

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



The 100Gbps Ethernet IP solution offers a fully integrated IEEE802.3-2015 compliant package for NIC (Network Interface Card) and Ethernet switching applications. As shown in the figure below, the 100Gbps Ethernet IP includes:

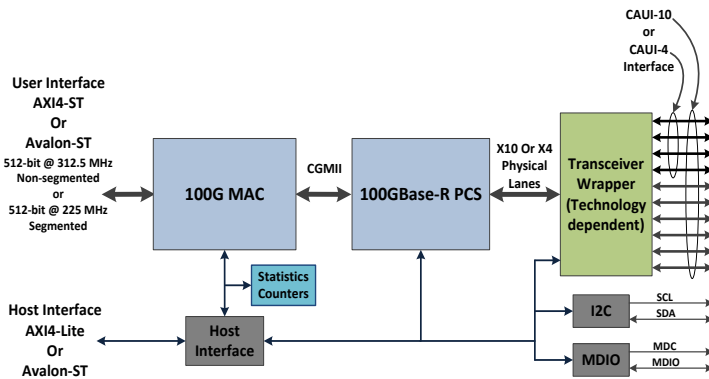
- 100Gbps MAC core with AXI-4 Streaming or Avalon Streaming user interface
- 100Gbps (100GBase-R) PCS core with support for **CAUI-4 (-C4 option)** and **CAUI-10 (-C10 option)** interfaces
- Technology dependent transceiver wrapper for Altera and/or Xilinx FPGAs
- Statistics counter block (for RMON and MIB)
- MDIO and I2C cores for optical module status and control

specifications) to enable Converged Enhanced Ethernet (CEE) applications like data center bridging that employ IEEE 802.1Qbb Priority Flow Control (PFC) to pause traffic based on the priority levels.

Features Overview

MAC Core Features

- Implements the full 802.3 specification with preamble/SFD generation, frame padding generation, CRC generation and checking on transmit and receive respectively.
- Implements 802.3bd specification with ability to generate and recognize PFC pause frames
- Implements reconciliation sublayer functionality with start and terminate control characters alignment, error control character and fault sequence insertion and detection.
- Implements a 320-bit CGMII interface operating at 312.5 MHz for 100G EMAC
- Implements Deficit Idle Count (DIC) mechanism to ensure maximum possible throughput at the transmit interface.
- Implements logic for padding of frames on the transmit path if the size of frame is less than 64 bytes.
- Implements fully automated XON and XOFF Pause Frame (802.3 Annex 31A) generation and termination providing flow control without user application intervention. Non PFC mode only.
- Pause frame generation additionally controllable by user application offering flexible traffic flow control.
- Support for VLAN tagged frames according to IEEE 802.1Q.
- Support any type of Ethernet Frames such as SNAP / LLC, Ethernet II/DIX or IP traffic.
- Discards frames with mismatching destination address on receive (Except Broadcast and Multicast frames).
- Supports programmable promiscuous mode to omit MAC destination address checking on receive EMAC.
- Optional multicast address filtering with 64-bit HASH Filtering table providing imperfect filtering to reduce load on higher layers.
- CRC-32 generation and checking at high speed using an efficient pipelined CRC calculation algorithm.
- Implements logic for optional padding removal on RX path for NIC applications or forwarding of unmodified data to the user interface.



A complete reference design using a synthesizable L2 (MAC level) packet generator/checker is also included to facilitate quick integration of the Ethernet IP in a user design. A GUI application interacts with the reference design's hardware elements through a UART interface (PCIe option is also available). A basic Linux PCIe driver/API is also provided for memory mapped read/write access to the internal registers. See **Appendix A** for details.

MAC and PCS cores are designed with 320-bit data path operating at 312.5MHz.

As the transceiver wrapper is included with the Ethernet IP solution, the line side directly connects the 10.3125Gbps (for CAUI-10 interface) or 25.78125Gbps (for CAUI-4 interface) FPGA transceivers to various optical modules including CFP, CFP2, CFP4, CXP and QSFP28.

