



ethernet alliance

40 GbE: What, Why & Its Market Potential

Contributors:

Gautam Chanda, Cisco Systems
Yinglin (Frank) Yang, CommScope, Inc.

November 2010



Table of Contents

Executive Summary.....	3
Introduction: The Business Case for 40GbE.....	4
40GbE: A Closer Look.....	5
Transceivers	5
Cables & Connectors.....	6
Use of Parallel Optics	7
Deploying 40GbE in the Network.....	7
What is involved.....	8
Don't Wait to Start Future-Proofing Your Network.....	9
About the Ethernet Alliance	9
References	10



Executive Summary

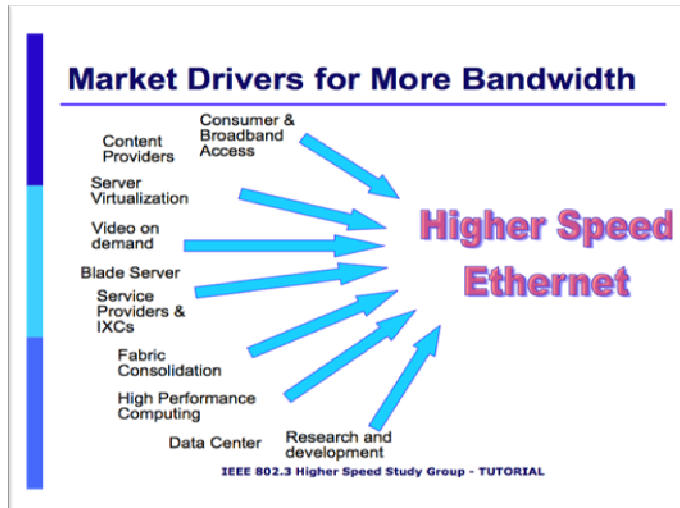


Figure 1. The current trends that are driving the demand for higher speed Ethernet.

The business case for 40GbE (Gigabit Ethernet) is becoming inescapably compelling. While 10GbE is still making its way into the data centers, CIOs and IT Managers must now consider how they are going to handle what's coming next (see Fig. 1): High-bandwidth applications such as server virtualization and cloud computing, fabric consolidation within the data center, and a demand for more high-performance computing among end-users. The need for faster data transfer rates is relentless and carries significant implications with regards to network productivity as well as OPEX costs.

So when the IEEE officially adopted IEEE Std. 802.3ba™ in June, 2010 paving the way for both 40GbE and 100GbE, it came not a moment too soon. The increased speed will allow networks to move newfound 10GbE resources to the access layer allowing the beefier 40GbE-enabled equipment to handle traffic at the aggregation and core layers. Based on analysts' forecast and the robust development efforts by OEMs, it is no longer a question as to *if* 40GbE will become an accepted part of the IT landscape, but *when* and *how*.

This white paper addresses the impending move to 40GbE, how it may change the network architecture, and what IT Managers can do now to prepare to migrate to the new standard.



Introduction: The Business Case for 40GbE

Since February 1980, when the first IEEE 802 standards committee convened, speeds in Ethernet delivery to all layers have made increasingly greater leaps over increasingly shorter intervals. Now, eight years to the month after the adoption of 10GbE, the IEEE has adopted 802.3ba, paving the way for 40GbE and 100GbE.

High-Speed Ethernet Trends (Source IEEE)

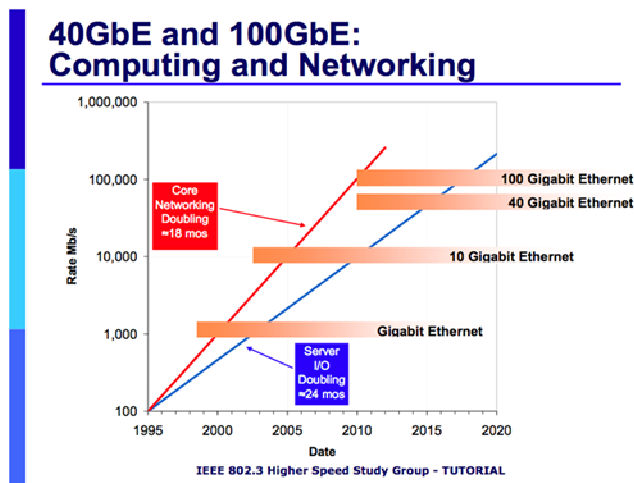


Figure 2. Represents leading edge deployment trends; general market deployment is typically delayed by several years.

