



Understanding Carrier Ethernet Service Assurance

Taking advantage of MEF's best practices for performance monitoring

— Part I —

An Introduction to Service Assurance and
Carrier Ethernet Service Performance Attributes

September 2016

Table of Contents

1. Introduction and Overview	4
1.1 Abstract	4
1.2 Target audience	4
1.3 Document Purpose and Scope	4
1.4 Executive Summary	4
2. Introduction to Service Assurance	5
3. Key Business Drivers	6
3.1 Service Level Agreements	6
3.1.1 Defining the service level quality	7
3.1.2 Measuring and maintaining the service level quality	7
3.2 The Application Experience	8
3.3 SLA Reporting and Portals	8
3.3.1 Service Performance Transparency	8
3.3.2 Performance Reporting Framework	9
3.4 The Third Network Vision	9
3.4.1 LSO – Synchronization needs for PM	11
4. Carrier Ethernet Service Performance	12
4.1 Why Performance Matters	12
4.1.1 Policing	12
4.1.2 Shaping	12
4.1.3 Classification	12
4.2 Performance Attributes – What to Measure	13
4.2.1 Defining Service Performance	13
5. Summary	21
6. Guide to the MEF Technical Specifications	22
7. About the MEF	24
8. Glossary and Terms	24
9. References	24
10. Acknowledgements	24

List of Figures

Figure 1 – Features Influencing Ethernet Services Customers’ Purchasing Decisions7

Figure 2 – Criticality of Customer Reporting9

Figure 3 – LSO Reference Architecture11

Figure 4 – Service Operations Functionality Components.....11

Figure 5 – Illustrative Use Case.....14

List of Tables

Table 1 - BizCo’s SLS Summary20

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 II: How to Measure Carrier Ethernet 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 the why and what of performance monitoring and assurance in the context of CE 2.0 services and touches on the evolving role of performance management with respect to 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 both monitor a network service and to maintain a particular service quality. This is a simple description to a very critical and complex topic. The MEF has tackled this topic to enable Carrier Ethernet providers to deliver high-value Carrier Ethernet services that are fully assured and can deliver the service quality required in today's competitive business service and carrier-to-carrier markets. In particular, the MEF has detailed service activation testing, performance management and fault management within its technical specifications.

This paper focuses on performance management beginning with the key business drivers. These drivers include service level agreements, the application experience and customer reporting with end-to-end performance visibility and their significance in driving both customer acquisition and customer retention.

Performance management became more complex and necessary as the industry began its transition from predictable, TDM-based networks to shared, packet-based networks. Once again its significance is well noted as the industry adopts SDN- and NFV- based principles along with cloud computing and the Third Network vision leveraging Lifecycle Services Orchestration (LSO).

Based on the relevant business drivers, the reader will benefit from understanding the key attributes of Carrier Ethernet performance and what to measure in terms of assuring the service quality of a Carrier Ethernet

service. By clarifying the standard performance metrics¹, 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.

Together with its companion paper “**Part II: How to Measure Carrier Ethernet Performance Attributes**”, 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. Introduction to Service Assurance

Service assurance has never been more critical for today’s evolving network service technologies. With the combination of automation, virtualization, growing global inter-connectedness and the transition to SDN, NFV and LSO, the services evolve to become far simpler for the end subscriber, but conversely far more varied and complex for the Communications Service Provider (CSP). This puts a great deal of importance on service assurance – ensuring offered services meet the pre-defined level of service quality – which is why assurance is one of the three tenets of the MEF’s Third Network Vision, alongside agility and orchestration.

Service assurance is primarily focused on processes and procedures involving the monitoring of a network service and its underlying infrastructure in an effort to maintain a particular service quality. The service quality may be very strictly defined based on service level agreements between the service provider and the subscriber or more loosely defined in the case of best-effort services, however there is an implicit, if not explicit, service quality objective. For example, even a best effort service is expected to be available and perform consistently most of the time. Although there may not be a defined agreement in place between provider and subscriber, there is typically an internal, provider-defined objective and set of processes for maintaining a pre-defined level of service quality. The service assurance domain covers a wide variety of disciplines including trouble-ticketing, testing, customer experience management, SLA monitoring, QoS management, fault and performance management.

The MEF has developed the concept of bandwidth profiles, establishing the policing of traffic through the Carrier Ethernet Network (CEN); introduced service level specifications (SLS), which define the committed level of service performance for service traffic; and incorporated the use of multiple classes of service (CoS) to prioritize service traffic. These elements of service assurance are combined with the following Carrier Ethernet service assurance functions described throughout the body of the MEF specifications:

- Service Activation Testing
- Performance Management
- Fault Management

The basis for much of the network service’s evolution in the global market is the flexibility and power of Carrier Ethernet 2.0. As of 2016, there is a global market of approximately 80 billion USD for Carrier Ethernet services and technologies. With over fifty specifications within the MEF canon covering service definitions, information models, implementation agreements and service catalogs, service assurance is woven in through much of the body of work.

¹ 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.

There are some specific specifications, which focus on service operations, administration and maintenance (OAM), and in particular fault and performance management. However critical elements of service assurance can be found throughout the MEF canon, which makes it more challenging to easily understand the service assurance constructs and best practices when approached from the topic itself.

The objective of this paper is to consolidate and abridge the pertinent details of service assurance, as it relates to Carrier Ethernet, specifically the key drivers and the relevant service performance attributes. Fault management will be addressed in later works.

This paper focuses on the performance management function, a proactive function which gathers and analyzes multiple data points across the CE service to understand the performance of the service and its underlying components and thereby detect any degradation or potential degradation in service quality so that steps can be taken, either manually or with automated policy control, to correct or amend the service to avoid such degradation and ensure the defined service quality is met. Performance management is primarily described as part of the MEF's Service OAM requirements which focus on the tools and processes for performance monitoring in addition to fault management.

Finally, as the global communications industry evolves, flexible and dynamic interconnectedness becomes of increasing importance. To enable the vision of ubiquitous, agile and dynamic multi-operator services, it is critical that the industry understands and adopts the best practices, from definitions of KPIs, instrumentation methods, through to the presentation and sharing of performance management to achieve the necessary service transparency on which this evolution will be based.

