

Based on announcements from vendors, enterprises and service providers, 100G system deployment is finally gaining real traction in the marketplace. The primary driver for this deployment is the customers' ceaseless demand for higher bandwidth. Various standard bodies are working to ratify the emerging 100G standards for transport and Ethernet, as well as optical interfaces. Due to their flexibility, FPGAs play a vital role for early adopters who want to design 100G systems today, prior to the standards being ratified.

Altera® Stratix® IV GT FPGAs solve the problem for both 100G transport and 100G Ethernet by providing integrated 11.3-Gbps transceivers in the 40-nm technology node. In addition, Altera's Stratix IV GX, Arria® II GX, and Arria II GZ FPGAs, and HardCopy® series ASICs can satisfy many other application needs in Optical Transport Networks (OTNs). These devices are ideal platforms for designing high-bandwidth systems, and provide a cost-effective and quick time-to-market solution.

Introduction

The increasing loads on today's networks are making it more difficult for vendors to deploy and manage their advanced systems. To accommodate the ever-increasing demand for greater bandwidth, OTNs have become the backbone for our next-generation networks. Fiber optics are quickly replacing copper wire and other media to become the fastest and most reliable media.

Two things are important in a network: speed and reliability. The network must be up all the time and it must be fast. Yet, the load on networks has increased tremendously. Data is a minor component of what the network carries. Voice, sound, and multimedia now form the major components carried through the network.

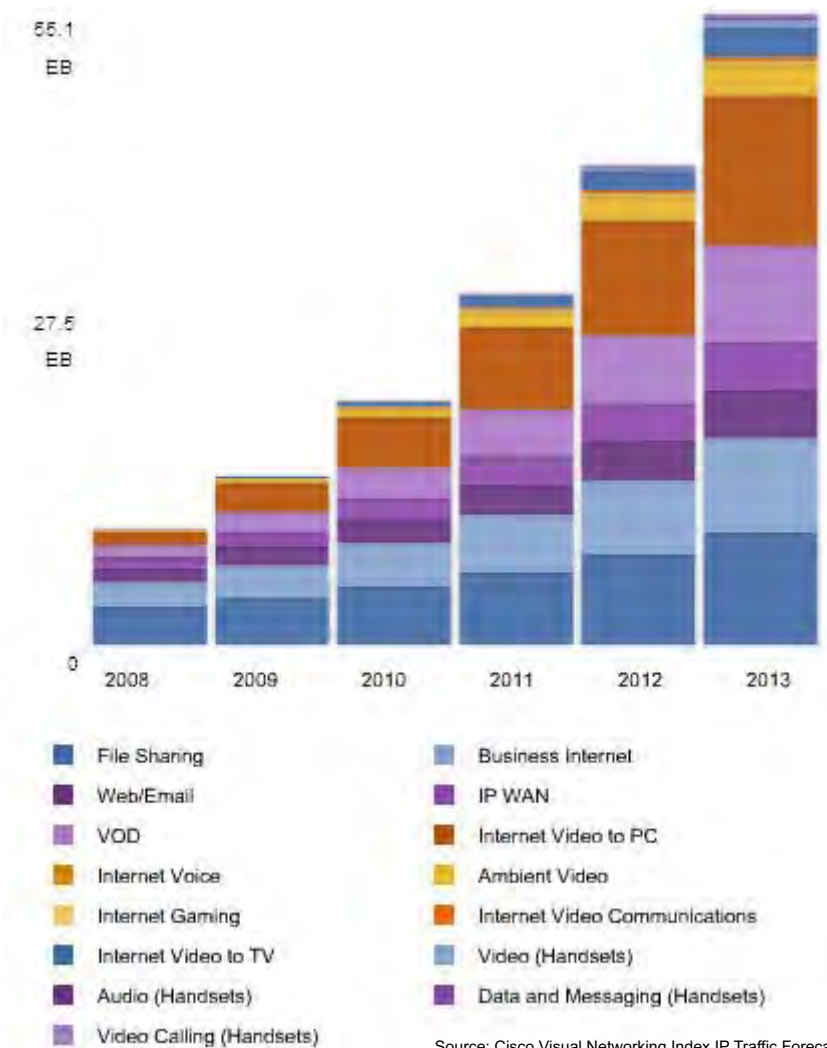
As [Figure 1](#) shows, total IP traffic will increase by a factor of six—nearly doubling every two years—from 2007 to 2012. By 2012, the annual rate of traffic will be 522 exabytes (10¹⁸ or half a zettabyte) per year. The main driver for this exponential growth will be high-definition video and high-speed broadband consumer applications.



