



WP464 (v1.0) June 30, 2015

PCI Express for UltraScale Architecture-Based Devices

By: Jason Lawley

From simple register access to moving hundreds of gigabits of data, the latest integrated block for PCI Express® in the UltraScale™ architecture enables diverse connectivity for next-generation systems.

ABSTRACT

Since the introduction of the PCI Express protocol, Xilinx has been the market leader in FPGA-based PCI Express solutions—from the soft IP logic-based FPGA solutions in the Virtex®-II Pro family, to the first integrated block for PCI Express in the Virtex-5 FPGA family, and its continued use in Virtex-6, Spartan®-6, and Xilinx® 7 series devices.

The Xilinx UltraScale™ architecture-based devices include the latest generation integrated block for PCI Express within a Xilinx FPGA, including support for up to sixteen lanes (x16) of PCI Express at 8.0 gigatransfers per second (GT/s) and up to eight (x8) lanes of 16.0GT/s (Gen 4).

This breadth of experience has provided Xilinx with the expertise to develop the easiest-to-use, most feature-rich, highest-performance PCI Express.

Integrated Block for PCIe in the UltraScale Architecture

Since its introduction by the PCI Special Interest Group (PCI-SIG®) in 2003, PCI Express has been the de facto standard for processor communications. Xilinx was the first programmable logic company with intellectual property to support the standard then—and has continued to offer leading-edge PCIe performance and features today.

The Gen3 link speed (8.0GT/s) was introduced in November 2010 and the Gen4 link speed (expected to be 16GT/s) will again double the effective data rate of PCIe. The Xilinx UltraScale architecture supports all link speeds from Gen1 (2.5GT/s) to Gen4 (16GT/s) when it becomes available. See [Table 1](#).

Table 1: PCIe Base Specification Details

PCI Specification	Link Speed	Encoding Scheme and Added Overhead	Maximum Theoretical Bandwidth ⁽¹⁾
Gen1	2.5GT/s	8B/10B = 20%	2.0Gb/s
Gen2	5.0GT/s	8B/10B = 20%	4.0Gb/s
Gen3	8.0GT/s	128B/130B = 1.5%	7.88Gb/s
Gen4	16.0GT/s	128B/130B = 1.5%	15.76Gb/s

Notes:

1. Achievable system bandwidth is less than effective bandwidth due to packet overhead, traffic overhead, and other system inefficiencies.

UltraScale architecture-based devices are composed of three main categories:

- UltraScale FPGAs
 - 20nm devices that support up to Gen3 x8.
 - Includes Kintex UltraScale and Virtex UltraScale families.
- UltraScale+™ FPGAs
 - 16nm FinFET devices that support up to Gen3 x16 and Gen4 x8.
 - Includes Kintex UltraScale+ and Virtex UltraScale+ families.
- UltraScale+ MPSoCs
 - Also 16nm FinFET devices, but they consist of both a Programmable Subsystem (PS) and a Programmable Logic (PL) region.
 - Includes Zynq UltraScale+ MPSoCs.
 - The PL region contains the same integrated block for PCIe that is in UltraScale+ FPGAs with support for up to Gen3 x16 and Gen4 x8. The integrated block for PCIe in the PS region supports up to Gen2 x4, and it also has a built-in PCIe DMA engine that can optionally be enabled by the user.

This white paper focuses on the integrated block for PCIe in the FPGA and PL region of the MPSoC.

[Table 2](#) summarizes the level of support for each family.