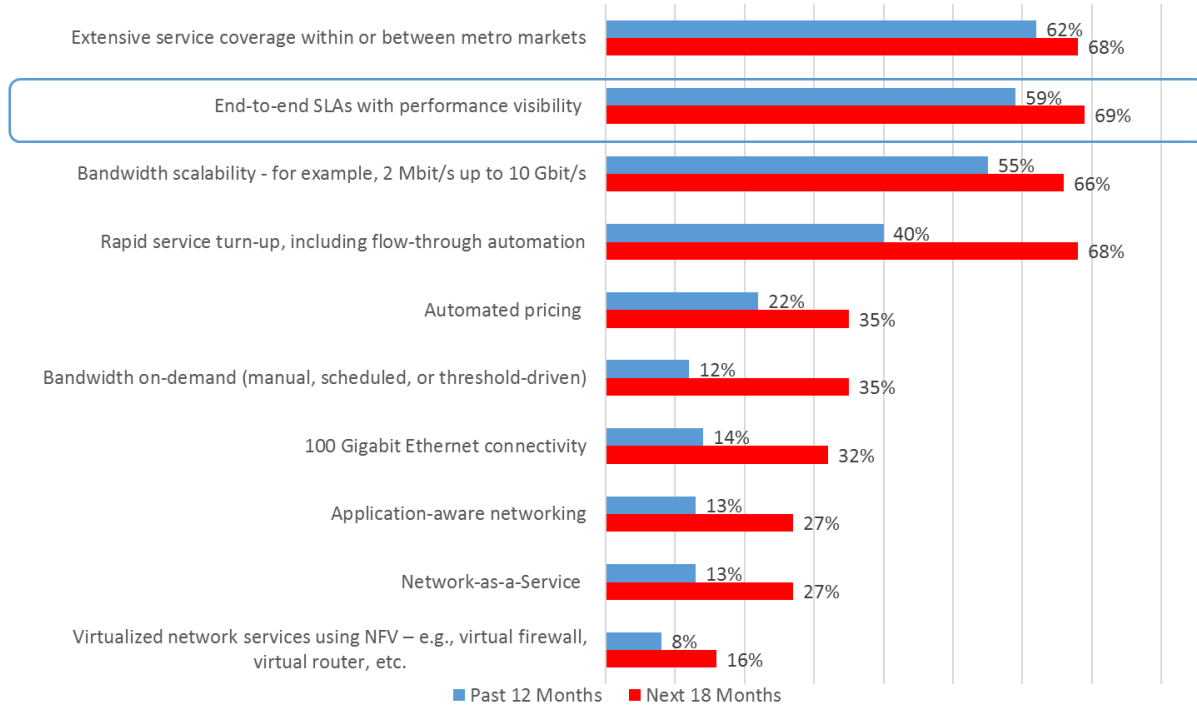
3. Key Business Drivers

3.1 Service Level Agreements

A Service Level Agreement or SLA is essentially a contract, implicit or otherwise, between provider and subscriber covering a potentially wide range of elements with respect to the service provided. In the context of Carrier Ethernet services, specifically the technical operations of the service, the key elements of the SLA comprise the service level specification or SLS. Critical to both retail business services and carrier to carrier or wholesale services is the SLS. In fact, recent surveys of the industry have demonstrated that end-to-end SLAs with performance objectives have been the most important feature affecting purchasing decisions.

Most Important Service Features Impacting Ethernet Services Customers’ Buying Decisions

What are the key / very important service features impacting your Ethernet services customers’ purchasing decisions in the past 12 months...and in the next 18 months?



MEF Survey – 2015 Carrier Ethernet Services Market Trends Report (Sep 2015)

Figure 1 – Features Influencing Ethernet Services Customers’ Purchasing Decisions

3.1.1 Defining the service level quality

A key concept of service assurance, which stems from the SLA, is that the service quality objectives must be defined. Some or all of these objectives may be documented in the SLA as part of the contract with the subscriber. Some objectives may be defined internally only, that is they are not guaranteed to the subscriber but they are used to manage the service quality to a particular level as part of the provider’s internal operations. The Service Level Specification (SLS) defines the end-to-end performance characteristics for the CEN, and is typically included as part of the technical section of a Carrier Ethernet SLA making the SLS a critical aspect of SLA contracts. In fact, as operators introduce technologies based on NFV and SDN, the foundational aspects of the service change significantly (virtual versus physical functionality, location, SDN controllers, etc.) in an effort to increase agility, yet the SLA and customer expectation of service quality remain unchanged. This customer expectation makes it critical to understand the performance monitoring best practices established by the MEF community.

3.1.2 Measuring and maintaining the service level quality

The second key concept of service assurance is the operational task of monitoring and measuring the on-going performance of the CEN and validating that in fact the objectives previously defined — whether part of the contract with the subscriber, or merely internal to the provider — are met. Part of this concept requires recognition that a comparison to an objective is a pass or fail test, that is either the objective (such as an average one-way frame delay objective) is either met or is not met.

Although an objective may not be defined as part of the SLS and hence the SLA, the actual performance metric may still be of interest to either the provider, subscriber or both. Calculating the specific performance metric enables the best understanding of the service's performance to support optimal decision-making. There is value in knowing both the pass/fail status of performance objectives; the granular visibility of a service's performance at a point in time; and a historical view of the service's performance. This historical view helps identify and forecast trends, detect and resolve technology or process issues and helps to spot changes or degradation that empower the Carrier Ethernet service's stakeholders.

3.2 The Application Experience

The end user experience for any network-based application is impacted significantly by the network's performance. Some elements of the network performance are limited by physical characteristics such as distance between end points (e.g. UNIs) which will impact latency, available transmission media (e.g. one pair of twisted copper versus fiber, versus microwave) which limits bandwidth impacting throughput, security (e.g. data scanning, encryption) which adds additional processing time at network end points and finally, the application's behavior itself (e.g. how it adapts to network conditions, chattiness, size of data and compute time at client and server locations) which typically becomes more prominent over constrained network resources. The result of these performance-impacting factors is subscriber frustration coupled with a potential loss of productivity.

