40/100 GigE Markets: 2010 and Beyond

Chapter One

January 2010

Table of Contents

Chapter Two: Higher-Speed Ethernet and Its Applications Drivers

2.1 Introduction: Massively Growing Needs from Data Centers and ISPs
2.2 Video On the Rise: The Business to Consumer Switch
2.2.1 Consumer Video: The Real Video Driver for 40/100 GigE
2.2.2 Business Video: Not What We Thought it Would Be
2.3 Data Centers and Enterprise Networks: Data Rate and Bandwidth Requirements
2.3.1 Aggregation and Faster Interfaces
2.3.2 Servers, Switches and Storage: Where Technologies Happen First
2.3.3 The Impact of Convergence in the Data Center and FCoE
2.4 High-Performance Computing: The People Who Brought Us 100 GigE
2.4.1 Bandwidth Requirements for HPC
2.4.2 InfiniBand versus Ethernet versus Fibre Channel?
2.5 Central Offices, POPs and Internet Exchanges
2.5.1 VSR OC-768: The Old-Fashioned Way to 40 Gbps
2.6 Long-Haul, Metro and Access Networks
2.6.1 Carrier 100G: What Will It Look Like?
2.7 Technical Requirements for 40G and 100G Networks
2.8 Design Architectures for 40G and 100G
2.9 Multiple PMD Types Supported by a Single Connector

Chapter Three: Evolution of Technology for 40G and 100G Networks

3.1 Introduction
3.2 40GBase-KR4 (1m Backplane) Technology
3.3 40GBase-CR4 and 100GBase-CR10 (10m Copper Cable) Technology
3.4 40GBase-SR4 and 100GBase-SR10 (100m LOMF) Technology
3.5 40GBase-LR4 (10km SM Fiber) Technology
3.6 100GBase-LR4 Technology
3.7 100GBase-ER4 Technology
3.8 Future Variants for 40G and 100G Ethernet
3.8.1 Future 40G Variants
3.8.2 Future 100G Variants
3.9 Basic Components for Higher-Speed Ethernet
3.9.1 Integration Imperative
3.9.2 Parallel vs. Serial Transmission
3.10 The Role for Optical Integration and Silicon Photonics
3.10.1 Materials Platforms for Optical Integration
3.11 Potential for Serial and Parallel Solutions for High-speed Optical Networks
3.12 Lasers, TOSAs and Modulators

Chapter Four: Evolution of Standards and MSAs Beyond 10-Gig

Chapter Five: Five-Year Forecasts of 40-Gbps and 100-Gbps Data Networks

Chapter One: Introduction

1.1 Background to this Report

CIR has been following the beyond 10G market for more than three years now, starting in 2007 when the first rumors surfaced that the IEEE would be working on a 100G project. Much has changed since then. For one thing, back in 2007, much was made of how 100G would represent a new start for carriers and that the long-time supporters of Ethernet would give no comfort to the old ways of doing things. 40G would be a thing of the past and that anyone in the carrier community who was old-fashioned enough to insist on 40G would have to go with the ultra-expensive SONET/SDH OC-768.

Within a year, it soon became clear that IEEE 802.3ba, Higher Speed Ethernet, would be culminating in a standard that actually covered both 40 *and* 100G, although this, in the end, had less to do with reactionary SONET/SDH supporters than with what could be practically and cost effectively achieved with today's technology. In addition, larger servers were increasingly in need of something better than 10 GigE and it was clear that they needed it now.

In CIR's previous 100 GigE report, we provided insight and analysis on the new directions for Higher Speed Ethernet and little has changed – or likely to change – in the general trends. However, as we enter 2010 we are entering the first year in which 40/100 GigE products will be available, if only in a limited way. The 40 and 100G Ethernet standard is expected to be published in June 2010 and pre-standard devices are anticipated in the first quarter 2010. In fact, it would not be surprising if the OFC/NFOEC 2010 turned out to be the "coming out" for 40 and 100G initial product offerings.

Bearing all this in mind, CIR believes it is time to take another look at the 40/100 G market and the new business opportunities that it will bring in its wake.