Table 2: PCIe Lane Width and Speed Support

UltraScale Architecture-Based Devices ⁽¹⁾	Integrated PCIe Type	Number of Blocks	Soft PCIe IP
Kintex® UltraScale	Gen3 x8	2–6	Gen3 x8
Kintex UltraScale+	Gen3 x16 Gen4 x8	0–5	
Zynq UltraScale+⁽²⁾	Gen3 x16 (PL) Gen4 x8 (PL) Gen2 x4 (PS)	— 0-5 (PL) 1 (PS)	
Virtex UltraScale	Gen3 x8	2–6	
Virtex UltraScale+	Gen3 x16 Gen4 x8	2–6	

Notes:

1. Links are to the associated Product Tables and Product Selection Guide.
2. Note that not all Zynq UltraScale+ MPSoC devices have an integrated PCIe block in the PL. See the [Zynq UltraScale+ MPSoC Product Tables and Product Selection Guide](#) for details.

Advanced Features

The integrated block for PCIe contains advanced features like Single Root I/O Virtualization, data straddling, and fast device configuration (Tandem), which allow customers to optimize their PCI Express solutions. More about these features can be found in the [PCIe Features for the UltraScale Architecture](#).

In addition to the integrated block for PCIe, Xilinx Alliance Partners Northwest Logic and PLDA provide Gen3x8 soft IP solutions that target UltraScale architecture-based devices. For more information, including additional documentation, videos, and a list of all Xilinx devices that support PCIe, go to the [PCIe product web page](#).

The scalable, optimized architecture of the integrated blocks for PCIe, along with the AXI4 user interfaces, allow easy migration and design reuse across all UltraScale architecture-based devices, from lower cost to ultra-high-performance applications.

The integrated block for PCIe can be configured using a simple GUI-based tool flow to set configuration options such as Endpoint, Root Port, Link Width, Link Speed, Device IDs, BAR register sizes, and numerous other options. Developers can select a variety of use modes, including IP Integrator (IPI) or standard RTL delivery. IPI, which is a Vivado® Design Suite tool, can be used to easily connect the integrated block to other IP or interconnect.

After users have customized the PCIe IP through the GUI, they have the option to generate a simple example design. This example design can be created from the configured IP and can be both implemented and simulated. Development boards such as the [KCU105](#) that have a PCIe interface can be targeted during IP generation to create an example design, which can be loaded quickly into hardware and tested.

In addition to easy-to-use development and implementation tools, Xilinx provides Targeted Reference Designs—fully validated and supported application examples—that accelerate the design schedule. These Targeted Reference Designs include all components of a PCIe design, such as DMA controllers, custom IP, device drivers, and software applications.

To learn more about Targeted Reference Designs for PCIe, see the specific evaluation kit instructions at:

<http://www.xilinx.com/products/boards-and-kits.html>

UltraScale Architecture Advantages

The UltraScale architecture has many features that make implementing a high-performance PCIe design possible. The items described in this section enable the PCIe design to operate at peak capacity and simplifies the design process.

Data Throughput and Performance

The PCI-SIG sets a goal of doubling the effective data throughput for each new generation of PCIe, and Gen3 and Gen4 are no exception. Effective data throughput (sometimes referred to as effective data transfer rate) is not the same thing as raw data transfer speed, such as 8GT/s or 16GT/s link speed. The effective data throughput rate is dependent on many variables, such as:

- Lane width
- Link speed
- System Maximum Payload Size and Maximum Read Request Size
- Encoding loss
- DMA scatter-gather overhead

For a description of the possible variables for effective data throughput and performance, see [WP350, Understanding Performance of PCI Express Systems](#).

The integrated block for PCIe in UltraScale+ devices is capable of sustained throughput of over 14GB/s per direction, when configured to operate as Gen3 x16 or Gen4 x8 in a system with a 256 byte system Maximum Payload Size. These data rates are achieved with an internal test design created specifically as a throughput test application that shows the maximum possible throughput for a PCIe core on a given system.

When a more realistic, real world scatter-gather DMA is used instead, the effective data throughput will decline slightly. A very reasonable effective data throughput that users can expect with a scatter-gather DMA is around 13GB/s per direction—but will vary based on the factors previously outlined above. A video demonstration of this performance can be found on the [PCI Express](#) landing page.

