# Implementation Agreement

# MEF 23.1

## Carrier Ethernet Class of Service – Phase 2

## January 2012

1. ABSTRACT ........................................................................................................1
2. TERMINOLOGY ........................................................................................................1
3. SCOPE ........................................................................................................8
4. COMPLIANCE LEVELS ........................................................................................................10
5. INTRODUCTION ........................................................................................................10
6. CLASS OF SERVICE MODEL AND OBJECTIVES (NORMATIVE) ........................................................................................................11
6.1 MOTIVATION AND BACKGROUND ON COS MODEL ........................................................................................................11
6.2 KEY TERMS AND HOW THEY RELATE TO EACH OTHER ........................................................................................................13
6.3 COS LABEL ........................................................................................................15
6.4 COS AND COLOR IDENTIFIERS ........................................................................................................15
6.4.1 CoS Identifier and Indication ........................................................................................................15
6.4.2 Color Identifier and Indication ........................................................................................................17
6.5 COS LABEL, COS IDENTIFICATION AND COLOR IDENTIFICATION REQUIREMENTS ........................................................................................................18
6.5.1 Default CoS Label for L2CP ........................................................................................................21
6.6 PERFORMANCE TIER AND ETHERNET NETWORK SECTIONS MODEL ........................................................................................................22
6.6.1 Performance Tier Model ........................................................................................................22
6.6.2 Ethernet Network Section Model ........................................................................................................24
6.7 PERFORMANCE ATTRIBUTES AND METRICS ........................................................................................................26
6.7.1 Frame Delay Performance ........................................................................................................27
6.7.2 Mean Frame Delay Performance ........................................................................................................28
6.7.3 Inter-Frame Delay Variation Performance ........................................................................................................28
6.7.4 Frame Delay Range Performance ........................................................................................................29
6.7.5 Frame Loss Ratio Performance ........................................................................................................30
6.7.6 Availability and Resiliency Performance ........................................................................................................30
6.8 BANDWIDTH PROFILE AND COLOR ........................................................................................................31
6.8.1 Bandwidth Profile Compliance ........................................................................................................32
6.8.2 Egress Bandwidth Profile Considerations ........................................................................................................33
6.9 EVC AND SERVICE TYPE APPLICABILITY ........................................................................................................33
6.10 OVC AND SERVICE TYPE APPLICABILITY ........................................................................................................34
6.11 COS LABEL MODEL ........................................................................................................34
6.11.1 Three CoS Label Model ........................................................................................................35
6.11.2 Performance Parameters ........................................................................................................38
6.11.3 CoS Performance Objectives Per Performance Tier ........................................................................................................41
6.11.4 PCP and DSCP Mapping ........................................................................................................46
6.11.4.1 UNI Mapping ........................................................................................................ 46
6.11.4.2 ENNI Mapping ........................................................................................................ 46

6.11.5 L2CP CoS Mapping....46
6.11.6 Requirement Sets....47

**7. REFERENCES....47**

**8. APPENDICES (INFORMATIVE)....49**

8.1 Potential Work Areas for Later Phases of MEF CoS IA....49
8.1.1 CoS Subset Mapping....49
8.1.2 Other Potential Work Areas....49
8.2 Performance Tier Model Derivation....49
8.2.1 Low Speed Link Considerations....50
8.3 Ethernet Network Section Model – Composing UNI-UNI Values....50
8.3.1 Mean delay....52
8.3.2 Loss ratio....52
8.3.3 Relationship for delay and delay range....53
8.3.4 Ethernet Network Section Recommendations....54
8.4 Key Applications to Derive Performance Requirements....55
8.4.1 Application-specific Performance Objectives....56
8.4.2 Derivation of CPOs from Application Performance Requirements....76
8.4.2.1 Mapping Applications to CoS Labels and Performance Tiers.... 77
8.4.2.2 Constraints on CPO Values.... 77
8.4.2.3 The CoS Performance Objective Compliance Tool.... 79
8.4.2.3.1 Performance Tier worksheets.... 79
8.4.2.3.2 CPO Summary worksheet.... 84
8.4.2.3.3 Application Mapping Summary Worksheet.... 85
8.5 Example PCP and DSCP Mapping at UNI for Multi-CoS EVCs Supporting Only Standard MEF Classes of Service....86
8.5.1 Example PCP Mappings....87
8.5.2 Example DSCP Mappings....88
8.6 Other Relevant Standards and Industry models....88
8.7 Burst Size and Shaper Considerations for ENNI....89
8.7.1 Burst Size and Burst Alignment....89
8.7.1.1 Representative Values for ΔCBS.... 91
8.7.1.2 Upper Bounds for Burst Sizes.... 92
8.7.2 Shaping Considerations for Burst Alignment....93

# List of Figures

Figure 1 – CoS IA Scope and Applicability.... 9
Figure 2 – CoS IA Motivation Example – ENNI Mapping.... 12
Figure 3 – Example Performance Tier for a Single MEN EVC.... 25
Figure 4 – Example Performance Tiers for a Multiple MEN EVC and OVCs.... 25

Figure 5 – CPO Summary worksheet ........ 85
Figure 6 – Application Mapping summary worksheet ........ 86
Figure 7 – Burst Alignment Example with Policing Points for Traffic Traversing the ENNI ..... 90
Figure 8 – Shaper Buffers and Transmission Buffer Relationship ........ 93
Figure 9 – Periodic Algorithm ........ 95
Figure 10 – Revised New Frame Algorithm ........ 96

# List of Tables

Table 1: Terminology and Definitions Table ........ 7
Table 2: CoS Labels and CoS ID Types in CoS IA ........ 36
Table 3: Color ID Values when CoS ID is Only EVC or OVC EP ........ 37
Table 4: CoS Identifiers and Color Identifiers ........ 38
Table 5: CoS Label H, M and L Parameter Values ........ 40
Table 6: PT1 (Metro PT) CPOs ........ 42
Table 9: PT4 (Global PT) CPOs ........ 45
Table 10: L2CP CoS Mapping ........ 46
Table 11: MEF – ITU Variable Mapping ........ 52
Table 12: Application list ........ 56
Table 13: VoIP Parameters ........ 57
Table 14: Interactive Video Parameters ........ 59
Table 15: Standard Definition Video Parameters ........ 60
Table 16: High Definition Video Parameters ........ 61
Table 17: Internet Streaming Audio/Video Parameters ........ 62
Table 18: Interactive Transaction Data Parameters ........ 63
Table 19: Interactive Gaming Parameters ........ 64
Table 20: Best Effort Parameters ........ 65
Table 21: Circuit Emulation Parameters ........ 66
Table 22: Point of Sale Transaction Parameters ........ 66
Table 23: Synchronous Replication Parameters ........ 67
Table 24: Asynchronous Replication Parameters ........ 68
Table 25: Network Attached Storage Parameters ........ 68
Table 26: Text and Graphics Terminal Parameters ........ 69
Table 27: T.38 Fax Parameters ........ 69
Table 28: Database Parameters – Hot Standby ........ 70
Table 29: Database Parameters – WAN Replication ........ 70
Table 30: Database Parameters – Client/Server ........ 71
Table 31: Financial Trading Parameters ........ 71
Table 32: CCTV Parameters ........ 72
Table 33: Telepresence Parameters ........ 73
Table 34: Mobile Backhaul Proposed CPOs ........ 74
Table 35: Summarized CPOs ........ 76

Table 36: Explicit Application Mapping for Derivation of CPOs ............................................. 77
Table 37: CPO Derivation Constraints ........................................................................................ 79
Table 38: PT1-4 CPO Derivation and Evaluation Spreadsheets ................................................ 84
Table 39: Example PCP Mapping for Multi-CoS Label EVC Supporting Only Standard CoS Labels at UNI – “Router-Application-Friendly” mapping.................................................. 87
Table 40: Example PCP Mapping for Multi-CoS Label EVC Supporting Only Standard CoS Labels at UNI – “Bridging-Application-Friendly” mapping ............................................... 88
Table 41: Example DSCP Mapping for Multi-CoS EVC Supporting Only Standard Classes of Service at UNI ....................................................................................................................... 88
Table 42: PCP Decoding (Adapted from [5]) .............................................................................. 89
Table 43: Representative Values for CBS Ranges ...................................................................... 92

# 1. Abstract

This Implementation Agreement (IA) specifies a set of Class of Service Names called CoS Labels that can be used by Operators, Service Providers and their Subscribers to indicate the performance expectations to be associated with a given set of frames that comprise a CoS Frame Set. This CoS IA includes standards for CoS and Color identification as well as performance objectives and supporting requirements. The CoS Labels are envisioned as a subset of all of the Class of Service Names an Operator may provide. The MEF CoS IA facilitates: Ethernet service interoperability and consistency between Operators, use of a common CoS Label set for Subscribers to utilize and use of performance objectives that support key applications. The terms CoS Label, CoS Name, CoS Frame Set and others are defined in Section 2 of this IA.

# 2. Terminology

This section defines the terms used in this document. In many cases, the normative definitions to terms are found in other documents. In these cases, the third column is used to provide the reference that is controlling. Note that a term may be defined differently in a document other than the controlling document. In this case, the definition from the controlling document is the one used in this document.

| Term | Definition | Reference |
|---|---|---|
| **Bandwidth Profile per CoS ID** | A Bandwidth Profile applied on a per-Class of Service Identifier basis. | [2] for Service Frames and [13] for ENNI Frames |
| **Bandwidth Profile per EVC** | A Bandwidth Profile applied on a per-EVC basis. | [2] |
| **Bandwidth Profile per OVC End Point** | A Bandwidth Profile applied on a per-OVC End Point basis. | [13] |
| **Bandwidth Profile per UNI** | A Bandwidth Profile applied on a per-UNI basis. | [2] |
| **Bandwidth Profile per VUNI** | A Bandwidth Profile applied on a per-VUNI basis. | [14] |
| **BWP** | Bandwidth Profile | [2] |
| **CBS** | Committed Burst Size | [2] |
| **CE** | Customer Edge | [2] |
| **CE-VLAN CoS** | Customer Edge VLAN CoS. Also C-Tag PCP. | [2] |
| **CE-VLAN Tag** | Customer Edge VLAN Tag. Also C-Tag. | [2] |

| Term | Definition | Reference |
|---|---|---|
| **CF** | Coupling Flag | [2] |
| **CIR** | Committed Information Rate | [2] |
| **Class of Service Identifier for Service Frames** | The mechanism and/or values of the parameters in the mechanism to be used to identify the CoS Name that applies to the frame at a given UNI | Derived from [2] |
| **Class of Service Identifier for EFO** | The mechanism and/or values of the parameters in the mechanism to be used to identify the CoS Name that applies to the frame at a given ENNI that maps to an OVC End Point. | Derived from [13] |
| **Class of Service Identifier for EFV** | The mechanism and/or values of the parameters in the mechanism to be used to identify the CoS Name that applies to the frame at a given ENNI that maps to a VUNI End Point. | Derived from [14] |
| **Class of Service Frame Set** | A set of Service or ENNI Frames that have a commitment from the Operator or Service Provider subject to a particular set of performance objectives. | This document |
| **Class of Service Label** | A CoS Name that is standardized in this document. Each CoS Label identifies four Performance Tiers where each Performance Tier contains a set of performance objectives and associated parameters. | This document |
| **Class of Service Name** | A designation given to one or more sets of performance objectives and associated parameters by the Service Provider or Operator. | This document |
| **Class of Service Performance Objective** | An objective for a given performance metric. | This document |
| **CM** | Color Mode | [2] |
| **Color Mode** | CM is a Bandwidth Profile parameter. The Color Mode parameter indicates whether the color-aware or color-blind property is employed by the Bandwidth Profile. It takes a value of “color-blind” or “color-aware” only. | [2] |
| **Color-aware** | A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each Service or ENNI Frame, indicated by the Color Identifier, is taken into account when determining the level of compliance for each Service Frame. | [2], [13] and this IA |
| **Color-blind** | A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each Service or ENNI Frame, if present, is ignored when determining the level of compliance for each Service or ENNI Frame. | Adapted from [2] and [13] |

| Term | Definition | Reference |
|---|---|---|
| **Color ID** | Color Identifier | This document |
| **Color Identifier for Service Frame** | The mechanism and/or values of the parameters in the mechanism used to identify the Color that applies to the frame at a given UNI. | This document |
| **Color Identifier for EFO** | The mechanism and/or values of the parameters in the mechanism used to identify the Color that applies to the frame at a given ENNI that maps to an OVC End Point. | This document |
| **Color Identifier for EFV** | The mechanism and/or values of the parameters in the mechanism used to identify the Color that applies to the frame at a given ENNI that maps to a VUNI End Point. | This document |
| **Committed Burst Size** | CBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of Service or ENNI Frames sent at the EI speed to remain CIR-conformant. | Adapted from [2] and [13] |
| **Committed Information Rate** | CIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of Service or ENNI Frames up to which the network delivers Service or ENNI Frames and meets the performance objectives defined by the CoS Service Attribute. | Adapted from [2] and [13] |
| **CoS** | Class of Service or Classes of Service | [2] |
| **CoS ID** | Class of Service Identifier | [2] |
| **CoS FS** | Class of Service Frame Set | This document |
| **Coupling Flag** | CF is a Bandwidth Profile parameter. The Coupling Flag allows the choice between two modes of operation of the Bandwidth Profile algorithm. It takes a value of 0 or 1 only. | Adapted from [2] |
| **CPO** | CoS Performance Objective | This document |
| **Customer Edge** | Equipment on the Subscriber side of the UNI. | [2] |
| **Customer Edge VLAN CoS** | The Priority Code Point bits in the IEEE 802.1Q Customer VLAN Tag in a Service Frame that is either tagged or priority tagged. Also C-Tag PCP. | [2] |
| **Customer Edge VLAN Tag** | The IEEE 802.1Q Customer VLAN Tag in a tagged Service Frame. Also C-Tag. | [2] |
| **C-Tag** | Subscriber VLAN Tag | [5] |
| **DEI** | Drop Eligible Indicator | [5] |
| **DSCP** | Differentiated Services Code Point | [2] |
| **EBS** | Excess Burst Size | [2] |

| Term | Definition | Reference |
|---|---|---|
| **EFO** | ENNI Frame that maps to OVC End Point | This document |
| **EFV** | ENNI Frame that maps to a VUNI End Point | This document |
| **Egress Bandwidth Profile** | A service attribute that specifies the length and arrival time characteristics of egress Service or ENNI Frames at the egress UNI or ENNI. | Adapted from [2] and [13] |
| **EI** | External Interface | [6] |
| **EIR** | Excess Information Rate | [2] |
| **E-LAN Service** | An Ethernet service type that is based on a Multipoint-to-Multipoint EVC. | [1] |
| **E-Line Service** | An Ethernet service type that is based on a Point-to-Point EVC. | [1] |
| **ENNI** | External Network Network Interface. An interface used to interconnect two MEN Operators | [6] |
| **ENNI Frame** | The first bit of the Destination Address to the last bit of the Frame Check Sequence of the Ethernet Frame transmitted across the ENNI | [13] |
| **ENS** | Ethernet Network Section | This document and [15] |
| **EPL** | Ethernet Private Line | [1] |
| **E-Tree Service** | An Ethernet service type that is based on a Rooted-Multipoint EVC. | [1] |
| **Ethernet Virtual Connection** | An association of two or more UNIs that limits the exchange of Service Frames to UNIs in the Ethernet Virtual Connection. | [2] |
| **Ethernet Network Section** | A set of one or more MENs, each under a single or collaborative jurisdictional responsibility, for the purpose of managing CPOs. | This document and [15] |
| **EVC** | Ethernet Virtual Connection | [2] |
| **EVPL** | Ethernet Virtual Private Line | [1] |
| **EVP-LAN** | Ethernet Virtual Private LAN | [1] |
| **Excess Burst Size** | EBS is a Bandwidth Profile parameter.  It limits the maximum number of bytes available for a burst of Service or ENNI Frames sent at the EI speed to remain EIR-conformant. | Adapted from [2] and [13] |

| Term | Definition | Reference |
|---|---|---|
| **Excess Information Rate** | EIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of Service or ENNI Frames up to which the network may deliver Service or Network Frames but without any performance objectives. | Adapted from [2] and [13] |
| **External Interface** | Either a UNI or an ENNI | [13] |
| **FD** | Frame Delay | [2] |
| **FDR** | Frame Delay Range | |
| **FLR** | Frame Loss Ratio | [2] |
| **Frame** | Short for Ethernet Frame | [2] |
| **Frame Delay** | The time required to transmit a Service or ENNI Frame from ingress EI to egress EI. | Adapted from [2] and [13] |
| **Frame Delay Performance** | A characterization of the delays experienced by different Service or ENNI Frames belonging to the same CoS Frame Set. | Adapted from [2] and [13] |
| **Frame Delay Range** | The difference between the observed percentile of delay at a target percentile and the observed minimum delay for the set of frames in time interval T. | [2] and this document |
| **Frame Delay Range Performance** | A characterization, based on Frame Delay Range, of the extent of delay variability experienced by different Service or ENNI Frames belonging to the same CoS Frame Set. | Adapted from [2] and [13] |
| **Frame Loss Ratio Performance** | Frame Loss Ratio is a characterization of the number of lost Service Frames or ENNI Frames between the ingress External Interface (EI) and the egress External Interface (EI). Frame Loss Ratio is expressed as a percentage. | Adapted from [2] and [13] |
| **IA** | Implementation Agreement | |
| **IFDV** | Inter Frame Delay Variation | [2] |
| **Ingress Bandwidth Profile** | A characterization of ingress Service or ENNI Frame arrival times and lengths at the ingress UNI or ENNI and a specification of disposition of each Service or ENNI Frame based on its level of compliance with the characterization. | Adapted from [2] and [13] |
| **Ingress Service Frame** | A Service Frame sent from the CE into the Service Provider network. | [2] |
| **Inter-Frame Delay Variation** | The difference in delay of two Service or ENNI Frames of the same CoS Frame Set. | Adapted from [2] and [13] |
| **Inter-Frame Delay Variation Performance** | A characterization, based on Inter-Frame Delay Variation, of the variation in the delays experienced by different Service or ENNI Frames belonging to the same CoS Frame Set. | Adapted from [2] and [13] |

| Term | Definition | Reference |
|---|---|---|
| **L2CP** | Layer 2 Control Protocol | [2] |
| **Layer 2 Control Protocol Service Frame** | A Service Frame that is used for Layer 2 Control, e.g., Spanning Tree Protocol. | [2] |
| **Layer 2 Control Protocol Tunneling** | The process by which a frame carrying a Layer 2 Control Protocol Service data unit is passed through the Service Provider or Operator network without being processed and is delivered to the proper EI(s). | Adapted from [2] and [20] |
| **MFD** | Mean Frame Delay | Adapted from [2] |
| **Mean Frame Delay** | The arithmetic mean, or average of delays experienced by Service or ENNI Frames belonging to the same CoS Frame Set. | Adapted from [2] and [13] |
| **MEN** | Metro Ethernet Network | [6] |
| **Metro Ethernet Network** | The Operator's or Service Provider's network providing Ethernet services. Synonymous with Carrier Ethernet Network (CEN) | [6] |
| **Multipoint-to-Multipoint EVC** | An EVC with two or more UNIs. A Multipoint-to-Multipoint EVC with two UNIs is different from a Point-to-Point EVC because one or more additional UNIs can be added to it. | [2] |
| **Operator** | Also Network Operator. The Administrative Entity of a MEN | Derived from [6] |
| **N/A** | Not Applicable | |
| **N/S** | Not Specified | |
| **Operator Virtual Connection** | An association of OVC End Points | [13] |
| **OVC** | Operator Virtual Connection | [13] |
| **OVC End Point** | An association of an OVC with a specific External Interface i.e., UNI, ENNI | [13] |
| **OVC EP** | OVC End Point | This document |
| $P_d$ | Frame Delay Performance percentile | Adapted from [13] |
| $P_v$ | Inter-Frame Delay Variation Performance percentile | Adapted from [13] |
| **PCP** | Priority Code Point | [5] |
| **Performance Tier** | A MEF CoS Performance Objectives (CPO) set | This document |

| Term | Definition | Reference |
|---|---|---|
| **Point-to-Point EVC** | An EVC with exactly 2 UNIs. | [2] |
| $P_r$ | A specific percentile of the Frame Delay Performance used in Frame Delay Range , where $P_r > 0$ | Adapted from [13] |
| **PT** | Performance Tier | This document |
| **Rooted-Multipoint EVC** | A multipoint EVC in which each UNI is designated as either a Root or a Leaf. Ingress Service Frames at a Root UNI can be delivered to one or more of any of the other UNIs in the EVC. Ingress Service Frames at a Leaf UNI can only be delivered to one or more Root UNIs in the EVC. | [2] |
| **S** | Subset of the ordered UNI pairs or a subset of the OVC End Point pairs | [2], [13] |
| **Service Frame** | An Ethernet frame transmitted across the UNI toward the Service Provider or an Ethernet frame transmitted across the UNI toward the Subscriber. | [2] |
| **Service Level Agreement** | The contract between the Subscriber and Service Provider specifying the agreed to service level commitments and related business agreements. | [2] |
| **Service Level Specification** | The technical specification of the service level being offered by either the Service Provider to the Subscriber in the case of an EVC service or by an Operator to a Service Provider in the case of an OVC. | Adapted from [2] and [13] |
| **Service Provider** | The organization providing Ethernet Service(s). | [2] |
| **S-Tag** | Service VLAN Tag | [5] |
| **SLA** | Service Level Agreement | [2] |
| **SLS** | Service Level Specification | [2] |
| **Subscriber** | The organization purchasing and/or using Ethernet Services. | [2] |
| **T** | A time interval that serves as a parameter for an SLS. | [2] |
| **UNI** | User Network Interface | [2] |
| **User Network Interface** | The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber. | [2] |
| **VLAN** | Virtual LAN | [3] |
| $\Delta t$ | A time interval much smaller than $T$ | [2] |

**Table 1: Terminology and Definitions Table**

# 3. Scope

Phase 1 of the CoS IA defined a set of 3 CoS Names called CoS Labels for UNI-to-UNI (i.e., EVC) services for both single MEN and multiple interconnected MENs administered by different Operators. In Phase 2 values for CoS Performance Objectives (CPOs) grouped in Performance Tier sets, Performance Parameters and 2 Performance metrics (MFD and FDR) are added.

The CoS Identification and Color Identification specifications in this IA are applicable at External Interfaces (EIs), which can be either UNI or ENNI, and the CPOs in this IA are applicable to CoS Frame Sets between the EIs. Phase 2 includes UNI-to-ENNI and ENNI-to-ENNI (i.e., OVC) services. Phase 2 also addresses Ethernet Network Sections associated with typical Operator domains that interconnect at ENNIs (e.g., concatenation of CPOs for OVCs to derive CPOs for EVCs).

A CoS model is defined that can be applied to the CoS Frame Sets of an EVC or OVC. The internal mechanisms for implementing the CoS IA are out of scope and left up to implementation.

Specification of all possible or likely CoS Names is also out of scope. This IA specifies a small set of CoS Labels (i.e., a Three CoS Label Model) that provides support for key applications. In addition, this IA specifies the usage of a subset of PCP and DSCP values (components of CoS Identifiers) with the defined CoS Labels while leaving some PCP and DSCP values available for Operator use. The Operator may use these additional values to map to MEF CoS Labels, internal CoS Names or additional Subscriber CoS Names, e.g., an Operator offers 2 CoS Names, in addition to the 3 MEF CoS Labels, at a UNI where the additional CoS Names are Operator specific. Operator specific CoS Names are out of scope in this IA. An Operator can implement any number (e.g., 3, 2, or 1) of the MEF CoS Labels (i.e., CoS ID values) for CoS Frame Set(s) of an EVC or OVC at an associated EI. An Operator can support the CPOs between EIs. Future Phases may specify additional MEF CoS Labels.

This IA includes the CoS Identifiers (CoS IDs) defined in [2], [13] and [14].

Phase 1 specified the CoS Label model structure including: 3 specified CoS Labels, Performance objectives placeholders, applicability of Bandwidth Profile options, and associated CoS ID parameters. In Phase 2 place holders for Frame Delay, Inter-Frame Delay Variation and Frame Loss Ratio CPOs are replaced with values and Mean Frame Delay and Frame Delay Range CPOs are added. This phase includes placeholders for Availability, High Loss Intervals, and Consecutive High Loss Intervals CPOs in preparation for later phases. Phase 2 elaborates on the relationship between CoS and Bandwidth Profile. Phase 2 also adds Performance metric associated Performance Parameters (e.g., Percentile ($P$), Time interval (T)) and specified values to allow determination of CPOs. A tunnelled L2CP specific default CoS Label is in scope.

Network control/signalling (beyond L2CP), operations and security aspects are out of scope. TRANS layer (defined in [17]) technology capabilities, used by Operators to indicate CoS Names and Color inside a network, are out of scope for this IA.

Where possible this IA will rely on CoS and performance related service attributes already defined in other MEF specifications. To further define CoS, this IA identifies, and where necessary constrains or extends, current MEF specifications. The IA also builds upon previous work in IEEE, ITU and IETF for consistency, fast development and facilitation of end-to-end CoS. This previous standards work includes CoS definitions for the IP layer, thus facilitating synergies between Ethernet and IP services and networks.

Figure 1 represents scope and applicability of the CoS IA to both UNI and ENNI, to Multipoint-to-Multipoint, Rooted-Multipoint and Point-to-Point EVCs and OVCs, and to both single and multiple MENs. While the Three CoS Label Model in Phase 2 is applicable to Multipoint-to-Multipoint and Rooted-Multipoint, their CPOs may be specified in a later phase.

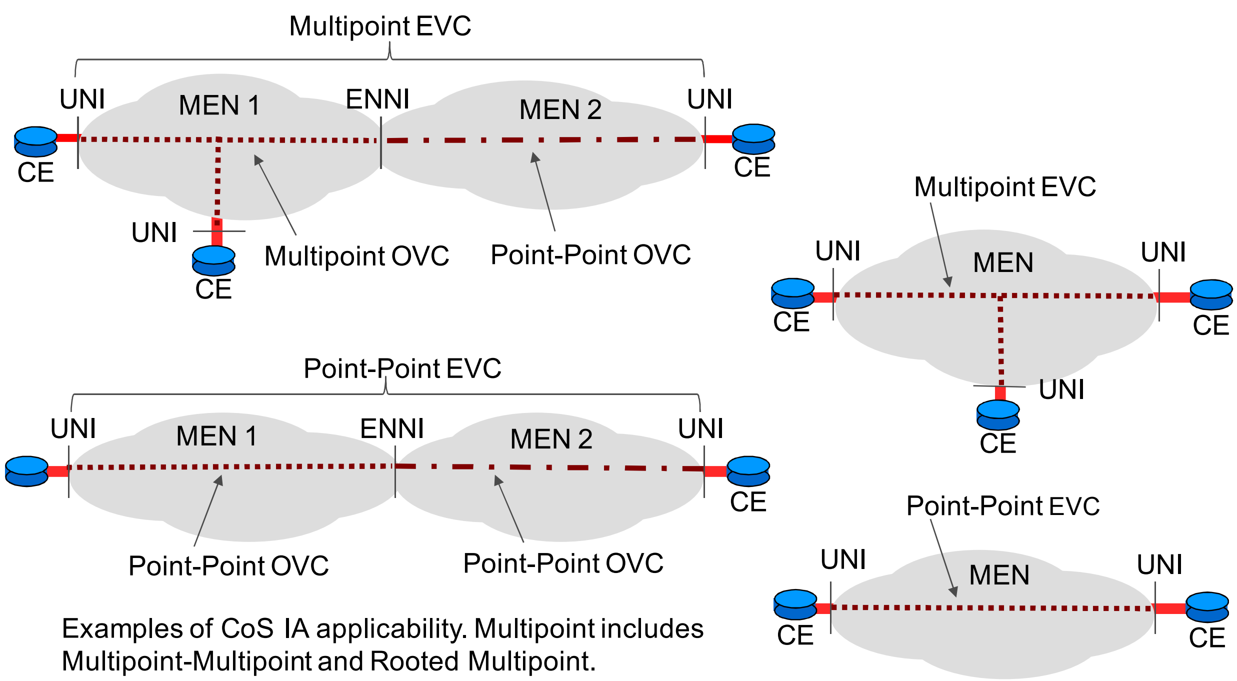


**Figure 1 – CoS IA Scope and Applicability**

With respect to the set of interfaces that are described as MEN External Interfaces in [6], the CoS IA will use the term External Interface (EI) to only include the UNI and ENNI for instances where UNI and ENNI share common characteristics (i.e., the associated statement applies to both UNI and ENNI).