The Markets for 40 and 100G: Bigger than First Thought

The need for speed is not a new one. Ten years ago skeptics were saying that we would never see Gigabit to the desktop, yet, today, that is just what we have. And, we had it five years ago. Within two or three years, it seems likely that GigE will be ubiquitous on all laptops

and desktops and 10 GigE will have become standard on all servers. Although the 40/100 GigE efforts were explicitly conceived with the idea that this technology would cater only for "power users" and heavily used backbone networks, there is actually no reason to suppose that 40/100 GigE won't make it all the way to the desktop one of these days.

In fact, the initial 100G efforts were user driven. 100G was needed in 2007 according to some proponents such as Google, Facebook, the Amsterdam Internet Exchange and Lawrence Livermore Labs. Now, both Google and Facebook are saying they need Terabit Ethernet! While companies running 10G right now are few and far between, there are power users such as these that have historically pushed the envelope and will continue to do so.

As we have already noted, the original thinking around the next generation of Ethernet supported a 100G standard only. This was, of course, because traditionally, Ethernet has taken x10 steps in speed. However, over the past couple of years, 40G technology has been seen as a way to get the post 10 GigE generation of products up and running.

Getting to market with 40G products which will eventually morph into 100G products makes sense because the 40G technology already exists so why not use it? OC-768 SONET has been around for more than seven years, the PCIe bus internal to the server has a 40G option and InfiniBand also has a 40G variant. The IEEE is very good at taking existing technologies and leveraging them so 40 GigE is no exception. Those close to the standards activity say that 40 GigE was mainly developed for servers and probably will not be used in the rest of the network, but the 100-meter reach will cover more than 90 percent of the channels in data centers and most LAN backbones, so CIR expects to see it in more than just servers.

Already the 40/10 GigE technology is showing signs of moving into markets that are much larger than the server mega-farms for which it was originally intended. And with real 40/100 G products entering the market soon, who will buy what and when are issues of immediate practical importance.

Has 100 GigE Killed SONET?

The death of SONET/SDH started to be declared just a few years after its birth, but has proved to be much exaggerated. Even now SONET/SDH transceiver/transponder

manufacturers continue to do a profitable business, although admittedly they treat this part of their efforts as very much a "cash cow" and do not put much marketing emphasis on it. CIR has always been at the forefront of those declaring SONET/SDH not dead yet. However, with the arrival of higher speed Ethernet, we believe it is finally time to read the last rites over SONET/SDH. The future, it seems, in the carrier networks will be one in which Ethernet is carried over an optical transport layer.

Ethernet data rates are usurping SONET ones, or will be very soon. There was a convergence at 10G and many thought SONET would die a slow death then, but OC-768 slowed that process. While it is really the only currently viable 40G end-to-end solution, OC-768 is too expensive for LANs so a cost-effective "Ethernetized" 40G has started to emerge rather quickly. This version, however, will not be used in the public network so OC-768 has its place – at least until long-reach 100G Ethernet appears, which is expected sometime in 2012. Nonetheless, in the past year we have seen clear indications that SONET/SDH is doomed and that we will have Ethernet everywhere eventually. Evidence of this is found in several places:

- The uptake of existing Carrier Ethernet at 10G is increasing. Until recently, Carrier Ethernet was close to an oxymoron

- The new-found close working relationship between the IEEE 802.3ba Task Force, the ITU-T OTN Working Group and the OIF. For much of the past two decades, the relationship between the ITU-T (the defenders of SONET) and the IEEE 802 project (defenders of Ethernet) was frosty

- The fact that there is no ITU-T project for any SONET/SDH standards beyond OC-768. As recently as two years ago, a post OC-768 era in SONET/SDH had not been completely ruled out. Carriers also talk increasingly about 100 Gbps backbones and (because there is nothing else available) they can only be referring to Ethernet deployment

All of this means that the Carrier Ethernet is no longer a niche market, but going forward a potentially huge market; raising the question as to what the 40/100 G products for this market will look like?

