicate the PCIe LTSSM state and sub-state.

General	Flow Control	Link 0: Function 0
Property		Value
General		
Bifurcation Select (Link_X_Lane)		1X8
Ref Clk Freq (MHz)		100
Simulation Reduce Timeout		0
Enabling LTSSM Polling Function		<input checked="" type="checkbox"/>

Figure 8.3. Checkbox to enable LTSSM Polling

Table 8.1. PCIe LTSSM State and Sub-State Definition

State Value	State Definition	Sub-state Value	Sub-state Definition
0	DETECT	0	DETECT_INACTIVE
		1	DETECT_QUIET
		2	DETECT_SPD_CHG0
		3	DETECT_SPD_CHG1
		4	DETECT_ACTIVE0
		5	DETECT_ACTIVE1
		6	DETECT_ACTIVE2
		7	DETECT_P1_TO_P0
		8	DETECT_P0_TO_P1_0
		9	DETECT_P0_TO_P1_1
		10	DETECT_P0_TO_P1_2
1	POLLING	0	POLLING_INACTIVE
		1	POLLING_ACTIVE_ENTRY
		2	POLLING_ACTIVE
		3	POLLING_CFG
		4	POLLING_COMP
		5	POLLING_COMP_ENTRY
		6	POLLING_COMP_EIOS
		7	POLLING_COMP_EIOS_ACK
		8	POLLING_COMP_IDLE
2	CONFIGURATION	0	CONFIGURATION_INACTIVE
		1	CONFIGURATION_US_LW_START
		2	CONFIGURATION_US_LW_ACCEPT
		3	CONFIGURATION_US_LN_WAIT
		4	CONFIGURATION_US_LN_ACCEPT
		5	CONFIGURATION_DS_LW_START
		6	CONFIGURATION_DS_LW_ACCEPT
		7	CONFIGURATION_DS_LN_WAIT
		8	CONFIGURATION_DS_LN_ACCEPT
		9	CONFIGURATION_COMPLETE
		10	CONFIGURATION_IDLE
3	L0	0	L0_INACTIVE
		1	L0_L0
		2	L0_TX_EL_IDLE
		3	L0_TX_IDLE_MIN
4	RECOVERY	0	RECOVERY_INACTIVE
		1	RECOVERY_RCVR_LOCK
		2	RECOVERY_RCVR_CFG
		3	RECOVERY_IDLE
		4	RECOVERY_SPEED0
		5	RECOVERY_SPEED1
		6	RECOVERY_SPEED2
		7	RECOVERY_SPEED3
		8	RECOVERY_EQ_PH0
		9	RECOVERY_EQ_PH1
		10	RECOVERY_EQ_PH2
		11	RECOVERY_EQ_PH3

State Value	State Definition	Sub-state Value	Sub-state Definition
5	DISABLED	0	DISABLED_INACTIVE
		1	DISABLED_0
		2	DISABLED_1
		3	DISABLED_2
		4	DISABLED_3
6	LOOPBACK	0	LOOPBACK_INACTIVE
		1	LOOPBACK_ENTRY
		2	LOOPBACK_ENTRY_EXIT
		3	LOOPBACK_EIOS
		4	LOOPBACK_EIOS_ACK
		5	LOOPBACK_IDLE
		6	LOOPBACK_ACTIVE
		7	LOOPBACK_EXIT0
		8	LOOPBACK_EXIT1
7	HOT_RESET	0	HOT_RESET_INACTIVE
		1	HOT_RESET_HOT_RESET
		2	HOT_RESET_MASTER_UP
		3	HOT_RESET_MASTER_DOWN
8	TX_LOS	0	TX_LOS_INACTIVE
		1	TX_LOS_IDLE
		2	TX_LOS_TO_L0
		3	TX_LOS_FTS0
		4	TX_LOS_FTS1
9	L1	0	L1_INACTIVE
		1	L1_IDLE
		2	L1_SUBSTATE
		3	L1_TO_L0
10	L2	0	L2_INACTIVE
		1	L2_IDLE
		2	L2_TX_WAKE0
		3	L2_TX_WAKE1
		4	L2_EXIT
		5	L2_SPEED

8.1.1.3. Data Transfer Debug

If data transfer fails, refer to the steps shown in [Figure 8.4](#).