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2010 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, and specific device designations are trademarks and/or service marks of Altera Corporation in the U.S. and other countries. All other words and logos identified as trademarks and/or service marks are the property of Altera Corporation or their respective owners. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.



Figure 1. Total Traffic Bandwidth Increases

Satisfying the Demand for High Bandwidth

End users do not want any break in their network service. They expect video conferences to have uninterrupted pictures and sound, just like television and telephones. OTNs are the only network backbone transport layer that is capable of supporting a 100 Gigabit Ethernet (GbE) LAN PHY, the standard of the next-generation Ethernet networks and the only standard that can provide the speed and reliability needed. Until new technology comes along, the OTN standard holds center stage, as it is the fastest and the most efficient. The high speed that OTNs provide is appreciable as well as scalable to meet future demands.

All forms of electronic communication utilize a packet or streams of packets, the information that a user wants to send, and a media, the type of transport on which the packets are carried. The faster the transport, the faster the packets will arrive. However, problems are occurring at the sending and receiving ends, where packets are arriving faster than they can be disbursed. Therefore, to increase efficiency, communication companies are putting in OTNs.

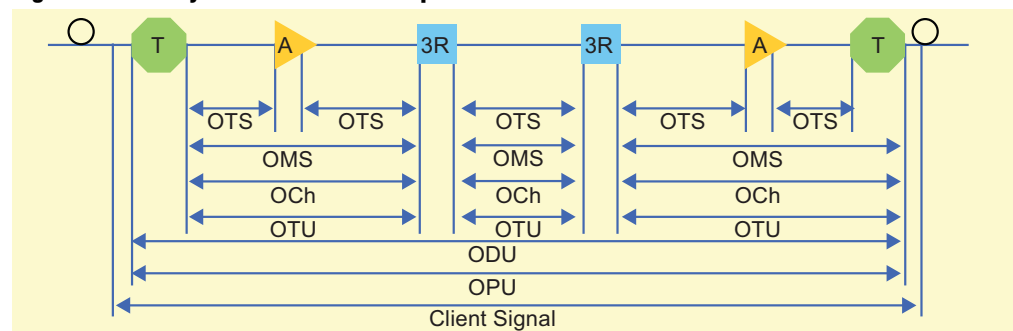
100G OTN (OTU4) Overview

By definition, 100G transport packets delivered via optical transport equipment are designed to expedite the transmission of any type of 100G data, which is encapsulated in either an OTN or Ethernet format. The overall traffic is distributed across metro, regional, or long-haul dense-wavelength division multiplexing (DWDM) networks. The current focus of the ITU study team is to leverage the existing 100G Ethernet specification, IEEE 802.3ba, and deploy 100G OTN over existing 40G and 10G infrastructures. This would result in satisfying the higher bandwidth requirement, reducing system complexity, and ultimately lowering cost, with a reduction of wavelengths to manage and increase overall spectral efficiency. Current 100G Ethernet deployment, by definition, covers a shorter distance than 100G transport, typically 40 km. 100G Ethernet and 100G transport have similar goals, to find a way to have faster connectivity to enable higher performance at a lower cost.

An OTN, consisting of networking capabilities and the protocol requirement to enable these capabilities, transmits information over optical media in a systematic manner. The focus of this paper is the transport and networking of Ethernet payloads over fiber optic cables. Established OTN mechanisms such as synchronous digital hierarchy (SDH) may fall within this broad definition, but our primary focus will remain on the LAN over WAN application, specifically the 40GbE and 100GbE applications (802.3ba). (For the purposes of this standardization and work plan, all new OTN functionality and related technologies are considered part of the Telecommunication Standardization Sector (ITU-T standards).)