Ethernet IP solution implements two user (application) side interfaces. The register configuration and control port is a 32-bit AXI4-Lite or Avalon-MM interface. A 512-bit non-segmented AXI-4 Streaming or Avalon Streaming bus at 312.5MHz is used to interface with the MAC block. Additionally, an interface wrapper is provided to support segmented interface operation at lower clock speeds.

100Gbps Ethernet IP supports advanced features like per-priority pause frames (compliant with 802.3bd

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



- Discards runt frames (less than 64 Byte) at the core's reconciliation sublayer.
- Implements logic for optional forwarding of the CRC field to user application interface.
- Implements logic for optional forwarding of received pause frames to the user application interface.
- Programmable frame maximum length providing support for any standard or proprietary frame length (e.g. 9K-Bytes Jumbo Frames).
- Status signals available with each Frame on the user interface providing information such as frame length, VLAN frame type indication and error information.
- Implements programmable internal CGMII Loop-back.
- Implements statistics indicators for frame traffic as well as errors (alignment, CRC, length) and pause frames.
- Implements statistics and event signals providing support for 802.3 basic and mandatory managed objects as well as IETF Management Information Database (MIB) package (RFC 2665) and Remote Network Monitoring (RMON) required in SNMP environments.
- Implements a streaming user application interface. The application interface is designed as a 512-bit non-segmented (start of a new frame on next 512-bit word) interface operating at 312.5MHz.
- An interface wrapper is provided for applications that implement a segmented (start of new frame within same 512-bit word with 64-bit alignment) bus. In segmented mode, the 512-bit bus operates at @ 225MHz for 100Gbps.
- Implements memory-mapped host controller interface for accessing the core's register file.
- Implements 66-bit block synchronization and Alignment Marker Lock machines as specified in 802.3ba specifications.
- Implements skew compensation logic in order to realign all the virtual lanes and reassemble an aggregate 100G stream (with all 64b/66b blocks in the correct order)
- Implements lane reordering to support reception of any virtual lane (VL) on any physical lane (PL).
- Implements BIP-8 insertion/checking per Virtual Lane on transmit/receive respectively.
- Implements Inter Packet Gap (IPG) Insertion/Deletion for Alignment marker and clock compensation while maintaining a minimum of 1 byte IPG.
- Implements programmable internal CGMII loop-back which directs traffic received from core's receive path back to transmit PCS.
- Implements Bit Error Rate (BER) monitor for monitoring excessive error ratio. In addition, the core implements various status and statistics required by the IEEE 802.3ba such as block synchronization status, AM lock status, lane de-skew and lane reordering status and BIP-8 error counters per virtual lane.

PCS Core Features (Common)

- Implements 100GBase-R PCS core compliant with IEEE 802.3ba Specifications.
- Implements a 320-bit CGMII interface operating at 312.5MHz for 100G Ethernet.
- Implements 64b/66b encoding/decoding for transmit and receive PCS.
- Implements 100G scrambling/descrambling using 802.3ba specified polynomial $1 + x^{39} + x^{58}$
- Implements Multi-Lane Distribution (MLD) across 20 Virtual Lanes (VLs)
- Implements periodic insertion of Alignment Marker (AM) on the transmit path and deletion on the receive path

PCS Core Features (CAUI-4 Option)

- Implements gear-box logic for Xilinx to convert 20 VLs of 66-bit blocks to 4 PLs of 160-bit data for line side CAUI-4 interface. The 160-bit interface operates at the transceiver reference clock of 161.1328125MHz.
- Implements gear-box logic for Altera to convert 20 VLs of 66-bit blocks to 4 PLs of 128-bit data for line side CAUI-4 interface. The 128-bit interface operates at the transceiver reference clock of 201.416015625MHz.
- Transceiver Wrappers for Xilinx GTY/GTZ transceivers and Altera GT transceivers.

PCS Core Features (CAUI-10 Option)

- Implements gear-box logic for Xilinx to convert 20 VLs of 66-bit blocks to 10 PLs of 40-bit data for line side CAUI-10 interface. The 40-bit interface operates at the transceiver reference clock of 257.8125MHz.
- Transceiver Wrappers for Xilinx GTX/GTH transceivers and Altera GX/GS transceivers.

