



ethernet alliance

Overview of Requirements and Applications for 40 Gigabit and 100 Gigabit Ethernet

Version 1.0 August 2007

Authors:

Mark Nowell, Cisco
Vijay Vusirikala, Infinera
Robert Hays, Intel



Executive Summary

Innovative and emerging server technologies including multi-core processing, virtualization, networked storage, and I/O convergence are contributing to a growing bandwidth requirement in the enterprise computing environment. Aggregating today's 10 Gigabit Ethernet (10GbE) compute devices and the future higher speed 40 Gigabit Ethernet (40GbE) interfaces in the data center requires 100 Gigabit Ethernet (100GbE) interfaces for the switch-to-switch interconnection. Similar aggregation effects are present in the carrier and service provider networks, where it has become increasingly challenging for service providers to accommodate customer requests for services due to the bandwidth constraints of their network core. 100GbE is the identified solution for providing the next generation of internet connectivity to continue to fuel the delivery of new services and content to the consumer and business customers.

Introduction

The Evolution of Ethernet Networks

Ethernet has evolved well beyond connecting PCs in an enterprise office. For over a quarter of a century, Ethernet has continued to expand its footprint and has demonstrated itself as a reliable and valuable networking standard. Evolving from the early days of a shared 10 Mb/s bus, Ethernet now operates 1000 times faster. Its deployment is widespread, covering enterprise LANs, broadband access, data center networking, and communications across metropolitan and wide area networks. In addition, Ethernet continues its expansion as the preferred carrier wire line vehicle by permitting wireless technologies, such as WiFi and WiMAX, to bridge into Ethernet networks.

Ethernet has become the networking communications technology of choice for so many applications largely because of the availability of cost-effective, reliable and interoperable networking products from a broad selection of vendors. Vendors' participation in and adherence to IEEE 802.3 standards has perpetuated the quality and value of Ethernet and led to ever-increasing interest in deploying Ethernet further into the end-to-end network.

