Understanding MEF's Service Activation Testing Power Play

November 2016

Table of Contents

Figures

Tables

1 Introduction and Overview

1.1 Abstract

The MEF introduced the MEF 48 (Service Activation Testing [SAT]) along with MEF 46 (Latching Loopback), and MEF 49 (SAT Control Protocol and PDU Formats) as a framework to help service providers with a new standardized approach to testing and troubleshooting Carrier Ethernet-based services. Collectively these standards are known as the MEF SAT Power Play. This white paper describes the portfolio of the SAT Power Play, how each element can be used either individually or collectively, and can be an indispensable resource to a service provider. Specifically, the paper discusses:

- What are the components of the MEF SAT Power Play.
- Why the specifications contained in the MEF SAT Power Play are more efficient than legacy test methodologies.
- How service providers can use the MEF SAT Power Play specifications to successfully turn-up or troubleshoot Carrier Ethernet 2.0 (CE 2.0) services, either individually, or in conjunction with each other.

1.2 Background

This paper targets technical managers, network designers, and operations personnel of Ethernet-based service providers worldwide. The main purpose for this white paper is to provide information for service providers to understand and use the latest Service Activation methodologies and tools. It also aims to explain how the various components of the SAT Power Play can be used individually, or together, to provide a more comprehensive testing solution using a variety of CE 2.0 Network Elements or Test Solutions.

The goal of this white paper is to introduce the SAT Power Play in the context of the related MEF service management technical specifications (TS). Special attention is given to the relationship between the MEF specifications and the work of other standards developing organizations (SDOs), specifically specifications such as RFC 2544 and ITU-T Y.1564.

1.3 Document Objective

This paper provides a summary view of the various SAT Power Play components including the MEF 48 SAT test methodology, the MEF 46 standardized Latching Loop mechanism, and the MEF 49 SAT Control Protocol and PDU Formats. These components are used to control various elements of the test infrastructure. Furthermore, the document details service provider use cases for utilizing these specifications in tandem with each other, or individually. These use cases can help the Service Provider develop a framework for testing CE 2.0 services in a more efficient and reportable manner.

1.4 Executive Summary

Service Activation is critical for a Service Provider to be able to launch a service for their Subscriber, to assure that all Service Acceptance Criteria (SAC), a subset of the Service Level Specification (SLS) technical documentation, have been met. The MEF has developed a variety of tools that allow Service Providers to not only meet those initial expectations, but also make the activation process part of a more efficient operational model called the Lifecycle Service Orchestration (LSO). These tools ensure that a service can meet the high demands of the current Carrier Ethernet Subscriber base, and be ready for the future of automated service creation.

The focus of this paper is to help the reader understand the benefits of an industry standardized mechanism for testing certified CE 2.0 services, and why previously created tools are no longer capable for meeting the expectations for the current Carrier Ethernet Subscriber. The validation and documentation of a verified service, in a manner that is standardized, allows the Service Provider, the Access Provider, and the Subscriber to be confident that a certified Carrier Ethernet 2.0 service is being delivered.

2 Service Activation Testing part of the Service Management Lifecycle

Figure 1 highlights the MEF Service Ops Life Cycle of a CE 2.0 service, specifically calling out step 5, the Service Activation Test step. In the context of this white paper, we will examine SAT as the part of the cycle after a provider has configured the service, but before the service is handed over to the Subscriber. This step is typically referred to as providing a "Birth Certificate" of a service, demonstrating all the characteristics of the service to the Subscriber as a means of ensuring that a service is what was ordered. The SAT step also provides a baseline for the Service OAM (SOAM) components of the Service Life Cycle to detect any issues that may develop after a service handoff.

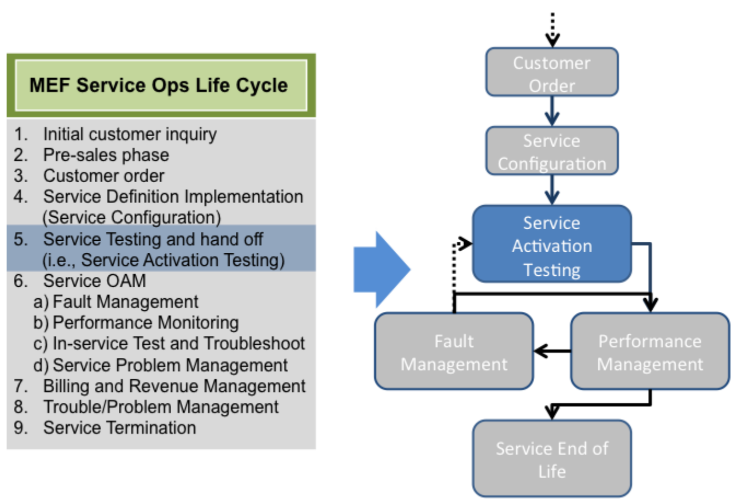

Figure 1 Service Management Life Cycle for Carrier Ethernet 2.0 Services

Note: the MEF's Technical & Operations Committee is working complementary areas including: Generic Service Life Cycle Process Model, Carrier Performance Reporting Framework, Ethernet Serviceability, Standardized Ethernet Service Order Specifications, Standardized Ethernet Product Catalog, etc.

The components of SAT often mimic the limits of the service when being used by a Subscriber, however they are different than the SOAM tools that are used after handoff for service monitoring. SAT provides mechanisms for testing the Bandwidth Profile, Frame Delay (FD), Inter Frame Delay Variation (IFDV), and Frame Loss Ratio (FLR), and other Service Attributes by sending traffic at volumes to match the service order. Conversely, SOAM

Performance Monitoring (PM) uses ITU-T G.8013/Y.1731 mechanisms, typically intermittent frames interlaced with Subscriber traffic, that can characterize the service while it is in use. Both techniques provide information on the Service Level Objectives, but in inherently different manners.

The origins of Service Activation Testing come from the needs of service providers to quantify and qualify service characteristics for Subscribers who felt their service was under performing. The service providers looked to a technical specification, IETF RFC 2544, that was originally drafted to qualify the performance of network elements, such as routers and switches. RFC 2544 suggests testing KPIs serially (versus in parallel) across a specific set of frame sizes (again serially versus interleaved). Service Providers began to apply these methodologies to testing an end-to-end service. As the volume of services grew over time, at the rate at which the services were being turned up these original methods were not practical, in both the length of time it took to test each service (2-24 hours) and the realistic application of traffic onto the network. Service Providers began to ask test vendors to develop a more efficient and realistic test procedure. ITU-T Y.1564 was the result of this work. It took the original principles of RFC 2544, but enhanced the functionality by providing parallel testing of the required KPIs, interleaving the set frame sizes, and limiting the test length to a user defined period between 15 minutes to 24 hours.

As the MEF evolved and developed CE 2.0 services that were carried across multiple Operators, the need to evolve and define testing methods, criteria, tools, and results across these CE 2.0 services became a higher priority. As a result, the MEF developed the components of the SAT Power Play to provide a more robust testing method, a more complete test record, and tools for automation that can be implemented by equipment vendors and utilized by Service providers.

2.1 Service Activation Testing Power Play Components

Ethernet networks are now servicing real-time and sensitive services. Subscriber traffic can be classified under three types: real-time, high-priority, and best effort. Each traffic type is affected differently by the network characteristics and must be groomed and shaped to meet their minimum performance objectives. As these CE 2.0 services can traverse multiple Service providers, it is imperative that standardized testing be performed to create uniformly understood results.

SAC are agreed upon between the service provider and Subscriber to guarantee that the minimum performance will be assured for the services provided.

