Chapter

2

Carrier Ethernet

by Abdul Kasim

In order to leverage the potential of Ethernet beyond the LAN, it had to be augmented with additional "carrier-class" characteristics; identifying and formalizing these detailed characteristics was, therefore, essential to enabling this role for Ethernet. This chapter focuses specifically on standardization and other efforts underway to develop a foundation for transforming LAN Ethernet into a Service Provider—based offering, henceforth referred to as *Carrier Ethernet (services).* Carrier Ethernet delivered over Service Provider networks across the MAN and WAN optimally enables next-generation packet applications.

The first fundamental step is defining Carrier Ethernet, what it precisely means and understanding the rationale for this definition. Also as fundamental, is an established reference framework—the context in which this definition applies, and the necessary elements that make up this context. In so doing, a common and consistent understanding as well as a "language" to describe Carrier Ethernet services is provided; with this as the basis, the attributes are discussed in greater detail (note: in the context of this book, only a sufficient overview can be reasonably provided), with selective discussions in a few areas that are deemed especially critical to enabling Carrier Ethernet.

Most of the standardization effort, especially at the service-level, has been carried on by the Metro Ethernet Forum (MEF) and so expectedly, this chapter devotes a significant part to the MEF-initiated development; but efforts by other standards bodies are also identified. This chapter also attempts to incorporate some commercial developments enabling Carrier Ethernet. Often, forward-looking entities—whether Service Providers or equipment manufacturers—are ahead of the standards bodies in terms of recognizing and addressing the practical issues that usually emerge when offering new services. A look at these issues and their respective solutions in the marketplace serves, therefore, to provide a better understanding of the actual status quo in the field.

Defining Carrier Ethernet

Although numerous efforts, both informal and formal (standards-based), have been undertaken to make Ethernet more viable as a technology and service beyond the LAN, the MEF has been instrumental in initiating a substantial formal effort to define Carrier Ethernet services (delivered by Service Providers). This definition was a prerequisite to developing a common understanding and a common objective in the delivery of such services.

Among the first steps undertaken was to define more precisely what such Ethernet services would entail, since, as noted in the previous chapter and repeated in Table 2.1, there are fundamental differences in providing Ethernet in the Service Provider network (broadly referred to as Carrier Ethernet) as opposed to providing Ethernet in the LAN. The context in which Carrier Ethernet services are defined is, therefore, the Service Provider networks and the several types of services already being delivered over these

TABLE 2.1 Ethernet in the LAN Versus Ethernet in a Service Provide Network (Spanning the MAN and WAN)

Dimension	Local Area Network	Service Provider Network
Geography/Reach	Usually less than 1–2 km; deployed in building(s) and small campuses	10–100 km and longer; deployed in a metro area or even across distant metro areas
Service Provider	Enterprise (IT group); implemented by internal IT group.	Service Provider (Carrier typically); services offered commercially for an initial and recurring cost
User of service	Enterprise	Enterprise
Number of end users/points (Scale)	In the tens/hundreds	Thousands or tens/hundreds of thousands
Bandwidth	10M/100M/1000M	1M and greater—up to 10,000M; usually in granular increments of 1M Aggregation required
Services offered (scope)	Enterprise data applications	Voice/TDM and data connectivity applications such as Internet Access, intra-metro connectivity
Delivery of Ethernet services	Over coax (CAT 5) and fiber; Best effort	Over a host of media, incumbent transport technologies, and with an associated service-level agreement (SLA)
Tolerance to failures (resiliency)	Generally reasonable because network is usually intra-enterprise and over a smaller physical area so failures can be addressed relatively quickly	Very low tolerance because failures usually have a larger impact—often on revenues and competitiveness
Manageability	Manageability possible with fairly simple tools given fewer number of users and applications within a smaller physical area (typically a building or campus) and the relatively higher tolerance to failure issues	Scale and scope of the Service Provider network in terms of the number of users and the geographical footprint introduces significant complexity necessitating sophisticated management tools and capabilities

networks. In fact, *Carrier Ethernet essentially encompasses the deterministic and other service delivery aspects for standardized Ethernet services*. This point is key because it highlights the focus on standardized Ethernet services and the specific *characteristics* of such services and not necessarily the underlying transport infrastructure itself. So what is Carrier Ethernet?

Carrier Ethernet: A Formal Definition

The MEF[1] has defined Carrier Ethernet as the "ubiquitous, standardized, Carrier-class *service* defined by *five attributes* that distinguish Carrier Ethernet from the familiar LAN based Ethernet." As depicted in Figure 2.1, these five attributes, in no particular order, are

1. Standardized services
2. Scalability
3. Reliability
4. Quality of Service (QoS)
5. Service management

Carrier Ethernet essentially augments traditional Ethernet, optimized for LAN deployment, with Carrier-class capabilities which make it optimal for deployment in Service Provider Access / Metro Area Networks and beyond, to the Wide Area Network. And conversely, from an end-user (enterprise) standpoint, Carrier Ethernet is a service that not only provides a standard Ethernet (or for that matter, a standardized non-Ethernet[2]) handoff but also provides the robustness, deterministic performance, management, and flexibility expected of Carrier-class services.

Fundamental to both Carrier Ethernet and LAN Ethernet is the fact that data is carried in an *Ethernet frame.* What this means is, in effect, an Ethernet frame originating at a device in the LAN, now continues to traverse across one or more Service Provider networks,[3] largely unaltered, and terminates at a device in a remote LAN. *One way to look at this transformation is that it essentially creates one larger Ethernet, spanning LANs, MANs, and may be even the WAN, albeit delivered as a service to the customer.* This transformation is shown in Figure 2.2, courtesy of the MEF, and illustrates the remarkable potential of Carrier Ethernet. The terms *UNI* and *NNI* in the figure denote standardized interface hand-offs between the enterprise customer and

[1] MEF is the preeminent nonprofit industry body focused solely on enabling Carrier Ethernet. The "Metro" reference in MEF is now a misnomer, however, and does not accurately reflect its charter and focus, which has long extended beyond the metro.
[2] Because it can, as will be seen later, also support non-Ethernet services (albeit over an Ethernet layer).
[3] The Service Provider networks could encompass both the MAN and the WAN.

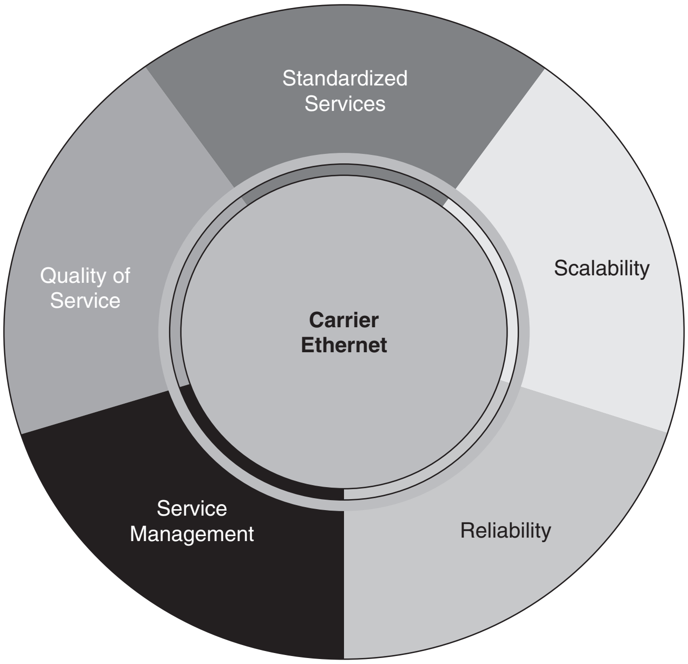

Figure 2.1 Attributes of Carrier Ethernet (Source: MEF)

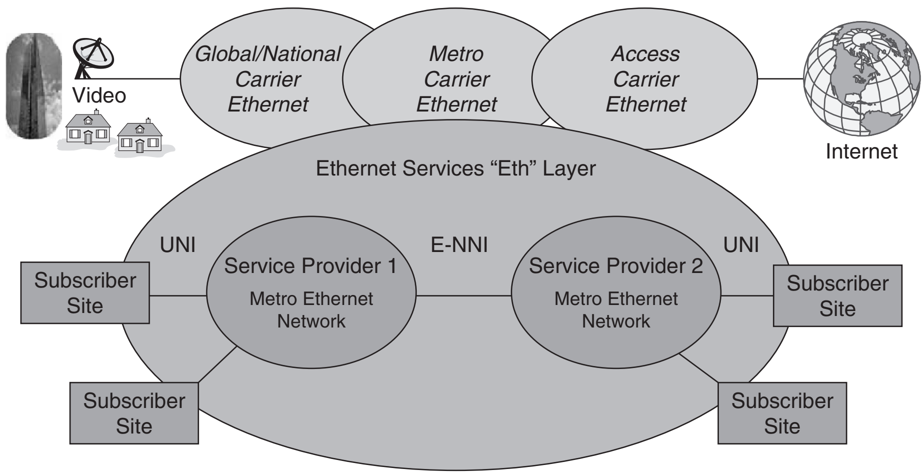

Figure 2.2 Carrier Ethernet spanning Access, Metro, and Wide Area Networks (Source: MEF)

the Service Provider network and between Service Providers (or Network Operators[4]), whose infrastructure is used to deliver the service, respectively, and are explained in more detail later in this chapter.

The Ethernet frame(s) may be transported as is, either natively and directly over a physical media or encapsulated and delivered over a variety of overlay networks built using different technologies. Each of these very different networking technology solutions, however, delivers[5] Carrier Ethernet services. It is critical to understand that the Carrier Ethernet attributes often manifest only partially in commercial solutions today because they exist at the network/transport/physical layers as opposed to the service layer[6]. This will become clear in rest of the Part II when the various commercial solutions currently employed to deliver Carrier Ethernet are discussed.

NOTE The focus in this book is primarily on delivering Carrier Ethernet services; the network and transport delivery infrastructure—the Carrier Ethernet solutions, provide the carrier-class attributes that enable commercial Carrier Ethernet services. Often, the term 'Carrier Ethernet' is interchangeably used to refer to both the Ethernet services and the underpinning enabling solution infrastructure.

The Carrier-class attributes are delivered differently by the various network solutions (for example, how reliability is offered in one solution versus another). This is largely a result of their respective geneses and subsequent evolution. It is important to also note that some of the Carrier Ethernet attributes in a solution existed pre-Carrier Ethernet (albeit at the transport layer and not at the service layer) and were, in fact, initial drivers for the use of respective solution. For example, SONET offered impressive resiliency to any failures in the fiber and/or equipment deployed in a ring topology, so it was adopted to support mission-critical voice services that required stringent SLAs.

Each of the Carrier Ethernet solutions and its respective evolution toward optimizing delivery in Service Provider networks is discussed in a fair amount of detail in Part II of the book.

Carrier Ethernet: The Attributes

The five attributes that define Carrier Ethernet essentially provide the additional capabilities necessary to use Ethernet in much the same way as the other preceding service provider technologies such as ATM and Frame Relay.[7] Each of these attributes is elaborated upon and its rationale highlighted in the sections that follow.

---

[4] A Network Operator is distinguished from a Service Provider by the fact that the former's infrastructure is employed in the delivery of Carrier Ethernet services; however, the service itself is commercially offered to the customers (usually on a subscription basis) by the Service Provider. Service Provider often lease infrastructure from network operators to deliver services.

[5] More accurately, as will be evident in Part II, the solutions strive to offer the attributes of Carrier Ethernet.

[6] Because at these lower layers inherently address only a subset of the higher-level service.

[7] Especially helpful today because Ethernet is largely being used as a substitute for Frame Relay and ATM.

Standardized Services This attribute essentially enables a Service Provider to deliver a host of both Packet and traditional TDM (see chapter 10 for more information on TDM) multi-point services in an efficient and deterministic manner over standardized equipment platforms. These services underpin the multitude of customer applications that are emerging across voice, data, and video. Specific components that define this attribute comprehensively are defined next.

- **Ubiquity** Carrier Ethernet enables ubiquitous Ethernet services provided via standardized equipment, independent of the underlying media and transport infrastructure. This is a critical prerequisite to extending Ethernet's appeal globally (similar to LAN Ethernet).
- **Ethernet Services** Carrier Ethernet supports two types of services: Point-to-Point (also referred to as *Ethernet Line* or *E-LINE*) and multipoint-to-multipoint Ethernet LAN (referred to as E-LAN) Ethernet services. These services are discussed in greater detail later in the chapter and are expected to provide the basis for all Ethernet services.
- **Circuit Emulation Services (CES)** Carrier Ethernet supports not only Ethernet-based services delivered across different transport technologies but also other (TDM) services transported over Carrier Ethernet itself. As noted previously, TDM services still remain an overwhelming contributor to Service Provider revenues and realistically need to be supported (and delivered over a converged Ethernet-based infrastructure). TDM-based voice applications especially need to be accommodated and characteristics of such applications such as synchronization and signaling need to be emulated.
- **Granularity and Quality of Services (QoS)** The services supported by Carrier Ethernet provide a wide choice and granularity of bandwidth and quality of service options. This flexibility is vital in Service Provider networks with its multitude of end users, each with slightly different application requirements and, typically, operating equipment from multiple vendors. QoS capability is crucial to enforcing the deterministic behavior of Carrier Ethernet.
- **Converged transport** Supports convergence of voice, data, and video services over a unified (Ethernet) transport and greatly simplifies the delivery, management, and addition of such services. Basically, all enterprise services and applications are now supported over a single Ethernet "pipe".

Scalability One fundamental difference between a LAN and a Service Provider network[8] is scale. In a Service Provider network, there are usually a hundredfold more end users and as a consequence, exponentially more connections for Ethernet-based

[8] Or multiple Service Provider or Network Operator networks, since several such entities could be involved in the delivery of an Ethernet service. A Network Operator owns the delivery infrastructure but may or may not be the one offering a service (or the Service Provider).

applications simply because it covers a larger geographic area. Carrier Ethernet solutions, therefore, scale across several dimensions simultaneously:

- **Users/endpoints** A Service Provider network supports hundreds of thousands of endpoints and millions of Ethernet users in an optimal fashion. Specifically, it supports the delivery of millions of Ethernet services with an appropriate level of performance or QoS.
- **Geographical reach** The services delivered can span access, metros, and beyond to encompass very large geographical distances and over a variety of infrastructures including Ethernet, WiFi, WiMax, TDM, SONET, and so on. As noted previously, the reach of such services can be augmented by employing multiple Service Providers' adjacent networks.
- **Applications** Current and emerging applications supporting a host of business, information, and entertainment applications and benefiting from the convergence of voice, data, and video. The landscape or breadth of application support is a vital driver for Carrier Ethernet.
- **Bandwidth** Bandwidth scales from 1M to 10G in granular increments of 1M, enabling a much more palatable solution to both the end user and Service Provider because end users only have to "pay for what is required" and Service Providers would possibly receive higher revenues.

As these dimensions scale collectively, they make for a formidable problem to deliver, isolate, troubleshoot, and in general, manage thousands of users and hundreds of thousands of services in a robust manner.

Reliability As Carrier Ethernet services are expected to support mission-critical applications on a wide scale, the ability to detect quickly and remotely any failures that may arise in the physical infrastructure or in the Ethernet services layer underlying these applications is essential. Specifically, the following aspects are addressed by Carrier Ethernet.

- **Service Resiliency** The impact of failures is localized and will not affect other customers and/or applications; Correlation among multiple errors will be quickly identified. Further, the process of troubleshooting and recovery from failures will be rapid and employ tools that will minimize operational expenditures for the Service Provider and any adverse impact on the end users.
- **Protection** Carrier Ethernet services provide an end-to-end service-level protection that encompasses protection against any failures in the underlying infrastructure employed in the delivery of the services. This means protection against failures in the end-to-end "path" of the service, as well as against any underlying physical link and node equipment failures.
- **Restoration** Carrier Ethernet provides similar or better recovery than SONET. The benchmark for resiliency in Service Provider networks has long been the

SONET sub-50 ms service restoration to support circuit-switched voice networks. As latency-sensitive voice and video applications are deployed over a Carrier Ethernet infrastructure, this SONET-like resiliency is a critical prerequisite. Techniques such as Spanning Tree Protocol (STP) and its variants, while feasible in the LAN, are simply not acceptable in large Service Provider networks because depending on the size and complexity of the network, recovery of failures employing these techniques takes in the range of several seconds to even minutes. Carrier Ethernet supports a host of latency-sensitive applications that are often critical to an enterprise (for instance, regular telephony services), and consequently offers better fault-tolerant and recovery mechanisms.

Quality of Service Providing Quality of Service (QoS) is necessary for Carrier Ethernet to be embraced as a substitute to ATM and Frame Relay and ultimately as a converged mechanism to deliver all services. QoS essentially conforms to a predefined level of performance expected by an application. As Carrier Ethernet supports delivery of critical enterprise applications that are commonly expected to adhere to certain performance levels, this QoS capability becomes essential.

The challenge to a Service Provider is significant given the fact that it has to *simultaneously* support individual QoS to typically thousands of applications and end users, using a limited set of resources (bandwidth, switching, and so on) whose availability varies with time.

Carrier Ethernet services providing QoS, encompass the following:

- **Performance Service Level Agreement (SLA)** There is the capability to provide the stringent end-to-end[9] SLAs necessary to provide a host of critical voice, video, and data services over a converged Ethernet infrastructure. Such SLAs are essential, and end users often demand them since they are already accustomed to such an assurance using the ATM, Frame Relay, or Private Line services, and it is only natural for them to expect the same of Ethernet services that support similar and next-generation applications.
- **SLA parameters** A set of configurable parameters allows a Service Provider to actually define the specific SLAs associated with a particular commercial service. These parameters provide significant latitude for defining numerous levels of service premiums. Further, these parameters although associated with a service, are enforced across the underlying infrastructure delivering that service.
- **Provisioning SLA** The QoS provides a hard performance guarantee based on the typical elements that define QoS in networks such as availability at a particular performance, packet loss, packet delay, and packet delay variation or jitter.

In a LAN with its abundant bandwidth and high performance, QoS is usually not an issue; the simple priority queuing capability using IEEE 802.1P/802.1D provides a “soft”

[9] End-to-end refers to the end points between which an Ethernet service is delivered.

QOS, but this is not sufficient in a Service Provider network, where with a multitude of users competing for shared resources (bandwidth, switching, and so on), complexity is at a totally different level. Different techniques, therefore, become necessary.

Service Management Managing a large number of customers stretched over a wider geographical area requires Service Providers to have a sophisticated capability for installing, troubleshooting, and upgrading Ethernet services cost effectively and quickly; engaging in a truck-roll each time there is an issue is simply cost-prohibitive and makes it infeasible to deliver Ethernet on a wider scale. Carrier Ethernet, in an attempt to address these issues, provides.

