

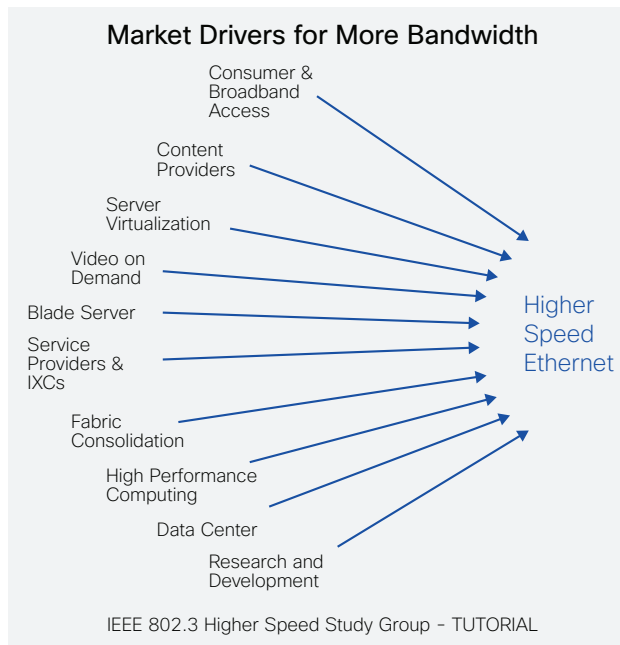
The Future Is 40 Gigabit Ethernet



Executive Summary

The business case for 40 Gigabit Ethernet is becoming inescapably compelling. While 10 Gigabit Ethernet 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: high-bandwidth applications such as server virtualization and cloud computing; fabric consolidation within the data center; and a greater demand for high-performance computing among end users (see Figure 1). The need for faster data transfer rates is relentless and carries significant implications with regard to network productivity as well as operating expenditure (OpEx) costs.

Figure 1. Current Trends Driving the Demand for Higher-Speed Ethernet



So when the IEEE officially adopted IEEE Std. 802.3ba in June 2015, paving the way for both 40 Gigabit Ethernet and 100 Gigabit Ethernet, it came not a moment too soon. The increased speed allows networks to move newfound 10 Gigabit Ethernet resources to the access layer, allowing the more powerful 40 Gigabit Ethernet-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** 40 Gigabit Ethernet will become an accepted part of the IT landscape, but **when** and **how**.

This report addresses the impending move to 40 Gigabit Ethernet, 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 40 Gigabit Ethernet

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. In 2016, eight years after the adoption of 10 Gigabit Ethernet, the IEEE has adopted 802.3ba, paving the way for 40 Gigabit Ethernet and 100 Gigabit Ethernet.

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 40 Gigabit Ethernet is a new generation of high-speed, high-demand, computing applications and technologies. These include the spreading deployment of virtual servers and cloud computing.

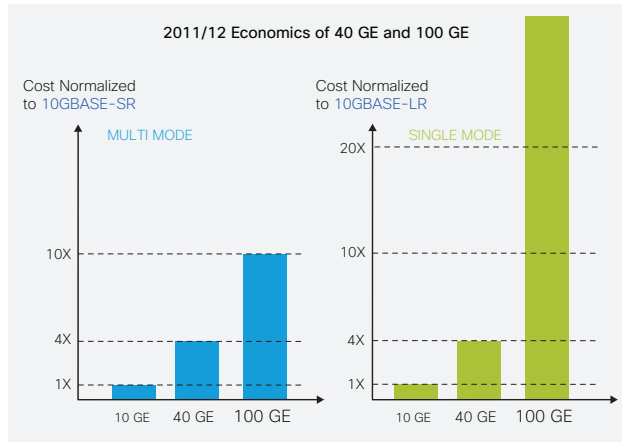
By the end of 2009, nearly one in five new servers were virtualized.¹ At the same time, the financial pressures of a challenging economy are forcing networks to look for ways to consolidate their resources in an effort to reduce OpEx and total cost of ownership.

With four times the capacity and the ability to cost effectively migrate to 100 Gigabit Ethernet (see Figure 2), 40 Gigabit Ethernet on multimode fiber is the next logical step in the evolution of the data network.

