



## **Technical Specification**

### **MEF 47**

# **Carrier Ethernet Services for Cloud Implementation Agreement**

**October, 2014**

## Disclaimer

The information in this publication is freely available for reproduction and use by any recipient and is believed to be accurate as of its publication date. Such information is subject to change without notice and the Metro Ethernet Forum (MEF) is not responsible for any errors. The MEF does not assume responsibility to update or correct any information in this publication. No representation or warranty, expressed or implied, is made by the MEF concerning the completeness, accuracy, or applicability of any information contained herein and no liability of any kind shall be assumed by the MEF as a result of reliance upon such information.

The information contained herein is intended to be used without modification by the recipient or user of this document. The MEF is not responsible or liable for any modifications to this document made by any other party.

The receipt or any use of this document or its contents does not in any way create, by implication or otherwise:

- a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any MEF member company which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- b) any warranty or representation that any MEF member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- c) any form of relationship between any MEF member companies and the recipient or user of this document.

Implementation or use of specific Metro Ethernet standards or recommendations and MEF specifications will be voluntary, and no company shall be obliged to implement them by virtue of participation in the Metro Ethernet Forum. The MEF is a non-profit international organization accelerating industry cooperation on Metro Ethernet technology. The MEF does not, expressly or otherwise, endorse or promote any specific products or services.

© The Metro Ethernet Forum 2014. All Rights Reserved.

## Table of Contents

1.	List of Contributing Members .....	1
2.	Abstract .....	1
3.	Terminology.....	1
4.	Introduction.....	4
5.	Scope.....	4
6.	Compliance Levels.....	4
7.	Numerical Prefix Conventions.....	5
8.	Carrier Ethernet for Cloud Service Model.....	5
8.1	Cloud Services Use Cases .....	7
8.1.1	Data Center Interconnect Use Cases .....	7
8.1.2	Data Center Access Use Cases.....	8
8.2	Applying MEF Service Definitions to Cloud Services .....	8
8.2.1	Ethernet Private Line Service.....	9
8.2.2	Ethernet Virtual Private Line Service .....	10
8.2.3	Ethernet Access Services .....	11
9.	Elastic Ethernet Services.....	11
9.1	Elastic Service Attributes .....	12
9.2	Additional Attributes and Parameters for Elastic Services .....	14
9.2.1	CIR and EIR Limiting Attributes.....	14
9.2.2	Scheduled Modification Parameters.....	15
9.2.3	Modification Frequency Limiting Attributes .....	16
9.3	Elastic Service Management Interface Performance Metrics.....	17
9.3.1	Total Modification Requests .....	17
9.3.2	Total Valid Requests .....	17
9.3.3	Total Accepted Requests.....	18
9.3.4	Total Fulfilled Requests .....	18
9.3.5	Components of an Elastic SLS.....	18
9.4	Elastic Service Modification Records .....	18
9.5	Elastic Service Management Interface .....	19
10.	UNI Requirements .....	19
11.	EVC Requirements .....	20
12.	References.....	22
Appendix A.	Carrier Ethernet Service for Cloud Examples.....	24
A.1	Virtual Machine Migration.....	24
A.2	Storage Replication.....	25
A.3	Cloud Bursting.....	25
Appendix B.	Carrier Ethernet Service OAM.....	27

**List of Figures**

Figure 1 – General Types of Cloud Use Cases .....	6
Figure 2 – Cloud Service Roles and Relationships.....	6
Figure 3 – Cloud Service Management Architecture .....	7
Figure 4 – DCI Use Case .....	8
Figure 5 – CP to CC Use Case (Point-to-Point) .....	8
Figure 6 – Example of Ethernet Private Line (EPL) Service.....	9
Figure 7 – Example of Ethernet Virtual Private Line (EVPL) Services.....	10
Figure 8 – Example including an Ethernet Access Service .....	11
Figure 9 – FM and PM Reference Model for DCI Use Case.....	27
Figure 10 – FM and PM Reference Model for DCA Use Case .....	27

**List of Tables**

Table 1 – Terminology and Definitions Table.....	4
Table 2 – Numerical Prefix Conventions.....	5
Table 3 – Elastic Service Attributes.....	12
Table 4 – Examples of Carrier Ethernet for Cloud Traffic Classes mapping to CoS Labels .....	21
Table 5 – One-way CPO examples for Point-to-Point Carrier Ethernet Service for Cloud .....	24

## 1. List of Contributing Members

The following members of the MEF participated in the development of this document and have requested to be included in this list.

Albis Technologies	Huawei Technologies Co., Ltd
Alcatel-Lucent	KDDI Corporation
Allstream	Omnitron Systems Technology, Inc.
Cable Television Labs	PLDT Corp. Business Solutions
China Telecom	RAD Data Communications
Ciena Corporation	Tata Communications
Cisco Systems	Tech 2000
Comcast	The Carrier Ethernet Academy

## 2. Abstract

This document identifies the requirements for MEF Ethernet Services and MEF External Interfaces (EIs such as UNIs) as well as a management interface for use in support of Cloud Services. This support includes elastic behavior of Ethernet Service attributes that can be modified during the lifetime of the service. Support for Cloud Services falls into two broad categories: 1) interconnection of a Cloud Provider's data centers (referred to as Data Center Interconnect – DCI), and 2) interconnection of Cloud Consumers (e.g. enterprises) and Cloud Provider data centers (referred to as Data Center Access – DCA). The services and requirements in this Implementation Agreement are based on the services defined in MEF 6.2 [5] and the attributes defined in MEF 10.3 [6] and this IA. Support of Cloud Services is addressed for a single Cloud Provider (CP) using one or more Carrier Ethernet Networks (CENs) and point-to-point Ethernet Services.

## 3. Terminology

This section defines the terms used in this document. In many cases, the normative definitions to terms are found in other documents. In these cases, the third column is used to provide the reference that is controlling in other MEF or external documents.

Term	Definition	Reference
Access Provider	A CEN Operator that offers the Ethernet Access Service type.	MEF 33 [14]
Carrier Ethernet for Cloud service	Carrier Ethernet Service tailored for use in Cloud Computing applications.	This IA
CB	Cloud Broker	This IA, based on NIST [3]
CBS	Committed Burst Size	MEF 10.3 [6]

CC	Cloud Consumer	This IA, based on NIST [3]
CEN	Carrier Ethernet Network	MEF 12.2 [7]
CIR	Committed Information Rate	MEF 10.3 [6]
CIR UB	CIR Upper Bound	This IA
{CIR+EIR} LB	Lower Bound for the sum CIR plus EIR	This IA
{CIR+EIR} UB	Upper Bound for the sum CIR plus EIR	This IA
Cloud Broker	An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between Cloud Providers, Cloud Consumers and Ethernet Cloud Carriers.	This IA, based on NIST [3]
Cloud Consumer	A person or organization that maintains a business relationship with, and uses service from, Cloud Providers. Also known as Subscriber (MEF 10.3 [6]).	This IA, based on NIST [3]
Cloud Service	A service provided to a Cloud Consumer using a shared pool of configurable resources (e.g., networks, servers, storage, applications, and services).	This IA, based on NIST [2]
Cloud Provider <sup>1</sup>	A person, organization or entity responsible for making cloud services available to Cloud Consumers.	This IA, based on NIST [3]
CoS	Class of Service	MEF 10.3 [6]
CoS Label	Class of Service Label	MEF 23.1 [11]
CPO	CoS Performance Objective	MEF 23.1 [11]
CP	Cloud Provider	This IA, based on NIST [3]
DCA	data center access (a use case addressed in this IA)	This IA
DCI	data center interconnect (a use case addressed in this IA)	This IA
Duration	A service modification request parameter indicating the amount of time a service attribute modification is to be in effect before reverting to its previous value.	This IA
EBS	Excess Burst Size	MEF 10.3 [6]
ECC	Ethernet Cloud Carrier	This IA
EIR	Excess Information Rate	MEF 10.3 [6]
EIR UB	EIR Upper Bound	This IA
Elastic	An adjective used to indicate the capability to modify an active service (e.g., “elastic service”) by changing the value of one or more service attributes (e.g., “elastic service attribute”).	This IA

<sup>1</sup> Cloud Provider (CP) and Cloud Service Provider (CSP) are terms used interchangeably in the industry. This document uses Cloud Provider which is formally defined in NIST [3].