The first step to ensuring the best application experience is assuring the performance of the underlying network. In the context of Carrier Ethernet, it is necessary to architect a Carrier Ethernet network service that provides the necessary bandwidth, QoS, prioritization (CoS) and guarantees (SLS based performance objectives) to support the application.

3.3 SLA Reporting and Portals

3.3.1 Service Performance Transparency

Consider the experience of waiting for a hotel elevator, which can often be frustrating, especially when in a hurry to check-out of the hotel and head to the airport. Now imagine you have no visibility on where the elevator is and no concept of when it may arrive to pick you up. Furthermore, there is no defined metric in terms of availability of the elevator or how quickly it will arrive and deliver you to the ground floor. This experience is similar to that of the CE services subscriber without a defined SLS as part of their SLA and no visibility of the service and its performance.

The result of this situation causes the subscriber frustration and anxiety that can lead to a poor impression of the service, and its provider, sowing the seeds for churn. Furthermore, the service desk will experience increased call volume based on subscriber perception as to whether there is a fault or performance degradation behind their perception. Evidence of the importance of this service performance transparency was demonstrated in results of a survey of Carrier Ethernet providers regarding the criticality of customer reporting in terms of both customer acquisition as well as customer retention (see figure). As a result, many CSPs have focused on delivering visibility of both their service level specifications and the performance of the CE service to empower the subscriber to understand real versus perceived performance issues and to act on that information (pause a backup, open a ticket, change routing, contact internal application owner, etc.).

How important is customer reporting to winning new business and retaining existing business?

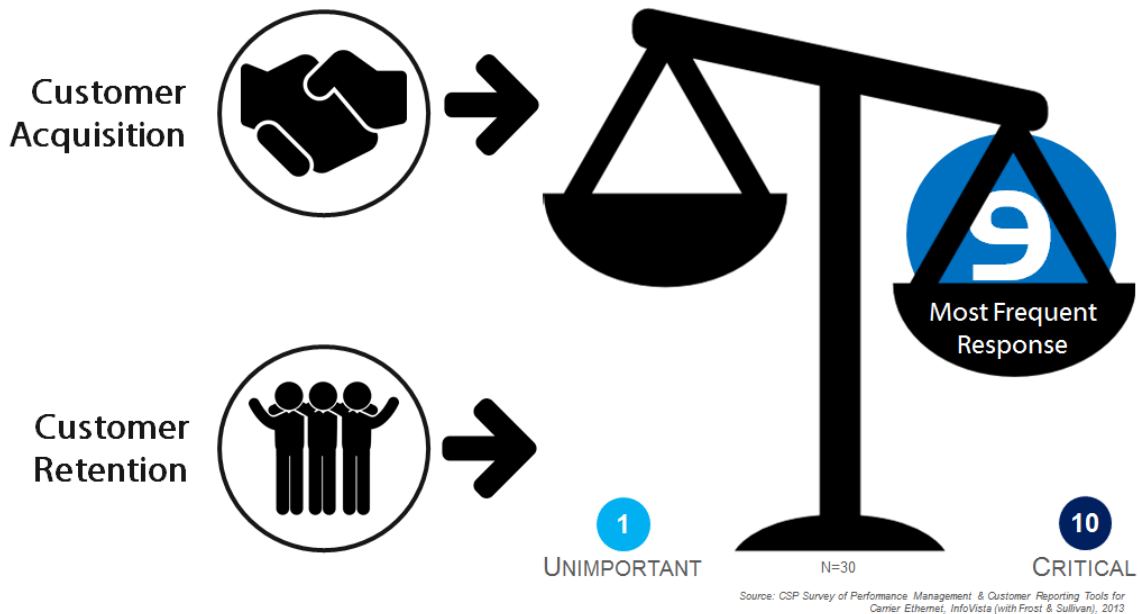


Figure 2 – Criticality of Customer Reporting

3.3.2 Performance Reporting Framework

To support the increasing need for service performance transparency for the CE subscriber, in particular in carrier-to-carrier services such as mobile backhaul, the MEF published specification 52 to enable a standardized access to key performance indicator data generated by the CE service provider’s network instrumentation and made available via an application programming interface to the subscriber. This API provides near-real time, read-only access to the service’s performance metrics (i.e. EVC or OVC based services) as specified by MEF.

The associated architecture includes a web-based portal available to the subscriber for retail services (e.g. E-Line, E-LAN and E-Tree services) and an interface specification for carrier-to-carrier (i.e. business-to-business) implementations enabling service providers to review performance of access or transit services (e.g. E-Access, E-transit) in near real time.

In both cases the subscriber is able to query for performance data based on the Carrier Ethernet service ID (i.e. EVC or OVC IDs). Performance data can be viewed in the portal or downloaded for offline analysis. Based on TM Forum’s SID, the MEF has provided an information model to support the description of the various performance metrics available so the subscriber (or service provider reviewing their access services) can query for the specific metrics based on what is available from the CE provider.

3.4 The Third Network Vision

The Third Network Vision is a critical concept in leading the way forward towards a ubiquitous, global, automated and dynamic set of network resources that work together to provide connectivity, not unlike the Internet, yet provide the service level guarantees that businesses rely on. Recent advances in SDN (software defined networking) and NFV (network functions virtualization) have provided many elements to support more dynamic and more automated network services.

The vision is based on three key tenets, Agility, Assurance and Orchestration. The ability to rapidly change to the demands of the subscriber, as well as be able to rapidly provision valuable network-based services is critical to the advancement of the digital economy and telecommunications, hence the significance of being agile. Furthermore, to provide this agility requires a new level of collaboration within a network and its constructs but also as importantly between networks, in particular networks from different operators so that services can be dynamically and automatically provisioned to address subscriber needs even when crossing administrative boundaries between operators. Finally, without assurance, the vision would simply be another “Internet”. The assurance element provides the key aspect that differentiates this vision from Internet-based services by delivering the necessary security and performance characteristics required. This assurance is required by the business services subscriber due to the criticality of their needs and is also a significant element of the value proposition for a service provider directly impacting the monetization of those services.