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

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 2013.²

Figure 2. Gigabit Ethernet for Multimode and Single-Mode



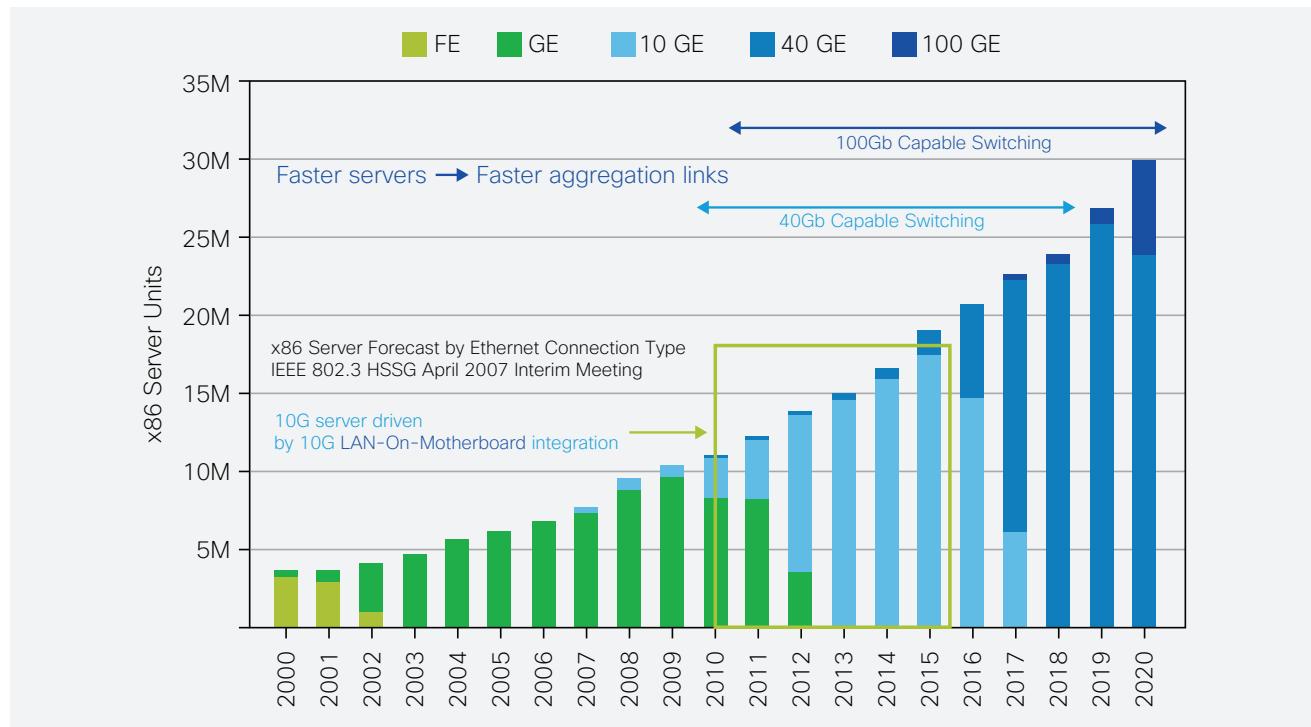
There has been some debate as to whether IT managers should hold off on deploying the 40 Gigabit Ethernet technology and bide their time waiting for 100 Gigabit Ethernet to become commercially available.

But that question is fast becoming moot because 40 Gigabit Ethernet provides design flexibility and cost advantage over 100 Gigabit Ethernet. 40 Gigabit Ethernet can be effectively deployed today in aggregation links in data center networks. By 2016, 40 Gigabit Ethernet will also be commonly applied to access links to connect servers, as Figure 3 indicates. Complementing 40 Gigabit Ethernet, 100 Gigabit Ethernet is a perfect choice for carrier service providers and core links in data centers.

For IT managers and CIOs who are intent on remaining competitive, 40 Gigabit Ethernet represents the best option for adding the required bandwidth. But without the proper planning, network operators may very well be caught unprepared when the time comes to make the move.

Let’s examine the process to thoroughly evaluate, plan, and specify a network upgrade to 40 Gigabit Ethernet. As Figure 4 illustrates, 40 Gigabit Ethernet is still a couple of years from wide-scale adoption, which gives IT managers and CIOs time to start their migration planning.

Figure 3. Projected Timeline Showing Mainstream Adoption of 40 Gigabit Ethernet-Capable Switching Equipment



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

Figure 4. Transceiver Form Factors Planned for First-Generation Implementation

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

Planned for 1st Generation
 Not Planned for 1st Generation

40 Gigabit Ethernet: A Closer Look

40 Gigabit Ethernet and 100 Gigabit Ethernet are Ethernet standards developed by the IEEE 802.3ba Task Force to 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 40 Gigabit Ethernet and 100 Gigabit Ethernet standards began in January 2008, and the standards were officially ratified in June 2010. At the heart of the 40 Gigabit Ethernet 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.