The normative content of this IA is in Section 6. This section provides motivation and background followed by specification of how CoS Identifiers and Color Identifiers are used. This includes the introduction of the terms CoS Label to represent CoS IA specified "classes" (i.e., MEF specified CoS Names) and Color Identification (Color ID) at the UNI and ENNI. Next are description and requirements for Frame Delay, Mean Frame Delay, Inter-Frame Delay Variation, Frame Delay Range and Frame Loss Ratio Performance Attribute objectives that are added in Phase 2. Next a short section provides the necessary Bandwidth Profile, including burst alignment, requirements in order to specify a CoS Label Model and associated CPOs. Additional Bandwidth Profile specification work might be required in future phases and/or other MEF specifications. After a description of CoS Label Model applicability to EVC and OVC, the CoS Label Model and associated Tables are specified. The tables provide the "classes" (i.e., CoS Labels), PCP and DSCP components of CoS Identification (i.e., "code points"), CoS Performance Objective values (CPOs) for each Performance Tier (PT) and Performance Attribute Parameter values. The tables are followed by a section on EI and L2CP mapping.

Finally there are several Appendices that provide background information, concatenation of Ethernet Network Sections, derivation of CPOs, use cases and preliminary direction for future phase work.

# 4. Compliance Levels

The requirements that apply to the MEF CoS are specified in the following sections. Items that are **REQUIRED** (contain the words **MUST** or **MUST NOT**) will be labeled as **[Rx]**. Items that are **RECOMMENDED** (contain the words **SHOULD** or **SHOULD NOT**) will be labeled as **[Dx]**. Items that are **OPTIONAL** (contain the words **MAY** or **OPTIONAL**) will be labeled as **[Ox].**

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in RFC 2119 [4]. All key words use upper case, bold text to distinguish them from other uses of the words. Any use of these key words (e.g., may and optional) without **[Rx], [Dx]** or **[Ox]** is not normative**.**

# 5. Introduction

Ethernet has its origins in providing local network connectivity and was not originally used to provide public services to Subscribers through Operators and Service Providers. With the introduction of Metro and Carrier Ethernet services, Service Providers and Operators started using this Ethernet "connectivity" technology to provide Ethernet "services". Various MEF specifications have added to IEEE 802 series standards in order to create a framework to define Ethernet services. This IA is motivated by the need to introduce and define specific "classes" or CoS Names called CoS Labels that will deliver a commitment for a particular level of performance for a set of Service or ENNI Frames (e.g., those belonging to a particular CoS Frame Set) from the Service Provider or Operator. This is to further develop Carrier Ethernet services that are interoperable and predictably support Subscriber applications. For example, Operators and Service Providers that connect MENs will be able to do so with a set of commonly understood CoS Labels, CoS IDs and CPOs in addition to any bilateral CoS Names they want to support.

This CoS IA normative language is primarily applicable to Subscribers, Service Providers and Operators who desire CoS Name interoperability across EIs. The requirements are developed based on the needs of Subscribers and their applications. Compliance with the CoS Labels in this IA does not limit an Operator from providing additional CoS Names using CoS Identifier values (e.g., PCP) that are left unused in this IA. Examples of additional CoS Names could include Operator defined CoS Names in addition to the specific MEF CoS Labels defined in this IA. Note that the CoS Performance Objective (CPO) and Parameter values are specified in this IA as maximums or minimums and thus do not limit Operators from providing conformant values that are less than the maximums or greater than the minimums. These other values could be described as more stringent, i.e., having more rigor or severity with respect to the standard or requirement value.

# 6. Class of Service Model and Objectives (Normative)

## 6.1 Motivation and background on CoS Model

Figure 2 illustrates the need for a standard CoS Label model for mapping at an ENNI which is one key motivation for this IA. The problem addressed is that the Operators of MEN 1 and MEN 2 may have different CoS Names and different methods and values to indicate the CoS Names. The figure illustrates how the use of the CoS IA can provide a common set of CoS Labels that the Operators can map frames into, to facilitate interworking. For example for a frame going from MEN 1 to MEN 2 whereby CoS Name Heart maps to MEF CoS Label M which then maps to CoS Name Paper in MEN 2. Similarly, for a frame going from MEN 1 to MEN 2 whereby CoS Name Square also maps to MEF CoS Label M and thus maps to CoS Name Paper in MEN 2. Finally, for a frame going from MEN 2 to MEN 1 whereby CoS Name Paper maps to MEF CoS Label M and thus maps to CoS Name Square in MEN 1. In this example and this IA, the transmitting MEN is responsible for mapping their internal CoS Names and Color to the MEF CoS Label and Color for the frame prior to transmitting across the ENNI, as per mutual agreement with the receiving MEN, so the receiving MEN can ensure compliance to the desired objectives within that MEN.

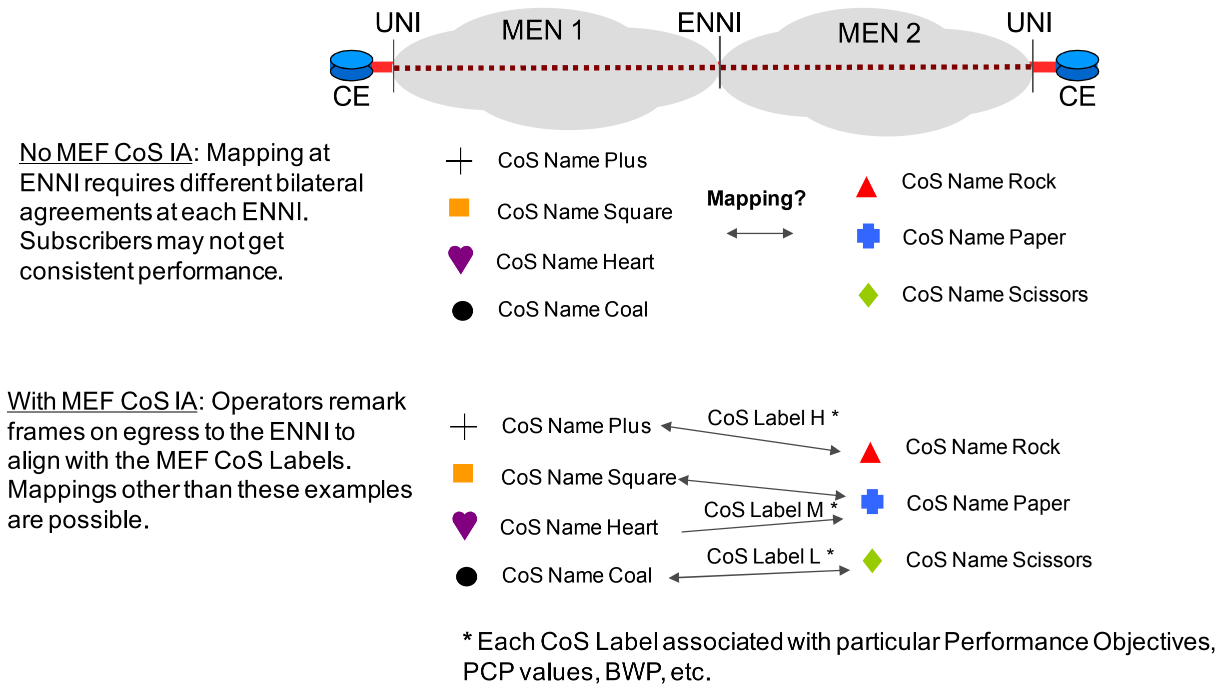


**Figure 2 – CoS IA Motivation Example – ENNI Mapping**

Note that in the figure above the 3 CoS Names used by the Operator (Rock, Paper, Scissors) may align with the CoS IA Three CoS Label Model. A case could be constructed where neither MEN complies with the CoS IA Three CoS Label Model at the UNIs in their MEN, but both map to the CoS IA Model at the ENNI. A Three CoS Label Model is specified in order to satisfy the competing needs of a diversity of applications, finding common needs among Operators, limited CoS Identifier and Color Identifier field value space (e.g., 8 possible PCP values) and ensuring sufficiently simple interoperability. CoS IA Phase 2 allows any combination of subsets of the 3 CoS Labels specified.

In addition, interconnection at the ENNI faces the challenge of providing UNI-to-UNI CoS with multiple Operators. Each Operator will provide a subset of the OVCs that make up the EVC. In addition to the need for CPOs associated with the UNI-to-UNI EVC, interworking and performance will be facilitated if each Operator has CPOs for their OVCs that are consistent with the EVC CPOs.

The CPOs for Phase 2 are specified as values associated with a CoS Label for a particular field of use or applicability called a Performance Tier. Note that "not-specified" (N/S) is a possible CPO value in this CoS IA. See Section 6.6.1 for more information.

## 6.2 Key Terms and How They Relate to Each Other

There are several key terms that are used throughout this document. This section defines these terms and describes how they relate to each other.

The key terms and definitions discussed in this section:

- **Class of Service Frame Set (CoS FS):** A set of Service Frames or ENNI Frames that have a commitment from the Operator or Service Provider subject to a particular set of performance objectives.
- **Class of Service Name (CoS Name):** A designation given to one or more sets of performance objectives and associated parameters by the Service Provider or Operator.
- **Class of Service Label (CoS Label):** A CoS Name that is standardized in this document. Each CoS Label identifies four Performance Tiers (see Section 6.4) where each Performance Tier contains a set of performance objectives and associated parameters.[1] CoS Label is further described in section 6.3.
- **Class of Service Identifier (CoS ID):** The mechanism (e.g., "EVC and PCP") and/or values of the parameters in the mechanism (e.g., PCP value of 3) to be used to identify the CoS Name that applies to the frame at a given EI. CoS ID is further described in sections 2 and 6.4.1.
- **Color Identifier (Color ID):** The mechanism (e.g., PCP, DEI) and/or values of the parameters in the mechanism (e.g., PCP value of 3) used to identify the Color that applies to the frame at a given EI. Color ID is further described in sections 2 and 6.4.2 and 6.5.

CoS IA Phase 1 [19] uses a different and more specific interpretation of CoS related terms than MEF 10.2 [2]. Section 6.9 in [2] defines CoS ID as identifying a CoS instance, but the examples in Section 6.9 in [2] make it clear that the intent is to identify a CoS Name. This second phase of CoS IA is defining CoS in still more detail. Usage of CoS instance in [13] has these same issues found in [2]. Therefore, this document defines and uses the terms CoS FS, CoS Name, CoS Label, CoS ID and Color ID as indicated above.

A Service Provider or Operator can use many CoS Names, each with several different sets of performance objectives and associated parameters. A key goal of this document is to standardize three CoS Names and the values for the sets of performance objectives and associated parameters. These three CoS Names are called CoS Labels and are designated H, M, and L. The sets of performance objectives and associated parameters for each label are called Performance Tiers. Knowing the CoS Label that applies to a given frame is not sufficient to know the performance objectives and associated parameters that apply to the frame. What is required is to know both the CoS Label and the Performance Tier that applies to the frame.

[1] In this document, the parameters have the same value across all Performance Tiers.

Since a CoS FS is a set of frames with performance objective values and associated parameter values, information in addition to the CoS Label and Performance Tier is required. In particular, the subset of ordered UNI pairs (S) is required when dealing with an EVC and the subset of ordered OVC End Point pairs (S) is required when dealing with an OVC. MEF 6.1 [1] requires EVC performance to be specified per CoS ID and for S. This IA requires specifying EVC performance per CoS ID for S and a PT. This is also required for OVC based services in [13] and [14].

Another goal of this document is to standardize how to decode a Service Frame or an ENNI Frame to identify which CoS Label applies to the frame. This is done in Sections 6.5 and 6.11.1.

As an example, consider a Multipoint-to-Multipoint EVC with UNIs in New York and Chicago. Three subsets of ordered UNI pairs could be specified:

- $S_1$: All ordered UNI pairs where both UNIs in each ordered pair are in New York,
- $S_2$: All ordered UNI pairs where both UNIs in each ordered pair are in Chicago, and
- $S_3$: All ordered UNI pairs where one UNI is in New York and the other UNI is in Chicago.

If two CoS IDs (e.g., $CoSID_1$ indicating CoS Label H and $CoSID_2$ indicating M) are supported on this EVC, the following six Class of Service FSs could be established for this EVC:

1. $\langle S_1, CoSID_1, PT_1 \rangle$: All Service Frames from UNI $i$ to UNI $j$ where $\langle i, j \rangle \in S_1$ whose CoS ID identifies CoS Label H.
2. $\langle S_2, CoSID_1, PT_1 \rangle$: All Service Frames from UNI $i$ to UNI $j$ where $\langle i, j \rangle \in S_2$ whose CoS ID identifies CoS Label H.
3. $\langle S_3, CoSID_1, PT_3 \rangle$: All Service Frames from UNI $i$ to UNI $j$ where $\langle i, j \rangle \in S_3$ whose CoS ID identifies CoS Label H.
4. $\langle S_1, CoSID_2, PT_1 \rangle$: All Service Frames from UNI $i$ to UNI $j$ where $\langle i, j \rangle \in S_1$ whose CoS ID identifies CoS Label M.
5. $\langle S_2, CoSID_2, PT_1 \rangle$: All Service Frames from UNI $i$ to UNI $j$ where $\langle i, j \rangle \in S_2$ whose CoS ID identifies CoS Label M.
6. $\langle S_3, CoSID_2, PT_3 \rangle$: All Service Frames from UNI $i$ to UNI $j$ where $\langle i, j \rangle \in S_3$ whose CoS ID identifies CoS Label M.

In this example, $PT_1$ is used for intra-Metro frames while $PT_3$ is used for inter-Metro frames. See section 6.6.1 for requirements relating to $\langle S, CoSID, PT \rangle$ for an EVC or OVC.

## 6.3 CoS Label

CoS Label is a term introduced in this IA and defined in sections 2 and 6.2. CoS Labels do not infer any specific implementation of network priority mechanisms (e.g., strict priority queuing, weighted fair queuing, etc.) in handling a frame. CoS Labels are H, M and L. These informally refer to High, Medium and Low. The order of the CoS Labels is based on the traffic classes in [5] and their associated PCP values.

CoS Label is independent of all Service Provider, Operator and other standards' CoS Names. Users of this IA, such as Operators and Service Providers, can assign any brand or marketing names desired to the MEF compliant CoS Labels for their own services.

## 6.4 CoS and Color Identifiers

For the purposes of this IA the terms identification and indication are used interchangeably.

### 6.4.1 CoS Identifier and Indication

At the UNI and the ENNI the CoS Name for each frame of the CoS Frame Set is indicated by a CoS Identifier (CoS ID). As specified in [2], [13] and [14] there are multiple CoS Identifiers (i.e., mechanisms specified for CoS Name identification) at the UNI and ENNI. Additional CoS Identifiers may be created in the future.

CoS Identifiers for a Service Frame at a UNI are defined in [2] section 6.8 and [13] section 7.5.3. Below are the 2 lists with abbreviated description that will be used in this IA.

From [2], Ethernet Virtual Connection Service Attributes - Class of Service Identifier Service Attribute:

- Based on EVC
- Based on Priority Code Point Field (i.e., EVC and C-Tag PCP value)
- Based on DSCP (i.e., EVC and DSCP value)
- Any of the above and based on L2CP[2].

[2] Methods to identify L2CP at a UNI are specified in [16]. DSCP is "one of the above" as per [2]. This is referring to use of DSCP for CoS ID of non-L2CP Service Frame and in that case CoS ID for L2CP frames is based on EVC + L2CP only (i.e., EVC+ DSCP otherwise).

From [13], OVC per UNI Service Attributes – Class of Service Identifiers:

- Based on OVC End Point
- Based on Priority Code Point Field (i.e., OVC End Point and C-Tag PCP value)
- Based on DSCP (i.e., OVC End Point and DSCP value).

Note that L2CP is not included for OVC per UNI since not included in [13]. Future phases of [13] may include L2CP. Future phases of CoS IA will align with these future phases of [13].

At an ENNI, the CoS Identifier for an ingress ENNI Frame, that is mapped to an OVC End Point and not mapped to or associated by a VUNI End Point, is as defined in [13] section 7.3.3 and Table 17. As per Section 7.3.3.1 of [13], the Class of Service Identifier in an ENNI Frame identifies the CoS Name (CoS Label if one of the CoS Names specified in this IA) for the receiving Operator. Note that in this IA the phrase "an ENNI Frame that maps to OVC End Point" (not to a VUNI End Point) can also be referred to as an EFO. Below is the abbreviated description of the CoS ID for EFO that will be used in this IA.

From [13], OVC End Point per ENNI Service Attributes – Class of Service Identifiers:

- S-Tag PCP Value (i.e., OVC and S-Tag PCP value).

This includes the case where all possible PCP values map to a single CoS Name thus yielding a single CoS ID for an OVC.

The CoS Identifier for an ingress ENNI Frame, that is mapped to a VUNI end point at an ENNI, is as defined in [14] section 7.3.1. Note that in this IA the phrase "an ENNI Frame that maps to a VUNI End Point" can also be referred to as an EFV.

Below is the abbreviated description of the CoS ID for EFV that will be used in this IA.

From [14], Service Attributes for an OVC End Point associated by the VUNI - VUNI Class of Service Identifiers:

- Based on OVC End Point (i.e., OVC End Point to which the frame is also mapped)
- Based on C-Tag Priority Code Point Field (i.e., OVC End Point and C-Tag PCP value)
- Based on DSCP (i.e., OVC End Point and DSCP value).

When CoS ID is EVC or OVC EP then there is only one CoS ID at the EI for all frames belonging to the EVC or OVC EP. The mapping of the EVC or OVC EP identification to the CoS Label is not in scope for this IA and thus is defined by mutual agreement between Service Provider and Subscriber at the UNI or between Operators at the ENNI. The purpose of specifying this case is to allow CoS Frame Sets that use this type of CoS ID to be compliant to the CPOs defined for the chosen CoS Label in this IA. Note that the PCP (or DSCP if untagged frame at UNI or EFV) may still be needed in the case of EVC or OVC EP CoS ID to indicate Color per frame. In this case PCP (or DSCP if untagged frame at UNI or EFV) values are not used for CoS ID. The term PCP indicates the presence of PCP in either tagged or priority tagged frames.

In this IA when the terms PCP or DSCP are used in conjunction with CoS ID it is short for the CoS Identifiers specified in [2], [13] and [14] (e.g., EVC and C-Tag PCP value, OVC End Point and C-Tag PCP value, OVC and S-Tag PCP value).

Specific values of PCP and DSCP for each CoS Label are specified in this IA, although additional values can also be mapped to each CoS Label.

See section 6.5 for specific requirements for CoS ID. When the CoS ID is based on L2CP type, a default CoS Label for tunnelled L2CP frames is recommended in this IA in section 6.5.1. The type of L2CP frames that are tunnelled are specified in [16]. Consistent with [13] and [14] this phase will not address L2CP at the ENNI or for ENNI Frames.

### 6.4.2 Color Identifier and Indication

Color Identifier (Color ID) is a Service Attribute introduced in Phase 1 of this IA that describes how the Service Frame or ENNI Frame indicates Color (e.g., Color Identifier can indicate a Yellow frame at an ENNI via the S-Tag PCP or DEI).

Color ID can be constrained by the choice of CoS ID in some cases. Color ID can be marked per frame or based on configuration per service (i.e., per CoS FS) when CoS ID is EVC or OVC EP (not per frame). In other words, Color ID based on EVC or OVC EP is only useful when there is no need to indicate Color per frame.

The PCP or DSCP may indicate Color and indicate CoS Label as common mechanisms of both Color ID and CoS ID. Note that when ingress Service Frames are untagged at the UNI, only DSCP, OVC EP or EVC can be used to indicate Color. Only DSCP can indicate Color per frame in this case. The use of DSCP in this IA is consistent with [10].

See section 6.5 for specific requirements for Color ID.

Note that Color indication can be critical, even in the case where the receiving Operator has not applied an Ingress Bandwidth Profile. This is because it can guide the receiving Operator on how to queue and schedule the frame.

## 6.5 CoS Label, CoS Identification and Color Identification Requirements

The following requirements address the specific CoS Label, CoS ID and Color ID requirements and associated Bandwidth Profile Color Modes for EIs. Unless otherwise stated, CoS ID for this IA applies to CoS Frame Sets at EIs as defined in [2], [13] and [14].

This IA provides the following specific CoS Name and CoS Label requirements:

**[R1]** The CoS ID for each frame in a CoS Frame Set at an EI **MUST** indicate the same CoS Name.

**[R2]** A CoS Label **MUST** be one of H, M, or L as per Table 4.

This IA does not allow multiple CoS Names or Labels for a single CoS Frame Set. Each CoS Frame Set must therefore map to a single CoS Name or Label. For example, a single CoS Frame Set cannot map to both CoS Label H and M.

This IA provides the following specific CoS and Color identification requirements.

With respect to Color Identification at the UNI:

**[R3]** Service Frame Color **MUST** be indicated using one of:
- EVC (all frames Green or all Yellow)
- OVC EP (all frames Green or all Yellow)
- C-Tag PCP value (i.e., CE-VLAN CoS in [2])
- DSCP value

[R3] means that at a UNI, a MEF CoS FS may have specific PCP or DSCP values as part of the CoS ID and Color ID as in Table 4 or may use only EVC or OVC End Point as the CoS ID as per Table 2. If EVC or OVC EP also provides Color ID then Color ID is not indicated per frame and must be all Green or all Yellow. If CoS ID based on EVC or OVC EP and Color ID is indicated per frame then Color ID will require use of PCP or DSCP as per Table 3. For the Subscriber to get the proper treatment of their frames the Subscriber needs to transmit frames with the Color ID indicated as in [R3].

With respect to Color Identification at the ENNI for EFV case (ENNI Frame mapped to VUNI):

**[R4]** When an ENNI Frame is mapped to a VUNI (EFV) by the receiver, the ENNI Frame Color **MUST** be indicated using one of:
- OVC EP (all frames Green or all Yellow)
- C-Tag PCP value (i.e., CE-VLAN CoS in [2])
- DSCP value

[R4] means that at an ENNI, when an ingress ENNI Frame is mapped to a VUNI end point (i.e., EFV), a MEF CoS FS may have specific PCP values as part of the CoS and Color ID as in Table 4 or may use only OVC End Point as the CoS ID as per Table 2. However, Color ID indication per frame requires the use of C-Tag PCP or DSCP to indicate Green vs. Yellow Color per frame. See [R9]. To get the proper treatment of the frames, the transmitting Operator needs to transmit frames to the receiving Operator with the Color ID indicated as in [R4] in the case of EFV.

With respect to Color Identification at the ENNI for EFO case (not mapped to VUNI):

**[R5]** When an ENNI Frame is mapped to an OVC End Point (EFO) by the receiver, any ENNI Frame Color **MUST** be indicated using one of:
- S-Tag PCP value
- S-Tag DEI.

[R5] means that at the ENNI, when an ingress ENNI Frame is mapped to an OVC EP and not mapped to a VUNI end point (i.e., EFO), the CoS ID for the frame includes the S-Tag PCP value and the Color ID for the frame can be indicated by either the S-Tag PCP value or the DEI. [R5] does not imply a requirement to support DEI. To get the proper treatment of the frames both Operators must choose the same alternative for Color ID from [R5] in the case of EFO.

With respect to relating CoS ID and Color ID at an EI:

**[R6]** When CoS ID is based on C-Tag PCP, any Color ID used **MUST** be based on the C-Tag PCP.

**[R7]** When CoS ID is based on DSCP, any Color ID used **MUST** be based on the DSCP.

**[R8]** When CoS ID is based on EVC or OVC EP at a UNI, any per frame Color ID used **MUST** be based on C-Tag PCP or DSCP using the values as per Table 3 .

**[R9]** When CoS ID is based on OVC EP at a VUNI (i.e., EFV), any per frame Color ID used **MUST** be based on C-Tag PCP or DSCP using the values as per Table 3.

[R8] and [R9] mean that when CoS ID is EVC or OVC EP and a potential for Yellow frames exists, the Color ID is based on PCP or DSCP.

With respect to applicability of the CoS Labels in this IA to various types of CoS Identifier including EVC or OVC EP:

**[R10]** When indicating one of the MEF CoS Labels, an Operator **MUST** use one of the CoS Identifier Types in Table 2 to indicate one of the MEF CoS Labels.

The intent of this requirement is to allow the Operator to apply the Performance Objectives and Parameters for the CoS Label for any CoS ID which a CoS Frame Set utilizes.

The ENNI Frame format is specified in [13]. With respect to IEEE 802.1ad-2005 ([5]) and the ENNI:

**[R11]** If IEEE DEI field is used to indicate Color it **MUST** be implemented as described in [5] in clause 9.7.

With respect to the Color Mode of Bandwidth Profile at the UNI:

**[D1]** For a given CoS Frame Set with a per CoS ID Ingress Bandwidth Profile, Color Mode parameter setting **SHOULD** be the same at all UNIs associated by the EVC supporting the CoS Frame Set.

Note that section 7.11.1 of [2] requires that the Ingress UNI Bandwidth Profile Color Mode parameter value of Color Blind be supported while Color Aware is optional.

With respect to Color indication at an EI:

**[R12]** When CoS ID is based on C-Tag PCP, the Color indication for a frame at an EI with no C-Tag **MUST** be determined by Subscriber/Service Provider agreement, i.e., all frames without a C-Tag are either Green or Yellow.

[R12] can include either Service Frames or EFV. The CoS ID requirements for these cases are in section 6.8.2 of [2] and section 7.3.1 of [14].

Note that sections 7.3.5-7.3.7 of [13] require that the Color Mode parameter of each Ingress and Egress Bandwidth Profile at the ENNI(s) be set to Color Aware mode for an EFO. However, [14] does not restrict the Color Mode parameter to Color Aware for an Ingress Bandwidth Profile for an EFV.

Note that Egress Bandwidth Profile requirements, such as Color Mode, are not included in this phase, but should be included in a later phase with intent they will be included when Multipoint CPOs are added.

With respect to the CoS and Color Identifiers at the EI:

**[D2]** When indicating an MEF CoS Label at an EI where a given frame's CoS Identifier includes C-Tag PCP, the PCP value **SHOULD** indicate the selected CoS Label and Color as per Table 4 column labeled *CoS and Color Identifiers, C-Tag PCP* and either *Color Green* or *Color Yellow*.

For example, at a UNI the Subscriber sender should set the PCP value to 5 for a frame associated with CoS Label H. The receiving Operator may map the ingress frame to their own internal mechanisms and values for indicating CoS Labels inside the MEN (outside the scope of this IA).

**[D3]** When indicating an MEF CoS Label at an EI where a given frame's CoS Identifier includes DSCP, the DSCP value **SHOULD** indicate the selected CoS Label and Color as per Table 4 column labeled *CoS and Color Identifiers, PHB (DSCP)* and either *Color Green* or *Color Yellow*.

**[R13]** When indicating an MEF CoS Label at an EI where a given frame's CoS Identifier includes S-Tag PCP and the Color Identifier also uses the S-Tag PCP, the S-Tag PCP value **MUST** indicate the selected CoS Label and Color as per the Table 4 column labeled *CoS and Color Identifiers, S-Tag PCP Without DEI Supported* and either *Color Green* or *Color Yellow*.

**[R14]** When indicating an MEF CoS Label at an EI where a given frame's CoS Identifier includes S-Tag PCP and the Color Identifier is DEI, the S-Tag PCP value **MUST** indicate the selected CoS Label as per the Table 4 column labeled *CoS and Color Identifiers* and *S-Tag PCP With DEI Supported* and the DEI **MUST** be used to identify the Color.

