Virtex-7 FPGA Gen3 Integrated Block for PCI Express v1.3

Product Guide

PG023 October 16, 2012





Table of Contents

SECTION I: SUMMARY

IP Facts

| A I. | | 4 | \sim | • |
|-------------|------|------------|-----------------|---------|
| (na | ntor | | ()\/ O r | \/_\\/ |
| CIIIa | μιςι | - . | Over | |

| Feature Summary | 10 |
|------------------------------------|----|
| Applications | 10 |
| Unsupported Features | 11 |
| Licensing and Ordering Information | 11 |

Chapter 2: Product Specification

| Standards Compliance | 12 |
|------------------------|----|
| Resource Utilization | 12 |
| Block Selection | 13 |
| Port Descriptions | 14 |
| Attribute Descriptions | 63 |
| Configuration Space | 67 |

Chapter 3: Designing with the Core

| General Design Guidelines | 74 |
|--|-----|
| Clocking | 74 |
| Resets | 77 |
| AXI4-Stream Interface Description | 78 |
| Clocking Requirements | 84 |
| Interface Operation | 85 |
| Power Management | 154 |
| Link Training: 2-Lane, 4-Lane, and 8-Lane Components 1 | L57 |
| Lane Reversal | L58 |

SECTION II: VIVADO DESIGN SUITE

XILINX.

| Chapter 4: Customizing and Generating the Core | |
|--|-----|
| Graphical User Interface (GUI) | 16 |
| Output Generation | 184 |

Chapter 5: Constraining the Core

| Required Constraints | 189 |
|---|-----|
| Device, Package, and Speed Grade Selections | 189 |
| Clock Frequencies | 190 |
| Clock Management | 190 |
| Clock Placement | 190 |
| Stacked Silicon Interconnect Devices | 190 |
| Transceiver Placement | 190 |
| I/O Standard and Placement | 195 |

Chapter 6: Detailed Example Design

| Directory and File Contents | 196 |
|-----------------------------|-----|
| Example Design | 196 |
| Demonstration Test Bench | 208 |
| Implementation | 224 |
| Simulation | 225 |

SECTION III: ISE DESIGN SUITE

| Chapter 7: Customizing and Generating the Core | |
|--|-----|
| GUI | 228 |
| Output Generation. | 254 |
| Chapter 8: Constraining the Core | |
| Required Constraints | 261 |
| Device, Package, and Speed Grade Selections | 261 |
| Clock Frequencies | 262 |
| Clock Management | 262 |
| Clock Placement | 262 |
| Stacked Silicon Interconnect Devices | 262 |
| Transceiver Placement | 262 |
| I/O Standard and Placement | 267 |
| | |

Chapter 9: Example Design and Model Test Bench for Endpoint Configuration

| Directory and File Contents | | 68 |
|-----------------------------|--|----|
|-----------------------------|--|----|

| Example Design | 268 |
|--------------------------|-----|
| Demonstration Test Bench | 280 |
| Implementation | 297 |
| Simulation | 298 |

Chapter 10: Example Design and Model Test Bench for Root Port Configuration

| Configurator Example Design | 301 |
|---|-----|
| Endpoint Model Test Bench for Root Port | 307 |

SECTION IV: APPENDICES

| Appendix A: Mi | grating |
|----------------|---------|
|----------------|---------|

Appendix B: Receiver Buffer Completion Space

Appendix C: Debugging

| Finding Help on Xilinx.com | 313 |
|-------------------------------------|-----|
| Contacting Xilinx Technical Support | 314 |
| Debug Tools | 315 |
| Hardware Debug | 318 |
| FPGA Configuration | 328 |
| Simulation Debug | 335 |

Appendix D: Attributes

Appendix E: Additional Resources

| Xilinx Resources | 381 |
|----------------------|-----|
| Solution Centers | 381 |
| References | 381 |
| Technical Support | 382 |
| Revision History | 382 |
| Notice of Disclaimer | 383 |



SECTION I: SUMMARY

IP Facts

Overview

Product Specification

Designing with the Core

IP Facts

Introduction

The LogiCORE[™] IP Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® core is a high-bandwidth, scalable, and reliable serial interconnect building block solution for use with all Virtex-7 XT and HT FPGAs except the XC7VX485T. The Integrated Block for PCI Express (PCIe®) solution supports 1-lane, 2-lane, 4-lane, and 8-lane Endpoint configurations, including Gen1 (2.5 GT/s), Gen2 (5.0 GT/s) and Gen3 (8 GT/s) speeds. It is compliant with PCI Express Base Specification, rev. 3.0. This solution supports the AXI4-Stream interface for the customer user interface.

PCI Express offers a serial architecture that alleviates many limitations of parallel bus architectures by using clock data recovery (CDR) and differential signaling. Using CDR (as opposed to source synchronous clocking) lowers pin count, enables superior frequency scalability, and makes data synchronization easier. PCI Express technology, adopted by the PCI-SIG® as the next generation PCI[™], is backward-compatible to the existing PCI software model.

With higher bandwidth per pin, low overhead, low latency, reduced signal integrity issues, and CDR architecture, the Integrated Block sets the industry standard for a high-performance, cost-efficient PCIe solution.

The Gen3 Integrated Block for PCIe solution is compatible with industry-standard application form factors such as the PCI Express Card Electromechanical (CEM) v3.0 and the PCI Industrial Computer Manufacturers Group (PICMG) 3.4 specifications.

| LogiCORE IP Facts Table | | | | |
|--|---|--|--|--|
| | Core Specifics | | | |
| Supported Device Family ⁽¹⁾ | Virtex-7 XT and HT ⁽²⁾ | | | |
| Supported User Interfaces | AXI4-Stream | | | |
| Resources | | | | |
| | Provided with Core | | | |
| Design Files | Verilog | | | |
| Example Design | Verilog | | | |
| Test Bench | Verilog | | | |
| Constraints File | ISE® Design Suite: UCF Vivado™ Design Suite: XDC | | | |
| Simulation Model | Verilog | | | |
| Supported S/W Drivers | N/A | | | |
| | Tested Design Flows ⁽³⁾ | | | |
| Design Entry | ISE Design Suite v14.3 Vivado Design Suite v2012.3 ⁽⁴⁾ | | | |
| Simulation | Mentor Graphics ModelSim Cadence Incisive Enterprise Simulator (IES) Synopsys VCS Xilinx ISim (for CORE Generator) Vivado Simulator | | | |
| Synthesis | Xilinx Synthesis Technology (XST) Vivado Synthesis | | | |
| Support | | | | |
| Provided by Xilinx @ <u>www.xilinx.com/support</u> | | | | |

Notes:

- 1. For a complete listing of supported devices, see the <u>release</u> <u>notes</u> for this core.
- 2. Except for the XC7VX485T.
- 3. For the supported versions of the tools, see the <u>Xilinx Design</u> <u>Tools: Release Notes Guide</u>.
- 4. Supports only 7 series devices.

EXILINX.

Features

The key features of the Virtex-7 FPGA Gen3 Integrated Block for PCI Express (8.0 GT/s) core are:

- High-performance, highly flexible, scalable, and reliable, general-purpose I/O core
 - Compliant with the PCI Express Base Specification, rev. 3.0
 - Compatible with conventional PCI software model
- GTH transceivers
 - 2.5 GT/s, 5.0 GT/s, and 8.0 GT/s line speeds
 - 1-lane, 2-lane, 4-lane, and 8-lane operation
- Endpoint configuration
- Multiple Function and Single-Root I/O Virtualization in the Endpoint configuration
 - 2 Physical Functions
 - 6 Virtual Functions
- Standardized user interface(s)
 - Compliant to AXI4-Stream
 - Separate Requester, Completion, and Message interfaces
 - Flexible Data Alignment
 - Parity generation and checking on AXI4-Stream interfaces
 - Easy-to-use packet-based protocol
 - Full-duplex communication enabling
 - Optional back-to-back transactions to enable greater link bandwidth utilization
 - Support for flow control of data and discontinuation of an in-process transaction in transmit direction
 - Support for flow control of data in receive direction
- Compliant with PCI and PCI Express power management functions
- Optional Tag Management feature
- Maximum transaction payload of up to 1024 bytes
- End-to-End Cyclic Redundancy Check (ECRC)
- Advanced Error Reporting (AER)
- Multi-Vector MSI for up to 32 vectors and MSI-X
- Atomic Operations and TLP Processing Hints



Chapter 1

Overview

The LogiCORE[™] IP Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® core is a reliable, high-bandwidth, scalable serial interconnect building block for use with Virtex-7 XT and HT FPGAs, except for the XC7VX485T. The core instantiates the Integrated Block found in Virtex-7 XT and HT FPGAs.

The Virtex-7 FPGA Gen3 Integrated Block for PCI Express is an IP core available with:

- the CORE Generator[™] tool in the ISE[®] Design Suite
- the IP Catalog in the Vivado[™] Design Suite

For detailed information about the core, see the <u>Virtex-7 FPGA Gen3 Integrated Block for</u> <u>PCI Express product page</u>.

Figure 1-1 shows the interfaces for the LogiCORE IP Virtex-7 FPGA Gen3 Integrated Block for PCI Express core.

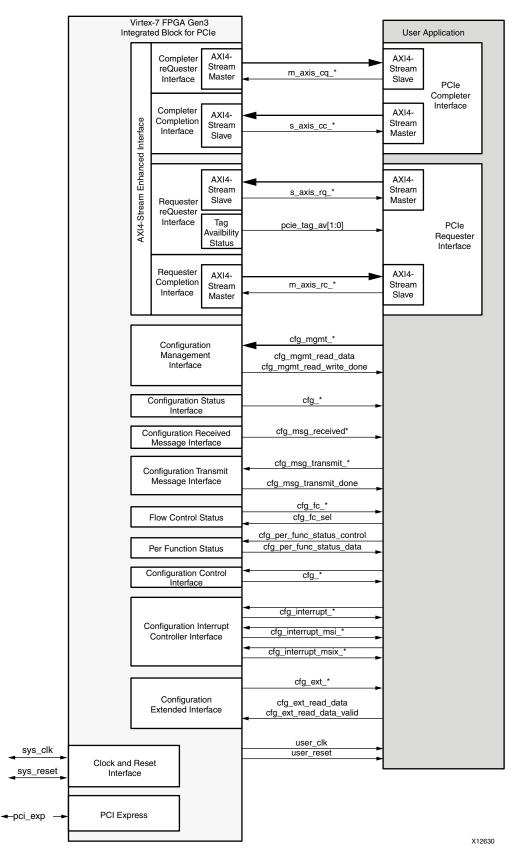


Figure 1-1: Virtex-7 FPGA Gen3 Integrated Block for PCI Express Interfaces

Feature Summary

The LogiCORE IP Virtex-7 FPGA Gen3 Integrated Block for PCI Express core is a high-bandwidth, scalable, and flexible general-purpose I/O core for use with most Virtex-7 XT and HT FPGAs. The GTH transceivers in the Integrated Block for PCI Express (PCIe®) solution support 1-lane, 2-lane, 4-lane, and 8-lane operation, running at 2.5 GT/s (Gen1), 5.0 GT/s (Gen2), and 8.0 GT/s (Gen3) line speeds. Endpoint configurations are supported.

The customer user interface is compliant with the AMBA® AXI4-Stream interface. This interface supports separate Requester, Completion, and Message interfaces. It allows for flexible data alignment and parity checking. Flow control of data is supported in the receive and transmit directions. The transmit direction additionally supports discontinuation of in-progress transactions. Optional back-to-back transactions utilize straddling to provide greater link bandwidth.

The Virtex-7 FPGA supports two physical functions and six virtual functions.

The core is compliant with PCI and PCI Express power management functions.

Note: For information about the availability of Root Port functionality, contact your Xilinx. See Technical Support, page 382.

Applications

The Integrated Block for PCI Express architecture enables a broad range of computing and communications target applications, emphasizing performance, cost, scalability, feature extensibility and mission-critical reliability. Typical applications include:

- Data communications networks
- Telecommunications networks
- Broadband wired and wireless applications
- Network interface cards
- Chip-to-chip and backplane interface cards
- Server add-in cards for various applications

Unsupported Features

The Integrated Block does not implement the Address Translation Service, but allows its implementation in external soft logic.

Switch ports and the Resizable BAR Extended Capability are not supported.

Licensing and Ordering Information

The LogiCORE IP Virtex-7 FPGA Gen3 Integrated Block for PCI Express core is provided at no additional cost with the Xilinx Vivado Design Suite and ISE Design Suite tools under the terms of the <u>Xilinx End User License</u>. Information about this and other Xilinx LogiCORE IP modules is available at the <u>Xilinx Intellectual Property</u> page. For information about pricing and availability of other Xilinx LogiCORE IP modules and tools, contact your <u>local Xilinx sales representative</u>.



Product Specification

Standards Compliance

The Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® solution is compatible with industry-standard application form factors such as the PCI Express Card Electromechanical (CEM) v3.0 and the PCI Industrial Computer Manufacturers Group (PICMG) 3.4 specifications.

Resource Utilization

Resources required for the Gen3 Integrated Block for PCIe® core have been estimated for the Virtex-7 FPGA (Table 2-1). These values were generated using the Xilinx® CORE Generator[™] tools, v14.3. They are derived from post-synthesis reports, and might change during MAP and PAR. The resources listed in Table 2-1 are for the default core configuration.

| Lanaa | GTHE2 | FF(1) | ··· · (1) | CMPS ⁽²⁾ | RX Completion | RX Request Buffer Size | TX Replay Buffer Size | Block RA | M Usage |
|-------|-------|--------|----------------------|---------------------|----------------------|---------------------------|--------------------------|----------|---------|
| Lanes | GINEZ | FF' / | | CIVIPS | Buffer Size (KB) | (KB) | | RAMB18 | RAMB36 |
| 1 | 1 | 566 | 832 | | 8 | | | 8 | |
| 1 | 1 | 200 | 052 | | 16 | | | 12 | |
| 2 | 2 | 957 | 1384 2500 | 128- | 8 | | | 8 | |
| 2 | 2 | | | | | 16 | 8 | 8 | 12 |
| 4 | 4 | 1740 2 | | 1024 | 8 | 0 | 0 | 8 | 5 |
| 4 | | | 1/40 | 2500 | - | 16 | | | 12 |
| 8 | 8 | 3399 | 4818 | - | 8 | | | 8 | |
| 0 | 0 | 5555 | 4010 | | 16 | | | 12 | |

| Table 2-1: | Gen3 Integrated Block for PCIe Resource Estimates |
|------------|---|
|------------|---|

Notes:

- 1. Numbers are for the default core configuration. Actual LUT and FF utilization values vary based on specific configurations.
- 2. Capability Maximum Payload Size (CMPS).

Block Selection

Table 2-2 lists which Gen3 Integrated Blocks for PCIe are available for use in FPGAs containing multiple Integrated Blocks. In some cases, not all Integrated Blocks can be used due to lack of bonded transceiver sites adjacent to the Integrated Block.

| Device Selection | | Inte | grated Block for | PCI Express Loca | ition |
|------------------|-------------------------------|------|------------------|------------------|-------|
| Device | Package | X0Y0 | X0Y1 | X0Y2 | X0Y3 |
| XC7VX330T | FFG1157 FFG1761 | Yes | Yes | | |
| XC7VX415T | FFG1157 FFG1158 FFG1927 | Yes | Yes | | |
| VCTVVEFOT | FFG1158 | | Yes | Yes | |
| XC7VX550T | FFG1927 | Yes | Yes | Yes | |
| | FFG1157 FFG1158 FFG1930 | | Yes | Yes | |
| XC7VX690T | FFG1761 FFG1926 FFG1927 | Yes | Yes | Yes | |
| XC7VX980T | FFG1926 FFG1928 FFG1933 | Yes | Yes | Yes | |
| | FFG1930 | | Yes | Yes | |
| | FLG1926 FLG1933 | Yes | Yes | Yes | |
| XC7VX1140T | FLG1928 | Yes | Yes | Yes | Yes |
| | FLG1930 | | Yes | Yes | |
| ХС7VН580Т | HCG1155 HCG1931 | Yes | | | |
| | HCG1932 | Yes | Yes | | |
| XC7VH870T | HCG1931 | | Yes | Yes | |
| ΛC/ΥΠΟ/ΟΙ | HCG1932 | Yes | Yes | Yes | |

| Table 2-2: | Available Integrated Blocks for PCI Express |
|------------|---|
|------------|---|

Port Descriptions

This section provides detailed port descriptions for each interface.

Core Interfaces

In addition to status and control interfaces, the core has these four required AXI4-Stream interfaces used to transfer and receive transactions:

- The Completer reQuest (CQ) interface through which all received requests from the link are delivered to the user application.
- The Completer Completion (CC) interface through which completions generated by the user application responses to the completer requests are transmitted. The user can process all Non-Posted transactions as split transactions. That is, it can continue to accept new requests on the Requester Completion interface while sending a completion for a request.
- The Requester reQuest (RQ) interface through which the user application generates requests to remote PCIe® devices.
- The Requester Completion (RC) interface through which the completions received from the link in response to the user's requests are presented to the user application.

The core also provides these interfaces:

- The Transmit Flow Control interface is used by the user application to request which flow control information the core provides. This interface provides the Posted/ Non-Posted Header Flow Control Credits, Posted/Non-Posted Data Flow Control Credits, the Completion Header Flow Control Credits, and the Completion Data Flow Control Credits to the user application based upon the setting flow control select input to the core.
- The Configuration Management interface is used to read and write to the Configuration Space Registers.
- The Configuration Status interface provides information on how the core is configured, such as the negotiated link width and speed, the power state of the core, and configuration errors.
- The Configuration Received Message interface indicates to the user logic that a decodable message from the link, the parameters associated with the data, and the type of message have been received.
- The Configuration Transmit Message interface is used by the user application to transmit messages to the core. The user application supplies the transmit message type and data information to the core, which responds with the Done signal.

- The Per Function Status interface provides status data as requested by the user application through the selected function.
- The Configuration Control Interface signals allow a broad range of information exchange between the user application and the core. The user application uses this interface to set the configuration space; indicate if a correctable or uncorrectable error has occurred; set the device serial number; set the Downstream Bus, Device, and Function Number; and receive per function configuration information. This interface also provides handshaking between the user application and the core when a Power State change or function level reset occurs.
- The Configuration Interrupt Controller interface allows the user application to set Legacy PCIe interrupts, MSI interrupts, or MSI-X interrupts. The core provides the interrupt status on the configuration interrupt sent and fail signals.
- The Configuration Extended interface allows the core to transfer configuration information with the user application when externally implemented configuration registers are implemented.
- The Clock and Reset interface, fundamental to the operation of the core, provides the system level clock and reset to the core as well as the user application clock and reset signal.
- The PCI Express interface contains all of the differential transmit and receive pairs.

Completer reQuest (CQ) Interface

Table 2-3 defines the ports in the CQ interface of the Gen3 Integrated Block for PCI Express core. In the Width column, DW denotes the configured data bus width (64, 128, or 256 bits)

| Port | Direction | Description | Width |
|-----------------|-----------|--|-------|
| m_axis_cq_tdata | Output | Transmit Data from the Completer reQuest Interface. Only the lower 128 bits are to be used when the interface width is 128 bits, and only the lower 64 bits are to be used when the interface width is 64 bits. Bits [255:128] are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [255:64] are set permanently to 0 when the interface width is configured as 64 bits. | DW/32 |
| m_axis_cq_tuser | Output | Completer reQuest User Data. This set of signals contains sideband information for the TLP being transferred. These signals are valid when m_axis_cq_tvalid is High. Table 2-4, page 18 describes the individual signals in this set. | 85 |

| Table 2-3: CQ In | erface Port Description | IS |
|------------------|-------------------------|----|
|------------------|-------------------------|----|

| Port | Direction | Description | Width |
|------------------|-----------|--|-------|
| m_axis_cq_tlast | Output | TLAST indication for Completer reQuest Data. The core asserts this signal in the last beat of a packet to indicate the end of the packet. When a TLP is transferred in a single beat, the core sets this bit in the first beat of the transfer. | 1 |
| m_axis_cq_tkeep | Output | TKEEP indication for Completer reQuest Data. The assertion of bit <i>i</i> of this bus during a transfer indicates to the client that Dword <i>i</i> of the m_axis_cq_tdata bus contains valid data. The core sets this bit to 1 contiguously for all Dwords starting from the first Dword of the descriptor to the last Dword of the payload. Thus, m_axis_cq_tdata is set to all 1s in all beats of a packet, except in the final beat when the total size of the packet is not a multiple of the width of the data bus (both in Dwords). This is true for both Dword-aligned and address-aligned modes of payload transfer. Bits [7:4] of this bus are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [7:2] are set permanently to 0 when the interface width is configured as 64 bits. | DW/32 |
| m_axis_cq_tvalid | Output | Completer reQuest Data Valid. The core asserts this output whenever it is driving valid data on the m_axis_cq_tdata bus. The core keeps the valid signal asserted during the transfer of a packet. The client application can pace the data transfer using the m_axis_cq_tready signal. | 1 |
| m_axis_cq_tready | Input | Completer reQuest Data Ready. Activation of this signal by the client logic indicates to the core that the client is ready to accept data. Data is transferred across the interface when both m_axis_cq_tvalid and m_axis_cq_tready are asserted in the same cycle. If the client deasserts the ready signal when m_axis_cq_tvalid is High, the core maintains the data on the bus and keeps the valid signal asserted until the client has asserted the ready signal. | 1 |

Table 2-3: CQ Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|----------------------|-----------|---|-------|
| pcie_cq_np_req | Input | Completer reQuest Non-Posted Request. This input is used by the client application to request the delivery of a Non-Posted request. The core implements a credit-based flow control mechanism to control the delivery of Non-Posted requests across the interface, without blocking Posted TLPs. This input to the core controls an internal credit count. The credit count is incremented in each clock cycle when pcie_cq_np_req is High, and decremented on the delivery of each Non-Posted request across the interface. The core temporarily stops delivering Non-Posted requests to the client when the credit count is zero. It continues to deliver any Posted TLPs received from the link even when the delivery of Non-Posted requests has been paused. The client application can either provide a one-cycle pulse on pcie_cq_np_req each time it is ready to receive a Non-Posted request, or can keep it High permanently if it does not need to exercise selective backpressure on Non-Posted requests. The assertion and deassertion of the pcie_cq_np_req signal does not need to be aligned with the packet transfers on the completer request interface. | 1 |
| pcie_cq_np_req_count | Output | Completer reQuest Non-Posted Request Count. This output provides the current value of the credit count maintained by the core for delivery of Non-Posted requests to the client. The core delivers a Non-Posted request across the completer request interface only when this credit count is non-zero. This counter saturates at a maximum limit of 32. Because of internal pipeline delays, there can be several cycles of delay between the core receiving a pulse on the pcie_cq_np_req input and updating the pcie_cq_np_req_count output in response. | 6 |

Table 2-3: CQ Interface Port Descriptions (Cont'd)

| Bit Index | Name | Width | Description |
|-----------|---------------|-------|---|
| 3:0 | first_be[3:0] | 4 | Byte enables for the first Dword of the payload. This field reflects the setting of the First_BE bits in the Transaction-Layer header of the TLP. For Memory Reads and I/O Reads, these four bits indicate the valid bytes to be read in the first Dword. For Memory Writes and I/O Writes, these bits indicate the valid bytes in the first Dword of the payload. For Atomic Operations and Messages with a payload, these bits are set to all 1s. This field is valid in the first beat of a packet, that is, when sop and m_axis_cq_tvalid are both High. |
| 7:4 | last_be[3:0] | 4 | Byte enables for the last Dword. This field reflects the setting of the Last_BE bits in the Transaction-Layer header of the TLP. For Memory Reads, these four bits indicate the valid bytes to be read in the last Dword of the block of data. For Memory Writes, these bits indicate the valid bytes in the ending Dword of the payload. For Atomic Operations and Messages with a payload, these bits are set to all 1s. This field is valid in the first beat of a packet, that is, when sop and m_axis_cq_tvalid are both High. |
| 39:8 | byte_en[31:0] | 32 | The client logic can optionally use these byte enable bits to determine the valid bytes in the payload of a packet being transferred. The assertion of bit <i>i</i> of this bus during a transfer indicates to the client that byte <i>i</i> of the m_axis_cq_tdata bus contains a valid payload byte. This bit is not asserted for descriptor bytes. Although the byte enables can be generated by client logic from information in the request descriptor (address and length) as well as the settings of the first_be and last_be signals, the client has the option to use these signals directly instead of generating them from other interface signals. When the payload size is more than two Dwords (eight bytes), the one bit on this bus for the payload is always contiguous. When the payload size is two Dwords or less, the one bit can be non-contiguous. For the special case of a zero-length memory write transaction defined by the PCI Express specifications, the byte_en bits are all 0s when the associated one-DW payload is being transferred. Bits [31:16] of this bus are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [31:8] are set permanently to 0 when the interface width is configured as 64 bits. |
| 40 | sop | 1 | Start of packet. This signal is asserted by the core in the first beat of a packet to indicate the start of the packet. Use of this signal is optional for the client. |

Table 2-4: Sideband Signal Descriptions in m_axis_cq_tuser

| Bit Index | Name | Width | Description |
|-----------|-----------------|-------|--|
| | | | This signal is asserted by the core in the last beat of a TLP, if it has detected an uncorrectable error while reading the TLP payload from its internal FIFO memory. The client application must discard the entire TLP when such an error is signaled by the core. |
| 41 | discontinue | 1 | This signal is never asserted when the TLP has no payload. It is asserted only in a cycle when m_axis_cq_tlast is High. |
| | | | When the core is configured as an Endpoint, the error is also reported by the core to the Root Complex to which it is attached, using Advanced Error Reporting (AER). |
| 42 | tph_present | 1 | This bit indicates the presence of a Transaction Processing Hint (TPH) in the request TLP being delivered across the interface. This bit is valid when sop and m_axis_cq_tvalid are both High. |
| 44:43 | tph_type[1:0] | 2 | When a TPH is present in the request TLP, these two bits provide the value of the PH[1:0] field associated with the hint. These bits are valid when sop and m_axis_cq_tvalid are both High. |
| 52:45 | tph_st_tag[7:0] | 8 | When a TPH is present in the request TLP, this output provides the 8-bit Steering Tag associated with the hint. These bits are valid when sop and m_axis_cq_tvalid are both High. |
| 84:53 | parity | 32 | Odd parity for the 256-bit transmit data. Bit <i>i</i> provides the odd parity computed for byte <i>i</i> of $m_axis_cq_tdata$. Only the lower 16 bits are to be used when the interface width is 128 bits, and only the lower 8 bits are to be used when the interface width is 64 bits. Bits [31:16] are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [31:8] are set permanently to 0 when the interface width is configured as 64 bits. |

Completer Completion (CC) Interface

Table 2-5 defines the ports in the CC interface of the Gen 3 Integrated Block for PCI Express core. In the Width column, DW denotes the configured data bus width (64, 128, or 256 bits).

Table 2-5: CC Interface Port Descriptions

| Port | Direction | Description | Width |
|------------------|-----------|--|-------|
| s_axis_cc_tdata | Input | Completer Completion Data bus. Completion data from the client application to the core. Only the lower 128 bits are to be used when the interface width is 128 bits, and only the lower 64 bits are to be used when the interface width is 64 bits. | DW |
| s_axis_cc_tuser | Input | Completer Completion User Data. This set of signals contain sideband information for the TLP being transferred. These signals are valid when s_axis_cc_tvalid is High. Table 2-6, page 21 describes the individual signals in this set. | 33 |
| s_axis_cc_tlast | Input | TLAST indication for Completer Completion Data. The client application must assert this signal in the last cycle of a packet to indicate the end of the packet. When the TLP is transferred in a single beat, the client must set this bit in the first cycle of the transfer. | 1 |
| s_axis_cc_tkeep | Input | TKEEP indication for Completer Completion Data. The assertion of bit <i>i</i> of this bus during a transfer indicates to the core that Dword <i>i</i> of the s_axis_cc_tdata bus contains valid data. The client must set this bit to 1 contiguously for all Dwords starting from the first Dword of the descriptor to the last Dword of the payload. Thus, s_axis_cc_tdata must be set to all 1s in all beats of a packet, except in the final beat when the total size of the packet is not a multiple of the width of the data bus (both in Dwords). This is true for both Dword-aligned and address-aligned modes of payload transfer. Bits [7:4] of this bus are not used by the core when the interface width is configured as 128 bits, and bits [7:2] are not used when the interface width is configured as 64 bits. | DW/32 |
| s_axis_cc_tvalid | Input | Completer Completion Data Valid. The client application must assert this output whenever it is driving valid data on the s_axis_cc_tdata bus. The client must keep the valid signal asserted during the transfer of a packet. The core paces the data transfer using the s_axis_cc_tready signal. | 1 |
| s_axis_cc_tready | Output | Completer Completion Data Ready. Activation of this signal by the core indicates that it is ready to accept data. Data is transferred across the interface when both s_axis_cc_tvalid and s_axis_cc_tready are asserted in the same cycle. If the core deasserts the ready signal when the valid signal is High, the client must maintain the data on the bus and keep the valid signal asserted until the core has asserted the ready signal. | 1 |

| Bit Index | Name | Width | Description |
|-----------|-------------|-------|---|
| 0 | discontinue | 1 | This signal can be asserted by the client application during a transfer if it has detected an error (such as an uncorrectable ECC error while reading the payload from memory) in the data being transferred and desires to abort the packet. The core nullifies the corresponding TLP on the link to avoid data corruption. The client can assert this signal during any cycle during the transfer. It can either choose to terminate the packet prematurely in the cycle where the error was signaled, or can continue until all bytes of the payload are delivered to the core. In the latter case, the core treats the error as sticky for the following beats of the packet, even if the client deasserts the discontinue signal before the end of the packet. The discontinue signal can be asserted only when $s_axis_cc_tvalid$ is High. The core samples this signal only when $s_axis_cc_tready$ is High. Thus, when asserted, it should not be deasserted until $s_axis_cc_tready$ is High. When the core is configured as an Endpoint, this error is also reported by the core to the Root Complex to which it is attached, using AER. |
| 32:1 | parity | 32 | Odd parity for the 256-bit data. When parity checking is enabled in the core, client logic must set bit <i>i</i> of this bus to the odd parity computed for byte <i>i</i> of s_axis_cc_tdata. Only the lower 16 bits are to be used when the interface width is 128 bits, and only the lower 8 bits are to be used when the interface width is 64 bits. On detection of a parity error, the core nullifies the corresponding TLP on the link and reports it as an Uncorrectable Internal Error. The parity bits can be permanently tied to 0 if parity check is not enabled in the core. |

| | Table 2-6: | Sideband Signal Descr | ptions in s | axis cc tuser |
|--|------------|-----------------------|-------------|---------------|
|--|------------|-----------------------|-------------|---------------|

Requester reQuest (RQ) Interface

Table 2-7 defines the ports in the RQ interface of the Gen3 Integrated Block for PCI Express core. In the Width column, DW denotes the configured data bus width (64, 128, or 256 bits).

| Table 2-7: | RQ Interface Port Descriptions |
|------------|---------------------------------------|
|------------|---------------------------------------|

| Port | Direction | Description | Width |
|------------------|-----------|--|-------|
| s_axis_rq_tdata | Input | Requester reQuest Data bus. This input contains the requester-side request data from the client application to the core. Only the lower 128 bits are to be used when the interface width is 128 bits, and only the lower 64 bits are to be used when the interface width is 64 bits. | DW |
| s_axis_rq_tuser | Input | Requester reQuest User Data. This set of signals contains sideband information for the TLP being transferred. These signals are valid when s_axis_rq_tvalid is High. Table 2-8, page 24 describes the individual signals in this set. | 60 |
| s_axis_rq_tlast | Input | TLAST Indication for Requester reQuest Data. The client application must assert this signal in the last cycle of a TLP to indicate the end of the packet. When the TLP is transferred in a single beat, the client must set this bit in the first cycle of the transfer. | 1 |
| s_axis_rq_tkeep | Input | TKEEP Indication for Requester reQuest Data. The assertion of bit <i>i</i> of this bus during a transfer indicates to the core that Dword <i>i</i> of the s_axis_rq_tdata bus contains valid data. The client must set this bit to 1 contiguously for all Dwords, starting from the first Dword of the descriptor to the last Dword of the payload. Thus, s_axis_rq_tdata must be set to all 1s in all beats of a packet, except in the final beat when the total size of the packet is not a multiple of the width of the data bus (both in Dwords). This is true for both Dword-aligned and address-aligned modes of payload transfer. Bits [7:4] of this bus are not used by the core when the interface width is configured as 128 bits, and bits [7:2] are not used when the interface width is configured as 64 bits. | DW/32 |
| s_axis_rq_tready | Output | Requester reQuest Data Ready. Activation of this signal by the core indicates that it is ready to accept data. Data is transferred across the interface when both s_axis_rq_tvalid and s_axis_rq_tready are asserted in the same cycle. If the core deasserts the ready signal when the valid signal is High, the client must maintain the data on the bus and keep the valid signal asserted until the core has asserted the ready signal. | 1 |
| s_axis_rq_tvalid | Input | Requester reQuest Data Valid. The client application must assert this output whenever it is driving valid data on the s_axis_rq_tdata bus. The client must keep the valid signal asserted during the transfer of a packet. The core paces the data transfer using the s_axis_rq_tready signal. | 1 |

| Port | Direction | Description | Width |
|---------------------|-----------|---|-------|
| pcie_rq_seq_num | Output | Requester reQuest TLP transmit sequence number. The client can optionally use this output to track the progress of the request in the core's transmit pipeline. To use this feature, the client must provide a sequence number for each request on the s_axis_rq_seq_num[3:0] bus. The core outputs this sequence number on the pcie_rq_seq_num[3:0] output when the request TLP has reached a point in the pipeline where a Completion TLP from the client cannot pass it. This mechanism enables the client to maintain ordering between Completions sent to the completer completion interface of the core and Posted requests sent to the requester request interface. Data on the pcie_rq_seq_num[3:0] output is valid when pcie_rq_seq_num_v1d is High. | 4 |
| pcie_rq_seq_num_vld | Output | Requester reQuest TLP transmit sequence number valid. This output is asserted by the core for one cycle when it has placed valid data on pcie_rq_seq_num[3:0]. | 1 |
| pcie_rq_tag | Output | Requester reQuest Non-Posted tag. When tag management for Non-Posted requests is performed by the core (AXISTEN_IF_ENABLE_CLIENT_TAG is 0), this output is used by the core to communicate the allocated tag for each Non-Posted request received from the client. The tag value on this bus is valid for one cycle when pcie_rq_tag_v1d is High. The client must copy this tag and use it to associate the completion data with the pending request. There can be a delay of several cycles between the transfer of the request on the s_axis_rq_tdata bus and the assertion of pcie_rq_tag_v1d by the core to provide the allocated tag for the request. The tags for requests are communicated on this bus in FIFO order, so the client can easily associate the tag value with the request it transferred. | 6 |
| pcie_rq_tag_vld | Output | Requester reQuest Non-Posted tag valid. The core asserts this output for one cycle when it has allocated a tag to an incoming Non-Posted request from the requester request interface and placed it on the pcie_rq_tag output. | 1 |

Table 2-7: RQ Interface Port Descriptions (Cont'd)

| Table 2-8: | Sideband Signal Descriptions in s_axis_rq_tuse | er |
|------------|--|----|
|------------|--|----|

| Bit Index | Name | Width | Description |
|-----------|------------------|-------|---|
| 3:0 | first_be[3:0] | 4 | Byte enables for the first Dword. This field must be set based on the desired value of the First_BE bits in the Transaction-Layer header of the request TLP. For Memory Reads, I/O Reads, and Configuration Reads, these four bits indicate the valid bytes to be read in the first Dword. For Memory Writes, I/O Writes, and Configuration Writes, these bits indicate the valid bytes in the first Dword of the payload. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. |
| 7:4 | last_be[3:0] | 4 | Byte enables for the last Dword. This field must be set based on the desired value of the Last_BE bits in the Transaction-Layer header of the TLP. For Memory Reads of two Dwords or more, these four bits indicate the valid bytes to be read in the last Dword of the block of data. For Memory Writes of two Dwords or more, these bits indicate the valid bytes in the last Dword of the payload. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. |
| 10:8 | addr_offset[2:0] | 3 | When the address-aligned mode is in use on this interface, the client application must provide the byte lane number where the payload data begins on the data bus, modulo 4, on this sideband bus. This enables the core to determine the alignment of the data block being transferred. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. When the requester request interface is configured in the Dword-alignment mode, this field must always be set to 0. |
| 11 | discontinue | 1 | This signal can be asserted by the client application during a transfer if it has detected an error in the data being transferred and desires to abort the packet. The core nullifies the corresponding TLP on the link to avoid data corruption. The client can assert this signal in any cycle during the transfer. It can either choose to terminate the packet prematurely in the cycle where the error was signaled, or can continue until all bytes of the payload are delivered to the core. In the latter case, the core treats the error as sticky for the following beats of the packet, even if the client deasserts the discontinue signal before the end of the packet. The discontinue signal can be asserted only when s_axis_rq_tready is High. Thus, when asserted, it should not be deasserted until s_axis_rq_tready is High. When the core is configured as an Endpoint, this error is also reported by the core to the Root Complex to which it is attached, using Advanced Error Reporting (AER). |

| Bit Index | Name | Width | Description |
|-----------|---------------------|-------|--|
| 12 | tph_present | 1 | This bit indicates the presence of a Transaction Processing Hint (TPH) in the request TLP being delivered across the interface. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. This bit must be permanently tied to 0 if the TPH capability is not in use. |
| 14:13 | tph_type[1:0] | 2 | When a TPH is present in the request TLP, these two bits provide the value of the PH[1:0] field associated with the hint. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. These bits can be set to any value if tph_present is set to 0. |
| 15 | tph_indirect_tag_en | 1 | When this bit is set, the core uses the lower bits of tph_st_tag[7:0] as an index into its Steering Tag Table, and insert the tag from this location in the transmitted request TLP. When this bit is 0, the core uses the value on tph_st_tag[7:0] directly as the Steering Tag. The core samples this bit in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. This bit can be set to any value if tph_present is set to 0. |
| 23:16 | tph_st_tag[7:0] | 8 | When a TPH is present in the request TLP, this output provides the 8-bit Steering Tag associated with the hint. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. These bits can be set to any value if tph_present is set to 0. |
| 27:24 | seq_num[3:0] | 4 | The client can optionally supply a 4-bit sequence number in this field to keep track of the progress of the request in the core's transmit pipeline. The core outputs this sequence number on its pcie_rq_seq_num[3:0] output when the request TLP has progressed to a point in the pipeline where a Completion TLP from the client is not able to pass it. The core samples this field in the first beat of a packet, when s_axis_rq_tvalid and s_axis_rq_tready are both High. This input can be hardwired to 0 when the client is not monitoring the pcie_rq_seq_num[3:0] output of the core. |
| 59:28 | parity | 32 | Odd parity for the 256-bit data. When parity checking is enabled in the core, client logic must set bit <i>i</i> of this bus to the odd parity computed for byte <i>i</i> of s_axis_rq_tdata. Only the lower 16 bits are to be used when the interface width is 128 bits, and only the lower 8 bits are to be used when the interface width is 64 bits. On detection of a parity error, the core nullifies the corresponding TLP on the link and reports it as an Uncorrectable Internal Error. These bits can be set to 0 if parity checking is disabled in the core. |

Table 2-8: Sideband Signal Descriptions in s_axis_rq_tuser (Cont'd)

Requester Completion (RC) Interface

Table 2-9 defines the ports in the RC interface of the Gen3 Integrated Block for PCI Express core. In the Width column, DW denotes the configured data bus width (64, 128, or 256 bits).

Table 2-9: RC Interface Port Descriptions

| Port | Direction | Description | Width |
|-----------------|-----------|--|-------|
| m_axis_rc_tdata | Output | Requester Completion Data bus. Transmit data from the Core requester completion interface to the client application. Only the lower 128 bits are used when the interface width is 128 bits, and only the lower 64 bits are used when the interface width is 64 bits. Bits [255:128] are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [255:64] are set permanently to 0 when the interface width is configured as 64 bits. | DW |
| m_axis_rc_tuser | Output | Requester Completion User Data. This set of signals contains sideband information for the TLP being transferred. These signals are valid when m_axis_rc_tvalid is High. Table 2-10, page 28 describes the individual signals in this set. | 75 |
| m_axis_rc_tlast | Output | TLAST indication for Requester Completion Data. The core asserts this signal in the last beat of a packet to indicate the end of the packet. When a TLP is transferred in a single beat, the core sets this bit in the first beat of the transfer. This output is used only when the straddle option is disabled. When the straddle option is enabled (for 256-bit interface), the core sets this output permanently to 0. | 1 |

| Port | Direction | Description | Width |
|------------------|-----------|---|-------|
| m_axis_rc_tkeep | Output | TKEEP indication for Requester Completion Data. The assertion of bit <i>i</i> of this bus during a transfer indicates to the client that Dword <i>i</i> of the m_axis_rc_tdata bus contains valid data. The core sets this bit to 1 contiguously for all Dwords starting from the first Dword of the descriptor to the last Dword of the payload. Thus, m_axis_cq_tdata is set to all 1s in all beats of a packet, except in the final beat when the total size of the packet is not a multiple of the width of the data bus (both in Dwords). This is true for both Dword-aligned and address-aligned modes of payload transfer. Bits [7:4] of this bus are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [7:2] are set permanently to 0 when the interface width is configured as 64 bits. These outputs are permanently set to all 1s when the interface width is 256 bits and the straddle option is enabled. The client logic must use the signals in m_axis_rc_tuser in that case to determine the start and end of Completion TLPs transferred over the interface. | DW/32 |
| m_axis_rc_tvalid | Output | Requester Completion Data Valid. The core asserts this output whenever it is driving valid data on the m_axis_rc_tdata bus. The core keeps the valid signal asserted during the transfer of a packet. The client application can pace the data transfer using the m_axis_rc_tready signal. | 1 |
| m_axis_rc_tready | Input | Requester Completion Data Ready. Activation of this signal by the client logic indicates to the core that the client is ready to accept data. Data is transferred across the interface when both m_axis_rc_tvalid and m_axis_rc_tready are asserted in the same cycle. If the client deasserts the ready signal when the valid signal is High, the core maintains the data on the bus and keeps the valid signal asserted until the client has asserted the ready signal. | 1 |

Table 2-9: RC Interface Port Descriptions (Cont'd)

| Bit Index | Name | Width | Description |
|-----------|---------------|-------|---|
| 31:0 | byte_en | 32 | The client logic can optionally use these byte enable bits to determine the valid bytes in the payload of a packet being transferred. The assertion of bit <i>i</i> of this bus during a transfer indicates to the client that byte <i>i</i> of the $m_axis_cq_tdata$ bus contains a valid payload byte. This bit is not asserted for descriptor bytes. Although the byte enables can be generated by client logic from information in the request descriptor (address and length), the client has the option to use these signals directly instead of generating them from other interface signals. The 1 bit in this bus for the payload of a TLP is always contiguous. Bits [31:16] of this bus are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [31:8] are set permanently to 0 when the interface width is configured as 64 bits. |
| 32 | is_sof_0 | 1 | Start of a first Completion TLP. For 64-bit and 128-bit interfaces, and for the 256-bit interface with no straddling, is_sof_0 is asserted by the core in the first beat of a packet to indicate the start of the TLP. On these interfaces, only a single TLP can be started in a data beat, and is_sof_1 is permanently set to 0. Use of this signal is optional for the client when the straddle option is not enabled. When the interface width is 256 bits and the straddle option is enabled, the core can straddle two Completion TLPs in the same beat. In this case, the Completion TLPs are not formatted as AXI4-Stream packets. The assertion of is_sof_0 indicates a Completion TLP starting in the beat. The first byte of this Completion TLP is in byte lane 0 if the previous TLP ended before this beat, or in byte lane 16 if the previous TLP continues in this beat. |
| 33 | is_sof_1 | 1 | Start of a second Completion TLP. This signal is used when the interface width is 256 bits and the straddle option is enabled, when the core can straddle two Completion TLPs in the same beat. The output is permanently set to 0 in all other cases. The assertion of is_sof_1 indicates a second Completion TLP starting in the beat, with its first bye in byte lane 16. The core starts a second TLP at byte position 16 only if the previous TLP ended in one of the byte positions 0-15 in the same beat; that is, only if $is_eof_0[0]$ is also set in the same beat. |
| 37:34 | is_eof_0[3:0] | 4 | End of a first Completion TLP and the offset of its last Dword. These outputs are used only when the interface width is 256 bits and the straddle option is enabled. The assertion of the bit $is_eof_0[0]$ indicates the end of a first Completion TLP in the current beat. When this bit is set, the bits $is_eof_0[3:1]$ provide the offset of the last Dword of this TLP. |

Table 2-10: Sideband Signal Descriptions in m_axis_rc_tuser

| Bit Index | Name | Width | Description |
|-----------|---------------|-------|--|
| 41:38 | is_eof_1[3:0] | 4 | End of a second Completion TLP and the offset of its last Dword. These outputs are used only when the interface width is 256 bits and the straddle option is enabled. The core can then straddle two Completion TLPs in the same beat. These outputs are reserved in all other cases. The assertion of $is_eof_1[0]$ indicates a second TLP ending in the same beat. When bit 0 of is_eof_1 is set, bits [3:1] provide the offset of the last Dword of the TLP ending in this beat. Because the second TLP can only end at a byte position in the range 27–31, $is_eof_1[3:1]$ can only take one of two values (6 or 7). The offset for the last byte of the second TLP can be determined from the starting address and length of the TLP, or from the byte enable signals $byte_en[31:0]$. If $is_eof_1[0]$ is High, the signals $is_eof_0[0]$ and is_sof_1 are also High in the same beat. |
| 42 | discontinue | 1 | This signal is asserted by the core in the last beat of a TLP, if it has detected an uncorrectable error while reading the TLP payload from its internal FIFO memory. The client application must discard the entire TLP when such an error is signaled by the core. This signal is never asserted when the TLP has no payload. It is asserted only in the last beat of the payload transfer; that is, when $is_eof_0[0]$ is High. When the straddle option is enabled, the core does not start a second TLP if it has asserted discontinue in a beat. When the core is configured as an Endpoint, the error is also reported by the core to the Root Complex to which it is attached, using Advanced Error Reporting (AER). |
| 74:43 | parity | 32 | Odd parity for the 256-bit transmit data. Bit <i>i</i> provides the odd parity computed for byte <i>i</i> of m_axis_rc_tdata. Only the lower 16 bits are used when the interface width is 128 bits, and only the lower 8 bits are used when the interface width is 64 bits. Bits [31:16] are set permanently to 0 by the core when the interface width is configured as 128 bits, and bits [31:8] are set permanently to 0 when the interface width is configured as 64 bits. |

Table 2-10: Sideband Signal Descriptions in m_axis_rc_tuser (Cont'd)

Transmit Flow Control Interface

Table 2-11 defines the ports in the Transmit Flow Control interface of the Integrated Block for PCI Express core.

| Port | Direction | Description | Width |
|-----------------|-----------|--|-------|
| pcie_tfc_nph_av | Output | Transmit flow control Non-Posted header credits available. This output indicates the currently available header credit for Non-Posted TLPs on the transmit side of the core. Client logic can check this output before transmitting a Non-Posted request on the requester request interface, to avoid blocking the interface when no credit is available. The encodings are: • 00: No credits available • 01: 1 credit available • 10: 2 credits available • 11: 3 or more credits available Because of pipeline delays, the value on this output does not include the credit consumed by the Non-Posted requests sent by the client in the last two clock cycles. The client logic must adjust the value on this output by the credit consumed by the Non-Posted requests it sent in the previous two clock cycles, if any. | 2 |
| pcie_tfc_npd_av | Output | Transmit flow control Non-Posted data credits available. This output indicates the currently available payload credit for Non-Posted TLPs on the transmit side of the core. Client logic can check this output before transmitting a Non-Posted request on the requester request interface, to avoid blocking the interface when no credit is available. The encodings are: • 00: No credits available • 01: 1 credit available • 10: 2 credits available • 11: 3 or more credits available Because of pipeline delays, the value on this output does not include the credit consumed by the Non-Posted requests sent by the client in the last two clock cycles. The client logic must adjust the value on this output by the credit consumed by the Non-Posted requests it sent in the previous two clock cycles, if any. | 2 |

| Table 2-11: Transmit Flow Control Interface Port Descriptions |
|---|
|---|

Configuration Management Interface

Table 2-12 defines the ports in the Configuration Management interface of the IntegratedBlock for PCI Express core.

| Port | Direction | Description | Width |
|-------------------------------|-----------|--|-------|
| cfg_mgmt_addr | Input | Read/Write Address. Address is in the Configuration and Management register space, and is Dword aligned. For accesses from the local management bus: for the address bits cfg_mgmt_addr[17:10], select the PCI Function associated with the configuration register; and for the bits cfg_mgmt_addr[9:0], select the register within the Function. The address bit cfg_mgmt_addr[18] must be set to zero (0) when accessing the PCI or PCI Express configuration registers, and to one (1) when accessing the local management registers. | 19 |
| cfg_mgmt_write | Input | Write Enable. Asserted for a write operation. Active High. | 1 |
| cfg_mgmt_write_data | Input | Write data. Write data is used to configure the Configuration and Management registers. | 32 |
| cfg_mgmt_byte_enable | Input | Byte Enable. Byte enable for write data, where cfg_mgmt_byte_enable[0] corresponds to cfg_mgmt_write_data[7:0], and so on. | 4 |
| cfg_mgmt_read | Input | Read Enable. Asserted for a read operation. Active High. | 1 |
| cfg_mgmt_read_data | Output | Read data out. Read data provides the configuration of the Configuration and Management registers. | 32 |
| cfg_mgmt_read_write_done | Output | Read/Write operation complete. Asserted for 1 cycle when operation is complete. Active High. | 1 |
| cfg_mgmt_type1_cfg_reg_access | Input | Type 1 RO, Write. When the core is configured in the Root Port mode, asserting this input during a write to a Type-1 PCI [™] Config Register forces a write into certain read-only fields of the register (see description of RC-mode Config registers). This input has no effect when the core is in the Endpoint mode, or when writing to any register other than a Type-1 Config Register. | 1 |

Table 2-12: Configuration Management Interface Port Descriptions

Configuration Status Interface

Table 2-13 defines the ports in the Configuration Status interface of the Integrated Block for PCI Express core.

| Port | Direction | Description | Width |
|----------------------|-----------|---|-------|
| cfg_phy_link_down | Output | Configuration Link Down. Status of the PCI Express link based on Physical Layer LTSSM. 1b: Link is Down (LinkUp state variable is 0b) 0b: Link is Up (LinkUp state variable is 1b) Note: Per the PCI Express Base Specification, rev. 3.0, LinkUp is 1b in the Recovery, L0, L0s, L1, and L2 cfg_ltssm states. In the Configuration state, LinkUp can be 0b or 1b. It is always 0b when the Configuration state is reached via Detect > Polling > Configuration. LinkUp is 1b if the configuration state is reached via | 1 |
| cfg_phy_link_status | Output | any other state transition. Configuration Link Status. Status of the PCI Express link. 00b: No receivers detected 01b: Link training in progress 10b: Link up, DL initialization in progress 11b: Link up, DL initialization completed | 2 |
| cfg_negotiated_width | Output | Configuration Link Status. Negotiated Link Width: PCI Express Link Status Register, Negotiated Link Width field. This field indicates the negotiated width of the given PCI Express Link and is valid when cfg_phy_link_status[1:0] == 11b (DL Initialization is complete). | 4 |
| cfg_current_speed | Output | Current Link Speed. This signal outputs the current link speed from Link Status register bits 1 down to 0. This field indicates the negotiated Link speed of the given PCI Express Link. 001b: 2.5 GT/s PCI Express Link 010b: 5.0 GT/s PCI Express Link 100b: 8.0 GT/s PCI Express Link | 3 |
| cfg_max_payload | Output | Max_Payload_Size. This signal outputs the maximum payload size from Device Control Register bits 7 down to 5. This field sets the maximum TLP payload size. As a Receiver, the user logic must handle TLPs as large as the set value. As a Transmitter, the user logic must not generate TLPs exceeding the set value. 000b: 128 bytes maximum payload size 001b: 256 bytes maximum payload size 010b: 512 bytes maximum payload size 011b: 1024 bytes maximum payload size 100b: 2048 bytes maximum payload size 101b: 4096 bytes maximum payload size | 3 |

Table 2-13: Configuration Status Interface Port Descriptions

| Port | Direction | Description | Width |
|--------------------------|-----------|--|-------|
| cfg_max_read_req | Output | Max_Read_Request_Size. This signal outputs the maximum read request size from Device Control register bits 14 down to 12. This field sets the maximum Read Request size for the user logic as a Requester. The user logic must not generate Read Requests with size exceeding the set value. 000b: 128 bytes maximum Read Request size 001b: 256 bytes maximum Read Request size 010b: 512 bytes maximum Read Request size 011b: 1024 bytes maximum Read Request size 100b: 2048 bytes maximum Read Request size 101b: 4096 bytes maximum Read Request size | 3 |
| cfg_function_status | Output | Configuration Function Status. These outputs indicate the states of the Command Register bits in the PCI configuration space of each Function. These outputs are used to enable requests and completions from the host logic. The assignment of bits is as follows: • Bit 0: Function 0 I/O Space Enable • Bit 1: Function 0 Memory Space Enable • Bit 2: Function 0 Bus Master Enable • Bit 2: Function 0 INTx Disable • Bit 4: Function 1 I/O Space Enable • Bit 5: Function 1 Memory Space Enable • Bit 5: Function 1 Memory Space Enable • Bit 6: Function 1 Bus Master Enable • Bit 7: Function 1 INTx Disable | 8 |
| cfg_vf_status | Output | Configuration Virtual Function Status. These outputs indicate the status of Virtual Functions, two bits each per Virtual Function. Bit 0: Virtual Function 0: Configured/Enabled by software Bit 1: Virtual Function 0: PCI Command Register, Bus Master Enable, and so on. | 12 |
| cfg_function_power_state | Output | Configuration Function Power State. These outputs indicate the current power state of the Physical Functions. Bits [2:0] capture the power state of Function 0, and bits [5:3] capture that of Function 1. The possible power states are: • 000: D0_uninitialized • 001: D0_active • 010: D1 • 100: D3_hot | 6 |

Table 2-13: Configuration Status Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|----------------------|-----------|---|-------|
| cfg_vf_power_state | Output | Configuration Virtual Function Power State. These outputs indicate the current power state of the Virtual Functions. Bits [2:0] capture the power state of Virtual Function 0, and bits [5:3] capture that of Virtual Function 1, and so on. The possible power states are: • 000: D0_uninitialized • 001: D0_active • 010: D1 • 100: D3_hot | 18 |
| cfg_link_power_state | Output | Current power state of the PCI Express link: • 00: L0 • 01: L0s • 10: L1 • 11: L2/Reserved | 2 |
| cfg_err_cor_out | Output | Correctable Error Detected. In the Endpoint mode, the core activates this output for one cycle when it has detected a correctable error and its reporting is not masked. In a multi-Function Endpoint, this is the logical OR of the correctable error status bits in the Device Status Registers of all Functions. In the Root Port mode, this output is activated on detection of a local correctable error, when its reporting is not masked. This output does not respond to any errors signaled by remote devices using PCI Express error messages. These error messages are delivered to the user through the message interface. | 1 |
| cfg_err_nonfatal_out | Output | Non-Fatal Error Detected. In the Endpoint mode, the core activates this output for one cycle when it has detected a non fatal error and its reporting is not masked. In a multi-Function Endpoint, this is the logical OR of the non fatal error status bits in the Device Status Registers of all Functions. In the Root Port mode, this output is activated on detection of a local non-fatal error, when its reporting is not masked. This output does not respond to any errors signaled by remote devices using PCI Express error messages. These error messages are delivered to the user through the message interface. | 1 |
| cfg_err_fatal_out | Output | Fatal Error Detected. In the Endpoint mode, the core activates this output for one cycle when it has detected a fatal error and its reporting is not masked. In a multi-Function Endpoint, this is the logical OR of the fatal error status bits in the Device Status Registers of all Functions. In the Root Port mode, this output is activated on detection of a local fatal error, when its reporting is not masked. This output does not respond to any errors signaled by remote devices using PCI Express error messages. These error messages are delivered to the user through the message interface. | 1 |

Table 2-13: Configuration Status Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-----------------|-----------|---|-------|
| cfg_local_error | Output | Local Error Conditions. Output is High when any of the local error conditions listed in Local Error Status Register occur. Each of the conditions can be selectively masked by setting the corresponding bits in the Local Interrupt Mask Register. | 1 |
| cfg_ltr_enable | Output | Latency Tolerance Reporting Enable. The state of this output reflects the setting of the LTR Mechanism Enable bit in the Device Control 2 Register of Physical Function 0. When the core is configured as an Endpoint user logic uses this output to enable the generation of LTR messages. This output is not to be used when the core is configured as a Root Port. | 1 |

Table 2-13: Configuration Status Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-----------------|-----------|--|-------|
| cfg_ltssm_state | Output | Current LTSSM State. Shows the current LTSSM state: Detect.Quiet 00 Detect.Active 01 Polling.Compliance 03 Polling.Configuration 04 Configuration.Linkwidth.Start 05 Configuration.Linkwidth.Accept 06 Configuration.Lanenum.Accept 07 Configuration.Lanenum.Wait 08 Configuration.Lanenum.Wait 08 Configuration.Complete 09 Configuration.Idle 0A Recovery.RcvrLock 0B Recovery.RcvrLock 0B Recovery.RcvrCfg 0D Recovery.RcvrCfg 0D Recovery.Idle 0E L0 10 Rx_L0s.Entry 11 Rx_L0s.Idle 12 Rx_L0s.Entry 11 Rx_L0s.Idle 12 Rx_L0s.Entry 14 Tx_L0s.Idle 15 Tx_L0s.FTS 16 L1.Entry 17 L1.Idle 18 L2.Idle 19 L2.TransmitWake 1A DISABLED = 6'h20; LOOPBACK_ENTRY_MASTER = 6'h21; LOOPBACK_CACTIVE_MASTER = 6'h22; LOOPBACK_ENTRY_SLAVE = 6'h23; LOOPBACK_EXIT_SLAVE = 6'h25; LOOPBACK_EXIT_SLAVE = 6'h25; LOOPBACK_EXIT_SLAVE = 6'h25; LOOPBACK_EXIT_SLAVE = 6'h26; HOT_RESET = 6'h27; RECOVERY_EQUALIZATION_PHASE0 = 6'h28; RECOVERY_EQUALIZATION_PHASE1 = 6'h29; RECOVERY_EQUALIZATION_PHASE3 = 6'h28; | 6 |

Table 2-13: Configuration Status Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-------------------------|-----------|---|-------|
| cfg_rcb_status | Output | RCB Status. Provides the setting of the Read Completion Boundary (RCB) bit in the Link Control Register of each Physical Function. In the Endpoint mode, bit 0 indicates the RCB for PF 0, and so on. In the RC mode, bit 0 indicates the RCB setting of the Link Control Register of the RP, bit 1 is reserved. For each bit, a value of 0 indicates an RCB of 64 bytes and a value of 1 indicates 128 bytes. | 2 |
| cfg_dpa_substate_change | Output | Dynamic Power Allocation Substate Change. In the Endpoint mode, the core generates a one-cycle pulse on one of these outputs when a Configuration Write transaction writes into the Dynamic Power Allocation Control Register to modify the DPA power state of the device. A pulse on bit 0 indicates such a DPA event for PF 0 and so on. These outputs are not active in the Root Port mode. | 2 |
| cfg_obff_enable | Output | Optimized Buffer Flush Fill Enable. This output reflects the setting of the OBFF Enable field in the Device Control 2 Register. 00: OBFF disabled 01: OBFF enabled using message signaling, Variation A 10: OBFF enabled using message signaling, Variation B 11: OBFF enabled using WAKE# signaling. | 2 |

Table 2-13: Configuration Status Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-----------------------------|-----------|--|-------|
| cfg_pl_status_change | Output | This output is used by the core in the Root Port mode to signal one of the following link training-related events: (a) The link bandwidth changed as a result of the change in the link width or operating speed and the change was initiated locally (not by the link partner), without the link going down. This interrupt is enabled by the Link Bandwidth Management Interrupt Enable bit in the Link Control Register. The status of this interrupt can be read from the Link Bandwidth Management Status bit of the Link Status Register; or (b) The link bandwidth changed autonomously as a result of the change in the link width or operating speed and the change was initiated by the remote node. This interrupt is enabled by the Link Autonomous Bandwidth Interrupt Enable bit in the Link Control Register. The status of this interrupt can be read from the Link Autonomous Bandwidth Status bit of the Link Status Register; or (c) The Link Equalization Request bit in the Link Status 2 Register was set by the hardware because it received a link equalization request from the remote node. This interrupt is enabled by the Link Equalization Request bit of the Link Status 2 Register. The p1_interrupt output is not active when the core is configured as an Endpoint. | 1 |
| cfg_tph_requester_enable | Output | Bit 0 of this output reflect the setting of the TPH Requester Enable bit [8] of the TPH Requester Control Register in the TPH Requester Capability Structure of Physical Function 0. Bit 1 corresponds to Physical Function 1. | 2 |
| cfg_tph_st_mode | Output | Bits [2:0] of this output reflect the setting of the ST Mode Select bits in the TPH Requester Control Register of Physical Function 0. Bits [5:3] reflect the setting of the same register field of PF 1. | 6 |
| cfg_vf_tph_requester_enable | Output | Each of the six bits of this output reflects the setting of the TPH Requester Enable bit 8 of the TPH Requester Control Register in the TPH Requester Capability Structure of the corresponding Virtual Function. | 6 |
| cfg_vf_tph_st_mode | Output | Bits [2:0] of this output reflect the setting of the ST Mode Select bits in the TPH Requester Control Register of Virtual Function 0. Bits [5:3] reflect the setting of the same register field of VF 1, and so on. | 18 |

Table 2-13: Configuration Status Interface Port Descriptions (Cont'd)

Configuration Received Message Interface

Table 2-14 defines the ports in the Configuration Received Message interface of the Integrated Block for PCI Express core.

| Port | Direction | Description | Width |
|-----------------------|-----------|--|-------|
| cfg_msg_received | Output | Configuration Received a Decodable Message. The core asserts this output for one or more consecutive clock cycles when it has received a decodable message from the link. The duration of its assertion is determined by the type of message. The core transfers any parameters associated with the message on the cfg_msg_data[7:0]output in one or more cycles when cfg_msg_received is High. Table 3-8, page 116 lists the number of cycles of cfg_msg_received assertion, and the parameters transferred on cfg_msg_data[7:0] in each cycle, for each type of message. The core inserts at least a one-cycle gap between two consecutive messages delivered on this interface. This output is active only when the AXISTEN_IF_ENABLE_RX_MSG_INTFC attribute is set. | 1 |
| cfg_msg_received_data | Output | This bus is used to transfer any parameters associated with the Received Message. The information it carries in each cycle for various message types is listed in Table 3-8, page 116. | 8 |
| cfg_msg_received_type | Output | Received message type. When cfg_msg_received is High, these five bits indicate the type of message being signaled by the core. The various message types are listed in Table 3-7, page 115. | 5 |

Configuration Transmit Message Interface

Table 2-15 defines the ports in the Configuration Transmit Message interface of the Integrated Block for PCI Express core.

| Port | Direction | Description | Width |
|-----------------------|-----------|--|-------|
| cfg_msg_transmit | Input | Configuration Transmit Encoded Message. This signal is asserted together with cfg_msg_transmit_type, which supplies the encoded message type and cfg_msg_transmit_data, which supplies optional data associated with the message, until cfg_msg_transmit_done is asserted in response. | |
| cfg_msg_transmit_type | Input | Configuration Transmit Encoded Message Type. Indicates the type of PCI Express message to be transmitted. Encodings supported are: • 000b: Latency Tolerance Reporting (LTR) • 001b: Optimized Buffer Flush/Fill (OBFF) • 010b: Set Slot Power Limit (SSPL) • 011b: Power Management (PM PME) • 100b -111b: Reserved | 3 |
| cfg_msg_transmit_data | Input | Configuration Transmit Encoded Message Data. Indicates message data associated with particular message type. 000b: LTR - cfg_msg_transmit_data[31] ≤ Snoop Latency Req., cfg_msg_transmit_data[28:26] ≤ Snoop Latency Scale, cfg_msg_transmit_data[25:16] ≤ Snoop Latency Value, cfg_msg_transmit_data[15] ≤ No-Snoop Latency Requirement, cfg_msg_transmit_data[12:10] ≤ No-Snoop Latency Scale, cfg_msg_transmit_data[9:0] ≤ No-Snoop Latency Value. 001b: OBFF - cfg_msg_transmit_data[3:0] ≤ OBFF Code. 010b: SSPL - cfg_msg_transmit_data[9:0] ≤ {Slot Power Limit Scale, Slot Power Limit Value}. 011b: PM_PME - cfg_msg_transmit_data[1:0] ≤ PF1, PF0; cfg_msg_transmit_data[9:4] ≤ VF5, VF4, VF3, VF2, VF1, VF0, where one or more PFs or VFs can signal PM_PME simultaneously. 100b - 111b: Reserved | 32 |
| cfg_msg_transmit_done | Output | Configuration Transmit Encoded Message Done. Asserted in response to cfg_mg_transmit assertion, for 1 cycle after the request is complete. | 1 |

Table 2-15: Configuration Transmit Message Interface Port Descriptions

Configuration Flow Control Interface

Table 2-16 defines the ports in the Configuration Flow Control interface of the Integrated Block for PCI Express core.

| Port | Direction | Description | Width |
|-------------|-----------|---|-------|
| cfg_fc_ph | Output | Posted Header Flow Control Credits. This output provides the number of Posted Header Flow Control Credits. This multiplexed output can be used to bring out various flow control parameters and variables related to Posted Header Credit maintained by the core. The flow control information to bring out on this core is selected by the cfg_fc_sel[2:0] input. | 8 |
| cfg_fc_pd | Output | Posted Data Flow Control Credits. This output provides the number of Posted Data Flow Control Credits. This multiplexed output can be used to bring out various flow control parameters and variables related to Posted Data Credit maintained by the core. The flow control information to bring out on this core is selected by the $cfg_fc_sel[2:0]$ input. | 12 |
| cfg_fc_nph | Output | Non-Posted Header Flow Control Credits. This output provides the number of Non-Posted Header Flow Control Credits. This multiplexed output can be used to bring out various flow control parameters and variables related to Non-Posted Header Credit maintained by the core. The flow control information to bring out on this core is selected by the cfg_fc_sel[2:0] input. | 8 |
| cfg_fc_npd | Output | Non-Posted Data Flow Control Credits. This output provides the number of Non-Posted Data Flow Control Credits. This multiplexed output can be used to bring out various flow control parameters and variables related to Non-Posted Data Credit maintained by the core. The flow control information to bring out on this core is selected by the $cfg_fc_sel[2:0]$ input. | 12 |
| cfg_fc_cplh | Output | Completion Header Flow Control Credits. This output provides the number of Completion Header Flow Control Credits. This multiplexed output can be used to bring out various flow control parameters and variables related to Completion Header Credit maintained by the core. The flow control information to bring out on this core is selected by the cfg_fc_sel[2:0] input. | 8 |
| cfg_fc_cpld | Output | Completion Data Flow Control Credits. This output provides the number of Completion Data Flow Control Credits. This multiplexed output can be used to bring out various flow control parameters and variables related to Completion Data Credit maintained by the core. The flow control information to bring out on this core is selected by the $cfg_fc_sel[2:0]$. | 12 |

Table 2-16: Configuration Flow Control Interface Port Descriptions

| Port | Direction | Description | Width |
|------------|-----------|--|-------|
| cfg_fc_sel | Input | Flow Control Informational Select. These inputs select the type of flow control to bring out on the cfg_fc_* outputs of the core. The various flow control parameters and variables that can be accessed for the different settings of these inputs are: 000: Receive credits currently available to the link partner 001: Receive credit limit 010: Receive credits consumed 011: Available space in receive buffer 100: Transmit credits available 101: Transmit credits consumed 111: Reserved Note: Infinite credit for transmit credits available (cfg_fc_sel == 3 'b100) is signaled as 8 'h80, 12 'h800 for header and data credits, respectively. For all other cfg_fc_sel selections, infinite credit is signaled as 8 'h00, 12 'h000, respectively, for header and data categories. | 3 |

Table 2-16: Configuration Flow Control Interface Port Descriptions (Cont'd)

Per Function Status Interface

Table 2-17 and Table 2-18 define the ports in the Function Status interface of the Integrated Block for PCI Express core.

| Table 2-17. Overview of Function Status Interface Port Descriptions | Table 2-17: | Overview of Function Status Interface Port Descriptions |
|---|-------------|--|
|---|-------------|--|

| Port | Direction | Description | Width |
|-----------------------------|-----------|---|-------|
| cfg_per_func_status_control | Input | Configuration Per Function Control. Controls information presented on the multi-function output cfg_per_func_status_data. Supported encodings are 000b, 001b, 010b, 011b, 100b, and 101b. All other encodings are reserved. | 3 |
| cfg_per_func_status_data | Output | Configuration Per Function Status Data. Provides a 16-bit status output for the selected function. Information presented depends on the values of cfg_per_func_status_data and cfg_per_function_number. | 16 |

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|---------------------------------------|-------|---|
| 0 | 0 | cfg_command_ io_enable | 1 | Configuration Command - I/O Space Enable: Command[0]. Endpoints: If 1, allows the device to receive I/O Space accesses. Otherwise, the core filters these and respond with an Unsupported Request. Root/Switch: Core takes no action based on this setting. If 0, user logic must not generate TLPs downstream. |
| 0 | 1 | cfg_command_ mem_enable | 1 | Configuration Command - Memory Space Enable: Command[1]. Endpoints: If 1, allows the device to receive Memory Space accesses. Otherwise, the core filters these and respond with an Unsupported Request. Root/Switch: Core takes no action based on this setting. If 0, user logic must not generate TLPs downstream. |
| 0 | 2 | cfg_command_ bus_master_en able | 1 | Configuration Command - Bus Master Enable: Command[2]. The core takes no action based on this setting; user logic must do that. Endpoints: When asserted, the user logic is allowed to issue Memory or I/O Requests (including MSI/X interrupts); otherwise it must not. Root and Switch Ports: When asserted, received Memory or I/O Requests might be forwarded upstream; otherwise they are handled as Unsupported Requests (UR), and for Non-Posted Requests a Completion with UR completion status is returned. |
| 0 | 3 | cfg_command_ interrupt_disa ble | 1 | Configuration Command - Interrupt Disable: Command[10]. When asserted, the core is prevented from asserting INTx interrupts. |
| 0 | 4 | cfg_command_ serr_en | 1 | Configuration Command - SERR Enable: Command[8]. When asserted, this bit enables reporting of Non-fatal and Fatal errors. Note that errors are reported if enabled either through this bit or through the PCI Express specific bits in the Device Control register. In addition, for a Root Complex or Switch, this bit controls transmission by the primary interface of ERR_NONFATAL and ERR_FATAL error messages forwarded from the secondary interface. |

Table 2-18: Detailed Function Status Interface Port Descriptions

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|---|-------|--|
| 0 | 5 | cfg_bridge_ser r_en | 1 | Configuration Bridge Control - SERR Enable: Bridge Ctrl[1]. When asserted, this bit enables the forwarding of Correctable, Non-fatal and Fatal errors (user must enforce that). |
| 0 | 6 | cfg_aer_ecrc_c heck_en | 1 | Configuration AER - ECRC Check Enable: AER_Cap_and_Ctl[8]. When asserted, this bit indicates that ECRC checking has been enabled by the host. |
| 0 | 7 | cfg_aer_ecrc_g en_en | 1 | Configuration AER - ECRC Generation Enable: AER_Cap_and_Ctl[6]. When asserted, this bit indicates that ECRC generation has been enabled by the host. |
| 0 | 15:08 | 0 | 8 | Reserved |
| 1 | 0 | cfg_dev_status _corr_err_dete cted | 1 | Configuration Device Status - Correctable Error Detected: Device_Status[0]. Indicates status of correctable errors detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. |
| 1 | 1 | cfg_dev_status _non_fatal_err_ detected | 1 | Configuration Device Status - Non-Fatal Error Detected: Device_Status[1]. Indicates status of Nonfatal errors detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. |
| 1 | 2 | cfg_dev_status _fatal_err_dete cted | 1 | Configuration Device Status - Fatal Error Detected: Device_Status[2]. Indicates status of Fatal errors detected. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. |
| 1 | 3 | cfg_dev_status _ur_detected | 1 | Configuration Device Status - Unsupported Request Detected: Device_Status[3]. Indicates that the core received an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. |
| 1 | 4 | cfg_dev_contr ol_corr_err_rep orting_en | 1 | Configuration Device Control - Correctable Error Reporting Enable: Device_Ctrl[0]. This bit, in conjunction with other bits, controls sending ERR_COR Messages. For a Root Port, the reporting of correctable errors is internal to the root; no external ERR_COR Message is generated. |

Table 2-18: Detailed Function Status Interface Port Descriptions

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|--|-------|---|
| 1 | 5 | cfg_dev_contr ol_non_fatal_re porting_en | 1 | Configuration Device Control - Non-Fatal Error Reporting Enable: Device_Ctrl[1]. This bit, in conjunction with other bits, controls sending ERR_NONFATAL Messages. For a Root Port, the reporting of correctable errors is internal to the root; no external ERR_NONFATAL Message is generated. |
| 1 | 6 | cfg_dev_contr ol_fatal_err_re porting_en | 1 | Configuration Device Control - Fatal Error Reporting Enable: Device_Ctrl[2]. This bit, in conjunction with other bits, controls sending ERR_FATAL Messages. For a Root Port, the reporting of correctable errors is internal to the root; no external ERR_FATAL Message is generated. |
| 1 | 7 | cfg_dev_contr ol_ur_err_repor ting_en | 1 | Configuration Device Control - UR Reporting Enable: Device_Ctrl[3]. This bit, in conjunction with other bits, controls the signaling of Unsupported Requests by sending Error Messages. |
| 1 | 10:08 | cfg_dev_contr ol_max_payloa d | 3 | Configuration Device Control - Max_Payload_Size: Device_Ctrl[7:5]. This field sets maximum TLP payload size. As a Receiver, the user logic must handle TLPs as large as the set value. As a Transmitter, the user logic must not generate TLPs exceeding the set value. 000b = 128 bytes max payload size 001b = 256 bytes max payload size 010b = 512 bytes max payload size 011b = 1024 bytes max payload size 100b = 2048 bytes max payload size |
| 1 | 11 | cfg_dev_contr ol_enable_ro | 1 | Configuration Device Control - Enable Relaxed Ordering: Device_Ctrl[4]. When asserted, the user logic is permitted to set the Relaxed Ordering bit in the Attributes field of transactions it initiates that do not require strong write ordering. |
| 1 | 12 | cfg_dev_contr ol_ext_tag_en | 1 | Configuration Device Control - Tag Field Enable: Device_Ctrl[8]. When asserted, enables user logic to use an 8-bit Tag field as a Requester. If deasserted, user logic is restricted to a 5-bit Tag field. Note that the core does not enforce the number of Tag bits used, either in outgoing request TLPs or incoming Completions. |

Table 2-18: Detailed Function Status Interface Port Descriptions

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|--|-------|--|
| 1 | 13 | cfg_dev_contr ol_no_snoop_e n | 1 | Configuration Device Control - Enable No Snoop: Device_Ctrl[11]. When asserted, the user logic is permitted to set the No Snoop bit in TLPs it initiates that do not require hardware enforced cache coherency. |
| 1 | 15:14 | 0 | 2 | Reserved |
| 2 | 2:00 | cfg_dev_contr ol_max_read_r eq | 3 | Configuration Device Control - Max_Read_Request_Size: Device_Ctrl[14:12]. This field sets the maximum Read Request size for the user logic as a Requester. The user logic must not generate Read Requests with size exceeding the set value. 000b = 128 bytes maximum Read Request size 001b = 256 bytes maximum Read Request size 010b = 512 bytes maximum Read Request size 011b = 1024 bytes maximum Read Request size 100b = 2048 bytes maximum Read Request size 101b = 4096 bytes maximum Read Request size |
| 2 | 3 | cfg_link_status _link_training | 1 | Configuration Link Status - Link Training: Link_Status[11]. Indicates that the Physical Layer LTSSM is in the Configuration or Recovery state, or that 1b was written to the Retrain Link bit but Link training has not yet begun. The core clears this bit when the LTSSM exits the Configuration/Recovery state. |
| 2 | 6:04 | cfg_link_status _current_speed | 3 | Configuration Link Status - Current Link Speed: Link_Status[1:0]. This field indicates the negotiated Link speed of the given PCI Express Link. 001b = 2.5 GT/s PCI Express Link 010b = 5.0 GT/s PCI Express Link 100b = 8.0 GT/s PCI Express Link |
| 2 | 10:07 | cfg_link_status _negotiated_wi dth | 4 | Configuration Link Status - Negotiated Link Width: Link_Status[7:4]. This field indicates the negotiated width of the given PCI Express Link (only widths up to x8 displayed). 0001b = x1 0010b = x2 0100b = x4 1000b = x8 |

| Table 2-18: | Detailed Function Status Interface Port Descriptions |
|-------------|--|
|-------------|--|

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|---|-------|---|
| 2 | 11 | cfg_link_status _bandwidth_st atus | 1 | Configuration Link Status - Link Bandwidth Management Status: Link_Status[14]. Indicates that either of the following has occurred without the Port transitioning through DL_Down status: A Link retraining has completed following a write of 1b to the Retrain Link bit. Note: This bit is Set following any write of 1b to the Retrain Link bit, including when the Link is in the process of retraining for some other reason. Hardware has changed Link speed or width to attempt to correct unreliable Link operation, either through an LTSSM timeout or a higher level process. This bit is set if the Physical Layer reports a speed or width change was initiated by the Downstream component that was not indicated as an autonomous change. |
| 2 | 12 | cfg_link_status _auto_bandwid th_status | 1 | Configuration Link Status - Link Autonomous Bandwidth Status: Link_Status[15]. Indicates the core has autonomously changed Link speed or width, without the Port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable Link operation. This bit must be set if the Physical Layer reports a speed or width change was initiated by the Downstream component that was indicated as an autonomous change. |
| 2 | 15:13 | 0 | 3 | Reserved |
| 3 | 1:00 | cfg_link_contr ol_aspm_contr ol | 2 | Configuration Link Control - ASPM Control: Link_Ctrl[1:0]. Indicates the level of ASPM supported, where: 00b = Disabled 01b = L0s Entry Enabled 10b = L1 Entry Enabled 11b = L0s and L1 Entry Enabled |
| 3 | 2 | cfg_link_contr ol_rcb | 1 | Configuration Link Control - RCB: Link_Ctrl[3]. Indicates the Read Completion Boundary value, where 0=64B, and 1=128B. |
| 3 | 3 | cfg_link_contr ol_link_disable | 1 | Configuration Link Control - Link Disable: Link_Ctrl[4]. When asserted, indicates the Link is disabled and directs the LTSSM to the Disabled state. |

| Table 2-18: | Detailed Function Status Interface Port Descriptions |
|-------------|--|
|-------------|--|

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|--|-------|---|
| 3 | 4 | cfg_link_contr ol_common_cl ock | 1 | Configuration Link Control - Common Clock Configuration: Link_Ctrl[6]. When asserted, indicates that this component and the component at the opposite end of this Link are operating with a distributed common reference clock. When deasserted, indicates they are operating with an asynchronous reference clock. |
| 3 | 5 | cfg_link_contr ol_extended_s ync | 1 | Configuration Link Control - Extended Synch: Link_Ctrl[7]. When asserted, forces the transmission of additional Ordered Sets when exiting the LOs state and when in the Recovery state. |
| 3 | 6 | cfg_link_contr ol_clock_pm_e n | 1 | Configuration Link Control - Enable Clock Power Management: Link_Ctrl[8]. For Upstream Ports that support a CLKREQ# mechanism, indicates: 0b = Clock power management disabled. 1b = The device is permitted to use CLKREQ#. The core will take no action based on the setting of this bit; external logic must implement this. |
| 3 | 7 | cfg_link_contr ol_hw_auto_wi dth_dis | 1 | Configuration Link Control - Hardware Autonomous Width Disable: Link_Ctrl[9]. When asserted, disables the core from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. |
| 3 | 8 | cfg_link_contr ol_bandwidth_i nt_en | 1 | Configuration Link Control - Link Bandwidth Management Interrupt Enable: Link_Ctrl[10]. When asserted, enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been set. The core will take no action based on the setting of this bit; user logic must create the interrupt. |
| 3 | 9 | cfg_link_contr ol_auto_bandw idth_int_en | 1 | Configuration Link Control - Link Autonomous Bandwidth Interrupt Enable: Link_Ctrl[11]. When asserted, this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been set. The core will take no action based on the setting of this bit; user logic must create the interrupt. |
| 3 | 10 | cfg_tph_reque ster_enable | 1 | TPH Requester Enable: bit [8] of the TPH Requester Control Register in the TPH Requester Capability Structure of the function. These bits are active only in the Endpoint mode. Indicates whether the software has enabled the device to generate requests with TPH Hints from the associated Function. |

Table 2-18: Detailed Function Status Interface Port Descriptions

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|--|-------|--|
| 3 | 13:11 | cfg_tph_steeri ng_tag_mode | 3 | TPH Steering Tag Mode: Reflect the setting of the ST Mode Select bits in the TPH Requester Control Register. These bits are active only in the Endpoint mode. They indicate the allowed modes for generation of TPH Hints by the corresponding Function. |
| 3 | 15:14 | 0 | 2 | Reserved |
| 4 | 3:00 | cfg_dev_contr ol2_cpl_timeou t_val | 4 | Configuration Device Control 2 - Completion Timeout Value: Device_Ctrl2[3:0]. This is the time range that the user logic should regard a Request's pending Completion as a Completion Timeout. The core will take no action based on this setting. 0000b = 50 μ s to 50 ms (default) 0001b = 50 μ s to 50 ms (default) 0010b = 1 ms to 100 μ s 0101b = 16 ms to 55 ms 0110b = 65 ms to 210 ms 1001b = 260 ms to 900 ms 1010b = 1 s to 3.5 s 1110b = 4 s to 13 s 1110b = 17 s to 64 s |
| 4 | 4 | cfg_dev_contr ol2_cpl_timeou t_dis | 1 | Configuration Device Control 2 - Completion Timeout Disable: Device_Ctrl2[4]. This should cause the user to disable their Completion Timeout counters. |
| 4 | 5 | cfg_dev_contr ol2_atomic_re quester_en | 1 | Configuration Device Control 2 - Atomic Requester Enable: Device_Ctrl2[6]. Applicable only to Endpoints and Root Ports; must be hardwired to 0b for other Function types. The Function is allowed to initiate AtomicOp Requests only if this bit and the Bus Master Enable bit in the Command register are both Set. This bit is required to be RW if the Endpoint or Root Port is capable of initiating AtomicOp Requests, but otherwise is permitted to be hardwired to 0b. This bit does not serve as a capability bit. This bit is permitted to be RW even if no AtomicOp Requester capabilities are supported by the Endpoint or Root Port. Default value of this bit is 0b. 32 nm |

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|--|-------|---|
| 4 | 6 | cfg_dev_contr ol2_ido_req_en | 1 | Configuration Device Control 2 - IDO Request Enable: Device_Ctrl2[8]. If this bit is Set, the Function is permitted to set the ID-Based Ordering (IDO) bit (Attribute[2]) of Requests it initiates (see Section 2.2.6.3 and Section 2.4). Endpoints, including RC Integrated Endpoints, and Root Ports are permitted to implement this capability. A Function is permitted to hardwire this bit to 0b if it never sets the IDO attribute in Requests. Default value of this bit is 0b. 32 nm |
| 4 | 7 | cfg_dev_contr ol2_ido_cpl_en | 1 | Configuration Device Control 2 - IDO Completion Enable: Device_Ctrl2[9]. If this bit is Set, the Function is permitted to set the ID-Based Ordering (IDO) bit (Attribute[2]) of Completions it returns (see Section 2.2.6.3 and Section 2.4). Endpoints, including RC Integrated Endpoints, and Root Ports are permitted to implement this capability. A Function is permitted to hardwire this bit to 0b if it never sets the IDO attribute in Requests. Default value of this bit is 0b. 32 nm |
| 4 | 8 | cfg_dev_contr ol2_ltr_en | 1 | Configuration Device Control 2 - LTR Mechanism Enable: Device_Ctrl2[10]. If this bit is Set, the Function is permitted to set the ID-Based Ordering (IDO) bit (Attribute[2]) of Completions it returns (see Section 2.2.6.3 and Section 2.4). Endpoints, including RC Integrated Endpoints, and Root Ports are permitted to implement this capability. A Function is permitted to hardwire this bit to 0b if it never sets the IDO attribute in Requests. Default value of this bit is 0b. 32 nm |
| 4 | 13:09 | cfg_dpa_subst ate | 5 | Dynamic Power Allocation Substate: Reflect the setting of the Dynamic Power Allocation Substate field in the DPA Control Register. |
| 4 | 15:14 | 0 | 1 | Reserved |
| 5 | 0 | cfg_root_contr ol_syserr_corr_ err_en | 1 | Configuration Root Control - System Error on Correctable Error Enable: Root_Control[0]. This bit enables the user logic to generate a System Error for reported Correctable Errors. |
| 5 | 1 | cfg_root_contr ol_syserr_non_ fatal_err_en | 1 | Configuration Root Control - System Error on Non-Fatal Error Enable: Root_Control[1]. This bit enables the user logic to generate a System Error for reported Non-Fatal Errors. |
| 5 | 2 | cfg_root_contr ol_syserr_fatal_ err_en | 1 | Configuration Root Control - System Error on Fatal Error Enable: Root_Control[2]. This bit enables the user logic to generate a System Error for reported Fatal Errors. |

| Table 2-18: | Detailed Function Status Interface Port Descriptions |
|-------------|--|
|-------------|--|

| cfg_per_func_ status_control [bit] | cfg_per_func_ status_data [bit/slice] | Status Output | Width | Description |
|--|---|--|-------|---|
| 5 | 3 | cfg_root_contr ol_pme_int_en | 1 | Configuration Root Control - PME Interrupt Enable: Root_Control[3]. This bit enables the user logic to generate an Interrupt for received PME Messages. |
| 5 | 4 | cfg_aer_rooter r_corr_err_repo rting_en | 1 | Configuration AER - Correctable Error Reporting Enable: AER_Root_Error_Command[0]. This bit enables the user logic to generate interrupts for reported Correctable Errors. |
| 5 | 5 | cfg_aer_rooter r_non_fatal_err _reporting_en | 1 | Configuration AER - Non Fatal Error Reporting Enable: AER_Root_Error_Command[1]. This bit enables the user logic to generate interrupts for reported Non-Fatal Errors. |
| 5 | 6 | cfg_aer_rooter r_fatal_err_rep orting_en | 1 | Configuration AER - Fatal Error Reporting Enable: AER_Root_Error_Command[2]. This bit enables the user logic to generate interrupts for reported Fatal Errors. |
| 5 | 7 | cfg_aer_rooter r_corr_err_rece ived | 1 | Configuration AER - Correctable Error Messages Received: AER_Root_Error_Status[0]. Indicates that an ERR_COR Message was received. |
| 5 | 8 | cfg_aer_rooter r_non_fatal_err _received | 1 | Configuration AER - Non-Fatal Error Messages Received: AER_Root_Error_Status[5]. Indicates that an ERR_NFE Message was received. |
| 5 | 9 | cfg_aer_rooter r_fatal_err_rece ived | 1 | Configuration AER - Fatal Error Messages Received: AER_Root_Error_Status[6]. Indicates that an ERR_FATAL Message was received. |
| 5 | 15:10 | 0 | 6 | Reserved |

Table 2-18: Detailed Function Status Interface Port Descriptions

Configuration Control Interface

Table 2-19 defines the ports in the Configuration Control interface of the Integrated Block for PCI Express core.