Ethernet Cloud Carrier	A CEN operator that provides Carrier Ethernet connectivity and transport of cloud services between Cloud Providers and Cloud Consumers and with other Ethernet Cloud Carriers.	This IA, based on NIST [3]
MI	Management Interface	This IA
Modification Frequency Limit	An attribute of an elastic service constraining the number of modification requests that can be made during a given time period.	This IA
Modification Interval Limit	An attribute of an elastic service specifying a minimum time between attribute modification requests.	This IA
NIST	National Institute of Standards and Technology	<a href="http://www.nist.gov/">http://www.nist.gov/</a>
Period	A service modification request parameter indicating an interval at which an elastic service modification request is scheduled to repeat.	This IA
PT	Performance Tier	MEF 23.1 [11]
Management Interface	An interface by which a Cloud Broker can request service attribute modifications from an ECC and be informed about the state of the service and the status of any requested attribute modifications.	This IA
Service Provider	The organization providing Ethernet Service(s). In this document the Service Provider is the ECC.	MEF 10.3 [6]
SLS	Service Level Specification	MEF 10.3 [6]
SP	Service Provider	MEF 10.3 [6]
Start Interval	A service modification request parameter indicating the time interval after the Start Time within which a requested service attribute modification can be made.	This IA
Start Time	A service modification request parameter indicating the time at which a requested service attribute modification can begin.	This IA
Subscriber	The organization purchasing and/or using Ethernet Services. In this document the Subscriber is the CC or CP.	MEF 10.3 [6]
Total Accepted Requests	The total number of modification requests accepted for an elastic service instance during a measurement interval.	This IA
Total Fulfilled Requests	The total number of modification requests fulfilled for an elastic service instance during a measurement interval.	This IA
Total Modification Requests	The total number of modification requests received for an elastic service instance during a measurement interval.	This IA.

Total Valid Requests	The total number of valid modification requests received for an elastic service instance during a measurement interval.	This IA
VM	virtual machine	NIST [3]

**Table 1 – Terminology and Definitions Table**

## 4. Introduction

The rapid growth of cloud services has created a market need for reliable and elastic connectivity between Cloud Provider (CP) data centers and between these CP data centers and Cloud Consumer (CC) locations. Initially this connectivity has been provided over the public Internet; however requirements for improved security, predictable and guaranteed performance, and control of data governance and regulatory compliance are difficult or impossible to realize over the public Internet. Carrier Ethernet Services provide a high quality alternative to the Internet for cloud service interconnection enabling strict control of access and conforming to a service level specification (SLS). This Carrier Ethernet for Cloud services implementation agreement addresses these requirements for enhanced connectivity.

## 5. Scope

This implementation agreement describes the applicability of MEF Ethernet Services to Cloud use cases, including CP data center interconnection and CC access to CP data centers. In particular it addresses the application of MEF 6.2 [5] EPL and EVPL services to these use cases and the associated service attribute requirements. CoS performance objectives (CPOs) are described that can support a range of cloud operations and use cases.

In addition, this IA specifies a set of elastic service attributes that can be adjusted during the lifetime of the service to meet the varying demands of cloud applications. This IA defines new service attributes for the management interface that is used to control elastic service attributes and report service status to the CC or CP (via the Cloud Broker).

## 6. Compliance Levels

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1]. All key words must be in upper case, bold text.

Items that are **REQUIRED** (contain the words **MUST** or **MUST NOT**) are labeled as [Rx] for required. Items that are **RECOMMENDED** (contain the words **SHOULD** or **SHOULD NOT**) are labeled as [Dx] for desirable. Items that are **OPTIONAL** (contain the words **MAY** or **OPTIONAL**) are labeled as [Ox] for optional.

A paragraph preceded by [CRa]< specifies a conditional mandatory requirement that **MUST** be followed if the condition(s) following the "<" have been met. For example, "[CR1]<[D38]" indicates that Conditional Mandatory Requirement 1 must be followed if Desirable Requirement 38 has been met. A paragraph preceded by [Cdb]< specifies a Conditional Desirable Requirement that **SHOULD** be followed if the condition(s) following the "<" have been met. A

paragraph preceded by [COc]< specifies a Conditional Optional Requirement that **MAY** be followed if the condition(s) following the “<” have been met.

## 7. Numerical Prefix Conventions

This document uses the prefix notation to indicate multiplier values as shown in Table 2.

Decimal		Binary	
Symbol	Value	Symbol	Value
<b>k</b>	$10^3$	<b>Ki</b>	$2^{10}$
<b>M</b>	$10^6$	<b>Mi</b>	$2^{20}$
<b>G</b>	$10^9$	<b>Gi</b>	$2^{30}$
<b>T</b>	$10^{12}$	<b>Ti</b>	$2^{40}$
<b>P</b>	$10^{15}$	<b>Pi</b>	$2^{50}$
<b>E</b>	$10^{18}$	<b>Ei</b>	$2^{60}$
<b>Z</b>	$10^{21}$	<b>Zi</b>	$2^{70}$
<b>Y</b>	$10^{24}$	<b>Yi</b>	$2^{80}$

Table 2 – Numerical Prefix Conventions

## 8. Carrier Ethernet for Cloud Service Model

Cloud services as defined by NIST ([2][3]) involve a number of parties: a CC uses cloud services purchased from a CP and interacts with these cloud services using a network connectivity service provided by a Cloud Carrier. This document describes the use of Carrier Ethernet Services in support of cloud services and therefore the Cloud Carrier is called an Ethernet Cloud Carrier (ECC). A Cloud Broker (CB) is responsible for negotiating service details with each party and coordinating the overall delivery of a cloud service.

For clarity each of these parties is treated as though it is a distinct entity; however, these parties can be considered roles and that one entity or organization can play multiple roles. For example a CP can also play the CB role, or one entity can be the ECC, CP, and CB.

There are two types of use cases covered by Carrier Ethernet for Cloud as shown in Figure 1. The first addresses the need for elastic MEF service connectivity between CP data centers. The second addresses providing elastic MEF service connectivity between CC locations and CP data centers.

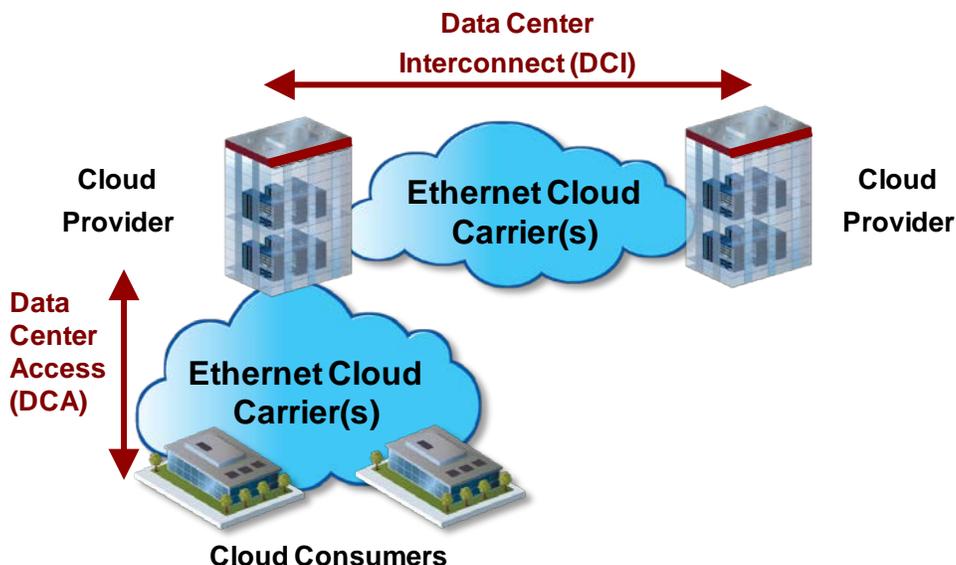


Figure 1 – General Types of Cloud Use Cases

The modification of Carrier Ethernet Service attributes to meet the needs of CPs and CCs is orchestrated by a CB responsible for managing the use, performance and delivery of cloud services. The CB negotiates relationships between CP, CC and ECC (see Figure 2).

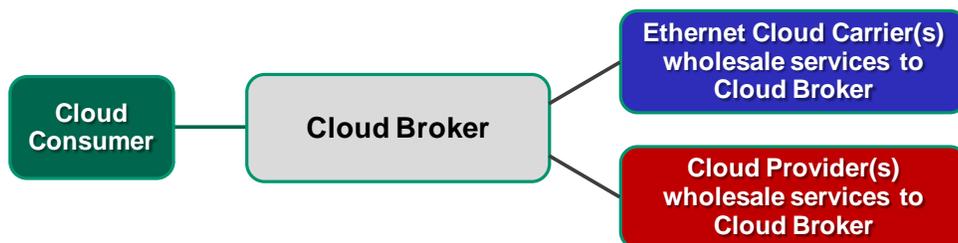
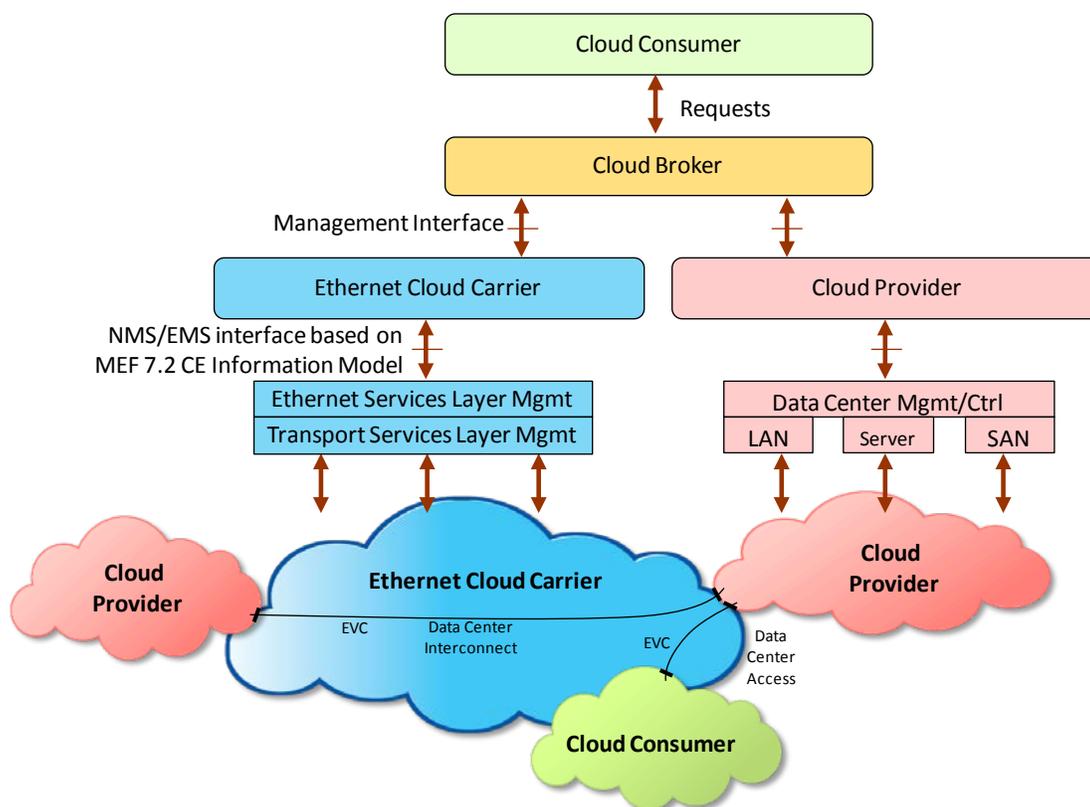


