

Xilinx Solutions Guide for PCI Express

User Guide

UG493 (v1.0) July 18, 2008





Xilinx is disclosing this user guide, manual, release note, and/or specification (the "Documentation") to you solely for use in the development of designs to operate with Xilinx hardware devices. You may not reproduce, distribute, republish, download, display, post, or transmit the Documentation in any form or by any means including, but not limited to, electronic, mechanical, photocopying, recording, or otherwise, without the prior written consent of Xilinx. Xilinx expressly disclaims any liability arising out of your use of the Documentation. Xilinx reserves the right, at its sole discretion, to change the Documentation without notice at any time. Xilinx assumes no obligation to correct any errors contained in the Documentation, or to advise you of any corrections or updates. Xilinx expressly disclaims any liability in connection with technical support or assistance that may be provided to you in connection with the Information.

THE DOCUMENTATION IS DISCLOSED TO YOU "AS-IS" WITH NO WARRANTY OF ANY KIND. XILINX MAKES NO OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED, OR STATUTORY, REGARDING THE DOCUMENTATION, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS. IN NO EVENT WILL XILINX BE LIABLE FOR ANY CONSEQUENTIAL, INDIRECT, EXEMPLARY, SPECIAL, OR INCIDENTAL DAMAGES, INCLUDING ANY LOSS OF DATA OR LOST PROFITS, ARISING FROM YOUR USE OF THE DOCUMENTATION.

© 2008 Xilinx, Inc. All rights reserved.

XILINX, the Xilinx logo, the Brand Window, and other designated brands included herein are trademarks of Xilinx, Inc. PCI, PCI-SIG, PCI EXPRESS, PCIE, PCI-X, PCI HOT PLUG, MINI PCI, EXPRESSMODULE, and the PCI, PCI-X, PCI HOT PLUG, and MINI PC design marks are trademarks, registered trademarks, and/or service marks of PCI-SIG. All other trademarks are the property of their respective owners.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
07/18/08	1.0	Initial Xilinx release.

Table of Contents

Preface: About This Guide

Guide Contents	5
Additional Support Resources	5
Typographical Conventions	5
Online Document	6

Xilinx Solutions Guide for PCI Express

Overview	7
Xilinx Solutions	7
Transaction Layer	8
Data Link Layer	8
Physical Layer	8
Configuration Management	9
Endpoint Block Plus Wrapper for PCI Express	9
Components of a Xilinx Design for PCI Express	10
Board Design	10
Hardware HDL Design	11
Driver Design	11
Software Application	11
Documentation for PCI Express Solutions	12
Endpoint Block Plus Wrapper for PCI Express	12
Soft IP Implementation for PCI Express	12
Endpoint PIPE Implementation for PCI Express	13
Release Notes and Known Issues	13
Application Notes and Reference Designs	13
XAPP859	13
XAPP1002	14
XAPP1022	14
XAPP1052	14
More Reference Designs	14
Other Reference Material	14
PCI Express Base Specification	14
PCI Specification	14
PCI Express Card Electro Mechanical Specification	14
PCI Express Developer Tools	15
FPGA Documentation	15
Virtex-5 Documentation	15
Virtex-4 Documentation	15
Virtex-II Pro Documentation	15
Spartan Documentation	15
Device Package User Guide	15
Customer Education Resources	16
License Information	16
Virtex-5 Integrated Block Known Issues	16
Virtex-5 Integrated Block Known Restrictions Matrix	17

Virtex-5 GTP/GTX Known Restrictions for PCI Express Matrix	19
Conclusion	22

About This Guide

This guide provides technical solutions for PCI Express®.

Guide Contents

This user guide contains the following sections:

- “Overview”
- “Xilinx Solutions”
- “Components of a Xilinx Design for PCI Express”
- “Documentation for PCI Express Solutions”
- “License Information”
- “Virtex-5 Integrated Block Known Restrictions Matrix”
- “Virtex-5 GTP/GTX Known Restrictions for PCI Express Matrix”
- “Conclusion”

Additional Support Resources

To search the database of silicon and software questions and answers, or to create a technical support case in WebCase, see the Xilinx website at:

<http://www.xilinx.com/support>.

Typographical Conventions

This document uses the following typographical conventions. An example illustrates each convention.

Convention	Meaning or Use	Example
<i>Italic font</i>	References to other documents	See the <i>Virtex-5 Configuration Guide</i> for more information.
	Emphasis in text	The address (F) is asserted <i>after</i> clock event 2.
<u>Underlined Text</u>	Indicates a link to a web page.	http://www.xilinx.com/virtex5

Online Document

The following conventions are used in this document:

Convention	Meaning or Use	Example
Blue text	Cross-reference link to a location in the current document	See the section “ Additional Support Resources ” for details. Refer to “ Overview ,” page 7 for details.
Red text	Cross-reference link to a location in another document	See Figure 2 in the <i>Virtex-5 Data Sheet</i>
Blue, underlined text	Hyperlink to a website (URL)	Go to http://www.xilinx.com for the latest documentation.