Table 2-19: Configuration Control Interface Port Descriptions

| Port | Direction | Description | Width |
|-------------------|-----------|--|-------|
| cfg_hot_reset_in | Input | Configuration Hot Reset In. In RP mode, assertion transitions LTSSM to hot reset state, active High. | 1 |
| cfg_hot_reset_out | Output | Configuration Hot Reset Out. In EP mode, assertion indicates that EP has transitioned to the hot reset state, active High. | 1 |

| Port | Direction | Description | Width |
|---------------------------------|-----------|---|-------|
| cfg_config_space_enable | Input | Configuration Configuration Space Enable. When this input is set to 0 in the Endpoint mode, the core generates a CRS Completion in response to Configuration Requests. This port should be held deasserted when the core configuration registers are loaded from the DRP due to a change in attributes. This prevents the core from responding to Configuration Requests before all the registers are loaded. This input can be High when the power-on default values of the Configuration Registers do not need to be modified before Configuration space enumeration. This input is not applicable for Root Port mode. | 1 |
| cfg_per_function_update_done | Output | Configuration per Function Update Complete. Asserted in response to cfg_per_function_output_request assertion, for one cycle after the request is complete. | 1 |
| cfg_per_function_number | Input | Configuration Per Function Target Function Number. The user provides the function number (0-7), where value 0–1 corresponds to PF0–1, and value 2–7 corresponds to VF0–5, and asserts cfg_per_function_output_request to obtain per function output values for the selected function. | 3 |
| cfg_per_function_output_request | Input | Configuration Per Function Output Request. When this port is asserted with a function number value on cfg_per_function_number, the core presents information on per-function configuration output pins and asserts cfg_update_done when complete. | 1 |
| cfg_dsn | Input | Configuration Device Serial Number. Indicates the value that should be transferred to the Device Serial Number Capability on PF0. Bits [31:0] are transferred to the first (Lower) Dword (byte offset 0x4h of the Capability), and bits [63:32] are transferred to the second (Upper) Dword (byte offset 0x8h of the Capability). If this value is not statically assigned, the user must pulse user_cfg_input_update after it is stable. | 64 |
| cfg_ds_bus_number | Input | Configuration Downstream Bus Number. Downstream Port Provides the bus number portion of the Requester ID (RID) of the Downstream Port. This is used in TLPs generated inside the core, such as UR Completions and Power-management messages; it does not affect TLPs presented on the TRN interface. Upstream Port No role. | 8 |

Table 2-19: Configuration Control Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|----------------------------------|-----------|---|-------|
| cfg_ds_device_number | Input | Configuration Downstream Device Number: Downstream Port Provides the device number portion of the RID of the Downstream Port. This is used in TLPs generated inside the core, such as UR Completions and Power-management messages; it does not affect TLPs presented on the TRN interface. Upstream Port No role. | 5 |
| cfg_ds_function_number | Input | Configuration Downstream Function Number. Downstream Port Provides the function number portion of the RID of the Downstream Port. This is used in TLPs generated inside the core, such as UR Completions and Power-management messages; it does not affect TLPs presented on the TRN interface. Upstream Port No role. | 3 |
| cfg_power_state_change_ack | Input | Configuration Power State Ack. The user must assert this input to the core for one cycle in response to the assertion of cfg_power_state_change_interrupt, when it is ready to transition to the low-power state requested by the configuration write request. The client can permanently hold this input High if it does not need to delay the return of the completions for the configuration write transactions, causing power-state changes. | 1 |
| cfg_power_state_change_interrupt | Output | Power State Change Interrupt. The core asserts this output when the power state of a Physical or Virtual Function is being changed to the D1 or D3 states by a write into its Power Management Control Register. The core holds this output High until the user asserts the cfg_power_state_change_ack input to the core. While cfg_power_state_change_interrupt remains High, the core does not return completions for any pending configuration read or write transaction received by the core. The purpose is to delay the completion for the configuration write transaction that caused the state change until the user is ready to transition to the low-power state. When cfg_power_state_change_interrupt is asserted, the Function number associated with the configuration write transaction is provided on the cfg_snp_function_number[7:0] output. When the client asserts cfg_power_state_change_ack, the new state of the Function that underwent the state change is reflected on cfg_function_power_state (for PFs) or the cfg_vf_power_state (for VFs) outputs of the core. | 1 |

Table 2-19: Configuration Control Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-----------------------|-----------|---|-------|
| cfg_err_cor_in | Input | Correctable Error Detected: The user can activate this input for one cycle to indicate a correctable error detected within the user logic that needs to be reported as an internal error through the PCI Express Advanced Error Reporting mechanism. In response, the core sets the Corrected Internal Error Status bit in the AER Correctable Error Status Register of all enabled Functions, and also sends an error message if enabled to do so. This error is not considered Function-specific. | 1 |
| cfg_err_uncor_in | Input | Uncorrectable Error Detected. The user can activate this input for one cycle to indicate a uncorrectable error detected within the user logic that needs to be reported as an internal error through the PCI Express Advanced Error Reporting mechanism. In response, the core sets the uncorrected Internal Error Status bit in the AER Uncorrectable Error Status Register of all enabled Functions, and also sends an error message if enabled to do so. This error is not considered Function-specific. | 1 |
| cfg_flr_done | Input | Function Level Reset Complete. The user must assert bit <i>i</i> of this bus when the reset operation of Function <i>i</i> completes. This causes the core to deassert cfg_flr_in_process for Function <i>i</i> and to re-enable configuration accesses to the Function. | 2 |
| cfg_vf_flr_done | Input | Function Level Reset for virtual Function is Complete. The user must assert bit <i>i</i> of this bus the reset operation of Virtual Function <i>i</i> completes. This causes the core to deassert cfg_vf_flr_in_process for Function <i>i</i> and to re-enable configuration accesses to the Virtual Function. | 6 |
| cfg_flr_in_process | Output | Function Level Reset In Process. The core asserts bit <i>i</i> of this bus when the host initiates a reset of Function <i>i</i> through its FLR bit in the configuration space. The core continues to hold the output High until the user sets the corresponding cfg_flr_done input for the corresponding Function to indicate the completion of the reset operation. | 2 |
| cfg_vf_flr_in_process | Output | Function Level Reset In Process for Virtual Function. The core asserts bit <i>i</i> of this bus when the host initiates a reset of Virtual Function <i>i</i> though its FLR bit in the configuration space. The core continues to hold the output High until the user sets the corresponding $cfg_vf_flr_done$ input for the corresponding Function to indicate the completion of the reset operation. | 6 |

Table 2-19: Configuration Control Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|---------------------------------|-----------|--|-------|
| cfg_req_pm_transition_l23_ready | Input | When the core is configured as an Endpoint, the user can assert this input to transition the power management state of the core to L23_READY (see Chapter 5 of the <i>PCI Express Specification</i> for a detailed description of power management). This is done after the PCI Functions in the core are placed in the D3 state and after the client acknowledges the PME_Turn_Off message from the Root Complex. Asserting this input causes the link to transition to the L2 state, and requires a hard reset to resume operation. This input can be hardwired to 0 if the link is not required to transition to L2. This input is not used in Root Complex mode. | 1 |
| cfg_link_training_enable | Input | This input must be set to 1 to enable the Link Training Status State Machine (LTSSM) to bring up the link. Setting it to 0 forces the LTSSM to stay in the Detect.Quiet state. | 1 |

Table 2-19: Configuration Control Interface Port Descriptions (Cont'd)

Configuration Interrupt Controller Interface

Table 2-20 defines the ports in the Configuration Interrupt Controller interface of the Integrated Block for PCI Express core.

| Table 2-20: | Configuration Interrupt Controller Interface Port Descriptions |
|-------------|--|
|-------------|--|

| Port | Direction | Description | Width |
|--------------------------|-----------|---|-------|
| cfg_interrupt_int | Input | Configuration INTx Vector. When the core is configured as EP, these four inputs are used by the client application to signal an interrupt from any of its PCI Functions to the RC using the Legacy PCI Express Interrupt Delivery mechanism of PCI Express. These four inputs correspond to INTA, INTB, INTC, and INTD. Asserting one of these signals causes the core to send out an Assert_INTx message, and deasserting the signal causes the core to transmit a Deassert_INTx message. | 4 |
| cfg_interrupt_sent | Output | Configuration INTx Sent. A pulse on this output indicates that the core has sent an INTx Assert or Deassert message in response to a change in the state of one of the cfg_interrupt_int inputs. | 1 |
| cfg_interrupt_pending | Input | Configuration INTx Interrupt Pending (active High). Per Function indication of a pending interrupt. cfg_interrupt_pending[0] corresponds to Function #0. | 2 |
| cfg_interrupt_msi_enable | Output | Configuration Interrupt MSI Function Enabled. Indicates that Message Signaling Interrupt (MSI) messaging is enabled per Function. | 2 |

| Port | Direction | Description | Width |
|----------------------------------|-----------|--|-------|
| cfg_interrupt_msi_vf_enable | Output | Configuration Interrupt MSI on VF Enabled. Indicates that MSI messaging is enabled, per Virtual Function. | 6 |
| cfg_interrupt_msi_int | Input | Configuration Interrupt MSI Vector. When the core is configured in the Endpoint mode to support MSI interrupts, these inputs are used to signal the 32 distinct interrupt conditions associated with a PCI Function (Physical or Virtual) from the user logic to the core. The Function number must be specified on the cfg_interrupt_msi_function_number input. After placing the Function number on the input cfg_interrupt_msi_function_number, the user logic must activate one of these signals for one cycle to transmit an interrupt. The user logic must not activate more than one of the 32 interrupt inputs in the same cycle. The core internally registers the interrupt condition on the 0-to-1 transition of any bit in cfg_interrupt_msi_int. After asserting an interrupt, the user logic must wait for the cfg_interrupt_msi_sent or cfg_interrupt_msi_fail indication from the core before asserting a new interrupt. | 32 |
| cfg_interrupt_msi_sent | Output | Configuration Interrupt MSI Interrupt Sent. The core generates a one-cycle pulse on this output to signal that an MSI interrupt message has been transmitted on the link. The user logic must wait for this pulse before signaling another interrupt condition to the core. | 1 |
| cfg_interrupt_msi_fail | Output | Configuration Interrupt MSI Interrupt Operation Failed. A one-cycle pulse on this output indicates that an MSI interrupt message was aborted before transmission on the link. The client must retransmit the MSI interrupt in this case. | 1 |
| cfg_interrupt_msi_mmenable | Output | Configuration Interrupt MSI Function Multiple Message Enable. When the core is configured in the Endpoint mode to support MSI interrupts, these outputs are driven by the "Multiple Message Enable" bits of the MSI Control Registers associated with Physical Functions. These bits encode the number of allocated MSI interrupt vectors for the corresponding Function. Bits [2:0] correspond to Physical Function 0. | 6 |
| cfg_interrupt_msi_pending_status | Input | Configuration Interrupt MSI Pending Status. These inputs provide the status of the MSI pending interrupts for the Physical Functions. The setting of these pins determines the value read from the MSI Pending Bits Register of the corresponding PF. Bits [31:0] belong to PF 0, bits [63:32] to PF 1. | 64 |
| cfg_interrupt_msi_mask_update | Output | Configuration Interrupt MSI Function Mask Updated. Asserted for one cycle when any enabled functions in the MSI Mask Register change value. | 1 |

| Table 2-20: | Configuration Interrupt Cont | roller Interface Port Descriptions (Cont'd) | |
|-------------|------------------------------|---|--|
| 10010 - 201 | | | |

| Port | Direction | Description | Width |
|------------------------------|-----------|--|-------|
| cfg_interrupt_msi_select | Input | Configuration Interrupt MSI Select. Values 0000b-0001b correspond to PF0-1 selection, and values 0010b-0111b correspond to VF0-5 selection. cfg_interrupt_msi_data[31:0] presents the value of the MSI Mask register from the selected function. When this input is driven to 1111b, cfg_interrupt_msi_data[17:0] presents the "Multiple Message Enable" bits of the MSI Control Registers associated with all Virtual Functions. These bits encode the number of allocated MSI interrupt vectors for the corresponding Function. cfg_interrupt_msi_data[2:0] correspond to Virtual Function 0, and so on. | 4 |
| cfg_interrupt_msi_data | Output | Configuration Interrupt MSI Data. The value presented depends on cfg_interrupt_msi_select. | 32 |
| cfg_interrupt_msix_enable | Output | Configuration Interrupt MSI-X Function Enabled. When asserted, indicates that the Message Signaling Interrupt (MSI-X) messaging is enabled, per Function. | 2 |
| cfg_interrupt_msix_mask | Output | Configuration Interrupt MSI-X Function Mask. Indicates the state of the Function Mask bit in the MSI-X Message Control field, per Function. | 2 |
| cfg_interrupt_msix_vf_enable | Output | Configuration Interrupt MSI-X on VF Enabled. When asserted, indicates that Message Signaling Interrupt (MSI-X) messaging is enabled, per Virtual Function. | 6 |
| cfg_interrupt_msix_vf_mask | Output | Configuration Interrupt MSI-X VF Mask. Indicates the state of the Function Mask bit in the MSI-X Message Control field, per Virtual Function. | 6 |
| cfg_interrupt_msix_address | Input | Configuration Interrupt MSI-X Address. When the core is configured to support MSI-X interrupts, this bus is used by the client logic to communicate the address to be used for an MSI-X message. | 64 |
| cfg_interrupt_msix_data | Input | Configuration Interrupt MSI-X Data. When the core is configured to support MSI-X interrupts, this bus is used by the client logic to communicate the data to be used for an MSI-X message. | 32 |

Table 2-20: Configuration Interrupt Controller Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-------------------------------|-----------|---|-------|
| cfg_interrupt_msix_int | Input | Configuration Interrupt MSI-X Data Valid. This signal indicates that valid information has been placed on the cfg_interrupt_msix_address[63:0] and cfg_interrupt_msix_data[31:0] buses, and the originating Function number has been placed on cfg_interrupt_function_number[3:0]. The core internally registers the associated address and data from cfg_interrupt_msix_address and cfg_interrupt_msix_data on the 0-to-1 transition of this valid signal. After asserting an interrupt, the client logic must wait for the cfg_interrupt_msix_fail indication from the core before asserting a new interrupt. | 1 |
| cfg_interrupt_msix_sent | Output | Configuration Interrupt MSI-X Interrupt Sent. The core generates a one-cycle pulse on this output to indicate that it has accepted the information placed on the cfg_interrupt_msix_address[63:0] and cfg_interrupt_msix_data[31:0] buses, and an MSI-X interrupt message has been transmitted on the link. The user application must wait for this pulse before signaling another interrupt condition to the core. | 1 |
| cfg_interrupt_msix_fail | Output | Configuration Interrupt MSI-X Interrupt Operation Failed. A one-cycle pulse on this output indicates that the interrupt controller has failed to transmit MSI-X interrupt on the link. The client must retransmit the MSI-X interrupt in this case. | 1 |
| cfg_interrupt_msi_attr | Input | Configuration Interrupt MSI/MSI-X TLP Attr. These bits provide the setting of the Attribute bits to be used for the MSI/MSI-X interrupt request. Bit 0 is the No Snoop bit, and bit 1 is the Relaxed Ordering bit. Bit 2 is the ID-Based Ordering bit. The core samples these bits on a 0-to-1 transition on cfg_interrupt_msi_int or cfg_interrupt_msix_int. | 3 |
| cfg_interrupt_msi_tph_present | Input | Configuration Interrupt MSI/MSI-X TPH Present. Indicates the presence of a Transaction Processing Hint (TPH) in the MSI/MSI-X interrupt request. The user application must set this bit while asserting cfg_interrupt_msi_int or cfg_interrupt_msix_int, if it includes a TPH in the MSI or MSI-X transaction. | 1 |
| cfg_interrupt_msi_tph_type | Input | Configuration Interrupt MSI/MSI-X TPH Type: When cfg_interrupt_msi_tph_present is 1 'b1, these two bits supply the two-bit type associated with the hint. The core samples these bits on a 0-to-1 transition on cfg_interrupt_msi_int or cfg_interrupt_msix_int. | 2 |

Table 2-20: Configuration Interrupt Controller Interface Port Descriptions (Cont'd)

| Port | Direction | Description | Width |
|-----------------------------------|-----------|--|-------|
| cfg_interrupt_msi_tph_st_tag | Input | Configuration Interrupt MSI/MSI-X TPH Steering Tag. When cfg_interrupt_msi_tph_present is 1 'b1, the Steering Tag associated with the Hint must be placed on cfg_interrupt_msi_tph_st_tag[7:0]. Setting cfg_interrupt_msi_tph_st_tag[8] to 1b activates the Indirect Tag mode. In the Indirect Tag mode, the core uses bits [5:0] of cfg_interrupt_msi_tph_st_tag as an index into its Steering Tag Table (STT) in the TPH Capability Structure (STT is limited to 64 entries per Function), and inserts the tag from this location in the transmitted request MSI/X TLP. Setting cfg_interrupt_msi_tph_st_tag[8] to 0b activates the Direct Tag mode. In the Direct Tag mode, the core inserts cfg_interrupt_msi_tph_st_tag[7:0] directly as the Tag in the transmitted MSI/X TLP. The core samples these bits on a 0-to-1 transition on any cfg_interrupt_msi_int bits or cfg_interrupt_msi_int bits or cfg_interrupt_msix_int. | 9 |
| cfg_interrupt_msi_function_number | Input | Configuration MSI/MSI-X Initiating Function. Indicates the Endpoint function number initiating the MSI or MSI-X transaction: • 0: PF0 • 1: PF1 • 2: VF0 • 3: VF1 • 4: VF2 • 7: VF5 | 3 |

Table 2-20: Configuration Interrupt Controller Interface Port Descriptions (Cont'd)

Configuration Extend Interface

Table 2-21 defines the ports in the Configuration Extend interface of the Integrated Block for PCI Express core.

| Port | Direction | Description | Width |
|---------------------------|-----------|---|-------|
| cfg_ext_read_received | Output | Configuration Extend Read Received. The core asserts this output when it has received a configuration read request from the link. When neither user-implemented legacy or extended configuration space is enabled, receipt of a configuration read results in a one-cycle assertion of this signal, together with valid cfg_ext_register_number and cfg_ext_function_number. When user-implemented legacy, extended configuration space, or both are enabled, for the cfg_ext_register_number ranges, 0x10-0x1f or 0x100-0x3ff, respectively, this signal is asserted, until user logic presents cfg_ext_read_data and cfg_ext_register_number ranges outside 0x10-0x1f or 0x100-0x3ff, receipt of a configuration read always results in a one-cycle assertion of this signal. | 1 |
| cfg_ext_write_received | Output | Configuration Extend Write Received. The core generates a one-cycle pulse on this output when it has received a configuration write request from the link. | 1 |
| cfg_ext_register_number | Output | Configuration Extend Register Number. The 10-bit address of the configuration register being read or written. The data is valid when cfg_ext_read_received or cfg_ext_write_received is High. | 10 |
| cfg_ext_function_number | Output | Configuration Extend Function Number. The 8-bit Function Number corresponding to the configuration read or write request. The data is valid when cfg_ext_read_received or cfg_ext_write_received is High. | 8 |
| cfg_ext_write_data | Output | Configuration Extend Write Data. Data being written into a configuration register. This output is valid when cfg_snp_write_received is High. | 32 |
| cfg_ext_write_byte_enable | Output | Configuration Extend Write Byte Enable. Byte enables for a configuration write transaction. | 4 |
| cfg_ext_read_data | Input | Configuration Extend Read Data. The user can provide data from an externally implemented configuration register to the core through this bus. The core samples this data on the next positive edge of the clock after it sets cfg_snp_read_received High, if the user has set cfg_snp_read_data_valid. | 32 |
| cfg_ext_read_data_valid | Input | Configuration Extend Read Data Valid. The user asserts this input to the core to supply data from an externally implemented configuration register. The core samples this input data on the next positive edge of the clock after it sets cfg_snp_read_received High. | 1 |

Table 2-21: Configuration Extend Interface Port Descriptions

Clock and Reset Interface

Table 2-22 defines the ports in the Clock and Reset interface of the Integrated Block for PCI Express core.

| Table 2-22: | Clock and Reset Interface Port Descriptions |
|-------------|---|
|-------------|---|

| Port | Direction | Description | Width |
|------------|-----------|--|-------|
| user_clk | Output | User clock output (62.5, 125, or 250 MHz). This clock has a fixed frequency and is configured in the CORE Generator™ software. | 1 |
| user_reset | Output | This signal is deasserted synchronously with respect to user_clk. It is deasserted and asserted asynchronously with sys_reset assertion. | 1 |
| sys_clk | Input | Reference clock. This clock has a selectable frequency of 100 MHz, 125 MHz, or 250 MHz. | 1 |
| sys_reset | Input | Fundamental reset input to the core (asynchronous, active-High). | 1 |

PCI Express Interface

The PCI Express (PCI_EXP) interface consists of differential transmit and receive pairs organized in multiple lanes. A PCI Express lane consists of a pair of transmit differential signals (pci_exp_txp, pci_exp_txn) and a pair of receive differential signals {pci_exp_rxp, pci_exp_rxn}. The 1-lane core supports only Lane 0, the 2-lane core supports lanes 0–1, the 4-lane core supports lanes 0-3, and the 8-lane core supports lanes 0–7. Transmit and receive signals of the PCI_EXP interface are defined in Table 2-23.

| Lane Number | Name | Direction | Description |
|----------------|--------------|-----------|---|
| 1-Lane Co | res | | |
| 0 | pci_exp_txp0 | Output | PCI Express Transmit Positive: Serial Differential Output 0 (+) |
| | pci_exp_txn0 | Output | PCI Express Transmit Negative: Serial Differential Output 0 (–) |
| | pci_exp_rxp0 | Input | PCI Express Receive Positive: Serial Differential Input 0 (+) |
| | pci_exp_rxn0 | Input | PCI Express Receive Negative: Serial Differential Input 0 (-) |
| 2-Lane Co | res | | |
| 0 | pci_exp_txp0 | Output | PCI Express Transmit Positive: Serial Differential Output 0 (+) |
| | pci_exp_txn0 | Output | PCI Express Transmit Negative: Serial Differential Output 0 (–) |
| | pci_exp_rxp0 | Input | PCI Express Receive Positive: Serial Differential Input 0 (+) |
| | pci_exp_rxn0 | Input | PCI Express Receive Negative: Serial Differential Input 0 (-) |

Table 2-23: PCI Express Interface Signals for 1-, 2-, 4- and 8-Lane Cores

| Lane Number | Name | Direction | Description |
|----------------|--------------|-----------|---|
| 1 | pci_exp_txp1 | Output | PCI Express Transmit Positive: Serial Differential Output 1 (+) |
| | pci_exp_txn1 | Output | PCI Express Transmit Negative: Serial Differential Output 1 (-) |
| | pci_exp_rxp1 | Input | PCI Express Receive Positive: Serial Differential Input 1 (+) |
| | pci_exp_rxn1 | Input | PCI Express Receive Negative: Serial Differential Input 1 (-) |
| -Lane Cor | res | <u>.</u> | • |
| 0 | pci_exp_txp0 | Output | PCI Express Transmit Positive: Serial Differential Output 0 (+) |
| | pci_exp_txn0 | Output | PCI Express Transmit Negative: Serial Differential Output 0 (–) |
| | pci_exp_rxp0 | Input | PCI Express Receive Positive: Serial Differential Input 0 (+) |
| | pci_exp_rxn0 | Input | PCI Express Receive Negative: Serial Differential Input 0 (-) |
| 1 | pci_exp_txp1 | Output | PCI Express Transmit Positive: Serial Differential Output 1 (+) |
| | pci_exp_txn1 | Output | PCI Express Transmit Negative: Serial Differential Output 1 (-) |
| | pci_exp_rxp1 | Input | PCI Express Receive Positive: Serial Differential Input 1 (+) |
| | pci_exp_rxn1 | Input | PCI Express Receive Negative: Serial Differential Input 1 (-) |
| 2 | pci_exp_txp2 | Output | PCI Express Transmit Positive: Serial Differential Output 2 (+) |
| | pci_exp_txn2 | Output | PCI Express Transmit Negative: Serial Differential Output 2 (–) |
| | pci_exp_rxp2 | Input | PCI Express Receive Positive: Serial Differential Input 2 (+) |
| | pci_exp_rxn2 | Input | PCI Express Receive Negative: Serial Differential Input 2 (-) |
| 3 | pci_exp_txp3 | Output | PCI Express Transmit Positive: Serial Differential Output 3 (+) |
| | pci_exp_txn3 | Output | PCI Express Transmit Negative: Serial Differential Output 3 (–) |
| | pci_exp_rxp3 | Input | PCI Express Receive Positive: Serial Differential Input 3 (+) |
| | pci_exp_rxn3 | Input | PCI Express Receive Negative: Serial Differential Input 3 (-) |
| B-Lane Cor | res | • • | |
| 0 | pci_exp_txp0 | Output | PCI Express Transmit Positive: Serial Differential Output 0 (+) |
| | pci_exp_txn0 | Output | PCI Express Transmit Negative: Serial Differential Output 0 (–) |
| | pci_exp_rxp0 | Input | PCI Express Receive Positive: Serial Differential Input 0 (+) |
| | pci_exp_rxn0 | Input | PCI Express Receive Negative: Serial Differential Input 0 (-) |
| 1 | pci_exp_txp1 | Output | PCI Express Transmit Positive: Serial Differential Output 1 (+) |
| | pci_exp_txn1 | Output | PCI Express Transmit Negative: Serial Differential Output 1 (-) |
| | pci_exp_rxp1 | Input | PCI Express Receive Positive: Serial Differential Input 1 (+) |
| | pci_exp_rxn1 | Input | PCI Express Receive Negative: Serial Differential Input 1 (-) |
| 2 | pci_exp_txp2 | Output | PCI Express Transmit Positive: Serial Differential Output 2 (+) |
| | pci_exp_txn2 | Output | PCI Express Transmit Negative: Serial Differential Output 2 (–) |
| | pci_exp_rxp2 | Input | PCI Express Receive Positive: Serial Differential Input 2 (+) |
| | pci_exp_rxn2 | Input | PCI Express Receive Negative: Serial Differential Input 2 (-) |

Table 2-23: PCI Express Interface Signals for 1-, 2-, 4- and 8-Lane Cores (Cont'd)

| Lane Number | Name | Direction | Description |
|----------------|--------------|-----------|---|
| 3 | pci_exp_txp3 | Output | PCI Express Transmit Positive: Serial Differential Output 3 (+) |
| | pci_exp_txn3 | Output | PCI Express Transmit Negative: Serial Differential Output 3 (–) |
| | pci_exp_rxp3 | Input | PCI Express Receive Positive: Serial Differential Input 3 (+) |
| | pci_exp_rxn3 | Input | PCI Express Receive Negative: Serial Differential Input 3 (-) |
| 4 | pci_exp_txp4 | Output | PCI Express Transmit Positive: Serial Differential Output 4 (+) |
| | pci_exp_txn4 | Output | PCI Express Transmit Negative: Serial Differential Output 4 (–) |
| | pci_exp_rxp4 | Input | PCI Express Receive Positive: Serial Differential Input 4 (+) |
| | pci_exp_rxn4 | Input | PCI Express Receive Negative: Serial Differential Input 4 (–) |
| 5 | pci_exp_txp5 | Output | PCI Express Transmit Positive: Serial Differential Output 5 (+) |
| | pci_exp_txn5 | Output | PCI Express Transmit Negative: Serial Differential Output 5 (–) |
| | pci_exp_rxp5 | Input | PCI Express Receive Positive: Serial Differential Input 5 (+) |
| | pci_exp_rxn5 | Input | PCI Express Receive Negative: Serial Differential Input 5 (–) |
| 6 | pci_exp_txp6 | Output | PCI Express Transmit Positive: Serial Differential Output 6 (+) |
| | pci_exp_txn6 | Output | PCI Express Transmit Negative: Serial Differential Output 6 (–) |
| | pci_exp_rxp6 | Input | PCI Express Receive Positive: Serial Differential Input 6 (+) |
| | pci_exp_rxn6 | Input | PCI Express Receive Negative: Serial Differential Input 6 (–) |
| 7 | pci_exp_txp7 | Output | PCI Express Transmit Positive: Serial Differential Output 7 (+) |
| | pci_exp_txn7 | Output | PCI Express Transmit Negative: Serial Differential Output 7 (–) |
| | pci_exp_rxp7 | Input | PCI Express Receive Positive: Serial Differential Input 7 (+) |
| | pci_exp_rxn7 | Input | PCI Express Receive Negative: Serial Differential Input 7 (–) |

Table 2-23: PCI Express Interface Signals for 1-, 2-, 4- and 8-Lane Cores (Cont'd)

Attribute Descriptions

Client Interface

Table 2-24 lists the configuration attributes controlling the operation of the client interface of the Gen3 Integrated Block for PCIe.

| Attribute Name | Туре | Description |
|---------------------------------|------------|--|
| USER_CLK2_FREQ | Integer | 0: Disable User Clock 1: 31.25 MHz 2: 62.50 MHz (default) 3: 125.00 MHz 4: 250.00 MHz 5: 500.00 MHz |
| PL_LINK_CAP_MAX_LINK_SPEED[2:0] | Bit vector | Defines the maximum speed of the PCIe link. • 001: 2.5 GT/s • 010: 5.0 GT/s • 100: 8.0 GT/s • others: Reserved |
| PL_LINK_CAP_MAX_LINK_WIDTH[3:0] | Bit vector | Maximum Link Width. Valid settings are: 0001b: x1 0010b: x2 0100b: x4 1000b: x8 All other encodings are reserved. This setting is propagated to all layers in the core. |
| C_DATA_WIDTH | Integer | Configures the width of the AXI4-Stream interfaces.64 bit interface128 bit interface256 bit interface |
| AXISTEN_IF_CQ_ALIGNMENT_MODE | String | Defines the data alignment mode for the completer request interface.FALSE: Dword-aligned ModeTRUE: Address-aligned Mode |
| AXISTEN_IF_CC_ALIGNMENT_MODE | String | Defines the data alignment mode for the completer completion interface.FALSE: Dword-aligned ModeTRUE: Address-aligned Mode |
| AXISTEN_IF_RQ_ALIGNMENT_MODE | String | Defines the data alignment mode for the requester request interface.FALSE: Dword-aligned ModeTRUE: Address-aligned Mode |
| AXISTEN_IF_RC_ALIGNMENT_MODE | String | Defines the data alignment mode for the requester completion interface.FALSE: Dword-aligned ModeTRUE: Address-aligned Mode |
| AXISTEN_IF_RC_STRADDLE | String | This attribute enables the straddle option on the requester completion interface.FALSE: Straddle option disabledTRUE: Straddle option enabled |

| Table 2-24: | Configuration | Attributes of the | Integrated Bloc | k Client Interface |
|-------------|---------------|-------------------|-----------------|--------------------|
|-------------|---------------|-------------------|-----------------|--------------------|

| Attribute Name | Туре | Description |
|--------------------------------|--------|---|
| AXISTEN_IF_RQ_PARITY_CHECK | String | This attribute enables parity checking on the requester request interface.FALSE: Parity check disabledTRUE: Parity check enabled |
| AXISTEN_IF_CC_PARITY_CHECK | String | This attribute enables parity checking on the completer completion interface.FALSE: Parity check disabledTRUE: Parity check enabled |
| AXISTEN_IF_ENABLE_RX_MSG_INTFC | String | This attribute controls how the core delivers a message received from the link. When this attribute is set to FALSE, the core delivers the received message TLPs on the completer request interface using the AXI4-Stream protocol. In this mode, the user can select the message types to receive using the AXISTEN_IF_ENABLE_MSG_ROUTE attributes. The receive message interface remains disabled in this mode. When this attribute is set to TRUE, the core internally decodes messages received from the link, and signals them to the user by activating the cfg_msg_received signal on the receive message interface. The core does not transfer any message TLPs on the completer request interface. The settings of the AXISTEN_ENABLE_MSG_ROUTE attributes have no effect on the operation of the receive message interface in this mode. |

Table 2-24: Configuration Attributes of the Integrated Block Client Interface (Cont'd)

| Attribute Name | Туре | Description |
|-----------------------------------|------------|--|
| AXISTEN_IF_ENABLE_MSG_ROUTE[17:0] | Bit vector | When the AXISTEN_IF_ENABLE_RX_MSG_INTFC attribute is set to 0, these attributes can be used to select the specific message types that the user wants to receive on the completer request interface. Setting a bit to 1 enables the delivery of the corresponding type of messages on the interface, and setting it to 0 results in the core filtering the message. Table 2-25 defines the attribute bit definitions corresponding to the various message types. |
| AXISTEN_IF_ENABLE_CLIENT_TAG | String | When this attribute is FALSE, tag management for Non-Posted transactions initiated from the requester request interface is performed by the Integrated Block. That is, for each Non-Posted request, the core allocates the tag for the transaction and communicates it to the client. Setting this attribute to TRUE disables the internal tag management, allowing the user to supply the tag to be used for each request. The user must present the Tag field in the Request descriptor header in the range 0–31 when the PF0_DEV_CAP_EXT_TAG_SUPPORTED attribute is FALSE, while the Tag field can be in the range 0–63 when the PF0_DEV_CAP_EXT_TAG_SUPPORTED attribute is TRUE. |

Table 2-24: Configuration Attributes of the Integrated Block Client Interface (Cont'd)

Table 2-25: AXISTEN_IF_ENABLE_MSG_ROUTE Attribute Bit Descriptions

| Bit Index | Message Type |
|-----------|------------------------------------|
| 0 | ERR_COR |
| 1 | ERR_NONFATAL |
| 2 | ERR_FATAL |
| 3 | Assert_INTA and Deassert_INTA |
| 4 | Assert_INTB and Deassert_INTB |
| 5 | Assert_INTC and Deassert_INTC |
| 6 | Assert_INTD and Deassert_INTD |
| 7 | PM_PME |
| 8 | PME_TO_Ack |
| 9 | PME_Turn_Off |
| 10 | PM_Active_State_Nak |
| 11 | Set_Slot_Power_Limit |
| 12 | Latency Tolerance Reporting (LTR) |
| 13 | Optimized Buffer Flush/Fill (OBFF) |
| 14 | Unlock |

| Bit Index | Message Type |
|-----------|---|
| 15 | Vendor_Defined Type 0 |
| 16 | Vendor_Defined Type 1 |
| 17 | Invalid Request, Invalid Completion, Page Request, PRG Response |

Table 2-25: AXISTEN_IF_ENABLE_MSG_ROUTE Attribute Bit Descriptions (Cont'd)

Configuration Space

The PCI configuration space consists of three primary parts, illustrated in Table 2-26. These include:

- Legacy PCI v3.0 Type 0/1 Configuration Space Header
 - Type 0 Configuration Space Header used by Endpoint applications (see Table 2-27)
 - Type 1 Configuration Space Header used by Root Port applications (see Table 2-28)
- Legacy Extended Capability Items
 - PCIe Capability Item
 - Power Management Capability Item
 - Message Signaled Interrupt (MSI) Capability Item
 - MSI-X Capability Item (optional)
- PCIe Capabilities
 - Advanced Error Reporting Extended Capability Structure (AER)
 - Alternate Requestor ID (ARI) (optional)
 - Device Serial Number Extended Capability Structure (DSN) (optional)
 - Power Budgeting Enhanced
 - Capability Header (PB) (optional)
 - Resizable BAR (RBAR) (optional)
 - Latency Tolerance Reporting (LTR) (optional)
 - Dynamic Power Allocation (DPA) (optional)
 - Single Root I/O Virtualization (SR-IOV) (optional)
 - Transaction Processing Hints (TPH) (optional)
 - Virtual Channel Extended Capability Structure (VC) (optional)
- PCIe Extended Capabilities

- Device Serial Number Extended Capability Structure (optional)
- Virtual Channel Extended Capability Structure (optional)
- Advanced Error Reporting Extended Capability Structure (optional)

The core implements up to four legacy extended capability items.

For more information about enabling this feature, see Chapter 7, Customizing and Generating the Core.

The core optionally implements up to ten PCI Express Extended Capabilities. The remaining PCI Express Extended Capability Space is available for users to implement. The starting address of the space available to users begins at 3DCh. If you choose to implement registers in this space, you can select the starting location of this space, and this space must be implemented in the User Application.

For more information about enabling this feature, see PCIe Extended Capabilities in Chapter 7.

| 31 | 16 | 15 | 0 | |
|------|------------|--------------------------------|-----------|---------------|
| Devi | ce ID | Vend | or ID | 000h |
| Sta | tus | Com | mand | 004h |
| | Class Code | | Rev ID | 008h |
| BIST | Header | Lat Timer | Cache Ln | 00Ch |
| | | | | 010h |
| | | | | 014h |
| | | | | 018h |
| | | | | 01Ch |
| | | pe Specific and Table 2-28) | | 020h |
| | | | | 024h |
| | | | | 028h |
| | | | | 02Ch |
| | | | | 030h |
| | | | CapPtr | 034h |
| | | | | 038h |
| | | Intr Pin | Intr Line | 03Ch |
| | Rese | erved | · | 040h- 07Ch |

Table 2-26:Common PCI Configuration Space Header

| 3 | 1 | 16 | 15 | | |
|------------------------------|---|--------------------------------|---------------|-----------|--|
| | PM Capabilit | у | NxtCap | РМ Сар | |
| | Data | Reserved | PMC | CSR | |
| | Reserved | | | | |
| Customizable ⁽¹⁾ | MSI Control | | NxtCap | MSI Cap | |
| | | Message Ad | dress (Lower) | | |
| | | Message Ad | dress (Upper) | | |
| | Reserved | | Messag | e Data | |
| | | Mas | k Bits | | |
| | | Pendi | ng Bits | | |
| | | Res | erved | | |
| Optional ⁽³⁾ | MSI-X Contro | bl | NxtCap | MSI-X Cap | |
| | Table Offset Table BIR | | | | |
| | PBA Offset PBA BIR | | | | |
| | Reserved | | | | |
| | PE Capability NxtCap PE Cap | | | | |
| | PCI Express Device Capabilities | | | | |
| | Device Statu | Device Status Device Control | | | |
| | PCI Express Link Capabilities | | | | |
| | Link Status | Link Status Link Control | | | |
| oot Port Only ⁽²⁾ | | Slot Ca | pabilities | | |
| | Slot Status | | Slot C | ontrol | |
| | Root Capabilit | Root Capabilities Root Control | | | |
| | Root Status | | | | |
| | PCI Express Device Capabilities 2 | | | | |
| | Device Status 2 Device Control 2 | | | ontrol 2 | |
| | PCI Express Link Capabilities 2 | | | | |
| | Link Status 2 | 2 | Link Co | ntrol 2 | |
| | Unimplemented Configuration Space (Returns 0x00000000) | | | | |

| | 31 | 16 | 15 | C | | |
|--|--|--------------|---|-------------------|--|--|
| Always Enabled | Next Cap | Cap. Ver. | PCI Express Exten | ded Cap. ID (AER) | | |
| | Uncorrectable Error Status Register | | | | | |
| | Uncorrectable Error Mask Register | | | | | |
| | Uncorrectable Error Severity Register | | | | | |
| | Correctable Error Status Register | | | | | |
| | Correctable Error Mask Register | | | | | |
| | Advanced Error Cap. & Control Register | | | | | |
| | Header Log Register 1 | | | | | |
| | Header Log Register 2 | | | | | |
| | Header Log Register 3 | | | | | |
| | Header Log Register 4 | | | | | |
| | | Res | erved | | | |
| Optional, Root Port only ⁽³⁾ | Root Error Command Register | | | | | |
| | Root Error Status Register | | | | | |
| | Error Source ID Register | | | | | |
| | Reserved | | | | | |
| Optional ⁽³⁾⁽⁴⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability - Alternate Requester ID (ARI) | | | |
| | Control | | Next Function | Function Groups | | |
| | Reserved | | | | | |
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability - DSN | | | |
| | PCI Express Device Serial Number (1st) | | | | | |
| | PCI Express Device Serial Number (2nd) | | | | | |
| | Reserved | | | | | |
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability - Power Budgeting Enhanced Capability Header | | | |
| | | Reserved | | DS | | |
| | Reserved Power Budget Data - State D0, D1, D3, | | | | | |
| | Power Budget Capability | | | | | |
| | Reserved | | | | | |
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability ID - Latency Tolerance Reporting (LTR) | | | |
| | No-Snoop | | Snoop | | | |

| | 31 | Le | 5 15 | 0 | | |
|-------------------------|-----------------------------------|--------------|---|---------------------------------------|--|--|
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | | ded Capability ID - ver Allocation | | |
| | Capability Register | | | | | |
| - | Latency Indicator | | | | | |
| | Co | ontrol | Status | | | |
| | Power Allocation Array Register 0 | | | | | |
| | Power Allocation Array Register 1 | | | | | |
| | Reserved | | | | | |
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability ID - Single Root I/O Virtualization (SR-IOV) | | | |
| | | Capabil | ity Register | | | |
| | SR-IOV Status (not supported) | | Control | | | |
| | Total VFs | | Initia | Initial VFs | | |
| | Function Dependency Link | | Numb | oer VFs | | |
| | VF Stride | | First VF Offset | | | |
| _ | VF Device ID | | Reserved | | | |
| | Supported Page Sizes | | | | | |
| _ | System Page Size | | | | | |
| - | VF Base Address Register 0 | | | | | |
| - | VF Base Address Register 1 | | | | | |
| _ | VF Base Address Register 2 | | | | | |
| _ | VF Base Address Register 3 | | | | | |
| _ | VF Base Address Register 4 | | | | | |
| | VF Base Address Register 5 | | | | | |
| | Reserved | | | | | |
| | Reserved | | | | | |
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability ID - Transaction Processing Hints (TPH) | | | |
| | | Capabil | ity Register | | | |
| F | Requester Control Register | | | | | |
| F | Res | erved | Steering Tag Upper | Steering Tag Lower | | |
| | Reserved | | | | | |

| 31 | | 16 15 | | (| 0 |
|-------------------------|--------------------------------------|--------------|--|---------------------------|--------------|
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | PCI Express Extended Capability ID - Secondary PCIe Extended Capability | | 300ł |
| | Lane Control (not supported) | | | | |
| | Reserved Lane Error Stat | | | Lane Error Status | 3081 |
| | Lane Equalization Control Register 0 | | | | |
| | Lane Equalization Control Register 1 | | | | |
| | Lane Equalization Control Register 2 | | | | |
| | Lane Equalization Control Register 3 | | | | |
| | Reserved | | | | 31CH 3BCH |
| Optional ⁽³⁾ | Next Cap | Cap. Ver. | | ss Extended ility - VC | 3C01 |
| | Port VC Capability Register 1 | | | | 3C41 |
| | Port VC Capability Register 2 | | | | 3C81 |
| | Port VC Status | | Port VC Control | | 3CC1 |
| | VC Resource Capability Register 0 | | | | |
| | VC Resource Control Register 0 | | | | |
| | VC Resource Status Register 0 | | | | |
| | Reserved | | | | 4001 FFF1 |

Notes:

- 1. The MSI Capability Structure varies depending on the selections in the CORE Generator tool GUI.
- 2. Reserved for Endpoint configurations (returns 0x0000000).
- 3. The layout of the PCI Express Extended Configuration Space (100h-FFFh) can change dependent on which optional capabilities are enabled. This table represents the Extended Configuration space layout when all optional extended capability structures, except RBAR, are enabled.
- 4. Enabled by default if the SR-IOV option is enabled.

| 31 | | 16 | 15 | 0 | |
|----|----------------------------------|-------------|---------------|-------------|-----|
| | Devi | ce ID | Vendor ID | | 00h |
| | Sta | itus | Com | mand | 04h |
| | | Class Code | | Rev ID | 08h |
| | BIST | Header | Lat Timer | Cache Ln | 0Ch |
| | | Base Addres | ss Register 0 | | 10h |
| | | Base Addres | ss Register 1 | | 14h |
| | | Base Addres | ss Register 2 | | 18h |
| | Base Address Register 3 | | | | 1Ch |
| | Base Address Register 4 | | | | 20h |
| | Base Address Register 5 | | | | 24h |
| | Cardbus CIS Pointer | | | | 28h |
| | Subsystem ID Subsystem Vendor ID | | | n Vendor ID | 2Ch |
| | Expansion ROM Base Address | | | | 30h |
| | Reserved CapPtr | | | | 34h |
| | Reserved | | | | 38h |
| Ν | lax Lat | Min Gnt | Intr Pin | Intr Line | 3Ch |

Table 2-27: Type 0 PCI Configuration Space Header

Table 2-28: Type 1 PCI Configuration Space Header

| 16 | 15 | 0 | |
|----------------------------------|---|--|---|
| vice ID | Vendor ID | | 00 |
| atus | Com | mand | 04 |
| Class Code | | Rev ID | 08 |
| Header | Lat Timer | Cache Ln | 00 |
| Base Addre | ss Register 0 | · | 10 |
| Base Addre | ss Register 1 | | 14 |
| Sub Bus Number | Second Bus Number | Primary Bus Number | 18 |
| ary Status | I/O Limit | I/O Base | 10 |
| ory Limit | Memory Base | | 20 |
| Prefetchable Memory Limit | | Prefetchable Memory Base | |
| Prefetchable Ba | se Upper 32 Bits | | 28 |
| Prefetchable Limit Upper 32 Bits | | | 20 |
| Jpper 16 Bits | I/O Base U | pper 16 Bits | 30 |
| Reserved | | CapPtr | 34 |
| Expansion ROM Base Address | | | 38 |
| e Control | Intr Pin | Intr Line | 30 |
| | vice ID tatus Class Code Header Base Addres Base Addres Base Addres Sub Bus Number lary Status ory Limit Memory Limit Prefetchable Ba Prefetchable Lin Upper 16 Bits Reserved | vice ID Venc tatus Com Class Code Header Lat Timer Base Address Register 0 Base Address Register 1 Sub Bus Number Second Bus Number lary Status I/O Limit ory Limit Memo e Memory Limit Prefetchable Prefetchable Base Upper 32 Bits Prefetchable Limit Upper 32 Bits Upper 16 Bits I/O Base U Reserved Expansion ROM Base Address | vice ID Vendor ID tatus Command Class Code Rev ID Class Code Lat Timer Cache Ln Base Address Register 0 Base Address Register 1 Base Address Register 1 Sub Bus Number Second Bus Number Primary Bus Number lary Status I/O Limit I/O Base ory Limit Memory Base e Memory Limit Prefetchable Memory Base Prefetchable Base Upper 32 Bits Prefetchable Limit Upper 32 Bits Upper 16 Bits I/O Base Upper 16 Bits Reserved CapPtr Expansion ROM Base Address |



Designing with the Core

This chapter includes guidelines and additional information to make designing with the core easier.

General Design Guidelines

See the 7 Series FPGAs Clocking Resources User Guide, 7 Series FPGAs GTX/GTH Transceivers User Guide, and 7 Series FPGAs SelectIO[™] Resources User Guide [Ref 3] for more information on clock, transceiver, and I/O placement rules. Pay special attention to clocking conflicts or errors that might result from choosing I/O and transceivers that do not follow these guides. Failure to adhere to these guides will result in build errors and data integrity errors.

Clocking

The input system clock signal of the Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® is called ref_clk. The core requires a 100 MHz, 125 MHz, or 250 MHz clock input. The clock frequency used must match the clock frequency selection in the CORE Generator[™] tool GUI. For more information, see the Answer Records at the Xilinx PCI Express Solution Center.

In a typical PCI Express solution, the PCI Express reference clock is a spread spectrum clock (SSC), provided at 100 MHz. In most commercial PCI Express systems, SSC cannot be disabled. For more information regarding SSC and PCI Express, see Section 4.3.7.1.1 of the *PCI Express Base Specification, rev. 3.0*.

Synchronous and Non-Synchronous Clocking

There are two ways to clock the PCI Express system:

- Using synchronous clocking, where a shared clock source is used for all devices.
- Using non-synchronous clocking, where each device has its own clock source. Spread spectrum clocking (SSC) and active state power management (ASPM) must not be used in systems with non-synchronous clocking.



IMPORTANT: The most commonly used clocking methodology is synchronous clocking. All add-in card designs must use synchronous clocking due to the characteristics of the provided reference clock. For devices using the Slot clock, the "Slot Clock Configuration" setting in the Link Status Register must be enabled in the CORE Generator tool GUI. See Clocking Requirements, page 84 and the 7 Series FPGAs GTX/GTH Transceivers User Guide for additional information regarding reference clock requirements.

For synchronous clocked systems, each link partner device shares the same clock source. Figure 3-1 and Figure 3-3 show a system using a 100 MHz reference clock. When using the 125 MHz or the 250 MHz reference clock option, an external PLL must be used to do a multiply of 5/4 and 5/2 to convert the 100 MHz clock to 125 MHz and 250 MHz, respectively, as illustrated in Figure 3-2 and Figure 3-4.

Even if the device is part of an embedded system, if the system uses commercial PCI Express root complexes or switches along with typical motherboard clocking schemes, synchronous clocking should still be used as shown in Figure 3-1 and Figure 3-2.

Figure 3-1 through Figure 3-4 illustrate high-level representations of the board layouts. Designers must ensure that proper coupling, termination, and so forth are used when laying out the board.

Note: Figure 3-1 through Figure 3-4 are high-level representations of the board layout. Ensure that proper coupling, termination, and so forth are used when laying out a board.

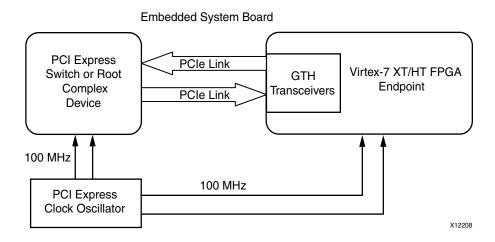


Figure 3-1: Embedded System Using 100 MHz Reference Clock

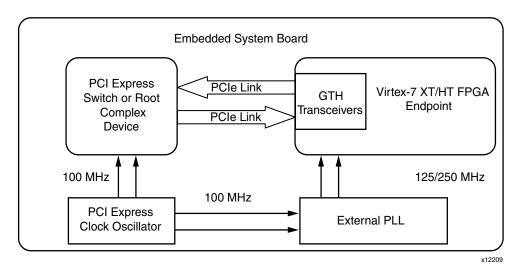


Figure 3-2: Embedded System Using 125/250 MHz Reference Clock

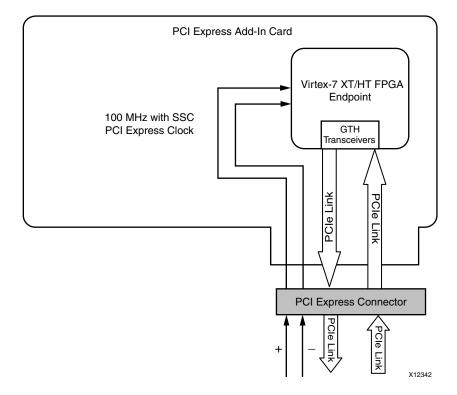


Figure 3-3: Open System Add-In Card Using 100 MHz Reference Clock

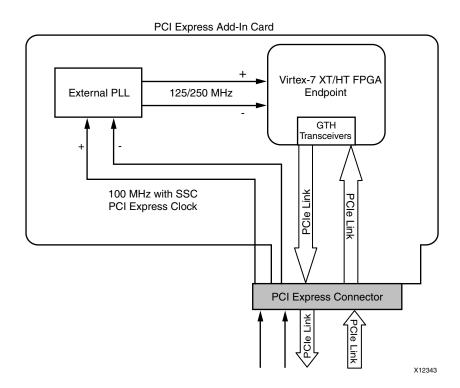


Figure 3-4: Open System Add-In Card Using 125/250 MHz Reference Clock

Resets

The LogiCORE[™] Gen3 Integrated Block for PCIe IP core resets the system using sys_reset, an asynchronous, active-Low reset signal asserted during the PCI Express Fundamental Reset. Asserting this signal causes a hard reset of the entire core, including the GTH transceivers. After the reset is released, the core attempts to link train and resume normal operation. In a typical Endpoint application, for example an add-in card, a sideband reset signal is normally present and should be connected to sys_reset. For Endpoint applications that do not have a sideband system reset signal, the initial hardware reset should be generated locally. Four reset events can occur in PCI Express:

- Cold Reset. A Fundamental Reset that occurs at the application of power. The sys_reset signal is asserted to cause the cold reset of the core.
- Warm Reset. A Fundamental Reset triggered by hardware without the removal and re-application of power. The sys_reset signal is asserted to cause the warm reset to the core.
- Hot Reset. In-band propagation of a reset across the PCI Express Link through the protocol, resetting the entire Endpoint device. In this case, sys_reset is not used. In the case of Hot Reset, the cfg_hot_reset_out signal is asserted to indicate the source of the reset.

• Function-Level Reset: In-band propagation of a reset across the PCI Express Link through the protocol, resetting only a specific function. In this case, the core asserts the bit of either cfg_flr_in_process and/or cfg_vf_flr_in_process that corresponds to the function being reset. Logic associated with the function being reset must assert the corresponding bit of cfg_flr_done or cfg_vf_flr_done to indicate it has completed the reset process. Support for function-level reset is indicated via the PF0_DEV_CAP_FUNCTION_LEVEL_RESET_CAPABLE parameter.

The User Application interface of the core has an output signal called user_reset. This signal is deasserted synchronously with respect to user_clk. The user_reset signal is asserted as a result of any of these conditions:

- Fundamental Reset: Occurs (cold or warm) due to assertion of sys_reset.
- PLL within the Core Wrapper: Loses lock, indicating an issue with the stability of the clock input.
- Loss of Transceiver PLL Lock: Any transceiver loses lock, indicating an issue with the PCI Express Link.

The user_reset signal is deasserted synchronously with user_clk after all of the listed conditions are resolved, allowing the core to attempt to train and resume normal operation.

Note: Systems designed to the PCI Express electromechanical specification provide a sideband reset signal, which uses 3.3V signaling levels—see the Virtex-7 FPGA data sheet to understand the requirements for interfacing to such signals.

AXI4-Stream Interface Description

This section provides a detailed description of the features, parameters, and signals associated with the client-side interfaces of the Gen3 Integrated Block for PCIe.

Overview of Features

Figure 3-5 illustrates the client-side interface of the Gen3 Integrated Block for PCIe.

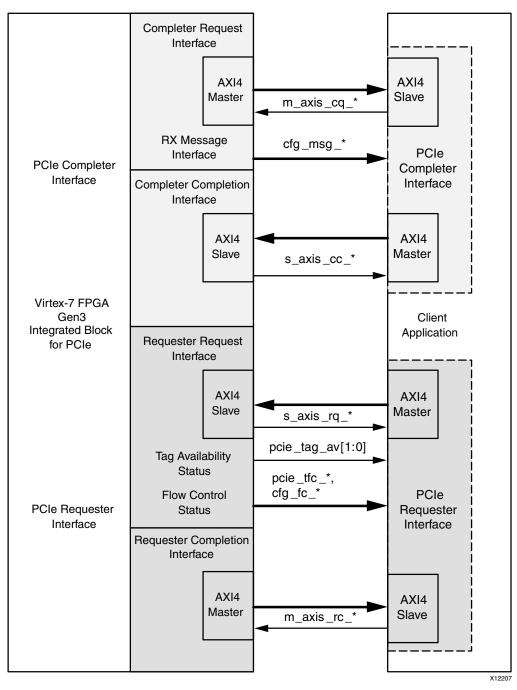


Figure 3-5: Block Diagram of Gen3 Integrated Block Client Interfaces

The interface is organized as four separate interfaces through which data can be transferred between the PCIe link and the client application:

- A PCIe Completer reQuest (CQ) interface through which requests arriving from the link are delivered to the client application.
- A PCIe Completer Completion (CC) interface through which the client application can send back responses to the completer requests. The client can process all Non-Posted

transactions as split transactions. That is, it can continue to accept new requests on the completer request interface while sending a completion for a request.

- A PCIe Requester reQuest (RQ) interface through which the client application can generate requests to remote PCIe devices attached to the link.
- A PCIe Requester Completion (RC) interface through which the Integrated Block returns the completions received from the link (in response to the client's requests as PCIe requester) to the client application.

Each of the four interfaces is based on the AMBA4® AXI4-Stream Protocol Specification. The width of these interfaces can be configured as 64, 128, or 256 bytes, and the user clock frequencies can be selected as 62.5, 125, or 250 MHz, depending on the number of lanes and PCIe generation chosen by the user. Table 3-1 lists the valid combinations of interface width and user clock frequency for the different link widths and link speeds supported by the Integrated Block. All four AXI4-Stream interfaces are configured with the same width in all cases.

In addition, the Integrated Block contains two interfaces through which status information is communicated to the PCIe master side of the client application:

- A flow control status interface that provides information on currently available transmit credit, so that the client application can schedule requests based on available credit.
- A tag availability status interface that provides information on the number of tags available to assign to Non-Posted requests, so that the client can schedule requests without the risk of being blocked by all tags being in use within the PCIe controller.

Finally, the Integrated Block also has a received-message interface which optionally provides indications to the user logic when a message is received from the link, rather than transferring the entire message to the client over the CQ interface.

| PCI Express Generation/ Maximum Link Speed | Maximum Link Width Capability | AXI4-Stream Interface Width | User Clock Frequency (MHz) |
|---|----------------------------------|--------------------------------|-------------------------------|
| | | 64 bits | 62.5 |
| | x1 | 64 bits | 125 |
| | | 64 bits | 250 |
| | | 64 bits | 62.5 |
| | x2 | 64 bits | 125 |
| Gen1 (2.5 GT/s) | 64 bits | 64 bits | 250 |
| | x4 64 bits 64 bits | 64 bits | 125 |
| | | 250 | |
| | .v9 | 64 bits | 250 |
| | x8 | 128 bits | 125 |

 Table 3-1:
 Data Width and Clock Frequency Settings for the Client Interfaces

| PCI Express Generation/ Maximum Link Speed | Maximum Link Width Capability | AXI4-Stream Interface Width | User Clock Frequency (MHz) |
|---|----------------------------------|--------------------------------|-------------------------------|
| | | 64 bits | 62.5 |
| | x1 | 64 bits | 125 |
| | | 64 bits | 250 |
| | | 64 bits | 125 |
| Gen2 (5.0 GT/s) | x2 64 bits | 250 | |
| | | 64 bits | 250 |
| | x4 | 128 bits | 125 |
| | 9 | 128 bits | 250 |
| | x8 | 256 bits | 125 |
| | x1 | 64 bits | 125 |
| | | 64 bits | 250 |
| | | 64 bits | 250 |
| Gen3 (8.0 GT/s) | x2 128 bits | 125 | |
| | | 128 bits | 250 |
| | x4 | 256 bits | 125 |
| | x8 | 256 bits | 250 |

| Tahle 2-1. | Data Width and Clock Frequency Settings for the Client Interfaces (Co | nt'd) |
|------------|---|--------|
| TUDIC J I. | Data whath and clock frequency settings for the cheft interfaces (co | nic uj |

Data Alignment Options

A transaction layer packet (TLP) is transferred on each of the AXI4-Stream interfaces as a descriptor followed by payload data (when the TLP has a payload). The descriptor has a fixed size of 16 bytes on the request interfaces and 12 bytes on the completion interfaces. On its transmit side (towards the link), the Integrated Block assembles the TLP header from the parameters supplied by the client application in the descriptor. On its receive side (towards the client), the Integrated Block extracts parameters from the headers of received TLP and constructs the descriptors for delivering to the client application. Each TLP is transferred as a packet, as defined in the AXI4-Stream Interface Protocol.

When a payload is present, there are two options for aligning the first byte of the payload with respect to the datapath.

- 1. Dword-aligned mode: In this mode, the descriptor bytes are followed immediately by the payload bytes in the next Dword position, whenever a payload is present.
- 2. Address-Aligned Mode: In this mode, the payload can begin at any byte position on the datapath. For data transferred from the Integrated Block to the client, the position of the first byte is determined as:

 $n = A \mod w$

where A is the memory or I/O address specified in the descriptor (for message and configuration requests, the address is taken as 0), and w is the configured width of the data bus in bytes. Any gap between the end of the descriptor and the start of the first byte of the payload is filled with null bytes.

For data transferred from the Integrated Block to the client application, the data alignment is determined based on the starting address where the data block is destined to in client memory. For data transferred from the client application to the Integrated Block, the client must explicitly communicate the position of the first byte to the Integrated Block using the tuser sideband signals when the address-aligned mode is in use.

In the address-aligned mode, the payload and descriptor are not allowed to overlap. That is, the transmitter begins a new beat to start the transfer of the payload after it has transmitted the descriptor. The transmitter fills any gaps between the last byte of the descriptor and the first byte of the payload with null bytes.

The CORE Generator tool customization GUI applies the data alignment option globally to all four interfaces. However, advanced users can select the alignment mode independently for each of the four AXI4-Stream interfaces. This is done by setting the corresponding alignment mode parameter, with the constraint that the Requester Completion (RC) interface can be set to the address-aligned mode only when the Requester reQuest (RQ) interface is configured in the address-aligned mode. See Interface Operation, page 85 for more details on address alignment and example diagrams.

Straddle Option on Requester Completion Interface

When the Requester Completion (RC) interface is configured for a width of 256 bits, depending on type of TLP and Payload size, there can be significant interface utilization inefficiencies, if a maximum of 1 TLP is allowed to start or end per interface beat. This inefficient use of RC interface can lead to overflow of the completion FIFO when Infinite Receiver Credits are advertized. The user must either a) Restrict the number of outstanding Non Posted requests, so as to keep the total number of completions received less than 64 and within the completion of the FIFO size selected, or b) Use the RC interface straddle option. See Figure 3-59 for waveforms showing this option.

The straddle option, available only on the 256-bit wide RC interface, is enabled through the CORE Generator tool customization GUI. See Interface Settings, page 231 for instructions on enabling the option in the GUI. When this option is enabled, the Integrated Block can start a new Completion TLP on byte lane 16 when the previous TLP has ended at or before byte lane 15 in the same beat. Thus, with this option enabled, it is possible for the Integrated Block to send two Completion TLPs entirely in the same beat on the RC interface, if neither of them has more than one Dword of payload.

The straddle setting is only available when the interface width is set to 256 bits and the RC interface is set to Dword-aligned mode.

Table 3-2 lists the valid combinations of interface width, addressing mode, and the straddle option.

| Interface Width | Alignment Mode | Straddle Option | Description |
|-----------------|-----------------|-----------------|--|
| 64 bits | Dword-aligned | Not applicable | 64-bit, Dword-aligned |
| 64 bits | Address-aligned | Not applicable | 64-bit, Address-aligned |
| 128 bits | Dword-aligned | Not applicable | 128-bit, Dword-aligned |
| 128 bits | Address-aligned | Not applicable | 128-bit, Address-aligned |
| 256 bits | Dword-aligned | Disabled | 256-bit, Dword-aligned, straddle disabled |
| 256 bits | Dword-aligned | Enabled | 256-bit, Dword-aligned, straddle enabled (only allowed for the Requester Completion interface) |
| 256 bits | Address-aligned | Not applicable | 256-bit, Address-aligned |

 Table 3-2:
 Valid Combinations of Interface Width, Alignment Mode, and Straddle

Receive Transaction Ordering

The Gen3 Integrated Block for PCIe contains logic on its receive side to ensure that TLPs received from the link and delivered on its completer request interface and requester completion interface do not violate the PCI Express transaction ordering constraints. The ordering actions performed by the Integrated Block are based on the following key rules:

 Posted requests must be able to pass Non-Posted requests on the Completer reQuest (CQ) interface. To enable this capability, the Integrated Block implements a flow control mechanism on the CQ interface through which client logic can control the flow of Non-Posted requests without affecting Posted requests. The client logic signals the availability of a buffer to receive a Non-Posted request by asserting the pcie_cq_np_req signal.

The Integrated Block delivers a Non-Posted request to the client only when the available credit is non-zero. The Integrated Block continues to deliver Posted requests while the delivery of Non-Posted requests has been paused for lack of credit. When no backpressure is applied by the credit mechanism for the delivery of Non-Posted requests, the Integrated Block delivers Posted and Non-Posted requests in the same order as received from the link. For more information on controlling the flow of Non-Posted requests, see Selective Flow Control for Non-Posted Requests, page 103.

- PCIe ordering requires that a completion TLP not be allowed to pass a Posted request, except in the following cases:
 - Completions with the Relaxed Ordering attribute bit set can pass Posted requests
 - Completions with the ID-based ordering bit set can pass a Posted request if the completion's Completer ID is different from the Posted request's Requestor ID.

The Integrated Block does not start the transfer of a Completion TLP received from the link on the Requester Completion (RC) interface until it has completely transferred all Posted TLPs that arrived before it, unless one of the two rules applies.

After a TLP has been transferred completely to the client side, it is the responsibility of the client application to enforce ordering constraints whenever needed.

Transmit Transaction Ordering

On the transmit side, the Integrated Block receives TLPs from the user on two different interfaces: the Requester reQuest (RQ) interface and the Completer Completion (CC) interface. The Integrated Block does not re-order transactions received from each of these interfaces. It is difficult to predict how the requester-side requests and completer-side completions are ordered in the transmit pipeline of the Integrated Block, after these have been multiplexed into a single traffic stream. In cases where completion TLPs must maintain ordering with respect to requests, client logic can supply a 4-bit sequence number with any request that needs to maintain strict ordering with respect to a Completion transmitted from the CC interface, on the seq_num[3:0] inputs within the s_axis_rq_tuser bus. The Integrated Block places this sequence number on its pcie_rq_seq_num[3:0] output and assert pcie_rq_seq_num_vld when the request TLP has reached a point in the transmit pipeline at which no new completion TLP from the client can pass it. This mechanism can be used in the following situations to maintain TLP order:

- The client logic requires ordering to be maintained between a request TLP and a completion TLP that follows it. In this case, client logic must wait for the sequence number of the requester request to appear on the pcie_rq_seq_num[3:0] output before starting the transfer of the completion TLP on the target completion interface.
- The client logic requires ordering to be maintained between a request TLP and MSI/MSI-X TLP signaled through the MSI Message interface. In this case, the client logic must wait for the sequence number of the requester request to appear on the pcie_rq_seq_num[3:0] output before signaling MSI or MSI-X on the MSI Message interface.

Clocking Requirements

All client interface signals of the Gen3 Integrated Block for PCIe are timed with respect to the user clock (user_clk), which can have a frequency of 62.5, 125, or 250 MHz, depending on the link speed and link width configured (see Table 3-1).

Interface Operation

This section describes the operation of the client-side interfaces of the Gen3 Integrated Block for PCIe.

Completer Interface

This interface maps the transactions (memory, I/O read/write, messages, Atomic Operations) received from the PCIe link into transactions on the Completer reQuest (CQ) interface based on the AXI4-Stream protocol. The completer interface consists of two separate interfaces, one for data transfers in each direction. Each interface is based on the AXI4-Stream protocol, and its width can be configured as 64, 128, or 256 bits. The CQ interface is for transfer of requests (with any associated payload data) to the client application, and the Completer Completion (CC) interface is for transferring the Completion data (for a Non-Posted request) from the client application for forwarding on the link. The two interfaces operate independently. That is, the Integrated Block can transfer new requests over the CQ interface while receiving a Completion for a previous request.

Completer Request Descriptor Formats

The Integrated Block transfers each request TLP received from the link over the CQ interface as an independent AXI4-Stream packet. Each packet starts with a descriptor and can have payload data following the descriptor. The descriptor is always 16 bytes long, and is sent in the first 16 bytes of the request packet. The descriptor is transferred during the first two beats on a 64-bit interface, and in the first beat on a 128-bit or 256-bit interface.

The formats of the descriptor for different request types are illustrated in Figure 3-6, Figure 3-7, Figure 3-8, and Figure 3-9. The format of Figure 3-6 applies when the request TLP being transferred is a memory read/write request, an I/O read/write request, or an Atomic Operation request. The format of Figure 3-7 is used for Vendor-Defined Messages (Type 0 or Type 1) only. The format of Figure 3-8 is used for all ATS messages (Invalid Request, Invalid Completion, Page Request, PRG Response). For all other messages, the descriptor takes the format of Figure 3-9.

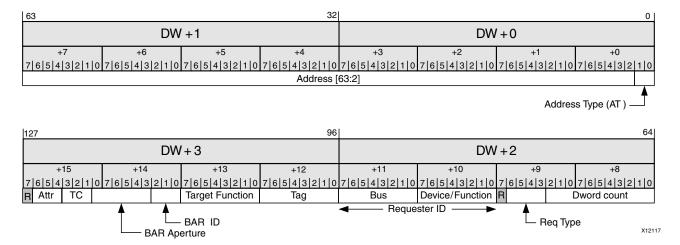


Figure 3-6: Completer Request Descriptor Format for Memory, I/O, and Atomic Op Requests

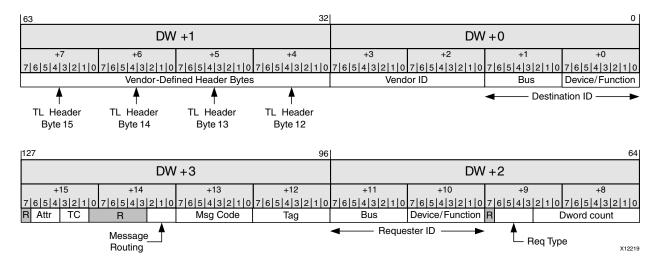


Figure 3-7: Completer Request Descriptor Format for Vendor-Defined Messages

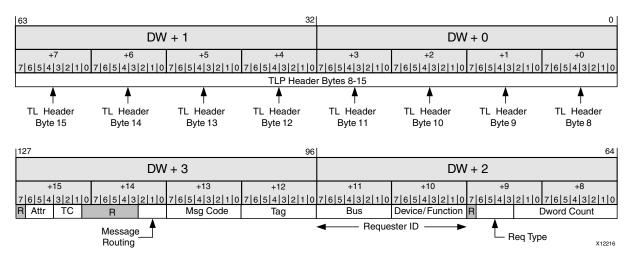


Figure 3-8: Completer Request Descriptor Format for ATS Messages

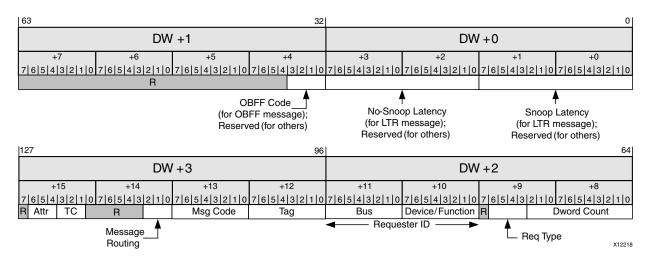




Table 3-3 describes the individual fields of the completer request descriptor.

| Table 3-3: | Completer Request Descriptor Fields |
|------------|--|
|------------|--|

| Bit Index | Field Name | Description |
|-----------|--------------|--|
| | | This field is defined for memory transactions and Atomic Operations only. It contains the AT bits extracted from the TL header of the request. |
| 1:0 | Address Type | 00: Address in the request is untranslated 01: Transaction is a Translation Request 10: Address in the request is a translated address |
| | | 11: Reserved |

| Bit Index | Field Name | Description |
|-----------|-----------------|---|
| 63:2 | Address | This field applies to memory, I/O, and Atomic Op requests. It provides the address from the TLP header. This is the address of the first Dword referenced by the request. The First_BE bits from m_axis_cq_tuser must be used to determine the byte-level address. When the transaction specifies a 32-bit address, bits [63:32] of this field are 0. |
| 74:64 | Dword Count | These 11 bits indicate the size of the block (in Dwords) to be read or written (for messages, size of the message payload). Its range is 0 - 256 Dwords. For I/O accesses, the Dword count is always 1. For a zero length memory read/write request, the Dword count is 1, with the First_BE bits set to all 0s. |
| 78:75 | Request Type | Identifies the transaction type. The transaction types and their encodings are listed in Table 3-4. |
| 95:80 | Requester ID | PCI Requester ID associated with the request. With legacy interpretation of RIDs, these 16 bits are divided into an 8-bit bus number [95:88], 5-bit device number [87:83], and 3-bit Function number [82:80]. When ARI is enabled, bits [95:88] carry the 8-bit bus number and [87:80] provide the Function number. When the request is a Non-Posted transaction, the client completer application must store this field and supply it back to the Integrated Block with the completion data. |
| 103:96 | Тад | PCIe Tag associated with the request. When the request is a Non-Posted transaction, the client logic must store this field and supply it back to the Integrated Block with the completion data. This field can be ignored for memory writes and messages. |
| 111:104 | Target Function | This field is defined for memory, I/O, and Atomic Op requests only. It provides the Function number the request is targeted at, determined by the BAR check. When ARI is in use, all 8 bits of this field are valid. Otherwise, only bits [106:104] are valid. Following are Target Function Value to PF/VF map mappings: • 0: PF0 • 1: PF1 • 64: VF0 • 65: VF1 • 66: VF2 • 67: VF3 • 68: VF4 • 69: VF5 |

Table 3-3: Completer Request Descriptor Fields (Cont'd)

| Bit Index | Field Name | Description |
|-----------|---------------------------|---|
| 114:112 | BAR ID | This field is defined for memory, I/O, and Atomic Op requests only. It provides the matching BAR number for the address in the request. 000: BAR 0 (VF-BAR 0 for VFs) 001: BAR 1 (VF-BAR 1 for VFs) 010: BAR 2 (VF-BAR 2 for VFs) 011: BAR 3 (VF-BAR 3 for VFs) 100: BAR 4 (VF-BAR 4 for VFs) 101: BAR 5 (VF-BAR 5 for VFs) 110: Expansion ROM Access For 64-bit transactions, the BAR number is given as the lower address of the matching pair of BARs (that is, 0, 2, or 4). |
| 120:115 | BAR Aperture | This 6-bit field is defined for memory, I/O, and Atomic Op requests only. It provides the aperture setting of the BAR matching the request. This information is useful in determining the bits to be used by the client in addressing its memory or I/O space. For example, a value of 12 indicates that the aperture of the matching BAR is 4K, and the client can therefore ignore bits [63:12] of the address. For VF BARs, the value provided on this output is based on the memory space consumed by a single VF covered by the BAR. |
| 123:121 | Transaction Class (TC) | PCIe Transaction Class (TC) associated with the request. When the request is a Non-Posted transaction, the client completer application must store this field and supply it back to the Integrated Block with the completion data. |
| 126:124 | Attributes | These bits provide the setting of the Attribute bits associated with the request. Bit 124 is the No Snoop bit and bit 125 is the Relaxed Ordering bit. Bit 126 is the ID-Based Ordering bit, and can be set only for memory requests and messages. When the request is a Non-Posted transaction, the client completer application must store this field and supply it back to the Integrated Block with the completion data. |
| 15:0 | Snoop Latency | This field is defined for LTR messages only. It provides the value of the 16-bit Snoop Latency field in the TLP header of the message. |
| 31:16 | No-Snoop Latency | This field is defined for LTR messages only. It provides the value of the 16-bit No-Snoop Latency field in the TLP header of the message. |
| 35:32 | OBFF Code | This field is defined for OBFF messages only. The OBFF Code field is used to distinguish between various OBFF cases: 1111b: "CPU Active" – System fully active for all device actions including bus mastering and interrupts 0001b: "OBFF" – System memory path available for device memory read/write bus master activities 0000b: "Idle" – System in an idle, low power state All other codes are reserved. |

Table 3-3: Completer Request Descriptor Fields (Cont'd)

| Bit Index | Field Name | Description |
|-----------|--------------------------|---|
| 111:104 | Message Code | This field is defined for all messages. It contains the 8-bit Message Code extracted from the TLP header. Appendix F of the <i>PCI Express Base Specification, rev. 3.0</i> provides a complete list of the supported Message Codes. |
| 114:112 | Message Routing | This field is defined for all messages. These bits provide the 3-bit Routing field r[2:0] from the TLP header. |
| 15:0 | Destination ID | This field applies to Vendor-Defined Messages only. When the message is routed by ID (that is, when the Message Routing field is 010 binary), this field provides the Destination ID of the message. |
| 63:32 | Vendor-Defined Header | This field applies to Vendor-Defined Messages only. It contains the bytes extracted from Dword 3 of the TLP header. |
| 63:0 | ATS Header | This field is applicable to ATS messages only. It contains the bytes extracted from Dwords 2 and 3 of the TLP header. |

Table 3-3: Completer Request Descriptor Fields (Cont'd)

Table 3-4: Transaction Types

| Request Type (binary) | Description | | |
|--------------------------|---|--|--|
| 0000 | Memory Read Request | | |
| 0001 | Memory Write Request | | |
| 0010 | I/O Read Request | | |
| 0011 | I/O Write Request | | |
| 0100 | Memory Fetch and Add Request | | |
| 0101 | Memory Unconditional Swap Request | | |
| 0110 | Memory Compare and Swap Request | | |
| 0111 | Locked Read Request (allowed only in Legacy Devices) | | |
| 1000 | Type 0 Configuration Read Request (on Requester side only) | | |
| 1001 | Type 1 Configuration Read Request (on Requester side only) | | |
| 1010 | Type 0 Configuration Write Request (on Requester side only) | | |
| 1011 | Type 1 Configuration Write Request (on Requester side only) | | |
| 1100 | Any message, except ATS and Vendor-Defined Messages | | |
| 1101 | Vendor-Defined Message | | |
| 1110 | ATS Message | | |
| 1111 | Reserved | | |

Completer Request Interface Operation

Figure 3-10 illustrates the signals associated with the completer request interface of the core. The core delivers each TLP on this interface as an AXI4-Stream packet. The packet starts with a 128-bit descriptor, followed by data in the case of TLPs with a payload.

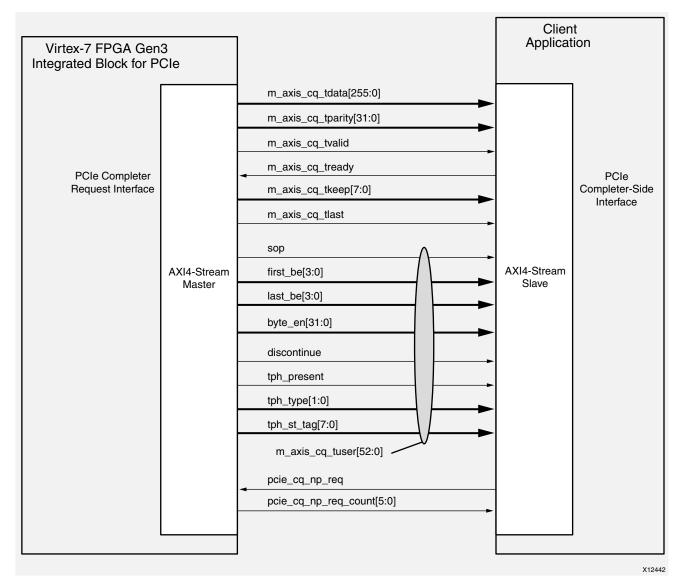


Figure 3-10: Completer Request Interface Signals

The completer request interface supports two distinct data alignment modes, selected by the attribute $AXISTEN_IF_CQ_ALIGNMENT_MODE$. In the Dword-aligned mode, the first byte of valid data appears in lane $n = (16 + A \mod 4) \mod w$, where:

- A is the byte-level starting address of the data block being transferred
- *w* is the width of the interface in bytes

In the address-aligned mode, the data always starts in a new beat after the descriptor has ended, and its first valid byte is on lane $n = A \mod w$, where w is the width of the interface in bytes. For memory, I/O, and Atomic Operation requests, address A is the address contained in the request. For messages, the address is always taken as 0 for the purpose of determining the alignment of its payload.