- **Unified management** This encompasses standardized vendor-independent capability to monitor, diagnose, and manage the delivery infrastructure. It is not unusual to deliver services across multiple Service Provider networks, each of which is often comprised of equipment from one or more manufacturers and is frequently subject to individual differences; hence, managing services across the different vendors' equipment using a common streamlined approach becomes paramount.
- **Carrier-class OAM** Carrier-class Operational, Administration, and Maintenance (OAM) capability that will integrate with existing Service Provider operational models. This covers a wide array of capabilities that enables life-cycle management at the service level. With Carrier Ethernet—based networks reaching tens of kilometers and thousands of subscribers, the need for sophisticated OAM features is apparent. Carrier Ethernet incorporates cutting-edge service creation and management techniques that exceed those of both enterprise Ethernet and the legacy telecom infrastructure.
- **Rapid Provisioning** The capability to provision new Ethernet services rapidly is a key departure from the long and protracted commissioning intervals for traditional TDM services. This capability translates into allowing granular increases in bandwidth to existing services; the addition of new services, each with a specific performance assurance (SLA); and the ability to enable these services remotely most of the time.

Carrier Ethernet leverages the established benefits of LAN Ethernet to the end users while simultaneously enabling Service Providers to offer a set of carrier-class attributes in a manner that is not only aligned with other services such as ATM, Frame Relay, and Private Line, but does so in a scalable, robust, and flexible manner that supports the next-generation of packet-based applications much more cost effectively . This ultimately translates into lower CAPital EXpenditures (CAPEX), lower OPerational EXpenditures (OPEX), and competitive positioning for Service Providers. Thus, Carrier Ethernet helps realize the compelling benefits to both end users and Service Providers as detailed in the Chapter 1.

Defining the attributes of Carrier Ethernet in greater detail and refining them further to be more relevant for next-generation applications is an ongoing effort; considerable progress has, however, been made.

Enabling Carrier Ethernet

Carrier Ethernet is increasingly being adopted by the Service Provider and enterprise end–user community not only as the default access solution (i.e., service connectivity is via Carrier Ethernet), but also one that is being employed end-to-end across the WAN. Service Provider networks are, in fact, evolving to deliver the consistent Carrier-class Ethernet services end users are coming to expect. Chapter 3 highlights the growing demand for Carrier Ethernet services worldwide.

Carrier Ethernet is, however, still a work in progress; in fact, it is still in its infancy and being more formally defined, refined, and continually augmented based on learning from real–life field deployments supporting emerging applications. If it is to achieve the success and dominance of its LAN variant, it has to not only incorporate these lessons rapidly in terms of new value-added features, but also standardize them.

Standardization of Carrier Ethernet is thus a key approach to enabling and, in fact, accelerating the deployment of Carrier Ethernet services.

Standards Bodies

There are several standard bodies that are involved, to varying degrees, in enabling Carrier Ethernet. These include the IEEE (primarily the 802 body), the Internet Engineering Task Force (IETF), the International Telecommunications Union (ITU), the Metro Ethernet Forum (MEF), and to a lesser degree, others such as the Tele Management Forum (TMF).

While the involvement of several bodies working in the same area may appear to be at cross purposes or at best, partially redundant and with the potential to introduce confusion, the reality has been different. These bodies have been—and are—working with a largely complementary focus, and where there has been some overlap, there has also been significant collaboration, with the net result actually expediting standardization efforts.

The IETF has traditionally had an IT orientation, while the ITU has focused on developing international standards to support the needs of national Service Providers (known as PTTs in most countries). The IEEE, of course, has focused on the 802 Ethernet standards at the physical and data-link layer. It is continuing its legacy work on Ethernet and extending it in two areas from the standpoint of Ethernet in the MAN and WAN: OAM and Architecture. The ITU is working across the spectrum, from service definition to service architecture to OAM and Ethernet interfaces.

These bodies were involved with LAN Ethernet and are now also focused on Carrier Ethernet given its role as a converged platform appealing to both Service Providers and end-user enterprises and spanning their traditional Service Provider and IT constituencies.

The Metro Ethernet Forum (MEF), unlike the others, was formed relatively recently (2001) and exclusively to advance the deployment of Carrier Ethernet. Consequently, it has been the most active body focused on enabling Carrier Ethernet as a well-defined service to support the next-generation of applications. And although the MEF's initial focus was the delivery of Carrier Ethernet in the metropolitan area (hence the "Metro" in MEF), it has now extended its charter well beyond and focuses on end-to-end Carrier Ethernet services spanning the MAN and the WAN. The MEF represents the first comprehensive effort to address all service delivery aspects as well as the testing necessary for confirmation. Figure 2.3 depicts the different MEF standards and their respective focus as of August 2007; these are continually being augmented as the MEF tackles new issues in its attempt to accelerate the deployment of Carrier Ethernet. While the MEF has a broader mandate than the other bodies (at least as far as Carrier Ethernet is concerned), it extensively builds and reuses the efforts of these bodies.

Figure 2.4 summarizes the different standards bodies' respective focus across four distinct areas with respect to Carrier Ethernet: Architecture, Services, Management, and Testing. The specific standards are identified in each of these areas; standards underway but not yet ratified are italicized.

It is clear that only the MEF is focused across the board in all four areas and is notably the only standards body testing and validating Carrier Ethernet. The IEEE 802 addresses some architectural aspects (in fact, it did so even pre-Carrier Ethernet) but has also added several new efforts. A key contribution has been in the area of Carrier Ethernet Management, especially link level and connectivity management. The ITU has been very active in the Architecture, Service, and Management areas; it has not only leveraged the efforts from the other standards bodies—for instance, the MEF for the Ethernet services definition, the IEEE for Management—but has also augmented it by, for instance, adding Performance Management to its Management standard to address the requirements of its constituency.

Specification	Scope
MEF 2	Requirements and Framework for Ethernet Service Protection
MEF 3	Circuit Emulation Service Definitions, Framework and Requirements in Metro Ethernet Networks
MEF 4	Metro Ethernet Network Architecture Framework Part 1: Generic framework
MEF 6	Metro Ethernet Services Definitions Phase 1
MEF 7	EMS-NMS Information Model
MEF 8	Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks
MEF 9	Abstract Test Suite for Ethernet Services at the UNI
MEF 10.1	Ethernet Services Attributes Phase 2
MEF 11	User Network Interface (UNI) Requirements and Framework
MEF 12	Metro Ethernet Network Architecture Framework Part 2: Ethernet Services Layer
MEF 13	User Network Interface (UNI) Type 1 Implementation Agreement
MEF 14	Abstract Test Suite for Ethernet Services at the UNI
MEF 15	Requirements for Management of Metro Ethernet Phase 1 Network Elements
MEF 16	Ethernet Local Management Interface
MEF 17	Service OAM Framework and Requirement
MEF 18	Abstract Test Suit for Circuit Emulation Services
MEF 19	Abstract Test Suit for UNI Type 1

Figure 2.3 MEF Standards specifications (Source: MEF)

Ethernet Standards Summary

Standards Body	Ethernet Services	Architecture/Control	Ethernet OAM	Ethernet Interfaces
IEEE	-	• 802.3 – MAC • 802.3ar – Congestion Management • 802.1D/Q – Bridges/VLAN • 802.17 - RPR • 802.1ad – Provider Bridges • .1ah – Provider Backbone Bridges (PBB) • .1ak – Multiple Registration Protocol • .1aj – Two Port MAC Relay • .1AE/af – MAC / Key Security • .1aq – Shortest Path Bridging • .1Qay – PBB – Traffic Engineering	• 802.3ah – EFM OAM • 802.1ag – CFM • 802.1AB - Discovery • 802.1ap – VLAN MIB	• 802.3 – PHYs • 802.3as - Frame Expansion
MEF	• MEF 10 – Service Attributes • MEF 3 – Circuit Emulation • MEF 6 – Service Definition • MEF 8 – PDH Emulation • MEF 9 – Test Suites • MEF 14 – Test Suites • Services Phase 2	• MEF 4 – Generic Architecture • MEF 2 – Protection Req & Framework • MEF 11 – UNI Req & Framework • MEF 12 - Layer Architecture	• MEF 7– EMS-NMS Info Model • MEF 15– NE Management Req • OAM Req & Framework • OAM Protocol – Phase 1 • Performance Monitoring	• MEF 13 - UNI Type 1 • MEF 16 – ELMI • E-NNI
ITU	• G.8011 – Services Framework • G.8011.1 – EPL Service • G.8011.2 – EVPL Service • G.asm – Service Mgmt Arch • G.smc – Service Mgmt Chnl	• G.8010 – Layer Architecture • G.8021 – Equipment Model • G.8010v2 – Layer Architecture • G.8021v2 – Equipment Model • Y.17ethmpls - ETH-MPLS Interwork	• Y.1730 – Ethernet OAM Req • Y.1731 – OAM Mechanisms • G.8031 – Protection • Y.17ethqos – QoS • Y.ethperf - Performance	• G.8012 – UNI/NNI • G.8012v2 – UNI/NNI
TMF	-	-	• TMF814 – EMS to NMS Model	-

Figure 2.4 Standards bodies and their respective areas of Carrier Ethernet focus (Source: MEF)

The detailed specifications are, of course, vastly outside the scope of this book but they are referenced in sufficient detail in the context of defining Carrier Ethernet services and its underlying five attributes. This represents the formalized (i.e., standardization) effort thus far toward enabling Carrier Ethernet. It must be noted that Carrier Ethernet standardization activity is relatively dynamic and frequently there continues to be new developments. It is, therefore, advisable to check the websites of the standards bodies (see bibliography) to get a sense of the latest progress.

A Service Architecture for Carrier Ethernet

Since Carrier Ethernet is essentially a commercial service offered by a Service Provider, it was vital to establish a clear and precise specification of what it entails. This was especially necessary because Ethernet in the LAN was not typically offered as a service but rather as a product/solution wherein the equipment was purchased, set up, and managed by the enterprise (IT group) itself. As such, there were generally no service-oriented expectations of the LAN Ethernet. Unfortunately, this was also largely the case with the Ethernet services that were initially (and to some extent are still being) offered by Service Providers. There were no formalized definitions and expectations of these services. The effort to change this only recently began in earnest and has been driven primarily by the MEF. Although still in the beginning stages, reasonably significant progress has been made.

Even before formalizing Carrier Ethernet services, however, it was necessary to establish a context for such services—a Service architecture—and to identify the necessary service components of such an architectural context. The MEF (and also the ITU) undertook this effort and developed a set of standard specifications for a generic Service Architecture that provides a common language for describing Ethernet services.

In MEF 6 and MEF 10.1, the MEF has established what an Ethernet service is, how a variety of subscriber services can be offered, and how these Ethernet services can be customized for certain performance and Service-Level Agreements (SLAs).

The MEF has also defined an overall framework to discuss Ethernet services—the Ethernet Service Model (ESM), which identifies the building blocks or service attributes of these services. (The Ethernet Service Model does not define the Ethernet service itself; this is done in an Ethernet Service Definition framework explained later.)

The Ethernet Service Model (ESM) The basic Service Provider architectural model defined by the MEF is shown in Figure 2.5. It has two main components:

- The Subscriber or customer equipment (CE)
- The Metro Ethernet Network (MEN) or more accurately, the Service Provider Ethernet Network (SEN)[10] This is owned/operated by a Service Provider.

Basically, the customer equipment is connected to the MEN [through a User Network Interface (UNI) which is explained in greater detail in the next section]. Any OSI layer 1 or 2 transport technology can be used as long as Ethernet frames are being handed off. The Subscriber or customer equipment is typically a router or a switch (an IEEE 802.1Q bridge).

A MEN itself consists of physical components (e.g., network elements, ports, etc.) and logical components (e.g., meters, policers, shapers, virtual switches, links, etc.). It can be owned and operated by multiple Service Providers and provides the underlying transport (SONET, WDM, RPR, etc.) to carry the Ethernet frames. It essentially connects geographically separated enterprise LANs across the MAN and WAN. The Carrier Ethernet service is actually provided by the Service Provider owning the MEN over an Ethernet Virtual Connection (or EVC, which is defined in a later section).

The MEF has more formally defined a three-layered model (also shown in Figure 2.5) for the MEN; the Application services (APP) layer supports end-user applications carried over Ethernet connectivity services provided at the Ethernet services (ETH) layer, and these connectivity services in turn are delivered over various transport/networking technologies in the Transport services (TRAN) layer. The key focus of the MEF and other standards bodies is the ETH layer; Carrier Ethernet is defined in this layer. The delivery of these Carrier Ethernet services can be over various media and the transport and networking technologies that make up the TRAN layer (the subject of Part II).

[10] Since the MEF has extended its focus beyond the metro and into the WAN, it is generally more accurate to label the Metro Ethernet Network (MEN) as the Service Provider Ethernet Network (SEN), which could support the MAN and/or WAN.

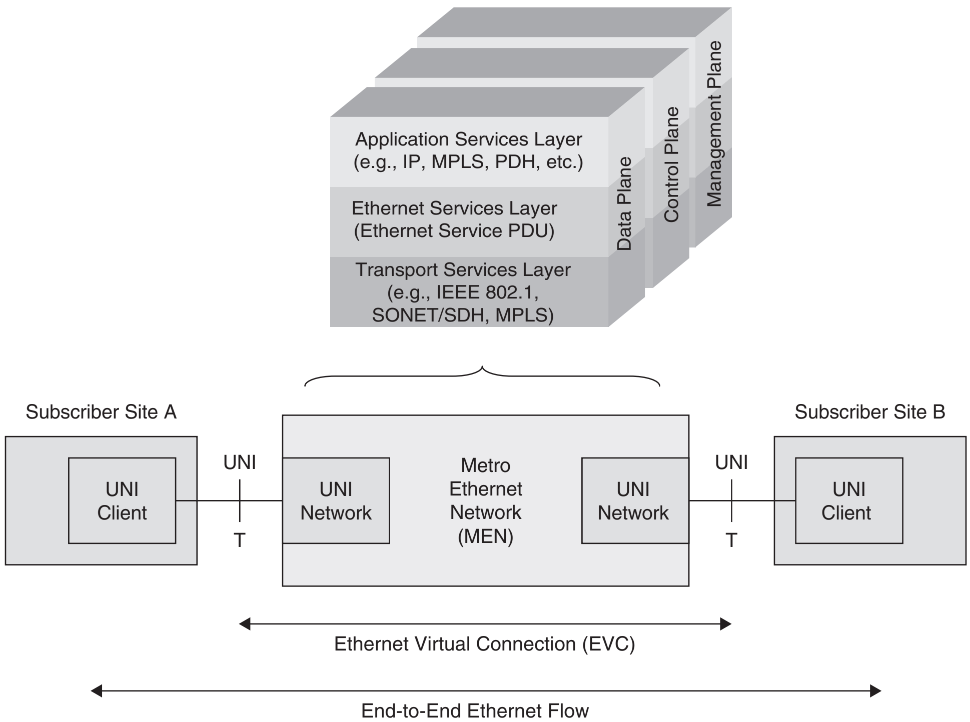

Figure 2.5 The basic Service Provider model for delivering Ethernet services (Source: MEF)

As will become evident in Part II, often the current—and evolving—attributes of Carrier Ethernet reside in the TRAN layer (depending on the specific technologies).

Each of the three layers has three associated operational planes: a Data plane, a Control plane, and a Management plane.

The Data plane, also referred to as the user/transport/forwarding plane, provides the functional elements required to steer the subscriber flow and supports the transport of subscriber traffic units among MEN Network Elements (NEs). The Control plane provides the functional elements that support distributed flow-management functions among NECs participating in the MEN data plane. The Control plane also provides the signaling mechanisms necessary to support distributed setup, supervision, and connection release operations, among other flow-control functions. The Management plane provides the functional elements that support Fault, Configuration (including flow and/or connection configuration), Account, Performance, and Security (FCAPS) functions, as well as any related Operations, Administration, and Maintenance (OAM) tools.

The three operational planes are generally well defined for the TRAN layer (numerous standards bodies have addressed it, and these are identified in Part II). For the ETH layer, the effort was, for the most part (except in the data plane), begun only recently. As will become evident in the rest of the book, *the control and management functions of the TRAN layer are often employed in delivering Carrier Ethernet currently*.

Ethernet services delivered over the MEN invariably have two key service attributes associated with them: the User Network Interface (UNI) and the Ethernet Virtual Connection (EVC).

User-Network Interface (UNI) The UNI is the interface used to interconnect a subscriber to an Ethernet Service Provider. The UNI also provides a reference point for demarcation between the MEN operator's (i.e., a Service Provider's) equipment that enables access to the MEN services and the subscriber access equipment. The demarcation point indicates the location where the responsibility of the Service Provider ends and where the responsibility of the subscriber begins. The UNI is a key Ethernet service attribute used to specify an Ethernet service.

Functionally, the UNI is an asymmetric, compound functional element that consists of a client side, referred to as the *UNI-C,* and a network side, referred to as the *UNI-N*. Thus, the term *UNI* is used to refer to these two functional elements and generically, to the data, management, and control plane functions associated with them.

UNI Client (UNI-C) The UNI-C represents all of the functions required to connect a subscriber to a MEN. Individual functions in a UNI-C are entirely in the subscriber domain, and may or may not be managed by the Service Provider/Network Operator. From the perspective of the MEN, the UNI-C supports the set of functions required to exchange data, control, and management plane information with the MEN subscriber. As such, the UNI-C includes functions associated with the Ethernet services infrastructure, the transport network infrastructure, and if present, application-specific components.

UNI Network (UNI-N) The UNI-N represents all of the functions required to connect a MEN to a MEN subscriber. The individual functions in a UNI-N are entirely in the Service Provider/Network Operator domain. From the perspective of the subscriber, the UNI-N supports the set of functions required to exchange data, control, and management plane information with the MEN. As such, the UNI-N includes functions associated with the Ethernet services infrastructure, the transport network infrastructure, and if present, application-specific components.

The MEF has defined a set of attributes to specify a UNI completely. These are listed at the end of the chapter (Figure 2.24).

Ethernet Virtual Connection (EVC) The *Ethernet Virtual Connection (EVC)* is a construct that performs two functions: One, it indicates the association of two or more UNIs for the purpose of delivering an Ethernet flow[11] between subscriber sites across the MEN. Two, an EVC prevents data transfers between subscriber sites that are not part of the

[11] An Ethernet flow represents a particular and potentially noncontiguous (e.g., consecutive Ethernet frames may belong to different flows) unidirectional stream of Ethernet frames that share a common treatment for the purpose of transfer steering across the MEN.

same EVC. The attributes associated with an EVC are shown in Figure 2.24 (at the end of the chapter) and are employed when specifying an Ethernet service.

NOTE There may be one or more subscriber flows mapped to a particular EVC. This capability enables an EVC to provide data privacy and security.