MEF provides standards to ensure that both parties are using measurements that are uniform and recognized. The MEF presents three standards that are used during service activation to meet this objective: MEF48, 46 and 49.

- **MEF 48** defines the method and parameters to measure in order to confirm the compliance of the performance of multiple streams with different agreed traffic priorities.
- **MEF 46** defines a loopback method to use as a far end device with a test instrument that executes MEF 48 Service Activation Method.

- **MEF 49** defines an in-band mechanism to perform unidirectional test. This method provides the ability to fully test the Bandwidth Profile at the far-end which is favorable when testing asymmetric links or isolating each direction of a service.

2.1.1 MEF 48: Service Activation Methodology

The MEF 48 specification defines business processes to ensure that the service is configured as specified and meet defined service acceptance criteria. MEF 48 is built on the foundation of ITU Y.1564. As Ethernet services are bursty in nature, MEF 48 defines three levels of conformance for traffic, to allow maximum utilization of the available bandwidth, while guaranteeing that the negotiated SLS will be met. These three compliance levels are assigned specific colors: green for committed traffic, yellow for excess traffic and red for discarded traffic.

- Green Traffic/Committed Information Rate (CIR): refers to bandwidth that is guaranteed available at all times for a specific service; for green traffic, performance objectives (i.e., key performance indicators or KPIs) are guaranteed to be met.
- Yellow Traffic/Excess Information Rate (EIR): refers to excess bandwidth above CIR that may be available depending on network loading and usage; for yellow traffic, performance objectives are not guaranteed to be met.
- Red Traffic/Discarded: refers to traffic that is above the CIR+EIR rate, and that cannot be forwarded without disrupting other services; red traffic is therefore discarded.

Traffic Class	Bandwidth	Performance Objective	KPI Level
Green Traffic	0 to CIR	Guaranteed Forwarding	KPIs are Guaranteed
Yellow Traffic	CIR to CIR+EIR	Best Effort	KPIs are Not Guaranteed
Red Traffic	>CIR+EIR	Discarded Traffic	Not Applicable

Figure 2 Three color compliance levels for service

MEF 48 defines attributes to measure and confirm that the service meets the set criteria for each specific traffic type. These attributes are: Information Rate (IR), Frame Loss Ratio (FLR), Frame Delay (FD), Inter-Frame Delay Variation (IFDV), and Frame Delay Range (FDR).

- Information rate (IR) (also known as Bandwidth). It refers to the maximum amount of data that can be forwarded. This measurement is a ratio of the total amount of traffic forwarded during a measurement window of one second. Bandwidth must be controlled, because of the multiple services sharing the link under test. Therefore, each service must be limited to avoid affecting another service. Generating traffic over the bandwidth limit usually leads to frame buffering, congestion and frame loss or service outages.
- Frame Loss Ratio (FLR). Frame loss can occur for numerous reasons, such as transmission errors or network congestion. Errors due to a physical phenomenon can occur during frame transmission, resulting in frames being discarded by networking devices such as switches and routers based on the frame check sequence field comparison. Network congestion also causes frames to be discarded, because networking devices must drop frames to avoid saturating a link in congestion conditions.

- Frame Delay (also known as Latency): Measurement of the time delay between a packet's transmission and its reception. Typically, this is a round-trip measurement, meaning that it includes both the near-end to far-end and the far-end to near-end directions. This measurement is critical for voice applications, as too much latency can affect call quality, leading to the perception of echoes, incoherent conversations or even dropped calls. Other real-time applications such as exchange trading and real-time video are susceptible to Frame Delay issues. Frame Delay must be measured in 1 of 2 ways: 1) Mean Frame Delay (MFD) is the arithmetic mean, or average of delays experienced by a set of frames OR 2) Frame Delay (FD) which has 2 parameters defined: the amount of delay and the percentile for frames that must meet that criteria.
- Inter-Frame Delay Variation (IFDV) (also known as Packet Jitter): The variability in arrival time between packet deliveries. As packets travel through a network, they are often queued and sent in bursts to the next hop. Random prioritization may occur, resulting in packet transmission at random rates. Packets are therefore received at irregular intervals. Real-time applications such as voice and video are especially sensitive to IFDV.
- Frame Delay Range (FDR): This is the range of delay that is observed at the endpoint as packets arrive. This measurement is an alternative way of looking at IFDV, and provides a similar function. Where IFDV is looking at the time gaps between the start of each packet, FDR looks at the difference between the delay of the current packet, and the minimum delay observed of all the packets. Only one method must be used between IFDV and FDR to calculate variation in packets arrival times.

MEF 48 Service Activation Test is executed, much like Y.1564, in two phases: Service Configuration and Service Performance:

- Service Configuration: Short duration tests are done for each stream, one after the other, to ensure they are not interfering with each other. This is to ensure that the stream is properly configured by validating it has connectivity with its far end peer. Tests are executed to confirm that SLS can be met: streams are filled up to Committed Information Rate (CIR), measurements are made to validate the service does not exceed its target KPI values, and thus meeting the SLS requirement. This test is run for each frame size that is selected individually. Typically, these tests last under a minute.

