Understanding Carrier Ethernet Service Assurance

Taking advantage of MEF's best practices for performance monitoring

— Part II —

How to Measure Carrier Ethernet Performance Attributes

September 2016

Table of Contents

List of Figures

List of Tables

1. Introduction and Overview

1.1 Abstract

The MEF has addressed multiple aspects of service assurance throughout its technical specifications to enable the necessary performance and service quality available with Carrier Ethernet 2.0 (CE 2.0) services.

The MEF's Service Management Life Cycle white paper describes the different stages in the life cycle of a CE 2.0 based service including performance monitoring and fault management. This paper together with its companion paper, "***Part I: An Introduction to Service Assurance and Carrier Ethernet Service Performance Attributes"***, provides an in-depth guide to performance monitoring and assurance across the body of work of the MEF enabling the reader an accessible, one-stop guide to understanding CE 2.0 performance monitoring best practices and use them to their full advantage.

1.2 Target audience

Communication Service Providers (CSPs) including telecom service providers, Internet service provider (ISP), cable operators/MSOs, cloud service providers, wireless network operators, network equipment providers, and OSS/LSO providers.

1.3 Document Purpose and Scope

This paper explores the details of how to measure and monitor performance in the context of CE 2.0 services, the Third Network vision and Lifecycle Service Orchestration (LSO).

Those familiar with some of the key technical specifications such as MEF 35.1 (Service OAM Performance Monitoring Implementation Agreement) as well as those unfamiliar with the work of the MEF or its many technical specifications will benefit from this holistic view covering the current MEF body of work with respect to the critical and complex topic of service assurance.

1.4 Executive Summary

Service assurance encompasses the procedures used to monitor a network service to maintain a particular service quality. This paper focuses on the practical considerations and methods for measuring Carrier Ethernet service performance. The reader is encouraged to read the companion paper, "***Part I: An Introduction to Service Assurance and Carrier Ethernet Service Performance Attributes"***, where the service performance attributes are defined.

Based on the relevant business drivers and performance attributes, the reader will benefit from understanding how to measure the service quality of Carrier Ethernet services based on MEF standards. There exists significant confusion and misconceptions with respect to performance management best practices for Carrier Ethernet services. This confusion is the result of multiple transport technologies, legacy tools, proprietary instrumentation, interoperability challenges as well as other related service assurance standards making it difficult to understand and leverage the best practices available to Carrier Ethernet service providers. By clarifying the what and the how, the industry can more quickly align to simplify management, accelerate delivery and turn-up of existing services as well as new service launches which together improve the critical time-to-revenue factor for providers.

Having the performance-oriented service assurance elements which are specified in over a dozen technical specifications centralized and accessible to the reader enables a more efficient and rapid implementation of the industry best practices for Carrier Ethernet 2.0 service assurance from the definition of service quality objectives, to the proper monitoring, presentation and sharing of the resulting service performance intelligence.

2. Carrier Ethernet Service Performance

The companion to this paper, ***"Part I: An Introduction to Service Assurance and Carrier Ethernet Service Performance Attributes"***, examined the defined Carrier Ethernet Service Performance Attributes. This paper defines how those attributes are measured.

2.1 Performance Attributes – How to Measure

2.1.1 Service OAM Overview

The Service OAM (operations administration and maintenance) framework includes the fault management, performance monitoring, auto-discovery and both intra- and inter-provider service OAM. It focuses on the Ethernet Services layer (ETH Layer) as opposed to the transport service layer (e.g. MPLS or SONET/SDH) or the Application Services Layer (e.g. IP, not to be confused with the subscriber's definition of the application such as an enterprise resource planning application).

Service OAM works in conjunction with both the control plane and management plane in an effort to measure performance and to monitor the data plane that carries traffic. The Service OAM framework defines the points in the network that must be maintained and groups them into Maintenance Entity Groups (MEGs). These groups contain end points (MEPs) and intermediate points (MIPs). Performance monitoring takes place at the MEP. Furthermore, it defines the protocol data units (PDUs) used for SOAM traffic and how that is handled through the CEN.

To handle the different administrative layers of a network service (e.g. the last mile connection, an access provider segment in a multi-operator service, the service provider segment and the subscriber segment), MEGs are defined with a level of 0 to 7 so that lower level MEGs allow SOAM traffic from higher level MEGs to pass through unaffected. For example, the customer domain MEG below would be defined as MEG level 6. The operator domains (in orange) could be MEG level 2 and that level 2 domain would pass the MEG level 6 SOAM traffic through unaffected so that metrics for the customer UNI to UNI performance attributes could be measured.