As illustrated by Figure 2, I/O data transfer rates within the access layer are doubling every 24 months while transfer rates at the core layer double approximately once every 18 months. A primary driver behind the push to 40GbE is a new generation of high-speed, high-demand, computing applications and technologies. These include the spreading deployment of virtual servers and cloud computing. By the end of 2009, nearly one in five new servers were virtualized.¹ At the same time, the financial pressures of a challenging economy are forcing networks to look for ways to consolidate their resources in an effort to reduce OPEX and total cost of ownership.

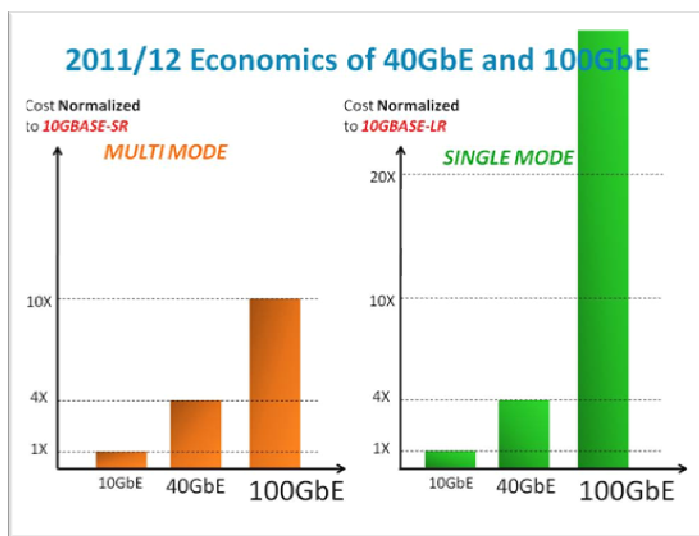


Figure 3. Implementation costs of 40GbE and 100GbE for Multimode and Single-Mode.

With four times the capacity and the ability to cost-effectively migrate to 100GbE (see Figure 3), 40GbE on multimode fiber is the next logical step in the evolution of the data network. The increase in urgency and the drop in prices have led analysts at the market research firm, the Dell’Oro Group, to predict that shipments in the 40-Gbps optical market will reach \$14.5 billion by 2013².

There has been some debate as to whether IT managers should hold off on deploying the 40GbE technology and bide their time waiting for 100GbE to become

commercially available, but that question is fast becoming moot because 40GbE provides design flexibility and cost advantage over 100GbE. 40GbE can be effectively deployed today in aggregation links in data center networks. By 2016 40GbE will also be commonly applied to access links to connect servers as Figure 4 indicates. Complementing 40GbE, 100GbE is a perfect choice for carrier service providers and core links in data centers.

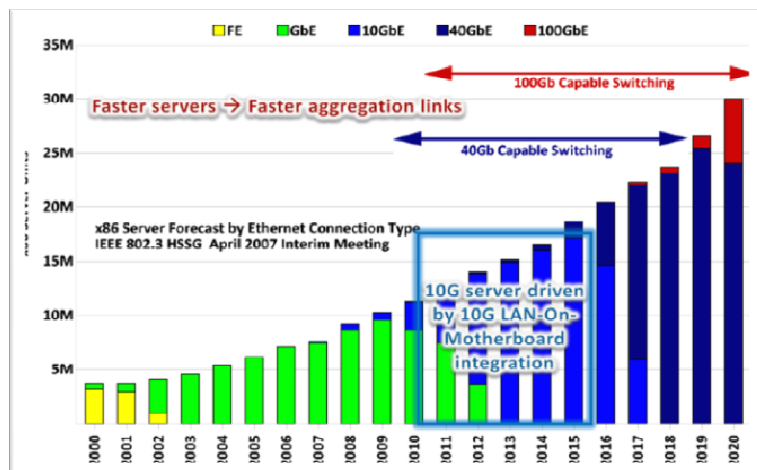


Figure 4. Projected timeline showing mainstream adoption of 40GbE-capable switching equipment.

