Carrier Ethernet for Delivery of Private Cloud Services

An overview of the MEF Ethernet Services for Cloud Service Providers and Cloud Service Concepts for Carrier Ethernet Service Providers

February 2012

Version 1.1

Table of Contents

List of Figures

1 Introduction

The wide area network (WAN) plays a critical role in the quality of experience a cloud service will provide to an end user. However, there has been little focus on the WAN by the cloud community since the WAN is assumed to be the Internet. Furthermore, enterprises have been reluctant to embrace mission-critical private cloud applications delivered over the Internet because of its security vulnerabilities, unpredictable performance, and compliance challenges regarding data governance.

Carrier Ethernet has become the wide area networking technology of choice for delivery of business-class services to the Enterprise with deployment in more than 100 countries. In 2011, the MEF (Metro Ethernet Forum) created a new project focusing the role of Carrier Ethernet networks and services in the delivery of private cloud services to enterprises. This paper is the first outcome of that project.

2 Purpose and Scope

The scope of the new MEF cloud project covers several phases. The first phase, covered by this paper, is to introduce the concept of delivering private cloud services via Carrier Ethernet wide area networks and services. The intention of this paper is to educate and initiate a dialog between the cloud industry's stakeholders (enterprise users, cloud service providers, standards development organizations and cloud fora) as the MEF develops new Ethernet service attributes to better align with the dynamic, on-demand nature of cloud services.

The paper discusses the reasons why this is an important and beneficial development for Cloud Service Providers, Cloud Consumers and communications service providers delivering Ethernet services. It provides several use cases of MEF standardized services to support cloud applications. Terminology used in this document is based on the U.S. National Institute of Standards and Technology (NIST) definitions for Cloud Services [1]. It is possible that the MEF will recommend some additions to these definitions to accommodate the concepts introduced by this work.

3 Why Use Carrier Ethernet to Deliver Private Cloud Services?

Today, the Internet is the predominant method to connect to cloud services which has served today's Cloud Consumers well. The Internet is the most ubiquitous type of WAN and is readily available in all markets. Because of this, the Cloud community has given little attention to WANs because it is assumed that the WAN would be the Internet. Larger Enterprises are hesitant to move mission-critical applications to the cloud when delivered via the Internet due to concerns with security, network performance, data governance and regulatory compliance.

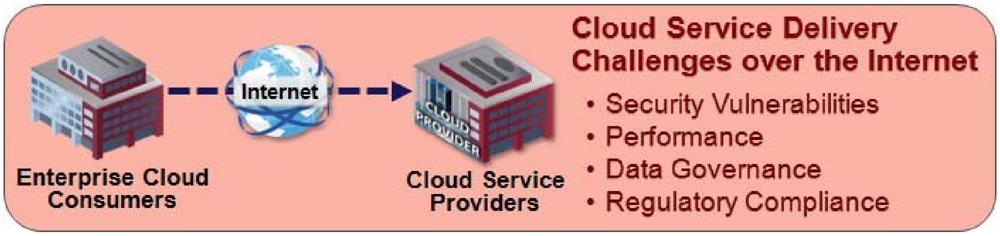

Figure 1: Current Challenges with Cloud Service Delivery

Carrier Ethernet is well suited to address these concerns. Communications service providers have a long history providing WAN services to enterprises. They can leverage this trusted relationship to deliver private cloud services in addition to their existing Ethernet WAN services. Ethernet service providers, such as telecommunications service providers, Ethernet Exchange providers and MSOs, can provide on-net data center co-location facilities to a Cloud Service Provider. This arrangement would result in better and more predictable cloud service performance, improved security and better control over their enterprise customer's data governance and regulatory compliance requirements.

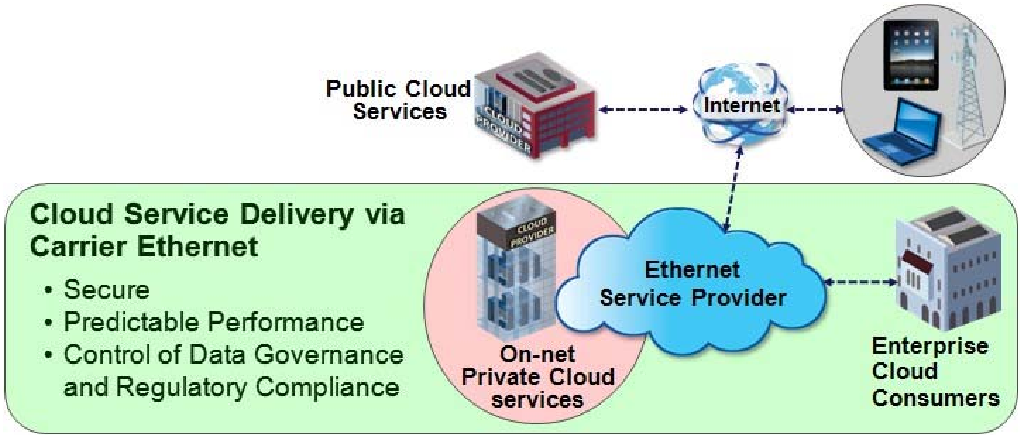

Figure 2: Ethernet Services to Public/Private Cloud Service Delivery

3.1 Benefits for each party

3.1.1 Cloud Service Providers

Carrier Ethernet provides an industry-standardized, business-class delivery mechanism for cloud-based applications. Using a single connection to an Ethernet Service Provider, the Cloud Service Provider can connect with its enterprise Cloud Consumers over a dedicated network connection or a Layer 2 VPN. The same Carrier Ethernet network can also provide connectivity to the Internet for delivery of public cloud services.

3.1.2 Ethernet Service Provider

Ethernet service providers, acting as a Cloud Carrier, have the opportunity to create new revenue streams by providing Cloud services. They can create strong differentiation and leverage customer loyalty by reinforcing their Ethernet services as the business class service of choice. They are also in a good position to align their network service SLAs with their Cloud service SLAs.

3.1.3 Cloud Consumers

Carrier Ethernet enables enterprise IT department management and control over the delivery of cloud-based services for their Cloud Consumers without the inherent security risks associated with the Internet. Carrier Ethernet addresses enterprise requirements for high performance with secure and controlled access to private cloud applications. It also facilitates compliance with corporate governance which has been a challenge for cloud services delivered over the Internet. For enterprise Cloud Consumers, Carrier Ethernet creates a predictable end user quality of experience because the private cloud service delivery is controlled from end to end.

4 Definitions and Terminology

As a preliminary step to describing the cloud taxonomy being adopted in today's market, it is important to understand the simple view that spawned the infamous "cloud" adjective for cloud computing and all related services and technology. The "cloud" is something the communications service providers have used for years and it essentially helps abstract the details of the technology (traditionally network technology) that the subscriber does not need to understand. For example, provider edge routers, optical networking equipment, Ethernet switches, etc. are abstracted by an image of a cloud. Subscriber's access to the network is depicted by lines joining the cloud to the subscriber sites.

The key evolution with the "cloud" in cloud computing is that additional domains beyond the network domain are also abstracted into the cloud. The subscriber continues to see only the service and service connection points and not the details of the technology infrastructure that is "inside the cloud". The service could be an application, compute, storage or network service or a combination of these.

4.1 MEF Definitions

4.1.1 MEF-defined Carrier Ethernet

Carrier Ethernet, unlike Ethernet used in the LAN, is differentiated by attributes uniquely required for wide area connectivity such as supporting multiple subscribers over a common network infrastructure and being able to troubleshoot network issues remotely. The MEF characterizes Carrier Ethernet by the following five attributes:

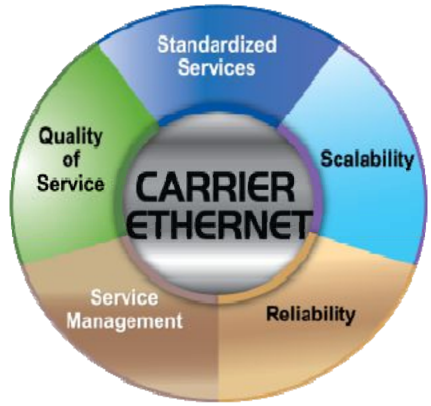

- Standardized Services
- Scalability
- Reliability
- Service Management
- Quality of Service

The MEF has defined three UNI-to-UNI Ethernet service types based on the type of connectivity they provide.

