



HIGH- PERFORMANCE NETWORKS

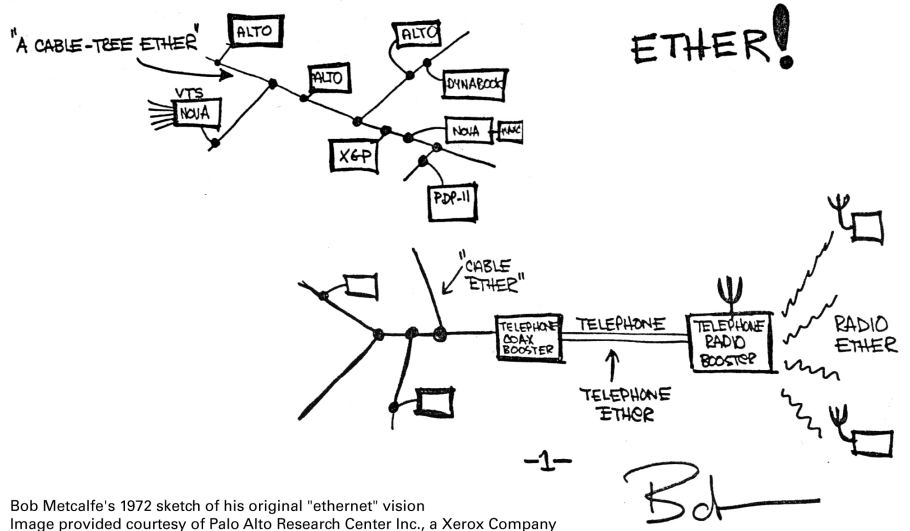
40 Gigabit and 100 Gigabit Ethernet Are Here!

The 40 Gigabit and 100 Gigabit Ethernet standards were adopted by the Institute of Electrical and Electronics Engineers (IEEE) 802.3 Working Group on June 17, 2010. These two revolutionary new Ethernet speeds open the door to significantly higher-capacity networks and enable networks to scale in ways that were previously impossible using 10 Gigabit Ethernet. This white paper presents an overview of the 40 Gigabit and 100 Gigabit Ethernet technologies and applications in service provider, data center, and enterprise networks.

INTRODUCTION

Since the conception of Ethernet by Bob Metcalfe and the team of engineers at XEROX PARC in the 1970s (see Figure 1), Ethernet technology has continued to evolve to meet the increasing bandwidth, media diversity, cost, and reliability demands of today's Internet. The latest Ethernet evolution increases its capacity from 10 Gigabit Ethernet (GbE) to 40 Gigabit and 100 Gigabit Ethernet. Ethernet and IP have become the underlying technologies that enable ubiquitous network access over a variety of wired and wireless media. From high-speed home broadband to wireless networks in remote villages. Internet traffic starts, transits, or ends on an Ethernet interface.

Figure 1.
"Here is some rough stuff on the Alto Aloha network," from a memo sent by Bob Metcalfe on May 22, 1973.

