

Carrier Ethernet

Ciena's Essential Series: Be in the know.

by
John Hawkins
with Earl Follis





Experience. Outcomes

Carrier Ethernet Networks

Published by

Ciena

7035 Ridge Rd.

Hanover, MD 21076

Copyright © 2016 by Ciena Corporation. All Rights Reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without the prior written permission of Ciena Corporation. For information regarding permission, write to: Ciena Experts Books 7035 Ridge Rd Hanover, MD 21076.

Trademarks: Ciena, all Ciena logos, and other associated marks and logos are trademarks and/or registered trademarks of Ciena Corporation both within and outside the United States of America, and may not be used without written permission.

LIMITATION OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

Manufactured in the United States of America

Publisher's Acknowledgments

We're proud of this book; please send us your comments at expertbooks@ciena.com

Some of the people who helped bring this book to market include the following:

Editorial, and Senior Project Editor:

Erin Malone

Layout and Graphics:

Gilbert Rafanan

Editor:

Nancy Sixsmith

Wil McLean



Carrier Ethernet

by John Hawkins
with Earl Follis



Contents

Executive Summary.....	5
The Ethernet Story.....	6
Ethernet Branches Beyond the LAN.....	7
Ethernet’s Relationship to IP.....	9
Ethernet Technology Primer.....	10
Ethernet Party Line.....	10
CSMA/CD to Switched Ethernet.....	12
Bridge Functions.....	13
Virtual LANs.....	15
Competitive Advantages of Ethernet in the LAN.....	16
The MEF and the Advent of Carrier Ethernet.....	17
MEF Working Goals and Objectives.....	17
Carrier Ethernet Attributes and Value Proposition.....	18
Structure of MEF Standards and Services.....	20
Key Carrier Ethernet Technology Elements.....	22
Encapsulation and Transport Techniques.....	22
SONET/SDH/OTN.....	22
MPLS.....	23
MPLS-TP/PBB/PBB-TE.....	23
Approaches for Added Resilience.....	24
Operations, Administration, and Maintenance Functions.....	26
Quality of Service.....	29
Making It All Scale.....	31
Carrier Ethernet Applications and Use Cases.....	32
Bank of Ether: Financial/Banking Enterprise.....	32
BioEther Healthcare: A Hospital/Clinic Multisite Enterprise.....	37
Ether State University: A Growing Multisite Education Institution.....	44
Carrier Ethernet Futures.....	50
Related Ethernet Developments.....	51
Precise Timing Distribution.....	52
Power over Ethernet.....	52
Software-Defined Networks and Network Functions Virtualization.....	52
Glossary/Acronyms Section.....	54



EXECUTIVE SUMMARY

Ethernet has been around for decades and has claimed the title of the superior Local Area Networking (LAN) technology of choice. The ubiquity and popularity of Ethernet technology continue to drive down the price of Ethernet components, so its economic advantages and market acceptance will likely continue for many years to come. Carrier Ethernet (CE) looks to build on that ubiquity, popularity, and economic advantage to provide a flexible framework, universal frame format, simple design, and enhanced Operations, Administration, and Maintenance (OAM) capabilities for high-performance networks ideally suited for both enterprises and carriers.

Ethernet bit rates have continued to increase, traditionally growing tenfold each time a new rate is defined. Gigabit Ethernet (GbE) interfaces are widely deployed in PCs and servers, and are 10 Gb/s in LAN backbones. With bit rates in the hundreds of Gb/s and extensions to the standard continuously under development at the MEF and the Institute of Electrical and Electronics Engineers (IEEE), Ethernet shows no sign of slowing. Because CE is based on Ethernet, the conversions to and/or from other types of Wide Area Networking (WAN) protocols are not necessary. As a WAN technology, CE is cheaper and faster than legacy technologies such as Asynchronous Transfer Mode (ATM), Synchronous Optical Network (SONET), and Frame Relay. CE also gives carriers and enterprises increased management capabilities while plotting out a road map that will support increasing speeds over the next few years.

This book covers the history of CE, including a discussion of how CE grew out of the Ethernet standard, how CE differs from traditional Ethernet, the maturing of CE via the MEF, the CE value proposition, and the evolution towards Software-Defined Networking (SDN) and Network Functions Virtualization (NFV) in a CE world. This book demonstrates how Ethernet has evolved into a robust and scalable solution for large networks, whether operated by service providers or enterprises. A few use cases will help define where and when one might consider implementing CE solutions. Though no book can answer all questions, this book should offer some enlightenment on the history, present, and future of CE.

This book covers:

- How Ethernet evolved
- Ethernet networking basics
- Competitive advantages of Ethernet
- Ethernet integration with optical networking technologies
- The MEF and its ongoing role with CE
- Why CE is a valued and viable option for carriers, network operators, and enterprise customers
- The future of CE

THE ETHERNET STORY

The term *Ethernet* has its roots in the Latin term *aether*, a mythical medium once believed to fill all 'empty space' on Earth while providing for magnetic wave transmission. The term was coined in a doctoral dissertation written by Bob Metcalfe and the original patent application Metcalfe filed while working at Xerox Palo Alto Research Center (PARC) with fellow researchers David Boggs, Butler Lampson, and Chuck Thacker in the early 1970s.

Modern networks provide high-bandwidth, high-speed access to mission-critical data locally, remotely, and in cloud-based environments. The most popular networking technology for LANs is Ethernet, and the vast majority of corporate LANs in the world run on it. But this was not always the case, as in the 1980s and early 1990s, there were a number of competing LAN technologies, including ARCNET, StarLAN, Fiber Distributed Data Interface (FDDI, an early fiber-based LAN ring-based technology), token bus, and token ring. By the mid-1990s, Ethernet was the clear winner of the LAN sweepstakes because of its native structural compatibility with Internet Protocol (IP) networks, plug-and-play simplicity, and (relatively) low cost.

In its original LAN-centric definition, Ethernet was envisioned as a technology to connect devices residing in the same general area, usually within a single building. The number of devices to be interconnected were few, and the distances between them relatively small. (The protocol itself was limited to a range of 100 meters.)

Ethernet Branches Beyond the LAN

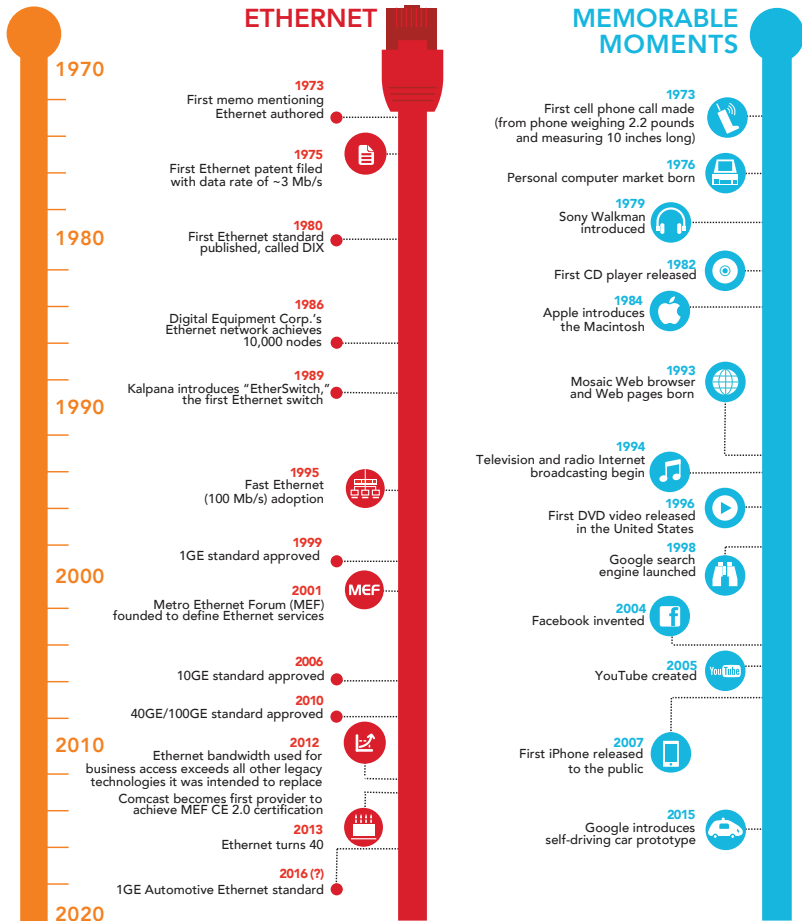
As the Internet grew, Ethernet was poised to play a role beyond the original single-building, short-distance applications originally envisioned. Entire campuses with hundreds or thousands of computers, servers, printers, and associated devices were being wired for Ethernet connectivity. More recently, WiFi networks (sometimes referred to as 'wireless Ethernet') have become commonplace. The original assumptions of the number of interconnected devices and the distances between them expanded considerably. With some evolutionary adaptations (the addition of Virtual Local Area Networks (VLANs) and faster versions of the Spanning Tree Protocol (SPT), for example), Ethernet was in the right place at the right time to evolve with these expanded requirements because it was a cost-effective yet high-performance networking technology that proved to be adaptable .

Other protocols were developed to interconnect these growing Ethernet 'islands' across longer distances. Data communications protocols such as the X.25, ISDN, Frame Relay, and ATM were devised and sold by carriers to provide connections across longer distances. In the early 2000s, *Carrier Ethernet* was introduced as a natural way to extend the Ethernet protocol to provide wide area network (WAN) connectivity. As its name implies, carriers appreciate CE because it enables a variety of network services that can be sold to other carriers (wholesale) and end-users (retail). These include mobile backhaul, business services, and data center interconnect. However, CE is not just for carriers. It is a feature-rich solution for all network operators that provide access to end-customers. This book discusses the establishment of CE as a networking technology and the benefits of migrating to CE-based networks in the modern network context. Figure 1 shows a timeline of the development of Ethernet alongside optical networking technologies and a few other memorable technology milestones.