Xilinx Solutions Guide for PCI Express

Overview

Xilinx provides a variety of solutions for PCI Express to enable customers to build PCI Express designs leveraging the flexibility of Xilinx FPGAs, while still meeting the demands of the PCI Express protocol. This guide describes these solutions to help designers understand which solutions are available for each FPGA. Also, this guide provides quick references to relevant product documentation and associated application notes. Additionally, known issue information specific to the Virtex-5 Integrated Block for PCI Express is included to help customers quickly determine whether the issue will impact their design.

Xilinx Solutions

The following table lists the [Xilinx solutions](#) for PCI Express along with the devices targeted and relevant documentation.

Table 1: Xilinx Solutions

FPGA Device	Lanes	IP Core	Documentation ²
Virtex [®] -5 LXT/SXT	1, 4, 8 ¹	Endpoint Block Plus Wrapper for PCI Express (Uses Virtex-5 FPGA Integrated Block)	DS551 , UG341 , GSG343 , UG197
		Soft-IP Implementation for PCI Express	DS506 , UG185 , GSG430
Virtex-4 FX	1, 4, 8	Soft-IP Implementation for PCI Express	DS506 , UG185 , GSG430
Virtex-II Pro	1, 4	Soft-IP Implementation for PCI Express	DS506 , UG185 , GSG430
Spartan [®] -3/-3A	1	Endpoint PIPE for PCI Express (Uses external NXP PHY)	DS321 , UG167 , GSG168

¹ The Block Plus wrapper will train down to a x2 implementation when configured as a x8 or x4 core assuming the link partner supports x2 operation. It cannot be configured as x2 initially. For more information regarding the Block Plus Wrapper, see the “Virtex-5 Integrated Block Known Restrictions Matrix” section.

² See “Documentation for PCI Express Solutions” for description of document types.

The Xilinx solutions for PCI Express allow customers to create PCI Express interfaces without having to have detailed understanding of the PCI Express architecture. The Xilinx

IP abstracts the PCI Express functionality to a simple LocalLink or handshaking interface allowing customers to transmit to and receive data from the PCI Express system.

The cores are architected based on the PCI Express specification and are divided into four major components. There are some subtle differences based on the solution but overall the function of each of these layers is the same. The components are:

- Transaction Layer
- Data Link Layer
- Physical Layer
- Configuration Management

The functions of the protocol layers, as defined by the PCI Express Base Specification, include generation and processing of Transaction Layer Packets (TLPs), flow control management, initialization and power management, data protection, error checking and retry, physical link interface initialization, maintenance and status tracking, serialization, deserialization and other circuitry for interface operation. Each layer is defined below.

Transaction Layer

The Transaction Layer is the upper layer of the PCI Express architecture, and its primary function is to accept, buffer, and disseminate Transaction Layer packets or TLPs. TLPs communicate information through the use of memory, IO, configuration, and message transactions. To maximize the efficiency of communication between devices, the Transaction Layer enforces PCI-compliant Transaction ordering rules, manages TLP buffer space via credit-based flow control and offers optional support for data poisoning.

Data Link Layer

The Data Link Layer acts as an intermediate stage between the Transaction Layer and the Physical Layer. Its primary responsibility is to provide a reliable mechanism for the exchange of TLPs between two components on a link. Services provided by the Data Link Layer include data exchange (TLPs), error detection and recovery, initialization services and the generation and consumption of Data Link Layer Packets (DLLPs). DLLPs are used to transfer information between Data Link Layers of two directly connected components on the link. DLLPs convey information.

Physical Layer

The Physical Layer interfaces the Data Link Layer with signaling technology for link data interchange and is subdivided into the Logical sub-block and the Electrical sub-block. The Logical sub-block is responsible for framing and de-framing of TLPs and DLLPs. It also implements the Link Training and Status State machine (LTSSM) which handles link initialization, training, and maintenance. Scrambling, de-scrambling and 8b/10b encoding and decoding of data is also performed in this sub-block. The Electrical sub-block defines the input and output buffer characteristics that interfaces the device to the PCIe link.

For the Endpoint PIPE solution, the Physical Layer exchanges information between the Data Link Layer and the external PHY component. Information received from the Data Link Layer is converted to an appropriate format and transmitted to the external PHY across the PXPIPE interface. In the same way, incoming data is received from the external PHY through the PXPIPE, and transferred onto the Data Link Layer. For the cores using the RocketIO Transceivers, the core's Physical layer instantiates and interfaces to the FPGA's transceivers.