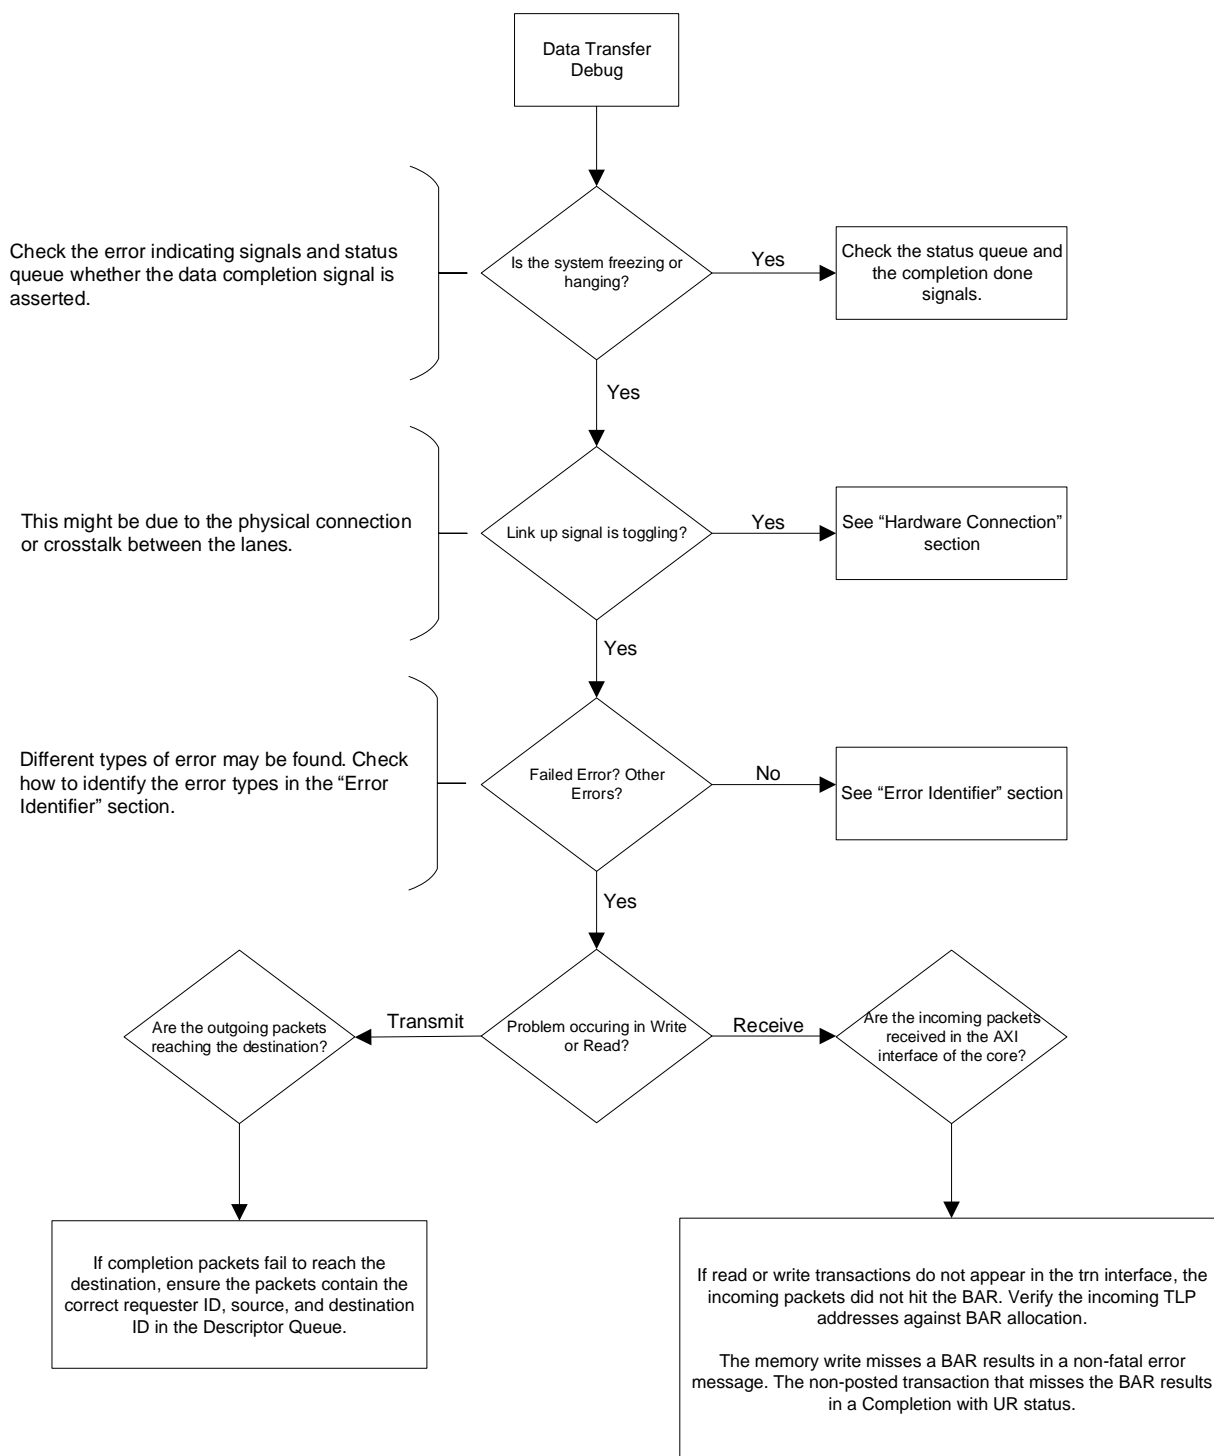


Figure 8.4. Data Transfer Issue Debugging Flow

8.1.1.4. FPGA Configuration Debug

Device initialization and configuration issues can be caused by not having the FPGA configured fast enough to enter link training and be recognized by the system. Prior to becoming operational, the FPGA goes through a sequence of states, including initialization, configuration, and