There are two basic rules that govern the delivery of Ethernet frames over an EVC. A service frame must never be delivered back to the UNI where it originated, and the Ethernet frame contents (including MAC addresses) must remain unchanged. The MEF has defined two types of EVCs: Point-to-Point or Multipoint-to-Multipoint. In a Point-to-Point EVC, exactly two UNIs must be associated with one another whereas in a Multipoint-to-Multipoint EVC, two or more UNIs must be associated with one another. Thus, an EVC can be used to construct a Layer 2 Private Line or a Layer 2 VPN[12] service.

Network to Network Interfaces (NNI) As noted in the reference Service Architecture, one or more Service Providers can be used to deliver Carrier Ethernet services. The demarcation or handoff between the Service Providers is referred to as the Network-to-Network Interfaces (NNIs). The MEF has defined several NNIs:

- **External Network-to-Network Interface (E-NNI)** An open interface used to interconnect two MEN Service Providers.
- **Internal Network-to-Network Interface (I-NNI)** An open interface used to interconnect network elements from a given MEN Service Provider.
- **Network Interworking Network-to-Network Interface (NI-NNI)** An open interface that supports the extension of transport facilities used to support Ethernet services and associated EVCs over an external transport network not directly involved in the end-to-end Ethernet service.
- **Service Interworking Network-to-Network Interface (SI-NNI)** An interface that supports the interworking of an MEF service with services provided via other service enabling technologies (e.g., Frame Relay, ATM, IP, etc.).

Defining Carrier Ethernet Services

Carrier Ethernet services are essentially connectivity services that employ Ethernet frames transported over the MEN using a host of different technologies such as SONET, WDM, MPLS, and so on. As shown in Figure 2.5, Ethernet services are delivered over an EVC provided by a Service Provider over a MEN, which is connected to the customer equipment (CE) via a standardized UNI. Thus, all Ethernet services will invariably have associated with them, one or more UNIs and one or more EVCs. The specific UNI and EVC attributes differentiate the specific services.

[12] Virtual Private Network (VPN) is a connectivity service between multiple points to multiple points.

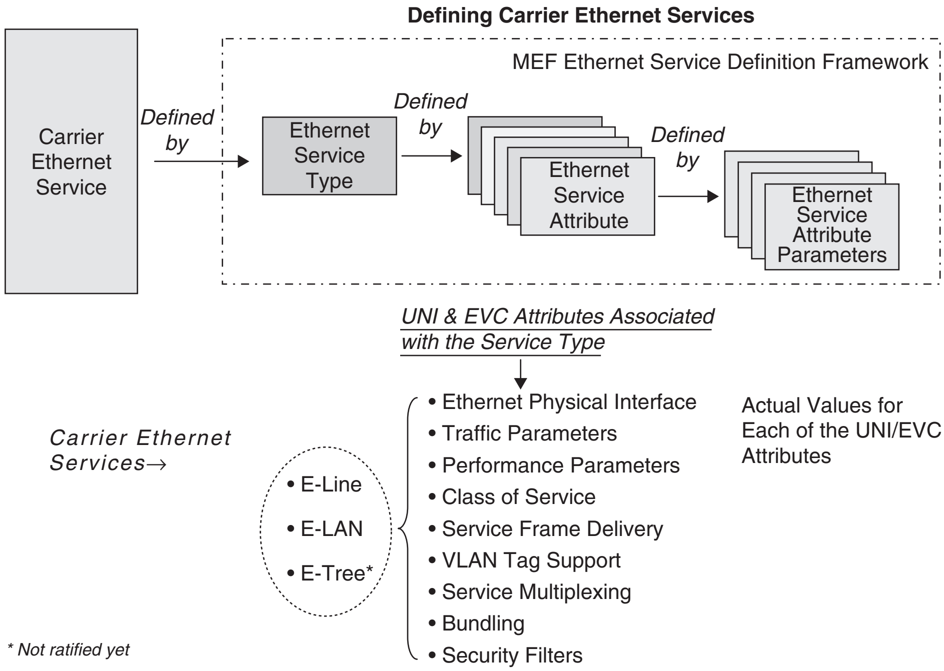

Figure 2.6 Defining Ethernet services

Carrier Ethernet services are defined from a subscriber perspective (and hence they're also referred to as "retail services"). As shown in Figure 2.6, the MEF has developed an Ethernet Services Definition Framework that defines any Carrier Ethernet service in terms of a predefined Ethernet service type. Each of these Ethernet service types (described next) are, in turn, defined by a set of Ethernet service attributes that define its capabilities. Some of these attributes apply to the UNI, others to the EVCs, and still others to both the UNI and EVCs associated with the service type. Specific parameters associated with each of these Ethernet service attributes ultimately define the Ethernet service fully.

This seemingly complicated approach is also illustrated in Figure 2.6, but it will become clearer when real-life examples are discussed later in the chapter. It is helpful to remember that every service is defined in terms of a service type and invariably has a set of UNI and EVC attributes[13] that will uniquely define it.

Before delving into the specific service types (which are defined in terms of the Ethernet service attributes), it is useful to understand these service attributes.

[13] Collectively referred to as the set of Ethernet service attributes.

Ethernet Service Attributes

The Ethernet service attributes are categorized into the following groups: Ethernet Physical interface, Traffic parameters, Performance parameters, Class of Service, Service frame delivery, VLAN tag support, Service Multiplexing, Bundling, and Security filters. Whether they apply to only the UNI or EVC or both is identified in the brief descriptions that follow.

Ethernet Physical Interface At the UNI, the Ethernet physical interface has several service attributes.

Physical Medium This UNI service attribute specifies the physical interface defined by the IEEE 802.3-2000 standard. Examples are 10BaseT, 100BaseSX, 1000BaseLX, and so on.

Speed This UNI service attribute specifies the standard Ethernet speed—either 10 Mbps, 100 Mbps, 1 Gbps, or 10 Gbps.

Mode This UNI service attribute specifies whether the UNI supports full or half duplex[14] and can provide auto-negotiation.

MAC Layer This UNI service attribute specifies which MAC layer is supported, i.e., as specified in the IEEE 802.3-2002.

Traffic Parameters/Bandwidth Profile The MEF has defined the Bandwidth Profile service attribute, which is associated with every Ethernet service and can be applied at the UNI or for an EVC. When there are multiple services associated with a UNI, there is a corresponding Bandwidth profile associated with each of these services.

A Bandwidth profile specifies a limit on the rate at which Ethernet frames can traverse the UNI associated with an Ethernet service. Bandwidth profiles enable both Service Providers and subscribers to optimize bandwidth and economics.

Service Providers have the ability to offer bandwidth in small increments and usually without having to add new physical interfaces. This means they can offer, engineer, and bill only the bandwidth needed by the subscriber for a specific service.

NOTE Multiple services can be offered over a subscriber UNI, and each of these services can have its own bandwidth profile.

Subscribers can purchase and pay for only the bandwidth they need. Furthermore, subscribers can be assured of a “committed” amount of bandwidth that meets certain performance objectives (usually specified in an SLA) and “excess” bandwidth that may not meet the SLA.

[14] Half duplex means transmission in one direction at any one time. Full duplex means transmission in both directions simultaneously; these are briefly discussed in Chapter 1.

Bandwidth Profile Traffic Parameters A Bandwidth profile associated with an Ethernet service consists of four traffic parameters: Committed Information Rate (CIR), Committed Burst Size (CBS), Excess Information Rate (EIR), and Excess Burst Size (EBS); in addition a service frame is associated with a Color Mode (CM). Together, these five parameters specify the bandwidth profile for a particular service:

Bandwidth Profile = <CIR, CBS, EIR, EBS, CM>

Committed Information Rate (CIR) CIR is the *average rate* up to which service frames are delivered as per the performance objectives (such as delay, loss, etc.) associated with the service; these service frames are referred to as being *CIR-conformant.*

The CIR value is always less than or equal to the UNI speed[15] and basically guarantees that the specified amount of bandwidth (or service frames) will be delivered according to a predetermined performance level. A CIR of zero indicates the service has neither bandwidth nor performance guarantees.

NOTE Independent of the CIR, the service frames are always sent at UNI speed.

Committed Burst Size (CBS) CBS is the limit on the maximum number, or bursts, of service frames in bytes allowed for incoming service frames so they are still CIR-conformant.

Excess Information Rate (EIR) The EIR specifies the average rate, greater or equal to the CIR, up to which service frames are admitted into the Service Provider network; these frames are said to be *EIR-conformant.* These frames are delivered without any performance guarantees and are not CIR-conformant; however, service frames that are not EIR-conformant are discarded.

Again, independent of the EIR, the service frames are always sent at the speed of the UNI (and hence, the EIR represents the average rate).

Excess Burst Size (EBS) The EBS is the limit on the maximum number, or bursts, of service frames in bytes allowed for incoming service frames so they are still EIR-conformant

Color Mode and Color Marking In addition to the bandwidth profile traffic parameters, there is also the concept of marking the service frames with a color. The color of a service frame is used to determine whether or not a particular service frame is in conformance with its bandwidth profile.

A service marked *green* is conformant with the CIR and CBS in the bandwidth profile. A green frame is always delivered per the performance SLA associated with the service.

[15] If multiple services are being delivered over a UNI, then the sum of the CIRs associated with individual services must be less than or equal to the UNI speed.

Yellow frames are out-of-bandwidth profile and will be delivered only if there are adequate bandwidth resources; if, on the other hand, the network is congested, then the frame is discarded. A *red* service frame is also out-of-bandwidth profile and is immediately discarded.

The Color Mode (CM) parameter specifies whether the UNI is operating in a color-aware or color-blind mode. When in a color-aware mode, the color associated with an incoming service frame is employed; in the color-blind mode, the color indication is ignored.

Bandwidth Profile Rate Enforcement The Bandwidth profile is enforced through a two–rate (committed or excess), three-color marker (green, yellow, or red) algorithm, referred to as the *trTCM algorithm;* this algorithm is usually implemented using a token bucket concept and is shown in Figure 2.7.

Two buckets, one referred to as the “committed” or C-bucket and the other referred to as the “excess” or E-bucket, are used. Initially, each of these buckets is full of tokens; the C-bucket has green tokens and the E-bucket has yellow tokens. As service frames enter the Service Provider network UNI, the same number of tokens in the C-bucket are removed (decreased). If, after this, there are green tokens in the C-bucket, then the service frame is CIR-conformant, colored green, and allowed in the network.

If no green tokens remain, however, then the E-bucket is checked to determine if any yellow tokens remain. If there are yellow tokens, then the service frame is EIR-conformant, colored yellow, and allowed in the network. If no yellow tokens are available, then the service frame is colored red and discarded.

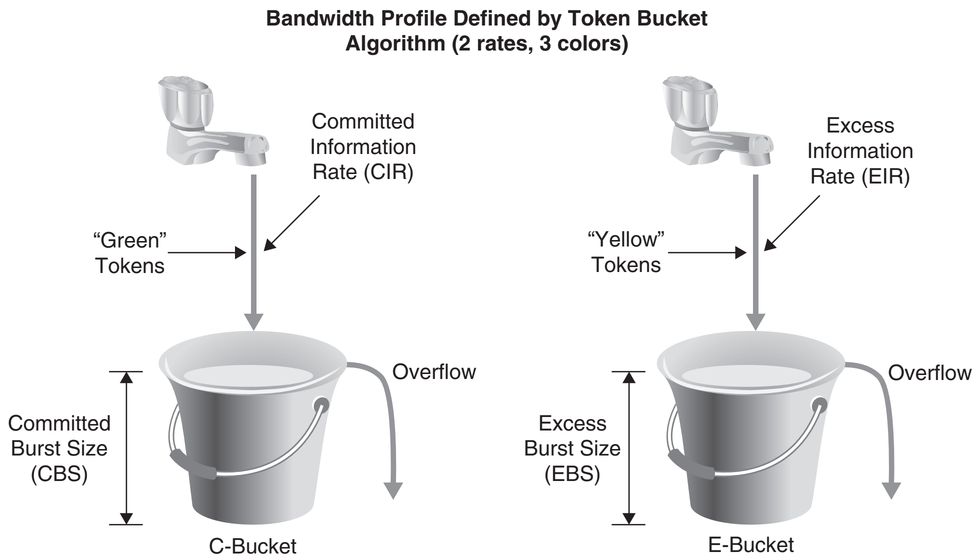

Figure 2.7 Enforcing a predefined bandwidth profile using the token bucket concept (Source: MEF)

The MEF has defined an additional capability whereby unused green tokens from the C-bucket may be added to the E-bucket as yellow tokens when checking EIR-conformance. If this capability is enabled, more yellow service frames are allowed in the Service Provider network.

Performance Parameters The performance parameters affect the service quality experienced by the subscriber and consist of the following.

Availability This is still being formalized by the MEF but essentially attempts to indicate the availability of a service at a predefined performance SLA.

Frame Delay This critical parameter can have an impact on real-time applications such as VoIP and is defined as the maximum delay measured for a percentile of successfully delivered CIR-conformant (green) service frames over a time interval.

The frame delay parameter is used in the CoS service attribute described shortly.

Frame Jitter This service attribute is also known as *delay variation* and is also critical in real-time applications such as VoIP or IP video. Such applications require a low and bounded delay variation to function seamlessly.

Frame Loss Frame loss is defined as the percentage of CIR-conformant (green) fames not delivered between UNIs over a measured interval. At this point, frame loss has been defined for only Point-to-Point EVCs.

NOTE The impact of frame loss depends on specific higher-layer applications. Usually such applications have the ability to recover from frame loss.

Class of Service (CoS) Class of Service (CoS) refers to the performance enforced on a set of similar services. A CoS can be associated with each of the Ethernet services offered but it is usually associated with a group of services. This association becomes especially useful when there are numerous services offered over a resource (e.g., a physical port) that cannot simultaneously support all these services and also meet their respective bandwidth profiles; in such a case, a relative priority between these services becomes necessary. A CoS essentially provides this.

The CoS is also useful because it enables Service Providers to model service demands realistically; customers are increasingly subscribing to services with very different performance demands, for example, Internet access and VoIP require different treatments. With CoS, Service Providers can offer the required level of service and also charge accordingly. It also gives subscribers flexibility.

Each CoS has performance parameters associated with it, and typically the Service Provider will enforce the specified performance. These parameters include bandwidth profile and also jitter, delay, and so on, which will be in the next section.

A CoS is identified using a CoS ID. The various CoS IDs are described in the following sections.

Physical Port Here a single CoS is provided per physical port. All traffic ingressing and egressing the port receives the same CoS. This is a very simple implementation of CoS, but it also affords the least flexibility; if a customer requires multiple CoSs for their traffic (VoIP and Internet access), then two separate ports would be required to enforce the appropriate CoS.

Customer Equipment VLAN (CE-VLAN or 802.1p) This CoS ID refers to the CoS (802.1p) bits in the IEEE 802.1Q tag in a tagged Ethernet service frame. These are usually referred to as the *priority bits.* Using this MEF-defined approach, up to eight classes of service can be provided. A bandwidth profile and performance parameters, which can be enforced by the Service Provider,[16] are associated with each CoS. The user-defined CE-VLAN value(s) may be mapped by a service provider to its own CoS and acted on accordingly.

DiffServ Code Points (DSCP)/IP Type of Service (ToS) The DSCP or IP ToS values in an IP header can be used to determine the CoS. IP ToS provides 8 CoS values, referred to as *IP precedence;* this is similar to the 802.1p bits in the VLAN tag of an Ethernet frame. DSCP, by contrast, specifies 64 different CoS values that correspond to a much more granular performance definition. In addition, DSCP provides a more robust capability that defines the performance over multiple hops in the network (referred to as *per-hop behaviors* or *PHBs*) and attempts to provide a QoS.

Types of Bandwidth Profiles There are three types of bandwidth profiles defined by the MEF; the initial focus has been on the ingress traffic only. Figure 2.8 illustrates the profiles.

- **Ingress bandwidth profile per ingress UNI** This profile provides rate enforcement for all Service Provider frames entering the UNI from subscriber to provider networks. This is useful when only a single service is supported at the UNI, i.e., the UNI is basically considered to be a pipe. The pipe's diameter (bandwidth profile) can be controlled by varying the CIR and EIR parameters. Rate enforcement is non discriminating and some frames may get more bandwidth than others.
- **Ingress bandwidth profile per EVC** This bandwidth profile provides more granular rate enforcement for all service frames entering the UNI that are associated with each EVC. This is useful when multiple services are supported at the UNI; if each EVC is considered to be a pipe inside of a larger UNI pipe, then the bandwidth profile of the EVC—or diameter of the pipe—can be controlled by varying CIR and EIR values.
- **Ingress bandwidth profile per CoS (or CE-VLAN CoS)** This bandwidth profile provides rate enforcement for all service frames belonging to each CoS associated with a particular EVC. The CoS is identified via a CoS identifier determined via the <EVC, CE-VLAN CoS> pair, so that this bandwidth profile applies to frames over a specific EVC with a particular CoS value or even a set of CoS values.

[16] Enforcement depends on whether the Service Provider is set up to handle the same CoS.

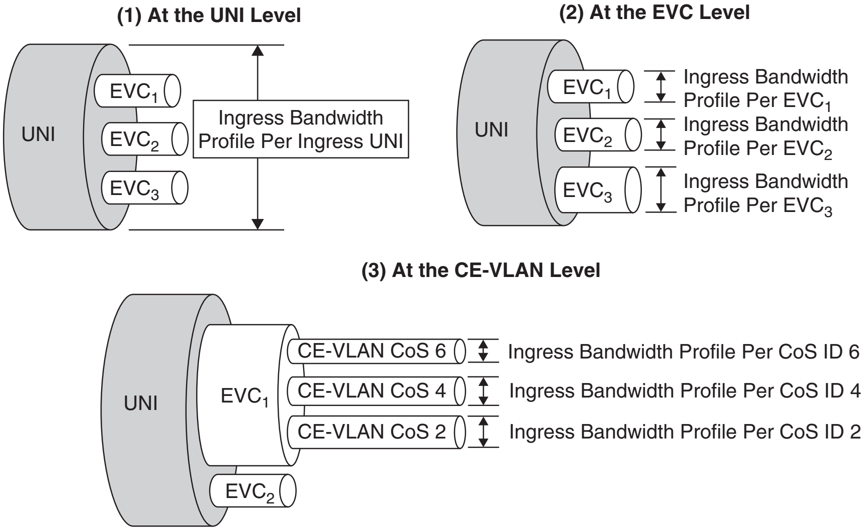

Figure 2.8 MEF-defined bandwidth profiles (Source: MEF)

Service Frame Delivery An EVC allows Ethernet service frames to be exchanged between UNIs that are connected via the same EVC. These may be data frames or control frames. A service provider can indicate what types of frames are supported and those that are not, and also the type of support provided, using four service frame delivery attributes. These are listed next.

Unicast Service Frame Delivery The unicast service frame is defined by the destination MAC address, which may be known (learnt by the network) or unknown. For each UNI pair, this EVC service attribute specifies whether unicast service frames are to be discarded, delivered conditionally (and the specific conditions), or unconditionally.