Figure 2 – Cloud Service Roles and Relationships

The CB role can be played by the CC, CP, ECC, or a third party (i.e., an independent broker). The responsibilities of the CB are essentially the same regardless of which entity is performing the role.

After a Carrier Ethernet for Cloud service is established the CB can modify elastic service attributes from time to time. The CB (acting for the CC or the CP) requests modification of elastic service attributes via a management interface provided by the ECC. The ECC (as SP) is responsible for coordinating attribute modification across one or more CENs. The CB interacts only with the ECC providing the Carrier Ethernet for Cloud service.



**Figure 3 – Cloud Service Management Architecture**

The management interface<sup>2</sup> used to modify elastic service attributes is shown in Figure 3. The CB can use this interface to request service attribute modifications, either immediately or at a specified time in the future. The ECC makes the requested modifications in the CEN Ethernet Services Layer and Transport Services Layer (MEF 4 [4]) as needed. The ECC responds with request fulfillment status reports. Service performance information is also communicated via the management interface. This can include fault notification as well as performance and usage statistics gathered by the ECC.

## 8.1 Cloud Services Use Cases

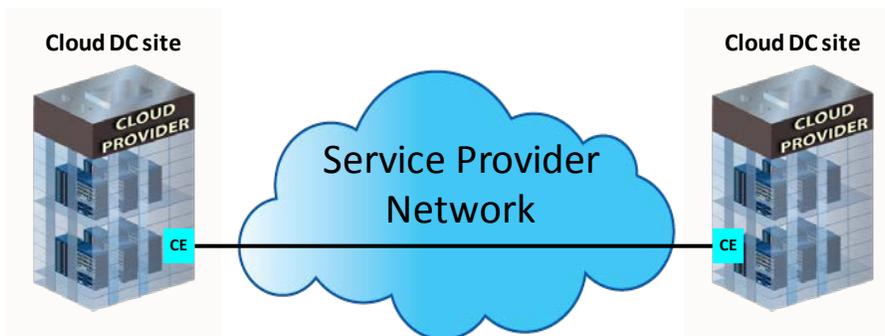
Wide Area Network (WAN) connectivity employed in support of Cloud services falls into two broad categories: connectivity used to interconnect the data centers of CPs and connectivity used to connect CCs to CP data centers. This section describes these use cases in general terms and section 8.2 describes the application of Carrier Ethernet services to these use cases in more detail.

### 8.1.1 Data Center Interconnect Use Cases

CP data centers are interconnected to support both normal levels of application traffic and occasional surges in traffic due to database transfers. The normal traffic can include application

<sup>2</sup> The management interface described in this IA may be part of a larger service management interface provided by the ECC to its Subscribers; however, this IA is focused only on those aspects of the management interface that relate to elastic service modification.

interactions between data centers and continuous synchronization of distributed/redundant databases. The occasional surges can include moving virtual machines from one data center to another or copying large datasets for backup. These occasional surges create temporary demand for increased bandwidth in the service interconnecting the data centers so as to enable faster completion of the transfer.

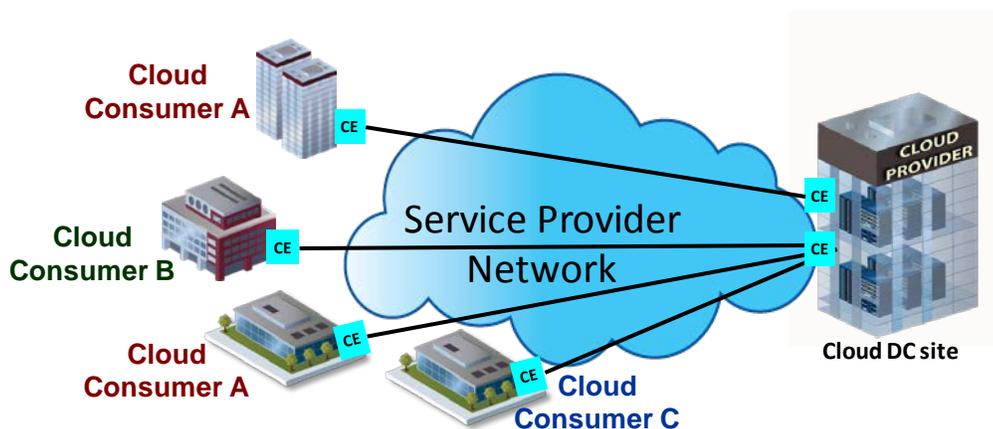


**Figure 4 – DCI Use Case**

This IA specifies the use of MEF E-Line services for this use case. Future work might include other MEF service types. In this use case the Subscriber is the CP operating the DC sites.

### 8.1.2 Data Center Access Use Cases

In the CP to CC use case services are used to connect one or more sites of a given CC as well as one or more CCs to a CP data center. Multiple CC sites are connected to the CP data center, as shown in Figure 5.



**Figure 5 – CP to CC Use Case (Point-to-Point)**

This IA specifies the use of MEF E-Line services for this use case. Future work might include other MEF service types. In this use case the Subscriber at one UNI in the EVC is the CP while the Subscriber at another UNI in the EVC is the CC.

## 8.2 Applying MEF Service Definitions to Cloud Services

This section specifies Carrier Ethernet for Cloud services. In addition to the baseline definition of MEF Services in MEF 6.2 [5], using service attributes defined in MEF 10.3 [6], this IA has

specified requirements for elastic services using additional attributes defined in this IA (see Section 9).

A CP can use VLAN-based services (EVPL) to provide connectivity to multiple CC locations with each cloud data center UNI. Using a single service multiplexed UNI at the data center the CP can provide one or more EVCs to a CC as well as connect to several CCs providing efficient scaling. VLAN-based services allow bandwidth profiles to be tailored to the specific needs of each CC.

- [R1] A Carrier Ethernet for Cloud service between MEF compliant UNIs **MUST** comply with the EVPL (Ethernet Virtual Private Line Service) (MEF 6.2 [5]) in terms of the service attributes for UNI and EVC, complemented by those specified in Section 9.5 and Section 11 in this IA.

While VLAN-based services offer efficient scaling, there are cases in which a cloud data center UNI can be dedicated to a single cloud consumer. Furthermore, DCI use cases can require dedicated bandwidth between cloud data centers. To address these cases a port-based EPL service is required.

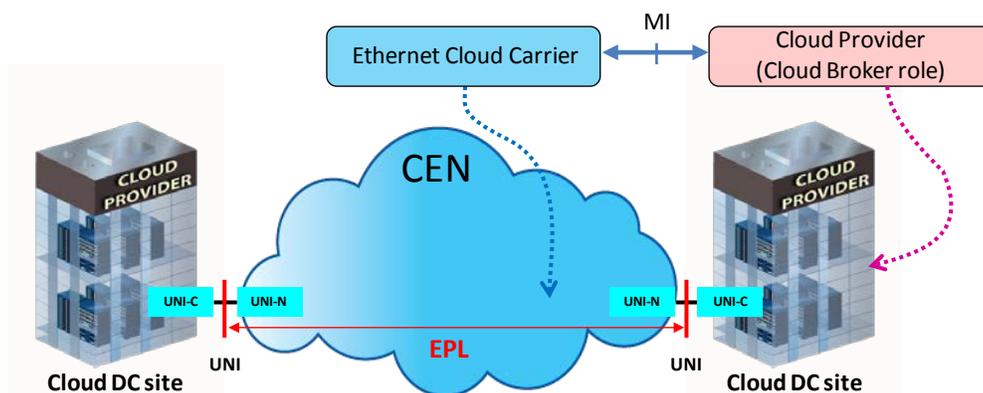
- [R2] A Carrier Ethernet for Cloud service between MEF compliant UNIs **MUST** comply with the EPL (Ethernet Private Line Service) (MEF 6.2 [5]) in terms of the service attributes for UNI and EVC, complemented by those specified in Section 9.5 and Section 11 in this IA.