wake-up. After programming the FPGA, a soft reset is required to configure the FPGA from flash. Performing the soft reset operation restarts the enumeration process.

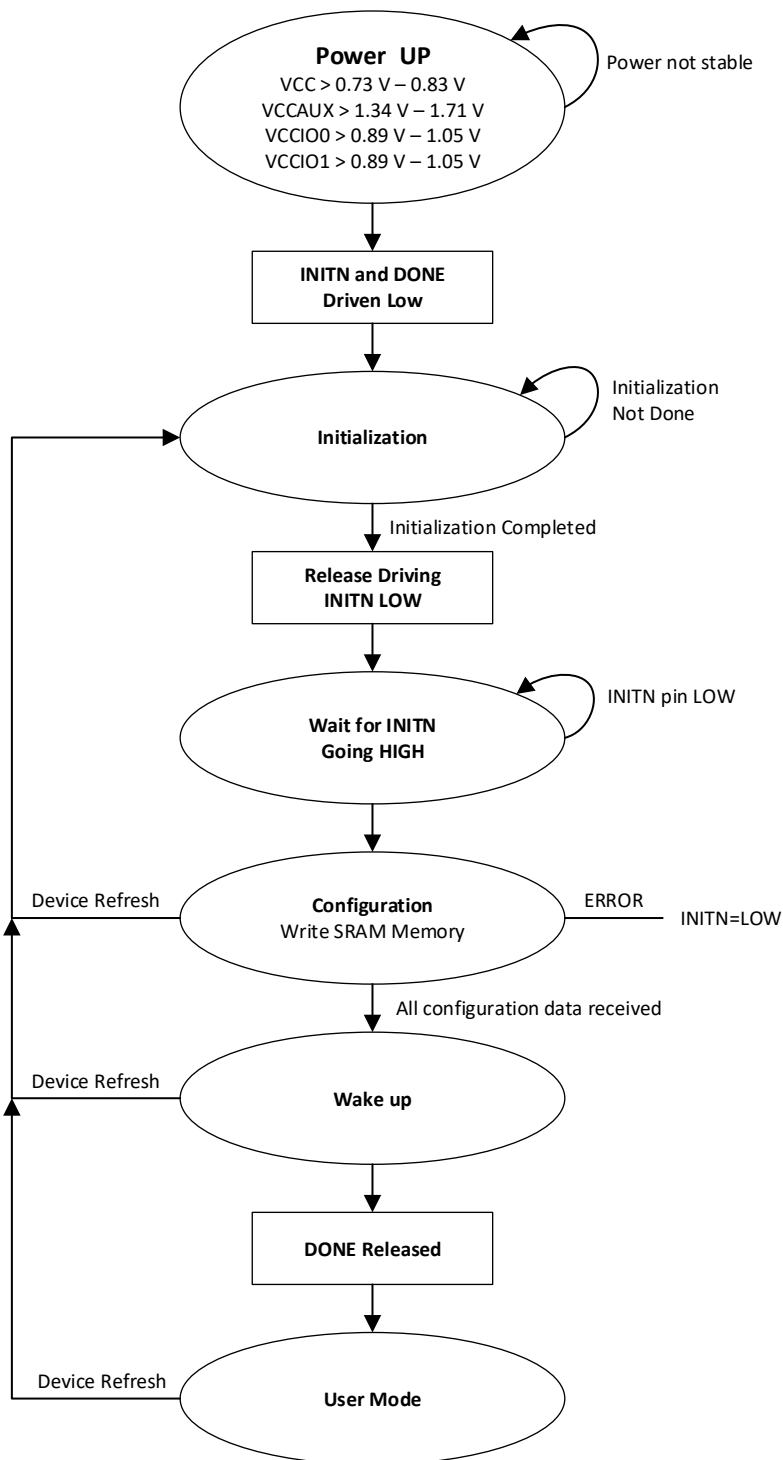


Figure 8.5. Debugging the FPGA Configuration Issues Flow

You can refer to the [sysCONFIG User Guide for Nexus Platform \(FPGA-TN-02099\)](#) document for more information on FPGA configuration.