The abundance of high-performance, low latency interconnect in the UltraScale architecture allows designers to create the wide high-performance data buses that are necessary to handle 28Gb/s of full bandwidth data.

Transceiver Advantages

The transceivers in the UltraScale architecture-based devices contain features that allow for very robust operation at PCIe data rates. These features include:

- Transmitter emphasis/equalization
- Auto-adaptive Equalization

The transmitter emphasis circuit is designed to overcome high-frequency channel insertion loss and is implemented as a 3-tap FIR filter. The 3 taps consist of a pre-, main-, and post-cursor taps. These taps are programmable and can support all of the various PCIe preset settings as well as link-partner specified tap settings. Typically, the user does not need to explicitly set the tap values because these values are set automatically by the PCIe Link Equalization protocol. The Continuous Time Linear Equalizer (CTLE) and Decision Feedback Equalizer (DFE) circuits in the GTH and GTY transceivers work together to compensate for up to 25dB of loss. The CTLE and DFE employ a fully auto-adaptive algorithm that continuously monitors the incoming signal, and optimally adjusts the frequency response of the high-pass filter function. This auto-adaptive feature lessens the burden on the user and solves the issue of over-equalization or under-equalization.

The DFE taps compensate for reflections and higher loss channels. This compensation is extremely useful when PCIe is used over a backplane, as commonly found in many wired communication and data center applications. For a detailed description on some of the advanced equalization features offered by the UltraScale architecture-based transceivers, see [WP458, Leveraging UltraScale FPGA Transceivers for High-Speed Serial I/O Connectivity](#).

Memory Bandwidth

Most PCIe applications use some type of memory for data buffering, typically DDR SDRAM.

The 13GB/s throughput example given for the scatter-gather DMA is actually a good fit for Xilinx DDR4 memory solutions. See [Figure 1](#).




Figure 1: Interfacing DDR4 Memory to UltraScale Architecture-based Devices

When determining memory bandwidth requirements, the designer should use a 2.5X bandwidth multiplier factor to account for both read and write directions and any additional overhead such as memory addressing.

For example, if transferring sustained 13GB/s from the PCIe link and all of this data is buffered in DDR4 memory, the designer can calculate the following to determine the memory bandwidth and/or interface width requirements.

Determining Memory Bandwidth Requirements

Total Memory Bandwidth Required for Sustained Transfers:

$$13\text{GB/s} * 2.5 = 32.5\text{GB/s}$$

Example: If using an 2133Mb/s DDR4-capable memory, the designer can calculate how wide the data interface has to be to keep up with 32.5GB/s.

Convert to Gb/s:

$$32.5\text{GB/s} * 8 \text{ bits/byte} = 260\text{Gb/s}$$

Calculate Required Interface Width for DDR4 Memory:

$$260\text{Gb/s} / 2133\text{Mb/s per pin} = \sim 122 \text{ pins}$$

This calculation shows that two standard 72-pin DDR4 interfaces operating at 2133Mb/s can keep up with full-duplex data from a x16 Gen3 PCIe link.

Devices that support slower DDR data rates, such as 1,600Mb/s, require additional pins and components.

For more information on UltraScale architecture-based memory solutions, see [WP454](#), *High-Performance, Lower-Power Memory Interfaces with UltraScale Architecture FPGAs*.

Scalable, Optimized AXI Interface

Xilinx's deployment of the AMBA®4 AXI4 specification allows for a consistent way to connect IP blocks while enabling better use of design resources. AXI4 allows the use and reuse of IP and enables easier integration across IP providers, all in support of Plug-and-Play FPGA design. See product guide [PG156](#), *UltraScale Architecture Gen3 Integrated Block for PCI Express* and product guide [PG194](#), *AXI Bridge for PCI Express Gen3 Subsystem*.

All PCIe solutions for UltraScale architecture-based devices are designed to the AMBA® AXI4 specification. Based on the PCIe cores used, the user has access to either the AXI4-Stream interface or the AXI4 Memory Mapped Interface.