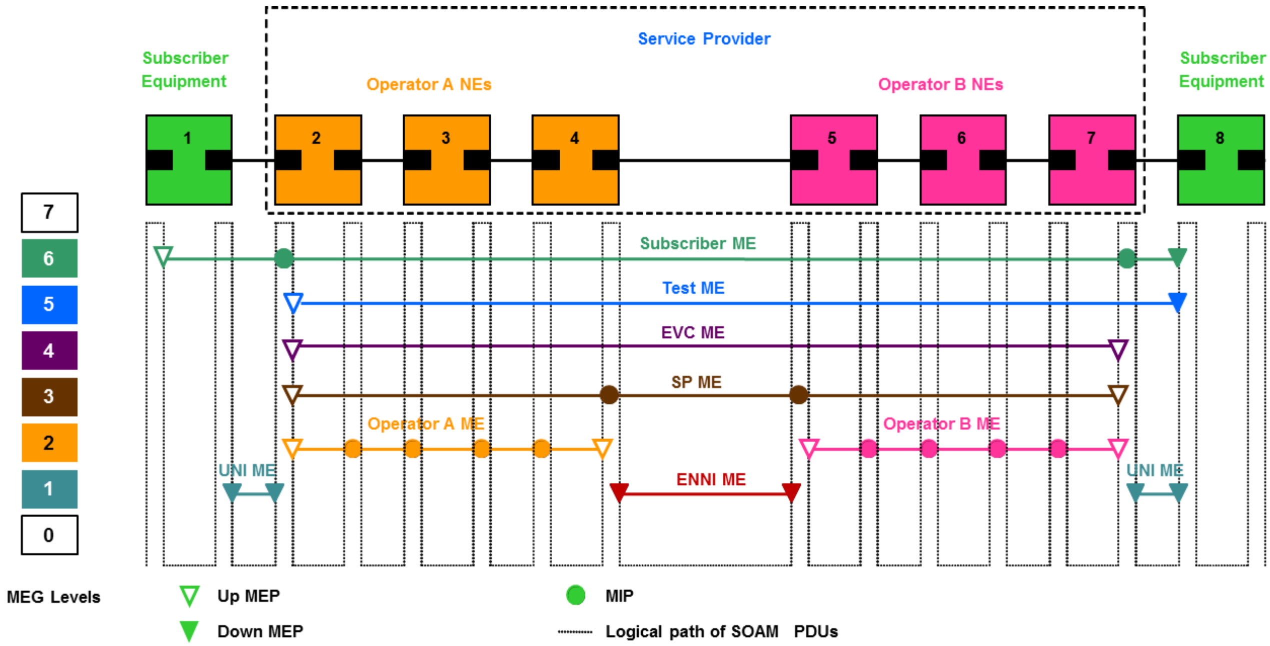

Figure 1 – Diagram of Example MEG Levels

2.1.2 MEF Service OAM and ITU-T G.8013/Y.1731

To help provide guidance on how to best implement the SOAM performance monitoring for CE services, the MEF created an implementation agreement (MEF 35.1). The ITU-T ***Recommendation G.8013/Y.1731*[1] *- OAM Functions and mechanisms for Ethernet-based networks***, (colloquially referred to as "Y.1731") provides the instrumentation to calculate each of the performance metrics that make up the SLS. Some confusion exists in the marketplace because of the fact that Y.1731 is partially implemented by many equipment vendors leading some operators to believe they are using MEF-defined performance monitoring, however though the MEF takes advantage of the various PM functions within Y.1731, the actual performance metrics and requirements for collection, measurement and storage are uniquely defined in MEF 35.1.

2.1.3 Performance Monitoring (PM) Measurements

The MEP (Maintenance Entity Group End Point) of an associated SOAM MEG level is able to both generate and respond to SOAM PDUs. The MIP (Maintenance Entity Group Intermediate Point) can only respond and is not used for performance monitoring functions.

The first key concept is that all measurements are performed at the network element layer (i.e. within the network equipment). The measurements, known officially as PM measurements may then be sent northbound to the Element management layer to an EMS (element management system) and in turn to the network management layer to an NMS (network management system). Alternatively, the PM measurements can be sent directly to the NMS (or queried/polled by the NMS). Given the recent work

[1] As of July 2011, the ITU-T Y.1731 recommendation was renamed G.8013/Y.1731; https://www.itu.int/rec/T-REC-Y.1731

of the MEF on LSO, the legacy terms EMS and NMS are beginning to transition to Element Control and Management (ECM) and Infrastructure Control and Management (ICM) respectively.

The MEF defines four performance monitoring (PM) solutions which refer to particular MEG types (point-to-point or multi-point), measurement techniques for loss measurements and the PM function(s) defined to generate the appropriate metrics for the solution. These solutions are outlined in the table below, where PM Solution 1 (PM-1) is mandatory, the others are optional.

Table 1 - PM Solutions Summary

PM Solution	MEG Type(s)	Measurement Technique for Loss	PM Function(s)	Mandatory or Optional
PM-1	point-to-point multipoint	Synthetic Testing	Single-Ended Delay Single-Ended Synthetic Loss	Mandatory
PM-2	point-to-point multipoint	n/a	Dual-Ended Delay	Optional
PM-3	point-to-point	Counting Service Frames	Single-Ended Service Loss	Optional
PM-4	point-to-point multipoint	Synthetic Testing	Dual-Ended Synthetic Loss	Optional

The PM-1 mandatory solution applies to all EVC (and OVC) service types and uses synthetic tests to measure loss. Synthetic testing uses test traffic to measure loss versus actually measuring real frame loss between two UNIs. Synthetic testing can provide visibility on network performance even when the network is not in use, helping establish early signs of performance degradation. In addition, the synthetic test functions are generally simpler computationally, however they add traffic to the network and ultimately do not exactly reflect real frame loss for qualified service frames.

2.1.4 One-way and two-way measurements versus single-ended and dual-ended messaging

The familiar TCP/IP tool: ICMP Ping enables a user to initiate a "Ping" test from a source IP address to a destination IP address and the command returns the round trip time for the ICMP packet to traverse from the source to the destination and back again. This common example may help in visualizing the terminology used throughout the MEF specifications with respect to the PM tools (i.e. those borrowed from Y.1731) and the PM Functions described in **section 2.1.5**.

