Best Practices: Multiple Classes of Service in Carrier Ethernet Mobile Backhaul Networks

April 2013

MEF Best Practices Document

Table of Contents

1 Introduction 3
1.1 Purpose of this Best Practices Paper 3
1.2 Recommendations Made in the Paper 3
1.3 Audience 3
1.4 Objectives of this Best Practices Paper 3
2 Assumptions 3
3 Ethernet Virtual Connection (EVC) Service Models 4
3.1 Single EVC Mobile Backhaul model 4
3.2 Redundant EVC Mobile Backhaul model 6
3.3 Multipoint EVC Mobile Backhaul model (three points) 8
3.4 Carrier Ethernet Class of Service Models 10
4 Recommendations 10
4.1 Carrier Ethernet Class of Service Recommendations 11
4.2 Ingress Bandwidth Profile Recommendations 11
4.3 EVC-based performance models and objectives 13
4.4 Performance Monitoring 16
4.5 Oversubscription 17
5 What Mobile Operators need for Multi-CoS Backhaul 17
6 Challenges 17
7 Conclusions 18
8 Glossary and Terms 18
9 References 21
10 About the MEF 21
11 Acknowledgements 21

Table of Figures

Figure 1: Example of Single EVC Backhaul 5
Figure 2: Example of Redundant EVC Backhaul 6
Figure 3: Example of Multipoint EVC Backhaul using three UNIs 9

Table of Tables

Table 1: Typical Characteristics of Single EVC Backhaul 6
Table 2: Typical Characteristics of Redundant EVC Backhaul 8
Table 3: Typical Characteristics of Multipoint EVC Backhaul 10
Table 4: Performance Set 1 Recommendations for Green (CIR and CBS conformant) frames 15
Table 5: Performance Set 2 Recommendations for Green (CIR conformant) frames 15
Table 6: Performance Set 3 Recommendations for Green (CIR conformant) frames 15

1 Introduction

1.1 Purpose of this Best Practices Paper

The deployment of Ethernet mobile backhaul networks has enabled the initial transition from voice dominant to data dominant mobile backhaul infrastructure and has lead to cost savings over legacy TDM deployment. However, most networks have been provisioned with a single Class of Service requiring a largely inefficient allocation of network resources in order to enable high quality of service. The resulting inefficiency is a significant cause for concern for Mobile Operators and Backhaul Providers. It is recognized as a major challenge and operating cost in the industry. The discussion on the impact and scale of these issues are highlighted in a companion MEF positioning paper "The Benefits of Multiple Classes of Service for Ethernet Mobile Backhaul" January 2012.

1.2 Recommendations Made in the Paper

This paper investigates and recommends best practices for deployment of Carrier Ethernet services provisioned with multiple classes of services. Recommendations also note the underlying rationale and potential impact of adoption together with possible consequences of not implementing the recommendation so that readers can be better positioned to make implementation decisions. The document describes ways for the Backhaul Provider to optimize the use of mobile backhaul bandwidth resulting in considerable cost savings while improving application-specific quality of service.

1.3 Audience

Whereas the companion Positioning paper focuses on the cost benefits for implementation, this paper is intended to provide senior technical management and network planners with overview guidance on how to implement multiple classes of services for Mobile backhaul. This paper is not exhaustive but is a precursor for more detailed investigation. This document recognizes that almost every Mobile Operator and Backhaul Provider have their own set of criteria and therefore the examples of best practices shown here should be interpreted as a guide rather than as a rigid set of rules.

1.4 Objectives of this Best Practices Paper

- To clearly define and describe both Single CoS (which is most prevalent today) and Multi-CoS Mobile Backhaul.
- To provide both guidance and implementation choices for multi-CoS mobile backhaul together with the implications of those choice for various scenarios
- To address the top challenges that Mobile Operators and Backhaul Providers face in moving to Multi-CoS Mobile Backhaul.

2 Assumptions

The recommendations in this paper are based on the following assumptions:

1. The Mobile Operator network and Backhaul Provider network are separate, i.e., the Mobile Operator buys backhaul service from the Backhaul Provider. It is recognized that this is not the case in many deployments. For such cases where it is an integrated network, then rate enforcement in the Backhaul Provider's network may not be used. That said, the control of performance and cost of the backhaul are equally important in either scenario.

2. Mobile Backhaul for LTE will require large backhaul bandwidths, while at the same time controlling the cost per bit.

   From NGMN White Paper: "Radio spectrum for mobile broadband is an expensive and limited resource, so backhaul should be generously provisioned to exceed cell throughput in most cases. At the same time, LTE needs to operate at a significantly lower cost per bit, so operators cannot afford to over-provision either. In this analysis, we assume backhaul should be provisioned to cope with all but the top 5% of cell throughputs (i.e. the 95 percentile of the cell throughput distribution)."

The advent of small cells is putting further pressure on backhaul cost per bit.

3. The three standardized Classes of Service defined in MEF 23.1 (H, M and L) are the ones offered by the Backhaul Provider. Combinations of 2-CoS may also be supported (H, M), (H, L). It is recognized that H+ Class of Service has been defined in MEF 22.1 for carrying synchronization traffic in some cases. The performance objectives for ‘H+’ are currently identical to the performance objectives for ‘H’, except for Frame Delay Range metric, which is undefined (assumed more stringent than H FDR). There is ongoing work to develop additional metrics and objectives that are optimized for synchronization traffic. As the H+ definition and new metric(s) mature, this paper can be updated.
4. Each Class of Service has its own performance guarantees

   Performance Tier 1 (metro scope) is the focus of this paper. Multiple models of SLS can be supported (three are recommended). The time period for Service Level Specification, T, is one month; typically, one month is the norm for Service Level Specifications (SLS).
5. Multi CoS backhaul from a given cell site can be supported with different EVC configurations; one (or more) multi-CoS EVCs or multiple single CoS EVCs.
6. The Backhaul Provider is able to support rate enforcement, per MEF 10.2, with CoS ID based on PCP value of the service frame. Optionally, CoS ID based on DSCP could be used in certain cases.
7. Preservation of service frame VLAN ID is normally desired by Mobile Operators (for administrative reasons), although there are use cases where configuration at the cell site could be simplified by always using the same VLAN ID(s), and doing the translation at the aggregation UNI. Similarly, preservation of the PCP field is normally not required (the Mobile Operator typically uses layer 3 devices to connect to the CEN, and these neglect the PCP marking of incoming service frames), although there may be cases where the Mobile Operator uses Layer 2 devices to connect directly to the MEN, and in these cases, the Mobile Operator may prefer to have the PCP values preserved.
8. The Ethernet Virtual Connection (EVC) to each cell site is used for all traffic types, including backhaul voice and data, packet synchronization (if used), cell-to-cell data and cell site management and control.