- E-Line for point-to-point connectivity
- E-LAN for multipoint-to-multipoint connectivity
- E-Tree for rooted-multipoint connectivity

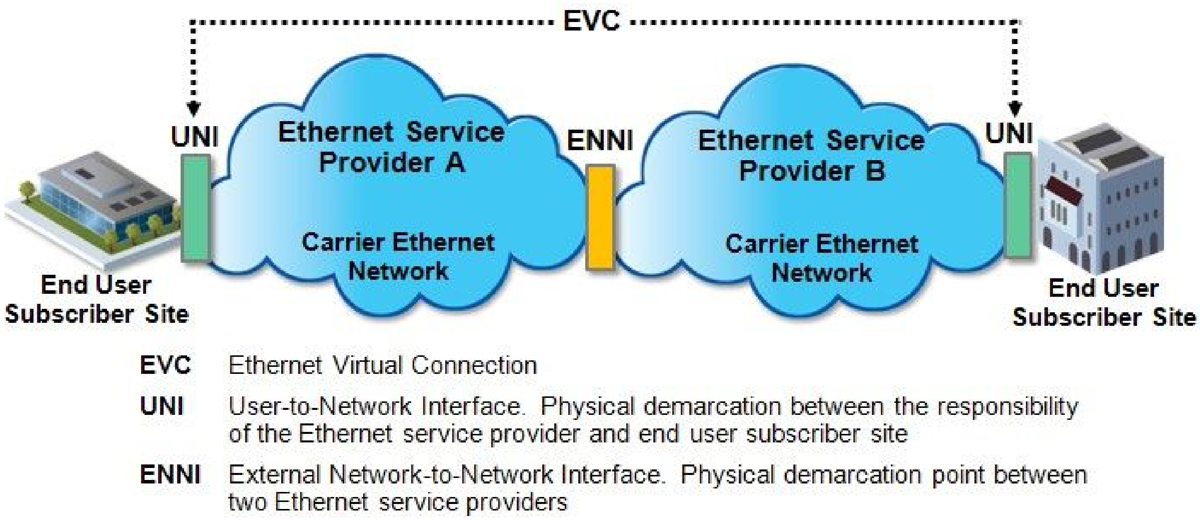

Figure 3: Carrier Ethernet Network showing EVC with UNI and ENNI service demarcation points

These service types use end-to-end Ethernet virtual connections (EVCs) between user-to-network (UNI) interfaces. The EVC is the logical representation of an Ethernet service instance. Each EVC can have a bandwidth profile which specifies the committed and excess information rates (CIR and EIR) to which traffic is transmitted to or received from the Ethernet service provider's network. The EVC may also support one or more classes of service and measurable Quality of Service (QoS) performance metrics, such as frame delay (latency) and frame loss, to accommodate various application performance requirements. The MEF has defined these Ethernet services independent of the underlying transport network technology.

A Carrier Ethernet network must support large scale metro network requirements and will often need to support numbers of subscribers (users) and applications that are often orders of magnitude higher than Ethernet users on a LAN. Given the large scalability required, the Carrier Ethernet network must be as reliable as the SONET/SDH networks that businesses have relied on for decades.

Service management addresses the ability to remotely provide OAM&P (Operations, Administration, Maintenance and Provisioning) of Ethernet services which is a fundamental requirement for delivering wide area networking services. Finally, the ability to deliver QoS performance assurances for metrics such as frame delay and frame loss as part of a service level agreement enable Ethernet service providers to differentiate their offerings.

4.1.2 MEF Ethernet Service Definitions

The MEF categorizes Ethernet services by their connectivity requirement, e.g., point-to-point or multipoint, and the interfaces type (UNI or ENNI). This is referred to as the Ethernet Service Type. The E-Line, E-LAN and E-Tree service types encompass all services that interconnect UNIs and represent point-to-point, multipoint-to-multipoint and rooted multipoint connectivity, respectively. The E-Access service type encompasses all services that interconnect an ENNI with at least one UNI.

The Ethernet Service Types are used to refer to all services within the category. Within each Ethernet Service Type, the MEF defines Ethernet services based on whether a single Ethernet service instance is supported at the UNI (port-based or Private services) or whether multiple Ethernet service instances are supported at the UNI (VLAN-based or Virtual Private services). Refer to Figure 4. One of the benefits of a VLAN-based service is the ability to multiplex multiple services onto a single UNI saving the cost of an additional Ethernet port on an Ethernet service subscriber's WAN attaching devices. Section 6 provides use cases for the different Ethernet services.

Service Type	Port-based Services	VLAN-Based Services
E-Line Point-to-Point, UNI-to-UNI	**Ethernet Private Line** (EPL)	**Ethernet Virtual Private Line** (EVPL)
E-Access UNI(s)-to-ENNI	**Access EPL**	**Access EVPL**
E-LAN Multipoint-to-Multipoint, UNIs-to-UNIs	**Ethernet Private LAN** (EP-LAN)	**Ethernet Virtual Private LAN** (EVP-LAN)
E-Tree Rooted-Multipoint, Root UNI(s)-to-Leaf UNIs	**Ethernet Private Tree** (EP-Tree)	**Ethernet Virtual Private Tree** (EVP-Tree)

Figure 4: MEF Ethernet Service Types with Port-based and VLAN-based Service definitions

4.2 NIST Definitions

NIST has been frequently referenced by other worldwide standards development organizations (SDOs) to help define the cloud taxonomy. NIST has a mandate from the U.S. Federal Government to accelerate the federal government's secure adoption of cloud computing by leading efforts to develop standards and guidelines in close consultation and collaboration with SDOs. NIST has a specific group, "The Architecture and Taxonomy Working Group", developing a reference architecture for the cloud [2]. The MEF has chosen to use the NIST model with some specific and relevant adaptation for the purpose of clarity and consistency in discussing the role of Carrier Ethernet within cloud computing and its related services.

4.2.1 NIST Cloud Terminology

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources, e.g., networks, servers, storage, applications and services, that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics listed below, three service models (refer to Figure 5) and four deployment models (refer to Figure 6).

- On-demand Self-service
- Broad Network Access
- Resource Pooling
- Rapid Elasticity
- Measured Service

The MEF is using the NIST models and is augmenting the terminology and definitions as they apply to Carrier Ethernet networks and services.

4.2.2 NIST Cloud Service Models

The next component to understanding the “cloud” is the classifications of cloud services - essentially what the Cloud Consumer needs from the Cloud Service Provider.

Consumer Type	Major Activities	Example Users
SaaS Software as a Service	Provides the applications for business process operations	Business users, software application administrators
PaaS Platform as a Service	Provides the development, testing, deployment and management of applications hosted in a cloud environment	Application developers, testers and administrators
IaaS Infrastructure as a Service	Provides the creation, management and monitoring of compute, storage and network resources for IT infrastructure operations	System developers, administrators and IT managers

Figure 5: NIST Cloud Service Models

4.2.3 NIST Cloud Service Deployment Models

Ambiguity frequently revolves around private versus public clouds, whether the operator of the cloud service is the enterprise itself or a 3rd party service provider, and whether the service is delivered over the Internet or a private data network. In Figure 6, the NIST service models are referenced with an enhancement suggested by the TMForum’s Enterprise Cloud Leadership Council (ECLC) to clarify who manages and hosts the cloud service [3].