1. **AXI4-Stream:** This interface splits/combines the data stream into Completer and Requester streams. This allows for optional features such as packet destraddling, data realignment, and completion tag management. See [Figure 2](#).




Figure 2: Enhanced AXI4-Stream Interface

2. **AXI4:** This is a memory-mapped interface for use with processor system based cores. This interface is the typical path for embedded designs. See Figure 3.




Figure 3: AXI4 Interface

PCIe Features for the UltraScale Architecture

The UltraScale architecture provides many features that give designers the ultimate in PCIe performance, flexibility, and ease of use.

Fast Initialization for the Integrated Block for PCIe

The *PCI Express Base Specification* requires the PCIe link to be ready to link train within 100ms after power is stable. This has traditionally been a challenge for large configurable devices (>100,000 logic cells) because it can take well over 100ms to configure a large device using common flash memory devices.

“Brute force” methods are traditionally used to resolve the 100ms requirement. Typically, designers use the fastest and widest flash memory devices available to achieve the necessary bandwidth to meet the configuration time requirement. Some cases require the use of multiple flash devices in conjunction with a CPLD to achieve the required bandwidth. While this can be the simplest method from a software perspective, it is often the most expensive due to increased BOM cost. This method also uses valuable I/O, especially when using wide input buses, and is quickly becoming obsolete as the size of Xilinx programmable devices has grown to two million logic cells and higher.

Beginning with the Virtex-6 family, Xilinx is the first FPGA company to provide multiple methods to meet this initialization requirement, each with different levels of complexity and expense.

Tandem and Tandem Field Update

The UltraScale architecture has two different flows that enable devices to meet the 100ms boot time requirement. These flows are named Tandem and Tandem Field Update. Both flows ensure that the PCIe interface is up and running so it can be enumerated into the system during initialization. Tandem Field Update has the added benefit that allows the device to be reconfigured over the PCIe link without bringing the PCIe link down.

Tandem

In this flow, there are two ways to initially configure the programmable device coming out of power on reset: Tandem PROM and Tandem PCIe.

Introduced in the Xilinx 7 series devices, the Tandem PROM method is the simplest and least expensive to implement. The user directs the implementation tools to create a two-stage bitstream via a simple software switch when building the PCIe core. The first stage of the bitstream contains just the configuration frames necessary to configure the integrated block for PCIe. When configured, a device STARTUP sequence occurs, and the PCIe link becomes active, thus easily satisfying the 100ms requirement. The remainder of the device configuration is then loaded while the PCIe enumeration/configuration system process is occurring. The two-stage bitstream method can use an inexpensive flash device to hold the bitstream(s). See [Figure 4](#).




Figure 4: Tandem PROM Method

The Tandem PCIe solution builds off the Tandem PROM technology and allows the user to load the second-stage bitstream via the PCIe link.

Tandem Field Update

Just like the Tandem method, Tandem Field Update allows the user to initially configure the device via Tandem PROM or Tandem PCIe. After the device is configured, the user can choose to download new device functionality over the PCIe link. The user can load as many designs over the PCIe link as needed. This is ideal for systems/designs that need to be updated in field. It is also useful for designers who are in the lab debugging and do not want to continue to reboot a PC every time a new device image needs to be loaded. See [Figure 5](#).

Tandem Field Update allows designers to select between either Tandem PROM or Tandem PCIe as an initial load. After the initial load is completed, designers can load any new logic over the PCIe bus. More information about this feature can be found in product guide [PG156](#), *UltraScale Architecture Gen3 Integrated Block for PCI Express v3.1*.