2.1.4.1 One-way vs. Two-way Measurements

There is a fundamental difference between a one-way measurement and a two-way measurement. Given two end points (A and Z), a one-way measurement will measure the value (e.g. delay) going from one end point to the other. For example, the one-way delay between A and Z can be measured from A to Z, and separately from Z to A. A one-way delay measurement for instance, is very useful as opposed to the round trip time common to a "Ping" test, given that the A-to-Z value is measured distinctly from the Z-to-A value which may demonstrate asynchronous results. In fact, the ITU-T defines one-way in the context of Ethernet frames in G.8001 as a unidirectional measurement. All of the MEF performance attributes are defined as one-way.

Conversely a two-way measurement is analogous to the "ping" test and is defined once again by ITU-T as a bidirectional or round-trip measurement. There are no round-trip-based performance attributes within the MEF specifications, however the round-trip frame delay can be used to estimate the one-way delay by assuming synchronous delay and dividing the round-trip result by two to approximate the one-way delay measurement. However, to achieve an accurate one-way measurement for delay, both the source (controller MEP) and destination (sink MEP) must have their clocks synchronized. Clock synchronization is not required for measuring one-way inter frame delay variation (IFDV) nor one-way frame delay range (FDR).

2.1.4.2 Single-ended vs. Dual-ended Messaging

ITU-T uses the terms single-ended and dual-ended for defining their PM tools. These terms differ slightly from the terms one-way and two-way measurements. The primary difference is that a measurement is a specific test result from the PM tool(s), that is, one-way delay measures the delay in one-direction from a source to a destination. The terms single-ended and dual-ended are not used to describe the measurement but rather how the measurement is achieved, specifically the messaging protocol.

For example, the ITU-T defines single-ended delay as a tool with a controller MEP that sends a particular message to a responder MEP that in turn replies with a related message back to the controller. This sounds a lot like the "Ping" test, however it isn't directly related to a measurement, it merely indicates that one MEP, the controller MEP, initiates the tool (sends the initial PDUs) and performs the calculations once the test is complete. The responder MEP doesn't get involved in any calculation in this instance, hence it is referred to as a 'single-ended' test. That test can provide a two-way measurement as well as a one-way measurement, however for the one-way measurement to be accurate in the case of a delay measurement, the two MEPs must have their clocks synchronized.

Conversely a dual-ended test requires that the controller MEP (initiating) sends the request PDU to the sink MEP (destination) and the sink MEP performs all of the calculations. There is no response or reply back to the controller MEP. In this case only a one-way measurement can be calculated, although academically speaking, a round-trip delay time could be estimated by doubling the one-way measurement but this is not done in practice, nor is it relevant to the MEF performance attributes which are all defined as one-way.

Typically, the choice to use single-ended versus dual-ended is based on a CSP's operational preferences. For example, dual-ended messaging means that the PM tools can be configured symmetrically on all devices. Furthermore, in multi-point networks, the test PDUs can be multicast to make more efficient use of the network. Otherwise, with the exception that dual-ended tests cannot provide two-way measurements, *single-ended and dual-ended messaging are equivalent* in terms of the measurements they can provide.

> ***Note****: The terms sink MEP and responder MEP are similar in that the controller MEP (initiating the tests) sends the test PDUs to the sink/responder, however there is a distinction in that a sink MEP is exclusive to dual-ended tests and there is no reply/response. A responder MEP on the other hand is exclusive to single-ended tests and responds to the controller MEP.*

2.1.5 Performance Monitoring (PM) Functions and Solutions

There are 4 PM Solutions defined which make use of five PM functions, single-ended delay, single-ended synthetic loss, dual-ended delay, single-ended service loss and dual-ended synthetic loss. PM-1 Solution is mandatory. PM-2, PM-3 and PM-4 Solutions are optional.

The relationship between each performance attribute and the associated PM functions which can provide measurements for the attribute is illustrated in the table below.

Table 2 - Performance Attributes and Relevant PM Functions

Performance Attribute	PM Function
One-way Frame Delay Performance	Single-ended delay (PM-1 Solution) Dual-ended delay (PM-2 Solution)
One-way Mean Frame Delay Performance	Single-ended delay (PM-1 Solution) Dual-ended delay (PM-2 Solution)
One-way Frame Delay Range Performance	Single-ended delay (PM-1 Solution) Dual-ended delay (PM-2 Solution)
One-way Inter-Frame Delay Variation Performance	Single-ended delay (PM-1 Solution) Dual-ended delay (PM-2 Solution)
One-way Frame Loss Ratio Performance	Single-ended synthetic loss (PM-1 Solution) Single-ended service loss (PM-3 Solution, point-to-point EVC/OVC only) Dual-ended synthetic loss (PM-4 Solution)
One-way Availability Performance	Single-ended synthetic loss (PM-1 Solution) Dual-ended synthetic loss (PM-4 Solution)
One-way Resiliency Performance expressed as High Loss Intervals	Single-ended synthetic loss (PM-1 Solution) Dual-ended synthetic loss (PM-4 Solution)
One-way Resiliency Performance expressed as Consecutive High Loss Intervals	Single-ended synthetic loss (PM-1 Solution) Dual-ended synthetic loss (PM-4 Solution)
One-way Group Availability Performance	Single-ended synthetic loss (PM-1 Solution) Dual-ended synthetic loss (PM-4 Solution)
Composite Performance Metric (CPM)	Either single-ended delay or dual-ended delay coupled with either single-ended synthetic loss, dual-ended synthetic loss or in the case of point-to-point EVC/OVC, single-ended service loss

The PM functions are described below, organized by their relevant PM solution.