Completer Memory Write Operation

The timing diagrams in Figure 3-11, Figure 3-12, and Figure 3-13 illustrate the Dword-aligned transfer of a memory write TLP received from the link across the Completer reQuest (CQ) interface, when the interface width is configured as 64, 128, and 256 bits, respectively. For illustration purposes, the starting Dword address of the data block being written into client memory is assumed to be (m * 32 + 1), for an integer m > 0. Its size is assumed to be n Dwords, for some n = k * 32 + 29, k > 0.

In both Dword-aligned and address-aligned modes, the transfer starts with the 16 descriptor bytes, followed immediately by the payload bytes. The m_axis_cq_tvalid signal remains asserted over the duration of the packet. The client can prolong a beat at any time by deasserting m_axis_cq_tready. The AXI4-Stream interface signals m_axis_cq_tkeep (one per Dword position) indicate the valid Dwords in the packet including the descriptor and any null bytes inserted between the descriptor and the payload. That is, the tkeep bits are set to 1 contiguously from the first Dword of the descriptor until the last Dword of the payload. During the transfer of a packet, the tkeep bits can be 0 only in the last beat of the packet, when the packet does not fill the entire width of the interface. The m_axis_cq_tlast signal is always asserted in the last beat of the packet.

The CQ interface also includes the First Byte Enable and the Last Enable bits in the m_axis_cq_tuser bus. These are valid in the first beat of the packet, and specify the valid bytes of the first and last Dwords of payload.

The m_axi_cq_tuser bus also provides several informational signals that can be used to simplify the logic associated with the client side of the interface, or to support additional features. The sop signal is asserted in the first beat of every packet, when its descriptor is on the bus. The byte enable outputs byte_en[31:0] (one per byte lane) indicate the valid bytes in the payload. The bits of byte_en are asserted only when a valid payload byte is in the corresponding lane (that is, not asserted for descriptor or padding bytes between the descriptor and payload). The asserted byte enable bits are always contiguous from the start of the payload, except when the payload size is two Dwords or less. For cases of one-Dword and two-Dword writes, the byte enables can be non-contiguous. Another special case is that of a zero-length memory write, when the Integrated Block transfers a one-Dword payload with all byte_en bits set to 0. Thus, the client logic can, in all cases, use the byte_en signals directly to enable the writing of the associated bytes into memory.

In the Dword-aligned mode, there can be a gap of zero, one, two, or three byte positions between the end of the descriptor and the first payload byte, based on the address of the first valid byte of the payload. The actual position of the first valid byte in the payload can be determined either from first_be[3:0] or byte_en[31:0] in the m_axis_cq_tuser bus.

When a Transaction Processing Hint is present in the received TLP, the Integrated Block transfers the parameters associated with the hint (TPH Steering Tag and Steering Tag Type) on signals within the m_axis_cq_tuser bus.

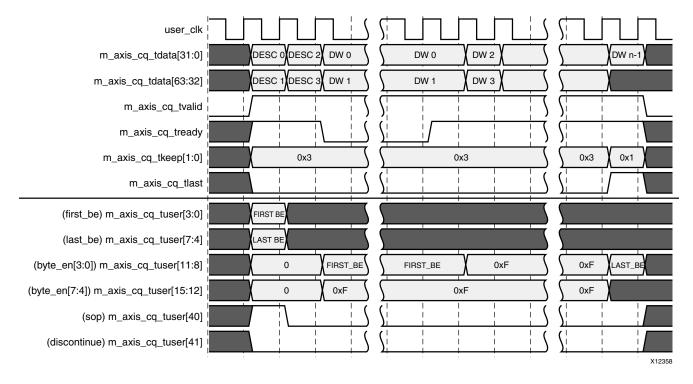


Figure 3-11: Memory Write Transaction on the Completer Request Interface (Dword-Aligned Mode, Interface Width = 64 Bits)

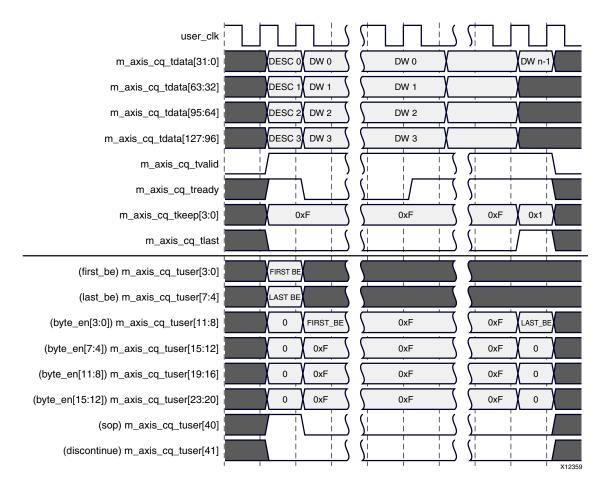


Figure 3-12: Memory Write Transaction on the Completer Request Interface (Dword-Aligned Mode, Interface Width = 128 Bits)

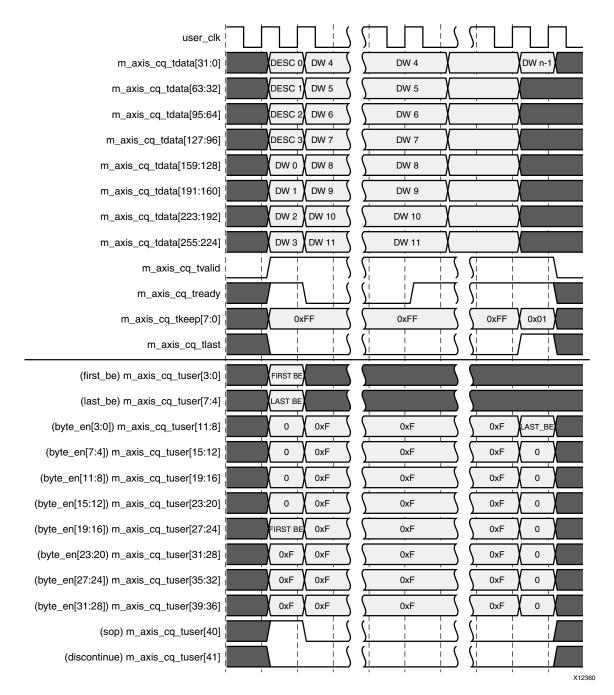


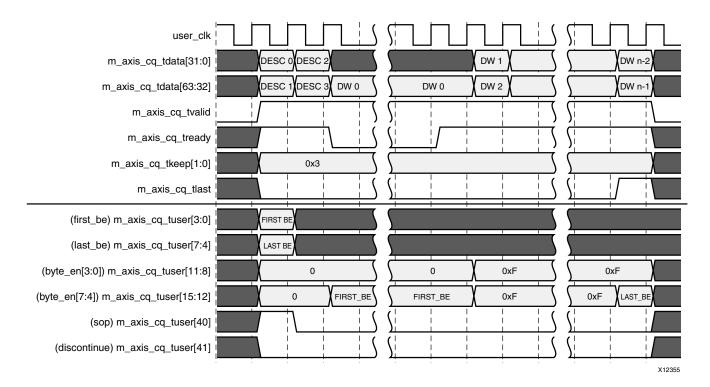
Figure 3-13: Memory Write Transaction on the Completer Request Interface (Dword-Aligned Mode, Interface Width = 256 Bits)

The timing diagrams in Figure 3-14, Figure 3-15, and Figure 3-16 illustrate the address-aligned transfer of a memory write TLP received from the link across the CQ interface, when the interface width is configured as 64, 128 and 256 bits, respectively. For the purpose of illustration, the starting Dword address of the data block being written into client memory is assumed to be (m * 32 + 1), for an integer m > 0. Its size is assumed to be n Dwords, for some n = k * 32 + 29, k > 0.

www.xilinx.com

In the address-aligned mode, the delivery of the payload always starts in the beat following the last byte of the descriptor. The first byte of the payload can appear on any byte lane, based on the address of the first valid byte of the payload. The keep outputs m_axis_cq_tkeep remain High in the gap between the descriptor and the payload. The actual position of the first valid byte in the payload can be determined either from the least significant bits of the address in the descriptor or from the byte enable bits byte_en[31:0] in the m_axis_cq_tuser bus.

For writes of two Dwords or less, the 1s on byte_en cannot be contiguous from the start of the payload. In the case of a zero-length memory write, the Integrated Block transfers a one-Dword payload with the byte_en bits all set to 0 for the payload bytes.





| user_clk | | | |
|---|---|------------------------------|-------------|
| m_axis_cq_tdata[31:0] | DESC 0 | DW 3 | DW n-2 |
| m_axis_cq_tdata[63:32] | DESC 1 DW 0 | DW 0 DW 4 | DW n-1 |
| m_axis_cq_tdata[95:64] | DESC 2 DW 1 | DW 1 DW 5 | |
| m_axis_cq_tdata[127:96] | DESC 3 DW 2 | DW 2 DW 6 | |
| m_axis_cq_tvalid | | | |
| m_axis_cq_tready | | | |
| m_axis_cq_tkeep[3:0] | 0xF |) 0xF | |
| m_axis_cq_tlast | | | |
| | | | |
| (first_be) m_axis_cq_tuser[3:0] | FIRST BE |) | |
| (first_be) m_axis_cq_tuser[3:0] (last_be) m_axis_cq_tuser[7:4] | FIRST BE | | |
| | |))) 0) 0xF | |
| (last_be) m_axis_cq_tuser[7:4] | LAST BE | 0 0 0xF | 0xF LAST_BE |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] | | <u>}∧</u> | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] | LAST BE 0 0 FIRST_BE | FIRST_BE 0xF | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] | LAST BE 0 0 FIRST_BE 0 0 0 0 |) FIRST_BE 0xF | OxF LAST_BE |
| <pre>(last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] (byte_en[15:12]) m_axis_cq_tuser[23:20]</pre> | LAST BE 0 0 FIRST_BE 0 0 0 0 |) FIRST_BE 0xF | OxF LAST_BE |

Figure 3-15: Memory Write Transaction on the Completer Request Interface (Address-Aligned Mode, Interface Width = 128 Bits)

| user_clk | | | |
|---|---|---|---|
| m_axis_cq_tdata[31:0] | DESC 0 DW 7 | | DW n-6 |
| m_axis_cq_tdata[63:32] | DESC 1 DW 0 DW 8 | | DW n-5 |
| m_axis_cq_tdata[95:64] | DESC 2 DW 1 DW 9 | | DW n-4 |
| m_axis_cq_tdata[127:96] | DESC 3 DW 2 DW 10 | DW 10 | DW n-3 |
| m_axis_cq_tdata[159:128] | DW 3 DW 11 | DW 11 | DW n-2 |
| m_axis_cq_tdata[191:160] | DW 4 DW 12 | DW 12 | DW n-1 |
| m_axis_cq_tdata[223:192] | DW 5 DW 13 | DW 13 | |
| m_axis_cq_tdata[255:224] | DW 6 DW 14 | DW 14 | |
| m_axis_cq_tvalid | | 5 5 5 | |
| m_axis_cq_tready | | | S I |
| m_axis_cq_tkeep[7:0] | 0xFF | 0xFF | 0xFF 0x3F |
| m_axis_cq_tlast | | | |
| | | / <u>/</u> / | |
| (first_be) m_axis_cq_tuser[3:0] | FIRST BE | | |
| (first_be) m_axis_cq_tuser[3:0] (last_be) m_axis_cq_tuser[7:4] | FIRST BE | | |
| | , ,,, | | |
| (last_be) m_axis_cq_tuser[7:4] | | | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] | | | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] | LAST BE | | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] | LAST BE | 0xF 0xF | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] (byte_en[15:12]) m_axis_cq_tuser[23:20] | LAST BE 0 0 0 0xF 0 FIRST BE 0xF 0 0xF 0 0xF | 0xF 0xF 0xF 0xF | |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] (byte_en[15:12]) m_axis_cq_tuser[23:20] (byte_en[19:16]) m_axis_cq_tuser[27:24] | LAST BE 0 0 0 0 0 FIRST BE 0 0xF 0 0xF 0 0xF 0 0xF 0 0xF 0 0xF | 0xF 0xF 0xF 0xF | 0xF 0xF 0xF 0xF 0xF 0xF |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] (byte_en[15:12]) m_axis_cq_tuser[23:20] (byte_en[19:16]) m_axis_cq_tuser[27:24] (byte_en[23:20) m_axis_cq_tuser[31:28] | LAST BE 0 0 0 FIRST BE 0 0xF | 0xF 0xF 0xF 0xF 0xF 0xF | 0xF 0xF 0xF 0xF 0xF 0xF 0xF |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] (byte_en[15:12]) m_axis_cq_tuser[23:20] (byte_en[19:16]) m_axis_cq_tuser[27:24] (byte_en[23:20) m_axis_cq_tuser[31:28] (byte_en[27:24]) m_axis_cq_tuser[35:32] | LAST BE 0 0 0 First BE 0 First BE 0 OxF 0 OxF | 0xF 0xF 0xF 0xF 0xF 0xF 0xF | 0xF 0xF 0xF 0xF 0xF 0xF 0xF 0xF 0xF |
| (last_be) m_axis_cq_tuser[7:4] (byte_en[3:0]) m_axis_cq_tuser[11:8] (byte_en[7:4]) m_axis_cq_tuser[15:12] (byte_en[11:8]) m_axis_cq_tuser[19:16] (byte_en[15:12]) m_axis_cq_tuser[23:20] (byte_en[19:16]) m_axis_cq_tuser[27:24] (byte_en[23:20) m_axis_cq_tuser[31:28] (byte_en[27:24]) m_axis_cq_tuser[35:32] (byte_en[31:28]) m_axis_cq_tuser[39:36] | LAST BE 0 0 0 First BE 0 First BE 0 OxF 0 OxF | 0xF 0xF 0xF 0xF 0xF 0xF 0xF | 0xF 0xF 0xF 0xF 0xF 0xF 0xF 0xF 0xF |

Figure 3-16: Memory Write Transaction on the Completer Request Interface (Address-Aligned Mode, Interface Width = 256 Bits)

Completer Memory Read Operation

A memory read request is transferred across the completer request interface in the same manner as a memory write request, except that the AXI4-Stream packet contains only the 16-byte descriptor. The timing diagrams in Figure 3-17, Figure 3-18, and Figure 3-19 illustrate the transfer of a memory read TLP received from the link across the completer request interface, when the interface width is configured as 64, 128, and 256 bits,

respectively. The packet occupies two consecutive beats on the 64-bit interface, while it is transferred in a single beat on the 128- and 256-bit interfaces. The m_axis_cq_tvalid signal remains asserted over the duration of the packet. The client can prolong a beat at any time by deasserting m_axis_cq_tready. The sop signal in the m_axis_cq_tuser bus is asserted when the first descriptor byte is on the bus.

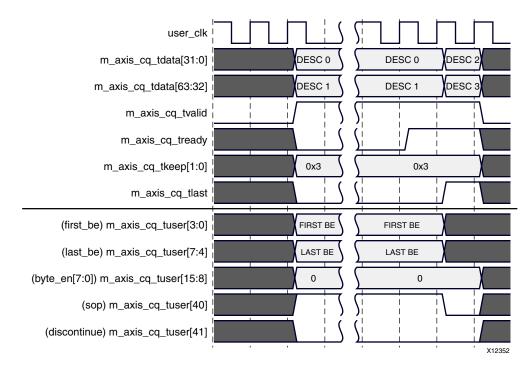


Figure 3-17: Memory Read Transaction on the Completer Request Interface (Interface Width = 64 Bits)

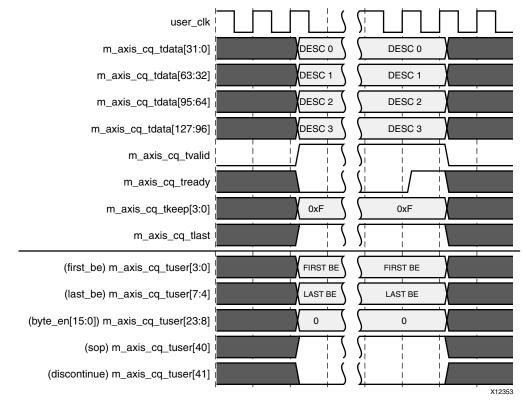


Figure 3-18: Memory Read Transaction on the Completer Request Interface (Interface Width = 128 Bits)

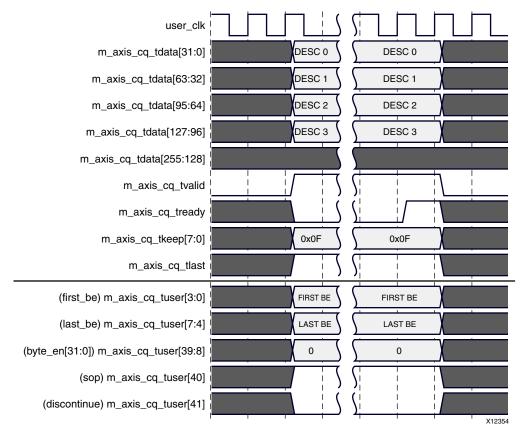


Figure 3-19: Memory Read Transaction on the Completer Request Interface (Interface Width = 256 Bits)

The byte enable bits associated with the read request for the first and last Dwords are supplied by the Integrated Block on the m_axis_cq_tuser sideband bus. These bits are valid when the first descriptor byte is being transferred, and must be used by the client to determine the byte-level starting address and the byte count associated with the request. For the special cases of one-Dword and two-Dword reads, the byte enables can be non-contiguous. The byte enables are contiguous in all other cases. A zero-length memory read is sent on the CQ interface with the Dword count field in the descriptor set to 1 and the first and last byte enables set to 0.

The client must respond to each memory read request with a Completion. The data requested by the read can be sent as a single Completion or multiple Split Completions. These Completions must be sent via the Completer Completion (CC) interface of the Integrated Block. The Completions for two distinct requests can be sent in any order, but the Split Completions for the same request must be in order. The operation of the CC interface is described in Completer Completion Interface Operation, page 104.

I/O Write Operation

The transfer of an I/O write request on the CQ interface is similar to that of a memory write request with a one-Dword payload. The transfer starts with the 128-bit descriptor, followed by the one-Dword payload. When the Dword-aligned mode is in use, the payload Dword

immediately follows the descriptor. When the address-alignment mode is in use, the payload Dword is supplied in a new beat after the descriptor, and its alignment in the datapath is based on the address in the descriptor. The First Byte Enable bits in the m_axis_cq_tuser indicate the valid bytes in the payload. The byte enable bits byte_en also provide this information.

Because an I/O write is a Non-Posted transaction, the client logic must respond to it with a Completion containing no data payload. The Completions for I/O requests can be sent in any order. Errors associated with the I/O write transaction can be signaled to the requester by setting the Completion Status field in the completion descriptor to CA (Completer Abort) or UR (Unsupported Request), as is appropriate. The operation of the Completer Completion interface is described in Completer Completion Interface Operation, page 104.

I/O Read Operation

The transfer of an I/O read request on the CQ interface is similar to that of a memory read request, and involves only the descriptor. The length of the requested data is always one Dword, and the First Byte Enable bits in m_axis_cq_tuser indicate the valid bytes to be read.

The client logic must respond to an I/O read request with a one-Dword Completion (or a Completion with no data in the case of an error). The Completions for two distinct I/O read requests can be sent in any order. Errors associated with an I/O read transaction can be signaled to the requester by setting the Completion Status field in the completion descriptor to CA (Completer Abort) or UR (Unsupported Request), as is appropriate. The operation of the Completer Completion interface is described in Completer Completion Interface Operation, page 104.

Atomic Operations on the Completer Request Interface

The transfer of an Atomic Op request on the completer request interface is similar to that of a memory write request. The payload for an Atomic Op can range from one Dword to eight Dwords, and its starting address is always aligned on a Dword boundary. The transfer starts with the 128-bit descriptor, followed by the payload. When the Dword-aligned mode is in use, the first payload Dword immediately follows the descriptor. When the address-alignment mode is in use, the payload starts in a new beat after the descriptor, and its alignment is based on the address in the descriptor. The m_axis_cq_tkeep output indicates the end of the payload. The byte_en signals in m_axis_cq_tuser also indicate the valid bytes in the payload. The First Byte Enable and Last Byte Enable bits in m_axis_cq_tuser should not be used for Atomic Operations.

Because an Atomic Operation is a Non-Posted transaction, the client logic must respond to it with a Completion containing the result of the operation. Errors associated with the operation can be signaled to the requester by setting the Completion Status field in the completion descriptor to Completer Abort (CA) or Unsupported Request (UR), as is appropriate. The operation of the Completer Completion interface is described in Completer Completion Interface Operation, page 104.

Message Requests on the Completer Request Interface

The transfer of a message on the CQ interface is similar to that of a memory write request, except that a payload might not always be present. The transfer starts with the 128-bit descriptor, followed by the payload, if present. When the Dword-aligned mode is in use, the payload immediately follows the descriptor. When the address-alignment mode is in use, the first Dword of the payload is supplied in a new beat after the descriptor, and always starts in byte lane 0. The client can determine the end of the end of the payload from the states of the m_axis_cq_tlast and m_axis_cq_tkeep signals. The byte_en signals in m_axis_cq_tuser also indicate the valid bytes in the payload. The First Byte Enable and Last Byte Enable bits in m_axis_cq_tuser should not be used for Message transactions.

The AXISTEN_IF_ENABLE_RX_MSG_INTFC parameter must be set to 0 to enable the delivery of messages through the CQ interface. When this parameter is set to 0, the component bits of the AXISTEN_IF_ENABLE_MSG_ROUTE[17:0] parameter can be used to select the specific message types that the user wants delivered over the CQ interface. Setting a parameter bit to 1 enables the delivery of the corresponding type of messages on the interface, and setting it to 0 results in the Integrated Block filtering the message.

When AXISTEN_IF_ENABLE_RX_MSG_INTFC is set to 1, no messages are delivered on the CQ interface. Indications of received message are instead sent through a dedicated receive message interface (see Receive Message Interface, page 115).

Aborting a Transfer

For any request that includes an associated payload, the Integrated Block can signal an error in the transferred payload by asserting the discontinue signal in the m_axis_cq_tuser bus in the last beat of the packet (along with m_axis_cq_tlast). This occurs when the Integrated Block has detected an uncorrectable error while reading data from its internal memories. The client application must discard the entire packet when it has detected discontinue asserted in the last beat of a packet. This condition is considered a fatal error in the Integrated Block.

Selective Flow Control for Non-Posted Requests

The PCI Express Base Specification requires that the Completer Request interface continue to deliver Posted transactions even when the client is unable to accept Non-Posted transactions. To enable this capability, the Integrated Block implements a credit-based flow control mechanism on the CQ interface through which client logic can control the flow of Non-Posted requests without affecting Posted requests. The client logic signals the availability of buffers for receive Non-Posted requests using the pcie_cq_np_req signal. The Gen3 Integrated Block delivers a Non-Posted request to the client only when the available credit is non-zero. The Integrated Block continues to deliver Posted requests while the delivery of Non-Posted requests has been paused for lack of credit. When no backpressure is applied by the credit mechanism for the delivery of Non-Posted requests, the Integrated Block delivers Posted and Non-Posted requests in the same order as received from the link.

The Integrated Block maintains an internal credit counter to track the credit available for Non-Posted requests on the completer request interface. The following algorithm is used to keep track of the available credit:

- On reset, the counter is set to 0.
- After the Integrated Block comes out of reset, in every clock cycle:
 - If pcie_cq_np_req is High and no Non-Posted request is being delivered this cycle, the credit count is incremented by 1, unless it has already reached its saturation limit of 32.
 - If pcie_cq_np_req is Low and a Non-Posted request is being delivered this cycle, the credit count is decremented by 1, unless it is already 0.
 - Otherwise, the credit count remains unchanged.
- The Integrated Block starts delivery of a Non-Posted TLP to the client only if the credit count is greater than 0.

The client application can either provide a one-cycle pulse on pcie_cq_np_req each time it is ready to receive a Non-Posted request, or can keep it permanently asserted if it does not need to exercise selective backpressure of Non-Posted requests. If the credit count is always non-zero, the Integrated Block delivers Posted and Non-Posted requests in the same order as received from the link. If it remains 0 for some time, Non-Posted requests can accumulate in the Integrated Block's FIFO. When the credit count becomes non-zero later, the Integrated Block first delivers the accumulated Non-Posted requests that arrived before Posted requests already delivered to the client, and then reverts to delivering the requests in the order received from the link.

The assertion and deassertion of the pcie_cq_np_req signal does not need to be aligned with the packet transfers on the completer request interface.

The client can monitor the current value of the credit count on the output pcie_cq_np_req_count [5:0]. The counter saturates at 32. Because of internal pipeline delays, there can be several cycles of delay between the Integrated Block receiving a pulse on the pcie_cq_np_req input and updating the pcie_cq_np_req_count output in response. Thus, when the client has adequate buffer space available, it should provide the credit in advance so that Non-Posted requests are not held up by the Integrated Block for lack of credit.

Completer Completion Interface Operation

Figure 3-20 illustrates the signals associated with the completer completion interface of the core. The core delivers each TLP on this interface as an AXI4-Stream packet.

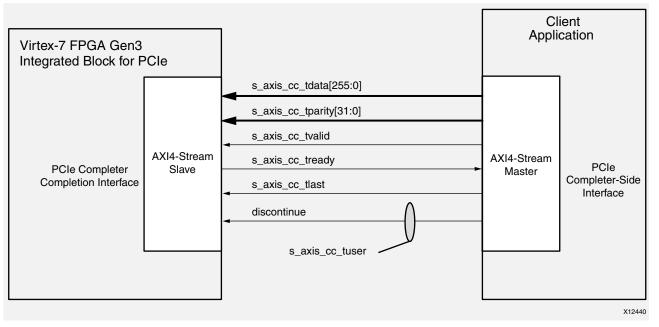


Figure 3-20: Completer Completion Interface Signals

The Gen3 Integrated Block for PCIe delivers each TLP on the Completer Completion (CC) interface as an AXI4-Stream packet. The packet starts with a 96-bit descriptor, followed by data in the case of Completions with a payload.

The CC interface supports two distinct data alignment modes, selected by the AXISTEN_IF_CC_ALIGNMENT_MODE parameter. In the Dword-aligned mode, the first byte of valid data must be presented in lane $n = (12 + A \mod 4) \mod w$, where A is the byte-level starting address of the data block being transferred (as conveyed in the Lower Address field of the descriptor) and w the width of the interface in bytes (8, 16, or 32). In the address-aligned mode, the data always starts in a new beat after the descriptor has ended. When transferring the Completion payload for a memory or I/O read request, its first valid byte is on lane $n = A \mod w$. For all other Completions, the payload is aligned with byte lane 0.

Completer Completion Descriptor Format

The client application sends completion data for a completer request to the CC interface of the Integrated Block as an independent AXI4-Stream packet. Each packet starts with a descriptor and can have payload data following the descriptor. The descriptor is always 12 bytes long, and is sent in the first 12 bytes of the completion packet. The descriptor is transferred during the first two beats on a 64-bit interface, and in the first beat on a 128-or 256-bit interface. When the client application splits the completion data for a request into multiple Split Completions, it must send each Split Completion as a separate AXI4-Stream packet, with its own descriptor.

The format of the completer completion descriptor is illustrated in Figure 3-21. The individual fields of the completer request descriptor are described in Table 3-5.

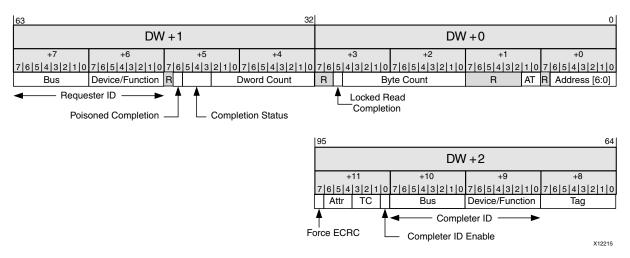


Figure 3-21: Completer Completion Descriptor Format

| Bit Index | Field Name | Description |
|-----------|---------------------------|--|
| 6:0 | Lower Address | For memory read Completions, this field must be set to the least significant 7 bits of the starting byte-level address of the memory block being transferred. For all other Completions, the Lower Address must be set to all zeros. |
| 9:8 | Address Type | This field is defined for Completions of memory transactions and Atomic Operations only. For these Completions, the client logic must copy the AT bits from the corresponding request descriptor into this field. This field must be set to 0 for all other Completions. |
| 28:16 | Byte Count | These 13 bits can have values in the range of 0 – 4096 bytes. If a Memory Read Request is completed using a single Completion, the Byte Count value indicates Payload size in bytes. This field must be set to 4 for I/O read Completions and I/O write Completions. The byte count must be set to 1 while sending a Completion for a zero-length memory read, and a dummy payload of 1 Dword must follow the descriptor. The byte count must be set to 0 when sending a UR or CA Completion. For each Memory Read Completion, the Byte Count field must indicate the remaining number of bytes required to complete the Request, including the number of bytes returned with the Completion. If a Memory Read Request is completed using multiple Completions, the Byte Count value for each successive Completion is the value indicated by the preceding Completion. The total number of bytes required to complete a Memory Read Request is calculated as shown in Table 3-6, page 108. |
| 29 | Locked Read Completion | This bit must be set when the Completion is in response to a Locked Read request. It must be set to 0 for all other Completions. |

| Bit Index | Field Name | Description |
|-----------|---|---|
| 42:32 | Dword Count | These 11 bits indicate the size of the payload of the current packet in Dwords. Its range is 0 - 1K Dwords. This field must be set to 1 for I/O read Completions and 0 for I/O write Completions. The Dword count must be set to 1 while sending a Completion for a zero-length memory read. The Dword count must be set to 0 when sending a UR or CA Completion. In all other cases, the Dword count must correspond to the actual number of Dwords in the payload of the current packet. |
| 45:43 | Completion Status | These bits must be set based on the type of Completion being sent. The only valid settings are: 000: Successful Completion 001: Unsupported Request (UR) 100: Completer Abort (CA) |
| 46 | Poisoned Completion | This bit can be used by the client to poison the Completion TLP being sent. This bit must be set to 0 for all Completions, except when the client has detected an error in the block of data following the descriptor and wants to communicate this information using the Data Poisoning feature of PCI Express. |
| 63:48 | Requester ID | PCI Requester ID associated with the request (copied by the client from the request). |
| 71:64 | Tag | PCIe Tag associated with the request (copied by the client from the request). |
| 79:72 | Target Function/ Device Number | Function number of the completer Function. The client must copy this value from the Target Function field of the descriptor of the corresponding request. When ARI is in use, all 8 bits of this field must be set to the target Function number. Otherwise, bits [74:72] must be set to the target Function number. When ARI is not in use, and the Integrated Block is configured as a Root Complex, the client must supply the 5-bit Device Number of the completer on bits [79:75]. When ARI is not used and the Integrated Block is configured as an Endpoint, the client can optionally supply a 5-bit Device Number of the completer on bits [79:75]. The client must set the Completer ID Enable bit in the descriptor if a Device Number is supplied on bits [79:75]. This value is used by the Integrated Block when sending the Completion TLP, instead of the stored value of the Device Number captured by the Integrated Block from Configuration Requests. |
| 87:80 | Completer Bus Number | Bus number associated with the completer Function. When the Integrated Block is configured as a Root Complex, the client must supply the 8-bit Bus Number of the completer in this field. When the Integrated Block is configured as an Endpoint, the client can optionally supply a Bus Number in this field. The client must set the Completer ID Enable bit in the descriptor if a Bus Number is supplied in this field. This value is used by the Integrated Block when sending the Completion TLP, instead of the stored value of the Bus Number captured by the Integrated Block from Configuration Requests. |

Table 3-5: Completer Completion Descriptor Fields (Cont'd)

| Bit Index | Field Name | Description |
|-----------|---------------------------|--|
| 88 | Completer ID Enable | The purpose of this field is to enable the client to supply the bus and device numbers to be used in the Completer ID. This field is applicable only to Endpoint configurations. If this field is 0, the Integrated Block uses the captured values of the bus and device numbers to form the Completer ID. If this input is 1, the Integrated Block uses the bus and device numbers supplied by the client in the descriptor to form the Completer ID. |
| 91:89 | Transaction Class (TC) | PCIe Transaction Class (TC) associated with the request. The client must copy this value from the TC field of the associated request descriptor. |
| 94:92 | Attributes | PCIe attributes associated with the request (copied from the request). Bit 92 is the No Snoop bit, bit 93 is the Relaxed Ordering bit, and bit 94 is the ID-Based Ordering bit. |
| 95 | Force ECRC | Force ECRC insertion. Setting this bit to 1 forces the Integrated Block to append a TLP Digest containing ECRC to the Completion TLP, even when ECRC is not enabled for the Function sending the Completion. |

Table 3-5: Completer Completion Descriptor Fields (Cont'd)

| Table 3-6: | Calculating Byte Count from Completer Request first_be[3:0], last_be[3:0], Dword |
|------------|--|
| Count[10:0 |)] |

| first_be[3:0] | last_be[3:0] | Total Byte Count |
|---------------|--------------|-------------------|
| 1xx1 | 0000 | 4 |
| 01x1 | 0000 | 3 |
| 1x10 | 0000 | 3 |
| 0011 | 0000 | 2 |
| 0110 | 0000 | 2 |
| 1100 | 0000 | 2 |
| 0001 | 0000 | 1 |
| 0010 | 0000 | 1 |
| 0100 | 0000 | 1 |
| 1000 | 0000 | 1 |
| 0000 | 0000 | 1 |
| xxx1 | 1xxx | Dword_count*4 |
| xxx1 | 01xx | (Dword_count*4)-1 |
| xxx1 | 001x | (Dword_count*4)-2 |
| xxx1 | 0001 | (Dword_count*4)-3 |
| xx10 | 1xxx | (Dword_count*4)-1 |
| xx10 | 01xx | (Dword_count*4)-2 |
| xx10 | 001x | (Dword_count*4)-3 |
| xx10 | 0001 | (Dword_count*4)-4 |
| ×100 | 1xxx | (Dword_count*4)-2 |

| first_be[3:0] | last_be[3:0] | Total Byte Count |
|---------------|--------------|-------------------|
| x100 | 01xx | (Dword_count*4)-3 |
| x100 | 001x | (Dword_count*4)-4 |
| x100 | 0001 | (Dword_count*4)-5 |
| 1000 | 1xxx | (Dword_count*4)-3 |
| 1000 | 01xx | (Dword_count*4)-4 |
| 1000 | 001x | (Dword_count*4)-5 |
| 1000 | 0001 | (Dword_count*4)-6 |

Table 3-6:Calculating Byte Count from Completer Request first_be[3:0], last_be[3:0], DwordCount[10:0] (Cont'd)

Completions with Successful Completion Status

The client must return a Completion to the CC interface of the Integrated Block for every Non-Posted request it receives from the completer request interface. When the request completes with no errors, the client must return a Completion with Successful Completion (SC) status. Such a Completion might or might not contain a payload, depending on the type of request. Furthermore, the data associated with the request can be broken up into multiple Split Completions when the size of the data block exceeds the maximum payload size configured. Client logic is responsible for splitting the data block into multiple Split Completion over the completer completion interface as a separate AXI4-Stream packet, with its own 12-byte descriptor.

In the example timing diagrams of this section, the starting Dword address of the data block being transferred (as conveyed in bits [6:2] of the Lower Address field of the descriptor) is assumed to be (m * 8 + 1), for an integer m. The size of the data block is assumed to be n Dwords, for some n = k * 32 + 28, k > 0.

The CC interface supports two data alignment modes: Dword-aligned and address-aligned. The timing diagrams in Figure 3-22, Figure 3-23, and Figure 3-24 illustrate the Dword-aligned transfer of a Completion from the client across the CC interface, when the interface width is configured as 64, 128, and 256 bits, respectively. In this case, the first Dword of the payload starts immediately after the descriptor. When the data block is not a multiple of four bytes, or when the start of the payload is not aligned on a Dword address boundary, the client must add null bytes to align the start of the payload on a Dword boundary and make the payload a multiple of Dwords. For example, when the data block starts at byte address 7 and has a size of 3 bytes, the client must add three null bytes before the first byte and two null bytes at the end of the block to make it two Dwords long. Also, in the case of non-contiguous reads, not all bytes in the data block returned are valid. In that case, the client must return the valid bytes in the proper positions, with null bytes added in gaps between valid bytes, when needed. The interface does not have any signals to indicate the valid bytes in the payload. This is not required, as the requester is responsible for keeping track of the byte enables in the request and discarding invalid bytes from the Completion.

In the Dword-aligned mode, the transfer starts with the 12 descriptor bytes, followed immediately by the payload bytes. The client must keep the s_axis_cc_tvalid signal asserted over the duration of the packet. The Integrated Block treats the deassertion of s_axis_cc_tvalid during the packet transfer as an error, and nullifies the corresponding Completion TLP transmitted on the link to avoid data corruption.

The client must also assert the s_axis_cc_tlast signal in the last beat of the packet. The Integrated Block can deassert s_axis_cc_tready in any cycle if it is not ready to accept data. The client must not change the values on the CC interface during a clock cycle that the Integrated Block has deasserted s_axis_cc_tready.

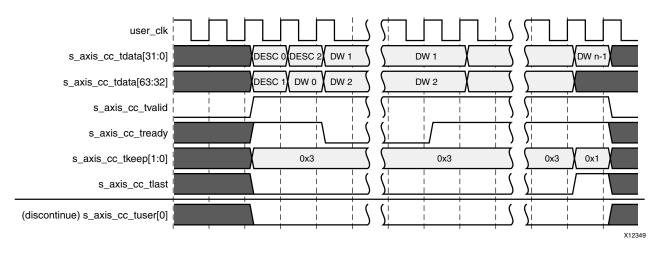


Figure 3-22: Transfer of a Normal Completion on the Completer Completion Interface (Dword-Aligned Mode, Interface Width = 64 Bits)

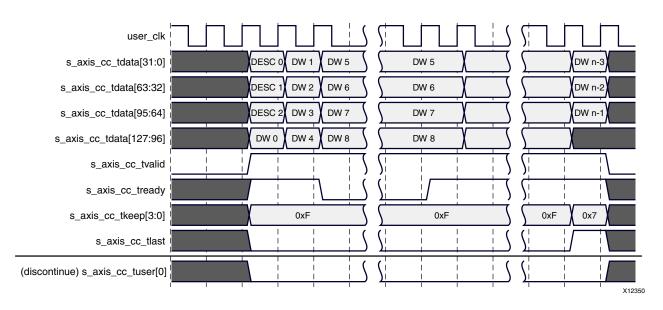
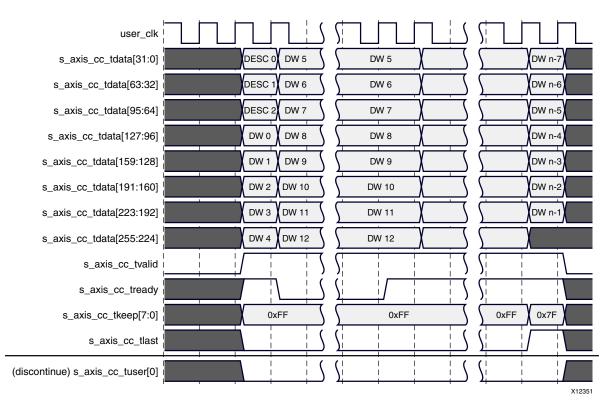
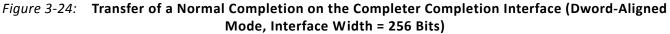


Figure 3-23: Transfer of a Normal Completion on the Completer Completion Interface (Dword-Aligned Mode, Interface Width = 128 Bits)

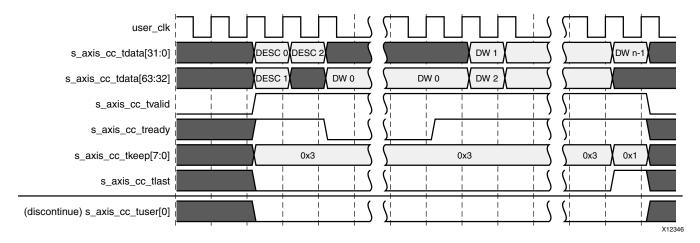
Gen3 Integrated Block for PCIe (v1.3) PG023 October 16, 2012 www.xilinx.com

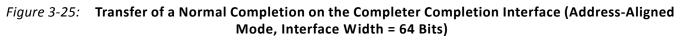




In the address-aligned mode, the delivery of the payload always starts in the beat following the last byte of the descriptor. For memory read Completions, the first byte of the payload can appear on any byte lane, based on the address of the first valid byte of the payload. For all other Completions, the payload must start in byte lane 0.

The timing diagrams in Figure 3-25, Figure 3-26, and Figure 3-27 illustrate the address-aligned transfer of a memory read Completion across the completer completion interface, when the interface width is configured as 64, 128, and 256 bits, respectively. For the purpose of illustration, the starting Dword address of the data block being transferred (as conveyed in bits [6:2] of the Lower Address field of the descriptor) is assumed to be (m * 8 + 1), for some integer m. The size of the data block is assumed to be n Dwords, for some n = k * 32 + 28, k > 0.





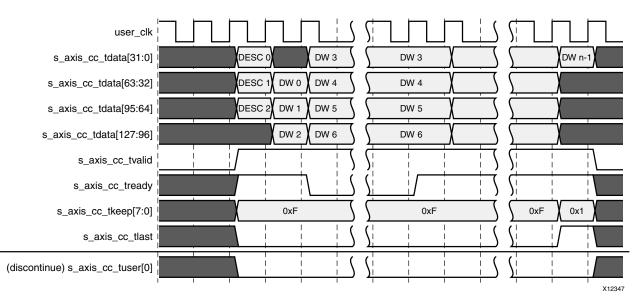


Figure 3-26: Transfer of a Normal Completion on the Completer Completion Interface (Address-Aligned Mode, Interface Width = 128 Bits)

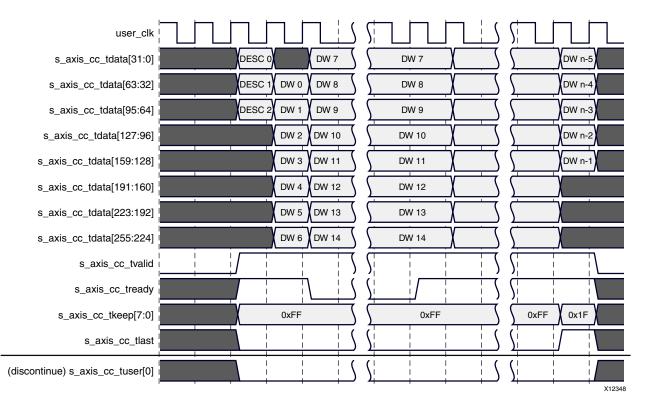


Figure 3-27: Transfer of a Normal Completion on the Completer Completion Interface (Address-Aligned Mode, Interface Width = 256 Bits)

Aborting a Completion Transfer

The client can abort the transfer of a Completion on the completer completion interface at any time during the transfer of the payload by asserting the discontinue signal in the s_axis_cc_tuser bus. The Integrated Block nullifies the corresponding TLP on the link to avoid data corruption.

The client can assert this signal in any cycle during the transfer, when the Completion being transferred has an associated payload. The client can either choose to terminate the packet prematurely in the cycle where the error was signaled (by asserting s_axis_cc_tlast), or can continue until all bytes of the payload are delivered to the Integrated Block. In the latter case, the Integrated Block treats the error as sticky for the following beats of the packet, even if the client deasserts the discontinue signal before reaching the end of the packet.

The discontinue signal can be asserted only when s_axis_cc_tvalid is High. The Integrated Block samples this signal when s_axis_cc_tvalid and s_axis_cc_tready are both asserted. Thus, after assertion, the discontinue signal should not be deasserted until s_axis_cc_tready is asserted.

When the Integrated Block is configured as an Endpoint, this error is reported by the Integrated Block to the Root Complex to which it is attached, as an Uncorrectable Internal Error using the Advanced Error Reporting (AER) mechanisms.

Completions with Error Status (UR and CA)

When responding to a request received on the completer request interface with an Unsupported Request (UR) or Completion Abort (CA) status, the client must send a three-Dword completion descriptor in the format of Figure 3-21, followed by five additional Dwords containing information on the request that generated the Completion. These five Dwords are necessary for the Integrated Block to log information about the request in its AER header log registers.

Figure 3-28 shows the sequence of information transferred when sending a Completion with UR or CA status. The information is formatted as an AXI4-Stream packet with a total of 8 Dwords, which are organized as follows:

- The first three Dwords contain the completion descriptor in the format of Figure 3-21.
- The fourth Dword contains the state of the following signals in m_axis_cq_tuser, copied from the request:
 - The First Byte Enable bits first_be[3:0] in m_axis_cq_tuser.
 - The Last Byte Enable bits last_be[3:0] in m_axis_cq_tuser.
 - Signals carrying information on Transaction Processing Hint: tph_present, tph_type[1:0], and tph_st_tag[7:0] in m_axis_cq_tuser.
- The four Dwords of the request descriptor received from the Integrated Block with the request.

| DW 1 | DW 0 |
|-----------------------------|-----------------------------|
| Completion Descriptor, DW 1 | Completion Descriptor, DW 0 |

| 63 32 | | | | | | |
|----------------------------|-----------------------|--|--|--------------------------|-----------------------------|--|
| DW 3 | | | | | DW 2 | |
| +7 7 6 5 4 3 2 1 0 | +6 7 6 5 4 3 2 1 0 | +5 +4 7 6 5 4 3 2 1 0 7 6 5 4 3 2 1 0 | | | Completion Descriptor, DW 2 | |
| R | | | | last_be | first_be | |
| tph_type[1:0]t tph_present | | | | | | |
| DW 5 | | | | DW 4 | | |
| Request Descriptor, DW 1 | | | | Request Descriptor, DW 0 | | |

| DW 7 | DW 6 |
|--------------------------|--------------------------|
| Request Descriptor, DW 3 | Request Descriptor, DW 2 |
| | X12232 |

Figure 3-28: Composition of the AXI4-Stream Packet for UR and CA Completions

The entire packet takes four beats on the 64-bit interface, two beats on the 128-bit interface, and a single beat on the 256-bit interface. The packet is transferred in an identical manner in both the Dword-aligned mode and the address-aligned mode, with the Dwords packed together. The client must keep the s_axis_cc_tvalid signal asserted over the duration of the packet. It must also assert the s_axis_cc_tlast signal in the last beat of the packet. The Integrated Block can deassert s_axis_cc_tready in any cycle if it is not ready to accept. The client must not change the values on the CC interface in any cycle that the Integrated Block has deasserted s_axis_cc_tready.

Receive Message Interface

The Gen3 Integrated Block for PCIe provides a separate receive-message interface which the client can optionally use to receive indications of messages received from the link. This interface is enabled by the AXISTEN_IF_ENABLE_RX_MSG_INTFC parameter. When the receive message interface is enabled, the Integrated Block signals the arrival of a message from the link by setting the cfg_msg_received_type[4:0] output to indicate the type of message (see Table 3-7) and pulsing the cfg_msg_received signal for one or more cycles. The duration of assertion of cfg_msg_received is determined by the type of message received (see Table 3-8). When cfg_msg_received is High, the Integrated Block transfers any parameters associated with the message on the bus 8 bits at a time on the bus cfg_msg_received_data. The parameters transferred on this bus in each cycle of cfg_msg_received assertion for various message types are listed in Table 3-8. For Vendor-Defined Messages, the Integrated Block transfers only the first Dword of any associated payload across this interface. When larger payloads are in use, the completer request interface should be used for the delivery of messages.

| cfg_msg_received_type[4:0] | Message Type |
|----------------------------|----------------|
| 0 | ERR_COR |
| 1 | ERR_NONFATAL |
| 2 | ERR_FATAL |
| 3 | Assert_INTA |
| 4 | Deassert_ INTA |
| 5 | Assert_INTB |
| 6 | Deassert_ INTB |
| 7 | Assert_INTC |
| 8 | Deassert_ INTC |
| 9 | Assert_INTD |
| 10 | Deassert_ INTD |
| 11 | PM_PME |
| 12 | PME_TO_Ack |
| 13 | PME_Turn_Off |

Table 3-7: Message Type Encoding on Receive Message Interface

| cfg_msg_received_type[4:0] | Message Type | |
|----------------------------|------------------------------------|--|
| 14 | PM_Active_State_Nak | |
| 15 | Set_Slot_Power_Limit | |
| 16 | Latency Tolerance Reporting (LTR) | |
| 17 | Optimized Buffer Flush/Fill (OBFF) | |
| 18 | Unlock | |
| 19 | Vendor_Defined Type 0 | |
| 20 | Vendor_Defined Type 1 | |
| 21 | ATS Invalid Request | |
| 22 | ATS Invalid Completion | |
| 23 | ATS Page Request | |
| 24 | ATS PRG Response | |
| 25 - 31 | Reserved | |

Table 3-7: Message Type Encoding on Receive Message Interface (Cont'd)

Table 3-8: Message Parameters on Receive Message Interface

| Message Type | Number of Cycles of cfg_msg_received Assertion | Parameter Transferred on cfg_msg_received_data[7:0] |
|---|--|--|
| ERR_COR, ERR_NONFATAL, ERR_FATAL | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| Assert_INTx, Deassert_INTx | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| PM_PME, PME_TO_Ack, PME_Turn_off, PM_Active_State_Nak | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| Set_Slot_Power_Limit | 6 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number Cycle 3: bits [7:0] of payload Cycle 4: bits [15:8] of payload Cycle 5: bits [23:16] of payload Cycle 6: bits [31:24] of payload |
| Latency Tolerance 6 Reporting (LTR) | | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number Cycle 3: bits [7:0] of Snoop Latency Cycle 4: bits [15:8] of Snoop Latency Cycle 5: bits [7:0] of No-Snoop Latency Cycle 6: bits [15:8] of No-Snoop Latency |
| Optimized Buffer Flush/Fill (OBFF) | 3 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number Cycle 3: OBFF Code |

| Message Type | Number of Cycles of cfg_msg_received Assertion | Parameter Transferred on cfg_msg_received_data[7:0] |
|------------------------|--|--|
| Unlock | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| Vendor_Defined Type 0 | 4 cycles when no data present, 8 cycles when data present. | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number Cycle 3: Vendor ID[7:0] Cycle 4: Vendor ID[15:8] Cycle 5: bits [7:0] of payload Cycle 6: bits [15:8] of payload Cycle 7: bits [23:16] of payload Cycle 8: bits [31:24] of payload |
| Vendor_Defined Type 1 | 4 cycles when no data present, 8 cycles when data present. | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number Cycle 3: Vendor ID[7:0] Cycle 4: Vendor ID[15:8] Cycle 5: bits [7:0] of payload Cycle 6: bits [15:8] of payload Cycle 7: bits [23:16] of payload Cycle 8: bits [31:24] of payload |
| ATS Invalid Request | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| ATS Invalid Completion | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| ATS Page Request | 2 | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |
| ATS PRG Response 2 | | Cycle 1: Requester ID, Bus Number Cycle 2: Requester ID, Device/Function Number |

Table 3-8: Message Parameters on Receive Message Interface (Cont'd)

Figure 3-29 is a timing diagram showing the example of a Set_Slot_Power_Limit message on the receive message interface. This message has an associated one-Dword payload. For this message, the parameters are transferred over six consecutive cycles. The following information appears on the cfg_msg_received_data bus in each cycle:

- Cycle 1: Bus number of Requester ID
- Cycle 2: Device/Function Number of Requester ID
- Cycle 3: Bits [7:0] of the payload Dword
- Cycle 4: Bits [15:8] of the payload Dword
- Cycle 5: Bits [23:16] of the payload Dword
- Cycle 6: Bits [31:24] of the payload Dword

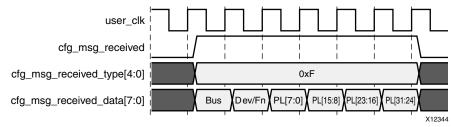


Figure 3-29: Receive Message Interface

The Integrated Block inserts a gap of at least one clock cycle between successive pulses on the cfg_msg_received output. There is no mechanism for the user to apply backpressure on the message indications delivered through the receive message interface. When using this interface, the client logic must always be ready to receive message indications.

Receive Message Interface Design Requirements

When configured as an Endpoint, the client application must implement one of the following:

- 1. The client application must issue Non-Posted Requests that result in Completions with the RO bit set.
- 2. The client application must not exceed the configured completion space.

This requirement ensures the RX Completion buffer does not overflow.

Requester Interface

The requester interface enables a client Endpoint application to initiate PCI transactions as a bus master across the PCIe link to the host memory. For Root Complexes, this interface is also used to initiate I/O and configuration requests. This interface can also be used by both Endpoints and Root Complexes to send messages on the PCIe link. The transactions on this interface are similar to those on the completer interface, except that the roles of the Gen3 Integrated Block for PCIe and the client application are reversed. Posted transactions are performed as single indivisible operations and Non-Posted transactions as split transactions.

The requester interface consists of two separate interfaces, one for data transfer in each direction. Each interface is based on the AXI4-Stream protocol, and its width can be configured as 64, 128, or 256 bits. The Requester reQuest (RQ) interface is for transfer of requests (with any associated payload data) from the client application to the Integrated Block, and the Requester Completion (RC) interface is used by the Integrated Block to deliver Completions received from the link (for Non-Posted requests) to the client application. The two interfaces operate independently. That is, the client can transfer new requests over the RQ interface while receiving a completion for a previous request.

Requester Request Interface Operation

On the RQ interface, the client delivers each TLP as an AXI4-Stream packet. The packet starts with a 128-bit descriptor, followed by data in the case of TLPs with a payload. Figure 3-30 shows the signals associated with the requester request interface.

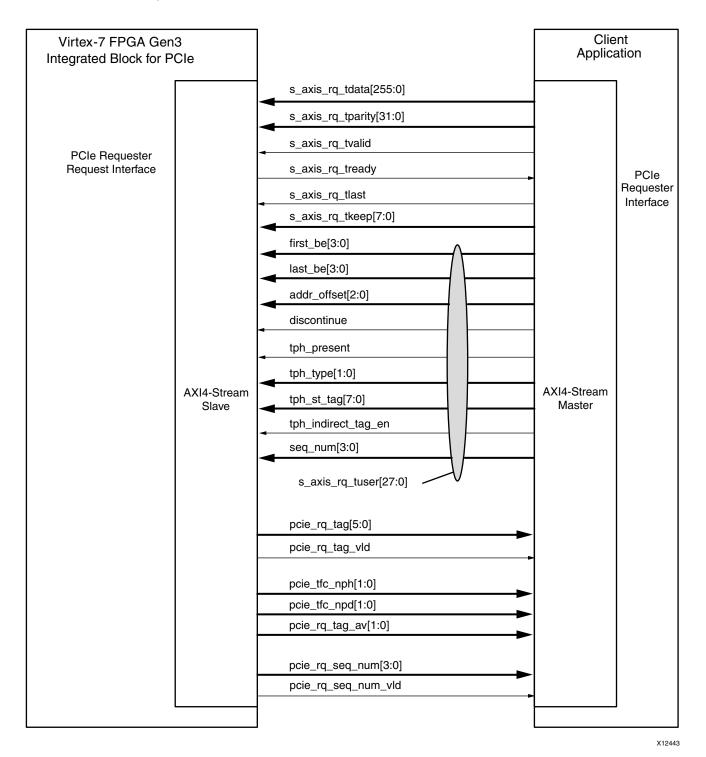


Figure 3-30: Requester Request Interface

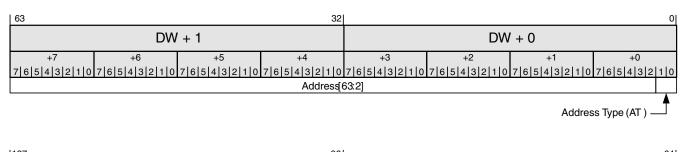
The RQ interface supports two distinct data alignment modes for transferring payloads, selected by the AXISTEN_IF_RQ_ALIGNMENT_MODE parameter. In the Dword-aligned mode, the client logic must provide the first Dword of the payload immediately after the last Dword of the descriptor. It must also set the bits in first_be[3:0] to indicate the valid bytes in the first Dword and the bits in last_be[3:0] (both part of the bus s_axis_rq_tuser) to indicate the valid bytes in the last Dword of the descriptor, and its first Dword can be in any of the possible Dword positions on the datapath. The client communicates the offset of the first Dword on the datapath using the addr_offset[2:0] signals in s_axis_rq_tuser. As in the case of the Dword-aligned mode, the client must also set the bits in first_be[3:0] to indicate the valid bytes in the last Dword of the rescriptor, and the bits in s_axis_rq_tuser. As in the case of the Dword-aligned mode, the client must also set the bits in first_be[3:0] to indicate the valid bytes in the first Dword and the bits in last_be[3:0] to indicate the valid bytes in the first Dword and the bits in first_be[3:0] to indicate the valid bytes in the first Dword and the bits in last_be[3:0] to indicate the valid bytes in the first Dword and the bits in last_be[3:0] to indicate the valid bytes in the last Dword of the payload.

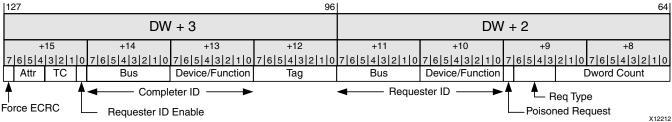
When the Transaction Processing Hint Capability is enabled in the Integrated Block, the client can provide an optional Hint with any memory transaction using the tph_* signals included in the s_axis_rq_tuser bus. To supply a Hint with a request, the client logic must assert tph_present in the first beat of the packet, and provide the TPH Steering Tag and Steering Tag Type on tph_st_tag[7:0] and tph_st_type[1:0], respectively. Instead of supplying the value of the Steering Tag to be used, the client also has the option of providing an indirect Steering Tag. This is done by setting the tph_indirect_tag[7:0], instead of the tag value. The Integrated Block then reads the tag stored in its Steering Tag Table associated with the requester Function at the offset specified in the index and inserts it in the request TLP.

Requester Request Descriptor Formats

The client must transfer each request to be transmitted on the link to the RQ interface of the Integrated Block as an independent AXI4-Stream packet. Each packet must start with a descriptor and can have payload data following the descriptor. The descriptor is always 16 bytes long, and must be sent in the first 16 bytes of the request packet. The descriptor is transferred during the first two beats on a 64-bit interface, and in the first beat on a 128-bit or 256-bit interface.

The formats of the descriptor for different request types are illustrated in Figure 3-31 through Figure 3-35. The format of Figure 3-31 applies when the request TLP being transferred is a memory read/write request, an I/O read/write request, or an Atomic Operation request. The format in Figure 3-32 is used for Configuration Requests. The format in Figure 3-33 is used for Vendor-Defined Messages (Type 0 or Type 1) only. The format in Figure 3-34 is used for all ATS messages (Invalid Request, Invalid Completion, Page Request, PRG Response). For all other messages, the descriptor takes the format shown in Figure 3-35.





Requester Request Descriptor Format for Memory, I/O, and Atomic Op Requests Figure 3-31:

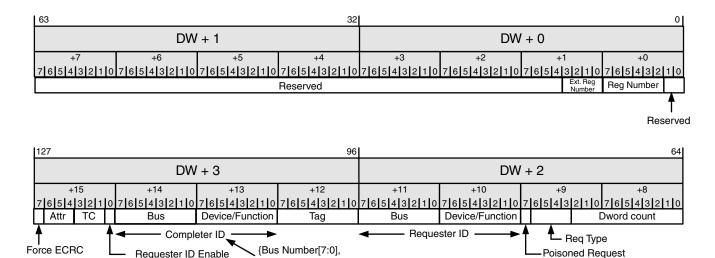


Figure 3-32: **Requester Request Descriptor Format for Configuration Requests**

Device Number[4:0], Function Number[2:0]}

Requester ID Enable

X12631

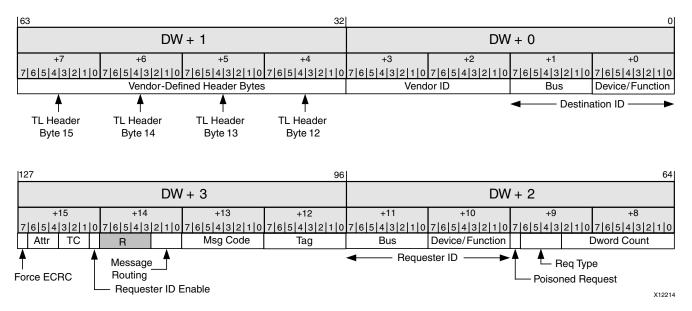


Figure 3-33: Requester Request Descriptor Format for Vendor-Defined Messages

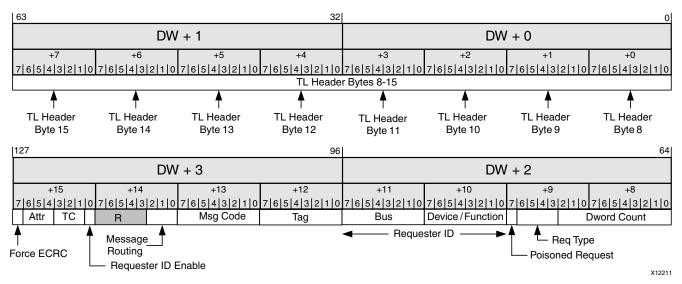


Figure 3-34: Requester Request Descriptor Format for ATS Messages

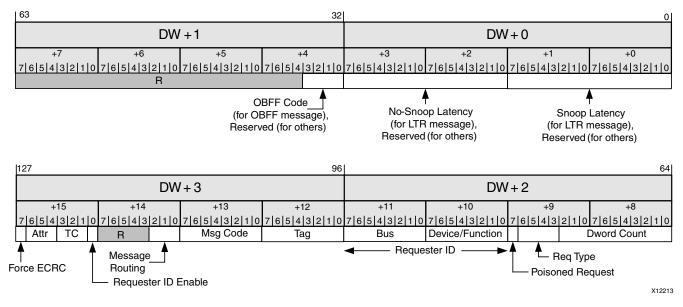


Figure 3-35: Requester Request Descriptor Format for all other Messages

Table 3-9 describes the individual fields of the completer request descriptor.

| Bit Index | Field Name | Description |
|-----------|--------------|---|
| 1:0 | Address Type | This field is defined for memory transactions and Atomic Operations only. The Integrated Block copies this field into the AT of the TL header of the request TLP. 00: Address in the request is untranslated 01: Transaction is a Translation Request 10: Address in the request is a translated address 11: Reserved |
| 63:2 | Address | This field applies to memory, I/O, and Atomic Op requests. This is the address of the first Dword referenced by the request. The client must also set the First_BE and Last_BE bits in s_axis_rq_tuser to indicate the valid bytes in the first and last Dwords, respectively. When the transaction specifies a 32-bit address, bits [63:32] of this field must be set to 0. |
| 74:64 | Dword Count | These 11 bits indicate the size of the block (in Dwords) to be read or written (for messages, size of the message payload). The valid range for Memory Write Requests is 0-256 Dwords. Memory Read Requests have a valid range of 1-1024 Dwords. For I/O accesses, the Dword count is always 1. For a zero length memory read/write request, the Dword count must be 1, with the First_BE bits set to all zeros. The Integrated Block does not check the setting of this field against the actual length of the payload supplied (for requests with payload), nor against the maximum payload size or read request size settings of the Integrated Block. |

Table 3-9: Requester Request Descriptor Fields

www.xilinx.com

| Bit Index | Field Name | Description |
|-----------|--|---|
| 78:75 | Request Type | Identifies the transaction type. The transaction types and their encodings are listed in Table 3-4. |
| 79 | Poisoned Request | This bit can be used by the client to poison the request TLP being sent. This feature is supported on all request types except Type 0 and Type 1 Configuration Write Requests. This bit must be set to 0 for all requests, except when the client has detected an error in the block of data following the descriptor and wants to communicate this information using the Data Poisoning feature of PCI Express. This feature is supported on all request types except Type 0 and Type 1 Configuration Write Requests. |
| 87:80 | Requester Function/Device Number | Function number of the Requester Function. When ARI is in use, all 8 bits of this field must be set to the Function number. Otherwise, bits [84:82] must be set to the completer Function number. When ARI is not in use, and the Integrated Block is configured as a Root Complex, the client must supply the 5-bit Device Number of the requester on bits [87:83]. When ARI is not use, and the Integrated Block is configured as an Endpoint, the client can optionally supply a 5-bit Device Number of the requester on bits [87:83]. The client must set the Requester ID Enable bit in the descriptor if a Device Number is supplied on bits [87:83]. This value is used by the Integrated Block when sending the Request TLP, instead of the stored value of the Device Number captured by the Integrated Block from Configuration Requests. |
| 95:88 | Requester Bus Number | Bus number associated with the requester Function. When the Integrated Block is configured as a Root Complex, the client must supply the 8-bit Bus Number of the requester in this field. When the Integrated Block is configured as an Endpoint, the client can optionally supply a Bus Number in this field. The client must set the Requester ID Enable bit in the descriptor if a Bus Number is supplied in this field. This value is used by the Integrated Block when sending the Request TLP, instead of the stored value of the Bus Number captured by the Integrated Block from Configuration Requests. |
| 103:96 | Tag | PCIe Tag associated with the request. For Posted transactions, the Integrated Block always uses the value from this field as the tag for the request. For Non-Posted transactions, the Integrated Block uses the value from this field if the AXISTEN_IF_ENABLE_CLIENT_TAG parameter is set (that is, when tag management is performed by the client). If this parameter is not set, tag management logic in the Integrated Block generates the tag to be used, and the value in the tag field of the descriptor is not used. |

Table 3-9: Requester Request Descriptor Fields (Cont'd)

| Bit Index | Field Name | Description |
|-----------|---------------------------|---|
| 119:104 | Completer ID | This field is applicable only to Configuration requests and messages routed by ID. For these requests, this field specifies the PCI Completer ID associated with the request (these 16 bits are divided into an 8-bit bus number, 5-bit device number, and 3-bit function number in the legacy interpretation mode. In the ARI mode, these 16 bits are treated as an 8-bit bus number + 8-bit Function number.). |
| 120 | Requester ID Enable | The purpose of this field is to enable the client to supply the bus and device numbers to be used in the Requester ID. This field is applicable only to Endpoints. If this field is 0, the Integrated Block uses the captured values of the bus and device numbers to form the Requester ID. If this input is 1, the Integrated Block uses the bus and device numbers supplied by the client in the descriptor to form the Requester ID. |
| 123:121 | Transaction Class (TC) | PCIe Transaction Class (TC) associated with the request. |
| 126:124 | Attributes | These bits provide the setting of the Attribute bits associated with the request. Bit 124 is the No Snoop bit and bit 125 is the Relaxed Ordering bit. Bit 126 is the ID-Based Ordering bit, and can be set only for memory requests and messages. The Integrated Block forces the attribute bits to 0 in the request sent on the link if the corresponding attribute is not enabled in the Function's PCI Express Device Control Register. |
| 127 | Force ECRC | Force ECRC insertion. Setting this bit to 1 forces the Integrated Block to append a TLP Digest containing ECRC to the Request TLP, even when ECRC is not enabled for the Function sending request. |
| 15:0 | Snoop Latency | This field is defined for LTR messages only. It provides the value of the 16-bit Snoop Latency field in the TLP header of the message. |
| 31:16 | No-Snoop Latency | This field is defined for LTR messages only. It provides the value of the 16-bit No-Snoop Latency field in the TLP header of the message. |
| 35:32 | OBFF Code | The OBFF Code field is used to distinguish between various OBFF cases: 1111b: "CPU Active" – System fully active for all device actions including bus mastering and interrupts 0001b: "OBFF" – System memory path available for device memory read/write bus master activities 0000b: "Idle" – System in an idle, low power state All other codes are reserved. |

| Table 3-9: | Requester R | leguest Descri | ptor Fields (Cont'd) |
|------------|--------------------|----------------|----------------------|
| | | | |

| Bit Index | Field Name | Description |
|-----------|--------------------------|--|
| 111:104 | Message Code | This field is defined for all messages. It contains the 8-bit Message Code to be set in the TL header. Appendix F of the <i>PCI Express Base Specification, rev. 3.0</i> provides a complete list of the supported Message Codes. |
| 114:112 | Message Routing | This field is defined for all messages. The Integrated Block copies these bits into the 3-bit Routing field r[2:0] of the TLP header of the Request TLP. |
| 15:0 | Destination ID | This field applies to Vendor-Defined Messages only. When the message is routed by ID (that is, when the Message Routing field is 010 binary), this field must be set to the Destination ID of the message. |
| 63:32 | Vendor-Defined Header | This field applies to Vendor-Defined Messages only. It is copied into Dword 3 of the TLP header. |
| 63:0 | ATS Header | This field is applicable to ATS messages only. It contains the bytes that the Integrated Block copies into Dwords 2 and 3 of the TLP header. |

Table 3-9: Requester Request Descriptor Fields (Cont'd)

Requester Memory Write Operation

In both Dword-aligned, the transfer starts with the sixteen descriptor bytes, followed immediately by the payload bytes. The client must keep the s_axis_rq_tvalid signal asserted over the duration of the packet. The Integrated Block treats the deassertion of s_axis_rq_tvalid during the packet transfer as an error, and nullifies the corresponding Request TLP transmitted on the link to avoid data corruption.

The client must also assert the $s_axis_rq_tlast$ signal in the last beat of the packet. The Integrated Block can deassert $s_axis_rq_tready$ in any cycle if it is not ready to accept data. The client must not change the values on the RQ interface during cycles when the Integrated Block has deasserted $s_axis_rq_tready$. The AXI4-Stream interface signals $m_axis_cq_tkeep$ (one per Dword position) must be set to indicate the valid Dwords in the packet including the descriptor and any null bytes inserted between the descriptor and the payload. That is, the tkeep bits must be set to 1 contiguously from the first Dword of the descriptor until the last Dword of the packet, when the packet does not fill the entire width of the interface.

The requester request interface also includes the First Byte Enable and the Last Enable bits in the s_axis_rq_tuser bus. These must be set in the first beat of the packet, and provides information of the valid bytes in the first and last Dwords of the payload.

The client must limit the size of the payload transferred in a single request to the maximum payload size configured in the Integrated Block, and must ensure that the payload does not cross a 4 Kbyte boundary. For memory writes of two Dwords or less, the 1s in first_be and last_be can be non-contiguous. For the special case of a zero-length memory write request, the client must provide a dummy one-Dword payload with first_be and

last_be both set to all 0s. In all other cases, the 1 bits in first_be and last_be must be
contiguous.

The timing diagrams in Figure 3-36, Figure 3-37, and Figure 3-38 illustrate the Dword-aligned transfer of a memory write request from the client across the requester request interface, when the interface width is configured as 64, 128, and 256 bits, respectively. For illustration purposes, the size of the data block being written into client memory is assumed to be *n* Dwords, for some n = k * 32 + 29, k > 0.

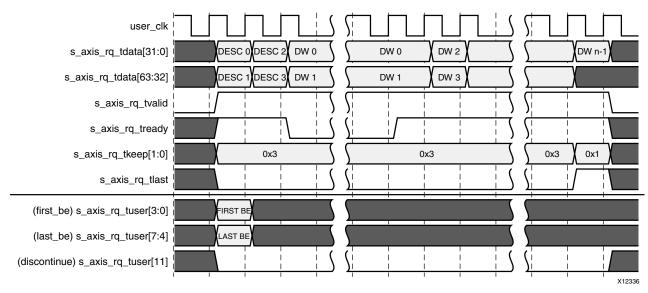


Figure 3-36: Memory Write Transaction on the Requester Request Interface (Dword-Aligned Mode, Interface Width = 64 Bits)

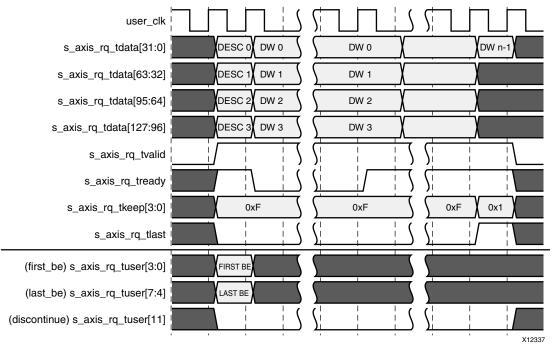


Figure 3-37: Memory Write Transaction on the Requester Request Interface (Dword-Aligned Mode, Interface Width = 128 Bits)

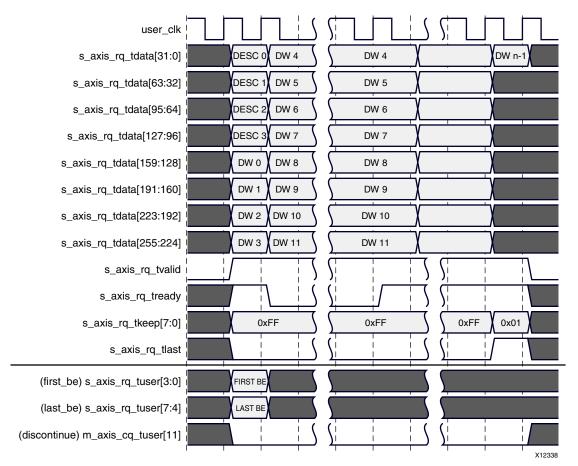


Figure 3-38: Memory Write Transaction on the Requester Request Interface (Dword-Aligned Mode, Interface Width = 256 Bits)

The timing diagrams in Figure 3-39, Figure 3-40, and Figure 3-41 illustrate the address-aligned transfer of a memory write request from the client across the RQ interface, when the interface width is configured as 64, 128, and 256 bits, respectively. For illustration purposes, the starting Dword offset of the data block being written into client memory is assumed to be (m * 32 + 1), for some integer m > 0. Its size is assumed to be n Dwords, for some n = k * 32 + 29, k > 0.

In the address-aligned mode, the delivery of the payload always starts in the beat following the last byte of the descriptor. The first Dword of the payload can appear at any Dword position. The client must communicate the offset of the first Dword of the payload on the datapath using the addr_offset[2:0] signal in s_axis_rq_tuser. The client must also set the bits in first_be[3:0] to indicate the valid bytes in the first Dword of the payload.

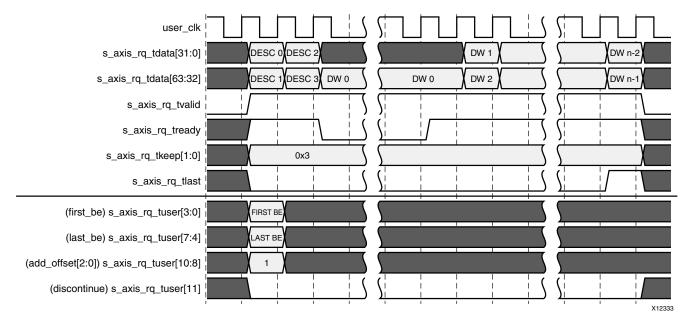


Figure 3-39: Memory Write Transaction on the Requester Request Interface (Address-Aligned Mode, Interface Width = 64 Bits)

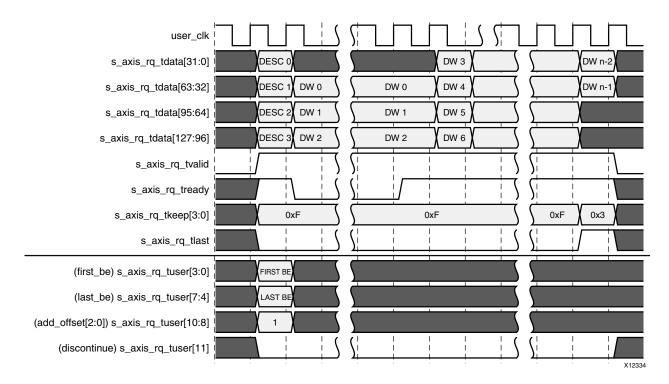
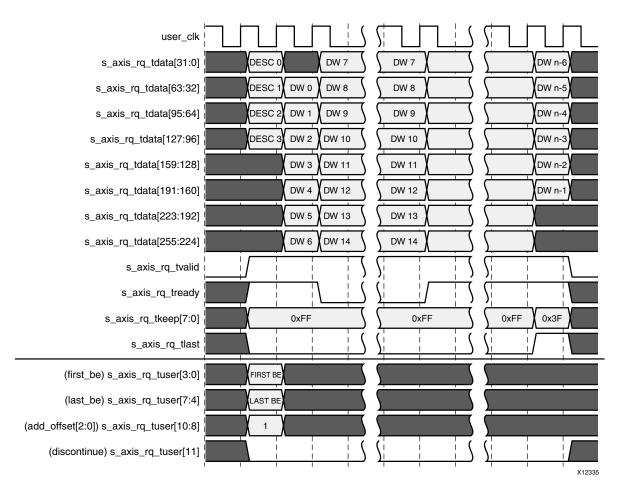
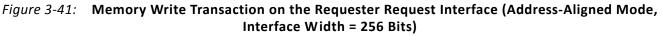


Figure 3-40: Memory Write Transaction on the Requester Request Interface (Address-Aligned Mode, Interface Width = 128 Bits)





Non-Posted Transactions with No Payload

Non-Posted transactions with no payload (memory read requests, I/O read requests, Configuration read requests) are transferred across the RQ interface in the same manner as a memory write request, except that the AXI4-Stream packet contains only the 16-byte descriptor. The timing diagrams in Figure 3-42, Figure 3-43, and Figure 3-44 illustrate the transfer of a memory read request across the RQ interface, when the interface width is configured as 64, 128, and 256 bits, respectively. The packet occupies two consecutive beats on the 64-bit interface, while it is transferred in a single beat on the 128- and 256-bit interfaces. The s_axis_rq_tvalid signal must remain asserted over the duration of the packet. The Integrated Block can deassert s_axis_rq_tready to prolong the beat. The s_axis_rq_tlast signal must be set in the last beat of the packet, and the bits in s_axis_rq_tkeep[7:0] must be set in all Dword positions where a descriptor is present.

The valid bytes in the first and last Dwords of the data block to be read must be indicated using first_be[3:0] and last_be[3:0], respectively. For the special case of a zero-length memory read, the length of the request must be set to one Dword, with both first_be[3:0] and last_be[3:0] set to all 0s. Additionally when in address-aligned

mode, addr_offset[2:0] in s_axis_rq_tuser specifies the desired starting alignment of data returned on the Requester Completion interface. The alignment is not required to be correlated to the address of the request.

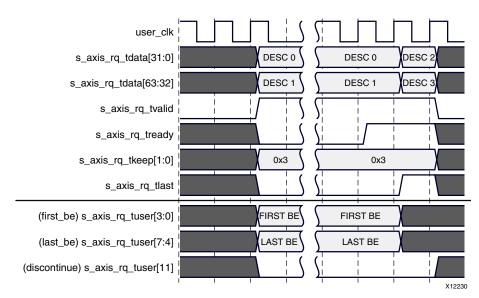


Figure 3-42: Memory Read Transaction on the Requester Request Interface (Interface Width = 64 Bits)

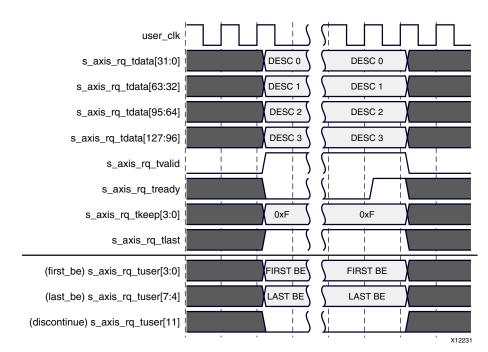


Figure 3-43: Memory Read Transaction on the Requester Request Interface (Interface Width = 128 Bits)

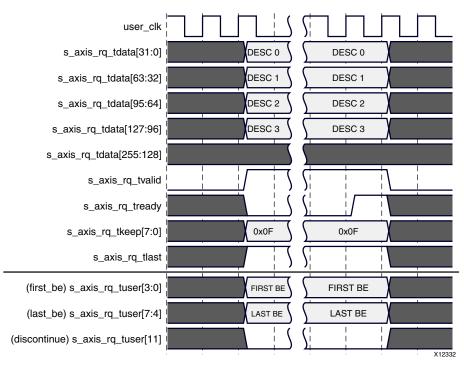


Figure 3-44: Memory Read Transaction on the Requester Request Interface (Interface Width = 256 Bits)

Non-Posted Transactions with a Payload

The transfer of a Non-Posted request with payload (an I/O write request, Configuration write request, or Atomic Operation request) is similar to the transfer of a memory request, with the following changes in how the payload is aligned on the datapath:

- In the Dword-aligned mode, the first Dword of the payload follows the last Dword of the descriptor, with no gaps between them.
- In the address-aligned mode, the payload must start in the beat following the last Dword of the descriptor. The payload can start at any Dword position on the datapath. The offset of its first Dword must be specified using the addr_offset[2:0] signal.

For I/O and Configuration write requests, the valid bytes in the one-Dword payload must be indicated using first_be[3:0]. For Atomic Operation requests, all bytes in the first and last Dwords are assumed valid.

Message Requests on the Requester Interface

The transfer of a message on the RQ interface is similar to that of a memory write request, except that a payload might not always be present. The transfer starts with the 128-bit descriptor, followed by the payload, if present. When the Dword-aligned mode is in use, the first Dword of the payload must immediately follow the descriptor. When the address-alignment mode is in use, the payload must start in the beat following the descriptor, and must be aligned to byte lane 0. The addr_offset input to the Integrated

Block must be set to 0 for messages when the address-aligned mode is in use. The Integrated Block determines the end of the payload from s_axis_rq_tlast and s_axis_rq_tkeep signals. The First Byte Enable and Last Byte Enable bits (first_be and last_be) are not used for message requests.

Aborting a Transfer

For any request that includes an associated payload, the client can abort the request at any time during the transfer of the payload by asserting the discontinue signal in the s_axis_rq_tuser bus. The Integrated Block nullifies the corresponding TLP on the link to avoid data corruption.

The client can assert this signal in any cycle during the transfer, when the request being transferred has an associated payload. The client can either choose to terminate the packet prematurely in the cycle where the error was signaled (by asserting s_axis_rq_tlast), or can continue until all bytes of the payload are delivered to the Integrated Block. In the latter case, the Integrated Block treats the error as sticky for the following beats of the packet, even if the client deasserts the discontinue signal before reaching the end of the packet.

The discontinue signal can be asserted only when s_axis_rq_tvalid is High. The Integrated Block samples this signal when s_axis_rq_tvalid and s_axis_rq_tready are both High. Thus, after assertion, the discontinue signal should not be deasserted until s_axis_rq_tready is High.

When the Integrated Block is configured as an Endpoint, this error is reported by the Integrated Block to the Root Complex it is attached to, as an Uncorrectable Internal Error using the Advanced Error Reporting (AER) mechanisms.

