The Functional Applications of Carrier Ethernet

Dmitri Kalintsev, Product Manager – Carrier Ethernet at Nextgen Networks

Abstract

This paper summarises the most common use cases of Carrier Ethernet services and describes the special needs of each particular case. The purpose of such description is to assist with more direct and accurate matching of customers' *functional requirements* with the *technical capabilities* of the available Carrier Ethernet services and products.

Introduction

Carrier Ethernet is a very popular telecommunication technology. To a large extent, this popularity is due to its high degree of flexibility, which allows Service Providers to create a variety of Carrier Ethernet products designed to meet a widest range of Enterprise and Business users' requirements.

On the flip side, the magnitude of this variety may sometimes be the source of customer confusion. In some cases, such confusion may lead to costly mistakes, when a wrong Carrier Ethernet product is chosen for the customer's intended purpose.

By developing the common understanding of the most prominent uses of the Carrier Ethernet services, the task of choosing the right service for the particular job is made less cumbersome, ultimately saving significant amounts of time and money for both the Service Provider and their Customer.

Carrier Ethernet Service Definitions and Attributes

Many Service Providers offer more than one type of Carrier Ethernet service. While all of these services are delivered using common User Network Interface (UNI) type – Ethernet, the *technical capabilities* of these services are different. To adequately describe these capabilities, a common language is required.

Such language is provided by the two MEF specification documents, which describe the Metro Ethernet Service types (MEF6.1) and their respective applicable Attributes (MEF10.1). The MEF6.1 provides definitions of three generic types of Carrier Ethernet services, based on their topology configuration:

1. Point to point (E-Line);
2. Multipoint to multipoint (E-LAN);
3. Rooted multipoint (E-Tree).

Furthermore, depending on whether a service is dedicated (**P**rivate) or multiplexed (**V**irtual **P**rivate), each of the above has two flavours, making it total of six: EPL, EVPL, EP-LAN, EVP-LAN, EP-Tree and EVP-Tree.

The MEF10.1 describes the Service Attributes, applicable either universally or to a particular service type. These Service Attributes are split into the following three groups:

1. Applicable to UNI;
2. Applicable to Ethernet Virtual Circuit (EVC);
3. Applicable to individual EVC per UNI.

The combined use of these service definitions with relevant service attributes allows unambiguous description of any Carrier Ethernet service offering.

From Service Definition to Functional Application

The MEF Service Definitions provide the framework for *standard* and *interoperable* Carrier Ethernet services created by individual Service Providers. A lot of flexibility in the way these services can be designed is provided, making possible an infinite range of innovative Ethernet-based products.

While flexibility is of an extreme importance for innovation and emerging uses, there is a clear value in a wide agreement on the Best Current Implementation Practices for the most common and well-understood *Functional Applications* of Carrier Ethernet.

The "Functional Application" is the next level down from the "Business Application". The MEF identified examples of "Business Applications" are the *Site-to-Site Access, Server Consolidation, Disaster Recovery and Business Continuity, Service Orientated Architecture, Software as a Service, Converged Networking* and *Internet Access*, each of which may have one or more functional applications mapped to it.

The "Functional Applications" proposed in this paper are the *Best Effort IP Transport, Converged IP Transport, Data Centre Interconnect, LAN Extension, Overlay Transport* and *SP Infrastructure Backhaul*. Each of these applications has its unique set of requirements for the standard Carrier Ethernet services used to deliver their required functionality. The following sections outline these requirements and map them to the relevant MEF Service Attributes.

Functional Applications in detail

Application 1: Best Effort IP Transport

The *Best Effort IP Transport* function is often used as the means to access services in best-effort IP "clouds", such as the Internet and simple L3 VPN services, as well as for the inexpensive basic IP connectivity between multiple customer sites.

This function supports both Point-to-Point (P2P) and Multipoint-to-Multipoint (M2M) topologies and can be implemented using the EPL, EVPL, EP-LAN and EVP-LAN service types, potentially with shared bandwidth.