When examining this vision, it is easy to understand the criticality of standardized methods for performance management, first and foremost to reduce internal operational costs incurred to address disparate and proprietary metrics, technology domains and the personnel to support non-standard operations. It becomes even more critical when considering the multi-operator scenario. To achieve the vision of the Third Network, operators must seamlessly interact to not only deliver but operate a network service for the life of the subscription to meet a consistent, well-defined service level. Practically speaking the only way to succeed with such an initiative, in particular in an automated fashion, is to agree on the metrics and KPIs used, how to collect that information and finally how to best communicate performance data across networks and across operators so that the service provider (that is, the entity that administers the subscriber-to-provider service level) can manage a multi-operator network as one.

Without such standards and the related adoption by the industry, the costs and inter-operator difficulties associated with proprietary or non-uniform methods of performance management will continue to persist. The Third Network’s vision to provide rapid, on-demand, dynamically provisioned multi-operator services will be inhibited. The result being that over-the-top players will continue to exploit the best-effort Internet for profitable service delivery without an assured monetizable alternative from CSPs.

3.4.1 LSO – Synchronization needs for PM

A critical enabler of the Third Network vision is Lifecycle Service Orchestration or LSO.

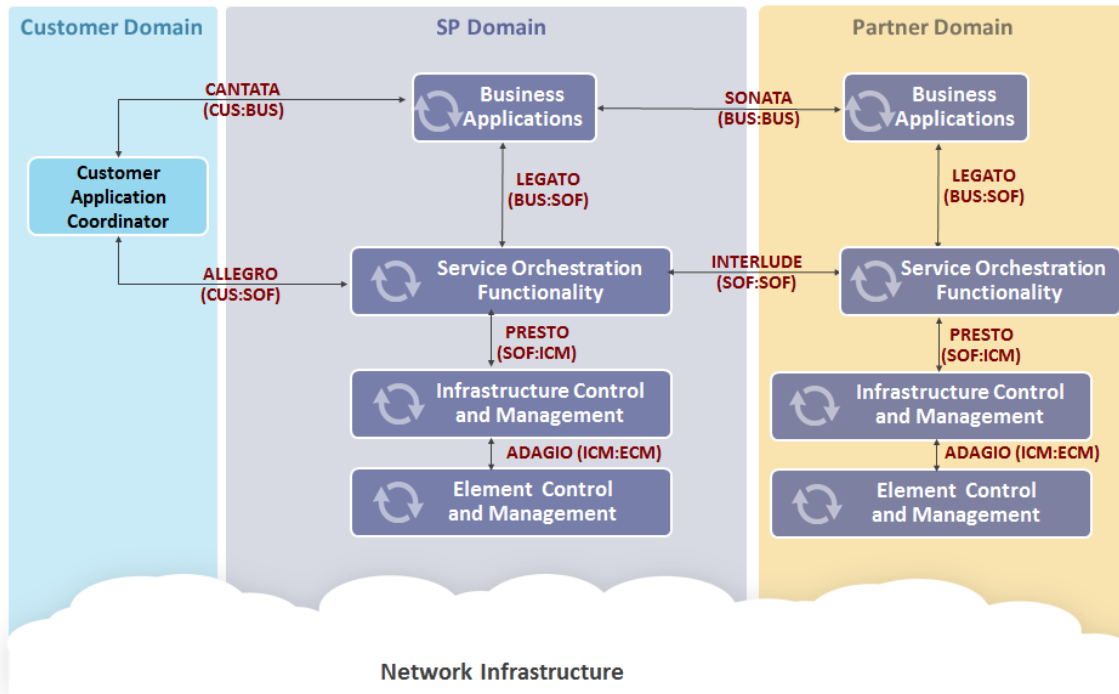


Figure 3 – LSO Reference Architecture

The LSO architecture is outside the scope of this paper, however it’s important to recognize the inter-dependencies of LSO functions, service assurance and performance monitoring.

The MEF has defined the Service Operations Functionality (SOF) within the LSO reference architecture as the set of service management layer functionality supporting an agile framework for streamlining and automating the service lifecycle in a sustainable fashion for coordinated management.