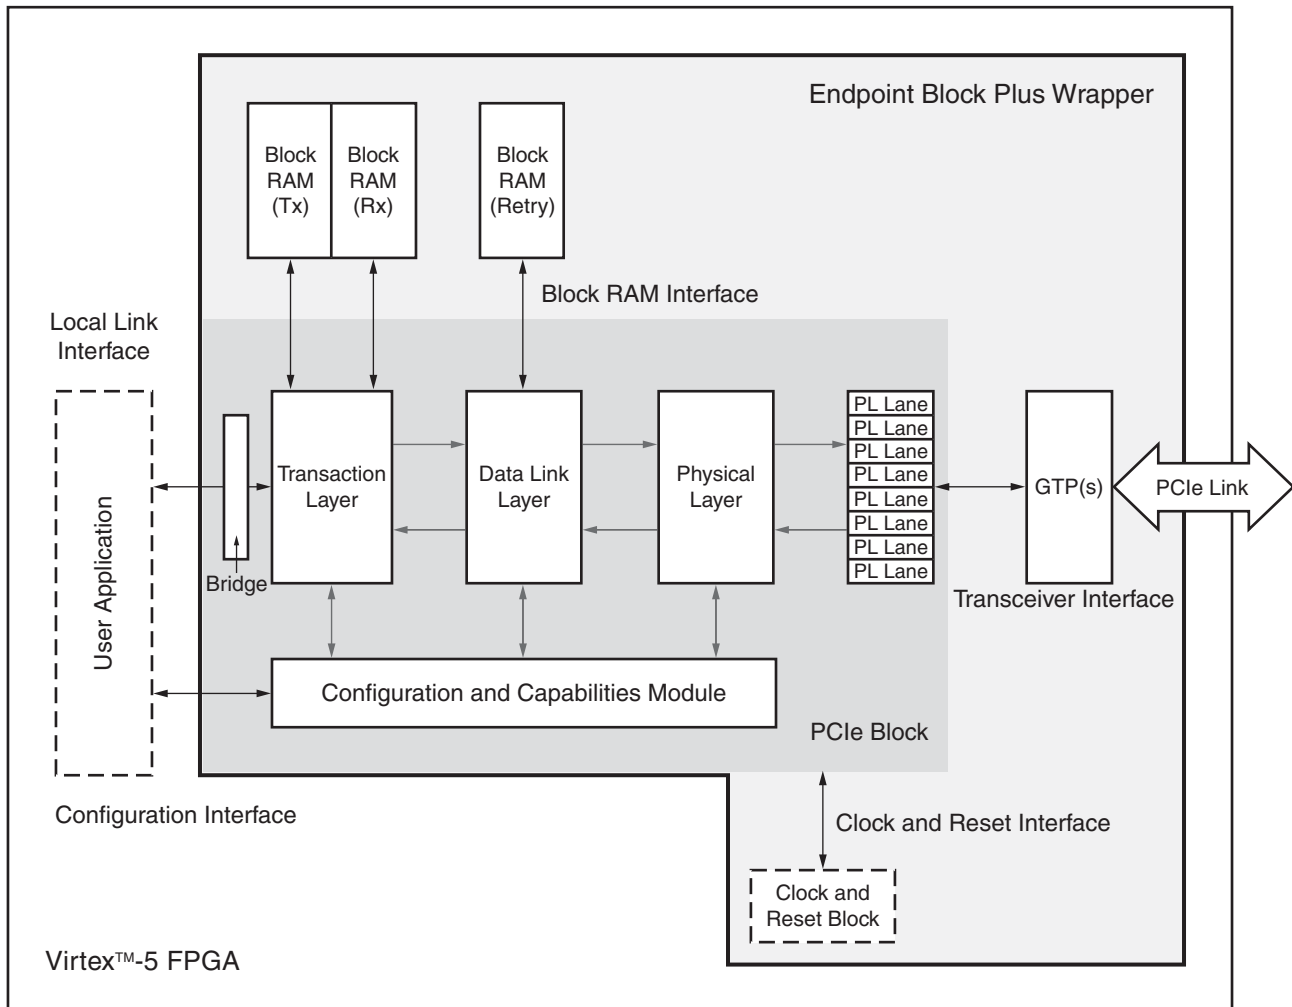
Configuration Management

The Configuration Management layer maintains the PCI Type0 Endpoint configuration space and supports the following features:

- Implements PCI configuration Space
- Supports configuration space accesses
- Power management functions
- Implements error reporting and status functionality
- Implements packet processing functions
- Receive
- Configuration reads and writes
- Transmit
- Completions with or without data
- TLM error messaging
- User error messaging
- Power management messaging/handshake
- Implements MSI and INTx interrupt emulation
- Implements the device

Endpoint Block Plus Wrapper for PCI Express

The Endpoint Block Plus Wrapper for PCI Express uses the Integrated Block for PCI Express to create a PCI Express Endpoint in the Virtex-5 FPGA. All users of the Virtex-5 Integrated Block for PCI Express should use the Block Plus Wrapper in their designs. The Virtex-5 Integrated Block implements the majority of the PCI Express architecture. The Block Plus core wraps around the block, providing a simple local link type interface to the block (similar to the Soft-IP core implementations). The Block Plus wrapper also performs many necessary functions and work-arounds to ease integration of the block into designs. See the “[Virtex-5 Integrated Block Known Issues](#)” section for more information on the block's known restrictions. [Figure 1](#) shows the basic components of the Endpoint Block Plus Wrapper.