The use of DEI for Color Identification, as described in [R14], may free up additional values of the S-Tag PCP, but may not be feasible in the near term unless the networking equipment supports it (e.g., older Ethernet equipment and MPLS do not support DEI or an equivalent). DEI values are not shown explicitly in Table 4.

As far as this IA is concerned PCP and DSCP values not in Table 4 can be used in any way the Operator desires. This IA only specifies a subset of possible CoS Identifier values at EIs and is not applicable to how CoS Name is identified internal to a MEN. In the Three CoS Label Model, three PCP values are left open for Operator use. If a subset of the three labels is used additional values are available. It is possible for an Operator to reuse the PCP CoS Identifier values in Table 4 inside the MEN, but is not constrained to do so. The intent of Phase 2 is for the CoS ID and Color ID values (e.g., PCP and DSCP) specified to apply at the EIs and for the CPOs and Parameter values to apply between EIs consistent with [2] and [13].

The Per Hop Behavior (PHB) column in Table 4 provides the DSCP values used as part of the CoS Identifier. The table includes Expedited Forwarding (EF), Assured Forwarding (AF) and Default PHBs.

### 6.5.1 Default CoS Label for L2CP

To ensure consistent performance of Subscriber L2CP traffic for MEF services, the CoS IA defines a default CoS Label for Subscriber L2CP so L2CP can be identified as a distinct CoS Frame Set. A distinct CoS ID and CoS FS for L2CP allows for a specific Bandwidth Profile for this L2CP traffic, unique from the Bandwidth Profile used for customer data plane. This Bandwidth Profile can be used to apply a rate limiter to L2CP traffic, ensuring unpredicted excessive bursts of L2CP traffic do not impair performance of data Service Frames that share this CoS Label. The choice for a default CoS Label is based on low loss and low delay performance requirements for tunnelled L2CP frames.

**[D4]** Tunnelled Subscriber L2CP traffic, as listed in [16], Section 8, **SHOULD** map to CoS Labels as defined in Table 10.

## 6.6 Performance Tier and Ethernet Network Sections Model

The specification of CPOs in Phase 2 requires introduction of additional constructs to allow for different CPOs for a range of 'field of use' or 'applicability'. The construct of a set of applicable CPOs and Parameters is called a Performance Tier (PT). The construct used to relate an EVC and its OVCs including the associated PTs and CPOs is called the Ethernet Network Section (ENS) Model. ENS is used here to refer to a set of one or more MENs under a single or collaborative jurisdictional responsibility for CPO purposes. ENS here is equivalent with Ethernet Subnetwork in [17]. See [18] for a full definition of Network Section in the context of IP networks.

### 6.6.1 Performance Tier Model

A MEF Performance Tier (PT) contains a set of MEF CoS Performance Objectives (CPOs). For a given EVC, a particular PT may be applied to the EVC and a different PT may apply to an OVC that is part of the EVC. Different PTs have different CPOs specified in this IA.

When an Operator chooses a PT that is most applicable to a given service, the Operator may base that choice on any criteria (e.g., distance, link speed). A particular service can be based on an EVC or OVC. Setting proper PT (i.e., CPO set) for OVCs requires a concept of CPOs for each OVC that comprises an EVC that are consistent with the EVC CPOs. There may be various rationales for a Service Provider or Operator to assign a particular PT to a particular CoS Frame Set. Examples of rationales include, but are not limited to: approximate distance of the path frames traverse between EIs, number of switching hops or speed of links traversed, including access links. Note that the speed and technology used for links is a factor in delay that can be significant. For example, for a 1500 byte frame the serialization delay on a 2 Mb/s link can be about 6 ms and the delay for certain multiple physical link bonding technologies and associated fragmentation and de-fragmentation can add several additional milliseconds. These link delays are not usually considered significant for 10 Mb/s and higher links.

In terms of the requirements of this IA, distance between EIs is not a performance-related parameter that must be measured and reported by an Operator. Distance is only used to derive CPOs in this IA. Therefore precise definitions regarding how to measure and report distances between EIs are not necessary. The CPOs for a given PT may be viewed as a set of CPOs for a particular 'field of use' or 'area of applicability' from the Operator point of view. *The Operator need not adhere to the distances used in the derivation of a PT in their use of a particular MEF PT.*

In deriving PT CPOs for CoS IA, applications were explicitly mapped to one or more CoS and PT. In MEN implementations, particular applications may be mapped differently. For example, a subset of the Mobile Backhaul traffic may have some of the smaller FD/MFD value requirements and these requirements may only be achievable in a particular PT set that is based on relatively low propagation (minimum) delay. CoS IA will not normatively make such application or service exclusions however.

This IA uses distance as the primary means of describing PTs and deriving minimum delays. The distances stated for each PT can be considered as approximate distance limits for a given CoS Frame Set only if the assumptions stated in section 8.2 are applicable to the CoS Frame Set. Below are the four PTs defined in this IA with the format: PT Number (PT Name) - Description (distance, derived propagation delay used in CPO constraints to establish a minimum per PT).

- PT1 (Metro PT) – derived from typical Metro distances (<250 km, 2 ms),
- PT2 (Regional PT) - derived from typical Regional distances (<1200 km, 8 ms),
- PT3 (Continental PT) - derived from typical National/Continental distances (<7000 km, 44 ms),
- PT4 (Global PT) – derived from typical Global/Intercontinental distances (<27500 km, 172 ms)

Appendix section 8.2 describes how PT sets were derived. Distances are not normative and are only used to provide per PT delay related PT CPO constraints. The intent is to provide a range of PT sets that address Carrier Ethernet Networks of different geographic coverage, design and scope. Thus a four PT model is adopted for MEF CoS Labels. CPO value sets are specified in a separate table per PT.

A single PT (i.e., CPO set) will be used for each subset of ordered pairs (S) on an EVC or OVC. Per [2] for an EVC the ordered pairs consist of UNI pairs. Per [13] for an OVC the ordered pairs consist of OVC End Point pairs. The following summarize characteristics of S in terms of CoS Frame Set consistent with [2] and [13]:

**[R15]** A given CoS Frame Set **MUST** be based on a single subset of ordered pairs (S), a single CoS ID and a single Performance Tier.

[R15] allows a given subset of ordered pairs (S) to be used in the basis of more than one CoS Frame Set.

**[R16]** For a given EVC, there **MUST** be only one CoS Frame Set based on a given S and a given CoS ID

**[R17]** For a given OVC, there **MUST** be only one CoS Frame Set based on a given S and a given CoS ID.

[R16] and [R17] ensure that for a given EVC or OVC, if there is a CoS Frame Set based on given triple <S, CoS ID, PT>, there cannot be another CoS Frame Set based on the triple <S, CoS ID, PT*> where PT does not equal PT*.

Consider the example in section 6.2 where 24 different CoS Frame Set variants are possible but only up to 6 Frame Sets can exist simultaneously.

Note that in this IA the Parameters for the Performance metrics have the same value across all Performance Tiers.

As described in section 6.6.2 the OVCs that make up the EVC will each map to a PT which may be the same or different for each OVC and the EVC.

When one of the defined PTs (PT1-PT4) is used, the CPO parameters and values are defined in Sections 6.11.2 and 6.11.3. When a CoS Name is used that is not a CoS Label, other PTs (which are out of scope for this IA) may also be used.

*Note that this IA does not constrain which PT a Service Provider or Operator assigns a particular CoS Frame Set to.*

### 6.6.2 Ethernet Network Section Model

When this IA is to be applied to an OVC that, along with other OVCs, comprise an EVC, the MENs associated with the OVCs are referred to as Ethernet Network Sections (ENSs). In CoS IA Phase 2 an ENS generally aligns with a MEN.

Note that the definition of delay in [2] and [13] includes the delay incurred in traversing each ENNI thus the calculated delay for the UNI-UNI using concatenated OVCs will be slightly overstated. See Appendix 8.3 for more information.

Each OVC of a multiple MEN EVC has separate per-OVC CPOs that need to have consistency with the UNI-to-UNI CPOs for the EVC. Each CoS Frame Set associated with an OVC is assigned a PT for its set of CPOs. The ENS Model is referring to the relationships the various OVC CPOs have with the EVC CPOs and to other OVC CPOs that comprise the EVC. It may be necessary to concatenate the OVC CPOs to verify consistency with EVC CPOs. An ENS Model concatenation method example and associated recommendations is provided in Appendix section 8.3 for a subset of Performance metrics based on the methods in [8]. Concatenation is sometimes described as accumulating or combining sections. Concatenation is part of composing the end-to-end (UNI-to-UNI) CPOs. Sectionalization or allocation is the inverse of concatenation. Appendix section 8.3 provides no direct method of calculating allocation but does provide guidance for an indirect approach based on iteration. Sectionalization facilitates establishing CoS Frame Set performance budgets for each Operator or domain.

The ability to sectionalize EVC CPOs and concatenate OVCs is motivated by several factors. These include:

- Typical administrative and network boundaries that exist between MENs at ENNIs and within Operator networks between administrative and technology domains (e.g., between access networks and Ethernet networks).
- Establishment of clear responsibilities for an appropriate budgeted part of the UNI-to-UNI CPO for each MEN and its Operator (or domain within a MEN).
- The need to specify and report CPO related SLS results (e.g., performance for each OVC) in an EVC that traverses multiple MENs.

Below is an illustrative set of PT and ENS use cases for point-to-point EVCs and OVCs.

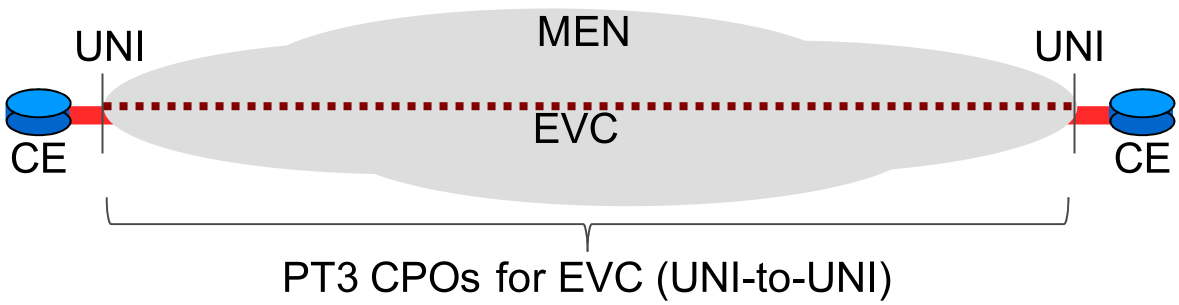


**Figure 3 – Example Performance Tier for a Single MEN EVC**

Figure 3 represents the simplest case, a point-point EVC in a single MEN. In this example, an EVC's CPOs utilize the PT3 set of CPOs for UNI-to-UNI SLS.

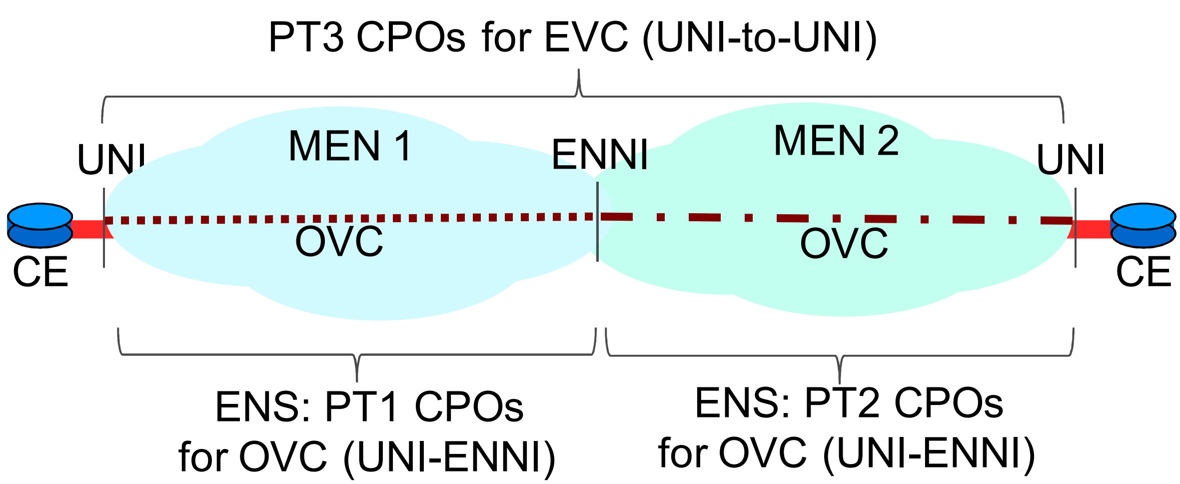


**Figure 4 – Example Performance Tiers for a Multiple MEN EVC and OVCs**

In Figure 4 the EVC traverses an ENNI that connects two MENs. The EVC will still have a UNI-to-UNI CPO set based on PT3 as represented by the bracket on top. The OVCs that comprise the EVC may have CPOs as represented by the bottom brackets. In this example the OVC in MEN1 (UNI-to-ENNI) uses PT1 and the OVC in MEN2 uses PT2 set of CPOs. Each of these OVCs is aligned with an ENS within an ENS Model that will relate OVC CPOs to the EVC CPOs using concatenation of OVC CPOs. Note that the OVC CPO values in PT1-4 in this IA are not likely to concatenate precisely to the EVC CPO values in PT1-4 tables in this IA. The methods, techniques and negotiations needed to arrive at acceptable objectives are beyond the scope of this IA. As stated previously, ENS Model includes both sectionalization and concatenation. While the example in Figure 4 is UNI-to-ENNI, a similar case can be constructed that includes ENNI-to-ENNI ENSs or the case of a multipoint EVC with a subset of ordered UNI pairs mapped to a PT.

The ENS Model could also be applied to scenarios in which a MEN that would appear from the outside as a single MEN is actually decomposed into multiple administrative based MENs. The CPOs for each of these component MENs can be composed into CPOs for the larger MEN using the ENS Model. An example of this would be a Service Provider that has subsidiaries that provide access service MENs on each end and a Wide-area Ethernet Network (WEN – see MEF 4 [6]) in the middle. These could be treated as three MENs for the purpose of setting CPOs. There could also be further subdivisions for performance within a MEN, but this is not in scope for CoS IA.

See Appendix section 8.3.4 for recommendations on how the apply the concatenation methods in section 8.3.

## 6.7 Performance Attributes and Metrics

Consistent with [2] and [13], Performance Objectives are defined such that they apply only to a Service or ENNI Frame when the frame is a Qualified Frame, which includes applicable Ingress Bandwidth Profile level of compliance of Green at the EI. In this IA, such frames are described as Qualified Frames. The preceding can be applied to both single and multiple-MEN EVCs. Bandwidth Profile compliance is defined further in section 6.8. Note that Phase 2 of CoS IA does not include CPOs for multipoint EVCs and OVCs so future phases may include additional considerations for multipoint.

Refer to [2] and [13] for complete definitions of Performance attributes, metrics and associated parameters.

Derivation of CPOs for this IA is found in Appendix section 8.4. The remainder of this section describes the Performance metrics and requirements for CPOs included in CoS IA Phase 2. Future phases of the CoS IA will align with future revisions of [2] and [13], if any.

Frame Delay (FD) and Mean Frame Delay (MFD) Performance form a pair for which this IA requires support for at least one. Either one or both of these two can apply to a given SLS. Similarly for Inter-Frame Delay Variation (IFDV) and Frame Delay Range (FDR) Performance, this IA requires support for at least one. Requirements below formalize this normatively. However, it should be noted that to support EVCs end-to-end with ENSs it is recommended that all Operators support the same choice of FD vs. MFD and IFDV vs. FDR. Furthermore for the case of ENS there are issues of sectionalization and concatenation to consider for Performance Objectives. See sections 6.6 and 8.3.

All included Performance metrics are one-way and therefore on a Point-to-Point EVC or OVC they apply to each direction and the Operator may elect to provide more stringent objectives than the CPO values in CoS IA for one or both directions.

### 6.7.1 Frame Delay Performance

Frame Delay (FD) Performance for subsets, S, of ordered UNI pairs in an EVC is defined in [2]. FD for subsets, S, of ordered OVC EP pairs at UNIs and ENNIs in an OVC is defined in [13].

Frame Delay for a Qualified Frame is the one-way delay that includes the delays encountered as a result of transmission across the ingress and egress UNIs and ENNIs (if present) as well as that introduced by the MEN. Note that FD Performance in [13] is defined using a Percentile ($P_d$) over a Time interval (T). In [2] it is defined using Percentile (*P)* over a Time interval (T). This IA will use $P_d$.

For EVCs there is an additional parameter (S) indicating a subset of the ordered UNI pairs. For OVCs there is an additional parameter (S) indicating a subset of the ordered OVC End Point pairs. While [2] and [13] do not specify values for objectives and parameters, $P_d$ or T, CoS IA specifies them for each CoS Label.

FD CPO requirements apply UNI-to-UNI, UNI-to-ENNI and ENNI-to-ENNI.

The intent of the IA is for an Operator to support at least one of FD or MFD for a CoS Frame Set that is associated with a CoS Label.

**[R18]** An SLS that is based on a MEF CoS Label **MUST** include at least one of either MFD or FD Performance as part of the SLS.

**[O1]** An SLS that is based on a MEF CoS Label **MAY** include both MFD and FD Performance as part of the SLS.

**[R19]** In an SLS that includes FD Performance and is based on a MEF CoS Label, the SLS **MUST** be specified per:

(1) FD Performance Objective for the associated *CoS Label* and *EVC/OVC Type* in Table 6, Table 7, Table 8, or Table 9, where Table selection is dependent on the PT selected; and

(2) specified $P_d$ and T Parameters for FD in Table 5

### 6.7.2 Mean Frame Delay Performance

Mean Frame Delay (MFD) Performance for subsets, S, of ordered UNI pairs in an EVC is defined in [2]. MFD for subsets, S, of ordered OVC EP pairs at UNIs and ENNIs in an OVC is defined in [13].

MFD is the arithmetic mean, or average of delays experienced by a set of frames that egress an EI as a result of an ingress frame at another EI in Time Interval (T). Further, these frames belong to the same CoS Frame Set.

Note that MFD Performance in [2] and [13] is defined using a Time interval (T). For EVCs there is an additional parameter (S) indicating a subset of the ordered UNI pairs. For OVCs there is an additional parameter (S) indicating a subset of the ordered OVC End Point pairs. While [2] and [13] do not specify values for the objective or parameter T, CoS IA will specify them for each CoS Label.

MFD CPO requirements apply UNI-to-UNI, UNI-to-ENNI and ENNI-to-ENNI.

**[R20]** In an SLS that includes MFD Performance and is based on a MEF CoS Label, the SLS **MUST** be specified per:

(1) MFD Performance Objective for the associated *CoS Label* and *EVC/OVC Type* in Table 6, Table 7, Table 8, or Table 9, where Table selection is dependent on the PT selected; and

(2) specified T Parameter for MFD in Table 5.

### 6.7.3 Inter-Frame Delay Variation Performance

Inter-Frame Delay Variation (IFDV) Performance for subsets, S, of ordered UNI pairs in an EVC is defined in [2]. IFDV for subsets, S, of ordered OVC EP pairs at UNIs and ENNIs in an OVC is defined in [13].

Inter-Frame Delay Variation Performance is defined in [13] as the Percentile ($P_v$) of the absolute values of the difference between the frame delays of Qualified Frame pairs under a list of specified conditions that includes parameters $\Delta t$ and T. In [2] the Percentile is defined using Percentile ($P$). This IA will use $P_v$.

For EVCs there is an additional parameter (S) indicating a subset of the ordered UNI pairs. For OVCs there is an additional parameter (S) indicating a subset of the ordered OVC End Point pairs. While [2] and [13] do not specify values for the objective or parameters $\Delta t$, $P_v$ or T, CoS IA specifies them.

IFDV CPO requirements will apply UNI-to-UNI, UNI-to-ENNI and ENNI-to-ENNI.

The intent of this IA is for an Operator to support at least one of FDR or IFDV on a compliant CoS Frame Set.

**[R21]** An SLS that is based on a MEF CoS Label **MUST** include at least one of either FDR or IFDV Performance as part of the SLS.

**[O2]** An SLS, that is based on a MEF CoS Label **MAY** include both FDR and IFDV Performance as part of the SLS.

**[R22]** In an SLS that includes IFDV Performance and is based on a MEF CoS Label, the SLS **MUST** be specified per:

(1) the IFDV Performance Objective for the associated *CoS Label* and *EVC/OVC Type* In Table 6, Table 7, Table 8, or Table 9, where Table selection is dependent on the PT selected; and

(2) specified $P_v$, $\Delta t$ and T Parameters for IFDV in Table 5.

### 6.7.4 Frame Delay Range Performance

Frame Delay Range (FDR) Performance for subsets, S, of ordered UNI pairs in an EVC is defined in [2]. FDR for subsets, S, of ordered OVC EP pairs at UNIs and ENNIs in an OVC is defined in [13].

FDR is described in detail in [2] and [13]. A simplified description is that FDR is the difference between the delay value at percentile ($P_r$) and the minimum delay value as mandated in [13]. In [2] FDR Performance is defined using Percentiles $P_x$ and $P_y$. This IA will use $P_r$ consistent with [13]. For EVCs there is an additional parameter (S) indicating a subset of the ordered UNI pairs. For OVCs there is an additional parameter (S) indicating a subset of the ordered OVC End Point pairs. While [2] and [13] do not specify values for the FDR Objective, Parameter $P_r$ or Time interval T, CoS IA specifies them.

FDR CPO requirements will apply UNI-to-UNI, UNI-to-ENNI or ENNI-to-ENNI.

**[R23]** In an SLS that includes FDR Performance and is based on a MEF CoS Label, the SLS **MUST** be specified per:

(1) the FDR Performance Objective for the associated *CoS Label* and *EVC/OVC Type* in Table 6, Table 7, Table 8, or Table 9, where Table selection is dependent

on the PT selected; and

(2) specified $P_r$ and T Parameters for FDR in Table 5.

MEN changes that alter delay such that delay is still within the SLS performance objectives for FD and MFD may lead to increases in FDR that cause it to miss FDR SLS objectives. For example, a topology change could increase or decrease the path distance thus increasing or decreasing the minimum delay during the interval T. This may increase the FDR for the interval T sufficiently to cause it to miss the FDR SLS. If this is a one-time event, however, the actual impact of the event at the application layer will be transient and may be insignificant. In such cases, the Service Provider and Subscriber or Service Provider and Operator may agree to ignore the FDR violation, especially if it can be shown that the impact of the topology change is the source of the miss or an IFDV objective, if one is specified, is met. This issue may need to be revisited in a later phase once MEF specifications include actual measurement and reporting of FDR and associated minimum delay.

### 6.7.5 Frame Loss Ratio Performance

Frame Loss Ratio (FLR) Performance for subsets, S, of ordered UNI pairs in an EVC is defined in [2]. FLR for subsets, S, of ordered OVC EP pairs at UNIs and ENNIs in an OVC is defined in [13].

FLR is defined in [2] and [13] as the ratio, expressed as a percentage, over a specified time interval (T), of the number of Qualified Frames not delivered divided by the total number of Qualified Frames that should have been delivered. For EVCs there is an additional parameter (S) indicating a subset of the ordered UNI pairs. For OVCs there is an additional parameter (S) indicating a subset of the ordered OVC End Point pairs. While [2] and [13] do not specify values for the objective or T, CoS IA specifies them.

FLR CPO requirements will apply UNI-to-UNI, UNI-to-ENNI and ENNI-to-ENNI.

**[R24]** In an SLS that is based on a MEF CoS Label, the SLS **MUST** be specified per:

(1) the FLR Performance Objective for the associated *CoS Label* and *EVC/OVC Type* in Table 6, Table 7, Table 8, or Table 9, where Table selection is dependent on the PT selected; and

(2) specified T Parameter for FLR in Table 5.

### 6.7.6 Availability and Resiliency Performance

Availability, High Loss Interval and Consecutive High Loss Interval performance are for a possible future phase. See [13] and [2].

## 6.8 BANDWIDTH PROFILE AND COLOR

[2] and [13] provide no requirements or guidelines for how the various Bandwidth Profile models should be applied in the various CoS ID options. For example, at the UNI the choice of “per UNI”, “per EVC” or “per CoS ID” Bandwidth Profile models are not constrained by the choice of CoS ID. For example, the choice of C-Tag PCP for CoS ID is very relevant when using a “per CoS ID” Bandwidth Profile, but the choice of C-Tag-PCP CoS ID does not preclude using a “per UNI” or “per EVC” Bandwidth profile model. The service specifications in [1] provide certain constraints for which Bandwidth Profile models are allowed for each MEF service. For example, [1] does not allow “per UNI” Ingress Bandwidth Profile or any form of Egress Bandwidth Profile for EPL Point-to-Point EVCs. For EVPL Point-to-Point EVCs all Bandwidth Profile models are allowed for Ingress Bandwidth Profile and only “per UNI” is an option for Egress Bandwidth Profile.

In Phase 2 this IA complements those requirements by recommending that the Bandwidth Profile granularity matches CoS ID granularity. Only when a single CoS ID is present at an EVC will a “per EVC” Bandwidth Profile ‘police’ at the granularity of CoS ID. For example, if multiple CoS IDs are mapped to an Ingress Bandwidth Profile “per EVC”, the Bandwidth Profile will not be able to ‘police’ Service Frames per CoS ID. This mismatch may allow too much traffic to be declared Green for some CoS IDs and not enough for others. To prevent this CoS IA mandates the use of “per CoS ID” as the Bandwidth Profile model at all Ingress Bandwidth Profiles.

Note that in the case of a UNI where there is All to One Bundling (e.g., EPL), the “per UNI” Bandwidth Profile model and the “per EVC” Bandwidth Profile model are equivalent. With a UNI where there is Service Multiplexing (e.g., a UNI with multiple EVPLs) the “per UNI” Ingress Bandwidth Profile model is allowed but will result in the same issues with not being able to ‘police’ Service Frames per CoS ID identified in the EVC example above.

[2] defines UNI Bandwidth Profile models for “per EVC” and “per CoS ID”. Section 6.4.1 of this IA describes the CoS IDs defined for the UNI in [2] and [13]. One of those CoS IDs is the EVC itself. Thus, “per EVC” and “per CoS ID”' Bandwidth Profiles are equivalent when CoS ID is derived from EVC and thus a single CoS ID.

Furthermore, if there are no CoS IDs based on Layer 2 Control Protocol type and the EVC is a CoS ID, these two models are also equivalent to the “per CoS ID” Bandwidth Profile model.

As described in more detail in 6.4.1, when CoS ID is EVC or OVC EP the mapping of the EVC or OVC EP identification to the CoS Label is not in scope for this IA and thus is defined by mutual agreement between Service Provider and Subscriber at the UNI or between Operators at the ENNI.

**[R25]** When Ingress Bandwidth Profiles are present, Ingress Bandwidth Profiles **MUST** utilize the “per CoS ID” model in [2] and [13] for MEF CoS Performance Objective compliance.

[R25] means that the Ingress Bandwidth Profile model and the CoS ID need to match to provide the best chance of delivering on the CPOs. Although the "per UNI", "per EVC" and "per OVC EP" Ingress Bandwidth Profile models are not explicitly addressed in the requirement they are addressed by the "per CoS ID" model for cases where that level of granularity matches the CoS ID. [R25] applies to both UNI and ENNI. Meeting [R25] is sufficient to satisfy the requirements for application of Bandwidth Profile in [1]. [R25] provides guidance to Operators in implementing the SLS associated with CPOs in this IA and provides guidance to Subscribers and vendors in supporting shaping per CoS ID at the CE.

For the case of focused overload of egress traffic at a UNI for a Multipoint-to-Multipoint or a Rooted-Multipoint EVC, MEF 10.2.1 [2] provides the option to exclude those discarded frames from the Availability performance. While multipoint, including this issue, are not in scope for this phase of CoS IA, future MEF documents may include addressing this issue for OVCs and for FLR as well.