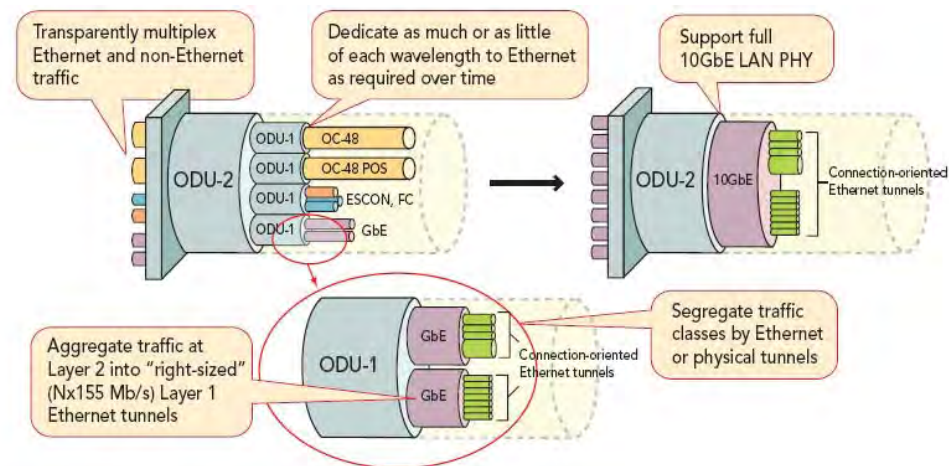
An OTN is composed of a set of optical network elements (shown in [Figure 2](#)) connected by optical fibre links, and able to provide transport, multiplexing, routing, management, supervision and survivability of optical channels carrying client signals, according to the requirements given in Recommendation G.872.

Figure 2. OTN Layers and Network Components



A distinguishing characteristic of the OTN is its provision of transport for any digital signal, independent of client-specific aspects (i.e., client independence). To do so according to the general functional modeling described in Recommendation G.805, the OTN boundary is placed across the optical channel/client adaptation to include the server specific processes and leave out the client specific processes, as shown in [Figure 3](#).

Figure 3. Aggregating Different Protocols on the Client Side Makes OTN a Cost-Effective Common Infrastructure



FPGAs play a vital role in the implementation of this flexible client adoption system across specific optical bandwidth channels. From the OTN implementation point of view, it is a data aggregation of various independent ports to provide the required bandwidth. Table 1 shows the data rates supported by current OTN standards. OTU4 will add a line rate of 100G.

Table 1. OTN Data Rates

G.709 Interface	Line Rate (Gbps)	SONET/SDH Rate	Line Rate (Gbps)
OTU1	2.666	OC-48/STM-16	2.488
OTU2	10.709	OC-192/STM-64	9.953
OTU3	43.018	OC-768/STM-256	39.813

Ethernet Frames over Transport

Ethernet currently is the dominant LAN technology in the private and enterprise sector, and emerging multiprotocol/multiservice Ethernet services are also offered over public transport networks. Defined by a set of IEEE 802 standards, public Ethernet services and frames over transport standards and implementation agreements are being debated in the ITU-T and other organizations. Ethernet can be described in the context of three major components: the service aspects, the network layer, and the physical layer.

Service Aspects

The public Ethernet service aspects (for service providers) include the different service markets, topology options, and ownership models. Public Ethernet services are defined by the ownership models employed and the type(s) of topologies used.

The topology options are categorized by the three types of services they support: line services, LAN services, and access services. Line services are point-to-point in nature and include services like Ethernet private and virtual lines. LAN services are multipoint-to-multipoint in nature, and include virtual LAN services. Access services are of hub-and-spoke nature, and enable single ISP/ASPs to serve multiple customers. (Due to their similarity from a public network perspective, line and access services may be essentially the same.)

The services provide different service qualities. A circuit-switched technology like SDH provides a guaranteed bit rate, while a packet-switched technology such as MPLS can provide various service qualities from best effort traffic to a guaranteed bit rate. Ethernet services can be provided for the Ethernet MAC layer or the Ethernet physical layer.