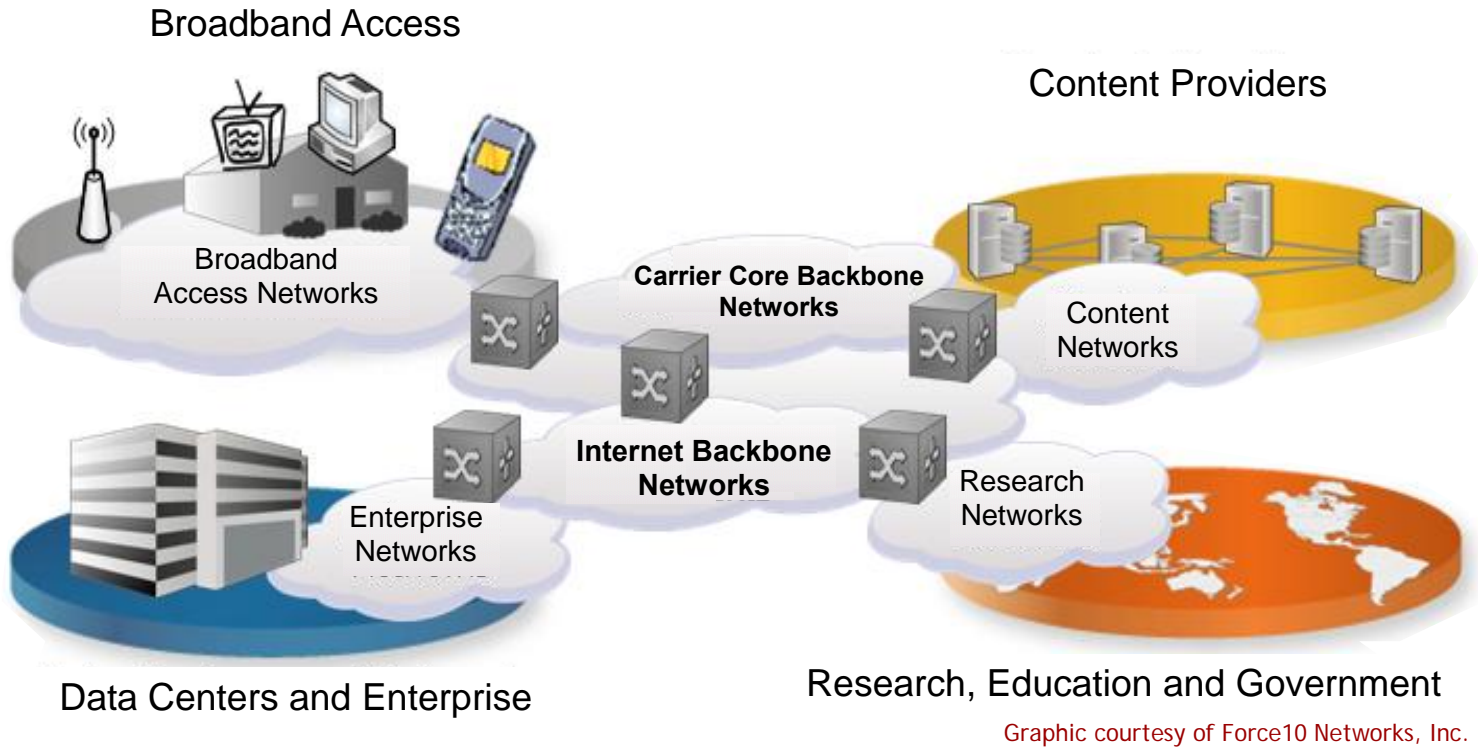


Figure 1— Ethernet’s Network Expansion

Network and Application Drivers for Higher Speed Ethernet

According to IDC, network equipment shipments continue to grow at an annual average rate of 17%. This growth is driven by the 75-125% bandwidth growth in Internet usage, electronic commerce, IPTV, VoIP, wireless communications, video-on-demand, online collaboration, etc. At some of the critical Internet aggregation points, as many as eight ‘lanes’ of 10 Gigabit Ethernet have been aggregated to handle the current bandwidth demand. Shortly, Ethernet at a rate of 10 gigabits per second will not be sufficient to cope with the bandwidth demands at these aggregation junctions.

In July 2006, the IEEE 802.3 working group formed a Higher Speed Study Group (HSSG) to investigate the needs and requirements for the next speed(s) of Ethernet technology. The HSSG brought together industry experts to analyze the trends and determine the scope of a project to specify the next speed or speeds for Ethernet. The goal was to try to determine what speed the market would need in the next 3 to 7 years.



Enterprise Computing and Data Center Networking

Market Trends

E-commerce, entertainment content distribution, high-performance clusters, Internet communications, etc. are increasing the usage of data center computing resources. To meet the demand, manufacturers are delivering emerging technologies including multi-core processing, virtualization, networked storage, and I/O convergence all which contribute to better performance, cost and power efficiency. These trends impact the future of Ethernet in the data center.

Multi-core Servers & Virtualization Trend

Like faster clock speeds in the past, multi-core processors are the main impetus providing today's performance gains in next generation servers. Multi-core processors are particularly valuable for multithreaded computing applications used in high-performance compute (HPC) clusters where many servers are networked together to execute large computing tasks through parallel processing. Data center managers are also taking advantage of the increased processing capabilities to reduce costs by consolidating workloads onto fewer, more highly utilized servers using virtualization technologies. The combination of improved total processing resources and increased usage of clustering and virtualization are driving the networking I/O bandwidth requirements beyond multi-GbE today, and 10GbE is starting to become an attractive option for servers. With performance on a Moore's Law curve and therefore doubling approximately every 24 months, 40GbE will be a logical next speed for servers in four years.

Networked Storage Trend

Networked storage is replacing the local hard disks traditionally found in servers with centralized storage shared by multiple servers over a network. Some benefits of networked storage are the following: broader data access, better reliability, lower maintenance costs, and improved disk utilization. Network storage solutions utilizing Ethernet include most Network Attached Storage (NAS) devices, iSCSI, and the proposed Fibre Channel over Ethernet (FCoE) standard. Disk I/O is one of the primary bandwidth consumers in servers and moving the disks out of the local chassis to a remote network drive increases the network I/O bandwidth requirements.