Transceivers

40 Gigabit Ethernet transceivers (see Figure 5) are being developed along several standard form factors. The C Form-Factor Pluggable (CFP) transceiver features 12 transmit and 12 receive 10-Gbps lanes to support one 100 Gigabit Ethernet port, or up to three 40 Gigabit Ethernet ports. Its larger size is suitable for the needs of single mode optics and can easily serve multimode optics or copper. The CXP transceiver form factor also provides 12 lanes in each direction, but is much smaller than the CFP and serves the needs of multimode optics and copper. The Quad Small-Form-Factor Pluggable (QSFP) is similar in size to the CXP. It provides four transmit and four receive lanes to support 40 Gigabit Ethernet applications for multimode fiber and copper today, and may serve

Figure 5. Transceiver Form Factors Planned for First-Generation Implementation

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

Planned for 1st Generation
 Not Planned for 1st Generation

single-mode in the future. Another future role for QSFP may be to serve 100 Gigabit Ethernet when lane rates increase to 25 Gbps.

Cables and Connectors

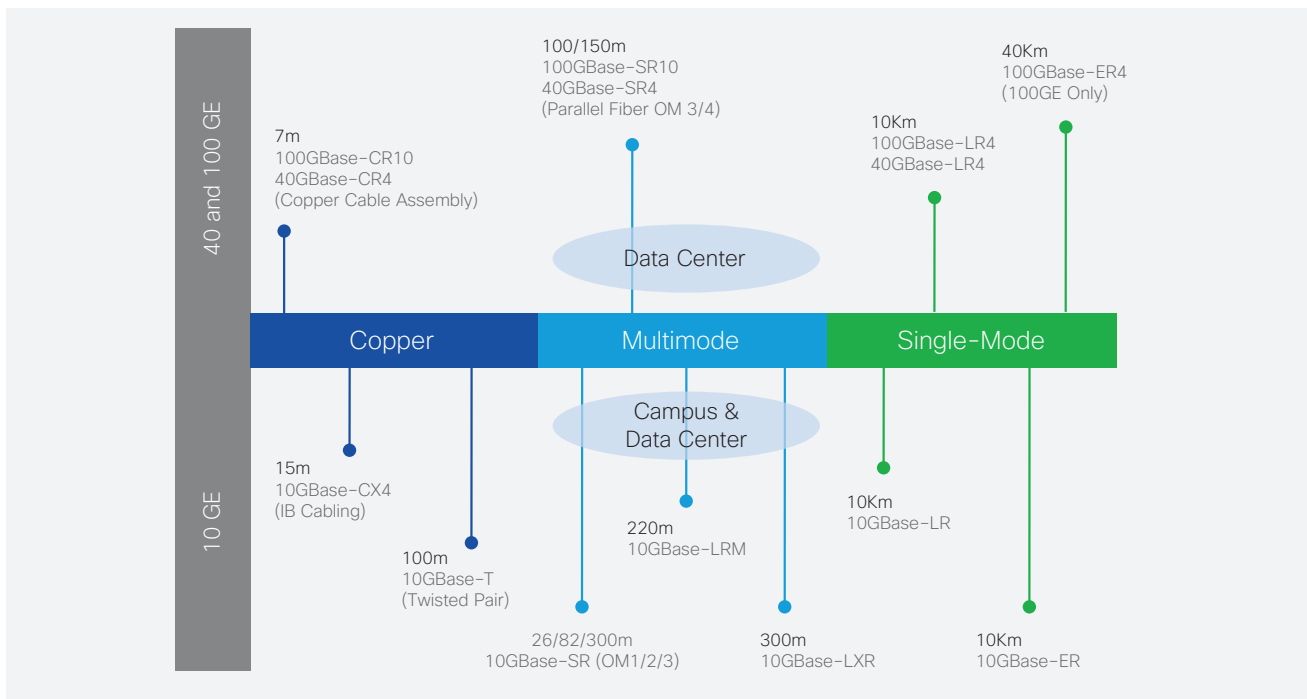
Cabling for 40 Gigabit Ethernet, summarized in Figure 6, can be optical fiber or copper. The supportable channel length depends on the cable and the transceiver type. With regard to connectors, the only significant change outlined in the 802.3ba standard is the use of MPO Multi-Fiber Push On (MPO)-type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

For data center environments operating at 40 Gbps or 100 Gbps, 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.