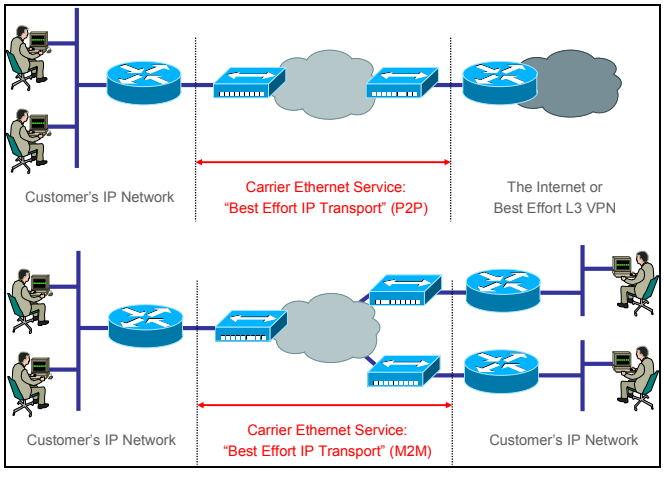

Figure 1: Best Effort IP Transport Applications

The Carrier Ethernet service used for the *Best Effort IP Transport* function is characterized by the following attributes:

- Support per-UNI or per-EVC bandwidth profiles (1 CoS), purely to control the access admission rate;

- Support for standard MTU sizes only (1518 bytes, from DA to FCS inclusive);
- No need to preserve customer's VLAN IDs and 802.1p CoS, as the service is Layer 3 (IP) centric;
- No need to tunnel Layer 2 Control Protocols (L2CP);
- Low number of MAC addresses (typically 1 per UNI), as the end-stations are behind a router.

Application 2: Converged IP Transport

The Converged IP Transport function is required when the customer are using a Carrier Ethernet service as the means to transport a mix of voice/video/interactive traffic with lower priority and best effort data, directly between customer's locations or as an access to a full-featured Layer 3 VPN.

Similarly to the Application 1 above, this function supports both P2P and M2M topologies and can be implemented using the EPL, EVPL, EP-LAN and EVP-LAN service types.

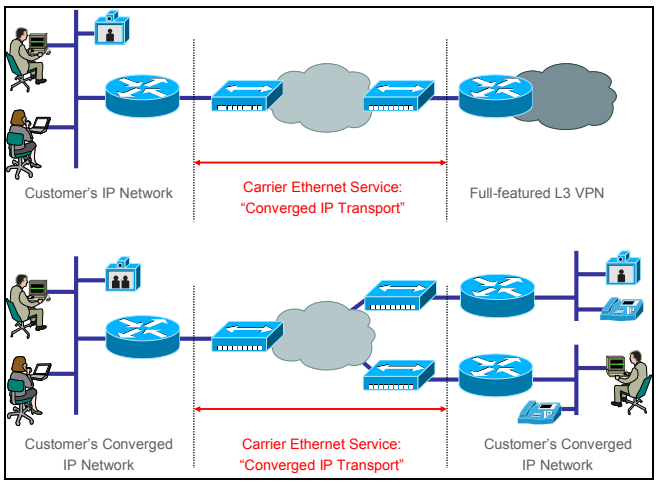

Figure 2: Converged IP Transport Applications

The Carrier Ethernet service used for the *Converged IP Transport* function is characterized by the following attributes:

- Support per-CoS (DSCP) bandwidth profiles (4 or more CoS) for differentiated treatment of voice/video/data;
- UNI capable of providing CIR of at least 768Kbit/s to avoid the head-of-line blocking due to serialisation delay;
- Support for standard MTU sizes only (1518/1522 bytes, from DA to FCS inclusive);
- No need to preserve customer's VLAN IDs and 802.1p CoS, as the service is Layer 3 (IP) centric;
- No need to tunnel Layer 2 Control Protocols;
- Low number of MAC addresses (typically 1 per UNI), as the end-stations are behind a router.

Application 3: Data Centre Interconnect

The Data Centre Interconnect function is used to provide high speed (1Gbit/s or more) Layer 2 connectivity between customer locations with high concentration of IT infrastructure (servers and storage), such as established Data Centres and/or major corporate buildings, serving as Head Office and/or Disaster Recovery.

This function is typically available in P2P configuration only and can be implemented using the EPL service type to ensure the maximum CE-VLAN ID/CoS and L2CP transparency required, with fully dedicated bandwidth.