Network Layer

The Ethernet network layer provides end-to-end transmission of Ethernet MAC frames between Ethernet end-points of individual services that are identified by their MAC addresses. Ethernet MAC layer services can be provided as line, LAN, and access services over circuit-switched technologies like SDH VCs and OTN ODUs, or over packet-switched technologies like MPLS and RPR. For the Ethernet LAN service, Ethernet MAC bridging can be performed within the public transport network in order to forward the MAC frames to the correct destination. Ethernet MAC services can be provided at any bit rate, because they are not bound to the physical data rates (i.e., 10 Mbps, 100 Mbps, 1 Gbps, 10 Gbps, 100 Gbps) defined by IEEE standards.

Physical Layer

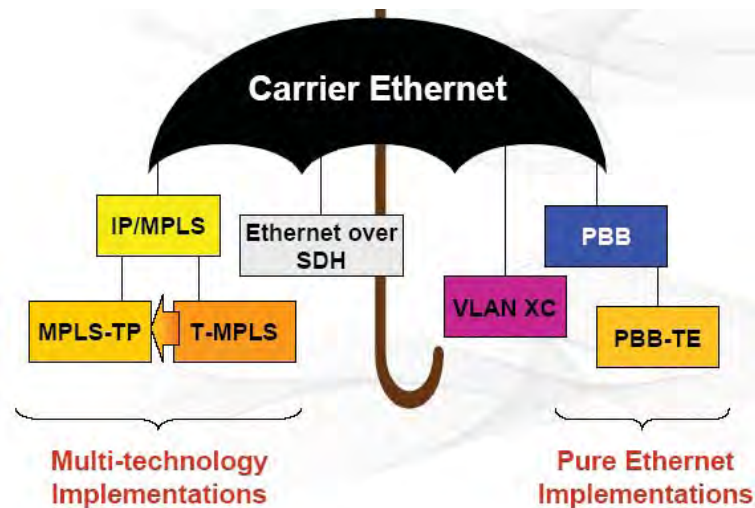
The IEEE has defined a distinct set of physical layer data rates for Ethernet with a set of interface options (electrical or optical). An Ethernet physical layer moves signals, such as a 10GbE WAN signal over an OTN or a 1GbE signal over SDH using transparent GFP mapping, transparently over a public transport network. Ethernet physical layer services are point-to-point only and are always at the standardized data rates. They are less flexible compared to Ethernet MAC layer services, but offer lower latencies.

Supporting OTN with Carrier-Class Ethernet Standards

Although Ethernet was originally designed to be used in a LAN environment, it now is used widely in backbone or metro area networks (MANs). Ethernet has been enhanced in several aspects, including high bit rate and long-reach interfaces, Ethernet-based access networks, virtualization of networks, scalability, backbone provider bridges, reliability in protection technologies, QoS traffic control and traffic conditioning, and higher bit rates so that it can carry a network operators' network. In addition, Ethernet can easily achieve multipoint-to-multipoint connectivity, which requires $n \times (n - 1) / 2$ connections in an existing point-to-point transport technology.

As shown in [Figure 4](#), Carrier Ethernet is an attempt to expand Ethernet beyond the borders of LAN and into the territory of WAN, and to explore the overall communication ecosystem. The aim is to provide customers with a WAN to connect sites together, in the same way that ATM, Frame Relay, and X.25 services from carriers did in the past. Carrier Ethernet is not about the Ethernet within LANs such as customers are used to seeing at our desks and in server rooms.

Figure 4. Multiple Protocols Tagged Under Carrier Ethernet



The process of converting Ethernet to Carrier Ethernet Transport started not long ago. So far, the ITU-T has provided options for building an Ethernet service-based carrier network system. ITU-T recommendations deal with Ethernet over Transport (EoT), where transport is a traditional carrier technology such as PDH, SDH, or OTN.