See Section 9.5 for the UNI Service Attributes and Section 11 for EVC Service Attributes from MEF 6.2 [5] as well as constraints, if any, as defined in this IA.

If a site involved in a Carrier Ethernet for Cloud service is reached via another Carrier Ethernet service provider, an Ethernet Access service as specified in MEF 33 [14] can be used. The scope of available Carrier Ethernet for Cloud service can be expanded if Access Providers support elastic behavior for Ethernet Access services.

### 8.2.1 Ethernet Private Line Service

In the data center interconnect (DCI) use case MEF 6.2 services can be used to provide dedicated bandwidth between CP data center sites. The DCI use case consists of a point-to-point EVC between data centers. In Figure 6 below, an EPL service is provided between two cloud data centers, each supporting an MEF compliant UNI-C Ethernet interface.



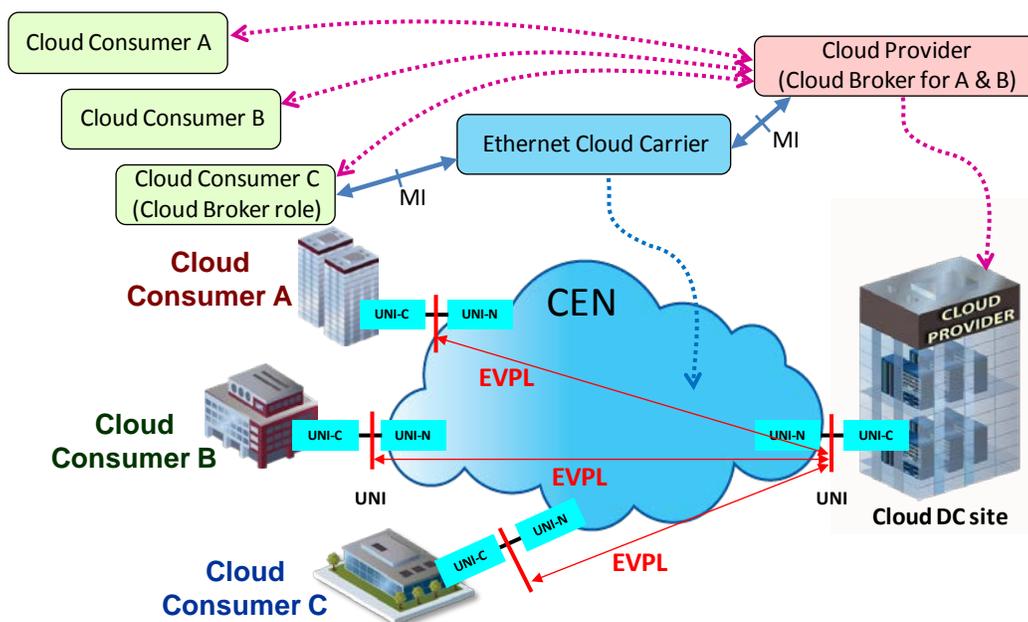
**Figure 6 – Example of Ethernet Private Line (EPL) Service**

To meet DCI needs the EPL service can provide elastic service behavior (e.g., the ability to change CIR) and flexible CoS mapping to allow the CP to adjust the service to meet varying traffic requirements between the Cloud DC sites.

A CB (shown as the CP in Figure 6) can be responsible for coordinating the management of CEN EVCs and cloud services. The elastic attributes of the EPL service are controlled via the management interface (interface labeled "MI" in Figure 6).

### 8.2.2 Ethernet Virtual Private Line Service

In Figure 7 below, EVPL services provide connectivity from one MEF compliant UNI-C Ethernet interface at the CP data center to multiple CC sites. Each EVPL connection has its own service attributes selected to meet that CP's and CC's cloud service requirements.



**Figure 7 – Example of Ethernet Virtual Private Line (EVPL) Services**

Each CC has access to elastic service behavior (e.g., the ability to modify CIR) and flexible CoS mapping enabling them to adjust their EVPL service attributes to meet their specific cloud service access needs, within the bounds specified in their service agreement.

A CB<sup>3</sup> (shown as either the CC or CP in Figure 7) is responsible for coordinating the management of CEN EVCs and cloud services. The elastic attributes of each EVPL service are controlled via a management interface. Some attribute modifications need to be coordinated among the participants in the Cloud service. For example, increasing the bandwidth of an EVC requires checking with the CP to verify that there is sufficient available bandwidth at their respective UNIs to support the requested increase and verifying that the increase is within the bounds of the service agreement. Similarly, adding a CE-VLAN ID to an EVPL might require selecting an identifier value that is available at both UNIs. The CB is responsible for this

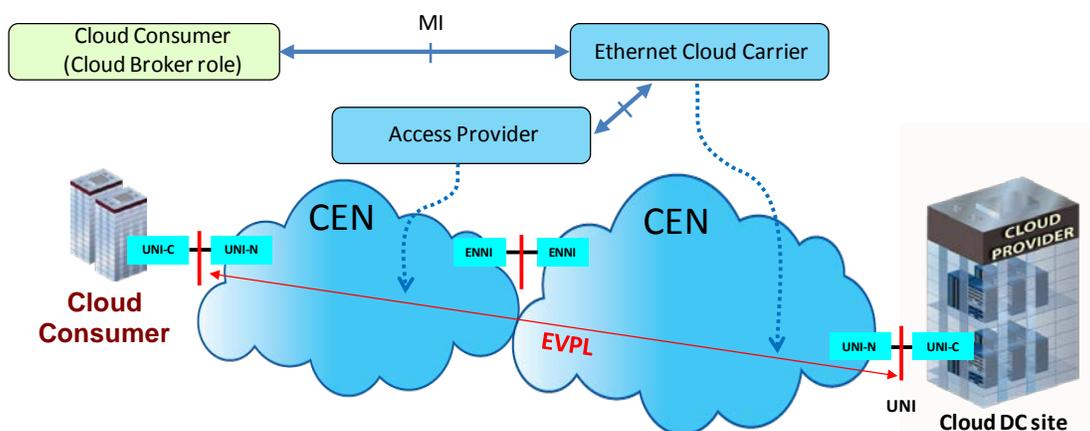
<sup>3</sup> The example in the figure shows three services with one of the CCs (CC 'C') playing the CB role for his own EVPL and the CP playing the CB role for the other two EVPLs (belonging to CC 'A' and CC 'B'). In general, various entities may play the CB role including the CC, the CP, a third party, or the ECC.

coordination among the participants, when needed, before making a service attribute modification request.

### 8.2.3 Ethernet Access Services

The ECC, shown in Figure 6 and Figure 7, can use one or more CENs to deliver the Carrier Ethernet for Cloud Services defined in this Implementation Agreement. An ENNI (MEF 26.1 [12]) is used to extend Ethernet services across multiple CENs when CCs' sites are not all served by a single CEN. Furthermore, as example, one of the CENs could be operated by the ECC responsible to deliver the Ethernet Service. In some cases a transit network providing ENNI-to-ENNI services can also be used to further extend the reach of a CP.

In Figure 8 below, an E-Access service (e.g., either A-EPL or A-EVPL) is used via a second CEN to support the EVPL service to a CC site. The ECC is responsible for coordinating modifications of UNI and EVC attributes in both CENs. In this scenario both CENs support elastic service; however, the interface used by the ECC to request E-Access Service modifications is outside the scope of this IA.



**Figure 8 – Example including an Ethernet Access Service**

The ECC can make modifications to the E-Access Service at any time. For example, if a modification to the EVC is scheduled to increase CIR at a given start time the ECC can modify the E-Access Service in advance to prepare for the scheduled EVC modification. Similarly, if a service modification request is received to reduce the CIR of the EVC the ECC can modify its OVC immediately and modify the E-Access Service to reduce CIR for that service at a later time.

## 9. Elastic Ethernet Services

Elastic Carrier Ethernet Services allow modification of selected EVC or UNI attributes. The CB can request modification of these service attributes to meet changing demands of the applications supported by an EPL or EVPL. If the requested service attribute modifications can be made then they are made during a short maintenance interval. The new service attribute values remain in effect for a specified duration or until the CB requests another modification. Service attribute modifications can be made on-demand (i.e., fulfillment requested immediately) or scheduled

(i.e., fulfillment requested at a designated time in the future). A scheduled modification request with a fixed duration can be periodic (i.e., recurring at a specified interval such as weekly or monthly).

Service attribute modification is requested via a management interface provided by the ECC to the CB (i.e., any party playing the CB role). Each modification request is validated to ensure it is well formed and conforms to the constraints associated with the service. If a request is valid the ECC determines whether or not the request can be fulfilled given the current state of the CEN. If the CEN can support the request the necessary configuration changes are made to fulfill the request and the CB is notified that the request has been completed via the management interface. If the request cannot be completed for any reason the CB is notified by an error message.