Service Deployment Type	Description	Operation with respect to the Enterprise	
		External	Internal
Private Cloud	Cloud infrastructure operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.	Hosted and operated by a 3rd party cloud service provider	Hosted and operated by the enterprise
Public Cloud	Cloud infrastructure made available to the general public or a large industry group and is owned by an organization selling cloud services.	Publically available and accessible over the Internet	NA
Community Cloud	Cloud infrastructure shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise.	Hosted and operated by 3rd party cloud service provider	Hosted and operated by the community
Hybrid Cloud	Cloud infrastructure composed of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).	Potentially a combination of both externally-based and internally-based operations.	

Figure 6: NIST Cloud Service Deployment Models

4.2.4 NIST Cloud Roles

To help define the potential role of the organization involved with the cloud service, it's helpful to look at the general roles (labeled "Actor" by NIST) related to cloud computing. Refer to Figure 7.

Actor	Definition
Cloud Consumer	The person or organization that maintains a business relationship with, and uses service from, Cloud Service Providers.
Cloud Service Provider	The person, organization or entity responsible for making a service available to Cloud Consumers.
Cloud Auditor	A party that can conduct an independent assessment of cloud services, information system operations, performance and security of the cloud implementation
Cloud Broker	An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between Cloud Service Providers and Cloud Consumers.
Cloud Carrier	The intermediary that provides connectivity and transport of cloud services from Cloud Service Providers to Cloud Consumers.

Figure 7: NIST Cloud Roles

4.3 MEF Augmentation of NIST Terminology

For the purpose of exploring a Carrier Ethernet cloud framework, the MEF has created additional terminology that defines specific use cases of the NIST definitions and provides additional clarification for the different NIST Actors. MEF augmentation to the NIST definitions is identified by underlined text in Figure 8.

Actor	MEF Augmented Definition
MEF Ethernet Cloud Carrier	The intermediary that provides <u>Carrier Ethernet</u> connectivity and transport of cloud services from Cloud Service Providers and Cloud Consumers <u>and with other Ethernet Cloud Carriers</u>.
NIST Cloud Carrier	The intermediary that provides connectivity and transport of cloud services from Cloud Service Providers to Cloud Consumers.
MEF Cloud Broker	An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between Cloud Service Providers, Cloud Consumers <u>and Ethernet Cloud Carriers</u>.
NIST Cloud Broker	An entity that manages the use, performance and delivery of cloud services, and negotiates relationships between Cloud Service Providers and Cloud Consumers.

Figure 8: MEF Terminology to augment NIST cloud definitions for Carrier Ethernet

4.3.1 Ethernet Cloud Carrier

The Ethernet Cloud Carrier definition is a specific instance of the NIST definition. Here, the Ethernet Cloud Carrier uses Carrier Ethernet as the wide area networking technology where NIST provides a generic definition independent of the underlying wide area networking technology used. If one assumes that all Cloud Carriers provide connectivity over the Internet, then it is not necessary to specify the networking technology because it would be the Internet protocol (IP). Carrier Ethernet can be used for creating private or virtual private network connections between Cloud Consumers and Cloud Service Providers or between Cloud Service Provider data centers where none of the traffic ever traverses the Internet. This approach is ideal for private cloud applications where the Cloud Consumer requires more deterministic network performance such as bounded network latency and loss, higher availability, different classes of service and more controlled network security.

The MEF definition further specifies that an Ethernet Cloud Carrier can also provide connectivity between other Ethernet Cloud Carriers in addition to connectivity between Cloud Consumers and Cloud Service Providers. This is because the Ethernet connectivity services are often delivered by more than one Ethernet Cloud Carrier. Such a scenario would be where the Cloud Consumer's location is accessible by one Ethernet Cloud Carrier and the Cloud Service Provider's data center is accessible by a different one. Each Ethernet Cloud Carrier would establish a peering agreement to interconnect their networks to provide the end-to-end connectivity between Cloud Consumer and Cloud Service Provider.

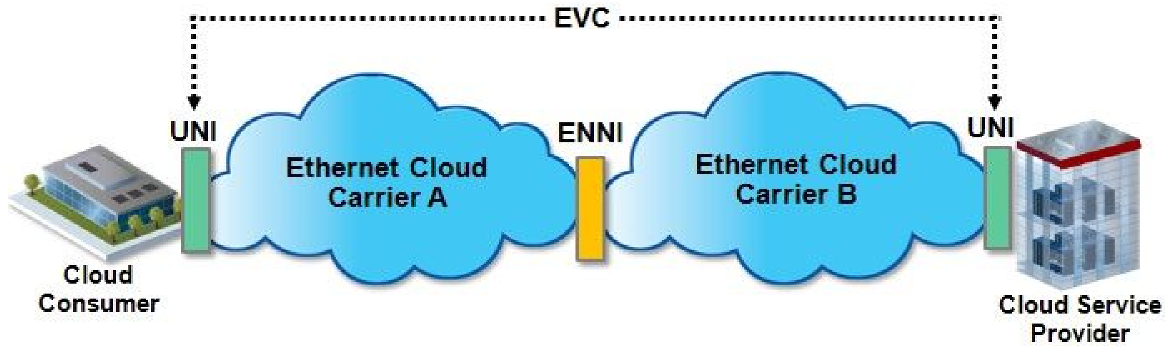

Figure 9: Scenario with Two Ethernet Cloud Carriers

4.3.2 Extension of Cloud Broker Definition

The MEF augments NIST's Cloud Broker definition by adding the Cloud Carrier (specifically the Ethernet Cloud Carrier) actor role to the definition. The Cloud Carrier is typically a separate entity from the Cloud Service Provider and therefore is one of the stakeholders involved in a brokered model. If one assumes that everyone uses the Internet to connect to cloud services, then the Cloud Carrier is assumed to be an Internet service provider and hence it is not noteworthy to mention it in the Cloud Broker definition. However, when alternative WAN technologies such as Carrier Ethernet are used to provide cloud connectivity, it is noteworthy to add the Cloud Carrier to the Cloud Broker definition as the MEF is proposing. Cloud Brokers could also act as the broker between Cloud Consumers, Cloud Service Providers and Cloud Carriers for private WAN connections between Cloud Consumers and Cloud Service Providers. An example of this would be an Ethernet Exchange Provider. A Cloud Broker could also be an Ethernet Cloud Carrier or Cloud Service Provider and broker between the three Actors while delivering the functions of the Cloud Carrier or Cloud Service Provider. Refer to Section 5 for more detailed examples.

5 Cloud Broker Use Cases

This section describes several use cases resulting in more detailed Cloud Broker definitions than those provided by NIST. The use cases will better explain how Ethernet Cloud Carriers and Cloud Service Providers can play the role of a Cloud Broker and the business relationship between the different NIST actors in a cloud ecosystem. In the use case figures, the solid line interconnecting the different boxes represents a business relationship and service level agreement (SLA) between the Actors involved.

5.1 Ethernet Cloud Carrier acting as a Cloud Broker

In this use case, the Ethernet Cloud Carrier has the business relationship with the Cloud Consumer and wholesales cloud services from one or more Cloud Service Providers (or could resell cloud services from its own cloud services organization). The Cloud Consumer has no business relationship, e.g., billing, service agreement, etc., with the Cloud Service Provider. The Ethernet Cloud Carrier in the model is acting in two different roles, namely Cloud Broker and Ethernet Cloud Carrier. In the telecoms industry, this type of Cloud Broker is analogous to a telecoms service provider delivering an Internet access service. In this case, the telecoms service provider is delivering and managing network connectivity and transport to the customer premises and provides Internet access like an Internet service provider (ISP).