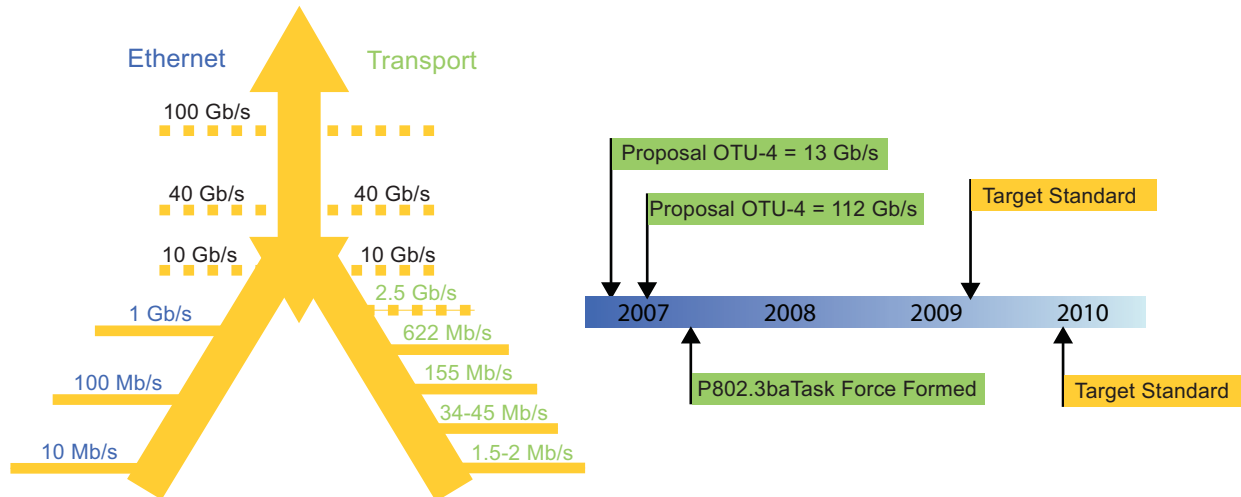
40G/100G Ethernet Architecture to Support OTN

IEEE 802.3ba is developing standards both for 40 Gbps and 100 Gbps. The current objectives are:

- Support full-duplex operation only
- Preserve the 802.3/Ethernet frame format using the 802.3 MAC standard
- Preserve the minimum and maximum frame sizes of the current 802.3 standard
- Support a BER of $\approx 10^{-12}$ at the MAC/PLS service interface
- Provide appropriate support for OTN
- Support a MAC data rate of 40 Gbps
- Provide physical-layer specifications which support 40-Gbps operation over:
 - ≥ 10 km on SMF
 - ≥ 100 m on OM3 MMF
 - ≥ 10 m over a copper cable assembly
 - ≥ 1 m over a backplane
- Support a MAC data rate of 100 Gbps
- Provide physical-layer specifications which support 100-Gbps operation over:
 - ≥ 40 km on SMF
 - ≥ 10 km on SMF
 - ≥ 100 m on OM3 MMF
 - ≥ 10 m over a copper cable assembly

As shown in Figure 5, this project is targeted to be complete by the middle of 2010. The industry activity for 100G deployment includes both transport and Ethernet. The transport and Ethernet standards are converging at a uniform speed of 100 Gbps, a process that started with 10G.

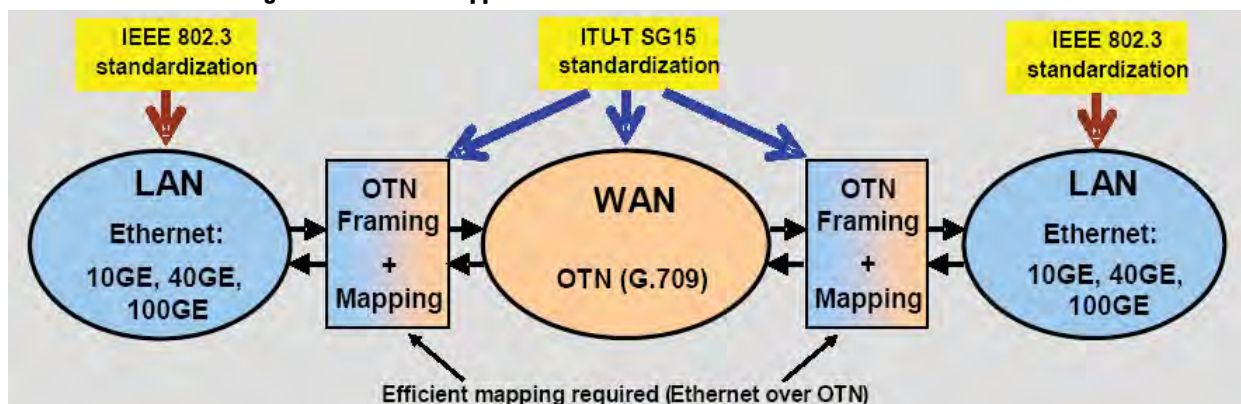
Figure 5. Convergence of Ethernet and Optical Transport Begins With 10 Gbps