How Converged Networks in the Data Center Will Drive 40/100 GigE

Somewhat ironically – since Ethernet was conceived as a LAN standard – Ethernet is not rising to the prominence in the data center that it seems to be destined for it in the carrier network; it does not seem as likely that it will replace Fibre Channel and InfiniBand. The reason is that the factors that led FC and IB technologies to be adopted in the first place have not disappeared. Lower latency and lossless channels are the two attributes that keep IT managers loyal to these technologies. Ethernet has not yet solved its problems with these characteristics, although there are standards being developed to address both (IEEE 802.1).

A comparison of the networking speeds of Ethernet, SONET, InfiniBand and Fibre Channel are given in the Exhibit 1-1 below.

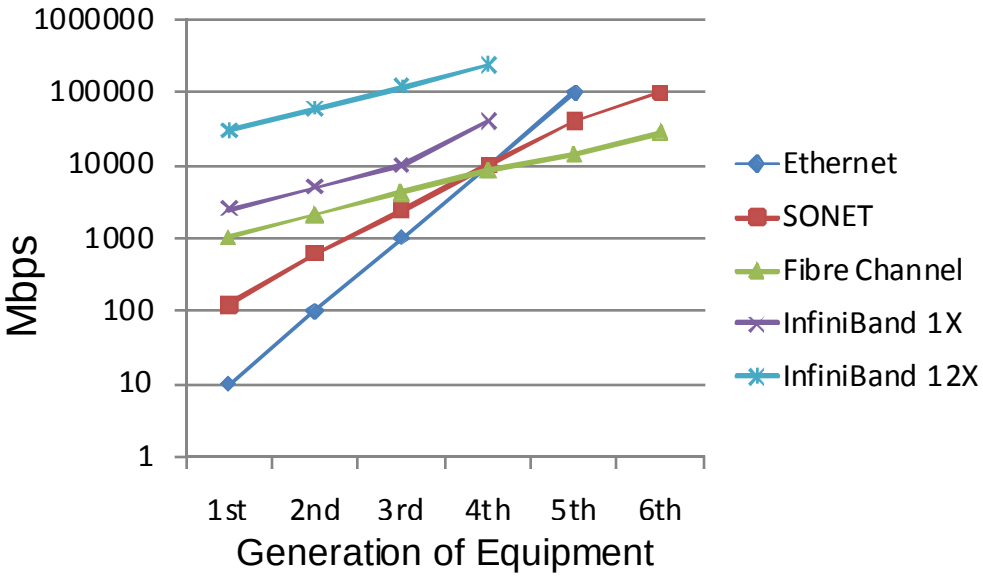

Exhibit 1-1: Comparison of Ethernet, SONET/SDH, FC and IB Data Rates

While Fibre Channel and InfiniBand are not expected to disappear from the data center, they *are* predicted to be transported over Ethernet outside of the SAN or server cluster. This will be accomplished by the work going on in T11.2 and IEEE 802.1 that are working on Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE). Today, the data

center has a complex network with at most times at least two networking protocols running – Ethernet and Fibre Channel. If there are server clusters then there might be a third with InfiniBand or proprietary clustering technology. This means that each server has at least three network cards (and probably six for redundancy). With FCoE and CEE, these three would be converged to one. Exhibit 1-2 shows the result – a much simpler network for the data center manager to administer.