It is up to the ECC to determine the dividing line between requests that will be treated as immediate and those that will be scheduled. The ECC can schedule maintenance on the management systems responsible for handling service modification requests and not accept either immediate or scheduled requests that would be executed during the maintenance period.

## 9.1 Elastic Service Attributes

The service attributes that can be modified are listed in Table 3 along with any limits on the range of modification (degree of elasticity). Service traffic flow can continue while modifying service attributes. Since some traffic impact might be unavoidable as network functions are reconfigured, service attributes are modified in a maintenance interval (MEF 10.3 [6]) during which performance monitoring (PM) is not included in the determination of the SLS. Thus potential traffic impact caused by network reconfiguration does not affect the SLS.

Service Attribute	Elasticity
UNI Service Attributes	
CE-VLAN ID/EVC Map	Add/remove CE-VLAN ID when EVPL
EVC per UNI Service Attributes	
Class of Service Identifier	Modify the set of CoS ID values that map to a given CoS Name
Ingress BWP per CoS ID	
CIR	Adjust CIR within bounds
EIR	Adjust EIR within bounds

**Table 3 – Elastic Service Attributes**

CCs can increase or decrease the number of VLANs connected to a CP data center, e.g., depending on which of their VLANs attach to cloud services at a given time. The CB can change the CE-VLAN ID/EVC Map by adding/removing a CE-VLAN ID to/from an EVPL at both UNIs. Since predicting future needs can be difficult, providing an elastic mapping of CE-VLAN ID to EVC is preferable to a fixed assignment. CE-VLAN ID translation is not provided by the ECC if more than one CE-VLAN is mapped to the EVC so each party is responsible for managing the CE-VLAN ID space at their UNI. The CB has to ensure that the CE-VLAN ID assignment is understood and accepted by both parties (CC and CP) on the EVPL.

**[D1]** If CE-VLAN ID Preservation is Enabled for an elastic service an elastic EVPL service **SHOULD** allow changing the CE-VLAN ID/EVC Map by adding or

removing one or more CE-VLAN IDs from the set mapped to the EVPL at both UNIs.

- [D2] If CE-VLAN ID Preservation is Disabled an elastic EVPL service **SHOULD** allow changing the CE-VLAN ID value mapped to the EVPL (but only one CE-VLAN ID can be mapped to the EVPL).

Different Cloud applications often require different CoS, for example, file transfer can use CoS Label L while video streaming might need CoS Label M and VM migration might require CoS Label H. An ECC can offer multiple CoS Names to enable the CC to use a variety of cloud applications.

- [D3] An elastic service **SHOULD** support more than 1 CoS Name.

For E-Line services the CB can modify the CoS used by the service by modifying the Class of Service Identifier service attribute for Data Service Frames (MEF 10.3 [6]). The set of CoS Names, the method of CoS identification (e.g., based on EVC, EVC+PCP, or EVC+IP) and the CoS Frame Sets <S, CoS Name, PT> (MEF 23.1 [11]) associated with the service are not elastic.

- [CD1]<[D3] An elastic service supporting multiple CoS Names **SHOULD** offer selectable performance objectives by allowing modification of the mapping of CoS ID value to CoS Name.

Mapping a CoS ID value to a CoS Name that was previously unmapped (i.e., previously no CoS ID value was mapped to the CoS Name) provides access to an additional CoS Frame Set with associated CPOs for the service. Removing the last mapping of a CoS ID value to a given CoS Name eliminates access to a CoS Frame set and its associated CPOs for the service. Note that while changing the CoS ID value to CoS Name mapping can cause changes in the amount of traffic assigned to a particular CoS Name this does not affect OAM configuration. For example, if PM is enabled for a given CoS Name this function is unaltered by changes in the traffic assigned to the CoS Name.

An elastic service is required to support modification of the service's Ingress Bandwidth Profile. MEF 6.2 [5] allows the Ingress Bandwidth Profile per CoS ID attribute value to be *No* or *Parameters* but the Ingress Bandwidth Profile can only be modified if it has specified parameters.

- [R3] An elastic service **MUST** specify Ingress Bandwidth Profile per CoS ID parameters (i.e., this attribute's value cannot be *No*).

The CB can change the Bandwidth Profile parameters CIR and EIR (within limits specified with the service) by requesting a new value associated with a particular CoS Name. The valid values that can be requested are within the bounds specified by the service ({CIR+EIR} Lower Bound, CIR Upper Bound, and {CIR+EIR} Upper Bound).

- [R4] An elastic service **MUST** allow setting the CIR for a CoS Name to a value within the range defined by the {CIR+EIR} Lower Bound and CIR Upper Bound defined for the service.

The ECC can place a constraint on CIR modification to limit the rate at which the value can increase, for example limiting each request to no more than a defined multiple of the current CIR value that can be less than the CIR Upper Bound. The ECC can also require that CIR values be expressed in particular increments, for example as specified in requirement 13 in MEF 13 [8],

and can reject a CIR request or round it up to the next higher valid value if the requested value does not conform to the increment rules.

- [R5] An elastic service **MUST** allow setting the EIR for a CoS Name to a value within the range defined by the {CIR+EIR} Lower Bound and {CIR+EIR} Upper Bound defined for the service.

CBS and EBS are determined by the ECC, based on their own algorithm, to support the agreed CIR and EIR and therefore these are not elastic attributes settable by the CB. However, the CBS and EBS attributes could be provided via the management interface. The algorithmic relationship between CIR and CBS and between EIR and EBS can also be provided to the subscriber as part of the elastic service description. Knowing these burst size attributes can be useful in determining the configuration of the subscriber's traffic shaping function. If necessary, the ECC can adjust the requested values for CIR and EIR to be consistent with these requirements.

The Egress Bandwidth Profile per Egress Equivalence Class Identifier and Egress Equivalence Class Identifier service attributes are not under the direct control of the CB. Therefore, if these service attributes are used, the ECC is responsible for making sure these attributes have values that are compatible with the Ingress BWP per CoS ID service attribute.

- [R6] If the Egress Equivalence Class Identifier and Egress Bandwidth Profile per Egress Equivalence Class Identifier service attributes are used, the ECC **MUST** adjust these attributes as necessary to maintain compatibility with the parameter values (e.g., CIR and CBS values) of the Ingress BWP per CoS ID at the opposite UNI.

## 9.2 Additional Attributes and Parameters for Elastic Services

The management interface for elastic services provides the means to modify selected service attributes. Limits can be placed on the modification of service attributes and these limits can be specified as a part of the Carrier Ethernet Service. These attributes include:

- CIR limiting attributes: CIR Upper Bound and {CIR+EIR} Lower Bound,
- EIR limiting attributes: {CIR+EIR} Lower Bound and {CIR+EIR} Upper Bound,
- Modification Frequency Limit, and
- Modification Interval Limit.

There are also parameters specific to a particular modification request. These parameters include:

- Start Time,
- Start Interval,
- Duration, and
- Period.

These attributes and parameters related to elastic services are described in the following sections.

### 9.2.1 CIR and EIR Limiting Attributes

The modification of CIR and EIR, for a given CoS Name, is constrained by three limits, an upper bound on CIR and an upper bound and a lower bound on the sum of CIR plus EIR. The CIR

Upper Bound Attribute (CIR UB) defines the maximum valid value that can be requested for CIR.

- [R7] An elastic service **MUST** specify a CIR Upper Bound in bits per second as the maximum valid value for CIR of a given CoS Name that can be requested for the Ingress Bandwidth Profile per CoS ID (MEF 10.3 [6]).

The {CIR+EIR} Upper Bound Attribute ({CIR+EIR} UB) constrains the maximum valid value that can be requested for EIR given the current or requested CIR value.

- [R8] An elastic service **MUST** specify a {CIR+EIR} Upper Bound in bits per second, constraining the maximum valid value of EIR of a given CoS Name that can be requested for a given Ingress Bandwidth Profile (MEF 10.3 [6]).

The {CIR+EIR} Lower Bound Attribute defines the minimum valid value that can be requested for the sum of CIR plus EIR.

- [R9] An elastic service **MUST** specify a {CIR+EIR} Lower Bound in bits per second constraining the minimum values for CIR and EIR of a given CoS Name that can be requested for the Ingress Bandwidth Profile per CoS ID (MEF 10.3 [6]).

## 9.2.2 Scheduled Modification Parameters

A service attribute modification can be scheduled to occur at a set time in the future. A scheduled modification can include a request to make a modification at a specified time or to make a modification immediately for a specified duration, implying a modification to revert to the previous attribute values when the duration ends. Three modification request parameters control the scheduling of service attribute modifications: Start Time, Start Interval, and Duration.

### 9.2.2.1 Start Time parameter

An elastic attribute modification request can include a Start Time parameter to indicate the time at which a requested attribute modification can be made.

- [R10] An elastic service modification request **MUST** support a Start Time parameter, specified to seconds in UTC, indicating the time at which the requested attribute modifications can begin.

The ECC can put a limit on the Start Time, for example, limiting scheduled requests to have a Start Time within six months or within one year of the current time. In addition, the ECC can treat a scheduled request as an immediate request if the Start Time is within a specified interval from the current time, e.g., if the Start Time is within an hour from the current time. The ECC can determine a policy for handling anomalous requests (e.g. a Start Time in the past), for example treating the request as an immediate request or rejecting the request.