Multicast Service Frame Delivery In this EVC service attribute, a range of destination MAC addresses are specified, and for each UNI pair, whether multicast service frames are to be discarded, delivered conditionally (and the specific conditions), or unconditionally.

Broadcast Frame Delivery The IEEE 802.3 defines the broadcast address as a destination MAC address of all 1s. For each UNI pair, this EVC service attribute specifies whether broadcast frames are to be discarded, delivered conditionally (and the specific conditions), or unconditionally,.

In general, all Ethernet services support unicast, multicast, and broadcast service frames.

Layer 2 Control Protocol Processing This service attribute can be applied at the UNI or per EVC. There are many Layer 2 control protocols that can be employed (such as IEEE 802.3x MAC control frames, IEEE 802.1x Port Authentication, Spanning Tree

Protocol, Link Aggregation Control Protocol, and so on). The Service Provider can, using this attribute, decide whether to process or discard these protocols at the UNI or pass them to the EVC to discard or tunnel them.

VLAN Tag Support VLAN tag support provides another important set of capabilities that affect service frame delivery and performance. This UNI service attribute allows Ethernet service frames to be 802.1Q tagged or untagged. They can also be used to determine how the frames should then be handled, and if tagged, whether the VLAN ID is used to determine frame delivery.

NOTE UNI pairs for an EVC may support different VLAN tags (one may support it, the other may not; this is useful in service multiplexing described in the next section).

When VLAN tags at the UNI are supported by the Service Provider, then the subscriber needs to knows this and also the action—preserved or discarded or stacked—if any, taken by the Service Provider.

Provider Versus Customer VLAN tag A Service Provider may add an additional VLAN tag to the incoming service frame header to separate from and preserve the customer's VLAN tag, using VLAN stacking (also referred to as *Q-in-Q*). The MEF has defined the term *Customer Edge VLAN ID (CE-VLAN ID)* to represents the customer's VLAN ID; this tag also contains the 802.1p field that the MEF has termed CE-VLAN CoS.

The MEF has defined two service attributes regarding CE-VLAN tag support: CE-VLAN ID preservation and CE-VLAN CoS preservation. The CE-VLAN tag consists of both the CE-VLAN ID and CE-VLAN CoS, so a service may preserve one, both, or neither.

The CE-VLAN ID preservation is an EVC service attribute that defines whether it is preserved across the EVC or not (if not, it is mapped to another value). This is useful for services such as LAN extension.

The CE-VLAN CoS preservation is an EVC service attribute that indicates whether the 802.1p bits are preserved across the EVC or not (if not, it is mapped to another value).

CE-VLAN IDs must be mapped when one UNI of a pair supports tagging whereas the other does not.

Service Multiplexing Service multiplexing provides the ability for a UNI (a physical interface) to support multiple EVCs and precludes the need for a separate physical interface to support each EVC. As illustrated in Figure 2.9, there are multiple EVCs between UNI A and other UNIs in a network (assume that UNI A is at a higher bandwidth physical interface than the other UNIs). By service multiplexing at UNI A, multiple EVCs can be accommodated without needing multiple physical interfaces at UNI A.

Service multiplexing reduces the CAPEX associated with deploying services because it reduces the physical equipment costs. One or fewer physical interfaces are required instead of many; likewise, this reduces the amount of ancillary equipment needed, such as cables. It also reduces the OPEX by enabling quick and remote provisioning of new services.

Bundling The bundling service attribute allows two or more CE VLAN IDs to be mapped to a single EVC at a UNI. These VLANs and the mapping specifics (i.e., which

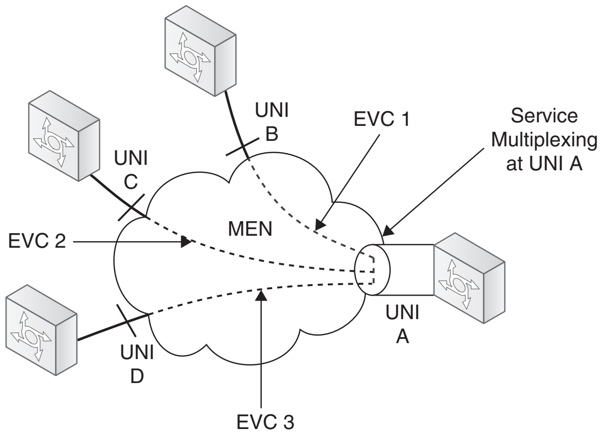

Figure 2.9 Service Multiplexing (Source: MEF)

VLANs map to which EVCs) should be agreed to by the end user and the service provider. A special case of bundling, *all to one bundling,* is enabled when all the VLAN IDs at a UNI are mapped to a single EVC.

Security Filters Security filters enable filtering of undesirable Ethernet frames to maintain security or traffic management. A very plausible case is one wherein a end-user subscriber wants only Ethernet frames originating from specific known sources (identified by the source MAC addresses) to be granted access; any other frames would be considered spurious and dropped, and the user alerted. This is akin to simple Access Control Lists (ACLs) at a UNI.

Ethernet Service Types

The Ethernet service type is essentially a generic Ethernet connectivity construct. The MEF has defined two basic service types:

- Ethernet Line (E-LINE)
- Ethernet LAN (E-LAN)

These two form umbrella categories, and any service can be created using these two categories, modifying only the specific attribute parameters. Therefore, any Ethernet service will be defined as an E-LINE or an E-LAN service, and it will have its own unique UNI and EVC attribute parameters. This should be clearer later in the chapter when we discuss common retail services.

NOTE A third service type called the *Ethernet Tree (E-Tree)* is also being considered by the MEF (possibly in MEF specification 6.1); since it is still in discussion, it is not presented here.

Ethernet Line (E-LINE) Service Any Ethernet service that is based on a Point-to-point Ethernet Virtual Connection (EVC) is designated as an Ethernet Line (E-LINE) service type. The Ethernet Line service is illustrated in Figure 2.10. An E-LINE service type can be used to create a broad range of Point-to-Point Ethernet services between two UNIs.

In its simplest form, an E-LINE service type can provide symmetrical bandwidth for data sent in either direction with no performance assurances, for example, best effort service between two 10 Mbps UNIs. In more sophisticated forms, an E-LINE service type may be between two UNIs at different speeds and may be defined with performance assurances such as CIR with an associated CBS, EIR with an associated EBS, delay, delay variation, and loss.

Service multiplexing may occur at neither, one, or both UNIs in the EVC. For example, more than one Point-to-Point EVC can be offered on the same physical port at one or both of the UNIs. An E-LINE service without any service multiplexing, for example, is very much like the common TDM-based private leased line service (where a UNI physical interface is required for each EVC) except that with an E-LINE service, the range of bandwidth and connectivity options is much greater.

Ethernet LAN (E-LAN) Service Any Ethernet service that is based upon a Multipoint-to-Multipoint Ethernet Virtual Connection (EVC) is designated as an Ethernet LAN (E-LAN) service type. The Ethernet LAN (E-LAN) service type is illustrated in Figure 2.11.

An E-LAN service connects two or more UNIs and service frames sent from one can be received at one or more of the other UNIs. In an E-LAN service, each UNI is connected to a multipoint EVC (even an E-LAN service connected to two UNIs is comprised of a multipoint EVC and hence, not an E-LINE service, which has a Point-to-Point EVC).

An E-LAN can be used to create a broad range of services. In its simplest form, an E-LAN service type can provide a best effort service with no performance assurances between the UNIs. In more sophisticated forms, an E-LAN service type may be defined with performance assurances such as CIR with an associated CBS and EIR with an associated EBS for a given CoS instance. The MEF has not defined service performance (delay, delay variation, and loss) attributes for the E-LAN service type.

For an E-LAN service type, Service multiplexing may occur at neither, one, or more of the UNIs in the EVC. For example, an E-LAN service type (Multipoint-to-Multipoint

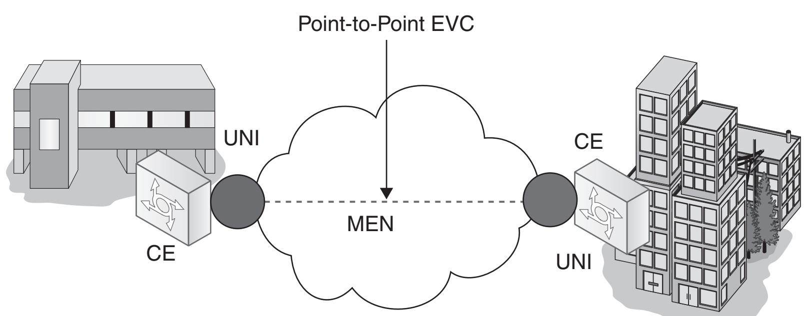

Figure 2.10 Ethernet Line (E-LINE) service type (Source: MEF)

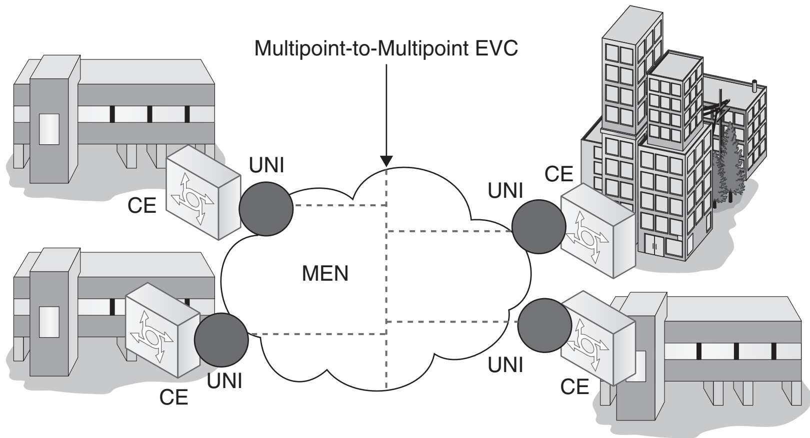

Figure 2.11 E-LAN service type using Multipoint-to-Multipoint EVC (Source: MEF)

EVC) and an E-LINE service type (Point-to-Point EVC) may be service multiplexed at the same UNI. In this example, the E-LAN service type may be used to interconnect other subscriber sites while the E-LINE service type is used to connect to the Internet with both services offered via EVC service multiplexing at the same UNI.

An E-LAN service may include a different bandwidth profile <CIR, CBS, EIR, EBS, CM> configured at each of the UNIs. An E-LAN service can also interconnect a large number of sites with much less complexity than legacy technologies such as Frame Relay and ATM. Furthermore, it can be used to create a broad range of services such as Private LAN and Virtual Private LAN service.

Ethernet Private and Virtual Connectivity Services Using the E-LINE and E-LAN service types, the MEF has also defined simple connectivity services based on whether they are port-based or VLAN-based. The port-based service, where all-to-one bundling is employed, is essentially providing a private service with dedicated bandwidth, while the VLAN-based service allows service multiplexing at a UNI to enable a virtual service, in which bandwidth is shared among multiple EVCs. This is detailed in Figure 2.12.

Connectivity Services

Ethernet Service Type	Dedicated (All to One Bundling)	Shared (Service Multiplexed)
E-LINE	Ethernet Private Line (EPL)	Ethernet Virtual Private Line (EVPL)
E-LAN	Ethernet Private Line (EPLAN)	Ethernet Virtual Private LAN (EVPLAN)

Figure 2.12 E-LINE and E-LAN Connectivity Services

Ethernet Private Line (EPL) Using E-LINE Service Type An Ethernet Private Line (EPL) service is specified using an E-LINE service type. EPL uses a Point-to-Point EVC between two UNIs and provides a high degree of transparency for service frames between the UNIs it interconnects, such that the service frame's header and payload are identical at both the source and destination UNI. The service also has an expectation of low frame delay, frame delay variation, and frame loss ratio. It does not allow for service multiplexing because a dedicated UNI (physical interface) is used for the service. Due to the amount of transparency in this service, there is no need for coordination between the subscriber and Service Provider on a detailed CE-VLAN ID/EVC map for each UNI because all service frames are mapped to a single EVC at the UNI. An EPL is depicted in Figure 2.13.

NOTE MEF 6.1 might incorporate a further distinction in the Ethernet Private Line; specifically, there is consensus to define two EPL service variants: EPL-T (EPL-Transport) and EPL-P (EPL-Packet). EPL-T would essentially be the EPL defined here. EPL-T would be enhanced, adding features such as multiple CoS parameters as well as bandwidth profile parameters.

Ethernet Virtual Private Line (EVPL) Using E-LINE Service Type An Ethernet Virtual Private Line (EVPL) is created using an E-LINE service type. An EVPL can be used to create services similar to the Ethernet Private Line (EPL) with some notable exceptions. First, an EVPL allows for service multiplexing at the UNI. This capability allows more than one EVC to be supported at the UNI whereas the EPL does not allow this. Second, an EVPL need not provide full transparency of service frames as with an EPL. Because service multiplexing is permitted, some service frames may be sent to one EVC while other service frames may be sent to other EVCs. An EVPL is also shown in Figure 2.13.

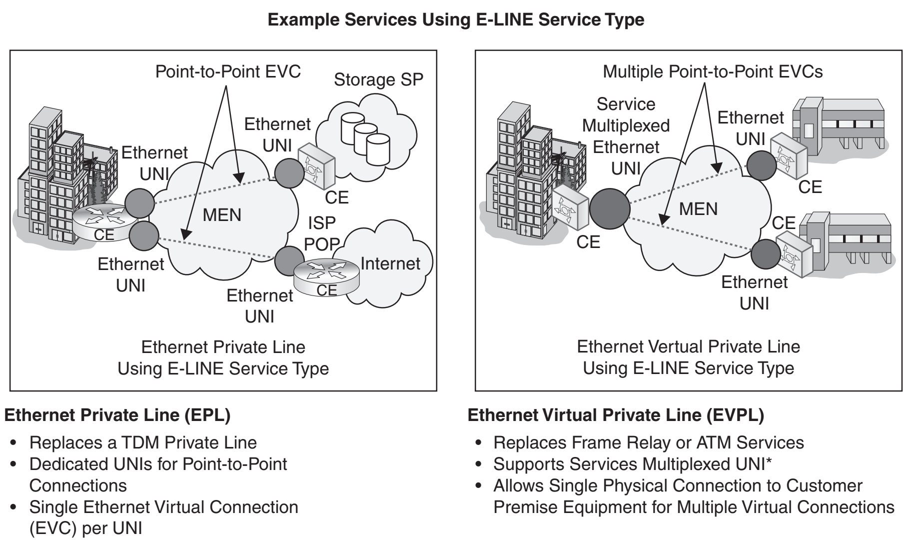

Figure 2.13 Examples of services using the E-LINE service type (Source: MEF)

Ethernet Private LAN (E-PLAN) Using E-LAN Service Type E-PLAN enables a wide area LAN Ethernet in which service multiplexing is allowed at the UNI. There is full transparency of service frames within an E-PLAN. This essentially creates a transparent LAN service that makes the Service Provider network one large Ethernet, as shown in Figure 2.14.

Ethernet Virtual Private LAN (EVPLAN) Using E-LAN Service Type or Layer 2 VPN Using an E-LAN service over a shared infrastructure, a transparent LAN service is created. This essentially makes the Service Provider network one large Ethernet and is typically used for applications such as intra-company connectivity services.

Sample Commercial Offerings Using Carrier Ethernet Services

Some common connectivity services delivered using Carrier Ethernet are briefly outlined next; higher level applications are increasingly employing Carrier Ethernet due to the many benefits that it affords.

LAN Extension Subscribers with multiple sites in a metro area often want to interconnect them at high speeds so all sites appear to be on the same LAN and have equivalent performance and access to resources such as servers and storage. This is referred to as *LAN extension.*

In essence, the LANs at each site are connected; this is simpler and cheaper than routing, although it may not scale as well in large networks. To connect only two sites, a Point-to-Point E-LINE service can be used. To connect three or more sites, the subscriber has the choice of using either multiple E-LINE services or an E-LAN service.

Figure 2.15 shows a four-site LAN extension created using an E-LAN service. Each of the sites/UNIs support CE-VLAN ID and CE-VLAN preservation so the subscriber's

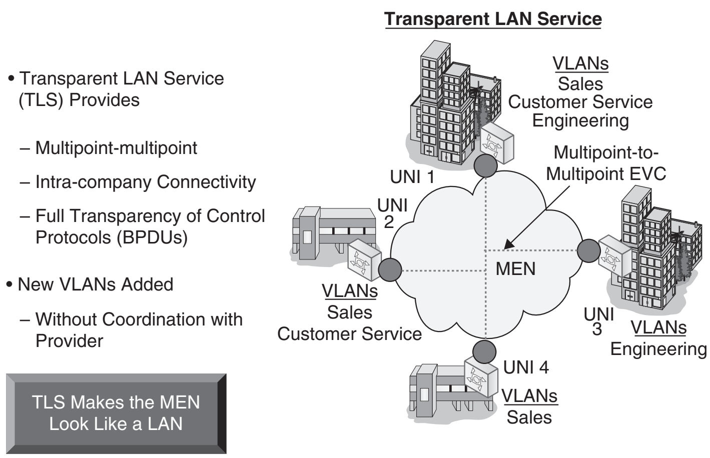

Figure 2.14 E-PLAN using E-LAN service type (Source: MEF)

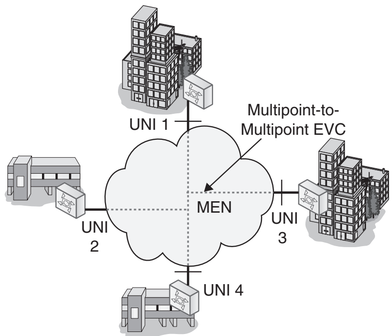

UNI Service Attribute	Service Attribute Values and Parameters
Physical Medium	IEEE 802.3–2002 Physical Interface
Speed	10 Mbps (all UNIs)
Mode	FDX Fixed Speed (all UNIs)
MAC Layer	IEEE 802.3–2002
Service Multiplexing	No
CE–VLAN ID/EVC Map	All CE-VLAN IDs Map to the Single EVC
Bundling	No
All to One Bundling	Yes
Ingress Bandwidth Profile Per Ingress UNI	All UNIs: CIR = 5 Mbps, CBS = 256 KB, EIR = 10 Mbps, EBS = 512 KB
Layer 2 Control Protocol Processing	Process IEEE 802.3 x MAC control frames
	Process Link Aggregation Control Protocol (LACP)
	Process IEEE 802.1 x Port Authentication
	Pass to EVC Generic Attribute Registration Protocol (GARP)
	Pass to EVC Spanning Tree Protocol
	Pass to EVC a Protocol Multicasted to all Bridges in a Bridged LAN

EVC Service Attribute	Service Attribute Values and Parameters
EVC Type	Multipoint-to-Multipoint
UNI List	UNI 1, UNI 2, UNI 3, UNI 4
CE-VLAN ID Preservation	Yes
CE-VLAN CoS Preservation	Yes
Unicast Frame Delivery	Deliver Unconditionally for each UNI Pair
Multicast Frame Delivery	Deliver Unconditionally for each UNI Pair
Broadcast Frame Delivery	Deliver Unconditionally for each UNI Pair
Layer 2 Control Protocol Processing	N/A–IEEE 802.3 x MAC Control Frames
	N/A–Link Aggregation Control Protocol
	N/A–IEEE 802.1 x Port Authentication
	Tunnel Generic Attribute Registration Protocol (GARP)
	Tunnel Spanning Tree Protocol (STP)
	Tunnel a Protocol Multicasted to all Bridges in a Bridged LAN
Service Performance	One CoS for all UNIs; Frame Delay < 30 ms, Frame Jitter: N/S, Frame Loss < 0.1%

Figure 2.15 LAN extension service using E-LINE service type (Source: MEF)

VLAN tag is not modified. In this case, the MEN appears as a single Ethernet segment in which any site can be a member of any VLAN. The advantage here is the subscriber can configure new CE-VLANs across these sites without involving the Service Provider. The service attributes are also shown in the figure.