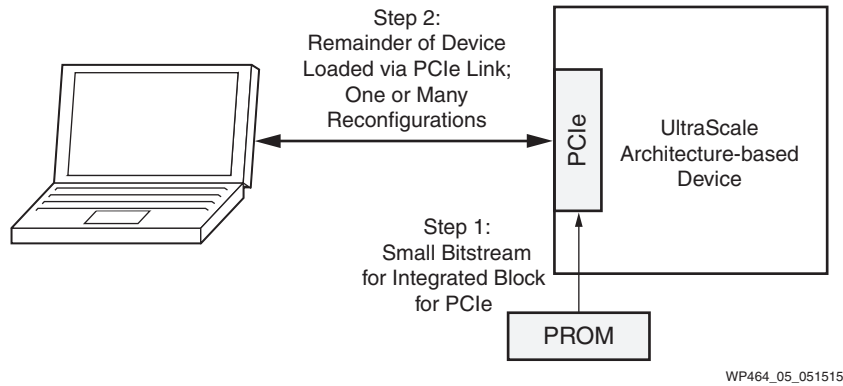


Figure 5: Tandem Field Update Tool Flow

Data Straddling for Performance

The integrated block for PCIe in UltraScale architecture-based devices has the highest throughput performance for any programmable-device-based PCIe solution on the market. Most of these solutions require the TLPs on the user interface to be received in an aligned manner, that is, only one packet can be in the data interface when the TLP ends. The next TLP cannot then be read from the core until the next clock cycle.

As the data rate continues to increase, so does the internal datapath. Gen4 x8 and Gen3 x16 designs will require a 512-bit datapath, making it imperative to limit wasteful cycles of data by allowing packets to be straddled.

Solutions without the ability to straddle data introduce gaps within the data stream, which in turn reduces overall data throughput. UltraScale devices have the ability to straddle packets (allow one TLP to end while another begins on the same clock cycle) on the user interface, thereby allowing the PCIe core to run at the full line rate. This is important for ultra high-end applications that require full line rate bandwidth. For applications that do not require extreme bandwidth and prefer aligned packets, the enhanced AXI-Stream interface has an optional alignment feature. See [Figure 6](#).