Given that the various types of server I/O - memory, storage & networking - are often limited by CPU performance capabilities which grows at a Moore's Law rate, storage bandwidth requirements tend to follow the same performance doubling seen in the server's network interface. Given this performance trend coupled with the increasing use of networked storage, 40GbE is anticipated to meet the upcoming networked storage bandwidth requirements of the data center.

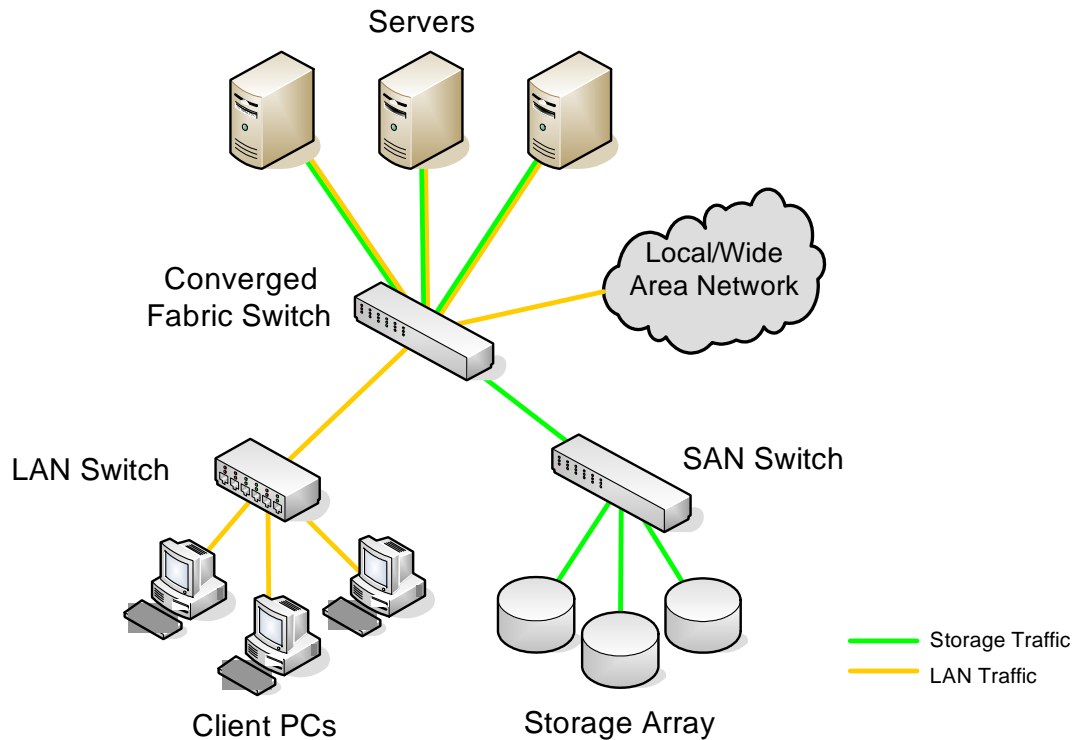


Figure 2— Network Storage and I/O Convergence

I/O Convergence Trend

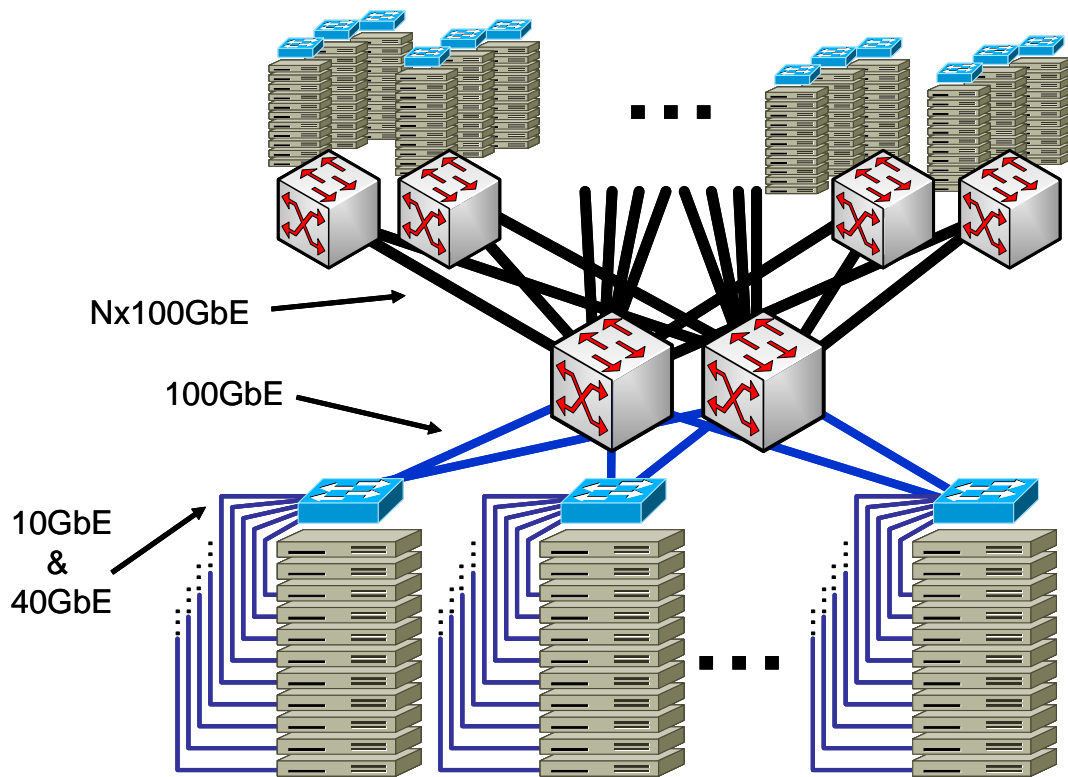
Today, many networked storage deployments connect servers to separate storage area networks and local area networks, duplicating much of the hardware infrastructure including adapter cards in the servers and switches in the data center. As 10 Gigabit Ethernet becomes more widely available, servers may be able to utilize a single converged Ethernet I/O interface. This single converged interface would carry servers' networking and storage I/O traffic over a common Ethernet connection. I/O convergence increases the demand on the converged I/O network connection, driving the need for future improvements in Ethernet beyond 10GbE. This bandwidth growth is similar to the growth for multi-core servers; and given the data centers requirement for solutions that are cost effective and power efficient, 40GbE



is likely the preferred next rate over higher speed alternatives such as 100GbE for these converged interfaces.

Data Center Network Aggregation Trend