Dedicated Internet Access (DIA) Dedicated Internet access enables subscribers to have a high-speed connection to the Internet to support their business objectives. An EVC can connect the subscriber's site to the local point-of presence (POP) of the Internet service provider (ISP) using a Point-to-Point E-LINE service.

If a customer is homed to multiple (say two) ISPs, as shown in the Figure 2.16, then a separate E-LINE would be used to connect each ISP. If the same UNI is expected to provide Internet access and other services, then a separate EVC would be used for each of the services.

At the ISP, service multiplexing is typically employed over a high-speed UNI to support multiple subscribers, so in effect, each subscriber appears to have a dedicated connection

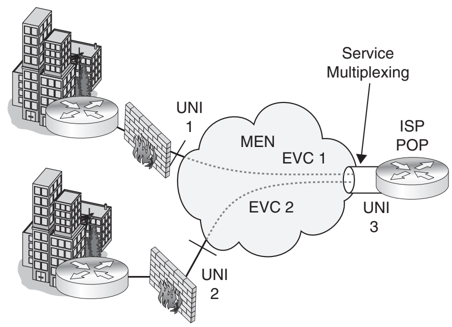

UNI Service Attribute	Service Attribute Values and Parameters
Physical Medium	IEEE 802.3–2002 Physical Interface
Speed	UNIs 1 and 2: 100 Mbps UNI 3: 1 Gbps
Mode	UNIs 1 and 2: 100 Mbps FDX Fixed UNI 3: 1 Gbps FDX
MAC Layer	IEEE 802.3–2002
Service Multiplexing	No at UNIs 1 and 2 Yes at UNI 3
CE-VLAN ID/ EVC Map	N/A Since Only Untagged Frames Used Over the EVC
Bundling	No
All to One Bundling	No
Ingress Bandwidth Profile Per EVC	UNIs 1 and 2: CIR = 50 Mbps, CBS = 2 MB, EIR = 100 Mbps, EBS = 4 MB UNI 3: CIR = 500 Mbps, CBS = 20 MB, EIR = 1 Gbps, EBS = 40 MB
Layer 2 Control Protocol Processing	Discard 802.3 x MAC Control Frames
	Discard Link Aggregation Control Protocol (LACP)
	Discard 802.1 x port Authentication
	Discard Generic Attribute Registration Protocol (GARP)
	Discard Spanning Tree Protocol
	Discard a Protocol Multicasted to all Bridges in a Bridged LAN

EVC Service Attribute	Service Attribute Values and Parameters
EVC Type	Point-to-Point
UNI List	EVC 1: UNI 1, UNI 3 EVC 2: UNI 2, UNI 3
CE-VLAN ID Preservation	No. Mapped VLAN ID for Use with Multi-homed ISPs (if required)
CE-VLAN CoS Preservation	No
Unicast Frame Delivery	Deliver Unconditionally for each UNI Pair
Multicast Frame Delivery	Deliver Unconditionally for each UNI Pair
Broadcast Frame Delivery	Deliver Unconditionally for each UNI Pair
Layer 2 Control Protocol Processing	N/A[3] –IEEE 802.3 x MAC Control Frames
	N/A–Link Aggregation Control Protocol (LACP)
	N/A–IEEE 802.1 x Port Authentication
	N/A–Generic Attribute Registration Protocol (GARP)
	N/A–Spanning Tree Protocol (STP)
Service Performance	Only 1 CoS Supported. Frame Delay < 30 ms (95th percentile), Frame Jitter: N/S[10], Frame Loss < 0.1%

Figure 2.16 Dedicated Internet access using E-LAN service type (Source: MEF)

to the ISP. In Figure 2.16, the ISP may have a 1GbE UNI, while the subscribers' UNIs 1 and 2 may be 100 Mbps. There is no service multiplexing at the subscriber UNI. The service attributes are also shown in the figure.

Other Commercial Applications of Carrier Ethernet Services Carrier Ethernet is increasingly being employed for several traditional and emerging applications that require carrier-class performance while minimizing the cost of delivery. A sample of some of the popular revenue-generating and value-added applications being enabled by Carrier Ethernet services includes packet video, VoIP and VoIP peering, Layer 2 VPNs, content peering, extranet connectivity, business continuity and disaster recovery, IP backbone expansion, and wireless backhaul.

Most of these are implemented over straightforward E-LINE services and, in some cases, over an E-LAN service. The simple but fairly encompassing nature of basic Ethernet services has enabled Service Providers to tailor a wide range of customized applications and generate new value-added and higher premium offerings literally every day. Currently, it is estimated that there are well over 500 different Ethernet services being offered by over 200 Carriers in the U.S. alone.

Carrier Ethernet: The Enablers

In this section, the developments of each of the attributes are presented. It is impractical to discuss these in any detail here, as several standards are involved, but they are highlighted with the progress to date. Some selective developments are, however, elaborated in reasonable depth given their importance in the enabling of Carrier Ethernet.

The MEF's prominent role in enabling Carrier Ethernet is shown in Figure 2.17, which highlights the specifications and specific attributes being addressed in the areas of Architecture, Management, Services, and Testing Measurement.

Standardized Services

The standardized services attribute requires support for Ethernet services and also for other prevailing services, notably TDM-based services over a Carrier Ethernet infrastructure (i.e., over an E-LINE/E-LAN service). There has been considerable effort spent on standardizing Ethernet services (as detailed in the previous section) in MEF 6 that encompassed setting up bandwidth profiles and traffic management. Major developments are outlined in Table 2.2.

Supporting other services, especially the TDM-based services, has also been addressed quite significantly by the MEF (and other bodies) and a reasonably detailed overview is provided next.

	Carrier Ethernet Attributes				
MEF Specs	Standardized Services	Service Management	Reliability	Quality of Service	Scalability
MEF 2			**Architecture Area**		
MEF 3	Service Area			Service Area	
MEF 4	**Architecture Area**				
MEF 6	Service Area			Service Area	Service Area
MEF 7		Management Area			
MEF 8	Service Area				
MEF 9	**Test & Measurement Area**		**Test & Measurement Area**		
MEF 10.1	Service Area			Service Area	Service Area
MEF 11	**Architecture Area**				
MEF 12	**Architecture Area**				**Architecture Area**
MEF 13	**Architecture Area**				
MEF 14	**Test & Measurement Area**		**Test & Measurement Area**	**Test & Measurement Area**	
MEF 15		Management Area			
MEF 16		Management Area			
MEF 17		Management Area			
MEF 18	**Test & Measurement Area**		**Test & Measurement Area**		
MEF 19	**Test & Measurement Area**		**Test & Measurement Area**		

Figure 2.17 The MEF specifications enabling Carrier Ethernet

TABLE 2.2 Standards Efforts Enabling Standardized Services

Key Components	Major Developments	Reference
Ubiquity	Standardization of UNI and traffic management for consistent delivery across different infrastructures	MEF 6, MEF 10.1, MEF 11
Ethernet services	Generic architecture and terminology developed; standardized Ethernet services defined in this context that form the basis for all Ethernet services	MEF 4, MEF 6, MEF 10.1
Circuit Emulation Services	Standardized circuit emulation services over Ethernet, along with performance requirements as well as practical implementation requirements	MEF 3, MEF 8, ITU-T, IETF PW3E
Granularity of bandwidth and QoS	A standard bandwidth profile for Ethernet services developed	MEF 6, MEF 10, MEF 11
Converged transport	A rich set of capabilities required for sophisticated implementation of converged enterprise and residential networks defined	MEF 12, ITU-8010

Circuit Emulation Services over Ethernet (CESoE) As noted previously, non packet services such as PDH and SONET/SDH account for a significant amount of customer demand in the market today, and Service Providers expect to leverage this opportunity. As Service Providers move to a Carrier Ethernet–based packet-optimized network infrastructure, they should still be able to provide these services. This requirement translates into being able to transport these synchronous Time Division Multiplexing (TDM) digital signals over an asynchronous Ethernet infrastructure. Or put another way, a TDM circuit-switched network should be emulated over this packet infrastructure and provide what is referred to as circuit emulation services (CES). In effect, these services tunnel customers' TDM traffic over the Ethernet network, as shown in Figure 2.18. The customers' source and destination TDM equipment on either end is unaware of this circuit emulation. Such CES typically run over standard E-LINE service.

With CES over Ethernet (CESoE), service providers can leverage the inherent advantages of Carrier Ethernet—flexibility, simplicity, and lowered OPEX,[17] while delivering legacy applications such as TDM voice and private lines (which still account for a very large proportion of revenues for most Service Providers). Thus, with CESoE, Service Providers can cost effectively offer a complete portfolio of emerging Ethernet services along with the legacy services, obtaining an approximately 30 percent savings in infrastructure costs, and OPEX can be realized by migrating to an unified Ethernet infrastructure.

[17] Ironically, this also extends the longevity of these legacy applications.

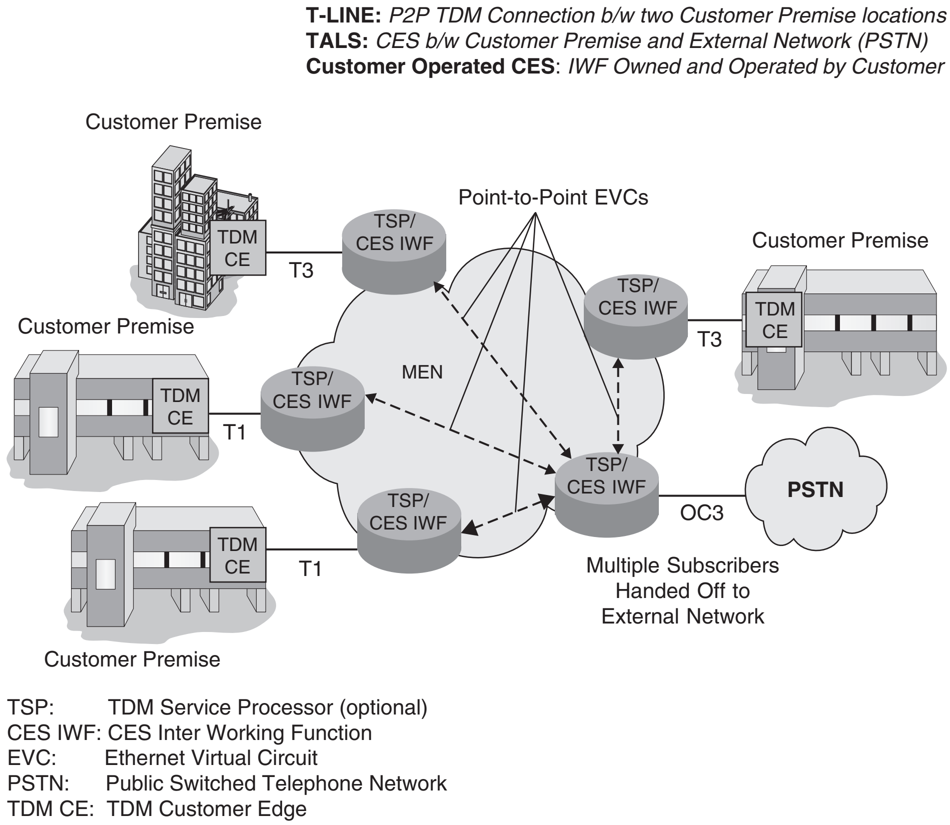

Figure 2.18 Circuit emulation over Ethernet (Source: MEF)

The MEF has provided the industry's first formal definition of CESoE that covers the ability to deliver both PDH services (e.g., N × 64 kbit/s, T1, E1, T3, and E3) and SONET/SDH services (STS-1, STS-3, STS-3c, STS-12, STS-12c, and European equivalents). The MEF 3 specifications address the types of CES that can be offered over a Service Provider–enabled Carrier Ethernet network (using EVCs) and also the requirements of these services. The specifications basically enable the support of traditional TDM handoffs to customer's voice equipment.

NOTE Voice is by far the dominating application requiring underlying TDM circuit-switched services.

The MEF 8 addresses the practical aspects of CES and provides precise instructions for implementing interoperable CES equipment that will conform to the performance requirements outlined for CES in ITU-T and the ANSI TDM standards. The ITU-T recommendation Y.1413 is very similar to the MEF 8, except for MPLS networks, and employs identical frame formats for payload and encapsulation so that the equipment supporting Y.1413 should also be capable of supporting MEF 8. The IETF has several drafts, including PsuedoWire Edge to Edge Emulation (PWE3) and CES over Packet Switch Network, that are similar to the MEF 8 but focus on IP/MPLS networks. Because the payload and encapsulation formats are identical, any equipment supporting these drafts will also support MEF 8.

The MEF has essentially defined three types of CESoE that are generically portrayed in Figure 2.19. In each case, the CES is based on a Point-to-Point connection between two inter-working functions labelled CES IWF; this CES IWF essentially provides a translation function with a TDM application interface on one side (customer equipment facing) and an Ethernet interface (Service Provider network facing).

There is also an optional TDM service processor (TSP) that consists of any TDM grooming function that may be required to convert the TDM service offered to the customer into a form that the CES IWF can accept. For example, the TSP may be a framer device, converting a fractional DS1 service offered to the customer into a N×64 kbit/s service for transport over the MEN.

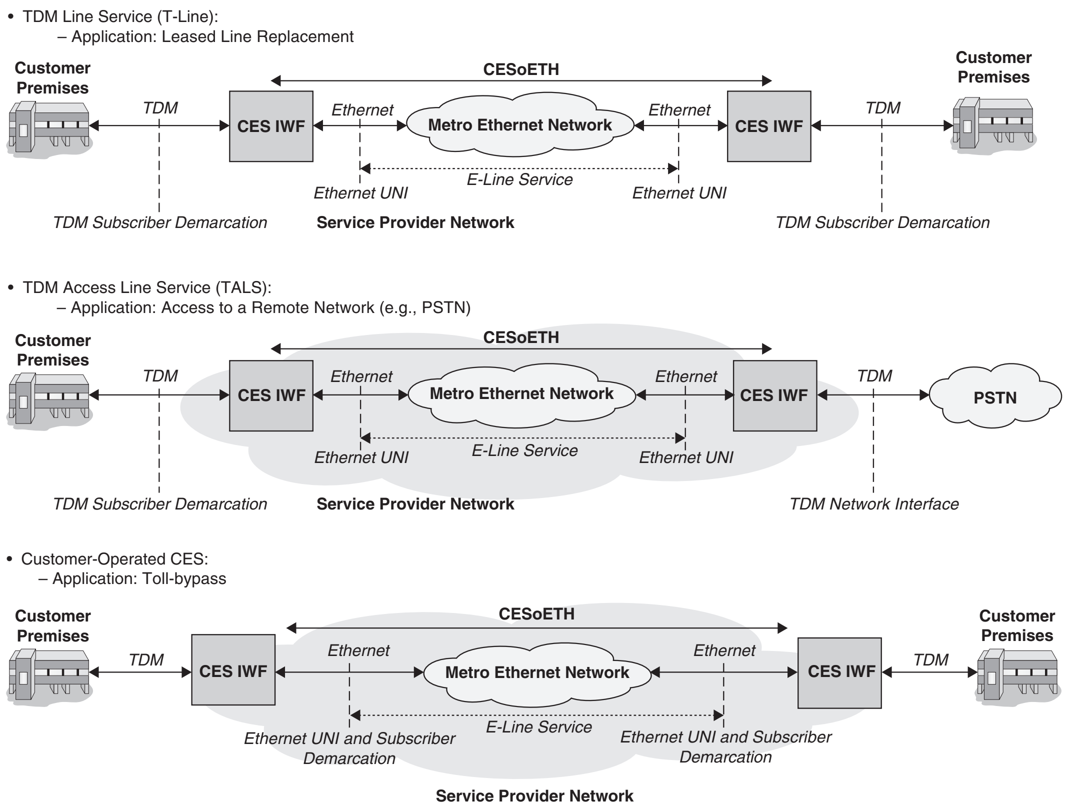

Figure 2.19 Types of CESoE defined by the MEF (Source: MEF)

The TSP and the CES IWF may physically reside in the Provider Edge (PE)[18] unit at the provider's nearest Point-Of-Presence or in a Service Provider–owned box in a customer location (e.g., a multi-tenant unit). From the architectural perspective, there is no difference between these alternatives.

TDM Line Service (T-Line) The basic TDM Line (T-Line) service is a Point-to-Point, constant bit rate service, similar to the traditional leased-line type of TDM service. However, service multiplexing may occur ahead of the CES inter-working functions, (e.g., aggregation of multiple emulated T1 lines into a single T3 or OC-3 link), creating a Multipoint-to-Point or even a Multipoint-to-Multipoint configuration, as shown in Figure 2.19.

The service multiplexing is carried out using standard TDM multiplexing techniques and is considered part of the TSP block, rather than the CES inter-working function. The TDM interface at the input of the CES inter-working function is the same as the output from the CES IWF at the opposite end of the emulated link. It is the TSP that may be used to multiplex (or de-multiplex) that TDM service into the actual TDM service provided to the customer. This allows a TDM service to a customer to be provided as a collection of emulated services at lower rates.

There are, therefore, three modes of operation: unstructured emulation, structured emulation, and multiplexing mode; in all three modes, the delivery of the TDM service employs Ethernet Virtual Connections (EVCs) as shown in the Figure 2.19.

In unstructured emulation, the service is Point-to-Point and will have the identical TDM handoff on either end. The CES capability should maintain integrity across the network. In a structured mode, the service is also Point-to-Point and will have identical TDM service handoffs on either end, except that the overhead and payload entering on one end will terminate the overhead at the near endpoint and transport the payload transparently to the other end where it is mapped to the same type of overhead and terminated. Examples of this are typical OC-3 services where the SONET overhead (SOH) is terminated locally and the payload transported and the overhead then added before terminating at the other end. The CES will maintain the integrity of the transport.