Tag Management for Non-Posted Transactions

The requester side of the Integrated Block maintains the state of all pending Non-Posted transactions (memory reads, I/O reads and writes, configuration reads and writes, Atomic Operations) initiated by the client, so that the completions returned by the targets can be matched against the corresponding requests. The state of each outstanding transaction is held in a Split Completion Table in the requester side of the interface, which has a capacity of 64 Non-Posted transactions. The returning Completions are matched with the pending requests using a 6-bit tag. There are two options for management of these tags.

• Internal Tag Management: This mode of operation is selected by setting the AXISTEN_IF_ENABLE_CLIENT_TAG parameter to FALSE, which is the default setting for the core. In this mode, logic within the Integrated Block is responsible for allocating the tag for each Non-Posted request initiated from the requester side. The Integrated Block maintains a list of free tags and assigns one of them to each request when the client initiates a Non-Posted transaction, and communicates the assigned tag value to the client via the output pcie_rq_tag[5:0]. The value on this bus is valid when the Integrated Block asserts pcie_rq_tag_v1d. The client logic must copy this tag so that any Completions delivered by the Integrated Block in response to the request can be matched to the request.

In this mode, logic within the Integrated Block checks for the Split Completion Table full condition, and backpressures a Non-Posted request from the client (using s_axis_rq_tready) if the total number of Non-Posted requests currently outstanding has reached its limit (64).

• External Tag Management: This mode of operation is selected by setting the AXISTEN_IF_ENABLE_CLIENT_TAG parameter to 1. In this mode, the client logic is responsible for allocating the tag for each Non-Posted request initiated from the requester side. The client logic must choose the tag value without conflicting with the tags of all other Non-Posted transactions outstanding at that time, and must communicate this chosen tag value to the Integrated Block via the request descriptor. The Integrated Block still maintains the outstanding requests in its Split Completion Table and matches the incoming Completions to the request, but does not perform any checks for the uniqueness of the tags, or for the Split Completion Table full condition.

When internal tag management is in use, the Integrated Block asserts pcie_rq_tag_vld for one cycle for each Non-Posted request, after it has placed its allocated tag on pcie_rq_tag[5:0]. There can be a delay of several cycles between the transfer of the request on the RQ interface and the assertion of pcie_rq_tag_vld by the Integrated Block to provide the allocated tag for the request. The client can, meanwhile, continue to send new requests. The tags for requests are communicated on the pcie_rq_tag bus in FIFO order, so it is easy for the client to associate the tag value with the request it transferred. A tag is reused when the EOF of the last completion of a split completion is accepted by the client.

Avoiding Head-of-Line Blocking for Posted Requests

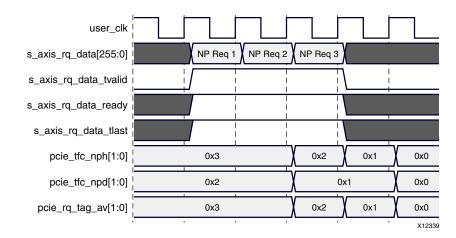
The Integrated Block can hold a Non-Posted request received on its RQ interface for lack of transmit credit or lack of available tags. This could potentially result in head-of-line (HOL) blocking for Posted transactions. The Integrated Block provides a mechanism for the client logic to avoid this situation through these signals:

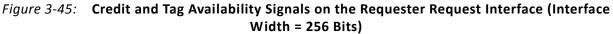
- pcie_tfc_nph_av[1:0]: These outputs indicate the Header Credit currently
 available for Non-Posted requests, where:
 - 00 = no credit available
 - 01 = 1 credit
 - 10 = 2 credits
 - 11 = 3 or more credits
- pcie_tfc_npd_av[1:0]: These outputs indicate the Data Credit currently available
 for Non-Posted requests, where:

- 00 = no credit available
- 01 = 1 credit
- 10 = 2 credits
- 11 = 3 or more credits

The client logic can optionally check these outputs before transmitting Non-Posted requests. Because of internal pipeline delays, the information on these outputs is delayed by two user clock cycles from the cycle in which the last byte of the descriptor is transferred on the RQ interface. Thus the client logic must adjust these values, taking into account any Non-Posted requests transmitted in the two previous clock cycles. Figure 3-45 illustrates the operation of these signals for the 256-bit interface. In this example, the Integrated Block initially had three Non-Posted Header Credits and two Non-Posted Data Credits, and had three free tags available for allocation. Request 1 from the client had a one-Dword payload, and therefore consumed one header and data credit each, and also one tag. Request 2 in the next clock cycle consumed one header credit, but no data credit. When the client presents Request 3 in the following clock cycle, it must adjust the available credit and available tag count by taking into account requests 1 and 2. If Request 3 consumes one header credit, both available credits are 0 two cycles later, as also the number of available tags.

Figure 3-46 and Figure 3-47 illustrate the timing of the credit and tag available signals for the same example, for interface width of 128 bits and 64 bits, respectively.





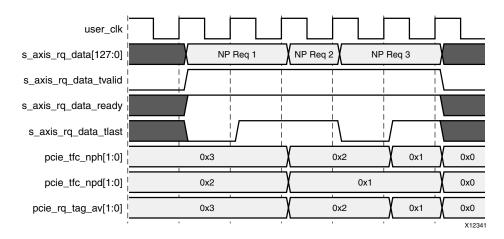
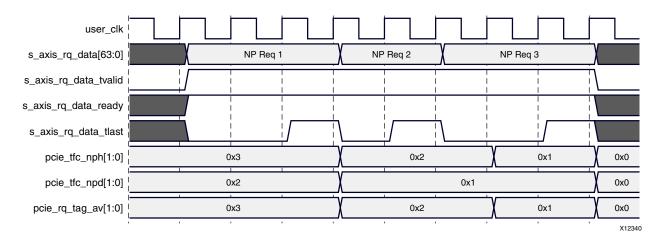
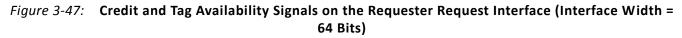


Figure 3-46: Credit and Tag Availability Signals on the Requester Request Interface (Interface Width = 128 Bits)





Maintaining Transaction Order

The Integrated Block does not change the order of requests received from the client on its requester interface when it transmits them on the link. In cases where the client would like to have precise control of the order of transactions sent on the RQ interface and the CC interface (typically to avoid Completions from passing Posted requests when using strict ordering), the Integrated Block provides a mechanism for the client to monitor the progress of a Posted transaction through its pipeline, so that it can determine when to schedule a Completion on the completer completion interface without the risk of passing a specific Posted request transmitted from the requester request interface,

When transferring a Posted request (memory write transactions or messages) across the requester request interface, the client can provide an optional 4-bit sequence number to the Integrated Block on its seq_num[3:0] input within s_axis_rq_tuser. The sequence

number must be valid in the first beat of the packet. The client can then monitor the pcie_rq_seq_num[3:0] output of the Integrated Block for this sequence number to appear. When the transaction has reached a stage in the internal transmit pipeline of the Integrated Block where a Completion cannot pass it, the Integrated Block asserts pcie_rq_seq_num_valid for one cycle and provides the sequence number of the Posted request on the pcie_rq_seq_num[3:0] output. Any Completions transmitted by the Integrated Block after the sequence number has appeared on pcie_rq_seq_num[3:0] cannot pass the Posted request in the internal transmit pipeline.

Requester Completion Interface Operation

Completions for requests generated by client logic are presented on the Integrated Block's Request Completion (RC) interface. See Figure 3-48 for an illustration of signals associated with the requester completion interface. When straddle is not enabled, the Integrated Block delivers each TLP on this interface as an AXI4-Stream packet. The packet starts with a 96-bit descriptor, followed by data in the case of Completions with a payload.

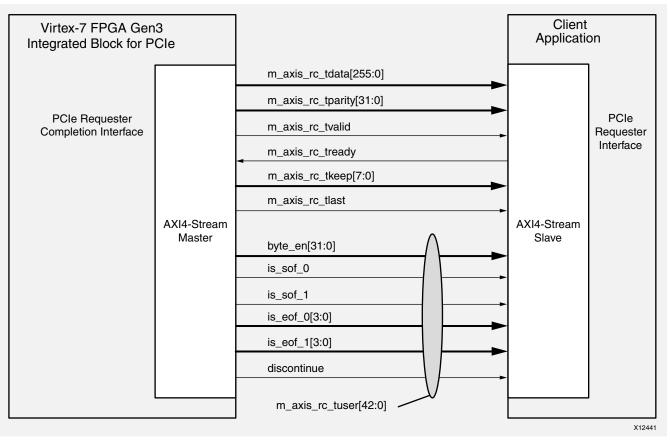


Figure 3-48: Requester Completion Interface

The RC interface supports two distinct data alignment modes for transferring payloads, selected by the AXISTEN_IF_RC_ALIGNMENT_MODE parameter. In the Dword-aligned mode, the Integrated Block transfers the first Dword of the Completion payload

immediately after the last Dword of the descriptor. In the address-aligned mode, the Integrated Block starts the payload transfer in the beat following the last Dword of the descriptor, and its first Dword can be in any of the possible Dword positions on the datapath. The alignment of the first Dword of the payload is determined by address offset provided by the client when it sent the request to the Integrated Block (that is, the setting of the addr_offset[2:0] input of the RQ interface). Thus, the address-aligned mode can be used on the RC interface only if the RQ interface is also configured to use the address-aligned mode.

Requester Completion Descriptor Format

The RC interface of the Integrated Block sends completion data received from the link to the client application as AXI4-Stream packets. Each packet starts with a descriptor and can have payload data following the descriptor. The descriptor is always 12 bytes long, and is sent in the first 12 bytes of the completion packet. The descriptor is transferred during the first two beats on a 64-bit interface, and in the first beat on a 128- or 256-bit interface. When the completion data is split into multiple Split Completions, the Integrated Block sends each Split Completion as a separate AXI4-Stream packet, with its own descriptor.

The format of the Requester Completion descriptor is illustrated in Figure 3-49. The individual fields of the RC descriptor are described in Table 3-10.

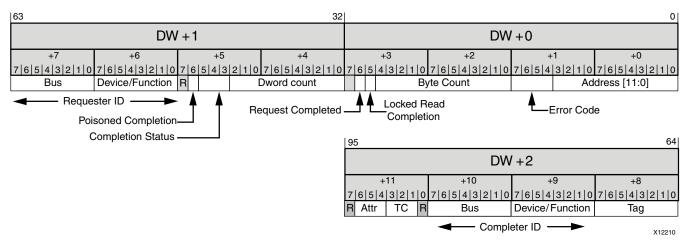


Figure 3-49: Requester Completion Descriptor Format

| Table 3-10: | Requester Completion Descriptor Fields | |
|-------------|---|--|
| | | |

| Bit Index | Field Name | Description |
|-----------|---------------|---|
| 11:0 | Lower Address | This field provides the 12 least significant bits of the first byte referenced by the request. The Integrated Block returns this address from its Split Completion Table, where it stores the address and other parameters of all pending Non-Posted requests on the requester side. When the Completion delivered has an error, only bits [6:0] of the address should be considered valid. This is a byte-level address. |

| Bit Index | Field Name | Description |
|-----------|---------------------------|---|
| 15:12 | Error Code | Completion error code. These three bits encode error conditions detected from error checking performed by the Integrated Block on received Completions. Its encodings are: 0000: Normal termination (all data received). 0001: The Completion TLP is Poisoned. 0010: Request terminated by a Completion with UR, CA or CRS status. 0011: Request terminated by a Completion with no data, or the byte count in the Completion was higher than the total number of bytes expected for the request. 0100: The current Completion being delivered has the same tag of an outstanding request, but its Requester ID, TC, or Attr fields did not match with the parameters of the outstanding request. 0101: Error in starting address. The low address bits in the Completion TLP header did not match with the starting address of the next expected byte for the request. 0110: Invalid tag. This Completion does not match the tags of any outstanding request. 1001: Request terminated by a Completion timeout. The other fields in the descriptor, except bit [30], the requester Function [55:48], and the tag field [71:64], are invalid in this case, because the descriptor does not correspond to a Completion TLP. 1000: Request terminated by a Function-Level Reset (FLR) targeted at the Function that generated the request. The other fields in the descriptor, except bit [30], the requester Function [55:48], and the tag field [71:64], are invalid in this case, because the descriptor does not correspond to a Completion TLP. |
| 28:16 | Byte Count | These 13 bits can have values in the range of 0 – 4096 bytes. If a Memory Read Request is completed using a single Completion, the Byte Count value indicates Payload size in bytes. This field must be set to 4 for I/O read Completions and I/O write Completions. The byte count must be set to 1 while sending a Completion for a zero-length memory read, and a dummy payload of 1 Dword must follow the descriptor. The byte count must be set to 0 when sending a UR or CA Completion. For each Memory Read Completion, the Byte Count field must indicate the remaining number of bytes required to complete the Request, including the number of bytes returned with the Completion. If a Memory Read Request is completed using multiple Completions, the Byte Count value for each successive Completion is the value indicated by the preceding Completion minus the number of bytes returned with the preceding Completion. |
| 29 | Locked Read Completion | This bit is set to 1 when the Completion is in response to a Locked Read request. It is set to 0 for all other Completions. |

Table 3-10: Requester Completion Descriptor Fields (Cont'd)

| Bit Index | Field Name | Description |
|-----------|---------------------------|---|
| 30 | Request Completed | The Integrated Block asserts this bit in the descriptor of the last Completion of a request. The assertion of the bit can indicate normal termination of the request (because all data has been received) or abnormal termination because of an error condition. The client logic can use this indication to clear its outstanding request status. When tags are assigned by the client, the client logic should not re-assign a tag allocated to a request until it has received a Completion Descriptor from the Integrated Block with a matching tag field and the Request Completed bit set to 1. |
| 42:32 | Dword Count | These 11 bits indicate the size of the payload of the current packet in Dwords. Its range is 0 - 1K Dwords. This field is set to 1 for I/O read Completions and 0 for I/O write Completions. The Dword count is also set to 1 while transferring a Completion for a zero-length memory read. In all other cases, the Dword count corresponds to the actual number of Dwords in the payload of the current packet. |
| 45:43 | Completion Status | These bits reflect the setting of the Completion Status field of the received Completion TLP. The valid settings are: 000: Successful Completion 001: Unsupported Request (UR) 010: Configuration Request Retry Status (CRS) 100: Completer Abort (CA) |
| 46 | Poisoned Completion | This bit is set to indicate that the Poison bit in the Completion TLP was set. Data in the packet should then be considered corrupted. |
| 63:48 | Requester ID | PCI Requester ID associated with the Completion. |
| 71:64 | Tag | PCIe Tag associated with the Completion. |
| 87:72 | Completer ID | Completer ID received in the Completion TLP. (These 16 bits are divided into an 8-bit bus number, 5-bit device number, and 3-bit function number in the legacy interpretation mode. In the ARI mode, these 16 bits must be treated as an 8-bit bus number + 8-bit Function number.). |
| 91:89 | Transaction Class (TC) | PCIe Transaction Class (TC) associated with the Completion. |
| 94:92 | Attributes | PCIe attributes associated with the Completion. Bit 92 is the No Snoop bit, bit 93 is the Relaxed Ordering bit, and bit 94 is the ID-Based Ordering bit. |

Table 3-10: Requester Completion Descriptor Fields (Cont'd)

Transfer of Completions with no Data

The timing diagrams in Figure 3-50, Figure 3-51, and Figure 3-52 illustrate the transfer of a Completion TLP received from the link with no associated payload across the RC interface, when the interface width is configured as 64, 128, and 256 bits, respectively. The timing diagrams in this section assume that the Completions are not straddled on the 256-bit interface. The straddle feature is described in Straddle Option for 256-Bit Interface, page 149.

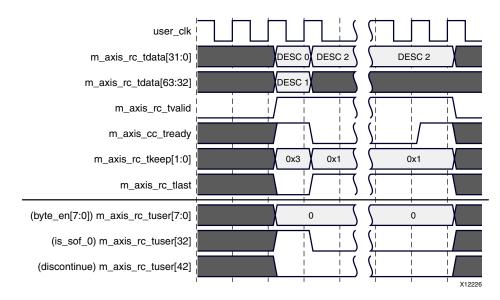


Figure 3-50: Transfer of a Completion with no Data on the Requester Completion Interface (Interface Width = 64 Bits)

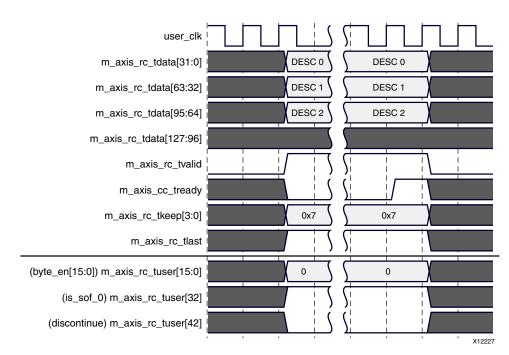


Figure 3-51: Transfer of a Completion with no Data on the Requester Completion Interface (Interface Width = 128 Bits)

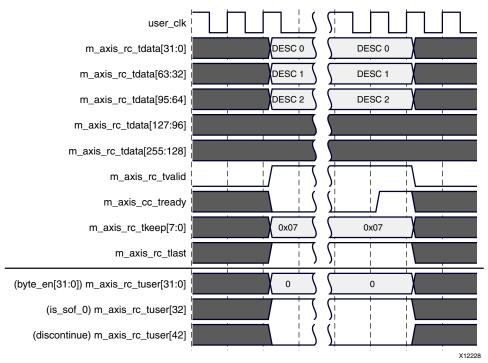


Figure 3-52: Transfer of a Completion with no Data on the Requester Completion Interface (Interface Width = 256 Bits)

The entire transfer of the Completion TLP takes only a single beat on the 256- and 128-bit interfaces, and two beats on the 64-bit interface. The Integrated Block keeps the m_axis_rc_tvalid signal asserted over the duration of the packet. The client can prolong a beat at any time by deasserting m_axis_rc_tready. The AXI4-Stream interface signals m_axis_rc_tkeep (one per Dword position) indicate the valid descriptor Dwords in the packet. That is, the tkeep bits are set to 1 contiguously from the first Dword of the descriptor until its last Dword. During the transfer of a packet, the tkeep bits can be 0 only in the last beat of the packet. The m_axis_cq_tlast signal is always asserted in the last beat of the packet.

The m_axi_cq_tuser bus also includes an is_sof_0 signal, which is asserted in the first beat of every packet. The client can optionally use this signal to qualify the start of the descriptor on the interface. No other signals within m_axi_cq_tuser are relevant to the transfer of Completions with no data, when the straddle option is not in use.

Transfer of Completions with Data

The timing diagrams in Figure 3-53, Figure 3-54, and Figure 3-55 illustrate the Dword-aligned transfer of a Completion TLP received from the link with an associated payload across the RC interface, when the interface width is configured as 64, 128, and 256 bits, respectively. For illustration purposes, the size of the data block being written into client memory is assumed to be *n* Dwords, for some n = k * 32 + 28, k > 0. The timing diagrams in this section assume that the Completions are not straddled on the 256-bit interface. The straddle feature is described in Straddle Option for 256-Bit Interface, page 149.

In the Dword-aligned mode, the transfer starts with the three descriptor Dwords, followed immediately by the payload Dwords. The entire TLP, consisting of the descriptor and payload, is transferred as a single AXI4-Stream packet. Data within the payload is always a contiguous stream of bytes when the length of the payload exceeds two Dwords. The positions of the first valid byte within the first Dword of the payload and the last valid byte in the last Dword can then be determined from the Lower Address and Byte Count fields of the Request Completion Descriptor. When the payload size is two Dwords or less, the valid bytes in the payload cannot be contiguous. In these cases, the client must store the First Byte Enable and the Last Byte Enable fields associated with each request sent out on the RQ interface and use them to determine the valid bytes in the completion payload. The client can optionally use the byte enable outputs byte_en[31:0] within the m_axi_cq_tuser bus to determine the valid bytes in the payloads.

The Integrated Block keeps the m_axis_rc_tvalid signal asserted over the entire duration of the packet. The client can prolong a beat at any time by deasserting m_axis_rc_tready. The AXI4-Stream interface signals m_axis_rc_tkeep (one per Dword position) indicate the valid Dwords in the packet including the descriptor and any null bytes inserted between the descriptor and the payload. That is, the tkeep bits are set to 1 contiguously from the first Dword of the descriptor until the last Dword of the payload. During the transfer of a packet, the tkeep bits can be 0 only in the last beat of the packet, when the packet does not fill the entire width of the interface. The m_axis_rc_tlast signal is always asserted in the last beat of the packet.

The m_axi_rc_tuser bus provides several informational signals that can be used to simplify the logic associated with the client side of the interface, or to support additional features. The is_sof_0 signal is asserted in the first beat of every packet, when its descriptor is on the bus. The byte enable outputs byte_en[31:0] (one per byte lane) indicate the valid bytes in the payload. These signals are asserted only when a valid payload byte is in the corresponding lane (it is not asserted for descriptor or null bytes). The asserted byte enable bits are always contiguous from the start of the payload, except when payload size is 2 Dwords or less. For Completion payloads of two Dwords or less, the 1s on byte_en might not be contiguous. Another special case is that of a zero-length memory read, when the Integrated Block transfers a one-Dword payload with the byte_en bits all set to 0. Thus, the client logic can, in all cases, use the byte_en signals directly to enable the writing of the associated bytes into memory.

The is_sof_1, is_eof_0[3:0], and is_eof_1[3:0] signals within the m_axis_rc_tuser bus are not to be used for 64-bit and 128-bit interfaces, and for 256-bit interfaces when the straddle option is not enabled.

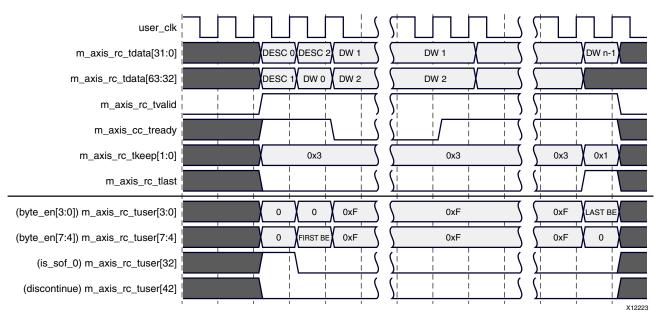
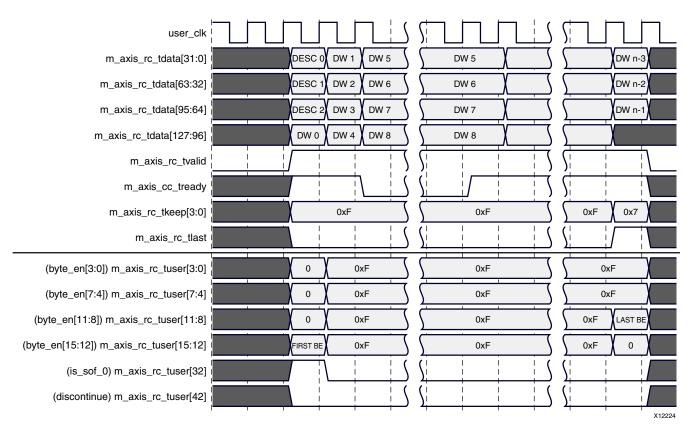


Figure 3-53: Transfer of a Completion with Data on the Requester Completion Interface (Dword-Aligned Mode, Interface Width = 64 Bits)





www.xilinx.com

| upper alle | | | |
|---|-----------------|----------|--------------|
| user_clk | | | |
| m_axis_rc_tdata[31:0] | DESC 0 DW 5 | DW 5 | DW n-7 |
| m_axis_rc_tdata[63:32] | DESC 1 DW 6 | | DW n-6 |
| m_axis_rc_tdata[95:64] | DESC 2 DW 7 | | DW n-5 |
| m_axis_rc_tdata[127:96] | DW 0 DW 8 | DW 8 | DW n-4 |
| m_axis_rc_tdata[159:128] | DW 1 DW 9 | DW 9 | DW n-3 |
| m_axis_rc_tdata[191:160] | DW 2 DW 10 | DW 10 | DW n-2 |
| m_axis_rc_tdata[223:192] | DW 3 DW 11 | DW 11 | DW n-1 |
| m_axis_rc_tdata[255:224] | DW 4 DW 12 | DW 12 | |
| m_axis_rc_tvalid | | <u>S</u> | |
| m_axis_cc_tready | | | S |
| m_axis_rc_tkeep[7:0] | 0xFF | 0xFF | 0xFF 0x7F |
| m_axis_rc_tlast | | <u>S</u> | |
| (byte_en[3:0]) m_axis_rc_tuser[3:0] | 0 0xFF | 0xFF | 0xFF |
| (byte_en[7:4]) m_axis_rc_tuser[7:4] | 0 0xFF | 0xFF | 0xFF |
| (byte_en[11:8]) m_axis_rc_tuser[11:8] | 0 0xFF | 0xFF | 0xFF |
| (byte_en[15:12]) m_axis_rc_tuser[15:12] | FIRST BE 0xFF | 0xFF | 0xFF |
| (byte_en[19:16]) m_axis_rc_tuser[19:16] | 0xFF | 0xFF | 0xFF |
| (byte_en[23:20]) m_axis_rc_tuser[23:20] | 0xFF | 0xFF | 0xFF |
| (byte_en[27:24]) m_axis_rc_tuser[27:24] | 0xFF | 0xFF | OXFF LAST BE |
| (byte_en[31:28]) m_axis_rc_tuser[31:28] | 0xFF | 0xFF | 0xFF 0 |
| (is_sof_0) m_axis_rc_tuser[32] | | <u>\</u> | <u></u> |
| (discontinue) m_axis_rc_tuser[42] | | | S |

Figure 3-55: Transfer of a Completion with Data on the Requester Completion Interface (Dword-Aligned Mode, Interface Width = 256 Bits)

The timing diagrams in Figure 3-56, Figure 3-57, and Figure 3-58 illustrate the address-aligned transfer of a Completion TLP received from the link with an associated payload across the RC interface, when the interface width is configured as 64, 128, and 256 bits, respectively. In the example timing diagrams, the starting Dword address of the data block being transferred (as conveyed in bits [6:2] of the Lower Address field of the descriptor) is assumed to be (m * 8 + 1), for an integer m. The size of the data block is assumed to be n Dwords, for some n = k * 32 + 28, k > 0. The straddle option is not valid for address-aligned transfers, so the timing diagrams assume that the Completions are not straddled on the 256-bit interface.

www.xilinx.com

In the address-aligned mode, the delivery of the payload always starts in the beat following the last byte of the descriptor. The first byte of the payload can appear on any byte lane, based on the address of the first valid byte of the payload. The tkeep bits are set to 1 contiguously from the first Dword of the descriptor until the last Dword of the payload. The alignment of the first Dword on the data bus is determined by the setting of the addr_offset[2:0] input of the requester request interface when the client sent the request to the Integrated Block. The client can optionally use the byte enable outputs byte_en[31:0] to determine the valid bytes in the payload.

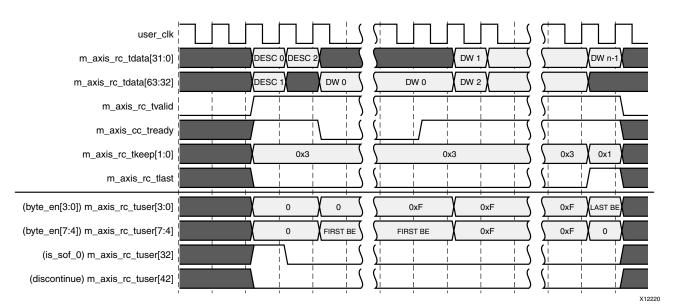


Figure 3-56: Transfer of a Completion with Data on the Requester Completion Interface (Address-Aligned Mode, Interface Width = 64 Bits)

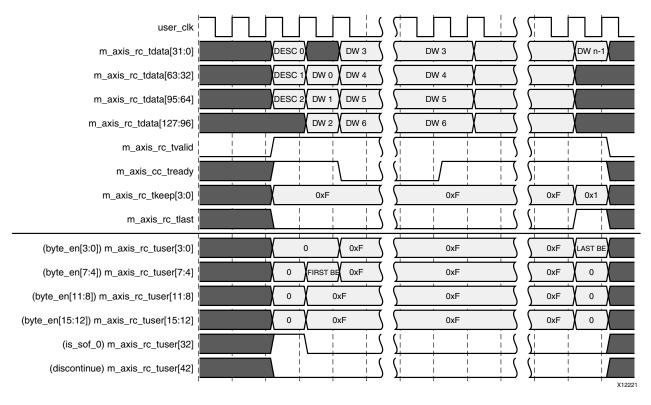
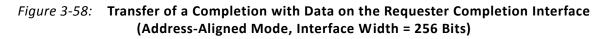


Figure 3-57: Transfer of a Completion with Data on the Requester Completion Interface (Address-Aligned Mode, Interface Width = 128 Bits)

| user_clk | | | | |
|---|--------|-----------|------------------|--|
| m_axis_rc_tdata[31:0] | DESC 0 | DW 7 | DW 7 | DW n-5 |
| m_axis_rc_tdata[63:32] | DESC 1 | W 0 DW 8 | | DW n-4 |
| m_axis_rc_tdata[95:64] | DESC 2 | W 1 DW 9 | DW 9 | DW n-3 |
| m_axis_rc_tdata[127:96] | | W 2 DW 10 | DW 10 | DW n-2 |
| m_axis_rc_tdata[159:128] | | W 3 DW 11 | DW 11 | DW n-1 |
| m_axis_rc_tdata[191:160] | | W 4 DW 12 | DW 12 | |
| m_axis_rc_tdata[223:192] | | W 5 DW 13 | DW 13 | |
| m_axis_rc_tdata[255:224] | | W 6 DW 14 | DW 14 | |
| m_axis_rc_tvalid | | | $\sum_{i=1}^{n}$ | |
| m_axis_cc_tready | | | <u>S</u> | S I |
| m_axis_rc_tkeep[7:0] | 0 | xFF | 0xFF | 0xFF 0x1F |
| m_axis_rc_tlast | | | <u><u> </u></u> | |
| (byte_en[3:0]) m_axis_rc_tuser[3:0] | 0 | 0xF | 0xF | 0xF |
| (byte_en[7:4]) m_axis_rc_tuser[7:4] | | ST BE 0xF | 0xF | 0xF |
| (byte_en[11:8]) m_axis_rc_tuser[11:8] | | 0xF | 0xF | 0xF |
| (byte_en[15:12]) m_axis_rc_tuser[15:12] | | 0xF | 0xF | 0xF |
| (byte_en[19:16]) m_axis_rc_tuser[19:16] | | 0xF | 0xF | 0xF LAST BE |
| (byte_en[23:20]) m_axis_rc_tuser[23:20] | | 0xF | 0xF | OxF 0 |
| (byte_en[27:24]) m_axis_rc_tuser[27:24] | | 0xF | 0xF | OxF 0 |
| (byte_en[31:28]) m_axis_rc_tuser[31:28] | | 0xF | 0xF | |
| (is_sof_0) m_axis_rc_tuser[32] | | | <u></u> | |
| (discontinue) m_axis_rc_tuser[42] | | | | I I I I I I I I X12222 |



Straddle Option for 256-Bit Interface

When the interface width is configured as 256 bits, the Integrated Block can start a new Completion transfer on the RC interface in the same beat when the previous Completion has ended on or before Dword position 3 on the data bus. This straddle option is enabled by setting the AXISTEN_IF_RC_STRADDLE parameter. The straddle option can be used only with the Dword-aligned mode.

When the straddle option is enabled, Completion TLPs are transferred on the RC interface as a continuous stream, with no packet boundaries (from an AXI4-Stream perspective). Thus, the m_axis_rc_tkeep and m_axis_rc_tlast signals are not useful in

determining the boundaries of Completion TLPs delivered on the interface (the Integrated Block sets m_axis_rc_tkeep to all 1s and m_axis_rc_tlast to 0 permanently when the straddle option is in use). Instead, delineation of TLPs is performed using the following signals provided within the m_axis_rc_tuser bus:

- is_sof_0: The Integrated Block drives this output High in a beat when there is at least one Completion TLP starting in the beat. The position of the first byte of this Completion TLP is determined as follows:
 - If the previous Completion TLP ended before this beat, the first byte of this Completion TLP is in byte lane 0.
 - If a previous TLP is continuing in this beat, the first byte of this Completion TLP is in byte lane 16. This is possible only when the previous TLP ends in the current beat, that is when is_eof_0[0] is also set.
- is_sof_1: The Integrated Block asserts this output in a beat when there are two Completion TLPs starting in the beat. The first TLP always starts at byte position 0 and the second TLP at byte position 16. The Integrated Block starts a second TLP at byte position 16 only if the previous TLP ended before byte position 16 in the same beat, that is only if is_eof_0[0] is also set in the same beat.
- is_eof_0[3:0]: These outputs are used to indicate the end of a Completion TLP and the position of its last Dword on the data bus. The assertion of the bit is_eof_0[0] indicates that there is at least one Completion TLP ending in this beat. When bit 0 of is_eof_0 is set, bits [3:1] provide the offset of the last Dword of the TLP ending in this beat. The offset for the last byte can be determined from the starting address and length of the TLP, or from the byte enable signals byte_en[31:0]. When there are two Completion TLPs ending in a beat, the setting of is_eof_0[3:1] is the offset of the last Dword of the first Completion TLP (in that case, its range is 0 through 3).
- is_eof_1[3:0]: The assertion of is_eof_1[0] indicates a second TLP ending in the same beat. When bit 0 of is_eof_1 is set, bits [3:1] provide the offset of the last Dword of the second TLP ending in this beat. Because the second TLP can start only on byte lane 16, it can only end at a byte lane in the range 27-31. Thus the offset is_eof_1[3:1] can only take one of two values: 6 or 7. If is_sof_1[0] is High, the signals is_eof_0[0] and is_sof_0 are also High in the same beat. If is_sof_1 is High, is_sof_0 is High. If is_eof_1 is High, is_eof_0 is High.

The timing diagram in Figure 3-59 illustrates the transfer of four Completion TLPs on the 256-bit RC interface when the straddle option is enabled. The first Completion TLP (COMPL 1) starts at Dword position 0 of Beat 1 and ends in Dword position 0 of Beat 3. The second TLP (COMPL 2) starts in Dword position 4 of the same beat. This second TLP has only a one-Dword payload, so it also ends in the same beat. The third and fourth Completion TLPs are transferred completely in Beat 4, because Completion 3 has only a one-Dword payload and Completion 4 has no payload.

| user_clk | → BEAT 1 | | | | | → BEAT 4 |
|---|----------|--------|------------------|---------|---|------------|
| m_axis_rc_tdata[31:0] | СОМРІ | | $\sum_{i=1}^{n}$ | COMPL 1 | COMPL 1 | COMPL 3 |
| m_axis_rc_tdata[63:32] | СОМРІ | | $\sum_{i=1}^{n}$ | COMPL 1 | X | COMPL 3 |
| m_axis_rc_tdata[95:64] | СОМРІ | | $\sum_{i=1}^{n}$ | COMPL 1 |) | COMPL 3 |
| m_axis_rc_tdata[127:96] | СОМРІ | | $\sum_{i=1}^{n}$ | COMPL 1 | | COMPL 3 |
| m_axis_rc_tdata[159:128] | СОМРІ | | | COMPL 1 | COMPL 2 | COMPL 4 |
| m_axis_rc_tdata[191:160] | COMPL | | $\sum_{i=1}^{n}$ | COMPL 1 | COMPL 2 | COMPL 4 |
| m_axis_rc_tdata[223:192] | COMPL | | \int | COMPL 1 | COMPL 2 | COMPL 4 |
| m_axis_rc_tdata[255:224] | СОМРІ | | $\sum_{i=1}^{n}$ | COMPL 1 | COMPL 2 | |
| m_axis_rc_tvalid | | 5 | $\sum_{i=1}^{n}$ | | | |
| m_axis_cc_tready | | | <u></u> | | | |
| (is_sof_0) m_axis_rc_tuser[32] | | | S | | | |
| (is_sof_1) m_axis_rc_tuser[33] | | | S | 1 | | |
| (is_eof_0[3:0]) m_axis_rc_tuser[37:34] | | 0 | $\sum_{i=1}^{n}$ | 0 | 0x1 | 0x7 |
| (is_eof_1[3:0]) m_axis_rc_tuser[41:38] | | 0 | $\sum_{i=1}^{n}$ | 0 | 0xF | 0xD |
| (byte_en[3:0]) m_axis_rc_tuser[3:0] | 0 | 0xF | $\sum_{i=1}^{n}$ | 0xF | LAST BE | 0 |
| (byte_en[7:4]) m_axis_rc_tuser[7:4] | 0 | 0xF | 5 | 0xF | | |
| (byte_en[11:8]) m_axis_rc_tuser[11:8] | 0 | 0xF | 5 | 0xF | Ϋ́, Ϋ́, | |
| (byte_en[15:12]) m_axis_rc_tuser[15:12] | FIRST | BE 0xF | Ś | 0xF | , i i i i i i i i i i i i i i i i i i i | FIRST BE |
| (byte_en[19:16]) | | 0xF | $\sum_{i=1}^{i}$ | 0xF | | 0 |
| (byte_en[23:20]) | | 0xF | | 0xF | | |
| (byte_en[27:24]) | | 0xF | \sum | 0xF | (0) | 0 |
| (byte_en[31:28]) | | 0xF | \sum | 0xF | FIRST BE | 0 |
| (discontinue) m_axis_rc_tuser[42] | | | | | | 1 1 X12229 |

Figure 3-59: **Transfer of Completion TLPs on the Requester Completion Interface with the Straddle Option Enabled**

Aborting a Completion Transfer

For any Completion that includes an associated payload, the Integrated Block can signal an error in the transferred payload by asserting the discontinue signal in the m_axis_rc_tuser bus in the last beat of the packet. This occurs when the Integrated Block has detected an uncorrectable error while reading data from its internal memories. The client application must discard the entire packet when it has detected the

www.xilinx.com

discontinue signal asserted in the last beat of a packet. This is also considered a fatal error in the Integrated Block.

When the straddle option is in use, the Integrated Block does not start a second Completion TLP in the same beat when it has asserted discontinue, aborting the Completion TLP ending in the beat.

Handling of Completion Errors

When a Completion TLP is received from the link, the Integrated Block matches it against the outstanding requests in the Split Completion Table to determine the corresponding request, and compares the fields in its header against the expected values to detect any error conditions. The Integrated Block then signals the error conditions in a 4-bit error code sent to the client as part of the completion descriptor. The Integrated Block also indicates the last completion for a request by setting the Request Completed bit (bit 30) in the descriptor. Table 3-11 defines the error conditions signaled by the various error codes.

| Error Code | Description |
|------------|--|
| 0000 | No errors detected. |
| 0001 | The Completion TLP received from the link was Poisoned. The client should discard any data that follows the descriptor. In addition, if the Request Completed bit in the descriptor is not set, the client should continue to discard the data subsequent completions for this tag until it receives a completion descriptor with the Request Completed bit set. On receiving a completion descriptor with the Request Completed bit set, the client can remove all state for the corresponding request. |
| 0010 | Request terminated by a Completion TLP with UR, CA, or CRS status. In this case, there is no data associated with the completion, and the Request Completed bit in the completion descriptor is set. On receiving such a Completion from the Integrated Block, the client can discard the corresponding request. |
| 0011 | Read Request terminated by a Completion TLP with incorrect byte count. This condition occurs when a Completion TLP is received with a byte count not matching the expected count. The Request Completed bit in the completion descriptor is set. On receiving such a completion from the Integrated Block, the client can discard the corresponding request. |
| 0100 | This code indicates the case when the current Completion being delivered has the same tag of an outstanding request, but its Requester ID, TC, or Attr fields did not match with the parameters of the outstanding request. The client should discard any data that follows the descriptor. In addition, if the Request Completed bit in the descriptor is not set, the client should continue to discard the data subsequent completions for this tag until it receives a completion descriptor with the Request Completed bit set. On receiving a completion descriptor with the Request Completed bit set, the client can remove all state associated with the request. |

| Error Code | Description |
|------------|---|
| 0101 | Error in starting address. The low address bits in the Completion TLP header did not match with the starting address of the next expected byte for the request. The client should discard any data that follows the descriptor. In addition, if the Request Completed bit in the descriptor is not set, the client should continue to discard the data subsequent Completions for this tag until it receives a completion descriptor with the Request Completed bit set. On receiving a completion descriptor with the Request Completed bit set, the client can discard the corresponding request. |
| 0110 | Invalid tag. This error code indicates that the tag in the Completion TLP did not match with the tags of any outstanding request. The client should discard any data following the descriptor. |
| 0111 | Invalid byte count. The byte count in the Completion was higher than the total number of bytes expected for the request. In this case, the Request Completed bit in the completion descriptor is also set. On receiving such a completion from the Integrated Block, the client can discard the corresponding request. |
| 1001 | Request terminated by a Completion timeout. This error code is used when an outstanding request times out without receiving a Completion from the link. The Integrated Block maintains a completion timer for each outstanding request, and responds to a completion timeout by transmitting a dummy completion descriptor on the requester completion interface to the client, so that the client can terminate the pending request, or retry the request. Because this descriptor does not correspond to a Completion TLP received from the link, only the Request Completed bit (bit 30), the tag field (bits [71: 64]) and the requester Function field (bits [55: 48]) are valid in this descriptor. |
| 1000 | Request terminated by a Function-Level Reset (FLR) targeting the Function that generated the request. In this case, the Integrated Block transmits a dummy completion descriptor on the requester completion interface to the client, so that the client can terminate the pending request. Because this descriptor does not correspond to a Completion TLP received from the link, only the Request Completed bit (bit 30), the tag field (bits [71:64]) and the requester Function field (bits [55:48]) are valid in this descriptor. |

Table 3-11: Encoding of Error Codes (Cont'd)

When the tags are managed internally by the Integrated Block, logic within the Integrated Block ensures that a tag allocated to a pending request is not re-used until either all the Completions for the request were received or the request was timed out.

When tags are managed by the client, however, the client must ensure that a tag assigned to a request is not re-used until the Integrated Block has signaled the termination of the request by setting the Request Completed bit in the completion descriptor. The client can close out a pending request on receiving a completion with a non-zero error code, but should not free the associated tag if the Request Completed bit in the completion descriptor is not set. Such a situation might occur when a request receives multiple split completions, one of which has an error. In this case, the Integrated Block can continue to receive Completions are incorrectly matched to a different request if its tag is re-assigned too soon. In some cases, the Integrated Block might have to wait for the request to time out even when a split completion is received with an error, before it can allow the tag to be re-used.

Power Management

The Gen3 Integrated Block core supports these power management modes:

- Active State Power Management (ASPM)
- Programmed Power Management (PPM)

Implementing these power management functions as part of the PCI Express design enables the PCI Express hierarchy to seamlessly exchange power-management messages to save system power. All power management message identification functions are implemented. The subsections in this section describe the user logic definition to support the above modes of power management.

For additional information on ASPM and PPM implementation, see the PCI Express Base Specification.

Active State Power Management

The Gen3 Integrated Block for PCIe advertises an N_FTS value of 255 to ensure proper alignment when exiting L0s. If the N_FTS value is modified, you must ensure enough FTS sequences are received to properly align and avoid transition into the Recovery state.

The Active State Power Management (ASPM) functionality is autonomous and transparent from a user-logic function perspective. The core supports the conditions required for ASPM. The integrated block supports ASPM L0s and not ASPM L1.

Note: This is not supported in non-synchronous clocking mode.

Programmed Power Management

To achieve considerable power savings on the PCI Express hierarchy tree, the core supports these link states of Programmed Power Management (PPM):

- L0: Active State (data exchange state)
- L1: Higher Latency, lower power standby state
- L3: Link Off State

The Programmed Power Management Protocol is initiated by the Downstream Component/Upstream Port.

PPM L0 State

The L0 state represents *normal* operation and is transparent to the user logic. The core reaches the L0 (active state) after a successful initialization and training of the PCI Express Link(s) as per the protocol.

PPM L1 State

These steps outline the transition of the core to the PPM L1 state:

- 1. The transition to a lower power PPM L1 state is always initiated by an upstream device, by programming the PCI Express device power state to D3-hot (or to D1 or D2, if they are supported).
- 2. The device power state is communicated to the user logic through the cfg_function_power_state output.
- 3. The core then throttles/stalls the user logic from initiating any new transactions on the user interface by deasserting s_axis_rq_tready. Any pending transactions on the user interface are, however, accepted fully and can be completed later.

There are two exceptions to this rule:

- The core is configured as an Endpoint and the User Configuration Space is enabled. In this situation, the user must refrain from sending new Request TLPs if cfg_function_power_state indicates non-D0, but the user can return Completions to Configuration transactions targeting User Configuration space.
- The core is configured as a Root Port. To be compliant in this situation, the user should refrain from sending new Requests if cfg_function_power_state indicates non-D0.
- 4. The core exchanges appropriate power management DLLPs with its link partner to successfully transition the link to a lower power PPM L1 state. This action is transparent to the user logic.
- 5. All user transactions are stalled for the duration of time when the device power state is non-D0, with the exceptions indicated in step 3.

PPM L3 State

These steps outline the transition of the Endpoint for PCI Express to the PPM L3 state:

- 1. The core negotiates a transition to the L23 Ready Link State upon receiving a PME_Turn_Off message from the upstream link partner.
- Upon receiving a PME_Turn_Off message, the core initiates a handshake with the user logic through cfg_power_state_change_interrupt (see Table 3-12) and expects a cfg_power_state_change_ack back from the user logic.

- 3. A successful handshake results in a transmission of the Power Management Turn-off Acknowledge (PME-turnoff_ack) Message by the core to its upstream link partner.
- 4. The core closes all its interfaces, disables the Physical/Data-Link/Transaction layers and is ready for *removal* of power to the core.

There are two exceptions to this rule:

- The core is configured as an Endpoint and the User Configuration Space is enabled. In this situation, the user must refrain from sending new Request TLPs if cfg_function_power_state indicates non-D0, but the user can return Completions to Configuration transactions targeting User Configuration space.
- The core is configured as a Root Port. To be compliant in this situation, the user should refrain from sending new Requests if cfg_function_power_state indicates non-D0.

Table 3-12: Power Management Handshaking Signals

| Port Name | Direction | Description |
|----------------------------------|-----------|---|
| cfg_power_state_change_interrupt | Output | Asserted if a power-down request TLP is received from the upstream device. After assertion, cfg_power_state_change_interrupt remains asserted until the user asserts cfg_power_state_change_ack. |
| cfg_power_state_change_ack | Input | Asserted by the User Application when it is safe to power down. |

Power-down negotiation follows these steps:

- 1. Before power and clock are turned off, the Root Complex or the Hot-Plug controller in a downstream switch issues a PME_Turn_Off broadcast message.
- 2. When the core receives this TLP, it asserts cfg_power_state_change_interrupt to the User Application and starts polling the cfg_power_state_change_ack input.
- When the User Application detects the assertion of cfg_to_turnoff, it must complete any packet in progress and stop generating any new packets. After the User Application is ready to be turned off, it asserts cfg_power_state_change_ack to the core. After assertion of cfg_power_state_change_ack, the User Application is committed to being turned off.
- The core sends a PME_TO_Ack when it detects assertion of cfg_power_state_change_ack.

Link Training: 2-Lane, 4-Lane, and 8-Lane Components

The 2-lane, 4-lane, and 8-lane Integrated Block for PCI Express can operate at less than the maximum lane width as required by the *PCI Express Base Specification*. Two cases cause core to operate at less than its specified maximum lane width, as defined in these subsections.

Link Partner Supports Fewer Lanes

When the 2-lane core is connected to a device that implements only 1 lane, the 2-lane core trains and operates as a 1-lane device using lane 0.

When the 4-lane core is connected to a device that implements 1 lane, the 4-lane core trains and operates as a 1-lane device using lane 0, as shown in Figure 3-60. Similarly, if the 4-lane core is connected to a 2-lane device, the core trains and operates as a 2-lane device using lanes 0 and 1.

When the 8-lane core is connected to a device that only implements 4 lanes, it trains and operates as a 4-lane device using lanes 0-3. Additionally, if the connected device only implements 1 or 2 lanes, the 8-lane core trains and operates as a 1- or 2-lane device.

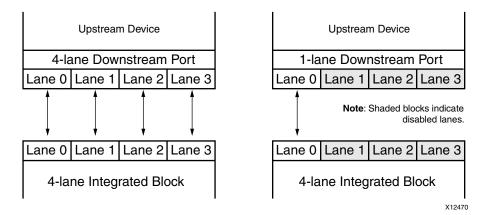


Figure 3-60: Scaling of 4-Lane Endpoint Block from 4-Lane to 1-Lane Operation

Lane Becomes Faulty

If a link becomes faulty after training to the maximum lane width supported by the core and the link partner device, the core attempts to recover and train to a lower lane width, if available. If lane 0 becomes faulty, the link is irrecoverably lost. If any or all of lanes 1–7 become faulty, the link goes into *recovery* and attempts to recover the largest viable link with whichever lanes are still operational.

For example, when using the 8-lane core, loss of lane 1 yields a recovery to 1-lane operation on lane 0, whereas the loss of lane 6 yields a recovery to 4-lane operation on lanes 0-3.

After recovery occurs, if the failed lane(s) becomes *alive* again, the core does not attempt to recover to a wider link width. The only way a wider link width can occur is if the link actually goes down and it attempts to retrain from scratch.

The user_clk clock output is a fixed frequency configured in the CORE Generator tool GUI. user_clk does not shift frequencies in case of link recovery or training down.

Lane Reversal

The Integrated Block supports limited lane reversal capabilities and therefore provides flexibility in the design of the board for the link partner. The link partner can choose to lay out the board with reversed lane numbers and the Integrated Block continues to link train successfully and operate normally. The configurations that have lane reversal support are x8 and x4 (excluding downshift modes). Downshift refers to the link width negotiation process that occurs when link partners have different lane width capabilities advertised. As a result of lane width negotiation, the link partners negotiate down to the smaller of the two advertised lane widths. Table 3-13 describes the several possible combinations including downshift modes and availability of lane reversal support.

| Integrated Block Negotiated Advertised Lane | | Lane Numbe (Endpoint Li | Lane Reversal | |
|--|-------|----------------------------|------------------|-------------------|
| Lane Width | Width | Endpoint | Link Partner | Supported |
| x8 | x8 | Lane 0 Lane 7 | Lane 7 Lane 0 | Yes |
| x8 | x4 | Lane 0 Lane 3 | Lane 7 Lane 4 | No ⁽¹⁾ |
| x8 | x2 | Lane 0 Lane 3 | Lane 7 Lane 6 | No ⁽¹⁾ |
| x4 | x4 | Lane 0 Lane 3 | Lane 3 Lane 0 | Yes |
| x4 | x2 | Lane 0 Lane 1 | Lane 3 Lane 2 | No ⁽¹⁾ |
| x2 | x2 | Lane 0 Lane 1 | Lane 1 Lane 0 | Yes |
| x2 | x1 | Lane 0 Lane 1 | Lane 1 | No ⁽¹⁾ |

Table 3-13: Lane Reversal Support

Notes:

1. When the lanes are reversed in the board layout and a downshift adapter card is inserted between the Endpoint and link partner, Lane 0 of the link partner remains unconnected (as shown by the lane mapping in this table) and therefore does not link train.



SECTION II: VIVADO DESIGN SUITE

Customizing and Generating the Core Constraining the Core Detailed Example Design



Customizing and Generating the Core

This chapter includes information on using the Vivado[™] IP Catalog to customize and generate the core.

Graphical User Interface (GUI)

The LogiCORE[™] IP Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® core is a fully configurable and highly customizable solution. The Gen3 Integrated Block for PCIe is customized using the Vivado IP Catalog.

Note: The screen captures in this chapter are conceptual representatives of their subjects and provide general information only. For the latest information, see the Vivado Design Suite.

Customizing the Core using the Vivado IP Catalog

The Vivado Design Suite IP Catalog for the Gen3 Integrated Block for PCIe consists two modes: Basic Mode and Advanced Mode. To select a mode, use the Mode drop-down list on the first page of the Customize IP dialog box.

Basic Mode

The Basic mode parameters are in the following pages:

- Basic
- Capabilities
- Identity Settings (PF0 IDs and PF1 IDs)
- Base Address Registers (PF0 and PF1)
- Legacy/MSI Capabilities

Basic

Basic Parameter Settings

The initial customization screen shown in Figure 4-1 is used to define the basic parameters for the core, including the component name, reference clock frequency, and silicon type.

| | Customize IP |
|---|--------------------------------------|
| Customize Virtex-7 FPGA Gen3 Integrated Block for PCI Express (1.3) by specifying IP Options. | |
| IP Options | |
| Virtex-7 FPGA Gen3 Integrated Block for PC | Ci Express |
| Component Name pcie3_7x_v1_3_0 | |
| Basic Capabilities PF0 IDs PF1 IDs PF0 BAR PF1 BAR | Legacy/MSI Cap |
| Mode Basic V | |
| Device / Port Type PCI Express Endpoint device | Generate Additional PCIE Constraints |
| PCle Block Location X0Y1 | PCIe Block Location : 🗹 X0Y1 🗌 X0Y2 |
| Number of Lanes | Maximum Link Speed |
| Lane Width XI 💌 | © 2.5 GT/s ○ 5.0 GT/s ○ 8.0 GT/s |
| AXI-ST Interface Width | AXI-ST Interface Frequency |
| AXI-ST Interface Width 64 bit 💌 | 62.5MHz |
| AXI-ST Alignment Mode | Enable AXISTd.EN Frame Straddle |
| Data Aligned Address Aligned | ☑ Enable Client Tag |
| Reference Clock Frequency (MHz) 100 MHz | |
| Xilinx Development Board None 💌 | |
| Silicon Revision | |
| Enable Pipe Simulation | |
| | |
| | |
| | |
| | |
| Show Advanced Options | |
| | OK Cancel |

Figure 4-1: Integrated Block for PCI Express Parameters

Component Name

Base name of the output files generated for the core. The name must begin with a letter and can be composed of these characters: a to z, 0 to 9, and "_."

Mode

Allows you to select the Basic or Advanced mode of the configuration of core.

PCIe Device / Port Type

Indicates the PCI Express logical device type.

PCIe Block Location

Selects from the available Integrated Blocks to enable generation of location-specific constraint files and pinouts. This selection is used in the default example design scripts.

This option is not available if a Xilinx Development Board is selected.

Generate Additional PCIe Constraints

Allows you to generate additional constraints files for other blocks available in the device.

This option is not available if a Xilinx Development Board is selected.

Number of Lanes

The Gen3 Integrated Block for PCIe requires the selection of the initial lane width. Figure 4-1 defines the available widths and associated generated core. Wider lane width cores can train down to smaller lane widths if attached to a smaller lane-width device. See Link Training: 2-Lane, 4-Lane, and 8-Lane Components, page 157 for more information.

| Lane Width | Product Generated |
|------------|--|
| x1 | 1-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| x2 | 2-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| x4 | 4-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| x8 | 8-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |

Table 4-1: Lane Width and Product Generated

Maximum Link Speed

The Gen3 Integrated Block for PCIe allows you to select the Maximum Link Speed supported by the device. Table 4-2 defines the lane widths and link speeds supported by the device. Higher link speed cores are capable of training to a lower link speed if connected to a lower link speed capable device.

Table 4-2: Lane Width and Link Speed

| Lane Width | Link Speed |
|------------|--------------------------|
| x1 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |
| x2 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |
| x4 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |
| x8 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |

AXI-ST Interface Width

The Gen3 Integrated Block for PCIe allows you to select the Interface Width, as defined in Table 4-3. The default interface width set in the Customize IP dialog box is the lowest possible interface width.

| Lane Width | Link Speed (Gb/s) | Interface Width (Bits) |
|------------|----------------------|------------------------|
| x1 | 2.5, 5.0, 8.0 | 64 |
| x2 | 2.5, 5.0 | 64 |
| x2 | 8.0 | 64, 128 |
| x4 | 2.5 | 64 |
| x4 | 5.0 | 64, 128 |
| x4 | 8.0 | 128, 256 |
| x8 | 2.5 | 64, 128 |
| x8 | 5.0 | 128 256 |
| x8 | 8.0 | 256 |

Table 4-3: Lane Width, Link Speed, and Interface Width

AXI-ST Interface Frequency

The frequency is set to 62.5Mhz.

AXI-ST Alignment Mode

When a payload is present, there are two options for aligning the first byte of the payload with respect to the datapath. See Data Alignment Options, page 81.

Requestor Completion Straddle

The Gen3 Integrated Block for PCIe provides an option to straddle packets on the Requestor Completion interface when the interface width is 256 bits. See Straddle Option for 256-Bit Interface, page 149.

Enable Client Tag

Enables you to use the client Tag.

Reference Clock Frequency

Selects the frequency of the reference clock provided on sys_clk. For important information about clocking the Gen3 Integrated Block for PCIe, see Clocking, page 64.

Xilinx Development Board

Selects the Xilinx Development Board to enable the generation of Xilinx Development Board-specific constraints files.

Silicon Type

Selects the silicon type.

Enable Pipe Simulation

When this box is checked, generates a core that can be simulated with pipe interfaces connected.

Capabilities

The Capabilities settings page is shown in Figure 4-2.

| ons | |
|---|--|
| ex-7 FPGA Gen3 Integrated Block for PCI Express | |
| ponent Name pcie3_7x_v1_3_0 | |
| Basic Capabilities PF0 IDs PF0 BAR Legacy/MSI Cap | |
| Physical Functions | |
| Enable Physical Function 0 Enable Physical Function 1 | |
| | |
| Device Capabilities Register PF 🗧 🛠 | Link Status Register |
| PF0 Max Payload Size 512 bytes T PF1 Max Payload Size 512 bytes T | Selects whether the device reference clock is provided by the connector (Synchronous) or generated via an onboard PLL(Asynchronous) |
| Extended Tag Field | Enable Slot Clock Configuration |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |



Enable Physical Function 0 and 1

The Gen3 Integrated Block for PCIe implements an additional Physical Function (PF).

The Integrated Block implements up to six Virtual Functions that are associated to either PF0 or PF1 (if enabled).

MPS

This field indicates the maximum payload size that the device or function can support for TLPs. This is the value advertised to the system in the Device Capabilities Register.

Extended Tag

This field indicates the maximum supported size of the Tag field as a Requester. The options are:

- When selected, 8-bit Tag field support
- When deselected, 5-bit Tag field support

Slot Clock Configuration

Enables the Slot Clock Configuration bit in the Link Status register. When you select this option, the link is synchronously clocked. For more information on clocking options, see Clocking, page 74.

Identity Settings (PF0 IDs and PF1 IDs)

The Identity Settings pages are shown in Figure 4-3 and Figure 4-4. These settings customize the IP initial values, class code, and Cardbus CIS pointer information. The page for Physical Function 1 (PF1) is only displayed when PF1 is enabled.

| | Customize I | P | |
|---|---|--------------------|----|
| Customize Virtex-7 FPGA Ge Integrated Block for PCI Exp by specifying IP Options. | 13 ress (1.3) | | |
| tions Tex-7 FPGA Gen3 | Integrated Block for PCI Express | | |
| mponent Name pcie3 7x v | | | |
| | DIDs PF0 BAR Legacy/MSI Cap | | |
| PF0 - ID Initial Values | | | \$ |
| Vendor ID 10 | E Range: 0000FFFF | | |
| Device ID 70 | 1 Range: 0000FFFF | | |
| Revision ID 00 | Range: 00FF | | |
| Subsystem Vendor ID 10 | | | |
| Subsystem ID 00 | 07 Range: 0000FFFF | | |
| | | | _ |
| Class Code | | | * |
| PF0 Use Class Code Lo | okup Assistant | | |
| Base Class Value | Simple communication controllers | | |
| Base Class | 05 | Range: 00FF | |
| Sub-Class/Interface Value | Generic XT compatible serial controller | | |
| Sub-Class | 80 | Range: 00FF | |
| Interface | 00 | Range: 00FF | |
| Class Code | 058000 | Range: 000000FFFFF | |
| | | | |
| | | | |
| | | | |
| | | | |
| Advanced Options | | | |
| | | | |

Figure 4-3: Identity Settings (PF0)

| tegrated Block for PCI Exp y specifying IP Options. | press (1.3) |
|--|--|
| | B Integrated Block for PCI Express |
| ponent Name pcie3_7x_ | v1_3_0 |
| Basic Capabilities PF(| 0 IDs PF1 IDs PF0 BAR PF1 BAR Legacy/MSI Cap |
| PF1 - ID Initial Values | |
| Vendor ID 10 |)EE |
| Device ID 70 | 011 Range: 0000FFFF |
| Revision ID 00 | 0 Range: 00FF |
| Subsystem Vendor ID 10 | NEE |
| Subsystem ID 00 | 007 Range: 0000FFFF |
| Subsystem iD | Normange: 0000PPP |
| Class Code | |
| PF1 Use Class Code Lo | ookup Assistant |
| Base Class Value | Simple communication controllers |
| Base Class | 05 Range: 00FF |
| Sub-Class/Interface Value | |
| Sub-Class | 80 Range: 00FF |
| nterface | 00 Range: 00FF |
| Class Code | 058000 Range: 00000FFFFF |
| | lininger sussessment in the |
| | |
| | |

Figure 4-4: Identity Settings (PF1)

PF0 ID Initial Values

- **Vendor ID:** Identifies the manufacturer of the device or application. Valid identifiers are assigned by the PCI Special Interest Group to guarantee that each identifier is unique. The default value, 10EEh, is the Vendor ID for Xilinx. Enter a vendor identification number here. FFFFh is reserved.
- **Device ID:** A unique identifier for the application; the default value, which depends on the configuration selected, is 70*<link speed><link width>*h. This field can be any value; change this value for the application.
- **Revision ID:** Indicates the revision of the device or application; an extension of the Device ID. The default value is 00h; enter values appropriate for the application.
- **Subsystem Vendor ID:** Further qualifies the manufacturer of the device or application. Enter a Subsystem Vendor ID here; the default value is 10EEh. Typically, this value is the same as Vendor ID. Setting the value to 0000h can cause compliance testing issues.
- **Subsystem ID:** Further qualifies the manufacturer of the device or application. This value is typically the same as the Device ID; the default value depends on the lane width and link speed selected. Setting the value to 0000h can cause compliance testing issues.

Class Code

The Class Code identifies the general function of a device, and is divided into three byte-size fields:

- **Base Class:** Broadly identifies the type of function performed by the device.
- **Sub-Class:** More specifically identifies the device function.
- **Interface:** Defines a specific register-level programming interface, if any, allowing device-independent software to interface with the device.

Class code encoding can be found at <u>www.pcisig.com</u>.

Class Code Look-up Assistant

The Class Code Look-up Assistant provides the Base Class, Sub-Class and Interface values for a selected general function of a device. This Look-up Assistant tool only displays the three values for a selected function. The user must enter the values in Class Code for these values to be translated into device settings.

Base Address Registers (PF0 and PF1)

The Base Address Registers (BARs) screens shown in Figure 4-5 and Figure 4-6 set the base address register space for the Endpoint configuration. Each BAR (0 through 5) configures the BAR Aperture Size and Control attributes of the Physical Function as described in Table D-1.

| | Customize IP | |
|--|---|----------|
| Customize Virtex-7 FPGA Gen3 Integrated Block for PCI Express (1.3) by specifying IP Options. | | |
| ^{ptions} rtex-7 FPGA Gen3 Integrated Block for PCI Exp | ress | |
| omponent Name pcie3_7x_v1_3_0 | | |
| Basic Capabilities PFC IDs PFO BAR Legacy/MSI Cap | | |
| Base Address Registers (BARs) serve two surposes. Initially, they serve as map. After the BIOS or OS determines what addresses to assign to the devite information to perform address decoding. BarO Type Memory G4 bit Prefetchable Size Unit Kilobytes Size Value 2 Value (Hex): FFFFF800 | a mechanism for the device to request blocks of address space in the system memory vice, the Base Address Registers are programmed with addresses and the device uses Barl Type IVA IVA Prefetchable Size Unit Klobytes Size Value Z Value (Hex): 000000000 | • |
| Bar2 Type N/A Image: Size Unit Kilobytes Value (Hex): 00000000 | Bar3 Type N/A Size Unit Kilobytes Value (Hex): 00000000 | 1.010 |
| Bar4 Type N/A Image: Size Unit Kilobytes Value (Hex): 50000000 | BarS Type N/A Size Unit Kilobytes Size Value Value (Hex): 00000000 | |
| Expansion Rom Size Unit Kilobytes Y Size Value 2 Y Value (Hexi: 00000000 | ,, | |
| Advanced Options | | OK Cance |

Figure 4-5: Base Address Register (PF0)

| | | Customize IP | × |
|------|--|--|--------|
| 1 | Customize Virtex-7 FPGA Gen3 Integrated Block for PCI Express (1.3) by specifying IP Options. | | |
| | tions | | |
| VI | tex-7 FPGA Gen3 Integrated Block for PCI Express | | |
| C | mponent Name pcie3_7x_v1_3_0 | | |
| | Basic Capabilities PF0 IDs PF1 IDs PF0 BAR PF1 BAR Legacy/MSI Cap |] | |
| | Base Address Registers (BARs) serve two purposes. Initially, they serve as a mech map. After the BIOS or 05 determines what addresses to assign to the device, the this information to perform address decoding. | anism for the device to request blocks of address space in the system memory e Base Address Registers are programmed with addresses and the device uses Bar1 Type N/A 		 64 bit Prefetchable Size Unit Klobytes 		 Size Value 2 Type N/A 		 64 bit Prefetchable Size Unit Klobytes 		 Size Value 2 | |
| | Value (Hex): 00000000 | Value (Hex): 00000000 | • |
| | Bar4 | Bar5 | |
| | Type N/A 💌 🛛 64 bit 🗌 Prefetchable | Type N/A Prefetchable | |
| | Size Unit Kilobytes Vize Value 2 | Size Unit: Kilobytes T Size Value 2 T | |
| | Value (Hex): 00000000 | Value (Hex): 00000000 | |
| | Expansion Rom Size Unit Kilobytes Visize Value 2 V Value (Hex): 00000000 | | • |
| Show | Advanced Options | ОК | Cancel |

Figure 4-6: Base Address Register (PF1)

Base Address Register Overview

The Gen3 Integrated Block for PCIe in Endpoint configuration supports up to six 32-bit BARs or three 64-bit BARs, and the Expansion read-only memory (ROM) BAR. The Gen3 Integrated Block for PCIe in Root Port configuration supports up to two 32-bit BARs or one 64-bit BAR, and the Expansion ROM BAR.

BARs can be one of two sizes:

- **32-bit BARs:** The address space can be as small as 16 bytes or as large as 2 gigabytes. Used for Memory to I/O.
- **64-bit BARs:** The address space can be as small as 128 bytes or as large as 8 exabytes. Used for Memory only.

All BAR registers share these options:

- **Checkbox:** Click the checkbox to enable the BAR; deselect the checkbox to disable the BAR.
- **Type:** BARs can either be I/O or Memory.
 - *I/O*: I/O BARs can only be 32-bit; the Prefetchable option does not apply to I/O BARs. I/O BARs are only enabled for the Legacy PCI Express Endpoint core.

- *Memory*: Memory BARs can be either 64-bit or 32-bit and can be prefetchable.
 When a BAR is set as 64 bits, it uses the next BAR for the extended address space and makes the next BAR inaccessible to the user.
- **Size:** The available Size range depends on the PCIe Device/Port Type and the Type of BAR selected. Table 4-4 lists the available BAR size ranges.

| Table 4-4: | BAR Size | Ranges for | Device | Configuration |
|------------|----------|------------|--------|---------------|
|------------|----------|------------|--------|---------------|

| PCIe Device / Port Type | BAR Type | BAR Size Range |
|-----------------------------|---------------|-------------------------|
| DCI Everage Endpoint | 32-bit Memory | 128 Bytes – 2 Gigabytes |
| PCI Express Endpoint | 64-bit Memory | 128 Bytes – 8 Exabytes |
| | 32-bit Memory | 16 Bytes – 2 Gigabytes |
| Legacy PCI Express Endpoint | 64-bit Memory | 16 Bytes – 8 Exabytes |
| | I/O | 16 Bytes – 2 Gigabytes |

- **Prefetchable:** Identifies the ability of the memory space to be prefetched.
- Value: The value assigned to the BAR based on the current selections.

For more information about managing the Base Address Register settings, see Managing Base Address Register Settings.

Expansion ROM Base Address Register

If selected, the Expansion ROM is activated and can be a value from 2 KB to 4 GB. According to the *PCI 3.0 Local Bus Specification*, the maximum size for the Expansion ROM BAR should be no larger than 16 MB. Selecting an address space larger than 16 MB can result in a non-compliant core.

Managing Base Address Register Settings

Memory, I/O, Type, and Prefetchable settings are handled by setting the appropriate GUI settings for the desired base address register.

Memory or I/O settings indicate whether the address space is defined as memory or I/O. The base address register only responds to commands that access the specified address space. Generally, memory spaces less than 4 KB in size should be avoided. The minimum I/O space allowed is 16 bytes; use of I/O space should be avoided in all new designs.

Prefetchability is the ability of memory space to be prefetched. A memory space is prefetchable if there are no side effects on reads (that is, data is not destroyed by reading, as from a RAM). Byte write operations can be merged into a single double word write, when applicable.

When configuring the core as an Endpoint for PCIe (non-Legacy), 64-bit addressing must be supported for all BARs (except BAR5) that have the prefetchable bit set. 32-bit addressing is permitted for all BARs that do not have the prefetchable bit set. The prefetchable

bit-related requirement does not apply to a Legacy Endpoint. The minimum memory address range supported by a BAR is 128 bytes for a PCI Express Endpoint and 16 bytes for a Legacy PCI Express Endpoint.

Disabling Unused Resources

For best results, disable unused base address registers to conserve system resources. A base address register is disabled by deselecting unused BARs in the GUI.

Legacy/MSI Capabilities

On this page, you set the Legacy Interrupt Settings and MSI Capabilities for all applicable Physical and Virtual Functions.

| ntegrated Block for PCI Express (1.3) ay specifying IP Options. tex-7 FPGA Gen3 Integrated Block for PCI Express nponent Name [pcie3_7v_v1_3_0 Basic Capabilities [PF0 IDs [PF0 BAR] Legacy/MSI Cap Legacy Interrupt Settings PF0 Interrupt PIN NONE | mponent Name pcie3_7x_v1_3_0 Basic Capabilities PF0 IDs PF0 BAB Lega | | |
|--|---|---|-------|
| Integrated Block for PCI Express Inponent Name Integrated Block for PCI Express Inponent Name Integration Integration | ions tex-7 FPGA Gen3 Integrated Bl mponent Name [pcle3_7%_v1_3_0 Basic Capabilities PF0 IDS PF0 BAP Lega | | |
| tex-7 FPGA Gen3 Integrated Block for PCI Express nponent Name [pcie3_7k_v1_3_0 Basic Capabilities PF0 IDs PF0 BAB Legacy/MSI Cap Legacy Interrupt Settings PF0 Interrupt PIN NONE MSI Capabilities | tex-7 FPGA Gen3 Integrated Bl mponent Name pcle3_7x_v1_3_0 Basic Capabilities PF0 IDs PF0 BAB Lega | | |
| nponent Name [pcie3_7k_V1_3_0 Basic Capabilities [PF0 IDs] PF0 BAB [Legacy/MSI Cap Legacy Interrupt Settings PF0 Interrupt PIN NONE MSI Capabilities : | mponent Name pcie3_7x_v1_3_0 Basic Capabilities PF0 IDs PF0 BAB Lega | | |
| Basic Capabilities PF0 IDs PF0 BAR Legacy/MSI Cap Legacy Interrupt Settings PF0 Interrupt PIN NONE | Basic Capabilities PF0 IDs PF0 BAB Lega | | |
| Legacy Interrupt Settings 22 PF0 Interrupt PIN NONE MSI Capabilities 22 | | | |
| PF0 Interrupt PIN NONE MSI Capabilities | | cy/MSI Cap | |
| PF0 Interrupt PIN NONE MSI Capabilities | Legacy Interrupt Settings | | * |
| MSI Capabilities : | | | |
| MSI Capabilities : | | | |
| ✓ PF0 Enable MSI Capability Structure PF0 Multiple Message Capable 1 vector | MSI Capabilities | | * |
| | PF0 Enable MSI Capability Structure | PF0 Multiple Message Capable 🛛 vector 🔄 | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | advanced Outlines | | |
| | avanceu options | | |
| dvanced Options | | | OK Ca |

Figure 4-7: Legacy/MSI Capabilities

Legacy Interrupt Settings

- **Enable INTX**: Enables the ability of the PCI Express function to generate INTx interrupts.
- **Interrupt PIN**: Indicates the mapping for Legacy Interrupt messages. A setting of "None" indicates no Legacy Interrupts are used.

MSI Capabilities

• Enable MSI Capability Structure: Indicates that the MSI Capability structure exists.

Note: Although it is possible not to enable MSI or MSI-X, the result would be a non-compliant core. The *PCI Express Base Specification* requires that MSI, MSI-X, or both be enabled.

- **64 bit Address Capable**: Indicates that the function can send a 64-bit Message Address.
- **Multiple Message Capable**: Selects the number of MSI vectors to request from the Root Complex.
- **Per Vector Masking Capable**: Indicates that the function supports MSI per-vector Masking.

Advanced Mode

In Advanced Mode, the GUI consists of the following pages:

- Basic
- Capabilities
- PF0 ID and PF1 ID
- PF0 BAR and PF1 BAR
- SRIOV Config (PF0 and PF1)
- PF0 SRIOV BARs and PF1 SRIVO BARs
- Legacy/MSI Capabilities
- MSI-X Capabilities
- Power Management
- Extended Capabilities 1 and Extended Capabilities 2

Basic

The Basic setting page is the same for both Basic or Advanced modes. See Basic, page 161.

Capabilities

The Capabilities settings for Advanced mode contains three additional parameters those for Basic mode. For a description of the basic mode settings, see Capabilities, page 164. The Advanced mode settings are described below.

| ex-7 FPGA Gen3 Integrated Block for PCI Express | |
|---|---|
| • | SRIOV BARs PF1 SRIOV BARs Legacy/MSI Cap MSIx Cap Power Management Extd. Capabilitir 4 |
| Physical Functions Physical Function 0 Finable Physical Function 0 Finable Physical Function 1 | |
| Device Capabilities Register PF PF0 Max Payload Size S12 bytes PF1 Max Payload Size S12 bytes Cextended Tag Field | Link Status Register Selects whether the device reference clock is provided by the connector (Synchronous) or generated via an onboard PLL(Asynchronous) Enable Slot Clock Configuration |
| ∃ SRIOV Capability | Function Level Reset |
| Device Capabilities Register 2 32-bit AtomicOp Completer Supported 44-bit AtomicOp Completer Supported 128-bit CAS Completer Supported DTPH Completer Supported 0BFF Supported 00 Not Supported | |
| | |

Figure 4-8: Capabilities Settings (Advanced Mode)

SRIOV Capabilities

Enables Single Root Port I/O Virtualization (SRIOV) Capabilities. The Integrated Block implements the Single Root Port I/O Virtualization PCIe extended capability. When this capability is enabled, the SRIOV capability is implemented for both PF0 and PF1 (if selected).

Function Level Reset

Indicates the Function Level Reset is enabled. The Integrated Block enables you to reset a specific device function. This mechanism is only applicable to Endpoint configurations.

Device Capabilities Registers 2

Device Capability Register 2 Settings: Specifies options for AtomicOps and TPH Completer Support. See the Device Capability Register 2 description in Chapter 7 of the PCI Express Base Specification for more information. These settings apply to both Physical Functions, if PF1 is enabled.

PF0 ID and PF1 ID

The Identity settings (PF0 and PF1 Initial ID) are the same for both Basic and Advanced modes. See Identity Settings (PF0 IDs and PF1 IDs), page 165.

PFO BAR and PF1 BAR

The PFO and PF1 BAR settings are the same for both Basic and Advanced modes. See Base Address Registers (PFO and PF1), page 168.

SRIOV Config (PF0 and PF1)

The SRIOV Config page is shows in Figure 4-9.

| ustomize Virtex-7 FP0 | SA Gen3 | | | ustomize IP | | | |
|--|--------------------|-----------------|-----------------------|----------------|---------------------|-------------------------|---------------------------|
| ntegrated Block for P by specifying IP Option | CI Express (1.3) | | | | | | |
| ons | | | | | | | |
| ex-7 FPGA G | en3 Integrated | Block for PC | I Express | | | | |
| nponent Name pcies | 3_7x_v1_3_0 | | | | | | |
| Basic Capabilities | PF0 IDs PF1 IDs PF | O BAR PF1 BAR S | RIOV Config PF0 SRIOV | BARS PF1 SRIOV | BARs Legacy/MSI Cap | MSIx Cap Power Manageme | nt 🗍 Extd. Capabiliti 🛛 🕯 |
| PF0 SRIOV Config | | \$ | PF1 SRIOV Config | | | | |
| Cap Version | 0 Ra | ange: 0F | Cap Version | 0 | Range: 0F | | |
| Number of PF0 VF's | 1 | | Number of PF1 VF's | 1 | | | |
| PF Dependency Link | 0000 Ra | ange: 0000FFFF | PF Dependency Link | 0001 | Range: 0000FFFF | | |
| First VF Offset | | ange: 0000FFFF | PF Dependency Link | 7F | Range: 0000FFFF | | |
| VF Device ID | 0000 Ra | ange: 0000FFFF | VF Device ID | 0000 | Range: 0000FFFF | | |
| Supported Page Size | 4КВ 💌 | | Supported Page Size | 4KB 💌 | | | |
| | | | | | - | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| dvanced Options | | | | | | | |

Figure 4-9: SRIOV Config (PF0 and PF1)

SRIOV Capability Version

Indicates the 4-bit SRIOV Capability Version for the Physical Function.

SRIOV Function Select

Indicates the number of Virtual Functions associated to the Physical Function. A maximum of six Virtual Functions are available to PF0 and PF1.

SRIOV Functional Dependency Link

Indicates the SRIOV Functional Dependency Link for the Physical Function. The programming model for a device can have vendor-specific dependencies between sets of Functions. The Function Dependency Link field is used to describe these dependencies.

SRIOV First VF Offset

Indicates the offset of the first Virtual Function (VF) for the Physical Function (PF). PF0 always resides at Offset 0 while PF1 always resides at Offset 1. There are 6 Virtual Functions available in the Virtex-7 Gen3 Integrated block for PCI Express, and the Virtual Functions reside at the function number range 64 - 69.

Virtual functions are mapped sequentially with VFs for PF0 taking precedence. For example, if PF0 has 2 Virtual Functions and PF1 has 3 Virtual Functions, the following mapping would occur:

| Physical Function | Virtual Function | Function Number Range |
|-------------------|------------------|--------------------------|
| PFO | VF0 | 64 |
| PF0 | VF1 | 65 |
| PF1 | VF0 | 66 |
| PF1 | VF1 | 67 |
| PF1 | VF2 | 68 |

Table 4-5: Example Virtual Function Mappings

The PFx_FIRST_VF_OFFSET is calculated by taking the first offset of the Virtual Function and subtracting that from the offset of the Physical Function.

PFx_FIRST_VF_OFFSET = (PFx first VF offset - PFx offset)

In the example above, the following offsets is used:

PF0_FIRST_VF_OFFSET = (64 - 0) = 64 PF1_FIRST_VF_OFFSET = (66 - 1) = 65

PF0 is always 64 assuming PF0 has 1 or more Virtual Functions. The initial offset for PF1 is a function of how many VF's are attached to PF0 and is defined in pseudo code below:

PF1_FIRST_VF_OFFSET = 63 + NUM_PF0_VFS

SRIOV VF Device ID

Indicates the 16-bit Device ID for all Virtual Functions associated with the Physical Function.

SRIOV Supported Page Size

Indicates the page size supported by the Physical Function. This Physical Function supports a page size of 2n+12, if bit *n* of the 32-bit register is set.

PF0 SRIOV BARs and PF1 SRIVO BARs

The SRIOV Base Address Registers (BARs) screens shown in Table 4-10 and Table 4-11 set the base address register space for the Endpoint configuration. Each BAR (0 through 5) configures the SRIOV BAR Aperture Size and SRIOV Control attributes as described in Table D-1.

| PF0 BAR PF1 BAR SRIOV Config PF0 | SRIOV BARS PF1 SRIOV BARS Legacy/MSI Cap MSIx | Con Rever Management Fuld Conshilition 1 Fuld 4 |
|---|---|--|
| | | . cap Fower Management Extu, capabilities 1 Extu, 4 7 |
| t addresses to assign to the device, the Ba | ism for the device to request blocks of address space ir ase Address Registers are programmed with addresses | n the system memory ; and the device uses |
| | Barl | 7 |
| Prefetchable Typ | e N/A 💌 🗆 64 bit 🗌 Prefetchable | |
| Size | Unit Kilobytes 💌 Size Value 2 | |
| Valu | ue 00000000 | |
| | Bar3 | |
| | | |
| | | |
| Valu | | |
| | Bar5 | - |
| Prefetchable | e N/A T Prefetchable | |
| ▼ Size | unit Kilobytes 👻 Size Value 2 💌 | |
| Valu | ue 00000000 | |
| | | |
| | addresses to assign to the device, the Bang. | addresses to assign to the device, the Base Address Registers are programmed with addresses Prefetchable Prefetchable Prefetchable Prefe |

Figure 4-10: **PFO SRIOV BARs Settings**

| | | Customize IP | |
|---|--|---|-------------------|
| Customize Virtex-7 F Integrated Block for by specifying IP Opti | PCI Express (1.3) | | t |
| ptions rtox-7 EDGA (| en3 Integrated Block for PC | 1 Everess | |
| | | n express | |
| omponent Name pc | | | |
| Basic Capabilitie | s PF0 IDs PF1 IDs PF0 BAR PF1 BAR S | RIOV Config PF0 SRIOV BARS PF1 SRIOV BARS Legacy/MSI Cap MSIx Cap Power Management Extd. | Capabiliti: 4 🕨 🗉 |
| map. After the BIOS | ers (BARs) serve two purposes. Initially, they s or OS determines what addresses to assign to erform address decoding. | serve as a mechanism for the device to request blocks of address space in the system memory the device, the Base Address Registers are programmed with addresses and the device uses | |
| 🛛 🗹 Bar0 | | Barl | |
| Type Memory | 🔄 🗆 64 bit 🗌 Prefetchable | Type N/A 🔽 🗆 64 bit 🗌 Prefetchable | |
| Size Unit Kilobyte | s 💌 Size Value 2 | Size Unit Kilobytes 💌 Size Value 2 | |
| Value FFFFF80 | 2 | Value 00000000 | |
| | | | |
| 🗆 Bar2 | | 🗆 Bar3 | |
| Type N/A | 💌 🗆 64 bit 🗌 Prefetchable | Type N/A 💽 🗆 64 bit 🗌 Prefetchable | |
| Size Unit Kilobyte | s 💌 Size Value 2 | Size Unit Kilobytes 💌 Size Value 2 | |
| Value 000000 | 0 | Value 00000000 | |
| Bar4 | | BarS | |
| Type N/A | 💌 🗆 64 bit 🔲 Prefetchable | Type N/A 💌 🗆 Prefetchable | |
| Size Unit Kilobyte | Size Value 2 | Size Unit Kilobytes 💌 Size Value 2 💌 | |
| Value 000000 | 0 | Value 00000000 | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| Advanced Options | | | |
| | | | OK Can |

Figure 4-11: **PF1 SRIOV BARs Settings**

SRIOV Base Address Register Overview

The Gen3 Integrated Block for PCIe in Endpoint configuration supports up to six 32-bit BARs or three 64-bit BARs. The Gen3 Integrated Block for PCIe in Root Port configuration supports up to two 32-bit BARs or one 64-bit BAR.