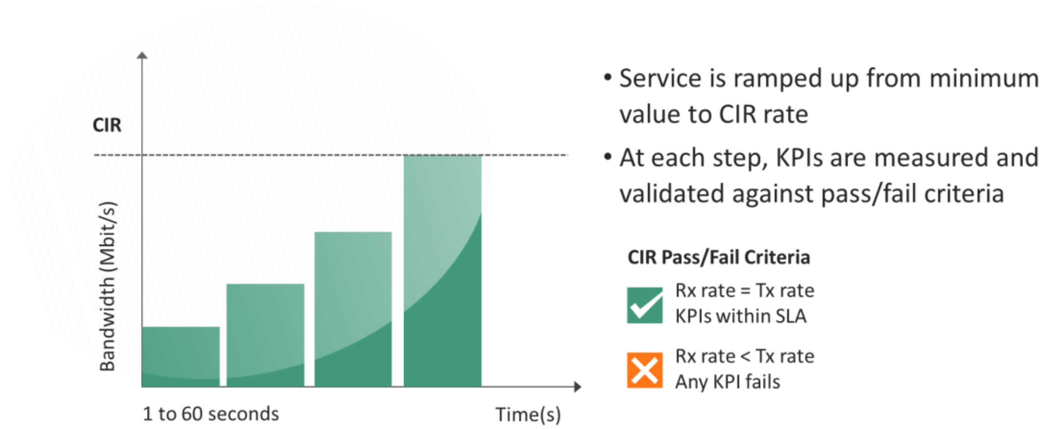

Figure 3 Confirming SLA: Traffic less than CIR and attribute met

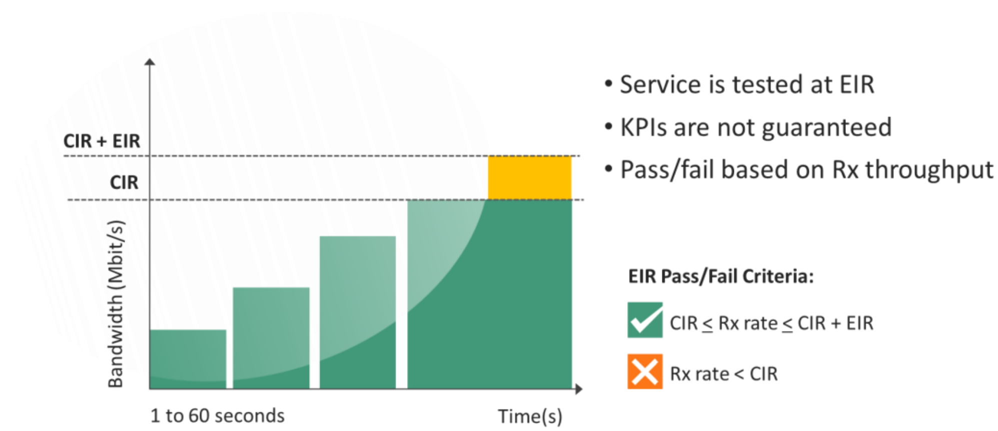

Figure 4 Behavior with EIR

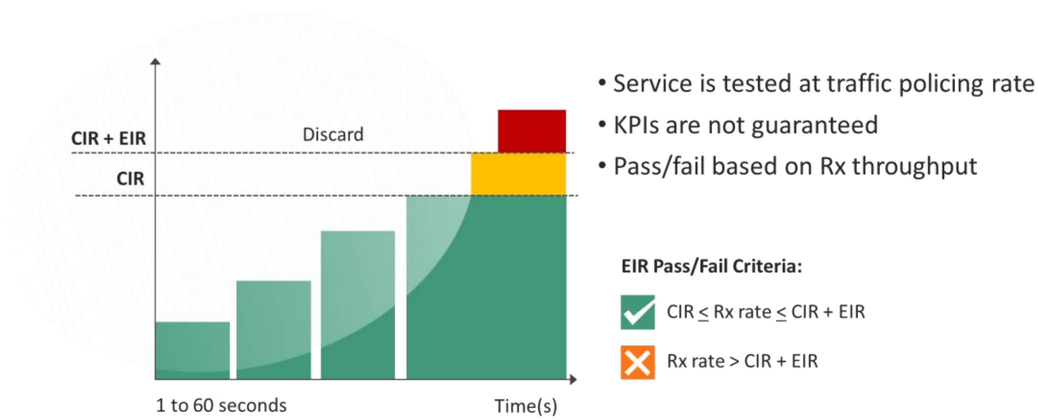

Figure 5 Traffic policing

- Service Performance: Tests are executed over a longer period: 15 min to 24 hours. The goal of that phase is to ensure that all services can be simultaneously activated and achieve Committed Information Rate while meeting all performance KPIs, thus meeting the SLS for each services. The tests are run using a single frame size, or a pattern of up to 8 frames, with varying standard sizes.

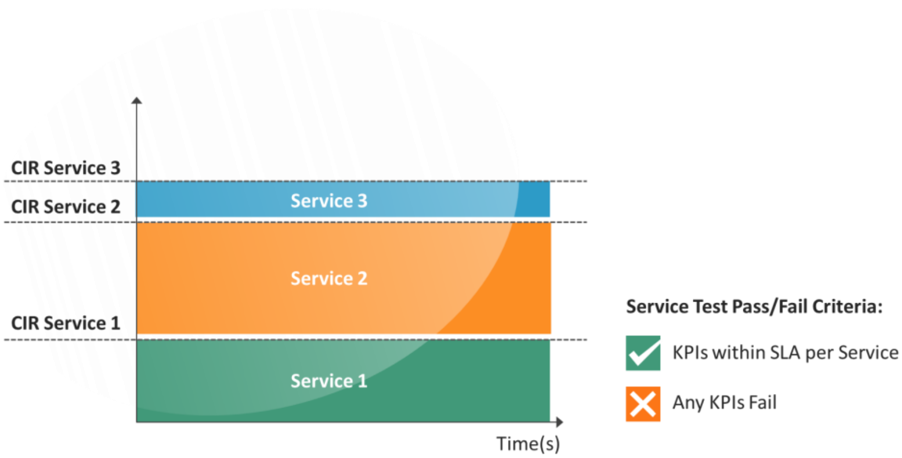

Figure 6 Performance test with all services activated at CIR

Once the configuration and performance tests are completed successfully, a test record is generated. This record becomes the "birth certificate" for the service. It would be used as a baseline to facilitate future performance comparisons.

As MEF 48 uses and extends test process and procedure based on the Ethernet Test Methodology defined by Y.1564, the reference table below is provided to help distinguish the differences between RFC 2544, Y.1564 and MEF 48.

	IETF RFC 2544	ITU-T Y.1564	MEF SAT
Testing Method	Sequential testing, tests cannot be run simultaneously	Tests run for all configured service simultaneously	Tests run simultaneously for an E-Access Service
Throughput	Maximum rate at which none of the offered frames are dropped by the device.	Maximum throughput must respect CIR + EIR	Maximum throughput must respect CIR + EIR
Frame Delay	Latency measured on 1 frame every 2 minutes	Measured on all flows simultaneously	Can be measured on all frames simultaneously with other performance attributes
Inter-Frame Delay Variation	Not defined	Measured on all flows simultaneously	Can be measured on all frames simultaneously with other performance attributes

Frame Loss	Frame Loss methodology measurement on lack of resources only	Frame Loss methodology measurement on lack of resources and service quality	Frame Loss methodology measurement on lack of resources and service quality
CoS Support	Not supported	Supported	Supported
Bandwidth Profile	Extendable to CIR/EIR	Confirms CIR/EIR (CBS/EBS)	Confirms CIR/EIR (CBS/EBS Optional)
Frame Size	Tests are run using a set of defined frame sizes independently	Service Configuration test is run using a set of defined frame sizes with the option of a custom and/or a defined Maximum Frame Size independently. Service Performance test is run using a single frame size, or a sequence of those frame sizes that is user determined (EMIX).	Service Configuration test is run using a set of defined frame sizes with the option of a custom and/or a defined Maximum Frame Size independently. Service Performance test is run using a single frame size, or a sequence of those frame sizes that is user determined (EMIX).
Service Attributes	Not Tested	Not Tested	Maximum Frame Size, CE-VLAN ID Preservation, CE-VLAN CoS Preservation and Unicast, Multicast and Broadcast delivery
Record MEF CE 2.0 Service Attributes	Not Supported	Not Supported	Supported
Defined Report Template	Not Supported	Supported	Supported

Table 1 – Differences between RFC 2544, Y.1564 and MEF 48

2.1.2 MEF 46: Latching Loopback

MEF 46 Latching Loopback defines the protocol and functionality to provide a loopback function. When performing Service Activation Testing, two functions are involved in the test: a Generator Test Function (GTF) and a Collector Test Function (CTF). If the service to be tested is only able to be tested in a roundtrip manner, a Latching Loopback Function (LLF) is required to redirect traffic back towards the initiating UNI where the CTF is located. The far end LLF can be located in a Network Interface Device (NID), an UNI-N, an ENNI, or an Ethernet Test Equipment (ETE). The initiating ETE is the Latching Loopback Controller (LLC) and the far end device is the Latching Loopback Responder (LLR).

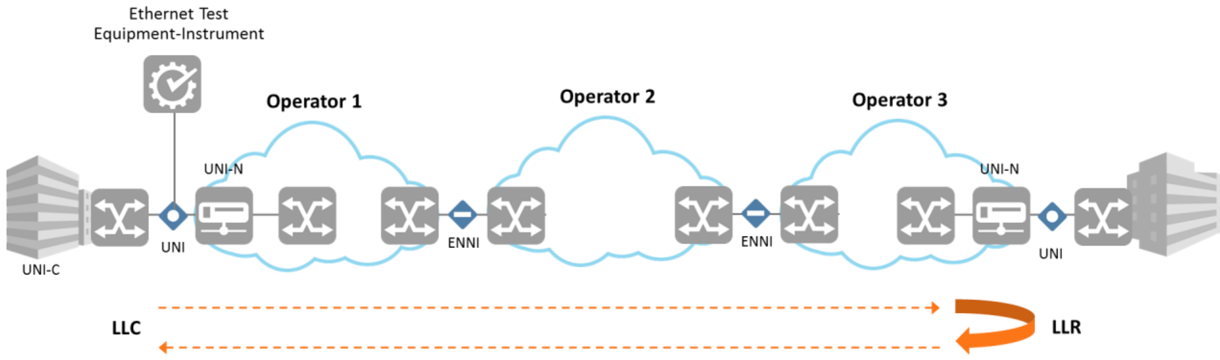

Figure 7 Example of a LLC and LLR

The LLC generates Ethernet frames that are sent to the LLR. Latching Loopback Messages (LLM) and Responses (LLR) are exchanged in-line with SOAM frames. The LLR has a Latching Loopback State Machine (LLSM) that processes the Test Frame and sends back an unaltered Ethernet Frame Payload. The LLSM has three states: Prohibited, Inactive and Active.