2.1.5.1 PM-1 Solution (for point-to-point and multipoint services)

Single-ended delay

This is based on initiating a test from a controller MEP to a responder MEP using ITU-T's single-ended ETH-DM tool defined in Y.1731. The PDUs used are DMM (delay measurement message) and the DMR (delay measurement response).

> ***Note**: Single-ended delay tests need clock synchronization between the controller and responder MEPs when used to provide accurate one-way Frame Delay Performance measurements for the purpose of calculating FD and MFD performance metrics. Otherwise, depending on the discrepancy in time-of-day between the two MEPS, the FD (and MFD) metrics will be suspect. Alternatively, in the absence of clock synchronization, these tests can provide two-way Frame Delay Performance measurements which can be used to estimate the one-way Frame Delay Performance.*

Single-ended synthetic loss

This is based on initiating a test from a controller MEP to a responder MEP using ITU-T's single-ended ETH-SLM tool defined in Y.1731. The PDUs used are SLM (synthetic loss message) and SLR (synthetic loss response). Synthetic loss is easily defined but may not represent the actual behavior experienced by real qualified service frames. In addition, a statistically significant quantity of tests is necessary to provide a reasonable estimate of frame loss.

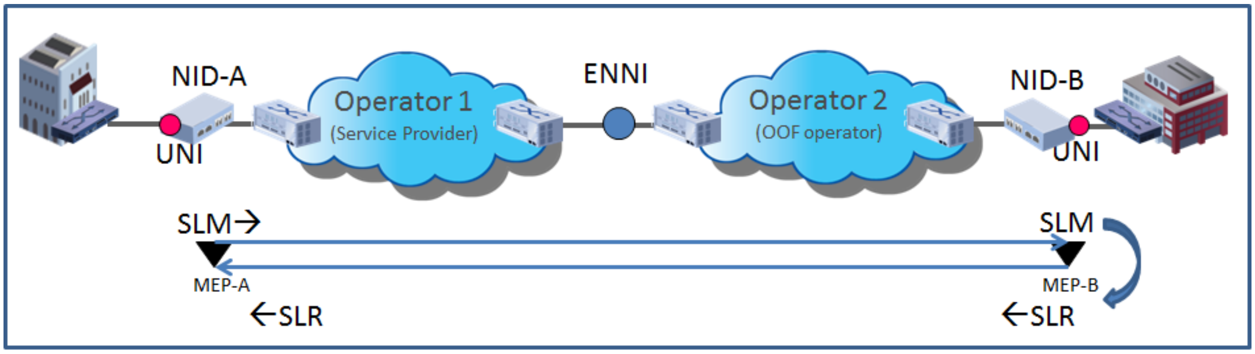

Figure 2 – Single-ended Synthetic Loss Measurement

Single-ended synthetic loss measurement example

In this example an SLM message is created at the controller MEP (MEP-A) which includes the CoS ID, and the local counter value containing the number of SLM messages sent. The SLR message is generated in response at the responder MEP (MEP-B) which includes the count of received messages. MEP-A then calculate the one-way frame loss performance based on the data in the SLR message.

Table 3 - PM-1 Solution's Supported Performance Attributes

	Supported Performance Attribute
PM-1 Solution	One-way Frame Delay Performance
	One-way Mean Frame Delay Performance
	One-way Frame Delay Range Performance
	One-way Inter-Frame Delay Variation Performance
	One-way Frame Loss Ratio Performance
	One-way Availability Performance
	One-way Resiliency Performance expressed as High Loss Intervals
	One-way Resiliency Performance expressed as Consecutive High Loss Intervals
	One-way Group Availability Performance

2.1.5.2 PM-2 Solution (for point-to-point and multipoint services)

Dual-ended delay

This is based on initiating a test from a controller MEP to a sink MEP using ITU-T's dual-ended ETH-DM tool defined in Y.1731. The PDU used is 1DM (single delay measurement).

> ***Note**: Dual-ended delay tests need clock synchronization between the controller and sink MEPs when used to provide accurate one-way Frame Delay Performance measurements for the purpose of calculating FD and MFD performance metrics. Otherwise, depending on the discrepancy in time-of-day between the two MEPS, the*

FD (and MFD) metrics will be suspect. Dual-ended delay tests cannot provide two-way Frame Delay measurements.

This PM solution is optional and applies to both point-to-point and multi-point services. It does not include any loss metrics, but uses the dual-ended delay PM function providing measurements for FD, MFD and FDR as well as IFDV based on the 1DM PDU. Note that 1DM PDUs can be multicast for efficient performance monitoring of multi-point services such that a single controller MEP can send the PDU to multiple sink MEPs.

Table 4 - PM-2 Solution's Supported Performance Attributes

PM-2 Solution	Supported Performance Attribute
	One-way Frame Delay Performance
	One-way Frame Delay Range Performance
	One-way Mean Frame Delay Performance
	One-way Inter-Frame Delay Variation Performance

2.1.5.3 PM-3 Solution (for point-to-point services)

Single-ended service loss

This is based on initiating a test from a controller MEP to a responder MEP using ITU-T's single-ended ETH-LM tool defined in Y.1731. The PDUs used are LMM (loss measurement message) and LMR (loss measurement response). This tool can provide the most accurate measurement of **real** frame loss between 2 UNIs as it is derived from **real** subscriber traffic. However, during low periods of usage, aggregate metrics (averages, peaks, etc.) may not reflect the actual performance of the service as compared to synthetic measurements with a statistically significant amount of test traffic. This solution can only be applied to a point-to-point EVC (or OVC) since all ingress service frames are destined for the other UNI as opposed to a multipoint service where ingress service frames may be destined for different UNIs.