3 Ethernet Virtual Connection (EVC) Service Models

This section presents three commonly used Mobile Backhaul EVC Models. These include: single point-to-point (p2p) EVC; redundant p2p EVCs; and multipoint EVC (three point). Other models may be used, as well, but they are not discussed here.

3.1 Single EVC Mobile Backhaul model

This is the basic model, where a Mobile Operator buys a single p2p EVC to connect the UNI at the cell site with the UNI at the aggregation site, where services from different cell sites are multiplexed. MEF refers to this service as an Ethernet Virtual Private Line Service, or EVPL. This model is usually used due to cost constraints on the part of the Mobile Operator. See Figure 1 below for an example.

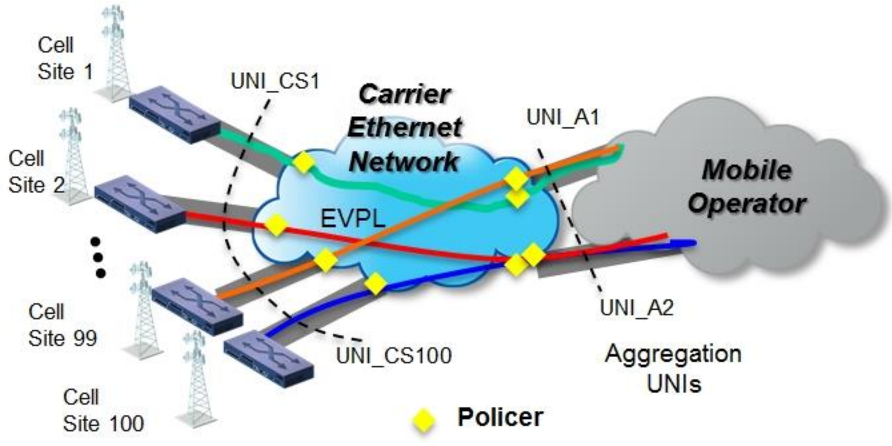

Figure 1: Example of Single EVC Backhaul

In this example, two UNIs are used for connecting at the aggregation site. A Point-to-Point EVC is used from each cell site to one of the aggregation UNIs. 50 EVCs are mapped to one aggregation UNI and the other 50 are mapped to the other aggregation UNI. It is assumed that each of the sets of 50 EVCs should be routed diversely within the CEN, and that the Mobile Operator and Backhaul Provider will work cooperatively to identify which cell sites should be mapped to which aggregation UNI, to minimize single point of failure impacts to the Mobile Operator service.

All traffic from each cell site travels over the p2p EVC back to the aggregation site. Note that X2 Hand-over (HO) traffic goes from cell site to aggregation site and back to the neighboring cell site. Each cell site requires x Mbps (upstream and downstream). X can range from 20 Mbps to 200+ Mbps. Emerging small cells may require less bandwidth than macro cells. Mobile backhaul traffic is often asymmetrical – i.e., higher rate downstream than upstream, and with small cell cost sensitivity the backhaul may become more asymmetrical as well (e.g., based on an asymmetric access technology such as DSL or DOCSIS). Actual UNI link utilization may be significantly less than the aggregate provisioned bandwidth at aggregation UNIs, since not all cell sites peak at same time.

Table 1 below summarizes the key characteristics of the Single EVC Backhaul scenario.

Attribute of Service	Typical Configuration	Comments
Cell Site UNI	100M, 1G optical or copper PHY	A single cell site UNI is typical. A second cell site UNI, for redundancy, may be used in certain specific cases, but is not shown in the above figure.
Aggregation Site UNI	10G optical PHY	1G also useful for sparse deployments
Aggregation Site Redundancy	Multiple Aggregation UNIs on different CEN switches	At least two different CEN switches are used for these UNIs
EVC Type	Point-to-Point	Ethernet Virtual Private Line (EVPL) service is used
EVC routing	Distributed EVCs possible	Single point of failure analysis is critical. Backhaul Provider and Mobile Operator may cooperate closely to route adjacent cell sites to different aggregation UNIs on different switches, to minimize catastrophic failure scenario.

Attribute of Service	Typical Configuration	Comments
Cell Site Bandwidth	20-200 Mbps	Macro sites that are used to aggregate other macros or small cells may exceed 200 Mbps. Small cells may be less than 20 Mbps.
Class of Service Label	H, M and/or L	The single EVC backhaul scenario
Rate enforcement	Single ingress Bandwidth Profile (policer) is applied per {EVC, CoS} at each UNI	Bandwidth values are normally symmetrical, although they could have different values up and down-stream.

Table 1: Typical Characteristics of Single EVC Backhaul

3.2 Redundant EVC Mobile Backhaul model

In this case, a Mobile Operator buys two p2p EVCs to connect the UNI at the cell site with a pair of UNIs at the aggregation site. This model is the same as the single EVC model, with one obvious exception: the number of EVCs, and therefore the protection philosophy. The redundant EVC model is normally used where the Mobile Operator favors reliability over cost. See

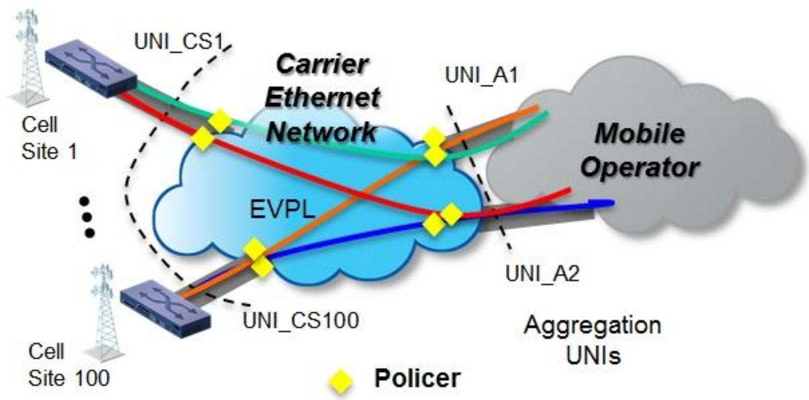

Figure 2 below for an example.