The continued proliferation of bandwidth-hungry apps such as email, web browsing, voice, and video traffic provided the original impetus for higher speeds in the LAN and WAN. More recently, emerging technologies such as storage virtualization and cloud computing have compounded the bandwidth requirements for network operators and enterprise customers. As users (especially businesses) have become more reliant on the cloud, the network has only become more critical to society, and reliance on the public Internet

THE HISTORY OF ETHERNET

A timeline of innovation in Ethernet and other memorable tech moments



SOURCES: MEF, Wikipedia, Heavy Reading, TIME, Business Insider, CBS News, University of North Carolina, Bright Hub, Tech Hive, Vertical Systems Group

Figure 1: Ethernet History

for access to data centers has been replaced with a hybrid (public and private) infrastructure. These trends are unlikely to fade any time soon, and the demand for ever-increasing bandwidth continues. Once again, Ethernet has evolved as a key solution to address this critical need.

Ethernet's Relationship to IP

Of course, the world of the Internet is driven by IP in its many applications. As soon as such an application (such as an email application) makes a request of the network, a series of data manipulations occur in the protocol stack of the requesting station. These applications use their own protocols to coordinate the exchange information with faraway hosts on the Internet (e.g., SMTP, FTP, and HTTP), but IP is universally used as the Network layer (often referred to as Layer 3) to communicate with other peers, peripherals, and so on. The OSI model shown in Figure 2 summarizes the layers in the stack.

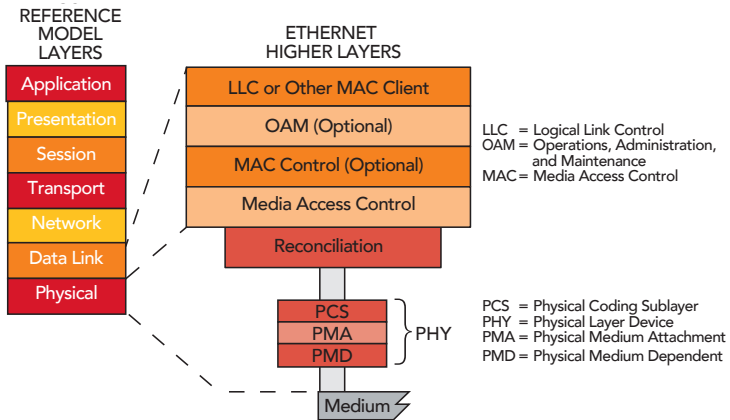


Figure 2: OSI Stack and Ethernet Model Layers

Note that Ethernet is part of both Layer 1 (the Physical (PHY) Layer) and Layer 2 (the Media Access Control (MAC) Layer). The PHY component defines the various physical implementations of Ethernet, such as 10BASE-T (10 Mb/s on twisted-pair copper) or 10GBASE-LR (10 Gb/s on long-reach fiber). These are defined for each data rate and media type of which there are dozens in the IEEE specs. As a Layer 2 protocol, Ethernet is considered the 'server layer' for IP, which (at Layer 3) can be referred to as a client. Ethernet receives requests from and provides responses to the IP layer above it. So Ethernet and IP work hand in hand.

This relationship has served Ethernet well as it transitioned to CE. The availability of CE services avoids significant complexities in mapping IP data into legacy telco services such as Frame Relay and ATM, and associates the development of CE services more closely with the ongoing development of

Ethernet. This book will explore this concept more fully when detailing the role of the MEF.

ETHERNET TECHNOLOGY PRIMER

Perhaps a natural question to ask at this point is, “what exactly *is* Ethernet?” The answer is multifaceted. At its root, Ethernet is a networking protocol that facilitates using some sort of physical medium, or PHY (e.g., twisted-pair copper, coax, fiber, and wireless spectrum) for data communications. The protocol is captured in a series of standards issued by the IEEE LAN/Metropolitan Area Network Standards Committee known as IEEE 802. The committee continues working to expand the standards and amplify the utility of Ethernet in modern networking applications.

The basic Ethernet frame type in use today is known as Ethernet Type II. Figure 3 shows how the Ethernet frame consists of the header (with Destination Address (DA), Source Address (SA), and protocol type), the user data or payload, and the frame check sequence or Cyclical Redundancy Check (CRC).

The Ethernet header includes a 6-byte DA (the intended recipient of the packet), a 6-byte SA (the device that sent the packet), and a 2-byte protocol type field. The CRC is 4 bytes long and is checked at every switch, and malformed or corrupted frames are dropped. According to the 802.3 standard, Ethernet frames must be at least 64 bytes long, so payloads must be at least 46 bytes in length. When actual user data is shorter than 46 bytes, the PHY will ‘pad’ the packet to make up the minimum size. The maximum length of an Ethernet packet is 1,518 bytes, with the maximum user data payload being 1,500 bytes. Later implementations allowed for so-called ‘jumbo’ frames up to 9,000 bytes long to facilitate certain types of large traffic flows such as file transfers and video links.


Ethernet Party Line

Think of early Ethernet as a party-line phone system that was common 50 years ago. In this system, multiple houses shared a single telephone wire, and only one phone user could talk on the telephone line at a time. In rural areas where the population was sparse and it was cost-prohibitive to run a separate

BASIC ETHERNET FRAME



VLAN FRAME



Q-IN-Q FRAME



MAC-IN-MAC FRAME




Figure 3: Ethernet Frame Formats

phone line to each house, party-line phone systems were a logical solution. To use a party-line phone, a caller would first pick up the phone and listen for a dial tone. If he or she did not hear a dial tone or heard other people talking, the caller would hang up and wait until the telephone line became available. If the caller began talking without listening to see if anyone else was talking, she might collide with a different conversation already in progress. Once the caller realized that someone else was already engaged in a conversation, she would have wait until that other conversation was terminated before placing her new call.

Early Ethernet worked in the same way: A network device first tried to transmit data then listened for a collision with other activity on the network segment. If a collision was detected, the sender automatically waited a random time period of a millisecond or two, then checked again to see if the network was available to send data. This network access scheme is called Carrier Sense, Multiple-Access, Collision Detection (CSMA/CD), explained below:

- **Carrier sense** means listening to the Ethernet line for a gap in network traffic to ascertain whether someone is already 'speaking' on the network.
- **Multiple access** means that many Ethernet devices can listen and transmit on the same physical network wire, although each device must take its turn so multiple conversations do not occur at the same time.
- **Collision detection** means the potential sender starts to transmit before the previous talker is done talking, although the sender does subsequently detect that the previous conversation is still place. Therefore, the sender stops talking in order to allow the previous conversation to finish.

CSMA/CD to Switched Ethernet

CSMA/CD works well as a network access scheme as long as links and number of nodes are kept small. However, CSMA/CD does not provide an effective network access mechanism as more devices are added to a CSMA/CD network or as the physical distance between those devices increases. Too many devices attempting to transmit data simultaneously results in contention for network access as each device attempts to

transmit data, detects a collision, waits a random number of milliseconds, and then retries the data transmission. As networks grow, this type of contention results in network latency. Similarly, as distance increases between the devices, the latency in the link can cause constant collisions and back-off periods, which also diminish throughput and increase latency. These challenges led to the birth of the modern switched Ethernet network. In the newer model, each device has its own network segment connecting it to an Ethernet switch port, which uses buffers to queue frames as needed, thus eliminating collisions and their resultant latency issues. So the previous party-line telephone analogy now gets upgraded to the modern telephone network, with a dedicated wire running between each transmitting device and a telephone switch which isolates the end-parties and connects them when the lines are free.

Similarly, with switched Ethernet networks, collisions are no longer an issue as the switched network segment to which a device is attached is never in use by other devices. Each device connects to a discrete port on an Ethernet switch, and only one device is attached to each port. Instead of all devices on an Ethernet LAN having to listen for and handle collisions, the Ethernet switch serves as the traffic cop for all devices connected to that switch, handling the consolidation of all traffic to and from devices residing on that switch. The switch recognizes the MAC address of each connected device and to which port it is connected. If one device talks to another on the same switch, the switch completes the connection between the two devices so their conversation can proceed without interruption. If the switch sees network traffic destined for a MAC address that is not connected to the same switch, that traffic is routed through the switch uplink port, which connects the switch to other switches in a data center or to a router that connects to other networks over WAN links.

Bridge Functions

Bridging occurs when one or more network segments are connected into a single aggregated network. Based on the IEEE 802.1d standard, bridging differs from routing in that routing connects and routes traffic between separate networks that remain distinctly separate, whereas bridging creates one combined MAC layer network. The most prevalent example of bridging occurs in Ethernet switches: the switch learns which source MAC addresses

reside on which physical switch ports and then forwards traffic based on the address residing in a forwarding table. If the DA is unknown to the switch, the data is flooded to all ports on that switch. This concept, known as multi-port bridging, serves as the basis for modern Ethernet network switches. The following bullets present additional details about the five most important bridge concepts for CE: learning, forwarding, filtering, flooding, and STP.

- **Learning:** Learning occurs as a switch monitors the SAs and corresponding port numbers in which each SA was observed. Once the SA/port number correlation is made within the switch software, that mapping is stored in a forwarding table for future reference (also called the Forwarding Information Base [FIB]). The switch constantly updates the forwarding table as traffic flows through the switch, in case the location of specific addresses changes from port to port on the switch. Address entries in the FIB also “age out” over time if no packets are sourced from the MAC address for a specified period of time to avoid the table growing exceedingly large.
- **Forwarding:** Forwarding consists of the Ethernet switch looking up the DA of an incoming Ethernet frame in the FIB and sending that frame to the port corresponding to that address. Forwarding, which is the basis of Ethernet switching, allows for network segmentation by port while supporting high speed throughput between ports on a switch.
- **Filtering:** When an Ethernet switch sees packets that show the source and destination devices residing on the same port, filtering occurs to avoid sending traffic back to its source segment. All clients residing on the same port see all traffic to/from the MAC addresses on that port, so no forwarding is required. Filtering discards any such packets, as well as all other malformed or partial packets.
- **Flooding:** If a switch sees a packet with an unknown DA (ie. there is no address entry for the DA in the FIB), the switch sends that packet to all ports on the switch (except originating the port) in an effort to deliver that network traffic to its intended recipient. Other situations in which flooding can occur include using broadcast packets (one

packet should be delivered to all ports by design) and multicast packets (a packet is addressed to one or more DAs that do not reside in the FIB).

- **Spanning Tree Protocol:** STP was introduced to prevent loops in bridged networks. Loops can occur when there is more than one active network path between any two network nodes. Such redundancy would cause frames to circulate endlessly, clogging (if not completely incapacitating) the bridging function. STP simply 'blocks' one of the redundant paths, so this circulation cannot occur. In a nod to network resiliency, STP enables the redundant path to act as a backup path by automatically re-enabling the blocked paths when the primary path fails. But the original STP was fairly inefficient, often taking on the order of tens or even hundreds of seconds to 're-converge.' More recently, Rapid STP and Multiple STP variants have been defined to speed up this process.