### 6.8.1 Bandwidth Profile Compliance

CoS IA Phase 1 provided limited specification of Bandwidth Profile (BWP) and CoS Performance Objective relationships and concentrated on providing the CoS Label Model and structure. This phase provides more detailed specifications of BWP including burst alignment.

Bandwidth Profile is important to this IA because it determines which frames ingress to a MEN or egress from a MEN at each EI and the frame's compliance with the Ingress Bandwidth Profile determines Color and applicability of SLS. Ingress Bandwidth Profiles apply to frames entering a MEN at an EI and Egress Bandwidth Profiles apply to frames exiting a MEN at an EI.

In CoS IA Phase 2 CoS ID, Bandwidth Profile and Color are used consistent with [2] for the UNI and [13] for the ENNI. Identification of Color can be used to indicate which frames are deemed to be within or outside of the SLS according to the Ingress Bandwidth Profile and the definition of Qualified Frames from [2] and [13]. Levels of Ingress Bandwidth Profile compliance are Green when fully compliant (compliant with CIR, CBS), Yellow when there is sufficient Ingress Bandwidth Profile compliance for transmission but without SLS Performance Objectives (compliant with EIR, EBS) and Red or discarded when not Ingress Bandwidth Profile compliant with either. Green and Yellow frames are identified as such in this IA. Note that the ITU terminology in [8] for Green is Discard Ineligible frames and for Yellow/Red it is Discard Eligible frames.

Note that Table 2 provides CIR, EIR and CF constraints.

As stated in [2] and [13] all performance metrics are defined such that they only apply to Qualified Frames.

**[R26]** At the UNI, when an Ingress BWP per CoS ID is present that meets the requirements of [2], an MEF compliant CoS Frame Set **MUST** use the parameters and value constraints in the *Bandwidth Profile Constraint* column per the associated *CoS Label* row in Table 2

**[R27]** At the ENNI, when an Ingress BWP per CoS ID is present that meets the requirements of [13] or [14], an MEF compliant CoS Frame Set **MUST** use the parameters and value constraints in the *Bandwidth Profile Constraint* column per the associated *CoS Label* row in Table 2.

When there is no Ingress Bandwidth Profile, implicit rate limiting is provided by the bandwidth limits of the EI Ethernet link. The requirements in this CoS IA for the case of no Ingress Bandwidth Profile apply. In particular, any frame successfully transmitted across the EI is declared Green unless the Color Identifier indicates that it is Yellow in which case it is declared Yellow.

The constraints for the Bandwidth Profile parameters shown in this IA are expressed as "equal to", "greater than" or "greater than or equal to" zero (e.g., CIR = 0, CIR >0, CIR ≥ 0). Bandwidth Profile parameters and values that are not specified are not constrained by this IA.

Note that [2] and [13] mandate the CBS and EBS be greater than or equal to the MTU Size. MEF 13 [12] mandates minimum CBS of 8 * MTU (12176 bytes) for UNI Type 1.1.

### 6.8.2 Egress Bandwidth Profile Considerations

For a future phase.

## 6.9 EVC and Service Type Applicability

Any of the MEF CoS Labels can be used with any type of EVC that is described in [2] or any Service Type that is described in [1]. In particular, Point-to-Point EVCs could use the same CoS Label as some Multipoint-to-Multipoint EVCs. Still, at the ENNI a specific implementation might serve these different service types using separate treatment (e.g., queues). MEF CoS IA is intended to be applicable to Point-to-Point, Multipoint-to-Multipoint and Rooted-Multipoint EVCs including the case where some or all are present simultaneously on a given EI. However, CPOs for Multipoint are to be determined in a later phase.

For example, serving an EVP-LAN might be more complex than an EVPL. A given pair of UNIs on a Multipoint-to-Multipoint EVC may communicate Service Frames using different paths within a MEN and among different Operator's MENs compared to the paths and network traversed by Service Frames from another pair of UNIs on the same EVC. This and the variability of traffic between UNI pairs within a given S (with >2 EIs) within compliance of the Ingress Bandwidth Profile can complicate meeting CoS Performance Attribute Objectives for Multipoint EVCs and OVCs. Careful use of multiple sets (S) can help to better characterize the traffic in a multipoint EVC or OVC.

In Phase 2, CPOs for a given MEF CoS Label and PT are provided for Point-to-Point and placeholders are provided for Multipoint (i.e., Multipoint-to-Multipoint and Rooted-Multipoint) EVC types as shown in Table 6, Table 7, Table 8, and Table 9. Point–to-Point EVCs (e.g., EVPL service) could have more stringent CPOs compared to Multipoint CPOs (when they are specified in a later phase).

Consistent with [2], the MEF CPOs apply between sets of ordered pairs of UNIs on the EVC that are allowed to exchange traffic. When the CPOs are applied to a set of two or more ordered pairs of UNIs, for Multipoint-to-Multipoint and Rooted-Multipoint EVCs, the performance is based on the worst pair's performance (in set S) as described in [2]. Different PTs may be applied to each S (e.g., S with shorter distance between them can get a lower PT than the pairs with a greater distance between them).

## 6.10 OVC and Service Type Applicability

Consistent with [13], the MEF CPOs apply between sets of ordered pairs of OVC End Points associated by the OVC. When the CPOs are applied to a set of two or more ordered pairs of OVC End Points for Multipoint-to-Multipoint OVCs, the performance is based on the worst pair's performance as described in [13]. Different PTs may be applied to each S (e.g., S with the shorter distance between OVC End Point pairs can get a lower PT than the pairs with a greater distance between them).

## 6.11 CoS Label Model

The CoS Label Model Tables provide normative information for each MEF CoS Label in a Three CoS Label Model. The Tables provide: CoS Label, CPOs, Bandwidth Profile constraints, CoS Identifier and Color Identifier. Only the PCP and DSCP CoS ID components are specified with values. All CPO requirements refer to UNI-to-UNI, UNI-to-ENNI and ENNI-to-ENNI performance in Phase 2.

In CoS IA, FD, MFD, IFDV, FDR and FLR CPOs are specified normatively as one of the following:

1. *Numeric values* expressed in milliseconds (ms) for FD, MFD, IFDV and FDR. FLR will be expressed as a decimal number representing a percentage.
2. *Unspecified* performance for a particular CPO for a given CoS Label via N/S

In Phase 1, CPOs were expressed in relative terms. In Phase 2 CPOs are expressed using values. In Phase 2, the Phase 1 CoS Model table has been divided into several tables. The Three CoS Model Table columns for Performance Attributes/Objectives and Rows for EVC Type from Phase 1 (see Table 2 in [19]) have been moved to new per-PT tables. Parameters are provided in separate tables. The tables are renamed Three CoS Label Model to be more precise in terminology.

Since this CoS IA supports a Three CoS Label Model and its subsets, there is a need for interworking or mapping between the subsets. For example, Operator of MEN 1 adopts all CoS Labels in the Three CoS Label Model and Operator of MEN 2 adopts a subset with 2 CoS Labels including CoS Labels H and L. If MEN 1 and MEN 2 are connected via an ENNI there is a need for mapping between the two models. No specific subset mapping is specified in Phase 2, but later phases may specify examples of this mapping. See section 8.1.1.

### 6.11.1 Three CoS Label Model

This model, as shown in Table 2, Table 3 and Table 4, specifies three MEF CoS Labels denoted by CoS Labels H, M and L. There is no restriction on how Operators may use the PCP (i.e., 4, 6, 7) and DSCP values not specified. However, there are additional restrictions on use of PCP values in [2] and [13] that are reiterated in Section 6.11.4.

Table 2 introduces the CoS Labels and specifies the Bandwidth Profile constraints and CoS ID types in Phase 2. Table 3 provides the PCP and DSCP values used for per frame Color ID when CoS ID type is EVC or OVC EP. Table 4 identifies the PCP and DSCP values to be used to identify the CoS Label and per frame Color ID when CoS ID type is PCP or DSCP. Note that the EVC or OVC part of the CoS Identifier is not explicitly shown for these cases. Further, Table 4 does not include a separate column for identifying the CoS Label when the CoS Identifier of the CoS Frame Set is only EVC or OVC EP, e.g., PCP values are not relevant to CoS ID or may not be present as in the case of an EVC with only untagged frames. See [R10] for the specific requirement that allows for cases of EVC or OVC EP as CoS ID with the CoS Labels. EVC and OVC EP CoS ID indication is not constrained by this IA.

Note that the DSCP and associated Per Hop Behavior (PHB) are provided. However, DSCP is what is actually used in the Service Frame. Additional CoS Identifiers may be specified in future phases of CoS IA.

The specific values for PCP in Table 4 were derived from [5] using Tables 6-4 and G-5 Priority Code Point Decoding. The table row used is "5P3D" scheme (5 traffic classes of which 3 also have drop eligibility PCP values). See Section 8.6 for table excerpts.

In [5] (Table 6-4 "5P3D" row) there is a traffic class called "Best Effort" which is associated with PCP=1 when not drop eligible and PCP=0 when drop eligible. In this IA CoS Label L is aligned with this traffic class in [5]. In terms of Bandwidth Profile note that CoS Label L allows CIR or EIR = 0. The special case of CIR = 0 effectively results in no CPOs for the Performance Attributes in this IA (i.e., Not-specified (N/S)) while the case of CIR > 0 will require conformance with CPOs. From a DSCP perspective CoS Label L is a combination of AF1 (for CIR>0) and Default (for CIR=0) classes.

| CoS Label | Ingress EI Bandwidth Profile Constraints[1] | CoS ID Types | | | Example Applications |
|---|---|---|---|---|---|
| | | EVC or OVC EP | PCP or DSCP | L2CP Related | |
| H | CIR>0; EIR≥0 [2] | See Table 3 | See Table 4 | See Section 6.5.1 & [16] | VoIP and Mobile Backhaul Control |
| M | CIR>0; EIR≥0 | See Table 3 | See Table 4 | See Section 6.5.1 & [16] | Near-Real-Time or Critical Data Apps |
| L | CIR≥0; EIR≥0 [3] | See Table 3 | See Table 4 | See Section 6.5.1 & [16] | Non-critical Data Apps |

[1] EBS and Color Mode Bandwidth Profile parameters are not addressed in this table.
[2] EIR is not constrained though EIR=0 assumed since this IA does not specify Color Yellow PCP and DSCP for CoS Label H. Relaxation of EIR constraint may be used in some situations for certain applications such as Mobile Backhaul.
[3] Both CIR and EIR = 0 is not allowed as this would result in no conformant Service or ENNI Frames under steady state operation.

**Table 2: CoS Labels and CoS ID Types in CoS IA**

| CoS Label | CoS ID Types | Color Identifiers[1] | | | |
|---|---|---|---|---|---|
| | | C-Tag PCP | | PHB (DSCP) | |
| | | Color Green | Color Yellow | Color Green | Color Yellow |
| H | EVC or OVC EP [2] | 5, 3 or 1 | N/S in Phase 2 | EF or AF (10, 26 or 46) | N/S in Phase 2 |
| M | EVC or OVC EP [2] | 5, 3 or 1 | 2 or 0 | EF or AF (10, 26 or 46) | AF (0, 12, 14, 28 or 30) |
| L | EVC or OVC EP [2] | 5, 3 or 1 | 2 or 0 | EF or AF (10, 26 or 46) | AF (0, 12, 14, 28 or 30) |

[1] Specifies only the PCP or DSCP values to be used for Color ID when CoS ID is limited to EVC or OVC EP. EVC and OVC End Point indication for CoS ID is not constrained by CoS IA.
[2] EVC or OVC EP CoS ID would be different to differentiate CoS Labels H, M and L for different CoS Frame Sets on a given EI

**Table 3: Color ID Values when CoS ID is Only EVC or OVC EP**

In Table 3 the PCP and DSCP values for each CoS Label include all of the values specified in Table 4 for that CoS Label. This is due to the values in Table 3 only indicating Color (not indicating CoS Label). This is possible only when the CoS ID is not indicated with the PCP or DSCP, but rather with the EVC or OVC EP alternative mechanisms. For example, consider a case of a service multiplexed UNI with two EVCs. CoS ID of EVC1 indicates CoS Label is H, while PCP and DSCP values are not used for CoS ID, but only for Color ID as needed. CoS ID of EVC2 indicates CoS Label is L and again the PCP and DSCP need only indicate Color as needed.

| CoS Label | CoS and Color Identifiers [1] | | | | | | |
|---|---|---|---|---|---|---|---|
| | C-Tag PCP | | PHB (DSCP) | | S-Tag PCP Without DEI Supported | | S-Tag PCP With DEI Supported |
| | Color Green | Color Yellow | Color Green | Color Yellow | Color Green | Color Yellow | |
| H | 5 | N/S in Phase 2 | EF (46) | N/S in Phase 2 | 5 | N/S in Phase 2 | 5 |
| M | 3 | 2 | AF31 (26) | AF32 (28) or AF33 (30) | 3 | 2 | 3 |
| L | 1 | 0 | AF11 (10) | AF12 (12), AF13 (14) or Default (0) | 1 | 0 | 1 |

[1] Full CoS Identifier includes EVC or OVC End Point. Table specifies only the PCP or DSCP values to be used with EVC or OVC End Point to specify a CoS ID. EVC and OVC End Point indication is not constrained by CoS IA.

**Table 4: CoS Identifiers and Color Identifiers**

Note that EVC and OVC EP are valid CoS IDs that are not included in Table 4, but can conform to the CPOs and Parameters for CoS Labels just as the CoS IDs above. See Table 2.

### 6.11.2 Performance Parameters

Table 5 specifies Performance Parameters as required to derive and specify CPOs. The CPOs in Table 6, Table 7, Table 8 and Table 9 are based on the Parameter values in Table 5. For a given CPO value an Operator may provide Parameter values less than the maximum or more than the minimum Parameter values (i.e., more stringent Parameter values) and comply with this IA. From a Service OAM Performance Monitoring (SOAM-PM) point of view, these Parameter values provide a basis for how the measurements are made for the CoS Frame Sets.

In Phase 2, Parameters associated with each Performance metric are stated separately for each CoS Label due to variances in Percentiles, though the values are uniform across PTs. Since Phase 2 CPO scope is limited to Point-to-Point EVC/OVC Types, Multipoint will be addressed in a later phase. There is no requirement that Parameters be uniform across CoS Labels, PTs, EVC/OVC Types or between Performance metrics with similar Parameters. For example the T associated with FLR may be different from the T associated with FD. However, there is a recommendation for uniformity across particular OVCs that comprise an EVC. See section 8.3.4.

Parameters may not be specified (i.e., N/S) in this IA when the associated CPOs are not specified (i.e., N/S).

| Performance Metric | Parameter Name | Parameter Values for CoS Label H | Parameter Values for CoS Label M | Parameter Values for CoS Label L |
|---|---|---|---|---|
| FD | Percentile ($P_d$) | ≥ 99.9th | ≥ 99th | ≥ 95th |
| | Time Interval (T) | ≤ Month | ≤ Month | ≤ Month |
| MFD | Time Interval (T) | ≤ Month | ≤ Month | ≤ Month |
| IFDV | Percentile ($P_v$) | ≥ 99.9th | ≥ 99th or N/S[1] | N/S |
| | Time Interval (T) | ≤ Month | ≤ Month or N/S[1] | N/S |
| | Pair Interval (Δt) | ≥ 1sec | ≥ 1sec or N/S[1] | N/S |
| FDR | Percentile ($P_r$) | ≥ 99.9th | ≥ 99th or N/S[1] | N/S |
| | Time Interval (T) | ≤ Month | ≤ Month or N/S[1] | N/S |
| FLR | Time Interval (T) | ≤ Month | ≤ Month | ≤ Month |
| Availability | TBD | TBD | TBD | TBD |
| High Loss Interval | TBD | TBD | TBD | TBD |
| Consecutive High Loss Interval | TBD | TBD | TBD | TBD |

[1] Parameters are N/S only when CPO is N/S
Note: each parameter value > 0

**Table 5: CoS Label H, M and L Parameter Values**

In this phase Performance Parameter values are stated within a single table. In future phases they may be stated per PT.

### 6.11.3 CoS Performance Objectives Per Performance Tier

Table 6, Table 7, Table 8 and Table 9 provide CPOs for each Performance metric per each CoS Label. Each Table provides CPOs for one PT of the four PTs. These are normative as per the requirements that refer to them. Note: Multipoint also includes Rooted Multipoint as per [2].

In the case of an EVC that is comprised of multiple OVCs, the EVC CPOs in Table 6, Table 7, Table 8 and Table 9 may not be met even if CoS Label mapping is aligned, such as when there is insufficient alignment of CBS between Operators and/or insufficient shaping at the ENNI. In other words, the EVC performance may be impacted enough to cause performance results that miss some CPOs for the EVC or create the need to utilize a less stringent PT. For informative guidance on these issues see Burst Size and Shaper Considerations for ENNI, Section 8.7.

| Performance Metric | CoS Label H | | CoS Label M | | CoS Label L[1] | | Applicability |
|---|---|---|---|---|---|---|---|
| | Pt-Pt | Multipt | Pt-Pt | Multipt | Pt-Pt | Multipt | |
| FD (ms) | ≤ 10 | TBD | ≤ 20 | TBD | ≤ 37 | TBD | At least one of either FD or MFD required |
| MFD (ms) | ≤ 7 | TBD | ≤ 13 | TBD | ≤ 28 | TBD | |
| IFDV (ms) | ≤ 3 | TBD | ≤ 8 or N/S [2] | TBD | N/S | TBD | At least one of either FDR or IFDV required |
| FDR (ms) | ≤ 5 | TBD | ≤ 10 or N/S [2] | TBD | N/S | TBD | |
| FLR (percent) | ≤ .01% i.e. $10^{-4}$ | TBD | ≤ .01% i.e. $10^{-4}$ | TBD | ≤ .1% i.e. $10^{-3}$ | TBD | |
| Availability | TBD | TBD | TBD | TBD | TBD | TBD | |
| High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |
| Consecutive High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |

[1] Ingress Bandwidth Profile parameters may be chosen such that no frames are subject to SLS.
[2] Compliant services may leave this objective not specified.

**Table 6: PT1 (Metro PT) CPOs**

| Performance Metric | CoS Label H | | CoS Label M | | CoS Label L[1] | | Applicability |
|---|---|---|---|---|---|---|---|
| | Pt-Pt | Multipt | Pt-Pt | Multipt | Pt-Pt | Multipt | |
| FD (ms) | ≤ 25 | TBD | ≤ 75 | TBD | ≤ 125 | TBD | At least one of either FD or MFD required |
| MFD (ms) | ≤ 18 | TBD | ≤ 30 | TBD | ≤ 50 | TBD | |
| IFDV (ms) | ≤ 8 | TBD | ≤ 40 or N/S[2] | TBD | N/S | TBD | At least one of either FDR or IFDV required |
| FDR (ms) | ≤ 10 | TBD | ≤ 50 or N/S[2] | TBD | N/S | TBD | |
| FLR (percent) | ≤ .01% i.e., $10^{-4}$ | TBD | ≤ .01% i.e., $10^{-4}$ | TBD | ≤ .1% i.e., $10^{-3}$ | TBD | |
| Availability | TBD | TBD | TBD | TBD | TBD | TBD | |
| High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |
| Consecutive High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |

[1] Ingress Bandwidth Profile parameters may be chosen such that no frames are subject to SLS.
[2] Compliant services may leave this objective not specified.

**Table 7: PT2 (Regional PT) CPOs**

| Performance Metric | CoS Label H | | CoS Label M | | CoS Label L[1] | | Applicability |
|---|---|---|---|---|---|---|---|
| | Pt-Pt | Multipt | Pt-Pt | Multipt | Pt-Pt | Multipt | |
| FD (ms) | ≤ 77 | TBD | ≤ 115 | TBD | ≤ 230 | TBD | At least one of either FD or MFD required |
| MFD (ms) | ≤ 70 | TBD | ≤ 80 | TBD | ≤ 125 | TBD | |
| IFDV (ms) | ≤ 10 | TBD | ≤ 40 or N/S [2] | TBD | N/S | TBD | At least one of either FDR or IFDV required |
| FDR (ms) | ≤ 12 | TBD | ≤ 50 or N/S [2] | TBD | N/S | TBD | |
| FLR (percent) | ≤ .025% i.e., $2.5x10^{-4}$ | TBD | ≤ .025% i.e., $2.5x10^{-4}$ | TBD | ≤ .1% i.e., $10^{-3}$ | TBD | |
| Availability | TBD | TBD | TBD | TBD | TBD | TBD | |
| High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |
| Consecutive High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |

[1] Ingress Bandwidth Profile parameters may be chosen such that no frames are subject to SLS.
[2] Compliant services may leave this objective not specified.

**Table 8: PT3 (Continental PT) CPOs**

| Performance Metric | CoS Label H | | CoS Label M | | CoS Label L[1] | | Applicability |
|---|---|---|---|---|---|---|---|
| | Pt-Pt | Multipt | Pt-Pt | Multipt | Pt-Pt | Multipt | |
| FD (ms) | ≤ 230 | TBD | ≤ 250 | TBD | ≤ 390 | TBD | At least one of either FD or MFD required |
| MFD (ms) | ≤ 200 | TBD | ≤ 220 | TBD | ≤ 240 | TBD | |
| IFDV (ms) | ≤ 32 | TBD | ≤ 40 or N/S [2] | TBD | N/S | TBD | At least one of either FDR or IFDV required |
| FDR (ms) | ≤ 40 | TBD | ≤ 50 or N/S [2] | TBD | N/S | TBD | |
| FLR (percent) | ≤ .05% i.e., $5x10^{-4}$ | TBD | ≤ .05% i.e., $5x10^{-4}$ | TBD | ≤ .1% i.e., $10^{-3}$ | TBD | |
| Availability | TBD | TBD | TBD | TBD | TBD | TBD | |
| High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |
| Consecutive High Loss Interval | TBD | TBD | TBD | TBD | TBD | TBD | |

[1] Ingress Bandwidth Profile parameters may be chosen such that no frames are subject to SLS.
[2] Compliant services may leave this objective not specified.

**Table 9: PT4 (Global PT) CPOs**

### 6.11.4 PCP and DSCP Mapping

#### 6.11.4.1 UNI Mapping

When the CoS ID is based on PCP or DSCP, full mapping of PCP or DSCP values at a UNI is required in [2] Section 6.8.2 or 6.8.3 to ensure that customer frames are not inadvertently discarded and to simplify configuration of customer equipment.

For a multi-CoS EVC that supports only the standard MEF CoS Labels as defined in this document, tables providing examples of full PCP and DSCP mapping at a UNI are located in Appendix Section 8.5. Providing the same CoS Label mapping on all UNIs for a given EVC will minimize customer confusion.

#### 6.11.4.2 ENNI Mapping

For a multi-CoS OVC full PCP mapping is required at an ENNI as per [13] (requirement number R83). Phase 2 does not provide examples of mappings, but future phases may.

### 6.11.5 L2CP CoS Mapping

Note: The methods for classifying L2CP CoS ID are defined in [16].

Table 10 defines the mapping of L2CP to a default CoS Label for each combination of multiple subscribed CoS Labels. In the case where only a single CoS Label is subscribed, L2CP shares this CoS Label with data Service Frames. The M CoS Label is chosen for L2CP whenever available, based on its superior loss performance, and a desire to keep it separate from real-time applications.

| Subscribed CoS Labels | Default L2CP CoS Label |
|---|---|
| H, M, L | M |
| H, M | M |
| H, L | H |
| M, L | M |

**Table 10: L2CP CoS Mapping**

### 6.11.6 Requirement Sets

There is a distinction between the normative content related to CoS ID and Color ID values (e.g., PCP values) in sections 6.4 and 6.11.1 and the normative content on CoS Performance Objective values (CPOs) and Parameter values in sections 6.7, 6.11.2 and 6.11.3. These two sets will be referred to as *CoS/Color ID* and *Performance*. The requirements for CoS/Color ID values in the former is not generally dependent on the CoS Performance requirements in the later, nor is the reverse a dependency.

For example, for a given CoS Frame Set the CPO values can meet the Performance requirement set even if the PCP marking values do not meet the CoS/Color ID requirement set. The reverse can also be true. This means that the PCP marking values can meet the CoS/Color requirements set while the CPO values do not meet the Performance set for a given CoS Frame Set.

The requirements in this IA up through section 6.5 are members of the CoS/Color ID Set and requirements after section 6.5 are members of the Performance Set.

Three cases of meeting these requirement sets exist: (A) meeting requirements in both sets; (B) meeting requirements in CoS/Color ID set but not Performance set; and (C) meeting requirements in the Performance set, but not the CoS/Color ID set. For Point-to-Point EVCs or OVCs all three cases are possible within the scope of this phase of CoS IA, but for Multipoint EVCs and OVCs only Case (B) is possible in this phase since Multipoint CPOs are out of scope.

The full benefit and value of CoS IA is achieved when both sets of requirements are met (case (A)). In the case of Point-to-Point, Case (B) is also useful when there are Performance Objectives, but they may not meet the CPOs in this IA (e.g., established prior to this IA).

## 7. References

[1] MEF Technical Specification MEF 6.1, “Ethernet Services Definitions - Phase 2”

[2] MEF Technical Specification, MEF 10.2, “Ethernet Services Attributes - Phase 2” and MEF 10.2.1, “Performance Attributes Amendment to MEF 10.2”

[3] IEEE 802.1Q – 2005, “Virtual Bridged Local Area Networks”

[4] RFC 2119, “Key words for use in RFCs to Indicate Requirement Levels”, S. Bradner

[5] IEEE 802.1ad-2005, “Virtual Bridged Local Area Networks – Amendment 4: Provider Bridges”

[6] MEF Technical Specification MEF 4, “Metro Ethernet Network Architecture Framework - Part 1: Generic Framework”

[7] Inter-provider Quality of Service, MIT Communications Futures Program, 2006

[8] ITU-T Recommendation Y.1541, “Network performance objectives for IP-based services”

[9] MEF Technical Specification MEF 3, "Circuit Emulation Service Definitions, Framework and Requirements in Metro Ethernet Networks"

[10] RFC 2597, "Assured Forwarding PHB Group", Heinanen

[11] ITU-T Recommendation I.356, "B-ISDN ATM layer cell transfer performance," March 2000

[12] MEF Technical Specification MEF 13, "UNI Type 1 Implementation Agreement"

[13] MEF Technical Specification MEF 26.1, "ENNI Phase 2"

[14] MEF Technical Specification MEF 28, "ENNI Support of UNI Tunnel Access and Virtual UNI (VUNI)"

[15] ITU-T Recommendation Y.1540, "Internet protocol data communication service – IP packet transfer and availability performance parameters"

[16] MEF Technical Specification MEF 6.1.1, "Layer 2 Control Protocol Handling Amendment to MEF 6.1"

[17] MEF Technical Specification MEF 12.1, "Carrier Ethernet Network Architecture Framework Part 2: Ethernet Services Layer - Basic Elements"

[18] ITU-T Recommendation Y.1563, "Ethernet frame transfer and availability performance"

[19] MEF Technical Specification MEF 23, "Carrier Ethernet Class of Service – Phase 1"

[20] MEF Technical Specification MEF 26.0.2, "OVC Layer 2 Control Protocol Tunneling Amendment to MEF 26"

# 8. Appendices (informative)

## 8.1 Potential Work Areas for Later Phases of MEF CoS IA

### 8.1.1 CoS Subset Mapping

CoS IA allows subsets of the three CoS Labels to be supported. It also allows for additional CoS Names to be supported beyond the three CoS Labels specified. Thus at the ENNI there is a need for Operators to map the 7 possible subsets of CoS Labels that may be supported by an Operator (i.e., H/M/L, H/L, H/M, M/L, H, M, L). Phase 2 does not specify how such mapping should be done but leaves it to the Operators, inclusive of the Service Provider, to negotiate. A rationale is that in most cases there are also non-MEF CoS Names involved in a given Operator's network and these must be accounted for in any mapping schema for the Operator to mark or remark frames for transmittal across an ENNI to another Operator. Later Phases of CoS IA may provide further guidance on mapping.