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



Licensing and Maintenance

- ***NO yearly maintenance fees for upgrades and bug fixes***
- Basic core licensing for a single vendor (either Xilinx or Altera) compiled (synthesized netlist) binary
- Option for vendor and device family agnostic source code (Verilog) license

Ordering Codes

HTK-100G-ETH-320-FPGA-C4: 100 Gigabit Ethernet Solution with 4 lane CAUI-4 PCS and optical side interface

HTK-100G-ETH-320-FPGA-C10: 100 Gigabit Ethernet Solution with 10 lane CAUI-10 PCS and optical side interface

Contact and Sales Information

For further information, contact sales representative at:

Phone: +1-301-528-8074

Email: sales@hiteksys.com

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



Resource Utilization

The core utilization summary for the 100G Ethernet solution is given in following tables. The Ethernet solution has been fully verified on different hardware platforms for both Altera and Xilinx FPGAs.

100G Ethernet IP - Resource Usage for Xilinx Devices

Device	User Interface (AXI4)	Priority Flow Control (PFC)	PCS Type	Slice LUTs	Slice Registers	BRAMs
UltraScale/ UltraScale+	512-Bit (Non-Segmented)	No	CAUI-4	51,837	62,022	18K = 4; 36K = 74
		Yes	CAUI-4	52,300	62,673	18K = 4; 36K = 74
		No	CAUI-10	43,891	58,365	18K = 4; 36K = 74
		Yes	CAUI-10	44,354	59,016	18K = 4; 36K = 74
7-Series	512-Bit (Non-Segmented)	No	CAUI-4	51,899	61,984	18K = 4; 36K = 74
		Yes	CAUI-4	52,349	62,635	18K = 4; 36K = 74
		No	CAUI-10	44,443	58,344	18K = 4; 36K = 74
		Yes	CAUI-10	44,893	58,995	18K = 4; 36K = 74

Note:

- Register based RMON statistics block adds additional 1948 Slice LUTs and 1807 Slice Registers.
- CAUI-4 supported for devices with 25Gbps transceivers only.

100G Ethernet IP - Resource Usage for Altera Devices

Device	User Interface (Avalon)	Priority Flow Control (PFC)	PCS Type	COMB. ALUTs	Registers	Memory Blocks
Arria 10	512-Bit (Non-Segmented)	No	CAUI-4	56,564	60,336	M20K = 149
		yes	CAUI-4	56,966	61,009	M20K = 149
		No	CAUI-10	37,804	55,349	M20K = 149
		yes	CAUI-10	38,206	56,022	M20K = 149
Stratix V	512-Bit (Non-Segmented)	No	CAUI-4	56,583	60,234	M20K = 149
		yes	CAUI-4	56,988	60,824	M20K = 149
		No	CAUI-10	37,802	55,163	M20K = 149
		yes	CAUI-10	38,297	55,753	M20K = 149

Note:

- Register based RMON statistics block adds additional 2000 Comb. ALUTs and 1800 Registers.
- CAUI-4 supported for devices with 25Gbps transceivers only.

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



Deliverables

- Ethernet wrapper design with:
 - Top level MAC and PCS wrappers (source files, Verilog) for user specific customizations
 - Compiled synthesizable binaries (Netlists) for MAC and PCS cores
 - Compiled synthesizable binaries (Netlists) for L2 packet generator and checker
 - Technology specific transceiver wrappers for the selected device family
 - Source code RTL (Verilog) for RMON and Register-File blocks
- UART/PCIe interface based reference design with:
 - Top level wrapper (source files, Verilog) for user specific customizations
 - Compiled synthesizable binaries (Netlists) for the I2C and MDIO cores
 - Compiled synthesizable binaries (Netlists) for the UART or PCIe host interface
- Encrypted RTL for MAC, PCS and packet generator/checker for simulation
- Constraint files and synthesis scripts for design compilation
- PCIe driver/API (source files, C) for Linux with the optional PCIe host interface
- GUI application (Linux only) for interfacing to the reference design
- Design guide(s) and user manual(s)