UG493_01_022508

Figure 1: Endpoint Block Plus Wrapper for PCI Express

Components of a Xilinx Design for PCI Express

A complete design for PCI Express requires more than just the core itself. This section describes which components are needed for a fully functioning design. Normally, a typical design for PCI Express consists of the following main components:

- Board Design
- Hardware HDL Design
- Driver Design
- Software Application

Board Design

The board can be an add-in card for plugging into a standard PCI Express slot on a commercial-type motherboard. Alternatively, it can be a fully embedded system in which the FPGA is on the same board as the link partner device.

Hardware HDL Design

The hardware design refers to the Verilog or VHDL application residing on the Xilinx FPGA. The most common design used with the Xilinx core is a bus master DMA design. This design contains control engines for the receive and transmit data path along with various registers and memory interfaces to store and retrieve data. The key to the bus master DMA design is that the endpoint device is responsible for moving data to (memory write operations) and from (memory read operations) system memory.

Driver Design

The driver design, normally written in C, is the link between the higher level software application and the hardware application. The driver contains various routines that are called by the software application used to communicate with the hardware via the PCI Express link. The driver resides in the kernel memory on the system. Tools are available to help develop drivers. Microsoft provides the [Windows Driver Development Kit](#) or DDK to assist in developing drivers for Windows. [Jungo](#) also provides a driver development kit called WinDriver that supports development on both Windows and Linux systems.

Software Application

The software application is most apparent to the user. The application can be written in any programming language and can be as simple as a small C program or as complex as a GUI-type application. The user interfaces with the software application, which will then invoke routines in the driver to perform the necessary data movements. Once completed, the hardware issues an interrupt informing the driver that the data movement is finished. The driver can invoke routines in the software application to inform the user that the request has completed.

Documentation for PCI Express Solutions

The Data Sheet, User Guide, and Getting Started Guide are the main documents associated with each solution, as shown in [Table 2](#). Along with these documents are other supporting documents, such as application notes, release notes, and answer records. This section describes the documents available for each core and provides its location.

Table 2: Documentation for PCI Express Solutions

Designation	Type and Description
DS	Data Sheet: provides a high-level description of the core, key features, and general information on how to interface to the core. It includes information on which ISE [®] software is supported by the current core version, as well as which Xilinx FPGA the core targets.
UG	User Guide: provides detailed descriptions of the interface and how to use the product. The User Guide contains waveforms to show interactions with the core and other important information needed to design with the product.
GSG	Getting Started Guide: describes how to generate the core using CORE Generator [™] , as well as how to simulate and implement the core in the Xilinx ISE software.

Endpoint Block Plus Wrapper for PCI Express

The Endpoint Block Plus Wrapper for PCI Express implements a x1, x4, or x8 PCI Express solution using the Virtex-5 Integrated Block for PCI Express and RocketIO[™] transceivers. The Block Plus wrapper makes all of the necessary connections for these components, and implements needed features, such as the BRAM for the Integrated Block, reset sequencing, clocking structure, and other modules to create a complete endpoint solution.

- Core Data Sheet: [DS551](#)
- Core Getting Started Guide: [GSG343](#)
- Core User Guide: [UG341](#)
- Virtex-5 Integrated Block User Guide: [UG197](#)

Users should primarily refer to UG341 when designing with the Endpoint Block Plus Wrapper. UG341 describes the interface to the wrapper and other needed information for designs. UG197 is the User Guide for the actual Integrated Block for PCI Express contained within the Endpoint Block Plus wrapper. UG197 should be used as a companion to UG341 and is mostly needed for debug purposes and referring to the Integrated Block's known restrictions. For more information on Known Restrictions when using the Integrated Block, please see "[Virtex-5 Integrated Block Known Restrictions Matrix](#)". Endpoint Block Plus Wrapper for PCI Express documentation is located at:

http://www.xilinx.com/support/documentation/ipbusinterface-i-o_pci-express_v5pciexpressblockplus.htm

Soft IP Implementation for PCI Express

The Endpoint for PCI Express Soft-IP implementation creates a fully functional endpoint design for use in the Virtex-5, Virtex-4, and Virtex-II Pro FPGA families. The Soft-IP design is purely fabric-based and does not use the Integrated Block for PCI Express when targeted

to a Virtex-5 device. The core supports x1, x4, and x8 operation using the device's RocketIO transceivers.

- Core Data Sheet: [DS506](#)
- Core Getting Started Guide: [GSG430](#)
- Core User Guide: [UG185](#)

All Endpoint Block Plus Wrapper for PCI Express documentation is located at:

http://www.xilinx.com/support/documentation/ipbusinterfacei-o_pci-express_do-di-pciexp.htm

Endpoint PIPE Implementation for PCI Express

The Endpoint PIPE Core for PCI Express provides a low-cost solution for use in the Spartan-3/-3E/-3A FPGA family. This core is used in conjunction with the external NXP PHY.