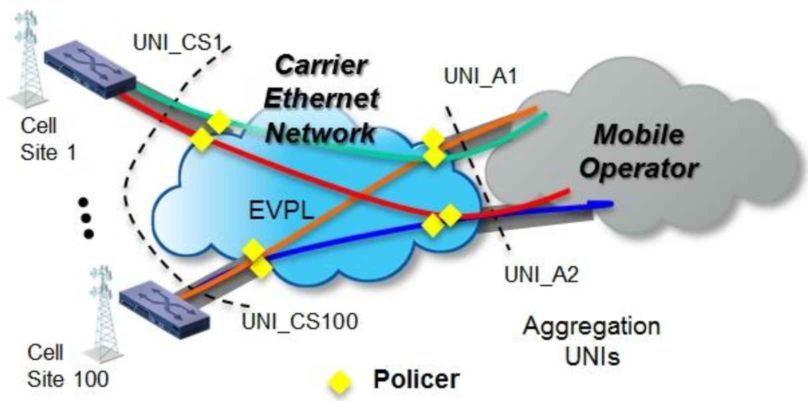

Figure 2: Example of Redundant EVC Backhaul

In this example, 100 EVCs are mapped to one aggregation UNI and the other 100 are mapped to the other aggregation UNI. Each aggregation UNI (and access) needs to be designed to support traffic from all cell sites in the event of a failure of one aggregation UNI. The pair of EVCs at each cell site should be diversely routed from each other, although it may not always be possible to achieve route diversity. The Mobile Operator is responsible for protection – detection and protection switching on a per cell site basis.

In normal operation, half of the EVCs are active on one path, and the other half active on the second path. If a failure is detected on the 'working' EVC, then the service is switched to the 'back-up' EVC. The Mobile Operator decides on the protection scheme (e.g., BFD or Ethernet CCM could be used to detect the failure), and is in complete control of it. Typically, 50 ms failover can be achieved.

Table 2 below summarizes the key characteristics of the Redundant EVC Backhaul model.

Attribute of Service	Typical Configuration	Comments
Cell Site UNI	100M, 1G optical or copper PHY	A single cell site UNI is typical. A second cell site UNI, for redundancy, may be used in certain specific cases, but is not shown in the above figure.
Aggregation Site UNI	Two 10G UNIs, optical or copper PHY using LAG for protection	1G is also useful for sparse deployments In dense deployments, 100G UNIs may be desired.
Aggregation Site Redundancy	Two CEN switches used; one for each UNI in the pair	
EVC Type	Point-to-Point	Ethernet Virtual Private Line (EVPL) service is used
EVC routing and protection	Mobile Operator is in control of protection across the pair of EVCs.	The two EVCs in the pair should be diversely routed in the CEN.
Cell Site Bandwidth	20-200 Mbps per macro base station	Macro sites that are used to aggregate other macros or small cells may exceed 200 Mbps. Small cells may be less than 20 Mbps.
Class of Service Label	H, M and/or L	
Rate enforcement	Ingress Bandwidth Profile (policer) is applied per {EVC, CoS} at each UNI. Alternatively, an aggregate policer for each CoS could be applied across the pair of EVCs at the cell site UNI. An aggregate Egress Bandwidth Profile (policer) may be applied for each CoS across the pair of EVCs at the cell site UNI.	The Mobile Operator should take care to manage the aggregate bandwidth such that the heartbeat on the working and back-up EVCs take priority over the data.

Table 2: Typical Characteristics of Redundant EVC Backhaul

3.3 Multipoint EVC Mobile Backhaul model (three points)

In this case, a Mobile Operator buys a single multipoint EVC to connect the UNI at the cell site with a pair of UNIs at the aggregation site. MEF refers to this service as an Ethernet Virtual Private LAN Service, or EVP-LAN. In this model, a bridging point is located in the CEN, enabling the three UNIs to interconnect. This model provides redundancy at the aggregation UNI and protection for a portion of the EVC – the level of the protection is dependent on the location of the bridging point in the CEN (the closer to the cell site, the broader the level of protection). Each aggregation UNI (and access) needs to be designed to support traffic from all cell sites in the event of a failure of one aggregation UNI. See Figure 3 below for an example.

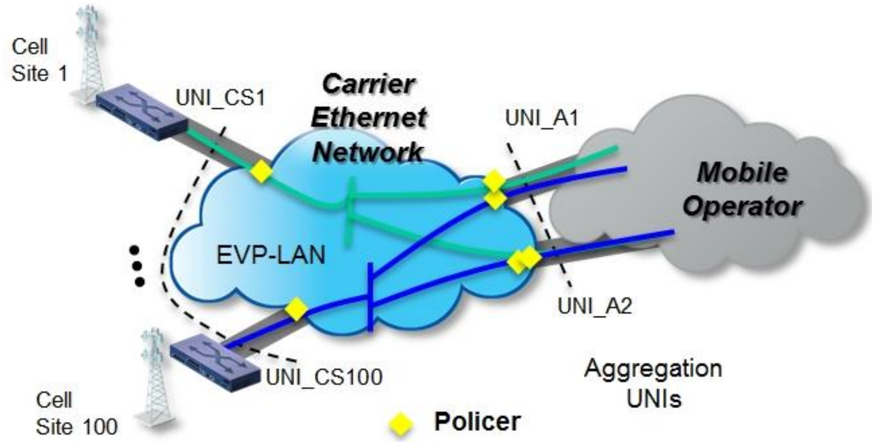

Figure 3: Example of Multipoint EVC Backhaul using three UNIs

In this example, the two legs of the EVC from the bridging point to the aggregation UNIs are assumed to be diversely routed (note: such diverse routing cannot always be achieved in the Backhaul Provider's CEN). As in the redundant EVC case, the Mobile Operator is responsible for protection.
In normal operation, and assuming the E-LAN service supports MAC learning and forwarding, only one path of the EVC carries the unicast traffic, since only the 'active' UNI responds to packets such as ARP, and the other path is there as a backup. If a failure is detected on the 'primary path', then the service is switched to the 'secondary path', since now the 'secondary' UNI will respond to packets such as ARP. The Mobile Operator decides on the protection scheme (e.g., VRRP, BFD, Ethernet CCM or a combination of these, could be used to detect the failure), and is in complete control of it. Typically, 50 ms failover can be achieved for point-to-point protection schemes overlaid over the multipoint EVC.

Table 3 below summarizes the key characteristics of the Multipoint EVC Backhaul model.

Attribute of Service	Typical Configuration	Comments
Cell Site UNI	100M, 1G optical or copper PHY	A single cell site UNI is typical. A second cell site UNI, for redundancy, may be used in certain specific cases, but is not shown in the above figure.
Aggregation Site UNI	Two 10G UNIs, optical or copper PHY using LAG for protection	1G is useful for sparse deployments. In dense deployments, 100G UNIs may be desired.
Aggregation Site Redundancy	Two CEN switches or routers used; one for each UNI in the pair	
EVC Type	Multipoint-to-Multipoint	Ethernet Virtual Private LAN (EVP-LAN) service is used
EVC Protection	Mobile Operator is in control of protection across the paths of the EVC, for example, using VRRP for determining the master router at the aggregation site; and BFD or CCM for failure detection of the eNodeB to aggregation site connectivity.	Single point of failure analysis is critical
Cell Site Bandwidth	20-200 Mbps	
Class of Service Label	H, M and/or L	
Rate enforcement	Ingress Bandwidth Profile (policer) is	The Mobile Operator should take care to

Attribute of Service	Typical Configuration	Comments
	applied per {EVC, CoS} at the UNI. An aggregate Egress Bandwidth Profile (policer) may be applied per {EVC, CoS} at the cell site UNI.	manage the aggregate bandwidth, such that the 'heartbeat' on the working and back-up paths is prioritized over the data.