SRIOV BARs can be one of two sizes:

- **32-bit BARs**: The address space can be as small as 16 bytes or as large as 2 gigabytes. Used for Memory to I/O.
- **64-bit BARs**: The address space can be as small as 128 bytes or as large as 8 exabytes. Used for Memory only.

All SRIOV BAR registers share these options:

• **Checkbox:** Click the checkbox to enable the BAR; deselect the checkbox to disable the BAR.

- **Type:** SRIOV BARs can either be I/O or Memory.
 - *I/O*: I/O BARs can only be 32-bit; the Prefetchable option does not apply to I/O BARs. I/O BARs are only enabled for the Legacy PCI Express Endpoint core.
 - *Memory*: Memory BARs can be either 64-bit or 32-bit and can be prefetchable.
 When a BAR is set as 64 bits, it uses the next BAR for the extended address space and makes the next BAR inaccessible to the user.
- **Size:** The available Size range depends on the PCIe® Device/Port Type and the Type of BAR selected. Table 4-6 lists the available BAR size ranges.

| PCIe Device / Port Type | BAR Type | BAR Size Range |
|-----------------------------|---------------|-------------------------|
| DCI Express Endpoint | 32-bit Memory | 128 Bytes – 2 Gigabytes |
| PCI Express Endpoint | 64-bit Memory | 128 Bytes – 8 Exabytes |
| | 32-bit Memory | 16 Bytes – 2 Gigabytes |
| Legacy PCI Express Endpoint | 64-bit Memory | 16 Bytes – 8 Exabytes |
| | I/O | 16 Bytes – 2 Gigabytes |

Table 4-6: SRIOV BAR Size Ranges for Device Configuration

- **Prefetchable:** Identifies the ability of the memory space to be prefetched.
- Value: The value assigned to the BAR based on the current selections.

For more information about managing the SRIOV Base Address Register settings, see Managing Base Address Register Settings.

Managing SRIOV Base Address Register Settings

Memory, I/O, Type, and Prefetchable settings are handled by setting the appropriate Customize IP dialog box settings for the desired base address register.

Memory or I/O settings indicate whether the address space is defined as memory or I/O. The base address register only responds to commands that access the specified address space. Generally, memory spaces less than 4 KB in size should be avoided. The minimum I/O space allowed is 16 bytes; use of I/O space should be avoided in all new designs.

Prefetchability is the ability of memory space to be prefetched. A memory space is prefetchable if there are no side effects on reads (that is, data is not destroyed by reading, as from a RAM). Byte write operations can be merged into a single double word write, when applicable.

When configuring the core as an Endpoint for PCIe (non-Legacy), 64-bit addressing must be supported for all SRIOV BARs (except BAR5) that have the prefetchable bit set. 32-bit addressing is permitted for all SRIOV BARs that do not have the prefetchable bit set. The prefetchable bit related requirement does not apply to a Legacy Endpoint. The minimum memory address range supported by a BAR is 128 bytes for a PCI Express Endpoint and 16 bytes for a Legacy PCI Express Endpoint.

Disabling Unused Resources

For best results, disable unused base address registers to conserve system resources. A base address register is disabled by deselecting unused BARs in the GUI.

Legacy/MSI Capabilities

This page is same as that of Basic mode. See Legacy/MSI Capabilities, page 172.

MSI-X Capabilities

The MSI-X Capabilities page is available in Advanced mode only.

| In table MSK: Capability Structure MSk: Table Settings Table Settings 000000000000000001FFFFFFF BAR Indicator MSk: Pending Bit Array (PBA) Settings PBA Offset 0000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Indicator PBA Offset 0000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 00000000001FFFFFF PBA Offset 000000000000000000000000000000000000 | | Customize IP | |
|--|------------------|--|----|
| <pre>rysect/ryse</pre> | Customize Virte | xx-7 FPGA Gen3 kr for PCI Express (1,3) | |
| ex-7 FPGA Gen3 Integrated Block for PCI Express ponent Name [prie3_7x_v1_3_0 Basic [Capabilities PF0 IDs PF0 BAR [Legacy/MSI Cap Power Management Extd. Capabilities.1 Extd. Capabilities.2 F0 Cability Structure MSk Table Settings Cability Structure MSk Table Settings Cability Structure MSk Table Settings Cability Structure MSk Pading Bt Array (PBA) Settings PBA Offset D00000001FFFFFFF PBA BAR Indicator PAF 0 | by specifying IP | Options. | |
| ponent Name prie3_7x_v1_3_0 Basic Capabilities PF0 IDs PF0 BAR Legacy/MSI Cap MSix Cap Power Management Extd. Capabilities-1 Extd. Capabilities-2 F0 Capability Structure MSix Table Settings A Trable Settings A Trable Settings A Trable Setting 0000000. IFFFFFFF Table Offset 00000000 0000000. IFFFFFFF A Table Offset 00000000 0000000. IFFFFFFF A Table Offset 00000000 0000000. IFFFFFFF A Table Offset 0000000 0000000. IFFFFFFF A Table Offset 0000000 0000000. IFFFFFFF A Table Offset 0000000 0000000. IFFFFFFFF A Table Offset 0000000 0000000. IFFFFFFF A Table Offset 0000000 00000000. IFFFFFFF A Table Offset 0000000 00000000. IFFFFFFF A Table Offset 0000000 00000000. IFFFFFFF A Table Offset 0000000 00000000000 00000000. IFFFFFFF A Table Offset 0000000 00000000. IFFFFFFF A Table Offset 0000000 0000000000000000000000000000 | | A Gen3 Integrated Block for PCI Express | |
| Bask Capabilities PF0 IDs PF0 BAR Legacy/MSI Cap MStx Cap Power Management Extd. Capabilities-2 F0 Charles State Capabilities PF0 IDs PF0 BAR Legacy/MSI Cap MStx Cap Power Management Extd. Capabilities-2 F0 Charles State Capabilities Structure MStx Table Settings Coocococo Coocococo Coocococo.lFFFFFFF BAR Indicator EAR 0 Charles State Capabilities Ca | | | |
| P0 2 C Enable MSk Capability Structure 3 MSk Table Settings 2 Table Sce 000 0007FF Table Offset 00000000 000000001FFFFFF BAR Indicator BAR 0 * MSk Pending Bit Array (PBA) Settings 2 PBA Offset 00000000 000000001FFFFFF PBA BAR Indicator BAF 0 * Vanced Options C C | | | |
| In table MSK: Capability Structure MSk: Table Settings Table Settings 000000000000000001FFFFFFF BAR Indicator MSk: Pending Bit Array (PBA) Settings PBA Offset 0000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Indicator PBA Offset 0000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 0000000001FFFFFFF PBA Offset 00000000001FFFFFFF PBA Offset 00000000001FFFFFF PBA Offset 000000000000000000000000000000000000 | | unintes Troibs Trobaix tegacyinorcap rost cap Tower nanagement Exter capabilities 1 Exter capabilities 2 | |
| MSk Table Ster 000 000.7FF Table Ster 00000000.000000.000000000000000000000 | PF0 | | \$ |
| Table Size 000 000.7FF Table Offset 00000000.1FFFFFFF BAR Indicator BAR 0 * | Enable MSI | bi Capability Structure | |
| Table Offset 000000001FFFFFF BAR Indicator BAR 0 VBM Pending Bit Array (PBA) Settings | | | \$ |
| BAR Indicator AR 0 v MSk Pending Bit Array (PBA) Settings PBA Offset 00000000.1FFFFFF PBA BAR Indicator BAF 0 v Kanced Options | Table Size | | |
| MSk Pending Bit Array (PBA) Settings PBA Offset 000000001FFFFFFF PBA BAR Indicator BAF 0 | | | |
| PBA Offset 00000000 0000000.1FFFFFF PBA BAR Indicator BAF 0 | BAR Indicato | r BAR 0 💌 | |
| PBA Offset 00000000 0000000.1FFFFFF PBA BAR Indicator BAF 0 | MSix Pending | a Bit Array (PBA) Settings | â |
| PBA BAR Indicator BAF 0 V | | | |
| Vanced Options | PBA BAR Indi | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | dvanced Optio | | |

Figure 4-12: MSIx Cap Settings

• Enable MSIx Capability Structure: Indicates that the MSI-X Capability structure exists.

Note: The Capability Structure needs at least one Memory BAR to be configured. You must maintain the MSI-X Table and Pending Bit Array in the application.

- **MSIx Table Settings**: Defines the MSI-X Table Structure.
 - Table Size: Specifies the MSI-X Table Size.
 - *Table Offset:* Specifies the Offset from the Base Address Register that points to the Base of the MSI-X Table.

- *BAR Indicator*: Indicates the Base Address Register in the Configuration Space that is used to map the function in the MSI-X Table onto Memory Space. For a 64-bit Base Address Register, this indicates the lower DWORD.
- **MSIx Pending Bit Array (PBA) Settings**: Defines the MSI-X Pending Bit Array (PBA) Structure.
 - *PBA Offset*: Specifies the Offset from the Base Address Register that points to the Base of the MSI-X PBA.
 - *PBA BAR Indicator*: Indicates the Base Address Register in the Configuration Space that is used to map the function in the MSI-X PBA onto Memory Space.

Power Management

The Power Management page shown in Figure 4-13 includes settings for the Power Management Registers, power consumption, and power dissipation options. These settings apply to both Physical Functions, if PF1 is enabled.

| ntegrated Block for PC ly specifying IP Options | 4 Gen3 I Express (1.3) | | | | |
|--|-------------------------------------|-----------------------------------|-----------------------------------|-------------------------|----|
| ons | • | | | | |
| ex-7 FPGA Ge | n3 Integrated Block for PC | CI Express | | | |
| ponent Name pcie3 | 7x_v1_3_0 | | | | |
| Basic Capabilities | PF0 IDs PF0 BAR Legacy/MSI Cap MSI: | x Cap Power Management Extd. Ca | pabilities-1 Extd. Capabilities-2 | | |
| Power Management Ri | egisters | | | | * |
| D1 Support | • | | | | |
| PME Support | | | | | \$ |
| | 3hot | | | | |
| | | | | | |
| BRAM Configuration Op | ations | | | | * |
| Performance Level | Posted Header/Data Credits | Non-posted Header/Data Credits | Completion Header/Data Credits | Total BRAMS Required | |
| Good | • 0x20/0xCC | 0x20/0x28 | 0x00/0x000 | 8/2 | |
| | 0x20/0x198 | 0x20/0x28 | 0x00/0x000 | 4/6 | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |

Figure 4-13: Page 12: Power Management Registers

• **D1 Support**: When selected, this option indicates that the function supports the D1 Power Management State. See section 3.2.3 of the *PCI Bus Power Management Interface Specification Revision 1.2.*

- **PME Support From**: When this option is selected, it indicates the power states in which the function can assert cfg_pm_wake. See section 3.2.3 of the *PCI Bus Power* Management Interface Specification Revision 1.2.
- **BRAM Configuration Options**: Can specify the number of receive block RAMs used for the solution. The table displays the number of receiver credits available for each packet type.

Extended Capabilities 1 and Extended Capabilities 2

The PCIe Extended Capabilities screens shown in Figure 4-14 and Figure 4-15 allow you to enable PCI Express Extended Capabilities. The Advanced Error Reporting Capability (offset 0x100h) is always enabled. The customization GUI sets up the link list based on the capabilities enabled. After enabling, you must configure the capability by setting the applicable attributes in the core top-level defined in Output Generation, page 184. See Appendix D, Attributes for parameters applicable to each capability.



Figure 4-14: Extended Capabilities 1

| | Customize IP | |
|--|--|----|
| Customize Virtex-7 FPGA Gen3 Integrated Block for PCI Express (1.3) by specifying IP Options. | | |
| ions- tex-7 FPGA Gen3 Integrated Block fo | or PCI Express | |
| mponent Name pcie3 7x v1 3 0 | | |
| | P MSIx Cap Power Management Extd. Capabilities-1 Extd. Capabilities-2 | |
| RBAR Capability | | \$ |
| | Extended Capability, which when enabled, adds a capability for Functions with BARs to report various | |
| Enable RBAR Capability(PF0) | | |
| PB Capability | | \$ |
| The Power Budgeting Capability allows a device to report for more details. | power consumption to the system. See Section 7.15 of the PCI Express Base Specification Revision 3.0 | |
| Enable PB Capability(PF0) | | |
| LTR Capability | | \$ |
| The Latency Tolerance Reporting Capability allows Endpoi Specification Revision 3.0 for more details | nts to report Latency requirements to the Root Complex. See Section 6.18 of the PCI Express Base | |
| Enable LTR Capability(PF0) | | |
| TPH Capability | | \$ |
| The Resizable BAR(RBAR) Capabilities is an optional PCIe B options for sizes of their memory mapped resources | Extended Capability, which when enabled, adds a capability for Functions with BARs to report various | |
| Enable TPH Capability(PF0) | | |
| User Defined Configuration Capabilities | | \$ |
| PCI Express Extended Configuration Space Enable | | |
| | | |
| | | |
| Advanced Options | | |

Figure 4-15: Extended Capabilities 2

Device Serial Number Capability

• **Device Serial Number Capability**: An optional PCIe Extended Capability containing a unique Device Serial Number. When this Capability is enabled, the DSN identifier must be presented on the Device Serial Number input pin of the port. This Capability must be turned on to enable the Virtual Channel and Vendor Specific Capabilities

Virtual Channel Capability

- **Virtual Channel Capability**: An optional PCIe Extended Capability which allows the user application to be operated in TCn/VC0 mode. Checking this allows Traffic Class filtering to be supported. This capability only exists for Physical Function 0.
- **Reject Snoop Transactions (Root Port Configuration Only)**: When enabled, any transactions for which the No Snoop attribute is applicable, but is not set in the TLP header, can be rejected as an Unsupported Request.

AER Capability

• **Enable AER Capability**: An optional PCIe Extended Capability that allows Advanced Error Reporting. This capability is always enabled.

Additional Optional Capabilities

- **Enable ARI**: An optional PCIe Extended Capability that allows Alternate Requestor ID. This capability is automatically enabled if SRIOV is enabled.
- **Enable PB**: An optional PCIe Extended Capability that implements the Power Budgeting Enhanced Capability Header.
- **Enable RBAR**: An optional PCIe Extended Capability that implements the Resizable BAR Capability.
- **Enable LTR**: An optional PCIe Extended Capability that implements the Latency Tolerance Reporting Capability.
- **Enable DPA**: An optional PCIe Extended Capability that implements Dynamic Power Allocation Capability.
- **Enable TPH**: An optional PCIe Extended Capability that implements Transaction Processing Hints Capability.

Output Generation

The Gen3 Integrated Block for PCIe example design directories and their associated files are defined in the sections that follow. Click a directory name in blue to go to the desired directory and its associated files.

| <project_directory></project_directory> |
|--|
| <project_directory>.srcs/</project_directory> |
| <pre>in sources_1</pre> |
| 🛅 ip |
| component_name_# |
| component_name_# |
| example_design |
| <pre>component_name_#>/<component_name_#>/simulation</component_name_#></pre> |
| 🚞 simulation/dsport |
| 🚞 simulation/functional |
| imulation/tests |
| <pre>component_name_#>/<component_name_#>/source</component_name_#></pre> |
| <pre>component_name_#>/sim</pre> |
| <pre>component_name_#>/synth</pre> |
| <pre>component_name_#>/source</pre> |

www.xilinx.com

example_design

The example_design directory contains the example design files provided with the core. Table 4-7 shows the directory contents for an Endpoint configuration core.

| Name | Description | |
|---|--|--|
| example_design | | |
| <pre>xilinx_pcie_3_0_ep_7x_01_lane_ gen1_xc7vx690t-ffg1761-3-PCIE_ X0Y0.xdc</pre> | Example design UCF. The file name varies by Device/ Port Type, lane width, maximum link speed, part, package, Integrated Block for PCI Express block location, and Xilinx Development Board selected. | |
| xilinx_pcie_3_0_ep_xt.v | Verilog top-level PIO example design file. | |
| <pre>pcie_app_7vx.v PIO_INTR_CTRL.v EP_MEM.v PIO.v\PIO_EP.v PIO_EP_MEM_ACCESS.v PIO_TO_CTRL.v PIO_RX_ENGINE.v PIO_TX_ENGINE.v</pre> | PIO example design files. | |

 Table 4-7:
 Example Design Directory: Endpoint Configuration

Back to Top

<component_name_#>/<component_name_#>/simulation

The simulation directory contains the simulation source files provided with the core.

simulation/dsport

The dsport directory contains the files for the Root Port model test bench.

 Table 4-8:
 dsport Directory: Endpoint Configuration

| Name | Description | |
|---|------------------------|--|
| <component_name_#>/<component_name_#>/simulation/dsport</component_name_#></component_name_#> | | |
| pcie_2_1_rp_v7.v | | |
| <pre>pci_exp_expect_tasks.v</pre> | | |
| pci_exp_usrapp_cfg.v | | |
| pci_exp_usrapp_com.v | Root Port model files. | |
| pci_exp_usrapp_pl.v | Root Fort model mes. | |
| pci_exp_usrapp_rx.v | | |
| pci_exp_usrapp_tx.v | | |
| <pre>xilinx_pcie_2_1_rport_v7.v</pre> | | |
| | | |

Back to Top

simulation/functional

The functional directory contains functional simulation scripts provided with the core.

| Table 4-9: | Functional | Directory |
|------------|------------|-----------|
| 10010 1 01 | | |

| Name | Description |
|--|--|
| <component_name_#>/<co< td=""><td><pre>omponent_name_#>/simulation/functional</pre></td></co<></component_name_#> | <pre>omponent_name_#>/simulation/functional</pre> |
| board.f | List of files for RTL simulations. |
| simulate_mti.do | Simulation script for ModelSim. |
| simulate_ncsim.sh | Simulation script for Cadence IES (Verilog only). |
| simulate_vcs.sh | Simulation script for VCS (Verilog only). |
| xilinx_lib_vcs.f | Points to the required SecureIP Model. |
| board_common.v (Endpoint configuration only) | Contains test bench definitions (Verilog only). |
| board.v | Top-level simulation module. |
| sys_clk_gen_ds.v (Endpoint configuration only) | System differential clock source. |
| sys_clk_gen.v | System clock source. |

Back to Top

simulation/tests

Note: This directory exists for Endpoint configuration only.

The tests directory contains test definitions for the example test bench.

| Name | Description | |
|--|--|--|
| <component_name_#>/<component_name_#>/simulation/tests</component_name_#></component_name_#> | | |
| sample_tests.v | Test definitions for example test bench. | |
| tests.v | | |
| | • | |

Back to Top

<component_name_#>/<component_name_#>/source

The source directory contains the generated core source files.

| Name | | Description | |
|--|--|--|--|
| <component_name_#>/<component_name_#>/source</component_name_#></component_name_#> | | | |
| pcie_3_0_7vx.v | | Verilog top-level core wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. | |

| Name | Description |
|--|--|
| pcie_top.v | AXI4-Stream solution wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| pcie_7vx.v | Solution Wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| pcie_init_ctrl_7vx.v | Initialization Controller for Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| pcie_pipe_pipeline.v pcie_pipe_lane.v pcie_pipe_misc.v | PIPE module for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| <pre>pcie_bram_7vx.v pcie_bram_7vx_16k.v pcie_bram_7vx_8k.v pcie_bram_7vx_cpl.v pcie_bram_7vx_rep.v pcie_bram_7vx_rep_8k.v pcie_bram_7vx_req.v</pre> | Block RAM module for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| gt_top.v | GTH wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| <pre>gt_wrapper.v pipe_clock.v pipe_drp.v pipe_rate.v pipe_reset.v pipe_sync.v pipe_user.v pipe_wrapper.v pipe_eq.v rxeq_scan.v qpll_drp.v qpll_reset.v qpll_wrapper.v</pre> | GTH module for the Virtex-7 FPGA GTH transceivers. |

Table 4-11: Source Directory (Cont'd)

Back to Top

<component_name_#>/sim

Table 4-12: **sim Directory**

| Name | Description | | | | |
|---|------------------------------------|--|--|--|--|
| <component_na< td=""><td>ame_#>/sim</td></component_na<> | ame_#>/sim | | | | |
| Component_name_#.v | Core top-level file for simulation | | | | |

<component_name_#>/synth

Table 4-13: synth Directory

| Name | Description | | | | |
|---|-----------------------------------|--|--|--|--|
| <pre><component_na< pre=""></component_na<></pre> | me_#/synth | | | | |
| Component_name_#.v | Core top-level file for synthesis | | | | |

<component_name_#>/source

Table 4-14: source Directory

| Name | Description | | | | |
|--|--|--|--|--|--|
| <component_nam< th=""><td>e_#>/source</td></component_nam<> | e_#>/source | | | | |
| <component_name>_pipe_clock.v</component_name> | Clock block used in the example design top level module. | | | | |



Constraining the Core

Required Constraints

The Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® solution requires the specification of timing and other physical implementation constraints to meet specified performance requirements for PCI Express. These constraints are provided with the Endpoint and Root Port solutions in a Xilinx Device Constraints (XDC) file. Pinouts and hierarchy names in the generated XDC correspond to the provided example design.

To achieve consistent implementation results, an XDC containing these original, unmodified constraints must be used when a design is run through the Xilinx tools. For additional details on the definition and use of an XDC or specific constraints, see the *Xilinx Libraries Guide* and/or *Command Line Tools User Guide*.

Constraints provided with the Integrated Block solution have been tested in hardware and provide consistent results. Constraints can be modified, but modifications should only be made with a thorough understanding of the effect of each constraint. Additionally, support is not provided for designs that deviate from the provided constraints.

Device, Package, and Speed Grade Selections

The device selection portion of the XDC informs the implementation tools which part, package, and speed grade to target for the design. Because Gen3 Integrated Block for PCIe cores are designed for specific part and package combinations, this section should not be modified by the designer.

The device selection section always contains a part selection line, but can also contain part or package-specific options. An example part selection line follows:

CONFIG PART = XC7VX690T-FFG1761-3

Clock Frequencies

See Chapter 3, Designing with the Core, for detailed information about clock requirements.

Clock Management

See Chapter 3, Designing with the Core, for detailed information about clock requirements.

Clock Placement

See the 7 Series FPGAs GTX/GTH Transceivers User Guide [Ref 3] for guidelines regarding clock resource selection.

See Chapter 3, Designing with the Core, for detailed information about clock requirements.

Stacked Silicon Interconnect Devices

Some Virtex-7 devices utilize stacked silicon interconnect (SSI) technology. The I/O and Integrated Block must remain on the same die when targeting an SSI device.

The sys_clk must be chosen to be in the same bank as the GTH transceiver it is connected to, or one bank above/below the GTH transceiver being used.

For more information, see the "Placement Information by Package" and "Placement Information by Device" appendices in the 7 Series FPGAs GTX/GTH Transceivers User Guide.

Transceiver Placement

These constraints select which transceivers to use and dictates the pinout for the transceiver differential pairs. For more information, see the "Placement Information by Package" appendix in the 7 Series FPGAs GTX/GTH Transceivers User Guide [Ref 3].

Table 5-1 through Table 5-8 list the supported transceiver locations available for supported Virtex-7 FPGA part and package combinations. The Vivado[™] IP Catalog provides an XDC for the selected part and package that matches the table contents. The following lists all devices with their associated tables containing transceiver locations:

- XC7VX330T: Table 5-1
- XC7VX415T: Table 5-2
- XC7VX550T: Table 5-3
- XC7VX690T: Table 5-4
- XC7VX980T: Table 5-5
- XC7VX1140T: Table 5-6
- XC7VH580T: Table 5-7
- XC7VH870T: Table 5-8

Table 5-1: Supported Transceiver Locations for the XC7VX330T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | | |
|---------|-------|---------------------------------------|---|--------|--------|--------|--------|--------|--------|--|--|--|--|
| | X0Y0 | X0Y11 | X0Y10 | X0Y9 | X0Y8 | X0Y7 | X0Y6 | X0Y5 | X0Y4 | | | | |
| FFG1157 | X0Y1 | X0Y23 | X0Y22 | X0Y21 | X0Y20 | X0Y19 | X0Y18 | X0Y17 | X0Y16 | | | | |
| FEGII21 | X0Y2 | | N/A | | | | | | | | | | |
| | X0Y3 | N/A | | | | | | | | | | | |
| | X0Y0 | X0Y11 | X0Y11 X0Y10 X0Y9 X0Y8 X0Y7 X0Y6 X0Y5 X0 | | | | | | | | | | |
| FFC1761 | X0Y1 | X0Y23 X0Y22 X0Y21 X0Y20 X0Y19 X0Y18 X | | | | | | | X0Y16 | | | | |
| FFG1761 | X0Y2 | | | | N, | /A | | | | | | | |
| | X0Y3 | | | | N, | /A | | | | | | | |

Table 5-2: Supported Transceiver Locations for the XC7VX415T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | | |
|---------|-------|--|---------------------------------------|--------|--------|--------|--------|--------|--------|--|--|--|--|
| | X0Y0 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | X1Y3 | X1Y2 | X1Y1 | X1Y0 | | | | |
| FFG1157 | X0Y1 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | X1Y15 | X1Y14 | X1Y13 | X1Y12 | | | | |
| LLGII21 | X0Y2 | | N/A | | | | | | | | | | |
| | X0Y3 | N/A | | | | | | | | | | | |
| | X0Y0 | X1Y7 | X1Y7 X1Y6 X1Y5 X1Y4 X1Y3 X1Y2 X1Y1 X1 | | | | | | | | | | |
| FFC11F9 | X0Y1 | X1Y19 X1Y18 X1Y17 X1Y16 X1Y15 X1Y14 X1Y1 | | | | | | | X1Y12 | | | | |
| FFG1158 | X0Y2 | | | | N, | /A | | | | | | | |
| | X0Y3 | | | | N | /A | | | | | | | |

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|
| FFG1927 | X0Y0 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | X1Y3 | X1Y2 | X1Y1 | X1Y0 |
| | X0Y1 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | X1Y15 | X1Y14 | X1Y13 | X1Y12 |
| FFG1927 | X0Y2 | | | | N, | /A | | | |
| | X0Y3 | | | | N, | /A | | | |

Table 5-2: Supported Transceiver Locations for the XC7VX415T (Cont'd)

Table 5-3: Supported Transceiver Locations for the XC7VX550T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|---------|-------------------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | | N/A | | | | | | | | | |
| FFG1158 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| LLGII20 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | N/A | | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFG1927 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1927 | X0Y2 X1Y35 | | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| _ | X0Y3 | | | | N, | /A | | | | | | |

Table 5-4: Supported Transceiver Locations for the XC7VX690T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|-----------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | | | | N, | /A | | | | | | |
| FFG1157 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1157 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | N/A | | | | | | | | | | |
| | X0Y0 | | N/A | | | | | | | | | |
| EEC11E9 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1158 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | N/A | | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFG1761 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1701 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | | | N, | /A | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| EEC1026 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1926 - | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | | | N, | /A | | | | | | |

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | |
|---------|-------|--------|----------|----------|--------|--------|--------|--------|--------|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| FFG1927 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| FFG1927 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | N/A | | | | | | | | | |
| | X0Y0 | N/A | | | | | | | | | |
| FFG1930 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| 1191920 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | | <u>.</u> | <u>.</u> | N, | /A | | * | | | |

Table 5-4: Supported Transceiver Locations for the XC7VX690T (Cont'd)

Table 5-5: Supported Transceiver Locations for the XC7VX980T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFC1026 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1926 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | N/A | | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFC1029 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1928 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | N/A | | | | | | | | | | |
| | X0Y0 | N/A | | | | | | | | | | |
| FFG1930 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1930 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | | | N, | /A | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| EEC1022 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1933 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | | | N, | /A | | | | | | |

| Table 5-6: | Supported | Transceiver | Locations | for t | the XC7VX1140T |
|------------|-----------|-------------|-----------|-------|----------------|
|------------|-----------|-------------|-----------|-------|----------------|

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|
| FFG1926 | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 |
| | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 |
| | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 |
| | X0Y3 | | | | N, | /A | | | |

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | |
| FLG1928 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | |
| FLG1920 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | |
| | X0Y3 | X1Y47 | X1Y46 | X1Y45 | X1Y44 | X1Y43 | X1Y42 | X1Y41 | X1Y40 | |
| | X0Y0 | | N/A | | | | | | | |
| FLG1930 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | |
| FLG1930 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | |
| | X0Y3 | N/A | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | |
| ELC1022 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | |
| FLG1933 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | |
| | X0Y3 | | | | N, | /A | | | | |

Table 5-6: Supported Transceiver Locations for the XC7VX1140T (Cont'd)

| Table 5-7: | Supported | Transceiver | Locations | for the | XC7VH580T |
|------------|-----------|-------------|-----------|---------|-----------|
|------------|-----------|-------------|-----------|---------|-----------|

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| HCG1155 | X0Y1 | | | | N, | /A | | | | | |
| псоттоо | X0Y2 | | N/A | | | | | | | | |
| | X0Y3 | | N/A | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| HCG1931 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| псотээт | X0Y2 | N/A | | | | | | | | | |
| | X0Y3 | N/A | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| HCG1932 | X0Y2 | | | | N | /A | | | | | |
| | X0Y3 | | | | N | /A | | | | | |

| Table 5-8: | Supported Transceiver Locations for the XC7VH870T |
|------------|---|
|------------|---|

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|----------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | N/A | | | | | | | | | | |
| 11661001 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| HCG1931 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | | N/A | | | | | | | | |

Gen3 Integrated Block for PCIe (v1.3) PG023 October 16, 2012

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|
| HCG1932 | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | |
| | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | |
| HCG1932 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | |
| | X0Y3 | | N/A | | | | | | | |

Table 5-8: Supported Transceiver Locations for the XC7VH870T

I/O Standard and Placement

This section controls the placement and options for I/Os belonging to the System (SYS) interface and PCI Express (PCI_EXP) interface of the core. NET constraints in this section control the pin location and I/O options for signals in the SYS group. Locations and options vary depending on which derivative of the core is used and should not be changed without fully understanding the system requirements.

For example:

set_property IOSTANDARD LVCMOS18 [get_ports sys_rst_n]
set_property LOC IBUFDS_GTE2_X0Y3 [get_cells refclk_ibuf]

INST constraints control placement of signals that belong to the PCI_EXP group. These constraints control the location of the transceiver(s) used, which implicitly controls pin locations for the transmit and receive differential pair.

For example:

```
set_property LOC GTXE2_CHANNEL_X0Y7 [get_cells {pcie_7x_v1_6_0_i/inst/inst/
gt_top_i/pipe_wrapper_i/pipe_lane[0].gt_wrapper_i/gtx_channel.gtxe2_channel_i}]
```



Detailed Example Design

Directory and File Contents

See Output Generation, page 184 for directory structure and file contents.

Example Design

This section provides an overview of the Virtex®-7 FPGA Gen Integrated Block for PCI Express® example design and instructions for generating the core. It also includes information about simulating and implementing the example design using the provided demonstration test bench.

For current information about generating, simulating, and implementing the core, see the Release Notes provided with the core, when it is generated using the Vivado[™] IP Catalog.

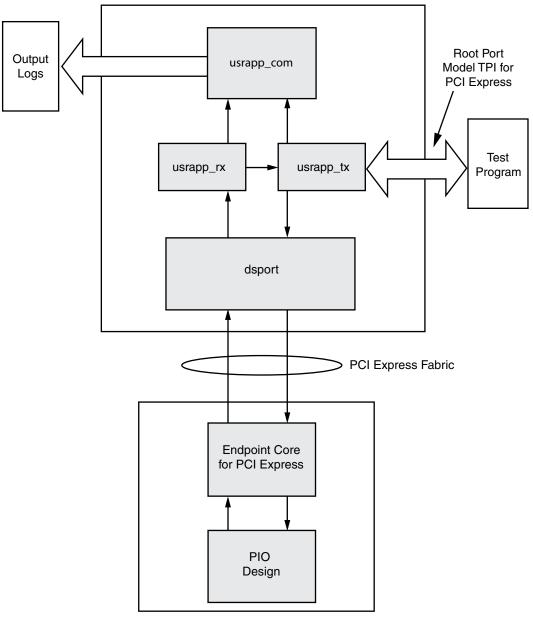
Integrated Block Endpoint Configuration Overview

The example simulation design for the Endpoint configuration of the integrated block consists of two discrete parts:

- The Root Port Model, a test bench that generates, consumes, and checks PCI Express bus traffic.
- The Programmed Input/Output (PIO) example design, a completer application for PCI Express. The PIO example design responds to Read and Write requests to its memory space and can be synthesized for testing in hardware.

Simulation Design Overview

For the simulation design, transactions are sent from the Root Port Model to the Integrated Block core (configured as an Endpoint) and processed by the PIO example design. Figure 6-1 illustrates the simulation design provided with the Integrated Block core. For more information about the Root Port Model, see Root Port Model Test Bench for Endpoint, page 208.



Endpoint DUT for PCI Express

X12471

Figure 6-1: Simulation Example Design Block Diagram

Implementation Design Overview

The implementation design consists of a simple PIO example that can accept read and write transactions and respond to requests, as illustrated in Figure 6-2. Source code for the example is provided with the core. For more information about the PIO example design, see Programmed Input/Output: Endpoint Example Design, page 199.

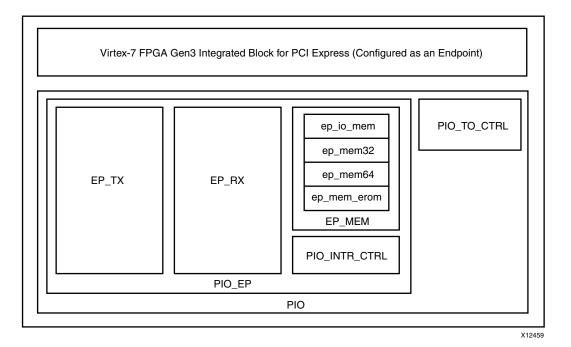


Figure 6-2: Implementation Example Design Block Diagram

Example Design Elements

The PIO example design elements include:

- Core wrapper
- An example Verilog HDL wrapper (instantiates the cores and example design)
- A customizable demonstration test bench to simulate the example design

The example design has been tested and verified with Vivado[™] Design Suite v2012.3 and these simulators:

- Mentor Graphics ModelSim
- Cadence Incisive Enterprise Simulator (IES)
- Synopsys VCS
- Vivado Simulator

For the supported versions of these tools, see the Xilinx Design Tools: Release Notes Guide.

www.xilinx.com

Programmed Input/Output: Endpoint Example Design

Programmed Input/Output (PIO) transactions are generally used by a PCI Express system host CPU to access Memory Mapped Input/Output (MMIO) and Configuration Mapped Input/Output (CMIO) locations in the PCI Express logic. Endpoints for PCI Express accept Memory and I/O Write transactions and respond to Memory and I/O Read transactions with Completion with Data transactions.

The PIO example design (PIO design) is included with the Gen3 Integrated Block for PCIe in Endpoint configuration generated by the CORE Generator[™] tool, which allows users to bring up their system board with a known established working design to verify the link and functionality of the board.

The PIO design Port Model is shared by the Gen3 Integrated Block for PCIe, Endpoint Block Plus for PCI Express, and Endpoint PIPE for PCI Express solutions. This appendix represents all the solutions generically using the name Endpoint for PCI Express (or Endpoint for PCIe[®]).

System Overview

The PIO design is a simple target-only application that interfaces with the Endpoint for the PCIe core Transaction (AXI4-Stream) interface and is provided as a starting point for customers to build their own designs. These features are included:

- Four transaction-specific 2 KB target regions using the internal FPGA block RAMs, providing a total target space of 8192 bytes
- Supports single Dword payload Read and Write PCI Express transactions to 32-/64-bit address memory spaces and I/O space with support for completion TLPs
- Utilizes the BAR ID[2:0] and Completer Request Descriptor[114:112] of the core to differentiate between TLP destination Base Address Registers
- Provides separate implementations optimized for 64-bit, 128-bit, and 256-bit AXI4-Stream interfaces

Figure 6-3 illustrates the PCI Express system architecture components, consisting of a Root Complex, a PCI Express switch device, and an Endpoint for PCIe. PIO operations move data *downstream* from the Root Complex (CPU register) to the Endpoint, and/or *upstream* from the Endpoint to the Root Complex (CPU register). In either case, the PCI Express protocol request to move the data is initiated by the host CPU.

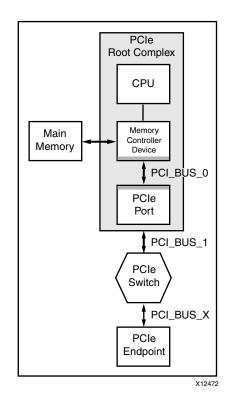


Figure 6-3: System Overview

Data is moved downstream when the CPU issues a store register to a MMIO address command. The Root Complex typically generates a Memory Write TLP with the appropriate MMIO location address, byte enables, and the register contents. The transaction terminates when the Endpoint receives the Memory Write TLP and updates the corresponding local register.

Data is moved upstream when the CPU issues a load register from a MMIO address command. The Root Complex typically generates a Memory Read TLP with the appropriate MMIO location address and byte enables. The Endpoint generates a Completion with Data TLP after it receives the Memory Read TLP. The Completion is steered to the Root Complex and payload is loaded into the target register, completing the transaction.

PIO Hardware

The PIO design implements a 8192 byte target space in FPGA block RAM, behind the Endpoint for PCIe. This 32-bit target space is accessible through single Dword I/O Read, I/O Write, Memory Read 64, Memory Write 64, Memory Read 32, and Memory Write 32 TLPs.

The PIO design generates a completion with one Dword of payload in response to a valid Memory Read 32 TLP, Memory Read 64 TLP, or I/O Read TLP request presented to it by the core. In addition, the PIO design returns a completion without data with successful status for I/O Write TLP request.

The PIO design can initiate:

- a Memory Read transaction when the received write address is 11 ' hEA8 and the write data is 32 ' hAAAA_BBBB, and Targeting the BAR0.
- a Legacy Interrupt when the received write address is 11 ' hEEC and the write data is 32 ' hCCCC_DDDD, and Targeting the BAR0.
- an MSI when the received write address is 11 'hEEC and the write data is 32 'hEEEE_FFFF, and Targeting the BAR0.
- an MSIx when the received write address is 11 ' hEEC and the write data is 32 ' hDEAD_BEEF, and Targeting the BARO.

The PIO design processes a Memory or I/O Write TLP with one Dword payload by updating the payload into the target address in the FPGA block RAM space.

Base Address Register Support

The PIO design supports four discrete target spaces, each consisting of a 2 KB block of memory represented by a separate Base Address Register (BAR). Using the default parameters, the CORE Generator tool produces a core configured to work with the PIO design defined in this section, consisting of:

- One 64-bit addressable Memory Space BAR
- One 32-bit Addressable Memory Space BAR

Users can change the default parameters used by the PIO design; however, in some cases they might need to change the back-end User Application depending on their system. See Changing IP Catalog Tool Default BAR Settings for information about changing the default CORE Generator tool parameters and the effect on the PIO design.

Each of the four 2 KB address spaces represented by the BARs corresponds to one of four 2 KB address regions in the PIO design. Each 2 KB region is implemented using a 2 KB dual-port block RAM. As transactions are received by the core, the core decodes the address and determines which of the four regions is being targeted. The core presents the TLP to the PIO design and asserts the appropriate bits of (BAR ID[2:0]), Completer Request Descriptor[114:112], as defined in Table 6-1.

| Block RAM | TLP Transaction Type | Default BAR | BAR ID[2:0] |
|-------------|--|---------------|-------------|
| ep_io_mem | I/O TLP transactions | Disabled | Disabled |
| ep_mem32 | 32-bit address Memory TLP transactions | 2 | 000b |
| ep_mem64 | 64-bit address Memory TLP transactions | 0-1 | 001b |
| ep_mem_erom | 32-bit address Memory TLP transactions destined for EROM | Expansion ROM | 110b |

| Table 6-1: | TLP Traffic Types |
|------------|-------------------|
|------------|-------------------|

Changing IP Catalog Tool Default BAR Settings

You can change the Vivado IP Catalog tool parameters and continue to use the PIO design to create customized Verilog source to match the selected BAR settings. However, because the PIO design parameters are more limited than the core parameters, consider these example design limitations when changing the default IP Catalog tool parameters:

- The example design supports one I/O space BAR, one 32-bit Memory space (that cannot be the Expansion ROM space), and one 64-bit Memory space. If these limits are exceeded, only the first space of a given type is active—accesses to the other spaces do not result in completions.
- Each space is implemented with a 2 KB memory. If the corresponding BAR is configured to a wider aperture, accesses beyond the 2 KB limit wrap around and overlap the 2 KB memory space.
- The PIO design supports one I/O space BAR, which by default is disabled, but can be changed if desired.

Although there are limitations to the PIO design, Verilog source code is provided so users can tailor the example design to their specific needs.

TLP Data Flow

This section defines the data flow of a TLP successfully processed by the PIO design.

The PIO design successfully processes single Dword payload Memory Read and Write TLPs and I/O Read and Write TLPs. Memory Read or Memory Write TLPs of lengths larger than one Dword are not processed correctly by the PIO design; however, the core does accept these TLPs and passes them along to the PIO design. If the PIO design receives a TLP with a length of greater than one Dword, the TLP is received completely from the core and discarded. No corresponding completion is generated.

Memory and I/O Write TLP Processing

When the Endpoint for PCIe receives a Memory or I/O Write TLP, the TLP destination address and transaction type are compared with the values in the core BARs. If the TLP passes this comparison check, the core passes the TLP to the Receive AXI4-Stream interface of the PIO design. The PIO design handles Memory writes and I/O TLP writes in different ways: the PIO design responds to *I/O writes* by generating a Completion Without Data (cpl), a requirement of the PCI Express specification.

Along with the start of packet, end of packet, and ready handshaking signals, the Completer Requester AXI4-Stream interface also asserts the appropriate (BAR ID[2:0]), Completer Request Descriptor[114:112] signal to indicate to the PIO design the specific destination BAR that matched the incoming TLP. On reception, the PIO design RX State Machine processes the incoming Write TLP and extracts the TLPs data and relevant address fields so that it can pass this along to the PIO design internal block RAM write request controller. Based on the specific BAR ID[2:0] signals asserted, the RX state machine indicates to the internal write controller the appropriate 2 KB block RAM to use prior to asserting the write enable request. For example, if an I/O Write Request is received by the core targeting BARO, the core passes the TLP to the PIO design and sets BAR ID[2:0] to 000b. The RX state machine extracts the lower address bits and the data field from the I/O Write TLP and instructs the internal Memory Write controller to begin a write to the block RAM.

In this example, the assertion of setting BAR ID[2:0] to 000b instructed the PIO memory write controller to access ep_mem0 (which by default represents 2 KB of I/O space). While the write is being carried out to the FPGA block RAM, the PIO design RX state machine deasserts m_axis_cq_tready, causing the Receive AXI4-Stream interface to stall receiving any further TLPs until the internal Memory Write controller completes the write to the block RAM. Deasserting m_axis_cq_tready in this way is not required for all designs using the core—the PIO design uses this method to simplify the control logic of the RX state machine.

Memory and I/O Read TLP Processing

When the Endpoint for PCIe receives a Memory or I/O Read TLP, the TLP destination address and transaction type are compared with the values programmed in the core BARs. If the TLP passes this comparison check, the core passes the TLP to the Receive AXI4-Stream interface of the PIO design.

Along with the start of packet, end of packet, and ready handshaking signals, the Completer Requester AXI4-Stream interface also asserts the appropriate BAR ID[2:0] signal to indicate to the PIO design the specific destination BAR that matched the incoming TLP. On reception, the PIO design state machine processes the incoming Read TLP and extracts the relevant TLP information and passes it along to the PIO design's internal block RAM read request controller.

Based on the specific BAR ID[2:0] signal asserted, the RX state machine indicates to the internal read request controller the appropriate 2 KB block RAM to use before asserting the read enable request. For example, if a Memory Read 32 Request TLP is received by the core targeting the default Mem32 BAR2, the core passes the TLP to the PIO design and sets BAR ID[2:0] to 010b. The RX state machine extracts the lower address bits from the Memory 32 Read TLP and instructs the internal Memory Read Request controller to start a read operation.

In this example, the setting BAR ID[2:0] to 010b instructs the PIO memory read controller to access the Mem32 space, which by default represents 2 KB of memory space. A notable difference in handling of memory write and read TLPs is the requirement of the receiving device to return a Completion with Data TLP in the case of memory or I/O read request.

While the read is being processed, the PIO design RX state machine deasserts m_axis_cq_tready, causing the Receive AXI4-Stream interface to stall receiving any further TLPs until the internal Memory Read controller completes the read access from the block RAM and generates the completion. Deasserting m_axis_cq_tready in this way is not required for all designs using the core. The PIO design uses this method to simplify the control logic of the RX state machine.

PIO File Structure

Table 6-2 defines the PIO design file structure. Based on the specific core targeted, not all files delivered by the Vivado IP Catalog are necessary, and some files might not be delivered. The major difference is that some of the Endpoint for PCIe solutions use a 32-bit user datapath, others use a 64-bit datapath, and the PIO design works with both. The width of the datapath depends on the specific core being targeted.

| File | Description |
|---------------------|--------------------------------|
| PIO.v | Top-level design wrapper |
| PIO_INTR_CTRL.v | PIO interrupt controller |
| PIO_EP.v | PIO application module |
| PIO_TO_CTRL.v | PIO turn-off controller module |
| PIO_RX_ENGINE.v | 32-bit Receive engine |
| PIO_TX_ENGINE.v | 32-bit Transmit engine |
| PIO_EP_MEM_ACCESS.v | Endpoint memory access module |
| PIO_EP_MEM.v | Endpoint memory |

Table 6-2: PIO Design File Structure

Three configurations of the PIO design are provided: PIO_64, PIO_128, and PIO_256 with 64-, 128-, and 256-bit AXI4-Stream interfaces, respectively. The PIO configuration generated depends on the selected Endpoint type (that is, Virtex-7 FPGA integrated block, PIPE, PCI Express, and Block Plus) as well as the number of PCI Express lanes and the interface width selected by the user. Table 6-3 identifies the PIO configuration generated based on your selection.

Table 6-3: PIO Configuration

| Core | x1 | x2 | x4 | x8 |
|-------------------------------------|--------|--------------------|--------------------------------|--|
| Virtex-7 FPGA Gen3 Integrated Block | PIO_64 | PIO_64, PIO_128 | PIO_64, PIO_128, PIO_256 | PIO_64, PIO_128 ⁽¹⁾ , PIO_256 |

Notes:

1. The core does not support 128-bit x8 8.0 Gb/s configuration and 500 MHz user clock frequency.

Figure 6-4 shows the various components of the PIO design, which is separated into four main parts: the TX Engine, RX Engine, Memory Access Controller, and Power Management Turn-Off Controller.

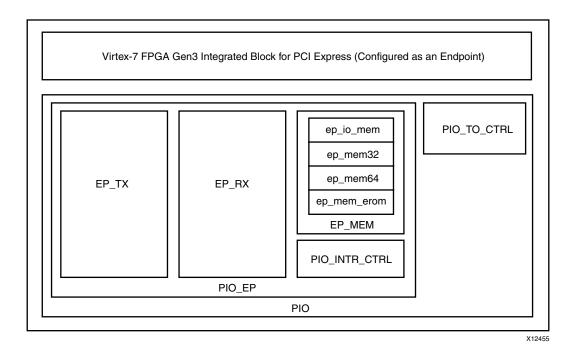


Figure 6-4: PIO Design Components

PIO Operation

PIO Read Transaction

Figure 6-5 depicts a Back-to-Back Memory Read request to the PIO design. The receive engine deasserts m_axis_rx_tready as soon as the first TLP is completely received. The next Read transaction is accepted only after compl_done_o is asserted by the transmit engine, indicating that Completion for the first request was successfully transmitted.

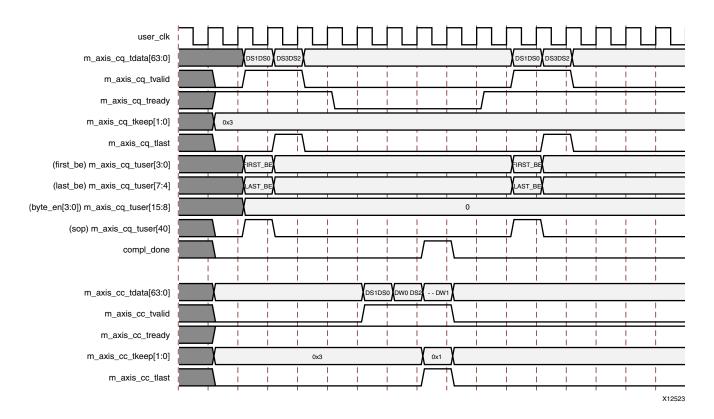


Figure 6-5: Back-to-Back Read Transactions

PIO Write Transaction

Figure 6-6 depicts a back-to-back Memory Write to the PIO design. The next Write transaction is accepted only after wr_busy_o is deasserted by the memory access unit, indicating that data associated with the first request was successfully written to the memory aperture.

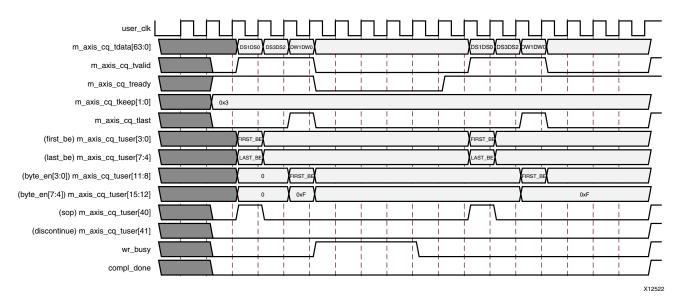


Figure 6-6: Back-to-Back Write Transactions

Device Utilization

Table 6-4 shows the PIO design FPGA resource utilization.

Table 6-4: PIO Design FPGA Resources

| Resources | Utilization |
|------------|-------------|
| LUTs | 300 |
| Flip-Flops | 500 |
| Block RAMs | 4 |

Demonstration Test Bench

Root Port Model Test Bench for Endpoint

The PCI Express Root Port Model is a robust test bench environment that provides a test program interface that can be used with the provided PIO design or with your design. The purpose of the Root Port Model is to provide a source mechanism for generating downstream PCI Express TLP traffic to stimulate the customer design, and a destination mechanism for receiving upstream PCI Express TLP traffic from the customer design in a simulation environment.

Source code for the Root Port Model is included to provide the model for a starting point for the user test bench. All the significant work for initializing the core configuration space, creating TLP transactions, generating TLP logs, and providing an interface for creating and verifying tests are complete, allowing you to dedicate efforts to verifying the correct functionality of the design rather than spending time developing an Endpoint core test bench infrastructure.

The Root Port Model consists of:

- Test Programming Interface (TPI), which allows the user to stimulate the Endpoint device for the PCI Express
- Example tests that illustrate how to use the test program TPI
- Verilog source code for all Root Port Model components, which allow you to customize the test bench

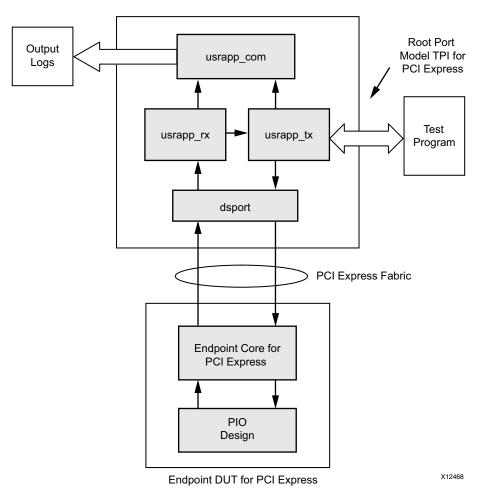


Figure 6-7 illustrates the illustrates the Root Port Model coupled with the PIO design.

Figure 6-7: Root Port Model and Top-Level Endpoint

Architecture

The Root Port Model consists of these blocks, illustrated in Figure 6-7:

- dsport (Root Port)
- usrapp_tx
- usrapp_rx
- usrapp_com (Verilog only)

The usrapp_tx and usrapp_rx blocks interface with the dsport block for transmission and reception of TLPs to/from the Endpoint Design Under Test (DUT). The Endpoint DUT consists of the Endpoint for PCIe and the PIO design (displayed) or customer design.

The usrapp_tx block sends TLPs to the dsport block for transmission across the PCI Express Link to the Endpoint DUT. In turn, the Endpoint DUT device transmits TLPs across the PCI Express Link to the dsport block, which are subsequently passed to the usrapp_rx block. The dsport and core are responsible for the data link layer and physical link layer processing when communicating across the PCI Express logic. Both usrapp_tx and usrapp_rx utilize the usrapp_com block for shared functions, for example, TLP processing and log file outputting. Transaction sequences or test programs are initiated by the usrapp_tx block to stimulate the Endpoint device's fabric interface. TLP responses from the Endpoint device are received by the usrapp_rx block. Communication between the usrapp_tx and usrapp_rx blocks allow the usrapp_tx block to verify correct behavior and act accordingly when the usrapp_rx block has received TLPs from the Endpoint device.

Simulating the Design

Simulation script files are provided with the model to facilitate simulation with various simulation tools:

- vsim -do simulate_mti.do
- ./simulate_ncsim.sh
- ./simulate_vcs.sh

The example simulation script files are located in this directory:

<project_dir>/<component_name>/simulation/functional

Instructions for simulating the PIO design using the Root Port Model are provided in Programmed Input/Output: Endpoint Example Design, page 199.

Scaled Simulation Timeouts

The simulation model of the Gen3 Integrated Block for PCIe uses scaled down times during link training to allow for the link to train in a reasonable amount of time during simulation. According to the *PCI Express Specification, rev. 3.0*, there are various timeouts associated with the link training and status state machine (LTSSM) states. The Gen3 Integrated Block for PCIe scales these timeouts by a factor of 256 in simulation, except in the Recovery Speed_1 LTSSM state, where the timeouts are not scaled.

Test Selection

Table 6-5 describes the tests provided with the Root Port Model, followed by specific sections for Verilog test selection.

Table 6-5: Root Port Model Provided Tests

| Test Name | Test in Verilog | Description |
|--------------------|-----------------|---|
| sample_smoke_test0 | Verilog | Issues a PCI Type 0 Configuration Read TLP and waits for the completion TLP; then compares the value returned with the expected Device/Vendor ID value. |
| sample_smoke_test1 | Verilog | Performs the same operation as sample_smoke_test0 but makes use of expectation tasks. This test uses two separate test program threads: one thread issues the PCI Type 0 Configuration Read TLP and the second thread issues the Completion with Data TLP expectation task. This test illustrates the form for a parallel test that uses expectation tasks. This test form allows for confirming reception of any TLPs from your design. Additionally, this method can be used to confirm reception of TLPs when ordering is unimportant. |

Verilog Test Selection

The Verilog test model used for the Root Port Model lets the user specify the name of the test to be run as a command line parameter to the simulator.

To change the test to be run, change the value provided to TESTNAME, which is defined in the test files sample_tests1.v and pio_tests.v. This mechanism is used for ModelSim. ISim uses the -testplusarg options to specify TESTNAME, for example: demo_tb.exe -gui -view wave.wcfg -wdb wave_isim -tclbatch isim_cmd.tcl -testplusarg TESTNAME=sample_smoke_test0.

Waveform Dumping

Table 6-6 describes the available simulator waveform dump file formats, provided in the simulator native file format. The same mechanism is used for ModelSim.

| Table 6-6: Simulator Dump File Forma | Table 6-6: | Simulator | Dump | File | Format |
|--------------------------------------|------------|-----------|------|------|--------|
|--------------------------------------|------------|-----------|------|------|--------|

| Simulator | Dump File Format |
|----------------------------------|------------------|
| Mentor Graphics ModelSim | .vcd |
| Synopsys VCS and Synopsys VCS_MX | .vpd |
| Cadence IES | .trn |

Verilog Flow

The Root Port Model provides a mechanism for outputting the simulation waveform to file by specifying the **+dump_all** command line parameter to the simulator.

Output Logging

When a test fails on the example or customer design, the test programmer debugs the offending test case. Typically, the test programmer inspects the wave file for the simulation and cross-reference this to the messages displayed on the standard output. Because this approach can be very time consuming, the Root Port Model offers an output logging mechanism to assist the tester with debugging failing test cases to speed the process.

The Root Port Model creates three output files (tx.dat, rx.dat, and error.dat) during each simulation run. Log files rx.dat and tx.dat each contain a detailed record of every TLP that was received and transmitted, respectively, by the Root Port Model.

0

TIP: With an understanding of the expected TLP transmission during a specific test case, you can isolate the failure.

The log file error.dat is used in conjunction with the expectation tasks. Test programs that utilize the expectation tasks generate a general error message to standard output. Detailed information about the specific comparison failures that have occurred due to the expectation error is located within error.dat.

Parallel Test Programs

There are two classes of tests are supported by the Root Port Model:

- Sequential tests. Tests that exist within one process and behave similarly to sequential programs. The test depicted in Test Program: pio_writeReadBack_test0, page 213 is an example of a sequential test. Sequential tests are very useful when verifying behavior that have events with a known order.
- Parallel tests. Tests involving more than one process thread. The test sample_smoke_test1 is an example of a parallel test with two process threads. Parallel tests are very useful when verifying that a specific set of events have occurred, however the order of these events are not known.

A typical parallel test uses the form of one command thread and one or more expectation threads. These threads work together to verify a device's functionality. The role of the command thread is to create the necessary TLP transactions that cause the device to receive and generate TLPs. The role of the expectation threads is to verify the reception of an expected TLP. The Root Port Model TPI has a complete set of expectation tasks to be used in conjunction with parallel tests.

Because the example design is a target-only device, only Completion TLPs can be expected by parallel test programs while using the PIO design. However, the full library of expectation tasks can be used for expecting any TLP type when used in conjunction with the customer's design (which can include bus-mastering functionality).

Test Description

The Root Port Model provides a Test Program Interface (TPI). The TPI provides the means to create tests by invoking a series of Verilog tasks. All Root Port Model tests should follow the same six steps:

- 1. Perform conditional comparison of a unique test name
- 2. Set up master timeout in case simulation hangs
- 3. Wait for Reset and link-up
- 4. Initialize the configuration space of the Endpoint
- 5. Transmit and receive TLPs between the Root Port Model and the Endpoint DUT
- 6. Verify that the test succeeded

Test Program: pio_writeReadBack_test0

```
1.
       else if(testname == "pio_writeReadBack_test1"
2.
       begin
       // This test performs a 32 bit write to a 32 bit Memory space and performs a read back
3
       TSK_SIMULATION_TIMEOUT(10050);
4.
5.
       TSK SYSTEM INITIALIZATION;
       TSK_BAR_INIT;
6.
       for (ii = 0; ii <= 6; ii = ii + 1) begin
7.
8.
          if (BAR_INIT_P_BAR_ENABLED[ii] > 2'b00) // bar is enabled
           case(BAR_INIT_P_BAR_ENABLED[ii])
9.
                   2'b01 : // IO SPACE
10
11.
                  begin
12.
                       $display("[%t] : NOTHING: to IO 32 Space BAR %x", $realtime, ii);
                  end
13.
                   2'b10 : // MEM 32 SPACE
14
15.
                     begin
16.
                      $display("[%t] : Transmitting TLPs to Memory 32 Space BAR %x",
17.
                                 $realtime, ii);
18.
              //-----
                                                  _____
19.
              // Event : Memory Write 32 bit TLP
20.
              //-----
                      DATA\_STORE[0] = 8'h04;
21.
22.
                        DATA\_STORE[1] = 8'h03;
23.
                        DATA\_STORE[2] = 8'h02;
                        DATA_STORE[3] = 8'h01;
24.
                        P_READ_DATA = 32'hffff_fff; // make sure P_READ_DATA has known initial value
25.
                        TSK_TX_MEMORY_WRITE_32(DEFAULT_TAG, DEFAULT_TC, 10'd1, BAR_INIT_P_BAR[ii][31:0] , 4'hF,
26.
       4'hF, 1'b0);
27.
                       TSK_TX_CLK_EAT(10);
28.
                       DEFAULT_TAG = DEFAULT_TAG + 1;
29.
                  //----
                                                    30.
                  // Event : Memory Read 32 bit TLP
31.
                  //-----
                        TSK_TX_MEMORY_READ_32(DEFAULT_TAG, DEFAULT_TC, 10'd1, BAR_INIT_P_BAR[ii][31:0], 4'hF,
32.
       4'hF);
33.
                        TSK_WAIT_FOR_READ_DATA;
34.
                        if (P_READ_DATA != {DATA_STORE[3], DATA_STORE[2], DATA_STORE[1], DATA_STORE[0] })
35.
                         begin
36.
                           $display("[%t] : Test FAILED --- Data Error Mismatch, Write Data %x != Read Data %x",
       $realtime,{DATA_STORE[3], DATA_STORE[2], DATA_STORE[1], DATA_STORE[0]}, P_READ_DATA);
37.
                         end
38.
                      else
39.
                        begin
40.
                           $display("[%t] : Test PASSED --- Write Data: %x successfully received", $realtime,
       P_READ_DATA);
41.
                        end
```

Expanding the Root Port Model

The Root Port Model was created to work with the PIO design, and for this reason is tailored to make specific checks and warnings based on the limitations of the PIO design. These checks and warnings are enabled by default when the Root Port Model is generated by the CORE Generator tool. However, these limitations can be disabled so that they do not affect the customer's design.

Because the PIO design was created to support at most one I/O BAR, one Mem64 BAR, and two Mem32 BARs (one of which must be the EROM space), the Root Port Model by default makes a check during device configuration that verifies that the core has been configured to meet this requirement. A violation of this check causes a warning message to be displayed as well as for the offending BAR to be gracefully disabled in the test bench. This check can be disabled by setting the pio_check_design variable to zero in the pci_exp_usrapp_tx.v file.

Root Port Model TPI Task List

The Root Port Model TPI tasks include these tasks, which are further defined in these tables.

- Table 6-7, Test Setup Tasks
- Table 6-8, TLP Tasks
- Table 6-9, BAR Initialization Tasks
- Table 6-10, Example PIO Design Tasks
- Table 6-11, Expectation Tasks

| Name | Input(s) | | Description | |
|---------------------------|-------------|-------|--|--|
| TSK_SYSTEM_INITIALIZATION | None | | Waits for transaction interface reset and link-up between the Root Port Model and the Endpoint DUT. This task must be invoked prior to the Endpoint core initialization. | |
| TSK_USR_DATA_SETUP_SEQ | None | | Initializes global 4096 byte DATA_STORE array entries to sequential values from zero to 4095. | |
| TSK_TX_CLK_EAT | clock count | 31:30 | Waits clock_count transaction interface clocks. | |
| TSK_SIMULATION_TIMEOUT | timeout | 31:0 | Sets master simulation timeout value in units of transaction interface clocks. This task should be used to ensure that all DUT tests complete. | |

Table 6-7: Test Setup Tasks

Table 6-8: TLP Tasks

| Name | Input(s) | | Description |
|----------------------------------|--|---|---|
| TSK_TX_TYPE0_CONFIGURATION_READ | tag_ reg_addr_ first_dw_be_ | 7:0 11:0 3:0 | Waits for transaction interface reset and link-up between the Root Port Model and the Endpoint DUT. This task must be invoked prior to Endpoint core initialization. |
| TSK_TX_TYPE1_CONFIGURATION_READ | tag_ reg_addr_ first_dw_be_ | 7:0 11:0 3:0 | Sends a Type 1 PCI Express Config Read TLP from Root Port Model to reg_addr_ of Endpoint DUT with tag_ and first_dw_be_ inputs. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_TYPE0_CONFIGURATION_WRITE | tag_ reg_addr_ reg_data_ first_dw_be_ | 7:0 11:0 31:0 3:0 | Sends a Type 0 PCI Express Config Write TLP from Root Port Model to reg_addr_ of Endpoint DUT with tag_ and first_dw_be_ inputs. Cpl returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_TYPE1_CONFIGURATION_WRITE | tag_ reg_addr_ reg_data_ first_dw_be_ | 7:0 11:0 31:0 3:0 | Sends a Type 1 PCI Express Config Write TLP from Root Port Model to reg_addr_ of Endpoint DUT with tag_ and first_dw_be_ inputs. Cpl returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MEMORY_READ_32 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ | 7:0 2:0 10:0 31:0 3:0 3:0 | Sends a PCI Express Memory Read TLP from Root Port to 32-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MEMORY_READ_64 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ | 7:0 2:0 10:0 63:0 3:0 3:0 | Sends a PCI Express Memory Read TLP from Root Port Model to 64-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MEMORY_WRITE_32 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ ep_ | 7:0 2:0 10:0 31:0 3:0 3:0 - | Sends a PCI Express Memory Write TLP from Root Port Model to 32-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. The global DATA_STORE byte array is used to pass write data to task. |

Table 6-8: TLP Tasks (Cont'd)

| Name | Input(s) | | Description |
|------------------------|---|---|---|
| TSK_TX_MEMORY_WRITE_64 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ ep_ | 7:0 2:0 10:0 63:0 3:0 3:0 - | Sends a PCI Express Memory Write TLP from Root Port Model to 64-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. The global DATA_STORE byte array is used to pass write data to task. |
| TSK_TX_COMPLETION | tag_ tc_ len_ comp_status_ | 7:0 2:0 10:0 2:0 | Sends a PCI Express Completion TLP from Root Port Model to the Endpoint DUT using global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_COMPLETION_DATA | tag_ tc_ len_ byte_count lower_addr comp_status ep_ | 7:0 2:0 10:0 11:0 6:0 2:0 - | Sends a PCI Express Completion with Data TLP from Root Port Model to the Endpoint DUT using global COMPLETE_ID_CFG as the completion ID. The global DATA_STORE byte array is used to pass completion data to task. |
| TSK_TX_MESSAGE | tag_ tc_ len_ data message_rtg message_code | 7:0 2:0 10:0 63:0 2:0 7:0 | Sends a PCI Express Message TLP from Root Port Model to Endpoint DUT. Completion returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MESSAGE_DATA | tag_ tc_ len_ data message_rtg message_code | 7:0 2:0 10:0 63:0 2:0 7:0 | Sends a PCI Express Message with Data TLP from Root Port Model to Endpoint DUT. The global DATA_STORE byte array is used to pass message data to task. Completion returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_IO_READ | tag_ addr_ first_dw_be_ | 7:0 31:0 3:0 | Sends a PCI Express I/O Read TLP from Root Port Model to I/O address addr_[31:2] of the Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_IO_WRITE | tag_ addr_ first_dw_be_ data | 7:0 31:0 3:0 31:0 | Sends a PCI Express I/O Write TLP from Root Port Model to I/O address addr_[31:2] of the Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |

Table 6-8: TLP Tasks (Cont'd)

| Name | Input(s) | | Description |
|------------------------|--|-----------------------------------|--|
| TSK_TX_BAR_READ | bar_index byte_offset tag_ tc_ | 2:0 31:0 7:0 2:0 | Sends a PCI Express one Dword Memory 32, Memory 64, or I/O Read TLP from the Root Port Model to the target address corresponding to offset byte_offset from BAR bar_index of the Endpoint DUT. This task sends the appropriate Read TLP based on how BAR bar_index has been configured during initialization. This task can only be called after TSK_BAR_INIT has successfully completed. CpID returned from the Endpoint DUT use the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_BAR_WRITE | bar_index byte_offset tag_ tc_ data_ | 2:0 31:0 7:0 2:0 31:0 | Sends a PCI Express one Dword Memory 32, Memory 64, or I/O Write TLP from the Root Port to the target address corresponding to offset byte_offset from BAR bar_index of the Endpoint DUT. This task sends the appropriate Write TLP based on how BAR bar_index has been configured during initialization. This task can only be called after TSK_BAR_INIT has successfully completed. |
| TSK_WAIT_FOR_READ_DATA | None | | Waits for the next completion with data TLP that was sent by the Endpoint DUT. On successful completion, the first Dword of data from the CpID is stored in the global P_READ_DATA. This task should be called immediately following any of the read tasks in the TPI that request Completion with Data TLPs to avoid any race conditions. By default this task locally times out and terminate the simulation after 1000 transaction interface clocks. The global cpld_to_finish can be set to zero so that local time out returns execution to the calling test and does not result in simulation timeout. For this case test programs should check the global cpld_to, which when set to one indicates that this task has timed out and that the contents of P_READ_DATA are invalid. |

Table 6-9: BAR Initialization Tasks

| Name | Input(s) | Description |
|----------------------|----------|---|
| TSK_BAR_INIT | None | Performs a standard sequence of Base Address Register initialization tasks to the Endpoint device using the PCI Express fabric. Performs a scan of the Endpoint's PCI BAR range requirements, performs the necessary memory and I/O space mapping calculations, and finally programs the Endpoint so that it is ready to be accessed. On completion, the user test program can begin memory and I/O transactions to the device. This function displays to standard output a memory and I/O table that details how the Endpoint has been initialized. This task also initializes global variables within the Root Port Model that are available for test program usage. This task should only be called after TSK_SYSTEM_INITIALIZATION. |
| TSK_BAR_SCAN | None | Performs a sequence of PCI Type 0 Configuration Writes and Configuration Reads using the PCI Express logic to determine the memory and I/O requirements for the Endpoint. The task stores this information in the global array BAR_INIT_P_BAR_RANGE[]. This task should only be called after TSK_SYSTEM_INITIALIZATION. |
| TSK_BUILD_PCIE_MAP | None | Performs memory and I/O mapping algorithm and allocates Memory 32, Memory 64, and I/O space based on the Endpoint requirements. This task has been customized to work in conjunction with the limitations of the PIO design and should only be called after completion of TSK_BAR_SCAN. |
| TSK_DISPLAY_PCIE_MAP | None | Displays the memory mapping information of the Endpoint core PCI Base Address Registers. For each BAR, the BAR value, the BAR range, and BAR type is given. This task should only be called after completion of TSK_BUILD_PCIE_MAP. |

Table 6-10: Example PIO Design Tasks

| Name | Input(s | s) | Description |
|------------------------|-----------|-----|---|
| TSK_TX_READBACK_CONFIG | None | | Performs a sequence of PCI Type 0 Configuration Reads to the Endpoint device's Base Address Registers, PCI Command Register, and PCIe Device Control Register using the PCI Express logic. This task should only be called after TSK_SYSTEM_INITIALIZATION. |
| TSK_MEM_TEST_DATA_BUS | bar_index | 2:0 | Tests whether the PIO design FPGA block RAM data bus interface is correctly connected by performing a 32-bit walking ones data test to the I/O or memory address pointed to by the input bar_index. For an exhaustive test, this task should be called four times, once for each block RAM used in the PIO design. |

| Table 6-10: | Example P | PIO Design | Tasks | (Cont'd) |
|-------------|-----------|------------|-------|----------|
| | Example 1 | IO Design | 10313 | |

| Name | Input(| s) | Description |
|-----------------------|---------------------|-------------|--|
| TSK_MEM_TEST_ADDR_BUS | bar_index nBytes | 2:0 31:0 | Tests whether the PIO design FPGA block RAM address bus interface is accurately connected by performing a walking ones address test starting at the I/O or memory address pointed to by the input bar_index. For an exhaustive test, this task should be called four times, once for each block RAM used in the PIO design. Additionally, the nBytes input should specify the entire size of the individual block RAM. |
| TSK_MEM_TEST_DEVICE | bar_index nBytes | 2:0 31:0 | Tests the integrity of each bit of the PIO design FPGA block RAM by performing an increment/decrement test on all bits starting at the block RAM pointed to by the input bar_index with the range specified by input nBytes. For an exhaustive test, this task should be called four times, once for each block RAM used in the PIO design. Additionally, the nBytes input should specify the entire size of the individual block RAM. |
| TSK_RESET | Reset | 0 | Initiates PERSTn. Forces the PERSTn signal to assert the reset. Use TSK_RESET (1'b1) to assert the reset and TSK_RESET (1'b0) to release the reset signal. |
| TSK_MALFORMED | malformed _bits | 7:0 | Control bits for creating malformed TLPs: 0001: Generate Malformed TLP for I/O Requests and Configuration Requests called immediately after this task 0010: Generate Malformed Completion TLPs for Memory Read requests received at the Root Port |

Table 6-11:Expectation Tasks

| Name | Input(| s) | Output | Description |
|--------------------|--|---|---------------|--|
| TSK_EXPECT_CPLD | traffic_class td ep attr length completer_id completer_status bcm byte_count requester_id tag address_low | 2:0 - 1:0 10:0 15:0 2:0 - 11:0 15:0 7:0 6:0 | Expect status | Waits for a Completion with Data TLP that matches traffic_class, td, ep, attr, length, and payload. Returns a 1 on successful completion; 0 otherwise. |
| TSK_EXPECT_CPL | traffic_class td ep attr completer_id completer_status bcm byte_count requester_id tag address_low | 2:0 - - 1:0 15:0 2:0 - 11:0 15:0 7:0 6:0 | Expect status | Waits for a Completion without Data TLP that matches traffic_class, td, ep, attr, and length. Returns a 1 on successful completion; 0 otherwise. |
| TSK_EXPECT_MEMRD | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - - 1:0 10:0 15:0 7:0 3:0 3:0 29:0 | Expect status | Waits for a 32-bit Address Memory Read TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |
| TSK_EXPECT_MEMRD64 | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - 1:0 10:0 15:0 7:0 3:0 3:0 3:0 61:0 | Expect status | Waits for a 64-bit Address Memory Read TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |

Table 6-11: Expectation Tasks (Cont'd)

| Name | Input | (s) | Output | Description |
|--------------------|--|--|---------------|---|
| TSK_EXPECT_MEMWR | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - 1:0 10:0 15:0 7:0 3:0 3:0 29:0 | Expect status | Waits for a 32-bit Address Memory Write TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |
| TSK_EXPECT_MEMWR64 | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - - 1:0 10:0 15:0 7:0 3:0 3:0 3:0 61:0 | Expect status | Waits for a 64-bit Address Memory Write TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |
| TSK_EXPECT_IOWR | td ep requester_id tag first_dw_be address data | - 15:0 7:0 3:0 31:0 31:0 | Expect status | Waits for an I/O Write TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |

Endpoint Model Test Bench for Root Port

The Endpoint model test bench for the Virtex-7 FPGAs Integrated Block for PCI Express in Root Port configuration is a simple example test bench that connects the Configurator example design and the PCI Express Endpoint model allowing the two to operate like two devices in a physical system. As the Configurator example design consists of logic that initializes itself and generates and consumes bus traffic, the example test bench only implements logic to monitor the operation of the system and terminate the simulation.

The Endpoint model test bench consists of:

- Verilog or VHDL source code for all Endpoint model components
- PIO slave design

Figure 6-7, page 209 illustrates the Endpoint model coupled with the Configurator example design.

Architecture

The Endpoint model consists of these blocks:

- PCI Express Endpoint (Gen3 Integrated Block for PCIe in Endpoint configuration) model.
- PIO slave design, consisting of:
 - PIO_RX_ENGINE
 - PIO_TX_ENGINE
 - PIO_EP_MEM
 - PIO_TO_CTRL

The PIO_RX_ENGINE and PIO_TX_ENGINE blocks interface with the ep block for reception and transmission of TLPs from/to the Root Port Design Under Test (DUT). The Root Port DUT consists of the Integrated Block for PCI Express configured as a Root Port and the Configurator Example Design, which consists of a Configurator block and a PIO Master design, or customer design.

The PIO slave design is described in detail in Programmed Input/Output: Endpoint Example Design, page 199.

Simulating the Design

A simulation script file is provided with the model to facilitate simulation with the Mentor Graphics ModelSim simulator:

• simulate_mti.do

The example simulation script files are located in this directory:

<project_dir>/<component_name>/simulation/functional

Instructions for simulating the Configurator example design with the Endpoint model are provided in Simulating the Example Design, page 225.

Note: For Cadence IES users, the work construct must be manually inserted into the cds.lib file:

DEFINE WORK WORK.

Scaled Simulation Timeouts

The simulation model of the Gen3 Integrated Block for PCIe uses scaled down times during link training to allow for the link to train in a reasonable amount of time during simulation.

According to the *PCI Express Specification, rev. 3.0*, there are various timeouts associated with the link training and status state machine (LTSSM) states. The Gen3 Integrated Block for PCIe scales these timeouts by a factor of 256 in simulation, except in the Recovery Speed_1 LTSSM state, where the timeouts are not scaled.

Waveform Dumping

Table 6-6 describes the available simulator waveform dump file format, which is provided in the simulator native file format.

Table 6-12: Simulator Dump File Format

| Simulator | Dump File Format |
|-----------|------------------|
| ModelSim | .vcd |

The Endpoint model test bench provides a mechanism for outputting the simulation waveform to file by specifying the +dump_all command line parameter to the simulator.

For example, the script file simulate_ncsim.sh (used to start the Cadence IES simulator) can indicate to the Endpoint model that the waveform should be saved to a file using this command line:

ncsim work.boardx01 +dump_all

Output Logging

The test bench outputs messages, captured in the simulation log, indicating the time at which these occur:

- user_reset deasserted
- user_lnk_up asserted
- cfg_done asserted by the Configurator
- pio_test_finished asserted by the PIO Master
- Simulation Timeout (if pio_test_finished or pio_test_failed never asserted)

PIPE MODE Simulation

The PIPE Simulation mode allows you to run the simulations without GT block, which speeds up simulations.

To run the simulations using the PIPE interface to speed up the simulation, generate the core using the **Enable PIPE simulation** option, as shown on the Basic page of the Customize IP dialog box described in Customizing the Core using the Vivado IP Catalog. With this option, the PIPE interface of the core top module in the PCIe example design is connected to PIPE interface of the model.



IMPORTANT: A new file $pcie3_7x_v1_3_gt_top_pipe.v$ is created in the simulation directory, and the file replaces the GT block for PIPE mode simulation.

To run simulations using GT block with the same core, define ENABLE_GT during run time so that the original GT block is instantiated in the core top module and simulations are run using the GT block. Comments are included in the simulation scripts to define which parameters need to be passed to run the simulations using GT block.



TIP: Implementation is always run with the GT block. The PIPE mode is only for simulation.

Implementation

Implementing the Example Design

To run the implementation on the generated core, go to the XCI file, right-click, and select open IP example design. A new Vivado tool window opens with the project name "example_project" within the project directory. In this new window, select the Run Synthesis and Run Implementation buttons and generate a bitstream either in sequence or any at a time. Selecting the Generate Bitstream button runs all steps: synthesis, implementation, and then bitstream. Selecting the Implementation button runs synthesis first and then implementation.

| Flow Navigator « Project Manager pried 7x v1 2.0 x | <u>File Edit Flow Tools Window La</u> | yout ⊻iew <u>H</u> elp | | Q= Search commands |
|---|---|--|--|---|
| Sources | 🯄 😂 🖩 🛤 🕫 🎼 🐘 🗙 ≽ 🕨 | 🐮 🚳 💥 ∑ 👒 😬 Default Layout 💿 🔭 🍋 | | Ready |
| Project Manager Project Sattings Proje | Flow Navigator | Project Manager - pcie3_7x_v1_2_0 | | × |
| Project Manager P | 🔍 🖾 🖶 | Sources _ D P × | Σ Project Summary \times | ㅁ ə × |
| | Project Manager | | | Dessages |
| Simulation Simulation Sattings Run Simulation RTL Analysis Synthesis Synthesis<td> 6 Project Settings 6 Add Sources 9 IP Catalog </td><td>. wilinx_pcie_3_0_7vx_ep (xilinx_pcie_3_0_7vx_ep.v)</td><td>Project name: pcie3_/X_V1_2_0 Product family: Virtex-7 Project part: <u>xc7vx690tffg1157-3</u></td><td>0 critical warnings</td> | 6 Project Settings 6 Add Sources 9 IP Catalog | . wilinx_pcie_3_0_7vx_ep (xilinx_pcie_3_0_7vx_ep.v) | Project name: pcie3_/X_V1_2_0 Product family: Virtex-7 Project part: <u>xc7vx690tffg1157-3</u> | 0 critical warnings |
| A RTL makysis A RTL makysis A RTL makysis A RTL makysis | 🍪 Simulation Settings 🍭 Run Simulation | Hierarchy IP Sources Libraries Compile Order | Part: xc7vx690tffg1157-3 Strategy: <u>Vivado Synthesis Defaults</u> | Part: xc7vx690tffg1157-3 Strategy: <u>Vivado Implementation Defaults</u> |
| Implementation Settings Name Part Constraints Strategy Host Statu Progress Stat Program and Debug Bitstream Settings Bitstream Settings Started 0% Owner Owner Owner Owner Owner Owner Owner Owner Start Program and Debug Bitstream Settings Generate Bitstream Start Owner Owner </td <td> RTL Analysis Open Elaborated Design Synthesis Synthesis Settings Run Synthesis </td> <td></td> <td></td> <td></td> | RTL Analysis Open Elaborated Design Synthesis Synthesis Settings Run Synthesis | | | |
| | Implementation Settings Run Implementation Implementation | Name Part Constraint ♥ ▷ synth_1 xc7vx690tffg1157-3 constraint xc7vx690tffg1157-3 constraint | s_1 Vivado Synthesis Defaults (Vivado Synthesis 2012) | Host Statu Progress Start Not started 0% 2012) Not started 0% |
| | | 🔲 Tcl Console 🔎 Messages 🔍 Log 🕒 Reports 🖏 Des | sign Runs | |

Figure 6-8: **Example Project**

Simulation

Simulating the Example Design

The example design provides a quick way to simulate and observe the behavior of the core.

Endpoint Configuration

The simulation environment provided with the LogiCORE[™] Gen3 Integrated Block for PCIe IP core in Endpoint configuration performs simple memory access tests on the PIO example design. Transactions are generated by the Root Port Model and responded to by the PIO example design.

- PCI Express Transaction Layer Packets (TLPs) are generated by the test bench transmit User Application (pci_exp_usrapp_tx). As it transmits TLPs, it also generates a log file, tx.dat.
- PCI Express TLPs are received by the test bench receive User Application (pci_exp_usrapp_rx). As the User Application receives the TLPs, it generates a log file, rx.dat.

For more information about the test bench, see Root Port Model Test Bench for Endpoint, page 208.

Setting Up for Simulation

To run the gate-level simulation, the Xilinx Simulation Libraries must be compiled for the user system. See the Compiling Xilinx Simulation Libraries (COMPXLIB) in the Xilinx ISE Synthesis and Verification Design Guide and the Xilinx ISE Software Manuals and Help. Documents can be downloaded from www.xilinx.com/support/software_manuals.htm.

Simulator Requirements

Virtex-7 FPGA designs require a Verilog LRM-IEEE 1364-2005 encryption-compliant simulator. This core supports this simulator:

- Mentor Graphics ModelSim
- Cadence IES
- Synopsys VCS
- Vivado Simulator

Running the Simulation

The simulation scripts provided with the example design support pre-implementation (RTL) simulation. The existing test bench can be used to simulate with a post-implementation version of the example design.