Validated/Ported Module List

1. **HiTech Global HTG-V7-X16PCIE-580 [HTG-728]**; Virtex-7 580 FPGA with GTZ (**CAUI-4**) Transceiver, Optical module with CFP2 and CFP4 interfaces.
(<http://www.hitechglobal.com/Boards/x16PCIExpress-Gen3.htm>)
2. **HiTech Global HTG-V7-OPTIC-690 [HTG-707]**; Virtex-7 (690) FPGA with GTH Transceiver, Optical module with CFP, QSFP+ and SFP+ interfaces.
(<http://www.hitechglobal.com/Boards/Virtex7-optic.htm>)
3. **Xilinx Kintex-7 FPGA KC724 Characterization Kit**; Kintex-7 K325T FPGA with GTX transceivers with BullsEye connector.
(<http://www.xilinx.com/products/boards-and-kits/ck-k7-kc724-g.html>)

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



A. Reference Design Details

A.1 Overview

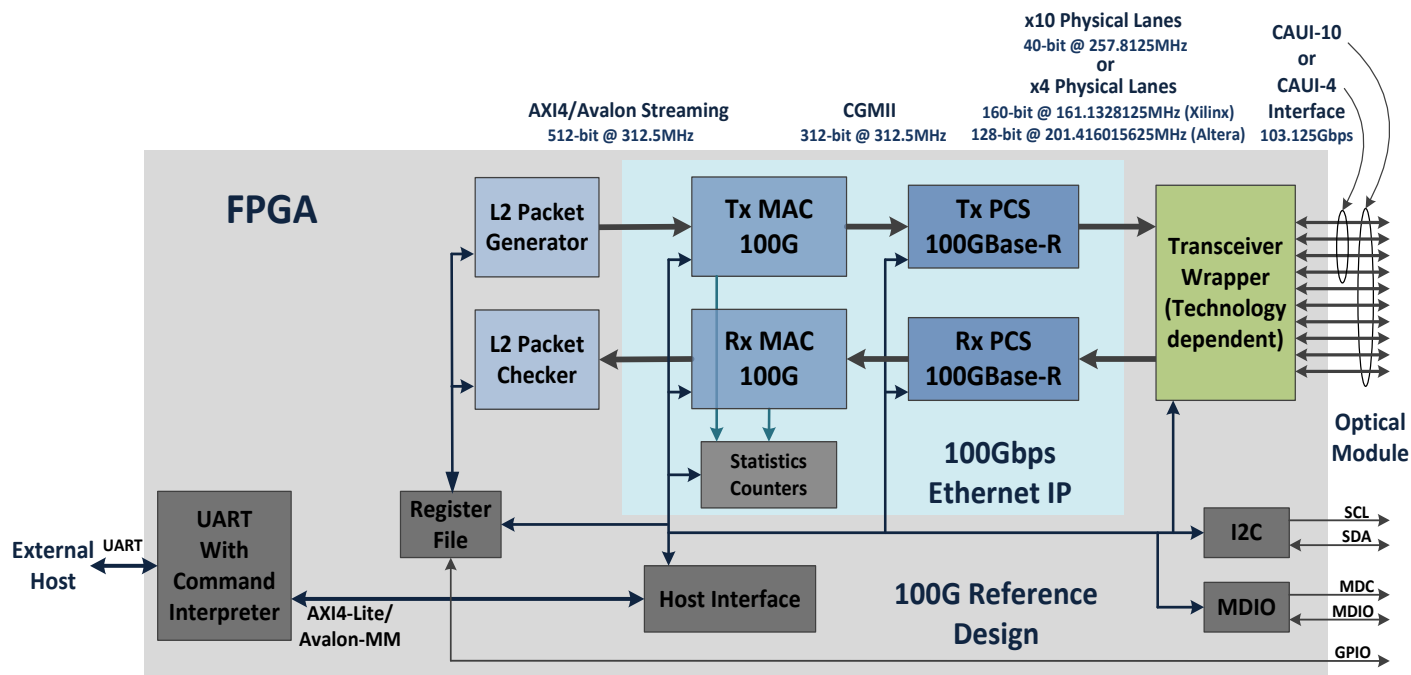
A 100Gbps reference design is included as part of the IP deliverable to facilitate quick L1 and L2 layer testing and verification of the 100Gbps Ethernet on target platform. The capability to run the L1 PRBS pattern and configure each transceiver independently can be used for a fast module bring-up in the lab and can also be used for factory diagnostics.