- Core Data Sheet: [DS321](#)
- Core Getting Started Guide: [GSG168](#)
- Core User Guide: [UG167](#)

All Endpoint PIPE for PCI Express documentation is located at:

http://www.xilinx.com/support/documentation/ipbusinterfacei-o_pci_do-di-pcie-pipe.htm

For information on the NXP External PHY, visit the NXP Web site at:

<http://www.nxp.com/>

Click on "All Datasheets" and look for PX1011A_PX1012A_2 (near the lower half of the page).

Release Notes and Known Issues

Known issues for all cores, including the cores for PCI Express, are described in the [IP Release Notes Guide](#).

Application Notes and Reference Designs

XAPP859

[XAPP859](#) provides a Bus Master DMA design that transfers data to and from the ML555 on-board DDR2 memory. The characteristics of this design are as follows:

- contains all the necessary components to interface with the Endpoint Block Plus Core to provide reliable data transfer
- is a good starting point for customers who are looking for a Bus Master DMA application
- includes a Jungo driver for Windows-XP along with a Windows-XP GUI application to exercise the design

XAPP1002

[XAPP1002](#) provides information on how to use ChipScope™ to debug common problems encountered when designing with the cores. This application note provides scripts that insert ChipScope ILA cores into the design for the user at key points to allow capturing of link traffic for analysis.

XAPP1022

[XAPP1022](#) provides an example memory driver application (only binary is available) that can be used to exercise the PIO example user application that comes with the core. This application note describes how to install the driver and use it to read and write data to the PIO example user application.

XAPP1052

[XAPP1052](#) provides information on the Xilinx Bus Master DMA Performance Demonstration. This application note includes a Verilog HDL Bus Master Performance design, Windows XP driver source code, and a GUI application to exercise the design. This design is helpful for users who want to gauge performance of their system. It also provides driver source so users can familiarize themselves with a real Window's driver for PCI Express.

More Reference Designs

Other reference designs can be found [here](#).

Other Reference Material

There are various other resources available beyond the Xilinx documentation that are useful when designing for PCI Express. One popular reference book is PCI Express System Architecture published by Mindshare, Inc. Also, the [PCI Special Interest Group](#) (PCISIG) Web site contains material available to registered members of the PCI Special Interest Group.

PCI Express Base Specification

The specification provides designers with information for understanding how to form and receive transaction layer packets. Reading the specification also helps designers understand the overall flow of traffic and other intricacies of the protocol, such as configuration, error reporting, interrupts, and so on.

PCI Specification

Although the actual protocol or method of moving data through the system is different than a traditional PCI system, the PCI Express specification is based on PCI. PCI Express is compatible with the PCI plug-and-play OS model.

PCI Express Card Electro Mechanical Specification

This document outlines requirements of add-in card design, including the mechanical and electrical requirements.

PCI Express Developer Tools

The PCISIG Web site provides a suite of tools that can be used to test boards and designs for functionality. For example, the PCIECV software package is available to run on Windows-XP. This software sends configuration transactions to the device verifying that it responds as expected. Other tools include the Clock Jitter Test software that allows users to verify that the clock provided by the connector is within specification.

FPGA Documentation

This section provides documentation that is helpful when designing with the core. All device documentation is located in the [Documentation Center](#). After the core user guide, the next most helpful document for PCI Express designers is the serial transceiver user guide for the targeted family. Other useful documents are the device data sheet (includes information on DC and switching characteristics) and the general device user guide.

Virtex-5 Documentation

Virtex-5 documentation, including the RocketIO GTP Transceiver User Guide (UG196), RocketIO GTX Transceiver User Guide (UG198), and PCB Designer's Guide (UG203), is located at:

<http://www.xilinx.com/support/documentation/virtex-5.htm>

Virtex-4 Documentation

Virtex-4 documentation, including the RocketIO Multi-Gigabit Transceiver User Guide (UG076), is located at:

<http://www.xilinx.com/support/documentation/virtex-4.htm>

Virtex-II Pro Documentation

Virtex-4 documentation, including the RocketIO Transceiver User Guide (UG024), is located at:

http://www.xilinx.com/support/documentation/virtex-ii_pro.htm

Spartan Documentation

The Endpoint PIPE Core for PCI Express targets various Spartan family devices.

- [Spartan-3A](#)
- [Spartan-3E](#)
- [Spartan-3](#)

Spartan family documentation is located in the [Documentation Center](#). Select the Device Tab, then the FPGA Device Family, and select the relevant Spartan device.

Device Package User Guide

The [Device Package User Guide](#) is a helpful resource for board design and layout. This document covers each Xilinx family and conveys important information, such as thermal information and reflow soldering process guidelines, and is compliant with the specifications jitter requirements.