8.1.2. Internal Register Read for Debug

If the above flowcharts do not capture the issues mentioned, you can read the 0x03114 (*vc_rx_status*) register, which indicates the Receive Buffer Parity/ECC Status where the error detection status can be seen.

The PCIe capability register addresses listed below also provide relevant information for debugging the PCIe issues:

- 0x47-0x44 (Device capability register) address for checking the different supported capabilities of the connected system.
- 0x49-0x48 (Device Control Register) address for checking the supported capabilities of the device.
- 0x4B-0x4A (Device Status Register) address for the device status. You can obtain the error status through this register address.

8.1.3. PCIe Loopback Test

The PCIe loopback test is a diagnostic feature specified by the PCIe Specification that can aid in debugging. The LTSSM Loopback is a state of the Link Training and Status State Machine (LTSSM), which is a mechanism for managing the link state of a serial bus such as PCI Express. In this state, the link partner can test its own transmitter and receiver by sending and receiving data packets without involving the link partner.

The LTSSM Loopback state can be entered from two different states: Configuration or Recovery. The entry into Loopback state is initiated by a Leader Loopback. Before register 0x2100 field is set to 1, all relevant registers containing Leader Loopback control options must be set to the desired values.

When *mgmt_tlb_debug_direct_to_loopback* = 1, no Leader Loopback control options may be changed. The LTSSM Loopback state has three substates: Entry, Active, and Exit. In Entry substate, both link partners wait for eight EIOS (End of Initialization Ordered Sets) before transitioning to Active substate. In Active substate, both link partners exchange data packets for testing purposes. In Exit substate, both link partners wait for eight EIOS before transitioning to Recovery or Configuration state depending on whether they received an Electrical Idle signal or not.

The LTSSM Loopback is useful for debugging and characterizing the performance of PCI Express links during Link Equalization Training which is a process of optimizing the signal quality between two link partners by adjusting various parameters such as amplitude, de-emphasis, preshoot, and jitter. For more information on the Loopback state, see [Table 2.5](#).

9. Design Consideration

9.1. DMA Based Design

DMA is not supported in 2024.1 release.

9.2. Non-DMA Based Design

To create a Non-DMA based design:

1. According to the PCIe IP configuration, select the appropriate clocking architecture. Refer to the [Clocking](#) section.
2. Select the proper data Interface type in IP according to the design requirement. Refer to the [General](#) section.
3. Initialize the register using LMMI interface/ AXI-L interface which are configure from the IP wizard. Refer to the [LMMI Interface](#) section.
4. Verify the TLP write and read Transactions. Refer to the [Transaction Layer Interface](#) section.
5. Verify the AXI-4 Stream write and read Transactions. See [AXI-4 Stream Data Interface](#) section.
6. Select the BAR's with BAR size according to the requirements. See [Base Address Register \(BAR\) \[0 to 5\]](#) section.

Appendix A. Resource Utilization

The Lattice PCIe IP core utilization report is provided in this section. You can check the resource utilized by the IP core and design top logic based on the available resource in the Avant-AT-G/X FPGA device.

[Table A.1](#) shows a sample resource utilization of the Lattice PCIe x8 IP Core on LAV-AT-G70-1LI with various link widths.

The resource utilization reports are generated with the Lattice Radiant tool version 2024.1 and with PCIe IP core version 1.1.0 generated for the following parameter settings. Note that x8 is generated from version 2.0.0. For x8 DMA variant, it is generated from version 2.1.0 using Lattice Radiant tool version 2024.2.

Table A.1. Lattice PCIe IP Core Resource Utilization

PCIe Core Config	Device Family	Map Resource Utilization								
		DMA				Non-DMA				
		LUT4	PFU Register	I/O Buffer	EBR	LUT4	PFU Register	I/O Buffer	EBR	Data Interface Type
1x1 EP	Avant-AT-G/X	—	—	—	—	5214	3143	5	37	AXI4_STREAM/ AXI-Lite
1x2 EP	Avant-AT-G/X	—	—	—	—	5522	3290	5	37	AXI4_STREAM/ AXI-Lite
1x4 EP	Avant-AT-G/X	—	—	—	—	6112	3568	5	37	AXI4_STREAM/ AXI-Lite
1x8 EP	Avant-AT-G/X	—	—	—	—	5433	2334	10	32	AXI4_STREAM/ AXI-Lite
1x1 EP	Avant-AT-G/X	—	—	—	—	1203	631	258	0	TLP/ LMMI
1x2 EP	Avant-AT-G/X	—	—	—	—	2285	966	35	0	TLP/ LMMI
1x4 EP	Avant-AT-G/X	—	—	—	—	2957	1266	35	0	TLP/ LMMI
1x8 EP	Avant-AT-G/X	—	—	—	—	1364	0	36	0	TLP/ LMMI
1x8 EP	Avant-AT-G/X	3414	0	12	0	—	—	—	—	AXI-MM/ LMMI

Note: Resource utilization differ with different configurations of the PCIe IP. The above resource utilization is provided for reference only. You can view the resource utilization under *Report > Map > Map Resource Usage*. To view the resource usage, you must run the Synthesize and Map Design.