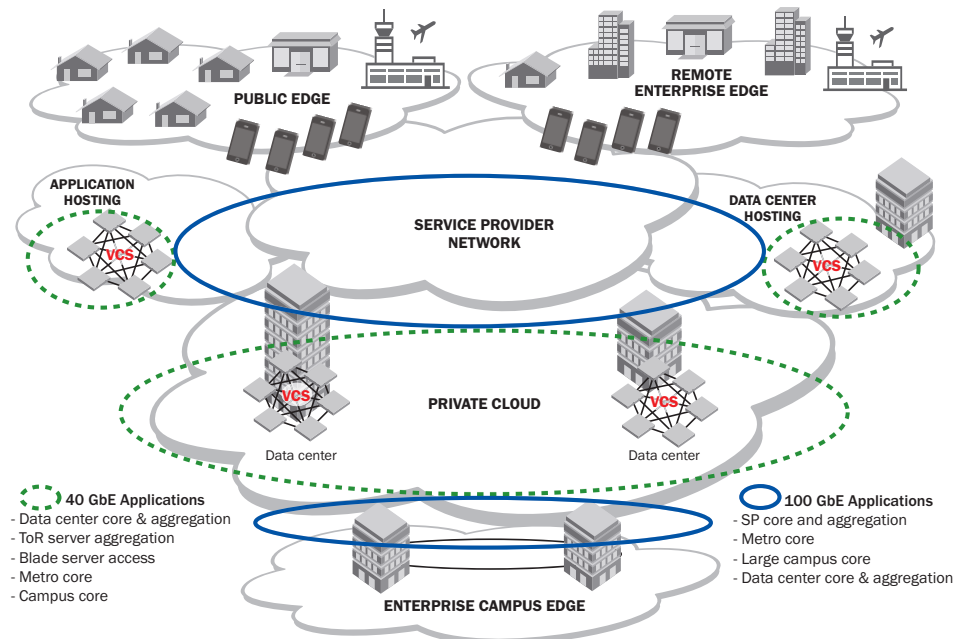


Bob Metcalfe's 1972 sketch of his original "ethernet" vision
Image provided courtesy of Palo Alto Research Center Inc., a Xerox Company

APPLICATIONS

40 Gigabit and 100 Gigabit Ethernet were designed with specific target applications and markets in mind, but these applications are not absolutes and the flexibility of the new speeds and interfaces allow each technology to be used in a variety of different applications. Capacity and redundancy are key requirements in all applications, and resiliency features such as Link Aggregation (LAG), Equal Cost Multi-Path (ECMP) routing and Multi-Protocol Label Switching (MPLS) are key factors to consider in these new applications.

Figure 2.
40 Gigabit (dashed) and 100 Gigabit (solid) Ethernet target applications.



40 Gigabit Ethernet Applications

Innovative data center technologies including virtualization, converged data center networks and cloud computing are changing traditional data center network designs. Modern data center networks now require much higher bandwidth capacities to meet the performance and flexibility requirements demanded by new applications. Initially 40 Gigabit Ethernet applications are targeted for short-reach data center core and aggregation layers or Top-of-Rack (ToR) server aggregation with copper cable or Multi-Mode Fiber (MMF) up to 125 m. Longer-reach interfaces for inter-data center, metro, and campus core networks with distances up to 10 km over Single-Mode Fiber (SMF) are also defined in the standard. 40 Gigabit Ethernet is considered the next logical speed for blade server access and server Network Interface Cards (NICs). As processor performance continues to follow Moore's Law and doubles about every 24 months, it is expected that servers will require 40 Gigabit Ethernet interfaces in the next few years.

100 Gigabit Ethernet Applications

100 Gigabit Ethernet applications, on the other hand, are targeted for SP core and aggregation networks, metro core, and large campus core networks. Driven by ubiquitous wired and mobile needs for both business and consumer access to high-bandwidth personalized content, these applications have the highest bandwidth and redundancy requirements for ultra-high-capacity Ethernet networks, and support distances up to 40 km over SMF. For data center core and aggregation applications, short-reach 100 Gigabit Ethernet technologies over copper cable and up to 125 m MMF are defined.

BACKGROUND

In 2006, the Institute of Electrical and Electronics Engineers (IEEE) 802.3 Working Group formed the Higher Speed Study Group (HSSG) to investigate the need for an Ethernet technology faster than 10 Gigabit Ethernet. The HSSG concluded that the capacity requirements driven by the ubiquity of the Internet and high-bandwidth content was rapidly outgrowing the capacities of networks that could be built with 10 Gigabit Ethernet. Two new Ethernet rates were proposed: 40 Gigabit Ethernet for server access and data center applications and 100 Gigabit Ethernet for network aggregation and backbone applications. In January 2008, the IEEE P802.3ba Task Force was formed to develop the 40 Gigabit and 100 Gigabit Ethernet standards, and on June 17, 2010, the IEEE approved the 802.3ba standard as an amendment to the IEEE 802.3-2008 specification.

