

Bandwidth Profiles for Ethernet Services

Ralph Santitoro

Abstract

This paper provides a comprehensive technical overview of bandwidth profiles for Ethernet services, based on the work of the Metro Ethernet Forum (MEF) Technical Committee. The paper is intended to help buyers, users, providers of Ethernet services, and equipment and semiconductor vendors understand the various types, characteristics and usage of bandwidth profiles as defined by the MEF. This paper will be updated as new work emerges from the MEF Technical Committee.

Introduction

Bandwidth profiles allow a service provider to bill for bandwidth usage and engineer their network resources to provide performance assurances for in-profile Service Frames. Bandwidth profiles enable a service provider to offer bandwidth to subscribers in increments less than the UNI (physical port) speed. Such granularity allows subscribers to purchase the bandwidth they need and allow service providers to price services more incrementally than TDM-based services.

“Bandwidth profiles ... allow subscribers to purchase the bandwidth they need and allow service providers to price services more incrementally”

Bandwidth profiles enable a service provider to offer multiple service instances per UNI and each service can have its own bandwidth profile. When compared to TDM-based services, this flexibility enables service providers to achieve higher profit margins with lower operational expenses while providing subscribers with more cost effective services.

Figure 1 illustrates the step function that occurs for TDM interfaces and non-Ethernet Layer 2 services as one increases bandwidth. The vertical axis indicates how the physical TDM interface changes as bandwidth increases. This requires the subscriber to replace equipment or interfaces cards as bandwidth needs cross bandwidth thresholds determined by the TDM digital hierarchy.

Similarly, Layer 2 services such as Frame Relay or ATM also have thresholds which require the subscriber to change equipment or service protocols once their bandwidth needs cross certain bandwidth thresholds. Frame Relay is predominantly used at speeds less than T1

(1.544 Mbps) although some providers offer it at T3 (45Mbps) speeds. ATM is predominantly used at T3 to OC-3/STM-1 speeds although some providers offer it at T1 or E1 (2.048Mbps) speeds. Finally, for higher bandwidth needs, SONET/SDH is the predominant physical layer technology used with the Packet over SONET (POS) protocol for packet-based services.




Figure 1: Service Bandwidth per TDM interface

Figure 2 illustrates an Ethernet service using a 100Mbps and 1Gbps Ethernet interface. In both cases, the same Ethernet protocol is used and hence as a subscriber’s bandwidth needs cross the 100Mbps threshold, a new interface card may be needed. Note that some implementations today support 10Mbps, 100Mbps and 1Gbps over the same interface card so subscriber equipment need not be replaced or upgraded.




Figure 2: Service Bandwidth per Ethernet interface

“Bandwidth profiles enable a service provider to offer multiple service instances per UNI ... to achieve higher profit margins with lower operational expenses...”

As subscriber bandwidth needs increase, Ethernet allows more granular bandwidth provisioning without introducing new protocols or networking technologies and often without adding new physical interfaces.

Ethernet Service Fundamentals

This section provides some fundamental definitions and terminology that will be used throughout this paper.

Ethernet User Network Interface

The basic model for Ethernet services is shown in Figure 3. Ethernet Service is provided by the Metro Ethernet Network (MEN) provider. Customer Equipment (CE) attaches to the network at the Ethernet User-Network Interface (UNI) using a standard 10Mbps, 100Mbps, 1Gbps or 10Gbps Ethernet interface.




Figure 3: Basic Model for Ethernet Service

Ethernet Virtual Connection

An Ethernet Virtual Connection (EVC) is defined by the MEF as ‘an instance of an association of two or more UNIs’, where the UNI is a standard Ethernet interface that is the point of demarcation between the Customer Equipment and service provider’s MEN. EVCs help conceptualize the service connectivity. The MEF has defined two types of EVCs, namely, Point-to-Point (refer to Figure 4) and Multipoint-to-Multipoint (refer to Figure 5).