10 Gigabit Ethernet deployments continue to grow as the compute bandwidth continues to grow and costs continue to decline due to the availability of lower cost 10GbE optical and copper PHY technologies. As the deployment of 10GbE on servers increases, there is a need for higher-speed switch uplinks for network aggregation in the data center. Several of the world's largest data center operators clearly articulated this requirement and the grim consequences on the ability of their infrastructure to serve the demanded content if these switch-to-switch aggregation capacities were not increased by at least an order of magnitude. A 100GbE rate was proposed to provide the bandwidth required to handle the increased traffic load. It is expected that 100GbE will be deployed in switch uplinks inside the data center as well as providing inter-building, inter-campus, MAN, and WAN connections for enterprise networks.



Graphic courtesy of Google, Inc.

Figure 3 – Data Center Network Aggregation



Carrier & Service Provider Networking

Market Trends

Traffic demand in the carrier backbone continues to grow rapidly (ranging from 75% to 125% per year¹), driven by a number of popular applications such as IPTV, video-on-demand services, remote storage, IP data transit, mobile broadband services, and VPN services, all facilitated by ultra-broadband access networks. Service providers ranging from wholesale carriers, cable multiple system operators (MSOs) to Internet service providers (ISP) and content delivery network (CDN) providers are forecasting a continuation of this sustained traffic growth. To meet these high-bandwidth needs, the client connections to the carrier network are rapidly migrating to 10GbE. Backbone network links typically require at least 4-10 times the bandwidth of the highest-speed user service interface to provide adequate performance. With the migration to 10GbE today, there is a need for a higher speed standard as soon as possible.