- Prohibited: This is the initial state of the LLSM. Latching Loopbacks messages addressed to a LLSM in that state will be discarded.
- Inactive: This is the state where latching loopback are permitted, and there is no loopback request currently active. An administrative action on the management interface of the LLR must be taken to change from prohibited to inactive.
- Active: This is the state where a loopback is active.

2.1.3 MEF 49: SAT Control Protocol and PDU Formats

The Latching Loopback method provided by MEF 46 allows measurement of MEF 48 Ethernet attributes for round-trip connection. MEF 49 defines protocol and functionality to provide uni-directional measurements of those attributes. It also provides a means for the far end device to send the measurement results to the local device by employing a CTF in the forward path (the direction from the "near end" where the test is initiated towards the "far end"), and a GTF in the backward path (the direction from the "far end" where the test is initiated towards the "near end"). MEF 49 also defines standard SAT Test PDU to enable interoperability between ETE vendors.

In MEF 49 the Controller End State Machine is responsible to initiate the test session with the far end Responder End State Machine. During the initiation, the controller provides attributes of the test session and then a unique ID is created and assigned to that test session. This ID allows the instantiation of multiple test sessions that can be executed concurrently.

For one test session, the Controller end can act as a GTF or as a CTF. The Responder will then act as the peer function needed. The GTF generates and transmits test frames and the CTF receives them. There are two test frames that are defined: Frame Loss PDU (FL-PDU) and Frame Delay PDU (FD-PDU). The FL-PDU are used to perform frame loss and throughput measurements, and is a new PDU defined by MEF 49. The FD-PDU are used to perform Frame Delay measurements, and utilizes the ITU-T G.8013/Y.1731 DMM/DMR message to avoid duplicating the delay mechanisms already built into a MEP.

The Controller and Responder have different state machines that are used to manage the execution of the test sequence that is required for the MEF 48 test. The following "stages" are defined for the Controller:

- Create: The controller transmits the parameters of the test session to the responder. The responder replies to the controller to confirm its capability to execute the test with the required parameters.
- Pending: When the controller is defined as a CTF, the controller will send a start instruction to the responder, as this last one is defined as the GTF.
- Start: This state is to start the transmission of the test frames
- Run: This state is the period while the tests are being executed.
- Stop: This state occurs to inform the that the test execution is completed
- Fetch: This state is used by the controller to fetch test results from the responder.
- Delete: This state is used by the controller to inform the responder that the result is no longer required and has to be deleted.
- Abort: This state is used to abort the test execution and, if required, send a stop notification to the responder.
- Status: This state is used by the controller to report the status of the test execution.

The Responder has less stages:

- Start: This state is to start the transmission of the test frames.
- Run: This state is the period while the tests are being executed.
- Fetch: This state is used to transmit test results to the controller.
- Delete: This state is used to delete test result.

MEF 49 was also designed to provide the Service Provider with a more complete understanding of each direction of the service. Attributes of a service can be defined uniquely in each direction, but without MEF 49, validating those attributes was limited to a round-trip fashion. Since MEF 49 is based on having a unique GTF and CTF in each direction of the service, each path can be measured independently. Subsequently, qualifying attributes such as Bandwidth Profile (BWP) at the far end of the service can only be done using MEF 49.

3 Making the most of the SAT Power Play: Uses Cases

The following sections will take a deeper look into application of the SAT Power Play components: how MEF 48, 46, & 49 can be used individually or matched together. The use cases will show how a Service Provider can be more successful in implementing a test strategy and deploying a service with confidence.

Let's look at a Carrier Ethernet service delivered to a Subscriber where the service is chained through different service providers or even transmission subnetworks. Without the use of MEF SAT Power Play components, each segment of the service must be individually qualified to ensure compliance to the contracted SLS, which increases effort, complexity and ultimately time to deliver the service to the Subscriber. It is critical at this point to use a sound methodology agreed and understood by all parties of the service chain to ensure proper end-to-end service delivery. The following Use Case strategies will focus on services where a Service Provider (SP) is providing service to a Subscriber that is served at one location in the SP's market. The other location is in a different market that is covered by an Access Provider (AP). The AP is providing an E-Access service to cover the last mile connection to the Subscriber.

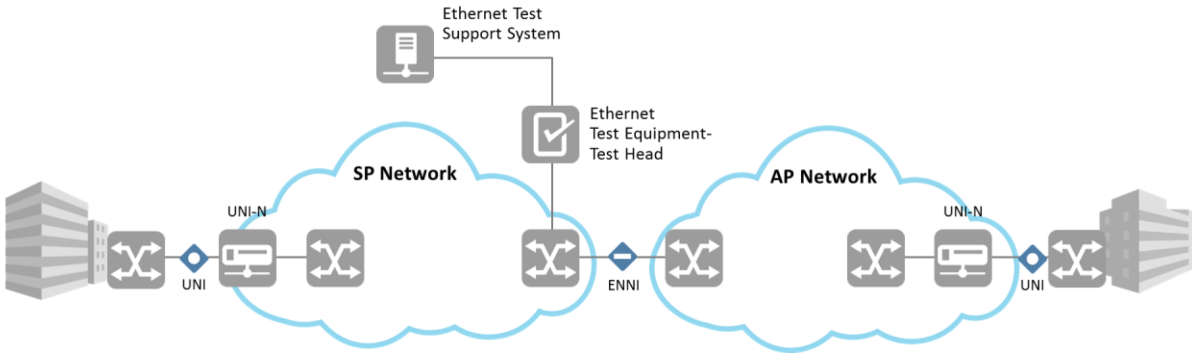

Figure 8 Sample Network Design for all Use Cases

3.1 MEF 48 (SAT)

MEF 48 SAT allows a SP to validate a service that they will be providing with the help of a AP, where the AP is providing the last mile of access to a location for the Subscriber that is outside the SPs service area. In the traditional sense a service provider could use a legacy testing philosophy such as Y.1564, however MEF 48 allows the SP to not only test performance metrics, but test other aspects of the Service Level Specification (SLS), that include the SAC.

Once the AP has handed over the E-Access Service to the SP, it is up to the SP to test the service as it initiates in the SP network, traverses the ENNI, and terminates at the Subscribers UNI. As part of the test, an ETE, which could be controlled locally or by an Ethernet Test Support System (ETSS), will initiate the MEF 48 functions. ETEs can take several different forms to support a wide variety of deployment models. An ETE can be dedicated hardware or it can be software functions that are embedded within the network elements.