Figure 4: Point-to-Point EVC




Figure 5: Multipoint-to-Multipoint EVC

Bandwidth Profiles

From a subscriber perspective, a Bandwidth Profile specifies the average rate of ‘committed’ and ‘excess’ Ethernet Service Frames allowed into the provider’s network at the UNI. Service Frames sent up to the ‘committed’ rate are allowed into the provider’s network and delivered per the service performance objectives, e.g., delay, loss and availability, specified in the Service Level Agreement (SLA) or Service Level Specification (SLS). These Service Frames are referred to as ‘in-profile’ or ‘conformant’ to the bandwidth profile.

Service Frames sent up to the ‘excess’ rate are allowed into the provider’s network but are delivered without any service performance objectives. These Service Frames are referred to as ‘out-of-profile’ or ‘non-conformant’ to the bandwidth profile.

Finally, Service Frames sent above the ‘excess’ rate are discarded.

“... a Bandwidth Profile specifies the average rate of ‘committed’ and ‘excess’ ... frames allowed into the provider’s network...”

Service Frame Color

A useful way to describe Service Frames when their average rate is in-profile or out-of-profile is through the use of colors. The ‘color’ of the service frame is used to identify the bandwidth profile conformance of a particular Service Frame.

A Service Frame is 'green' if it is conformant with 'committed' rate of the bandwidth profile. A Service Frame is 'yellow' if it is not conformant with the 'committed' rate but conformant with the 'excess' rate of the bandwidth profile. Finally, a service frame is 'red' if it is conformant with neither the 'committed' nor 'excess' rates of the bandwidth profile.

Significance of Color

Green Service Frames are delivered per the service performance objectives specified for the service, e.g., delay, loss, etc. Green Service Frames, in general, should not be discarded because they are in-profile and conformant with the bandwidth profile.

Yellow Service Frames are out-of-profile but are typically not immediately discarded. However, yellow Service Frames are not delivered per the service performance objectives and may get discarded by the network under different conditions, e.g., network congestion.

Red Service Frames are also out-of-profile and are immediately discarded.

The MEF Technical Committee has not yet specified how color marking is indicated.

"The 'color' of the service frame is used to identify the bandwidth profile conformance of a particular Service Frame."

Bandwidth Profile Parameters

The MEF has defined bandwidth profiles that consist of the following four parameters:

- CIR (Committed Information Rate)
- CBS (Committed Burst Size)
- EIR (Excess Information Rate)
- EBS (Excess Burst Size)
- CM (Color Mode)

CIR and CBS

The Committed Information Rate (CIR) is the average rate up to which Service Frames are delivered per the service performance objectives. Such frames are referred to as being 'CIR-conformant'. The CIR is an average rate because all Service Frames are always sent at the UNI speed, e.g., 10Mbps, and not at the CIR, e.g., 2Mbps. CBS is the maximum number of bytes allowed for incoming Service Frames to still be CIR-conformant. CIR-conformant Service Frames are colored green.

Service Frames whose average rate is greater than the CIR are not CIR-conformant and are either colored yellow (if EIR is non-zero) or are discarded (if EIR=0).

A non-zero CIR may be specified to be less than or equal to the UNI speed. If multiple bandwidth profiles are applied at the UNI, the sum of all CIRs in each bandwidth profile must be less than or equal to the UNI speed.

A CIR of zero indicates that the service provides no bandwidth or performance assurances for delivery of subscriber Service Frames. This is often referred to as a 'best effort' service whereby all Service Frames are yellow (eligible to be discarded).

The significance of CIR and CBS is illustrated in the following two simple examples for a service that provides a 10Mbps UNI, a CIR=2Mbps (=2,000,000 bps) and a CBS=2KB (=2000 bytes). To simplify the discussion, this example service specifies EIR=0 and EBS=0.

CIR and CBS Example 1

Two 1518 byte Service Frames are sent back to back. The first service frame depletes 1518 bytes of the initial 2KB CBS in the token bucket leaving 422 bytes remaining. This service frame is in-profile (green) and delivered per the performance parameters specified by the service. The second 1518 byte Service Frame needs more than the 422 bytes remaining in the bucket and therefore is out-of-profile and is immediately discarded.