Table 3: Typical Characteristics of Multipoint EVC Backhaul

3.4 Carrier Ethernet Class of Service Models

In a previous paper entitled "The Benefits of Multiple Classes of Service for Ethernet Mobile Backhaul" [1], the single and multi-CoS models were explained at a high level and compared primarily on the basis of cost. The reader is referred to that paper for the rationale as to why multi-CoS makes sense. This paper focuses on the technical recommendations as to how to implement multi-CoS Mobile Backhaul solutions.

In the multi-CoS model, the Mobile Operator works with the Backhaul Provider on a multi-CoS policy. Such factors as application mapping to Carrier Ethernet Classes of Service, bandwidth requirements, burst size requirements, performance requirements and cost are all included in the policy.

CoS Label 'H' is normally included in the multi-CoS offer, but it is limited in bandwidth ('H' should be more expensive and supports smaller burst sizes, which are not optimal for data services). 'H' typically has a CIR only (no EIR) bandwidth component. Buffering inside the network tends to be small for this queue, to control delay and delay variation. Because of that, the Committed Burst Size (CBS) parameter of the Ingress Bandwidth Profile is also small, to minimize bursts per flow causing congestion in the small queues. A cautious 'under-subscription' approach is prudent for H CoS, allowing for speed up on the link, e.g., aggregate CIR for H CoS on a 10G UNI might be considerably less than 5 Gbps.

CoS Labels 'M' and 'L' are designed for data applications that are more tolerant to delay variation and require larger burst capability. 'M' CoS is generally for apps with stringent loss requirements (less stringent delay related requirements) and 'L' CoS is generally for apps that are more tolerant to delay and loss (or apps that are economically driven to lowest cost transport). Having both M and L CoS in the policy is very useful for differentiation. Such differentiation could be based on traffic types (e.g., video streaming, Internet) or customer type (e.g., Customer X just exceeded his byte count for the month, and now all his frames use the lowest CoS).

CoS definitions must be flexible to support different mapping options for the Mobile Operator. For example, one Mobile Operator might want to map mobile applications (i.e., 9 QCI priorities for LTE) into two or three Carrier Ethernet Classes of Service. The nine QCI priorities are specified by 3GPP [4], and replicated in Table 16 of MEF 22.1 [2]. Another Mobile Operator might want to differentiate based on customer, e.g., Subscriber A pays more than B per month, so data service is prioritized for the month. Subscriber C has exceeded his bandwidth limit for the month, so his traffic gets lowest priority for the remainder of the month. Another Mobile Operator may differentiate based on both.

4 Recommendations

This section provides recommendations to the Backhaul Provider and Mobile Operator involved in Mobile Backhaul. These are meant as 'Best Practice' guidelines for MEF services, and provide a standardized approach to multi-CoS Mobile Backhaul. Actual implementations can vary from these recommendations, as needed. This document may be updated as the multi-CoS practice matures.

Specific recommendations are shown as [RC1], [RC2], etc.

4.1 Carrier Ethernet Class of Service Recommendations

There are three standard MEF Class of Service Labels, as defined in MEF 23.1 [3]. These are 'H', 'M' and 'L'. It is assumed that a Backhaul Provider offering multi-CoS Mobile Backhaul would support all three. Of course, any two-CoS model using standard MEF CoS may be used for a given application. An additional CoS, H+, is defined in MEF 22.1 [2] for packet synchronization traffic. While this may be used, since completing the specification of performance objectives for H+ is still in discussion within MEF 22.2, we will assume that packet synchronization frames, if used, will use the H CoS. In the future, H+ may be the preferred CoS for synchronization traffic in certain use cases.

There are two basic options for identifying the Class of Service of an ingress service frame at the UNI for a multi-CoS EVC, as specified in MEF 10.2 [5]. One approach is a layer 2 method, using the PCP value in the VLAN-tagged header; the other is layer 3 based, using the DSCP value in an IPv4 packet, which could be useful for Mobile Operators who don't need to bother to map the DSCP field to the PCP field in the customer edge routers. Since the PCP method is more suited to Ethernet services, we focus on that here.

The following summarizes the recommendations for the three MEF Classes of Service.

Carrier Ethernet Class of Service Recommendations:

[RC1] It is strongly recommended to support at least one of the following CoS Label sets for Mobile Backhaul: {H, M, L}, or {H, M} or {H, L}.

> Note: In the above recommendation, H class is included to support control signaling, voice & synchronization guarantees of tight control of delay and delay variation. Possible consequences of not implementing this recommendation could be poor performance for those applications requiring tight control of delay and delay variation.

[RC2] It is strongly recommended to support a multiple CoS Label set on one EVC, i.e., a multi-CoS EVC.

> Note: Implementing Multi-CoS reduces the number of EVCs required. Single CoS EVCs may be acceptable, but in order to implement Multi-CoS backhaul, at least one EVC for each CoS will be required, which results in increased administrative overhead.

[RC3] It is strongly recommended to support PCP-based classification for a multi-CoS EVC, as per [MEF 22.1, D16].

> Note: The recommendation is based on the fact that Carrier Ethernet deployments typically use PCP rather than DSCP for traffic differentiation. Possible consequences of not implementing this recommendation could be not matching the Mobile Operators requirements, as well as a lack of consistency in classification mechanisms of multi-CoS EVCs (since many Backhaul Providers may not offer layer 3 classification, PCP-based classification would be the most common). Implementations may support DSCP-based classification for a multi-CoS EVC, as per [MEF 22.1, O12].

4.2 Ingress Bandwidth Profile Recommendations

MEF 10.2 [5] and MEF 23.1 specify applying a separate Ingress Bandwidth Profile (ingress policer) per CoS ID for a multi-CoS EVC. This section recommends Ingress Bandwidth Profile parameters for Mobile Backhaul. Both Backhaul Providers and Mobile Operators should review these recommendations carefully, since the Mobile Operator will need to shape its traffic into each CoS.

The Ingress Bandwidth Profile has six parameters that can be used to configure a given service. The algorithm is defined in section 7.11.1 of MEF 10.2 [5], and is based on a dual token bucket methodology.

- Committed Information Rate (CIR) – the long-term average rate, in bits per second, which the subscriber should expect to achieve from the service.