References

- [PCIe x8 IP Release Notes \(FPGA-RN-02061\)](#)
- [Avant-G/X PCIe Host DMA Driver Software User Guide \(FPGA-TN-02405\)](#)
- [CertusPro-NX and Avant-G/X PCIe Basic Memory-Mapped Host Driver \(Non-DMA\) User Guide \(FPGA-TN-02387\)](#)
- [Avant-E web page](#)
- [Avant-G web page](#)
- [Avant-X web page](#)
- PCI Express Base Specification, Revision 4.0
- PCI Local Interface Specification Revision 3.0
- PCI Bus Power Management Interface Specification Revision 1.2

Other references:

- [Lattice Insights](#) for Lattice Semiconductor training courses and learning plans
- [Lattice Radiant](#) FPGA design software

Technical Support Assistance

Submit a technical support case through www.latticesemi.com/techsupport.

For frequently asked questions, refer to the Lattice Answer Database at www.latticesemi.com/Support/AnswerDatabase.

Revision History

Revision 1.3, IP v2.2.0, March 2025

Section	Change Summary
All	Updated IP version in the cover page.
Introduction	Updated the IP core and Radiant version in Table 1.1. Summary of the PCIe x8 IP .
Functional Description	Added DMA Performance and DMA Bypass Interface sections.
IP Parameter Description	<ul style="list-style-type: none"> Updated Figure 3.1. Attributes in the General Tab, Figure 3.2. Attributes in the Optional Port Tab, and Figure 3.10. Attributes in PCIe Capability. Updated Table 3.2. Optional Port Attributes to remove rows. Updated section name to PCI Express Capability.
Example Design	<ul style="list-style-type: none"> Updated Table 6.1. PCIe x8 IP Configuration Supported by the Example Design to remove <i>Use TLP Interface</i> row and update Non-DMA Design column values for <i>Bifurcation Selection</i>, <i>Data Interface Type</i>, and <i>Number of Physical Function</i>. Updated Overview of the Example Design and Features description. Updated Figure 6.2. Components within the Harden DMA Design, DMA Design section content, and OSC document reference under Non-DMA Design section in Example Design Components. Updated section name to Simulating the Example Design and moved Running Functional Simulation (Non-DMA) and Running Functional Simulation (DMA) under this section (previously from the Designing with the IP section); updated section content as well including adding additional steps and figures. Added Generating the Example Design for Hardware section for DMA and non-DMA designs. Updated Limitations of the Example Design section content.
Debugging	Updated Link Training Debug section content including adding Figure 8.2. Link Training Issue Debugging Flow and Table 8.1. PCIe LTSSM State and Sub-State Definition .
References	Added Avant PCIe Driver document references.

Revision 1.2, IP v2.1.0, December 2024

Section	Change Summary
All	Added IP version to the cover page and revision history.
Introduction	<ul style="list-style-type: none"> Updated Table 1.1. Summary of the PCIe x8 IP to change to All Avant-AT-G and Avant-AT-X family in Targeted Devices, added IP Changes row, and updated content as well, and updated IP core and Radiant software version. Updated SRIOV bullet item to add future release for non-DMA text and removed future release info for Hardened high-performance bullet point in Features section. Added Radiant 2024.1 SP1 information in Soft IP section. Updated part numbers in Table 1.2. Ordering Part Number. Renamed IP Validation Summary section to Hardware Support and updated section content. Added IP Support Summary Hardware Support and Speed Grade Supported section.
Functional Description	<ul style="list-style-type: none"> Added a note for bridge core support in the DMA Support section. Added a note for <code>usr_lmml_rst_n_i</code> once asserted in the LMML Interface section.
IP Parameter Description	Updated Figure 3.1. Attributes in the General Tab and Table 3.1. General Tab Attributes Description to add and remove rows and updated Selectable Values and Parameter columns.
Signal Description	<ul style="list-style-type: none"> Updated Table 4.7. Lattice Memory Mapped Interface Ports to add note for <code>usr_lmml_rst_n_i</code>. Updated Table 4.11. DMA AXI4 Manager Write Interface to update values in Clock Domain column and change <code>m0_axi4</code> to <code>m0_axi</code> across the table. Added Table 4.15. DMA AXI4 Interrupt Interface and Unused Interface section.