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

Figure 6. Cabling Alternatives for 10 Gigabit Ethernet and 40/100 Gigabit Ethernet



traffic per channel. To get multiple lanes, the 40 and 100 Gigabit Ethernet standard uses parallel optics, as indicated in Figure 7.

The 40 Gigabit Ethernet 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 100 Gigabit Ethernet 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 10 fibers are dedicated to traffic while the two fibers on either end remain dark.

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

Deploying 40 Gigabit Ethernet in the Network

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

Figure 7. Layout Showing Ethernet Channel Distribution for 10/40/100 Gigabit Ethernet Using Multimode Fiber


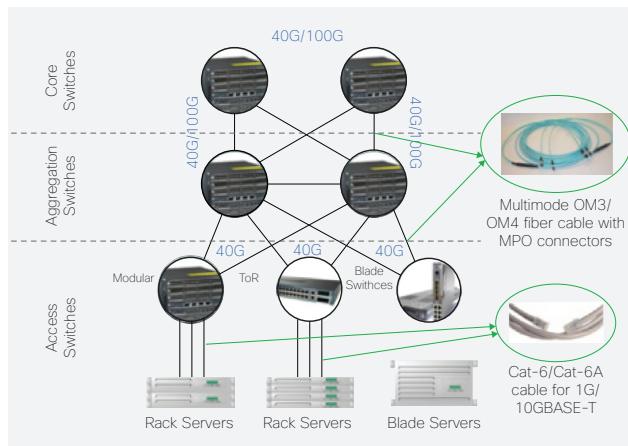


Figure 8. 40/100 Gigabit Ethernet Deployment Throughout the Data Center Network



“With enterprises beginning to deploy 10 Gigabit Ethernet uplinks on client side switches in the wiring closet, as well as 10 Gigabit Ethernet 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

“For example, at present, enterprises must deploy complicated fat tree or spanning tree architectures to aggregate 10 Gigabit Ethernet using the same speed as both the aggregator and the aggregated. We believe that 40 Gigabit Ethernet and 100 Gigabit Ethernet 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 40 Gigabit Ethernet is broad applications and design flexibility. Considering the productivity gains and decrease in OpEx, migrating to 40 Gigabit Ethernet will prove very cost effective for those who do it correctly.

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

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

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

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

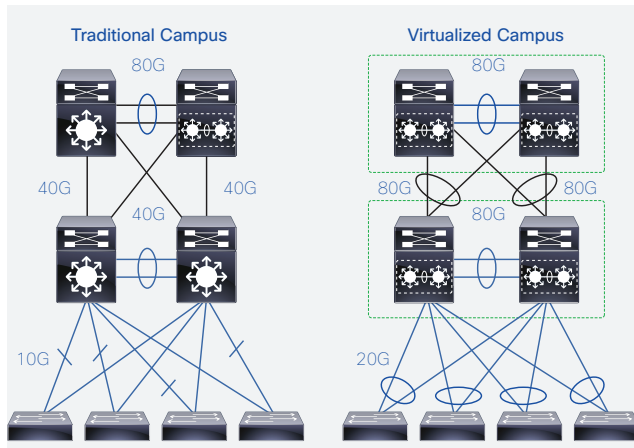
- What factors determine when it is more effective to deploy a 40 Gigabit Ethernet physical layer compared to aggregations of 10 Gigabit Ethernet channels?
- What installed hardware or cabling, if any, will need to be replaced or reconfigured?
- What are the capabilities of various 40 Gigabit Ethernet transmission alternatives?

Uses Cases

Campus Aggregation and Core

As links to access switches move to 10 Gigabit Ethernet, customers need interfaces greater than 10 Gigabit to link the aggregation and core layers together (Figure 9).


Figure 9. Campus Deployment of a 40 Gigabit Ethernet Line Card



Enterprise and Service Delivery Data Center: Aggregation and Core

For switch uplinks supporting higher bandwidths as end hosts move toward 10 Gigabit Ethernet. This is the principal market for 40 Gigabit Ethernet (Figure 10).

Figure 10. Data Center Deployment of a 40 Gigabit Ethernet Line Card



Conclusion

Don't wait to start preparing your network for future technology.

The risk in delaying migration planning is that you may underestimate what is involved in evaluating and selecting the best migration path. While it is true that upgrading from 10 Gigabit Ethernet to 40 Gigabit Ethernet 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 40 Gigabit Ethernet environment, but what implications that transformation will have on the eventual migration to 100 Gigabit Ethernet and beyond. By thoroughly considering these issues now and developing an implementation plan, IT can help ensure a smooth evolution of their networks.

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