### 8.1.2 Other Potential Work Areas

Future phases may also address Multipoint CPOs and Parameters, additional Bandwidth Profile Parameter guidance and additional Performance Metrics such as Availability, HLI and CHLI.

## 8.2 Performance Tier Model Derivation

Assumptions for PTs:

- PT distances represent the path a frame would traverse and thus drive associated propagation delay minimums for FD/MFD/FDR
- Though number of switch hops generally increases with longer distance PTs, hops will not be quantified
- For simplicity, PT CPOs are expressed as constants based on the maximum distance for the PT rather than formulas with distance variables
- PTs are derived with certain distance and application assignments
- PTs can be arbitrarily assigned to given services by Operators based on factors in or outside the scope of this IA
- All links, including access links, will have a link speed of at least 10 Mb/s, with the notion that a given service may utilize a "higher" PT for slower links based on Operator discretion.

  

A four PT model was chosen to allow for sufficient granularity and cover range from small area networks and applications to global. This IA uses distance as the primary means of describing PTs. Below are the four PTs defined in this IA with the format: PT Number (PT Name) - Description (distance, derived propagation delay used in CPO constraints to establish a minimum per PT).

- PT1 (Metro PT) – derived from Metro distances (<250 km, 2 ms*)
- PT2 (Regional PT) - derived from Regional distances (<1200 km, 8 ms*)
- PT3 (Continental PT) - derived from National/Continental distances (<7000 km, 44 ms*)
- PT4 (Global PT) – derived from Global/Intercontinental distances (<27500 km, 172 ms*)
    - Based on I.356 [11].

*Minimum MFD based on distance * .005 ms/km * 1.25 where distance is in kilometers (km), .005 ms/km propagation delay and 1.25 is route/airline distance ratio. Distance is difficult to ascertain in real-networks as path (i.e., circuit) distance is unknown or may vary due to routing or other path changes (e.g., dynamic control protocols). In real MENs there may be additional delays (e.g., switch hops, buffering, shaping, serialization for low speed links). Note: FD > MFD.

An Operator's Ethernet service compliance with this IA does not depend on adherence to PT distances. As stated in the normative sections, a given service may utilize a particular PT for reasons other than EI to EI distance of the service.

### 8.2.1 Low Speed Link Considerations

Delay can be significantly impacted by low speed access or links in a MEN. In CoS Phase 2 this is accounted for by the choice of PT for a service or UNI pair within a service. This is simpler than a Low Speed Factor that is applied to the CPO per CoS Label. For example, if a service would otherwise utilize PT1 CPOs it could utilize PT2 due to its use of sub-10Mb/s low speed links in the access between the NID and the core of the MEN. Additional low speed performance considerations are contained in [8] and [7].

## 8.3 Ethernet Network Section Model – Composing UNI-UNI Values

ITU-T Recommendation Y.1541 [8] defines methods for concatenating performance objectives or measurements associated with network sections, thus combining their performance to estimate the complete path (i.e., composing). This Appendix reproduces the equations using MEF variables where possible and uses MEF terminology whereby ENS replaces the term network sections used in [8].

While these methods are applicable to both objective setting and measurements, the methods are often not needed for measurements if ENS (e.g., UNI-ENNI for the OVCs) and end-to-end (e.g., UNI-UNI for the EVC) measurements are available.

When combining the metrics based on percentiles, it is a gross over-estimate to simply add the performance values for each ENS. However, there may be circumstances when even this over-estimate will suffice. For example, consider two ENSs, each of which has FDR of 2 ms. If the Subscriber is satisfied with 4 ms, simple addition could suffice. If the Subscriber requires 3 ms, then simple addition is not sufficient.

As mentioned in section 6.6.2, this IA provides no direct method of calculating allocation but the concatenation methods can be used to evaluate proposed OVC ENS CPOs against an EVC CPOs and through iteration adjust EVC or OVC objectives to guide the determination of OVC CPOs. Iteration is practical based on a small range/set of potential CPOs for the OVCs under consideration and a small number of ENS (i.e., usually 2-4).

The following table illustrates the mapping used, to the extent possible. Note that many ITU-T variables do not have a counterpart in MEF and that [8] does not address a metric equivalent to the MEF IFDV.

| Metric/Parameter | MEF 23.1/26.1 | Y.1541 | Notes |
|---|---|---|---|
| UNI-UNI One-way Delay Distribution | | $T$ | No MEF equivalent |
| SLS Interval | $T$ | | No ITU-T equivalent |
| Subset of ordered UNI pairs | $S$ | | No ITU-T equivalent |
| $k^{th}$ Network Section | | $k$ | No MEF equivalent |
| Mean One-way Delay | $\bar{\mu}_{TS}$ | $\mu_k$ | |
| Variance of One-way Delay | | $\sigma_k^2$ | No MEF equivalent |
| Probability or Percentile of interest | $P_d$ or $P_r$ | $p$ | $P_d$ for Frame Delay and or $P_r$ for Frame Delay Range |
| Delay at Percentile | $\bar{d}_{TdS}$ , $\bar{d}_{TrS}$ or $\bar{d}_{TRS}$ | $t_k, t$ | Frame Delay ($d$), or Frame Delay (r), Frame Delay Range (R); $t_k$ & $t$ are values of delay used in the Steps below |
| Skewness | | $\gamma_k$ | No MEF equivalent |
| Third moment | | $\omega_k$ | No MEF equivalent |

| Metric/Parameter | MEF 23.1/26.1 | Y.1541 | Notes |
|---|---|---|---|
| Value of the standard normal distrib. at *p* | | $x_p$ | No MEF equivalent |
| Loss Ratio | $FLR_{T,S}$ | $IPLR_k$ | |

**Table 11: MEF – ITU Variable Mapping**

### 8.3.1 Mean delay

For the Mean Frame Delay (MFD), or $\overline{\mu}_{TS}$ performance parameter (grouped with performance metrics in this IA), the UNI-UNI performance is the sum of the means contributed by Ethernet Network Sections.

$$\overline{\mu}_{TS} = \overline{\mu}_{TS1} + \overline{\mu}_{TS2} + \overline{\mu}_{TS3} + ... + \overline{\mu}_{TSn}$$

The units of $\overline{\mu}_{TS}$ values are seconds.

Note that the definition of delay in MEF per [2] and [13] includes the delay incurred in traversing the External Interface thus the calculated delay for the UNI-UNI using this concatenation method will be overstated. The sum of per ENS delays will be greater than the UNI-UNI delay. In general this overstatement is likely to be small in terms of modeling objectives and in terms of measurements may not be feasible to capture precisely as defined. This is not addressed in CoS IA Phase 2.

### 8.3.2 Loss ratio

For the Frame Loss Ratio ( $FLR_{T,E}$ ) performance metric, the UNI-UNI performance may be estimated by inverting the probability of successful frame transfer across *n* Ethernet Network Sections (*En*), as follows:

$$FLR_{T,E} = 1 - \{(1 - FLR_{T,E1}) \times (1 - FLR_{T,E2}) \times (1 - FLR_{T,E3}) \times ... \times (1 - FLR_{T,En})\}$$

This relationship does not have limits on the parameter values, so it is preferred over other approximations, such as the simple sum of loss ratios.

The units of $FLR_{T,S}$ values are lost Qualified Frames per total Qualified Frames sent. This is equivalent to MEF FLR except that it is not expressed as a percentage.

### 8.3.3 Relationship for delay and delay range

The relationship for estimating the UNI-UNI Frame Delay ($\bar{d}_{TdS}$) or the Frame Delay Range ($\bar{d}_{TRS}$) performance from the Ethernet Network Section values must recognize their sub-additive nature and is difficult to estimate accurately without considerable information about the individual delay distributions. If, for example, characterizations of independent delay distributions are known or measured, they may be convolved to estimate the combined distribution. This detailed information will seldom be shared among Operators, and may not be available in the form of a continuous distribution. As a result, the UNI-UNI delay estimation may have accuracy limitations.

The relationship for combining Frame Delay at $P_d$, or Frame Delay Range (i.e., delay at $P_r$ less minimum delay) values is given below. Note that $P_d$ parameter value is equal to $P_r$ parameter value for this IA for a given CoS Label and PT.

The problem under consideration can be stated as follows: estimate the quantile $\bar{d}_{TRS}$ of the UNI-UNI Frame Delay Range $T$ as defined by the condition:

$\Pr(T < \bar{d}_{TRS}) = p$ where $p = P_r/100$ for UNI-UNI Frame Delay Range.

A similar relation for UNI-UNI Frame Delay would be based on $\bar{d}_{TdS}$ and $p=P_d/100$.

When using the methods below to calculate Frame Delay Range, the calculations are based on using the difference between the delay and the minimum delay. In other words, all delay values are normalized by removing the minimum delay observed over T.

**Step 1**

Measure the mean and variance for the delay for each of $n$ Ethernet Network Sections. Estimate the mean and variance of the UNI-UNI delay by summing the means and variances of the component distributions.

$$\bar{\mu}_{TS} = \sum_{k=1}^{n} \bar{\mu}_{TSk}$$

$$\sigma^2 = \sum_{k=1}^{n} \sigma_k^2$$

**Step 2**

Measure the quantiles for each delay component at the probability of interest, e.g., $P_d = 99.9$ and $p = 0.999$. Estimate the corresponding skewness and third moment using the formula shown below, where $x_{0.999} = 3.090$ is the value satisfying $\Phi(x_{0.999}) = 0.999$ where $\Phi$ denotes the standard normal (mean 0, variance 1) distribution function. Note that $x_{0.999}$ is an example based on 99.9th percentile. This IA also recommends use of other percentiles including 95th and 99th which yield $x_{0.95}$= 1.645 and $x_{0.99}$= 2.33.

$$\gamma_k = 6 \cdot \frac{x_P - \frac{t_k - \bar{\mu}_{TSk}}{\sigma_k}}{1 - x_P^2}$$

where $t_k$ represents delay at $x_P$ based on $P_d/100$ for Frame Delay or where $t_k$ represents delay less minimum delay at $x_P$ based on $P_r/100$ for Frame Delay Range.

$$\omega_k = \gamma_k \cdot \sigma_k^3$$

Assuming independence of the delay distributions, the third moment of the UNI-UNI delay is just the sum of the Ethernet Network Section third moments.

$$\omega = \omega_1 + \omega_2 + \omega_3 + ... = \sum_{k=1}^{n} \omega_k$$

The UNI-UNI skewness is computed by dividing by $\sigma^3$ as shown below.

$$\gamma = \omega / \sigma^3$$

**Step 3**

The estimate of the 99.9-th percentile ( $p = 0.999$ ) of UNI-UNI delay, $\bar{d}_{TdS}$ or $\bar{d}_{TRS}$ is represented by $t$ as follows:.

$$t = \bar{\mu}_{TS} + \sigma \cdot \left\{ x_P - \frac{\gamma}{6}\left(1 - x_P{}^2\right) \right\}$$

where $t$ represents $\bar{d}_{TdS}$ at $x_P$ based on $P_d/100$ for Frame Delay or where $t$ represents $\bar{d}_{TRS}$ at $x_P$ based on $P_r/100$ for Frame Delay Range.

### 8.3.4 Ethernet Network Section Recommendations

Below are recommendations for how to apply the concatenation methods in section 8.3.

- Suggest that the choice of MFD and/or FD metrics be the same for each OVC that comprises the EVC and the same for the EVC CPOs.
- Suggest that the FDR Performance be used for each OVC that comprises the EVC and for the EVC CPOs.

- Suggest that the choice of Parameter values for the Performance metrics from Table 5 be the same for each OVC that comprises the EVC and the same for the EVC.
- Suggest that the boundaries for the SLS time interval T be aligned for each OVC that comprises the EVC and the same for the EVC CPOs.

## 8.4 Key Applications to Derive Performance Requirements

The intent of the CoS IA is to provide sufficient CoS Labels and Performance Objectives to efficiently support the vast majority of well-known applications. Identification of the applications supported, quantification of CPOs, specification of associated parameters (e.g., *P*, *T*, etc) and mapping to CoS Labels is described in this section.

Application mapping is for the purpose of determining the quantitative Performance Objectives for each CoS Label. It is not intended to mandate how an Operator, Service Provider or Subscriber maps a particular instance of an application. For example, a Subscriber could map some VoIP for certain types of communication to CoS Label L and other VoIP to CoS Label H if desired. This IA will be constructed such that VoIP (of the high-quality type defined in this appendix) will be supported in the CoS Label it is mapped to if the Operator conforms with this IA for that CoS Label. The mapping that will be developed is for showing how the CoS Performance Objectives are derived and not meant to imply a requirement for application mapping in actual implementations.

Similar to Application mapping, L2CP needs to be mapped to CoS Labels. There may be different CoS Labels for different L2CP types. At a minimum, there is a need to specify a CoS Label that meets the L2CP application requirements.

The applications considered in the process of generating CPOs and mapping requirements to CoS Labels are shown in Table 12. The applications fall into three general user segments: Consumer, Business, and Mobile. The user segments are not mutually exclusive, and many applications are aligned with more than one segment.

| Application | Consumer | Business | Mobile |
|---|---|---|---|
| VoIP Data | X | X | X |
| Interactive Video (Video Conferencing) | X | X | ? |
| VoIP and Video Signaling | X | X | X |
| Web Browsing | X | X | X |
| IPTV Data Plane | X | X | ? |
| IPTV Control Plane | X | X | ? |
| Streaming Media | X | X | X |
| Interactive Gaming | X | | X |

| Application | Consumer | Business | Mobile |
|---|---|---|---|
| Best Effort | X | X | X |
| Circuit Emulation | | X | X |
| Telepresence | | X | |
| Remote Surgery (Video) | | X | |
| Remote Surgery (Control) | | X | |
| Telehealth (Hi-res image file transfer) | | X | |
| Email | X | X | X |
| Broadcast Engineering (Pro Video over IP) | | X | |
| CCTV | X | X | X |
| Financial/Trading | | X | |
| Database | | X | |
| Real Time Fax over IP | X | X | |
| Store and Forward Fax over IP | X | X | |
| SANs (Synchronous Replication) | | X | |
| SANs (Asynchronous Replication) | | X | |
| Wide Area File Services | | X | |
| Network Attached Storage | X | X | |
| Text Terminals (telnet, ssh) | | X | |
| Graphics Terminals (Thin Clients) | | X | |
| Point of Sale Transactions | | X | |
| E-Commerce (Secure transactions) | X | X | X |
| Mobile Backhaul System Requirements | | | X |

**Table 12: Application list**

### 8.4.1 Application-specific Performance Objectives

Each of the applications listed in Table 12 was researched to determine the performance requirements associated with the application and the corresponding application-specific Performance Objectives associated with MEN Performance metrics. The requirements for application performance are usually specified from end-to-end. Since the MEN of interest may only be a portion of the end-to-end network which can also include customer network segments and endpoint devices, allocation or budgeting of the objective is generally required as the application-specific Performance Objectives are quantified. In addition, application level requirements for zero loss frequently assume the use of a loss recovery mechanism such as TCP operating above the MEN.

Table 13 through Table 33 show the requirements compiled for each application. Each table comprises two or three general sections. The top section provides application-level requirements and supporting measurement parameters compiled directly from the available sources. The second section maps the application level requirements to application-specific Performance Objective values for each MEN Performance metric and applies the appropriate parameters to each metric. The third section (if present) provides supplementary information about the application.

Application requirements were compiled from a variety of public sources. The first and most desirable category for source references is standards-based. Where standards-based references are unavailable, industry-based Best Practices are used, as well as vendor-specific and product-specific information. The sources for all application requirements are provided in their respective tables.

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **One-way delay** | < 150 ms preferred<br>< 400 ms limit | G.1010,<br>TS 22.105 | Total mouth-to-ear, includes encoding, decoding, and all buffering in addition to network delays. |
| | | < 150 ms | TR-126 | |
| | **Delay variation** | < 1 ms | G.1010,<br>TS 22.105 | Total mouth-to-ear, achieved using de-jitter buffer in receiver. |
| | **Meas. Params.** | $T \approx 1$ minute<br>$P = 0.999$ | Y.1541<br>Y.1541 | Suggested value (section 5.3.2)<br>Table 1/Y.1541 |
| **Appl. Perform. Objectives** | **FLR** | < 3e-2 | G.1010,<br>TS 22.105 | Assumes use of a packet loss concealment algorithm to minimize effect of packet loss. |
| | **FD** | < 125 ms preferred<br>< 375 ms limit | See text | $P_d = 0.999$ |
| | **FDR** | < 50 ms | Y.1541 | $P_r = 0.999$ |
| | **MFD** | < 100 ms preferred<br>< 350 ms limit | See text | |
| | **IFDV** | < 40 ms | | $P_v = 0.999$ |
| **Info** | **Bit rates** | 4 to 64 kbps | G.1010 | |
| | **Frame sizes** | ≤ 200 bytes | | 200 bytes based on G.711 with 20 ms frames. Most other codecs result in equal or smaller frame sizes. |
| | **Availability** | ≥ 99.99% | TR-NWT-000418,<br>TA-NWT-000909 | Bellcore standard for the PSTN (quoted from TR-126). |

**Table 13: VoIP Parameters**

The values in Table 13 provide an example of how application level requirements are mapped to application-specific Performance Objectives. The preferred value for one way delay for VoIP is 150 ms. The scope of this parameter includes everything between the talker's mouth and the listener's ear – the microphone, analog-to-digital conversion, speech encoding, buffering and framing, network delays, dejitter buffering, decoding, digital-to-analog conversion, and the speaker which converts the decoded analog signal to sound waves. Of all these elements, only network delays are within the scope of the MEN.

Typical non-network delays are identified and summed with guidance from ITU-T Recommendations G.114 and Y.1541. Per G.114, the buffering and framing delays associated with a G.711 encoder with 20 ms voice frames is 20.125 ms. Using Table VII.2/Y.1541 in Appendix VII of Y.1541 for guidance, a dejitter buffer of 50 ms is assumed and half of that value (25 ms) is allocated as its contribution to mean delay. A total of 5 ms is used for the contributions of other processes and equipment, for a total non-network contribution of approximately 50 ms to mean delay. The resulting Mean Frame Delay that can be allocated to the MEN as a Performance metric is 100 ms.

Frame Delay is mapped using a similar process. In this case, all non-network sources of delay except for the dejitter buffer are subtracted from the application parameter. The dejitter buffer acts to "smooth out" the variation in received voice frames resulting from network jitter. As a result, frames that arrive at the receiver with minimum delay are held in the dejitter buffer for its maximum duration, and frames arriving at the receiver at the maximum end of the jitter range are forwarded immediately, with no added delay in the dejitter buffer. Since the non-network delays (not including the dejitter buffer) total approximately 25 ms, the preferred value of 150 ms for one way application delay maps to a Frame Delay (at $P_d = 0.999$, close to the maximum value) of approximately 125 ms.

Application level parameters are mapped to Performance Objectives in Table 14 through Table 33 using the process described in the above example. Where source data is available, recommended measurement parameter values are also provided.

Real-time and streaming applications typically make use of a dejitter buffer such as that described above in the VoIP example. For those applications, frames which do not arrive at the dejitter buffer within a delay window corresponding to the length of the buffer are likely to be discarded. As a result, there is an implicit relationship between the percentile valued parameters used to define maximum delay or jitter ($P_d$ for Frame Delay, $P_v$ for Inter Frame Delay Variation and $P_r$ for Frame Delay Range) and the Frame Loss Ratio for those types of applications, since frames which arrive too late to be accepted into the dejitter buffer are effectively lost to the application. The relationship is:

$$P_r \text{ (or } P_v \text{ or } P_d) = 1 - \text{FLR}.$$

For real-time and streaming applications in the tables below, the above relationship has been used to derive $P_r$ or $P_v$ if recommended values for the parameters are not directly available from the source documentation.

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **One-way delay** | < 150 ms preferred<br>< 400 ms limit | G.1010,<br>TS 22.105 | Total user-to-user, includes encoding, decoding, and all buffering in addition to network delays. |
| | **Delay variation** | < 1 ms | G.1010 | Total user-to-user, achieved using de-jitter buffer in receiver. |
| | **Meas. Params.** | $T \approx$ 1 minute<br>$P$ = 0.999 | Y.1541<br>Y.1541 | Suggested value (section 5.3.2)<br>Table 1/Y.1541 |
| **Appl. Perform. Objectives** | **FLR** | < 1e-2 | G.1010,<br>TS 22.105 | Assumes use of a packet loss concealment algorithm to minimize effect of packet loss. |
| | **FD** | < 125 ms preferred<br>< 325 ms limit | | $P_d$ = 0.999 |
| | **MFD** | < 100 ms preferred<br>< 350 ms limit | | Network and de-jitter delays similar to VoIP case<br>H.264 supports sub-frame encoding/decoding delays (20 ms used for conversion) |
| | **FDR** | < 50 ms | Y.1541 | $P_r$ = 0.999 |
| | **IFDV** | < 40 ms | | $P_v$ = 0.999 |
| **Info** | **A/V synch** | < 80 ms | G.1010 | |
| | | < 100 ms | TS 22.105 | |
| | **Bit rates** | 16 to 384 kbps | G.1010 | |
| | | 32 to 384 kbps | TS 22.105 | |
| | | Up to ≈ 2 Mbps | H.264 | Configurable to 2 Mbps in current applications |
| | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | | | Not specified |

**Table 14: Interactive Video Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **Latency** | < 200 ms | TR-126 | Further detail unspecified in source, interpreted as upper bound on network delay. |
| | **Jitter** | < 50 ms | TR-126 | |
| | **Packet Loss Rate** | < 5.26E-6 | TR-126 | End-to-end application layer objective. Minimum value from TR-126 Tables 12 and 13. Assumes no or minimal loss concealment (tolerable loss rates may be higher depending on degree and quality of STB loss concealment). |
| **Appl. Perform. Objectives** | **FLR** | < 1E-3 | Y.1541 Amd. 3 | Network objective assuming Application Layer Forward Error Correction (AL-FEC) sufficient to provide application layer packet loss rate objective. |
| | **FDR** | < 50 ms | Y.1541 | Assumes AL-FEC sufficient to provide application layer packet loss rate objective. $P_r$ = 0.999* |
| | **MFD** | < 100 ms | See Notes | Encoding delay not included. Allow 100 ms for de-jitter buffer, decoding and AL-FEC delays. |
| | **FD** | < 125 ms | | $P_d$ = 0.999* |
| | **IFDV** | < 40 ms | | $P_v$ = 0.999* |
| **Info** | **Bit rates (MPEG-2)** | 3 to 5 Mbps | TR-126 | From TR-126 Table 12 |
| | **Bit rates (MPEG-4)** | 1.75 to 3 Mbps | TR-126 | From TR-126 Table 13 |
| | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | ≥ 99.99% | TR-126 | |

*No direct reference for percentiles, but dejitter buffering is required

**Table 15: Standard Definition Video Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **Latency** | < 200 ms | TR-126 | Further detail unspecified in source, interpreted as upper bound on network delay. |
| | **Jitter** | < 50 ms | TR-126 | |
| | **Packet Loss Rate** | < 1.16E-6 | TR-126 | End-to-end application layer objective. Minimum value from TR-126 Tables 14 and 15. Assumes some loss concealment. |
| **Appl. Perform. Objectives** | **FLR** | < 1E-3 | Y.1541 Amd 3 | Network objective assuming AL-FEC sufficient to provide application layer packet loss rate objective. |
| | **FDR** | < 50 ms | Y.1541 | Assumes AL-FEC sufficient to provide application layer packet loss rate objective. $P_r = 0.999$* |
| | **MFD** | < 100 ms | See Notes | Encoding delay not included. Allow 100 ms for de-jitter buffer, decoding and AL-FEC delays. |
| | **FD** | < 125 ms | | $P_d = 0.999$* |
| | **IFDV** | < 40 ms | | $P_v = 0.999$* |
| **Info** | **Bit rates (MPEG-2)** | 15 to 18.1 Mbps | TR-126 | From TR-126 Table 14 |
| | **Bit rates (MPEG-4)** | 8 to 12 Mbps | TR-126 | From TR-126 Table 15 |
| | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | ≥ 99.99% | TR-126 | |

*No direct reference for percentiles, but dejitter buffering is required

**Table 16: High Definition Video Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **Delay** | < 10 s | G.1010, TS 22.105 | Further detail unspecified in source, interpreted as time from request to initiation of playout. |
| | **Delay Variation** | << 1 ms | G.1010 | Value specified in G.1010 for audio as parameter at ear (post de-jitter buffer). Unspecified for video. |
| **Appl. Perform. Objectives** | **FDR** | < 2 s | TS 22.105 | Transport path, implies a 2 s de-jitter buffer. $P_r$ values unspecified in source. |
| | **FLR** | < 1% | G.1010 | |
| | **MFD** | | | Not specified |
| | **FD** | | | Not specified |
| | **IFDV** | < 1.5 s | | $P_v = 0.99$* |
| **Info** | **Bit rates (audio)** | 16 to 128 kbps | G.1010 | |
| | | 5 to 128 kbps | TS 22.105 | |
| | **Bit rates (video)** | 16 to 384 kbps | G.1010 | |
| | | 20 to 384 kbps | TS 22.105 | |
| | | Up to 2+ Mbps | | Measured video playout rates |
| | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | | | Not specified |

*No direct reference for percentiles, but dejitter buffering is required

**Table 17: Internet Streaming Audio/Video Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **One way delay** | < 250 ms | G.1010, TS 22.105 | Telemetry/two-way control/command and control category. |
| | **IPTV control plane response** | < 200 ms | TR-126 | Set-top box (STB) command processing - time interval between the remote control action (button push) and GUI update.<br>May include middleware server processing time for some functions. |
| | **Channel change response** | < 2 s | TR-126 | Remote button to stable video on new channel. |
| | **Delay Variation** | N.A. | G.1010, TS 22.105 | |
| | **Loss** | 0 | G.1010, TS 22.105 | |
| **Appl. Perform. Objectives** | **FDR** | N.A. | G.1010, TS 22.105 | |
| | **FLR** | 1e-3 | G.1010, TS 22.105 | Assumes TCP or other loss recovery |
| | **MFD** | < 75 ms | | Uses STB command processing with middleware server processing as worst case.<br>Allocates 50 ms to combined STB/middleware server processing, 150 ms to round trip delay. |
| | **FD** | N.A. | | |
| | **IFDV** | N.A. | | |
| **Info** | **Bit rates** | < 1 kbps | G.1010 | |
| | | < 28.8 kbps | TS 22.105 | |
| | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | ≥ 99.99% | TR-126 | Same as VoIP and SD/HD Video data plane requirements. |

**Table 18: Interactive Transaction Data Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **One way delay** | < 200 ms | G.1010 | TR-126 refers to this value as "likely too high." |
| | | < 75 ms preferred | TS 22.105 | |
| | | < 50 ms objective | TR-126 | Includes application layer (game server and game client) and network layer delays. |
| | **Delay Variation** | N.A. | G.1010, TS 22.105 | |
| | | < 10 ms objective | TR-126 | |
| | Loss | 0 | G.1010 | |
| **Appl. Perform. Objectives** | **FDR** | < 10 ms objective | TR-126 | |
| | **MFD** | < 40 ms objective | | TR-126 does not provide typical client/server delays. 10 ms used as a strawman value for the combination. |
| | **FLR** | 1e-3 | G.1010 | Assumes TCP or other loss recovery |
| | **FD** | < 50 ms objective | | |
| | **IFDV** | < 8 ms objective | | |
| **Info** | **Data** | < 1 KB | G.1010, TR-126 | Data per transaction. |
| | **Bit rates** | < 60 kbps | TS 22.105 | |
| | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | | | Not specified |