Figure 10: Ethernet Cloud Carrier acting as a Cloud Broker

5.2 Cloud Service Provider acting as a Cloud Broker

In this use case, the Cloud Service Provider has the business relationship with the Cloud Consumer and wholesales network connectivity and transport services from one or more Ethernet Cloud Carriers. The Cloud Consumer has no business relationship, e.g., billing, service agreement, etc., with the Ethernet Cloud Carrier. The Cloud Service Provider in the model is acting in two different roles, namely Cloud Broker and Cloud Service Provider. In the telecoms industry, this type of Cloud Broker is analogous to an ISP delivering Internet access service to their customers and wholesaling the access network connection from the local incumbent telecoms service provider who owns the access lines to the customer premises. In effect, the ISP is delivering and managing Internet access and providing network connectivity and transport to the customer premises.

Figure 11: Cloud Service Provider acting as a Cloud Broker

5.3 Generalized Cloud Broker

In this use case, the Cloud Broker has the business relationship with the Cloud Consumer and wholesales network connectivity and transport services from one or more Ethernet Cloud Carriers and cloud services from one or more Cloud Service Providers. The Cloud Consumer has no business relationship, e.g., billing, service agreement, etc., with either the Ethernet Cloud Carrier or Cloud Service Provider. In the telecoms industry, this type of Cloud Broker is analogous to a mobile virtual network operator (MVNO) who sells cellular wireless voice and data services to their customers even though they do not own or operate a cellular wireless network. However from the customer's perspective, they are buying cellular wireless voice and data services from a cellular wireless network operator.

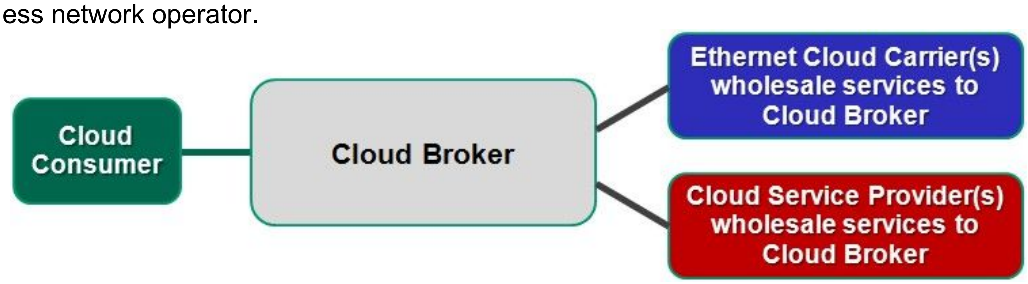

Figure 12: Cloud Broker providing broker service for Ethernet Cloud Carrier and Cloud Service Provider

6 MEF Carrier Ethernet-to-Cloud Focus Areas and Use Cases

The MEF is focusing on the use of MEF standardized Ethernet services for two distinct applications related to cloud computing. Refer to Figure 13.

A. Enterprise Cloud Consumer business locations connecting to a Cloud Service Provider via a Carrier Ethernet network managed by an Ethernet Cloud Carrier
B. Interconnecting Cloud Service Provider data centers via a Carrier Ethernet network managed by an Ethernet Cloud Carrier[1]

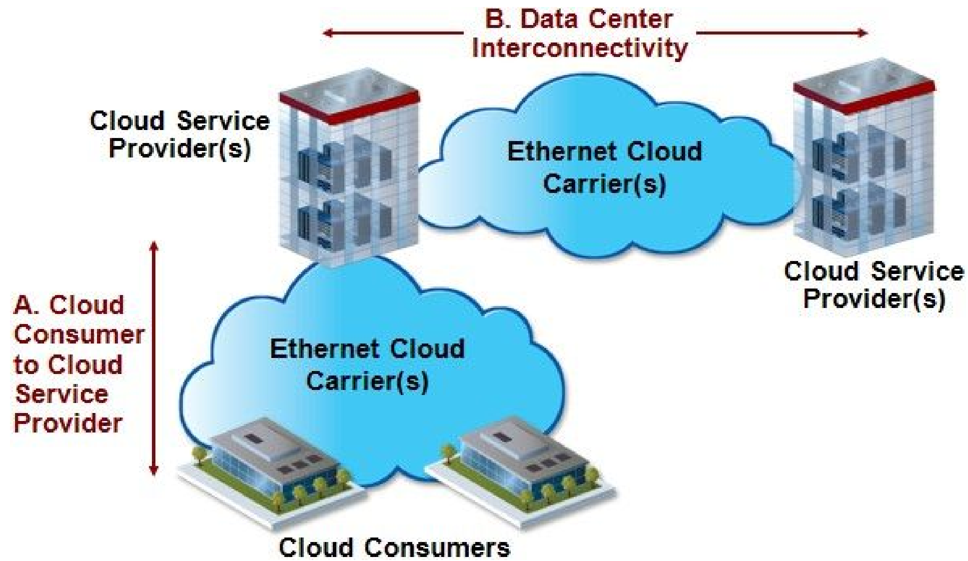

Figure 13: MEF Carrier Ethernet-to-Cloud Focus Areas

Figure 13 illustrates two different Ethernet Cloud Carriers performing different functions, namely, Cloud Consumer to Cloud Service Provider interconnectivity and Data Center interconnectivity. However, these two functions could be performed by the same Ethernet Cloud Carrier. Furthermore, each Ethernet Cloud Carrier could be decomposed into multiple Ethernet Cloud Carriers if a single Ethernet Cloud Carrier cannot provide the end-to-end connectivity. For example, if the Cloud Consumer’s business location is not geographically near the Cloud Service Provider’s data center, two or more Ethernet Cloud Carriers may need to interconnect their networks to provide the end-to-end WAN connection.

Different use cases of MEF Ethernet services are illustrated in the remainder of this section. To simplify the diagrams, a single Ethernet Cloud Carrier is used to provide the WAN connectivity between the Cloud Consumer and the Cloud Service Provider. Note that the Ethernet Cloud Carrier or Cloud Services Provider could also act as a Cloud Broker as discussed in Section 5.

6.1 E-Line Service Type

The MEF E-Line service type refers to a class of Ethernet services that use a Point-to-Point Ethernet Virtual Connection (EVC) interconnecting two UNIs. E-Line services provide connectivity analogous to other familiar point-to-point WAN services such as TDM private lines, Frame Relay, ATM, SONET, and wavelength services.

6.1.1 Ethernet Private Line (EPL)

MEF EPL services are applicable to use cases where a Cloud Consumer requires a dedicated Ethernet WAN connection to a Cloud Service Provider at a data center. Each UNI at the Cloud Consumer and Cloud Service Provider sites support a single EVC. An EPL service does not require any VLAN ID coordination with the Ethernet Cloud Carrier since all VLAN IDs are passed transparently across the network. In Figure 14, two Cloud Consumers (A and B) each use an EPL service to connect to the Cloud Service Provider’s primary data

[1] Multiple Cloud Service Providers tenants may be co-located at a given data center. The MEF will consider this in their work to determine if there are unique Ethernet service and external interface (UNI / ENNI) requirements between single tenant and multi-tenant scenarios.

center. A third EPL service is used to interconnect the Cloud Service Provider's primary and secondary data centers.