For IT Managers and CIOs who are intent on remaining competitive, 40GbE represents an optimal option for adding the required bandwidth but without the proper planning, network operators may very well be caught flat-footed when the time comes to make the move.

This article explains the process to thoroughly evaluate, plan, and specify a network upgrade to 40GbE. As Figure 4 illustrates, 40GbE is still a couple of years from wide scale adoption, but IT Managers and CIOs should begin their migration planning now.

40GbE: A Closer Look

40GbE and 100GbE are Ethernet standards developed by the IEEE 802.3ba Task Force that support sending Ethernet frames at 40 and 100 gigabits per second. They also address physical layer specifications for communication across backplanes, copper cabling, multimode fiber and single-mode fiber. Official development of the 40GbE/100GbE standard began in January 2008 and was officially ratified in June 2010.

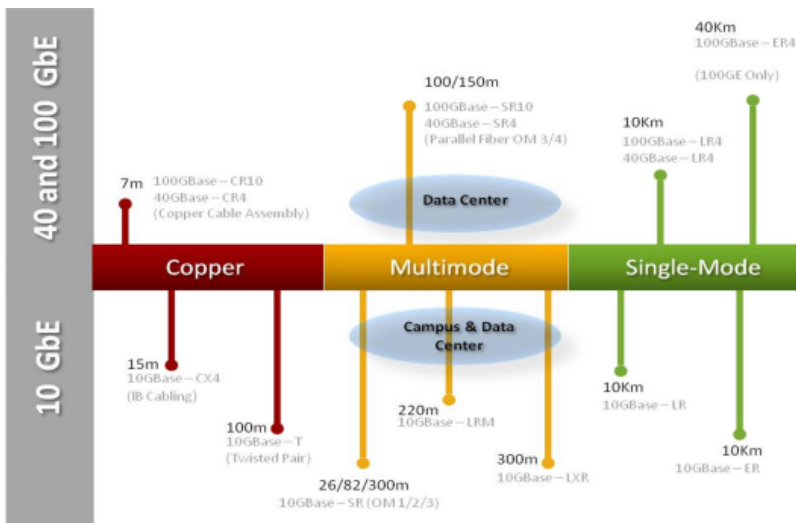
At the heart of the 40GbE network layer is a pair of transceivers connected by a cable, for example OM4 or OM3 fiber cable. The transceivers, in turn, are plugged into either network servers or a variety of components including interface cards and switches.

		Planned for 1 st Generation		Not Planned for 1 st Generation	
Media	Reach	Speed	CFP	QSFP	CXP
Single-mode	10Km	100G	Planned	Not Planned	Not Planned
		40G	Planned	Future?	Not Planned
	40Km	100G	Planned	Not Planned	Not Planned
Multimode (OM3)	100m	100G	Planned	Future?	Planned
		40G	Planned	Planned	Not Planned
Multimode (OM4)	150m	100G	Planned	Future?	Planned
		40G	Planned	Planned	Not Planned
Copper	3-7m	100G	Planned	Future?	Planned
		40G	Planned	Planned	Not Planned

Figure 5. Transceiver form factors planned for 1st generation implementation.

Transceivers: 40GbE transceivers (Figure 5) are being developed along several standard form factors. The CFP (C form-factor pluggable) transceiver features twelve transmit and twelve receive 10Gb/s lanes to support one 100GbE port, or up to three 40GbE ports. Its larger size is suitable for the needs of single-mode optics and can easily serve multimode optics or copper as well. The CXP transceiver form factor also provides twelve lanes in each direction but is much smaller than the CFP and serves the needs of multimode optics and copper. The QSFP (quad small-form-factor pluggable) is similar in size to the CXP and provides four transmit and four receive lanes to support 40GbE applications for multimode fiber and copper today and may serve single-mode in the future. Another future role for the QSFP may be to serve 100GE when lane rates increase to 25Gb/s.