Virtual LANs

As networks grew, it became desirable to subdivide the physical switches into virtual switches to limit the effects of flooding and simplify the STP function. A 4-byte VLAN tag (or Q-tag, after the 802.1Q standard that defined it), was inserted into the Ethernet header. Two bytes simply identified that the frame was now a Q-tagged frame; the remaining bytes allowed for the definition of up to eight classes of service (3 bits, known as the Priority Code Point [PCP] bits), and 4092 VLANs (12 bits, also known as the VLAN ID). One final bit is used (although rarely) if a frame is 'discard eligible' (the Discard Eligibility Indicator [DEI]).

Of interest to this discussion are the 4094 VLANs, which limit the capability of the switch to forward or flood traffic with unknown DAs to only those ports identified with a given VLAN ID. This limits the unnecessary broadcast of data to end-points where it is not needed and simply consumes bandwidth and network resources (also known as 'broadcast storms'). In a subsequent edition of the 802.1 standard, a second Q-tag was defined for use as a second layer of VLANs with its own VLAN IDs and P-bits. Known as Q-in-Q (refer to Figure 3: Ethernet Frame Formats 2), this approach has been popular with operators who allow the end-user to define the 'customer' VLAN ID (or CVID or 'inner tag')

while using the provider VLAN ID (or PVID or 'outer tag') to identify individual customers or services on the carrier infrastructure. Sometime later, the IEEE 802.1ah standard, known as Provider Backbone Bridging (PBB) enabled an additional MAC header to be added to the frame (again, see Figure 3: Ethernet Frame Formats²), providing the ultimate in flexibility for large Ethernet networks with virtually unlimited scale. In addition, PBB provides clear separation between the service provider and customer networks because each has a dedicated set of MAC addresses (and associated FIBs). When an Ethernet frame reaches the Ethernet User Network Interface (UNI), the service provider MAC address is added to the customer's Ethernet frame. The service provider network switch then checks this MAC address against its FIB, and forwarding, filtering, learning, etc. go on as usual. This is an added advantage, because only switches at the edge of the provider network need to be PBB-enabled. Switches in the core of the network switch on a standard MAC header (in this case, the service provider header) so any IEEE 802.1 Ethernet switch will suffice.

Competitive Advantages of Ethernet in the LAN

One major benefit of Ethernet equipment is how well it gets along with most networking gear. Built-in features ensure that when different variants of Ethernet (e.g., two different PHY data rates) are interconnected, Ethernet attempts to use the most powerful features available to provide the best possible performance. Specialized negotiation and sensing processes identify the common denominator between different hardware configurations and protocol formats. Ethernet has established a well-earned reputation for being 'plug-and-play' in terms of network interoperability.

Another key factor in the early adoption of Ethernet was its low cost when compared with competing technologies. The simplicity of the protocol and ease of day-to-day management made it popular with IT departments worldwide and helped settle the early 'protocol wars' with technologies such as Token Ring, Token Bus, and FDDI. As the personal computer industry took off, Ethernet was a nonproprietary solution whose time had come. This led to what has become known as the Ethernet 'virtuous cycle.' As Ethernet became popular, it gave incentive to components and PC manufacturers to increase their volume of production for chipsets, modules, and cards. As these volumes rose, the per-unit cost was driven down, because R&D costs and

manufacturing startup costs could be amortized over millions of units. And as prices came down, Ethernet was adopted even further, including in residential, data center, and multitenant applications.

THE MEF AND THE ADVENT OF CARRIER ETHERNET

The MEF¹ was established in 2001 by a coalition of networking industry vendors with the goal of developing a new set of standards principally for optical-based, metropolitan carrier CE, used to connect enterprise LANs. Although the scope of the MEF's work has evolved over the years to be global in nature (i.e., well beyond its initial focus on the metro network), it remains very active in developing and revising CE standards as line speeds increase and additional features are added. In addition of the MEF also provides a certification program for CE products and services to ensure compliance to the standards. Finally, the MEF provides certification programs for networking professionals looking to expand their knowledge and skills to support CE products and services.

MEF Working Goals and Objectives

MEF's mission is to "accelerate the worldwide adoption of Carrier-class Ethernet networks and services." The MEF pursues that mission via four specific work goals that are overseen by four standing committees:

- **Technical Committee:** Develops technical specifications and implementation processes for CE services and architecture that ensures interoperability of CE hardware and software. The MEF Technical Committee continues to develop and manage more than 50 specifications related to CE interoperability and implementation. Many of these specifications are in their second or third iteration as part of the ongoing refinement of CE specifications and standards by the MEF.
- **Certification Committee:** Develops certification programs for CE-related hardware, software, services, and networking professionals. The MEF Certification Committee administers CE test suites based on equipment and services specifications developed by the MEF Technical Committee.

¹Originally known as 'Metro Ethernet Forum,' today it is simply 'MEF' reflecting its global role beyond the metro.

- **Service Operations Committee:** Develops standardized processes for purchasing, selling, delivering, and managing CE services. The MEF Services Operations Committee was formed in 2013 to implement this significant work task.
- **Marketing Committee:** Educates stakeholders and raises awareness of the MEF via development of CE use case studies, marketing collateral, audio and video clips, and white papers that clarify and communicate the MEF’s goals. The MEF Marketing Committee organizes CE webinars, videos, and conferences that educate interested parties on CE specifications and the work of the MEF.

Carrier Ethernet Attributes and Value Proposition

One of the significant early accomplishments of the MEF was to capture and formalize the vocabulary used when describing Ethernet services and attributes. For instance, there used to be many definitions of exactly what ‘carrier’ or ‘carrier-class’ Ethernet should be or deliver in terms of service quality, reliability, and so on. Figure 4 shows what became a popular diagram used to describe the five attributes of CE:

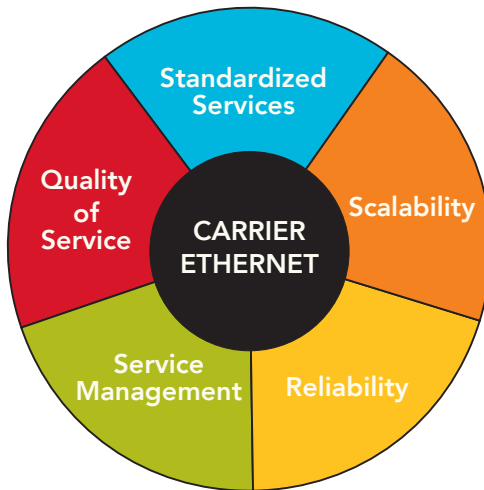


Figure 4: Five Carrier Ethernet Attributes

- **Standardized Services:** The need for standards-based solutions that result in predictable and repeatable use of the service. MEF services rely on IEEE standards, for example, to define physical details such as line rates, encodings, and packet sizes
- **Scalability:** The capability to support such services for a large variety of uses (business, residential, mobile) over very large distances (beyond the LAN) and at higher and higher data rates as they become technically feasible
- **Reliability:** The capability of the network to detect and recover from failures and meet the most demanding availability requirements (50 ms recovery time is a well-established metric in the industry)
- **Quality of Service:** Enables a wide choice of options in terms of performance metrics required to fulfill Service-Level Agreements (SLAs), including those appropriate for voice, video, data, and mobile services
- **Service Management:** The capability to visualize the infrastructure, roll out services, diagnose problem areas, and carry out the day-to-day management of a network

But beyond the basic vocabulary describing what qualifies as a CE service, the actual service definitions (captured in the MEF 6 specs) and their attributes (captured in the MEF 10 specs) clarify in great detail what each service should be called and how it should behave. Throughout its history, the MEF has been careful to steer clear of dictating specific implementation details, leaving them to vendors and operators. The MEF's key role has been to articulate the industry consensus on how CE services should behave, and the attributes necessary to build, deploy, and ultimately buy/sell them. MEF certifications are intended to reduce the risk of the adoption of MEF services and promote their uniform usage across the spectrum of telecom offerings.

The drivers for CE adoption include the push to reduce carrier networking costs by implementing converged networks that combine business, residential, and wireless traffic to create economies of scale. Enterprises and

other network operators can also realize substantial savings by converging corporate networks in the same way CE supports the convergence of carrier networks.

Structure of MEF Standards and Services

The MEF Ethernet Services Model (Figure 5) defines a set of building blocks that are then used to assemble a variety of CE services. For instance, the UNI is formally defined along with its attributes (e.g. unique ID, physical layer type, etc). Ethernet Virtual Connections (EVCs) are also defined with attributes such as VLAN tag preservation, Class of Service (CoS) preservation, and various performance attributes. Connecting two UNIs via an EVC creates an E-Line service type. By using these building blocks in this way, service providers can construct a rich set of CE services that can be understood by buyers and sellers alike using consistent and standardized nomenclature.




Figure 5: MEF Service Model

MEF documents also differentiate between operators and service providers. An operator administers a CE network, but sells services to other operators; service providers sell them to end-users. Sometimes these are referred to as wholesalers and retailers, respectively. Formal definitions of Operator Virtual Connections (OVCs), as compared with EVCs, are captured in MEF 26.1, which defines External Network-to-Network Interfaces (ENNIs) and MEF 51, which defines OVCs. These additional building blocks allow for creative and flexible services to be offered from the CE menu (Figure 6).

	Service Type	Port-based Service	VLAN-aware Service
EVC Services	E-Line P2P EVC	Ethernet Private Line (EPL) 	Ethernet Virtual Private Line (EVPL)
	E-LAN MP2MP EVC	Ethernet Private LAN (EP-LAN) 	Ethernet Virtual Private LAN (EVP-LAN)
	E-Tree RMP EVC	Ethernet Private Tree (EP-Tree) 	Ethernet Virtual Private Tree (EVP-Tree)
OVC Services	E-Access	Access EPL 	Access EVPL
		Access E-Line P2P OVC with 1 UNI and 1 ENNI	
	E-Transit	Access E-LAN MP2MP OVC with UNI(s) and 1 ENNI 	
		Transit E-Line P2P OVC with 2 ENNIs	
		Transit E-LAN MP2MP OVC with ENNI(s) 	

EVC: Ethernet Virtual Connection
OVC: Operator Virtual Connection
P2P: Point-to-Point

MP2MP: Multipoint-to-Multipoint
RMP: Rooted MultiPoint

EPL: Ethernet Private Line
EVPL: Ethernet Virtual Private Line

Figure 6: MEF Service Types

The ultimate benefit of MEF activities, however, has to do with their certification program. Both CE equipment and services can be certified as compliant to the various specifications by a third-party technical evaluation using hundreds of test cases. By certifying their service offerings, providers achieve several goals, including demonstrating a serious commitment to offering standards-compliant services and a willingness to prove this commitment to a neutral party. Regardless of whether the customer is an end-user or another operator, buying and selling certified services avoids a lot of complex one-on-one contractual negotiations that otherwise would be needed.