- Committed Burst Size (CBS) – the short-term burst, in bytes, allowed on the UNI for traffic to be conformant.
- Excess Information Rate (EIR) - the additional long-term average rate, in bits per second, which the subscriber may be able to get from the service, depending on congestion conditions within the network. CIR+EIR can be thought of as the peak information rate.
- Excess Burst Size (EBS) - the additional short-term burst, in bytes, allowed on the UNI.
- Color Mode (CM) – 'Color-Aware' mode means that the Ingress Bandwidth Profile can differentiate between 'Green' and 'Yellow' frames marked by the subscriber. 'Color Blind' mode means that the Ingress Bandwidth Profile cannot determine the color of the incoming service frame.
- Coupling Flag (CF) – a useful parameter when in Color Aware mode; CF allows the overflow tokens from the CBS bucket to be used for the EBS bucket.

These parameters may not be fully orderable (e.g., some Backhaul Providers may fix the CBS to a certain value, and only offer CM='Color Blind' and CF=0 functionality). In this document, the focus is on the first four parameters listed above. Services based on Color Aware mode may be part of a future study.

For CoS Label 'H', {CIR, CBS} are the only parameters specified in MEF 23.1 [3], and the only ones normally used in practice. The result is 'Green' frames only in the network, for which performance objectives apply.

For CoS Label 'M', {CIR, CBS, EIR, EBS) may be specified. Typically, CoS Label 'M' is more weighted to CIR than EIR. The additional EIR may help certain Mobile Operators get a little bit more out of their service, trading off loss guarantees for the 'Yellow' frames. Only the 'Green' frames are counted against the performance guarantees.

For CoS Label 'L', {CIR, CBS, EIR, EBS) may be specified. Typically, CoS Label 'L' is more weighted to EIR than CIR. The CIR provides a baseline guarantee of performance, and the additional EIR provides burst to possibly much higher levels.
Note that CIR-only services are simple to define, and carry with them performance guarantees for all conformant traffic. 'CIR + EIR' services provide added flexibility for the Mobile Operator for M and L CoS; the Mobile Operator can get more bandwidth at lower cost, taking the risk of a performance hit. The Mobile Operator has a choice to shape to CIR only, or somewhere between CIR and CIR+EIR. Some Mobile Operators are considering the use of CIR+EIR backhaul services.

Ingress Bandwidth Profile Recommendations:

[RC4] It is recommended to support the CIR-only (EIR=0) for CoS Label 'H'.

Note: This guarantees that all traffic sent through the network will be 'Green' frames, with guaranteed delivery. Possible consequences of not implementing this recommendation could be poor performance for the non-guaranteed 'Yellow' frames in the H CoS.

[RC5] It is recommended to support the option of CIR and EIR > 0 for CoS Labels 'M' and 'L'.

Note: This is recommended in order to meet many Mobile Operators' requirements to optimize use of physical bandwidth. Possible consequences of not implementing this recommendation could be lack of meeting of some Mobile Operator requirements.

[RC6] It is recommended to support 'Color Blind' mode of operation of the Ingress Bandwidth Profile

Note: This is recommended in order to meet the deployments of most Backhaul Providers, and to provide consistency.

[RC7] It is recommended to support at least the minimum {CBS, EBS} value (8 * MTU), as specified in Requirement 36 of the UNI Type 1 Implementation Agreement [7].

Note: This minimum value is to ensure good application performance. Providers may support greater than 8*MTU size. Possible consequences of not implementing this recommendation (i.e., small burst size) could be very poor performance for TCP type applications, for example causing very slow file/data transfer.

Shaping before policing is important to achieving optimal performance with the CoS and bandwidth available. Mobile Operators should shape their traffic onto the UNI to conform to the Backhaul Provider's bandwidth profile. Lack of a shaping ability could result in frames discarded by the BWP policer and significantly lower TCP throughput. The Backhaul Provider may opt to provide a larger CBS value, as appropriate, for a given Mobile Operator. It is recognized that there is active research and discussion ongoing about the proper burst size values for access networks, and in particular, mobile access networks. This recommendation may be re-assessed based on this ongoing work.

4.3 EVC-based Performance Models and Objectives

This subsection briefly describes the performance metrics and models that apply to the three standard Carrier Ethernet Class of Service Labels – H, M and L – for Mobile Backhaul, and makes recommendations on the Service Level Specification. For a detailed definition of the metrics, please see section 6.9 of MEF 10.2 [5]. A brief summary of the metrics and related parameters is provided here, for reader clarity.

Frame Delay (FD): This is a one-way definition from ingress UNI to egress UNI, first bit in to last bit out. The FD metric includes the following basic parameters: T (Time period of the SLS, typically 1 month), O (Objective, in ms), P (Percentile for which the objective applies).

Mean Frame Delay (MFD): An objective is set for the average frame delays over T.

Frame Delay Range (FDR): An objective is set for the difference between the percentile delay and the minimum delay over T.

Inter Frame Delay Variation (IFDV): An objective is set for the difference in delay between a pair of selected frames. Normally a short time interval, like 1 second, is used for selecting frame pairs. As in FD, this is a percentile-based objective over T.

Frame Loss Ratio (FLR): An objective is set for percentage of frames lost during T. For example, an FLR of 0.1% means that up to 1 in 1,000 frames may be lost over the month.

Note that these metrics normally require the specification of S (the set of ordered UNI pairs). For this subsection, the assumption is point-to-point EVCs, where S is the full set of ordered UNI pairs.

Note also that a percentile-based objective can allow for a reasonable 'upper bound' on FD, FDR or IFDV, effectively subtracting out anomalous results. For example, a percentile of 99% allows for 1% of the delays to exceed the objective.

The Mobile Operator should understand and define performance metrics per MEF attribute definitions, and reflect these in the SLS, to simplify discussion with CEN Backhaul Operators - a bad scenario is when a Mobile Operator complains of poor 'jitter' performance and the Backhaul Provider says it is good – they are likely using different definitions of the metrics. Availability (and its relationship to the other performance metrics) should be included in the SLS (even though, availability objectives are not specified in MEF 23.1, the effect of availability state should be included in the SLS for the other metrics). See MEF 10.2.1 [6] for the definition of the availability metric, and its relationship to the other performance metrics. In general, when the service is in unavailable state the objectives for the other metrics do not apply. Availability, Resiliency and other metrics and associated parameters and objectives may be defined in future versions of MEF 23.1.

MEF 23.1 [3] specifies objective (and related parameter) values for each of the three standard CoS by Performance Tiers (PT), loosely thought of as geographic tiers. PT1 has the tightest objectives and is the focus of this recommendation, assuming a 'metro scope' backhaul scenario. See Tables 5 and 6 in MEF 23.1 [3]. For regional backhaul services, a different PT may apply.