This PM solution is optional and applies only to point-to-point services given that it counts the actual service frames that egress from the responder MEP and compares this to the number of frames sent at ingress to the controller MEP.

Table 5 - PM-3 Solution's Supported Performance Attributes

PM-3 Solution	Supported Performance Attribute (Point-to-point EVC/OVC only)
	One-way Frame Loss Ratio Performance[2]

[2] FLR is used in the calculation for resiliency (HLI, CHLI) and Availability, thus the reader may question why these attributes are not supported by PM-3 Solution along with FLR. The primary reason is that the One-way Availability Performance attribute is meant to provide continuous, proactive monitoring. Because PM-3 solution requires traffic on the network to take a measurement, there is no way to measure availability if the service is idle (no traffic), therefore PM-3 can only be used for one-way Frame Loss Ratio performance.

2.1.5.4 PM-4 Solution (for point-to-point and multipoint services)

Dual-ended synthetic loss

This is based on initiating a test from a controller MEP to a sink MEP using ITU-T's single-ended ETH-SLM tool defined in Y.1731. The PDU used is 1SL (single synthetic loss). Synthetic loss is easily defined but may not represent the actual behavior experienced by real qualified service frames. In addition, a statistically significant quantity of tests is necessary to have a reasonable estimate of frame loss.

This PM solution is optional and applies to both point-to-point and multipoint services. It is based strictly on dual-ended synthetic loss measurements, that is, true one-way measurements from a controller MEP to one or more sink MEPs. The 1SL PDUs can be unicast or multicast for efficient performance monitoring of multipoint services.

Table 6 - PM-4 Solution's Supported Performance Attributes

	Supported Performance Attribute
PM-4 Solution	One-way Frame Loss Ratio Performance
	One-way Availability Performance
	One-way Resiliency Performance expressed as High Loss Intervals
	One-way Resiliency Performance expressed as Consecutive High Loss Intervals
	One-way Group Availability Performance

2.1.6 PM Sessions and Measurement Intervals

A performance monitoring session or PM session is a term used to describe the specific PM solution (and its PM function), such as single-ended delay between two MEPs for a specific CoS. Performance measurements are partitioned into fixed-length Measurement Intervals. The SLS objectives are defined over a period of time T. Typically T is made up of multiple measurement intervals.

The MEF does not dictate the specific measurement intervals that should be used by the industry. Such decisions must be made by each CPS, based on the specific service, its underlying transport technology and available instrumentation as well as the market conditions that govern expected or common SLAs which subscribers expect.

2.1.7 Proactive versus on-demand monitoring

These PM tools can be initiated through a PM session on an on-going basis which is generally the assumed method for providing SLS compliance reporting and more broadly SLA and performance reporting to both internal stakeholders such as engineering, network operations and service delivery as well as external stakeholders, namely the subscriber.

Proactive on-going sessions deliver comprehensive monitoring at the expense of PM session traffic. Therefore, proactive sessions can be supplemented with on-demand sessions to minimize the total amount of PM session traffic. An on-demand PM session could be instantiated for specific testing purposes, or where a performance attribute is not specified as a performance objective within the SLS but the attribute is of interest to engineering or other stakeholders. For example, One-way Inter-Frame Delay Variation may not be specified in the SLS but the CSP may wish to monitor this metric using an on-demand PM session.

2.1.8 Required instrumentation

As defined in MEF 35.1, the calculation of performance metrics is the responsibility of the network element layer. To successfully support the implementation agreement, the network equipment must support the various PM tools based on ITU-T Y.1731 and be able to perform the various calculations for the associated performance metrics. Although not mandatory, the MEF has provided a SNMP-based MIB in MEF specification 36.1 that reflects the requirements defined in the MEF 35.1 implementation agreement. If one's network equipment is not using this MIB, then the defined metrics must be accessed using an EMS or an equipment-vendor-defined method. MEPs must be defined in the network elements.

Typically, that is done in one of the following types of network element:

Switch/router
Often a multi-service device with both layer 2 and layer 3 processing capabilities, these may be used inside the operator's network and in some cases at the customer edge. Due to their greater functional capabilities, these elements are often costlier.

Network Interface Device (NID)
These are special purposed elements designed to terminate CE services and perform the end-to-end performance metric calculations. Some provide both layer 2 and layer 3 termination services. Given their specialized nature, they are often used to easily define all of the UNI attributes of an EVC as well as setup the MEPs. They can be reasonably small and in fact recently many vendors have introduced small-form pluggable (SFP) modules that can be inserted into other network elements helping reduce the power and space requirements.

Test equipment
Various specialized test equipment including handheld and larger devices provide what is typically the most comprehensive set of available performance metric calculations and generally provide excellent support of the MEF defined best practices. However, as they are dedicated to testing, they are not generally available for operational, that is proactive PM.