Aggregation of the Access - The power of large numbers

As the Internet continues to influence numerous aspects of our daily lives, there is a constant, unabated growth in the bandwidth requirements placed on the aggregation of the diverse access networks. The number of residential users connected to the web is constantly increasing, as well as the bandwidth being demanded by these users as they consume and generate their own content. Adding to this growth is the development and availability of additional methods of accessing the Internet such as WiFi or WiMax networking and 3G mobile services. As people continue to demand more bandwidth and as the number of people online continues to grow, this multiplicative effect is causing serious challenges on how to adequately handle the aggregate traffic in the network.

This continued progression of capabilities and bandwidth demands of access applications will soon be constrained unless there is an adequate ability to accommodate these aggregate loads in the network. When examining aggregation solutions, it is necessary to enable a sufficient bandwidth increase to address the capacity constraints while avoiding specifying solutions which are not economically feasible or

1. G. Coffman and A. M. Odlyzko, 'Growth of the Internet', Optical Fiber Telecommunications IV B: Systems and Impairments, I. P. Kaminow and T. Li, eds. Academic Press, 2002, pp. 17-56.



cause carriers to prematurely replace their infrastructure equipment. Ethernet's order of magnitude increase has been proven numerous times to provide an appropriate balance between cost and performance for these critical aggregation applications.