Increasingly, certification is being requested, if not required, in Requests for Proposals (RFPs) between operators. The advantages are considerable in terms of long-term expense avoidance (from all those “bespoke,” or customized services) and time to market. Buyers and sellers can quickly identify the requirements using industry-accepted MEF service definitions and associated attributes, and thus get on with business quickly and efficiently.

One of the keys to achieving service-level certification is to work with systems providers who have achieved certifications for their gear using the same test procedures used at the service level. Such collaboration is facilitated by MEF and increasingly common in the industry.

KEY CARRIER ETHERNET TECHNOLOGY ELEMENTS

Carrier Ethernet has evolved into a set of services that rely on a variety of technology elements to achieve the five key attributes discussed previously. Although the attributes do not define specific implementations, some of them are addressed here.

Encapsulation and Transport Techniques

To transport an Ethernet service (an E-Line or E-LAN) across a specific infrastructure (for example, SONET/SDN , Optical Transport Networking [OTN], or MPLS), many techniques exist to encapsulate Ethernet frames into various transport infrastructure types. This allows an adaptation of the Ethernet user data and headers into the infrastructure-specific protocol data unit. Once across the network infrastructure, the Ethernet frame can be de-encapsulated and delivered to the destination in its native form.

SONET/SDH/OTN

Given its popularity as a carrier infrastructure, Layer 1 protocols such as SONET/SDH and OTN have been used to transport CE services in many operator networks. Techniques such as Generic Framing Procedure (GFP) allow Ethernet and other client protocols to be mapped to these synchronous transport infrastructures. OTN defines specific Optical Data Unit (ODU) types appropriate for the various Ethernet line rates so they can be transparently transported. These approaches enable Ethernet services to be multiplexed across these infrastructures with other native services such as video, voice, and Fiber Channel (FC) while avoiding ‘forklift upgrades,’ even in the face of significant Ethernet service growth. This approach also allows Ethernet services to be fully protected with sub-50 ms restoration provided by the underlying Layer 1 protocol.

MPLS

Multiprotocol Label Switching (MPLS) has been widely deployed in core networks to carry packet-based traffic between core routers. MPLS simplified much of the address lookup procedure required by IP and sped up the transition from TDM transport links to packet-based links. In time, MPLS was also used to transport Ethernet frames by emulating a Layer 2 network (L2VPN) or a Layer 3 routed network (L3VPN). This method of transport has become increasingly popular because MPLS allows for fast restoration times, multiservice multiplexing, and efficient bandwidth engineering. Some operators have begun to extend the MPLS 'cloud' to the edges of their networks, but are finding that the complexity of the protocol and lack of certain management tools (see the OAM section ahead) to be key challenges.

MPLS-TP/PBB/PBB-TE

Recently, MPLS-Transport Protocol (MPLS-TP) was devised to both simplify some of the MPLS constructs and create the management tools required by operators to administer their networks on a day-to-day basis. These include functions such as connection verification, fault monitoring, and in-band control/management.

As discussed earlier, QinQ and PBB techniques allowed a separation between network operator and end-user addressing domains. They simplified the bridging functions and limited the effects of broadcast storms and network failures. However, the principle attribute of transport systems used by major operators has to do with selecting appropriate paths which packets take through a network. In many cases, connection-oriented routes are preferred because they provide stable, predictable behavior especially in large networks. Provider Backbone Bridging – Traffic Engineering (PBB-TE) and MPLS-TP provide mechanisms to remove the notion of flooding of unknowns and STP (which are troublesome from a security and scalability point of view). Instead, FIB tables are explicitly populated by a centralized management or control entity. As a result, the operator can create prescribed, predetermined paths, resulting in totally predictable network behavior. These protocols have the additional advantage of provisioning services, in the same way SONET/SDH were provisioned for decades in transport networks. Given the success of SONET/SDH, the familiar operational model is a significant advantage for

service providers, who can then maximize network utilization and minimize the cost per bit carried. In addition, security is increased because any misconfiguration or frame leakage is minimized and localized to a small portion of the network.

Approaches for Added Resilience

As discussed earlier, STP was designed to prevent bridge loops in Ethernet networks. Such loops could otherwise cause endless frame circulation and result in total chaos within the basic Ethernet flooding and learning algorithm. STP also allows spare links in the network to become automatic backup links in case of a link failure. Although this works well in theory, the algorithm tends to be slow to 'converge' when such failures occur (on the order of tens of seconds or even minutes, depending on the size of the spanning tree). This results in dropped traffic and poor Quality of Service (QoS). Alternatives were devised by the IEEE including Rapid Spanning Tree and Multiple Spanning Trees. These have reduced convergence times, but still cannot match the de-facto expectation of sub-50ms restoration observed in traditional SONET/SDH networks.

Another alternate resilience technique is to employ the Link Aggregation Protocol (LACP; IEEE 802.3ad) using multiple Ethernet physical media across multiple ports in parallel to increase link redundancy and availability. LACP allows for faster connections between devices managed as a single connection, load sharing, and load balancing among the individual links within a logical connection. A failure in one of the links would result in partial connectivity via the remaining healthy links (albeit the overall bandwidth capacity would be reduced). However, while LACP failover times can be sub-second, they do not meet the traditional sub-50 ms benchmark expected by typical transport network operators.

Several approaches can provide such protection to Ethernet services. One is to map Ethernet frames into SONET/SDH or OTN frames, as previously discussed. Ethernet Protocol Data Units (PDUs) can be mapped as client frames into the appropriate SONET/SDH/OTN frame structure and carried transparently across this TDM infrastructure. The underlying (Layer 1) protection mechanisms inherent in the SONET/SDH/OTN protocols are thus able to provide sub-50 ms protection in the event of a failure.

With connection-oriented protocols such as PBB-TE and MPLS-TP, connection-oriented tunnels can be backed up by one or more protection tunnels (also connection-oriented). The failover from the primary tunnel to the backup tunnel can be signaled via a 'heartbeat' that, when missing, indicates a loss of connectivity (refer to the following discussion on IEEE 802.1 ag). This failover can be accomplished in well under the prescribed 50 ms expectation of transport operators.

Yet another approach involves using ITU-T G.8032, as it provides ring-based connectivity and protection. It also does not depend on STP for loop avoidance since it uses its own approach to avoid loops. Given the prevalence of ring-based topologies, especially in metropolitan network scenarios, G.8032 is an Ethernet ring-protection technology that provides a simple, cost-efficient, and scalable solution for rapid service restoration that delivers SONET/SDH-grade resilience at reasonable cost. G.8032 technology also allows service providers to future-proof their networks, giving them significant capital expenditure advantages. G.8032 technology:

- Supports flexible deployment models in access, metro, and core network segments including rings, subrings, and the multiple connections between them
- Supports deterministic sub-50 ms service restoration for all services and service types being transported over the ring
- Addresses service providers' SLA needs (such as fault, performance, and protection tolerances)
- Can interwork with existing networking solutions such as MPLS, MPLS-TP, VPLS, PBB-TE, and PBB
- Efficiently supports multiple service types (such as E-LINE, E-TREE, and E-LAN) and key applications
- Supports efficient resource bandwidth utilization
- Supports traditional transport operational procedures involving force switch and do-not-revert

Operations, Administration, and Maintenance Functions

Operations, Administration, and Maintenance (OAM) functionality is a hallmark of the legacy TDM-based services that CE increasingly replaces in the modern network. Ethernet OAM is therefore required to match those capabilities to provide visibility and assurance to the correct function and services performance. These are key to delivering on the service management attribute of CE, which is especially important because Ethernet services traversing the WAN can span hundreds to thousands of kilometers.

As a result, a rich OAM toolkit has recently been added to the Ethernet protocol standards to allow for OAM functions that provide fault and performance management tools, which were not part of the original LAN technology. Table 1 lists several important new OAM functions.

OAM Standard	Standards Body	Description
802.1ag Connectivity Fault Management	IEEE	Specifies continuity check, loopback and link trace protocols to detect, locate, and isolate network faults. These can be applied to various maintenance domains corresponding to end-users, service-providers or operators. Continuity Check Messages (CCMs) can be used to detect faults between defined points (end-points or intermediate points) in the network. Similarly, Loopback Messages (LBMs) can be used to confirm that specific maintenance points are reachable, and Link Trace Messages (LTMs) can isolate where a fault has occurred.

OAM Standard	Standards Body	Description
Y.1731 OAM Functions and Mechanisms for Ethernet-based Networks	ITU-T	Supports virtually identical fault management functions as 802.1ag, but adds tools for performance monitoring including Frame Loss Ratio (FLR), Frame Delay (often called latency), and Frame Delay Variation (FDV) (often called jitter). These are important metrics for verifying customer SLAs and supporting high-performance services such as VoIP, video delivery, and data center interconnect.
802.1ah Ethernet in the First Mile	IEEE	Defines mechanisms for monitoring and troubleshooting Ethernet access links. Specifically, it defines tools for discovery, remote failure indication, local loopbacks, and status and performance monitoring. The most common is associated with a specific feature known as 'dying gasp' whereby a node can quickly send a message upstream, indicating it has suffered an unrecoverable local failure such as loss of power.