Finally, in the multiplexed mode, multiple lower-rate transparent services are multiplexed at a specific service endpoint on the network into a higher digital hierarchy. Similarly, a higher rate service may be decomposed into several lower rate services. For example, a customer may have several sites—a head office with a full DS1 connection and several satellites with fractional DS1 connections, as shown previously in Figure 2.16. The same architecture can be used for multiplexing of other rate services, for example, several full DS1 services onto a single DS3 or multiplexing of VT-1.5s into an STS-1.

In order to attain some efficiency between mapping the TDM hierarchy signals into an Ethernet frame, the recommended bandwidth granularity is 100 kbits/s. TDM multiplexing of signals is also possible and a higher aggregated signal handed to the IWF for transport; at the other end, an identical de-multiplexing will occur.

[18] Provider Edge (PE) denotes the Carrier POP or CO.

TDM Access Line Service (TALS) TALS is almost identical to the T-Line service except that it is a Multipoint-to-Point service, and one or more ends of the TDM service handoff is to an external network (such as a PSTN, as shown in Figure 2.19). A common example of such a service is when the Service Provider Ethernet network is an access to an external network. As with T-Line service, it can be operated in similar modes and should ensure that it maintains integrity of the signal on an end-to-end basis.

Customer-Operated CES In this type of CESoE, the IWF is actually owned and/or provided by the customer and the customer only subscribes to a typical E-LINE service from the Service Provider. Usually in such a scenario, the Service Provider is expected to provide a stringent SLA with tighter definitions of parameters such as packet delay, variation in packet delay, and packet loss, to accommodate the TDM service.

NOTE From a Service Provider's standpoint, the CESoE is actually just an Ethernet service.

The MEF 3 has also defined performance expectations for the Service Provider network delivering the CESoE to ensure that toll-grade voice quality is maintained. Specifically, it identifies the following four Class of Service (CoS) characteristics: Ethernet frame delay, Ethernet frame delay variation (jitter), Ethernet frame loss, and network availability. These parameters should conform to values consistent with those in a typical TDM environment (i.e., five 9s network availability, less than 10 ms jitter, and so on).

Implementation Support MEF 8 provides further detail on implementing the requirements specified in MEF 3 when supporting PDH services over a MEN/Service Provider Ethernet Network (SEN). In so doing, the specification is attempting to address the inherent challenges of transporting TDM signals. The technical challenges faced by CESoE primarily stem from replicating constant bit rate TDM services over a variable bit rate Ethernet infrastructure. These challenges include packetization, frame delay variation, clock recovery, and synchronization and TDM performance monitoring.

In particular, five functions are specified to ensure interoperability: connectivity, timing, signaling, MEN performance criteria, and MEN services OAM. MEF 8 focuses especially on timing and signaling issues. The specification also augments the performance characteristics defined in MEF 3.

Timing/Synchronization Synchronization is an important consideration in any circuit emulation scheme and the clock of the incoming signal (into the IWF) and outgoing signal (to the IWF) should be synchronized (i.e., the frequency should be the same). There are four options for this clock:

- **TDM line timing** Use the clock from the incoming TDM line.
- **External timing** Use an external reference clock source.
- **Free run timing** Use a free-running oscillator.
- **Ethernet line timing** Recover the clock from the Ethernet interface.

The last option, Ethernet line timing, covers all methods where information is extracted from the Ethernet, including *adaptive timing,* where the clock is recovered from data in the CESoE frames and the arrival time of the frames, and *differential timing,* where the clock is recovered from a combination of data contained in the CESoE frames and knowledge of a reference clock common to both the SEN-bound and TDM-bound IWFs.

For maximum applicability, it is recommended that CESoE implementations should support at least TDM line external and adaptive timing to enable the implementation to be used in the majority of timing scenarios. Synchronization (and jitter and wander) requirements are placed on a CESoE implementation by the MEF 8 and should conform to the ITU-T recommendations G.823 and G.824 for E1/E3 and DS-1/DS-3, respectively.

Signaling CE applications interconnected over a CESoE service may exchange signaling in addition to TDM data. The typical example is telephony applications that exchange their state (e.g., off-hook/on-hook) in addition to TDM data carrying PCM-encoded voice.

With structure-agnostic emulation, signaling is not required to intercept or process CE signaling. Signaling is embedded in the TDM data stream, and hence it is carried end-to-end across the emulated circuit.

With structure-aware emulation, transport of Common Channel Signaling (CCS) may be achieved by carrying the signaling channel with the emulated service (e.g., channel 23 for DS1 or channel 16 for E1). However, Channel Associated Signaling (CAS), such as DS1 Robbed Bit Signaling or E1 CAS, requires knowing the relationship of the timeslot to the trunk multiframe structure. This is indicated by the framing bits, which may not be preserved by N×64 kbit/s basic service.

MEF 8 describes a generic method for extending the N×64 kbit/s basic service by carrying CE signaling (CAS or CCS) in separate signaling packets that are independent of the TDM circuit type. This method may be used in situations where the individual 64 kbit/s channels are selected from multiple TDM circuits or picked off a TDM bus rather than from a specific TDM circuit; it also saves SEN bandwidth.

Scalability

One of the major requirements of Carrier Ethernet is to scale to meet the needs of Service Provider offerings. The limitations imposed by the QinQ (IEEE 802.1ad, stacked VLANs) allow for only 4094 VLANs/service instances in a service area (based on the 12 bits used in the VLAN ID field for this purpose). However, this is inadequate to support the kind of scale required by the MEF. Key standards developments are noted in Table 2.3.

Provider Backbone Bridging Provider Backbone Bridging (PBB) or IEEE 802.1ah addresses the service scaling limitations in native Ethernet networks by enabling millions of service instances in a serving area through the creative use of the MAC address.

TABLE 2.3 Standards Efforts Enabling Scalability

Key Components	Major Developments	Reference
Millions of users/endpoints	Extended the addressable space for users and architecture and framework for scaling services defined	IEEE 802.1ah,IEEE 802.1d, MEF 6, MEF 10.1
Geographic reach/applications	Provided for MAC encapsulation (MAC-in-MAC) to enable substantial Layer 2 scalability	IEEE 802.1ah, IEEE 802.1QAy
Bandwidth granularity	Defined how the bandwidth profile parameters can be set from 1M to 10G in granular increments	MEF 11

Essentially, PBB employs an additional Service Provider 16 bit MAC address[19] that corresponds to the ingress Ethernet ports of the Service Provider edge device and basically encapsulates the end user's MAC (this is also referred to as *MAC-in-MAC*). The outer MAC address is used to forward the Ethernet frames across the Service Provider network, and this much larger physical address space (approx 2^{16}) allows for a more scalable network than the traditional one with VLAN IDs—where even with the QinQ scheme, stacking a Service Provider VLAN tag over the customer VLAN tag, only 4094 service instances are supported.

The MAC-in-MAC significantly improves scalability and also provides some security by separating the customer and Service Provider address space. It also precludes a MAC address explosion and the need for learning substantially more end-user MAC addresses in the Service Provider's core infrastructure (switches and so on). Minimizing the number of MAC addresses that need to be learned also reduces the aging out and relearning of MAC addresses, enhancing end-to-end performance, and in general, making the network more stable as far as forwarding Ethernet frames is concerned.

The IEEE 802.1ah efforts to standardize PBB should be consulted for more updated information on PBB.

Provider Bridge Transport (PBT)/PBB with Traffic Engineering (PBB-TE) While Carrier Ethernet can be delivered over numerous transport technologies, such as SONET and MPLS (see Part II for a compendium), one option is to deliver it over native Ethernet (see Chapter 13 on bridging and switching). However, native Ethernet itself has been limited as a plausible transport technology especially as Carrier Ethernet services were enabled on a wide area basis (beyond the access networks and stretching well into the core and beyond). One key hurdle is the inherent best-effort approach of LAN Ethernet, which is ill-suited in a service that supports time-sensitive applications.

With the emergence of a new standardization effort, namely the Provider Bridge Transport (PBT), a more deterministic Ethernet is being attempted. In fact, PBT

[19] The Ethernet frame size is now correspondingly augmented; the devices in the network should be able to support this.

aims to provide the connection-oriented features of TDM to the hitherto connectionless Ethernet. The IEEE has undertaken this effort—also referred to as the *Provider Backbone Bridge with Traffic Engineering (PBB-TE)*—since it is essentially a variation of the IEEE 802.1ah PBB standard. In fact, PBT also employs a MAC-in-MAC forwarding scheme from PBB and also distributes the bridging tables using the control plane. PBT, however, does not use some of the features defined in PBB such as broadcasting and MAC learning and does not support the Spanning Tree Protocol.

PBT basically provisions Point-to-Point Ethernet paths that are engineered across Service Provider Ethernet networks. These paths provide traffic engineering (and are referred to as *PBB-TE*) and allow for setting up QoS to meet predefined SLAs across the service provider WAN. PBT operates by adding configured routes to the standard PBB network.

In addition, 50 ms recovery can also be provided to meet the industry expectation of service provider networks. In conjunction with Ethernet OAM standards (discussed later in this chapter), proactive fault management can also be incorporated for these Ethernet paths. Because PBT transport can be independent of the service carried over this transport, it can be used to Carrier non-Ethernet services as well.

Given that existing technologies such as MPLS are more established (especially in the core of Service Provider networks), the need for PBT is being questioned in some quarters; while proponents claim compelling CAPEX and OPEX savings vis-à-vis MPLS, the incumbency of MPLS (i.e., already deployed and depreciating) may make it harder to displace, especially in existing networks. In green field networks, however, there may be a better opportunity for PBT.

More information on PBT can be obtained from the appropriately noted references.

Reliability

While MEF has defined service-level reliability and its components' service resiliency, protection, and less than 50 ms restoration, several of the underlying transport solutions employed to deliver Carrier Ethernet, particularly SONET and RPR, have established a high level of reliability in Service Provider networks. MEF 2 allows the MEN to leverage any underlying transport layer protection type if it can enable end-to-end service protection.

Table 2.4 identifies key standards based developments that are incorporating Reliability.

Two protection types, 1+1 and *M*:*N*, have been defined. In the 1+1 approach, duplicate traffic is the norm, and in the case of a failure/protection event, one stream of traffic is still available (unless the failure is catastrophic). In the case of *M*:*N*, *N* working resources are provided protection using *M* protection sources.

Four different protection mechanisms have also been defined: Aggregate Link and Node Protection (ALNP) to protect against local link/node failure; End-to-End Path Protection (EEPP), where redundancy is provided for the primary path on an end-to-end basis; MP2MP protection of E-LAN services including Rapid Spanning Tree Protocol (RSTP) and link redundancy; and finally, link protection based on link aggregation, where one or more Ethernet links connected between the same nodes can be aggregated.

TABLE 2.4 Standards Efforts Enabling Reliability

Key Components	Major Developments	Reference
Service resiliency	Less than 50 ms resiliency has been defined as a critical requirement	MEF 2, IEEE 802.1ag
Protection	Defined broad framework for hop by hop and end-to-end service-level protection Defined four protection mechanisms and also allowed leveraged end-to-end service protection available at the transport layer	MEF 2
Restoration	Different levels of restoration have been defined to afford a wide variety of application requirements	MEF 2, IEEE 802.1QAy (PBT/PBB-TE)

MEF 2 has provided for supporting a wide variety of restoration times, from less than 5 seconds to the less than 50 ms range, in order to support the wide variety of applications and their corresponding requirements. The MEF 2 also allows end users to choose a variety of protection parameters for a Carrier Ethernet service. These protection parameters must be applicable on a per service or a group level. Any of the ETH layer protection mechanisms in MEF 2 should be able to work in conjunction with the lower layer (transport) protection mechanisms.

Quality of Service

As mentioned earlier, Provide Bridge Transport (PBT) can provide deterministic transport of Ethernet services, and hence QoS much like other underlying transport used to deliver Carrier Ethernet. This is, in fact, a critical requirement of Carrier Ethernet, one that needs to be addressed well before the market begins to embrace it more wholeheartedly. At the ETH layer, MEF 10 has undertaken a significant amount of effort toward defining and implementing QoS to ensure rigorous SLAs.

Table 2.5 notes some of the key developments in the standards bodies' with regard to Carrier Ethernet QoS.

TABLE 2.5 Standards Efforts Enabling Quality of Service

Key Components	Major Developments	Reference
Wide choice of granularity and QoS options	Different levels of granular bandwidth defined; also bandwidth profile defined for providing different class of services	MEF 6, MEF 10.1
End-to-end performance SLAs	Defined how some traffic is delivered with strict SLAs while other traffic is delivered with best effort; traffic management algorithm to ensure SLA	MEF 3, MEF 7, MEF 10.1 IEEE 802.1Qay
Provisioning based on SLA components—CIR, frame loss, delay, and jitter	Defined bandwidth profile capability that enables provisioning traffic based on SLA attributes	MEF 3, MEF 10.1, IEEE 802.1ag, ITU Y.1731

Specifically, MEF 10.1 has defined a Bandwidth profile and also identified specific service-related performance parameters. It has also defined the algorithm to enforce QoS or performance by ensuring conformance to the bandwidth profile. This was discussed at some length earlier in the chapter.

Standardized Management

LAN Ethernet was most lacking in the area of standardized management; consequently, this has been the focus of considerable work in Ethernet's transformation to Carrier Ethernet. Specifically, Ethernet OAM has had to be developed from the ground up. The developments in this area are discussed at some length; Table 2.6 notes some of the developments in the standards bodies.

MEF 7 focuses on standardizing for Service Provider Element and Network Management Systems (EMS/NMS) to provision, configure, and fault manage Carrier Ethernet services. It also defines OAM at the Ethernet services layer; however, it does not define OAM at the transport link/network layers, and it complements the work done in the ITU, IEEE, and IETF at the transport data-link and network layers based on G.809.

MEF 7 also provides a framework and concepts for managing and monitoring flows across an end-to-end connectionless network, and it also provides mechanisms to perform node discovery, establish connectivity, monitor CoS, and detect service impairments.

Ethernet Operations, Administration, and Maintenance (OAM) One of the key prerequisites to wide-scale Ethernet service deployment in Service Provider networks is a comprehensive Operations, Administration, and Maintenance (OAM) capability. The need to support hundreds of thousands of customers who are already accustomed to the fairly stringent SLAs for ATM, Frame Relay, and private line services means that significant new management capabilities are necessary for Carrier Ethernet; Ethernet has traditionally been weak in this respect and the relatively lower demands for OAM within an enterprise LAN (often within a building) were easily

TABLE 2.6 Standards Efforts Enabling Standardized Management

Key Components	Major Developments	Reference
Unified management	Defined a framework to monitor and manage flows across a connectionless network (at the Ethernet layer)	MEF 7, MEF 15, ITU G.809
Carrier-class OAM	Carrier-class link level management and end-to-end service-level management defined	MEF 7, IEEE 802.3ah, IEEE 802.1ag, Y.1731
Rapid services provisioning	Defined local management interface to enable rapid provisioning and management	MEF 16

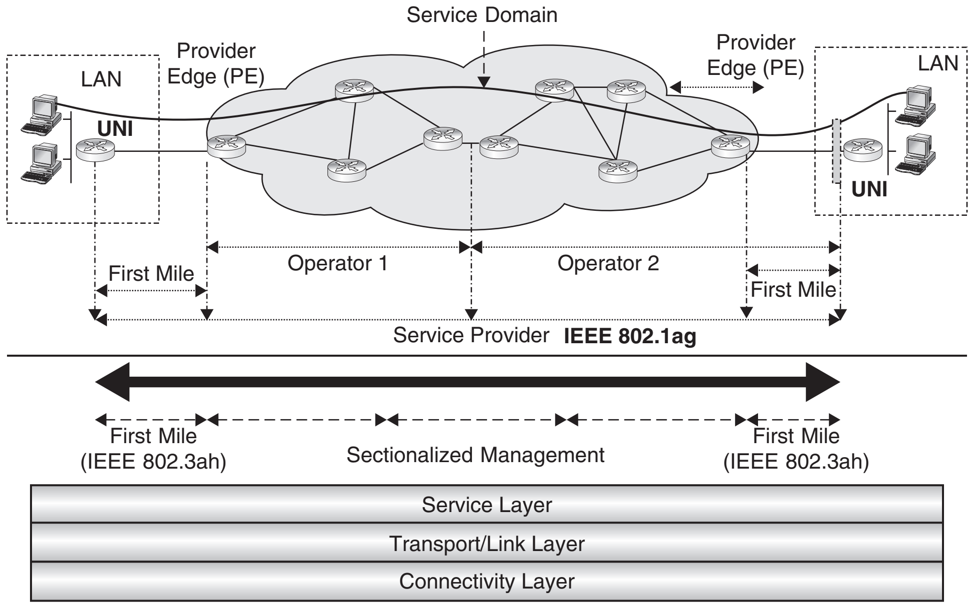

Figure 2.20 Ethernet OAM—a layered perspective

managed by the use of (less than efficient) Layer 3 protocols such as Simple Network Management Protocol (SNMP).[20]

As Carrier Ethernet is accelerating as a Carrier-class service delivered over multiple large and complex Service Provider networks, its OAM capabilities have to offer sophisticated tools to provision individual Ethernet services, monitor their performance, and identify and manage any issues quickly across such network topologies. Ultimately, this will lead to reducing the total cost of ownership, which is a prerequisite before Carrier Ethernet can meaningfully attempt to become a mass market service.

In attempting to define a comprehensive OAM capability for Carrier Ethernet, a layered approach is conceptually employed[21] to align with the layered nature of Service Provider networks used to deliver Carrier Ethernet. Each of the OAM layers delineate the different focus and functionality of the respective layer in the context of delivering Carrier Ethernet. This is shown in Figure 2.20. The three-layered OAM approach focuses on the service layer, the network/connectivity layer, and the transport/data-link layer. The OAM at each of the layers is independent of the other layers; however, they all employ standard Ethernet frames as the means of OAM-related communication.

[20] Ironically, these management protocols would not be usable without the Ethernet (layer 2) being operational. This scenario is somewhat ridiculous—when there is an issue in Layer 2, then the higher layer–based (i.e., Layer 3) management protocol is useless, defeating the very purpose of having a management capability.

[21] There is not yet a formally defined OAM-layered model available, but the ones employed are generally close.

Service Layer OAM at the service layer provides the capability to manage the entire Carrier Ethernet *service* being offered, i.e., a service instance represented as a uniquely identified Ethernet Virtual Circuit (EVC) offered between two or more customer UNIs. This end-to-end domain of the service—basically the customer domain—is ultimately what matters to the end-user experience, so here the OAM is focused on ensuring the service being offered is compliant with any agreed upon SLAs. The OAM, therefore, provides the ability to monitor the performance of a service continually, independent of the underlying network infrastructure. In addition, it also provides the capability to provision customer devices for services with specific performance and operational profiles. Both the IEEE 802.1ag and Y.1731 focus on service layer fault management, while Y.1731 augments with performance monitoring. The MEF specification 16 standardizes around the capability to provision the customer premise equipment by a service provider.