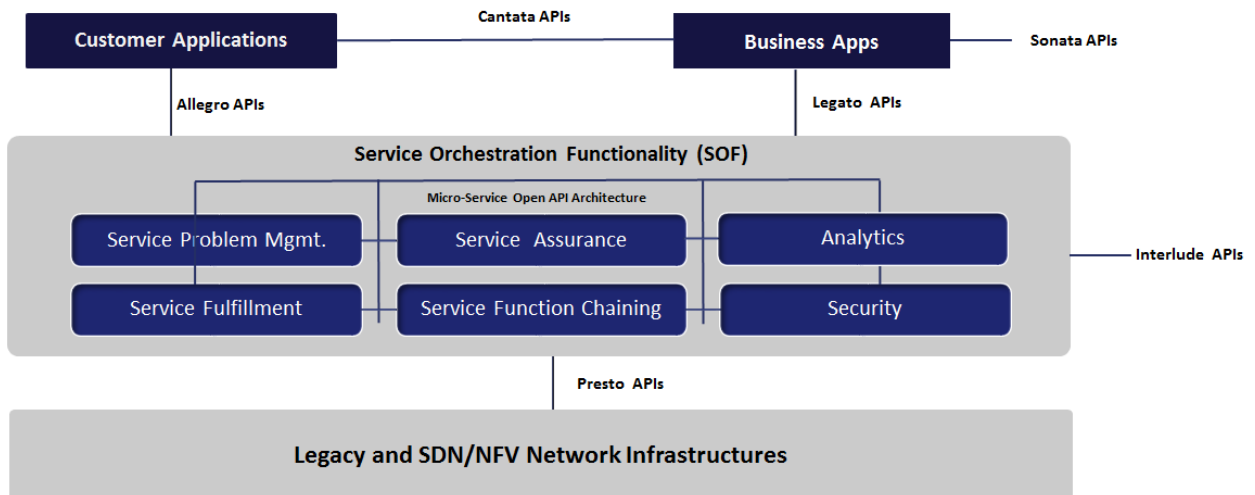


Figure 4 – Service Operations Functionality Components

For example, the CE service's defined service level specifications must be configured, ultimately by either direct definition of the subscriber or by a mapping between product needs and SLS parameters, handled by the LSO. Secondly, the network elements (both physical and virtual) must be provided the relevant configuration information and then begin collecting the appropriate performance metrics to support LSO's service quality management and to some extent, service problem management. As the configuration, and ultimate operation of performance monitoring requires understanding of the maintenance entity group levels, as well as the defined end points (MEPs), the performance management function must have an understanding and mapping of the collected metrics, the respective service, subscriber (or carrier in an OVC context), SLS, performance reporting parameters and the relevant network elements.

This synchronization of service and subscriber context with the relevant performance data is not directly addressed as part of the LSO reference architecture at this time but as operators incorporate LSO along with both SDN and NFV concepts into their networks and services, the ability to synchronize PM data within the LSO context must be easily performed, in a standardized manner, to enable the rapid and practical transition of current and legacy OSS/BSS systems to an LSO-based environment.

4. Carrier Ethernet Service Performance

4.1 Why Performance Matters

Shared resources, such as packet-based (versus circuit-based) networks, are constrained, thus affecting how well traffic is able to flow through the network. Various controls must therefore be put in place to ensure the integrity of the network.

With the introduction of NFV and SDN, network resources are further virtualized and decomposed. To properly and dynamically automate the allocation of resources, a performance feedback loop is necessary so as to advise the controllers, orchestrators and similar functions of the present state and anticipated future state of performance enabling those functions to best adapt resource allocation to meet the defined policies and subscriber requirements of the service.

4.1.1 Policing

When traffic is presented for delivery on the network, the network may not have sufficient resources to deliver that traffic while meeting the expected level of service quality (e.g. the SLS). In this case, the excess traffic may be delivered in a best-effort fashion. Ultimately, if more traffic is presented than the network can deliver, some traffic must be dropped. The determination of which traffic has guaranteed performance, which is delivered best-effort, and which is dropped, is a key QoS function and is performed by the policer.

4.1.2 Shaping

As network traffic is often variable over time, rather than dropping traffic because the network is not able to handle delivery at that time, the network can buffer, or queue the traffic for a limited time and then send it when resources are available.

4.1.3 Classification

Some network traffic is more important than other traffic either due to criticality to the subscriber or the nature of the application itself. Therefore, classifying network traffic into particular classes of service enables

the network to deliver each class based on the needs of that network traffic. Furthermore, policing and shaping of the traffic for each class of service can be performed if desired.

In legacy TDM-based networking environments, performance was very predictable, in part because the resources were dedicated to specific applications (like voice) and the transport was deterministic. As a result, performance management was mostly concerned with the service availability. Ultimately due to the fact that packet-based networks are constrained, and multiple applications share the same constrained resource, some traffic must be buffered, some must be dropped, some must be sent with no guarantees and some must be prioritized over others, possibly with performance guarantees.

With the growth of cloud computing (Gartner says worldwide public cloud service market will reach \$204 Billion in 2016)² far more application traffic is traversing the WAN. As a result, WAN performance and in particular CE service performance is more critical than ever before. Coupled with the dynamic and varied application behavior, it is critical for service providers to manage their CE services' performance to monitor for performance degradation (or potential performance degradation) to assure they are respecting any agreement with the subscriber, and to minimize service disruption. In addition, the subscriber has a vested interest in understanding the CE service's performance so as to isolate performance degradation to either the WAN itself or enterprise-internal areas such as the LAN, IT systems and the applications themselves.

4.2 Performance Attributes – What to Measure

4.2.1 Defining Service Performance

In today's packet-oriented, shared data networking model, multiple types of applications (real-time, bursty, asymmetrical, interactive and high bandwidth) compete for limited resources and as a result the end user experience may vary greatly. The primary elements of a CE network service that impact the user experience and hence the perceived value of the network service include the size of the network in terms of bandwidth, service availability, response time of the network, consistency (including variation in response times), and loss which may lead to slower response times due to re-transmissions or even transaction failure.

4.2.1.1 Bandwidth Profiles

Bandwidth profiles can police traffic for an entire UNI (least granular), for each EVC or for each EVC's Class of Service (CoS) at the UNI (most granular)³. Bandwidth is a prime user-experience-impacting element. The bandwidth profile defines the committed information rate (CIR) in bits per second which is the average rate of service frame traffic the CEN will accept and mark as 'green', meaning the SLS will be applied to those frames. As packet-based traffic is rarely consistent, QoS mechanisms are required to police the incoming traffic.

***Note:** The MEF does not define any specification related to the shaping of traffic. Bandwidth profiles serve only to police the traffic. Traffic shaping may be performed by the subscriber, often at the port level but this is outside the scope of the MEF's work.*

In essence, the green frames are service assured. The bandwidth profile also defines the excess information rate (EIR), which is the average rate of extra service frame traffic in bits per second that the CEN will accept (over and above the committed rate); however, the frames are marked 'yellow' meaning that the SLS will not

² <http://www.gartner.com/newsroom/id/3188817>

³ MEF 6.2 eliminated the per UNI and per EVC bandwidth profile service attributes, keeping only the bandwidth profile per CoS service attribute – the latter being the most granular.

apply to those frames and they may be dropped at a later point. These service frames, although accepted for transport through the network, are not assured.

Bandwidth profiles also specify the maximum burst size for committed and excess traffic known as the Committed Burst Size (CBS) and Excess Burst Size (EBS) respectively. These represent the number of bytes of service frames that can burst at the line rate of the UNI and still be marked as ‘green’ (based on CBS) or ‘yellow’ (based on EBS). Lastly the bandwidth profile defines the presence of a coupling flag and the color mode, two additional parameters which are out of scope for this paper [see MEF 10.3].