*Table 1: Today's Rich Carrier Ethernet OAM 'Tool Kit'
(continued on next page)*

OAM Standard	Standards Body	Description
RFC 2544/ Y.1564 Ethernet Service Activation Test Methodology	IETF/ITU-T	Both standards define Ethernet-specific test methodologies to provide verification of key service attributes at the time service is set up (often called a service 'birth certificate'). They define a virtual test generator and reflector functions to allow for throughput, latency, jitter, and bandwidth profile measurements that are important to the SLA. The generator function generates test patterns that are injected into the network, while the reflector creates a loopback to allow for two-way measurements. While both of these standards are similar, Y.1564 is the more recent and improved standard, widely adopted in the latest gear. These functions can be built-in ('in-skin') to the network elements or incorporated into dedicated (often hand-held) test equipment for use in the field.
RFCs 4379, 6371, 6428 (and others) Label Switched Path (LSP) Ping/TR, + OAM Framework for MPLS-based Transport	IETF	This series of RFCs creates MPLS-specific OAM functions to verify connectivity at the MPLS level (such as detecting and isolating path failures or label mismatches). MPLS-TP introduced these functions to provide for monitoring paths, identifying defects, and alerting the Network Management System (NMS) when critical events occur.

OAM Standard	Standards Body	Description
RFC 5357 Two-Way Active Measurement Protocol (TWAMP)	IETF	Enables two-way performance measurements with TCP/IP and a more accurate time-stamping technique than is available with conventional Ping/Traceroute functions. TWAMP can measure packet loss, latency, jitter, duplication, and out-of-order scenarios, among other things.
IEEE 802.1AB Link Layer Discovery Protocol (LLDP)	IEEE	Performs functions to discover the identity and capabilities of network devices in an Ethernet network. Information that can be discovered includes node name, port name/description, VLAN name, management IP address, and other system capabilities supported. Network and element management systems can then query the device to reveal and act on the information.

Table 1: Today's Rich Carrier Ethernet OAM 'Tool Kit'

Quality of Service

It is fair enough to describe a service behavior under ideal circumstances. The challenge comes when the network is placed under stress, such as when traffic loads stress various sections of the infrastructure. How does the network react to excess traffic loads and still maintain its contractual SLAs with end-users? Dropped traffic and excessive delays can result if the network is not architected to handle such situations.

The hop-by-hop behavior of traditional Ethernet leads to lack of determinism. In other words, because the service provider does not know how a frame travels across the network, it is impossible to predict how it will be affected by traffic congestion. In a CE environment, QoS and traffic engineering techniques (such as those provided by connection-oriented Ethernet) allow

prioritization of traffic streams to ensure that SLA parameters are maintained across the network.

This means providing end-to-end performance based on Committed Information Rate (CIR), frame loss, delay, and delay variation targets. When congestion occurs in a network, Ethernet switches will need to store packets in a queue until they can be scheduled for forwarding to their destination. A variety of techniques and approaches can help decide which frames get priority, but most require frames to be classified based on some policy (based on type of traffic, SLAs, etc.). Each type of traffic is stored in separate queues that are then scheduled onto the appropriate output port based on a scheduling 'discipline' such as weighted round robin, random early discard, weighted fair queuing, and others.

The result is that service providers can match the requirements for voice, video, and data over converged business and residential networks. Creating hierarchies of QoS levels allows delivery of appropriate resources to both service and application-level needs. For instance, within a given customer's service (such as an E-Line), voice, video and HTTP traffic can be differentiated and treated appropriately, given the application need for end-to-end delay or bandwidth capacity, for example (Figure 7).

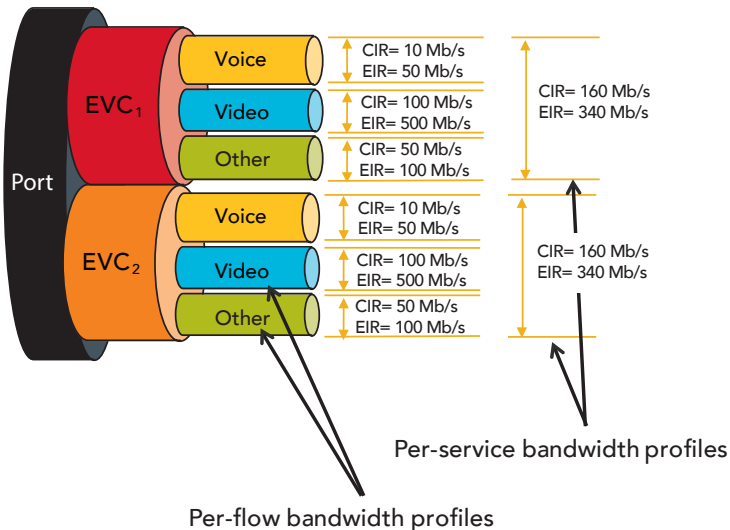


Figure 7: Ethernet Service Bandwidth Profiles

Since connection-oriented Ethernet can be used to specify the path a frame takes across the network, service providers can traffic-engineer their networks (or reserve bandwidth and queue resources). MPLS-TP and PBB-TE deliver 'hard QoS,' meeting bandwidth reservation and customer SLAs without overprovisioning the network capacity.

Making It All Scale

These various tools and techniques allow the CE provider to build a solid infrastructure that is efficient, resilient, and manageable. But will it scale to thousands of service instances across great distances? And, as bandwidth capacity continues to grow, will the network be able to keep up?

Most of the scalability challenges at the protocol level have been addressed with techniques such as QinQ, PBB, hierarchical QoS, and connection oriented Ethernet techniques. CE does not suffer from an address-space limitation (which can be in the many millions), and are rarely the limiting factors with modern silicon techniques. CPU-intensive functions (table lookups, address translations) can be a limiting factor, but modern scaling limitations have to do with operational procedures and human factors that could prevent the rapid deployment of CE services. For this challenge, a new generation of networks are emerging that use automation, scripted procedures, and software-based intelligent policies for service activation and ongoing management.

Techniques such as Zero Touch Provisioning, for example, allow thousands of switches and their associated services to be deployed in a short period of time often avoiding contact by field personnel via expensive truck rolls. The procedure enables a switch to identify itself to the network upon power-up, download its latest configurations and service attributes from a secure server, and bring itself online in a matter of seconds.

SDN has added to the list of options of services possible in a CE network by creating services on demand and service levels that match the dynamic nature of modern cloud-based applications. Bandwidth on demand, route optimization, and network 'defragmentation' are examples of these automated and creative service constructs. Each avoids a significant operational challenge for service providers and, hence, obstacles to scaling networks in the future.

CARRIER ETHERNET APPLICATIONS AND USE CASES

Nothing speaks to the success of CE like its continued adoption rate in the market. A few fictitious, but not unusual, sample business cases follow. The proposed CE-based solutions are also examples only.

Bank of Ether: Financial/Banking Enterprise

Business Situation:

A bank—let’s call it Bank of Ether—manages two private data centers and a growing number of bank branches in a large metropolitan market. Bank of Ether requires a large connection between the primary and secondary data centers to exchange data and maintain duplicate images of customer records in an active-active database application. In addition, the company requires branch-level access to the data centers in which application servers are housed to manage day-to-day banking activity.




Figure 8: Bank of Ether Situation Analysis

Table 2 suggests concerns for the Bank of Ether's IT director.

Requirement	Details
Physical Access/ Locations	<ul style="list-style-type: none"> • HQ campus with ~500 corporate employees • 20 regionally distributed branch offices (~200 mile radius from HQ) • Centralized data-processing center at HQ campus (primary data center) with redundant mirroring site (secondary data center) ~50 miles from HQ • Current connectivity via T3/SDH infrastructure is cost-limited to <150 Mb/s
Applications Support	<ul style="list-style-type: none"> • Database access for real-time financial transactions/ reporting (high-speed daily banking transactions, external handoff to Federal Reserve, monthly/quarterly reporting use); applications requiring low-latency data transmission • Third-party Web-hosting infrastructure for bank client account access applications • Internal HR/training network requires heavy video use, including support of video conferencing and distance learning applications at all branches • Operations/back-office secure infrastructure (file exchange, internal reporting)
Security	<ul style="list-style-type: none"> • As with all financial institutions, user and corporate financial data security is a high priority; primary concerns include: <ul style="list-style-type: none"> ○ Loss/theft of financial/user data or assets ○ Protection from internal loss/theft ○ Intentional or unintentional data access interruption
Growth Potential	<ul style="list-style-type: none"> • Historically an acquisition-intensive business; added branches and developer sites likely

Table 2: Concerns for the Bank of Ether's IT director.

Bank of Ether requires an advanced network to support an ambitious set of demanding applications. However, its IT department is a small, perennially budget-constrained operation, in which low-touch operations are a high priority. CE services, as well as managed services (such as VoIP support, firewall maintenance, and mail server support) are frequently of great interest in such scenarios because they help outsource many functions that might otherwise overwhelm the IT department. This example will focus on the network connectivity services that keep the branches, HQ (and primary data center), and secondary data center online and available 24/7/365.

Of course, among the most critical considerations is service availability. This example assumes there is at least one provider able to support most, if not all, of the physical locations required. Although this may not always be the case, there are options to employ legacy services (TDM-based) as well as DSL- or HFC-based services in real-world scenarios, but it is beyond the scope of this document to cover all possible variants. CE services are defined for any physical media or transport technology.

Table 3 shows key buyer considerations.

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Bandwidth	4	<ul style="list-style-type: none"> Increasing use of video applications at branches suggests branch connectivity in the 100-200 Mb/s range. Database synchronization between HQ and offsite data centers depends on the size of datasets involved (assume 5 Gb/s). Application development and testing can drive a requirement for >1 Gb/s.

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Performance/ QoS	5	<ul style="list-style-type: none"> • Branch office connectivity requires multiservice (ranging from database interconnect to best-effort branch Internet access) and multi-QoS support with separated traffic types (security). • Latency is critical for real-time database applications.
Reliability	4	<ul style="list-style-type: none"> • The main link between HQ and the secondary data center should be a fully protected link, diversely routed. • Branch connectivity should be highly reliable with protected links, but subject to cost considerations. Critical applications could be protected via legacy TDM services.
Reach/Availability	3	<ul style="list-style-type: none"> • Services are needed in all physical locations. Branches may be served with copper access, but growth potential (both in bandwidth and physical locations) makes fiber or wireless last-mile access preferable.
Security, Governance, Regulatory Compliance	5	<ul style="list-style-type: none"> • Although end-user and application-level security are critical, the network can provide a layer of traffic separation to minimize the effect of attacks, loss of sensitive data, and unauthorized intrusion (insiders or outsiders).