TECHNOLOGY OVERVIEW

The 802.3ba standard specifies extensible architectures that accommodate 40 Gigabit Ethernet, 100 Gigabit Ethernet, and a variety of physical layer media. These flexible architectures support physical layer specifications that are technically feasible and cost effective using today's technology—as well as future technology—without rewriting the standard. They also enable vendors to reuse their existing technology and intellectual property.

The following objectives were adopted by the IEEE 802.3 Working Group to ensure that 40 Gigabit and 100 Gigabit Ethernet would be seamlessly compatible with existing Ethernet networks:

- Support full duplex only
- Preserve the 802.3 Ethernet frame format
- Preserve the 802.3 minimum and maximum frame sizes
- Support a Bit Error Rate (BER) better than or equal to 10^{-12}

To ensure that 40 Gigabit and 100 Gigabit Ethernet can be carried over metro and wide area optical networks, an objective was adopted to provide support for Optical Transport Network (OTN). The IEEE has worked closely with the Telecommunication Standardization Sector (ITU-T) SG15 to define interoperable Ethernet and optical transport standards. As defined in the ITU-T G.709 hierarchy (Amendment 3, October 2009), 40 Gigabit Ethernet can be transported in OTU3 and 100 Gigabit Ethernet can be transported in OTU4.

Finally, a number of physical layer objectives (see Table 1) were chosen to provide specifications for a diverse combination of reaches and media, so that many new Ethernet applications could be supported. One of the key problems to be solved was choosing the best solution for signaling that uses available technology while still meeting market needs and cost targets. The initial electrical interfaces are either four lanes of 10 Gbps for 40 Gigabit Ethernet or ten lanes of 10 Gbps for 100 Gigabit Ethernet. The standard is flexible and will support two x 50 Gbps lanes or one x 100 Gbps lane in the future as technology evolves.

Table 1.
Summary of reach objectives and physical layer specifications.

Physical Layer Reach	1 m Backplane	7 m Copper Cable	100 m OM3, 125 m OM4 MMF	10 km SMF	40 km SMF
40 Gigabit Ethernet Target Applications: Servers, Data Center, Campus, Metro, Backbone					
Name	40GBASE-KR4	40GBASE-CR4	40GBASE-SR4	40GBASE-LR4	
Signaling	4 x 10 Gbps	4 x 10 Gbps	4 x 10 Gbps	4 x 10 Gbps	
Media	Copper backplane	Twinax cable	MPO MMF	Duplex SMF	
Module/Connector	Copper backplane	QSFP module, CX4 interface	QSFP module	QSFP module, CFP interface	
Availability	No known development	2010	2010	QSFP: 2011–2012 CFP: 2010	
100 Gigabit Ethernet Target Applications: Data Center, Campus, Metro, Backbone, WAN					
Name		100GBASE-CR10	100GBASE-SR10	100GBASE-LR4	100GBASE-ER4
Signaling		10 x 10 Gbps	10 x 10 Gbps	4 x 25 Gbps	4 x 25 Gbps
Media		Twinax cable	MPO MMF	Duplex SMF	Duplex SMF
Module/Connector		QSFP module, CX4 interface	CXP module, CFP module	CFP module	CFP module
Availability		2010	2010	2010	2011–2012

MEDIA INTERFACES

As with any new generation of technology, one design goal was to leverage as much existing technology as possible. By minimizing the number of new interfaces, the interfaces become less expensive and take advantage of volume production and simplicity. To meet this design goal, only three media modules will be used in the first generation of 40 Gigabit and 100 Gigabit Ethernet: QSFP, CXP, and CFP.

40 Gigabit Ethernet QSFP Modules

Created to replace four SFP modules in short-reach data center applications, the Quad Small Form Factor Pluggable (QSFP) module, shown in Figure 3, supports optical and electrical interfaces that have low power consumption. This module is used for a range of Ethernet, Fibre Channel, and InfiniBand applications including 40GBASE-CR4, 40GBASE-SR4, and 40GBASE-LR4 (in 2011). The QSFP is the preferred optical module for 40 Gigabit Ethernet because of its small form factor and low cost. The QSFP module has the same faceplate size as the XFP but is slightly shorter (8.5 mm high x 18.35 mm wide x 52.4 mm long).