Section	Change Summary
Example Design	<ul style="list-style-type: none"> Updated DMA Design and Non-DMA Design columns in Table 6.1. PCIe x8 IP Configuration Supported by the Example Design. Updated DMA Design section content including adding Figure 6.2. Components within the Harden DMA Design in Example Design Components. Updated DMA Design section content in Design Test Case Examples. Updated Limitations of the Example Design section content including adding non-DMA and Harden DMA info.
Designing with the IP	<ul style="list-style-type: none"> Updated Generating and Instantiating the IP section to add non-DMA and Harden DMA info, including updating steps and adding related figures. Updated section name to Running Functional Simulation (Non-DMA) and section content including updating steps and related figures. Added Running Functional Simulation (DMA) section.
Appendix A. Resource Utilization	Moved section to Appendix, updated section content including updating Table A.1. Lattice PCIe IP Core Resource Utilization to add 1x8 EP for AXI-MM/LMMI row.
References	Added document reference for PCIe x8 release notes.

Revision 1.1, June 2024

Section	Change Summary
Acronyms in This Document	Rearranged some acronyms in alphabetical order.
Introduction	<ul style="list-style-type: none"> Updated IP core version to 1.1.0 in Table 1.1. Summary of the PCIe x8 IP. Moved Resource Utilization from Appendix A to this section, updated IP core version to 1.1.0, added note that x8 is generated from version 2.0.0, and updated Table 1.4. Lattice PCIe IP Core Resource Utilization to add rows for 1x8 EP AXI4 Stream/AXI-Lite and 1x8 EP TLP/LMMI; updated row values for 1x2 and 1x4 EP TLP/LMMI, including removing DMA EBR values. Removed future release text in some bullet points and added Future Radiant releases in Radiant 2024.1 bullet point in Features section. Removed Interrupt ports of Hard DMA Core and Linkup Status ports bullet points and added DMA is supported only in x8 mode bullet point in Hard IP Limitations section.

Section	Change Summary
Functional Description	<ul style="list-style-type: none"> Updated the following in PCIe IP Architecture Overview: <ul style="list-style-type: none"> Removed soft logic information in the first paragraph. Adjusted bullet points in this section. Changed link0_clk_i to sys_clk_i in the Clock Interface bullet point. Removed future release text in <i>Available if Hard DMA Core is enabled</i>. Updated APSM L0s section to add link to Table 5.12. Itssm_nfts Register 0x50 for mgmt_tlb_itssm_nfts_to_extend. Updated entire DMA Support section content, including additional sections and figures. Updated Figure 2.12. Non-DMA Write Operation (TLP Data Interface). Updated the following in Hard IP Interface: <ul style="list-style-type: none"> Updated Figure 2.20. TLP Memory Read Operation for Link0 (x4 Lane), Figure 2.21. TLP Memory Read Operation for Link0 (x2 Lane), and Figure 2.22. TLP Memory Read Operation for Link0 (x1 Lane). Updated the following in LMMI Interface: <ul style="list-style-type: none"> Updated text to: <i>The data transaction, through the LMMI, only starts when both usr_lmml_request_i and usr_lmml_ready_o are asserted high and When both usr_lmml_request_i and usr_lmml_ready_o are asserted high, usr_lmml_wr_rdn_i, and usr_lmml_offset_i must be valid and describe the transaction to execute; if the transaction is a write as indicated by usr_lmml_wr_rdn_i being asserted to high, usr_lmml_wdata_i must also be valid.</i> Updated section information for LMMI Write Operations and LMMI Read Operation, including Figure 2.31. LMMI Write Operation and Figure 2.32. LMMI Read Operation. Updated Figure 2.34. PCIe to AXI4-Stream Transaction for x1, Figure 2.35. PCIe to AXI4-Stream Transaction for x2, Figure 2.36. PCIe to AXI4-Stream Transaction for x8, Figure 2.37. AXI4-Stream to PCIe Transaction for x1, Figure 2.38. AXI4-Stream to PCIe Transaction for x2, and Figure 2.39. AXI4-Stream to PCIe Transaction for x8. Updated Multi-Protocol Support section content, including adding Figure 2.40.
IP Parameter Description	Removed table note in Table 3.16. Function 1-3 Tab.
Signal Description	<ul style="list-style-type: none"> Updated the following in Table 4.1. Clock and Reset Ports: <ul style="list-style-type: none"> Updated link0_perst_n_i description field to change link0_rst_n_i to link0_rst_usr_n_i. Updated link0_rst_n_i port and description fields to change link0_rst_n_i to link0_rst_usr_n_i. Changed link0_clk_i port to sys_clk_i. Changed link0_clk_src_o to link0_clk_usr_o port and link0_clk_i to sys_clk_i in the description field. Updated port and description fields for link0_pl_up_o, link0_dl_up_o, and link0_tl_up_o to add link to the port name. Removed sys_clk_i row. Updated port and description fields content in Table 4.2. PHY Interface Descriptions, including removing LINK – range [0, 3] in table note. Updated Table 4.4. TLP Transmit Credit Interface Ports to remove table notes and note in link0_tx_credit_nh_o[11:0]. Updated Table 4.6. TLP Receive Credit Interface Ports to remove table notes and note in link0_rx_credit_nh_i[11:0] and link0_rx_credit_nh_inf_i. Updated Table 4.7. Lattice Memory Mapped Interface Ports to update port and description fields for usr_lmml_offset_i. Updated note in Power Management Interface to add information that it is supported in 2024.1 and later. Updated axi4l_awaddr_i and axi4l_araddr_i to change [19:0] to [23:0] in Table 4.8. AXI-L Configuration Interface Ports.