Stratix IV FPGAs Pave the Way for 100GbE OTN Designs

The current industry trend is to use Ethernet over WDM for packet transport and IP/MPLS/Ethernet for data transport. Altera’s 40-nm Stratix IV FPGA family is well positioned to meet the performance and system bandwidth requirements for 100G Ethernet and Transport system designs. Stratix IV GT FPGAs provide the highest density with integrated 11.3-Gbps transceivers, which are essential for implementing 100GbE/Fibre Channel/RPR MAC functions into a single device, as well as handling key functions like forward error correction (FEC), mapping, and framing of OTN packets. The OTU-4 standards for 100GbE use enhanced FEC (EFEC), which must be designed with specific algorithms to make sure that the optical bandwidth can be used to its fullest extent. Due to its superior fabric performance, Stratix IV GT FPGAs can handle the EFEC functions as well, so they are an ideal platform for OTN system designed for algorithm implementation and testing. Figure 6 shows how customers are using Stratix IV GT FPGAs to implement all of the above functions when designing 100GbE OTN boxes.

Figure 6. 100G OTN Application: LAN over WAN



OTN and the Need for Universal Client Ports

OTN encompasses a variety of optical network elements, each of which provides a unique role in efficiently transporting traffic. Independent voice, video, data, and storage networks have evolved to a form a common backbone network that is served by OTN. OTN equipment must map many different types of traffic (Ethernet, SONET/SDH, ESCON, Fiber Channel, and video) into this common backbone.

As optical equipment manufacturers continue to reduce costs and leverage components across multiple platforms, flexible solutions for mapping various client ports have gained adoption. FPGAs have become a central component delivering “universal client ports,” which can be configured to support any mix of client interfaces. This allows a single component to be used in multiple applications efficiently.

Scalable Support for OTN

Altera has a portfolio of products optimized across the entire OTN hierarchy, as shown in [Table 2](#).

Table 2. Altera Device Portfolio

OTU Type	Bit Rate	Client Interfaces	Chip/Module Interface	Altera Device
OTU1	2.5 Gbps	GbE, OC-3/48, STM-1/16, Fibre Channel	OC-48	Arria II GX, Arria II GZ
OTU2	10 Gbps	10GbE WAN, 10GbE LAN, OC-48/192/STM-16/64 10G Fibre Channel	SFI-4.1, SFI-4.2	Arria II GX, Arria II GZ, Stratix IV GX, HardCopy IV
OTU3	40 Gbps	40GbE, OC-768/STM-256	SFI-5.1, SFI-5.2	Stratix IV GX, HardCopy IV, Stratix IV GT
OTU4	100 Gbps	40GbE, 100GbE	SFI-S, MLD	Stratix IV GT

As volumes have ramped, OTN1 and OTN2 applications have become increasingly cost and power sensitive. As shown in [Table 3](#), Altera’s Arria II GX and Arria II GZ FPGAs with embedded transceivers deliver the functionality needed to implement OTN1 and OTN2 muxponders and cross connects with power and cost efficiency.