Figure 3.
Finisar quadwire Multi-Mode
Optical Fiber QSFP.

100 Gigabit Ethernet CXP Modules

The CXP module, shown in Figure 4, was designed as the replacement for the SNAP-12 module (used primarily for InfiniBand) and was created for high-density, short-reach, data center applications. The CXP is very compact but cannot dissipate the power of long-reach modules and uses a 24-fiber parallel multi-mode cable. CXP modules are used for 100GBASE-CR10, 100GBASE-SR10, and InfiniBand 12X QDR. The CXP is slightly larger than the XFP (27 mm wide x 45 mm long).



Figure 4.
Finisar C.wire MMF CXP.

40 Gigabit and 100 Gigabit Ethernet CFP Modules

The CFP, shown in Figure 5, is a new media module that was designed for longer-reach applications, with up to 24 watts of power dissipation. Its dense electrical connectors and integrated, riding heat sink enable a range of interfaces. This module is used for 40GBASE-SR4, 40GBASE-LR4, 100GBASE-SR10, 100GBASE-LR4, and 100GBASE-ER4. While a 100GBASE-SR10 module does not dissipate much power and can fit in a CXP module, many vendors prefer to have one media module form factor and may even support 100GBASE-CR10 in a CFP module.




Figure 5.
Finisar CFP.

The CFP module is large enough to dissipate the power consumed by the 10:4 gearbox, which converts ten electrical lanes of 10 Gbps for four electrical lanes of 25 Gbps for single-mode fiber applications. These 25 Gbps electrical lanes drive four lasers at 25 Gbps, which consume considerable power as well. The CFP is also large enough to house an external Wavelength Division Multiplexer (WDM) within the module, so that all four wavelengths are multiplexed onto one fiber. It is about twice as wide as the XENPAK (14 mm high x 82 mm wide x 145 mm long). The 100GBASE-R4 40 km long-reach interface is not expected to be available until 2011 or 2012.

FROM STANDARD TO SHIPPING INTERFACES

System component and optical module vendors started developing components a few years ago to enable equipment vendors to begin building 40 Gigabit and 100 Gigabit Ethernet interfaces. Several test equipment and media module technology demonstrations were conducted in 2009. Switch and router vendors also announced 100 Gbps slot capacities and the first 100 Gigabit Ethernet interfaces in 2009. As the 802.3ba standard is now finalized, the industry expects many new interface announcements from network hardware vendors and a variety of interfaces to be available on the market in the summer and fall of 2010.

WHAT'S NEXT?

As higher-speed electrical, optical, and system component technologies are developed, Ethernet will continue to evolve—becoming cheaper and supporting higher-density line cards. Even though the 40 Gigabit and 100 Gigabit Ethernet standards have just been adopted, work on the second generation of technology using 25 Gbps electrical signaling between the ASICs and media modules has already begun.

The second generation of 100 Gigabit Ethernet will use four lanes of 25 Gbps electrical signaling. The standardization of the electrical interfaces is being developed by the Optical Internetworking Forum (OIF) and the IEEE, and Fibre Channel and InfiniBand are expected to reference this seminal work. Defining 25 Gbps signaling requires fundamental technology advances and the transition to compact four x 25 Gbps lanes is still years away. While single-mode optical interfaces are already defined at 25 Gbps, the IEEE will need to define four x 25 Gbps optical interfaces. The QSFP will need to be redesigned to support 25 Gbps signaling, and a new single-mode module, known as the CFP2, will need to be defined. The next generations of 100 Gigabit Ethernet are not expected to be available until 2014.

One new form factor for 40 Gigabit Ethernet is already being defined in IEEE P802.3bg. The new interface will be known as 40GBASE-FR and is designed to connect to OC-768 optical modules in OTN equipment at up to 2 km on SMF. 40GBASE-FR uses a 4:1 gearbox that converts four lanes of 10 Gbps to one lane of 40 Gbps. Because of the power consumption of the gearbox and the 40 Gbps laser, 40GBASE-FR will probably appear in the CFP media module first.