### 9.2.2.2 Start Interval parameter

The scheduling of service attribute value modifications and the fulfillment of these requests is handled by the ECC, and coordinated as necessary with access service providers. The fulfillment time required to complete the requested attribute modifications can vary due to operational constraints. Furthermore, subscribers can have varying requirements for how soon after the Start Time the modified service attributes are needed. To allow for flexibility in network operations

and subscriber requirements a time interval within which a scheduled attribute modification is expected to be completed can be specified using a Start Interval parameter.

- [R11] An elastic service modification request **MUST** support a Start Interval parameter to indicate the acceptable interval after the Start Time during which the service attribute modifications can be made.

Providing the ability to specify an acceptable interval in which the attribute modifications can be made allows the CB to indicate their requirements and can enable the ECC to optimize the use of their resources by scheduling modification actions within the allowed interval. A notification can be sent via the management interface to inform the CB when the attribute modifications have been completed. If the Start Interval parameter is not provided or is set to a value lower than the ECC's expected fulfillment time for the attribute modifications the requested modification is initiated at the Start Time (i.e., there is no flexibility for the SP to make the modification at a later time).

#### 9.2.2.3 Duration parameter

If the CB requires service attribute modifications for a specific time interval and wants the service attributes to revert to their previous values after that interval, a Duration parameter can be used to indicate the required time interval.

- [D4] An elastic service modification request **SHOULD** support a Duration parameter indicating the time interval for which the requested attribute modifications will remain in effect before automatically reverting to their previous values.

The Duration is measured from the time the CB is notified of successful service modification and is expected to correspond to the time interval recorded in the billing record for the attribute modification, unless the interval is cut short by subsequent events, e.g., a further modification request or a network failure. The Duration can be specified in days, hours, minutes, seconds or whatever form the ECC chooses to support. If the Duration parameter is not included in a service modification request the modified service attributes will remain in effect until changed by another modification request.

#### 9.2.2.4 Period parameter

If the CB requires a modification request with a limited duration to be scheduled on a regular basis, a Period parameter can be used to enable periodic scheduling with a single request.

- [D5] An elastic service modification request **SHOULD** support a Period parameter (e.g., daily weekly, or monthly) indicating the time interval at which the requested attribute modification request is to repeat.

### 9.2.3 Modification Frequency Limiting Attributes

To protect a network from potential disruption due to a high level of service modification activity, an ECC can set a limit on the frequency of elastic service attribute modification requests. Two kinds of limit can be used: a limit on the number of requests in a certain time period and a limit on the interval between successive requests.

- [R12] An elastic service **MUST** include a Modification Frequency Limit attribute, specified in requests per time period (e.g., per day, week, or month), that sets a limit on the general frequency of elastic service attribute modification.
- [R13] An elastic service **MUST** include a Modification Interval Limit attribute, specified as a time interval (e.g., in minutes or hours), that sets a limit on the minimum time between attribute modification requests.

There can be more requests made than actual attribute modifications since some requests might be denied leading to alternate requests or one request can override or cancel another. It is impossible to limit the number of requests made by a CB; however, an ECC can limit or cancel elastic service for a CB who transmits modification requests at an excessive frequency and thereby impacts the ability of other customers to obtain service via the management interface.

### 9.3 Elastic Service Management Interface Performance Metrics

Elastic services introduce a new service interface (the management interface) and new dynamic actions related to the service instance (modification of selected service attributes). The performance of these new service actions can be measured and reported<sup>4</sup> as a part of the elastic service.

#### 9.3.1 Total Modification Requests

The first of a series of elastic service performance measures is the total number of modification requests for a service instance. This provides a measure of the level of activity of the CB responsible for requesting service attribute modifications.

- [R14] An ECC **MUST** measure the total number of modification requests, Total Modification Requests (TMR), for a given elastic service received at the management interface over the period T used for other service performance measurements associated with the SLS.

The TMR measures all requests received, including malformed requests and requests that do not meet the validity constraints associated with the elastic service.

#### 9.3.2 Total Valid Requests

The second elastic service performance measure is the total number of valid modification requests for a service instance. This provides a measure of the adherence of the CB's requests to the constraints associated with the elastic service.

- [R15] An ECC **MUST** measure the total number of valid modification requests, Total Valid Requests (TVR), for a given elastic service received at the management interface over the period T used for other service performance measurements associated with the SLS.

The TVR measures the number of requests received that are well-formed and meet the validity constraints associated with the elastic service, including frequency limits.

---

<sup>4</sup> Often the ratios of the metrics defined in this section may be more interesting than the metrics themselves. These ratios can be derived from the raw counts, so the management interface is defined to convey the raw counts and the CB or Subscriber can calculate other interesting values from that data.

### 9.3.3 Total Accepted Requests

The third elastic service performance measure is the total number of accepted modification requests for a service instance. A valid modification request might not be accepted due to the state of the network when the request is made. For example, faults in the network or an unexpected level of demand can leave the network unable to satisfy a valid request. This provides a measure of the ability of the network to handle modification requests for the service.

- [R16]** An ECC **MUST** measure the total number of accepted modification requests, Total Accepted Requests (TAR), for a given elastic service received at the management interface over the period T used for other service performance measurements associated with the SLS.

The TAR measures the number of valid requests that the network is capable of satisfying.

### 9.3.4 Total Fulfilled Requests

The fourth elastic service performance measure is the total number of modification requests fulfilled within the specified Start Interval for a service instance. An accepted modification request might not be fulfilled due to a fault in the service attribute modification system or processes. This provides a measure of success in fulfilling modification requests for the service.

- [R17]** An ECC **MUST** measure the total number of fulfilled modification requests, Total Fulfilled Requests (TFR), for a given elastic service received at the management interface over the period T used for other service performance measurements associated with the SLS.

The TFR measures the number of requests that are successfully fulfilled.

### 9.3.5 Components of an Elastic SLS

The elastic service management interface performance metrics can be included in an SLS agreed between the CB and ECC. For example, the percent of valid requests accepted (TAR/TVR) or the percent of accepted requests fulfilled (TFR/TAR) might be useful measures of elastic service quality.

## 9.4 Elastic Service Modification Records

Service modification records can be required to support usage-based billing for elastic services. The time at which each modification is made and the modified service attribute value(s) can be included in a modification record and these records provided to the ECC's operations systems and to customers as part of a billing process. A modification record is a list of {time, attribute, value}, where time is measured to the second in UTC, for a fulfilled service attribute modification request.

- [O1]** An elastic service **MAY** record and report a modification record including the time of modification and new attribute values for each fulfilled service attribute modification request.

Modification records are not required to be reported to the CB via the management interface. Furthermore, modification records might not be needed if an elastic service is offered at a flat rate or if there is a cap on charges after a specified limit is reached.

## 9.5 Elastic Service Management Interface

The management interface accepts requests for modification of elastic service attributes and provides information related to elastic services on behalf of the ECC. Traditional service information is not required to be made available via the management interface since avenues for this information to be provided to the Subscriber already exist.

The information provided via the management interface is primarily information directly related to service attribute modification requests or otherwise related to elastic service operation. The following requirements apply to the management interface.

- [R18] The management interface **MUST** be able to accept requests for elastic service attribute modifications.
- [R19] The management interface **MUST** provide an indication of whether or not a service attribute modification request is accepted.
- [R20] The management interface **MUST** provide a notification when a service attribute modification request has been fulfilled.
- [R21] The management interface **MUST** provide the performance metrics specified in section 9.3.
- [D6] The management interface **SHOULD** provide the current values of the elastic service attributes.
- [D7] The management interface **SHOULD** provide the current values of CBS and EBS.
- [O2] The management interface **MAY** provide the current status of outstanding (scheduled) service attribute modification requests.
- [R22] The management interface **MUST** provide a means to cancel (erase) all outstanding (scheduled) service attribute modification requests.

## 10. UNI Requirements

UNI requirements for Ethernet Services are specified in MEF 6.2 [5]. This section specifies additional UNI requirements or constraints for Carrier Ethernet for Cloud Services.

Requirements specified for the UNI apply to both the CP UNI and CC UNI, unless specified otherwise.

Data center interconnect services can bundle multiple VLANs on a single EVC between a pair of data center sites.

- [R23] An elastic EVPL service **MUST** have Bundling Enabled at both UNI's.

Cloud services can involve transfer of large amounts of data, and in low error rate environments this can be made more efficient by using larger frame sizes. MEF 6.2 requires a UNI MTU of at least 1522 bytes and recommends at least 1600 bytes. MEF 20 [10] R73 includes a desirable

requirement for 2000 bytes. Other frame sizes can be useful, such as 9600 to support applications desiring jumbo frames.

- [D8] The ECC **SHOULD** support UNI Max Service Frame Size values greater than or equal to 2000.

An ECC can support UNI implementations to enable service that is resilient to some UNI failure scenarios. One example is Link Aggregation [16] for port protection or line card protection. In addition, there is the option to have multiple UNIs to the same CP site where UNIs can be on the same NE or different NEs. Typically, it is expected that the CP site might have more complex implementations than CC sites.

- [D9] The ECC **SHOULD** support (i.e., be capable of enabling) 2-Link Aggregation for UNI Resiliency as defined in MEF 10.3 [6].