CIR and CBS Example 2

A 1518 byte Service Frame is sent. 8ms later, another 1518 byte Service Frame is sent. The first Service Frame depletes 1518 bytes of the initial 2KB CBS in the token bucket leaving 422 bytes remaining. Before the second Service Frame arrives, the tokens in the bucket are replenished at a rate of CIR/8 bytes/second = 250KB/sec in this example. During the 8ms before the second Service Frame arrives, 2000 bytes (tokens) have been added to the bucket. Therefore, the second Service Frame's 1518 bytes will not deplete the bucket and hence the second service Frame will also be in-profile (green).

The value of CBS will depend upon the type of applications or traffic that the service is targeting to support. For example, for a service designed to support bursty TCP-based data applications, CBS will be much larger than for a service supporting more constant rate UDP-based applications such as VoIP.

“CIR-conformant Service Frames ... are delivered per the service performance objectives specified for the service.”

EIR and EBS

The Excess Information Rate (EIR) specifies the average rate, greater than or equal to the CIR, up to which Service Frames are admitted into the provider’s network. Note that these Service Frames are not CIR-conformant and are hence delivered without any performance objectives. The EIR is an average rate because all Service Frames are sent at the UNI speed, e.g., 10Mbps, and not at the EIR, e.g., 8Mbps. EBS is the maximum number of bytes allowed for incoming Service Frames to be EIR-conformant. EIR-conformant Service Frames are colored yellow.

Service Frames whose average rate is greater than the EIR are not EIR-conformant and are discarded.

The EIR may be specified to be less than or equal to the UNI speed. A non-zero EIR must always be greater than the CIR.

EIR-conformant Service Frames may be delivered but without any service performance objectives.




Conformance	Color	Service Frame Delivery
CIR Conformant		Service Frames green and delivered per the performance objectives specified in the SLA/SLS.
EIR Conformant		Service Frames are yellow and may be delivered but with no performance assurances.
None		Service Frames are red and dropped.

Table 1: Service Frame Color Summary

Color Mode (CM)

The Color Mode parameter specifies whether the UNI is operating in color-aware or color-blind mode. This is described in further detail in the section titled “Color Blind and Color Aware UNIs”.

Bandwidth Profile Rate Enforcement

From an implementation perspective, the bandwidth profile rates are enforced through an algorithm which is commonly implemented via a token bucket algorithm. The MEF Technical Committee has defined a two rate, three color marker (trTCM) algorithm which can be implemented via two token buckets.

One bucket, referred to as the ‘Committed’ or ‘C’ bucket, is used to determine CIR-conformant, in-profile Service Frames while a second bucket, referred to as the ‘Excess’ or ‘E’ bucket, is used to determine EIR-conformant, excess Service Frames.

Each token bucket consists of a bucket of bytes referred to as ‘tokens’. Initially, each token bucket is full of tokens. As Service Frames enter the provider’s network, the trTCM decrements the number of tokens in the C bucket (green tokens) by the number of bytes received from the service frame. If green tokens still remain, then the Service Frame is CIR-conformant, colored green and allowed into the provider’s network.

If no green tokens remain, then a second E bucket is checked to determine if any E bucket tokens (yellow tokens) remain. If yellow tokens are available, then the Service Frame is colored yellow and allowed into the provider’s network. If no yellow tokens are available, then the Service Frame is declared red and discarded. Refer to Figure 6.

The MEF has defined an additional, optional capability of the trTCM whereby unused green tokens from the C bucket may be added to the E bucket as yellow tokens when checking EIR-conformance. When this capability is enabled, when operating in color-aware mode, more yellow Service Frames are allowed into the service provider’s network. The MEF currently has no quantitative data describing the implications of this. This capability is not expected to be specified in an SLA to the subscriber and may be used by the provider in marketing their service capabilities.

© The Metro Ethernet Forum 2004. Any reproduction of this document, or any portion thereof, shall contain the following statement: "Reproduced with permission of the Metro Ethernet Forum." No user of this document is authorized to modify any of the information contained herein.