Customer Education Resources

Xilinx provides a two day training course for users of our solutions for PCI Express. For more information, visit the customer education Web site at:

<http://www.xilinx.com/support/training/abstracts/pci-express.htm>

License Information

There are three levels of licensing for the Endpoint for PCI Express solutions. These are:

- "Simulation Only - Available by default for all Endpoint cores generated using CORE Generator. This license allows the user to customize the core through a CORE Generator software customization GUI and generate a Unisim-based model for functional simulation.
- "Full System Hardware Evaluation - allows the user to do everything that can be done with the Fully Licensed IP core, including configure place and route, simulate, estimate timing and program a Xilinx FPGA device. The core will cease to function after 8 hours and requires the FPGA to be configured again before the core will resume operation. These license are normally valid for 3 months.
- "Full License - Allows users to generate and implement the IP in their designs. The full license does not grant access to the IP source code.

The Block Plus Endpoint Wrapper used to implement the Virtex-5 Integrated Block for PCI Express is provided free of charge to all licensed Xilinx ISE customers, there is no "Full System Hardware Evaluation" version of this core. Although, there is no cost for the Block Plus Endpoint Plus wrappers uses must still register and obtain a license to enable core generation. The license can be found by visiting the product page.

For more information on licensing please visit this page. License installation help can be found here.

Virtex-5 Integrated Block Known Issues

There are three main components to the Endpoint Block Plus Wrapper for PCI Express:

- Virtex-5 FPGA Integrated Block for PCI Express
- Virtex-5 FPGA GTP or GTX Transceivers
- Block Plus Wrapper FPGA fabric logic

Each of these components may have its own set of known issues.

[Table 3](#), [Table 4](#), [Table 5](#), and [Table 6](#) provide a quick overview of the Virtex-5 Integrated Block and GTP/GTX Known Restrictions when designing for PCI Express. Detailed information of the Integrated Blocks Known Restrictions can be found in the Known Restriction section of UG197. For Block Plus Wrapper known issues, please refer to the wrappers current release notes as pointed to by the [IP Release Notes Guide](#).

Issues listed in [Table 3](#) and [Table 5](#), are considered low impact and not a concern for the majority of applications. The "Comments" section for each issues describes where the problem may occur so designers can decide quickly if further investigation is needed by looking more closely at the description. [Table 4](#) and [Table 5](#) list the Known Restrictions that have been fixed in the Block Plus wrapper for PCI Express.

Virtex-5 Integrated Block Known Restrictions Matrix

Table 3 provides a quick overview of the Virtex-5 Integrated Block Known Restrictions.

Table 3: Xilinx Virtex-5 FPGA Integrated Block Known Restrictions Not Fixed in the Block Plus Wrapper

Issue Name (UG197)	Area of Impact	Comments
Continuous Deassertion of LLKTXCONFIGREADY	Error Handling	This issue is exposed when receiving an undefined message with data with length greater than two or receiving a malformed configuration request with formal 2'b11. By definition of the specification receiving such illegal or undefined TLPs (see description name or UG197) should result in a Fatal error message that in most cases would halt the system. The PCIe block does not produce a fatal error message and the block will hang and cause the system to halt. In a normally working system, the block should not receive such TLPs. The occurrence of this condition is rare.
Transmitting Completion TLP with Completer Abort Status	Error Handling	This exposure exists if the application requires the transmission of a completion with the Completer Abort bit set. In this case, the PCIe block will incorrectly send a Non-Fatal Error message; in some cases, this can cause the system to halt. Applications can avoid this issue by not generating a Completion packet with the Completer Abort (CA) bit set and should instead return a completion with the status of Unsupported Request (UR).
Link Retrain Due to an Absence of Update FC DLLPs	Link Training	This issue may occur if the link partner advertises infinite header and data credits on all three packet types: Posted, Non-Posted, and Completions. Advertising initial infinite credits on all three packet types can cause the link to retrain due to an absence of Flow Control DLLP updates. No known silicon advertises infinite credits for Posted, Non-Posted and Completion types.
Automatic Transmission of PME_TO_Ack Message	Power Management	The PCI Express block sends the PME_to_ACK Message TLP automatically instead of under application control. This limits the amount of time for internal house keeping before the device is powered off. Most applications do not require any or much house keeping. For applications that do require some housekeeping, prior to power down, they can consider alternative message mechanisms in higher layer software.
64 Packet Threshold on Posted Packets Passing Non-Posted and Completion Packets in TX direction	TLP Traffic	In a normal PCI Express End Point application, it is rare that a Non-Posted or Completion packet would be passed by 64 Posted transactions, because Non-Posted or Completion transactions are usually processed by the same element as the Posted transactions (usually a memory controller). If this condition does occur, it can have multiple types of impact (refer to UG197).
REPLAY_NUM Rollover in LTSSM State TX.L0s	Power Management	Active State Power Management (ASPM) is OFF by default on power-on. If the BIOS or system does not turn on ASPM in the end-point device, the system will not be affected.
ACK Ignored When Followed by IDLE Ordered Set	Power Management	Active State Power Management (ASPM) is off by default on power-on. If the BIOS or system does not turn on ASPM in the downstream peer device, the system will not be affected. If ASPM is required, the ACK to the device gets dropped resulting in a TLP replay (due to ACK timeout), which can result in a performance impact and, in worst cases, a link retrain. When the integrated block replays the packet, it should result in another ACK. If the ACK is lost repeatedly, the link will go into RECOVERY and then return to the L0 state.
Access to Unimplemented Configuration Space	Driver/Software	System software should not access unimplemented configuration space. Compliant system software accesses the configuration space by walking the capability list.