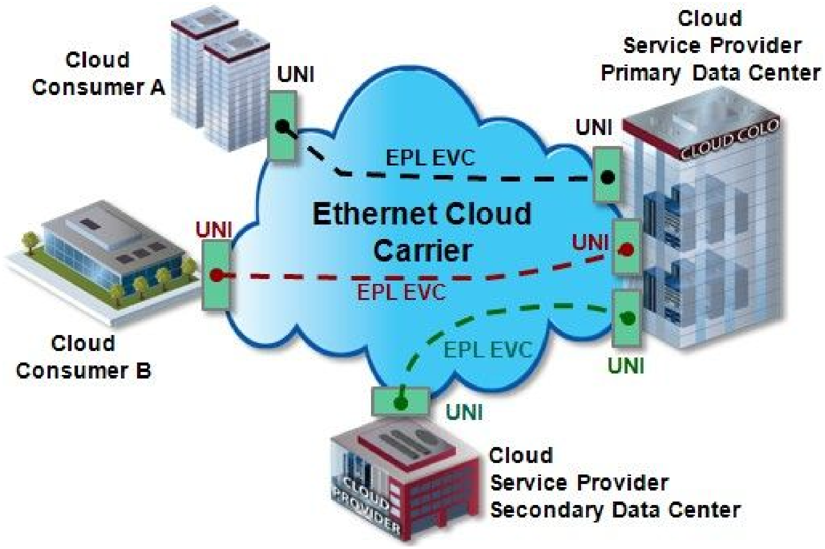

Figure 14: EPLs to connect Cloud Consumers to Cloud Service Provider and data center interconnect

6.1.2 Ethernet Virtual Private Line (EVPL)

MEF EVPL services are applicable to use cases where a Cloud Consumer or Cloud Service Provider wants to multiplex EVCs (services) onto the same UNI. This enables them to save the cost of an extra Ethernet port on their attaching equipment but often saves service cost because the Ethernet Cloud Carrier doesn't need to install a separate port for the additional EVC as would be the case with an EPL service. However, the Cloud Service Provider will need to coordinate customer VLAN IDs (C-VLAN IDs) to ensure that Cloud Consumers A and B do not use overlapping C-VLAN IDs.

In Figure 15, two Cloud Consumers (A and B) each use an EVPL service to connect to the data center in which their Cloud Service Provider is located. At the data center, the EVCs from each EVPL service are multiplexed onto the same UNI (UNI C). Because the EVPL service supports EVC multiplexing at a UNI, additional EVCs can be added to UNI C at the data center to connect to additional Cloud Consumer sites, connect to the Internet or even connect to another Cloud Service Provider data center. Furthermore, additional EVCs can be added at the UNIs at Cloud Consumers A and B for other services, e.g., Internet Access, site-to-site Ethernet L2 VPN, etc.

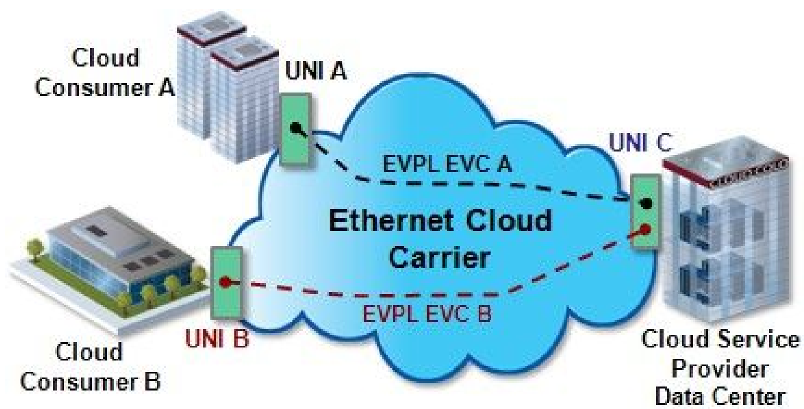

Figure 15: EVPLs from two Cloud Consumers multiplexed at Cloud Service Provider's UNI

In Figure 16, Cloud Consumer B uses an EVPL service to connect to the Internet and another EVPL service to connect to the Cloud Service Provider's data center. At the Cloud Consumer's UNI, the EVCs for each EVPL service are multiplexed onto the same UNI (UNI B).

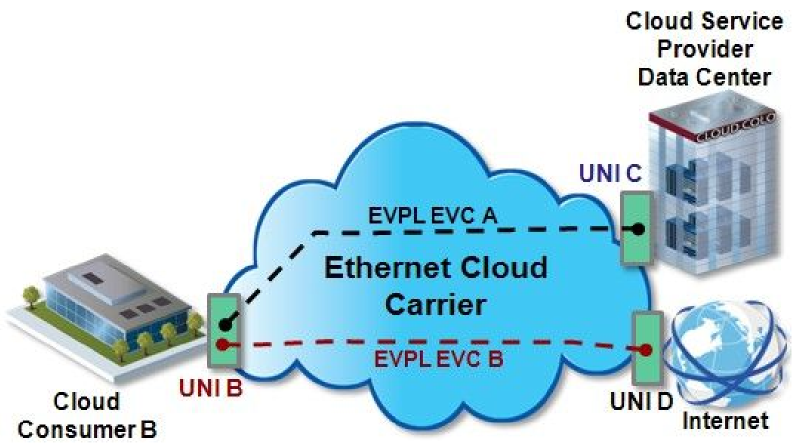

Figure 16: EVPL Use Case with 2 EVCs multiplexed at Cloud Consumer UNI

6.2 E-Tree Service Type

The MEF E-Tree service type refers to a class of Ethernet services that use a Rooted Multipoint EVC. This type of EVC has two different types of UNIs, namely, Leaf UNIs and Root UNIs. An E-Tree service type supports data exchange between Root UNIs but prohibits Leaf UNI data exchange. Leaf UNIs can only exchange data with Root UNIs. Since E-Tree services prohibit data exchange between Leaf UNIs, different Cloud Consumers can connect to the same E-Tree service EVC with the necessary traffic isolation to ensure no data interchange occurs between different Cloud Consumer organizations. The Cloud Service Provider can add additional Cloud Consumers (Leaf UNIs) to the existing E-Tree EVC.

6.2.1 EP-Tree

MEF EP-Tree services can be used to connect different Cloud Consumer locations to a common Cloud Service Provider data center. In Figure 17, Cloud Consumers A, B and C connect to two Cloud Service Provider data centers. Leaf UNIs are used at the Cloud Consumer sites and Root UNIs are used at the data center sites.

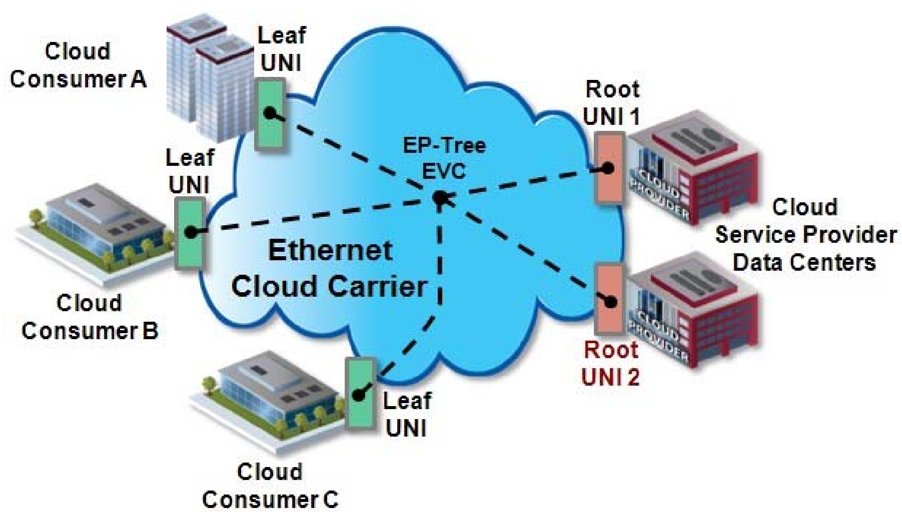

Figure 17: EP-Tree to interconnect Cloud Consumers with Cloud Service Provider Data Centers

The EP-Tree provides the traffic separation and isolation between Cloud Consumers A, B and C and enables the Cloud Services Provider to add additional Cloud Consumers to the same EP-Tree. If C-VLAN IDs are used, then they must be coordinated with the different Cloud Consumers to ensure no overlapping values are used. Cloud Consumers A, B and C have access to either of the Cloud Service Provider's data centers connected by Root UNIs 1 and 2. Since Root UNIs enable intercommunication, each Root UNI can be used to interconnect a different data center facilitating data replication, disaster recovery and cloud bursting applications.