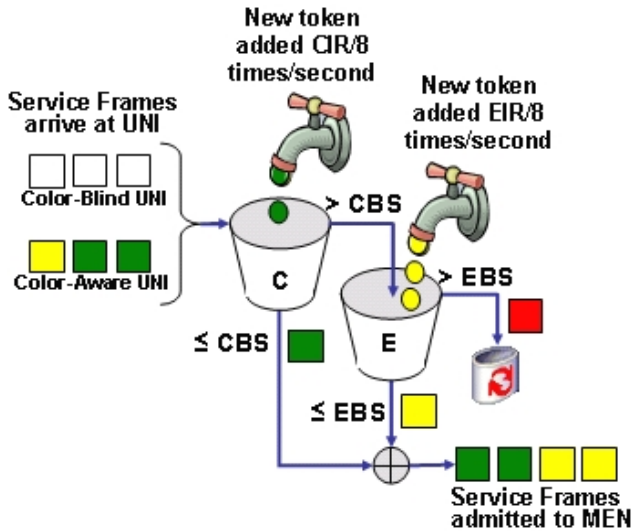


Figure 6: MEF trTCM algorithm

“... bandwidth profile rates are enforced through ... a two rate three color marker (trTCM) algorithm.”

Color Blind and Color Aware UNIs

A color blind UNI is one where the TBA ignores any color indication that the subscriber may have marked in their Service Frames. For example, the color may be marked in the Subscriber Frame via the user priority (802.1p) bits in the IEEE 802.1Q tag [5]. These are referred to in the MEF as the CE-VLAN CoS bits.

A color aware UNI is one where the TBA uses the color indication that the subscriber marked in their Service Frames. For example, enterprise networks typically use the IP Differentiated Services (DiffServ) [6] architecture to provide QoS in their networks. Their IP packets are marked with a DiffServ Code Point (DSCP) [7] indicating the packet color and class of service (CoS). The CE device can map the color and CoS indicated by the DSCP to a CE-VLAN CoS (802.1p) value in the Service Frame to convey to the provider which Service Frames may be discarded under congestion. The TBA then uses this pre-marked color information to make rate enforcement decisions.

For example, suppose a Service Frame is marked yellow by the subscriber’s CE device based on a yellow DSCP value mapped to the CE-VLAN CoS value. The TBA would then only check this Service Frame’s conformance with the E bucket and bypass the C bucket. As with the color blind UNI, if no yellow tokens are available, then the Service Frame is declared red and discarded. If yellow tokens are available, as with the color blind UNI,

then the Service Frame is admitted into the provider’s network.

Bandwidth Profile Service Attribute

The MEF has defined three Bandwidth Profiles service attributes. (Note that the MEF technical committee has initially focused on defining ingress bandwidth profiles). The Ingress Bandwidth Profile per Ingress UNI service attribute applies for all Service Frames at a UNI. The Ingress Bandwidth Profile per EVC applies to all Service Frames sent over a particular EVC. The Ingress Bandwidth Profile per CoS Identifier service attribute applies to all Service Frames within an EVC identified via the customer’s IEEE 802.1p user priority bit values in the IEEE 802.1Q field. In MEF terminology, the 802.1p field is referred to the Customer Edge VLAN CoS or CE-VLAN CoS.

In summary, the Bandwidth Profiles service attributes are as follows. Each bandwidth profile consists of the for traffic parameters <CIR, CBS, EIR, EBS, CM>.

- Ingress Bandwidth Profile per Ingress UNI
- Ingress Bandwidth Profile Per EVC
- Ingress Bandwidth Profile Per CoS

The next sections describe these three bandwidth profiles in more detail including examples of how they can be applied and their affect on service frame rate enforcement.

Ingress Bandwidth Profile Per Ingress UNI

The Ingress Bandwidth Profile Per Ingress UNI provides rate enforcement for all Service Frames ingressing the UNI from subscriber to provider networks. Since the rate enforcement is non-discriminating, some Service Frames may get more bandwidth while others may get less. Refer to Figure 7.

This bandwidth profile is useful when only a single service is supported at the UNI. If the UNI is considered a pipe, one can think of this bandwidth profile as setting the opening diameter of the UNI pipe. Refer to Figure 7. By varying the CIR and EIR parameters, one adjusts the UNI pipe diameter.