MEF 48 will certify in a short period of time the configuration of the service proving the availability of capacity for the end-to-end EVC, and ultimately proving that all of the forwarding elements between the two endpoints are correctly configured regardless of the provider domain, using the Service Configuration test. Additionally, MEF 48 SAT also introduces a Service Performance test, which proves compliance to the SAC for the end-to-end service over a longer period, which amounts to a stress test for all elements. However, without a standard mechanism on the far end AP UNI to test to, the SP must be in contact with the AP to provide a mechanism to loop the test traffic back towards the originating ETE. The AP will have to instantiate a loopback either on the UNI or behind the CPE using some type of test equipment. Most CPEs that are installed today have the ability to provide the looping mechanism needed for testing, and can be configured remotely (either through CLI, SNMP, or Web UI) so a truck roll is not necessarily required by the AP to put in an external piece of testing hardware.

At this point the SP will continue to run the MEF 48 test to meet the SAC, working through both the Service Configuration and Service Performance Tests. The testing will comprise a combination of both testing attributes that are part of the SAC (Bandwidth Profile, FLR, Frame Delay, etc.) and testing methodologies that are inherent to the SP (which standard frame sizes to test during Service Configuration, the combination of frames sizes to use for EMIX during Service Performance, how long to test, etc.).

Once the testing is complete to the satisfaction of the SP, the SP can then communicate to the AP to remove the loopback that was put in place for the testing. The SP can then develop the "birth certificate" SAT Record, and provide that as part of the Service Delivery to the Subscriber. An Example of a SAT Record can be found in MEF 48, Annex B.

3.2 MEF 48 + MEF 46 (SAT + Latching Loopback)

The combination of MEF 48 and MEF 46 enables operators to significantly reduce the costs associated to activating and troubleshooting services. An ETE located at the UNI/ENNI can be leveraged to perform MEF 48 SAT towards the UNI at the far end of the service, proving the entire service chain. A traditional method to perform SAT is to employ devices such as Ethernet Test Equipment as the loopback device, or manually employ a loopback function on a NID. This ultimately increases the cycle time and requires coordination between the GTF endpoint and the loopback endpoint. MEF 46 plays a key role by reducing reliance on additional tools. The ETE can remotely activate a loopback by generating and targeting the proper endpoint with a specific Latching Loopback request message (LLM), execute the SAT process and normalize the service after completion, thus eliminating the wait time and effectively eliminating the truck roll. The End-to-End service can be quickly and efficiently activated in minutes rather than hours or days and with lower costs.

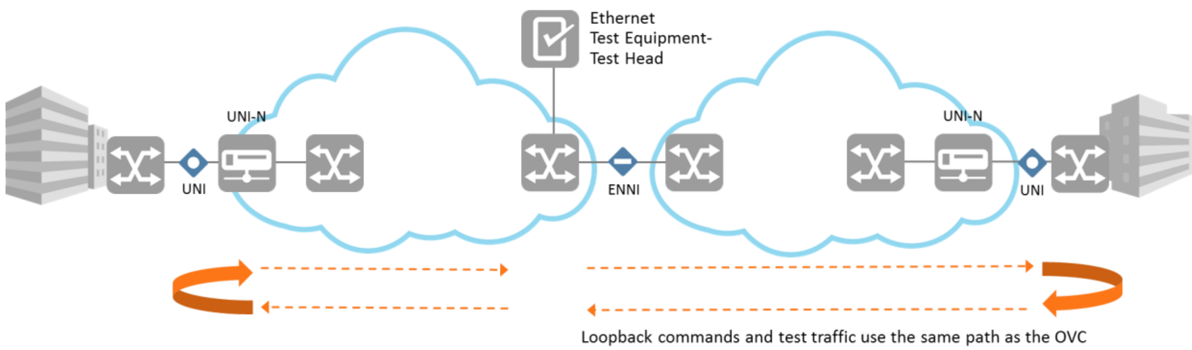

Figure 9 Example of MEF 48 + MEF 46

Even with the most thorough Service Activation process, networks degrade over time and the introduction of new equipment, human error or simple faults will impact the performance of the end-to-end service. MEF 48 and MEF 46 also contribute significantly to help troubleshoot services efficiently. The traditional approach relies on analyzing a significant number of counters, estimating regions of concerns and dispatching ETE to troubleshoot domains. This could also require coordination from multiple teams or even service providers which significantly impact the MTTR. Additionally, such methods could result in service outage for the Subscriber as the service is placed out-of-service for testing, all leading to an unsatisfied Subscriber. An operator leveraging MEF 48 and 46 can greatly improve the troubleshooting process, by making sound use of the MEF 46 LLB capability. As in the same scenario as above, an ETE located at the UNI/ENNI can be used to perform a full activation test using MEF 48 to certify availability of network resources on the end-to-end service chain as well as compliance to SLS. However, the addition of LLB brings additional flexibility where a remote UNI or ENNI can be remotely placed in loopback for testing. The operator can then look at performance from an end-to-end perspective and then target the prior ENNI, force it into loopback, then perform another SAT to that particular element. This segmentation method allows the operator to view compliance and performance from the Ethernet Test location to successive segments and identify quickly and easily the segment that is contributing to the degradation of the service. Leveraging MEF 46 LLB accelerates the process by eliminating the need to dispatch additional tools or reliance on multiple teams.

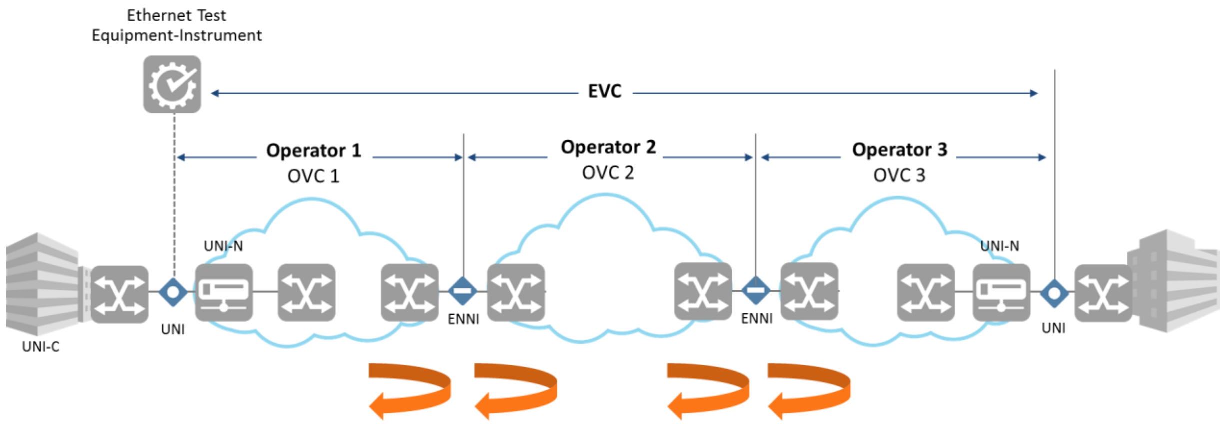

Figure 10 Troubleshooting Carrier Ethernet Services with MEF 46

Using MEF 48 and MEF 46 is a sound and efficient approach to tackle the delivery of end-to-end Ethernet Services. While any loopback has its limitation, it is a simple and often valid method of proving compliance in both directions of the service chain. As a first step, it will optimize the entire process by enabling operators to quickly activate and troubleshoot service with loopback and resort to additional methods for unidirectional testing only in the event of unexplainable failures or complex problems. MEF 48 and MEF 46 will help operators considerably simplify process and ultimately reduce costs and delays in delivering the best end-to-end services to their subscribers.