Table 3: Key Buyer Considerations (continued on the following page)

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Cost	3	<ul style="list-style-type: none"> • Cost is always a consideration, but Bank of Ether tends to place high importance on performance and reliability over cost.
Ease of Operation	3	<ul style="list-style-type: none"> • Bank of Ether wants to provide its customers with reliable and convenient financial services, not become an IT department.

Table 3: Key Buyer Considerations

Bank of Ether has several CE service options at its disposal. In this case, the bank chose (as shown in Figure 9) an Ethernet Private Line (EPL) solution with 10 Gb/s links to connect the primary and secondary data centers, given the critical nature of the data exchanged between them and the desire to future-proof the network against the bank’s growth plans. A fully protected, diverse router service is recommended.

The branches are connected via an E-Tree service with 20 Mb/s CIR and 200 Mb/s Excess Information Rate (EIR) EVCs connecting to the HQ router. Up to four CoS mappings are available, depending on the application supported (ranging from Internet access to video conferencing, and critical bank database access). Lesser-quality CoS EVCs are chosen for the less-critical applications to keep network expenses down.

The bank selected an E-Line Virtual Private Line (EVPL) service configured for 500 Mb/s CIR and 1GbE EIR to connect the corporate firewall/router at HQ with the Internet and the bank’s chosen ISP. Internet traffic from the branches and HQ is funneled through this link, along with several managed services the bank arranges with outside providers (such as mail servers and VoIP service). Each of these can be segregated within the enterprise via VLANs and mapped to individual EVCs by the service provider if required.

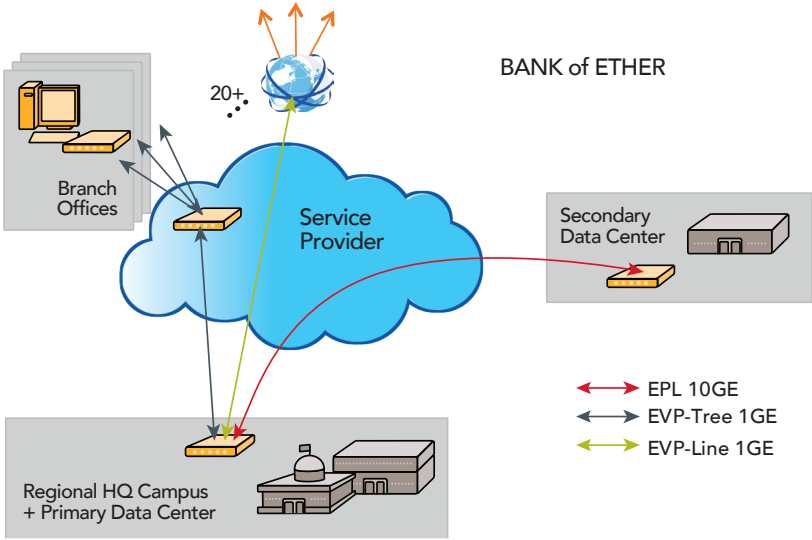


Figure 9: Bank of Ether Proposed Solution

A final consideration for Bank of Ether is the increased use of creative network services provided by managed service providers that bundle virtual functions with connectivity services. For example, the firewall managed by Bank of Ether’s IT department at HQ could easily be included in a managed network service, avoiding an additional appliance and the associated costs of managing its day-to-day operations, periodic upgrades, licensing costs, and so on. The approach to cost savings is facilitated by an efficient and high-performance CE network that avoids every-other-year network upgrades that might otherwise be required to chase the latest application. Dynamic bandwidth allocation to the various locations is another popular approach to cover the predictable and periodic high-bandwidth demand (such as nightly backups or quarterly all-hands video conferences).

BioEther Healthcare: A Hospital/Clinic Multisite Enterprise

BioEther Healthcare, which is a rural provider of healthcare services across a large and distributed area, consists of 10 small-town and neighborhood clinics that host two to four practitioners each, and two hospital facilities in regional population centers that host several dozen practitioners. The service area extends several hundred miles from the main hospitals, which are 150 miles apart (Figure 10).

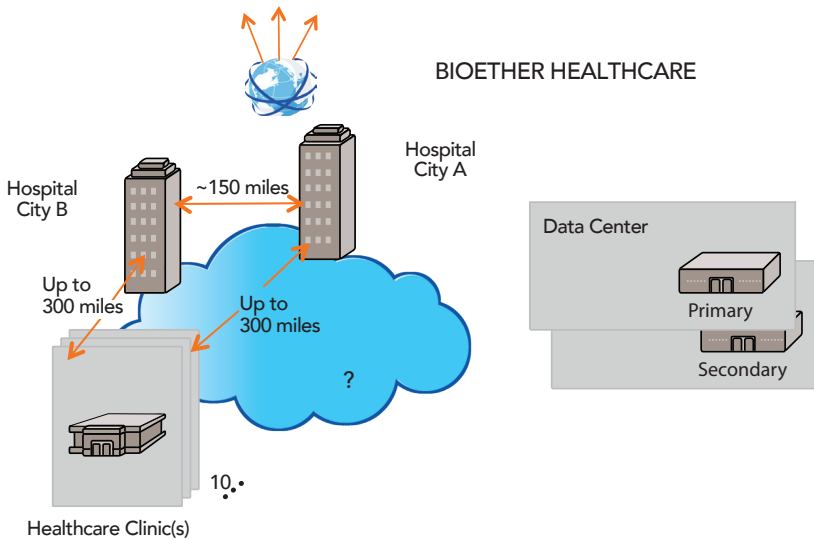


Figure 10: BioEther Healthcare Situation Analysis

BioEther has four IT department employees who handle day-to-day operations, from PC/software/printer management to procurement of IT applications. Specialized medical equipment is supported by service contracts with the suppliers. The company's cloud-based applications include healthcare data records management (including patient records), billing systems, hospital/clinic management (procurement, etc.), unified communications (VoIP, email, video conferencing), and legal/regulatory compliance applications. These applications are hosted in a private cloud—operated by a third party on behalf of the company—in two redundant data center facilities.

In addition, each clinic relies on high-bandwidth connections to the main hospitals for high-definition image/scan exchanges and real-time specialist provider consultations. Table 4 summarizes BioEther's requirements.

Requirement	Details
Physical Access/Locations	<ul style="list-style-type: none"> • Two large hospitals that operate 24/7, with ~250 employees and 30 doctors/specialists. • Ten clinics, distributed throughout the service, have two to four doctors each and ~10 support staff (nurses, admins, etc.). They operate weekdays from 7am to 5pm. • Current clinic connectivity via TDM infrastructure is limited to < 10 Mb/s • Connectivity between clinics is not seen as a high priority. • Reliable connection to the cloud-based, third-party managed data centers is a high priority.
Applications Support	<ul style="list-style-type: none"> • Most applications are outsourced to third-party providers who host a variety of healthcare management systems used for patient record maintenance, procurement, billing, back-office systems, and so on. They have varying bandwidth requirements (<100 Mb/s each), but high reliability and security requirements. • Internal HR/training network requires heavy video use between clinics and hospitals during daylight hours. • Operations/back-office requires a secure infrastructure (file exchange, internal reporting, etc.).

Table 4: BioEther's Requirements (continued on the following page)

Requirement	Details
Security	<ul style="list-style-type: none"> • As with all healthcare institutions, patient data security is a high priority. Primary concerns include: <ul style="list-style-type: none"> ○ Loss/theft of financial/user data or assets ○ Regulatory compliance ○ Data exchange with partner organizations (diagnostic labs, insurance companies, etc.) ○ Intentional or unintentional data access interruption
Growth Potential	<ul style="list-style-type: none"> • While the region is experiencing some growth in the two main population centers, the rural service area has been stable.

Table 4: BioEther's Requirements

BioEther's critical requirements appear to be twofold: reliable and high-performance access between the clinics and hospitals, and data security for patient healthcare records stored with the cloud service provider. Reliable connectivity to the hospitals may be mission-critical because life-saving consultations with specialists would otherwise require travel that could take hours.

As evidenced by its use of outsourced applications hosting and support, the IT department is lean and would be limited in its capability to manage a complex networking environment. Hence, Carrier-managed services are of great interest.

Given the rural nature of the required connections, the earlier assumption about service availability may not be practical, and multiple operators may need to be engaged in different portions of the coverage area to secure the desired connectivity. BioEther may also need to avail itself of legacy TDM-based services as well as DSL- or HFC-based services in some clinic situations.

Table 5 lists several buyer considerations.

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Bandwidth	5	<ul style="list-style-type: none"> • Video exchanges between clinics and hospitals are critical to the company's ability to provide quality care. These links should therefore accommodate bandwidth on the order of 100 Mb/s. • It is important that the clinics be able to rapidly exchange very large images/scans with the hospitals (in seconds, not hours). Because these files can be several GBs in size, a 1 Gb/s capability is recommended.
Bandwidth	5	<ul style="list-style-type: none"> • Connections between the hospitals and the cloud service provider data center require less bandwidth, but must take into account that the clinics derive their access from these same links (via the hospitals).
Performance/ QoS	3	<ul style="list-style-type: none"> • Clinic connectivity requires multiservice support (ranging from HD video for training/consultations to unified communications and best-effort internet access). Multi-QoS support with separated traffic types (security) is necessary. • Latency is important for video and scan/image exchanges.

Table 5: Key Buyer Considerations (continued on the following page)

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Reliability	5	<ul style="list-style-type: none"> • Because connectivity to the hospitals is of potentially life-saving importance, dual connections (one to each hospital) should be considered. • Hosted applications, while not as critical, are nonetheless important to the overall functioning of the hospital system.
Reach/ Availability	4	<ul style="list-style-type: none"> • Services are needed in all physical locations. Clinics may need to avail themselves of whatever physical connectivity is available, given the rural nature of the service area. As with other instances, fiber or wireless last-mile access is preferable.
Security, Governance, Regulatory Compliance	5	<ul style="list-style-type: none"> • Regulatory compliance is a major consideration in all healthcare applications. • Data security and protection from outside intrusion/interruption are critical.
Cost	3	<ul style="list-style-type: none"> • Although cost is always a consideration, BioEther tends to place more importance on performance and reliability.
Ease of Operation	5	<ul style="list-style-type: none"> • Limited IT resources mean the company is keen to outsource as many of its hosting and networking functions as possible.