Cloud application uses cases considered in this phase of the document are not expected to need a synchronization service, so the Synchronous Mode attribute can be disabled.

- [D10] A UNI at a CC or CP site **SHOULD** have Synchronous Mode set to Disabled.

## 11. EVC Requirements

EVC requirements for E-Line Services are specified in MEF 6.2 [5]. In general, Cloud services place no additional constraints on EVC attributes as specified in MEF 6.2 [5]. This section focuses on CoS performance objectives and recommends an approach to defining the SLS for a Carrier Ethernet Service supporting Cloud applications.

DCI services often require more stringent CoS performance objectives than DCA services. For example, DCI services might need to support live migration of VMs, real-time database synchronization, or storage networks, all of which require low delay and frame loss guarantees. DCA services can include secure and reliable access to cloud-based applications, but these seldom include stringent performance requirements. However, if a DCA service is used to connect an enterprise data center to a CP data center, it is possible that the SLS will look more like that of a typical DCI service. In any case, the CPOs defined in MEF 23.1 [11], or later revisions, are expected to be sufficient for Cloud applications.

There are three standard MEF CoS Labels defined in MEF 23.1 [11] – ‘H’, ‘M’ and ‘L’. It is expected that an ECC offering multi-CoS Carrier Ethernet for Cloud service would support all three. An ECC can also define their own CoS Names with associated CPOs and Performance Tiers. In some cases a two-CoS model (e.g., H and L or H and M) or a single CoS (e.g., H only) might be sufficient for a Carrier Ethernet for Cloud service.

- [D11] A Carrier Ethernet for Cloud service **SHOULD** support one of the following CoS Label sets: {H, M, L}, or {H, M} or {H, L}, using the Labels specified in MEF 23.1 [11].

A CP might need a certain number of traffic classes between cloud data center sites. A CEN might be capable of supporting a certain number of CoS Names. If this is less than the number of traffic classes required by the CP it is possible for the CP to aggregate traffic classes requiring similar service performance in to lesser number of CoS Names. The CoS ID for the CoS Name can be defined with more than 1 PCP or DSCP (MEF 10.3 [6]) which allows multiple traffic classes to get the same forwarding treatment in the CEN.

Table 4 provides an example mapping for Carrier Ethernet for Cloud traffic classes into 3 and 2 MEF standard CoS Names consistent with MEF 23.1 [11], i.e., CoS Labels H/M/L. CoS Labels are the names for the CoS for which CoS ID and Color ID types and values, Bandwidth Profile constraints, CPO values and parameter values are specified (MEF 23.1 [11]). In the table the following traffic classes are used:

- Critical Control – data center control traffic that is essential to the correct operation of DC network or services,
- Data Synchronization – traffic used in rapid (near real-time) update of data, e.g., database replication or VM migration,
- Streaming Media – traffic that can require limited delay variation and low frame loss (e.g., retransmission is not an option),
- Storage Migration – traffic used to transfer potentially large datasets but without strict timing requirements,
- Interactive WWW – traffic used to access Cloud-based applications, and
- Background – traffic that is deferrable (lowest priority, best effort service).

CoS Labels Defined in MEF 23.1 [11]	Generic Traffic Classes mapping to CoS Labels		
	3 CoS Labels	2 CoS Labels	2 CoS Labels
<b>High (H)</b>	Data Synchronization, and Critical Control	Data Synchronization, Critical Control and Streaming Media	Data Synchronization, and Critical Control
<b>Medium (M)</b>	Streaming Media and Storage Migration		Streaming media, Storage Migration, Interactive WWW and Background
<b>Low (L)</b>	Interactive WWW and Background	Storage Migration, Interactive WWW and Background	

**Table 4 – Examples of Carrier Ethernet for Cloud Traffic Classes mapping to CoS Labels**

The names of the traffic classes used in Table 4 are meant to represent a non-exhaustive set of generic traffic classes that could apply in Cloud services use cases.

- [D12] The mapping for supporting the entire set of traffic classes used generally for Carrier Ethernet for Cloud services **SHOULD** be based on the mapping of Generic Traffic Classes to CoS Names defined in Table 4.

A CP could also use multiple EVCs, with each EVC providing the CoS ID for a separate CoS Name for different traffic classes. In this case the CP needs the ability to classify the different traffic classes to different sets of CE-VLANs with EVC based CoS ID. The CEN can then map the traffic to different EVCs at the UNI with the CE-VLAN to EVC map. Different EVCs might also be appropriate if each traffic class requires different ingress bandwidth profile but are mapped to same CoS Name.

MEF 23.1 [11] specifies objective (and related parameter) values for each of the three standard CoS Labels by Performance Tiers (PT), loosely thought of as geographic tiers. PT1 has the tightest objectives, assuming a ‘metro scope’ scenario. See Tables 5 and 6 in MEF 23.1 [11]. For regional or global services, a different PT might apply. MEF 23.1, or a later revision, can be used as the foundation for defining the SLS for a Carrier Ethernet for Cloud service. MEF 23.1 [11] section 8.4 provides example performance requirements for a variety of applications, many of which are directly or indirectly related to Cloud services. Note that tighter objectives than the bounds given in MEF 23.1 [11] for a given PT might be appropriate.

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

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

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

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

The choice of any of these models allows a reasonable upper bound on FD, depending on percentile, while allowing some flexibility in the SLS offering.

- [D13] A Carrier Ethernet for Cloud SLS **SHOULD**, at a minimum, include {FLR, FD or MFD, FDR or IFDV} with objectives and parameters as specified in MEF 23.1[11].
- [D14] A Carrier Ethernet for Cloud SLS **SHOULD** use one of the three models described above, with objectives and parameters as specified in MEF 23.1[11] in the SLS for CoS Label H.
- [D15] A Carrier Ethernet for Cloud SLS **SHOULD** use one of the three models described above, with objectives and parameters as specified in MEF 23.1[11] in the SLS for CoS Label M.
- [D16] A Carrier Ethernet for Cloud SLS **SHOULD** use one of the three models described above, with objectives and parameters as specified in MEF 23.1[11] in the SLS for CoS Label L; except IFDV and FDR need not be specified.

Performance Attributes for which CPOs are not specified in MEF 23.1 [11] include Availability, HLI and CHLI.

## 12. References

- [1] Internet Engineering Task Force RFC 2119. Key words for use in RFCs to Indicate Requirement Levels, March 1997.
- [2] Special Publication 800-145, *The NIST Definition of Cloud Computing*, September 2011, <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>
- [3] Special Publication 500-292, *NIST Cloud Computing Reference Architecture*, September 2011, [http://www.nist.gov/customcf/get\\_pdf.cfm?pub\\_id=909505](http://www.nist.gov/customcf/get_pdf.cfm?pub_id=909505)
- [4] MEF 4, “Metro Ethernet Network Architecture Framework - Part 1: Generic Framework”, May 2004.
- [5] MEF 6.2, Ethernet Services Definitions, August 2014.

- [6] MEF 10.3, Ethernet Services Attributes Phase 3, October 2013.
- [7] MEF 12.2, “Carrier Ethernet Network Architecture Framework Part 2: Ethernet Services Layer”, May 2014.
- [8] MEF 13, “User Network Interface (UNI) Type 1 Implementation Agreement”, November, 2005.
- [9] MEF 17, "Service OAM Requirements & Framework – Phase 1", April 2007.
- [10] MEF 20, "User Network Interface (UNI) Type 2 Implementation Agreement", July 2008.
- [11] MEF 23.1, “Carrier Ethernet Class of Service – Phase 2”, January 2012.
- [12] MEF 26.1, "External Network Network Interface (ENNI) – Phase 2", January 2012.
- [13] MEF 30.1, "Service OAM Fault Management Implementation Agreement: Phase 2", April 2013.
- [14] MEF 33, “Ethernet Access Services Definition”, January 2012.
- [15] MEF 35, “Service OAM Performance Monitoring Implementation Agreement”, April 2012.
- [16] IEEE Std 802.1AX<sup>TM</sup> – 2008, IEEE Standard for Local and metropolitan area networks – Link Aggregation, November 2008.

## Appendix A. Carrier Ethernet Service for Cloud Examples

This appendix provides a few example use cases that can be supported by elastic Carrier Ethernet Services in cloud service scenarios.

Table 5 provides examples of one-way CPOs that might be required for Point-to-Point Carrier Ethernet Service for Cloud. The table also contains an indication of the bandwidth profiles (CIR and EIR) for each use case. The attributes for an elastic service would have to be set to accommodate the (temporary) requirements of the use case the service is intended to support. For example, to support the use cases shown, the CIR UB attribute could be set to the maximum bandwidth shown in the table to allow the CB to adjust CIR up to this value when required and reduce the CIR when the maximum value is not needed.