At minimum, the network equipment needs to support measurement intervals of 15 minutes as well as on-demand durations of 1 minute to 15 minutes in 1 minute intervals. Furthermore, the network equipment must be able to store 8 hours of historic (past) measurement data for each PM session for any on-going (versus on-demand) PM sessions. This history must be stored per measurement interval, that is, it cannot be aggregated. Ideally, the network equipment should be able to store up to 24 hours of historic data.

2.1.8.1 Measurement Bins

To alleviate some storage and additional processing, the use of measurement bins is implemented meaning the individual calculations made are not uniquely stored but are counted as part of a set of intervals using measurement bins. Measurement bins are used to record the results of frame delay and inter-frame delay variation measurements. These results can determine if the performance objectives for FD (frame delay), FDR (frame delay range) and IFDV (inter-frame delay variation) are met. Implementations (for all frame delay performance attributes) must support at least two bins and MEF recommends at least ten.

For example, bins may be created for frame delay with two intervals: (1) between 0ms and 100ms and (2) 100 ms and higher, and a counter is used for each bin to count the number of times a FD measurement was found within the respective interval. Separate bins and associated counters are used for each measurement interval.

2.1.9 Threshold Crossing Alerts (TCAs)

The MEF also defines the ability to generate threshold crossing alerts northbound to the EMS or NMS based on configured threshold levels for a particular PM session. This feature is optional. TCA reporting can be stateless which means each and every time a threshold is crossed at least once during a single measurement interval an alert is generated. Alternatively, TCAs can be stateful such that an alert is received only for the first measurement interval the threshold is crossed and no further alerts will be generated until a measurement interval occurs where the measurement is below the threshold, thus clearing the alert. An alert is received for the 1st measurement interval in a sequence of intervals where the configured threshold was crossed. Note that a change in availability state (from available to unavailable or vice versa) is required functionality for the one-way availability performance attribute and therefore there is no need to use TCA for availability.

2.1.10 Performance Objectives and CoS

Carrier Ethernet support multiple classes of service. The typical classes are H (high), M (medium) and L (low). Performance objectives may be defined for each CoS defined in the EVC, as well as the bandwidth profile making each CoS configurable to the specific needs of the network traffic it needs to carry.

Furthermore, given the technology agnostic approach of CE services, an operator can implement a CE service over various technologies, access speeds and geographic distances. A performance objective that is defined for a particular class of services is known as a CoS Performance objective or CPO. A Performance Tier (PT) is a set of CPOs and may apply to the EVC as well as any OVCs that are part of the EVC. MEF defines four PTs, based primarily on distance between UNIs with an associated one-way mean frame delay.

- PT1 (Metro) <250km, MFD=2ms
- PT2 (Regional) <1200km, MFD=8ms
- PT3 (Continental) <7000km, MFD=44ms
- PT4 (Global) <27500km, MFD=172ms

As a result, each PT has a unique set of CPOs that is reflective of the distance across the network, in particular with respect to delay measurements which are impacted by distance. An operator can alternatively select the appropriate PT based on link speed (i.e. for low speed, copper-based connectivity). Essentially the PTs can be considered a group of CPOs for a particular use case. That use case may be based on distance, link speed, link technology or some other set of parameters. For each PT defined, the MEF provides the respective CPOs which can be used based on the required PT and the specific CoS (H – high, M – medium and L – Low). [please refer to MEF 23.2, Tables 9 through 12].

2.1.11 Practical Considerations and Terminology

2.1.11.1 Terminology Caveats

There is a lot of seemingly synonymous terminology within the realm of performance monitoring, and the MEF specifications in particular. It is therefore helpful to recognize the sometimes subtle yet critical differences in terms. The following table highlights some of these key differences:

Table 7 - Performance Terminology

Term	Description	Notes
Performance *Attribute*[3]	A specific characteristic of a Carrier Ethernet service's performance.	Carrier Ethernet services have a number of defined service attributes including performance attributes. All of the service attributes are specified in MEF 10.3 "Ethernet Service Attributes Phase 3". Service attributes define many aspects of the service such as the type of service, the UNIs involved, bandwidth profiles, etc.
Performance *Metric*[4]	Defined in MEF 10.3 as a quantitative characterization of Service Frame delivery quality.	This term is essentially synonymous with performance attribute in the context of defining the various performance attributes such as one-way Frame Delay Performance. In theory, the performance metric is defined using actual service frames, that is, the real traffic within the Carrier Ethernet service. In practice, many of the performance metrics are calculated using test traffic.
Performance *Objective*	The specific objective related to a performance attribute that is defined in the SLS.	Without an objective defined, a service can still be monitored and the performance measured to provide information on the performance *attributes* however, there is no agreement in place between subscriber and provider to deliver the service according to a particular performance *attribute*(s) unless an *objective* is defined within the SLS.

[3] This paper uses the term 'performance attribute' when referring to the definition of the attribute as opposed to a specific instance of the attribute.
[4] This paper uses the term 'performance metric' when referring to a specific instance of a performance attribute; i.e. the associated parameters are defined and a result is calculated.

Term	Description	Notes
Performance *Measurement*	The measurement of a service's performance. These measurements are used to calculate the performance metric. Many performance metrics require multiple measurements to calculate the performance metric of interest.	In most cases, the measurements used to calculate a performance metric can only approximate the true value of said performance metric. The reason being that most measurements are based on test frames which are generally more practical than measuring actual service frames. In addition, such test measurements can be simpler to implement and in some cases are the only available measurement option for a performance attribute (such as One-way Frame Loss Ratio Performance in a multipoint-to-multipoint network).