Table 3. Altera’s Arria II GX Transceiver Protocols for OTN Applications

Standards	Data Rate in Gbps
SGMII	1.25
GbE	1.25
10G Ethernet (XAUI)	3.125
SONET OC-3/OC-12/OC-48	0.155, 0.622, 2.488

The flexibility of Arria II GX and Arria II GZ FPGAs offers several benefits over fixed standard product solutions:

- Support of emerging mapping technologies, such as ODU0, which is necessary for GE mapping into OTN
- Configuration to support a various mix of client-side interfaces, allowing for multiple implementations with the same device
- Support for Multiple FEC and EFEC techniques within the same device, simply by reconfiguring the FPGA

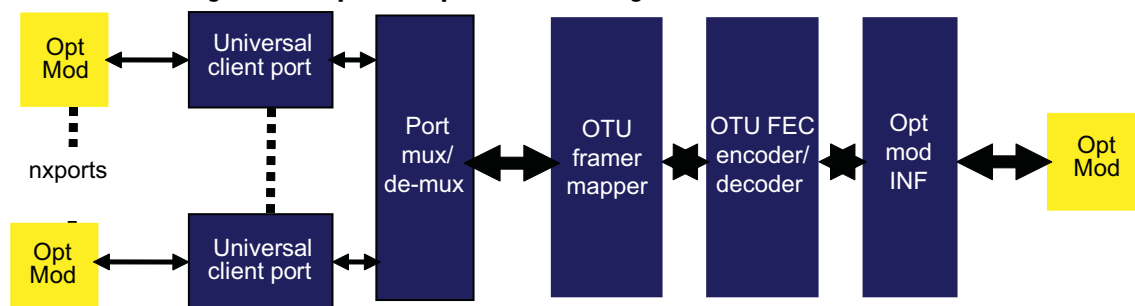
40G Muxponder Design in Single FPGAs

The primary focus of muxponders (multiplexing transponders) is to aggregate multiple, lower rate client signals into a higher rate wavelength carrier. This strategy makes more efficient use of the WDM spectrum rather than having individual lower speed clients allocated to independent wavelengths.

Industry analysts have forecasted that 40G optical ports will grow dramatically through 2013. The increasing throughput of 40G OTN equipment places a greater demand on improved FEC techniques to allow for sending the signals greater distances. Because these EFEC standards require significantly more logic capacity to implement, integrating universal client ports, mappers, framers, and EFEC into a single device is challenging at 40 Gbps.

However, Altera's Stratix IV GX family has been architected to support a full 40G muxponder in a single device, as shown in [Figure 7](#). Stratix IV GX FPGAs provide efficient implementation for a comprehensive mix of data, storage, TDM, and video protocols including GbE, Fibre Channel (1G, 2G, 4G), SONET (OC-N), and SDH (STM-N) by providing the necessary density and fabric performance. All the listed protocols are either supported directly from Altera or through a partner ecosystem.

Figure 7. Muxponder Implementation Using Stratix IV FPGAs



The Stratix IV GX family supports up to 32 transceivers with clock data recovery (CDR) and data rates from 600 Mbps to 8.5 Gbps, plus up to an additional 16 transceivers with CDR, supporting data rates from 600 Mbps to 6.5 Gbps. Stratix IV GX FPGAs also support up to 530K logic elements (LEs), which can support a full 40G muxponder application. [Table 4](#) shows Stratix IV GX transceiver support for various OTN data rates.

Table 4. Stratix IV GX Transceiver Protocols for OTN Applications

Protocol	Data Rate
OIF CEI-6G	4.976 Gbps–6.375 Gbps
Interlaken	3.125 Gbps–6.375 Gbps
10-GbE XAU1	3.125 Gbps
HiGig	3.75 Gbps
SFI-5	2.488 Gbps–3.125 Gbps
GbE	1.25 Gbps
SDH/SONET OC-12	622 Mbps
SDH/SONET OC-48	2.488 Gbps
Fibre Channel	1.0625 Gbps, 2.125 Gbps, 4.25 Gbps, 8.5 Gbps