Table 5: Key Buyer Considerations

BioEther Healthcare has several CE service options at its disposal. One approach would be to employ EVPL services with multi-QoS application mappings from each clinic to both hospitals in a hub-and-spoke arrangement. (Note that an operator may be able to leverage ring-based topologies, if available, to provide this redundant connectivity, but this is an implementation detail that is largely transparent to the client). 100 Mb/s CIR with 1 Gb/s EIR is thought to be sufficient to accommodate the variety of applications required, up to and including the large file exchanges described previously (Figure 11).

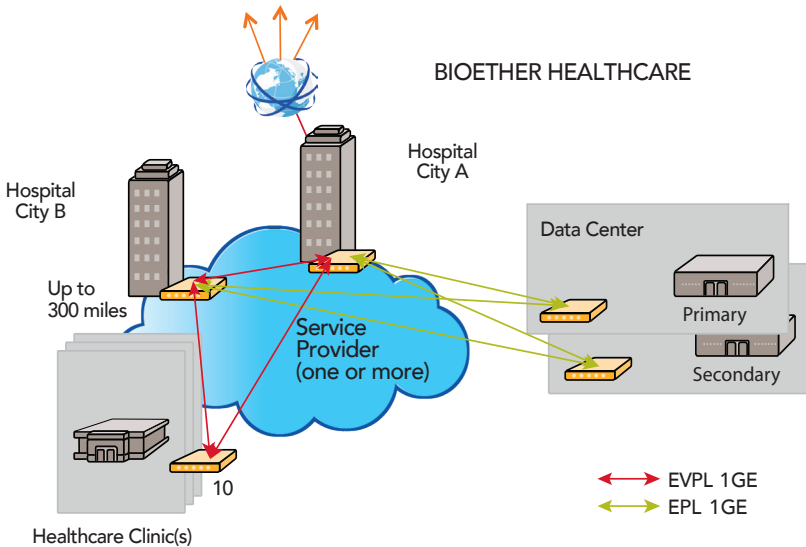


Figure 11: BioEther Healthcare Recommended Solution

Alternatively, an E-Tree service might be feasible for clinic-to-hospital communications, but would require a single operator to provide the service to all (or most) of the clinics. This may or may not be feasible, given operator service areas. It is possible that a combination of E-Tree and E-Access services (for last-mile connectivity between two operators across an E-NNI interface) might be possible, again depending on coverage areas. Details of costing and service availability would probably drive this decision.

For lower-QoS service types (Internet access, some hosted applications), it may be more cost-effective to connect individual clinics to the Internet or cloud-hosted applications directly (instead of routing them back through the main hospitals). This decision is also cost-driven, but may be easier to manage on a day-to-day basis via the single connection back to the hospitals.

Connections between the hospitals and the cloud service provider locations are the critical links to day-to-day operations of both hospitals and clinics. Therefore, a fully protected, dual 1 Gb/s (perhaps via LACP or ring-based G.8032) is advisable. Though not explicitly stated, connectivity between the two hospitals also seems to be important, for which an EVPL with multi-QoS support on the order of 1 Gb/s is reasonable (Figure 10).

Ether State University: A Growing Multisite Education Institution

Ether State University is a growing university system with distributed campuses across a major metropolitan area. Four satellite facilities (with one or two buildings and a few dozen employees serving several hundred students) and one large campus (dozens of buildings and hundreds of employees serving thousands of students) require connection to a statewide university learning and research consortium (Figure 12: Ether State University Situation Analysis2).

Cloud-hosted applications include various infrastructure systems (employee/student records, admissions, finance, procurement, research databases, employee file storage, and application hosting). Various high-performance computing projects are accessed via the university system's large supercomputer lab in a neighboring city.

Video instruction between the satellite locations and the main campus is an area of growing focus for both day and evening programs. Students may take classes taught locally at one of the satellite sites transmitted live from the main campus or streamed on a time-delayed basis from a cloud-hosted learning solution.

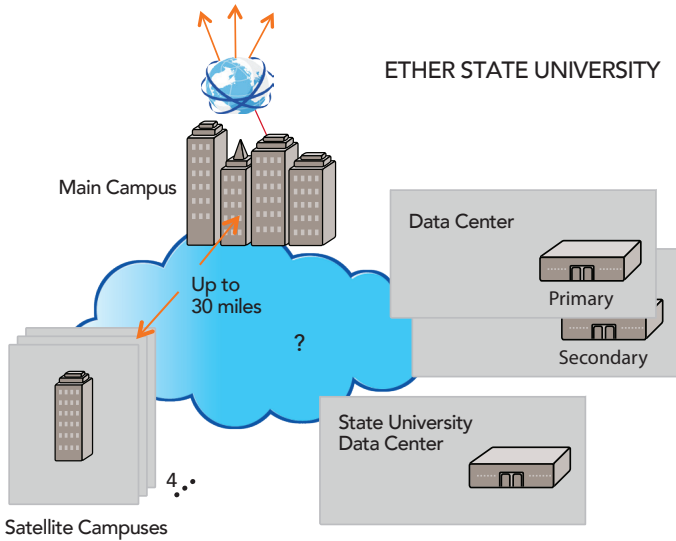


Figure 12: Ether State University Situation Analysis

Student Internet access from the classrooms, libraries, labs, and other study areas puts increasing demands on the current infrastructure, given heavy file sharing and video applications across a growing number of devices, many of which are mobile. These are supported across the campus-wide WiFi infrastructure.

Table 6 summarizes Ether State University's requirements.

Requirement	Details
Physical Access/ Locations	<ul style="list-style-type: none"> • One large campus environment supporting multiple buildings, labs and libraries, in addition to a large WiFi-based access infrastructure. Hundreds of employees and thousands of students use the system daily. • Four satellite learning centers used for remote learning for both full- and part-time students. All satellites are within 30 miles of the main campus. Each supports a few local administrative employees and visiting teaching staff, but most instruction occurs remotely from the main campus. Such instruction is often broadcast from the main campus to multiple satellite facilities. The satellites rarely require connectivity to one another. • Reliable connection to the university data center in a nearby city (<100 miles) is a high priority because research projects tend to require high-bandwidth, low-latency connectivity. This location is outside of the footprint of any operator that serves the main campus and its satellites. The service offer has to rely on multiple operators to comprise the end-to-end service.
Applications Support	<ul style="list-style-type: none"> • Distance learning is the principal application requiring main campus-to-satellite campus connectivity. The university provides both live and on-demand instruction, and courses are stored in a public cloud data center managed by a third party specializing in distance learning. • Most applications are also outsourced to third-party providers who host a variety of student records, email, VoIP, and university administration functions; along with general file storage, servers, and so on. Growing demands for Internet access from student and faculty mobile devices are increasingly stressing the university network.

Requirement	Details
Security	<ul style="list-style-type: none"> • Secure data handling for some applications is important. Primary concerns include: <ul style="list-style-type: none"> ○ Loss/theft of student data or assets ○ Attempts, malicious or otherwise, to hack into the university infrastructure ○ Intentional or unintentional data access interruption
Growth Potential	<ul style="list-style-type: none"> • The university is constantly expanding its programs throughout the area and beyond (including proposed international satellite campuses).

Table 6: Required Attributes for Ether State University

Implementation Considerations:

Ether State University’s main requirements are the efficient distribution of video for distance learning applications and reliable connectivity to its data center for research projects. Distance learning may be carried out live or on a time-delayed basis.

While the university IT department is sizeable, most of its attention is focused on maintaining the physical infrastructure of the various campuses and supporting the various day-to-day demands of the students, faculty, administration, and research staff. The university’s preference is to outsource management of services to maintain a lean and predictable IT budget. [Note that we could have chosen to discuss a “build your own” approach for such a network, but here, we are choosing and “outsource the network” approach. Either is possible and feasible with CE]

Ethernet services are available from multiple providers across the metropolitan area and beyond. Table 7 outlines several buyer considerations.

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Bandwidth	4	<ul style="list-style-type: none"> • Video exchanges between satellites and the main campus are important to provide distance learning. The links should be capable of two-way live interactions, as well as one-way streamed video transmission. These links should therefore accommodate bandwidth on the order of 100 Mb/s. • Bandwidth between the main campus and the state university system data center is important for research purposes, requiring 10 Gb/s connectivity. • Growing student and staff Internet access requirements should accommodate at least 1 Gb/s of traffic.
Performance/ QoS	3	<ul style="list-style-type: none"> • Multi-QoS support with separated traffic types is viewed as important to maximize the quality of instruction video applications (as compared with students streaming videos from the Internet, for instance). • Bandwidth and latency are important for both distance learning video applications and research applications hosted in the state university system data center.

Buyer Considerations	Importance (1-5, Low-High)	Recommendation Rationale
Reliability	3	<ul style="list-style-type: none"> • Interruptions to the distance learning or university applications access can be inconvenient and result in significant customer dissatisfaction. • Interruptions in research applications communications can be costly to both the university and the statewide university system.
Reach/ Availability	2	<ul style="list-style-type: none"> • Given the good coverage in the metropolitan area from a variety of providers, reach and availability should not be a challenge.
Security, Governance, Regulatory Compliance	4	<ul style="list-style-type: none"> • Data security and protection from outside intrusion/interruption are important.
Cost	3	<ul style="list-style-type: none"> • The university wants to outsource cloud-based applications wherever possible.
Ease of Operation	3	<ul style="list-style-type: none"> • Given a choice, university administration has preferred to outsource as much of its hosting and networking functions as possible.

Table 7: Key Buyer Considerations

Chosen Solution: Given the small number of sites involved and the need for application flexibility, a port-based E-LAN solution should be workable here. Perhaps the state university data center is connected to the service provider datacenters via separate EPL, EVPL, or E-LAN connections so, that the Cloud Data Center is also acting as the firewall for IN and OUT content. Providers could also offer Network-as-a-Service solutions, including managed functions such as encryption, firewalls, and WAN optimizers, because the IT department seems keen on outsourcing labor-intensive functions.

It may be prudent to design this network with the future in mind; student-generated, WiFi, and mobile-based video content is predicted to grow over the lifetime of the network. GE connections with separate QoS for institutional traffic (versus student-generated loads) should be used. This solution would incorporate individual 1GE service to each satellite campus. The state university data center is also likely to house increasing amounts of application data vital to the research departments within the university, so a 10GE link is recommended to be shared among the various departments, (Figure 13).