Three SLS models, in no particular order, are described below (others may be possible, but these are recommended):

Model 1: Frame Delay, Frame Delay Range, Frame Loss Ratio {FD, FDR, FLR}

Model 2: Frame Delay, Inter Frame Delay Variation, Frame Loss Ratio {FD, IFDV, FLR}

Model 3: Mean Frame Delay, Frame Delay Range, Frame Loss Ratio {MFD, FDR, FLR}

The choice of either of these models allows a reasonable upper bound on FD, depending on the percentile, while allowing some flexibility in the SLS offering. Note: MEF 22.1 [R39] requires the inclusion of the FDR metric if the CoS Name will be carrying synchronization traffic (i.e., CoS Label H in this paper). Since there are continuing discussions in MEF re: the appropriate metric(s) for packet sync traffic, the need for inclusion is not quite clear. Backhaul Providers are encouraged to add an FDR metric to their SLS, as required by the Mobile Operator.

Service (EVC-based) Performance Recommendations:

[RC8] It is strongly recommended that the SLS use the following metrics, at minimum {FLR, FD or MFD, FDR or IFDV}. Other metrics may also be defined for p2p EVCs – Multipoint EVCs are outside the scope of these performance recommendations.

Note: This recommendation (which employs standardized MEF metrics) ensures that the SLS contains performance objectives that result in predictable delay, delay variation and loss performance for the Mobile Operator. This use of common, standardized definitions and performance objectives can result in a commonly agreed performance testing methodology, and enable a common approach for Mobile Operators partnering with many Backhaul Providers, greatly reducing implementation complexity and integration time. Possible consequences of not implementing this recommendation could be inconsistent metric definitions and performance objectives, resulting in finger pointing when the service isn't working to the expectations of the Mobile Operator.

[RC9] CoS Label 'H': It is strongly recommended to support Performance Sets 1, 2 or 3, as described in Table 4: Performance Set 1 Recommendations for Green (CIR and CBS conformant) frames, Table 5 and Table 6 below (other metrics may also be specified). If Performance Set 2 is used, and the Mobile Operator requires packet synchronization traffic, Backhaul Providers are encouraged to add an FDR metric to their SLS.

Note: Each of the three performance sets provide a predictable cap to the delay and delay variation and loss performance objectives. Possible consequences of not implementing this recommendation are that the upper limits of delays are not set with consequences such as unexpected or undesirable performance implications similar to that stated in [RC8].

[RC10] CoS Label 'M': It is strongly recommended to support one of the three performance sets described below in the SLS; IFDV and FDR need not be specified for CoS Label 'M', since the recommendation is to always support CoS Label 'H', which provides the delay variation guarantees (other metrics may also be specified).

Note: The rationale is similar to [RC9]. Possible consequences of not implementing this recommendation are similar to that stated in [RC8]

[RC11] CoS Label 'L': It is recommended to support one of the three performance sets described below in the SLS; IFDV and FDR need not be specified.

Note: The rationale is similar to [RC9]. Possible consequences of not implementing this recommendation are similar to that stated in [RC8].

Performance Set 1 – Frame Delay (FD) oriented SLS, with FDR and FLR				
Performance Metric	Parameters	Class of Service		
		H	M	L
Frame Delay	Time Period (T)	1 month	1 month	1 month
	Objective (ms)	≤10	≤20	≤37
	Percentile (%)	99.9	99	95
Frame Delay Range	Time Period (T)	1 month	1 month	N/S
	Objective (ms)	≤5	≤10 or N/S	N/S
	Percentile (%)	99.9	99	N/S
Frame Loss Ratio	Time Period (T)	1 month	1 month	1 month
	Objective (%)	≤.01	≤.01	≤.1

Table 4: Performance Set 1 Recommendations for Green (CIR and CBS conformant) frames

Performance Set 2 – Frame Delay (FD) oriented SLS, with IFDV and FLR				
Performance Metric	Parameters	Class of Service		
		H	M	L
Frame Delay	Time Period (T)	1 month	1 month	1 month
	Objective (ms)	≤10	≤20	≤37
	Percentile (%)	99.9	99	95
Inter Frame Delay Variation	Time Period (T)	1 month	1 month	N/S
	Objective (ms)	≤3	≤8	N/S
	Percentile (%)	99.9	99	N/S
	Δt (seconds)	1	1	N/S
Frame Loss Ratio	Time Period (T)	1 month	1 month	1 month
	Objective (%)	≤.01	≤.01	≤.1

Table 5: Performance Set 2 Recommendations for Green (CIR conformant) frames

Performance Set 3 – Mean Frame Delay (MFD) oriented SLS, with FDR and FLR				
Performance Metric	Parameters	Class of Service		
		H	M	L
Mean Frame Delay	Time Period (T)	1 month	1 month	1 month
	Objective (ms)	≤7	≤13	≤28
Frame Delay Range	Time Period (T)	1 month	1 month	N/S
	Objective (ms)	≤5	≤10	N/S
	Percentile (%)	99.9	99	N/S
Frame Loss Ratio	Time Period (T)	1 month	1 month	1 month
	Objective (%)	≤.01	≤.01	≤.1

Table 6: Performance Set 3 Recommendations for Green (CIR conformant) frames

4.4 Performance Monitoring

The Backhaul Provider should implement Service OAM Fault Management (SOAM-FM) per MEF 30 [8], and Service OAM Performance Monitoring (SOAM-PM) per MEF 35 [9]. Services with EIR only do not require PM. MEF 30 specifies the SOAM infrastructure that is required by SOAM-PM, as well. Among other aspects, MEF 30 defines the Maintenance Entity Group (MEG) levels, the positioning and orientation of the Maintenance End Points (MEPs). MEF 35 defines the PM solutions that use the appropriate MEPs and MEG levels.

PM Solution 1 uses synthetic test frames for measuring the metrics (delay, loss, availability) from a single end. The protocols recommended are Delay Measurement Message and Response (DMM/DMR) version 1, and Synthetic Loss Measurement and Response (SLM/SLR). These are specified in ITU-T Y.1731, Rev 1 [10]. A separate instance of DMM version 1 (for delay measurements) and SLM (for loss measurements) are required to perform simultaneous PM on different CoS within the same EVC.

Performance Monitoring should have the widest scope possible. For Mobile Backhaul applications, the EVC MEPs used for PM by the Backhaul Provider are normally located in network devices at the Mobile Operator premises (cell sites and aggregation sites).

Furthermore, the Mobile Operator and Backhaul Provider should measure performance according to the SLS definitions, to ensure that both the Mobile Operator and the Backhaul Provider are in agreement with results of the measurements.

Service OAM PM Recommendations:

[RC12] It is recommended that the Backhaul Provider use an EVC MEG for performance monitoring of the service.

Note: This is consistent with MEF recommendations as with [RC8]. This approach matches the reach of the entire EVC which may not be the case compared to the placement of an Operator MEG. Possible consequences of not implementing this recommendation could be non-standard monitoring functionality making discussions about actual performance problematic.

[RC13] It is recommended that the EVC MEG MEP be implemented in the Network Element hosting the UNI, i.e., at the Mobile Operator site, for widest scope possible.