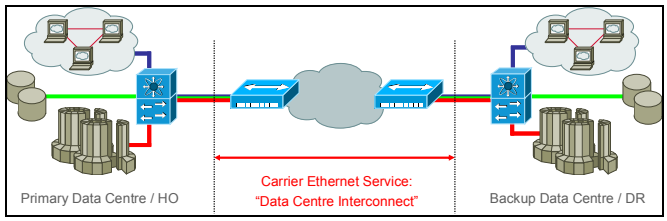

Figure 3: Data Centre Interconnect Application

The Carrier Ethernet service used for the *Data Centre Interconnect* function is characterized by the following attributes:

- Support per UNI bandwidth profiles;
- UNI capable of providing CIR of at least 1Gbit/s;
- Support for Expanded (802.3as) or Jumbo MTU to enable the most effective bandwidth utilisation;
- Preserve customer's VLAN IDs and CoS;
- Transparency to customer's Layer 2 Control Protocols;
- Point to point only with no MAC learning.

Application 4: LAN Extension

The LAN Extension function is utilised by Enterprise customers wishing to extend their campus switched LAN over longer distances or to join two or more geographically separate LAN segments.

By its nature, this function supports M2M topology, but can also be used in P2P configuration and can be implemented using the EP-LAN and EPL service types. Service multiplexing is disabled to ensure correct function of customer's Layer 2 Control Protocols, such as Spanning Tree Protocol.

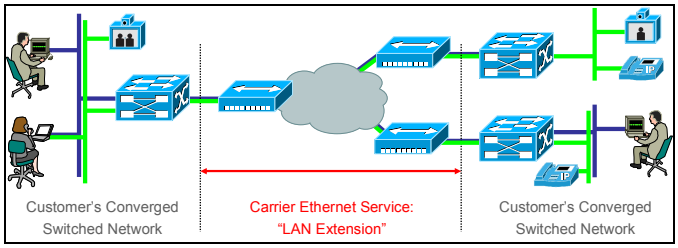

Figure 4: LAN Extension Application

The Carrier Ethernet service used for the *LAN Extension* function is characterized by the following attributes:

- Support per-CoS (802.1p) bandwidth profiles (4 or more CoS) for differentiated treatment of voice/video/data;
- UNI capable of providing CIR of at least 768Kbit/s to avoid the head-of-line blocking due to serialisation delay;
- Support for slightly oversize MTU (minimum of 1526 bytes [DA to FCS]) to support CE-VLAN transparency;
- Preserve customer's VLAN IDs and 802.1p CoS, as the service is Layer 2 (Ethernet) centric;
- Transparency to customer's Layer 2 Control Protocols to ensure best possible LAN behaviour emulation;
- Support for large number of MAC addresses (EP-LAN implementation), as end-stations MAC addresses are directly exposed.

Application 5: Overlay Transport

The Overlay Transport function is utilised by System Integrators, Customer Service Providers or Large Enterprises, who implement their own "Carrier Networks", based on the transport & traffic engineering technologies such as MPLS and PBB/PBB-TE (802.1ah/802.1Qay).

This function is predominantly deployed in P2P configuration, with M2M available optionally and can be implemented using the EPL and EP-LAN service types with fully dedicated bandwidth.

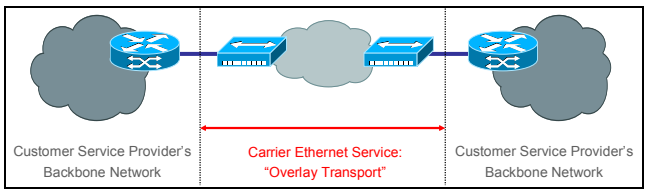

Figure 5: Overlay Transport Application

The Carrier Ethernet service used for the *Overlay Transport* function is characterized by the following attributes:

- Support Per-UNI bandwidth profiles (1 CoS), as this class of customers typically implements their own overlay QoS model;
- Support for Expanded (802.3as) or Jumbo MTU to accommodate necessary overlay encapsulation overheads;
- Preserve customer's VLAN IDs and CoS to provide the required flexibility and independence;
- Transparency to customer's Layer 2 Control Protocols;
- Support for Link and Service Ethernet OAM to facilitate effective overlay troubleshooting process;
- Medium MAC address requirements (or no MAC learning).

Application 6: SP Infrastructure Backhaul

The Service Provider (SP) Infrastructure Backhaul function is utilised by Customer Service Providers to implement backhaul connectivity between their end-user access devices, such as DSLAM, CMTS, OLT, Node-B and so on and a centralized service aggregation or gateway, such as Remote Access Server (RAS), Radio Access Network Controller (RAN NC), etc.