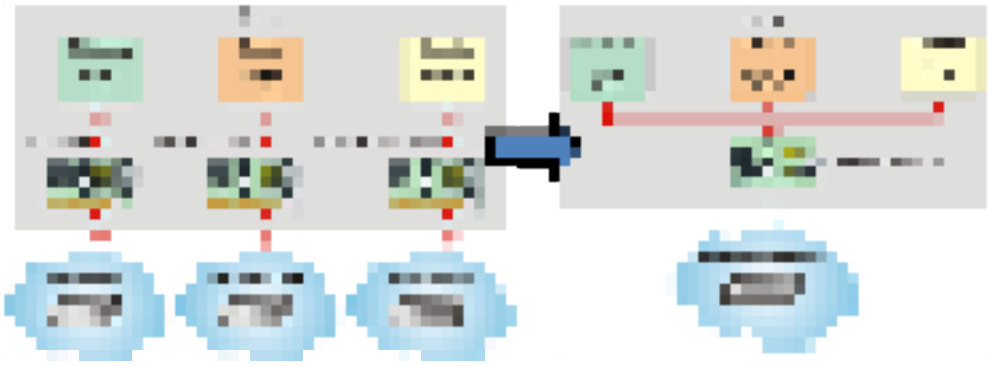

Source: Fibre Channel Industry Association

Exhibit 1-2: Fibre Channel, InfiniBand and Ethernet in the Data Center

While this consolidated network will make it easier for the data center manager, it will drive networking speeds up more rapidly because now three different networks are vying for the same I/O port. And this is precisely why these trends are so important to the future of 40/100 GigE; servers will need to have a minimum of 10G network cards and will probably need 40G in short order due to other factors like virtualization.

The impact of virtualization and converged data center networks on the technology and volume of optical components will be discussed in detail in this report. But given all of the above it is clear, that the latest trends in the data center, are likely to be very friendly to 40/100 GigE.

1.2 Objectives of this Report

This report is intended to provide in-depth analysis of the 40 and 100G Ethernet components markets and technologies in the light of all the latest developments set out above. The

specific purpose of this report is therefore to, first and foremost, spotlight likely new business opportunities for components, subsystems and modules firms. However, we also discuss what the impact of these newer networking technologies will be on service providers and equipment manufacturers.

1.3 Scope of this Report

Components covered in this report will be the same as those that have been covered in the last two CIR reports on 40 and 100G. These are detailed in Exhibit 1-3 below.

Exhibit 1-3

Component, Module and Subsystem Products Covered in this Report

Product Category	Focus for High-Speed Networking Products
Lasers, transmitters and TOSAs	High-speed lasers and laser arrays including VCSEL arrays Vertically integrated laser assemblies
Detectors, receivers and ROSAs	High-speed and coherent detectors and detector arrays Integrated receiver assemblies
WDM and parallel optics	CWDM for high-speed modules High-speed parallel optics using ribbon fiber, etc.
Dispersion compensation and error correction	Optical Dispersion Compensation, Electronic Dispersion Compensation (EDC) and Forward Error Correction (FEC)
Electronics	MACs, PHYs and PMDs Use of new materials and processes including 65-nm and 45-nm bulk CMOS, SiGe and SOI process technology
Modules	MSAs for next-generation networks
Media and cabling	MMF, SMF, active optical and copper cabling

In developing opinions about and forecasts for these devices, CIR, by default, must analyze equipment that uses these devices and also where that equipment is being installed. To that end, CIR must speak with some equipment manufacturers and service providers to verify demand-side economics. However, this report does not attempt to quantify the market for networking equipment, except in the context of component application.

1.4 Methodology and Information Sources for this Report

CIR prides itself on its independent analysis and level-headed forecasting. We interview representatives from not only the prominent components suppliers, but equipment manufacturers, service providers, standards development personnel and even end users when appropriate. Facts are gathered, analyzed and checked with common sense and extensive secondary research.

Secondary research for this report consisted of reviewing many sources including corporate Web sites, research journals, SEC reports, standards information, industry trade shows, industry conferences, industry marketing groups and previous CIR reports.

1.5 Plan of this Report

Chapter Two will detail Higher Speed Ethernet's applications drivers, which include video, HPC, data centers along with needs in the public network. Also discussed are other networking technologies and their relevance and how they fit with Ethernet. In Chapter Three, technologies that will enable 40 and 100G networks will be featured. This includes everything from high-speed backplanes to long-reach optics. In addition, different variants are analyzed along with how they will most likely be implemented.

Chapter Four is dedicated to explaining the standards and MSAs and their significance to optical components. Forecasts are contained in Chapter Five. The forecasts will include permutations by variant (reach), MSA form factor and speed.