The UART (normally through an onboard USB-to-UART converter chip) based 100G Ethernet reference design can be seamlessly ported to various COTS FPGA networking and evaluation modules (see section for the list of verified modules). A GUI application controls the register read/writes to the FPGA through a UART core with integrated command interpreter. Both Linux and Windows platforms are supported for the UART based interface control.

This reference design can also be used on custom embedded design where the FPGA connects to the host processor via a PCIe interface. For the PCIe control interface, GUI application is hosted on a Linux platform (as PCIe driver/API is provided for Linux OS only).

A.2 Functional Description

Following figure shows the connectivity and the elements of the 100G Ethernet IP reference design. Usually the UART interface from the FPGA connects to an external (can be on the same module as well) USB-UART converter. A Linux or Windows host (through a USB port) running the GUI application is used to configure and control the 100G Ethernet. I2C, MDIO and GPIO interfaces included in the reference design can be used to control any optical module on the target platform including the 300Pin MSA (I2C), CXP (I2C) and CFP/CFP2 (MDIO) MSA compliant modules.



For L1 (physical layer verification and testing) GUI application provides an interface to independently control and configure all 10 10.3125Gbps (for CAUI-10 interface) or 4 25.78125Gbps (for CAUI-4 interface) transceivers used for 100G

100 Gigabit Ethernet IP Solution

Product Brief (HTK-100G-ETH-320-FPGA-Cxx)



Ethernet transport. User can configure the transceivers to run various PRBS pattern and configure various transceivers parameters like transmit voltage, transmit pre-emphasis, receive equalization and receive gain.

For L2 testing, GUI application uses the 100Gbps packet generator/checker inside the FPGA to generate and check MAC frames up to full line rate. The packet generator supports a basic rate control mechanism to control the packet/data rate on the interface. The generator can be configured for fixed size as well as pseudo random packet size packet transmission. An incrementing counter is used as payload for the MAC frames. The checker on the receive side verifies the payload of receive MAC frames and reports error in the payload.

A comprehensive set of transmit and receive counters in the MAC core provide a detailed view of the packet statistics including various error types.

Following is a snapshot for the GUI application for the L2 packet test results screen.

Address (Hex)	Register Name	Data
00000	FPGA:REFD:COMMON:REVISION	0x0
00010	FPGA:REFD:COMMON:SYS_GPIO_CTRL	0x0
00014	FPGA:REFD:COMMON:SYS_GPIO_STAT	0x0
000b0	FPGA:REFD:ETH-100g:ETH_GPIO_CTRL	0x0
000b4	FPGA:REFD:ETH-100g:ETH_GPIO_STAT	0x0
01000	FPGA:I2C:BLOCK-00:REVISION	0x0
01004	FPGA:I2C:BLOCK-00:I2C_XFER_REQ	0x0
01008	FPGA:I2C:BLOCK-00:I2C_CLK_DIV	0x0
0100c	(Multiple 2/2 visible)	0x0
	FPGA:I2C:BLOCK-00:I2C_RDY_INTR	0x0
	FPGA:I2C:BLOCK-00:RDY_INTR_MASK	0x0
01010	FPGA:I2C:BLOCK-00:I2C_INSTR	0x0
01014	FPGA:I2C:BLOCK-00:I2C_RDATA	0x0
01018	(Multiple 3/3 visible)	0x0
	FPGA:I2C:BLOCK-00:I2C_BYTE_NACK	0x0
	FPGA:I2C:BLOCK-00:I2C_BUS_LOST	0x0
	FPGA:I2C:BLOCK-00:I2C_RDY	0x0
02000	FPGA:MDIO:BLOCK-00:REVISION	0x0