3.3 MEF 48 + MEF 49 (SAT + SAT PDU)

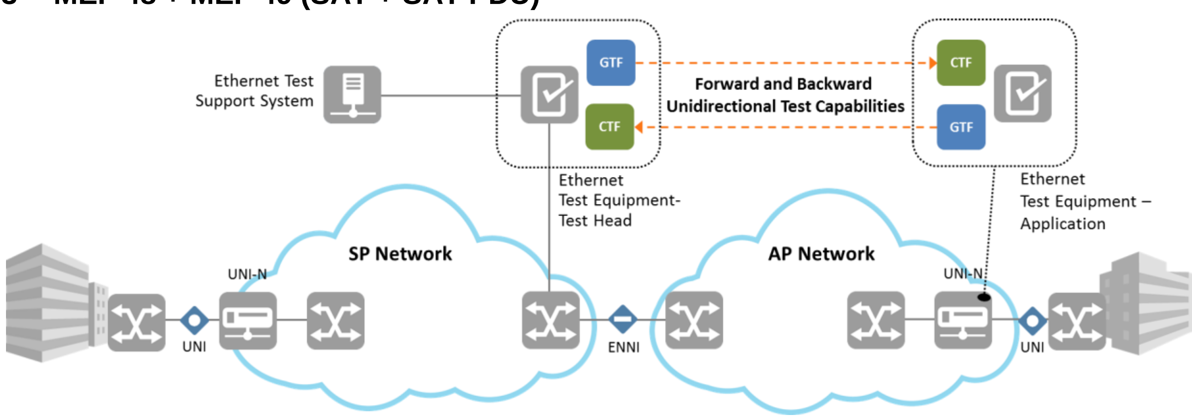

Figure 11 Example of MEF 48 + MEF 49

A Service Provider has requested a service from an Access Provider to complete an order for a customer. The customer is requesting a larger bandwidth downloading to the far end location, than what is needed in the reverse path.

In order to validate the service from the SP's ENNI location to the remote end of the service at the Access Provider UNI, the most complete way to test all attributes of a service is by measuring each direction of the service, including the Bandwidth Profile (BWP) in each direction. The MEF 49 SAT PDU function enhances the capability of the MEF 48 methodology to provide a more complete analysis of the end to end service. Most Service Providers/Access Providers do not typically deploy the same vendor equipment, whether it is a CPE providing an Ethernet Test Equipment - Application (ETE-A) function, or a test vendor based Ethernet Test Equipment – Test Head (ETE-TH), or an ETSS. Having dissimilar equipment could have compatibility issues when trying to instantiate a Loopback function for a test, and for providing unidirectional results. By installing equipment at both ends that are MEF 49 (SAT PDU) enabled, it allows for an exchange of information about how to perform and measure a test, regardless of the vendor equipment. Concurrently, the need to validate both directions of the service is imperative in order to validate the asymmetrical attributes that the customer has requested.

When initiating a MEF 48 test, while using the MEF 49 protocol, the Service Provider can initialize a test to the remote ETE-A on the AP's UNI. This allows the SP to setup a GTF on the SP Test Equipment (in this case an ETE-TH) to the CTF on the AP's vendor equipment (in this case an embedded ETE-A) in the "forward" direction. There is also an equivalent GTF/CTF for the opposite "backward" direction. By providing unique test results for each direction of the EVC, the Service Provider can more accurately diagnose where trouble spots are in the network. This solution also allows a Service Provider to be able to test an asymmetrical service, or a service that provides different attributes in either direction.

Simply by having the MEF 49 protocol on the far end AP UNI the amount of time and resource involvement is drastically reduced. The SP is allowed to run the test without the required assistance of the AP. This level of automation plays into the ability to easily turn up a new service, but also allow for quick turnaround on upgraded services. These services could be totally orchestrated by the LSO.

3.4 Using the SAT Power Play Tools after a Service Change

After the CE 2.0 certification of a service(s) of a well-defined EVC or OVC on a production network, a SP has a big challenge to face: to keep all the new service instances activated on the entire network domain at the same quality as the one that has been once certified by MEF accredited lab, see picture.

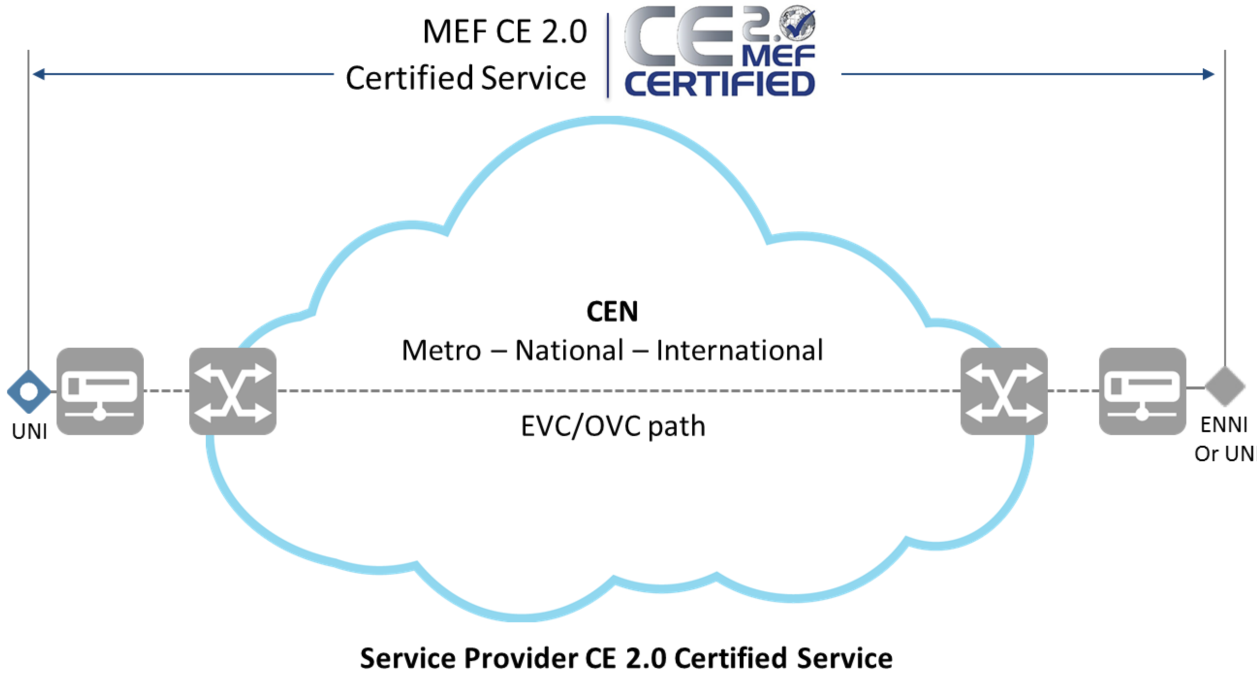

Figure 12 Example of Validating a Service Modification with MEF 48

EVCs or OVCs can be activated on infrastructures that can differ from the ones of the certified one for a number of reasons that are related to the living nature of the network; just to mention some examples, not exhaustive:

- Upgrades of software releases. This is the usual situation in the real networks due to the natural lifecycle of the software;
- SPs can always decide to change or update their switches/routers for any reason;
- Different technology domains. SPs can have different technologies covering their end-to-end network domain between the EVC/OVC External Interfaces (e.g. different access technologies);

Any differences like the ones listed above, with reference to the EVC/OVC once certified, could possibly have a negative impact on the end-to-end service requirements, and SPs look for methodologies and tools to manage network evolution maintaining the end-to-end service performance.

A possible use of SAT for a SP is to support and speed up the internal validation process during the usual engineering activities to qualify network changes like the ones listed before. The following examples should clarify how SAT can assist SPs that have already acquired CE 2.0 certification on a service instantiated on a particular network technology domain between EIs.