**Table 19: Interactive Gaming Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **Web browsing response time** | < 2 s/page preferred<br>< 4 s/page acceptable | G.1010,<br>TR-126 | Multiple round trip delays for most web pages imply requirement for MFD of less than 100 ms to meet 4 s response time.<br>Typical page size of ≈10 kbytes specified. Current page sizes range from ≈20 kbytes to >1 Mbyte. |
| | | < 4 s/page | TS 22.105 | Multiple round trip delays for most web pages imply requirement for MFD of less than 100 ms to meet 4 s response time. |
| | **Transaction services (e.g., e-commerce)** | < 2 s preferred<br>< 4 s acceptable | G.1010 | Multiple round trip delays for most web pages imply requirement for MFD of less than 100 ms to meet 4 s response time. |
| | | < 4 s | TS 22.105 | Multiple round trip delays for most web pages imply requirement for MFD of less than 100 ms to meet 4 s response time. |
| **Appl. Perform. Objectives** | **FDR** | N.A. | G.1010,<br>TR-126,<br>TS 22.105 | |
| | **FLR** | N.A. | G.1010,<br>TS 22.105 | |
| | **MFD** | N.A. | | Not specified |
| | **FD** | N.A. | | |
| | **IFDV** | N.A. | | |
| **Info** | **Frame sizes** | ≤ 1500 bytes | | |
| | **Availability** | | | Not specified |

**Table 20: Best Effort Parameters**

| Category | Param. | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **FD** | 25 ms | MEF 3 | $P_d$ = 99.9999% |
| | **Packet loss** | 1e-5 to 1e-7 | MEF 3 | Dependent on TDM service |
| | **Jitter** | 10 ms | MEF 3 | $P$ = 99.9999% |
| **Appl. Perform. Objectives** | **FLR** | 1E-6 | | |
| | **FDR** | 15 ms | Inferred from IFDV | $P_r$ = 99.9% |
| | **MFD** | 20 ms | Inferred from FD, IFDV | |
| | **IFDV** | 10 ms | MEF 8 | $P_v$ = 99.9%, $\Delta t$ = 900s, $T$ = 3600s |
| | **FD** | 25 ms | MEF 3 | $P_d$ = 99.9999% |
| | | | | |

**Table 21: Circuit Emulation Parameters**

Circuit Emulation is further defined in [9].

| Category | Param. | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **Delay** | < 2 s preferred<br>< 4 s acceptable | G.1010 | Transaction services |
| | **Packet loss** | 0 | G.1010 | Transaction services<br>Application level requirement |
| | **Jitter** | N.A. | G.1010 | Transaction services |
| **Appl. Perform. Objectives** | **FLR** | 1e-3 | Y.1541 Class 3 | |
| | **FDR** | Not specified | | |
| | **MFD** | 1 s | | |
| | **IFDV** | Not specified | | |
| | **FD** | 2 s | | |

**Table 22: Point of Sale Transaction Parameters**

| Category | Param. | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | 10 ms | IBM/Cisco SAN Multiprotocol Routing IBM Redbook SG24-7543-01 | Round trip Includes jitter |
| | | 5 ms | EMC SRDF Connectivity Guide | Best practice |
| | | 15 ms | IBM/Brocade SAN Multiprotocol Routing IBM Redbook SG24-7544-01 | Referring to iSCSI implementation |
| | **Packet loss** | 0.1% limit 0.01% rec. | EMC SRDF Connectivity Guide | Network requirement |
| | | 0.01% rec. | IBM SAN Multiprotocol Routing IBM Redbook SG24-7321-00 | Network requirement |
| | **Jitter** | 25% of latency or 25 ms | EMC SRDF Connectivity Guide | Use the lower value |
| **Appl. Perform. Objectives** | **FLR** | ≤ 1e-4 | | |
| | **FDR** | ≤ 1.25 ms | | 25% of 5 ms (one way) |
| | **MFD** | ≤ 3.75 ms | | 75% of 5 ms (one way) |
| | **IFDV** | ≤ 1 ms | | |
| | **FD** | ≤ 5 ms | | |

**Table 23: Synchronous Replication Parameters**

| Category | Parameter | Value | Source | Notes |
| --- | --- | --- | --- | --- |
| **Appl. Req's.** | **RTT** | 80 ms | IBM SAN Volume Controller Configuration Guide<br>IBM Redbook SC23-6628-02 | Round trip, Includes jitter<br>SVC version 4.1.1 or higher |
| | | 200 ms | EMC SRDF Connectivity Guide | Round trip |
| | **Packet loss** | 1% limit<br>0.01% rec. | EMC SRDF Connectivity Guide | Network requirement |
| | | 0.01% rec. | IBM SAN Multiprotocol Routing<br>IBM Redbook SG24-7321-00 | Network requirement |
| | **Jitter** | 25% of latency or 25 ms | EMC SRDF Connectivity Guide | Use the lower value |
| **Appl. Perform. Objectives** | **FLR** | ≤ 1e-4 | | |
| | **FD** | ≤ 40 ms | | |
| | **MFD** | ≤ 30 ms | | 75% of 40 ms (one way) |
| | **FDR** | ≤ 10 ms | | 25% of 40 ms (one way) |
| | **IFDV** | ≤ 8 ms | | |

**Table 24: Asynchronous Replication Parameters**

| Category | Parameter | Value | Source | Notes |
| --- | --- | --- | --- | --- |
| **Appl. Req's.** | **Delay** | 15 s preferred<br>60 s acceptable | G.1010 bulk data | Time for entire file to transfer |
| | **Packet loss** | 0 | G.1010 bulk data | Application level requirement |
| | **Jitter** | N.A. | G.1010 bulk data | |
| **Appl. Perform. Objectives** | **FLR** | ≤ 1e-3 | Y.1541 Class 4 | Assumes reliable delivery protocol (e.g., TCP) |
| | **FDR** | Unspecified | | |
| | **MFD** | ≤ 1 s | Y.1541 Class 4 | |
| | **IFDV** | Unspecified | | |
| | **FD** | Unspecified | | |

**Table 25: Network Attached Storage Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **One way delay** | < 200 ms | G.1010 | |
| | **Packet loss** | 0 | G.1010 | At application layer |
| | **Jitter** | N.A. | G.1010 | |
| **Appl. Perform. Objectives** | **FLR** | 1e-3 | Y.1541 Class 3 | Assumes TCP |
| | **FDR** | Unspecified | | |
| | **MFD** | < 200 ms | | |
| | **IFDV** | Unspecified | | |
| | **FD** | Unspecified | | |

**Table 26: Text and Graphics Terminal Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **One-way delay** | < 400 ms | G.1010 | VoIP "acceptable" value |
| | **Delay variation** | < 1 ms | G.1010, TS 22.105 | Achieved using de-jitter buffer in T.38 gateway |
| **Appl. Perform. Objectives** | **FLR** | < 3e-2 | G.1010, TS 22.105 | RTP, UDPTL, TCP all provide protection against frame loss |
| | **FDR** | < 50 ms | Y.1541 | $P_r$ = 0.999 |
| | **MFD** | < 350 ms | | From VoIP "acceptable" value |
| | **IFDV** | < 40 ms | Y.1541 | $P_v$ = 0.999 |
| | **FD** | < 400 ms | Y.1541 | From VoIP "acceptable" value $P_d$ = 0.999 |

**Table 27: T.38 Fax Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | < 3 ms | IBM System Storage Business Continuity Planning Guide | Synchronous copy / replication |
| | | ≤ 10 ms | Oracle Configuration Best Practices | Synchronous multiple log writer (LGWR) process |
| | | ≤ 12 ms | Oracle9i Data Guard Best Practice | Physical standby database distance |
| | | ≤ 100 ms | Active/Active clusters in SQL Server | Server Clustering |
| | **Packet loss** | 0 | G.1010 | Transaction Service |
| | **Jitter** | N.A | G.1010 Transaction services | |
| **Appl. Perform. Objectives** | **FLR** | 1e-5 | Y.1541 TCP Performance | |
| | **FD** | ≤ 5 ms | | |
| | **MFD** | N/S | | |
| | **FDR** | N/S | | |
| | **IFDV** | N/S | | |

**Table 28: Database Parameters – Hot Standby**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | ≤ 100 ms | Oracle9i Data Guard: Primary Site and Network Cfg BP | Asynchronous LGWR process |
| | | 100 ms | Active/Active clusters in SQL Server | Server Clustering |
| | **Packet loss** | 1e-5 | Y.1541 TCP Performance | |
| | **Jitter** | N.A | G.1010 Transaction services | |
| **Appl. Perform. Objectives** | **FLR** | $10^{-5}$ | Y.1541 TCP Performance | |
| | **FD** | ≤ 50 ms | | |
| | **MFD** | N/S | | |
| | **FDR** | N/S | | |
| | **IFDV** | N/S | | |

**Table 29: Database Parameters – WAN Replication**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | ≤ 300 ms | Oracle On Demand Reference Guide | End user to Oracle hosted servers |
| | | ≤ 2 s | G.1010 Transaction services | Preferred < 2 s Acceptable < 4 s |
| | | ≤ 7 s | Zona Research | eCommerce threshold (abandon rate) |
| | **Packet loss** | ≤ 0.1% | Oracle On Demand Reference Guide | End user to Oracle hosted servers |
| | | zero | G.1010 Transaction services | |
| | **Jitter** | N.A. | G.1010 Transaction services | |
| **Appl. Perform. Objectives** | **FLR** | 1e-3 | Y.1541 Class 3 (Transaction data, interactive) | Assumes TCP |
| | **FD** | N/S | | |
| | **MFD** | ≤ 1 s | G.1010 Transaction services | |
| | **FDR** | N/S | | |
| | **IFDV** | N/S | | |

**Table 30: Database Parameters – Client/Server**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | ≤ 1 s | SEC Regulation NMS Self-Help | |
| | | < 1 s | SEC Regulation NMS Intermarket Sweep Order Workflow | |
| | **Packet loss** | Extremely low | Cisco Trading Floor Architecture | |
| | **Jitter** | N/S | | |
| **Appl. Perform. Objectives** | **FLR** | 1e-5 | | |
| | **FD** | N/S | | |
| | **MFD** | ≤ 2 ms | | |
| | **FDR** | N/S | | |
| | **IFDV** | N/S | | |
| **Info** | **Availability** | 99.999% | | Various sources |

**Table 31: Financial Trading Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | ≤ 500 ms | Various, use cases | Based on 250ms (one way) PTZ requirement |
| | | ≤ 80 ms | Cisco Video Surveillance Best Practice | between client viewing station and VSOM |
| | **Packet loss** | ≤ 0.01% | MPEGIF | Based on MPEG-4 with Simple Profile |
| | **Jitter** | < 1 ms | G.1010 | Total user-to-user, achieved using de-jitter buffer in receiver. |
| **Appl. Perform. Objectives** | **FLR** | < 1e-2 | G.1010, TS 22.105 | Assumes use of a packet loss concealment algorithm to minimize effect of packet loss. |
| | **FD** | ≤ 150 ms (MPEG-4) ≤ 200 ms (MJPEG) | | Based on 250ms for PTZ, leaves 100ms for MPEG-4 encoding / decoding, 50ms for MJPEG encoding / decoding |
| | **MFD** | N/S | | |
| | **FDR** | 50 ms | Y.1541 | $P_r = 0.999$ |
| | **IFDV** | N/S | | |
| **Info** | **Availability** | | | |

**Table 32: CCTV Parameters**

| Category | Parameter | Value | Source | Notes |
|---|---|---|---|---|
| **Appl. Req's.** | **RTT** | ≤ 300 ms | Cisco TelePresence (1) | 240 ms Service Provider budget |
| | | ≤ 300 ms | Polycom (2) | Video endpoints and multipoint server delay is in addition |
| | **Packet loss** | ≤ 0.05% | Cisco TelePresence (1) | 0.025% Service Provider budget |
| | | ≤ 0.1% | Polycom (2) | Average over 5-minute interval |
| | **Jitter** | ≤ 10 ms | Cisco TelePresence (1) | |
| | | ≤ 40 ms | Polycom (2) | |
| **Appl. Perform. Objectives** | **FD** | ≤ 120 ms | Cisco TelePresence (1) | $P_d$ = 0.999 |
| | **MFD** | ≤ 110 ms | Cisco TelePresence (1)<br>Polycom (2) | = 120 – 10 ms<br>= 150 – 40 ms |
| | **FLR** | ≤ 0.025% | Cisco TelePresence (1)<br>Service Provider budget | |
| | **FDR** | ≤ 40 ms | Polycom jitter | $P_r$ = 0.999 |
| | **IFDV** | ≤ 10 ms | Cisco TelePresence (1) | $P_v$ = 0.9999 |
| | **Bandwidth** | 15 Mbps | Cisco TelePresence (1) | |

**Table 33: Telepresence Parameters**

The CPO ranges proposed relative to Mobile Backhaul are listed separately in Table 34. These CPO ranges map values associated with H, M, and L required classes as developed jointly between the CoS and Mobile Backhaul projects. Note that the driver for the requirements in the CoS Label H are often based on MBH for the older Mobile technologies (2G and 3G). For example, due to the tight control/signaling requirements when Ethernet MBH is inserted in the 3G UMTS RAN between the NodeB and the RNC (e.g., soft handover).

| **CoS Label** | **Example CoS Performance Objectives for each Metric#** | | | | | |
|---|---|---|---|---|---|---|
| | **MFD*** | **FD*** | **FDR** | **IFDV** | **FLR** | **Availability^** |
| **H** | 7 ms | 10 ms | 5 ms | 3 ms | $10^{-4}$ | TBD |
| **M** | 13 ms | 20 ms | 10 ms | 8 ms | $10^{-4}$ | TBD |
| **L** | 28 ms | 37 ms | N/S | N/S | $10^{-3}$ | TBD |

Notes:

Values are not recommendations for or reflections of actual services from contributing companies but rather represent reasonable industry values based on a wide range of MBH requirement sources, wide variety of applications, on any possible 2G-4G technologies. Less stringent values could be used for certain technologies or under certain mix of services/applications or network assumptions. Values will evolve (to more or less stringent values) as technologies mature and relational constraints between attributes are better understood and applied, and when SP field experiences will be available**.** SPs are free to provide CPOs that are more stringent for their specific services based on their field experience**.**

# Per MEF 10.2, Objectives in this table will not include periods declared Unavailable per the evolving MEF Availability attribute. Additional transient outage attributes may also be exclusions if adopted in the future (e.g., Consecutive High Loss Interval).

* MFD and FD Objectives assume geographic area/scope of limited size/distance (i.e., a Metro Performance Tier)

^ Availability metric is added as a Placeholder for MBH Phase 2 and CoS IA Phase 2. Values are TBD in future phase.

**Table 34: Mobile Backhaul Proposed CPOs**

All of the applications and their respective Performance Objectives are summarized in Table 35. Not all applications from the list in Table 12 are represented in Table 35. The remote control aspects of remote surgery and the IP-based transport of professional video were applications for which no clear guidance was found.

| Application | FD | MFD | FLR | FDR | IFDV |
|---|---|---|---|---|---|
| VoIP Data | 125 ms pref<br>375 ms limit<br>$P_d$ = 0.999 | 100 ms pref<br>350 ms limit | 3e-2 | 50 ms<br>$P_r$ = 0.999 | 40 ms<br>$P_v$ = 0.999 |
| Video Conferencing Data | 125 ms pref<br>375 ms limit<br>$P_d$ = 0.999 | 100 ms pref<br>350 ms limit | 1e-2 | 50 ms<br>$P_r$ = 0.999 | 40 ms<br>$P_v$ = 0.999 |
| VoIP and Videoconf Signaling | Not specified | 250 ms pref | 1e-3 | Not specified | Not specified |
| IPTV Data Plane | 125 ms<br>$P_d$ = 0.999 | 100 ms | 1e-3 | 50 ms<br>$P_r$ = 0.999 | 40 ms<br>$P_v$ = 0.999 |
| IPTV Control Plane | Not specified | 75 ms | 1e-3 | Not specified | Not specified |
| Streaming Media | Not specified | Not specified | 1e-2 | 2 s | 1.5 s<br>$P_v$ = 0.99 |
| Interactive Gaming | 50 ms | 40 ms | 1e-3 | 10 ms | 8 ms |
| Circuit Emulation | 25 ms<br>$P_d$ = .999999 | 20 ms | 1e-6 | 15 ms<br>$P_r$ = .999 | 10 ms<br>$P_v$ = .999,<br>$\Delta t$ = 900s,<br>$T$ = 3600s |
| Telepresence, includes: Remote Surgery (Video) | 120 ms<br>$P_d$ = 0.999 | 110 ms | 2.5e-4 | 40 ms<br>$P_r$ = 0.999 | 10 ms |
| Financial/Trading | Unknown | 2 ms | 1e-5 | Unknown | Unknown |
| CCTV | 150 ms (MPEG-4)<br>200 ms (MJPEG)<br>$P_d$=0.999 | Not specified | 1e-2 | 50 ms<br>$P_r$ = 0.999 | Not specified |
| Database (Hot Standby) | 5 ms | Not specified | 1e-5 | Unknown | Unknown |
| Database (WAN Replication) | 50 ms | Not specified | 1e-5 | Unknown | Unknown |
| Database (Client/Server) | Not specified | 1 s | 1e-3 | Not specified | Not specified |

| Application | FD | MFD | FLR | FDR | IFDV |
|---|---|---|---|---|---|
| T.38 Fax | 400 ms<br>$P_d = 0.999$ | 350 ms | 3e-2 | 50 ms<br>$P_r = 0.999$ | 40 ms<br>$P_v = 0.999$ |
| SANs (Synchronous Replication) | 5 ms | 3.75 ms | 1e-4 | 1.25 ms | 1 ms |
| SANs (Asynchronous Replication)* | 40 ms | 30 ms | 1e-4 | 10 ms | 8 ms |
| Network Attached Storage | Not specified | 1 s | 1e-3 | Not specified | Not specified |
| Text and Graphics Terminals | Not specified | 200 ms | 1e-3 | Not specified | Not specified |
| Point of Sale Transactions | 2 s | 1 s | 1e-3 | Not specified | Not specified |
| Best Effort, includes:<br>Email<br>Store/Forward Fax<br>WAFS<br>Web Browsing<br>File Transfer (including hi-res image file transfer)<br>E-Commerce | Not specified | Not specified | Not specified | Not specified | Not specified |
| Mobile Backhaul H | 10 ms | 7 ms | 1e-4 | 5 ms | 3 ms |
| Mobile Backhaul M | 20 ms | 13 ms | 1e-4 | 10 ms | 8 ms |
| Mobile Backhaul L | 37 ms | 28 ms | 1e-3 | Not specified | Not specified |

**Table 35: Summarized CPOs**

### 8.4.2 Derivation of CPOs from Application Performance Requirements

The values for CoS Performance Objectives (CPOs) are derived using multiple criteria. First, the set of applications described in section 8.4.1 is mapped into CoS Labels and Performance Tiers to determine the set of application-specific Performance Objectives applicable for each case. Candidate CPO values are determined which meet the Performance Objectives for most or all of the applications mapped into a CoS Label/Performance Tier combination. Ideally, all of the application-specific Performance Objectives will be satisfied for each application mapped into a specific CoS Label/Performance Tier combination – however, given the limited number of CoS Labels in the 3-CoS Label model and the breadth of the applications considered, this is not always possible.

Second, a set of statistical and other constraints are applied to the candidate CPO values to make sure that they maintain the correct relationships to each other across CoS Labels, across Performance Tiers, and between the CPOs within a single CoS Label/Performance Tier. The candidate CPO values are modified as necessary to meet the constraints while still satisfying the application-specific Performance Objectives.

#### 8.4.2.1 Mapping Applications to CoS Labels and Performance Tiers

Table 36 below is a table representing the explicit mapping of the applications in the tables in Section 8.4.1 above to the MEF 3 CoS Label Model. This mapping is informative for the purpose of derivation of CPOs, and does not constrain any mapping of actual applications to CoS Labels or Performance Tiers by Subscribers or Operators.

| **CoS Label** | **H** | | | | **M** | | | | **L** | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **Performance Tier** | **1** | **2** | **3** | **4** | **1** | **2** | **3** | **4** | **1** | **2** | **3** | **4** |
| VoIP | X | X | X | X | | | | | | | | |
| VoIP & videoconf signaling | | | | | X | X | X | X | | | | |
| Videoconf data | | | | | X | X | X | X | | | | |
| IPTV data | | | | | X | X | X | | | | | |
| IPTV control | | | | | X | X | X | | | | | |
| Streaming media | | | | | | | | | X | X | X | X |
| Interactive gaming | X | X | | | X | X | | | | | | |
| SANs synch replication | | | | | X | | | | | | | |
| SANs asynch replication | | | | | X | | | | | | | |
| Network attached storage | | | | | | | | | X | X | X | X |
| Text & graphics terminals | | | | | | | | | X | X | X | X |
| T.38 fax over IP | | | | | X | X | X | X | | | | |
| Database hot standby | | | | | X | | | | | | | |
| Database WAN replication | | | | | X | | | | | | | |
| Database client/server | | | | | | | | | X | X | X | X |
| Financial/Trading | X | | | | | | | | | | | |
| CCTV | | | | | X | X | X | X | | | | |
| Telepresence | X | X | X | | | | | | | | | |
| Circuit Emulation | X | | | | | | | | | | | |
| Mobile BH H | X | | | | | | | | | | | |
| Mobile BH M | | | | | X | | | | | | | |
| Mobile BH L | | | | | | | | | X | | | |

**Table 36: Explicit Application Mapping for Derivation of CPOs**

#### 8.4.2.2 Constraints on CPO Values

The set of CPOs for each class in a given tier is derived initially from the objectives of one or more applications, subject to minimum FD/MFD values implied by the distance range of that tier.

The following constraints on CPOs are required to avoid a statistical contradiction:

- FDR > FD – MFD
- MFD < FD
- IFDV < FDR

Also, assuming that the distribution of delays has a long tail to the right:

- FD – MFD >> .5 FDR (.5 represents a symmetric distribution)

We also apply two constraints to ensure consistency between the values for FD and FDR and the estimated maximum Propagation Delay PD associated with each performance tier, calculated as described in Section 8.2. When the percentile parameter $P_d = P_r$, then the Minimum Delay (MinD) associated with a given CoS Label/Performance Tier can be calculated as MinD = FD – FDR. This value MinD should be no less than PD. MinD should also not be significantly higher than PD. The first constraint is satisfied by:

- FD – FDR ≥ PD.

The second constraint is satisfied if the CPO values meet either of two tests. The first test scales PD by a ratio and then compares it to MinD. The second test, which prevents the constraint from becoming too severe for very low propagation delays, adds a fixed offset to PD before comparing it to MinD. The second test is expressed as:

- (FD – FDR ≤ PD * 1.5) OR (FD – FDR ≤ PD + 20ms)

Finally, for PT constraints we assume that CPOs should never improve as tier number increases and that the MFD for each PT must exceed the estimated maximum propagation delay for the PT.

Below is a tabular summary of the various constraints that are applied to the Application driven performance objectives in order to derive CPOs.

| Statistical and Inter-CoS Label Constraints | Notes |
|---|---|
| H CoS Label CPOs ≤ all other CoS Label CPOs, except H FLR ≥ M FLR | For all in-scope metrics CPO (assumes Parameters are consistent across CoS Labels) |
| FD – MFD >> .5 FDR * | Where .5 represents a symmetric distribution |
| MFD < FD | |
| FDR > FD – MFD * | |
| IFDV < FDR | |
| FD – FDR ≥ PD | PD = estimated max Propagation Delay for a given PT |
| (FD – FDR ≤ PD * 1.5) OR | PD = estimated max Propagation Delay for a |

| Statistical and Inter-CoS Label Constraints | Notes |
|---|---|
| (FD – FDR ≤ PD + 20ms) | given PT |

- *Note: can be combined into various forms, e.g., MFD + .5 FDR << FD < MFD + FDR.

| PT Constraints | Notes |
|---|---|
| PTm CPO ≤ PTn CPO | Where m<n (assumes Parameters are consistent across PTs. Includes all in-scope CPOs.) |
| PT1 MFD > 2 ms | Estimated max Propagation Delay for PT1 |
| PT2 MFD > 8 ms | Estimated max Propagation Delay for PT2 |
| PT3 MFD > 44 ms | Estimated max Propagation Delay for PT3 |
| PT4 MFD > 172 ms | Estimated max Propagation Delay for PT4 |

| Standards and Other Constraints | Notes |
|---|---|
| MEF CPOs ≤ Y.1541 IP QoS Class Objectives<br>CoS Label H PT1-3 for ITU QoS Class 0, 2<br>CoS Label H PT4 for ITU QoS Class 1<br>CoS Label M PT1-4 for ITU QoS Class 3<br>CoS Label L PT1-4 for ITU QoS Class 4 | Includes MFD (IPTD) and FLR (IPLR). Where PT1, PT2, PT3 comparable to National and PT4 comparable to Global |
| PT1 (Metro) ≤ CPOs for MBH | Not including any synchronization-only driven objectives that could be developed. These are for future phase |
| CPOs and Parameters will be expressed as maximum or minimum values (not ranges) | |

**Table 37: CPO Derivation Constraints**

#### 8.4.2.3 The CoS Performance Objective Compliance Tool

The CoS Performance Objective Compliance Tool is a Microsoft Excel spreadsheet used to test candidate CPO values against the application-specific Performance Objectives and the constraints identified above. The tool comprises a worksheet for each Performance Tier as well as two summary worksheets. The first worksheet summarizes all CPO values in one table and displays whether they meet the constraint tests. The second summary worksheet shows how the CPO values compare to the mapped application-specific Performance Objectives.

##### 8.4.2.3.1 Performance Tier worksheets

There are a total of four Performance Tier worksheets, one for each PT. At the bottom left of the table for each tier is a set of proposed CPO values (MFD, FDR, FLR, FD, and IFDV) for each class (H, M, L) in the 3-CoS Label model. The tool checks the compliance of each set of class objectives against the Application Performance Metrics objectives contained in the upper part of the table; the result of the compliance checks is displayed to the right of the application objective values.

In its current form, the definition of compliance used in the tool is as follows.

1. Each CPO value is compared to the corresponding Application Objective (AO) value. If the CPO value is less stringent than the AO value, it is considered Not Compliant; otherwise, the CPO value is considered Compliant. Two types of compliance are defined: Loose and Tight. If the AO value is within a (configurable) range of the CPO value, it is considered Tight compliance; otherwise it is Loose compliance. As an example, if an AO for MFD is 50% higher (less stringent) than the corresponding CPO, it is considered Loose compliance. An Unspecified or Unknown application objective also results in Loose compliance.

2. The compliance results for the set of CPO values for a class as compared to an application's requirements are then combined as follows:

   a. If any CPO value is Not Compliant, the overall compliance of the class to that application's requirements is considered "Bad."

   b. If any CPO value for the class yields Loose compliance, the overall compliance of the CPOs to that application's requirements is considered "OK" (which may be interpreted as "overkill," i.e., the stringency of the CPO is greater than required by the application).

   c. Otherwise, the overall compliance of the CPOs for the class to that application's requirements is considered "Good."

The spreadsheet based tables below illustrate the derivation of CPOs per PT. The derivation was based on a visual basic macro incorporated in the spreadsheet to provide a best fit for the application objectives into the 3 CoS Labels. In addition the constraints above were applied. (Note that the figures below are illustrative of the process used to derive the CPOs, and that the specific values may not reflect the normative CPO values in this document.)