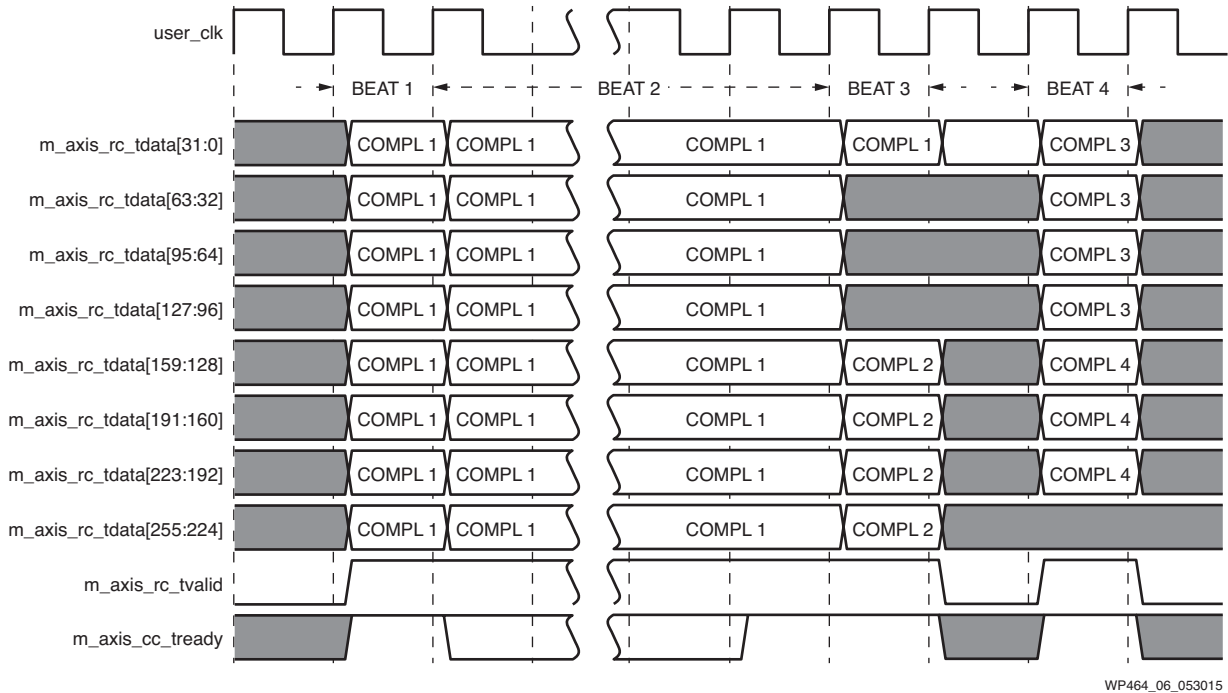


Figure 6: UltraScale Architecture Straddled Cycle

In addition to support for straddled packets, the UltraScale architecture-based devices also have features to improve overall performance such as improved user control for credit allocation schemes as well as new flow control capabilities that give the user more granular control over Posted and Non-Posted traffic.

Tag Management for Read Requests

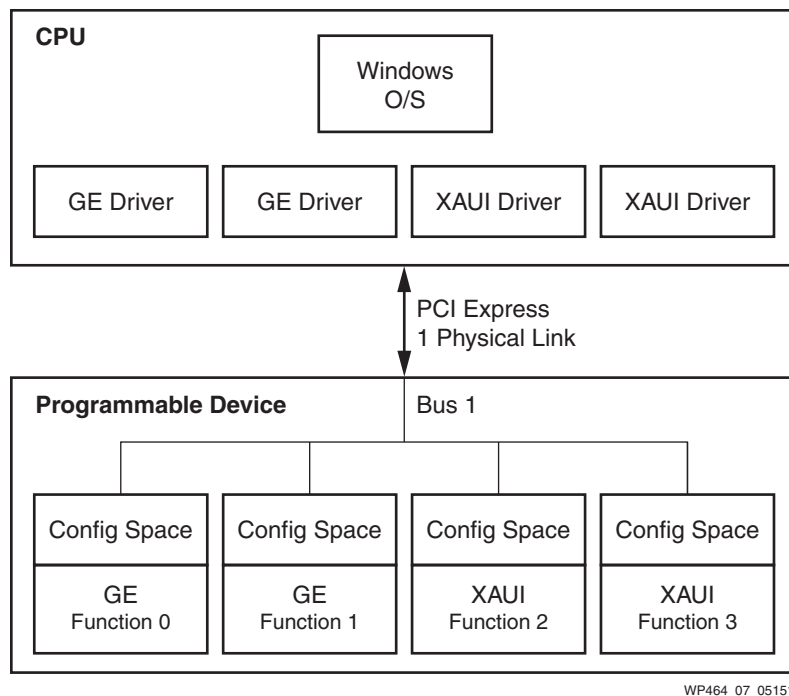
One of the difficult tasks that a designer must undertake when transmitting Read Request TLPs that are larger than a typical system Read Completion Boundary size of 64 bytes is the handling of multiple completions and completions that are returned out-of-order. Typically, the designer must store the tags for outgoing read requests, and then reconcile and manage those tags with the incoming completion TLPs. In addition, the designer must also monitor for error conditions, such as completion time-outs.

Tag management is a necessary feature for Bus-Mastering DMA designs that send Read Requests, or in other words, “pull” data from a producer. This is done by managing the tags for outgoing read-requests and reconciling the incoming completions to these tags. The PCIe solution for UltraScale devices optionally provides this tag management feature, greatly simplifying the design requirements for DMA designers.

Multiple Functions

The PCIe solution has the ability to operate as a multifunction device. This type of device has several functions all sharing a single PCIe link. Each function has its own PCIe Configuration Header space; thus, from a host-system software perspective, each function appears as an individual PCIe device on its own PCIe link. This greatly simplifies device driver development and portability

because the driver developer can create a single driver and replicate it for each hardware function. See [Figure 7](#).