Connectivity Layer An Ethernet service is usually provided by a Service Provider over a physical network infrastructure; this infrastructure could belong to and be managed by one or more providers (or operators), each employing different network technologies to deliver Carrier Ethernet services (e.g., SONET, WDM, native Ethernet, MPLS, etc.). The OAM in this layer is concerned with the connectivity between the network elements that underpin the service delivery. In Figure 2.20, this encompasses the elements that exist between the boundaries of the Service Provider network (which, of course, could be comprised of networks belonging to multiple independent operators) and typically notated as being between the Provider Edge (PE) devices. Providing the capability to detect, troubleshoot, and *proactively* manage any issues emerging at this layer essentially means providing the ability to sectionalize any segment in the network quickly; thus an issue can be narrowed to a specific point in the infrastructure and quickly homed in on. Any issues at this layer will invariably have an impact on the higher service layer, and the specific impact (i.e., which service instances have been affected) on the management infrastructure needs to be identified. The IEEE 802.1ag and Y.1731 standards focus on this layer.

Transport/Data-Link Layer At the Data-Link layer, the OAM is focused on providing the capability to manage a single physical data link between two Ethernet interfaces; such links, of course, make up the network infrastructure, but the OAM capabilities on this layer are restricted to only individual physical links and include the ability to troubleshoot any issues employing loopbacks and monitor performance effectively.

Any impact on this layer manifests in possible issues at the higher (connectivity and service) layers, and robust capabilities to monitor, troubleshoot, and identify any issues are vital. The key standard in this area, the IEEE 802.3ah, focuses on the access link (first/last mile) of native Ethernet access networks. Multiple transport solutions for Ethernet can be employed, such as SONET, WDM, etc., and there are well-established OAM standards for these respective solutions.

Standards Work Key standard bodies such as the ITU, MEF, and IEEE and their respective standards/specifications are focused on developing OAM capabilities across

TABLE 2.7 Ethernet OAM Layers, Functionality, and Standards

Layer	Standards	Key Functions
Service	ITU Y.1731 IEEE 802.1ag MEF Spec 7, MEF Spec 16 (E-LMI)	Discovery Continuity check Loopback AIS/RDI (alarm indication signal/remote defect indicator) Traceroute Performance management
Connectivity	ITU-T Y.1731 IEEE 802.1ag	Discovery Continuity check Loopback AIS/RDI (alarm indication signal/remote defect indicator) Traceroute
Transport/ data-link	IEEE 802.3ah Misc. transport standards	Discovery Link monitoring Remote failure indication Remote, local loopback Fault isolation Performance monitoring

Source: ADVA Optical Networking

the three layers as shown in Table 2.7. Some of the key functions provided by the different standards at each of the layers are also briefly discussed.

Some of these functions and standards are focused beyond a single layer (IEEE 802.1ag/Y.1731, for example, is applicable to both the service and connectivity layers). Also, there are multiple standards in some of the layers. Generally, there is alignment between the respective standards efforts such that they mutually reinforce each other and do not conflict. Thus, both IEEE 802.1ag and ITU Y.1731 provide similar capabilities to monitor the service end-to-end.

- **Discovery** This function enables auto discovery and exchange of information pertaining to OAM and other capabilities between peer entities in a network (or on a link in the transport layer).
- **Continuity check** This function allows for continuous monitoring of a path (multiple hops) or a link between two endpoints using a periodic "I am alive" message exchange.
- **Loopback** This common function provides the ability to test whether a physical/virtual circuit is operating correctly. It essentially sends and receives a set of Ethernet frames to a remote point in the Service Provider network. If the remote location is a physical Ethernet port/facility, then the loopback will be intrusive (i.e., will impact regular data flow); a nonintrusive loopback can be initiated on a per-service instance (i.e., a specific EVC) basis.
- **Alarm Indication Signal (AIS) and Remote Defect Indicator (RDI)** This provides the capability of generating only one alarm message when an issue is

detected and ensures that other devices that receive the same alarm suppress duplicate notifications (very much like what's available on SONET and ATM).

- **Traceroute** This is a simple "ping-like" function that basically tests a specific multihop path across a Service Provider network (and likely across multiple operators' domains).
- **Performance monitoring** This functionality allows measurement of specific SLA parameters relating to a particular service instance (EVC) such as delay encountered, loss packets, jitter (or differential delay), and availability over a period. These measurements are on an end-to-end basis and closely reflect the performance of an actual service.

At the data-link layer, performance monitoring is limited across a physical link. This section briefly introduces some of the key OAM standards and highlights their essential characteristics.

IEEE 802.1ag Connectivity Fault Management (CFM) The IEEE 802.1ag is expected to be ratified in 2007 and enables Service Providers to manage individual EVCs, representing specific Ethernet services. Such management will be on an end-to-end basis across the network(s) over which the service is delivered. As such, this would require all the underlying equipment involved (and belonging to one or more operators) to also support the IEEE 802.1ag standard. The IEEE 802.1ag is closely aligned with the work on fault management from the Y.1731 standard.

The IEEE 802.1ag separates the Service Provider network—the one delivering the end-user service—into maintenance domains, which are each essentially managed/administered independently. These domains are typically hierarchical and encompass the three distinct entities that are involved in delivering a service: the customers using the service, the Service Provider delivering the service, and the operators whose networks may be used to deliver the service. Such a framework is useful in quickly homing in on—and resolving an issue.

The IEEE 802.1ag uses normal Ethernet frames to communicate between the different devices, with the only distinction being the use of a special Ethernet MAC address identifying it as an 802.1ag message. There are fours categories of messages that are employed to troubleshoot and manage Ethernets:

- **Continuity check messages (CCMs)** These are "I am alive" heartbeat messages that are issued periodically to identify any loss of service between two (equipment/devices) endpoints or intermediate points. Any erratic behavior in these messages could enable preemptively addressing any emerging issues.
- **Link trace messages** These messages are used by a Service Provider to track a specific path between two pieces of equipment/devices traversing through the intervening devices. This hop-by-hop approach is useful in identifying whether a data path exists or not.
- **Loopback messages** These allow a Service Provider to validate connectivity (either on a service or a circuit basis) to a particular maintenance point to

determine whether it is reachable or not, without particularly worrying about the intermediate nodes.

- **Alarm indication signal (AIS) messages** These messages are used to indicate that there is a fault in the network.

ITU Y.1731 The ITU Standards Group (SG) 13, in Recommendation Y.1731, identifies the OAM functions in an Ethernet network that are needed to allow fault management and performance monitoring. Fault management is closely aligned with the fault management capability of the IEEE 802.1ag (and hence includes capabilities such as discovery, continuity checks, loopbacks, link trace, etc.). However, the Y.1731 augments this with performance monitoring as well.

Performance monitoring allows the measurement of typical SLA parameters around error counts and delay measurements such as loss of Ethernet frames, delay between frames, variation between consecutive delays (also known as jitter), and other information such as link up or down, throughput, and so on.

Currently, the Y.1731 standard supports performance monitoring only for address Point-to-Point connectivity at this time (multipoint connectivity is expected in the next phase).

The ITU group is working closely with IEEE 802.1ag group to ensure alignment and preclude any conflicting approaches.

IEEE 802.3ah Ethernet over First Mile The IEEE 802.3ah OAM is also known as Ethernet First Mile (EFM) OAM and provides OAM between the Ethernet ports at the CPE and the Provider Edge (the "first mile"), which is deployed over a physical IEEE 802.3 medium (copper, fiber, or PON). In fact, the IEEE 802.3ah also addresses the PHY (physical) layer characteristics for the different media in the first mile; the OAM part of the IEEE 802.3ah is, however, independent of the physical layer. EFM OAM is the first standards-based effort to ensure Ethernet devices in Service Provider networks have an inherent management capability.

The IEEE 802.3ah was ratified in 2004 and was expected to complement existing protocols such as SNMP that were otherwise being employed for management purposes.

The EFM OAM also uses Ethernet frames (albeit with a specific destination MAC address and the Ethernet type/length field to identify EFM-related frames uniquely (PDUs). It is also an in-band protocol (i.e., it uses the same bandwidth as the data frames) and is characterized as a slow protocol; it is not required for normal operation and typically uses about 10 frames per second.

The EFM OAM addresses some fundamental aspects necessary when deploying Ethernet over the first/last mile:

- **Link monitoring** Gives the Service Provider visibility of the first mile physical connection through periodic heartbeat messages. In case of any issues on this link, the Service Provider is immediately notified with pertinent information.
- **Fault signaling** Enables a device to convey to its peer at the remote location that severe conditions such as link failure (noted because it can no longer receive any signal) or a dying gasp (when the remote device is about to be powered down and operationally unavailable) have occurred.

- **Remote loopback** Enables a loopback to be initiated from one entity to a remote peer entity to ensure the quality of the intervening Ethernet circuit (specific tests for delay, jitter, and so on, can also be measured).
- **MIB variable retrieval** Provides a management information base (MIB), which is a database of management variables and typically includes all performance and error statistics maintained on an Ethernet link. The IEEE 802.3ah Ethernet OAM provides a read-only access remote MIB (and does not allow the variables to be set).

Organization Specific Extensions The IEEE 802.3ah OAM also allows equipment vendors to extend Ethernet OAM capabilities through organizational-specific PDUs to support additional capabilities, such as extending OAM messages beyond one link and monitoring other equipment performance parameters, that will contribute to offering more robust Ethernet services.

Ethernet-Local Management Interface (E-LMI) The E-LMI, specified in the MEF 16, defines the protocol to communicate service-level information to enable the automatic configuration of the Customer Premise Equipment (CPE). This ability allows the Service Provider to ensure, remotely, that the CPE is set up correctly to support a specific Ethernet service, rather than have the enterprise administrator configure it. Basically, the entire configuration for a specific service is downloaded into the CPE from the provider edge device using E-LMI. Specifically, the E-LMI provides the following capabilities to a Service Provider:

- Add or delete an Ethernet Virtual Circuit (i.e., an Ethernet service instance) in the CPE.
- Inform the status of an already configured EVC, specifically whether it is available or not.
- Verify the integrity of the link between the Provider Edge (PE) and the CPE.
- Ensure that the UNI and EVC attributes are correctly passed to the CPE.

Carrier Ethernet: Field Realities

While Carrier Ethernet is being embraced quite aggressively—evident by the number of Service Providers offering these services and also by the promising growth predicted, it is important to note that it (Carrier Ethernet) still accounts for a relatively small portion of the addressable market. In fact, a study by the Vertical Systems Group (VSG) indicates that it makes up less than 5 percent of business service spending on telecom services.

As Carrier Ethernet services are beginning to grow, they will invariably have to address the Small and Medium Enterprises (SME) that make up the larger part of the enterprise market opportunity and are represented graphically by the lower part of the pyramid in Figure 2.21. Essentially, this segment of the market is comprised of a much larger number of (relatively smaller) end customers as compared to the initially

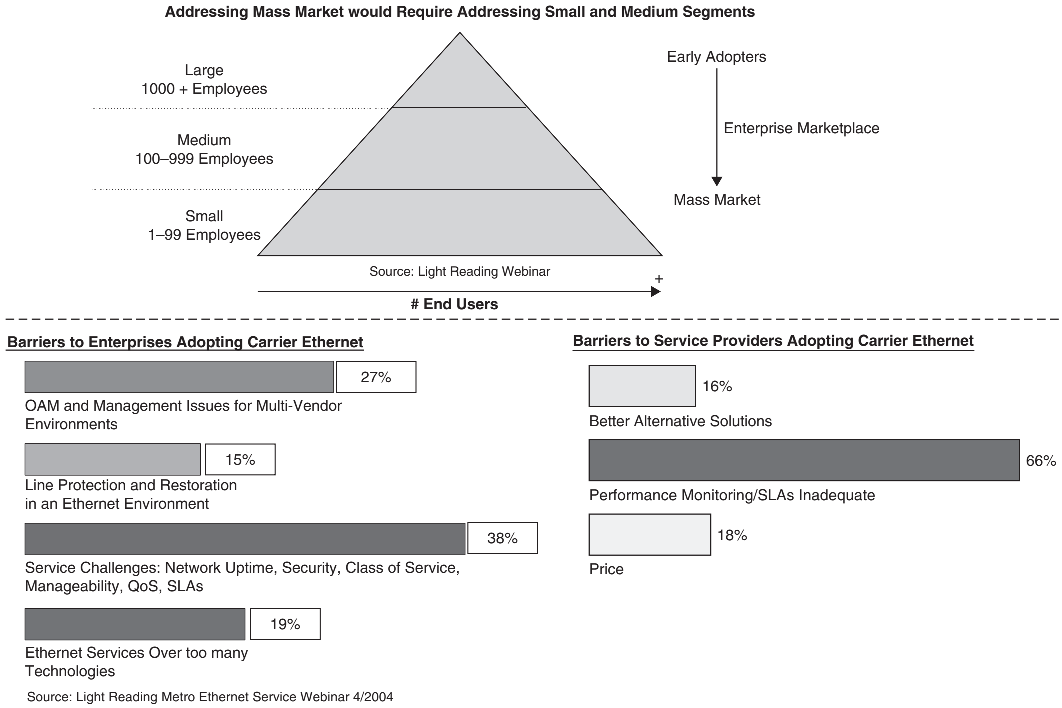

Figure 2.21 Key barriers to Carrier Ethernet today

addressed market represented by the top of the pyramid. Addressing the needs of this customer base economically is, therefore, a prerequisite to fueling Carrier Ethernet toward becoming a true mass market service. There are, however, several—some unique—challenges to addressing this segment of the market. These and other issues are discussed next, followed by an overview of the industry response.

Current Challenges in Delivering Carrier Ethernet

The challenges in delivering Carrier Ethernet, especially to the SMEs, are noted in the different studies depicted in Figure 2.21. The key issues are distilled in the sections that follow.

Availability of Carrier Ethernet Services While SMEs are increasingly aware of the benefits of Ethernet services, a big issue is the availability of such services. Specifically, the following issues pose a barrier to subscribing to Carrier Ethernet services.

Fiber Availability Carrier Ethernet services are being delivered significantly in a native fashion over a fiber infrastructure. Given that, according to Vertical Systems

Group, only about 11 percent of the buildings in the U.S. are connected to fiber. This definitely limits coverage, especially because most of the enterprises (and notably the SMEs) are the ones that occupy such buildings. The significant cost of laying fiber (actually the cost of regulatory approvals and delays) has slowed down the process considerably, although as discussed in Chapter 16, Service Providers are finally beginning to proceed fairly aggressively.

Lack of Availability at All Locations It is not uncommon for SMEs to have multiple offices physically served by different Service Provider networks because they are located physically across more than one Service Provider's footprint (e.g., a SME has offices at location A, B, and C, and the respective telecom services are being delivered, perhaps due to regulatory and/or competitive constraints, by Service Providers 1, 2, and 3, respectively. Note that locations can be in the same or different cities and/or countries).

Such SMEs frequently do not have Ethernet services offered at each of their locations (perhaps they are not served by fiber or there could be other competitive and economic reasons) and consequently, do not subscribe to Ethernet services at all.

Lack of Key Carrier Attributes The SME customer base (over 95 percent as stated earlier) is often served by legacy ATM, Frame Relay, and private line services today to support their voice, data, and video applications. While they recognize the value of Carrier Ethernet (see Chapter 1 for a detailed listing of the benefits), there is some hesitancy to migrate to Carrier Ethernet due to the (perceived[22]) lack of service features that are deemed important and that they're accustomed to with their legacy services. This feature deficiency is depicted in Figure 2.19 as well, the most important being the lack of service-level agreement (SLA) monitoring and, more generally, OAM capabilities.

SLA Monitoring As SMEs are considering Carrier Ethernet as the convergent access, and hence relying on it for their mission-critical applications as well (storage backup, voice, etc.), it is imperative that SME customers have the assurance that the underlying Carrier Ethernet services are performing according to stringent SLA requirements. With private line and other technologies, they have that capability; with Carrier Ethernet, the absence of such SLA measurement capabilities precludes its adoption.

Lack of OAMs As Service Providers are required to deliver Carrier Ethernet to a substantially greater number of individual customers (the SMEs), they have to make it an economically viable offering with adequate profitability margins. In order to provide Carrier Ethernet to a mass market, the capital expenditures (CAPEX) and, more importantly, their operational expenditures[23] (OPEX) need to be addressed. A significant contributor

[22] As will be evident in Part II, numerous Ethernet solutions do offer most of the Carrier-class attributes.

[23] It has been estimated that over 70 percent of the total cost of ownership of delivering a service is comprised of the OPEX. Hence, a reduction in OPEX has considerably higher impact on reducing cost.

to OPEX for Ethernet services currently is the largely manual and time-consuming effort required to manage these services—from provisioning to fault notification, troubleshooting to managing any issues that emerge. The key reason for this situation is the lack of sophisticated tools and features to manage Ethernet services, especially on an end-to-end service level; as noted in the previous section of this chapter, this Carrier Ethernet services' shortcoming has been the focus of the initial standards efforts.

Bandwidth Demand Curve While newer bandwidth intensive applications such as video make Carrier Ethernet the natural solution for enterprises, it must be pointed out that a significant portion of the addressable market—nearly 95 percent according to Vertical Systems Group—is served by T1 (64 Kbps) connections. While bandwidth demand at these SMEs is indeed growing, it is not quite jumping to 10 Mbps—the typical Carrier Ethernet service offering. Even though Carrier Ethernet is designed to offer bandwidth anywhere from 1 Mbps in fine increments of 1M or even less (in fact, this is considered one of its big advantages as noted in Chapter 1), less than 10 M is not in reality being offered. This speaks to enforcing the Carrier Ethernet defined UNI to leverage the market opportunity.

Economics One big advantage of carrier Ethernet services is the economics for both the Service Providers and enterprise end users. However, as these services are currently being delivered over numerous underlying technologies (refer to Part II for a discussion on these), the economics may be less attractive (as opposed to delivering native Ethernet). Further, because pricing of Carrier Ethernet services is combined with other application services such as Internet access, the true cost of Carrier Ethernet is hard to discern.

Interoperability In a LAN Ethernet, enterprise customers have come to expect that any device, from any manufacturer with a standard Ethernet port, can be easily deployed in their LAN. A similar expectation is assumed by Service Providers when it comes to Carrier Ethernet; after all, their networks are akin to a LAN and any Carrier Ethernet equipment from multiple vendors deployed over these networks should inter-work and provide consistent services and, of course, offer the features and tools to provision and manage these services.