The pre-implementation simulation consists of these components:

- Verilog model of the test bench
- Verilog RTL example design
- The Verilog model of the Virtex-7 FPGA Gen3 Integrated Block for PCI Express
- 1. To run the simulation, go to this directory:

<project_dir>/<component_name>/simulation/functional

- 2. Launch the simulator and run the script that corresponds to the user simulation tool:
 - **ModelSim** > do simulate_mti.do
 - VCS > ./simulate_vcs.sh
 - IUS > ./simulate_ncsim.sh



SECTION III: ISE DESIGN SUITE

Customizing and Generating the Core

Constraining the Core

Example Design and Model Test Bench for Endpoint Configuration

Example Design and Model Test Bench for Root Port Configuration



Customizing and Generating the Core

This chapter includes information on using the ISE® Design Suite to customize and generate the core.

GUI

The LogiCORE[™] IP Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® core is a fully configurable and highly customizable solution. The Gen3 Integrated Block for PCIe is customized using the CORE Generator[™] tool.

Note: The screen captures in this chapter are conceptual representatives of their subjects and provide general information only. For the latest information, see the CORE Generator tool.

Customizing the Core using the CORE Generator Tool

The CORE Generator tool GUI for the Gen3 Integrated Block for PCIe consists of these screens:

- Page 1: Basic Parameter Settings
- Page 2: Interface Settings
- Pages 3 and 4: Identity Settings (PF0 and PF1)
- Pages 5 and 6: Base Address Registers (PF0 and PF1)
- Pages 7 and 8: SRIOV Config (PF0 and PF1)
- Pages 9 and 10: SRIOV Base Address Registers (PF0 and PF1)
- Page 11: Interrupt Capabilities (all PFs and VFs)
- Page 12: Power Management Registers
- Pages 13 and 14: PCIe Extended Capabilities

Basic Parameter Settings

The initial customization screen shown in Figure 7-1 is used to define the basic parameters for the core, including the component name, reference clock frequency, and silicon type.

www.xilinx.com

| | en3 Integrated Bl | lock for PCI Express | | | |
|--|---|--|-------------------------------|--|-------------------------|
| cuments | | | | | |
| | Vir | tex-7 FPG | iA Gen3 | | |
| di CORE | Integ | rated Blo | ck for PCI | | |
| | | Expres | | | illing and the stars in |
| | | 5 , 7 | | | xilinx.com:ip:pcie3_7x |
| | Name pcie3_ | 7x_v1_1 | | | |
| PCIe Device | | | | | |
| | | | ction of the Device / Port Ty | De la companya de la | |
| Device / Po | rt Type PCI Ex | opress Endpoint device | • | | |
| Pinout Selec | ction | | | | |
| -Xilinx Dev | relopment Boa | irds | | | |
| Generate | Xilinx Develop | oment Board specific U | JCF | | |
| Xilinx Dev | velopment Boa | ard None 💌 | | | |
| PCIe Block | Location Sele | ection | | | |
| Selects fr | om available F | PCIe Block locations for | r a part-package combinatio | n which determines Pinout | |
| | k Location X0 | | a part-package combinatio | in which determines findul. | |
| | a no contert frie | | | | |
| Generate A | dditional PCIe | Constraints | | | |
| Enables ge | neration of add | ditional constraints files | 5 | | |
| | ocation : | VOX N | □ xoyı | □ X0Y2 | |
| PCIe Block I | Lucation . | | | | |
| | | z v | | | |
| Reference (| Clock Frequen | | ction of the reference clock | frequency | |
| Reference (The Integra | Clock Frequent | °CI Express allows sele | ction of the reference clock | frequency | |
| Reference (The Integra Frequency | Clock Frequent ted Block for F (MHz) 100 MH | °CI Express allows sele | ction of the reference clock | frequency | |
| Reference (The Integra Frequency Link Status | Clock Frequent ted Block for F (MHz) 100 MH Register | PCI Express allows sele | | | |
| Reference (The Integra Frequency Link Status Selects whe | Clock Frequence ted Block for F (MHz) 100 MH Register ether the device | PCI Express allows sele | | frequency ynchronous) or generated via an onb | pard PLL (Asynchronous) |
| Reference (The Integra Frequency Link Status Selects whe | Clock Frequent ted Block for F (MHz) 100 MH Register | PCI Express allows sele | | | pard PLL (Asynchronous) |
| Reference (The Integra Frequency Link Status Selects whe | Clock Frequence ted Block for F (MHz) 100 MH Register ether the device | PCI Express allows sele | | | Dard PLL (Asynchronous) |
| Reference (The Integra Frequency Link Status Selects whe Enable S | Clock Frequence ted Block for F (MHz) 100 MH Register ether the device flot Clock Conf | PCI Express allows sele | | | pard PLL (Asynchronous) |
| Reference (The Integra Frequency Link Status Selects whe Enable S Silicon | Clock Frequence ted Block for F (MHz) 100 MH Register ether the device flot Clock Conf | Cl Express allows sele tz ce reference clock is p iguration | | | pard PLL (Asynchronous) |
| Reference (The Integra Frequency Link Status Selects whe Enable S Silicon | Clock Frequence ted Block for F (MHz) 100 MH Register ether the device flot Clock Conf | Cl Express allows sele tz ce reference clock is p iguration | | | Dard PLL (Asynchronous) |
| Reference (The Integra Frequency Link Status Selects whe Enable S Silicon | Clock Frequence ted Block for F (MHz) 100 MH Register ether the device flot Clock Conf | Cl Express allows sele tz ce reference clock is p iguration | | | Dard PLL (Asynchronous) |

Figure 7-1: **Page 1: Integrated Block for PCI Express Parameters**

Component Name

Base name of the output files generated for the core. The name must begin with a letter and can be composed of these characters: a to z, 0 to 9, and "_."

PCIe Device / Port Type

Indicates the PCI Express logical device type.

Xilinx Development Board

Selects the Xilinx Development Board to enable the generation of Xilinx Development Board specific constraints files.

Integrated Block for PCIe Location

Selects from the available Integrated Blocks to enable generation of location-specific constraint files and pinouts. This selection is used in the default example design scripts.

This option is not available if a Xilinx Development Board is selected.

Additional Constraints

Allows generation of additional constraints files for other blocks available in the device.

This option is not available if a Xilinx Development Board is selected.

Reference Clock Frequency

Selects the frequency of the reference clock provided on sys_clk. For information about clocking the Gen3 Integrated Block for PCIe, see Clocking, page 74.

Slot Clock Configuration

Enables the Slot Clock Configuration bit in the Link Status register. Selecting this option means the link is synchronously clocked. See Clocking, page 74 for more information on clocking options.

Silicon Type

Selects the silicon type.

Interface Settings

The Interface Settings page is shown in Figure 7-2.

| | rtex-7 FPGA grated Blocl | | |
|--|------------------------------|--|--|
| | Express | | xilinx.com:ip:pcie3_7x |
| Number of Lanes | | | |
| | | n initial lane width be selected. Wider lane width con y the lane width that is necessary for the design. | res can train down to smaller lane widths if |
| Lane Width X1 | | , are take main that the cost of the action. | |
| Link Speed | | | AXI-ST Interface Frequency |
| | PCI Express allows selection | in of the Maximum Link Speed supported by the dev | |
| © 2.5 GT/s | C 5.0 GT/s | C 8.0 GT/s | |
| Interface Width | | | 62.5MHz |
| | PCI Express allows selection | n of Intorface Width | |
| 64-bit | C 128-bit | C 256-bit | |
| Receive Completion Stra | ddle Enable | Alignment Mode | |
| Enable AXISTEN Fram | | | Address Aligned |
| | | | , and ess mighted |
| Physical Functions | | | |
| Enable Physical Funct | ion 0 | Enable Physical Function 1 | |
| Device Capabilities Regi | ster – PF | | |
| PF0 Max Payload Size 5 | 12 bytes 💌 PF1 Max Pay | rload Size 512 bytes 💌 | |
| Extended Tag Field | | | |
| Function Level Reset | | | |
| Enables Function Leve | l Reset Capability | | |
| □ Function Level Rese | t | | |
| | | | |
| SRIOV Capability | /irtualization | | |
| SRIOV Capability | | | |
| SRIOV Capability Enables Single Root I/O N 「 SRIOV Capability Enal | | | |

Figure 7-2: Page 2: Interface Settings

Number of Lanes

The Gen3 Integrated Block for PCIe requires the selection of the initial lane width. Figure 7-1 defines the available widths and associated generated core. Wider lane width cores can train down to smaller lane widths if attached to a smaller lane-width device. See Link Training: 2-Lane, 4-Lane, and 8-Lane Components, page 157 for more information.

| Lane Width | Product Generated |
|------------|--|
| xl | 1-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| x2 | 2-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| x4 | 4-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| x8 | 8-Lane Virtex-7 FPGA Gen3 Integrated Block for PCI Express |

Table 7-1: Lane Width and Product Generated

Link Speed

The Gen3 Integrated Block for PCIe allows the selection of Maximum Link Speed supported by the device. Table 7-2 defines the lane widths and link speeds supported by the device. Higher link speed cores are capable of training to a lower link speed if connected to a lower link speed capable device.

Table 7-2: Lane Width and Link Speed

| Lane Width | Link Speed |
|------------|--------------------------|
| x1 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |
| x2 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |
| x4 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |
| x8 | 2.5 Gb/s, 5 Gb/s, 8 Gb/s |

Interface Width

The Gen3 Integrated Block for PCIe allows the selection of Interface Width, as defined in Table 7-3. The default interface width set in the CORE Generator tool GUI is the lowest possible interface width.

| Table 7-3: | Lane Width, Link Speed, and Interface Width |
|------------|---|
|------------|---|

| Lane Width | Link Speed (Gb/s) | Interface Width (Bits) |
|------------|----------------------|------------------------|
| x1 | 2.5, 5.0, 8.0 | 64 |
| x2 | 2.5, 5.0 | 64 |
| x2 | 8.0 | 64, 128 |
| x4 | 2.5 | 64 |
| x4 | 5.0 | 64, 128 |
| x4 | 8.0 | 128, 256 |
| x8 | 2.5 | 64, 128 |
| x8 | 5.0 | 128 256 |
| x8 | 8.0 | 256 |

Requestor Completion Straddle

The Gen3 Integrated Block for PCIe provides an option to straddle packets on the Requestor Completion interface when the interface width is 256 bits. See Straddle Option for 256-Bit Interface, page 149.

Alignment Mode

When a payload is present, there are two options for aligning the first byte of the payload with respect to the datapath. See Data Alignment Options, page 81.

Physical Function 1 Enable

The Gen3 Integrated Block for PCIe implements an additional Physical Function.

MPS

This field indicates the maximum payload size that the device or function can support for TLPs. This is the value advertised to the system in the Device Capabilities Register.

Extended Tag

This field indicates the maximum supported size of the Tag field as a Requester. Options are 8-bit Tag field support (when selected) or 5-bit Tag field support (when deselected).

SRIOV Enable

The Integrated Block implements the Single Root Port I/O Virtualization PCIe extended capability. When this capability is enabled, the SRIOV capability is implemented for both PF0 and PF1 (if selected).

Number of Functions

The Integrated Block implements up to six Virtual Functions that are associated to either PF0 or PF1 (if enabled).

Function Level Reset

The Integrated Block enables you to reset a specific device function. This mechanism is only applicable to Endpoint configurations.

Identity Settings (PF0 and PF1)

The Identity Settings pages are shown in Figure 7-3 and Figure 7-4. These settings customize the IP initial values, class code, and Cardbus CIS pointer information. The page for Physical Function 1 (PF1) is only displayed when PF1 is enabled.

| ic <u>èr</u> e II | | c-7 FPGA Gen3 ted Block for PCI | | |
|---------------------|----------------|---------------------------------------|----------|-----------------------|
| | _ | Express | | xilinx.com:ip:pcie3_7 |
| 9F0 - ID Initial Va | lues | | | |
| Vendor ID | 10EE | Range: 0000FFFF | | |
| Device ID | 7011 | Range: 0000FFFF | | |
| Revision ID | 00 | Range: 00FF | | |
| Subsystem Vend | dor ID 10EE | Range: 0000FFFF | | |
| Subsystem ID | 0007 | Range: 0000FFFF | | |
| PF0 - Class Code | 1 | | | |
| Base Class | 05 | Range: 00FF | | |
| Sub-Class | 80 | Range: 00FF | | |
| Interface | 00 | - Range: 00FF | | |
| Class Code | 058000 | (Hex) | | |
| Class Code Loo | okup Assistant | | | |
| Must enter valu | | | | |
| Base Class | 1.00 | nple communication controllers | × | |
| Base Class | 07 | | in a | |
| | 1 | neric XT compatible serial controller | <u> </u> | |
| Sub-Class | 00 | | | |
| intenace | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

Figure 7-3: Page 3: Identity Settings (PF0)

| | | Express | | xilinx.com:ip:pcie3_7 |
|---------------------|-----------------|---------------------------------------|---|-----------------------|
| F1 - ID Initial Val | ues | | | |
| /endor ID | 10EE | | | |
| Device ID | 7011 | Range: 0000FFFF | | |
| Revision ID | 00 | Range: 00FF | | |
| Subsystem Vend | or ID 10EE | | | |
| Subsystem ID | 0007 | Range: 0000FFFF | | |
| F1 - Class Code | | | | |
| ase Class | 05 | Range: 00FF | | |
| Sub-Class | 80 | Range: 00FF | | |
| nterface | 00 | Range: 00FF | | |
| lass Code | 058000 | (Hex) | | |
| Class Code Loo | kup Assistant – | | | |
| Must enter valu | ies above. | | | |
| Base Class | Sin | nple communication controllers | - | |
| Base Class | 07 | h | | |
| Sub-Class/Inter | face Value Ge | neric XT compatible serial controller | • | |
| Sub-Class | 00 | h | | |
| Interface | 00 | h | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

Figure 7-4: Page 4: Identity Settings (PF1)

ID Initial Values

- **Vendor ID:** Identifies the manufacturer of the device or application. Valid identifiers are assigned by the PCI Special Interest Group to guarantee that each identifier is unique. The default value, 10EEh, is the Vendor ID for Xilinx. Enter a vendor identification number here. FFFFh is reserved.
- **Device ID:** A unique identifier for the application; the default value, which depends on the configuration selected, is 70<*link speed*><*link width*>h. This field can be any value; change this value for the application.
- **Revision ID:** Indicates the revision of the device or application; an extension of the Device ID. The default value is 00h; enter values appropriate for the application.

- **Subsystem Vendor ID:** Further qualifies the manufacturer of the device or application. Enter a Subsystem Vendor ID here; the default value is 10EEh. Typically, this value is the same as Vendor ID. Setting the value to 0000h can cause compliance testing issues.
- **Subsystem ID:** Further qualifies the manufacturer of the device or application. This value is typically the same as the Device ID; the default value depends on the lane width and link speed selected. Setting the value to 0000h can cause compliance testing issues.

Class Code

The Class Code identifies the general function of a device, and is divided into three byte-size fields:

- **Base Class:** Broadly identifies the type of function performed by the device.
- **Sub-Class:** More specifically identifies the device function.
- **Interface:** Defines a specific register-level programming interface, if any, allowing device-independent software to interface with the device.

Class code encoding can be found at <u>www.pcisig.com</u>.

Class Code Look-up Assistant

The Class Code Look-up Assistant provides the Base Class, Sub-Class and Interface values for a selected general function of a device. This Look-up Assistant tool only displays the three values for a selected function. The user must enter the values in Class Code for these values to be translated into device settings.

Base Address Registers (PF0 and PF1)

The Base Address Registers (BARs) screens shown in Figure 7-5 and Figure 7-6 set the base address register space for the Endpoint configuration. Each BAR (0 through 5) configures the BAR Aperture Size and Control attributes of the Physical Function, as described in Table D-1.

| | Virtex-7 FPGA | Gen3 | |
|-------------------------|--|--------------------|--|
| CRE | Integrated Block | for PCI | |
| | Express | | xilinx.com.ip:pcie3_7. |
| PEO Base Ar | ddress Registers | | |
| Base Addre system me | ess Registers (BARs) serve two purposes. | nes what addresses | as a mechanism for the device to request blocks of address space in the to assign to the device. the Base Address Registers are programmed with ofing. |
| BAR 0 Opt | tions | | BAR 1 Options |
| 🔽 Bar0 | Type Memory 💌 🗆 64 bit | □ Prefetchable | □ Barl Type N/A 🔄 □ 64 bit □ Prefetchable |
| | Size 2 💌 Kilobytes 💌 | | Size 2 💌 Kilobytes 💌 |
| , | Value FFFFF800 (Hex) | | Value 00000000 (Hex) |
| BAR 2 Opt | tions | | BAR 3 Options |
| □ Bar2 | Type N/A 🗾 🗖 64 bit | □ Prefetchable | □ Bar3 Type N/A □ □ 64 bit □ Prefetchable |
| | Size 2 💌 Kilobytes 💌 | | Size 2 👻 Kilobytes 💌 |
| , | Value 00000000 (Hex) | | Value 00000000 (Hex) |
| BAR 4 Opt | tions | | BAR 5 Options |
| □ Bar4 | Type N/A 🔽 🗖 64 bit | | □ Bar5 Type N/A □ □ Prefetchable |
| | Size 2 🝸 Kilobytes 🝸 | | Size 2 Xilobytes X |
| | Value 00000000 (Hex) | | Value 00000000 (Hex) |
| Expansion | n ROM Base Address Register | | |
| 15. | sion Rom Size 2 Skiloby | tes 🔻 | |
| i Expans | Value 00000000 (Hex) | | |
| | Value 0000000 (Hex) | | |

Figure 7-5: Page 5: Base Address Register (PF0)

| RE | | 7 FPGA ed Block | | | | | |
|---------|--|--------------------|---------------------|-----------------|--------------|----------------|---------------------|
| | | xpress | | | | | xilinx.com:ip:pcie3 |
| Base A | Address Registers | | | | | | |
| tem m | ress Registers (BARs) ser lemory map. After the Bl s and the device uses thi | IOS or OS determ | ines what addresses | to assign to th | | | |
| AR 0 Op | | | | BAR 1 Opti | ions | | |
| Bar0 | Type Memory 💌 | 🖵 64 bit | Prefetchable | □ Barl T | Type N/A | ▼ 64 bit | □ Prefetchable |
| | Size 2 | Kilobytes 💌 | | S | Size 2 | 👻 Kilobytes 👻 | 1 |
| | Value FFFFF800 | (Hex) | | V | /alue 000000 | 00 (Hex) | |
| AR 2 Op | ptions | | | BAR 3 Opti | ions | | |
| Bar2 | Type N/A 💌 | 🗖 64 bit | □ Prefetchable | □ Bar3 T | Type N/A | 👻 🗖 64 bit | □ Prefetchable |
| | Size 2 | Kilobytes 👻 | | S | Size 2 | 👻 Kilobytes 👻 | l |
| | Value 00000000 | (Hex) | | V | /alue 000000 | 00 (Hex) | |
| AR 4 Op | ptions | | | BAR 5 Opti | ions | | |
| Bar4 | Type N/A 💌 | 🗖 64 bit | □ Prefetchable | □ Bar5 T | ype N/A | ▼ Prefetchable | |
| | Size 2 | Kilobytes 👻 | | S | Size 2 | 👻 Kilobytes 👻 | l – |
| | Value 00000000 | (Hex) | | v | /alue 000000 | 00 (Hex) | |
| pansio | on ROM Base Address Re | gister | | | | | |
| Expar | nsion Rom Size 2 | 👻 Kiloby | rtes 🗾 | | | | |
| | Value 0000 | 0000 (Hex) | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |

Figure 7-6: Page 6: Base Address Register (PF1)

Base Address Register Overview

The Gen3 Integrated Block for PCIe in Endpoint configuration supports up to six 32-bit BARs or three 64-bit BARs, and the Expansion ROM BAR. The Gen3 Integrated Block for PCIe in Root Port configuration supports up to two 32-bit BARs or one 64-bit BAR, and the Expansion ROM BAR.

BARs can be one of two sizes:

- **32-bit BARs:** The address space can be as small as 16 bytes or as large as 2 gigabytes. Used for Memory to I/O.
- **64-bit BARs:** The address space can be as small as 128 bytes or as large as 8 exabytes. Used for Memory only.

All BAR registers share these options:

- **Checkbox:** Click the checkbox to enable the BAR; deselect the checkbox to disable the BAR.
- **Type:** BARs can either be I/O or Memory.
 - I/O: I/O BARs can only be 32-bit; the Prefetchable option does not apply to I/O BARs. I/O BARs are only enabled for the Legacy PCI Express Endpoint core.
 - Memory: Memory BARs can be either 64-bit or 32-bit and can be prefetchable.
 When a BAR is set as 64 bits, it uses the next BAR for the extended address space and makes the next BAR inaccessible to the user.
- **Size:** The available Size range depends on the PCIe® Device/Port Type and the Type of BAR selected. Table 7-4 lists the available BAR size ranges.

| PCle Device / Port Type | BAR Type | BAR Size Range |
|-----------------------------|---------------|-------------------------|
| PCI Express Endpoint | 32-bit Memory | 128 Bytes – 2 Gigabytes |
| | 64-bit Memory | 128 Bytes – 8 Exabytes |
| | 32-bit Memory | 16 Bytes – 2 Gigabytes |
| Legacy PCI Express Endpoint | 64-bit Memory | 16 Bytes – 8 Exabytes |
| | I/O | 16 Bytes – 2 Gigabytes |

Table 7-4: BAR Size Ranges for Device Configuration

- **Prefetchable:** Identifies the ability of the memory space to be prefetched.
- Value: The value assigned to the BAR based on the current selections.

For more information about managing the Base Address Register settings, see Managing Base Address Register Settings.

Expansion ROM Base Address Register

If selected, the Expansion ROM is activated and can be a value from 2 KB to 4 GB. According to the *PCI 3.0 Local Bus Specification*, the maximum size for the Expansion ROM BAR should be no larger than 16 MB. Selecting an address space larger than 16 MB might result in a non-compliant core.

Managing Base Address Register Settings

Memory, I/O, Type, and Prefetchable settings are handled by setting the appropriate GUI settings for the desired base address register.

Memory or I/O settings indicate whether the address space is defined as memory or I/O. The base address register only responds to commands that access the specified address space. Generally, memory spaces less than 4 KB in size should be avoided. The minimum I/ O space allowed is 16 bytes; use of I/O space should be avoided in all new designs. Prefetchability is the ability of memory space to be prefetched. A memory space is prefetchable if there are no side effects on reads (that is, data is not destroyed by reading, as from a RAM). Byte write operations can be merged into a single double word write, when applicable.

When configuring the core as an Endpoint for PCIe (non-Legacy), 64-bit addressing must be supported for all BARs (except BAR5) that have the prefetchable bit set. 32-bit addressing is permitted for all BARs that do not have the prefetchable bit set. The prefetchable bit related requirement does not apply to a Legacy Endpoint. The minimum memory address range supported by a BAR is 128 bytes for a PCI Express Endpoint and 16 bytes for a Legacy PCI Express Endpoint.

Disabling Unused Resources

For best results, disable unused base address registers to conserve system resources. A base address register is disabled by deselecting unused BARs in the GUI.

SRIOV Config (PF0 and PF1)

The SRIOV Config screen is shown in Figure 7-7.

| SRIOV Capability Version |
|--|
| e: 0F SRIOV Cap Version 0 Range: 0F |
| Virtual Function Select |
| Number of PF1 VF's |
| SRIOV Functional Dependency Link |
| FFFF SRIOV PF Dependency Link 0000 Range: 0000FFF |
| SRIOV First VF Offset |
| FFFF SRIOV PF Dependency Link 0000 Range: 0000FFFF |
| SRIOV VF Device ID |
| FFFF SRIOV VF Device ID 0000 Range: 0000FFFF |
| SRIOV Supported Page Size |
| SRIOV Supported Page Size 4KB |
| |
| |
| |
| |
| |
| |
| |

Figure 7-7: Page 7: SRIOV Config (PF0 and PF1)

SRIOV Capability Version

Indicates the 4-bit SRIOV Capability Version for the Physical Function.

SRIOV Function Select

Indicates the number of Virtual Functions associated to the Physical Function. A maximum of six Virtual Functions are available to PF0 and PF1.

SRIOV Functional Dependency Link

Indicates the SRIOV Functional Dependency Link for the Physical Function. The programming model for a device can have vendor-specific dependencies between sets of Functions. The Function Dependency Link field is used to describe these dependencies.

SRIOV First VF Offset

Indicates the offset of the first Virtual Function for the Physical Function.

SRIOV VF Device ID

Indicates the 16-bit Device ID for all Virtual Functions associated with the Physical Function.

SRIOV Supported Page Size

Indicates the page size supported by the Physical Function. This Physical Function supports a page size of 2n+12, if bit *n* of the 32-bit register is set.

SRIOV Base Address Registers (PF0 and PF1)

The SRIOV Base Address Registers (BARs) screens shown in Figure 7-8 and Figure 7-9 set the base address register space for the Endpoint configuration. Each BAR (0 through 5) configures the SRIOV BAR Aperture Size and SRIOV Control attributes as described in Table D-1.

| OPF | | | x-7 FPG | | | |
|-------------|---------|---------------|-------------------------------|----------------|--|---------------------------|
| CRE | In | tegra | | ck for PCI | | |
| | | | Expres | S | | xilinx.com:ip:pcie3_2 |
| PFO SRIOV (| Base A | ddress Regist | ers | | | |
| system me | emory r | map. After th | e BIOS or OS dete | | a mechanism for the device to request assign to the device. the Base Address ng. | |
| PF0 SRIOV | | | | | PF0 SRIOV BAR 1 Options | |
| 🔽 Bar0 | туре [| Memory | ▼ 🗆 64 bit | □ Prefetchable | □ Barl Type N/A 🔄 □ 64 | bit F Prefetchable |
| | Size [| 2 | Kilobytes | - | Size 2 💌 Kiloby | /tes 💌 |
| | Value | FFFFF800 | (Hex) | | Value 00000000 (Hex) | |
| PF0 SRIOV | V BAR 2 | 2 Options | | | PF0 SRIOV BAR 3 Options | |
| □ Bar2 | туре [| N/A | ▼ □ 64 bit | | □ Bar3 Type N/A 		 □ 64 | bit 🔽 Prefetchable |
| | Size [| 2 | Kilobytes |] | Size 2 Kilob | /tes 👻 |
| | Value | 00000000 | (Hex) | | Value 00000000 (Hex) | |
| PF0 SRIOV | V BAR 4 | Options | | | PF0 SRIOV BAR 5 Options | |
| □ Bar4 | Type [| N/A | ▼ Г 64 bit | | □ Bar5 Type N/A □ □ Pre | fetchable |
| | Size [| | - Kilobytes - |] | Size 2 Kiloby | rtes 💌 |
| , | Value | 00000000 | (Hex) | | Value 00000000 (Hex) | |
| | | | | | | |

Figure 7-8: Page 8: SRIOV Base Address Register (PF0)

| RE | 1 | 2023 | | | | Gen3 | | | | | |
|--------------------|-----------------|-----------|------------------|------------------------------|-------|----------------|--------------|-------|-----------|--|---|
| av. | II | iteg | | ea Bie Expre | | c for PCI | | | | | |
| | | | | Expre | 55 | | | | | | xilinx.com:ip:pcie3 |
| se Addr stern m | ess Re emory | map. Aft | ARs) s er the | erve two pur BIOS or OS d | eterm | | to assign to | | | e to request blocks of se Address Registers | address space in the are programmed with |
| 1 SRIC | V BAR | 0 Options | | | | | PF1 SRIC | V BAR | 1 Options | | |
| Bar0 | | Memory | | ☐ 64 bit | | | □ Barl | Туре | N/A | ✓ 64 bit | |
| | Size | 2 | * | Kilobytes | • | | | Size | | 👻 Kilobytes 👻 | |
| | Value | FFFFF80 | 0 | (Hex) | | | | Value | 00000000 | (Hex) | |
| 1 SRIC | V BAR | 2 Options | | | | | PF1 SRIC | V BAR | 3 Options | | |
| Bar2 | Туре | N/A | ¥ | □ 64 bit | | | □ Bar3 | Туре | N/A | <u>▼</u> 64 bit | |
| | Size | | * | Kilobytes | Ŧ | | | Size | | 👻 Kilobytes 👻 | |
| | Value | 0000000 | 00 | (Hex) | | | | Value | 00000000 | (Hex) | |
| 1 SRIC | V BAR | 4 Options | | | | | PF1 SRIC | V BAR | 5 Options | | |
| Bar4 | Туре | N/A | Ψ. | □ 64 bit | | □ Prefetchable | □ Bar5 | Туре | N/A | T Prefetchable | |
| | | 2 | | Kilobytes | Ŧ | | | Size | | 👻 Kilobytes 👻 | |
| | Value | 0000000 | 00 | (Hex) | | | | Value | 00000000 | (Hex) | |
| | | | | | | | | | | | |

Figure 7-9: Page 9: SRIOV Base Address Register (PF1)

SRIOV Base Address Register Overview

The Gen3 Integrated Block for PCIe in Endpoint configuration supports up to six 32-bit BARs or three 64-bit BARs. The Gen3 Integrated Block for PCIe in Root Port configuration supports up to two 32-bit BARs or one 64-bit BAR.

SRIOV BARs can be one of two sizes:

- **32-bit BARs:** The address space can be as small as 16 bytes or as large as 2 gigabytes. Used for Memory to I/O.
- **64-bit BARs:** The address space can be as small as 128 bytes or as large as 8 exabytes. Used for Memory only.

All SRIOV BAR registers share these options:

- **Checkbox:** Click the checkbox to enable the BAR; deselect the checkbox to disable the BAR.
- **Type:** SRIOV BARs can either be I/O or Memory.
 - I/O: I/O BARs can only be 32-bit; the Prefetchable option does not apply to I/O BARs. I/O BARs are only enabled for the Legacy PCI Express Endpoint core.
 - Memory: Memory BARs can be either 64-bit or 32-bit and can be prefetchable.
 When a BAR is set as 64 bits, it uses the next BAR for the extended address space and makes the next BAR inaccessible to the user.
- **Size:** The available Size range depends on the PCIe® Device/Port Type and the Type of BAR selected. Table 7-5 lists the available BAR size ranges.

| PCIe Device / Port Type | BAR Type | BAR Size Range |
|-----------------------------|---------------|-------------------------|
| PCI Express Endpoint | 32-bit Memory | 128 Bytes – 2 Gigabytes |
| | 64-bit Memory | 128 Bytes – 8 Exabytes |
| | 32-bit Memory | 16 Bytes – 2 Gigabytes |
| Legacy PCI Express Endpoint | 64-bit Memory | 16 Bytes – 8 Exabytes |
| | I/O | 16 Bytes – 2 Gigabytes |

Table 7-5: SRIOV BAR Size Ranges for Device Configuration

- **Prefetchable:** Identifies the ability of the memory space to be prefetched.
- Value: The value assigned to the BAR based on the current selections.

For more information about managing the SRIOV Base Address Register settings, see Managing Base Address Register Settings.

Managing SRIOV Base Address Register Settings

Memory, I/O, Type, and Prefetchable settings are handled by setting the appropriate GUI settings for the desired base address register.

Memory or I/O settings indicate whether the address space is defined as memory or I/O. The base address register only responds to commands that access the specified address space. Generally, memory spaces less than 4 KB in size should be avoided. The minimum I/ O space allowed is 16 bytes; use of I/O space should be avoided in all new designs.

Prefetchability is the ability of memory space to be prefetched. A memory space is prefetchable if there are no side effects on reads (that is, data is not destroyed by reading, as from a RAM). Byte write operations can be merged into a single double word write, when applicable.

When configuring the core as an Endpoint for PCIe (non-Legacy), 64-bit addressing must be supported for all SRIOV BARs (except BAR5) that have the prefetchable bit set. 32-bit

addressing is permitted for all SRIOV BARs that do not have the prefetchable bit set. The prefetchable bit related requirement does not apply to a Legacy Endpoint. The minimum memory address range supported by a BAR is 128 bytes for a PCI Express Endpoint and 16 bytes for a Legacy PCI Express Endpoint.

Disabling Unused Resources

For best results, disable unused base address registers to conserve system resources. A base address register is disabled by deselecting unused BARs in the GUI.

Interrupt Capabilities

The Interrupt Capabilities screen shown in Figure 7-10 sets the Legacy Interrupt Settings and MSI Capabilities for all applicable Physical and Virtual Functions. The MSI-X Capabilities screen shown in Figure 7-11 sets the MSI-X Capabilities.

| tex-7 FPGA Gen3 Integrated Block for PCI Express | | |
|---|--|-------------------------------|
| uments | | |
| Virtex-7 FPGA G | ien3 | |
| 👯 🕺 Integrated Block f | or PCI | |
| Express | | xilinx.com:ip:pcie3_7x:1 |
| Interrupt Capabilities | | |
| The PCI Express Base Specification 3.0 requires Legacy. | . MSI. and/or MSI-X interrupts capability to be implemented in a device. ustomize the capability list as needed for the application | |
| Legacy Interrupt Settings | 1 3 11 | |
| PF0 Interrupt PIN NONE | PF1 Interrupt PIN NONE | |
| MSI Capabilities | | |
| F PF0 Enable MSI Capability Structure | ✓ PF1 Enable MSI Capability Structure | |
| PF0 Multiple Message Capable 🛽 vector 💌 | PF1 Multiple Message Capable 1 vector 💌 | |
| VF0 Multiple Message Capable 🛽 vector 🛛 💌 | | |
| | VF1 Multiple Message Capable 1 vector 💌 | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | < Back Page 10 of 14 Next > Generate | e <u>C</u> ancel <u>H</u> elp |

Figure 7-10: Page 10: Interrupt Capabilities

| x Capabilities | xilinx.com.ip:pcie3 |
|--|---|
| F0 Enable MSIx Capability Structure MSIx Table Settings Table Size 000 Range: 0000.7FF Table Offset 0000000 Range: 0000000.1FFFFFFF BAR Indicator BAR 0 MSIx Pending Bit Array (PBA) Settings PBA Offset 0000000 Range: 00000000.1FFFFFFF PBA BAR Indicator BAR 0 | PF1 Enable MSIx Capability Structure MSIx Table Settings Table Size 000 Range: 000.7FF Table Offset 0000000 Range: 0000000.1FFFFFFF BAR Indicator BAR 0 PBA Offset 0000000 Range: 00000000 Range: 0000000.1FFFFFFF PBA Affset 00000000 Range: 00000000.1FFFFFFF PBA BAR Indicator BAR 0 |
| F0-VF0 MSIx Table Settings Table Size 000 Range: 0007FF Table Offset 00000000 Range: 000000001FFFFFFF BAR Indicator BAR 0 MSIx Pending Bit Array (PBA) Settings PBA Offset 0000000 Range: 000000001FFFFFFF PBA BAR Indicator BAR 0 | PF1-VF1 MSIx Table Settings Table Size 000 Range: 000.7FF Table Offset 0000000 Range: 0000000.1FFFFFF BAR Indicator BAR 0 MSIx Pending Bit Array (PBA) Settings PBA Offset 00000000 Range: 00000000.1FFFFFF PBA Affset 00000000 Range: 00000000.1FFFFFFF PBA Affset 00000000 Range: 00000000.1FFFFFFF |

Figure 7-11: Page 11: MSI-X Interrupt Capabilities

Legacy Interrupt Settings

- **Enable INTX**: Enables the ability of the PCI Express function to generate INTx interrupts.
- **Interrupt PIN**: Indicates the mapping for Legacy Interrupt messages. A setting of "None" indicates no Legacy Interrupts are used.

MSI Capabilities

• Enable MSI Capability Structure: Indicates that the MSI Capability structure exists.

Note: Although it is possible not to enable MSI or MSI-X, the result would be a non-compliant core. The *PCI Express Base Specification* requires that MSI, MSI-X, or both be enabled.

• **64 bit Address Capable**: Indicates that the function can send a 64-bit Message Address.

- **Multiple Message Capable**: Selects the number of MSI vectors to request from the Root Complex.
- **Per Vector Masking Capable**: Indicates that the function supports MSI per-vector Masking.

MSI-X Capabilities

• Enable MSIx Capability Structure: Indicates that the MSI-X Capability structure exists.

Note: This Capability Structure needs at least one Memory BAR to be configured. The user must maintain the MSI-X Table and Pending Bit Array in the user application.

- MSIx Table Settings: Defines the MSI-X Table Structure.
 - Table Size: Specifies the MSI-X Table Size.
 - *Table Offset*: Specifies the Offset from the Base Address Register that points to the Base of the MSI-X Table.
 - *BAR Indicator*: Indicates the Base Address Register in the Configuration Space that is used to map the function MSI-X Table, onto Memory Space. For a 64-bit Base Address Register, this indicates the lower DWORD.
- **MSIx Pending Bit Array (PBA) Settings**: Defines the MSI-X Pending Bit Array (PBA) Structure.
 - *PBA Offset*: Specifies the Offset from the Base Address Register that points to the Base of the MSI-X PBA.
 - *PBA BAR Indicator*: Indicates the Base Address Register in the Configuration Space that is used to map the function MSI-X PBA, onto Memory Space.

Power Management Registers

The Power Management Registers screen shown in Figure 7-12 includes settings for the Power Management Registers, power consumption, and power dissipation options. These settings apply to both Physical Functions, if PF1 is enabled.

| uments | | | | |
|------------------------|-----------------------|---------------------------|---------------------------|-------------------------|
| | irtex-7 FPG | 57 - 17 C C C C | | |
| icere Inte | egrated Bloc | k for PCI | | |
| | Express | 5 | | xilinx.com:ip:pcie3_ |
| Power Management R | egisters | | | |
| D1 Support | | | | |
| ME Support from: | | | | |
| | not | | | |
| Device Capabilities Re | egister 2 | | | |
| 32-bit AtomicOp Ci | - | | | |
| 64-bit AtomicOp Ci | | | | |
| 128-bit CAS Compl | | | | |
| TPH Completer Sup | | | | |
| OBFF Supported 00 N | | | | |
| · · · · | Salation Contraction | | | |
| BRAM Configuration O | | 9 | | |
| Performance Level | Posted Header/Data | Non-posted Header/Data | Completion Header/Data | Total BRAMS Required |
| | Credits | Credits | Credits | 18K/36K |
| Good | 0x20/0x | 0×20/0× | 0x20/0x | 8/2 |
| C Extreme | 0x20/0x | 0×20/0× | 0×20/0× | 4/6 |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

Figure 7-12: Page 12: Power Management Registers

- **D1 Support**: When selected, this option indicates that the function supports the D1 Power Management State. See section 3.2.3 of the *PCI Bus Power Management Interface Specification Revision 1.2.*
- **PME Support From**: When this option is selected, it indicates the power states in which the function can assert cfg_pm_wake. See section 3.2.3 of the *PCI Bus Power Management Interface Specification Revision 1.2.*
- **Device Capability Register 2 Settings**: Specifies options for AtomicOps and TPH Completer Support. See the Device Capability Register 2 description in Chapter 7 of the

PCI Express Base Specification for more information. These settings apply to both Physical Functions, if PF1 is enabled.

• **BRAM Configuration Options**: Can specify the number of receive block RAMs used for the solution. The table displays the number of receiver credits available for each packet type.

PCIe Extended Capabilities

The PCIe Extended Capabilities screens shown in Figure 7-13 and Figure 7-14 allow you to enable PCI Express Extended Capabilities. The Advanced Error Reporting Capability (offset 0x100h) is always enabled. The customization GUI sets up the link list based on the capabilities enabled. After enabling, you must configure the capability by setting the applicable attributes in the core top-level defined in Output Generation, page 254. See Appendix D, Attributes for parameters applicable to each capability.

| | Virtex-7 F | PGA Gen3 | |
|---|---|---|---|
| CRE | Integrated | Block for PCI | |
| | | ress | |
| | Exh | 1655 | xilinx.com:ip:pcie3_ |
| Cle Exter | nded Capabilities | | |
| Device S | Serial Number Capability | | |
| | rice Serial Number (DSN) Capa presented on the Device Seria | | t contains a unique Device Serial Number. This identifier |
| F Enabl | le DSN Capability(PF0) | ☐ Enable DSN | Capability(PF1) |
| Virtual C | hannel Capability | | |
| | | n optional PCIe Extended Capability. which wher default Virtual Channel (VC0). Checking this allow | n enabled. allows the port to support functionality beyond ws Traffic Class (TC) filtering to be supported |
| □ Enabl | le VC Capability(PF0) | | |
| DPA Cap | ability | | |
| | | ism to allocate power dynamically for devices v the PCI Express Base Specification Revision 3.0 | |
| Enabl | le DPA Capability(PF0) | ☐ Enable DPA | Capability(PF1) |
| | | | |
| AER Cap | abilities | | |
| The Adv | | apability is an optional PCIe Extended Capability. Inting | which when enabled. |
| The Adv allows a | anced Error Reporting(AER) Ca | | which when enabled. |
| The Adv allows a II Enabl | anced Error Reporting(AER) Ca dvanced error control and rep | orting | |
| The Adv allows a I Enabl | anced Error Reporting(AER) Ca dvanced error control and rep le AER Capability(PF0) le AER Capability(PF1) | F ECRC Check Capable(PF0) | F ECRC Gen Capable(PF0) |
| The Adv allows a I Enabl Enabl ARI Capa The Alte when en | anced Error Reporting(AER) Ca dvanced error control and rep le AER Capability(PF0) le AER Capability(PF1) ability mative Routing ID-Interpretatio labled, allows a device to supp | F ECRC Check Capable(PF0) | ECRC Gen Capable(PF0) ECRC Gen Capable(PF1) apability. which field vector (Bus |
| The Adv allows a Enabl Enabl ARI Capa The Alte when en Number | anced Error Reporting(AER) Ca dvanced error control and rep le AER Capability(PF0) le AER Capability(PF1) ability mative Routing ID-Interpretatio labled, allows a device to supp | Ting ECRC Check Capable(PF0) ECRC Check Capable(PF1) (ARI) Capability is an optional PCIe Extended C ort up to 256 functions by reducing the ID from 3 aber) to a 2 field vector (Bus Number, Function N | ECRC Gen Capable(PF0) ECRC Gen Capable(PF1) apability. which field vector (Bus |
| The Adv allows a Enabl Enabl ARI Capa The Alte when en Number | anced Error Reporting(AER) Ca dvanced error control and rep le AER Capability(PF0) le AER Capability(PF1) ability mative Routing ID-Interpretatio iabled, allows a device to supp . Device Number. Function Nur le ARI Capability(PF0) | Ting ECRC Check Capable(PF0) ECRC Check Capable(PF1) (ARI) Capability is an optional PCIe Extended C ort up to 256 functions by reducing the ID from 3 aber) to a 2 field vector (Bus Number, Function N | ECRC Gen Capable(PF0) ECRC Gen Capable(PF1) apability. which B field vector (Bus Sumber) |
| The Adv allows a F Enabl F Enabl ARI Capa The Alte when en Number F Enabl RBAR Ca The Res | anced Error Reporting(AER) Ca dvanced error control and rep le AER Capability(PF0) le AER Capability(PF1) ability mative Routing ID-Interpretatio abled. allows a device to supp . Device Number. Function Nur le ARI Capability(PF0) apability izable BAR(RBAR) Capabilities capability for Functions with BA | Ting ECRC Check Capable(PF0) ECRC Check Capable(PF1) (ARI) Capability is an optional PCIe Extended C ort up to 256 functions by reducing the ID from 3 aber) to a 2 field vector (Bus Number, Function N | ECRC Gen Capable(PF0) ECRC Gen Capable(PF1) apability. which 5 field vector (Bus Jumber) Capability(PF1) hen enabled. |
| The Adv allows a F Enabl E Enabl ARI Capa The Alte when en Number RBAR Ca RBAR Ca The Res adds a c | anced Error Reporting(AER) Ca dvanced error control and rep le AER Capability(PF0) le AER Capability(PF1) ability mative Routing ID-Interpretatio abled. allows a device to supp . Device Number. Function Nur le ARI Capability(PF0) apability izable BAR(RBAR) Capabilities capability for Functions with BA | T ECRC Check Capable(PF0) E ECRC Check Capable(PF0) E ECRC Check Capable(PF1) n (ARI) Capability is an optional PCIe Extended C oft up to 256 functions by reducing the ID from 3 nber) to a 2 field vector (Bus Number, Function N E Enable ARI s an optional PCIe Extended Capability, which wi as to report various options for sizes of their me | ECRC Gen Capable(PF0) ECRC Gen Capable(PF1) apability. which 5 field vector (Bus Jumber) Capability(PF1) hen enabled. |

Figure 7-13: Page 13: PCIe Extended Capabilities

| Virtex-7 FPGA Gen3 Integrated Block for PCI Express wire wire PCIE Stended Capabilities 2 wire PCIE Stended Capabilities 2 mice PCIE Stended Capability (PF0) mice The fower Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the PCI Express Base Specification Revision 3.0 for more details mice PL Capability The Capability(PF0) mice the Root Complex. See Section 6.18 of the PCI Express Base Specification Revision 3.0 for more details mable the PCI Express Base Specification Revision 3.0 for more details PL Capability The Capability(PF0) mable the Root Complex. See Section 6.18 of the PCI Express Base Specification Revision 3.0 for more details mable the PCI Express Base Specification Revision 3.0 for more details PL Capability The Capability is a feature which facilitates optimization of request targeting Memory Space. The TLP Processing Mitris Capability is required for all functions capabile di generating request TLP's with TM-1. See section 6.1.7 of the PCI Express Base Specification Space Enable (SFF-FFF: Dword Aligned) User Defined Configuration Space Enable (SFF-FFF: Dword Aligned) mable TPL Capability FD | Virtex-7 FPGA 0 | Gen3 Integrated Block for PCI Ex | kpress | | | . × |
|--|-----------------|----------------------------------|-----------------------------------|-------------------------------------|-----------------------------|------------|
| Integrated Block for PLS subset subset PS Automatical States PS Castended Capabilities 2 PS PS PS Capability The Power Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the Cipaper See Section 2.15 of the PCI spatiation (PRO) Image: Cipaper Section 2.15 of the PCI spatiality (PFO) The Capability The able PB Capability(PFC) Image: Cipaper Section 2.15 of the PCI spatiality allows Endpoints to report Latency requirements to the Root Complex. Section 1.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Section 2.15 of the PCI Express Base Specification 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Section 2.15 of the PCI Express Section 2.15 of the PCI Express Section 2.15 o | Documents | | | | | |
| Integrated Block for PLS subset subset PS Automatical States PS Castended Capabilities 2 PS PS PS Capability The Power Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the Cipaper See Section 2.15 of the PCI spatiation (PRO) Image: Cipaper Section 2.15 of the PCI spatiality (PFO) The Capability The able PB Capability(PFC) Image: Cipaper Section 2.15 of the PCI spatiality allows Endpoints to report Latency requirements to the Root Complex. Section 1.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification Revision 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Base Specification 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Section 2.15 of the PCI Express Base Specification 3.0 for more details Image: Cipaper Section 2.15 of the PCI Express Section 2.15 of the PCI Express Section 2.15 of the PCI Express Section 2.15 o | | Virtex-7 | FPGA Gen3 | | | |
| Cbc Extended Capabilities 2 PC Capability The Power Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the Cickpress Base Specification Revision 3.0 for more details. Chc Extended Capability (FF0) C Enable PB Capability(FF1) UR Capability The Latency Tolerance Reporting Capability allows Endpoints to report Latency requirements to the Root Complex. See Section 61.8 of the PCI Express Base Specification Revision 3.0 for more details Chable TR Capability is required for all functions capable of generating request turp's with TPH. See section 61.0 of the PCI Express Base Specification Scapable of generating request Turp's with TPH. See section 61.0 of the PCI Express Base Specification Scapable of generating request Turp's with TPH. See section 61.0 of the PCI Express Base Specification Capability (FF1) User Defined Configuration Capabilities C Enable TPH Capability(FF0) Che Express Extended Configuration Space Enable (3FF-FFF. Dword Aligned) | IndiCORE | | | | | |
| PCIE Extended Capability 2 P PG Capability The Power Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the PCI Express Base Specification Revision 3.0 for more details. P P Enable PE Capability (FPO) P Enable PB Capability (FF1) P URC Capability The Latency Tolerance Reporting Capability allows Endpoints to report Latency requirements to the Root Complex. See Section 5.18 of the PCI Express Base Specification Revision 3.0 for more details P P Enable PB Capability The Latency Tolerance Reporting Capability allows Endpoints to report Latency requirements to the Root Complex. See Section 5.18 of the PCI Express Base Specification Revision 3.0 for more details P P Enable ITR Capability is required for all functions capable of generating request TLP's with TFH. See section 6.17 of the PCI Express Base Specification S.0 for more details P Enable TFH Capability(FF0) P Enable TFH Capability(FF1) User Defined Configuration Capabilities P CI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) P CI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) P | Ly | | | | | |
| P8 Capability The Power Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the PCI Express Base Specification Revision 3.0 for more details. C Explores Base Specification Revision 3.0 for more details. C Rable P8 Capability(FF0) C Enable P8 Capability(FF1) UR Capability The Capability F0 C enable TPH Capability (F60) C enable TPH Capability(F61) User Defined Configuration Capabilities P PCI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) | | | (press | | xilinx.com.ip:pcie3_7x:1 | 1 |
| The Power Budgeting Capability allows a device to report power consumption to the system. See Section 7.15 of the PCI Express Base Specification Revision 3.0 for more details. | PCle Exten | ded Capabilities 2 | | | | 7 |
| PCI Express Base Specification Revision 3.0 for more details. | PB Capab | pility | | | | |
| LTR Capability The Latency Tolerance Reporting Capability allows Endpoints to report Latency requirements to the Root Complex. See Section 6.18 of the PCI Express Base Specification Revision 3.0 for more details □ Enable LTR Capability(FF0) THP Capability TLP Processing Hints (TPH) is a feature which facilitates optimization of requests targeting Memory Space. The TLP Processing Hints Capability is required for all functions capable of generating request TLP's with TPH. See section 6.17 of the PCI Express Base Specification 3.0 for more details □ Enable TPH Capability(FF0) □ Enable TPH Capability(FF0) □ Enable TPH Capability(FF1) User Defined Configuration Capabilities □ □ PCI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) | | | | sumption to the system. See Section | n 7.15 of the | |
| The Latency Tolerance Reporting Capability allows Endpoints to report Latency requirements to the Root Complex. See Section 6.18 of the PCI Express Base Specification Revision 3.0 for more details The Capability TLP Cocessing Hints (TPH) is a feature which facilitates optimization of requests targeting Memory Space. The TLP Processing Hints Capability is required for all functions capable of generating request TLP's with TPH. See section 6.17 of the PCI Express Base Specification 3.0 for more details C Enable TPH Capability (PFO) User Defined Configuration Capabilities PCI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) | 🗆 Enable | PB Capability(PF0) | | □ Enable PB Capability(PF1) | | |
| See Section 6.18 of the PCI Express Base Specification Revision 3.0 for more details | LTR Capa | bility | | | | |
| TPH Capability TLP Processing Hints (TPH) is a feature which facilitates optimization of requests targeting Memory Space. The TLP Processing Hints Capability is required for all functions capable of generating request TLP's with TPH. See section 6.17 of the PCI Express Base Specification 3.0 for more details Tenable TPH Capability(FF0) User Defined Configuration Capabilities PCI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) | | | | | Complex. | |
| TLP Processing Hints (TPH) is a feature which facilitates optimization of requests targeting Memory Space. The TLP Processing Hints Capability is required for all functions capable of generating request TLP's with TPH. See section 6.17 of the PCI Express Base Specification 3.0 for more details | 🗆 Enable | ELTR Capability(PF0) | | | | |
| Hints Capability is required for all functions capable of generating request TLP's with TPH. See section 6.17 of the PCI Express Base Specification 3.0 for more details | TPH Capa | ability | | | | |
| User Defined Configuration Capabilities PCI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) | Hints Cap | ability is required for all fu | inctions capable of generating re | | | |
| PCI Express Extended Configuration Space Enable (3FF-FFF: Dword Aligned) | 🗆 Enable | e TPH Capability(PF0) | | □ Enable TPH Capability(PF1) | | |
| | User Defi | ned Configuration Capabili | ities | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | T PCI Ex | press Extended Configurat | tion Space Enable (3FF-FFF: Dwor | d Aligned) | | |
| < Back Page 14 of 14 Next > Generate Gancel Help | | | | | | |
| ح Rack Page 14 of 14 <u>Next</u> <u>Generate</u> <u>Cancel</u> <u>Help</u> | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Gancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | |
| < Back Page 14 of 14 Next > Generate Cancel Help | | | | | | - |
| | | | | < Back Page 14 of 14 | Mext > Generate Cancel Help | \Box_{j} |

Figure 7-14: Page 14: PCIe Extended Capabilities 2

Device Serial Number Capability

• **Device Serial Number Capability**: An optional PCIe Extended Capability containing a unique Device Serial Number. When this Capability is enabled, the DSN identifier must be presented on the Device Serial Number input pin of the port. This Capability must be turned on to enable the Virtual Channel and Vendor Specific Capabilities

Virtual Channel Capability

• **Virtual Channel Capability**: An optional PCIe Extended Capability which allows the user application to be operated in TCn/VC0 mode. Checking this allows Traffic Class filtering to be supported. This capability only exists for Physical Function 0.

• **Reject Snoop Transactions (Root Port Configuration Only)**: When enabled, any transactions for which the No Snoop attribute is applicable, but is not set in the TLP header, can be rejected as an Unsupported Request.

AER Capability

• **Enable AER Capability**: An optional PCIe Extended Capability that allows Advanced Error Reporting. This capability is always enabled.

Additional Optional Capabilities

- **Enable ARI**: An optional PCIe Extended Capability that allows Alternate Requestor ID. This capability is automatically enabled if SRIOV is enabled.
- **Enable PB**: An optional PCIe Extended Capability that implements the Power Budgeting Enhanced Capability Header.
- **Enable RBAR**: An optional PCIe Extended Capability that implements the Resizable BAR Capability.
- **Enable LTR**: An optional PCIe Extended Capability that implements the Latency Tolerance Reporting Capability.
- **Enable DPA**: An optional PCIe Extended Capability that implements Dynamic Power Allocation Capability.
- **Enable TPH**: An optional PCIe Extended Capability that implements Transaction Processing Hints Capability.

Output Generation

The Gen3 Integrated Block for PCIe example design directories and their associated files are defined in the sections that follow. Click a directory name to go to the desired directory and its associated files.

c) <project directory> Top-level project directory; name is user-defined 🛅 <project directory>/<component name> Core release notes readme file <component name>/doc Product documentation 🛅 <component name>/example_design Verilog or VHDL design files <component name>/implement **i** Implementation script files implement/results Contains implement script results implement/xst Contains synthesis results, when XST is chosen as the synthesis tool <component name>/source Core source files <component name>/simulation Simulation scripts simulation/dsport (for Endpoint configuration only) Root Port Bus Functional Model simulation/functional Functional simulation files

simulation/tests (for Endpoint configuration only) Test command files

<project directory>

The project directory contains all the CORE Generator tool project files.

| Table 7-6: Project Direct |
|---------------------------|
|---------------------------|

| Name | Description | |
|---|---|--|
| <project_dir></project_dir> | | |
| <component_name>.xco</component_name> | CORE Generator software project-specific option file; can be used as an input to the CORE Generator tool. | |
| <component_name>_flist.txt</component_name> | List of files delivered with core. | |
| <component_name>.{veo vho}</component_name> | Verilog or VHDL instantiation template. | |
| <component_name>_xmdf.tcl</component_name> | Xilinx standard IP Core information file used by Xilinx design tools. | |
| <component_name>_synth.{v vhd}</component_name> | Verilog/VHDL top-level solution wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. This file receives user-specific settings and sets parameters accordingly | |

Back to Top

<project directory>/<component name>

The component name directory contains the release notes in the readme file provided with the core, which can include tool requirements, updates, and issue resolution.

Table 7-7: Component Name Directory

| Name | Description | |
|---|---------------------|--|
| <project_dir>/<component_name></component_name></project_dir> | | |
| <pre>pcie3_7x_v1_1_readme.txt</pre> | Release notes file. | |

Back to Top

<component name>/doc

The doc directory contains a redirect PDF, which points to this document on the Xilinx website.

Table 7-8: Doc Directory

| Name | Description | |
|---|---|--|
| <project_dir>/<component_name>/doc</component_name></project_dir> | | |
| pg023_v7_pcie_gen3.pdf | Virtex-7 FPGA Gen3 Integrated Block for PCI Express Product Guide | |
| Pack to Tan | 1 | |

<component name>/example_design

The example_design directory contains the example design files provided with the core. Table 3-9 shows the directory contents for an Endpoint configuration core.

| Name | Description | |
|---|--|--|
| <project_dir>/<component_name>/example_design</component_name></project_dir> | | |
| <pre>xilinx_pcie_3_0_ep_7x_01_lane_ gen1_xc7vx690t-ffg1761-3-PCIE_ X0Y0.ucf</pre> | Example design UCF. The file name varies by Device/ Port Type, lane width, maximum link speed, part, package, Integrated Block for PCI Express block location, and Xilinx Development Board selected. | |
| xilinx_pcie_3_0_ep_xt.v | Verilog top-level PIO example design file. | |
| <pre>pcie_app_7vx.v PIO_INTR_CTRL.v EP_MEM.v[hd] PIO.v[hd]\PIO_EP.v[hd] PIO_EP_MEM_ACCESS.v[hd] PIO_TO_CTRL.v[hd] PIO_RX_ENGINE.v[hd] PIO_TX_ENGINE.v[hd]</pre> | PIO example design files. | |

 Table 7-9:
 Example Design Directory: Endpoint Configuration

Back to Top

<component name>/implement

The implement directory contains the core implementation script files.

Table 7-10: Implement Directory

| Name | Description | |
|---|---------------------------------------|--|
| <project_dir>/<component_name>/implement</component_name></project_dir> | | |
| implement.bat | DOS and Linux implementation scripts. | |
| implement.sh | bos and linux implementation scripts. | |
| xilinx_pcie_3_0_7vx_ep.prj | XST file list for the core. | |
| xilinx_pcie_3_0_7vx_ep.xst | XST command file. | |
| xilinx_pcie_3_0_7vx_ep.xcf | XST synthesis constraints file. | |

implement/results

The results directory is created by the implement script. The implement script results are placed in the results directory.

Table 7-11: Results Directory

| Name | Description |
|---|-------------|
| <project_dir>/<component_name>/implement/results</component_name></project_dir> | |
| Implement script result files. | |

Back to Top

implement/xst

The xst directory is created by the XST script. The synthesis results are placed in the xst directory.

Table 7-12: XST Results Directory

| | Name | Description |
|---|------|-------------|
| <project_dir>/<component_name>/implement/xst</component_name></project_dir> | | |
| XST result files. | | |
| Pack to Top | | |

<component name>/source

The source directory contains the generated core source files.

| Name | Description |
|--|--|
| <project_dir>/<compor< td=""><td>nent_name>/source</td></compor<></project_dir> | nent_name>/source |
| pcie_3_0_7vx.v | Verilog top-level core wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| pcie_top.v | AXI4-Stream solution wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| pcie_7vx.v | Solution Wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| pcie_init_ctrl_7vx.v | Initialization Controller for Virtex-7 FPGA Gen3 Integrated Block for PCI Express |
| <pre>pcie_pipe_pipeline.v pcie_pipe_lane.v pcie_pipe_misc.v</pre> | PIPE module for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| <pre>pcie_bram_7vx.v pcie_bram_7vx_16k.v pcie_bram_7vx_8k.v pcie_bram_7vx_cpl.v pcie_bram_7vx_rep.v pcie_bram_7vx_rep_8k.v pcie_bram_7vx_req.v</pre> | Block RAM module for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| gt_top.v | GTH wrapper for the Virtex-7 FPGA Gen3 Integrated Block for PCI Express. |
| <pre>gt_wrapper.v pipe_clock.v pipe_drp.v pipe_rate.v pipe_reset.v pipe_sync.v pipe_user.v pipe_wrapper.v pipe_eq.v rxeq_scan.v qpll_drp.v qpll_reset.v qpll_wrapper.v</pre> | GTH module for the Virtex-7 FPGA GTH transceivers. |

<component name>/simulation

The simulation directory contains the simulation source files provided with the core.

simulation/dsport

The dsport directory contains the files for the Root Port model test bench.

| Name | Description | |
|---|------------------------|--|
| <project_dir>/<component_name>/simulation/dsport</component_name></project_dir> | | |
| pcie_2_1_rp_v7.v | | |
| pci_exp_expect_tasks.v | | |
| pci_exp_usrapp_cfg.v | | |
| pci_exp_usrapp_com.v | | |
| pci_exp_usrapp_pl.v | Root Port model files. | |
| pci_exp_usrapp_rx.v | | |
| pci_exp_usrapp_tx.v | | |
| <pre>xilinx_pcie_2_1_rport_v7.v</pre> | | |
| test_interface.vhd | | |

Back to Top

simulation/functional

The functional directory contains functional simulation scripts provided with the core.

| Name | Description |
|---|---|
| <project_dir>/<com< td=""><td>ponent_name>/simulation/functional</td></com<></project_dir> | ponent_name>/simulation/functional |
| board.f | List of files for RTL simulations. |
| simulate_mti.do | Simulation script for ModelSim. |
| simulate_ncsim.sh | Simulation script for Cadence IES (Verilog only). |
| simulate_vcs.sh | Simulation script for VCS (Verilog only). |
| simulate_isim.sh | Simulation script for ISim (Verilog only). |
| xilinx_lib_vcs.f | Points to the required SecureIP Model. |
| board_common.v (Endpoint configuration only) | Contains test bench definitions (Verilog only). |
| board.v | Top-level simulation module. |
| sys_clk_gen_ds.v (Endpoint configuration only) | System differential clock source. |
| sys_clk_gen.v | System clock source. |

Table 7-15: Functional Directory

simulation/tests

Note: This directory exists for Endpoint configuration only.

The tests directory contains test definitions for the example test bench.

Table 7-16: Tests Directory

| simulation/tests |
|----------------------------------|
| , india 01011, 00000 |
| initions for example test bench. |
| |
| i |



Constraining the Core

Required Constraints

The Virtex®-7 FPGA Gen3 Integrated Block for PCI Express® solution requires the specification of timing and other physical implementation constraints to meet specified performance requirements for PCI Express. These constraints are provided with the Endpoint and Root Port solutions in a user constraints file (UCF). Pinouts and hierarchy names in the generated UCF correspond to the provided example design.

To achieve consistent implementation results, a UCF containing these original, unmodified constraints must be used when a design is run through the Xilinx tools. For additional details on the definition and use of a UCF or specific constraints, see the *Xilinx Libraries Guide* and/or *Command Line Tools User Guide*.

Constraints provided with the Integrated Block solution have been tested in hardware and provide consistent results. Constraints can be modified, but modifications should only be made with a thorough understanding of the effect of each constraint. Additionally, support is not provided for designs that deviate from the provided constraints.

Device, Package, and Speed Grade Selections

The device selection portion of the UCF informs the implementation tools which part, package, and speed grade to target for the design. Because Gen3 Integrated Block for PCIe cores are designed for specific part and package combinations, this section should not be modified by the designer.

The device selection section always contains a part selection line, but can also contain part or package-specific options. An example part selection line follows:

CONFIG PART = XC7VX690T-FFG1761-3

Clock Frequencies

See Chapter 3, Designing with the Core, for detailed information about clock requirements.

Clock Management

See Chapter 3, Designing with the Core, for detailed information about clock requirements.

Clock Placement

See the 7 Series FPGAs GTX/GTH Transceivers User Guide [Ref 3] for guidelines regarding clock resource selection.

See Chapter 3, Designing with the Core, for detailed information about clock requirements.

Stacked Silicon Interconnect Devices

Some Virtex-7 devices utilize stacked silicon interconnect (SSI) technology. The I/O and Integrated Block must remain on the same die when targeting an SSI device.

The sys_clk must be chosen to be in the same bank as the GTH transceiver it is connected to, or one bank above/below the GTH transceiver being used.

For more information, see the "Placement Information by Package" and "Placement Information by Device" appendices in the 7 Series FPGAs GTX/GTH Transceivers User Guide.

Transceiver Placement

These constraints select which transceivers to use and dictates the pinout for the transceiver differential pairs. For more information, see the "Placement Information by Package" appendix in the 7 Series FPGAs GTX/GTH Transceivers User Guide [Ref 3].

Table 8-1 through Table 8-8 list the supported transceiver locations available for supported Virtex-7 FPGA part and package combinations. The CORE Generator tool provides a UCF for the selected part and package that matches the table contents. The following lists all devices with their associated tables containing transceiver locations:

- XC7VX330T: Table 8-1
- XC7VX415T: Table 8-2
- XC7VX550T: Table 8-3
- XC7VX690T: Table 8-4
- XC7VX980T: Table 8-5
- XC7VX1140T: Table 8-6
- XC7VH580T: Table 8-7
- XC7VH870T: Table 8-8

Table 8-1: Supported Transceiver Locations for the XC7VX330T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | | |
|---------|-----------------|--|--------|--------|--------|--------|--------|--------|--------|--|--|--|--|
| | X0Y0 | X0Y11 | X0Y10 | X0Y9 | X0Y8 | X0Y7 | X0Y6 | X0Y5 | X0Y4 | | | | |
| FFG1157 | X0Y1 | X0Y23 | X0Y22 | X0Y21 | X0Y20 | X0Y19 | X0Y18 | X0Y17 | X0Y16 | | | | |
| FEGIIS | X0Y2 N/A | | | | | | | | | | | | |
| | X0Y3 | | N/A | | | | | | | | | | |
| | X0Y0 | X0Y11 | X0Y10 | X0Y9 | X0Y8 | X0Y7 | X0Y6 | X0Y5 | X0Y4 | | | | |
| FFG1761 | X0Y1 | X0Y23 X0Y22 X0Y21 X0Y20 X0Y19 X0Y18 X0 | | | | | | | X0Y16 | | | | |
| FLG1/01 | X0Y2 | | | | N, | /A | | | | | | | |
| | X0Y3 | | | | N, | /A | | | | | | | |

Table 8-2: Supported Transceiver Locations for the XC7VX415T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|---------|-------|---|--------|--------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | X1Y3 | X1Y2 | X1Y1 | X1Y0 | | | |
| FFG1157 | X0Y1 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | X1Y15 | X1Y14 | X1Y13 | X1Y12 | | | |
| FEGII21 | X0Y2 | | N/A | | | | | | | | | |
| | X0Y3 | | N/A | | | | | | | | | |
| | X0Y0 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | X1Y3 | X1Y2 | X1Y1 | X1Y0 | | | |
| FFG1158 | X0Y1 | X0Y1 X1Y19 X1Y18 X1Y17 X1Y16 X1Y15 X1Y14 | | | | | | | X1Y12 | | | |
| LL01120 | X0Y2 | | | | N, | /A | | | | | | |
| | X0Y3 | | | | N, | /A | | | | | | |

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|---------|-------|--------|---|--------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | X1Y3 | X1Y2 | X1Y1 | X1Y0 | | | |
| FFG1927 | X0Y1 | X1Y19 | X1Y19 X1Y18 X1Y17 X1Y16 X1Y15 X1Y14 X1Y13 X1Y12 | | | | | | | | | |
| FFG1927 | X0Y2 | | N/A | | | | | | | | | |
| | X0Y3 | | | | N, | /A | | | | | | |

Table 8-2: Supported Transceiver Locations for the XC7VX415T (Cont'd)

Table 8-3: Supported Transceiver Locations for the XC7VX550T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|
| | X0Y0 | | | | N, | /A | | | | | |
| FFG1158 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| FEGIISO | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | N/A | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| EEC1027 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| FFG1927 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | | · | | N, | /A | · | | | | |

Table 8-4: Supported Transceiver Locations for the XC7VX690T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|---------|-------|--------|---|----------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | | | | N, | /A | | | | | | |
| FFG1157 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FEGII5/ | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | N/A | | | | | | | | | |
| | X0Y0 | | | | N, | /A | | | | | | |
| FFG1158 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1156 | X0Y2 | X1Y35 | X1Y35 X1Y34 X1Y33 X1Y32 X1Y31 X1Y30 X1Y29 X1Y28 | | | | | | | | | |
| | X0Y3 | | N/A | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFG1761 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1701 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | | | N, | /A | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFG1926 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1920 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | <u>.</u> | <u>.</u> | N, | /A | | * | | | | |

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | |
|---------|-------|--------|----------|----------|--------|--------|--------|--------|--------|--|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | |
| FFG1927 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1927 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | N/A | | | | | | | | | |
| | X0Y0 | | N/A | | | | | | | | | |
| EEC1020 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | |
| FFG1930 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | |
| | X0Y3 | | <u>.</u> | <u>.</u> | N, | /A | | * | | | | |

Table 8-4: Supported Transceiver Locations for the XC7VX690T (Cont'd)

Table 8-5: Supported Transceiver Locations for the XC7VX980T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | | | |
|---------|-------|--------|---|--------|--------|--------|--------|--------|--------|--|--|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | | |
| FFG1926 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | | |
| FFG1920 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | | |
| | X0Y3 | | | | N, | /A | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | | |
| FFG1928 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | | |
| FFG1928 | X0Y2 | X1Y35 | X1Y35 X1Y34 X1Y33 X1Y32 X1Y31 X1Y30 X1Y29 X1Y28 | | | | | | | | | | |
| | X0Y3 | | N/A | | | | | | | | | | |
| | X0Y0 | | N/A | | | | | | | | | | |
| FFG1930 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | | |
| FFG1930 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | | |
| | X0Y3 | | | | N, | /A | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | | | |
| FFG1933 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | | | |
| 1101333 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | | | |
| | X0Y3 | | | | N, | /A | | | | | | | |

Table 8-6: Supported Transceiver Locations for the XC7VX1140T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 |
| FFG1926 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 |
| FFG1920 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 |
| | X0Y3 | | | | N, | /A | | | |

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| FLG1928 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| FLG1920 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | X1Y47 | X1Y46 | X1Y45 | X1Y44 | X1Y43 | X1Y42 | X1Y41 | X1Y40 | | |
| | X0Y0 | | N/A | | | | | | | | |
| FLG1930 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| FLG1950 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | N/A | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| FLG1933 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| FLG1955 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 | | |
| | X0Y3 | | | | N | /A | | | | | |

Table 8-6: Supported Transceiver Locations for the XC7VX1140T (Cont'd)

| Table 8-7: | Supported | Transceiver | Locations | for the | XC7VH580T |
|------------|-----------|-------------|-----------|---------|-----------|
|------------|-----------|-------------|-----------|---------|-----------|

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 | | |
|---------|-------|--------|--------|----------|--------|--------|--------|--------|----------|--|--|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| HCG1155 | X0Y1 | | N/A | | | | | | | | |
| псотто | X0Y2 | | N/A | | | | | | | | |
| | X0Y3 | | N/A | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| HCG1931 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| HCG1951 | X0Y2 | N/A | | | | | | | | | |
| | X0Y3 | N/A | | | | | | | | | |
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 | | |
| HCG1932 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 | | |
| nc01952 | X0Y2 | | * | <u>.</u> | N, | /A | * | | <u>.</u> | | |
| | X0Y3 | | | | N | /A | | | | | |

Table 8-8: Supported Transceiver Locations for the XC7VH870T

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|
| | X0Y0 | | | | N, | /A | | | |
| HCG1931 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 |
| HCG1931 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 |
| | X0Y3 | | | 1 | N, | /A | | 1 | |

Gen3 Integrated Block for PCIe (v1.3) PG023 October 16, 2012

| Package | Block | Lane 0 | Lane 1 | Lane 2 | Lane 3 | Lane 4 | Lane 5 | Lane 6 | Lane 7 |
|---------|-------|--------|--------|--------|--------|--------|--------|--------|--------|
| | X0Y0 | X1Y11 | X1Y10 | X1Y9 | X1Y8 | X1Y7 | X1Y6 | X1Y5 | X1Y4 |
| HCG1932 | X0Y1 | X1Y23 | X1Y22 | X1Y21 | X1Y20 | X1Y19 | X1Y18 | X1Y17 | X1Y16 |
| HCG1932 | X0Y2 | X1Y35 | X1Y34 | X1Y33 | X1Y32 | X1Y31 | X1Y30 | X1Y29 | X1Y28 |
| | X0Y3 | N/A | | | | | | | |

Table 8-8: Supported Transceiver Locations for the XC7VH870T

I/O Standard and Placement

This section controls the placement and options for I/Os belonging to the core System (SYS) interface and PCI Express (PCI_EXP) interface. NET constraints in this section control the pin location and I/O options for signals in the SYS group. Locations and options vary depending on which derivative of the core is used and should not be changed without fully understanding the system requirements.

For example:

```
NET "sys_rt_n" IOSTANDARD = LVCMOS18| PULLUP | NODELAY;
INST "refclk_ibuf" LOC = IBUFDS_GT2_X0Y7;
```

INST constraints control placement of signals that belong to the PCI_EXP group. These constraints control the location of the transceiver(s) used, which implicitly controls pin locations for the transmit and receive differential pair.

For example:

```
INST "core_i/gt_top_i/pipe_wrapper_i/pipe_lane[0].gt_wrapper_i/
gth_channel.gthe2_channel_i" LOC = GTHE2_CHANNEL_X1Y11;
```



Chapter 9

Example Design and Model Test Bench for Endpoint Configuration

Directory and File Contents

See Output Generation, page 184 for directory structure and file contents.

Example Design

This section provides an overview of the Virtex®-7 FPGA Integrated Block for PCI Express® Gen3 example design and instructions for generating the core. It also includes information about simulating and implementing the example design using the provided demonstration test bench.

For current information about generating, simulating, and implementing the core, see the Release Notes provided with the core, when it is generated using the CORE Generator[™] tool available in the ISE® Design Suite.

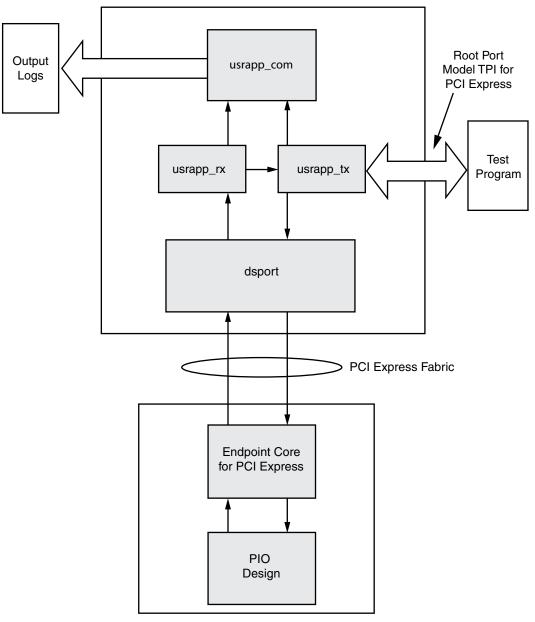
Integrated Block Endpoint Configuration Overview

The example simulation design for the Endpoint configuration of the integrated block consists of two discrete parts:

- The Root Port Model, a test bench that generates, consumes, and checks PCI Express bus traffic.
- The Programmed Input/Output (PIO) example design, a completer application for PCI Express. The PIO example design responds to Read and Write requests to its memory space and can be synthesized for testing in hardware.

Simulation Design Overview

For the simulation design, transactions are sent from the Root Port Model to the Integrated Block core (configured as an Endpoint) and processed by the PIO example design. Figure 9-1 illustrates the simulation design provided with the Integrated Block core. For more information about the Root Port Model, see Root Port Model Test Bench for Endpoint, page 280.



Endpoint DUT for PCI Express

X12471

Figure 9-1: Simulation Example Design Block Diagram

Implementation Design Overview

The implementation design consists of a simple PIO example that can accept read and write transactions and respond to requests, as illustrated in Figure 9-2. Source code for the example is provided with the core. For more information about the PIO example design, see Programmed Input/Output: Endpoint Example Design, page 271.

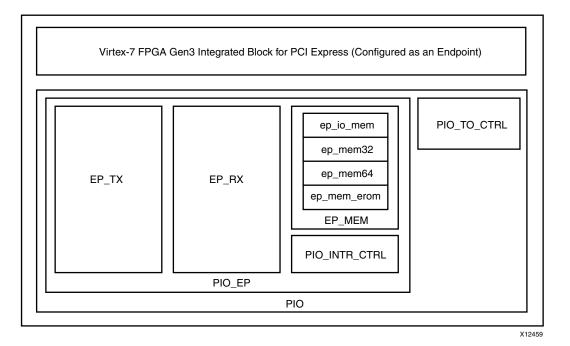


Figure 9-2: Implementation Example Design Block Diagram

Example Design Elements

The PIO example design elements include:

- Core wrapper
- An example Verilog HDL wrapper (instantiates the cores and example design)
- A customizable demonstration test bench to simulate the example design

The example design has been tested and verified with Xilinx® ISE® Design Suite v14.3 and this simulator:

- Mentor Graphics ModelSim
- Cadence IES
- Synopsys VCS
- Xilinx ISim

For the supported versions of these tools, see the Release Notes Guide.

Programmed Input/Output: Endpoint Example Design

Programmed Input/Output (PIO) transactions are generally used by a PCI Express[®] system host CPU to access Memory Mapped Input/Output (MMIO) and Configuration Mapped Input/Output (CMIO) locations in the PCI Express logic. Endpoints for PCI Express accept Memory and I/O Write transactions and respond to Memory and I/O Read transactions with Completion with Data transactions.

The PIO example design (PIO design) is included with the Gen3 Integrated Block for PCIe in Endpoint configuration generated by the CORE Generator[™] tool, which allows users to bring up their system board with a known established working design to verify the link and functionality of the board.

The PIO design Port Model is shared by the Gen3 Integrated Block for PCIe, Endpoint Block Plus for PCI Express, and Endpoint PIPE for PCI Express solutions. This appendix represents all the solutions generically using the name Endpoint for PCI Express (or Endpoint for PCIe[®]).

System Overview

The PIO design is a simple target-only application that interfaces with the Endpoint for the Transaction (AXI4-Stream) interface of the PCIe core and is provided as a starting point for customers to build their own designs. These features are included:

- Four transaction-specific 2 KB target regions using the internal FPGA block RAMs, providing a total target space of 8192 bytes
- Supports single Dword payload Read and Write PCI Express transactions to 32-/64-bit address memory spaces and I/O space with support for completion TLPs
- Utilizes the BAR ID[2:0] and Completer Request Descriptor[114:112] of the core to differentiate between TLP destination Base Address Registers
- Provides separate implementations optimized for 64-bit, 128-bit, and 256-bit AXI4-Stream interfaces

Figure 9-3 illustrates the PCI Express system architecture components, consisting of a Root Complex, a PCI Express switch device, and an Endpoint for PCIe. PIO operations move data *downstream* from the Root Complex (CPU register) to the Endpoint, and/or *upstream* from the Endpoint to the Root Complex (CPU register). In either case, the PCI Express protocol request to move the data is initiated by the host CPU.

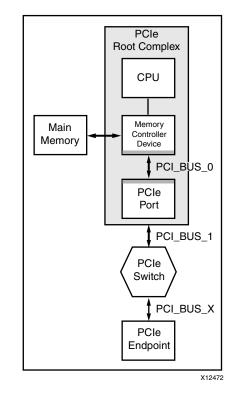


Figure 9-3: System Overview

Data is moved downstream when the CPU issues a store register to a MMIO address command. The Root Complex typically generates a Memory Write TLP with the appropriate MMIO location address, byte enables, and the register contents. The transaction terminates when the Endpoint receives the Memory Write TLP and updates the corresponding local register.

Data is moved upstream when the CPU issues a load register from a MMIO address command. The Root Complex typically generates a Memory Read TLP with the appropriate MMIO location address and byte enables. The Endpoint generates a Completion with Data TLP after it receives the Memory Read TLP. The Completion is steered to the Root Complex and payload is loaded into the target register, completing the transaction.

PIO Hardware

The PIO design implements a 8192 byte target space in FPGA block RAM, behind the Endpoint for PCIe. This 32-bit target space is accessible through single Dword I/O Read, I/O Write, Memory Read 64, Memory Write 64, Memory Read 32, and Memory Write 32 TLPs.

The PIO design generates a completion with one Dword of payload in response to a valid Memory Read 32 TLP, Memory Read 64 TLP, or I/O Read TLP request presented to it by the core. In addition, the PIO design returns a completion without data with successful status for I/O Write TLP request.

The PIO design can initiate:

- a Memory Read transaction when the received write address is 11'h7AA and the write data is 32'hAAAA_BBBB
- a Legacy Interrupt when the received write address is 11 ' h7BB and the write data is 32 ' hCCCC_DDDD
- an MSI when the received write address is 11 'h7BB and the write data is 32 'hEEEE_FFFF
- an MSIx when the received write address is 11 ' h7BB and the write data is 32 ' hDEAD_BEEF.

The PIO design processes a Memory or I/O Write TLP with one Dword payload by updating the payload into the target address in the FPGA block RAM space.

Base Address Register Support

The PIO design supports four discrete target spaces, each consisting of a 2 KB block of memory represented by a separate Base Address Register (BAR). Using the default parameters, the CORE Generator tool produces a core configured to work with the PIO design defined in this section, consisting of:

- One 64-bit addressable Memory Space BAR
- One 32-bit Addressable Memory Space BAR

Users can change the default parameters used by the PIO design; however, in some cases they might need to change the back-end User Application depending on their system. See Changing CORE Generator Tool Default BAR Settings for information about changing the default CORE Generator tool parameters and the effect on the PIO design.

Each of the four 2 KB address spaces represented by the BARs corresponds to one of four 2 KB address regions in the PIO design. Each 2 KB region is implemented using a 2 KB dual-port block RAM. As transactions are received by the core, the core decodes the address and determines which of the four regions is being targeted. The core presents the TLP to the PIO design and asserts the appropriate bits of (BAR ID[2:0]), Completer Request Descriptor[114:112], as defined in Table 9-1.

| Block RAM | TLP Transaction Type | Default BAR | BAR ID[2:0] |
|-------------|--|---------------|-------------|
| ep_io_mem | I/O TLP transactions | Disabled | Disabled |
| ep_mem32 | 32-bit address Memory TLP transactions | 2 | 000b |
| ep_mem64 | 64-bit address Memory TLP transactions | 0-1 | 001b |
| ep_mem_erom | 32-bit address Memory TLP transactions destined for EROM | Expansion ROM | 110b |

| Table 9-1: | TLP Traffic Types |
|------------|-------------------|
|------------|-------------------|

Changing CORE Generator Tool Default BAR Settings

You can change the CORE Generator tool parameters and continue to use the PIO design to create customized Verilog source to match the selected BAR settings. However, because the PIO design parameters are more limited than the core parameters, consider these example design limitations when changing the default CORE Generator tool parameters:

- The example design supports one I/O space BAR, one 32-bit Memory space (that cannot be the Expansion ROM space), and one 64-bit Memory space. If these limits are exceeded, only the first space of a given type is active—accesses to the other spaces do not result in completions.
- Each space is implemented with a 2 KB memory. If the corresponding BAR is configured to a wider aperture, accesses beyond the 2 KB limit wrap around and overlap the 2 KB memory space.
- The PIO design supports one I/O space BAR, which by default is disabled, but can be changed if desired.

Although there are limitations to the PIO design, Verilog source code is provided so users can tailor the example design to their specific needs.

TLP Data Flow

This section defines the data flow of a TLP successfully processed by the PIO design.