2.1.11.2 Definitions and settings

There are many decisions required by the operator configuring performance objectives and defining the SLS. In particular, these include the time period applicable to the performance objective defined within the SLS, e.g. availability must be a minimum of 99.99% over a month long period.

2.1.11.2.1 Defining the SLS

The CSP must define the SLS by identifying the following:

- Which (if any) performance attributes will be used (currently 10 including One-way Group Availability and CPM). If none are used, then the remaining items are irrelevant.
- One (or more) performance objective(s) for each of the chosen performance attributes (recall that different objectives can be defined for the same performance attribute as long as one or more of their parameters differ, such as 80% of all service frames mush have FD ≤ 100ms and 90% of all service frames must have FD ≤ 200ms.)
- A set of UNI pairs within the EVC for which the objective is applied (these are uni-directional and may be a subset of the total permutations of UNI pairs in the EVC)
- The related parameters for the objective(s) including but not limited to:
    - The time period T to use for measuring adherence to each objective.
    - The small time period Δt, for measuring FLR used to determine available time and therefore whether any of the performance objectives are defined, as well as measuring FLR and Availability if they are used in the SLS.
    - The number of consecutive time intervals to use for calculating One-way Availability (n)
    - The FLR that determines unavailability (C)
    - The CoS ID(s) related to the objective
    - The relevant percentile(s) for each objective

2.1.11.2.2 Configuring the performance measurements

Furthermore, the operator must define PDU sizes for the test frames, decide upon which PM solution(s) to use to calculate the performance metric along with the measurement bins and related intervals for bin-based counters. Finally, for all synthetic tests, the operator must determine the number of test PDUs to be sent ensuring a statistically valid quantity, as well as how frequently they are sent. Both of which have strong bearing on the ability to approximate the real service frame performance.

Many of these decisions may in fact be dictated by the available instrumentation (e.g. the capabilities and adherence to standards within the network equipment itself). Many of the considerations and caveats, as well as suggested default values for configuring the performance monitoring are included in the MEF specifications themselves but are outside the scope of this paper.

2.1.11.3 Collection and reporting

After defining the objectives and configuring the network to deliver the necessary instrumentation, the Carrier Ethernet provider must also ensure they have the ability to collect, aggregate and display the resulting performance intelligence to the various stakeholders including network operations (in near real time), engineering and planning (for both troubleshooting, design work, right-sizing, management reporting), quality assurance and customer care and of course the end subscriber, whether a retail business enterprise or a carrier or wireless operator in the way of SLA reporting and end-to-end performance visibility.

A number of elements support the successful Carrier Ethernet provider in achieving these capabilities:

a. Instrumentation and management systems supporting the MEF standards (e.g. supporting MEF 35.1/36.1 where appropriate)
b. Capacity and scale to address real-time collection, query and display of service performance metrics for large carrier-based networks
c. Role-based interface into the performance data to support the specific needs of each stakeholder both inside the operator as well as external to the operator, e.g. subscriber-oriented SLA reports,
d. Tailored reporting, preferably with APIs to enable on-demand report generation and analysis
e. Metric normalization to support legacy instrumentation that may be migrating to MEF-based standards but not yet available in existing areas of the network
f. Service-based, multi-tenancy to map the collected metrics to the appropriate service (EVC or OVC) and customer while maintaining secure access to data to authorized users
g. Direct, off-line access to performance data to enable subscribers to download their performance metrics
h. Automated provisioning of collection and reporting in-line with Service Operations Functionality within the framework of LSO

3. Summary

Assurance is one of the three tenets of the Third Network vision and truly distinguishes the concept of the Third Network from that of the best-effort Internet. Thus service assurance is a driving force within the work of the MEF. From creating a service OAM layer independent from the underlying network transport technology, to defining the performance attributes necessary to monitor and manage a Carrier Ethernet network, the fifty-plus specifications of the MEF often touch on this complex topic. From service activation testing, performance management to fault management, there is a broad body of work available to help the Carrier Ethernet provider implement Carrier Ethernet service assurance best practices.

It begins with setting objectives and expectations between the provider of the service and the subscriber of the service, covered in the companion paper, "***Part I: An Introduction to Service Assurance and Carrier Ethernet Service Performance Attributes"*** and then performing the necessary monitoring of the network performance in order to determine whether the objectives are met. Furthermore, the successful provider needs to deliver the service performance visibility to enhance the subscriber experience and aid in both subscriber-based and provider-based decision making. From the MEF's definition of performance attributes, the service level specification and the PM solutions to the recently published Carrier Ethernet Performance Reporting Framework specification (MEF 52), Carrier Ethernet providers can best deliver a high-value, differentiated business class CE 2.0 service.

The MEF's best practices leverage considerable work from the industry including the use of ITU-T's Recommendation G.8013/Y.1731 and the TM Forum's SID model to name a few, and have developed a rich set of attributes, guidelines and best practices to create resilient, valuable, business-class connectivity services. The greater the level of adoption by all members of the industry (equipment providers, LSO/OSS solution providers and Carrier Ethernet service providers), the faster the time to revenue the industry can enjoy. Particularly in terms of launching Carrier Ethernet products with the end-to-end SLAs the market is demanding. Furthermore, in the context of multi-operator services, as each operator aligns on the same metrics, and begins to align on the display and sharing of the performance data through standard APIs, the more effective and transparent multi-operator service management becomes by agreeing on what to measure, how to measure and how to share those measurements.