4.2.1.2 Service Level Specifications

The service level specification (SLS) is a key component of the SLA and deals with technical (that is, engineering-based) specifications versus broader business-level specifications such as mean time to respond to a customer reported incident. Specifically, it is the technical specification of the service level being offered by the service provider to the subscriber.

It is not mandatory to provide any performance objectives for an SLS, the result of which is a pure, best-effort service, however at least one performance objective should be provided using the MEF-defined performance attributes⁴.

In an effort to discriminate between different types of applications as well as their relative priority to the business, multiple classes of service can be defined within an EVC in which case the SLS defines performance objectives for each class of service.

To best understand the MEF’s performance management concepts, review the simplified use case below which will be used to help illustrate the different performance attributes which follow.




Figure 5 – Illustrative Use Case

In this use case, a subscriber, BizCo has one point-to-point EVC, an E-Line. The E-Line is provided by EZProvider, BizCo’s selected CSP. BizCo uses this E-Line to connect their headquarters to their remote data center. The UNI at the headquarters is UNI_A and the one at the data center is UNI_Z.

BizCo negotiates an SLA with EZProvider covering a number of aspects including the minimum network service quality. The specific requirements of this service quality are defined within the service level

⁴ This paper uses the term ‘performance attribute’ when referring to the definition of the attribute as opposed to a specific instance of the attribute.

specification of the Carrier Ethernet E-Line along with the bandwidth profile. We will assume a single CoS for simplicity.

4.2.1.3 Performance Attributes

Frames that are within the CIR/CBS defined by the bandwidth profile are marked green and the SLS applies to those frames. That is, these frames are service assured. Frames that are within the EIR/EBS are marked yellow and are delivered into the network but are exempt from the SLS. Frames outside of CIR+EIR are dropped.

Recall that the concept of service assurance includes both the definition of the service quality and the measuring and monitoring of the service to validate if it is meeting (or exceeding) that service quality objective. As such the SLS defines an objective for one or more of the performance attributes that follow. Per the illustrative use case, where EZProvider is delivering one point-to-point service to BizCo, the validation is straightforward; each attribute is measured between the end points of the EVC, and if the measurement meets or exceeds the objective defined in the SLS, then the SLA is met. If not, the SLA would fail.

Each of the attributes is defined in one direction, i.e., from UNI_A to UNI_Z and vice versa. Where it appears to get more complicated is with multipoint services, however these are easily handled by defining the relevant set of ordered UNI pairs (that is UNI_A to UNI_Z , UNI_Z to UNI_A , etc.) for which the SLS applies. This definition may include all ordered UNI pairs or a subset of them. The performance metric(s) are calculated for each of the ordered UNI pairs defined in the set. If any of those pairs fails to meet the objectives for one or more of the performance attributes, then the SLA for the entire multipoint service would fail. For example, in a three site E-LAN, there is a total of six ordered UNI pairs. If a performance objective defined within the SLS is specified (e.g. one-way frame delay performance) for the entire set of all six ordered pairs and five of them meet the objective but one of them fails, then the entire E-LAN service fails the SLA. The concept of a chain being as strong as its weakest link is analogous to evaluating whether a multipoint service passes a performance objective; that is, a multipoint service is as strong as its worst performing ordered UNI pair.

BizCo's E-Line service from EZProvider is defined with a CIR of 20Mbps and an EIR of 10Mbps. For this use case we won't detail the specific committed burst size (CBS) and excess burst size (EBS). The service level specification will apply to all of the frames within the CIR (and CBS). Any excess traffic, that is the additional 10Mbps, will be delivered in a best effort manner. Should BizCo submit more than the combined 30Mbps of traffic, the additional traffic will be dropped. Note that for this use case there are exactly 2 ordered UNI pairs, namely $\{UNI_A, UNI_Z\}$ and $\{UNI_Z, UNI_A\}$

There are nine⁵ performance attributes related to an EVC that are defined within the SLS and all are specified by CoS with different performance objectives defined for each CoS in the case of multiple CoS. The nine attributes are listed below grouped using the high level concepts of latency, variation and loss.

4.2.1.3.1 Delay (latency)

1. One-way Frame Delay Performance

One-way Frame Delay is the one-way delay from the 1st bit of the ingress service frame to the last bit of the service frame at egress. The attribute is expressed with an objective that a certain percentage of service frames must not exceed the defined one-way delay.

⁵ There is a 10th attribute known as the Composite Performance Metric (CPM) which will be covered separately in this paper.

For example, a performance objective could be defined as 100ms for the 95th percentile of all service frames for a month between UNI_A and UNI_Z of an EVPL service, i.e. 95% of all service frames during the month must have frame delay less than or equal to 100ms.

Frame delay measures both frame processing time as well as latency across the network and is very relevant to the behavior and resulting responsiveness of many applications.

BizCo has defined a one-way frame delay performance objective of 150ms for the 95th percentile of all service frames in a 24-hour period, midnight-to-midnight. This objective is applied to both ordered UNI pairs, in other words, the 150ms objective must be met in both directions, from HQ to the data center and from the data center to HQ.

Note: One-way delay may be approximated by a two-way or round-trip measurement and similarly for the other one-way metrics (2) and (3).

Note: The Carrier Ethernet Class of Service (CoS) – Phase 2 Implementation Agreement requires that for any SLS that is based on a particular CoS (which would be defined by the MEF CoS label), that SLA must include either the FD or MFD performance.

2. One-way Mean Frame Delay Performance

Similar to one-way frame delay performance, this metric calculates the mean average of the frame delay for all service frames defined across a set of UNI pairs and defines the objective for this mean value based on a defined time interval.

For example, a performance objective could be defined as 80ms mean for all service frames between UNI_A and UNI_Z of an EVPL service.

This may be a more appropriate metric to allow some variability in latency but for non-real-time applications will define a reasonable expectation of application responsiveness based on the CE service.

BizCo is happy with the FD performance objective and does not require a MFD objective as part of their SLS.