WP464_07_051515

Figure 7: Multifunction Devices

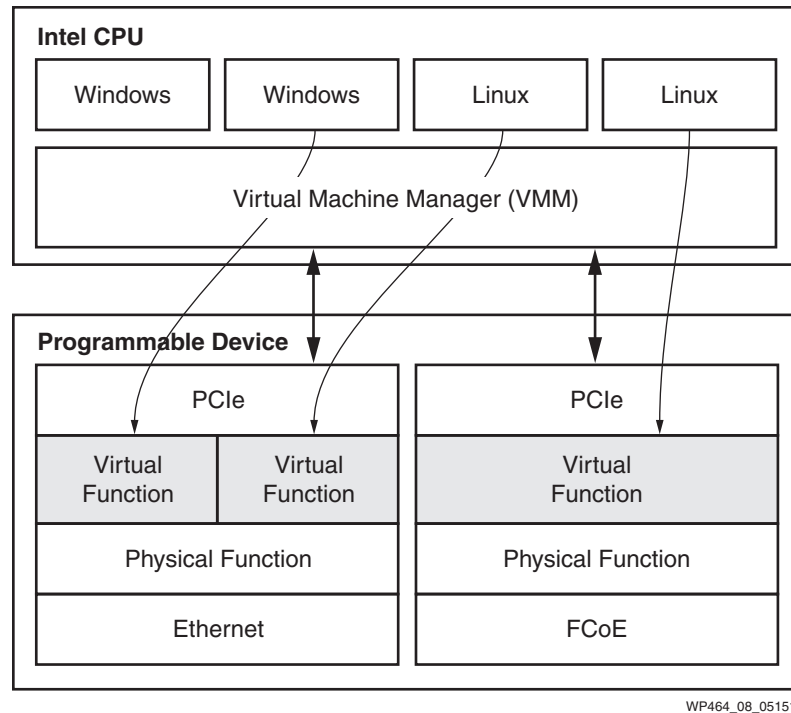
UltraScale and UltraScale+ devices have 2 and 4 physical functions respectively, implemented completely in the integrated block for PCIe.

Single Root I/O Virtualization

The integrated block for PCIe in UltraScale devices has up to 2 physical and 6 virtual functions built in to the block itself. UltraScale+ devices extend this functionality with up to 4 physical functions and 252 virtual functions.

SR-IOV allows for multiple guests (operating systems) running on a single root (CPU subsystem) to access I/O devices without the software penalty incurred in virtualized systems that do not support SR-IOV. Similar to how multifunction devices provide an individual configuration space for each physical function, SR-IOV works by providing a virtual function (virtual configuration space) for each guest operating system accessing the I/O device. Thus, each guest operating system has its own "view" of the I/O device.

Adapters that support SR-IOV have shown vast improvements in I/O efficiency in virtualized environments. Not only has SR-IOV become a widely adopted standard within the Enterprise IT market (Data Center), but it is beginning to see inroads within the Communications and Storage Networking markets as well. See [Figure 8](#).



WP464_08_051515

Figure 8: SR-IOV Virtual Configuration Space

Built-in MSI-X Table

MSI-X interrupts have two major advantages over MSI (message signal interrupts). First, the number of interrupt vectors that can be supported increases from 32 in MSI to 2048 in MSI-X. Second, the MSI-X interrupt vector can be steered to different locations that are stored in a table.

UltraScale FPGAs implement MSI-X by having the user build and manage the MSI-X table in the programmable logic. UltraScale+ devices simplify this operation by providing the option to implement the table in the integrated block for PCIe, thus simplifying the solution for the user.

Advanced Error Reporting and End-to-End CRC

Advanced Error Reporting (AER) is an optional feature that provides more granularity and control for the types of errors that can occur in a PCIe-based system. In non-AER PCIe-based systems, only three types of errors are defined: fatal, non-fatal, and correctable. In most cases, the three defined error types do not give enough information to the system to recover gracefully from an error. With AER enabled, the system software can determine the exact cause of a particular error and attempt to recover if possible.