The Standardization Efforts While the standardization efforts have made significant headway, this is a work in progress and still very much in the early stages. Functionalities such as Network to Network Interface (NNI)—which defines the handoff between two Service Providers and is a key requirement for wholesale Carrier Ethernet services or when the delivery infrastructure is leased from one or more network operators—needs to be formally defined before Carrier Ethernet will be deployed more aggressively in the WAN. This effort is still underway at the standards bodies.

Recent Industry Response to Challenges

Given the economic and competitive attractiveness of Carrier Ethernet services, the industry has naturally embarked on addressing some of the challenges noted above.

Two specific ones—Intelligent Demarcation and the MEF Certification Program, are discussed here, and illustrate the considerable and effective effort in this regard.

Intelligent Ethernet Demarcation

One industry response to enabling the acceleration of Carrier Ethernet has been to introduce a new class of intelligent *Ethernet demarcation devices (EDDs)*.[24] These are also referred to as *network termination units* or *network termination elements (NTEs)*.

NOTE While these devices may be standalone, as is the case currently, this functionality may well be integrated into other edge devices (such as switches, routers, ADMs, etc.) as well.

These devices typically reside at the Customer Premise (CP) or Customer Edge (CE), and in addition to serving as a physical point of separation between the Service Provider and user networks (typically a LAN),[25] they provide three key functional capabilities:

- A standardized Ethernet UNI
- Ethernet OAM
- Media/protocol conversion

The standardized Ethernet UNI essentially provides the MEF-defined capabilities that include an IEEE 802.3 handoff (PHY), provisioning, and enforcing a bandwidth profile for any EVCs initiated there, along with the CoS, and any service multiplexing necessary.

The Ethernet OAM provides end-to-end visibility of the Ethernet service(s) and the associated performance SLAs. It also encompasses sophisticated and proactive fault notification so that any potential issues can be addressed remotely before they manifest more broadly. The OAM provides troubleshooting tools to enable such a capability. Most of this is based on the IEEE 802.1ag and ITU Y.1731 standards and can also measure typical SLA components (such as delay, jitter, frame loss, etc.) on a per-service (EVC).

The OAM capability, in addition to ensuring that the Ethernet services are being delivered per the SLAs, also reduces the Service Provider OPEX by providing the ability to address most of the typical service issues remotely (and thereby precluding expensive truck rolls).

Finally, the media conversion capability provides a standardized UNI to the customer while supporting a host of last/first mile transport technologies and media to

[24] This is not particularly unique; earlier technologies such as Private Line and ATM/Frame Relay addressed similar barriers to wide–scale deployment by introducing demarcation devices. It was almost natural that Carrier Ethernet followed suit.

[25] This physical demarcation between the Service Provider and the subscriber/customer also signifies where the responsibility of a Service Provider ends in terms of identifying and resolving any issues. Anything beyond the EDD (toward the customer) is the responsibility of the customer.

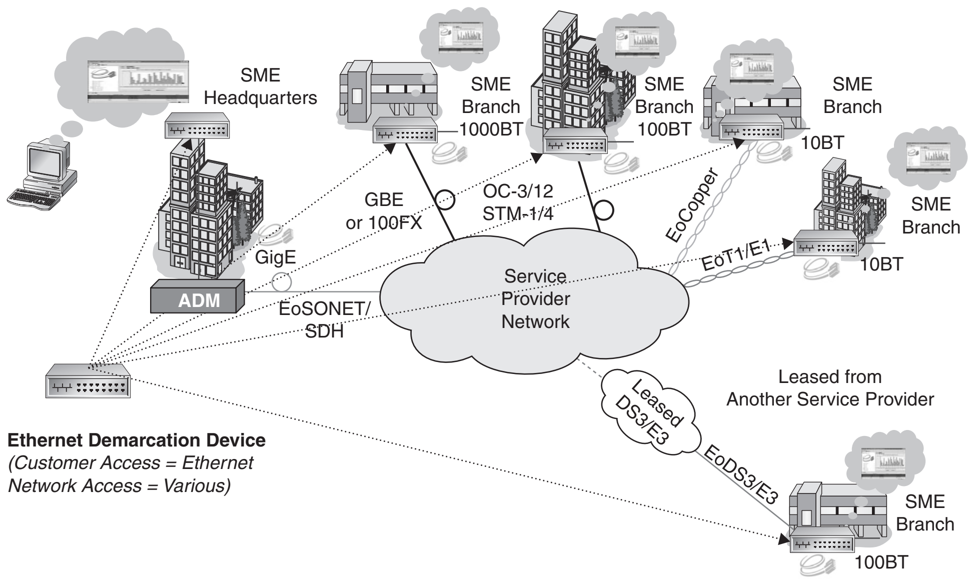

Figure 2.22 Use of Ethernet demarcation to provide Carrier Ethernet services to the mass market (Source: ADVA Optical Networking)

the Service Provider network (such as T1, DS-3, OC-3, OC-12, etc); thus, the Ethernet handoff is "converted" to whatever the last mile transport technology is.[26] What this means is that now a standardized Carrier Ethernet handoff can be provided to a customer independent of the last mile infrastructure.

Figure 2.22 depicts the use of Ethernet demarcation in a real-life scenario. In this example, a reasonably large enterprise customer with several physical locations, each of which are served by different last/first mile infrastructures, requires Carrier Ethernet services. Some of the locations are served by old SONET ADMs that have no Ethernet capability. By introducing Ethernet demarcation, such issues are addressed and the customer is provided a standardized Ethernet UNI, with the same look and feel at all locations.

Thus, Ethernet demarcation is enabling the delivery of Carrier Ethernet services despite the challenges of fiber shortage and the presence of a host of last mile infrastructures that may not always be amenable to delivery of such services (e.g., older SONET ADMs are usually not equipped with Ethernet interfaces). Further, it is important to note that Ethernet demarcation devices also enable Ethernet services quickly (i.e., speed to market), relatively easily (i.e., it is easy to augment current last mile technology solutions with a standalone EDD), and ultimately cost effectively. The IEEE 802.1aj (two-port relay) effort is considering standardizing such a functionality.

[26] Employing standard techniques such as encapsulating using GFP (Generic Frame Protocol) for carrying over SONET, etc.

The MEF Certification Program

Carrier Ethernet is designed to scale from a local to a ubiquitous worldwide service that could span thousands of offices and hundreds of countries and where everything works together harmoniously. It is therefore no small matter for Service Providers to offer Carrier Ethernet services that would enable mission critical applications, and delivered over a variety of transport technologies (discussed in Part II). Often Service Providers would have to cooperate with other Service Providers/Network Operators to offer these Carrier Ethernet services across the MAN/WAN and this would almost invariably entail equipment from several vendors.

Conversely, to the enterprise user, Carrier Ethernet services must work as simply as plugging in an Ethernet cable and powering up.

The MEF Certification Program was conceived explicitly to address the underlying challenges that inherently exist in simplifying the deployment, while ensuring consistency (of Carrier Ethernet services) in a multi-vendor environment. In so doing, the goal is to accelerate the deployment of Carrier Ethernet.

The program commenced in April 2005 and essentially consists of a series of thorough tests providing evidence for end-users, service providers and manufacturers alike, that products and services are compliant to published MEF specifications.

It initially certified equipment (systems) that it delivers MEF-compliant Ethernet services. This program subsequently also began certifying that Service Provider–delivered Ethernet services are also consistent with the MEF Carrier Ethernet specifications.

NOTE The MEF does not conduct the certification directly but rather works with an independent testing entity, Iometrix, for conducting the actual testing and validating compliance.

Thus, whether it is a Service Provider evaluating equipment for delivering Carrier Ethernet or end users assessing Carrier Ethernet services, knowing that the underlying equipment or service is MEF-compliant expedites deployment. Specifically, the MEF certification program offers the following benefits to the three main constituents that drive Carrier Ethernet:

Enterprises end users:

- Provides a common basis/terminology to meaningfully compare services from different Service Providers.
- Empowers informed decisions regarding equipment/CPE purchases and minimize risk.
- Assures that Ethernet services perform according to pre defined specifications and standards.
- Ultimately benefits from the efficiencies and cost savings to the Service Providers, which are usually passed on to the end users.

Service Providers:

- Immediate assurance that vendor's equipment complies to MEF specifications.
- Saves money and time on complex testing between vendors, especially on global accounts.
- Establishes solid foundation for Carrier Ethernet ubiquity and interoperability.
- Removes confusion caused by proprietary names and descriptions
- Conformance to MEF 9 allows customers to specify their service requirements unambiguously using standards.

Equipment vendors/Manufacturers:

- Globally recognized interoperability standard improves approval process
- Increases tender opportunities and competitiveness.
- Independent validation of function and conformance that their equipment is MEF compliant; this helps with positioning and deployment at Service Provider customers.
- Dramatically reduces testing costs, time-to-market, as well as installation time.
- Provides a performance and behaviour benchmark.
- It forms the basis for RFP requests and helps manufacturers focus on their features that distinguish them from competition

The MEF certification program was rolled out in two phases and has focused on MEF 9 and MEF 14:

- **Phase 1** The focus here was on equipment and systems that deliver Carrier Ethernet, specifically on whether they are compliant with the MEF-defined services. Thus far hundreds of systems from over 45 vendors have been certified for MEF 9 (Abstract Test Suite for Ethernet Services at the UNI; the Ethernet services are defined in MEF 6); certification for MEF 14 (Service Quality) is also now underway and numerous vendors — over 35, have already been certified as well.
- **Phase 2** This is focused on ensuring that the Carrier Ethernet services offered by Service Providers are compliant with the MEF specifications. The first set of over 15 Service Providers was certified for MEF 9; this guarantees that the E-LINE and E-LAN from these Service Providers will be compliant with MEF. Eleven Service Providers have been certified for MEF 14 as well.

The certification program is extending to testing Traffic Management (MEF 10.1) according to the definitions in MEF 7.

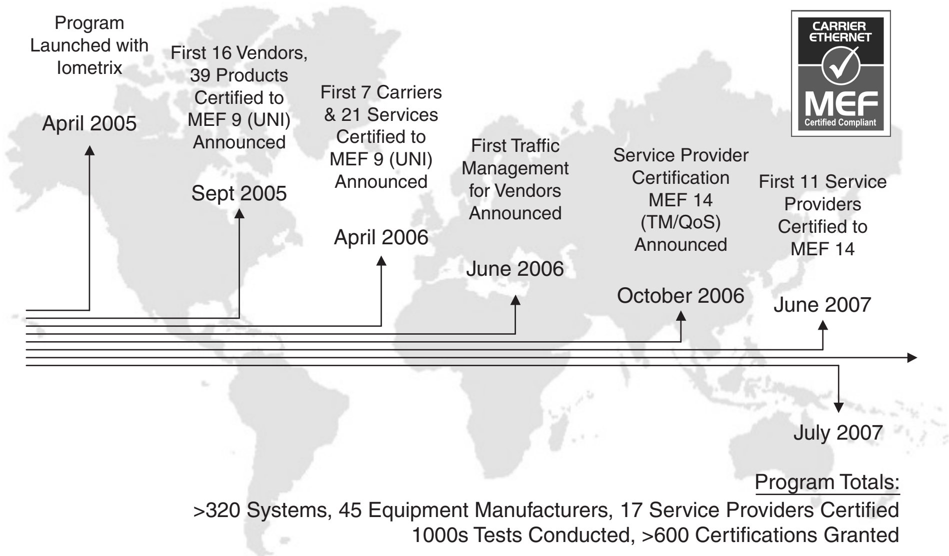

Figure 2.23 The MEF Certification Program—key milestones to date (Source: MEF)

Figure 2.23 depicts the extent of progress made in the MEF Certification Program, as of July 2007. To date 17 Service Providers and 45 equipment vendors with 320 systems have been certified.

MEF has also recently introduced two new technical test specifications, MEF 18 (Abstract test suite for CES over Ethernet services) and MEF 19 (Abstract test suite for UNI Type 1), and is working on developing potentially test suites for E-NNI and LMI.

Other Carrier Ethernet requirements—One Service Provider's perspective

While the MEF has made considerable strides in the realm of identifying and refining the Carrier Ethernet attributes, it is, of course, a work in progress. Emerging applications, field experience, and new network constraints/requirements continually push the boundary and need to be addressed if Carrier Ethernet is to dominate the market.

Here a brief overview of the requirements that Carrier Ethernet faces — or will face shortly, is provided and is based on the experience and insights of one of the foremost Carrier Ethernet Service Providers, Verizon[27]. That most of these requirements are already being actively addressed (or at least being considered) by the MEF vividly demonstrates the unprecedented participation—and influence, of Service Providers in MEF.

[27] Verizon identifies these requirements in the context of its four key drivers for Carrier Ethernet—Business Ethernet services, Broadband access, Residential video service transport, and wireless backhaul.

Standardized Services

- A common set of Class of Service (CoS) definitions and associated performance guarantees, bandwidth increments for each CoS and maximum frame service need to be developed (some work has already begun in MEF).
- Equipment (both customer and network) needs to incorporate configuration management as defined in the E-LMI specification (MEF 16). This would enable such equipment to automate, and hence simplify configuring the increasingly sophisticated services.
- Timing synchronization specifications have gained some urgency as carriers are beginning to migrate TDM services to converged Ethernet networks. (MEF has begun addressing in the mobile backhaul project).

Scalability

- A standard Carrier to Carrier interconnection specification is required (The E-NNI effort has commenced by the MEF)
- A standardized access interconnection for emerging access methods is required (MEF has begun developing the Service Node Interface, SNI)
- A dynamic control plane solution is required to enable automated provisioning.
- Overcome the VLAN/MAC limitations (actively being addressed by the PBB, MPLS etc as noted previously in the chapter)
- Efficiently forcing customer specific traffic only on some backbone links (Multiple Registration Protocol, MRP, per IEEE 802.1ak has begun focusing on this capability).

Reliability

- Standardized SNI required to provide an efficient method of introducing resilient access solutions in the metro. (MEF work underway)
- While fault management has been well defined, it is yet to be implemented in commercial equipment solutions.

Quality of Service

- Topology discovery tools in Network Management Systems to support Connection Admission Control (CAC) in an Ethernet network.

NOTE CAC is usually provided by Service Provider Provisioning systems, and so the Network Management systems should coordinate with Provisioning systems to ensure the delivery of stringent QoS.

- A distributed control plane is required to support large Carrier Ethernet networks.
- CoS awareness in Layer 1 transport devices in access networks is required to preclude any speed mismatches.

Standardized Management

- Standards available but need to be implemented in commercial solutions (most of them support pre-standard versions)
- Require enhanced version of Link Aggregation that distributes Ethernet frames based on VLAN ID (not just on MAC/IP addresses; this forces service-related and associated OAM frames are pinned to same links through out the network).

As should be clear, some challenges are already being identified by forward looking Service Providers such as Verizon; in most cases, it must be noted that proprietary solutions have been adopted in the interim to address the challenges, as to not impede the progress of Carrier Ethernet deployment.

The service attributes, their respective definition and parameters for the MEF-defined UNI and EVC, which were discussed earlier in this chapter, are shown in Figure 2-24.

EVC Attributes

Service Attribute	Service Attribute Parameters
EVC Type	Point-to-Point or Multipoint-to-Multipoint
UNI List	A list of UNIs (identified via the UNI Identifier service attribute) used with the EVC
CE-VLAN ID Preservation	Yes or No. Specifies whether customer VLAN ID is preserved or not.
CE-VLAN CoS Preservation	Yes or No. Specifies whether customer VLAN CoS (802.1p) is preserved or not.
Unicast Service Frame Delivery	Specifies whether unicast frames are Discarded, Delivered Unconditionally or Delivered Conditionally
Multicast Service Frame Delivery	Specifies whether multicast frames are Discarded, Delivered Unconditionally or Delivered Conditionally
Broadcast Service Frame Delivery	Specifies whether broadcast frames are Discarded, Delivered Unconditionally or Delivered Conditionally
Layer 2 Control Protocol Processing	Discard or Tunnel per Protocol
Service Performance	Specifies the Frame Delay, Frame Jitter and Frame Loss per EVC or frames within an EVC Identified via their CE-VLAN CoS (802.1p) value

UNI Attributes

Service Attribute	Service Attribute Parameters
UNI Identifier	A string used to identity of a UNI, e.g., NYCBldg12Rm102Slot22Port3
Physical Medium	Standard Ethernet PHY
Speed	10 Mbps, 100 Mbps, 1 Gbps or 10 Gbps
Mode	Full Duplex or Auto Negotiation
MAC Layer	IEEE 802.3-2002
Service Multiplexing	Yes or No. Defines whether multiple services can be on the UNI
UNI EVC ID	A string used identify an EVC, e.g., NYCBldg1Rm102Slot22Port3EVC3
CE-VLAN ID/EVC Map	Mapping table of customer VLAN IDs to EVC
Max. Number of EVCs	The maximum number of EVCs allowed per UNI
Bundling	No or Yes. Specifies that one or more customer VLAN IDs are mapped to an EVC at the UNI
All to One Bundling	No or Yes (all customer VLAN IDs are mapped to an EVC at the UNI).
Ingress Bandwidth Profile Per Ingress UNI	None or <CIR, CBS, EIR, EBS>. This Bandwidth profile applies to all frames across the UNI.
Ingress Bandwidth Profile Per EVC	None or <CIR, CBS, EIR, EBS>. This Bandwidth profile applies to all frames over particular EVC.
Ingress Bandwidth Profile Per CoS ID	None or <CIR, CBS, EIR, EBS>. This Bandwidth profile applies to all frames marked with a particular CoS ID over an EVC.
Layer 2 Control Protocol Processing	Discard, Peer or Pass to EVC per protocol

Figure 2.24 EVC and UNI service attributes and definitions (source: MEF)

They essentially represent how sophisticated the UNI and EVC can potentially be at the current time; of course, in time, one should expect these to evolve to accommodate new requirements imposed by forthcoming applications.

References

1. Extensively used MEF material in this chapter is reproduced with permission of the Metro Ethernet Forum.
2. MEF technical specifications 2, 3, 4, 6, 7, 8, 10.1, 11, 12, 13, 14, 15, 16, 17, 18, 19: www.metroethernetforum.org.
3. IEEE 802 specifications: www.getIEEE802.org.
4. ITU specifications: www.itu.int/publications/default.aspx.
5. IETF specifications: www.ietf.org
6. "Business Ethernet—The Game Plan for 2007," MEF, Vertical Systems Group, January 2007.
7. Provider Backbone Bridging (PBB) and Provider Backbone Transport, Nortel: www.nortel.com (http://www2.nortel.com/go/solution_assoc.jsp?segId=0&parId=0&catId=0&rend_id=17102&contOid=100188013&prod_id=55120).
8. Carrier Ethernet: A Reality Check by Stuart Elby, Haidar Chamas, William Bjorkman, Vincent Alesi, Verizon, NFOEC, March 2007.