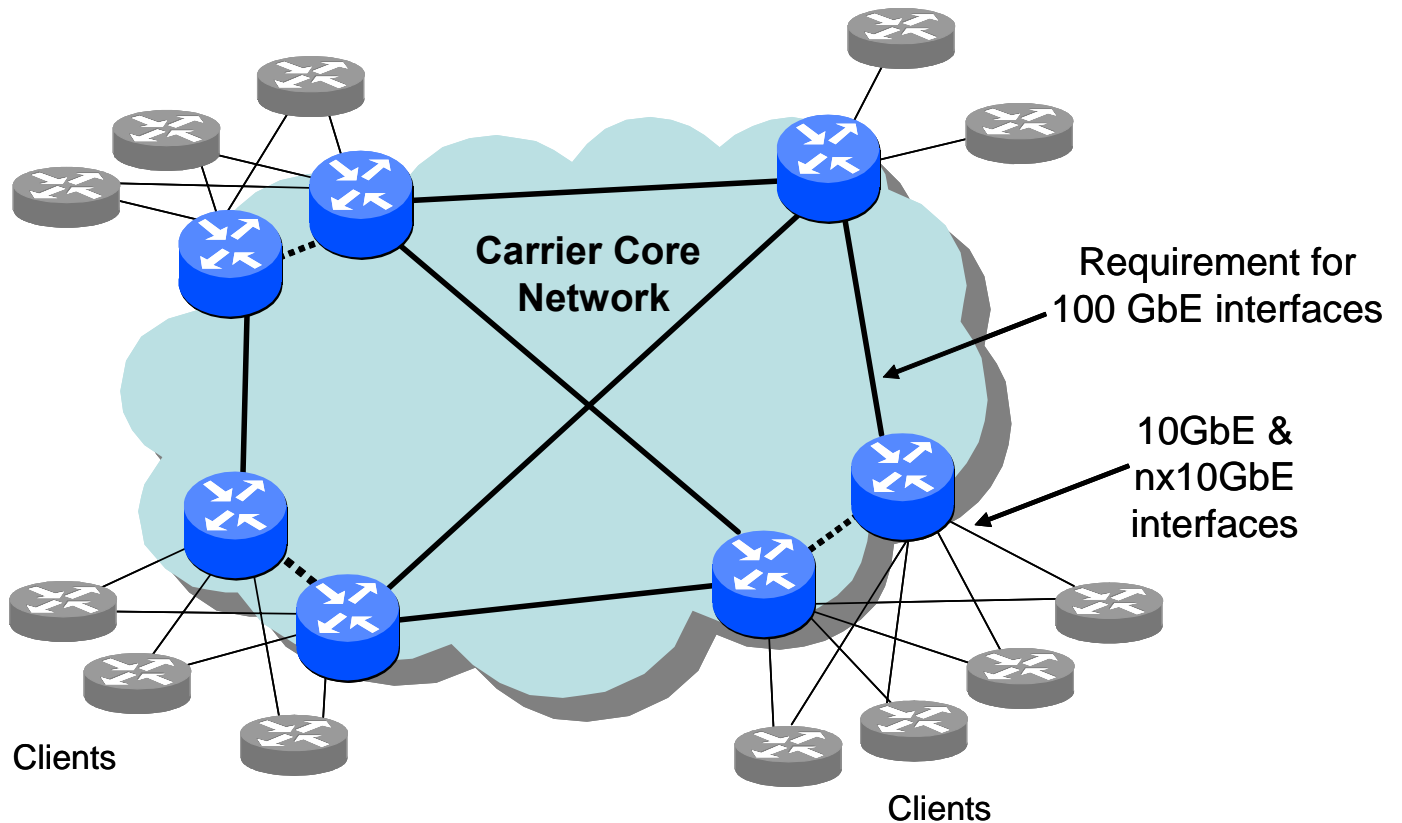


Figure 4 – Carrier/Service Provider Core Network

Ethernet in Carrier Networks Trend

While the vast majority of core router interfaces are currently at 10G (10G Packet-over-SONET - POS - or 10GbE), 40G OC-768/STM-256 POS interface deployment is steadily growing due to its availability as a solution to today's aggregation bandwidth constraints. The increasing trend of Ethernet being an end-to-end network solution and the understanding that a 40G interface will not be sufficient to address the aggregation demand has resulted in the 100GbE rate being identified as the next high-speed network operator interface. The core network needs to evolve to a higher rate, and operators are looking to Ethernet as the new 'fat pipe' for the core of their backbone; positioning 100GbE to complete the transition to an end-to-end all-Ethernet infrastructure.



Direction for Higher-Speed Ethernet

Application Requirements

An analysis of the requirements for enterprise computing and network aggregation highlighted the 'one size fits all' approach would not work as it had in the past. A key finding by the HSSG was the divergence in bandwidth growth rates between the network core (network aggregation) and the server network (computing I/O) as illustrated in the figure below.