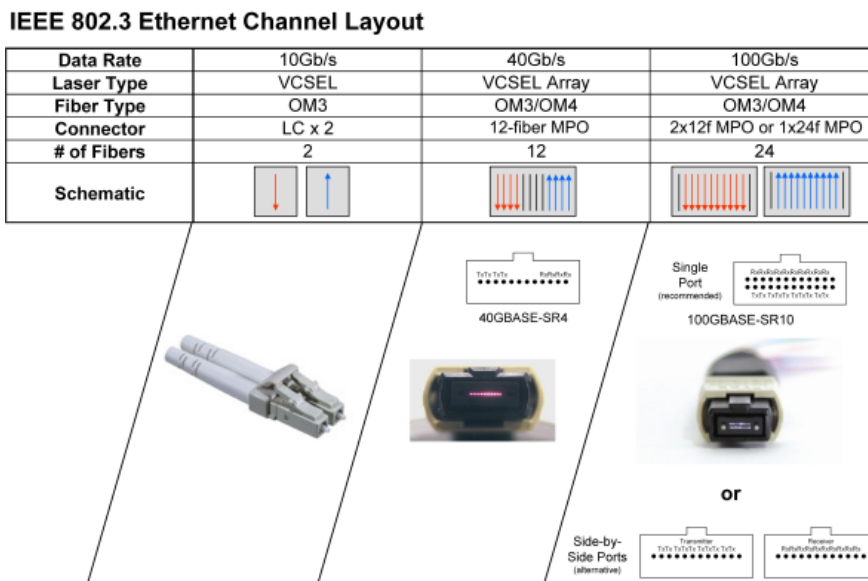
Figure 6. Cabling Alternatives 10GbE and 40/100GbE



Cables & Connectors: Cabling for 40GbE, summarized in Figure 6, can be optical fiber or copper. The supportable channel length depends on the cable and the transceiver type. With regards to connectors, the only significant change outlined in the 802.3ba standard is the use of MPO (Multi-Fiber Push On) type connectors at the multimode transceivers to support the multi-fiber parallel optics channels.

For data center environments operating at 40GbE or 100GbE, OM3 and OM4 multimode cabling is generally recommended because its reach supports a wider range of deployment configurations compared to copper solutions, and the cost is lower compared to single-mode solutions.

Figure 7. Layout showing Ethernet channel distribution for 10/40/100GbE using multimode fiber



Use of Parallel Optics

Traditionally, the Ethernet standard has relied upon duplex fiber cabling with each channel using one fiber to transmit and the other to receive. However, the 802.3ab standard requires multiple lanes of traffic per channel. To do this, the 40/100GbE standard uses parallel optics as indicated in Figure 7. The 40GbE specification calls for a 12-fiber cabling solution with each channel featuring four dedicated transmit fibers and four dedicated receive fibers. The middle four fibers remain unused, or dark. The 100GbE solution specifies 24 fibers divided into two 12-fiber arrays with one array dedicated to transmit and the other dedicated to receive. In each array, the middle ten fibers are dedicated to traffic while the two fibers on either end remain dark.

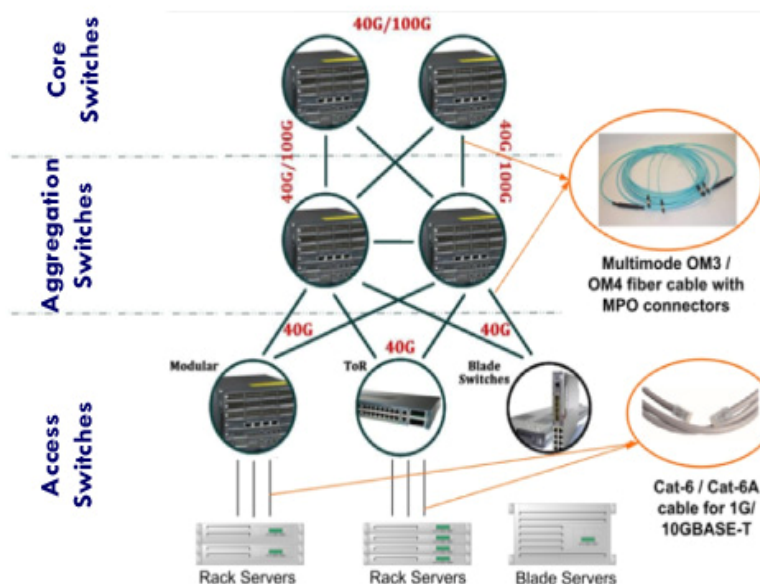
For 100GbE several interface variants have been described with the preferred option being a single 24-fiber MPO connector. Alternately, two 12-fiber connectors can be positioned either vertically or side by side to make up the channel.