The PIO design successfully processes single Dword payload Memory Read and Write TLPs and I/O Read and Write TLPs. Memory Read or Memory Write TLPs of lengths larger than one Dword are not processed correctly by the PIO design; however, the core does accept these TLPs and passes them along to the PIO design. If the PIO design receives a TLP with a length of greater than one Dword, the TLP is received completely from the core and discarded. No corresponding completion is generated.

Memory and I/O Write TLP Processing

When the Endpoint for PCIe receives a Memory or I/O Write TLP, the TLP destination address and transaction type are compared with the values in the core BARs. If the TLP passes this comparison check, the core passes the TLP to the Receive AXI4-Stream interface of the PIO design. The PIO design handles Memory writes and I/O TLP writes in different ways: the PIO design responds to *I/O writes* by generating a Completion Without Data (cpl), a requirement of the PCI Express specification.

Along with the start of packet, end of packet, and ready handshaking signals, the Completer Requester AXI4-Stream interface also asserts the appropriate (BAR ID[2:0]), Completer Request Descriptor[114:112] signal to indicate to the PIO design the specific destination BAR that matched the incoming TLP. On reception, the RX State Machine of the PIO design processes the incoming Write TLP and extracts the TLPs data and relevant address fields so that it can pass this along to the PIO design internal block RAM write request controller. Based on the specific BAR ID[2:0] signals asserted, the RX state machine indicates to the internal write controller the appropriate 2 KB block RAM to use prior to asserting the write enable request. For example, if an I/O Write Request is received by the core targeting BAR0, the core passes the TLP to the PIO design and sets BAR ID[2:0] to 000b. The RX state machine extracts the lower address bits and the data field from the I/O Write TLP and instructs the internal Memory Write controller to begin a write to the block RAM.

In this example, the assertion of setting BAR ID[2:0] to 000b instructed the PIO memory write controller to access ep_mem0 (which by default represents 2 KB of I/O space). While the write is being carried out to the FPGA block RAM, the PIO design RX state machine deasserts m_axis_cq_tready, causing the Receive AXI4-Stream interface to stall receiving any further TLPs until the internal Memory Write controller completes the write to the block RAM. Deasserting m_axis_cq_tready in this way is not required for all designs using the core—the PIO design uses this method to simplify the control logic of the RX state machine.

Memory and I/O Read TLP Processing

When the Endpoint for PCIe receives a Memory or I/O Read TLP, the TLP destination address and transaction type are compared with the values programmed in the core BARs. If the TLP passes this comparison check, the core passes the TLP to the Receive AXI4-Stream interface of the PIO design.

Along with the start of packet, end of packet, and ready handshaking signals, the Completer Requester AXI4-Stream interface also asserts the appropriate BAR ID[2:0] signal to indicate to the PIO design the specific destination BAR that matched the incoming TLP. On reception, the PIO design state machine processes the incoming Read TLP and extracts the relevant TLP information and passes it along to the PIO design's internal block RAM read request controller.

Based on the specific BAR ID[2:0] signal asserted, the RX state machine indicates to the internal read request controller the appropriate 2 KB block RAM to use before asserting the read enable request. For example, if a Memory Read 32 Request TLP is received by the core targeting the default Mem32 BAR2, the core passes the TLP to the PIO design and sets BAR ID[2:0] to 010b. The RX state machine extracts the lower address bits from the Memory 32 Read TLP and instructs the internal Memory Read Request controller to start a read operation.

In this example, the setting BAR ID[2:0] to 010b instructs the PIO memory read controller to access the Mem32 space, which by default represents 2 KB of memory space. A notable difference in handling of memory write and read TLPs is the requirement of the receiving device to return a Completion with Data TLP in the case of memory or I/O read request.

While the read is being processed, the PIO design RX state machine deasserts m_axis_cq_tready, causing the Receive AXI4-Stream interface to stall receiving any further TLPs until the internal Memory Read controller completes the read access from the block RAM and generates the completion. Deasserting m_axis_cq_tready in this way is not required for all designs using the core. The PIO design uses this method to simplify the control logic of the RX state machine.

PIO File Structure

Table 9-2 defines the PIO design file structure. Based on the specific core targeted, not all files delivered by the CORE Generator tool are necessary, and some files might not be delivered. The major difference is that some of the Endpoint for PCIe solutions use a 32-bit user datapath, others use a 64-bit datapath, and the PIO design works with both. The width of the datapath depends on the specific core being targeted.

| File | Description |
|---------------------|--------------------------------|
| PIO.v | Top-level design wrapper |
| PIO_INTR_CTRL.v | PIO interrupt controller |
| PIO_EP.v | PIO application module |
| PIO_TO_CTRL.v | PIO turn-off controller module |
| PIO_RX_ENGINE.v | 32-bit Receive engine |
| PIO_TX_ENGINE.v | 32-bit Transmit engine |
| PIO_EP_MEM_ACCESS.v | Endpoint memory access module |
| PIO_EP_MEM.v | Endpoint memory |

Table 9-2: PIO Design File Structure

Three configurations of the PIO design are provided: PIO_64, PIO_128, and PIO_256 with 64-, 128-, and 256-bit AXI4-Stream interfaces, respectively. The PIO configuration generated depends on the selected Endpoint type (that is, Virtex-7 FPGA integrated block, PIPE, PCI Express, and Block Plus) as well as the number of PCI Express lanes and the interface width selected by the user. Table 9-3 identifies the PIO configuration generated based on your selection.

Table 9-3: PIO Configuration

| Core | x1 | x2 | x4 | x8 |
|-------------------------------------|--------|--------------------|--------------------------------|--|
| Virtex-7 FPGA Gen3 Integrated Block | PIO_64 | PIO_64, PIO_128 | PIO_64, PIO_128, PIO_256 | PIO_64, PIO_128 ⁽¹⁾ , PIO_256 |

Notes:

1. The core does not support 128-bit x8 8.0 Gb/s configuration and 500 MHz user clock frequency.

Figure 9-4 shows the various components of the PIO design, which is separated into four main parts: the TX Engine, RX Engine, Memory Access Controller, and Power Management Turn-Off Controller.

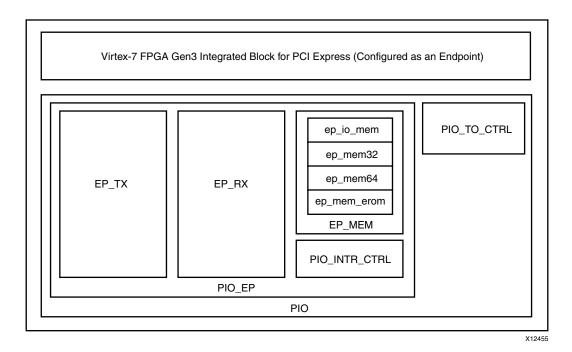


Figure 9-4: PIO Design Components

PIO Operation

PIO Read Transaction

Figure 9-5 depicts a Back-to-Back Memory Read request to the PIO design. The receive engine deasserts m_axis_rx_tready as soon as the first TLP is completely received. The next Read transaction is accepted only after compl_done_o is asserted by the transmit engine, indicating that Completion for the first request was successfully transmitted.

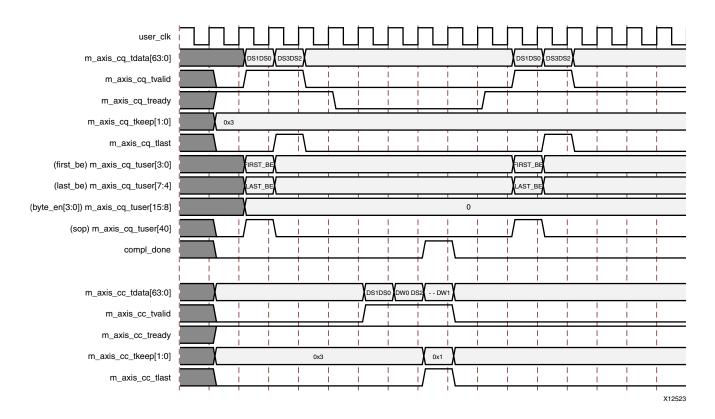


Figure 9-5: Back-to-Back Read Transactions

PIO Write Transaction

Figure 9-6 depicts a back-to-back Memory Write to the PIO design. The next Write transaction is accepted only after wr_busy_o is deasserted by the memory access unit, indicating that data associated with the first request was successfully written to the memory aperture.

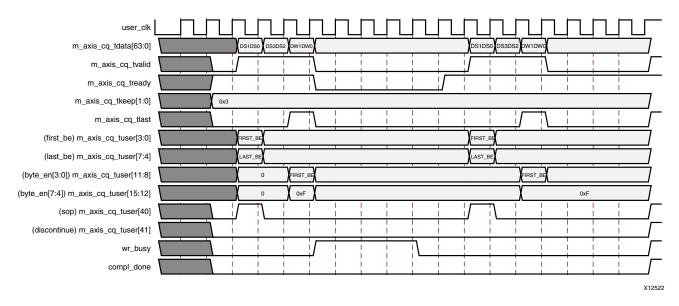


Figure 9-6: Back-to-Back Write Transactions

Device Utilization

Table 9-4 shows the PIO design FPGA resource utilization.

Table 9-4: PIO Design FPGA Resources

| Resources | Utilization |
|------------|-------------|
| LUTs | 300 |
| Flip-Flops | 500 |
| Block RAMs | 4 |

Demonstration Test Bench

Root Port Model Test Bench for Endpoint

The PCI Express Root Port Model is a robust test bench environment that provides a test program interface that can be used with the provided PIO design or with your design. The purpose of the Root Port Model is to provide a source mechanism for generating downstream PCI Express TLP traffic to stimulate the customer design, and a destination mechanism for receiving upstream PCI Express TLP traffic from the customer design in a simulation environment.

Source code for the Root Port Model is included to provide the model for a starting point for the user test bench. All the significant work for initializing the configuration space of the core, creating TLP transactions, generating TLP logs, and providing an interface for creating and verifying tests are complete, allowing the user to dedicate efforts to verifying the correct functionality of the design rather than spending time developing an Endpoint core test bench infrastructure.

The Root Port Model consists of:

- Test Programming Interface (TPI), which allows the user to stimulate the Endpoint device for the PCI Express
- Example tests that illustrate how to use the test program TPI
- Verilog source code for all Root Port Model components, which allow the user to customize the test bench

Figure 9-7 illustrates the illustrates the Root Port Model coupled with the PIO design.

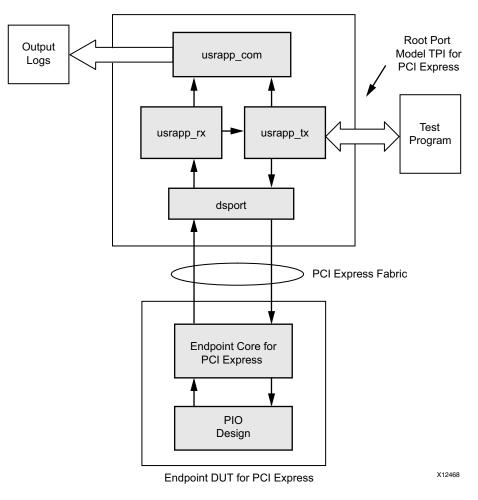


Figure 9-7: Root Port Model and Top-Level Endpoint

Architecture

The Root Port Model consists of these blocks, illustrated in Figure 9-7:

- dsport (Root Port)
- usrapp_tx
- usrapp_rx
- usrapp_com (Verilog only)

The usrapp_tx and usrapp_rx blocks interface with the dsport block for transmission and reception of TLPs to/from the Endpoint Design Under Test (DUT). The Endpoint DUT consists of the Endpoint for PCIe and the PIO design (displayed) or customer design.

The usrapp_tx block sends TLPs to the dsport block for transmission across the PCI Express Link to the Endpoint DUT. In turn, the Endpoint DUT device transmits TLPs across the PCI Express Link to the dsport block, which are subsequently passed to the usrapp_rx block. The dsport and core are responsible for the data link layer and physical link layer processing when communicating across the PCI Express logic. Both usrapp_tx and usrapp_rx utilize the usrapp_com block for shared functions, for example, TLP processing and log file outputting. Transaction sequences or test programs are initiated by the usrapp_tx block to stimulate the Endpoint device's fabric interface. TLP responses from the Endpoint device are received by the usrapp_rx block. Communication between the usrapp_tx and usrapp_rx blocks allow the usrapp_tx block to verify correct behavior and act accordingly when the usrapp_rx block has received TLPs from the Endpoint device.

Simulating the Design

Simulation script files are provided with the model to facilitate simulation with various simulation tools:

- vsim -do simulate_mti.do
- ./simulate_vcs.sh
- ./simulate_ncsim.sh
- ./simulate_isim.sh

The example simulation script file is located in this directory:

<project_dir>/<component_name>/simulation/functional

Instructions for simulating the PIO design using the Root Port Model are provided in Programmed Input/Output: Endpoint Example Design, page 271.

Scaled Simulation Timeouts

The simulation model of the Gen3 Integrated Block for PCIe uses scaled down times during link training to allow for the link to train in a reasonable amount of time during simulation. According to the *PCI Express Specification, rev. 3.0*, there are various timeouts associated with the link training and status state machine (LTSSM) states. The Gen3 Integrated Block scales these timeouts by a factor of 256 in simulation, except in the Recovery Speed_1 LTSSM state, where the timeouts are not scaled.

Test Selection

Table 9-5 describes the tests provided with the Root Port Model, followed by specific sections for VHDL and Verilog test selection.

Table 9-5: Root Port Model Provided Tests

| Test Name | Test in Verilog | Description |
|--------------------|---------------------|---|
| sample_smoke_test0 | Verilog and VHDL | Issues a PCI Type 0 Configuration Read TLP and waits for the completion TLP; then compares the value returned with the expected Device/Vendor ID value. |
| sample_smoke_test1 | Verilog | Performs the same operation as sample_smoke_test0 but makes use of expectation tasks. This test uses two separate test program threads: one thread issues the PCI Type 0 Configuration Read TLP and the second thread issues the Completion with Data TLP expectation task. This test illustrates the form for a parallel test that uses expectation tasks. This test form allows for confirming reception of any TLPs from your design. Additionally, this method can be used to confirm reception of TLPs when ordering is unimportant. |

Verilog Test Selection

The Verilog test model used for the Root Port Model lets the user specify the name of the test to be run as a command line parameter to the simulator.

To change the test to be run, change the value provided to TESTNAME, which is defined in the test files sample_tests1.v and pio_tests.v. This mechanism is used for ModelSim. ISim uses the -testplusarg options to specify TESTNAME, for example: demo_tb.exe -gui -view wave.wcfg -wdb wave_isim -tclbatch isim_cmd.tcl -testplusarg TESTNAME=sample_smoke_test0.

Waveform Dumping

Table 9-6 describes the available simulator waveform dump file formats, provided in the simulator native file format. The same mechanism is used for ModelSim.

| Simulator | Dump File Format |
|---------------------------------|------------------|
| Mentor Graphics ModelSim | .vcd |
| Synopsys VCS or Synopsys VCS_MX | .vpd |
| Cadence IES | .trn |
| Xilinx ISE Simulator (ISim) | .wbd |

Table 9-6: Simulator Dump File Format

Verilog Flow

The Root Port Model provides a mechanism for outputting the simulation waveform to file by specifying the **+dump_all** command line parameter to the simulator.

Output Logging

When a test fails on the example or customer design, the test programmer debugs the offending test case. Typically, the test programmer inspects the wave file for the simulation and cross-reference this to the messages displayed on the standard output. Because this approach can be very time consuming, the Root Port Model offers an output logging mechanism to assist the tester with debugging failing test cases to speed the process.

The Root Port Model creates three output files (tx.dat, rx.dat, and error.dat) during each simulation run. Log files rx.dat and tx.dat each contain a detailed record of every TLP that was received and transmitted, respectively, by the Root Port Model.



TIP: With an understanding of the expected TLP transmission during a specific test case, you can more easily isolate the failure.

The log file error.dat is used in conjunction with the expectation tasks. Test programs that utilize the expectation tasks generate a general error message to standard output. Detailed information about the specific comparison failures that have occurred due to the expectation error is located within error.dat.

Parallel Test Programs

There are two classes of tests are supported by the Root Port Model:

- Sequential tests. Tests that exist within one process and behave similarly to sequential programs. The test depicted in Test Program: pio_writeReadBack_test0, page 286 is an example of a sequential test. Sequential tests are very useful when verifying behavior that have events with a known order.
- Parallel tests. Tests involving more than one process thread. The test sample_smoke_test1 is an example of a parallel test with two process threads. Parallel tests are very useful when verifying that a specific set of events have occurred, however the order of these events are not known.

A typical parallel test uses the form of one command thread and one or more expectation threads. These threads work together to verify a device's functionality. The role of the command thread is to create the necessary TLP transactions that cause the device to receive and generate TLPs. The role of the expectation threads is to verify the reception of an expected TLP. The Root Port Model TPI has a complete set of expectation tasks to be used in conjunction with parallel tests.

Because the example design is a target-only device, only Completion TLPs can be expected by parallel test programs while using the PIO design. However, the full library of expectation tasks can be used for expecting any TLP type when used in conjunction with the customer's design (which can include bus-mastering functionality).

Test Description

The Root Port Model provides a Test Program Interface (TPI). The TPI provides the means to create tests by invoking a series of Verilog tasks. All Root Port Model tests should follow the same six steps:

- 1. Perform conditional comparison of a unique test name
- 2. Set up master timeout in case simulation hangs
- 3. Wait for Reset and link-up
- 4. Initialize the configuration space of the Endpoint
- 5. Transmit and receive TLPs between the Root Port Model and the Endpoint DUT
- 6. Verify that the test succeeded

Test Program: pio_writeReadBack_test0

```
else if(testname == "pio_writeReadBack_test1"
1
2.
       begin
3.
       // This test performs a 32 bit write to a 32 bit Memory space and performs a read back
       TSK SIMULATION TIMEOUT(10050):
4.
       TSK_SYSTEM_INITIALIZATION;
5
6.
       TSK BAR INIT;
      for (ii = 0; ii <= 6; ii = ii + 1) begin
7.
           if (BAR_INIT_P_BAR_ENABLED[ii] > 2'b00) // bar is enabled
8.
9
           case(BAR_INIT_P_BAR_ENABLED[ii])
                   2'b01 : // IO SPACE
10.
11.
                  begin
                      $display("[%t] : NOTHING: to IO 32 Space BAR %x", $realtime, ii);
12.
13.
                  end
14.
                   2'b10 : // MEM 32 SPACE
15.
                     begin
                      display("[%t] : Transmitting TLPs to Memory 32 Space BAR %x",
16.
17.
                                $realtime, ii);
18.
              //-----
              // Event : Memory Write 32 bit TLP
19.
20.
               //-----
21.
                       DATA\_STORE[0] = 8'h04;
                       DATA STORE [1] = 8'h03;
22.
23.
                       DATA STORE [2] = 8'h02;
                       DATA_STORE[3] = 8'h01;
24.
25.
                       P_READ_DATA = 32'hffff_fff; // make sure P_READ_DATA has known initial value
                       TSK_TX_MEMORY_WRITE_32(DEFAULT_TAG, DEFAULT_TC, 10'd1, BAR_INIT_P_BAR[ii][31:0], 4'hF,
26.
       4'hF, 1'b0);
27.
                       TSK_TX_CLK_EAT(10);
                      DEFAULT_TAG = DEFAULT_TAG + 1;
28.
29.
                  //-----
                                   _____
30.
                  // Event : Memory Read 32 bit TLP
31.
                  //---
                                              _____
32.
                       TSK TX MEMORY READ 32(DEFAULT TAG, DEFAULT TC, 10'd1, BAR INIT P BAR[ii][31:0], 4'hF,
       4'hF);
33.
                       TSK_WAIT_FOR_READ_DATA;
                       if (P_READ_DATA != {DATA_STORE[3], DATA_STORE[2], DATA_STORE[1], DATA_STORE[0] })
34.
35.
                         begin
                          $display("[%t] : Test FAILED --- Data Error Mismatch, Write Data %x != Read Data %x",
36.
       $realtime,{DATA_STORE[3], DATA_STORE[2], DATA_STORE[1], DATA_STORE[0]}, P_READ_DATA);
37.
                        end
38
                      else
39.
                       begin
40.
                          $display("[%t] : Test PASSED --- Write Data: %x successfully received", $realtime,
       P READ DATA);
41
                       end
```

Expanding the Root Port Model

The Root Port Model was created to work with the PIO design, and for this reason is tailored to make specific checks and warnings based on the limitations of the PIO design. These checks and warnings are enabled by default when the Root Port Model is generated by the CORE Generator tool. However, these limitations can be disabled so that they do not affect the customer's design.

Because the PIO design was created to support at most one I/O BAR, one Mem64 BAR, and two Mem32 BARs (one of which must be the EROM space), the Root Port Model by default makes a check during device configuration that verifies that the core has been configured to meet this requirement. A violation of this check causes a warning message to be displayed as well as for the offending BAR to be gracefully disabled in the test bench. This check can be disabled by setting the pio_check_design variable to zero in the pci_exp_usrapp_tx.v file.

Root Port Model TPI Task List

The Root Port Model TPI tasks include these tasks, which are further defined in these tables.

- Table 9-7, Test Setup Tasks
- Table 9-8, TLP Tasks
- Table 9-9, BAR Initialization Tasks
- Table 9-10, Example PIO Design Tasks
- Table 9-11, Expectation Tasks

Table 9-7: Test Setup Tasks

| Name | Input(s | 5) | Description |
|---------------------------|-------------|-------|--|
| TSK_SYSTEM_INITIALIZATION | None | | Waits for transaction interface reset and link-up between the Root Port Model and the Endpoint DUT. This task must be invoked prior to the Endpoint core initialization. |
| TSK_USR_DATA_SETUP_SEQ | None | | Initializes global 4096 byte DATA_STORE array entries to sequential values from zero to 4095. |
| TSK_TX_CLK_EAT | clock count | 31:30 | Waits clock_count transaction interface clocks. |
| TSK_SIMULATION_TIMEOUT | timeout | 31:0 | Sets master simulation timeout value in units of transaction interface clocks. This task should be used to ensure that all DUT tests complete. |

Table 9-8: TLP Tasks

| Name | Input(s) | | Description |
|---------------------------------|-----------------------------------|--------------------|---|
| TSK_TX_TYPE0_CONFIGURATION_READ | tag_ reg_addr_ first_dw_be_ | 7:0 11:0 3:0 | Waits for transaction interface reset and link-up between the Root Port Model and the Endpoint DUT. This task must be invoked prior to Endpoint core initialization. |
| TSK_TX_TYPE1_CONFIGURATION_READ | tag_ reg_addr_ first_dw_be_ | 7:0 11:0 3:0 | Sends a Type 1 PCI Express Config Read TLP from Root Port Model to reg_addr_ of Endpoint DUT with tag_ and first_dw_be_ inputs. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |

Table 9-8: TLP Tasks (Cont'd)

| Name | Input(s) | | Description |
|----------------------------------|--|---|---|
| TSK_TX_TYPE0_CONFIGURATION_WRITE | tag_ reg_addr_ reg_data_ first_dw_be_ | 7:0 11:0 31:0 3:0 | Sends a Type 0 PCI Express Config Write TLP from Root Port Model to reg_addr_ of Endpoint DUT with tag_ and first_dw_be_ inputs. Cpl returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_TYPE1_CONFIGURATION_WRITE | tag_ reg_addr_ reg_data_ first_dw_be_ | 7:0 11:0 31:0 3:0 | Sends a Type 1 PCI Express Config Write TLP from Root Port Model to reg_addr_ of Endpoint DUT with tag_ and first_dw_be_ inputs. Cpl returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MEMORY_READ_32 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ | 7:0 2:0 10:0 31:0 3:0 3:0 | Sends a PCI Express Memory Read TLP from Root Port to 32-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MEMORY_READ_64 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ | 7:0 2:0 10:0 63:0 3:0 3:0 | Sends a PCI Express Memory Read TLP from Root Port Model to 64-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MEMORY_WRITE_32 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ ep_ | 7:0 2:0 10:0 31:0 3:0 3:0 - | Sends a PCI Express Memory Write TLP from Root Port Model to 32-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. The global DATA_STORE byte array is used to pass write data to task. |
| TSK_TX_MEMORY_WRITE_64 | tag_ tc_ len_ addr_ last_dw_be_ first_dw_be_ ep_ | 7:0 2:0 10:0 63:0 3:0 3:0 - | Sends a PCI Express Memory Write TLP from Root Port Model to 64-bit memory address addr_ of Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. The global DATA_STORE byte array is used to pass write data to task. |
| TSK_TX_COMPLETION | tag_ tc_ len_ comp_status_ | 7:0 2:0 10:0 2:0 | Sends a PCI Express Completion TLP from Root Port Model to the Endpoint DUT using global COMPLETE_ID_CFG as the completion ID. |

Table 9-8: TLP Tasks (Cont'd)

| Name | Input(s) | | Description |
|------------------------|---|---|--|
| TSK_TX_COMPLETION_DATA | tag_ tc_ len_ byte_count lower_addr comp_status ep_ | 7:0 2:0 10:0 11:0 6:0 2:0 - | Sends a PCI Express Completion with Data TLP from Root Port Model to the Endpoint DUT using global COMPLETE_ID_CFG as the completion ID. The global DATA_STORE byte array is used to pass completion data to task. |
| TSK_TX_MESSAGE | tag_ tc_ len_ data message_rtg message_code | 7:0 2:0 10:0 63:0 2:0 7:0 | Sends a PCI Express Message TLP from Root Port Model to Endpoint DUT. Completion returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_MESSAGE_DATA | tag_ tc_ len_ data message_rtg message_code | 7:0 2:0 10:0 63:0 2:0 7:0 | Sends a PCI Express Message with Data TLP from Root Port Model to Endpoint DUT. The global DATA_STORE byte array is used to pass message data to task. Completion returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_IO_READ | tag_ addr_ first_dw_be_ | 7:0 31:0 3:0 | Sends a PCI Express I/O Read TLP from Root Port Model to I/O address addr_[31:2] of the Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_IO_WRITE | tag_ addr_ first_dw_be_ data | 7:0 31:0 3:0 31:0 | Sends a PCI Express I/O Write TLP from Root Port Model to I/O address addr_[31:2] of the Endpoint DUT. CpID returned from the Endpoint DUT uses the contents of global COMPLETE_ID_CFG as the completion ID. |
| TSK_TX_BAR_READ | bar_index byte_offset tag_ tc_ | 2:0 31:0 7:0 2:0 | Sends a PCI Express one Dword Memory 32, Memory 64, or I/O Read TLP from the Root Port Model to the target address corresponding to offset byte_offset from BAR bar_index of the Endpoint DUT. This task sends the appropriate Read TLP based on how BAR bar_index has been configured during initialization. This task can only be called after TSK_BAR_INIT has successfully completed. CpID returned from the Endpoint DUT use the contents of global COMPLETE_ID_CFG as the completion ID. |

Table 9-8: TLP Tasks (Cont'd)

| Name | Input(s) | | Description |
|------------------------|--|-----------------------------------|--|
| TSK_TX_BAR_WRITE | bar_index byte_offset tag_ tc_ data_ | 2:0 31:0 7:0 2:0 31:0 | Sends a PCI Express one Dword Memory 32, Memory 64, or I/O Write TLP from the Root Port to the target address corresponding to offset byte_offset from BAR bar_index of the Endpoint DUT. This task sends the appropriate Write TLP based on how BAR bar_index has been configured during initialization. This task can only be called after TSK_BAR_INIT has successfully completed. |
| TSK_WAIT_FOR_READ_DATA | None | | Waits for the next completion with data TLP that was sent by the Endpoint DUT. On successful completion, the first Dword of data from the CpID is stored in the global P_READ_DATA. This task should be called immediately following any of the read tasks in the TPI that request Completion with Data TLPs to avoid any race conditions. By default this task locally times out and terminate the simulation after 1000 transaction interface clocks. The global cpld_to_finish can be set to zero so that local time out returns execution to the calling test and does not result in simulation timeout. For this case test programs should check the global cpld_to, which when set to one indicates that this task has timed out and that the contents of P_READ_DATA are invalid. |

Table 9-9:BAR Initialization Tasks

| Name | Input(s) | Description |
|----------------------|----------|---|
| TSK_BAR_INIT | None | Performs a standard sequence of Base Address Register initialization tasks to the Endpoint device using the PCI Express fabric. Performs a scan of the Endpoint's PCI BAR range requirements, performs the necessary memory and I/O space mapping calculations, and finally programs the Endpoint so that it is ready to be accessed. On completion, the user test program can begin memory and I/O transactions to the device. This function displays to standard output a memory and I/O table that details how the Endpoint has been initialized. This task also initializes global variables within the Root Port Model that are available for test program usage. This task should only be called after TSK_SYSTEM_INITIALIZATION. |
| TSK_BAR_SCAN | None | Performs a sequence of PCI Type 0 Configuration Writes and Configuration Reads using the PCI Express logic to determine the memory and I/O requirements for the Endpoint. The task stores this information in the global array BAR_INIT_P_BAR_RANGE[]. This task should only be called after TSK_SYSTEM_INITIALIZATION. |
| TSK_BUILD_PCIE_MAP | None | Performs memory and I/O mapping algorithm and allocates Memory 32, Memory 64, and I/O space based on the Endpoint requirements. This task has been customized to work in conjunction with the limitations of the PIO design and should only be called after completion of TSK_BAR_SCAN. |
| TSK_DISPLAY_PCIE_MAP | None | Displays the memory mapping information of the Endpoint core PCI Base Address Registers. For each BAR, the BAR value, the BAR range, and BAR type is given. This task should only be called after completion of TSK_BUILD_PCIE_MAP. |

Table 9-10: Example PIO Design Tasks

| Name | Input(s) | | Description |
|------------------------|---------------|--|---|
| TSK_TX_READBACK_CONFIG | None | | Performs a sequence of PCI Type 0 Configuration Reads to the Endpoint device's Base Address Registers, PCI Command Register, and PCIe Device Control Register using the PCI Express logic. This task should only be called after TSK_SYSTEM_INITIALIZATION. |
| TSK_MEM_TEST_DATA_BUS | bar_index 2:0 | | Tests whether the PIO design FPGA block RAM data bus interface is correctly connected by performing a 32-bit walking ones data test to the I/O or memory address pointed to by the input bar_index. For an exhaustive test, this task should be called four times, once for each block RAM used in the PIO design. |

| | Table 9-10: | Example P | IO Design | Tasks <i>(Cont'd)</i> |
|--|-------------|-----------|-----------|-----------------------|
|--|-------------|-----------|-----------|-----------------------|

| Name | Input(| s) | Description |
|-----------------------|---------------------|-------------|--|
| TSK_MEM_TEST_ADDR_BUS | bar_index nBytes | 2:0 31:0 | Tests whether the PIO design FPGA block RAM address bus interface is accurately connected by performing a walking ones address test starting at the I/O or memory address pointed to by the input bar_index. For an exhaustive test, this task should be called four times, once for each block RAM used in the PIO design. Additionally, the nBytes input should specify the entire size of the individual block RAM. |
| TSK_MEM_TEST_DEVICE | bar_index nBytes | 2:0 31:0 | Tests the integrity of each bit of the PIO design FPGA block RAM by performing an increment/decrement test on all bits starting at the block RAM pointed to by the input bar_index with the range specified by input nBytes. For an exhaustive test, this task should be called four times, once for each block RAM used in the PIO design. Additionally, the nBytes input should specify the entire size of the individual block RAM. |
| TSK_RESET | Reset | 0 | Initiates PERSTn. Forces the PERSTn signal to assert the reset. Use TSK_RESET (1'b1) to assert the reset and TSK_RESET (1'b0) to release the reset signal. |
| TSK_MALFORMED | malformed _bits | 7:0 | Control bits for creating malformed TLPs: 0001: Generate Malformed TLP for I/O Requests and Configuration Requests called immediately after this task 0010: Generate Malformed Completion TLPs for Memory Read requests received at the Root Port |

Table 9-11:Expectation Tasks

| Name | Input(| Input(s) | | Description |
|--------------------|--|---|------------------|--|
| TSK_EXPECT_CPLD | traffic_class td ep attr length completer_id completer_status bcm byte_count requester_id tag address_low | 2:0 - 1:0 10:0 15:0 2:0 - 11:0 15:0 7:0 6:0 | Expect status | Waits for a Completion with Data TLP that matches traffic_class, td, ep, attr, length, and payload. Returns a 1 on successful completion; 0 otherwise. |
| TSK_EXPECT_CPL | traffic_class td ep attr completer_id completer_status bcm byte_count requester_id tag address_low | 2:0 - - 1:0 15:0 2:0 - 11:0 15:0 7:0 6:0 | Expect status | Waits for a Completion without Data TLP that matches traffic_class, td, ep, attr, and length. Returns a 1 on successful completion; 0 otherwise. |
| TSK_EXPECT_MEMRD | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - - 1:0 10:0 15:0 7:0 3:0 3:0 29:0 | Expect status | Waits for a 32-bit Address Memory Read TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |
| TSK_EXPECT_MEMRD64 | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - - 1:0 10:0 15:0 7:0 3:0 3:0 61:0 | Expect status | Waits for a 64-bit Address Memory Read TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. |

Table 9-11: Expectation Tasks (Cont'd)

| Name | Input | (s) | Output | Description | |
|--------------------|--|--|------------------|---|--|
| TSK_EXPECT_MEMWR | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - - 1:0 10:0 15:0 7:0 3:0 3:0 29:0 | Expect status | Waits for a 32-bit Address Memory Write TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. | |
| TSK_EXPECT_MEMWR64 | traffic_class td ep attr length requester_id tag last_dw_be first_dw_be address | 2:0 - - 1:0 10:0 15:0 7:0 3:0 3:0 3:0 61:0 | Expect status | Waits for a 64-bit Address Memory Write TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. | |
| TSK_EXPECT_IOWR | td ep requester_id tag first_dw_be address data | - 15:0 7:0 3:0 31:0 31:0 | Expect status | Waits for an I/O Write TLP with matching header fields. Returns a 1 on successful completion; 0 otherwise. This task can only be used in conjunction with Bus Master designs. | |

Endpoint Model Test Bench for Root Port

The Endpoint model test bench for the Virtex-7 FPGAs Integrated Block for PCI Express in Root Port configuration is a simple example test bench that connects the Configurator example design and the PCI Express Endpoint model allowing the two to operate like two devices in a physical system. As the Configurator example design consists of logic that initializes itself and generates and consumes bus traffic, the example test bench only implements logic to monitor the operation of the system and terminate the simulation.

The Endpoint model test bench consists of:

- Verilog or VHDL source code for all Endpoint model components
- PIO slave design

Architecture

The Endpoint model consists of these blocks:

- PCI Express Endpoint (Gen3 Integrated Block for PCIe in Endpoint configuration) model.
- PIO slave design, consisting of:
 - PIO_RX_ENGINE
 - PIO_TX_ENGINE
 - PIO_EP_MEM
 - PIO_TO_CTRL

The PIO_RX_ENGINE and PIO_TX_ENGINE blocks interface with the ep block for reception and transmission of TLPs from/to the Root Port Design Under Test (DUT). The Root Port DUT consists of the Integrated Block for PCI Express configured as a Root Port and the Configurator Example Design, which consists of a Configurator block and a PIO Master design, or customer design.

The PIO slave design is described in detail in Programmed Input/Output: Endpoint Example Design, page 271.

Simulating the Design

A simulation script file is provided with the model to facilitate simulation with the Mentor Graphics ModelSim simulator:

• simulate_mti.do

The example simulation script files are located in this directory:

```
<project_dir>/<component_name>/simulation/functional
```

Note: For Cadence IES users, the work construct must be manually inserted into the cds.lib file:

DEFINE WORK WORK.

Scaled Simulation Timeouts

The simulation model of the Gen3 Integrated Block for PCIe uses scaled down times during link training to allow for the link to train in a reasonable amount of time during simulation. According to the *PCI Express Specification, rev. 3.0*, there are various timeouts associated with the link training and status state machine (LTSSM) states. The Gen3 Integrated Block for PCIe scales these timeouts by a factor of 256 in simulation, except in the Recovery Speed_1 LTSSM state, where the timeouts are not scaled.

Waveform Dumping

Table 9-12 describes the available simulator waveform dump file format, which is provided in the simulator native file format.

| Table 9-12: | Simulator | Dump | File | Format |
|-------------|-----------|------|------|--------|
|-------------|-----------|------|------|--------|

| Simulator | Dump File Format |
|-----------|------------------|
| ModelSim | .vcd |

The Endpoint model test bench provides a mechanism for outputting the simulation waveform to file by specifying the +dump_all command line parameter to the simulator.

For example, the script file simulate_ncsim.sh (used to start the Cadence IES simulator) can indicate to the Endpoint model that the waveform should be saved to a file using this command line:

ncsim work.boardx01 +dump_all

Output Logging

The test bench outputs messages, captured in the simulation log, indicating the time at which these occur:

- user_reset deasserted
- user_lnk_up asserted
- cfg_done asserted by the Configurator
- pio_test_finished asserted by the PIO Master
- Simulation Timeout (if pio_test_finished or pio_test_failed never asserted)

PIPE MODE Simulation

The PIPE Simulation mode allows you to run the simulations without GT block, which speeds up simulations.

To run the simulations using the PIPE interface to speed up the simulation, generate the core using the **Enable PIPE simulation** option, as shown on the Basic page of the Customize IP dialog box described in Customizing the Core using the Vivado IP Catalog. With this option, the PIPE interface of the core top module in the PCIe example design is connected to PIPE interface of the model.



IMPORTANT: A new $pcie3_7x_v1_3_gt_top_pipe.v$ file is created in the simulation directory, and the file replaces the GT block for PIPE mode simulation.

To run simulations using GT block with the same core, define ENABLE_GT during run time so that the original GT block is instantiated in the core top module and simulations are run using the GT block. Comments are included in the simulation scripts to define which parameters need to be passed in order to run the simulations using GT block.



TIP: Implementation is always run with the GT block. The PIPE mode is only for simulation.

Implementation

Implementing the Example Design

After generating the core, the netlists and the example design can be processed using the Xilinx implementation tools. The generated output files include scripts to assist in running the Xilinx software.

To implement the example design:

Open a command prompt or terminal window and type:

Windows

```
ms-dos> cd <project_dir>\<component_name>\implement
ms-dos> implement.bat
```

Linux

```
% cd <project_dir>/<component_name>/implement
% ./implement.sh
```

These commands execute a script that synthesizes, builds, maps, and place-and-routes the example design, and then generates a post-par simulation model for use in timing simulation. The resulting files are placed in the results directory and execute these processes:

- 1. Removes data files from the previous runs.
- 2. Synthesizes the example design using XST based on the flow settings in the Project Options window.

- 3. ngdbuild: Builds a Xilinx design database for the example design.
 - Inputs:

Part-Package-Speed Grade selection:

XC7V485T-FFG1157-3

Example design UCF:

xilinx_pcie_2_1_ep_7x_01_lane_gen1_xc7v485t-ffg1157-3-PCIE_X0Y0.ucf

- 4. map: Maps design to the selected FPGA using the constraints provided.
- 5. par: Places cells onto FPGA resources and routes connectivity.
- 6. trce: Performs static timing analysis on design using constraints specified.
- 7. netgen: Generates a logical Verilog HDL representation of the design and an SDF file for post-layout verification.

These FPGA implementation related files are generated in the results directory:

- routed.v Verilog functional model.
- routed.sdf Timing model Standard Delay File.
- mapped.mrp Xilinx map report.
- routed.par
 Xilinx place and route report.
- routed.twr Xilinx timing analysis report.

Simulation

Simulating the Example Design

The example design provides a quick way to simulate and observe the behavior of the core.

Endpoint Configuration

The simulation environment provided with the LogiCORE[™] IP Gen3 Integrated Block for PCIe core in Endpoint configuration performs simple memory access tests on the PIO example design. Transactions are generated by the Root Port Model and responded to by the PIO example design.

- PCI Express Transaction Layer Packets (TLPs) are generated by the test bench transmit User Application (pci_exp_usrapp_tx). As it transmits TLPs, it also generates a log file, tx.dat.
- PCI Express TLPs are received by the test bench receive User Application (pci_exp_usrapp_rx). As the User Application receives the TLPs, it generates a log file, rx.dat.

For more information about the test bench, see Root Port Model Test Bench for Endpoint, page 280.

Setting Up for Simulation

To run the gate-level simulation, the Xilinx Simulation Libraries must be compiled for the user system. See the Compiling Xilinx Simulation Libraries (COMPXLIB) in the Xilinx ISE Synthesis and Verification Design Guide and the Xilinx ISE Software Manuals and Help. Documents can be downloaded from www.xilinx.com/support/software_manuals.htm.

Simulator Requirements

Virtex-7 FPGA designs require a Verilog LRM-IEEE 1364-2005 encryption-compliant simulator. This core supports these simulators:

- Mentor Graphics ModelSim
- Cadence IES
- Synopsys VCS
- Xilinx ISim

Running the Simulation

The simulation scripts provided with the example design support pre-implementation (RTL) simulation. The existing test bench can be used to simulate with a post-implementation version of the example design.

The pre-implementation simulation consists of these components:

- Verilog model of the test bench
- Verilog RTL example design
- The Verilog model of the Gen3 Integrated Block for PCIe
- 1. To run the simulation, go to this directory:

<project_dir>/<component_name>/simulation/functional

2. Launch the simulator and run the script that corresponds to the user simulation tool:

ModelSim > do simulate_mti.do

vsim -do simulate_mti.do

- ./simulate_ncsim.sh
- ./simulate_vcs.sh
- ./simulate_isim.sh



Example Design and Model Test Bench for Root Port Configuration

Note: For information about the availability of Root Port functionality, contact Xilinx. See Technical Support, page 382.

Configurator Example Design

The Configurator example design, included with the Xilinx® Virtex-7 FPGA Gen3 Integrated Block for PCI Express® in Root Port configuration generated by the CORE Generator[™] tool, is a synthesizable, lightweight design that demonstrates the minimum setup required for the integrated block in Root Port configuration to begin application-level transactions with an Endpoint.

System Overview

PCI Express devices require setup after power-on, before devices in the system can begin application specific communication with each other. Minimally, two devices connected via a PCI Express Link must have their Configuration spaces initialized and be enumerated to communicate.

Root Ports facilitate PCI Express enumeration and configuration by sending Configuration Read (CfgRd) and Write (CfgWr) TLPs to the downstream devices such as Endpoints and Switches to set up the configuration spaces of those devices. When this process is complete, higher-level interactions, such as Memory Reads (MemRd TLPs) and Writes (MemWr TLPs), can occur within the PCI Express System.

The Configurator example design described herein performs the configuration transactions required to enumerate and configure the Configuration space of a single connected PCI Express Endpoint and allow application-specific interactions to occur.

Configurator Example Design Hardware

The Configurator example design consists of four high-level blocks:

• Root Port: The Gen3 Integrated Block PCIe in Root Port configuration.

- Configurator Block: Logical block which interacts with the configuration space of a PCI Express Endpoint device connected to the Root Port.
- Configurator ROM: Read-only memory that sources configuration transactions to the Configurator Block.
- PIO Master: Logical block which interacts with the user logic connected to the Endpoint by exchanging data packets and checking the validity of the received data. The data packets are limited to a single DWORD and represent the type of traffic that would be generated by a CPU.

Note: The Configurator Block and Configurator ROM, and Root Port are logically grouped in the RTL code within a wrapper file called the Configurator Wrapper.

The Configurator example design, as delivered, is designed to be used with the PIO Slave example included with Xilinx Endpoint cores and described in Chapter 9, Example Design and Model Test Bench for Endpoint Configuration. The PIO Master is useful for simple bring-up and debugging, and is an example of how to interact with the Configurator Wrapper. The Configurator example design can be modified to be used with other Endpoints.

Figure 10-1 shows the various components of the Configurator example design.

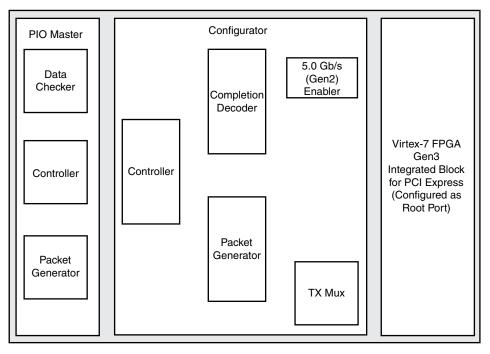


Figure 10-1: **Configurator Example Design Components**

Figure 10-2 shows how the blocks are connected in an overall system view.

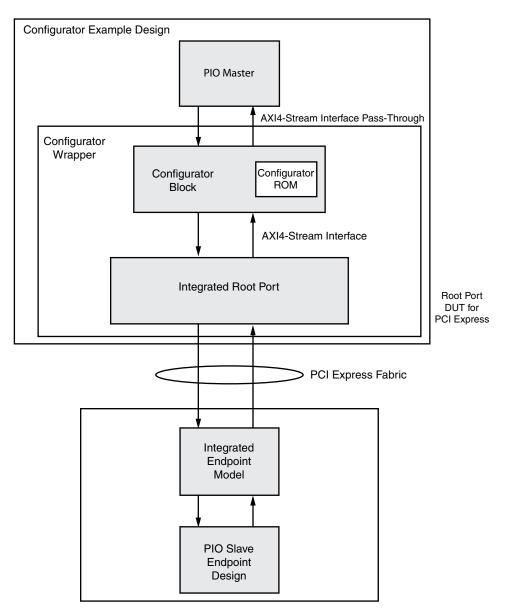


Figure 10-2: Configurator Example Design

Configurator Block

The Configurator Block generates CfgRd and CfgWr TLPs and presents them to the AXI4-Stream interface of the integrated block in Root Port configuration. The TLPs that the Configurator Block generates are determined by the contents of the Configurator ROM.

The generated configuration traffic is predetermined by the designer to address their particular system requirements. The configuration traffic is encoded in a memory-initialization file (the Configurator ROM) which is synthesized as part of the Configurator. The Configurator Block and the attached Configurator ROM is intended to be usable a part of a real-world embedded design.

The Configurator Block steps through the Configuration ROM file and sends the TLPs specified therein. Supported TLP types are Message, Message w/Data, Configuration Write (Type 0), and Configuration Read (Type 0). For the Configuration packets, the Configurator Block waits for a Completion to be returned before transmitting the next TLP. If the Completion TLP fields do not match the expected values, PCI Express configuration fails. However, the Data field of Completion TLPs is ignored and not checked

Note: There is no completion timeout mechanism in the Configurator Block, so if no Completion is returned, the Configurator Block waits forever.

The Configurator Block has these parameters, which can be altered by the user:

- **TCQ**: Clock-to-out delay modeled by all registers in design.
- **EXTRA_PIPELINE**: Controls insertion of an extra pipeline stage on the Receive AXI4-Stream interface for timing.
- **ROM_FILE**: File name containing configuration steps to perform.
- **ROM_SIZE**: Number of lines in ROM_FILE containing data (equals number of TLPs to send/2).
- **REQUESTER_ID**: Value for the Requester ID field in outgoing TLPs.

When the Configurator Block design is used, all TLP traffic must pass through the Configurator Block. The user design is responsible for asserting the start_config input (for one clock cycle) to initiate the configuration process when user_lnk_up has been asserted by the core. Following start_config, the Configurator Block performs whatever configuration steps have been specified in the Configuration ROM. During configuration, the Configurator Block controls the core's AXI4-Stream interface. Following configuration, all AXI4-Stream traffic is routed to/from the User Application, which in the case of this example design is the PIO Master. The end of configuration is signaled by the assertion of finished_config. If configuration is unsuccessful for some reason, failed_config is also asserted.

If used in a system that supports PCIe v2.1 5.0 Gb/s links, the Configurator Block begins its process by attempting to up-train the link from 2.5 Gb/s to 5.0 Gb/s. This feature is enabled depending on the LINK_CAP_MAX_LINK_SPEED parameter on the Configurator Wrapper.

The Configurator does not support the user throttling received data on the Receive AXI4-Stream interface. Because of this, the Root Port inputs which control throttling are not included on the Configurator Wrapper. These signals are m_axis_rx_tready and rx_np_ok. This is a limitation of the Configurator Example Design and not of the Integrated Block for PCI Express in Root Port configuration. This means that the user design interfacing with the Configurator Example Design must be able to accept received data at line rate.

Configurator ROM

The Configurator ROM stores the necessary configuration transactions to configure a PCI Express Endpoint. This ROM interfaces with the Configurator Block to send these transactions over the PCI Express link.

The example ROM file included with this design shows the operations needed to configure a Gen3 Integrated Endpoint Block for PCIe and PIO Example Design.

The Configurator ROM can be customized for other Endpoints and PCI Express system topologies. The unique set of configuration transactions required depends on the Endpoint that interacts with the Root Port. This information can be obtained from the documentation provided with the Endpoint.

The ROM file follows the format specified in the Verilog specification (IEEE 1364-2001) section 17.2.8, which describes using the \$readmemb function to pre-load data into a RAM or ROM. Verilog-style comments are allowed.

The file is read by the simulator or synthesis tool and each memory value encountered is used as a single location in memory. Digits can be separated by an underscore character (_) for clarity without constituting a new location.

Each configuration transaction specified uses two adjacent memory locations—the first location specifies the header fields, while the second location specifies the 32-bit data payload. (For CfgRd TLPs and Messages without data, the data location is unused but still present.) In other words, header fields are on even addresses, while data payloads are on odd addresses.

For headers, Messages and CfgRd/CfgWr TLPs use different fields. For all TLPs, two bits specify the TLP type. For Messages, Message Routing and Message Code are specified. For CfgRd/CfgWr TLPs, Function Number, Register Number, and first DWORD Byte-Enable are specified. The specific bit layout is shown in the example ROM file.

PIO Master

The PIO Master demonstrates how a user-application design might interact with the Configurator Block. It directs the Configurator Block to bring up the link partner at the appropriate time, and then (after successful bring-up) generates and consumes bus traffic. The PIO Master performs writes and reads across the PCI Express Link to the PIO Slave Example Design (from the Endpoint core) to confirm basic operation of the link and the Endpoint.

The PIO Master waits until user_lnk_up is asserted by the Root Port. It then asserts start_config to the Configurator Block. When the Configurator Block asserts finished_config, the PIO Master writes and reads to/from each BAR in the PIO Slave design. If the readback data matches what was written, the PIO Master asserts its pio_test_finished_output. If there is a data mismatch or the Configurator Block fails to

configure the Endpoint, the PIO Master asserts its pio_test_failed output. The PIO Master's operation can be restarted by asserting its pio_test_restart input for one clock cycle.

Configurator File Structure

Table 10-1 defines the Configurator example design file structure.

| File | Description |
|----------------------------|--|
| xilinx_pcie_2_1_rport_7x.v | Top-level wrapper file for Configurator example design |
| cgator_wrapper.v | Wrapper for Configurator and Root Port |
| cgator.v | Wrapper for Configurator sub-blocks |
| cgator_cpl_decoder.v | Completion decoder |
| cgator_pkt_generator.v | Configuration TLP generator |
| cgator_tx_mux.v | Transmit AXI4-Stream muxing logic |
| cgator_gen2_enabler.v | 5.0 Gb/s directed speed change module |
| cgator_controller.v | Configurator transmit engine |
| cgator_cfg_rom.data | Configurator ROM file |
| pio_master.v | Wrapper for PIO Master |
| pio_master_controller.v | TX and RX Engine for PIO Master |
| pio_master_checker.v | Checks incoming User-Application Completion TLPs |
| pio_master_pkt_generator.v | Generates User-Application TLPs |

Table 10-1: Example Design File Structure

The hierarchy of the Configurator example design is:

- xilinx_pcie_2_1_rport_7x
 - cgator_wrapper
 - pcie_2_1_rport_7x (in the source directory)
 This directory contains all the source files for the Integrated Block for PCI Express in Root Port Configuration.
 - cgator
 - cgator_cpl_decoder
 - cgator_pkt_generator
 - cgator_tx_mux
 - cgator_gen2_enabler
 - cgator_controller
 This directory contains <cgator_cfg_rom.data> (specified by ROM_FILE)*

- pio_master
 - pio_master_controller
 - pio_master_checker
 - pio_master_pkt_generator

Note: cgator_cfg_rom.data is the default name of the ROM data file. You can override this by changing the value of the ROM_FILE parameter.

Configurator Example Design Summary

The Configurator example design is a synthesizable design that demonstrates the capabilities of the Gen3 Integrated Block for PCIe when configured as a Root Port. The example is provided via the CORE Generator tool and uses the Endpoint PIO example as a target for PCI Express enumeration and configuration. The design can be modified to target other Endpoints by changing the contents of a ROM file.

Endpoint Model Test Bench for Root Port

The Endpoint model test bench for the Gen3 Integrated Block for PCIe in Root Port configuration is a simple example test bench that connects the Configurator example design and the PCI Express Endpoint model allowing the two to operate like two devices in a physical system. As the Configurator example design consists of logic that initializes itself and generates and consumes bus traffic, the example test bench only implements logic to monitor the operation of the system and terminate the simulation.

The Endpoint model test bench consists of:

- Verilog or VHDL source code for all Endpoint model components
- PIO slave design

Figure 10-2, page 303 illustrates the Endpoint model coupled with the Configurator example design.

Architecture

The Endpoint model consists of these blocks:

- PCI Express Endpoint (Gen3 Integrated Block for PCIe in Endpoint configuration) model.
- PIO slave design, consisting of:
 - PIO_RX_ENGINE
 - PIO_TX_ENGINE

- PIO_EP_MEM
- PIO_TO_CTRL

The PIO_RX_ENGINE and PIO_TX_ENGINE blocks interface with the ep block for reception and transmission of TLPs from/to the Root Port Design Under Test (DUT). The Root Port DUT consists of the Integrated Block for PCI Express configured as a Root Port and the Configurator Example Design, which consists of a Configurator block and a PIO Master design, or customer design.

The PIO slave design is described in detail in Programmed Input/Output: Endpoint Example Design, page 271.

Simulating the Design

Three simulation script files are provided with the model to facilitate simulation with Synopsys VCS and VCS MX, Cadence IES, and Mentor Graphics ModelSim simulators:

- simulate_vcs.sh (Verilog only)
- simulate_ncsim.sh (Verilog only)
- simulate_mti.do

The example simulation script files are located in this directory:

<project_dir>/<component_name>/simulation/functional



TIP: For Cadence IES users, the work construct must be manually inserted into the cds.lib file: DEFINE WORK WORK.

Scaled Simulation Timeouts

The simulation model of the Gen3 Integrated Block for PCIe uses scaled down times during link training to allow for the link to train in a reasonable amount of time during simulation. According to the *PCI Express Specification, rev. 2.1* [Ref 2], there are various timeouts associated with the link training and status state machine (LTSSM) states. The Gen3 Integrated Block for PCIe scales these timeouts by a factor of 256 in simulation, except in the Recovery Speed_1 LTSSM state, where the timeouts are not scaled.

Waveform Dumping

Table 10-2 describes the available simulator waveform dump file formats, each of which is provided in the simulators native file format. The same mechanism is used for VCS and ModelSim.

| Simulator | Dump File Format | | | |
|-------------------------|------------------|--|--|--|
| Synopsys VCS and VCS MX | .vpd | | | |
| ModelSim | .vcd | | | |
| Cadence IES | .trn | | | |

Table 10-2: Simulator Dump File Format

The Endpoint model test bench provides a mechanism for outputting the simulation waveform to file by specifying the +dump_all command line parameter to the simulator.

For example, the script file simulate_ncsim.sh (used to start the Cadence IES simulator) can indicate to the Endpoint model that the waveform should be saved to a file using this command line:

ncsim work.boardx01 +dump_all

Output Logging

The test bench outputs messages, captured in the simulation log, indicating the time at which these occur:

- user_reset deasserted
- user_lnk_up asserted
- cfg_done asserted by the Configurator
- pio_test_finished asserted by the PIO Master
- Simulation Timeout (if pio_test_finished or pio_test_failed never asserted)



SECTION IV: APPENDICES

Migrating

Receiver Buffer Completion Space

Debugging

Attributes

Additional Resources



Appendix A

Migrating

For information on migrating to the Vivado[™] Design Suite, see UG911, *Vivado Design Suite Migration Methodology Guide* [Ref 4].



Receiver Buffer Completion Space

Table B-1 defines the completion space reserved in the receiver buffer by the core. The values differ depending on the different Capability Max Payload Size settings of the core and the performance level that you select. If you decide to not have the TLP Digests (ECRC) removed from the incoming packet stream, the TLP Digests (ECRC) must be accounted for as part of the data payload. Values are credits, expressed in Hexadecimal.

| Table B-1: | Buffer Allocation | |
|------------|-------------------|--|
| | | |

. ..

..

| Capabilities Max | Performance Level: Good | | Performance Level: Extreme | |
|------------------|-------------------------|------------------------|----------------------------|------------------------|
| Payload Size | Completion Header | Completion Data | Completion Header | Completion Data |
| 128,256,512 | 20H | 1F0H | 20H | ЗЕОН |
| 1024 | 20H | ЗЕОН | 20H | 3E0H |



Appendix C

Debugging

This appendix provides information on using resources available on the Xilinx Support website, available debug tools, and a step-by-step process for debugging designs that use the Virtex®-7 FPGA Gen3 Integrated Block for PCI Express®. This appendix uses flow diagrams to guide you through the debug process.

The following information is found in this appendix:

- Finding Help on Xilinx.com
- Contacting Xilinx Technical Support
- Debug Tools
- Hardware Debug
- Simulation Debug

See Solution Centers in Appendix E for information helpful to the debugging progress.

Finding Help on Xilinx.com

To help in the design and debug process when using the Virtex-7 FPGA, the Xilinx Support webpage (<u>www.xilinx.com/support</u>) contains key resources such as Product documentation, Release Notes, Answer Records, and links to opening a Technical Support case.

Documentation

This Product Guide is the main document associated with the Gen3 Integrated Block for PCIe. This guide along with documentation related to all products that aid in the design process can be found on the Xilinx Support webpage. Documentation is sorted by product family at the main support page or by solution at the Documentation Center.

To see the available documentation by device family:

Navigate to <u>www.xilinx.com/support</u>.

To see the available documentation by solution:

XILINX,

- Navigate to <u>www.xilinx.com/support</u>.
- Select the **Documentation** tab located at the top of the webpage.
- This is the Documentation Center where Xilinx documentation is sorted by Devices, Boards, IP, Design Tools, Doc Type, and Topic.

Release Notes and Known Issues

Known issues for all cores, including the Gen3 Integrated Block for PCIe, are described in the IP Release Notes Guide.

Answer Records

Answer Records include information on commonly encountered problems, helpful information on how to resolve these problems, and any known issues with a product. Answer Records are created and maintained daily to provide you with up-to-date information on Xilinx products. Answer Records can be found by searching the Answers Database.

To use the Answers Database Search:

- Navigate to <u>www.xilinx.com/support</u>. The Answers Database Search is located at the top of this webpage.
- Enter keywords in the provided search field and select **Search**.
 - Examples of searchable keywords are product names, error messages, or a generic summary of the issue encountered.
 - To see all answer records directly related to the Gen3 Integrated Block for PCIe, search for the phrase "Virtex-7 FPGA Gen3 Integrated Block for PCI Express."

Contacting Xilinx Technical Support

Xilinx provides premier technical support for customers encountering issues that requires additional assistance.

To contact Technical Support:

- Navigate to <u>www.xilinx.com/support</u>.
- Open a WebCase by selecting the WebCase link located under **Support Quick Links**.

When opening a WebCase, include:

• Target FPGA including package and speed grade

- All applicable ISE® Design Suite, synthesis (if not XST), and simulator software versions
- The XCO file created during generation of the LogiCORE[™] IP wrapper
 - This file is located in the directory targeted for the CORE Generator™ tool project

Additional files might be required based on the specific issue. See the relevant sections in this debug guide for further information on specific files to include with the WebCase.

Debug Tools

There are many tools available to debug PCI Express design issues. This section indicates which tools are useful for debugging the various situations encountered.

Example Design

Xilinx® Endpoint for PCI Express products come with a synthesizable back-end application called the PIO design that has been tested and is proven to be interoperable in available systems. The design appropriately handles all incoming one Endpoint read and write transactions. It returns completions for non-posted transactions and updates the target memory space for writes. For more information, see Programmed Input/Output: Endpoint Example Design, page 271.

ChipScope Pro Tool

The ChipScope[™] Pro tool inserts logic analyzer, bus analyzer, and virtual I/O software cores directly into the user design. The ChipScope Pro tool allows the user to set trigger conditions to capture application and Integrated Block port signals in hardware. Captured signals can then be analyzed through the ChipScope Pro Logic Analyzer tool. For detailed information on the ChipScope Pro tool, visit <u>www.xilinx.com/chipscope</u>.

Link Analyzers

Third-party link analyzers show link traffic in a graphical or text format. Lecroy, Agilent, and Vmetro are companies that make common analyzers available today. These tools greatly assist in debugging link issues and allow users to capture data which Xilinx support representatives can view to assist in interpreting link behavior.

Third-Party Software Tools

This section describes third-party software tools that can be useful in debugging.

LSPCI (Linux)

LSPCI is available on Linux platforms and allows users to view the PCI Express device configuration space. LSPCI is usually found in the /sbin directory. LSPCI displays a list of devices on the PCI buses in the system. See the LSPCI manual for all command options. Some useful commands for debugging include:

lspci -x -d [<vendor>]:[<device>]

This displays the first 64 bytes of configuration space in hexadecimal form for the device with vendor and device ID specified (omit the -d option to display information for all devices). The default Vendor/Device ID for Xilinx cores is 10EE:6012. Here is a sample of a read of the configuration space of a Xilinx device:

Included in this section of the configuration space are the Device ID, Vendor ID, Class Code, Status and Command, and Base Address Registers.

lspci -xxxx -d [<vendor>]:[<device>]

This displays the extended configuration space of the device. It can be useful to read the extended configuration space on the root and look for the Advanced Error Reporting (AER) registers. These registers provide more information on why the device has flagged an error (for example, it might show that a correctable error was issued because of a replay timer timeout).

• lspci -k

Shows kernel drivers handling each device and kernel modules capable of handling it (works with kernel 2.6 or later).

PCItree (Windows)

PCItree can be downloaded at <u>www.pcitree.de</u> and allows the user to view the PCI Express device configuration space and perform one DWORD memory writes and reads to the aperture.

The configuration space is displayed by default in the lower right corner when the device is selected, as shown in Figure C-1.

| direct select: bus: dev: func: | how INT routing | highest | EXIT |
|--|--|---|---|
| · · · · · · · · · · · · · · · · · · · | show Nem Nap | busnr: S | About |
| ⇒ host CPU 0.00.0 Host/PCI: Bridge D 0.00.0 C→1 (1) PCI/PCI: Bri 1.00.0 C→1 (1) PCI/PCI: Bri 1.00.0 Cther: Benory Co 0.02.0 VGA: PC Compatible 0.26.0 Universal Host Con 0.26.1 Universal Host Con 0.26.7 c. serial hus Devi 0.27.0 c. Nultimedia 808 0.26.0 C→4 (4) PCI/PCI: Bri 0.26.4 C→3 (3) PCI/PCI: Bri 0.28.5 C→2 (2) PCI/PCI: Bri 0.29.0 Universal Host Con 0.29.1 Universal Host Con 0.29.1 Universal Host Con 0.29.7 c. serial hus Devi 0.29.7 c. serial bus Devi 0.29.7 Con Serial Bost Con 0.29.7 Con Serial Bo | VID: XIOKE XILINX COPP DID: X506F no device name four SubVID: X10EE XILINX SubID: X506F no-name rev.: X00 XB<-INTA edit ConfReg: Nr of Conf edit ConfReg Vrite ConfReg Vrite ConfReg Config Space Dump: (type 1 xs) S06F 10EE <00 : DID VID 0010 0007 <04 : Stat Cad 0580 00000 <08 : BaseClass SubC 0000 0008 <0C : BIST Header Las F74F F800 <10 : BAR 0 sem 32b1 | | me found no (B<-INTA# of ConfRegs (E 16 C 64 use BIOS int fresh dump: 1 XS) of so SubClass sader LatTim |
| 0.31.2 c. Mass Storage Co 0.31.3 SHBus: Serial Bus 0.31.5 c. Mass Storage Co | 0000 0000 <20 | : BAR 4 : BAR 5 : Cardbus : SubID 2 : Exp_ROM : reserve : reserve | SubVendorID 1 BAR d d |

Figure C-1: PCItree with Read of Configuration Space

HWDIRECT (Windows)

HWDIRECT can be purchased at <u>www.eprotek.com</u> and allows the user to view the PCI Express device configuration space as well as the extended configuration space (including the AER registers on the root).

| and the part of the second sec | - | | |
|--|---|--|--|
| 881 × 10 1 0 | | | |
| ACTACING 20 (Process) Action 10 (Action 10 (Ac | | | |

Figure C-2: HWDIRECT with Read of Configuration Space

PCI-SIG Software Suites

PCI-SIG® software suites such as PCIE-CV can be used to test compliance with the specification. This software can be downloaded at <u>www.pcisig.com</u>.

Hardware Debug

Hardware issues can range from device recognition issues to problems seen after hours of testing. This section provides debug flow diagrams for some of the most common issues experienced by users. Endpoints that are shaded gray indicate that more information can be found in sections after Figure C-3.

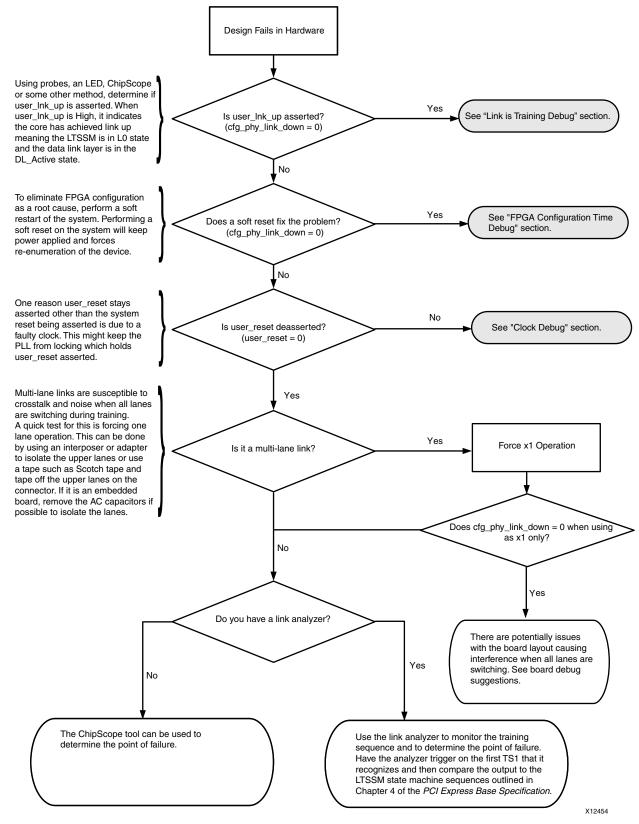


Figure C-3: Design Fails in Hardware Debug Flow Diagram

www.xilinx.com

FPGA Configuration Time 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. Section 6.6 of *PCI Express Base Specification, rev. 3.0* states two rules that might be impacted by FPGA Configuration Time:

- A component must enter the LTSSM Detect state within 20 ms of the end of the Fundamental reset.
- A system must guarantee that all components intended to be software visible at boot time are ready to receive Configuration Requests within 100 ms of the end of Conventional Reset at the Root Complex.

These statements basically mean the FPGA must be configured within a certain finite time, and not meeting these requirements could cause problems with link training and device recognition.

Configuration can be accomplished using an onboard PROM or dynamically using JTAG. When using JTAG to configure the device, configuration typically occurs after the Chipset has enumerated each peripheral. After configuring the FPGA, a soft reset is required to restart enumeration and configuration of the device. A soft reset on a Windows based PC is performed by going to **Start > Shut Down** and then selecting **Restart**.

To eliminate FPGA configuration as a root cause, the designer should perform a soft restart of the system. Performing a soft reset on the system keeps power applied and forces re-enumeration of the device. If the device links up and is recognized after a soft reset is performed, the FPGA configuration is most likely the issue. Most typical systems use ATX power supplies which provide some margin on this 100 ms window as the power supply is normally valid before the 100 ms window starts. For more information on FPGA configuration, see FPGA Configuration, page 328.

Link is Training Debug

Figure C-4 shows the flowchart for link trained debug.

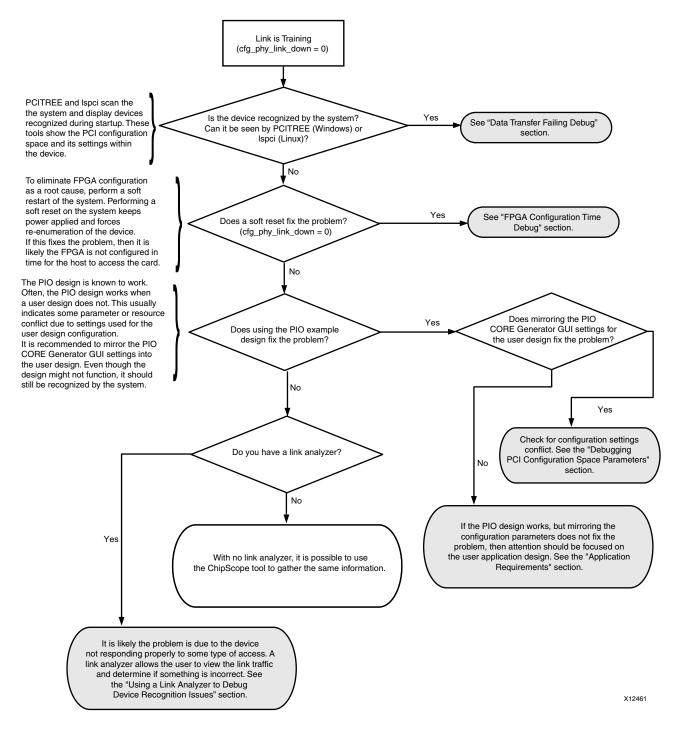


Figure C-4: Link Trained Debug Flow Diagram

Clock Debug

One reason to not deassert the user_reset signal is that the FPGA logic PLL (MMCM) and Transceiver PLL have not locked to the incoming clock. To verify lock, monitor the transceiver RXPLLLKDET output and the MMCM LOCK output. If the PLLs do not lock as

www.xilinx.com

expected, it is necessary to ensure the incoming reference clock meets the requirements in the 7 Series FPGAs GTX/GTH Transceivers User Guide [Ref 3]. The REFCLK signal should be routed to the dedicated reference clock input pins on the serial transceiver, and the user design should instantiate the IBUFDS_GTE2 primitive in the user design. See the 7 Series FPGAs GTX/GTH Transceivers User Guide for more information on PCB layout requirements, including reference clock requirements.

Reference clock jitter can potentially close both the TX and RX eyes, depending on the frequency content of the phase jitter. Therefore, as clean a reference clock as possible must be maintained. Reduce crosstalk on REFCLK by isolating the clock signal from nearby high-speed traces. Maintain a separation of at least 25 mils from the nearest aggressor signals. The PCI Special Interest Group website provides other tools for ensuring the reference clocks are compliant to the requirements of the *PCI Express Specification*: www.pcisig.com/specifications/pciexpress/compliance/compliance_library

Debugging PCI Configuration Space Parameters

Often, a user application fails to be recognized by the system, but the Xilinx PIO Example design works. In these cases, the user application is often using a PCI configuration space setting that is interfering with the system systems ability to recognize and allocate resources to the card.

Xilinx solutions for PCI Express handle all configuration transactions internally and generate the correct responses to incoming configuration requests. Chipsets have limits as to the amount of system resources it can allocate and the core must be configured to adhere to these limitations.

The resources requested by the Endpoint are identified by the BAR settings within the Endpoint configuration space. The user should verify that the resources requested in each BAR can be allocated by the chipset. I/O BARs are especially limited so configuring a large I/O BAR typically prevents the chipset from configuring the device. Generate a core that implements a small amount of memory (approximately 2 KB) to identify if this is the root cause.

The Class Code setting selected in the CORE Generator tool GUI can also affect configuration. The Class Code informs the Chipset as to what type of device the Endpoint is. Chipsets might expect a certain type of device to be plugged into the PCI Express slot and configuration might fail if it reads an unexpected Class Code. The BIOS could be configurable to work around this issue.

Use the PIO design with default settings to rule out any device allocation issues. The PIO design default settings have proven to work in all systems encountered when debugging problems. If the default settings allow the device to be recognized, then change the PIO design settings to match the intended user application by changing the PIO configuration the CORE Generator tool GUI. Trial and error might be required to pinpoint the issue if a link analyzer is not available.

Using a link analyzer, it is possible to monitor the link traffic and possibly determine when during the enumeration and configuration process problems occur.

Application Requirements

During enumeration, it is possible for the chipset to issue TLP traffic that is passed from the core to the backend application. A common oversight when designing custom backend applications is to not have logic which handles every type incoming request. As a result, no response is created and problems arise. The PIO design has the necessary backend functions to respond correctly to any incoming request. It is the responsibility of the application to generate the correct response. These packet types are presented to the application:

- Requests targeting the Expansion ROM (if enabled)
- Message TLPs
- Memory or I/O requests targeting a BAR
- All completion packets

The PIO design, can be used to rule out any of these types of concerns, as the PIO design responds to all incoming transactions to the user application in some way to ensure the host receives the proper response allowing the system to progress. If the PIO design works, but the custom application does not, some transaction is not being handled properly.

The ChipScope tool should be implemented on the wrapper Receive AXI4-Stream interface to identify if requests targeting the backend application are drained and completed successfully.

Using a Link Analyzer to Debug Device Recognition Issues

In cases where the link is up $(cfg_phy_link_down = 0)$, but the device is not recognized by the system, a link analyzer can help solve the issue. It is likely the FPGA is not responding properly to some type of access. The link view can be used to analyze the traffic and see if anything looks out of place.

To focus on the issue, it might be necessary to try different triggers. Here are some trigger examples:

- Trigger on the first INIT_FC1 and/or UPDATE_FC in either direction. This allows the analyzer to begin capture after link up.
- The first TLP normally transmitted to an Endpoint is the Set Slot Power Limit Message. This usually occurs before Configuration traffic begins. This might be a good trigger point.
- Trigger on Configuration TLPs.
- Trigger on Memory Read or Memory Write TLPs.

Data Transfer Failing Debug

Figure C-5 shows the flowchart for data transfer debug.

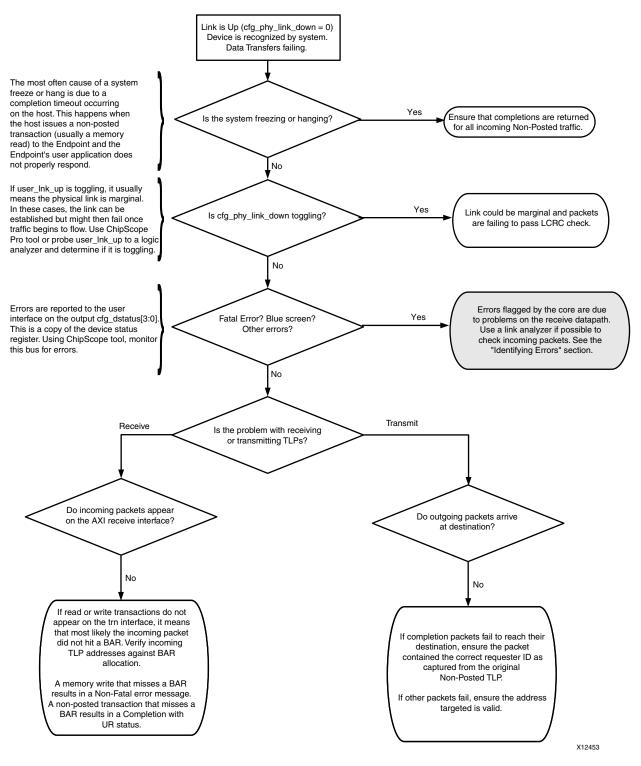


Figure C-5: Data Transfer Debug Flow Diagram

Identifying Errors

Hardware symptoms of system lock up issues are indicated when the system hangs or a blue screen appears (PC systems). The *PCI Express Base Specification, rev. 3.0* requires that error detection be implemented at the receiver. A system lock up or hang is commonly the result of a Fatal Error and is reported in bit 2 of the receiver Device Status register. Using the ChipScope tool, monitor the device status register of the core to see if a fatal error is being reported.

A fatal error reported at the Root complex implies an issue on the transmit side of the EP. The Root Complex Device Status register can often times be seen using PCITree (Windows) or LSPCI (Linux). If a fatal error is detected, see the Transmit section. A Root Complex can often implement Advanced Error Reporting, which further distinguishes the type of error reported. AER provides valuable information as to why a certain error was flagged and is provided as an extended capability within a devices configuration space. Section 7.10 of the *PCI Express Base Specification, rev. 3.0* provides more information on AER registers.

Transmit

Fatal Error Detected on Root or Link Partner

Check to make sure the TLP is correctly formed and that the payload (if one is attached) matches what is stated in the header length field. The Endpoints device status register does not report errors created by traffic on the transmit channel.

Monitor the AXI4-Stream signals to verify all traffic is being initiated correctly (see Port Descriptions, page 14).

Fatal Error Not Detected

Ensure that the address provided in the TLP header is valid. The kernel mode driver attached to the device is responsible for obtaining the system resources allocated to the device. In a Bus Mastering design, the driver is also responsible for providing the application with a valid address range. System hangs or blue screens might occur if a TLP contains an address that does not target the designated address range for that device.

Receive

System lock up conditions due to issues on the receive channel of the PCI Express core are often result of an error message being sent upstream to the root. Error messages are only sent when error reporting is enabled in the Device Control register.

A fatal condition is reported if any of these events occur:

- Training Error
- DLL Protocol Error
- Flow Control Protocol Error
- Malformed TLP
- Receiver Overflow

Non-Fatal Errors

This subsection lists conditions reported as Non-Fatal errors. See the *PCI Express Base Specification, rev. 3.0* for more details.

If the error is being reported by the root, the AER registers can be read to determine the condition that led to the error. Use a tool such as HWDIRECT, discussed in Third-Party Software Tools, page 315, to read the root AER registers. Chapter 7 of the *PCI Express Base Specification* defines the AER registers. If the error is signaled by the Endpoint, debug ports are available to help determine the specific cause of the error.

Correctable Non-Fatal errors are:

- Receiver Error
- Bad TLP
- Bad DLLP
- Replay Timeout
- Replay NUM Rollover

The first three errors listed above are detected by the receiver and are not common in hardware systems. The replay error conditions are signaled by the transmitter. If an ACK is not received for a packet within the allowed time, it is replayed by the transmitter. Throughput can be reduced if many packets are being replayed, and the source can usually be determined by examining the link analyzer or ChipScope tool captures.

Uncorrectable Non-Fatal errors are:

- Poisoned TLP
- Received ECRC Check Failed

- Unsupported Request (UR)
- Completion Timeout
- Completer Abort
- Unexpected Completion
- ACS Violation

An unsupported request usually indicates that the address in the TLP did not fall within the address space allocated to the BAR. This often points to an issue with the address translation performed by the driver. Ensure also that the BAR has been assigned correctly by the root at start-up. LSPCI or PCItree discussed in Third-Party Software Tools, page 315 can be used to read the BAR values for each device.

A completion timeout indicates that no completion was returned for a transmitted TLP and is reported by the requester. This can cause the system to hang (could include a blue screen on Windows) and is usually caused when one of the devices locks up and stops responding to incoming TLPs. If the root is reporting the completion timeout, the ChipScope tool can be used to investigate why the User Application did not respond to a TLP (for example, the User Application is busy, there are no transmit buffers available, or s_axis_tx_tready is deasserted). If the Endpoint is reporting the Completion timeout, a link analyzer would show the traffic patterns during the time of failure and would be useful in determining the root cause.

Next Steps

If the debug suggestions listed previously do not resolve the issue, open a support case to have the appropriate Xilinx expert assist with the issue.

To create a technical support case in WebCase, see the Xilinx website at:

www.xilinx.com/support/clearexpress/websupport.htm

Items to include when opening a case:

- Detailed description of the issue and results of the steps listed above.
- Attach ChipScope tool VCD captures taken in the steps above.

To discuss possible solutions, use the Xilinx User Community:

forums.xilinx.com/xlnx/

FPGA Configuration

This section discusses how to configure the Virtex-7 FPGA so that the device can link up and be recognized by the system. This information is provided for you to choose the correct FPGA configuration method for the system and verify that it works as expected.

This section discusses how specific requirements of the PCI Express Base Specification and PCI Express Card Electromechanical Specification apply to FPGA configuration.

RECOMMENDED: Where appropriate, Xilinx recommends that you read the actual specifications for detailed information.

This section is divided into four subsections:

- Configuration Terminology. Defines terms used in this chapter.
- Configuration Access Time. Several specification items govern when an Endpoint device needs to be ready to receive configuration accesses from the host (Root Complex).
- Board Power in Real-World Systems. Understanding real-world system constraints related to board power and how they affect the specification requirements.
- Recommendations. Describes methods for FPGA configuration and includes sample problem analysis for FPGA configuration timing issues.

Configuration Terminology

In this section, these terms are used to differentiate between FPGA configuration and configuration of the PCI Express® device:

- Configuration of the FPGA. FPGA configuration is used.
- Configuration of the PCI Express device. After the link is active, configuration is used.

Configuration Access Time

In standard systems for PCI Express, when the system is powered up, configuration software running on the processor starts scanning the PCI Express bus to discover the machine topology.

The process of scanning the PCI Express hierarchy to determine its topology is referred to as the *enumeration process*. The root complex accomplishes this by initiating configuration transactions to devices as it traverses and determines the topology.

All PCI Express devices are expected to have established the link with their link partner and be ready to accept configuration requests during the enumeration process. As a result,

there are requirements as to when a device needs to be ready to accept configuration requests after power up; if the requirements are not met, this occurs:

- If a device is not ready and does not respond to configuration requests, the root complex does not discover it and treats it as non-existent.
- The operating system does not report the device's existence and the user's application is not able to communicate with the device.

Choosing the appropriate FPGA configuration method is key to ensuring the device is able to communicate with the system in time to achieve link up and respond to the configuration accesses.

Configuration Access Specification Requirements

Two PCI Express specification items are relevant to configuration access:

 Section 6.6 of PCI Express Base Specification, rev. 3.0 states "A system must guarantee that all components intended to be software visible at boot time are ready to receive Configuration Requests within 100 ms of the end of Fundamental Reset at the Root Complex." For detailed information about how this is accomplished, see the specification; it is beyond the scope of this discussion.

Xilinx compliance to this specification is validated by the PCI Express-CV tests. The <u>PCI</u> <u>Special Interest Group (PCI-SIG)</u> provides the PCI Express Configuration Test Software to verify the device meets the requirement of being able to receive configuration accesses within 100 ms of the end of the fundamental reset. The software, available to any member of the PCI-SIG, generates several resets using the in-band reset mechanism and PERST# toggling to validate robustness and compliance to the specification.

 Section 6.6 of PCI Express Base Specification rev. 3.0 defines three parameters necessary "where power and PERST# are supplied." The parameter T_{PVPERL} applies to FPGA configuration timing and is defined as:

T_{PVPERL} - PERST# must remain active at least this long after power becomes valid.

The PCI Express Base Specification does not give a specific value for T_{PVPERL} – only its meaning is defined. The most common form factor used by designers with the Integrated Block core is an ATX-based form factor. The PCI Express Card Electromechanical Specification focuses on requirements for ATX-based form factors. This applies to most designs targeted to standard desktop or server type motherboards. Figure C-6 shows the relationship between Power Stable and PERST#.