This function is typically implemented using a P2P or a Tree topology and can be implemented using the EPL, EVPL, EP-Tree and EVP-Tree service types; with EP-LAN and EVP-LAN being an option.

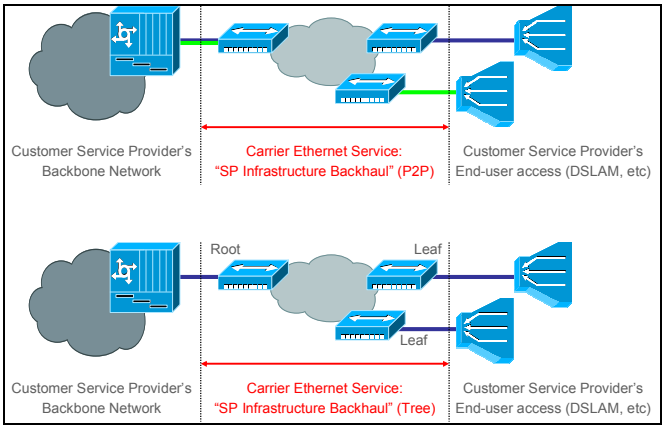

Figure 6: SP Infrastructure Backhaul Application

The Carrier Ethernet service used for the *SP Infrastructure Backhaul* function is characterized by the following attributes:

- Support per-CoS bandwidth profiles (4 or more CoS) for differentiated treatment of signalling / voice / data on M2M and E-Tree EVCs or multiplexed UNIs;
- Provide for Expanded (802.3as) or Jumbo MTU to support the efficient encapsulation (e.g. Arberg);
- Support for Link and Service Ethernet OAM to facilitate effective overlay troubleshooting process;
- Potentially high MAC address requirements (depending on the overlay implementation).

Appendix A: The Service Attributes

This Section describes the Service Attributes referred to throughout the rest of the paper, and their relevant meaning.

Bandwidth Profiles

The Bandwidth Profiles are discussed in the Section 7.11 of the MEF10.1:

"A Bandwidth Profile is a method [of] characterizing Service Frames for the purpose of rate enforcement or policing. There are two types of Bandwidth Profile. An Ingress Bandwidth Profile is used to regulate the amount of ingress traffic at a particular UNI, while an Egress Bandwidth Profile is used to regulate the amount of egress traffic at a particular UNI. The Ingress Bandwidth Profile is described in Section 7.11.2. The Egress Bandwidth Profile is described in Section 7.11.3."

Maximum Transmission Unit Size (MTU)

The MTU size Service Attribute is applicable independently to UNI and EVC; but from customer's perspective they perceive the combined effect as an ability to send Ethernet frames of a certain size over the provided service. The MTU Service Attribute is described in the MEF10.1 Section 6.10 (EVC MTU) and Section 7.4 (UNI MTU).

The standard MTU size which any Ethernet service must support is 1518 bytes, from Destination Address to FCS field, inclusive. Support for larger MTU sizes (Expanded, up to 2000 bytes as defined by the IEEE 802.3as) or jumbo (9000 bytes or more) may be required to carry overlay encapsulation overheads (such as 802.1ad, 802.1ah/802.1Qay, MPLS, PPPoE, IPSec, etc) or to facilitate the optimum bandwidth utilisation for high speed links.

Transparency to Customer VLAN ID (CE-VLAN ID Preservation)

The CE-VLAN ID preservation, observed by a customer as the service's "transparency" to customer-imposed VLAN ID tags, is described in the MEF10.1 Section 6.6.1. According to it, the egress frame(s), resulting from an original ingress frame to a service with CE-VLAN ID preservation attribute will have the same VLAN ID as the original frame if it had a VLAN tag and will be untagged if the original frame had no VLAN tag.

Transparency to Customer's CoS (CE-VLAN CoS Preservation)

The CE-VLAN CoS Preservation is described in the MEF10.1 Section 6.6.2:

*"In an EVC with CE-VLAN CoS Preservation, an egress Service Frame resulting from an ingress Service Frame that contains a CE-VLAN CoS **MUST** have the identical CE-VLAN CoS."*

This attribute applies to both Layer 2 CoS (based on 802.1p bits), and to Layer 3 CoS (based on DSCP code point).