4. Guide to the MEF Technical Specifications

A summary of each specification that defines elements of performance management is provided for the reader that wishes to investigate the technical details of each particular specification.

MEF 6.2 - EVC Ethernet Services Definitions Phase 3

Defines the various CE services and their attributes (E-Line, E-LAN, etc.) including the EVC performance as well as envelopes and token sharing service attributes that relate to bandwidth profiles.

MEF 7.2 - Carrier Ethernet Information Model

This specification defines the EMS-NMS information model such that management system developers can create interoperable CE service management systems.

MEF 10.3 - Ethernet Services Attributes Phase 3

Defines the service attributes for end-to-end Carrier Ethernet services, including the definition of bandwidth profiles and the performance attributes used to define the SLS.

MEF 10.3.1 - Composite Performance Metric (CPM) Amendment to MEF 10.3

Defines an additional performance attribute; the composite performance metric (CPM).

MEF 15 - Requirements for Management of Metro Ethernet Phase 1 Network Elements

Original definition of statistics to be collected by network equipment.

MEF 17 - Service OAM Framework and Requirements
As Carrier Ethernet services are technology agnostic, a common OAM definition is needed, regardless of the transport layer, hence the Service OAM (SOAM) architectures was created as part of this specification. Performance attributes defined within this specification have evolved in as part of more recent specifications.

MEF 22.2 – Mobile Backhaul Phase 3 Implementation Agreement
CoS Mapping and CPOs and performance recommendations for mobile backhaul networks.

MEF 23.2 – Class of Service Phase 3 Implementation Agreement
Defines 3 specific classes of service (H, M and L), and introduces the concept of performance tiers (PT). Specific values for performance objectives related to each PT and CoS are defined, as well as recommended values for specific applications and their CoS. The framework presented can be used to extend the MEF CoS model from the original 3 classes (H, M and L).

MEF 26.2 - External Network to Network Interface (ENNI) and Operator Service Attributes
This specification focuses on ENNI and defines the service attributes for inter-provider Carrier Ethernet services. It includes the performance metrics and service level specification (SLS) for operator virtual connections (OVC), as well as the definition of bandwidth profiles for inter-provider services.

MEF 33 - Ethernet Access Service Definition
This specification defines the Ethernet Access Services, that is, ENNI to UNI Carrier Ethernet services in contrast to the EVC-based (UNI to UNI) services which are defined in MEF 6.2.

MEF 35.1 - Service OAM Performance Monitoring Implementation Agreement
Implementation agreement on performance monitoring; in essence the MEF's best practices for performance monitoring of CE services on which a significant portion of this white paper is based.

MEF 36.1 - Service OAM SNMP MIB for Performance Monitoring and MEF 39 - SOAM Performance Monitoring YANG Module
MEF 36.1 is the instantiation of MEF 35.1 using the Simple Network Management Protocol (SNMP)'s MIB structure. MEF 39 reflects the MEF 35 (predecessor to 35.1) using YANG models.

MEF 41 – Generic Token Bucket Algorithm
A specification of the Generic Token Bucket Algorithm based on the Bandwidth Profile Algorithm which is specified in MEF 10.3.

MEF 51 - OVC Services Definition
This specification focuses on Operator Virtual Connection (OVC) Services based mainly on the service attributes defined in MEF 26.1 (predecessor to 26.2); there are also some service attributes defined in this document that go beyond MEF 26.1.

MEF 52 - Carrier Ethernet Performance Reporting Framework
Definition of the framework to support CE service subscribers with access to KPI performance data and statistics through a standardized interface from their service provider. The specification borrows from the TM Forum frameworks.

MEF 55 - Lifecycle Service Orchestration (LSO): Reference Architecture and Framework
Details the LSO reference architecture and places the concept of service quality management, customer performance reporting, performance events defined as part of service problem management.

5. About the MEF

MEF is the driving force enabling Third Network services for the digital economy and the hyper-connected world. The Third Network concept combines the agility and ubiquity of the Internet with the assurance and security of CE 2.0 (Carrier Ethernet 2.0). Third Network services provide an on-demand, orchestrated, and assured experience with user-directed control over network resources and cloud connectivity. Third Network services are delivered over automated, virtualized, and interconnected networks powered by LSO, SDN, NFV, and CE 2.0.

MEF leverages its global 210+ network operators and technology vendor community, builds upon the robust $80 billion Carrier Ethernet market, and provides a practical evolution to the Third Network with LSO, SDN, and NFV implementations that build upon a CE 2.0 foundation. For more information, see www.MEF.net.

6. Glossary and Terms

A glossary of terms used in this document can be found online at MEF.net.

7. References

Source	Link
MEF	The Third Network, Vision and Strategy (Based on Network as a Service Principles)
MEF	CE 2.0 Service Management Life Cycle White Paper
MEF	Published Technical Specifications
MEF	Understanding Carrier Ethernet Service Assurance - Part I: An Introduction to Service Assurance and Carrier Ethernet Service Performance Attributes
ITU-T	Recommendation G.8013/Y.1731: Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks

8. Acknowledgements

Editor: Christopher Cullan, InfoVista

Principal Authors:
- Christopher Cullan, InfoVista

Contributors:
- David Ball, Cisco
- Yoav Cohen, RAD
- Eric Leser, Charter Communications
- John Siwko, AT&T
- Marjory Sy, PLDT