Section	Change Summary
	<ul style="list-style-type: none"> Removed Soft DMA Core information in AXI-4 Stream Data Interface. Updated section content and name from DMA Interrupt Interface to DMA Interface, including adding Table 4.11. DMA AXI4 Manager Write Interface to Table 4.14. DMA AXI4 Subordinate Read Interface.
Register Description	<ul style="list-style-type: none"> Updated Table 5.2. Hard PCIe Core Register Mapping to update Start Byte Offset column values for Quad 0 MPPHY; updated Start and End Byte Offset columns for Quad 1 MPPHY. Added DMA Configuration Space Registers section.
Example Design	<ul style="list-style-type: none"> Updated Table 6.1. PCIe x8 IP Configuration Supported by the Example Design to DMA Design column values to Not supported in 2024.1 release. Updated DMA Design Example not supported information to 2024.1 release in Example Design Components and Design Test Case Examples sections. Updated bullet point in Limitations of the Example Design to include ModelSim 2024.1 not supported information.
Designing with the IP	Updated Design Implementation and Timing Constraints sections to update Lattice Radiant Software user guide version to 2024.1.
Appendix A. Resource Utilization	Updated Table 1.4. Lattice PCIe IP Core Resource Utilization to add rows for 1x8 EP AXI4 Stream/AXI-Lite and 1x8 EP TLP/LMMI; updated row values for 1x2 and 1x4 EP TLP/LMMI.

Revision 1.0, December 2023

Section	Change Summary
All	<ul style="list-style-type: none"> Changed document status to Production release. Changed all lnk signal references to <i>link</i> across the document.
Introduction	<ul style="list-style-type: none"> Updated Table 1.1. Summary of the PCIe x8 IP to change IP Core and Radiant software version and remove ModelSim information. Updated the following in Features section: <ul style="list-style-type: none"> Added information that some features are available in future release. Updated supported lane configurations to <i>Radiant 2023.2</i> and <i>future release</i>. Changed PCIe x4 core support to <i>1-8 PF</i>. Removed Soft IP features. Updated Hard IP Limitations to remove bullet item for burst or consecutive access for PCIe config. Updated Soft IP to remove Data Interfaces bullet item. Updated Table 1.2. Ordering Part Number to add OPNs. Merged Hardware Evaluation content with IP Validation Summary section. Moved Ordering Part Number under Licensing and Ordering Information section. Updated Naming Conventions to remove [LINK] index bullet item.
Functional Description	<ul style="list-style-type: none"> Updated PCIe IP Architecture Overview to remove Link 1 to 3 bullet items for link layer cores. Updated Figure 2.2. Lattice 8-lane SERDES/PCS + PCIe Hard-IP to change information that Quad 1 and DMA Core are available in future release. Updated AXI interface to AXI-S Rx and AXI-S Tx in Figure 2.3. Lattice PCIe x8 Core Hard IP. Updated Clocking to remove clk_usr_ps90_i information. Updated Reset to remove c_apb_preset_n_i and change reset signal from usr_lmmi_resetrn_i to usr_cfg_reset_n_i. Updated Table 2.9. Base Address to Enable Interrupt and Table 2.10. Legacy Interrupt Register to change the base address. Removed Register Address and Lane Register Address in LMMI Read Operation. Updated AXI-L Interface to remove AHB-Lite and Register Address information and other content related to PCIe AXI-L, including diagrams. Removed AXI-S Interface, Mapping of the PCIe Config Register to AXI-S Subordinate Interface, and DMA Enabled sections.