6.2.2 EVP-Tree

MEF EVP-Tree services can address the same applications as EP-Tree services described in Section 6.2.1 with the added benefit of supporting multiple services on the same UNI.

EVP-Tree services are used when a Cloud Consumer or Cloud Service Provider wants to multiplex EVCs (services) onto the same UNI. This capability enables them to save the cost of extra ports on the equipment attaching to the Ethernet Cloud Carrier. Additional cost reductions are possible because the Ethernet Cloud Carrier doesn't need to install a separate port for the additional EVC on their equipment as would be the case with an EP-Tree service. However, the Cloud Service Provider will need to coordinate customer VLAN IDs (C-VLAN IDs) to ensure that each EVC on a given UNI does not use overlapping C-VLAN IDs.

Figure 18 provides a use case where an EVP-Tree service is used to interconnect different Cloud Consumers to the Cloud Service Provider's data centers similar to the EP-Tree service use case illustrated in Figure 17. However, EVP-Tree service supports multiple services on a given UNI. In Figure 18, Cloud Consumer A has an additional EVPL EVC providing a connection to the Internet on the same UNI as the EVP-Tree EVC connecting to the Cloud Service Provider data centers. The Cloud Service Provider also has an EVPL EVC on Root UNI 1 providing a connection to the Internet which could be used to deliver public cloud services or support cloud bursting over the Internet to another data center when local compute and storage resources are fully consumed.

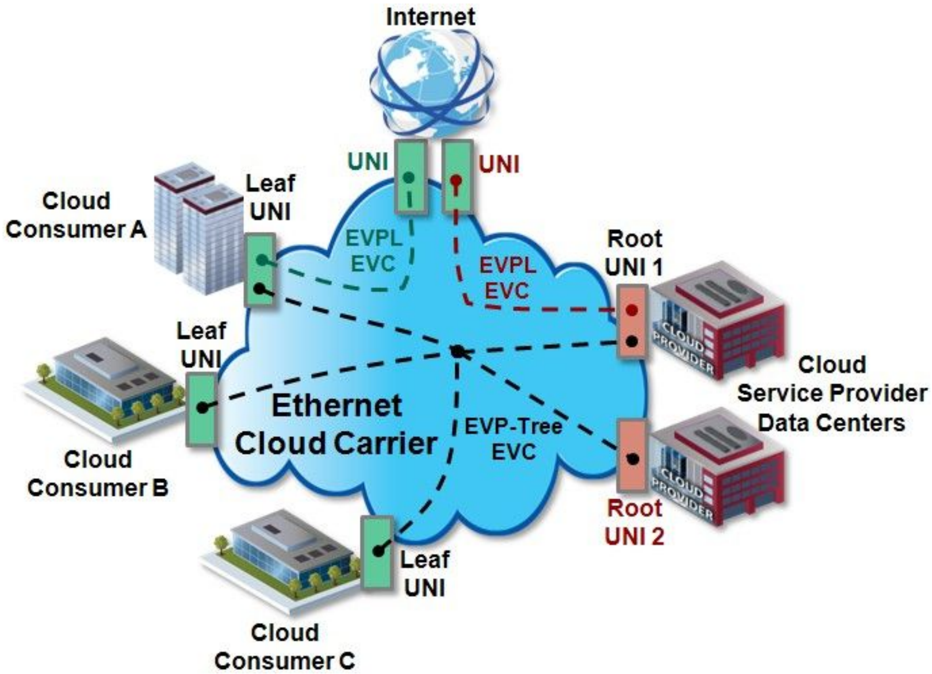

Figure 18: Example EVP-Tree use case with EVPLs used to provide Internet connectivity

6.3 E-LAN Service Type

The MEF E-LAN service type refers to a class of Ethernet services that use a Multipoint-to-Multipoint EVC. This type of EVC operates most similarly to an Ethernet LAN (with augmentation to support WAN service requirements) by providing a broadcast domain where any UNI can exchange data with any other UNI. Unlike E-Tree services, E-LAN services do not restrict data exchange between any of their UNIs. Therefore, only different sites of the same Cloud Consumer organization should be connected to the same E-LAN service EVC because data exchange cannot be restricted.

Note that with an E-LAN service, all Cloud Consumer data can be exchanged and will be visible at all UNIs in the EVC including the Cloud Service Provider who will often be a 3rd-party organization. For information security protection, when using an E-LAN service, data exchanged on the E-LAN EVC should be dedicated to the cloud services provided by the Cloud Service Provider.

E-LAN services facilitate data replication, disaster recovery and cloud bursting applications between Cloud Service Provider data centers connected to the E-LAN service EVC. Furthermore, Cloud Consumers connected to the same E-LAN service EVC as the data centers have full access to both primary and backup data centers which could simplify the Cloud Consumer's disaster recovery business requirements.

Finally, E-LAN services enable the Cloud Consumer organization to add additional sites (UNIs) to the same E-LAN EVC. The newly added sites will get the same level of connectivity and data exchange capabilities as all of the other connected sites.

6.3.1 EP-LAN

MEF EP-LAN services can be used to connect different Cloud Consumer locations to one or more Cloud Service Provider data centers. EP-LAN services can only support a single service at each UNI. In Figure 19, a Cloud Consumer has three sites (Sites 1, 2 and 3) that connect to a Cloud Service Provider's primary and secondary data centers. The secondary data center could be used for cloud bursting applications to provide additional virtual machine (VM) or data storage capacity to the primary data center. It could also be used as a disaster recovery data center. Note that if C-VLAN IDs are used, then they must be coordinated with the different Cloud Consumers sites to ensure no overlapping values are used.

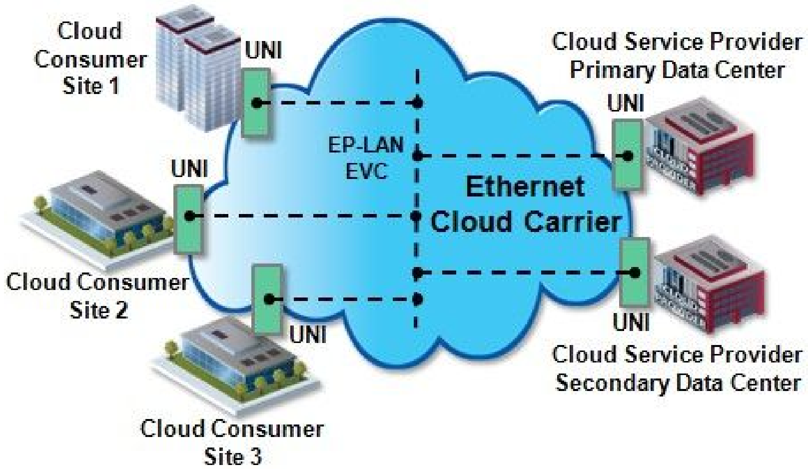

Figure 19: EP-LAN interconnecting Cloud Consumer sites with Cloud Service Provider data centers

6.3.2 EVP-LAN

MEF EVP-LAN services can be used to connect different Cloud Consumer locations to one or more Cloud Service Provider data centers similarly to the EP-LAN example in Section 6.3.1. However EVP-LAN services can support multiple Ethernet services (EVCs) at each UNI. The EVP-LAN use case in Figure 20 provides a similar connectivity scenario as with EP-LAN scenario in Figure 19 with three Cloud Consumer sites (Sites 1, 2 and 3) that connect to a Cloud Service Provider's primary and secondary data centers. The secondary data center could be used for cloud bursting applications or as a disaster recovery data center. However, additional Ethernet services (EVCs) can be multiplexed on the UNIs of the EVP-LAN EVC. In Figure 20, an EVPL EVC is multiplexed on Cloud Consumer Site 1 providing Internet connectivity from the same UNI (port). Another EVPL EVC is multiplexed on one of the Cloud Service Provider's data center UNIs to provide Internet connectivity. This EVPL could be used to deliver Public Cloud services or enable mobile users to connect to the private cloud application over an encrypted IP VPN over the Internet.