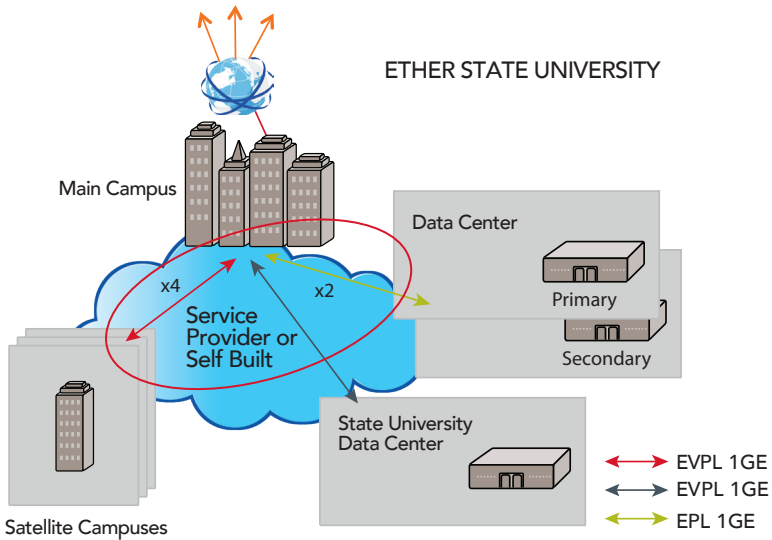


Figure 13: Ether State University Recommended Solution

If a provider has a ring-based solution (which is dependent on the provider's physical infrastructure and topology), a shared G.8032 solution would provide significant benefit in terms of quick restoration in case of failure and the ability to share fiber capacity among the various sites in a dynamic fashion. Enterprises may not care to specify this implementation detail, but it is listed as an example of the flexibility of infrastructure and topology types.

CARRIER ETHERNET FUTURES

Already, the IEEE is working on future data rates to support the needs of data center, enterprise, and operator environments. Its roadmap includes data rates driven by market needs to support ever-increasing uptake of video-based and

cloud-hosted applications. Figure 14 shows the view from the Ethernet Alliance (a key supporter of IEEE standards efforts). Note that several “non-traditional” rates are envisioned (ie not multiples of 10, per Ethernet tradition). These are targeted at specific scenarios and cost points driven by the increasing use within data center and enterprise scenarios.




Figure 14: Ethernet Alliance View of the Ethernet Roadmap

Related Ethernet Developments

While this book has concentrated on CE developments, it is important to note that Ethernet has evolved to cover a variety of areas, from audio-visual to industrial and military applications. These applications often do not involve changes to the protocol but may add specific capabilities that go beyond basic network connectivity provided by the base standards. For enterprises and service providers two Ethernet extensions are particularly useful: precise timing distribution and power over Ethernet.

Precise Timing Distribution

Beginning with the transition from 2G to 3G mobile networks, CE became a viable and increasingly popular approach for connecting cell towers to mobile telephone switching offices. Earlier backhaul technologies provided both connectivity and precise timing references to the base stations, which in turn used the information to synchronize key functions such as call handoffs between towers and spectrum coordination between calls. SONET/SDH, for example, has historically been a popular backhaul technology; because it was a synchronous protocol, it could simply pass along the appropriate timing information to the cell towers for their synchronization purposes.

But with the transition to Ethernet (an asynchronous packet technology) for backhaul, the synchronization information has to be distributed by some other means. In some regions, GPS units have been used for this purpose, but this is a costly option, and operators have pursued alternative approaches. Synchronous Ethernet, or SyncE, is International Telecommunications Union-Telecommunication Standards Bureau (ITU-T) protocol that uses Ethernet's physical layer clocking to provide frequency synchronization. In more modern 4G networks, in which phase and time-of-day are also required, Precision Timing Protocol, or IEEE 1588v2, can be used. In reality, a combination of GPS, SyncE, and 1588v2 is not uncommon, depending on application and operator preferences.

Power over Ethernet

The practice of using an Ethernet cable to provide power to a peripheral device (such as printers, scanners) originated in the consumer environment, but soon found a use in enterprise and carrier networks as well. Power and data are carried in the same cable, thus avoiding a dedicated (and expensive) power source for cameras, wireless radios, and similar devices. IEEE 802.3at, commonly known as PoE+, provides for up to 25 watts of power at distances up to 100 meters.

Software-Defined Networks and Network Functions Virtualization

The advent of SDN has created new opportunities to rethink the way networks are operated, especially in data center contexts. Network operators have proposed new centralized control mechanisms (formerly carried out by distributed control planes such as STP) to more efficiently and cost-effectively

run their networks. Protocols such as OpenFlow enable forwarding tables within routers and switches to be programmed based on whatever algorithms these centralized algorithms deem optimal. Modeling languages such as YANG, used in combination with the NETCONF interface protocol, allow creation of open interfaces to hardware and software entities within the network so applications can readily observe and/or control them. This opens up the network to a level of innovation that promises to revolutionize the market in the next few years.

In addition, the concept of NFV, whereby formerly 'appliance'-based applications (such as routers, firewalls, and loadbalancers) can be implemented in software (or 'virtually') and deployed in a data center, saving the hardware, associated deployment and ongoing maintenance costs. Both enterprises and carriers are poised to leverage SDN and NFV to not only save money, but also innovate at an unprecedented level in the near future.

GLOSSARY/ACRONYMS SECTION

ATM:	Asynchronous Transfer Mode
b/s:	Bits per second
CE:	Carrier Ethernet
CEN:	Carrier Ethernet Network. A network that supports MEF services
CCM:	Continuity Check Message
CIR:	Committed Information Rate
CoS:	Class of Service
CRC:	Cyclical Redundancy Check
CSMA/CD:	Carrier sense, multiple access/collision detection
DA:	Destination Address
DEI:	Discard Eligibility Indicator
EIR:	Excess Information Rate
EVPL:	Ethernet Virtual Private Line service
EMS:	Element Management System
ENNI:	External Network-to-Network Interface
EPL:	Ethernet Private Line
EVC:	Ethernet Virtual Connection
FC:	Fibre Channel
FDV:	Frame Delay Variation (commonly "jitter")
FIB:	Forwarding Information Base
FLR:	Frame Loss Ratio
G.709:	ITU-T recommendation for interfaces for the Optical Transport Network
GbE:	Gigabit Ethernet (10 GbE = ten gigabit Ethernet, 100 GbE = hundred Gigabit Ethernet)
Gb/s:	Gigabits per second
GFP:	Generic Framing Procedure
IEEE:	Institute of Electrical and Electronics Engineers
IETF:	Internet Engineering Task Force
IP:	Internet Protocol
ITU-T:	International Telecommunications Union – Telecommunication Standardization Bureau
LACP:	Link Aggregation Protocol
LAN:	Local Area Network
LBM:	Loopback Message
LSP:	Label Switched Path
LTM:	Linked Trace Message
MAC:	Media Access Control

MAN:	Metropolitan Area Network
Mb/s:	Megabits per second
MEF:	Formerly "Metro Ethernet Forum" and today known as simply "MEF"; organization that originated the Carrier Ethernet trend and established its service specifications
MPLS:	Multi-Protocol Label Switching
NFV:	Network Functions Virtualization
NMS:	Network Management System
OAM:	Operations, Administration, and Maintenance
OC-n:	Optical Carrier Level n (1, 3, 12, 48, 192, 768)
ODU:	Optical Data Unit
Operator:	The entity that administers a Carrier Ethernet network
ODU:	Optical Data Unit
OTN:	Optical Transport Networking (see G.709)
OVC:	Operator Virtual Connection
P2P:	Point-to-Point
PBB:	Provider Backbone Bridging
PCP:	Priority Code Point
PDU:	Protocol Data Unit
PHY:	Physical
QoS:	Quality of service
RFC:	Request for Comment: IETF's designation for a standard
RFP:	Request for Proposal
SA:	Source Address
SDH:	Synchronous Digital Hierarchy
SDN:	Software-Defined Network
Service Provider:	Provides end-to-end service to end-users; might not be a network operator
SLA:	Service-Level Agreement
SONET:	Synchronous Optical Network
STP:	Spanning Tree Protocol
Tb/s:	Terabits per second
TDM:	Time-Division Multiplexing
TWAMP:	Two Way Active Measurement Protocol
UNI:	User Network Interface
VLAN:	Virtual Local Area Network
VoIP:	Voice over IP
WAN:	Wide Area Network



John Hawkins

Senior Advisor, Product &
Technical Marketing

John Hawkins is Senior Advisor for Product and Technical Marketing at Ciena where he supports the company's Packet Networking product line. Joining Ciena in 2009, John has held positions in business development and product management functions. He also led the business development and promotion of Ciena's E-Suite family of packet switch modules. John is active in networking industry conferences and standards development, and currently serves as co-chair of the MEF's Certification Committee.

Prior to joining Ciena, John was with Nortel where he was responsible for Carrier Ethernet product management. His projects included a number of innovative technology building blocks within Nortel's Optical Ethernet portfolio. John began his career at GE as an IC designer, and later product manager in the Aerospace Division.

With over 25 years of technical experience, John is a frequent contributor to the Metro Ethernet Forum and sought-after industry speaker. John holds a BSEE from North Carolina State, an MS in Telecommunications from Southern Methodist University, as well as an MBA from Duke University.

“Carrier Ethernet has fundamentally changed the way enterprises, government agencies, operators and service providers build networks. This guide provides an excellent introduction, not only to Carrier Ethernet as a technology, but to the account of just how far we’ve come and where we’re ultimately headed as an industry.”

Nan Chem, *President,*

MEF

As businesses grow reliant on cloud services, networks have become the critical lynchpin to enterprises, and reliance on the public Internet for access to data centers has been replaced with a hybrid (public and private) infrastructure. These trends are unlikely to subside as demand for ever-increasing bandwidth is expected to grow unabated.

At the same time, Ethernet has evolved to address these critical network requirements. In the early 2000s, Carrier Ethernet (CE) was introduced as a natural way to extend the Ethernet protocol from local area networks to deliver wide area network connectivity. Carrier Ethernet is now crucial to network services, including mobile backhaul, business and data-center interconnect services, working as a feature-rich solution for all types of network operators who provide Ethernet access to end-user customers.

This book offers an entry-level overview of Carrier Ethernet, intended for anyone seeking a better understanding of this core networking technology, along with insights into advances such as SDN and NFV, which impact future CE expansion.



**Ciena's
Essential Series:**

Be in the know.