The integrated block for PCIe in UltraScale and UltraScale+ devices optionally performs automatic end-to-end CRC (ECRC) checking and generation, when enabled by the user. Ports are accessible to control the error generation and flags if an ECRC error is detected. The ECRC checking and generation logic no longer needs to be implemented into the user's design.

AER and ECRC are used in applications where high reliability and high availability are key driving factors. These features are commonly used in market segments like Aerospace and Defense, Banking and Finance, Communications, and Storage.

Atomic Operations

Atomic Operations introduces three new TLP types that are intended to improve system performance and latency by creating standard synchronization primitives such as mutexes and spin-locks, directly over the I/O bus, in this case, PCIe. This is helpful in any system with multiple producers and consumers, for example, a multi-CPU system. The target application space for this feature is in co-processing and hardware acceleration adapters. UltraScale and UltraScale+ devices fully support Atomic Operations.

Features that Enable High Performance PCIe Applications

The integrated block for PCIe in the UltraScale architecture contains many features that enable better system performance. See [Table 3](#).

Table 3: PCIe Features by Device

UltraScale	UltraScale+	Both
Data straddling on the 256-bit completion interface	Data straddling on all 512-bit interfaces and on the 256-bit completion interface	Four high performance AXI4-Streaming interfaces optimized for high performance designs
Built-in tag management for up to 32 outstanding read requests	Built-in tag management for up to 256 outstanding read requests	Parity protection on AXI4-Streaming interfaces
16KB completion buffer space	32KB completion buffer space with up to 256 completions	ECC protection on all internal buffer memory
Built-in multifunction and SR-IOV with 2 physical functions and 6 virtual functions	Built-in multifunction and SR-IOV with 4 physical functions and 252 virtual functions	Atomic Operation Transactions
	Built-in MSI-X table	Address Translation Services (ATS)
		TLP Processing Hints Capability (TPH)
		Relaxed ordering support on the receive path

Other Advanced Features

In addition to SR-IOV and Atomic Operations, UltraScale and UltraScale+ devices support many of the ECNs introduced in the latest version of the *PCI Express Base Specification*. Many are supported directly by the block without any user intervention:

- Extended Tag Field Enable
- Internal Error Reporting
- ASPM Optionality

For detailed information on these features, go to: [PG156](#), *UltraScale Architecture Gen3 Integrated Block for PCI Express*.

Conclusion

The integrated block for PCIe in UltraScale architecture-based devices marks the fourth generation of integrated PCIe within a Xilinx device family. Drawing on such broad experience, Xilinx has developed the easiest to use, most feature-rich, and highest performing PCIe solution for programmable devices available. The optimized architecture and scalable AXI4 interconnect allows users the ability to seamlessly reuse and migrate existing designs across the UltraScale and UltraScale+ families. Features such as PCIe Gen3 and Gen4, x16 link widths, straddled packets, and SR-IOV allow designers to achieve bandwidth and system performance never before imagined. With simple software tool flows and Targeted Reference Designs, designers can easily customize the integrated block for PCIe and accelerate time-to-market for their application.

Additional Information

[PG194](#), *AXI Bridge for PCI Express Gen3 Subsystem*

[Release Notes](#), *UltraScale FPGA Gen3 Integrated Block for PCI Express*

[Release Notes](#), *AXI Bridge for PCI Express Gen3*

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
06/30/2015	1.0	Initial Xilinx release.

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 Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; 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 such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS APPLICATIONS RELATED TO: (I) THE DEPLOYMENT OF AIRBAGS, (II) CONTROL OF A VEHICLE, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR, OR (III) USES THAT COULD LEAD TO DEATH OR PERSONAL INJURY. CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN SUCH APPLICATIONS.