Section	Change Summary
IP Parameter Description	<ul style="list-style-type: none"> Updated the following in Table 3.1. General Tab Attributes Description: <ul style="list-style-type: none"> Added 16.0G selectable values for Link 0 Target Link Speed. Changed selectable value for Link 0 Number of physical functions to 1-8 and parameter to <i>LINK0_NUM_FUNCTIONS</i> = {1..8}. Added information for Enable DMA Support that attribute is not supported in 2023.2 release. Removed Merge Register and Subordinate Interface row. Added table note in Table 3.16. Function 1-3 Tab. Removed RX TLP Destination Base Address section.
Signal Description	<ul style="list-style-type: none"> Merged content of Clock Interface and Reset Interface, including tables, and moved to Clock and Reset Interface section. Added note in Table 4.1. Clock and Reset Ports that 250 MHz is supported in 2024 release. Updated content of PHY Interface, Power Management Interface, and AXI-L Configuration Interface section. Updated <i>lnk0_tx_data_i</i> and <i>lnk0_tx_data_p_i</i> in Table 4.3. TLP Transmit Interface Ports. Updated content of Table 4.5. TLP Receive Interface Ports and Table 4.7. Lattice Memory Mapped Interface Ports. Updated description of <i>link0_rx_credit_nh_i</i>[11:0] and <i>link0_rx_credit_nh_inf_i</i> and added note 2 in Table 4.6. TLP Receive Credit Interface Ports. Updated AXI-4 Stream Data Interface section and the following items in Table 4.9. AXI-4 Stream Transmitter Interface Ports: <ul style="list-style-type: none"> Changed <i>m0_tready_i</i> to <i>m0_tready_i</i>. Changed <i>m0_tvalid_o</i> to <i>m0_tvalid_o</i>. Updated 1x2 and 1x1 values for <i>m0_tstrb_o</i> [NUM_LANES*8-1:0] port. Updated 1x2 and 1x1 values for <i>m0_tkeep_o</i> [NUM_LANES*8-1:0] port. Added note for <i>m0_tdest_o</i>[3:0]. Updated port names from <i>m[LINK]_tlast_o</i> and <i>m[LINK]_tid_o</i> [7:0] to <i>m0_tlast_o</i> and <i>m0_tid_o</i> [7:0]. Updated ports, including updating 1x2 and 1x1 values and added note in <i>s0_tdata_i</i>[NUM_LANES*64-1:0] in Table 4.10. AXI-4 Stream Receiver Interface Ports. Removed Legacy Interrupt Interface section.
Register Description	<ul style="list-style-type: none"> Updated content including table name in Table 5.2. Hard PCIe Core Register Mapping. Updated Table 5.3. CSR Values Recommended for EP Applications to add 0x4 in Offset values. Added note and changed section name to <i>mgmt_tlb</i> (0x4_2000) section. Updated Table 5.7. <i>ltssm_cfg</i> Register 0x3c to add that 8G and 16G for [11:8] <i>lw_start_updn_rate_en</i> is available for future release. Added note and reference to table note in the following tables: <ul style="list-style-type: none"> Table 5.13. <i>ltssm_ds_initial_auto</i> Register 0x54 Table 5.18. <i>ltssm_compliance_toggle</i> Register 0x68 Table 5.20. <i>ltssm_link</i> Register 0x80Table 5.44. <i>debug_loopback_control</i> Register 0x104 Added note and reference to note for Target Link Speed in <i>eq_status_table_control</i> Register 0x3d8. Changed section name to <i>mgmt_ptl</i> (0x4_3000) and <i>mgmt_ftl</i> (0x4_4000). Updated <i>mgmt_ftl_mf</i>[7:1] (0x4_5000,0x4_6000,0x4_7000,0x4_8000,0x4_9000,0x4_A000,0x4_B000) to change section name and update content of Table 5.187. Base Address for <i>mgmt_ftl_mf</i>. Updated Table 5.188. Function Register 0x8 to change the values of the Description column for [31:1] and [0] fields. Updated Table 5.212. PCI Express Capability to add 1001 and 1010 in 43-42 Addr is for future release and added 0100 (16GT/s) value for 4F-4C addr.

Section	Change Summary
	<ul style="list-style-type: none"> Updated Table 5.218. Vendor-Specific Extended Capability to add that 8G and 16G in 163-160 addr are for future release. Added ATS Extended Capability and Resizable BAR Capability sections. Removed pcie_11(0x0F000) and Soft IP Configuration, Control, and Status Registers sections.
Example Design	<ul style="list-style-type: none"> Updated DMA Design in the section to add that support is available in next release. Updated DMA Design and Non-DMA Design support and added table note 2 in Table 6.1. PCIe x8 IP Configuration Supported by the Example Design. Added note and reference to note and changed AXI/AHB-Lite to TLP in Overview of the Example Design and Features. Updated OSC Module document link and title in Example Design Components. Updated content in Limitations of the Example Design section.
Designing with the IP	<ul style="list-style-type: none"> Updated Figure 7.1. Module/IP Block Wizard, Figure 7.2. IP Configuration, and Figure 7.3. Check Generated Result to change to pcie_x8. Updated content of Running Functional Simulation section.
Design Consideration	Added this section.
Appendix A. Resource Utilization	Moved Resource Utilization from Introduction to Appendix section and updated Table A.1. Lattice PCIe IP Core Resource Utilization.
References	Updated Avant webpage links.

Revision 0.80, October 2023

Section	Change Summary
All	Initial preliminary release.



www.latticesemi.com