Case 1: the SP needs to update equipment at the UNI and/or NNI (or even change the software release) vs the ones used during the certification by the MEF accredited lab.

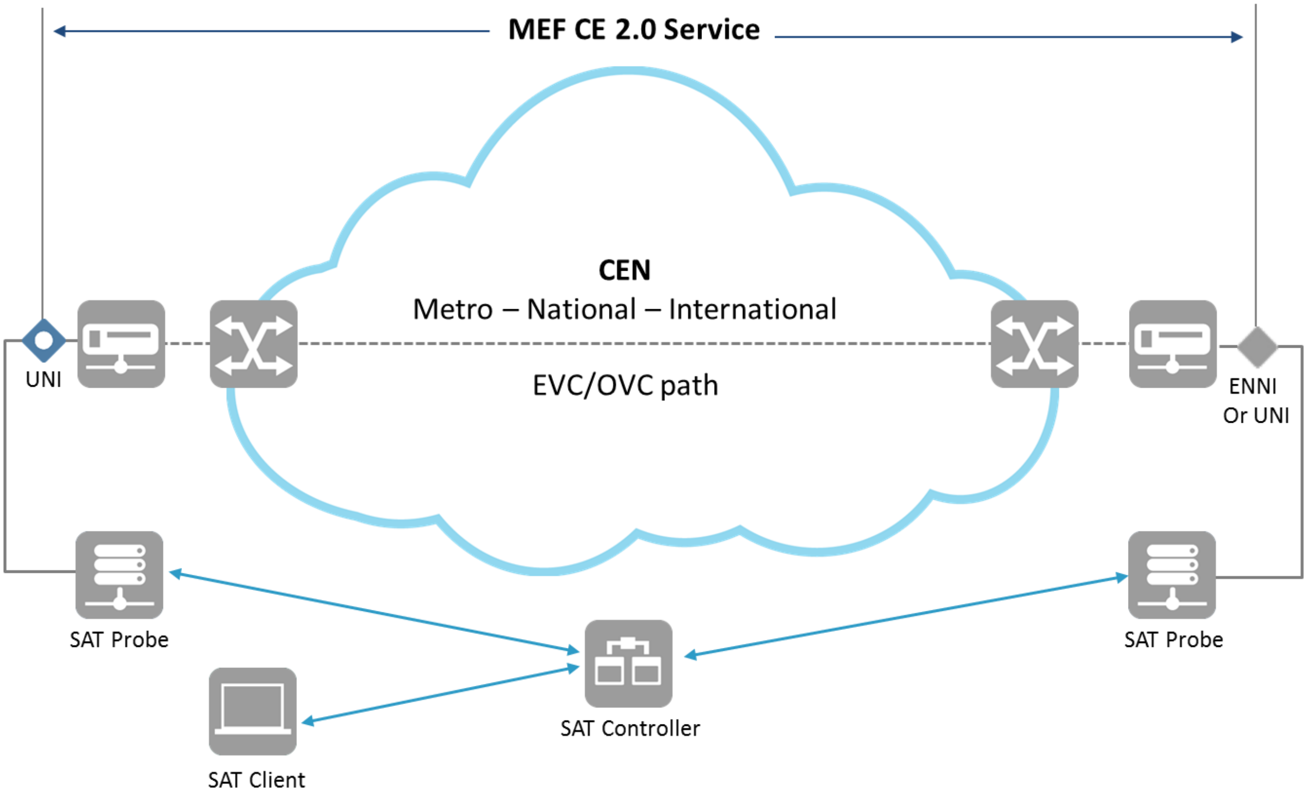

Figure 13 Case 1: Testing a Carrier Ethernet Service with MEF 48

Case 2: SP wants to deliver the same service, once certified by the MEF accredited lab, to a Subscriber reached by a different access technology. A dedicated test session should assure SP that all the service requirements and performances are met in this new network chain.

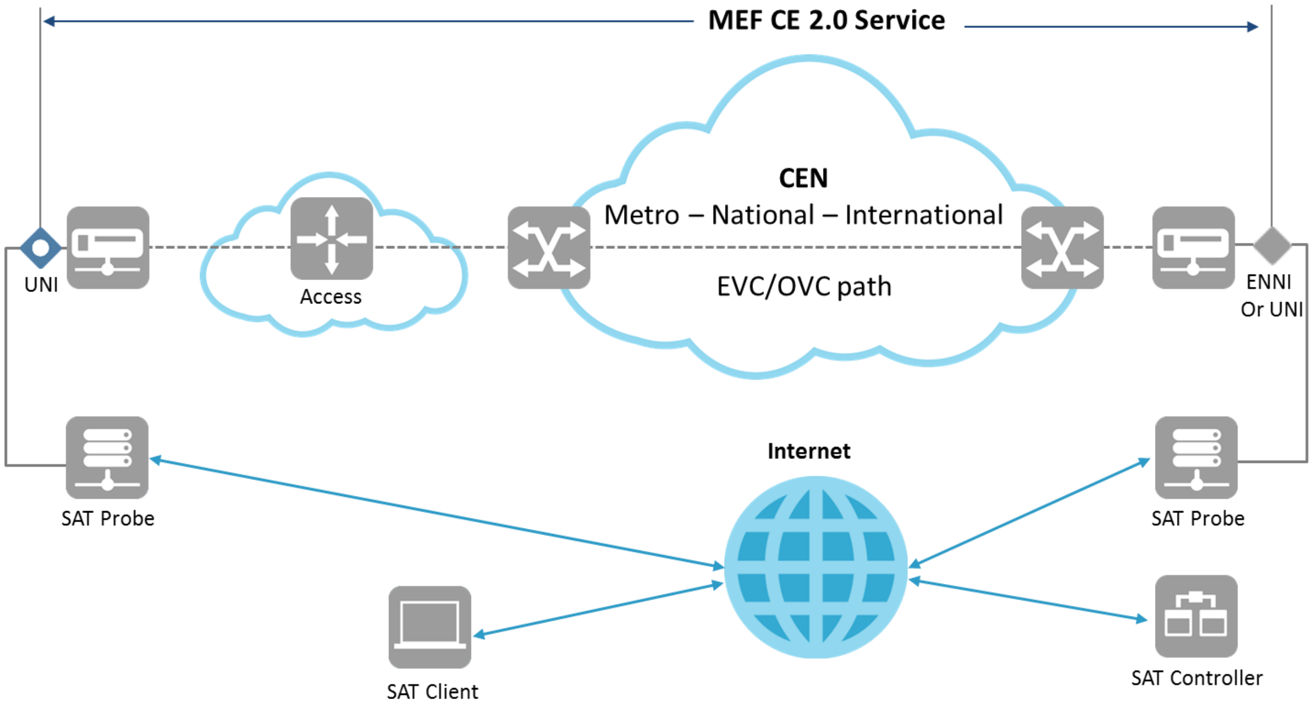

Figure 14 Case 2: Service Validation with MEF 48

For both cases, the use of SAT tools in the SP's lab test beds can support SPs to verify the compliance of the new network chain to the service requirements. SAT probes should be able to run comprehensive test suites, tailored to the service, and support SPs in possible troubleshooting.

3.5 Comparing SAT Power Play Use Cases

Below is a quick reference matrix of features built into the various SAT Power Play components to aid in understanding what combinations can be used for a service provider's test scenario.