Deploying 40GbE in the Network

The new 40GbE technology will, most likely, first be deployed in the data center as shown in Figure 8. This will help alleviate bottlenecks in the layer that connects access switches to aggregation switches.

“With enterprises beginning to deploy 10GbE uplinks on client side switches in the wiring closet, as well as 10GbE direct server connections, the aggregation of these links is becoming a bottleneck in the network until higher speeds are widely available,” said Alan Weckel, director, Ethernet Switch & Enterprise Telephony Market Research at Dell’Oro Group.³

Figure 8. 40/100GbE deployment throughout the data center network.



“For example, at present enterprises must deploy complicated fat tree or spanning tree architectures to aggregate 10GbE using the same speed as both the aggregator and the aggregated. We believe that 40GbE and 100GbE will be critical to meeting the increasing demands for bandwidth in the data center,” Weckel explained.²

What’s Involved?



One of the most attractive characteristics of 40GbE is broad applications and design flexibility. Considering the productivity gains and decrease in OPEX, migrating to 40GbE will prove very cost-effective for those who do it right.

In migrating to 40GbE, some networks will be able to use their current 10GbE switch chassis and just upgrade their line cards and transceivers. Deploying a CFP form factor transceiver will provide the flexibility to migrate from 40GbE to 100GbE.

With regards to cabling, OM3 or OM4 is optimal for the 40GbE or 100GbE data center environment. The major difference is in the maximum span distances. In a 10GbE network, OM3 fiber can span up to 300m while OM4 supports even longer channels. In a 40GbE or 100GbE environment, OM3 can be used up to 100m and OM4 up to 150m according to the IEEE802.3ba standard. For applications approaching 150m, the cable should be terminated with low loss connectors.

There are a number of alternative upgrade paths leading from 10GbE to 40GbE. To make an optimal decision for your network here are some key issues to consider:

What factors determine when it is more effective to deploy a 40GbE physical layer compared to aggregations of 10GbE channels?

What installed hardware or cabling, if any, will need to be replaced or reconfigured?

What are the capabilities of various 40GbE transmission alternatives?

Don't Wait to Start Future-Proofing Your Network

The risk in delaying migration planning is the potential to underestimate what is involved in evaluating and selecting the best migration path. While it is true that upgrading from 10GbE to 40GbE should be a relatively smooth process, today's network manager must be thinking a step beyond. That means considering not only how best to transform today's legacy systems to a 40GbE environment, but what implications that transformation will have on the eventual migration to 100GbE and beyond. By thoroughly considering these issues now and developing an implementation plan, those in IT can ensure a future-proof evolution of their networks.

It helps to remember that the pre-planning process need not involve any immediate purchases. With the 802.3ba specifications now in place, networks can safely plan without fearing that the landscape will change dramatically.



About Ethernet Alliance

The Ethernet Alliance is a community of Ethernet end users, system and component vendors, industry experts and university and government professionals who are committed to the continued success and expansion of Ethernet. The Ethernet Alliance brings Ethernet standards to life by supporting activities that span from incubation of new Ethernet technologies to interoperability demonstrations, certification and education.

References

¹ Worldwide Quarterly Server Virtualization Tracker, International Data Corporation, April 2010

² InformationWeek, January 29th, 2009, reported by W. David Gardner

³ Ethernet Alliance[®] Congratulates IEEE on the Ratification of 40 and 100 Gigabit Ethernet Standard; Announces Demonstration and Interoperability Plans, Press Release, June 21, 2010