Note: Performance Monitoring by the Backhaul Provider will have the widest scope possible (UNI to UNI), and so results will more closely align with the backhaul performance experienced by the Mobile Operator. Possible consequences of not implementing this recommendation could be monitoring scope differences leading to a wider variance in measured values for the metrics.

[RC14] It is recommended that the Subscriber MEG MEP be implemented in the UNI-C function of the Mobile Operator equipment (MEG level 6).

Note: This creates a close correlation between Mobile Operator and Backhaul Provider performance monitoring of the service, since the monitoring points are located close to each other (only separated by the UNI link). Possible consequences of not implementing this recommendation could be Mobile Operator inability to monitor with the same methodology as the Backhaul Provider, with possibly a wider scope (if test probes are located further away from the UNI) leading to a wider variance in measured values for the metrics.

[RC15] It is recommended that proactive PM be implemented, per MEF 35, using PM Solution 1.

Note: This widely supported standard promotes consistent measurements that monitor performance of all metrics in the SLS. Possible consequences of not implementing this recommendation could be more complicated and inconsistent methodology for monitoring the metrics specified in the SLS.

[RC16] It is recommended that proactive PM be implemented for any CoS Frame Set with CIR > 0.

Note: This ensures that transient issues can be tracked that could be missed with reactive monitoring. Possible consequences of not implementing this recommendation could be inability to continuously monitor performance for a CoS that requires performance guarantees.

This means that, at a minimum, CoS Labels 'H' and 'M' will be continuously monitored. If CoS Label 'L' has CIR >0, it also will be continuously monitored.

4.5 Oversubscription

Oversubscription policies vary by Backhaul Provider, and therefore the Mobile Operators should inquire as to what policies are implemented.

5 What Mobile Operators need for Multi-CoS Backhaul

This section briefly lists some key requirements on the equipment the Mobile Operator uses to connect to the Carrier Ethernet Mobile Backhaul service. Cell site and Aggregation site devices that connect to the CEN should have the following functionality:

- Support for a MEF compliant UNI, as specified in MEF 13 [7] and MEF 20 [11]. These UNI Implementation Agreements essentially allow support for any of the Ethernet PHYs specified in IEEE 802.3-2008 [12], except for the Ethernet PON PHYs.
- Ability to classify traffic by Class of Service, with ability to mark the Ethernet VLAN tag PCP value on service frames heading into the CEN appropriately for the multi-CoS policy. For example, a service frame with a PCP value of 5 as marked by the Mobile Operator will be classified into the CEN H CoS.
- Ability to shape traffic per {EVC, CoS} to conform to the Ingress Bandwidth Profile of each CoS.
- Ability to monitor performance for each CoS frame set {EVC, CoS} in accordance with metrics and methodologies described in sections 4.3 and 4.4 of this document.

6 Challenges

There are some challenges in the deployment of successful multi-CoS Mobile Backhaul solutions. These include:

- Simplicity in using multi-CoS. It may help to agree on a standardized set of 'packages' to help simplify the decision process. One example is a three CoS Label {H, M, L} package, with 10% of aggregate CIR in H CoS, 20% in M CoS, and 70% in L CoS. Each Backhaul Provider can develop its own package. This could be a topic for a future revision of this Best Practices document.
- Working out the SLA's for each CoS. Getting agreement on the details and nuances of a single CoS has been quite challenging, sometimes taking months to agree on metric definitions used in the SLS, and on measurement methodology. For multi-CoS, the effort is even more. If the Backhaul Providers and Mobile Operators can agree on the details as recommended in this document, it should help in coming to closure on the SLS more quickly and in a standardized way.
- Flexibility in changing the bandwidth requirements per cell site. It is obvious that not all cell sites peak at the same time. If Mobile Operators had the ability to change bandwidth per day (or per hour) on a given EVC, they could have more control over their backhaul network and potentially lower their backhaul costs and improve performance. MEF is currently working on a Dynamic

Responsive Ethernet (DRE) project that may be useful in the future for Mobile Backhaul scenarios.

- Close cooperation between Mobile Operator and Backhaul Providers. It's important for the Backhaul Provider to provide the assurance to the Mobile Operator that the network will meet their reliability and performance requirements. This can be accomplished through a combination of: MEF Performance Sets as recommended in this document; emerging new MEF metrics and performance objectives (e.g., Availability and Resiliency); and, where these metrics are insufficient, an understanding of how the Backhaul Provider builds in the needed performance and reliability into his network.
- Engineering of the EVCs, collecting performance and utilization stats, managing backhaul bandwidth together.
- Connecting to multiple Backhaul Providers in a given metro or regional area. An Ethernet Exchange model may help simplify backhaul in cases where Mobile Operators are connecting to several, smaller Backhaul Providers, and the backhaul services are not yet standardized.
- A multipoint EVC service model supporting LTE X2 interface, and maybe others services, is still under development. A recommendation may be made in a future phase.

7 Conclusions

This Best Practices document makes recommendations to the Backhaul Provider and the Mobile Operator for standardizing multi-CoS arrangements for Mobile Backhaul.

8 Glossary and Terms

This section defines some of the key MEF terms used in this document, and provides a reference to the appropriate MEF specification for more detailed definitions.

Term	Definition	Reference
Availability	A measure of the percentage of time that a service is useable	MEF 10.2 [5]
Bandwidth Profile	A Bandwidth Profile is a method characterizing Service Frames for the purpose of rate enforcement or policing.	MEF 10.2 [5]
CEN	Carrier Ethernet Network	
Backhaul Provider	Service provider who provides the enabling wire-line and wireless access service on behalf of the Mobile Operator	
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.	MEF 23.1 [3]
Class of Service Identifier (CoS ID)	Information derivable from the: a) EVC to which the Service Frame is mapped b) combination of the EVC to which the Service Frame is mapped and a set of one or more CE-VLAN CoS values, c) combination of the EVC to which the Service Frame is mapped and a set of one or more DSCP values, or d) combination of the EVC to which the Service Frame is mapped and a set of one or more tunneled Layer 2 Control Protocols.	MEF 10.2 [5]