The CPOs for PT1 are primarily driven by the MBH application.

| | |
|---|---|
| (yellow) | =Unspecified application objective |
| (brown) | =Unknown application objective |

**Application Performance Attributes** / **MEF CPOs Compliance**

| Application Attributes | Application | Context | CIR-only? | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) | H | M | L |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Consumer Applications | VoIP | PE-PE* | FALSE | 100 | 50 | 3.E-02 | 125 | 40 | OK | OK | OK |
| | VoIP and Videoconf Signaling | PE-PE* | FALSE | 250 | | 1.E-03 | 250 | | OK | OK | OK |
| | Video Conferencing Data | PE-PE* | FALSE | 100 | 50 | 1.E-02 | 125 | 40 | OK | OK | OK |
| | IPTV data plane | PE-PE* | FALSE | 100 | 50 | 1.E-03 | 125 | 40 | OK | OK | OK |
| | IPTV control plane | PE-PE* | FALSE | 75 | | 1.E-03 | | | OK | OK | OK |
| | Streaming media | PE-PE* | FALSE | | 2000 | 1.E-02 | | 1500 | OK | OK | OK |
| | Interactive gaming | PE-PE* | FALSE | 40 | 10 | 1.E-03 | 50 | 8 | OK | OK | Bad |
| Business Applications | SANs (Synchronous Replication) | PE-PE* | FALSE | 3.75 | 1.25 | 1.E-04 | 5 | 1 | Bad | Bad | Bad |
| | SANs (Asynchronous Replication) | PE-PE* | FALSE | 30 | 10 | 1.E-04 | 40 | 8 | OK | OK | Bad |
| | Network Attached Storage | PE-PE* | FALSE | 1000 | | 1.E-03 | | | OK | OK | OK |
| | Text and Graphics Terminals | PE-PE* | FALSE | 200 | | 1.E-03 | | | OK | OK | OK |
| | T.38 Real-time Fax over IP | PE-PE* | FALSE | 350 | 50 | 3.E-02 | 400 | 40 | OK | OK | OK |
| | Database (Hot Standby) | PE-PE* | FALSE | | | 1.E-05 | 5 | | Bad | Bad | Bad |
| | Database (WAN Replication) | PE-PE* | FALSE | | | 1.E-05 | 50 | | Bad | Bad | Bad |
| | Database (Client-Server) | PE-PE* | FALSE | 1000 | | 1.E-03 | | | OK | OK | OK |
| | Financial/Trading | PE-PE* | FALSE | 2 | | 1.E-05 | | | Bad | Bad | Bad |
| | CCTV | PE-PE* | FALSE | | 50 | 1.E-02 | 150 | | OK | OK | OK |
| | Telepresence (includes Remote Surgery video) | PE-PE* | FALSE | 110 | 18 | 3.E-04 | 120 | 10 | OK | OK | Bad |
| | Circuit Emulation | PE-PE* | FALSE | 20 | 15 | 1.E-06 | 25 | 10 | Bad | Bad | Bad |
| MBH Applications | MBH H | PE-PE* | FALSE | 7 | 5 | 1.E-04 | 10 | 3 | Good | Bad | Bad |
| | MBH M | PE-PE* | FALSE | 13 | 10 | 1.E-04 | 20 | 8 | OK | Good | Bad |
| | MBH L | PE-PE* | FALSE | 28 | 16 | 1.E-03 | 37 | 14 | OK | OK | Good |

**MEF CPOs (PT1)**

| MEF CoS Parameter Objectives (CPOs) (PT1, e.g., Metro) | Description (MEF Example Suggested Applications) | MEF CoS | CIR-only | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) |
|---|---|---|---|---|---|---|---|---|
| | Sync, Voice, Near-RT | H | FALSE | 7 | 5 | 1.E-04 | 10 | 3 |
| | Control/Signaling, Data | M | FALSE | 13 | 10 | 1.E-04 | 20 | 8 |
| | Data, Background | L | FALSE | 28 | 16 | 1.E-03 | 37 | 14 |

**Statistical Constraints**

| | MFD<FD | IFDV<FDR | FD<MFD+FDR | FD>MFD+FDR/2 | (FD-FDR > CRD) AND ((FD-FDR < CRD+Offset) OR [Minimum Delay Test (aka "Bob Test" |
|---|---|---|---|---|---|
| H | Good | Good | Good | Good | Good |
| M | Good | Good | Good | Good | Good |
| L | Good | Good | Good | Good | Good |

**Non-Statistical Constraints**

As stringent as Y.1541

| | MFD/IPTD | FLR/IPLR |
|---|---|---|
| H | Good | Good |
| M | Good | Good |
| L | Good | Good |

As stringent as higher tiers

| | MFD | FDR | FLR | FD | IFDV |
|---|---|---|---|---|---|
| H | Good | Good | Good | Good | Good |
| M | Good | Good | Good | Good | Good |
| L | Good | Good | Good | Good | Good |

| | MFD | FDR | FLR | FD | IFDV |
|---|---|---|---|---|---|
| H<=M (FLR:H>=M) | Good | Good | Good | Good | Good |
| H<=L | Good | Good | Good | Good | Good |

MFD > Calculated route distance

| | MFD | Air D | CRD km | CRD ms |
|---|---|---|---|---|
| H | Good | 250 | 312.5 | 1.5625 |
| M | Good | | | |
| L | Good | | | |

As stringent as MBH (PT1 only)

| | MFD | FDR | FLR | FD | IFDV |
|---|---|---|---|---|---|
| H | Good | Good | Good | Good | Good |
| M | Good | Good | Good | Good | Good |
| L | Good | Good | Good | Good | Good |

The following chart illustrates derivation of PT2 objectives.

Yellow = Unspecified application objective
Brown = Unknown application objective

| Application Attributes | Application | Context | CIR-only? | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) | H | M | L |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | Application Performance Attributes | | | | | MEF CPOs Compliance | | |
| Consumer Applications | VoIP | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 3.E-02 | 1.E+02 | 4.E+01 | OK | Good | Bad |
| | VoIP and Videoconf Signaling | PE-PE* | FALSE | 3.E+02 | | 1.E-03 | 3.E+02 | | OK | OK | OK |
| | Video Conferencing Data | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 1.E-02 | 1.E+02 | 4.E+01 | OK | Good | Bad |
| | IPTV data plane | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 1.E-03 | 1.E+02 | 4.E+01 | OK | Good | Bad |
| | IPTV control plane | PE-PE* | FALSE | 8.E+01 | | 1.E-03 | | | OK | OK | Good |
| | Streaming media | PE-PE* | FALSE | | 2.E+03 | 1.E-02 | | 2.E+03 | OK | OK | OK |
| | Interactive gaming | PE-PE* | FALSE | 4.E+01 | 1.E+01 | 1.E-03 | 5.E+01 | 8.E+00 | OK | Bad | Bad |
| Business Applications | SANs (Synchronous Replication) | PE-PE* | FALSE | 4.E+00 | 1.E+00 | 1.E-04 | 5.E+00 | 1.E+00 | Bad | Bad | Bad |
| | SANs (Asynchronous Replication) | PE-PE* | FALSE | 3.E+01 | 1.E+01 | 1.E-04 | 4.E+01 | 8.E+00 | OK | Bad | Bad |
| | Network Attached Storage | PE-PE* | FALSE | 1.E+03 | | 1.E-03 | | | OK | OK | OK |
| | Text and Graphics Terminals | PE-PE* | FALSE | 2.E+02 | | 1.E-03 | | | OK | OK | OK |
| | T.38 Real-time Fax over IP | PE-PE* | FALSE | 4.E+02 | 5.E+01 | 3.E-02 | 4.E+02 | 4.E+01 | OK | Good | Bad |
| | Database (Hot Standby) | PE-PE* | FALSE | | | 1.E-05 | 5.E+00 | | Bad | Bad | Bad |
| | Database (WAN Replication) | PE-PE* | FALSE | | | 1.E-05 | 5.E+01 | | Bad | Bad | Bad |
| | Database (Client-Server) | PE-PE* | FALSE | 1.E+03 | | 1.E-03 | | | OK | OK | OK |
| | Financial/Trading | PE-PE* | FALSE | 2.E+00 | | 1.E-05 | | | Bad | Bad | Bad |
| | CCTV | PE-PE* | FALSE | | 5.E+01 | 1.E-02 | 2.E+02 | | OK | OK | Bad |
| | Telepresence (includes Remote Surgery video) | PE-PE* | FALSE | 1.E+02 | 2.E+01 | 3.E-04 | 1.E+02 | 1.E+01 | OK | Bad | Bad |
| | Circuit Emulation | PE-PE* | FALSE | 2.E+01 | 2.E+01 | 1.E-06 | 3.E+01 | 1.E+01 | Bad | Bad | Bad |
| MBH Applications | MBH H | PE-PE* | FALSE | 6.E+00 | 3.E+00 | 1.E-05 | 8.E+00 | 2.E+00 | Bad | Bad | Bad |
| | MBH M | PE-PE* | FALSE | 1.E+01 | 1.E+01 | 1.E-05 | 2.E+01 | 8.E+00 | Bad | Bad | Bad |
| | MBH L | PE-PE* | FALSE | 3.E+01 | 2.E+01 | 1.E-03 | 4.E+01 | 1.E+01 | OK | Bad | Bad |

**MEF CoS Parameter Objectives (CPOs)**
**(PT2, e.g., Regional)**

MEF CPOs (PT2)

| Description (MEF Example Suggested Applications) | MEF CoS | CIR-only | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) |
|---|---|---|---|---|---|---|---|
| Sync, Voice, Near-RT | H | FALSE | 18 | 10 | 1.E-04 | 25 | 8 |
| Control/Signaling, Data | M | FALSE | 30 | 50 | 1.E-04 | 75 | 40 |
| Data, Background | L | FALSE | 50 | 100 | 1.E-03 | 125 | 80 |

**Statistical Constraints**

| | CoS | MFD<FD | IFDV<FDR | FD<MFD+FDR | FD>MFD+FDR/2 | | (FD-FDR > CRD) AND ((FD-FDR < CRD+Offset) OR |
|---|---|---|---|---|---|---|---|
| | H | Good | Good | Good | Good | | Good |
| | M | Good | Good | Good | Good | | Good |
| | L | Good | Good | Good | Good | | Good |

[Minimum Delay Test (aka "Bob Test"

**Non-Statistical Constraints**

| | CoS | MFD/IPTD | | FLR/IPLR | | |
|---|---|---|---|---|---|---|
| **As stringent as Y.1541** | H | Good | | Good | | |
| | M | Good | | Good | | |
| | L | Good | | Good | | |

| | CoS | MFD | FDR | FLR | FD | IFDV |
|---|---|---|---|---|---|---|
| **As stringent as higher tiers** | H | Good | Good | Good | Good | Good |
| | M | Good | Good | Good | Good | Good |
| | L | Good | Good | Good | Good | Good |

| | MFD | FDR | FLR | FD | IFDV |
|---|---|---|---|---|---|
| **H<=M (FLR:H>=M)** | Good | Good | Good | Good | Good |
| **H<=L** | Good | Good | Good | Good | Good |

| | CoS | MFD | Air D | CRD km | CRD ms |
|---|---|---|---|---|---|
| **MFD > Calculated route distance** | H | Good | 1200 | 1500 | 7.5 |
| | M | Good | | | |
| | L | Good | | | |

Likewise, the following chart illustrates derivation of PT3 objectives.

=Unspecified application objective
=Unknown application objective

| Application Attributes | Application | Context | CIR-only? | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) | MEF CPOs Compliance H | M | L |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Consumer Applications | VoIP | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 3.E-02 | 1.E+02 | 4.E+01 | OK | OK | Bad |
| | VoIP and Videoconf Signaling | PE-PE* | FALSE | 3.E+02 | | 1.E-03 | 3.E+02 | | OK | OK | OK |
| | Video Conferencing Data | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 1.E-02 | 1.E+02 | 4.E+01 | OK | OK | Bad |
| | IPTV data plane | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 1.E-03 | 1.E+02 | 4.E+01 | OK | OK | Bad |
| | IPTV control plane | PE-PE* | FALSE | 8.E+01 | | 1.E-03 | | | OK | Bad | Bad |
| | Streaming media | PE-PE* | FALSE | | 2.E+03 | 1.E-02 | | 2.E+03 | OK | OK | OK |
| | Interactive gaming | PE-PE* | FALSE | 4.E+01 | 1.E+01 | 1.E-03 | 5.E+01 | 8.E+00 | Bad | Bad | Bad |
| Business Applications | SANs (Synchronous Replication) | PE-PE* | FALSE | 4.E+00 | 1.E+00 | 1.E-04 | 5.E+00 | 1.E+00 | Bad | Bad | Bad |
| | SANs (Asynchronous Replication) | PE-PE* | FALSE | 3.E+01 | 1.E+01 | 1.E-04 | 4.E+01 | 8.E+00 | Bad | Bad | Bad |
| | Network Attached Storage | PE-PE* | FALSE | 1.E+03 | | 1.E-03 | | | OK | OK | OK |
| | Text and Graphics Terminals | PE-PE* | FALSE | 2.E+02 | | 1.E-03 | | | OK | OK | OK |
| | T.38 Real-time Fax over IP | PE-PE* | FALSE | 4.E+02 | 5.E+01 | 3.E-02 | 4.E+02 | 4.E+01 | OK | OK | Bad |
| | Database (Hot Standby) | PE-PE* | FALSE | | | 1.E-05 | 5.E+00 | | Bad | Bad | Bad |
| | Database (WAN Replication) | PE-PE* | FALSE | | | 1.E-05 | 5.E+01 | | Bad | Bad | Bad |
| | Database (Client-Server) | PE-PE* | FALSE | 1.E+03 | | 1.E-03 | | | OK | OK | OK |
| | Financial/Trading | PE-PE* | FALSE | 2.E+00 | | 1.E-05 | | | Bad | Bad | Bad |
| | CCTV | PE-PE* | FALSE | | 5.E+01 | 1.E-02 | 2.E+02 | | OK | OK | Bad |
| | Telepresence (includes Remote Surgery video) | PE-PE* | FALSE | 1.E+02 | 2.E+01 | 3.E-04 | 1.E+02 | 1.E+01 | OK | Bad | Bad |
| | Circuit Emulation | PE-PE* | FALSE | 2.E+01 | 2.E+01 | 1.E-06 | 3.E+01 | 1.E+01 | Bad | Bad | Bad |
| MBH Applications | MBH H | PE-PE* | FALSE | 6.E+00 | 3.E+00 | 1.E-05 | 8.E+00 | 2.E+00 | Bad | Bad | Bad |
| | MBH M | PE-PE* | FALSE | 1.E+01 | 1.E+01 | 1.E-05 | 2.E+01 | 8.E+00 | Bad | Bad | Bad |
| | MBH L | PE-PE* | FALSE | 3.E+01 | 2.E+01 | 1.E-03 | 4.E+01 | 1.E+01 | Bad | Bad | Bad |

MEF CPOs (PT3)

| MEF CoS Parameter Objectives (CPOs) (PT3, e.g., National) | Description (MEF Example Suggested Applications) | MEF CoS | CIR-only | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) |
|---|---|---|---|---|---|---|---|---|
| | Sync, Voice, Near-RT | H | FALSE | 70 | 12 | 2.5E-04 | 77 | 10 |
| | Control/Signaling, Data | M | FALSE | 80 | 50 | 2.5E-04 | 115 | 40 |
| | Data, Background | L | FALSE | 125 | 165 | 1.0E-03 | 230 | 130 |

| Constraint | CoS | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Statistical Constraints** | | MFD<FD | IFDV<FDR | FD<MFD+FDR | FD>MFD+FDR/2 | | (FD-FDR > CRD) AND ((FD-FDR < CRD+Offset) OR | |
| | H | Good | Good | Good | Good | | Good | [Minimum Delay Test (aka "Bob Test" |
| | M | Good | Good | Good | Good | | Good | |
| | L | Good | Good | Good | Good | | Good | |
| **Non-Statistical Constraints** | | | MFD/IPTD | | FLR/IPLR | | | |
| As stringent as Y.1541 | H | | Good | | Good | | | |
| | M | | Good | | Good | | | |
| | L | | Good | | Good | | | |
| | | | MFD | FDR | FLR | FD | IFDV | |
| As stringent as higher tiers | H | | Good | Good | Good | Good | Good | |
| | M | | Good | Good | Good | Good | Good | |
| | L | | Good | Good | Good | Good | Good | |
| | | | MFD | FDR | FLR | FD | IFDV | |
| H<=M (FLR:H>=M) | | | Good | Good | Good | Good | Good | |
| H<=L | | | Good | Good | Good | Good | Good | |
| | | | MFD | Air D | CRD km | CRD ms | | |
| MFD > Calculated route distance | H | | Good | 7000 | 8750 | 43.75 | | |
| | M | | Good | | | | | |

Finally, the following chart illustrates the derivation of PT4 objectives.

(yellow) =Unspecified application objective
(brown) =Unknown application objective

| Application Attributes | Application | Context | CIR-only? | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) | MEF CPOs Compliance H | M | L |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | **Application Performance Attributes** | | | | | | | |
| Consumer Applications | VoIP | PE-PE* | FALSE | 4.E+02 | 5.E+01 | 3.E-02 | 4.E+02 | 4.E+01 | OK | OK | Bad |
| | VoIP and Videoconf Signaling | PE-PE* | FALSE | 3.E+02 | | 1.E-03 | 3.E+02 | | OK | OK | Bad |
| | Video Conferencing Data | PE-PE* | FALSE | 3.E+02 | 5.E+01 | 1.E-02 | 4.E+02 | 4.E+01 | OK | OK | Bad |
| | IPTV data plane | PE-PE* | FALSE | 1.E+02 | 5.E+01 | 1.E-03 | 1.E+02 | 4.E+01 | Bad | Bad | Bad |
| | IPTV control plane | PE-PE* | FALSE | 8.E+01 | | 1.E-03 | | | Bad | Bad | Bad |
| | Streaming media | PE-PE* | FALSE | | 2.E+03 | 1.E-02 | | 2.E+03 | OK | OK | OK |
| | Interactive gaming | PE-PE* | FALSE | 4.E+01 | 1.E+01 | 1.E-03 | 5.E+01 | 8.E+00 | Bad | Bad | Bad |
| Business Applications | SANs (Synchronous Replication) | PE-PE* | FALSE | 4.E+00 | 1.E+00 | 1.E-04 | 5.E+00 | 1.E+00 | Bad | Bad | Bad |
| | SANs (Asynchronous Replication) | PE-PE* | FALSE | 3.E+01 | 1.E+01 | 1.E-04 | 4.E+01 | 8.E+00 | Bad | Bad | Bad |
| | Network Attached Storage | PE-PE* | FALSE | 1.E+03 | | 1.E-03 | | | OK | OK | OK |
| | Text and Graphics Terminals | PE-PE* | FALSE | 2.E+02 | | 1.E-03 | | | OK | Bad | Bad |
| | T.38 Real-time Fax over IP | PE-PE* | FALSE | 4.E+02 | 5.E+01 | 3.E-02 | 4.E+02 | 4.E+01 | OK | OK | Bad |
| | Database (Hot Standby) | PE-PE* | FALSE | | | 1.E-05 | 5.E+00 | | Bad | Bad | Bad |
| | Database (WAN Replication) | PE-PE* | FALSE | | | 1.E-05 | 5.E+01 | | Bad | Bad | Bad |
| | Database (Client-Server) | PE-PE* | FALSE | 1.E+03 | | 1.E-03 | | | OK | OK | OK |
| | Financial/Trading | PE-PE* | FALSE | 2.E+00 | | 1.E-05 | | | Bad | Bad | Bad |
| | CCTV | PE-PE* | FALSE | | 5.E+01 | 1.E-02 | 2.E+02 | | Bad | Bad | Bad |
| | Telepresence (includes Remote Surgery video) | PE-PE* | FALSE | 1.E+02 | 2.E+01 | 3.E-04 | 1.E+02 | 1.E+01 | Bad | Bad | Bad |
| | Circuit Emulation | PE-PE* | FALSE | 2.E+01 | 2.E+01 | 1.E-06 | 3.E+01 | 1.E+01 | Bad | Bad | Bad |
| MBH Applications | MBH H | PE-PE* | FALSE | 6.E+00 | 3.E+00 | 1.E-05 | 8.E+00 | 2.E+00 | Bad | Bad | Bad |
| | MBH M | PE-PE* | FALSE | 1.E+01 | 1.E+01 | 1.E-05 | 2.E+01 | 8.E+00 | Bad | Bad | Bad |
| | MBH L | PE-PE* | FALSE | 3.E+01 | 2.E+01 | 1.E-03 | 4.E+01 | 1.E+01 | Bad | Bad | Bad |

**MEF CPOs (PT4)**

| MEF CoS Parameter Objectives (CPOs) (PT4, e.g., Global) | Description (MEF Example Suggested Applications) | MEF CoS | CIR-only | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) |
|---|---|---|---|---|---|---|---|---|
| | Sync, Voice, Near-RT | H | FALSE | 200 | 40 | 5.E-04 | 230 | 32 |
| | Control/Signaling, Data | M | FALSE | 220 | 50 | 5.E-04 | 250 | 40 |
| | Data, Background | L | FALSE | 240 | 200 | 1.E-03 | 390 | 160 |

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| **Statistical Constraints** | | MFD<FD | IFDV<FDR | FD<MFD+FDR | FD>MFD+FDR/2 | | (FD-FDR > CRD) AND ((FD-FDR < CRD+Offset) OR | |
| | H | Good | Good | Good | Good | | Good | [Minimum Delay Test (aka "Bob Test" |
| | M | Good | Good | Good | Good | | Good | |
| | L | Good | Good | Good | Good | | Good | |
| **Non-Statistical Constraints** | | | MFD/IPTD | | FLR/IPLR | | | |
| As stringent as Y.1541 | H | | Good | | Good | | | |
| | M | | Good | | Good | | | |
| | L | | Good | | Good | | | |
| | | | MFD | FDR | FLR | FD | IFDV | |
| As stringent as higher tiers | H | | NA | NA | NA | NA | NA | |
| | M | | NA | NA | NA | NA | NA | |
| | L | | NA | NA | NA | NA | NA | |
| | | | MFD | FDR | FLR | FD | IFDV | |
| H<=M (FLR:H>=M) | | | Good | Good | Good | Good | Good | |
| H<=L | | | Good | Good | Good | Good | Good | |
| | | | MFD | Air D | CRD km | CRD ms | | |
| MFD > Calculated route distance | H | | Good | 27500 | 34375 | 171.875 | | |
| | M | | Good | | | | | |
| | L | | Good | | | | | |
| | | | (FD-FDR > CRD) AND ((FD-FDR < CRD+Offset) OR (FD-FDR < CRD*Ratio)) | | | | | |
| Minimum Delay Test (aka "Bob test") | H | | Good | | | | | |
| | M | | Good | | | | | |
| | L | | Good | | | | | |

**Table 38: PT1-4 CPO Derivation and Evaluation Spreadsheets**

#### 8.4.2.3.2 CPO Summary worksheet

The CPO Summary worksheet displays numerical values for all CPOs (even for those CPOs defined as "Not Specified" in Table 6 through Table 9) and shows the results of the constraint tests applied to those CPO values. Figure 5 shows the summary displays.

| | MEF CoS | CIR-only | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) |
|---|---|---|---|---|---|---|---|
| | | | (See MEF 10.2, Section 6.9 for definitions) | | | | |
| PT1 | H | FALSE | 7 | 5 | 1.0E-04 | 10 | 3 |
| | M | FALSE | 13 | 10 | 1.0E-04 | 20 | 8 |
| | L | FALSE | 28 | 16 | 1.0E-03 | 37 | 14 |
| PT2 | H | FALSE | 18 | 10 | 1.0E-04 | 25 | 8 |
| | M | FALSE | 30 | 50 | 1.0E-04 | 75 | 40 |
| | L | FALSE | 50 | 100 | 1.0E-03 | 125 | 80 |
| PT3 | H | FALSE | 70 | 12 | 2.5E-04 | 77 | 10 |
| | M | FALSE | 80 | 50 | 2.5E-04 | 115 | 40 |
| | L | FALSE | 125 | 165 | 1.0E-03 | 230 | 130 |
| PT4 | H | FALSE | 200 | 40 | 5.0E-04 | 230 | 32 |
| | M | FALSE | 220 | 50 | 5.0E-04 | 250 | 40 |
| | L | FALSE | 240 | 200 | 1.0E-03 | 390 | 160 |

PT comparison

| | MEF CoS | MFD (ms) | FDR (ms) | FLR (ratio) | FD (ms) | IFDV (ms) | Minimum Delay Test | MinD = FD-FDR | Propagation Delay (ms) | Shaping delay budget factor | Implied values: Serialization Delay (ms) | Implied values: Queuing Delay + Shaping Delay budget (ms) | Implied values: Shaping Delay from budget (ms) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| PT1 | H | Good | Good | Good | Good | Good | Good | 5 | 2 | 0.50 | 3.2 | 4.8 | 2.4 |
| | M | Good | Good | Good | Good | Good | Good | 10 | 2 | 0.50 | 3.2 | 14.8 | 7.4 |
| | L | Good | Good | Good | Good | Good | Good | 21 | 2 | 0.50 | 3.2 | 31.8 | 15.9 |
| PT2 | H | Good | Good | Good | Good | Good | Good | 15 | 8 | 0.50 | 3.2 | 13.8 | 6.9 |
| | M | Good | Good | Good | Good | Good | Good | 25 | 8 | 0.50 | 3.2 | 63.8 | 31.9 |
| | L | Good | Good | Good | Good | Good | Good | 25 | 8 | 0.50 | 3.2 | 113.8 | 56.9 |
| PT3 | H | Good | Good | Good | Good | Good | Good | 65 | 44 | 0.50 | 3.2 | 29.8 | 14.9 |
| | M | Good | Good | Good | Good | Good | Good | 65 | 44 | 0.50 | 3.2 | 67.8 | 33.9 |
| | L | Good | Good | Good | Good | Good | Good | 65 | 44 | 0.50 | 3.2 | 182.8 | 91.4 |
| PT4 | H | Good | Good | Good | Good | Good | Good | 190 | 172 | 0.50 | 3.2 | 54.8 | 27.4 |
| | M | Good | Good | Good | Good | Good | Good | 200 | 172 | 0.50 | 3.2 | 74.8 | 37.4 |
| | L | Good | Good | Good | Good | Good | Good | 190 | 172 | 0.50 | 3.2 | 214.8 | 107.4 |

Minimum Delay Test (aka "Bob Test") = (MinD >= PD) AND ((MinD <= PD+Offset) OR (MinD <= PD*Ratio))
Offset 20
Ratio 1.5

**Figure 5 – CPO Summary worksheet**

#### 8.4.2.3.3 Application Mapping Summary Worksheet

The Application Mapping summary worksheet contains several tables. The lower table defines the explicit mapping of applications to CoS Label/Performance Tier combinations used to test the CPO values. An 'X' in a cell maps the application in the cell's row to the CoS Label/Performance Tier in the cell's column. The right side of the table includes a summary of the application-specific Performance Objectives for each application. The upper left table shows how well the mapped application-specific Performance Objectives match the CPO values, using the criteria described for the Performance Tier worksheets in Section 8.4.2.3.1 above. The upper right table provides a summary of how well the application-specific Performance Objectives match the CPO values for all applications, CoS Labels and Performance Tiers, both mapped and unmapped. Figure 6 shows the application mapping tables.