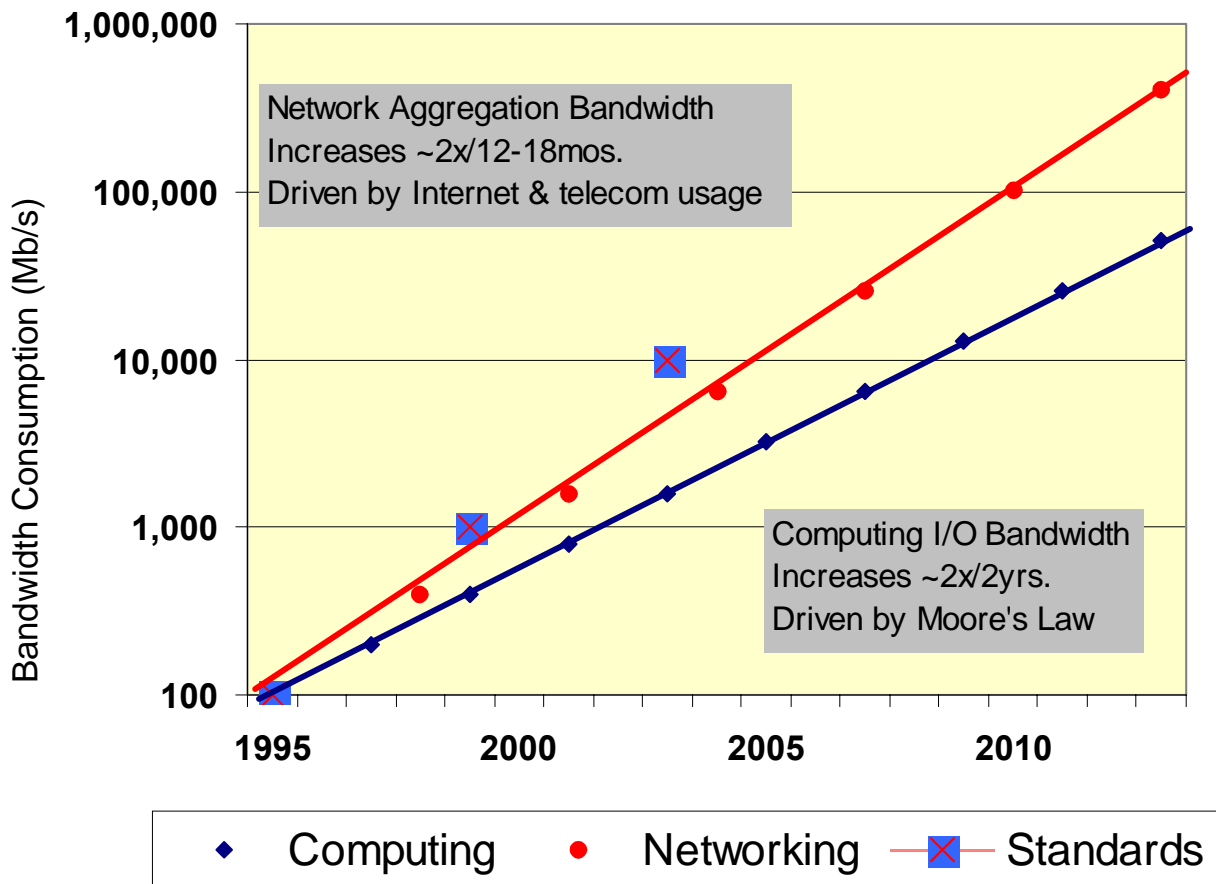


Figure 5 – Ethernet Bandwidth Demand Trends



The server network bandwidth is influenced by CPU, host bus and memory performance which typically follow Moore’s Law. On the other hand, network aggregation bandwidth requirements are influenced by server performance as well as the growth in users and access points, and higher data rates needed by rich media content. Statistics from the past 5 years show Internet backbone bandwidth doubling every 12-18 months and the computing bandwidth I/O doubling approximately every 24 months. The effect is resulting in higher bandwidth growth in the network aggregation core than in the server computing network.

In recognition of these underlying bandwidth drivers, the HSSG recommended developing a specification that includes both 40 Gb/s and 100 Gb/s rates to address the distinct markets’ needs. The computing and server market will benefit from a smaller bandwidth step which is more in line with the key driving technology of projected CPU bandwidth I/O. The capacity strain felt in core networking and data center aggregation applications will benefit from the traditional 10x bandwidth step which has been a proven and an effective step in matching the cost/performance requirements of the aggregation equipment.

There are also different physical interface (PHY) requirements for each application. The 40 Gigabit Ethernet (40GbE) rate will include a family of physical layer solutions to cover distances inside the data center up to 100m for inclusion in a full range of server form factors including blade, rack, and pedestal. The 100 Gigabit Ethernet (100GbE) rate will include distances and media appropriate for data center networking as well as service provider inter-connection for intra-office and inter-office applications. The proposed PHYs are shown below in Table 1.

	40GbE	100GbE
1m backplane	√	
10m copper	√	√
100m MMF	√	√
10km SMF		√
40km SMF		√

Table 1 – Proposed HSSG PHYs



The HSSG's recommendation to specify two MAC rates, each with the appropriate physical layer interfaces, comprehensively covers the upcoming needs of the enterprise computing, data center aggregation and carrier/service provider core network.

Conclusion

The IEEE 802.3 working group recently approved the project authorization request (PAR) for a higher speed Ethernet amendment to IEEE 802.3 standard. The working group approved the PAR forwarded to it by the HSSG. The next step is for the PAR to be approved by the IEEE 802 executive committee, the IEEE Standards Association (IEEE-SA), New Standards Committee (NesCom), and the IEEE-SA Standards Board. Once the PAR has been approved by those bodies, the HSSG will cease to exist and a new IEEE 802.3 standards project task force will be formed. The task force will begin to create a draft standard for these new rates and physical layers. While no definitive standards development timeline has been agreed to, it is anticipated that the 40GbE and 100GbE standard may be completed in 2010.

References

For further information on the IEEE 802.3 standards processes, visit: <http://www.ieee802.org/>