4.2.1.3.2 Variation

1. One-way Frame Delay Range Performance

This attribute defines an objective for the variation in frame delay over a period of time. The objective defines the upper limit of variation between the minimum frame delay measured over a period of time and a specified percentage of all frame delay measurements over that same period. This objective is the one-way frame delay range (FDR). It can be calculated by taking one-way frame delay measurements between two UNIs over the specified time period and for each of those measurements, subtracting the minimum delay over that same period. The one-way frame delay range performance is the specified percentile of all of the calculated differences.

For example, a CSP may define one-way frame delay range performance objective as 100 ms over 24 hours using the 95th percentile of the differences between each frame delay measurement and the minimum. Therefore, the 95th percentile (i.e. 95%) of all FD measurements need to be no more than 100ms greater than the minimum of all FD measurements made over the 24 hours.

This objective helps define a particular consistency of the service's latency.

BizCo is not too concerned with consistency because the bulk of the traffic between these two sites is data oriented, not interactive voice or video. As a result, BizCo has negotiated a FDR of 30 ms over the same 24-hour period as defined for their FD objective. The FDR will be based on difference between the 95th percentile of FD measurements and the minimum FD. The objective is defined in both directions.

2. One-way Inter-Frame Delay Variation Performance

Often considered synonymous with jitter, One-way Inter-Frame Delay variation measures the absolute difference in frame delay between pairs of service frames. A large difference may be difficult for real-time video interactive applications (e.g. telepresence). The objective is the maximum inter-frame delay variation for a given percentage of pairs of service frames over a particular time interval.

For example, a CSP may define the IFDV performance objective as 10 ms over 1 hour using the 95th percentile inter-frame delay variation measurements. Therefore, the 95th percentile of all IFDV measurements made over the hour needs to be no more than 10ms.

Applications that are sensitive to buffering, for example voice traffic require a very small inter frame delay variation.

***Note:** that one-way frame delay range can be used as an alternative to IFDV based on their similarity; one measuring the maximum deviation from a percentile of FD calculations compared with its minimum versus the difference between the one-way frame delay of a single pair of service frames.*

IFDV is similar to FDR in the sense that it is a measure of variability in delay. As a result, it could be considered somewhat redundant depending on the specific use case to define FDR and IFDV but they are unique calculations and could be required in some situations.

Let's assume that BizCo later realizes that voice traffic is growing between the two sites and as a result they update their SLS with EZProvider for the E-Line to include an IFDV objective of 5 ms for the 80th percentile over a 24-hour period. Once again, this objective is set for both ordered UNI pairs.

***Note:** The Carrier Ethernet Class of Service (CoS) – Phase 2 Implementation Agreement requires that for any SLS that is based on a particular CoS (which would be defined by the MEF CoS label), that SLA must include either the FDR or IFDV performance.*

4.2.1.3.3 Loss and Availability

1. One-way Frame Loss Ratio Performance

This metric essentially measures the frames going in at one UNI and how many of them as a percentage of those sent, are not successfully delivered through the other UNI in a UNI pair. The performance objective is defined as a maximum value of this metric for a set of UNI pairs over a particular time interval.

Loss of service frames results in higher-layer protocols or possibly the application itself to have to retransmit data across the CE service resulting in application delays and performance degradation.

To cover archiving and critical DB replication services, FLR is crucial to BizCo, however there are no financial transactions or similar applications that could result in massive business impact caused by loss of service frames resulting in higher application layer retransmits, therefore they have negotiated the FLR with EZProvider of 0.1% over each 24-hour period.

2. One-way Availability Performance

One-way Availability Performance is based on Service Frame loss during a sequence of consecutive small time intervals. This is a unique attribute in that it is not a measurement of pure up-time (i.e. the network is up and delivering 'some' traffic) but rather a measurement of how much traffic is regularly making its way across the network. As such it is a very end-user relevant measure of availability because ultimately even if the network is up, if an application is suffering too much frame loss it becomes unusable which, for practical purposes, is similar to the network being completely 'down'.

In particular, this metric uses the concept of high-loss intervals (HLI) and consecutive high loss intervals (CHLI). The reader may refer to MEF specifications 10.3 and 26.2 for more information about HLI.

For example, a CSP may define 99.99% availability for a month using a threshold for frame loss of 20% meaning that the service is considered unavailable when frame loss is higher than 20% for a number of consecutive intervals. Therefore, after determination of unavailable time based on the consecutive intervals of frame loss, the total unavailable time for the service must be $\leq 0.01\%$ of the month or ≤ 4.38 minutes.

With a solid FLR in place, BizCo wants to ensure the service is highly available over a longer period of time, negotiating a 99.99% availability objective over a 24-hour period, with the service considered unavailable should frame loss exceed 40% for 10 or more consecutive intervals of 1 minute each.

3. One-way Resiliency Performance expressed as High Loss Intervals.

A high loss interval is a defined, small period of time over which there is frame loss above a defined threshold. A performance objective is defined to be the maximum number of such intervals over a particular period of time (e.g. an hour).

Similar to *errored seconds*, this metric indicates an issue in service delivery resulting in re-transmits and ultimately application responsiveness issues.

BizCo negotiates a maximum count of 40 high loss intervals within a 24-hour period. A high loss interval will be based on assessing small time intervals of 1 minute where the FLR is greater than 40%. By definition, high loss intervals have to be counted only during 'Available Time'. In other words, if a period of time is considered unavailable due to the one-way availability performance which is based on Consecutive High Loss Intervals (see the following section), then any high loss intervals within that unavailable time will not impact this performance objective.

4. One-way Resiliency Performance expressed as Consecutive High Loss Intervals

Using the same calculation as (3), this metric defines an objective for the number of consecutive high loss intervals, or consecutive small periods of time that have a high loss interval. Due to the consecutive nature of loss, this metric defines more of a lengthy period of loss and is likely to have a greater impact on application

experience as multiple re-transmission or complete application transactions may have to be made to achieve a response.

In effect, 10 small chunks of time that have a high loss interval in an hour is different from 10 consecutive small chunks of time.