	LAN DCI Use Case	SAN DCI Use Case	DCA Use Case
	VM Migration (individual)	Storage Replication	Cloud Burst (Enterprise to CP)
Ingress BWP per CoS ID (MEF 6.2)			
CIR	≤3Gb/s	≤10Gb/s	≤1Gb/s
CBS	≥ Max Service Frame Size	≥ Max Service Frame Size	≥ Max Service Frame Size
EIR	0	0	<1Gb/s
EBS	N/A	N/A	≥ Max service frame
CoS Performance Objectives (MEF 10.3, 23.1)			
CoS Name	'H'	'M'	'H'
FD	≤5ms	≤40ms	≤5ms
MFD	≤3.75ms	≤30ms	≤3.75ms
IFDV	≤1ms	≤8ms	≤1ms
FDR	≤1.25ms	≤10ms	≤1.25ms
FLR	≤10 <sup>-4</sup>	≤10 <sup>-4</sup>	≤10 <sup>-4</sup>
Holding Time			
"High CIR" Duration	minutes	10s minutes ↔ hour	10s minutes

BWP = Bandwidth Profile      CoS = Class of Service      EVC = Ethernet Virtual Connection  
 CIR = Committed Information Rate      FD = Frame Delay      DCI = Data Center Interconnect  
 CBS = Committed Burst Size      MFD = Mean Frame Delay      DCA = Data Center Access  
 EIR = Excess Information Rate      IFDV = Inter-Frame Delay Variation      EBS = Excess Burst Size  
 FDR = Frame Delay Range      FLR = Frame Loss Ratio

Note: CPO values taken from MEF 23.1 clause 8.4.1 Table 23 (for VM Migration and Cloud Burst), Table 24 (for Storage Replication) as examples. The "Holding Time" information provides a rough idea of how long the modified attribute (in these examples a higher CIR) is expected to be required before reverting back to the original value.

**Table 5 – One-way CPO examples for Point-to-Point Carrier Ethernet Service for Cloud**

### A.1 Virtual Machine Migration

CPs perform Virtual Machine (VM) migrations (i.e., moving a live VM from one location to another) for reasons such as:

- improving performance of a VM cluster,
- load balancing, or

- server maintenance.

In general, the DC System Administrator initiates a migration via a VM Cluster Manager (VCM) which manages the source and target Hypervisors (this could be on-demand or scheduled for a maintenance window).

Migration is done over a TCP connection on a dedicated management VLAN and can involve migrating many VMs in parallel (e.g., to off-load a blade server prior to an upgrade).

The entire memory associated with a VM is initially copied followed by iterative updates of those data blocks whose content has changed since the last update.

An individual VM transfer size can range between 1-10GB, use up to 3Gb/s, take under 2 minutes, require essentially zero packet loss (FLR >>> 99.95%) and is round trip time limited (this is vendor specific, e.g., 10ms).

A VM migration can require an elastic service request to increase CIR by 1-3Gbps for the duration of the migration interval. In an on-demand scenario this modification can be required to be fulfilled in a few minutes and can last just for the duration of the migration activity, on the order of several minutes (i.e., a request to reduce CIR by the same amount can be received several minutes later, after the VM migration is completed).

## A.2 Storage Replication

A CP often needs to maintain multiple copies of data either to provide higher performance for data access or to support disaster recovery. Two scenarios of storage replication will be explored<sup>5</sup>:

- Asynchronous replication based on Fiber Channel extension over IP (FCIP)
- Asynchronous replication based on IP storage extension

Use of FCIP can extend native FC to North American inter-city distances (e.g., >1000km). Normally DCI will be provided using an EPL service, giving the CP full control over mapping customer VLANs to/from the DCI service without having to coordinate with the ECC. An EPL provides elastic point-to-point bandwidth with QoS guarantees to assure performance of data center operations. When storage replication is required the bandwidth of the service can be increased to expedite the transfer of a large amount of data, for example increasing CIR or {CIR+EIR} to 10Gb/s. Once the transfer is completed the bandwidth can be reduced to the level required for continuous operation. As storage replication flows can last minutes to hours, it is a good candidate for a scheduled set-up of an EVC

## A.3 Cloud Bursting

Cloud Bursting is the on-demand relocation of an Enterprise's applications from a private cloud infrastructure to a public cloud infrastructure to handle temporary spikes in demand. Cloud Bursting allows the Enterprise to take advantage of additional resources (CPUs, storage, and network bandwidth) available in the public cloud infrastructure located in CP data centers across

<sup>5</sup> Note that for synchronous replication FC is typically mapped into SONET/SDH or OTN Private Lines and will not be considered here.

a WAN. Cloud Bursting enables Enterprises to offload excess demand to CPs to address short term peaks in demand or to load balance between two or more data centers when local cloud computing resources are low or become unavailable. Cloud Bursting results in a seamless quality of experience for the CC's end user in that they can be unaware that they can be connecting to a different data center while accessing cloud applications.

In this use case it is assumed that the CC already has an Internet access service via the ECC (EVC A) to the CP's DC and the Internet.

EVC: CIR = <0>, EIR = <200> Mb/s, 'L' CoS, best effort

At the end of a business quarter for the Enterprise and for bookkeeping purposes, the Enterprise IT department arranges with the CP to schedule a VM offload. The VM uses 3GB consumed RAM and 5GB data store. The Infrastructure as a Service (IaaS) Cloud computing request is scheduled for a two-week period following the quarter.

IaaS: 2 CPU, 6GB RAM, 10GB data store / 2 weeks

The EVC bandwidth is increased between the Enterprise CC office and the CP DC via the ECC. During the cloud burst, the EVC is configured to increase bandwidth (EVC: CIR = <900> Mb/s, EIR = <0>) to replicate databases need for the CC applications. The EVC is also reconfigured to use CoS Label 'H' to allow VM migration to complete in under 2 minutes with less than 4 seconds pause time (note the data store replication timeframe is limited by the disk read speed at the Enterprise, e.g., 20k IOPS and a frame size of 2kB, so 40MB/s or 320Mb/s).

Following the VM migration and associated database replication, the EVC is re-configured to use CoS Label 'L' and the bandwidth is reduced to a level used for end user access (i.e., users within the Enterprise CC location) to the CP.

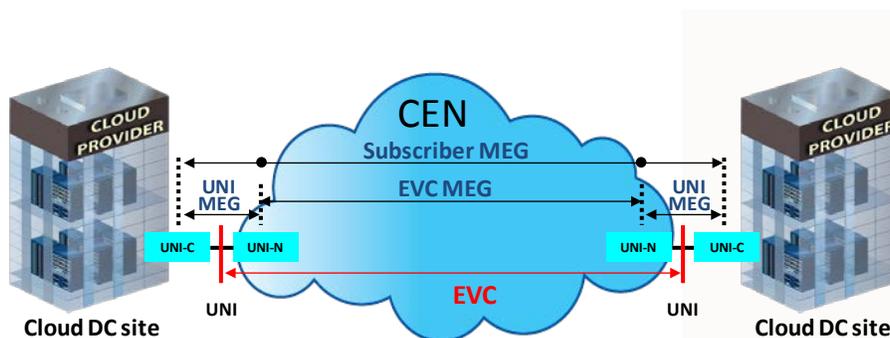
EVC: CIR = <100> Mb/s, EIR = <0>

## Appendix B. Carrier Ethernet Service OAM

This appendix describes the OAM model for Fault Monitoring (FM) and Performance Monitoring (PM) for a Carrier Ethernet Service for Cloud across a single CEN domain. In addition to Service OAM (MEF 17 [9]), Link OAM (MEF 20 [10]) is also specified for use across a UNI.

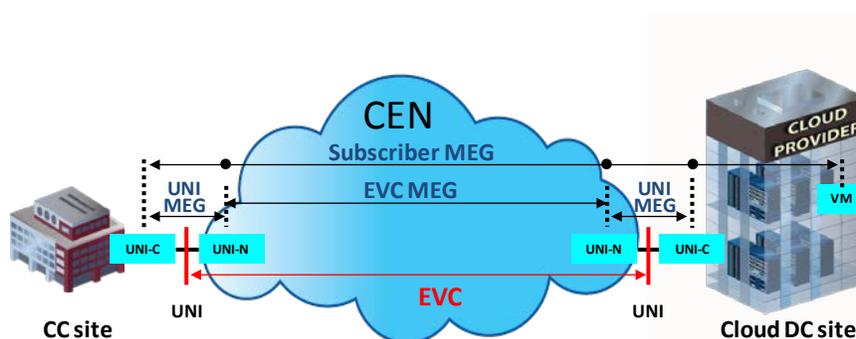
OAM is a term used in this IA to collectively refer to Link OAM (MEF 20 [10]) and Service OAM (MEF 17 [9], MEF 30.1 [13], and MEF 35 [15]). ECCs, CCs, and CPs can implement OAM to monitor their services and service interfaces.

An example FM and PM reference model for a DCI use case is illustrated in Figure 9 below. The figure shows the reference model for Service and SOAM for FM as well as PM.



**Figure 9 – FM and PM Reference Model for DCI Use Case**

Figure 10 shows an example FM and PM reference model for a DCI use case. The DCA service is a Point-to-Point EVC between a CC site and a CP data center site. Note that in this case the Subscriber MEG can extend from the CC UNI-C to the CP UNI-C or (as shown in the figure) all the way to a VM within the CP data center.



**Figure 10 – FM and PM Reference Model for DCA Use Case**

SOAM (MEF 17 [9] and MEF 30.1 [13]) is used on the different service components (UNI, EVC) by CPs and CCs as well as the ECC. Methods describing how to perform PM measurements can be found in MEF 35 [15].