Table 3: Xilinx Virtex-5 FPGA Integrated Block Known Restrictions Not Fixed in the Block Plus Wrapper

Issue Name (UG197)	Area of Impact	Comments
Receive TLPs with Illegal Payload Length	Handling RX malformed TLPs	If the downstream port sends a malformed TLP (that does not have a payload that is DWORD aligned), then the block will not generate a Fatal Error as required. If TLPs with Illegal Payload lengths are being sent, this is typically an indication of poor link stability or a bad connected device (that is sending malformed TLPs). TLPs with Illegal Payload lengths will be NACKed and replayed by the link partner. If the packet continues to be flawed, multiple NACKs will cause the link to go back into link training (which may correct the problem). However, this condition should not occur in a normally operating system and is an indication of larger system problems.
Receiving PM_PME_TO_Ack Messages	Power Management/Error Handling	PME_TO_ACK (ACK to a Power Management Event turn off request) messages will be ignored and the link will continue to operate as normal. It is non-compliant for the link partner to be transmitting these messages to an Endpoint so the occurrence of this is rare.
Loopback Slave Mode Considerations	Loopback Testing	User applications should be aware that when the connected component initiates an exit from Loopback that not all data will be returned. Since Loopback is only used for test and verification and is not used during normal link operation, this should have little impact on designs.
Returning to L1 from L0 in D3hot state	Power Management and Disabling Drivers	This problem is most notably seen when the device's driver is disabled. Upon doing so, Windows will place the device in the D3 hot state. While in D3 hot, if a TLP is received by the endpoint, it will transition back to L0 to process the TLP and then remain in L0 instead of returning to L1. This will not result in loss of data and eventually the software will transition the device back to D0 and normal link operations will return.
Receipt of Ignored Messages	Message TLPs	"Ignored Messages" used in PCI Express specification v1.0a are no longer supported. Receiving these messages is rare. The core passes these messages which can be dropped by the user application.

Table 4 lists the Known Restrictions referred to in UG197 that are fixed in the Block Plus Wrapper for PCI Express.

Table 4: Virtex-5 Integrated Block Known Restrictions Fixed in the Block Plus Wrapper

Issue Name (UG197)	Version Fixed
TX Transmission Issues Due to Lack of Data Credits	1.6.1
64-Packet Threshold for Completion Streaming on RX Interface	v1.3
Reset Considerations in LTSSM Polling State	v1.3
Invalid Cycles in LLKRXPREferredTYPE Signal	v1.3
Receipt of Unsupported Configuration Requests and Poisoned	v1.3
Receipt of back-to-back ACK DLLPs	v1.8

Virtex-5 GTP/GTX Known Restrictions for PCI Express Matrix

Table 5 describes known restrictions due to the GTP (Virtex-5 SXT, LXT Family) or GTX (Virtex-5 FXT Family) transceivers that impact PCI Express.

Table 5: Virtex-5 GTP/GTX Known Restrictions for PCI Express Not Fixed in the Block Plus Wrapper

Issue Name	Area of Impact	Description	Comments	Affects GTP (LXT and SXT)	Affects GTX (FXT)
Transceiver Exit from Rx.L0s	Power Management	When the GTP transceiver exits from Rx.L0s, up to eight symbol times of valid data following the SKP symbol in the FTS-FTS-COM-SKP channel bonding sequence will be ignored by the block. This is a result of a delay in the assertion of the RXCHANISALIGNED signal from the GTP. If the packets are replayed by the upstream port and correctly ACKed by the integrated block, the link will stay up. However, there is a chance that if the original missed data was an ACK, that the upstream component sends the ACK, goes back to L0s, wakes up to send the ACK, and then returns to L0s. If this repeats continuously, this eventually triggers a link retrain. However, no data is lost and the link eventually recovers, causing minimal to no impact on safe operation.	This issue would only arise if Active State Power Management is in use. ASPM is OFF by default on power-on. If the BIOS or system does not turn on ASPM in the end-point device, the system will not be affected.	Yes	Yes
TX Path Entering L0s and Packet Received	Power Management	When the TX path of core is entering L0s and a packet is transmitted by the link partner, the transition back to L0 to send an ACK is very long causing the link partner to replay the packet. It results in two ACKs sent by the core (one for the original packet and one for the replayed packet). This is due to the time it takes for the GTP to return to L0 from L0s. This action is still compliant to the specification as two ACKs for the same packet can be legally sent. This could cause minor impacts on performance.	This issue would only arise if Active State Power Management is in use. ASPM is OFF by default on power-on. If the BIOS or system does not turn on ASPM in the end-point device, the system will not be affected.	Yes	Yes