To understand the significance of this type of bandwidth profile, let’s look at two extreme examples of CIR and EIR values. If CIR=UNI speed, then all Service Frames are in-profile (green). This configuration would be analogous to a private line service with the bandwidth

equal to the UNI speed. If CIR=0, then all Service Frames are out-of-profile. With this configuration, the EIR value becomes significant. If the $EIR \leq UNI$ speed, then all Service Frames conformant with the EIR are yellow and allowed into the network. All Service Frames that are non-conformant with the EIR (if $EIR < UNI$ speed), are discarded (red). Note that $CIR=0, EIR \leq UNI$ speed is a common definition for a ‘best effort’ service.

“The Ingress Bandwidth Profile Per Ingress UNI provides rate enforcement for all Service Frames ingressing the UNI...”




Figure 7: Ingress Bandwidth profile per Ingress UNI

Ingress Bandwidth Profile Per EVC

The Ingress Bandwidth Profile Per EVC provides rate enforcement for all Service Frames ingressing the UNI that are associated with a particular EVC as illustrated in Figure 8. This bandwidth profile provides more granular bandwidth management for each service instance (EVC) at the UNI.

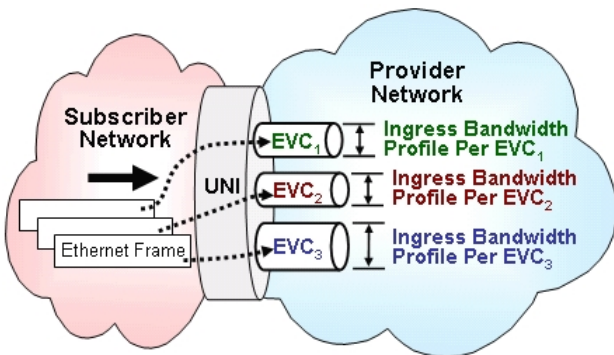


Figure 8: Ingress Bandwidth Profile Per EVC

This bandwidth profile is useful when multiple services are supported at the UNI. If each EVC is considered to be a smaller pipe inside of the larger UNI pipe, one can think of this bandwidth profile as setting the opening diameter

for each EVC pipe. By varying the CIR and EIR parameters, one adjusts the EVC pipe diameter.

To understand the significance of this type of bandwidth profile, let’s look at an example of EVCs used for two different services with different CIR and EIR values. Let’s assume EVC₁ provides an E-Line service (VPN) between two enterprise locations and EVC₂ provides an E-Line service to an ISP for Internet access. In this example, the UNI speed is 100Mbps, EVC₁ has a CIR=10 Mbps and an EIR=100Mbps and EVC₂ has a CIR=40Mbps and an EIR=100Mbps. The E-Line service constructed with EVC₁ would be analogous to a virtual private line service that allows bursting up to the UNI speed and all Service Frames ≤ 10Mbps are admitted into the MEN as in-profile (green) Service Frames. The E-Line service constructed with EVC₂ is used for Internet access and allows bursting up to the UNI speed and all Service Frames ≤ 40Mbps are admitted into the MEN as in-profile (green) Service Frames.

Note that the sum of CIRs for all EVCs must be ≤ the UNI speed.

“The Ingress Bandwidth Profile Per EVC provides rate enforcement for all Service Frames ... associated with a particular EVC”

Ingress Bandwidth Profile Per CoS

The Ingress Bandwidth Profile Per CoS provides rate enforcement for all Service Frames for a given class of service. The class of service is identified via a CoS Identifier determined via the <EVC, CE-VLAN CoS> pair so this bandwidth profile applies to Service Frames over an EVC with a particular CE-VLAN CoS value or set of values. The MEF has defined CE-VLAN CoS as the user priority (802.1p) bits in the IEEE 802.1Q tag. Figure 9 illustrates an example how a separate Ingress Bandwidth Profile Per CoS can be applied to Service Frames sent over EVC₁ marked with CE-VLAN CoS 7, 6 and CE-VLAN CoS 4, 3, 2, 1, 0.