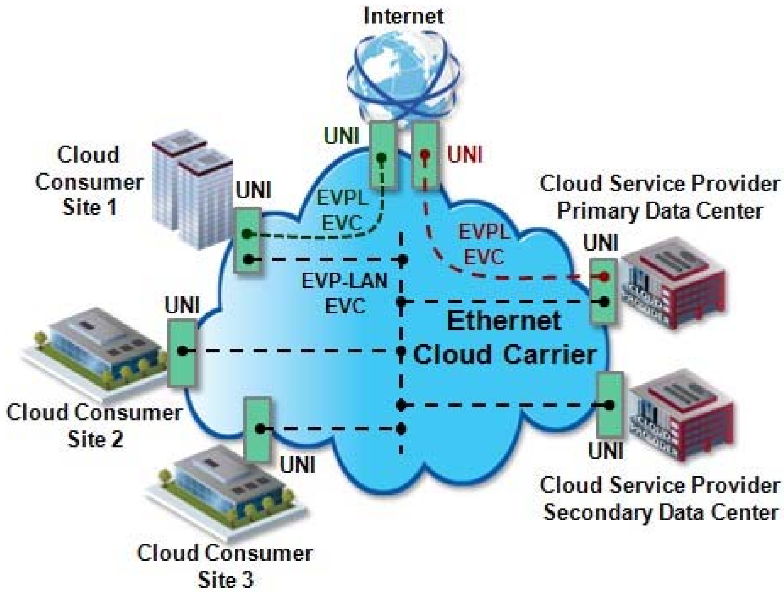

Figure 20: EVP-LAN between Cloud Consumers and Cloud Service Provider and EVPLs for Internet access

6.4 E-Access Service Type

The E-Access service type defines Ethernet services that consist of one or more UNIs and one ENNI. The MEF currently has created two Ethernet service definitions for the E-Access service type, namely, Access EPL and Access EVPL which are services which provide point-to-point connectivity between a UNI and an ENNI.

The standard MEF UNI supports untagged or single-tagged (C-VLAN tag) Ethernet frames. The standard MEF ENNI can support stacked VLAN tagged Ethernet frames where one of the tags is the C-VLAN tag (customer VLAN tag) and the other an S-VLAN tag (service VLAN tag). This tag stacking at the ENNI enables overlapping C-VLAN IDs from different Cloud Consumers to be used. This eliminates the need for the Cloud Service Provider to coordinate C-VLAN IDs with each Cloud Consumer customer as in the use case with EVPL services described in Section 6.1.2.

The fundamental difference between the E-Line services (EPL and EVPL) and Access EPL and Access EVPL services is that E-Line services interconnect two UNIs while Access EPL and Access EVPL interconnect a UNI and an ENNI. The MEF defines a UNI as an Ethernet interface that can multiplex Ethernet services typically from the same subscriber (Cloud Consumers) but if the UNI is owned by the Cloud Service Provider or data center provider, it can multiplex Ethernet services from different subscribers (Cloud Consumers).

6.4.1 Access EPL

MEF Access EPL services are applicable to use cases where a Cloud Consumer requires a dedicated Ethernet UNI for a WAN connection to a Cloud Service Provider data center. This may appear similar to the EPL services use case described in Section 6.1.1 but it has some important differences. An Access EPL uses an ENNI at one of the service endpoints unlike an EPL which uses a UNI at both service endpoints. The ENNI enables the Cloud Service Provider to multiplex several Access EPL services onto a single port where an EPL service cannot support this capability.

Access EPL services, like EPL services, do not require C-VLAN ID coordination between the Cloud Consumer and Cloud Service Provider. Since the ENNI supports service multiplexing, it is similar to an EVPL service UNI which also supports service multiplexing but with one important difference: An EVPL service requires the Cloud Consumers and Cloud Service Provider to coordinate C-VLAN IDs to ensure there are no overlapping C-VLAN IDs among the Cloud Consumers whose services are multiplexed on the UNI at the data center. Since an Access EPL uses an ENNI supporting service multiplexing and VLAN tag stacking, the Cloud Service Provider can support multiple, overlapping C-VLAN IDs from Access EPLs from each Cloud Consumer. This is possible because each Cloud Consumer can be identified by the S-VLAN ID.

In Figure 21, two Cloud Consumers (A and B) each use an Access EPL service to connect to the Cloud Service Provider's data center. The UNI at each Cloud Consumer location is dedicated for this application. The ENNI at the Cloud Service Provider data center multiplexes each Access EPL onto a common port and breaks out the services into two additional Access EPLs that connect to the appropriate data center infrastructure to which Cloud Consumers A and B require connectivity. Note that the Access EPLs shown inside the data center could be implemented inside a single networking device providing the ENNI and UNI on different ports on the device.

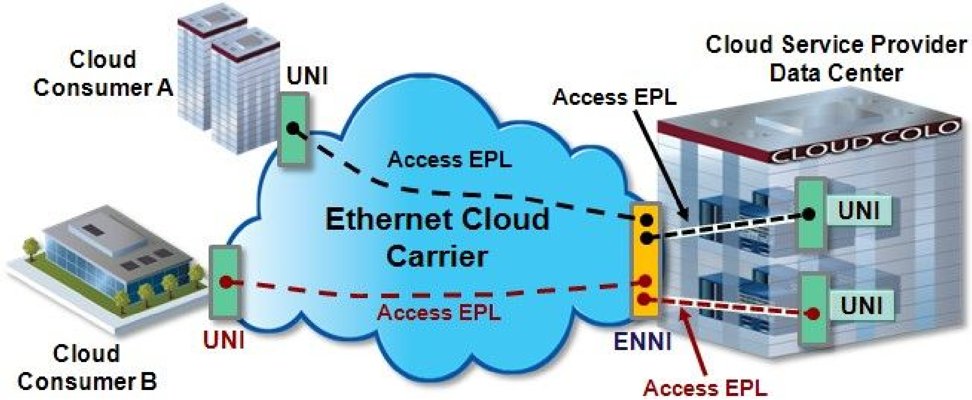

Figure 21: Access EPLs to connect Cloud Consumers at ENNI at Cloud Service Provider Data Center

6.4.2 Access EVPL

MEF Access EVPL services are similar to Access EPL with some important differences. One difference is that Access EVPL services support multiple services on their UNI service endpoints similarly to how EVPL services described in Section 6.1.2 support this. Because services are multiplexed, the Cloud Consumer needs to consider which C-VLAN IDs are used for each service. At the ENNI, as with an Access EPL, no C-VLAN ID coordination is required for each Cloud Consumer's Access EVPL.

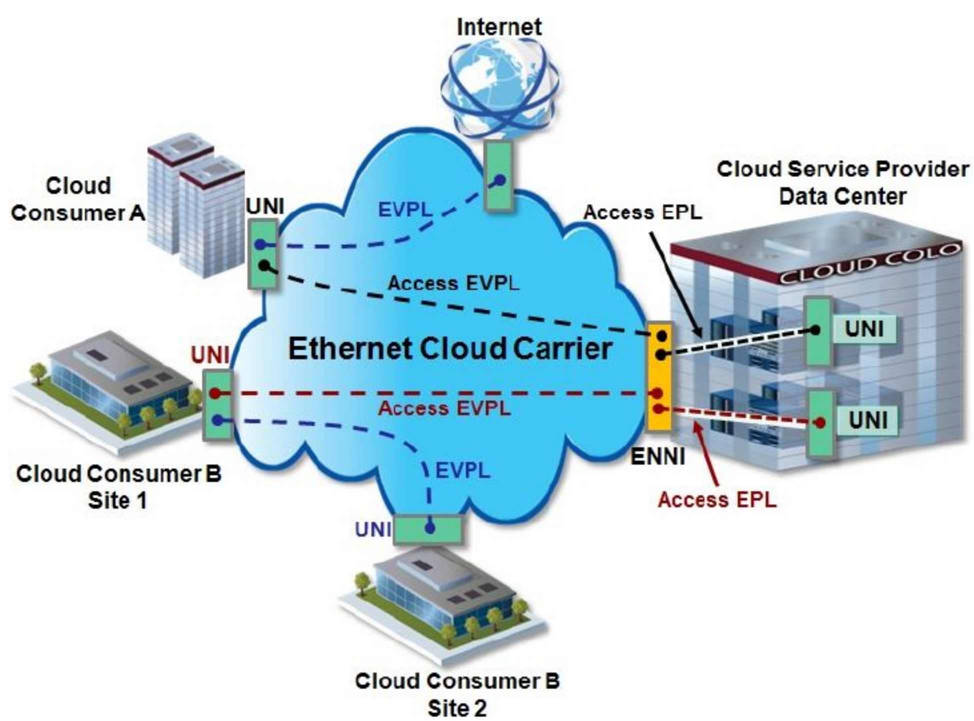

Figure 22: Access EVPLs interconnecting Cloud Consumers to a Cloud Service Provider Data Center