100G OTN Designs with Stratix IV GT FPGAs

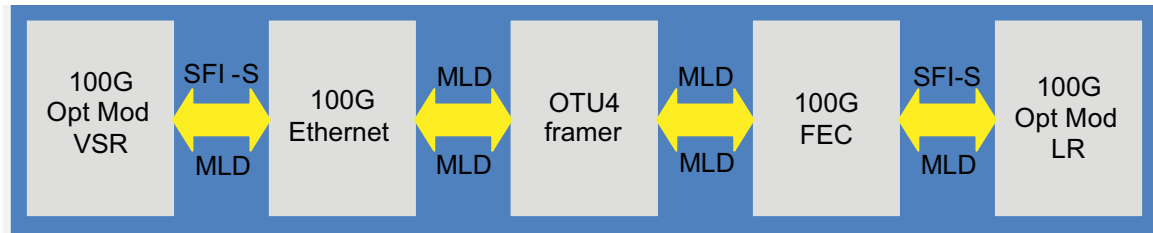
Recent standards activities are aligning around OTN4 for 100G optical transport. These applications demand a combination of high-speed transceivers and 10G transceivers to support the necessary throughput requirements, as well as core performance and logic density to handle the complex processing needed for managing the 100G data traffic. Stratix IV GT FPGAs enable a direct connection to the 10G optical interface on the client side, as well as a direct connection to 100G CFP or QSFP modules on the network side. This is a key benefit as it eliminates the use of external PHY devices and simplifies the overall system complexity. Table 5 lists the protocol support for Stratix IV GT FPGAs.

Table 5. Stratix IV GT Transceiver Protocols for OTN Applications

Protocol	Data Rate (Gbps per Lane)
100G IEEE 802.3ba	10.3125
10G Fibre Channel	10.3125
10G IEEE 802.3ae	10.3125
40G IEEE 802.3ba	10.3125
OTN-2	9.9–11.3
OTN-2 G.709	10.7
OTN-3	9.9–11.3
OTN-4	9.9–11.3
OTN-4 MLD	9.9–11.3
SFI-S (includes SFI-5.2)	9.9–11.3
SONET/SDH OC192/STM-64	9.9–11.3

In addition, the devices support bonded interfaces such as MLD and SFI-S for chip-to-module and chip-to-chip connections. Equipment manufacturers can develop early versions of OTN4 muxponders, transponders, and regenerators as the standards within the ITU and OIF evolve. Figure 8 shows how a 100G WAN over LAN OTN system can be implemented using Stratix IV GT and GX devices.

Figure 8. 100G WAN over LAN OTN System Using Stratix IV FPGAs



Conclusion

Altera's current 40-nm FPGA portfolio along with its partner ecosystem is well suited for both the emerging 100G OTU4 standards as well as the traditional OTN solution with line rates from 2.5 Gbps (OTU1) to 10 Gbps (OTU4). Altera's Stratix IV GX, Arria II GX, and Arria II GZ FPGAs, and HardCopy series ASICs cover the entire spectrum of Optical Transport applications like MSPPs, P-OTN, and Carrier Ethernet transport.

In an area of evolving standards and new protocol deployment, such as OTU4, which requires data rates of 10G and high density device for complete 100G solution, Altera's Stratix IV GT devices are the only FPGAs available today that are architected to meet the 100G system requirements.

Further Information

- Stratix IV FPGA 40G/100G IP Solutions:
www.altera.com/products/devices/stratix-fpgas/stratix-iv/transceivers/stxiv-40g-100g-ip-solutions.html
- Arria II FPGAs: Cost-Optimized, Lowest Power 6G Transceiver FPGAs:
www.altera.com/products/devices/arria-fpgas/arria-ii-gx/aiigx-index.jsp
- Wireline Transmission Technologies:
www.altera.com/end-markets/wireline/transmission/wil-transmission.html
- ITU-T SG13 and SG15:
www.itu.int
- IEEE 802.1 WG and IEEE 802.3 WG:
www.ieee802.org
- IETF:
www.ietf.org
- Metro Ethernet Forum:
www.metroethernetforum.org

Acknowledgements

- Rishi Chugh, Product Marketing Manager, Components Product Marketing, Altera Corporation
- Kevin Cackovic, Strategic Marketing Manager, Communication Business Division, Altera Corporation

Document Revision History

Table 6 shows the revision history for this document.

Table 6. Document Revision History

Date	Version	Changes
July 2010	1.1	<ul style="list-style-type: none"> ■ Updated Table 3, Further Information. ■ Minor text edits.
October 2009	1.0	Initial release.