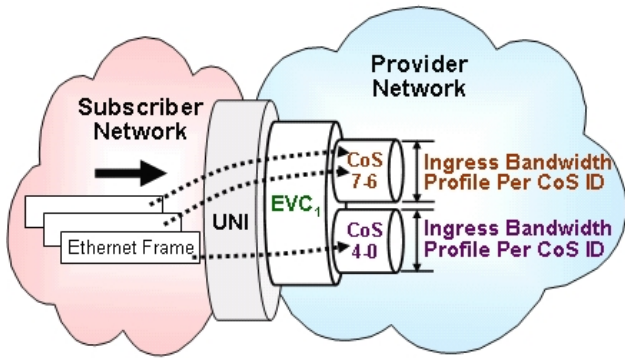


Figure 9: Ingress Bandwidth Profile Per CoS

“The Ingress Bandwidth Profile Per CoS provides rate enforcement for all Service Frames for a given EVC ... identified via the CoS Identifier”

Summary

Bandwidth profiles for Ethernet services allow service providers to sell bandwidth in increments less than the UNI speed, manage their network resources more predictably and provide performance assurances for in-profile traffic. Subscribers can thus purchase bandwidth they need, when they need it more incrementally than TDM-based services.

The MEF has defined three types of bandwidth profiles, each defining different amounts of bandwidth granularity that can be applied to one or more services being offered at the UNI.

Appendix

Terminology

Term	Definition
ATM	Asynchronous Transfer Mode
CBS	Committed Burst Size
CoS	Class of Service
CE	Customer Edge equipment
CBS	Committed Burst Size
CIR	Committed Information Rate
EBS	Excess Burst Size
EIR	Excess Information Rate
EVC	Ethernet Virtual Circuit
LAN	Local Area Network
MEN	Metro Ethernet Network
MPLS	Multi-protocol Label Switching
OAM&P	Operations, Administration, Management and Provisioning
POS	Packet over SONET
SDH	Synchronous Digital Hierarchy

Term	Definition
SONET	Synchronous Optical Network
QoS	Quality of Service
UNI	User-Network Interface
VLAN	Virtual LAN

References and Resources

- [1] MEF 5, “Traffic Management Specification – Phase I”, <http://www.metroethernetforum.org/PDFs/Standards/MEF5.pdf>
- [2] MEF “Metro Ethernet Services – A Technical Overview” <http://www.metroethernetforum.org/metro-ethernet-services.pdf>
- [3] MEF “Metro Ethernet Services Overview” presentation, <http://www.metroethernetforum.org/Metro-Ethernet-Services-Overview.ppt>
- [4] MEF 1 “Ethernet Services Model, Phase 1”, <http://www.metroethernetforum.org/PDFs/Standards/MEF1.pdf>
- [5] “Virtual Bridged Local Area Networks”, <http://standards.ieee.org/reading/ieee/std/lanman/802.1Q-1998.pdf>
- [6] IETF RFC 2475, “An Architecture for Differentiated Services”, <http://www.ietf.org/rfc/rfc2475.txt>
- [7] IETF RFC 2474, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers”, <http://www.ietf.org/rfc/rfc2474.txt>
- [8] Metro Ethernet Forum Membership Application, <http://www.metroethernetforum.org/MEFMemberApplication10-19-01.pdf>

Disclaimer

Some technical details may change in due course (by 75% vote) and this paper will be updated as deemed necessary to reflect such changes. The paper does not necessarily represent the views of the author or his commercial affiliations.

About the Metro Ethernet Forum

The Metro Ethernet Forum (MEF) is a non-profit organization dedicated to accelerating the adoption of optical Ethernet as the technology of choice in metro networks worldwide.

The Forum is comprised of leading service providers, major incumbent local exchange carriers, top network equipment vendors and other prominent networking companies that share an interest in metro Ethernet. As of November 2004, the MEF had over 60 members.

About the Author

Ralph Santituro, Director of Network Architecture at Nortel Networks, provides best practice guidelines for broadband service delivery over metro Ethernet networks. Mr. Santituro has actively participated in the MEF Technical Committee's work on Ethernet service definitions and traffic management since December 2001. Mr. Santituro is co-author of three MEF Ethernet service technical specifications and co-chair of the MEF Technical Marketing Committee.

Mr. Santituro can be contacted at +1 805-527-3024 or rsantito@nortelnetworks.com