Additions to the Ethernet standards likely to be made in future include:

- 40 Gbps over 40 km SMF
- 100 Gbps backplane Ethernet with four x 25 Gbps signaling
- Duplex MMF specifications that use WDM or gearboxes
- Four x 25 Gbps MMF and copper form factors

CONCLUSION

One thing is certain: no one is asking for networks to be slower. The “killer app” is not a single application; it is actually ubiquitous, high-bandwidth, on-demand personalized content. The proliferation of high-bandwidth applications to millions and even billions of users will continue to drive higher-bandwidth requirements. The industry will focus on making 40 Gigabit and 100 Gigabit Ethernet interfaces more economical and available in higher-density line cards first, before jumping to the next speed. More experience and more choices will become available as the second generations of 40 Gigabit and 100 Gigabit Ethernet are developed. It is clear that Ethernet will continue to evolve as network and bandwidth requirements change to meet future requirements to scale and reduce costs.

BROCADE 100 GIGABIT ETHERNET IP/MPLS SOLUTIONS

Brocade recently announced the MLXe Core Routers for 100 GbE to its portfolio. The Brocade MLXe delivers processing and throughput capacity so that service providers can address rapid broadband growth, cloud shared services, anytime anywhere access, and IP service convergence. The Brocade MLXe integrates multi-port 100 GbE, multi Terabit trunks, and comprehensive 100 GbE packet processors to ensure wire-speed performance. The router provides one platform, one software, and five times the throughput—scaling from the service provider core through the data center network—features that significantly reduce network complexity.



Figure 6.
Brocade MLXe
Core Routers and
100 GbE module.

The Brocade MLXe is optimized for IP/MPLS Ethernet deployments, providing symmetric scaling with chassis options that include 4-, 8-, 16-, and 32-slot systems, shown in Figure 6. They offer industry-leading, wire-speed port capacity without compromising the performance of advanced capabilities such as IPv6, MPLS, and MPLS Virtual Private Networks (VPNs). For example, the Brocade MLXe-32 delivers data-forwarding performance in excess of 6 Terabits per second (Tbps) today and scales to 15.36 Tbps, enough capacity to future proof networks for years to come.

The Brocade 2-port 100 GbE module, also shown in Figure 6, is the first high-density 100 GbE module in a routing platform that uses true 100 Gbps packet processing technology. Using an advanced sixth-generation, network processor-based architecture, each 100 GbE port is powered by 100 Gbps input and output packet processing ASICs, a 100 Gbps traffic manager ASIC, and 200 Gbps of data slot capacity into the router's Clos switching fabric. The use of 100 Gbps processors and a 100 Gbps data path through the router enable each port to run at full 100 Gbps wire speed without compromising performance.

For more details on Brocade 100 GbE solutions, see a White Paper entitled, *Introducing the Brocade MLXe: Brocade's High-Density 100 Gigabit Ethernet IP/MPLS Router* on www.brocade.com. For more about the Brocade MLXe, see the Data Sheet for the Brocade MLX Series on www.brocade.com.

ABOUT BROCADE

Brocade provides innovative, end-to-end network solutions that help the world's leading organizations transition smoothly to a virtualized world where applications and information can reside anywhere. These solutions deliver the unique capabilities for a more flexible IT infrastructure with unmatched simplicity, non-stop networking, optimized applications, and investment protection. As a result, organizations in a wide range of industries can achieve their most critical business objectives with greater simplicity and a faster return on investment.

For more information about Brocade products and solutions, visit www.brocade.com.

Corporate Headquarters

San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters

Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters

Singapore
T: +65-6538-4700
apac-info@brocade.com

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 11/10 GA-WP-1498-01

Brocade, the B-wing symbol, BigIron, DCFM, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, Turbolron, and Wingspan are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, Extraordinary Networks, MyBrocade, VCS, and VDX are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

**BROCADE**