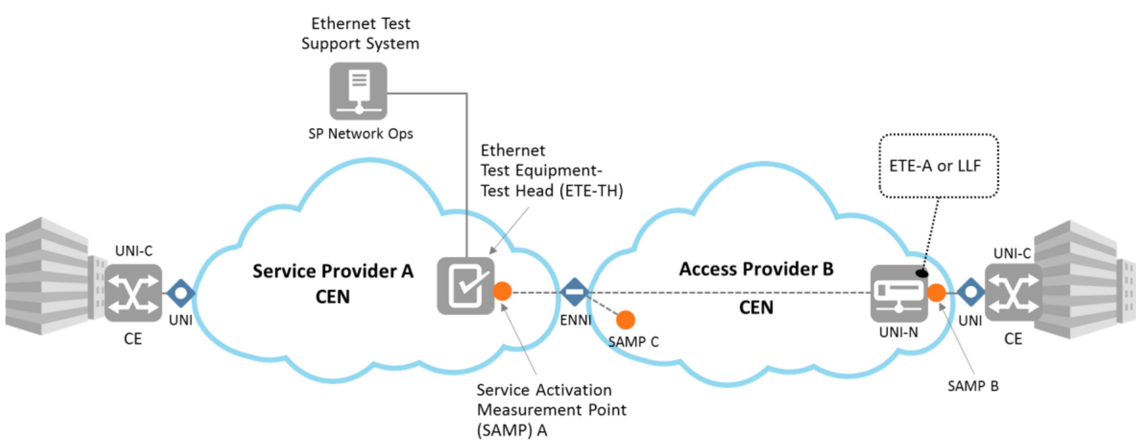

Figure 15 Typical Network Topology

Feature	MEF 48 + 46 SAT + LL	MEF 48 + 49 SAT + SAT PDU
Reporting of Full Suite of Required MEF KPIs and Service Attributes – In a round-trip fashion[1]	X	X
Reporting of Full Suite of Required MEF KPIs and Service Attributes - In each direction of the service		X
Reporting of Frame Delay KPI as Two Way	X	X
Reporting of Frame Delay KPI as Two Way and One Way[2]		X
Symmetrical Service (Service has the same attributes for each direction)	X	X
Asymmetrical Service (Service has different attributes for each direction)		X
Automated Loopback - End-to-end automation of SAT test. ETE/ETSS can initiate the far end ETE as part of the test.	X	
Automated Unidirectional Test - End-to-end automation of SAT test. ETE/ETSS can initiate the far end ETE as part of the test.		X
[1]Testing attributes with a Latching Loopback can only be tested at the "near end" where both the GTF and CTF reside. [2]One Way Frame Delay requires both ETEs be synchronized using a supported timing mechanism		

Table 2 – MEF Specification Capability Matrix

The following matrix provides the pros & cons of each use case:

Use Case	Pros	Cons
MEF 48 (SAT)	• Reporting of Required Full Suite of MEF KPIs and Service Attributes • Simplest way to have MEF SAT Tested Service	• Far end must be manually looped up, typically requiring interaction between a Service Provider and an Access Provider to perform a test. • KPIs and attributes cannot be measured at the far-end
MEF 48 + 46 (SAT + LL)	• Reporting of Required Most of the MEF KPIs and Service Attributes[1] • Allows the Service Provider to perform a SAT test without involving the Access Provider in the test.	• Only Two Way KPIs are available; One Way KPIs and attributes cannot be measured at the far-end • Cannot isolate the direction of issues that might discovered during testing. • Cannot test each direction of a service, including asymmetrical services that have different attributes in each direction.
MEF 48 + 49 (SAT + SAT PDU)	• Reporting of Full Suite of MEF KPIs and Service Attributes for each direction of the service independently. • Allows the Service Provider to perform a SAT test without involving the Access Provider in the test. • Can provide One Way Delay measurements when the near-end and far-end SAMPs are using synchronized timing mechanisms	• Most complex use case as an ETE function must be available on both ends of a service

Table 3 – Pros and Cons Use Case Matrix

[1]The Bandwidth Profile (BWP) at the far end LL point cannot be measured.

4 Summary and Conclusions

Service Testing, as part of activating a service or troubleshooting a service that is already in service, is a critical component in providing the quality associated with a MEF CE 2.0 service. With the MEF SAT Power Play combination of tools, a Service Provider can have the confidence that the service they are providing to a Subscriber meets the MEF guidelines as well as documents the characteristics and attributes. The MEF is continuing to build on this service testing base as it looks to the future. The components of the SAT Power Play will be expanded upon to include more service types and service attributes in the near future.

The MEF has taken the fundamentals from well-established practices in the legacy methods of RFC 2544 and the modern implementation ITU-T Y.1564 to create a rich testing practice in MEF 48 SAT. When combined with a new and standardized approach on both controlling the far end of a service, and even allowing that far end to be involved in the measurement, the MEF has developed a complete solution that allows Service Providers the building blocks for an automated test solution. These building blocks provide the foundation testing methodology as the industry migrates to the MEF's Third Network vision, as to where the Service Testing can be orchestrated within the LSO.

5 Guide to the MEF Technical Specifications

A summary is provided of each specification that defines elements of performance management. The reader can consult each particular specification for further details.

MEF 10.3 - Ethernet Services Attributes Phase 3

Defines the performance attributes used to specify end-to-end CE services.

MEF 23.2 - Carrier Ethernet Class of Service - Phase 2

This Implementation Agreement (IA) specifies a set of Class of Service Names called CoS Labels that can be used by Operators, Service Providers and their Subscribers to indicate the performance expectations associated with a given set of frames that comprise a CoS Frame Set. This CoS IA includes standards for CoS and Color identification as well as performance objectives and supporting requirements.

MEF 26.2 – ENNI and Operator Service Attributes

Defines the performance attributes used to specify inter-provider CE 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 46 - Latching Loopback Protocol and Functionality

This specification describes the use cases, functionality and protocol for a Latching Loopback feature in Ethernet Equipment. Latching Loopback can be controlled to return selected Ethernet frames in the direction from which they came.

MEF 48 - Carrier Ethernet Service Activation Testing

This specification defines the methods, procedures, and reporting mechanisms for testing a MEF 33 service that has been purchased by a Service Provider from an Access Provider

MEF 49 / 49.0.1 - Service Activation Testing Control Protocol and PDU Formats

The control protocol provides the ability to configure and control the SAT steps and to fetch test results at the completion of the test. The requirements defined within this document are based on the SAT process as defined in Section 10 of MEF 48.

6 About the MEF

The MEF is the defining body and driving force behind the global market for Carrier Ethernet. MEF's flagship work is CE 2.0, including specifications and related certification programs for services, equipment and professionals (MEF-CECP 2.0). An industry alliance consisting of more than 225 member organizations, the MEF operates through a powerful collaborative framework of service providers, network solutions suppliers and other stakeholders to achieve its CE 2.0 development and globalization objectives.

Building on thirteen years of success and widespread adoption of CE 2.0, MEF is now focused on defining Lifecycle Service Orchestration with APIs for existing network, NFV and SDN implementations enabling Agile, Assured and Orchestrated Network as a Service. For more details, visit MEF.net.

7 Glossary and Terms

A glossary of terms used in this document can be found online at http://www.mef.net/carrier-ethernet/terms-used-in-mef-specifications.

8 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	MEF Technical Specifications
ITU-T	Y.1564: Ethernet Service Activation Test Methodology
IETF	RFC 2544: Benchmarking Methodology for Network Interconnect Devices

9 Acknowledgements

Editor: Eric Leser, Charter Communications

Principal Authors:

- Eric Leser, Charter Communications
- Bruno Giguere, MEF
- Isabelle Morency, Iometrix

Contributors:

- Mario Bianchetti, Telecom Italia
- Irene Chang, Zayo Managed Services
- Thierno Diallo, EXFO
- Yves Gagnon, EXFO
- Bruno Giguere, MEF
- Isabelle Morency, Iometrix