Term	Definition	Reference
Class of Service Label	A CoS Name that is standardized by MEF. MEF 23.1 defines the following three CoS Labels: 'H', 'M' and 'L'. Each CoS Label identifies four Performance Tiers, where each Performance Tier contains a set of performance objectives and associated parameters.	MEF 23.1 [3]
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.	MEF 10.2 [5]
Color-aware	A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each Service Frame, indicated by the Color Identifier, is taken into account when determining the level of compliance for each Service Frame.	MEF 10.2 [5]
Color-blind	A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each Service Frame, if present, is ignored when determining the level of compliance for each Service Frame.	MEF 10.2 [5]
Committed Burst Size (CBS)	CBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of Service Frames sent at the UNI speed to remain CIR conformant.	MEF 10.2 [5]
Committed Information Rate (CIR)	CIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of Service Frames up to which the network delivers Service Frames and meets the performance objectives defined by the CoS Service Attribute	MEF 10.2 [5]
Ethernet Virtual Connection (EVC)	An association of two or more UNIs that limits the exchange of Service Frames to UNIs in the Ethernet Virtual Connection.	MEF 10.2 [5]
Excess Burst Size (EBS)	EBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of Service Frames sent at the UNI speed to remain EIR conformant.	MEF 10.2 [5]
Excess Information Rate (EIR)	EIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of Service Frames up to which the network delivers Service Frames, but without any performance objectives.	MEF 10.2 [5]
Frame Delay (FD)	The time required to transmit a Service Frame from ingress UNI to egress UNI.	MEF 10.2 [5]
Frame Delay Range (FDR)	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.	MEF 10.2 [5] MEF 23.1 [3]
Frame Loss Ratio (FLR)	Frame Loss Ratio is a characterization of the number of lost Service Frames between the ingress and the egress UNI. FLR is expressed as a percentage.	MEF 10.2 [5]
Ingress Bandwidth Profile	A characterization of ingress Service Frame arrival times and lengths at the ingress UNI and a specification of disposition of each Service Frame based on its level of compliance with the characterization.	MEF 10.2 [5]

Term	Definition	Reference
Inter-Frame Delay Variation (IFDV)	The difference in delay of two Service Frames of the same CoS Frame Set.	MEF 10.2 [5]
Mean Frame Delay (MFD)	The arithmetic mean, or average of delays experienced by Service Frames belonging to the same CoS Frame Set.	MEF 10.2 [5]
Maintenance End Point (MEP)	Maintenance association End Point, or equivalently, MEG end point. An actively managed SOAM entity associated with a specific service instance that can generate and receive SOAM PDUs and track any responses. It is an end point of a single MEG, and is an endpoint of a separate Maintenance Entity for each of the other MEPs in the same MEG.	IEEE 802.1Q [13] ITU-T Y.1731 [10]
Maintenance Entity (ME)	A point-to-point relationship between two MEPs within a single MEG.	IEEE 802.1Q [13]
Maintenance Entity Group (MEG)	A set of MEs that exist in the same administrative boundary, with the same MEG Level and MEG ID.	ITU-T Y.1731 [10]
MEG Level	A small integer in a field in a SOAM PDU that is used, along with the VID in the VLAN tag, to identify to which Maintenance Association among those associated with the SOAM frame's VID, and thus to which ME, a SOAM PDU belongs. The MEG Level determines the MPs a) that are interested in the contents of a SOAM PDU, and b) through which the frame carrying that SOAM PDU is allowed to pass. This term is equivalent to MD Level, which is used in [IEEE 802.1Q-2011].	ITU-T Y.1731 [10]
Mobile Operator	Service Provider who is responsible for and has a service agreement with the end user customer (whether consumer or Enterprise buyer)	
Performance Monitoring (PM)	Performance Monitoring involves the collection of data concerning the performance of the network.	MEF 35 [9]
PM Function	A MEP capability specified for performance monitoring purposes (e.g., Single-Ended Delay, Single-Ended Synthetic Loss).	MEF 35 [9]
PM Solution	A PM Solution is a set of related requirements that, when implemented, allow a given set of performance metrics to be measured using a given set of PM functions.	MEF 35 [9]
Priority Code Point (PCP)	A three-bit field in the VLAN Tag header that indicates priority of the Ethernet frame.	IEEE 802.1Q [13]
Proactive	OAM actions that are carried on continuously to permit timely reporting of fault and/or performance status.	MEF 35 [9]
Service Frame	An Ethernet frame transmitted across the UNI toward the Service Provider or an Ethernet frame transmitted across the UNI toward the Subscriber.	MEF 10.2 [5]
Service Level Specification (SLS)	The technical specification of the service level being offered by the Service Provider to the Subscriber.	MEF 10.2 [5]
Service OAM (SOAM)	Service Operations, Administration, and Maintenance	MEF 17 [14]

Term	Definition	Reference
Switch	A network device used in the Carrier Ethernet network that forwards traffic for a given Ethernet Virtual Connection (EVC) onto the appropriate link(s) to reach the desired destination(s).	This document
User Network Interface (UNI)	The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber.	MEF 10.2 [5]

9 References

[1] “The Benefits of Multiple Classes of Service for Ethernet Mobile Backhaul”, MEF, January 2012

[2] MEF 22.1, “Mobile Backhaul, Phase 2, Implementation Agreement”, January 2012

[3] MEF 23.1, “Carrier Ethernet Classes of Service, Phase 2, Implementation Agreement”, January 2012

[4] 3GPP, TS 23.203

[5] MEF 10.2, “Ethernet Service Attributes, Phase 2”, October 2009

[6] MEF 10.2.1, “Performance attributes Amendment to MEF 10.2”, January 2011

[7] MEF 13, “User Network Interface (UNI) Type 1 Implementation Agreement, November 2005

[8] MEF 30, “Service OAM Fault Management Implementation Agreement”, January 2011

[9] MEF 35, “Service OAM Performance Monitoring Implementation Agreement”, April 2012

[10] ITU-T G.8013/Y.1731, “OAM functions and mechanisms for Ethernet based networks”, July 2011

[11] MEF 20, “User Network Interface (UNI) Type 2 Implementation Agreement, July 2008

[12] IEEE 802.3-2008, “Part 3: Carrier sense multiple access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications”, December 2008

[13] IEEE 802.1Q-2011, “Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks”, August 2011

[14] MEF 17, “Service OAM Requirements & Framework – Phase 1”, April 2007

10 About the MEF

The MEF is a global industry alliance comprising of more than 210+ companies, including telecommunications service providers, cable MSOs, network equipment/software manufacturers, systems integrators and testing organizations. The MEF develops technical specifications and implementation agreements to promote interoperability and deployment of Carrier Ethernet worldwide. All activities are member lead supported by MEF staff. The MEF has driven the record-setting double digit growth of Carrier Ethernet worldwide since its inception in 2004. Membership growth continues at record levels. For more information visit www.metroethernetforum.org for more information.

11 Acknowledgements

Editor: Bill Bjorkman, MEF
Contributors: Bill Rembert, AT&T, Carsten Rossenhoevel, EANTC, Chris Purdy, CENX, Dudu Bercovich, Ceragon, Ghassan Semaan, Verizon, Mark Fishburn, MEF, Phil Tilley, Alcatel-Lucent, Raghu Ranganathan, Ciena
Ralph Santitoro, Fujitsu, Rami Yaron, Dragon Wave, Scott Mansfield, Ericsson, Tim Pearson, Sprint, Yoav Cohen, RAD