Using the same one-way frame loss ratio threshold used for the one-way availability performance objective of 40% and the small time interval of 1 minute, BizCo will negotiate the number of consecutive time intervals as 5, and the count of such consecutive intervals must be no more than 6 in a 24-hour period. In other words, EZProvider must ensure that over a 24-hour period, there are no more than 6 periods of consecutive high loss, meaning loss of more than 40% in a minute for 5 consecutive minutes.

5. One-way Group Availability Performance

For multipoint-to-multipoint services a given set of UNI pairs is defined where at least a minimum number of such sets is available. The service provider can define the minimum number as well as the associated sets of UNI pairs.

As this use case is a point-to-point service, one-way group availability is not applicable.

Below is a table summarizing the example service level specification including the specific performance objectives defined for BizCo's E-Line service provided by EZProvider.

Table 1 - BizCo's SLS Summary

Performance Metric ⁶	UNI _A to UNI _Z ⁷	UNI _Z to UNI _A ⁸	Time Interval (T) ⁹	Small Time Interval (Δt) ¹⁰	Percentile
FD (ms)	≤ 150	≤ 150	24h	1m	95 th
FDR (ms)	≤ 30	≤ 30	24h	1m	95 th
MFD (ms)	—	—	—	—	NA
IFDV (ms)	≤ 5	≤ 5	24h	1m	80 th
FLR (% of frames)	≤ 0.1	≤ 0.1	24h	1m	NA
Availability (% of time)	≥ 99.99	≥ 99.99	24h	1m	NA
Resiliency (HLI) (No. intervals in time period)	≤ 40	≤ 40	24h	1m	NA
Resiliency (CHLI) (No. intervals in time period)	≤ 6	≤ 6	24h	1m	NA
Group Availability (% of time)	NA	NA	NA	NA	NA

4.2.1.3.4 Maintenance Intervals

For each performance attribute, an objective may be set as part of the SLS. At least one such objective should be defined. Any and all objectives require the definition of a time period over which the objective must be satisfied. Let's represent the value of that time period as T. T can consist of *Available* time (as per the One-way Availability Performance attribute), *Unavailable* time and *Maintenance* Interval time.

Periods of time that are defined as maintenance periods are exempt from performance measurements. That is, any measurements recorded during a maintenance interval must be excluded from the performance metric calculations. Maintenance intervals must be agreed to between the service provider and the subscriber. Typically, that implies the intervals are pre-scheduled (such as scheduled change window), but that doesn't preclude both parties from designating a period of time as a maintenance interval due to an emergent situation.

Furthermore, performance measurements are only valid during *available* time. The determination of the current measurement interval's availability state cannot be made until a future interval, as defined by the

⁶ In our example, each attribute had a single objective defined in the SLS, however multiple objectives may be defined for one or more attributes, such as adding 80% of frames must have an FD less than 100ms, in addition to the 150ms objective for 95% of the frames.

⁷ This use case is a point-to-point service; for multipoint services different objectives may be defined for different sets of UNI pairs, for example 2 sites in the same city may have a lower FD objective defined than the same 2 sites in relation to a 3rd site overseas.

⁸ In this use case, the objectives are set to the same value in both directions of the E-Line, however this is not a requirement.

⁹ The time interval T must be the same for all performance objectives in a SLS.

¹⁰ Although the small time interval is not directly part of FD, FDR, MFD, IFDV and FLR, it is still required in the sense that it is used to determine available time and the other performance attributes are only defined during available time.

parameters configured for the one-way availability performance objective by the CSP. Therefore, other performance metric calculations must be temporarily stored until they are deemed **included** because the interval they measured was available or **excluded** because the interval they measured was marked unavailable. This temporary storage is mandatory for FLR metrics and is recommended for all other performance metrics.

As a result of maintenance intervals being excluded from all performance metric calculations, the objective is met based on the time period T, **excluding** any maintenance interval. This exception enables the provider and subscriber to agree to emergent and/or standard maintenance windows such that those windows will not violate the SLA.

***Note:** Given that the performance attributes other than availability are meaningless when the service is unavailable, the performance measurements are only valid during Available time and must be excluded when the service is deemed Unavailable. That said, there may still be interest in the performance measurements during maintenance intervals for informational/reporting purposes.*

4.2.1.3.5 Composite Performance Metric (CPM)

Most recently the MEF added a new performance attribute to address circumstances where certain performance objectives such as FD, IFDV and FLR can be met over a long period (e.g. a month) yet the related performance measurements can be very poor over short periods (e.g. 30 seconds). CPM is intended track all metrics continuously in a similar fashion that availability tracks frame loss and as a result, determine the proportion time that the performance was acceptable. Subscribers wish to have a greater amount of granularity in terms of the performance metrics such as FD and IFDV even though the overall time period (T) is large.

Certain services such as mobile backhaul, are sensitive to such short periods of degradation. Synchronization traffic is a key use case in mobile backhaul where the one-way frame delay is critical and in particular the inter-frame delay variation must be minimized.

The one-way CPM indicates the frequency in which an EVC meets (or exceeds) the frame delay, inter-frame delay variation and frame loss service performance over a specified time interval. Each of these three elements can be included or excluded from the CPM calculation. By excluding frame delay and inter-frame delay variation, the CPM is equivalent to the one-way availability performance attribute.

5. 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 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 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. With the MEF's definition of performance attributes, the service level specification and

the recently published Carrier Ethernet Performance Reporting Framework specification (MEF 52), Carrier Ethernet providers can begin to 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 Recommendation G.8013/Y.1731 (commonly referred to as Y.1731¹¹) and the TM Forums 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 companion to this paper, "**Part II: How to Measure Carrier Ethernet Performance Attributes**" explores the specifics and practical considerations for performing the measurements of the performance attributes.

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.

6. 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.

¹¹ 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>

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.

7. 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.

8. Glossary and Terms

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

9. 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 II: How to Measure Carrier Ethernet Performance Attributes
ITU-T	Recommendation G.8013/Y.1731: Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks

10. Acknowledgements

Editor: Christopher Cullan, InfoVista

Principal Authors:

- Christopher Cullan, InfoVista

Contributors:

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