In Figure 22, Cloud Consumer A uses an Access EVPL and EVPL service multiplexed on the same UNI port. The Access EVPL is used to connect to the Cloud Service Provider's data center and the EVPL is used to connect to the Internet. Cloud Consumer B uses an Access EVPL from one of their business locations (Site 1) to connect to the Cloud Service Provider data center and an EVPL to interconnect their two business locations (Sites 1 and 2).

7 Static vs. "On-Demand" Ethernet Service Bandwidth

Today's Ethernet WAN services are predominantly static and do not meet the five characteristics of cloud services (section 4.2.1). By static, we mean that an Ethernet service is purchased based on a static amount of network bandwidth defined at ingress or egress of Ethernet UNI or ENNI service end points. This bandwidth is managed via what the MEF calls a bandwidth profile in which a committed information rate (CIR), specified in Mbps, is purchased. The CIR provides a fixed amount of bandwidth for the duration of the service agreement. When additional bandwidth is required, a higher CIR is purchased and the service agreement is updated. The process for doing this could take days, weeks or even months and does not currently address the "On-demand Self-service, Rapid Elasticity and Measured Service" characteristics of cloud services.

8 Dynamic Ethernet Service Attributes and Service Management

The MEF is developing new Technical Committee work to define the dynamic Ethernet service requirements to match the on-demand requirements of cloud services. The MEF is actively engaging the cloud service provider and standards development community to develop these requirements. The following areas are currently being researched.

- Ability to dynamically add CIR bandwidth to an existing Ethernet service EVC for a time period after which, the CIR would revert back to its previous value and the network resources associated with the higher CIR are released
- Ability to instantiate a new Ethernet service EVC on existing UNIs associated with another EVC
- Define a method to trigger a dynamic change in an elastic service attribute of an Ethernet service EVC
- Define the timeframe (e.g., minutes or seconds) in which elastic Ethernet service attributes (e.g., a CIR change) need to be performed including the Ethernet Cloud Carrier's operational and business support systems' (OSS/BSS) requirements to support usage-based Ethernet services
- Service management orchestration between the Cloud Service Provider data center resources and the Ethernet Cloud Carrier's network resources
- Determine inter-relationship between Cloud Service Provider service level objectives for compute and storage resources and Ethernet Cloud Carrier's network resources

9 Summary

The wide area network is critical to meet the requirements for delivering external private cloud and hybrid cloud services. Enterprises often cannot rely on the Internet to provide connectivity to their mission-critical Private Cloud applications delivered via remote data centers due to the Internet's security vulnerabilities, unpredictable performance, and compliance challenges regarding data governance.

Carrier Ethernet is the fastest growing wide area networking technology for delivery of business-class services to Enterprises. The MEF's work in standardizing Ethernet services with a rich set of service attributes has enabled the market to flourish. The MEF sees Cloud Computing and Cloud Services as having a massive impact on the Ethernet services market. To ensure that the advantages of Ethernet services can be used by the new cloud-centric IT services model, the MEF is engaging the Cloud Service Provider market to educate, understand and continue to evolve Carrier Ethernet networks and services. The MEF is also engaging the cloud community to identify their unique dynamic and elastic WAN service requirements to enable communications service providers to offer Ethernet services that optimally support the delivery of private and hybrid cloud services.

10 Glossary and Terms

A glossary of terms used in this document can be found online at http://MetroEthernetForum.org/glossary. Other terms used in this document are listed below.

Term	Description
Access EPL	Access Ethernet Private Line
Access EVPL	Access Ethernet Virtual Private Line
CIR	Committed Information Rate
Cloud bursting	The ability of a cloud application to utilize additional remote public or external private cloud compute and storage resources in the event that local resources are fully consumed when a Cloud Consumer's demand exceeds a Cloud Service Provider's local resource capacity.
Communications Service Provider	A service provider that transports information electronically. The term encompasses public and private companies in the telecom (landline and wireless), Internet, cable, satellite, and managed services businesses
E-Access	Ethernet Access service type
E-LAN	Ethernet LAN service type
E-Line	Ethernet Line service type
ENNI	External Network-to-Network Interface
EPL	Ethernet Private Line service
EP-LAN	Ethernet Private LAN service
EP-Tree	Ethernet Private Tree service
E-Tree	Ethernet Tree service type
EVC	Ethernet Virtual Connection
EVPL	Ethernet Virtual Private Line service
EVP-LAN	Ethernet Virtual Private LAN service
EVP-Tree	Ethernet Virtual Private Tree service
IaaS	Infrastructure as a Service
MSO	Multiple Systems Operator. A cable television operator that offers telecommunications services.
NIST	U.S. National Institute of Standards and Technology
OAM&P	Operations, Administration, Maintenance and Provisioning
PaaS	Platform as a Service
SaaS	Software as a Service
SLA	Service Level Agreement
UNI	User-to-Network Interface

11 Feedback on the Paper

The MEF appreciates any feedback provided on this paper to help us best address the cloud community WAN requirements using Carrier Ethernet networks and services.
Please send comments to Cloud-Feedback@MetroEthernetForum.net.

12 References and Resources

[1] "The NIST Definition of Cloud Computing", Special Publication 800-145, National Institute of Standards and Technology U.S. Department of Commerce, Peter Mell and Timothy Grance, September 2011.

[2] "NIST Cloud Computing Reference Architecture", National Institute of Standards and Technology U.S. Department of Commerce, F. Liu, J. Tong, J. Mao, R. B. Bohn, J. V. Messina, M. L. Badger, D. M. Leaf, September 2011.

[3] "Enterprise-Grade External Compute IaaS Requirements v1.00 (Final)", ECLC VPC Team, October 2010.

[4] "Carrier Ethernet Essentials", Fujitsu Network Communications, Ralph Santitoro, February 2011

13 About the MEF

The MEF is a global industry alliance comprising more than 180 leading organizations including telecommunications service providers, cable MSOs, network equipment and semiconductor manufacturers, software and testing organizations. The MEF's mission is to accelerate the worldwide adoption of Carrier-class Ethernet networks and services. The MEF develops Carrier Ethernet technical specifications and implementation agreements to promote interoperability and deployment of Carrier Ethernet worldwide. http://MetroEthernetForum.org.

14 Acknowledgements

Editor: Ralph Santitoro, Fujitsu

Contributors: Christopher Cullan, InfoVista; Yoav Cohen; Enrique Hernandez-Valencia, Alcatel-Lucent; Steve Holmgren, at&t; Mehmet Toy, Comcast; Lucy Yong, Huawei; Ellis Reid, Intune Networks; Mark Fishburn , MEF; Andy Mayer, Applied Communications Science.

Special Acknowledgement: Margaret Chiosi, at&t, for her leadership in the MEF cloud project.