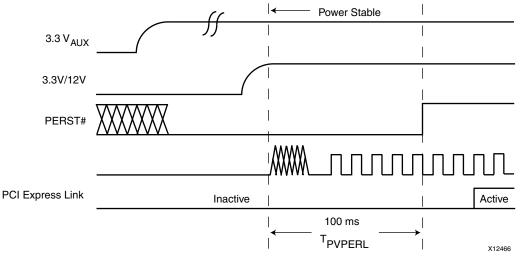


Figure C-6: Power Up

Section 2.6.2 of the *PCI Express Card Electromechanical Specification, rev. 3.0* defines T_{PVPREL} as a minimum of 100 ms, indicating that from the time power is stable the system reset is asserted for at least 100 ms (as shown in Table C-1).

Table C-1: T_{PVPERL} Specification

| Symbol | Parameter | Min | Max | Units |
|---------------------|---------------------------------|-----|-----|-------|
| T _{PVPERL} | Power stable to PERST# inactive | 100 | | ms |

From Figure C-6 and Table C-1, it is possible to obtain a simple equation to define the FPGA configuration time as follows:

FPGA Configuration Time
$$\leq T_{PWRVLD} + T_{PVPERL}$$
Equation C-1

Given that T_{PVPERL} is defined as 100 ms minimum, this becomes:

FPGA Configuration Time $\leq T_{PWRVLD} + 100 \text{ ms}$ Equation C-2

Note: Although T_{PWRVLD} is included in Equation C-2, it has yet to be defined in this discussion because it depends on the type of system in use. The Board Power in Real-World Systems section defines T_{PWRVLD} for both ATX-based and non ATX-based systems.

FPGA configuration time is only relevant at cold boot; subsequent warm or hot resets do not cause reconfiguration of the FPGA. If the design appears to be having issues due to FPGA configuration, the user should issue a warm reset as a simple test, which resets the system, including the PCI Express link, but keeps the board powered. If the issue does not appear, the issue could be FPGA configuration time related.

Board Power in Real-World Systems

Several boards are used in PCI Express systems. The *ATX Power Supply Design* specification, endorsed by Intel, is used as a guideline and for this reason followed in the majority of mother boards and 100% of the time if it is an Intel-based motherboard. The relationship between power rails and power valid signaling is described in the <u>ATX 12V Power Supply</u> <u>Design Guide</u>. Figure C-7, redrawn here and simplified to show the information relevant to FPGA configuration, is based on the information and diagram found in section 3.3 of the *ATX 12V Power Supply Design Guide*. For the entire diagram and definition of all parameters, see the *ATX 12V Power Supply Design Guide*.

Figure C-7 shows that power stable indication from Figure C-6 for the PCI Express system is indicated by the assertion of PWR_OK. PWR_OK is asserted High after some delay when the power supply has reached 95% of nominal.

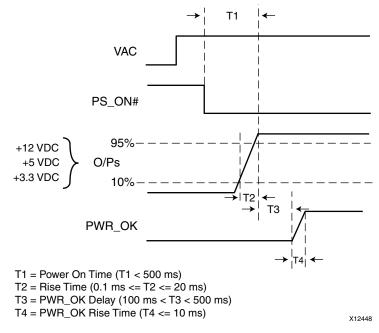


Figure C-7: **ATX Power Supply**

Figure C-7 shows that power is valid before PWR_OK is asserted High. This is represented by T3 and is the PWR_OK delay. The *ATX 12V Power Supply Design Guide* defines PWR_OK as 100 ms < T3 < 500 ms, indicating that from the point at which the power level reaches 95% of nominal, there is a minimum of at least 100 ms but no more than 500 ms of delay before PWR_OK is asserted. Remember, according to the *PCI Express Card Electromechanical Specification*, the PERST# is guaranteed to be asserted a minimum of 100 ms from when power is stable indicated in an ATX system by the assertion of PWR_OK.

Again, the FPGA configuration time equation is:

FPGA Configuration Time
$$\leq T_{PWRVLD} + 100 \text{ ms}$$
Equation C-3

 T_{PWRVLD} is defined as PWR_OK delay period; that is, T_{PWRVLD} represents the amount of time that power is valid in the system before PWR_OK is asserted. This time can be added to the amount of time the FPGA has to configure. The minimum values of T2 and T4 are negligible and considered zero for purposes of these calculations. For ATX-based motherboards, which represent the majority of real-world motherboards in use, T_{PWRVLD} can be defined as:

$$100 \text{ ms} \leq \text{T}_{PWRVLD} \leq 500 \text{ ms}$$

Equation C-4

This provides these requirements for FPGA configuration time in both ATX and non-ATX-based motherboards:

- FPGA Configuration Time ≤ 200 ms (for ATX based motherboard)
- FPGA Configuration Time ≤ 100 ms (for non-ATX based motherboard)

The second equation for the non-ATX based motherboards assumes a T_{PWRVLD} value of 0 ms because it is not defined in this context. Designers with non-ATX based motherboards should evaluate their own power supply design to obtain a value for T_{PWRVLD}.

This section assumes that the FPGA power (V_{CCINT}) is stable before or at the same time that PWR_OK is asserted. If this is not the case, additional time must be subtracted from the available time for FPGA configuration.

RECOMMENDED: Avoid designing add-in cards with staggered voltage regulators with long delays.

Hot Plug Systems

Hot Plug systems generally employ the use of a Hot-Plug Power Controller located on the system motherboard. Many discrete Hot-Plug Power Controllers extend T_{PVPERL} beyond the minimum 100 ms. Add-in card designers should consult the Hot-Plug Power Controller data sheet to determine the value of T_{PVPERL} . If the Hot-Plug Power Controller is unknown, then a T_{PVPERL} value of 100 ms should be assumed.

Recommendations

RECOMMENDED: For minimum FPGA configuration time, use the BPI configuration mode with a parallel NOR flash, which supports high-speed synchronous read operation.

In addition, an external clock source can be supplied to the external master configuration clock (EMCCLK) pin to ensure a consistent configuration clock frequency for all conditions. See the 7 Series FPGAs Configuration User Guide [Ref 3] for descriptions of the BPI configuration mode and EMCCLK pin. This section discusses these recommendations and includes sample analysis of potential issues that might arise during FPGA configuration.

FPGA Configuration Times for Virtex-7 Devices

During power up, the FPGA configuration sequence is performed in three steps:

- 1. Wait for power on reset (POR) for all voltages (V_{CCINT}, V_{CCAUX}, and VCCO_0) in the FPGA to trip, referred to as POR Trip Time.
- 2. Wait for completion (deassertion) of INIT_B to allow the FPGA to initialize before accepting a bitstream transfer.

Note: As a general rule, steps 1 and 2 require \leq 50 ms

- 3. Wait for assertion of DONE, the actual time required for a bitstream to transfer depends on:
 - Bitstream size
 - Clock (CCLK) frequency
 - Transfer mode (and data bus width) from the flash device
 - SPI = Serial Peripheral Interface (x1, x2, or x4)
 - BPI = Byte Peripheral Interface (x8 or x16)

Bitstream transfer time can be estimated using this equation.

Bitstream transfer time = (bitstream size in bits)/(CCLK frequency)/ (data bus width in bits) Equation C-5

For detailed information about the configuration process, see the 7 Series FPGAs Configuration User Guide [Ref 3].

Sample Problem Analysis

This section presents data from an ASUS PL5 system to demonstrate the relationships between Power Valid, FPGA Configuration, and PERST#. Figure C-8 shows a case where the Endpoint failed to be recognized due to an FPGA configuration time issue. Figure C-9 shows a successful FPGA configuration with the Endpoint being recognized by the system.

Failed FPGA Recognition

Figure C-8 illustrates an example of a cold boot where the host failed to recognize the Xilinx FPGA. Although a second PERST# pulse assists in allowing more time for the FPGA to configure, the slowness of the FPGA configuration clock (2 MHz) causes configuration to complete well after this second deassertion. During this time, the system enumerated the bus and did not recognize the FPGA.

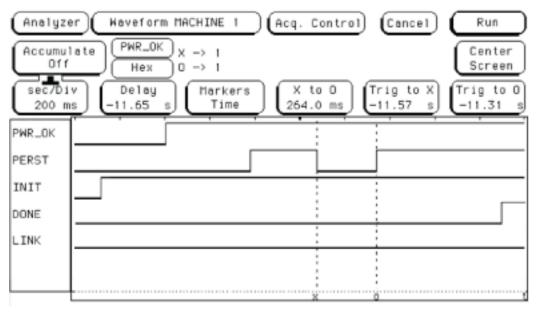


Figure C-8: Host Fails to Recognize FPGA Due to Slow Configuration Time

Successful FPGA Recognition

Figure C-9 illustrates a successful cold boot test on the same system. In this test, the CCLK was running at 50 MHz, allowing the FPGA to configure in time to be enumerated and recognized. The figure shows that the FPGA began initialization approximately 250 ms before PWR_OK. DONE going High shows that the FPGA was configured even before PWR_OK was asserted.

| Analyze | er (Haveform MACHINE 1) (Acq. Control) (Cancel) | Run |
|-----------------|---|-----------------------|
| Accumul | $\begin{array}{c} \text{late} \\ \hline \\ \text{Hex} \\ \text{Hex} \\ \end{array} \begin{array}{c} 0 \\ \rightarrow \\ 1 \end{array} \end{array} $ | Center Screen |
| sec/Di 200 m | | Trig to 0 384.0 ms |
| PWR_OK | | |
| PERST | | |
| INIT | | |
| DONE | | |
| CLK | | |
| | | |
| | x t o | |

Figure C-9: Host Successfully Recognizes FPGA

Workarounds for Closed Systems

For failing FPGA configuration combinations, designers might be able to work around the issue in closed systems or systems where they can guarantee behavior. These options are not recommended for products where the targeted end system is unknown.

1. Check if the motherboard and BIOS generate multiple PERST# pulses at start-up. This can be determined by capturing the signal on the board using an oscilloscope. This is similar to what is shown in Figure C-8. If multiple PERST# pulses are generated, this typically adds extra time for FPGA configuration.

Define $T_{PERSTPERIOD}$ as the total sum of the pulse width of PERST# and deassertion period before the next PERST# pulse arrives. Because the FPGA is not power cycled or reconfigured with additional PERST# assertions, the $T_{PERSTPERIOD}$ number can be added to the FPGA configuration equation.

```
FPGA Configuration Time \leq T_{PWRVLD} + T_{PERSTPERIOD} + 100 \text{ ms}Equation C-6
```

2. In closed systems, it might be possible to create scripts to force the system to perform a warm reset after the FPGA is configured, after the initial power up sequence. This resets the system along with the PCI Express subsystem allowing the device to be recognized by the system.

Simulation Debug

This section provides simulation debug flow diagrams for some of the most common issues experienced by users. Endpoints that are shaded gray indicate that more information can be found in sections after Figure C-10.

ModelSim Debug

Figure C-10 shows the flowchart for ModelSim debug.

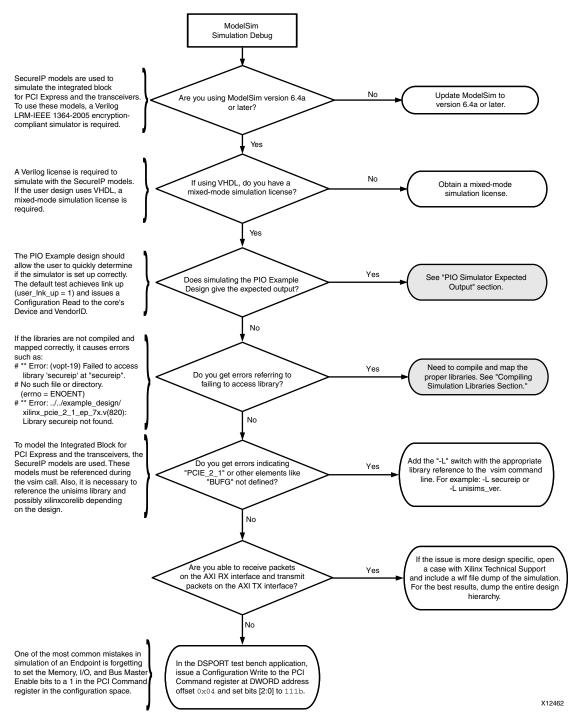


Figure C-10: ModelSim Debug Flow Diagram

Compiling Simulation Libraries

Use the compxlib command to compile simulation libraries. This tool is delivered as part of the Xilinx software. For more information see the ISE tool manuals and specifically the *Development System Reference Guide* under the section titled compxlib.

www.xilinx.com

Assuming the Xilinx and ModelSim environments are set up correctly, this is an example of compiling the SecureIP and UNISIM libraries for Verilog into the current directory:

```
compxlib -s mti_se -arch virtex7 -l verilog -lib secureip -lib unisims -dir ./
```

There are many other options available for compxlib described in the *Command Line Tools* User Guide.

Compxlib produces a modelsim.ini file containing the library mappings. In ModelSim, to see the current library mappings type **vmap** at the prompt. The mappings can be updated in the ini file, or to map a library at the ModelSim prompt, type:

```
vmap [<logical_name>] [<path>]
```

For example:

Vmap unisims_ver C:\my_unisim_lib

Next Steps

If the debug suggestions listed previously do not resolve the issue, a support case should be opened to have the appropriate Xilinx expert assist with the issue.

To create a technical support case in WebCase, see the Xilinx website at:

www.xilinx.com/support/clearexpress/websupport.htm

Items to include when opening a case:

- Detailed description of the issue and results of the steps listed above.
- Attach a VCD or WLF dump of the simulation.

To discuss possible solutions, use the Xilinx User Community:

forums.xilinx.com/xlnx/



Appendix D

Attributes

Table D-1 defines the attributes in the LogiCORE[™] IP Virtex-7 FPGA Gen3 Integrated Block for PCI Express core.

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|---|---------------------------|
| Core Basic | | | | |
| CRM_CORE_CLK_FREQ_500 | 1 | STRING | Core Clock Frequency. When this attribute is TRUE, the core clock input frequency is 500 MHz; otherwise, the core clock frequency is 250 MHz. | Y |
| CRM_USER_CLK_FREQ | 2 | HEX | User clock frequency. Valid settings are: • 00b: 62.5 MHz • 01b: 125 MHz • 10b: 250 MHz • 11b: Reserved | Y |
| AXISTEN_IF_WIDTH | 2 | HEX | AXI4-Stream Enhanced Interface Width. Valid settings are: • 00b: 64b • 01b: 128b • 10b: 256b • 11b: Reserved | Y |
| AXISTEN_IF_CQ_ALIGNMENT_MODE | 1 | STRING | AXI4-Stream Enhanced Interface CQ Alignment. Determines the data alignment mode for the CQ interface: 0: Dword-aligned mode 1: Address-aligned mode | Y |
| AXISTEN_IF_CC_ALIGNMENT_MODE | 1 | STRING | AXI4-Stream Enhanced Interface CC Alignment. Determines the data alignment mode for the CC interface: 0: Dword-aligned mode 1: Address-aligned mode | Y |
| AXISTEN_IF_RQ_ALIGNMENT_MODE | 1 | STRING | AXI4-Stream Enhanced Interface RQ Alignment. Determines the data alignment mode for the RQ interface: 0: Dword-aligned mode 1: address-aligned mode | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode |
|--------------------------------|-------------------|----------------|--|--------------------------|
| AXISTEN_IF_RC_ALIGNMENT_MODE | 1 | STRING | AXI4-Stream Enhanced Interface RC Alignment. Determines the data alignment mode for the RC interface: 0: Dword-aligned mode 1: address-aligned mode | Y |
| AXISTEN_IF_RC_STRADDLE | 1 | STRING | Received AXISTEN Frame Straddle. When this attribute is TRUE, received requester completion AXISTEN frames are enabled to straddle single cycle transfers when AXISTEN_IF_WIDTH is configured to 256 bits. When this attribute is FALSE, the straddle feature is disabled. | Y |
| AXISTEN_IF_ENABLE_RX_MSG_INTFC | 1 | STRING | Received AXISTEN message interface enable. When this attribute is set to 0, received messages are delivered through the CQ interface. When this attribute is set to 1, these messages are delivered through the receive message interface. | Y |
| AXISTEN_IF_ENABLE_MSG_ROUTE | 18 | HEX | Received AXISTEN message routing. Enables the routing of message TLPs to the user through the AXI4-Stream CQ interface. A bit value of 1 enables routing of the message TLP to the user. Messages are always decoded by the message decoder. • Bit 0: ERR_COR • Bit 1: ERR_NONFATAL • Bit 2: ERR_FATAL • Bit 2: ERR_FATAL • Bit 3: Assert_INTA and Deassert_INTA • Bit 4: Assert_INTB and Deassert_INTB • Bit 5: Assert_INTC and Deassert_INTC • Bit 6: Assert_INTD and Deassert_INTD • Bit 7: PM_PME • Bit 8: PME_TO_Ack • Bit 9: PME_Turn_Off • Bit 10: PM_Active_State_Nak • Bit 11: Set_Slot_Power_Limit • Bit 12: Latency Tolerance Reporting (LTR) • Bit 13: Optimized Buffer Flush/Fill (OBFF) • Bit 14: Unlock • Bit 15: Vendor_Defined Type 0 • Bit 16: Vendor_Defined Type 1 • Bit 17: Invalid Request, Invalid Completion, Page, Request, PRG Response | Y |
| AXISTEN_IF_RQ_PARITY_CHK | 1 | STRING | Requestor Request Parity Check. • TRUE: Parity check is enabled • FALSE: Parity check is disabled | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|--|---------------------------|
| AXISTEN_IF_CC_PARITY_CHK | 1 | STRING | Completer Completion Parity Check. TRUE: Parity check is enabled FALSE: Parity check is disabled | Y |
| AXISTEN_IF_ENABLE_CLIENT_TAG | 1 | STRING | AXI4-Stream Enhanced Interface Tag management option for the RQ interface.FALSE: Tags are managed by the coreTRUE: Tags are managed externally by the client | Y |
| Power Management | | | | |
| PM_ASPMLOS_TIMEOUT | 16 | HEX | LOS Timeout Limit Register. Timeout value for transitioning to the LOS power state. If the transmit side has been idle for this interval, the core transmits the idle sequence on the link and transitions the state of the link to LOS. Contains the timeout value (in units of 4 ns) for transitioning to the LOS power state. Setting this attribute to 0 permanently disables the transition to the LOS power state. | Y |
| PM_L1_REENTRY_DELAY | 32 | HEX | L1 State Re-entry Delay Register. Time the core waits before it re-enters the L1 state, if its link partner transitions the link to L0 while all Functions of the core are in the D3 power state. The core changes the power state of the link from L0 to L1 if no activity is detected on both transmit and receive sides before this interval, while all Functions are in the D3 state and the link is in the L0 state. Setting this register to 0 disables re-entry to the L1 state if the link partner returns the link to L0 from L1 when all Functions of the core are in the D3 state. This register controls only the re-entry to L1. The initial transition to L1 always occurs when all Functions of the core are set to the D3 state. | Y |
| PM_ASPML1_ENTRY_DELAY | 20 | HEX | ASPM L1 Entry Timeout Delay Register. Contains the timeout value (in units of 4 ns) for transitioning to the L1 power state. Setting this register to 0 permanently disables the transition to the L1 power state. | Y |
| PM_ENABLE_SLOT_POWER_CAPTURE | 1 | STRING | When this attribute is set to TRUE and the core is configured as an Endpoint, the core captures the Slot Power Limit Value and Slot Power Limit Scale parameters from a Set_Slot_Power_Limit message received in the Device Capabilities Register. When this attribute is 0, the capture is disabled. | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|---|---------------------------|
| PM_PME_SERVICE_TIMEOUT_DELAY | 20 | HEX | Specifies the timeout delay for retransmission of PM_PME messages. The value is in units of microseconds. | Y |
| PM_PME_TURNOFF_ACK_DELAY | 16 | HEX | Time in microseconds between the core receiving a PME_Turn_Off message TLP and sending a PME_TO_Ack response to it. This field must be set to a non-zero value to enable the core to send the response. Setting this field to 0 suppresses the response of the core to the PME_Turn_Off message, so that you can transmit the PME_TO_Ack message through the AXI4-Stream interface. | Y |
| Physical Layer | | • | | |
| PL_UPSTREAM_FACING | 1 | STRING | Physical Layer Mode. TRUE specifies an upstream-facing port. FALSE specifies a downstream-facing port. This setting is propagated to all layers in the core. | Y |
| PL_LINK_CAP_MAX_LINK_WIDTH | 4 | HEX | Maximum Link Width. Valid settings are: • 0001b: x1 • 0010b: x2 • 0100b: x4 • 1000b: x8 All other encodings are reserved. This setting is propagated to all layers in the core. | Y |
| PL_LINK_CAP_MAX_LINK_SPEED | 3 | HEX | Maximum Link Speed. Valid settings are: 001b: Gen1 (2.5 GT/s) 010b: Gen2 (5.0 GT/s) 100b: Gen3 (8.0 GT/s) All other encodings are reserved. This setting is propagated to all layers in the core. | Y |
| PL_DISABLE_GEN3_DC_BALANCE | 1 | STRING | Disable Gen3 DC Balance. Disables transmission of special symbols when set to TRUE. | Y |
| PL_DISABLE_EI_INFER_IN_L0 | 1 | STRING | When this attribute is set to TRUE, the inferring of electrical idle in the L0 state is disabled. Electrical idle is inferred when no flow control updates and no SKP sequences are received within an interval of 128 µs. Do not set this attribute to TRUE during normal operation. Use it for system debug. | Y |
| PL_N_FTS_COMCLK_GEN1 | 8 | DECIMAL | Sets the number of FTS OS, advertised in the TS1 Ordered Sets, when the Link Configuration register shows that a common clock source is selected. | Y |

 PL_LANE0_EQ_CONTROL
 16
 HEX
 Equalization Control Register in the Secondary PCI Express Extended Capability Header.

 PL_LANE0_EQ_CONTROL
 16
 HEX
 Bit [3:0]: Downstream Port Transmitter Preset

 Bit [6:4]: Downstream Port Receiver Preset Hint
 Bit [11:8]: Upstream Port Transmitter Preset

 Bit [11:8]: Upstream Port Receiver Preset Hint
 Bit [14:12]: Upstream Port Receiver Preset Hint

| XI | LIN | 1X° | |
|----|-----|-----|--|
| | | | |

PL_N_FTS_GEN1

PL_N_FTS_GEN2

PL_N_FTS_GEN3

PL_N_FTS_COMCLK_GEN2

PL_N_FTS_COMCLK_GEN3

PL_DISABLE_UPCONFIG_CAPABLE

Attribute Name

Σ.

Attribute Description

Sets the number of FTS OS, advertised in the

Configuration register shows that a common

Sets the number of FTS OS, advertised in the

Configuration register shows that a common

Sets the number of FTS OS, advertised in the

Configuration register shows that a common

Sets the number of FTS OS, advertised in the

Configuration register shows that a common

Sets the number of FTS OS, advertised in the

Configuration register shows that a common

When set to TRUE, this attribute disables the

upconfigure capability. When set to FALSE, it

Lane #0 Equalization Control Register. Sets the appropriate lane-specific entry in the

TS1 Ordered Sets, when the Link

enables the upconfigure capability.

clock source is not selected.

clock source is not selected.

clock source is selected.

clock source is not selected.

clock source is selected.

Applicable

to RP Mode?

Y

Y

Υ

Y

Y

Y

Y

Table D-1: Virtex-7 FPGA Gen3 Integrated Block for PCI Express Core Attribute Descriptions (Cont'd)

Attribute Type

DECIMAL

DECIMAL

DECIMAL

DECIMAL

DECIMAL

STRING

Number

of Bits

8

8

8

8

8

1

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---------------------|-------------------|----------------|--|---------------------------|
| PL_LANE1_EQ_CONTROL | 16 | HEX | Lane #1 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |
| PL_LANE2_EQ_CONTROL | 16 | HEX | Lane #2 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |
| PL_LANE3_EQ_CONTROL | 16 | HEX | Lane #3 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---------------------|-------------------|----------------|--|---------------------------|
| PL_LANE4_EQ_CONTROL | 16 | HEX | Lane #4 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |
| PL_LANE5_EQ_CONTROL | 16 | HEX | Lane #5 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |
| PL_LANE6_EQ_CONTROL | 16 | HEX | Lane #6 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|----------------------------------|-------------------|----------------|--|---------------------------|
| PL_LANE7_EQ_CONTROL | 16 | HEX | Lane #7 Equalization Control Register. Sets the appropriate lane-specific entry in the Equalization Control Register in the Secondary PCI Express Extended Capability Header. Bit [3:0]: Downstream Port Transmitter Preset Bit [6:4]: Downstream Port Receiver Preset Hint Bit [11:8]: Upstream Port Transmitter Preset Bit [14:12]: Upstream Port Receiver Preset Hint | Y |
| PL_EQ_BYPASS_PHASE23 | 1 | STRING | Bypass Equalization Phases 2 and 3. When this attribute is TRUE and PL_UPSTREAM_FACING is FALSE, optional EQ Phases are bypassed. | Y |
| PL_EQ_ADAPT_ITER_COUNT | 5 | HEX | Link Partner Transmitter Adaptive Equalization Iteration Count. When in EQ Phase 2 for EP and Phase 3 for RP, this attribute contains the maximum number of iterations of Adaptive Equalization attempted before the Recovery.EQ phase is exited. The supported range is 2–31. | Y |
| PL_EQ_ADAPT_REJECT_RETRY_COUNT | 2 | HEX | Equalization Adaptation Proposal Rejection Retry Count. This attribute contains the number of times new remote transmitter coefficients are proposed after being continuously rejected by the link partner, per lane. The supported range is 0–3. | Y |
| PL_EQ_SHORT_ADAPT_PHASE | 1 | STRING | Shorten the Receive Adaptation Phase. When this attribute is set to TRUE, EQ Phase 2 for EP and Phase 3 for RP return the received TX Preset OR Coefficients as the new proposed settings. This attribute can be used for simulation speed-up and debug purposes. | Y |
| PL_EQ_ADAPT_DISABLE_COEFF_CHECK | 1 | STRING | Disable checks on Received Coefficients. When this attribute is set to TRUE, received coefficient checking is disabled (no rejection) during EQ Phase 3 for EP and EQ Phase 2 for RP. | Y |
| PL_EQ_ADAPT_DISABLE_PRESET_CHECK | 1 | STRING | Disable checks on Received Presets. When this attribute is set to TRUE, received preset checking is disabled (no rejection) during EQ Phase 3 for EP and EQ Phase 2 for RP. | Y |

E XILINX.

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---------------------------|-------------------|----------------|--|---------------------------|
| PL_SIM_FAST_LINK_TRAINING | 1 | STRING | When this attribute is set to TRUE, the link training time is shortened to facilitate fast simulation of the design, especially at the gate level. Enabling this attribute has these effects: The 1 ms, 2 ms, 12 ms, 24 ms, 32 ms, and 48 ms timeout intervals in the LTSSM are shortened by a factor of 500. In the Polling.Active state of the LTSSM, only 16 training sequences are required to be transmitted (instead of 1024) to make the transition to the Configuration state. This must be set to FALSE for non-simulation use. | Y |
| Link Layer | | | | |
| LL_ACK_TIMEOUT_EN | 1 | STRING | When this attribute is TRUE, the Ack/Nak Latency Timer is enabled to use the user-defined LL_ACK_TIMEOUT value (or combined with the built-in value, depending on LL_ACK_TIMEOUT_FUNC). When this attribute is FALSE, the built-in value is used. | Y |
| LL_ACK_TIMEOUT | 9 | HEX | Sets a user-defined timeout for the Ack/Nak Latency Timer to force any pending ACK or NAK DLLPs to be transmitted. Refer to LL_ACK_TIMEOUT_EN and LL_ACK_TIMEOUT_FUNC to see how this value is used. The units for this attribute are in symbol times, which is 4 ns at Gen1 speeds (2.5 GT/s) and 2 ns at Gen2. speeds (5.0 GT/s). | Y |
| LL_ACK_TIMEOUT_FUNC | 2 | DECIMAL | Defines how LL_ACK_TIMEOUT is to be used, if enabled with LL_ACK_TIMEOUT_EN (otherwise, this attribute is not used). 0: No effect 1: Add LL_ACK_TIMEOUT to the built-in table value. 2: Subtract LL_ACK_TIMEOUT from the built-in table value. The LL_ACK_TIMEOUT walue should follow these rules for any Width at Gen1/2/3 speeds: if and only if MPS > 512B: the allowed Range is 1 ≤ RANGE ≤ 64d if and only if MPS ≤ 512B: the allowed Range is 1 ≤ RANGE ≤ 32d | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|--|---------------------------|
| LL_REPLAY_TIMEOUT_EN | 1 | STRING | When this attribute is TRUE, the Replay Timer is enabled to use the user-defined LL_REPLAY_TIMEOUT value (or combined with the built-in value, depending on LL_REPLAY_TIMEOUT_FUNC). When this attribute is FALSE, the built-in value is used. | Y |
| LL_REPLAY_TIMEOUT | 9 | HEX | Sets a user-defined timeout for the Replay Timer to force cause the retransmission of unacknowledged TLPs. See LL_REPLAY_TIMEOUT_EN and LL_REPLAY_TIMEOUT_FUNC to see how this value is used. The unit for this attribute is in symbol times, which is 4 ns at Gen1 speeds (2.5 GT/s) and 2 ns at Gen2 speeds (5.0 GT/s). | Y |
| LL_REPLAY_TIMEOUT_FUNC | 2 | DECIMAL | Defines how to use LL_REPLAY_TIMEOUT, if enabled with LL_REPLAY_TIMEOUT_EN (otherwise, this attribute is not used). 0: No effect 1: Add LL_REPLAY_TIMEOUT to the built-in table value. 2: Subtract LL_REPLAY_TIMEOUT from the built-in table value. The LL_REPLAY_TIMEOUT value should follow these rules for any Width at Gen1/2/3 speeds: if and only if MPS > 512B: The allowed Range is 1 < RANGE < 64d if and only if MPS < 512B: the allowed Range is 1 < RANGE < 32d | Y |
| Transaction Layer | 1 | 1 | | 1 |
| TL_PF_ENABLE_REG | 1 | STRING | Function #1 Enable. Function #1 is enabled when this attribute is set to TRUE. | Y |
| TL_CREDITS_CD | 12 | HEX | Receiver Credit Limit for Completion Data. Units are credits. Supported values are: • 3E0h: 15,872 bytes • 1F0h: 7,936 bytes • 000h: Infinite Must be set to 000h when PL_UPSTREAM_FACING is TRUE. | Y |
| TL_CREDITS_CH | 8 | HEX | Receiver Credit Limit for Completion Header. Units are number of TLPs. Supported values are: 20h: 32 TLPs 00h: Infinite credits Must be set to 00h when PL_UPSTREAM_FACING is TRUE. | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|-----------------------|-------------------|----------------|--|---------------------------|
| TL_CREDITS_NPD | 12 | HEX | Credit Limit for Non Posted Data. Units are credits. Supported values are: • 28h: 640 bytes | Y |
| TL_CREDITS_NPH | 8 | HEX | Receiver Credit Limit for Non-Posted Header. Units are number of TLPs. Supported values are: • 20h: 32 TLPs | Y |
| TL_CREDITS_PD | 12 | HEX | Receiver Credit Limit for Posted Data. Units are credits. Supported values are: 198h: 6,528 bytes CCh: 3,264 bytes | Y |
| TL_CREDITS_PH | 8 | HEX | Receiver Credit Limit for Posted Header. Units are number of TLPs. Supported values are: • 20h: 32 TLPs | Y |
| TL_COMPL_TIMEOUT_REG0 | 24 | HEX | Completion Timeout Limit Register #0. This register contains the timeout value used to detect a completion timeout event for a Request originated by the core from its AXI4-Stream master interface, when sub-range 1 (value 0101) is programmed in the Device Control 2 Register. The valid range for this parameter is 16 ms to 55 ms, which can be specified in units of 4 ns cycles. For example, the value of 12,500,000 in this parameter corresponds to 50 ms (12,500,000 * 4 ns). | Y |
| TL_COMPL_TIMEOUT_REG1 | 28 | HEX | Completion Timeout Limit Register 1. This register contains the timeout value used to detect a completion timeout event for a Request originated by the core from its AXI4-Stream master interface, when sub-range 2 (value 0110) is programmed in the Device Control 2 register. The valid range for this parameter is 65 ms to 210ms, which can be specified in units of 4 ns cycles. For example, the value of 25,000,000 in this parameter corresponds to 100 ms (25,000,000 * 4 ns). | Y |
| TL_LEGACY_MODE_ENABLE | 1 | STRING | Enable Legacy Endpoint Mode. When this attribute and PL_UPSTREAM_FACING are TRUE, the core is configured as a Legacy Endpoint, where it forwards Locked Transactions through the AXI4-Stream target interface, and generates Locked Completions in response to them. If this attribute is TRUE and PL_UPSTREAM_FACING is FALSE, you can generate Locked Read Transactions as a Master. When this attribute is FALSE, the core is configured as a PCIe Endpoint. | Y |

| TL_ENABLE_MESSAGE_RID_CHECK_ ENABLE | 1 | STRING | Messages with an RID mismatch are handled as Unsupported Requests and discarded within the core. When this attribute is FALSE, the core does not check the RIDs of received Message TLPs and forwards them to the AXISTEN interface. This attribute is applicable only when the core is configured as an Endpoint. |
|---|----|--------|---|
| TL_LEGACY_CFG_EXTEND_INTERFACE_ ENABLE | 1 | STRING | Configuration Legacy Space Extend Interface Enable. When this attribute is TRUE, all received Configuration Type0 Transactions in the register address range $0 \ge 10-0 \ge 1f$ for every enabled function, are steered to the CFGEXT interface. |
| TL_EXTENDED_CFG_EXTEND_ INTERFACE_ENABLE | 1 | STRING | Configuration Extended Space Extend Interface Enable. When this attribute is TRUE, all received Configuration Type0 Transactions in the register address range $0 \times 100-0 \times 3 \text{ ff}$ for every enabled function are steered to the CFGEXT interface. |
| Legacy PCI Header | | | |
| PFx_DEVICE_ID | 16 | HEX | Device ID. Individual Device IDs for each physical function (PF). |
| PFx_REVISION_ID | 8 | HEX | Revision ID. Individual Revision IDs for each PF and associated virtual functions (VFs). |
| PFx_CLASS_CODE | 24 | HEX | Class Code. Code identifying basic function, subclass, and applicable programming interface. This value is transferred to the Class Code, Sub-Class Code, and Programming Interface. VFs must return the same Class Code as the corresponding PF. |
| PFx_SUBSYSTEM_ID | 16 | HEX | Subsystem ID. Individual Subsystem ID for each PF and associated VFs. |
| PFx_INTERRUPT_PIN | 3 | HEX | Interrupt Pin Register. This attribute indicates the mapping for legacy Interrupt Messages. Valid values are 1 INTA, 2 INTB, 3 INTC, 4 INTD. Zero indicates no legacy interrupt messages used. This attribute does not apply to VFs. |

www.xilinx.com

Table D-1: Virtex-7 FPGA Gen3 Integrated Block for PCI Express Core Attribute Descriptions (Cont'd)

Attribute Type

Number

of Bits

Attribute Name

Attribute Description

Enable Requestor ID Checking of Received TLPs. When this attribute is set to TRUE, the core checks the requestor ID (RID) of the incoming Message TLPs that are routed by ID, against the RIDs of its enabled Functions. Applicable

to RP Mode?

Y

Ν

Ν

Y

Υ

Y

Υ

Y

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|---|---------------------------|
| PFx_INTERRUPT_LINE | 8 | HEX | Interrupt Line Register. This attribute identifies the IRQx input of the interrupt controller at the Root Complex that is activated by this Function interrupt. 00h: IRQ0 0Fh: IRQ15 FFh: Unknown or not connected This attribute does not apply to VFs. | Y |
| PFx_BIST_REGISTER | 8 | HEX | BIST Control Register. Sets the 8 bits in the BIST register. This attribute does not apply to VFs. | Y |
| PFx_CAPABILITY_POINTER | 8 | HEX | Capability Pointer. Next capability pointer at 34h in each PF. | Y |
| VF0_CAPABILITY_POINTER | 8 | HEX | Capability Pointer. Next capability pointer at 34h in VF0. | N |
| BARs | | | | |
| PFx_BAR0_CONTROL | 3 | HEX | BAR0 Control. Specifies the configuration of BAR 0. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable 110: Part of 64-bit memory BAR 0-1, non-prefetchable 111: Part of 64-bit memory BAR 0-1, prefetchable | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR0_APERTURE_SIZE | 5 | HEX | BAR0 Aperture. This attribute specifies the aperture of the 32-bit BAR 0 or 64-bit BAR 0-1. For Endpoints, the encodings are: • 00000: 128 bytes • 00001: 256 bytes • 00010: 512 bytes • 00101: 1 Kbytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00101: 4 Kbytes • 00101: 6 Kbytes • 01001: 64 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01001: 128 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 2 Mbytes • 01111: 4 Mbytes • 101001: 8 Mbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10011: 256 Mbytes • 10101: 26 Mbytes • 10101: 4 Gbytes • 10111: 1 Gbytes • 11001: 4 Gbytes • 11101: 28 Gbytes • 11101: 128 Gbytes • 11111: 256 Gbytes | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---------------------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR0_APERTURE_SIZE (Continued) | 5 | HEX | For Root Ports, the encodings are: • 00000: 4 bytes • 00001: 8 bytes • 00010: 16 bytes • 00011: 32 bytes • 00100: 84 bytes • 00101: 128 bytes • 00101: 256 bytes • 01000: 1 Kbytes • 01001: 2 Kbytes • 01001: 2 Kbytes • 01010: 4 Kbytes • 01010: 4 Kbytes • 01101: 32 Kbytes • 01101: 64 Kbytes • 01111: 128 Kbytes • 01111: 128 Kbytes • 10000: 256 Kbytes • 10001: 512 Kbytes • 10010: 1 Mbytes • 10010: 1 Mbytes • 10010: 2 Mbytes • 10101: 8 Mbytes • 10101: 8 Mbytes • 10101: 8 Mbytes • 10101: 128 Mbytes • 10111: 32 Mbytes • 10111: 32 Mbytes • 10111: 32 Mbytes • 11001: 128 Mbytes • 11001: 2 Gbytes • 11101: 2 Gbytes • 11101: 2 Gbytes • 11101: 4 Gbytes • 11111: 8 Gbytes | Y |
| PFx_BAR1_CONTROL | 3 | HEX | BAR1 Control. Specifies the configuration of BAR 1 when it is configured as a 32-bit BAR. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR1_APERTURE_SIZE | 5 | HEX | BAR1 Aperture. This attribute specifies the aperture of BAR 1 when it is configured as a 32-bit BAR. For Endpoints, the valid encodings are: • 00000: 128 bytes • 00001: 256 bytes • 00010: 512 bytes • 00010: 512 bytes • 00100: 2 Kbytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01010: 64 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01001: 512 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 2 Mbytes • 01111: 4 Mbytes • 10000: 8 Mbytes • 10001: 64 Kbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10101: 256 Mbytes • 10101: 256 Mbytes • 10101: 512 Mbytes • 10110: 512 Mbytes • 10111: 1 Gbytes | Y |

E XILINX.

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---------------------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR1_APERTURE_SIZE (Continued) | 5 | HEX | For Root Ports, the valid encodings are: • 00000: 4 bytes • 00011: 3 bytes • 00010: 16 bytes • 00101: 32 bytes • 00101: 32 bytes • 00101: 128 bytes • 00101: 256 bytes • 01001: 2 Kbytes • 01001: 2 Kbytes • 01001: 4 Kbytes • 01010: 4 Kbytes • 01101: 32 Kbytes • 01101: 32 Kbytes • 01111: 128 Kbytes • 01111: 128 Kbytes • 01111: 128 Kbytes • 10001: 512 Kbytes • 10001: 512 Kbytes • 10010: 1 Mbytes • 10010: 2 Kbytes • 10011: 2 Mbytes • 10011: 32 Kbytes • 10011: 32 Kbytes • 10011: 4 Mbytes • 10011: 512 Kbytes • 10101: 4 Mbytes • 10101: 4 Mbytes • 10101: 16 Mbytes • 10111: 32 Mbytes • 10111: 32 Mbytes • 11001: 128 Mbytes | Y |
| PFx_BAR2_CONTROL | 3 | HEX | BAR2 Control. Specifies the configuration of BAR 2. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable 110: Part of 64-bit memory BAR 2-3, non-prefetchable 111: Part of 64-bit memory BAR 2-3, prefetchable | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR2_APERTURE_SIZE | 5 | HEX | BAR2 Aperture. Specifies the aperture of the 32-bit BAR 2 or 64-bit BAR2–3. The encodings are: • 00000: 128 bytes • 00011: 256 bytes • 00010: 512 bytes • 00101: 1 Kbytes • 00101: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01001: 6 Kbytes • 01001: 6 Kbytes • 01001: 128 Kbytes • 01011: 256 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01111: 4 Mbytes • 01111: 4 Mbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10101: 256 Mbytes • 10101: 256 Mbytes • 10101: 266 Mbytes • 10101: 266 Mbytes • 10101: 276 Mbytes • 10101: 4 Gbytes • 10111: 1 Gbytes • 11001: 4 Gbytes • 11001: 4 Gbytes • 11101: 32 Gbytes • 11101: 64 Gbytes | Ν |
| PFx_BAR3_CONTROL | 3 | HEX | 11111: 256 Gbytes BAR3 Control. Specifies the configuration of BAR 3 when it is configured as a 32-bit BAR. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable | Ν |

E XILINX.

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR3_APERTURE_SIZE | 5 | HEX | BAR3 Aperture. Specifies the aperture of BAR 3 when it is configured as a 32-bit BAR. The valid encodings are: • 00000: 128 bytes • 00001: 256 bytes • 00010: 512 bytes • 00010: 512 bytes • 00100: 2 Kbytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01000: 32 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01010: 512 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 2 Mbytes • 01111: 4 Mbytes • 10000: 8 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10101: 1 Gbytes • 10111: 1 Gbytes • 11000: 2 Gbytes | N |
| PFx_BAR4_CONTROL | 3 | HEX | BAR4 Control. Specifies the configuration of BAR 4. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable 110: Part of 64-bit memory BAR 4-5, non-prefetchable 111: Part of 64-bit memory BAR 4-5, prefetchable | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR4_APERTURE_SIZE | 5 | HEX | BAR4 Aperture. Specifies the aperture of the 32-bit BAR 4 or 64-bit BAR4-5. The encodings are: • 00000: 128 bytes • 00011: 256 bytes • 00010: 512 bytes • 00101: 1 Kbytes • 00101: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01001: 32 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01011: 256 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01111: 4 Mbytes • 01111: 4 Mbytes • 10001: 8 Mbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10101: 4 Gbytes • 10101: 4 Gbytes • 11001: 4 Gbytes • 11101: 4 Gbytes | Ν |
| PFx_BAR5_CONTROL | 3 | HEX | BAR5 Control: Specifies the configuration of BAR 3 when it is configured as a 32-bit BAR. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable | Ν |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|--------------------------|-------------------|----------------|--|---------------------------|
| PFx_BAR5_APERTURE_SIZE | 5 | HEX | BAR5 Aperture. Specifies the aperture of BAR 5 when it is configured as a 32-bit BAR. The valid encodings are: • 00000: 128 bytes • 00001: 256 bytes • 00010: 512 bytes • 00010: 512 bytes • 00100: 2 Kbytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00110: 8 Kbytes • 00111: 16 Kbytes • 01000: 32 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01001: 512 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01111: 4 Mbytes • 10100: 8 Mbytes • 10001: 16 Mbytes • 10001: 64 Mbytes • 10011: 256 Mbytes • 10101: 256 Mbytes • 10111: 1 Gbytes • 10111: 1 Gbytes | Ν |
| PFx_EXPANSION_ROM_ENABLE | 1 | STRING | PFx Expansion ROM BAR Enable. This bit must be set to enable the Expansion ROM BAR associated with the Function. | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---------------------------------|-------------------|----------------|--|---------------------------|
| PFx_EXPANSION_ROM_APERTURE_SIZE | 5 | HEX | PFx Expansion ROM BAR Aperture Size. Encoding is as follows: 00000-00001: 256B 00010: 512 B 00011: 1 KB 00100: 2 KB 00101: 4 KB 00110: 8 KB 00110: 8 KB 01000: 32 KB 01001: 64 KB 01001: 128 KB 01101: 126 KB 01101: 1 MB 01110: 2 MB 01111: 4 MB 10000: 8 MB 10001: 16 MB 10010-11111: Reserved | Y |
| PCI Express Capability | | | | |
| PFx_DEV_CAP_MAX_PAYLOAD_SIZE | 3 | HEX | Capability Max Payload Size. This field indicates the maximum payload size that the Function can support for TLPs. Encodings are: • 000b: 128-byte maximum payload size • 001b: 256-byte maximum payload size • 010b: 512-byte maximum payload size • 011b: 1024-byte maximum payload size | Y |
| PF0_DEV_CAP_EXT_TAG_SUPPORTED | 1 | STRING | Extended Tags support. • FALSE: 5-bit tag • TRUE: 8-bit tag | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|--|-------------------|----------------|--|---------------------------|
| PF0_DEV_CAP_ENDPOINT_L0S_LATENCY | 3 | DECIMAL | Endpoint L0s Acceptable Latency. Records the latency the Endpoint can withstand on transitions from the L0s state to L0. Valid settings are: • 0h: Less than 64 ns • 1h: 64 ns to 128 ns • 2h: 128 ns to 256 ns • 3h: 256 ns to 512 ns • 4h: 512 ns to 1 µs • 5h: 1 µs to 2 µs • 6h: 2 µs to 4 µs • 7h: More than 4 µs For Endpoints only. Must be 0h for other devices. | Y |
| PF0_DEV_CAP_ENDPOINT_L1_LATENCY | 3 | DECIMAL | Endpoint L1 Acceptable Latency. Records the latency that the Endpoint can withstand on transitions from the L1 state to L0 (if the L1 state is supported). Valid settings are: • 0h: Less than 1 µs • 1h: 1 µs to 2 µs • 2h: 2 µs to 4 µs • 3h: 4 µs to 8 µs • 4h: 8 µs to 16 µs • 5h: 16 µs to 32 µs • 6h: 32 µs to 64 µs • 7h: Greater than 64 µs For Endpoints only. Must be 0h for other devices. | Y |
| PF0_DEV_CAP_FUNCTION_LEVEL_RESET_ CAPABLE | 1 | STRING | Function Level Reset. Set to TRUE when the device has Function-Level Reset capability. | |
| PF0_LINK_CAP_ASPM_SUPPORT | 2 | DECIMAL | Active State PM Support. Indicates the level of active state power management supported by the selected PCI Express Link, encoded as follows: • 00b: No ASPM • 01b: L0s supported • 10b: L1 supported • 11b: L0s and L1 entry supported. | Y |
| PF0_LINK_STATUS_SLOT_CLOCK_CONFIG | 1 | STRING | Slot Clock Configuration. Set to TRUE if the device uses a clock provided on the slot connector; otherwise set to FALSE if the device uses an independent clock irrespective of the presence of a reference on the connector. | Y |

Gen3 Integrated Block for PCIe (v1.3) PG023 October 16, 2012

EXILINX.

| Table D_1. | Virtox-7 EDGA Gon2 Integrated Block for E | PCI Express Core Attribute Descriptions (Cor | nt'd) |
|------------|---|--|-------|
| Tuble D-1. | Virtex-/ FPGA Gens integrated block for F | PCI Express Core Attribute Descriptions (Cor | πια |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---|-------------------|----------------|--|---------------------------|
| PF0_LINK_CAP_L0S_EXIT_LATENCY_ COMCLK_GEN1 | 3 | DECIMAL | Sets the exit latency from the L0s state to be applied (at 2.5 GT/s) when a common clock is used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_L0S_EXIT_LATENCY_ COMCLK_GEN2 | 3 | DECIMAL | Sets the exit latency from the L0s state to be applied (at 5 GT/s) when a common clock is used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_LOS_EXIT_LATENCY_ COMCLK_GEN3 | 3 | DECIMAL | Sets the exit latency from the L0s state to be applied (at 8 GT/s) when a common clock is used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_LOS_EXIT_LATENCY_GEN1 | 3 | DECIMAL | Sets the exit latency from the L0s state to be applied (at 2.5 GT/s) when separate clocks are used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_LOS_EXIT_LATENCY_GEN2 | 3 | DECIMAL | Sets the exit latency from the L0s state to be applied (at 5 GT/s) when separate clocks are used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_LOS_EXIT_LATENCY_GEN3 | 3 | DECIMAL | Sets the exit latency from the L0s state to be applied (at 8 GT/s) when separate clocks are used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_L1_EXIT_LATENCY_ COMCLK_GEN1 | 3 | DECIMAL | Sets the exit latency from the L1 state to be applied (at 2.5 GT/s) when a common clock is used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_L1_EXIT_LATENCY_ COMCLK_GEN2 | 3 | DECIMAL | Sets the exit latency from the L1 state to be applied (at 5 GT/s) when a common clock is used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_L1_EXIT_LATENCY_ COMCLK_GEN3 | 3 | DECIMAL | Sets the exit latency from the L1 state to be applied (at 8 GT/s) when a common clock is used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_L1_EXIT_LATENCY_GEN1 | 3 | DECIMAL | Sets the exit latency from the L1 state to be applied (at 2.5 GT/s) when separate clocks are used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_LINK_CAP_L1_EXIT_LATENCY_GEN2 | 3 | DECIMAL | Sets the exit latency from the L1 state to be applied (at 5 GT/s) when separate clocks are used. The latency is transferred to the Link Capabilities register. | Y |

www.xilinx.com

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode |
|--|-------------------|----------------|--|--------------------------|
| PF0_LINK_CAP_L1_EXIT_LATENCY_GEN3 | 3 | DECIMAL | Sets the exit latency from the L1 state to be applied (at 8 GT/s) where separate clocks are used. The latency is transferred to the Link Capabilities register. | Y |
| PF0_DEV_CAP2_CPL_TIMEOUT_DISABLE | 1 | STRING | Completion Timeout Disable Capable. When this attribute is set to TRUE, bit 4 is set to indicate that the associated Function supports the capability to turn off its Completion timeout. This attribute should be set to FALSE under normal operating conditions. | Y |
| PF0_DEV_CAP2_32B_ATOMIC_ COMPLETER_SUPPORT | 1 | STRING | 32-bit AtomicOp Completer Supported. When this attribute is set to TRUE, bit 7 is set to include FetchAdd, Swap, and CAS AtomicOps optional capabilities. | Y |
| PF0_DEV_CAP2_64B_ATOMIC_ COMPLETER_SUPPORT | 1 | STRING | 64-bit AtomicOp Completer Supported. When this attribute is set to TRUE, bit 8 is set to include FetchAdd, Swap, and CAS AtomicOps optional capabilities. | Y |
| PF0_DEV_CAP2_128B_CAS_ATOMIC_ COMPLETER_SUPPORT | 1 | STRING | 128-bit CAS Completer Supported. When this attribute is set to TRUE, bit 9 is set to enable optional capabilities. | Y |
| PF0_DEV_CAP2_LTR_SUPPORT | 1 | STRING | LTR Mechanism Supported. When this attribute is set to TRUE, bit 11 is set to indicate support for the optional Latency Tolerance Reporting (LTR) mechanism. | Y |
| PF0_DEV_CAP2_TPH_COMPLETER_ SUPPORT | 1 | STRING | TPH Completer Supported. This attribute sets bit 12 to indicate Completer support for TPH. Supported encodings are: 0b: TPH and Extended TPH Completer not supported 1b: TPH Completer supported | Y |
| PF0_DEV_CAP2_OBFF_SUPPORT | 2 | HEX | OBFF Supported. This attribute sets bits [19:18]. Encodings are: 00b: OBFF not supported 01b: OBFF supported using Message signaling only 10b: OBFF supported using WAKE# signaling only (not supported) 11b: OBFF supported using WAKE# and Message signaling (not supported) | Y |
| MSI Capability | | 1 | | |
| PFx_MSI_CAP_NEXTPTR | 8 | HEX | MSI Capability's Next Capability Offset pointer to the next item in the capabilities list, or 00h if this is the final capability. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode |
|-----------------------------|-------------------|----------------|--|--------------------------|
| zFx_MSI_CAP_MULTIMSGCAP | 3 | DECIMAL | Multiple Message Capable. This attribute sets bits [19:17] in the Control register. Each MSI function can request up to 32 unique messages. System software can read this field to determine the number of messages requested. The number of Messages requested are encoded as follows: • 0h: 1 vector • 1h: 2 vectors • 2h: 4 vectors • 3h: 8 vectors • 4h: 16 vectors • 5h: 32 vectors • 6h, 7h: Reserved | N |
| MSI-X Capability | | I | | |
| PFx_MSIX_CAP_NEXTPTR | 8 | HEX | MSI-X Capability's Next Capability Offset. This attribute contains the pointer to the next item in the capabilities list, or 00h if this is the final capability. | N |
| zFx_MSIX_CAP_PBA_BIR | 3 | DECIMAL | MSI-X Pending Bit Array BIR. This value is transferred to the MSI-X PBA BIR field. It is set to 0 if MSI-X is disabled. | N |
| zFx_MSIX_CAP_PBA_OFFSET | 29 | HEX | MSI-X Pending Bit Array Offset. This value is transferred to the MSI-X PBA Offset field. This attribute is set to 0 if MSI-X is disabled. | N |
| zFx_MSIX_CAP_TABLE_BIR | 3 | DECIMAL | MSI-X Table BIR. This value is transferred to the MSI-X Table BIR field. This attribute is set to 0 if MSI-X is disabled. | N |
| zFx_MSIX_CAP_TABLE_OFFSET | 29 | HEX | MSI-X Table Offset. This value is transferred to the MSI-X Table Offset field. This attribute is set to 0 if MSI-X is disabled. | N |
| zFx_MSIX_CAP_TABLE_SIZE | 11 | HEX | MSI-X Table Size. This value is transferred to the MSI-X Message Control[10:0] field. This attribute is set to 0 if MSI-X is disabled. The table must be implemented in user logic. | N |
| PM Capability | | | | |
| zFx_PM_CAP_ID | 8 | HEX | PM Capability ID. This attribute indicates that the capability structure is for Power Management. | N |
| zFx_PM_CAP_NEXTPTR | 8 | HEX | PM Capability Next Cap Pointer. Contains the pointer to the next PCI Capability Structure. | N |
| PF0_PM_CAP_PMESUPPORT_D3HOT | 1 | STRING | PME Support for D3hot State. This attribute sets bit 14 of the PMC register when TRUE. All Functions assume the value is programmed into PF0. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|--------------------------|-------------------|----------------|---|---------------------------|
| PF0_PM_CAP_PMESUPPORT_D1 | 1 | STRING | PME Support for D1 State. This attribute sets bit 12 of the PMC register when TRUE. All functions assume the value is programmed into PF0. | N |
| PF0_PM_CAP_PMESUPPORT_D0 | 1 | STRING | PME Support for D0 State. This attribute sets bit 11 of the PMC register when TRUE. All functions assume the value is programmed into PF0. | N |
| PF0_PM_CAP_SUPP_D1_STATE | 1 | STRING | D1_Support for D0 State. This attribute sets bit 9 of the PMC register when TRUE. All functions assume the value is programmed into PF0. | N |
| zFx_PM_CAP_VER_ID | 3 | HEX | Version of PM Specification. Indicates which version of the PCI Bus Power Management Specifications that the Function implements. | N |
| PF0_PM_CSR_NOSOFTRESET | 1 | STRING | No_Soft_Reset. Power Management CSR [3], No Soft Reset bit. All functions assume value programmed into PF0. | N |
| RBAR Capability | | | | |
| PFx_RBAR_CAP_ENABLE | 1 | STRING | Enable Resizable BAR Capability. When this attribute is set to TRUE, RBAR Capability functionality is enabled. | N |
| PFx_RBAR_CAP_VER | 4 | HEX | RBAR Capability Version. | N |
| PFx_RBAR_CAP_NEXTPTR | 12 | HEX | Resizable BAR Capability's Next Capability Offset pointer to the next item in the capabilities list, or 000h if this is the final capability. | N |
| PFx_RBAR_NUM | 3 | НЕХ | The number of Resizable BARs in the cap structure. The value is transferred to the Resizable BAR Control Register 0, bits [7:5]. | N |
| PFx_RBAR_CAP_INDEX0 | 3 | HEX | BAR Index Value for Resizable BAR Control Register 0. This value must be the lowest BAR Index of the resizable BARs. | N |
| PFx_RBAR_CAP_SIZE0 | 20 | HEX | BAR Size Supported Vector for Resizable BAR Capability Register 0. Bits [3:0] and [31:24] should always be set to 0. | N |
| PFx_RBAR_CAP_INDEX1 | 3 | HEX | BAR Index Value for Resizable BAR Control Register 1. This attribute is set to 0 if one or fewer BARs can be resized. This value should not be lower than the value on RBAR_CAP_INDEX0. | N |
| PFx_RBAR_CAP_SIZE1 | 20 | HEX | BAR Size Supported Vector for Resizable BAR Capability Register 1. Bits [3:0] and [31:24] should always be set to 0. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode |
|--------------------------------|-------------------|----------------|--|--------------------------|
| PFx_RBAR_CAP_INDEX2 | 3 | HEX | BAR Index Value for Resizable BAR Control Register 2. This attribute is set to 0 if two or fewer BARs can be resized. This value should not be lower than the value on RBAR_CAP_INDEX0, 1. | N |
| PFx_RBAR_CAP_SIZE2 | 20 | HEX | BAR Size Supported Vector for Resizable BAR Capability Register 2. Bits [3:0] and [31:24] should always be set to 0. | N |
| Root Port Capability | I | I | · | 1 |
| DNSTREAM_LINK_NUM | 8 | HEX | Used in downstream facing mode only. Specifies the link number that this device advertises in TS1 and TS2 during link training. | Y |
| DSN Capability | | | | |
| PFx_DSN_CAP_NEXTPTR | 12 | HEX | Device Serial Number's Capability's Next Capability Offset Pointer. This attribute contains the pointer to the next item in the capabilities list, or 000h if this is the final capability. | N |
| VC Capability (PF0 only) | | 1 | | |
| PF0_VC_CAP_VER | 4 | HEX | VC Capability Version | Ν |
| PF0_VC_CAP_NEXTPTR | 12 | HEX | VC Next Capability Pointer | N |
| AER Capability | | | 1 | |
| PFx_AER_CAP_NEXTPTR | 12 | HEX | AER's Next Capability Offset Pointer. This attribute contains the pointer to the next item in the capabilities list, or 000h if this is the final capability. | Y |
| PFx_AER_CAP_ECRC_CHECK_CAPABLE | 1 | STRING | This attribute indicates that the core can check ECRC. The value is transferred to bit 7 of the AER Capabilities and Control register. | Y |
| PFx_AER_CAP_ECRC_GEN_CAPABLE | 1 | STRING | This attribute indicates that the core can generate ECRC. The value is transferred to bit 5 of the AER Capabilities and Control register. | Y |
| ARI Capability | | | | |
| ARI_CAP_ENABLE | 1 | STRING | Enable ARI Capability. When this attribute is FALSE, the legacy interpretation of PCI RID {8-bit Bus number, 5-bit device number, 3-bit Function number} is enabled. When this attribute is TRUE, an alternate interpretation of PCI RID {8-bit Bus number, 8-bit Function number} is enabled. | N |
| zFx_ARI_CAP_NEXTPTR | 12 | HEX | ARI Next Capability Offset. Bits [31:20] of the ARI Extended Capability Header Register. | N |
| PF0_ARI_CAP_VER | 4 | HEX | ARI Capability Version. Bits [19:16] of the ARI Extended Capability Header Register. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|----------------------------------|-------------------|----------------|--|---------------------------|
| PFx_ARI_CAP_NEXT_FUNC | 8 | HEX | ARI Next Function. Bits [15:8] of the ARI Capability and Control Registers. This attribute points to the next Physical Function in the device. | N |
| Power Budget Enhanced Capability | | | | |
| PFx_PB_CAP_NEXTPTR | 12 | HEX | Power Budget Next Capability Offset. Bits [31:20] of the ARI Extended Capability Header Register. | N |
| PFx_PB_CAP_VER | 4 | HEX | Power Budget Capability Version. Bits [19:16] of the PB Extended Capability Header Register. | N |
| PFx_PB_CAP_SYSTEM_ALLOCATED | 1 | STRING | Power Budget System Allocated. This attribute sets bit 0 in the Power Budget Capability register. This bit is set to indicate that the device power specified by this Power Management Capability Structure is included in the system power budget. When this bit set, the software must exclude the device power reported by this Capability Structure from power calculations, when making power budgeting decisions. | N |
| LTR Capability (PF0 only) | 1 | I | · | J |
| PF0_LTR_CAP_NEXTPTR | 12 | HEX | LTR Next Capability Offset. Bits [31:20] of the LTR Extended Capability Header Register. | N |
| PF0_LTR_CAP_VER | 4 | HEX | LTR Capability Version. Bits [19:16] of the LTR Extended Capability Header Register. | N |
| PF0_LTR_CAP_MAX_SNOOP_LAT | 10 | HEX | LTR Capability Maximum Snoop Latency. Bits [9:0] of the LTR Max Snoop/Max No-Snoop Latency register. System software sets this attribute. Non-default values should not be set under normal operating conditions. | N |
| PF0_LTR_CAP_MAX_NOSNOOP_LAT | 10 | HEX | LTR Capability Max No Snoop Latency. Bits [25:16] of the LTR Max Snoop/Max No-Snoop Latency Register. The use model is for system software to set this field, and non-default values should not be set under normal operating conditions. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---|-------------------|----------------|--|---------------------------|
| LTR_TX_MESSAGE_ON_LTR_ENABLE | 1 | STRING | When this attribute is TRUE, the core automatically transmits an LTR message whenever the LTR Mechanism Enable bit in the Device Control 2 Register changes from 0 to 1, with the parameters specified in the LTR Snoop/No-Snoop Latency Register. When this bit is TRUE, the core also transmits an LTR message whenever the LTR Mechanism Enable bit is cleared, if these conditions are both true: The core sent at least one LTR message since the LTR Mechanism Enable bit was last set. The most recent LTR message transmitted by the core had at least one of the Requirement bits set. The core sets the Requirement bits in this LTR message to 0. When this bit is 0, the core does not, by itself, send any LTR messages in response to state changes of the LTR Mechanism Enable bit. You can monitor the state of the cfg_ltr_enable output of the core and transmit LTR messages through the AXI4-Stream interface in response to its state changes. | Ν |
| LTR_TX_MESSAGE_ON_FUNC_POWER_ STATE_CHANGE | 1 | STRING | When this attribute is TRUE, the core automatically transmits an LTR message when all Functions in the core have transitioned to a non-D0 power state, provided that the following conditions are both true: The core sent at least one LTR message since the Data Link layer last transitioned from down to up state. The most recent LTR message transmitted by the core had at least one of the Requirement bits set. The core sets the Requirement bits in this LTR message to 0. When this bit 12 is 0, the core does not, by itself, send any LTR messages in response to Function Power State changes. You can monitor the cfg_function_power_state outputs of the core and transmit LTR messages through the AXI4-Stream interface in response to changes in their states. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---|-------------------|----------------|---|---------------------------|
| LTR_TX_MESSAGE_MINIMUM_INTERVAL | 10 | HEX | This attribute specifies the minimum spacing between LTR messages transmitted by the core in microseconds. The PCI Express Specifications recommend sending no more than two LTR messages within a 500 microsecond interval. The core waits for the minimum delay specified by this field after sending an LTR message before transmitting a new LTR message. | Ν |
| DPA Capability | | 1 | | |
| PFx_DPA_CAP_NEXTPTR | 12 | HEX | DPA Next Capability Offset. Bits [31:20] of the DPA Extended Capability Header Register. | N |
| PFx_DPA_CAP_VER | 4 | HEX | DPA Capability Version. Bits [19:16] of the DPA Extended Capability Header Register. | N |
| PFx_DPA_CAP_SUB_STATE_CONTROL_EN | 1 | STRING | DPA Substate Control Enable. This attribute sets bit 8 in the DPA Control and Status register. This bit enables the Substate Control field. This bit is initialized to 1 by the hardware on a hard reset or a Function-Level Reset. Software can clear this bit by writing a 1 to this bit position, but cannot set this bit directly through a configuration write. Clearing this bit disables the Substate Control field, thus preventing further substate transitions for this Function. | Ν |
| PFx_DPA_CAP_SUB_STATE_CONTROL | 5 | НЕХ | DPA Substate Control. This attribute sets bits [20:16] in the DPA Control and Status register. Used by software to configure the Function substate. Software writes the substate value in this field to initiate a substate transition. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION0 | 8 | HEX | DPA Substate Power State Allocation 0. This attribute contains the power allocation for the first DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION1 | 8 | HEX | DPA Substate Power State Allocation 1. This attribute contains the power allocation for the second DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|---|-------------------|----------------|--|---------------------------|
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION2 | 8 | HEX | DPA Substate Power State Allocation 2. This attribute contains the power allocation for the third DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION3 | 8 | HEX | DPA Substate Power State Allocation 3. This attribute contains the power allocation for the fourth DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION4 | 8 | HEX | DPA Substate Power State Allocation 4. This attribute contains the power allocation for the fifth DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION5 | 8 | HEX | DPA Substate Power State Allocation 5. This field contains the power allocation for the sixth DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION6 | 8 | HEX | DPA Substate Power State Allocation 6. This attribute contains the power allocation for the seventh DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |
| PFx_DPA_CAP_SUB_STATE_POWER_ ALLOCATION7 | 8 | HEX | DPA Substate Power State Allocation 7. This attribute contains the power allocation for the eighth DPA substate covered by this register. This value, when multiplied by the Power Allocation Scale programmed in the DPA Capability register, provides the power associated with the corresponding substate in watts. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|-------------------------------|-------------------|----------------|---|---------------------------|
| SRIOV Capability (PF0 only) | | | | |
| SRIOV_CAP_ENABLE | 1 | STRING | Enable SRIOV Capability. The Single Root I/O Virtualization feature is enabled when this attribute is set to TRUE. ARI_CAP_ENABLE must also be set to TRUE to enable this feature. | N |
| PFx_SRIOV_CAP_NEXTPTR | 12 | HEX | SRIOV Next Capability Offset. Bits [31:20] of the SRIOV Extended Capability Header Register. | N |
| PFx_SRIOV_CAP_VER | 4 | HEX | SRIOV Capability Version. Bits [19:16] of the SRIOV Extended Capability Header Register. | N |
| PFx_SRIOV_CAP_INITIAL_VF | 16 | HEX | Initial Number of VFs. This attribute contains the initial number of VFs configured for PF0. | N |
| PFx_SRIOV_CAP_TOTAL_VF | 16 | HEX | Total Number of VFs. This attribute contains the total number of VFs configured for PF0. The value must be equal to PF0_SRIOV_CAP_INITIAL_VF. | N |
| PFx_SRIOV_FUNC_DEP_LINK | 16 | HEX | Physical Function Dependency Link. This attribute is used to specify dependencies between PFs. The programming model for a Device can have vendor-specific dependencies between sets of Functions. The Function Dependency Link field is used to describe these dependencies. | N |
| PFx_SRIOV_FIRST_VF_OFFSET | 16 | HEX | Offset of first VF, relative to PF0. The default value depends on the setting of VF_MODE. The offset is such that the Routing IDs of the VFs are mapped in the range: • PF0 VFs: 64-69 • PF1 VFs: 63-68 | N |
| PFx_SRIOV_VF_DEVICE_ID | 16 | HEX | VF Device ID assigned to device. This field contains the Device ID that should be presented for every VF to the single root I/O virtualization (SR-IOV). | N |
| PFx_SRIOV_SUPPORTED_PAGE_SIZE | 32 | HEX | Page Size Supported by Device. This field indicates the page sizes supported by the PF. This PF supports a page size of $2n+12$, if bit <i>n</i> is set. | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------|-------------------|----------------|---|---------------------------|
| PFx_SRIOV_BAR0_CONTROL | 3 | HEX | BAR0 Control. Specifies the configuration of BAR 0. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable 110: Part of 64-bit memory BAR 0-1, non-prefetchable 111: Part of 64-bit memory BAR 0-1, prefetchable | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|--|---------------------------|
| PFx_SRIOV_BAR0_APERTURE_SIZE | 5 | HEX | BAR0 Aperture. Specifies the aperture of the 32-bit BAR 0 or 64-bit BAR 0–1. The encodings are: • 00000: 128 bytes • 00011: 28 bytes • 00010: 512 bytes • 00101: 1 Kbytes • 00101: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01001: 6 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01001: 128 Kbytes • 01011: 256 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01111: 4 Mbytes • 01111: 4 Mbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10101: 4 Mbytes • 10101: 4 Gbytes • 10111: 1 Gbytes • 11001: 4 Gbytes • 11001: 4 Gbytes • 11001: 4 Gbytes • 11101: 4 Gbytes | Ν |
| PFx_SRIOV_BAR1_CONTROL | 3 | HEX | BAR1 Control - Specifies the configuration of BAR 1 when it is configured as a 32-bit BAR. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable | N |

E XILINX.

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|--|---------------------------|
| PFx_SRIOV_BAR1_APERTURE_SIZE | 5 | HEX | BAR1 Aperture. Specifies the aperture of BAR 1 when it is configured as a 32-bit BAR. The valid encodings are: • 00000-00011: 1 Kbytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01001: 32 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01010: 128 Kbytes • 01010: 512 Kbytes • 01101: 1 Mbytes • 01111: 4 Mbytes • 01111: 4 Mbytes • 10000: 8 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10101: 256 Mbytes • 10101: 256 Mbytes • 10101: 128 Mbytes • 10101: 256 Mbytes • 10111: 1 Gbytes • 11000: 2 Gbytes | N |
| PFx_SRIOV_BAR2_CONTROL | 3 | HEX | BAR2 Control. Specifies the configuration of BAR 2. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable 110: Part of 64-bit memory BAR 0–1, non-prefetchable 111: Part of 64-bit memory BAR 0–1, prefetchable | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|---|---------------------------|
| PFx_SRIOV_BAR2_APERTURE_SIZE | 5 | HEX | BAR2 Aperture. Specifies the aperture of the 32-bit BAR 2 or 64-bit BAR2–3. The encodings are: • 00000: 128 bytes • 00011: 256 bytes • 00010: 512 bytes • 00101: 1 Kbytes • 00101: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01001: 64 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01010: 512 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 2 Mbytes • 01111: 4 Mbytes • 10001: 8 Mbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10101: 4 Gbytes • 10101: 4 Gbytes • 11001: 4 Gbytes • 11001: 4 Gbytes • 11101: 46 Gbytes • 11101: 128 Gbytes | Ν |
| PFx_SRIOV_BAR3_CONTROL | 3 | HEX | BAR3 Control. Specifies the configuration of BAR 3 when it is configured as a 32-bit BAR. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable | N |

E XILINX.

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|---|---------------------------|
| PFx_SRIOV_BAR3_APERTURE_SIZE | 5 | HEX | BAR3 Aperture. Specifies the aperture of BAR 3 when it is configured as a 32-bit BAR. The valid encodings are: • 00000: 128 bytes • 00001: 256 bytes • 00010: 512 bytes • 00010: 512 bytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00101: 8 Kbytes • 00101: 8 Kbytes • 01000: 32 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01010: 512 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01110: 2 Mbytes • 01111: 4 Mbytes • 10000: 8 Mbytes • 10001: 16 Mbytes • 10010: 32 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 256 Mbytes • 10101: 256 Mbytes • 10111: 1 Gbytes • 10111: 1 Gbytes | Ν |
| PFx_SRIOV_BAR4_CONTROL | 3 | HEX | BAR4 Control. Specifies the configuration of BAR 4. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable 110: Part of 64-bit memory BAR 4–5, non-prefetchable 111: Part of 64-bit memory BAR 4–5, prefetchable | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|---|---------------------------|
| PFx_SRIOV_BAR4_APERTURE_SIZE | 5 | HEX | BAR4 Aperture. Specifies the aperture of the 32-bit BAR 4 or 64-bit BAR4–5. The encodings are: • 00000: 128 bytes • 00011: 256 bytes • 00010: 512 bytes • 00101: 1 Kbytes • 00101: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01001: 6 Kbytes • 01001: 6 Kbytes • 01001: 128 Kbytes • 01001: 128 Kbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01101: 1 Mbytes • 01111: 4 Mbytes • 01111: 4 Mbytes • 10001: 8 Mbytes • 10001: 16 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10011: 512 Mbytes • 10011: 64 Mbytes • 10011: 64 Mbytes • 10101: 128 Mbytes • 10101: 256 Mbytes • 10101: 256 Mbytes • 10101: 256 Mbytes • 10101: 4 Gbytes • 11001: 4 Gbytes • 11001: 4 Gbytes • 11101: 16 Gbytes • 11101: 128 Gbytes • 11101: 128 Gbytes | Ν |
| PFx_SRIOV_BAR5_CONTROL | 3 | HEX | 11111: 256 Gbytes BAR5 Control. Specifies the configuration of BAR 3 when it is configured as a 32-bit BAR. The encodings are: 000: Disabled 001: 32-bit I/O BAR 010-011: Reserved 100: 32-bit memory BAR, non-prefetchable 101: 32-bit memory BAR, prefetchable | N |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|------------------------------|-------------------|----------------|--|---------------------------|
| PFx_SRIOV_BAR5_APERTURE_SIZE | 5 | HEX | BAR5 Aperture. Specifies the aperture of BAR 5 when it is configured as a 32-bit BAR. The valid encodings are: • 00000: 128 bytes • 00001: 256 bytes • 00010: 512 bytes • 00010: 512 bytes • 00100: 2 Kbytes • 00100: 2 Kbytes • 00101: 4 Kbytes • 00111: 16 Kbytes • 01000: 32 Kbytes • 01001: 64 Kbytes • 01001: 128 Kbytes • 01010: 512 Kbytes • 01101: 1 Mbytes • 01110: 2 Mbytes • 01111: 4 Mbytes • 10000: 8 Mbytes • 10001: 16 Mbytes • 10011: 64 Mbytes • 10101: 128 Mbytes • 101011: 128 Mbytes • 101011: 128 Mbytes • 101011: 128 Mbytes | N |
| TPH Requester Capability | 1 | | | 1 |
| zFx_TPHR_CAP_NEXTPTR | 12 | HEX | TPHR Next Capability Offset. Bits [31:20] of the TPHR Extended Capability Header Register. | N |
| zFx_TPHR_CAP_VER | 4 | HEX | TPHR Capability Version. Bits [19:16] of the TPHR Extended Capability Header Register. | N |
| zFx_TPHR_CAP_INT_VEC_MODE | 1 | STRING | Interrupt Vector Mode Supported. Bit 1 in the TPH Requester Capability Register. A setting of TRUE indicates that the Function supports the Interrupt Vector Mode for TPH Steering Tag generation. In the Interrupt Vector Mode, Steering Tags are attached to MSI/MSI-X interrupt requests. The Steering Tag for each interrupt request is selected by the MSI/ MSI-X interrupt vector number. | Ν |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|--------------------------------|-------------------|----------------|--|---------------------------|
| zFx_TPHR_CAP_DEV_SPECIFIC_MODE | 1 | STRING | Device Specific Mode Supported. Bit 2 in the TPH Requester Capability Register. A setting of TRUE indicates that the Function supports the Device-Specific Mode for TPH Steering Tag generation. In this mode, the Steering Tags are supplied by the client for each request through the AXI4-Stream interface. The client typically chooses the Steering Tag values from the ST Table, but is not required to do so. | N |
| zFx_TPHR_CAP_ST_TABLE_LOC | 2 | HEX | Steering Tag Table Location. Bits [10:0] in the TPH Requester Capability Register. This field indicates if a Steering Tag table is implemented for this Function and its location, if present. 00b: Steering Tag table is not present 01b: Steering Tag table is in the TPH Requester Capability Structure 10b: Steering Tag values stored in the MSI-X table in used memory space 11b: Reserved | N |
| zFx_TPHR_CAP_ST_TABLE_SIZE | 11 | HEX | Steering Tag Table Size. Sets bits [26:16] in the TPH Requester Capability Register. The value indicates the maximum number of Steering Tag table entries the Function can use. Software reads this field to determine the Steering Tag Table Size N, which is encoded as N - 1. There is an upper limit of 64 entries when the Steering Tag table is located in the TPH Requester Capability structure. There is an upper limit of 2048 entries when the Steering Tag table is located in the MSI-X table. | N |
| zFx_TPHR_CAP_ST_MODE_SEL | 3 | HEX | Steering Tag Mode Select. Sets bits [2:0] in the TPH Requester Control Register. Selects the Steering Tag mode of operation. Valid encodings are: 000b: No ST Mode 001b: Interrupt Vector Mode 010b: Device-Specific Mode All other encodings are reserved. | Ν |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|--------------------------------------|-------------------|----------------|---|---------------------------|
| zFx_TPHR_CAP_ENABLE | 1 | STRING | TPH Requester Enable. Sets bit 8 in the TPH Requester Control Register. This attribute controls the ability to issue Request TLPs using either TPH. Encodings are: FALSE (0): Function operating as a Requester is not permitted to issue Requests with TPH or Extended TPH. TRUE (1): Function operating as a Requester is permitted to issue Requests with TPH and is not permitted to issue Requests with Extended TPH. The use model is for system software to set this field, and non-default values should not be set under normal operating conditions. | Ν |
| Gen3 PCS | | | | |
| GEN3_PCS_AUTO_REALIGN | 2 | HEX | Gen3 PCS Sync Header Error Control. Valid encodings are: O: Any RxSyncHeader error causes the RX path to realign 1: Any two RxSyncHeader errors over 8 data beats causes the RX path to realign 2: Any two RxSyncHeader errors over 64 data beats causes the RX path to realign | Y |
| GEN3_PCS_RX_ELECIDLE_INTERNAL | 1 | STRING | Gen3 PCS Rx Electrical Idle Internal. When this attribute is TRUE, per lane Rx Electrical Idle is generated internally (the pipe_rxn_elec_idle input is ignored). When this attribute is FALSE, per lane Rx Electrical Idle pipe_rxn_elec_idle is used at Gen3 speeds. | Y |
| Functional Test and Debug | | | | |
| PL_DISABLE_SCRAMBLING | 1 | STRING | When this attribute is TRUE, the scrambler and descrambler in the Physical Layer are disabled at Gen1 and Gen2 speeds. This attribute is used for test and debug only. | Y |
| Pin Characterization Mode and Spares | | 1 | 1 | 1 |
| TEST_MODE_PIN_CHAR | 1 | STRING | Pin Characterization Test mode. When this attribute is TRUE, Input to Output paths are enabled for pin characterization. | Y |
| SPARE_BIT0 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT1 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT2 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT3 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT4 | 1 | DECIMAL | Spare attribute | Y |

| Attribute Name | Number of Bits | Attribute Type | Attribute Description | Applicable to RP Mode? |
|----------------|-------------------|----------------|-----------------------|---------------------------|
| SPARE_BIT5 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT6 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT7 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BIT8 | 1 | DECIMAL | Spare attribute | Y |
| SPARE_BYTE0 | 8 | HEX | Spare attribute | Y |
| SPARE_BYTE1 | 8 | HEX | Spare attribute | Y |
| SPARE_BYTE2 | 8 | HEX | Spare attribute | Y |
| SPARE_BYTE3 | 8 | HEX | Spare attribute | Y |
| SPARE_WORD0 | 32 | HEX | Spare attribute | Y |
| SPARE_WORD1 | 32 | HEX | Spare attribute | Y |
| SPARE_WORD2 | 32 | HEX | Spare attribute | Y |
| SPARE_WORD3 | 32 | HEX | Spare attribute | Y |



Appendix E

Additional Resources

Xilinx Resources

For support resources such as Answers, Documentation, Downloads, and Forums, see the Xilinx Support website at:

www.xilinx.com/support.

For a glossary of technical terms used in Xilinx documentation, see:

www.xilinx.com/company/terms.htm.

Solution Centers

See the <u>Xilinx Solution Centers</u> for support on devices, software tools, and intellectual property at all stages of the design cycle. Topics include design assistance, advisories, and troubleshooting tips. The Solution Center specific to the Gen3 Integrated Block for PCIe core is located at <u>Xilinx PCI Express Solution Center</u>.

References

These documents provide supplemental material useful with this product guide:

- 1. AMBA AXI4-Stream Protocol Specification
- 2. PCI-SIG® Documentation (www.pcisig.com/specifications)
 - PCI Express Base Specification, rev. 3.0
 - PCI Express Card Electromechanical (CEM) Specification 3.0
- 3. Virtex-7 FPGA Documentation (www.xilinx.com/support/documentation/7_series.htm)
 - DS183, Virtex-7 FPGAs Data Sheet: DC and Switching Characteristics
 - UG470, 7 Series FPGAs Configuration User Guide

- UG471, 7 Series FPGAs SelectIO Resources User Guide
- UG472, 7 Series FPGAs Clocking Resources User Guide
- UG476, 7 Series FPGAs GTX/GTH Transceivers User Guide
- 4. Vivado[™] Design Suite 2012.3 Documentation

Technical Support

Xilinx provides technical support at <u>www.xilinx.com/support</u> for this LogiCORE[™] IP product when used as described in the product documentation. Xilinx cannot guarantee timing, functionality, or support of product if implemented in devices that are not defined in the documentation, if customized beyond that allowed in the product documentation, or if changes are made to any section of the design labeled DO NOT MODIFY.

See the IP Release Notes Guide (<u>XTP025</u>) for more information on this core. For each core, there is a master Answer Record that contains the Release Notes and Known Issues list for the core being used. The following information is listed for each version of the core:

- New Features
- Resolved Issues
- Known Issues

Revision History

The following table shows the revision history for this document.

| Date | Version | Revision |
|----------|---------|-------------------------|
| 04/24/12 | 1.0 | Initial Xilinx release. |

| Date | Version | Revision |
|----------|---------|---|
| 07/25/12 | 1.1 | Updated core to v1.2 and ISE Design Suite version to 14.2. Added support for Vivado Design Suite 2012.2. Added Endpoint Model Test Bench for Root Port in Chapter 9. |
| 10/16/12 | 1.2 | Updated core to v1.3, ISE Design Suite to 14.3 and Vivado Design Suite to 2012.3. New screenshots and descriptions in Chapter 4, Customizing and Generating the Core. Removed XC7VH290T Updated description for cfg_mgmt_addr, cfg_ltssm_state, and cfg_interrupt_msi_int. Added Table 2-18, and made major updates to Table 2-24. Added Target Function Value to PF/VF map mappings (Table 3-3). Added note to contact Xilinx regarding Root Port configuration availability. Added PIPE MODE Simulation section in Chapter 6 and Chapter 9. Added new PIO write address and write data numbers in Programmed Input/Output: Endpoint Example Design. Updated description for PFx_SRIOV_FIRST_VF_OFFSET attribute (Table D-1). |

Notice of Disclaimer

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at http://www.xilinx.com/warranty.htm; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: http://www.xilinx.com/warranty.htm#critapps.

© Copyright 2012 Xilinx, Inc. Xilinx, the Xilinx logo, Artix, ISE, Kintex, Spartan, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. PCI, PCIe and PCI Express are trademarks of PCI-SIG and used under license. All other trademarks are the property of their respective owners.