Table 5: Virtex-5 GTP/GTX Known Restrictions for PCI Express Not Fixed in the Block Plus Wrapper

Issue Name	Area of Impact	Description	Comments	Affects GTP (LXT and SXT)	Affects GTX (FXT)
Reporting 8B/10B Error	Link Traffic	The GTP transceiver is required to signal an 8B/10B error by setting RXSTATUS[2:0] to 3'b100, RXDATA[7:0] to 8'hFE and RXCHARISK to 1. The transceiver correctly sets RXSTATUS, but does not indicate the correct values on RXDATA and RXCHARISK. As a result, an 8B/10B error manifests itself as an LCRC error and the integrated block transmits a NAK DLLP. This does not cause any fatal errors, but results in a non-compliant response. The link partner should re-send the packet in response to the NAK. If the 8b/10b error continues to occur, the link will retrain.	In most functioning systems 8b/10b errors are rare. Although they may occur, since they manifest to the Integrated Block as an LCRC failure, it will NAK the packet which results in the link partner resending the packet. This may have slight impacts on performance, but since the occurrence is rare, there have been no reported concerns due to this issue.	Yes	No

Table 5: Virtex-5 GTP/GTX Known Restrictions for PCI Express Not Fixed in the Block Plus Wrapper

Issue Name	Area of Impact	Description	Comments	Affects GTP (LXT and SXT)	Affects GTX (FXT)
Simulating Electrical Idle	Simulation	During simulation, it is necessary at certain times for each link partner to simulate driving electrical idle. The Xilinx GTP transceivers expect to receive either a 1, 0, or X on the RXp and RXn lines, and not a High-Z to represent electrical idle. If the link partner model transmits an High-Z to represent electrical idle, the design will fail to link up.	This issue only affects simulation and can be worked around. To work around this problem, you can either set the link partner model to drive a 1, 0, or X for electrical idle, or insert some logic in the testbench to "trap" the High-Z and replace it with an 1, 0, or X. This is similar to what is performed in (Xilinx Answer 29294) for a similar problem when interfacing to the Xilinx downstream port model for simulation.	Yes	Yes
Simulating GTP Reset	Simulation	In simulation when the GTP is in reset, the TX outputs are held to the last logic value driven instead of driving both TXp and TXn to either X, 1, or 0. This might cause the link partner model to incorrectly interpret the behavior and see it as potential START of packet.	This has been reported using the Denali BFM and it causes the Denali model to report errors as it interprets these values as potential START of packet. To work around this issue, turn off the Denali error reporting when GTPGTX is in reset. Although this has been reported only by customers using the Denali BFM, it could affect other BFMs as well.	Yes	Yes

Table 6 lists Known Restrictions due to the GTP/GTX transceivers that are fixed in the Block Plus Wrapper for PCI Express.

Table 6: Virtex-5 GTP/GTX Known Restrictions Fixed in the Block Plus Wrapper

Issue Name	Description	Version Fixed
Channel Bonding on Exit from L0s	When the link transitions from L0s to L0, there is an issue with channel bonding on a multi-lane link. A work-around for this issue is already included in the Block Plus Wrapper for PCI Express. During implementation, MAP might report that latches are being used on a path similar to: "ep/BU2/U0/pcie_ep0/pcie_blk/SIO/.pcie_gt_wrapper_i/icdrreset <7:0>". These latches are expected and are necessary for proper operation of the core.	v1.3
Electrical Idle response to TXELECIDLE in the absence of clocks	TXELECIDLE assertion will transition the TXP and TXN outputs to Electrical Idle, only if USRCLK is active. Whenever the GTP/GTX is reset the PLLs are reset causing the clock to disappear. Hence the outputs TXN and TXP rely on AC Caps to slowly discharge to their common mode voltage. During this discharge, the link partner might not recognize TXN and TXP to be in electrical idle state and can cause potential link training issues. The work around for this is to use the GTP/GTX TXPOWERDOWN signals, which do not require clocks, to drive the TXN and TXP to common mode voltage. This workaround is implemented in the Block Plus wrapper.	v1.3

Conclusion

Xilinx provides various solutions for PCI Express to meet demands of customer designs. These include the Block Plus Wrapper for PCI Express, Soft-IP Endpoint Implementation for PCI Express, and the Endpoint PIPE for PCI Express. This guide helps customers understand the devices targeted by each solution and the associated documentation for each solution.