Transparency to Customer's Layer 2 Control Protocols (UNI/EVC Control Protocol Processing)

The Layer 2 Control Protocol (L2CP) Processing is discussed in the MEF10.1 Sections 6.7 (at EVC) and 7.13 (at UNI). The definition of the "Layer 2 Control Protocol Service Frames" is provided in the Section 6.5.1.4.

While an EVC and a UNI have their independent L2CP processing attributes, the combined effect from customer's perspective is the ability to use particular L2CPs across the Ethernet service.

Note: Frames of many L2CPs are untagged, which dictates the use of non-multiplexed UNI (*thus no* EVPL, EVP-LAN or EVP-Tree).

Link and Service Ethernet OAM (L-OAM & S-OAM)

The Operations, Administration and Maintenance (OAM) capabilities of Carrier Ethernet services fall into two major categories: *Link* and *Service*.

The Link OAM functionality is described by the *IEEE 802.3-2005 Clause 57* (802.3ah), with the specific implementation requirements of a MEF-compliant UNI documented in the *UNI Type 2 Implementation Agreement* (MEF20).

The Service OAM functionality is described by the *IEEE 802.1ag* and *ITU-T Rec. Y.1731*, with the specific implementation requirements documented (for future compliance) in the *Service OAM Framework and Requirements* (MEF17).

The detailed description of the functionality provided by the Link and Service OAM is outside of the scope of this document.

MAC Limits

When an Ethernet service provides a bridge function, it always includes a MAC learning function too. The source MAC addresses of service frames entering such service are recorded in a forwarding table of a finite capacity. Service Providers control the number of MAC addresses that a particular service can learn. When the limit is reached, the further learning is stopped until inactive entries expire (typically 5 minutes). The forwarding process to the destinations which could not be added to the forwarding table is performed as if these destinations were "unknown" (i.e. flooding).

The MEF framework does not include an equivalent Service Attribute.

Appendix B: Quick Reference Table

Application	MEF Type(s)	Service Attributes
Best Effort IP Transport	EPL, EVPL, EP-LAN, EVP-LAN	No special requirements
Converged IP Transport	EPL, EVPL, EP-LAN, EVP-LAN	- 4+ CoS - 768+ Kbit/s UNI
Data Centre Interconnect	EPL	- 1+ Gbit/s UNI - Large MTU (2000+ bytes) - CE-VLAN ID+CoS and L2CP transparency
LAN Extension	EPL, EP-LAN	- 4+ CoS - 768+ Kbit/s UNI - CE-VLAN ID+CoS and L2CP transparency - Large # of MACs
Overlay Transport	EPL, EP-LAN	- Large MTU (2000+ bytes) - CE-VLAN ID+CoS and L2CP transparency - L-OAM & S-OAM
SP Infrastructure Backhaul	EPL, EVPL, EP-Tree, EVP-Tree (EP-LAN & EVP-LAN Optional)	- 4+ CoS - Large MTU (2000+ bytes) - L-OAM & S-OAM

Appendix C: Nextgen Networks' Services Map

Nextgen Networks offers three major Carrier Ethernet products:

- Ethernet over SDH (EoSDH);
- Point to point Ethernet over MPLS (EoMPLS P2P); and
- Multipoint to multipoint Ethernet over MPLS (VPLS).

The table below maps these Carrier Ethernet Services to the Functional Applications discussed above.

Functional Application	Nextgen Product	Access Method
Best Effort IP Transport	VPLS EoSDH P2P Switched Ethernet	Fibre Microwave SHDSL EoE1 ADSL
Converged IP Transport	VPLS EoSDH P2P Switched Ethernet	Fibre Microwave SHDSL EoE1
Data Centre Interconnect	EoSDH P2P Switched Ethernet	Fibre
LAN Extension	VPLS EoSDH P2P Switched Ethernet	Fibre Microwave SHDSL
Overlay Transport	VPLS EoSDH P2P Switched Ethernet	Fibre Microwave SHDSL
SP Infrastructure Backhaul	VPLS EoSDH P2P Switched Ethernet	Fibre

Table 1: Nextgen Ethernet Services to Functional Applications map

Note: the list of the functional applications discussed in this paper is by no means comprehensive and discusses only *the most common* scenarios.