**Merge of actual and desired states**

| Category | Application | PT1 CoS H | PT1 CoS M | PT1 CoS L | PT2 CoS H | PT2 CoS M | PT2 CoS L | PT3 CoS H | PT3 CoS M | PT3 CoS L | PT4 CoS H | PT4 CoS M | PT4 CoS L |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Real time | VoIP | OK | | | OK | | | OK | | | OK | | |
| Interactive | VoIP and videoconf signaling | | OK | | | OK | | | OK | | | OK | |
| Real time | Videoconf data | | OK | | | Good | | | OK | | | OK | |
| Near real time | IPTV data | | OK | | | Good | | | OK | | | | |
| Interactive | IPTV control | | OK | | | OK | | | Bad | | | | |
| Streaming | Streaming media | | | OK | | | OK | | | OK | | | OK |
| Low delay | Interactive gaming | OK | OK | | OK | Bad | | | | | | | |
| Very low delay | SANs synchronous replication | | Bad | | | | | | | | | | |
| Low delay | SANs asynchronous replication | | OK | | | | | | | | | | |
| Best effort | Network attached storage | | | OK | | | OK | | | OK | | | OK |
| Best effort | Text and graphics terminals | | | OK | | | OK | | | OK | | | Bad |
| Near real time | T.38 fax over IP | | OK | | | Good | | | OK | | | OK | |
| Very low delay | Database hot standby | | Bad | | | | | | | | | | |
| Low delay | Database WAN replication | | OK | | | | | | | | | | |
| Best effort | Database client/server | | | OK | | | OK | | | OK | | | OK |
| Very low delay | Financial/Trading | Bad | | | | | | | | | | | |
| Near real time | CCTV | | OK | | | OK | | | OK | | | Bad | |
| Real time | Telepresence | OK | | | OK | | | OK | | | | | |
| Real time | Circuit Emulation | Bad | | | | | | | | | | | |
| Very low delay | MBH H | Good | | | | | | | | | | | |
| Very low delay | MBH M | | Good | | | | | | | | | | |
| Low delay | MBH L | | | Good | | | | | | | | | |

**Applications to CoS Levels (Current state)**

| PT1 CoS H | PT1 CoS M | PT1 CoS L | PT2 CoS H | PT2 CoS M | PT2 CoS L | PT3 CoS H | PT3 CoS M | PT3 CoS L | PT4 CoS H | PT4 CoS M | PT4 CoS L |
|---|---|---|---|---|---|---|---|---|---|---|---|
| OK | OK | OK | OK | Good | Bad | OK | OK | Bad | OK | OK | Bad |
| OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | Bad |
| OK | OK | OK | OK | Good | Bad | OK | OK | Bad | OK | OK | Bad |
| OK | OK | OK | OK | Good | Bad | OK | OK | Bad | Bad | Bad | Bad |
| OK | OK | OK | OK | OK | Good | OK | Bad | Bad | Bad | Bad | Bad |
| OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK |
| OK | OK | Bad | OK | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | OK | Bad | OK | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK |
| OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | Bad | Bad |
| OK | OK | OK | OK | Good | Bad | OK | OK | Bad | OK | OK | Bad |
| Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | OK | Bad | OK | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK | OK |
| Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | OK | OK | OK | OK | Bad | OK | OK | Bad | Bad | Bad | Bad |
| OK | OK | Bad | OK | Bad | Bad | OK | Bad | Bad | Bad | Bad | Bad |
| Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| Good | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | Good | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |
| OK | OK | Good | OK | Bad | Bad | Bad | Bad | Bad | Bad | Bad | Bad |

**Applications to CoS Levels (Desired state)**

| Category | Application | Very low jitter (<< 50 ms) CoS H PT1 | PT2 | PT3 | PT4 | Low jitter (50 ms) CoS M PT1 | PT2 | PT3 | PT4 | Non-critical data CoS L PT1 | PT2 | PT3 | PT4 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Real time | VoIP | X | X | X | X | | | | | | | | |
| Interactive | VoIP and videoconf signaling | | | | | X | X | X | X | | | | |
| Real time | Videoconf data | | | | | X | X | X | X | | | | |
| Near real time | IPTV data | | | | | X | X | X | | | | | |
| Interactive | IPTV control | | | | | X | X | X | | | | | |
| Streaming | Streaming media | | | | | | | | | X | X | X | X |
| Low delay | Interactive gaming | X | X | | | X | X | | | | | | |
| Very low delay | SANs synchronous replication | | | | | X | | | | | | | |
| Low delay | SANs asynchronous replication | | | | | X | | | | | | | |
| Best effort | Network attached storage | | | | | | | | | X | X | X | X |
| Best effort | Text and graphics terminals | | | | | | | | | X | X | X | X |
| Near real time | T.38 fax over IP | | | | | X | X | X | X | | | | |
| Very low delay | Database hot standby | | | | | X | | | | | | | |
| Low delay | Database WAN replication | | | | | X | | | | | | | |
| Best effort | Database client/server | | | | | | | | | X | X | X | X |
| Very low delay | Financial/Trading | X | | | | | | | | | | | |
| Near real time | CCTV | | | | | X | X | X | X | | | | |
| Real time | Telepresence | X | X | X | | | | | | | | | |
| Real time | Circuit Emulation | X | | | | | | | | | | | |
| Very low delay | MBH H | X | | | | | | | | | | | |
| Very low delay | MBH M | | | | | X | | | | | | | |
| Low delay | MBH L | | | | | | | | | X | | | |

**Performance Attributes**

| MFD | FDR | FLR | FD | IFDV |
|---|---|---|---|---|
| 100 | 50 | 0.03 | 125 | 40 |
| 250 | | 0.001 | 250 | |
| 100 | 50 | 0.01 | 125 | 40 |
| 100 | 50 | 0.001 | 125 | 40 |
| 75 | | 0.001 | | |
| | 2000 | 0.01 | | 1500 |
| 40 | 10 | 0.001 | 50 | 8 |
| 3.75 | 1.25 | 0.0001 | 5 | 1 |
| 30 | 10 | 0.0001 | 40 | 8 |
| 1000 | | 0.001 | | |
| 200 | | 0.001 | | |
| 350 | 50 | 0.03 | 400 | 40 |
| | | 1E-05 | 5 | |
| | | 1E-05 | 50 | |
| 1000 | | 0.001 | | |
| 2 | | 1E-05 | | |
| | 50 | 0.01 | 150 | |
| 110 | 18 | 0.0003 | 120 | 10 |
| 20 | 15 | 1E-06 | 25 | 10 |
| 6 | 3 | 1E-05 | 8 | 2 |
| 13 | 10 | 1E-05 | 20 | 8 |
| 28 | 16 | 0.001 | 37 | 14 |

=Unspecified application objective

=Unknown application objective

**Figure 6 – Application Mapping summary worksheet**

## 8.5 Example PCP and DSCP Mapping at UNI for Multi-CoS EVCs Supporting Only Standard MEF Classes of Service

The CoS IA requires that all PCP (or DSCP) values that may occur in any service deployment are to be supported in some way by the service. Several alternatives exist. For example, any specific MEN service may support additional CoS Names beyond those defined in this IA, and PCP (or DSCP) values not specified as CoS Identifiers in the CoS IA may be mapped to a CoS Name provided as an addition to the CoS IA defined CoS Labels. Alternatively, a service may include at least one additional CoS Name intended specifically to handle frames not associated (by PCP/DSCP value) with a defined CoS Identifier. If a specific MEN service supports *only* the CoS Labels defined by this IA, there needs to be a mapping of all possible PCP (or DSCP) values to one of the CoS Labels defined in the CoS IA or to a CoS defined in [2] called "Discard" which simply discards all frames that are classified as such.

This section provides example mappings for this case assuming no "Discard" CoS Name. Note that in some cases the use of a "Discard" CoS with only the PCP and DSCP values specified in Table 4 may be the simplest way to negotiate markings. In this case all PCP and DSCP values not shown in Table 4 would be discarded at the EI.

### 8.5.1 Example PCP Mappings

The following tables provide examples of full mappings of PCP at a UNI for multi-CoS Label EVCs that support only standard MEF CoS Labels.

Table 39 shows an example mapping in which PCP value 5 is assumed to be handled by CE routers as "EF" traffic. This may be a common approach in handling low latency traffic based on a PCP marking – particularly when using (for instance) IP Routers.

| MEF CoS Label Combination Supported on EVC | PCP Mapping per Class of Service Label - Color Blind Mode | | |
|---|---|---|---|
| | H | M | L |
| {H + M + L} | 5 | 2-4, 6, 7 | 0, 1 |
| {H + M} | 5 | 0-4, 6, 7 | N/A |
| {H + L} | 5 | N/A | 0-4, 6, 7 |
| {M + L} | N/A | 2-7 | 0, 1 |

**Table 39: Example PCP Mapping for Multi-CoS Label EVC Supporting Only Standard CoS Labels at UNI – "Router-Application-Friendly" mapping**

Table 40 shows a similar mapping that may apply in an application that bases choices of PCP values on the assumption of Ethernet CE bridges forwarding based on strict priority. In this case, higher PCP values would be handled at a higher priority. This mapping works in an application where very-high priority traffic is (by nature) very low volume (possibly less than 1 percent of the total traffic volume). This mapping is needed, for instance, if the application is not necessarily able to distinguish traffic that is carried natively in Ethernet over the local LAN from traffic that may be carried by a MEN service.

| MEF CoS Label Combination Supported on EVC | PCP Mapping per Class of Service Label - Color Blind Mode | | |
|---|---|---|---|
| | H | M | L |
| {H + M + L} | 4-7 | 2,3 | 0, 1 |
| {H + M} | 4-7 | 0-3 | N/A |
| {H + L} | 4-7 | N/A | 0-3 |

| | | | |
|---|---|---|---|
| {M + L} | N/A | 2-7 | 0, 1 |

**Table 40: Example PCP Mapping for Multi-CoS Label EVC Supporting Only Standard CoS Labels at UNI – "Bridging-Application-Friendly" mapping**

### 8.5.2 Example DSCP Mappings

The following table provides an example of a full mapping of DSCP values at a UNI for multi-CoS Label EVCs that support only standard MEF CoS Labels.

| MEF CoS Combination Supported on EVC | DSCP Mapping per Class of Service – Color Blind Mode | | |
|---|---|---|---|
| | H | M | L |
| {H + M + L} | 40-47 | 16-39, 48-63 | 0-15 |
| {H + M} | 40-47 | 0-39, 48-63 | N/A |
| {H + L} | 40-47 | N/A | 0-39, 48-63 |
| {M + L} | N/A | 16-63 | 0-15 |

**Table 41: Example DSCP Mapping for Multi-CoS EVC Supporting Only Standard Classes of Service at UNI**

## 8.6 Other Relevant Standards and Industry Models

This section excerpts information from relevant standards that may be helpful in reading this document.

Below are excerpted tables from Section 6 and Annex G (informative) of [5]. Specifically this IA used the 5P3D row PCP values (bottom row on the excerpt below) from Table 6-4 for the CoS Identifier PCP values in Table 4 because 5 Priorities (i.e., classes) is the closest match to the 3 CoS Label Model. There is no row in the table for a smaller number of Priorities than 5P3D. Note that in Table G-2 of [3] the VO (voice) class specifies 10ms latency and jitter.

| **PCP Allocation** | | **PCP Values and Traffic Classes** | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **# PCP Priorities** | **# PCP Drop Eligible** | **PCP = 7** | **6** | **5** | **4** | **3** | **2** | **1** | **0** |

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 8 | 0 | IEEE Traffic Class = 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| **5** | **3** | IEEE Traffic Class = 7 | 6 | 4 | 4 DE | 2 | 2 DE | 0 | 0 DE |

**Table 42: PCP Decoding (Adapted from [5])**

## 8.7 BURST SIZE AND SHAPER CONSIDERATIONS FOR ENNI

### 8.7.1 Burst Size and Burst Alignment

A Service Provider ought to ensure Operator alignment on Committed Burst Size across an ENNI in order to avoid exceeding frame loss objectives in boundary situations where loss performance is close to exceeding the loss objective. For example, consider the situation shown in Figure 7.

This example depicts a single-CoS (e.g., H CoS) point-to-point EVC stitched from two OVCs crossing an ENNI between the operator networks MEN- 1 and MEN-2. The CIR values for the Ingress Bandwidth Profile at UNI-1 for OVC-1 and for the Ingress Bandwidth Profile at the ENNI for OVC-2 are the same ignoring the different overheads for Service vs. ENNI Frames[3].

The CBS value for the Ingress Bandwidth Profile at the UNI-1 for OVC-1 is CBS=x. The CBS value for the Ingress Bandwidth Profile at the ENNI for OVC-2 is CBS=y. The burst sizes may differ between the OVCs ($x \neq y$) ignoring the different overheads for Service vs. ENNI Frames. In particular, if $x > y$ (perhaps due to inability of an Operator to customize the CBS value for a given OVC), then traffic flowing from UNI-1 to UNI-2 may experience frame loss due to policing at the ENNI Ingress Bandwidth Profile. This frame loss, when added to loss due to other factors, may cause FLR objectives to be exceeded.

[3] In this example, CIR and CBS is slightly higher for an ENNI if this ENNI will allow same user traffic load.

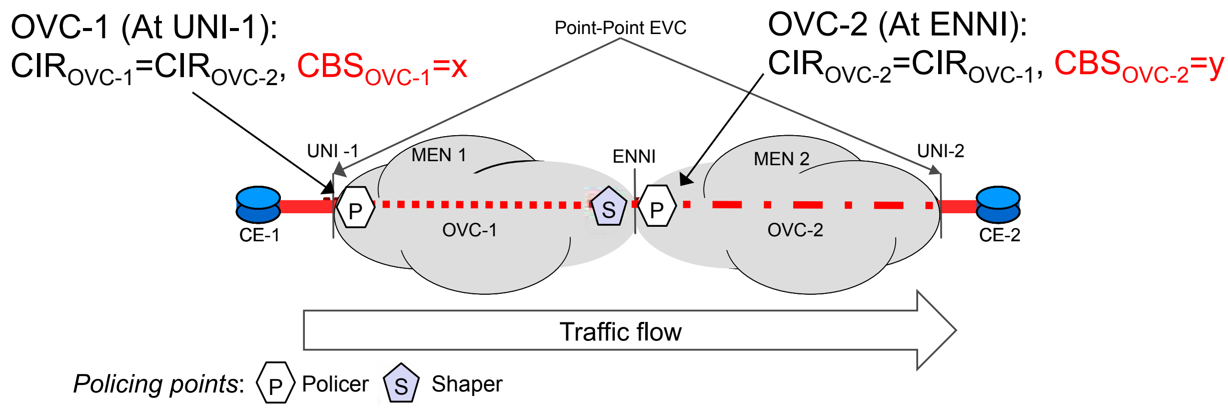


**Figure 7 – Burst Alignment Example with Policing Points for Traffic Traversing the ENNI**

If the Frame Loss Ratio is to be small, the Service Provider ought to ensure:

- Consistency between Operators on burst size on the respective OVCs (i.e., MEN-1 CBS x ≤ MEN-2 CBS y); or
- Shaping in one direction at the ENNI, in order to mitigate the difference between burst sizes between OVCs. Note that even if MEN-1 CBS and MEN-2 CBS are equal, the effects of frame delay variation may result in loss at the ENNI. Therefore, shaping may be needed even in the case of equal burst sizes (x=y).

The shaping option allows the burst size in MEN-2 to be less than that of MEN-1, within a certain range determined by the shaping parameters.[4] ENNI related shaping may occur either at the egress of MEN-1, or at the ingress of MEN-2.[5] Shaping may also occur at the CE. Only shaping at the egress of MEN-1 is addressed in this document as depicted in Figure 7. Care must be taken in the selection of shaping parameters in order not to violate delay requirements of the EVC in its Performance Tier, due to the added delay of the shaper.

For reference, example shaper algorithms for implementation at the egress of MEN-1 at the ENNI are given in Section 8.7.2.

Referring to Figure 7:

- Let *ΔCBS* be defined as the difference $CBS_{OVC-1} - CBS_{OVC-2}$, where $CBS_{OVC-1}$ is the CBS at UNI-1, and $CBS_{OVC-2}$ is the CBS at the ingress on the OVC-2 side of the ENNI;
- Let *CIR* be the CIR of OVC-1 at UNI-1 (assume $CIR_{OVC-1} = CIR_{OVC-2}$), in bits/sec;

[4] Compare with the recommendation for CE egress traffic shaping in Section 10.3 of MEF 10.2 [2].

[5] If shaping is performed at the ingress of MEN-2, ingress policing at the ENNI may be optional. However, shaping at the ingress of MEN-2 is out of scope for this document.

- Let $R_O$ be the effective[6] UNI information rate of the OVC (e.g., OVC-1), in bits/sec;
    - Where $R_O$ = (average frame size / (average frame size + 20)) * UNI line rate;
- Let $D_S$ be the upper bound on delay (maximum waiting time) at the shaper for a given Performance Tier and CoS combination, in sec, defined as $D_S = B_S * FDR$, where $B_S$ is the "Shaper Budget" factor (e.g., $B_S = 0.5$) for indicating the amount of the total Frame Delay Range (FDR) objective allocated to shaper delay versus other queuing delay; $B_S$ can vary by CoS and Performance Tier.
- The shaper in MEN-1 does not buffer any frames that were declared Yellow by the Ingress Bandwidth Profile at UNI-1.

In cases where $\Delta CBS$ is positive, a Service Provider ought to ensure alignment of burst parameters among the Operators across an ENNI by use of a shaper by the Operator of MEN-1 at its egress at the ENNI, configured such that $\Delta CBS$ is as follows.

At an ENNI where CBS is specified for a given Class of Service Frame Set with a given CIR and Performance Tier, it is recommended that $\Delta CBS$ satisfy the following equation:

$$\Delta CBS \leq (1/8) * (CIR * R_O * D_S)/(R_O - CIR) \text{ (in bytes)}$$

This equation provides guidance, but due to delay variation in MEN-1 and/or other factors, may not always be sufficient.

As an example of applying this equation, consider an EVC with CoS H, CIR=10 Mbps, in PT1; further assume $B_S = 0.5$ and an average frame size of 500 bytes. Then, for a UNI line rate of 100 Mbps, the UNI information rate $R_O$ is 96.15 Mbps. Then $D_S$ = 2.5 msec, and $\Delta CBS$ = 3488 bytes. If $CBS_{OVC-1}$ = 33 KB and $CBS_{OVC-2}$ = 30 KB, the OVCs comprising the EVC conform to the equation above since $CBS_{OVC-1} – CBS_{OVC-2} \leq 3488$ bytes.

Values of $\Delta CBS$ for representative values of $CIR$ and $R_O$ are given in Table 43 in Section 8.7.1.1.

#### 8.7.1.1 Representative Values for ΔCBS

Values of $\Delta CBS$ (as defined in Section 8.7.1) for representative values of $CIR$ and $R_O$ are shown in Table 43. In this table, $B_S$ is assumed to be 0.5 for all PT/CoS Label combinations. The values of $R_O$ listed in the table correspond to UNI line rates of 10 Mbps, 100 Mbps, 1000 Mbps, and 10,000 Mbps respectively, for an average frame size of 500 bytes.

[6] *CIR* is defined in terms of MAC DA through FCS, not counting IFG and preamble; whereas a utilized line rate includes start of frame delimiter, IFG and preamble. The information rate $R_O$ must be expressed in terms of an average frame size (e.g., 500KB).

| Perf. Tier | CoS Label | FDR (ms) | Shaping Delay (ms) | ΔCBS (bytes) | | | |
|---|---|---|---|---|---|---|---|
| | | | | CIR=1 Mbps, $R_0$=9.615 Mbps | CIR=10 Mbps, $R_0$=96.15 Mbps | CIR=100 Mbps, $R_0$=961.5 Mbps | CIR=1000 Mbps, $R_0$=9615 Mbps |
| PT1 | H | 5 | 2.5 | 349 | 3,488 | 34,877 | 348,772 |
| | M | 10 | 5.0 | 698 | 6,975 | 69,754 | 697,545 |
| | L | 16 | 8.0 | 1,116 | 11,161 | 111,607 | 1,116,071 |
| PT2 | H | 10 | 5.0 | 698 | 6,975 | 69,754 | 697,545 |
| | M | 50 | 25.0 | 3,488 | 34,877 | 347,772 | 3,487,723 |
| | L | 100 | 50.0 | 6,975 | 69,754 | 697,545 | 6,975,446 |
| PT3 | H | 12 | 6.0 | 837 | 8,371 | 83,705 | 837,054 |
| | M | 50 | 25.0 | 3,488 | 34,877 | 348,772 | 3,487,723 |
| | L | 165 | 82.5 | 11,509 | 115,095 | 1,150,949 | 11,509,487 |
| PT4 | H | 40 | 20.0 | 2,790 | 27,902 | 279,018 | 2,790,179 |
| | M | 50 | 25.0 | 3,488 | 34,877 | 348,772 | 3,487,723 |
| | L | 200 | 100.0 | 13,951 | 139,509 | 1,395,089 | 13,950,893 |

**Table 43: Representative Values for CBS Ranges**

#### 8.7.1.2 Upper Bounds for Burst Sizes

The shaping delay as defined in Section 8.7.1 is a function of the difference in CBS values ($CBS_{OVC-1}$ - $CBS_{OVC-2}$), yet is insensitive to the absolute CBS values; e.g., $CBS_{OVC-1}$ = 66KB & $CBS_{OVC-2}$ = 60KB has similar shaping delay as $CBS_{OVC-1}$ = 12KB & $CBS_{OVC-2}$ = 6KB.

However, the absolute CBS values have an impact on egress transmission buffer sizing. For example, for the same 6KB difference in burst size, $CBS_{OVC-1}$ = 66KB & $CBS_{OVC-2}$ = 60KB will require more transmission buffer than $CBS_{OVC-1}$ = 12KB & $CBS_{OVC-2}$ = 6KB. Figure 8 below shows the relationship between the shaper buffers and a transmission buffer associated with the transmission link.

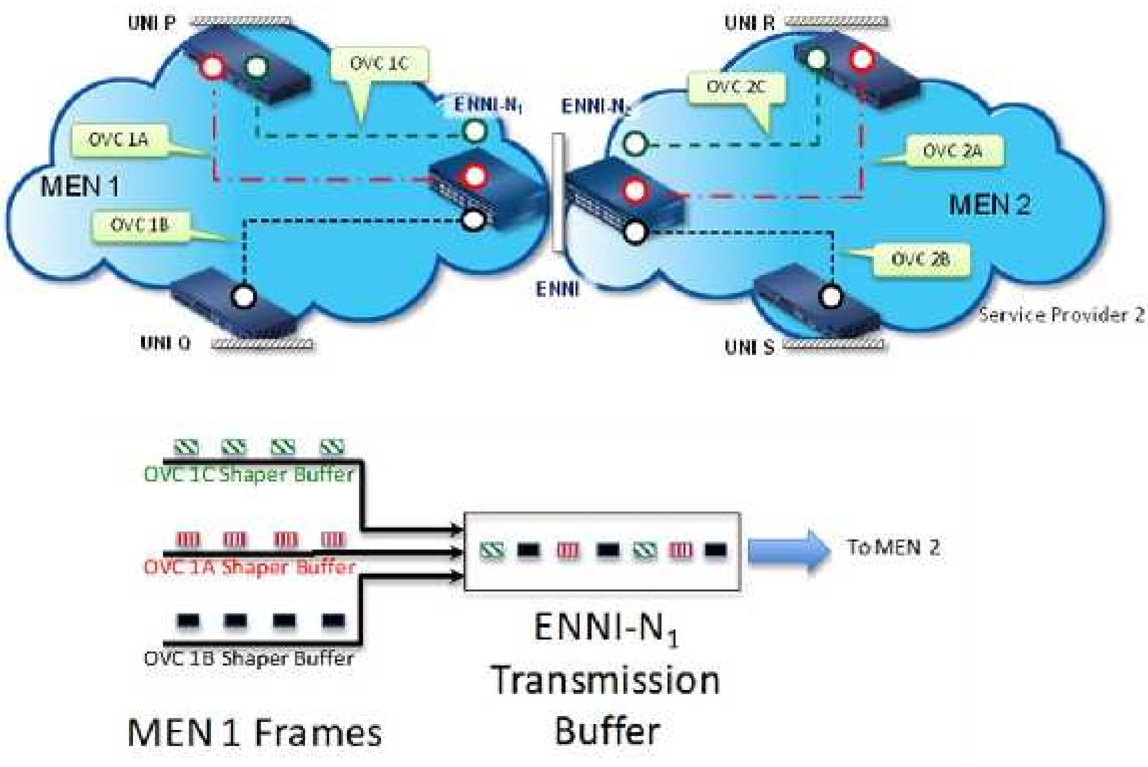


**Figure 8 – Shaper Buffers and Transmission Buffer Relationship**

In order to enable appropriate transmission buffer sizing, a Service Provider ought to address how the Operators configure absolute CBS values (in addition to adhering to the *ΔCBS* guidance in section 8.7.1).

Quantitative guidance on setting upper bounds on burst sizes is beyond the scope of this phase of CoS IA.

### 8.7.2 Shaping Considerations for Burst Alignment

This section presents example shaper algorithms in support of the ENNI Burst Size and Burst Alignment guidance in Section 8.7.1.

Section 10.3 of MEF 10.2 [2] ("Traffic Shaping") describes a pair of single-bucket shaper algorithms for implementation at the CE: The first example algorithm ("Periodic Algorithm") is intended to be run every $\Delta t$ seconds, where $\Delta t$ = the token bucket refresh rate; the second example algorithm ("New Frame Algorithm") is designed to be run every time a new frame arrives at the shaper.

Similarly, we define a pair of example algorithms for implementation at the egress of MEN-1 at the ENNI. The New Frame Algorithm is updated to be Color Aware, so that it handles Yellow frames coming from MEN-1. In the example New Frame Algorithm, no Yellow frames are ever placed in the shaper buffer[7] (and will be dropped if required). Thus, no Yellow frame will ever delay a Green frame due to shaping. There may be Yellow frames ahead of a Green frame in the transmission buffer, but that is no different from current practice.

The following parameters are used in the example algorithms (using the notation from [2]):

- CIR = the shaping rate of Green frames (average output rate of the shaper); Specifically, this equals the CIR of the Ingress Bandwidth Profile of MEN-2 in Figure 7;
- CBS = the shaping burst size of Green frames (maximum output burst of the shaper); Specifically, this equals the CBS of the Ingress Bandwidth Profile of MEN-2 in Figure 7;
- CBS* = the accepted burst of Green frames (maximum shaper buffer size for Green frames),
- CBS* ≥ CBS, which means the shaper accepts larger bursts at its input and generates smaller bursts at its output,
- EIR = the shaping rate of Yellow frames (average output rate of the shaper); Specifically, this equals the EIR of the Ingress Bandwidth Profile of MEN-2 in Figure 7;
- EBS = the shaping burst of Yellow frames (maximum output burst of the shaper); Specifically, this equals the EBS of the Ingress Bandwidth Profile of MEN-2 in Figure 7.

The following notation is used in the example algorithms (following the definitions in [2]):

- B(t) = the instantaneous shaper buffer occupancy in bytes,
- C(t) = the instantaneous value of the tokens in the Committed token bucket,
- E(t) = the instantaneous value of the tokens in the Excess token bucket,
- L = the length of the frame at the head of the shaper buffer, and
- THS = a configured buffer threshold such that the difference between THS and the shaper's buffer size, CBS*, is large enough to hold a maximum sized frame.

Note that C(t) and E(t) are assumed to be updated with additional tokens by a separate process run at a period equal to the token refresh rate $\Delta t$, i.e.,

[7] Note that we differentiate between the shaper buffer, and the transmission buffer (outgoing link queue). Frames taken from the head of the shaper buffer are enqueued on the transmission buffer and transmitted at line rate. We assume that the transmission buffer remains unchanged from the current situation.

$C(t) = \min(CBS, C(t) + (CIR/8) * \Delta t)$ and $E(t) = \min(EBS, E(t) + (EIR/8) * \Delta t)$.

Because no Yellow frames are ever placed in the shaper buffer, our example Periodic Algorithm is essentially unchanged from the MEF 10.2 example algorithm:

```
while((L <= C(t)) && (B(t) > 0))
{
   C(t) = C(t) - L;
   B(t) = B(t) - L;
   send the frame at the head of the shaper buffer to the transmission buffer;
//Should be declared green
}
```

**Figure 9 – Periodic Algorithm**

The revision of the New Frame Algorithm from [2] to handle transmission of Yellow frames is shown below. We add the following notation:

- LNF = the length of the newly-arrived frame.

```
if(B(t) == 0) // If shaper buffer is empty
{
    if(new frame color is Yellow)
    {
        if(LNF <= E(t))
        {
            E(t) = E(t) - LNF;
            send new frame to transmission buffer;
        }
        else
        {
            discard new frame;
        }
    }
    else // new frame is Green
    {
        if(LNF <= C(t))
        {
            C(t) = C(t) - LNF;
            send new frame to transmission buffer;
        }
        else
        {
            place new frame in shaper buffer;
            B(t) = B(t) + LNF;
        }
    }
}
else  // shaper buffer is not empty
{
    while(L <= C(t)) && (B(t) > 0))
    {
       C(t) = C(t) - L;
       B(t) = B(t) - L;
       send the frame at the head of the shaper buffer to transmission buffer; // All
frames already in shaper buffer are Green
    }
    if((new frame color is Green) && (B(t) <= THS))
    {
        place new frame in shaper buffer;
        B(t) = B(t) + L;
    }
    else // new frame is Yellow, or no room in shaper buffer for another Green frame
    {
       discard new frame;
    }
}
```

**Figure 10 – Revised New Frame Algorithm**